IEEE Standard for Ethernet

SECTION FOUR

This section includes Clause 44 through Clause 55 and Annex 44A through Annex 55B.

Contents

44. Introduction to 10 Gb/s baseband network

44.1 Overview

44.1.1 Scope

10 Gigabit Ethernet uses the IEEE 802.3 MAC sublayer, connected through a 10 Gigabit Media Independent Interface (XGMII) to Physical Layer entities such as 10GBASE-SR, 10GBASE-LX4, 10GBASE-CX4, 10GBASE-LRM, 10GBASE-LR, 10GBASE-ER, 10GBASE-SW, 10GBASE-LW, 10GBASE-EW, and 10GBASE-T.

10 Gigabit Ethernet extends the IEEE 802.3 MAC beyond 1000 Mb/s to 10 Gb/s. The bit rate is faster and the bit times are shorter—both in proportion to the change in bandwidth. The minimum packet transmission time has been reduced by a factor of ten. A rate control mode (see 4.2.3.2.2) is added to the MAC to adapt the average MAC data rate to the SONET/SDH data rate for WAN-compatible applications of this standard. Achievable topologies for 10 Gb/s operation are comparable to those found in 1000BASE-X full duplex mode and equivalent to those found in WAN applications.

10 Gigabit Ethernet is defined for full duplex mode of operation only.

44.1.2 Objectives

The following are the objectives of 10 Gigabit Ethernet:

a) Support the full duplex Ethernet MAC.
b) Provide 10 Gb/s data rate at the XGMII.
c) Support LAN PMDs operating at 10 Gb/s, and WAN PMDs operating at SONET STS-192c/SDH VC-4-64c rate.
d) Support cable plants using cabled optical fiber compliant with ISO/IEC 11801:1995.
e) Allow for a nominal network extent of up to 40 km.
f) Support operation over a twinaxial cable assembly for wiring closet and data center applications.
g) Support operation over selected copper media from ISO/IEC 11801:2002.
h) Support a BER objective of 10^{-12}.

44.1.3 Relationship of 10 Gigabit Ethernet to the ISO OSI reference model

10 Gigabit Ethernet couples the IEEE 802.3 MAC to a family of 10 Gb/s Physical Layers. The relationships among 10 Gigabit Ethernet, the IEEE 802.3 MAC, and the ISO Open System Interconnection (OSI) reference model are shown in Figure 44–1.

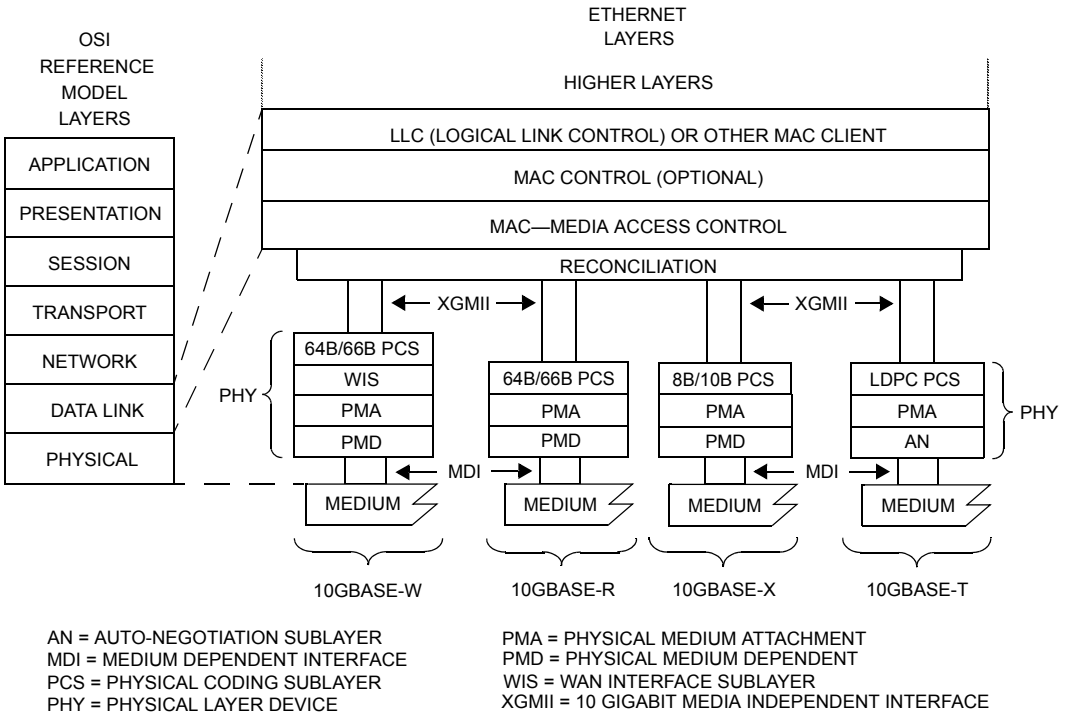

Figure 44–1—Architectural positioning of 10 Gigabit Ethernet

It is important to note that, while this specification defines interfaces in terms of bits, octets, and frames, implementations may choose other data-path widths for implementation convenience. The only exceptions are as follows:

a) The XGMII, which, when implemented at an observable interconnection port, uses a four octet-wide data path as specified in Clause 46.
b) The management interface, which, when physically implemented as the MDIO/MDC (Management Data Input/Output and Management Data Clock) at an observable interconnection port, uses a bit-wide data path as specified in Clause 45.
c) The PMA Service Interface, which, when physically implemented as the XSBI (10 Gigabit Sixteen Bit Interface) at an observable interconnection port, uses a 16-bit-wide data path as specified in Clause 51.
d) The MDI as specified in Clause 53 for 10GBASE-LX4, in Clause 54 for 10GBASE-CX4, in Clause 55 for 10GBASE-T, in Clause 68 for 10GBASE-LRM, and in Clause 52 for other PMD types.

44.1.4 Summary of 10 Gigabit Ethernet sublayers

44.1.4.1 Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII)

The 10 Gigabit Media Independent Interface (Clause 46) provides an interconnection between the Media Access Control (MAC) sublayer and Physical Layer entities (PHY). This XGMII supports 10 Gb/s operation through its 32-bit-wide transmit and receive data paths. The Reconciliation Sublayer provides a mapping between the signals provided at the XGMII and the MAC/PLS service definition.

While the XGMII is an optional interface, it is used extensively in this standard as a basis for functional specification and provides a common service interface for Clause 47, Clause 48, Clause 49, and Clause 55.

44.1.4.2 XGMII Extender Sublayer (XGXS) and 10 Gigabit Attachment Unit Interface (XAUI)

The 10 Gigabit Attachment Unit Interface (Clause 47) provides an interconnection between two XGMII Extender sublayers to increase the reach of the XGMII. This XAUI supports 10 Gb/s operation through its four-lane, differential-pair transmit and receive paths. The XGXS provides a mapping between the signals provided at the XGMII and the XAUI.

44.1.4.3 Management interface (MDIO/MDC)

The MDIO/MDC management interface (Clause 45) provides an interconnection between MDIO Manageable Devices (MMD) and Station Management (STA) entities.

44.1.4.4 Physical Layer signaling systems

This standard specifies a family of Physical Layer implementations. The generic term 10 Gigabit Ethernet refers to any use of the 10 Gb/s IEEE 802.3 MAC (the 10 Gigabit Ethernet MAC) coupled with any IEEE 802.3 10GBASE Physical Layer implementation. Table 44–1 specifies the correlation between nomenclature and clauses. Implementations conforming to one or more nomenclatures shall meet the requirements of the corresponding clauses.

Table 44–1—Nomenclature and clause correlation

Nomenclature	Clause										
	48	49	50	51	52			53	54	55	68
	8B/10B PCS & PMA	64B/66B PCS	WIS	Serial PMA	850 nm Serial PMD	1310 nm Serial PMD	1550 nm Serial PMD	1310 nm WDM PMD	4-Lane electrical PMD	Twisted-pair PCS & PMA	1310 nm Serial MMF PMD
10GBASE-SR		M[a]		M	M						
10GBASE-SW		M	M	M	M						
10GBASE-LX4	M							M			
10GBASE-CX4	M								M		
10GBASE-LR		M		M		M					
10GBASE-LW		M	M	M		M					
10GBASE-ER		M		M			M				
10GBASE-EW		M	M	M			M				
10GBASE-T										M	
10GBASE-LRM		M		M							M

[a]M = Mandatory

The term 10GBASE-X, specified in Clause 48, Clause 53, and Clause 54, refers to a specific family of Physical Layer implementations based upon 8B/10B data coding method. The 10GBASE-X family of Physical Layer implementations is composed of 10GBASE-LX4 and 10GBASE-CX4.

The term 10GBASE-R refers to a specific family of Physical Layer implementations. The 10GBASE-R family of Physical Layer implementations based upon 64B/66B data coding method is composed of 10GBASE-SR, 10GBASE-LR, 10GBASE-ER, and 10GBASE-LRM.

The term 10GBASE-W, specified in Clause 49 through Clause 52, refers to a specific family of Physical Layer implementations based upon STS-192c/SDH VC-4-64c encapsulation of 64B/66B encoded data. The 10GBASE-W family of Physical Layer standards has been adapted from the ATIS-0600416.1999(R2010) (SONET STS-192c/SDH VC-4-64c) Physical Layer specifications. The 10GBASE-W family of Physical Layer implementations is composed of 10GBASE-SW, 10GBASE-LW, and 10GBASE-EW.

All 10GBASE-R and 10GBASE-W PHY devices share a common PCS specification (see Clause 49). The 10GBASE-W PHY devices also require the use of the WAN Interface Sublayer, (WIS) (Clause 50).

The term 10GBASE-T, specified in Clause 55, refers to a specific Physical Layer implementation based upon 64B/65B data coding placed in a low density parity check (LDPC) frame that is mapped to a 128 double-square (DSQ128) constellation for transmission on 4-pair, twisted-pair copper cabling.

Physical Layer device specifications are contained in Clause 52, Clause 53, Clause 54, Clause 55, and Clause 68.

Annex 44A contains diagrams of the data flow between the MAC and the MDI, as well as information on the relation between data valid signals and loopback.

44.1.4.5 WAN Interface Sublayer (WIS), type 10GBASE-W

The WIS provides a 10GBASE-W device with the capability to transmit and receive IEEE 802.3 MAC frames within the payload envelope of a SONET STS-192c/SDH VC-4-64c frame.

44.1.5 Management

Managed objects, attributes, and actions are defined for all 10 Gigabit Ethernet components. These items are defined in Clause 30.

44.2 State diagrams

State diagrams take precedence over text.

The conventions of 1.2 are adopted, along with the extensions listed in 21.5.

44.3 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementers must conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices. Table 44–2 contains the values of maximum sublayer round-trip (sum of transmit and receive) delay in bit time as specified in 1.4 and pause_quanta as specified in 31B.2.

Equation (44–1) specifies the calculation of bit time per meter of fiber or electrical cable based upon the parameter *n*, which represents the ratio of the speed of electromagnetic propagation in the fiber or electrical cable to the speed of light in a vacuum. The value of *n* should be available from the fiber or electrical cable manufacturer, but if no value is known then a conservative delay estimate can be calculated using a default value of $n = 0.66$. The speed of light in a vacuum is $c = 3 \times 10^8$ m/s. Table 44–3 can be used to convert fiber or electrical cable delay values specified relative to the speed of light or in nanoseconds per meter.

$$\text{cable delay} = \frac{10^{10}}{nc} \text{ BT/m} \tag{44–1}$$

Table 44–2—Round-trip delay constraints (informative)

Sublayer	Maximum (bit time)	Maximum (pause_quanta)	Notes
MAC, RS, and MAC Control	8192	16	See 46.1.4.
XGXS and XAUI	4096	8	Round-trip of 2 XGXS and trace for both directions. See 47.2.2.
10GBASE-X PCS and PMA	2048	4	See 48.5
10GBASE-R PCS	3584	7	See 49.2.15.
WIS	14336	28	See 50.3.7.
LX4 PMD	512	1	Includes 2 m of fiber. See 53.2.
CX4 PMD	512	1	See 54.3.
Serial PMA and PMD (except LRM)	512	1	Includes 2 m of fiber. See 52.2.
LRM PMA and PMD	9216	18	Includes 2 m of fiber. See 68.2.
10GBASE-T PHY	25 600	50	See 55.11.

Table 44–3—Conversion table for cable delays

Speed relative to *c*	ns/m	BT/m
0.40	8.33	83.3
0.50	6.67	66.7
0.51	6.54	65.4
0.52	6.41	64.1
0.53	6.29	62.9
0.54	6.17	61.7
0.55	6.06	60.6
0.56	5.95	59.5
0.57	5.85	58.5
0.58	5.75	57.5
0.5852	5.70	57.0
0.59	5.65	56.5
0.60	5.56	55.6
0.61	5.46	54.6

Table 44–3—Conversion table for cable delays *(continued)*

Speed relative to *c*	ns/m	BT/m
0.62	5.38	53.8
0.63	5.29	52.9
0.64	5.21	52.1
0.65	5.13	51.3
0.654	5.10	51.0
0.66	5.05	50.5
0.666	5.01	50.1
0.67	4.98	49.8
0.68	4.90	49.0
0.69	4.83	48.3
0.7	4.76	47.6
0.8	4.17	41.7
0.9	3.70	37.0

44.4 Protocol implementation conformance statement (PICS) proforma

The supplier of a protocol implementation that is claimed to conform to any part of IEEE Std 802.3, Clause 45 through Clause 55 and Clause 68, demonstrates compliance by completing a protocol implementation conformance statement (PICS) proforma.

A completed PICS proforma is the PICS for the implementation in question. The PICS is a statement of which capabilities and options of the protocol have been implemented. A PICS is included at the end of each clause as appropriate. Each of the 10 Gigabit Ethernet PICS conforms to the same notation and conventions used in 100BASE-T (see 21.6).

45. Management Data Input/Output (MDIO) Interface

45.1 Overview

This clause defines the logical and electrical characteristics of an extension to the two signal Management Data Input/Output (MDIO) Interface specified in Clause 22.

The purpose of this extension is to provide the ability to access more device registers while still retaining logical compatibility with the MDIO interface defined in Clause 22. Clause 22 specifies the MDIO frame format and uses an ST code of 01 to access registers. In this clause, additional registers are added to the address space by defining MDIO frames that use an ST code of 00.

This extension to the MDIO interface is applicable to the following:

— Implementations that operate at speeds of 10 Gb/s and above.
— Implementations of 10PASS-TS and 2BASE-TL subscriber network Physical Layer devices.
— Implementations of 10, 100, or 1000 Mb/s with additional management functions beyond those defined in Clause 22.

The MDIO electrical interface is optional. Where no physical embodiment of the MDIO exists, provision of an equivalent mechanism to access the registers is recommended.

Throughout this clause, an “a.b.c” format is used to identify register bits, where “a” is the device address, “b” is the register address, and “c” is the bit number within the register.

45.1.1 Summary of major concepts

The following are major concepts of the MDIO Interface:

a) Preserve the management frame structure defined in 22.2.4.5.
b) Define a mechanism to address more registers than specified in 22.2.4.5.
c) Define ST and OP codes to identify and control the extended access functions.
d) Provide an electrical interface specification that is compatible with common digital CMOS ASIC processes.

45.1.2 Application

This clause defines a management interface between Station Management (STA) and the sublayers that form a Physical Layer device (PHY) entity. Where a sublayer, or grouping of sublayers, is an individually manageable entity, it is known as an MDIO Manageable Device (MMD). This clause allows a single STA, through a single MDIO interface, to access up to 32 PHYs (defined as PRTAD in the frame format defined in 45.3) consisting of up to 32 MMDs as shown in Figure 45–1. The MDIO interface can support up to a maximum of 65 536 registers in each MMD.

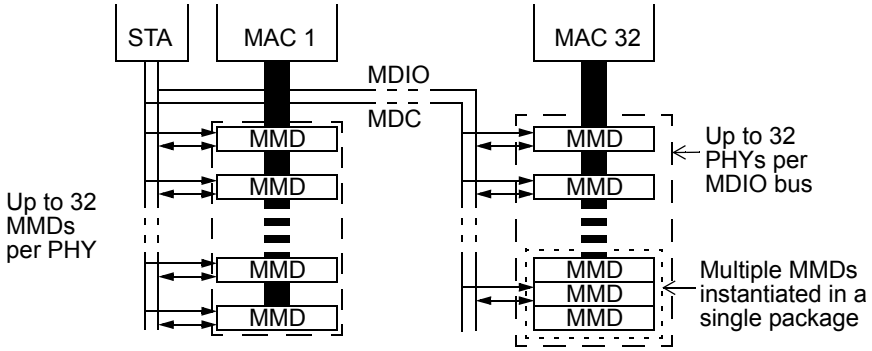

Figure 45–1—DTE and MMD devices

45.2 MDIO Interface Registers

The management interface specified in Clause 22 provides a simple, two signal, serial interface to connect a Station Management entity and a managed PHY for providing access to management parameters and services. The interface is referred to as the MII management interface.

The MDIO interface is based on the MII management interface, but differs from it in several ways. The MDIO interface uses indirect addressing to create an extended address space allowing a much larger number of registers to be accessed within each MMD. The MDIO address space is orthogonal to the MII management interface address space. The mechanism for the addressing is defined in 45.3. The MDIO electrical interface operates at lower voltages than those specified for the MII management interface. The electrical interface is specified in 45.4. For cases where a single entity combines Clause 45 MMDs with Clause 22 registers, then the Clause 22 registers may be accessed using the Clause 45 electrical interface and the Clause 22 management frame structure. The list of possible MMDs is shown in Table 45–1. The PHY XS and DTE XS devices are the two partner devices used to extend the interface that sits immediately below the Reconciliation Sublayer. For 10 Gigabit Ethernet, the interface extenders are defined as the XGXS devices. For 10PASS-TS and 2BASE-TL, control and monitoring of the TC sublayer is defined in the TC MMD. For 10, 100 and 1000 Mb/s PHYs, further management capability is defined in the Clause 22 extension MMD.

Table 45–1—MDIO Manageable Device addresses

Device address	MMD name
0	Reserved
1	PMA/PMD
2	WIS
3	PCS
4	PHY XS
5	DTE XS
6	TC

Table 45–1—MDIO Manageable Device addresses *(continued)*

Device address	MMD name
7	Auto-Negotiation
8	Separated PMA (1)
9	Separated PMA (2)
10	Separated PMA (3)
11	Separated PMA (4)
12 through 28	Reserved
29	Clause 22 extension
30	Vendor specific 1
31	Vendor specific 2

10PASS-TS and 2BASE-TL each have two port subtypes, 10PASS-TS-O, 10PASS-TS-R, 2BASE-TL-O and 2BASE-TL-R. Hereafter, referred to generically as -O and -R. The -O subtype corresponds to the port located at the service provider end of a subscriber link (the central office end). The -R subtype corresponds to the port located at the subscriber end of a subscriber link (the remote end). See 61.1 for more information.

Some register behavior may differ based on the port subtype. In the case where a register's behavior or definition differs between port subtypes, it is noted in the register description and in the bit definition tables (denoted by "O:" and "R:" in the R/W column).

The Clause 22 extension MMD allows new features to be added to 10, 100, and 1000 Mb/s PHYs beyond those already defined in Clause 22.

If a device supports the MDIO interface it shall respond to all possible register addresses for the device and return a value of zero for undefined and unsupported registers. Writes to undefined registers and read-only registers shall have no effect. The operation of an MMD shall not be affected by writes to reserved and unsupported register bits, and such register bits shall return a value of zero when read.

In the case of two registers that together form a 32-bit counter, whenever the most significant 16-bit register of the counter is read, the 32-bit counter value is latched into the register pair, the value being latched before the contents of the most significant 16 bits are driven on the MDIO interface and the contents of both registers is cleared to all zeros. A subsequent read from the least significant 16-bit register will return the least significant 16 bits of the latched value, but will not change the contents of the register pair. Writing to these registers has no effect. Counters that adhere to this behavior are marked in their bit definition tables with the tag "MW = Multi-word".

To ensure compatibility with future use of reserved bits and registers, the Management Entity should write to reserved bits with a value of zero and ignore reserved bits on read.

Some of the bits within MMD registers are defined as latching low (LL) or latching high (LH). When a bit is defined as latching low and the condition for the bit to be low has occurred, the bit shall remain low until after it has been read via the management interface. Once such a read has occurred, the bit shall assume a value based on the current state of the condition it monitors. When a bit is defined as latching high and the condition for the bit to be high has occurred, the bit shall remain high until after it has been read via the management interface. Once such a read has occurred, the bit shall assume a value based on the current state of the condition it monitors.

For multi-bit fields, the lowest numbered bit of the field in the register corresponds to the least significant bit of the field.

Figure 45–2 describes the signal terminology used for the MMDs.

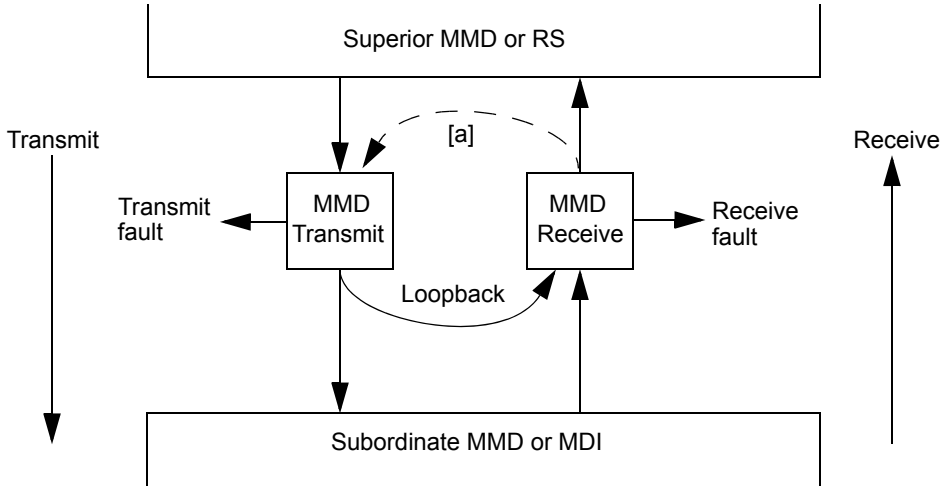

Figure 45–2—MMD signal terminology

[a] Direction of the optional PHY XS loopback

Each MMD contains registers 5 and 6, as defined in Table 45–2. Bits read as a one in this register indicate which MMDs are instantiated within the same package as the MMD being accessed. Bit 5.0 is used to indicate that Clause 22 functionality has been implemented within a Clause 45 electrical interface device. Bit 6.13 indicates that Clause 22 functionality is extended using the Clause 45 electrical interface through MMD 29. The definition of the term package is vendor specific and could be a chip, module, or other similar entity.

Table 45–2—Devices in package registers bit definitions

Bit(s)[a]	Name	Description	R/W[b]
m.6.15	Vendor-specific device 2 present	1 = Vendor-specific device 2 present in package 0 = Vendor-specific device 2 not present in package	RO
m.6.14	Vendor-specific device 1 present	1 = Vendor-specific device 1 present in package 0 = Vendor-specific device 1 not present in package	RO
m.6.13	Clause 22 extension present	1 = Clause 22 extension present in package 0 = Clause 22 extension not present in package	RO
m.6.12:0	Reserved	Value always 0	RO
m.5.15:12	Reserved	Value always 0	RO
m.5.11	Separated PMA (4) present	1 = Separated PMA[c] (4) present in package 0 = Separated PMA (4) not present in package	RO
m.5.10	Separated PMA (3) present	1 = Separated PMA (3) present in package 0 = Separated PMA (3) not present in package	RO

Table 45–2—Devices in package registers bit definitions *(continued)*

Bit(s)[a]	Name	Description	R/W[b]
m.5.9	Separated PMA (2) present	1 = Separated PMA (2) present in package 0 = Separated PMA (2) not present in package	RO
m.5.8	Separated PMA (1) present	1 = Separated PMA (1) present in package 0 = Separated PMA (1) not present in package	RO
m.5.7	Auto-Negotiation present	1 = Auto-Negotiation present in package 0 = Auto-Negotiation not present in package	RO
m.5.6	TC present	1 = TC present in package 0 = TC not present in package	RO
m.5.5	DTE XS present	1 = DTE XS present in package 0 = DTE XS not present in package	RO
m.5.4	PHY XS present	1 = PHY XS present in package 0 = PHY XS not present in package	RO
m.5.3	PCS present	1 = PCS present in package 0 = PCS not present in package	RO
m.5.2	WIS present	1 = WIS present in package 0 = WIS not present in package	RO
m.5.1	PMD/PMA present	1 = PMA/PMD present in package 0 = PMA/PMD not present in package	RO
m.5.0	Clause 22 registers present	1 = Clause 22 registers present in package 0 = Clause 22 registers not present in package	RO

[a]m = address of MMD accessed (see Table 45–1)
[b]RO = Read only
[c]Separated PMAs are defined in 45.2.1

45.2.1 PMA/PMD registers

For devices operating at 40 Gb/s or higher speeds, the PMA may be instantiated as multiple sublayers (see 83.1.4 for how MMD addresses are allocated to multiple PMA sublayers). A PMA sublayer that is packaged with the PMD is addressed as MMD 1. More addressable instances of PMA sublayers, each one separated from lower addressable instances, may be implemented and addressed as MMD 8, 9, 10, and 11 where MMD 8 is the closest to the PMD and MMD 11 is the furthest from the PMD. The addresses and functions of all registers in MMD 8, 9, 10 and 11 are defined identically to MMD 1, except registers m.5 and m.6 as defined in Table 45–2.

The assignment of registers in the PMA/PMD is shown in Table 45–3.

Table 45–3—PMA/PMD registers

Register address	Register name	Subclause
1.0	PMA/PMD control 1	45.2.1.1
1.1	PMA/PMD status 1	45.2.1.2
1.2, 1.3	PMA/PMD device identifier	45.2.1.3
1.4	PMA/PMD speed ability	45.2.1.4
1.5, 1.6	PMA/PMD devices in package	45.2.1.5
1.7	PMA/PMD control 2	45.2.1.6
1.8	PMA/PMD status 2	45.2.1.7
1.9	PMA/PMD transmit disable	45.2.1.8
1.10	PMD receive signal detect	45.2.1.9
1.11	PMA/PMD extended ability	45.2.1.10
1.12	10G-EPON PMA/PMD ability	45.2.1.11
1.13	40G/100G PMA/PMD extended ability	45.2.1.12
1.14, 1.15	PMA/PMD package identifier	45.2.1.13
1.16	EEE capability	45.2.1.14
1.17 through 1.29	Reserved	
1.30	10P/2B PMA/PMD control	45.2.1.15
1.31	10P/2B PMA/PMD status	45.2.1.16
1.32	10P/2B link partner PMA/PMD control[a]	45.2.1.17
1.33	10P/2B link partner PMA/PMD status[a]	45.2.1.18
1.34, 1.35	Reserved	
1.36	10P/2B link loss counter	45.2.1.19
1.37	10P/2B RX SNR margin	45.2.1.20
1.38	10P/2B link partner RX SNR margin[a]	45.2.1.21
1.39	10P/2B line attenuation	45.2.1.22
1.40	10P/2B link partner line attenuation[a]	45.2.1.23
1.41	10P/2B line quality thresholds	45.2.1.24
1.42	2B link partner line quality thresholds[a]	45.2.1.25
1.43	10P FEC correctable errors counter	45.2.1.26
1.44	10P FEC uncorrectable errors counter	45.2.1.27
1.45	10P link partner FEC correctable errors[a]	45.2.1.28
1.46	10P link partner FEC uncorrectable errors[a]	45.2.1.29
1.47	10P electrical length	45.2.1.30
1.48	10P link partner electrical length[a]	45.2.1.31
1.49	10P PMA/PMD general configuration[a]	45.2.1.32
1.50	10P PSD configuration[a]	45.2.1.33
1.51, 1.52	10P downstream data rate configuration[a]	45.2.1.34

Table 45–3—PMA/PMD registers *(continued)*

Register address	Register name	Subclause
1.53	10P downstream Reed-Solomon configuration[a]	45.2.1.35
1.54, 1.55	10P upstream data rate[a]	45.2.1.36
1.56	10P upstream Reed-Solomon configuration[a]	45.2.1.37
1.57, 1.58	10P tone group	45.2.1.38
1.59, 1.60, 1.61, 1.62, 1.63	10P tone control parameters[a]	45.2.1.39
1.64	10P tone control action[a]	45.2.1.40
1.65, 1.66, 1.67	10P tone status	45.2.1.41
1.68	10P outgoing indicator bits	45.2.1.42
1.69	10P incoming indicator bits	45.2.1.43
1.70	10P cyclic extension configuration	45.2.1.44
1.71	10P attainable downstream data rate	45.2.1.45
1.72 through 1.79	Reserved	
1.80	2B general parameter	45.2.1.46
1.81 through 1.88	2B PMD parameters	45.2.1.47
1.89	2B code violation errors counter	45.2.1.48
1.90	2B link partner code violation errors[a]	45.2.1.49
1.91	2B errored seconds counter	45.2.1.50
1.92	2B link partner errored seconds[a]	45.2.1.51
1.93	2B severely errored seconds counter	45.2.1.52
1.94	2B link partner severely errored seconds[a]	45.2.1.53
1.95	2B LOSW counter	45.2.1.54
1.96	2B link partner LOSW[a]	45.2.1.55
1.97	2B unavailable seconds counter	45.2.1.56
1.98	2B link partner unavailable seconds[a]	45.2.1.57
1.99	2B state defects	45.2.1.58
1.100	2B link partner state defects[a]	45.2.1.59
1.101	2B negotiated constellation	45.2.1.60
1.102 through 1.109	2B extended PMD parameters	45.2.1.61
1.110 through 1.128	Reserved	
1.129	10GBASE-T status	45.2.1.62
1.130	10GBASE-T pair swap and polarity	45.2.1.63
1.131	10GBASE-T TX power backoff and PHY short reach setting	45.2.1.64
1.132	10GBASE-T test mode	45.2.1.65
1.133	10GBASE-T SNR operating margin channel A	45.2.1.66
1.134	10GBASE-T SNR operating margin channel B	45.2.1.67
1.135	10GBASE-T SNR operating margin channel C	45.2.1.68
1.136	10GBASE-T SNR operating margin channel D	45.2.1.69

Table 45–3—PMA/PMD registers *(continued)*

Register address	Register name	Subclause
1.137	10GBASE-T minimum margin channel A	45.2.1.70
1.138	10GBASE-T minimum margin channel B	45.2.1.71
1.139	10GBASE-T minimum margin channel C	45.2.1.72
1.140	10GBASE-T minimum margin channel D	45.2.1.73
1.141	10GBASE-T RX signal power channel A	45.2.1.74
1.142	10GBASE-T RX signal power channel B	45.2.1.75
1.143	10GBASE-T RX signal power channel C	45.2.1.76
1.144	10GBASE-T RX signal power channel D	45.2.1.77
1.145 through 1.146	10GBASE-T skew delay	45.2.1.78
1.147	10GBASE-T fast retrain status and control register	45.2.1.79
1.148 through 1.149	Reserved	
1.150	BASE-R PMD control	45.2.1.80
1.151	BASE-R PMD status	45.2.1.81
1.152	BASE-R LP coefficient update, lane 0	45.2.1.82
1.153	BASE-R LP status report, lane 0	45.2.1.83
1.154	BASE-R LD coefficient update, lane 0	45.2.1.84
1.155	BASE-R LD status report, lane 0	45.2.1.85
1.156	BASE-R PMD status 2	45.2.1.86
1.157	BASE-R PMD status 3	45.2.1.87
1.158 through 1.159	Reserved	
1.160	1000BASE-KX control	45.2.1.88
1.161	1000BASE-KX status	45.2.1.89
1.162 through 1.164	PMA overhead control 1, 2, and 3 registers	45.2.1.90
1.165, 1.166	PMA overhead status 1 and 2 registers	45.2.1.91
1.167 through 1.169	Reserved	
1.170	BASE-R FEC ability	45.2.1.92
1.171	BASE-R FEC control	45.2.1.93
1.172 through 1.173	10GBASE-R FEC corrected blocks counter	45.2.1.94
1.174 through 1.175	10GBASE-R FEC uncorrected blocks counter	45.2.1.95
1.176 through 1.178	Reserved	
1.179	CAUI-4 chip-to-module recommended CTLE	45.2.1.96
1.180 through 1.183	CAUI-4 chip-to-chip transmitter equalization, receive direction, lane 0 through lane 3	45.2.1.97, 45.2.1.98
1.184 through 1.187	CAUI-4 chip-to-chip transmitter equalization, transmit direction, lane 0 through lane 3	45.2.1.99, 45.2.1.100
1.188 through 1.199	Reserved	
1.200	RS-FEC control register	45.2.1.101

Table 45–3—PMA/PMD registers *(continued)*

Register address	Register name	Subclause
1.201	RS-FEC status register	45.2.1.102
1.202, 1.203	RS-FEC corrected codewords counter	45.2.1.103
1.204, 1.205	RS-FEC uncorrected codewords counter	45.2.1.104
1.206	RS-FEC lane mapping register	45.2.1.105
1.207 through 1.209	Reserved	
1.210 through 1.217	RS-FEC symbol error counter, lane 0 to 3	45.2.1.106, 45.2.1.107
1.218 through 1.229	Reserved	
1.230 through 1.249	RS-FEC BIP error counter, lane 0 to 19	45.2.1.108, 45.2.1.109
1.250 through 1.269	RS-FEC PCS lane mapping, lane 0 to 19	45.2.1.110, 45.2.1.111
1.270 through 1.279	Reserved	
1.280 through 1.283	RS-FEC PCS alignment status 1 through 4	45.2.1.112 to 45.2.1.115
1.284 through 1.299	Reserved	
1.300 through 1.339	BASE-R FEC corrected blocks counter, lanes 0 through 19	45.2.1.116
1.340 through 1.699	Reserved	
1.700 through 1.739	BASE-R FEC uncorrected blocks counter, lanes 0 through 19	45.2.1.117
1.740 through 1.1099	Reserved	
1.1100	BASE-R LP coefficient update, lane 0 (copy)	45.2.1.82
1.1101 through 1.1109	BASE-R LP coefficient update, lanes 1 through 9	45.2.1.118
1.1110 through 1.1199	Reserved	
1.1200	BASE-R LP status report, lane 0 (copy)	45.2.1.83
1.1201 through 1.1209	BASE-R LP status report, lanes 1 through 9	45.2.1.119
1.1210 through 1.1299	Reserved	
1.1300	BASE-R LD coefficient update, lane 0 (copy)	45.2.1.84
1.1301 through 1.1309	BASE-R LD coefficient update, lanes 1 through 9	45.2.1.120
1.1310 through 1.1399	Reserved	
1.1400	BASE-R LD status report, lane 0 (copy)	45.2.1.85
1.1401 through 1.1409	BASE-R LD status report, lanes 1 through 9	45.2.1.121
1.1410 through 1.1449	Reserved	
1.1450 through 1.1453	PMD training pattern, lanes 0 to 3	45.2.1.122
1.1454 through 1.1499	Reserved	
1.1500	Test-pattern ability	45.2.1.123
1.1501	PRBS pattern testing control	45.2.1.124
1.1502 through 1.1509	Reserved	
1.1510	Square wave testing control	45.2.1.125

Table 45–3—PMA/PMD registers *(continued)*

Register address	Register name	Subclause
1.1511 through 1.1599	Reserved	
1.1600 through 1.1609	PRBS Tx error counters, lane 0 through lane 9	45.2.1.126
1.1610 through 1.1699	Reserved	
1.1700 through 1.1709	PRBS Rx error counters, lane 0 through lane 9	45.2.1.127
1.1710 through 1.1799	Reserved	
1.1800	TimeSync PMA/PMD capability	45.2.1.128
1.1801 through 1.1804	TimeSync PMA/PMD transmit path data delay	45.2.1.129
1.1805 through 1.1808	TimeSync PMA/PMD receive path data delay	45.2.1.130
1.1809 through 1.32767	Reserved	
1.32 768 through 1.65 535	Vendor specific	

[a]Register is defined only for -O port types and is reserved for -R ports.

45.2.1.1 PMA/PMD control 1 register (Register 1.0)

The assignment of bits in the PMA/PMD control 1 register is shown in Table 45–4. The default value for each bit of the PMA/PMD control 1 register has been chosen so that the initial state of the device upon power up or completion of reset is a normal operational state without management intervention.

Table 45–4—PMA/PMD control 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.0.15	Reset	1 = PMA/PMD reset 0 = Normal operation	R/W SC
1.0.14	Reserved	Value always 0	RO
1.0.13	Speed selection (LSB)	1.0.6 1.0.13 1 1 = bits 5:2 select speed 1 0 = 1000 Mb/s 0 1 = 100 Mb/s 0 0 = 10 Mb/s	R/W
1.0.12	Reserved	Value always 0	RO
1.0.11	Low power	1 = Low-power mode 0 = Normal operation	R/W
1.0.10:7	Reserved	Value always 0	RO
1.0.6	Speed selection (MSB)	1.0.6 1.0.13 1 1 = bits 5:2 select speed 1 0 = 1000 Mb/s 0 1 = 100 Mb/s 0 0 = 10 Mb/s	R/W

Table 45–4—PMA/PMD control 1 register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
1.0.5:2	Speed selection	5 4 3 2 1 x x x = Reserved x 1 x x = Reserved 0 0 1 1 = 100 Gb/s 0 0 1 0 = 40 Gb/s 0 0 0 1 = 10PASS-TS/2BASE-TL 0 0 0 0 = 10 Gb/s	R/W
1.0.1	PMA remote loopback	1 = Enable PMA remote loopback mode 0 = Disable PMA remote loopback mode	R/W
1.0.0	PMA local loopback	1 = Enable PMA local loopback mode 0 = Disable PMA local loopback mode	R/W

[a]R/W = Read/Write, SC = Self-clearing, RO = Read only

45.2.1.1.1 Reset (1.0.15)

Resetting a PMA/PMD is accomplished by setting bit 1.0.15 to a one. This action shall set all PMA/PMD registers to their default states. As a consequence, this action may change the internal state of the PMA/PMD and the state of the physical link. This action may also initiate a reset in any other MMDs that are instantiated in the same package. This bit is self-clearing, and a PMA/PMD shall return a value of one in bit 1.0.15 when a reset is in progress; otherwise, it shall return a value of zero. A PMA/PMD is not required to accept a write transaction to any of its registers until the reset process is completed. The control and management interface shall be restored to operation within 0.5 s from the setting of bit 1.0.15. During a reset, a PMD/PMA shall respond to reads from register bits 1.0.15 and 1.8.15:14. All other register bits should be ignored.

NOTE—This operation may interrupt data communication. The data path of a PMD, depending on type and temperature, may take many seconds to run at optimum error ratio after exiting from reset or low-power mode.

45.2.1.1.2 Low power (1.0.11)

A PMA/PMD may be placed into a low-power mode by setting bit 1.0.11 to a one. This action may also initiate a low-power mode in any other MMDs that are instantiated in the same package. The low-power mode is exited by resetting the PMA/PMD. The behavior of the PMA/PMD in transition to and from the low-power mode is implementation specific and any interface signals should not be relied upon. While in the low-power mode, the device shall, as a minimum, respond to management transactions necessary to exit the low-power mode. The default value of bit 1.0.11 is zero.

NOTE—This operation will interrupt data communication. The data path of a PMD, depending on type and temperature, may take many seconds to run at optimum error ratio after exiting from reset or low-power mode.

45.2.1.1.3 Speed selection (1.0.13, 1.0.6, 1.0.5:2)

For devices operating at 10 Mb/s, 100 Mb/s, or 1000 Mb/s the speed of the PMA/PMD may be selected using bits 13 and 6. The speed abilities of the PMA/PMD are advertised in the PMA/PMD speed ability register. These two bits use the same definition as the speed selection bits defined in Clause 22.

For devices not operating at 10 Mb/s, 100 Mb/s, or 1000 Mb/s, the speed of the PMA/PMD may be selected using bits 5 through 2. When bits 5 through 2 are set to 0000 the use of a 10G PMA/PMD is selected. More specific selection is performed using the PMA/PMD control 2 register (Register 1.7) (see 45.2.1.6). The

speed abilities of the PMA/PMD are advertised in the PMA/PMD speed ability register. A PMA/PMD may ignore writes to the PMA/PMD speed selection bits that select speeds it has not advertised in the PMA/PMD speed ability register. It is the responsibility of the STA entity to ensure that mutually acceptable speeds are applied consistently across all the MMDs on a particular PHY.

The PMA/PMD speed selection defaults to a supported ability.

When set to 0001, bits 5:2 select the use of the 10PASS-TS or 2BASE-TL PMA/PMD. More specific mode selection is performed using the 10P/2B PMA control register (45.2.1.15).

When bits 5 through 2 are set to 0010 the use of a 40G PMA/PMD is selected; when set to 0011 the use of a 100G PMA/PMD is selected. More specific selection is performed using the PMA/PMD control 2 register (Register 1.7) (see 45.2.1.6.3).

45.2.1.1.4 PMA remote loopback (1.0.1)

The PMA shall be placed in a remote loopback mode of operation when bit 1.0.1 is set to a one. When bit 1.0.1 is set to a one, the PMA shall accept data on the receive path and return it on the transmit path.

The remote loopback function is optional for all port types, except 2BASE-TL and 10PASS-TS, which do not support loopback. A device's ability to perform the remote loopback function is advertised in the remote loopback ability bit of the related speed-dependent status register. A PMA that is unable to perform the remote loopback function shall ignore writes to this bit and shall return a value of zero when read. For 40/100 Gb/s operation, the remote loopback functionality is detailed in 83.5.9. For 40/100 Gb/s operation, the remote loopback ability bit is specified in the 40G/100G PMA/PMD extended ability register.

The default value of bit 1.0.1 is zero.

45.2.1.1.5 PMA local loopback (1.0.0)

The PMA shall be placed in a local loopback mode of operation when bit 1.0.0 is set to a one. When bit 1.0.0 is set to a one, the PMA shall accept data on the transmit path and return it on the receive path.

The local loopback function is mandatory for the 1000BASE-KX, 10GBASE-KR, 10GBASE-X, 40GBASE-KR4, 40GBASE-CR4, and 100GBASE-CR10 port type and optional for all other port types, except 2BASE-TL, 10PASS-TS, and 10/1GBASE-PRX, which do not support loopback. A device's ability to perform the local loopback function is advertised in the local loopback ability bit of the related speed-dependent status register. A PMA that is unable to perform the local loopback function shall ignore writes to this bit and shall return a value of zero when read. For 10 Gb/s operation, the local loopback functionality is detailed in 48.3.3 and 51.8. For 40/100 Gb/s operation, the local loopback functionality is detailed in 83.5.8. For 10/40/100 Gb/s operation, the local loopback ability bit is specified in the PMA/PMD status 2 register.

The default value of bit 1.0.0 is zero.

NOTE—The signal path through the PMA that is exercised in the loopback mode of operation is implementation specific, but it is recommended that the signal path encompass as much of the PMA circuitry as is practical. The intention of providing this loopback mode of operation is to permit a diagnostic or self-test function to perform the transmission and reception of a PDU, thus testing the transmit and receive data paths. Other loopback signal paths may be enabled using loopback controls within other MMDs.

45.2.1.2 PMA/PMD status 1 register (Register 1.1)

The assignment of bits in the PMA/PMD status 1 register is shown in Table 45–5. All the bits in the PMA/PMD status 1 register are read only; therefore, a write to the PMA/PMD status 1 register shall have no effect.

Table 45–5—PMA/PMD status 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.1.15:10	Reserved	Value always 0	RO
1.1.9	PIASA	PMA ingress AUI stop ability	RO
1.1.8	PEASA	PMA egress AUI stop ability	RO
1.1.7	Fault	1 = Fault condition detected 0 = Fault condition not detected	RO
1.1.6:3	Reserved	Value always 0	RO
1.1.2	Receive link status	1 = PMA/PMD receive link up 0 = PMA/PMD receive link down	RO/LL
1.1.1	Low-power ability	1 = PMA/PMD supports low-power mode 0 = PMA/PMD does not support low-power mode	RO
1.1.0	Reserved	Value always 0	RO

[a]RO = Read only, LL = Latching low

45.2.1.2.1 PMA ingress AUI stop ability (1.1.9)

If bit 1.1.9 is set to one, then the PMA is indicating that the PMA sublayer attached by the ingress AUI is permitted to stop signaling during LPI. If the bit is set to zero, then the PMA is indicating that the PMA sublayer attached by the ingress AUI is not permitted to stop signaling during LPI. If the PMA sublayer attached by the ingress AUI does not support EEE capability or is not capable to stop signaling, then this bit has no effect.

45.2.1.2.2 PMA egress AUI stop ability (1.1.8)

If bit 1.1.8 is set to one, then the PMA is indicating that the PMA sublayer attached by the egress AUI is permitted to stop signaling during LPI. If the bit is set to zero, then the PMA is indicating that the PMA sublayer attached by the egress AUI is not permitted to stop signaling during LPI. If the PMA sublayer attached by the egress AUI does not support EEE capability or is not capable to stop signaling, then this bit has no effect.

45.2.1.2.3 Fault (1.1.7)

Fault is a global PMA/PMD variable. When read as a one, bit 1.1.7 indicates that either (or both) the PMA or the PMD has detected a fault condition on either the transmit or receive paths. When read as a zero, bit 1.1.7 indicates that neither the PMA nor the PMD has detected a fault condition. For 10/40/100 Gb/s operation, bit 1.1.7 is set to a one when either of the fault bits (1.8.11, 1.8.10) located in register 1.8 are set to a one. For 10PASS-TS or 2BASE-TL operations, when read as a one, a fault has been detected and more detailed information is conveyed in 45.2.1.19, 45.2.1.42, 45.2.1.43, and 45.2.1.58.

45.2.1.2.4 Receive link status (1.1.2)

When read as a one, bit 1.1.2 indicates that the PMA/PMD receive link is up. When read as a zero, bit 1.1.2 indicates that the PMA/PMD receive link is down. The receive link status bit shall be implemented with latching low behavior.

While a 10PASS-TS or 2BASE-TL PMA/PMD is initializing, this bit shall indicate receive link down (see 45.2.1.16).

45.2.1.2.5 Low-power ability (1.1.1)

When read as a one, bit 1.1.1 indicates that the PMA/PMD supports the low-power feature. When read as a zero, bit 1.1.1 indicates that the PMA/PMD does not support the low-power feature. If a PMA/PMD supports the low-power feature, then it is controlled using the low-power bit 1.0.11.

45.2.1.3 PMA/PMD device identifier (Registers 1.2 and 1.3)

Registers 1.2 and 1.3 provide a 32-bit value, which may constitute a unique identifier for a particular type of PMA/PMD. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the device manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number.

A PMA/PMD may return a value of zero in each of the 32 bits of the PMA/PMD device identifier to indicate that a unique identifier as described above is not provided.

The format of the PMA/PMD device identifier is specified in 22.2.4.3.1.

NOTE—The use of only 22 bits of the OUI as described here has been deprecated by the IEEE Registration Authority. The definition of vendor-specific device identifiers for other applications is expected to use the full 24 bits to accommodate the use of either an OUI or Company ID.

45.2.1.4 PMA/PMD speed ability (Register 1.4)

The assignment of bits in the PMA/PMD speed ability register is shown in Table 45–6.

Table 45–6—PMA/PMD speed ability register bit definitions

Bit(s)	Name	Description	R/W[a]
1.4.15:10	Reserved for future speeds	Value always 0	RO
1.4.9	100G capable	1 = PMA/PMD is capable of operating at 100 Gb/s 0 = PMA/PMD is not capable of operating as 100 Gb/s	RO
1.4.8	40G capable	1 = PMA/PMD is capable of operating at 40 Gb/s 0 = PMA/PMD is not capable of operating as 40 Gb/s	RO
1.4.7	10/1G capable	1 = PMA/PMD is capable of operating at 10 Gb/s downstream and 1 Gb/s upstream 0 = PMA/PMD is not capable of operating at 10 Gb/s downstream and 1 Gb/s upstream.	RO
1.4.6	10M capable	1 = PMA/PMD is capable of operating at 10 Mb/s 0 = PMA/PMD is not capable of operating as 10 Mb/s	RO
1.4.5	100M capable	1 = PMA/PMD is capable of operating at 100 Mb/s 0 = PMA/PMD is not capable of operating at 100 Mb/s	RO
1.4.4	1000M capable	1 = PMA/PMD is capable of operating at 1000 Mb/s 0 = PMA/PMD is not capable of operating at 1000 Mb/s	RO
1.4.3	Reserved	Value always 0	RO
1.4.2	10PASS-TS capable	1 = PMA/PMD is capable of operating as 10PASS-TS 0 = PMA/PMD is not capable of operating as 10PASS-TS	RO

Table 45–6—PMA/PMD speed ability register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
1.4.1	2BASE-TL capable	1 = PMA/PMD is capable of operating as 2BASE-TL 0 = PMA/PMD is not capable of operating as 2BASE-TL	RO
1.4.0	10G capable	1 = PMA/PMD is capable of operating at 10 Gb/s 0 = PMA/PMD is not capable of operating at 10 Gb/s	RO

[a]RO = Read only

45.2.1.4.1 100G capable (1.4.9)

When read as a one, bit 1.4.9 indicates that the PMA/PMD is able to operate at a data rate of 100 Gb/s. When read as a zero, bit 1.4.9 indicates that the PMA/PMD is not able to operate at a data rate of 100 Gb/s.

45.2.1.4.2 40G capable (1.4.8)

When read as a one, bit 1.4.8 indicates that the PMA/PMD is able to operate at a data rate of 40 Gb/s. When read as a zero, bit 1.4.8 indicates that the PMA/PMD is not able to operate at a data rate of 40 Gb/s.

45.2.1.4.3 10/1G capable (1.4.7)

When read as a one, bit 1.4.7 indicates that the PMA/PMD is able to operate at a data rate of 10 Gb/s in the downstream direction and 1 Gb/s in the upstream direction. When read as a zero, bit 1.4.7 indicates that the PMA/PMD is not able to operate at a data rate of 10 Gb/s in the downstream direction and 1 Gb/s in the upstream direction.

45.2.1.4.4 10M capable (1.4.6)

When read as a one, bit 1.4.6 indicates that the PMA/PMD is able to operate at a data rate of 10 Mb/s. When read as a zero, bit 1.4.6 indicates that the PMA/PMD is not able to operate at a data rate of 10 Mb/s.

45.2.1.4.5 100M capable (1.4.5)

When read as a one, bit 1.4.5 indicates that the PMA/PMD is able to operate at a data rate of 100 Mb/s. When read as a zero, bit 1.4.5 indicates that the PMA/PMD is not able to operate at a data rate of 100 Mb/s.

45.2.1.4.6 1000M capable (1.4.4)

When read as a one, bit 1.4.4 indicates that the PMA/PMD is able to operate at a data rate of 1000 Mb/s. When read as a zero, bit 1.4.4 indicates that the PMA/PMD is not able to operate at a data rate of 1000 Mb/s.

45.2.1.4.7 10PASS-TS capable (1.4.2)

When read as a one, bit 1.4.2 indicates that the PMA/PMD is able to operate as 10PASS-TS. When read as a zero, bit 1.4.2 indicates that the PMA/PMD is not able to operate as 10PASS-TS.

45.2.1.4.8 2BASE-TL capable (1.4.1)

When read as a one, bit 1.4.1 indicates that the PMA/PMD is able to operate as 2BASE-TL. When read as a zero, bit 1.4.1 indicates that the PMA/PMD is not able to operate as 2BASE-TL.

45.2.1.4.9 10G capable (1.4.0)

When read as a one, bit 1.4.0 indicates that the PMA/PMD is able to operate at a data rate of 10 Gb/s. When read as a zero, bit 1.4.0 indicates that the PMA/PMD is not able to operate at a data rate of 10 Gb/s.

45.2.1.5 PMA/PMD devices in package (Registers 1.5 and 1.6)

The PMA/PMD devices in package registers are defined in Table 45–2.

45.2.1.6 PMA/PMD control 2 register (Register 1.7)

The assignment of bits in the PMA/PMD control 2 register is shown in Table 45–7.

45.2.1.6.1 PMA ingress AUI stop enable (1.7.9)

If bit 1.7.9 is set to 1 then the PMA may stop the ingress direction AUI signaling during LPI otherwise it shall keep active signaling on that AUI. If the PMA does not support EEE capability or is not able to stop the ingress direction AUI signaling (see 45.2.1.2.1) then this bit has no effect.

45.2.1.6.2 PMA egress AUI stop enable (1.7.8)

If bit 1.7.8 is set to 1 then the PMA may stop the egress direction AUI signaling during LPI otherwise it shall keep active signaling on that AUI. If the PMA does not support EEE capability or is not able to stop the egress direction AUI signaling (see 45.2.1.2.2) then this bit has no effect.

45.2.1.6.3 PMA/PMD type selection (1.7.5:0)

The PMA/PMD type of the PMA/PMD shall be selected using bits 5 to 0. The PMA/PMD type abilities of the PMA/PMD are advertised in bits 9 and 7 through 0 of the PMA/PMD status 2 register; the PMA/PMD extended ability register; and the 40G/100G PMA/PMD extended ability register. A PMA/PMD shall ignore writes to the PMA/PMD type selection bits that select PMA/PMD types it has not advertised. It is the responsibility of the STA entity to ensure that mutually acceptable MMD types are applied consistently across all the MMDs on a particular PHY.

The PMA/PMD type selection defaults to a supported ability.

45.2.1.7 PMA/PMD status 2 register (Register 1.8)

The assignment of bits in the PMA/PMD status 2 register is shown in Table 45–8. All the bits in the PMA/PMD status 2 register are read only; a write to the PMA/PMD status 2 register shall have no effect.

45.2.1.7.1 Device present (1.8.15:14)

When read as <10>, bits 1.8.15:14 indicate that a device is present and responding at this register address. When read as anything other than <10>, bits 1.8.15:14 indicate that no device is present at this register address or that the device is not functioning properly.

45.2.1.7.2 Transmit fault ability (1.8.13)

When read as a one, bit 1.8.13 indicates that the PMA/PMD has the ability to detect a fault condition on the transmit path. When read as a zero, bit 1.8.13 indicates that the PMA/PMD does not have the ability to detect a fault condition on the transmit path.

Table 45–7—PMA/PMD control 2 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.7.15:10	Reserved	Value always 0	RO
1.7.9	PIASE	PMA ingress AUI stop enable	R/W
1.7.8	PEASE	PMA egress AUI stop enable	R/W
1.7.7:6	Reserved	Value always 0	RO
1.7.5:0	PMA/PMD type selection	5 4 3 2 1 0 1 1 x x x x = reserved 1 0 1 1 1 1 = 100GBASE-SR4 PMA/PMD 1 0 1 1 1 0 = 100GBASE-CR4 PMA/PMD 1 0 1 1 0 1 = 100GBASE-KR4 PMA/PMD 1 0 1 1 0 0 = 100GBASE-KP4 PMA/PMD 1 0 1 0 1 1 = 100GBASE-ER4 PMA/PMD 1 0 1 0 1 0 = 100GBASE-LR4 PMA/PMD 1 0 1 0 0 1 = 100GBASE-SR10 PMA/PMD 1 0 1 0 0 0 = 100GBASE-CR10 PMA/PMD 1 0 0 1 1 x = reserved 1 0 0 1 0 1 = 40GBASE-ER4 PMA/PMD 1 0 0 1 0 0 = 40GBASE-FR PMA/PMD 1 0 0 0 1 1 = 40GBASE-LR4 PMA/PMD 1 0 0 0 1 0 = 40GBASE-SR4 PMA/PMD 1 0 0 0 0 1 = 40GBASE-CR4 PMA/PMD 1 0 0 0 0 0 = 40GBASE-KR4 PMA/PMD 0 1 1 1 1 1 = 10/1GBASE-PRX-U4 0 1 1 1 1 0 = 10GBASE-PR-U4 0 1 1 1 0 1 = 10/1GBASE-PRX-D4 0 1 1 1 0 0 = 10GBASE-PR-D4 0 1 1 0 1 1 = reserved 0 1 1 0 1 0 = 10GBASE-PR-U3 0 1 1 0 0 1 = 10GBASE-PR-U1 0 1 1 0 0 0 = 10/1GBASE-PRX-U3 0 1 0 1 1 1 = 10/1GBASE-PRX-U2 0 1 0 1 1 0 = 10/1GBASE-PRX-U1 0 1 0 1 0 1 = 10GBASE-PR-D3 0 1 0 1 0 0 = 10GBASE-PR-D2 0 1 0 0 1 1 = 10GBASE-PR-D1 0 1 0 0 1 0 = 10/1GBASE-PRX-D3 0 1 0 0 0 1 = 10/1GBASE-PRX-D2 0 1 0 0 0 0 = 10/1GBASE-PRX-D1 0 0 1 1 1 1 = 10BASE-T PMA/PMD 0 0 1 1 1 0 = 100BASE-TX PMA/PMD 0 0 1 1 0 1 = 1000BASE-KX PMA/PMD 0 0 1 1 0 0 = 1000BASE-T PMA/PMD 0 0 1 0 1 1 = 10GBASE-KR PMA/PMD 0 0 1 0 1 0 = 10GBASE-KX4 PMA/PMD 0 0 1 0 0 1 = 10GBASE-T PMA 0 0 1 0 0 0 = 10GBASE-LRM PMA/PMD 0 0 0 1 1 1 = 10GBASE-SR PMA/PMD 0 0 0 1 1 0 = 10GBASE-LR PMA/PMD 0 0 0 1 0 1 = 10GBASE-ER PMA/PMD 0 0 0 1 0 0 = 10GBASE-LX4 PMA/PMD 0 0 0 0 1 1 = 10GBASE-SW PMA/PMD 0 0 0 0 1 0 = 10GBASE-LW PMA/PMD 0 0 0 0 0 1 = 10GBASE-EW PMA/PMD 0 0 0 0 0 0 = 10GBASE-CX4 PMA/PMD	R/W

[a]R/W = Read/Write, RO = Read only

Table 45–8—PMA/PMD status 2 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.8.15:14	Device present	15 14 1 0 = Device responding at this address 1 1 = No device responding at this address 0 1 = No device responding at this address 0 0 = No device responding at this address	RO
1.8.13	Transmit fault ability	1 = PMA/PMD has the ability to detect a fault condition on the transmit path 0 = PMA/PMD does not have the ability to detect a fault condition on the transmit path	RO
1.8.12	Receive fault ability	1 = PMA/PMD has the ability to detect a fault condition on the receive path 0 = PMA/PMD does not have the ability to detect a fault condition on the receive path	RO
1.8.11	Transmit fault	1 = Fault condition on transmit path 0 = No fault condition on transmit path	RO/LH
1.8.10	Receive fault	1 = Fault condition on receive path 0 = No fault condition on receive path	RO/LH
1.8.9	Extended abilities	1 = PMA/PMD has extended abilities listed in register 1.11 0 = PMA/PMD does not have extended abilities	RO
1.8.8	PMD transmit disable ability	1 = PMD has the ability to disable the transmit path 0 = PMD does not have the ability to disable the transmit path	RO
1.8.7	10GBASE-SR ability	1 = PMA/PMD is able to perform 10GBASE-SR 0 = PMA/PMD is not able to perform 10GBASE-SR	RO
1.8.6	10GBASE-LR ability	1 = PMA/PMD is able to perform 10GBASE-LR 0 = PMA/PMD is not able to perform 10GBASE-LR	RO
1.8.5	10GBASE-ER ability	1 = PMA/PMD is able to perform 10GBASE-ER 0 = PMA/PMD is not able to perform 10GBASE-ER	RO
1.8.4	10GBASE-LX4 ability	1 = PMA/PMD is able to perform 10GBASE-LX4 0 = PMA/PMD is not able to perform 10GBASE-LX4	RO
1.8.3	10GBASE-SW ability	1 = PMA/PMD is able to perform 10GBASE-SW 0 = PMA/PMD is not able to perform 10GBASE-SW	RO
1.8.2	10GBASE-LW ability	1 = PMA/PMD is able to perform 10GBASE-LW 0 = PMA/PMD is not able to perform 10GBASE-LW	RO
1.8.1	10GBASE-EW ability	1 = PMA/PMD is able to perform 10GBASE-EW 0 = PMA/PMD is not able to perform 10GBASE-EW	RO
1.8.0	PMA local loopback ability	1 = PMA has the ability to perform a local loopback function 0 = PMA does not have the ability to perform a local loopback function	RO

[a]RO = Read only, LH = Latching high

45.2.1.7.3 Receive fault ability (1.8.12)

When read as a one, bit 1.8.12 indicates that the PMA/PMD has the ability to detect a fault condition on the receive path. When read as a zero, bit 1.8.12 indicates that the PMA/PMD does not have the ability to detect a fault condition on the receive path.

45.2.1.7.4 Transmit fault (1.8.11)

When read as a one, bit 1.8.11 indicates that the PMA/PMD has detected a fault condition on the transmit path. When read as a zero, bit 1.8.11 indicates that the PMA/PMD has not detected a fault condition on the transmit path. Detection of a fault condition on the transmit path is optional and the ability to detect such a condition is advertised by bit 1.8.13. A PMA/PMD that is unable to detect a fault condition on the transmit path shall return a value of zero for this bit. The description of the transmit fault function for the various PMA/PMDs is given in Table 45–9.

Table 45–9—Transmit fault description location

PMA/PMD	Description location
10GBASE-KR	72.6.8
10GBASE-LRM	68.4.8
10GBASE-S, 10GBASE-L, 10GBASE-E	52.4.8
10GBASE-LX4	53.4.10
10GBASE-CX4	54.5.10
10GBASE-T	55.4.2.2
10GBASE-KX4	71.6.10
40GBASE-KR4	84.7.10
40GBASE-CR4, 100GBASE-CR10	85.7.10
40GBASE-SR4, 100GBASE-SR10	86.5.10
40GBASE-LR4, 40GBASE-ER4	87.5.10
40GBASE-FR	89.5.8
100GBASE-KP4	94.3.8
100GBASE-KR4	93.7.10
100GBASE-CR4	92.7.10
100GBASE-SR4	95.5.10
100GBASE-LR4, 100GBASE-ER4	88.5.10

The transmit fault bit shall be implemented with latching high behavior.

The default value of bit 1.8.11 is zero.

45.2.1.7.5 Receive fault (1.8.10)

When read as a one, bit 1.8.10 indicates that the PMA/PMD has detected a fault condition on the receive path. When read as a zero, bit 1.8.10 indicates that the PMA/PMD has not detected a fault condition on the

receive path. Detection of a fault condition on the receive path is optional and the ability to detect such a condition is advertised by bit 1.8.12. A PMA/PMD that is unable to detect a fault condition on the receive path shall return a value of zero for this bit. The description of the receive fault function for the various PMA/PMDs is given in Table 45–10.

Table 45–10—Receive fault description location

PMA/PMD	Description location
10GBASE-KR	72.6.9
10GBASE-LRM	68.4.9
10GBASE-S, 10GBASE-L, 10GBASE-E	52.4.9
10GBASE-LX4	53.4.11
10GBASE-CX4	54.5.11
10GBASE-T	55.4.2.4
10GBASE-KX4	71.6.11
40GBASE-KR4	84.7.11
40GBASE-CR4, 100GBASE-CR10	85.7.11
40GBASE-SR4, 100GBASE-SR10	86.5.11
40GBASE-LR4, 40GBASE-ER4	87.5.11
40GBASE-FR	89.5.9
100GBASE-KP4	94.3.9
100GBASE-KR4	93.7.11
100GBASE-CR4	92.7.11
100GBASE-SR4	95.5.11
100GBASE-LR4, 100GBASE-ER4	88.5.11

The receive fault bit shall be implemented with latching high behavior.

45.2.1.7.6 PMA/PMD extended abilities (1.8.9)

When read as a one, bit 1.8.9 indicates that the PMA/PMD has extended abilities listed in register 1.11. When read as a zero, bit 1.8.9 indicates that the PMA/PMD does not have extended abilities.

45.2.1.7.7 PMD transmit disable ability (1.8.8)

When read as a one, bit 1.8.8 indicates that the PMD is able to perform the transmit disable function. When read as a zero, bit 1.8.8 indicates that the PMD is not able to perform the transmit disable function. If a PMD is able to perform the transmit disable function, then it is controlled using the PMD transmit disable register.

45.2.1.7.8 10GBASE-SR ability (1.8.7)

When read as a one, bit 1.8.7 indicates that the PMA/PMD is able to support a 10GBASE-SR PMA/PMD type. When read as a zero, bit 1.8.7 indicates that the PMA/PMD is not able to support a 10GBASE-SR PMA/PMD type.

45.2.1.7.9 10GBASE-LR ability (1.8.6)

When read as a one, bit 1.8.6 indicates that the PMA/PMD is able to support a 10GBASE-LR PMA/PMD type. When read as a zero, bit 1.8.6 indicates that the PMA/PMD is not able to support a 10GBASE-LR PMA/PMD type.

45.2.1.7.10 10GBASE-ER ability (1.8.5)

When read as a one, bit 1.8.5 indicates that the PMA/PMD is able to support a 10GBASE-ER PMA/PMD type. When read as a zero, bit 1.8.5 indicates that the PMA/PMD is not able to support a 10GBASE-ER PMA/PMD type.

45.2.1.7.11 10GBASE-LX4 ability (1.8.4)

When read as a one, bit 1.8.4 indicates that the PMA/PMD is able to support a 10GBASE-LX4 PMA/PMD type. When read as a zero, bit 1.8.4 indicates that the PMA/PMD is not able to support a 10GBASE-LX4 PMA/PMD type.

45.2.1.7.12 10GBASE-SW ability (1.8.3)

When read as a one, bit 1.8.3 indicates that the PMA/PMD is able to support a 10GBASE-SW PMA/PMD type. When read as a zero, bit 1.8.3 indicates that the PMA/PMD is not able to support a 10GBASE-SW PMA/PMD type.

45.2.1.7.13 10GBASE-LW ability (1.8.2)

When read as a one, bit 1.8.2 indicates that the PMA/PMD is able to support a 10GBASE-LW PMA/PMD type. When read as a zero, bit 1.8.2 indicates that the PMA/PMD is not able to support a 10GBASE-LW PMA/PMD type.

45.2.1.7.14 10GBASE-EW ability (1.8.1)

When read as a one, bit 1.8.1 indicates that the PMA/PMD is able to support a 10GBASE-EW PMA/PMD type. When read as a zero, bit 1.8.1 indicates that the PMA/PMD is not able to support a 10GBASE-EW PMA/PMD type.

45.2.1.7.15 PMA local loopback ability (1.8.0)

When read as a one, bit 1.8.0 indicates that the PMA is able to perform the local loopback function. When read as a zero, bit 1.8.0 indicates that the PMA is not able to perform the local loopback function. If a PMA is able to perform the local loopback function, then it is controlled using the PMA local loopback bit 1.0.0.

45.2.1.8 PMD transmit disable register (Register 1.9)

The assignment of bits in the PMD transmit disable register is shown in Table 45–11. The transmit disable functionality is optional and a PMD's ability to perform the transmit disable functionality is advertised in the PMD transmit disable ability bit 1.8.8. A PMD that does not implement the transmit disable functionality shall ignore writes to the PMD transmit disable register and may return a value of zero for all bits. A PMD

device that operates using a single lane and has implemented the transmit disable function shall use bit 1.9.0 to control the function. Such devices shall ignore writes to bits 1.9.10:1 and return a value of zero for those bits when they are read. The description of the transmit disable function for the various PMA/PMDs is given in Table 45–12.

NOTE—Disabling the transmitter on one or more lanes stops the entire link from carrying data.

Table 45–11—PMD transmit disable register bit definitions

Bit(s)	Name	Description	R/W[a]
1.9.15:11	Reserved	Value always 0	RO
1.9.10	PMD transmit disable 9	1 = Disable output on transmit lane 9 0 = Enable output on transmit lane 9	R/W
1.9.9	PMD transmit disable 8	1 = Disable output on transmit lane 8 0 = Enable output on transmit lane 8	R/W
1.9.8	PMD transmit disable 7	1 = Disable output on transmit lane 7 0 = Enable output on transmit lane 7	R/W
1.9.7	PMD transmit disable 6	1 = Disable output on transmit lane 6 0 = Enable output on transmit lane 6	R/W
1.9.6	PMD transmit disable 5	1 = Disable output on transmit lane 5 0 = Enable output on transmit lane 5	R/W
1.9.5	PMD transmit disable 4	1 = Disable output on transmit lane 4 0 = Enable output on transmit lane 4	R/W
1.9.4	PMD transmit disable 3	1 = Disable output on transmit lane 3 0 = Enable output on transmit lane 3	R/W
1.9.3	PMD transmit disable 2	1 = Disable output on transmit lane 2 0 = Enable output on transmit lane 2	R/W
1.9.2	PMD transmit disable 1	1 = Disable output on transmit lane 1 0 = Enable output on transmit lane 1	R/W
1.9.1	PMD transmit disable 0	1 = Disable output on transmit lane 0 0 = Enable output on transmit lane 0	R/W
1.9.0	Global PMD transmit disable	1 = Disable transmitter output 0 = Enable transmitter output	R/W

[a]RO = Read only, R/W = Read/Write

45.2.1.8.1 PMD transmit disable 9 (1.9.10)

When bit 1.9.10 is set to a one, the PMD shall disable output on lane 9 of the transmit path. When bit 1.9.10 is set to zero, the PMD shall enable output on lane 9 of the transmit path.

The default value for bit 1.9.10 is zero.

NOTE—Transmission will not be enabled when this bit is set to zero unless the global PMD transmit disable bit is also zero.

Table 45–12—Transmit disable description location

PMA/PMD	Description location
10GBASE-KR	72.6.5
10GBASE-LRM	68.4.7
Other 10GBASE-R	52.4.7
10GBASE-LX4	53.4.7
10GBASE-CX4	54.5.6
10GBASE-T	55.4.2.3
10GBASE-KX4	71.6.6
40GBASE-KR4	84.7.6
40GBASE-CR4 and 100GBASE-CR10	85.7.6
40GBASE-SR4 and 100GBASE-SR10	86.5.7
40GBASE-LR4 and 40GBASE-ER4	87.5.7
40GBASE-FR	89.5.6
100GBASE-KP4	94.3.6.6
100GBASE-KR4	93.7.6
100GBASE-CR4	92.7.6
100GBASE-SR4	95.5.7
100GBASE-LR4 and 100GBASE-ER4	88.5.7

45.2.1.8.2 PMD transmit disable 4, 5, 6, 7, 8 (1.9.5, 1.9.6, 1.9.7, 1.9.8, 1.9.9)

These bits are defined similarly to bit 1.9.10 for lanes 4, 5, 6, 7, and 8, respectively.

45.2.1.8.3 PMD transmit disable 3 (1.9.4)

When bit 1.9.4 is set to a one, the PMD shall disable output on lane 3 of the transmit path. When bit 1.9.4 is set to a zero, the PMD shall enable output on lane 3 of the transmit path.

The default value for bit 1.9.4 is zero.

NOTE—Transmission will not be enabled when this bit is set to a zero unless the global PMD transmit disable bit is also zero.

45.2.1.8.4 PMD transmit disable 2 (1.9.3)

When bit 1.9.3 is set to a one, the PMD shall disable output on lane 2 of the transmit path. When bit 1.9.3 is set to a zero, the PMD shall enable output on lane 2 of the transmit path.

The default value for bit 1.9.3 is zero.

NOTE—Transmission will not be enabled when this bit is set to a zero unless the global PMD transmit disable bit is also zero.

45.2.1.8.5 PMD transmit disable 1 (1.9.2)

When bit 1.9.2 is set to a one, the PMD shall disable output on lane 1 of the transmit path. When bit 1.9.2 is set to a zero, the PMD shall enable output on lane 1 of the transmit path.

The default value for bit 1.9.2 is zero.

NOTE—Transmission will not be enabled when this bit is set to a zero unless the global PMD transmit disable bit is also zero.

45.2.1.8.6 PMD transmit disable 0 (1.9.1)

When bit 1.9.1 is set to a one, the PMD shall disable output on lane 0 of the transmit path. When bit 1.9.1 is set to a zero, the PMD shall enable output on lane 0 of the transmit path.

The default value for bit 1.9.1 is zero.

NOTE—Transmission will not be enabled when this bit is set to a zero unless the global PMD transmit disable bit is also zero.

45.2.1.8.7 Global PMD transmit disable (1.9.0)

When bit 1.9.0 is set to a one, the PMD shall disable output on the transmit path. When bit 1.9.0 is set to a zero, the PMD shall enable output on the transmit path.

For single wavelength PMD types, transmission will be disabled when this bit is set to one. When this bit is set to zero, transmission is enabled.

For multiple wavelength or lane PMD types, transmission will be disabled on all lanes when this bit is set to one. When this bit is set to zero, the lanes are individually controlled by their corresponding transmit disable bits 1.9.4:1.

The default value for bit 1.9.0 is zero.

45.2.1.9 PMD receive signal detect register (Register 1.10)

The assignment of bits in the PMD receive signal detect register is shown in Table 45–13. The 10G PMD receive signal detect register is mandatory. PMD types that use only a single lane indicate the status of the receive signal detect using bit 1.10.0 and return a value of zero for bits 1.10.10:1. PMD types that use multiple wavelengths or lanes indicate the status of each lane in bits 1.10.10:1 and the logical AND of those bits in bit 1.10.0.

Table 45–13—PMD receive signal detect register bit definitions

Bit(s)	Name	Description	R/W[a]
1.10.15:11	Reserved	Value always 0	RO
1.10.10	PMD receive signal detect 9	1 = Signal detected on receive lane 9 0 = Signal not detected on receive lane 9	RO
1.10.9	PMD receive signal detect 8	1 = Signal detected on receive lane 8 0 = Signal not detected on receive lane 8	RO

Table 45–13—PMD receive signal detect register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
1.10.8	PMD receive signal detect 7	1 = Signal detected on receive lane 7 0 = Signal not detected on receive lane 7	RO
1.10.7	PMD receive signal detect 6	1 = Signal detected on receive lane 6 0 = Signal not detected on receive lane 6	RO
1.10.6	PMD receive signal detect 5	1 = Signal detected on receive lane 5 0 = Signal not detected on receive lane 5	RO
1.10.5	PMD receive signal detect 4	1 = Signal detected on receive lane 4 0 = Signal not detected on receive lane 4	RO
1.10.4	PMD receive signal detect 3	1 = Signal detected on receive lane 3 0 = Signal not detected on receive lane 3	RO
1.10.3	PMD receive signal detect 2	1 = Signal detected on receive lane 2 0 = Signal not detected on receive lane 2	RO
1.10.2	PMD receive signal detect 1	1 = Signal detected on receive lane 1 0 = Signal not detected on receive lane 1	RO
1.10.1	PMD receive signal detect 0	1 = Signal detected on receive lane 0 0 = Signal not detected on receive lane 0	RO
1.10.0	Global PMD receive signal detect	1 = Signal detected on receive 0 = Signal not detected on receive	RO

[a]RO = Read only

45.2.1.9.1 PMD receive signal detect 9 (1.10.10)

When bit 1.10.10 is read as a one, a signal has been detected on lane 9 of the PMD receive path. When bit 1.10.10 is read as a zero, a signal has not been detected on lane 9 of the PMD receive path.

45.2.1.9.2 PMD receive signal detect 4, 5, 6, 7, 8 (1.10.5, 1.10.6, 1.10.7, 1.10.8, 1.10.9)

These bits are defined similarly to bit 1.10.10 for lanes 4, 5, 6, 7, and 8, respectively.

45.2.1.9.3 PMD receive signal detect 3 (1.10.4)

When bit 1.10.4 is read as a one, a signal has been detected on lane 3 of the PMD receive path. When bit 1.10.4 is read as a zero, a signal has not been detected on lane 3 of the PMD receive path.

45.2.1.9.4 PMD receive signal detect 2 (1.10.3)

When bit 1.10.3 is read as a one, a signal has been detected on lane 2 of the PMD receive path. When bit 1.10.3 is read as a zero, a signal has not been detected on lane 2 of the PMD receive path.

45.2.1.9.5 PMD receive signal detect 1 (1.10.2)

When bit 1.10.2 is read as a one, a signal has been detected on lane 1 of the PMD receive path. When bit 1.10.2 is read as a zero, a signal has not been detected on lane 1 of the PMD receive path.

45.2.1.9.6 PMD receive signal detect 0 (1.10.1)

When bit 1.10.1 is read as a one, a signal has been detected on lane 0 of the PMD receive path. When bit 1.10.1 is read as a zero, a signal has not been detected on lane 0 of the PMD receive path.

45.2.1.9.7 Global PMD receive signal detect (1.10.0)

When bit 1.10.0 is read as a one, a signal has been detected on all the PMD receive paths. When bit 1.10.0 is read as a zero, a signal has not been detected on at least one of the PMD receive paths.

Single wavelength PMD types indicate the status of their receive path signal using this bit.

Multiple wavelength or multiple lane PMD types indicate the global status of the lane-by-lane signal detect indications using this bit. This bit is read as a one when all the lane signal detect indications are one; otherwise, this bit is read as a zero.

45.2.1.10 PMA/PMD extended ability register (Register 1.11)

The assignment of bits in the PMA/PMD extended ability register is shown in Table 45–14. All of the bits in the PMA/PMD extended ability register are read only; a write to the PMA/PMD extended ability register shall have no effect.

Table 45–14—PMA/PMD Extended Ability register bit definitions

Bit(s)	Name	Description	R/W[a]
1.11.15:11	Reserved	Value always 0	RO
1.11.10	40G/100G extended abilities	1 = PMA/PMD has 40G/100G extended abilities listed in register 1.13 0 = PMA/PMD does not have 40G/100G extended abilities	RO
1.11.9	P2MP ability	1 = PMA/PMD has P2MP abilities listed in register 1.12 0 = PMA/PMD does not have P2MP abilities	RO
1.11.8	10BASE-T ability	1 = PMA/PMD is able to perform 10BASE-T 0 = PMA/PMD is not able to perform 10BASE-T	RO
1.11.7	100BASE-TX ability	1 = PMA/PMD is able to perform 100BASE-TX 0 = PMA/PMD is not able to perform 100BASE-TX	RO
1.11.6	1000BASE-KX ability	1 = PMA/PMD is able to perform 1000BASE-KX 0 = PMA/PMD is not able to perform 1000BASE-KX	RO
1.11.5	1000BASE-T ability	1 = PMA/PMD is able to perform 1000BASE-T 0 = PMA/PMD is not able to perform 1000BASE-T	RO
1.11.4	10GBASE-KR ability	1 = PMA/PMD is able to perform 10GBASE-KR 0 = PMA/PMD is not able to perform 10GBASE-KR	RO
1.11.3	10GBASE-KX4 ability	1 = PMA/PMD is able to perform 10GBASE-KX4 0 = PMA/PMD is not able to perform 10GBASE-KX4	RO

Table 45–14—PMA/PMD Extended Ability register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
1.11.2	10GBASE-T ability	1 = PMA/PMD is able to perform 10GBASE-T 0 = PMA/PMD is not able to perform 10GBASE-T	RO
1.11.1	10GBASE-LRM ability	1 = PMA/PMD is able to perform 10GBASE-LRM 0 = PMA/PMD is not able to perform 10GBASE-LRM	RO
1.11.0	10GBASE-CX4 ability	1 = PMA/PMD is able to perform 10GBASE-CX4 0 = PMA/PMD is not able to perform 10GBASE-CX4	RO

[a]RO = Read only

45.2.1.10.1 40G/100G extended abilities (1.11.10)

When read as a one, bit 1.11.10 indicates that the PMA/PMD has 40G/100G extended abilities listed in register 1.13. When read as a zero, bit 1.11.10 indicates that the PMA/PMD does not have 40G/100G extended abilities.

45.2.1.10.2 P2MP ability (1.11.9)

When read as a one, bit 1.11.9 indicates that the PMA/PMD has P2MP abilities listed in register 1.12. When read as a zero, bit 1.11.9 indicates that the PMA/PMD does not have P2MP abilities.

45.2.1.10.3 10BASE-T ability (1.11.8)

When read as a one, bit 1.11.8 indicates that the PMA/PMD is able to operate as a 10BASE-T PMA/PMD type. When read as a zero, bit 1.11.8 indicates that the PMA/PMD is not able to operate as a 10BASE-T PMA/PMD type.

45.2.1.10.4 100BASE-TX ability (1.11.7)

When read as a one, bit 1.11.7 indicates that the PMA/PMD is able to operate as a 100BASE-TX PMA/PMD type. When read as a zero, bit 1.11.7 indicates that the PMA/PMD is not able to operate as a 100BASE-TX PMA/PMD type.

45.2.1.10.5 1000BASE-KX ability (1.11.6)

When read as a one, bit 1.11.6 indicates that the PMA/PMD is able to operate as 1000BASE-KX. When read as a zero, bit 1.11.6 indicates that the PMA/PMD is not able to operate as 1000BASE-KX.

45.2.1.10.6 1000BASE-T ability (1.11.5)

When read as a one, bit 1.11.5 indicates that the PMA/PMD is able to operate as a 1000BASE-T PMA/PMD type. When read as a zero, bit 1.11.5 indicates that the PMA/PMD is not able to operate as a 1000BASE-T PMA/PMD type.

45.2.1.10.7 10GBASE-KR ability (1.11.4)

When read as a one, bit 1.11.4 indicates that the PMA/PMD is able to operate as 10GBASE-KR. When read as a zero, bit 1.11.4 indicates that the PMA/PMD is not able to operate as 10GBASE-KR.

45.2.1.10.8 10GBASE-KX4 ability (1.11.3)

When read as a one, bit 1.11.3 indicates that the PMA/PMD is able to operate as 10GBASE-KX4. When read as a zero, bit 1.11.3 indicates that the PMA/PMD is not able to operate as 10GBASE-KX4.

45.2.1.10.9 10GBASE-T ability (1.11.2)

When read as a one, bit 1.11.2 indicates that the PMA/PMD is able to operate as a 10GBASE-T PMA/PMD type. When read as a zero, bit 1.11.2 indicates that the PMA/PMD is not able to operate as a 10GBASE-T PMA/PMD type.

45.2.1.10.10 10GBASE-LRM ability (1.11.1)

When read as a one, bit 1.11.1 indicates that the PMA/PMD is able to operate as 10GBASE-LRM. When read as a zero, bit 1.11.1 indicates that the PMA/PMD is not able to operate as 10GBASE-LRM.

45.2.1.10.11 10GBASE-CX4 ability (1.11.0)

When read as a one, bit 1.11.0 indicates that the PMA/PMD is able to operate as a 10GBASE-CX4 PMA/PMD type. When read as a zero, bit 1.11.0 indicates that the PMA/PMD is not able to operate as a 10GBASE-CX4 PMA/PMD type.

45.2.1.11 10G-EPON PMA/PMD ability register (Register 1.12)

The assignment of bits in the 10G-EPON PMA/PMD ability register is shown in Table 45–15. All of the bits in the 10G-EPON PMA/PMD ability register are read only; a write to the 10G-EPON PMA/PMD ability register shall have no effect.

Table 45–15—10G-EPON PMA/PMD ability register bit definitions

Bit(s)	Name	Description	R/W[a]
1.12.15	Reserved	Value always 0	RO
1.12.14	10GBASE-PR-D4 ability	1 = PMA/PMD is able to perform 10GBASE-PR-D4 0 = PMA/PMD is not able to perform 10GBASE-PR-D4	RO
1.12.13	10GBASE-PR-U4 ability	1 = PMA/PMD is able to perform 10GBASE-PR-U4 0 = PMA/PMD is not able to perform 10GBASE-PR-U4	RO
1.12.12	10/1GBASE-PRX-D4 ability	1 = PMA/PMD is able to perform 10/1GBASE-PRX-D4 0 = PMA/PMD is not able to perform 10/1GBASE-PRX-D4	RO
1.12.11	10/1GBASE-PRX-U4 ability	1 = PMA/PMD is able to perform 10/1GBASE-PRX-U4 0 = PMA/PMD is not able to perform 10/1GBASE-PRX-U4	RO
1.12.10	10/1GBASE-PRX-D1 ability	1 = PMA/PMD is able to perform 10/1GBASE-PRX-D1 0 = PMA/PMD is not able to perform 10/1GBASE-PRX-D1	RO
1.12.9	10/1GBASE-PRX-D2 ability	1 = PMA/PMD is able to perform 10/1GBASE-PRX-D2 0 = PMA/PMD is not able to perform 10/1GBASE-PRX-D2	RO
1.12.8	10/1GBASE-PRX-D3 ability	1 = PMA/PMD is able to perform 10/1GBASE-PRX-D3 0 = PMA/PMD is not able to perform 10/1GBASE-PRX-D3	RO
1.12.7	10GBASE-PR-D1 ability	1 = PMA/PMD is able to perform 10GBASE-PR-D1 0 = PMA/PMD is not able to perform 10GBASE-PR-D1	RO

Table 45–15—10G-EPON PMA/PMD ability register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
1.12.6	10GBASE-PR-D2 ability	1 = PMA/PMD is able to perform 10GBASE-PR-D2 0 = PMA/PMD is not able to perform 10GBASE-PR-D2	RO
1.12.5	10GBASE-PR-D3 ability	1 = PMA/PMD is able to perform 10GBASE-PR-D3 0 = PMA/PMD is not able to perform 10GBASE-PR-D3	RO
1.12.4	10/1GBASE-PRX-U1 ability	1 = PMA/PMD is able to perform 10/1GBASE-PRX-U1 0 = PMA/PMD is not able to perform 10/1GBASE-PRX-U1	RO
1.12.3	10/1GBASE-PRX-U2 ability	1 = PMA/PMD is able to perform 10/1GBASE-PRX-U2 0 = PMA/PMD is not able to perform 10/1GBASE-PRX-U2	RO
1.12.2	10/1GBASE-PRX-U3 ability	1 = PMA/PMD is able to perform 10/1GBASE-PRX-U3 0 = PMA/PMD is not able to perform 10/1GBASE-PRX-U3	RO
1.12.1	10GBASE-PR-U1 ability	1 = PMA/PMD is able to perform 10GBASE-PR-U1 0 = PMA/PMD is not able to perform 10GBASE-PR-U1	RO
1.12.0	10GBASE-PR-U3 ability	1 = PMA/PMD is able to perform 10GBASE-PR-U3 0 = PMA/PMD is not able to perform 10GBASE-PR-U3	RO

[a]RO = Read only

45.2.1.11.1 10GBASE-PR-D4 ability (1.12.14)

When read as a one, bit 1.12.14 indicates that the PMA/PMD is able to operate as a 10GBASE-PR-D4 PMA/PMD type. When read as a zero, bit 1.12.14 indicates that the PMA/PMD is not able to operate as a 10GBASE-PR-D4 PMA/PMD type.

45.2.1.11.2 10GBASE-PR-U4 ability (1.12.13)

When read as a one, bit 1.12.13 indicates that the PMA/PMD is able to operate as a 10GBASE-PR-U4 PMA/PMD type. When read as a zero, bit 1.12.13 indicates that the PMA/PMD is not able to operate as a 10GBASE-PR-U4 PMA/PMD type.

45.2.1.11.3 10/1GBASE-PRX-D4 ability (1.12.12)

When read as a one, bit 1.12.12 indicates that the PMA/PMD is able to operate as a 10/1GBASE-PRX-D4 PMA/PMD type. When read as a zero, bit 1.12.12 indicates that the PMA/PMD is not able to operate as a 10/1GBASE-PRX-D4 PMA/PMD type.

45.2.1.11.4 10/1GBASE-PRX-U4 ability (1.12.11)

When read as a one, bit 1.12.11 indicates that the PMA/PMD is able to operate as a 10/1GBASE-PRX-U4 PMA/PMD type. When read as a zero, bit 1.12.11 indicates that the PMA/PMD is not able to operate as a 10/1GBASE-PRX-U4 PMA/PMD type.

45.2.1.11.5 10/1GBASE-PRX-D1 ability (1.12.10)

When read as a one, bit 1.12.10 indicates that the PMA/PMD is able to operate as a 10/1GBASE-PRX-D1 PMA/PMD type. When read as a zero, bit 1.12.10 indicates that the PMA/PMD is not able to operate as a 10/1GBASE-PRX-D1 PMA/PMD type.

45.2.1.11.6 10/1GBASE-PRX-D2 ability (1.12.9)

When read as a one, bit 1.12.9 indicates that the PMA/PMD is able to operate as a 10/1GBASE-PRX-D2 PMA/PMD type. When read as a zero, bit 1.12.9 indicates that the PMA/PMD is not able to operate as a 10/1GBASE-PRX-D2 PMA/PMD type.

45.2.1.11.7 10/1GBASE-PRX-D3 ability (1.12.8)

When read as a one, bit 1.12.8 indicates that the PMA/PMD is able to operate as a 10/1GBASE-PRX-D3 PMA/PMD type. When read as a zero, bit 1.12.8 indicates that the PMA/PMD is not able to operate as a 10/1GBASE-PRX-D3 PMA/PMD type.

45.2.1.11.8 10GBASE-PR-D1 ability (1.12.7)

When read as a one, bit 1.12.7 indicates that the PMA/PMD is able to operate as a 10GBASE-PR-D1 PMA/PMD type. When read as a zero, bit 1.12.7 indicates that the PMA/PMD is not able to operate as a 10GBASE-PR-D1 PMA/PMD type.

45.2.1.11.9 10GBASE-PR-D2 ability (1.12.6)

When read as a one, bit 1.12.6 indicates that the PMA/PMD is able to operate as a 10GBASE-PR-D2 PMA/PMD type. When read as a zero, bit 1.12.6 indicates that the PMA/PMD is not able to operate as a 10GBASE-PR-D2 PMA/PMD type.

45.2.1.11.10 10GBASE-PR-D3 ability (1.12.5)

When read as a one, bit 1.12.5 indicates that the PMA/PMD is able to operate as a 10GBASE-PR-D3 PMA/PMD type. When read as a zero, bit 1.12.5 indicates that the PMA/PMD is not able to operate as a 10GBASE-PR-D3 PMA/PMD type.

45.2.1.11.11 10/1GBASE-PRX-U1 ability (1.12.4)

When read as a one, bit 1.12.4 indicates that the PMA/PMD is able to operate as a 10/1GBASE-PRX-U1 PMA/PMD type. When read as a zero, bit 1.12.4 indicates that the PMA/PMD is not able to operate as a 10/1GBASE-PRX-U1 PMA/PMD type.

45.2.1.11.12 10/1GBASE-PRX-U2 ability (1.12.3)

When read as a one, bit 1.12.3 indicates that the PMA/PMD is able to operate as a 10/1GBASE-PRX-U2 PMA/PMD type. When read as a zero, bit 1.12.3 indicates that the PMA/PMD is not able to operate as a 10/1GBASE-PRX-U2 PMA/PMD type.

45.2.1.11.13 10/1GBASE-PRX-U3 ability (1.12.2)

When read as a one, bit 1.12.2 indicates that the PMA/PMD is able to operate as a 10/1GBASE-PRX-U3 PMA/PMD type. When read as a zero, bit 1.12.2 indicates that the PMA/PMD is not able to operate as a 10/1GBASE-PRX-U3 PMA/PMD type.

45.2.1.11.14 10GBASE-PR-U1 ability (1.12.1)

When read as a one, bit 1.12.1 indicates that the PMA/PMD is able to operate as a 10GBASE-PR-U1 PMA/PMD type. When read as a zero, bit 1.12.1 indicates that the PMA/PMD is not able to operate as a 10GBASE-PR-U1 PMA/PMD type.

45.2.1.11.15 10GBASE-PR-U3 ability (1.12.0)

When read as a one, bit 1.12.0 indicates that the PMA/PMD is able to operate as a 10GBASE-PR-U3 PMA/PMD type. When read as a zero, bit 1.12.0 indicates that the PMA/PMD is not able to operate as a 10GBASE-PR-U3 PMA/PMD type.

45.2.1.12 40G/100G PMA/PMD extended ability register (Register 1.13)

The assignment of bits in the 40G/100G PMA/PMD extended ability register is shown in Table 45–16. All of the bits in the PMA/PMD extended ability register are read only; a write to the PMA/PMD extended ability register shall have no effect.

Table 45–16—40G/100G PMA/PMD extended ability register bit definitions

Bit(s)	Name	Description	R/W[a]
1.13.15	PMA remote loopback ability	1 = PMA has the ability to perform a remote loopback function 0 = PMA does not have the ability to perform a remote loopback function	RO
1.13.14	100GBASE-CR4 ability	1 = PMA/PMD is able to perform 100GBASE-CR4 0 = PMA/PMD is not able to perform 100GBASE-CR4	RO
1.13.13	100GBASE-KR4 ability	1 = PMA/PMD is able to perform 100GBASE-KR4 0 = PMA/PMD is not able to perform 100GBASE-KR4	RO
1.13.12	100GBASE-KP4 ability	1 = PMA/PMD is able to perform 100GBASE-KP4 0 = PMA/PMD is not able to perform 100GBASE-KP4	RO
1.13.11	100GBASE-ER4 ability	1 = PMA/PMD is able to perform 100GBASE-ER4 0 = PMA/PMD is not able to perform 100GBASE-ER4	RO
1.13.10	100GBASE-LR4 ability	1 = PMA/PMD is able to perform 100GBASE-LR4 0 = PMA/PMD is not able to perform 100GBASE-LR4	RO
1.13.9	100GBASE-SR10 ability	1 = PMA/PMD is able to perform 100GBASE-SR10 0 = PMA/PMD is not able to perform 100GBASE-SR10	RO
1.13.8	100GBASE-CR10 ability	1 = PMA/PMD is able to perform 100GBASE-CR10 0 = PMA/PMD is not able to perform 100GBASE-CR10	RO
1.13.7	100GBASE-SR4 ability	1 = PMA/PMD is able to perform 100GBASE-SR4 0 = PMA/PMD is not able to perform 100GBASE-SR4	RO
1.13.6	Reserved	Value always 0	RO
1.13.5	40GBASE-ER4 ability	1 = PMA/PMD is able to perform 40GBASE-ER4 0 = PMA/PMD is not able to perform 40GBASE-ER4	RO
1.13.4	40GBASE-FR ability	1 = PMA/PMD is able to perform 40GBASE-FR 0 = PMA/PMD is not able to perform 40GBASE-FR	RO
1.13.3	40GBASE-LR4 ability	1 = PMA/PMD is able to perform 40GBASE-LR4 0 = PMA/PMD is not able to perform 40GBASE-LR4	RO
1.13.2	40GBASE-SR4 ability	1 = PMA/PMD is able to perform 40GBASE-SR4 0 = PMA/PMD is not able to perform 40GBASE-SR4	RO
1.13.1	40GBASE-CR4 ability	1 = PMA/PMD is able to perform 40GBASE-CR4 0 = PMA/PMD is not able to perform 40GBASE-CR4	RO
1.13.0	40GBASE-KR4 ability	1 = PMA/PMD is able to perform 40GBASE-KR4 0 = PMA/PMD is not able to perform 40GBASE-KR4	RO

[a]RO = Read only

45.2.1.12.1 PMA remote loopback ability (1.13.15)

When read as a one, bit 1.13.15 indicates that the PMA is able to perform the remote loopback function. When read as a zero, bit 1.13.15 indicates that the PMA is not able to perform the remote loopback function. If a PMA is able to perform the remote loopback function, then it is controlled using the PMA remote loopback bit 1.0.1 (see 45.2.1.1.4).

45.2.1.12.2 100GBASE-CR4 ability (1.13.14)

When read as a one, bit 1.13.14 indicates that the PMA/PMD is able to operate as a 100GBASE-CR4 PMA/PMD type. When read as a zero, bit 1.13.14 indicates that the PMA/PMD is not able to operate as a 100GBASE-CR4 PMA/PMD type.

45.2.1.12.3 100GBASE-KR4 ability (1.13.13)

When read as a one, bit 1.13.13 indicates that the PMA/PMD is able to operate as a 100GBASE-KR4 PMA/PMD type. When read as a zero, bit 1.13.13 indicates that the PMA/PMD is not able to operate as a 100GBASE-KR4 PMA/PMD type.

45.2.1.12.4 100GBASE-KP4 ability (1.13.12)

When read as a one, bit 1.13.12 indicates that the PMA/PMD is able to operate as a 100GBASE-KP4 PMA/PMD type. When read as a zero, bit 1.13.12 indicates that the PMA/PMD is not able to operate as a 100GBASE-KP4 PMA/PMD type.

45.2.1.12.5 100GBASE-ER4 ability (1.13.11)

When read as a one, bit 1.13.11 indicates that the PMA/PMD is able to operate as a 100GBASE-ER4 PMA/PMD type. When read as a zero, bit 1.13.11 indicates that the PMA/PMD is not able to operate as a 100GBASE-ER4 PMA/PMD type.

45.2.1.12.6 100GBASE-LR4 ability (1.13.10)

When read as a one, bit 1.13.10 indicates that the PMA/PMD is able to operate as a 100GBASE-LR4 PMA/PMD type. When read as a zero, bit 1.13.10 indicates that the PMA/PMD is not able to operate as a 100GBASE-LR4 PMA/PMD type.

45.2.1.12.7 100GBASE-SR10 ability (1.13.9)

When read as a one, bit 1.13.9 indicates that the PMA/PMD is able to operate as a 100GBASE-SR10 PMA/PMD type. When read as a zero, bit 1.13.9 indicates that the PMA/PMD is not able to operate as a 100GBASE-SR10 PMA/PMD type.

45.2.1.12.8 100GBASE-CR10 ability (1.13.8)

When read as a one, bit 1.13.8 indicates that the PMA/PMD is able to operate as a 100GBASE-CR10 PMA/PMD type. When read as a zero, bit 1.13.8 indicates that the PMA/PMD is not able to operate as a 100GBASE-CR10 PMA/PMD type.

45.2.1.12.9 100GBASE-SR4 ability (1.13.7)

When read as a one, bit 1.13.7 indicates that the PMA/PMD is able to operate as a 100GBASE-SR4 PMA/PMD type. When read as a zero, bit 1.13.7 indicates that the PMA/PMD is not able to operate as a 100GBASE-SR4 PMA/PMD type.

45.2.1.12.10 40GBASE-ER4 ability (1.13.5)

When read as a one, bit 1.13.5 indicates that the PMA/PMD is able to operate as a 40GBASE-ER4 PMA/PMD type. When read as a zero, bit 1.13.5 indicates that the PMA/PMD is not able to operate as a 40GBASE-ER4 PMA/PMD type.

45.2.1.12.11 40GBASE-FR ability (1.13.4)

When read as a one, bit 1.13.4 indicates that the PMA/PMD is able to operate as a 40GBASE-FR PMA/PMD type. When read as a zero, bit 1.13.4 indicates that the PMA/PMD is not able to operate as a 40GBASE-FR PMA/PMD type.

45.2.1.12.12 40GBASE-LR4 ability (1.13.3)

When read as a one, bit 1.13.3 indicates that the PMA/PMD is able to operate as a 40GBASE-LR4 PMA/PMD type. When read as a zero, bit 1.13.3 indicates that the PMA/PMD is not able to operate as a 40GBASE-LR4 PMA/PMD type.

45.2.1.12.13 40GBASE-SR4 ability (1.13.2)

When read as a one, bit 1.13.2 indicates that the PMA/PMD is able to operate as a 40GBASE-SR4 PMA/PMD type. When read as a zero, bit 1.13.2 indicates that the PMA/PMD is not able to operate as a 40GBASE-SR4 PMA/PMD type.

45.2.1.12.14 40GBASE-CR4 ability (1.13.1)

When read as a one, bit 1.13.1 indicates that the PMA/PMD is able to operate as a 40GBASE-CR4 PMA/PMD type. When read as a zero, bit 1.13.1 indicates that the PMA/PMD is not able to operate as a 40GBASE-CR4 PMA/PMD type.

45.2.1.12.15 40GBASE-KR4 ability (1.13.0)

When read as a one, bit 1.13.0 indicates that the PMA/PMD is able to operate as a 40GBASE-KR4 PMA/PMD type. When read as a zero, bit 1.13.0 indicates that the PMA/PMD is not able to operate as a 40GBASE-KR4 PMA/PMD type.

45.2.1.13 PMA/PMD package identifier (Registers 1.14 and 1.15)

Registers 1.14 and 1.15 provide a 32-bit value, which may constitute a unique identifier for a particular type of package that the PMA/PMD is instantiated within. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the package manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A PMA/PMD may return a value of zero in each of the 32 bits of the package identifier.

A non-zero package identifier may be returned by one or more MMDs in the same package. The package identifier may be the same as the device identifier.

The format of the package identifier is specified in 22.2.4.3.1.

45.2.1.14 EEE capability (Register 1.16)

This register is used to indicate the capability of the PMA/PMD to support EEE functions for each PHY type. The assignment of bits in the EEE capability register is shown in Table 45–17.

Table 45–17—EEE capability register bit definitions

Bit(s)	Name	Description	R/W[a]
1.16.15:12	Reserved	Value always 0	RO
1.16.11	100GBASE-CR4 deep sleep	1 = EEE deep sleep is supported for 100GBASE-CR4 0 = EEE deep sleep is not supported for 100GBASE-CR4	RO
1.16.10	100GBASE-KR4 deep sleep	1 = EEE deep sleep is supported for 100GBASE-KR4 0 = EEE deep sleep is not supported for 100GBASE-KR4	RO
1.16.9	100GBASE-KP4 deep sleep	1 = EEE deep sleep is supported for 100GBASE-KP4 0 = EEE deep sleep is not supported for 100GBASE-KP4	RO
1.16.8	100GBASE-CR10 deep sleep	1 = EEE deep sleep is supported for 100GBASE-CR10 0 = EEE deep sleep is not supported for 100GBASE-CR10	RO
1.16.7:2	Reserved	Value always 0	RO
1.16.1	40GBASE-CR4 deep sleep	1 = EEE deep sleep is supported for 40GBASE-CR4 0 = EEE deep sleep is not supported for 40GBASE-CR4	RO
1.16.0	40GBASE-KR4 deep sleep	1 = EEE deep sleep is supported for 40GBASE-KR4 0 = EEE deep sleep is not supported for 40GBASE-KR4	RO

[a] RO = Read only

45.2.1.14.1 100GBASE-CR4 EEE deep sleep supported (1.16.11)

If the device supports EEE deep sleep operation for 100GBASE-CR4 as defined in 92.1, this bit shall be set to a one; otherwise this bit shall be set to a zero.

45.2.1.14.2 100GBASE-KR4 EEE deep sleep supported (1.16.10)

If the device supports EEE deep sleep operation for 100GBASE-KR4 as defined in 93.1, this bit shall be set to a one; otherwise this bit shall be set to a zero.

45.2.1.14.3 100GBASE-KP4 EEE deep sleep supported (1.16.9)

If the device supports EEE deep sleep operation for 100GBASE-KP4 as defined in 94.1, this bit shall be set to a one; otherwise this bit shall be set to a zero.

45.2.1.14.4 100GBASE-CR10 EEE deep sleep supported (1.16.8)

If the device supports EEE deep sleep operation for 100GBASE-CR10 as defined in 85.1, this bit shall be set to a one; otherwise this bit shall be set to a zero.

45.2.1.14.5 40GBASE-CR4 EEE deep sleep supported (1.16.1)

If the device supports EEE deep sleep operation for 40GBASE-CR4 as defined in 85.1, this bit shall be set to a one; otherwise this bit shall be set to a zero.

45.2.1.14.6 40GBASE-KR4 EEE deep sleep supported (1.16.0)

If the device supports EEE deep sleep operation for 40GBASE-KR4 as defined in 84.1, this bit shall be set to a one; otherwise this bit shall be set to a zero.

45.2.1.15 10P/2B PMA/PMD control register (Register 1.30)

The assignment of bits in the 10P/2B PMA control register is shown in Table 45–18.

Table 45–18—10P/2B PMA control register bit definitions

Bit(s)	Name	Description	R/W[a]
1.30.15	PMA/PMD link control	1 = begin initialization, enable link (-R default) 0 = force link down (-O default)	R/W
1.30.14	STFU	Silence the far unit 1 = send silence command 0 = silence command inactive (default)	R/W
1.30.13:8	Silence time	Silence time = 10 × (value of bits + 1) seconds	R/W
1.30.7	Port subtype select	1 = port operates as an -O subtype 0 = port operates as a -R subtype	R/W
1.30.6	Handshake cleardown	1 = send cleardown command 0 = idle (default)	R/W, SC
1.30.5	Ignore incoming handshake	1 = PMA/PMD does not respond to handshake tones 0 = PMA/PMD responds to handshake tones	R/W
1.30.4:0	PMA/PMD type selection	4 3 2 1 0 0 0 0 0 0 = 10PASS-TS PMA/PMD type 0 0 0 0 1 = 2BASE-TL PMA/PMD type 0 0 0 1 0 = 2BASE-TL or 10PASS-TS (-R only) 0 0 0 1 1 = 2BASE-TL preferred, or 10PASS-TS (-O only) 0 0 1 0 0 = 10PASS-TS preferred, or 2BASE-TL (-O only) all other values are reserved	R/W

[a]R/W = Read/Write, SC = Self-clearing

45.2.1.15.1 PMA/PMD link control (1.30.15)

The STA may enable the PMA/PMD link and initiate link initialization by writing this bit to a one. While link is initializing or up, this bit shall remain a one and writing a one to this bit shall be ignored. The STA may force the link down by writing a zero to this bit. While this bit is set to zero, the PHY shall not send G.994.1 handshake tones. For -O subtypes, upon link drop or MMD reset, the PMA/PMD shall set these bits to zero. For -R subtypes, upon link drop or MMD reset, the PMA/PMD shall set these bits to one.

45.2.1.15.2 STFU (1.30.14)

When this bit is set to a one, the PMA/PMD sends a message to the link partner instructing it to be silent for the silence time (see 45.2.1.15.3). Writing to this bit is valid only when the PMA/PMD link status bits in the PMA/PMD status register (see 45.2.1.16.4) are set to “link is down (ready)”. Writes are otherwise ignored.

This bit clears to zero when the silence command is sent, or upon the execution of an MMD reset.

45.2.1.15.3 Silence time (1.30.13:8)

The value of these bits sets the silence time conveyed in a STFU operation (see 45.2.1.15.2). The silence time is encoded according to Equation (45–1), where x is the decimal value of the bits:

$$\text{time} = 10 \times (x + 1) \text{ seconds} \quad (45–1)$$

45.2.1.15.4 Port subtype select (1.30.7)

This register bit selects the port subtype for PMA/PMD operation. The bit defaults to a supported mode. The PHY shall ignore writes that select an unsupported mode (see 45.2.1.15).

Changing this bit alters the fundamental operation of the PMA/PMD, therefore, writes to change this bit shall be ignored if the link is up or initializing (see 45.2.1.16.4).

45.2.1.15.5 Handshake cleardown (1.30.6)

Setting this bit to a one shall cause the PMA/PMD to issue a G.994.1 cleardown command to the link partner (see 61.4.3). The PMA/PMD shall clear this bit to zero after the cleardown command has been sent or upon MMD reset. If the PMA/PMD link is not in the "link down (ready)" state (see 45.2.1.16.4), writes to this register shall be ignored.

45.2.1.15.6 Ignore incoming handshake (1.30.5)

When set to a one, the PMA/PMD shall not respond to received handshake tones (see 61.4.3). When set to a zero, the PMA/PMD shall respond to received handshake tones normally, according to 61.4.3 and G.994.1. Upon MMD reset, this bit shall be cleared to zero.

45.2.1.15.7 PMA/PMD type selection (1.30.4:0)

The PMA/PMD type of a 10P/2B PHY may be selected using bits 4 through 0. A PHY shall ignore writes to the type selection bits that select PMA/PMD types it has not advertised in the speed ability register. It is the responsibility of the STA entity to ensure that mutually acceptable port types are applied consistently across all the MMDs on a particular PHY.

A value of 00010 may be set in -R subtype PMA/PMDs that have both 2BASE-TL and 10PASS-TS capability set in the PMA/PMD speed ability register. The PMA/PMD type of the -R is set upon link initialization by the -O.

Values of 00011 and 00100 may be set in -O subtype PMA/PMDs that have both 2BASE-TL and 10PASS-TS capabilities set in the PMA/PMD speed ability register. These values indicate whether the -R is set to 10PASS-TS or 2BASE-TL respectively. If the -R is not capable of the "preferred" mode, the -R is set to 10PASS-TS or 2BASE-TL respectively.

The selection is advertised during link initialization G.994 handshake.

The PMA/PMD type selection defaults to a supported ability.

45.2.1.16 10P/2B PMA/PMD status register (Register 1.31)

The assignment of bits in the 10P/2B PMA/PMD status register is shown in Table 45–19.

45.2.1.16.1 Data rate (1.31.15:5)

These bits indicate the current bit rate of an operational PMA/PMD link. These bits shall be set to all zeros when the link is down or initializing.

45.2.1.16.2 CO supported (1.31.4)

This bit indicates that the PMA/PMD supports operation as a -O subtype. This bit is set to a one when the capability is supported and zero otherwise. This bit reflects the signal PMA_PMD_type in 61.3.2.1.

Table 45–19—10P/2B PMA/PMD status register bit definitions

Bit(s)	Name	Description	R/W[a]
1.31.15:5	Data rate	Current operating bit rate of the PMD *n* = the value of the bits Data rate = 64*n* kb/s	RO
1.31.4	CO supported	1 = -O subtype operation supported 0 = -O subtype operation not supported	RO
1.31.3	CPE supported	1 = -R subtype operation supported 0 = -R subtype operation not supported	RO
1.31.2:0	PMA/PMD link status	2 1 0 0 0 0 = link is down (not ready) 1 0 0 = link is down (ready) 0 0 1 = link is initializing 0 1 0 = link is up, 10PASS-TS 0 1 1 = link is up, 2BASE-TL all other values reserved	RO

[a]RO = Read only

45.2.1.16.3 CPE supported (1.31.3)

This bit indicates that the PMA/PMD supports operation as a -R subtype. This bit is set to a one when the capability is supported and zero otherwise. This bit reflects the signal PMA_PMD_type in 61.3.2.1.

45.2.1.16.4 PMA/PMD link status (1.31.2:0)

The overall state of the PMA/PMD link is reflected in bits 2:0. After the PMA/PMD is linked to the remote PHY, the PHY shall set these bits to indicate the PMA/PMD port type that is linked (010 for 10PASS-TS and 011 for 2BASE-TL). The corresponding signal, PMA_received_synchronized, is defined in 61.3.2.1.

While the link is initializing, these bits shall be set to 001.

When read as 000, these bits shall indicate that PMA/PMD link is down and the PMA/PMD is not detecting handshake tones from a link partner. This state is known as “not ready”.

When read as 100, these bits shall indicate that the PMA/PMD link is down and the PMA/PMD is detecting handshake tones from a link partner. This state is known as “ready”.

45.2.1.17 Link partner PMA/PMD control register (Register 1.32)

The 10PASS-TS-O and 2BASE-TL-O PMA/PMDs allow access to certain register values of their link partner via the local MDIO interface. A summary of link partner parameters that may be sent or retrieved is provided in Table 45–20.

The Link partner PMA/PMD control register allows the -O STA to control the transmission and retrieval of parameters from its -R link partner.

The -R STA may read values exchanged by the -O STA in the local register counterpart to the link partner register. For example, the -O 10P link partner electrical length register will be populated with the contents of the -R 10P electrical length register upon a successful “Get link partner parameters” command. Similarly, the -R 10P/2B line quality thresholds register will contain the values sent by the -O in the 10P/2B link partner line quality thresholds register, after a successful “Send link partner parameters” command.

Table 45–20—Link partner PMA/PMD registers and PMA/PMD register duals

	Register type	Link partner register	Local register counterpart	
	PHY subtype	-O	-O	-R
Link partner register name		**Address and access type[a]**		
10P/2B link partner RX SNR margin		1.38 RO	1.37 RO	1.37 RO
10P/2B link partner line attenuation		1.40 RO	1.39 RO	1.39 RO
10P/2B link partner line quality thresholds		1.42 R/W	1.41 R/W	n/a
10P link partner FEC correctable errors		1.45 RO	1.43 RO	1.43 RO
10P link partner FEC uncorrectable errors		1.46 RO	1.44 RO	1.44 RO
10P link partner electrical length		1.48 RO	1.47 RO	1.47 RO
2B link partner code violation errors		1.90 RO	1.89 RO	1.89 RO
2B link partner errored seconds		1.92 RO	1.91 RO	1.91 RO
2B link partner severely errored seconds		1.94 RO	1.93 RO	1.93 RO
2B link partner LOSW		1.96 RO	1.95 RO	1.95 RO
2B link partner unavailable seconds		1.98 RO	1.97 RO	1.97 RO
2B link partner state defects register		1.100 RO	1.99 RO	1.99 RO

[a]R/W = Read/Write, RO = read only, n/a = undefined

The link partner registers listed in Table 45–20 have the same behavior upon being read or reset as their local register counterparts.

This register is defined for -O port subtypes only.

Bit definitions for the Link partner PMA/PMD control register are found in Table 45–21.

Table 45–21—Link partner PMA/PMD control register bit definitions

Bit(s)	Name	Description	R/W[a]
1.32.15	Get link partner parameters	1 = get link partner parameters 0 = operation complete, ready	R/W, SC
1.32.14	Reserved	Value always 0	RO
1.32.13	Send link partner parameters	1 = send link partner parameters 0 = operation complete, ready	R/W, SC
1.32.12:0	Reserved	Value always 0	RO

[a]R/W = Read/Write, SC = Self-clearing, RO = read only

45.2.1.17.1 Get link partner parameters (1.32.15)

When this bit is set to a one, the -O PHY updates its link partner registers shown in Table 45–20 with values from the link partner. While the operation is in progress, the PHY shall keep the bit set as one. If the "Get link partner parameters" operation does not complete within 10 seconds, its result shall be marked as "failed" (see 45.2.1.18) and the operation marked as "complete". After completion of the operation or upon reset, the PHY shall reset the bit to zero. A write to this bit when link is down shall cause the result to be marked as "failed" and the operation marked as "complete".

45.2.1.17.2 Send link partner parameters (1.32.13)

When this bit is set to a one, the -O PHY sends the contents of the 2B link partner line quality thresholds register (see 45.2.1.24) to the link partner. While the operation is in progress, the PHY shall keep the bit set as one. The "Send link partner parameters" operation must complete within 10 seconds, or its result shall be marked as "failed" (see 45.2.1.17) and the operation marked as "complete". After completion of the operation or upon reset, the PHY shall reset the bit to zero. A write to this bit when link is down shall cause the result to be marked as "failed" and the operation marked as "complete".

45.2.1.18 Link partner PMA/PMD status register (Register 1.33)

The Link partner PMA/PMD status register reflects the result of the operations that are performed using the Link Partner PMA/PMD control register (1.32).

This register is defined for -O port subtypes only.

Bit definitions for the Link partner PMA/PMD status register are found in Table 45–22.

Table 45–22—Link Partner PMA/PMD status register bit definitions

Bit(s)	Name	Description	R/W[a]
1.33.15	Reserved	Value always 0	RO
1.33.14	Get link partner result	1 = operation failed 0 = operation successful	RO, LH
1.33.13	Reserved	Value always 0	RO
1.33.12	Send link partner result	1 = operation failed 0 = operation successful	RO, LH
1.33.11:0	Reserved	Value always 0	RO

[a]RO = Read only, LH = Latching high

45.2.1.18.1 Get link partner result (1.33.14)

After a "Get link partner parameters" operation terminates, this bit reflects the result of the operation. If the operation did not complete successfully, the PHY shall set this bit to a one. Upon being read or a reset, the PHY shall set the bit to zero.

The definition of an unsuccessful "Get link partner parameters" operation is unspecified and left to the implementation.

45.2.1.18.2 Send link partner result (1.33.12)

After a "Send link partner parameters" operation terminates, this bit reflects the result of the operation. If the operation did not complete successfully, the PHY shall set this bit to a one. Upon being read or a reset, the PHY resets the bit to zero.

The definition of an unsuccessful "Send link partner parameters" operation is unspecified and left to the implementation.

45.2.1.19 10P/2B PMA/PMD link loss register (Register 1.36)

The 10P/2B PMA/PMD link loss register is a 16 bit counter that contains the number of times the PMA/PMD link is lost. Link is considered lost when the PMA_receive_synchronized signal (see 61.3.2.1) transitions from up to down. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the MMD reset. These bits shall be held at all ones in the case of overflow. The assignment of bits in the 10P/2B PMA/PMD link loss register is shown in Table 45–23.

Table 45–23—10P/2B PMA/PMD link loss register bit definitions

Bit(s)	Name	Description	R/W[a]
1.36.15:0	Link lost events	The bytes of the counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.1.20 10P/2B RX SNR margin register (Register 1.37)

For further information on 2BASE-TL SNR margin, see 63.3. For 10PASS-TS SNR margin, see 62.3.

The bit definitions for the 10P/2B RX SNR margin register are found in Table 45–24.

Table 45–24—10P/2B RX SNR margin register bit definition

Bit(s)	Name	Description	R/W[a]
1.37.15:8	Reserved	Value always 0	RO
1.37.7:0	RX SNR margin	Value of SNR margin in dB	RO

[a]RO = Read only

45.2.1.21 10P/2B link partner RX SNR margin register (Register 1.38)

The 10P/2B link partner RX SNR margin register provides access to the link partner's receive SNR margin. The contents of this register are refreshed when the STA activates the "Get link partner parameter" command in the Link partner PMA/PMD control register (see 45.2.1.17) and the command completes successfully (see 45.2.1.18.1).

This register is defined for -O port subtypes only.

The bit definitions for this register are found in Table 45–24.

45.2.1.22 10P/2B line attenuation register (Register 1.39)

This register reports the line attenuation as measured by the PMA/PMD. For more information, see the reference documents in 63.3 and 62.3.

The bit definitions for the 10P/2B line attenuation register are found in Table 45–25.

Table 45–25—10P/2B line attenuation register bit definitions

Bit(s)	Name	Description	R/W[a]
1.39.15:0	Line attenuation	The value of the line attenuation in dB, as perceived by the local PMD.	RO

[a]RO = Read only

45.2.1.23 10P/2B link partner line attenuation register (Register 1.40)

The 10P/2B link partner line attenuation register provides access to the link partner's perceived line attenuation margin. The contents of this register are refreshed when the STA activates the "Get link partner parameter" command in the Link partner PMA/PMD control register (see 45.2.1.17) and the command completes successfully (see 45.2.1.18.1).

This register is defined for -O port subtypes only.

The bit definitions for this register are found in Table 45–25.

45.2.1.24 10P/2B line quality thresholds register (Register 1.41)

The 10P/2B line quality thresholds register sets the target environment for the 10PASS-TS/2BASE-TL connection. The line quality is defined by the SNR margin and line attenuation values.

Bit definitions for the 10P/2B line quality threshold register are found in Table 45–26.

Table 45–26—10P/2B line quality thresholds register bit definition

Bit(s)	Name	Description	R/W[a]
1.41.15:8	Loop attenuation threshold	Attenuation threshold in dB	O: R/W R: RO
1.41.7:4	SNR margin threshold	SNR margin threshold in dB	O: R/W R: RO
1.41.3:0	Reserved	Value always 0	RO

[a]R/W = Read/Write, RO = Read only

45.2.1.24.1 Loop attenuation threshold (1.41.15:8)

These bits set the loop attenuation threshold for 2BASE-TL PMA/PMDs. Writing to these bits on a 10PASS-TS PMA/PMD shall have no effect. The threshold value is in units of dB. For more information on the loop attenuation threshold, see 63.2.2.3.

45.2.1.24.2 SNR margin threshold (1.41.7:4)

These bits set the SNR margin threshold for 10PASS-TS and 2BASE-TL PMA/PMDs. The threshold is expressed in units of dB. For more information of the SNR margin threshold, see 63.2.2.3 for 2BASE-TL and Section 10 of the document referenced in 62.1.3 for 10PASS-TS.

45.2.1.25 2B link partner line quality thresholds register (Register 1.42)

The 2B link partner line quality thresholds register allows the -O STA to set its -R link partner's line quality thresholds. The contents of this register are transmitted to the -R when the STA activates the "Send link partner parameter" command in the Link Partner PMA/PMD control register (see 45.2.1.17) and the command completes successfully (see 45.2.1.18.2).

This register is defined for -O port subtypes only.

The bit definitions for this register are found in Table 45–26.

45.2.1.26 10P FEC correctable errors counter (Register 1.43)

The 10P FEC correctable errors counter is a 16 bit counter that contains the number of FEC codewords that have been received and corrected. For more information on 10PASS-TS FEC, see 62.2.4.2. These bits shall be reset to all zeros upon execution of the MMD reset and upon being read. The assignment of bits in the 10P FEC correctable error counter is shown in Table 45–27.

Table 45–27—10P FEC correctable errors counter register bit definitions

Bit(s)	Name	Description	R/W[a]
1.43.15:0	Correctable codewords [15:0]	The bytes of the counter	RO

[a]RO = Read only

45.2.1.27 10P FEC uncorrectable errors counter (Register 1.44)

The 10P FEC uncorrectable errors counter is a 16 bit counter that contains the number of FEC codewords that have been received and are uncorrectable. For more information on 10PASS-TS FEC, see 62.2.4.2. These bits shall be reset to all zeros upon execution of the MMD reset and upon being read. The assignment of bits in the 10P FEC uncorrectable error counter is shown in Table 45–28.

Table 45–28—10P FEC uncorrectable errors counter bit definitions

Bit(s)	Name	Description	R/W[a]
1.44.15:0	Uncorrectable codewords [15:0]	The bytes of the counter	RO

[a]RO = Read only

45.2.1.28 10P link partner FEC correctable errors register (Register 1.45)

The 10P link partner FEC correctable errors register provides the -O STA with a snapshot of the -R link partner's FEC correctable errors counter. Because this register is not a counter, its value will increment only

when refreshed. The contents of this register are refreshed when the STA activates the "Get link partner parameter" command in the Link Partner PMA/PMD control register (see 45.2.1.17) and the command completes successfully (see 45.2.1.18.1).

This register is defined for -O port subtypes only.

The bit definitions for this register are found in Table 45–27 and 45.2.1.26.

45.2.1.29 10P link partner FEC uncorrectable errors register (Register 1.46)

The 10P link partner FEC uncorrectable errors register provides the -O STA a snapshot of the -R link partner's FEC uncorrectable errors counter. Because this register is not a counter, its value will increment only when refreshed. The contents of this register are refreshed when the STA activates the "Get link partner parameter" command in the Link Partner PMA/PMD control register (see 45.2.1.17) and the command completes successfully (see 45.2.1.18.1).

This register is defined for -O port subtypes only.

The bit definitions for this register are found in Table 45–28 and 45.2.1.27.

45.2.1.30 10P electrical length register (Register 1.47)

The bit definitions for the 10P electrical length register are found in Table 45–29.

Table 45–29—10P electrical length register bit definitions

Bit(s)	Name	Description	R/W[a]
1.47.15:0	Electrical length	The electrical length of the medium (in meters), as perceived at the local PMD	RO

[a]RO = Read only

45.2.1.30.1 Electrical length (1.47.15:0)

After the link is established, these bits contain the measured electrical length (in meters) of the medium as measured by the PMD. If the link is down or the PMD is unable to determine the electrical length, these bits shall be set to all ones (see 62.3.2).

45.2.1.31 10P link partner electrical length register (Register 1.48)

The 10P link partner electrical length register provides access to the link partner's electrical length measurement. The contents of this register are refreshed when the STA activates the "Get link partner parameter" command in the Link Partner PMA/PMD control register (see 45.2.1.17) and the command completes successfully (see 45.2.1.18.1).

This register is defined for -O port subtypes only.

The bit definitions for this register are found in Table 45–29 and 45.2.1.30.1.

45.2.1.32 10P PMA/PMD general configuration register (Register 1.49)

The 10P PMA/PMD general configuration register is defined for -O port types only.

The 10P PMA/PMD general configuration register bit definitions are found in Table 45–30.

Table 45–30—10P PMA/PMD general configuration register bit definitions

Bit(s)	Name	Description	R/W[a]
1.49.15:8	Reserved	Value always 0	RO
1.49.7:0	TX window length	Transmit window length	O: R/W R: undefined

[a]RO = Read only, R/W = Read/Write

45.2.1.32.1 TX window length (1.49.7:0)

Bits 7:0 control the PMD transmit window length within the cyclic prefix and suffix in units of number of samples, as defined in 62.3.2.

45.2.1.33 10P PSD configuration register (Register 1.50)

This register is defined for -O port subtypes only.

The 10P PSD configuration register bit definitions may be found in Table 45–31.

Table 45–31—10P PSD configuration register bit definitions

Bit(s)	Name	Description	R/W[a]
1.50.15:9	Reserved	Value always 0	RO
1.50.8	PBO disable	1 = PBO disabled 0 = PBO normal operation	O: R/W R: undefined
1.50.7:0	Reserved	Value always 0	RO

[a]RO = Read only, R/W = Read/Write

45.2.1.33.1 PBO disable (1.50.8)

Setting this bit to a one disables UPBO for performance testing purposes. Refer to 62.3.4.4.

45.2.1.34 10P downstream data rate configuration (Registers 1.51, 1.52)

These registers are defined for -O port subtypes only.

The bit definitions for the 10P downstream data rate configuration registers are found in Table 45–32.

Table 45–32—10P downstream data rate configuration register bit definitions

Bit(s)	Name	Description	R/W[a]
1.51.15:0	Minimum downstream data rate	Sets the minimum required downstream payload data rate *M* = value of bits Data rate = *M* × 64000 b/s	O: R/W R: undefined
1.52.15:0	Maximum downstream data rate	Sets the maximum downstream payload data rate *M* = value of bits Data rate = *M* × 64000 b/s	O: R/W R: undefined

[a]R/W = Read/Write

45.2.1.35 10P downstream Reed-Solomon configuration (Register 1.53)

This register is defined for -O port subtypes only.

The bit definitions for 10P downstream Reed-Solomon configuration are found in Table 45–33.

Table 45–33—10P downstream Reed-Solomon configuration register bit definitions

Bit(s)	Name	Description	R/W[a]
1.53.15:1	Reserved	Value always 0	RO
1.53.0	RS codeword length	1 = codeword length of 144 0 = codeword length of 240	O: R/W R: undefined

[a]RO = Read only, R/W = Read/Write

45.2.1.35.1 RS codeword length (1.53.0)

This bit selects the Reed-Solomon forward error correction codeword length used in the downstream direction. For more information, see 62.2.4.2.

45.2.1.36 10P upstream data rate configuration (Registers 1.54, 1.55)

These registers are defined for -O port subtypes only.

The bit definitions for 10P upstream data rate configuration are found in Table 45–34.

45.2.1.37 10P upstream 10P upstream Reed-Solomon configuration register (Register 1.56)

This register is defined for -O port subtypes only.

The bit definitions for the 10P upstream Reed-Solomon configuration are found in Table 45–35.

Table 45–34—10P upstream data rate configuration register bit definitions

Bit(s)	Name	Description	R/W[a]
1.54.15:0	Minimum upstream data rate	Sets the required upstream payload data rate *M* = value of bits Data rate = *M* × 64000 b/s	O: R/W R: undefined
1.55.15:0	Maximum upstream data rate	Sets the maximum upstream payload data rate *M* = value of bits Data rate = *M* × 64000 b/s	O: R/W R: undefined

[a]R/W = Read/Write

Table 45–35—10P upstream Reed-Solomon configuration register bit definitions

Bit(s)	Name	Description	R/W[a]
1.56.15:1	Reserved	Value always 0	RO
1.56.0	RS codeword length	1 = codeword length = 144 0 = codeword length = 240	O: R/W R: undefined

[a]RO = Read only, R/W = Read/Write

45.2.1.37.1 RS codeword length (1.56.0)

This bit selects the Reed-Solomon forward error correction codeword length used in the upstream direction. For more information, see 62.2.4.2.

45.2.1.38 10P tone group registers (Registers 1.57, 1.58)

10PASS-TS operates by modulating 4096 individual tones across the transmission spectrum. Each tone can be assigned a PSD level, desired SNR margin and transmission direction (downstream or upstream). To reduce the complexity of addressing individual tones, tones are addressed by group. The STA sets the lower and upper tones in a group, sets the parameters for that group, and issues a command to activate those parameters for that group. See 62.3.4.7 for details on the mechanism that transfers tone information across the link to and from the -R link partner.

This register allows the STA to specify the range of tones to control. The bit definitions for the 10P tone group register are defined in Table 45–36.

Table 45–36—10P tone group register bit definitions

Bit(s)	Name	Description	R/W[a]
1.57.15:0	Lower tone	The number of the lower frequency tone in the group. Valid when ≤ the Upper tone.	R/W
1.58.15:0	Upper tone	The number of the higher frequency tone in the group. Valid when ≥ the Lower tone.	R/W

[a]R/W = Read/Write

45.2.1.39 10P tone control parameters (Registers 1.59, 1.60, 1.61, 1.62, 1.63)

These registers allow the STA to specify parameters for the tones selected in the 10P tone group registers. These values do not take effect until the corresponding activation commands are issued in the 10P tone control action register. The bit definitions for the 10P tone control parameters are shown in Table 45–37.

These registers are defined for -O port subtypes only.

Table 45–37—10P tone control parameters register bit definitions

Bit(s)	Name	Description	R/W[a]
1.59.15	Tone active	1 = selected tones are active 0 = selected tones are disabled	R/W
1.59.14	Tone direction	1 = selected tones assigned to upstream communication 0 = selected tones assigned to downstream communication	R/W
1.59.13:5	Max SNR margin	Assigns the maximum SNR margin the PMD may achieve *M* = value of bits Max SNR margin = *M*/4 dB	R/W
1.59.4:0	Reserved	Value always 0	RO
1.60:15:9	Reserved	Value always 0	RO
1.60.8:0	Target SNR margin	Assigns the target SNR margin for the selected tones *M* = value of bits Target SNR margin = *M*/4 dB	R/W
1.61:15:9	Reserved	Value always 0	RO
1.61.8:0	Min SNR margin	Assigns the minimum SNR margin for the selected tones *M* = value of bits Min SNR margin = *M*/4 dB	R/W
1.62.15:9	Reserved	Value always 0	RO
1.62.8:0	PSD level	Assigns a TX PSD level for the selected tones in dBm/Hz *P* = value of bits (2's complement) PSD level = *P*/4 – 100 dBm/Hz	R/W
1.63.15:9	Reserved	Value always 0	RO
1.63.8:0	USPBO reference	Assigns the level of the USPBO reference at the points represented by the selected tones *P* = value of bits (2's complement) PSD level = *P*/4 – 100 dBm/Hz	R/W

[a]R/W = Read/Write, RO = Read only

45.2.1.39.1 Tone active (1.59.15)

These bits are used to control the activity of the selected tones. When the "Change tone activity" command is issued (1.64.4), the selected tones will be either activated or deactivated based on the value set in these bits.

45.2.1.39.2 Tone direction (1.59.14)

These bits are used to control the direction of the selected tones. When the "Change tone direction" command is issued (1.64.3), the selected tones will adopt the direction set in these bits.

45.2.1.39.3 Max SNR margin (1.59.13:5)

These bits control the maximum SNR margin for the selected tones. When the "Change SNR margin" command is issued (1.64.2), the PMA/PMD will use the value set in these bits in calculations related to maximum SNR margin. The SNR margin is in units of dB, derived by dividing the value of bits 13:5 by 4.

45.2.1.39.4 Target SNR margin (1.60.8:0)

These bits control the target SNR margin for the selected tones. When the "Change SNR margin" command is issued (1.64.2), the PMA/PMD will use the value set in these bits in calculations related to target SNR margin. The target SNR margin is in units of dB, derived by dividing the value of bits 8:0 by 4.

45.2.1.39.5 Minimum SNR margin (1.61.8:0)

These bits control the minimum SNR margin for the selected tones. When the "Change SNR margin" command is issued (1.64.2), the PMA/PMD will use the value set in these bits in calculations related to minimum SNR margin. The minimum SNR margin is in units of dB, derived by dividing the value of bits 8:0 by 4.

45.2.1.39.6 PSD level (1.62.8:0)

These bits control the transmit PSD level of the selected tones. When the "Change PSD level" command is issued (1.64.1), the PMA/PMD will set the PSD level of the selected tones to according to Equation (45–2), where *x* is the value of bits 8:0:

$$power = \frac{x}{4} - 100\frac{\text{dBm}}{\text{Hz}} \tag{45–2}$$

45.2.1.39.7 USPBO reference (1.63.8:0)

These bits control the reference level for the upstream power back-off function of the PMA/PMD. When the "Change USPBO reference PSD" command (1.64.0) is issued, the portion of the USPBO reference curve that corresponds to the selected tones is changed to the value specified by these bits. The USPBO reference level is specified according to Equation (45–3), where *x* is the value of bits 8:0:

$$power = \frac{x}{4} - 100\frac{\text{dBm}}{\text{Hz}} \tag{45–3}$$

45.2.1.40 10P tone control action register (Register 1.64)

The operations in this register apply to the tones selected in the 10P tone group registers (1.57, 1.58).

This register is defined for -O port subtypes only.

The bit definitions for the 10P tone control action register are shown in Table 45–38.

Table 45–38—10P tone control action register bit definitions

Bit(s)	Name	Description	R/W[a]
1.64.15:6	Reserved	Value always 0	RO
1.64.5	Refresh tone status	1 = refresh selected tones for the 10P tone status registers 0 = ready, operation complete	R/W, SC
1.64.4	Change tone activity	1 = activate tone active setting as in tone control parameter register 0 = ready, operation complete	R/W, SC
1.64.3	Change tone direction	1 = activate tone direction setting as in tone control parameter register 0 = ready, operation complete	R/W, SC
1.64.2	Change SNR margin	1 = activate min, max and target SNR margin settings as in tone control parameter register 0 = ready, operation complete	R/W, SC
1.64.1	Change PSD level	1 = activate PSD level setting as in tone control parameter register 0 = ready, operation complete	R/W, SC
1.64.0	Change UPBO reference PSD	1 = activate UPBO reference PSD settings as in tone control parameter register 0 = ready, operation complete	R/W, SC

[a]RO = Read only, R/W = Read/Write, SC = Self-clearing

45.2.1.40.1 Refresh tone status (1.64.5)

When this bit is set to a one, the tone status information from the local and link partner is gathered so that it may be read using the 10P tone status registers (1.65, 1.66, and 1.67). While the tones are being refreshed, this bit shall remain set as one. This bit shall be reset to zero when the refresh operation is over or upon reset.

NOTE—Refreshing a large number of tones may take a long time to complete.

45.2.1.40.2 Change tone activity (1.64.4)

When this bit is set to a one, the selected tones are enabled or disabled according to the assignment in the tone active bit of the 10P tone control parameters register (1.59.15). While the tones are being activated/deactivated, this bit shall remain set as one. This bit shall be reset to zero when the operation is over or upon reset.

45.2.1.40.3 Change tone direction (1.64.3)

When this bit is set to a one, the transmission direction of selected tones is changed according to the assignment in the tone direction bit of the 10P tone control parameters register (1.59.14). While the tones are being assigned, this bit shall remain set as one. This bit shall be reset to zero when the operation is over or upon reset.

45.2.1.40.4 Change SNR margin (1.64.2)

When this bit is set to a one, the SNR margin parameters for the selected tones are loaded according the assignment in the Minimum, Target and Maximum SNR margin bits of the 10P tone control parameters

register (1.59.13:5, 1.60.8:0, 1.61.8:0). While the parameters are being loaded, this bit shall remain set as one. This bit shall be reset to zero when the operation is over or upon reset.

45.2.1.40.5 Change PSD level (1.64.1)

When this bit is set to a one, the PSD level for the selected tones is set according to the value in the PSD level bits of the 10P tone control parameters register (1.62.8:0). While the PSD is being set, this bit shall remain set as one. This bit shall be reset to zero when the operation is over or upon reset.

45.2.1.40.6 Change USPBO reference PSD (1.64.0)

When this bit is set to a one, the upstream power back-off reference PSD level for the selected tones is set according to the value in the USPBO PSD reference bits of the 10P tone control parameters register (1.63.8:0). While the reference PSD is being set, this bit shall remain set as one. This bit shall be reset to zero when the operation is over or upon reset.

45.2.1.41 10P tone status registers (Registers 1.65, 1.66, 1.67)

The 10P tone status registers allow the STA to query the status of any individual tone in the link. The values read in the 10P tone status register correspond to the tone whose number is set in the "Lower tone" of the 10P tone group registers (see 45.2.1.38).

The status of some tones is read from the link partner. Because the constant update of these values would be a strain on channel resources, these values are only updated for selected tones when the "Refresh tone table" command is issued in the 10P tone control action register (1.64).

The 10P tone status registers bit definitions are given in Table 45–39.

Table 45–39—10P tone status registers bit definitions

Bit(s)	Name	Description	R/W[a]
1.65.15	Refresh status	1 = tone entry has been refreshed 0 = tone entry has not been refreshed since last read	RO
1.65.14	Active	1 = tone is disabled 0 = tone is active	RO
1.65.13	Direction	1 = tone is assigned to upstream communication 0 = tone is assigned to downstream communication	RO
1.65.12:8	Reserved	Value always 0	RO
1.65.7:0	RX PSD	PSD of the tone at the receiver in dBm/Hz	RO
1.66.15:8	TX PSD	PSD of the tone at the transmitter in dBm/Hz	RO
1.66.7:3	Bit load	The number of bits currently loaded on the tone	RO
1.66.2:0	Reserved	Value always 0	RO
1.67.15:10	Reserved	Value always 0	RO
1.67.9:0	SNR margin	Current SNR margin for the tone R = value of bits SNR Margin = R/4 dB	RO

[a]RO = Read only

45.2.1.41.1 Refresh status (1.65.15)

When read as a one, bit 1.65.15 indicates that the values for this tone table have not been read since the last "Refresh tone status" command was issued by the 10P tone control action register (1.64). Upon reading this bit or upon reset, the bit shall be reset to zero.

45.2.1.41.2 Active (1.65.14)

When read as a one, this bit indicates that the selected tone is disabled (i.e., powered off and not carrying data).

45.2.1.41.3 Direction (1.65.13)

When read as a one, this bit indicates that the selected tone is assigned to upstream communication. When read as a zero, the tone is assigned to downstream communication.

45.2.1.41.4 RX PSD (1.65.7:0)

These bits report the PSD of the selected tone as perceived at the receiver in units of dBm/Hz.

45.2.1.41.5 TX PSD (1.66.15:8)

These bits report the PSD of the selected tone as output by the transmitter in units of dBm/Hz.

45.2.1.41.6 Bit load (1.66.7:3)

These bits report the number of bits encoded on the selected tone.

45.2.1.41.7 SNR margin (1.67.9:0)

These bits report the current SNR margin for the selected tone, as perceived by the receiver, in units of dB. The value of the SNR margin is obtained by dividing the decimal value of bits 9:0 by 4.

45.2.1.42 10P outgoing indicator bits status register (Register 1.68)

The 10P outgoing indicator bits status register conveys the current state of the indicator bits being sent to the link partner (See 62.3.4.7). The bit definitions for the 10P indicator bits status register are shown in Table 45–40.

Table 45–40—10P outgoing indicator bits status register bit definition

Bit(s)	Name	Description	R/W[a]
1.68.15:9	Reserved	Value always 0	RO
1.68.8	LoM	1 = received signal below SNR margin threshold 0 = normal state	RO
1.68.7	lpr	1 = power supply voltage invalid 0 = normal state	RO
1.68.6	po	1 = PMA/PMD is being powered off 0 = normal state	RO
1.68.5	Rdi	1 = severely errored frames have been received 0 = normal state	RO
1.68.4	los	1 = signal power is lower than the threshold 0 = normal state	RO

Table 45–40—10P outgoing indicator bits status register bit definition *(continued)*

Bit(s)	Name	Description	R/W[a]
1.68.3:2	Reserved	Value always 0	RO
1.68.1	fec-s	1 = corrected errors have been detected in the received FEC block of slow data 0 = normal state	RO
1.68.0	be-s	1 = non-corrected errors have been detected in the received block of slow data 0 = normal state	RO

[a]RO = Read only

NOTE—These bit refer to “slow” data. 10PASS-TS uses the slow data channel as referenced in T1.424. The name is kept here to simplify a comparison between the two documents.

45.2.1.42.1 LoM (1.68.8)

When read as a one, this bit indicates that the PMA/PMD is receiving a signal whose SNR margin is below the set threshold (see 45.2.1.24). The status of this condition is reported to the link partner PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.42.2 lpr (1.68.7)

When read as a one, this bit indicates that the PMA/PMD is not receiving sufficient power supply input for proper operation. The specific conditions that cause this bit to be set are implementation specific. The status of this condition is reported to the link partner PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.42.3 po (1.68.6)

When read as a one, this bit indicates that the PMA/PMD has been instructed to power off. The specific conditions that cause this bit to be set are implementation specific. The status of this condition is reported to the link partner PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.42.4 Rdi (1.68.5)

When read as a one, this bit indicates that the PMA/PMD has received PMA/PMD frames containing severe errors. The status of this condition is reported to the link partner PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.42.5 los (1.68.4)

When read as a one, this bit indicates that the PMA/PMD is not receiving a valid signal. The status of this condition is reported to the link partner PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.42.6 fec-s (1.68.1)

When read as a one, this bit indicates that the PMA/PMD is receiving FEC blocks with one or more correctable errors. The status of this condition is reported to the link partner PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.42.7 be-s (1.68.0)

When read as a one, this bit indicates that the PMA/PMD is receiving FEC blocks with one or more uncorrectable errors. The status of this condition is reported to the link partner PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.43 10P incoming indicator bits status register (Register 1.69)

The 10P indicator bits status register conveys the current state of the indicator bits being received from the link partner's PMA (see 62.3.4.7). The bit definitions for the 10P incoming indicator bits status register are shown in Table 45–41.

Table 45–41—10P incoming indicator bits status register bit definition

Bit(s)	Name	Description	R/W[a]
1.69.15:9	Reserved	Value always 0	RO
1.69.8	LoM	1 = received signal below SNR margin threshold 0 = normal state	RO
1.69.7	Flpr	1 = power supply voltage invalid 0 = normal state	RO
1.69.6	Fpo	1 = PMA/PMD is being powered off 0 = normal state	RO
1.69.5	Rdi	1 = severely errored frames have been received 0 = normal state	RO
1.69.4	Flos	1 = signal power is lower than the threshold 0 = normal state	RO
1.69.3:2	Reserved	Value always 0	RO
1.69.1	Ffec-s	1 = corrected errors have been detected in the received FEC block of slow data 0 = normal state	RO
1.69.0	Febe-s	1 = non-corrected errors have been detected in the received block of slow data 0 = normal state	RO

[a]RO = Read only

NOTE—These bit refer to "slow" data. 10PASS-TS uses the slow data channel as referenced in T1.424. The name is kept here to simplify a comparison between the two documents.

45.2.1.43.1 LoM (1.69.8)

When read as a one, this bit indicates that the link partner PMA/PMD is receiving a signal whose SNR margin is below the set threshold (see 45.2.1.24). The status of this condition is reported to the local PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.43.2 Flpr (1.69.7)

When read as a one, this bit indicates that the link partner PMA/PMD is not receiving sufficient power supply input for proper operation. The specific conditions that cause this bit to be set are implementation specific. The status of this condition is reported to the local PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.43.3 Fpo (1.69.6)

When read as a one, this bit indicates that the link partner PMA/PMD has been instructed to power off. The specific conditions that cause this bit to be set are implementation specific. The status of this condition is reported to the local PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.43.4 Rdi (1.69.5)

When read as a one, this bit indicates that the link partner PMA/PMD has received PMA/PMD frames containing severe errors. The status of this condition is reported to the local PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.43.5 Flos (1.69.4)

When read as a one, this bit indicates that the link partner PMA/PMD has is not receiving a valid signal. The status of this condition is reported to the local PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.43.6 Ffec-s (1.69.1)

When read as a one, this bit indicates that the link partner PMA/PMD is receiving FEC blocks with one or more correctable errors. The status of this condition is reported to the local PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.43.7 Febe-s (1.69.0)

When read as a one, this bit indicates that the link partner PMA/PMD is receiving FEC blocks with one or more uncorrectable errors. The status of this condition is reported to the local PMA/PMD via the 10PASS-TS indicator bits.

45.2.1.44 10P cyclic extension configuration register (Register 1.70)

The 10P cyclic extension configuration register controls the length of the cyclic extension for the 10P PMD. For more information, see 62.3.4.2. The value of the cyclic extension is equal to the decimal value set in bits 15:0. Values of decimal 10, 20 and 40 are valid. Writes to set any other values shall be ignored. Upon reset, the PMD shall set these bits to a decimal value of 20.

The bit definitions for this register are shown in Table 45–42.

Table 45–42—10P cyclic extension configuration register bit definitions

Bit(s)	Name	Description	R/W[a]
1.70.15:0	Cyclic extension	The value of the cyclic extension	O: R/W R: RO

[a]R/W = Read/Write, RO = Read only

45.2.1.45 10P attainable downstream data rate register (Register 1.71)

The 10P attainable downstream data rate register reports the data rate that the -R link partner measures to be the highest data rate for downstream transmission. The data rate is encoded as 1 kb/s times the decimal value of the register bits 15:0. The value of the register bits are not valid until after link is “up” (see 45.2.1.16.4).

The bit definitions for this register are found in Table 45–43.

Table 45–43—10P attainable downstream data rate register bit definitions

Bit(s)	Name	Description	R/W[a]
1.71.15:0	Attainable downstream data rate	Data rate in 1 kb/s increments	RO

[a]RO = Read only

45.2.1.46 2B general parameter register (Register 1.80)

The 2B general parameter register controls various parameters for the operation of the 2BASE-TL PMA/PMD.

This register is read only for -R ports which may be read so the STA may know the mode selected by the -O port. The selected parameters on the -O are sent to the -R link partner on link initialization. For more information on how these parameters are passed across the physical link using G.994.1 signaling, see 61B.3.2.

The bit definitions for the 2B general parameter register are found in Table 45–44.

Table 45–44— 2B general parameter register bit definition

Bit(s)	Name	Description	R/W[a]
1.80.15	Reserved	Value always 0	RO
1.80.14:10	PMMS target margin	margin = 14:10 - 10dB	R/W
1.80.9	Line probing control	1 = use line probing 0 = do not use line probing (default)	R/W
1.80.8	Noise environment	1 = current condition 0 = worst case (default)	R/W
1.80.7:2	Reserved	Value always 0	RO
1.80.1:0	Region	Selects the regional annex to operate under 00 = Annex A 01 = Annex B 10 = Annex C 11 = reserved, writes ignored	O: R/W R: RO

[a]R/W = Read/Write, RO = Read only

45.2.1.46.1 PMMS target margin (1.80.14:10)

The PMMS target margin specified in bits 14:10 specifies the noise margin that the PMMS procedure tries to attain. The margin is expressed in dB as the decimal value of bits 14:10 minus 10 dB. The margin specified is measured either against either worst case or current line conditions, based on the value set in bit 1.80.8.

The PMMS margin value is transferred during 2BASE-TL initialization via the worst case PMMS margin bits in Table 61B–57 and Table 61B–43, or the current condition PMMS margin bits in Table 61B–48 and Table 61B–44.

45.2.1.46.2 Line probing control (1.80.9)

When set to a one, this bit tells the PMA/PMD to perform line probing the next time link is initialized. When set to a zero, the PMA/PMD does not use line probing. Line probing causes the PMA/PMD to select probe duration and the link data rate. For more information, see the documents referenced in 63.3.

45.2.1.46.3 Noise environment (1.80.8)

This bit controls the reference noise used during line probing. When set to a one, the noise environment is based on the current line conditions. When set to a zero, the noise environment is based on worst case models. For more information, see the documents referenced in 63.3.

45.2.1.46.4 Region (1.80.1:0)

These bits select the regional annex that is used for the operation of the 2BASE-TL PMA/PMD. These annexes refer to clauses in documents referenced by the 2BASE-TL PMA/PMD specification. These are not annexes in IEEE Std 802.3. For details on each annex, see the document referenced in 63.1.3.

45.2.1.47 2B PMD parameters registers (Registers 1.81 through 1.88)

The 2B PMD parameters registers set the transmission parameters for an individual 2BASE-TL PMA/PMD link. When the link is initialized, these parameters are used by the link partner PMA/PMDs in an attempt to achieve specified settings.

These registers allow one to specify a single fixed data rate or up to four data rate ranges at the -O PMA/PMD. An additional set of four data ranges are found in the 2B extended PMD parameters registers (1.102 through 1.109). Bit descriptions for the 2B extended PMD parameters registers are found in 45.2.1.61. Together these sets allow up to eight data rate ranges to be specified.

If at least one data rate range is specified with different Minimum and Maximum data rates, the link is trained with the highest attainable rate. If line probing is enabled, the highest rate is determined by the result of line probing and the “Data rate step” value is ignored. If line probing is disabled, the minimum and maximum rate, “Data rate step” and “Power” values are used to determine the highest attainable rate.

In the case of a single fixed rate specified (Minimum data rate1 equals Maximum data rate1, Data rate step[1:8] set to zero, Minimum/Maximum data rate[2:8] set to zero), the link is trained at the specified rate.

When multiple ranges are specified, the PMD selects the first attainable range, starting sequentially from the first range.

Since writing to this register does not have an immediate effect, reading this register returns the desired parameters, which are not necessarily the current operating parameters.

For more information on how these parameters are passed across the physical link using G.994.1 signaling (see 61B.3.2).

The bit definitions for the 2B PMD parameters register are found in Table 45–45.

Table 45–45— 2B PMD parameters registers bit definition

Bit(s)	Name	Description	R/W[a]
1.81.15	Reserved	Value always 0	RO
1.81.14:8	Min data rate1	Min data rate of the first range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.81.7	Reserved	Value always 0	RO
1.81.6:0	Max data rate1	Max data rate of the first range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.82.15:14	Reserved	Value always 0	RO
1.82.13:7	Data rate step1	Data rate step of the first range *n* = value of the bits, *n* valid 1 to 86 Data Rate = 64*n* kb/s	O: R/W R: RO
1.82.6:2	Power1	*x* = multiple of 0.5 dBm to add to 5 dBm offset Power = (5 + 0.5*x*) dBm	O: R/W R: RO
1.82.1:0	Constellation1	11 = reserved 10 = 16-TCPAM 01 = 32-TCPAM 00 = automatically set by the PHY	O: R/W R: RO
1.83.15	Reserved	Value always 0	RO
1.83.14:8	Min data rate2	Min data rate of the second range *n* = value of the bits, *n* valid 3 to 89 Data Rate =64*n* kb/s	O: R/W R: RO
1.83.7	Reserved	Value always 0	RO
1.83.6:0	Max data rate2	Max data rate of the second range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.84.15:14	Reserved	Value always 0	RO
1.84.13:7	Data rate step2	Data rate step of the second range *n* = value of the bits, *n* valid 1 to 86 Data Rate = 64*n* kb/s	O: R/W R: RO
1.84.6:2	Power2	*x* = multiple of 0.5 dBm to add to 5 dBm offset Power = (5 + 0.5*x*) dBm	O: R/W R: RO
1.84.1:0	Constellation2	11 = reserved 10 = 16-TCPAM 01 = 32-TCPAM 00 = automatically set by the PHY	O: R/W R: RO
1.85.15	Reserved	Value always 0	RO
1.85.14:8	Min data rate3	Min data rate of the third range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.85.7	Reserved	Value always 0	RO

Table 45–45— 2B PMD parameters registers bit definition *(continued)*

Bit(s)	Name	Description	R/W[a]
1.85.6:0	Max data rate3	Max data rate of the third range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.86.15:14	Reserved	Value always 0	RO
1.86.13:7	Data rate step3	Data rate step of the third range *n* = value of the bits, *n* valid 1 to 86 Data Rate = 64*n* kb/s	O: R/W R: RO
1.86.6:2	Power3	*x* = multiple of 0.5 dBm to add to 5 dBm offset Power = (5 + 0.5*x*) dBm	O: R/W R: RO
1.86.1:0	Constellation3	11 = reserved 10 = 16-TCPAM 01 = 32-TCPAM 00 = automatically set by the PHY	O: R/W R: RO
1.87.15	Reserved	Value always 0	RO
1.87.14:8	Min data rate4	Min data rate of the fourth range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.87.7	Reserved	Value always 0	RO
1.87.6:0	Max data rate4	Max data rate of the fourth range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.88.15:14	Reserved	Value always 0	RO
1.88.13:7	Data rate step4	Data rate step of the fourth range *n* = value of the bits, *n* valid 1 to 86 Data Rate = 64*n* kb/s	O: R/W R: RO
1.88.6:2	Power4	*x* = multiple of 0.5 dBm to add to 5 dBm offset Power = (5 + 0.5*x*) dBm	O: R/W R: RO
1.88.1:0	Constellation4	11 = reserved 10 = 16-TCPAM 01 = 32-TCPAM 00 = automatically set by the PHY	O: R/W R: RO

[a]R/W = Read/Write, RO = Read only

45.2.1.47.1 Minimum data rate (1.81, 1.83, 1.85, 1.87. Bits 14:8)

Bits 14:8 in registers 1.81 through 1.87 set the minimum data rate for each of the four ranges. Valid values for these bits are decimal 3 through 89, writes to set an invalid value shall be ignored. The rate is expressed in units of kb/s and is derived by multiplying the decimal value of bits 14:8 by 64.

45.2.1.47.2 Max data rate (1.81, 1.83, 1.85, 1.87. Bits 6:0)

Bits 6:0 in registers 1.81 through 1.87 set the maximum data rate for each of the four ranges. Valid values for these bits are decimal 3 through 89, writes to set an invalid value shall be ignored. The rate is expressed in units of kb/s and is derived by multiplying the decimal value of bits 6:0 by 64.

45.2.1.47.3 Data rate step (1.82, 1.84, 1.86, 1.88. Bits 13:7)

Bits 13:7 in registers 1.82 through 1.88 set the granularity used by the PMA/PMD when determining the line rate. Valid values for these bits are decimal 1 through 86, writes to set an invalid value shall be ignored. The data rate step is expressed in units of kb/s and is derived by multiplying the decimal value of bits 13:7 by 64.

45.2.1.47.4 Power (1.82, 1.84, 1.86, 1.88. Bits 6:2)

Bits 6:2 in registers 1.82 through 1.88 set the allowed power level for each data rate range. The power levels set in these bits override those of the annex selected in the 2B general parameter register (1.80). The power level is expressed in units of dBm and is derived by Equation (45–4), where *x* equals the value of bits 6:2.

$$power = \left(5 + \frac{x}{2}\right) \text{ dBm} \tag{45–4}$$

45.2.1.47.5 Constellation (1.82, 1.84, 1.86, 1.88. Bits 1:0)

Bits 1:0 in registers 1.82 through 1.88 set the allowed constellation for each data rate range. Setting a value of 10 or 01 restricts the constellation to 16- or 32-TCPAM respectively. When set to a value of 00, the PMD automatically determines the constellation during initialization. Attempts to set a value of 11 shall be ignored.

45.2.1.48 2B code violation errors counter (Register 1.89)

The 2B code violation errors counter is a 16-bit counter that contains the number of the 2BASE-TL CRC anomalies. See 63.2.2.3 for more information. These bits shall be set to all zeros upon an MMD reset and upon being read.

Bit definitions for the 2B code violation errors counter are found in Table 45–46.

Table 45–46—2B code violation errors counter bit definitions

Bit(s)	Name	Description	R/W[a]
1.89.15:0	Code violations [15:0]	The bytes of the counter	RO

[a]RO = Read only

45.2.1.49 2B link partner code violations register (Register 1.90)

The 2B link partner code violations register provides the -O STA with a snapshot of the -R link partner's 2B code violations counter. Because this register is not a counter, its value will be updated only when refreshed. The contents of this register are refreshed when the STA activates the "Get link partner parameter" command in the Link Partner PMA/PMD control register (see 45.2.1.17) and the command completes successfully (see 45.2.1.18.1).

This register is defined for -O port subtypes only.

The bit definitions for this register are found in Table 45–46.

45.2.1.50 2B errored seconds counter (Register 1.91)

This 8-bit counter contains the number of errored seconds (see 63.2.2.3) These bits shall be set to all zeros when the register is read by management or upon an MMD reset.

Bit definitions for the 2B errored seconds counter are found in Table 45–47.

Table 45–47—2B errored seconds counter bit definitions

Bit(s)	Name	Description	R/W[a]
1.91.15:8	Reserved	Value always 0	RO
1.91.7:0	Errored seconds [7:0]	The byte of the counter	RO

[a]RO = Read only

45.2.1.51 2B link partner errored seconds register (Register 1.92)

The 2B link partner errored seconds register provides the -O STA with a snapshot of the -R link partner's 2B errored seconds counter. Because this register is not a counter, its value will be updated only when refreshed. The contents of this register are refreshed when the STA activates the "Get link partner parameter" command in the Link Partner PMA/PMD control register (see 45.2.1.17) and the command completes successfully (see 45.2.1.18.1).

This register is defined for -O port subtypes only.

The bit definitions for this register are found in Table 45–47.

45.2.1.52 2B severely errored seconds counter (Register 1.93)

This 8-bit counter contains the number severely errored seconds (see 63.2.2.3). These bits shall be set to all zeros when the register is read by management or upon an MMD reset.

Bit definitions for the 2B severely errored seconds register are found in Table 45–48.

Table 45–48—2B severely errored counter register bit definitions

Bit(s)	Name	Description	R/W[a]
1.93.15:8	Reserved	Value always 0	RO
1.93.7:0	Severely errored seconds [7:0]	The byte of the counter	RO

[a]RO = Read only

45.2.1.53 2B link partner severely errored seconds register (Register 1.94)

The 2B link partner severely errored seconds register provides the -O STA with a snapshot of the -R link partner's 2B severely errored seconds counter. Because this register is not a counter, its value will be updated only when refreshed. The contents of this register are refreshed when the STA activates the "Get link partner parameter" command in the Link Partner PMA/PMD control register (see 45.2.1.17) and the command completes successfully (see 45.2.1.18.1).

This register is defined for -O port subtypes only.

The bit definitions for this register are found in Table 45–48.

45.2.1.54 2B LOSW counter (Register 1.95)

This 8-bit counter contains the number of loss of sync seconds (see 63.2.2.3). These bits shall be set to all zeros when the register is read by management or upon an MMD reset.

Bit definitions for the 2B LOSW counter are found in Table 45–49.

Table 45–49—2B LOSW counter bit definitions

Bit(s)	Name	Description	R/W[a]
1.95.15:8	Reserved	Value always 0	RO
1.95.7:0	Loss of sync seconds [7:0]	The byte of the counter	RO

[a]RO = Read only

45.2.1.55 2B link partner LOSW register (Register 1.96)

The 2B link partner LOSW register provides the -O STA with a snapshot of the -R link partner's 2B LOSW counter. Because this register is not a counter, its value will be updated only when refreshed. The contents of this register are refreshed when the STA activates the "Get link partner parameter" command in the Link partner PMA/PMD control register (see 45.2.1.17) and the command completes successfully (see 45.2.1.18.1).

This register is defined for -O port subtypes only.

The bit definitions for this register are found in Table 45–49.

45.2.1.56 2B unavailable seconds counter (Register 1.97)

This 8-bit counter contains the number of unavailable seconds (see 63.2.2.3). These bits shall be set to all zeros when the register is read by management or upon an MMD reset.

Bit definitions for the 2B unavailable seconds counter are found in Table 45–50.

Table 45–50—2B unavailable seconds counter bit definitions

Bit(s)	Name	Description	R/W[a]
1.97.15:8	Reserved	Value always 0	RO
1.97.7:0	unavailable seconds [7:0]	The byte of the counter	RO

[a]RO = Read only

45.2.1.57 2B link partner unavailable seconds register (Register 1.98)

The 2B link partner unavailable seconds register provides the -O STA with a snapshot of the -R link partner's 2B unavailable seconds counter. Because this register is not a counter, its value will be updated only when refreshed. The contents of this register are refreshed when the STA activates the "Get link partner parameter" command in the Link Partner PMA/PMD control register (see 45.2.1.17) and the command completes successfully (see 45.2.1.18.1).

This register is defined for -O port subtypes only.

The bit definitions for this register are found in Table 45–50.

45.2.1.58 2B state defects register (Register 1.99)

The 2B state defects register is used to communicate defect states from the 2BASE-TL PMD (see 63.2.2.3). The thresholds for these defects are set using the 2B line quality threshold register (see 45.2.1.24). The register bits are cleared to zero when read by the STA or upon MMD reset. On a -R PMA/PMD, these bits are also cleared to zero upon the successful reception of a "Get link partner parameters" command (see 45.2.1.16.1).

Bit definitions for the 2B state defects register are found in Table 45–51.

Table 45–51—2B state defects register bit definitions

Bit(s)	Name	Description	R/W[a]
1.99.15	Segment defect	1 = segment defect detected 0 = normal condition	RO, LH
1.99.14	SNR margin defect	1 = SNR margin defect detected 0 = normal condition	RO, LH
1.99.13	Loop attenuation defect	1 = loop attenuation defect detected 0 = normal condition	RO, LH
1.99.12	Loss of sync word	1 = loss of sync word detected 0 = normal condition	RO, LH
1.99.11:0	Reserved	Value always 0	RO

[a]RO = Read only, LH = Latching high

45.2.1.58.1 Segment defect (1.99.15)

When read as a one, this bit indicates that the local PMA/PMD has detected a segment defect.

45.2.1.58.2 SNR margin defect (1.99.14)

When read as a one, this bit indicates that the local PMA/PMD has received a signal whose SNR is below the set threshold (see 45.2.1.24).

45.2.1.58.3 Loop attenuation defect (1.99.13)

When read as a one, this bit indicates that the PMA/PMD has detected that the loop attenuation is below the set threshold (see 45.2.1.24).

45.2.1.58.4 Loss of sync word (1.99.12)

When read as a one, this bit indicates that the PMA/PMD has lost PMA/PMD frame sync.

45.2.1.59 2B link partner state defects register (Register 1.100)

The 2B link partner state defects register provides the -O STA with a snapshot of the -R link partner's 2B state defects register. The contents of this register are refreshed when the STA activates the "Get link partner parameter" command in the Link Partner PMA/PMD control register (see 45.2.1.17) and the command completes successfully (see 45.2.1.18.1).

This register is defined for -O port subtypes only.

The bit definitions for this register are found in Table 45–51.

45.2.1.60 2B negotiated constellation register (Register 1.101)

The bit definitions for this register are shown in Table 45–52.

Table 45–52—2B register bit definition

Bit(s)	Name	Description	R/W[a]
1.101.15:2	Reserved	Value always 0	RO
1.101.1:0	Negotiated constellation	11 = reserved 10 = 16-TCPAM 01 = 32-TCPAM 00 = undetermined	RO

[a]RO = Read only

45.2.1.60.1 Negotiated constellation (1.101.1:0)

These bits report the resulting constellation that was obtained after initialization. For more information on configuring 2BASE-TL PMA/PMD link initialization, see the 2B PMD parameter registers (see 45.2.1.47). When read as 10 or 01, the constellation has been set as either 16- or 32-TCPAM respectively. When read as 00, the local PMD has not arrived at a constellation with its link partner (as may be the case while link is down or initializing, after reset or upon a failed initialization).

45.2.1.61 2B extended PMD parameters registers (Registers 1.102 through 1.109)

The 2B extended PMD parameters registers define four additional data range sets to be used in conjunction with the 2B PMD parameters registers when additional PMD configuration detail is desired. For a complete description of the use of these registers, see 45.2.1.47.

Bit definitions for these registers can be found in Table 45–53.

Table 45–53— 2B extended PMD parameters registers bit definition

Bit(s)	Name	Description	R/W[a]
1.102.15	Reserved	Value always 0	RO
1.102.14:8	Min data rate5	Min data rate of the fifth range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.102.7	Reserved	Value always 0	RO
1.102.6:0	Max data rate5	Max data rate of the fifth range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.103.15:14	Reserved	Value always 0	RO
1.103.13:7	Data rate step5	Data rate step of the fifth range *n* = value of the bits, *n* valid 1 to 86 Data Rate = 64*n* kb/s	O: R/W R: RO
1.103.6:2	Power5	*x* = multiple of 0.5 dBm to add to 5 dBm offset Power = (5 + 0.5*x*) dBm	O: R/W R: RO
1.103.1:0	Constellation5	11 = reserved 10 = 16-TCPAM 01 = 32-TCPAM 00 = automatically set by the PHY	O: R/W R: RO
1.104.15	Reserved	Value always 0	RO
1.104.14:8	Min data rate6	Min data rate of the sixth range *n* = value of the bits, *n* valid 3 to 89 Data Rate =64*n* kb/s	O: R/W R: RO
1.104.7	Reserved	Value always 0	RO
1.104.6:0	Max data rate6	Max data rate of the sixth range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.105.15:14	Reserved	Value always 0	RO
1.105.13:7	Data rate step6	Data rate step of the sixth range *n* = value of the bits, *n* valid 1 to 86 Data Rate = 64*n* kb/s	O: R/W R: RO
1.105.6:2	Power6	*x* = multiple of 0.5 dBm to add to 5 dBm offset Power = (5 + 0.5*x*) dBm	O: R/W R: RO
1.105.1:0	Constellation6	11 = reserved 10 = 16-TCPAM 01 = 32-TCPAM 00 = automatically set by the PHY	O: R/W R: RO
1.106.15	Reserved	Value always 0	RO
1.106.14:8	Min data rate7	Min data rate of the seventh range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.106.7	Reserved	Value always 0	RO

Table 45–53— 2B extended PMD parameters registers bit definition *(continued)*

Bit(s)	Name	Description	R/W[a]
1.106.6:0	Max data rate7	Max data rate of the seventh range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.107.15:14	Reserved	Value always 0	RO
1.107.13:7	Data rate step7	Data rate step of the seventh range *n* = value of the bits, *n* valid 1 to 86 Data Rate = 64*n* kb/s	O: R/W R: RO
1.107.6:2	Power7	*x* = multiple of 0.5 dBm to add to 5 dBm offset Power = (5 + 0.5*x*) dBm	O: R/W R: RO
1.107.1:0	Constellation7	11 = reserved 10 = 16-TCPAM 01 = 32-TCPAM 00 = automatically set by the PHY	O: R/W R: RO
1.108.15	Reserved	Value always 0	RO
1.108.14:8	Min data rate8	Min data rate of the eighth range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.108.7	Reserved	Value always 0	RO
1.108.6:0	Max data rate8	Max data rate of the eighth range *n* = value of the bits, *n* valid 3 to 89 Data Rate = 64*n* kb/s	O: R/W R: RO
1.109.15:14	Reserved	Value always 0	RO
1.109.13:7	Data rate step8	Data rate step of the eighth range *n* = value of the bits, *n* valid 1 to 86 Data Rate = 64*n* kb/s	O: R/W R: RO
1.109.6:2	Power8	*x* = multiple of 0.5 dBm to add to 5 dBm offset Power = (5 + 0.5*x*) dBm	O: R/W R: RO
1.109.1:0	Constellation8	11 = reserved 10 = 16-TCPAM 01 = 32-TCPAM 00 = automatically set by the PHY	O: R/W R: RO

[a]R/W = Read/Write, RO = Read only

45.2.1.61.1 Minimum data rate (1.102, 1.104, 1.106, 1.108. Bits 14:8)

Bits 14:8 in registers 1.102, 1.104, 1.106, and 1.108 set the minimum data rate for each of the four extended ranges. Valid values for these bits are decimal 3 through 89. writes to set an invalid value shall be ignored. The rate is expressed in units of kb/s and is derived by multiplying the decimal value of bits 14:8 by 64.

45.2.1.61.2 Max data rate (1.102, 1.104, 1.106, 1.108. Bits 6:0)

Bits 6:0 in registers 1.102, 1.104, 1.106, and 1.108 set the maximum data rate for each of the four extended ranges. Valid values for these bits are decimal 3 through 89, writes to set an invalid value shall be ignored. The rate is expressed in units of kb/s and is derived by multiplying the decimal value of bits 6:0 by 64.

45.2.1.61.3 Data rate step (1.103, 1.105, 1.107, 1.109. Bits 13:7)

Bits 13:7 in registers 1.103, 1.105, 1.107 and 1.109 set the granularity used by the PMA/PMD when determining the line rate. Valid values for these bits are decimal 1 through 86, writes to set an invalid value shall be ignored. The data rate step is expressed in units of kb/s and is derived by multiplying the decimal value of bits 13:7 by 64.

45.2.1.61.4 Power (1.103, 1.105, 1.107, 1.109. Bits 6:2)

Bits 6:2 in registers 1.103, 1.105, 1.107 and 1.109 set the allowed power level for each extended data rate range. The power levels set in these bits override those of the annex selected in the 2B general parameter register (1.80). The power level is expressed in units of dBm and is derived by Equation (45–5), where *x* equals the value of bits 6:2.

$$\text{power} = \left(5 + \frac{x}{2}\right) \text{dBm} \tag{45–5}$$

45.2.1.61.5 Constellation (1.103, 1.105, 1.107, 1.109. Bits 1:0)

Bits 1:0 in registers 1.103, 1.105, 1.107 and 1.109 set the allowed constellation for each extended data rate range. Setting a value of 10 or 01 restricts the constellation to 16- or 32-TCPAM respectively. When set to a value of 00, the PMD automatically determines the constellation during initialization. Attempts to set a value of 11 shall be ignored.

45.2.1.62 10GBASE-T status (Register 1.129)

The assignments of in the 10GBASE-T status register is shown in Table 45–54.

Table 45–54—10GBASE-T status register bit definitions

Bit(s)	Name	Description	R/W[a]
1.129.15:1	Reserved	Value always 0	RO
1.129.0	LP information valid	1 = Link partner information is valid 0 = Link partner information is invalid	RO

[a]RO = Read only

45.2.1.62.1 LP information valid (1.129.0)

When read as a one, bit 1.129.0 indicates that the startup protocol defined in 55.4.2.5 has been completed, and that the contents of bits 1.130.11:0, 1.131.15:10, 1.145.14:8, 1.146.14:8, and 1.146.6:0, which are established during the startup protocol, are valid. When read as a zero, bit 1.129.0 indicates that the startup process has not been completed, and that the contents of these bits that are established during the startup protocol are invalid. A 10GBASE-T PMA shall return a value of zero in bit 1.129.0 if PMA link_status=FAIL.

45.2.1.63 10GBASE-T pair swap and polarity register (Register 1.130)

Register 1.130 reflects the status of the pair-to-pair connectivity between the local device and the link partner (see Table 45–55). The polarity of each individual pair is also indicated.

Table 45–55—10GBASE-T pair swap and polarity register bit definitions

Bit(s)	Name	Description	R/W[a]
1.130.15:12	Reserved	Value always 0	RO
1.130.11	Pair D polarity	1 = Polarity of pair D is reversed 0 = Polarity of pair D is not reversed	RO
1.130.10	Pair C polarity	1 = Polarity of pair C is reversed 0 = Polarity of pair C is not reversed	RO
1.130.9	Pair B polarity	1 = Polarity of pair B is reversed 0 = Polarity of pair B is not reversed	RO
1.130.8	Pair A polarity	1 = Polarity of pair A is reversed 0 = Polarity of pair A is not reversed	RO
1.130.7:2	Reserved	Value always 0	RO
1.130.1:0	MDI/MDI-X connection	1 0 1 1 = No crossover 1 0 = Reserved 0 1 = Reserved 0 0 = Pair A/B and pair C/D crossover	RO

[a]RO = Read only

45.2.1.63.1 Pair D polarity (1.130.11)

When read as zero, bit 1.130.11 indicates that the polarity within pair D is not reversed. When read as one, bit 1.130.11 indicates that the polarity of pair D is reversed.

45.2.1.63.2 Pair C polarity (1.130.10)

When read as zero, bit 1.130.10 indicates that the polarity within pair C is not reversed. When read as one, bit 1.130.10 indicates that the polarity of pair C is reversed.

45.2.1.63.3 Pair B polarity (1.130.9)

When read as zero, bit 1.130.9 indicates that the polarity within pair B is not reversed. When read as one, bit 1.130.9 indicates that the polarity of pair B is reversed.

45.2.1.63.4 Pair A polarity (1.130.8)

When read as zero, bit 1.130.8 indicates that the polarity within pair A is not reversed. When read as one, bit 1.130.8 indicates that the polarity of pair A is reversed.

45.2.1.63.5 MDI/MDI-X connection (1.130.1:0)

Bits 1.130.1:0 indicate the resolution of the automatic MDI/MDI-X configuration defined in 55.4.4. This function is intended to eliminate the need for crossover cables between similar devices. The automatic configuration method used shall comply with 40.4.4.1 and 40.4.4.2.

45.2.1.64 10GBASE-T TX power backoff and PHY short reach setting (Register 1.131)

The complete assignment of bits in the 10GBASE-T TX power backoff and short reach mode settings register is shown in Table 45–56.

Table 45–56—10GBASE-T TX power backoff and PHY short reach setting register bit definitions

Bit(s)	Name	Description	R/W[a]
1.131.15:13	Link partner TX power backoff setting	15 14 13 1 1 1 = 14 dB 1 1 0 = 12 dB 1 0 1 = 10 dB 1 0 0 = 8 dB 0 1 1 = 6 dB 0 1 0 = 4 dB 0 0 1 = 2 dB 0 0 0 = 0 dB	RO
1.131.12:10	TX power backoff setting	12 11 10 1 1 1 = 14 dB 1 1 0 = 12 dB 1 0 1 = 10 dB 1 0 0 = 8 dB 0 1 1 = 6 dB 0 1 0 = 4 dB 0 0 1 = 2 dB 0 0 0 = 0 dB	RO
1.131.9:1	Reserved	Value always 0	RO
1.131.0	Short reach mode	1 = PHY is operating in short reach mode 0 = PHY is not operating in short reach mode	R/W

[a]RO = Read only, R/W = Read/Write

45.2.1.64.1 10GBASE-T TX power backoff settings (1.131.15:10)

The 10GBASE-T TX power backoff settings reflects the TX power backoff selected during the startup negotiation process.The startup negotiation process and all TX power backoff settings are defined in 55.4.2.5 and 55.4.5.1. If LP information valid bit, 1.129.0, is set to one then bits 1.131.15:13 indicates the TX power backoff setting of the link partner and bits 1.131.12:10 indicates the TX power backoff setting of the local device.

45.2.1.64.2 PHY short reach mode (1.131.0)

The short reach mode of the 10GBASE-T PHY provides a means for operation on a cable plant that has parametric performance equivalent to 30 m of Class F and Class E_A cabling as defined in 55.5.4.5. If bit 1.131.0 is a one, the PHY is in short reach mode. If bit 1.131.0 is a zero the PHY is operating in normal mode. The default value for this bit is zero.

45.2.1.65 10GBASE-T test mode register (Register 1.132)

The assignment of bits in the 10GBASE-T test mode register is shown in Table 45–57. The default values for each bit should be chosen so that the initial state of the device upon power up or reset is a normal operational state without management intervention.

45.2.1.65.1 Test mode control (1.132.15:13)

Transmitter test mode operations defined by bits 1.132.15:13, are described in 55.5.2 and Table 55–12. The default value for bits 1.132.15:13 is zero.

Table 45–57—10GBASE-T test mode register bit definitions

Bit(s)	Name	Description	R/W[a]
1.132.15:13	Test mode control	15 14 13 1 1 1 = Test mode 7 1 1 0 = Test mode 6 1 0 1 = Test mode 5 1 0 0 = Test mode 4 0 1 1 = Test mode 3 0 1 0 = Test mode 2 0 0 1 = Test mode 1 0 0 0 = Normal operation	R/W
1.132.12:10	Transmitter test frequencies	12 11 10 1 1 1 = Reserved 1 1 0 = Dual tone 5 1 0 1 = Dual tone 4 1 0 0 = Dual tone 3 0 1 1 = Reserved 0 1 0 = Dual tone 2 0 0 1 = Dual tone 1 0 0 0 = Reserved	R/W
1.132.9:0	Reserved	Value always 0	RO

[a]R/W = Read/Write, RO = Read only

45.2.1.65.2 Transmitter test frequencies (1.132.12:10)

When test mode 4 is selected by setting bits 1.132.15:13 to one, zero, zero respectively, bits 1.132.12:10 select the transmit test frequency as shown in Table 45–57. Detailed use and operation of these transmitter test frequencies is described in 55.5.2.

45.2.1.66 SNR operating margin channel A register (Register 1.133)

Register 1.133 contains the current SNR operating margin measured at the slicer input for channel A for the 10GBASE-T PMA. It is reported with 0.1 dB of resolution to an accuracy of 0.5 dB within the range of –12.7 dB to 12.7 dB. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.67 SNR operating margin channel B register (Register 1.134)

Register 1.134 contains the current SNR operating margin measured at the slicer input for channel B for the 10GBASE-T PMA. It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –12.7 dB to 12.7 dB. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.68 SNR operating margin channel C register (Register 1.135)

Register 1.135 contains the current SNR operating margin measured at the slicer input for channel C for the 10GBASE-T PMA. It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –12.7 dB to 12.7 dB. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.69 SNR operating margin channel D register (Register 1.136)

Register 1.136 contains the current SNR operating margin measured at the slicer input for channel D for the 10GBASE-T PMA. It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –12.7 dB to 12.7 dB. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.70 Minimum margin channel A register (Register 1.137)

The Minimum margin channel A register contains a latched copy of the lowest value observed in the SNR operating margin channel A register (1.133) since the last read. It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –12.7 dB to 12.7 dB. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.71 Minimum margin channel B register (Register 1.138)

The Minimum margin channel B register contains a latched copy of the lowest value observed in the SNR operating margin channel B register (1.134) since the last read. It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –12.7 dB to 12.7 dB. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.72 Minimum margin channel C register (Register 1.139)

The Minimum margin channel C register contains a latched copy of the lowest value observed in the SNR operating margin channel C register (1.135) since the last read. It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –12.7 dB to 12.7 dB. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.73 Minimum margin channel D register (Register 1.140)

The Minimum margin channel D register contains a latched copy of the lowest value observed in the SNR operating margin channel D register (1.136) since the last read. It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –12.7 dB to 12.7 dB. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.74 RX signal power channel A register (Register 1.141)

The RX signal power channel A register is read only and contains the receive signal power measured at the MDI during training as described in 55.4.3.1. The RX signal power should reflect the power measured when the device transitions out of the state PMA_Training_Init_M or PMA_Training_Init_S (as appropriate, see 55.4.6.1), when the link partner is transmitting with PBO_tx = 4 (8 dB power backoff). It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –20 dBm to 5.5 dBm. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.75 RX signal power channel B register (Register 1.142)

The RX signal power channel B register is read only and contains the receive signal power measured at the MDI during training as described in 55.4.3.1. The RX signal power should reflect the power measured when the device transitions out of the state PMA_Training_Init_M or PMA_Training_Init_S (as appropriate, see 55.4.6.1), when the link partner is transmitting with PBO_tx = 4 (8 dB power backoff). It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –20 dBm to 5.5 dBm. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.76 RX signal power channel C register (Register 1.143)

The RX signal power channel C register is read only and contains the receive signal power measured at the MDI during training as described in 55.4.3.1. The RX signal power should reflect the power measured when the device transitions out of the state PMA_Training_Init_M or PMA_Training_Init_S (as appropriate, see 55.4.6.1), when the link partner is transmitting with PBO_tx = 4 (8 dB power backoff). It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –20 dBm to 5.5 dBm. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.77 RX signal power channel D register (Register 1.144)

The RX signal power channel D register is read only and contains the receive signal power measured at the MDI during training as described in 55.4.3.1. The RX signal power should reflect the power measured when the device transitions out of the state PMA_Training_Init_M or PMA_Training_Init_S (as appropriate, see 55.4.6.1), when the link partner is transmitting with PBO_tx = 4 (8 dB power backoff). It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –20 dBm to 5.5 dBm. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.78 10GBASE-T skew delay register (Registers 1.145 and 1.146)

The skew delay register reports the current skew delay on each of the pair with respect to physical pair A (see Table 45–58). It is reported with 1.25 ns resolution to an accuracy of 2.5 ns. The number reported is in two's complement notation with positive values representing delay and negative values representing advance with respect to physical pair A. If the delay exceed the maximum amount that can be represented by the range (–80 ns to +78.75 ns), the field displays the maximum respective value. The value shall be updated at least once per second.

Table 45–58—10GBASE-T skew delay register bit definitions

Bit(s)	Name	Description	R/W[a]
1.145.15	Reserved	Value always 0	RO
1.145.14:8	Skew delay B	Skew delay for pair B	RO
1.145.7:0	Reserved	Value always 0	RO
1.146.15	Reserved	Value always 0	RO
1.146.14:8	Skew delay D	Skew delay for pair D	RO
1.146.7	Reserved	Value always 0	RO
1.146.6:0	Skew delay C	Skew delay for pair C	RO

[a]RO = Read only

45.2.1.79 10GBASE-T fast retrain status and control register (Register 1.147)

Table 45–59—10GBASE-T fast retrain status and control register bit definitions

Bit(s)	Name	Description	R/W[a]
1.147.15:11	LP fast retrain count	Counts the number of fast retrains requested by the link partner	RO/NR
1.147.10:6	LD fast retrain count	Counts the number of fast retrains requested by the local device	RO/NR
1.147.5	Reserved	Value always 0	RO
1.147.4	Fast retrain ability	1 = Fast retrain capability is supported 0 = Fast retrain capability is not supported	RO
1.147.3	Fast retrain negotiated	1 = Fast retrain capability was negotiated 0 = Fast retrain capability was not negotiated	RO
1.147.2:1	Fast retrain signal type	11 = Reserved 10 = PHY signals Link Interruption during fast retrain 01 = PHY signals Local Fault during fast retrain 00 = PHY signals IDLE during fast retrain	R/W
1.147.0	Fast retrain enable	1 = Fast retrain capability is enabled 0 = Fast retrain capability is disabled	R/W

[a] RO = Read only, R/W = Read/Write, NR = Non Roll-over

45.2.1.79.1 LP fast retrain count (1.147.15:11)

These bits map to fr_rx_counter as defined in 55.4.5.1. The counter is a 5-bit count of the number of 10GBASE-T fast retrains requested by the link partner. These bits shall be reset to all zeros when read or upon execution of the PMA reset. These bits shall be held at all ones in the case of overflow.

45.2.1.79.2 LD fast retrain count (1.147.10:6)

These bits map to fr_tx_counter as defined in 55.4.5.1. The counter is a 5-bit count of the number of 10GBASE-T fast retrains requested by the local device. These bits shall be reset to all zeros when read or upon execution of the PMA reset. These bits shall be held at all ones in the case of overflow.

45.2.1.79.3 Fast retrain ability (1.147.4)

When read as a one, bit 1.147.4 indicates that the PHY supports fast retrain, as defined in 55.4.2.5.15. When read as a zero, bit 1.147.4 indicates that the PHY does not support fast retrain.

45.2.1.79.4 Fast retrain negotiated (1.147.3)

When read as a one, bit 1.147.3 indicates that the PHY negotiated fast retrain, as defined in 55.4.2.5.15 during the most recent auto-negotiation. This is the condition where both the local device indicated fast retrain ability (bit 7.32.1 is one) and the link partner indicated fast retrain ability (bit 7.33.1 is one). When read as a zero, bit 1.147.3 indicates that the PHY did not negotiate fast retrain. See 45.2.7.10.6.

45.2.1.79.5 Fast retrain signal type (1.147.2:1)

For PHYs that support fast retrain, these bits map to fr_sigtype as defined in 55.3.6.2.2. When Fast retrain signal type is set to 00, the PMA sends IDLE characters on the receive path during fast retrain. When Fast

retrain signal type is set to 01, the PMA sends Local Fault on the receive path during fast retrain. When Fast retrain signal type is set to 10, the PMA sends Link Interruption on the receive path during fast retrain.

45.2.1.79.6 Fast retrain enable (1.147.0)

For PHYs that support fast retrain, this bit controls fr_enable as defined in 55.4.5.1. When PMA reset is executed, this bit is set to one.

NOTE—Setting this bit to zero while a link is up will cause the PHY to stop supporting fast retrain, and the link will drop if the link partner initiates a fast retrain.

45.2.1.80 BASE-R PMD control register (Register 1.150)

The BASE-R PMD control register is used for 10GBASE-KR and other PHY types using the PMDs described in Clause 72, Clause 84, Clause 85, Clause 92, Clause 93, or Clause 94. The assignment of bits in the BASE-R PMD control register is shown in Table 45–60.

Table 45–60—BASE-R PMD control register

Bit(s)	Name	Description	R/W[a]
1.150.15:2	Reserved	Value always 0	RO
1.150.1	Training enable	1 = Enable the BASE-R start-up protocol 0 = Disable the BASE-R start-up protocol	R/W
1.150.0	Restart training	1 = Restart BASE-R start-up protocol 0 = Normal operation	R/W SC

[a]R/W = Read/Write, SC = Self-clearing, RO = Read only

45.2.1.80.1 Restart training (1.150.0)

This bit maps to the state variable mr_restart_training as defined in 72.6.10.3.1.

45.2.1.80.2 Training enable (1.150.1)

This bit maps to the state variable mr_training_enable as defined in 72.6.10.3.1.

45.2.1.81 BASE-R PMD status register (Register 1.151)

The BASE-R PMD status register is used for 10GBASE-KR and other PHY types using the PMDs described in Clause 72, Clause 84, Clause 85, Clause 92, Clause 93, or Clause 94. The assignment of bits in the BASE-R PMD status register is shown in Table 45–61.

Table 45–61—BASE-R PMD status register

Bit(s)	Name	Description	R/W[a]
1.151.15	Training failure 3	1 = Training failure has been detected for lane 3 0 = Training failure has not been detected for lane 3	RO
1.151.14	Start-up protocol status 3	1 = Start-up protocol in progress for lane 3 0 = Start-up protocol complete for lane 3	RO
1.151.13	Frame lock 3	1 = Training frame delineation detected for lane 3 0 = Training frame delineation not detected for lane 3	RO
1.151.12	Receiver status 3	1 = Receiver trained and ready to receive data for lane 3 0 = Receiver training for lane 3	RO
1.151.11	Training failure 2	1 = Training failure has been detected for lane 2 0 = Training failure has not been detected for lane 2	RO
1.151.10	Start-up protocol status 2	1 = Start-up protocol in progress for lane 2 0 = Start-up protocol complete for lane 2	RO
1.151.9	Frame lock 2	1 = Training frame delineation detected for lane 2 0 = Training frame delineation not detected for lane 2	RO
1.151.8	Receiver status 2	1 = Receiver trained and ready to receive data for lane 2 0 = Receiver training for lane 2	RO
1.151.7	Training failure 1	1 = Training failure has been detected for lane 1 0 = Training failure has not been detected for lane 1	RO
1.151.6	Start-up protocol status 1	1 = Start-up protocol in progress for lane 1 0 = Start-up protocol complete for lane 1	RO
1.151.5	Frame lock 1	1 = Training frame delineation detected for lane 1 0 = Training frame delineation not detected for lane 1	RO
1.151.4	Receiver status 1	1 = Receiver trained and ready to receive data for lane 1 0 = Receiver training for lane 1	RO
1.151.3	Training failure 0	1 = Training failure has been detected for lane 0 0 = Training failure has not been detected for lane 0	RO
1.151.2	Start-up protocol status 0	1 = Start-up protocol in progress for lane 0 0 = Start-up protocol complete for lane 0	RO
1.151.1	Frame lock 0	1 = Training frame delineation detected for lane 0 0 = Training frame delineation not detected for lane 0	RO
1.151.0	Receiver status 0	1 = Receiver trained and ready to receive data for lane 0 0 = Receiver training for lane 0	RO

[a]RO = Read only

45.2.1.81.1 Receiver status 0 (1.151.0)

This bit maps to the state variable rx_trained as defined in 72.6.10.3.1.

45.2.1.81.2 Frame lock 0 (1.151.1)

This bit maps to the state variable frame_lock as defined in 72.6.10.3.1.

45.2.1.81.3 Start-up protocol status 0 (1.151.2)

This bit maps to the state variable training as defined in 72.6.10.3.1.

45.2.1.81.4 Training failure 0 (1.151.3)

This bit maps to the state variable training_failure as defined in 72.6.10.3.1.

45.2.1.81.5 Receiver status 1, 2, 3 (1.151.4, 1.151.8, 1.151.12)

These bits are defined identically to 1.151.0 for lanes 1, 2, and 3 respectively.

45.2.1.81.6 Frame lock 1, 2, 3 (1.151.5, 1.151.9, 1.151.13)

These bits are defined identically to 1.151.1 for lanes 1, 2, and 3 respectively.

45.2.1.81.7 Start-up protocol status 1, 2, 3 (1.151.6, 1.151.10, 1.151.14)

These bits are defined identically to 1.151.2 for lanes 1, 2, and 3 respectively.

45.2.1.81.8 Training failure 1, 2, 3 (1.151.7, 1.151.11, 1.151.15)

These bits are defined identically to 1.151.3 for lanes 1, 2, and 3 respectively.

45.2.1.82 BASE-R LP coefficient update, lane 0 register (Register 1.152)

The BASE-R LP coefficient update, lane 0 register is used for 10GBASE-KR and other PHY types using the PMDs described in Clause 72, Clause 84, Clause 85, Clause 92, Clause 93, or Clause 94. The BASE-R LP coefficient update, lane 0 register reflects the contents of the first 16-bit word of the training frame most recently received from the control channel for lane 0 or for a single-lane PHY.

The assignment of bits in the BASE-R LP coefficient update, lane 0 register is shown in Table 45–62. Normally the bits in this register are read only; however, when training is disabled by setting low bit 1 in the BASE-R PMD control register, the BASE-R LP coefficient update, lane 0 register becomes writeable.

A copy of this register may be implemented at address 1.1100 to assist PHY access for devices using post-read-increment-address access for a multi-lane PCS. If implemented, all accesses to the copy have identical behavior as accesses to the original register.

45.2.1.82.1 Preset (1.152.13)

The preset control bit requests that the coefficients be set to a state where equalization is turned off. The function and values of the preset bit is defined in 72.6.10.2.3.1.

45.2.1.82.2 Initialize (1.152.12)

The initialize control is sent to request that the coefficients be set to configure the transmit equalizer to its INITIALIZE state. The function and values of the initialize bit is defined in 72.6.10.2.3.2.

45.2.1.82.3 Coefficient (k) update (1.152.5:0)

Each coefficient, k, is assigned a 2-bit field describing a requested update. Three request encodings are defined: increment, decrement, and hold. The valid range for k is –1 to +1 where k = 0 denotes the main, or gain, tap. The function and values of the coefficient (k) update bits are defined in 72.6.10.2.3.3.

Table 45–62—BASE-R LP coefficient update, lane 0 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.152.15:14	Reserved	Value always 0	RO
1.152.13	Preset	1 = Preset coefficients 0 = Normal operation	R/W
1.152.12	Initialize	1 = Initialize coefficients 0 = Normal operation	R/W
1.152.11:6	Reserved	Value always 0	RO
1.152.5:4	Coefficient (+1) update	5 4 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold	R/W
1.152.3:2	Coefficient (0) update	3 2 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold	R/W
1.152.1:0	Coefficient (–1) update	1 0 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold	R/W

[a]R/W = Read/Write, RO = Read only

45.2.1.83 BASE-R LP status report, lane 0 register (Register 1.153)

The BASE-R LP status report, lane 0 register is used for 10GBASE-KR and other PHY types using the PMDs described in Clause 72, Clause 84, Clause 85, Clause 92, Clause 93, or Clause 94. The BASE-R LP status report, lane 0 register reflects the contents of the second 16-bit word of the training frame most recently received from the control channel for lane 0 or for a single-lane PHY

The assignment of bits in the BASE-R LP status report, lane 0 register is shown in Table 45–63.

A copy of this register may be implemented at address 1.1200 to assist PHY access for devices using post-read-increment-address access for a multi-lane PCS. If implemented, all accesses to the copy have identical behavior as the original register.

Table 45–63—BASE-R LP status report, lane 0 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.153.15	Receiver ready	1 = The LP receiver has determined that training is complete and is prepared to receive data 0 = The LP receiver is requesting that training continue	RO
1.153.14:6	Reserved	Value always 0	RO
1.153.5:4	Coefficient (+1) status	5 4 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated	RO
1.153.3:2	Coefficient (0) status	3 2 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated	RO
1.153.1:0	Coefficient (–1) status	1 0 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated	RO

[a]RO = Read only

45.2.1.83.1 Receiver ready (1.153.15)

The function and values for the receiver ready bit are defined in 72.6.10.2.4.4.

45.2.1.83.2 Coefficient (k) status (1.153.5:0)

The function and values for the coefficient status bits are defined in 72.6.10.2.4.5.

45.2.1.84 BASE-R LD coefficient update, lane 0 register (Register 1.154)

The BASE-R LD coefficient update, lane 0 register is used for 10GBASE-KR and other PHY types using the PMDs described in Clause 72, Clause 84, Clause 85, Clause 92, Clause 93, or Clause 94. The BASE-R LD coefficient update, lane 0 register reflects the contents of the first 16-bit word of the outgoing training frame as defined by the LD receiver adaptation process in 72.6.10.2.5 for lane 0 or for a single-lane PHY.

The assignment of bits in the BASE-R LD coefficient update, lane 0 register is shown in Table 45–64.

A copy of this register may be implemented at address 1.1300 to assist PHY access for devices using post-read-increment-address access for a multi-lane PCS. If implemented, all accesses to the copy have identical behavior as the original register.

45.2.1.84.1 Preset (1.154.13)

The function and values of the preset bit is defined in 72.6.10.2.3.1.

45.2.1.84.2 Initialize (1.154.12)

The function and values of the initialize bit is defined in 72.6.10.2.3.2.

45.2.1.84.3 Coefficient (k) update(1.154.5:0)

The function and values of the coefficient (k) update bits are defined in 72.6.10.2.3.3.

Table 45–64—BASE-R LD coefficient update, lane 0 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.154.15:14	Reserved	Value always 0	RO
1.154.13	Preset	1 = Preset coefficients 0 = Normal operation	RO
1.154.12	Initialize	1 = Initialize coefficients 0 = Normal operation	RO
1.154.11:6	Reserved	Value always 0	RO
1.154.5:4	Coefficient (+1) update	5 4 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold	RO
1.154.3:2	Coefficient (0) update	3 2 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold	RO
1.154.1:0	Coefficient (–1) update	1 0 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold	RO

[a]RO = Read only

45.2.1.85 BASE-R LD status report, lane 0 register (Register 1.155)

The BASE-R LD status report, lane 0 register is used for 10GBASE-KR and other PHY types using the PMDs described in Clause 72, Clause 84, Clause 85, Clause 92, Clause 93, or Clause 94. The BASE-R LD status report, lane 0 register reflects the contents of the second 16-bit word of the current outgoing training frame, as defined in the training state diagram in Figure 72–5 for lane 0 or for a single-lane PHY.

The assignment of bits in the BASE-R LD status report, lane 0 register is shown in Table 45–65.

A copy of this register may be implemented at address 1.1400 to assist PHY access for devices using post-read-increment-address access for a multi-lane PCS. If implemented, all accesses to the copy have identical behavior as the original register.

Table 45–65—BASE-R LD status report, lane 0 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.155.15	Receiver ready	1 = The LD receiver has determined that training is complete and is prepared to receive data 0 = The LD receiver is requesting that training continue	RO
1.155.14:6	Reserved	Value always 0	RO
1.155.5:4	Coefficient (+1) status	5 4 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated	RO
1.155.3:2	Coefficient (0) status	3 2 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated	RO
1.155.1:0	Coefficient (–1) status	1 0 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated	RO

[a]RO = Read only

45.2.1.85.1 Receiver ready (1.155.15)

The function and values for the receiver ready bit are defined in 72.6.10.2.4.4.

45.2.1.85.2 Coefficient (k) status (1.155.5:0)

The function and values for the coefficient status bits are defined in 72.6.10.2.4.5.

45.2.1.86 BASE-R PMD status 2 register (Register 1.156)

The BASE-R PMD status 2 register is used for 100GBASE-CR10 and other PHY types using the PMDs described in Clause 72, Clause 84, or Clause 85 over more than 4 lanes. The assignment of bits in the BASE-R PMD status 2 register is shown in Table 45–66.

Table 45–66—BASE-R PMD status 2 register

Bit(s)	Name	Description	R/W[a]
1.156.15	Training failure 7	1 = Training failure has been detected for lane 7 0 = Training failure has not been detected for lane 7	RO
1.156.14	Start-up protocol status 7	1 = Start-up protocol in progress for lane 7 0 = Start-up protocol complete for lane 7	RO
1.156.13	Frame lock 7	1 = Training frame delineation detected for lane 7 0 = Training frame delineation not detected for lane 7	RO

Table 45–66—BASE-R PMD status 2 register *(continued)*

1.156.12	Receiver status 7	1 = Receiver trained and ready to receive data for lane 7 0 = Receiver training for lane 7	RO
1.156.11	Training failure 6	1 = Training failure has been detected for lane 6 0 = Training failure has not been detected for lane 6	RO
1.156.10	Start-up protocol status 6	1 = Start-up protocol in progress for lane 6 0 = Start-up protocol complete for lane 6	RO
1.156.9	Frame lock 6	1 = Training frame delineation detected for lane 6 0 = Training frame delineation not detected for lane 6	RO
1.156.8	Receiver status 6	1 = Receiver trained and ready to receive data for lane 6 0 = Receiver training for lane 6	RO
1.156.7	Training failure 5	1 = Training failure has been detected for lane 5 0 = Training failure has not been detected for lane 5	RO
1.156.6	Start-up protocol status 5	1 = Start-up protocol in progress for lane 5 0 = Start-up protocol complete for lane 5	RO
1.156.5	Frame lock 5	1 = Training frame delineation detected for lane 5 0 = Training frame delineation not detected for lane 5	RO
1.156.4	Receiver status 5	1 = Receiver trained and ready to receive data for lane 5 0 = Receiver training for lane 5	RO
1.156.3	Training failure 4	1 = Training failure has been detected for lane 4 0 = Training failure has not been detected for lane 4	RO
1.156.2	Start-up protocol status 4	1 = Start-up protocol in progress for lane 4 0 = Start-up protocol complete for lane 4	RO
1.156.1	Frame lock 4	1 = Training frame delineation detected for lane 4 0 = Training frame delineation not detected for lane 4	RO
1.156.0	Receiver status 4	1 = Receiver trained and ready to receive data for lane 4 0 = Receiver training for lane 4	RO

[a]RO = Read only

45.2.1.86.1 Receiver status 4, 5, 6, 7 (1.156.0, 1.156.4, 1.156.8, 1.156.12)

These bits are defined identically to 1.151.0 for lanes 4, 5, 6, and 7 respectively.

45.2.1.86.2 Frame lock 4, 5, 6, 7 (1.156.1, 1.156.5, 1.156.9, 1.156.13)

These bits are defined identically to 1.151.1 for lanes 4, 5, 6, and 7 respectively.

45.2.1.86.3 Start-up protocol status 4, 5, 6, 7 (1.156.2, 1.156.6, 1.156.10, 1.156.14)

These bits are defined identically to 1.151.2 for lanes 4, 5, 6, and 7 respectively.

45.2.1.86.4 Training failure 4, 5, 6, 7 (1.156.3, 1.156.7, 1.156.11, 1.156.15)

These bits are defined identically to 1.151.3 for lanes 4, 5, 6, and 7 respectively.

45.2.1.87 BASE-R PMD status 3 register (Register 1.157)

The BASE-R PMD status 3 register is used for 100GBASE-CR10 and other PHY types using the PMDs described in Clause 72, Clause 84, or Clause 85 over more than 8 lanes. The assignment of bits in the BASE-R PMD status 3 register is shown in Table 45–67.

Table 45–67—BASE-R PMD status 3 register

Bit(s)	Name	Description	R/W[a]
1.157.15:8	Reserved	Value always 0	RO
1.157.7	Training failure 9	1 = Training failure has been detected for lane 9 0 = Training failure has not been detected for lane 9	RO
1.157.6	Start-up protocol status 9	1 = Start-up protocol in progress for lane 9 0 = Start-up protocol complete for lane 9	RO
1.157.5	Frame lock 9	1 = Training frame delineation detected for lane 9 0 = Training frame delineation not detected for lane 9	RO
1.157.4	Receiver status 9	1 = Receiver trained and ready to receive data for lane 9 0 = Receiver training for lane 9	RO
1.157.3	Training failure 8	1 = Training failure has been detected for lane 8 0 = Training failure has not been detected for lane 8	RO
1.157.2	Start-up protocol status 8	1 = Start-up protocol in progress for lane 8 0 = Start-up protocol complete for lane 8	RO
1.157.1	Frame lock 8	1 = Training frame delineation detected for lane 8 0 = Training frame delineation not detected for lane 8	RO
1.157.0	Receiver status 8	1 = Receiver trained and ready to receive data for lane 8 0 = Receiver training for lane 8	RO

[a]RO = Read only

45.2.1.87.1 Receiver status 8, 9 (1.157.0, 1.157.4)

These bits are defined identically to 1.151.0 for lanes 8 and 9 respectively.

45.2.1.87.2 Frame lock 8, 9 (1.157.1, 1.157.5)

These bits are defined identically to 1.151.1 for lanes 8 and 9 respectively.

45.2.1.87.3 Start-up protocol status 8, 9 (1.157.2, 1.157.6)

These bits are defined identically to 1.151.2 for lanes 8 and 9 respectively.

45.2.1.87.4 Training failure 8, 9 (1.157.3, 1.157.7)

These bits are defined identically to 1.151.3 for lanes 8 and 9 respectively.

45.2.1.88 1000BASE-KX control register (Register 1.160)

The assignment of bits in the 1000BASE-KX control register is shown in Table 45–68.

Table 45–68—1000BASE-KX control register

Bit(s)	Name	Description	R/W[a]
1.160.15:1	Reserved	Value always 0	RO
1.160.0	PMD transmit disable	1 = Disable transmitter output 0 = Enable transmitter output	R/W

[a]RO = Read only, R/W = Read/Write

45.2.1.88.1 PMD transmit disable (1.160.0)

This bit disables the 1000BASE-KX transmitter as defined in 70.6.5.

45.2.1.89 1000BASE-KX status register (Register 1.161)

The assignment of bits in the 1000BASE-KX status register is shown in Table 45–69.

Table 45–69—1000BASE-KX status register

Bit(s)	Name	Description	R/W[a]
1.161.15:14	Reserved	Value always 0	RO
1.161.13	Transmit fault ability	1 = PMA/PMD has the ability to detect a fault condition on the transmit path 0 = PMA/ PMD does not have the ability to detect a fault condition on the transmit path	RO
1.161.12	Receive fault ability	1 = PMA/PMD has the ability to detect a fault condition on the receive path 0 = PMA/ PMD does not have the ability to detect a fault condition on the receive path	RO
1.161.11	Transmit fault	1 = Fault condition on transmit path 0 = No fault condition on transmit path	LH
1.161.10	Receive fault	1 = Fault condition on receive path 0 = No fault condition on receive path	LH
1.161.9	Reserved	Value always 0	RO
1.161.8	PMD transmit disable ability	1 = PMD has the ability to disable the transmit path 0 = PMD does not have the ability to disable the transmit path	RO
1.161.7:1	Reserved	Value always 0	RO
1.161.0	Signal detect signal from PMD	1 = PMD has asserted signal detect 0 = PMD has not asserted signal detect	RO

[a]RO = Read only, LH = Latching high

45.2.1.89.1 PMD transmit fault ability (1.161.13)

When read as a one, bit 1.161.13 indicates that the PMA/PMD has the ability to detect a fault condition on the transmit path. When read as a zero, bit 1.161.13 indicates that the PMA/PMD does not have the ability to detect a fault condition on the transmit path.

45.2.1.89.2 PMD receive fault ability (1.161.12)

When read as a one, bit 1.161.12 indicates that the PMA/PMD has the ability to detect a fault condition on the receive path. When read as a zero, bit 1.161.12 indicates that the PMA/PMD does not have the ability to detect a fault condition on the receive path.

45.2.1.89.3 PMD transmit fault (1.161.11)

When read as a one, bit 1.161.11 indicates that the PMA/PMD has detected a fault condition on the transmit path. When read as a zero, bit 1.161.11 indicates that the PMA/PMD has not detected a fault condition on the transmit path. Detection of a fault condition on the transmit path is optional and the ability to detect such a condition is advertised by bit 1.161.13. A PMA/PMD that is unable to detect a fault condition on the transmit path shall return a value of zero for this bit. The transmit fault bit shall be implemented with latching high behavior.

The default value of bit 1.161.11 is zero.

45.2.1.89.4 PMD receive fault (1.161.10)

When read as a one, bit 1.161.10 indicates that the PMA/PMD has detected a fault condition on the receive path. When read as a zero, bit 1.161.10 indicates that the PMA/PMD has not detected a fault condition on the receive path. Detection of a fault condition on the receive path is optional and the ability to detect such a condition is advertised by bit 1.161.12. A PMA/PMD that is unable to detect a fault condition on the receive path shall return a value of zero for this bit. The receive fault bit shall be implemented with latching high behavior.

The default value of bit 1.161.10 is zero.

45.2.1.89.5 PMD transmit disable ability (1.161.8)

When read as a one, bit 1.161.8 indicates that the PMD is able to perform the transmit disable function. When read as a zero, bit 1.161.8 indicates that the PMD is not able to perform the transmit disable function.

45.2.1.89.6 1000BASE-KX signal detect (1.161.0)

The PMD signal detect function is optional see 70.6.4. The 1000BASE-X PCS requires signal detect to be one before synchronization can occur. If the signal detect function is not implemented this bit is set to one.

45.2.1.90 PMA overhead control 1, 2, and 3 registers (Register 1.162 through 1.164)

Assignment of bits in the PMA overhead control 1, 2, and 3 registers is shown in Table 45–70. These bits shall be reset to the default values indicated in Table 45–70 upon PHY reset. For the 100GBASE-KP4 PHY the use of these registers is specified in 94.2.2.3 and 94.2.3.1.

Table 45–70—PMA overhead control 1, 2, and 3 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.162.15:13	Reserved	Value always 0	RO
1.162.12:8	PMA transmit overhead sequence 0	Sequence of overhead groups for lane 0 Default = 00110	R/W
1.162.7:0	PMA transmit overhead pattern	Bit pattern for 8-bit transmit overhead group Default = 01100110	R/W
1.163.15	Reserved	Value always 0	RO
1.163.14:10	PMA transmit overhead sequence 3	Sequence of overhead groups for lane 3 Default = 11001	R/W
1.163.9:5	PMA transmit overhead sequence 2	Sequence of overhead groups for lane 2 Default = 10101	R/W
1.163.4:0	PMA transmit overhead sequence 1	Sequence of overhead groups for lane 1 Default = 01010	R/W
1.164.15:8	Reserved	Value always 0	RO
1.164.7:0	PMA receive overhead pattern	Bit pattern for 8-bit receive overhead group Default = 01100110	R/W

[a]R/W = Read/Write, RO = Read only

45.2.1.91 PMA overhead status 1 and 2 registers (Register 1.165, 1.166)

Assignment of bits in the PMA overhead status 1 and 2 registers is shown in Table 45–71. These bits shall be reset to all zeros upon PHY reset. For the 100GBASE-KP4 PHY the use of these registers is specified in 94.2.3.1.

Table 45–71—PMA overhead status 1 and 2 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.165.15:12	Reserved	Value always 0	RO
1.165.11:6	PMA receive overhead sequence 1	Sequence of overhead groups for lane 1	RO
1.165.5:0	PMA receive overhead sequence 0	Sequence of overhead groups for lane 0	RO
1.166.15:12	Reserved	Value always 0	RO
1.166.11:6	PMA receive overhead sequence 3	Sequence of overhead groups for lane 3	RO
1.166.5:0	PMA receive overhead sequence 2	Sequence of overhead groups for lane 2	RO

[a]RO = Read only

45.2.1.92 BASE-R FEC ability register (Register 1.170)

The assignment of bits in the BASE-R FEC ability register is shown in Table 45–72.

Table 45–72—BASE-R FEC ability register bit definitions

Bit(s)	Name	Description	R/W[a]
1.170.15:2	Reserved	Value always 0	RO
1.170.1	BASE-R FEC error indication ability	A read of 1 in this bit indicates that the sublayer is able to report FEC decoding errors to the PCS layer	RO
1.170.0	BASE-R FEC ability	A read of 1 in this bit indicates that the sublayer supports FEC	RO

[a]RO Read only

45.2.1.92.1 BASE-R FEC ability (1.170.0)

When read as a one, this bit indicates that the sublayer supports forward error correction (FEC). When read as a zero, the sublayer does not support forward error correction.

45.2.1.92.2 BASE-R FEC error indication ability (1.170.1)

When read as a one, this bit indicates that the BASE-R FEC sublayer is able to indicate decoding errors to the PCS layer (see 74.8.3). When read as a zero, the BASE-R FEC sublayer is not able to indicate decoding errors to the PCS layer. BASE-R FEC error indication is controlled by the FEC enable error indication bit in the BASE-R FEC control register (see 45.2.1.93.2).

45.2.1.93 BASE-R FEC control register (Register 1.171)

The assignment of bits in the BASE-R FEC control register is shown in Table 45–73.

Table 45–73—BASE-R FEC control register bit definitions

Bit(s)	Name	Description	R/W[a]
1.171.15:2	Reserved	Value always 0	RO
1.171.1	FEC enable error indication	A write of 1 to this bit configures the FEC decoder to indicate errors to the PCS layer	R/W
1.171.0	FEC enable	A write of 1 to this bit enables FEC A write of 0 to this bit disables FEC	R/W

[a]R/W = Read/Write, RO Read only

45.2.1.93.1 FEC enable (1.171.0)

When written as a one, this bit enables FEC for the BASE-R PHY. When written as a zero, FEC is disabled in the BASE-R PHY. This bit shall be set to zero upon execution of PHY reset.

45.2.1.93.2 FEC enable error indication (1.171.1)

This bit enables the BASE-R FEC decoder to indicate decoding errors to the upper layers (PCS) through the sync bits for the BASE-R PHY in the Local Device. When written as a one, this bit enables indication of decoding errors through the sync bits to the PCS layer. When written as zero the error indication function is disabled. Writes to this bit are ignored and reads return a zero if the BASE-R FEC does not have the ability to indicate decoding errors to the PCS layer (see 45.2.1.92.2 and 74.8.3).

45.2.1.94 10GBASE-R FEC corrected blocks counter (Register 1.172, 1.173)

The assignment of bits in the 10GBASE-R FEC corrected blocks counter register is shown in Table 45–74. See 74.8.4.1 for a definition of this register. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.172, 1.173 are used to read the value of a 32-bit counter. When registers 1.172 and 1.173 are used to read the 32-bit counter value, the register 1.172 is read first, the value of the register 1.173 is latched when (and only when) register 1.172 is read and reads of register 1.173 returns the latched value rather than the current value of the counter.

Table 45–74—10GBASE-R FEC corrected blocks counter register bit definitions

Bit(s)	Name	Description	R/W[a]
1.172.15:0	FEC corrected blocks lower	FEC_corrected_blocks_counter[15:0]	RO, NR
1.173.15:0	FEC corrected blocks upper	FEC_corrected_blocks_counter[31:16]	RO, NR

[a]RO = Read only, NR = Non Roll-over

For a multi-PCS lane PHY, this register may be a copy of register 1.300 BASE-R FEC corrected blocks counter, lane 0. If implemented, all accesses to the copy shall have identical behavior as the original register.

45.2.1.95 10GBASE-R FEC uncorrected blocks counter (Register 1.174, 1.175)

The assignment of bits in the 10GBASE-R FEC uncorrected blocks counter register is shown in Table 45–75. See 74.8.4.2 for a definition of this register. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.174, 1.175 are used to read the value of a 32-bit counter. When registers 1.174 and 1.175 are used to read the 32-bit counter value, the register 1.174 is read first, the value of the register 1.175 is latched when (and only when) register 1.174 is read and reads of register 1.175 returns the latched value rather than the current value of the counter.

Table 45–75—10GBASE-R FEC uncorrected blocks counter register bit definitions

Bit(s)	Name	Description	R/W[a]
1.174.15:0	FEC uncorrected blocks lower	FEC_uncorrected_blocks_counter[15:0]	RO, NR
1.175.15:0	FEC uncorrected blocks upper	FEC_uncorrected_blocks_counter[31:16]	RO, NR

[a]RO = Read only, NR = Non Roll-over

For a multi-PCS lane PHY, this register may be a copy of register 1.700 BASE-R FEC uncorrected blocks counter, lane 0. If implemented, all accesses to the copy shall have identical behavior as the original register.

45.2.1.96 CAUI-4 chip-to-module recommended CTLE register (Register 1.179)

The assignment of bits in the CAUI-4 chip-to-module recommended CTLE register is shown in Table 45–76.

Table 45–76—CAUI-4 chip-to-module recommended CTLE register bit definitions

Bit(s)	Name	Description	R/W[a]
1.179.15:5	Reserved	Value always 0	RO
1.179.4:1	Recommended CTLE peaking	4 3 2 1 1 1 x x = Reserved 1 0 1 x = Reserved 1 0 0 1 = 9 dB 1 0 0 0 = 8 dB 0 1 1 1 = 7 dB 0 1 1 0 = 6 dB 0 1 0 1 = 5 dB 0 1 0 0 = 4 dB 0 0 1 1 = 3 dB 0 0 1 0 = 2 dB 0 0 0 1 = 1 dB 0 0 0 0 = Reserved	R/W
1.179.0	Reserved	Value always 0	RO

[a]R/W = Read/Write, RO = Read only

45.2.1.96.1 Recommended CTLE peaking (1.179.4:1)

The value of these bits sets the CTLE peaking value recommended by a host that implements the optional CAUI-4 chip-to-module interface defined in Annex 83E (see 83E.3.1.6). The module may optionally use this information to adjust its CTLE setting.

45.2.1.97 CAUI-4 chip-to-chip transmitter equalization, receive direction, lane 0 register (Register 1.180)

The assignment of bits in the CAUI-4 chip-to-chip transmitter equalization, receive direction, lane 0 register is shown in Table 45–77. The transmitter, receive direction, is the transmitter that sends data towards the PCS.

45.2.1.97.1 Request flag (1.180.15)

The value of this bit indicates the value of the variable *Request_flag* in the lane 0 CAUI-4 receiver in the receive direction (see 83D.3.3.2). This indicates whether the CAUI-4 chip-to-chip device is issuing a request to change the remote transmitter equalization in the CAUI-4 chip-to-chip lane 0 transmitter in the receive direction. If a lane 0 CAUI-4 receiver in the receive direction is not present in the package, then the value returned for this bit should be zero.

Table 45–77—CAUI-4 chip-to-chip transmitter equalization, receive direction, lane 0 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.180.15	Request flag	1 = Change in equalization is requested 0 = No change in equalization is requested	RO
1.180.14:12	Post-cursor request	14 13 12 1 1 1 Reserved 1 1 0 Reserved 1 0 1 *Requested_eq_c1* = 5 (c(1) ratio –0.25) 1 0 0 *Requested_eq_c1* = 4 (c(1) ratio –0.2) 0 1 1 *Requested_eq_c1* = 3 (c(1) ratio –0.15) 0 1 0 *Requested_eq_c1* = 2 (c(1) ratio –0.1) 0 0 1 *Requested_eq_c1* = 1 (c(1) ratio –0.05) 0 0 0 *Requested_eq_c1* = 0 (c(1) ratio 0)	RO
1.180.11:10	Pre-cursor request	11 10 1 1 *Requested_eq_cm1* = 3 (c(–1) ratio –0.15) 1 0 *Requested_eq_cm1* = 2 (c(–1) ratio –0.1) 0 1 *Requested_eq_cm1* = 1 (c(–1) ratio –0.05) 0 0 *Requested_eq_cm1* = 0 (c(–1) ratio 0)	RO
1.180.9:7	Post-cursor remote setting	9 8 7 1 1 1 Reserved 1 1 0 Reserved 1 0 1 *Remote_eq_c1* = 5 (c(1) ratio –0.25) 1 0 0 *Remote_eq_c1* = 4 (c(1) ratio –0.2) 0 1 1 *Remote_eq_c1* = 3 (c(1) ratio –0.15) 0 1 0 *Remote_eq_c1* = 2 (c(1) ratio –0.1) 0 0 1 *Remote_eq_c1* = 1 (c(1) ratio –0.05) 0 0 0 *Remote_eq_c1* = 0 (c(1) ratio 0)	R/W
1.180.6:5	Pre-cursor remote setting	6 5 1 1 *Remote_eq_cm1* = 3 (c(–1) ratio –0.15) 1 0 *Remote_eq_cm1* = 2 (c(–1) ratio –0.1) 0 1 *Remote_eq_cm1* = 1 (c(–1) ratio –0.05) 0 0 *Remote_eq_cm1* = 0 (c(–1) ratio 0)	R/W
1.180.4:2	Post-cursor local setting	4 3 2 1 1 1 Reserved 1 1 0 Reserved 1 0 1 *Local_eq_c1* = 5 (c(1) ratio –0.25) 1 0 0 *Local_eq_c1* = 4 (c(1) ratio –0.2) 0 1 1 *Local_eq_c1* = 3 (c(1) ratio –0.15) 0 1 0 *Local_eq_c1* = 2 (c(1) ratio –0.1) 0 0 1 *Local_eq_c1* = 1 (c(1) ratio –0.05) 0 0 0 *Local_eq_c1* = 0 (c(1) ratio 0)	R/W
1.180.1:0	Pre-cursor local setting	1 0 1 1 *Local_eq_cm1* = 3 (c(–1) ratio –0.15) 1 0 *Local_eq_cm1* = 2 (c(–1) ratio –0.1) 0 1 *Local_eq_cm1* = 1 (c(–1) ratio –0.05) 0 0 *Local_eq_cm1* = 0 (c(–1) ratio 0)	R/W

[a]R/W = Read/Write, RO = Read only

45.2.1.97.2 Post-cursor request (1.180.14:12)

The value of these bits indicates the value of the variable *Requested_eq_c1* in the lane 0 CAUI-4 receiver in the receive direction (see 83D.3.3.2). When *Request_flag* is equal to 1, this value indicates the ratio of the

post-cursor coefficient c(1), which is requested for the transmitter equalization in the CAUI-4 chip-to-chip lane 0 transmitter in the receive direction.

45.2.1.97.3 Pre-cursor request (1.180.11:10)

The value of these bits indicates the value of the variable *Requested_eq_cm1* in the lane 0 CAUI-4 receiver in the receive direction (see 83D.3.3.2). When *Request_flag* is equal to 1, this value indicates the ratio of the pre-cursor coefficient c(–1), which is requested for the transmitter equalization in the CAUI-4 chip-to-chip lane 0 transmitter in the receive direction.

45.2.1.97.4 Post-cursor remote setting (1.180.9:7)

The value of these bits sets the variable *Remote_eq_c1* for the lane 0 CAUI-4 receiver in the receive direction (see 83D.3.3.2). This is used by a CAUI-4 receiver that implements the optional transmitter equalization feedback as an indication of the ratio of the post-cursor coefficient c(1) being used in lane 0 of the CAUI-4 transmitter in the receive direction (see 83D.3.1.1). It may be used to generate values for the request flag and the request bits. If a lane 0 CAUI-4 receiver in the receive direction is not present in the package, then these bits have no effect.

45.2.1.97.5 Pre-cursor remote setting (1.180.6:5)

The value of these bits sets the variable *Remote_eq_cm1* for the lane 0 CAUI-4 receiver in the receive direction (see 83D.3.3.2). This is used by a CAUI-4 receiver that implements the optional transmitter equalization feedback as an indication of the ratio of the pre-cursor coefficient c(–1) being used in lane 0 of the CAUI-4 transmitter in the receive direction (see 83D.3.1.1). It may be used to generate values for the request flag and the request bits. If a lane 0 CAUI-4 receiver in the receive direction is not present in the package, then these bits have no effect.

45.2.1.97.6 Post-cursor local setting (1.180.4:2)

The value of these bits sets the variable *Local_eq_c1* for the lane 0 CAUI-4 transmitter in the receive direction (see 83D.3.1.1 and Table 83D–3), which controls the weight of the transmitter equalization post-cursor coefficient c(1). If a lane 0 CAUI-4 transmitter in the receive direction is not present in the package, then these bits have no effect.

45.2.1.97.7 Pre-cursor local setting (1.180.1:0)

The value of these bits sets the variable *Local_eq_cm1* for the lane 0 CAUI-4 transmitter in the receive direction (see 83D.3.1.1 and Table 83D–2), which controls the weight of the transmitter equalization pre-cursor coefficient c(–1). If a lane 0 CAUI-4 transmitter in the receive direction is not present in the package, then these bits have no effect.

45.2.1.98 CAUI-4 chip-to-chip transmitter equalization, receive direction, lane 1 through lane 3 registers (Registers 1.181, 1.182, 1.183)

The CAUI-4 chip-to-chip transmitter equalization, receive direction, lane 1 through lane 3 registers are defined similarly to register 1.180 (which is used for lane 0, see 45.2.1.97) but for lanes 1 through 3, respectively. The transmitter, receive direction, is the transmitter that sends data towards the PCS.

45.2.1.99 CAUI-4 chip-to-chip transmitter equalization, transmit direction, lane 0 register (Register 1.184)

The assignment of bits in the CAUI-4 chip-to-chip transmitter equalization, transmit direction, lane 0 register is shown in Table 45–78. The transmitter, transmit direction, is the transmitter that sends data towards the PMD.

Table 45–78—CAUI-4 chip-to-chip transmitter equalization, transmit direction, lane 0 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.184.15	Request flag	1 = Change in equalization is requested 0 = No change in equalization is requested	RO
1.184.14:12	Post-cursor request	14 13 12 1 1 1 Reserved 1 1 0 Reserved 1 0 1 *Requested_eq_c1* = 5 (c(1) ratio –0.25) 1 0 0 *Requested_eq_c1* = 4 (c(1) ratio –0.2) 0 1 1 *Requested_eq_c1* = 3 (c(1) ratio –0.15) 0 1 0 *Requested_eq_c1* = 2 (c(1) ratio –0.1) 0 0 1 *Requested_eq_c1* = 1 (c(1) ratio –0.05) 0 0 0 *Requested_eq_c1* = 0 (c(1) ratio 0)	RO
1.184.11:10	Pre-cursor request	11 10 1 1 *Requested_eq_cm1* = 3 (c(–1) ratio –0.15) 1 0 *Requested_eq_cm1* = 2 (c(–1) ratio –0.1) 0 1 *Requested_eq_cm1* = 1 (c(–1) ratio –0.05) 0 0 *Requested_eq_cm1* = 0 (c(–1) ratio 0)	RO
1.184.9:7	Post-cursor remote setting	9 8 7 1 1 1 Reserved 1 1 0 Reserved 1 0 1 *Remote_eq_c1* = 5 (c(1) ratio –0.25) 1 0 0 *Remote_eq_c1* = 4 (c(1) ratio –0.2) 0 1 1 *Remote_eq_c1* = 3 (c(1) ratio –0.15) 0 1 0 *Remote_eq_c1* = 2 (c(1) ratio –0.1) 0 0 1 *Remote_eq_c1* = 1 (c(1) ratio –0.05) 0 0 0 *Remote_eq_c1* = 0 (c(1) ratio 0)	R/W
1.184.6:5	Pre-cursor remote setting	6 5 1 1 *Remote_eq_cm1* = 3 (c(–1) ratio –0.15) 1 0 *Remote_eq_cm1* = 2 (c(–1) ratio –0.1) 0 1 *Remote_eq_cm1* = 1 (c(–1) ratio –0.05) 0 0 *Remote_eq_cm1* = 0 (c(–1) ratio 0)	R/W
1.184.4:2	Post-cursor local setting	4 3 2 1 1 1 Reserved 1 1 0 Reserved 1 0 1 *Local_eq_c1* = 5 (c(1) ratio –0.25) 1 0 0 *Local_eq_c1* = 4 (c(1) ratio –0.2) 0 1 1 *Local_eq_c1* = 3 (c(1) ratio –0.15) 0 1 0 *Local_eq_c1* = 2 (c(1) ratio –0.1) 0 0 1 *Local_eq_c1* = 1 (c(1) ratio–0.05) 0 0 0 *Local_eq_c1* = 0 (c(1) ratio 0)	R/W
1.184.1:0	Pre-cursor local setting	1 0 1 1 *Local_eq_cm1* = 3 (c(–1) ratio –0.15) 1 0 *Local_eq_cm1* = 2 (c(–1) ratio –0.1) 0 1 *Local_eq_cm1* = 1 (c(–1) ratio –0.05) 0 0 *Local_eq_cm1* = 0 (c(–1) ratio 0)	R/W

[a]R/W = Read/Write, RO = Read only

45.2.1.99.1 Request flag (1.184.15)

The value of this bit indicates the value of the variable *Request_flag* in the lane 0 CAUI-4 receiver in the transmit direction (see 83D.3.3.2). This indicates whether the CAUI-4 chip-to-chip device is issuing a request to change the remote transmitter equalization in the CAUI-4 chip-to-chip lane 0 transmitter in the transmit direction. If a lane 0 CAUI-4 receiver in the transmit direction is not present in the package, then the value returned for this bit should be zero.

45.2.1.99.2 Post-cursor request (1.184.14:12)

The value of these bits indicates the value of the variable *Requested_eq_c1* in the lane 0 CAUI-4 receiver in the transmit direction (see 83D.3.3.2). When *Request_flag* is equal to 1, this value indicates the ratio of the post-cursor coefficient c(1), which is requested for the transmitter equalization in the CAUI-4 chip-to-chip lane 0 transmitter in the transmit direction.

45.2.1.99.3 Pre-cursor request (1.184.11:10)

The value of these bits indicates the value of the variable *Requested_eq_cm1* in the lane 0 CAUI-4 receiver in the transmit direction (see 83D.3.3.2). When *Request_flag* is equal to 1, this value indicates the ratio of the pre-cursor coefficient c(–1), which is requested for the transmitter equalization in the CAUI-4 chip-to-chip lane 0 transmitter in the transmit direction.

45.2.1.99.4 Post-cursor remote setting (1.184.9:7)

The value of these bits sets the variable *Remote_eq_c1* for the lane 0 CAUI-4 receiver in the transmit direction (see 83D.3.3.2). This is used by a CAUI-4 receiver that implements the optional transmitter equalization feedback as an indication of the ratio of the post-cursor coefficient c(1) being used in lane 0 of the CAUI-4 transmitter in the transmit direction (see 83D.3.1.1). It may be used to generate values for the request flag and the request bits. If a lane 0 CAUI-4 receiver in the transmit direction is not present in the package, then these bits have no effect.

45.2.1.99.5 Pre-cursor remote setting (1.184.6:5)

The value of these bits sets the variable *Remote_eq_cm1* for the lane 0 CAUI-4 receiver in the transmit direction (see 83D.3.3.2). This is used by a CAUI-4 receiver that implements the optional transmitter equalization feedback as an indication of the ratio of the pre-cursor coefficient c(–1) being used in lane 0 of the CAUI-4 transmitter in the transmit direction (see 83D.3.1.1). It may be used to generate values for the request flag and the request bits. If a lane 0 CAUI-4 receiver in the transmit direction is not present in the package, then these bits have no effect.

45.2.1.99.6 Post-cursor local setting (1.184.4:2)

The value of these bits sets the variable *Local_eq_c1* for the lane 0 CAUI-4 transmitter in the transmit direction (see 83D.3.1.1 and Table 83D–3), which controls the weight of the transmitter equalization post-cursor coefficient c(1). If a lane 0 CAUI-4 transmitter in the transmit direction is not present in the package, then these bits have no effect.

45.2.1.99.7 Pre-cursor local setting (1.184.1:0)

The value of these bits sets the variable *Local_eq_cm1* for the lane 0 CAUI-4 transmitter in the transmit direction (see 83D.3.1.1 and Table 83D–2), which controls the weight of the transmitter equalization pre-cursor coefficient c(–1). If a lane 0 CAUI-4 transmitter in the transmit direction is not present in the package, then these bits have no effect.

45.2.1.100 CAUI-4 chip-to-chip transmitter equalization, transmit direction, lane 1 through lane 3 registers (Registers 1.185, 1.186, 1.187)

The CAUI-4 chip-to-chip transmitter equalization, transmit direction, lane 1 through lane 3 registers are defined similarly to register 1.184 (which is used for lane 0, see 45.2.1.99) but for lanes 1 through 3, respectively. The transmitter, transmit direction, is the transmitter that sends data towards the PMD.

45.2.1.101 RS-FEC control register (Register 1.200)

The assignment of bits in the RS-FEC control register is shown in Table 45–79.

Table 45–79—RS-FEC control register bit definitions

Bit(s)	Name	Description	R/W[a]
1.200.15:2	Reserved	Value always 0	RO
1.200.1	FEC bypass indication enable	1 = FEC decoder does not indicate errors to the PCS 0 = FEC decoder indicates errors to the PCS layer	R/W
1.200.0	FEC bypass correction enable	1 = FEC decoder performs error detection without error correction 0 = FEC decoder performs error detection and error correction	R/W

[a]R/W = Read/Write, RO = Read only

45.2.1.101.1 FEC bypass indication enable (1.200.1)

This bit enables the RS-FEC decoder to bypass error indication to the upper layers (PCS) through the sync bits for the BASE-R PHY in the Local Device. When set to a one, this bit enables bypass of the error indication. When set to a zero, errors are indicated to the PCS through the sync bits. Writes to this bit are ignored and reads return a zero if the RS-FEC does not have the ability to bypass indicating decoding errors to the PCS layer (see 91.5.3.3).

45.2.1.101.2 FEC bypass correction enable (1.200.0)

When this bit is set to one the Reed-Solomon decoder performs error detection without error correction (see 91.5.3.3). When this bit is set to zero, the decoder also performs error correction. Writes to this bit are ignored and reads return a zero if the RS-FEC does not have the ability to bypass correction.

45.2.1.102 RS-FEC status register (Register 1.201)

The assignment of bits in the RS-FEC status register is shown in Table 45–80.

45.2.1.102.1 PCS align status (1.201.15)

When read as a one, bit 1.201.15 indicates that the RS-FEC described in Clause 91 has locked and aligned all transmit PCS lanes. When read as a zero, bit 1.201.15 indicates that the RS-FEC has not locked and aligned all transmit PCS lanes. A device that implements the RS-FEC status register but does not implement a separated RS-FEC shall return a one for bit 1.201.15.

Table 45–80—RS-FEC status register bit definitions

Bit(s)	Name	Description	R/W[a]
1.201.15	PCS align status	1 = FEC encoder has locked and aligned all PCS lanes 0 = FEC encoder has not locked and aligned all PCS lanes	RO
1.201.14	RS-FEC align status	1 = RS-FEC receive lanes locked and aligned 0 = RS-FEC receive lanes not locked and aligned	RO
1.201.13:12	Reserved	Value always 0	RO
1.201.11	FEC AM lock 3	1 = RS-FEC receive lane 3 locked and aligned 0 = RS-FEC receive lane 3 not locked and aligned	RO
1.201.10	FEC AM lock 2	1 = RS-FEC receive lane 2 locked and aligned 0 = RS-FEC receive lane 2 not locked and aligned	RO
1.201.9	FEC AM lock 1	1 = RS-FEC receive lane 1 locked and aligned 0 = RS-FEC receive lane 1 not locked and aligned	RO
1.201.8	FEC AM lock 0	1 = RS-FEC receive lane 0 locked and aligned 0 = RS-FEC receive lane 0 not locked and aligned	RO
1.201.7:3	Reserved	Value always 0	RO
1.201.2	RS-FEC high SER	1 = FEC errors have exceeded threshold 0 = FEC errors have not exceeded threshold	RO/LH
1.201.1	FEC bypass indication ability	1 = FEC decoder has the ability to bypass error indication 0 = FEC decoder does not have the ability to bypass error indication	RO
1.201.0	FEC bypass correction ability	1 = FEC decoder has the ability to bypass error correction 0 = FEC decoder does not have the ability to bypass error correction	RO

[a]RO = Read only, LH = Latching high

45.2.1.102.2 RS-FEC align status (1.201.14)

When read as a one, bit 1.201.14 indicates that the RS-FEC described in Clause 91 has locked and aligned all receive RS-FEC lanes. When read as a zero, bit 1.201.14 indicates that the RS-FEC has not locked and aligned all receive RS-FEC lanes.

45.2.1.102.3 FEC AM lock 3 (1.201.11)

When read as a one, bit 1.201.11 indicates that the RS-FEC described in Clause 91 has locked and aligned lane 3 of the PMA service interface. When read as a zero, bit 1.201.11 indicates that the RS-FEC has not locked and aligned lane 3 of the PMA service interface. This bit reflects the state of amps_lock[3] (see 91.5.3.1).

45.2.1.102.4 FEC AM lock 2 (1.201.10)

When read as a one, bit 1.201.10 indicates that the RS-FEC described in Clause 91 has locked and aligned lane 2 of the PMA service interface. When read as a zero, bit 1.201.10 indicates that the RS-FEC has not locked and aligned lane 2 of the PMA service interface. This bit reflects the state of amps_lock[2] (see 91.5.3.1).

45.2.1.102.5 FEC AM lock 1 (1.201.9)

When read as a one, bit 1.201.9 indicates that the RS-FEC described in Clause 91 has locked and aligned lane 1 of the PMA service interface. When read as a zero, bit 1.201.9 indicates that the RS-FEC has not locked and aligned lane 1 of the PMA service interface. This bit reflects the state of amps_lock[1] (see 91.5.3.1).

45.2.1.102.6 FEC AM lock 0 (1.201.8)

When read as a one, bit 1.201.8 indicates that the RS-FEC described in Clause 91 has locked and aligned lane 0 of the PMA service interface. When read as a zero, bit 1.201.8 indicates that the RS-FEC has not locked and aligned lane 0 of the PMA service interface. This bit reflects the state of amps_lock[0] (see 91.5.3.1).

45.2.1.102.7 RS-FEC high SER (1.201.2)

When FEC_bypass_indication_enable is set to one, this bit is set to one if the number of RS-FEC symbol errors in a window of 8192 codewords exceeds the threshold (see 91.5.3.3) and is set to zero otherwise. The bit is set to zero if FEC_bypass_indication_enable is set to zero. This bit shall be implemented with latching high behavior.

45.2.1.102.8 FEC bypass indication ability (1.201.1)

The Reed-Solomon decoder may have the option to perform error detection without error indication (see 91.5.3.3) to reduce the delay contributed by the RS-FEC sublayer. This bit is set to one to indicate that the decoder has this ability to bypass error indication. The bit is set to zero if this ability is not supported.

45.2.1.102.9 FEC bypass correction ability (1.201.0)

The Reed-Solomon decoder may have the option to perform error detection without error correction (see 91.5.3.3) to reduce the delay contributed by the RS-FEC sublayer. This bit is set to one to indicate that the decoder has this ability to bypass error correction. The bit is set to zero if this ability is not supported.

45.2.1.103 RS-FEC corrected codewords counter (Register 1.202, 1.203)

The assignment of bits in the RS-FEC corrected codewords counter register is shown in Table 45–81. See 91.6.8 for a definition of this register. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.202, 1.203 are used to read the value of a 32-bit counter. When registers 1.202 and 1.203 are used to read the 32-bit counter value, the register 1.202 is read first, the value of the register 1.203 is latched when (and only when) register 1.202 is read, and reads of register 1.203 return the latched value rather than the current value of the counter.

Table 45–81—RS-FEC corrected codewords counter register bit definitions

Bit(s)	Name	Description	R/W[a]
1.202.15:0	FEC corrected codewords lower	FEC_corrected_cw_counter[15:0]	RO, NR
1.203.15:0	FEC corrected codewords upper	FEC_corrected_cw_counter[31:16]	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.1.104 RS-FEC uncorrected codewords counter (Register 1.204, 1.205)

The assignment of bits in the RS-FEC uncorrected codewords counter register is shown in Table 45–82. See 91.6.9 for a definition of this register. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.204, 1.205 are used to read the value of a 32-bit counter. When registers 1.204 and 1.205 are used to read the 32-bit counter value, the register 1.204 is read first, the value of the register 1.205 is latched when (and only when) register 1.204 is read, and reads of register 1.205 return the latched value rather than the current value of the counter.

Table 45–82—RS-FEC uncorrected codewords counter register bit definitions

Bit(s)	Name	Description	R/W[a]
1.204.15:0	FEC uncorrected codewords lower	FEC_uncorrected_cw_counter[15:0]	RO, NR
1.205.15:0	FEC uncorrected codewords upper	FEC_uncorrected_cw_counter[31:16]	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.1.105 RS-FEC lane mapping register (Register 1.206)

The assignment of bits in the RS-FEC lane mapping register is shown in Table 45–83. When the RS-FEC detects and locks the RS-FEC for PMA service interface lane 0, the detected RS-FEC lane number is recorded bits 1:0 in this register. Similarly, the detected RS-FEC lane numbers for PMA service lanes 1, 2, and 3 are recorded in register bits 3:2, 5:4, and 7:6, respectively. The contents of the RS-FEC lane mapping register bits 7:0 are valid when RS-FEC align status (1.201.14) is set to one and are invalid otherwise.

Table 45–83—RS-FEC lane mapping register

Bit(s)	Name	Description	R/W[a]
1.206.15:8	Reserved	Value always 0	RO
1.206.7:6	RS-FEC lane 3 mapping	RS-FEC lane mapped to PMA lane 3	RO
1.206.5:4	RS-FEC lane 2 mapping	RS-FEC lane mapped to PMA lane 2	RO
1.206.3:2	RS-FEC lane 1 mapping	RS-FEC lane mapped to PMA lane 1	RO
1.206.1:0	RS-FEC lane 0 mapping	RS-FEC lane mapped to PMA lane 0	RO

[a]RO = Read only

45.2.1.106 RS-FEC symbol error counter lane 0 (Register 1.210, 1.211)

The assignment of bits in the RS-FEC symbol error counter lane 0 register is shown in Table 45–84. Symbol errors detected in FEC lane 0 are counted and shown in register 1.210.15:0 and 1.211.15:0. See 91.6.11 for a definition of this register. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.210, 1.211 are used to read the value of a 32-bit counter. When registers 1.210 and 1.211 are used to read the 32-bit counter value, the register 1.210 is read first, the value of the register 1.211 is latched when (and only

when) register 1.210 is read, and reads of register 1.211 return the latched value rather than the current value of the counter.

Table 45–84—RS-FEC symbol error counter register bit definitions

Bit(s)	Name	Description	R/W[a]
1.210.15:0	FEC symbol errors, lane 0 lower	FEC_symbol_error_counter_0[15:0]	RO, NR
1.211.15:0	FEC symbol errors, lane 0 upper	FEC_symbol_error_counter_0[31:16]	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.1.107 RS-FEC symbol error counter lane 1 through 3 (Register 1.212, 1.213, 1.214, 1.215, 1.216, 1.217)

The behavior of the RS-FEC symbol error counters, lane 1 through 3 is identical to that described for FEC lane 0 in 45.2.1.106. Errors detected in each FEC lane are counted and shown in the corresponding register. FEC lane 1, lower 16 bits are shown in register 1.212; FEC lane 1, upper 16 bits are shown in register 1.213; FEC lane 2, lower 16 bits are shown in register 1.214; through register 1.217 for FEC lane 3, upper 16 bits.

45.2.1.108 RS-FEC BIP error counter lane 0 (Register 1.230)

The assignment of bits in the RS-FEC BIP error counter lane 0 is shown in Table 45–159. The RS-FEC described in Clause 91 calculates a BIP value for each PCS lane (see 91.5.2.4, 91.6.3). Errors detected in PCS lane 0 are counted and shown in register 1.230.15:0. The 16-bit counter shall be reset to all zeros when register 1.230 is read or upon PMA/PMD reset. The 16-bit counter shall be held at all ones in the case of overflow. A device that does not implement a separated RS-FEC shall return a zero for all bits in the RS-FEC BIP error counter, lane 0 register.

Table 45–85—RS-FEC BIP error counter lane 0 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.230.15:0	BIP error counter, lane 0	Errors detected by BIP in PCS lane 0	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.1.109 RS-FEC BIP error counter, lane 1 through 19 (Registers 1.231 through 1.249)

The behavior of the RS-FEC BIP error counters, lane 1 through 19 is identical to that described for lane 0 in 45.2.1.108. Errors detected in each PCS lane are counted and shown in register bits 15:0 in the corresponding register. PCS lane 1 is shown in register 1.231; PCS lane 2 is shown in register 1.232; through register 1.249 for PCS lane 19.

45.2.1.110 RS-FEC PCS lane 0 mapping register (Register 1.250)

The assignment of bits in the RS-FEC PCS lane 0 mapping register is shown in Table 45–86. When the RS-FEC instance of the multi-lane PCS described in Clause 82 detects and locks the alignment marker for service interface lane 0, the detected PCS lane number is recorded in this register. The contents of the Lane 0 mapping register is valid when the transmit PCS align status bit (register 1.201.15) is set to one and is

invalid otherwise (see 45.2.1.102). A device that does not implement a separated RS-FEC shall return a zero for all bits in the RS-FEC PCS lane 0 mapping register.

Table 45–86—RS-FEC PCS lane 0 mapping register bit definitions

Bit(s)	Name	Description	R/W[a]
1.250.15:5	Reserved	Value always 0	RO
1.250.4:0	Lane 0 mapping	PCS lane received in service interface lane 0	RO

[a]RO = Read only

45.2.1.111 RS-FEC PCS lanes 1 through 19 mapping registers (Registers 1.251 through 1.269)

The definition of Lanes 1 through 19 mapping registers is identical to that described for lane 0 in 45.2.1.110. The lane mapping for lane 1 is in register 1.251; lane 2 is in register 1.252; etc.

45.2.1.112 RS-FEC PCS alignment status 1 register (Register 1.280)

The assignment of bits in the RS-FEC PCS alignment status 1 register is shown in Table 45–136. All the bits in the RS-FEC PCS alignment status 1 register are read only; a write to the RS-FEC PCS alignment status 1 register shall have no effect. A device that does not implement a separated RS-FEC shall return a zero for all bits in the RS-FEC PCS alignment status 1 register. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

Table 45–87—RS-FEC PCS alignment status 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.280.15:8	Reserved	Value always 0	RO
1.280.7	Block 7 lock	1 = Lane 7 is locked 0 = Lane 7 is not locked	RO
1.280.6	Block 6 lock	1 = Lane 6 is locked 0 = Lane 6 is not locked	RO
1.280.5	Block 5 lock	1 = Lane 5 is locked 0 = Lane 5 is not locked	RO
1.280.4	Block 4 lock	1 = Lane 4 is locked 0 = Lane 4 is not locked	RO
1.280.3	Block 3 lock	1 = Lane 3 is locked 0 = Lane 3 is not locked	RO

Table 45–87—RS-FEC PCS alignment status 1 register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
1.280.2	Block 2 lock	1 = Lane 2 is locked 0 = Lane 2 is not locked	RO
1.280.1	Block 1 lock	1 = Lane 1 is locked 0 = Lane 1 is not locked	RO
1.280.0	Block 0 lock	1 = Lane 0 is locked 0 = Lane 0 is not locked	RO

[a]RO = Read only

45.2.1.112.1 Block 7 lock (1.280.7)

When read as a one, bit 1.280.7 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 7. When read as a zero, bit 1.280.7 indicates that the RS-FEC transmit function lane 7 has not achieved block lock. This bit reflects the state of block_lock[7] (see 91.5.2.1).

45.2.1.112.2 Block 6 lock (1.280.6)

When read as a one, bit 1.280.6 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 6. When read as a zero, bit 1.280.6 indicates that the RS-FEC transmit function lane 6 has not achieved block lock. This bit reflects the state of block_lock[6] (see 91.5.2.1).

45.2.1.112.3 Block 5 lock (1.280.5)

When read as a one, bit 1.280.5 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 5. When read as a zero, bit 1.280.5 indicates that the RS-FEC transmit function lane 5 has not achieved block lock. This bit reflects the state of block_lock[5] (see 91.5.2.1).

45.2.1.112.4 Block 4 lock (1.280.4)

When read as a one, bit 1.280.4 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 4. When read as a zero, bit 1.280.4 indicates that the RS-FEC transmit function lane 4 has not achieved block lock. This bit reflects the state of block_lock[4] (see 91.5.2.1).

45.2.1.112.5 Block 3 lock (1.280.3)

When read as a one, bit 1.280.3 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 3. When read as a zero, bit 1.280.3 indicates that the RS-FEC transmit function lane 3 has not achieved block lock. This bit reflects the state of block_lock[3] (see 91.5.2.1).

45.2.1.112.6 Block 2 lock (1.280.2)

When read as a one, bit 1.280.2 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 2. When read as a zero, bit 1.280.2 indicates that the RS-FEC transmit function lane 2 has not achieved block lock. This bit reflects the state of block_lock[2] (see 91.5.2.1).

45.2.1.112.7 Block 1 lock (1.280.1)

When read as a one, bit 1.280.1 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 1. When read as a zero, bit 1.280.1 indicates that the RS-FEC transmit function lane 1 has not achieved block lock. This bit reflects the state of block_lock[1] (see 91.5.2.1).

45.2.1.112.8 Block 0 lock (1.280.0)

When read as a one, bit 1.280.0 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 0. When read as a zero, bit 1.280.0 indicates that the RS-FEC transmit function lane 0 has not achieved block lock. This bit reflects the state of block_lock[0] (see 91.5.2.1).

45.2.1.113 RS-FEC PCS alignment status 2 register (Register 1.281)

The assignment of bits in the RS-FEC PCS alignment status 2 register is shown in Table 45–137. All the bits in the RS-FEC PCS alignment status 2 register are read only; a write to the RS-FEC PCS alignment status 2 register shall have no effect. A device that does not implement a separated RS-FEC shall return a zero for all bits in the RS-FEC PCS alignment status 2 register. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

Table 45–88—RS-FEC PCS alignment status 2 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.281.15:12	Reserved	Value always 0	RO
1.281.11	Block 19 lock	1 = Lane 19 is locked 0 = Lane 19 is not locked	RO
1.281.10	Block 18 lock	1 = Lane 18 is locked 0 = Lane 18 is not locked	RO
1.281.9	Block 17 lock	1 = Lane 17 is locked 0 = Lane 17 is not locked	RO
1.281.8	Block 16 lock	1 = Lane 16 is locked 0 = Lane 16 is not locked	RO
1.281.7	Block 15 lock	1 = Lane 15 is locked 0 = Lane 15 is not locked	RO
1.281.6	Block 14 lock	1 = Lane 14 is locked 0 = Lane 14 is not locked	RO
1.281.5	Block 13 lock	1 = Lane 13 is locked 0 = Lane 13 is not locked	RO
1.281.4	Block 12 lock	1 = Lane 12 is locked 0 = Lane 12 is not locked	RO
1.281.3	Block 11 lock	1 = Lane 11 is locked 0 = Lane 11 is not locked	RO

Table 45–88—RS-FEC PCS alignment status 2 register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
1.281.2	Block 10 lock	1 = Lane 10 is locked 0 = Lane 10 is not locked	RO
1.281.1	Block 9 lock	1 = Lane 9 is locked 0 = Lane 9 is not locked	RO
1.281.0	Block 8 lock	1 = Lane 8 is locked 0 = Lane 8 is not locked	RO

[a]RO = Read only

45.2.1.113.1 Block 19 lock (1.281.11)

When read as a one, bit 1.281.11 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 19. When read as a zero, bit 1.281.11 indicates that the RS-FEC transmit function lane 19 has not achieved block lock. This bit reflects the state of block_lock[19] (see 91.5.2.1).

45.2.1.113.2 Block 18 lock (1.281.10)

When read as a one, bit 1.281.10 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 18. When read as a zero, bit 1.281.10 indicates that the RS-FEC transmit function lane 18 has not achieved block lock. This bit reflects the state of block_lock[18] (see 91.5.2.1).

45.2.1.113.3 Block 17 lock (1.281.9)

When read as a one, bit 1.281.9 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 17. When read as a zero, bit 1.281.9 indicates that the RS-FEC transmit function lane 17 has not achieved block lock. This bit reflects the state of block_lock[17] (see 91.5.2.1).

45.2.1.113.4 Block 16 lock (1.281.8)

When read as a one, bit 1.281.8 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 16. When read as a zero, bit 1.281.8 indicates that the RS-FEC transmit function lane 16 has not achieved block lock. This bit reflects the state of block_lock[16] (see 91.5.2.1).

45.2.1.113.5 Block 15 lock (1.281.7)

When read as a one, bit 1.281.7 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 15. When read as a zero, bit 1.281.7 indicates that the RS-FEC transmit function lane 15 has not achieved block lock. This bit reflects the state of block_lock[15] (see 91.5.2.1).

45.2.1.113.6 Block 14 lock (1.281.6)

When read as a one, bit 1.281.6 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 14. When read as a zero, bit 1.281.6 indicates that the RS-FEC transmit function lane 14 has not achieved block lock. This bit reflects the state of block_lock[14] (see 91.5.2.1).

45.2.1.113.7 Block 13 lock (1.281.5)

When read as a one, bit 1.281.5 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 13. When read as a zero, bit 1.281.5 indicates that the RS-FEC transmit function lane 13 has not achieved block lock. This bit reflects the state of block_lock[13] (see 91.5.2.1).

45.2.1.113.8 Block 12 lock (1.281.4)

When read as a one, bit 1.281.4 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 12. When read as a zero, bit 1.281.4 indicates that the RS-FEC transmit function lane 12 has not achieved block lock. This bit reflects the state of block_lock[12] (see 91.5.2.1).

45.2.1.113.9 Block 11 lock (1.281.3)

When read as a one, bit 1.281.3 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 11. When read as a zero, bit 1.281.3 indicates that the RS-FEC transmit function lane 11 has not achieved block lock. This bit reflects the state of block_lock[11] (see 91.5.2.1).

45.2.1.113.10 Block 10 lock (1.281.2)

When read as a one, bit 1.281.2 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 10. When read as a zero, bit 1.281.2 indicates that the RS-FEC transmit function lane 10 has not achieved block lock. This bit reflects the state of block_lock[10] (see 91.5.2.1).

45.2.1.113.11 Block 9 lock (1.281.1)

When read as a one, bit 1.281.1 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 9. When read as a zero, bit 1.281.1 indicates that the RS-FEC transmit function lane 9 has not achieved block lock. This bit reflects the state of block_lock[9] (see 91.5.2.1).

45.2.1.113.12 Block 8 lock (1.281.0)

When read as a one, bit 1.281.0 indicates that the RS-FEC transmit function has achieved block lock for service interface lane 8. When read as a zero, bit 1.281.0 indicates that the RS-FEC transmit function lane 8 has not achieved block lock. This bit reflects the state of block_lock[8] (see 91.5.2.1).

45.2.1.114 RS-FEC PCS alignment status 3 register (Register 1.282)

The assignment of bits in the RS-FEC PCS alignment status 3 register is shown in Table 45–89. All the bits in the RS-FEC PCS alignment status 3 register are read only; a write to the RS-FEC PCS alignment status 3 register shall have no effect. A device that does not implement a separated RS-FEC shall return a zero for all bits in the RS-FEC PCS alignment status 3 register. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

45.2.1.114.1 Lane 7 aligned (1.282.7)

When read as a one, bit 1.282.7 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 7. When read as a zero, bit 1.282.7 indicates that the RS-FEC transmit function lane 7 has not achieved alignment marker lock. This bit reflects the state of am_lock[7] (see 91.5.2.2).

Table 45–89—RS-FEC PCS alignment status 3 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.282.15:8	Reserved	Value always 0	RO
1.282.7	Lane 7 aligned	1 = Lane 7 alignment marker is locked 0 = Lane 7 alignment marker is not locked	RO
1.282.6	Lane 6 aligned	1 = Lane 6 alignment marker is locked 0 = Lane 6 alignment marker is not locked	RO
1.282.5	Lane 5 aligned	1 = Lane 5 alignment marker is locked 0 = Lane 5 alignment marker is not locked	RO
1.282.4	Lane 4 aligned	1 = Lane 4 alignment marker is locked 0 = Lane 4 alignment marker is not locked	RO
1.282.3	Lane 3 aligned	1 = Lane 3 alignment marker is locked 0 = Lane 3 alignment marker is not locked	RO
1.282.2	Lane 2 aligned	1 = Lane 2 alignment marker is locked 0 = Lane 2 alignment marker is not locked	RO
1.282.1	Lane 1 aligned	1 = Lane 1 alignment marker is locked 0 = Lane 1 alignment marker is not locked	RO
1.282.0	Lane 0 aligned	1 = Lane 0 alignment marker is locked 0 = Lane 0 alignment marker is not locked	RO

[a]RO = Read only

45.2.1.114.2 Lane 6 aligned (1.282.6)

When read as a one, bit 1.282.6 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 6. When read as a zero, bit 1.282.6 indicates that the RS-FEC transmit function lane 6 has not achieved alignment marker lock. This bit reflects the state of am_lock[6] (see 91.5.2.2).

45.2.1.114.3 Lane 5 aligned (1.282.5)

When read as a one, bit 1.282.5 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 5. When read as a zero, bit 1.282.5 indicates that the RS-FEC transmit function lane 5 has not achieved alignment marker lock. This bit reflects the state of am_lock[5] (see 91.5.2.2).

45.2.1.114.4 Lane 4 aligned (1.282.4)

When read as a one, bit 1.282.4 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 4. When read as a zero, bit 1.282.4 indicates that the RS-FEC transmit function lane 4 has not achieved alignment marker lock. This bit reflects the state of am_lock[4] (see 91.5.2.2).

45.2.1.114.5 Lane 3 aligned (1.282.3)

When read as a one, bit 1.282.3 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 3. When read as a zero, bit 1.282.3 indicates that the RS-FEC transmit function lane 3 has not achieved alignment marker lock. This bit reflects the state of am_lock[3] (see 91.5.2.2).

45.2.1.114.6 Lane 2 aligned (1.282.2)

When read as a one, bit 1.282.2 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 2. When read as a zero, bit 1.282.2 indicates that the RS-FEC transmit function lane 2 has not achieved alignment marker lock. This bit reflects the state of am_lock[2] (see 91.5.2.2).

45.2.1.114.7 Lane 1 aligned (1.282.1)

When read as a one, bit 1.282.1 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 1. When read as a zero, bit 1.282.1 indicates that the RS-FEC transmit function lane 1 has not achieved alignment marker lock. This bit reflects the state of am_lock[1] (see 91.5.2.2).

45.2.1.114.8 Lane 0 aligned (1.282.0)

When read as a one, bit 1.282.0 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 0. When read as a zero, bit 1.282.0 indicates that the RS-FEC transmit function lane 0 has not achieved alignment marker lock. This bit reflects the state of am_lock[0] (see 91.5.2.2).

45.2.1.115 RS-FEC PCS alignment status 4 register (Register 1.283)

The assignment of bits in the RS-FEC PCS alignment status 4 register is shown in Table 45–90. All the bits in the RS-FEC PCS alignment status 4 register are read only; a write to the RS-FEC PCS alignment status 4 register shall have no effect. A device that does not implement a separated RS-FEC shall return a zero for all bits in the RS-FEC PCS alignment status 4 register. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

Table 45–90—RS-FEC PCS alignment status 4 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.283.15:12	Reserved	Value always 0	RO
1.283.11	Lane 19 aligned	1 = Lane 19 alignment marker is locked 0 = Lane 19 alignment marker is not locked	RO
1.283.10	Lane 18 aligned	1 = Lane 18 alignment marker is locked 0 = Lane 18 alignment marker is not locked	RO
1.283.9	Lane 17 aligned	1 = Lane 17 alignment marker is locked 0 = Lane 17 alignment marker is not locked	RO
1.283.8	Lane 16 aligned	1 = Lane 16 alignment marker is locked 0 = Lane 16 alignment marker is not locked	RO
1.283.7	Lane 15 aligned	1 = Lane 15 alignment marker is locked 0 = Lane 15 alignment marker is not locked	RO
1.283.6	Lane 14 aligned	1 = Lane 14 alignment marker is locked 0 = Lane 14 alignment marker is not locked	RO
1.283.5	Lane 13 aligned	1 = Lane 13 alignment marker is locked 0 = Lane 13 alignment marker is not locked	RO
1.283.4	Lane 12 aligned	1 = Lane 12 alignment marker is locked 0 = Lane 12 alignment marker is not locked	RO

Table 45–90—RS-FEC PCS alignment status 4 register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
1.283.3	Lane 11 aligned	1 = Lane 11 alignment marker is locked 0 = Lane 11 alignment marker is not locked	RO
1.283.2	Lane 10 aligned	1 = Lane 10 alignment marker is locked 0 = Lane 10 alignment marker is not locked	RO
1.283.1	Lane 9 aligned	1 = Lane 9 alignment marker is locked 0 = Lane 9 alignment marker is not locked	RO
1.283.0	Lane 8 aligned	1 = Lane 8 alignment marker is locked 0 = Lane 8 alignment marker is not locked	RO

[a]RO = Read only

45.2.1.115.1 Lane 19 aligned (1.283.11)

When read as a one, bit 1.283.11 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 19. When read as a zero, bit 1.283.11 indicates that the RS-FEC transmit function lane 19 has not achieved alignment marker lock. This bit reflects the state of am_lock[19] (see 91.5.2.2).

45.2.1.115.2 Lane 18 aligned (1.283.10)

When read as a one, bit 1.283.10 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 18. When read as a zero, bit 1.283.10 indicates that the RS-FEC transmit function lane 18 has not achieved alignment marker lock. This bit reflects the state of am_lock[18] (see 91.5.2.2).

45.2.1.115.3 Lane 17 aligned (1.283.9)

When read as a one, bit 1.283.9 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 17. When read as a zero, bit 1.283.9 indicates that the RS-FEC transmit function lane 17 has not achieved alignment marker lock. This bit reflects the state of am_lock[17] (see 91.5.2.2).

45.2.1.115.4 Lane 16 aligned (1.283.8)

When read as a one, bit 1.283.8 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 16. When read as a zero, bit 1.283.8 indicates that the RS-FEC transmit function lane 16 has not achieved alignment marker lock. This bit reflects the state of am_lock[16] (see 91.5.2.2).

45.2.1.115.5 Lane 15 aligned (1.283.7)

When read as a one, bit 1.283.7 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 15. When read as a zero, bit 1.283.7 indicates that the RS-FEC transmit function lane 15 has not achieved alignment marker lock. This bit reflects the state of am_lock[15] (see 91.5.2.2).

45.2.1.115.6 Lane 14 aligned (1.283.6)

When read as a one, bit 1.283.6 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 14. When read as a zero, bit 1.283.6 indicates that the RS-FEC transmit function lane 14 has not achieved alignment marker lock. This bit reflects the state of am_lock[14] (see 91.5.2.2).

45.2.1.115.7 Lane 13 aligned (1.283.5)

When read as a one, bit 1.283.5 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 13. When read as a zero, bit 1.283.5 indicates that the RS-FEC transmit function lane 13 has not achieved alignment marker lock. This bit reflects the state of am_lock[13] (see 91.5.2.2).

45.2.1.115.8 Lane 12 aligned (1.283.4)

When read as a one, bit 1.283.4 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 12. When read as a zero, bit 1.283.4 indicates that the RS-FEC transmit function lane 12 has not achieved alignment marker lock. This bit reflects the state of am_lock[12] (see 91.5.2.2).

45.2.1.115.9 Lane 11 aligned (1.283.3)

When read as a one, bit 1.283.3 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 11. When read as a zero, bit 1.283.3 indicates that the RS-FEC transmit function lane 11 has not achieved alignment marker lock. This bit reflects the state of am_lock[11] (see 91.5.2.2).

45.2.1.115.10 Lane 10 aligned (1.283.2)

When read as a one, bit 1.283.2 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 10. When read as a zero, bit 1.283.2 indicates that the RS-FEC transmit function lane 10 has not achieved alignment marker lock. This bit reflects the state of am_lock[10] (see 91.5.2.2).

45.2.1.115.11 Lane 9 aligned (1.283.1)

When read as a one, bit 1.283.1 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 9. When read as a zero, bit 1.283.1 indicates that the RS-FEC transmit function lane 9 has not achieved alignment marker lock. This bit reflects the state of am_lock[9] (see 91.5.2.2).

45.2.1.115.12 Lane 8 aligned (1.283.0)

When read as a one, bit 1.283.0 indicates that the RS-FEC transmit function has achieved alignment marker lock for service interface lane 8. When read as a zero, bit 1.283.0 indicates that the RS-FEC transmit function lane 8 has not achieved alignment marker lock. This bit reflects the state of am_lock[8] (see 91.5.2.2).

45.2.1.116 BASE-R FEC corrected blocks counter, lanes 0 through 19

(Register 1.300, 1.301, 1.302, 1.303, 1.304, 1.305, 1.306, 1.307, 1.308, 1.309, 1.310, 1.311, 1.312, 1.313, 1.314, 1.315, 1.316, 1.317, 1.318, 1.319, 1.320, 1.321, 1.322, 1.323, 1.324, 1.325, 1.326, 1.327, 1.328, 1.329, 1.330, 1.331, 1.332, 1.333, 1.334, 1.335, 1.336, 1.337, 1.338, 1.339)

For multi-PCS lane BASE-R PHYs, the even-numbered registers in this set are defined similarly to register 1.172 (see 45.2.1.94) but for lanes 0 through 19 respectively of multi-PCS lane PHYs. The odd-numbered

registers in this set are defined similarly to register 1.173 (see 45.2.1.94) but for lanes 0 through 19 respectively of multilane PHYs. Note that the lane numbers refer to the service interface lanes and do not necessarily correspond to the PCS lane numbers.

Registers corresponding to lanes that are not used for the selected PHY shall return all zeros.

45.2.1.117 BASE-R FEC uncorrected blocks counter, lanes 0 through 19

(Register 1.700, 1.701, 1.702, 1.703, 1.704, 1.705, 1.706, 1.707, 1.708, 1.709, 1.710, 1.711, 1.712, 1.713, 1.714, 1.715, 1.716, 1.717, 1.718, 1.719, 1.720, 1.721, 1.722, 1.723, 1.724, 1.725, 1.726, 1.727, 1.728, 1.729, 1.730, 1.731, 1.732, 1.733, 1.734, 1.735, 1.736, 1.737, 1.738, 1.739)

For multi-PCS lane BASE-R PHYs, the even-numbered registers in this set are defined similarly to register 1.174 (see 45.2.1.95) but for lanes 0 through 19 respectively of multi-PCS lane PHYs. The odd-numbered registers in this set are defined similarly to register 1.175 (see 45.2.1.95) but for lanes 0 through 19 respectively of multilane PHYs. Note that the lane numbers refer to the service interface lanes and do not necessarily correspond to the PCS lane numbers.

Registers corresponding to lanes that are not used for the selected PHY shall return all zeros.

45.2.1.118 BASE-R LP coefficient update register, lanes 1 through 9

(Register 1.1101, 1.1102, 1.1103, 1.1104, 1.1105, 1.1106, 1.1107, 1.1108, 1.1109)

For multi-PCS lane BASE-R PHYs, these registers are defined similarly to register 1.152 (which is used for lane 0, see 45.2.1.82) but for lanes 1 through 9 respectively of multi-PCS lane PHYs.

Registers corresponding to lanes that are not used for the selected PHY shall return all zeros.

45.2.1.119 BASE-R LP status report register, lanes 1 through 9

(Register 1.1201, 1.1202, 1.1203, 1.1204, 1.1205, 1.1206, 1.1207, 1.1208, 1.1209)

For multi-PCS lane BASE-R PHYs, these registers are defined similarly to register 1.153 (which is used for lane 0, see 45.2.1.83) but for lanes 1 through 9 respectively of multi-PCS lane PHYs.

Registers corresponding to lanes that are not used for the selected PHY shall return all zeros.

45.2.1.120 BASE-R LD coefficient update register, lanes 1 through 9

(Register 1.1301, 1.1302, 1.1303, 1.1304, 1.1305, 1.1306, 1.1307, 1.1308, 1.1309)

For multi-PCS lane BASE-R PHYs, these registers are defined similarly to register 1.154 (which is used for lane 0, see 45.2.1.84) but for lanes 1 through 9 respectively of multi-PCS lane PHYs.

Registers corresponding to lanes that are not used for the selected PHY shall return all zeros.

45.2.1.121 BASE-R LD status report register, lanes 1 through 9

(Register 1.1401, 1.1402, 1.1403, 1.1404, 1.1405, 1.1406, 1.1407, 1.1408, 1.1409)

For multi-PCS lane BASE-R PHYs, these registers are defined similarly to register 1.155 (which is used for lane 0, see 45.2.1.85) but for lanes 1 through 9 respectively of multi-PCS lane PHYs.

Registers corresponding to lanes that are not used for the selected PHY shall return all zeros.

45.2.1.122 PMD training pattern lanes 0 through 3 (Register 1.1450 through 1.1453)

The assignment of bits in the PMD training pattern lane 0 register is shown in Table 45–91. The assignment of bits in the PMD training pattern lanes 1 through 3 registers are defined similarly to lane 0. Register 1.1450 controls the PMD training pattern for PMD lane 0; register 1.1451 controls the PMD training pattern for PMD lane 1; etc.

Table 45–91—PMD training pattern lane 0 bit definitions

Bit(s)	Name	Description	R/W[a]
1.1450.15:13	Reserved	Value always 0	RO
1.1450.12:11	Polynomial identifier	Identifier (0, 1, 2, or 3) selecting polynomial for PRBS	R/W
1.1450.10:0	Seed	11 bit, binary seed for sequence	R/W

[a]R/W = Read/Write, RO = Read only

Register bits 12:11 contain a 2-bit identifier that selects the polynomial used for training in the particular PMD lane according to the definition in 92.7.12. The polynomial identifier for each lane should be unique; two lanes having the same identifier could impair operation of the PMD control function. The default identifiers are (binary): for lane 0, 00; for lane 1, 01; for lane 2, 10; for lane 3, 11. Register bits 10:0 contain the 11-bit seed for the sequence, where register bit 0 gives seed bit S0; register bit 1 gives seed bit S1; etc., through register bit 10 gives seed bit S10. The default seeds are (binary, S0 is left-most bit): for lane 0, 10101111110; for lane 1, 11001000101; for lane 2, 11100101101; for lane 3, 11110110110. This produces the following initial output (hexadecimal representation where the hex symbols are transmitted from left to right and the most significant bit of each hex symbol is transmitted first): for lane 0, fbf1cb3e; for lane 1, fbb1e665; for lane 2, f3fdae46; for lane 3, f2ffa46b.

45.2.1.123 Test-pattern ability (Register 1.1500)

The test-pattern ability register is used for PHY types that implement square wave testing and PRBS testing in the PMA. These functions are described in 83.5.10. The assignment of bits in the test-pattern ability register is shown in Table 45–92.

Table 45–92—Test-pattern ability register bit definitions

Bit(s)	Name	Description	R/W[a]
1.1500.15:13	Reserved	Value always 0	RO
1.1500.12	Square wave test ability	1 = Square wave testing supported 0 = Square wave testing not supported	RO
1.1500.11:6	Reserved	Value always 0	RO
1.1500.5	PRBS9 Tx generator ability	1 = PRBS9 transmit direction pattern generator supported 0 = PRBS9 transmit direction pattern generator not supported	RO
1.1500.4	PRBS9 Rx generator ability	1 = PRBS9 receive direction pattern generator supported 0 = PRBS9 receive direction pattern generator not supported	RO

Table 45–92—Test-pattern ability register bit definitions *(continued)*

1.1500.3	PRBS31 Tx generator ability	1 = PRBS31 Transmit direction pattern generator supported 0 = PRBS31 Tx direction pattern generator not supported	RO
1.1500.2	PRBS31 Tx checker ability	1 = PRBS31 Transmit direction pattern checker supported 0 = PRBS31 Tx direction pattern checker not supported	RO
1.1500.1	PRBS31 Rx generator ability	1 = PRBS31 Receive direction pattern generator supported 0 = PRBS31 Rx direction pattern generator not supported	RO
1.1500.0	PRBS31 Rx checker ability	1 = PRBS31 Receive direction pattern checker supported 0 = PRBS31 Rx direction pattern checker not supported	RO

[a]RO = Read only

If square wave testing is supported and this register is implemented then bit 1.1500.12 shall be set to 1. The square wave test is controlled by register 1.1510 (see 45.2.1.125).

If PRBS9 pattern testing is supported and this register is implemented then bit 1.1500.5 shall indicate the generation ability in the transmit direction and bit 1.1500.4 shall indicate the generation ability in the receive direction. The PRBS pattern test is controlled by register 1.1501 (see 45.2.1.124).

If PRBS31 pattern testing is supported and this register is implemented then bit 1.1500.3 shall indicate the generation ability in the transmit direction and bit 1.1500.1 shall indicate the generation ability in the receive direction. Bit 1.1500.2 shall indicate the checker ability in the transmit direction and bit 1.1500.0 shall indicate the checker ability in the receive direction. The PRBS pattern test is controlled by register 1.1501 (see 45.2.1.124).

45.2.1.124 PRBS pattern testing control (Register 1.1501)

The PRBS pattern testing control register is used for PHY types that implement PRBS pattern testing in the PMA. This function is described in 83.5.10. The assignment of bits in the PRBS pattern testing control register is shown in Table 45–93.

Table 45–93—PRBS pattern testing control register bit definitions

Bit(s)	Name	Description	R/W[a]
1.1501.15:12	Reserved	Value always 0	RO
1.1501.11	Transmitter linearity test pattern enable	1 = Enable transmitter linearity test-pattern 0 = Disable transmitter linearity test-pattern	R/W
1.1501.10	QPRBS13 pattern enable	1 = Enable QPRBS13 test-pattern 0 = Disable QPRBS13 test-pattern	R/W
1.1501.9	JP03B pattern enable	1 = Enable JP03B test-pattern 0 = Disable JP03B test-pattern	R/W
1.1501.8	JP03A pattern enable	1 = Enable JP03A test-pattern 0 = Disable JP03A test-pattern	R/W
1.1501.7	PRBS31 pattern enable	1 = Enable PRBS31 test-pattern 0 = Disable PRBS31 test-pattern	R/W

Table 45–93—PRBS pattern testing control register bit definitions *(continued)*

1.1501.6	PRBS9 pattern enable	1 = Enable PRBS9 test-pattern 0 = Disable PRBS9 test-pattern	R/W
1.1501.5:4	Reserved	Value always 0	RO
1.1501.3	Tx generator enable	1 = Enable transmit direction test-pattern generator 0 = Disable transmit direction test-pattern generator	R/W
1.1501.2	Tx checker enable	1 = Enable transmit direction test-pattern checker 0 = Disable transmit direction test-pattern checker	R/W
1.1501.1	Rx generator enable	1 = Enable receive direction test-pattern generator 0 = Disable receive direction test-pattern generator	R/W
1.1501.0	Rx checker enable	1 = Enable receive direction test-pattern checker 0 = Disable receive direction test-pattern checker	R/W

[a]R/W = Read/Write, RO = Read only

Register 1.1501, bit 7 enables testing with the PRBS31 pattern defined in 83.5.10. Register 1.1501, bit 6 enables testing with the PRBS9 pattern defined in 83.5.10. The assertion of register 1.1501 bits 7 and 6 is mutually exclusive. If both bits are asserted the behavior is undefined. The assertion of register 1.1501, bits 7 and 6 works in conjunction with register 1.1501, bits 3:0. If none of the bits 3:0 are asserted then bits 7 and 6 have no effect.

Register 1.1501, bit 3 enables PRBS generation in the transmit direction. Register 1.1501, bit 2 enables PRBS checking in the transmit direction. Register 1.1501, bit 1 enables PRBS generation in the receive direction. Register 1.1501, bit 0 enables PRBS checking in the receive direction. If neither of the bits 7 and 6 are asserted then bits 3:0 have no effect.

Register 1.1501 bit 8 enables testing with the JP03A pattern defined in 94.2.9.1 for 100GBASE-KP4 PMA/PMD. Register 1.1501 bit 9 enables testing with the JP03B pattern defined in 94.2.9.2 for 100GBASE-KP4 PMA/PMD. Register field 1.1501 bit 10 enables testing with the QPRBS13 pattern defined in 94.2.9.3 for 100GBASE-KP4 PMA/PMD. Register field 1.1501 bit 11 enables the transmitter linearity test pattern defined in 94.2.9.4 for 100GBASE-KP4 PMA/PMD. The assertion of bits 1.1501.8, 1.1501.9, 1.1501.10, 1.1501.11 are mutually exclusive. If more than one bit is asserted, the behavior is undefined. The assertion of 1.1501.8, 1.1501.9, 1.501.10, and 1.1501.11 operates in conjunction with register 1.1501 bit 3 for 100GBASE-KP4 PMA/PMD. For other PMA/PMD types or if bit 1.1501.3 is not asserted, then 1.1501.8, 1.1501.9, 1.501.10, and 1.1501.11 have no effect.

45.2.1.125 Square wave testing control (Register 1.1510)

The square wave testing control register is used for PHY types that implement transmit square wave testing in the PMA. This function is described in 83.5.10. The assignment of bits in the square wave testing control register is shown in Table 45–94.

Register 1.1510, bits 0 through 9 enable square wave output on PMA lanes 0 through 9 respectively. Lanes for which a square wave pattern is not enabled act as determined by other registers.

Table 45–94—Square wave testing control register bit definitions

Bit(s)	Name	Description	R/W[a]
1.1510.15:10	Reserved	Value always 0	RO
1.1510.9	Lane 9 SW enable	1 = Enable square wave on lane 9 0 = Disable square wave on lane 9	R/W
1.1510.8	Lane 8 SW enable	1 = Enable square wave on lane 8 0 = Disable square wave on lane 8	R/W
1.1510.7	Lane 7 SW enable	1 = Enable square wave on lane 7 0 = Disable square wave on lane 7	R/W
1.1510.6	Lane 6 SW enable	1 = Enable square wave on lane 6 0 = Disable square wave on lane 6	R/W
1.1510.5	Lane 5 SW enable	1 = Enable square wave on lane 5 0 = Disable square wave on lane 5	R/W
1.1510.4	Lane 4 SW enable	1 = Enable square wave on lane 4 0 = Disable square wave on lane 4	R/W
1.1510.3	Lane 3 SW enable	1 = Enable square wave on lane 3 0 = Disable square wave on lane 3	R/W
1.1510.2	Lane 2 SW enable	1 = Enable square wave on lane 2 0 = Disable square wave on lane 2	R/W
1.1510.1	Lane 1 SW enable	1 = Enable square wave on lane 1 0 = Disable square wave on lane 1	R/W
1.1510.0	Lane 0 SW enable	1 = Enable square wave on lane 0 0 = Disable square wave on lane 0	R/W

[a]R/W = Read/Write, RO = Read only

45.2.1.126 PRBS Tx pattern testing error counter (Register 1.1600, 1.1601, 1.1602, 1.1603, 1.1604, 1.1605, 1.1606, 1.1607, 1.1608, 1.1609)

The PRBS Tx pattern testing error counter registers are used for PHY types that implement PRBS Tx pattern testing in the PMA. This function is described in 83.5.10. The assignment of bits in the PRBS Tx pattern testing error counter registers are shown in Table 45–95. Register 1.1600 contains the PRBS Tx pattern testing error counter for lane 0, register 1.1601 contains the PRBS Tx pattern testing error counter for lane 1, and registers 1.1602 through 1.1609 contain the PRBS Tx pattern testing error counters for lanes 2 through 9 respectively. Counters corresponding to lanes that are not implemented in a PMA shall read all zeros.

Table 45–95—PRBS Tx pattern testing error counter

Bit(s)	Name	Description	R/W[a]
1.1600[b].15:0	Error counter		RO, NR

[a]RO = Read only, NR = Non Roll-over.
[b]All instances of address 1.1600 also apply to addresses 1.1601 through 1.1609.

The PRBS Tx pattern testing error counter is a 16-bit counter as defined in 83.5.10. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the PMA reset. These bits shall be held at all ones in the case of overflow.

45.2.1.127 PRBS Rx pattern testing error counter (Register 1.1700, 1.1701, 1.1702, 1.1703, 1.1704, 1.1705, 1.1706, 1.1707, 1.1708, 1.1709)

The PRBS Rx pattern testing error counter registers are used for PHY types that implement PRBS Rx pattern testing in the PMA. This function is described in 83.5.10. The assignment of bits in the PRBS Rx pattern testing error counter registers is identical to the PRBS Tx pattern testing error counter as shown in Table 45–95. Register 1.1700 contains the PRBS Rx pattern testing error counter for lane 0, register 1.1701 contains the PRBS Rx pattern testing error counter for lane 1, and registers 1.1702 through 1.1709 contain the PRBS Rx pattern testing error counters for lanes 2 through 9 respectively. Counters corresponding to lanes that are not implemented in a PMA shall read all zeros.

The PRBS Rx pattern testing error counter is a 16-bit counter as defined in 83.5.10. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the PMA reset. These bits shall be held at all ones in the case of overflow.

45.2.1.128 TimeSync PMA/PMD capability (Register 1.1800)

The TimeSync PMA/PMD capability register (see Table 45–96) indicates the capability of the PMA/PMD to report the transmit and receive data delay, stored in registers 1.1801 and 1.1805 through 1.1808 respectively.

Table 45–96—TimeSync PMA/PMD capability

Bit(s)	Name	Description	R/W[a]
1.1800.15:2	Reserved	Value always 0	RO
1.1800.1	TimeSync transmit path data delay	1 = PMA/PMD provides information on transmit path data delay in registers 1.1801 through 1.1804 0 = PMA/PMD does not provide information on transmit path data delay	RO
1.1800.0	TimeSync receive path data delay	1 = PMA/PMD provides information on receive path data delay in registers 1.1805 through 1.1808 0 = PMA/PMD does not provide information on receive path data delay	RO

[a]RO = Read only.

45.2.1.129 TimeSync PMA/PMD transmit path data delay (Registers 1.1801, 1.1802, 1.1803, 1.1804)

The TimeSync PMA/PMD transmit path data delay register contains the maximum (Registers 1.1801, 1.1802, see Table 45–97) and minimum (Registers 1.1803, 1.1804, see Table 45–97) values of the transmit

path data delay. The transmit path data delay is expressed in units of ns. The values contained in these registers are valid when the link is established, as indicated by bit 2 in Register 1.1 (see 45.2.1.2.4).

Table 45–97—TimeSync PMA/PMD transmit path data delay register

Bit(s)	Name	Description	R/W[a]
1.1801.15:0	Maximum PMA/PMD transmit path data delay, lower	PMA/PMD_delay_TX_max [15:0]	RO, MW
1.1802.15:0	Maximum PMA/PMD transmit path data delay, upper	PMA/PMD_delay_TX_max [31:16]	RO, MW
1.1803.15:0	Minimum PMA/PMD transmit path data delay, lower	PMA/PMD_delay_TX_min [15:0]	RO, MW
1.1804.15:0	Minimum PMA/PMD transmit path data delay, upper	PMA/PMD_delay_TX_min [31:16]	RO, MW

[a]RO = Read only, MW = Multi-word.

45.2.1.130 TimeSync PMA/PMD receive path data delay (Registers 1.1805, 1.1806, 1.1807, 1.1808)

The TimeSync PMA/PMD receive path data delay register contains the maximum (Registers 1.1805, 1.1806, see Table 45–98) and minimum (Registers 1.1807, 1.1808, see Table 45–98) values of the receive path data delay. The receive path data delay is expressed in units of ns. The values contained in these registers are valid when the link is established, as indicated by bit 2 in Register 1.1 (see 45.2.1.2.4).

Table 45–98—TimeSync PMA/PMD receive path data delay register

Bit(s)	Name	Description	R/W[a]
1.1805.15:0	Maximum PMA/PMD receive path data delay, lower	PMA/PMD_delay_RX_max [15:0]	RO, MW
1.1806.15:0	Maximum PMA/PMD receive path data delay, upper	PMA/PMD_delay_RX_max [31:16]	RO, MW
1.1807.15:0	Minimum PMA/PMD receive path data delay, lower	PMA/PMD_delay_RX_min [15:0]	RO, MW
1.1808.15:0	Minimum PMA/PMD receive path data delay, upper	PMA/PMD_delay_RX_min [31:16]	RO, MW

[a]RO = Read only, MW = Multi-word.

45.2.2 WIS registers

The assignment of registers in the WIS is shown in Table 45–99. For the WIS octet fields, bit 8 of the corresponding field in the WIS frame maps to the lowest numbered bit of the field in the register.

Table 45–99—WIS registers

Register address	Register name	Subclause
2.0	WIS control 1	45.2.2.1
2.1	WIS status 1	45.2.2.2
2.2, 2.3	WIS device identifier	45.2.2.3
2.4	WIS speed ability	45.2.2.4
2.5, 2.6	WIS devices in package	45.2.2.5
2.7	10G WIS control 2	45.2.2.6
2.8	10G WIS status 2	45.2.2.7
2.9	10G WIS test-pattern error counter	45.2.2.8
2.10 through 2.13	Reserved	
2.14, 2.15	WIS package identifier	45.2.2.9
2.16 through 2.32	Reserved	
2.33	10G WIS status 3	45.2.2.10
2.34 through 2.36	Reserved	
2.37	10G WIS far end path block error count	45.2.2.11
2.38	Reserved	
2.39 through 2.46	10G WIS J1 transmit	45.2.2.12
2.47 through 2.54	10G WIS J1 receive	45.2.2.13
2.55, 2.56	10G WIS far end line BIP errors	45.2.2.14
2.57, 2.58	10G WIS line BIP errors	45.2.2.15
2.59	10G WIS path block error count	45.2.2.16
2.60	10G WIS section BIP error count	45.2.2.17
2.61 through 2.63	Reserved	
2.64 through 2.71	10G WIS J0 transmit	45.2.2.18
2.72 through 2.79	10G WIS J0 receive	45.2.2.19
2.80 through 2.1799	Reserved	
2.1800	TimeSync WIS capability	45.2.2.20
2.1801 through 2.1804	TimeSync WIS transmit path data delay	45.2.2.21

Table 45–99—WIS registers *(continued)*

Register address	Register name	Subclause
2.1805 through 2.1808	TimeSync WIS receive path data delay	45.2.2.22
2.1809 through 2.32767	Reserved	
2.32 768 through 2.65 535	Vendor specific	

45.2.2.1 WIS control 1 register (Register 2.0)

The assignment of bits in the WIS control 1 register is shown in Table 45–100. The default value for each bit of the WIS control 1 register should be chosen so that the initial state of the device upon power up or reset is a normal operational state without management intervention.

Table 45–100—WIS control 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
2.0.15	Reset	1 = WIS reset 0 = Normal operation	R/W SC
2.0.14	Loopback	1 = Enable loopback mode 0 = Disable loopback mode	R/W
2.0.13	Speed selection	1 = Operation at 10 Gb/s and above 0 = Unspecified	R/W
2.0.12	Reserved	Value always 0	RO
2.0.11	Low power	1 = Low-power mode 0 = Normal operation	R/W
2.0.10:7	Reserved	Value always 0	RO
2.0.6	Speed selection	1 = Operation at 10 Gb/s and above 0 = Unspecified	R/W
2.0.5:2	Speed selection	5 4 3 2 1 x x x = Reserved x 1 x x = Reserved x x 1 x = Reserved 0 0 0 1 = Reserved 0 0 0 0 = 10 Gb/s	R/W
2.0.1:0	Reserved	Value always 0	RO

[a]R/W = Read/Write, SC = Self-clearing, RO = Read only

45.2.2.1.1 Reset (2.0.15)

Resetting a WIS is accomplished by setting bit 2.0.15 to a one. This action shall set all to their default states. As a consequence, this action may change the internal state of the WIS and the state of the physical link.

This action may also initiate a reset in any other MMDs that are instantiated in the same package. This bit is self-clearing, and a WIS shall return a value of one in bit 2.0.15 when a reset is in progress and a value of zero otherwise. A WIS is not required to accept a write transaction to any of its registers until the reset process is completed. The reset process shall be completed within 0.5 s from the setting of bit 2.0.15. During a reset, a WIS shall respond to reads from register bits 2.0.15 and 2.8.15:14. All other register bits should be ignored.

NOTE—This operation may interrupt data communication.

45.2.2.1.2 Loopback (2.0.14)

The WIS shall be placed in a loopback mode of operation when bit 2.0.14 is set to a one. When bit 2.0.14 is set to a one, the WIS shall ignore all data presented to it by the PMA sublayer. When bit 2.0.14 is set to a one, the WIS shall accept data on the transmit path and return it on the receive path. For 10 Gb/s operation, the detailed behavior of the WIS during loopback is specified in 50.3.9.

The default value of bit 2.0.14 is zero.

NOTE—The signal path through the WIS that is exercised in the loopback mode of operation is implementation specific, but it is recommended that the signal path encompass as much of the WIS circuitry as is practical. The intention of providing this loopback mode of operation is to permit a diagnostic or self-test function to perform the transmission and reception of a PDU, thus testing the transmit and receive data paths. Other loopback signal paths may be enabled using loopback controls within other MMDs.

45.2.2.1.3 Low power (2.0.11)

A WIS may be placed into a low-power mode by setting bit 2.0.11 to a one. This action may also initiate a low-power mode in any other MMDs that are instantiated in the same package. The low-power mode is exited by resetting the WIS. The behavior of the WIS in transition to and from the low-power mode is implementation specific and any interface signals should not be relied upon. While in the low-power mode, the device shall, as a minimum, respond to management transactions necessary to exit the low-power mode. The default value of bit 2.0.11 is zero.

45.2.2.1.4 Speed selection (2.0.13, 2.0.6, and 2.0.5:2)

Speed selection bits 2.0.13 and 2.0.6 shall both be written as a one. Any attempt to change the bits to an invalid setting shall be ignored. These two bits are set to one in order to make them compatible with Clause 22.

The speed of the WIS may be selected using bits 5 through 2. The speed abilities of the WIS are advertised in the WIS speed ability register. A WIS may ignore writes to the WIS speed selection bits that select speeds it has not advertised in the WIS speed ability register. It is the responsibility of the STA entity to ensure that mutually acceptable speeds are applied consistently across all the MMDs on a particular PHY.

The WIS speed selection defaults to a supported ability.

45.2.2.2 WIS status 1 register (Register 2.1)

The assignment of bits in the WIS status 1 register is shown in Table 45–101. All the bits in the WIS status 1 register are read only; a write to the WIS status 1 register shall have no effect.

Table 45–101—WIS status 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
2.1.15:8	Reserved	Value always 0	RO
2.1.7	Fault	1 = Fault condition 0 = No fault condition	RO/LH
2.1.6:3	Reserved	Value always 0	RO
2.1.2	Link status	1 = WIS link up 0 = WIS link down	RO/LL
2.1.1	Low-power ability	1 = WIS supports low-power mode 0 = WIS does not support low-power mode	RO
2.1.0	Reserved	Value always 0	RO

[a]RO = Read only, LH = Latching high, LL = Latching low

45.2.2.2.1 Fault (2.1.7)

When read as a one, bit 2.1.7 indicates that the WIS has detected a fault condition. When read as a zero, bit 2.1.7 indicates that the WIS has not detected a fault condition. The fault bit shall be implemented with latching high behavior.

The default value of bit 2.1.7 is zero.

45.2.2.2.2 Link status (2.1.2)

When read as a one, bit 2.1.2 indicates that the WIS receive link is up. When read as a zero, bit 2.1.2 indicates that the WIS receive link is down. The link status bit shall be implemented with latching low behavior.

45.2.2.2.3 Low-power ability (2.1.1)

When read as a one, bit 2.1.1 indicates that the WIS supports the low-power feature. When read as a zero, bit 2.1.1 indicates that the WIS does not support the low-power feature. If a WIS supports the low-power feature, then it is controlled using the low-power bit in the WIS control register.

45.2.2.3 WIS device identifier (Registers 2.2 and 2.3)

Registers 2.2 and 2.3 provide a 32-bit value, which may constitute a unique identifier for a particular type of WIS. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the device manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A WIS may return a value of zero in each of the 32 bits of the WIS device identifier.

The format of the WIS device identifier is specified in 22.2.4.3.1.

45.2.2.4 WIS speed ability (Register 2.4)

The assignment of bits in the WIS speed ability register is shown in Table 45–102.

Table 45–102—WIS speed ability register bit definitions

Bit(s)	Name	Description	R/W[a]
2.4.15:1	Reserved for future speeds	Value always 0	RO
2.4.0	10G capable	1 = WIS is capable of operating at 10 Gb/s 0 = WIS is not capable of operating at 10 Gb/s	RO

[a]RO = Read only

45.2.2.4.1 10G capable (2.4.0)

When read as a one, bit 2.4.0 indicates that the WIS is able to operate at a data rate of 10 Gb/s (9.58 Gb/s payload rate). When read as a zero, bit 2.4.0 indicates that the WIS is not able to operate at a data rate of 10 Gb/s (9.58 Gb/s payload rate).

45.2.2.5 WIS devices in package (Registers 2.5 and 2.6)

The WIS devices in package registers are defined in Table 45–2.

45.2.2.6 10G WIS control 2 register (Register 2.7)

The assignment of bits in the 10G WIS control 2 register is shown in Table 45–103. The default value for each bit of the 10G WIS control 2 register should be chosen so that the initial state of the device upon power up or reset is a normal operational state without management intervention.

Table 45–103—10G WIS control 2 register bit definitions

Bit(s)	Name	Description	R/W[a]
2.7.15:6	Reserved	Value always 0	RO
2.7.5	PRBS31 receive test-pattern enable	1 = Enable PRBS31 test-pattern mode on the receive path 0 = Disable PRBS31 test-pattern mode on the receive path	R/W
2.7.4	PRBS31 transmit test-pattern enable	1 = Enable PRBS31 test-pattern mode on the transmit path 0 = Disable PRBS31 test-pattern mode on the transmit path	R/W
2.7.3	Test-pattern selection	1 = Select square wave test pattern 0 = Select mixed-frequency test pattern	R/W
2.7.2	Receive test-pattern enable	1 = Enable test-pattern mode on the receive path 0 = Disable test-pattern mode on the receive path	R/W
2.7.1	Transmit test-pattern enable	1 = Enable test-pattern mode on the transmit path 0 = Disable test-pattern mode on the transmit path	R/W
2.7.0	PCS type selection	1 = Select 10GBASE-W PCS type 0 = Select 10GBASE-R PCS type	R/W

[a]RO = Read only, R/W = Read/Write

45.2.2.6.1 PRBS31 receive test-pattern enable (2.7.5)

If the WIS supports the optional PRBS31 (see 49.2.8) pattern testing advertised in bit 2.8.1 and the mandatory receive test-pattern enable bit (2.7.2) is not one, setting bit 2.7.5 to a one shall set the receive path of the WIS into the PRBS31 test-pattern mode. Setting bit 2.7.5 to a zero shall disable the PRBS31 test-pattern mode on the receive path of the WIS. The behavior of the WIS when in PRBS31 test-pattern mode is specified in 50.3.8.2.

45.2.2.6.2 PRBS31 transmit test-pattern enable (2.7.4)

If the WIS supports the optional PRBS31 pattern testing advertised in bit 2.8.1 and the mandatory transmit test-pattern enable bit (2.7.1) is not one, then setting bit 2.7.4 to a one shall set the transmit path of the WIS into the PRBS31 test-pattern mode. Setting bit 2.7.4 to a zero shall disable the PRBS31 test-pattern mode on the transmit path of the WIS. The behavior of the WIS when in PRBS31 test-pattern mode is specified in 50.3.8.2.

45.2.2.6.3 Test-pattern selection (2.7.3)

Bit 2.7.3 controls the type of pattern sent by the transmitter when in test-pattern mode. Setting bit 2.7.3 to a one shall select the square wave test pattern. Setting bit 2.7.3 to a zero shall select the mixed-frequency test pattern. The details of the test patterns are specified in Clause 50.

45.2.2.6.4 Receive test-pattern enable (2.7.2)

Setting bit 2.7.2 to a one shall set the receive path of the WIS into the test-pattern mode. Setting bit 2.7.2 to a zero shall disable the test-pattern mode on the receive path of the WIS. The behavior of the WIS when in test-pattern mode is specified in Clause 50.

45.2.2.6.5 Transmit test-pattern enable (2.7.1)

Setting bit 2.7.1 to a one shall set the transmit path of the WIS into the test-pattern mode. Setting bit 2.7.1 to a zero shall disable the test-pattern mode on the transmit path of the WIS. The behavior of the WIS when in test-pattern mode is specified in Clause 50.

45.2.2.6.6 PCS type selection (2.7.0)

Setting bit 2.7.0 to a one shall enable the 10GBASE-W logic and set the speed of the WIS-PMA interface to 9.95328 Gb/s. Setting bit 2.7.0 to a zero shall disable the 10GBASE-W logic, set the speed of the PCS-PMA interface to 10.3125 Gb/s and bypass the data around the 10GBASE-W logic. A WIS that is only capable of supporting 10GBASE-W operation and is unable to support 10GBASE-R operation shall ignore values written to this bit and shall return a value of one when read. It is the responsibility of the STA entity to ensure that mutually acceptable MMD types are applied consistently across all the MMDs on a particular PHY.

45.2.2.7 10G WIS status 2 register (Register 2.8)

The assignment of bits in the 10G WIS status 2 register is shown in Table 45–104. All the bits in the 10G WIS status 2 register are read only; a write to the 10G WIS status 2 register shall have no effect.

45.2.2.7.1 Device present (2.8.15:14)

When read as <10>, bits 2.8.15:14 indicate that a device is present and responding at this register address. When read as anything other than <10>, bits 2.8.15:14 indicate that no device is present at this address or that the device is not functioning properly.

Table 45–104—10G WIS status 2 register bit definitions

Bit(s)	Name	Description	R/W[a]
2.8.15:14	Device present	15 14 1 0 = Device responding at this address 1 1 = No device responding at this address 0 1 = No device responding at this address 0 0 = No device responding at this address	RO
2.8.13:2	Reserved	Value always 0	RO
2.8.1	PRBS31 pattern testing ability	1 = WIS is able to support PRBS31 pattern testing 0 = WIS is not able to support PRBS31 pattern testing	RO
2.8.0	10GBASE-R ability	1 = WIS is able to support 10GBASE-R port types 0 = WIS is not able to support 10GBASE-R port types	RO

[a]RO = Read only

45.2.2.7.2 PRBS31 pattern testing ability (2.8.1)

When read as a one, bit 2.8.1 indicates that the WIS is able to support PRBS31 pattern testing. When read as a zero, bit 2.8.1 indicates that the WIS is not able to support PRBS31 pattern testing. If the WIS is able to support PRBS31 pattern testing, then the pattern generation and checking is controlled using bits 2.7.5:4.

45.2.2.7.3 10GBASE-R ability (2.8.0)

When read as a one, bit 2.8.0 indicates that the WIS is able to bypass the WIS logic and adjust the XSBI interface speed to support 10GBASE-R port types. When read as a zero, bit 2.8.0 indicates that the WIS is not able to bypass the WIS logic and cannot support 10GBASE-R port types.

45.2.2.8 10G WIS test-pattern error counter register (Register 2.9)

The assignment of bits in the 10G WIS test-pattern error counter register is shown in Table 45–105. This register is only required when the PRBS31 pattern generation capability is supported.

Table 45–105—10G WIS test-pattern error counter register bit definitions

Bit(s)	Name	Description	R/W[a]
2.9.15:0	Test-pattern error counter	Error counter	RO

[a]RO = Read only

The test-pattern error counter is a sixteen bit counter that contains the number of errors received during a pattern test. These bits shall be reset to all zeros when the test-pattern error counter is read by the management function or upon execution of the WIS reset. These bits shall be held at all ones in the case of overflow. The test-pattern methodology is described in 49.2.8.

45.2.2.9 WIS package identifier (Registers 2.14 and 2.15)

Registers 2.14 and 2.15 provide a 32-bit value, which may constitute a unique identifier for a particular type of package that the WIS is instantiated within. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the package manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A WIS may return a value of zero in each of the 32 bits of the WIS package identifier.

A non-zero package identifier may be returned by one or more MMDs in the same package. The package identifier may be the same as the device identifier.

The format of the WIS package identifier is specified in 22.2.4.3.1.

45.2.2.10 10G WIS status 3 register (Register 2.33)

The assignment of bits in the 10G WIS status 3 register is shown in Table 45–106. All the bits in the 10G WIS status 3 register are read only; a write to the 10G WIS status 3 register shall have no effect.

Table 45–106—10G WIS status 3 register bit definitions

Bit(s)	Name	Description	R/W[a]
2.33.15:12	Reserved	Value always 0	RO
2.33.11	SEF	Severely errored frame	RO/LH
2.33.10	Far end PLM-P/LCD-P	1 = Far end path label mismatch / Loss of code-group delineation 0 = No far end path label mismatch / Loss of code-group delineation	RO/LH
2.33.9	Far end AIS-P/LOP-P	1 = Far end path alarm indication signal / Path loss of pointer 0 = No far end path alarm indication signal / Path loss of pointer	RO/LH
2.33.8	Reserved	Value always 0	RO
2.33.7	LOF	1 = Loss of frame flag raised 0 = Loss of frame flag lowered	RO/LH
2.33.6	LOS	1 = Loss of signal flag raised 0 = Loss of signal flag lowered	RO/LH
2.33.5	RDI-L	1 = Line remote defect flag raised 0 = Line remote defect flag lowered	RO/LH
2.33.4	AIS-L	1 = Line alarm indication flag raised 0 = Line alarm indication flag lowered	RO/LH
2.33.3	LCD-P	1 = Path loss of code-group delineation flag raised 0 = Path loss of code-group delineation flag lowered	RO/LH

Table 45–106—10G WIS status 3 register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
2.33.2	PLM-P	1 = Path label mismatch flag raised 0 = Path label mismatch flag lowered	RO/LH
2.33.1	AIS-P	1 = Path alarm indication signal raised 0 = Path alarm indication signal lowered	RO/LH
2.33.0	LOP-P	1 = Loss of pointer flag raised 0 = Loss of pointer flag lowered	RO/LH

[a]RO = Read only, LH = Latching high

45.2.2.10.1 SEF (2.33.11)

When read as a one, bit 2.33.11 indicates that the SEF flag has been raised by the WIS. When read as a zero, bit 2.33.11 indicates that the SEF flag is lowered. The SEF bit shall be implemented with latching high behavior.

The SEF functionality implemented by the WIS is described in 50.3.2.5.

45.2.2.10.2 Far end PLM-P/LCD-P (2.33.10)

When read as a one, bit 2.33.10 indicates that the far end path label mismatch/loss of code-group delineation flag has been raised. When read as a zero, bit 2.33.10 indicates that the far end path label mismatch/loss of code-group delineation flag is lowered. The far end PLM-P/LCD-P bit shall be implemented with latching high behavior.

The far end path label mismatch/loss of code-group delineation functionality implemented by the WIS is described in 50.3.2.5.

45.2.2.10.3 Far end AIS-P/LOP-P (2.33.9)

When read as a one, bit 2.33.9 indicates that the far end path alarm indication signal/path loss of pointer flag has been raised by the WIS. When read as a zero, bit 2.33.9 indicates that the far end path alarm indication signal/path loss of pointer flag is lowered. The far end AIS-P/LOP-P bit shall be implemented with latching high behavior.

The far end path alarm indication signal/path loss of pointer functionality implemented by the WIS is described in 50.3.2.5.

45.2.2.10.4 LOF (2.33.7)

When read as a one, bit 2.33.7 indicates that the loss of frame flag has been raised. When read as a zero, bit 2.33.7 indicates that the loss of frame flag is lowered. The LOF bit shall be implemented with latching high behavior.

The LOF functionality implemented by the WIS is described in 50.3.2.5.

45.2.2.10.5 LOS (2.33.6)

When read as a one, bit 2.33.6 indicates that the loss of signal flag has been raised. When read as a zero, bit 2.33.6 indicates that the loss of signal flag is lowered. The LOS bit shall be implemented with latching high behavior.

The LOS functionality implemented by the WIS is described in 50.3.2.5.

45.2.2.10.6 RDI-L (2.33.5)

When read as a one, bit 2.33.5 indicates that the line remote defect flag has been raised. When read as a zero, bit 2.33.5 indicates that the line remote defect flag is lowered. The RDI-L bit shall be implemented with latching high behavior.

The RDI-L functionality implemented by the WIS is described in 50.3.2.5.

45.2.2.10.7 AIS-L (2.33.4)

When read as a one, bit 2.33.4 indicates that the line alarm indication flag has been raised. When read as a zero, bit 2.33.4 indicates that the line alarm indication flag is lowered. The AIS-L bit shall be implemented with latching high behavior.

The AIS-L functionality implemented by the WIS is described in 50.3.2.5.

45.2.2.10.8 LCD-P (2.33.3)

When read as a one, bit 2.33.3 indicates that the loss of code-group delineation flag has been raised. When read as a zero, bit 2.33.3 indicates that the loss of code-group delineation flag is lowered. The LCD-P bit shall be implemented with latching high behavior.

The loss of code-group delineation functionality implemented by the WIS is described in 50.3.2.5.

45.2.2.10.9 PLM-P (2.33.2)

When read as a one, bit 2.33.2 indicates that the path label mismatch flag has been raised. When read as a zero, bit 2.33.2 indicates that the path label mismatch flag is lowered. The PLM-P bit shall be implemented with latching high behavior.

The PLM-P functionality implemented by the WIS is described in 50.3.2.5.

45.2.2.10.10 AIS-P (2.33.1)

When read as a one, bit 2.33.1 indicates that the path alarm indication signal has been raised. When read as a zero, bit 2.33.1 indicates that the path alarm indication signal is lowered. The AIS-P bit shall be implemented with latching high behavior.

The path alarm indication signal functionality implemented by the WIS is described in 50.3.2.5.

45.2.2.10.11 LOP-P (2.33.0)

When read as a one, bit 2.33.0 indicates that the loss of pointer flag has been raised. When read as a zero, bit 2.33.0 indicates that the loss of pointer flag is lowered. The LOP-P bit shall be implemented with latching high behavior.

The LOP-P functionality implemented by the WIS is described in 50.3.2.5.

45.2.2.11 10G WIS far end path block error count (Register 2.37)

The assignment of bits in the 10G WIS far end path block error count register is shown in Table 45–107.

Table 45–107—10G WIS far end path block error count register bit definitions

Bit(s)	Name	Description	R/W[a]
2.37.15:0	Far end path block error count	Far end path block error count	RO

[a]RO = Read only,

The 10G WIS far end path block error count is incremented by one whenever a far end path block error, defined in Annex 50A, is detected as described in 50.3.2.5. The counter wraps around to zero when it is incremented beyond its maximum value of 65 535. It is cleared to zero when the WIS is reset.

45.2.2.12 10G WIS J1 transmit (Registers 2.39 through 2.46)

The assignment of octets in the 10G WIS J1 transmit registers is shown in Table 45–108.

Table 45–108—10G WIS J1 transmit 0–15 register bit definitions

Bit(s)	Name	Description	R/W[a]
2.46.15:8	J1 transmit 15	Transmitted path trace octet 15	R/W
2.46.7:0	J1 transmit 14	Transmitted path trace octet 14	R/W
2.45.15:8	J1 transmit 13	Transmitted path trace octet 13	R/W
2.45.7:0	J1 transmit 12	Transmitted path trace octet 12	R/W
2.44.15:8	J1 transmit 11	Transmitted path trace octet 11	R/W
2.44.7:0	J1 transmit 10	Transmitted path trace octet 10	R/W
2.43.15:8	J1 transmit 9	Transmitted path trace octet 9	R/W
2.43.7:0	J1 transmit 8	Transmitted path trace octet 8	R/W
2.42.15:8	J1 transmit 7	Transmitted path trace octet 7	R/W
2.42.7:0	J1 transmit 6	Transmitted path trace octet 6	R/W
2.41.15:8	J1 transmit 5	Transmitted path trace octet 5	R/W
2.41.7:0	J1 transmit 4	Transmitted path trace octet 4	R/W
2.40.15:8	J1 transmit 3	Transmitted path trace octet 3	R/W

Table 45–108—10G WIS J1 transmit 0–15 register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
2.40.7:0	J1 transmit 2	Transmitted path trace octet 2	R/W
2.39.15:8	J1 transmit 1	Transmitted path trace octet 1	R/W
2.39.7:0	J1 transmit 0	Transmitted path trace octet 0	R/W

[a]R/W = Read/Write

The first transmitted path trace octet is J1 transmit 15, which contains the delineation octet. The default value for the J1 transmit 15 octet is 137 (hexadecimal 89). The last transmitted path trace octet is J1 transmit 0. The default value for the J1 transmit 0 through 14 octets is 0. The transmitted path trace is described in 50.3.2.1.

45.2.2.13 10G WIS J1 receive (Registers 2.47 through 2.54)

The assignment of octets in the 10G WIS J1 receive registers is shown in Table 45–109.

Table 45–109—10G WIS J1 receive 0–15 register bit definitions

Bit(s)	Name	Description	R/W[a]
2.54.15:8	J1 receive 15	Received path trace octet 15	RO
2.54.7:0	J1 receive 14	Received path trace octet 14	RO
2.53.15:8	J1 receive 13	Received path trace octet 13	RO
2.53.7:0	J1 receive 12	Received path trace octet 12	RO
2.52.15:8	J1 receive 11	Received path trace octet 11	RO
2.52.7:0	J1 receive 10	Received path trace octet 10	RO
2.51.15:8	J1 receive 9	Received path trace octet 9	RO
2.51.7:0	J1 receive 8	Received path trace octet 8	RO
2.50.15:8	J1 receive 7	Received path trace octet 7	RO
2.50.7:0	J1 receive 6	Received path trace octet 6	RO
2.49.15:8	J1 receive 5	Received path trace octet 5	RO
2.49.7:0	J1 receive 4	Received path trace octet 4	RO
2.48.15:8	J1 receive 3	Received path trace octet 3	RO

Table 45–109—10G WIS J1 receive 0–15 register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
2.48.7:0	J1 receive 2	Received path trace octet 2	RO
2.47.15:8	J1 receive 1	Received path trace octet 1	RO
2.47.7:0	J1 receive 0	Received path trace octet 0	RO

[a]RO = Read only

The first received path trace octet is J1 receive 15. The last received path trace octet is J1 receive 0. The received path trace is described in 50.3.2.4.

45.2.2.14 10G WIS far end line BIP errors (Registers 2.55 and 2.56)

The assignment of octets in the 10G WIS far end line BIP errors registers is shown in Table 45–110.

Table 45–110—10G WIS far end line BIP errors 0–1 register bit definitions

Bit(s)	Name	Description	R/W[a]
2.56.15:0	WIS far end line BIP errors 0	Least significant word of the WIS far end line BIP errors counter	RO
2.55.15:0	WIS far end line BIP errors 1	Most significant word of the WIS far end line BIP errors counter	RO

[a]RO = Read only

The 10G WIS far end line BIP Errors register pair reflects the contents of the far end line BIP errors counter (as described in 50.3.11.3) that is incremented on each WIS frame by the number of far end line BIP errors reported by the far end, as described in 50.3.2.5. Whenever the most significant 16 bit register of the counter (2.55) is read, the 32 bit counter value is latched into the register pair, with the most significant bits appearing in 2.55 and the least significant 16 bits appearing in 2.56, the value being latched before the contents of 2.55 (the most significant 16 bits) are driven on the MDIO interface. A subsequent read from register 2.56 will return the least significant 16 bits of the latched value, but will not change the register contents. Writes to these registers have no effect.

NOTE—These counters do not follow the behavior described in 45.2 for 32-bit counters.

45.2.2.15 10G WIS line BIP errors (Registers 2.57 and 2.58)

The assignment of octets in the 10G WIS line BIP errors registers is shown in Table 45–111.

The 10G WIS line BIP errors register pair reflects the contents of the line BIP errors counter (as described in 50.3.11.3) that is incremented on each WIS frame by the number of line BIP errors detected on the incoming data stream, as described in 50.3.2.5. Whenever the most significant 16 bit register of the counter (2.57) is read, the 32 bit counter value is latched into the register pair, with the most significant bits appearing in 2.57 and the least significant 16 bits appearing in 2.58, the value being latched before the contents of 2.57 (the most significant 16 bits) are driven on the MDIO interface. A subsequent read from register 2.58 will return the least significant 16 bits of the latched value, but will not change the register contents. Writes to these registers have no effect.

Table 45–111—10G WIS line BIP errors 0–1 register bit definitions

Bit(s)	Name	Description	R/W[a]
2.58.15:0	WIS line BIP errors 0	Least significant word of the WIS line BIP errors counter	RO
2.57.15:0	WIS line BIP errors 1	Most significant word of the WIS line BIP errors counter	RO

[a]RO = Read only

NOTE—These counters do not follow the behavior described in 45.2 for 32-bit counters.

45.2.2.16 10G WIS path block error count (Register 2.59)

The assignment of bits in the 10G WIS path block error count register is shown in Table 45–112.

Table 45–112—10G WIS path block error count register bit definitions

Bit(s)	Name	Description	R/W[a]
2.59.15:0	Path block error count	Path block error counter	RO

[a]RO = Read only

45.2.2.16.1 Path block error count (2.59.15:0)

The path block error count is incremented by one whenever a B3 parity error (defined in Annex 50A) is detected, as described in 50.3.2.5. The counter wraps around to zero when it is incremented beyond its maximum value of 65 535. It is cleared to zero when the WIS is reset.

45.2.2.17 10G WIS section BIP error count (Register 2.60)

The assignment of bits in the 10G WIS section BIP error count register is shown in Table 45–113.

Table 45–113—10G WIS section BIP error count register bit definitions

Bit(s)	Name	Description	R/W[a]
2.60.15:0	Section BIP error count	Section BIP error count	RO

[a]RO = Read only

45.2.2.17.1 Section BIP error count (2.60.15:0)

The section BIP error count is incremented by the number of section BIP errors detected within each WIS frame, as described in 50.3.2.5. The counter wraps around to zero when it is incremented beyond its maximum value of 65 535. It is cleared to zero when the WIS is reset.

45.2.2.18 10G WIS J0 transmit (Registers 2.64 through 2.71)

The assignment of octets in the 10G WIS J0 transmit registers is shown in Table 45–114.

Table 45–114—10G WIS J0 transmit 0–15 register bit definitions

Bit(s)	Name	Description	R/W[a]
2.71.15:8	J0 transmit 15	Transmitted section trace octet 15	R/W
2.71.7:0	J0 transmit 14	Transmitted section trace octet 14	R/W
2.70.15:8	J0 transmit 13	Transmitted section trace octet 13	R/W
2.70.7:0	J0 transmit 12	Transmitted section trace octet 12	R/W
2.69.15:8	J0 transmit 11	Transmitted section trace octet 11	R/W
2.69.7:0	J0 transmit 10	Transmitted section trace octet 10	R/W
2.68.15:8	J0 transmit 9	Transmitted section trace octet 9	R/W
2.68.7:0	J0 transmit 8	Transmitted section trace octet 8	R/W
2.67.15:8	J0 transmit 7	Transmitted section trace octet 7	R/W
2.67.7:0	J0 transmit 6	Transmitted section trace octet 6	R/W
2.66.15:8	J0 transmit 5	Transmitted section trace octet 5	R/W
2.66.7:0	J0 transmit 4	Transmitted section trace octet 4	R/W
2.65.15:8	J0 transmit 3	Transmitted section trace octet 3	R/W
2.65.7:0	J0 transmit 2	Transmitted section trace octet 2	R/W
2.64.15:8	J0 transmit 1	Transmitted section trace octet 1	R/W
2.64.7:0	J0 transmit 0	Transmitted section trace octet 0	R/W

[a]R/W = Read/Write

The J0 transmit octets allow a receiver to verify its continued connection to the WIS transmitter. The first transmitted section trace octet is J0 transmit 15, which contains the delineation octet. The default value for the J0 transmit 15 octet is 137 (hexadecimal 89). The last transmitted section trace octet is J0 transmit 0. The default value for the J0 transmit 0 through 14 octets is 0. The transmitted section trace is described in 50.3.2.3.

45.2.2.19 10G WIS J0 receive (Registers 2.72 through 2.79)

The assignment of octets in the 10G WIS J0 receive registers is shown in Table 45–115.

Table 45–115— 10G WIS J0 receive 0–15 register bit definitions

Bit(s)	Name	Description	R/W[a]
2.79.15:8	J0 receive 15	Received section trace octet 15	RO
2.79.7:0	J0 receive 14	Received section trace octet 14	RO
2.78.15:8	J0 receive 13	Received section trace octet 13	RO
2.78.7:0	J0 receive 12	Received section trace octet 12	RO
2.77.15:8	J0 receive 11	Received section trace octet 11	RO
2.77.7:0	J0 receive 10	Received section trace octet 10	RO
2.76.15:8	J0 receive 9	Received section trace octet 9	RO
2.76.7:0	J0 receive 8	Received section trace octet 8	RO
2.75.15:8	J0 receive 7	Received section trace octet 7	RO
2.75.7:0	J0 receive 6	Received section trace octet 6	RO
2.74.15:8	J0 receive 5	Received section trace octet 5	RO
2.74.7:0	J0 receive 4	Received section trace octet 4	RO
2.73.15:8	J0 receive 3	Received section trace octet 3	RO
2.73.7:0	J0 receive 2	Received section trace octet 2	RO
2.72.15:8	J0 receive 1	Received section trace octet 1	RO
2.72.7:0	J0 receive 0	Received section trace octet 0	RO

[a]RO = Read only

The first received section trace octet is J0 receive 15. The last received section trace octet is J0 receive 0. The J0 receive octets allow a WIS receiver to verify its continued connection to the intended transmitter. The received section trace is described in 50.3.2.4.

45.2.2.20 TimeSync WIS capability (Register 2.1800)

The TimeSync WIS capability register (see Table 45–116) indicates the capability of the WIS to report the transmit and receive data delay, stored in registers 2.1801 through 2.1804 and 2.1805 through 2.1808, respectively.

Table 45–116—TimeSync WIS capability

Bit(s)	Name	Description	R/W[a]
2.1800.15:2	Reserved	Value always 0	RO
2.1800.1	TimeSync transmit path data delay	1 = WIS provides information on transmit path data delay in registers 2.1801 through 2.1804 0 = WIS does not provide information on transmit path data delay	RO
2.1800.0	TimeSync receive path data delay	1 = WIS provides information on receive path data delay in registers 2.1805 through 2.1808 0 = WIS does not provide information on receive path data delay	RO

[a]RO = Read only

45.2.2.21 TimeSync WIS transmit path data delay (Registers 2.1801, 2.1802, 2.1803, 2.1804)

The TimeSync WIS transmit path data delay register contains the maximum (Registers 2.1801, 2.1802, see Table 45–117) and minimum (Registers 2.1803, 2.1804, see Table 45–117) values of the transmit path data delay. The transmit path data delay is expressed in units of ns. The values contained in these registers are valid when the link is established, as indicated by bit 2 in Register 1.1 (see 45.2.1.2.4).

Table 45–117—TimeSync WIS transmit path data delay register

Bit(s)	Name	Description	R/W[a]
2.1801.15:0	Maximum WIS transmit path data delay, lower	WIS_delay_TX_max [15:0]	RO, MW
2.1802.15:0	Maximum WIS transmit path data delay, upper	WIS_delay_TX_max [31:16]	RO, MW
2.1803.15:0	Minimum WIS transmit path data delay, lower	WIS_delay_TX_min [15:0]	RO, MW
2.1804.15:0	Minimum WIS transmit path data delay, upper	WIS_delay_TX_min [31:16]	RO, MW

[a]RO = Read only, MW = Multi-word

45.2.2.22 TimeSync WIS receive path data delay (Registers 2.1805, 2.1806, 2.1807, 2.1808)

The TimeSync WIS receive path data delay register contains the maximum (Registers 2.1805, 2.1806, see Table 45–118) and minimum (Registers 2.1807, 2.1808, see Table 45–118) values of the receive path data delay. The receive path data delay is expressed in units of ns. The values contained in these registers are valid when the link is established, as indicated by bit 2 in Register 1.1 (see 45.2.1.2.4).

Table 45–118—TimeSync WIS receive path data delay register

Bit(s)	Name	Description	R/W[a]
2.1805.15:0	Maximum WIS receive path data delay, lower	WIS_delay_RX_max [15:0]	RO, MW
2.1806.15:0	Maximum WIS receive path data delay, upper	WIS_delay_RX_max [31:16]	RO, MW
2.1807.15:0	Minimum WIS receive path data delay, lower	WIS_delay_RX_min [15:0]	RO, MW
2.1808.15:0	Minimum WIS receive path data delay, upper	WIS_delay_RX_min [31:16]	RO, MW

[a]RO = Read only, MW = Multi-word

45.2.3 PCS registers

The assignment of registers in the PCS is shown in Table 45–119.

Table 45–119—PCS registers

Register address	Register name	Subclause
3.0	PCS control 1	45.2.3.1
3.1	PCS status 1	45.2.3.2
3.2, 3.3	PCS device identifier	45.2.3.3
3.4	PCS speed ability	45.2.3.4
3.5, 3.6	PCS devices in package	45.2.3.5
3.7	PCS control 2	45.2.3.6
3.8	PCS status 2	45.2.3.7
3.9 through 3.13	Reserved	
3.14, 3.15	PCS package identifier	45.2.3.8
3.16 through 3.19	Reserved	
3.20	EEE control and capability	45.2.3.9
3.21	Reserved	
3.22	EEE wake error counter	45.2.3.10
3.23	Reserved	
3.24	10GBASE-X PCS status	45.2.3.11
3.25	10GBASE-X PCS test control	45.2.3.12
3.26 through 3.31	Reserved	
3.32	BASE-R and 10GBASE-T PCS status 1	45.2.3.13
3.33	BASE-R and 10GBASE-T PCS status 2	45.2.3.14
3.34 through 3.37	10GBASE-R PCS test pattern seed A	45.2.3.15
3.38 through 3.41	10GBASE-R PCS test pattern seed B	45.2.3.16
3.42	BASE-R PCS test pattern control	45.2.3.17
3.43	BASE-R PCS test pattern error counter	45.2.3.18

Table 45–119—PCS registers *(continued)*

Register address	Register name	Subclause
3.44	BER high order counter	45.2.3.19
3.45	Errored blocks high order counter	45.2.3.20
3.46 through 3.49	Reserved	
3.50 through 3.53	Multi-lane BASE-R PCS alignment status 1 through 4	45.2.3.21
3.54 through 3.59	Reserved	
3.60	10P/2B capability	45.2.3.25
3.61	10P/2B PCS control	45.2.3.26
3.62, 3.63	10P/2B PME available	45.2.3.27
3.64, 3.65	10P/2B PME aggregate	45.2.3.28
3.66	10P/2B PAF RX error counter	45.2.3.29
3.67	10P/2B PAF small fragment counter	45.2.3.30
3.68	10P/2B PAF large fragments counter	45.2.3.31
3.69	10P/2B PAF overflow counter	45.2.3.32
3.70	10P/2B PAF bad fragments counter	45.2.3.33
3.71	10P/2B PAF lost fragments counter	45.2.3.34
3.72	10P/2B PAF lost starts of fragments counter	45.2.3.35
3.73	10P/2B PAF lost ends of fragments counter	45.2.3.36
3.74	10GBASE-PR and 10/1GBASE-PRX FEC ability	45.2.3.37
3.75	10GBASE-PR and 10/1GBASE-PRX FEC control	45.2.3.38
3.76, 3.77	10/1GBASE-PRX and 10GBASE-PR corrected FEC codewords counter	45.2.3.39
3.78, 3.79	10/1GBASE-PRX and 10GBASE-PR uncorrected FEC codewords counter	45.2.3.40
3.80	10GBASE-PR and 10/1GBASE-PRX BER monitor timer control	45.2.3.41
3.81	10GBASE-PR and 10/1GBASE-PRX BER monitor status	45.2.3.42
3.82	10GBASE-PR and 10/1GBASE-PRX BER monitor threshold control	45.2.3.43
3.83 through 3.199	Reserved	
3.200 through 3.219	BIP error counters, lanes 0 through 19	45.2.3.44
3.220 through 3.399	Reserved	
3.400 through 3.419	PCS lane mapping registers, lanes 0 through 19	45.2.3.46
3.420 through 3.1799	Reserved	
3.1800	TimeSync PCS capability	45.2.3.48
3.1801 through 3.1804	TimeSync PCS transmit path data delay	45.2.3.49
3.1805 through 3.1808	TimeSync PCS receive path data delay	45.2.3.50
3.1809 through 3.32767	Reserved	
3.32 768 through 3.65 535	Vendor specific	

45.2.3.1 PCS control 1 register (Register 3.0)

The assignment of bits in the PCS control 1 register is shown in Table 45–120. The default value for each bit of the PCS control 1 register should be chosen so that the initial state of the device upon power up or reset is a normal operational state without management intervention.

Table 45–120—PCS control 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.0.15	Reset	1 = PCS reset 0 = Normal operation	R/W SC
3.0.14	Loopback	1 = Enable loopback mode 0 = Disable loopback mode	R/W
3.0.13	Speed selection	13 6 1 1 = bits 5:2 select speed 0 x = unspecified x 0 = unspecified	R/W
3.0.12	Reserved	Value always 0	RO
3.0.11	Low power	1 = Low-power mode 0 = Normal operation	R/W
3.0.10	Clock stop enable	1 = The PHY may stop the clock during LPI 0 = Clock not stoppable	R/W
3.0.9:7	Reserved	Value always 0	RO
3.0.6	Speed selection	13 6 1 1 = bits 5:2 select speed 0 x = unspecified x 0 = unspecified	R/W
3.0.5:2	Speed selection	5 4 3 2 1 x x x = Reserved 0 1 1 x = Reserved 0 1 0 1 = Reserved 0 1 0 0 = 100 Gb/s 0 0 1 1 = 40 Gb/s 0 0 1 0 = 10/1 Gb/s 0 0 0 1 = 10PASS-TS/2BASE-TL 0 0 0 0 = 10 Gb/s	R/W
3.0.1:0	Reserved	Value always 0	RO

[a]RO = Read only, R/W = Read/Write, SC = Self-clearing

45.2.3.1.1 Reset (3.0.15)

Resetting a PCS is accomplished by setting bit 3.0.15 to a one. This action shall set all PCS registers to their default states. As a consequence, this action may change the internal state of the PCS and the state of the physical link. This action may also initiate a reset in any other MMDs that are instantiated in the same package. This bit is self-clearing, and a PCS shall return a value of one in bit 3.0.15 when a reset is in progress and a value of zero otherwise. A PCS is not required to accept a write transaction to any of its

registers until the reset process is completed. The reset process shall be completed within 0.5 s from the setting of bit 3.0.15. During a reset, a PCS shall respond to reads from register bits 3.0.15 and 3.8.15:14. All other register bits should be ignored.

NOTE—This operation may interrupt data communication.

45.2.3.1.2 Loopback (3.0.14)

When the 10GBASE-T or the 10GBASE-R mode of operation is selected for the PCS using the PCS type selection field (3.7.2:0), the PCS shall be placed in a loopback mode of operation when bit 3.0.14 is set to a one. When bit 3.0.14 is set to a one, the 10GBASE-R or 10GBASE-T PCS shall accept data on the transmit path and return it on the receive path. The specific behavior of the 10GBASE-R PCS during loopback is specified in 49.2. The specific behavior for the 10GBASE-T PCS during loopback is specified in 55.3.7.3. For all other port types, the PCS loopback functionality is not applicable and writes to this bit shall be ignored and reads from this bit shall return a value of zero.

The default value of bit 3.0.14 is zero.

NOTE—The signal path through the PCS that is exercised in the loopback mode of operation is implementation specific, but it is recommended that the signal path encompass as much of the PCS circuitry as is practical. The intention of providing this loopback mode of operation is to permit a diagnostic or self-test function to perform the transmission and reception of a PDU, thus testing the transmit and receive data paths. Other loopback signal paths may be enabled using loopback controls within other MMDs.

45.2.3.1.3 Low power (3.0.11)

A PCS may be placed into a low-power mode by setting bit 3.0.11 to a one. This action may also initiate a low-power mode in any other MMDs that are instantiated in the same package. The low-power mode is exited by resetting the PCS. The behavior of the PCS in transition to and from the low-power mode is implementation specific and any interface signals should not be relied upon. While in the low-power mode, the device shall, as a minimum, respond to management transactions necessary to exit the low-power mode. The default value of bit 3.0.11 is zero.

45.2.3.1.4 Clock stop enable (3.0.10)

If bit 3.0.10 is set to 1 then the PHY may stop the receive xMII clock while it is signaling LPI otherwise it shall keep the clock active. If the PHY does not support EEE capability or is not able to stop the receive clock then this bit has no effect (see 22.2.2.2, 35.2.2.10, and 46.3.2.4).

45.2.3.1.5 Speed selection (3.0.13, 3.0.6, 3.0.5:2)

Speed selection bits 3.0.13 and 3.0.6 shall both be written as a one. Any attempt to change the bits to an invalid setting shall be ignored. These two bits are set to one in order to make them compatible with Clause 22.

The speed of the PCS may be selected using bits 5 through 2. The speed abilities of the PCS are advertised in the PCS speed ability register. A PCS may ignore writes to the PCS speed selection bits that select speeds it has not advertised in the PCS speed ability register. It is the responsibility of the STA entity to ensure that mutually acceptable speeds are applied consistently across all the MMDs on a particular PHY.

The PCS speed selection defaults to a supported ability.

The speed selection bits 3.0.5:2, when set to 0001, select the use of the 10PASS-TS and 2BASE-TL PCS.

45.2.3.2 PCS status 1 register (Register 3.1)

The assignment of bits in the PCS status 1 register is shown in Table 45–121. All the bits in the PCS status 1 register are read only; a write to the PCS status 1 register shall have no effect.

Table 45–121—PCS status 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.1.15:12	Reserved	Value always 0	RO
3.1.11	Tx LPI received	1 = Tx PCS has received LPI 0 = LPI not received	RO/LH
3.1.10	Rx LPI received	1 = Rx PCS has received LPI 0 = LPI not received	RO/LH
3.1.9	Tx LPI indication	1 = Tx PCS is currently receiving LPI 0 = PCS is not currently receiving LPI	RO
3.1.8	Rx LPI indication	1 = Rx PCS is currently receiving LPI 0 = PCS is not currently receiving LPI	RO
3.1.7	Fault	1 = Fault condition detected 0 = No fault condition detected	RO
3.1.6	Clock stop capable	1 = The MAC may stop the clock during LPI 0 = Clock not stoppable	RO
3.1.5:3	Reserved	Value always 0	RO
3.1.2	PCS receive link status	1 = PCS receive link up 0 = PCS receive link down	RO/LL
3.1.1	Low-power ability	1 = PCS supports low-power mode 0 = PCS does not support low-power mode	RO
3.1.0	Reserved	Value always 0	RO

[a]RO = Read only, LL = Latching low, LH = Latching high

45.2.3.2.1 Transmit LPI received (3.1.11)

When read as a one, bit 3.1.11 indicates that the transmit PCS has received LPI signaling one or more times since the register was last read. When read as a zero, bit 3.1.11 indicates that the PCS has not received LPI signaling. This bit shall be implemented with latching high behavior.

45.2.3.2.2 Receive LPI received (3.1.10)

When read as a one, bit 3.1.10 indicates that the receive PCS has received LPI signaling one or more times since the register was last read. When read as a zero, bit 3.1.10 indicates that the PCS has not received LPI signaling. This bit shall be implemented with latching high behavior.

45.2.3.2.3 Transmit LPI indication (3.1.9)

When read as a one, bit 3.1.9 indicates that the transmit PCS is currently receiving LPI signals. When read as a zero, bit 3.1.9 indicates that the PCS is not currently receiving LPI signals. The behavior if read during a state transition is undefined.

45.2.3.2.4 Receive LPI indication (3.1.8)

When read as a one, bit 3.1.8 indicates that the receive PCS is currently receiving LPI signals. When read as a zero, bit 3.1.8 indicates that the PCS is not currently receiving LPI signals. The behavior if read during a state transition is undefined.

45.2.3.2.5 Fault (3.1.7)

When read as a one, bit 3.1.7 indicates that the PCS has detected a fault condition on either the transmit or receive paths. When read as a zero, bit 3.1.7 indicates that the PCS has not detected a fault condition. For 10 Gb/s operation, bit 3.1.7 is read as a one when either of the fault bits (3.8.11, 3.8.10) located in register 3.8 are read as a one. For 10BASE-TS or 2BASE-TL operation, this bit shall become a one when any 10P/2B PCS registers indicate a fault (see 45.2.3.29 through 45.2.3.36).

45.2.3.2.6 Clock stop capable (3.1.6)

If bit 3.1.6 is set to one then the RS may stop the transmit xMII clock while it is signaling LPI otherwise it shall keep the clock active. If the RS does not support EEE capability or is not able to stop the transmit clock then this bit has no effect (see 22.2.2.6, 35.2.2.6, and 46.3.2.4).

45.2.3.2.7 PCS receive link status (3.1.2)

When read as a one, bit 3.1.2 indicates that the PCS receive link is up. When read as a zero, bit 3.1.2 indicates that the PCS receive link is down. When a 10/40/100GBASE-R, 10GBASE-W, or 10GBASE-T mode of operation is selected for the PCS using the PCS type selection field (3.7.2:0), this bit is a latching low version of bit 3.32.12. When a 10GBASE-X mode of operation is selected for the PCS using the PCS type selection field (3.7.2:0), this bit is a latching low version of bit 3.24.12. The receive link status bit shall be implemented with latching low behavior.

45.2.3.2.8 Low-power ability (3.1.1)

When read as a one, bit 3.1.1 indicates that the PCS supports the low-power feature. When read as a zero, bit 3.1.1 indicates that the PCS does not support the low-power feature. If a PCS supports the low-power feature then it is controlled using the low-power bit 3.0.11.

45.2.3.3 PCS device identifier (Registers 3.2 and 3.3)

Registers 3.2 and 3.3 provide a 32-bit value, which may constitute a unique identifier for a particular type of PCS. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the device manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A PCS may return a value of zero in each of the 32 bits of the PCS device identifier.

The format of the PCS device identifier is specified in 22.2.4.3.1.

45.2.3.4 PCS speed ability (Register 3.4)

The assignment of bits in the PCS speed ability register is shown in Table 45–122.

Table 45–122—PCS speed ability register bit definitions

Bit(s)	Name	Description	R/W[a]
3.4.15:4	Reserved for future speeds	Value always 0	RO
3.4.3	100G capable	1 = PCS is capable of operating at 100 Gb/s 0 = PCS is not capable of operating at 100 Gb/s	RO
3.4.2	40G capable	1 = PCS is capable of operating at 40 Gb/s 0 = PCS is not capable of operating at 40 Gb/s	RO
3.4.1	10PASS-TS/2BASE-TL capable	1 = PCS is capable of operating as the 10P/2B PCS 0 = PCS is not capable of operating as the 10P/2B PCS	RO
3.4.0	10G capable	1 = PCS is capable of operating at 10 Gb/s 0 = PCS is not capable of operating at 10 Gb/s	RO

[a]RO = Read only

45.2.3.4.1 10G capable (3.4.0)

When read as a one, bit 3.4.0 indicates that the PCS is able to operate at a data rate of 10 Gb/s. When read as a zero, bit 3.4.0 indicates that the PCS is not able to operate at a data rate of 10 Gb/s.

45.2.3.4.2 10PASS-TS/2BASE-TL capable

When read as a one, this bit indicates that the PCS is able to operate as the 10PASS-TS/2BASE-TL PCS, as specified in Clause 61.

45.2.3.4.3 40G capable (3.4.2)

When read as a one, bit 3.4.2 indicates that the PCS is able to operate at a data rate of 40 Gb/s. When read as a zero, bit 3.4.2 indicates that the PCS is not able to operate at a data rate of 40 Gb/s.

45.2.3.4.4 100G capable (3.4.3)

When read as a one, bit 3.4.3 indicates that the PCS is able to operate at a data rate of 100 Gb/s. When read as a zero, bit 3.4.3 indicates that the PCS is not able to operate at a data rate of 100 Gb/s.

45.2.3.5 PCS devices in package (Registers 3.5 and 3.6)

The PCS devices in package registers are defined in Table 45–2.

45.2.3.6 PCS control 2 register (Register 3.7)

The assignment of bits in the PCS control 2 register is shown in Table 45–123. The default value for each bit of the PCS control 2 register should be chosen so that the initial state of the device upon power up or reset is a normal operational state without management intervention.

45.2.3.6.1 PCS type selection (3.7.2:0)

The PCS type shall be selected using bits 2 through 0. The PCS type abilities of the PCS are advertised in bits 3.8.5:0. A PCS shall ignore writes to the PCS type selection bits that select PCS types it has not advertised in the PCS status 2 register. It is the responsibility of the STA entity to ensure that mutually

Table 45–123—PCS control 2 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.7.15:3	Reserved	Value always 0	RO
3.7.2:0	PCS type selection	2 1 0 1 1 x = reserved 1 0 1 = Select 100GBASE-R PCS type 1 0 0 = Select 40GBASE-R PCS type 0 1 1 = Select 10GBASE-T PCS type 0 1 0 = Select 10GBASE-W PCS type 0 0 1 = Select 10GBASE-X PCS type 0 0 0 = Select 10GBASE-R PCS type	R/W

[a]RO = Read only, R/W = Read/Write

acceptable MMD types are applied consistently across all the MMDs on a particular PHY. The PCS type selection defaults to a supported ability.

45.2.3.7 PCS status 2 register (Register 3.8)

The assignment of bits in the PCS status 2 register is shown in Table 45–124. All the bits in the PCS status 2 register are read only; a write to the PCS status 2 register shall have no effect.

Table 45–124—PCS status 2 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.8.15:14	Device present	15 14 1 0 = Device responding at this address 1 1 = No device responding at this address 0 1 = No device responding at this address 0 0 = No device responding at this address	RO
3.8.13:12	Reserved	Value always 0	RO
3.8.11	Transmit fault	1 = Fault condition on transmit path 0 = No fault condition on transmit path	RO/LH
3.8.10	Receive fault	1 = Fault condition on the receive path 0 = No fault condition on the receive path	RO/LH
3.8.9:6	Reserved	Value always 0	RO
3.8.5	100GBASE-R capable	1 = PCS is able to support 100GBASE-R PCS type 0 = PCS is not able to support 100GBASE-R PCS type	RO
3.8.4	40GBASE-R capable	1 = PCS is able to support 40GBASE-R PCS type 0 = PCS is not able to support 40GBASE-R PCS type	RO
3.8.3	10GBASE-T capable	1 = PCS is able to support 10GBASE-T PCS type 0 = PCS is not able to support 10GBASE-T PCS type	RO
3.8.2	10GBASE-W capable	1 = PCS is able to support 10GBASE-W PCS type 0 = PCS is not able to support 10GBASE-W PCS type	RO

Table 45–124—PCS status 2 register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
3.8.1	10GBASE-X capable	1 = PCS is able to support 10GBASE-X PCS type 0 = PCS is not able to support 10GBASE-X PCS type	RO
3.8.0	10GBASE-R capable	1 = PCS is able to support 10GBASE-R PCS types 0 = PCS is not able to support 10GBASE-R PCS types	RO

[a]RO = Read only, LH = Latching high

45.2.3.7.1 Device present (3.8.15:14)

When read as <10>, bits 3.8.15:14 indicate that a device is present and responding at this register address. When read as anything other than <10>, bits 3.8.15:14 indicate that no device is present at this address or that the device is not functioning properly.

45.2.3.7.2 Transmit fault (3.8.11)

When read as a one, bit 3.8.11 indicates that the PCS has detected a fault condition on the transmit path. When read as a zero, bit 3.8.11 indicates that the PCS has not detected a fault condition on the transmit path. The transmit fault bit shall be implemented with latching high behavior.

The default value of bit 3.8.11 is zero.

45.2.3.7.3 Receive fault (3.8.10)

When read as a one, bit 3.8.10 indicates that the PCS has detected a fault condition on the receive path. When read as a zero, bit 3.8.10 indicates that the PCS has not detected a fault condition on the receive path. The receive fault bit shall be implemented with latching high behavior.

The default value of bit 3.8.10 is zero.

45.2.3.7.4 100GBASE-R capable (3.8.5)

When read as a one, bit 3.8.5 indicates that the PCS is able to support the 100GBASE-R PCS type. When read as a zero, bit 3.8.5 indicates that the PCS is not able to support the 100GBASE-R PCS type.

45.2.3.7.5 40GBASE-R capable (3.8.4)

When read as a one, bit 3.8.4 indicates that the PCS is able to support the 40GBASE-R PCS type. When read as a zero, bit 3.8.4 indicates that the PCS is not able to support the 40GBASE-R PCS type.

45.2.3.7.6 10GBASE-T capable (3.8.3)

When read as a one, bit 3.8.3 indicates that the PCS is able to support the 10GBASE-T PCS type. When read as a zero, bit 3.8.3 indicates that the PCS is not able to support the 10GBASE-T PCS type.

45.2.3.7.7 10GBASE-W capable (3.8.2)

When read as a one, bit 3.8.2 indicates that the 64B/66B PCS is able to support operation in a 10GBASE-W PHY (that is, supports operation with a WIS). When read as a zero, bit 3.8.2 indicates that the 64B/66B PCS is not able to support operation with a WIS in a 10GBASE-W PHY.

NOTE—This bit does not indicate that the PCS is performing the functionality contained in the WIS. This bit indicates whether the 64B/66B PCS would be able to support a WIS if it were to be attached.

45.2.3.7.8 10GBASE-X capable (3.8.1)

When read as a one, bit 3.8.1 indicates that the PCS is able to support the 10GBASE-X PCS type. When read as a zero, bit 3.8.1 indicates that the PCS is not able to support the 10GBASE-X PCS type.

45.2.3.7.9 10GBASE-R capable (3.8.0)

When read as a one, bit 3.8.0 indicates that the PCS is able to support operation in a 10GBASE-R PHY. When read as a zero, bit 3.8.0 indicates that the PCS is not able to support operation in a 10GBASE-R PHY.

45.2.3.8 PCS package identifier (Registers 3.14 and 3.15)

Registers 3.14 and 3.15 provide a 32-bit value, which may constitute a unique identifier for a particular type of package that the PCS is instantiated within. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the package manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A PCS may return a value of zero in each of the 32 bits of the PCS package identifier.

A non-zero package identifier may be returned by one or more MMDs in the same package. The package identifier may be the same as the device identifier.

The format of the package identifier is specified in 22.2.4.3.1.

45.2.3.9 EEE control and capability (Register 3.20)

This register is used to indicate the capability of the PCS to support EEE functions for each PHY type. The assignment of bits in the EEE capability register is shown in Table 45–125.

Table 45–125—EEE control and capability register bit definitions

Bit(s)	Name	Description	R/W[a]
3.20.15:14	Reserved	Value always 0	RO
3.20.13	100GBASE-R deep sleep	1 = EEE deep sleep is supported for 100GBASE-R 0 = EEE deep sleep is not supported for 100GBASE-R	RO
3.20.12	100GBASE-R fast wake	1 = EEE fast wake is supported for 100GBASE-R 0 = EEE fast wake is not supported for 100GBASE-R	RO
3.20.11:10	Reserved	Value always 0	RO
3.20.9	40GBASE-R deep sleep	1 = EEE deep sleep is supported for 40GBASE-R 0 = EEE deep sleep is not supported for 40GBASE-R	RO
3.20.8	40GBASE-R fast wake	1 = EEE fast wake is supported for 40GBASE-R 0 = EEE fast wake is not supported for 40GBASE-R	RO

Table 45–125—EEE control and capability register bit definitions *(continued)*

3.20.7	Reserved	Value always 0	RO
3.20.6	10GBASE-KR EEE	1 = EEE is supported for 10GBASE-KR 0 = EEE is not supported for 10GBASE-KR	RO
3.20.5	10GBASE-KX4 EEE	1 = EEE is supported for 10GBASE-KX4 0 = EEE is not supported for 10GBASE-KX4	RO
3.20.4	1000BASE-KX EEE	1 = EEE is supported for 1000BASE-KX 0 = EEE is not supported for 1000BASE-KX	RO
3.20.3	10GBASE-T EEE	1 = EEE is supported for 10GBASE-T 0 = EEE is not supported for 10GBASE-T	RO
3.20.2	1000BASE-T EEE	1 = EEE is supported for 1000BASE-T 0 = EEE is not supported for 1000BASE-T	RO
3.20.1	100BASE-TX EEE	1 = EEE is supported for 100BASE-TX 0 = EEE is not supported for 100BASE-TX	RO
3.20.0	LPI_FW	1 = Fast wake mode is used for LPI function 0 = Deep sleep is used for LPI function	R/W

[a]R/W = Read/Write, RO = Read only

45.2.3.9.1 100GBASE-R EEE deep sleep supported (3.20.13)

If the PCS supports EEE deep sleep operation for 100GBASE-R, this bit shall be set to a one; otherwise this bit shall be set to a zero.

45.2.3.9.2 100GBASE-R EEE fast wake supported (3.20.12)

If the PCS supports EEE fast wake operation for 100GBASE-R, this bit shall be set to a one; otherwise this bit shall be set to a zero.

45.2.3.9.3 40GBASE-R EEE deep sleep supported (3.20.9)

If the PCS supports EEE deep sleep operation for 40GBASE-R, as defined in 78.1, this bit shall be set to a one; otherwise this bit shall be set to a zero.

45.2.3.9.4 40GBASE-R EEE fast wake supported (3.20.8)

If the PCS supports EEE fast wake operation for 40GBASE-R, as defined in 78.1, this bit shall be set to a one; otherwise this bit shall be set to a zero.

45.2.3.9.5 10GBASE-KR EEE supported (3.20.6)

If the device supports EEE operation for 10GBASE-KR as defined in 72.1, this bit shall be set to one.

45.2.3.9.6 10GBASE-KX4 EEE supported (3.20.5)

If the device supports EEE operation for 10GBASE-KX4 as defined in 71.2, this bit shall be set to one.

45.2.3.9.7 1000BASE-KX EEE supported (3.20.4)

If the device supports EEE operation for 1000BASE-KX as defined in 70.1, this bit shall be set to one.

45.2.3.9.8 10GBASE-T EEE supported (3.20.3)

If the device supports EEE operation for 10GBASE-T as defined in 55.1.3.3, this bit shall be set to one.

45.2.3.9.9 1000BASE-T EEE supported (3.20.2)

If the device supports EEE operation for 1000BASE-T as defined in 40.1.3, this bit shall be set to one.

45.2.3.9.10 100BASE-TX EEE supported (3.20.1)

If the device supports EEE operation for 100BASE-TX as defined in 24.1.1, this bit shall be set to one.

45.2.3.9.11 LPI_FW (3.20.0)

If the device supports fast wake as defined in 78.5, this bit selects fast wake or deep sleep operation. Setting 3.20.0 to one selects fast wake, setting to zero selects deep sleep. This bit is ignored by devices that do not support fast wake, and this bit defaults to one for devices that support fast wake.

45.2.3.10 EEE wake error counter (Register 3.22)

This register is used by PHY types that support EEE to count wake time faults where the PHY fails to complete its normal wake sequence within the time required for the specific PHY type. The definition of the fault event to be counted is defined for each PHY and may occur during a refresh or a wake-up as defined by the PHY. This 16-bit counter shall be reset to all zeros when the EEE wake error counter is read by the management function or upon execution of the PCS reset. This counter shall be held at all ones in the case of overflow.

45.2.3.11 10GBASE-X PCS status register (Register 3.24)

The assignment of bits in the 10GBASE-X PCS status register is shown in Table 45–126. All the bits in the 10GBASE-X PCS status register are read only; a write to the 10GBASE-X PCS status register shall have no effect. A PCS device that does not implement 10GBASE-X shall return a zero for all bits in the 10GBASE-X PCS status register. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

Table 45–126—10GBASE-X PCS status register bit definitions

Bit(s)	Name	Description	R/W[a]
3.24.15:13	Reserved	Value always 0	RO
3.24.12	10GBASE-X lane alignment status	1 = 10GBASE-X PCS receive lanes aligned 0 = 10GBASE-X PCS receive lanes not aligned	RO
3.24.11	Pattern testing ability	1 = 10GBASE-X PCS is able to generate test patterns 0 = 10GBASE-X PCS is not able to generate test patterns	RO
3.24.10:4	Reserved	Value always 0	RO

Table 45–126—10GBASE-X PCS status register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
3.24.3	Lane 3 sync	1 = Lane 3 is synchronized 0 = Lane 3 is not synchronized	RO
3.24.2	Lane 2 sync	1 = Lane 2 is synchronized 0 = Lane 2 is not synchronized	RO
3.24.1	Lane 1 sync	1 = Lane 1 is synchronized 0 = Lane 1 is not synchronized	RO
3.24.0	Lane 0 sync	1 = Lane 0 is synchronized 0 = Lane 0 is not synchronized	RO

[a]RO = Read only

45.2.3.11.1 10GBASE-X receive lane alignment status (3.24.12)

When read as a one, bit 3.24.12 indicates that the 10GBASE-X PCS has synchronized and aligned all four receive lanes. When read as a zero, bit 3.24.12 indicates that the 10GBASE-X PCS has not synchronized and aligned all four receive lanes.

45.2.3.11.2 Pattern testing ability (3.24.11)

When read as a one, bit 3.24.11 indicates that the 10GBASE-X PCS is able to generate test patterns. When read as a zero, bit 3.24.11 indicates that the 10GBASE-X PCS is not able to generate test patterns. If the 10GBASE-X PCS is able to generate test patterns, then the functionality is controlled using the transmit test-pattern enable bit in register 3.25.

45.2.3.11.3 Lane 3 sync (3.24.3)

When read as a one, bit 3.24.3 indicates that the 10GBASE-X PCS receive lane 3 is synchronized. When read as a zero, bit 3.24.3 indicates that the 10GBASE-X PCS receive lane 3 is not synchronized.

45.2.3.11.4 Lane 2 sync (3.24.2)

When read as a one, bit 3.24.2 indicates that the 10GBASE-X PCS receive lane 2 is synchronized. When read as a zero, bit 3.24.2 indicates that the 10GBASE-X PCS receive lane 2 is not synchronized.

45.2.3.11.5 Lane 1 sync (3.24.1)

When read as a one, bit 3.24.1 indicates that the 10GBASE-X PCS receive lane 1 is synchronized. When read as a zero, bit 3.24.1 indicates that the 10GBASE-X PCS receive lane 1 is not synchronized.

45.2.3.11.6 Lane 0 sync (3.24.0)

When read as a one, bit 3.24.0 indicates that the 10GBASE-X PCS receive lane 0 is synchronized. When read as a zero, bit 3.24.0 indicates that the 10GBASE-X PCS receive lane 0 is not synchronized.

45.2.3.12 10GBASE-X PCS test control register (Register 3.25)

The assignment of bits in the 10GBASE-X PCS test control register is shown in Table 45–127. The default value for each bit of the 10GBASE-X PCS test control register should be chosen so that the initial state of the device upon power up or reset is a normal operational state without management intervention.

Table 45–127—10GBASE-X PCS test control register bit definitions

Bit(s)	Name	Description	R/W[a]
3.25.15:3	Reserved	Value always 0	RO
3.25.2	Transmit test-pattern enable	1 = Transmit test pattern enabled 0 = Transmit test pattern not enabled	R/W
3.25.1:0	Test pattern select	1 0 1 1 = Reserved 1 0 = Mixed-frequency test pattern 0 1 = Low-frequency test pattern 0 0 = High-frequency test pattern	R/W

[a]RO = Read only, R/W = Read/Write

45.2.3.12.1 Transmit test-pattern enable (3.25.2)

When bit 3.25.2 is set to a one, pattern testing is enabled on the transmit path. When bit 3.25.2 is set to a zero, pattern testing is disabled on the transmit path. Pattern testing is optional, and the ability of the 10GBASE-X PCS to generate test patterns is advertised by the pattern testing ability bit in register 3.24. A 10GBASE-X PCS that does not support the generation of test patterns shall ignore writes to this bit and always return a value of zero. The default of bit 3.25.2 is zero.

45.2.3.12.2 Test pattern select (3.25.1:0)

The test pattern to be used when pattern testing is enabled using bit 3.25.2 is selected using bits 3.25.1:0. When bits 3.25.1:0 are set to <10>, the mixed-frequency test pattern shall be selected for pattern testing. When bits 3.25.1:0 are set to <01>, the low-frequency test pattern shall be selected for pattern testing. When bits 3.25.1:0 are set to <00>, the high-frequency test pattern shall be selected for pattern testing. The test patterns are defined in Annex 48A.

45.2.3.13 BASE-R and 10GBASE-T PCS status 1 register (Register 3.32)

The assignment of bits in the BASE-R and 10GBASE-T PCS status 1 register is shown in Table 45–128. All the bits in the BASE-R and 10GBASE-T PCS status 1 register are read only; a write to the BASE-R and 10GBASE-T PCS status 1 register shall have no effect. A PCS device that does not implement BASE-R or the 10GBASE-T shall return a zero for all bits in the BASE-R and 10GBASE-T PCS status 1 register. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits. The contents of register 3.32 are undefined when the BASE-R PCS or the 10GBASE-T PCS is operating in seed test-pattern mode, PRBS31 test-pattern mode, or PRBS9 test-pattern mode.

Table 45–128—BASE-R and 10GBASE-T PCS status 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.32.15:13	Reserved	Value always 0	RO
3.32.12	BASE-R and 10GBASE-T receive link status	1 = BASE-R or 10GBASE-T PCS receive link up 0 = BASE-R or 10GBASE-T PCS receive link down	RO
3.32.11:4	Reserved	Value always 0	RO
3.32.3	10GBASE-R PRBS9 pattern testing ability	1 = PCS is able to support PRBS9 pattern testing 0 = PCS is not able to support PRBS9 pattern testing	RO
3.32.2	10GBASE-R PRBS31 pattern testing ability	1 = PCS is able to support PRBS31 pattern testing 0 = PCS is not able to support PRBS31 pattern testing	RO
3.32.1	BASE-R and 10GBASE-T PCS high BER	1 = BASE-R or 10GBASE-T PCS reporting a high BER 0 = BASE-R or 10GBASE-T PCS not reporting a high BER	RO
3.32.0	BASE-R and 10GBASE-T PCS block lock	1 = BASE-R or 10GBASE-T PCS locked to received blocks 0 = BASE-R or 10GBASE-T PCS not locked to received blocks	RO

[a]RO = Read only

45.2.3.13.1 BASE-R and 10GBASE-T receive link status (3.32.12)

When read as a one, bit 3.32.12 indicates that the PCS is in a fully operational state. When read as a zero, bit 3.32.12 indicates that the PCS is not fully operational. This bit is a reflection of the PCS_status variable defined in 49.2.14.1 for 10GBASE-R, in 55.3.6.1 for 10GBASE-T and in 82.3.1 for 40/100GBASE-R.

45.2.3.13.2 PRBS9 pattern testing ability (3.32.3)

When read as a one, bit 3.32.3 indicates that the PCS is able to support PRBS9 pattern testing of its transmitter. When read as a zero, bit 3.32.3 indicates that the PCS is not able to support PRBS9 pattern testing of its transmitter. If the PCS is able to support PRBS9 pattern testing of its transmitter then the pattern generation is controlled using bit 3.42.6.

45.2.3.13.3 PRBS31 pattern testing ability (3.32.2)

When read as a one, bit 3.32.2 indicates that the PCS is able to support PRBS31 pattern testing. When read as a zero, bit 3.32.2 indicates that the PCS is not able to support PRBS31 pattern testing. If the PCS is able to support PRBS31 pattern testing then the pattern generation and checking is controlled using bits 3.42.5:4.

45.2.3.13.4 BASE-R and 10GBASE-T PCS high BER (3.32.1)

For BASE-R, when read as a one, bit 3.32.1 indicates that the 64B/66B receiver is detecting a BER of $\geq 10^{-4}$. When read as a zero, bit 3.32.1 indicates that the 64B/66B receiver is detecting a BER of $< 10^{-4}$. This bit is a direct reflection of the state of the hi_ber variable in the 64B/66B state diagram and is defined in 49.2.13.2.2 for 10GBASE-R and in 82.2.19.2.2 for 40/100GBASE-R.

For 10GBASE-T, when read as a one, bit 3.32.1 indicates that the 64B/65B receiver is detecting a BER of $\geq 10^{-4}$. When read as a zero, bit 3.32.1 indicates that the 64B/65B receiver is detecting a BER of $< 10^{-4}$. This bit is a direct reflection of the state of the hi_lfer variable in the 64B/65B state diagram and is defined in 55.3.6.1.

45.2.3.13.5 BASE-R and 10GBASE-T PCS block lock (3.32.0)

When read as a one, bit 3.32.0 indicates that the 64B/66B receiver for BASE-R or the 64B/65B receiver for the 10GBASE-T has block lock. When read as a zero, bit 3.32.0 indicates that the 64B/66B receiver for BASE-R or the 64B/65B receiver for the 10GBASE-T has not achieved block lock. This bit is a direct reflection of the state of the block_lock variable in the 64B/66B state diagram and is defined in 49.2.13.2.2 for 10GBASE-R and in 82.2.19.2.2 for 40/100GBASE-R. For the 10GBASE-T PCS the block_lock variable in the 64B/65B state diagram is defined in 55.3.2.3. For a multi-lane PCS, this bit indicates that the receiver has both block lock and alignment for all lanes and is identical to 3.50.12 (see 45.2.3.21.1).

45.2.3.14 BASE-R and 10GBASE-T PCS status 2 register (Register 3.33)

The assignment of bits in the BASE-R and 10GBASE-T PCS status 2 register is shown in Table 45–129. All the bits in the BASE-R and 10GBASE-T PCS status 2 register are read only; a write to the BASE-R and 10GBASE-T PCS status 2 register shall have no effect. A PCS device that does not implement BASE-R or 10GBASE-T shall return a zero for all bits in the BASE-R and 10GBASE-T PCS status 2 register. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits. The contents of register 3.33 are undefined when the BASE-R or the 10GBASE-T PCS is operating seed test-pattern mode, PRBS31 test-pattern mode, or PRBS9 test-pattern mode.

Table 45–129—BASE-R and 10GBASE-T PCS status 2 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.33.15	Latched block lock	1 = BASE-R or 10GBASE-T PCS has block lock 0 = BASE-R or 10GBASE-T PCS does not have block lock	RO/LL
3.33.14	Latched high BER	1 = BASE-R or 10GBASE-T PCS has reported a high BER 0 = BASE-R or 10GBASE-T PCS has not reported a high BER	RO/LH
3.33.13:8	BER	BER counter	RO/NR
3.33.7:0	Errored blocks	Errored blocks counter	RO/NR

[a]RO = Read only, LL = Latching low, LH = Latching high, NR = Non Roll-over

45.2.3.14.1 Latched block lock (3.33.15)

When read as a one, bit 3.33.15 indicates that the 10/40/100GBASE-R or the 10GBASE-T PCS has achieved block lock. When read as a zero, bit 3.33.15 indicates that the 10/40/100GBASE-R or the 10GBASE-T PCS has lost block lock.

The latched block lock bit shall be implemented with latching low behavior.

This bit is a latching low version of the 10/40/100GBASE-R and 10GBASE-T PCS block lock status bit (3.32.0).

45.2.3.14.2 Latched high BER (3.33.14)

When read as a one, bit 3.33.14 indicates that the 10/40/100GBASE-R or the 10GBASE-T PCS has detected a high BER. When read as a zero, bit 3.33.14 indicates that the 10/40/100GBASE-R or the 10GBASE-T PCS has not detected a high BER.

The latched high BER bit shall be implemented with latching high behavior.

This bit is a latching high version of the 10/40/100GBASE-R and 10GBASE-T PCS high BER status bit (3.32.1).

45.2.3.14.3 BER (3.33.13:8)

The BER counter is a six bit count as defined by the ber_count variable in 49.2.14.2 and 82.2.19.2.4 for 10/40/100GBASE-R and defined by the lfer_count variable in 55.3.6.2 for 10GBASE-T. These bits shall be reset to all zeros when the BASE-R and 10GBASE-T PCS status 2 register is read by the management function or upon execution of the PCS reset. If the BER high order counter, 3.44 (see 45.2.3.19) is not implemented then these bits shall be held at all ones in the case of overflow.

45.2.3.14.4 Errored blocks (3.33.7:0)

The errored blocks counter is an eight bit count defined by the errored_block_count counter specified in 49.2.14.2 for 10GBASE-R, in 82.3.1 for 40/100GBASE-R and defined by the errored_block_count variable in 55.3.6.2 for 10GBASE-T. These bits shall be reset to all zeros when the errored blocks count is read by the management function or upon execution of the PCS reset. If the Errored blocks high order counter, 3.45 (see 45.2.3.20) is not implemented then these bits shall be held at all ones in the case of overflow.

45.2.3.15 10GBASE-R PCS test pattern seed A (Registers 3.34 through 3.37)

The assignment of bits in the 10GBASE-R PCS test pattern seed A registers is shown in Table 45–130. This register is only required when the 10GBASE-R capability is supported. If both 10GBASE-R and 10GBASE-W capability is supported, then this register may either ignore writes and return zeros for reads when in 10GBASE-W mode or may function as defined for 10GBASE-R. For each seed register, seed bits are assigned to register bits in order with the lowest numbered seed bit for that register being assigned to register bit 0.

Table 45–130—10GBASE-R PCS test pattern seed A 0-3 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.37.15:10	Reserved	Value always 0	RO
3.37.9:0	Test pattern seed A 3	Test pattern seed A bits 48-57	R/W
3.36.15:0	Test pattern seed A 2	Test pattern seed A bits 32-47	R/W
3.35.15:0	Test pattern seed A 1	Test pattern seed A bits 16-31	R/W
3.34.15:0	Test pattern seed A 0	Test pattern seed A bits 0-15	R/W

[a]RO = Read only, R/W = Read/Write

The A seed for the pseudo random test pattern is held in registers 3.34 through 3.37. The test-pattern methodology is described in 49.2.8.

45.2.3.16 10GBASE-R PCS test pattern seed B (Registers 3.38 through 3.41)

The assignment of bits in the 10GBASE-R PCS test pattern seed B registers is shown in Table 45–131. This register is only required when the 10GBASE-R capability is supported. If both 10GBASE-R and 10GBASE-W capability is supported, then this register may either ignore writes and return zeros for reads when in 10GBASE-W mode or may function as defined for 10GBASE-R. For each seed register, seed bits are assigned to register bits in order with the lowest numbered seed bit for that register being assigned to register bit 0.

Table 45–131—10GBASE-R PCS test pattern seed B 0-3 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.41.15:10	Reserved	Value always 0	RO
3.41.9:0	Test pattern seed B 3	Test pattern seed B bits 48-57	R/W
3.40.15:0	Test pattern seed B 2	Test pattern seed B bits 32-47	R/W
3.39.15:0	Test pattern seed B 1	Test pattern seed B bits 16-31	R/W
3.38.15:0	Test pattern seed B 0	Test pattern seed B bits 0-15	R/W

[a]RO = Read only, R/W = Read/Write

The B seed for the pseudo random test pattern is held in registers 3.38 through 3.41. The test-pattern methodology is described in 49.2.8.

45.2.3.17 BASE-R PCS test-pattern control register (Register 3.42)

The assignment of bits in the BASE-R PCS test-pattern control register is shown in Table 45–132. This register is only required when the BASE-R capability is supported. If both BASE-R and 10GBASE-W capability is supported, then this register may either ignore writes and return zeros for reads when in 10GBASE-W mode or may function as defined for BASE-R. PRBS9, PRBS31, pseudo random, and square wave test patterns are defined for 10GBASE-R PCS only. Scrambled idle test patterns are defined for 40/100GBASE-R PCS only. The test-pattern methodology is described in 49.2.8 and 82.2.11.

Table 45–132—BASE-R PCS test-pattern control register bit definitions

Bit(s)	Name	Description	R/W[a]
3.42.15:8	Reserved	Value always 0	RO
3.42.7	Scrambled idle test-pattern enable	1 = Enable scrambled idle test-pattern mode 0 = Disable scrambled idle test-pattern mode	R/W
3.42.6	10GBASE-R PRBS9 transmit test-pattern enable	1 = Enable PRBS9 test-pattern mode on the transmit path 0 = Disable PRBS9 test-pattern mode on the transmit path	R/W

Table 45–132—BASE-R PCS test-pattern control register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
3.42.5	10GBASE-R PRBS31 receive test-pattern enable	1 = Enable PRBS31 test-pattern mode on the receive path 0 = Disable PRBS31 test-pattern mode on the receive path	R/W
3.42.4	10GBASE-R PRBS31 transmit test-pattern enable	1 = Enable PRBS31 test-pattern mode on the transmit path 0 = Disable PRBS31 test-pattern mode on the transmit path	R/W
3.42.3	Transmit test-pattern enable	1 = Enable transmit test pattern 0 = Disable transmit test pattern	R/W
3.42.2	Receive test-pattern enable	1 = Enable receive test-pattern testing 0 = Disable receive test-pattern testing	R/W
3.42.1	Test-pattern select	1 = Square wave test pattern 0 = Pseudo random test pattern	R/W
3.42.0	Data pattern select	1 = Zeros data pattern 0 = LF data pattern	R/W

[a]RO = Read only, R/W = Read/Write

45.2.3.17.1 Scrambled idle test-pattern enable (3.42.7)

When bit 3.42.7 is set to a one, scrambled idle pattern testing is enabled. When bit 3.42.7 is set to zero, scrambled idle pattern testing is disabled.

The default value for bit 3.42.3 is zero.

45.2.3.17.2 10GBASE-R PRBS9 transmit test-pattern enable (3.42.6)

If the PCS supports the optional PRBS9 pattern testing (indicated by bit 3.32.3), the mandatory transmit test-pattern enable bit (3.42.3) is not one, and the optional PRBS31 transmit test-pattern enable bit (3.42.4) is not one, then when bit 3.42.6 is set to one the PCS shall transmit PRBS9. When bit 3.42.6 is set to zero, the PCS shall not generate PRBS9. The PRBS9 test-pattern is specified in 68.6.1. The behavior of the PCS when in PRBS9 test-pattern mode is specified in Clause 49.

45.2.3.17.3 10GBASE-R PRBS31 receive test-pattern enable (3.42.5)

If the PCS supports the optional PRBS31 pattern testing advertised in bit 3.32.2 and the mandatory receive test-pattern enable bit (3.42.2) is not one, setting bit 3.42.5 to a one shall set the receive path of the PCS into the PRBS31 test-pattern mode. The number of errors received during a PRBS31 pattern test are recorded in register 3.43. Setting bit 3.42.5 to a zero shall disable the PRBS31 test-pattern mode on the receive path of the PCS. The behavior of the PCS when in PRBS31 test-pattern mode is specified in Clause 49.

45.2.3.17.4 10GBASE-R PRBS31 transmit test-pattern enable (3.42.4)

If the PCS supports the optional PRBS31 pattern testing advertised in bit 3.32.2 and the mandatory transmit test-pattern enable bit (3.42.3) is not one, then setting bit 3.42.4 to a one shall set the transmit path of the PCS into the PRBS31 test-pattern mode. Setting bit 3.42.4 to a zero shall disable the PRBS31 test-pattern mode on the transmit path of the PCS. The behavior of the PCS when in PRBS31 test-pattern mode is specified in Clause 49.

45.2.3.17.5 Transmit test-pattern enable (3.42.3)

When bit 3.42.3 is set to a one, pattern testing is enabled on the transmit path. When bit 3.42.3 is set to a zero, pattern testing is disabled on the transmit path.

The default value for bit 3.42.3 is zero.

45.2.3.17.6 Receive test-pattern enable (3.42.2)

When bit 3.42.2 is set to a one, pattern testing is enabled on the receive path. When bit 3.42.2 is set to a zero, pattern testing is disabled on the receive path.

The default value for bit 3.42.2 is zero.

45.2.3.17.7 Test-pattern select (3.42.1)

When bit 3.42.1 is set to a one, the square wave test pattern is used for pattern testing. When bit 3.42.1 is set to a zero, the pseudo random test pattern is used for pattern testing.

The default value for bit 3.42.1 is zero.

45.2.3.17.8 Data pattern select (3.42.0)

When bit 3.42.0 is set to a one, the zeros data pattern is used for pattern testing. When bit 3.42.0 is set to a zero, the LF data pattern is used for pattern testing.

The default value for bit 3.42.0 is zero.

45.2.3.18 BASE-R PCS test-pattern error counter register (Register 3.43)

The assignment of bits in the BASE-R PCS test-pattern error counter register is shown in Table 45–133. This register is only required when the BASE-R capability is supported. If both BASE-R and 10GBASE-W capability is supported, then this register may either ignore writes and return zeros for reads when in 10GBASE-W mode, or may function as defined for BASE-R.

Table 45–133—BASE-R PCS test-pattern error counter register bit definitions

Bit(s)	Name	Description	R/W[a]
3.43.15:0	Test-pattern error counter	Error counter	RO

[a]RO = Read only

The test-pattern error counter is a sixteen bit counter that contains the number of errors received during a pattern test. These bits shall be reset to all zeros when the test-pattern error counter is read by the management function or upon execution of the PCS reset. These bits shall be held at all ones in the case of overflow. The test-pattern methodology is described in 49.2.12 and 82.2.11. This counter will count either block errors or bit errors dependent on the test mode (see 49.2.12).

45.2.3.19 BER high order counter (Register 3.44)

The assignment of bits in the BER high order counter register is shown in Table 45–134. This register is mandatory when the 40/100GBASE-R capability is supported and optional for other PHY types that implement register 3.33.

Table 45–134—BER high order counter register bit definitions

Bit(s)	Name	Description	R/W[a]
3.44.15:0	BER high order	Bits 21:6 of BER counter	RO

[a]RO = Read only

Bits 15:0 of this register are concatenated with bits 13:8 of register 3.33 (see 45.2.3.14) to read the value of a twenty-two bit BER counter. When registers 3.44 and 3.33 are used to read the 22-bit counter value, the register 3.33 is read first (indicating the lower 6 bits). The remaining 16 bits shall be latched when register 3.33 is read and reads to register 3.44 return the latched value. The counter continues to function as defined in 45.2.3.14.3 regardless of the state of the latched register. The 22-bit counter shall be reset to all zeros when register 3.33 is read or upon PCS reset. The 22-bit counter shall be held at all ones in the case of overflow.

45.2.3.20 Errored blocks high order counter (Register 3.45)

The assignment of bits in the Errored blocks high order counter register is shown in Table 45–135. This register is mandatory when the 40/100GBASE-R capability is supported and optional for other PHY types that implement register 3.33. If this register is implemented then bit 3.45.15 shall be set to 1.

Table 45–135—Errored blocks high order counter register bit definitions

Bit(s)	Name	Description	R/W[a]
3.45.15	High order counter present	Always reads as 1 if this register is implemented	RO
3.45.14	Reserved	Value always 0	RO
3.45.13:0	Errored blocks high order	Bits 21:8 of errored blocks counter	RO

[a]RO = Read only

Bits 13:0 of this register are concatenated with bits 7:0 of register 3.33 (see 45.2.3.14) to read the value of a twenty-two bit errored blocks counter. When registers 3.45 and 3.33 are used to read the 22-bit counter value, the register 3.33 is read first (indicating the lower 8 bits). The remaining 14 bits shall be latched when register 3.33 is read and reads to register 3.45 return the latched value. The counter continues to function as defined in 45.2.3.14.4 regardless of the state of the latched register. The 22-bit counter shall be reset to all zeros when register 3.33 is read or upon PCS reset. The 22-bit counter shall be held at all ones in the case of overflow.

45.2.3.21 Multi-lane BASE-R PCS alignment status 1 register (Register 3.50)

The assignment of bits in the multi-lane BASE-R PCS alignment status 1 register is shown in Table 45–136. All the bits in the multi-lane BASE-R PCS alignment status 1 register are read only; a write to the multi-lane

BASE-R PCS alignment status 1 register shall have no effect. A PCS device that does not implement multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 1 register. A device that implements multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 1 register that are not required for the PCS configuration. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

Table 45–136—Multi-lane BASE-R PCS alignment status 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.50.15:13	Reserved	Value always 0	RO
3.50.12	PCS lane alignment status	1 = PCS receive lanes locked and aligned 0 = PCS receive lanes not locked and aligned	RO
3.50.11:8	Reserved	Value always 0	RO
3.50.7	Block 7 lock	1 = Lane 7 is locked 0 = Lane 7 is not locked	RO
3.50.6	Block 6 lock	1 = Lane 6 is locked 0 = Lane 6 is not locked	RO
3.50.5	Block 5 lock	1 = Lane 5 is locked 0 = Lane 5 is not locked	RO
3.50.4	Block 4 lock	1 = Lane 4 is locked 0 = Lane 4 is not locked	RO
3.50.3	Block 3 lock	1 = Lane 3 is locked 0 = Lane 3 is not locked	RO
3.50.2	Block 2 lock	1 = Lane 2 is locked 0 = Lane 2 is not locked	RO
3.50.1	Block 1 lock	1 = Lane 1 is locked 0 = Lane 1 is not locked	RO
3.50.0	Block 0 lock	1 = Lane 0 is locked 0 = Lane 0 is not locked	RO

[a]RO = Read only

45.2.3.21.1 Multi-lane BASE-R PCS alignment status (3.50.12)

When read as a one, bit 3.50.12 indicates that the PCS has locked and aligned all receive lanes. When read as a zero, bit 3.50.12 indicates that the PCS has not locked and aligned all receive lanes.

45.2.3.21.2 Block 7 lock (3.50.7)

When read as a one, bit 3.50.7 indicates that the PCS receiver has achieved block lock for service interface lane 7. When read as a zero, bit 3.50.7 indicates that the PCS receiver lane 7 has not achieved block lock. This bit reflects the state of block_lock[7] (see 82.2.19.2.2).

45.2.3.21.3 Block 6 lock (3.50.6)

When read as a one, bit 3.50.6 indicates that the PCS receiver has achieved block lock for service interface lane 6. When read as a zero, bit 3.50.6 indicates that the PCS receiver lane 6 has not achieved block lock. This bit reflects the state of block_lock[6] (see 82.2.19.2.2).

45.2.3.21.4 Block 5 lock (3.50.5)

When read as a one, bit 3.50.5 indicates that the PCS receiver has achieved block lock for service interface lane 5. When read as a zero, bit 3.50.5 indicates that the PCS receiver lane 5 has not achieved block lock. This bit reflects the state of block_lock[5] (see 82.2.19.2.2).

45.2.3.21.5 Block 4 lock (3.50.4)

When read as a one, bit 3.50.4 indicates that the PCS receiver has achieved block lock for service interface lane 4. When read as a zero, bit 3.50.4 indicates that the PCS receiver lane 4 has not achieved block lock. This bit reflects the state of block_lock[4] (see 82.2.19.2.2).

45.2.3.21.6 Block 3 lock (3.50.3)

When read as a one, bit 3.50.3 indicates that the PCS receiver has achieved block lock for service interface lane 3. When read as a zero, bit 3.50.3 indicates that the PCS receiver lane 3 has not achieved block lock. This bit reflects the state of block_lock[3] (see 82.2.19.2.2).

45.2.3.21.7 Block 2 lock (3.50.2)

When read as a one, bit 3.50.2 indicates that the PCS receiver has achieved block lock for service interface lane 2. When read as a zero, bit 3.50.2 indicates that the PCS receiver lane 2 has not achieved block lock. This bit reflects the state of block_lock[2] (see 82.2.19.2.2).

45.2.3.21.8 Block 1 lock (3.50.1)

When read as a one, bit 3.50.1 indicates that the PCS receiver has achieved block lock for service interface lane 1. When read as a zero, bit 3.50.1 indicates that the PCS receiver lane 1 has not achieved block lock. This bit reflects the state of block_lock[1] (see 82.2.19.2.2).

45.2.3.21.9 Block 0 lock (3.50.0)

When read as a one, bit 3.50.0 indicates that the PCS receiver has achieved block lock for service interface lane 0. When read as a zero, bit 3.50.0 indicates that the PCS receiver lane 0 has not achieved block lock. This bit reflects the state of block_lock[0] (see 82.2.19.2.2).

45.2.3.22 Multi-lane BASE-R PCS alignment status 2 register (Register 3.51)

The assignment of bits in the multi-lane BASE-R PCS alignment status 2 register is shown in Table 45–137. All the bits in the multi-lane BASE-R PCS alignment status 2 register are read only; a write to the multi-lane BASE-R PCS alignment status 2 register shall have no effect. A PCS device that does not implement multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 2 register. A device that implements multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 2 register that are not required for the PCS configuration. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

Table 45–137—Multi-lane BASE-R PCS alignment status 2 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.51.15:12	Reserved	Value always 0	RO
3.51.11	Block 19 lock	1 = Lane 19 is locked 0 = Lane 19 is not locked	RO
3.51.10	Block 18 lock	1 = Lane 18 is locked 0 = Lane 18 is not locked	RO
3.51.9	Block 17 lock	1 = Lane 17 is locked 0 = Lane 17 is not locked	RO
3.51.8	Block 16 lock	1 = Lane 16 is locked 0 = Lane 16 is not locked	RO
3.51.7	Block 15 lock	1 = Lane 15 is locked 0 = Lane 15 is not locked	RO
3.51.6	Block 14 lock	1 = Lane 14 is locked 0 = Lane 14 is not locked	RO
3.51.5	Block 13 lock	1 = Lane 13 is locked 0 = Lane 13 is not locked	RO
3.51.4	Block 12 lock	1 = Lane 12 is locked 0 = Lane 12 is not locked	RO
3.51.3	Block 11 lock	1 = Lane 11 is locked 0 = Lane 11 is not locked	RO
3.51.2	Block 10 lock	1 = Lane 10 is locked 0 = Lane 10 is not locked	RO
3.51.1	Block 9 lock	1 = Lane 9 is locked 0 = Lane 9 is not locked	RO
3.51.0	Block 8 lock	1 = Lane 8 is locked 0 = Lane 8 is not locked	RO

[a]RO = Read only

45.2.3.22.1 Block 19 lock (3.51.11)

When read as a one, bit 3.51.11 indicates that the PCS receiver has achieved block lock for service interface lane 19. When read as a zero, bit 3.51.11 indicates that the PCS receiver lane 19 has not achieved block lock. This bit reflects the state of block_lock[19] (see 82.2.19.2.2).

45.2.3.22.2 Block 18 lock (3.51.10)

When read as a one, bit 3.51.10 indicates that the PCS receiver has achieved block lock for service interface lane 18. When read as a zero, bit 3.51.10 indicates that the PCS receiver lane 18 has not achieved block lock This bit reflects the state of block_lock[18] (see 82.2.19.2.2).

45.2.3.22.3 Block 17 lock (3.51.9)

When read as a one, bit 3.51.9 indicates that the PCS receiver has achieved block lock for service interface lane 17. When read as a zero, bit 3.51.9 indicates that the PCS receiver lane 17 has not achieved block lock This bit reflects the state of block_lock[17] (see 82.2.19.2.2).

45.2.3.22.4 Block 16 lock (3.51.8)

When read as a one, bit 3.51.8 indicates that the PCS receiver has achieved block lock for service interface lane 16. When read as a zero, bit 3.51.8 indicates that the PCS receiver lane 16 has not achieved block lock This bit reflects the state of block_lock[16] (see 82.2.19.2.2).

45.2.3.22.5 Block 15 lock (3.51.7)

When read as a one, bit 3.51.7 indicates that the PCS receiver has achieved block lock for service interface lane 15. When read as a zero, bit 3.51.7 indicates that the PCS receiver lane 15 has not achieved block lock This bit reflects the state of block_lock[15] (see 82.2.19.2.2).

45.2.3.22.6 Block 14 lock (3.51.6)

When read as a one, bit 3.51.6 indicates that the PCS receiver has achieved block lock for service interface lane 14. When read as a zero, bit 3.51.6 indicates that the PCS receiver lane 14 has not achieved block lock This bit reflects the state of block_lock[14] (see 82.2.19.2.2).

45.2.3.22.7 Block 13 lock (3.51.5)

When read as a one, bit 3.51.5 indicates that the PCS receiver has achieved block lock for service interface lane 13. When read as a zero, bit 3.51.5 indicates that the PCS receiver lane 13 has not achieved block lock This bit reflects the state of block_lock[13] (see 82.2.19.2.2).

45.2.3.22.8 Block 12 lock (3.51.4)

When read as a one, bit 3.51.4 indicates that the PCS receiver has achieved block lock for service interface lane 12. When read as a zero, bit 3.51.4 indicates that the PCS receiver lane 12 has not achieved block lock This bit reflects the state of block_lock[12] (see 82.2.19.2.2).

45.2.3.22.9 Block 11 lock (3.51.3)

When read as a one, bit 3.51.3 indicates that the PCS receiver has achieved block lock for service interface lane 11. When read as a zero, bit 3.51.3 indicates that the PCS receiver lane 11 has not achieved block lock This bit reflects the state of block_lock[11] (see 82.2.19.2.2).

45.2.3.22.10 Block 10 lock (3.51.2)

When read as a one, bit 3.51.2 indicates that the PCS receiver has achieved block lock for service interface lane 10. When read as a zero, bit 3.51.2 indicates that the PCS receiver lane 10 has not achieved block lock This bit reflects the state of block_lock[10] (see 82.2.19.2.2).

45.2.3.22.11 Block 9 lock (3.51.1)

When read as a one, bit 3.51.1 indicates that the PCS receiver has achieved block lock for service interface lane 9. When read as a zero, bit 3.51.1 indicates that the PCS receiver lane 9 has not achieved block lock This bit reflects the state of block_lock[9] (see 82.2.19.2.2).

45.2.3.22.12 Block 8 lock (3.51.0)

When read as a one, bit 3.51.0 indicates that the PCS receiver has achieved block lock for service interface lane 8. When read as a zero, bit 3.51.0 indicates that the PCS receiver lane 8 has not achieved block lock This bit reflects the state of block_lock[8] (see 82.2.19.2.2).

45.2.3.23 Multi-lane BASE-R PCS alignment status 3 register (Register 3.52)

The assignment of bits in the multi-lane BASE-R PCS alignment status 3 register is shown in Table 45–138. All the bits in the multi-lane BASE-R PCS alignment status 3 register are read only; a write to the multi-lane BASE-R PCS alignment status 3 register shall have no effect. A PCS device that does not implement multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 3 register. A device that implements multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 3 register that are not required for the PCS configuration. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

Table 45–138—Multi-lane BASE-R PCS alignment status 3 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.52.15:8	Reserved	Value always 0	RO
3.52.7	Lane 7 aligned	1 = Lane 7 alignment marker is locked 0 = Lane 7 alignment marker is not locked	RO
3.52.6	Lane 6 aligned	1 = Lane 6 alignment marker is locked 0 = Lane 6 alignment marker is not locked	RO
3.52.5	Lane 5 aligned	1 = Lane 5 alignment marker is locked 0 = Lane 5 alignment marker is not locked	RO
3.52.4	Lane 4 aligned	1 = Lane 4 alignment marker is locked 0 = Lane 4 alignment marker is not locked	RO
3.52.3	Lane 3 aligned	1 = Lane 3 alignment marker is locked 0 = Lane 3 alignment marker is not locked	RO
3.52.2	Lane 2 aligned	1 = Lane 2 alignment marker is locked 0 = Lane 2 alignment marker is not locked	RO
3.52.1	Lane 1 aligned	1 = Lane 1 alignment marker is locked 0 = Lane 1 alignment marker is not locked	RO
3.52.0	Lane 0 aligned	1 = Lane 0 alignment marker is locked 0 = Lane 0 alignment marker is not locked	RO

[a]RO = Read only

45.2.3.23.1 Lane 7 aligned (3.52.7)

When read as a one, bit 3.52.7 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 7. When read as a zero, bit 3.52.7 indicates that the PCS receiver lane 7 has not achieved alignment marker lock. This bit reflects the state of am_lock[7] (see 82.2.19.2.2).

45.2.3.23.2 Lane 6 aligned (3.52.6)

When read as a one, bit 3.52.6 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 6. When read as a zero, bit 3.52.6 indicates that the PCS receiver lane 6 has not achieved alignment marker lock. This bit reflects the state of am_lock[6] (see 82.2.19.2.2).

45.2.3.23.3 Lane 5 aligned (3.52.5)

When read as a one, bit 3.52.5 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 5. When read as a zero, bit 3.52.5 indicates that the PCS receiver lane 5 has not achieved alignment marker lock. This bit reflects the state of am_lock[5] (see 82.2.19.2.2).

45.2.3.23.4 Lane 4 aligned (3.52.4)

When read as a one, bit 3.52.4 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 4. When read as a zero, bit 3.52.4 indicates that the PCS receiver lane 4 has not achieved alignment marker lock. This bit reflects the state of am_lock[4] (see 82.2.19.2.2).

45.2.3.23.5 Lane 3 aligned (3.52.3)

When read as a one, bit 3.52.3 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 3. When read as a zero, bit 3.52.3 indicates that the PCS receiver lane 3 has not achieved alignment marker lock. This bit reflects the state of am_lock[3] (see 82.2.19.2.2).

45.2.3.23.6 Lane 2 aligned (3.52.2)

When read as a one, bit 3.52.2 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 2. When read as a zero, bit 3.52.2 indicates that the PCS receiver lane 2 has not achieved alignment marker lock. This bit reflects the state of am_lock[2] (see 82.2.19.2.2).

45.2.3.23.7 Lane 1 aligned (3.52.1)

When read as a one, bit 3.52.1 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 1. When read as a zero, bit 3.52.1 indicates that the PCS receiver lane 1 has not achieved alignment marker lock. This bit reflects the state of am_lock[1] (see 82.2.19.2.2).

45.2.3.23.8 Lane 0 aligned (3.52.0)

When read as a one, bit 3.52.0 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 0. When read as a zero, bit 3.52.0 indicates that the PCS receiver lane 0 has not achieved alignment marker lock. This bit reflects the state of am_lock[0] (see 82.2.19.2.2).

45.2.3.24 Multi-lane BASE-R PCS alignment status 4 register (Register 3.53)

The assignment of bits in the multi-lane BASE-R PCS alignment status 4 register is shown in Table 45–139. All the bits in the multi-lane BASE-R PCS alignment status 4 register are read only; a write to the multi-lane BASE-R PCS alignment status 4 register shall have no effect. A PCS device that does not implement multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 4 register. A device that implements multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 4 register that are not required for the PCS configuration. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

Table 45–139—Multi-lane BASE-R PCS alignment status 4 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.53.15:12	Reserved	Value always 0	RO
3.53.11	Lane 19 aligned	1 = Lane 19 alignment marker is locked 0 = Lane 19 alignment marker is not locked	RO
3.53.10	Lane 18 aligned	1 = Lane 18 alignment marker is locked 0 = Lane 18 alignment marker is not locked	RO
3.53.9	Lane 17 aligned	1 = Lane 17 alignment marker is locked 0 = Lane 17 alignment marker is not locked	RO
3.53.8	Lane 16 aligned	1 = Lane 16 alignment marker is locked 0 = Lane 16 alignment marker is not locked	RO
3.53.7	Lane 15 aligned	1 = Lane 15 alignment marker is locked 0 = Lane 15 alignment marker is not locked	RO
3.53.6	Lane 14 aligned	1 = Lane 14 alignment marker is locked 0 = Lane 14 alignment marker is not locked	RO
3.53.5	Lane 13 aligned	1 = Lane 13 alignment marker is locked 0 = Lane 13 alignment marker is not locked	RO
3.53.4	Lane 12 aligned	1 = Lane 12 alignment marker is locked 0 = Lane 12 alignment marker is not locked	RO
3.53.3	Lane 11 aligned	1 = Lane 11 alignment marker is locked 0 = Lane 11 alignment marker is not locked	RO
3.53.2	Lane 10 aligned	1 = Lane 10 alignment marker is locked 0 = Lane 10 alignment marker is not locked	RO
3.53.1	Lane 9 aligned	1 = Lane 9 alignment marker is locked 0 = Lane 9 alignment marker is not locked	RO
3.53.0	Lane 8 aligned	1 = Lane 8 alignment marker is locked 0 = Lane 8 alignment marker is not locked	RO

[a]RO = Read only

45.2.3.24.1 Lane 19 aligned (3.53.11)

When read as a one, bit 3.53.11 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 19. When read as a zero, bit 3.53.11 indicates that the PCS receiver lane 19 has not achieved alignment marker lock. This bit reflects the state of am_lock[19] (see 82.2.19.2.2).

45.2.3.24.2 Lane 18 aligned (3.53.10)

When read as a one, bit 3.53.10 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 18. When read as a zero, bit 3.53.10 indicates that the PCS receiver lane 18 has not achieved alignment marker lock. This bit reflects the state of am_lock[18] (see 82.2.19.2.2).

45.2.3.24.3 Lane 17 aligned (3.53.9)

When read as a one, bit 3.53.9 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 17. When read as a zero, bit 3.53.9 indicates that the PCS receiver lane 17 has not achieved alignment marker lock. This bit reflects the state of am_lock[17] (see 82.2.19.2.2).

45.2.3.24.4 Lane 16 aligned (3.53.8)

When read as a one, bit 3.53.8 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 16. When read as a zero, bit 3.53.8 indicates that the PCS receiver lane 16 has not achieved alignment marker lock. This bit reflects the state of am_lock[16] (see 82.2.19.2.2).

45.2.3.24.5 Lane 15 aligned (3.53.7)

When read as a one, bit 3.53.7 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 15. When read as a zero, bit 3.53.7 indicates that the PCS receiver lane 15 has not achieved alignment marker lock. This bit reflects the state of am_lock[15] (see 82.2.19.2.2).

45.2.3.24.6 Lane 14 aligned (3.53.6)

When read as a one, bit 3.53.6 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 14. When read as a zero, bit 3.53.6 indicates that the PCS receiver lane 14 has not achieved alignment marker lock. This bit reflects the state of am_lock[14] (see 82.2.19.2.2).

45.2.3.24.7 Lane 13 aligned (3.53.5)

When read as a one, bit 3.53.5 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 13. When read as a zero, bit 3.53.5 indicates that the PCS receiver lane 13 has not achieved alignment marker lock. This bit reflects the state of am_lock[13] (see 82.2.19.2.2).

45.2.3.24.8 Lane 12 aligned (3.53.4)

When read as a one, bit 3.53.4 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 12. When read as a zero, bit 3.53.4 indicates that the PCS receiver lane 12 has not achieved alignment marker lock. This bit reflects the state of am_lock[12] (see 82.2.19.2.2).

45.2.3.24.9 Lane 11 aligned (3.53.3)

When read as a one, bit 3.53.3 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 11. When read as a zero, bit 3.53.3 indicates that the PCS receiver lane 11 has not achieved alignment marker lock. This bit reflects the state of am_lock[11] (see 82.2.19.2.2).

45.2.3.24.10 Lane 10 aligned (3.53.2)

When read as a one, bit 3.53.2 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 10. When read as a zero, bit 3.53.2 indicates that the PCS receiver lane 10 has not achieved alignment marker lock. This bit reflects the state of am_lock[10] (see 82.2.19.2.2).

45.2.3.24.11 Lane 9 aligned (3.53.1)

When read as a one, bit 3.53.1 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 9. When read as a zero, bit 3.53.1 indicates that the PCS receiver lane 9 has not achieved alignment marker lock. This bit reflects the state of am_lock[9] (see 82.2.19.2.2).

45.2.3.24.12 Lane 8 aligned (3.53.0)

When read as a one, bit 3.53.8 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 0. When read as a zero, bit 3.53.8 indicates that the PCS receiver lane 0 has not achieved alignment marker lock. This bit reflects the state of am_lock[8] (see 82.2.19.2.2).

45.2.3.25 10P/2B capability register (3.60)

The 10P/2B capability register reports which functions are supported by the PCS. This register is present at the PCS layer for each PHY. The bit definitions of the 10P/2B capability register are shown in Table 45–140.

Table 45–140—10P/2B capability register bit definitions

Bit(s)	Name	Description	R/W[a]
3.60.15:13	Reserved	Value always 0	RO
3.60.12	PAF available	1 = PAF supported 0 = PAF not supported	RO
3.60.11	Remote PAF supported	1 = link partner supports PAF 0 = link partner does not support PAF	RO
3.60.10:0	Reserved	Value always 0	RO

[a]RO = Read only

45.2.3.25.1 PAF available (3.60.12)

This bit indicates that the PHY supports the PME aggregation function. The PHY sets this bit to a one when the capability is supported and zero otherwise. This bit reflects the signal PAF_available in 61.2.3.

45.2.3.25.2 Remote PAF supported (3.60.11)

This bit indicates that the remote, link-partner PHY supports the PME aggregation function. The PHY sets this bit to a one when the capability is supported and zero otherwise. This bit does not accurately report the capability of the remote PCS until a remote discovery operation has been completed by the -O PHY. In this case, this bit is set if the “Ethernet bonding” NPar(2) bit is set in the capabilities exchange message received from the other device. See 61.4.7, which discusses use of G.994.1 to access remote registers.

45.2.3.26 10P/2B PCS control register (Register 3.61)

The assignment of bits in the 10P/2B PCS control register is shown in Table 45–141.

Table 45–141—10B/2B PCS control register bit definitions

Bit(s)	Name	Description	R/W[a]
3.61.15	MII receive during transmit	1 = MII can TX/RX simultaneously 0 = MII cannot TX/RX simultaneously (default)	R/W
3.61.14	TX_EN and CRS infer a collision	1 = MII uses TX_EN and CRS to infer a collision 0 = MII uses COL to indicate a collision (default)	R/W

Table 45–141—10B/2B PCS control register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
3.61.13:1	Reserved	Value always 0	RO
3.61.0	PAF enable	1 = use PAF 0 = do not use PAF	O: R/W R: RO

[a]R/W = Read/Write, RO = Read only

45.2.3.26.1 MII receive during transmit (3.61.15)

This register bit is used to tell the PHY-MAC rate matching function if the MAC is capable of receiving frames from the PHY while the MAC is transmitting (i.e., sending frames to the PHY). The variable tx_rx_simultaneously for the PHY-MAC Rate-Matching function takes on the value of this bit as defined in 61.2.3.

45.2.3.26.2 TX_EN and CRS infer a collision (3.61.14)

This bit is set by the STA to tell the MAC-PHY rate matching function that the MAC-PHY interface does not have a separate collision signal but instead infers a collision when TX_EN and CRS are asserted simultaneously. The variable crs_and_tx_en_infer_col in the PHY-MAC Rate-Matching function takes on the value of this bit as in 61.2.3. This bit will default to a supported mode, and writes to unsupported modes will be ignored.

45.2.3.26.3 PAF enable (3.61.0)

Setting this bit to a one shall activate the PME aggregation function of the PCS when the link is established. Writes to this bit while link is up or initializing (see 45.2.1.15) or if the PAF is not supported shall be ignored. When link is established, handshake indicates the use of PAF to the -R PHY. This bit reflects the signal PAF_enable in 61.2.3.

45.2.3.27 10P/2B PME available (Registers 3.62 and 3.63)

The 10P/2B PME available registers are used to indicate which PMEs in the aggregation group are available to be attached to the queried PCS. A PME is marked as unavailable if the PME does not support PME aggregation or if the PME is currently marked to be aggregated with another PMD. For a device that does not support aggregation of multiple PMEs, a single bit of this register shall be set to one and all other bits cleared to zero.

These registers may be writeable for -R ports. For PMEs that may be accessed through more than one MII, the availability is limited such that no PME may be mapped to more than one MII prior to enabling the links. In this case, the reset state of the 10P/2B PME available registers shall reflect the capabilities of the device, the management entity should reset appropriate bits to meet the restriction described.

If the -R device is not capable of aggregating PMEs to multiple MIIs then these registers may be read only.

The 10P/2B PME available register shall be available per PCS. For example, a package implementing four PMEs and one MII would have only one set of 10P/2B PME available registers, addressed by a read or write to 3.62 and 3.63 on any of those PHYs.

For more information, see 61.2.2.8.3.

The assignment of bits in the 10P/2B PME available registers is shown in Table 45–142.

Table 45–142—10P/2B PME available register bit definitions

Bit(s)	Name	Description	R/W[a]
3.62.15:0	PME [p = 31:16] available	For each bit in the sequence: 1 = PME[p] is available for aggregating 0 = PME[p] is unavailable	O: RO R: R/W
3.63.15:0	PME [p = 15:0] available	For each bit in the sequence: 1 = PME[p] is available for aggregating 0 = PME[p] is unavailable	O: RO R: R/W

[a]RO = Read only, R/W = Read/Write

45.2.3.28 10P/2B PME aggregate registers (Registers 3.64 and 3.65)

The 10P/2B PME aggregate registers are used to select PMEs for aggregation. Attempts to activate aggregation with an unavailable PME (see 45.2.3.27) are ignored. The PCS shall use PME aggregation if one or more bits are set to a one and if PME aggregation is supported.

The 10P/2B PME aggregate register shall be available per PCS. For example, a package implementing four PMEs and one MII would have only one set of 10P/2B PME aggregate registers, accessed by a read or write to 3.64, 3.65 on any of those PHYs.

Upon MMD reset, these registers shall be reset to all zeros.

For more information, see 61.2.2.8.3.

The assignment of bits for the 10P/2B PME aggregate registers is shown in Table 45–143.

Table 45–143—10P/2B PME aggregate register bit definitions

Bit(s)	Name	Description	R/W[a]
3.64.15:0	Aggregate with PME [p = 31:16]	For each bit in the sequence: 1 = activate aggregation with PME[p] 0 = deactivate aggregation with PME[p]	R/W
3.65.15:0	Aggregate with PME [p = 15:0]	For each bit in the sequence: 1 = activate aggregation with PME[p] 0 = deactivate aggregation with PME[p]	R/W

[a]R/W = Read/Write

45.2.3.29 10P/2B PAF RX error register (Register 3.66)

The 10P/2B PAF RX error register is a 16 bit counter that contains the number of fragments that have been received across the gamma interface with RxErr asserted. The corresponding signal, TC_PAF_RxErrorReceived, is defined in 61.2.3. This counter is inactive when the PAF is unsupported or

disabled. Upon disabling the PAF, the register retains its previous value. These bits shall be reset to all zeros when the 10P/2B PAF RX error register is read by the management function or upon execution of the MMD reset. These bits shall be held at all ones in the case of overflow. The assignment of bits in the 10P/2B PAF RX error register is shown in Table 45–144.

Table 45–144—10P/2B PAF RX error register bit definitions

Bit(s)	Name	Description	R/W[a]
3.66.15:0	PAF RX errors[15:0]	The bytes of the counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.3.30 10P/2B PAF small fragments register (Register 3.67)

The 10P/2B PAF small fragments register is a 16 bit counter that contains the number of small fragments that have been received across the gamma interface. The corresponding signal, TC_PAF_FragmentTooSmall, is defined in 61.2.3. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the MMD reset. This counter is inactive when the PAF is unsupported or disabled. Upon disabling the PAF, the register retains its previous value. These bits shall be held at all ones in the case of overflow. The assignment of bits in the 10P/2B PAF small fragment register is shown in Table 45–145.

Table 45–145—10P/2B PAF small fragments register bit definitions

Bit(s)	Name	Description	R/W[a]
3.67.15:0	PAF small fragments[15:0]	The bytes of the counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.3.31 10P/2B PAF large fragments register (Register 3.68)

The 10P/2B PAF large fragments register is a 16 bit counter that contains the number of large fragments that have been received across the gamma interface. The corresponding signal, TC_PAF_FragmentTooLarge, is defined in 61.2.3. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the MMD reset. This counter is inactive when the PAF is unsupported or disabled. Upon disabling the PAF, the register retains its previous value. These bits shall be held at all ones in the case of overflow. The assignment of bits in the 10P/2B PAF large fragments register is shown in Table 45–146.

Table 45–146—10P/2B PAF large fragments register bit definitions

Bit(s)	Name	Description	R/W[a]
3.68.15:0	PAF large fragments[15:0]	The bytes of the counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.3.32 10P/2B PAF overflow register (Register 3.69)

The 10P/2B PAF overflow register is a 16 bit counter that contains the number of fragments that have been received across the gamma interface which would have caused the receive buffer to overflow. The corresponding signal, TC_PAF_Overflow, is defined in 61.2.3. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the MMD reset. This counter is inactive when the PAF is unsupported or disabled. Upon disabling the PAF, the register retains its previous value. These bits shall be held at all ones in the case of overflow. The assignment of bits in the 10P/2B PAF overflow register is shown in Table 45–147.

Table 45–147—10P/2B PAF overflow register bit definitions

Bit(s)	Name	Description	R/W[a]
3.69.15:0	PAF overflow fragments[15:0]	The bytes of the counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.3.33 10P/2B PAF bad fragments register (Register 3.70)

The 10P/2B PAF bad fragments register is a 16 bit counter that contains the number of bad fragments that have been received across the gamma interface. The corresponding signal, TC_PAF_BadFragmentReceived, is defined in 61.2.3. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the MMD reset. This counter is inactive when the PAF is unsupported or disabled. Upon disabling the PAF, the register retains its previous value. These bits shall be held at all ones in the case of overflow. The assignment of bits in the 10P/2B PAF bad fragments register is shown in Table 45–148.

Table 45–148—P10P/2B AF bad fragments register bit definitions

Bit(s)	Name	Description	R/W[a]
3.70.15:0	PAF bad fragments[15:0]	The bytes of the counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.3.34 10P/2B PAF lost fragments register (Register 3.71)

The 10P/2B PAF lost fragments register is a 16 bit counter that contains the number of gaps in the sequence of fragments that have been received across the gamma interface. The corresponding signal, TC_PAF_LostFragment, is defined in 61.2.3.

These bits shall be reset to all zeros when the register is read by the management function or upon execution of the MMD reset. This counter is inactive when the PAF is unsupported or disabled. Upon disabling the PAF, the register retains its previous value. These bits shall be held at all ones in the case of overflow. The assignment of bits in the 10P/2B PAF lost fragments register is shown in Table 45–149.

Table 45–149—10P/2B PAF lost fragments register bit definitions

Bit(s)	Name	Description	R/W[a]
3.71.15:0	PAF lost fragments[15:0]	The bytes of the counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.3.35 10P/2B PAF lost starts of fragments register (Register 3.72)

The 10P/2B PAF lost starts of fragments register is a 16-bit counter that contains the number of missing start of fragment indicators expected by the frame assembly function. The corresponding signal, TC_PAF_LostStart, is defined in 61.2.3. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the MMD reset. This counter is inactive when the PAF is unsupported or disabled. Upon disabling the PAF, the register retains its previous value. These bits shall be held at all ones in the case of overflow. The assignment of bits in the 10P/2B PAF lost starts of fragments register is shown in Table 45–150.

Table 45–150—10P/2B PAF lost starts of fragments register bit definitions

Bit(s)	Name	Description	R/W[a]
3.72.15:0	PAF lost starts of fragments[15:0]	The bytes of the counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.3.36 10P/2B PAF lost ends of fragments register (Register 3.73)

The 10P/2B PAF lost ends of fragments register is a 16 bit counter that contains the number of missing end of fragment indicators expected by the frame assembly function. The corresponding signal, TC_PAF_LostEnd, is defined in 61.2.3. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the MMD reset. This counter is inactive when the PAF is unsupported or disabled. Upon disabling the PAF, the register retains its previous value. These bits shall be held at all ones in the case of overflow. The assignment of bits in the 10P/2B PAF lost ends of fragments register is shown in Table 45–151.

Table 45–151—10P/2B PAF lost ends of fragments register bit definitions

Bit(s)	Name	Description	R/W[a]
3.73.15:0	PAF lost ends of fragments[15:0]	The bytes of the counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.3.37 10GBASE-PR and 10/1GBASE-PRX FEC ability register (Register 3.74)

The assignment of bits in the 10GBASE-PR and 10/1GBASE-PRX FEC ability register is shown in Table 45–152.

Table 45–152—10GBASE-PR and 10/1GBASE-PRX FEC ability register bit definitions

Bit(s)	Name	Description	R/W[a]
3.74.15:1	Reserved	Value always 0	RO
3.74.0	10 Gb/s FEC ability	A read of 1 in this bit indicates that the PCS supports the 10/1GBASE-PRX or 10GBASE-PR 10 Gb/s FEC (always reads as 1 for 10/1GBASE-PRX or 10GBASE-PR).	RO

[a]RO = Read only

45.2.3.38 10GBASE-PR and 10/1GBASE-PRX FEC control register (Register 3.75)

The assignment of bits in the 10GBASE-PR FEC control register is shown in Table 45–153.

Table 45–153—10GBASE-PR and 10/1GBASE-PRX FEC control register bit definitions

Bit(s)	Name	Description	R/W[a]
3.75.15:2	Reserved	Value always 0	RO
3.75.1	Enable FEC error indication	A write of 1 to this bit configures the 10 Gb/s FEC decoder to indicate uncorrectable codeword errors to the higher layer. In a 10/1GBASE-PRX OLT, this bit is undefined.	R/W
3.75.0	10 Gb/s FEC enable	Always reads as 1 for 10/1GBASE-PRX or 10GBASE-PR since 10 Gb/s FEC is always enabled.	RO

[a]RO = Read only, R/W = Read/Write

45.2.3.38.1 FEC enable error indication (3.75.1)

This bit instructs the 10 Gb/s FEC decoder component of the 10GBASE-PR and 10/1GBASE-PRX PCS to indicate decoding errors to the upper layers (see 45.2.3.38 and 76.3.3.3).

When written as a one, the receiving PCS invalidates 66 bit blocks received in uncorrectable FEC codewords. As a consequence, the receiving MAC discards any packet which includes data that was received in an uncorrectable FEC codeword (even though the packet itself might or might not contain errors).

When written as a zero, the receiving PCS does not modify 66 bit blocks received in uncorrectable FEC codewords. As a consequence, the receiving MAC performs regular processing on a packet that includes data that was received in an uncorrectable FEC codeword (though the packet itself may contain errors which might or might not be detected by the MAC FCS)

45.2.3.38.2 10 Gb/s FEC Enable (3.75.0)

This bit indicates whether 10 Gb/s FEC is enabled in the 10GBASE-PR and 10/1GBASE-PRX PCS and always reads as one.

The register for enabling and disabling forward error correction in the 10/1GBASE-PRX upstream is specified in 45.2.8.3.

45.2.3.39 10/1GBASE-PRX and 10GBASE-PR corrected FEC codewords counter (Register 3.76, 3.77)

The assignment of bits in the 10/1GBASE-PRX and 10GBASE-PR corrected FEC codewords counter register is shown in Table 45–154. See 76.3.3.3.2 for a definition of this counter. These bits shall be reset to

all zeros when the register is read by the management function or upon PCS reset. These bits shall be held at all ones in the case of overflow.

Table 45–154—10GBASE-PR corrected FEC codewords counter register bit definitions

Bit(s)	Name	Description	R/W[a]
3.76.15:0	corrected FEC codewords lower	corrected_FEC_codewords_counter[15:0]	RO, MW, NR
3.77.15:0	corrected FEC codewords upper	corrected_FEC_codewords_counter[31:16]	RO, MW, NR

[a]RO = Read only, MW = Multi-word, NR = Non Roll-over

45.2.3.40 10/1GBASE-PRX and 10GBASE-PR uncorrected FEC codewords counter (Register 3.78, 3.79)

The assignment of bits in the 10/1GBASE-PRX and 10GBASE-PR uncorrected FEC codewords counter register is shown in Table 45–155. See 76.3.3.3.2 for a definition of this counter. These bits shall be reset to all zeros when the register is read by the management function or upon PCS reset. These bits shall be held at all ones in the case of overflow.

Table 45–155—10GBASE-PR uncorrected FEC codewords counter register bit definitions

Bit(s)	Name	Description	R/W[a]
3.78.15:0	uncorrected FEC codewords lower	uncorrected_FEC_codewords_counter[15:0]	RO, MW, NR
3.79.15:0	uncorrected FEC codewords lower	uncorrected_FEC_codewords_counter[32:16]	RO, MW, NR

[a]RO = Read only, MW = Multi-word, NR = Non Roll-over

45.2.3.41 10GBASE-PR and 10/1GBASE-PRX BER monitor timer control register (Register 3.80)

The assignment of bits in the 10GBASE-PR and 10/1GBASE-PRX BER monitor timer control register is shown in Table 45–156. This register is defined only when 10GBASE-PR or 10/1GBASE-PRX ONU capability is supported. The 10G-EPON BER monitor is described in 76.3.3.4.

Table 45–156—10GBASE-PR and 10/1GBASE-PRX BER monitor timer control register bit definitions

Bit(s)	Name	Description	R/W[a]
3.80.7:0	10G-EPON BER monitor timer	Duration (in units of 5 microseconds) of the timer used by the 10G-EPON BER monitor function. Default value is 25 (i.e., 125 microseconds). A value of 0 indicates that the BER monitor function is disabled.	R/W

[a]R/W = Read/Write

45.2.3.42 10GBASE-PR and 10/1GBASE-PRX BER monitor status (Register 3.81)

The assignment of bits in the 10GBASE-PR and 10/1GBASE-PRX BER monitor status register is shown in Table 45–157. This register is defined only when 10GBASE-PR or 10/1GBASE-PRX ONU capability is supported.

Table 45–157—10GBASE-PR and 10/1GBASE-PRX BER monitor status register bit definitions

Bit(s)	Name	Description	R/W[a]
3.81.7:2	Reserved	Value always 0	RO
3.81.1	Latched high BER	1 = 10GBASE-PR or 10/1GBASE-PRX PCS reported a high BER. 0 = 10GBASE-PR or 10/1GBASE-PRX PCS did not report a high BER.	RO, LH
3.81.0	high BER	1 = 10GBASE-PR or 10/1GBASE-PRX PCS reporting a high BER. 0 = 10GBASE-PR or 10/1GBASE-PRX PCS not reporting a high BER.	RO

[a]RO Read only, LH = Latching high

45.2.3.42.1 10GBASE-PR and 10/1GBASE-PRX PCS high BER (3.81.0)

In the 10GBASE-PR and 10/1GBASE-PRX PCS, when read as a one, bit 3.81.0 indicates that the receiver is detecting a BER greater than the configurable threshold (high BER state). When read as a zero, bit 3.81.0 indicates that the receiver is detecting a BER lower than the configurable threshold (low BER state). This bit mirrors the state of the hi_ber variable, defined in 76.3.3.4.

45.2.3.42.2 10GBASE-PR and 10/1GBASE-PRX PCS latched high BER (3.81.1)

In the 10GBASE-PR and 10/1GBASE-PRX PCS, when read as a one, bit 3.81.1 indicates that the receiver detected a BER greater than the configurable threshold (high BER state). When read as a zero, bit 3.81.1 indicates that the receiver detected BER lower than the configurable threshold (low BER state).

This bit is a latching high version of the 10GBASE-PR and 10/1GBASE-PRX high BER status bit (3.81.0).

45.2.3.43 10GBASE-PR and 10/1GBASE-PRX BER monitor threshold control (Register 3.82)

The assignment of bits in the 10GBASE-PR and 10/1GBASE-PRX BER monitor threshold control register is shown in Table 45–158. This register is defined only when 10GBASE-PR or 10/1GBASE-PRX ONU capability is supported. The 10G-EPON BER monitor is described in 76.3.3.4.

Table 45–158—10GBASE-PR and 10/1GBASE-PRX BER monitor threshold control register bit definitions

Bit(s)	Name	Description	R/W[a]
3.82.15:0	10G-EPON BER monitor threshold	Number of sync header errors within a timer interval that triggers a high BER condition for the 10G-EPON BER monitor function. Default value is 1600. A value of 0 indicates that the BER monitor function is disabled.	R/W

[a]R/W = Read/Write

45.2.3.44 BIP error counter lane 0 (Register 3.200)

The assignment of bits in the BIP error counter lane 0 is shown in Table 45–159. The multi-lane PCS described in Clause 82 calculates a BIP value for each PCS lane (see 82.2.8, 82.2.15). Errors detected in PCS lane 0 are counted and shown in register 3.200.15:0. The 16-bit counter shall be reset to all zeros when register 3.200 is read or upon PCS reset. The 16-bit counter shall be held at all ones in the case of overflow.

Table 45–159—BIP error counter, lane 0 register bit definitions

Bit(s)	Name	Description	R/W[a]
3.200.15:0	BIP error counter, lane 0	Errors detected by BIP in PCS lane 0	RO

[a]RO = Read only

45.2.3.45 BIP error counter, lanes 1 through 19 (Registers 3.201 through 3.219)

The behavior of the BIP error counters, lanes 1 through 19 is identical to that described for PCS lane 0 in 45.2.3.44. Errors detected in each PCS lane are counted and shown in register bits 15:0 in the corresponding register. PCS lane 1 is shown in register 3.201; PCS lane 2 is shown in register 3.202; through register 3.219 for PCS lane 19.

45.2.3.46 Lane 0 mapping register (Register 3.400)

The assignment of bits in the Lane 0 mapping register is shown in Table 45–160. When the multi-lane PCS described in Clause 82 detects and locks the alignment marker for service interface lane 0, the detected PCS lane number is recorded in this register. The contents of the Lane 0 mapping register is valid when Lane 0 aligned bit (3.52.0) is set to one and is invalid otherwise.

Table 45–160—Lane 0 mapping register bit definitions

Bit(s)	Name	Description	R/W[a]
3.400.15:5	Reserved	Value always 0	RO
3.400.4:0	Lane 0 mapping	PCS lane received in service interface lane 0	RO

[a]RO = Read only

45.2.3.47 Lanes 1 through 19 mapping registers (Registers 3.401 through 3.419)

The definition of lanes 1 through 19 mapping registers is identical to that described for lane 0 in 45.2.3.46. The lane mapping for lane 1 is in register 3.401; lane 2 is in register 3.402; etc.

45.2.3.48 TimeSync PCS capability (Register 3.1800)

The TimeSync PCS capability register (see Table 45–161) indicates the capability of the PCS to report the transmit and receive data delay, stored in registers 3.1801 through 3.1804 and 3.1805 through 3.1808, respectively.

Table 45–161—TimeSync PCS capability

Bit(s)	Name	Description	R/W[a]
3.1800.15:2	Reserved	Value always 0	RO
3.1800.1	TimeSync transmit path data delay	1 = PCS provides information on transmit path data delay in registers 3.1801 through 3.1804 0 = PCS does not provide information on transmit path data delay	RO
3.1800.0	TimeSync receive path data delay	1 = PCS provides information on receive path data delay in registers 3.1805 through 3.1808 0 = PCS does not provide information on receive path data delay	RO

[a]RO = Read only

45.2.3.49 TimeSync PCS transmit path data delay (Registers 3.1801, 3.1802, 3.1803, 3.1804)

The TimeSync PCS transmit path data delay register contains the maximum (Registers 3.1801, 3.1802, see Table 45–162) and minimum (Registers 3.1803, 3.1804, see Table 45–162) values of the transmit path data delay. The transmit path data delay is expressed in units of ns. The values contained in these registers are valid when the link is established, as indicated by bit 2 in Register 1.1 (see 45.2.1.2.4).

Table 45–162—TimeSync PCS transmit path data delay register

Bit(s)	Name	Description	R/W[a]
3.1801.15:0	Maximum PCS transmit path data delay, lower	PCS_delay_TX_max [15:0]	RO, MW
3.1802.15:0	Maximum PCS transmit path data delay, upper	PCS_delay_TX_max [31:16]	RO, MW
3.1803.15:0	Minimum PCS transmit path data delay, lower	PCS_delay_TX_min [15:0]	RO, MW
3.1804.15:0	Minimum PCS transmit path data delay, upper	PCS_delay_TX_min [31:16]	RO, MW

[a]RO = Read only, MW = Multi-word

45.2.3.50 TimeSync PCS receive path data delay (Registers 3.1805, 3.1806, 3.1807, 3.1808)

The TimeSync PCS receive path data delay register contains the maximum (Registers 3.1805, 3.1806, see Table 45–163) and minimum (Registers 3.1807, 3.1808, see Table 45–163) values of the receive path data

delay. The receive path data delay is expressed in units of ns. The values contained in these registers are valid when the link is established, as indicated by bit 2 in Register 1.1 (see 45.2.1.2.4).

Table 45–163—TimeSync PCS receive path data delay register

Bit(s)	Name	Description	R/W[a]
3.1805.15:0	Maximum PCS receive path data delay, lower	PCS_delay_RX_max [15:0]	RO, MW
3.1806.15:0	Maximum PCS receive path data delay, upper	PCS_delay_RX_max [31:16]	RO, MW
3.1807.15:0	Minimum PCS receive path data delay, lower	PCS_delay_RX_min [15:0]	RO, MW
3.1808.15:0	Minimum PCS receive path data delay, upper	PCS_delay_RX_min [31:16]	RO, MW

[a]RO = Read only, MW = Multi-word

45.2.4 PHY XS registers

The assignment of registers in the PHY XS is shown in Table 45–164.

Table 45–164—PHY XS registers

Register address	Register name	Subclause
4.0	PHY XS control 1	45.2.4.1
4.1	PHY XS status 1	45.2.4.2
4.2, 4.3	PHY XS device identifier	45.2.4.3
4.4	PHY XS speed ability	45.2.4.4
4.5, 4.6	PHY XS devices in package	45.2.4.5
4.7	Reserved	
4.8	PHY XS status 2	45.2.4.6
4.9 through 4.13	Reserved	
4.14, 4.15	PHY XS package identifier	45.2.4.7
4.16 through 4.19	Reserved	
4.20	EEE capability	45.2.4.8
4.21	Reserved	
4.22	EEE wake error counter	45.2.4.9
4.23	Reserved	
4.24	10G PHY XGXS lane status	45.2.4.10
4.25	10G PHY XGXS test control	45.2.4.11
4.26 through 4.1799	Reserved	
4.1800	TimeSync PHY XS capability	45.2.4.12
4.1801 through 4.1804	TimeSync PHY XS transmit path data delay	45.2.4.13
4.1805 through 4.1808	TimeSync PHY XS receive path data delay	45.2.4.14

Table 45–164—PHY XS registers *(continued)*

Register address	Register name	Subclause
4.1809 through 4.32767	Reserved	
4.32 768 through 4.65 535	Vendor specific	

45.2.4.1 PHY XS control 1 register (Register 4.0)

The assignment of bits in the PHY XS control 1 register is shown in Table 45–165. The default value for each bit of the PHY XS control 1 register should be chosen so that the initial state of the device upon power up or reset is a normal operational state without management intervention.

Table 45–165—PHY XS control 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
4.0.15	Reset	1 = PHY XS reset 0 = Normal operation	R/W SC
4.0.14	Loopback	1 = Enable loopback mode 0 = Disable loopback mode	R/W
4.0.13	Speed selection	1 = Operation at 10 Gb/s and above 0 = Unspecified	R/W
4.0.12	Reserved	Value always 0	RO
4.0.11	Low power	1 = Low-power mode 0 = Normal operation	R/W
4.0.10	Clock stop enable	1 = The PHY XS may stop the clock during LPI 0 = Clock not stoppable	R/W
4.0.9	XAUI stop enable	1 = The PHY XS may stop XAUI signals during LPI 0 = XAUI not stoppable	R/W
4.0.8:7	Reserved	Value always 0	RO
4.0.6	Speed selection	1 = Operation at 10 Gb/s and above 0 = Unspecified	R/W
4.0.5:2	Speed selection	5 4 3 2 1 x x x = Reserved x 1 x x = Reserved x x 1 x = Reserved 0 0 0 1 = Reserved 0 0 0 0 = 10 Gb/s	R/W
4.0.1:0	Reserved	Value always 0	RO

[a]R/W = Read/Write, SC = Self-clearing, RO = Read only

45.2.4.1.1 Reset (4.0.15)

Resetting a PHY XS is accomplished by setting bit 4.0.15 to a one. This action shall set all PHY XS registers to their default states. As a consequence, this action may change the internal state of the PHY XS and the state of the physical link. This action may also initiate a reset in any other MMDs that are instantiated in the same package. This bit is self-clearing, and a PHY XS shall return a value of one in bit 4.0.15 when a reset is in progress and a value of zero otherwise. A PHY XS is not required to accept a write transaction to any of its registers until the reset process is completed. The reset process shall be completed within 0.5 s from the setting of bit 4.0.15. During a reset, a PHY XS shall respond to reads from register bits 4.0.15 and 4.8.15:14. All other register bits should be ignored.

NOTE—This operation may interrupt data communication.

45.2.4.1.2 Loopback (4.0.14)

The PHY XS shall be placed in a loopback mode of operation when bit 4.0.14 is set to a one. When bit 4.0.14 is set to a one, the PHY XS shall accept data on the receive path and return it on the transmit path. The direction of the loopback path for the PHY XS is opposite to all other MMD loopbacks.

The loopback function is optional. A device's ability to perform the loopback function is advertised in the loopback ability bit of the related speed-dependent status register. A PHY XS that is unable to perform the loopback function shall ignore writes to this bit and return a value of zero when read. For 10 Gb/s operation, the loopback functionality is detailed in 48.3.3 and the loopback ability bit is specified in the 10G PHY XGXS Lane status register.

The default value of bit 4.0.14 is zero.

NOTE—The signal path through the PHY XS that is exercised in the loopback mode of operation is implementation specific, but it is recommended that the signal path encompass as much of the PHY XS circuitry as is practical. The intention of providing this loopback mode of operation is to permit a diagnostic or self-test function to perform the transmission and reception of a PDU, thus testing the transmit and receive data paths. Other loopback signal paths may be enabled using loopback controls within other MMDs.

45.2.4.1.3 Low power (4.0.11)

A PHY XS may be placed into a low-power mode by setting bit 4.0.11 to a one. This action may also initiate a low-power mode in any other MMDs that are instantiated in the same package. The low-power mode is exited by resetting the PHY XS. The behavior of the PHY XS in transition to and from the low-power mode is implementation specific and any interface signals should not be relied upon. While in the low-power mode, the device shall, as a minimum, respond to management transactions necessary to exit the low-power mode. The default value of bit 4.0.11 is zero.

45.2.4.1.4 Clock stop enable (4.0.10)

If bit 4.0.10 is set to 1, then the PHY XS may stop the transmit direction xMII clock while it is signaling LPI, otherwise it shall keep the clock active. If the PHY XS does not support EEE capability or is not able to stop the transmit clock, then this bit has no effect (see 46.3.2.4).

45.2.4.1.5 XAUI stop enable (4.0.9)

If bit 4.0.9 is set to 1, then the PHY XS may stop signaling on the XAUI in the receive direction during LPI, otherwise the PHY XS shall keep the XAUI signals active. If the PHY XS does not support EEE capability or is not able to stop the receive path XAUI signals, then this bit has no effect.

45.2.4.1.6 Speed selection (4.0.13, 4.0.6, 4.0.5:2)

Speed selection bits 4.0.13 and 4.0.6 shall both be written as a one. Any attempt to change the bits to an invalid setting shall be ignored. These two bits are set to one in order to make them compatible with Clause 22.

The speed of the PHY XS may be selected using bits 5 through 2. The speed abilities of the PHY XS are advertised in the PHY XS speed ability register. A PHY XS may ignore writes to the PHY XS speed selection bits that select speeds it has not advertised in the PHY XS speed ability register. It is the responsibility of the STA entity to ensure that mutually acceptable speeds are applied consistently across all the MMDs on a particular PHY.

The PHY XS speed selection defaults to a supported ability.

45.2.4.2 PHY XS status 1 register (Register 4.1)

The assignment of bits in the PHY XS status 1 register is shown in Table 45–166. All the bits in the PHY XS status 1 register are read only; a write to the PHY XS status 1 register shall have no effect.

Table 45–166—PHY XS status 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
4.1.15:12	Reserved	Value always 0	RO
4.1.11	Tx LPI received	1 = Tx PHY XS has received LPI 0 = LPI not received	RO/LH
4.1.10	Rx LPI received	1 = Rx PHY XS has received LPI 0 = LPI not received	RO/LH
4.1.9	Tx LPI indication	1 = Tx PHY XS is currently receiving LPI 0 = PCS is not currently receiving LPI	RO
4.1.8	Rx LPI indication	1 = Rx PHY XS is currently receiving LPI 0 = PCS is not currently receiving LPI	RO
4.1.7	Fault	1 = Fault condition detected 0 = No fault condition detected	RO
4.1.6	Clock stop capable	1 = The attached PHY may stop the clock during LPI 0 = Clock not stoppable	RO
4.1.5:3	Reserved	Value always 0	RO
4.1.2	PHY XS transmit link status	1 = The PHY XS transmit link is up 0 = The PHY XS transmit link is down	RO/LL
4.1.1	Low-power ability	1 = PHY XS supports low-power mode 0 = PHY XS does not support low-power mode	RO
4.1.0	Reserved	Value always 0	RO

[a]RO = Read only, LL = Latching low, LH = Latching high

45.2.4.2.1 Transmit LPI received (4.1.11)

When read as a one, bit 4.1.11 indicates that the transmit PHY XS has received LPI signaling one or more times since the register was last read. When read as a zero, bit 4.1.11 indicates that the PHY XS has not received LPI signaling. This bit shall be implemented with latching high behavior.

45.2.4.2.2 Receive LPI received (4.1.10)

When read as a one, bit 4.1.10 indicates that the receive PHY XS has received LPI signaling one or more times since the register was last read. When read as a zero, bit 4.1.10 indicates that the PHY XS has not received LPI signaling. This bit shall be implemented with latching high behavior.

45.2.4.2.3 Transmit LPI indication (4.1.9)

When read as a one, bit 4.1.9 indicates that the transmit PHY XS is currently receiving LPI signals. When read as a zero, bit 4.1.9 indicates that the PHY XS is not currently receiving LPI signals. The behavior if read during a state transition is undefined.

45.2.4.2.4 Receive LPI indication (4.1.8)

When read as a one, bit 4.1.8 indicates that the receive PHY XS is currently receiving LPI signals. When read as a zero, bit 4.1.8 indicates that the PHY XS is not currently receiving LPI signals. The behavior if read during a state transition is undefined.

45.2.4.2.5 Fault (4.1.7)

When read as a one, bit 4.1.7 indicates that the PHY XS has detected a fault condition on either the transmit or receive paths. When read as a zero, bit 4.1.7 indicates that the PHY XS has not detected a fault condition. Bit 4.1.7 is set to a one when either of the fault bits (4.8.11, 4.8.10) located in register 4.8 are set to a one.

45.2.4.2.6 Clock stop capable (4.1.6)

If bit 4.1.6 is set to one then the PHY XS is indicating that the attached PHY is permitted to stop the receive direction xMII clock while it is signaling LPI. If the bit is set to zero then the PHY XS is indicating that the attached PHY is not permitted to stop the receive xMII clock while it is signaling LPI. If the attached PHY does not support EEE capability or is not able to stop the receive direction xMII clock then this bit has no effect (see 46.3.2.4).

45.2.4.2.7 PHY XS transmit link status (4.1.2)

When read as a one, bit 4.1.2 indicates that the PHY XS transmit link is aligned. When read as a zero, bit 4.1.2 indicates that the PHY XS transmit link is not aligned. The transmit link status bit shall be implemented with latching low behavior.

For 10 Gb/s operation, bit 4.1.2 is a latching low version of bit 4.24.12.

45.2.4.2.8 Low-power ability (4.1.1)

When read as a one, bit 4.1.1 indicates that the PHY XS supports the low-power feature. When read as a zero, bit 4.1.1 indicates that the PHY XS does not support the low-power feature. If a PHY XS supports the low-power feature then it is controlled using the low-power bit in the PHY XS control register.

45.2.4.3 PHY XS device identifier (Registers 4.2 and 4.3)

Registers 4.2 and 4.3 provide a 32-bit value, which may constitute a unique identifier for a PHY XS. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the device manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A PHY XS may return a value of zero in each of the 32 bits of the PHY XS device identifier.

The format of the PHY XS device identifier is specified in 22.2.4.3.1.

45.2.4.4 PHY XS speed ability (Register 4.4)

The assignment of bits in the PHY XS speed ability register is shown in Table 45–167.

Table 45–167—PHY XS speed ability register bit definitions

Bit(s)	Name	Description	R/W[a]
4.4.15:1	Reserved for future speeds	Value always 0	RO
4.4.0	10G capable	1 = PHY XS is capable of operating at 10 Gb/s 0 = PHY XS is not capable of operating at 10 Gb/s	RO

[a]RO = Read only

45.2.4.4.1 10G capable (4.4.0)

When read as a one, bit 4.4.0 indicates that the PHY XS is able to operate at a data rate of 10 Gb/s. When read as a zero, bit 4.4.0 indicates that the PHY-XS is not able to operate at a data rate of 10 Gb/s.

45.2.4.5 PHY XS devices in package (Registers 4.5 and 4.6)

The PHY XS devices in package registers are defined in Table 45–2.

45.2.4.6 PHY XS status 2 register (Register 4.8)

The assignment of bits in the PHY XS status 2 register is shown in Table 45–168. All the bits in the PHY XS status 2 register are read only; a write to the PHY XS status 2 register shall have no effect.

Table 45–168—PHY XS status 2 register bit definitions

Bit(s)	Name	Description	R/W[a]
4.8.15:14	Device present	15 14 1 0 = Device responding at this address 1 1 = No device responding at this address 0 1 = No device responding at this address 0 0 = No device responding at this address	RO
4.8.13:12	Reserved	Value always 0	RO
4.8.11	Transmit fault	1 = Fault condition on transmit path 0 = No fault condition on transmit path	RO/LH

Table 45–168—PHY XS status 2 register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
4.8.10	Receive fault	1 = Fault condition on receive path 0 = No fault condition on receive path	RO/LH
4.8.9:0	Reserved	Value always 0	RO

[a]RO = Read only, LH = Latching high

45.2.4.6.1 Device present (4.8.15:14)

When read as <10>, bits 4.8.15:14 indicate that a device is present and responding at this register address. When read as anything other than <10>, bits 4.8.15:14 indicate that no device is present at this register address or that the device is not functioning properly.

45.2.4.6.2 Transmit fault (4.8.11)

When read as a one, bit 4.8.11 indicates that the PHY XS has detected a fault condition on the transmit path. When read as a zero, bit 4.8.11 indicates that the PHY XS has not detected a fault condition on the transmit path. The transmit fault bit shall be implemented with latching high behavior.

The default value for bit 4.8.11 is zero.

45.2.4.6.3 Receive fault (4.8.10)

When read as a one, bit 4.8.10 indicates that the PHY XS has detected a fault condition on the receive path. When read as a zero, bit 4.8.10 indicates that the PHY XS has not detected a fault condition on the receive path. The receive fault bit shall be implemented with latching high behavior.

The default value of bit 4.8.10 is zero.

45.2.4.7 PHY XS package identifier (Registers 4.14 and 4.15)

Registers 4.14 and 4.15 provide a 32-bit value, which may constitute a unique identifier for a particular type of package that the PHY XS is instantiated within. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the package manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A PHY XS may return a value of zero in each of the 32 bits of the PHY XS package identifier.

A non-zero package identifier may be returned by one or more MMDs in the same package. The package identifier may be the same as the device identifier.

The format of the PHY XS package identifier is specified in 22.2.4.3.1.

45.2.4.8 EEE capability (Register 4.20)

This register is used to indicate the capability of the PHY XS to support EEE functions. The assignment of bits in the EEE capability register is shown in Table 45–169.

Table 45–169—EEE capability register bit definitions

Bit(s)	Name	Description	R/W[a]
4.20.15:5	Reserved	Value always 0	RO
4.20.4	PHY XS EEE	1 = EEE is supported for PHY XS 0 = EEE is not supported for PHY XS	RO
4.20.3:1	Reserved	Value always 0	RO
4.20.0	XAUI stop capable	1 = The DTE XS may stop XAUI signals during LPI 0 = XAUI signals not stoppable	RO

[a] RO = Read only

45.2.4.8.1 PHY XS EEE supported (4.20.4)

If the device supports EEE operation for PHY XS as defined in 48.2, this bit shall be set to one.

45.2.4.8.2 XAUI stop capable (4.20.0)

If bit 4.20.0 is set to one, then the PHY XS is indicating that the attached DTE XS is permitted to stop transmitting XAUI signals during LPI. If the bit is set to zero then the PHY XS is indicating that the attached DTE XS is not permitted to stop transmitting XAUI signals during LPI. If the DTE XS does not support EEE capability or is not able to stop the transmit direction XAUI, then this bit has no effect.

45.2.4.9 EEE wake error counter (Register 4.22)

This register is used by PHY XS that support EEE to count wake time faults where the PHY XS fails to complete its normal wake sequence after a period of quiescence for XAUI transmit signals. The fault event to be counted may occur during a refresh or a wake-up. This 16-bit counter shall be reset to all zeros when the EEE wake error counter is read by the management function or upon execution of the PHY XS reset. This counter shall be held at all ones in the case of overflow.

45.2.4.10 10G PHY XGXS lane status register (Register 4.24)

The assignment of bits in the 10G PHY XGXS lane status register is shown in Table 45–170. All the bits in the 10G PHY XGXS lane status register are read only; a write to the 10G PHY XGXS lane status register shall have no effect.

45.2.4.10.1 PHY XGXS transmit lane alignment status (4.24.12)

When read as a one, bit 4.24.12 indicates that the PHY XGXS has synchronized and aligned all four transmit lanes. When read as a zero, bit 4.24.12 indicates that the PHY XGXS has not synchronized and aligned all four transmit lanes.

Table 45–170—10G PHY XGXS lane status register bit definitions

Bit(s)	Name	Description	R/W[a]
4.24.15:13	Reserved	Value always 0	RO
4.24.12	PHY XGXS lane alignment status	1 = PHY XGXS transmit lanes aligned 0 = PHY XGXS transmit lanes not aligned	RO
4.24.11	Pattern testing ability	1 = PHY XGXS is able to generate test patterns 0 = PHY XGXS is not able to generate test patterns	RO
4.24.10	PHY XGXS loopback ability	1 = PHY XGXS has the ability to perform a loopback function 0 = PHY XGXS does not have the ability to perform a loopback function	RO
4.24.9:4	Reserved	Value always 0	RO
4.24.3	Lane 3 sync	1 = Lane 3 is synchronized 0 = Lane 3 is not synchronized	RO
4.24.2	Lane 2 sync	1 = Lane 2 is synchronized 0 = Lane 2 is not synchronized	RO
4.24.1	Lane 1 sync	1 = Lane 1 is synchronized 0 = Lane 1 is not synchronized	RO
4.24.0	Lane 0 sync	1 = Lane 0 is synchronized 0 = Lane 0 is not synchronized	RO

[a]RO = Read only

45.2.4.10.2 Pattern testing ability (4.24.11)

When read as a one, bit 4.24.11 indicates that the 10G PHY XGXS is able to generate test patterns. When read as a zero, bit 4.24.11 indicates that the 10G PHY XGXS is not able to generate test patterns. If the 10G PHY XGXS is able to generate test patterns, then the functionality is controlled using the transmit test-pattern enable bit in register 4.25.

45.2.4.10.3 PHY XS loopback ability (4.24.10)

When read as a one, bit 4.24.10 indicates that the PHY XGXS is able to perform the loopback function. When read as a zero, bit 4.24.10 indicates that the PHY XGXS is not able to perform the loopback function. If a 10G PHY XGXS is able to perform the loopback function, then it is controlled using the PHY XGXS loopback bit 4.0.14.

45.2.4.10.4 Lane 3 sync (4.24.3)

When read as a one, bit 4.24.3 indicates that the 10G PHY XGXS transmit lane 3 is synchronized. When read as a zero, bit 4.24.3 indicates that the 10G PHY XGXS transmit lane 3 is not synchronized.

45.2.4.10.5 Lane 2 sync (4.24.2)

When read as a one, bit 4.24.2 indicates that the 10G PHY XGXS transmit lane 2 is synchronized. When read as a zero, bit 4.24.2 indicates that the 10G PHY XGXS transmit lane 2 is not synchronized.

45.2.4.10.6 Lane 1 sync (4.24.1)

When read as a one, bit 4.24.1 indicates that the 10G PHY XGXS transmit lane 1 is synchronized. When read as a zero, bit 4.24.1 indicates that the 10G PHY XGXS transmit lane 1 is not synchronized.

45.2.4.10.7 Lane 0 sync (4.24.0)

When read as a one, bit 4.24.0 indicates that the 10G PHY XGXS transmit lane 0 is synchronized. When read as a zero, bit 4.24.0 indicates that the 10G PHY XGXS transmit lane 0 is not synchronized.

45.2.4.11 10G PHY XGXS test control register (Register 4.25)

The assignment of bits in the 10G PHY XGXS test control register is shown in Table 45–171. The default value for each bit of the 10G PHY XGXS test control register should be chosen so that the initial state of the device upon power up or reset is a normal operational state without management intervention.

Table 45–171—10G PHY XGXS test control register bit definitions

Bit(s)	Name	Description	R/W[a]
4.25.15:3	Reserved	Value always 0	RO
4.25.2	Receive test-pattern enable	1 = Receive test pattern enabled 0 = Receive test pattern not enabled	R/W
4.25.1:0	Test-pattern select	1 0 1 1 = Reserved 1 0 = Mixed-frequency test pattern 0 1 = Low-frequency test pattern 0 0 = High-frequency test pattern	R/W

[a]RO = Read only, R/W = Read/Write

45.2.4.11.1 10G PHY XGXS test-pattern enable (4.25.2)

When bit 4.25.2 is set to a one, pattern testing is enabled on the receive path. When bit 4.25.2 is set to a zero, pattern testing is disabled on the receive path. Pattern testing is optional, and the ability of the 10G PHY XGXS to generate test patterns is advertised by the pattern testing ability bit in register 4.24. A 10G PHY XGXS that does not support the generation of test patterns shall ignore writes to this bit and always return a value of zero. The default of bit 4.25.2 is zero.

45.2.4.11.2 10G PHY XGXS test-pattern select (4.25.1:0)

The test pattern to be used when pattern testing is enabled using bit 4.25.2 is selected using bits 4.25.1:0. When bits 4.25.1:0 are set to <10>, the mixed-frequency test pattern shall be selected for pattern testing. When bits 4.25.1:0 are set to <01>, the low-frequency test pattern shall be selected for pattern testing. When bits 4.25.1:0 are set to <00>, the high-frequency test pattern shall be selected for pattern testing. The test patterns are defined in Annex 48A.

45.2.4.12 TimeSync PHY XS capability (Register 4.1800)

The TimeSync PHY XS capability register (see Table 45–172) indicates the capability of the PHY XS to report the transmit and receive data delay, stored in registers 4.1801 through 4.1804 and 4.1805 through 4.1808, respectively.

Table 45–172—TimeSync PHY XS capability

Bit(s)	Name	Description	R/W[a]
4.1800.15:2	Reserved	Value always 0	RO
4.1800.1	TimeSync transmit path data delay	1 = PHY XS provides information on transmit path data delay in registers 4.1801 through 4.1804 0 = PHY XS does not provide information on transmit path data delay	RO
4.1800.0	TimeSync receive path data delay	1 = PHY XS provides information on receive path data delay in registers 4.1805 through 4.1808 0 = PHY XS does not provide information on receive path data delay	RO

[a]RO = Read only

45.2.4.13 TimeSync PHY XS transmit path data delay (Registers 4.1801, 4.1802, 4.1803, 4.1804)

The TimeSync PHY XS transmit path data delay register contains the maximum (Registers 4.1801, 4.1802, see Table 45–173) and minimum (Registers 4.1803, 4.1804, see Table 45–173) values of the transmit path data delay. The transmit path data delay is expressed in units of ns. The values contained in these registers are valid when the link is established, as indicated by bit 2 in Register 1.1 (see 45.2.1.2.4).

Table 45–173—TimeSync PHY XS transmit path data delay register

Bit(s)	Name	Description	R/W[a]
4.1801.15:0	Maximum PHY XS transmit path data delay, lower	PHY_XS_delay_TX_max [15:0]	RO, MW
4.1802.15:0	Maximum PHY XS transmit path data delay, upper	PHY_XS_delay_TX_max [31:16]	RO, MW
4.1803.15:0	Minimum PHY XS transmit path data delay, lower	PHY_XS_delay_TX_min [15:0]	RO, MW
4.1804.15:0	Minimum PHY XS transmit path data delay, upper	PHY_XS_delay_TX_min [31:16]	RO, MW

[a]RO = Read only, MW = Multi-word

45.2.4.14 TimeSync PHY XS receive path data delay (Registers 4.1805, 4.1806, 4.1807, 4.1808)

The TimeSync PHY XS receive path data delay register contains the maximum (Registers 4.1805, 4.1806, see Table 45–174) and minimum (Registers 4.1807, 4.1808, see Table 45–174) values of the receive path data delay. The receive path data delay is expressed in units of ns. The values contained in these registers are valid when the link is established, as indicated by bit 2 in Register 1.1 (see 45.2.1.2.4).

Table 45–174—TimeSync PHY XS receive path data delay register

Bit(s)	Name	Description	R/W[a]
4.1805.15:0	Maximum PHY XS receive path data delay, lower	PHY_XS_delay_RX_max [15:0]	RO, MW
4.1806.15:0	Maximum PHY XS receive path data delay, upper	PHY_XS_delay_RX_max [31:16]	RO, MW
4.1807.15:0	Minimum PHY XS receive path data delay, lower	PHY_XS_delay_RX_min [15:0]	RO, MW
4.1808.15:0	Minimum PHY XS receive path data delay, upper	PHY_XS_delay_RX_min [31:16]	RO, MW

[a]RO = Read only, MW = Multi-word

45.2.5 DTE XS registers

The assignment of registers in the DTE XS is shown in Table 45–175.

Table 45–175—DTE XS registers

Register address	Register name	Subclause
5.0	DTE XS control 1	45.2.5.1
5.1	DTE XS status 1	45.2.5.2
5.2, 5.3	DTE XS device identifier	45.2.5.3
5.4	DTE XS speed ability	45.2.5.4
5.5, 5.6	DTE XS devices in package	45.2.5.5
5.7	Reserved	
5.8	DTE XS status 2	45.2.5.6
5.9 through 5.13	Reserved	
5.14, 5.15	DTE XS package identifier	45.2.5.7
5.16 through 5.19	Reserved	
5.20	EEE capability	45.2.5.8
5.21	Reserved	
5.22	EEE wake error counter	45.2.5.9
5.23	Reserved	
5.24	10G DTE XGXS lane status	45.2.5.10
5.25	10G DTE XGXS test control	45.2.5.11

Table 45–175—DTE XS registers *(continued)*

Register address	Register name	Subclause
5.26 through 5.1799	Reserved	
5.1800	TimeSync DTE XS capability	45.2.5.12
5.1801 through 5.1804	TimeSync DTE XS transmit path data delay	45.2.5.13
5.1805 through 5.1808	TimeSync DTE XS receive path data delay	45.2.5.14
5.1809 through 5.32767	Reserved	
5.32 768 through 5.65 535	Vendor specific	

45.2.5.1 DTE XS control 1 register (Register 5.0)

The assignment of bits in the DTE XS control 1 register is shown in Table 45–176. The default value for each bit of the DTE XS control 1 register should be chosen so that the initial state of the device upon power up or reset is a normal operational state without management intervention.

Table 45–176—DTE XS control 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
5.0.15	Reset	1 = DTE XS reset 0 = Normal operation	R/W SC
5.0.14	Loopback	1 = Enable loopback mode 0 = Disable loopback mode	R/W
5.0.13	Speed selection	1 = Operation at 10 Gbp/s and above 0 = Unspecified	R/W
5.0.12	Reserved	Value always 0	RO
5.0.11	Low power	1 = Low-power mode 0 = Normal operation	R/W
5.0.10	Clock stop enable	1 = The DTE XS may stop the clock during LPI 0 = Clock not stoppable	R/W
5.0.9	XAUI stop enable	1 = The DTE XS may stop XAUI signals during LPI 0 = XAUI not stoppable	R/W
5.0.8:7	Reserved	Value always 0	RO
5.0.6	Speed selection	1 = Operation at 10 Gb/s and above 0 = Unspecified	R/W

Table 45–176—DTE XS control 1 register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
5.0.5:2	Speed selection	5 4 3 2 1 x x x = Reserved x 1 x x = Reserved x x 1 x = Reserved 0 0 0 1 = Reserved 0 0 0 0 = 10 Gb/s	R/W
5.0.1:0	Reserved	Value always 0	RO

[a]R/W = Read/Write, SC = Self-clearing, RO = Read only

45.2.5.1.1 Reset (5.0.15)

Resetting a DTE XS is accomplished by setting bit 5.0.15 to a one. This action shall set all DTE XS registers to their default states. As a consequence, this action may change the internal state of the DTE XS and the state of the physical link. This action may also initiate a reset in any other MMDs that are instantiated in the same package. This bit is self-clearing, and a DTE XS shall return a value of one in bit 5.0.15 when a reset is in progress and a value of zero otherwise. A DTE XS is not required to accept a write transaction to any of its registers until the reset process is completed. The reset process shall be completed within 0.5 s from the setting of bit 5.0.15. During a reset, a DTE XS shall respond to reads to register bits 5.0.15 and 5.8.15:14. All other register bits should be ignored.

NOTE—This operation may interrupt data communication.

45.2.5.1.2 Loopback (5.0.14)

The DTE XS shall be placed in a loopback mode of operation when bit 5.0.14 is set to a one. When bit 5.0.14 is set to a one, the DTE XS shall accept data on the transmit path and return it on the receive path. For 10 Gb/s operation, the specific behavior of a DTE XS during loopback is specified in 48.3.3.

The default value of bit 5.0.14 is zero.

NOTE—The signal path through the DTE XS that is exercised in the loopback mode of operation is implementation specific, but it is recommended that the signal path encompass as much of the DTE XS circuitry as is practical. The intention of providing this loopback mode of operation is to permit a diagnostic or self-test function to perform the transmission and reception of a PDU, thus testing the transmit and receive data paths. Other loopback signal paths may be enabled using loopback controls within other MMDs.

45.2.5.1.3 Low power (5.0.11)

A DTE XS may be placed into a low-power mode by setting bit 5.0.11 to a one. This action may also initiate a low-power mode in any other MMDs that are instantiated in the same package. The low-power mode is exited by resetting the DTE XS. The behavior of the DTE XS in transition to and from the low-power mode is implementation specific and any interface signals should not be relied upon. While in the low-power mode, the device shall, as a minimum, respond to management transactions necessary to exit the low-power mode. The default value of bit 5.0.11 is zero.

45.2.5.1.4 Clock stop enable (5.0.10)

If bit 5.0.10 is set to 1, then the DTE XS may stop the receive xMII clock while it is signaling LPI, otherwise it shall keep the clock active. If the DTE XS does not support EEE capability or is not able to stop the receive clock, then this bit has no effect (see 46.3.2.4).

45.2.5.1.5 XAUI stop enable (5.0.9)

If bit 5.0.9 is set to 1, then the DTE XS may stop signaling on the XAUI in the transmit direction during LPI, otherwise the DTE XS shall keep the XAUI signals active. If the DTE XS does not support EEE capability or is not able to stop the transmit path XAUI signals, then this bit has no effect.

45.2.5.1.6 Speed selection (5.0.13, 5.0.6, 5.0.5:2)

Speed selection bits 5.0.13 and 5.0.6 shall both be written as a one. Any attempt to change the bits to an invalid setting shall be ignored. These two bits are set to one in order to make them compatible with Clause 22.

The speed of the DTE XS may be selected using bits 5 through 2. The speed abilities of the DTE XS are advertised in the DTE XS speed ability register. A DTE XS may ignore writes to the DTE XS speed selection bits that select speeds it has not advertised in the DTE XS speed ability register. It is the responsibility of the STA entity to ensure that mutually acceptable speeds are applied consistently across all the MMDs on a particular PHY.

The DTE XS speed selection defaults to a supported ability.

45.2.5.2 DTE XS status 1 register (Register 5.1)

The assignment of bits in the DTE XS status 1 register is shown in Table 45–177. All the bits in the DTE XS status 1 register are read only; a write to the DTE XS status 1 register shall have no effect.

Table 45–177—DTE XS status 1 register bit definitions

Bit(s)	Name	Description	R/W[a]
5.1.15:12	Reserved	Value always 0	RO
5.1.11	Tx LPI received	1 = Tx PCS has received LPI 0 = LPI not received	RO/LH
5.1.10	Rx LPI received	1 = Rx PCS has received LPI 0 = LPI not received	RO/LH
5.1.9	Tx LPI indication	1 = Tx PCS is currently receiving LPI 0 = PCS is not currently receiving LPI	RO
5.1.8	Rx LPI indication	1 = Rx PCS is currently receiving LPI 0 = PCS is not currently receiving LPI	RO
5.1.7	Fault	1 = Fault condition detected 0 = No fault condition detected	RO
5.1.6	Clock stop capable	1 = The MAC may stop the clock during LPI 0 = Clock not stoppable	RO
5.1.5:3	Reserved	Value always 0	RO
5.1.2	DTE XS receive link status	1 = The DTE XS receive link is up 0 = The DTE XS receive link is down	RO/LL

Table 45–177—DTE XS status 1 register bit definitions *(continued)*

Bit(s)	Name	Description	R/W[a]
5.1.1	Low-power ability	1 = DTE XS supports low-power mode 0 = DTE XS does not support low-power mode	RO
5.1.0	Reserved	Value always 0	RO

[a]RO = Read only, LL = Latching low, LH = Latching high

45.2.5.2.1 Transmit LPI received (5.1.11)

When read as a one, bit 5.1.11 indicates that the transmit DTE XS has received LPI signaling one or more times since the register was last read. When read as a zero, bit 5.1.11 indicates that the DTE XS has not received LPI signaling. This bit shall be implemented with latching high behavior.

45.2.5.2.2 Receive LPI received (5.1.10)

When read as a one, bit 5.1.10 indicates that the receive DTE XS has received LPI signaling one or more times since the register was last read. When read as a zero, bit 5.1.10 indicates that the DTE XS has not received LPI signaling. This bit shall be implemented with latching high behavior.

45.2.5.2.3 Transmit LPI indication (5.1.9)

When read as a one, bit 5.1.9 indicates that the transmit DTE XS is currently receiving LPI signals. When read as a zero, bit 5.1.9 indicates that the DTE XS is not currently receiving LPI signals. The behavior if read during a state transition is undefined.

45.2.5.2.4 Receive LPI indication (5.1.8)

When read as a one, bit 5.1.8 indicates that the receive DTE XS is currently receiving LPI signals. When read as a zero, bit 5.1.8 indicates that the DTE XS is not currently receiving LPI signals. The behavior if read during a state transition is undefined.

45.2.5.2.5 Fault (5.1.7)

When read as a one, bit 5.1.7 indicates that the DTE XS has detected a fault condition on either the transmit or receive paths. When read as a zero, bit 5.1.7 indicates that the DTE XS has not detected a fault condition. Bit 5.1.7 is set to a one when either of the fault bits (5.8.11, 5.8.10) located in register 5.8 are set to a one.

45.2.5.2.6 Clock stop capable (5.1.6)

If bit 5.1.6 is set to one, then the DTE XS is indicating that the attached RS is permitted to stop the transmit xMII clock while it is signaling LPI. If the bit is set to zero, then the DTE XS is indicating that the attached RS is not permitted to stop the transmit xMII clock while it is signaling LPI. If the RS does not support EEE capability or is not able to stop the transmit direction xMII clock, then this bit has no effect (see 46.3.2.4).

45.2.5.2.7 DTE XS receive link status (5.1.2)

When read as a one, bit 5.1.2 indicates that the DTE XS receive link is aligned. When read as a zero, bit 5.1.2 indicates that the DTE XS receive link is not aligned. The receive link status bit shall be implemented with latching low behavior.

For 10 Gb/s operation, this bit is a latching low version of bit 5.24.12.

45.2.5.2.8 Low-power ability (5.1.1)

When read as a one, bit 5.1.1 indicates that the DTE XS supports the low-power feature. When read as a zero, bit 5.1.1 indicates that the DTE XS does not support the low-power feature. If a DTE XS supports the low-power feature then it is controlled using the low-power bit in the DTE XS control register.

45.2.5.3 DTE XS device identifier (Registers 5.2 and 5.3)

Registers 5.2 and 5.3 provide a 32-bit value, which may constitute a unique identifier for a DTE XS. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the device manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A DTE XS may return a value of zero in each of the 32 bits of the DTE XS device identifier.

The format of the DTE XS device identifier is specified in 22.2.4.3.1.

45.2.5.4 DTE XS speed ability (Register 5.4)

The assignment of bits in the DTE XS speed ability register is shown in Table 45–178.

Table 45–178— DTE XS speed ability register bit definitions

Bit(s)	Name	Description	R/W[a]
5.4.15:1	Reserved for future speeds	Value always 0	RO
5.4.0	10G capable	1 = DTE XS is capable of operating at 10 Gb/s 0 = DTE XS is not capable of operating at 10 Gb/s	RO

[a]RO = Read only

45.2.5.4.1 10G capable (5.4.0)

When read as a one, bit 5.4.0 indicates that the DTE XS is able to operate at a data rate of 10 Gb/s. When read as a zero, bit 5.4.0 indicates that the DTE XS is not able to operate at a data rate of 10 Gb/s.

45.2.5.5 DTE XS devices in package (Registers 5.5 and 5.6)

The DTE XS devices in package registers are defined in Table 45–2.

45.2.5.6 DTE XS status 2 register (Register 5.8)

The assignment of bits in the DTE XS status 2 register is shown in Table 45–179. All the bits in the DTE XS status 2 register are read only; a write to the DTE XS status 2 register shall have no effect.

45.2.5.6.1 Device present (5.8.15:14)

When read as <10>, bits 5.8.15:14 indicate that a device is present and responding at this register address. When read as anything other than <10>, bits 5.8.15:14 indicate that no device is present at this register address or that the device is not functioning properly.

Table 45–179—DTE XS status 2 register bit definitions

Bit(s)	Name	Description	R/W[a]
5.8.15:14	Device present	15 14 1 0 = Device responding at this address 1 1 = No device responding at this address 0 1 = No device responding at this address 0 0 = No device responding at this address	RO
5.8.13:12	Reserved	Value always 0	RO
5.8.11	Transmit fault	1 = Fault condition on transmit path 0 = No fault condition on transmit path	RO/LH
5.8.10	Receive fault	1 = Fault condition on receive path 0 = No fault condition on receive path	RO/LH
5.8.9:0	Reserved	Value always 0	RO

[a]RO = Read only, LH = Latching high

45.2.5.6.2 Transmit fault (5.8.11)

When read as a one, bit 5.8.11 indicates that the DTE XS has detected a fault condition on the transmit path. When read as a zero, bit 5.8.11 indicates that the DTE XS has not detected a fault condition on the transmit path. The transmit fault bit shall be implemented with latching high behavior.

The default value of bit 5.8.11 is zero.

45.2.5.6.3 Receive fault (5.8.10)

When read as a one, bit 5.8.10 indicates that the DTE XS has detected a fault condition on the receive path. When read as a zero, bit 5.8.10 indicates that the DTE XS has not detected a fault condition on the receive path. The receive fault bit shall be implemented with latching high behavior.

The default value of bit 5.8.10 is zero.

45.2.5.7 DTE XS package identifier (Registers 5.14 and 5.15)

Registers 5.14 and 5.15 provide a 32-bit value, which may constitute a unique identifier for a particular type of package that the DTE XS is instantiated within. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the package manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A DTE XS may return a value of zero in each of the 32 bits of the DTE XS package identifier.

A non-zero package identifier may be returned by one or more MMDs in the same package. The package identifier may be the same as the device identifier.

The format of the DTE XS package identifier is specified in 22.2.4.3.1.

45.2.5.8 EEE capability (Register 5.20)

This register is used to indicate the capability of the DTE XS to support EEE functions. The assignment of bits in the EEE capability register is shown in Table 45–180.

Table 45–180—EEE capability register bit definitions

Bit(s)	Name	Description	R/W[a]
5.20.15:5	Reserved	Value always 0	RO
5.20.4	DTE XS EEE	1 = EEE is supported for DTE XS 0 = EEE is not supported for DTE XS	RO
5.20.3:1	Reserved	Value always 0	RO
5.20.0	XAUI stop capable	1 = The PHY XS may stop XAUI signals during LPI 0 = XAUI signals not stoppable	RO

[a] RO = Read only

45.2.5.8.1 PHY XS EEE supported (5.20.4)

If the device supports EEE operation for DTE XS as defined in 48.2, this bit shall be set to one.

45.2.5.8.2 XAUI stop capable (5.20.0)

If bit 5.20.0 is set to one, then the DTE XS is indicating that the attached PHY XS is permitted to stop signaling the XAUI in the receive direction during LPI. If the bit is set to zero, then the DTE XS is indicating that the attached PHY XS is not permitted to stop signaling the XAUI in the receive direction during LPI. If the PHY XS does not support EEE capability or is not able to stop the receive direction XAUI, then this bit has no effect.

45.2.5.9 EEE wake error counter (Register 5.22)

This register is used by DTE XS that support EEE to count wake time faults where the DTE XS fails to complete its normal wake sequence after a period of quiescence for XAUI receive signals. The fault event to be counted may occur during a refresh or a wake-up. This 16-bit counter shall be reset to all zeros when the EEE wake error counter is read by the management function or upon execution of the DTE XS reset. This counter shall be held at all ones in the case of overflow.

45.2.5.10 10G DTE XGXS lane status register (Register 5.24)

The assignment of bits in the 10G DTE XGXS lane status register is shown in Table 45–181. All the bits in the 10G DTE XGXS lane status register are read only; a write to the 10G DTE XGXS lane status register shall have no effect.

45.2.5.10.1 DTE XGXS receive lane alignment status (5.24.12)

When read as a one, bit 5.24.12 indicates that the DTE XGXS has synchronized and aligned all four receive lanes. When read as a zero, bit 5.24.12 indicates that the DTE XGXS has not synchronized and aligned all four receive lanes.

Table 45–181—10G DTE XGXS lane status register bit definitions

Bit(s)	Name	Description	R/W[a]
5.24.15:13	Reserved	Value always 0	RO
5.24.12	DTE XGXS lane alignment status	1 = DTE XGXS receive lanes aligned 0 = DTE XGXS receive lanes not aligned	RO
5.24.11	Pattern testing ability	1 = DTE XGXS is able to generate test patterns 0 = DTE XGXS is not able to generate test patterns	RO
5.24.10	Ignored	Value 0 or 1, writes ignored	RO
5.24.9:4	Reserved	Value always 0	RO
5.24.3	Lane 3 sync	1 = Lane 3 is synchronized 0 = Lane 3 is not synchronized	RO
5.24.2	Lane 2 sync	1 = Lane 2 is synchronized 0 = Lane 2 is not synchronized	RO
5.24.1	Lane 1 sync	1 = Lane 1 is synchronized 0 = Lane 1 is not synchronized	RO
5.24.0	Lane 0 sync	1 = Lane 0 is synchronized 0 = Lane 0 is not synchronized	RO

[a]RO = Read only

45.2.5.10.2 Pattern testing ability (5.24.11)

When read as a one, bit 5.24.11 indicates that the 10G DTE XGXS is able to generate test patterns. When read as a zero, bit 5.24.11 indicates that the 10G DTE XGXS is not able to generate test patterns. If the 10G DTE XGXS is able to generate test patterns then the functionality is controlled using the transmit test-pattern enable bit in register 5.25.

45.2.5.10.3 Ignored (5.24.10)

So that a single device can implement either register 4.24 or register 5.24, bit 5.24.10 can return either a one or a zero and should be ignored.

45.2.5.10.4 Lane 3 sync (5.24.3)

When read as a one, bit 5.24.3 indicates that the XGXS receive lane 3 is synchronized. When read as a zero, bit 5.24.3 indicates that the XGXS receive lane 3 is not synchronized.

45.2.5.10.5 Lane 2 sync (5.24.2)

When read as a one, bit 5.24.2 indicates that the XGXS receive lane 2 is synchronized. When read as a zero, bit 5.24.2 indicates that the XGXS receive lane 2 is not synchronized.

45.2.5.10.6 Lane 1 sync (5.24.1)

When read as a one, bit 5.24.1 indicates that the XGXS receive lane 1 is synchronized. When read as a zero, bit 5.24.1 indicates that the XGXS receive lane 1 is not synchronized.

45.2.5.10.7 Lane 0 sync (5.24.0)

When read as a one, bit 5.24.0 indicates that the XGXS receive lane 0 is synchronized. When read as a zero, bit 5.24.0 indicates that the XGXS receive lane 0 is not synchronized.

45.2.5.11 10G DTE XGXS test control register (Register 5.25)

The assignment of bits in the 10G DTE XGXS test control register is shown in Table 45–182. The default value for each bit of the 10G DTE XGXS test control register should be chosen so that the initial state of the device upon power up or reset is a normal operational state without management intervention.

Table 45–182—10G DTE XGXS test control register bit definitions

Bit(s)	Name	Description	R/W[a]
5.25.15:3	Reserved	Value always 0	RO
5.25.2	Transmit test-pattern enable	1 = Transmit test pattern enabled 0 = Transmit test pattern not enabled	R/W
5.25.1:0	Test-pattern select	1 0 1 1 = Reserved 1 0 = Mixed-frequency test pattern 0 1 = Low-frequency test pattern 0 0 = High-frequency test pattern	R/W

[a]RO = Read only, R/W = Read/Write

45.2.5.11.1 10G DTE XGXS test-pattern enable (5.25.2)

When bit 5.25.2 is set to a one, pattern testing is enabled on the transmit path. When bit 5.25.2 is set to a zero, pattern testing is disabled on the transmit path. Pattern testing is optional, and the ability of the 10G DTE XGXS to generate test patterns is advertised by the pattern testing ability bit in register 5.24. A 10G DTE XGXS that does not support the generation of test patterns shall ignore writes to this bit and always return a value of zero. The default of bit 5.25.2 is zero.

45.2.5.11.2 10G DTE XGXS test-pattern select (5.25.1:0)

The test pattern to be used when pattern testing is enabled using bit 5.25.2 is selected using bits 5.25.1:0. When bits 5.25.1:0 are set to <10>, the mixed-frequency test pattern shall be selected for pattern testing. When bits 5.25.1:0 are set to <01>, the low-frequency test pattern shall be selected for pattern testing. When bits 5.25.1:0 are set to <00>, the high-frequency test pattern shall be selected for pattern testing. The test patterns are defined in Annex 48A.

45.2.5.12 TimeSync DTE XS capability (Register 5.1800)

The TimeSync DTE XS capability register (see Table 45–183) indicates the capability of the DTE XS to report the transmit and receive data delay, stored in registers 5.1801 through 5.1804 and 5.1805 through 5.1808, respectively

Table 45–183—TimeSync DTE XS capability

Bit(s)	Name	Description	R/W[a]
5.1800.15:2	Reserved	Value always 0	RO
5.1800.1	TimeSync transmit path data delay	1 = DTE XS provides information on transmit path data delay in registers 5.1801 through 5.1804 0 = DTE XS does not provide information on transmit path data delay	RO
5.1800.0	TimeSync receive path data delay	1 = DTE XS provides information on receive path data delay in registers 5.1805 through 5.1808 0 = DTE XS does not provide information on receive path data delay	RO

[a]RO = Read only

45.2.5.13 TimeSync DTE XS transmit path data delay (Registers 5.1801, 5.1802, 5.1803, 5.1804)

The TimeSync DTE XS transmit path data delay register contains the maximum (Registers 5.1801, 5.1802, see Table 45–184) and minimum (Registers 5.1803, 5.1804, see Table 45–184) values of the transmit path data delay. The transmit path data delay is expressed in units of ns. The values contained in these registers are valid when the link is established, as indicated by bit 2 in Register 1.1 (see 45.2.1.2.4).

Table 45–184—TimeSync DTE XS transmit path data delay register

Bit(s)	Name	Description	R/W[a]
5.1801.15:0	Maximum DTE XS transmit path data delay, lower	DTE_XS_delay_TX_max [15:0]	RO, MW
5.1802.15:0	Maximum DTE XS transmit path data delay, upper	DTE_XS_delay_TX_max [31:16]	RO, MW
5.1803.15:0	Minimum DTE XS transmit path data delay, lower	DTE_XS_delay_TX_min [15:0]	RO, MW
5.1804.15:0	Minimum DTE XS transmit path data delay, upper	DTE_XS_delay_TX_min [31:16]	RO, MW

[a]RO = Read only, MW = Multi-word

45.2.5.14 TimeSync DTE XS receive path data delay (Registers 5.1805, 5.1806, 5.1807, 5.1808)

The TimeSync DTE XS receive path data delay register contains the maximum (Registers 5.1805, 5.1806, see Table 45–185) and minimum (Registers 5.1807, 5.1808, see Table 45–185) values of the receive path data delay. The receive path data delay is expressed in units of ns. The values contained in these registers are valid when the link is established, as indicated by bit 2 in Register 1.1 (see 45.2.1.2.4).

Table 45–185—TimeSync DTE XS receive path data delay register

Bit(s)	Name	Description	R/W[a]
5.1805.15:0	Maximum DTE XS receive path data delay, lower	DTE_XS_delay_RX_max [15:0]	RO, MW
5.1806.15:0	Maximum DTE XS receive path data delay, upper	DTE_XS_delay_RX_max [31:16]	RO, MW
5.1807.15:0	Minimum DTE XS receive path data delay, lower	DTE_XS_delay_RX_min [15:0]	RO, MW
5.1808.15:0	Minimum DTE XS receive path data delay, upper	DTE_XS_delay_RX_min [31:16]	RO, MW

[a]RO = Read only, MW = Multi-word

45.2.6 TC registers

The assignment of registers in the TC MMD is shown in Table 45–186.

Table 45–186—TC registers

Register address	Register name	Subclause
6.0	TC control	45.2.6.1
6.1	Reserved	
6.2, 6.3	TC device identifier	45.2.6.2
6.4	TC speed ability	45.2.6.3
6.5, 6.6	TC devices in package	45.2.6.4
6.7 through 6.13	Reserved	
6.14, 6.15	TC package identifier	45.2.6.5
6.16	10P/2B aggregation discovery control[a]	45.2.6.6
6.17	10P/2B aggregation and discovery status[a]	45.2.6.7
6.18, 6.19, 6.20	10P/2B aggregation discovery code[a]	45.2.6.8
6.21	10P/2B link partner PME aggregate control[a]	45.2.6.9
6.22, 6.23	10P/2B link partner PME aggregate data[a]	45.2.6.10
6.24	10P/2B TC CRC error counter	45.2.6.11
6.25, 6.26	10P/2B TPS-TC coding violations counter	45.2.6.12
6.27	10P/2B TC indications	45.2.6.13
6.28 through 6.1799	Reserved	
6.1800	TimeSync TC capability	45.2.6.14

Table 45–186—TC registers *(continued)*

Register address	Register name	Subclause
6.1801 through 6.1804	TimeSync TC transmit path data delay	45.2.6.15
6.1805 through 6.1808	TimeSync TC receive path data delay	45.2.6.16
6.1809 through 6.32767	Reserved	
6.32 768 through 6.65 535	Vendor specific	

[a]Register is defined only for -O port types and is reserved for -R ports

45.2.6.1 TC control register (Register 6.0)

The assignment of bits in the TC control register is shown in Table 45–187. The default value for each bit of the TC control register has been chosen so that the initial state of the device upon power up or completion of reset is a normal operational state without management intervention.

Table 45–187—TC control register bit definitions

Bit(s)	Name	Description	R/W[a]
6.0.15	Reset	1 = TC reset 0 = Normal operation	R/W SC
6.0.14	Reserved	Value always 0	RO
6.0.13	Speed selection	13 6 1 1 = bits 5:2 select speed 0 x = unspecified x 0 = unspecified	R/W
6.0.12:7	Reserved	Value always 0	RO
6.0.6	Speed selection	13 6 1 1 = bits 5:2 select speed 0 x = unspecified x 0 = unspecified	R/W
6.0.5:2	Speed selection	5 4 3 2 1 x x x = Reserved x 1 x x = Reserved x x 1 x = Reserved 0 0 0 1 = 10PASS-TS/2BASE-TL 0 0 0 0 = Reserved	R/W
6.0.1:0	Reserved	Value always 0	RO

[a]R/W = Read/Write, SC = Self-clearing, RO = Read only

45.2.6.1.1 Reset (6.0.15)

Resetting a TC is accomplished by setting bit 6.0.15 to a one. This action shall set all TC registers to their default states. As a consequence, this action may change the internal state of the TC and the state of the physical link. This action may also initiate a reset in any other MMDs that are instantiated in the same

package. This bit is self-clearing, and a TC shall return a value of one in bit 6.0.15 when a reset is in progress; otherwise, it shall return a value of zero. A TC is not required to accept a write transaction to any of its registers until the reset process is completed. The control and management interface shall be restored to operation within 0.5 s from the setting of bit 6.0.15. During a reset, a TC shall respond to reads from register bit 6.0.15.

NOTE—This operation may interrupt data communication. The data path of a TC, depending on type and temperature, may take many seconds to run at optimum error ratio after exiting from reset.

45.2.6.1.2 Speed selection (6.0.13, 6.0.6, 6.0.5:2)

Speed selection bits 6.0.13 and 6.0.6 shall both be written as a one. Any attempt to change the bits to an invalid setting shall be ignored. These two bits are set to one in order to make them compatible with Clause 22.

The operating mode of the TC may be selected using bits 5 through 2. The abilities of the TC are advertised in the TC speed ability register. A TC may ignore writes to the TC speed selection bits that select speeds it has not advertised in the TC speed ability register. It is the responsibility of the STA entity to ensure that mutually acceptable speeds are applied consistently across all the MMDs on a particular PHY.

The speed selection bits 6.0.5:2, when set to 0001, select the use of the 10PASS-TS and 2BASE-TL TC.

The TC speed selection defaults to a supported ability.

45.2.6.2 TC device identifier (Registers 6.2 and 6.3)

Registers 6.2 and 6.3 provide a 32-bit value, which may constitute a unique identifier for a particular type of TC. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the device manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A TC may return a value of zero in each of the 32 bits of the TC device identifier.

The format of the TC device identifier is specified in 22.2.4.3.1.

45.2.6.3 TC speed ability (Register 6.4)

The assignment of bits in the TC speed ability register is shown in Table 45–188.

Table 45–188—TC speed ability register bit definitions

Bit(s)	Name	Description	R/W[a]
6.4.15:1	Reserved for future speeds	Value always 0	RO
6.4.1	10PASS-TS/2BASE-TL capable	1 = TC is capable of operating as the 10P/2B TC 0 = TC is not capable of operating as the 10P/2B TC	RO
6.4.0	Reserved	Value always 0	RO

[a]RO = Read only

45.2.6.3.1 10PASS-TS/2BASE-TL capable (6.4.1)

When read as a one, this bit indicates that the TC is able to operate as the 10PASS-TS/2BASE-TL TC, as specified in Clause 61.

45.2.6.4 TC devices in package registers (Registers 6.5, 6.6)

The TC devices in package registers are defined in Table 45–2.

45.2.6.5 TC package identifier registers (Registers 6.14, 6.15)

Registers 6.14 and 6.15 provide a 32-bit value, which may constitute a unique identifier for a particular type of package that the TC MMD is instantiated within. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the package manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A TC may return a value of zero in each of the 32 bits of the package identifier.

A non-zero package identifier may be returned by one or more MMDs in the same package. The package identifier may be the same as the device identifier.

The format of the package identifier is specified in 22.2.4.3.1.

45.2.6.6 10P/2B aggregation discovery control register (Register 6.16)

The 10P/2B aggregation discovery control register allows the STA of an -O port to determine the aggregation capabilities of an -R link-partner.

The 10P/2B aggregation discovery control register shall be implemented as a unique register for each TC MMD in a package. For example, a package implementing four PHYs would have four independent instances of the 10P/2B aggregation discovery control register, accessed by a read or write to each PHY.

For information on the encoding of this function on the physical link, see 61.4.7.

This register is defined for -O port subtypes only. The register bit definitions for the 10P/2B aggregation discovery control register are shown in Table 45–189.

Table 45–189—10P/2B aggregation discovery control register bit definitions

Bit(s)	Name	Description	R/W[a]
6.16.15:2	Reserved	Value always 0	RO
6.16.1:0	Discovery operation	01 = Ready (default) 00 = Set if clear 11 = Clear if same 10 = Get	R/W

[a]RO = Read only, R/W = Read/Write

45.2.6.6.1 Discovery operation (6.16.1:0)

The Discovery operation bits are used to query and manipulate the remote discovery register. The remote discover register is not a Clause 45 object, but a variable of the PME aggregation PCS function on -R ports.

The Discovery operation makes use of G.994.1 handshaking messages, therefore valid only when the link status is down (i.e., neither Initializing nor Up). Attempts to perform an operation while the link is Initializing or Up shall be ignored.

The default state of these bits is "Ready". The bits shall indicate "Ready" any time the PME aggregation function is capable of performing an operation on the remote discovery register. If PAF is not supported, the discovery operation bits shall indicate "Ready" and ignore writes. These bits shall return to the "Ready" state upon MMD Reset.

If the STA sets the bits to "Get," the PME aggregation function queries the remote discovery register and returns its contents to the aggregation discovery code register.

If the STA sets the bits to "Set if clear," the PME aggregation function passes a message to the -R PCS instructing it to set the remote discovery register to the contents of the aggregation discovery code register, but only if the remote discovery register is all zeros.

If the STA sets the bits to "Clear if same," the PME aggregation function passes a message to the -R PCS instructing it to clear the remote discovery register, but only if the contents of the remote discovery register currently match the contents of the aggregation discovery code register.

While the requested operation is in progress, the PHY maintains the operation value in the bits. After the operation is complete, the PHY shall set the bits to indicate "Ready". If the operation does not complete within a 255 second time-out, the discovery operation result bit (6.17.0) will be set to "1" (operation unsuccessful), and the discovery operation bits will be set to "Ready".

45.2.6.7 10P/2B aggregation and discovery status register (Register 6.17)

The 10P/2B aggregation and discovery status register is defined for -O port subtypes only.

The assignment of bits in the 10P/2B aggregation and discovery status register is shown in Table 45–190.

Table 45–190—10P/2B aggregation and discovery status register bit definitions

Bit(s)	Name	Description	R/W[a]
6.17.15:2	Reserved	Value always 0	RO
6.17.1	Link partner aggregate operation result	1 = operation unsuccessful 0 = operation completed successfully (default)	RO, LH
6.17.0	Discovery operation result	1 = operation unsuccessful 0 = discovery operation completed successfully (default)	RO, LH

[a]RO = Read only, LH = Latching high

45.2.6.7.1 Link partner aggregate operation result (6.17.1)

When a link partner aggregate operation is complete, the PHY sets this bit to indicate the result of the operation. A "1" indicates that the operation could not be completed. This may be for a variety of reasons:

a) PMA/PMD link status is initializing or up.
b) The link partner is not present or not responding.

If PAF is not supported, this bit shall remain set to zero.

45.2.6.7.2 Discovery operation result (6.17.0)

When a discovery operation is complete, the PHY sets this bit to indicate the result of the operation. A "1" indicates that the operation could not be completed. This may be for a variety of reasons:

a) PMA/PMD link status is initializing or up.
b) A "Set if clear" operation was requested but the remote discovery register was not clear.
c) A "Clear if same" operation was requested but the remote discovery register did not match the aggregation discovery code register.
d) The link partner is not present or not responding.

If PAF is not supported, this bit shall read as zero.

45.2.6.8 10P/2B aggregation discovery code (Registers 6.18, 6.19, 6.20)

The 10P/2B aggregation discovery code registers store the value of the remote_discovery_register exchanged with the -R link partner.

This register is defined for -O port subtypes only.

These registers shall be implemented as unique registers for each TC MMD in a package. For example, a package implementing four TCs would have four independent instances of the 10P/2B aggregation discovery code registers, accessed by a read or write to each PHY.

For information on the encoding of this function on the physical link, please see 61.4.7.

The assignment of bits for the 10P/2B aggregation discovery code registers are shown in Table 45–191.

Table 45–191—10P/2B aggregation discovery code bit definitions

Bit(s)	Name	Description	R/W[a]
6.18.15:0	Code [47:32]	The two most significant octets of the aggregation discovery code	R/W
6.19.15:0	Code [31:16]	The two middle octets of the aggregation discovery code	R/W
6.20.15:0	Code [15:0]	The two least significant octets of the aggregation discovery code	R/W

[a]R/W = Read/Write

45.2.6.9 10P/2B link partner PME aggregate control register (Register 6.21)

The 10P/2B link partner PME aggregate control register allows the STA of an -O port to read and write the remote PME_Aggregate_register (see 61.2.2.8.3).

The 10P/2B link partner PME aggregate control register shall be implemented as a unique register for each TC MMD in a package. For example, a package implementing four TCs would have four independent instances of the 10P/2B link partner PME aggregate control register, accessed by a read or write to each PHY.

This register is defined for -O port subtypes only.

The register bit definitions for the 10P/2B link partner PME aggregate control register are shown in Table 45–192.

Table 45–192—10P/2B link partner PME aggregate control register bit definitions

Bit(s)	Name	Description	R/W[a]
6.21.15:2	Reserved	Value always 0	RO
6.21.1:0	Link partner aggregate operation	01 = Ready (default) 00 = Set 11 = invalid 10 = Get	R/W

[a]RO = Read only, R/W = Read/Write

45.2.6.9.1 Link partner aggregate operation (6.21.1:0)

The Link partner aggregate operation bits are used to query and manipulate the remote PME_Aggregate_register. This operation makes use of G.994.1 handshaking messages and therefore must be performed only when the link status is down (i.e., neither Initializing nor Up). Attempts to perform an operation while the link is Initializing or Up shall be ignored.

The default state of these bits is "Ready." The bits shall indicate "Ready" any time the PME aggregation function is capable of performing an operation on the remote PME_Aggregate_register. If PAF is not supported, the link partner aggregate operation bits shall indicate "Ready" ignore writes. These bits shall return to the "Ready" state upon MMD Reset.

If the STA sets the bits to "Get," the PME aggregation function queries the remote PME_Aggregate_register and returns its contents to the 10P/2B link partner PME aggregate data register (see 45.2.6.10).

If the STA sets the bits to "Set," the PME aggregation function passes a message to the -R PCS instructing it to set the bit location in the remote PME_Aggregate_register corresponding to the TC on which the message was received to the contents of bit 0 of the 10P/2B link partner PME aggregate data register.

While the requested operation is in progress, the PHY maintains the operation value in the bits. After the operation is complete, the PHY shall set the bits to indicate "Ready".

45.2.6.10 10P/2B link partner PME aggregate data (Registers 6.22, 6.23)

The 10P/2B link partner PME aggregate data registers store the data for the link partner aggregate operation. This register either contains the result of a "Get" operation, the data sent in a "Set" operation, or all zeros following an MMD reset.

These registers are defined for -O port subtypes only.

These registers shall be implemented as unique registers for each TC MMD in a package. For example, a package implementing four TCs would have four independent instances of the registers, accessed by a read or write to each PHY.

The assignment of bits for the 10P/2B link partner PME aggregate data registers are shown in Table 45–193.

Table 45–193—10P/2B link partner PME aggregate data registers bit definitions

Bit(s)	Name	Description	R/W[a]
6.22.15:0	Data[31:16]	The two most significant octets of the link partner PME aggregate data	R/W
6.23.15:0	Data[15:0]	The two least significant octets of the link partner PME aggregate data	R/W

[a]R/W = Read/Write

45.2.6.11 10P/2B TC CRC error register (Register 6.24)

The 10P/2B TC CRC error register is a 16 bit counter that contains the number of TC frames received with the TC_CRC_error primitive asserted, defined in 61.2.3. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the MMD reset. These bits shall be held at all ones in the case of overflow. The assignment of bits in the 10P/2B TC CRC error register are shown in Table 45–194.

Table 45–194—10P/2B TC CRC error register bit definitions

Bit(s)	Name	Description	R/W[a]
6.24.15:0	CRC errors[15:0]	The bytes of the counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.6.12 10P/2B TPS-TC coding violations counter (Registers 6.25, 6.26)

The 10P/2B TPS-TC coding violations counter is a 32-bit counter that contains the number of 64/64-octet encapsulation errors, defined in 61.3.3.1. This counter increments for each 64/65-octet received with the TC_coding_error signal asserted. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the MMD reset. These bits shall be held at all ones in the case of overflow. The assignment of bits in the 10P/2B TPS-TC coding violations counter are shown in Table 45–195.

Table 45–195—10P/2B TPS-TC coding violations counter bit definitions

Bit(s)	Name	Description	R/W[a]
6.25.15:0	Coding violations[31:16]	The high order bytes of the counter	RO, MW
6.26.15:0	Coding violations[15:0]	The low order bytes of the counter	RO, MW

[a]RO = Read only, MW = Multi-word

45.2.6.13 10P/2B TC indications register (Register 6.27)

The 10P/2B TC indications register reflects the state of the TC sync detect state diagram and the state of the link partner TC sync detect state diagram (if present) (see 61.3.3.5). The assignment of bits in the 10P/2B TC indications register is shown in Table 45–196.

Table 45–196—10P/2B TC indications register bit definitions

Bit(s)	Name	Description	R/W[a]
6.27.15:9	Reserved	Value always 0	RO
6.27.8	Local TC synchronized	1 = TC_synchronized is TRUE 0 = TC_synchronized is FALSE	RO
6.27.7:1	Reserved	Value always 0	RO
6.27.0	Remote TC synchronized	1 = remote_TC_out_of_sync is FALSE 0 = remote_TC_out_of_sync is TRUE	RO

[a]RO = Read only

45.2.6.13.1 Local TC synchronized (6.27.8)

This bit is read as a one when the TC_synchronized variable in the TC sync detect state diagram is TRUE (see 61.3.3.8). In all other cases, this bit is read as zero.

45.2.6.13.2 Remote TC synchronized (6.27.0)

This bit is read as a one when the remote_TC_out_of_sync variable in the link partner TC sync detect state diagram is FALSE (see 61.3.3.8). In all other cases, this bit is read as zero.

45.2.6.14 TimeSync TC capability (Register 6.1800)

The TimeSync TC capability register (see Table 45–197) indicates the capability of the TC to report the transmit and receive data delay, stored in registers 6.1801 through 6.1804 and 6.1805 through 6.1808, respectively.

Table 45–197—TimeSync TC capability

Bit(s)	Name	Description	R/W[a]
6.1800.15:2	Reserved	Value always 0	RO
6.1800.1	TimeSync transmit path data delay	1 = TC provides information on transmit path data delay in registers 6.1801 through 6.1804 0 = TC does not provide information on transmit path data delay	RO
6.1800.0	TimeSync receive path data delay	1 = TC provides information on receive path data delay in registers 6.1805 through 6.1808 0 = TC does not provide information on receive path data delay	RO

[a]RO = Read only

45.2.6.15 TimeSync TC transmit path data delay (Registers 6.1801, 6.1802, 6.1803, 6.1804)

The TimeSync TC transmit path data delay register contains the maximum (Registers 6.1801, 6.1802, see Table 45–198) and minimum (Registers 6.1803, 6.1804, see Table 45–198) values of the transmit path data delay. The transmit path data delay is expressed in units of ns. The values contained in these registers are valid when the link is established, as indicated by bit 2 in Register 1.1 (see 45.2.1.2.4).

Table 45–198—TimeSync TC transmit path data delay register

Bit(s)	Name	Description	R/W[a]
6.1801.15:0	Maximum TC transmit path data delay, lower	TC_delay_TX_max [15:0]	RO, MW
6.1802.15:0	Maximum TC transmit path data delay, upper	TC_delay_TX_max [31:16]	RO, MW
6.1803.15:0	Minimum TC transmit path data delay, lower	TC_delay_TX_min [15:0]	RO, MW
6.1804.15:0	Minimum TC transmit path data delay, upper	TC_delay_TX_min [31:16]	RO, MW

[a]RO = Read only, MW = Multi-word

45.2.6.16 TimeSync TC receive path data delay (Registers 6.1805, 6.1806, 6.1807, 6.1808)

The TimeSync TC receive path data delay register contains the maximum (Registers 6.1805, 6.1806, see Table 45–199) and minimum (Registers 6.1807, 6.1808, see Table 45–199) values of the receive path data delay. The receive path data delay is expressed in units of ns. The values contained in these registers are valid when the link is established, as indicated by bit 2 in Register 1.1 (see 45.2.1.2.4).

Table 45–199—TimeSync TC receive path data delay register

Bit(s)	Name	Description	R/W[a]
6.1805.15:0	Maximum TC receive path data delay, lower	TC_delay_RX_max [15:0]	RO, MW
6.1806.15:0	Maximum TC receive path data delay, upper	TC_delay_RX_max [31:16]	RO, MW
6.1807.15:0	Minimum TC receive path data delay, lower	TC_delay_RX_min [15:0]	RO, MW
6.1808.15:0	Minimum TC receive path data delay, upper	TC_delay_RX_min [31:16]	RO, MW

[a]RO = Read only, MW = Multi-word

45.2.7 Auto-Negotiation registers

The assignment of registers in the Auto-Negotiation (AN) MMD is shown in Table 45–200.

Table 45–200—Auto-Negotiation MMD registers

Register address	Register name	Subclause
7.0	AN control	45.2.7.1
7.1	AN status	45.2.7.2
7.2, 7.3	AN device identifier	45.2.7.3
7.4	Reserved	
7.5, 7.6	AN devices in package	45.2.7.4
7.7 through 7.13	Reserved	
7.14, 7.15	AN package identifier	45.2.7.5
7.16 through 7.18	AN advertisement	45.2.7.6
7.19 through 7.21	AN LP Base Page ability	45.2.7.7
7.22 through 7.24	AN XNP transmit	45.2.7.8
7.25 through 7.27	AN LP XNP ability	45.2.7.9
7.28 through 7.31	Reserved	
7.32	10GBASE-T AN control	45.2.7.10
7.33	10GBASE-T AN status	45.2.7.11
7.34 through 7.47	Reserved	
7.48	Backplane Ethernet, BASE-R copper status	45.2.7.12
7.49 through 7.59	Reserved	
7.60	EEE advertisement	45.2.7.13

Table 45–200—Auto-Negotiation MMD registers *(continued)*

Register address	Register name	Subclause
7.61	EEE LP ability	45.2.7.14
7.62 through 7.32 767	Reserved	
7.32 768 through 7.65 535	Vendor specific	

45.2.7.1 AN control register (Register 7.0)

The assignment of bits in the AN control register is shown in Table 45–201. The default value for each bit of the AN control register has been chosen so that the initial state of the device upon power up or completion of reset is a normal operational state without management intervention.

Table 45–201—AN control register bit definitions

Bit(s)	Name	Description	R/W[a]
7.0.15	AN reset	1 = AN reset 0 = AN normal operation	R/W SC
7.0.14	Reserved	Value always 0	RO
7.0.13	Extended Next Page control	1 = Extended Next Pages are enabled 0 = Extended Next Pages are disabled	R/W
7.0.12	Auto-Negotiation enable	1 = enable Auto-Negotiation process 0 = disable Auto-Negotiation process	R/W
7.0.11:10	Reserved	Value always 0	RO
7.0.9	Restart Auto-Negotiation	1 = Restart Auto-Negotiation process 0 = Auto-Negotiation in process, disabled, or not supported	R/W SC
7.0.8:0	Reserved	Value always 0	RO

[a]R/W = Read/Write, SC = Self-clearing, RO = Read only

A device that supports multiple port types may implement both Clause 22 control register operation and Clause 45 control register operation. Some control functions have been duplicated in both definitions. The register bits to control these functions are simply echoed in both locations; any reads or writes to these bits behave identically whether made through the Clause 22 location or the Clause 45 location.

45.2.7.1.1 AN reset (7.0.15)

Resetting AN is accomplished by setting bit 7.0.15 to a one. This action shall set all AN registers to their default states. As a consequence, this action may change the internal state of AN and the state of the physical link. This action may also initiate a reset in any other MMDs that are instantiated in the same package. This bit is self-clearing, and AN shall return a value of one in bit 7.0.15 when a reset is in progress and a value of zero otherwise. AN is not required to accept a write transaction to any of its registers until the reset process is complete. The reset process shall be completed within 0.5 s from the setting of bit 7.0.15. During an AN reset, AN shall respond to reads from register bit 7.0.15. All other register bits should be ignored.

The default value for bit 7.0.15 is zero. See 22.2.4.1.1.

NOTE—This operation may interrupt data communication.

45.2.7.1.2 Extended Next Page control (7.0.13)

When bit 7.0.13 is set to one, Extended Next Page(s) are exchanged if the device(s) is/are capable. If a device intends to enable the exchange of Extended Next Page, it shall set bit 7.0.13 to one. A device may choose not to exchange Extended Next Page by setting bit 7.0.13 to zero. The Extended Next Page function is defined in 28.2.3.4. Setting of this bit shall have no effect if Extended Next Page ability bit 7.16.12 is set to zero.

The default value for bit 7.0.13 is one if bit 7.16.12 is one. Otherwise, bit 7.0.13 defaults to zero.

45.2.7.1.3 Auto-Negotiation enable (7.0.12)

The Auto-Negotiation function shall be enabled by setting bit 7.0.12 to a one. If bit 7.0.12 is set to one, then speed selection bits 1.0.13, 1.0.6, and bits 1.0.5:2 in PHY control register 1 and PHY type selection bits 1.7.3:0 in PHY control register 2 shall have no effect on the link configuration, and the Auto-Negotiation process determines the link configuration. If bit 7.0.12 is cleared to zero, then bits 1.0.13, 1.0.6, and bits 1.0.5:2 in control register 1 and type selection bit 1.7.3:0 in control register 2 determines the link configuration regardless of the prior state of the link configuration and the Auto-Negotiation process. Bit 7.0.12 is a copy of bit 0.12 in register 0 if present (see 22.2.4.1.4).

The default value of bit 7.0.12 is one, unless the PHY reports via bit 7.1.3 or 1.3 if present (see 22.2.4.2.12) that it lacks the ability to perform Auto-Negotiation, in which case the default value of bit 7.0.12 is zero (see 22.2.4.1.4).

45.2.7.1.4 Restart Auto-Negotiation (7.0.9)

If the PMA/PMD reports (via bit 7.1.3) that it lacks the ability to perform Auto-Negotiation, or if Auto-Negotiation is disabled, the PMA/PMD shall return a value of zero in bit 7.0.9 and any attempt to write a one to bit 7.0.9 shall be ignored.

Otherwise, the Auto-Negotiation process shall be restarted by setting bit 7.0.9 to one. This bit is self-clearing, and a PMA/PMD shall return a value of one in bit 7.0.9 until the Auto-Negotiation process has been initiated. If Auto-Negotiation was completed prior to this bit being set, the process shall be reinitiated. The Auto-Negotiation process shall not be affected by clearing this bit to zero. Bit 7.0.9 is a copy of 0.9 in register 0, if present (see 22.2.4.1.7).

The default value for 7.0.9 is zero (see 22.2.4.1.7).

45.2.7.2 AN status (Register 7.1)

The assignment of bits in the AN status register is shown in Table 45–202. All the bits in the AN status register are read only; therefore, a write to the AN status register shall have no effect.

Table 45–202—AN status register

Bit(s)	Name	Description	R/W[a]
7.1.15:10	Reserved	Value always 0	RO
7.1.9	Parallel detection fault	1 = A fault has been detected via the parallel detection function. 0 = A fault has not been detected via the parallel detection function.	RO LH
7.1.8	Reserved	Value always 0	RO
7.1.7	Extended Next Page status	1 = Extended Next Page format is used 0 = Extended Next Page is not allowed	RO

Table 45–202—AN status register *(continued)*

Bit(s)	Name	Description	R/W[a]
7.1.6	Page received	1 = A page has been received 0 = A page has not been received	RO LH
7.1.5	Auto-Negotiation complete	1 = Auto-Negotiation process completed 0 = Auto-Negotiation process not completed	RO
7.1.4	Remote fault	1 = remote fault condition detected 0 = no remote fault condition detected	RO LH
7.1.3	Auto-Negotiation ability	1 = PHY is able to perform Auto-Negotiation 0 = PHY is not able to perform Auto-Negotiation	RO
7.1.2	Link status	1 = Link is up 0 = Link is down	RO LL
7.1.1	Reserved	Value always 0	RO
7.1.0	Link partner Auto-Negotiation ability	1 = LP is able to perform Auto-Negotiation 0 = LP is not able to perform Auto-Negotiation	RO

[a]RO = Read only, LH = Latching high, LL = Latching low

45.2.7.2.1 Parallel detection fault (7.1.9)

The parallel detection Fault bit (7.1.9) shall be set to one to indicate that more than one of 1000BASE-KX or 10GBASE-KX4 PMAs have indicated link_status=OK when the autoneg_wait_timer expires. The parallel detection fault bit shall be reset to zero on a read of the AN status register (Register 7.1).

45.2.7.2.2 Extended Next Page status (7.1.7)

When set to one, bit 7.1.7 indicates that both the local device and the link partner have indicated support for Extended Next Page. When set to zero, bit 7.1.7 indicates that Extended Next Page shall not be used. If bit 7.0.13 or bit 7.16.12 are set to zero then this bit shall not be set to one.

45.2.7.2.3 Page received (7.1.6)

The Page received bit (7.1.6) shall be set to one to indicate that a new link codeword has been received and stored in the AN LP Base Page ability registers 7.19 to 7.21 or AN LP XNP ability registers 7.25 to 7.27. The contents of the AN LP Base Page ability registers 7.19 to 7.21 are valid when bit 7.1.6 is set the first time during the Auto-Negotiation. The Page received bit shall be reset to zero on a read of the AN status register (Register 7.1) or if present, the Auto-Negotiation expansion register 6 (see 28.2.4.1.5). This bit is a copy of bit 6.1 in register 6, if present (see 28.2.4.1.5).

45.2.7.2.4 Auto-Negotiation complete (7.1.5)

When read as a one, bit 7.1.5 indicates that the Auto-Negotiation process has been completed, and that the contents of the Auto-Negotiation register 7.16 and 7.19 are valid. When read as a zero, bit 7.1.5 indicates that the Auto-Negotiation process has not been completed, and that the contents of 7.19, 7.22 through 7.27, and 7.33 registers are as defined by the current state of the Auto-Negotiation protocol, or as written for manual configuration. A PMA/PMD shall return a value of zero in bit 7.1.5 if Auto-Negotiation is disabled by clearing bit 7.0.12. A PMA/PMD shall also return a value of zero in bit 7.1.5 if it lacks the ability to perform Auto-Negotiation. Bit 7.1.5 is a copy of bit 1.5 in register 1, if present (see 22.2.4).

45.2.7.2.5 Remote fault (7.1.4)

When read as one, bit 7.1.4 indicates that a remote fault condition has been detected. The type of fault as well as the criteria and method of fault detection is AN specific. The remote fault bit shall be implemented

with a latching function, such that the occurrence of a remote fault causes the bit 7.1.4 to become set and remain set until it is cleared. Bit 7.1.4 shall be cleared each time register 7.1 is read via the management interface, and shall also be cleared by a AN reset. Bit 7.1.4 is a copy of bit 1.4 in register 1, if present (see 22.2.4).

45.2.7.2.6 Auto-Negotiation ability (7.1.3)

When read as a one, bit 7.1.3 indicates that the PMA/PMD has the ability to perform Auto-Negotiation. When read as a zero, bit 7.1.3 indicates that the PMA/PMD lacks the ability to perform Auto-Negotiation. Bit 7.1.3 is a copy of bit 1.3 in register 1, if present (see 22.2.4).

45.2.7.2.7 Link status (7.1.2)

When read as a one, bit 7.1.2 indicates that the PMA/PMD has determined that a valid link has been established. When read as a zero, bit 7.1.2 indicates that the link has been invalid after this bit was last read. Bit 7.1.2 is set to one when the variable link_status equals OK and is cleared to zero when the variable link_status equals FAIL. The link status bit shall be implemented with a latching function, such that the occurrence of a link_status equals FAIL condition causes the link status bit to become cleared and remain cleared until it is read via the management interface. Bit 7.1.2 shall be cleared upon AN reset. This status indication is intended to support the management attribute defined in 30.5.1.1.4, aMediaAvailable.

45.2.7.2.8 Link partner Auto-Negotiation ability (7.1.0)

The link partner Auto-Negotiation ability bit shall be set to one to indicate that the link partner is able to participate in the Auto-Negotiation function. This bit shall be reset to zero if the link partner is not Auto-Negotiation able.

45.2.7.3 Auto-Negotiation device identifier (Registers 7.2 and 7.3)

Registers 7.2 and 7.3 provide a 32-bit value, which may constitute a unique identifier for an Auto-Negotiation. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the device manufacturer by the IEEE, plus a six-bit model number, plus a 4-bit revision number. An Auto-Negotiation may return a value of zero in each of the 32 bits of the Auto-Negotiation device identifier.

The format of the Auto-Negotiation device identifier is specified in 22.2.4.3.1.

45.2.7.4 AN devices in package (Registers 7.5 and 7.6)

The AN devices in package registers are defined in Table 45–2.

45.2.7.5 AN package identifier (Registers 7.14 and 7.15)

Registers 7.14 and 7.15 provide a 32-bit value, which may constitute a unique identifier for a particular type of package that the Auto-Negotiation is instantiated within. The identifier shall be composed of the 3rd through 24th bits of the OUI assigned to the package manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. An Auto-Negotiation MMD may return a value of zero in each of the 32 bits of the Auto-Negotiation package identifier.

A non-zero package identifier may be returned by one or more MMDs in the same package. The package identifier may be the same as the device identifier.

The format of the Auto-Negotiation package identifier is specified in 22.2.4.3.1.

45.2.7.6 AN advertisement register (7.16, 7.17, and 7.18)

If the BP AN ability bit (7.48.0) in the BP Ethernet status register is set to zero then only sixteen bits of the AN advertisement register are used and they are defined in 28.2.1.2.

If the BP AN ability bit (7.48.0) in the BP Ethernet status register is set to one then all forty eight bits are used and they are defined in 73.6.

The Selector field (7.16.4:0) is set to the IEEE 802.3 code as specified in Annex 28A. The Acknowledge bit (7.16.14) is set to zero.

The technology ability field, as defined in 28.2.1.2 and 73.6, represents the technologies supported by the local device. Only bits representing supported technologies may be set. Management may clear bits in the technology ability field and restart Auto-Negotiation to negotiate an alternate common mode.

The management entity initiates renegotiation with the link partner using alternate abilities by setting the Restart Auto-Negotiation bit (7.0.9) in the AN control register to one.

Any writes to this register prior to completion of Auto-Negotiation, as indicated by bit 7.1.5, should be followed by a renegotiation for the new values to take effect. Once Auto-Negotiation has completed, software may examine this register along with the LP Base Page ability register to determine the highest common denominator technology.

If the Auto-Negotiation advertisement register (Register 4) is present, (see 28.2.4.1.3), then this register is a copy of the Auto-Negotiation advertisement register (Register 4). In this case, reads to the AN advertisement register (7.16) report the value of the Auto-Negotiation advertisement register (Register 4); writes to the AN advertisement register (7.16) cause a write to occur to the Auto-Negotiation advertisement register.

The Base Page value is transferred to mr_adv_ability when register 7.16 is written. Therefore, if used, registers 7.17 and 7.18 should be written before 7.16.

Table 45–203—AN advertisement register bit definitions

Bit(s)	Name	Description	R/W[a]
7.16.15	Next Page	See 28.2.1.2 and 73.6.9	R/W
7.16.14	Acknowledge	Value always 0, writes ignored	RO
7.16.13	Remote fault	See 28.2.1.2 and 73.6.7	R/W
7.16.12:5	D12:D5	See 28.2.1.2 and 73.6	R/W
7.16.4:0	Selector field	See Annex 28A	R/W
7.17.15:0	D31:D16	See 73.6	R/W
7.18.15:0	D47:D32	See 73.6	R/W

[a]R/W = Read/Write, RO = Read only

NOTE 1—Auto-Negotiation that supports a 48-bit Base Page uses registers 7.16, 7.17, and 7.18 for storing Base Page information; Auto-Negotiation that supports a 16-bit Base Page only uses register 7.16 for storing Base Page information.

NOTE 2—Clause 37 1000BASE-X Auto-Negotiation is controlled through Clause 22 registers.

45.2.7.7 AN LP Base Page ability register (7.19, 7.20, and 7.21)

All of the bits in the AN LP Base Page ability register are read only. A write to the AN LP Base Page ability register shall have no effect.

Register 7.19 is a copy of register 5, if present (see 28.2.4.1).

When registers 7.20 and 7.21 are used, the value of the registers 7.20 and 7.21 is latched when register 7.19 is read and reads of registers 7.20 and 7.21 return the latched value rather than the current value.

Bit 7.19.12 is reserved for Backplane Ethernet port types as they do not use Extended Next Pages.

Table 45–204—AN LP Base Page ability register bit definitions

Bit(s)	Name	Description	R/W[a]
7.19.15:0	D15:D0	See 28.2.1.2 and 73.6	RO
7.20.15:0	D31:D16	See 73.6	RO
7.21.15:0	D47:D32	See 73.6	RO

[a]RO = Read only

NOTE—Auto-Negotiation that supports a 48-bit Base Page uses registers 7.19, 7.20, and 7.21 for storing LP Base Page information; Auto-Negotiation that supports a 16-bit Base Page only use register 7.19 for storing LP Base Page information.

45.2.7.8 AN XNP transmit register (7.22, 7.23, and 7.24)

If the BP AN ability bit (7.48.0) in the BP Ethernet status register is set to zero then Extended Next Pages are enabled by setting bit 7.0.13 to one and the AN XNP transmit register contains the Next Page link codeword to be transmitted as defined in 28.2.3.4.

If the BP AN ability bit (7.48.0) is set to one then Next Page transmission is always enabled and the register contains the AN LD Next Page link codeword of the BP Ethernet PHY as defined in 73.7.7.1. Even though Backplane Ethernet does not use Extended Next Pages XNP is still used in the register name.

On power-up or AN reset, this register shall contain the default value, which represents a Message Page with the message code set to Null Message. This value may be replaced by any valid Extended Next Page message code that the device intends to transmit.

A write to register 7.23 or 7.24 does not set mr_next_page_loaded. Only a write to register 7.22 sets mr_next_page_loaded true as described in 28.2.4.1.8. Therefore registers 7.23 and 7.24 register should be written before register 7.22.

Table 45–205—AN XNP transmit register bit definitions

Bit(s)	Name	Description	R/W[a]
7.22.15	Next Page	See 28.2.3.4 and 73.7.7.1	R/W
7.22.14	Reserved	Value always 0	RO
7.22.13	Message Page	See 28.2.3.4 and 73.7.7.1	R/W
7.22.12	Acknowledge 2	See 28.2.3.4 and 73.7.7.1	R/W
7.22.11	Toggle	See 28.2.3.4 and 73.7.7.1	RO
7.22.10:0	Message/Unformatted Code Field	See 28.2.3.4 and 73.7.7.1	R/W
7.23.15:0	Unformatted Code Field 1	See 28.2.3.4 and 73.7.7.1	R/W
7.24.15:0	Unformatted Code Field 2	See 28.2.3.4 and 73.7.7.1	R/W

[a]R/W = Read/Write, RO = Read only

45.2.7.9 AN LP XNP ability register (7.25, 7.26, and 7.27)

AN LP XNP ability register (registers 7.25, 7.26, and 7.27) store link partner Extended Next Pages as shown in Table 45–206. All of the bits in the AN LP XNP ability register are read only. A write to the AN LP XNP ability register shall have no effect.

The value of registers 7.26 and 7.27 is latched when register 7.25 is read and reads of registers 7.26 and 7.27 return the latched value rather than the current value.

NOTE—If this register is used to store multiple link partner Extended Next Pages, the previous value of this register is assumed to be stored by a management entity that needs the information overwritten by subsequent link partner Extended Next Pages.

Table 45–206—AN LP XNP ability register bit definitions

Bit(s)	Name	Description	R/W[a]
7.25.15	Next Page	See 28.2.3.4 and 73.7.7.1	RO
7.25.14	Acknowledge	See 28.2.3.4 and 73.7.7.1	RO
7.25.13	Message Page	See 28.2.3.4 and 73.7.7.1	RO
7.25.12	Acknowledge 2	See 28.2.3.4 and 73.7.7.1	RO
7.25.11	Toggle	See 28.2.3.4 and 73.7.7.1	RO
7.25.10:0	Message/Unformatted Code Field	See 28.2.3.4 and 73.7.7.1	RO
7.26.15:0	Unformatted Code Field 1	See 28.2.3.4 and 73.7.7.1	RO
7.27.15:0	Unformatted Code Field 2	See 28.2.3.4 and 73.7.7.1	RO

[a]RO = Read only

45.2.7.10 10GBASE-T AN control register (Register 7.32)

The assignment of bits in the 10GBASE-T AN control register is shown in Table 45–207. The default values for each bit of the 10GBASE-T AN control register should be chosen so that the initial state of the device upon power up or reset is a normal operational state without management intervention.

Table 45–207—10GBASE-T AN control register

Bit(s)	Name	Description	R/W[a]
7.32.15	MASTER-SLAVE manual config enable	1=Enable MASTER-SLAVE manual configuration 0=Disable MASTER-SLAVE manual configuration Default value is 0	R/W
7.32.14	MASTER-SLAVE config value	1=Configure PHY as MASTER 0=Configure PHY as SLAVE Default value is 0	R/W
7.32.13	Port type	1=Multiport device 0=single-port device	R/W
7.32.12	10GBASE-T ability	1 = Advertise PHY as 10GBASE-T capable. 0 = Do not advertise the PHY as 10GBASE-T capable	R/W
7.32.11:3	Reserved	Value always 0	RO
7.32.2	LD PMA training reset request	1 = Local device requests that link partner reset PMA training PRBS every frame 0 = Local device requests that link partner run PMA training PRBS continuously	R/W
7.32.1	Fast retrain ability	1 = Advertise PHY as 10GBASE-T fast retrain capable 0 = Do not advertise PHY as 10GBASE-T fast retrain capable	R/W
7.32.0	LD loop timing ability	1 = Advertise PHY as capable of loop timing 0 = Do not advertise PHY as capable of loop timing	R/W

[a]R/W = Read/Write, RO = Read only

45.2.7.10.1 MASTER-SLAVE manual config enable (7.32.15)

If bit 7.32.15 is set to one then MASTER-SLAVE config value bit 7.32.14 is used to determine if the PHY operates as MASTER or SLAVE. If bit 7.32.15 is set to zero the PHY shall use Auto-Negotiation to determine the selection of MASTER or SLAVE. This bit is a copy of 9.12, if present (see 40.5.1.1), and any write to either 7.32.15 or 9.12 is reflected in both. Additional information regarding the resolution and selection of MASTER-SLAVE and bit 7.32.15 is contained in 55.6.2. The default value for bit 7.32.15 is zero.

45.2.7.10.2 MASTER-SLAVE config value (7.32.14)

Bit 7.32.14 is used to select MASTER or SLAVE operation if MASTER-SLAVE manual config bit 7.32.15 is set to one. If bit 7.32.14 is set to one the PHY shall operate as MASTER. If bit 7.32.14 is set to zero the PHY shall operate as SLAVE. This bit is a copy of 9.11 (see 40.5.1.1), if present, and any write to either 7.32.14 or 9.11 is reflected in both. Additional information regarding the resolution and selection of MASTER-SLAVE and bit 7.32.14 is contained in 55.6.2. The default value for bit 7.32.14 is zero.

45.2.7.10.3 Port type (7.32.13)

Bit 7.32.13 is to be used to indicate the preference during Auto-Negotiation to operate as MASTER (multiport device) or as SLAVE (single-port device) if the MASTER-SLAVE manual configuration enable bit 7.32.15 is zero. If bit 7.32.13 is set to one a preference to operate as MASTER is indicated. If bit 7.32.13 is set to zero a preference to operate as SLAVE is indicated. This bit is a copy of 9.10 (see 40.5.1.1), if present, and any write to either 7.32.13 or 9.10 is reflected in both. Subclause 55.6.2 contains additional information regarding the resolution and selection of MASTER-SLAVE and this bit.

45.2.7.10.4 10GBASE-T capability (7.32.12)

Bit 7.32.12 is to be used to select whether or not Auto-Negotiation advertises the ability to operate as a 10GBASE-T PHY. If bit 7.32.12 is set to one the PHY shall advertise 10GBASE-T PHY capability. If bit 7.32.12 is set to zero the PHY shall not advertise 10GBASE-T PHY capability.

45.2.7.10.5 LD PMA training reset request (7.32.2)

If bit 7.32.2 is set to one then the local device expects the link partner to reset the PMA training PRBS for every PMA training frame. If bit 7.32.2 is zero then the local device expects link partner to run PMA training PRBS continuously through every PMA training frame.

45.2.7.10.6 Fast retrain ability

Bit 7.32.1 is used to select whether or not Auto-Negotiation advertises the ability to support 10GBASE-T fast retrain. If bit 7.32.1 is set to one, the PHY shall advertise fast retrain ability. If bit 7.32.1 is set to zero, the PHY shall not advertise fast retrain ability.

45.2.7.10.7 LD loop timing ability (7.32.0)

Bit 7.32.0 is to be used to select whether or not Auto-Negotiation advertises the ability to perform loop timing. If bit 7.32.0 is set to one the PHY shall advertise loop timing capability. If bit 7.32.0 is set to zero the PHY shall not advertise loop timing capability.

45.2.7.11 10GBASE-T AN status register (Register 7.33)

The assignment of bits in the 10GBASE-T AN status register is shown in Table 45–208. All the bits in the 10GBASE-T AN status register are read only; a write shall have no effect.

Table 45–208—10GBASE-T AN status register

Bit(s)	Name	Description	R/W[a]
7.33.15	MASTER-SLAVE configuration fault	1 = MASTER-SLAVE configuration fault detected 0 = No MASTER-SLAVE configuration fault detected	RO LH SC
7.33.14	MASTER-SLAVE configuration resolution	1 = Local PHY configuration resolved to MASTER 0 = Local PHY configuration resolved to SLAVE	RO
7.33.13	Local receiver status	1 = Local receiver OK 0 = Local receiver not OK	RO
7.33.12	Remote receiver status	1 = Remote receiver OK 0 = Remote receiver not OK	RO
7.33.11	Link partner 10GBASE-T capability	1 = Link partner is able to operate as 10GBASE-T 0 = Link partner is not able to operate as 10GBASE-T This bit is guaranteed to be valid only when the Page received bit (7.1.6) has been set to 1.	RO
7.33.10	LP loop timing ability	1 = Link partner is capable of loop timing 0 = Link partner is not capable of loop timing	RO
7.33.9	LP PMA training reset request	1 = Link partner requests that local device reset PMA training PRBS every frame 0 = Link partner requests that local device run PMA training PRBS continuously	RO
7.33.8:2	Reserved	Value always 0	RO
7.33.1	Fast retrain ability	1 = Link partner is capable of 10GBASE-T fast retrain 0 = Link partner is not capable of 10GBASE-T fast retrain	RO
7.33.0	Reserved	Value always 0	RO

[a]RO = Read only, SC = Self-clearing, LH = Latching high

45.2.7.11.1 MASTER-SLAVE configuration fault (7.33.15)

MASTER-SLAVE configuration fault bit 7.33.15 shall be set in the event that determination of the MASTER-SLAVE cannot be successfully concluded. MASTER-SLAVE configuration fault, as well as the criteria and method of fault detection, is PHY specific. Additional information regarding the resolution and selection of MASTER-SLAVE and bit 7.33.15 for 10GBASE-T in contained in 55.6.2. The MASTER-SLAVE configuration fault bit 7.33.15 shall be cleared each time the 10GBASE-T status register 7.33 is read via the management interface and shall be cleared by a 10GBASE-T PMA reset. Bit 7.33.15 shall self clear upon Auto-Negotiation enable.

45.2.7.11.2 MASTER-SLAVE configuration resolution (7.33.14)

Bit 7.33.14 is determined by the MASTER-SLAVE configuration resolution function described in 55.6.2. If the MASTER-SLAVE configuration resolution bit 7.33.14 is set to one and the Auto-Negotiation complete bit 7.1.5 is set and if MASTER-SLAVE configuration fault bit 7.33.15 in the 10GBASE-T status register is

zero, then MASTER mode of operation has been selected. If the MASTER-SLAVE configuration resolution bit 7.33.14 is set to zero and the Auto-Negotiation complete bit 7.1.5 is set and if MASTER-SLAVE configuration fault bit 7.33.15 in the 10GBASE-T status register is zero, then SLAVE mode of operation has been selected.

45.2.7.11.3 Local receiver status (7.33.13)

Local receiver status bit 7.33.13 shall be set if the local receiver is OK as defined in 55.2.2.7. If the local receiver status bit 7.33.13 is zero, the local receiver is not OK as defined in 55.2.2.7.2.

45.2.7.11.4 Remote receiver status (7.33.12)

Remote receiver status bit 7.33.13 shall be set if the remote receiver status is OK as defined in 55.2.2.8. If the local receiver status bit 7.33.13 is zero, the local receiver is not OK as defined in 55.2.2.8.2.

45.2.7.11.5 Link partner 10GBASE-T capability (7.33.11)

The bit is only valid when page receive bit 7.1.6 in is set to one. When read as a one, bit 7.33.11 indicates that the link partner is able to operate as 10GBASE-T. When read as a zero, bit 7.33.11 indicates that the link partner is not able to operate as 10GBASE-T.

45.2.7.11.6 Link partner loop timing ability (7.33.10)

When read as a one, bit 7.33.10 indicates that the link partner has the ability to support loop timing as specified in 55.1.3. When read as a zero, bit 7.33.10 indicates that the link partner lacks the ability to support loop timing.

45.2.7.11.7 Link partner PMA training reset request (7.33.9)

If bit 7.33.9 is set to one then the link partner is expecting the local device to reset the PMA training PRBS for every PMA training frame. If bit 7.33.9 is zero then the link partner expects the local device to run PMA training PRBS continuously through every PMA training frame.

45.2.7.11.8 Fast retrain ability (7.33.1)

When read as a one, bit 7.33.1 is used to indicate that the link partner has the ability to support the fast retrain capability as specified in 55.4.2.5.15. When read as a zero, bit 7.33.1 indicates that the PHY lacks the ability to support the fast retrain capability.

45.2.7.12 Backplane Ethernet, BASE-R copper status (Register 7.48)

The assignment of bits in the Backplane Ethernet, BASE-R copper status register is shown in Table 45–209.

Table 45–209—Backplane Ethernet, BASE-R copper status register (Register 7.48) bit definitions

Bit(s)	Name	Description	R/W[a]
7.48.15:12	Reserved	Value always 0	RO
7.48.11	100GBASE-CR4	1 = PMA/PMD is negotiated to perform 100GBASE-CR4 0 = PMA/PMD is not negotiated to perform 100GBASE-CR4	RO

Table 45–209—Backplane Ethernet, BASE-R copper status register (Register 7.48) bit definitions *(continued)*

7.48.10	100GBASE-KR4	1 = PMA/PMD is negotiated to perform 100GBASE-KR4 0 = PMA/PMD is not negotiated to perform 100GBASE-KR4	RO
7.48.9	100GBASE-KP4	1 = PMA/PMD is negotiated to perform 100GBASE-KP4 0 = PMA/PMD is not negotiated to perform 100GBASE-KP4	RO
7.48.8	100GBASE-CR10	1 = PMA/PMD is negotiated to perform 100GBASE-CR10 0 = PMA/PMD is not negotiated to perform 100GBASE-CR10	RO
7.48.7	Reserved	Value always 0	RO
7.48.6	40GBASE-CR4	1 = PMA/PMD is negotiated to perform 40GBASE-CR4 0 = PMA/PMD is not negotiated to perform 40GBASE-CR4	RO
7.48.5	40GBASE-KR4	1 = PMA/PMD is negotiated to perform 40GBASE-KR4 0 = PMA/PMD is not negotiated to perform 40GBASE-KR4	RO
7.48.4	BASE-R FEC negotiated	1 = PMA/PMD is negotiated to perform BASE-R FEC 0 = PMA/PMD is not negotiated to perform BASE-R FEC	RO
7.48.3	10GBASE-KR	1 = PMA/PMD is negotiated to perform 10GBASE-KR 0 = PMA/PMD is not negotiated to perform 10GBASE-KR	RO
7.48.2	10GBASE-KX4	1 = PMA/PMD is negotiated to perform 10GBASE-KX4 or CX4 0 = PMA/PMD is not negotiated to perform 10GBASE-KX4/CX4	RO
7.48.1	1000BASE-KX	1 = PMA/PMD is negotiated to perform 1000BASE-KX 0 = PMA/PMD is not negotiated to perform 1000BASE-KX	RO
7.48.0	BP AN ability	If a Backplane, BASE-R copper PHY type is implemented, this bit is set to 1	RO

[a]RO = Read only

45.2.7.12.1 BASE-R FEC negotiated (7.48.4)

When the Auto-Negotiation process has completed as indicated by the AN complete bit (7.1.5), bit 7.48.4 indicates that BASE-R FEC operation has been negotiated. This bit is set only if a BASE-R PHY supporting FEC operation has also been negotiated.

45.2.7.12.2 Negotiated Port Type (7.48.1, 7.48.2, 7.48.3, 7.48.5, 7.48.6, 7.48.8, 7.48.9, 7.48.10, 7.48.11)

When the AN process has been completed as indicated by the AN complete bit, these bits (1000BASE-KX, 10GBASE-KX4, 10GBASE-KR, 40GBASE-KR4, 40GBASE-CR4, 100GBASE-CR10, 100GBASE-KP4, 100GBASE-KR4, 100GBASE-CR4) indicate the negotiated port type. Only one of these bits is set depending on the priority resolution function. System developers need to distinguish between parallel detection of 10GBASE-KX4 and 10GBASE-CX4 based on the MDI and media type present.

45.2.7.12.3 Backplane Ethernet, BASE-R copper AN ability (7.48.0)

If a Backplane, BASE-R copper PHY type is implemented, this bit shall be set to 1.

When read as a one, bit 7.48.0 indicates that the PMA/PMD has the ability to perform Backplane Ethernet, BASE-R copper AN. When read as a zero, bit 7.48.0 indicates that the PMA/PMD lacks the ability to perform Backplane Ethernet BASE-R copper AN.

45.2.7.13 EEE advertisement (Register 7.60)

This register defines the EEE advertisement that is sent in the Unformatted Next Page following a EEE technology message code as defined in 28C.12 or sent in the unformatted code field of Message Next Page with EEE technology message code as defined in 73A.4 or sent as part of the 10GBASE-T and 1000BASE-T technology message code as defined in 28C.11. The assignment of bits in the EEE advertisement register and the correspondence with the bits in the Next Page messages are shown in Table 45–210.

Bits 10:0 of register 7.60 map to bits U10 through U0 respectively of the Unformatted Next Page following a EEE technology message code as defined in 28C.12. Bits 15:0 of register 7.60 map to bits U15 through U0 respectively of the unformatted code field of Message Next Page with EEE technology message code as defined in 73A.4. Bits 3:1 of register 7.60 also map to bits U24 through U22 respectively of the 10GBASE-T and 1000BASE-T technology message code as defined in 28C.11. Devices using Clause 28 auto-negotiation may ignore bits defined for Clause 73 auto-negotiation, and devices using Clause 73 auto-negotiation may ignore bits defined for Clause 28 auto-negotiation.

Table 45–210—EEE advertisement register (Register 7.60) bit definitions

Bit(s)	Name	Description	Clause reference; Next Page bit number	R/W[a]
7.60.15:14	Reserved	Value always 0		RO
7.60.13	100GBASE-CR4 EEE	1 = Advertise that the 100GBASE-CR4 has EEE deep sleep capability 0 = Do not advertise that the 100GBASE-CR4 has EEE deep sleep capability	73.7.7.1; U13	R/W
7.60.12	100GBASE-KR4 EEE	1 = Advertise that the 100GBASE-KR4 has EEE deep sleep capability 0 = Do not advertise that the 100GBASE-KR4 has EEE deep sleep capability	73.7.7.1; U12	R/W
7.60.11	100GBASE-KP4 EEE	1 = Advertise that the 100GBASE-KP4 has EEE deep sleep capability 0 = Do not advertise that the 100GBASE-KP4 has EEE deep sleep capability	73.7.7.1; U11	R/W
7.60.10	100GBASE-CR10 EEE	1 = Advertise that the 100GBASE-CR10 has EEE deep sleep capability 0 = Do not advertise that the 100GBASE-CR10 has EEE deep sleep capability	73.7.7.1; U10	R/W
7.60.9	Reserved	Value always 0		RO
7.60.8	40GBASE-CR4 EEE	1 = Advertise that the 40GBASE-CR4 has EEE deep sleep capability 0 = Do not advertise that the 40GBASE-CR4 has EEE deep sleep capability	73.7.7.1; U8	R/W
7.60.7	40GBASE-KR4 EEE	1 = Advertise that the 40GBASE-KR4 has EEE deep sleep capability 0 = Do not advertise that the 40GBASE-KR4 has EEE deep sleep capability	73.7.7.1; U7	R/W
7.60.6	10GBASE-KR EEE	1 = Advertise that the 10GBASE-KR has EEE capability 0 = Do not advertise that the 10GBASE-KR has EEE capability	73.7.7.1; U6	R/W

Table 45–210—EEE advertisement register (Register 7.60) bit definitions *(continued)*

7.60.5	10GBASE-KX4 EEE	1 = Advertise that the 10GBASE-KX4 has EEE capability 0 = Do not advertise that the 10GBASE-KX4 has EEE capability	73.7.7.1; U5	R/W
7.60.4	1000BASE-KX EEE	1 = Advertise that the 1000BASE-KX has EEE capability 0 = Do not advertise that the 1000BASE-KX has EEE capability	73.7.7.1; U4	R/W
7.60.3	10GBASE-T EEE	1 = Advertise that the 10GBASE-T has EEE capability 0 = Do not advertise that the 10GBASE-T has EEE capability	28.2.3.4.1; U3 / 28.2.3.4.2; U24	R/W
7.60.2	1000BASE-T EEE	1 = Advertise that the 1000BASE-T has EEE capability 0 = Do not advertise that the 1000BASE-T has EEE capability	28.2.3.4.1; U2 / 28.2.3.4.2; U23	R/W
7.60.1	100BASE-TX EEE	1 = Advertise that the 100BASE-TX has EEE capability 0 = Do not advertise that the 100BASE-TX has EEE capability	28.2.3.4.1; U1 / 28.2.3.4.2; U22	R/W
7.60.0	Reserved	Value always 0		RO

[a]R/W = Read/Write, RO = Read only

45.2.7.13.1 100GBASE-CR4 EEE supported (7.60.13)

Support for EEE deep sleep operation for 100GBASE-CR4, as defined in 92.1, shall be advertised if this bit is set to one. Support for EEE deep sleep operation should only be advertised if it is supported on all sublayers of the PHY as well as physical instantiations of the PMA service interface as appropriate.

45.2.7.13.2 100GBASE-KR4 EEE supported (7.60.12)

Support for EEE deep sleep operation for 100GBASE-KR4, as defined in 93.1, shall be advertised if this bit is set to one. Support for EEE deep sleep operation should only be advertised if it is supported on all sublayers of the PHY as well as physical instantiations of the PMA service interface as appropriate.

45.2.7.13.3 100GBASE-KP4 EEE supported (7.60.11)

Support for EEE deep sleep operation for 100GBASE-KP4, as defined in 94.1, shall be advertised if this bit is set to one. Support for EEE deep sleep operation should only be advertised if it is supported on all sublayers of the PHY as well as physical instantiations of the PMA service interface as appropriate.

45.2.7.13.4 100GBASE-CR10 EEE supported (7.60.10)

Support for EEE deep sleep operation for 100GBASE-CR10, as defined in 85.1, shall be advertised if this bit is set to one. Support for EEE deep sleep operation should only be advertised if it is supported on all sublayers of the PHY as well as physical instantiations of the PMA service interface as appropriate.

45.2.7.13.5 40GBASE-CR4 EEE supported (7.60.8)

Support for EEE deep sleep operation for 40GBASE-CR4, as defined in 85.1, shall be advertised if this bit is set to one. Support for EEE deep sleep operation should only be advertised if it is supported on all sublayers of the PHY as well as physical instantiations of the PMA service interface as appropriate.

45.2.7.13.6 40GBASE-KR4 EEE supported (7.60.7)

Support for EEE deep sleep operation for 40GBASE-KR4, as defined in 84.1, shall be advertised if this bit is set to one. Support for EEE deep sleep operation should only be advertised if it is supported on all sublayers of the PHY as well as physical instantiations of the PMA service interface as appropriate.

45.2.7.13.7 10GBASE-KR EEE supported (7.60.6)

If the device supports EEE operation for 10GBASE-KR as defined in 72.1, and EEE operation is desired, this bit shall be set to one.

45.2.7.13.8 10GBASE-KX4 EEE supported (7.60.5)

If the device supports EEE operation for 10GBASE-KX4 as defined in 71.2, and EEE operation is desired, this bit shall be set to one.

45.2.7.13.9 1000BASE-KX EEE supported (7.60.4)

If the device supports EEE operation for 1000BASE-KX as defined in 70.1, and EEE operation is desired, this bit shall be set to one.

45.2.7.13.10 10GBASE-T EEE supported (7.60.3)

If the device supports EEE operation for 10GBASE-T as defined in 55.1.3.3, and EEE operation is desired, this bit shall be set to one.

45.2.7.13.11 1000BASE-T EEE supported (7.60.2)

If the device supports EEE operation for 1000BASE-T as defined in 40.2.11, and EEE operation is desired, this bit shall be set to one.

45.2.7.13.12 100BASE-TX EEE supported (7.60.1)

If the device supports EEE operation for 100BASE-TX as defined in 25.5, and EEE operation is desired, this bit shall be set to one.

45.2.7.14 EEE link partner ability (Register 7.61)

All of the bits in the EEE LP ability register are read-only. A write to the EEE LP ability register shall have no effect. When the AN process has been completed, this register shall reflect the contents of the link

partner's EEE advertisement register. The assignment of bits in the EEE link partner ability register and the correspondence with the bits in the Next Page messages are shown in Table 45–211.

Table 45–211—EEE link partner ability (Register 7.61) bit definitions

Bit(s)	Name	Description	Clause reference; Next Page bit number	R/W[a]
7.61.15:14	Reserved	Value always 0		RO
7.61.13	100GBASE-CR4 EEE	1 = Link partner is advertising EEE deep sleep capability for 100GBASE-CR4 0 = Link partner is not advertising EEE deep sleep capability for 100GBASE-CR4	73.7.7.1; U13	RO
7.61.12	100GBASE-KR4 EEE	1 = Link partner is advertising EEE deep sleep capability for 100GBASE-KR4 0 = Link partner is not advertising EEE deep sleep capability for 100GBASE-KR4	73.7.7.1; U12	RO
7.61.11	100GBASE-KP4 EEE	1 = Link partner is advertising EEE deep sleep capability for 100GBASE-KP4 0 = Link partner is not advertising EEE deep sleep capability for 100GBASE-KP4	73.7.7.1; U11	RO
7.61.10	100GBASE-CR10 EEE	1 = Link partner is advertising EEE deep sleep capability for 100GBASE-CR10 0 = Link partner is not advertising EEE deep sleep capability for 100GBASE-CR10	73.7.7.1; U10	RO
7.61.9	Reserved	Value always 0		RO
7.61.8	40GBASE-CR4 EEE	1 = Link partner is advertising EEE deep sleep capability for 40GBASE-CR4 0 = Link partner is not advertising EEE deep sleep capability for 40GBASE-CR4	73.7.7.1; U8	RO
7.61.7	40GBASE-KR4 EEE	1 = Link partner is advertising EEE deep sleep capability for 40GBASE-KR4 0 = Link partner is not advertising EEE deep sleep capability for 40GBASE-KR4	73.7.7.1; U7	RO
7.61.6	10GBASE-KR EEE	1 = Link partner is advertising EEE capability for 10GBASE-KR 0 = Link partner is not advertising EEE capability for 10GBASE-KR	73.7.7.1; U6	RO
7.61.5	10GBASE-KX4 EEE	1 = Link partner is advertising EEE capability for 10GBASE-KX4 0 = Link partner is not advertising EEE capability for 10GBASE-KX4	73.7.7.1; U5	RO

Table 45–211—EEE link partner ability (Register 7.61) bit definitions *(continued)*

7.61.4	1000BASE-KX EEE	1 = Link partner is advertising EEE capability for 1000BASE-KX 0 = Link partner is not advertising EEE capability for 1000BASE-KX	73.7.7.1; U4	RO
7.61.3	10GBASE-T EEE	1 = Link partner is advertising EEE capability for 10GBASE-T 0 = Link partner is not advertising EEE capability for 10GBASE-T	28.2.3.4.1; U3 / 55.6.1; U24	RO
7.61.2	1000BASE-T EEE	1 = Link partner is advertising EEE capability for 1000BASE-T 0 = Link partner is not advertising EEE capability for 1000BASE-T	28.2.3.4.1; U2 / 55.6.1; U23	RO
7.61.1	100BASE-TX EEE	1 = Link partner is advertising EEE capability for 100BASE-TX 0 = Link partner is not advertising EEE capability for 100BASE-TX	28.2.3.4.1; U1 / 55.6.1; U22	RO
7.61.0	Reserved	Value always 0		RO

[a]RO = Read only

The definitions for the contents of the EEE LP ability register are given by the definitions for the contents on the link partner's EEE advertisement register, 7.60 (see 45.2.7.13).

45.2.8 Clause 22 extension registers

As new management features are added to 10, 100, and 1000 Mb/s PHYs, more register space is required beyond that defined in Clause 22. The Clause 22 extension MMD provides this space. This MMD is defined only for 10, 100, and 1000 Mb/s PHYs. Since these PHYs do not segment their management by their sublayers, all management extensions to these PHYs will appear in the Clause 22 extension MMD.

The assignment of registers in the Clause 22 extension MMD is shown in Table 45–212.

Table 45–212—Clause 22 extension registers

Register address	Register name	Subclause
29.0 through 29.4	Reserved	
29.5, 29.6	Clause 22 extension devices in package	45.2.8.1
29.7	FEC capability	45.2.8.2
29.8	FEC control	45.2.8.3
29.9	FEC buffer head coding violation counter	45.2.8.4
29.10	FEC corrected blocks counter	45.2.8.5
29.11	FEC uncorrected blocks counter	45.2.8.6
29.12 through 29.32 767	Reserved	

45.2.8.1 Clause 22 extension devices in package registers (Registers 29.5, 29.6)

The Clause 22 extension devices in package registers are defined in Table 45–2.

45.2.8.2 FEC capability register (Register 29.7)

The assignment of bits in the FEC capability register is shown in Table 45–213.

Table 45–213—FEC capability register bit definitions

Bit(s)	Name	Description	R/W[a]
29.7.15:1	Reserved	Value always 0	RO
29.7.0	FEC capable	1 = FEC supported 0 = FEC unsupported	RO

[a]RO = Read only

45.2.8.2.1 FEC capable (29.7.0)

When read as a one, this bit indicates that the PHY supports forward error correction. When read as a zero, the PHY does not support forward error correction.

45.2.8.3 FEC control register (Register 29.8)

The assignment of bits in the FEC control register is shown in Table 45–214.

Table 45–214—FEC control register bit definitions

Bit(s)	Name	Description	R/W[a]
29.8.15:1	Reserved	Value always 0	RO
29.8.0	FEC enable	1 = FEC enabled 0 = FEC disabled	R/W

[a]RO = Read only, R/W = Read/Write

45.2.8.3.1 FEC enable (29.8.0)

When written as a one, this bit enables the PHY's forward error correction. When written as a zero, FEC is disabled. This bit shall be set to zero upon execution of a PHY reset.

45.2.8.4 FEC buffer head coding violation counter (Register 29.9)

The assignment of bits in the FEC buffer head coding violation counter register is shown in Table 45–215. See 65.2.3.6.1 for a definition of this register. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow.

Table 45–215—FEC buffer head coding violation counter register bit definitions

Bit(s)	Name	Description	R/W[a]
29.9.15:0	FEC buffer head coding violation counter	Error counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.8.5 FEC corrected blocks counter (Register 29.10)

The assignment of bits in the FEC corrected blocks counter register is shown in Table 45–216. See 65.2.3.6.2 for a definition of this register. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow.

Table 45–216—FEC corrected blocks counter register bit definitions

Bit(s)	Name	Description	R/W[a]
29.10.15:0	FEC corrected blocks counter	Error counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.8.6 FEC uncorrected blocks counter (Register 29.11)

The assignment of bits in the FEC uncorrected blocks counter register is shown in Table 45–217. See 65.2.3.6.3 for a definition of this register. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow.

Table 45–217—FEC uncorrected blocks counter register bit definitions

Bit(s)	Name	Description	R/W[a]
29.11.15:0	FEC uncorrected blocks counter	Error counter	RO, NR

[a]RO = Read only, NR = Non Roll-over

45.2.9 Vendor specific MMD 1 registers

The assignment of registers in the vendor specific MMD 1 is shown in Table 45–218. A vendor specific MMD may have a device address of either 30 or 31. It is recommended that the device address is configurable and that the configuration is performed by some means other than via the MDIO.

Table 45–218—Vendor specific MMD 1 registers

Register address	Register name
30.0, 30.1	Vendor specific
30.2, 30.3	Vendor specific MMD 1 device identifier

Table 45–218—Vendor specific MMD 1 registers *(continued)*

Register address	Register name
30.4 through 30.7	Vendor specific
30.8	Vendor specific MMD 1 status register
30.9 through 30.13	Vendor specific
30.14, 30.15	Vendor specific MMD 1 package identifier
30.16 through 30.65 535	Vendor specific

45.2.9.1 Vendor specific MMD 1 device identifier (Registers 30.2 and 30.3)

Registers 30.2 and 30.3 provide a 32-bit value, which may constitute a unique identifier for a particular type of vendor-specific device. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the device manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A vendor-specific device may return a value of zero in each of the 32 bits of the vendor specific MMD 1 device identifier.

The format of the vendor specific MMD 1 device identifier is specified in 22.2.4.3.1.

45.2.9.2 Vendor specific MMD 1 status register (Register 30.8)

The assignment of bits in the vendor specific MMD 1 status register is shown in Table 45–219. All the bits in the vendor specific MMD 1 status register are read only; a write to the vendor specific MMD 1 status register shall have no effect.

Table 45–219—Vendor specific MMD 1 status register bit definitions

Bit(s)	Name	Description	R/W[a]
30.8.15:14	Device present	15 14 1 0 = Device responding at this address 1 1 = No device responding at this address 0 1 = No device responding at this address 0 0 = No device responding at this address	RO
30.8.13:0	Reserved	Value always 0	RO

[a]RO = Read only

45.2.9.2.1 Device present (30.8.15:14)

When read as <10>, bits 30.8.15:14 indicate that a device is present and responding at this register address. When read as anything other than <10>, bits 30.8.15:14 indicate that no device is present at this register address or that the device is not functioning properly.

45.2.9.3 Vendor specific MMD 1 package identifier (Registers 30.14 and 30.15)

Registers 30.14 and 30.15 provide a 32-bit value, which may constitute a unique identifier for a particular type of package that the vendor specific MMD 1 is instantiated within. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the package manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A vendor specific MMD 1 may return a value of zero in each of the 32 bits of the vendor specific MMD 1 package identifier.

A non-zero package identifier may be returned by one or more MMDs in the same package. The package identifier may be the same as the device identifier.

The format of the vendor specific MMD 1 package identifier is specified in 22.2.4.3.1.

45.2.10 Vendor specific MMD 2 registers

The assignment of registers in the vendor specific MMD 2 is shown in Table 45–220. A vendor specific MMD may have a device address of either 30 or 31. It is recommended that the device address is configurable and that the configuration is performed by some means other than via the MDIO.

Table 45–220—Vendor specific MMD 2 registers

Register address	Register name	Subclause
31.0, 31.1	Vendor specific	
31.2, 31.3	Vendor specific MMD 2 device identifier	45.2.10.1
31.4 through 31.7	Vendor specific	
31.8	Vendor specific MMD 2 status register	45.2.10.2
31.9 through 31.13	Vendor specific	
31.14, 30.15	Vendor specific MMD 2 package identifier	45.2.10.3
31.16 through 31.65 535	Vendor specific	

45.2.10.1 Vendor specific MMD 2 device identifier (Registers 31.2 and 31.3)

Registers 31.2 and 31.3 provide a 32-bit value, which may constitute a unique identifier for a particular type of vendor-specific device. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the device manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A vendor-specific device may return a value of zero in each of the 32 bits of the vendor specific MMD 2 device identifier.

The format of the vendor specific MMD 2 device identifier is specified in 22.2.4.3.1.

45.2.10.2 Vendor specific MMD 2 status register (Register 31.8)

The assignment of bits in the vendor specific MMD 2 status register is shown in Table 45–221. All the bits in the vendor specific MMD 2 status register are read only; a write to the vendor specific MMD status register shall have no effect.

Table 45–221—Vendor specific MMD 2 status register bit definitions

Bit(s)	Name	Description	R/W[a]
31.8.15:14	Device present	15 14 1 0 = Device responding at this address 1 1 = No device responding at this address 0 1 = No device responding at this address 0 0 = No device responding at this address	RO
31.8.13:0	Reserved	Value always 0	RO

[a]RO = Read only

45.2.10.2.1 Device present (31.8.15:14)

When read as <10>, bits 31.8.15:14 indicate that a device is present and responding at this register address. When read as anything other than <10>, bits 31.8.15:14 indicate that no device is present at this register address or that the device is not functioning properly.

45.2.10.3 Vendor specific MMD 2 package identifier (Registers 31.14 and 31.15)

Registers 31.14 and 31.15 provide a 32-bit value, which may constitute a unique identifier for a particular type of package that the vendor specific MMD is instantiated within. The identifier shall be composed of the 3rd through 24th bits of the Organizationally Unique Identifier (OUI) assigned to the package manufacturer by the IEEE, plus a six-bit model number, plus a four-bit revision number. A vendor specific MMD may return a value of zero in each of the 32 bits of the package identifier.

A non-zero package identifier may be returned by one or more MMDs in the same package. The package identifier may be the same as the device identifier.

The format of the vendor specific MMD 2 package identifier is specified in 22.2.4.3.1.

45.3 Management frame structure

The MDIO interface frame structure is compatible with the one defined in 22.2.4.5 such that the two systems can co-exist on the same MDIO bus. The electrical specification for the MDIO interface is incompatible to that defined in 22.2.4.5; therefore, if the two systems are to co-exist on the same bus, a voltage translation device is required (see Annex 45A). The extensions that are used for MDIO indirect register accesses are specified in Table 45–222.

Table 45–222—Extensions to management frame format for indirect access

	Management frame fields							
Frame	**PRE**	**ST**	**OP**	**PRTAD**	**DEVAD**	**TA**	**ADDRESS / DATA**	**IDLE**
Address	1...1	00	00	PPPPP	EEEEE	10	AAAAAAAAAAAAAAAA	Z
Write	1...1	00	01	PPPPP	EEEEE	10	DDDDDDDDDDDDDDDD	Z

Table 45–222—Extensions to management frame format for indirect access *(continued)*

	Management frame fields							
Frame	**PRE**	**ST**	**OP**	**PRTAD**	**DEVAD**	**TA**	**ADDRESS / DATA**	**IDLE**
Read	1...1	00	11	PPPPP	EEEEE	Z0	DDDDDDDDDDDDDDDD	Z
Post-read-increment-address	1...1	00	10	PPPPP	EEEEE	Z0	DDDDDDDDDDDDDDDD	Z

Each MMD shall implement a sixteen bit address register that stores the address of the register to be accessed by data transaction frames. The address register shall be overwritten by address frames. At power up or device reset, the contents of the address register are undefined.

Write, read, and post-read-increment-address frames shall access the register whose address is stored in the address register. Write and read frames shall not modify the contents of the address register.

Upon receiving a post-read-increment-address frame and having completed the read operation, the MMD shall increment the address register by one. For the case where the MMD's address register contains 65 535, the MMD shall not increment the address register.

Implementations that incorporate several MMDs within a single component shall implement separate address registers so that the MMD's address registers operate independently of one another.

45.3.1 IDLE (idle condition)

The idle condition on MDIO is a high-impedance state. All three state drivers shall be disabled and the MMD's pull-up resistor will pull the MDIO line to a one.

45.3.2 PRE (preamble)

At the beginning of each transaction, the station management entity shall send a sequence of 32 contiguous one bits on MDIO with 32 corresponding cycles on MDC to provide the MMD with a pattern that it can use to establish synchronization. An MMD shall observe a sequence of 32 contiguous one bits on MDIO with 32 corresponding cycles on MDC before it responds to any transaction.

45.3.3 ST (start of frame)

The start of frame for indirect access cycles is indicated by the <00> pattern. This pattern assures a transition from the default one and identifies the frame as an indirect access. Frames that contain the ST=<01> pattern defined in Clause 22 shall be ignored by the devices specified in Clause 45.

45.3.4 OP (operation code)

The operation code field indicates the type of transaction being performed by the frame. A <00> pattern indicates that the frame payload contains the address of the register to access. A <01> pattern indicates that the frame payload contains data to be written to the register whose address was provided in the previous address frame. A <11> pattern indicates that the frame is read operation. A <10> pattern indicates that the frame is a post-read-increment-address operation.

45.3.5 PRTAD (port address)

The port address is five bits, allowing 32 unique port addresses. The first port address bit to be transmitted and received is the MSB of the address. A station management entity must have a priori knowledge of the appropriate port address for each port to which it is attached, whether connected to a single port or to multiple ports.

45.3.6 DEVAD (device address)

The device address is five bits, allowing 32 unique MMDs per port. The first device address bit transmitted and received is the MSB of the address.

45.3.7 TA (turnaround)

The turnaround time is a 2 bit time spacing between the device address field and the data field of a management frame to avoid contention during a read transaction. For a read or post-read-increment-address transaction, both the STA and the MMD shall remain in a high-impedance state for the first bit time of the turnaround. The MMD shall drive a zero bit during the second bit time of the turnaround of a read or post-read-increment-address transaction. During a write or address transaction, the STA shall drive a one bit for the first bit time of the turnaround and a zero bit for the second bit time of the turnaround. Figure 22–15 shows the behavior of the MDIO signal during the turnaround field of a read or post-read-increment-address transaction.

45.3.8 ADDRESS / DATA

The address/data field is 16 bits. For an address cycle, it contains the address of the register to be accessed on the next cycle. For the data cycle of a write frame, the field contains the data to be written to the register. For a read or post-read-increment-address frame, the field contains the contents of the register. The first bit transmitted and received shall be bit 15.

45.4 Electrical interface

45.4.1 Electrical specification

The electrical characteristics of the MDIO interface are shown in Table 45–223. The MDIO uses signal levels that are compatible with devices operating at a nominal supply voltage of 1.2 V. More information on the electrical interface is given in Annex 45A. Voltage translators between the Clause 22 electrical interface and the Clause 45 electrical interface are described in 45A.3 and 45A.4.

NOTE—It is possible to implement the MDIO electrical interface using open drain buffers and a resistive pull-up to a V_{DD} of 1.2 V (see 45A.1).

Table 45–223—MDIO electrical interface characteristics

Symbol	Parameter	Condition	Min.	Max.
V_{IH}	Input high voltage		0.84 V	1.5 V
V_{IL}	Input low voltage		–0.3 V	0.36 V
V_{OH}	Output high voltage	I_{OH} = –100 uA	1.0 V	1.5 V

Table 45–223—MDIO electrical interface characteristics *(continued)*

Symbol	Parameter	Condition	Min.	Max.
V_{OL}	Output low voltage	I_{OL} = 100 uA	–0.3 V	0.2 V
I_{OH}[a]	Output high current	V_{OH} = 1.0 V		–4 mA
I_{OL}	Output low current	V_{OL} = 0.2 V	+4 mA	
C_i	Input capacitance			10 pF
C_L	Total capacitive load			470 pF

[a]I_{OH} parameter is not applicable to open drain drivers.

45.4.2 Timing specification

MDIO is a bidirectional signal that can be sourced by the Station Management Entity (STA) or the MMD. When the STA sources the MDIO signal, the STA shall provide a minimum of 10 ns of setup time and a minimum of 10 ns of hold time referenced to the rising edge of MDC, as shown in Figure 45–3, measured at the MMD.

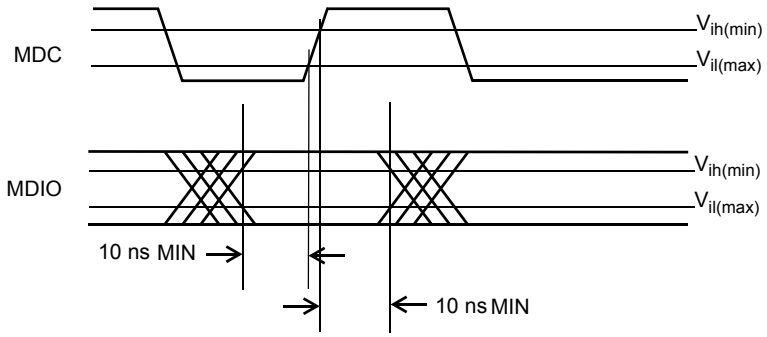

Figure 45–3—MDIO sourced by STA

When the MDIO signal is sourced by the MMD, it is sampled by the STA synchronously with respect to the beginning of the rising edge of MDC. The clock to output delay from the MMD, as measured at the STA, shall be a minimum of 0 ns, and a maximum of 300 ns, as shown in Figure 45–4.

The timing specification for the MDC signal is given in 22.2.2.13.

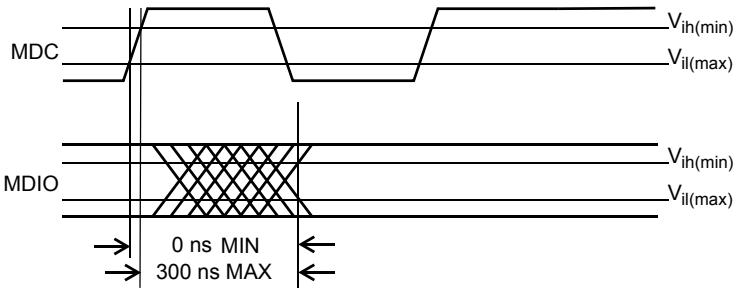

Figure 45–4—MDIO sourced by MMD

45.5 Protocol implementation conformance statement (PICS) proforma for Clause 45, Management Data Input/Output (MDIO) interface[1]

45.5.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 45, MDIO interface, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

45.5.2 Identification

45.5.2.1 Implementation identification

Supplier[1]	
Contact point for inquiries about the PICS[1]	
Implementation Name(s) and Version(s)[1,3]	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)[2]	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

45.5.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 45, Management Data Input/Output (MDIO) Interface
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [ ] Yes [ ] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2015.)	

Date of Statement	

[1] *Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

45.5.2.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*PMA	Implementation of PMA/PMD MMD	45.2.1		O	Yes [ ] No [ ]
*WIS	Implementation of WIS MMD	45.2.2		O	Yes [ ] No [ ]
*PCS	Implementation of PCS MMD	45.2.3		O	Yes [ ] No [ ]
*PX	Implementation of PHY XS MMD	45.2.4		O	Yes [ ] No [ ]
*DX	Implementation of DTE XS MMD	45.2.5		O	Yes [ ] No [ ]
*VSA	Implementation of Vendor specific MMD 1	45.2.9		O	Yes [ ] No [ ]
*VSB	Implementation of Vendor specific MMD 2	45.2.10		O	Yes [ ] No [ ]
*TC	Implementation of the TC MMD	45.2.6		10P*2B:M	Yes [ ] No [ ]
*CTT	Implementation of the Clause 22 extension MMD	45.2.8		O	Yes [ ] No [ ]
*ODB	Open drain buffer	45.4.1		O	Yes [ ] No [ ]
*AN	Implementation of Auto-Negotiation MMD	45.2.7		O	Yes [ ] No [ ]

45.5.3 PICS proforma tables for the Management Data Input Output (MDIO) interface

45.5.3.1 MDIO signal functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
SF1	MDIO uses three-state drivers	45.4.1		M	Yes [ ]

45.5.3.2 PMA/PMD MMD options

Item	Feature	Subclause	Value/Comment	Status	Support
*ALB	Implementation of PMA local loopback function	45.2.1.1.5		PMA:O	Yes [ ] No [ ] N/A [ ]
*RLB	Implementation of PMA remote loopback function	45.2.1.1.4		PMA:O	Yes [ ] No [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
*PLF	Implementation of fault detection	45.2.1.7		PMA:O	Yes [ ] No [ ] N/A [ ]
*ALP	Implementation of PMA/PMD low power ability	45.2.1.2.5		PMA:O	Yes [ ] No [ ] N/A [ ]
*PTD	Implementation of transmit disable function	45.2.1.8		PMA:O	Yes [ ] No [ ] N/A [ ]
*10T	Implementation of the 10GBASE-T PMA	45.2.1.6		PMA:O	Yes [ ] No [ ]
*40G	Implementation of 40 Gb/s PMA/PMD	45.2.1.4		PMA:O	Yes [ ] No [ ]
*100G	Implementation of 100 Gb/s PMA/PMD	45.2.1.4		PMA:O	Yes [ ] No [ ]
*10M	Implementation of 10 Mb/s PMA/PMD	45.2.1.4		PMA:O	Yes [ ] No [ ]
*100M	Implementation of 100 Mb/s PMA/PMD	45.2.1.4		PMA:O	Yes [ ] No [ ]
*1G	Implementation of 1000 Mb/s PMA/PMD	45.2.1.4		PMA:O	Yes [ ] No [ ]
*10P	Implementation of the 10PASS-TS PMA/PMD	45.2.1.4		PMA:O	Yes [ ] No [ ]
*2B	Implementation of the 2BASE-TL PMA/PMD	45.2.1.4		PMA:O	Yes [ ] No [ ]
*10G	Implementation of 10 Gb/s PMA/PMD	45.2.1.4		PMA:O	Yes [ ] No [ ]
*KX	Implementation of 1000BASE-KX PMA/PMD	45.2.1.6		PMA:O	Yes [ ] No [ ]
*KX4	Implementation of 10GBASE-KX4 PMA/PMD	45.2.1.6		PMA:O	Yes [ ] No [ ]
*KR	Implementation of 10GBASE-KR PMA/PMD	45.2.1.6		PMA:O	Yes [ ] No [ ]
*PXAR	Implementation of PMA/PMD Extended Ability Register	45.2.1.10	Required for certain abilities	PMA:O	Yes [ ] No [ ]
*40XAR	Implementation of 40G/100G PMA/PMD Extended Ability Register	45.2.1.12	Required for certain abilities	PMA:O	Yes [ ] No [ ]
*MMD8	Implementation of separately addressable instance, address 8	45.2.1	All register definitions as for address 1	PMA:O	Yes [ ] No [ ]
*MMD9	Implementation of separately addressable instance, address 9	45.2.1	All register definitions as for address 1	PMA:O	Yes [ ] No [ ]
*MMD10	Implementation of separately addressable instance, address 10	45.2.1	All register definitions as for address 1	PMA:O	Yes [ ] No [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
*MMD11	Implementation of separately addressable instance, address 11	45.2.1	All register definitions as for address 1	PMA:O	Yes [ ] No [ ]
*FEC-R	Implementation of BASE-R FEC	45.2.1.92		PMA:O	Yes [ ] No [ ]
*RS-FEC	Implementation of RS-FEC	45.2.1.102		PMA:O	Yes [ ] No [ ]

45.5.3.3 PMA/PMD management functions

Item	Feature	Subclause	Value/Comment	Status	Support
MM1	Device responds to all register addresses for that device	45.2		PMA:M	Yes [ ] N/A [ ]
MM2	Writes to undefined and read-only registers have no effect	45.2		PMA:M	Yes [ ] N/A [ ]
MM3	Operation is not affected by writes to reserved and unsupported bits.	45.2		PMA:M	Yes [ ] N/A [ ]
MM4	Reserved and unsupported bits return a value of zero	45.2		PMA:M	Yes [ ] N/A [ ]
MM5	Latching low bits remain low until after they have been read via the management interface	45.2		PMA:M	Yes [ ] N/A [ ]
MM6	Latching low bits assume correct value once read via the management interface	45.2	Correct value is based upon current state	PMA:M	Yes [ ] N/A [ ]
MM7	Latching high bits remain high until after they have been read via the management interface	45.2		PMA:M	Yes [ ] N/A [ ]
MM8	Latching high bits assume correct value once read via the management interface	45.2	Correct value is based upon current state	PMA:M	Yes [ ] N/A [ ]
MM9	Action on reset	45.2.1.1.1	Reset the registers of the entire device to default values and set bit 15 of the Control register to one	PMA:M	Yes [ ] N/A [ ]
MM10	Return 1 until reset completed	45.2.1.1.1		PMA:M	Yes [ ] N/A [ ]
MM11	Control and management interfaces are restored to operation within 0.5 s of reset	45.2.1.1.1		PMA:M	Yes [ ] N/A [ ]
MM12	Responds to reads of bit 15 and 1.8.15:14 during reset	45.2.1.1.1		PMA:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
MM13	Device responds to transactions necessary to exit low-power mode while in low-power state	45.2.1.1.2		PMA:M	Yes [ ] N/A [ ]
MM14	Invalid writes to speed selection bits are ignored	45.2.1.1.3		PMA:M	Yes [ ] N/A [ ]
MM15	PMA is set into local loopback mode when bit 0 is set to a one	45.2.1.1.5		ALB:M	Yes [ ] N/A [ ]
MM16	PMA transmit data is returned on receive path when in local loopback	45.2.1.1.5		ALB:M	Yes [ ] N/A [ ]
MM17	PMA ignores writes to this bit if it does not support local loopback.	45.2.1.1.5		PMA*!ALB:M	Yes [ ] N/A [ ]
MM18	PMA returns a value of zero when read if it does not support local loopback.	45.2.1.1.5		PMA*!ALB:M	Yes [ ] N/A [ ]
MM19	PMA is set into remote loopback mode when bit 1 is set to a one	45.2.1.1.4		RLB:M	Yes [ ] N/A [ ]
MM20	PMA receive data is returned on transmit path when in remote loopback	45.2.1.1.4		RLB:M	Yes [ ] N/A [ ]
MM21	PMA ignores writes to this bit if it does not support remote loopback.	45.2.1.1.4		PMA*!RLB:M	Yes [ ] N/A [ ]
MM22	PMA returns a value of zero when read if it does not support remote loopback.	45.2.1.1.4		PMA*!RLB:M	Yes [ ] N/A [ ]
MM23	Writes to status 1 register have no effect	45.2.1.2		PMA:M	Yes [ ] N/A [ ]
MM24	Receive link status implemented with latching low behavior	45.2.1.2.4		PMA:M	Yes [ ] N/A [ ]
MM25	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.1.3		PMA:M	Yes [ ] N/A [ ]
MM26	PMA/PMD type is selected using bits 5:0	45.2.1.6.3		PMA:M	Yes [ ] N/A [ ]
MM27	PMA/PMD ignores writes to type selection bits that select types that it has not advertised	45.2.1.6.3		PMA:M	Yes [ ] N/A [ ]
MM28	Writes to the status 2 register have no effect	45.2.1.7		PMA:M	Yes [ ] N/A [ ]
MM29	PMA/PMD returns a value of zero for transmit fault if it is unable to detect a transmit fault	45.2.1.7.4		PMA:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
MM30	Transmit fault is implemented using latching high behavior	45.2.1.7.4		PMA*PLF:M	Yes [ ] N/A [ ]
MM31	PMA/PMD returns a value of zero for receive fault if it is unable to detect a receive fault	45.2.1.7.5		PMA*!PLF:M	Yes [ ] N/A [ ]
MM32	Receive fault is implemented using latching high behavior	45.2.1.7.5		PMA*PLF:M	Yes [ ] N/A [ ]
MM33	Writes to register 9 are ignored by device that does not implement transmit disable	45.2.1.8		PMA*!PTD:M	Yes [ ] N/A [ ]
MM34	Single wavelength device uses bit 1.9.0 for transmit disable	45.2.1.8		PMA*PTD:M	Yes [ ] N/A [ ]
MM35	Single wavelength device ignores writes to bits 10:1 and returns a value of zero for them	45.2.1.8		PTD:M	Yes [ ] N/A [ ]
MM36	Bits 1 to 10, set to 1 disables transmission on corresponding lane	45.2.1.8		PTD:M	Yes [ ] No [ ] N/A [ ]
MM37	Bits 1 to 10, set to 0 enables transmission on corresponding lane	45.2.1.8		PTD:M	Yes [ ] No [ ] N/A [ ]
MM38	Setting bit 0 to a one disables transmission	45.2.1.8.7		PMA*PTD:M	Yes [ ] No [ ] N/A [ ]
MM39	Setting bit 0 to a zero enables transmission	45.2.1.8.7	Only if all lane transmit disables are zero	PMA*PTD:M	Yes [ ] No [ ] N/A [ ]
MM40	Receive signal detect register behaves as described	45.2.1.9		PMA:M	Yes [ ] N/A [ ]
MM41	Writes to the extended ability register have no effect	45.2.1.10		PMA:M	Yes [ ] N/A [ ]
MM42	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.1.11		PMA:M	Yes [ ] N/A [ ]
MM43	Writes to the 40G/100G PMA/PMD extended ability register have no effect	45.2.1.12		PMA:M	Yes [ ] N/A [ ]
MM44	Bit indicates link down while initializing	45.2.1.2.4		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM45	Bit remains a one and writing a one is ignored when link is up or initializing	45.2.1.11		PMA*10P*2B:M	Yes [ ] N/A [ ]
*MM46	PRBS pattern testing implemented	45.2.1.123	Ability indicated in register 1.1500	PMA:O	Yes [ ] N/A [ ]
MM47	Square wave testing implemented	45.2.1.123	Bit 1.1500.12 is set to one	PMA:O	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
MM48	Counters corresponding to lanes that are not implemented return all zeros	45.2.1.126		MM46:M	Yes [ ] N/A [ ]
MM49	Counters are reset to zero by read or PMA reset	45.2.1.126		MM46:M	Yes [ ] N/A [ ]
MM50	Counters are held at all ones in the case of overflow	45.2.1.126		MM46:M	Yes [ ] N/A [ ]
MM51	Bit set to zero upon reset or upon link down	45.2.1.15.1	-O subtypes only	PMA*10P*2B:M	Yes [ ] N/A [ ]
MM52	Bit set to one upon reset or upon link down	45.2.1.15.1	-R subtypes only	PMA*10P*2B:M	Yes [ ] N/A [ ]
MM53	Handshake tones not sent while bit is set to zero	45.2.1.15.1		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM54	Writes to set unsupported modes or when link is not down are ignored	45.2.1.15.4		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM55	Setting bit to one to one issues a cleardown command	45.2.1.15.5		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM56	MMD clears bit to zero when cleardown command is issued or on reset	45.2.1.15.5		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM57	Writes ignored if link is not in “Link down (ready)” state	45.2.1.15.5	Link state described in 45.2.1.16.4	PMA*10P*2B:M	Yes [ ] N/A [ ]
MM58	PMA/PMD does not respond to handshake tones while bit is set to one	45.2.1.15.6		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM59	PMA/PMD responds to handshake tones properly when bit is set to zero	45.2.1.15.6		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM60	Bit set to zero upon MMD reset	45.2.1.15.6		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM61	Writes to set unsupported modes are ignored	45.2.1.15.7		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM62	Bits zero when link is down or initializing	45.2.1.16.1		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM63	Bits set indicate linked port type or link status	45.2.1.16.4		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM64	Bits indicate 001 while link is initializing	45.2.1.16.4		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM65	Bits indicate 000 when link is down and handshake tones are not detected	45.2.1.16.4		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM66	Bits indicate 100 when link is down and handshake tones are detected	45.2.1.16.4		PMA*10P*2B:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
MM67	Bit held as one during operation, clears to zero after	45.2.1.17.1		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM68	Result = failed after 10 second timeout	45.2.1.17.1		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM69	Writes to one while link is down are marked completed and failed	45.2.1.17.1		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM70	Bit held as one during operation, clears to zero after	45.2.1.17.2		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM71	Result = failed after 10 second timeout	45.2.1.17.2		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM72	Writes to one while link is down are marked completed and failed	45.2.1.17.2		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM73	Bit set to result of the “Get” operation	45.2.1.18.1		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM74	Bit set to zero on read or reset	45.2.1.18.1		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM75	Bit set to result of the “Send” operation	45.2.1.18.2		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM76	Bits are reset to zero when read or on reset	45.2.1.19		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM77	Bits are held to all ones upon counter overflow	45.2.1.19		PMA*10P*2B:M	Yes [ ] N/A [ ]
MM78	Bits are reset to zero when read or reset	45.2.1.26		PMA*10P:M	Yes [ ] N/A [ ]
MM79	Bits are reset to zero when read or reset	45.2.1.27		PMA*10P:M	Yes [ ] N/A [ ]
MM80	Bits are held at all ones when PHY cannot determine value	45.2.1.30.1	ex: While link is down	PMA*10P:M	Yes [ ] N/A [ ]
MM81	Bit remain as one while tones are being refreshed	45.2.1.40.1		PMA*10P:M	Yes [ ] N/A [ ]
MM82	Bit set to zero when operation completes or upon reset	45.2.1.40.1		PMA*10P:M	Yes [ ] N/A [ ]
MM83	Bit remain as one while tones are being activated/deactivated	45.2.1.40.2		PMA*10P:M	Yes [ ] N/A [ ]
MM84	Bit set to zero when operation completes or upon reset	45.2.1.40.2		PMA*10P:M	Yes [ ] N/A [ ]
MM85	Bit remain as one while tone direction is being changed	45.2.1.40.3		PMA*10P:M	Yes [ ] N/A [ ]
MM86	Bit set to zero when operation completes or upon reset	45.2.1.40.3		PMA*10P:M	Yes [ ] N/A [ ]
MM87	Bit remain as one while SNR margins parameters are loaded	45.2.1.40.4		PMA*10P:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
MM88	Bit set to zero when operation completes or upon reset	45.2.1.40.4		PMA*10P:M	Yes [ ] N/A [ ]
MM89	Bit remain as one while PSD level is set	45.2.1.40.5		PMA*10P:M	Yes [ ] N/A [ ]
MM90	Bit set to zero when operation completes or upon reset	45.2.1.40.5		PMA*10P:M	Yes [ ] N/A [ ]
MM91	Bit remain as one while reference PSD level is set	45.2.1.40.6		PMA*10P:M	Yes [ ] N/A [ ]
MM92	Bit set to zero when operation completes or upon reset	45.2.1.40.6		PMA*10P:M	Yes [ ] N/A [ ]
MM93	Bits are reset to zero when read or upon reset	45.2.1.41.1		PMA*10P:M	Yes [ ] N/A [ ]
MM94	Bits read as zero	45.2.1.42.6		PMA*10P:M	Yes [ ] N/A [ ]
MM95	Bits read as zero	45.2.1.42.7		PMA*10P:M	Yes [ ] N/A [ ]
MM96	Writes to an invalid value are ignored	45.2.1.44	Valid values are decimal 10, 20, or 40	PMA*10P:M	Yes [ ] N/A [ ]
MM97	Bits set to default value on MMD reset	45.2.1.44	Default value is decimal 20	PMA*10P:M	Yes [ ] N/A [ ]
MM98	Writes to set an invalid value are ignored	45.2.1.47.1	Valid values are decimal 3 through 89	PMA*10P:M	Yes [ ] N/A [ ]
MM99	Writes to set an invalid value are ignored	45.2.1.47.2	Valid values are decimal 3 through 89	PMA*10P:M	Yes [ ] N/A [ ]
MM100	Writes to set an invalid value are ignored	45.2.1.47.3	Valid values are decimal 1through 86	PMA*10P:M	Yes [ ] N/A [ ]
MM101	Writes to set an invalid value are ignored	45.2.1.47.5	Invalid value is 11	PMA*10P:M	Yes [ ] N/A [ ]
MM102	Bits set to zero when read or reset	45.2.1.48		PMA*2B:M	Yes [ ] N/A [ ]
MM103	Bits set to zero when read or reset	45.2.1.50		PMA*2B:M	Yes [ ] N/A [ ]
MM104	Bits set to zero when read or reset	45.2.1.52		PMA*2B:M	Yes [ ] N/A [ ]
MM105	Bits set to zero when read or reset	45.2.1.54		PMA*2B:M	Yes [ ] N/A [ ]
MM106	Bits set to zero when read or reset	45.2.1.56		PMA*2B:M	Yes [ ] N/A [ ]
MM107	Writes to set an invalid value are ignored	45.2.1.61.1	Valid values are decimal 3 through 89	PMA*10P:M	Yes [ ] N/A [ ]
MM108	Writes to set an invalid value are ignored	45.2.1.61.2	Valid values are decimal 3 through 89	PMA*10P:M	Yes [ ] N/A [ ]
MM109	Writes to set an invalid value are ignored	45.2.1.61.3	Valid values are decimal 1 through 86	PMA*10P:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
MM110	Writes to set an invalid value are ignored	45.2.1.61.5	Invalid value is 11	PMA*10P:M	Yes [ ] N/A [ ]
MM111	Bit set to zero if PMA link_status=FAIL	45.2.1.62.1		PMA*10T:M	Yes [ ] N/A [ ]
MM112	Skew delay register update rate.	45.2.1.78	At least once per second	PMA*10T:M	Yes [ ] N/A [ ]
MM113	A PMA/PMD that is unable to detect a transmit fault returns a value of zero for the transmit fault bit.	45.2.1.89.3		KX:M	Yes [ ] N/A [ ]
MM114	The transmit fault bit is implemented with latching high behavior.	45.2.1.89.3		KX:M	Yes [ ] N/A [ ]
MM115	A PMA/PMD that is unable to detect a receive fault returns a value of zero for the receive fault bit.	45.2.1.89.4		KX:M	Yes [ ] N/A [ ]
MM116	The receive fault bit is implemented with latching high behavior.	45.2.1.89.4		KX:M	Yes [ ] N/A [ ]
MM117	FEC enable is set to zero upon execution of PHY reset	45.2.1.93.1		FEC-R:M	Yes [ ] N/A [ ]
MM118	RS-FEC counters are reset when read or upon PHY reset.	45.2.1.103, 45.2.1.104, 45.2.1.105, 45.2.1.107, 45.2.1.108, 45.2.1.109		RS-FEC:M	Yes [ ] N/A [ ]
MM119	RS-FEC counters are held at all ones in the case of overflow	45.2.1.103, 45.2.1.104, 45.2.1.105, 45.2.1.107, 45.2.1.108, 45.2.1.109		RS-FEC:M	Yes [ ] N/A [ ]
MM120	FEC corrected blocks counters are reset when read or upon PHY reset.	45.2.1.94, 45.2.1.116		FEC-R:M	Yes [ ] N/A [ ]
MM121	FEC corrected blocks counters are held at all ones in the case of overflow	45.2.1.94, 45.2.1.116		FEC-R:M	Yes [ ] N/A [ ]
MM122	FEC uncorrected blocks counters are reset when read or upon PHY reset.	45.2.1.95, 45.2.1.117		FEC-R:M	Yes [ ] N/A [ ]
MM123	FEC uncorrected blocks counters are held at all ones in the case of overflow	45.2.1.95, 45.2.1.117		FEC-R:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
MM124	Backplane Ethernet AN ability bit is set to 1 for 1000BASE-KX, 10GBASE-KX4 and 10GBASE-KR PHYs.	45.2.7.12.3		KX:M KX4:M KR:M	Yes [ ] N/A [ ]
MM125	Writes to this register have no effect	45.2.1.11		PMA:M	Yes [ ] No [ ] N/A [ ]

45.5.3.4 WIS options

Item	Feature	Subclause	Value/Comment	Status	Support
*WPT	Implementation of PRBS31 pattern testing	45.2.2		WIS:O	Yes [ ] No [ ] N/A [ ]

45.5.3.5 WIS management functions

Item	Feature	Subclause	Value/Comment	Status	Support
WM1	Device responds to all register addresses for that device	45.2		WIS:M	Yes [ ] N/A [ ]
WM2	Writes to undefined and read-only registers have no effect	45.2		WIS:M	Yes [ ] N/A [ ]
WM3	Operation is not affected by writes to reserved and unsupported bits.	45.2		WIS:M	Yes [ ] N/A [ ]
WM4	Reserved and unsupported bits return a value of zero	45.2		WIS:M	Yes [ ] N/A [ ]
WM5	Latching low bits remain low until after they have been read via the management interface	45.2		WIS:M	Yes [ ] N/A [ ]
WM6	Latching low bits assume correct value once read via the management interface	45.2	Correct value is based upon current state	WIS:M	Yes [ ] N/A [ ]
WM7	Latching high bits remain high until after they have been read via the management interface	45.2		WIS:M	Yes [ ] N/A [ ]
WM8	Latching high bits assume correct value once read via the management interface	45.2	Correct value is based upon current state	WIS:M	Yes [ ] N/A [ ]
WM9	Action on reset	45.2.2.1.1	Reset the registers of the entire device to default values and set bit 15 of the Control register to one	WIS:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
WM10	Return 1 until reset completed	45.2.2.1.1		WIS:M	Yes [ ] N/A [ ]
WM11	Reset completes within 0.5 s	45.2.2.1.1		WIS:M	Yes [ ] N/A [ ]
WM12	Responds to reads of bits 2.0.15 and 2.8.15:14 during reset	45.2.2.1.1		WIS:M	Yes [ ] N/A [ ]
WM13	Loopback mode	45.2.2.1.2	Whenever bit 2.0.14 is set to a one	WIS:M	Yes [ ] N/A [ ]
WM14	Data received from PMA ignored during loopback	45.2.2.1.2		WIS:M	Yes [ ] N/A [ ]
WM15	Transmit data returned on receive path during loopback	45.2.2.1.2		WIS:M	Yes [ ] N/A [ ]
WM16	Device responds to transactions necessary to exit low-power mode while in low- power state	45.2.2.1.3		WIS:M	Yes [ ] N/A [ ]
WM17	Speed selection bits 13 and 6 are written as one	45.2.2.1.4		WIS:M	Yes [ ] N/A [ ]
WM18	Invalid writes to speed selection bits are ignored	45.2.2.1.4		WIS:M	Yes [ ] N/A [ ]
WM19	Writes to status 1 register have no effect	45.2.2.2		WIS:M	Yes [ ] N/A [ ]
WM20	Fault bit implemented using latching high behavior	45.2.2.2.1		WIS:M	Yes [ ] N/A [ ]
WM21	Link status bit implemented using latching low behavior	45.2.2.2.2		WIS:M	Yes [ ] N/A [ ]
WM22	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.2.3		WIS:M	Yes [ ] N/A [ ]
WM23	Setting bit 2.7.5 to a one enables PRBS31 receive pattern testing if bit 2.8.1 is a one and bit 2.7.2 is not a one	45.2.2.6.1		WIS* WPT:M	Yes [ ] N/A [ ]
WM24	Setting bit 2.7.5 to a zero disables PRBS31 receive pattern testing	45.2.2.6.1		WIS* WPT:M	Yes [ ] N/A [ ]
WM25	Setting bit 2.7.4 to a one enables PRBS31 transmit pattern testing if bit 2.8.1 is a one and bit 2.7.1 is not a one	45.2.2.6.2		WIS* WPT:M	Yes [ ] N/A [ ]
WM26	Setting bit 2.7.4 to a zero disables PRBS31 transmit pattern testing	45.2.2.6.2		WIS* WPT:M	Yes [ ] N/A [ ]
WM27	Setting bit 3 to one selects the square wave test pattern	45.2.2.6.3		WIS:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
WM28	Setting bit 3 to zero selects the pseudo random test pattern	45.2.2.6.3		WIS:M	Yes [ ] N/A [ ]
WM29	Setting bit 2 to one enables receive pattern testing	45.2.2.6.4		WIS:M	Yes [ ] N/A [ ]
WM30	Setting bit 2 to zero disables receive pattern testing	45.2.2.6.4		WIS:M	Yes [ ] N/A [ ]
WM31	Setting bit 1 to one enables transmit pattern testing	45.2.2.6.5		WIS:M	Yes [ ] N/A [ ]
WM32	Setting bit 1 to zero disables transmit pattern testing	45.2.2.6.5		WIS:M	Yes [ ] N/A [ ]
WM33	Setting bit 0 to a one enables 10GBASE-W logic and sets interface speed	45.2.2.6.6		WIS:M	Yes [ ] N/A [ ]
WM34	Setting bit 0 to a zero disables 10GBASE-W logic, sets interface speed. and bypasses data	45.2.2.6.6		WIS:O	Yes [ ] N/A [ ]
WM35	Writes to bit are ignored by WIS not capable of supporting 10GBASE-R	45.2.2.6.6		WIS:M	Yes [ ] N/A [ ]
WM36	Bit returns one when read if WIS is not capable of supporting 10GBASE-R	45.2.2.6.6		WIS:M	Yes [ ] N/A [ ]
WM37	Writes to status 2 register have no effect	45.2.2.7		WIS:M	Yes [ ] N/A [ ]
WM38	Counter is reset to all zeros when read or reset	45.2.2.8		WIS* WPT:M	Yes [ ] N/A [ ]
WM39	Counter is held at all ones at overflow	45.2.2.8		WIS* WPT:M	Yes [ ] N/A [ ]
WM40	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.2.9		WIS:M	Yes [ ] N/A [ ]
WM41	Writes to Status 3 register have no effect	45.2.2.10		WIS:M	Yes [ ] N/A [ ]
WM42	SEF bit implemented using latching high behavior	45.2.2.10.1		WIS:M	Yes [ ] N/A [ ]
WM43	Far end PLM-P/LCD-P bit implemented using latching high behavior	45.2.2.10.2		WIS:M	Yes [ ] N/A [ ]
WM44	Far end AIS-P/LOP-P bit implemented using latching high behavior	45.2.2.10.3		WIS:M	Yes [ ] N/A [ ]
WM45	LOF bit implemented using latching high behavior	45.2.2.10.4		WIS:M	Yes [ ] N/A [ ]
WM46	LOS bit implemented using latching high behavior	45.2.2.10.5		WIS:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
WM47	RDI-L bit implemented using latching high behavior	45.2.2.10.6		WIS:M	Yes [ ] N/A [ ]
WM48	AIS-L bit implemented using latching high behavior	45.2.2.10.7		WIS:M	Yes [ ] N/A [ ]
WM49	LCD-P bit implemented using latching high behavior	45.2.2.10.8		WIS:M	Yes [ ] N/A [ ]
WM50	PLM-P bit implemented using latching high behavior	45.2.2.10.9		WIS:M	Yes [ ] N/A [ ]
WM51	AIS-P bit implemented using latching high behavior	45.2.2.10.10		WIS:M	Yes [ ] N/A [ ]
WM52	LOP-P bit implemented using latching high behavior	45.2.2.10.11		WIS:M	Yes [ ] N/A [ ]

45.5.3.6 PCS options

Item	Feature	Subclause	Value/Comment	Status	Support
*100CR	Implementation of 100GBASE-R PCS	45.2.3.7		PCS:O	Yes [ ] No [ ] N/A [ ]
*40CR	Implementation of 40GBASE-R PCS	45.2.3.7		PCS:O	Yes [ ] No [ ] N/A [ ]
*CR	Implementation of 10GBASE-R PCS	45.2.3		PCS:O	Yes [ ] No [ ] N/A [ ]
*XCR	Implementation of 10/40/ 100GBASE-R PCS	45.2.3.7		100CR:M 40CR:M 10CR:M	Yes [ ] No [ ] N/A [ ]
*CT	Implementation of the 10GBASE-T PCS	45.2.3		PCS:O	Yes [ ] No [ ] N/A [ ]
*CX	Implementation of 10GBASE-X PCS	45.2.3		PCS:O	Yes [ ] No [ ] N/A [ ]
*XP	Implementation of 10GBASE-X pattern testing	45.2.3		CX:O	Yes [ ] No [ ] N/A [ ]
*PPT	Implementation of PRBS31 pattern testing	45.2.3		PCS:O	Yes [ ] No [ ] N/A [ ]
*PTT	Implementation of PRBS9 pattern testing	45.2.3		PCS:O	Yes [ ] No [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
*EPC	Implementation of the 10BASE-TS/2BASE-TL PCS	45.2.3.25		PCS:O	Yes [ ] No [ ] N/A [ ]
*PAF	Implementation of the PME aggregation function	45.2.3.25		EPC:O	Yes [ ] No [ ] N/A [ ]
*CPR	Implementation of 10GBASE-PR or 10/1GBASE-PRX PCS	45.2.3		PCS:O	Yes [ ] No [ ] N/A [ ]
*EEE	Implementation of EEE			O	Yes [ ] N/A [ ]

45.5.3.7 PCS management functions

Item	Feature	Subclause	Value/Comment	Status	Support
RM1	Device responds to all register addresses for that device	45.2		PCS:M	Yes [ ] N/A [ ]
RM2	Writes to undefined and read-only registers have no effect	45.2		PCS:M	Yes [ ] N/A [ ]
RM3	Operation is not affected by writes to reserved and unsupported bits	45.2		PCS:M	Yes [ ] N/A [ ]
RM4	Reserved and unsupported bits return a value of zero	45.2		PCS:M	Yes [ ] N/A [ ]
RM5	Latching low bits remain low until after they have been read via the management interface	45.2		PCS:M	Yes [ ] N/A [ ]
RM6	Latching low bits assume correct value once read via the management interface	45.2	Correct value is based upon current state	PCS:M	Yes [ ] N/A [ ]
RM7	Latching high bits remain high until after they have been read via the management interface	45.2		PCS:M	Yes [ ] N/A [ ]
RM8	Latching high bits assume correct value once read via the management interface	45.2	Correct value is based upon current state	PCS:M	Yes [ ] N/A [ ]
RM9	Action on reset	45.2.3.1.1	Reset the registers of the entire device to default values and set bit 15 of the Control register to one	PCS:M	Yes [ ] N/A [ ]
RM10	Return 1 until reset completed	45.2.3.1.1		PCS:M	Yes [ ] N/A [ ]
RM11	Reset completes within 0.5 s	45.2.3.1.1		PCS:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
RM12	Device responds to reads of register bits 3.0.15 and 3.5.15:14 during reset	45.2.3.1.1		PCS:M	Yes [ ] N/A [ ]
RM13	Loopback mode	45.2.3.1.2	Whenever bit 3.0.14 is set to a one	PCS:M	Yes [ ] N/A [ ]
RM14	Transmit data is returned on the receive path during loopback	45.2.3.1.2		PCS:M	Yes [ ] N/A [ ]
RM15	Writes to loopback bit are ignored when operating at 10 Gb/s with port type selections other than 10GBASE-R or 10GBASE-T	45.2.3.1.2		PCS:M	Yes [ ] N/A [ ]
RM16	Loopback bit returns zero when operating at 10 Gb/s with port type selections other than 10GBASE-R or 10GBASE-T	45.2.3.1.2		PCS:M	Yes [ ] N/A [ ]
RM17	Device responds to transactions necessary to exit low-power mode while in low- power state	45.2.3.1.3		PCS:M	Yes [ ] N/A [ ]
RM18	Speed selection bits 13 and 6 are written as one	45.2.3.1.5		PCS:M	Yes [ ] N/A [ ]
RM19	Invalid writes to speed selection bits are ignored	45.2.3.1.5		PCS:M	Yes [ ] N/A [ ]
RM20	Writes to PCS status 1 register have no effect	45.2.3.2		PCS:M	Yes [ ] N/A [ ]
RM21	Receive link status implemented using latching low behavior	45.2.3.2.7		PCS:M	Yes [ ] N/A [ ]
RM22	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.3.3		PCS:M	Yes [ ] N/A [ ]
RM23	PCS type is selected using bits 1 through 0	45.2.3.6.1		PCS:M	Yes [ ] N/A [ ]
RM24	Writes to the type selection bits that select types that have not been advertised are ignored	45.2.3.6.1		PCS:M	Yes [ ] N/A [ ]
RM25	Writes to PCS status 2 register have no effect	45.2.3.7		PCS:M	Yes [ ] N/A [ ]
RM26	Transmit fault implemented with latching high behavior	45.2.3.7.2		PCS:M	Yes [ ] N/A [ ]
RM27	Receive fault implemented with latching high behavior	45.2.3.7.3		PCS:M	Yes [ ] N/A [ ]
RM28	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.3.8		PCS:M	Yes [ ] N/A [ ]
RM29	Writes to 10GBASE-X PCS status register have no effect	45.2.3.11		PCS* CX:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
RM30	Register returns zero if the PCS does not implement the 10GBASE-X port type	45.2.3.11		PCS* !CX:M	Yes [ ] N/A [ ]
RM31	EEE capability indicated for each port type	45.2.3.9		EEE:M	Yes [ ] N/A [ ]
RM32	EEE wake error counter behavior as specified	45.2.3.10		EEE:M	Yes [ ] N/A [ ]
RM33	Writes to bit are ignored and reads return a value of zero	45.2.3.12.1		PCS* PX:M	Yes [ ] N/A [ ]
RM34	Setting the bits to <10> selects the mixed frequency pattern	45.2.3.12.2		PCS* PX:M	Yes [ ] N/A [ ]
RM35	Setting the bits to <01> selects the low-frequency pattern	45.2.3.12.2		PCS* PX:M	Yes [ ] N/A [ ]
RM36	Setting the bits to <00> selects the high-frequency pattern	45.2.3.12.2		PCS* PX:M	Yes [ ] N/A [ ]
RM37	Writes to BASE-R and 10GBASE-T PCS status 1 register have no effect	45.2.3.13		CR:M CT:M	Yes [ ] N/A [ ]
RM38	Reads from BASE-R and 10GBASE-T PCS status 1 register return zero for PCS that does not support 10GBASE-R	45.2.3.13		CR:M CT:M	Yes [ ] N/A [ ]
RM39	Writes to BASE-R and 10GBASE-T PCS status 2 register have no effect	45.2.3.14		XCR:M CT:M	Yes [ ] N/A [ ]
RM40	Reads from BASE-R and 10GBASE-T PCS status 2 register return zero for PCS that does not support 10/40/ 100GBASE-R or 10GBASE-T	45.2.3.14		XCR:M CT:M	Yes [ ] N/A [ ]
RM41	Latched block lock implemented with latching low behavior	45.2.3.14.1		XCR:M CT:M	Yes [ ] N/A [ ]
RM42	Latched high BER implemented with latching high behavior	45.2.3.14.2		XCR:M CT:M	Yes [ ] N/A [ ]
RM43	BER counter resets to all zeros on read or reset	45.2.3.14.3		XCR:M CT:M	Yes [ ] N/A [ ]
RM44	BER counter holds at all ones at overflow	45.2.3.14.3		!RM53:M	Yes [ ] N/A [ ]
RM45	Errored blocks counter implemented as a non roll over counter	45.2.3.14.4		!RM57:M	Yes [ ] N/A [ ]
RM46	Errored blocks counter resets to all zeros on read	45.2.3.14.4		XCR:M CT:M	Yes [ ] N/A [ ]
RM47	Setting bit 3.42.5 to a one enables PRBS31 receive pattern testing if bit 3.32.2 is a one and bit 3.42.2 is not a one	45.2.3.17.3		PCS* PPT:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
RM48	Setting bit 3.42.5 to a zero disables PRBS31 receive pattern testing	45.2.3.17.3		PCS* PPT:M	Yes [ ] N/A [ ]
RM49	Setting bit 3.42.4 to a one enables PRBS31 transmit pattern testing if bit 3.32.2 is a one and bit 3.42.3 is not a one	45.2.3.17.4		PCS* PPT:M	Yes [ ] N/A [ ]
RM50	Setting bit 3.42.4 to a zero disables PRBS31 transmit pattern testing	45.2.3.17.4		PCS* PPT:M	Yes [ ] N/A [ ]
RM51	Test-pattern error counter resets to all zeros on read or reset	45.2.3.18		XCR:M	Yes [ ] N/A [ ]
RM52	Test-pattern error counter holds at all ones at overflow	45.2.3.18		XCR:M	Yes [ ] N/A [ ]
*RM53	BER high order counter implemented	45.2.3.19		CR:O 40CR:M 100CR:M	Yes [ ] N/A [ ]
RM54	High order bits latched on read to 3.33	45.2.3.19		RM53:M	Yes [ ] N/A [ ]
RM55	Counter reset on read to 3.33 or PCS reset	45.2.3.19		RM53:M	Yes [ ] N/A [ ]
RM56	Counter held at all ones in the case of an overflow	45.2.3.19		RM53:M	Yes [ ] N/A [ ]
*RM57	Errored blocks high order counter implemented	45.2.3.20		CR:O 40CR:M 100CR:M	Yes [ ] N/A [ ]
RM58	Register bit 3.45.15 set to 1	45.2.3.20		RM57:M	Yes [ ] N/A [ ]
RM59	High order bits latched on read to 3.33	45.2.3.20		RM57:M	Yes [ ] N/A [ ]
RM60	Counter reset on read to 3.33 or PCS reset	45.2.3.20		RM57:M	Yes [ ] N/A [ ]
RM61	Counter held at all ones in the case of an overflow	45.2.3.20		RM57:M	Yes [ ] N/A [ ]
RM62	Bit indicates fault when any PCS register indicates fault	45.2.3.2.5	If subtype is supported	PCS*EPC:O	Yes [ ] No [ ] N/A [ ]
RM63	Writes to multi-lane BASE-R PCS alignment status 1 register have no effect	45.2.3.21		XCR:M	Yes [ ] N/A [ ]
RM64	Non multi-lane BASE-R device shall return all zeros	45.2.3.21		!XCR:M	Yes [ ] N/A [ ]
RM65	Device shall return a zero for lanes not required	45.2.3.21		XCR:M	Yes [ ] N/A [ ]
RM66	Writes to multi-lane BASE-R PCS alignment status 2 register have no effect	45.2.3.22		XCR:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
RM67	Non multi-PCS lane BASE-R device shall return all zeros	45.2.3.22		!XCR:M	Yes [ ] N/A [ ]
RM68	Device shall return a zero for lanes not required	45.2.3.22		XCR:M	Yes [ ] N/A [ ]
RM69	Writes to multi-lane BASE-R PCS alignment status 3 register have no effect	45.2.3.23		XCR:M	Yes [ ] N/A [ ]
RM70	Non multi-PCS lane BASE-R device shall return all zeros	45.2.3.23		!XCR:M	Yes [ ] N/A [ ]
RM71	Device shall return a zero for lanes not required	45.2.3.23		XCR:M	Yes [ ] N/A [ ]
RM72	Writes to multi-lane BASE-R PCS alignment status 4 register have no effect	45.2.3.24		XCR:M	Yes [ ] N/A [ ]
RM73	Non multi-PCS lane BASE-R device shall return all zeros	45.2.3.24		!XCR:M	Yes [ ] N/A [ ]
RM74	Device shall return a zero for lanes not required	45.2.3.24		XCR:M	Yes [ ] N/A [ ]
RM75	Counters reset on read to 3.200 through 3.219 or PCS reset	45.2.3.44		XCR:M	Yes [ ] N/A [ ]
RM76	Counters held at all ones in the case of an overflow	45.2.3.44		XCR:M	Yes [ ] N/A [ ]
RM77	Writing this bit to a one activates the PAF when link is established	45.2.3.26.3		PCS*PAF:M	Yes [ ] N/A [ ]
RM78	Writes to bit are ignored while link is active or initializing or if PAF is not supported	45.2.3.26.3		PCS*PAF:M	Yes [ ] N/A [ ]
RM79	Bits indicate device capabilities upon reset	45.2.3.27		PCS*PAF:M	Yes [ ] N/A [ ]
RM80	Writes to the register through any PCS MMD in the same package effect the register equally	45.2.3.27		PCS*PAF:M	Yes [ ] N/A [ ]
RM81	Single bit set to one and all others cleared to zero when device does not support aggregation of multiple PMEs	45.2.3.27		PCS:M	Yes [ ] N/A [ ]
RM82	PME aggregation used if one or more bits set	45.2.3.28		PCS*PAF:M	Yes [ ] N/A [ ]
RM83	Registers set to all zeros upon reset	45.2.3.28		PCS*PAF:M	Yes [ ] N/A [ ]
RM84	Writes to the register through any PCS MMD in the same package effect the register equally	45.2.3.28		PCS*PAF:M	Yes [ ] N/A [ ]
RM85	Bits reset to zero when read or upon MMD reset	45.2.3.29		PCS*PAF:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
RM86	Bits held to one upon overflow	45.2.3.29		PCS*PAF:M	Yes [ ] N/A [ ]
RM87	Bits reset to zero when read or upon MMD reset	45.2.3.30		PCS*PAF:M	Yes [ ] N/A [ ]
RM88	Bits held to one upon overflow	45.2.3.30		PCS*PAF:M	Yes [ ] N/A [ ]
RM89	Bits reset to zero when read or upon MMD reset	45.2.3.31		PCS*PAF:M	Yes [ ] N/A [ ]
RM90	Bits held to one upon overflow	45.2.3.31		PCS*PAF:M	Yes [ ] N/A [ ]
RM91	Bits reset to zero when read or upon MMD reset	45.2.3.32		PCS*PAF:M	Yes [ ] N/A [ ]
RM92	Bits held to one upon overflow	45.2.3.32		PCS*PAF:M	Yes [ ] N/A [ ]
RM93	Bits reset to zero when read or upon MMD reset	45.2.3.33		PCS*PAF:M	Yes [ ] N/A [ ]
RM94	Bits held to one upon overflow	45.2.3.33		PCS*PAF:M	Yes [ ] N/A [ ]
RM95	Bits reset to zero when read or upon MMD reset	45.2.3.34		PCS*PAF:M	Yes [ ] N/A [ ]
RM96	Bits held to one upon overflow	45.2.3.34		PCS*PAF:M	Yes [ ] N/A [ ]
RM97	Bits reset to zero when read or upon MMD reset	45.2.3.35		PCS*PAF:M	Yes [ ] N/A [ ]
RM98	Bits held to one upon overflow	45.2.3.35		PCS*PAF:M	Yes [ ] N/A [ ]
RM99	Bits reset to zero when read or upon MMD reset	45.2.3.36		PCS*PAF:M	Yes [ ] N/A [ ]
RM100	Bits held to one upon overflow	45.2.3.36		PCS*PAF:M	Yes [ ] N/A [ ]
RM101	Setting bit 3.42.6 to a one enables PRBS9 transmit pattern testing if bit 3.32.3 is a one and bit 3.42.3 is not a one and bit 3.42.4 is not a one	45.2.3.17.2		PCS* PTT:M	Yes [ ] N/A [ ]
RM102	Setting bit 3.42.6 to a zero disables PRBS9 transmit pattern testing	45.2.3.17.2		PCS* PTT:M	Yes [ ] N/A [ ]
RM103	Corrected FEC codeword counter is reset when read or upon PHY reset	45.2.3.39		CPR:M	Yes [ ] No [ ] N/A [ ]
RM104	Corrected FEC codeword counter is held at all ones in the case of overflow	45.2.3.39		CPR:M	Yes [ ] No [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
RM105	Uncorrected FEC codeword counter is reset when read or upon PHY reset	45.2.3.40		CPR:M	Yes [ ] No [ ] N/A [ ]
RM106	Uncorrected FEC codeword counter is held at all ones in the case of overflow	45.2.3.40		CPR:M	Yes [ ] No [ ] N/A [ ]

45.5.3.8 Auto-Negotiation options

Item	Feature	Subclause	Value/Comment	Status	Support
*AT	Implementation of 10GBASE-T Auto-Negotiation	45.2.7		AN:O	Yes [ ] No [ ] N/A [ ]
*AB	Implementation of Backplane Ethernet Auto-negotiation	45.2.7		AN:M	Yes [ ] No [ ] N/A [ ]
*ABN	Implementation of Next Page support for Backplane Ethernet Auto-negotiation	45.2.7		AB:M	Yes [ ] No [ ] N/A [ ]
*EEE	Implementation of EEE			O	Yes [ ] N/A [ ]

45.5.3.9 Auto-Negotiation management functions

Item	Feature	Subclause	Value/Comment	Status	Support
AM1	Device responds to all register addresses for that device	45.2		AN:M	Yes [ ] N/A [ ]
AM2	Writes to undefined and read-only registers have no effect	45.2		AN:M	Yes [ ] N/A [ ]
AM3	Operation is not affected by writes to reserved and unsupported bits	45.2		AN:M	Yes [ ] N/A [ ]
AM4	Reserved and unsupported bits return a value of zero	45.2		AN:M	Yes [ ] N/A [ ]
AM5	Latching low bits remain low until after they have been read via the management interface	45.2		AN:M	Yes [ ] N/A [ ]
AM6	Latching low bits assume correct value once read via the management interface	45.2	Correct value is based upon current state	AN:M	Yes [ ] N/A [ ]
AM7	Latching high bits remain high until after they have been read via the management interface	45.2		AN:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
AM8	Latching high bits assume correct value once read via the management interface	45.2	Correct value is based upon current state	AN:M	Yes [ ] N/A [ ]
AM9	Action on reset	45.2.7.1.1	Reset the registers of the entire device to default values and set bit 15 of the Control register to one	AN:M	Yes [ ] N/A [ ]
AM10	Return 1 until reset completed	45.2.7.1.1		AN:M	Yes [ ] N/A [ ]
AM11	Reset completes within 0.5 s	45.2.7.1.1		AN:M	Yes [ ] N/A [ ]
AM12	Device responds to reads of register bit 7.0.15	45.2.7.1.1	All other register bits should be ignored	AN:M	Yes [ ] N/A [ ]
AM13	Bit set enables Extended Next Page exchange	45.2.7.1.2		AN:M	Yes [ ] N/A [ ]
AM14	Bit 7.0.13 has no effect unless bit 7.16.12 is set	45.2.7.1.2		AN:M	Yes [ ] N/A [ ]
AM15	Auto-Negotiation enabled by setting bit 7.0.12	45.2.7.1.3		AN:M	Yes [ ] N/A [ ]
AM16	Bits 1.0.13, 1.0.6, 1.0.5:2, and 1.7.3:0 have no effect and Auto-Negotiation determines link configuration	45.2.7.1.3		AN:M	Yes [ ] N/A [ ]
AM17	Writes are ignored and bit is set to zero if Auto-Negotiation is disabled or unsupported	45.2.7.1.4		AN:M	Yes [ ] N/A [ ]
AM18	Setting bit to one restarts Auto-Negotiation if supported and enabled	45.2.7.1.4		AN:M	Yes [ ] N/A [ ]
AM19	Bit set until Auto-Negotiation started then self cleared to zero	45.2.7.1.4		AN:M	Yes [ ] N/A [ ]
AM20	Setting this bit restarts Auto-Negotiation if previously completed	45.2.7.1.4		AN:M	Yes [ ] N/A [ ]
AM21	Setting the bit to zero has no effect on Auto-Negotiation process	45.2.7.1.4		AN:M	Yes [ ] N/A [ ]
AM22	Writing the bit to one is ignored if 7.1.3 = 0 or Auto-Negotiation is disabled	45.2.7.1.4		AN:M	Yes [ ] N/A [ ]
AM23	Writes to register ignored	45.2.7.2		AN:M	Yes [ ] N/A [ ]
AM24	Bit set to zero when bits 7.0.13 or 7.16.12 are set to zero	45.2.7.2.2		AN:M	Yes [ ] N/A [ ]
AM25	Extended Next Page not used when bit set to zero	45.2.7.2.2		AN:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
AM26	Bit set to one with link codeword received and stored	45.2.7.2.3		AN:M	Yes [ ] N/A [ ]
AM27	Bit set to zero upon read of register 7.1	45.2.7.2.3		AN:M	Yes [ ] N/A [ ]
AM28	Bit set to zero when bit 7.0.12 is set to zero	45.2.7.2.4		AN:M	Yes [ ] N/A [ ]
AM29	Bit set to zero if device does not support Auto-Negotiation	45.2.7.2.4		AN:M	Yes [ ] N/A [ ]
AM30	Bit remains set until cleared	45.2.7.2.5		AN:M	Yes [ ] N/A [ ]
AM31	Bit cleared to zero upon reset or read of register	45.2.7.2.5		AN:M	Yes [ ] N/A [ ]
AM32	Bit remains set until cleared by register read	45.2.7.2.7		AN:M	Yes [ ] N/A [ ]
AM33	Bit cleared upon AN reset	45.2.7.2.7		AN:M	Yes [ ] N/A [ ]
AM34	Bit set to one if link partner is able to participate in Auto-Negotiation	45.2.7.2.8		AN:M	Yes [ ] N/A [ ]
AM35	Bit set to zero if link partner is not able to participate in Auto-Negotiation	45.2.7.2.8		AN:M	Yes [ ] N/A [ ]
AM36	Unique Identifier is composed of 22 bits of OUI, model number and revision	45.2.7.5		AN:M	Yes [ ] N/A [ ]
AM37	Prior to Auto-Negotiation bits 7.16.12:5 accurately reflect device capability	45.2.7.6		AN:M	Yes [ ] N/A [ ]
AM38	Write has no effect	45.2.7.7		AN:M	Yes [ ] N/A [ ]
AM39	At power up or AN reset message code set to null message	45.2.7.8		AN:M	Yes [ ] N/A [ ]
AM40	Write has no effect	45.2.7.9		AN:M	Yes [ ] N/A [ ]
AM41	Auto-Negotiation used to determine MASTER or SLAVE when bit set to zero	45.2.7.10.1		AN:M	Yes [ ] N/A [ ]
AM42	Operates as MASTER when bit set to one	45.2.7.10.2		AN:M	Yes [ ] N/A [ ]
AM43	Operates as SLAVE when bit set to zero	45.2.7.10.2		AN:M	Yes [ ] N/A [ ]
AM44	Advertise 10GBASE-T PHY capability when bit is set to one	45.2.7.10.4		AN:M	Yes [ ] N/A [ ]
AM45	10GBASE-T PHY capability not advertised when bit is set to zero	45.2.7.10.4		AN:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
AM46	Advertise loop timing capability when bit is set to one	45.2.7.10.7		AN:M	Yes [ ] N/A [ ]
AM47	Loop timing capability not advertised when bit is set to zero	45.2.7.10.7		AN:M	Yes [ ] N/A [ ]
AM48	Write has no effect	45.2.7.11		AN:M	Yes [ ] N/A [ ]
AM49	Bit set if MASTER-SLAVE cannot be successfully concluded	45.2.7.11.1		AN:M	Yes [ ] N/A [ ]
AM50	Bit set to zero on read of register 7.33	45.2.7.11.1		AN:M	Yes [ ] N/A [ ]
AM51	Bit set to zero by 10GBASE-T PMA reset	45.2.7.11.1		AN:M	Yes [ ] N/A [ ]
AM52	Bit set to zero when Auto-Negotiation is enabled	45.2.7.11.1		AN:M	Yes [ ] N/A [ ]
AM53	Bit set to one if local receiver is OK as defined in 55.2.2.7	45.2.7.11.3		AN:M	Yes [ ] N/A [ ]
AM54	Bit is set to one if the remote receiver status is OK as defined in 55.2.2.8	45.2.7.11.4		AN:M	Yes [ ] N/A [ ]
AM55	Parallel detection fault (7.1.9)	45.2.7.2.1	ONE to indicate more than one Backplane Ethernet PMA indicates link_status of OK when autoneg_wait_timer expires	AB:M	Yes [ ] N/A [ ]
AM56	Parallel detection fault clearing	45.2.7.2.1	ZERO on read of the AN status register.	AN:M	Yes [ ] N/A [ ]
AM57	Bit 7.48.0 set to 1	45.2.7.12.3	Set to 1 if KX, KX4, or KR PHY is implemented	AB:M	Yes [ ] N/A [ ]
AM58	EEE capability in advertisement register for each port type	45.2.7.13		AN:EEE:M	Yes [ ] N/A [ ]
AM59	EEE LP advertisement register reflects link partner's capabilities	45.2.7.14		AN:EEE:M	Yes [ ] N/A [ ]
AM60	Writes to EEE LP advertisement register have no effect	45.2.7.14		AN:EEE:M	Yes [ ] N/A [ ]

45.5.3.10 PHY XS options

Item	Feature	Subclause	Value/Comment	Status	Support
*PL	Implementation of loopback	45.2.4		PX:O	Yes [ ] No [ ] N/A [ ]
*PT	Implementation of pattern testing	45.2.4		PX:O	Yes [ ] No [ ] N/A [ ]

45.5.3.11 PHY XS management functions

Item	Feature	Subclause	Value/Comment	Status	Support
PM1	Device responds to all register addresses for that device	45.2		PX:M	Yes [ ] N/A [ ]
PM2	Writes to undefined and read-only registers have no effect	45.2		PX:M	Yes [ ] N/A [ ]
PM3	Operation is not affected by writes to reserved and unsupported bits	45.2		PX:M	Yes [ ] N/A [ ]
PM4	Reserved and unsupported bits return a value of zero	45.2		PX:M	Yes [ ] N/A [ ]
PM5	Latching low bits remain low until after they have been read via the management interface	45.2		PX:M	Yes [ ] N/A [ ]
PM6	Latching low bits assume correct value once read via the management interface	45.2	Correct value is based upon current state	PX:M	Yes [ ] N/A [ ]
PM7	Latching high bits remain high until after they have been read via the management interface	45.2		PX:M	Yes [ ] N/A [ ]
PM8	Latching high bits assume correct value once read via the management interface	45.2	Correct value is based upon current state	PX:M	Yes [ ] N/A [ ]
PM9	Action on reset	45.2.4.1.1	Reset the registers of the entire device to default values and set bit 15 of the Control register to one	PX:M	Yes [ ] N/A [ ]
PM10	Return 1 until reset completed	45.2.4.1.1		PX:M	Yes [ ] N/A [ ]
PM11	Reset completes within 0.5 s	45.2.4.1.1		PX:M	Yes [ ] N/A [ ]
PM12	Device responds to reads of bits 4.0.15 and 4.8.15:14 during reset	45.2.4.1.1		PX:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
PM13	Loopback mode	45.2.4.1.2	Whenever bit 4.0.14 is set to a one	PX*PL:M	Yes [ ] N/A [ ]
PM14	Receive data is returned on transmit path during loopback	45.2.4.1.2		PX*PL:M	Yes [ ] N/A [ ]
PM15	Writes to loopback bit are ignored and reads return zero	45.2.4.1.2		PX*!PL:M	Yes [ ] N/A [ ]
PM16	Device responds to transactions necessary to exit low-power mode while in low-power state	45.2.4.1.3		PX:M	Yes [ ] N/A [ ]
PM17	Speed selection bits 13 and 6 are written as one	45.2.4.1.6		PX:M	Yes [ ] N/A [ ]
PM18	Invalid writes to speed selection bits are ignored	45.2.4.1.6		PX:M	Yes [ ] N/A [ ]
PM19	Writes to status 1 register have no effect	45.2.4.2		PX:M	Yes [ ] N/A [ ]
PM20	Transmit link status implemented using latching low behavior	45.2.4.2.7		PX:M	Yes [ ] N/A [ ]
PM21	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.4.3		PX:M	Yes [ ] N/A [ ]
PM22	Writes to status 2 register have no effect	45.2.4.6		PX:M	Yes [ ] N/A [ ]
PM23	Transmit fault implemented with latching high behavior	45.2.4.6.2		PX:M	Yes [ ] N/A [ ]
PM24	Receive fault implemented with latching high behavior	45.2.4.6.3		PX:M	Yes [ ] N/A [ ]
PM25	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.4.7		PX:M	Yes [ ] N/A [ ]
PM26	Writes to 10G PHY XGXS Lane status register have no effect	45.2.4.10		PX:M	Yes [ ] N/A [ ]
PM27	Writes to bit are ignored and reads return a value of zero	45.2.4.11.1		PX*!PT:M	Yes [ ] N/A [ ]
PM28	Setting the bits to <10> selects the mixed frequency pattern	45.2.4.11.2		PX*PT:M	Yes [ ] N/A [ ]
PM29	Setting the bits to <01> selects the low-frequency pattern	45.2.4.11.2		PX*PT:M	Yes [ ] N/A [ ]
PM30	Setting the bits to <00> selects the high-frequency pattern	45.2.4.11.2		PX*PT:M	Yes [ ] N/A [ ]

45.5.3.12 DTE XS options

Item	Feature	Subclause	Value/Comment	Status	Support
*DT	Implementation of pattern testing	45.2.5		DX:O	Yes [ ] No [ ] N/A [ ]

45.5.3.13 DTE XS management functions

Item	Feature	Subclause	Value/Comment	Status	Support
DM1	Device responds to all register addresses for that device	45.2		DX:M	Yes [ ] N/A [ ]
DM2	Writes to undefined and read-only registers have no effect	45.2		DX:M	Yes [ ] N/A [ ]
DM3	Operation is not affected by writes to reserved and unsupported bits	45.2		DX:M	Yes [ ] N/A [ ]
DM4	Reserved and unsupported bits return a value of zero	45.2		DX:M	Yes [ ] N/A [ ]
DM5	Latching low bits remain low until after they have been read via the management interface	45.2		DX:M	Yes [ ] N/A [ ]
DM6	Latching low bits assume correct value once read via the management interface	45.2	Correct value is based upon current state	DX:M	Yes [ ] N/A [ ]
DM7	Latching high bits remain high until after they have been read via the management interface	45.2		DX:M	Yes [ ] N/A [ ]
DM8	Latching high bits assume correct value once read via the management interface	45.2	Correct value is based upon current state	DX:M	Yes [ ] N/A [ ]
DM9	Action on reset	45.2.5.1.1	Reset the registers of the entire device to default values and set bit 15 of the Control register to one	DX:M	Yes [ ] N/A [ ]
DM10	Return 1 until reset completed	45.2.5.1.1		DX:M	Yes [ ] N/A [ ]
DM11	Reset completes within 0.5 s	45.2.5.1.1		DX:M	Yes [ ] N/A [ ]
DM12	Device responds to reads of bits 5.0.15 and 5.8.15:14 during reset	45.2.5.1.1		DX:M	Yes [ ] N/A [ ]
DM13	Loopback mode	45.2.5.1.2	Whenever bit 5.0.14 is set to a one	DX:M	Yes [ ] N/A [ ]
DM14	Transmit data is returned on receive path during loopback	45.2.5.1.2		DX:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
DM15	Device responds to transactions necessary to exit low-power mode while in low-power state	45.2.5.1.3		DX:M	Yes [ ] N/A [ ]
DM16	Speed selection bits 13 and 6 are written as one	45.2.5.1.6		DX:M	Yes [ ] N/A [ ]
DM17	Invalid writes to speed selection bits are ignored	45.2.5.1.6		DX:M	Yes [ ] N/A [ ]
DM18	Writes to status 1 register have no effect	45.2.5.2		DX:M	Yes [ ] N/A [ ]
DM19	Receive link status implemented using latching low behavior	45.2.5.2.7		DX:M	Yes [ ] N/A [ ]
DM20	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.5.3		DX:M	Yes [ ] N/A [ ]
DM21	Writes to status 2 register have no effect	45.2.5.6		DX:M	Yes [ ] N/A [ ]
DM22	Transmit fault implemented with latching high behavior	45.2.5.6.2		DX:M	Yes [ ] N/A [ ]
DM23	Receive fault implemented with latching high behavior	45.2.5.6.3		DX:M	Yes [ ] N/A [ ]
DM24	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.5.7		DX:M	Yes [ ] N/A [ ]
DM25	Writes to 10G DTE XGXS Lane status register have no effect	45.2.5.10		DX:M	Yes [ ] N/A [ ]
DM26	Writes to bit are ignored and reads return a value of zero	45.2.5.11.1		DX*!DT:M	Yes [ ] N/A [ ]
DM27	Setting the bits to <10> selects the mixed frequency pattern	45.2.5.11.2		DX*DT:M	Yes [ ] N/A [ ]
DM28	Setting the bits to <01> selects the low-frequency pattern	45.2.5.11.2		DX*DT:M	Yes [ ] N/A [ ]
DM29	Setting the bits to <00> selects the high-frequency pattern	45.2.5.11.2		DX*DT:M	Yes [ ] N/A [ ]

45.5.3.14 Vendor specific MMD 1 management functions

Item	Feature	Subclause	Value/Comment	Status	Support
VSA1	Device responds to all register addresses for that device	45.2		VSA:M	Yes [ ] N/A [ ]
VSA2	Writes to undefined and read-only registers have no effect	45.2		VSA:M	Yes [ ] N/A [ ]
VSA3	Operation is not affected by writes to reserved and unsupported bits	45.2		VSA:M	Yes [ ] N/A [ ]
VSA4	Reserved and unsupported bits return a value of zero	45.2		VSA:M	Yes [ ] N/A [ ]
VSA5	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.9.1		VSA:M	Yes [ ] N/A [ ]
VSA6	Writes to status register have no effect	45.2.9.2		VSA:M	Yes [ ] N/A [ ]
VSA7	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.9.3		VSA:M	Yes [ ] N/A [ ]

45.5.3.15 Vendor specific MMD 2 management functions

Item	Feature	Subclause	Value/Comment	Status	Support
VSB1	Device responds to all register addresses for that device	45.2		VSB:M	Yes [ ] N/A [ ]
VSB2	Writes to undefined and read-only registers have no effect	45.2		VSB:M	Yes [ ] N/A [ ]
VSB3	Operation is not affected by writes to reserved and unsupported bits	45.2		VSB:M	Yes [ ] N/A [ ]
VSB4	Reserved and unsupported bits return a value of zero	45.2		VSB:M	Yes [ ] N/A [ ]
VSB5	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.10.1		VSB:M	Yes [ ] N/A [ ]
VSB6	Writes to status register have no effect	45.2.10.2		VSB:M	Yes [ ] N/A [ ]
VSB7	Unique identifier is composed of 22 bits of OUI, model number and revision	45.2.10.3		VSB:M	Yes [ ] N/A [ ]

45.5.3.16 Management frame structure

Item	Feature	Subclause	Value/Comment	Status	Support
MF1	Device has implemented sixteen bit address register	45.3		M	Yes [ ]
MF2	Address register is overwritten by address frames	45.3		M	Yes [ ]
MF3	Write, read, and post-read-increment-address frames access the register whose address is held in the address register	45.3		M	Yes [ ]
MF4	Write and read frames do not modify the address register	45.3		M	Yes [ ]
MF5	Post-read-increment-address frames increment the address register by one unless the address register contains 65 535	45.3		M	Yes [ ]
MF6	Components containing several MMDs implement separate address registers	45.3		M	Yes [ ]
MF7	Tri state drivers are disabled during idle	45.3.1		M	Yes [ ]
MF8	STA sources 32 contiguous ones at the beginning of each transaction	45.3.2		M	Yes [ ]
MF9	MMD observes 32 contiguous ones at the beginning of each transaction	45.3.2		M	Yes [ ]
MF10	Frames containing ST=<01> sequence are ignored	45.3.3		M	Yes [ ]
MF11	STA tri state driver is high impedence during first bit of TA during read or post-read-increment-address frames	45.3.7		M	Yes [ ]
MF12	MMD tri state driver is high impedence during first bit of TA during read or post-read-increment-address frames	45.3.7		M	Yes [ ]
MF13	MMD tri state driver drives a zero bit during second bit of TA during read or post-read-increment-address frames	45.3.7		M	Yes [ ]
MF14	STA tri state driver drives a one bit followed by a zero bit for the TA during write or address frames	45.3.7		M	Yes [ ]
MF15	First bit transmitted and received is bit 15	45.3.8		M	Yes [ ]

45.5.3.17 TC management functions

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	Device responds to all register addresses for that device	45.2		TC:M	Yes [ ] N/A [ ]
TC2	Writes to undefined and read only register have no effect	45.2		TC:M	Yes [ ] N/A [ ]
TC3	Operation is not affected by writes to reserved and unsupported bits	45.2		TC:M	Yes [ ] N/A [ ]
TC4	Reserved and unsupported bits return a value of zero	45.2		TC:M	Yes [ ] N/A [ ]
TC5	Setting bit to a one sets all TC registers to their default states	45.2.6.1.1		TC:M	Yes [ ] N/A [ ]
TC6	Bit reads one while reset is in progress otherwise reads zero	45.2.6.1.1		TC:M	Yes [ ] N/A [ ]
TC7	Control and management interface is restored within 0.5s from setting bit to a one	45.2.6.1.1		TC:M	Yes [ ] N/A [ ]
TC8	During reset, TC responds to reads from bit	45.2.6.1.1		TC:M	Yes [ ] N/A [ ]
TC9	Writes that would select an unsupported speed are ignored	45.2.6.1.2		TC:M	Yes [ ] N/A [ ]
TC10	Identifier composed properly	45.2.6.2		TC:M	Yes [ ] N/A [ ]
TC11	Identifier composed properly	45.2.6.5		TC:M	Yes [ ] N/A [ ]
TC12	Register is unique across all PCS MMDs in a package	45.2.6.6		TC:M	Yes [ ] N/A [ ]
TC13	Operation ignored when link is up or initializing	45.2.6.6.1		TC:M	Yes [ ] N/A [ ]
TC14	Bits indicate “Ready” when PAF is capable	45.2.6.6.1		TC:M	Yes [ ] N/A [ ]
TC15	Writes ignored and “Ready” indicated if PAF is unsupported	45.2.6.6.1		TC:M	Yes [ ] N/A [ ]
TC16	Bits indicate “Ready” when operation is complete or upon reset	45.2.6.6.1		TC:M	Yes [ ] N/A [ ]
TC17	Bits read as zero if PAF is unsupported	45.2.6.7.1		TC:M	Yes [ ] N/A [ ]
TC18	Bits read as zero if PAF is unsupported	45.2.6.7.2		TC:M	Yes [ ] N/A [ ]
TC19	Register is unique across all PCS MMDs in a package	45.2.6.8		TC:M	Yes [ ] N/A [ ]
TC20	Register is unique across all PCS MMDs in a package	45.2.6.9		TC:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
TC21	Operation ignored when link is up or initializing	45.2.6.9.1		TC:M	Yes [ ] N/A [ ]
TC22	Bits indicate “Ready” when PAF is capable	45.2.6.9.1		TC:M	Yes [ ] N/A [ ]
TC23	Writes ignored and “Ready” indicated if PAF is unsupported	45.2.6.9.1		TC:M	Yes [ ] N/A [ ]
TC24	Bits indicate “Ready” when operation is complete or on reset	45.2.6.9.1		TC:M	Yes [ ] N/A [ ]
TC25	Register is unique across all PCS MMDs in a package	45.2.6.10		TC:M	Yes [ ] N/A [ ]
TC26	Bits reset to zero when read or reset	45.2.6.11		TC:M	Yes [ ] N/A [ ]
TC27	Bits held to one upon overflow	45.2.6.11		TC:M	Yes [ ] N/A [ ]
TC28	Bits reset to zero when read or reset	45.2.6.12		TC:M	Yes [ ] N/A [ ]
TC29	Bits held to one upon overflow	45.2.6.12		TC:M	Yes [ ] N/A [ ]

45.5.3.18 Clause 22 extension options

Item	Feature	Subclause	Value/Comment	Status	Support
*FEC	Implementation of PHY FEC	65.4.4.6		CTT:O	Yes [ ] No [ ] N/A [ ]

45.5.3.19 Clause 22 extension management functions

Item	Feature	Subclause	Value/Comment	Status	Support
CT1	Device responds to all register addresses for that device	45.2		CTT:M	Yes [ ] N/A [ ]
CT2	Writes to undefined and read only register have no effect	45.2		CTT:M	Yes [ ] N/A [ ]
CT3	Operation is not affected by writes to reserved and unsupported bits	45.2		CTT:M	Yes [ ] N/A [ ]
CT4	Reserved and unsupported bits return a value of zero	45.2		CTT:M	Yes [ ] N/A [ ]
CT5	Bits set to zero upon PHY reset	45.2.8.3.1		CTT*FEC:M	Yes [ ] N/A [ ]
CT6	Bits reset to all zeros when the register is read by the management function or upon PHY reset.	45.2.8.4		CTT*FEC:M	Yes [ ] N/A [ ]
CT7	Bits held at all ones in the case of overflow	45.2.8.4		CTT*FEC:M	Yes [ ] N/A [ ]
CT8	Bits reset to all zeros when the register is read by the management function or upon PHY reset.	45.2.8.5		CTT*FEC:M	Yes [ ] N/A [ ]
CT9	Bits held at all ones in the case of overflow	45.2.8.5		CTT*FEC:M	Yes [ ] N/A [ ]
CT10	Bits reset to all zeros when the register is read by the management function or upon PHY reset.	45.2.8.6		CTT*FEC:M	Yes [ ] N/A [ ]
CT11	Bits held at all ones in the case of overflow	45.2.8.6		CTT*FEC:M	Yes [ ] N/A [ ]

45.5.3.20 Signal timing characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
ST1	MDIO setup and hold time	45.4.2	Setup min = 10 ns; Hold min = 10 ns per Figure 45–3	M	Yes [ ]
ST2	MDIO clock to output delay	45.4.2	Min = 0 ns; Max = 300 ns per Figure 45–4	M	Yes [ ]
ST3	MDC min high/low time	45.4.2	160 ns	M	Yes [ ]
ST4	MDC min period	45.4.2	400 ns	M	Yes [ ]

45.5.3.21 Electrical characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
EC1	V_{OH}	45.4.1	≥ 1.0 V (I_{OH} = –100 uA) ≤ 1.5 V (I_{OH} = –100 uA)	M	Yes [ ]
EC2	V_{OL}	45.4.1	≥ –0.3 V (I_{OL} = 100 uA) ≤ 0.2 V (I_{OL} = 100 uA)	M	Yes [ ]
EC3	V_{IH}	45.4.1	0.84 V ≤ V_{IH} ≤ 1.5 V	M	Yes [ ]
EC4	V_{IL}	45.4.1	–0.3 ≤ V_{IL} ≤ 0.36 V	M	Yes [ ]
EC5	Input capacitance for MDIO	45.4.1	≤ 10 pF	M	Yes [ ]
EC6	Total capacitive load	45.4.1	≤ 470 pF	M	Yes [ ]
EC7	I_{OH}	45.4.1	≤ –4 mA at V_{OH} = 1.0 V	!ODB:M	Yes [ ] N/A[ ]
EC8	I_{OL}	45.4.1	≥ +4 mA at V_{OL} = 0.2 V	M	Yes [ ]

46. Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII)

46.1 Overview

This clause defines the logical and electrical characteristics for the Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII) between Ethernet media access controllers and various PHYs. Figure 46–1 shows the relationship of the RS and XGMII to the ISO/IEC (IEEE) OSI reference model.

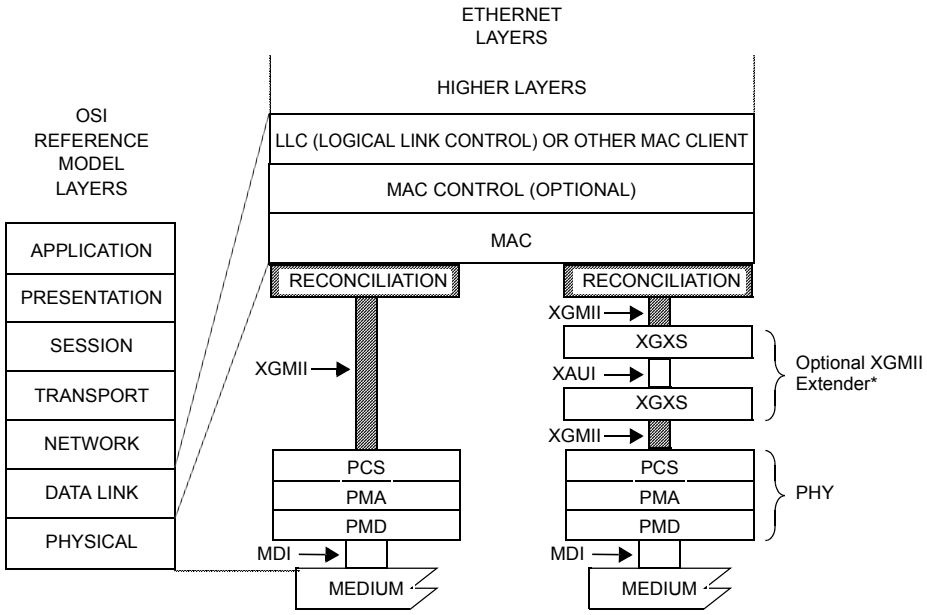

MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XAUI = 10 GIGABIT ATTACHMENT UNIT INTERFACE
XGMII = 10 GIGABIT MEDIA INDEPENDENT INTERFACE
XGXS = XGMII EXTENDER SUBLAYER

*specified in Clause 47

Figure 46–1—XGMII relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

The purpose of the XGMII is to provide a simple, inexpensive, and easy-to-implement interconnection between the Media Access Control (MAC) sublayer and the Physical Layer (PHY). The 10 Gigabit Attachment Unit Interface (XAUI) may optionally be used to extend the operational distance of the XGMII with reduced pin count (see Clause 47).

The RS adapts the bit serial protocols of the MAC to the parallel encodings of 10 Gb/s PHYs. Though the XGMII is an optional interface, it is used extensively in this standard as a basis for specification. The 10 Gb/s Physical Coding Sublayer (PCS) is specified to the XGMII, so if not implemented, a conforming implementation shall behave functionally as if the RS and XGMII were implemented.

The XGMII has the following characteristics:

a) It is capable of supporting 10 Gb/s operation.
b) Data and delimiters are synchronous to clock reference.

c) It provides independent 32-bit-wide transmit and receive data paths.
d) It uses signal levels compatible with common digital ASIC processes.
e) It provides for full duplex operation only.

46.1.1 Summary of major concepts

The following are the major concepts of XGMII:

a) The XGMII is functionally similar to the MII defined in Clause 22 and GMII defined in Clause 35 as they all define an interface allowing independent development of MAC and PHY logic.
b) The RS converts between the MAC serial data stream and the parallel data paths of the XGMII.
c) The RS maps the signal set provided at the XGMII to the PLS service primitives provided at the MAC.
d) Each direction of data transfer is independent and serviced by data, control, and clock signals.
e) The RS generates continuous data or control characters on the transmit path and expects continuous data or control characters on the receive path.
f) The RS participates in link fault detection and reporting by monitoring the receive path for status reports that indicate an unreliable link, and generating status reports on the transmit path to report detected link faults to the DTE on the remote end of the connecting link.
g) When the XGMII is optionally extended with XAUI, two XGMII interfaces logically exist (see Figure 46–1). The transmit path signals are from the RS to the DTE (top) XGXS of the XAUI via one XGMII and from the PHY (bottom) XGXS to the PCS via the other XGMII. The receive path signals are from the PCS to the PHY XGXS of the XAUI via one XGMII and from the DTE XGXS to the RS via the other XGMII. The descriptions of the XGMII as between the RS and the PCS are, therefore, equally applicable between the RS and the DTE XGXS or the PHY XGXS and the PCS.
h) The XGMII may also support Low Power Idle (LPI) signaling for PHY types supporting Energy-Efficient Ethernet (EEE) (see Clause 78).

46.1.2 Application

This clause applies to the interface between the MAC and PHY. The physical implementation of the interface is primarily intended as a chip-to-chip (integrated circuit to integrated circuit) interface implemented with traces on a printed circuit board. The XGMII may also be used in other ways, for example, as a logical interface between ASIC logic modules within an integrated circuit.

This interface is used to provide media independence so that an identical media access controller may be used with all 10GBASE PHY types.

46.1.3 Rate of operation

The XGMII supports only the 10 Gb/s MAC data rate as defined within this clause. Operation at 10 Mb/s and 100 Mb/s is supported by the MII defined in Clause 22 and operation at 1000 Mb/s by the GMII defined in Clause 35.

PHYs that provide an XGMII shall support the 10 Gb/s MAC data rate. 10GBASE-X and 10GBASE-R PHYs operate at a 10 Gb/s data rate. 10GBASE-W PHYs operate at the STS-192/VC-4-64c line rate of 9.95328 Gb/s, mapping the encoded data stream at a 9.58464 Gb/s payload rate. On transmit, this mapping is performed by discarding Idle control characters corresponding to the stretched interpacket gap created by the MAC in this mode of operation, and on receive, by adding interpacket gap Idle control characters as required to adapt to the XGMII RX_CLK rate.

46.1.4 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementers must conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices. The maximum cumulative MAC Control, MAC and RS round-trip (sum of transmit and receive) delay shall meet the values specified in Table 46–1. Bit time is specified in 1.4, and pause_quanta is specified in 31B.2.

Table 46–1—Round-trip delay constraints

Sublayer	Maximum (bit time)	Maximum (pause_quanta)
MAC, RS, and MAC Control	8192	16

46.1.5 Allocation of functions

The allocation of functions at the XGMII balances the need for media independence with the need for a simple and cost-effective interface. The bus width and signaling rate are applicable to short distance chip-to-chip interconnect with printed circuit board trace lengths electrically limited to approximately 7 cm. The XGMII (like the MII and GMII) maximizes media independence by cleanly separating the Data Link and Physical Layers of the ISO (IEEE) seven-layer reference model. This allocation also recognizes that implementations can benefit from a close coupling between the PLS sublayer or PCS and the PMA sublayer.

46.1.6 XGMII structure

The XGMII is composed of independent transmit and receive paths. Each direction uses 32 data signals (TXD<31:0> and RXD<31:0>), four control signals (TXC<3:0> and RXC<3:0>), and a clock (TX_CLK and RX_CLK). Figure 46–2 depicts a schematic view of the RS inputs and outputs.

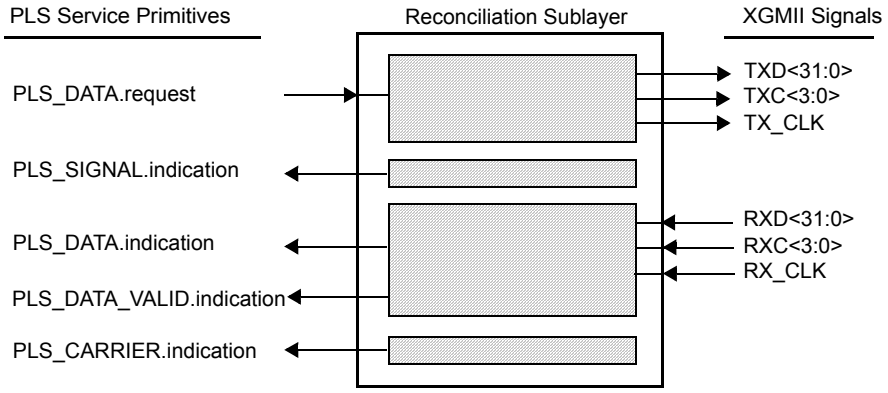

Figure 46–2—Reconciliation Sublayer (RS) inputs and outputs

The 32 TXD and four TXC signals shall be organized into four data lanes, as shall the 32 RXD and four RXC signals (see Table 46–2). The four lanes in each direction share a common clock—TX_CLK for transmit and RX_CLK for receive. The four lanes are used in round-robin sequence to carry an octet stream. On transmit, each eight PLS_DATA.request transactions represent an octet transmitted by the MAC. The

first octet is aligned to lane 0, the second to lane 1, the third to lane 2 the fourth to lane 3, then repeating with the fifth to lane 0, etc. Delimiters and interframe idle characters are encoded on the TXD and RXD signals with the control code indicated by assertion of TXC and RXC, respectively.

Table 46–2—Transmit and receive lane associations

TXD RXD	TXC RXC	Lane
<7:0>	<0>	0
<15:8>	<1>	1
<23:16>	<2>	2
<31:24>	<3>	3

46.1.7 Mapping of XGMII signals to PLS service primitives

The Reconciliation Sublayer (RS) shall map the signals provided at the XGMII to the PLS service primitives defined in Clause 6. The PLS service primitives provided by the RS and described here behave in exactly the same manner as defined in Clause 6. Full duplex operation only is implemented at 10 Gb/s; therefore, PLS service primitives supporting CSMA/CD operation are not mapped through the RS to the XGMII. The mapping is changed if EEE capability is supported (see 78.3). This behavior and restrictions are the same as described in 22.7, with the details of the signaling described in 46.3. LPI_REQUEST shall not be set to ASSERT unless the attached link has been operational for at least one second (i.e., link_status = OK, according to the underlying PCS/PMA).

EEE capability requires the use of the MAC defined in Annex 4A for simplified full duplex operation (with carrier sense deferral). This provides full duplex operation but uses the carrier sense signal to defer transmission when the PHY is in its low power state.

Mappings for the following primitives are defined for 10 Gb/s operation:

PLS_DATA.request
PLS_DATA.indication
PLS_CARRIER.indication
PLS_SIGNAL.indication
PLS_DATA_VALID.indication

46.1.7.1 Mapping of PLS_DATA.request

46.1.7.1.1 Function

Map the primitive PLS_DATA.request to the XGMII signals TXD<31:0>, TXC<3:0>, and TX_CLK.

46.1.7.1.2 Semantics of the service primitive

PLS_DATA.request (OUTPUT_UNIT)

The OUTPUT_UNIT parameter can take one of three values: ONE, ZERO, or DATA_COMPLETE. It represents a single data bit. The DATA_COMPLETE value signifies that the Media Access Control sublayer has no more data to output.

46.1.7.1.3 When generated

This primitive is generated by the MAC sublayer to request the transmission of a single data bit on the physical medium or to stop transmission.

46.1.7.1.4 Effect of receipt

The OUTPUT_UNIT values are conveyed to the PHY by the signals TXD<31:0> and TXC<3:0> on each TX_CLK edge. Each PLS_DATA.request transaction shall be mapped to a TXD signal in sequence (TXD<0>, TXD<1>,... TXD<31>, TXD<0>) as described in 46.2. After 32 PLS_DATA.request transactions from the MAC sublayer (four octets of eight PLS_DATA.request transactions each), the RS requests transmission of 32 data bits by the PHY. The first octet of preamble shall be converted to a Start control character and aligned to lane 0. The TXD<31:0> and TXC<3:0> shall be generated by the RS for each 32 bit-times of the MAC sublayer.

The DATA_COMPLETE value shall be mapped to a Terminate control character encoded on the next eight TXD signals in sequence after the last data octet; and is transferred to the PHY at the next TX_CLK edge. This may be on the same TX_CLK edge as the last data octet or the subsequent TX_CLK edge. When the Terminate control character is in lane 0, 1, or 2, the lanes following in sequence are encoded with an Idle control character.

46.1.7.2 Mapping of PLS_DATA.indication

46.1.7.2.1 Function

Map the primitive PLS_DATA.indication to the XGMII signals RXD<31:0>, RXC<3:0> and RX_CLK.

46.1.7.2.2 Semantics of the service primitive

PLS_DATA.indication (INPUT_UNIT)

The INPUT_UNIT parameter can take one of two values: ONE or ZERO. It represents a single data bit.

46.1.7.2.3 When generated

The INPUT_UNIT values are derived from the signals RXC<3:0> and RXD<31:0> received from the PHY on each edge of RX_CLK. Each primitive generated to the MAC sublayer entity corresponds to a PLS_DATA.request issued by the MAC at the remote end of the link connecting two DTEs. For each RXD<31:0> during frame reception, the RS shall generate 32 PLS_DATA.indication transactions until the end of frame (Terminate control character), where 0, 8, 16, or 24 PLS_DATA.indication transactions will be generated from the RXD<31:0> containing the Terminate. During frame reception, each RXD signal shall be mapped in sequence into a PLS_DATA.indication transaction (RXD<0>, RXD<1>,... RXD<31>, RXD<0>) as described in 46.2.

The RS shall convert a valid Start control character to a preamble octet prior to generation of the associated PLS_DATA.indication transactions. The RS shall not generate any PLS_DATA.indication primitives for a Terminate control character. To assure robust operation, the value of the data transferred to the MAC may be changed by the RS as required by XGMII error indications (see 46.3.3). Sequence ordered sets are not indicated to the MAC (see 46.3.4).

46.1.7.2.4 Effect of receipt

The effect of receipt of this primitive by the MAC sublayer is unspecified.

46.1.7.3 Mapping of PLS_CARRIER.indication

10 Gb/s operation supports full duplex operation only. The RS never generates this primitive for PHYs that do not support EEE or Link Interruption.

For PHYs that support EEE capability, CARRIER_STATUS is set in response to LPI_REQUEST as shown in Figure 46–13. For PHYs that support Link Interruption, CARRIER_STATUS may be set in response to link_fault. CARRIER_STATUS is set to CARRIER_ON if LPI_CARRIER_STATUS is TRUE or if link_fault is Link Interruption. CARRIER_STATUS is otherwise set to CARRIER_OFF. The deferral mechanism based upon the Link Interruption signal may be enabled or disabled by management.

46.1.7.4 Mapping of PLS_SIGNAL.indication

10 Gb/s operation supports full duplex operation only. The RS never generates this primitive.

46.1.7.5 Mapping of PLS_DATA_VALID.indication

46.1.7.5.1 Function

Map the primitive PLS_DATA_VALID.indication to the XGMII signals RXC<3:0> and RXD<31:0>.

46.1.7.5.2 Semantics of the service primitive

PLS_DATA_VALID.indication (DATA_VALID_STATUS)

The DATA_VALID_STATUS parameter can take one of two values: DATA_VALID or DATA_NOT_VALID. The DATA_VALID value indicates that the INPUT_UNIT parameter of the PLS_DATA.indication primitive contains a valid data of an incoming frame. The DATA_NOT_VALID value indicates that the INPUT_UNIT parameter of the PLS_DATA.indication primitive does not contain valid data of an incoming frame.

46.1.7.5.3 When generated

The PLS_DATA_VALID.indication service primitive shall be generated by the RS whenever the DATA_VALID_STATUS parameter changes from DATA_VALID to DATA_NOT_VALID or vice versa.

DATA_VALID_STATUS shall assume the value DATA_VALID when a PLS_DATA.indication transaction is generated in response to reception of a Start control character on lane 0 if the prior RXC<3:0> and RXD<31:0> contained four Idle characters or a Sequence ordered set. DATA_VALID_STATUS shall assume the value DATA_NOT_VALID when RXC of the current lane in sequence is asserted for anything except an Error control character. In the absence of errors, DATA_NOT_VALID is caused by a Terminate control character. When DATA_VALID_STATUS changes from DATA_VALID to DATA_NOT_VALID because of a control character other than Terminate, the RS shall ensure that the MAC will detect a FrameCheckError prior to indicating DATA_NOT_VALID to the MAC (see 46.3.3.1).

46.1.7.5.4 Effect of receipt

The effect of receipt of this primitive by the MAC sublayer is unspecified.

46.2 XGMII data stream

Packets transmitted through the XGMII shall be transferred within the XGMII data stream. The data stream is a sequence of bytes, where each byte conveys either a data octet or control character. The parts of the data stream are shown in Figure 46–3.

<inter-frame><preamble><sfd><data><efd>

Figure 46–3—XGMII data stream

For the XGMII, transmission and reception of each bit and mapping of data octets to lanes shall be as shown in Figure 46–4.

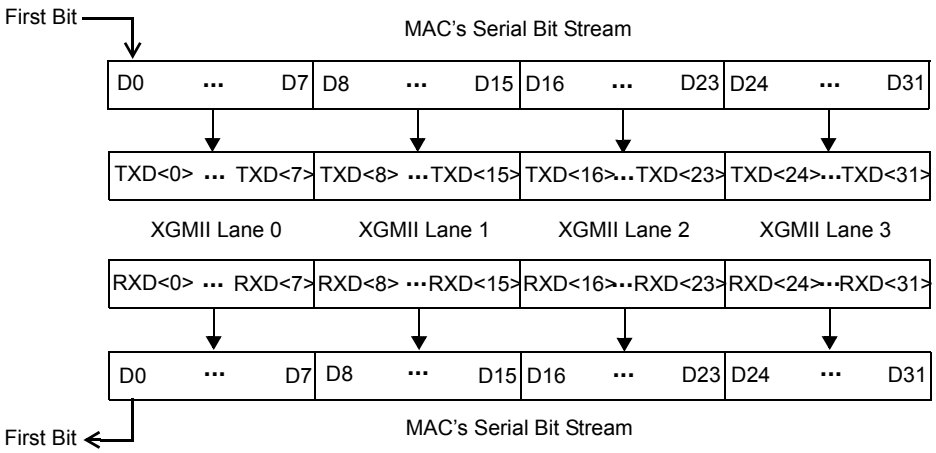

Figure 46–4—Relationship of data lanes to MAC serial bit stream

46.2.1 Inter-frame <inter-frame>

The inter-frame <inter-frame> period on an XGMII transmit or receive path is an interval during which no frame data activity occurs. The <inter-frame> corresponding to the MAC interpacket gap begins with the Terminate control character, continues with Idle control characters and ends with the Idle control character prior to a Start control character. The length of the interpacket gap may be changed between the transmitting MAC and receiving MAC by one or more functions (e.g., RS lane alignment, PHY clock rate compensation, or 10GBASE-W data rate adaptation functions). The minimum IPG at the XGMII of the receiving RS is five octets.

The signaling of link status information logically occurs in the <inter-frame> period (see 46.3.4). Subclause 46.3.3 describes frame processing when signaling of link status information is initiated or terminated.

46.2.2 Preamble <preamble> and start of frame delimiter <sfd>

The preamble <preamble> begins a frame transmission by a MAC as specified in 4.2.5 and when generated by a MAC consists of 7 octets with the following bit values:

10101010 10101010 10101010 10101010 10101010 10101010 10101010

The Start control character indicates the beginning of MAC data on the XGMII. On transmit, the RS converts the first data octet of preamble transferred from the MAC into a Start control character. On receive, the RS will convert the Start control character into a preamble data octet. The start control character is aligned to lane 0 of the XGMII by the RS on transmit and by the PHY on receive.

The start of frame delimiter <sfd> indicates the start of a frame and immediately follows the preamble. The bit value of <sfd> at the XGMII is unchanged from the Start Frame Delimiter (SFD) specified in 4.2.6 and is the bit sequence:

10101011

The preamble and SFD are shown previously with their bits ordered for serial transmission from left to right. As shown, the left-most bit of each octet is the LSB of the octet and the right-most bit of each octet is the MSB of the octet.

The preamble and SFD are transmitted through the XGMII as octets sequentially ordered on the lanes of the XGMII. The first preamble octet is replaced with a Start control character and it is aligned to lane 0, the second octet on lane 1, the third on lane 2 and the fourth on lane 3, and the four octets are transferred on the next edge of TX_CLK. The fifth octet is assigned to lane 0 with subsequent octets sequentially assigned to the lanes with the SFD assigned to lane 3. The XGMII <preamble> and <sfd> are as follows:

Lane 0	**Lane 1**	**Lane 2**	**Lane 3**
Start	10101010	10101010	10101010
10101010	10101010	10101010	10101011

46.2.3 Data <data>

The data <data> in a well-formed frame shall consist of a set of data octets.

46.2.4 End of frame delimiter <efd>

Assertion of TXC with the appropriate Terminate control character encoding of TXD on a lane constitutes an end of frame delimiter <efd> for the transmit data stream. Similarly, assertion of RXC with the appropriate Terminate control character encoding of RXD constitutes an end of frame delimiter for the receive data stream. The XGMII shall recognize the end of frame delimiter on any of the four lanes of the XGMII.

46.2.5 Definition of Start of Packet and End of Packet Delimiters

For the purposes of Clause 30, the Start of Packet delimiter is defined as the Start control character, and the End of Packet delimiter is defined as the end of the last sequential data octet preceding the Terminate control character or other control character causing a change from DATA_VALID to DATA_NOT_VALID. (See 46.1.7.5.2 and 30.3.2.1.5.)

46.3 XGMII functional specifications

The XGMII is designed to make the differences among the various media and transceiver combinations transparent to the MAC sublayer. The selection of logical control signals and the functional procedures are all designed to this end.

NOTE—No XGMII loopback is defined, but XGMII signals are specified such that transmit signals may be connected to receive signals to create a loopback path. To do this, TXD<0> is connected to RXD<0> ... TXD<31> to RXD<31>, TXC<0> to RXC<0> ... TXC<3> to RXC<3>, and TXCLK to RXCLK. Such a loopback does not test the Link Fault Signaling state diagram, nor any of the error handling functions of the receive RS.

46.3.1 Transmit

46.3.1.1 TX_CLK (10 Gb/s transmit clock)

TX_CLK is a continuous clock used for operation at 10 Gb/s. TX_CLK provides the timing reference for the transfer of the TXC<3:0> and TXD<31:0> signals from the RS to the PHY. The values of TXC<3:0> and TXD<31:0> shall be sampled by the PHY on both the rising edge and falling edge of TX_CLK. TX_CLK is sourced by the RS.

The TX_CLK frequency shall be 156.25 MHz ±0.01%, one-sixty-fourth of the MAC transmit data rate.

NOTE—For EEE capability, TX_CLK may be halted according to 46.3.1.5.

46.3.1.2 TXC<3:0> (transmit control)

TXC<3:0> indicate that the RS is presenting either data or control characters on the XGMII for transmission. The TXC signal for a lane shall be de-asserted when a data octet is being sent on the corresponding lane and asserted when a control character is being sent. In the absence of errors, the TXC signals are de-asserted by the RS for each octet of the preamble (except the first octet that is replaced with a Start control character) and remain de-asserted while all octets to be transmitted are presented on the lanes of the XGMII. TXC<3:0> are driven by the RS and shall transition synchronously with respect to both the rising and falling edges of TX_CLK. Table 46–3 specifies the permissible encodings of TXD and TXC for a XGMII transmit lane. Additional requirements apply for proper code sequences and in which lanes particular codes are valid (e.g., Start control character is to be aligned to lane 0).

A PHY with EEE capability shall interpret the combination of TXC and TXD as shown in Table 46–3 as an assertion of LPI. Transition into and out of the LPI state is shown in Figure 46–7.

46.3.1.3 TXD<31:0> (transmit data)

TXD is a bundle of 32 data signals organized into four lanes of eight signals each (TXD<7:0>, TXD<15:8>, TXD<23:16>, and TXD<31:24>) that are driven by the RS. Each lane is associated with a TXC signal as shown in Table 46–2 and shall be encoded as shown in Table 46–3. TXD<31:0> shall transition synchronously with respect to both the rising and falling edges of TX_CLK. For each high or low TX_CLK transition, data and/or control are presented on TXD<31:0> to the PHY for transmission. TXD<0> is the least significant bit of lane 0, TXD<8> the least significant bit of lane 1, TXD<16> the least significant bit of lane 2, and TXD<24> the least significant bit of lane 3.

Assertion on a lane of appropriate TXD values when TXC is asserted will cause the PHY to generate code-groups associated with either Idle, Start, Terminate, Sequence, or Error control characters. While the TXC of a lane is de-asserted, TXD of the lane is used to request the PHY to generate code-groups corresponding to the data octet value of TXD. An example of normal frame transmission is illustrated in Figure 46–5.

Table 46–3—Permissible encodings of TXC and TXD

TXC	TXD	Description	PLS_DATA.request parameter
0	00 through FF	Normal data transmission	ZERO, ONE (eight bits)
1	00 through 05	Reserved	—
1	06	Only valid on all four lanes simultaneously to request LPI	No applicable parameter (normal interframe)
1	07	Idle	No applicable parameter (Normal inter-frame)
1	08 through 9B	Reserved	—
1	9C	Sequence (only valid in lane 0)	No applicable parameter (Inter-frame status signal)
1	9D through FA	Reserved	—
1	FB	Start (only valid in lane 0)	No applicable parameter, replaces first eight ZERO, ONE of a frame (preamble octet)
1	FC	Reserved	—
1	FD	Terminate	DATA_COMPLETE
1	FE	Transmit error propagation	No applicable parameter
1	FF	Reserved	—
NOTE—Values in TXD column are in hexadecimal, most significant bit to least significant bit (i.e., <7:0>).			

Figure 46–6 shows the behavior of TXD and TXC during an example transmission of a frame propagating an error.

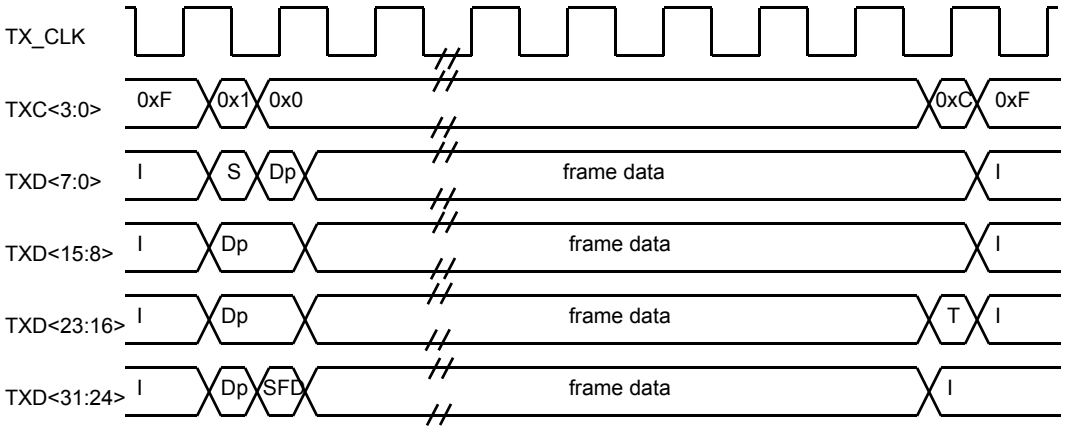

I: Idle control character, S: Start control character, Dp: preamble Data octet, T: Terminate control character

Figure 46–5—Normal frame transmission

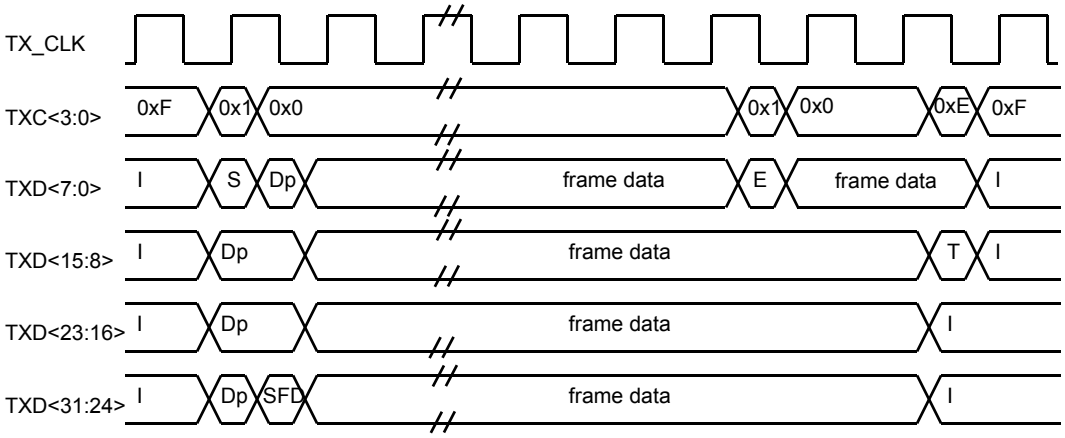

I: Idle control character, S: Start control character, Dp: preamble Data octet, T: Terminate control character, E: Error control character

Figure 46–6—Transmit Error Propagation

46.3.1.4 Start control character alignment

On transmit, it may be necessary for the RS to modify the length of the <inter-frame> in order to align the Start control character (first octet of preamble) on lane 0. This shall be accomplished in one of the following two ways:

1) A MAC implementation may incorporate this RS function into its design and always insert additional idle characters to align the start of preamble on a four byte boundary. Note that this will reduce the effective data rate for certain packet sizes separated with minimum inter-frame spacing.

2) Alternatively, the RS may maintain the effective data rate by sometimes inserting and sometimes deleting idle characters to align the Start control character. When using this method the RS must maintain a Deficit Idle Count (DIC) that represents the cumulative count of idle characters deleted or inserted. The DIC is incremented for each idle character deleted, decremented for each idle character inserted, and the decision of whether to insert or delete idle characters is constrained by bounding the DIC to a minimum value of zero and maximum value of three. Note that this may result in inter-frame spacing observed on the transmit XGMII that is up to three octets shorter than the minimum transmitted inter-frame spacing specified in Clause 4; however, the frequency of shortened inter-frame spacing is constrained by the DIC rules. The DIC is only reset at initialization and is applied regardless of the size of the IPG transmitted by the MAC sublayer. An equivalent technique may be employed to control RS alignment of the Start control character provided that the result is the same as if the RS implemented DIC as described.

46.3.1.5 Transmit direction LPI transition

LPI operation and the LPI client are described in 78.1. The RS requests the PHY to transition to the LPI state by asserting TXC and setting TXD to 0x06 (in all lanes). The RS maintains the same state for these signals for the entire time that the PHY is to remain in the LPI state.

The RS may halt TX_CLK at any time more than 128 clock cycles after the start of the LPI state as shown in Figure 46–7 if the clock stop capable bit of the attached sublayer is asserted (see 45.2.3.2.1 and 45.2.5.2.1). It is the responsibility of the management entity to ensure that the RS does not halt the TX_CLK if the attached device does not have its stop clock capable bit set. The RS shall restart TX_CLK so that at least one positive transition occurs before it deaserts LPI.

The RS asserts TXC and asserts IDLE on lanes 0–3 in order to make the PHY transition out of the LPI state. The RS should not present a start code for valid transmit data until after the wake-up time specified for the PHY.

Figure 46–7 shows the behavior of TXC and TXD<7:0> during the transition into and out of the LPI state.

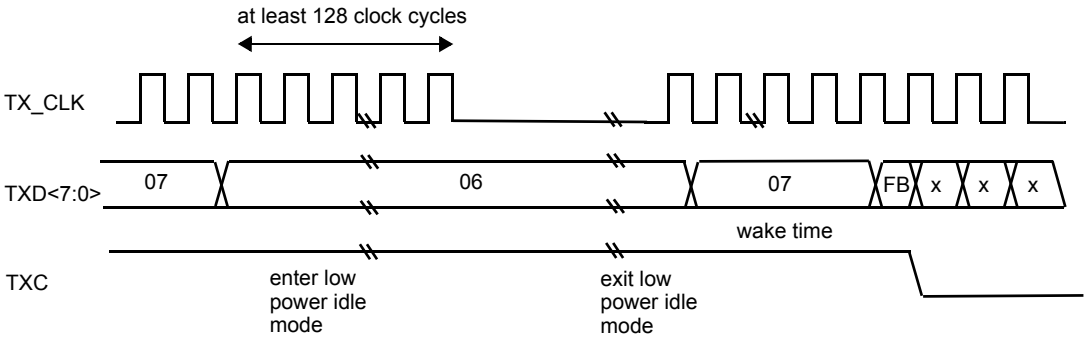

NOTE—TXC and TXD are shown for one lane, all four lanes behave identically during LPI.

Figure 46–7—LPI transition

Table 46–3 summarizes the permissible encodings of TXD<31:0>, TXC<3:0>.

46.3.2 Receive

46.3.2.1 RX_CLK (receive clock)

RX_CLK is a continuous clock that provides the timing reference for the transfer of the RXC<3:0> and RXD<31:0> signals from the PHY to the RS. RXC<3:0> and RXD<31:0> shall be sampled by the RS on both the rising and falling edge of RX_CLK. RX_CLK is sourced by the PHY.

The frequency of RX_CLK may be derived from the received data or it may be that of a nominal clock (e.g., TX_CLK). When the received data rate at the PHY is within tolerance, the RX_CLK frequency shall be 156.25 MHz ±0.01%, one-sixty-fourth of the MAC receive data rate.

There is no need to transition between the recovered clock reference and a nominal clock reference on a frame-by-frame basis. If loss of received signal from the medium causes a PHY to lose the recovered RX_CLK reference, the PHY shall source the RX_CLK from a nominal clock reference. Transitions from nominal clock to recovered clock or from recovered clock to nominal clock shall not decrease the time between adjacent edges of RX_CLK.

NOTE—This standard neither requires nor assumes a guaranteed phase relationship between the RX_CLK and TX_CLK signals. For EEE capability, RX_CLK may be halted according to 46.3.2.4.

46.3.2.2 RXC<3:0> (receive control)

RXC<3:0> indicate that the PHY is presenting either recovered and decoded data or control characters on the XGMII. The RXC signal for a lane shall be de-asserted when a data octet is being received on the corresponding lane and asserted when a control character is being received. In the absence of errors, the RXC signals are de-asserted by the PHY for each octet of the preamble (except the first octet that is replaced with a Start control character) and remain de-asserted while all octets to be received are presented on the lanes of the XGMII. RXC<3:0> are driven by the PHY and shall transition synchronously with respect to both the rising and falling edges of RX_CLK. Table 46–4 specifies the permissible encodings of RXD and RXC for a XGMII receive lane. Additional requirements apply for proper code sequences and in which lanes particular codes are valid (e.g., Start control character is to be aligned to lane 0).

Figure 46–8 shows the behavior of RXC<3:0> during an example frame reception with no errors.

Table 46–4—Permissible lane encodings of RXD and RXC

RXC	RXD	Description	PLS_DATA.indication parameter
0	00 through FF	Normal data reception	ZERO, ONE (eight bits)
1	00 through 05	Reserved	—
1	06	Only valid on all four lanes simultaneously to indicate LP_IDLE is asserted	No applicable parameter (Normal interframe)
1	07	Idle	No applicable parameter (Normal inter-frame)
1	08 through 9B	Reserved	—
1	9C	Sequence (only valid in lane 0)	No applicable parameter (Inter-frame status signal)
1	9D through FA	Reserved	—
1	FB	Start (only valid in lane 0)	No applicable parameter, first eight ZERO, ONE of a frame (a preamble octet)
1	FC	Reserved	—
1	FD	Terminate	No applicable parameter (Start of inter-frame)
1	FE	Receive error	No applicable parameter
1	FF	Reserved	—

NOTE—Values in RXD column are in hexadecimal, most significant bit to least significant bit (i.e., <7:0>).

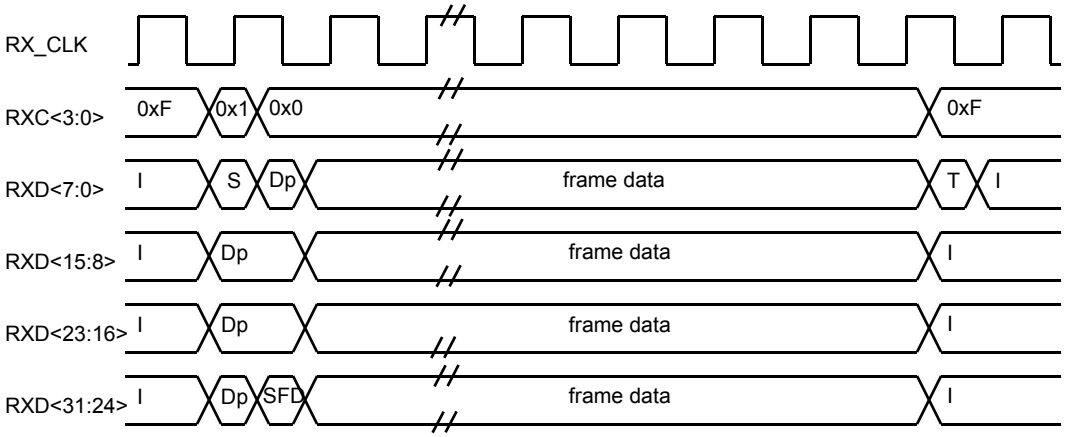

I: Idle control character, S: Start control character, Dp: preamble Data octet, T: Terminate control character

Figure 46–8—Basic frame reception

46.3.2.3 RXD (receive data)

RXD is a bundle of 32 data signals (RXD<31:0>) organized into four lanes of eight signals each (RXD<7:0>, RXD<15:8>, RXD<23:16>, and RXD<31:24>) that are driven by the PHY. Each lane is associated with a RXC signal as shown in Table 46–2 and shall be decoded by the RS as shown in Table 46–4. RXD<31:0> shall transition synchronously with respect to both the rising and falling edges of RX_CLK. For each high or low RX_CLK transition, received data and/or control are presented on RXD<31:0> for mapping by the RS. RXD<0> is the least significant bit of lane 0, RXD<8> the least significant bit of lane 1, RXD<16> the least significant bit of lane 2, and RXD<24> the least significant bit of lane 3. Figure 46–8 shows the behavior of RXD<31:0> during frame reception.

While the RXC of a lane is de-asserted, RXD of the lane is used by the RS to generate PLS_DATA.indication transactions. Assertion on a lane of appropriate RXD values when RXC is asserted indicates to the RS the Start control character, Terminate control character, Sequence control character, or Error control character that drive its mapping functions.

RXC of a lane is asserted with the appropriate Error control character encoding on RXD of the lane to indicate an error was detected somewhere in the frame presently being transferred from the PHY to the RS (e.g., a coding error, or any error that the PHY is capable of detecting, and that may otherwise be undetectable at the MAC sublayer).

The effect of an Error control character on the RS is defined in 46.3.3.1. Figure 46–9 shows the behavior of RXC and RXD during the reception of an example frame with an error.

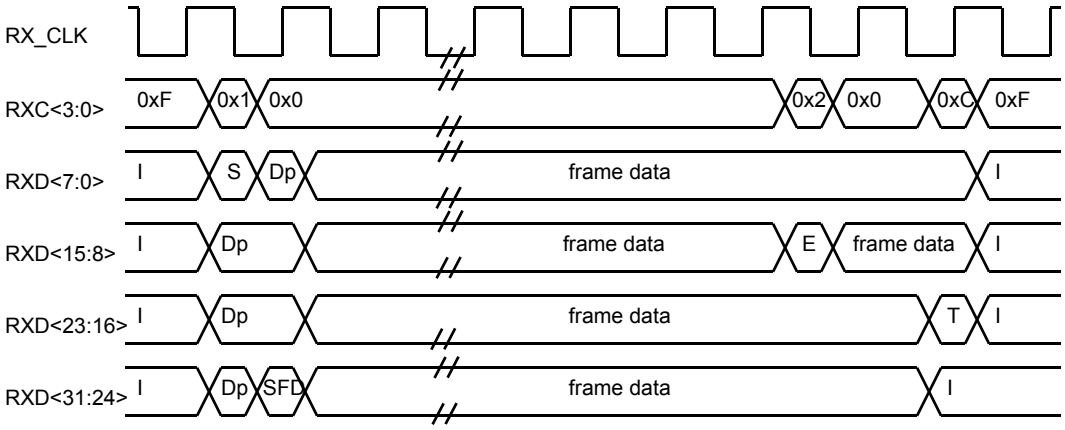

I: Idle control character, S: Start control character, Dp: preamble Data octet, T: Terminate control character, E: Error control character

Figure 46–9—Reception with error

46.3.2.4 Receive direction LPI transition

LPI operation and the LPI client are described in 78.1. When the PHY receives signals from the link partner to indicate transition into the low power state, it indicates this to the RS by asserting RXC and setting RXD to 0x06 (in all lanes). The PHY maintains these signals in this state while it remains in the LPI state. When the PHY receives signals from the link partner to indicate transition out of the LPI state, it indicates this to the RS by asserting RXC and asserting idle on lanes 0–3 to return to a normal interframe state. The RS shall interpret the LPI coding as shown in Table 46–4.

The PHY or DTE XS may halt RX_CLK at any time more than 128 clock cycles after the start of the LPI state as shown in Figure 46–10 if the clock stop enable bit is asserted (see 45.2.3.1.4 and 45.2.5.1.4). The PHY shall restart RX_CLK so that at least one positive transition occurs before it deasserts LPI.

Figure 46–10 shows the behavior of RXC and RXD<7:0> during LPI transitions.

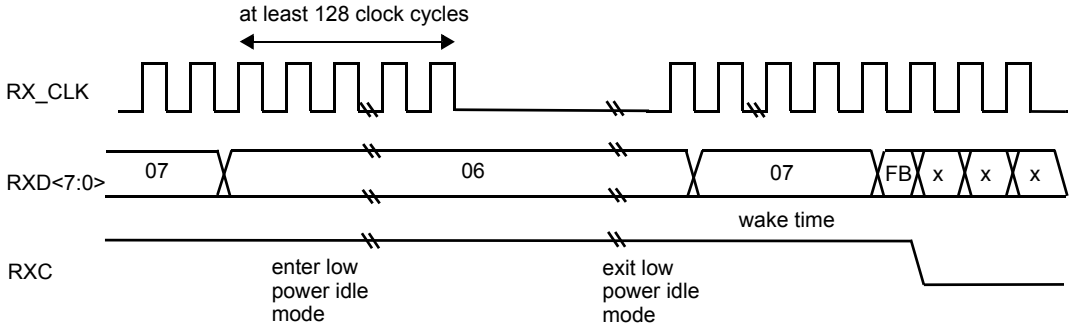

NOTE 1—RXC and RXD are shown for one lane, all 4 lanes behave identically during LPI.
NOTE 2—In some instances, LPI may be followed by characters other than IDLE during wake time.

Figure 46–10—LPI transition

46.3.3 Error and fault handling

46.3.3.1 Response to error indications by the XGMII

If, during frame reception (i.e., when DATA_VALID_STATUS = DATA_VALID), a control character other than a Terminate control character is signaled on a received lane, the RS shall ensure that the MAC will detect a FrameCheckError in that frame. This requirement may be met by incorporating a function in the RS that produces a received frame data sequence delivered to the MAC sublayer that is guaranteed to not yield a valid CRC result, as specified by the frame check sequence algorithm (see 3.2.8). This data sequence may be produced by substituting data delivered to the MAC. The RS generates eight PLS_DATA.indication primitives for each Error control character received within a frame, and may generate eight PLS_DATA.indication primitives to ensure FrameCheckError when a control character other than Terminate causes the end of the frame.

Other techniques may be employed to respond to a received Error control character provided that the result is that the MAC sublayer behaves as though a FrameCheckError occurred in the received frame.

46.3.3.2 Conditions for generation of transmit Error control characters

If, during the process of transmitting a frame, it is necessary to request that the PHY deliberately corrupt the contents of the frame in such a manner that a receiver will detect the corruption with the highest degree of probability, then an Error control character may be asserted on a transmit lane by the appropriate encoding of the lane's TXD and TXC signals.

46.3.3.3 Response to received invalid frame sequences

The 10 Gb/s PCS is required to either preserve the column alignment of the transmitting RS, or align the Start control character to lane 0. The RS shall not indicate DATA_VALID to the MAC for a Start control character received on any other lane. Error free 10 Gb/s operation will not change the SFD alignment in lane

3. A 10 Gb/s MAC/RS implementation is not required to process a packet that has an SFD in a position other than lane 3 of the column following the column containing the Start control character.

46.3.4 Link fault signaling

Link fault signaling operates between the remote RS and the local RS. Faults detected between the remote RS and the local RS are received by the local RS as Local Fault. Only an RS originates Remote Fault signals.

Sublayers within the PHY are capable of detecting faults that render a link unreliable for communication. Upon recognition of a fault condition a PHY sublayer indicates Local Fault status on the data path. When this Local Fault status reaches an RS, the RS stops sending MAC data or LPI, and continuously generates a Remote Fault status on the transmit data path (possibly truncating a MAC frame being transmitted). When Remote Fault or Link Interruption status is received by an RS, the RS stops sending MAC data or LPI, and continuously generates Idle control characters. When the RS no longer receives fault status messages, it returns to normal operation, sending MAC data or LPI.

Status is signaled in a four byte Sequence ordered set as shown in Table 46–5. The PHY indicates Local Fault with a Sequence control character in lane 0 and data characters of 0x00 in lanes 1 and 2 plus a data character of 0x01 in lane 3. The RS indicates a Remote Fault with a Sequence control character in lane 0 and data characters of 0x00 in lanes 1 and 2 plus a data character of 0x02 in lane 3. Though most fault detection is on the receive data path of a PHY, in some specific sublayers, faults can be detected on the transmit side of the PHY. This is also indicated by the PHY with a Local Fault status.

For operation with links that may be temporarily interrupted, optional detection of a third fault condition, Link Interruption, is provided. Link Interruption is indicated by the PHY receive function by continuously sending the Link Interruption ordered set as defined in Table 46–5.

Table 46–5—Sequence ordered sets

Lane 0	Lane 1	Lane 2	Lane 3	Description
Sequence	0x00	0x00	0x00	Reserved
Sequence	0x00	0x00	0x01	Local Fault
Sequence	0x00	0x00	0x02	Remote Fault
Sequence	0x00	0x00	0x03	Link Interruption
Sequence	≥ 0x00	≥ 0x00	≥ 0x04	Reserved
NOTE—Values in Lane 1, Lane 2, and Lane 3 columns are in hexadecimal, most significant bit to least significant bit (i.e., <7:0>). The link fault signaling state diagram allows future standardization of reserved Sequence ordered sets for functions other than link fault indications				

The RS reports the fault status of the link. Local Fault indicates a fault detected on the receive data path between the remote RS and the local RS. Remote Fault indicates a fault on the transmit path between the local RS and the remote RS. The RS shall implement the link fault signaling state diagram (see Figure 46–11).

46.3.4.1 Conventions

The notation used in the state diagram follows the conventions of 21.5. The notation ++ after a counter indicates it is to be incremented.

46.3.4.2 Variables and counters

The link fault signaling state diagram uses the following variables and counters:

col_cnt
: A count of the number of columns received not containing a fault_sequence. This counter increments at RX_CLK rate (on both the rising and falling clock transitions) unless reset.

fault_sequence
: A new column received on RXC<3:0> and RXD<31:0> comprising a Sequence ordered set of four bytes and consisting of a Sequence control character in lane 0 and a seq_type in lanes 1, 2, and 3 indicating either Local Faut, Remote Fault, or Link Interruption.

last_seq_type
: The seq_type of the previous Sequence ordered set received
Values:Local Fault; 0x00 in lane 1, 0x00 in lane 2, 0x01 in lane 3.
Remote Fault; 0x00 in lane 1, 0x00 in lane 2, 0x02 in lane 3.
Link Interruption; 0x00 in lane 1, 0x00 in lane 2, 0x03 in lane 3.

link_fault
: An indicator of the fault status.
Values:OK; No fault.
Local Fault; fault detected by the PHY.
Remote Fault; fault detection signaled by the remote RS.
Link Interruption; link temporarily unavailable, signaled by the PHY.

reset
: Condition that is true until such time as the power supply for the device that contains the RS has reached the operating region.
Values:FALSE: The device is completely powered and has not been reset (default).
TRUE: The device has not been completely powered or has been reset.

seq_cnt
: A count of the number of received Sequence ordered sets of the same type.

seq_type
: The value received in the current Sequence ordered set
Values:Local Fault; 0x00 in lane 1, 0x00 in lane 2, 0x01 in lane 3.
Remote Fault; 0x00 in lane 1, 0x00 in lane 2, 0x02 in lane 3.
Link Interruption; 0x00 in lane 1, 0x00 in lane 2, 0x03 in lane 3.

46.3.4.3 State diagram

The link fault signaling state diagram specifies the RS monitoring of RXC<3:0> and RXD<31:0> for Sequence ordered sets. The variable link_fault is set to indicate the value of a received Sequence ordered set when four fault_sequences containing the same fault value have been received with each pair of fault sequences separated by less than 128 columns and no intervening fault_sequences of a different fault value.

The variable link_fault is set to OK following any interval of 128 columns not containing a Remote Fault, Local Fault, or Link Interruption Sequence ordered set.

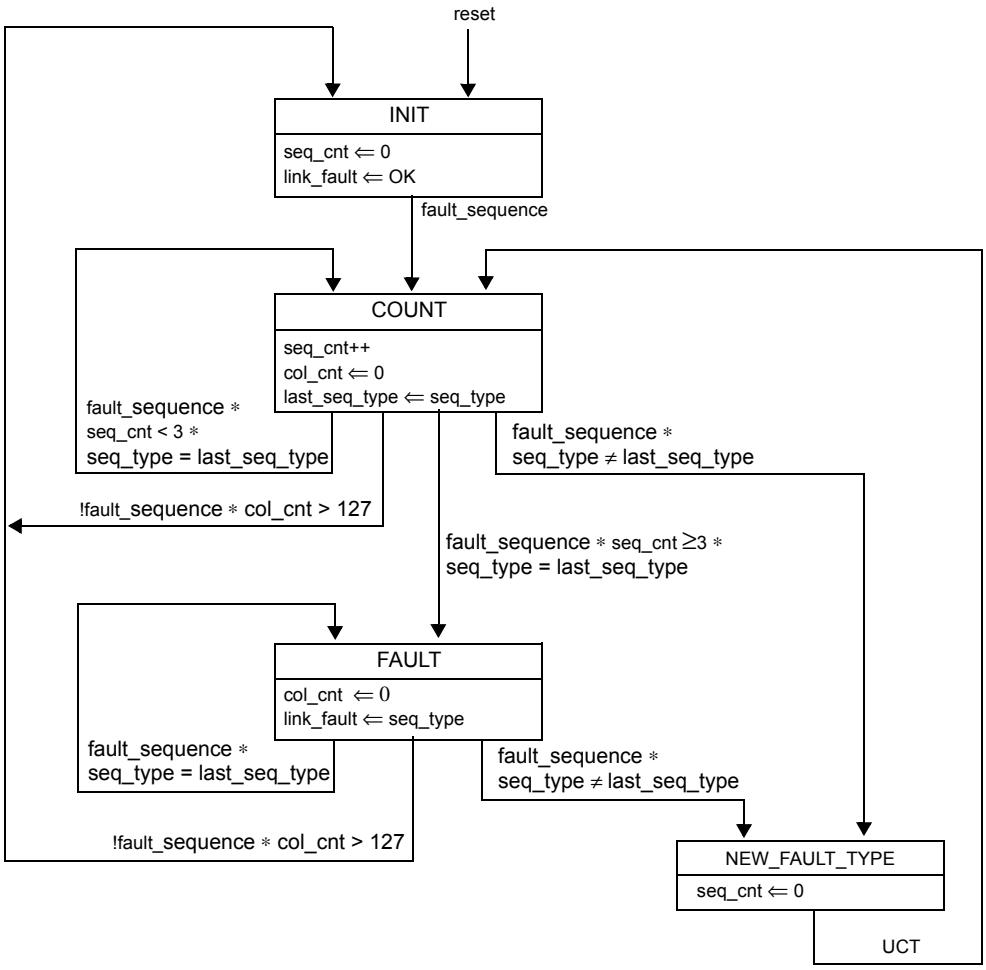

Figure 46–11—Link fault signaling state diagram

The RS output onto TXC<3:0> and TXD<31:0> is controlled by the variable link_fault.

a) link_fault = OK
   The RS shall send MAC frames as requested through the PLS service interface. In the absence of MAC frames, the RS shall generate Idle control characters.
b) link_fault = Local Fault
   The RS shall continuously generate Remote Fault Sequence ordered sets.
c) link_fault = Remote Fault or link_fault = Link Interruption
   The RS shall continuously generate Idle control characters.

46.4 LPI assertion and detection

Certain PHYs support Energy-Efficient Ethernet (see Clause 78). PHYs with EEE capability support LPI assertion and detection. LPI operation and the LPI client are described in 78.1. LPI signaling allows the RS to signal to the PHY and to the link partner that a break in the data stream is expected and components may use this information to enter power-saving modes that require additional time to resume normal operation. Similarly, it allows the LPI client to understand that the link partner has sent such an indication.

The LPI assertion and detection mechanism fits conceptually between the PLS Service Primitives and the XGMII signals as shown in Figure 46–12.

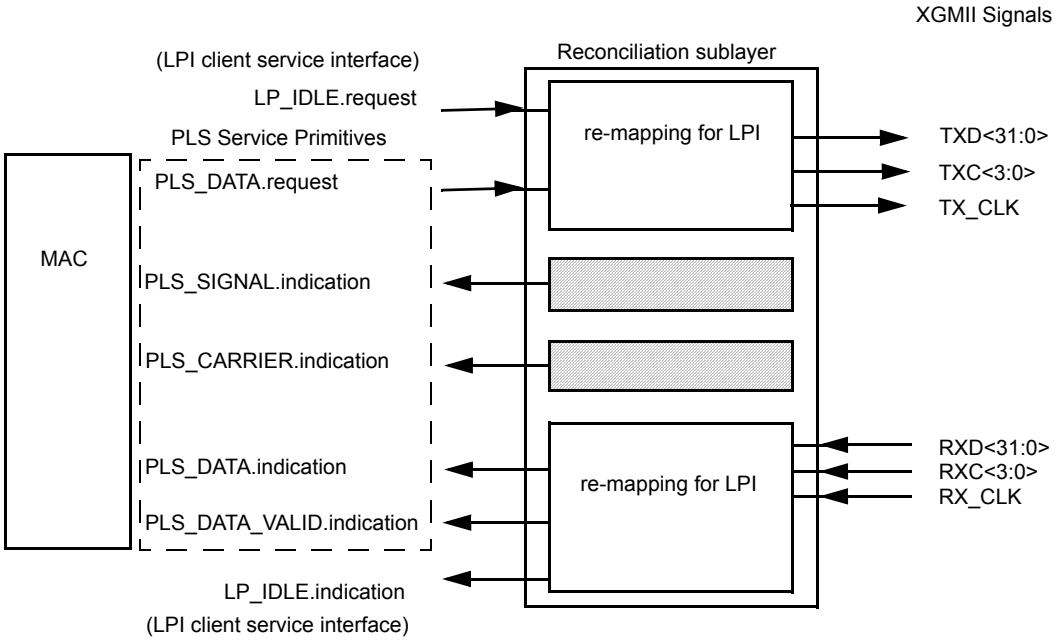

Figure 46–12—LPI assertion and detection mechanism

The definition of TXC<3:0> and TXD<31:0> is derived from the state of PLS_DATA.request (46.1.7), except when it is overridden by an assertion of LP_IDLE.request.

Similarly, RXC<3:0> and RXD<31:0> are mapped to PLS_DATA.indication except when LP_IDLE is detected

PLS_CARRIER.indication(CARRIER_STATUS) will be set to CARRIER_ON when the link is in LPI mode. See 46.1.7.3.

The timing of PLS_CARRIER.indication when used for the LPI function is controlled by the LPI transmit state diagram.

46.4.1 LPI messages

LP_IDLE.indication(LPI_INDICATION)

A primitive that indicates to the LPI client that the PHY has detected the assertion or de-assertion of LPI from the link partner.

Values:DEASSERT: The link partner is operating with normal interframe behavior (default).

ASSERT: The link partner has asserted LPI.

LP_IDLE.request(LPI_REQUEST)

The LPI_REQUEST parameter can take one of two values: ASSERT or DE-ASSERT. ASSERT initiates the signaling of LPI to the link partner. DE-ASSERT stops the signaling of LPI to the link partner. The effect of receipt of this primitive is undefined if link_status is not OK (see 28.2.6.1.1) or within 1 s of the change of link_status to OK.

46.4.2 Transmit LPI state diagram

The operation of LPI in the PHY requires that the MAC does not send valid data for a time after LPI has been de-asserted as governed by resolved Transmit T_{w_sys} defined in 78.4.2.3.

This wake-up time is enforced by the transmit LPI state diagram using CARRIER_SENSE.indication. The implementation shall conform to the behavior described by the transmit LPI state diagram shown in Figure 46–13.

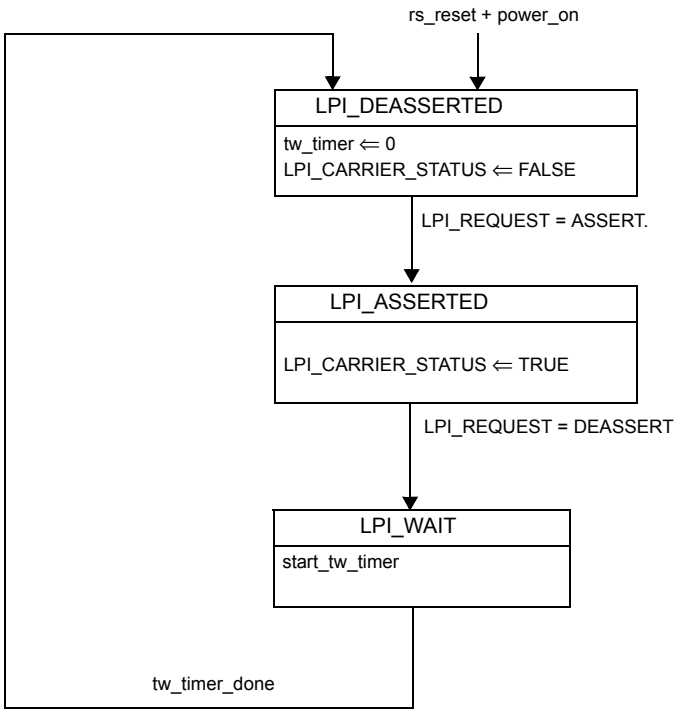

Figure 46–13—Transmit LPI state diagram

46.4.2.1 Variables and counters

The transmit LPI state diagram uses the following variables and counters:

LPI_CARRIER_STATUS
: The LPI_CARRIER_STATUS variable indicates how the CARRIER_STATUS parameter is controlled by the LPI_REQUEST parameter. The LPI_CARRIER_STATUS is either TRUE or FALSE as determined by the Transmit LPI state diagram in Figure 46–13.

power_on
: Condition that is true until such time as the power supply for the device that contains the RS has reached the operating region.
Values:FALSE: The device is completely powered (default).
TRUE: The device has not been completely powered.

rs_reset
: Used by management to control the resetting of the RS.
Values:FALSE: Do not reset the RS (default).
TRUE: Reset the RS.

tw_timer
: A timer that counts, in microseconds, the time since the de-assertion of LPI. The terminal count of the timer is the value of the resolved $T_{w\_sys\_tx}$ as defined in 78.2. If DTE XS XAUI stop enable bit is asserted (5.0.9), the terminal count of the timer is the value of the resolved $T_{w\_sys\_tx}$ as defined in 78.2 plus additional time equal to $T_{w\_sys\_tx} - T_{w\_sys\_rx}$ for the XGXS as shown in Table 78–4.

The signal tw_timer_done is asserted when tw_timer reaches its terminal count.

46.4.3 Considerations for transmit system behavior

The transmit system should expect that egress data flow will be halted for at least resolved $T_{w\_sys\_tx}$ (see 78.2) time, in microseconds, after it requests the de-assertion of LPI. Buffering and queue management should be designed to accommodate this.

46.4.4 Considerations for receive system behavior

The mapping function of the Reconciliation Sublayer shall continue to signal IDLE on PLS_DATA.indicate while it is detecting LP_IDLE on the XGMII. The receive system should be aware that data frames may arrive at the XGMII following the de-assertion of LPI_INDICATION with a delay corresponding to the link partner's resolved $T_{w\_sys\_rx}$ (as specified in 78.5) time, in microseconds.

If the PHY XS XAUI stop enable bit (4.0.9) is asserted, the PHY XS may stop signaling on the XAUI in the receive direction to conserve energy. The receiver should negotiate an additional 9.5 μs for the remote T_{w_sys} (equal to $T_{w\_sys\_tx} - T_{w\_sys\_rx}$ for the XGXS as shown in Table 78–4) before setting the PHY XS XAUI stop enable bit.

46.5 XGMII electrical characteristics

The electrical characteristics of the XGMII are specified such that the XGMII can be applied within a variety of 10 Gb/s equipment types. The electrical specifications are selected for an integrated circuit to integrated circuit application. The electrical characteristics specified in this clause apply to all XGMII signals.

When implemented as a chip-to-chip interface, the XGMII uses High Speed Transceiver Logic (HSTL), specified for a 1.5 volt output buffer supply voltage. XGMII chip-to-chip signals shall comply with EIA/JEDEC Standard EIA/JESD8-6 using Class I, output buffers. Output impedance shall be greater than 38Ω to assure acceptable overshoot and undershoot performance in an unterminated interconnection.

The thresholds and parametric values for HSTL XGMII signals are shown in the informative Figure 46–14 and informative Table 46–6.

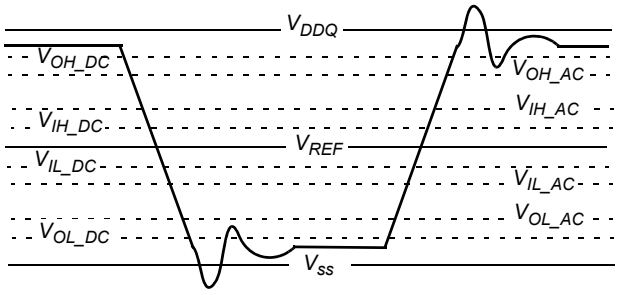

Figure 46–14—Electrical characteristics (informative)

The HSTL specification shows a number of termination options for different applications. Informative Figure 46–15 illustrates a possible XGMII circuit topology. Unterminated interconnection is recommended.

The XGMII chip-to-chip signals shall meet the timing requirements shown in Figure 46–16. All XGMII timing measurements shall be made at the XGMII driver output with the optional termination shown in Figure 46–15 and with a capacitive load from all sources of 10pF and are specified relative to the $V_{IL\_AC(max)}$ and $V_{IH\_AC(min)}$ thresholds.

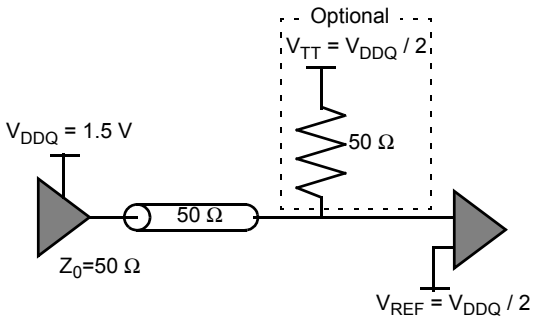

Figure 46–15—Circuit topology example (informative)

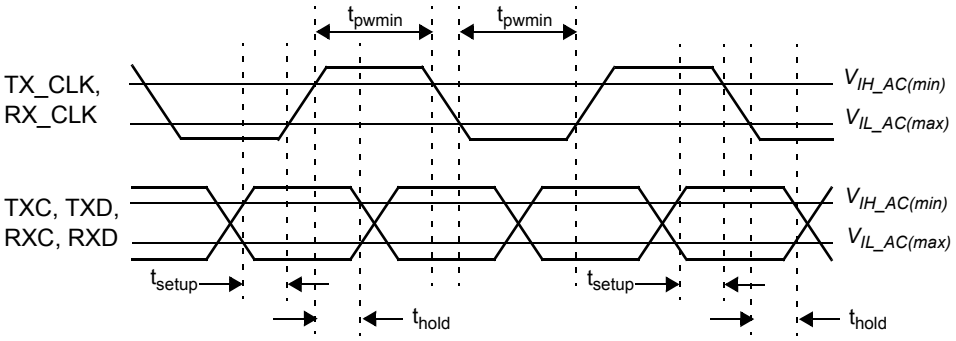

Symbol	Driver	Receiver	Units
t_{setup}	960	480	ps
t_{hold}	960	480	ps
t_{pwmin}	2.5	—	ns

Figure 46–16—TX_CLK and RX_CLK timing parameters

Table 46–6—DC and AC specifications (informative)

Symbol	Parameter	Min	Nom	Max	Units
V_{DDQ}	Output supply voltage	1.4	1.5	1.6	V
V_{REF}	Input reference voltage	0.68	0.75	0.90	V
V_{IH_DC}	DC input logic high	V_{REF}+0.10	—	V_{DDQ}+0.3	V
V_{IL_DC}	DC input logic low	–0.30	—	V_{REF}–0.1	V
V_{IH_AC}	AC input logic high	V_{REF}+0.20	—	—	V
V_{IL_AC}	AC input logic low	—	—	V_{REF}–0.20	V

46.6 Protocol implementation conformance statement (PICS) proforma for Clause 46, Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII)[2]

46.6.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 46, Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII), shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

46.6.2 Identification

46.6.2.1 Implementation identification

Supplier[1]	
Contact point for inquiries about the PICS[1]	
Implementation Name(s) and Version(s)[1,3]	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)[2]	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

46.6.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 46, Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII)
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [ ] Yes [ ] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2015.)	

Date of Statement	

[2]*Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

46.6.2.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*PHY	PHY support of XGMII	46.2, 46.3		O	Yes [ ] No [ ]
*RS	Reconciliation Sublayer support of XGMII	46.2, 46.3		O	Yes [ ] No [ ]
*XGE	XGMII electrical interface	46.4		O	Yes [ ] No [ ]
*LPI	Implementation of LPI	46.1.7		O	Yes [ ] No [ ]

46.6.3 PICS proforma Tables for Reconciliation Sublayer and 10 Gigabit Media Independent Interface

46.6.3.1 General

Item	Feature	Subclause	Value/Comment	Status	Support
G1	PHY support of MAC data rate	46.1.3	Support MAC data rate of 10 Gb/s	PHY:M	Yes [ ] N/A [ ]
G2	Cumulative MAC Control, MAC and RS round-trip delay	46.1.4	Per Table 46–1	RS:M	Yes [ ] N/A [ ]
G3	Lane structure	46.1.6	Per Table 46–2	M	Yes [ ]

46.6.3.2 Mapping of PLS service primitives

Item	Feature	Subclause	Value/Comment	Status	Support
PL1	Mapping to Clause 6	46.1.7	RS implements mapping to Clause 6 PLS service primitives	RS:M	Yes [ ] N/A [ ]
PL2	Mapping of PLS_DATA.requests	46.1.7.1.4	In sequence TXD<0> to TXD<31>	RS:M	Yes [ ] N/A [ ]
PL3	Start control character creation	46.1.7.1.4	First octet of preamble converted to Start control character	RS:M	Yes [ ] N/A [ ]
PL4	TXD and TXC generation	46.1.7.1.4	For each 32 PLS_DATA.requests	RS:M	Yes [ ] N/A [ ]
PL5	Terminate control character creation	46.1.7.1.4	DATA_COMPLETE causes creation of Terminate control character in next lane in sequence	RS:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
PL6	Mapping RXD to PLS_DATA.incicates	46.1.7.2.3	Create PLS_DATA.increments in sequence from RXD<0> to RXD<31>	RS:M	Yes [ ] N/A [ ]
PL7	PLS_DATA.indication generation	46.1.7.2.3	Generate 32 PLS_DATA.indications for each RXD<0:31> until Terminate then generating 0, 8, 16, or 24	RS:M	Yes [ ] N/A [ ]
PL8	Start control character conversion	46.1.7.2.3	Convert valid Start control character to preamble before generating PLS_DATA.indications	RS:M	Yes [ ] N/A [ ]
PL9	Terminate control character	46.1.7.2.3	No PLS_DATA.indications generated	RS:M	Yes [ ] N/A [ ]
PL10	PLS_DATA_VALID.indicatio n generation	46.1.7.5.3	On change of value of DATA_VALID_STATUS	RS:M	Yes [ ] N/A [ ]
PL11	DATA_VALID_STATUS	46.1.7.5.3	Value of DATA_VALID on a lane 0 Start control character preceded by four idles or a Sequence ordered set	RS:M	Yes [ ] N/A [ ]
PL12	DATA_VALID_STATUS	46.1.7.5.3	Value of DATA_NOT_VALID on any control character but Error	RS:M	Yes [ ] N/A [ ]
PL13	Frame not ending with Terminate control character	46.1.7.5.3	Ensure MAC detects CRC error	RS:M	Yes [ ] N/A [ ]

46.6.3.3 Data stream structure

Item	Feature	Subclause	Value/Comment	Status	Support
DS1	Frame transfer	46.2	Within XGMII data stream	RS:M	Yes [ ] N/A [ ]
DS2	Bit mapping	46.2	Per Figure 46–4	RS:M	Yes [ ] N/A [ ]
DS3	Content of <data>	46.2.3	Consist of data octets	RS:M	Yes [ ] N/A [ ]
DS4	Recognition of <efd>	46.2.4	Terminate recognized in any lane	RS:M	Yes [ ] N/A [ ]

46.6.3.4 LPI functions

Item	Feature	Subclause	Value/Comment	Status	Support
L1	Assertion of LPI in Tx direction	46.3.1.2	As defined in Table 46–3	LPI:M	Yes [ ] N/A [ ]
L2	Assertion of LPI in Rx direction	46.3.2.2	As defined in Table 46–4	LPI:M	Yes [ ] N/A [ ]
*L3	TX_CLK stoppable during LPI	46.3.1.5	At least 128 cycles after LPI assertion	LPI:O	Yes [ ] No [ ]
L4	TX_CLK restart before LPI deassert	46.3.1.5	At least 1 positive edge before LPI deassertion	L3:M	Yes [ ] N/A [ ]
L5	RX_CLK stoppable during LPI	46.3.2.4		LPI:O	Yes [ ] No [ ]

46.6.3.5 Link Interruption

Item	Feature	Subclause	Value/Comment	Status	Support
LINT	Detection of Link Interruption	46.3.4		O	Yes [ ] No [ ]
LINT1	CARRIER_STATUS response to Link Interruption	46.1.7.3	Set to CARRIER_ON if link_fault is Link Interruption	LINT:O	Yes [ ] No [ ]

46.6.3.6 XGMII signal functional specifications

NOTE—An XGXS adjacent to an RS exhibits the characteristics of a PHY for the items in this subclause, and an XGXS adjacent to a PCS exhibits the characteristics of an RS for the items in this subclause.

Item	Feature	Subclause	Value/Comment	Status	Support
FS1	TX_CLK active edges	46.3.1.1	TXD and TXC sampled on both edges of TX_CLK	XGE:M	Yes [ ] N/A [ ]
FS2	TX_CLK frequency	46.3.1.1	156.25 MHz ± 0.01%	XGE:M	Yes [ ] N/A [ ]
FS3	TXC assertion and de-assertion	46.3.1.2	De-asserted for data, asserted for control character	RS:M	Yes [ ] N/A [ ]
FS4	TXC clock	46.3.1.2	Synchronous to TX_CLK	XGE:M	Yes [ ] N/A [ ]
FS5	TXD encoding	46.3.1.3	Per Table 46–3	RS:M	Yes [ ] N/A [ ]
FS6	TXD clock	46.3.1.3	Synchronous to TX_CLK	XGE:M	Yes [ ] N/A [ ]
FS7	Start alignment	46.3.1.4	Start control character aligned to lane 0	RS:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
FS8	RX_CLK active edges	46.3.2.1	RXD and RXC sampled on both edges of RX_CLK	XGE:M	Yes [ ] N/A [ ]
FS9	RX_CLK frequency	46.3.2.1	156.25 MHz ± 0.01% when received data rate is within tolerance	XGE:M	Yes [ ] N/A [ ]
FS10	Loss of receive signal	46.3.2.1	Source RX_CLK from nominal clock	PHY:M	Yes [ ] N/A [ ]
FS11	Transition between clock sources	46.3.2.1	No decrease of RX_CLK period when switching sources	PHY:M	Yes [ ] N/A [ ]
FS12	RXC assertion and de-assertion	46.3.2.2	De-asserted for data, asserted for control character	PHY:M	Yes [ ] N/A [ ]
FS13	RXC clock	46.3.2.2	Synchronous to RX_CLK	XGE:M	Yes [ ] N/A [ ]
FS14	RXD decoding	46.3.2.3	Per Table 46–4	RS:M	Yes [ ] N/A [ ]
FS15	RXD clock	46.3.2.3	Synchronous to RX_CLK	XGE:M	Yes [ ] N/A [ ]
FS16	Received Error control character	46.3.3.1	RS cause MAC FrameCheckError	RS:M	Yes [ ] N/A [ ]
FS17	DATA_VALID assertion	46.3.3.3	RS not assert DATA_VALID unless Start control character in lane 0	RS:M	Yes [ ] N/A [ ]

46.6.3.7 Link fault signaling state diagram

Item	Feature	Subclause	Value/Comment	Status	Support
LF1	Link fault signaling state diagram	46.3.4	Implement per Figure 46–11	RS:M	Yes [ ] N/A [ ]
LF2	link_fault = OK and MAC frames	46.3.4.3	RS services MAC frame transmission requests	RS:M	Yes [ ] N/A [ ]
LF3	link_fault = OK and no MAC frames	46.3.4.3	In absence of MAC frames, RS transmits Idle control characters	RS:M	Yes [ ] N/A [ ]
LF4	link_fault = Local Fault	46.3.4.3	RS transmits continuous Remote Fault Sequence ordered sets	RS:M	Yes [ ] N/A [ ]
LF5	link_fault = Remote Fault	46.3.4.3	RS transmits continuous Idle control characters	RS:M	Yes [ ] N/A [ ]

46.6.3.8 Electrical characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
EC1	Referenced standard for signals	46.5	Signals to be compliant with EIA/JESD8-6 using Class I, output buffers	XGE:M	Yes [ ] N/A [ ]
EC2	Output impedance	46.5	Signal output impedance of at least 38 Ω	XGE:M	Yes [ ] N/A [ ]
EC3	Signal timing	46.5	Per Figure 46–16	XGE:M	Yes [ ] N/A [ ]
EC4	Signal measurement	46.5	At driver output with optional termination per Figure 46–15 and with a capacitive load from all sources of 10 pF	XGE:M	Yes [ ] N/A [ ]

47. XGMII Extender Sublayer (XGXS) and 10 Gigabit Attachment Unit Interface (XAUI)

47.1 Overview

This clause defines the functional and electrical characteristics for the optional XGMII Extender Sublayer (XGXS) and 10 Gigabit Attachment Unit Interface (XAUI). Figure 47–1 shows the relationships of the XGMII, XGMII Extender, XGXS, and XAUI.

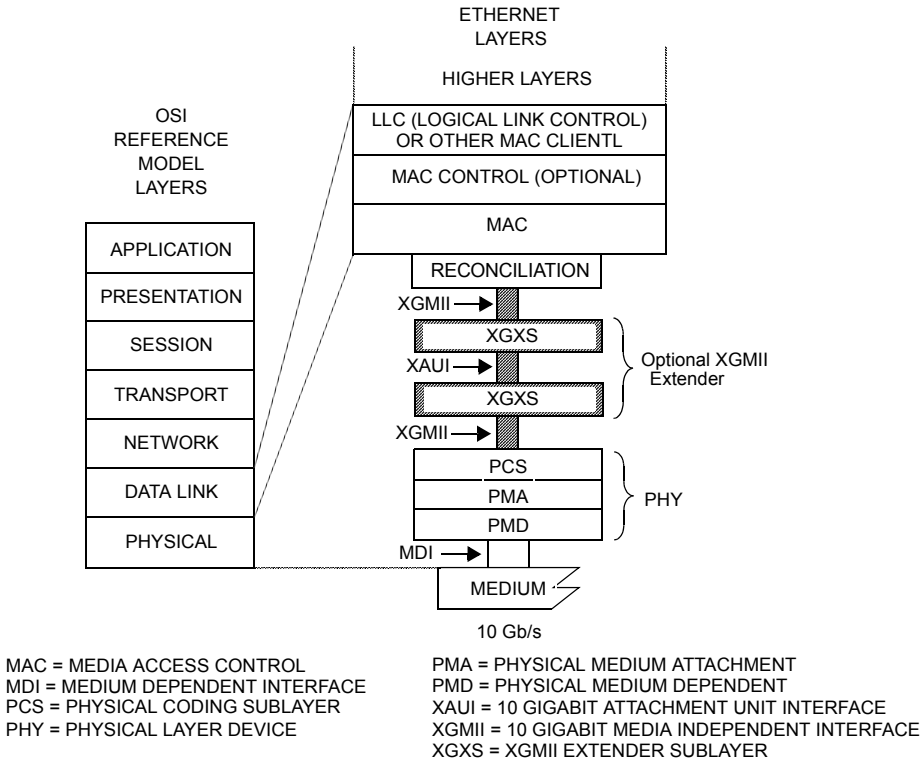

Figure 47–1—XAUI and XGXS relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

The purpose of the XGMII Extender, which is comprised of an XGXS at the RS end (DTE XGXS), an XGXS at the PHY end (PHY XGXS) and a XAUI between them, is to extend the operational distance of the XGMII and to reduce the number of interface signals. Applications include extending the physical separation possible between MAC and PHY components in a 10 Gigabit Ethernet system distributed across a circuit board.

An XGMII Extender with the optional Energy-Efficient Ethernet (EEE) capability (see 78.3) may enter a low power state to conserve energy during periods of low link utilization. The ability to support transition to a low power state is indicated by register 4.20.0 (for a PHY XS) or 5.20.0 (for a DTE XS). Transition to the low power state is enabled by register 4.0.9 (for a PHY XS) or 5.0.9 (for a DTE XS). The assertion of Low Power Idle (LPI) at the XGMII is encoded in the transmitted symbols. Detection of LPI encoding in the received symbols is indicated as LPI at the XGMII. When LPI is received on the transmit XGMII, an energy-efficient XGMII Extender sends sleep symbols, then, if enabled, ceases transmission and deactivates XAUI transmit signals to conserve energy. When the receiver sees the sleep symbols it transitions to a quiet state. The XGMII Extender periodically transmits during the quiet period to allow the attached XGMII Extender to refresh its receiver state (e.g., timing recovery, adaptive filter coefficients) and thereby track

long-term variation in the timing of the link or the underlying channel characteristics. If, during the quiet or refresh periods, normal interframe idle is asserted at the transmit XGMII, the XGMII Extender reactivates transmit functions and initiates transmission. This transmission will be detected by the attached XGMII Extender, causing it to also exit the low power state.

The optional XGMII Extender has the following characteristics:

a) Simple signal mapping to the XGMII
b) Independent transmit and receive data paths
c) Four lanes conveying the XGMII 32-bit data and control
d) Differential signaling with low voltage swing
e) Self-timed interface allows jitter control to the PCS
f) Shared technology with other 10 Gb/s interfaces
g) Shared functionality with other 10 Gb/s Ethernet blocks
h) Utilization of 8B/10B coding
i) Optionally extend LPI signaling to PHY for EEE
j) Optionally conserve energy during periods of low utilization

47.1.1 Summary of major concepts

The following is a list of the major concepts of XGXS and XAUI:

a) The optional XGMII Extender can be inserted between the Reconciliation Sublayer and the PHY to transparently extend the physical reach of the XGMII and reduce the interface pin count.
b) The XGMII is organized into four lanes with each lane conveying a data octet or control character on each edge of the associated clock. The source XGXS converts bytes on an XGMII lane into a self clocked, serial, 8B/10B encoded data stream. Each of the four XGMII lanes is transmitted across one of the four XAUI lanes.
c) The source XGXS converts XGMII Idle control characters (interframe) into an 8B/10B code sequence. The destination XGXS recovers clock and data from each XAUI lane and deskews the four XAUI lanes into the single-clock XGMII.
d) The destination XGXS adds to or deletes from the interframe as needed for clock rate disparity compensation prior to converting the interframe code sequence back into XGMII Idle control characters.
e) The XGXS uses the same code and coding rules as the 10GBASE-X PCS and PMA specified in Clause 48.

47.1.2 Application

This clause applies to the XGMII between the MAC and PHY. The implementation of the optional XGMII Extender is primarily intended as a chip-to-chip (integrated circuit to integrated circuit) interface implemented with traces on a printed circuit board. Where the XGMII is electrically limited to distances of approximately 7 cm, the XGMII Extender allows distances up to approximately 50 cm.

47.1.3 Rate of operation

The XGMII Extender supports the 10 Gb/s data rate of the XGMII. The 10 Gb/s MAC data stream is converted into four lanes at the XGMII (by the Reconciliation Sublayer for transmit or the PHY for receive). The byte stream of each lane is 8B/10B encoded by the XGXS for transmission across the XAUI at a nominal rate of 3.125 GBd. The XGXS at the PHY end of the XGMII Extender (PHY XGXS) and the XGXS at the RS end (DTE XGXS) may operate on independent clocks.

47.1.4 Allocation of functions

The XGMII Extender is transparent to the Reconciliation Sublayer and PHY device, and operates symmetrically with similar functions on the DTE transmit and receive data paths. The XGMII Extender is logically composed of two XGXSs interconnected with a XAUI data path in each direction. One XGXS acts as the source to the XAUI data path in the DTE transmit path and as the destination in the receive path. The other XGXS is the destination in the transmit path and source in the receive path. Each XAUI data path is composed of four serial lanes. All specifications for the XGMII Extender are written assuming conversion from XGMII to XAUI and back to XGMII, but other techniques may be employed provided that the result is that the XGMII Extender operates as if all specified conversions had been made. One example of this is the use of the optional XAUI with the 10GBASE-LX4 8B/10B PHY, where the XGXS interfacing to the Reconciliation Sublayer provides the PCS and PMA functionality required by the PHY. An XGXS layer is not required at the PHY end of the XAUI in this case. However, means may still be required to remove jitter introduced on the XAUI in order to meet PHY jitter requirements.

47.1.5 Global signal detect function

Global signal detect is mandatory for EEE capability, otherwise it is optional and its definition is beyond the scope of this standard. When global signal detect is not implemented, the value of SIGNAL_DETECT shall be set to OK for purposes of management and signaling of the primitive.

For EEE capability, the global signal detect function shall control the PMA SIGNAL_DETECT parameter. The SIGNAL_DETECT parameter can take on one of two values, OK or FAIL, indicating whether the XGXS is detecting electrical energy at the XAUI receiver (OK) or not (FAIL). When SIGNAL_DETECT = FAIL, PMA parameter rx_lane<3:0> is undefined.

47.1.6 Global transmit disable function

Global transmit disable is mandatory for EEE capability. The transmit disable function shall turn off all transmitter lanes after tx_mode changes to QUIET within a time and voltage level specified in 47.3.3.2. The transmit disable function shall turn on all transmitter lanes after tx_mode changes to DATA within a time and voltage level specified in 47.3.3.2.

47.2 Functional specifications

At the source side of a XAUI, the XGXS takes XGMII data streams striped over four lanes as its input, maps XGMII data and control characters into XAUI code-groups, and encodes them for transmission. At the destination end of the XAUI, the XGXS decodes the code-groups, deskews the four lanes, compensates for clock rate disparity, and maps the XAUI code-groups back into XGMII data and control characters. Each XGXS is bidirectional, having both source and destination functionality. Figure 47–2 depicts a schematic representation of the XGXS inputs and outputs.

47.2.1 PCS and PMA functionality

The XGXS shall meet all mandatory portions of 48.2 and 48.3, and may meet any optional portions of 48.2 and 48.3. Since the PHY XGXS operates with the XGMII below the XAUI, the transmit requirements of 48.2 and 48.3 apply to the PHY XGXS receive requirements and the receive requirements apply to the PHY XGXS transmit functionality.

47.2.2 Delay constraints

The XGMII Extender shall meet the delay constraints in 48.5. The contribution of the XAUI interconnect is included in these delay constraints.

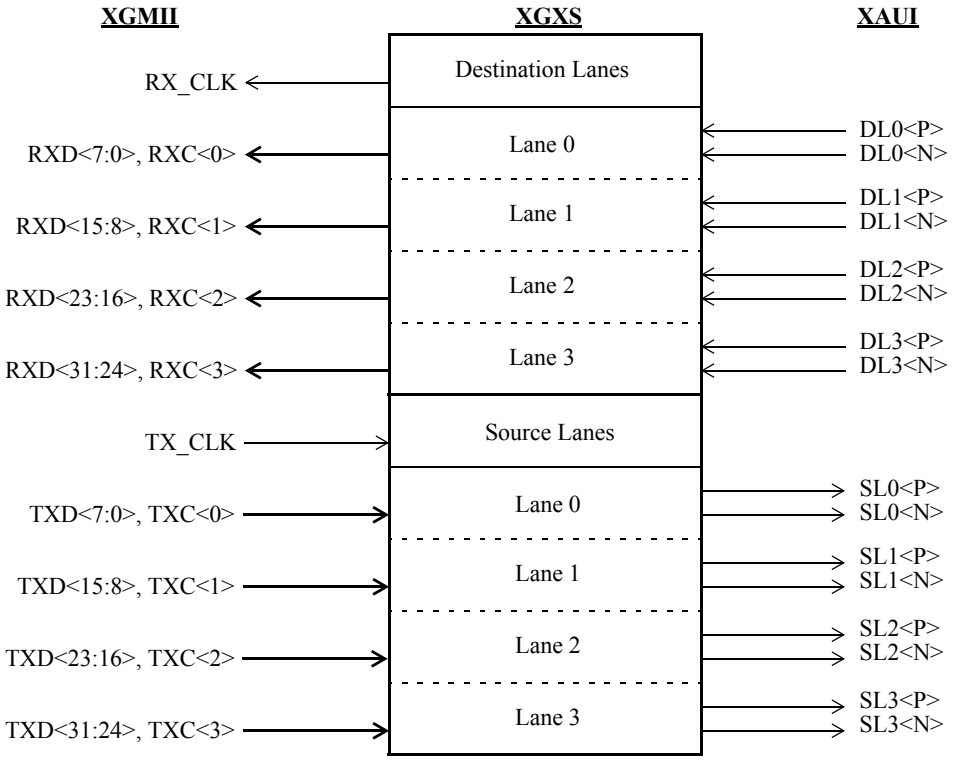

Figure 47–2—XGXS inputs and outputs

47.3 XAUI Electrical characteristics

The electrical characteristics of the XGMII Extender are specified such that it can be applied within a variety of 10 Gb/s Ethernet equipment types. There are two interface types associated with the XGMII Extender: the XGMII and XAUI. The electrical characteristics of both are selected for an integrated circuit to integrated circuit application. The electrical characteristics for the XGMII are specified in 46.5. The electrical characteristics for XAUI are specified in this subclause. Unless specified otherwise, the electrical characteristics defined in this subclause are applicable to all valid sequences of code-groups.

47.3.1 Signal levels

The XAUI is a low swing AC-coupled differential interface. AC-coupling allows for interoperability between components operating from different supply voltages. Low swing differential signaling provides noise immunity and improved electromagnetic interference (EMI). Differential signal swings defined in 47.3 depend on several factors, such as transmitter pre-equalization and transmission line losses.

47.3.2 Signal paths

The XAUI signal paths are point-to-point connections. Each path corresponds to a XAUI lane and is comprised of two complementary signals making a balanced differential pair. There are four differential

paths in each direction for a total of eight pairs, or sixteen connections. The signal paths are intended to operate up to approximately 50 cm over controlled impedance traces on standard FR4 printed circuit boards (PCBs).

47.3.3 Driver characteristics

The XAUI driver characteristics are summarized in Table 47–1. The XAUI Baud shall be 3.125 GBd ±100 ppm. The corresponding Baud period is nominally 320 ps.

Table 47–1—Driver characteristics

Parameter	Value	Units
Baud rate tolerance	3.125 GBd ± 100 ppm	GBd ppm
Unit interval nominal	320	ps
Differential amplitude maximum	1600	mV_{p-p}
Absolute output voltage limits maximum minimum	 2.3 –0.4	 V V
Differential output return loss minimum	[See Equation (47–1)]	dB
Output jitter Near-end maximums Total jitter Deterministic jitter Far-end maximums Total jitter Deterministic jitter	 ± 0.175 peak from the mean ± 0.085 peak from the mean ± 0.275 peak from the mean ± 0.185 peak from the mean	 UI UI UI UI

47.3.3.1 Load

The load is 100 Ω ± 5% differential to 2.5 GHz for these measurements, unless otherwise noted.

47.3.3.2 Amplitude and swing

Driver differential output amplitude shall be less than 1600 mV_{p-p} including any transmit equalization. DC-referenced logic levels are not defined since the receiver is AC-coupled. Absolute driver output voltage shall be between –0.4 V and 2.3 V with respect to ground. See Figure 47–3 for an illustration of absolute driver output voltage limits and definition of differential peak-to-peak amplitude.

For EEE capability, the transmitter lane's differential peak-to-peak output voltage shall be less than 30 mV within 500 ns of tx_quiet being asserted. Furthermore, the transmitter lane's differential peak-to-peak output voltage shall be greater than 800 mV within 500 ns of tx_quiet being de-asserted.

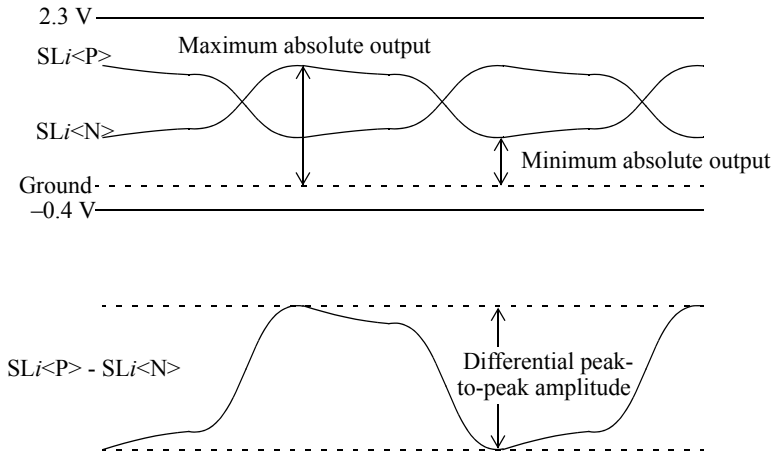

Figure 47–3—Driver output voltage limits and definitions
[L*i*<P> and L*i*<N> are the positive and negative sides of the differential signal pair for lane *i* (*i* = 0, 1, 2, 3)]

47.3.3.3 Transition time

Differential transition times between 60 and 130 ps are recommended, as measured between the 20% and 80% levels. Shorter transitions may result in excessive high-frequency components and increase EMI and crosstalk. The upper recommended limit of 130 ps corresponds to a sine wave at the half Baud.

47.3.3.4 Output impedance

For frequencies from 312.5 MHz to 3.125 GHz, the differential return loss of the driver shall exceed Equation (47–1). Differential return loss includes contributions from on-chip circuitry, chip packaging, and any off-chip components related to the driver. This output impedance requirement applies to all valid output levels. The reference impedance for differential return loss measurements is 100 Ω.

$$s_{11} = -10 \text{ dB for } 312.5 \text{ MHz} < \text{Freq } (f) < 625 \text{ MHz, and} \quad (47\text{–}1)$$
$$-10 + 10\log(f/625) \text{ dB for } 625 \text{ MHz} \leq \text{Freq } (f) \leq 3.125 \text{ GHz}$$

where f is frequency in MHz.

47.3.3.5 Driver template and jitter

The driver shall satisfy either the near-end eye template and jitter requirements, or the far-end eye template and jitter requirements. The eye templates are given in Figure 47–4 and Table 47–2. The template measurement requirements are specified in 47.4.2. The jitter requirements at the near end are for a maximum total jitter of ± 0.175 UI peak from the mean and a maximum deterministic component of ± 0.085 UI peak from the mean. The far-end requirements are for a maximum total jitter of ± 0.275 UI peak from the mean and a maximum deterministic component of ± 0.185 UI peak from the mean. Note that these values assume symmetrical jitter distributions about the mean. If a distribution is not symmetrical, its peak-to-peak total jitter value must be less than these total jitter values to claim compliance to the template requirements per the methods of 47.4.2. Jitter specifications include all but 10^{-12} of the jitter population. The maximum random jitter is equal to the maximum total jitter minus the actual deterministic jitter. Jitter measurement requirements are described in 47.4.3.

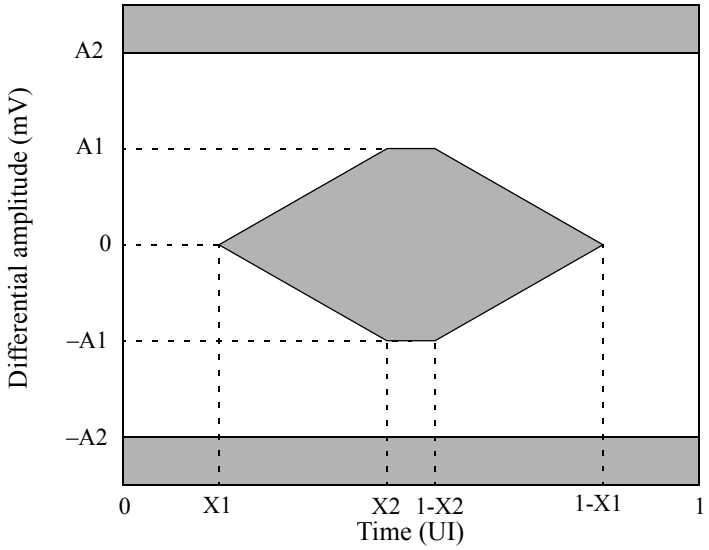

Figure 47–4—Driver template

Table 47–2—Driver template intervals

Symbol	Near-end value	Far-end value	Units
X1	0.175	0.275	UI
X2	0.390	0.400	UI
A1	400	100	mV
A2	800	800	mV

47.3.4 Receiver characteristics

Receiver characteristics are summarized in Table 47–3 and detailed in the following subclauses.

47.3.4.1 Bit error ratio

The receiver shall operate with a BER of better than 10^{-12} in the presence of a reference input signal as defined in 47.3.4.2.

47.3.4.2 Reference input signals

Reference input signals to a XAUI receiver have the characteristics determined by compliant XGXSs and XAUI drivers. Reference input signals satisfy the far-end template given in Figure 47–4 and Table 47–2 when the signal source impedance is 100 Ω ±5%. The template measurement requirements are specified in 47.4.2. Note that the input signal might not meet this template when this load is replaced by the actual receiver. Signal jitter does not exceed the jitter tolerance requirements specified in 47.3.4.6.

Table 47–3—Receiver characteristics

Parameter	Value	Units
Baud rate tolerance	3.125 ±100	GBd ppm
Unit interval (UI) nominal	320	ps
Receiver coupling	AC	
EEE Signal Detect deactivation time (T_{SD}) from active to LPI quiet	750	ns
EEE Signal Detect activation time (T_{SA}) from LPI quiet to active	750	ns
Return loss[a] differential common-mode	 10 6	 dB dB
Jitter amplitude tolerance[b]	0.65	UI_{p-p}

[a]Relative to 100 Ω differential and 25 Ω common-mode. See 47.3.4.5 for input impedance details.
[b]See 47.3.4.6 for jitter tolerance details.

47.3.4.3 Input signal amplitude

XAUI receivers shall accept differential input signal amplitudes produced by compliant transmitters connected without attenuation to the receiver. Note that this may be larger than the 1600 mV_{p-p} maximum of 47.3.3.2 due to actual driver and receiver input impedances. The minimum input amplitude is defined by the far-end driver template and the actual receiver input impedance. Note that the far-end driver template is defined using a well controlled load impedance. The minimum signal amplitude into an actual receiver may vary from the minimum template height due to the actual receiver input impedance. Since the XAUI receiver is AC-coupled to the XAUI, the absolute voltage levels with respect to the receiver ground are dependent on the receiver implementation.

47.3.4.4 AC-coupling

The XAUI receiver shall be AC-coupled to the XAUI to allow for maximum interoperability between various 10 Gb/s components. AC-coupling is considered to be part of the receiver for the purposes of this specification unless explicitly stated otherwise. It should be noted that there may be various methods for AC-coupling in actual implementations.

47.3.4.5 Input impedance

Receiver input impedance shall result in a differential return loss better than 10 dB and a common-mode return loss better than 6 dB from 100 MHz to 2.5 GHz. This includes contributions from on-chip circuitry, the chip package and any off-chip components related to the receiver. AC-coupling components are included in this requirement. The reference impedance for return loss measurements is 100 Ω for differential return loss and 25 Ω for common-mode.

47.3.4.6 Jitter tolerance

The XAUI receiver shall have a peak-to-peak total jitter amplitude tolerance of at least 0.65 UI. This total jitter is composed of three components: deterministic jitter, random jitter, and an additional sinusoidal jitter. Deterministic jitter tolerance shall be at least 0.37 UI_{p-p}. Tolerance to the sum of deterministic and random jitter shall be at least 0.55 UI_{p-p}. The XAUI receiver shall tolerate an additional sinusoidal jitter with any

frequency and amplitude defined by the mask of Figure 47–5. This additional component is intended to ensure margin for low-frequency jitter, wander, noise, crosstalk and other variable system effects. Jitter specifications include all but 10^{-12} of the jitter population. Jitter tolerance test requirements are specified in 47.4.3.

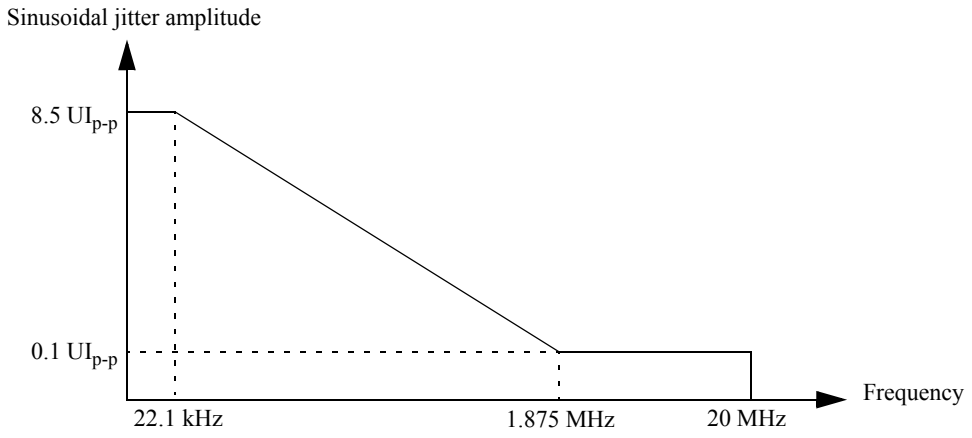

Figure 47–5—Single-tone sinusoidal jitter mask

47.3.4.7 EEE receiver timing

For EEE capability, the receiver shall meet the timing requirements shown in Table 47–3 for Signal_Detect activation and deactivation.

47.3.5 Interconnect characteristics

The XAUI is primarily intended as a point-to-point interface of up to approximately 50 cm between integrated circuits using controlled impedance traces on low-cost printed circuit boards (PCBs). Informative loss and jitter budgets are presented in Table 47–4 to demonstrate the feasibility of standard FR4 epoxy PCBs. The performance of an actual XAUI interconnect is highly dependent on the implementation. The compliance interconnect limit of 47.4.1 represents the median performance of a range of interconnect designs. The range included designs from 46 to 56 cm in total length, having trace widths of 0.125 to 0.300 mm, and using different grades and thicknesses of FR4. Interconnect configurations ranged from single-board designs to systems of two daughter cards mating to a backplane through high-speed electrical connectors.

47.3.5.1 Characteristic impedance

The recommended differential characteristic impedance of circuit board trace pairs is 100 Ω ± 10% from 100 MHz to 2.5 GHz.

47.3.5.2 Connector impedance

The recommended impedance of any connectors, such as used between circuit board subsystems, is 100 Ω ± 30%.

Table 47–4—Informative XAUI loss, skew and jitter budget

	Loss (dB)[a]	Differential skew (ps_{p-p})	Total jitter (UI_{p-p})[c]	Deterministic jitter (UI_{p-p})[c]
Driver	0	15	0.35	0.17
Interconnect	7.5	60	0.20	0.20
Other[b]	4.5		0.10	0.10
Total	12.0	75	0.65	0.47

[a]Budgetary loss in height of eye opening.
[b]Includes such effects as crosstalk, noise, and interaction between jitter and eye height.
[c]Jitter specifications include all but 10^{-12} of the jitter population.

47.4 Electrical measurement requirements

47.4.1 Compliance interconnect definition

The compliance interconnect is a 100 Ω differential system specified with respect to transmission magnitude response and intersymbol interference (ISI) loss. The compliance interconnect limits have been chosen to allow a realistic differential interconnect of about 50 cm length on FR4 epoxy PCB. See 47.3.5 for a more detailed description of the target XAUI interconnect. The transmission magnitude response, $|s_{21}|$, of the compliance interconnect in dB satisfies Equation (47–2).

$$|s_{21}| \le |s_{21}|_{limit} = -20\log(e) \times [a_1\sqrt{f} + a_2 f + a_3 f^2] \quad (47\text{–}2)$$

where f is frequency in Hz, $a_1 = 6.5 \times 10^{-6}$, $a_2 = 2.0 \times 10^{-10}$, and $a_3 = 3.3 \times 10^{-20}$. This limit applies from DC to 3.125 GHz. The magnitude response above 3.125 GHz does not exceed –11.4 dB. The ISI loss, defined as the difference in magnitude response between two frequencies, is greater than 4.0 dB between 312.5 MHz and 1.5625 GHz. The magnitude response and ISI loss limits are illustrated in Figure 47–6.

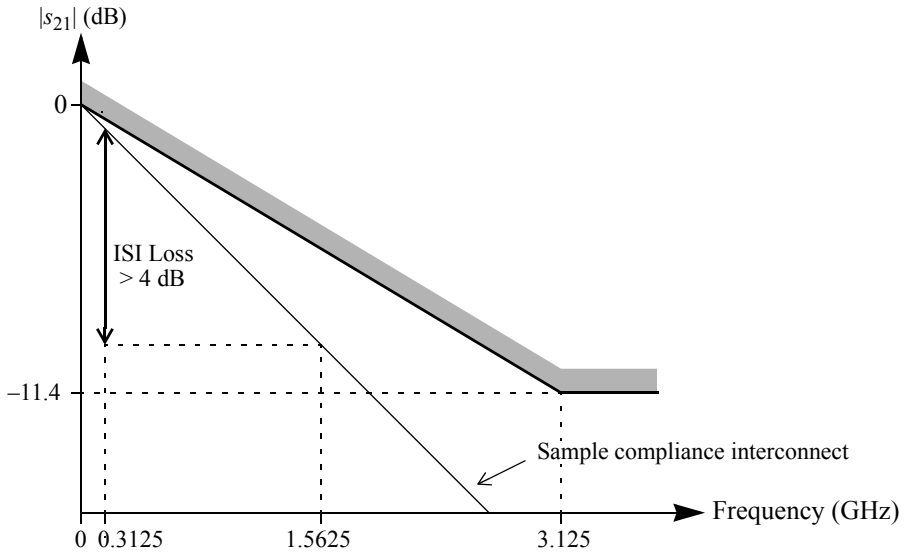

Figure 47–6—Compliance interconnect magnitude response and ISI loss

47.4.2 Eye template measurements

For the purpose of eye template measurements, the effect of a single-pole high-pass filter with a 3 dB point at 1.875 MHz is applied to the jitter. See 48B.1.3 for an explanation of this technique. The data pattern for template measurements is the CJPAT pattern defined in Annex 48A. All XAUI lanes are active in both the transmit and receive directions, and opposite ends of the link use asynchronous clocks. The amount of data represented in the data eye must be adequate to ensure a bit error ratio of less than 10^{-12}. The eye template is measured with AC-coupling and centered at 0 volts differential. The left and right edges of the template are aligned with the mean zero crossing points of the measured data eye, as illustrated in Figure 47–7. The near-end load for this test is specified in 47.3.3.1. The far-end template is measured at the end of the compliance interconnect specified in 47.4.1. The far-end load for the compliance link is specified in 47.3.3.1.

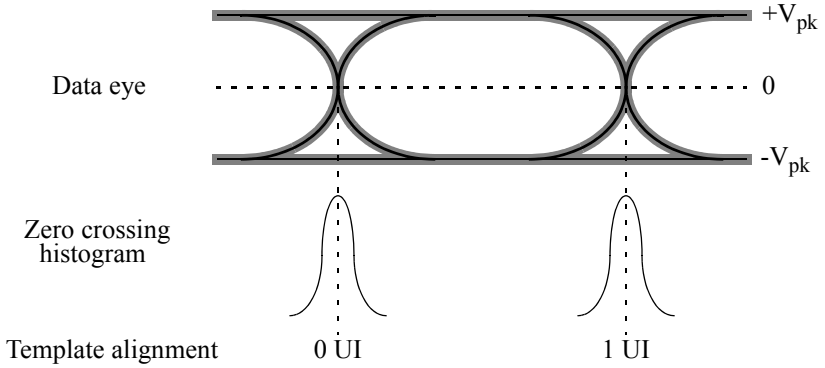

Figure 47–7—Eye template alignment

47.4.3 Jitter test requirements

For the purpose of jitter measurement, the effect of a single-pole high-pass filter with a 3 dB point at 1.875 MHz is applied to the jitter. The data pattern for jitter measurements is the CJPAT pattern defined in Annex 48A. All four lanes of XAUI are active in both directions, and opposite ends of the link use asynchronous clocks. Jitter is measured with AC-coupling and at 0 volts differential. Jitter measurement for the transmitter (or for calibration of a jitter tolerance setup) shall be performed with a test procedure resulting in a BER bathtub curve such as that described in Annex 48B.

47.4.3.1 Transmit jitter

Transmit near-end jitter is measured at the driver output when terminated into the load specified in 47.3.3.1. Far-end jitter is measured at the end of a compliance interconnect specified in 47.4.1. The far-end load for the compliance link is specified in 47.3.3.1.

47.4.3.2 Jitter tolerance

Jitter tolerance is measured at the receiver using a jitter tolerance test signal. This signal is obtained by first producing the required sum of deterministic and random jitter defined in 47.3.4.6 and then adjusting the signal amplitude until the data eye contacts the 6 points of the driver's template shown in Figure 47–4 and Table 47–2. Note that for this to occur, the test signal must have vertical waveform symmetry about the average value and have horizontal symmetry (including jitter) about the mean of the zero crossing. If these symmetries are not achieved, then some portions of the test signal will encroach into the template and provide overstress of the receiver, and/or some points of the template may not be contacted, resulting in understress of the receiver. Eye template measurement requirements are given in 47.4.2. Random jitter is calibrated using a high-pass filter with a low-frequency corner of 20 MHz and 20 dB/decade rolloff below this. The required sinusoidal jitter specified in 47.3.4.6 is then added to the signal and the far-end load is replaced by the receiver being tested.

47.5 Environmental specifications

All equipment subject to this clause shall conform to the requirements of 14.7 and applicable sections of ISO/IEC 11801:1995.

47.6 Protocol implementation conformance statement (PICS) proforma for Clause 47, XGMII Extender (XGMII) and 10 Gigabit Attachment Unit Interface (XAUI)[3]

47.6.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 47, XGMII Extender Sublayer (XGXS) and 10 Gigabit Attachment Unit Interface (XAUI), shall complete the following protocol implementation conformance statement (PICS) proforma. A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

47.6.2 Identification

47.6.2.1 Implementation identification

Supplier[1]	
Contact point for inquiries about the PICS[1]	
Implementation Name(s) and Version(s)[1,3]	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)[2]	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

47.6.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 47, XGMII Extender Sublayer (XGXS) and 10 Gigabit Attachment Unit Interface (XAUI)
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [ ] Yes [ ] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2015.)	
Date of Statement	

[3]*Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

47.6.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
XGE	XGMII compatibility interface	46, 47.1.4	Compatibility interface is supported	O	Yes [ ] No [ ]
MD	MDIO	47.2	Registers and interface supported	O	Yes [ ] No [ ]
LPI	Implementation of LPI	47.1		O	Yes [ ] No [ ]

47.6.4 PICS Proforma tables for XGXS and XAUI

47.6.4.1 Compatibility considerations

Item	Feature	Subclause	Value/Comment	Status	Support
C1	Environmental specifications	47.5		M	Yes [ ]

47.6.4.2 XGXS and XAUI functions

Item	Feature	Subclause	Value/Comment	Status	Support
F1	XGXS PCS and PMA requirements	47.2.1	Meet all mandatory requirements of 48.2 and 48.3	M	Yes [ ]
F2	XGXS PCS and PMA options	47.2.1	Meet optional clock rate compensation in unencoded data stream per 48.2.4.2.3	O	Yes [ ] No [ ]
F3	XGMII Extender delay	47.2.2	Meets delay constraints of 48.5	M	Yes [ ]

47.6.4.3 Electrical characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
E1	XAUI Baud	47.3.3	3.125 GBd ± 100 ppm	M	Yes [ ]
E2	Driver output amplitude	47.3.3.2	Less than 1600 mV_{p-p}	M	Yes [ ]
E3	Driver output swing	47.3.3.2	Between –0.4 and +2.3 V	M	Yes [ ]
E4	Driver output impedance	47.3.3.4	s_{11} = –10 dB for 312.5 MHz < Freq (f) < 625 MHz, and –10 + 10log(f/625) dB for 625 MHz ≤ Freq (f)≤ 3.125 GHz (f is frequency in MHz)	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
E5	Driver near-end template and jitter	47.3.3.5	At driver output	O.1	Yes [ ] No [ ]
E6	Driver far-end template and jitter	47.3.3.5	At end of compliance interconnect	O.1	Yes [ ] No [ ]
E7	Receiver bit error ratio	47.3.4	Less than 10^{-12}	M	Yes [ ]
E8	Receive input amplitude tolerance	47.3.4.3	May be larger than 1600 mV_{p-p}	M	Yes [ ]
E9	Receiver coupling	47.3.4.4	AC-coupled	M	Yes [ ]
E10	Receiver input impedance	47.3.4.5	At least 10 dB differential and 6 dB common-mode return loss	M	Yes [ ]
E11	Total jitter tolerance	47.3.4.6	At least 0.65 UI	M	Yes [ ]
E12	Deterministic jitter tolerance	47.3.4.6	At least 0.37 UI	M	Yes [ ]
E13	Tolerance to sum of deterministic and random jitter	47.3.4.6	At least 0.55 UI	M	Yes [ ]
E14	Additional sinusoidal jitter tolerance	47.3.4.6	Per Figure 47–5	M	Yes [ ]
E15	Jitter measurement	47.4.3	Meet BER bathtub curve, See Annex 48B	M	Yes [ ]

47.6.4.4 LPI functions

Item	Feature	Subclause	Value/Comment	Status	Support
LP-01	Global signal detect	47.1.5	Meets the requirements of 47.1.5	LPI:M	Yes [ ] No [ ]
LP-02	Global transmit disable	47.1.6	Meets the requirements of 47.1.6	LPI:M	Yes [ ] No [ ]
LP-03	Transmit amplitude	47.3.3.2	Meets the requirements of 47.3.3.2	LPI:M	Yes [ ] No [ ]
LP-04	Signal detect timing	47.3.4	Meets the requirements of Table 47–3	LPI:M	Yes [ ] No [ ]

48. Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer, type 10GBASE-X

48.1 Overview

This clause specifies the Physical Coding Sublayer (PCS) and the Physical Medium Attachment (PMA) sublayer that are common to a family of 10 Gb/s Physical Layer implementations, collectively known as 10GBASE-X. The 10GBASE-X PHY family consists of 10GBASE-CX4 (see Clause 54), 10GBASE-KX4 (see Clause 71), and 10GBASE-LX4 (see Clause 53).

The 10GBASE-X PCS and PMA sublayers are also utilized by the XGXS specified in Clause 47.

10GBASE-X PCS and PMA sublayers map the interface characteristics of the PMD sublayer (including MDI) to the services expected by the Reconciliation Sublayer (RS) and the logical and electrical characteristics of the 10 Gigabit Media Independent Interface (XGMII). Although the XGMII is optional, it is used as the basis for the definition of the 10GBASE-X PCS and PMA sublayers.

10GBASE-X assumes the use of the MDIO interface and register set for communication between PHY and Station Management (STA) entities, see Clause 45.

10GBASE-X has the following characteristics:

a) The capability of supporting 10 Gb/s operation at the XGMII and RS
b) Clock references embedded in all data and control code-groups
c) Data paths consisting of independent serial links called lanes
d) Independent four-lane-wide transmit and receive data paths
e) Simple signal mapping to the XGMII and RS
f) Full duplex operation
g) Shared technology with other 10 Gb/s interfaces
h) Shared functionality with other 10 Gb/s Ethernet blocks

48.1.1 Objectives

The following are the objectives of 10GBASE-X:

a) Support the IEEE 802.3 MAC
b) Provide a data rate of 10 Gb/s at the XGMII
c) Support cable plants using cabled optical fiber compliant with second edition of ISO/IEC 11801:1995
d) Support a BER objective of 10^{-12}
e) Support the optional XAUI
f) Support link fault and error indications

48.1.2 Relationship of 10GBASE-X to other standards

Figure 48–1 depicts the relationships among the 10GBASE-X sublayers (shown shaded), the IEEE 802.3 MAC and RS, and the IEEE 802.2 LLC.

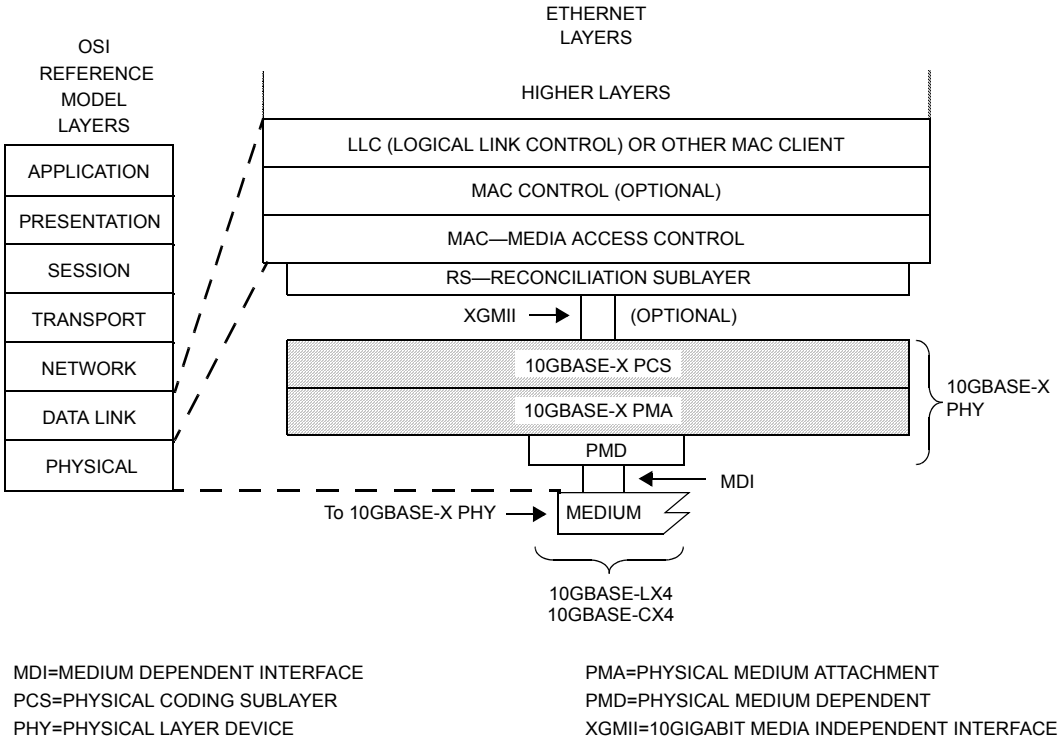

Figure 48–1—10GBASE-X PCS and PMA relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE Ethernet Model

48.1.3 Summary of 10GBASE-X sublayers

The following provides an overview of the 10GBASE-X sublayers.[4]

48.1.3.1 Physical Coding Sublayer (PCS)

The interface between the PCS and the RS is the XGMII as specified in Clause 46. The 10GBASE-X PCS provides services to the XGMII in a manner analogous to how the 1000BASE-X PCS provides services to the 1000 Mb/s GMII.

The 10GBASE-X PCS provides all services required by the XGMII and in support of the 10GBASE-X PMA, including:

a) Encoding of 32 XGMII data bits and 4 XGMII control bits to four parallel lanes conveying 10-bit code-groups each, for communication with the underlying PMA.
b) Decoding of four PMA parallel lanes, conveying 10-bit code-groups each, to 32 XGMII data bits and 4 XGMII control bits.
c) Synchronization of code-groups on each lane to determine code-group boundaries.

[4] The 10GBASE-X PHY consists of that portion of the Physical Layer between the MDI and XGMII consisting of the PCS, PMA, and PMD sublayers. The 10GBASE-X PHY is roughly analogous to the 1000BASE-X PHY.

d) Deskew of received code-groups from all lanes to an alignment pattern.
e) Support of the MDIO interface and register set as specified in Clause 45 to report status and enable control of the PCS.
f) Conversion of XGMII Idle control characters to (from) a randomized sequence of code-groups to enable serial lane synchronization, clock rate compensation and lane-to-lane alignment.
g) Clock rate compensation protocol.
h) Link Initialization based on the transmission and reception of the Idle sequence.
i) Link status reporting for fault conditions.

48.1.3.2 Physical Medium Attachment (PMA) sublayer

The PMA provides a medium-independent means for the PCS to support the use of a range of serial-bit-oriented physical media. The 10GBASE-X PMA performs the following functions:

a) Mapping of transmit and receive code-groups between the PCS and PMA via the PMA service interface.
b) Serialization (deserialization) of code-groups for transmission (reception) on the underlying serial PMD.
c) Clock recovery from the code-groups supplied by the PMD.
d) Mapping of transmit and receive bits between the PMA and PMD via the PMD service interface.
e) Direct passing of signal_detect from the PMD to the PCS through the PMA via the PMD and PMA service interfaces.

48.1.3.3 Physical Medium Dependent (PMD) sublayer

10GBASE-X uses the PMD sublayer and MDI specified in Clause 53, Clause 54, and Clause 71. The 10GBASE-CX4, 10GBASE-KX4, and 10GBASE-LX4 perform the following functions:

a) Transmission of quad serial bit streams on the underlying medium.
b) Reception of quad serial bit streams on the underlying medium.

48.1.4 Rate of operation

The 10GBASE-X PCS and PMA support the 10 Gb/s MAC data rate. The line rate of each of four PMA lanes is 3.125 GBd ± 100 ppm.

48.1.5 Allocation of functions

PCS and PMA functions directly map onto the 10GBASE-X PMD, MDI and medium which attach, in turn, to another 10GBASE-X PHY. In addition, 10GBASE-X PCS and PMA functions embodied in the XGXS described in Clause 47 may be used to attach to alternate 10 Gb/s PHYs such as 10GBASE-R or 10GBASE-W.

The longer interconnect distances afforded through the specification of a self-clocked serial architecture enable significant implementation flexibility while imposing a requirement on those implementations to ensure sufficient signal fidelity over the link. The implementer of this standard is expected to meet the required specifications in this and related clauses through implementation methods not specified by this standard.

Certain PHYs support Energy-Efficient Ethernet (EEE) (see Clause 78). PHYs that support EEE (see 78.3) use Low Power Idle (LPI) signaling to allow systems on both sides of the link to save power during periods of low link utilization. LPI signaling may optionally be used by XGXS to extend the EEE function to attached PHYs. Both PHY and DTE XGXS may optionally use LPI signaling to control the shutdown of signals on the XAUI to reduce power for PHY attachments.

48.1.6 Inter-sublayer interfaces

There are a number of interfaces employed by 10GBASE-X. Some (such as the PMA service interface) use an abstract service model to define the operation of the interface. Figure 48–2 depicts the relationship and mapping of the services provided by all of the interfaces relevant to 10GBASE-X.

Multiple optional physical instantiations of the PCS service interface have been defined. One is the XGMII described in Clause 46. The other is the interface to the XGXS described in Clause 47.

Physical instantiations of the 10GBASE-X PMA and PMD service interfaces are not defined in this standard.

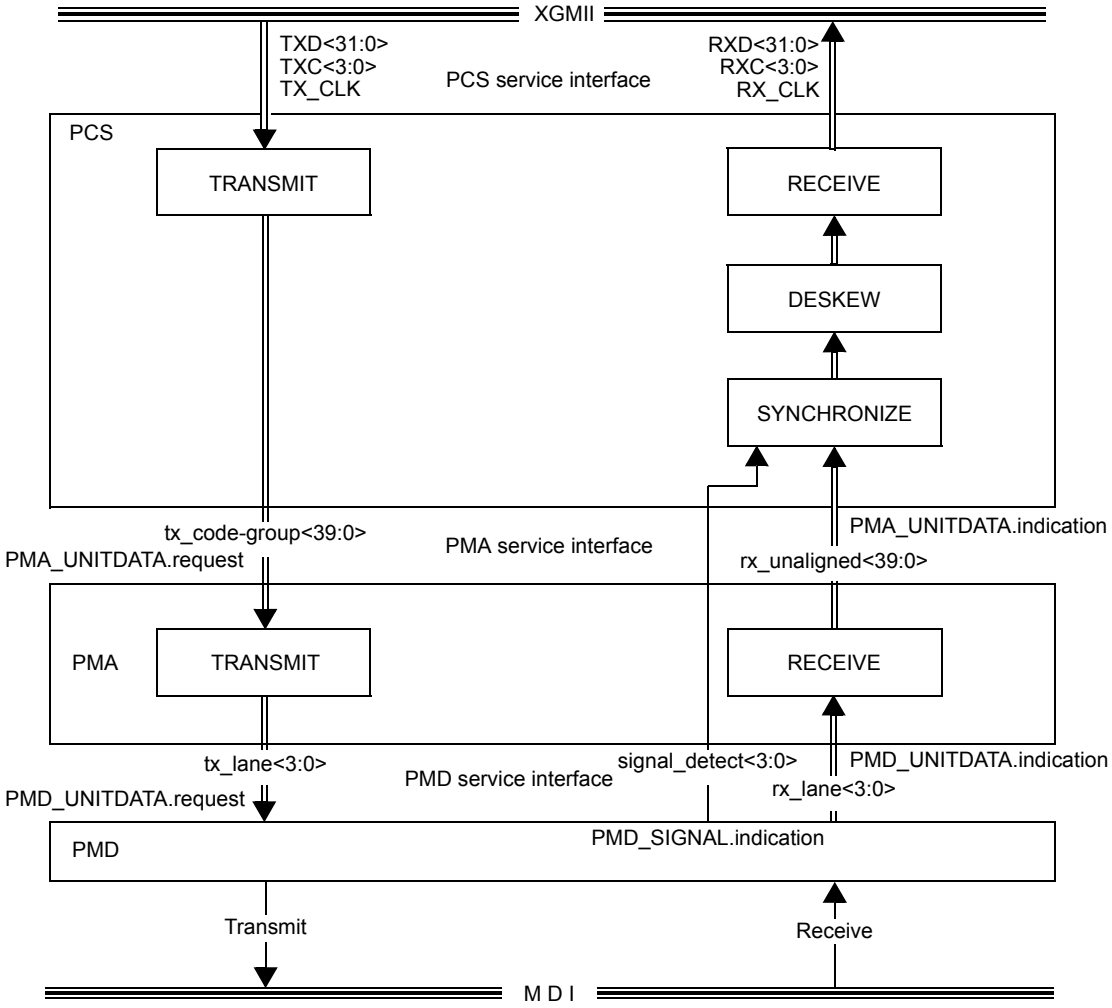

Figure 48–2—Functional block diagram

48.1.7 Functional block diagram

Figure 48–2 provides a functional block diagram of the 10GBASE-X PHY.

48.1.8 Special symbols

/x/ The code-group x is represented by preceding and following slash characters.

||y|| Four code-groups, one each in lanes 0 through 3 inclusive, synchronous to each other and arranged in a column identified by the value y, is represented by preceding and following double bar characters.

48.2 Physical Coding Sublayer (PCS)

48.2.1 PCS service interface (XGMII)

The PCS service interface allows the 10GBASE-X PCS to transfer information to and from the PCS client. The PCS client is the RS defined in Clause 46, or the XGXS defined in Clause 47. An instantiation of the PCS service interface is the XGMII defined in Clause 46.

In the transmit direction, the 10GBASE-X PCS accepts packets from the PCS client on the XGMII. Due to the continuously signaled nature of the underlying PMA, and the encoding performed by the PCS, the 10GBASE-X PCS maps XGMII data and control characters into a code-group stream. In the receive direction, the PCS decodes the code-group stream received from the PMA, maps the code-groups to XGMII data and control characters and forwards the character stream to the XGMII to the PCS client for further processing.

48.2.2 Functions within the PCS

The PCS includes the Transmit, Receive, Synchronization, and Deskew processes for 10GBASE-X. The PCS shields the RS (and MAC) from the specific nature of the underlying channel.

When communicating with the XGMII, the PCS uses, in each direction, 32 data signals (TXD <31:0> and RXD <31:0>), four control signals (TXC <3:0> and RXC <3:0>), and a clock (TX_CLK and RX_CLK).

When communicating with the PMA, the PCS uses the data signals tx_code-group <39:0> in the transmit direction and rx_unaligned <39:0> in the receive direction. Each set of data signals conveys four lanes of 10-bit code-groups. At the PMA service interface, code-group alignment, lane-to-lane deskew, and provision for PHY clock rate compensation are made possible by embedding special non-data code-groups in the idle stream. The PCS provides the functions necessary to map packets between the XGMII format and the PMA service interface format.

The tx_code-group and rx_unaligned signals are organized into four lanes in a manner similar to that of the XGMII. On transmit, the first PCS code-group is aligned to lane 0, the second to lane 1, the third to lane 2, the fourth to lane 3, then repeating with the fifth to lane 0, etc. This lane-oriented organization extends through the PMA to the PMD service interface. (See Table 48–1.)

The PCS Transmit process continuously generates code-groups based upon the TXD <31:0> and TXC <3:0> signals on the XGMII, sending them to the PMA service interface via the PMA_UNITDATA.request primitive.

The PCS Synchronization process continuously accepts unaligned and unsynchronized code-groups via the PMA_UNITDATA.indication primitive, obtains 10-bit code-group synchronization, and conveys synchronized 10-bit code-groups to the PCS Deskew process via the SYNC_UNITDATA.indicate message. The PCS Synchronization process sets the lane_sync_status <3:0> flags to indicate whether the PMA is functioning dependably (as well as can be determined without exhaustive error-ratio analysis).

Table 48–1—Transmit and receive lane associations

Lane	XGMII TXD RXD	XGMII TXC RXC	PMA tx_code-group rx_unaligned	PMD tx_lane rx_lane
0	<7:0>	<0>	<9:0>	<0>
1	<15:8>	<1>	<19:10>	<1>
2	<23:16>	<2>	<29:20>	<2>
3	<31:24>	<3>	<39:30>	<3>

The PCS Deskew process continuously accepts synchronized code-groups via the SYNC_UNITDATA.indicate message, aligns the code-groups to remove skew between the lanes that has been introduced by the link, and conveys aligned and synchronized code-groups to the PCS Receive process via the ALIGN_UNITDATA.indicate message. The PCS Deskew process asserts the align_status flag to indicate that the PCS has successfully deskewed and aligned code-groups on all PCS lanes. The PCS Deskew process attempts deskew and alignment whenever the align_status flag is de-asserted. The PCS Deskew process is otherwise idle.

The PCS Receive process continuously accepts code-groups from the PMA service interface via the ALIGN_UNITDATA.indicate message. The PCS Receive process monitors these code-groups and generates RXD and RXC on the XGMII. All code-groups received that represent idle are replaced with Idle characters prior to forwarding to the XGMII.

The PCS Transmit and Receive processes provide support for Link status reporting, which supports the transmit fault and receive fault conditions.

All PCS processes are described in detail in the state diagrams in 48.2.6.2.

48.2.3 Use of code-groups

The transmission code used by the PCS, referred to as 8B/10B, is identical to that specified in Clause 36. The PCS maps XGMII characters into 10-bit code-groups, and vice versa, using the 8B/10B block coding scheme. Implicit in the definition of a code-group is an establishment of code-group boundaries by a PCS Synchronization process. The 8B/10B transmission code as well as the rules by which the PCS ENCODE and DECODE functions generate, manipulate, and interpret code-groups are specified in 36.2.4. A 10GBASE-X PCS shall meet the requirements specified in 36.2.4.1 through 36.2.4.6, 36.2.4.8, and 36.2.4.9. PCS lanes are independent of one another. All code-group rules specified in 36.2.4 are applicable to each lane. The mapping of XGMII characters to PCS code-groups is specified in Table 48–2. The mapping of PCS code-groups to XGMII characters is specified in Table 48–3. PHYs that support EEE are able to transmit and receive LPI characters.

Figure 48–3 illustrates the mapping of an example XGMII character stream into a PCS code-group stream. Figure 48–4 illustrates the mapping during LPI.

The relationship of code-group bit positions to XGMII, PCS and PMA constructs and PMD bit transmission order, exemplified for lane 0, is illustrated in Figure 48–5.

XGMII																		
T/RXD<7:0>	I	I	S	Dp	D	D	D	---	D	D	D	D	I	I	I	I	I	I
T/RXD<15:8>	I	I	Dp	Dp	D	D	D	---	D	D	D	T	I	I	I	I	I	I
T/RXD<23:16>	I	I	Dp	Dp	D	D	D	---	D	D	D	I	I	I	I	I	I	I
T/RXD<31:24>	I	I	Dp	Ds	D	D	D	---	D	D	D	I	I	I	I	I	I	I

PCS																		
LANE 0	K	R	S	Dp	D	D	D	---	D	D	D	D	A	R	R	K	K	R
LANE 1	K	R	Dp	Dp	D	D	D	---	D	D	D	T	A	R	R	K	K	R
LANE 2	K	R	Dp	Dp	D	D	D	---	D	D	D	K	A	R	R	K	K	R
LANE 3	K	R	Dp	Ds	D	D	D	---	D	D	D	K	A	R	R	K	K	R

Legend:
Dp represents a data character containing the preamble pattern
Ds represents a data character containing the SFD pattern

Figure 48–3—XGMII character stream to PCS code-group stream mapping example

XGMII	LPI sleep/quiet/refresh											wake time						
T/RXD<7:0>	I	I	LI	LI	LI	LI	LI	---	LI	LI	LI	I	I	---	I	S	Dp	D
T/RXD<15:8>	I	I	LI	LI	LI	LI	LI	---	LI	LI	LI	I	I	---	I	Dp	Dp	D
T/RXD<23:16>	I	I	LI	LI	LI	LI	LI	---	LI	LI	LI	I	I	---	I	Dp	Dp	D
T/RXD<31:24>	I	I	LI	LI	LI	LI	LI	---	LI	LI	LI	I	I	---	I	Dp	Ds	D
	active time														active time			

PCS																		
LANE 0	K	R	R	K	R	A	K	---	LL	R	K	K	A	---	R	S	Dp	D
LANE 1	K	R	LL	K	R	A	K	---	R	R	K	K	A	---	R	Dp	Dp	D
LANE 2	K	R	R	K	LL	A	LL	---	R	R	LL	K	A	---	R	Dp	Dp	D
LANE 3	K	R	R	LL	R	A	K	---	R	LL	K	K	A	---	R	Dp	Ds	D

Legend:
LI represents the data character containing the XGMII LPI pattern (06)
LL represents the LPI indication code-group /D20.5/
Dp represents a data character containing the preamble pattern
Ds represents a data character containing the SFD pattern

Figure 48–4—XGMII and PCS mapping example with optional LPI

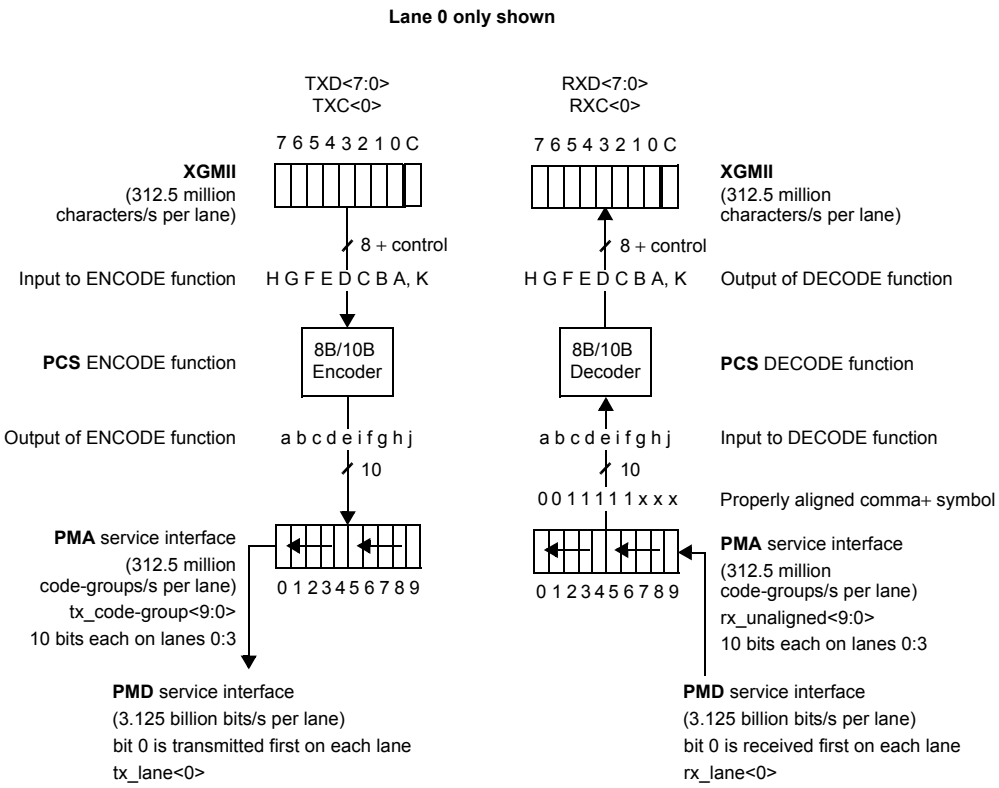

Figure 48–5—PCS reference diagram

48.2.4 Ordered sets and special code-groups

10GBASE-X PCS ordered sets consist of combinations of special and data code-groups (defined as a column of code-groups). All ordered sets are four code-groups in length and begin in lane 0. In addition to ordered sets, the PCS defines several special code-groups for control purposes. Ordered sets and special code-groups provide the following capabilities:

a) PCS Synchronization process ability to obtain bit and code-group synchronization
b) Packet delineation
c) Synchronization between the transmitter and receiver circuits at opposite ends of a link
d) Deskew of received code-groups from all serial lanes to an alignment pattern
e) Clock rate compensation protocol
f) Link status reporting protocol
g) Column fill
h) Error indication

Table 48–4 lists the defined ordered sets and special code-groups.

48.2.4.1 Data (/D/)

A data code-group, when not used to distinguish or convey information for a defined ordered set, conveys one octet of data between the XGMII and the PCS. Any data code-group can be followed by any other data code-group. Data code-groups are encoded and decoded but not interpreted by the PCS.

Table 48–2—XGMII character to PCS code-group mapping

XGMII TXC	XGMII TXD	PCS code-group	Description
0	00 through FF	Dxx.y	Normal data transmission
1	06	K28.0 or K28.3 or K28.5 or D20.5[a]	Assert LPI
1	07	K28.0 or K28.3 or K28.5	Idle in ||I||
1	07	K28.5	Idle in ||T||
1	9C	K28.4	Sequence
1	FB	K27.7	Start
1	FD	K29.7	Terminate
1	FE	K30.7	Error
1	Other value in Table 36–2	See Table 36–2	Reserved XGMII character
1	Any other value	K30.7	Invalid XGMII character
NOTE—Values in TXD column are in hexadecimal.			

[a]Insertion of /D20.5/ is per the rules described in 48.2.4.2.

Table 48–3—PCS code-group to XGMII character mapping

XGMII RXC	XGMII RXD	PCS code-group	Description
0	00 through FF	Dxx.y	Normal data reception
1	06	K28.0 or K28.3 or K28.5 or D20.5[a]	Assert LPI
1	07	K28.0 or K28.3 or K28.5	Idle in ||I||
1	07	K28.5	Idle in ||T||
1	9C	K28.4	Sequence
1	FB	K27.7	Start
1	FD	K29.7	Terminate
1	FE	K30.7	Error
1	FE	Invalid code-group	Received code-group
1	See Table 36–2	Other valid code-group	Received reserved code-group
NOTE—Values in RXD column are in hexadecimal.			

[a]Insertion of /D20.5/ is per the rules described in 48.2.4.2.

Table 48–4—Defined ordered sets and special code-groups

Code	Ordered set	Number of code-groups	Encoding
||I||	**Idle**		**Substitute for XGMII Idle**
||K||	Sync column	4	/K28.5/K28.5/K28.5/K28.5/
||R||	Skip column	4	/K28.0/K28.0/K28.0/K28.0/
||A||	Align column	4	/K28.3/K28.3/K28.3/K28.3/
	Encapsulation		
||S||	Start column	4	/K27.7/Dx.y/Dx.y/Dx.y/[a]
||T||	Terminate column	4	Terminate code-group in any lane
$\|\|T_0\|\|$	Terminate in Lane 0	4	/K29.7/K28.5/K28.5/K28.5/
$\|\|T_1\|\|$	Terminate in Lane 1	4	/Dx.y/K29.7/K28.5/K28.5/[a]
$\|\|T_2\|\|$	Terminate in Lane 2	4	/Dx.y/Dx.y/K29.7/K28.5/[a]
$\|\|T_3\|\|$	Terminate in Lane 3	4	/Dx.y/Dx.y/Dx.y/K29.7/[a]
	Control		
/E/	Error code-group	1	/K30.7/
	Link Status		
||Q||	Sequence ordered set	4	/K28.4/Dx.y/Dx.y/Dx.y/[a]
||LF||	Local Fault signal	4	/K28.4/D0.0/D0.0/D1.0/
||RF||	Remote Fault signal	4	/K28.4/D0.0/D0.0/D2.0/
||LINT||	Link Interruption signal	4	/K28.4/D0.0/D0.0/D3.0/
||Qrsvd||	Reserved	4	!||LF|| and !||RF|| and !||LINT||
	Reserved		
||Fsig||	Signal ordered set	4	/K28.2/Dx.y/Dx.y/Dx.y/[a,b]

[a]/Dx.y/ indicates any data code-group.
[b]Reserved for INCITS T11.

48.2.4.2 Idle (||I||) and Low Power Idle (||LPIDLE||)

Idle ordered sets (||I||) are transmitted in full columns continuously and repetitively whenever the XGMII is idle (TXD <31:0>=0x07070707 and TXC <3:0>=0xF). ||I|| provides a continuous fill pattern to establish and maintain lane synchronization, perform lane-to-lane deskew and perform PHY clock rate compensation. ||I|| is emitted from, and interpreted by, the PCS.

A sequence of ||I|| ordered sets consists of one or more consecutively transmitted ||K||, ||R||, or ||A|| ordered sets, as defined in Table 48–4. Rules for ||I|| ordered set sequencing shall be as follows:

a) ||I|| sequencing starts with the first column following a ||T||.
b) The first ||I|| following ||T|| alternates between ||A|| or ||K|| except if an ||A|| is to be sent and less than *r* [see item d)] columns have been sent since the last ||A||, a ||K|| is sent instead.

c) ||R|| is chosen as the second ||I|| following ||T||.

d) Each ||A|| is sent after *r* non-||A|| columns where *r* is a randomly distributed number between 16 and 31, inclusive. The corresponding minimum spacing of 16 non-||A|| columns between two ||A|| columns provides a theoretical 85-bit deskew capability.

e) When not sending an ||A||, either ||K|| or ||R|| is sent with a random uniform distribution between the two.

f) Whenever sync_status=OK, all ||I|| received during idle are translated to XGMII Idle control characters for transmission over the XGMII. All other !||I|| received during idle are mapped directly to XGMII data or control characters on a lane by lane basis, with the following exceptions for PHYs with EEE capability:
   1) /D20.5/ (LPI) being detected in any lane and the rest of the lanes in the same column being detected /K/ only or /R/ only, which will result in reporting LP_IDLE characters in all lanes.
   2) ||A|| being detected and /D20.5/ (LPI) being detected in any lane of the previous column and the rest of the lanes in the previous column being detected /K/ only or /R/ only, which will result in reporting LP_IDLE characters in all lanes.

The purpose of randomizing the ||I|| sequence is to reduce 10GBASE-X electromagnetic interference (EMI) during idle. The randomized ||I|| sequence produces no discrete spectrum. Both ||A|| spacing as well as ||K||, ||R||, or ||A|| selection shall be based on the generation of a random integer *r* generated by a PRBS based on one of the 7th order polynomials listed in Figure 48–6. ||A|| spacing is set to the next generated value of *r*. The rate of generation of *r* is once per column, 312.5 MHz ± 100 ppm. Once the ||A|| spacing count goes to zero (A_CNT=0), ||A|| is selected for transmission at the next opportunity during the Idle sequence. ||K|| and ||R|| selection follows the value of code_sel, which is continuously set according to the even or odd value of *r*. The method of generating the random integer *r* is left to the implementer. PCS Idle randomizer logic is illustrated in Figure 48–6.

||LPIDLE|| is coded in the same manner as ||I|| except that the /D20.5/ code-group replaces one code-group in each ||K|| or ||R|| (not ||A||) column with a random uniform distribution across the lanes. Insertion of /D20.5/ does not alter the distribution of ||A||, ||K|| or ||R||. Clock compensation may be performed during LPI according to the rules described in 48.2.4.2.3.

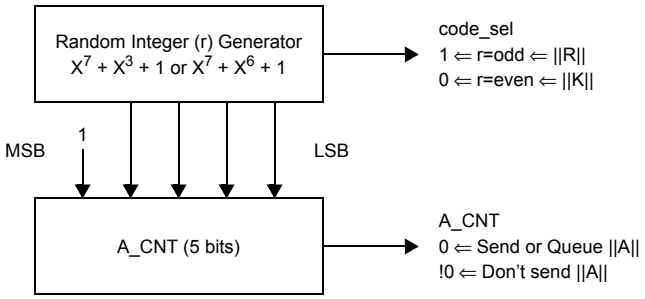

Figure 48–6—PCS idle randomizer

48.2.4.2.1 Sync ||K||

Code-group synchronization is the process by which the receiver detects code-group boundaries in the incoming bit stream of each lane. The detection of the comma pattern in the incoming bit stream identifies a code-group boundary. The proper alignment of a comma used for code-group synchronization is depicted in Figure 48–5. The Sync or ||K|| ordered set included in the PCS Idle sequence guarantees a sufficient frequency of commas in each lane. The comma pattern is defined in 36.2.4.9. The Sync ordered set is

defined in Table 48–4. /K/ code-groups are interpreted per lane by the PCS Synchronization process, which is described in 48.2.6.2.2. Detection of both comma+ and comma- variants of /K/ shall be required.

48.2.4.2.2 Align ||A||

Skew is introduced between lanes by both active and passive elements of a 10GBASE-X link. The PCS deskew function compensates for all lane-to-lane skew observed at the receiver. The Align or ||A|| ordered set consists of a unique special code-group, also known as Align or /A/ in each lane. /A/ is not used in any other ordered set. The definition of a 10GBASE-X ordered set guarantees that /A/ code-groups are simultaneously initiated on all lanes at the transmitter, resulting in minimal lane-to-lane skew at the transmitter. The Align ordered set is defined in Table 48–4. Allowable skew for all link elements shall be as specified in Table 48–5.

Table 48–5—Skew budget

Skew Source	Occurences	Skew	Total Skew
PMA Tx	1	1 UI[a]	1 UI
PCB	2	1 UI	2 UI
Medium	1	<18 UI	<18 UI
PMA Rx[b]	1	20 UI	20 UI
Total			<41 UI

[a]UI represents unit interval. For 10GBASE-X, 1 UI = 320 ps.
[b]Includes deserialization function, physical deserializer skew and clock boundary transition.

48.2.4.2.3 Skip ||R||

The 10GBASE-X PHY allows for multiple clock domains along a single link. The Skip or ||R|| ordered set is included in the PCS Idle sequence to allow for clock rate compensation for the case of multiple clock domains. Clock rate compensation may be performed via insertion or removal of either Idle characters in the unencoded data stream or ||R|| in the encoded Idle stream. Any ||R|| may be removed. ||R|| may be inserted anywhere in the Idle stream with the exception of the first column following ||T||.

When clock compensation is done in the unencoded data stream, rules for idle insertion and deletion shall be as follows:

a) Idle insertion or deletion occurs in groups of four consecutive Idle characters.
b) Idle characters are added following idle or ordered sets.
c) Idle characters are not added while data is being received.
d) The four characters following a Terminate control character are not deleted.
e) Sequence ordered sets are deleted to adapt between clock rates.
f) Sequence ordered set deletion occurs only when two consecutive sequence ordered sets have been received and deletes only one of the two.
g) Only idles are inserted for clock compensation.

The disparity of the /R/ code-group is neutral, allowing its removal or insertion without affecting the current running disparity of the serial stream. The correct current running disparity version of /R/ must be inserted in each lane during Skip insertion (see Table 36–2). The Skip ordered set is defined in Table 48–4.

For EEE capability, a column containing three /R/ code-groups and one /D20.5/ code-group may be inserted or deleted in the same manner as four /R/ code-groups.

48.2.4.3 Encapsulation

The Start and Terminate ordered sets correspond to columns containing the XGMII Start and Terminate control characters, respectively.

48.2.4.3.1 Start ||S||

The Start or ||S|| ordered set directly maps to the XGMII Start control character in lane 0 followed by any three data characters in XGMII lanes 1 through 3. Normally, the three data characters will be the preamble pattern, but the PCS neither checks nor alters their contents. ||S|| indicates to the PCS that a packet has been initiated. The Start ordered set is defined in Table 48–4.

48.2.4.3.2 Terminate ||T||

The Terminate or ||T|| ordered set directly maps to the XGMII Terminate control character located in any lane, preceded by data characters if Terminate is not in lane 0, and followed by Idle characters if Terminate is not in lane 3. ||T|| indicates to the PCS that a packet has been terminated. All XGMII control characters following ||T|| are translated by the PCS until the recognition of the next XGMII Start control character. The PCS considers the MAC interpacket gap (IPG) to have begun with the reception of ||T||. The Terminate ordered set is defined in Table 48–4.

Unrecognized running disparity errors which propagate to any Idle code-groups in ||T|| or to the column following ||T|| are indicated as /E/ in the preceding column in the same lane in which the errors were recognized. All ||I|| ordered sets are selected to ensure that propagated code violations are recognized and not propagated further.

The cvtx_terminate function is used to convert all XGMII Idle control characters in the same column as the XGMII Terminate control character to /K/ code-groups. The cvrx_terminate function is used to convert all /K/ code-groups in ||T|| to XGMII Idle control characters.

48.2.4.4 Error /E/

The Error code-group is directly mapped to the XGMII Error control character. /E/ may also be generated by the PCS client to indicate a transmission error to its peer entity or deliberately corrupt the contents of the frame in such a manner that a receiver will detect the corruption with the highest degree of probability. Error is signaled per lane since code-violations are detected on a per lane basis. The Error code-group is defined in Table 48–4.

The presence of /E/ or any invalid code-group on the medium denotes an error condition. 10GBASE-X elements that detect code-group violations shall replace the invalid code-group with /E/ prior to retransmission.

48.2.4.5 Link status

Link status reporting uses the Sequence ordered-set to transport the transmit fault and receive fault link status conditions.

48.2.4.5.1 Sequence ||Q||

The Sequence or ||Q|| ordered set directly maps to the XGMII Sequence control character on lane 0 followed by three data characters in XGMII lanes 1 through 3. ||Q|| indicates to the PCS that a link status message has been initiated. The PCS Receive process may also initiate Sequence ordered-sets upon detection of a link status condition. Sequence ordered-sets are always sent over the PMA service interface in the column that follows an ||A|| ordered-set. The Sequence ordered-sets do not otherwise interfere with the randomized ||I|| sequence. Sequence ordered-sets corresponding to Local Fault signal and Remote Fault signal are specified in Table 48–4.

48.2.5 Management function requirements

The 10GBASE-X PCS supports a set of required and optional management objects to permit it to be controlled by the Station Management entity (STA). Access to management objects within the 10GBASE-X sublayer is accomplished by means of a set of registers within the MDIO register space as defined in 45.2.4 and 45.2.5. The details of the register bit allocations and general usage are given in Clause 45. Table 48–6, Table 48–7, and Table 48–8 describe how the PCS state diagram variables map to management register bits. If an MDIO interface is provided (see Clause 45), they are accessed via that interface. If not, it is recommended that an equivalent access be provided.

Table 48–6—State diagram variable to management register mapping for PCS

State diagram variable	Management register bit
reset	3.0.15 Reset
sync_status	3.24.12 10GBASE-X lane alignment status
lane_sync_status<3>	3.24.3 Lane 3 sync
lane_sync_status<2>	3.24.2 Lane 2 sync
lane_sync_status<1>	3.24.1 Lane 1 sync
lane_sync_status<0>	3.24.0 Lane 0 sync

Table 48–7—State diagram variable to management register mapping for PHY XS

State diagram variable	Management register bit
reset	4.0.15 Reset
align_status	4.8.10 Receive fault
sync_status	4.24.12 PHY XGXS lane alignment status
lane_sync_status<3>	4.24.3 Lane 3 sync
lane_sync_status<2>	4.24.2 Lane 2 sync
lane_sync_status<1>	4.24.1 Lane 1 sync
lane_sync_status<0>	4.24.0 Lane 0 sync

Table 48–8—State diagram variable to management register mapping for DTE XS

State diagram variable	Management register bit
reset	5.0.15 Reset
align_status	5.8.10 Receive fault
sync_status	5.24.12 PHY XGXS lane alignment status
lane_sync_status<3>	5.24.3 Lane 3 sync
lane_sync_status<2>	5.24.2 Lane 2 sync
lane_sync_status<1>	5.24.1 Lane 1 sync
lane_sync_status<0>	5.24.0 Lane 0 sync

48.2.6 Detailed functions and state diagrams

The body of this clause is comprised of state diagrams, including the associated definitions of variables, constants, and functions. Should there be a discrepancy between a state diagram and descriptive text, the state diagram prevails.

The notation used in the state diagrams in this clause follows the conventions in 21.5. State diagram variables follow the conventions of 21.5.2 except when the variable has a default value. Variables in a state diagram with default values evaluate to the variable default in each state where the variable value is not explicitly set.

Timeless states are employed as an editorial convenience to facilitate the distribution of transition conditions from prior states. No actions are taken within these states. Exit conditions are evaluated for timeless states. There is one timeless state. It is PCS Receive state RECEIVE.

48.2.6.1 State variables

48.2.6.1.1 Notation conventions

/x/

Denotes the constant code-group specified in 48.2.6.1.2 (valid code-groups must follow the rules of running disparity as per 36.2.4.5 and 36.2.4.6).

[/x/]

Denotes the latched received value of the constant code-group (/x/) specified in 48.2.6.1.2 and conveyed by the SYNC_UNITDATA.indicate message described in 48.2.6.1.7.

||y||

Denotes the column of constant code-groups in lanes 0 through 3, inclusively, specified in 48.2.6.1.2 (valid code-groups must follow the rules of running disparity as per 36.2.4.5 and 36.2.4.6).

[||y||]

Denotes the latched received value of the column of constant code-groups in lanes 0 through 3, inclusively (||y||), specified in 48.2.6.1.2 and conveyed by the SYNC_UNITDATA.indicate message described in 48.2.6.1.7 or the ALIGN_UNITDATA.indicate message described in 48.2.6.1.7.

48.2.6.1.2 Constants

/A/

The /K28.3/ code-group used in the Idle Align function specified in 48.2.4.2.2.

||A||

The column of four identical Idle Align code-groups corresponding to the Idle Align function specified in 48.2.4.2.2.

/COMMA/

The set of special code-groups which include a comma as specified in 36.2.4.9 and listed in Table 36–2.

/D/

The set of 256 code-groups corresponding to valid data, as specified in 48.2.4.1 and listed in Table 36–1a.

||D||

The column of four Data code-groups present during packet reception.

/Dx.y/

One of the set of 256 code-groups corresponding to valid data, as specified in 48.2.4.1 and listed in Table 36-1.

/E/

The /K30.7/ code-group corresponding to the Error function specified in 48.2.4.4.

||I||

The column of four identical Idle code-groups corresponding to the Idle function specified in 48.2.4.2. Also used to represent the corresponding XGMII control characters.

||IDLE||

Alias for ||I||.

/INVALID/

The set of invalid data or special code-groups, as specified in 36.2.4.6.

/K/

The /K28.5/ code-group used in the Idle Sync function specified in 48.2.4.2.1. Also used in the Terminate function specified in 48.2.4.3.2.

||K||

The column of four identical Idle Sync code-groups corresponding to the Idle Sync function specified in 48.2.4.2.1. For EEE capability, one lane of ||K|| is replaced by /D20.5/ during the assertion of LPI as defined in 48.2.4.2.

/Kx.y/

One of the set of 12 code-groups corresponding to valid special code-groups, as specified in Table 36–2.

LFAULT

A vector of bits RXD<31:0> and RXC<3:0> containing a Local Fault sequence ordered set. The Local Fault sequence ordered set is defined in 46.3.4.

||Q||

The column of four code-groups corresponding to the Sequence function specified in 48.2.4.5.1. Also used to represent the corresponding set of XGMII control and data characters.

/R/

The /K28.0/ code-group used for the Idle Skip function specified in 48.2.4.2.3.

||R||

The column of four identical Idle Skip code-groups corresponding to the Idle Skip function specified in 48.2.4.2.3. For EEE capability, one lane of ||R|| is replaced by /D20.5/ during the assertion of LPI as defined in 48.2.4.2.

||S||

The column of code-groups including the Start code-group as specified in 48.2.4.3.1. Also used to represent the corresponding set of XGMII control and data characters.

||T||

The column of code-groups including the Terminate code-group as specified in 48.2.4.3.2. Also used to represent the corresponding XGMII control and data characters.

The following constants are used only for the EEE capability:

||LPIDLE||

The column consisting of three /K/ characters and one of /D20.5/, or three /R/ characters and one /D20.5/, or a column of ||A|| preceded by a column containing three /K/ characters and one /D20.5/ or three /R/ characters and one /D20.5/ as specified in 48.2.4.2.

48.2.6.1.3 Variables

align_column <39:0>

A vector of bits represented by the most recently received column of aligned 10-bit code-groups on all four lanes from the PCS Deskew process. For lane 0, the element align_column <0> corresponds to the least recently received (oldest) rx_bit of the code-group; align_column <9> corresponds to the most recently received (newest) rx_bit of the code-group. The same bit aging is applicable to the bits in align_column <39:10> with respect to rx_lane <3:1>. Code-group to lane assignment is specified in Table 48–1.

align_status

A parameter set by the PCS Deskew process to reflect the status of the lane-to-lane code-group alignment.

Values: FAIL; The deskew process is not complete.
OK; All lanes are synchronized and aligned.

NOTE—For EEE capability, this variable is affected by the LPI receive state diagram. Without EEE capability this variable is identical to deskew_align_status controlled by the deskew state diagram.

code_sel

A Boolean derived from a uniformly distributed random integer *r* generated by a PRBS based on a 7th order polynomial.

Values: 0; LSB of random number is zero.
1; LSB of random number is one.

deskew_error

A Boolean used by the PCS Deskew process to indicate that a lane-to-lane alignment error has been detected.

Values: FALSE; /A/ not recognized in any lane or recognized in all lanes simultaneously.
TRUE; /A/ recognized in fewer than all lanes.

enable_cgalign

A Boolean that indicates the enabling and disabling of code-group comma alignment. The code-group boundary may be changed whenever code-group comma alignment is enabled. This process is known as code-group alignment.

Values: FALSE; Code-group alignment is disabled.
TRUE; Code-group alignment is enabled.

enable_deskew

A Boolean that indicates the enabling and disabling of the deskew process. Code-groups may be discarded whenever deskew is enabled. This process is known as code-group slipping.

Values: FALSE; Deskew is disabled.
TRUE; Deskew is enabled.

IDLE

A vector of bits RXD<31:0> and RXC<3:0> containing Idle. Idle is defined in Table 46–4.

lane_sync_status <3:0>

A parameter set by the PCS Synchronization process to reflect the status of the link for each lane as viewed by the receiver. lane_sync_status <n> represents lane_sync_status on lane n where n=0:3.

Values: FAIL; The receiver is not synchronized to the code-group boundary.
OK; The receiver is synchronized to the code-group boundary.

next_ifg

Controls the ||IDLE|| pattern immediately following the next frame. It is used to ensure an equal and deterministic presence of both ||A|| and ||K||

Values: A; The first ||IDLE|| following the end of the next frame will be ||A||.
K; The first ||IDLE|| following the end of the next frame will be ||K||.

reset

Condition that is true until such time as the power supply for the device that contains the PCS has reached the operating region. The condition is also true when the device has low-power mode set via Control register bit 4.0.11 or 5.0.11. The condition is also true when a reset request is detected via Control register bit 4.0.15 or 5.0.15.

Values: FALSE; The device is completely powered and has not been reset (default).
TRUE; The device has not been completely powered or has been reset.

NOTE—Reset evaluates to its default value in each state where it is not explicitly set.

RX

Alias for RXD <31:0> and RXC <3:0> representing the XGMII Receive Data and Control signals.

RXC <3:0>

Receive Control signals of the XGMII as specified in Clause 46. Set by the PCS Receive process.

RXD <31:0>

Receive Data signals of the XGMII as specified in Clause 46. Set by the PCS Receive process.

rx_lane <3:0>

A vector of bits representing the serial lanes used to convey data from the PMD to the PMA via the PMD_UNITDATA.indication service primitive as specified in Clause 53 or Clause 54.

Bit values: ZERO; Data bit is a logical zero.
ONE; Data bit is a logical one.

rx_unaligned <39:0>

A vector of bits represented by the most recently received column of unaligned 10-bit code-groups on all four lanes from the PMA. For lane 0, the element rx_unaligned <0> is the least recently received (oldest) rx_bit from the PMD; rx_unaligned <9> is the most recently received (newest) rx_bit from the PMD. The same bit aging is applicable to the bits in rx_unaligned <39:10> with respect to rx_lane <3:1>. Code-group to lane assignment is specified in Table 48–1.

signal_detect <3:0>

A parameter set continuously and directly from the PMD_SIGNAL.indication(signal_detect <3:0>) primitive to reflect the status of the incoming link signal for each lane. Used by the PCS Synchronization process to validate the data received on rx_unaligned <39:0>. Signal_detect <n> represents signal_detect on lane n where n=0:3.

Values: FAIL; A signal is not present on the lane.
OK; A signal is present on the lane.

sync_code-group <39:0>

A vector of bits represented by the most recently received column of unaligned 10-bit code-groups on all four lanes from the PCS Synchronization process. For lane 0, the element sync_code-group <0> corresponds to the least recently received (oldest) rx_bit; sync_code-group <9> corresponds to the most recently received (newest) rx_bit. The same bit aging is applicable to the bits in sync_code-group<39:10> with respect to rx_lane <3:1>. Code-group to lane assignment is specified in Table 48–1.

sync_status

A Boolean that represents the following behavior: For all n in lane_sync_status<n>.

Values: FAIL; At least one lane is not in sync.
OK; All lanes are in sync.

TQMSG

A vector of bits representing the last link status message received over the XGMII once the link status message is recognized. Used by the PCS Transmit process to load tx_code-group <39:0>.

tx_code-group <39:0>

A vector of bits representing a column of four aligned 10-bit code-groups which has been prepared for transmission by the PCS Transmit process. This vector is conveyed to the PMA as the parameter of a PMA_UNITDATA.request(tx_lane <3:0>) service primitive. For lane 0, the element tx_code-group <0> is the first tx_bit transmitted; tx_code-group <9> is the last tx_bit transmitted. The same bit aging is applicable to the bits in tx_code-group <39:10> with respect to tx_lane <3:1>. Code-group to lane assignment is specified in Table 48–1.

TX

Alias for either TXD <31:0> and TXC <3:0> representing the XGMII Transmit Data and Control signals, or the Local Fault ordered set as defined in 46.3.4 when a fault condition is detected on the transmit path.

TXC <3:0>

Transmit Control signals of the XGMII as specified in Clause 46. Interpreted by the PCS Transmit process.

TXD <31:0>

Transmit Data signals of the XGMII as specified in Clause 46. Interpreted by the PCS Transmit process.

tx_lane <3:0>

A vector of bits representing the serial lanes used to convey data from the PMA to the PMD via the PMD_UNITDATA.request service primitive as specified in Clause 53 or Clause 54.

Bit values: ZERO; Data bit is a logical zero.
ONE; Data bit is a logical one.

The following variables are used only for the EEE capability:

deskew_align_status

Variable used by the deskew state diagram to reflect the status of the lane-to-lane code-group alignment.

Values: FAIL; The deskew process is not complete.
OK; All lanes are synchronized and aligned.

rx_lpi_active

A Boolean variable that is set to TRUE when the receiver is in a low power state and set to FALSE when it is in an active state and is not restricted by the LPI receive state diagram.

rx_quiet

A Boolean variable set to TRUE while in the RX_QUIET state and set to FALSE otherwise. When this variable is TRUE it indicates that receive PCS and PMD may power-down nonessential functions.

tx_quiet

A Boolean variable set to TRUE when the transmitter is in the TX_QUIET state and set to FALSE otherwise. When set to TRUE, the PMD will disable the transmitter as described in 71.6.6. When this variable is TRUE it indicates that transmit PCS and PMD may power-down nonessential functions.

48.2.6.1.4 Functions

check_end

Prescient Terminate function used by the PCS Receive process to set the RXD<31:0> and RXC<3:0> signals to indicate Error if a running disparity error was propagated to any Idle code-groups in ||T||, or to the column following ||T||. The XGMII Error control character is returned in all lanes less than n in ||T||, where n identifies the specific Terminate ordered-set $||T_n||$, for which a running disparity error or any code-groups other than /A/ or /K/ are recognized in the column following ||T||. The XGMII Error control character is also returned in all lanes greater than n in the column prior to ||T||, where n identifies the specific Terminate ordered-set $||T_n||$, for which a running disparity error or any code group other than /K/ is recognized in the corresponding lane of ||T||. For all other lanes the value set previously is retained.

cvrx_terminate

Conversion function used by the PCS Receive process when Terminate is indicated to convert all /K/ code-groups to Idle control characters signaled via RX. Conversion is performed for all lanes.

cvtx_terminate

Conversion function used by the PCS Transmit process when Terminate is indicated to convert all Idle control characters signaled via TX to /K/ code-groups. Conversion is performed for all lanes.

DECODE ([||y||])

Consists of four independent synchronous processes, one each per lane. In the PCS Receive process, this function takes as its arguments the latched value of align_column ([||y||]) and the current running disparity, and returns XGMII RX as specified in 48.2.3 and 48.2.4. When decoding

||T||, the returned XGMII RX value is further modified by the cvrx_terminate and check_end functions, the result of the check_end function takes priority over the result of the cvrx_terminate function. DECODE also updates the current running disparity per the running disparity rules outlined in 36.2.4.4.

ENCODE(TX)
: Consists of four independent synchronous processes, one each per lane. In the PCS Transmit process, this function takes as its argument the XGMII TX signals and the current running disparity for each lane, and returns four corresponding 10-bit code-groups as specified in 48.2.3 and 48.2.4. When encoding ||T||, the XGMII TX values are modified by the result of the cvtx_terminate function. ENCODE also updates the current running disparity per Tables 36-1 or 36–2.

Q_det
: Function to determine the need to transmit sequence ordered sets. If TX=||Q|| then Q_det is set to true and TQMSG is set to the result of ENCODE(TX). Q_det remains true until set to false by the PCS transmit source state diagram. In the event that this function and the state diagram both attempt to modify Q_det, the setting of Q_det by this function to true will take priority.

signal_detectCHANGE <3:0>
: In the PCS Synchronization process, this function monitors the signal_detect parameter on a per lane basis for a state change. The function is set upon state change detection, which is required to detect signal_detect changes which occur asynchronously to PUDI. signal_detectCHANGE <n> represents signal_detectCHANGE on lane n where n=0:3.

  Values: TRUE; The output of this function changes to true when the function detects a change in signal_detect and stays true until the false condition is satisfied.
  FALSE; The output of this function changes to false when the LOSS_OF_SYNC state of the PCS synchronization state diagram is entered.

48.2.6.1.5 Counters

A_CNT
: A 5-bit down counter used to control Align code-group spacing. Loaded with a value between 16 and 31, inclusive, by a uniformly distributed random integer *r* generated by a PRBS based on a 7th order polynomial as described in 48.2.4.2. A_CNT is decremented once per PUDR. The count remains at zero until ||A|| is transmitted, at which time a new value is loaded.

good_cgs
: A 2-bit consecutive valid code-groups received counter.

The following counter is used only for the EEE capability:

wake_error_counter
: A counter that is incremented each time that the LPI receive state diagram enters the RX_WTF state indicating that a wake time fault has been detected. The counter is reflected in register 3.22 (see 45.2.3.10)

48.2.6.1.6 Timers

The following timers are used only for the EEE capability.

rx_tq_timer
: This timer is started when the PCS receiver enters the RX_SLEEP state. The timer terminal count is set to T_{QR}. When the timer reaches terminal count it will set the rx_tq_timer_done = TRUE.

rx_tw_timer
: This timer is started when the PCS receiver enters the RX_WAKE state. The timer terminal count shall be set to a value no larger than the maximum value given for T_{WR} in Table 48–10. When the timer reaches terminal count it will set the rx_tw_timer_done = TRUE.

rx_wf_timer
: This timer is started when the PCS receiver enters the RX_WTF state, indicating that the receiver has encountered a wake time fault. The rx_wf_timer allows the receiver an additional period in which to synchronize or return to the QUIET state before a link failure is indicated. The timer terminal count is set to T_{WTF}. When the timer reaches terminal count it will set the rx_wf_timer_done = TRUE.

tx_ts_timer
: This timer is started when the PCS transmitter enters the TX_SLEEP state. The timer terminal count is set to T_{SL}. When the timer reaches terminal count it will set the tx_ts_timer_done = TRUE.

tx_tq_timer
: This timer is started when the PCS transmitter enters the TX_QUIET state. The timer terminal count is set to T_{QL}. When the timer reaches terminal count it will set the tx_tq_timer_done = TRUE.

tx_tr_timer
: This timer is started when the PCS transmitter enters the TX_REFRESH state. The timer terminal count is set to T_{UL}. When the timer reaches terminal count it will set the tx_tr_timer_done = TRUE.

48.2.6.1.7 Messages

AUDI([||y||])
: Alias for ALIGN_UNITDATA.indicate(parameters).

ALIGN_UNITDATA.indicate([align_column <39:0>])
: A signal sent by the PCS Deskew process to the PCS Receive process conveying the latched value of the indicated column of code-groups over each lane ([||y||]) (see 48.2.6.2.3).

PMA_UNITDATA.indication(rx_unaligned <39:0>)
: A signal sent by the PMA Receive process to the PCS Synchronization process conveying the next code-group set received over each lane of the medium (see 48.3.2.2).

PMA_UNITDATA.request(tx_code-group <39:0>)
: A signal sent by the PCS Transmit process conveying the next code-group set for all lanes ready for transmission over the medium (see 48.3.2.1).

PMD_SIGNAL.indication(signal_detect <3:0>)
: Indicates the status of the incoming link signal. A signal mapped to the PMD_SIGNAL.indication(SIGNAL_DETECT) service primitive specified in Clause 53 or Clause 54. signal_detect <n> is set to the same value for all lanes n where n=0:3.

  Values: FAIL; A signal is not present on the lane.
  OK; A signal is present on the lane.

PUDI
: Alias for PMA_UNITDATA.indication(parameters).

PUDR
: Alias for PMA_UNITDATA.request(parameters).

SUDI
: Alias for SYNC_UNITDATA.indicate(parameters).

SYNC_UNITDATA.indicate(sync_code-group <39:0>)
: A signal sent by the PCS Synchronization process to the PCS Deskew process conveying code-groups over each lane (see 48.2.6.2.2).

The following messages are used only for the EEE capability:

PMD_RXQUIET.request(rx_quiet)
: A signal sent by the PCS LPI receive state diagram to the PMD. When TRUE this indicates that the receiver is in a quiet state and is not expecting incoming data.

PMD_TXQUIET.request(tx_quiet)
: A signal sent by the PCS LPI transmit state diagram to the PMD. When TRUE this indicates that the transmitter is in a quiet state and may cease to transmit a signal on the medium.

48.2.6.2 State diagrams

48.2.6.2.1 Transmit

The PCS shall implement its Transmit process as depicted in Figure 48–7, including compliance with the associated state variables as specified in 48.2.6.1. This state diagram makes exactly one transition for each transmitted ordered set that is processed.

The Transmit Source process determines whether XGMII data and control information should be passed through to the PMA for serialization or converted before passing to the PMA. In all cases, XGMII data and control information is encoded before it is passed to the PMA. Data and control information may be sourced either directly from the XGMII or generated by the PCS depending on whether packet or Idle is being sourced by the XGMII. The recognition of the XGMII Start, Terminate, Idle and Sequence control characters is used to determine whether packet, Idle, or link status is being sourced by the XGMII.

The detection of a link status condition including the receipt of link status messages over the XGMII causes the PCS Transmit process to generate link status messages interspersed in an Idle sequence.

The PCS Transmit process continuously sources tx_code-group<39:0> to the PMA. The Transmit process determines the proper code-group to source on each lane based on running disparity requirements.

48.2.6.2.2 Synchronization

The PCS shall implement four Synchronization processes as depicted in Figure 48–8 including compliance with the associated state variables as specified in 48.2.6.1. The Synchronization process is responsible for determining whether the underlying receive channel is ready for operation. Failure of the underlying channel typically causes the PMA client to suspend normal actions. A Synchronization process operates independently on each lane, and synchronization is complete only when synchronization is acquired on all lanes. The synchronization process described in the following paragraphs applies to each lane.

The PCS Synchronization process continuously accepts code-groups via the PMA_UNITDATA.indication primitive and conveys received code-groups to the PCS Deskew process via the SYNC_UNITDATA.indicate message.

When in the LOSS_OF_SYNC state, the PCS may attempt to realign its current code-group boundary to one which coincides with the code-group boundary defined by a comma (see 36.2.4.9). This process is referred to in this document as code-group alignment.

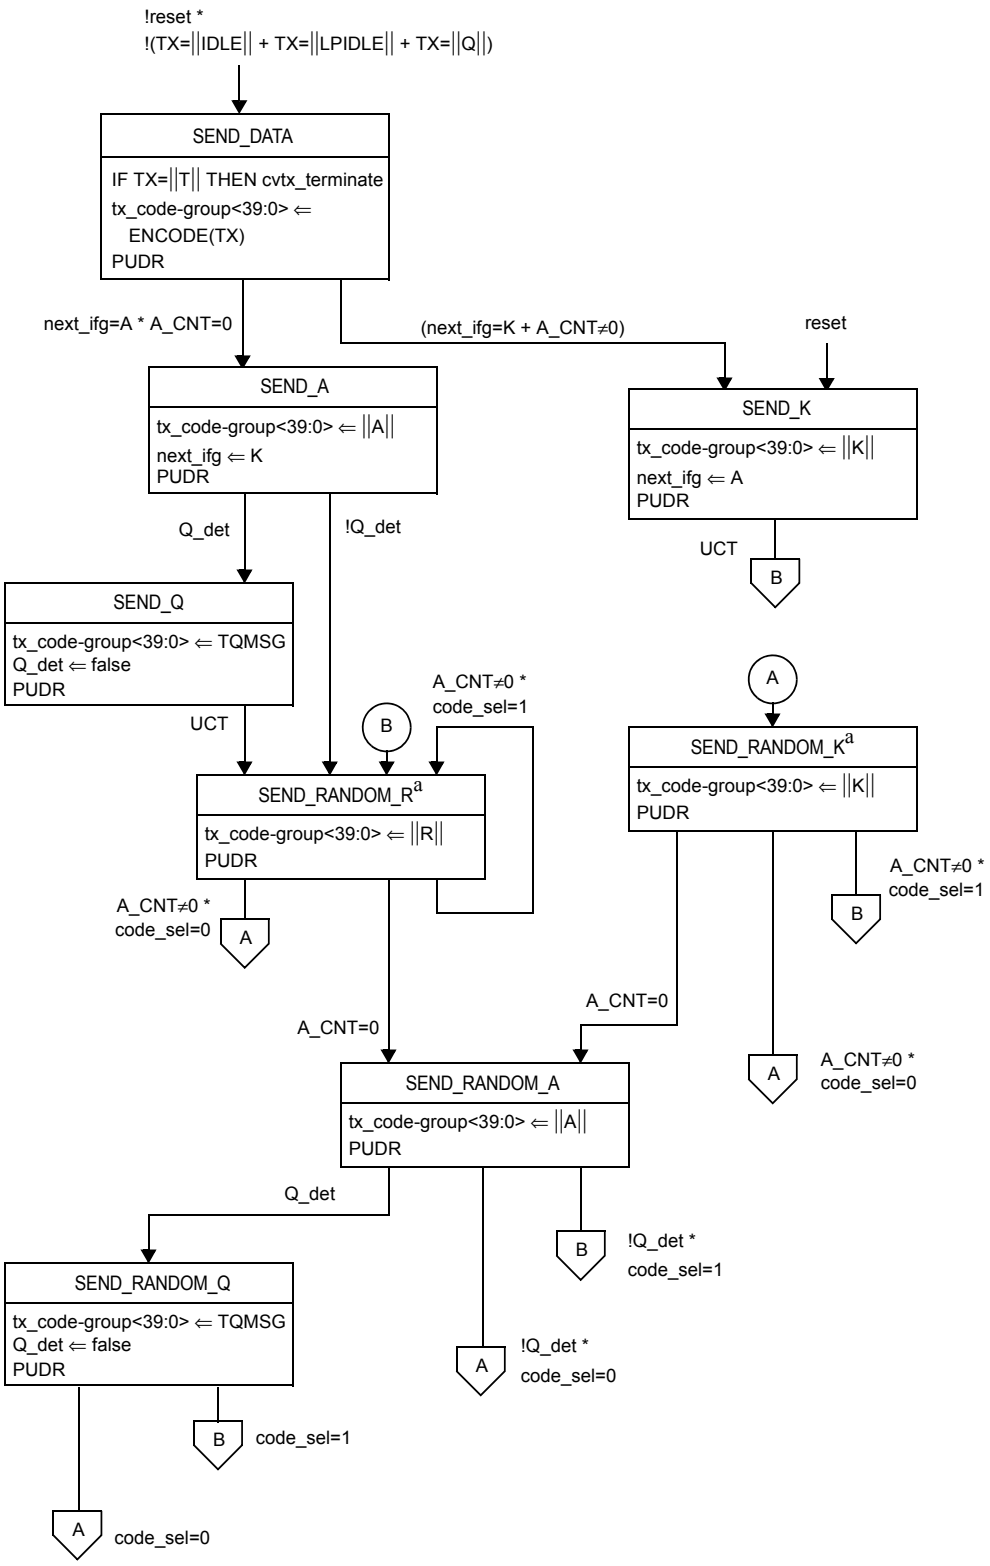

[a] If TX=||LPIDLE|| one lane is replaced by /D20.5/ as defined in 48.2.4.2.

NOTE—The state diagram makes exactly one transition for each transmit code-group processed.

Figure 48–7—PCS transmit source state diagram

Once synchronization is acquired, the Synchronization process tests received code-groups in sets of four code-groups and employs multiple sub-states, effecting hysteresis, to move between the SYNC_ACQUIRED_1 and LOSS_OF_SYNC states. The Synchronization process sets the lane_sync_status <3:0> flags to indicate whether the PMA is functioning dependably (as well as can be determined without exhaustive error-ratio analysis). Whenever any PMA lane is not operating dependably, as indicated by the setting of lane_sync_status <3:0>, the deskew_align_status flag is set to FAIL.

48.2.6.2.3 Deskew

The PCS shall implement the Deskew process as depicted in Figure 48–9 including compliance with the associated state variables as specified in 48.2.6.1. The Deskew process is responsible for determining whether the underlying receive channel is capable of presenting coherent data to the XGMII. The Deskew process asserts the deskew_align_status flag to indicate that the PCS has successfully deskewed and aligned code-groups on all lanes. The Deskew process attempts deskew and alignment whenever the deskew_align_status flag is de-asserted. The Deskew process is otherwise idle. For the EEE capability the relationship between align_status and deskew_align_status is given by Figure 48–12, otherwise align_status is identical to deskew_align_status. Whenever the align_status flag is set to FAIL the condition is indicated as a link_status=FAIL condition in the status register bit 4.1.2 or 5.1.2.

Once alignment is acquired, the Deskew process tests received columns and employs multiple sub-states, effecting hysteresis, to move between the ALIGN_ACQUIRED_1 and LOSS_OF_ALIGNMENT states. These states monitor the link for continued alignment, tolerate alignment inconsistencies due to a reasonably low BER, and restart the Deskew process if alignment can not be reliably maintained.

48.2.6.2.4 Receive

The PCS shall implement its Receive process as depicted in Figure 48–10, including compliance with the associated state variables as specified in 48.2.6.1 and including the optional EEE capability.

The PCS Receive process continuously performs the DECODE function on the received code-groups from the PCS Deskew Process via the ALIGNED_UNITDATA.indicate message. The PCS Receive process generates the receive clock signal of the XGMII (RX_CLK) as specified in Clause 46. State transitions in the PCS Receive state diagram that generate the data and control characters (RXD<31:0> and RXC<3:0>) on the XGMII occur synchronous to RX_CLK. The Receive process operates in the following two modes:

a) Data mode during packet reception including Start and Terminate. Additionally, Data mode is active whenever !||I|| columns are received during the Idle sequence or !||I|| or !||Q|| columns are received during the Fault sequence signifying either an error or unusual or unsupported indication. Valid code-groups received while in Data mode are mapped to corresponding XGMII data or control characters regardless of whether or not the control characters are valid XGMII control characters. Invalid or Error code-groups are mapped directly to XGMII Error control characters. All code-groups are mapped on a lane by lane basis.
b) Idle mode during idle reception excluding Start and Terminate. Idle mode is active whenever ||I|| is received during idle reception. ||I|| is translated to XGMII Idle control characters.

48.2.6.2.5 LPI state diagrams

A PCS that supports the EEE capability shall implement the LPI transmit and receive processes as shown in Figure 48–11 and Figure 48–12. The transmit LPI state diagram controls tx_quiet, which disables the transmitter when TRUE. The receive LPI state diagram controls align_status during LPI and synchronizes the receive state diagram with the end of the LPI.

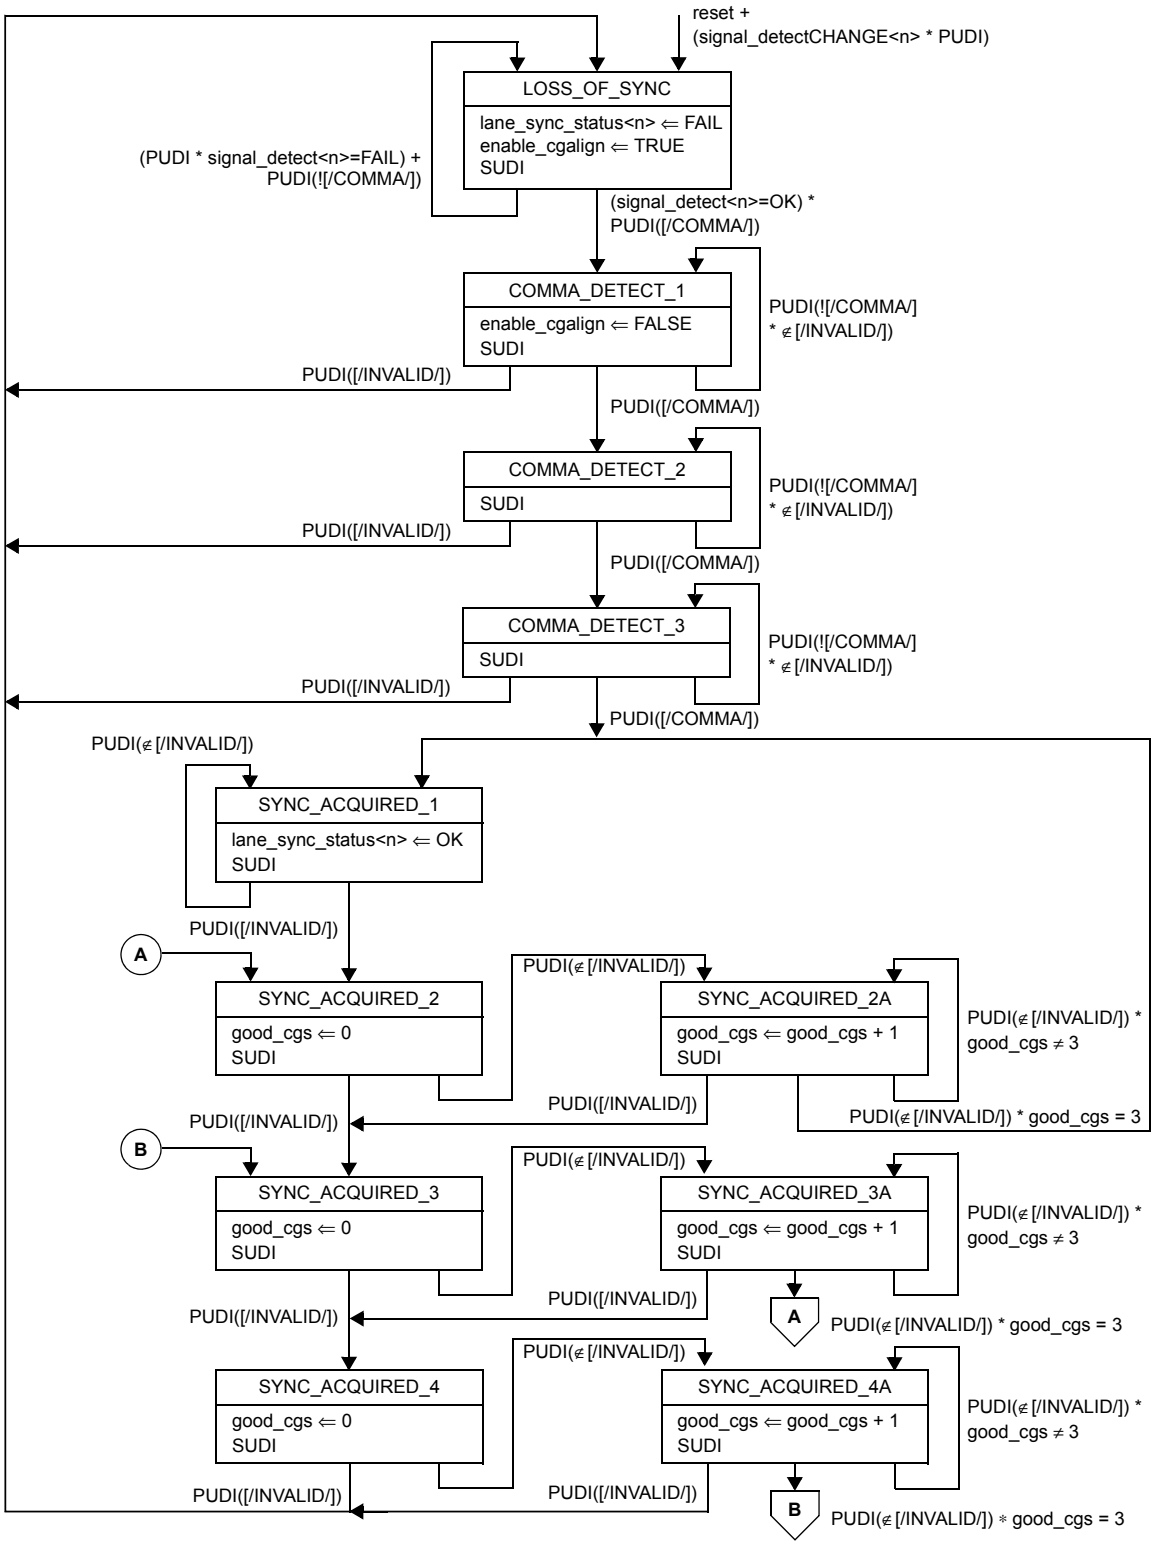

NOTE— lane_sync_status<n>, signal_detect<n> and signal_detectCHANGE<n>, refer to the number of the received lane n where n=0:3

Figure 48–8—PCS synchronization state diagram

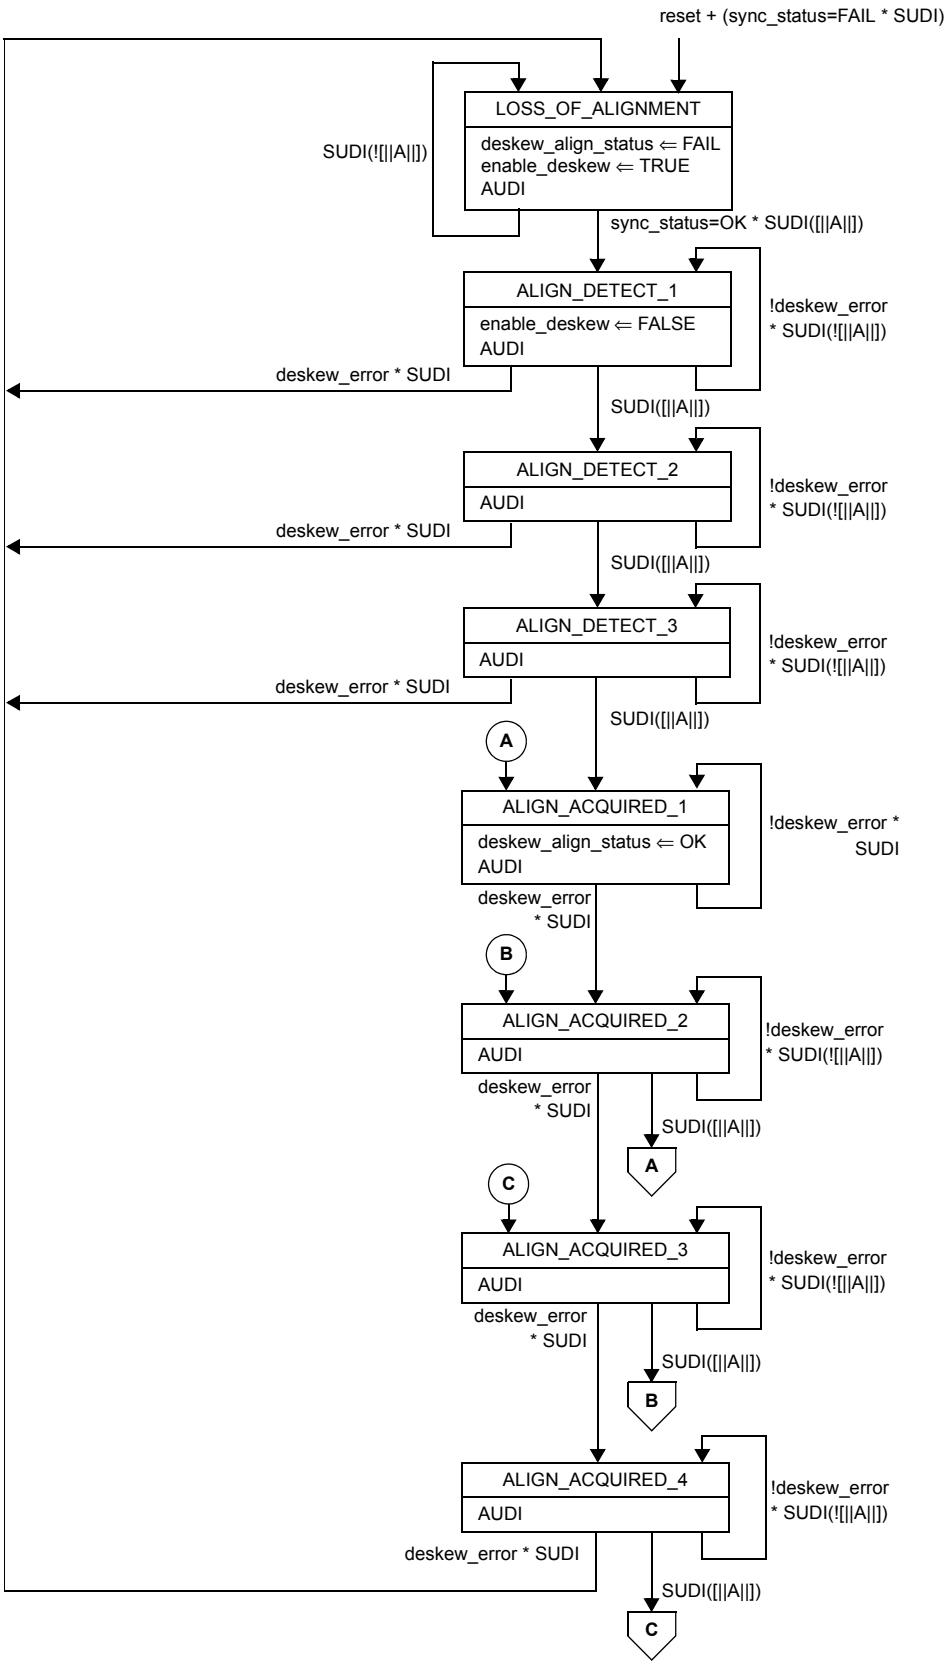

Figure 48–9—PCS deskew state diagram

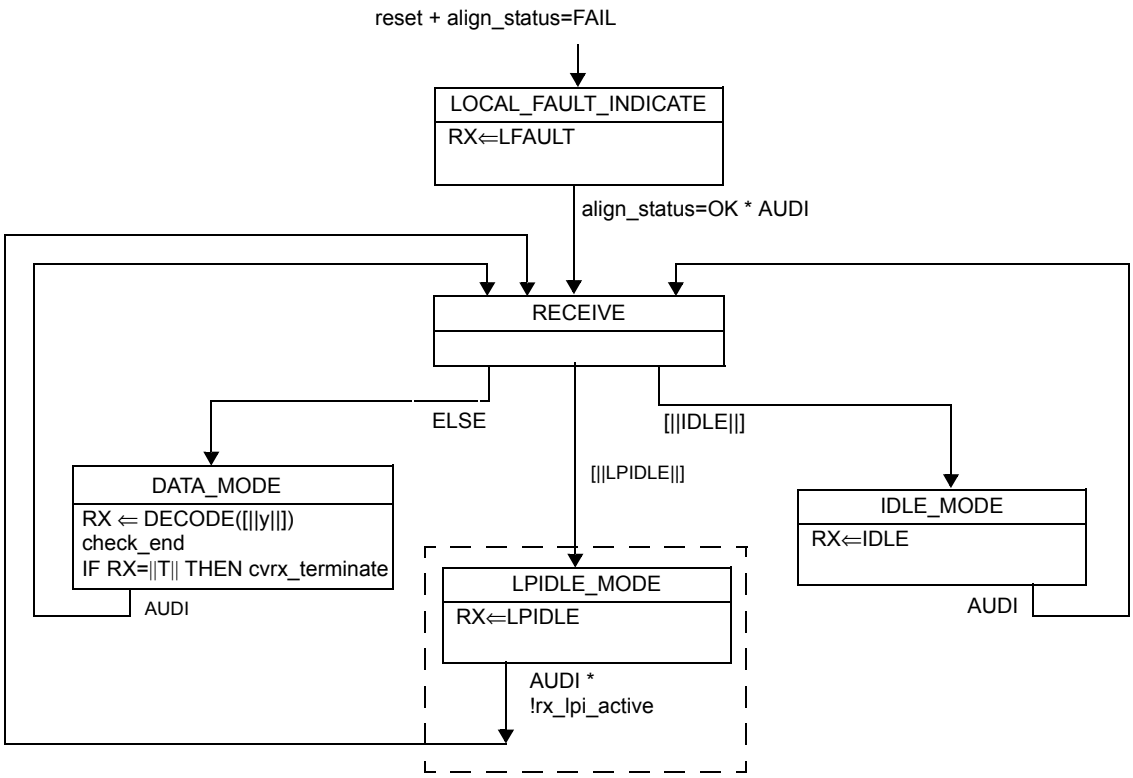

NOTE—Optional state to support LPI is shown inside the dotted box. The transition to the optional state is only possible with EEE capability.

Figure 48–10—PCS receive state diagram

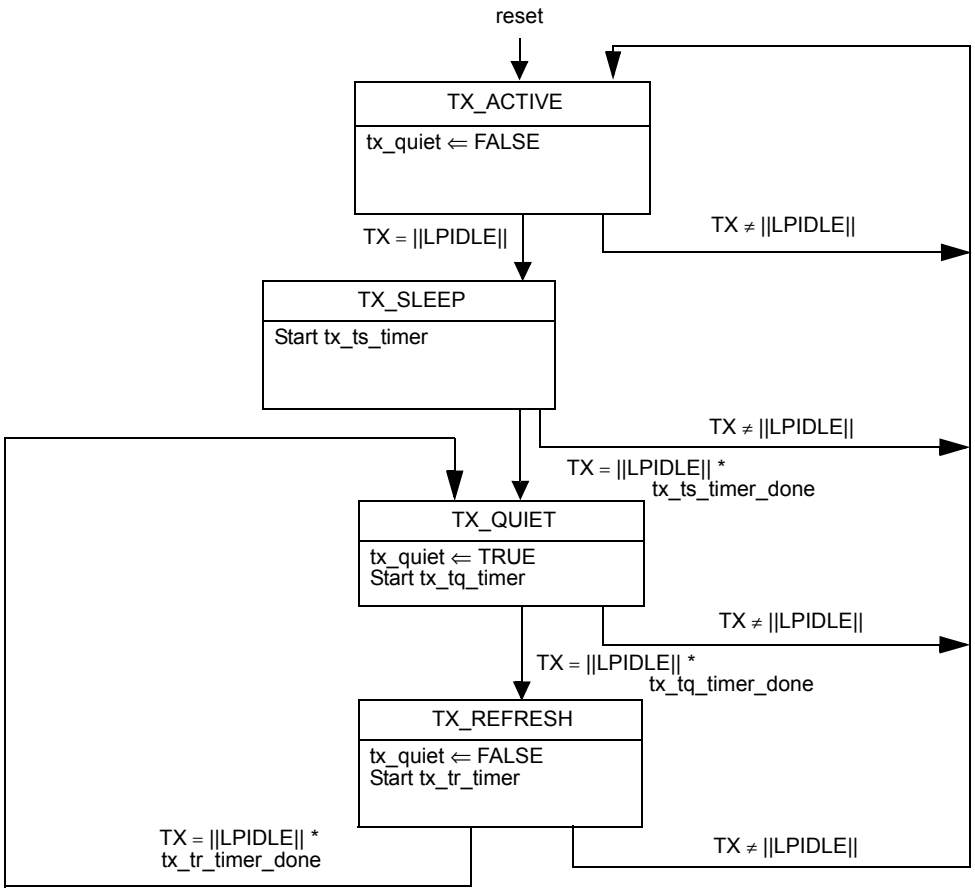

Figure 48–11—LPI Transmit state diagram

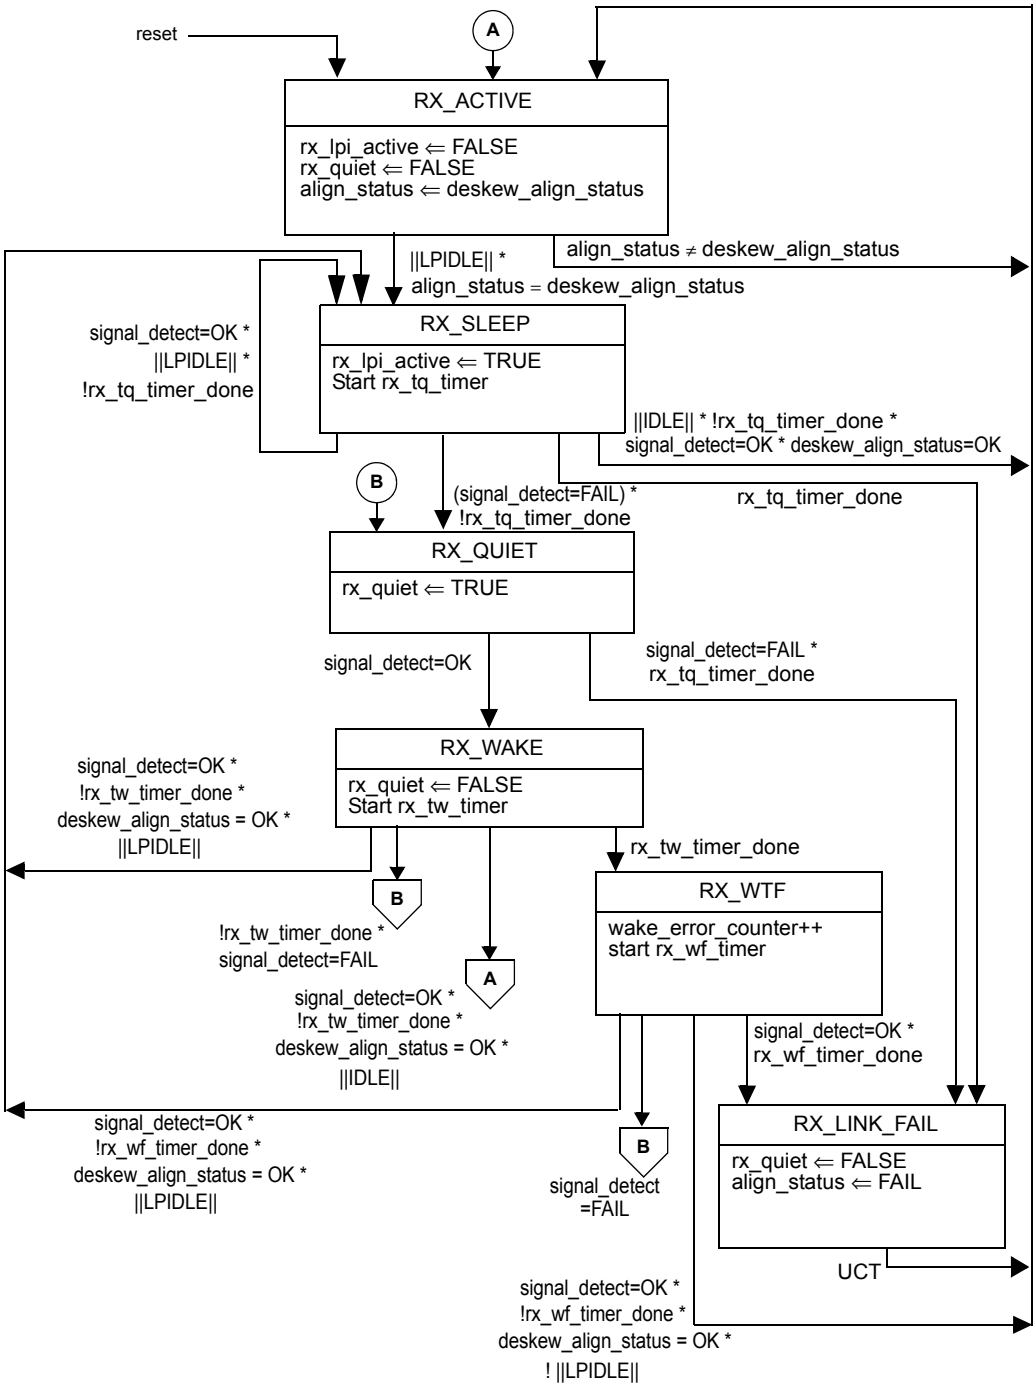

Figure 48–12—LPI Receive state diagram

The LPI functions shall use timer values for these state diagrams as shown in Table 48–9 for transmit and Table 48–10 for receive.

Table 48–9—Transmitter LPI timing parameters

Parameter	Description	Min	Max	Units
T_{SL}	Local Sleep Time from entering the TX_SLEEP state to when tx_quiet is set to TRUE	19.9	20.1	µs
T_{QL}	Local Quiet Time from when tx_quiet is set to TRUE to entry into the TX_REFRESH state	2.5	2.6	ms
T_{UL}	Local Refresh Time from entry into the TX_REFRESH state to entry into the TX_QUIET state	19.9	20.1	µs

Table 48–10—Receiver LPI timing parameters

Parameter	Description	Min	Max	Units
T_{QR}	The time the receiver waits for signal detect to be set to OK while in the RX_SLEEP and RX_QUIET states before asserting rx_fault	3	4	ms
T_{WR}	Time the receiver waits in the RX_WAKE state before indicating a wake time fault	—	9	µs
T_{WTF}	Wake time fault recovery time	—	1	ms

48.2.6.2.6 LPI status and management

For EEE capability, the PCS indicates to the management system that LPI is currently active in the receive and transmit directions using the status variable shown in Table 48–11.

Table 48–11—MDIO status indications

MDIO status variable	Register name	Register address	Note
Tx LPI received	PCS status register 1	3.1.11	Latched version of 3.1.9
Rx LPI received	PCS status register 1	3.1.10	Latched version of 3.1.8
Tx LPI indication	PCS status register 1	3.1.9	TRUE when not in state TX_ACTIVE
Rx LPI indication	PCS status register 1	3.1.8	TRUE when not in state RX_ACTIVE
Tx LPI received	PHY XS status register 1	4.1.11	Latched version of 4.1.9
Rx LPI received	PHY XS status register 1	4.1.10	Latched version of 4.1.8
Tx LPI indication	PHY XS status register 1	4.1.9	TRUE when not in state TX_ACTIVE
Rx LPI indication	PHY XS status register 1	4.1.8	TRUE when not in state RX_ACTIVE
Tx LPI received	DTE XS status register 1	5.1.11	Latched version of 5.1.9
Rx LPI received	DTE XS status register 1	5.1.10	Latched version of 5.1.8
Tx LPI indication	DTE XS status register 1	5.1.9	TRUE when not in state TX_ACTIVE
Rx LPI indication	DTE XS status register 1	5.1.8	TRUE when not in state RX_ACTIVE

48.2.6.3 Initialization process

Link initialization involves the completion of the PCS Synchronization and Deskew processes and the ability to transmit and receive code-groups via the PCS Transmit and Receive processes, respectively. The status register link_status flag is set to OK whenever the align_status flag is set to OK and no errors preventing link operation are present in the PCS or PMA.

48.2.6.4 Link status reporting

Link status reporting involves detection of link status conditions and the signaling of link fault status. The purpose of link status reporting is to quickly identify and convey link status conditions to the RS which can take the necessary action to activate (deactivate) the link via the setting of the RS link_fault 2-bit variable. Link status reporting and MAC packet transmission is mutually exclusive.

48.2.6.4.1 Link status detection

10GBASE-X link status conditions include signal and deskew status conditions. Link status conditions include Local Fault and Remote Fault signals. A receive fault is recognized by the PCS Receive process whenever align_status=FAIL. Other fault conditions are not detected by the PCS or PMA and are detected only by the RS. A fault condition may also be recognized by any 10GBASE-X process upon detection of an error condition, which prevents continued reliable operation, but this is beyond the scope of this standard.

48.2.6.4.2 Link status signaling

Link status signaling follows the detection or recognition of a link status condition and involves the generation of Sequence ordered-sets (||Q||) by the PCS Transmit process. Link status signaling involves the transmission of ||Q|| following ||A|| transmission in the Idle sequence as specified in 48.2.4.2.

48.2.6.4.3 Link status messages

A 10GBASE-X link status message is a Sequence ordered-set. Ordered-sets associated with link status messages are specified in Table 48–4. Link status messages detected by the PCS Receive process are forwarded to the XGMII.

48.2.7 Auto-Negotiation for Backplane Ethernet

The following requirements apply to a PCS used with a 10GBASE-KX4 PMD. Support for the Auto-Negotiation process defined in Clause 73 is mandatory. The PCS shall support the primitive AN_LINK.indication(link_status) (see 73.9). The parameter link_status shall take the value FAIL when align_status=FAIL and the value OK when align_status=OK. The primitive shall be generated when the value of link_status changes.

48.3 Physical Medium Attachment (PMA) sublayer

The PMA is specified in the form of a service interface to the PCS. These services are described in an abstract manner and do not imply any particular implementation. The PMA service interface supports the exchange of code-group information between PCS entities. The PMA converts code-groups into bits and passes these to the PMD, and vice versa.

48.3.1 Functions within the PMA

The PMA comprises the PMA Transmit process and PMA Receive process. Figure 48–5 depicts the mapping of the 36-bit-wide data and control path of the XGMII to the forty-bit-wide code-groups of the PMA service interface, and on to the four lane serial PMD service interface.

NOTE—Strict adherence to manufacturer-supplied guidelines for the operation and use of PMA serializer components is required to meet the jitter specifications of Clause 47, Clause 53, and Clause 54. The supplied guidelines should address the quality of power supply filtering associated with the transmit clock generator, and also the purity of the reference clock fed to the transmit clock generator.

48.3.1.1 PMA transmit process

The PMA Transmit process passes data unaltered (except for serializing) from the PCS directly to the PMD. Upon receipt of the PMA_UNITDATA.request primitive, the PMA shall individually serialize the four aligned 10-bit code-groups, one from each of four lanes, and transmit them to the PMD in the form of forty PMD_UNITDATA.request primitives, 10 each on four lanes.

Within each lane for each 10-bit code-group, the lowest numbered bit of the PMA_UNITDATA.indication parameter corresponds to the first bit transmitted to the PMD and the highest numbered bit of the PMA_UNITDATA.indication parameter corresponds to the last bit transmitted to the PMD. There is no numerical significance ascribed to the bits within a PMA code-group; that is, the code-group is simply a bit pattern that has some predefined interpretation. PMA to PMD bit and lane association is illustrated in Table 48–1. PMA to PMD bit ordering is illustrated in Figure 48–5.

PMA_UNITDATA.request primitives shall be generated with a frequency of 312.5 MHz ± 100 ppm. An internal clock multiplier unit multiplies the frequency of the 10-bit code-group based PMA_UNITDATA.request primitives by a factor of 10 to serialize the latched data out of the PMA and into the PMD.

48.3.1.2 PMA receive process

The PMA Receive process passes data unaltered (except for deserializing) from the PMD directly to the PCS. Upon receipt of forty PMD_UNITDATA.indication primitives, 10 in succession from each of four lanes, the PMA shall assemble the bits received into a 40-bit vector representing four 10-bit unaligned code-groups and pass that vector to the PCS as the parameter of four PMA_UNITDATA.indication primitives.

Within each lane for each 10-bit code-group, the lowest numbered bit of the PMA_UNITDATA.indication parameter corresponds to the first bit received from the PMD and the highest numbered bit of the PMA_UNITDATA.indication parameter corresponds to the last bit received from the PMD. There is no numerical significance ascribed to the bits within a PMA code-group; that is, the code-group is simply a bit pattern that has some predefined interpretation. Receive code-group alignment is performed by the PCS and is not applicable to PMA. PMA to PMD bit and lane association is illustrated in Table 48–1. PMA to PMD bit ordering is illustrated in Figure 48–5.

PMA_UNITDATA.indication primitives shall be generated with a frequency of 312.5 MHz ± 100 ppm. The line rate of each of four PMD lanes is 3.125 GBd ± 100 ppm. The serial data stream of the PMD includes an embedded clock that nominally operates at a frequency 10 times higher than that of a single lane code-group stream. The PMA Receive process shall recover a clock from a valid 8B/10B received data stream if the stream is within tolerance.

48.3.2 Service interface

The following primitives are defined:

PMA_UNITDATA.request(tx_code-group <39:0>)
PMA_UNITDATA.indication(rx_unaligned <39:0>)

48.3.2.1 PMA_UNITDATA.request

This primitive defines the transfer of data in the form of aligned code-groups from the PCS to the PMA. PMA_UNITDATA.request is generated by the PCS Transmit process.

48.3.2.1.1 Semantics of the service primitive

PMA_UNITDATA.request(tx_code-group <39:0>)

The data conveyed by PMA_UNITDATA.request is the tx_code-group <39:0> parameter defined in 48.2.6.1.3.

48.3.2.1.2 When generated

The PCS continuously sends tx_code-group <39:0> to the PMA at a nominal rate of 312.5 MHz ± 100 ppm, as governed by frequency and tolerance of XGMII TX_CLK.

48.3.2.1.3 Effect of receipt

Upon receipt of this primitive, the PMA generates a series of forty PMD_UNITDATA.request primitives, 10 for each of the four PMD lanes, tx_lane <3:0>, requesting the four lane serial transmission of tx_code-group <39:0> to the PMD.

48.3.2.2 PMA_UNITDATA.indication

This primitive defines the transfer of data in the form of code-groups from the PMA to the PCS. PMA_UNITDATA.indication is used by the PCS Synchronization process.

48.3.2.2.1 Semantics of the service primitive

PMA_UNITDATA.indication(rx_unaligned <39:0>)

The data conveyed by PMA_UNITDATA.indication is the rx_unaligned <39:0> parameter defined in 48.2.6.1.3.

48.3.2.2.2 When generated

The PMA continuously sends rx_unaligned <39:0> to the PCS at a nominal rate of 312.5 MHz ± 100 ppm, as governed by frequency and tolerance of PMD_UNITDATA.indication.

48.3.2.2.3 Effect of receipt

Upon receipt of this primitive, the PCS Synchronization process attempts to achieve code-group synchronization on each lane (see 48.2.6.2.2).

48.3.3 Loopback mode

Loopback mode shall be provided for the 10GBASE-X PMA and DTE XGXS, and optionally for the PHY XGXS, as specified in this subclause, by the transmitter and receiver of a device as a test function to the device. When Loopback mode is selected, transmission requests passed to the transmitter are shunted directly to the receiver, overriding any signal detected by the receiver on its attached link. A device is explicitly placed in Loopback mode (i.e., Loopback mode is not the normal mode of operation of a device). Loopback applies to all lanes as a group (i.e., the lane 0 transmitter is directly connected to the lane 0 receiver, the lane 1 transmitter is directly connected to the lane 1 receiver, etc.). The method of implementing Loopback mode is not defined by this standard.

Control of the Loopback function may be supported through the MDIO management interface of Clause 45 or equivalent.

NOTE—The signal path that is exercised in the Loopback mode is implementation specific, but it is recommended that this signal path encompass as much of the circuitry as is practical. The intention of providing this Loopback mode of operation is to permit diagnostic or self-test functions to test the transmit and receive data paths using actual data. Other loopback signal paths may also be enabled independently using loopback controls within other devices or sublayers.

48.3.3.1 Receiver considerations

Entry into or exit from Loopback mode may result in a temporary loss of synchronization.

48.3.3.2 Transmitter considerations

While in Loopback mode, the transmitter output is not defined.

48.3.4 Test functions

A limited set of test functions may be provided as an implementation option for testing of the transmitter function or for testing of an attached receiver.

Some test functions that are not defined by this standard may be provided by certain implementations. Compliance with the standard is not affected by the provision or exclusion of such functions by an implementation. Random jitter test patterns for 10GBASE-X are specified in Annex 48A. Test-pattern capability and selection is optional and supported via MDIO register bits defined in Clause 45. Jitter Test methodology for 10GBASE-X is specified in Annex 48B. It is recommended that the capability for generating the mixed-frequency, low frequency and high frequency patterns is implemented in the PCS.

A typical test function is the ability to transmit invalid code-groups within an otherwise valid PHY bit stream. Certain invalid PHY bit streams may cause a receiver to lose word and/or bit synchronization. See ANSI/INCITS 450 (FC-PI-4), subclause 5.4 for a more detailed discussion of receiver and transmitter behavior under various test conditions.

48.4 Compatibility considerations

There is no requirement for a compliant device to implement or expose any of the interfaces specified for the PCS or PMA. Implementations of an XGMII shall comply with the requirements as specified in Clause 46.

48.5 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control

sublayer, and PHY implementers must conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices.

The sum of transmit and receive delay contributed by the 10GBASE-X PCS shall be no more than 2048 BT.

The reference point for all MDI measurements is the 50% point of the mid-cell transition corresponding to the reference bit, as measured at the MDI.

48.6 Environmental specifications

All equipment subject to this clause shall conform to the requirements of 14.7 and applicable sections of ISO/IEC 11801:1995.

48.7 Protocol implementation conformance statement (PICS) proforma for Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer, type 10GBASE-X[5]

48.7.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 48, Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer, type 10GBASE-X, shall complete the following protocol implementation conformance statement (PICS) proforma. A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

48.7.2 Identification

48.7.2.1 Implementation identification

Supplier[1]	
Contact point for inquiries about the PICS[1]	
Implementation Name(s) and Version(s)[1, 3]	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)[2]	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

48.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 48, Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer, type 10GBASE-X
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [ ] Yes [ ] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2015.)	
Date of Statement	

[5]*Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

48.7.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
MD	MDIO	45, 48.1.3.1	Registers and interface supported	O	Yes [ ] No [ ]
XGXS	Support of XAUI/XGXS	47, 48.1.5		O	Yes [ ] No [ ]
XGE	XGMII compatibility interface	46, 48.1.3.1	Compatibility interface is supported	O	Yes [ ] No [ ]
LX4	Support of 10GBASE-LX4 PMD	53, 48.1.3.3		O	Yes [ ] No [ ]
CX4	Support of 10GBASE-CX4 PMD	54, 48.1.3.3		O	Yes [ ] No [ ]
LPI	Implementation of LPI	48.2.3		O	Yes [ ] No [ ]

48.7.4 PICS proforma tables for the PCS and PMA sublayer, type 10GBASE-X

48.7.4.1 Compatibility considerations

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	Jitter test patterns	48.3.4	As per Annex 48A	O	Yes [ ] No [ ]
CC2	Environmental specifications	48.6		M	Yes [ ]

48.7.4.2 PCS functions

Item	Feature	Subclause	Value/Comment	Status	Support
CG	Code-group usage	48.2.3	PCS support of 8B/10B code-groups	M	Yes [ ]
IOS	||I|| sequencing rules	48.2.4.2	All rules apply	M	Yes [ ]
PRBS	Random integer generator	48.2.4.2	$X^7 + X^3 + 1$ or $X^7 + X^6 + 1$. Used for ||A|| spacing. Optional and ancillary use for testing	M	Yes [ ]
CMA	Comma detection	48.2.4.2.1	comma+ and comma- for /K/	M	Yes [ ]
CKCU	Clock rate compensation in unencoded idle stream	48.2.4.2.3	Meets the requirements of 48.2.4.2.3	O	Yes [ ] No [ ]
ERR	Error indication	48.2.4.4	Replacement of invalid code-groups with /E/	M	Yes [ ]
TSD	Transmit state diagrams	48.2.6.2.1	Meets the requirements of Figures 48–7	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
SSD	Synchronization state diagram	48.2.6.2.2	Meets the requirements of Figure 48–8	M	Yes [ ]
DSD	Deskew state diagrams	48.2.6.2.3	Meets the requirements of Figure 48–9	M	Yes [ ]
RSD	Receive state diagrams	48.2.6.2.4	Meets the requirements of Figure 48–10	M	Yes [ ]
AN1*	Support for use with a 10GBASE-KX4 PMD	48.2.7	AN technology dependent interface described in Clause 73	O	Yes [ ]
AN2	AN_LINK.indication primitive	48.2.7	Support of the primitive AN_LINK.indication(link_status), when the PCS is used with 10GBASE-KX4 PMD	AN1:M	Yes [ ]
AN3	link_status parameter	48.2.7	Takes the value OK or FAIL, as described in 48.2.7	AN1:M	Yes [ ]
AN4	Generation of AN_LINK.indication primitive	48.2.7	Generated when the value of link_status changes	AN1:M	Yes [ ]

48.7.4.3 PMA Functions

Item	Feature	Subclause	Value/Comment	Status	Support
PMAT	Transmit function	48.3.1.1	PMA_UNITDATA.request	M	Yes [ ]
TXRT	Transmit rate	48.3.1.1	3.125 GBd ± 100 ppm	M	Yes [ ]
PMAR	Receive function	48.3.1.2	PMA_UNITDATA.indication	M	Yes [ ]
RXRT	Receive rate	48.3.1.2	3.125 GBd ± 100 ppm	M	Yes [ ]
CDR	Clock and data recovery	48.3.1.2	Required if line rate is within 3.125 GBd ± 100 ppm	M	Yes [ ]

48.7.4.4 Interface functions

Item	Feature	Subclause	Value/Comment	Status	Support
SKEW	Allowable lane skew	48.2.4.2.2	Table 48–5	M	Yes [ ]
LBXPMA	10GBASE-X PMA Loopback mode	48.3.3		M	Yes [ ]
LBDTE	DTE XGXS Loopback mode	48.3.3		M	Yes [ ]
LBPHY	PHY XGXS Loopback mode	48.3.3		O	Yes [ ]
DLY	Delay constraints	48.5	2048 BT	M	Yes [ ]

48.7.4.5 LPI functions

Item	Feature	Subclause	Value/Comment	Status	Support
LP-01	Receive state diagrams	48.2.6.2	Support additions to Figure 48–10 for LPI operation	LPI:M	Yes [ ] No [ ]
LP-02	LPI transmit state diagrams	48.2.6.2.5	Meet the requirements of Figure 48–11	LPI:M	Yes [ ] No [ ]
LP-03	LPI receive state diagrams	48.2.6.2.5	Meet the requirements of Figure 48–12	LPI:M	Yes [ ] No [ ]
LP-04	LPI transmit timing	48.2.6.2.5	Meet the requirements of Table 48–9	LPI:M	Yes [ ] No [ ]
LP-05	LPI receive timing	48.2.6.2.5	Meet the requirements of Table 48–10	LPI:M	Yes [ ] No [ ]

49. Physical Coding Sublayer (PCS) for 64B/66B, type 10GBASE-R

49.1 Overview

49.1.1 Scope

This clause specifies the Physical Coding Sublayer (PCS) that is common to a family of 10 Gb/s Physical Layer implementations, known as 10GBASE-R. This PCS can connect directly to one of the 10GBASE-R Physical Layers: 10GBASE-SR, 10GBASE-LR, 10GBASE-ER, 10GBASE-LRM, and 10GBASE-KR. Alternatively, this PCS can connect to a WAN Interface Sublayer (WIS), which will produce the 10GBASE-W encoding (10GBASE-R encoded data stream encapsulated into frames compatible with SONET and SDH networks) for transport by the 10GBASE-W Physical Layers: 10GBASE-SW, 10GBASE-LW, and 10GBASE-EW. The term 10GBASE-R is used when referring generally to Physical Layers using the PCS defined here.

The 10GBASE-R is based on a 64B/66B code. The 64B/66B code supports data and control characters, while maintaining robust error detection.

10GBASE-R PCS maps the interface characteristics of the WIS when present, and the PMA sublayer to the services expected by the Reconciliation and XGXS sublayers. 10GBASE-R can be extended to support any full duplex medium requiring only that the medium be compliant at the PMA level. 10GBASE-R PCS may be attached through the PMA sublayer to a LAN PMD sublayer supporting a data rate of 10 Gb/s or it may be attached to a WAN PMD through the WIS and PMA sublayers. When attached to a WAN sublayer, this PCS adapts the data stream to the WAN data rate.

49.1.2 Objectives

The following are the objectives of 10GBASE-R:

a) Support the full duplex Ethernet MAC.
b) Provide 10 Gb/s data rate at the XGMII.
c) Support LAN PMDs operating at 10 Gb/s and WAN PMDs operating at SONET STS-192c/SDH VC-4-64c rate.
d) Support cable plants using cabled optical fiber compliant with ISO/IEC 11801:1995 as specified in Clause 52 and Clause 68.
e) Allow for a nominal network extent of up to 40 km.
f) Support a BER objective of 10^{-12}.

49.1.3 Relationship of 10GBASE-R to other standards

Figure 49–1 depicts the relationships among the 10GBASE-R sublayers (shown shaded), the Ethernet MAC and reconciliation layers, and the higher layers.

49.1.4 Summary of 10GBASE-R and 10GBASE-W sublayers

The following subclauses provide an overview of the 10GBASE-R and 10GBASE-W sublayers.[6] Figure 49–1 depicts the relationship between the 10GBASE-R PCS and its associated sublayers.

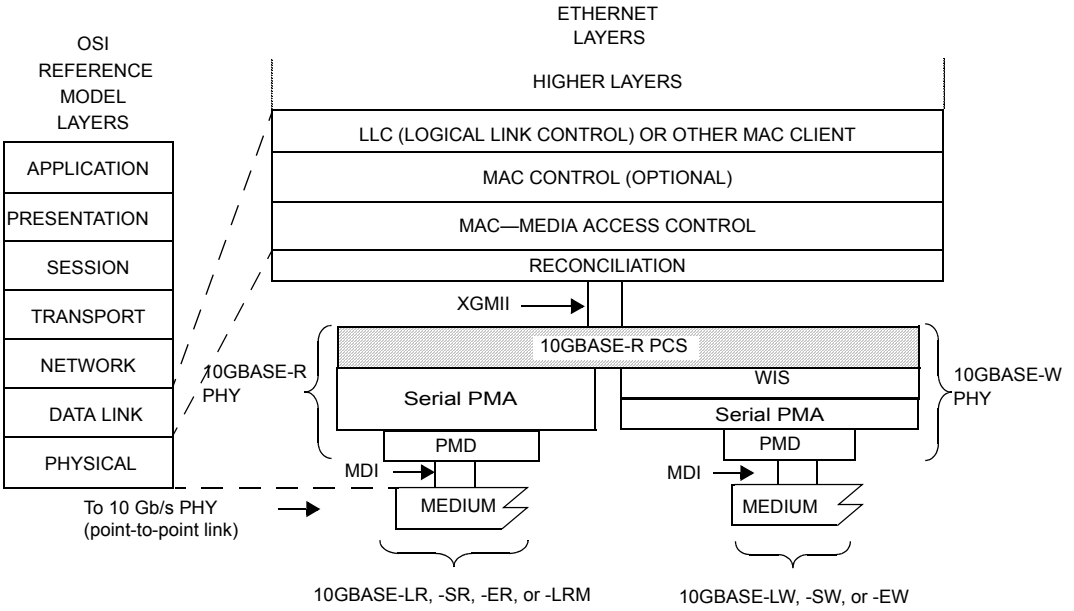

MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
WIS = WAN INTERFACE SUBLAYER
XGMII = 10 GIGABIT MEDIA INDEPENDENT INTERFACE

NOTE—The PMD sublayers are mutually independent.

PMD TYPES:
Medium:
E = PMD FOR FIBER—1550 nm WAVELENGTH
L = PMD FOR FIBER—1310 nm WAVELENGTH
S = PMD FOR FIBER—850 nm WAVELENGTH
M = PMD WITH DISPERSION COMPENSATION FOR MULTI-MODE FIBER
Encoding:
R = 64B/66B ENCODED WITHOUT WIS
W = 64B/66B ENCODED WITH WIS

Figure 49–1—10GBASE-R PCS relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

49.1.4.1 Physical Coding Sublayer (PCS)

The PCS service interface is the 10 Gigabit Media Independent Interface (XGMII), which is defined in Clause 46. The XGMII provides a uniform interface to the Reconciliation Sublayer for all 10 Gb/s PHY implementations (e.g., not only 10GBASE-R but also other types of 10 Gigabit PHY entities).

The 10GBASE-R PCS provides all services required by the XGMII, including the following:

a) Encoding (decoding) of eight XGMII data octets to (from) 66-bit blocks (64B/66B).
b) Transferring encoded data to (from) the PMA in 16 bit transfers.
c) When connected to a WAN PMD, deleting (inserting) idles to compensate for the rate difference between the MAC and PMD.

[6] The 10GBASE-R PHY consists of that portion of the Physical Layer between the MDI and XGMII consisting of the PCS, PMA, and PMD sublayers. The 10GBASE-R PHY is roughly analogous to the 100BASE-X PHY.

d) Determining when a functional link has been established and informing the management entity via the MDIO when the PHY is ready for use.

49.1.4.2 WAN Interface Sublayer (WIS)

The WIS provides a medium-independent means for the PCS to operate over WAN links. It creates a 10GBASE-W encoding by encapsulating the encoded data steam from the 10GBASE-R PCS in frames compatible with SONET and SDH transmission formats. The WIS is specified in Clause 50.

49.1.4.3 Physical Medium Attachment (PMA) sublayer

The PMA provides a medium-independent means for the PCS to support the use of a range of physical media. The 10GBASE-R PMA performs the following functions:

a) Mapping of transmit and receive data streams between the PCS or WIS and PMA via the PMA service interface.
b) Serialization (deserialization) of bits for transmission (reception) on the underlying serial PMD.
c) Recovery of clock from the received data stream.
d) Mapping of transmit and receive bits between the PMA and PMD via the PMD service interface.
e) Optionally provides data loopback at the PMA service interface.

The PMA is specified in Clause 51.

49.1.4.4 Physical Medium Dependent (PMD) sublayer

The PMD and its media are specified in Clause 52 and Clause 68.

The MDI, logically subsumed within each PMD subclause, is the actual medium attachment for the various supported media.

49.1.4.5 Bit ordering across 10GBASE-R and 10GBASE-W sublayers.

The ordering of bits and octets for the case when no WIS is present between the PCS and PMA is shown in Figure 49–2. The diagram depicts the bit mappings for the data path at the service interfaces.

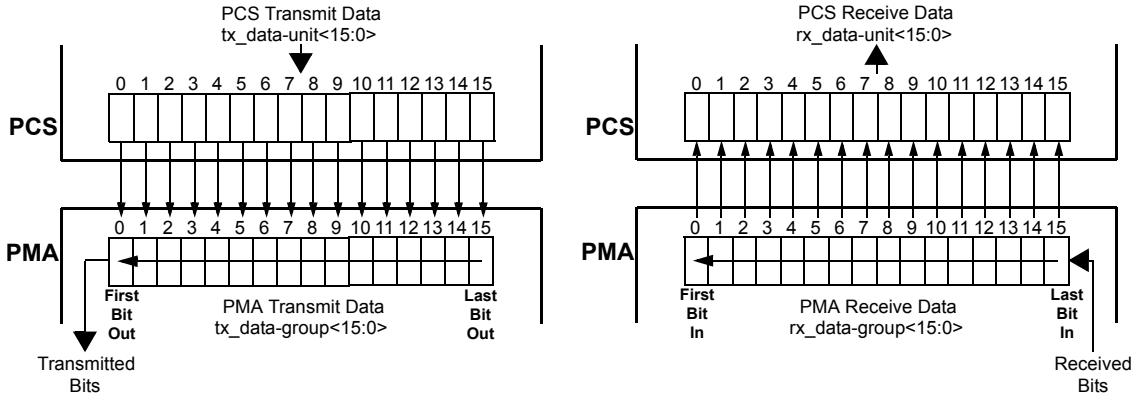

Figure 49–2—10GBASE-R PHY transmission order

The bit/octet ordering when a WIS is interposed between the PCS and PMA is shown in Figure 49–3. The details of WIS frame generation, as well as internal interfaces, have been omitted. Note that since SONET

convention is to send the most significant bit (MSB) of each octet first and Ethernet convention is to send the least significant bit (LSB) of each octet first, the LSBs from the PCS octets are mapped into the MSBs of the WIS octets. It should be clear from the diagram that the Ethernet order of transmission (i.e., least-significant bit first), and hence the error detecting properties of the Ethernet CRC, are preserved in spite of the different bit ordering conventions followed by the WIS. Also, the SONET/SDH bit labeling conventions are different from the usual IEEE 802.3 bit labeling. The bits of a SONET/SDH octet are labeled from 1 to 8 with bit 1 being the MSB. Ethernet conventions label bits of an n-bit field from 0 to n–1 with bit 0 being the LSB. Figure 49–3 shows the results of these conventions. For example, tx_data-unit<0> through tx_data-unit<7> map to bits 1 through 8 respectively of a WIS octet.

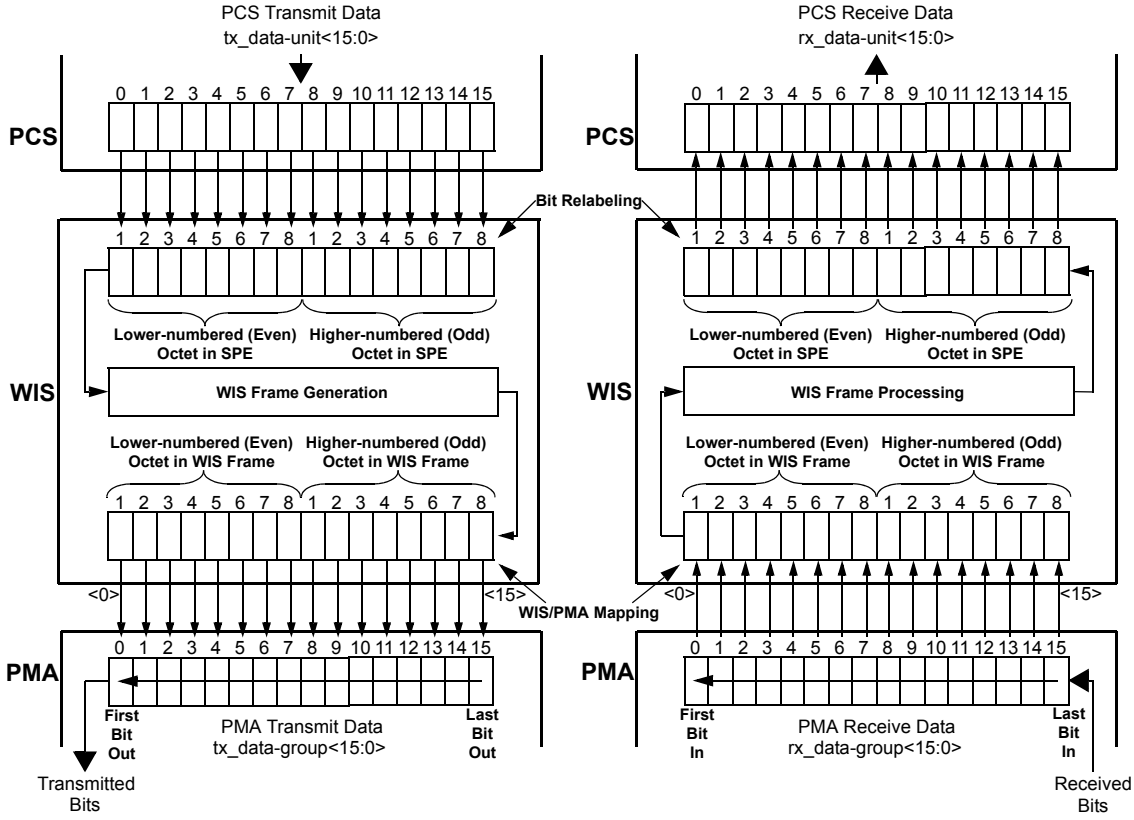

Figure 49–3—10GBASE-W PHY Transmission order

49.1.5 Inter-sublayer interfaces

There are a number of interfaces employed by 10GBASE-R. Some (such as the PMA service interface) use an abstract service model to define the operation of the interface. The PCS service interface is the XGMII that is defined in Clause 46. The XGMII has an optional physical instantiation. An optional physical instantiation of the PMA service interface has also been defined (see Clause 51). It is called XSBI (10 Gigabit Sixteen Bit Interface). Figure 49–4 depicts the relationship and mapping of the services provided by all of the interfaces relevant to 10GBASE-R.

The upper interface of the PCS may connect to the Reconciliation Sublayer through the XGMII or the PCS may connect to an XGXS sublayer. The XGXS and the Reconciliation Sublayer provide the same service interface to the PCS. The lower interface of the PCS may connect to the WIS to support a WAN PMD or to the PMA sublayer to support a 10GBASE-R LAN PMD. The WIS and PMA interfaces are functionally

equivalent except for data rate. When the PCS is connected directly to a LAN PMA, the nominal rate of the PMA service interface is 644.53 Mtransfers/s, which provides capacity for the MAC data rate of 10 Gb/s. When the PCS is connected to a WAN PMA, the nominal rate of the WIS service interface is 599.04 Mtransfers/s and the MAC uses IFS stretch mode to ensure that there will be enough idle time that the PCS can delete idles to adjust to the lower rate. Since the data rates are different, WIS and PMA interface connections pose somewhat different constraints. The PCS shall support connection to either a WIS or to a PMA and may optionally support both.

If the optional Energy-Efficient Ethernet (EEE) capability is supported (see Clause 78, 78.3) then the interface with the PMA sublayer (or FEC sublayer) includes rx_mode and tx_mode to control power states in lower sublayers and energy_detect that indicates whether the PMD sublayer has detected a signal at the receiver. If the PHY includes an FEC sublayer, the interface includes rx_lpi_active to indicate that the LPI receive state diagram is not in RX_ACTIVE state.

It is important to note that, while this specification defines interfaces in terms of bits, octets, and frames, implementations may choose other data-path widths for implementation convenience.

49.1.6 Functional block diagram

Figure 49–4 provides a functional block diagram of the 10GBASE-R PHY.

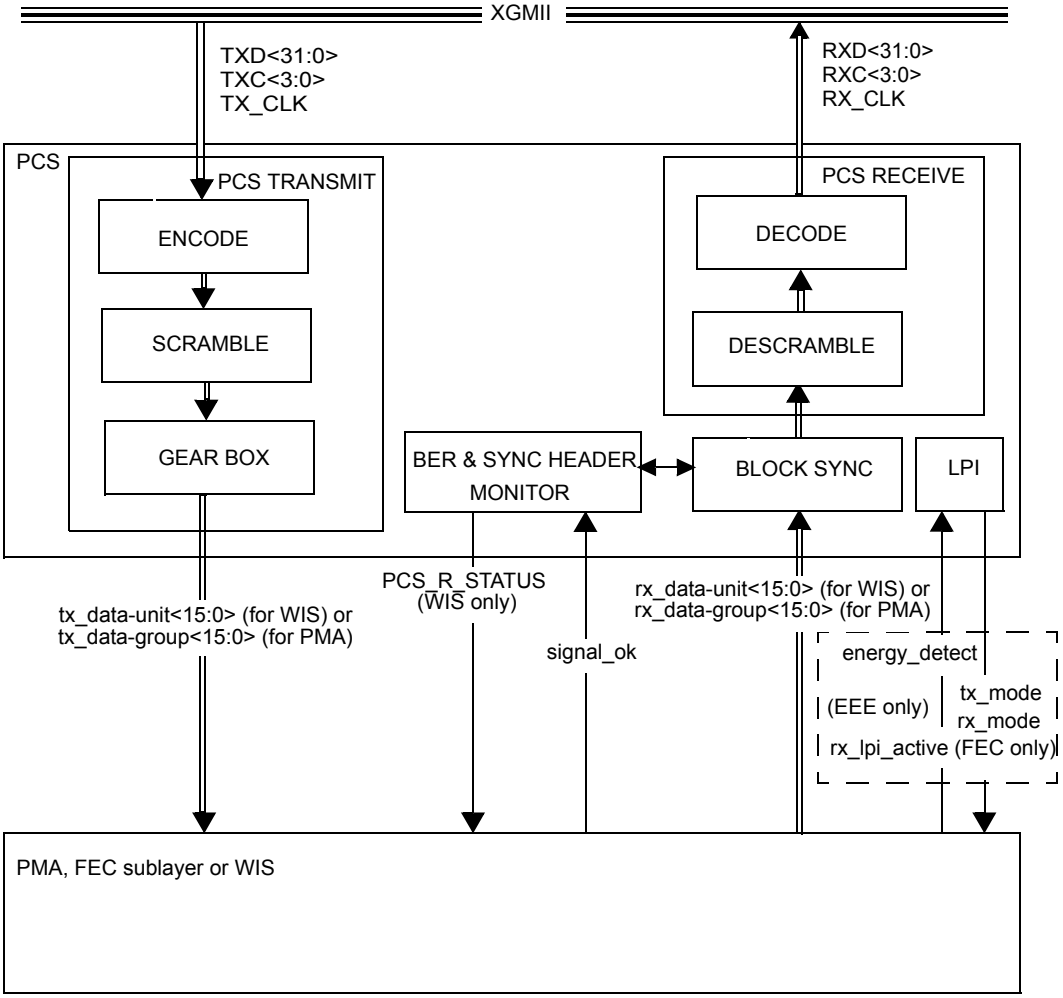

Figure 49–4—Functional block diagram

49.2 Physical Coding Sublayer (PCS)

49.2.1 PCS service interface (XGMII)

The PCS service interface allows the 10GBASE-R PCS to transfer information to and from a PCS client. A PCS client is generally the Reconciliation Sublayer or an XGXS sublayer. The PCS Interface is precisely defined as the 10 Gigabit Media Independent Interface (XGMII) in Clause 46.

49.2.2 Functions within the PCS

The PCS comprises the PCS Transmit, Block Synchronization, PCS Receive, and BER monitor processes for 10GBASE-R. The PCS shields the Reconciliation Sublayer (and MAC) from the specific nature of the underlying channel. The PCS transmit channel and receive channel can each operate in normal mode or, when not attached to a WIS, test-pattern mode. When the PCS is attached to a WIS, the WIS provides the test-pattern functionality.

When communicating with the XGMII, the PCS uses a four octet-wide, synchronous data path, with packet delimiting being provided by transmit control signals (TXCn = 1) and receive control signals (RXCn = 1). When communicating with the PMA or WIS, the PCS uses a 16-bit wide, synchronous data path that conveys 16 encoded bits. Alignment to 64B/66B block is performed in the PCS. The WIS and PMA sublayers operate independent of block and packet boundaries. The PCS provides the functions necessary to map packets between the XGMII format and the PMA service interface format.

When the transmit channel is in normal mode, the PCS Transmit process continuously generates blocks based upon the TXD <31:0> and TXC <3:0> signals on the XGMII. The Gearbox function of the PCS Transmit process then packs the resulting bits into 16-bit transmit data-units. Transmit data-units are sent to the PMA or WIS service interface via the PMA_UNITDATA.request or WIS_UNITDATA.request primitive, respectively. When the WIS is present, the PCS Transmit process also adapts between the XGMII and WIS data rates by deleting idle characters.

When the transmit channel is in test-pattern mode, a test pattern is packed into the transmit data-units that are sent to the PMA service interface via the PMA_UNITDATA.request primitive.

When the receive channel is in normal mode, the PCS Synchronization process continuously monitors PMA_SIGNAL.indication(SIGNAL_OK) or WIS_SIGNAL.indication(SIGNAL_OK). When SIGNAL_OK indicates OK, then the PCS Synchronization process accepts data-units via the PMA_UNITDATA.indication primitive or the WIS_UNITDATA.indication primitive. It attains block synchronization based on the 2-bit synchronization headers and conveys received blocks to the PCS Receive process. The PCS Synchronization process sets the sync_status flag to indicate whether the PCS has obtained synchronization.

When the PCS Synchronization process has obtained synchronization, the BER monitor process monitors the signal quality asserting hi_ber if excessive errors are detected. When sync_status is asserted and hi_ber is de-asserted, the PCS Receive process continuously accepts blocks. The PCS Receive process monitors these blocks and generates RXD <31:0> and RXC <3:0> on the XGMII. When the WIS is present, the PCS Receive process also adapts between the WIS and XGMII data rates by inserting idle characters. The PCS Receive process also sends the WIS_SIGNAL.request(PCS_R_STATUS) primitive to the WIS to indicate its status. The primitive is sent when the value of PCS_R_STATUS changes. The value of PCS_R_STATUS shall be FAIL when the Receive state diagram is in the RX_INIT state and OK otherwise.

When the receive channel is in test-pattern mode, the BER monitor process is disabled. The Receive process will be held in the RX_INIT state. The received bits will be compared to the test pattern and errors counted.

A PCS that supports direct connection to a PMA shall provide transmit test-pattern mode for the square wave and pseudo-random patterns, and shall provide receive test-pattern mode for the pseudo-random pattern. The PCS may provide support for the PRBS9 transmit test pattern. It may provide support for the PRBS31 test pattern. Support of the optional PRBS31 test pattern shall include both the transmit and the receive capability for that pattern. Test-pattern mode is activated separately for transmit and receive. A PCS that supports direct connection to a PMA shall support transmit test-pattern mode and receive test-pattern mode operating simultaneously (if applicable) so as to support loopback testing. Test-pattern mode is provided by the WIS when a WIS is present.

49.2.3 Use of blocks

The PCS maps XGMII signals into 66-bit blocks, and vice versa, using a 64B/66B coding scheme. The synchronization headers of the blocks allow establishment of block boundaries by the PCS Synchronization process. Blocks are unobservable and have no meaning outside the PCS. The PCS functions ENCODE and DECODE generate, manipulate, and interpret blocks as provided by the rules in 49.2.4.

49.2.4 64B/66B transmission code

The PCS uses a transmission code to improve the transmission characteristics of information to be transferred across the link and to support transmission of control and data characters. The encodings defined by the transmission code ensure that sufficient transitions are present in the PHY bit stream to make clock recovery possible at the receiver. The encoding also preserves the likelihood of detecting any single or multiple bit errors that may occur during transmission and reception of information. In addition, the synchronization headers of the code enable the receiver to achieve block alignment on the incoming PHY bit stream. The 64B/66B transmission code specified for use in this standard has a high transition density and is a run-length-limited code.

The relationship of block bit positions to XGMII, PMA, and other PCS constructs is illustrated in Figure 49–5 for transmit and Figure 49–6 for receive. These figures illustrate the processing of a block containing 8 data octets. See 49.2.4.3 for information on how blocks containing control characters are mapped. Note that the sync header is generated by the encoder and bypasses the scrambler.

49.2.4.1 Notation conventions

For values shown as binary, the leftmost bit is the first transmitted bit.

64B/66B encodes 8 data octets or control characters into a block. Blocks containing control characters also contain a block type field. Data octets are labeled D_0 to D_7. Control characters other than /O/, /S/ and /T/ are labeled C_0 to C_7. The control character for ordered set is labeled as O_0 or O_4 since it is only valid on the first octet of the XGMII. The control character for start is labeled as S_0 or S_4 for the same reason. The control character for terminate is labeled as T_0 to T_7.

Two consecutive XGMII transfers provide eight characters that are encoded into one 66-bit transmission block. The subscript in the above labels indicates the position of the character in the eight characters from the XGMII transfers.

Contents of block type fields, data octets and control characters are shown as hexadecimal values. The LSB of the hexadecimal value represents the first transmitted bit. For instance, the block type field 0x1e is sent from left to right as 01111000. The bits of a transmitted or received block are labeled TxB<65:0> and RxB<65:0> respectively where TxB<0> and RxB<0> represent the first transmitted bit. The value of the sync header is shown as a binary value. Binary values are shown with the first transmitted bit (the LSB) on the left.

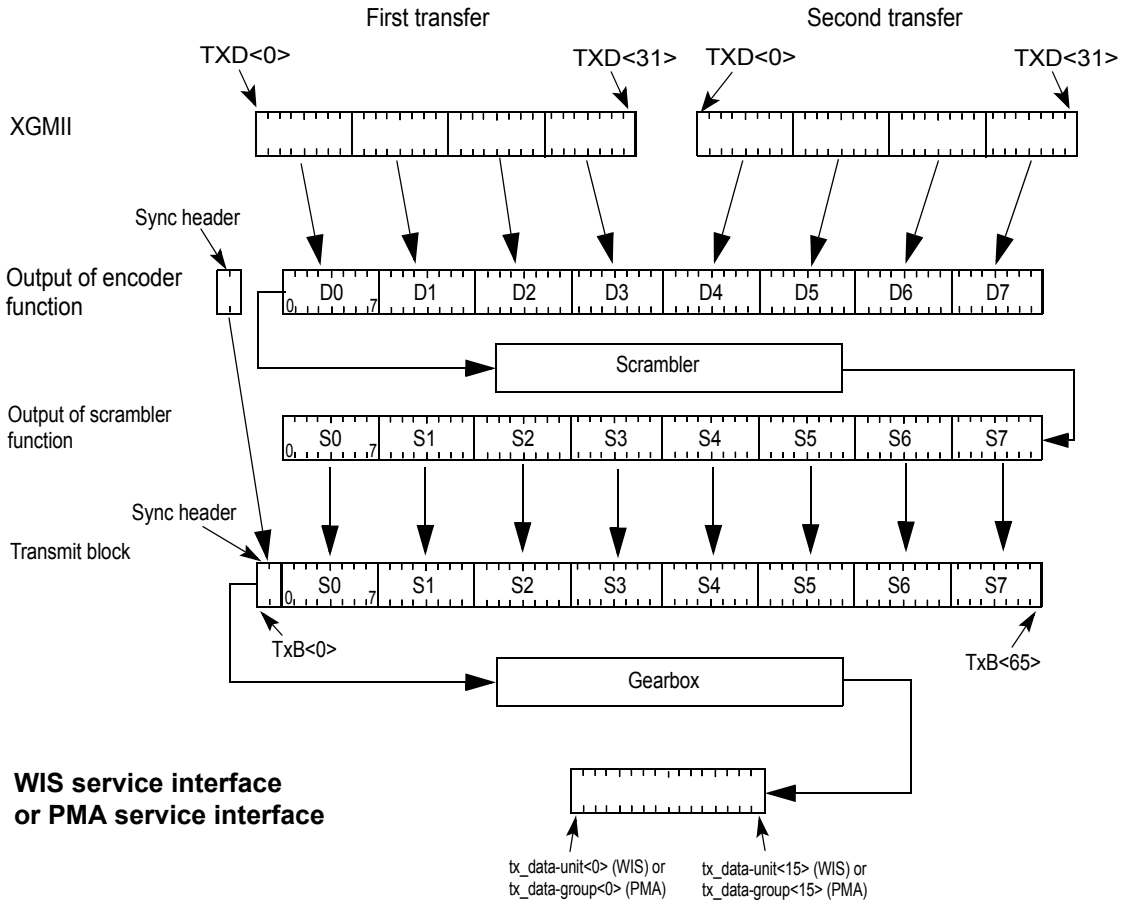

Figure 49–5—PCS Transmit bit ordering

49.2.4.2 Transmission order

Block bit transmission order is illustrated in Figure 49–5 and Figure 49–6. Note that these figures show the mapping from XGMII to 64B/66B block for a block containing eight data characters.

49.2.4.3 Block structure

Blocks consist of 66 bits. The first two bits of a block are the synchronization header (sync header). Blocks are either data blocks or control blocks. The sync header is 01 for data blocks and 10 for control blocks. Thus, there is always a transition between the first two bits of a block. The remainder of the block contains the payload. The payload is scrambled and the sync header bypasses the scrambler. Therefore, the sync header is the only position in the block that always contains a transition. This feature of the code is used to obtain block synchronization.

Data blocks contain eight data characters. Control blocks begin with an 8-bit block type field that indicates the format of the remainder of the block. For control blocks containing a Start or Terminate character, that character is implied by the block type field. Other control characters are encoded in a 7-bit control code or a 4-bit O Code. Each control block contains eight characters.

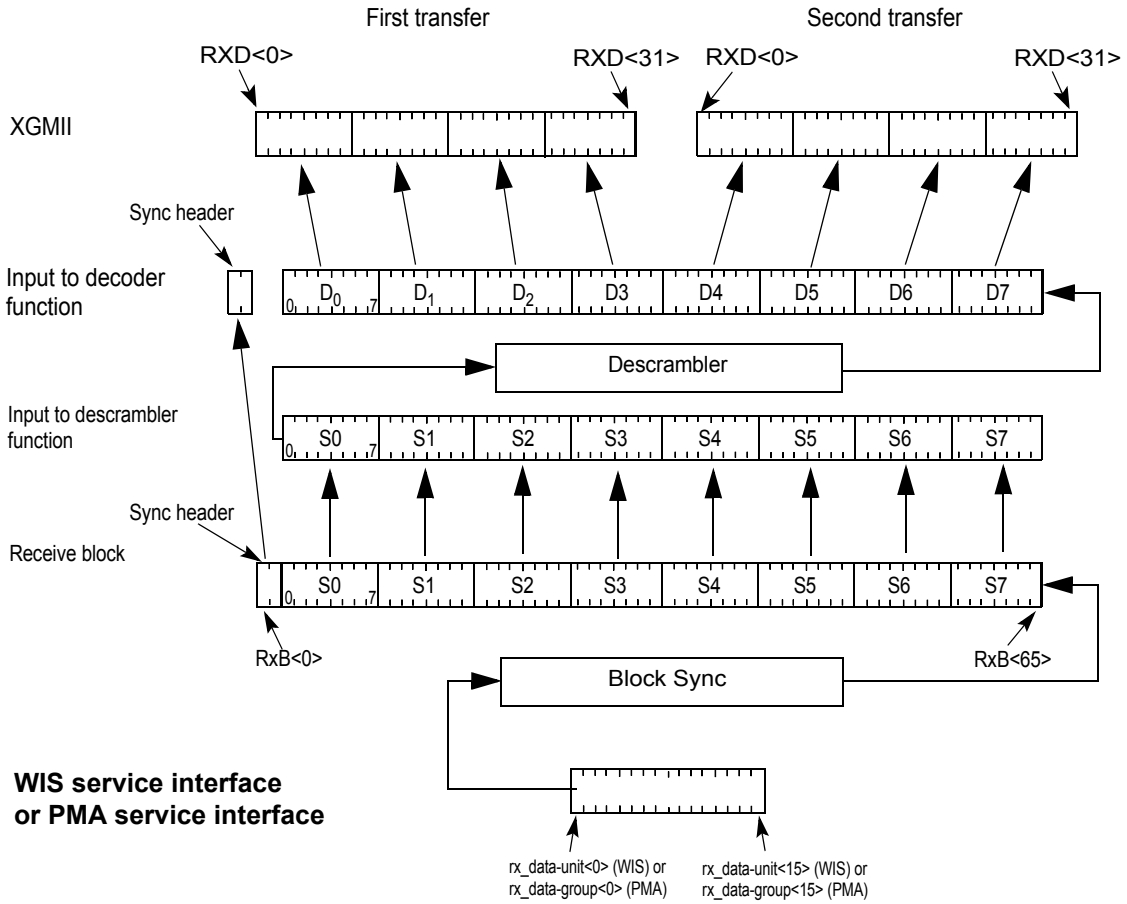

Figure 49–6—PCS Receive bit ordering

The format of the blocks is as shown in Figure 49–7. In the figure, the column labeled Input Data shows, in abbreviated form, the eight characters used to create the 66-bit block. These characters are either data characters or control characters and, when transferred across the XGMII interface, the corresponding TXC or RXC bit is set accordingly. Within the Input Data column, D_0 through D_7 are data octets and are transferred with the corresponding TXC or RXC bit set to zero. All other characters are control octets and are transferred with the corresponding TXC or RXC bit set to one. The single bit fields (thin rectangles with no label in the figure) are sent as zero and ignored upon receipt.

Bits and field positions are shown with the least significant bit on the left. Hexadecimal numbers are shown in normal hexadecimal. For example the block type field 0x1e is sent as 01111000 representing bits 2 through 9 of the 66 bit block. The least significant bit for each field is placed in the lowest numbered position of the field.

All unused values of block type field[7] are reserved.

[7]The block type field values have been chosen to have a 4-bit Hamming distance between them. The only unused value that maintains the Hamming distance is 0x00.

49.2.4.4 Control codes

The same set of control characters are supported by the XGMII and the 10GBASE-R PCS. The representations of the control characters are the control codes. XGMII encodes a control character into an octet (an eight bit value). The 10GBASE-R PCS encodes the start and terminate control characters implicitly by the block type field. The 10GBASE-R PCS encodes the ordered set control codes using a combination of the block type field and a 4-bit O code for each ordered set. The 10GBASE-R PCS encodes each of the other control characters into a 7-bit C code.

The control characters and their mappings to 10GBASE-R control codes and XGMII control codes are specified in Table 49–1. All XGMII and 10GBASE-R control code values that do not appear in the table shall not be transmitted and shall be treated as an error if received. The ability to transmit or receive Low Power Idle (LPI) is required for PHYs that support EEE (see Clause 78). If EEE is not supported, LPI shall not be transmitted and shall be treated as an error if received.

Input Data	**Sync**	**Block Payload**							
Bit Position: / **Data Block Format:**	0 1	2							65
$D_0\ D_1\ D_2\ D_3/D_4\ D_5\ D_6\ D_7$	01	D_0	D_1	D_2	D_3	D_4	D_5	D_6	D_7
Control Block Formats:		**Block Type Field**							
$C_0\ C_1\ C_2\ C_3/C_4\ C_5\ C_6\ C_7$	10	0x1e	C_0	C_1	C_2	C_3	C_4	C_5	C_6	C_7
$C_0\ C_1\ C_2\ C_3/O_4\ D_5\ D_6\ D_7$	10	0x2d	C_0	C_1	C_2	C_3	O_4	D_5	D_6	D_7
$C_0\ C_1\ C_2\ C_3/S_4\ D_5\ D_6\ D_7$	10	0x33	C_0	C_1	C_2	C_3		D_5	D_6	D_7
$O_0\ D_1\ D_2\ D_3/S_4\ D_5\ D_6\ D_7$	10	0x66	D_1	D_2	D_3	O_0		D_5	D_6	D_7
$O_0\ D_1\ D_2\ D_3/O_4\ D_5\ D_6\ D_7$	10	0x55	D_1	D_2	D_3	O_0	O_4	D_5	D_6	D_7
$S_0\ D_1\ D_2\ D_3/D_4\ D_5\ D_6\ D_7$	10	0x78	D_1	D_2	D_3	D_4		D_5	D_6	D_7
$O_0\ D_1\ D_2\ D_3/C_4\ C_5\ C_6\ C_7$	10	0x4b	D_1	D_2	D_3	O_0	C_4	C_5	C_6	C_7
$T_0\ C_1\ C_2\ C_3/C_4\ C_5\ C_6\ C_7$	10	0x87		C_1	C_2	C_3	C_4	C_5	C_6	C_7
$D_0\ T_1\ C_2\ C_3/C_4\ C_5\ C_6\ C_7$	10	0x99	D_0		C_2	C_3	C_4	C_5	C_6	C_7
$D_0\ D_1\ T_2\ C_3/C_4\ C_5\ C_6\ C_7$	10	0xaa	D_0	D_1		C_3	C_4	C_5	C_6	C_7
$D_0\ D_1\ D_2\ T_3/C_4\ C_5\ C_6\ C_7$	10	0xb4	D_0	D_1	D_2		C_4	C_5	C_6	C_7
$D_0\ D_1\ D_2\ D_3/T_4\ C_5\ C_6\ C_7$	10	0xcc	D_0	D_1	D_2	D_3		C_5	C_6	C_7
$D_0\ D_1\ D_2\ D_3/D_4\ T_5\ C_6\ C_7$	10	0xd2	D_0	D_1	D_2	D_3	D_4		C_6	C_7
$D_0\ D_1\ D_2\ D_3/D_4\ D_5\ T_6\ C_7$	10	0xe1	D_0	D_1	D_2	D_3	D_4	D_5		C_7
$D_0\ D_1\ D_2\ D_3/D_4\ D_5\ D_6\ T_7$	10	0xff	D_0	D_1	D_2	D_3	D_4	D_5	D_6	

Figure 49–7—64B/66B block formats

49.2.4.5 Ordered sets

Ordered sets are used to extend the ability to send control and status information over the link such as remote fault and local fault status. Ordered sets consist of a control character followed by three data characters. Ordered sets always begin on the first octet of the XGMII. 10 Gigabit Ethernet uses one kind of ordered set: the sequence ordered set (see 46.3.4). The sequence ordered set control character is denoted /Q/. An

additional ordered set, the signal ordered set, has been reserved and it begins with another control code. The 4-bit O field encodes the control code. See Table 49–1 for the mappings.

49.2.4.6 Valid and invalid blocks

A block is invalid if any of the following conditions exists:

a) The sync field has a value of 00 or 11.
b) The block type field contains a reserved value.
c) Any control character contains a value not in Table 49–1.
d) Any O code contains a value not in Table 49–1.
e) The set of eight XGMII characters does not have a corresponding block format in Figure 49–7.

49.2.4.7 Idle (/I/) and Low Power Idle (/LI/)

Idle control characters (/I/) are transmitted when idle control characters are received from the XGMII. Idle characters may be added or deleted by the PCS to adapt between clock rates. /I/ insertion and deletion shall occur in groups of 4. /I/s may be added following idle or ordered sets. They shall not be added while data is being received. When deleting /I/s, the first four characters after a /T/ shall not be deleted.

To communicate LPI, LPI control character /LI/ is sent continuously in place of /I/. LPI control characters are transmitted when LPI control characters are received from the XGMII. LPI characters may be added or deleted by the PCS to adapt between clock rates in a similar manner to idle control characters. /LI/ insertion and deletion shall occur in groups of four. /LI/s may only be added following other LPI characters.

Table 49–1—Control codes

Control character	Notation	XGMII Control Code	10GBASE-R Control Code	10GBASE-R O Code	8B/10B Code[a]
idle	/I/	0x07	0x00		K28.0 or K28.3 or K28.5
LPI	/LI/	0x06	0x06		K28.0 with D20.5 in one row or K28.3 or K28.5 with D20.5 in one row[b]
start	/S/	0xfb	Encoded by block type field		K27.7
terminate	/T/	0xfd	Encoded by block type field		K29.7
error	/E/	0xfe	0x1e		K30.7
Sequence ordered set	/Q/	0x9c	Encoded by block type field plus O code	0x0	K28.4
reserved0	/R/[c]	0x1c	0x2d		K28.0
reserved1		0x3c	0x33		K28.1
reserved2	/A/	0x7c	0x4b		K28.3
reserved3	/K/	0xbc	0x55		K28.5

Table 49–1—Control codes *(continued)*

Control character	Notation	XGMII Control Code	10GBASE-R Control Code	10GBASE-R O Code	8B/10B Code[a]
reserved4		0xdc	0x66		K28.6
reserved5		0xf7	0x78		K23.7
Signal ordered set[d]	/Fsig/	0x5c	Encoded by block type field plus O code	0xF	K28.2

[a]For information only. The 8B/10B code is specified in Clause 36. Usage of the 8B/10B code for 10 Gb/s operation is specified in Clause 48.
[b]See 48.2.4.2.
[c]The codes for /A/, /K/, and /R/ are used on the XAUI interface to signal idle. They are not present on the XGMII when no errors have occurred, but certain bit errors cause the XGXS to send them on the XGMII.
[d]Reserved for INCITS T11 Fibre Channel use.

49.2.4.8 Start (/S/)

The start control character (/S/) indicates the start of a packet. This delimiter is only valid on the first octet of the XGMII (TXD<0:7> and RXD<0:7>). Receipt of an /S/ on any other octet of TxD indicates an error. Block type field values implicitly encode an /S/ as the fifth or first character of the block. These are the only characters of a block on which a start can occur.

49.2.4.9 Terminate (/T/)

The terminate control character (/T/) indicates the end of a packet. Since packets may be any length, the /T/ can occur on any octet of the XGMII interface and within any character of the block. The location of the /T/ in the block is implicitly encoded in the block type field. A valid end of packet occurs when a block containing a /T/ is followed by a control block that does not contain a /T/.

49.2.4.10 ordered set (/O/)

The ordered set control characters (/O/) indicate the start of an ordered set. There are two kinds of ordered sets: the sequence ordered set and the signal ordered set (which is reserved). When it is necessary to designate the control character for the sequence ordered set specifically, /Q/ will be used. /O/ is only valid on the first octet of the XGMII. Receipt of an /O/ on any other octet of TXD indicates an error. Block type field values implicitly encode an /O/ as the first or fifth character of the block. The 4-bit O code encodes the specific /O/ character for the ordered set.

Sequence ordered sets may be deleted by the PCS to adapt between clock rates. Such deletion shall only occur when two consecutive sequence ordered sets have been received and shall delete only one of the two. Only Idles may be inserted for clock compensation. Signal ordered sets are not deleted for clock compensation.

49.2.4.11 Error (/E/)

The /E/ is sent whenever an /E/ is received. It is also sent when invalid blocks are received. The /E/ allows physical sublayers such as the XGXS and PCS to propagate received errors. See R_BLOCK_TYPE and T_BLOCK_TYPE function definitions in 49.2.13.2.3 for further information.

49.2.5 Transmit process

The transmit process generates blocks based upon the TXD<31:0> and TXC<3:0> signals received from the XGMII. Two XGMII data transfers are encoded into each block. It takes 4.125 PMA_UNITDATA or WIS_UNITDATA transfers to send a block of data. Therefore, if the PCS is connected to an XGMII and PMA sublayer where the ratio of their transfer rates is exactly 16:33, then the transmit process does not need to perform rate adaptation. Where the XGMII and PMA sublayer data rates are not synchronized to that ratio, the transmit process will need to insert idles, delete idles, or delete sequence ordered sets to adapt between the rates. The WIS data rate is always slower than the XGMII data rate and a PCS connected to a WIS must delete idles or sequence ordered sets to adapt between rates.

The transmit process generates blocks as specified in the transmit process state diagram. The contents of each block are contained in a vector tx_coded<65:0>, which is passed to the scrambler. tx_coded<1:0> contains the sync header and the remainder of the bits contain the block payload.

49.2.6 Scrambler

The payload of the block is scrambled with a self-synchronizing scrambler. The scrambler shall produce the same result as the implementation shown in Figure 49–8. This implements the scrambler polynomial:[8]

$$G(x) = 1 + x^{39} + x^{58} \tag{49–1}$$

There is no requirement on the initial value for the scrambler. The scrambler is run continuously on all payload bits. The sync header bits bypass the scrambler.

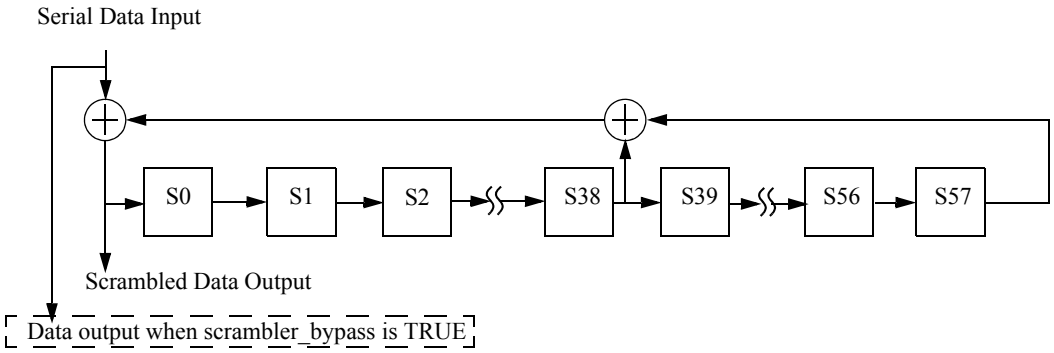

NOTE—Scrambler_bypass is only required to support EEE capability.

Figure 49–8—Scrambler

To aid block synchronization in the receiver for EEE capability when Clause 74 FEC is in use, when scrambler_bypass is TRUE the PCS shall pass the unscrambled data from the scrambler input rather than the scrambled data from the scrambler output. The scrambler shall continue to advance normally.

[8]The convention here, which considers the most recent bit into the scrambler to be the lowest order term, is consistent with most references and with other scramblers shown in this standard. Some references consider the most recent bit into the scrambler to be the highest order term and would therefore identify this as the inverse of the polynomial in Equation (49–1). In case of doubt, note that the conformance requirement is based on the representation of the scrambler in the figure rather than the polynomial equation.

49.2.7 Gearbox

The gearbox adapts between the 66-bit width of the blocks and the 16-bit width of the PMA or WIS interface. It receives the 66-bit blocks. When the transmit channel is operating in normal mode, the gearbox sends 16 bits of transmit data at a time via WIS_UNITDATA.request or PMA_UNITDATA.request primitives. The UNITDATA.request primitives are fully packed with bits. For example, if one block happened to start with the sync header on bits 0 and 1 of a PMA_UNITDATA.request, then the last two bits of that block would be on bits 0 and 1 of a PMA_UNITDATA.request and the next block would begin with a sync header on bits 2 and 3 of that PMA_UNITDATA.request. When a PMA_UNITDATA.request or WIS_UNITDATA.request contains bits from two blocks, then the bits from the first block shall be placed in the lowest numbered bits of tx_data-group<15:0>. The bits shall be packed into the tx_data-group in sequence with the lowest numbered bit of the block going into the lowest numbered bit of the part of tx_data-group<15:0> bits containing bits from that block (see Figure 49–5).

The gearbox functionality is necessary when the optional PMA compatibility interface, XSBI, is implemented since that interface passes data over a 16-bit wide path. It is also necessary when connecting to a WIS since the WIS processes the data stream with 8-bit granularity. When neither the WIS nor the XSBI is implemented, the internal data-path width between the PCS and PMA is an implementation choice. Depending on the path width, the gearbox functionality may not be necessary.

49.2.8 Test-pattern generators

When the transmit channel is operating in test-pattern mode, it sends 16 bits of test pattern at a time via PMA_UNITDATA.request primitives. When the PCS allows direct connection to the PMA, the test-pattern generator shall be implemented. The test-pattern generator does not apply to a PCS, which only supports connection to the WIS. A PCS which supports both WIS and direct PMA attachment may reject or allow an attempt to activate a transmit test-pattern mode when a WIS is attached.

There are two types of required transmit test patterns: square wave and pseudo-random. The square wave pattern is intended to aid in conducting certain transmitter tests. It is not intended for receiver tests and the receiver is not expected to receive this test pattern. The pseudo-random test-pattern mode is suitable for receiver tests and for certain transmitter tests. There is an optional PRBS9 transmit test pattern that may be used for some transmitter tests. There is also an optional PRBS31 test pattern, which may be used for some transmit and receiver tests. When this option is supported, both the PRBS31 test-pattern generator and the PRBS31 test-pattern checker shall be provided. See 52.9 and 68.6 for recommendations on the appropriate pattern for tests.

When square wave pattern is selected, the PCS will send a repeating pattern of n ones followed by n zeros where n may be any number between 4 and 11 inclusive. The value of n is an implementation choice.

When pseudo-random pattern is selected, the test pattern is generated by the scrambler using the seeds loaded through the MDIO registers and the selected data pattern. The scrambler is loaded with a seed or its inverse at the start of a block every 128 blocks. The seeds are loaded in the following pattern:

Seed A
Seed A Invert
Seed B
Seed B Invert

Invert indicates that the seed is inverted for that load. Either 64 zeros or the 64-bit encoding for two Local Fault ordered sets can be selected as the data pattern. After loading Seed A or Seed B, the scrambler input shall be driven with the data pattern. After loading Seed A Invert or Seed B Invert, the scrambler input shall be driven with the inverse of the data pattern. While in pseudo-random test-pattern mode, the sync headers will be the control sync header, 10. Thus the pseudo-random test pattern is a series of blocks with the control

sync header and a pseudo-random payload. The characteristics of the pseudo-random test pattern can be varied by varying the seed values and data input.

When the optional PRBS31 mode is selected, the PRBS31 pattern generator sends 16 bits of PRBS31 test pattern at a time via PMA_UNITDATA.request primitives. The PRBS31 test pattern is the output of a Pseudo-Random Bit Sequence of order 31 (PRBS31) generator. The PRBS31 pattern generator shall produce the same result as the implementation shown in Figure 49–9. This implements the inverted version of the bit stream produced by the polynomial shown in Equation (49–2):[9]

$$G(x) = 1 + x^{28} + x^{31} \tag{49–2}$$

The initial value of the PRBS31 pattern generator shall not be all zeros. It may be any other value.

The optional PRBS9 pattern is defined in Table 68–6.

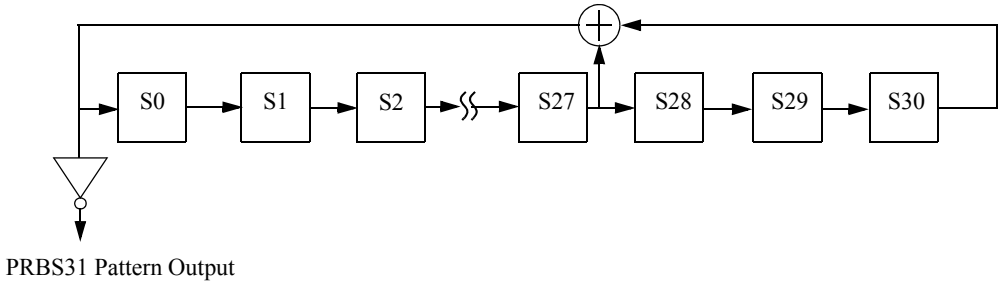

Figure 49–9—PRBS31 pattern generator

49.2.9 Block synchronization

When the receive channel is operating in normal mode, the block synchronization function receives data via 16-bit PMA_UNITDATA.request or WIS_UNITDATA.request primitives. It shall form a bit stream from the primitives by concatenating requests with the bits of each primitive in order from rx_data-group<0> to rx_data-group<15> (see Figure 49–6). It obtains lock to the 66-bit blocks in the bit stream using the sync headers and outputs 66-bit blocks. Lock is obtained as specified in the block lock state diagram shown in Figure 49–14.

If EEE is not supported then block_lock is identical to rx_block_lock. Otherwise the relationship between block_lock and rx_block_lock is given by Figure 49–13.

49.2.10 Descrambler

The descrambler processes the payload to reverse the effect of the scrambler using the same polynomial. It shall produce the same result as the implementation shown in Figure 49–10.

49.2.11 Receive process

The receive process decodes blocks to produce RXD<31:0> and RXC<3:0> for transmission to the XGMII. Two XGMII data transfers are decoded from each block. Where the XGMII and PMA sublayer data rates are not synchronized to a 16:33 ratio, the receive process will insert idles, delete idles, or delete sequence

[9]The convention here, which considers the most recent bit into the PRBS generator to be the lowest order term, is consistant with most references and with other scramblers shown in this standard. Some references consider the most recent bit into the PRBS generator to be the highest order term and would therefore identify this as the inverse of the polynomial in Equation (49–2). In case of doubt, note that the conformance requirement is based on the representation in the figure rather than the polynomial equation.

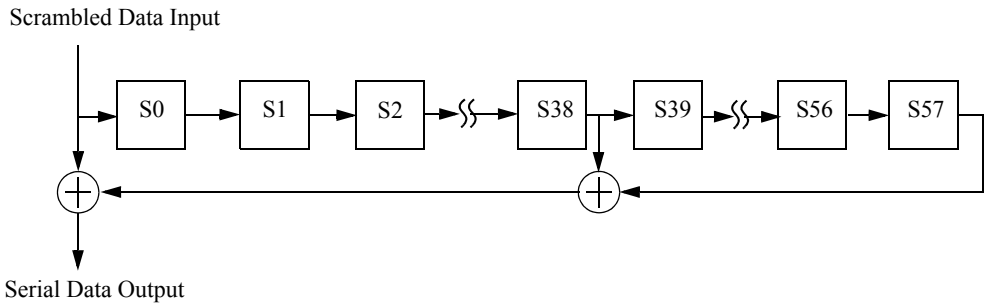

Figure 49–10—Descrambler

ordered sets to adapt between rates. The WIS data rate is always slower than the XGMII data rate and a PCS connected to a WIS will insert idles to adapt between the rates.

The receive process decodes blocks as specified in the receive state diagram shown in Figure 49–17.

49.2.12 Test-pattern checker

When the PCS allows direct connection to the PMA, the pseudo-random test-pattern checker shall be implemented. The pseudo-random test-pattern checker does not apply to a PCS, which only supports connection to the WIS. When the PRBS31 test-pattern option is supported, the PRBS31 receive test-pattern checker shall be implemented. A PCS that supports both WIS and direct PMA attachment may reject or allow an attempt to activate a receive test-pattern mode when a WIS is attached.

When the receive channel is operating in pseudo-random test-pattern mode, the pseudo-random test-pattern checker checks the bits received via PMA_UNITDATA.indication primitives.

The pseudo-random test-pattern checker utilizes the lock state diagram and the descrambler operating as they do during normal data reception. The hi_ber state diagram is disabled during receive test-pattern mode. When block_lock is true and the pseudo-random receive test-pattern mode is active, the pseudo-random test-pattern checker observes the output from the descrambler. When the output of the descrambler is the data pattern or its inverse, a match is detected. Since the transmitter's scrambler is loaded with a seed value every 128 blocks and the receiver's descrambler is running normally, a mismatch will be detected once every 128 blocks in the absence of errors. The pseudo-random test-pattern checker will count 128-block windows. When operating in pseudo-random test pattern, the test-pattern error counter counts blocks with a mismatch corrected to remove the effect of loading a new seed. The first block with a mismatch in a window shall not increment the test-pattern error counter. Any subsequent block with a mismatch in a window indicates an error and shall increment the test-pattern error counter.

When the receive channel is operating in PRBS31 test mode, the PRBS31 pattern checker checks the bits received in PMA_UNITDATA.indication primitives relative to the PRBS31 test pattern.

The PRBS31 pattern error checker is self-synchronizing. It compares each bit received to the result of the PRBS31 generator based on the prior 31 bits received. It shall produce the same result as the implementation shown in Figure 49–11. When no errors occur, the PRBS31 pattern error signal will be zero. When an isolated bit error occurs, it will cause the PRBS31 pattern error signal to go high three times; once when it is

received and once when it is at each tap. The test-pattern error counter shall increment once for each bit time that the PRBS31 pattern error signal is high.

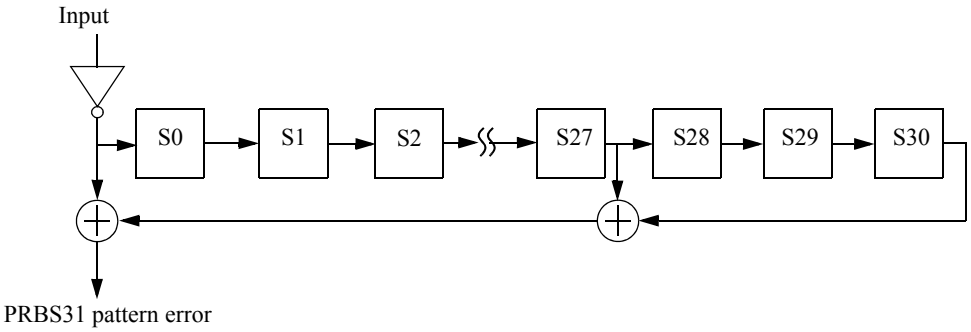

Figure 49–11—PRBS31 pattern checker

Note that the test-pattern error counter behavior is dependent on the test-pattern mode. In pseudo-random test mode, it is counting block errors. In PRBS31 test mode, it is counting bit errors at the PRBS31 pattern checker output.

49.2.13 Detailed functions and state diagrams

49.2.13.1 State diagram conventions

The body of this subclause is comprised of state diagrams, including the associated definitions of variables, constants, and functions. Should there be a discrepancy between a state diagram and descriptive text, the state diagram prevails.

The notation used in the state diagrams follows the conventions of 21.5. State diagram timers follow the conventions of 14.2.3.2. The notation ++ after a counter or integer variable indicates that its value is to be incremented.

49.2.13.2 State variables

49.2.13.2.1 Constants

EBLOCK_R<71:0>
: 72 bit vector to be sent to the XGMII interface containing /E/ in all the eight character locations.

EBLOCK_T<65:0>
: 66 bit vector to be sent to the PMA containing /E/ in all the eight character locations.

LBLOCK_R<71:0>
: 72 bit vector to be sent to the XGMII interface containing two Local Fault ordered sets. The Local Fault ordered set is defined in 46.3.4.

LBLOCK_T<65:0>
: 66 bit vector to be sent to the PMA containing two Local Fault ordered sets.

49.2.13.2.2 Variables

ber_test_sh
: Boolean variable that is set true when a new sync header is available for testing and false when BER_TEST_SH state is entered. A new sync header is available for testing when the Block Sync process has accumulated enough bits from the PMA or the WIS to evaluate the header of the next block

block_lock
: Boolean variable that is set true when receiver acquires block delineation

NOTE—If the EEE capability is supported, then this variable is affected by the LPI receive state diagram. If the EEE capability is not supported then this variable is identical to rx_block_lock controlled by the lock state diagram.

hi_ber
: Boolean variable which is asserted true when the ber_cnt exceeds 16 indicating a bit error ratio $>10^{-4}$

reset
: Boolean variable that controls the resetting of the PCS. It is true whenever a reset is necessary including when reset is initiated from the MDIO, during power on, and when the MDIO has put the PCS into low-power mode.

r_test_mode
: Boolean variable that is asserted true when the receiver is in test-pattern mode.

rx_block_lock
: Variable used by the lock state diagram to reflect the status of the code-group delineation. This variable is set TRUE when the receiver acquires block delineation.

rx_coded<65:0>
: Vector containing the input to the 64B/66B decoder. The format for this vector is shown in Figure 49–7. The leftmost bit in the figure is rx_coded<0> and the rightmost bit is rx_coded<65>.

rx_raw<71:0>
: Vector containing two successive XGMII transfers. RXC<0> through RXC<3> for the first transfer are placed in rx_raw<0> through rx_raw<3>, respectively. RXC<0> through RXC<3> for the second transfer are placed in rx_raw<4> through rx_raw<7>, respectively. RXD<0> through RXD<31> for the first transfer are placed in rx_raw<8> through rx_raw<39>, respectively. RXD<0> through RXD<31> for the second transfer are placed in rx_raw<40> through rx_raw<71>, respectively.

sh_valid
: Boolean indication that is set true if received block rx_coded has valid sync header bits. That is, sh_valid is asserted if rx_coded<0> ≠ rx_coded<1> and de-asserted otherwise.

signal_ok
: Boolean variable that is set based on the most recently received value of PMA_SIGNAL.indication(SIGNAL_OK) or WIS_SIGNAL.indication(SIGNAL_OK). It is true if the value was OK and false if the value was FAIL.

slip_done
: Boolean variable that is asserted true when the SLIP requested by the Block Lock state diagram has been completed indicating that the next candidate block sync position can be tested.

test_sh
: Boolean variable that is set true when a new sync header is available for testing and false when TEST_SH state is entered. A new sync header is available for testing when the Block Sync process has accumulated enough bits from the PMA or the WIS to evaluate the header of the next block

tx_coded<65:0>
: Vector containing the output from the 64B/66B encoder. The format for this vector is shown in Figure 49–7. The leftmost bit in the figure is tx_coded<0> and the rightmost bit is tx_coded<65>.

tx_raw<71:0>
: Vector containing two successive XGMII transfers. TXC<0> through TXC<3> for the first transfer are placed in tx_raw<0> through tx_raw<3>, respectively. TXC<0> through TXC<3> for the second transfer are placed in tx_raw<4> through tx_raw<7>, respectively. TXD<0> through TXD<31> for the first transfer are placed in tx_raw<8> through tx_raw<39>, respectively. TXD<0> through TXD<31> for the second transfer are placed in tx_raw<40> through tx_raw<71>, respectively.

The following variables are used only for the EEE capability:

energy_detect
: A Boolean variable sent from the PMD that is set to TRUE when signal energy is detected at the receiver and is set to FALSE otherwise

rx_lpi_active
: A Boolean variable that is set to TRUE when the receiver is in a low power state and set to FALSE when it is in an active state and capable of receiving data.

rx_mode
: A variable set to QUIET while the receiver is in the RX_QUIET state and set to DATA otherwise

tx_mode
: A variable set to QUIET when the transmitter is in the TX_QUIET state, set to ALERT when the transmitter is in the TX_ALERT state, and set to DATA otherwise. When set to QUIET, the PMD disables the transmitter as described in 72.6.5. When set to ALERT, the PMD transmits a repeating pattern of eight ones and eight zeroes as described in 72.6.2. When set to DATA the PMD passes data as normal.

scrambler_bypass
: This Boolean variable is used to bypass the Tx PCS scrambler in order to assist rapid synchronization following low power idle. When set to TRUE, the PCS will pass the unscrambled data from the scrambler input rather than the scrambled data from the scrambler output. The scrambler will continue to operate normally, shifting input data into the delay line. When scrambler_bypass is set to FALSE the PCS will pass scrambled data from the scrambler output.

scr_bypass_enable
: A Boolean variable used to indicate to the transmit LPI state diagram that the scrambler bypass option is required. The PHY shall set scr_bypass_enable = TRUE if Clause 74 FEC is in use. The PHY shall set scr_bypass_enable = FALSE if this FEC is not in use.

49.2.13.2.3 Functions

DECODE(rx_coded<65:0>)
: Decodes the 66-bit vector returning rx_raw<71:0> that is sent to the XGMII. The DECODE function shall decode the block as specified in 49.2.4.

ENCODE(tx_raw<71:0>)
: Encodes the 72-bit vector returning tx_coded<65:0> of which tx_coded<63:0> is sent to the scrambler. The two high order sync bits bypass the scrambler. The ENCODE function shall encode the block as specified in 49.2.4.

R_BLOCK_TYPE = {C, S, T, D, E, LI}
: This function classifies each 66-bit rx_coded vector as belonging to one of the following types depending on its contents.

  Values: C; The vector contains a sync header of 10 and one of the following:
  a) A block type field of 0x1e and eight valid control characters other than /E/; and, if the EEE capability is supported, zero or four of the characters are /LI/;
  b) A block type field of 0x2d or 0x4b, a valid O code, and four valid control characters;
  c) A block type field of 0x55 and two valid O codes.

  LI; For EEE capability, the LI type is supported where the vector contains a sync header of 10, a block type field of 0x1e, and eight control characters of 0x06 (/LI/).

  S; The vector contains a sync header of 10 and one of the following:
  a) A block type field of 0x33 and four valid control characters;
  b) A block type field of 0x66 and a valid O code;
  c) A block type field of 0x78.

T; The vector contains a sync header of 10, a block type field of 0x87, 0x99, 0xaa, 0xb4, 0xcc, 0xd2, 0xe1 or 0xff and all control characters are valid.
D; The vector contains a sync header of 01.
E; The vector does not meet the criteria for any other value.

A valid control character is one containing a 10GBASE-R control code specified in Table 49–1. A valid O code is one containing an O code specified in Table 49–1.
NOTE—A PCS that does not support EEE classifies vectors containing one or more /LI/ control characters as type E.

R_TYPE(rx_coded<65:0>)
Returns the R_BLOCK_TYPE of the rx_coded<65:0> bit vector.

R_TYPE_NEXT
Prescient end of packet check function. It returns the R_BLOCK_TYPE of the rx_coded vector immediately following the current rx_coded vector.

SLIP
Causes the next candidate block sync position to be tested. The precise method for determining the next candidate block sync position is not specified and is implementation dependent. However, an implementation shall ensure that all possible bit positions are evaluated.

T_BLOCK_TYPE = {C, S, T, D, E, LI}
This function classifies each 72-bit tx_raw vector as belonging to one of the following types depending on its contents.

Values: C; The vector contains one of the following:
a) eight valid control characters other than /O/, /S/, /T/ and /E/; and, if the EEE capability is supported, zero or four of the characters are /LI/;
b) one valid ordered set and four valid control characters other than /O/, /S/ and /T/;
c) two valid ordered sets.
LI; For EEE capability, this vector contains eight /LI/ characters.
S; The vector contains an /S/ in its first or fifth character, any characters before the S character are valid control characters other than /O/, /S/ and /T/ or form a valid ordered set, and all characters following the /S/ are data characters.
T; The vector contains a /T/ in one of its characters, all characters before the /T/ are data characters, and all characters following the /T/ are valid control characters other than /O/, /S/ and /T/.
D; The vector contains eight data characters.
E; The vector does not meet the criteria for any other value.

A tx_raw character is a control character if its associated TXC bit is asserted. A valid control character is one containing an XGMII control code specified in Table 49–1. A valid ordered set consists of a valid /O/ character in the first or fifth characters and data characters in the three characters following the /O/. A valid /O/ is any character with a value for O code in Table 49–1.
NOTE—A PCS that does not support EEE, classifies vectors containing one or more /LI/ control characters as type E.

T_TYPE(tx_raw<71:0>)
Returns the T_BLOCK_TYPE of the tx_raw<71:0> bit vector.

49.2.13.2.4 Counters

ber_cnt
Count up to a maximum of 16 of the number of invalid sync headers within the current 125 μs period.

sh_cnt
: Count of the number of sync headers checked within the current 64 block window.

sh_invalid_cnt
: Count of the number of invalid sync headers within the current 64 block window.

The following counter is used only for the EEE capability.

wake_error_counter
: A counter that is incremented each time that the LPI receive state diagram enters the RX_WTF state indicating that a wake time fault has been detected. The counter is reflected in register 3.22 (see 45.2.3.10).

49.2.13.2.5 Timers

State diagram timers follow the conventions of 14.2.3.2.

125us_timer
: Timer that is triggered every 125 μs +1%, –25%.

The following timers are used only for the EEE capability:

one_us_timer
: A timer used to count approximately 1 μs intervals. The timer terminal count is set to T_{1U}. When the timer reaches terminal count it will set the one_us_timer_done = TRUE.

rx_tq_timer
: This timer is started when the PCS receiver enters the RX_SLEEP state. The timer terminal count is set to T_{QR}. When the timer reaches terminal count it will set the rx_tq_timer_done = TRUE.

rx_tw_timer
: This timer is started when the PCS receiver enters the RX_WAKE state. The timer terminal count shall be set to a value no larger than the maximum value given for T_{WR} in Table 49–3. When the timer reaches terminal count it will set the rx_tw_timer_done = TRUE.

rx_wf_timer
: This timer is started when the PCS receiver enters the RX_WTF state, indicating that the receiver has encountered a wake time fault. The rx_wf_timer allows the receiver an additional period in which to synchronize or return to the QUIET state before a link failure is indicated. The timer terminal count is set to T_{WTF}. When the timer reaches terminal count it will set the rx_wf_timer_done = TRUE.

tx_ts_timer
: This timer is started when the PCS transmitter enters the TX_SLEEP state. The timer terminal count is set to T_{SL}. When the timer reaches terminal count it will set the tx_ts_timer_done = TRUE.

tx_tq_timer
: This timer is started when the PCS transmitter enters the TX_QUIET state. The timer terminal count is set to T_{QL}. When the timer reaches terminal count it will set the tx_tq_timer_done = TRUE.

tx_tw_timer
: This timer is started when the PCS transmitter enters the TX_WAKE state. The timer terminal count is set to T_{WL}. When the timer reaches terminal count it will set the tx_tw_timer_done = TRUE.

49.2.13.3 State diagrams

The Lock state diagram shown in Figure 49–14 determines when the PCS has obtained lock to the received data stream. The BER Monitor state diagram shown in Figure 49–15 monitors the received signal for high bit error ratio.

The Transmit state diagram shown in Figure 49–16 controls the encoding of transmitted blocks. It makes exactly one transition for each transmit block processed. Though the Transmit state diagram sends Local Fault ordered sets when reset is asserted, the scrambler and gearbox may not be operational during reset. Thus, the Local Fault ordered sets may not appear on the WIS or PMA service interface.

The Receive state diagram shown in Figure 49–17 controls the decoding of received blocks. It makes exactly one transition for each receive block processed.

The PCS shall perform the functions of Lock, BER Monitor, Transmit, and Receive as specified in these state diagrams, including the optional EEE capability if implemented.

49.2.13.3.1 LPI state diagrams

A PCS that supports the EEE capability shall implement the LPI transmit and receive processes as shown in Figure 49–12 and Figure 49–13. The transmit LPI state diagram controls tx_mode, which disables the transmitter when it is set to QUIET. The receive LPI state diagram controls block_lock during LPI and signals the end of LPI to the receive state diagram.

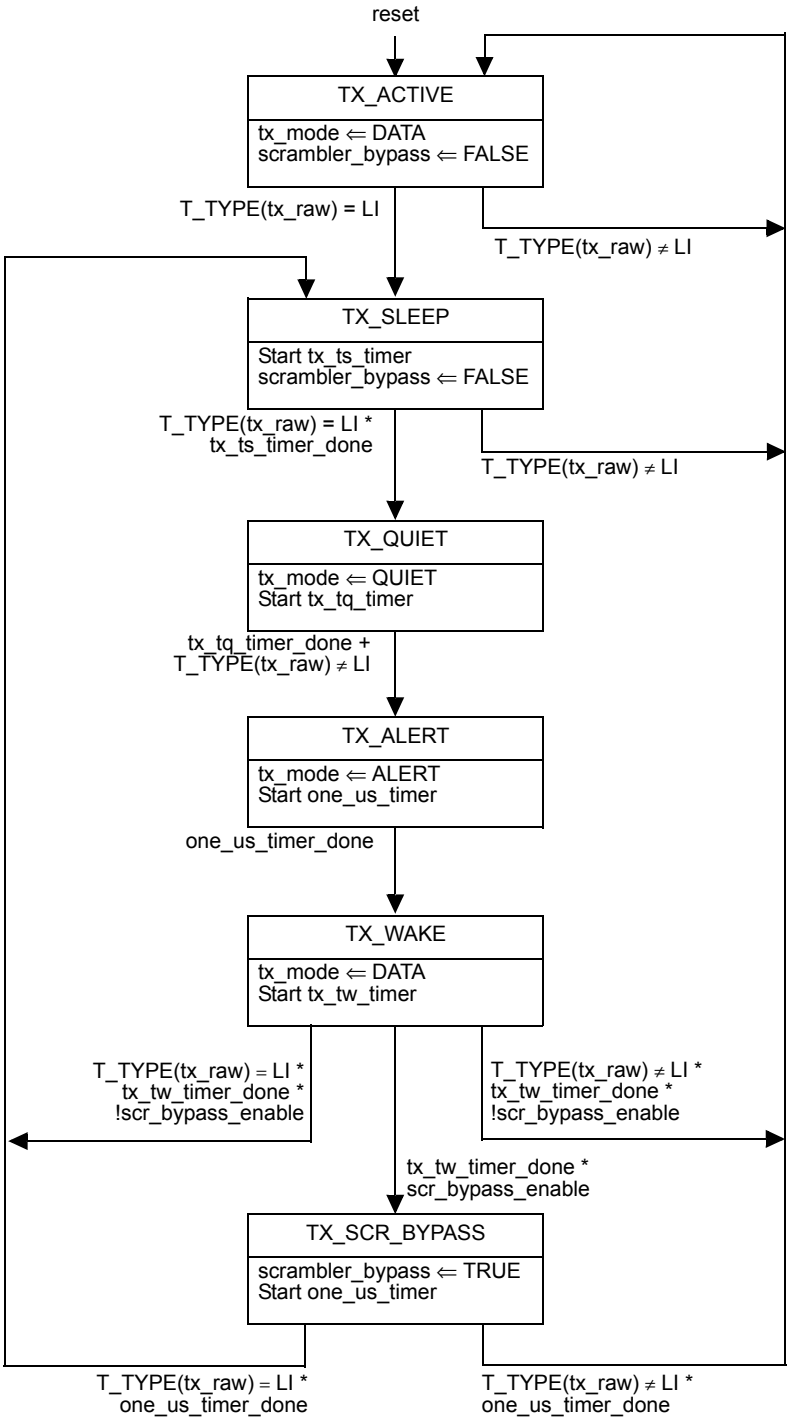

Figure 49–12—LPI Transmit state diagram

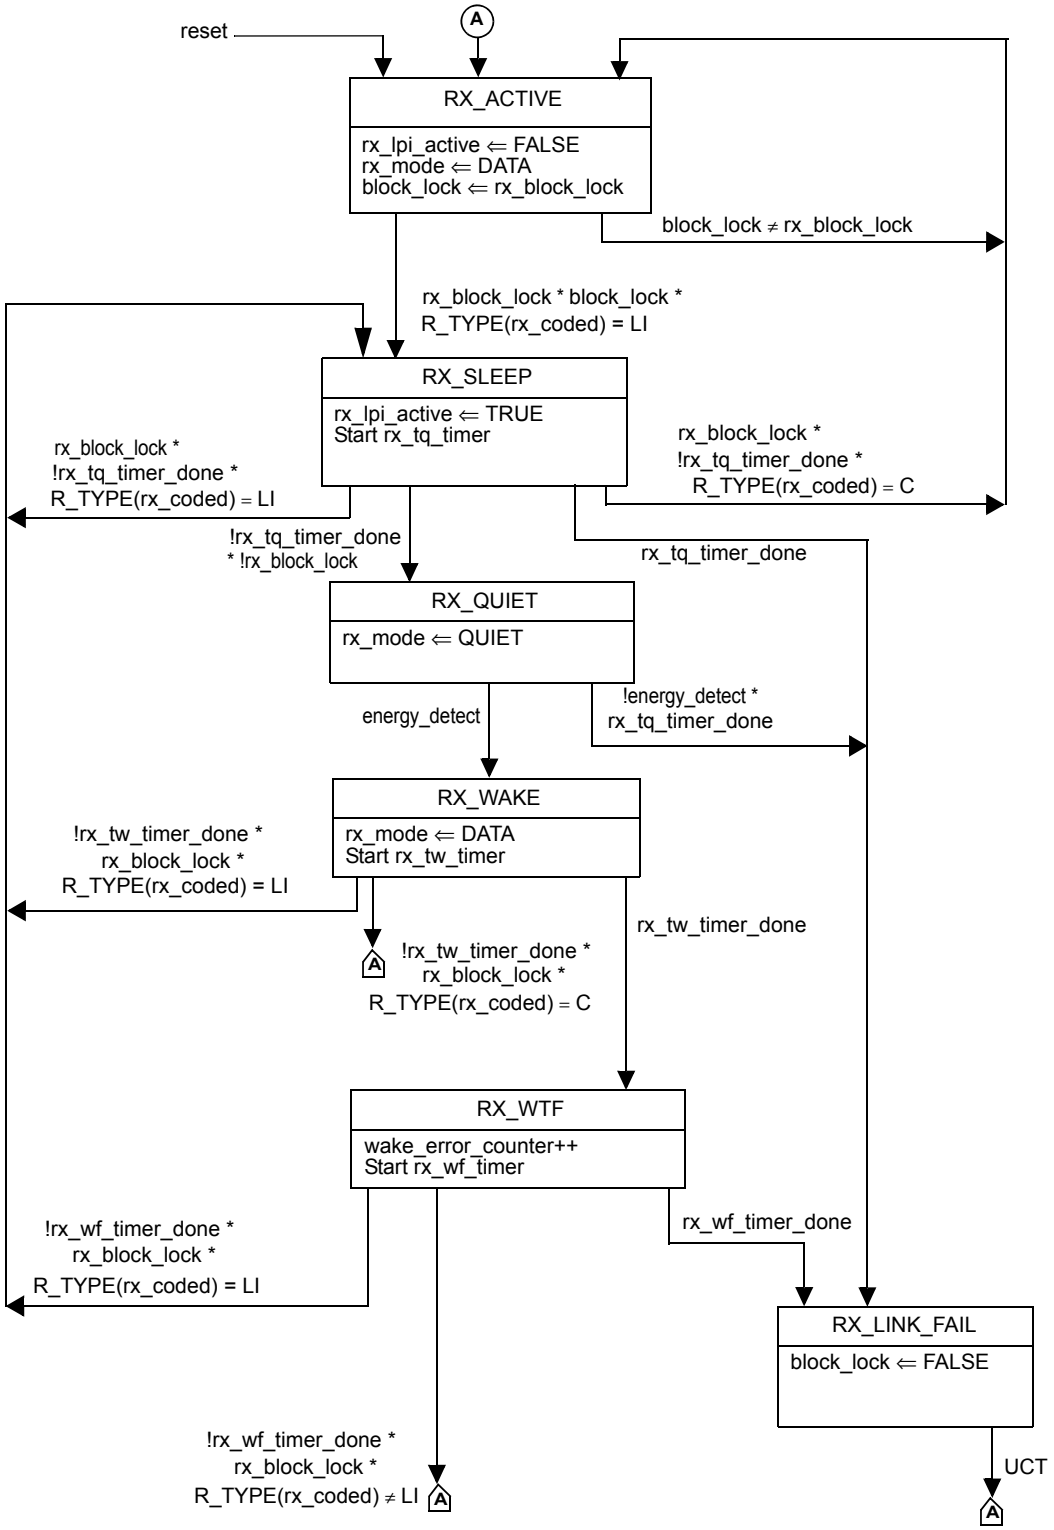

Figure 49–13—LPI Receive state diagram

Following a period of LPI, the receiver is required to achieve block synchronization within the wake-up time specified (see Figure 49–13). The implementation of the block synchronization state diagram should use techniques to ensure that block lock is achieved with minimal numbers of slip attempts. When the Clause 74 FEC is enabled, the receiver may use the knowledge that the link partner's transmitter will bypass the scrambler as part of the wake sequence. The idle sequence following this event will form a fixed pattern for the duration of the wake period.

The LPI functions shall use timer values for these state diagrams as shown in Table 49–2 for transmit and Table 49–3 for receive.

Table 49–2—Transmitter LPI timing parameters

Parameter	Description	Min	Max	Units
T_{SL}	Local Sleep Time from entering the TX_SLEEP state to when tx_mode is set to QUIET	4.9	5.1	µs
T_{QL}	Local Quiet Time from when tx_mode is set to QUIET to entry into the TX_ALERT state	1.7	1.8	ms
T_{WL}	Time spent in the TX_WAKE state	10.9	11.1	µs
T_{1U}	Time spent in the TX_ALERT and TX_SCR_BYPASS states	1.1	1.3	µs

Table 49–3—Receiver LPI timing parameters

Parameter	Description	Min	Max	Units
T_{QR}	The time the receiver waits for energy_detect to be set to TRUE while in the RX_SLEEP and RX_QUIET states before asserting receive fault	2	3	ms
T_{WR}	Time the receiver waits in the RX_WAKE state before indicating a wake time fault (when scr_bypass_enable = FALSE)	—	11.5	µs
T_{WR}	Time the receiver waits in the RX_WAKE state before indicating a wake time fault (when scr_bypass_enable = TRUE)	—	13.7	µs
T_{WTF}	Wake time fault recovery time	—	10	ms

49.2.14 PCS Management

The following objects apply to PCS management. If an MDIO Interface is provided (see Clause 45), they are accessed via that interface. If not, it is recommended that an equivalent access be provided.

49.2.14.1 Status

PCS_status:
: Indicates whether the PCS is in a fully operational state. It is only true if block_lock is true and hi_ber is false. This status is reflected in MDIO register 3.32.12. A latch low view of this status is reflected in MDIO register 3.1.2 and a latch high of the inverse of this status, Receive fault, is reflected in MDIO register 3.8.10.

block_lock:
: Indicates the state of the block_lock variable. This status is reflected in MDIO register 3.32.0. A latch low view of this status is reflected in MDIO register 3.33.15.

hi_ber:

Indicates the state of the hi_ber variable. This status is reflected in MDIO register 3.32.1. A latch high view of this status is reflected in MDIO register 3.33.14.

Rx LPI indication:

For EEE capability, this variable indicates the current state of the receive LPI function. This flag is set to TRUE (register bit set to one) when the LPI receive state diagram is in any state other than RX_ACTIVE. This status is reflected in MDIO register 3.1.8. A latch high view of this status is reflected in MDIO register 3.1.10 (Rx LPI received).

Tx LPI indication:

For EEE capability, this variable indicates the current state of the transmit LPI function. This flag is set to TRUE (register bit set to one) when the LPI transmit state diagram is in any state other than TX_ACTIVE. This status is reflected in MDIO register 3.1.9. A latch high view of this status is reflected in MDIO register 3.1.11 (Tx LPI received).

49.2.14.2 Counters

The following counters are reset to zero upon read and upon reset of the PCS. When they reach all ones, they stop counting. Their purpose is to help monitor the quality of the link.

ber_count:

6-bit counter that counts each time BER_BAD_SH state is entered. This counter is reflected in MDIO register bits 3.33.13:8. Note that this counter counts a maximum of 16 counts per 125 μs since the BER_BAD_SH can be entered a maximum of 16 times per 125 μs window.

errored_block_count:

8-bit counter. When the receiver is in normal mode, errored_block_count counts once for each time RX_E state is entered. This counter is reflected in MDIO register bits 3.33.7:0.

test_pattern_error_count:

16-bit counter. When the receiver is in test-pattern mode, the test_pattern_error_count counts errors as described in 49.2.12. This counter is reflected in MDIO register bits 3.43.15:0.

49.2.14.3 Test mode control

tx_test_mode:

Boolean variable controlling transmit channel operating mode. When false, the transmit channel operates in normal mode. When true, the transmit channel operates in test-pattern mode.

rx_test_mode:

Boolean variable controlling receive channel operating mode. When false, the receive channel operates in normal mode. When true, the receive channel operates in test-pattern mode.

49.2.14.4 Loopback

The PCS shall be placed in Loopback mode when the Loopback bit in MDIO register 3.0.14 is set to a logic one. In this mode, the PCS shall accept data on the transmit path from the XGMII and return it on the receive path to the XGMII. In addition, the PCS shall transmit a continuous stream of 0x00FF data words to the PMA or WIS sublayer, and shall ignore all data presented to it by the PMA or WIS sublayer.

49.2.15 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementers must conform to certain delay maxima, and that network planners and

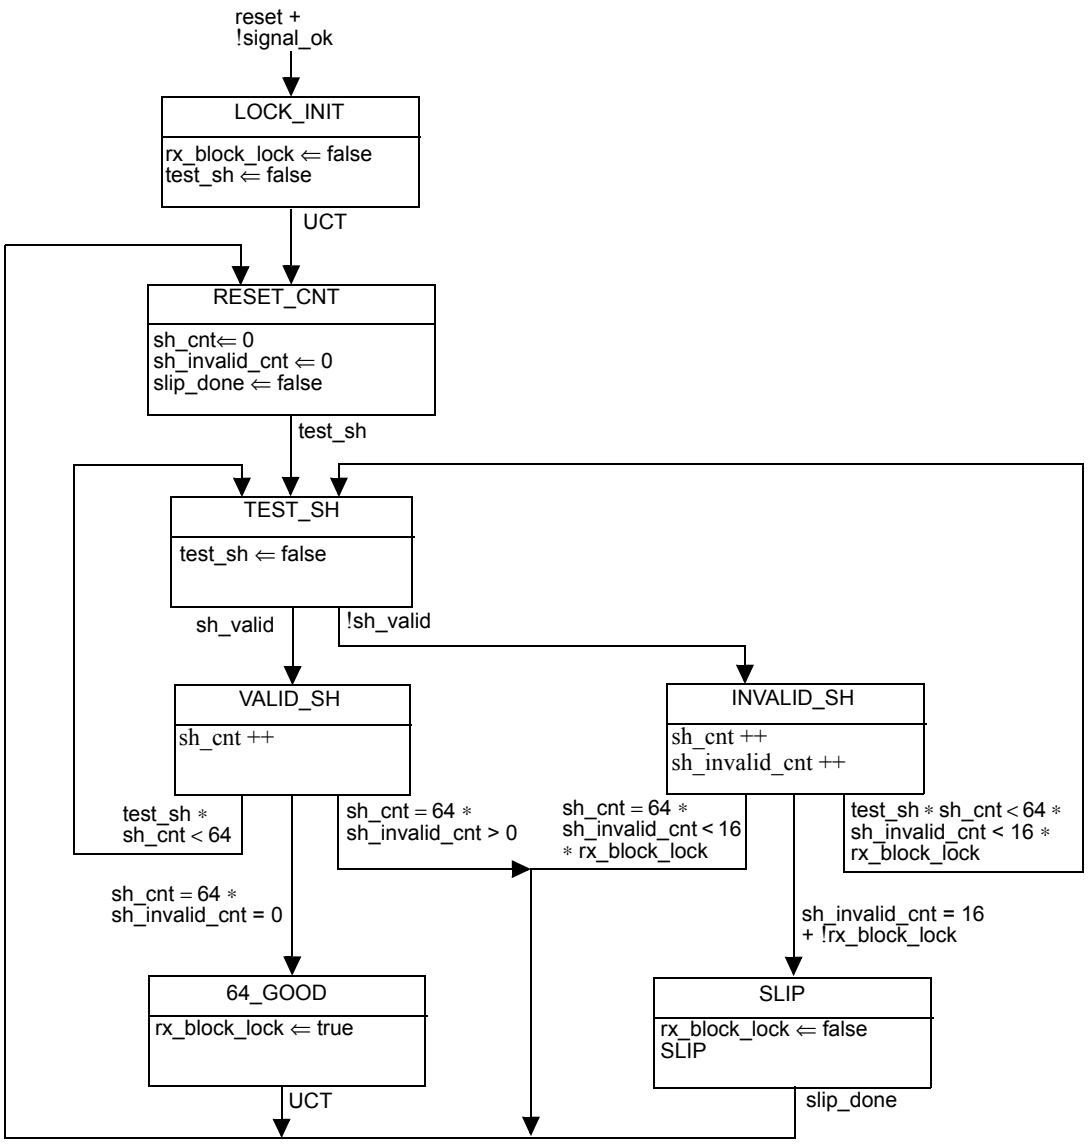

Figure 49–14—Lock state diagram

administrators conform to constraints regarding the cable topology and concatenation of devices. The sum of transmit and receive delay contributed by the 10GBASE-R PCS shall be no more than 3584 BT.

49.2.16 Auto-Negotiation for Backplane Ethernet

The following requirements apply to a PCS used with a 10GBASE-KR PMD. Support for the Auto-Negotiation process defined in Clause 73 is mandatory. The PCS shall support the primitive AN_LINK.indication(link_status) (see 73.9). The parameter link_status shall take the value FAIL when PCS_status=false and the value OK when PCS_status=true. The primitive shall be generated when the value of link_status changes.

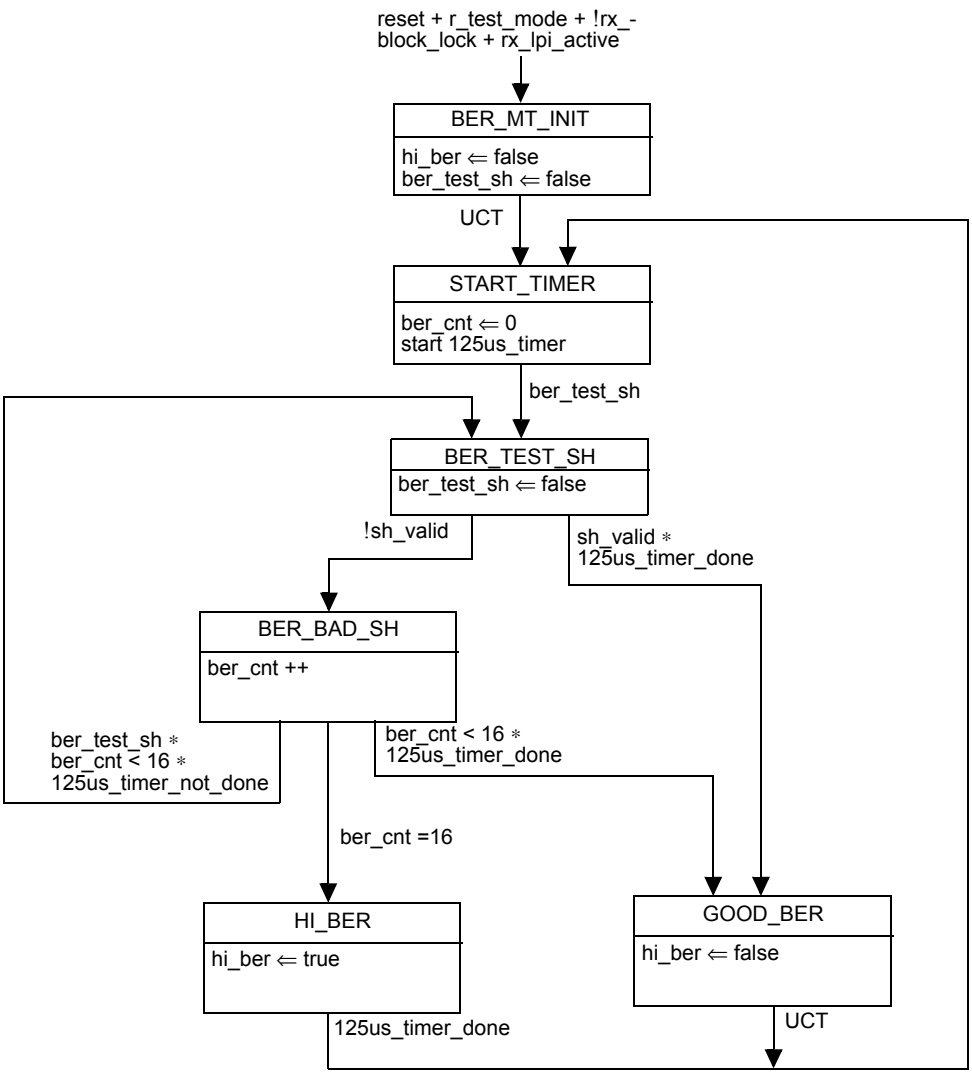

Figure 49–15—BER monitor state diagram

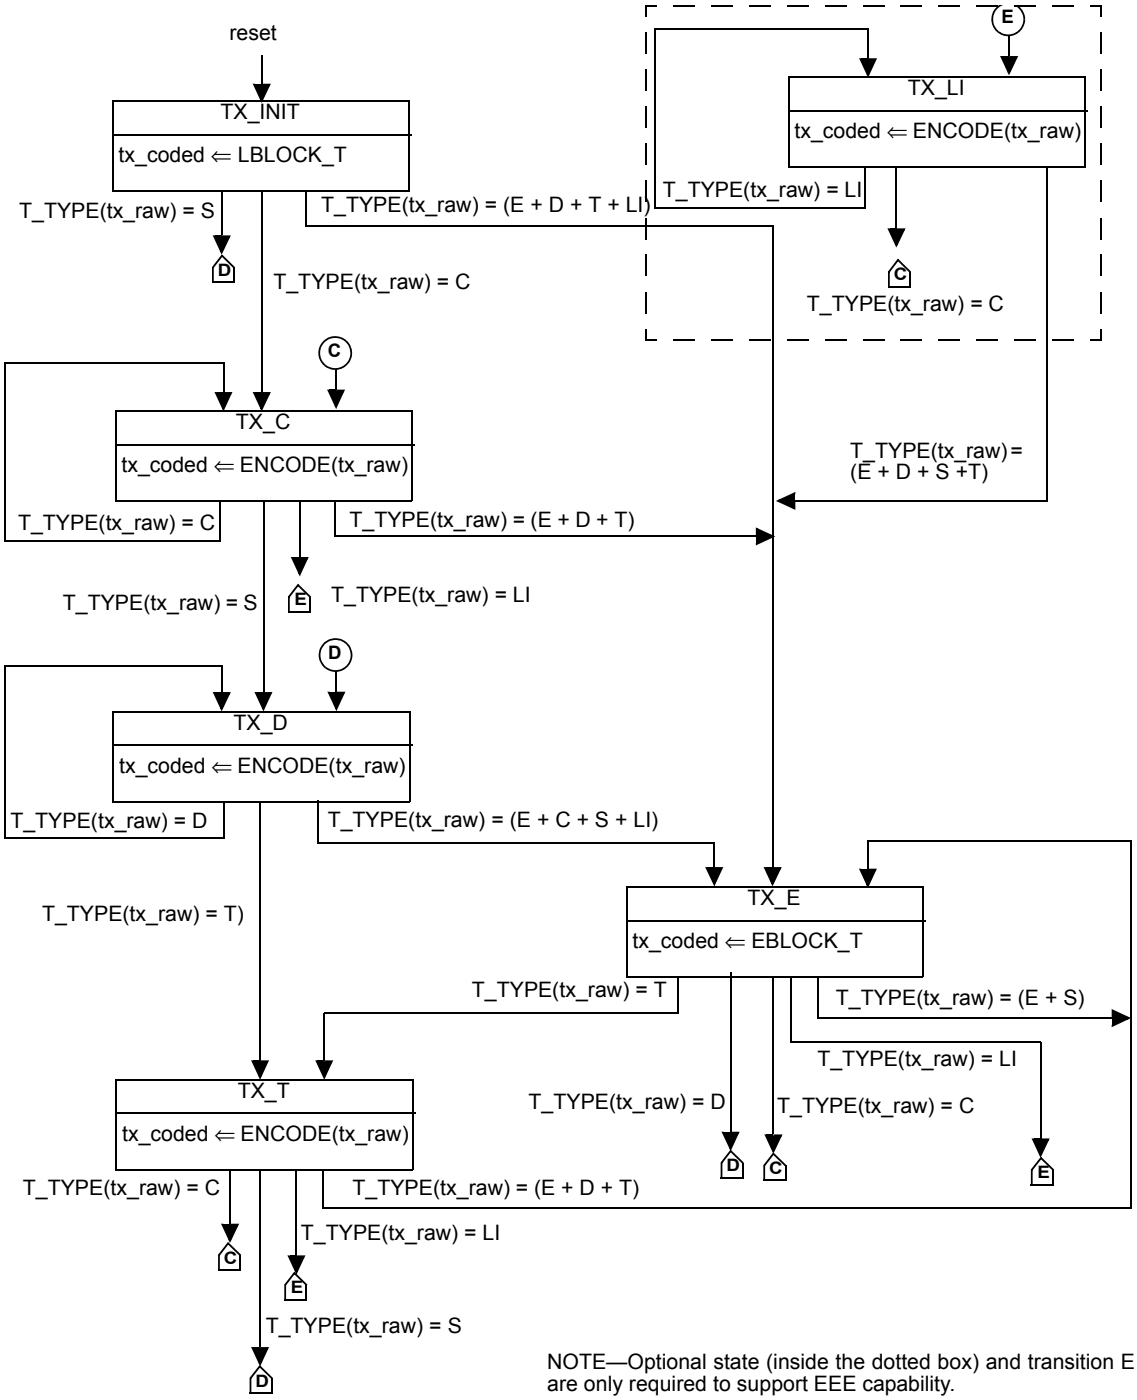

Figure 49–16—Transmit state diagram

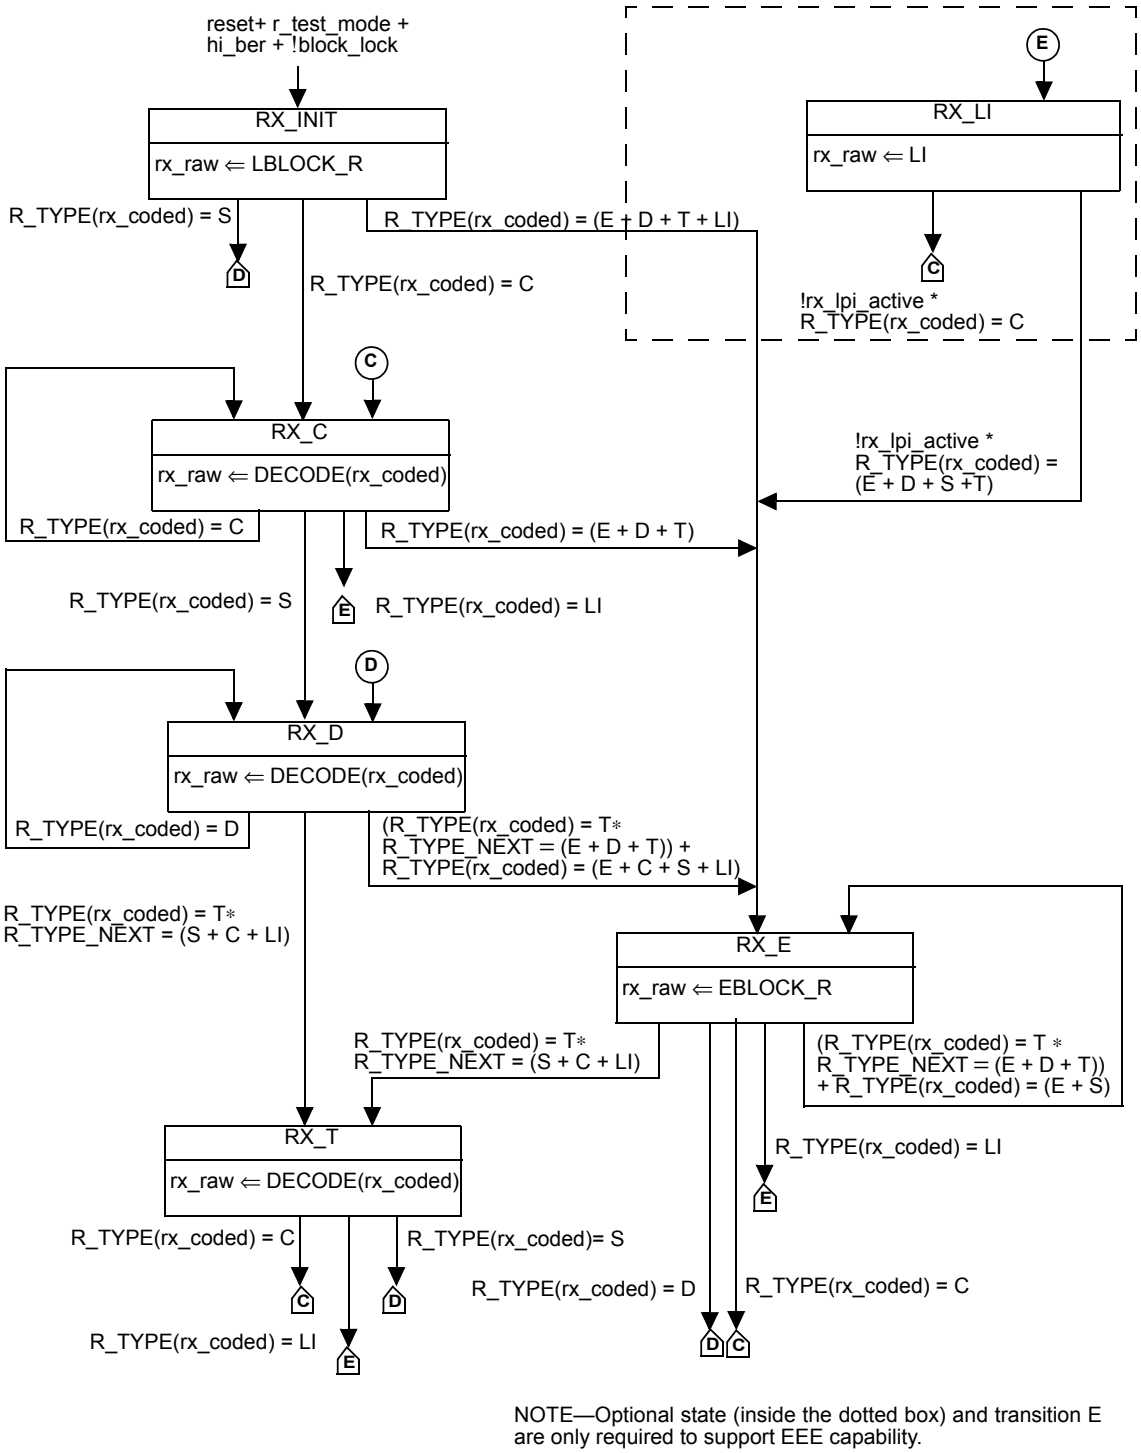

Figure 49–17—Receive state diagram

49.3 Protocol implementation conformance statement (PICS) proforma for Clause 49, Physical Coding Sublayer (PCS) type 10GBASE-R[10]

49.3.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 49, Physical Coding Sublayer (PCS), type 10GBASE-R, shall complete the following protocol implementation conformance statement (PICS) proforma. A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

49.3.2 Identification

49.3.2.1 Implementation identification

Supplier[1]	
Contact point for inquiries about the PICS[1]	
Implementation Name(s) and Version(s)[1,3]	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)[2]	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

49.3.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 49, Physical Coding Sublayer (PCS), type 10GBASE-R
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [ ] Yes [ ] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2015.)	
Date of Statement	

[10] *Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

49.3.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
XSBI	XBSI compatibility interface	51, 49.1.5	Compatibility interface is supported	O	Yes [ ] No [ ]
XGE	XGMII compatibility interface	46, 49.1.5	Compatibility interface is supported	O	Yes [ ] No [ ]
MD	MDIO	45, 49.2.14	Registers and interface supported	O	Yes [ ] No [ ]
*WIS	Supports operation with a WIS	49.1.5		O.1	Yes [ ] No [ ]
PMA	Supports operation directly connected to a PMA	49.1.5		O.1	Yes [ ] No [ ]
*JTM	Supports test-pattern mode	49.2.2		PMA:M	Yes [ ] No [ ] N/A[ ]
LPI	Implementation of LPI	49.2.4.4		O	Yes [ ] No [ ]

49.3.4 PICS Proforma Tables for PCS, type 10GBASE-R

49.3.4.1 Coding rules

Item	Feature	Subclause	Value/Comment	Status	Support
C1	Encoder implements the code as specified	49.2.13.2.3 and 49.2.4		M	Yes [ ] No [ ]
C2	Decoder implements the code as specified	49.2.13.2.3 49.2.4		M	Yes [ ] No [ ]
C3	IDLE control code insertion and deletion	49.2.4.7	Insertion or Deletion in groups of 4 /I/s	M	Yes [ ] No [ ]
C4	IDLE control code deletion	49.2.4.7	When deleting /I/s, the first four characters after a /T/ shall not be deleted.	M	Yes [ ] No [ ]
C5	Sequence ordered set deletion	49.2.4.10	Only one whole ordered set of two consecutive sequence ordered sets may be deleted	M	Yes [ ] No [ ]

49.3.4.2 Scrambler and Descrambler

Item	Feature	Subclause	Value/Comment	Status	Support
S1	Scrambler	49.2.6	Performs as shown in Figure 49–8	M	Yes [ ] No [ ]
S2	Descrambler	49.2.10	Performs as shown in Figure 49–10	M	Yes [ ] No [ ]

49.3.5 Test-pattern modes

Item	Feature	Subclause	Value/Comment	Status	Support
JT1	Square wave and pseudo-random transmit test-pattern generators are implemented	49.2.8	Performs as in 49.2.8	JTM:M	Yes [ ] No [ ] N/A[ ]
JT2	Pseudo-random receive test-pattern checker is implemented	49.2.12	Performs as in 49.2.12	JTM:M	Yes [ ] No [ ] N/A[ ]
JT3	Transmit and receive test-pattern modes can operate simultaneously	49.2.2		JTM:M	Yes [ ] No [ ] N/A[ ]
JT4	Reject transmit test-pattern mode when WIS is attached	49.2.8		JTM:O	Yes [ ] No [ ] N/A[ ]
JT5	Reject receive test-pattern mode when WIS is attached	49.2.12		JTM:O	Yes [ ] No [ ] N/A[ ]
*JT6	Support for PRBS31 test pattern	49.2.8		JTM:O	Yes [ ] No [ ] N/A[ ]
JT7	PRBS31 test-pattern generator is implemented	49.2.8	Performs as in 49.2.8	JT6:M	Yes [ ] No [ ] N/A[ ]
JT8	PRBS31 test-pattern checker is implemented	49.2.12	Performs as in 49.2.12	JT6:M	Yes [ ] No [ ] N/A[ ]
*JT9	Support for PRBS9 transmit test pattern	49.2.8		JTM:O	Yes [ ] No [ ] N/A[ ]
JT10	PRBS9 transmit test pattern is implemented	49.2.8	Performs as in 49.2.8	JT9:M	Yes [ ] No [ ] N/A[ ]

49.3.5.1 Bit order

Item	Feature	Subclause	Value/Comment	Status	Support
B1	Transmit bit order	49.2.7	Placement of bits into tx_data-group<15:0> as specified in 49.2.7	M	Yes [ ] No [ ]
B2	Receive bit order	49.2.9	Placement of bits from rx_data-group<15:0> into blocks as specified in 49.2.9	M	Yes [ ] No [ ]

49.3.6 Management

Item	Feature	Subclause	Value/Comment	Status	Support
M1	Alternate access to PCS Management objects is provided	49.2.14		O	Yes [ ] No [ ]

49.3.6.1 State diagrams

Item	Feature	Subclause	Value/Comment	Status	Support
SM1	Lock	49.2.13.3	Meets the requirements of Figure 49–14	M	Yes [ ] No [ ]
SM2	BER Monitor	49.2.13.3	Meets the requirements of Figure 49–15	M	Yes [ ] No [ ]
SM3	Transmit	49.2.13.3	Meets the requirements of Figure 49–16	M	Yes [ ] No [ ]
SM4	Receive	49.2.13.3	Meets the requirements of Figure 49–17	M	Yes [ ] No [ ]
SM5	PCS_R_Status	49.2.2	PCS_R_STATUS is FAIL in RX_INIT state; OK otherwise	M	Yes [ ] No [ ]

49.3.6.2 WIS

Item	Feature	Subclause	Value/Comment	Status	Support
W1	Supports WIS_SIGNAL.request	49.2.2	PCS_R_STATUS is FAIL when in RX_INIT state	WIS:M	Yes [ ] No [ ] N/A[ ]

49.3.6.3 Loopback

Item	Feature	Subclause	Value/Comment	Status	Support
L1	Supports Loopback	49.2.14.4	Performs as in 49.2.14.4	M	Yes [ ] No [ ] N/A[ ]

49.3.6.4 Delay Constraints

Item	Feature	Subclause	Value/Comment	Status	Support
TIM1	PCS Delay Constraint	49.2.15	No more than 3584 BT for sum of transmit and receive path delays	M	Yes [ ] No [ ]

49.3.6.5 Auto-Negotiation for Backplane Ethernet functions

Item	Feature	Subclause	Value/Comment	Status	Support
AN1*	Support for use with a 10GBASE-KR PMD	49.2.16	AN technology dependent interface described in Clause 73	O	Yes [ ]
AN2	AN_LINK.indication primitive	49.2.16	Support of the primitive AN_LINK.indication(link_status), when the PCS is used with 10GBASE-KR PMD	AN1:M	Yes [ ]
AN3	link_status parameter	49.2.16	Takes the value OK or FAIL, as described in 49.2.16	AN1:M	Yes [ ]
AN4	Generation of AN_LINK.indication primitive	49.2.16	Generated when the value of link_status changes	AN1:M	Yes [ ]

49.3.6.6 LPI functions

Item	Feature	Subclause	Value/Comment	Status	Support
LP-01	Insertion and deletion of LPIs in groups of four	49.2.4.7		LPI:M	Yes [ ] No [ ]
LP-02	Unscrambled data transmitted when scrambler_bypass = TRUE	49.2.6		LPI:M	Yes [ ] No [ ]
LP-03	Scrambler continues to operate as normal when scrambler_bypass = TRUE	49.2.6		LPI:M	Yes [ ] No [ ]
LP-04	scr_bypass_enable = TRUE when FEC is in use	49.2.13.2.2		LPI:M	Yes [ ] No [ ]
LP-05	Transmit state diagrams	49.2.13.3	Support additions to Figure 49–16 for LPI operation	LPI:M	Yes [ ] No [ ]
LP-06	Receive state diagrams	49.2.13.3	Support additions to for LPI operation	LPI:M	Yes [ ] No [ ]
LP-07	LPI transmit state diagrams	49.2.13.3.1	Meets the requirements of Figure 49–12	LPI:M	Yes [ ] No [ ]
LP-08	LPI receive state diagrams	49.2.13.3.1	Meets the requirements of Figure 49–13	LPI:M	Yes [ ] No [ ]
LP-09	LPI transmit timing	49.2.13.3.1	Meets the requirements of Table 49–2	LPI:M	Yes [ ] No [ ]
LP-10	LPI receive timing	49.2.13.3.1	Meets the requirements of Table 49–3	LPI:M	Yes [ ] No [ ]

50. WAN Interface Sublayer (WIS), type 10GBASE-W

50.1 Overview

The WAN Interface Sublayer (WIS) is an optional PHY sublayer that may be used to create a 10GBASE-W PHY that is data-rate and format compatible with the SONET STS-192c transmission format defined by ANSI, as well as the Synchronous Digital Hierarchy (SDH) VC-4-64c container specified by ITU. The purpose of the WIS is to allow 10GBASE-W equipment to generate Ethernet data streams that may be mapped directly to STS-192c or VC-4-64c streams at the PHY level, without requiring MAC or higher-layer processing. The WIS therefore specifies a subset of the logical frame formats in the SONET and SDH standards. In addition, the WIS constrains the effective data throughput at its service interface to the payload capacity of STS-192c / VC-4-64c, i.e., 9.58464 Gb/s. Multiplexed SONET/SDH formats are not supported.

The WIS does not render a 10GBASE-W PHY compliant with either SONET or SDH at any rate or format. A 10GBASE-W interface is not intended to interoperate directly with interfaces that comply with SONET or SDH standards, or other synchronous networks. Such interoperation would require full conformance to the optical, electrical, and logical requirements specified by SONET or SDH, and is outside the scope and intent of this standard. Operation over electrically multiplexed payloads of a transmission network is outside the scope of this standard.

From the perspective of the 10 Gb/s MAC layer, a 10GBASE-W PHY does not appear different (in either the functions or service interface) from a PHY without a WIS, with the exception of sustained data rate. However, a 10GBASE-W interface may interoperate only with another 10GBASE-W interface.

50.1.1 Scope

This clause specifies the functions, features, services, and protocol of the WIS. The WIS may be used with any of the PCS, PMA, and PMD sublayers that are defined for 10GBASE-W, as shown in Figure 50–1, it is placed between the PCS and PMA sublayers within the 10GBASE-W PHY. The WIS is common to all members of the family of 10GBASE-W WAN-compatible PHY implementations. There are currently three embodiments within this family: 10GBASE-SW, 10GBASE-LW, and 10GBASE-EW.

The definition of the WIS is based on the subset of signaling rates and data formats standardized by ATIS-0600416.1999(R2010), which in turn is based on ATIS-0900105.2008. The WIS maps the encoded Ethernet data received (transmitted) from (to) the PCS into a frame structure that has the same format as that defined by ATIS-0600416.1999(R2010), implementing a minimal number of the standard SONET overhead fields and functions. The WIS does not adhere to the electrical and optical aspects of SONET specified by ATIS-0600416.1999(R2010), as it is intended to be used with PHYs that conform to the corresponding parameters defined by the 10GBASE-W standard.

The WIS shall meet all requirements of ATIS-0600416.1999(R2010) except those that are specifically excluded by this clause. The following sections of ATIS-0600416.1999(R2010) are excluded in their entirety:

a) Section 5 (Jitter)
b) Section 6 (Synchronization)
c) Section 7.2.2 (VT1.5 rate—Electrical Interface)
d) Section 7.4.2 (VT1.5 rate)
e) Section 7.6 (Performance and failure alarm monitoring)
f) Section 7.7 (Performance monitoring functions)
g) Annex A (SONET VT1.5 Line Interface Common Criteria)
h) Annex B (SONET maintenance signals for the NI)

i) Annex C (Receiver jitter tolerance and transfer)

For convenience and clarity of explanation, this clause employs the notational conventions, nomenclature, and acronyms of SONET as defined by ATIS-0600416.1999(R2010). This should not, however, be taken as implying that the WIS excludes compatibility with ITU-defined SDH rates and formats.

50.1.2 Objectives

The following are the objectives for the WIS:

a) To support the full duplex mode of operation of the Ethernet MAC
b) To support the PCS, PMA and PMD sublayers defined for 10GBASE-W
c) To provide a 9.95328 Gb/s effective data rate at the service interface presented by the PMA sublayer, conforming with the requirements of SONET STS-192c and SDH VC-4-64c frame rates
d) To implement the framing, scrambling, and defect/anomaly detection to allow minimal compatibility with the requirements of both SONET and SDH transmission networks
e) To preserve the full duplex behavior and BER objectives of the PCS and PMD sublayers with which it may be used

NOTE—This clause will always use the term “WIS frames” when referring to the SONET-compatible frame format that is generated and terminated by the WIS, in order to distinguish such frames from those generated and terminated by the MAC. The WIS has no involvement with, or visibility into, the MAC framing processes. The “WIS frames” handled by the WIS are not propagated to higher layers via the WIS service interface; they are only exchanged between WIS peer entities on either side of a point-to-point link, and are unobservable and have no meaning outside WIS entities.

50.1.3 Relationship to other sublayers

Figure 50–1 depicts the relationships among the WIS (shown shaded), the 10 Gb/s MAC and Reconciliation Sublayers, the 10GBASE-R PCS and PMA, the ISO/IEC 8802-2 LLC, and the ISO/IEC Open System Interconnection (OSI) reference model.

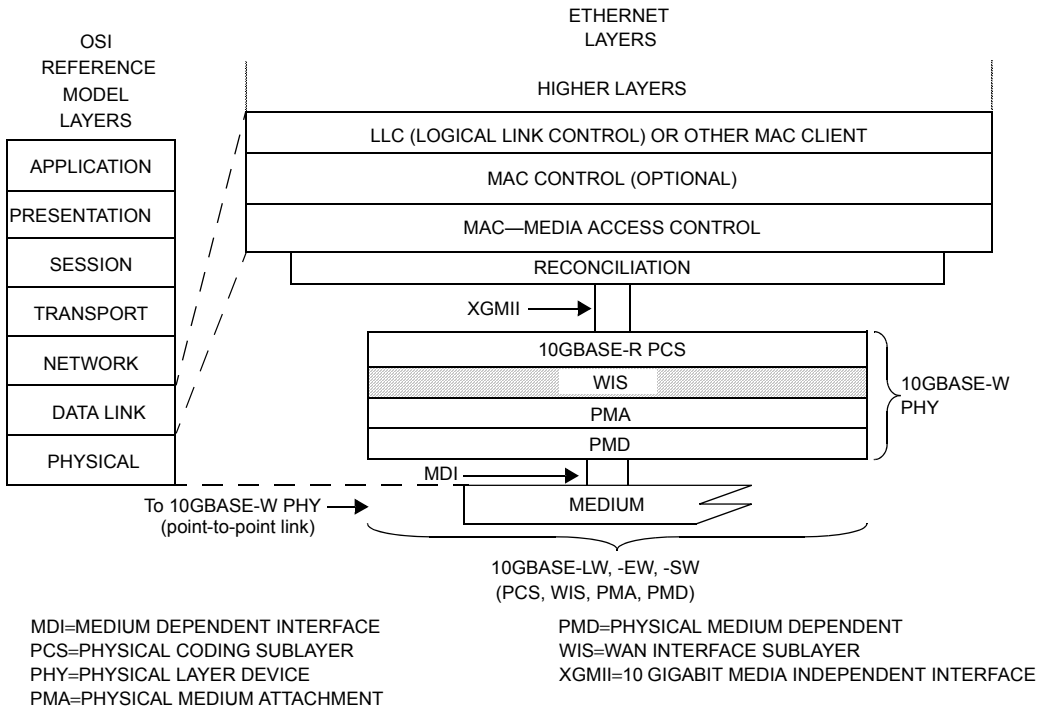

Figure 50–1—WIS relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

50.1.4 Summary of functions

The following provides a summary of the principal functions implemented by the WIS.

In the transmit direction (i.e., when transferring data from the PCS to the PMA), the WIS performs the following functions:

a) Mapping of data-units received from the PCS via the WIS Service Interface to the payload capacity of the Synchronous Payload Envelope (SPE) defined for STS-192c
b) Addition of Path Overhead and fixed stuff octets to generate the actual SPE
c) Creation of frames consisting of Line Overhead and Section Overhead octets plus the SPE, and the generation and insertion of Section, Line, and Path Bit Interleaved Parity (BIP)
d) Scrambling of the generated WIS frames
e) Transmission of these frames to the PMA sublayer via the PMA Service Interface

In the receive direction, the functions performed by the WIS include the following:

f) Reception of data from the PMA sublayer via the PMA Service Interface
g) Delineation of octet boundaries as well as STS-192c frame boundaries within the unaligned data stream from the PMA
h) Descrambling of the payload and overhead fields within the incoming frames

i) Processing of the pointer field within the Line Overhead, and delineation of the boundaries of the Synchronous Payload Envelopes (SPE) within the received WIS frames
j) Generation and checking of Bit Interleaved Parity (BIP) within the Section, Line, and Path Overheads
k) Removal of Line, Section, and Path Overhead columns, as well as fixed stuff columns, in order to extract the actual payload field
l) Handling of errors and exception conditions detected within the incoming WIS frame stream, and reporting these errors to Layer Management
m) Mapping of octets extracted from the payload capacity of the incoming SPE to data-units that are passed to the PCS via the WIS Service Interface

50.1.5 Sublayer interfaces

A WIS Service Interface is provided to allow the WIS to transfer information to and from the 10GBASE-R PCS, which is the sole WIS client. An abstract service model is used to define the operation of this interface. In addition, the WIS utilizes the service interface provided by the PMA sublayer to transfer information to and from the PMA. This standard defines these interfaces in terms of bits, octets, data-units and signals; however, implementers may choose other data-path widths and other control mechanisms for implementation convenience, provided that the logical models of the service interfaces are adhered to.

NOTE—The PMA service interface may be optionally instantiated as an actual physical interface, referred to as the XSBI; in this case, the WIS must also implement the client portion of the XSBI, as defined in 51.4, and conform to its data-path widths and control mechanisms.

50.1.6 Functional block diagram

Figure 50–2 provides a functional block diagram of the WIS.

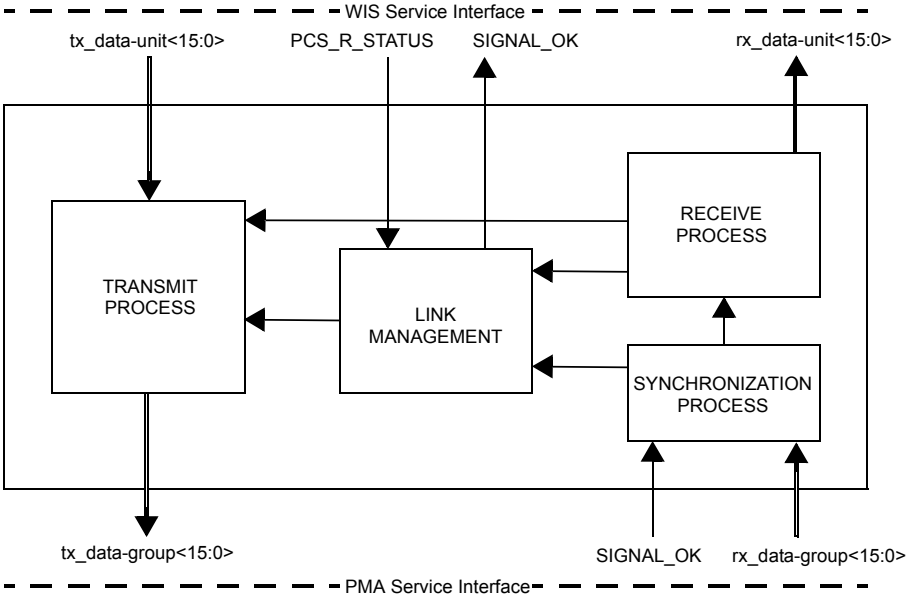

Figure 50–2—Functional block diagram

50.1.7 Notational conventions

The state diagrams within the body of this clause follow the notations and conventions of 21.5. State diagram timers follow the conventions of 14.2.3.2.

The Path, Line, and Section Overhead fields handled by the WIS are referenced using the standard terminology given in Section 4.2 of ATIS-0600416.1999(R2010). Nomenclature relating to SONET Defects and Anomalies that must be processed by the WIS, as well as other definitions and abbreviations required by this clause, are defined by Section 3 of ATIS-0600416.1999(R2010).

The labeling and transmission order of bits within an octet is different between the SONET/SDH and Ethernet protocol families. SONET numbers bits within an octet from 1 to 8 inclusive, with the bit labeled 1 being the most-significant bit (MSB) and the bit labeled 8 being the least-significant bit (LSB); when serialized, bits are transmitted in order from MSB to LSB (i.e., the bit numbered 1 is transmitted first and the bit numbered 8 is transmitted last). The Ethernet numbering scheme, on the other hand, normally numbers bits from 0 to 7 inclusive, with bit 0 being the LSB and bit 7 being the MSB; bits are transmitted from LSB to MSB. This clause employs the SONET numbering scheme throughout, with the exception of the WIS service interface, which utilizes the Ethernet numbering scheme. Special provisions are made when mapping payload data into the WIS frame (see 50.3.1.1 and 50.3.1.2) to ensure that the reversed transmission order does not impact the properties of the Ethernet FCS.

50.2 WIS Service Interface

The WIS Service Interface is provided to allow the 10GBASE-R PCS to transfer information to and from the WIS. These services are defined in an abstract manner and do not imply any particular implementation. The WIS Service Interface supports the exchange of data-units between PCS entities on either side of a 10GBASE-W link using request and indicate primitives. Data-units are mapped into WIS frames by the WIS and passed to the PMA, and vice versa.

The following primitives are defined within the WIS Service Interface:

WIS_UNITDATA.request(tx_data-unit<15:0>)
WIS_UNITDATA.indication(rx_data-unit<15:0>)
WIS_SIGNAL.request(PCS_R_STATUS)
WIS_SIGNAL.indication(SIGNAL_OK)

50.2.1 WIS_UNITDATA.request

This primitive defines the transfer of data, in the form of constant-width data-units, from the PCS to the WIS. The data supplied via WIS_UNITDATA.request is mapped by the WIS Transmit process into the payload capacity of the outgoing WIS frame stream.

50.2.1.1 Semantics of the service primitive

WIS_UNITDATA.request(tx_data-unit<15:0>)

The data conveyed by WIS_UNITDATA.request is a 16-bit vector representing a single data-unit which has been prepared for transmission by the 10GBASE-R PCS Transmit process. The element tx_data-unit<0> shall be interpreted as the least significant bit of the vector, and tx_data-unit<7:0> shall be interpreted as the least significant octet of the vector. The least significant octet of the vector is assumed to be generated first by the PCS; when the WIS transfers the provided data to the PMA, it is mapped such that the least significant octet is also transmitted first.

50.2.1.2 When generated

The 10GBASE-R PCS sends tx_data-unit<15:0> to the WIS at a nominal rate of 599.04 Mtransfers/s, corresponding to the STS-192c payload rate of 9.58464 Gb/s.

NOTE—The mapping of a constant-rate data stream into the payload capacity of a SONET-compatible SPE may require buffering in order to absorb the effects of periodically inserting overhead fields. In addition, the WIS and PCS may require some form of clock synchronization between the data-unit generation process and the normal SONET framing process. Both of these mechanisms are considered to be implementation-specific and outside the scope of this standard.

50.2.1.3 Effect of receipt

Upon receipt of this primitive, the WIS Transmit process maps the data conveyed by the tx_data-unit<15:0> parameter into the payload of the transmitted WIS frame stream, adds overhead octets as necessary, scrambles the data, and transfers the result to the PMA via one or more PMA_UNITDATA.request primitives.

50.2.2 WIS_UNITDATA.indication

This primitive defines the transfer of received data, in the form of constant-width data-units, from the WIS to the PCS. WIS_UNITDATA.indication is generated by the WIS Receive process in response to WIS frame data received from the PMA.

50.2.2.1 Semantics of the service primitive

WIS_UNITDATA.indication(rx_data-unit<15:0>)

The rx_data-unit<15:0> parameter is a 16-bit vector that represents the data-unit transferred by the WIS to the 10GBASE-R PCS. The element rx_data-unit<0> shall be interpreted as the least significant bit of the vector, and rx_data-unit<7:0> shall be interpreted as the least significant octet of the vector. The least significant octet of the vector is assumed to be processed first by the PCS; when the WIS obtains the data from the PMA, it is mapped such that the least significant octet was also received first.

50.2.2.2 When generated

The WIS sends one rx_data-unit<15:0> to the 10GBASE-R PCS whenever it has delineated exactly 16 bits of valid payload information from the incoming data stream received from the PMA sublayer. The nominal rate of generation of the WIS_UNITDATA.indication primitive is 599.04 Mtransfers/s, as governed by the recovered bit clock and the reduction in data rate imposed by the removal and discard of overhead fields from the incoming WIS frame stream. This corresponds to the STS-192c payload rate of 9.58464 Gb/s.

NOTE—The removal of overhead fields creates variations in data rate that may have to be removed by suitable buffering before data are transferred to the PCS. These buffering and rate adaptation mechanisms are implementation dependent and outside the scope of this standard.

50.2.2.3 Effect of receipt

The effect of receipt of this primitive by the WIS client is unspecified by the WIS.

50.2.3 WIS_SIGNAL.request

This primitive is sent by the 10GBASE-R PCS to the WIS to indicate when the PCS has acquired or lost frame delineation, i.e., synchronization with respect to the codewords present in the received data transferred from the WIS to the PCS.

50.2.3.1 Semantics of the service primitive

WIS_SIGNAL.request(PCS_R_STATUS)

The PCS_R_STATUS parameter can take one of two values: OK or FAIL. A value of OK denotes that the PCS is successfully delineating valid codewords from the incoming data stream transferred to it by the WIS via the WIS_UNITDATA.indication primitive. A value of FAIL denotes that errors have been detected by the PCS that prevent synchronization to these codewords.

50.2.3.2 When generated

The 10GBASE-R PCS generates the WIS_SIGNAL.request primitive to the WIS whenever there is a change in the value of the PCS_R_STATUS parameter.

50.2.3.3 Effect of receipt

The receipt of this primitive causes the WIS to identify the presence or absence of a Loss of Code-group Delineation condition, and to report this condition to the remote WIS entity via an LCD-P defect indication, as described in 50.3.5.3.

50.2.4 WIS_SIGNAL.indication

This primitive is sent by the WIS to the PCS to indicate the status of the Receive process. WIS_SIGNAL.indication is generated by the WIS Receive process in order to propagate the detection of severe error conditions (e.g., no valid signal being received from the PMA sublayer) to the PCS.

50.2.4.1 Semantics of the service primitive

WIS_SIGNAL.indication(SIGNAL_OK)

The SIGNAL_OK parameter can take one of two values: OK or FAIL. A value of OK denotes that the WIS Receive process is successfully delineating valid payload information from the incoming data stream received from the PMA sublayer, and this payload information is being presented to the PCS via the WIS_UNITDATA.indication primitive. A value of FAIL denotes that errors have been detected by the Receive process that prevent valid data from being presented to the PCS, as described in 50.3.5; in this case the WIS_UNITDATA.indication primitive and its associated rx_data-unit<15:0> parameter are meaningless.

50.2.4.2 When generated

The WIS generates the WIS_SIGNAL.indication primitive to the 10GBASE-R PCS whenever there is a change in the value of the SIGNAL_OK parameter.

50.2.4.3 Effect of receipt

The effect of receipt of this primitive by the WIS client is unspecified by the WIS.

50.3 Functions within the WIS

As shown in Figure 50–2, the WIS comprises the WIS Transmit and WIS Receive processes for 10GBASE-W, together with a Synchronization process and a Link Management function. Figure 50–3 depicts a conceptual view of the functions within each of the Transmit and Receive processes. Note that the figure does not show the interfaces to the Link Management function.

The WIS Transmit process accepts fixed-width tx_data-units from the PCS via the WIS Service Interface, and maps them into the payload capacity of the transmitted WIS frame stream. Fixed stuff octets are added, together with a set of Path Overhead octets, to create a Synchronous Payload Envelope (SPE). Line and Section Overhead octets are combined with the SPE and then scrambled using a frame-synchronous

scrambler to produce the final transmitted WIS frame. The WIS continuously generates one SONET-compatible WIS frame, comprising overhead fields, fixed stuff and payload, every 125 microseconds. No gaps are present between WIS frames. The data produced by the Transmit process, depicted as tx_data-group<15:0> in Figure 50–2, are passed to the PMA via the PMA Service Interface.

The WIS Synchronization process accepts data from the PMA (via the PMA Service Interface, depicted as rx_data-group<15:0> in Figure 50–2) and performs an alignment operation to delineate both octet and frame boundaries within the received data stream. Aligned and framed data are passed to the WIS Receive process (depicted as sync_bits<15:0> in Figure 50–3), where Section and Line Overhead octets are extracted from the WIS frames and processed after de-scrambling the frame data. The payload pointer within the Line Overhead is used to delineate the start and end of the received SPE, and the Path Overhead is extracted from the SPE and processed. Finally, the fixed stuff is removed from the SPE as well, and the resulting data stream is conveyed to the PCS via the WIS Service Interface. Severe errors detected by the Synchronization and Receive processes (e.g., loss of WIS frame synchronization) cause a WIS_SIGNAL.indication primitive to be sent to the PCS with a parameter of FAIL (see 50.2.4).

The WIS Link Management function implements the management registers defined in 45.2.2. These registers allow the local Station Management entity to control and monitor the operational status of the WIS Transmit and Receive processes. In addition, SONET Defects and Anomalies detected during the processing of the incoming data are sent to the management registers for reporting.

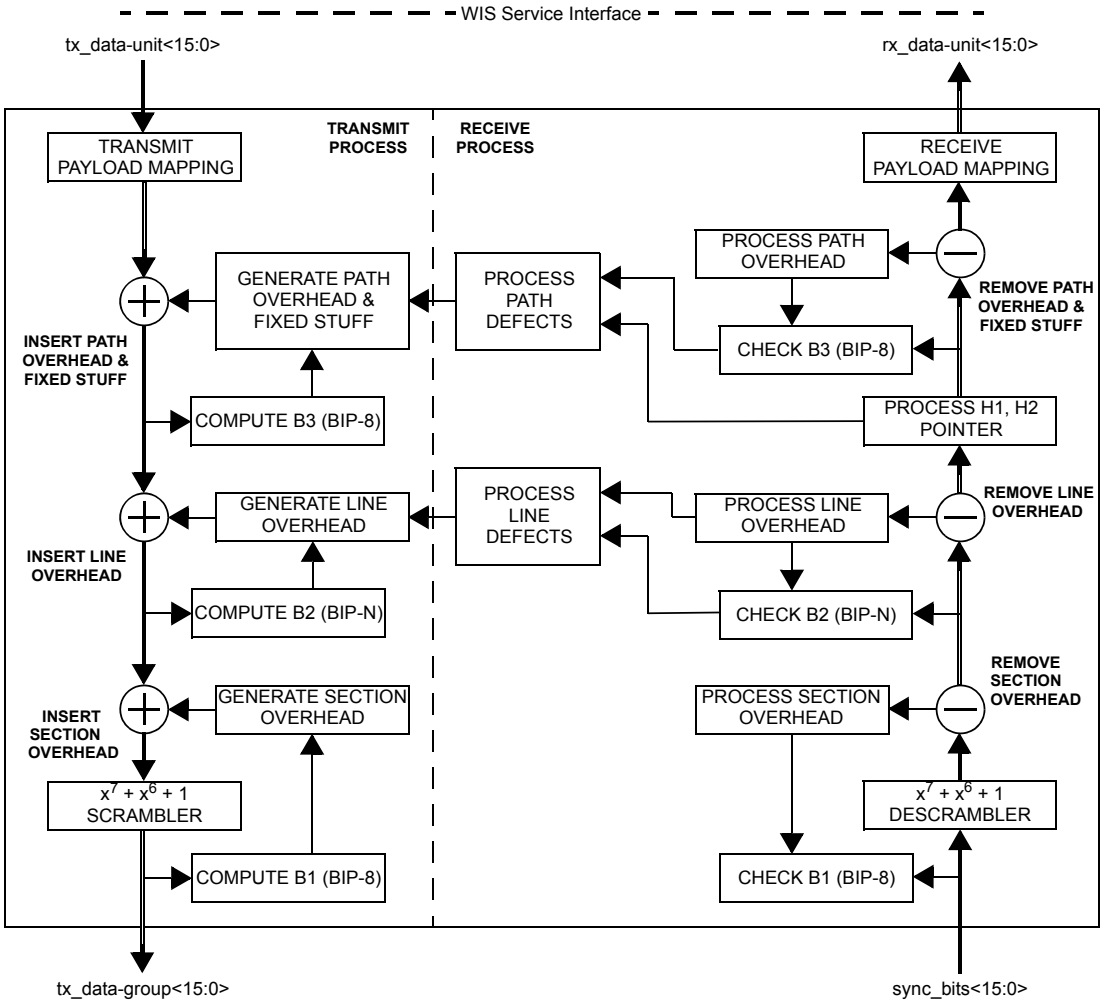

Figure 50–3—WIS Transmit and Receive processes

50.3.1 Payload mapping and data-unit delineation

The WIS maps the tx_data-unit<15:0> and rx_data-unit<15:0> parameters that are transferred via its service interface to/from the payload capacity of a standard STS-192c Synchronous Payload Envelope (SPE) structure. This structure is shown in Figure 50–4 for informative purposes. It is represented as a two-dimensional array with 9 rows and 16 704 columns, each row consisting of one octet of Path Overhead (added and removed as per 50.3.2.1 and 50.3.2.4), 63 octets of fixed stuff, and 16 640 octets of actual payload capacity.

The total payload capacity of the SPE comprises 149 760 octets (per WIS frame). As depicted in Figure 50–4, the octets of the payload capacity are numbered from left to right, starting at 0 (zero). The transmission order is also from left to right, i.e., lower-numbered octets are transmitted before higher-numbered octets.

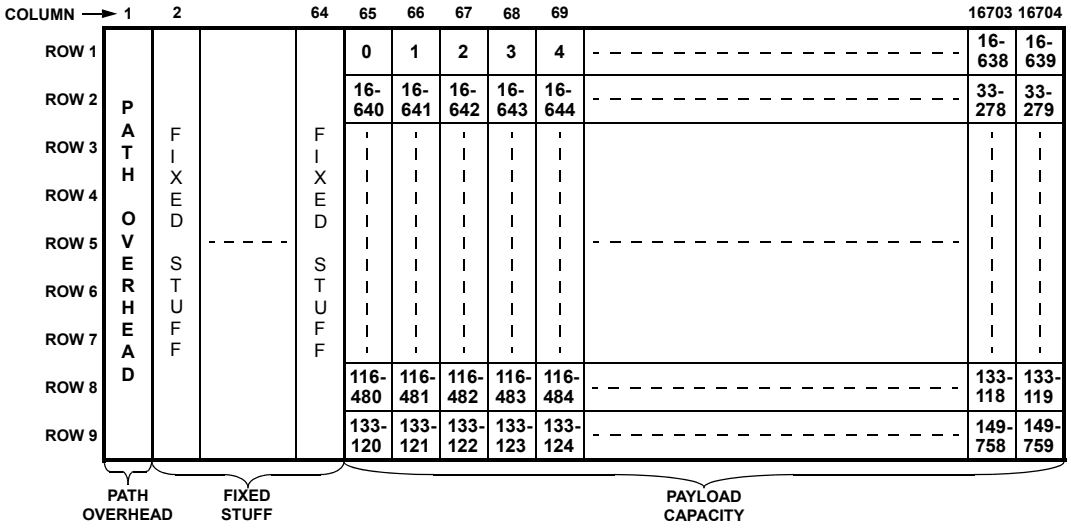

Figure 50–4—Structure of Synchronous Payload Envelope (informative)

50.3.1.1 Transmit payload mapping

The WIS Transmit function accepts a continuous stream of data-units from the PCS, conveyed by the tx_data-unit<15:0> parameter of the WIS_UNITDATA.request primitive of the WIS Service Interface. These data units shall be mapped to the SPE payload capacity and passed to the frame generation process described in 50.3.2 for insertion of the fixed stuff and Path Overhead columns. For each data-unit, tx_data-unit<7:0> shall be mapped to a lower-numbered (even) octet in the SPE payload capacity, and tx_data-unit<15:8> shall be mapped to the next higher-numbered (odd) octet.

A bit relabeling function is performed to map the bit numbering and ordering conventions followed by the WIS Service Interface to those implemented by the Transmit process (which follows the bit numbering and ordering conventions of SONET/SDH). The relabeling function shall be performed as shown in Figure 50–5. Bit 0 of tx_data-unit<15:0> is renumbered as bit 1 of the lower-numbered (even) octet within the SONET SPE, bit 15 of tx_data-unit<15:0> is renumbered as bit 8 of the higher-numbered octet within the SONET SPE, and the rest of the bits are renumbered in corresponding sequence.

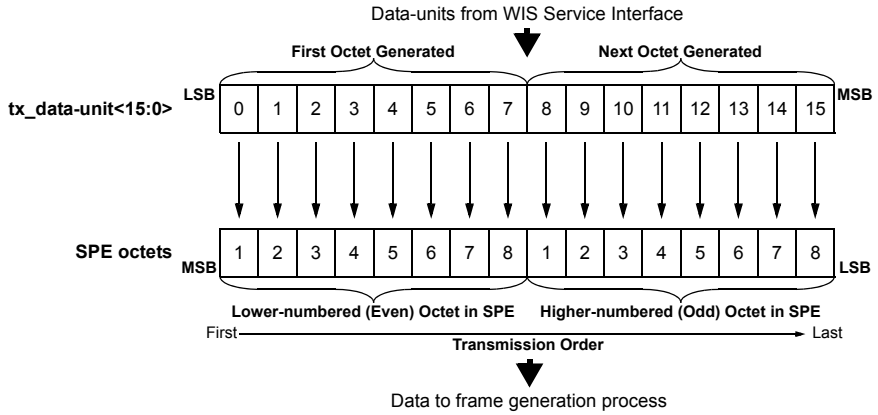

Figure 50–5—Transmit bit relabeling function

NOTE—The transmit bit relabeling function is required because SONET transmission order is from MSB to LSB (i.e., bit 1 to bit 8 of each octet, in the SONET numbering convention) while the Ethernet transmission order is from LSB to MSB (i.e., bit 0 to bit 7 of each octet, in the Ethernet numbering convention). It is necessary to maintain the expected transmission order of Ethernet in order to preserve the error detection properties of the Ethernet FCS. Therefore, the renumbering scheme causes the payload portion of the SPE to be transmitted from LSB to MSB with reference to incoming data accepted from the WIS Service Interface, but the Path, Line, and Section Overhead octets are transmitted from MSB to LSB as required by SONET.

50.3.1.2 Receive payload mapping

The WIS Receive function accepts a continuous stream of octets from the frame reception process described in 50.3.2, corresponding to the payload capacity of the SPE as illustrated in Figure 50–4. It shall map the octets in the SPE payload capacity to a continuous stream of data-units that are supplied to the PCS via the rx_data-unit<15:0> parameter of the WIS_UNITDATA.indication primitive of the WIS Service Interface. Within each data-unit, rx_data-unit<7:0> shall be mapped to a lower-numbered (even) octet in the SPE payload capacity, and rx_data-unit<15:8> shall be mapped to the next higher-numbered (odd) octet.

A bit relabeling function is performed to map the bit numbering and ordering conventions followed by the Receive process (which follows the bit numbering and ordering conventions of SONET/SDH) to those required by the WIS Service Interface. The relabeling function shall be performed as shown in Figure 50–6. Bit 1 of the lower-numbered (even) octet within the received SONET SPE is renumbered as bit 0 of rx_data-unit<15:0> at the WIS Service Interface, bit 8 of the higher-numbered octet within the SONET SPE is renumbered as bit 15 of rx_data-unit<15:0>, and the rest of the bits are renumbered in corresponding sequence.

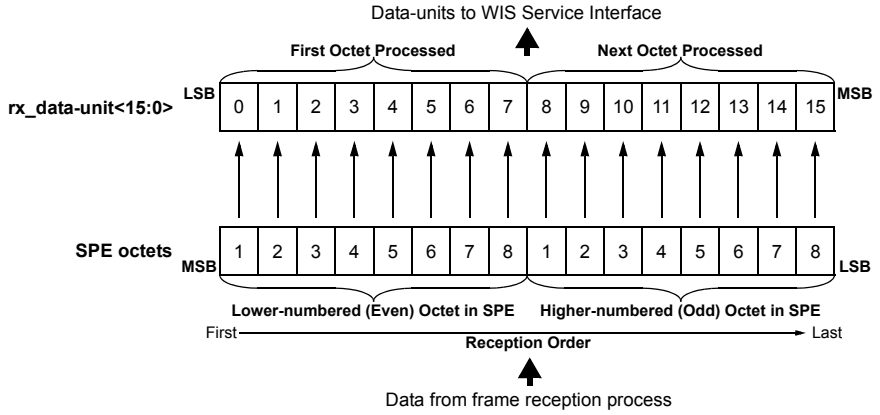

Figure 50–6—Receive bit relabeling function

NOTE—The receive bit relabeling function is required for the same reasons as the transmit bit relabeling function described in 50.3.1.1, i.e., to preserve the error detection properties of the Ethernet FCS in the face of the different transmission order required by SONET framing.

50.3.2 WIS frame generation

As part of the Transmit process, the WIS encapsulates the payload generated by the payload mapping function within a series of WIS frames. The Receive process performs the reverse operation, extracting payload from the incoming WIS frame stream and submitting it to the payload mapping function. SONET compatibility in the WIS follows ATIS-0600416.1999(R2010), with the exception of the superseding specifications in this document.

Frame generation in both the transmit and receive directions is performed in three stages. Figure 50–3 illustrates the sequence of stages. The relevant octets in the Path Overhead portion of the frame are shown in Figure 50–7, and Figure 50–8 illustrates the Section and Line Overhead generated by the WIS.

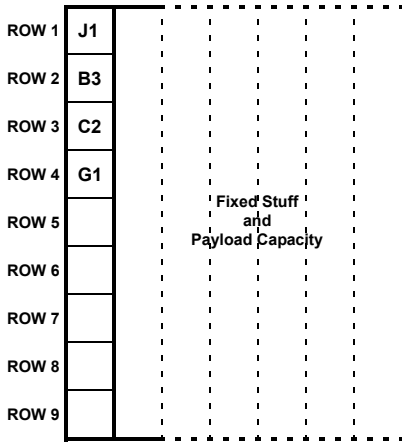

NOTE—The Path Overhead comprises 9 octets in total. Only 4 octets are defined for the WIS. Octets that are undefined and unused by the WIS are indicated as blank boxes in this figure.

Figure 50–7—Structure of Path Overhead Generated by WIS (informative)

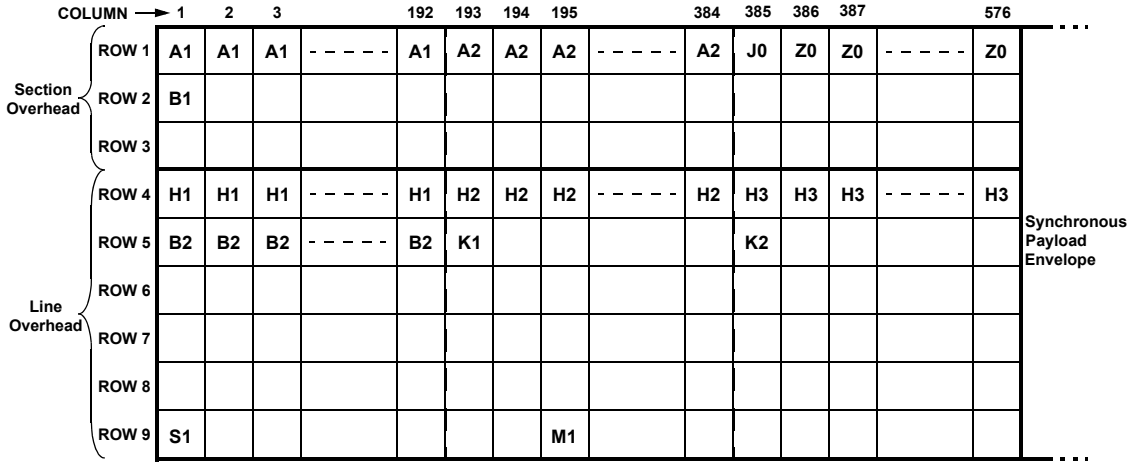

NOTE—The Section and Line Overhead comprise 5184 octets in total (576 x 9). Only 1349 octets are defined for the WIS. Octets that are undefined and unused by the WIS are indicated as blank boxes in this figure.

Figure 50–8—Structure of Section and Line Overhead Generated by WIS (informative)

NOTE—Figures 50–8 and 50–7 are provided for informative purposes only. In the event of any discrepancy observed between the figures and the normative portions of ATIS-0600416.1999(R2010), the latter shall take precedence.

50.3.2.1 Transmit Path Overhead insertion

The WIS Transmit process shall insert Path Overhead fields as defined in Section 4.2 of ATIS-0600416.1999(R2010), and specified in Table 50–1 of this standard. For the fields where the "Coding" column of Table 50–1 contains "per 416," the field is inserted according to the specifications of ATIS-0600416.1999(R2010). For the fields where the "Coding" column of Table 50–1 contains a specific value or "see text," this document supersedes the corresponding values in Table 1, "SONET overhead at NIs" in the ATIS document.

Table 50–1—STS Path Overhead

Overhead octet	Function	Usage	Coding (bits 1..8)
B3	STS Path error monitoring (Path BIP-8)	supported	per 416
C2	STS Path signal label	specified value	00011010
F2	Path user channel	unsupported	00000000
G1	Path status	supported	per 416
H4	Multiframe indicator	unsupported	00000000
J1	STS Path trace	specified value	see text
N1	Tandem connection maintenance/Path data channel	unsupported	00000000
Z3-Z4	Reserved for Path growth	unsupported	00000000
NOTE—SONET/SDH and IEEE 802.3 differ in bit ordering conventions. The values in this table follow SONET bit ordering, in which bit index values range from 1 to 8, from left to right, and bit 8 is the least significant bit.			

NOTE—The value assigned to the C2 octet corresponds to the provisional Path signal label value assigned for 10 Gigabit Ethernet. The fixed stuff octets that follow the Path Overhead column in the SPE are transmitted as 00000000 and are not checked on receive.

The J1 octet shall transport a 16-octet continuously repeating Path Trace Message that is formatted as defined by Section 5 and Annex A of ANSI T1.269-2000. Each successive octet of the Path Trace Message, starting from the first, is placed in the J1 octet of a successive WIS frame; after all 16 octets have been transmitted in this way, the process repeats. The Station Management entity may modify the Path Trace Message being transmitted via the Link Management register facilities provided in 50.3.11.1. If the Station Management entity is not currently transmitting a Path Trace Message, a default Path Trace Message consisting of 15 octets of zeros and a header octet formatted according to Section 5 of ANSI T1.269-2000 shall be transmitted.

NOTE—The bit representation of the header octet of the default Path Trace Message is 10001001 (89 hexadecimal). This value is provided here for informational purposes only; ANSI T1.269-2000 shall take precedence in case of any discrepancy.

50.3.2.2 Transmit Line Overhead insertion

The WIS Transmit process shall insert Line Overhead fields as defined in Section 4.2 of ATIS-0600416.1999(R2010), and specified in Table 50–2 of this document. For the fields where the "Coding" column of Table 50–2 contains "per 416," the field is inserted according to the specifications of ATIS-0600416.1999(R2010). For the fields where the "Coding" column of Table 50–2 contains a specific value or "see text," this document supersedes the corresponding values in Table 1, "SONET overhead at NIs" in the ATIS document. In addition, Line Overhead octets not listed in the latter table shall be set to 00000000 by the Transmit process.

The H1 and H2 pointer octets shall be set by the Transmit process, in accordance with the pointer mechanism defined by ATIS-0600416.1999(R2010), to indicate a constant pointer value of 522 decimal, and shall also indicate a concatenated payload.

The K1 octet shall be set to 00000000 by the Transmit process. In addition, bits 1 to 5 inclusive of the K2 octet shall be set to 00000 binary. Bits 6 through 8 inclusive shall either transport an RDI-L encoding, as defined by Section 7.4.1 of ATIS-0600416.1999(R2010), or be set to 000 binary. No other encodings of the K2 octet are allowed to be transmitted by the WIS Transmit process.

Table 50–2—Line Overhead

Overhead octet	Function	Usage	Coding (bits 1..8)
B2	Line error monitoring (Line BIP-1536)	supported	per 416
D4-D12	Line Data Communications Channel (DCC)	unsupported	00000000
E2	Orderwire	unsupported	00000000
H1-H2	Pointer	specified value	See text.
H3	Pointer action	specified value	00000000
K1, K2	Automatic protection switch (APS) channel and Line Remote Defect Identifier (RDI-L)	specified value	See text.
M0	STS-1 Line Remote Error Indication (REI)	unsupported	00000000
M1	STS-*N* Line Remote Error Indication (REI)	supported	per 416
S1	Synchronization messaging	unsupported	00001111
Z1	Reserved for Line growth	unsupported	00000000
Z2	Reserved for Line growth	unsupported	00000000
NOTE—SONET/SDH and IEEE 802.3 differ in bit ordering conventions. The values in this table follow SONET bit ordering, in which bit index values range from 1 to 8, from left to right, and bit 8 is the least significant bit.			

NOTE—The setting of the constant portions of the K1 and K2 octets denotes that the interface containing the WIS is acting as a working channel. The RDI-L encoding within the K2 octet serves to support the fault processing specified in 50.3.2.5. The setting of the S1 octet signifies that the transmit clock should not be used for synchronization. The bit numbering for these octets follows the SONET convention, in which bit index values range from 1 to 8, from left to right, and bit 8 is the least significant bit.

50.3.2.3 Transmit Section Overhead insertion

The WIS Transmit process shall insert Section Overhead fields as defined in Section 4.2 of ATIS-0600416.1999(R2010), and specified in Table 50–3 of this document. For the fields where the "Coding" column of Table 50–3 contains "per 416," the field is inserted according to the specifications of ATIS-0600416.1999(R2010). For the fields where the "Coding" column of Table 50–3 contains a specific value or "see text," this document supersedes the corresponding values in Table 1, "SONET overhead at NIs" in the ATIS document. In addition, Section Overhead octets not listed in the latter table shall be set to 00000000 by the Transmit process.

The J0 octet shall transport a 16-octet continuously repeating Section Trace Message that is formatted as defined by Section 5 and Annex A of ANSI T1.269-2000. Each successive octet of the Section Trace Message, starting from the first, is placed in the J0 octet of a successive WIS frame; after all 16 octets have

been transmitted in this way, the process repeats. The Station Management entity may modify the Section Trace Message being transmitted via the Link Management register facilities provided in 50.3.11.1. If the Station Management entity is not currently transmitting a Section Trace Message, a default Section Trace Message consisting of 15 octets of zeros and a header octet formatted according to Section 5 of ANSI T1.269-2000 shall be transmitted.

NOTE—The bit representation of the header octet of the default Section Trace Message is 10001001 (89 hexadecimal). This value is provided here for informational purposes only; ANSI T1.269-2000 shall take precedence in case of any discrepancy.

Table 50–3—Section Overhead

Overhead octet	Function	Usage	Coding (bits 1..8)
A1	Frame alignment	supported	per 416
A2	Frame alignment	supported	per 416
B1	Section error monitoring (Section BIP-8)	supported	per 416
D1-D3	Section Data Communications Channel (DCC)	unsupported	00000000
E1	Orderwire	unsupported	00000000
F1	Section User Channel	unsupported	00000000
J0	Section trace	specified value	See text
Z0	Reserved for Section growth	unsupported	11001100
NOTE 1—The bit representations of the octet values assigned to A1 and A2 are 11110110 (F6 hexadecimal) and 00101000 (28 hexadecimal), respectively. These values are provided here for informational purposes only. ATIS-0600416.1999(R2010) shall take precedence in case of any discrepancy. NOTE 2—SONET/SDH and IEEE 802.3 differ in bit ordering conventions. The values in this table follow SONET bit ordering, in which bit index values range from 1 to 8, from left to right, and bit 8 is the least significant bit.			

50.3.2.4 Receive Path, Line, and Section Overhead extraction

The WIS Receive process shall extract Path, Line and Section Overhead fields as defined in Section 4.2 of ATIS-0600416.1999(R2010), and specified in Table 50–1, Table 50–2, and Table 50–3 of this document. For the fields where the "Coding" columns of Table 50–1, Table 50–2, and Table 50–3 contain "per 416," the field is extracted according to the specifications of ATIS-0600416.1999(R2010). For the fields where the "Coding" columns contain a specific value or "see text," this document supersedes the corresponding values in Table 1, "SONET overhead at NIs" in the ATIS document. Overhead octets marked as unsupported in Table 50–1, Table 50–2, and Table 50–3 shall be ignored by the Receive process. In addition, overhead octets not listed in Table 1, "SONET overhead at NIs" in ATIS-0600416.1999(R2010) shall also be ignored by the Receive process.

The J0 octet in the received WIS frames shall be interpreted as transporting a 16-octet continuously repeating Section Trace Message. This Section Trace Message is extracted from the incoming WIS frame stream and passed to the Station Management entity via dedicated registers within the WIS MDIO register space (50.3.11.1). The WIS Receive process is not required to delineate Section Trace Message boundaries or process them in any way.

Extraction of valid Section Trace Message data shall begin after the WIS Receive process has successfully synchronized to the incoming WIS frame stream. Each successive J0 octet received thereafter is placed in a successive octet of the WIS J0 receive register set, until 16 J0 octets have been received, after which the process repeats from the first octet of the register set. As the incoming Section Trace Message is 16 octets in size, the contents of the WIS J0 receive register set will remain static as long as the same message is being received. Extraction of Section Trace Messages in this fashion is performed continuously as long as valid WIS frames are being received. The WIS Receive process is not required to interpret or process the extracted message in any way.

The J1 octet in the received WIS frames shall be interpreted as transporting a 16-octet continuously repeating Path Trace Message. This Path Trace Message is extracted from the incoming WIS frame stream and passed to the Station Management entity via dedicated registers within the WIS MDIO register space (50.3.11.1). The WIS Receive process is not required to delineate Path Trace Message boundaries or process them in any way.

Extraction of valid Path Trace Message data shall begin after the WIS Receive process has successfully synchronized to the incoming WIS frame stream. Each successive J1 octet received thereafter is placed in a successive octet of the WIS J1 receive register set, until 16 J1 octets have been received, after which the process repeats from the first octet of the register set. As the incoming Path Trace Message is 16 octets in size, the contents of the WIS J1 receive register set will remain static as long as the same message is being received. Extraction of Path Trace Messages in this fashion is performed continuously as long as valid WIS frames are being received. The WIS Receive process is not required to interpret or process the extracted message in any way.

Bits 6 through 8, inclusive, of the K2 octet in the received Line Overhead shall be interpreted by the WIS Receive process as carrying the Line Alarm Indication Signal (AIS-L) or Line Remote Defect Indication (RDI-L) signals (see 50.3.2.5). All other encodings of these bits shall be ignored by the Receive process.

NOTE—Values of overhead fields encoded by the Transmit process within a WIS entity at one end of a link may be changed by intervening Line and Section equipment before it reaches the peer WIS entity at the other end of the link. In particular, the H1, H2, and H3 fields may take on values different from those originally encoded. The receive process must be able to handle these situations in accordance with ATIS-0600416.1999(R2010), and therefore must contain a pointer interpreter function compliant with the latter. Bits 5 and 6 of the H1 octet are ignored by the receiver.

50.3.2.5 Fault processing

Defects and anomalies detected by the Receive process are classified as defined in Section 7.1 of ATIS-0600416.1999(R2010). Section, Line, and Path defects and anomalies listed in Table 50–4 of this document shall be detected and processed as defined by Sections 7.3, 7.4.1, and 7.5 of ATIS-0600416.1999(R2010). Defects and anomalies not listed in Table 50–4 are ignored.

NOTE—The description of the LOP-P defect in Section 7.5 of ATIS-0600416.1999(R2010) is in error, and deviates from the underlying definitive reference, which is ANSI T1.231-1997. The reader should therefore consult section 8.1.2.4.1 of ANSI T1.231-1997 for the specification of the LOP-P defect until ATIS-0600416.1999(R2010) has been corrected.

The WIS shall additionally support the PLM-P (Path Label Mismatch) and LCD-P (Loss of Code-group Delineation) defects. The detection and reporting of the PLM-P defect follows Section 7.5 of ATIS-0600416.1999(R2010). The LCD-P defect shares the same coding value and reporting method as the PLM-P defect, but is detected according to 50.3.5.3.

The Receive process shall detect a Loss Of Signal (LOS) defect as defined by Section 7.2.1 of ATIS-0600416.1999(R2010) without using any services provided by the PMA or PMD sublayers for this purpose. The parameters *T* and *T'* in Section 7.2.1 of ATIS-0600416.1999(R2010) shall both be set to a value ranging between 2.3 and 100 microseconds.

Table 50–4—WIS supported Near end events and Far end reports

	Physical media	Section		Line		Path	
	Defect	Anomaly	Defect	Anomaly	Defect	Anomaly	Defect
Near end	LOS	BIP-N(S)	SEF/LOF	BIP-N(L)	AIS-L	BIP-N(P)	LOP-P AIS-P
Far end	N/A	N/A	N/A	REI-L	RDI-L	REI-P	ERDI-P

NOTE—Section 7.2.2, "VT1.5 rate–Electrical Interface," Section 7.4.2 "VT1.5 Rate," Section 7.6, "Performance and Failure Monitoring," and Section 7.7 "Performance Monitoring Functions" of ATIS-0600416.1999(R2010) are not applicable and are not supported. Also, the AIS-L defect is only processed and reported by the WIS Receive process; it is never transmitted by the WIS Transmit process. In addition, a PLM-P defect is indistinguishable from an LCD-P defect at the far end, as the two defects are coded identically in the outgoing WIS frame stream; however, as both defects indicate a payload defect, this is of no consequence. They are detected separately for local reporting purposes.

50.3.3 Scrambling

The WIS shall implement a frame-synchronous scrambler within the Transmit process, and shall also implement a frame-synchronous descrambler within the Receive processes, both of sequence length 127 and as specified by Section 10.3 of ATIS-0900105.2008. The use of the scrambler is intended to provide for DC balance (i.e., an equal number of 1's and 0's on average) and for sufficient transition density within the bit stream transmitted on the medium. The scrambler is run over each bit of the entire WIS frame prior to transmission to the PMA, with the exception of all the A1, A2, J0 and Z0 octets (576 octets in total, per WIS frame), which are transmitted and received unscrambled. The scrambling function implemented by the Transmit process is identical in logical construction to the corresponding descrambling function in the Receive process.

Figure 50–9 illustrates the functional diagram of the frame-synchronous scrambler, in bit-serial form.

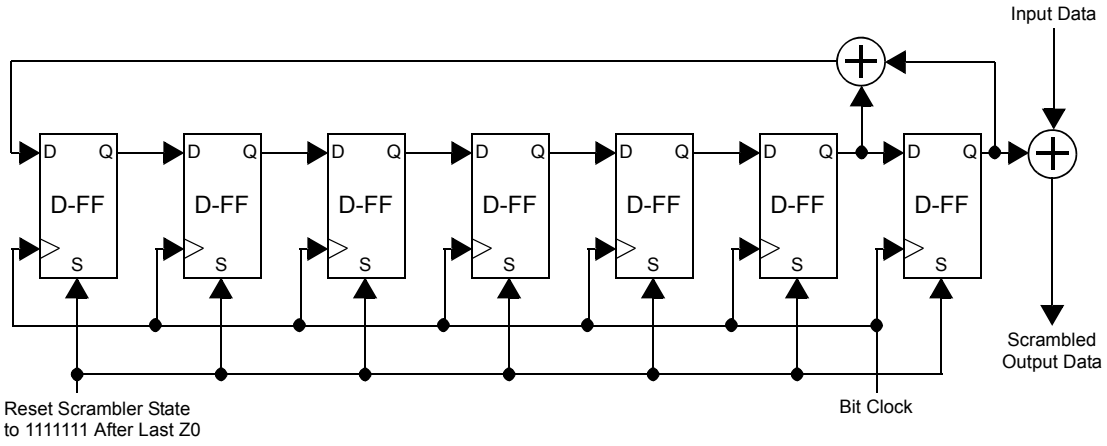

Figure 50–9—Scrambling function (informative)

NOTE—The diagram in Figure 50–9 is intended for reference and informative purposes only. Implementations may select any conforming means of realizing the scrambler function described herein. If a discrepancy is observed between the descriptions and diagram in this section and ATIS-0900105.2008, the latter takes precedence.

50.3.3.1 Scrambler polynomial

The scrambler polynomial used follows that specified in Section 10.3 of ATIS-0900105.2008. The scrambler state is reset, as specified in the latter, to binary 1111111 prior to processing the octet immediately following the last Z0 octet in every WIS frame.

50.3.3.2 Scrambler bit ordering

The bit ordering considerations for the frame-synchronous scrambler in the WIS are illustrated in Figure 50–10. As shown, the scrambler in effect processes the transmitted (received) WIS frame octet-by-octet, starting with the most-significant (leftmost) bit of the first octet following the last Z0 octet in the WIS frame. The scrambler then proceeds bit-by-bit until the least-significant bit has been reached, after which it continues with the most-significant bit of the second octet.

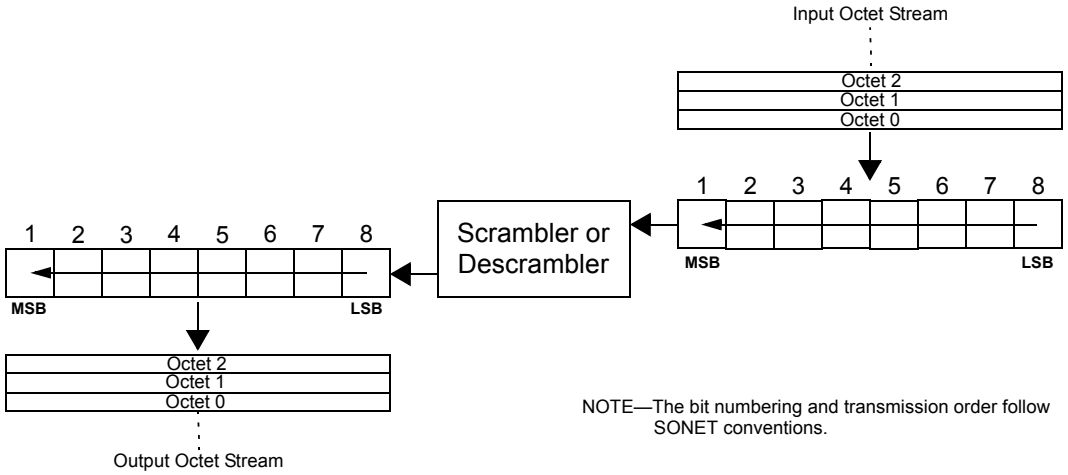

Figure 50–10—Scrambler bit ordering (informative)

NOTE—The above description and diagram are informative.

50.3.4 Octet and frame delineation

The WIS is required to delineate both octet and WIS frame boundaries within the stream of received data words presented by the PMA via its service interface to the Receive process. Delineation of these boundaries and alignment of the received data are done prior to performing the descrambling function (as opposed to the SPE delineation process, which is done after descrambling the received data). The delineation process employs the A1 and A2 octets in the Section Overhead portion of the received WIS frame stream, and searches for valid transitions between A1 and A2 patterns.

The WIS shall implement the synchronization process described in 50.4.2 to delineate octet and WIS frame boundaries. Failure to achieve or maintain synchronization shall result in Severely Errored Frame (SEF) or Loss Of Frame (LOF) Defects being indicated as described in 50.3.2.5.

50.3.5 Error propagation

The WIS Receive and Synchronization processes detect specific errors during reception that prevent delineation of valid data from the incoming WIS frame stream, and cause these errors to be propagated to the 10GBASE-R PCS. Error propagation shall be performed by generating a WIS_SIGNAL.indication primitive with the SIGNAL_OK parameter set to FAIL, regardless of the actual contents of any payload that may be extracted from the incoming WIS frame stream. A subsequent WIS_SIGNAL.indication primitive

with SIGNAL_OK set to OK shall be issued only after the WIS determines that none of these particular errors exist, subject to the error propagation timing below.

In addition, the WIS Receive process detects a Loss of Code-group Delineation condition, and report it to the far end WIS as an LCD-P defect by the method described in 50.3.2.5.

50.3.5.1 Propagated errors

The following conditions shall be detected as errors and reported to the 10GBASE-R PCS via the mechanism of 50.3.5:

a) The Synchronization process is not in the SYNC state, as defined by the state diagram of 50.4.2.
b) A PLM-P defect is detected (50.3.2.5).
c) An AIS-P defect is detected (50.3.2.5).
d) An LOP-P defect is detected (50.3.2.5).

NOTE—The error propagation mechanism may also be used to indicate when the WIS is unable to supply valid data to the PCS for any other reason, such as an internal error, a reset condition, or when disabled via layer management. This is considered to be implementation specific and is outside the scope of this clause.

50.3.5.2 Error propagation timing

Propagation of errors to the PCS according to the mechanism of 50.3.5 begins as soon as possible after the detection of one or more of the error conditions specified in 50.3.5.1. Error propagation shall terminate, and valid data shall be transferred to the PCS, within 125 microseconds of the removal of all of the error conditions in 50.3.5.1.

50.3.5.3 Loss of Code-group Delineation

The WIS Receive process shall identify a Loss of Code-group Delineation condition if the 10GBASE-R PCS signals, via the WIS_SIGNAL.request(PCS_R_STATUS) primitive of the WIS Service Interface (50.2.3), that synchronization has been lost and valid code-groups are no longer being delineated from the received payload stream being passed to the PCS. The LCD-P defect shall be reported to the far end WIS if this condition persists continuously for at least 3 ms. In addition, re-initialization of the WIS shall cause the LCD-P defect to be reported, with the same timing specification, until the PCS signals that valid code-groups are being delineated. The 3 millisecond timer is not allowed to start timing the false condition of WIS_SIGNAL.request(PCS_R_STATUS) while WIS_SIGNAL.indication is false, thus inhibiting the reporting of the LCD-P defect whenever the WIS is unable to supply valid received data to the PCS.

The WIS Receive process shall identify the absence of a Loss of Code-group Delineation condition if the 10GBASE-R PCS signals, also via the WIS_SIGNAL.request(PCS_R_STATUS) primitive, that synchronization has been regained and valid code-groups are being delineated from the received payload stream. Reporting of the LCD-P defect to the far end WIS shall terminate if the condition has been absent continuously for at least 1 ms.

50.3.6 Mapping between WIS and PMA

As a client of the PMA sublayer, the WIS utilizes the PMA Service Interface to transfer transmitted (received) data-groups to (from) the PMA sublayer. The PMA Service Interface is described in 51.2. Note that the PMA Service Interface may optionally be instantiated as a physical interface, referred to as the XSBI. In this case, the WIS must implement the client portion of the XSBI, including the logical formats and protocols as well as the physical signals and electrical timing, conforming with the XSBI specification in 51.4.

The WIS Transmit process transfers a continuous stream of 16-bit data-groups, containing the transmit WIS frame data generated as in 50.3.2, to the PMA via the PMA Service Interface. The octet boundaries in the transmit WIS frames shall be aligned to those in the data-groups, and the first data-group transmitted shall contain the first two octets in the Section Overhead (i.e., the first two A1 octets, as illustrated in Figure 50–8). Figure 50–11 depicts the transmission order required from the PMA for the data-groups. Each 16-bit data-group shall be transmitted from left to right, with lower-numbered bits being transmitted before higher-numbered bits. The data-groups themselves shall be transmitted in the sequence in which they are presented to the PMA by the WIS Transmit process.

A similar ordering is required during the reception process, as shown in Figure 50–11. Successive received bits shall be ordered from left to right in the 16-bit data-groups that are passed to the WIS (i.e., lower-numbered bits are received before higher-numbered bits), and the data-groups must be presented in the sequence in which they were received. The PMA sublayer is not required to align the data being presented to the WIS on any boundary.

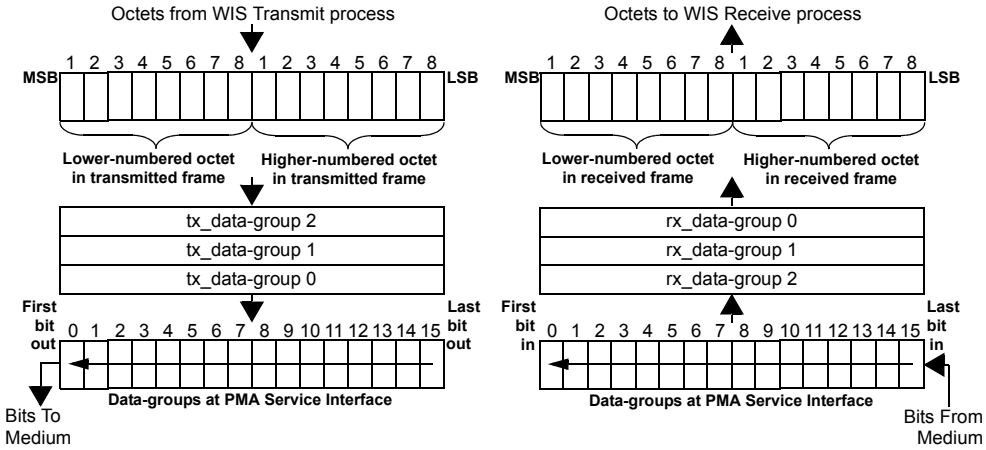

Figure 50–11—Transmission order

The WIS also utilizes the PMA_SIGNAL.indication primitive received from the PMA sublayer to determine when the PMA is unable to provide valid data to the WIS via its service interface. This primitive is used to unlock the state diagram implemented in the Synchronization process and force it to re-synchronize, as described in 50.4.

NOTE—The PMA_SIGNAL.indication primitive received from the PMA differs from the Loss of Signal (LOS) defect described in 50.3.2.5. LOS is a status condition that is reported to Layer Management for error monitoring purposes, but otherwise does not affect the internal functioning of the Receive Process or the Synchronization Process. The PMA primitive, however, acts as a control signal that directly affects all parts of the WIS receive functionality.

50.3.7 WIS data delay constraints

The sum of the transmit and receive data delays for any implementation of the WIS shall not exceed 14336 BT. Transmit data delay is measured from the input of a given unit of data by the PCS at the WIS service interface to the presentation of the same unit of data by the WIS to the PMA at the PMA service interface. Receive data delay is measured from the input of a given unit of data by the PMA at the PMA service interface to the presentation of the same unit of data by the WIS to the PCS at the WIS service interface. The time required to insert or process any necessary overhead or stuff octets must be included as part of the data delay incurred by the WIS. No constraint is placed on the individual values of the transmit and receive data delays for a given implementation, provided their sum falls within the above limit.

50.3.8 WIS test-pattern generator and checker

The WIS shall incorporate a test-pattern generator and a test-pattern checker to permit in-circuit testing using test patterns. These serial test patterns allow the 10GBASE-W PMA and PMD sublayers to be tested for compliance while in a system environment. Three patterns have been defined for testing: a fixed square wave pattern, an optional PRBS31 pattern, and a framed mixed frequency pattern. The patterns may be implemented at a bit or frame level and may be used for transmitter testing. The test-pattern checker has the ability to synchronize to the PRBS31 and mixed-frequency test patterns and report bit errors detected within the payload to the Station Management entity. The test-pattern checker is not required to synchronize to or analyze the square wave pattern.

The transmit and receive portions of the WIS can be placed in test-pattern mode separately by means of the WIS Transmit test-pattern enable and WIS Receive test-pattern enable control bits supported within its management registers (see 45.2.2.6). In addition, the three different types of transmit test pattern (square-wave, PRBS31 and mixed frequency) may be selected by means of the test-pattern control bits. (The test-pattern receiver only operates in PRBS31 and mixed-frequency test-pattern modes.) Errors detected during PRBS31 testing are recorded and provided to the Station Management entity via the 10G WIS test-pattern error counter register (45.2.2.8), while errors detected during mixed frequency testing are recorded and provided to the Station Management entity via the 10G WIS Section BIP Error Count register (45.2.2.17), the 10G WIS Line BIP Errors register (45.2.2.15), and the 10G WIS Path Block Error Count register (45.2.2.16). The WIS Service Interface is inoperative when either the WIS Transmit process or WIS Receive process is placed in test-pattern mode, regardless of the type of test pattern selected.

When the WIS transmit function is operating in test-pattern mode, the test-pattern generator produces a continuous test pattern that is sent 16 bits at a time to the underlying PMA sublayer via PMA_UNITDATA.request primitives. When the WIS receive function is operating in test-pattern mode, 16-bit data-groups received from the underlying PMA sublayer by means of PMA_UNITDATA.indication primitives are accepted and processed by the test-pattern checker.

50.3.8.1 Square wave test pattern

In the square-wave test-pattern mode, the WIS Transmit process is disabled or otherwise prevented from processing data, and a square wave as defined in 52.9.1 shall be continuously transferred to the PMA via the PMA Service Interface. No checking is performed on the 16-bit data-groups received from the PMA sublayer in this mode.

50.3.8.2 PRBS31 test pattern

The PRBS31 test-pattern mode is optional. When the transmit portion of the WIS is operating in PRBS31 test-pattern mode, the WIS Transmit process is disabled or otherwise prevented from processing data, and the output of a Pseudo-Random Bit Sequence (PRBS) generator shall be continuously transferred to the PMA via the PMA Service Interface. The PRBS generator functionality is described in 49.2.8.

When the receive portion of the WIS is operating in PRBS31 test-pattern mode, the WIS Receive Process is disabled or otherwise prevented from processing data, and a PRBS pattern checker shall check the bits received from the PMA via the PMA Service Interface. The PRBS checker functionality is described in 49.2.12.

If no errors have occurred, the test-pattern error signal from the pattern checker will be zero. An isolated bit error will cause the test-pattern error signal to be asserted three times; once when it is received and once when it is at each tap. The 10G WIS test-pattern error counter register (45.2.2.8) shall be incremented whenever the test-pattern error signal is asserted.

50.3.8.3 Mixed-frequency test pattern

In the mixed-frequency test-pattern mode, the WIS Transmit process is utilized to generate a transmitted test pattern (if activated by the WIS Transmit test-pattern enable control bit) and the WIS Synchronization and Receive processes are utilized to check a received test pattern (if separately activated by the WIS Receive test-pattern enable control bit). Figure 50–12 provides a functional diagram of the WIS when operating in this mode.

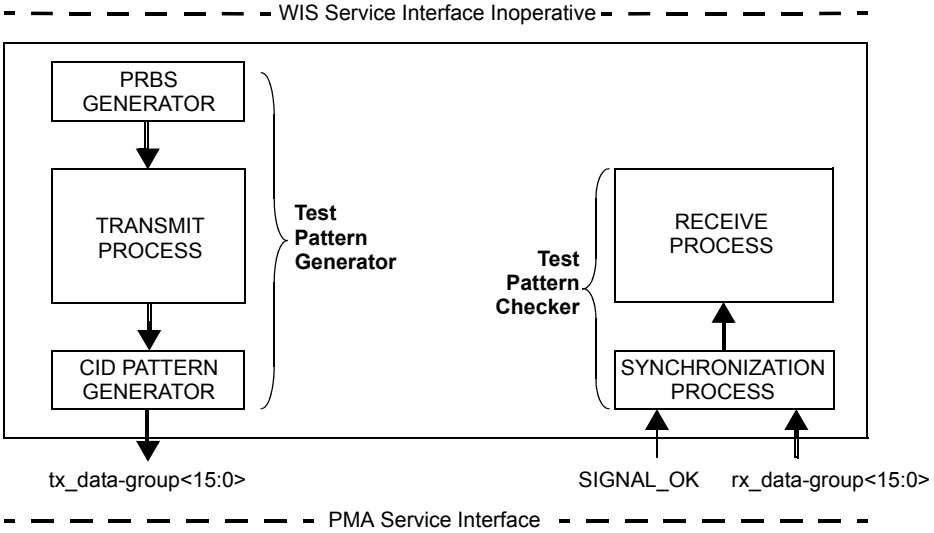

Figure 50–12—Mixed-frequency test-pattern mode operation

The mixed-frequency test pattern has the Test Signal Structure (TSS) described by 50.3.8.3.1, with a Consecutive Identical Digit (CID) pattern described by 50.3.8.3.2. The pattern is generated using the WIS Transmit process, with the SPE payload capacity being filled with a (2^{23}– 1) pseudo-random bit sequence (PRBS) as defined by ITU-T Recommendation O.150, 1996, in conjunction with a CID pattern generator that overwrites the last nine Z0 octets of the Section Overhead (see Figure 50–8) with an alternating pattern. The WIS Synchronization and Receive processes operate normally. Bit errors in the received pattern shall be detected by means of the Section, Line and Path BIP checking facilities, and reported to Station Management by means of the corresponding error counters within the MDIO register space (see 50.3.8).

NOTE 1—The PRBS is substituted for the payload data that would normally be sent to the WIS Transmit process from the 10GBASE-R PCS. The CID pattern is selected to stress the lock range of the receiver circuitry, and is placed in the Z0 octet locations as these are not scrambled by the WIS Transmit process. The values of the Z0 octets are required to be ignored by the WIS Receive process (see 50.3.2.4), and therefore no change to the latter is necessary.

The WIS Transmit, Receive and Synchronization processes shall function as described in 50.3.2, 50.3.3, and 50.3.4, with the exception that fault processing as per 50.3.2.5 and error propagation as per 50.3.5 shall not be carried out, the Z0 octets shall be overwritten with a CID pattern as specified, and the data output by the PRBS generator shall be used as the data source in place of the PCS.

NOTE 2—The implied re-use of the WIS Transmit, Receive, and Synchronization processes to implement the mixed-frequency test pattern functions is not a required attribute; implementations may select any physical means of realizing the test-pattern transmit and receive functionality, provided that the external behavior of the WIS conforms to that described herein.

50.3.8.3.1 Test Signal Structure (TSS)

The TSS is defined as a repeating set of two consecutive WIS frames, i.e., 300 672 octets in length, and shall be formatted as shown in Figure 50–13. The first frame of the set has a CID pattern set to all zeros, and the second frame has the pattern set to all ones. The PRBS generator shall be reset to all ones prior to the start of the payload capacity of each SPE contained within the TSS. In addition, the PRBS generator output shall be placed directly into the payload capacity of the first SPE, and inverted before being placed into the second SPE. Both WIS frames include the Section, Line, and Path Overhead and fixed stuff as normally generated by the Transmit process. All overhead octets except for J1 shall be set to their default values.

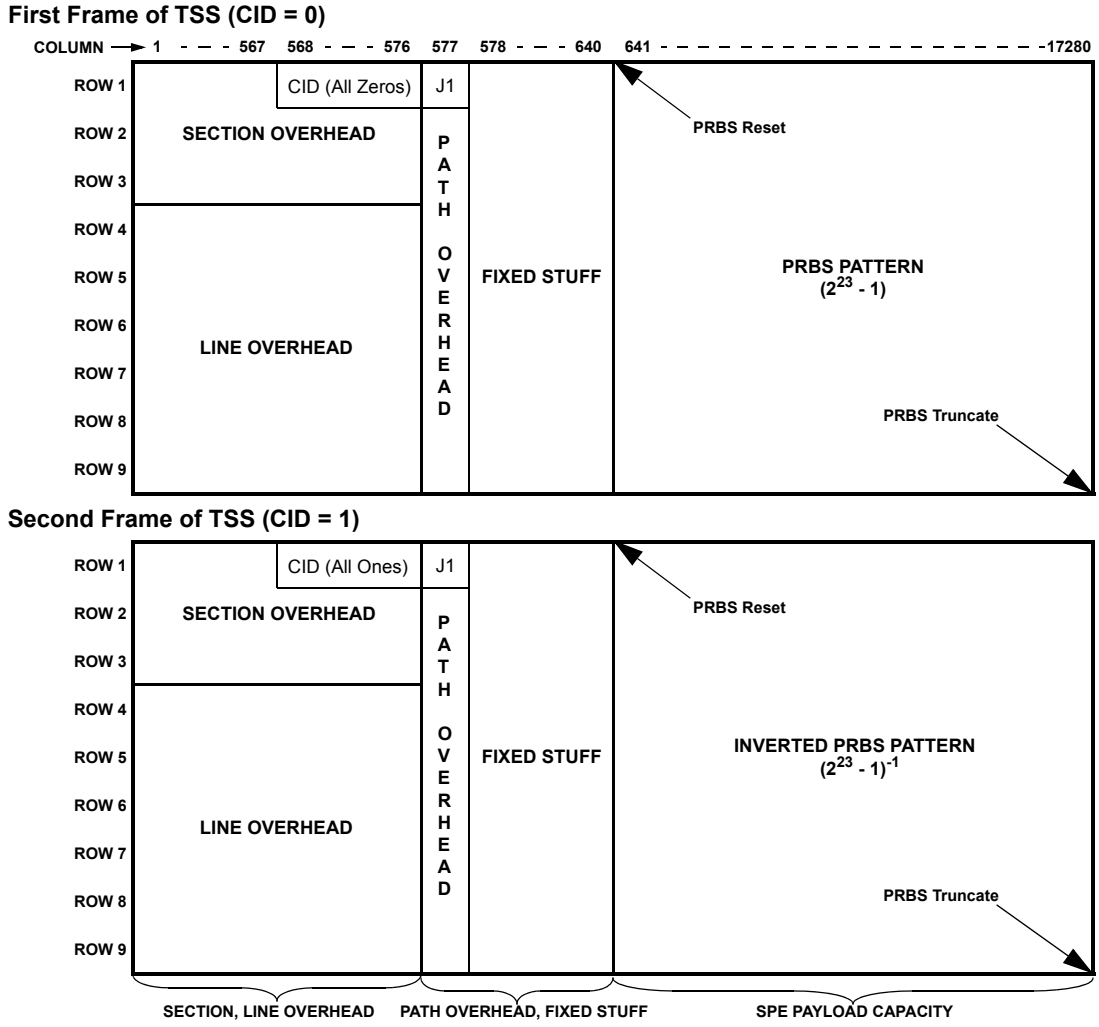

Figure 50–13—Test Signal Structure

As shown in Figure 50–13, the J1 octet of the Path Overhead is located immediately following the last octet of the CID pattern, and is hence the point of maximum stress with respect to the 10GBASE-W PMA and PMD. The J1 octet may therefore be programmed via the 10G WIS J1 transmit Registers (45.2.2.12) to a fixed value that, after scrambling, provides optimal stress to the PMA/PMD transmitter and receiver. If not used for this purpose, it shall be set to a value of 89 hexadecimal in both the WIS frames comprising the TSS.

NOTE—Standard SONET test equipment may not support the WIS test-pattern generator and checker as the PRBS is re-synchronized on every WIS frame, instead of the free-running PRBS described in ITU-T Recommendation O.172. Pattern-based bit error ratio testers will require 300 672 bytes of pattern storage memory to hold the complete TSS.

50.3.8.3.2 Continuous Identical Digits

The CID pattern shall comprise 9 octets (72 bits) overwriting the Z0 octets at the end of the first row of the Section Overhead, as shown in Figure 50–14. The pattern shall alternate in consecutive WIS frames; in the first frame of the TSS, the pattern consists of 72 bits of zeros, and in the second frame the pattern consists of 72 bits of all ones.

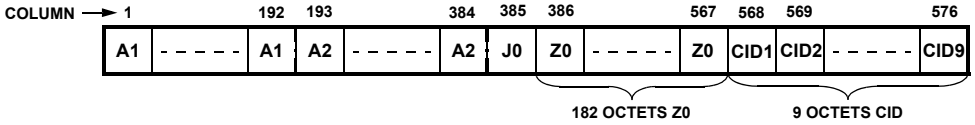

Figure 50–14—CID pattern structure

As the CID pattern is not scrambled, the string of ones (and, alternately, the string of zeros) forming the pattern produce transition-free areas within the TSS that stress the transmitter and receiver. Errors (eye closure at the transmitter or sampling errors at the receiver) are reflected as bit errors that are detected and reported via the BIP checking functionality in the WIS.

50.3.9 Loopback

The WIS is placed in Loopback mode when the Loopback bit in the WIS Control 1 register (45.2.2.1.2) is set to a logic one. In this mode, the WIS shall accept data on the transmit path from the 10GBASE-R PCS and return it on the receive path to the 10GBASE-R PCS. In addition, the WIS shall transmit a constant pattern to the PMA sublayer, and shall ignore all data presented to it by the PMA sublayer. The pattern output to the PMA transmit path at this time shall consist of a sequence of 8 logic zero bits and 8 logic one bits, forming the 16-bit word 00-FF hexadecimal. No SONET overhead or fixed stuff is output to the PMA at this time.

NOTE—The signal path through the WIS that is exercised in the Loopback mode of operation is implementation specific, but it is recommended that this signal path encompass as much of the WIS circuitry as is practical. The intention of providing this Loopback mode of operation is to permit diagnostic or self-test functions to test the transmit and receive data paths using actual data, while ensuring that remote entities do not interpret this test data as valid information. Other loopback signal paths may also be enabled independently using loopback controls within other devices or sublayers.

50.3.10 Link status

The WIS link status, as reported in the Link status bit of the WIS status 1 register (see 45.2.2.2.2), shall be set to down if a PLM-P, AIS-P, or LOP-P defect is detected (50.3.5.1), an AIS-L defect is detected (50.3.2.5), or the Synchronization process is not in the SYNC state (50.4.2). Otherwise, the WIS link status is set to up.

50.3.11 Management interface

The WIS supports a set of required and optional management objects to permit it to be controlled by the Station Management entity (STA). Access to management objects within the WIS is accomplished by means of a set of registers within the MDIO register space as defined in 45.2.2, which are implemented by the Link Management function depicted in Figure 50–2.

50.3.11.1 Management registers

The WIS management interface function shall utilize the following dedicated management registers:

a) WIS Control 1 register (Register 0)
b) WIS status 1 register (Register 1)
c) WIS Control 2 register (Register 7)
d) WIS status 2 register (Register 8)
e) WIS test-pattern error counter register (Register 9)
f) WIS Status 3 register (Register 33)
g) WIS J0 transmit register (Registers 64, 65, 66, 67, 68, 69, 70, 71)
h) WIS J0 receive register (Registers 72, 73, 74, 75, 76, 77, 78, 79)
i) WIS Far End Path Block Error Count register (Register 37)
j) WIS J1 transmit register (Registers 39, 40, 41, 42, 43, 44, 45, 46)
k) WIS J1 receive register (Registers 47, 48, 49, 50, 51, 52, 53, 54)
l) WIS Far End Line BIP Errors register pair (Registers 55 and 56)
m) WIS Line BIP Errors register pair (Register 57 and 58)
n) WIS Path Block Error Count register (Register 59)
o) WIS Section BIP Error Count register (Register 60)

These registers are implemented within the Device address space assigned to the WIS. The details of the register bit allocations and general usage are given in Clause 45. Note that Clause 45 also specifies several registers and register bits that are generic to every Device (e.g., Speed Ability and Identifier registers), and these are not covered by this clause.

NOTE—Clause 45 provides for additional vendor-specific registers within the address space assigned to the WIS. The presence, format, and function of these registers are implementation specific.

If the optional MDIO interface and associated register set are not implemented for the WIS, then equivalent capabilities are required to be provided.

50.3.11.2 WIS managed object class

The WIS Managed Object Class is defined in 30.8.1.

50.3.11.3 Management support objects

The following counters shall be provided to facilitate support of the WIS managed object class. If an MDIO interface is provided for the WIS, these counters are accessed via the WIS Far End Line BIP Errors and WIS Line BIP Errors registers, as described in 45.2.2.14 and 45.2.2.15 respectively. If no MDIO interface is implemented, these counters are to be accessible by equivalent means.

Far End Line BIP Errors

A 32-bit counter that is incremented by the number of Line BIP errors detected by the far-end WIS and reported via the M1 octet present in each received WIS frame (see 50.3.2.5). This counter has a maximum increment rate of 2 040 000 counts per second. The counter is cleared to zero when the WIS is reset, and wraps around to zero when it is incremented beyond the maximum count value. The contents of this counter are reflected in the WIS Far End Line BIP Errors register pair.

Line BIP Errors

A 32-bit counter that is incremented by the number of Line BIP errors detected by the Receive process while processing each received WIS frame (see 50.3.2.5). This counter has a maximum increment rate of 12 288 000 counts per second. The counter is cleared to zero when the WIS is reset, and wraps around to zero when it is incremented beyond the maximum count value. The contents of this counter are reflected in the WIS Line BIP Errors register pair.

50.4 Synchronization state diagram

The WIS Synchronization process shall implement the state diagrams specified in this subclause and meet the corresponding state diagram interface requirements, also specified here. Additional requirements to this state diagrams are made in 50.3.4. In the case of any ambiguity between the text and the state diagrams, the state diagrams shall take precedence.

The notation used in the state diagrams is described in 21.5, and the conventions for state diagram timers in 14.2.3.2. The notation ++ after a counter or integer variable indicates that its value is to be incremented by 1.

50.4.1 State diagram variables

Variables of the form name<x:0> indicate arrays or vectors of bits; for such variables, "<x>" indexes an element or set of elements in the array or vector, where "x" may be as follows:

— Any integer or set of integers.
— Any variable that takes on integer values.

In the variable and function descriptions in 50.4.1.2 and 50.4.1.3, the parameters *f*, *i*, *j*, *k*, *m*, and *n* are implementation-specific. Maximum and minimum values for these parameters are specified in 50.4.3.

50.4.1.1 Constants

A1

An octet value (bits 1:8) of 11110110 as assigned to the A1 framing character within the SONET Section Overhead, as specified in Table 1 of Section 4.2 of ATIS-0600416.1999(R2010). Used to obtain octet and WIS frame alignment.

A2

An octet value (bits 1:8) of 00101000 as assigned to the A2 framing character within the SONET Section Overhead, as specified in Table 1 of Section 4.2 of ATIS-0600416.1999(R2010). Used to obtain octet and WIS frame alignment.

Hunt_Pattern

A sequence of *i* consecutive A1 octets.

Presync_Pattern

A sequence of *j* consecutive A1 octets, immediately followed by a sequence of *k* consecutive A2 octets.

Sync_Pattern

A sequence of *f* consecutive A1 octets immediately followed by a sequence of *f* consecutive A2 octets.

50.4.1.2 Variables

power_on

A condition that is true until such time as the power supply to the device containing the WIS has reached the operating region.

Values: FALSE; The device is completely powered and ready to operate.
TRUE; The device has not been completely powered.

reset

A Boolean variable that is true when a hardware reset has been applied to the WIS, and false otherwise.

Values: TRUE; A hardware reset is being applied.
FALSE; A hardware reset is not being applied.

signal_fail

A condition that is true whenever the underlying PMA sublayer indicates by means of the PMA_SIGNAL.indication primitive that valid data are not being supplied to the WIS. Once set to TRUE, signal_fail remains set until the PMA sublayer explicitly indicates that valid data are being recovered and supplied to the WIS, after which it is set to FALSE.

Values: FALSE; A PMA_SIGNAL.indication(SIGNAL_OK) primitive was received with SIGNAL_OK set to OK, indicating that the PMA has begun to provide valid data to the WIS Receive process.
TRUE; A PMA_SIGNAL.indication(SIGNAL_OK) primitive was received with SIGNAL_OK set to FAIL, indicating that the PMA is unable to provide valid data to the WIS Receive process.

in_HUNT

Boolean variable, set by the Primary Synchronization state diagram to indicate to the Interval Pattern Search state diagram to indicate that it has entered, or returned to, the HUNT state.

Values: TRUE; The Primary Synchronization state diagram is presently in the HUNT state.
FALSE; The Primary Synchronization state diagram is not in the HUNT state.

sync_start

Boolean variable, set by the Primary Synchronization state diagram to trigger the operation of the Interval Pattern Search state diagram.

Values: TRUE; The Interval Pattern Search state diagram should begin scanning for framing sequences.
FALSE; The Interval Pattern Search state diagram should not begin scanning.

50.4.1.3 Functions

found_Hunt

For each bit input to the Synchronization process, this function indicates whether the Hunt_Pattern pattern has been detected in the bit string formed by concatenating the current bit with the (i * 8 – 1) consecutive previously input bits. If the number of bits previously input is less than (i * 8 – 1), this function outputs a FALSE value. Note that this function inspects its input on a bit-by-bit basis.

Values: TRUE; The (i * 8) bits examined so far, consisting of the current bit plus previously input bits, matches Hunt_Pattern.
FALSE; The pattern is not matched, or less than (i * 8) bits have been input so far.

found_Presync

For the last (*j* + *k)* sets of 8 bits (1 octet) input to the Synchronization process, this function indicates whether the Presync_Pattern pattern has been matched. If the number of octets previously input is less than (*j* + *k*), or the string match is still proceeding, this function outputs a WAIT value. Note that this function inspects its input on an octet-by-octet basis.

Values: TRUE; The octets examined so far, consisting of the current octet plus previously input bits, matches Presync_Pattern.
FALSE; The pattern is not matched.
WAIT; Insufficient data has been received so far to complete the match or declare failure.

found_Sync

For the last 2*f* sets of 8 bits (1 octet) input to the Synchronization process, this function indicates whether the Sync_Pattern pattern has been matched. If the number of octets previously input is less than 2*f*, the function outputs a FALSE value. Note that this function inspects its input on an octet-by-octet basis.

Values: TRUE; The octets examined so far, consisting of the current octet plus previously input octets, matches Sync_Pattern.
FALSE; The pattern is not matched, or insufficient octets have been received so far.

50.4.1.4 Counters

good_sync_cnt

Count of the number of consecutive candidate Sync_Pattern pattern locations that contain a valid Sync_Pattern pattern. This counter is controlled by the Interval Pattern Search state diagram, and used to report the number of valid framing locations to the Primary Synchronization state diagram.

bad_sync_cnt

Count of the number of consecutive candidate Sync_Pattern pattern locations that have failed to contain a valid Sync_Pattern pattern. This counter is controlled by the Interval Pattern Search state diagram, and used to report the number of mismatched framing locations to the Primary Synchronization state diagram.

octet_cnt

Count of octets input to the Synchronization process; always increments for every octet input. This counter is forced to zero in specific states of the Interval Pattern Search state diagram, but for all other states it increments by 1 for each octet received. It is used to force the state diagram to wait in a given state until a required number of octets have been received.

50.4.2 State diagram

The Primary Synchronization state diagram is depicted in Figure 50–15, and the Interval Pattern Search state diagram in Figure 50–16. The WIS shall implement the Synchronization process according to these state diagrams, including compliance with the associated state variables as specified in 50.4.1. The Synchronization process is responsible for detecting octet and WIS frame boundaries and also for determining whether the underlying receive channel is ready for operation. Failure to achieve synchronization, as indicated by the Synchronization state diagram not being in the SYNC state, shall cause the WIS Receive process to suspend normal operation and propagate the error to the PCS as described in 50.3.5.

The function of the Primary Synchronization state diagram is to track the four phases through which the Synchronization process must progress in order to achieve and maintain lock to the incoming frame stream. In the first phase, the state diagram performs a bit-by-bit hunt over the input data for a string of A1 octets.

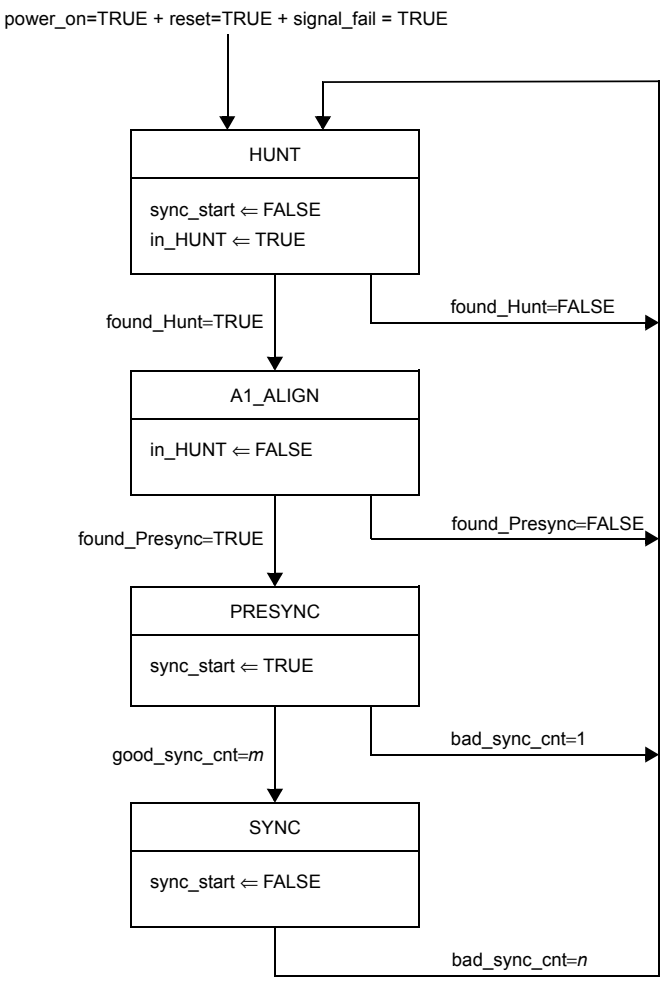

Figure 50–15—Primary Synchronization state diagram

Successful location of several A1 octets indicates both a candidate octet and a candidate frame boundary; this is validated by scanning for additional A1 octets followed immediately by another string of A2 octets, which corresponds to the framing pattern present at the start of every WIS frame. The candidate frame boundary is now confirmed by searching for several consecutive such patterns, each separated by a frame size; if successful, frame lock is deemed to have been achieved. The state diagram then remains in the frame lock state until a number of consecutive candidate pattern locations have failed to match the A1/A2 pattern. This behavior implements the necessary hysteresis required to reject transient bit errors.

The Interval Pattern Search state diagram is activated when the Synchronization process has successfully located the first occurrence of the framing pattern, and performs the actual function of searching for subsequent framing patterns that are located 155 520 octets apart, and counting failed and successful matches.

The functions fount_Hunt, found_Presync, and found_Sync implement bit and octet scanners that search through the incoming data stream for the respective data patterns that must be matched.

NOTE—Implementations may choose any conforming means of realizing the state diagram described herein, provided that the external behavior of the Synchronization process is unchanged.

50.4.3 Parameter values

Table 50–5 identifies the minimum and maximum values permissible for the parameters *f*, *i*, *j*, *k*, *m*, and *n* that are used in the preceding state diagram descriptions. Implementations shall set these parameters to values within the limits specified in the table.

Table 50–5—Minimum and maximum parameter values

Parameters	Minimum values	Maximum values	Purpose
f	2	192	Controls width of Sync_Pattern pattern
i	1	192	Controls width of Hunt_Pattern pattern
j	16	190	Controls width of Presync_Pattern pattern
k	16	192	Controls width of Presync_Pattern pattern
m	4	8	Controls hysteresis for SYNC state entry
n	1	8	Controls hysteresis for SYNC state exit

NOTE—The ranges for the parameters stipulated in Table 50–5 are intended to provide the implementer with latitude in creating implementations of the Synchronization process. However, adherence to the stipulated minimum values for the parameters will result in an implementation that provides the minimum time-to-frame (5 WIS frame periods) and required bit error tolerance at the normal BER (10^{-12}).

50.5 Environmental specifications

All equipment subject to this clause shall conform to the requirements of 14.7 and applicable sections of ISO/IEC 11801:1995.

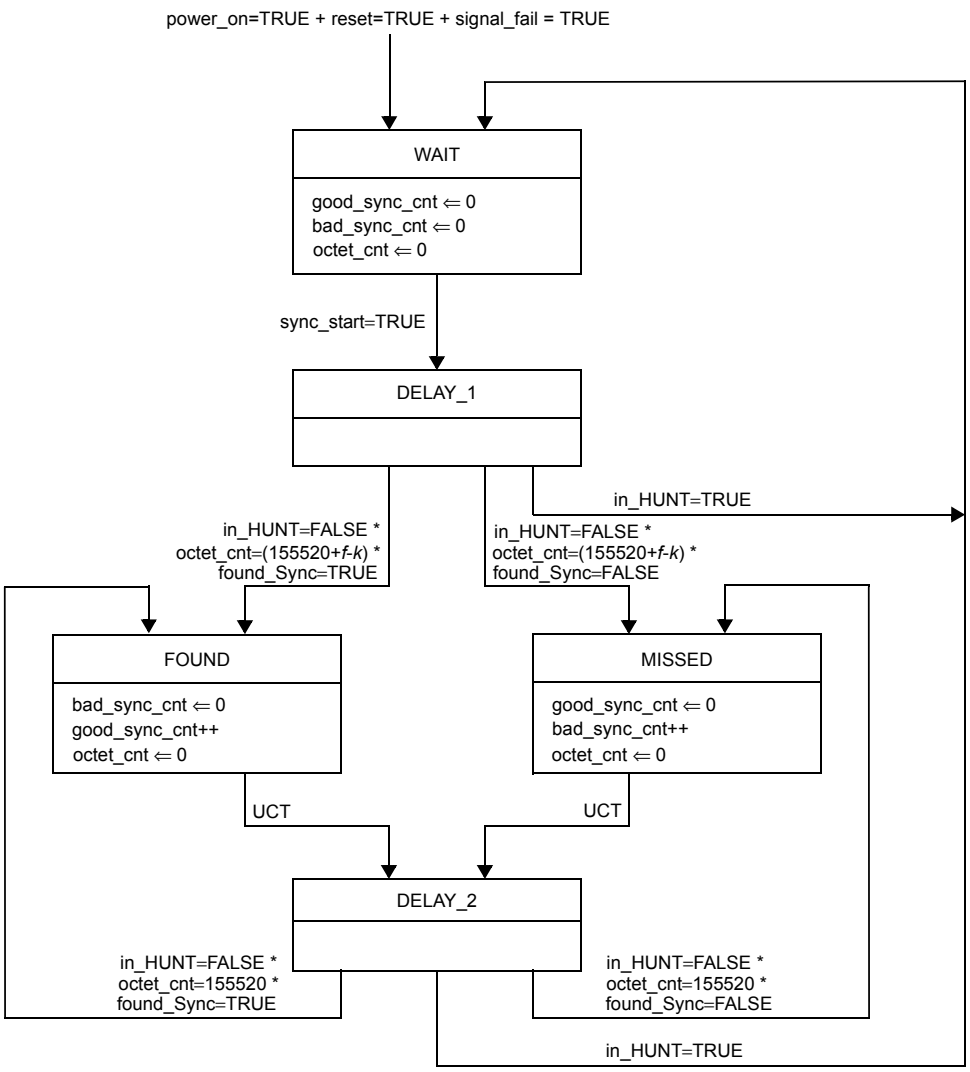

Figure 50–16—Interval Pattern Search state diagram

50.6 Protocol implementation conformance statement (PICS) proforma for Clause 50, WAN Interface Sublayer (WIS), type 10GBASE-W[11]

50.6.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 50, WAN Interface Sublayer, type 10GBASE-W, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the same, can be found in Clause 21.

50.6.2 Identification

50.6.2.1 Implementation identification

Supplier[1]	
Contact point for inquiries about the PICS[1]	
Implementation Name(s) and Version(s)[1,3]	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)[2]	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

50.6.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, WAN Interface Sublayer (WIS), type 10GBASE-W
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [ ] Yes [ ] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2015.)	

Date of Statement	

[11] *Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

50.6.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
XSBI	XSBI compatibility interface	51, 50.3.6	Compatibility interface supported	O	Yes [ ] No [ ]
*MD	MDIO	50.3.11.1	Registers and interface supported	O	Yes [ ] No [ ]
*PRBS	PRBS31 Test-pattern mode	50.3.8.2	Registers and functionality supported	O	Yes [ ] No [ ]

50.6.4 PICS proforma tables for the WAN Interface Sublayer (WIS), type 10GBASE-W

50.6.4.1 Compatibility considerations

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	Meets requirements of ATIS-0600416.1999(R2010)	50.1.1	Except as specifically excluded	M	
CC2	Precedence of ATIS-0600416.1999(R2010) and ANSI T1.269-2000 in case of discrepancies	50.3.2, 50.3.2.1, 50.3.2.3, 50.3.3	Unless specifically overridden by Clause 50	M	
CC3	Environmental specifications	50.5		M	Yes [ ]

50.6.4.2 WIS transmit functions

Item	Feature	Subclause	Value/Comment	Status	Support
WT1	Bit significance of data-unit vector from 10GBASE-R PCS	50.2.1.1	Data-unit<0> is interpreted as the least-significant bit	M	Yes [ ]
WT2	Payload mapping of octets from data-units to SPE by Transmit process	50.3.1.1	Tx_data-unit<7:0> mapped to lower-numbered (even) octet and tx_data-unit<15:8> to next higher-numbered (odd) octet	M	Yes [ ]
WT3	Bit relabeling during transmit mapping	50.3.1.1	Ethernet FCS error detection properties to be maintained	M	Yes [ ]
WT4	Path Overhead generation	50.3.2.1		M	Yes [ ]
WT5	J1 octet in Path Overhead	50.3.2.1	Transports a 16-octet repeating Path Trace Message	M	Yes [ ]
WT6	Default Path Trace Message	50.3.2.1	Consists of 15 octets of zeros and a header octet	M	Yes [ ]
WT7	Line Overhead generation	50.3.2.2		M	Yes [ ]
WT8	Unspecified Line Overhead octets	50.3.2.2	Set to 00000000 binary if not specifically listed	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
WT9	H1 and H2 octets in Line Overhead	50.3.2.2	Set to 522 decimal and also indicates concatenated payload	M	Yes [ ]
WT10	K1 octet in Line Overhead	50.3.2.2	Set to 00000000 binary	M	Yes [ ]
WT11	Bits 1 to 5 of K2 octet in Line Overhead	50.3.2.2	Set to 00000 binary	M	Yes [ ]
WT12	Bits 6 to 8 of K2 octet in Line Overhead	50.3.2.2	Set to either an RDI-L encoding, or 000 binary	M	Yes [ ]
WT13	Section Overhead generation	50.3.2.3		M	Yes [ ]
WT14	Unspecified Section Overhead octets	50.3.2.3	Set to 00000000 binary if not specifically listed	M	Yes [ ]
WT15	J0 octet in Section Overhead	50.3.2.3	Transports a 16-octet repeating Section Trace Message	M	Yes [ ]
WT16	Default Section Trace Message	50.3.2.3	Consists of 15 octets of zeros and a header octet	M	Yes [ ]
WT17	Frame scrambler	50.3.3		M	Yes [ ]
WT18	Octet ordering to PMA	50.3.6	First data-group transmitted contains the first two A1 octets	M	Yes [ ]
WT19	Bit ordering to PMA	50.3.6	Ethernet FCS error detection properties to be maintained	M	Yes [ ]
WT20	Data-group ordering to PMA	50.3.6		M	Yes [ ]
WT21	Transmit data delay constraint	50.3.7	Sum of transmit and receive data delays not to exceed 14336 BT	M	Yes [ ]

50.6.4.3 WIS receive functions

Item	Feature	Subclause	Value/Comment	Status	Support
WR1	Bit significance of data-unit vector to 10GBASE-R PCS	50.2.2.1	Data-unit<0> is interpreted as the least-significant bit	M	Yes [ ]
WR2	Payload mapping of octets from SPE to data-units	50.3.1.2	Rx_data-unit<7:0> mapped from lower-numbered (even) octet and rx_data-unit<15:8> from next higher-numbered (odd) octet	M	Yes [ ]
WR3	Bit relabeling during receive mapping	50.3.1.2	Ethernet FCS error detection properties to be maintained	M	Yes [ ]
WR4	Section, Line, and Path Overhead extraction and processing	50.3.2.4		M	Yes [ ]
WR5	Unsupported overhead octets	50.3.2.4	Ignored if marked as unsupported in Tables 50–1, 50–2, and 50–3	M	Yes [ ]
WR6	Unlisted overhead octets	50.3.2.4	Ignored if not listed in Table 1 of ATIS-0600416.1999(R2010)	M	Yes [ ]
WR7	Interpretation of J1 octet in Path Overhead	50.3.2.4	Interpreted as carrying a 16-byte repeating Path Trace Message	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
WR8	Extraction of valid Path Trace Message data	50.3.2.4	After receiver accomplishes synchronization	M	Yes [ ]
WR9	Interpretation of bits 6 through 8 of K2 octet in Line Overhead	50.3.2.4	Interpreted as carrying AIS-L or RDI-L	M	Yes [ ]
WR10	Unsupported encodings of bits 6 through 8 of K2 octet in Line Overhead	50.3.2.4	Ignored by receiver	M	Yes [ ]
WR11	Interpretation of J0 octet in Path Overhead	50.3.2.4	Interpreted as carrying a 16-byte repeating Section Trace Message	M	Yes [ ]
WR12	Extraction of valid Section Trace Message data	50.3.2.4	After receiver accomplishes synchronization	M	Yes [ ]
WR13	Detection of Section, Line and Path Defects and Anomalies	50.3.2.5	Only Defects and Anomalies listed in Table 50–4 processed	M	Yes [ ]
WR14	Support for PLM-P and LCD-P defects	50.3.2.5		M	Yes [ ]
WR15	Detection of LOS defect	50.3.2.5	No services provided by the PMA or PMD are used	M	Yes [ ]
WR16	LOS defect detection parameters	50.3.2.5	Parameters *T* and *T'* between 2.3 and 100 microseconds	M	Yes [ ]
WR17	Frame descrambler	50.3.3		M	Yes [ ]
WR18	Bit, octet and data-group ordering from PMA	50.3.6	Ethernet FCS error detection properties to be maintained	M	Yes [ ]
WR19	Delineation of octet and WIS frame boundaries	50.3.4	According to Synchronization process	M	Yes [ ]
WR20	Reporting of errors on synchronization failure	50.3.4	SEF or LOF defects reported on synchronization failure	M	Yes [ ]
WR21	Receive data delay constraint	50.3.7	Sum of transmit and receive data delays not to exceed 14336 BT	M	Yes [ ]

50.6.4.4 State diagrams

Item	Feature	Subclause	Value/Comment	Status	Support
SD1	Synchronization state diagram	50.4	Meets the requirements of Figures 50–15 and 50–16	M	Yes [ ]
SD2	Synchronization process	50.4.2		M	Yes [ ]
SD3	Receive process suspended upon failure to achieve synchronization	50.4.2		M	Yes [ ]
SD4	Synchronization parameters	50.4.3	Meets the requirements of Table 50–5	M	Yes [ ]

50.6.4.5 Error notification

Item	Feature	Subclause	Value/Comment	Status	Support
EN1	Notification of fault presence to PCS	50.3.5	WIS_SIGNAL.indication SIGNAL_OK set to FAIL	M	Yes [ ]
EN2	Removal of fault notification to PCS in the absence of faults	50.3.5	WIS_SIGNAL.indication SIGNAL_OK set to OK	M	Yes [ ]
EN3	Faults notified to PCS	50.3.5.1	Detects and reports the following errors: Synchronization process not in SYNC state, PLM-P defect, AIS-P defect, and LOP-P defect	M	Yes [ ]
EN4	PCS fault notification timing	50.3.5.2	Error propagation terminates, and valid data transferred, within 125 microseconds after error conditions removed	M	Yes [ ]
EN5	Presence of Loss of Code-group Delineation condition	50.3.5.3	When indicated by the 10GBASE-R PCS	M	Yes [ ]
EN6	Reporting of Loss of Code-group Delineation presence	50.3.5.3	Reported only if condition persists for more than 3 ms	M	Yes [ ]
EN7	Reporting with respect to WIS re-initialization	50.3.5.3	Report Loss of Code-group Delineation on re-initialization	M	Yes [ ]
EN8	Absence of Loss of Code-group Delineation condition	50.3.5.3	When indicated by the 10GBASE-R PCS	M	Yes [ ]
EN9	Reporting of Loss of Code-group Delineation absence	50.3.5.3	Reported only if condition is absent for more than 1 ms	M	Yes [ ]

50.6.4.6 Management registers and functions

Item	Feature	Subclause	Value/Comment	Status	Support
MR1	Management registers	50.3.11.1	If management registers are not physically implemented, equivalent capabilities must be provided	MD:M	Yes [ ]
MR2	Management support objects	50.3.11.3	If management interface is not physically implemented, equivalent capabilities must be provided	MD:M	Yes [ ]
MR3	Loopback function	50.3.9	Data accepted on the transmit path from the 10GBASE-R PCS returned via the receive path	M	Yes [ ]
MR4	Data transmitted to PMA during loopback	50.3.9	The 16-bit pattern 00-FF hexadecimal must be transmitted	M	Yes [ ]
MR5	Data received from PMA during loopback	50.3.9	All data received from PMA during loopback must be ignored	M	Yes [ ]
MR6	WIS link status function	50.3.10	Set to up or down as defined	M	Yes [ ]

50.6.4.7 WIS test-pattern generator and checker

Item	Feature	Subclause	Value/Comment	Status	Support
TP1	Test-pattern generator and checker	50.3.8	Separate generator and checker functionality for transmit and receive	M	Yes [ ]
TP2	Data pattern used in square-wave test-pattern mode	50.3.8.1	Data pattern as specified in 52.9.1	M	Yes [ ]
TP3	PRBS31 generator functionality	50.3.8.2		PRBS:M	Yes [ ]
TP4	PRBS31 checker functionality	50.3.8.2		PRBS:M	Yes [ ]
TP5	Bit error detection and reporting by PRBS31 checker	50.3.8.2	If management interface is not physically implemented, equivalent reporting capabilities must be provided	PRBS:M	Yes [ ]
TP6	Bit error detection in received mixed-frequency test pattern	50.3.8.3	Detection by Section, Line and Path BIP checking facilities	M	Yes [ ]
TP7	Functionality of WIS Transmit, Receive and Synchronization processed in mixed-frequency test-pattern mode	50.3.8.3	With the exceptions that fault processing and error propagation are not carried out, the Z0 octets are overwritten with CID pattern, and a PRBS used in place of the PCS data	M	Yes [ ]
TP8	Format of TSS	50.3.8.3.1		M	Yes [ ]
TP9	PRBS generator reset in TSS	50.3.8.3.1	Reset to all-ones prior to start of SPE	M	Yes [ ]
TP10	PRBS generator invert in TSS	50.3.8.3.1	Invert PRBS pattern in second SPE	M	Yes [ ]
TP11	Overhead octets in TSS	50.3.8.3.1	Set to defaults except for J1	M	Yes [ ]
TP12	Default value for J1 in TSS	50.3.8.3.1	Set to 89 hexadecimal if not used for providing stress	M	Yes [ ]
TP13	CID pattern used in TSS	50.3.8.3.2	72 bits of ones and 72 bits of zeros	M	Yes [ ]
TP14	Alternating value for CID pattern	50.3.8.3.2	72 bits of zeros in first frame and 72 bits of ones in second	M	Yes [ ]

51. Physical Medium Attachment (PMA) sublayer, type Serial

51.1 Overview

This clause defines the functional characteristics for the Physical Media Attachment (PMA) used in 10GBASE-R and 10GBASE-W. Figure 51–1 depicts the relationships of the serial PMA (shown shaded) with other sublayers and the ISO/IEC Open System Interconnection (OSI) reference model.

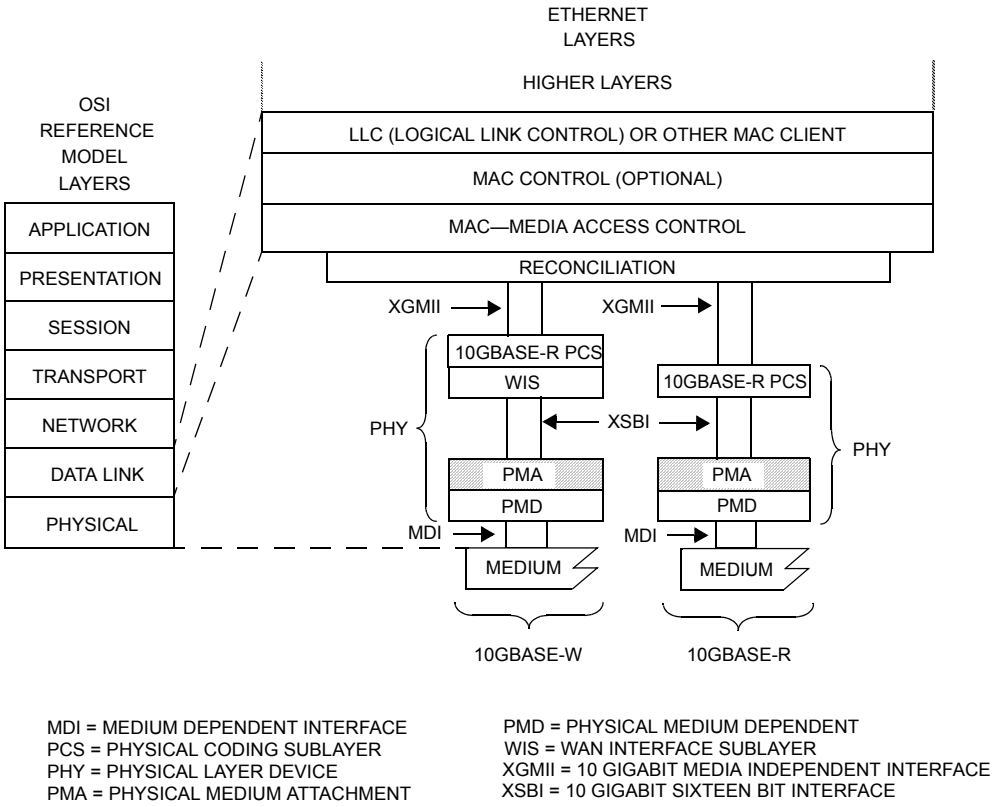

Figure 51–1—Relationship of serial PMA to PCS, WIS, and PMD

The purpose of the serial PMA is to attach the PMD of choice to its client, i.e., the PCS or WIS sublayer.

51.1.1 Scope

This clause specifies the functions, features, and services of the serial PMA sublayer, including electrical and timing specifications for the 10 Gigabit sixteen bit interface (XSBI). Specifications for the serial input and output of the PMA are left to the implementer with conformance and consideration of the serial PMD type. The serial PMA is used between the PMD and the PMA client as shown in Figure 51–1.

51.1.2 Summary of functions

The following is a summary of the principal functions implemented by the PMA.

In the transmit direction (i.e., transmitting data from the PMA client to the PMD), the PMA performs the following functions:

a) Provide transmit source clock to PMA client.
b) Serialization of 16-bit data to serial bit stream.
c) Transmission of serial data to PMD.

In the receive direction (i.e., serial data from PMD to the PMA client), the PMA performs the following functions:

a1) Bit clock recovery of serial data from PMD.
b1) Provide receive clock to PMA client.
c1) Deserialization of serial data to16-bit parallel data.
d1) Transmission of parallel data to PMA client.
e1) Provide link status information.

51.2 PMA Service Interface

The Serial PMA provides a Service Interface to the 10G BASE-R PCS or WIS sublayer, i.e., the PMA client. These services are described in an abstract manner and do not imply any particular implementation. The PMA Service Interface shall support the exchange of data-groups between the PMA and the PMA client. The PMA converts data-groups into bits and passes these to the PMD, and vice versa. It also generates an additional status indication for use by its client.

The following primitives are defined:

PMA_UNITDATA.request(tx_data-group<15:0>)

PMA_UNITDATA.indication(rx_data-group<15:0>)

PMA_SIGNAL.indication(SIGNAL_OK)

PMA_RXMODE.request(rx_mode)

PMA_TXMODE.request(tx_mode)

PMA_ENERGY.indication(energy_detect)

51.2.1 PMA_UNITDATA.request

This primitive defines the transfer of data (in the form of data-groups) from the PMA client to the PMA. PMA_UNITDATA.request is generated by the PMA client's transmit process.

51.2.1.1 Semantics of the service primitive

PMA_UNITDATA.request(tx_data-group<15:0>)

The data conveyed by PMA_UNITDATA.request is a 16 bit vector representing a single data-unit which has been prepared for transmission by the PMA client.

51.2.1.2 When generated

The PMA client continuously sends tx_data-group<15:0> to the PMA at a nominal clock rate of 644.53125 MHz and 622.08 MHz in 10GBASE-R and 10GBASE-W operations, respectively.

51.2.1.3 Effect of receipt

Upon receipt of this primitive, the PMA generates a series of sixteen PMD_UNITDATA.request primitives, requesting transmission of the indicated tx_data-group bit to the PMD.

51.2.2 PMA_UNITDATA.indication

This primitive defines the transfer of data (in the form of data-groups) from the PMA to its client. PMA_UNITDATA.indication is used by the client's synchronization process.

51.2.2.1 Semantics of the service primitive

PMA_UNITDATA.indication(rx_data-group<15:0>)

The data conveyed by PMA_UNITDATA.indication is a 16 bit vector representing a single data-unit which has been prepared for transmission by the PMA receive process to the PMA client.

51.2.2.2 When generated

The PMA continuously sends one rx_data-group<15:0> to the PMA client corresponding to the receipt of each data-group of sixteen PMD_UNITDATA.indication primitives received from the PMD. The PMA sends the rx_data-group<15:0> at a nominal rate derived from the recovered bit clock.

51.2.2.3 Effect of receipt

The effect of receipt of this primitive by the client is unspecified by the PMA sublayer.

51.2.3 PMA_SIGNAL.indication

This primitive is sent by the PMA to its client to indicate the status of the receive process. PMA_SIGNAL.indication is generated by the PMA receive process to propagate the detection of severe error conditions (e.g. no valid signal being received from the PMD sublayer) to the PMA client.

51.2.3.1 Semantics of the service primitive

PMA_SIGNAL.indication (SIGNAL_OK)

The SIGNAL_OK can take one of two values: OK or FAIL. A value of FAIL denotes that invalid data is being presented to the PMA client. A value of OK does not guarantee valid data is being presented to the PMA client.

51.2.3.2 When generated

The PMA generates a PMA_SIGNAL.indication primitive to the PMA client whenever there is change in the value of the SIGNAL_OK parameter.

51.2.3.3 Effect of receipt

The effect of receipt of this primitive by the client is unspecified by the PMA sublayer.

51.2.4 PMA_RXMODE.request

This primitive is generated by the PCS Receive Process for EEE capability (see 78.3) to indicate when the PMA and PMD receive functions may go into a low power mode, see 49.3.6.6. Without EEE capability, the primitive is never invoked and the PMA behaves as if rx_mode = DATA.

51.2.4.1 Semantics of the service primitive

PMA_RXMODE.request(rx_mode)

The rx_mode parameter takes on one of two values: QUIET or DATA.

51.2.4.2 When generated

The PCS generates this primitive to indicate the low power mode of the receive path.

51.2.4.3 Effect of receipt

When received the PMA receive is configured appropriately for the indicated state and the value is propagated to PMD_RX_MODE.request(rx_mode). When rx_mode is DATA, the PMA operates normally. When rx_mode is QUIET, the PMA may go into a low power mode.

51.2.5 PMA_TXMODE.request

This primitive is generated by the PCS Transmit Process for EEE capability to invoke the appropriate PMA and PMD transmit EEE states, see 49.2.13.3.1. Without EEE capability, the primitive is never invoked and the PMA behaves as if tx_mode = DATA.

51.2.5.1 Semantics of the service primitive

PMA_TXMODE.request(tx_mode)

The tx_mode parameter takes on one of three values: QUIET, ALERT, or DATA.

51.2.5.2 When generated

The PCS generates this primitive to indicate the low power mode of the transmit path.

51.2.5.3 Effect of receipt

When received the PMA transmit is configured appropriately for the indicated state and the value is propagated to PMD_TX_MODE.request(tx_mode). When tx_mode is DATA, the PMA operates normally. When tx_mode is QUIET, the PMA may go into a low power mode. When tx_mode is ALERT, the PMA operation is not defined.

51.2.6 PMA_ENERGY.indication

This primitive is sent by the PMA to its client to indicate the status of the receive process for EEE capability. PMA_ENERGY.indication is generated by the PMA receive process to propagate the energy detection indication from the PMD to the PMA client.

51.2.6.1 Semantics of the service primitive

PMA_ENERGY.indication(energy_detect)

The energy_detect parameter is Boolean and reflects the state of the PMD_SIGNAL.-indication(SIGNAL_OK) received from the PMD.

51.2.6.2 When generated

The PMA generates this primitive whenever there is a change in the value of the SIGNAL_OK parameter.

51.2.6.3 Effect of receipt

The effect of receipt of this primitive by the client is unspecified by the PMA sublayer.

51.3 Functions within the PMA

The PMA comprises the PMA Transmit and PMA Receive processes for 10GBASE-W and 10GBASE-R. The PMA Transmit process serializes the tx_data-groups and passes them to the PMD for transmission on the underlying medium. Similarly, the PMA Receive process deserializes received data from the PMD and presents the data as rx_data-groups to the PMA client. The PMA receiver continuously conveys sixteen-bit data-groups to the PMA client, independent of data-group alignment.

NOTE—Strict adherence to manufacturer-supplied guidelines for the operation and use of PMA serializer components is required to meet the jitter specifications of the respective PMD clause. The supplied guidelines should address the quality of power supply filtering associated with the transmit clock generator, and also the purity of reference clock fed to the transmit clock generator.

51.3.1 PMA transmit function

The PMA Transmit function passes data unaltered (except for serializing) from the PMA client directly to the PMD. Upon receipt of a PMA_UNITDATA.request primitive, the PMA Transmit function shall serialize the sixteen bits of the tx_data-group<15:0> parameter and transmit them to the PMD in the form of sixteen successive PMD_UNITDATA.request primitives.

51.3.2 PMA receive function

The PMA Receive function passes data unaltered from the PMD directly to the PMA client. Upon receipt of sixteen successive PMD_UNITDATA.indication primitives, the PMA shall assemble the sixteen received bits into a single sixteen-bit value and pass that value to the PMA client as the rx_data-group<15:0> parameter of the primitive PMA_UNITDATA.indication. The PMA receive function does not align rx_data-group<15:0> to the original tx_data-group<15:0> from the remote end of the link.

51.3.3 Delay Constraints

The PMA receives a one bit data stream from the PMD and presents a sixteen bit wide data unit to the PMA client. Received bits from the PMD are buffered to facilitate proper deserialization of the rx_data-group<15:0> to the PMA client. These functions necessitate an internal PMA delay of at least sixteen unit intervals. In practice, this serial to parallel conversion may necessitate even longer delays of the incoming data stream.

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementers will conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices. The sum of transmit and receive delay constraints for the serial PMA/PMD sublayer shall meet the requirements as specified in the respective PMD clause. The serial PMA/PMD sublayer includes the serial PMA, the serial PMD, and 2 m of fiber.

51.4 Sixteen-Bit Interface (XSBI)

A physical instantiation of the sixteen-bit PMA service interface (XSBI) is defined to provide compatibility among devices designed by different manufacturers for the operation of PHY implementations in either the 10GBASE-R family or the 10GBASE-W family. There is no requirement for an integrated compliant device to implement or expose the XSBI. A XSBI implementation is described in 51.4 through 51.9. Though the XSBI is an optional interface, it is used extensively in this standard as a basis for specification. The PMA is specified to the XSBI interface, so if the XSBI is not physically implemented, a conforming implementation shall behave as if the XSBI functions were implemented.

Table 51–1 defines terms used within the description of the XSBI interface and the rest of this clause. For the XSBI instantiation of the PMA service interface, the mapping of bits between the PMA service interface and the XSBI physical interface is defined for transmit in Table 51–2 and for receive in Table 51–3. Figure 51–2 depicts the relationship between the XSBI physical instantiation and the service interface provided by the PMA to its client.

NOTE—Document OIF SFI-4-01.0, an implementer's agreement, was used as a basis for the development of the XSBI instantiation.

Table 51–1—Terms and definitions

Term	Definition
10GBASE-R nominal baud rate	10.3125 GBd
10GBASE-W nominal baud rate	9.95328 GBd
Data signals, data-group, clock signal	Operating rate is 1/16 of nominal baud rate
Differential signal name used without <P,N> (e.g., PMA_TX_CLK)	= signal<P> minus signal <N> (e.g., = PMA_TX_CLK<P-N>)

Table 51–2—Transmit bit mapping

	Name	Bit numbers															
PMA service interface	tx_data-group	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
XSBI physical interface	xsbi_tx	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

Table 51–3—Receive bit mapping

	Name	Bit numbers															
PMA service interface	rx_data-group	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
XSBI physical interface	xsbi_rx	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

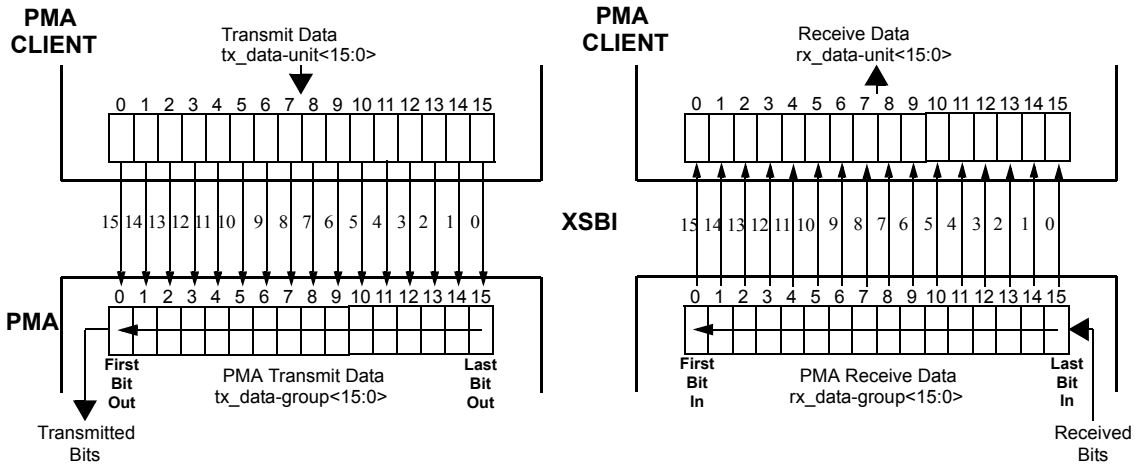

Figure 51–2—XSBI physical interface mapping into PMA service interface

As depicted in Figure 51–3, the XSBI connects the PMA to its client, the PCS or WIS sublayer. It is equipped for full duplex transmission of data-groups with associated data-group clock signals. The PMA client provides data-groups on xsbi_tx<15:0> to the XSBI transmit function which latches the data on the rising edge of the PMA_TX_CLK. See section 51.6.1.2 for details. The PMA_TX_CLK signal is derived from PMA_TXCLK_SRC as provided by the PMA. An internal Transmit Clock Generation Unit (TXCGU) uses REFCLK to generate the internal bit clock that is used to serialize the latched data out of the PMA outputs.

The PMA Receive function accepts serial data from the PMD and extracts a bit clock and recovered data from the serial inputs in the Receive Clock Recovery Unit (RXCRU). The recovered data is deserialized and conveyed to the PMA client on xsbi_rx<15:0>. The rising edge of the recovered clock, PMA_RX_CLK, which is at 1/16 the bit rate, is used by the PMA to send the received 16-bit data-groups to the PMA client.

The PMA_SIGNAL.indication is a function of PMD_SIGNAL.indication, the Sync_Err signal and the optional PMA loopback signal. These signals can be seen in context in Figure 44A–7 and Figure 51–3. In the case of PMA loopback being inactive, the PMA_SIGNAL will indicate a FAIL whenever the PMD_SIGNAL.indication indicates a FAIL. The PMA_SIGNAL will also indicate a FAIL when Sync_Err is valid, i.e., PMA unable to recover clock from the incoming data stream. If the PMA loopback function is implemented and activated, the PMA_SIGNAL.indication will ignore the PMD_SIGNAL.indication and behave as if PMD_SIGNAL.indication is valid.

PMA_RX_CLK is derived from the serial input data whenever possible. When there is no valid input signal or the optional Synchronization Error (Sync_Err) asserted or other conditions where the PMA cannot derive the clock from the serial input data, a valid PMA_RX_CLK is provided as described in 51.7.2.

51.4.1 Required signals

In the event this XSBI is made accessible, the signals listed in Table 51–4 are provided, with the meanings described elsewhere in this subclause.

xsbi_tx<15:0>
: The 16-bit parallel low-voltage differential signaling (LVDS) data presented to the XSBI for serialization and transmission onto the media. Bit xsbi_tx<15> shall be transmitted first, followed by xsbi_tx<14> through xsbi_tx<0>.

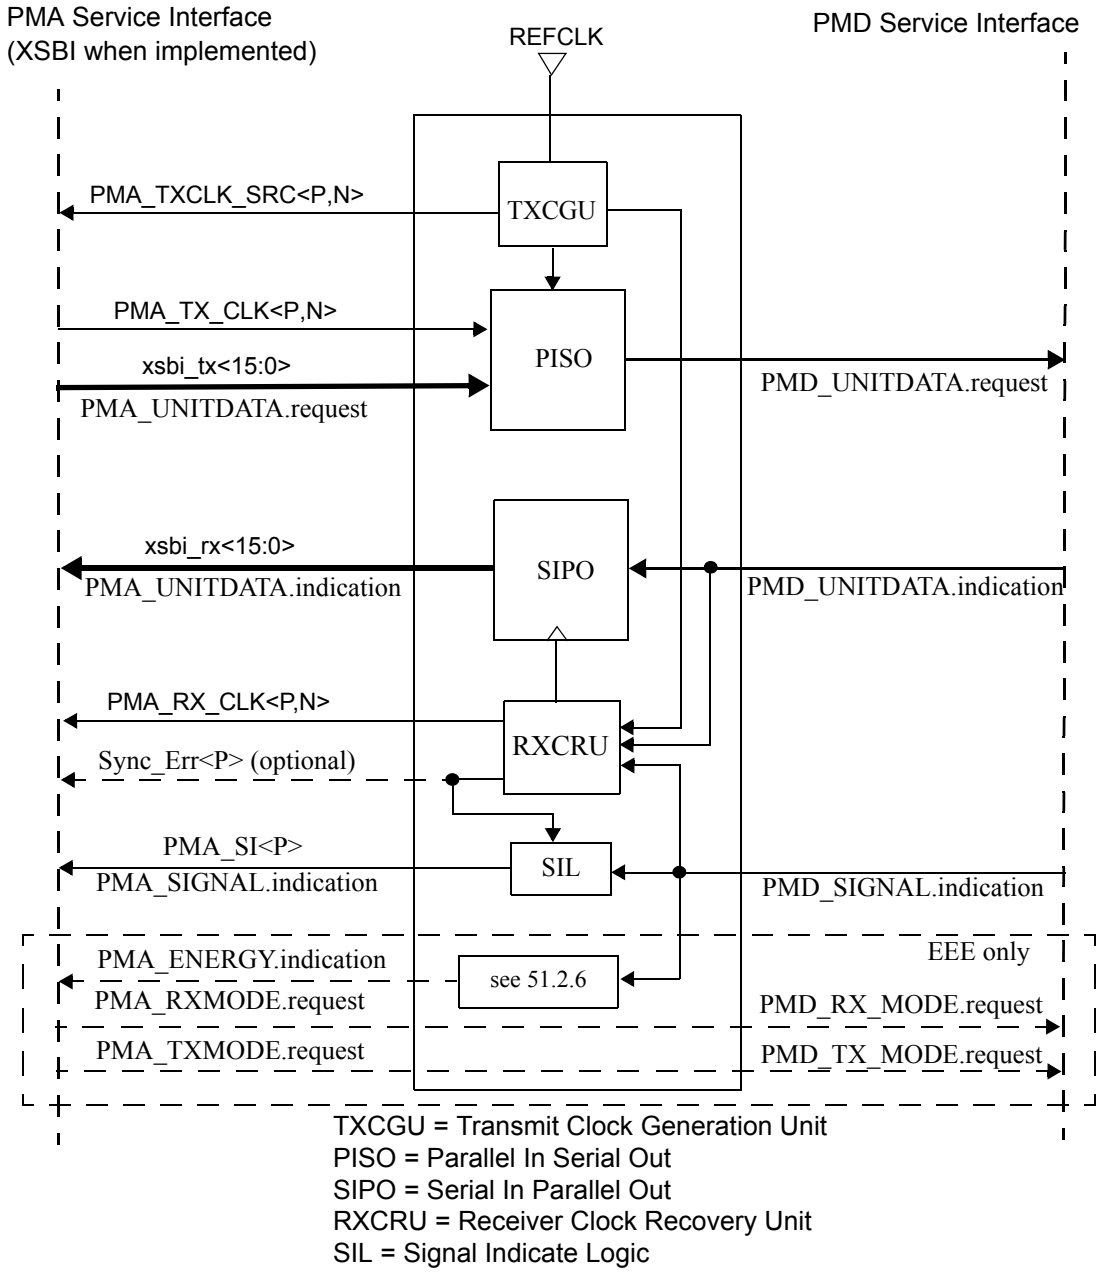

Figure 51–3—XSBI reference diagram

PMA_TX_CLK<P,N>

The LVDS transmit data-group clock. This data-group clock is used to latch data into the PMA for transmission. The PMA_TX_CLK<P,N> shall be derived from PMA_TXCLK_SRC<P,N>. The rising edge of PMA_TX_CLK is used to latch data into the PMA for transmission. Refer to 51.6.1.2 for details.

Table 51–4—XSBI required signals

Symbol	Signal name	Signal type	Active level
xsbi_tx<15:0>	Transmit Data	Input	H
PMA_TX_CLK<P,N>	Transmit Clock	Input	↑
PMA_TXCLK_SRC<P,N>	Transmit Clock Source	Output	either edge
xsbi_rx<15:0>	Receive Data	Output	H
PMA_RX_CLK<P,N>	Receive Clock	Output	↑
PMA_SI<P>	PMA Signal Indicate	Output	H

PMA_TXCLK_SRC<P,N>
: The LVDS transmit data-group clock source. The PMA derives the PMA_TXCLK_SRC<P,N> from the local REFCLK. The PMA_TXCLK_SRC<P,N> is used by the PMA client to derive PMA_TX_CLK<P,N>.

xsbi_rx<15:0>
: The 16-bit parallel LVDS data transmitted by the PMA and presented to the PMA client for further processing. The PMA deserializes the incoming data to 16-bit wide data groups, xsbi_rx<15:0>. Bit xsbi_rx<15> shall be received first, followed by xsbi_rx<14> through xsbi_rx<0>. Data alignment is NOT performed by the PMA.

PMA_RX_CLK<P,N>
: The LVDS receive data-group clock. The PMA derives the PMA_RX_CLK<P,N> from the line rate when the serial data from the PMD is present, or the REFCLK when the serial data from the PMD is absent, e.g. PMD_SI invalid or Sync_Err asserted. The xsbi_rx<15:0> data groups are presented to the PMA client on the rising edge of PMA_RX_CLK<P,N>.

PMA_SI<P>
: A logic 0 on PMA Signal Indicate (PMA_SI) indicates that the receiver is unable to properly recover the receive data. The indicator is a function of the PMD_SIGNAL.indication status, the Sync_Err function and the optional PMA loopback signal. This signal is the physical instantiation of the SIGNAL_OK parameter as described in the PMA_SIGNAL.indication primitive. This signal shall be compliant with the EIA/JESD8-B Interface Standards for Nominal 3V/3.3V Supply Digital Integrated Circuits.

51.4.2 Optional Signals

Sync_Err<P>
: This signal is used to indicate the inability of the PMA to recover the clock from the serial data stream. A logic high indicates that there is a synchronization error. A logic low does not guarantee synchronization. This signal shall be compliant with the EIA/JESD8-B Interface Standards for Nominal 3V/3.3V Supply Digital Integrated Circuits.

energy_detect
: If the optional Energy-Efficient Ethernet (EEE) function is supported (see Clause 78) then the XSBI interface includes energy_detect as described in 51.2.

rx_quiet
: If the optional EEE function is supported (see Clause 78) then the XSBI interface includes rx_quiet as described in 51.2.

tx_quiet

If the optional EEE function is supported (see Clause 78) then the XSBI interface includes tx_quiet as described in 51.2.

51.5 General electrical characteristics of the XSBI

In the event this XSBI is made accessible, this subclause specifies the general electrical characteristics of the XSBI.

NOTE—All LVDS AC and DC parameters shall conform to the TIA/EIA-644 LVDS specification with the exception of those parameters specified by the tables of this clause.

51.5.1 DC characteristics

Table 51–5 documents the required DC parametric attributes of all inputs and outputs of the XSBI. The signal specification is based on low voltage differential signals (LVDS) as described in ANSI/TIA-644 LVDS specifications. Unless otherwise stated, all terms and test conditions related to electrical parameters (AC or DC) are as specified in the TIA/EIA-644 LVDS document.

Table 51–5—Electrical specifications

Symbol	Parameter	Conditions	Min	Max	Units	Note
$\|V_{OD}\|$	Output differential voltage	Differential load, $R_{load} = 100\Omega \pm 1\%$	250	400	mV	
V_{OS}	Output offset voltage	Differential load, $R_{load} = 100\Omega \pm 1\%$	1125	1375	mV	
R_O	Output impedance, single ended		40	140	Ω	[a]
ΔR_O	R_O mismatch			10	%	*
$\|\Delta V_{OD}\|$	Change in VOD between “0” and “1”	Differential load, $R_{load} = 100\Omega \pm 1\%$		50	mV	
$\|\Delta V_{OS}\|$	Change in VOS between “0” and “1”	Differential load, $R_{load} = 100\Omega \pm 1\%$		50	mV	
V_I	Input voltage range	|Vgpd| < 50mV	900	1600	mV	[b]
V_{ID}	Input differential voltage	|Vgpd| < 50mV	100	600	mV	
R_{IN}	Receiver differential input impedance		70	130	Ω	*
t_R, t_F	20%-80% rise and fall times		100	400	ps	*

[a]Modified from or not specified in TIA/EIA-644 document.
[b]$|V_{gpd}|$ is the ground potential differential between PMA client and PMA.

51.5.2 Valid signal levels

All AC measurements are made from the cross-over level of the clock to the valid input or output data levels as shown in Figure 51–4. The CLK and DATA is defined to be the differential signal, i.e., <P-N>, of a clock or data signal.

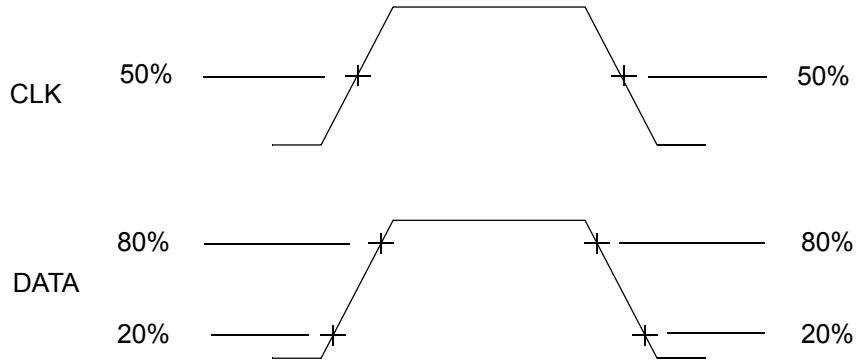

Figure 51–4—Input/output valid level for AC measurements

51.5.3 Rise and fall time definition

The rise and fall time definition for PMA_TX_CLK, PMA_TXCLK_SRC, PMA_RX_CLK, and DATA is shown below in Figure 51–5. All signals are measured differentially.

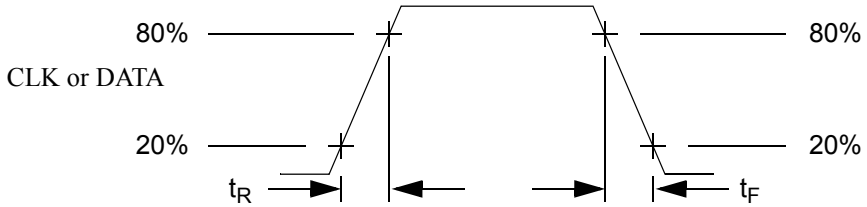

Figure 51–5—Rise and fall time definition

51.5.4 Output load

All AC measurements are assumed to have a differential output load of 100 Ω ± 1%.

51.6 XSBI transmit interface electrical characteristics

In the event this XSBI is made accessible, the electrical characteristics of the XSBI transmit interface are specified in this subclause. A XSBI receiver shall meet the specifications in sections 51.6.1 through 51.6.2.

NOTE—The following approach is taken for positioning clocks relative to the data. For both the PMA and PMA client drivers, the <P-N> clock edges are aligned to the data edges, to allow simplification of macro design. For both the PMA and PMA client receivers, the <P-N> clock edges are centered on the data bit, to allow simplification of macro design. The implementation to meet these requirements may be achieved on the system board. This can be done with either a delay of the clocks or by exchanging the positive and negative signals of the differential clock outputs.

51.6.1 XSBI transmit interface timing

The XSBI transmit interface timing specification consists of two separate specifications. The timing specifications at the PMA client output and the PMA input are described in 51.6.1.1 and 51.6.1.2, respectively. Data is latched into the PMA using rising edge of PMA_TX_CLK. The relationship between

PMA_TX_CLK and PMA_TXCLK_SRC are described in 51.6.2. All transitions in Figure 51–6 and Figure 51–7 are specified from the PMA_TX_CLK reference level, to valid input signal levels (refer to 51.5.2).

51.6.1.1 PMA client output timing

Figure 51–6 and Table 51–6 detail the XSBI timing requirements for transmit data relative to PMA_TX_CLK at the PMA client outputs.

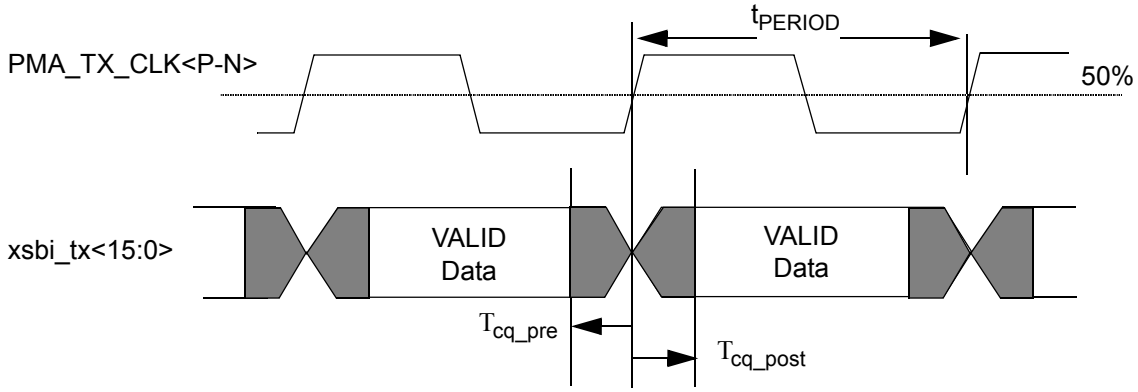

Figure 51–6—PMA client output timing

Table 51–6—PMA client output timing specifications

Parameter	Description	Min	Typ	Max	Units
t_{PERIOD}	PMA_TX_CLK Period 10GBASE-R 10GBASE-W		 1.55151 1.60751		ns[a]
T_{cq_pre}	data invalid window before ↑ PMA_TX_CLK<N-P>			200	ps
T_{cq_post}	data invalid window after ↑ PMA_TX_CLK<N-P>			200	ps
t_{DUTY}	PMA_TX_CLK Duty Cycle	40		60	%

[a]Period time value =1/644.53125 MHz for 10GBASE-R, similarly = 1/622.08 MHz for 10GBASE-W.

51.6.1.2 PMA input timing

Figure 51–7 and Table 51–7 detail the XSBI timing requirements for transmit data relative to the PMA_TX_CLK at the PMA inputs.

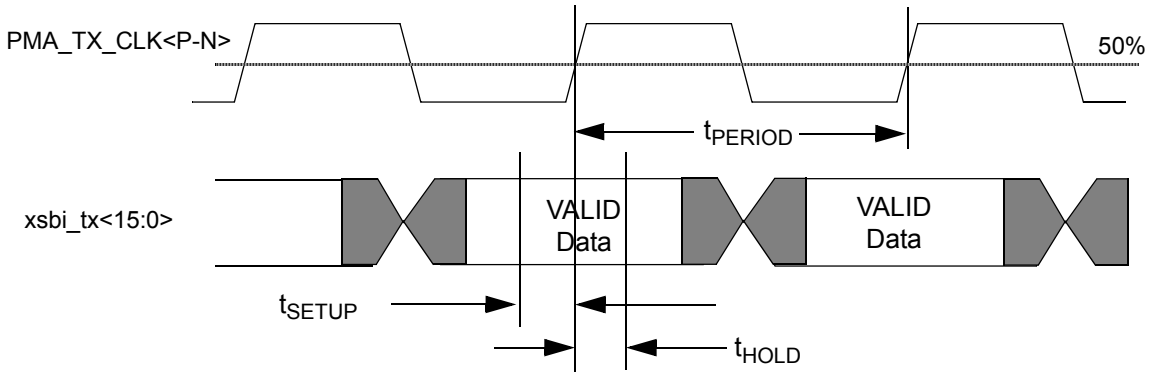

Figure 51–7—PMA input timing

Table 51–7—PMA input timing specifications

Parameter	Description	Min	Typ	Max	Units
t_{PERIOD}	PMA_TX_CLK period 10GBASE-R 10GBASE-W		 1.55151 1.60751		ns
t_{SETUP}	Data Setup to ↑ PMA_TX_CLK	300			ps
t_{HOLD}	Data Hold from ↑ PMA_TX_CLK	300			ps
t_{DUTY}	PMA_TX_CLK duty cycle	40		60	%

51.6.2 XSBI PMA_TX_CLK and PMA_TXCLK_SRC Specification

Table 51–8 specifies the XSBI interface transmit frequency. The frequency of xsbi_tx<15:0> and PMA_TX_CLK are derived from the frequency of PMA_TXCLK_SRC.

Table 51–9 specifies the phase variation between PMA_TX_CLK and PMA_TXCLK_SRC and the jitter requirements between PMA_TX_CLK and PMA_TXCLK_SRC. This specification describes the maximum delay variation through the PMA client (PMA_TXCLK_SRC to PMA_TX_CLK and xsbi_tx<15:0>). The absolute value of the delay is not specified, as it is not relevant to the operation of the interface.

Table 51–8—Transmit source clock specification

Parameter	Description	Value
$f_{PMA\_TXCLK\_SRC}$	PMA_TXCLK_SRC frequency 10GBASE-W 10GBASE-R	 622.08 MHz ± 20 ppm 644.53125 MHz ± 100 ppm

Table 51–9—Transmitter clocks specification

Parameter	Description	Condition	Value(max)
TD	Variation of PMA_TXCLK_SRC to PMA_TX_CLK delay	frequency < 10 kHz	2ns(p-p)
CJ	Jitter between PMA_TXCLK_SRC and PMA_TX_CLK	frequency >10 kHz	175ps(p-p)

NOTE—The ± 20 ppm clock tolerance is not intended to interoperate directly with interfaces that comply with SONET of SDH standards, or other synchronous networks. Operation over electrically multiplexed payloads of a transmission network is outside the scope of this standard.

The TD parameter is defined to be the peak-to-peak variation of the phase offset measured in time between an arbitrary clock edge of the PMA_TXCLK_SRC clock and a selected clock edge on the PMA_TX_CLK. Letting "X" be the offset parameter, then $TD = X_{max} - X_{min}$.

51.7 XSBI receive interface electrical characteristics

In the event this XSBI is made accessible, the electrical characteristics of the XSBI receive interface are specified in this subclause. A XSBI receiver shall meet the specifications in 51.7.1 through 51.7.2.

NOTE—The following approach is taken for positioning clocks relative to the data. For both the PMA and PMA client drivers, the <P-N> clock edges are aligned to the data edges, to allow simplification of macro design. For both the PMA and PMA client receivers, the <P-N> clock edges are centered on the data bit, to allow simplification of macro design. The implementation to meet these requirements may be achieved on the system board. This can be done with either a delay of the clocks or by exchanging the positive and negative signals of the differential clock outputs.

51.7.1 XSBI receive interface timing

The XSBI receive interface timing specification consists of two separate specifications. The timing specifications at the PMA output and the PMA client input are described in 51.7.1.1 and 51.7.1.2, respectively. Data is latched into the PMA client using the rising edge of PMA_RX_CLK. All transitions in Figure 51–8 and Figure 51–9 are specified from the PMA_RX_CLK reference level, to valid output signal levels (refer to 51.5.2).

51.7.1.1 PMA output timing

Figure 51–8 and Table 51–10 detail the XSBI timing requirements for receive data relative to the PMA_RX_CLK at the PMA outputs.

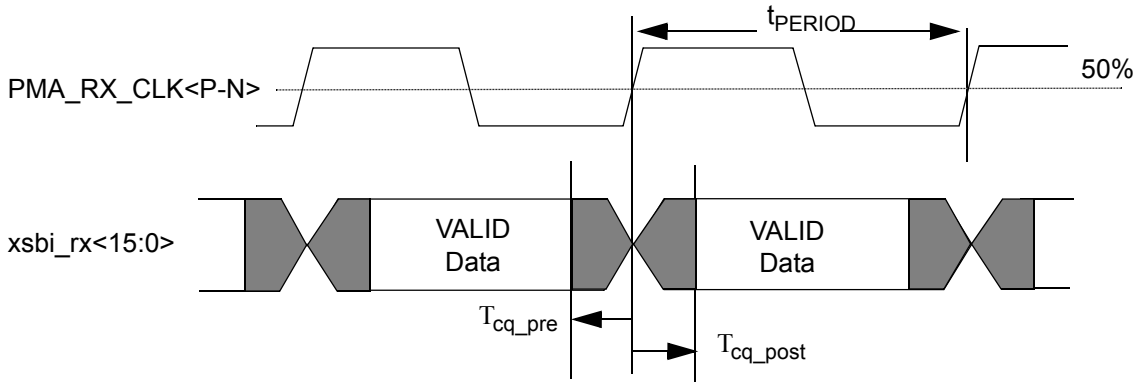

Figure 51–8—PMA output timing

Table 51–10—PMA output timing specifications

Parameter	Description	Min	Typical	Max	Units
t_{PERIOD}	PMA_RX_CLK period 10GBASE-R 10GBASE-W	—	 1.55151 1.60751	—	ns
T_{cq_pre}	data invalid window before ↑ PMA_RX_CLK<N-P>	—	—	200	ps
T_{cq_post}	data invalid window after ↑ PMA_RX_CLK<N-P>	—	—	200	ps
t_{DUTY}	PMA_RX_CLK duty cycle	45	—	55	%

51.7.1.2 PMA client input timing

Figure 51–9 and Table 51–11 detail the XSBI timing requirements for receive data relative to the PMA_RX_CLK at the PMA client inputs.

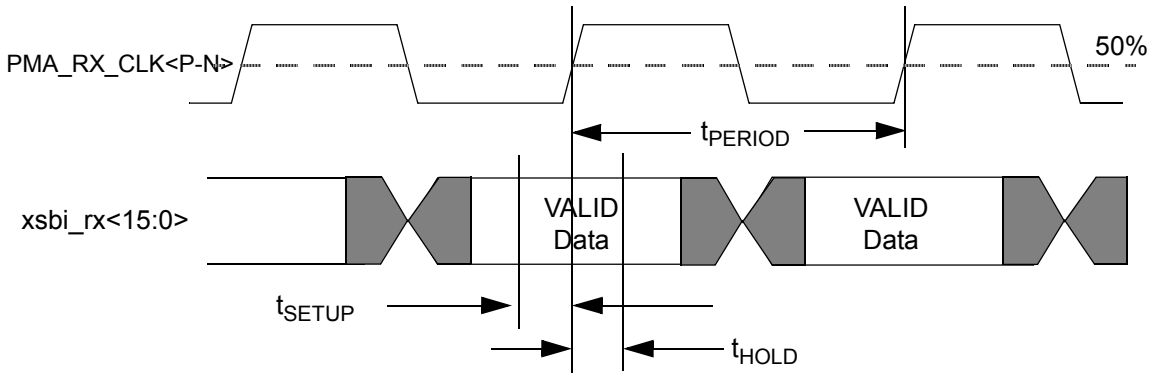

Figure 51–9—PMA client input timing

Table 51–11—PMA client input timing specifications

Parameter	Description	Min	Typical	Max	Units
t_{DUTY}	PMA_RX_CLK duty cycle	45	—	55	%
t_{SETUP}	Data Setup before ↑PMA_RX_CLK	300	—	—	ps
t_{HOLD}	Data Hold after ↑PMA_RX_CLK	300	—	—	ps

51.7.2 XSBI PMA_RX_CLK specification

Table 51–12 defines the PMA_RX_CLK frequency variation during PMA_SI valid and invalid conditions. PMA_RX_CLK is derived from the serial input data whenever possible. Under a loss-of-signal (LOS) or the optional Synchronization Error (Sync_Err) or other conditions where the PMA cannot derive the clock from the serial input data, a valid PMA_RX_CLK shall be provided with frequency characteristics as defined in Table 51–12. During the transitions from nominal clock to recovered clock or from recovered clock to nominal clock, the period and duty cycle requirements do not apply. However, the minimum pulse width should not change during the transitions between clock sources. During the transitions, the PMA_RX_CLK pulse width shall not be less than the minimum that is calculated by the period times the duty cycle as defined in Table 51–10 and Table 51–12.

51.8 PMA loopback mode (optional)

PMA loopback is optional. If PMA loopback is implemented it shall conform to the requirements of this subclause, i.e., 51.8.

Table 51–12—Receiver clock specification

PMA_SI	Parameter	Description	Value
valid	$f_{PMA\_RX\_CLK}$	PMA_RX_CLK frequency 10GBASE-W 10GBASE-R	 622.08 MHz ± 20 ppm 644.53125 MHz ± 100 ppm
invalid	$f_{PMA\_RX\_CLK}$	PMA_RX_CLK frequency 10GBASE-W 10GBASE-R	 622.08 MHz ± 2500 ppm 644.53125 MHz ± 2500 ppm

If a Clause 45 MDIO is supported, then this function maps to the PMA loopback function as specified in 45.2.1.1.5. A device is placed in Loopback mode when the loopback bit in the PMA/PMD Control register 1 is set to a logic one. A device is removed from Loopback mode when this bit is set to a logic zero.

When loopback mode is selected, transmission requests passed to the transmitter are shunted directly to the receiver, overriding any signal detected by the receiver on its attached link. A device is explicitly placed in loopback mode (i.e., loopback mode is not the normal mode of operation of a device). The method of implementing loopback mode is not defined by this standard.

NOTE—Loopback mode may be implemented either in the parallel or the serial circuitry of a device.

51.9 Environmental specifications

All equipment subject to this clause shall conform to the requirements of 14.7 and applicable sections of ISO/IEC 11801:1995.

51.10 Protocol implementation conformance statement (PICS) proforma for Clause 51, Physical Medium Attachment (PMA) sublayer, type Serial[12]

51.10.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 51, PMA Interface Sublayer, type Serial, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the same, can be found in Clause 21.

51.10.2 Identification

51.10.2.1 Implementation identification

Supplier[1]	
Contact point for inquiries about the PICS[1]	
Implementation Name(s) and Version(s)[1,3]	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)[2]	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

51.10.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 51, PMA Interface Sublayer (PMA), type Serial
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [ ] Yes [ ] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2015.)	

Date of Statement	

[12]*Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

51.10.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
LPI	Implementation of LPI	51.2		O	Yes [ ] No [ ]
*XSBI	XSBI interface	51.4	Physical instantiation of PMA service interface	O	Yes [ ] No [ ]
M1	XSBI functionality	51.4	Behave functionally as if the XSBI is implemented	M	Yes [ ]
M2	PMA Service Interface	51.2	Supports data group transfer between PMA and PMA client	M	Yes [ ]
M3	PMA_SI	51.4.1	EIA/JESD8-B compliance	M	Yes [ ]
M4	Sync_Err	51.4.2	Local PMA synchronization indicator and EIA/ JESD8-B compliance	O	Yes [ ] No [ ]
M5	Local PMA loopback	51.8		O	Yes [ ] No [ ]
MD	MDIO interface	51.8	Registers and interface supported	O	Yes [ ] No [ ]

51.10.4 PICS proforma tables for the PMA Interface Sublayer, type Serial

51.10.4.1 Compatibility considerations

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	Environmental specifications	51.9		M	Yes [ ]

51.10.4.2 PMA transmit functions

Item	Feature	Subclause	Value/Comment	Status	Support
PT1	Serialization and transmission of datagroup	51.3.1		M	Yes [ ]
PT2	Order of transmission	51.4.1	Bit xsbi_tx<15> shall be transmitted first	M	Yes [ ]
PT3	PMA_TX_CLK	51.4.1	PMA_TX_CLK derived from PMA_TXCLK_SRC	XSBI:M	Yes [ ] N/A[ ]
PT4	LVDS electrical compliance	51.5	Conformance to TIA/EIA644 LVDS specifications and to Table 51–5	XSBI:M	Yes [ ] N/A[ ]
PT5	Transmit electrical specifications	51.6	Electrical and timing specifications	XSBI:M	Yes [ ] N/A[ ]

51.10.4.3 PMA receive functions

Item	Feature	Subclause	Value/Comment	Status	Support
PR1	Deserialization and transmissions of datagroup to PMA client	51.3.2		M	Yes [ ]
PR2	Order of reception	51.4.1	Bit xsbi_rx<15> shall be received first	M	Yes [ ]
PR3	LVDS electrical compliance	51.5	conformance to TIA/EIA644 LVDS specifications and to Table 51–5	XSBI:M	Yes [ ] N/A:[ ]
PR4	Transmit electrical specifications	51.7	Electrical and timing specifications	XSBI:M	Yes [ ] N/A[ ]
PR5	Valid PMA_RX_CLK	51.7.2	A valid clock is provided as defined in Table 51–12	XSBI:M	Yes [ ] N/A[ ]
PR6	PMA_RX_CLK pulse width during valid and invalid transitions	51.7.2	Clock pulse width is greater than or equal to the period times the duty cycle	XSBI:M	Yes [ ] N/A[ ]

51.10.4.4 PMA delay constraints

Item	Feature	Subclause	Value/Comment	Status	Support
PD1	Maximum delay for PMA/ PMD functions	51.3.3	Meets the requirements of the respective PMD clause	M	Yes [ ]

52. Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-S (short wavelength serial), 10GBASE-L (long wavelength serial), and 10GBASE-E (extra long wavelength serial)

52.1 Overview

This clause specifies the PMDs and baseband media, including both single and multimode optical fiber, for the following 10GBASE serial PHYs, shown in Table 52–1.

Table 52–1—10GBASE serial PHYs

Name	Description
10GBASE-SR	850 nm serial LAN PHY
10GBASE-LR	1310 nm serial LAN PHY
10GBASE-ER	1550 nm serial LAN PHY
10GBASE-SW	850 nm serial WAN PHY
10GBASE-LW	1310 nm serial WAN PHY
10GBASE-EW	1550 nm serial WAN PHY

In order to form a complete Physical Layer, each PMD is combined with the appropriate physical sublayers indicated in Table 52–2 and optionally with the management functions that may be accessible through the management interface defined in Clause 45.

Table 52–2—PMD type and associated clauses

Associated clause	10GBASE-R	10GBASE-W
46—RS	Required	Required
46—XGMII	Optional	Optional
47—XGXS and XAUI	Optional	Optional
49—Type R PCS	Required	Required
50—WIS	N/A	Required
51—Serial PMA	Required	Required

Figure 52–1 depicts the relationships of the serial PMD (shown shaded) with other sublayers and the ISO/IEC Open System Interconnection (OSI) reference model.

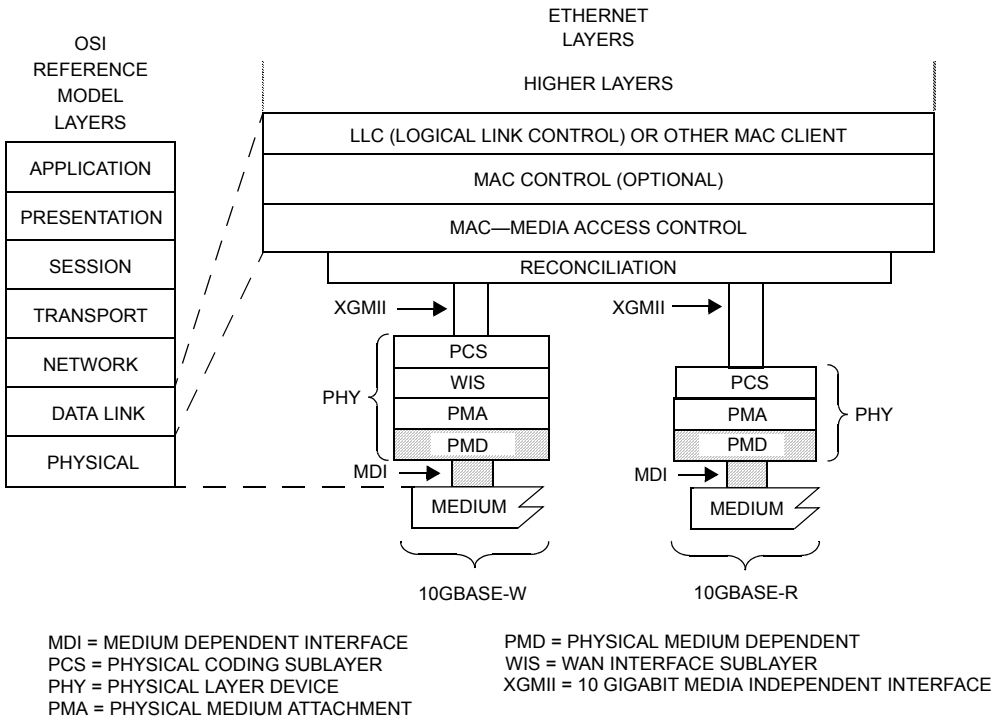

Figure 52–1—10GBASE-S, -L, and -E PMDs relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

52.1.1 Physical Medium Dependent (PMD) sublayer service interface

The following specifies the services provided by the 10GBASE-R and 10GBASE-W PMDs. These PMD sublayer service interfaces are described in an abstract manner and do not imply any particular implementation.

The PMD Service Interface supports the exchange of encoded and scrambled 64B/66B blocks between the PMA and PMD entities. In the case of the 10GBASE-W PMD types, these blocks are framed and scrambled by the WIS as described in Clause 50. The PMD translates the serialized data of the PMA to and from signals suitable for the specified medium.

The following primitives are defined:

PMD_UNITDATA.request

PMD_UNITDATA.indication

PMD_SIGNAL.indication

NOTE—Primitives are described in 1.2.2.

52.1.1.1 PMD_UNITDATA.request

This primitive defines the transfer of a serial data stream from the PMA to the PMD.

52.1.1.1.1 Semantics of the service primitive

PMD_UNITDATA.request(tx_bit)

The data conveyed by PMD_UNITDATA.request is a continuous stream of bits. The tx_bit parameter can take one of two values: ONE or ZERO.

52.1.1.1.2 When generated

The PMA continuously sends the appropriate stream of bits to the PMD for transmission on the medium, at a nominal 10.3125 GBd signaling speed for 10GBASE-R PMD types and 9.95328 GBd signaling speed for 10GBASE-W PMDs.

52.1.1.1.3 Effect of receipt

Upon receipt of this primitive, the PMD converts the specified stream of bits into the appropriate signals on the MDI.

52.1.1.2 PMD_UNITDATA.indication

This primitive defines the transfer of data (in the form of serialized data) from the PMD to the PMA.

52.1.1.2.1 Semantics of the service primitive

PMD_UNITDATA.indication(rx_bit)

The data conveyed by PMD_UNITDATA.indication is a continuous stream of bits. The rx_bit parameter can take one of two values: ONE or ZERO.

52.1.1.2.2 When generated

The PMD continuously sends stream of bits to the PMA corresponding to the signals received from the MDI.

52.1.1.2.3 Effect of receipt

The effect of receipt of this primitive by the client is unspecified by the PMD sublayer.

52.1.1.3 PMD_SIGNAL.indication

This primitive is generated by the PMD to indicate the status of the signal being received from the MDI.

52.1.1.3.1 Semantics of the service primitive

PMD_SIGNAL.indication(SIGNAL_DETECT)

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL, indicating whether the PMD is detecting light at the receiver (OK) or not (FAIL). When SIGNAL_DETECT = FAIL, PMD_UNITDATA.indication(rx_bit) is undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that PMD_UNITDATA.indication(rx_bit) is known good. It is possible for a poor quality link to provide sufficient light for a SIGNAL_DETECT = OK indication and still not meet the 10^{-12} BER objective.

52.1.1.3.2 When generated

The PMD generates this primitive to indicate a change in the value of SIGNAL_DETECT. If the MDIO interface is implemented, then PMD_global_signal_detect shall be continuously set to the value of SIGNAL_DETECT.

52.1.1.3.3 Effect of receipt

The effect of receipt of this primitive by the client is unspecified by the PMD sublayer.

52.2 Delay constraints

An upper bound to the delay through the PMA and PMD is required for predictable operation of the MAC Control PAUSE operation. The PMA and PMD shall incur a round-trip delay (transmit and receive) of not more than 512 bit-times, or 1 pause_quantum, including 2 m of fiber. A description of overall system delay constraints and the definitions for bit-times and pause_quanta can be found in 44.3.

52.3 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that provide control and status information for and about the PMD. If MDIO is implemented, it shall map MDIO control variables to PMD control variables as shown in Table 52–3, and MDIO status variables to PMD status variables as shown in Table 52–4.

Table 52–3—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/ bit number	PMD control variable
Reset	Control register 1	1.0.15	PMD_reset
Global Transmit Disable	Transmit disable register	1.9.0	PMD_global_transmit_disable

Table 52–4—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Fault	Status register 1	1.1.7	PMD_fault
Transmit fault	Status register 2	1.8.11	PMD_transmit_fault
Receive fault	Status register 2	1.8.10	PMD_receive_fault
Global PMD Receive signal detect	Receive signal detect register	1.10.0	PMD_global_signal_detect

52.4 PMD functional specifications

The 10GBASE-R and 10GBASE-W PMDs perform the Transmit and Receive functions that convey data between the PMD service interface and the MDI.

52.4.1 PMD block diagram

For purposes of system conformance, the PMD sublayer is standardized at test points TP2 and TP3 as shown in Figure 52–2. The optical transmit signal is defined at the output end of a patch cord (TP2), between 2 and 5 m in length, of a type consistent with the link type connected to the transmitter. Unless specified otherwise, all transmitter measurements and tests defined in 52.9 are made at TP2. The optical receive signal is defined at the output of the fiber optic cabling (TP3) connected to the receiver. Unless specified otherwise, all receiver measurements and tests defined in 52.9 are made at TP3.

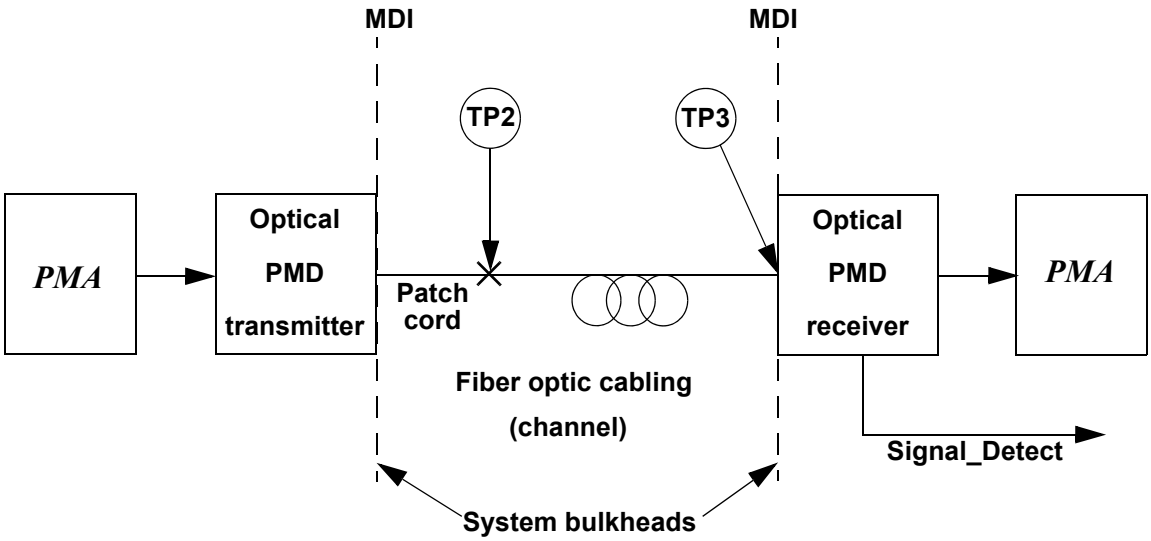

Figure 52–2—Block diagram

52.4.2 PMD Transmit function

The PMD Transmit function shall convey the bits requested by the PMD service interface message PMD_UNITDATA.request(tx_bit) to the MDI according to the optical specifications in this clause. The higher optical power level shall correspond to tx_bit = ONE.

52.4.3 PMD Receive function

The PMD Receive function shall convey the bits received from the MDI according to the optical specifications in this clause to the PMD service interface using the message PMD_UNITDATA.indication(rx_bit). The higher optical power level shall correspond to rx_bit = ONE.

52.4.4 PMD Signal Detect function

The PMD Signal Detect function shall report to the PMD service interface, using the message PMD_SIGNAL.indication(SIGNAL_DETECT) which is signaled continuously. PMD_SIGNAL.indication is intended to be an indicator of optical signal presence. If the MDIO interface is implemented, then

PMD_global_signal_detect (1.10.0) shall be continuously set to the value of SIGNAL_DETECT as described in 45.2.1.9.7.

The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 52–5. The PMD receiver is not required to verify whether a compliant 10GBASE-R or 10GBASE-W signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

As an unavoidable consequence of the requirements for the setting of the SIGNAL_DETECT parameter, implementations must provide adequate margin between the input optical power level at which the SIGNAL_DETECT parameter is set to OK, and the inherent noise level of the PMD due to cross talk, power supply noise, etc.

Table 52–5—SIGNAL_DETECT value definition

Receive conditions	Signal Detect value
Input_optical_power ≤ -30 dBm average power	FAIL
Input_optical_power ≥ Receiver sensitivity (max) in OMA in Table 52–9, Table 52–13, or Table 52–17 AND compliant 10GBASE-R or 10GBASE-W signal input	OK
All other conditions	Unspecified

Various implementations of the Signal Detect function are permitted by this standard, including implementations that generate the SIGNAL_DETECT parameter values in response to the amplitude of the modulation of the optical signal and implementations that respond to the average optical power of the modulated optical signal.

52.4.5 PMD_reset function

If the MDIO interface is implemented, and if PMD_reset is asserted, the PMD shall be reset as defined in 45.2.1.1.1.

52.4.6 PMD_fault function

If the MDIO is implemented, PMD_fault is the logical OR of PMD_receive_fault, PMD_transmit_fault and any other implementation specific fault.

52.4.7 PMD_global_transmit_disable function

The PMD_global_transmit_disable function is optional. When asserted, this function shall turn off the optical transmitter so that it meets the requirements of the average launch power of OFF Transmitter in Table 52–7, Table 52–12, or Table 52–16.

If a PMD_transmit_fault (optional) is detected, then the PMD_global_transmit_disable function should also be asserted.

If the MDIO interface is implemented, then this function shall map to the PMD_global_transmit_disable bit as specified in 45.2.1.8.7.

NOTE—PMD Transmit Disable 0 is not used for serial PMDs.

52.4.8 PMD_transmit_fault function

The PMD_transmit_fault function is optional. The faults detected by this function are implementation specific, but should not include the assertion of the PMD_global_transmit_disable function.

If a PMD_transmit_fault (optional) is detected, then the PMD_global_transmit_disable function should also be asserted.

If the MDIO interface is implemented, then this function shall be mapped to the PMD_transmit_fault bit as specified in 45.2.1.7.4.

52.4.9 PMD_receive_fault function

The PMD_receive_fault function is optional. PMD_receive_fault is the logical OR of NOT SIGNAL_DETECT and any implementation specific fault.

If the MDIO interface is implemented, then this function shall contribute to PMA/PMD receive fault bit as specified in 45.2.1.7.5.

52.5 PMD to MDI optical specifications for 10GBASE-S

The operating ranges for 10GBASE-S are defined in Table 52–6. A 10GBASE-S compliant PMD supports multimode fiber media types listed in Table 52–6 according to the specifications defined in 52.14. A PMD that exceeds the operational range requirement while meeting all other optical specifications is considered compliant (e.g., a 400 MHz km 50 μm solution operating at 80 m meets the operating range requirement of 2 m to 66 m).

Table 52–6—10GBASE-S operating range for each optical fiber type

Fiber type	Minimum modal bandwidth @ 850 nm (MHz•km)	Operating range (m)
62.5 μm MMF	160	2 to 26
	200	2 to 33
50 μm MMF	400	2 to 66
	500	2 to 82
	2000	2 to 300
	4700	2 to 400

52.5.1 10GBASE-S transmitter optical specifications

The 10GBASE-S transmitter shall meet the specifications defined in Table 52–7 per measurement techniques defined in 52.9. It shall also meet the center wavelength, maximum spectral width and minimum optical modulation amplitude as defined in Table 52–8.

Table 52–7—10GBASE-S transmit characteristics

Description	10GBASE-SW	10GBASE-SR	Unit
Signaling speed (nominal)	9.95328	10.3125	GBd
Signaling speed variation from nominal (max)	± 20	± 100	ppm
Center wavelength (range)	840 to 860		nm
RMS spectral width[a] (max)	See footnote[b]		
Average launch power (max)	See footnote[c]		
Average launch power[d] (min)	–7.3		dBm
Launch power (min) in OMA	See footnote[b]		
Average launch power of OFF transmitter[e] (max)	–30		dBm
Extinction ratio (min)	3		dB
$RIN_{12}OMA$ (max)	–128		dB/Hz
Optical Return Loss Tolerance (max)	12		dB
Encircled flux	See footnote[f]		
Transmitter eye mask definition[g] A {X1, X2, X3, Y1, Y2, Y3}	{0.25, 0.40, 0.45, 0.25, 0.28, 0.40}		
Transmitter eye mask definition[g] B {X1, X2, X3, Y1, Y2, Y3} Hit ratio 5×10^{-5} per sample	{0.235, 0.395, 0.45, 0.235, 0.265, 0.4}		
Transmitter and dispersion penalty[h] (max)	3.9 dB		dB

[a]RMS spectral width is the standard deviation of the spectrum.
[b]Trade-offs are available between spectral width, center wavelength and minimum optical modulation amplitude. See Figure 52–3 and Table 52–8.
[c]The 10GBASE-S launch power shall be the lesser of the Hazard Level 1 safety limit as defined by 52.10.2 or the average receive power (max) defined by Table 52–9.
[d]Average launch power (min) is informative and not the principal indicator of signal strength. A transmitter with launch power below this value cannot be compliant; however, a value above this does not ensure compliance.
[e]Examples of an OFF transmitter are: no power supplied to the PMD, laser shutdown for safety conditions, activation of a PMD_global_transmit_disable or other optional transmitter shut down conditions.
[f]The encircled flux at 19 µm shall be greater than or equal to 86% and the encircled flux at 4.5 µm shall be less than or equal to 30% when measured into Type A1a (50/125 µm multimode) fiber per ANSI/TIA/EIA-455-203-2001.
[g]Either transmitter eye mask definition A or B may be used. A transmitter is not required to comply with both definitions.
[h]TDP(max) and OMA(min) are at the respective wavelength and spectral width as specified in Table 52–8.

The tradeoffs between center wavelength, maximum RMS spectral width and minimum optical modulation amplitude are defined in Table 52–8 and are shown graphically in the informative Figure 52–3. In Table 52–8 the minimum optical modulation amplitude (dBm) allowed by this standard is given for the various ranges of center wavelength and spectral width.

Table 52–8—Minimum 10GBASE-S optical modulation amplitude (dBm) as a function of center wavelength and spectral width

Center wavelength (nm)	RMS Spectral width (nm)								
	Up to 0.05	0.05 to 0.1	0.1 to 0.15	0.15 to 0.2	0.2 to 0.25	0.25 to 0.3	0.3 to 0.35	0.35 to 0.4	0.4 to 0.45
840 to 842	–4.2	–4.2	–4.1	–4.1	–3.9	–3.8	–3.5	–3.2	–2.8
842 to 844	–4.2	–4.2	–4.2	–4.1	–3.9	–3.8	–3.6	–3.3	–2.9
844 to 846	–4.2	–4.2	–4.2	–4.1	–4.0	–3.8	–3.6	–3.3	–2.9
846 to 848	–4.3	–4.2	–4.2	–4.1	–4.0	–3.8	–3.6	–3.3	–2.9
848 to 850	–4.3	–4.2	–4.2	–4.1	–4.0	–3.8	–3.6	–3.3	–3.0
850 to 852	–4.3	–4.2	–4.2	–4.1	–4.0	–3.8	–3.6	–3.4	–3.0
852 to 854	–4.3	–4.2	–4.2	–4.1	–4.0	–3.9	–3.7	–3.4	–3.1
854 to 856	–4.3	–4.3	–4.2	–4.1	–4.0	–3.9	–3.7	–3.4	–3.1
856 to 858	–4.3	–4.3	–4.2	–4.1	–4.0	–3.9	–3.7	–3.5	–3.1
858 to 860	–4.3	–4.3	–4.2	–4.2	–4.1	–3.9	–3.7	–3.5	–3.2

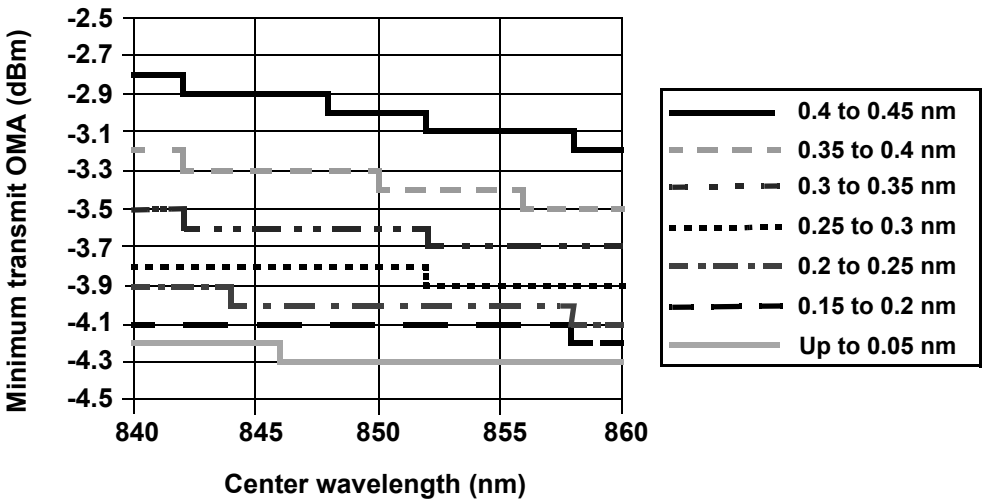

Figure 52–3—Triple tradeoff curve for 10GBASE-S (informative)

52.5.2 10GBASE-S receive optical specifications

The 10GBASE-S receiver shall meet the specifications defined in Table 52–9 per measurement techniques defined in 52.9.

Table 52–9—10GBASE-S receive characteristics

Description	10GBASE-SW	10GBASE-SR	Unit
Signaling speed (nominal)	9.95328	10.3125	GBd
Signaling speed variation from nominal (max)	± 20	± 100	ppm
Center wavelength (range)	840 to 860		nm
Average receive power[a] (max)	–1.0		dBm
Average receive power[b] (min)	–9.9		dBm
Receiver sensitivity (max) in OMA[c]	0.077 (–11.1)		mW (dBm)
Receiver Reflectance (max)	–12		dB
Stressed receiver sensitivity in OMA[d] [e](max)	0.18 (–7.5)		mW (dBm)
Vertical eye closure penalty[f] (min)	3.5		dB
Stressed eye jitter[g] (min)	0.3		UI pk-pk
Receive electrical 3 dB upper cutoff frequency (max)	12.3		GHz

[a]The receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having a power level equal to the average receive power (max) plus at least 1 dB.
[b]Average receive power (min) is informative and not the principal indicator of signal strength. A received power below this value cannot be compliant; however, a value above this does not ensure compliance.
[c]Receiver sensitivity is informative.
[d]Measured with conformance test signal at TP3 (see 52.9.9.3) for BER = 10^{-12}.
[e]The stressed sensitivity values in the table are for system level BER measurements which include the effects of CDR circuits. It is recommended that at least 0.4 dB additional margin be allocated if component level measurements are made without the effect of CDR circuits
[f]Vertical eye closure penalty is a test condition for measuring stressed receiver sensitivity. It is not a required characteristic of the receiver.
[g]Stressed eye jitter is a test condition for measuring stressed receiver sensitivity. It is not a required characteristic of the receiver.

52.5.3 10GBASE-S link power budgets (informative)

Example power budgets and penalties for a 10GBASE-S channel are shown in Table 52–10.

52.6 PMD to MDI optical specifications for 10GBASE-L

The operating range for 10GBASE-L PMDs is defined in Table 52–11. A 10GBASE-L compliant PMD supports Types B1.1 and B1.3 single-mode fibers according to the specifications defined in 52.14. A PMD that exceeds the operational range requirement while meeting all other optical specifications is considered compliant (e.g., operating at 12.5 km meets the minimum range requirement of 2 m to 10 km).

Table 52–10—10GBASE-S link power budgets[a,b]

Parameter	62.5 μm MMF		50 μm MMF				Unit
Modal bandwidth as measured at 850 nm	160	200	400	500	2000	4700	MHz•km
Power budget	7.3	7.3	7.3	7.3	7.3	7.3	dB
Operating distance	26	33	66	82	300	400	m
Channel insertion loss[c, d]	1.6	1.6	1.7	1.8	2.6	2.9	dB
Allocation for penalties	4.7	4.8	5.1	5.0	4.7	4.4	dB
Additional insertion loss allowed[e]	1.0	0.8	0.5	0.5	0.0	0.0	dB

[a]Budget numbers are rounded to nearest 0.1 dB.
[b]Link penalties are used for link budget calculations. They are not requirements and are not meant to be tested.
[c]Operating distances used to calculate the channel insertion loss are the maximum values specified in Table 52–6.
[d]The specifications for a wavelength of 840 nm and a spectral width of 0.29 nm in Table 52–8 is used to calculate channel insertion loss, allocation for penalties, and additional insertion loss allowed.
[e]This portion of the link budget is permitted to be used to overcome insertion loss higher than the "Channel insertion loss" value and in some cases may be less than the value shown.

Table 52–11—10GBASE-L operating range

PMD Type	Nominal Wavelength (nm)	Minimum range
10GBASE-L	1310	2 m to 10 km

52.6.1 10GBASE-L transmitter optical specifications

The 10GBASE-L transmitter shall meet the specifications defined in Table 52–12 per measurement techniques defined in 52.9.

Table 52–12—10GBASE-L transmit characteristics

Description	10GBASE-LW	10GBASE-LR	Unit
Signaling speed (nominal)	9.95328	10.3125	GBd
Signaling speed variation from nominal (max)	± 20	± 100	ppm
Center wavelength (range)	1260 to 1355		nm
Side Mode Suppression Ratio (min)	30		dB
Average launch power (max)	0.5		dBm
Average launch power[a] (min)	–8.2		dBm
Launch power (min) in OMA minus TDP[b]	–6.2		dBm
Optical Modulation Amplitude[c] (min)	–5.2		dBm
Transmitter and dispersion penalty (max)	3.2		dB
Average launch power of OFF transmitter[d] (max)	–30		dBm
Extinction ratio (min)	3.5		dB
$RIN_{12}OMA$ (max)	–128		dB/Hz
Optical Return Loss Tolerance (max)	12		dB
Transmitter Reflectance[e] (max)	–12		dB
Transmitter eye mask definition[f] A {X1, X2, X3, Y1, Y2, Y3}	{0.25, 0.40, 0.45, 0.25, 0.28, 0.40}		
Transmitter eye mask definition[f] B {X1, X2, X3, Y1, Y2, Y3} Hit ratio 5×10^{-5} per sample	{0.235, 0.395, 0.45, 0.235, 0.265, 0.4}		

[a]Average launch power (min) is informative and not the principal indicator of signal strength. A transmitter with launch power below this value cannot be compliant; however, a value above this does not ensure compliance.
[b]TDP is transmitter and dispersion penalty.
[c]Even if the TDP < 1 dB, the OMA(min) must exceed this value.
[d]Examples of an OFF transmitter are: no power supplied to the PMD, laser shutdown for safety conditions, activation of a PMD_global_transmit_disable or other optional transmitter shut down conditions.
[e]Transmitter reflectance is defined looking into the transmitter.
[f]Either transmitter eye mask definition A or B may be used. A transmitter is not required to comply with both definitions.

52.6.2 10GBASE-L receive optical specifications

The 10GBASE-L receiver shall meet the specifications defined in Table 52–13 per measurement techniques defined in 52.9.

Table 52–13—10GBASE-L receive characteristics

Description	10GBASE-LW	10BGASE-LR	Unit
Signaling speed (nominal)	9.95328	10.3125	GBd
Signaling speed variation from nominal (max)	± 20	± 100	ppm
Center wavelength (range)	1260 to 1355		nm
Average receive power[a] (max)	0.5		dBm
Average receive power[b] (min)	–14.4		dBm
Receiver sensitivity (max) in OMA[c]	0.055 (–12.6)		mW (dBm)
Receiver Reflectance (max)	–12		dB
Stressed receiver sensitivity (max) in OMA[d, e]	0.093 (–10.3)		mW (dBm)
Vertical eye closure penalty[f] (min)	2.2		dB
Stressed eye jitter[g] (min)	0.3		UI pk-pk
Receive electrical 3 dB upper cutoff frequency (max)	12.3		GHz

[a]The receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having a power level equal to the average receive power (max) plus at least 1 dB.
[b]Average receive power (min) is informative and not the principal indicator of signal strength. A received power below this value cannot be compliant; however, a value above this does not ensure compliance.
[c]Receiver sensitivity is informative.
[d]Measured with conformance test signal at TP3 (see 52.9.9.3) for BER = 10^{-12}.
[e]The stressed sensitivity values in the table are for system level BER measurements, which include the effects of CDR circuits. It is recommended that at least 0.4 dB additional margin be allocated if component level measurements are made without the effect of CDR circuits.
[f]Vertical eye closure penalty is a test condition for measuring stressed receiver sensitivity. It is not a required characteristic of the receiver.
[g]Stressed eye jitter is a test condition for measuring stressed receiver sensitivity. It is not a required characteristic of the receiver.

52.6.3 10GBASE-L link power budgets (informative)

An example power budget and penalties for a 10GBASE-L channel are shown in Table 52–14.

Table 52–14—10GBASE-L link power budget[a,b]

Parameter	10GBASE-L	Unit
Power budget	9.4	dB
Operating distance	10	km
Channel insertion loss[c, d]	6.2	dB
Allocation for penalties	3.2	dB
Additional insertion loss allowed	0.0	dB

[a]Budget numbers are rounded to nearest 0.1 dB.
[b]Link penalties are used for link budget calculations. They are not requirements and are not meant to be tested.
[c]The channel insertion loss is calculated using the maximum distance specified in Table 52–11 and cabled optical fiber attenuation of 0.4 dB/km at 1310 nm plus an allocation for connection and splice loss given in 52.14.2.1.
[d]A transmitter wavelength of 1260 nm with a TDP of 3 dB is used to calculate channel insertion loss, and allocation for penalties in this table.

52.7 PMD to MDI optical specifications for 10GBASE-E

The operating range for 10GBASE-E PMDs is defined in Table 52–15. A 10GBASE-E compliant PMD supports Types B1.1 and B1.3 single-mode fibers according to the specifications defined in 52.14. A PMD which exceeds the operational range requirement while meeting all other optical specifications is considered compliant (e.g., operating at 42.5 km meets the minimum range requirement of 2 m to 30 km).

Table 52–15—10GBASE-E operating range

PMD Type	Nominal Wavelength (nm)	Minimum Range
10GBASE-E	1550	2 m to 30 km
		2 m to 40 km[a]

[a]Links longer than 30 km for the same link power budget are considered engineered links. Attenuation for such links needs to be less than the minimum specified for B1.1 or B1.3 single-mode fiber.

Link attenuation requirements are specified in 52.14.3.

52.7.1 10GBASE-E transmitter optical specifications

The 10GBASE-E transmitter shall meet the specifications defined in Table 52–16 per measurement techniques defined in 52.9.

Table 52–16—10GBASE-E transmit characteristics

Description	10GBASE-EW	10GBASE-ER	Unit
Signaling speed (nominal)	9.95328	10.3125	GBd
Signaling speed variation from nominal (max)	± 20	± 100	ppm
Center wavelength (range)	1530 to 1565		nm
Side Mode Suppression Ratio (min)	30		dB
Average launch power (max)	4.0		dBm
Average launch power[a] (min)	–4.7		dBm
Launch power (min) in OMA minus TDP[b]	–2.1		dBm
Average launch power of OFF transmitter[c] (max)	–30		dBm
Optical Modulation Amplitude[d] (min)	–1.7		dBm
Transmitter and dispersion penalty (max)	3.0		dB
Extinction ratio (min)	3		dB
$RIN_{21}OMA$[e] (max)	–128		dB/Hz
Optical Return Loss Tolerance (max)	21		dB
Transmitter eye mask definition[f] A {X1, X2, X3, Y1, Y2, Y3}	{0.25, 0.40, 0.45, 0.25, 0.28, 0.40}		
Transmitter eye mask definition[f] B{X1, X2, X3, Y1, Y2, Y3} Hit ratio 5×10^{-5} per sample	{0.235, 0.395, 0.45, 0.235, 0.265, 0.4}		

[a]Average launch power (min) is informative and not the principal indicator of signal strength. A transmitter with launch power below this value cannot be compliant; however, a value above this does not ensure compliance.
[b]TDP is transmitter and dispersion penalty.
[c]Examples of an OFF transmitter are: no power supplied to the PMD, laser shutdown for safety conditions, activation of a PMD_global_transmit_disable or other optional transmitter shut-down conditions.
[d]Even if the TDP < 0.4 dB, the OMA(min) must exceed this value.
[e]RIN measurement is made with a return loss at 21 dB.
[f]Either transmitter eye mask definition A or B may be used. A transmitter is not required to comply with both definitions.

52.7.2 10GBASE-E receive optical specifications

The 10GBASE-E receiver shall meet the specifications defined in Table 52–17 per measurement techniques defined in 52.9.

Table 52–17—10GBASE-E receive characteristics

Description	10GBASE-EW	10GBASE-ER	Unit
Signaling speed (nominal)	9.95328	10.3125	GBd
Signaling speed variation from nominal (max)	± 20	± 100	ppm
Center wavelength (range)	1530 to 1565		nm
Average receive power (max)	–1.0		dBm
Average receive power[a] (min)	–15.8		dBm
Maximum receive power (for damage)	4.0		dBm
Receiver sensitivity (max) in OMA[b]	0.039 (–14.1)		mW (dBm)
Receiver Reflectance (max)	–26		dB
Stressed receiver sensitivity (max) in OMA[c,d]	0.074 (-11.3)		mW (dBm)
Vertical eye closure penalty[e] (min)	2.7		dB
Stressed eye jitter (min)[f]	0.3		UI pk-pk
Receive electrical 3 dB upper cutoff frequency (max)	12.3		GHz

[a]Average receive power (min) is informative and not the principal indicator of signal strength. A received power below this value cannot be compliant; however, a value above this does not ensure compliance.
[b]Receiver sensitivity is informative.
[c]Measured with conformance test signal at TP3 (see 52.9.9.3) for BER = 10^{-12}.
[d]The stressed sensitivity values in the table are for system level BER measurements which include the effects of CDR circuits. It is recommended that at least 0.4 dB additional margin be allocated if component level measurements are made without the effects of CDR circuits.
[e]Vertical eye closure penalty is a test condition for measuring stressed receiver sensitivity. It is not a required characteristic of the receiver.
[f]Stressed eye jitter is a test condition for measuring stressed receiver sensitivity. It is not a required characteristic of the receiver.

52.7.3 10GBASE-E link power budgets (informative)

Example power budgets and penalties for a 10GBASE-E channel are shown in Table 52–18.

52.8 Jitter specifications for 10GBASE-R and 10GBASE-W

Acceptable transmitted jitter is achieved by compliance with 52.9.7, transmitter optical waveform, and 52.9.10, transmitter and dispersion penalty. Tolerance to received jitter is achieved by compliance to 52.9.9, stressed receiver conformance, which includes sinusoidal jitter as in the following subclause. There is no jitter transfer specification.

Table 52–18—10GBASE-E link power budgets[a,b]

Parameter	10GBASE-E		Unit
Power budget	15.0		dB
Operating distance	30	40[c]	km
Channel insertion loss[d, e]	10.9	10.9	dB
Maximum Discrete Reflectance (max)	–26		dB
Allocation for penalties	3.6	4.1	dB
Additional insertion loss allowed	0.5	0.0	dB

[a]Budget numbers are rounded to nearest 0.1 dB.
[b]Link penalties are built into the transmitter specifications by testing the PMD with a maximum dispersion fiber.
[c]Links longer than 30 km are considered engineered links. Attenuation for such links needs to be less than that guaranteed by B1.1 or B1.3 single-mode fiber.
[d]Operating distances used to calculate the channel insertion loss are the maximum values specified in Table 52–15
[e]A wavelength of 1565 nm and 3dB transmitter and dispersion penalty (TDP) is used to calculate channel insertion loss, and allocation for penalties.

52.8.1 Sinusoidal jitter for receiver conformance test

The sinusoidal jitter is used to test receiver jitter tolerance. Sinusoidal jitter may vary over a magnitude range as required to accurately calibrate a stressed eye per 52.9.9. The range is limited by the constraints of Table 52–19.

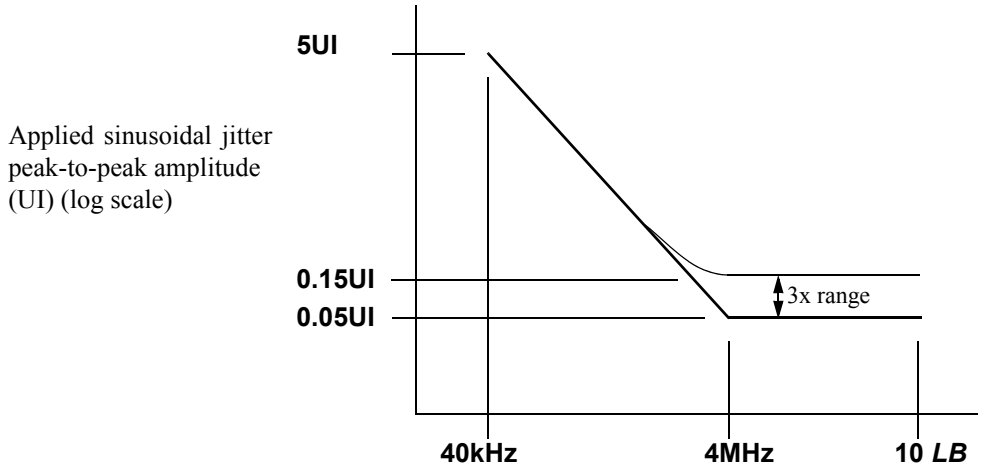

Figure 52–4—Mask of the sinusoidal component of jitter tolerance (informative)

Table 52–19—Applied sinusoidal jitter

Frequency range	Sinusoidal jitter ($UI_{pk\ to\ pk}$)
$f < 40$kHz	NA
40 kHz $< f \leq$ 4MHz	$2 \times 10^5 / f + S - 0.05$[a]
4 MHz $< f <$ 10 LB[b]	$0.05 \leq S \leq 0.15$ [a]

[a]S is the magnitude of sine jitter actually used in the calibration of the stressed eye per the methods of 52.9.9.3.
[b]LB = Loop Bandwidth; Upper frequency bound for added sine jitter should be at least 10 times the loop bandwidth of the receiver being tested.

52.9 Optical measurement requirements

All optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified.

52.9.1 Test patterns

Compliance is to be achieved in normal operation. Two types of test pattern are used, square wave (52.9.1.2) and other (52.9.1.1). Test patterns are as in Table 52–22 for 10GBASE-R and in 50.3.8 for 10GBASE-W unless specified otherwise.

NOTE—Test patterns for specific optical tests are designed to emulate system operation, which would entail passing valid 10GBASE-R or 10GBASE-W data.

52.9.1.1 Test-pattern definition

Patterns 1, 2, and 3 are defined in Table 52–21. Pattern 3 is optional.

For 10GBASE-R, two test pattern segments are specified, in two variants, “normal” (n) and “inverted” (i). Both are 8448 bits long. They may be generated dynamically by the 58 bit scrambler and “control block” sync header generation defined in 49.2.6, and using the scrambler starting seeds specified in Table 52–20 and the method of generation in 49.2.8. The segments are assembled into patterns, each containing four segments, as described in Table 52–21. Each may be held as a static pattern in test equipment, or generated or detected dynamically using the methods of 49.2.8 and 49.2.12 respectively, or otherwise.

Each segment contains a sync header transition every 66 bits.

Table 52–20—Pattern segments

Segments	Seed[57:0][a]
A_n	0x3C8B44DCAB6804F
B_n	0x34906BB85A38884

[a]The “invert” segments A_i and B_i are generated using the inverted seeds for A_n and B_n, respectively.

Table 52–21—Test patterns

Pattern	Pattern for 10GBASE-R	Pattern for 10GBASE-W
1	$B_n B_i B_n B_i$	Mixed[a]
2	$A_n A_i A_n A_i$	
3	PRBS31[b]	PRBS31[c]

[a]The patterns for 10GBASE-W are defined in 50.3.8.
[b]This is the test-pattern checker defined in 49.2.12.
[c]This is the test-pattern checker defined in 50.3.8.2.

The test pattern 1 is generated with the data input mode programmed to select LF data input. The test pattern 2 is generated with the data input mode programmed to select all zero data input.

Pattern 1 represents typical scrambled data while pattern 2 represents a less typical pattern that could happen by chance and is thought to be more demanding of the transmission process including the clock recovery subsystem. Both patterns are balanced over their length of 33 792 bits.

NOTE—While other test methods and patterns could be used it is the implementer's responsibility to ensure that measurements carried out with the specified patterns achieve the requirements specified.

Table 52–22—Test-pattern definitions and related subclauses

Test	Pattern	Related subclause
Average optical power	1 or 3	52.9.3
OMA (modulated optical power)	Square	52.9.5
Extinction ratio	1 or 3	52.9.4
Transmit eye	1 or 3	52.9.7
Receive upper cutoff frequency	1 or 3	52.9.11
$RIN_x OMA$	Square	52.9.6
Wavelength, spectral width	1 or 3	52.9.2
Side mode suppression ratio	1 or 3	—
Vertical eye closure penalty calibration	2 or 3	52.9.9
Receiver sensitivity	1 or 3	52.9.9
Receiver overload	1 or 3	—
Stressed receive conformance	2 or 3	52.9.9
Transmitter and dispersion penalty	2 or 3	52.9.10

52.9.1.2 Square wave pattern definition

A pattern consisting of four to eleven consecutive ones followed by an equal run of zeros may be used as a square wave. These patterns have fundamental frequencies between approximately 452 MHz (10GBASE-W) and 1289 MHz (10GBASE-R).

52.9.2 Center wavelength and spectral width measurements

The center wavelength and spectral width (RMS) shall be measured using an optical spectrum analyzer per TIA-455-127-A under modulated conditions using an appropriate PRBS or a valid 10GBASE-R or 10GBASE-W signal, OC-192 signal, STM-64 signal, or another representative test pattern.

52.9.3 Average optical power measurements

Average optical power shall be measured using the methods specified in TIA/EIA-455-95. This measurement may be made with the node transmitting test pattern 1 or 3 or a valid 10GBASE-R or 10GBASE-W signal, OC-192 signal, STM-64 signal, or another representative test pattern.

52.9.4 Extinction ratio measurements

Extinction ratio shall be measured using the methods specified in IEC 61280-2-2. This parameter is to be assured during system operation. However, measurements with an appropriate PRBS ($2^{23} – 1$ or $2^{31} – 1$) or a valid 10GBASE-R or 10GBASE-W or OC-192c or STM-64 signal will give equivalent results.

52.9.5 Optical modulation amplitude (OMA) test procedure

OMA is the difference in optical power for the nominal “1” and “0” levels of the optical signal, OMA in Figure 52–6, using waveform averaging or histogram means. OMA should be measured for a node transmitting the square wave pattern defined in 52.9.1.

The recommended technique for measuring optical modulation amplitude is illustrated in Figure 52–5. Optionally, a fourth-order Bessel-Thomson filter as specified in 52.9.7 can be used after the O/E converter. The measurement system consisting of the O/E converter, the optional filter and the oscilloscope has the following requirements:

a) The bandwidth of the measurement system shall be at least 3/T where T is the time at high or low (00001111 giving approximately T = 400 ps and 7.5 GHz as an example); and
b) The measurement system is calibrated at the appropriate wavelength for the transmitter under test.

With the device under test transmitting the square wave described above, use the following procedure to measure optical modulation amplitude:

c) Configure the test equipment as illustrated in Figure 52–5.
d) Measure the mean optical power P_1 of the logic “1” as defined over the center 20% of the time interval where the signal is in the high state. (See Figure 52–6.)
e) Measure the mean optical power P_0 of the logic “0” as defined over the center 20% of the time interval where the signal is in the low state. (See Figure 52–6.)
f) $OMA = P_1 – P_0$.

A method of approximating OMA is shown in Figure 52–11.

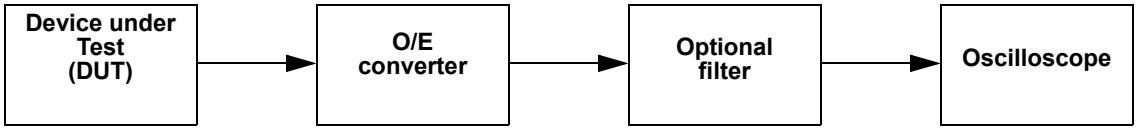

Figure 52–5—Recommended test equipment for measurement of optical modulation amplitude

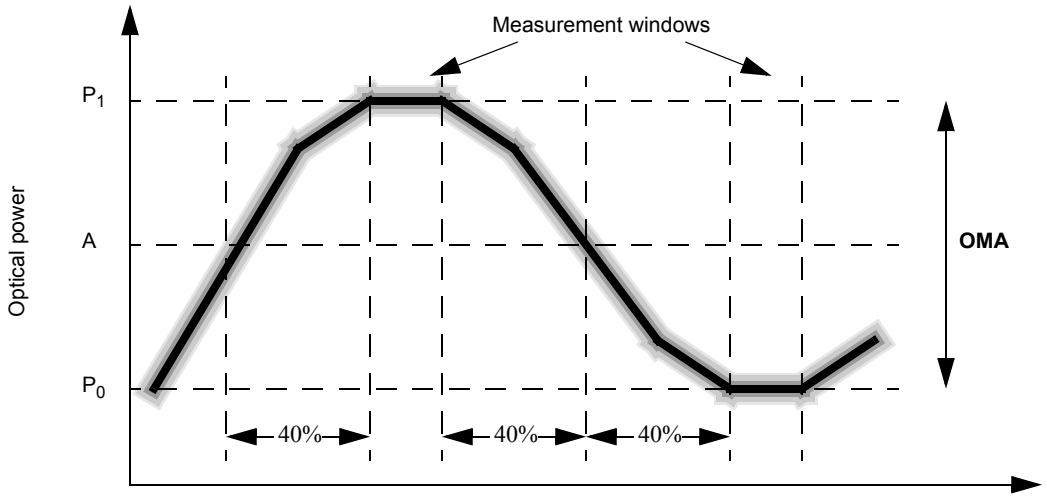

Figure 52–6—Optical modulation amplitude waveform measurement

52.9.6 Relative intensity noise optical modulation amplitude (RIN_xOMA) measuring procedure

This procedure describes a component test that may not be appropriate for a system level test depending on the implementation. If used, the procedure shall be performed as described in 52.9.6.1, 52.9.6.2, and 52.9.6.3.

52.9.6.1 General test description

The test arrangement is shown in Figure 52–7. The optical path between the device under test (DUT) and the detector has a single discrete reflection with the specified optical return loss as seen by the DUT.

Both the OMA power and noise power are measured by AC-coupling the O/E converter into the high-frequency electrical power meter. If needed, an amplifier may be used to boost the signal to the power meter. A low-pass filter is used between the photodetector and the power meter to limit the noise measured to the passband appropriate to the data rate of interest. In order to measure the noise, the modulation to the DUT is turned off.

52.9.6.2 Component descriptions

Optical path: The optical path and detector combination must be configured for a single dominant reflection with an optical return loss specified in Table 52–7 for 10GBASE-S, Table 52–12 for 10GBASE-L, and Table 52–16 for 10GBASE-E. (The optical return loss may be determined by the method of FOTP-107.) The length of the test cable is not critical but should be in excess of 2 m.

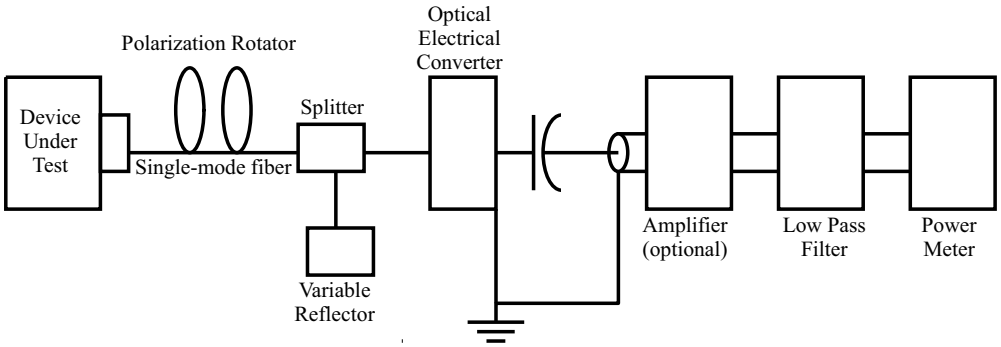

Figure 52–7—RIN$_x$OMA measurement setup

Polarization rotator: The polarization rotator is required to be capable of transforming an arbitrary orientation elliptically polarized wave into a fixed orientation linearly polarized wave.

O/E converter (and amplifier): If necessary, the noise may be amplified to a level consistent with accurate measurement by the power meter.

Filter: The upper –3 dB limit of the measurement apparatus is to be approximately equal to the bit rate (i.e., 10 GHz). The total filter bandwidth used in the RIN calculation is required to take the low-frequency cutoff of the DC-blocking capacitor into consideration. The low-frequency cutoff is recommended to be less than 1 MHz.

The filter should be placed in the circuit as the last component before the power meter so that any high-frequency noise components generated by the detector/amplifier are eliminated. If the power meter used has a very wide bandwidth, care should be taken to ensure that the filter does not lose its rejection at extremely high frequencies.

Power meter: The RMS electrical power meter should be capable of being zeroed in the absence of input optical power to remove any residual noise.

52.9.6.3 Test Procedure

Use the following procedure to test relative intensity noise optical modulation amplitude:

a) With the DUT disconnected, zero the power meter;
b) Connect the DUT, turn on the laser, and ensure that the laser is not modulated;
c) Operate the polarization rotator while observing the power meter output to maximize the noise read by the power meter. Note the maximum power, P_N;
d) Turn on the modulation to the laser using the square wave pattern of 52.9.1 and note the power measurement, P_M. It may be necessary to change or remove the effective reflection to obtain an accurate reading;
e) Calculate RIN from the observed electrical signal power and noise power by use of Equation (52–1).

$$RIN_xOMA = 10\log\frac{P_N}{BW \times P_M} \text{ (dB/Hz)} \quad (52–1)$$

where:

RIN_xOMA= Relative Intensity Noise referred to optical modulation amplitude measured with x dB reflection,

P_N = Electrical noise power in watts with modulation off,

P_M = Electrical power in watts with modulation on,

BW = Low-pass bandwidth of filter – high-pass bandwidth of DC-blocking capacitor [noise bandwidth of the measuring system (Hz)].

For testing multimode components or systems, the polarization rotator is removed from the setup and the single-mode fiber replaced with a multimode fiber. Step c) of the test procedure is eliminated.

52.9.7 Transmitter optical waveform

The required transmitter pulse shape characteristics are specified in the form of a mask of the transmitter eye diagram as shown in Figure 52–8. Compliance is to be assured during system operation. However, measurements with pattern 3 or 1 defined in 52.9.1, or other patterns such as a 2^{23}– 1 PRBS or a valid 10GBASE-R, or 10GBASE-W, or OC-192c, or STM-64 signal are likely to give very similar results. Measurements should be made as defined by IEC 61280-2-2 with the eye mask definition A coordinates, or as defined by 86.8.3.2 with the eye mask definition B coordinates. The two sets of coordinates (A and B) are given in Table 52–7, Table 52–12 or Table 52–16 as appropriate. A transmitter is not required to comply with both definitions. Measurement at 10.3125 GBd shall qualify for 10GBASE-W and 10GBASE-R use, measurement at 9.95328 GBd shall qualify for 10GBASE-W use only.

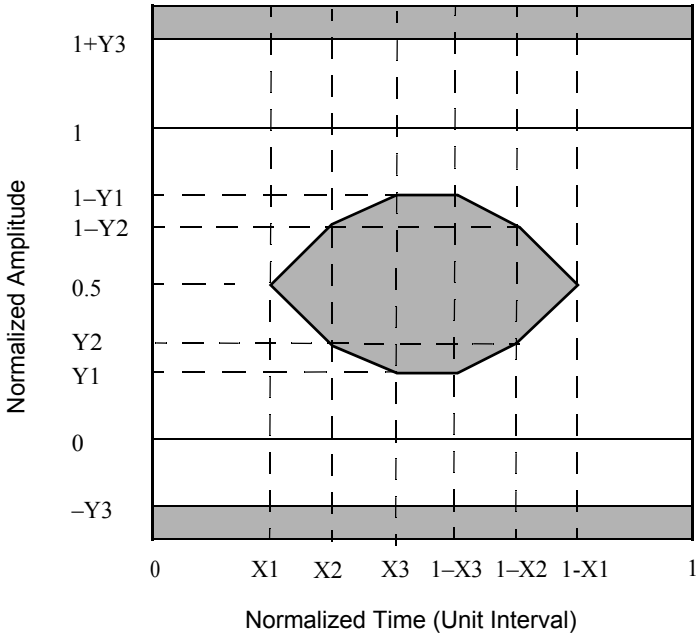

Figure 52–8—Transmitter eye mask definition

Normalized amplitudes of 0.0 and 1.0 represent the amplitudes of logic ZERO and ONE respectively. Normalized times of 0 and 1 on the unit interval scale are to be determined by the eye crossing means measured at the average value of the optical eye pattern.

A clock recovery unit (CRU) should be used to trigger the scope for mask measurements as shown in Figure 52–9. It should have a high-frequency corner bandwidth of less than or equal to 4 MHz and a slope of –20 dB/decade.

The eye shall be measured with respect to the mask of the eye using a receiver with a fourth-order Bessel-Thomson response having a transfer function given by Equation (52–2) and Equation (52–3):

$$H(y) = \frac{105}{105 + 105y + 45y^2 + 10y^3 + y^4} \tag{52–2}$$

where:

$$y = 2.114p\ ; \quad p = \frac{j\omega}{\omega_r}\ ; \quad \omega_r = 2\pi f_r\ ; \quad f_r = 7.5\text{GHz} \tag{52–3}$$

and where the filter response vs. frequency range for this fourth-order Bessel-Thomson receiver is defined in ITU-T G.691, along with the allowed tolerances (STM-64 values) for its physical implementation.

This Bessel-Thomson receiver is not intended to represent the noise filter used within a compliant optical receiver, but is intended to provide uniform measurement conditions at the transmitter.

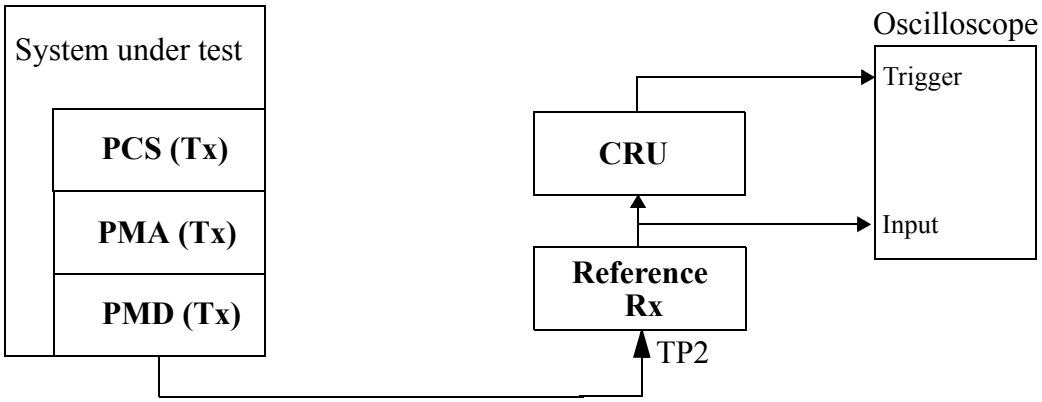

Figure 52–9—Transmitter optical waveform test block diagram

The fourth-order Bessel-Thomson filter is reactive. In order to suppress reflections, a 6 dB attenuator may be required at the filter input and/or output.

52.9.8 Receiver sensitivity measurements

Receiver sensitivity, which is defined for an ideal input signal, is informative and testing is not required. If measured, the test signal should have negligible impairments such as intersymbol interference (ISI), rise/fall times, jitter and RIN. Instead, the normative requirement for receivers is stressed receiver sensitivity. Stressed sensitivity is measured with a conditioned input signal where both vertical eye closure and jitter have been added according to 52.9.9.

52.9.9 Stressed receiver conformance test

Stressed receiver tolerance testing shall be performed in accordance with the requirements of 52.9.9.1, 52.9.9.2, 52.9.9.3, and 52.9.9.4.

Receivers must operate with BER less than 10^{-12} when tested with a conditioned input signal that combines vertical eye closure and jitter according to this clause.

The measurements in this subclause are performed with asynchronous data flowing out of the optical transmitter of the system under test. The output data pattern from the transmitter of the system under test is to be the same pattern defined for this measurement in 52.9.1.

52.9.9.1 Stressed receiver conformance test block diagram

A block diagram for the receiver conformance test is shown in Figure 52–10. A suitable pattern generator is used to continuously generate a signal or test pattern according to 52.9.1. The optical test signal is conditioned (stressed) using the methodology, as defined in 52.9.9.3, while applying sinusoidal jitter, as defined in 52.8.1. The receiver of the system under test is tested for conformance by enabling the error counter on the receiving side. As defined in section 49.2.12 and 50.3.8, the PCS or WIS is capable of detecting the data pattern and reporting any errors received. The optical power penalty for the stressed eye is intended to be similar to its vertical eye closure penalty. This is not necessarily the same as the highest TDP anticipated in service, but represents a standardized test condition for the receiver.

A suitable test set is needed to characterize and verify that the signal used to test the receiver has the appropriate characteristics. The test fiber called out for LW/LR and EW/ER and the transversal filter called out for SW/SR are not needed to characterize the receiver input signal; nor are they used during testing.

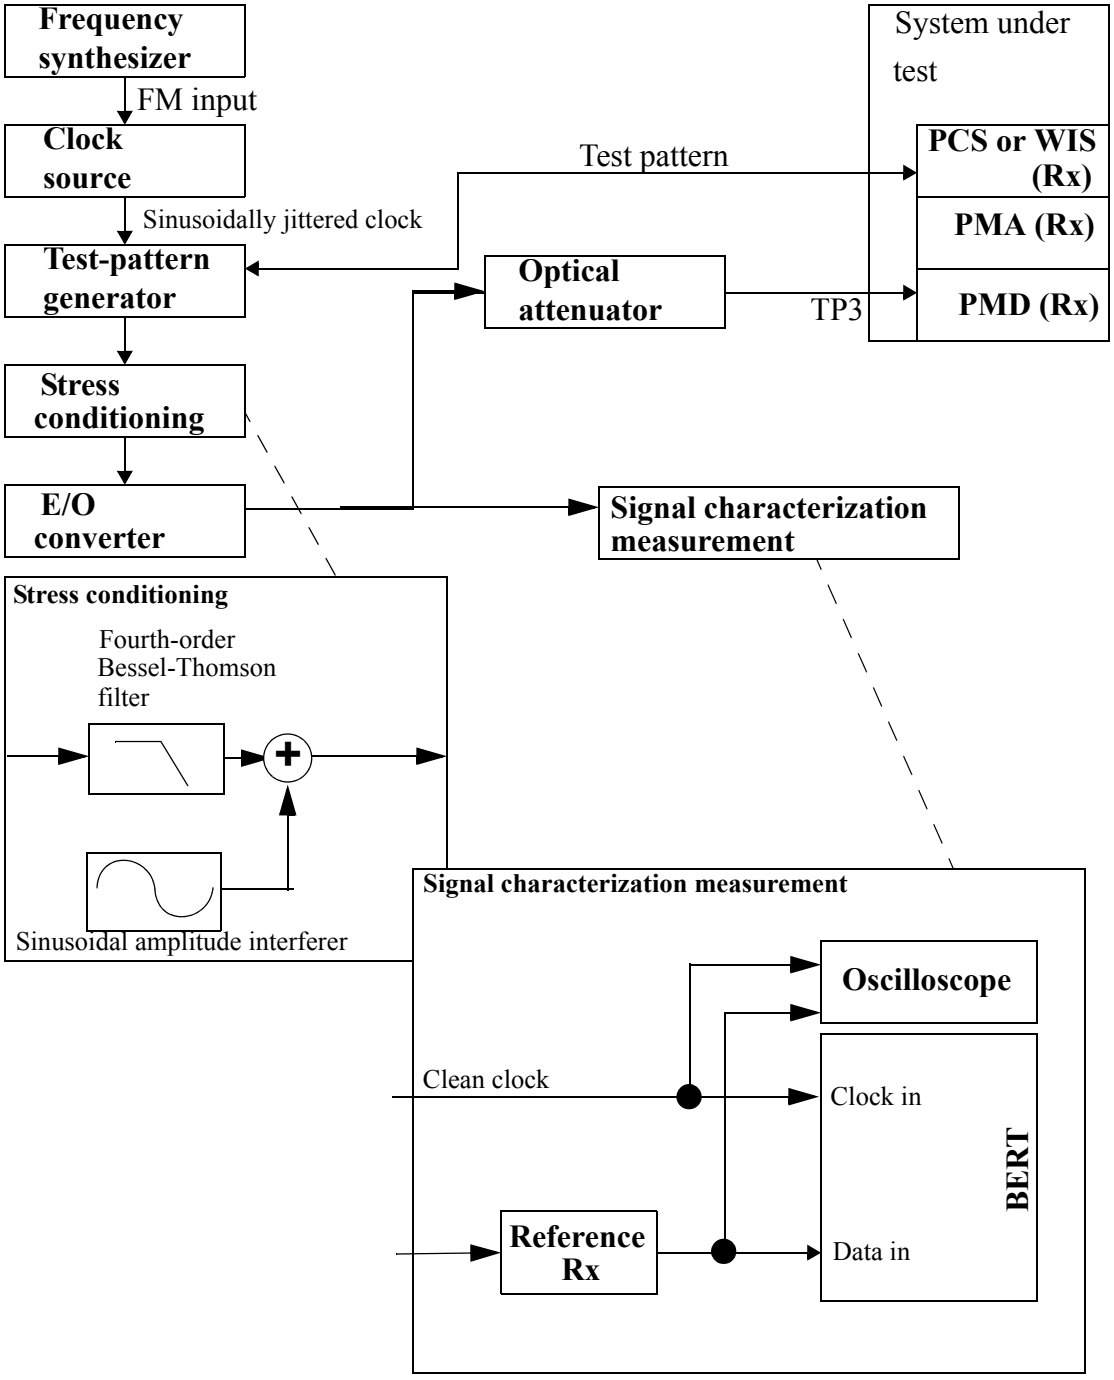

Figure 52–10—Stressed receiver conformance test block diagram

The fourth-order Bessel-Thomson filter is used to create ISI-induced vertical eye closure penalty (VECP). The sinusoidal amplitude interferer causes additional eye closure, but in conjunction with the slowed edge rates from the filter, also causes jitter. The nature of the jitter is intended to emulate instantaneous bit shrinkage that can occur with DDJ. This type of jitter cannot be created by simple phase modulation. The sinusoidal phase modulation represents other forms of jitter and also verifies that the receiver under test can

track low-frequency jitter. The frequency of the sinusoidal interference may be set at any frequency between 100 MHz and 2 GHz, although be careful to avoid a harmonic relationship between the sinusoidal interference, the sinusoidal jitter, the data rate and the pattern repetition rate.

For improved visibility for calibration, it is imperative that the Bessel-Thomson filter and all other elements in the signal path (cables, DC blocks, E/O converter, etc.) have wide and smooth frequency response and linear phase response throughout the spectrum of interest. Baseline wander and overshoot and undershoot should be minimized. If this is achieved, then data dependent effects should be minimal, and short data patterns can be used for calibration with the benefit of providing much improved trace visibility on sampling oscilloscopes. Actual patterns for testing the receiver shall be as specified in Table 52–22.

To further improve visibility for calibration, random noise effects, such as RIN and random clock jitter, should also be minimized. A small amount of residual noise and jitter from all sources is unavoidable, but should be less than 0.25 UI peak-peak of jitter at the 10^{-12} points.

The Bessel-Thomson filter and the E/O converter should have the appropriate frequency response to result in the appropriate level of initial ISI eye closure before the sinusoidal terms are added. The E/O converter should be fast and linear such that the waveshape and edge rates are predominantly controlled or limited by the electrical circuitry. Electrical summing requires high linearity of all elements including the E/O converter. Summing with an optical coupler after the E/O converter is an option that eases linearity requirements, but requires a second source for the interfering signal, will complicate settings of extinction ratio, and will add more RIN.

The vertical and horizontal eye closures to be used for receiver conformance testing are verified using an optical reference receiver with a 7.5 GHz fourth order ideal Bessel-Thomson response. Use of G.691 tolerance filters may significantly degrade this calibration. Care should be taken to ensure that all the light from the fiber is collected by the fast photodetector and that there is negligible mode selective loss, especially in the optical attenuator and the optical coupler, if used.

The clock output from the clock source in Figure 52–10 will be modulated with the sinusoidal jitter. To use an oscilloscope to calibrate the final stressed eye jitter that includes the sinusoidal jitter component, a separate clock source (clean clock of Figure 52–11) is required that is synchronized to the source clock, but not modulated with the jitter source.

52.9.9.2 Parameter definitions

The primary parameters of the conformance test signal are vertical eye closure penalty (VECP) and stressed eye jitter, J. Vertical eye closure penalty is measured at the time center of the eye (halfway between 0 and 1 on the unit interval scale as defined in 52.9.7) and is calculated relative to the measured OMA value. J is defined from the 0.5th to the 99.5th percentile of the jitter histogram. J is measured at the average optical power, which can be obtained with AC-coupling. The values of these components are defined by their histogram results.

The vertical eye closure penalty is given by Equation (52–4).

$$\text{Vertical eye closure penalty [dB, optical]} = 10 \times \log\frac{OMA}{A_O} \tag{52–4}$$

where:

A_O is the amplitude of the eye opening from the 99.95th percentile of the lower histogram to the 0.05th percentile of the upper histogram, and

OMA is the normal amplitude without ISI, as shown in Figure 52–11.

52.9.9.3 Stressed receiver conformance test signal characteristics and calibration

The conformance test signal is used to validate that the PMD receiver meets BER requirements with near worst case waveforms including pulse width shrinkage, power, simulated channel penalties, and a swept frequency sinusoidal jitter contribution applied at TP3.

Signal characteristics are described below along with a suggested approach for calibration.

The test signal includes vertical eye closure and high-probability jitter components. For this test, these two components are defined by peak values that include all but 0.1% for VECP and all but 1% for jitter of their histograms. Histograms should include at least 10 000 hits, and should be about 1%-width in the direction not being measured. Residual low-probability noise and jitter should be minimized—that is, the outer slopes of the final histograms should be as steep as possible down to very low probabilities.

The following steps describe a suggested method for calibrating a stressed eye generator:

1) Set the signaling speed of the test-pattern generator to satisfy the requirements of 52.5.2, 52.6.2, or 52.7.2.
2) Turn on the calibration pattern. A short pattern may be used for calibration if the conditions described in 52.9.9.1 are met, but this increases the risk that the longer test pattern used during testing will overstress the device under test. In any case, a pattern shorter than PRBS10 is not recommended.
3) Set the extinction ratio to approximately the Extinction Ratio (min) value given in 52.5.1, 52.6.1, or 52.7.1. Sinusoidal interference and jitter signals should be turned off at this point. If optical summing is used, ER may need to be adjusted after the sinusoidal interference signal is added below.
4) Measure the OMA of the test signal (without attenuation). OMA is measured per the method in 52.9.5 using the square wave pattern.
5) The requirements for vertical eye closure and jitter of the stressed eye test signal are given by the Vertical eye closure penalty (VECP) and Stressed eye jitter (J) values given in Table 52–9 for 10GBASE-S, Table 52–13 for 10GBASE-L, or Table 52–17 for 10GBASE-E.

   There are three components involved in calibration for vertical closure and J. These are a linear phase filter, sinusoidal interference, and sinusoidal jitter.

   Without sinusoidal jitter or sinusoidal interference, greater than two thirds of the vertical eye closure penalty value should be created by use of a linear phase, low jitter filter (such as Bessel-Thomson). The filter should be tested with the prescribed test patterns to verify that residual jitter and baseline wander are small, not to exceed 0.25 UI peak-peak. If not, the stress may be more than desired, leading to conservative results. However, compensation is not allowed. Once done, revert to the calibration pattern, if different than the final test pattern specified in Table 52–22.

   Any remaining vertical eye closure required must be created with sinusoidal interference or sinusoidal jitter.

   Sinusoidal jitter (phase modulation) must be added per the template of Table 52–19. For calibration purposes, sinusoidal jitter frequencies must be well within the flat portion of the template greater than 4 MHz.

   Iterate the settings for sinusoidal interference and/or jitter until all constraints are met, including jitter (J), vertical closure (VECP), and that sinusoidal jitter above 4 MHz is as specified in Table 52–19.

Verify that the optical power penalty for the stressed eye (relative to the reference transmitter per 52.9.10) is greater than or equal to VECP.

To emulate the effects of DCD or data-dependent jitter, at least 5 ps peak-peak of pulse shrinkage jitter should have been achieved.

6) Decrease the amplitude with the optical attenuator until the OMA complies with the OMA values given in Table 52–9 for 10GBASE-S, Table 52–13 for 10GBASE-L, Table 52–17 for 10GBASE-E. If high linearity exists, then the sinusoidal interference should not change the OMA value. OMA can be approximated with histograms as suggested in Figure 52–11. However, the normative definition for OMA is as given in 52.9.5.
7) For testing, turn on the actual required test pattern(s) per 52.9.1.

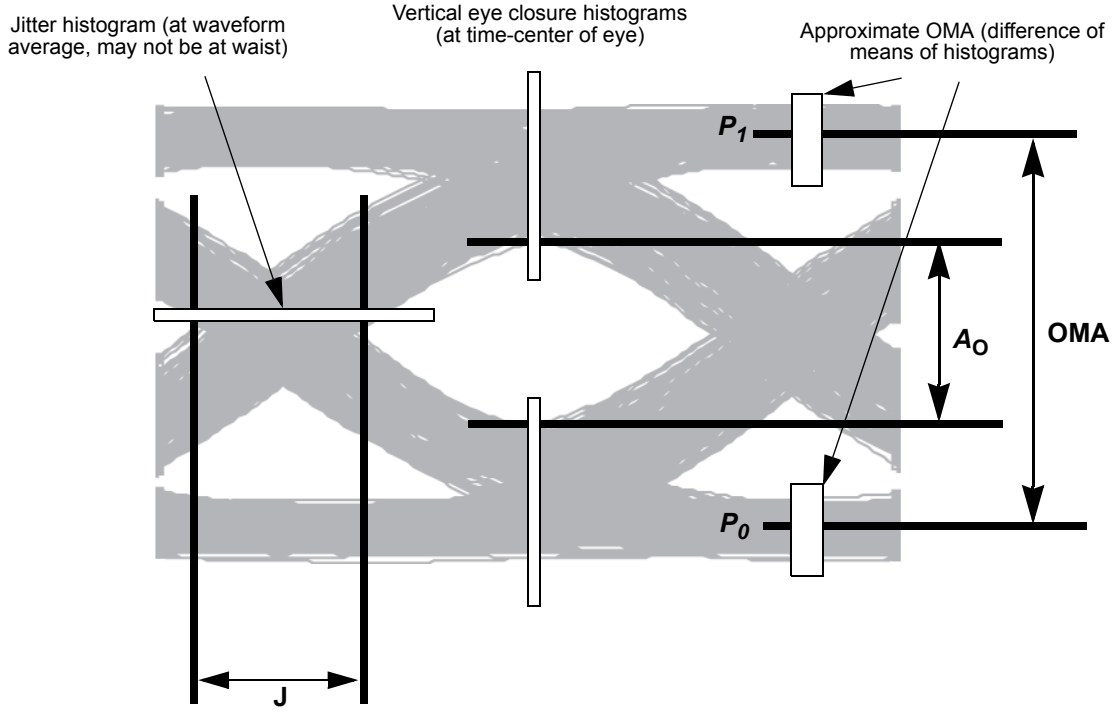

Figure 52–11—Required characteristics of the conformance test signal at TP3

Care should be taken when characterizing the signal used to make receiver tolerance measurements. In the case of a transmit jitter measurement, excessive and/or uncalibrated noise/jitter in the test system makes it more difficult to meet the specification and may have a negative impact on yield but will not effect interoperability. Running the receiver tolerance test with a signal that is under-stressed may result in the deployment of non-compliant receivers. Care should be taken to minimize and/or correct for the noise/jitter introduced by the reference receiver, filters, oscilloscope, and BERT. While the details of measurement and test equipment are beyond the scope of this standard it is recommended that the implementer fully characterize the test equipment and apply appropriate guard bands to ensure that the receive input signal meets the specified requirements.

52.9.9.4 Stressed receiver conformance test procedure

The test apparatus is set up as described in 52.9.9.1, 52.9.9.2 and 52.9.9.3. The sinusoidal jitter is then stepped across the frequency and amplitude range specified in 52.8.1 while monitoring BER at the receiver. The BER is to be compliant for all jitter frequencies. This method does not result in values for jitter

contributed by the receiver. It does, however, guarantee that a receiver meeting the requirements of this test will operate with the worst-case optical input.

52.9.10 Transmitter and dispersion penalty measurement

The transmitter and dispersion penalty (TDP) measurement tests for transmitter impairments with chromatic effects for a transmitter to be used with single-mode fiber, and for transmitter impairments with modal (not chromatic) dispersion effects for a transmitter to be used with multimode fiber.

The setup for measurement of transmitter and dispersion penalty is shown in Figure 52–12 and consists of a reference transmitter, the transmitter under test, a controlled optical reflection, an optical attenuator, a test fiber, a reference receiver, a transversal filter for 10GBASE-S, and a bit-error ratio tester. For 10GBASE-S, the polarization rotator shown in Figure 52–12 is removed from the setup and the single-mode fiber replaced with a multimode fiber. All BER and sensitivity measurements are made with the test patterns in 52.9.1.

52.9.10.1 Reference transmitter requirements

The reference transmitter is a high-quality instrument-grade device, which can be implemented by a CW laser modulated by a high-performance modulator. It should meet the following basic requirements:

a) The rise/fall times should be less than 30 ps at 20% to 80%.
b) The output optical eye is symmetric and passes the eye mask test of 52.9.7.
c) In the center 20% region of the eye, the worst case vertical eye closure penalty as defined in 52.9.9.3 is less than 0.5 dB.
d) Jitter less than 0.20 UI peak-peak.
e) RIN should be minimized to less than –136 dB/Hz

52.9.10.2 Channel requirements

The transmitter is tested using an optical channel that meets the requirements listed in Table 52–23.

Table 52–23—Transmitter compliance channel specifications

PMD Type	Dispersion[a] (ps/nm)		Insertion loss[b]	Optical return loss[c] (max)
	Minimum	Maximum		
10GBASE-S	N/A	N/A	Minimum	See ORLT in Table 52–7
10GBASE-L	$0.2325 \cdot \lambda \cdot [1 - (1324 / \lambda)^4]$	$0.2325 \cdot \lambda \cdot [1 – (1300 / \lambda)^4]$	Minimum	21 dB
10GBASE-E	0 (maximum)	$0.93 \cdot \lambda \cdot [1 – (1300 / \lambda)^4]$	Minimum	See ORLT in Table 52–16

[a]The dispersion is measured for the wavelength of the device under test. The coefficient assumes 10 km for 10GBASE-L and 40 km for 10GBASE-E.
[b]There is no intent to stress the sensitivity of the BERT's optical receiver.
[c]The optical return loss is applied at TP2.

A 10GBASE-L/E transmitter is to be compliant with a total dispersion at least as negative as the "minimum dispersion" and at least as positive as the "maximum dispersion" columns specified in Table 52–23 for the

wavelength of the device under test. This may be achieved with channels consisting of fibers with lengths chosen to meet the dispersion requirements.

To verify that the fiber has the correct amount of dispersion, the measurement method defined in ANSI/TIA/EIA-455-175A-92 may be used. The measurement is made in the linear power regime of the fiber.

The channel for 10GBASE-S is a 2 m to 5 m patch cord meeting the requirements in Table 52–25.

The channel provides a maximum optical return loss specified in Table 52–23. The state of polarization of the back reflection is adjusted to create the greatest RIN. The methods of 52.9.6.2 and 52.9.6.3 may be used.

52.9.10.3 Reference receiver requirements

The reference receiver should have the bandwidth given in 52.9.7. The sensitivity of the reference receiver should be limited by Gaussian noise. The receiver should have minimal threshold offset, deadband, hysteresis, baseline wander, deterministic jitter or other distortions. Decision sampling should be instantaneous with minimal uncertainty and setup/hold properties.

The nominal sensitivity of the reference receiver, S, is measured in OMA using the set up of Figure 52–12 without the test fiber and with the transversal filter removed. The sensitivity S must be corrected for any significant reference transmitter impairments including any vertical eye closure. It should be measured while sampling at the eye center or corrected for off-center sampling. It is calibrated at the wavelength of the transmitter under test.

For all transmitter and dispersion penalty measurements, determination of the center of the eye is required. Center of the eye is defined as the time halfway between the left and right sampling points within the eye where the measured BER is greater than or equal to 1×10^{-3}.

For 10GBASE-S the reference receiver is followed by a transversal filter with two equal amplitude paths with a differential delay of 55 ps. For 10GBASE-L/E the transversal filter is not used.

The clock recovery unit (CRU) used in the TDP measurement has a corner frequency of less than or equal to 4 MHz and a slope of 20 dB/decade. When using a clock recovery unit as a clock for BER measurements, passing of low-frequency jitter from the data to the clock removes this low-frequency jitter from the measurement.

52.9.10.4 Test procedure

To measure the transmitter and dispersion penalty (TDP) the following procedure shall be used. The decision threshold amplitude is defined to occur at the average signal level. The sampling instant is displaced from the eye center by $\pm$ 0.05 UI. The following procedure is repeated for early and late decision and the larger TDP value is used:

a) Configure the test equipment as described above and illustrated in Figure 52–12.
b) Adjust the attenuation of the optical attenuator to obtain a BER of 1×10^{-12}.
c) Record the optical power in OMA at the input to the reference receiver, P_DUT, in dBm.
d) If P_DUT is larger than S, the transmitter and dispersion penalty (TDP) for the transmitter under test is the difference between P_DUT and S, TDP = P_DUT – S. Otherwise the transmitter and dispersion penalty is zero, TDP = 0.

It is to be ensured that the measurements are made in the linear power regime of the fiber.

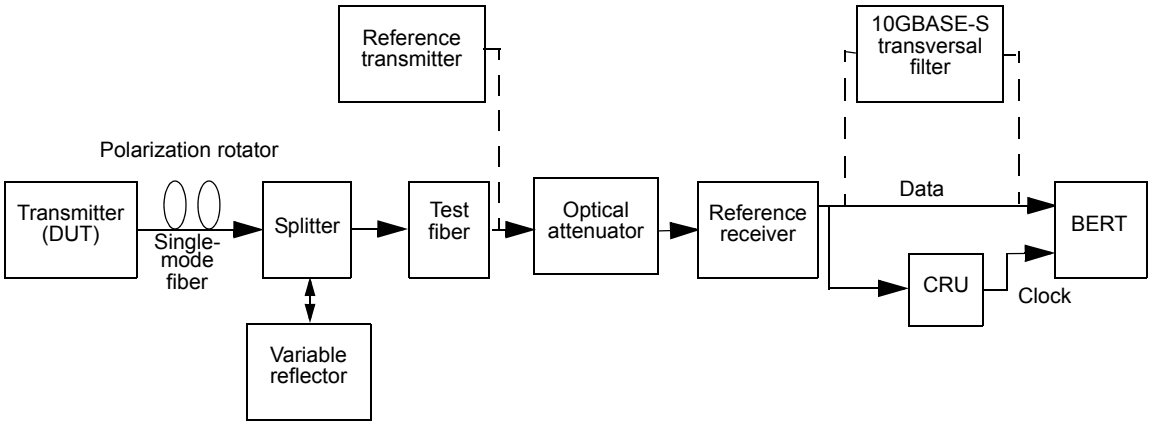

Figure 52–12—Test setup for measurement of transmitter and dispersion penalty

52.9.11 Measurement of the receiver 3 dB electrical upper cutoff frequency

The receiver 3 dB electrical upper cutoff frequency may be measured as described in Figure 52–13 and the list below. The test setup is shown in Figure 52–13. The test uses two optical sources and an optical combiner. One source is modulated by a digital data signal. The other, approximately linear, source is modulated with an analog signal. The analog and digital signals should be asynchronous. The recommended patterns are test patterns 1 or 3 of 52.9.1. An appropriate PRBS or a valid 10GBASE-R or 10GBASE-W signal, OC-192 signal, STM-64 signal, or another representative test pattern may be used.

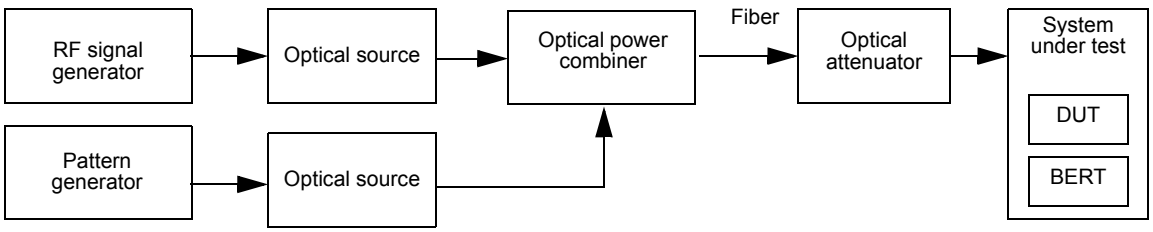

Figure 52–13—Test setup for receiver bandwidth measurement

The 3 dB upper cutoff frequency may be measured using the following steps:

a) Calibrate the frequency response characteristics of the test equipment including the analog radio frequency (RF) signal generator and laser sources. The digital optical source meets the requirements of this clause.

b) Configure the test equipment as shown in Figure 52–13. Take care to minimize changes to the signal path which could affect the system frequency response after the calibration in step a. With the RF modulation turned off, and both optical sources turned on, set the Optical Modulation Amplitude to a level that approximates the stressed receiver sensitivity level in Table 52–9 for 10GBASE-S, Table 52–13 for 10GBASE-L, and Table 52–17 for 10GBASE-E.

c) Turn on the RF modulation while maintaining the same average optical power and OMA established in step b).

d) Measure the necessary RF modulation amplitude (in electrical dBm) required to achieve a constant BER (e.g., 10^{-8}) for a number of frequencies.

e) The receiver 3 dB electrical upper cutoff frequency is that frequency where the corrected RF modulation amplitude (the measured amplitude in "d" corrected with the calibration data in "a") increases by 3 dB (electrical).

52.10 Environmental specifications

52.10.1 General safety

All equipment meeting this standard shall conform to IEC-60950-1.

52.10.2 Laser safety

10GBASE-R and 10GBASE-W optical transceivers shall conform to Hazard Level 1 laser requirements as defined in IEC 60825-1 and IEC 60825-2, under any condition of operation. This includes single fault conditions whether coupled into a fiber or out of an open bore.

Conformance to additional laser safety standards may be required for operation within specific geographic regions.

Laser safety standards and regulations require that the manufacturer of a laser product provide information about the product's laser, safety features, labeling, use, maintenance, and service. This documentation is required to explicitly define requirements and usage restrictions on the host system necessary to meet these safety certifications.[13]

52.10.3 Installation

Sound installation practice, as defined by applicable local codes and regulations, shall be followed in every instance in which such practice is applicable.

52.11 Environment

Normative specifications in this clause shall be met by a system integrating a 10GBASE-R or 10GBASE-W over the life of the product while the product operates within the manufacturer's range of environmental, power, and other specifications.

It is recommended that manufacturers indicate in the literature associated with the PHY the operating environmental conditions to facilitate selection, installation, and maintenance.

It is recommended that manufacturers indicate, in the literature associated with the components of the optical link, the distance and operating environmental conditions over which the specifications of this clause will be met.

52.11.1 Electromagnetic emission

A system integrating a 10GBASE-R or 10GBASE-W PMD shall comply with applicable local and national codes for the limitation of electromagnetic interference.

[13]A host system that fails to meet the manufacturers requirements and/or usage restrictions may emit laser radiation in excess of the safety limits of one or more safety standards. In such a case, the host manufacturer is required to obtain its own laser safety certification.

52.11.2 Temperature, humidity, and handling

The optical link is expected to operate over a reasonable range of environmental conditions related to temperature, humidity, and physical handling (such as shock and vibration). Specific requirements and values for these parameters are considered to be beyond the scope of this standard.

52.12 PMD labeling requirements

It is recommended that each PHY (and supporting documentation) be labeled in a manner visible to the user, with at least the applicable safety warnings and the applicable port type designation (e.g., 10GBASE-ER).

Labeling requirements for Hazard Level 1 lasers are given in the laser safety standards referenced in 52.10.2.

52.13 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 52–14.

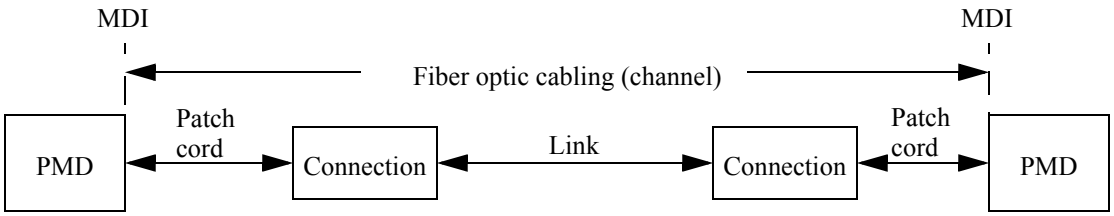

Figure 52–14—Fiber optic cabling model

The channel insertion loss is given in Table 52–24. A channel may contain additional connectors or other optical elements as long as the optical characteristics of the channel, such as attenuation, dispersion, reflections, polarization mode dispersion and modal bandwidth meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with ANSI/TIA/EIA-526-14A/method B, and ANSI/TIA/EIA-526-7/method A-1. The fiber optic cabling model (channel) defined here is the same as a simplex fiber optic link segment. The term channel is used here for consistency with generic cabling standards.

Table 52–24—Fiber optic cabling (channel)

Description	62.5 μm MMF		50 μm MMF				Type B1.1, B1.3 SMF			Unit
Nominal wavelength	**850[a]**						**1310[b]**	**1550**		**nm**
Modal bandwidth (min)	**160**	**200**	**400**	**500**	**2000**	**4700**	**N/A**	**N/A**		**MHz•km**
Operating distance (max)	26 m	33 m	66 m	82 m	300 m	400 m	10 km	30 km	40 km	
Channel insertion loss (max) [c,d,e]	2.6	2.5	2.2	2.3	2.6	2.9	6.0	11.0[f]		dB
Channel insertion loss (min)	0	0	0	0	0	0	0	5		dB
Dispersion (max)	N/A	N/A	N/A	N/A	N/A	N/A	N/A	546	728	ps/nm
DGD_max[g]	N/A	N/A	N/A	N/A	N/A	N/A	10	19		ps
Optical return loss	N/A	N/A	N/A	N/A	N/A	N/A	N/A	21		dB

[a]Channel insertion loss at 850 nm includes cable, connectors, and splices.
[b]Channel insertion loss at 1310 nm includes cable, connectors, and splices.
[c]These channel insertion loss numbers are based on the nominal wavelength.
[d]Operating distances used to calculate channel insertion loss are those listed in this table.
[e]Maximum attenuation given in Table 52–25.
[f]Channel insertion loss at 1550 nm includes cable, connectors and splices.
[g]Differential Group Delay (DGD) is the time difference between the fractions of a pulse that are transmitted in the two principal states of polarization of an optical signal. DGD_max is the maximum differential group delay that the system must tolerate.

52.14 Characteristics of the fiber optic cabling (channel)

The 10GBASE-R and 10GBASE-W fiber optic cabling shall meet the specifications defined in Table 52–24. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together.

52.14.1 Optical fiber and cable

The fiber optic cable shall meet the requirements of IEC 60793-2 or the requirements of Table 52–25 where they differ for fiber types A1a (50/125 μm multimode), A1b (62.5/125 μm multimode), B1.1 (dispersion un-shifted single mode), or B1.3 (low water peak single mode).

NOTE—It is believed that for 10GBASE-E, type B4 fiber with positive dispersion may be substituted for B1.1 or B1.3. A link using B4 fiber with negative dispersion should be validated for compliance at TP3.

52.14.2 Optical fiber connection

An optical fiber connection, as shown in Figure 52–14, consists of a mated pair of optical connectors.

52.14.2.1 Connection insertion loss

The insertion loss is specified for a connection, which consists of a mated pair of optical connectors.

The maximum link distances for multimode fiber are calculated based on an allocation of 1.5 dB total connection and splice loss. For example, this allocation supports three connections with an insertion loss equal to 0.5 dB (or less) per connection, or two connections (as shown in Figure 52–14) with an insertion

Table 52–25—Optical fiber and cable characteristics

Description	62.5 μm MMF	50 μm MMF		Type B1.1, B1.3 SMF		Unit
Nominal fiber specification wavelength	**850**	**850**		**1310**	**1550**	**nm**
Cabled optical fiber attenuation (max)	3.5	3.5		0.4[a] or 0.5[b]	see footnote[c]	dB/km
Modal Bandwidth (min)	160[d] or 200[d]	400[d] or 500[d] or 2000[e]	4700[f]	N/A		MHz km
Zero dispersion wavelength (λ_0)	$1320 \le \lambda_0 \le 1365$	$1295 \le \lambda_0 \le 1320$[g]	$1295 \le \lambda_0 \le 1340$[f]	$1300 \le \lambda_0 \le 1324$		nm
Dispersion slope (max) (S_0)	0.11 for $1320 \le \lambda_0 \le 1348$ and $0.001(1458-\lambda_0)$ for $1348 \le \lambda_0 \le 1365$	0.11 for $1300 \le \lambda_0 \le 1320$ and $0.001(\lambda_0-1190)$ for $1295 \le \lambda_0 \le 1300$[g]	0.105 for $1295 \le \lambda_0 \le 1310$ and $0.000375(1590-\lambda_0)$ for $1310 \le \lambda_0 \le 1340$	0.093		ps / nm^2 km

[a]For the single-mode case, the 0.4 dB/km attenuation for optical fiber cables is defined in ITU-T G.652.
[b]For the single-mode case, the 0.5 dB/km attenuation is provided for Outside Plant cable as defined in ANSI/TIA 568-C.3. Using 0.5 dB/km may not support operation at 10 km.
[c]Attenuation for 1550 nm links is based on the fiber channel and is specified in 52.14.3.
[d]Overfilled launch bandwidth per IEC 60793-1-41 or ANSI/TIA/EIA 455-204-2000.
[e]Effective modal bandwidth for fiber meeting IEC 60793-2-10 Type A1a.2 when used with sources meeting the wavelength (range) and encircled flux specifications of Table 52–7.
[f]Effective modal bandwidth, zero dispersion wavelength and dispersion slope for OM4 fibers are specified in IEC 60793-2-10 Type A1a.3.
[g]Cabled optical fiber with 400 or 500 or 2000 MHz km minimum Modal Bandwidth may alternatively comply with the zero dispersion wavelength and dispersion slope specifications for 4700 MHz km minimum Modal Bandwidth fiber.

loss of 0.75 dB per connection. Connections with different loss characteristics may be used provided the requirements of Table 52–24 are met.

The maximum link distances for single-mode fiber are calculated based on an allocation of 2 dB total connection and splice loss at 1310 nm for 10GBASE-L, and 2 dB for 30 km total connection and splice loss at 1550 nm for 10GBASE-E.

52.14.2.2 Maximum discrete reflectance

The maximum discrete reflectance for 10GBASE-S shall be less than –20 dB.

The maximum discrete reflectance for 10GBASE-L and 10GBASE-E shall be less than –26 dB.

52.14.3 10GBASE-E attenuator management

The 10GBASE-E channel shall have an attenuation between 5 dB and 11 dB. If required, an attenuator can be added to comply with this specification. The ideal channel attenuation is 8 dB. Figure 52–15 graphically shows the compliant region.

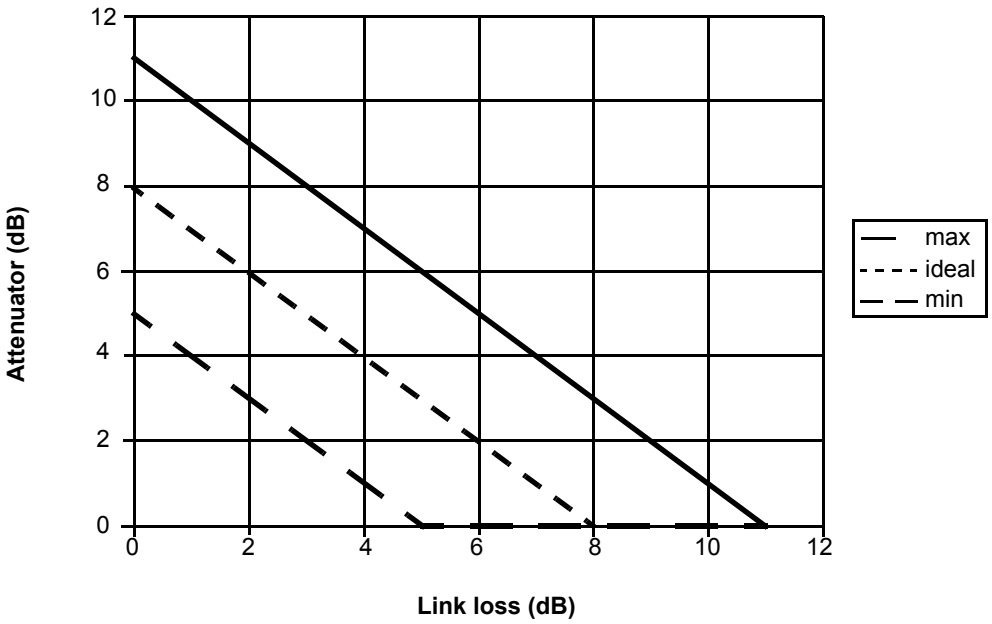

Figure 52–15—10GBASE-E attenuator management (informative)

52.14.4 Medium Dependent Interface (MDI) requirements

The 10GBASE-R and 10GBASE-W PMD is coupled to the fiber optic cabling at the MDI. The MDI is the interface between the PMD and the "fiber optic cabling" (as shown in Figure 52–14). Examples of an MDI include the following:

a) Connectorized fiber pigtail.
b) PMD receptacle.

When the MDI is a connector plug and receptacle connection, it shall meet the interface performance specifications of the following:

c) IEC 61753-1-1—Fibre optic interconnecting devices and passive component performance standard—Part 1-1: General and guidance—Interconnecting devices (connectors).
d) IEC 61753-021-2—Fibre optic passive components performance standard—Part 021-2: Fibre optic connectors terminated on single-mode fibre for Category C—Controlled environment, performance Class S.
e) IEC 61753-022-2—Fibre optic passive components performance standard—Part 022-2: Fibre optic connectors terminated on multimode fibre for Category C—Controlled environment, performance Class M.

NOTE—Compliance testing is performed at TP2 and TP3 as defined in 52.4.1, not at the MDI.

52.15 Protocol implementation conformance statement (PICS) proforma for Clause 52, Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-S (short wavelength serial), 10GBASE-L (long wavelength serial), and 10GBASE-E (extra long wavelength serial)[14]

52.15.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 52, Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-R and 10GBASE-W, shall complete the following protocol implementation conformance statement (PICS) proforma. A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

52.15.2 Identification

52.15.2.1 Implementation identification

Supplier[1]	
Contact point for inquiries about the PICS[1]	
Implementation Name(s) and Version(s)[1,3]	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)[2]	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

52.15.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 52, Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-R and 10GBASE-W
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [ ] Yes [ ] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2015.)	

Date of Statement	

[14] *Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

52.15.2.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*MD	MDIO capability	52.3	Registers and interface supported	O	Yes [ ] No [ ]
*SR	10GBASE-SR PMD	52.5	Device supports shortwave (850 nm) operation LAN PHY	O	Yes [ ] No [ ]
*LR	10GBASE-LR PMD	52.6	Device supports longwave (1310 nm) operation LAN PHY	O	Yes [ ] No [ ]
*ER	10GBASE-ER PMD	52.7	Device supports extra long wave (1550 nm) operation LAN PHY	O	Yes [ ] No [ ]
*SW	10GBASE-SW PMD	52.5	Device supports shortwave (850 nm) operation WAN PHY	O	Yes [ ] No [ ]
*LW	10GBASE-LW PMD	52.6	Device supports longwave (1310 nm) operation WAN PHY	O	Yes [ ] No [ ]
*EW	10GBASE-EW PMD	52.7	Device supports extra long wave (1550 nm) operation WAN PHY	O	Yes [ ] No [ ]
*INS	Installation / Cable	52.4.1	Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]
DLY	Delay constraints	52.2	Device conforms to delay constraints	M	Yes [ ]

52.15.3 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and baseband medium, types 10GBASE-R and 10GBASE-W

52.15.3.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
FS1	Transmit function	52.4.2	Conveys bits from PMD service interface to MDI	M	Yes [ ]
FS2	Transmitter optical signal	52.4.2	Higher optical power transmitted is a logic 1	M	Yes [ ]
FS3	Receive function	52.4.3	Conveys bits from MDI to PMD service interface	M	Yes [ ]
FS4	Receiver optical signal	52.4.3	Higher optical power received is a logic 1	M	Yes [ ]
FS5	Signal detect function	52.4.4	Mapping to PMD service interface	M	Yes [ ]
FS6	Signal detect parameter	52.4.4	Generated according to Table 52–5	M	Yes [ ]

52.15.3.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
MD1	Management register set	52.3	As defined in Table 52–3 and Table 52–4	MD:M	Yes [ ] N/A [ ]
MD2	PMD_reset bit	52.4.5	As defined in 45.2.1.1.1	MD:M	Yes [ ] N/A[ ]
MD3	PMD_global_transmit_disable bit	52.4.7	As defined in 45.2.1.8.7	MD:O	Yes [ ] No [ ] N/A [ ]
MD4	PMD_transmit_fault bit	52.4.8	As defined in 45.2.1.7.4	MD:O	Yes [ ] No [ ] N/A [ ]
MD5	PMA/PMD receive fault bit	52.4.9	Contributes to receive fault bit as defined in 45.2.1.7.5	MD:O	Yes [ ] No [ ] N/A [ ]
MD6	PMD_signal_detect bit	52.4.4	As defined in 45.2.1.9.7	MD:M	Yes [ ] N/A [ ]

52.15.3.3 PMD to MDI optical specifications for 10GBASE-SR

Item	Feature	Subclause	Value/Comment	Status	Support
SR1	10GBASE-SR transmitter	52.5.1	Meets specifications in Table 52–8 and Table 52–7	SR:M	Yes [ ] N/A [ ]
SR2	10GBASE-SR receiver	52.5.2	Meets specifications in Table 52–9	SR:M	Yes [ ] N/A [ ]

52.15.3.4 PMD to MDI optical specifications for 10GBASE-SW

Item	Feature	Subclause	Value/Comment	Status	Support
SW1	10GBASE-SW transmitter	52.5.1	Meets specifications in Table 52–8 and Table 52–7	SW:M	Yes [ ] N/A [ ]
SW2	10GBASE-SW receiver	52.5.2	Meets specifications in Table 52–9	SW:M	Yes [ ] N/A [ ]

52.15.3.5 PMD to MDI optical specifications for 10GBASE-LR

Item	Feature	Subclause	Value/Comment	Status	Support
LR1	10GBASE-LR transmitter	52.6.1	Meets specifications in Table 52–12	LR:M	Yes [ ] N/A [ ]
LR2	10GBASE-LR receiver	52.6.2	Meets specifications in Table 52–13	LR:M	Yes [ ] N/A [ ]

52.15.3.6 PMD to MDI optical specifications for 10GBASE-LW

Item	Feature	Subclause	Value/Comment	Status	Support
LW1	10GBASE-LW transmitter	52.6.1	Meets specifications in Table 52–12	LW:M	Yes [ ] N/A [ ]
LW2	10GBASE-LW receiver	52.6.2	Meets specifications in Table 52–13	LW:M	Yes [ ] N/A [ ]

52.15.3.7 PMD to MDI optical specifications for 10GBASE-ER

Item	Feature	Subclause	Value/Comment	Status	Support
ER1	10GBASE-ER transmitter	52.7.1	Meets specifications in Table 52–16	ER:M	Yes [ ] N/A [ ]
ER2	10GBASE-ER receiver	52.7.2	Meets specifications in Table 52–17	ER:M	Yes [ ] N/A [ ]

NOTE— Link attenuation requirements are specified in 52.14.3.

52.15.3.8 PMD to MDI optical specifications for 10GBASE-EW

Item	Feature	Subclause	Value/Comment	Status	Support
EW2	10GBASE-EW transmitter	52.7.1	Meets specifications in Table 52–16	EW:M	Yes [ ] N/A [ ]
EW2	10GBASE-EW receiver	52.7.2	Meets specifications in Table 52–17	EW:M	Yes [ ] N/A [ ]

52.15.3.9 Optical measurement requirements

Item	Feature	Subclause	Value/Comment	Status	Support
OM1	Measurement cable	52.9	2–5 m in length	M	Yes [ ]
OM2	Center wavelength and spectral width measurement	52.9.2	Measured using an optical spectrum analyzer per TIA-455-127-A under modulated conditions	M	Yes [ ]
OM3	Average optical power	52.9.3	Measured using the methods specified in TIA/EIA-455-95	M	Yes [ ]
OM4	Extinction ratio	52.9.4	Measured using the methods specified in IEC 61280-2-2	M	Yes [ ]
OM5	OMA measurement bandwidth	52.9.5	At least 3/T	M	Yes [ ]
OM6	RIN_XOMA measurement procedure	52.9.6	As described in 52.9.6.1, 52.9.6.2, and 52.9.6.3	M	Yes [ ]
OM7	Transmit eye	52.9.7	10.3125 GBd shall qualify for 10GBASE-W and 10GBASE-R use, measurement at 9.95328 GBd shall qualify for 10GBASE-W use only	M	Yes [ ]
OM8	Stressed receiver conformance test	52.9.9	Performed in accordance with the requirements of 52.9.9.1, 52.9.9.2, 52.9.9.3, and 52.9.9.4	M	Yes [ ]
OM9	Transmitter and dispersion penalty measurement	52.9.10.4		M	Yes [ ]

52.15.3.10 Characteristics of the fiber optic cabling and MDI

Item	Feature	Subclause	Value/Comment	Status	Support
FO1	Fiber optic cabling	52.14	Specified in Table 52–24	INS:M	Yes [ ] N/A [ ]
FO2	Optical fiber characteristics	52.14.1	Meet the requirements of IEC 60793-2 and the requirements of Table 52–25	INS:M	Yes [ ] N/A [ ]
FO3	Maximum discrete reflectance-multimode fiber	52.14.2.2	Less than –20 dB	INS:M	Yes [ ] N/A [ ]
FO4	Maximum discrete reflectance -single-mode fiber	52.14.2.2	Less than –26 dB	INS:M	Yes [ ] N/A [ ]
FO5	10GBASE-E attenuation management	52.14.3	An attenuation between 5 dB and 11 dB	INS:M	Yes [ ] N/A [ ]
FO6	MDI requirements	52.14.4	Meet the interface performance specifications listed	INS:M	Yes [ ] N/A [ ]

52.15.3.11 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
ES1	General safety	52.10.1	Conforms to IEC-60950-1	M	Yes [ ]
ES2	Laser safety —IEC Hazard Level 1	52.10.2	Conform to Hazard Level 1 laser requirements defined in IEC 60825-1 and IEC 60825-2	M	Yes [ ]
ES3	Installation	52.10.3	Follow applicable local codes and regulations	M	Yes [ ]

52.15.3.12 Environment

Item	Feature	Subclause	Value/Comment	Status	Support
EC1	Electromagnetic interference	52.11.1	Comply with applicable local and national codes for the limitation of electromagnetic interference	M	Yes [ ]

53. Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-LX4

53.1 Overview

This clause specifies the 10GBASE-LX4 PMD (including MDI) and the baseband medium for both multimode and single-mode optical fiber. When forming a complete PHY, a PMD shall be connected to the appropriate sublayers (see Table 53–1) and with the management functions that are accessible through the Management Interface defined in Clause 45, all of which are hereby incorporated by reference.

Table 53–1—LX4 PMD type and associated Physical Layer clauses

Associated clause	10GBASE-LX4
46—RS and	Required
XGMII[a]	Optional
47—XGXS and XAUI	Optional
48—10GBASE-X PCS/PMA	Required
49—10GBASE-R PCS	n/a
50—10GBASE-W WIS	n/a

[a]The XGMII is an optional interface. However, if the XGMII is not implemented, a conforming implementation must behave functionally as though the RS and XGMII were present.

Figure 53–1 shows the relationship of the PMD and MDI sublayers to the ISO/IEC (IEEE) OSI reference model.

53.1.1 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 10GBASE-LX4 PMD. The service interface for this PMD is described in an abstract manner and do not imply any particular implementation. The PMD Service Interface supports the exchange of encoded data between peer PMA entities. The PMD translates the encoded data to and from signals suitable for the specified medium.

The following PMD service primitives are defined:

PMD_UNITDATA.request

PMD_UNITDATA.indication

PMD_SIGNAL.indication

53.1.2 PMD_UNITDATA.request

This primitive defines the transfer of data (in the form of encoded 8B/10B characters) from the PMA to the PMD.

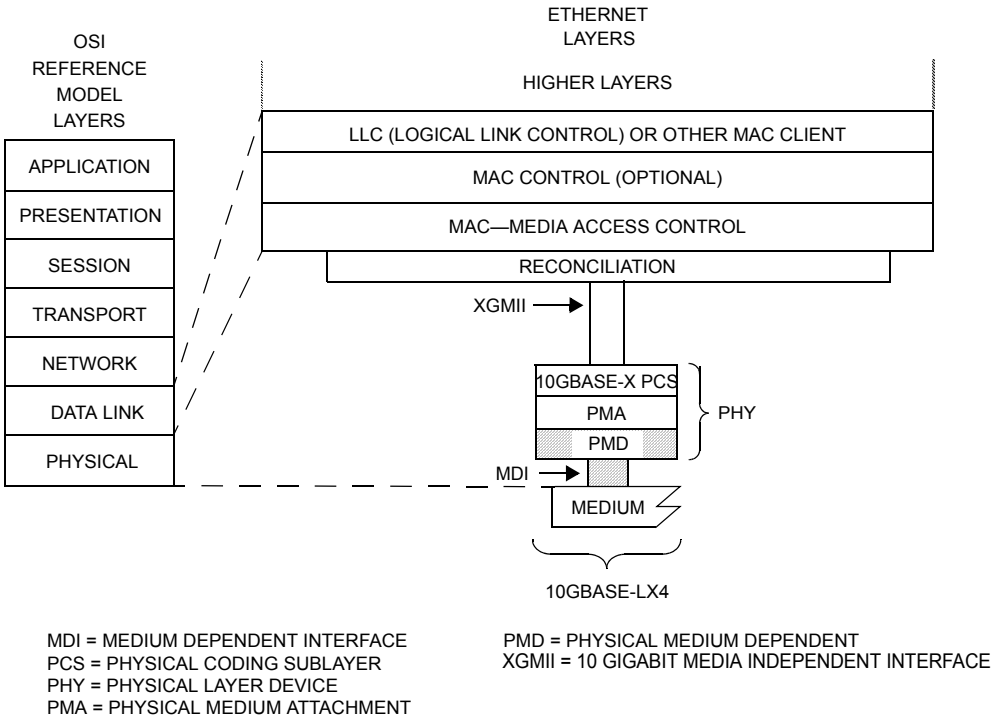

Figure 53–1—10GBASE-LX4 PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

53.1.2.1 Semantics of the service primitive

PMD_UNITDATA.request (tx_bit <0:3>)

The data conveyed by PMD_UNITDATA.request is a continuous sequence of four parallel code-group streams, one stream for each lane. The tx_bit <0:3> correspond to the bits in the tx_lane<0:3> bit streams. Each bit in the tx_bit parameter can take one of two values: ONE or ZERO.

53.1.2.2 When generated

The PMA continuously sends four parallel code-group streams to the PMD at a nominal signaling speed of 3.125 GBd.

53.1.2.3 Effect of Receipt

Upon receipt of this primitive, the PMD converts the specified stream of bits into the appropriate signals on the MDI.

53.1.3 PMD_UNITDATA.indication

This primitive defines the transfer of data (in the form of encoded 8B/10B characters) from the PMD to the PMA.

53.1.3.1 Semantics of the service primitive

PMD_UNITDATA.indication (rx_bit <0:3>)

The data conveyed by PMD_UNITDATA.indication is a continuous sequence of four parallel encoded bit streams. The rx_bit<0:3> correspond to the bits in the rx_lane<0:3> bit streams. Each bit in the rx_bit parameter can take one of two values: ONE or ZERO.

53.1.3.2 When generated

The PMD continuously sends stream of bits to the PMA corresponding to the signals received from the MDI.

53.1.3.3 Effect of receipt

The effect of receipt of this primitive by the client is unspecified by the PMD sublayer.

53.1.4 PMD_SIGNAL.indication

This primitive is generated by the PMD to indicate the status of the signals being received from the MDI.

53.1.4.1 Semantics of the service primitive

PMD_SIGNAL.indication (SIGNAL_DETECT)

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, rx_bit is undefined, but consequent actions based on PMD_UNITDATA.indication, where necessary, interpret rx_bit as a logic ZERO.

NOTE—SIGNAL_DETECT = OK does not guarantee that rx_bit is known to be good. It is possible for a poor quality link to provide sufficient light for a SIGNAL_DETECT = OK indication and still not meet the 10^{-12} BER objective.

53.1.4.2 When generated

The PMD generates this primitive to indicate a change in the value of SIGNAL_DETECT.

53.1.4.3 Effect of receipt

The effect of receipt of this primitive by the client is unspecified by the PMD sublayer.

53.2 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementers must conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices. The sum of transmit and receive delay contributed by the 10GBASE-LX4 PMD shall be no more than 512 BT (including 2 m of fiber).

53.3 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that provide control and status information for and about the PMD. Mapping of MDIO control variables to PMD control variables is shown in Table 53–2. Mapping of MDIO status variables to PMD status variables is shown in Table 53–3.

Table 53–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/ bit number	PMD control variable
Reset	Control register 1	1.0.15	PMD_reset
Global transmit disable	Transmit disable register	1.9.0	Global_PMD_transmit_disable
Transmit disable 3	Transmit disable register	1.9.4	PMD_transmit_disable_3
Transmit disable 2	Transmit disable register	1.9.3	PMD_transmit_disable_2
Transmit disable 1	Transmit disable register	1.9.2	PMD_transmit_disable_1
Transmit disable 0	Transmit disable register	1.9.1	PMD_transmit_disable_0

Table 53–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Local fault	Status register 1	1.1.7	PMD_fault
Transmit fault	Status register 2	1.8.11	PMD_transmit_fault
Receive fault	Status register 2	1.8.10	PMD_receive_fault
Global PMD signal detect	Receive signal detect register	1.10.0	Global_PMD_signal_detect
PMD signal detect 3	Receive signal detect register	1.10.4	PMD_signal_detect_3
PMD signal detect 2	Receive signal detect register	1.10.3	PMD_signal_detect_2
PMD signal detect 1	Receive signal detect register	1.10.2	PMD_signal_detect_1
PMD signal detect 0	Receive signal detect register	1.10.1	PMD_signal_detect_0

53.4 PMD functional specifications

The 10GBASE-LX4 PMD performs the Transmit and Receive functions which convey data between the PMD service interface and the MDI plus various management functions if the optional MDIO is implemented.

53.4.1 PMD block diagram

The PMD block diagram is shown in Figure 53–2. For purposes of system conformance, the PMD sublayer is standardized at the points described in this subclause. The optical transmit signal is defined at the output end of a patch cord (TP2), between 2 m and 5 m in length, of a type consistent with the link type connected

to the transmitter receptacle defined in 53.14.2. If a single-mode fiber offset-launch mode-conditioning patch cord is used, the optical transmit signal is defined at the end of this single-mode fiber offset-launch mode-conditioning patch cord at TP2. Unless specified otherwise, all transmitter measurements and tests defined in 53.9.1 through 53.9.8 are made at TP2. The optical receive signal is defined at the output of the fiber optic cabling (TP3) at the MDI (see 53.14.3). Unless specified otherwise, all receiver measurements and tests defined in 53.9.9 through 53.9.15 are made at TP3.

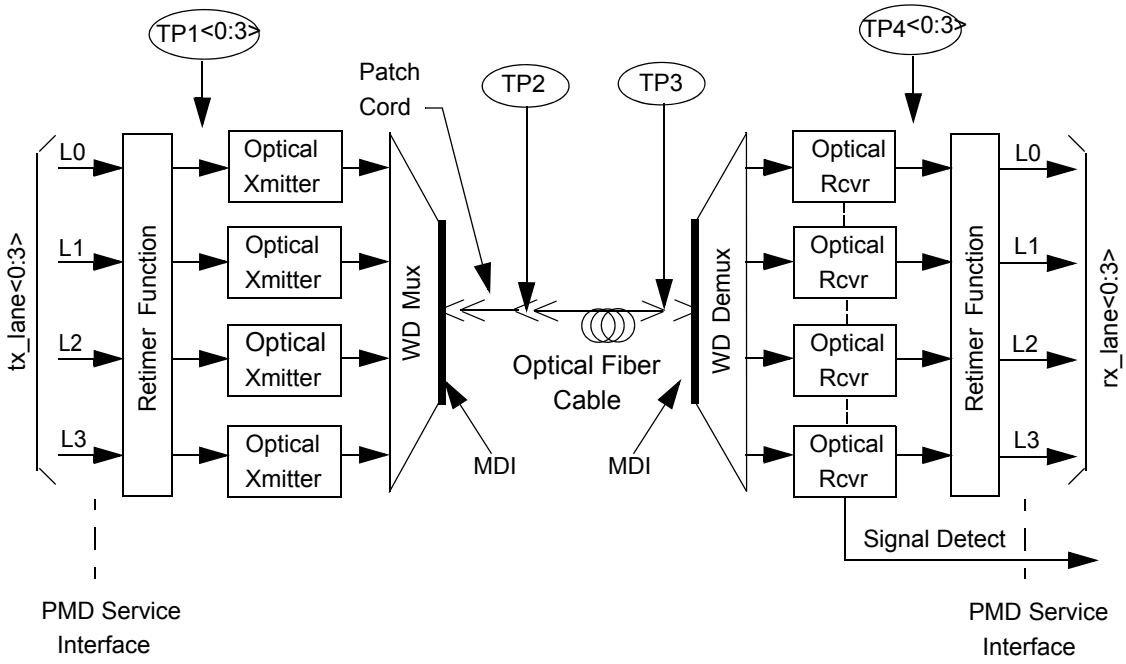

WD = Wavelength Division

NOTE—Specification of the retimer function is beyond the scope of this standard; however, a retimer may be required to ensure compliance at test points TP2 and TP3.

Figure 53–2—Block diagram for LX4 PMD transmit/receive paths

TP1 <0:3> and TP4 <0:3> are informative reference points that may be useful to implementers for testing components (these test points will not typically be testable in an implemented system).

53.4.2 PMD transmit function

The PMD Transmit function shall convert the four electronic bit streams requested by the PMD service interface message PMD_UNITDATA.request (tx_bit<0:3>) into four separate optical signal streams. The four optical signal streams shall then be wavelength division multiplexed and delivered to the MDI, all according to the transmit optical specifications in this clause. The higher optical power level in each signal stream shall correspond to a tx_bit = ONE.

53.4.3 PMD receive function

The PMD Receive function shall demultiplex the composite optical signal stream received from the MDI into four separate optical signal streams. The four optical signal streams shall then be converted into four electronic bit streams for delivery to the PMD service interface using the message PMD_UNITDATA.indication (rx_bit<0:3>), all according to the receive optical specifications in this clause. The higher optical power level in each signal stream shall correspond to a rx_bit = ONE.

The PMD shall convey the bits received from the PMD_UNITDATA.request(tx_bit<0:3>) service primitive to the PMD service interface using the message PMD_UNITDATA.indication(rx_bit<0:3>), where rx_bit<0:3> = tx_bit<0:3>.

53.4.4 Global PMD signal detect function

The Global PMD Receive Signal OK function shall report the state of SIGNAL_DETECT via the PMD service interface. The SIGNAL_DETECT parameter is signaled continuously, while the PMD_SIGNAL.indication message is generated when a change in the value of SIGNAL_DETECT occurs.

SIGNAL_DETECT shall be a global indicator of the presence of optical signals on all four lanes. The PMD receiver is not required to verify whether a compliant 10GBASE-LX4 signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

Table 53–4—SIGNAL_DETECT value definition

Receive conditions	Receive Signal OK value
For any lane; Input_optical_power ≤ –30 dBm	FAIL
For all lanes; [(Input_optical_power ≥ Receiver sensitivity (max) in OMA in Table 53–8) AND (compliant 10GBASE-LX4 signal input)]	OK
All other conditions	Unspecified

As an unavoidable consequence of the requirements for the setting of the SIGNAL_DETECT parameter, implementations must provide adequate margin between the input optical power level at which the SIGNAL_DETECT parameter is set to OK, and the inherent noise level of the PMD due to crosstalk, power supply noise, etc.

Various implementations of the Signal Detect function are permitted by this standard, including implementations that generate the SIGNAL_DETECT parameter values in response to the amplitude of the modulation of the optical signal and implementations that respond to the average optical power of the modulated optical signal.

53.4.5 PMD lane by lane signal detect function

Various implementations of the Signal Detect function are permitted by this standard. When the MDIO is implemented, each PMD_signal_detect_n, where n represents the lane number in the range 0:3, value shall be continuously set in response to the amplitude of the average optical power of the modulated optical signal on its associated lane, according to the requirements of Table 53–4.

53.4.6 PMD reset function

If the MDIO interface is implemented, and if PMD_reset is asserted, the PMD shall be reset as defined in 45.2.1.1.1.

53.4.7 Global PMD transmit disable function

The Global_PMD_transmit_disable function is optional and allows all of the optical transmitters to be disabled.

a) When a Global_PMD_transmit_disable variable is set to ONE, this function shall turn off all of the optical transmitters so that the each transmitter meets the requirements of the Average Launch Power of the OFF Transmitter in Table 53–7.
b) If a PMD_fault is detected, then the PMD may set the Global_PMD_transmit_disable to ONE, turning off the optical transmitter in each lane.

53.4.8 PMD lane by lane transmit disable function

The PMD_transmit_disable function is optional and allows the optical transmitters in each lane to be selectively disabled.

a) When a PMD_transmit_disable_n variable is set to ONE, this function shall turn off the optical transmitter associated with that variable so that the transmitter meets the requirements of the Average Launch Power of the OFF Transmitter in Table 53–7.
b) If a PMD_fault is detected, then the PMD may set each PMD_transmit_disable_n to ONE, turning off the optical transmitter in each lane.

If the optional PMD_lane_by_lane_transmit_disable function is not implemented in MDIO, an alternative method shall be provided to independently disable each transmit lane.

53.4.9 PMD fault function

If the MDIO is implemented, and the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set PMD_fault to ONE.

53.4.10 PMD transmit fault function (optional)

If the MDIO is implemented, and the PMD has detected a local fault on any transmit lane, the PMD shall set the PMD_transmit_fault variable to ONE.

53.4.11 PMD receive fault function (optional)

If the MDIO is implemented, and the PMD has detected a local fault on any receive lane, the PMD shall set the PMD_receive_fault variable to ONE.

53.5 Wavelength-division multiplexed-lane assignments

The wavelengths for each multiplexed lane of the 10GBASE-LX4 PMD shall be as defined in Table 53–5.

53.6 Operating ranges for 10GBASE-LX4 PMD

The operating ranges for 10GBASE-LX4 PMD are defined in Table 53–6. A 10GBASE-LX4 compliant transceiver supports all media types listed in Table 53–6 (i.e., 50 μm and 62.5 μm multimode fiber, and

Table 53–5—Wavelength-division-multiplexed lane assignments

Lane	Wavelength ranges	PMD Service Interface transmit bit stream	PMD Service Interface receive bit stream
L_0	1269.0 – 1282.4 nm	tx_lane<0>	rx_lane<0>
L_1	1293.5 – 1306.9 nm	tx_lane<1>	rx_lane<1>
L_2	1318.0 – 1331.4 nm	tx_lane<2>	rx_lane<2>
L_3	1342.5 – 1355.9 nm	tx_lane<3>	rx_lane<3>

10 µm single-mode fiber) according to the specifications defined in 53.14. A transceiver that exceeds the operational range requirement while meeting all other optical specifications is considered compliant (e.g., a single-mode solution operating at 10 500 m meets the minimum range requirement of 2 m to 10 000 m).

Table 53–6—Operating range for 10GBASE-LX4 PMD over each optical fiber type

Fiber type	Modal bandwidth at 1300 nm (min. overfilled launch) (MHz•km)	Minimum range (meters)
62.5 µm MMF	500	2 to 300
50 µm MMF	400	2 to 240
50 µm MMF	500	2 to 300
10 µm SMF	n/a	2 to 10 000

For MMF links, to ensure that the specifications of Table 53–6 are met with MMF links, the wavelength-division-multiplexed 10GBASE-LX4 transmitter outputs shall be coupled through a single-mode fiber offset-launch mode-conditioning patch cord, as defined in 38.11.4.

53.7 PMD to MDI optical specifications for 10GBASE-LX4

53.7.1 Transmitter optical specifications

The 10GBASE-LX4 transmitters shall meet the specifications defined in Table 53–7 per measurement techniques defined in 53.9.

Table 53–7—Transmit characteristics for the 10GBASE-LX4 over each optical fiber type

Description	62.5 µm and 50 µm MMF	10 µm SMF	Units
Transmitter type	Longwave Laser		
Signaling speed per lane (nominal)	3.125 ± 100 ppm		GBd
Lane wavelengths (range)	1269.0 – 1282.4 1293.5 – 1306.9 1318.0 – 1331.4 1342.5 – 1355.9		nm
Trise/Tfall (max. 20–80 % response time)	120		ps
Side-mode suppression ratio (SMSR), (min)	0.0		dB
RMS spectral width (max)[a]	0.62		nm
Average launch power, four lanes (max)	5.5		dBm
Average launch power, per lane (max)	–0.5		dBm
Optical Modulation Amplitude (OMA), per lane (max)	750 (–1.25)		µW (dBm)
Optical Modulation Amplitude (OMA), per lane (min)[b]	211 (–6.75)	237 (–6.25)	µW (dBm)
Extinction Ratio (min)	3.5		dB
Average launch power of OFF transmitter, per lane (max)	–30		dBm
RIN_{12} (OMA)	–120		dB/Hz

[a]For sources with a distribution of two discrete modes, the spectral window containing 90% of source spectral power (max) equals 1.0 nm and the spectral window containing 99% of source spectral power (max) equals 1.4 nm. For sources with a continuous spectral distribution, the spectral window containing 90% of source spectral power (max) equals 2.0 nm and the spectral window containing 99% of source spectral power (max) equals 3.2 nm. The spectral window specifications for the continuous case is consistent with a Gaussian spectral distribution having an RMS spectral width of 0.62 nm.

[b]Transmitter values are specified at TP2. For MMF, the values in the above table include a 0.5dB attenuation for the offset launch patch cord.

53.7.2 Receive optical specifications

The 10GBASE-LX4 receiver shall meet the specifications defined in Table 53–8 per measurement techniques defined in 53.9. The receive sensitivity includes the extinction ratio penalty.

Table 53–8—10GBASE-LX4 receive characteristics

Description	62.5 μm MMF 50 μm MMF	10 μm SMF	Unit
Signaling speed per lane (nominal)	3.125 ± 100 ppm		GBd
Lane wavelengths (range)	1269.0 – 1282.4 1293.5 – 1306.9 1318.0 – 1331.4 1342.5 – 1355.9		nm
Average receive power, four lanes (max)	5.5		dBm
Average receive power, per lane (max)	–0.5		dBm
Return loss (min)	12		dB
Receive sensitivity (OMA), per lane	37.4 (–14.25)	35.9 (–14.45)	μW (dBm)
Stressed receive sensitivity (OMA),[a,b] per lane	89.0 (–10.5)	45.7 (–13.4)	μW (dBm)
Vertical eye closure penalty,[c] per lane	3.7	1.1	dB
Receive electrical 3 dB upper cutoff frequency, per lane (max)	3750		MHz

[a]Measured with conformance test signal at TP3 (see 53.9.14) for BER = 10^{-12}.
[b]The stressed sensitivity values in the table are for system level BER measurements which include the effects of CDR circuits. It is recommended that at least 0.4dB additional margin be allocated if component level measurements are made without the effect of CDR circuits.
[c]Vertical eye closure penalty is a test condition for measuring stressed receive sensitivity. It is not a required characteristic of the receiver.

53.7.3 Worst case 10GBASE-LX4 link power budget and penalties (informative)

The worst case power budget and link penalties for a 10GBASE-LX4 channel are shown in Table 53–9.

53.8 Jitter specifications for each lane of the 10GBASE-LX4 PMD

53.8.1 Transmit jitter specification

The transmit jitter is tested using a bit error ratio tester[15] (BERT), where the tester horizontally scans the eye opening across the centerline at a virtual TP3 (hereafter referred to as simply TP3) and measures the bit error ratio at each point in time. The plot of BER as a function of sampling time is called the “bathtub curve.”

All points on the BER “bathtub curve” shall have an eye opening greater than specified by the bit error ratio tester (BERT) mask at the respective BER when measured using the jitter compliance test methodology for

[15]In this document, BERT is intended to include any test measurement device that can generate a BER bathtub curve. Here, it is implied that the BERT includes the Reference Receiver, filters, PLLs, etc.

Table 53–9—Worst case 10GBASE-LX4 link power budget and penalties[a]

Parameter	62.5 µm MMF	50 µm MMF		10 µm SMF	Unit
Modal bandwidth as measured at 1300 nm (minimum, overfilled launch)	500	400	500	n/a	MHz•km
Link power budget	7.5	7.5	7.5	8.2	dB
Operating distance[b]	300	240	300	10 000	m
Channel insertion loss[c, d]	2.0	1.9	2.0	6.2	dB
Allocation of penalties	5.0	5.5	5.5	1.9	dB
Additional insertion loss allowed[e]	0.5	0.1	0.0	0.1	dB

[a]Link penalties are used for link budget calculations. They are not requirements and are not meant to be tested.
[b]An offset launch patch cord is required for MMF.
[c]The channel insertion loss is calculated using the maximum distance values specified in Table 53–6 plus an allocation of 1.5dB for MMF and 2.0dB for SMF for connection and splice loss. Channel insertion loss is specified from TP2 to TP3 in Figure 53–2. The total insertion loss for cabled multimode optical fiber, including the attenuation of the offset launch patch cord, is allowed to be 0.5 dB higher than shown in the table.
[d]A wavelength of 1269 nm, a minimum receiver bandwidth of 2550 MHz, and a DCD_DJ of 14 ps is used to calculate lane insertion loss, link power penalties, and unallocated margin.
[e]This portion of the link budget is permitted to be used to overcome insertion loss higher than the "Channel insertion loss" value.

the transmitter as defined in 53.9.10. The BERT mask (see example in Figure 53–3) is specified by the following polynomial:

$$\log_{10}(BER) \le A - B\left(\frac{t - 0.5W}{\sigma}\right)^2 \quad (53–1)$$

$$\log_{10}(BER) \le A - B\left(\frac{1 - t - 0.5W}{\sigma}\right)^2 \quad (53–2)$$

where:

$$A = -1.75, B = \frac{\log_{10}(e)}{2} \approx 0.217$$

and the variables "W" and "σ" for the polynomial are specified in Table 53–10. The variables "W" and "σ" are the effective DJ and RJ respectively although the DJ and RJ values are not normative in the standard. In the above equations the variable "t" is time specified in unit intervals (UI). The BER mask is defined for $10^{-12} <$ BER $< 10^{-6}$. The DJ and RJ values do not need to be individually met, the required mask is defined by the formulas above.

Table 53–10—BERT mask specifications

PMD	W (UI pk to pk)	σ (UI rms)
10GBASE-LX4	0.3	0.015

.

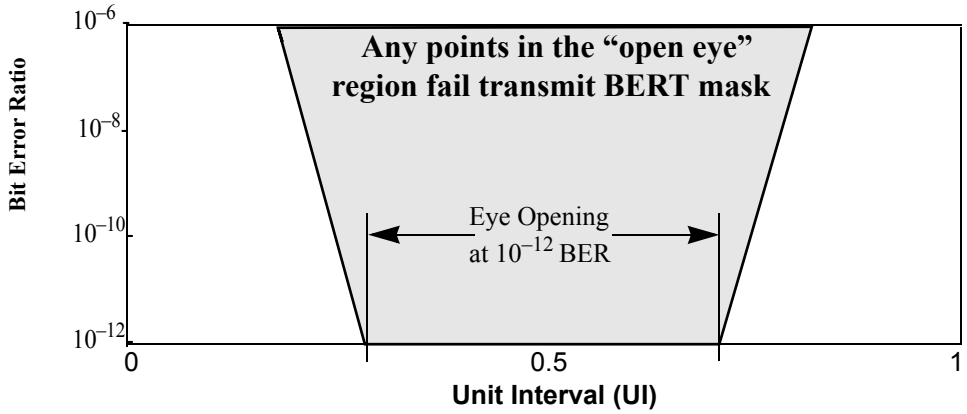

Figure 53–3—Example transmit BER mask at TP3 (informative)

53.8.1.1 Channel requirements for transmit jitter testing

The optical channel for 10GBASE-LX4 shall

a) Have an ITU-T G.652 fiber or fibers with lengths chosen to have a total dispersion larger than specified in Table 53–12 for the wavelength of the device under test. To verify that the fiber has the correct amount of dispersion, use the measurement method defined in ANSI/TIA/EIA-455-175A-92.
b) Meet these requirements in the linear regime of the single-mode fiber.
c) Meet the requirements in Table 53–14 for the case of a 2 to 5 meter offset launch patch cord.

The transmitter shall be tested for single-mode fiber use using the single-mode simulation channel defined in 53.9.10.2. The transmitter shall also be tested for multimode fiber use using the multimode simulation channel defined in 53.9.10.2.

53.8.1.2 Test pattern requirements for transmit jitter testing

Test patterns for 10GBASE-LX4 are specified in Annex 48A.

53.8.2 Receive jitter tolerance specification

The jitter compliance methodology for the receiver is defined in Annex 48B. The receiver shall operate at a BER less than 10^{-12} when tested with an input signal defined in 53.8.2.1 through 53.8.2.1.

The input signal used for jitter tolerance is comprised of the appropriate test pattern, the stressed receiver conformance test signal, the minimum specified input jitter (TJ and DJ) and added sinusoidal jitter.

53.8.2.1 Input jitter for receiver jitter test

The input jitter used to test receiver jitter tolerance shall meet the requirements of the receiver input jitter mask defined by the following equations and Table 53–10.

$$\log_{10}(BER) \geq A - B\left(\frac{t - 0.5W}{\sigma}\right)^2 \quad (53\text{–}3)$$

$$\log_{10}(BER) \geq A - B\left(\frac{1 - t - 0.5W}{\sigma}\right)^2 \quad (53\text{–}4)$$

where:

$$A = -1.75, B = \frac{\log_{10}(e)}{2} \approx 0.217$$

The mask is defined for a BER range of $10^{-12} < BER < 10^{-6}$. An explanation of the variables W, σ, and t can be found in 53.8.1.

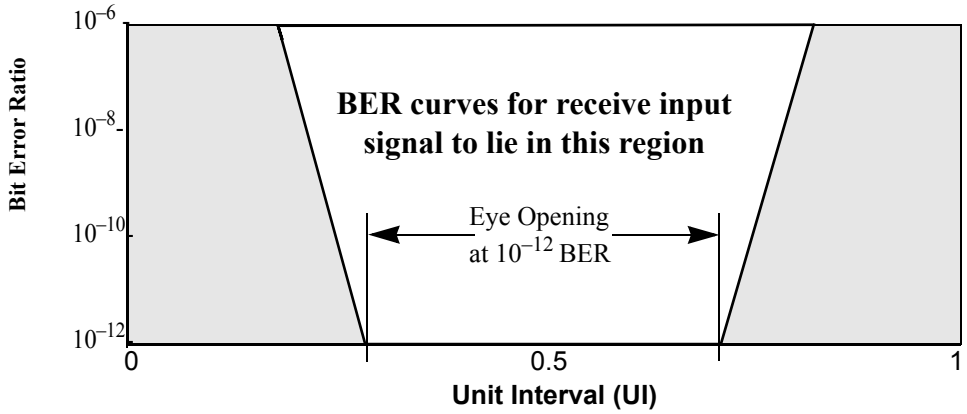

Figure 53–4—Input jitter mask for receiver test (informative)

The random jitter (RJ) component of the input signal shall have uniform spectral content over the measurement frequency range of 18.75kHz to 1.5GHz.

The test method for verification of the input jitter is defined in 53.8.1. A Clock Recovery Unit shall be used for verification of the input jitter. It shall have a high-frequency corner of less than or equal to 1.875MHz and a slope of 20dB/decade. The low-frequency corner corresponds to the point at which the PLL must begin to track low-frequency jitter. The filter used for RX input signal characterization shall be a fourth-order Bessel-Thomson filter as specified in section 53.9.7.

53.8.2.2 Added sinusoidal jitter for receiver jitter test

The sinusoidal jitter used to test receiver jitter tolerance shall meet the requirements of Table 53–11. Sinusoidal jitter shall be added to the test signal that complies with 53.8.2.1.

Table 53–11—Applied sinusoidal jitter

Frequency range	Sinusoidal amplitude jitter (UI pk-pk)
f <18.75 kHz	NA
18.75 kHz < f < 1.875 MHz	$\frac{93750}{f}$
1.875 MHz < f < 10 LB[a]	0.05

[a]LB = Loop Bandwidth; Recommended minimum value of upper bound of 0.05 UI added range is 10 times loop bandwidth.

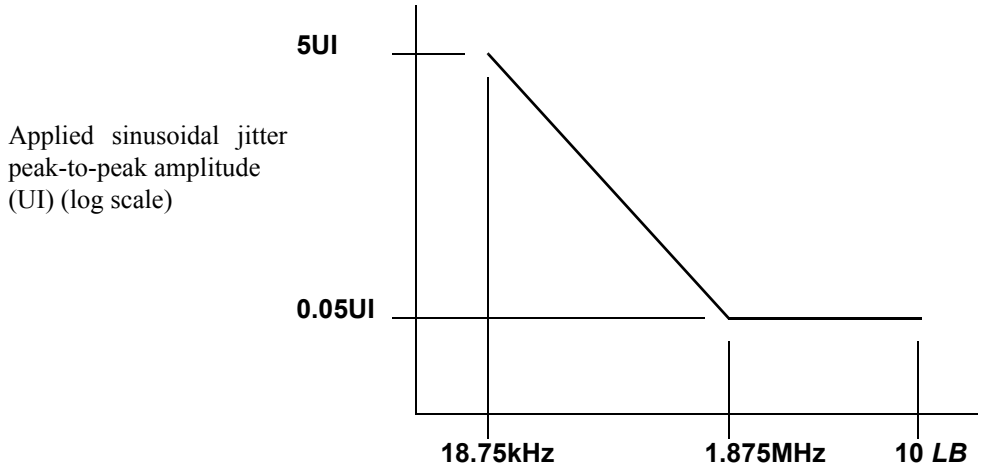

Figure 53–5—Mask of the sinusoidal component of jitter tolerance (informative)

53.9 Optical measurement requirements

All optical measurements shall be made through a short patch cable, between 2 m and 5 m in length. If a single-mode fiber offset-launch mode-conditioning patch cord is used, the optical transmit signal is defined at the output end (TP2) of the single-mode fiber offset-launch mode-conditioning patch cord.

It is recommended that wavelength measurements, RIN measurements, optical power measurements, source spectral window measurements, extinction ratio measurements, optical modulation amplitude measurements, and transmit rise/fall characteristics be performed using a low-frequency square-wave pattern (K28.7).

It is recommended that the receiver 3dB electrical upper cutoff frequency characterization be performed using CRPAT or another reasonable mixed frequency pattern.

Furthermore, it is recommended that the transmitter optical waveform, receive sensitivity measurements, transmitter jitter conformance tests, stressed receiver conformance tests, and jitter tolerance measurements be characterized using CJPAT.

The recommendation of specific test patterns is provided in an effort to meet the overall objective of ensuring compliance of the standard under any and all valid data patterns.

53.9.1 Wavelength range measurements

The wavelength ranges of each wavelength lane shall be measured with an optical spectrum analyzer (OSA) or equivalent device over the wavelength range specified in Table 53–8; with the following conditions:

a) The resolution bandwidth equal to the spectral window values for the particular source type as specified in Table 53–8, and
b) The channel under test is modulated using valid 10GBASE-LX4 signals.

53.9.2 Optical power measurements

The absolute optical power of each channel shall be measured using the methods in TIA/EIA-455-95, with the sum of the optical power from all of the channels not under test below –30 dBm, per the test set-up in Figure 53–6.

.

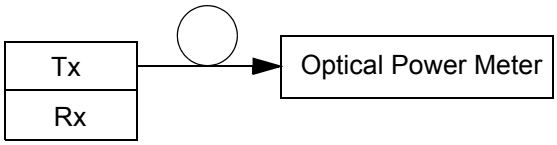

Figure 53–6—Optical power measurement test set-up

53.9.3 Source spectral window measurements

The source spectral window shall be measured for each channel individually with the sum of the optical power from all of the channels not under test below –30dBm, per the test set-up in Figure 53–7. The channel under test shall be modulated using valid 10GBASE-LX4 signals.

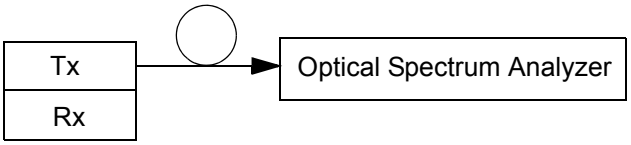

Figure 53–7—Source spectral window measurement test set-up

53.9.4 Extinction ratio measurements

Extinction ratio shall be measured using the methods specified in IEC 61280-2-2. The extinction ratio is measured under fully modulated conditions.

53.9.5 Optical Modulation Amplitude (OMA) measurements

The OMA measurement methodology is defined in 52.9.5 with the exception that each channel will be tested individually and the sum of the optical power from all of the channels not under test shall be below –30 dBm.

53.9.6 Relative Intensity Noise [RIN_{12}(OMA)]

The RIN measurement methodology is defined in 52.9.6 with the exception that each channel will be tested individually, the sum of the optical power from all of the channels not under test shall be below –30 dBm, and the upper –3 dB limit of the measurement apparatus is to be approximately equal to the bit rate (i.e., 3.125 GHz).

53.9.7 Transmitter optical waveform (transmit eye)

The required transmitter pulse shape characteristics are specified in the form of a mask of the transmitter eye diagram as shown in Figure 53–8. Measurements shall be made as per IEC 61280-2-2. The transmit mask is not used for response time and jitter specification.

Normalized amplitudes of 0.0 and 1.0 represent the amplitudes of logic ZERO and ONE respectively.

The eye shall be measured with respect to the mask of the eye using a fourth-order Bessel-Thomson filter having a transfer function given by Equation (53–5):

$$H(p) = \frac{105}{105 + 105\,y + 45\,y^2 + 10\,y^3 + y^4} \tag{53–5}$$

where:

$$y = 2.114\,p; \qquad p = \frac{j\omega}{\omega_r}; \qquad \omega_r = 2\pi f_r \; ; \qquad f_r = 2.344 \text{ GHz}$$

and where the filter response vs. frequency range for this fourth-order Bessel-Thomson filter is defined in ITU-T G.957, along with the allowed tolerances (STM-16 values) for its physical implementation.

This Bessel-Thomson filter is not intended to represent the noise filter used within an optical receiver, but is intended to provide uniform measurement conditions at the transmitter.

The fourth-order Bessel-Thomson filter is reactive. In order to suppress reflections, a 6 dB attenuator may be required at the filter input and/or output.

A Clock Recovery Unit shall be used to trigger the scope for mask measurements. It shall have a low-frequency corner of less than or equal to 1.875 MHz and a slope of 20 dB/decade.

For each lane, the transmit optical waveform is tested with the receive section in operation on all four lanes and with the transmit lanes not under test in operation. CJPAT, CRPAT, or valid 8B/10B encoded data, may be sent to the receive section of the transmitter under test. The data being received must be asynchronous to the transmitted data.

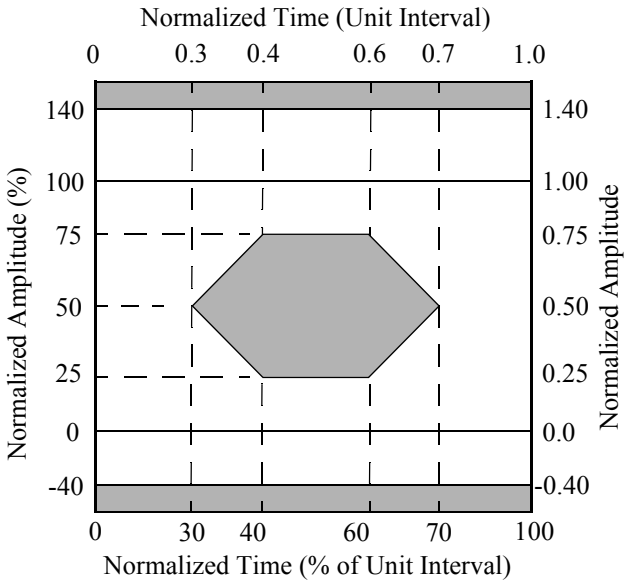

Figure 53–8—Transmitter eye mask definition

53.9.8 Transmit rise/fall characteristics

Optical response time specifications are based on unfiltered waveforms. Some lasers have overshoot and ringing on the optical waveforms, which, if unfiltered, reduce the accuracy of the measured 20–80% response times. For the purpose of standardizing the measurement method, measured waveforms shall conform to the mask defined in Figure 53–8. If a filter is needed to conform to the mask, the filter response shall be removed using Equation (53–6):

$$T_{rise,\,fall} = \sqrt{(T_{rise,\,fall\_measured})^2 - (T_{rise,\,fall\_filter})^2} \quad (53–6)$$

where the filter may be different for rise and fall. Any filter shall have an impulse response equivalent to a fourth-order Bessel-Thomson filter. The fourth-order Bessel-Thomson filter defined in 53.9.7 may be a convenient filter for this measurement; however, its low bandwidth adversely impacts the accuracy of the $T_{rise,fall}$ measurements.

53.9.9 Receive sensitivity measurements

The stressed receive sensitivity shall be measured using the conformance test signal at TP3, as specified in 53.9.14 and meet the conditions specified in Table 53–8.

53.9.10 Transmitter jitter conformance (per lane)

53.9.10.1 Block diagram and general description of test set up

A block diagram for the transmitter jitter conformance test is shown in Figure 53–9. The transmitter (Tx) of the system under test is tested for conformance by continuously generating the test pattern specified in Annex 48A. Depending on the port type, a "Test Fiber" is added to the channel so that the jitter can be measured at a virtual TP3 (hereafter simply referred to simply as TP3) and thus include dispersion and other chromatic and channel induced penalties.

A "Reference Optical Filter" selects the wavelength channel under test in the presence of all channels operating.

A "Reference Receiver (Rx)" converts the optical signal to the electrical domain for input to the BERT.

For 10GBASE-LX4 (single-mode) the receiver shall have a fourth-order Bessel-Thomson response with the transfer function specified in 53.9.7.

Since there is no known way to create a channel using multimode fiber that would yield consistent results, the fiber is omitted from TP2 to TP3.

For 10GBASE-LX4 (multimode), the receiver shall have a fourth-order Bessel-Thomson response with the transfer function specified in 53.9.7 followed by a transversal filter with two equal amplitude paths with a differential delay of 157ps.

The PLL is used in the jitter measurement. It shall have a corner frequency of less than or equal to 1.875MHz and a slope of 20dB/decade. When using a Clock Recovery Unit as a clock for BER measurements, passing of low-frequency jitter from the data to the clock removes this low-frequency jitter from the measurement. The corner frequency corresponds to the point at which the PLL must begin to track this low-frequency jitter.

Jitter shall be measured at the average value of the overall waveform. This can be accomplished with AC-coupling to ground and measuring at ground.

The measurement in this section shall be satisfied with asynchronous data flowing into all four optical receiver channels of the system under test.

The method by which the test pattern is loaded into, or generated by, the BERT is outside the scope of this document.

For each lane, the transmit jitter is tested with the receive section in operation on all four lanes and with the transmit lanes not under test in operation. CJPAT, CRPAT, or valid 8B/10B encoded data, may be sent to the receive section of the transmitter under test. The data being received must be asynchronous to the transmitted data.

Annex 48B contains both theoretical and practical information on jitter testing.

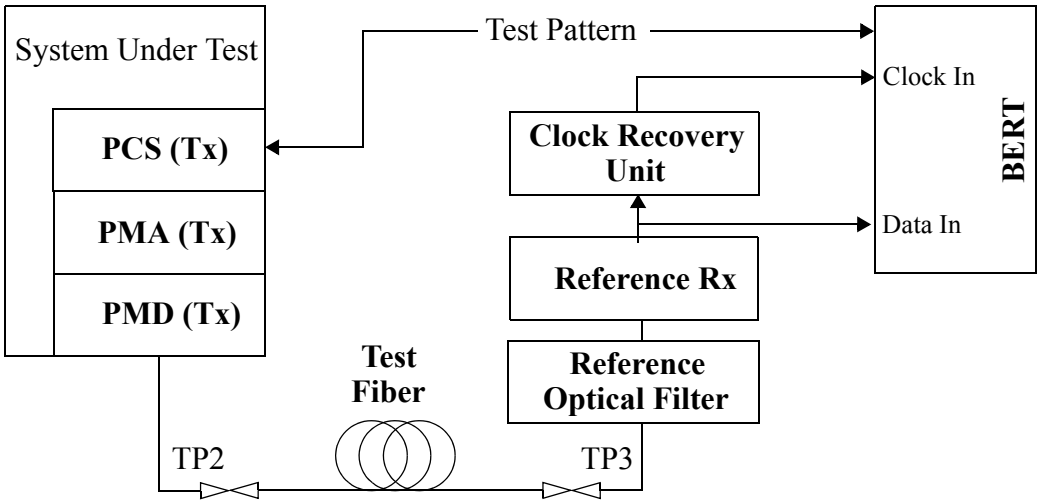

Figure 53–9—Transmit jitter test block diagram

53.9.10.2 Channel requirements for transmit jitter testing

The optical channel used to test the transmitter shall meet the requirements listed in Table 53–12.

Table 53–12—Transmit jitter compliance channel specifications

PMD Type	Minimum dispersion[a] (ps/nm)		Insertion loss[b]	Minimum Return loss[c]
	Minimum	Maximum		
10GBASE-LX4 (multimode)	N/A	N/A	Minimum	See Table 53–8
10GBASE-LX4 (single-mode)	$0.2325 \cdot \lambda \cdot [1 - (1300 / \lambda)^4]$	$0.2325 \cdot \lambda \cdot [1 - (1324 / \lambda)^4]$	Minimum	See Table 53–8

[a]The dispersion is measured for the wavelength of the device under test. The coefficient assumes 10 km for the case of single-mode fiber.
[b]There is no intent to stress the sensitivity of the BERT's optical receiver.
[c]The return loss is applied at TP2. Note that the referenced return losses are maximums. Here, the return loss is a minimum.

For 10GBASE-LX4 (single-mode), the transmitter shall be compliant with dispersion at least as negative as the “minimum dispersion” and at least as positive as the “maximum dispersion” columns. This shall be achieved using ITU-T G.652 fiber or fibers with lengths chosen to have a total dispersion larger than specified in Table 53–12 for the wavelength of the device under test. To verify that the fiber has the correct amount of dispersion, the measurement method defined in ANSI/TIA/EIA-455-175A-92 may be used. The channel shall meet these requirements in the linear regime of the fiber. The channel shall provide an optical back reflection specified in Table 53–7 for 10GBASE-LX4. The state of polarization of the back reflection shall be adjusted to create the greatest RIN. The methods of 52.9.6.2 and 52.9.6.3 may be used.

The channel for 10GBASE-LX4 (multimode) is a 2 to 5 meter patch cord meeting the requirements in Table 53–12. The channel shall provide back reflection to the transmitter at –12 dB. The polarization adjustment does not apply, as noted in 52.9.6.3 item c).

53.9.10.3 Transmit jitter test procedure

After setting up the test as described above, the BERT is scanned horizontally across the center of the eye, from 0 unit intervals (UI) to 1 unit interval, while measuring the error ratio to develop a BER “bathtub curve” as described in 53.8.1.1. This clause also specifies the BER mask. It is not necessary to measure the error ratio in the center of the eye where the BER is less than 10^{-12} or at the edges where the BER is worse than 10^{-4}. The mask natively includes those components of jitter (random, deterministic, bounded) historically measured independently.

A Clock Recovery Unit having a low-frequency corner of less than or equal to 1.875 MHz and a slope of 20 dB/decade shall be used to generate the reference clock for transmit jitter measurements.

NOTE—The specifications for the Reference Optical Filter, Reference Rx, Clock Recovery Unit, and BERT, except as specified above, are outside the scope of this document. The Reference Rx and Clock Recovery Unit are intended to provide consistent and repeatable measurements, not to represent the worst case receiver. It should also be noted that a poorer grade of test equipment will force a greater burden onto the system-under-test to meet specifications. Similarly, a better grade of test equipment will ease the development and manufacture of the system and system components. It is expected that trade-offs needed to optimize the overall cost of development, manufacture and test will change over time and are best left to the implementer.

53.9.11 Receive sensitivity measurements

The receiver sensitivity, which is defined for an ideal input signal is informative. Receivers are tested with a conditioned input signal where both vertical eye closure and jitter have been added according to 53.9.12. When tested according to 53.9.12, the stressed sensitivity shall meet the specifications in Table 53–8 for 10GBASE-LX4.

53.9.12 Stressed receiver conformance test

The measurements in this subclause shall be satisfied with asynchronous data flowing out of the optical transmitter of the system under test. This data shall be consistent with normal signal properties and content.

53.9.12.1 Block diagram of stressed receiver tolerance test set up

A block diagram for the receiver jitter conformance test is shown in Figure 53–10. The receiver of the system under test is tested for conformance by putting the PCS in test mode as specified in Annex 48A. A suitable pattern generator is used to continuously generate a test pattern defined in Annex 48A. The optical test signal is conditioned (stressed) using the stressed receiver methodology defined in 53.9.14 and applying sinusoidal jitter as defined in 53.8.2.2. As defined in section Annex 48A, the PCS is capable of detecting the data pattern and reporting any errors received.

A BER test set or other suitable test set is needed to characterize and verify that the signal used to test the receiver has the appropriate characteristics. This measurement is made using the same method as the transmit jitter test method described in 52.8. The "Test Fiber" called out for LX4 is not needed to characterize the receiver input signal. A "Clock Recovery Unit" meeting the requirements of 53.8.2.1 shall be used.

Jitter shall be calibrated at the average value of the overall optical waveform. This can be accomplished by AC-coupling.

The method by which the test pattern is loaded into, or generated by, the BERT is outside the scope of this document.

Annex 48B contains both theoretical and practical information on jitter testing.

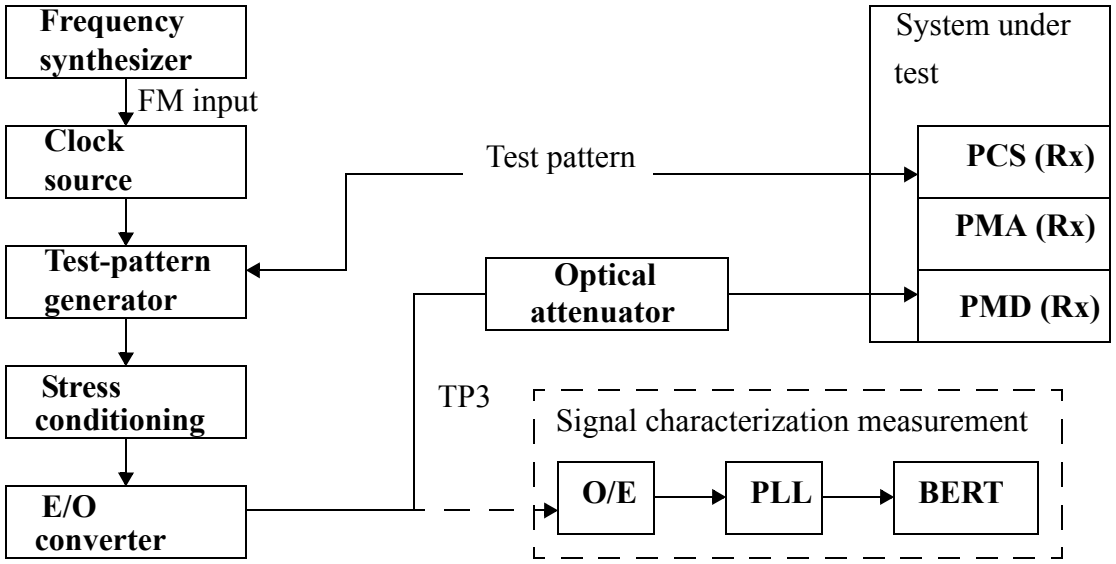

Figure 53–10—Receive jitter test block diagram

53.9.12.2 Stressed receiver conformance test procedure

A BER test set, or other suitable test set, is needed to characterize and verify that the signal used to test the receiver has the appropriate characteristics. This measurement is made using the same method as the transmit jitter test method described in 53.9.10. The "Test Fiber" called out for LX4 (single-mode) and the transversal filter called out for LX4 (multimode) are not needed to characterize the receiver input signal. A Clock Recovery Unit meeting the requirements of 53.8.2.1 shall be used.

53.9.12.3 Characterization of receiver input signal

Care should be taken when characterizing the signal used to make receiver tolerance measurements. In the case of a transmit jitter measurement, excessive and/or uncalibrated noise/jitter in the test system makes it more difficult to meet the specification and may have a negative impact on yield, but will not effect interoperability. In the case of the receiver input calibration measurement, however, excessive noise/jitter will result in an input signal that does not meet the mask requirements defined in 53.8.2.1. Running the receiver tolerance test with a signal that is under-stressed may result in the deployment of non-compliant

receivers. Care should be taken to minimize the noise/jitter introduced by the reference O-E, filters and BERT and/or to correct for this noise. While the details of a BER scan measurement and test equipment are beyond the scope of this document, it is recommended that the implementer fully characterize their test equipment and apply appropriate guard bands to ensure that the RX input signal meets the mask specified in 53.8.2.1.

53.9.12.4 Jitter tolerance test procedure

Set up the test apparatus as described in 53.9.12.1 and adjust the optical input power to the receiver under test to meet the requirements of 53.8.2.1. The sinusoidal jitter is then swept across the frequency and amplitude range specified in 53.8.2.1 while monitoring BER at the receiver. This method does not result in values for TJ, DJ, or RJ contributed by the receiver. It does, however, guarantee that a receiver meeting the requirements of this test will operate with the worst-case optical input

For each lane, the receive jitter tolerance is tested with the transmit section in operation on all four lanes and with the receive lanes not under test in operation. CJPAT, CRPAT, or valid 8B/10B encoded data, may be sent from the transmit section of the receiver under test. The data being transmitted must be asynchronous to the received data.

53.9.13 Measurement of the receiver 3 dB electrical upper cutoff frequency

The receiver cutoff frequency measurement shall be performed on each wavelength channel independently using a laser source with its output wavelength within the specified wavelength range of the channel to be tested. The test setup is shown in Figure 53–11. The test is performed with a tunable laser that is suitable for analog signal transmission. The laser is modulated by a digital data signal. In addition to the digital modulation the laser is modulated with an analog signal. The analog and digital signals should be asynchronous. The data pattern to be used for this test is the short continuous random test pattern defined in Annex 48A. The frequency response of the laser must be sufficient to allow it to respond to both the digital modulation and the analog modulation. The laser should be biased so that it remains linear when driven by the combined signals. The test may use two optical sources and an optical combiner as defined in 52.9.11.

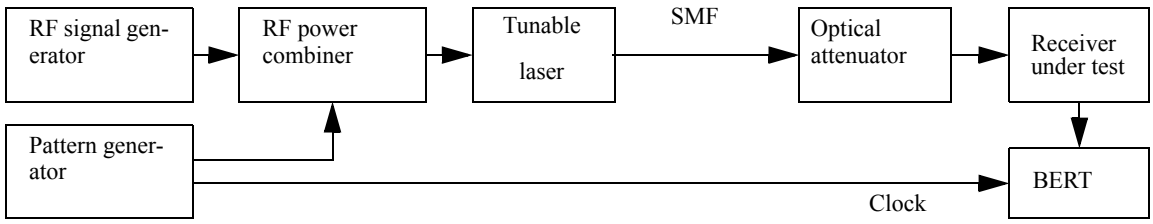

Figure 53–11—Test setup for receiver bandwidth measurement

The 3 dB upper cutoff frequency shall be measured using the following steps:

a) Calibrate the frequency response characteristics of the test equipment including the analog radio frequency (RF) signal generator, RF power combiner, and laser source, with the optical source meeting the requirements of this clause.
b) Configure the test equipment as shown in Figure 53–11. Take care to minimize changes to the signal path that could affect the system frequency response after the calibration in step a). Connect the laser output with no RF modulation applied to the receiver under test through an optical attenuator and

taking into account the extinction ratio of the source, set the optical power to a level that approximates the stressed receive sensitivity level in Table 53–8.

c) Locate the center of the eye with the BERT. Turn on the RF modulation while maintaining the same average optical power established in step b).

d) Measure the necessary RF modulation amplitude (in dBm) required to achieve a constant BER (e.g., 10^{-8}) for a number of frequencies.

e) The receiver 3 dB electrical upper cutoff frequency is that frequency where the corrected RF modulation amplitude (the measured amplitude in "d" corrected with the calibration data in "a") increases by 3 dB (electrical). If necessary, interpolate between the measured response values.

53.9.14 Conformance test signal at TP3 for receiver testing

Receivers being tested for conformance to the stressed receive sensitivity requirements of 53.9.9 and the total jitter requirements of 53.9.10 shall be tested using a conformance test signal at TP3 conforming to the requirements described in Figure 53–12. It is recommended that the conformance test signal be generated using a short continuous random test pattern as defined in Annex 48A. The conformance test signal is conditioned by applying deterministic jitter (DJ) and intersymbol interference (ISI). The conditioned conformance test signal is shown schematically in Figure 53–12. The horizontal eye closure (reduction of pulse width) caused by the duty cycle distortion (DCD) component of DJ shall be no less than 14 ps.

The vertical eye closure penalty shall be greater than or equal to the value specified in Table 53–8. The DJ cannot be added with a simple phase modulation, which does not account for the DCD component of DJ.

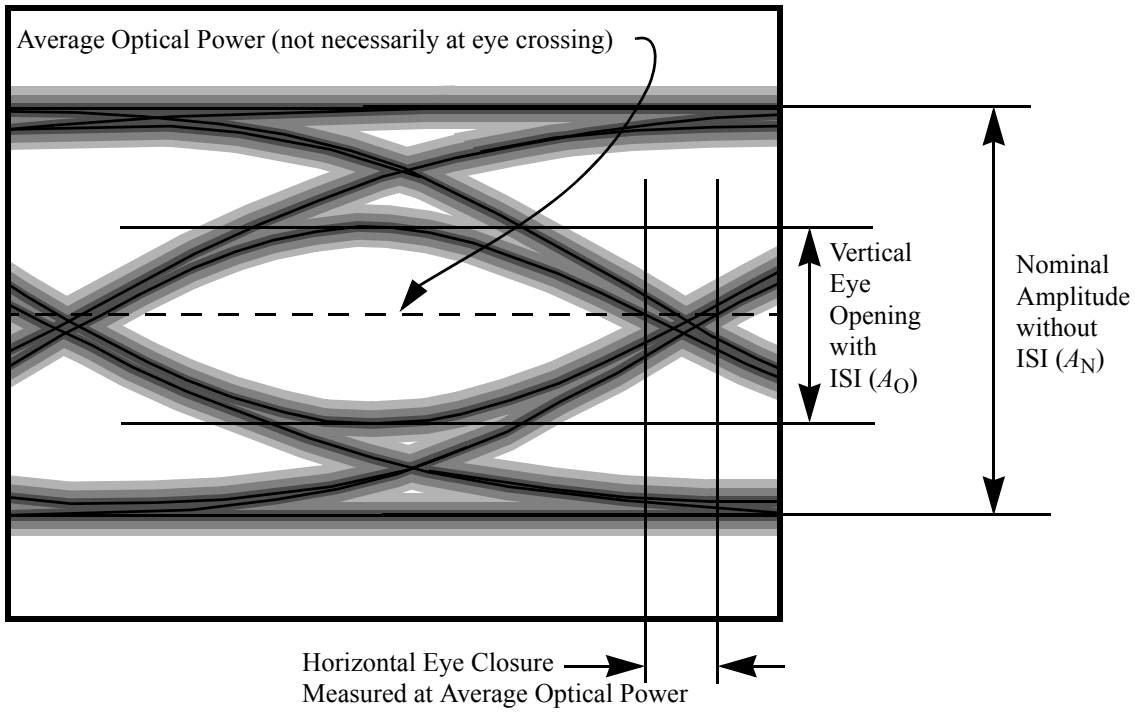

Figure 53–12—Required characteristics of the conformance test signal at TP3

The vertical eye closure penalty is given by:

$$\text{Vertical eye closure penalty [dB]} = 10 \times \log \frac{A_O}{A_N}$$

where, A_O is the amplitude of the eye opening and A_N is the normal amplitude without ISI, as measured in Figure 53–12.

Figure 53–13 shows a possible test set up for producing the conformance test signal at TP3. The coaxial cable is adjusted in length to produce the correct DCD component of DJ. Since the coaxial cable can produce the incorrect ISI, a limiting amplifier is used to restore fast rise and fall times. A Bessel-Thomson filter is selected to produce the minimum ISI induced eye closure as specified per Figure 53–12. This conditioned signal is used to drive a high bandwidth, tunable, wavelength tunable source.

Figure 53–13 shows this function being performed by a tunable source in combination with an external optical modulator. However, any other method capable of this combined function will suffice. Similarly, the remaining sources supply to their respective channels modulated signals at specific wavelengths, as specified in 53.9.15. This could be accomplished with tunable or fixed sources at the wavelengths required. The vertical and horizontal eye closures to be used for receiver conformance testing are verified using a fast photodetector and amplifier coupled to the oscilloscope input through a filter. The combined filtering effect of the photodetector, amplifier, and filter shall be a fourth-order Bessel-Thomson filter of 2.34 GHz bandwidth. Special care should be taken to ensure that all the light from the fiber is collected by the fast photodetector and that there is negligible mode selective loss, especially in the optical attenuator.

The source for the channel under test shall be set to supply a signal at the output of the optical multiplexer at the minimum OMA with all other remaining channels set to the maximum OMA. Each channel is to be tested with its adjacent channels set at the near end of their wavelength range. This is to occur sequentially, as described in 53.9.15. The channel under test is to be tuned over its wavelength range during a given measurement to account for wavelength dependent losses within the channel.

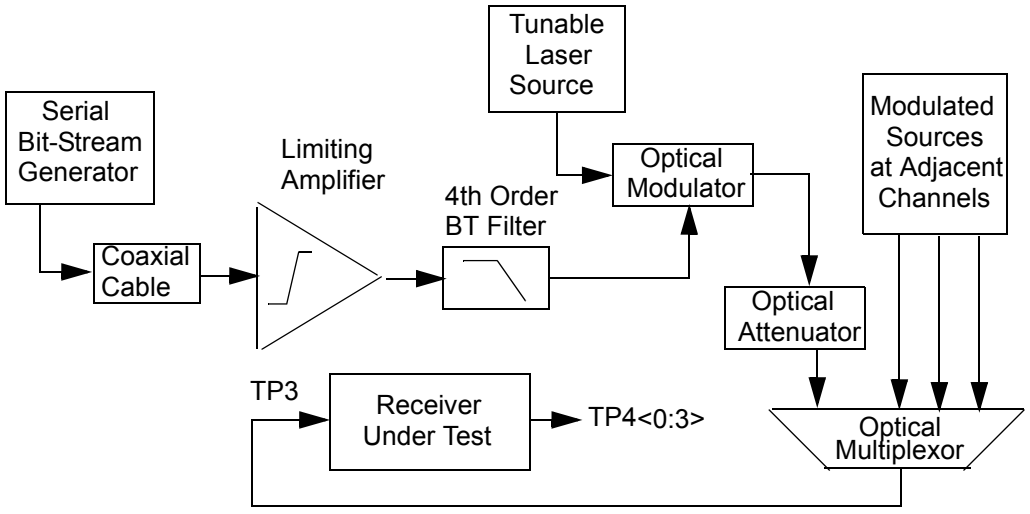

Figure 53–13—Setup for generating receiver conformance test signal at TP3

53.9.15 Receiver test suite for WDM conformance testing

The receiver tests requiring the TP3 conformance test signal are performed on a per channel basis and shall meet the following test conditions:

a) All channels are modulated simultaneously, using valid 10GBASE-LX4 signals.
b) The center wavelengths of channels adjacent to the channel under test are tuned to the edge of their wavelength band nearest the channel under test.
c) When setting the wavelength of the channels adjacent to the channel under test, the center wavelength of the adjacent channels are set within 0.5 nm of the edge of that channel's wavelength band while remaining within that channel's wavelength band.
d) In the case of the interior channels, which have two adjacent channels, each adjacent channel is tuned individually and receiver testing is done twice, once for each adjacent channel.
e) The non-adjacent channels are to be tuned to the center of their respective wavelength ranges.
f) The transmitter of the transceiver under test is operating with valid test patterns as defined in Annex 48A.

These conditions are summarized graphically in Figure 53–14 for each channel under test.

a) Channel L0 under test

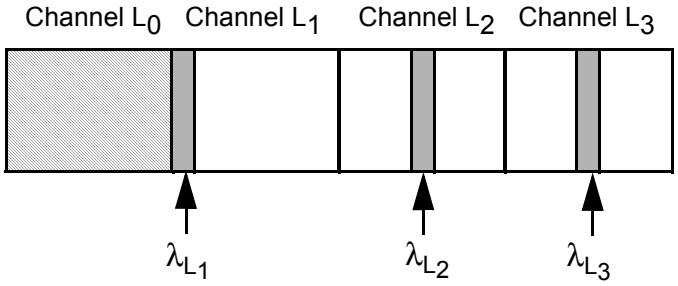

Channel L1 under test (L0 in extreme position)

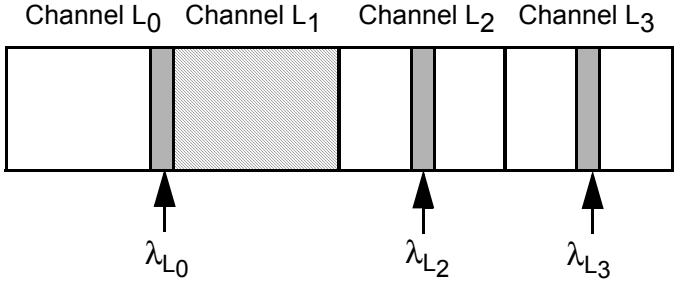

Channel L1 under test (L2 in extreme position)

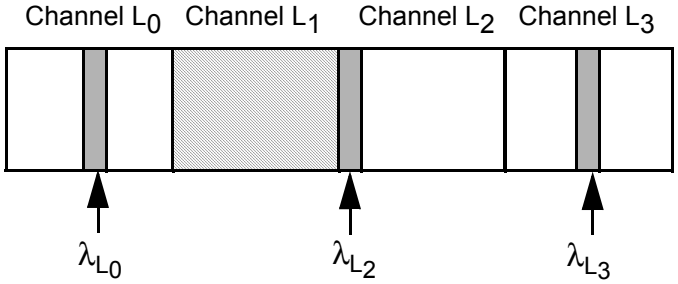

Channel L2 under test (L1 in extreme position)

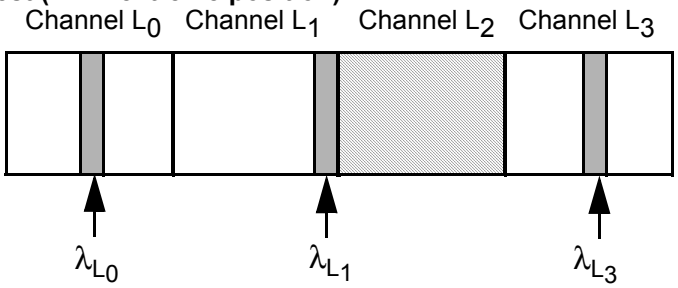

Figure 53–14—Channel test conditions for conformance test at TP3

Channel L2 under test (L3 in extreme position)

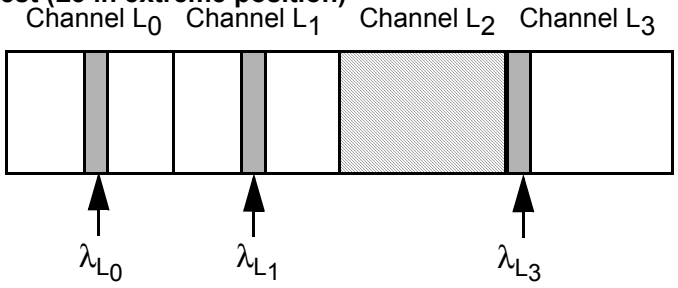

Channel L3 under test (L2 in extreme position)

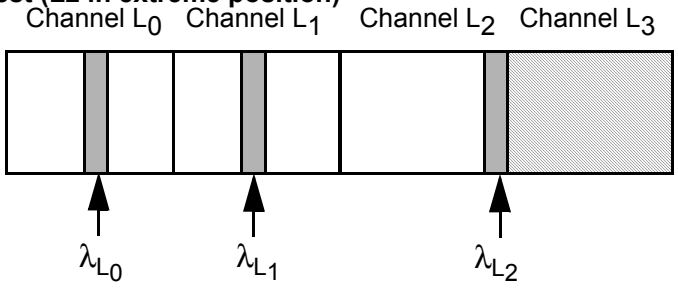

Figure 53–14—Channel test conditions for conformance test at TP3 *(continued)*

NOTE—The two interior channels (L_1, L_2) require two different wavelength configurations since they have two adjacent channels. Therefore, there will be twice as many tests to perform on these channels as on exterior channels L_0, L_3.

53.10 Environmental specifications

53.10.1 General safety

All equipment meeting this standard shall conform to IEC-60950-1.

53.10.2 Laser safety

The 10GBASE-LX4 optical transceivers shall be Hazard Level 1 laser certified under any condition of operation in conformance to IEC 60825-1 and IEC 60825-2. This includes single fault conditions whether coupled into a fiber or out of an open bore.

Conformance to additional laser safety standards may be required for operation within specific geographic regions.

Laser safety standards and regulations require that the manufacturer of a laser product provide information about the product's laser, safety features, labeling, use, maintenance, and service. This documentation explicitly define requirements and usage restrictions on the host system necessary to meet these safety certifications.[16]

[16]A host system that fails to meet the manufacturers requirements and/or usage restrictions may emit laser radiation in excess of the safety limits of one or more safety standards. In such a case, the host manufacturer is required to obtain its own laser safety certification.

53.10.3 Installation

It is recommended that proper installation practices, as defined by applicable local codes and regulation, be followed in every instance in which such practices are applicable.

53.11 Environment

Normative specifications in this clause shall be met by a system integrating a 10GBASE-LX4 PMD over the life of the product while the product operates within the manufacturer's range of environmental, power, and other specifications.

It is recommended that manufacturers indicate in the literature associated with the PHY the operating environmental conditions to facilitate selection, installation, and maintenance.

It is recommended that manufacturers indicate, in the literature associated with the components of the optical link, the distance and operating environmental conditions over which the specifications of this clause will be met.

53.11.1 Electromagnetic emission

A system integrating a 10GBASE-LX4 PMD shall comply with applicable local, national, and international codes for the limitation of electromagnetic interference.

53.11.2 Temperature, humidity, and handling

The optical link is expected to operate over a reasonable range of environmental conditions related to temperature, humidity, and physical handling (such as shock and vibration). Specific requirements and values for these parameters are considered to be beyond the scope of this standard.

53.12 PMD labeling requirements

It is recommended that each PHY (and supporting documentation) be labeled in a manner visible to the user with at least the following parameters, according to the PMD-MDI type.

PMD MDI type 10GBASE-LX4:

a) 10GBASE-LX4;
b) Applicable safety warnings.

Labeling requirements for Hazard Level 1 lasers are given in the laser safety standards referenced in 53.10.2.

53.13 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 53–15.

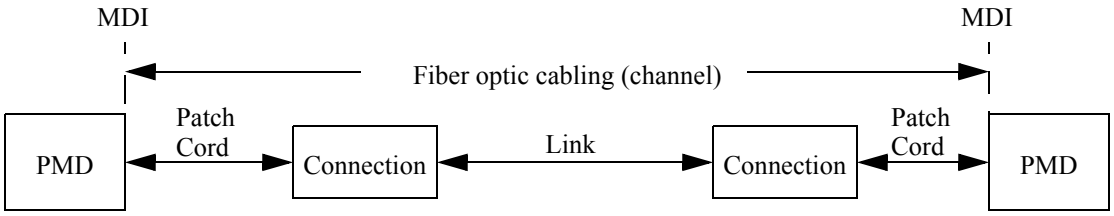

Figure 53–15—Fiber optic cabling model

The channel insertion loss is given in Table 53–13. Insertion loss measurements of installed fiber cables are made in accordance with ANSI/TIA/EIA-526-14A/method B, and ANSI/TIA/EIA-526-7/method A-1. The fiber optic cabling model (channel) defined here is the same as a simplex fiber optic link segment. The term channel is used here for consistency with generic cabling standards.

Table 53–13—Fiber optics cabling (channel) characteristics

Description	62.5 μm MMF	50 μm MMF		SMF	Unit
Modal bandwidth at 1300nm (min; overfilled launch)	**500**	**400**	**500**	**n/a**	**MHz•km**
Operating distance (max)	300	240	300	10 000	m
Channel insertion loss (max)[a]	2.5	2.0	2.0	6.6	dB

[a]These channel insertion loss numbers are calculated using a wavelength of 1300nm for MMF and 1310nm for SMF. An offset launch patch cord is assumed. Channel insertion loss is specified from TP2 to TP3 in Figure 53–2. The total insertion loss, when including the attenuation of the offset launch patch cord, is allowed to be 0.5dB higher than shown in the table.

53.14 Characteristics of the fiber optic cabling (channel)

The 10GBASE-LX4 fiber optic cabling shall meet the specifications defined in Table 53–13. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together. The fiber optic cabling spans from one MDI to another MDI, as shown in Figure 53–15.

53.14.1 Optical fiber and cable

The fiber optic cable requirements are satisfied by the fibers specified in IEC 60793-2:1992. Types A1a (50/125 μm multimode), A1b (62.5/125 μm multimode), and B1.1 (dispersion un-shifted single-mode) and B1.3 (low water peak single-mode) with the exceptions noted in Table 53–14.

Table 53–14—Optical fiber and cable characteristics

Description	62.5 μm MMF	50 μm MMF	Type B1.1, B1.3 SMF	Unit
Nominal fiber specification wavelength	1300	1300	1310	nm
Cabled optical fiber attenuation (max)	1.5	1.5	0.4[a] or 0.5[b]	dB/km
Modal Bandwidth	500	400	N/A	MHz km
(min; overfilled launch)		500	N/A	MHz km
Zero dispersion wavelength (λ_0)	$1320 \le \lambda 0 \le 1365$	$1295 \le \lambda 0 \le 1320$	$1300 \le \lambda 0 \le 1324$	nm
Dispersion slope (max) (S_0)	0.11 for $1320 \le \lambda_0 \le 1348$ and $0.001(1458 - \lambda_0)$ for $1348 \le \lambda_0 \le 1365$	0.11 for $1300 \le \lambda_0 \le 1320$ and $0.001(\lambda_0 - 1190)$ for $1295 \le \lambda_0 \le 1300$	0.093	ps/nm^2 km

[a]For the single-mode case, the 0.4 dB/km attenuation for optical fiber cables is defined in ITU-T G.652.
[b]For the single-mode case, the 0.5 dB/km attenuation is provided for Outside Plant cable as defined in ANSI/TIA 568-C.3. Using 0.5 dB/km may not support operation at 10 km.

53.14.2 Optical fiber connection

An optical fiber connection as shown in Figure 53–15 consists of a mated pair of optical connectors. The 10GBASE-LX4 PMD is coupled to the fiber optic cabling through a connector plug into the MDI optical receptacle, as shown in 53.14.3.

53.14.2.1 Connection insertion loss

The insertion loss is specified for a connection that consists of a mated pair of optical connectors.

The maximum link distances for multimode fiber are calculated based on an allocation of 1.5 dB total connection and splice loss. For example, this allocation supports three connections with an average insertion loss equal to 0.5 dB (or less) per connection. Connections with different loss characteristics may be used provided the requirements of Table 53–13 and Table 53–14 are met.

The maximum link distances for single-mode fiber are calculated based on an allocation of 2.0 dB total connection and splice loss. For example, this allocation supports four connections with an average insertion loss per connection of 0.5 dB. Connections with different loss characteristics may be used provided the requirements of Table 53–13 and Table 53–14 are met.

53.14.2.2 Connection return loss

The return loss for multimode connections shall be greater than 20 dB.

The return loss for single-mode connections shall be greater than 26 dB.

53.14.3 Medium Dependent Interface (MDI)

The 10GBASE-LX4 PMD is coupled to the fiber optic cabling at the MDI. The MDI is the interface between the PMD and the fiber optic cabling. Examples of an MDI include the following:

a) Connectorized fiber pigtail.
b) PMD receptacle.

When connected to the fiber optic cabling, the MDI shall meet the interface performance specifications of:

c) IEC 61753-1-1—Fibre optic interconnecting devices and passive component performance standard—Part 1-1: General and guidance—Interconnecting devices (connectors).
d) IEC 61753-021-2—Fibre optic passive component performance standard—Part 021-2:Fibre optic connectors terminated on single-mode fibre for Category C—Controlled environment, performance Class S.
e) IEC 61753-022-2—Fibre optic passive component performance standard—Part 022-2:Fibre optic connectors terminated on multimode fibre for Category C—Controlled environment, performance Class M.

NOTE—Compliance testing is performed at TP2 and TP3 as defined in 53.4.1, not at the MDI.

53.15 Protocol implementation conformance statement (PICS) proforma for Clause 53, Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-LX4[17]

53.15.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 53, Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-LX4 (Long Wavelength Laser), shall complete the following protocol implementation conformance statement (PICS) proforma. A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

53.15.2 Identification

53.15.2.1 Implementation identification

Supplier[1]	
Contact point for inquiries about the PICS[1]	
Implementation Name(s) and Version(s)[1,3]	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)[2]	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

53.15.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 53, Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-LX4
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [ ] Yes [ ] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2015.)	

Date of Statement	

[17] *Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

53.15.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
LX4	10GBASE-LX4 PMD	53.1	Device supports long wavelength operation (1269–1356 nm)	O/1	Yes [ ] No [ ]
*INS	Installation / cable	53.13	Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]
OFP	Single-mode offset-launch mode-conditioning patch cord	53.6	Items marked with OFP include installation practices and cable specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]
TP1	Standardized reference point TP1 exposed and available for testing	53.4.1	This point may be made available for use by implementers to certify component conformance	O	Yes [ ] No [ ]
TP4	Standardized reference point TP4 exposed and available for testing	53.4.1	This point may be made available for use by implementers to certify component conformance	O	Yes [ ] No [ ]
DC	Delay constraints	53.2	Device conforms to delay constraints	M	Yes [ ]
*MD	MDIO capability	53.3	Registers and interface supported	O	Yes [ ] No [ ]

53.15.4 PICS proforma tables for 10GBASE-LX4 and baseband medium

53.15.4.1 PMD Functional specifications

Item	Feature	Sub clause	Value/Comment	Status	Support
FN1	Integration with 10GBASE-X PCS and PMA and management functions	53.1		M	Yes [ ]
FN2	Transmit function	53.4.2	Convey bits requested by PMD_UNITDATA.request() to the MDI	M	Yes [ ]
FN3	Optical multiplexing and delivery to the MDI	53.4.2	Optically multiplexes the four optical signal streams for delivery to the MDI	M	Yes [ ]
FN4	Mapping between optical signal and logical signal for transmitter	53.4.2	Higher optical power is a one	M	Yes [ ]
FN5	Receive function	53.4.3	Convey bits received from the MDI to PMD_UNITDATA.indication(rx_bit<0:3>)	M	Yes [ ]
FN6	Conversion of four optical signals to four electrical signals	53.4.3	Converts the four optical signal streams into four electrical bit streams for delivery to the PMD service	M	Yes [ ]
FN7	Mapping between optical signal and logical signal for receiver	53.4.3	Higher optical power is a one	M	Yes [ ]
FN8	Receive function behavior	53.4.3	Conveys bits from PMD service primitive to the PMD service interface	M	Yes [ ]
FN9	Global Signal Detect function	53.4.4	Report to the PMD service interface the message PMD_SIGNAL.indication(SIGNAL_DETECT)	M	Yes [ ]
FN10	Global Signal Detect behavior	53.4.4	SIGNAL_DETECT is a global indicator of the presence of optical signals on all four lanes	M	Yes [ ]
FN11	Lane-by-Lane Signal Detect function	53.4.5	Sets PMD_signal_detect_n values on a lane-by-lane basis per requirements of Table 53–4	MD:O	Yes [ ] No [ ] N/A [ ]
FN12	PMD_reset function	53.4.6	Resets the PMD sublayer	MD:O	Yes [ ] No [ ] N/A [ ]

53.15.4.2 PMD to MDI optical specifications for 10GBASE-LX4

Item	Feature	Subclause	Value/Comment	Status	Support
PMS1	Wavelength division multiplexed lane assignment	53.5	Device supports passbands defined in Table 53–5	M	Yes [ ] N/A [ ]
PMS2	Transmitter meets specifications in Table 53–7	53.7.1	Per measurement techniques in 53.9	M	Yes [ ] N/A [ ]
PMS3	Receiver meets specifications in Table 53–8	53.7.2	Per measurement techniques in 53.9	M	Yes [ ] N/A [ ]

53.15.4.3 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
MR1	Management register set	53.3		MD:M	Yes [ ] N/A [ ]
MR2	Global transmit disable function	53.4.7	Disables all of the optical transmitters with the Global_PMD_transmit_disable variable	MD:O	Yes [ ] No [ ] N/A [ ]
MR3	PMD_lane_by_lane_transmit_disable function	53.4.8	Disables the optical transmitter on the lane associated with the PMD_transmit_disable_n variable	MD:O.2	Yes [ ] No [ ] N/A [ ]
MR4	PMD_lane_by_lane_transmit_disable	53.4.8	Disables each optical transmitter independently if MR3 = NO	O.2	Yes [ ] No [ ]
MR5	PMD_fault function	53.4.9	Sets PMD_fault to a logical 1 if any local fault is detected	MD:O	Yes [ ] No [ ] N/A [ ]
MR6	PMD_transmit_fault function	53.4.10	Sets PMD_transmit_fault to a logical 1 if a local fault is detected on any transmit lane	MD:O	Yes [ ] No [ ] N/A [ ]
MR7	PMD_receive_fault function	53.4.11	Sets PMD_receive_fault to a logical 1 if a local fault is detected on any receive lane	MD:O	Yes [ ] No [ ] N/A [ ]

53.15.4.4 Jitter specifications

Item	Feature	Subclause	Value/Comment	Status	Support
JS1	Transmit jitter	53.8.1	Meet BER “bathtub curve” specifications	M	Yes [ ]
JS2	Channel transmit jitter	53.8.1.1	As described in steps a) through c) in 53.8.1.1	M	Yes [ ]
JS3	Channel transmit jitter	53.8.1.1	Also tested on MMF as defined in 53.9.10.1	M	Yes [ ]
JS4	Receive jitter	53.8.2	BER less than 10^{-12}	M	Yes [ ]
JS5	Receive jitter	53.8.2.1	Meets requirements of the receiver input jitter mask	M	Yes [ ]
JS6	Receive jitter	53.8.2.1	Uniform spectral content over the measurement frequency range of 18.75kHz to 1.5GHz	M	Yes [ ]
JS7	Receive jitter	53.8.2.1	Using a Clock Recovery Unit	M	Yes [ ]
JS8	Receive jitter	53.8.2.1	Using a low-frequency corner of less than or equal to 1.875MHz and a slope of 20dB/decade	M	Yes [ ]
JS9	Receive jitter	53.8.2.1	Using fourth-order Bessel-Thomson filter	M	Yes [ ]
JS10	Receive jitter	53.8.2.2	Meets the requirements of Table 53–11	M	Yes [ ]
JS11	Receive jitter	53.8.2.2	Sinusoidal jitter added to the test signal in compliance with 53.8.2.1	M	Yes [ ]

53.15.4.5 Optical measurement requirements

Item	Feature	Subclause	Value/Comment	Status	Support
OM1	Length of patch cord used for measurements	53.9	2 to 5 m	M	Yes [ ]
OM2	Wavelength ranges	53.9.1	Wavelengths fall within ranges specified inTable 53–5, and under modulated conditions using valid 10GBASE-X signals	M	Yes [ ]
OM3	Optical power measurements	53.9.2	Per TIA/EIA-455-95	M	Yes [ ]
OM4	Source spectral window measurements	53.9.3	Individually measured per test setup in Figure 53–7, with all other channels below –30 dBm	M	Yes [ ]
OM5	Source spectral window measurements	53.9.3	Under modulated conditions using valid 10GBASE-X signals	M	Yes [ ]
OM6	Extinction ratio measurements	53.9.4	Per IEC 61280-2-2	M	Yes [ ]
OM7	OMA measurements	53.9.5	Each channel tested individually per methodology defined in 52.9.5	M	Yes [ ]
OM8	$RIN_{12}OMA$	53.9.6	Each channel tested individually per methodology defined in 52.9.6		Yes [ ]
OM9	Transmit eye	53.9.7	Per IEC 61280-2-2	M	Yes [ ]
OM10	Transmit eye mask measurement conditions	53.9.7	Using fourth-order Bessel-Thomson filter	M	Yes [ ]
OM11	Transmit eye mask measurement conditions	53.9.7	Using a Clock Recovery Unit to trigger the scope	M	Yes [ ]
OM12	Transmit eye mask measurement conditions	53.9.7	Using a low-frequency corner of less than or equal to 1.875MHz and a slope of 20dB/decade	M	Yes [ ]
OM13	Transmit rise/fall characteristics conditions	53.9.8	Waveforms conform to mask in Figure 53–8, measured from 20% to 80%, using a patch cord	M	Yes [ ]
OM14	Transmit rise/fall characteristics conditions	53.9.8	Removed mask conforming filter mathematically	M	Yes [ ]
OM15	Transmit rise/fall characteristics conditions	53.9.8	Mask filters use a fourth-order Bessel-Thomson filter	M	Yes [ ]
OM16	Receive sensitivity measurement conditions	53.9.9	Using conformance test at TP3 and meeting conditions specified in Table 53–8	M	Yes [ ]
OM17	Transmit jitter conformance measurement conditions	53.9.10.1	Using a fourth-order Bessel-Thomson filter for single-mode fiber	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
OM18	Transmit jitter conformance measurement conditions	53.9.10.1	Using a fourth-order Bessel-Thomson filter followed by a transversal filter with 2 equal amplitude paths with a differential delay of 157 ps for multimode fiber	M	Yes [ ]
OM19	Transmit jitter conformance measurement conditions	53.9.10.1	Using a low-frequency corner of less than or equal to 1.875MHz and a slope of 20dB/decade	M	Yes [ ]
OM20	Transmit jitter conformance measurement conditions	53.9.10.1	Measured at the average value of the overall waveform	M	Yes [ ]
OM21	Transmit jitter conformance measurement conditions	53.9.10.1	Asynchronous data flowing in all four optical receiver channels	M	Yes [ ]
OM22	Transmit jitter conformance measurement conditions	53.9.10.2	Meets requirements listed in Table 53–12	M	Yes [ ]
OM23	Transmit jitter conformance measurement conditions	53.9.10.2	For single-mode fiber; compliant with dispersion at least as negative as the “minimum dispersion” and at least as positive as the “maximum dispersion”	M	Yes [ ]
OM24	Transmit jitter conformance measurement conditions	53.9.10.2	Achieved using ITU-T G.652 fiber	M	Yes [ ]
OM25	Transmit jitter conformance measurement conditions	53.9.10.2	Using the linear regime of the single-mode fiber	M	Yes [ ]
OM26	Transmit jitter conformance measurement conditions	53.9.10.2	Provide an optical back reflection specified in Table 53–7	M	Yes [ ]
OM27	Transmit jitter conformance measurement conditions	53.9.10.2	Back reflection adjusted to create the greatest RIN	M	Yes [ ]
OM28	Transmit jitter conformance measurement conditions	53.9.10.2	For multimode fiber, back reflection set to –12 dB	M	Yes [ ]
OM29	Transmit jitter conformance measurement conditions	53.9.10.3	Using a low-frequency corner of less than or equal to 1.875MHz and a slope of 20dB/decade	M	Yes [ ]
OM30	Receiver sensitivity	53.9.11	Meet the specifications in Table 53–8	M	Yes [ ]
OM31	Stressed receiver conformance conditions	53.9.12	Asynchronous data flowing out of the optical transmitter of the system under test	M	Yes [ ]
OM32	Stressed receiver conformance conditions	53.9.12	Data is consistent with normal signal properties and content	M	Yes [ ]
OM33	Stressed receiver conformance conditions	53.9.12.1	Using a Clock Recovery Unit meeting the requirements of 53.8.2.1	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
OM34	Stressed receiver conformance conditions	53.9.12.1	Calibrated at the average value of the overall optical waveform	M	Yes [ ]
OM35	Stressed receiver conformance conditions	53.9.12.3	Using a Clock Recovery Unit meeting the requirements of 53.8.2.1	M	Yes [ ]
OM36	Receiver 3dB electrical upper cutoff frequency	53.9.13	Performed on each channel independently using a laser source with its output wavelength within the specified wavelength range of the channel to be tested	M	Yes [ ]
OM37	Receiver 3dB electrical upper cutoff frequency	53.9.13	As described in steps a) through e) of 53.9.13	M	Yes [ ]
OM38	Compliance test signal at TP3	53.9.14	Meets the requirements of Figure 53–12	M	Yes [ ]
OM39	Compliance test signal at TP3	53.9.14	DJ eye closure no less than 14 ps	M	Yes [ ]
OM40	Compliance test signal at TP3	53.9.14	Vertical eye-closure penalty meets requirements of Table 53–8	M	Yes [ ] N/A [ ]
OM41	Compliance test signal at TP3	53.9.14	Bandwidth of photodetector > 2.34GHz, and couple through fourth-order Bessel-Thomson filter	M	Yes [ ]
OM42	Receiver WDM conformance conditions	53.9.15	As described in steps a) through f) of 53.9.15	M	Yes [ ]
OM43	General safety	53.10.1	Conform to IEC-60950-1	M	Yes [ ]
OM44	Laser safety	53.10.2	Hazard Level 1	M	Yes [ ]
OM45	Compliance with all requirements over the life of the product	53.11		M	Yes [ ]
OM46	Compliance with applicable local and national codes for the limitation of electromagnetic interference	53.11.1		M	Yes [ ]

53.15.4.6 Characteristics of the fiber optic cabling

Item	Feature	Subclause	Value/Comment	Status	Support
LI1	Fiber optic cabling	53.13	Meets specifications in Table 53–13	INS:M	Yes [ ] N/A [ ]
LI2	Return loss for multimode connections	53.14.2.2	> 20 dB	INS:M	Yes [ ] No [ ] N/A [ ]
LI3	Return loss for single-mode connections	53.14.2.2	> 26 dB	INS:M	Yes [ ] No [ ] N/A [ ]
LI4	MDI	53.14.3	IEC 61753-1-1 and IEC 61753-3-2	M	Yes [ ]

54. Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-CX4

54.1 Overview

This clause specifies the 10GBASE-CX4 PMD (including MDI) and the baseband medium. In order to form a complete PHY (Physical Layer device), a PMD is combined with the appropriate sublayers (see Table 54–1) and with the management functions, which are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 54–1—Physical Layer clauses associated with the 10GBASE-CX4 PMD

Associated clause	10GBASE-CX4
46—XGMII[a]	Optional
47—XGXS and XAUI	Optional
48—10GBASE-X PCS/PMA	Required

[a]The XGMII is an optional interface. However, if the XGMII is not implemented, a conforming implementation must behave functionally as though the RS and XGMII were present.

Figure 54–1 shows the relationship of the 10GBASE-CX4 PMD sublayers and MDI to the ISO/IEC Open System Interconnection (OSI) reference model.

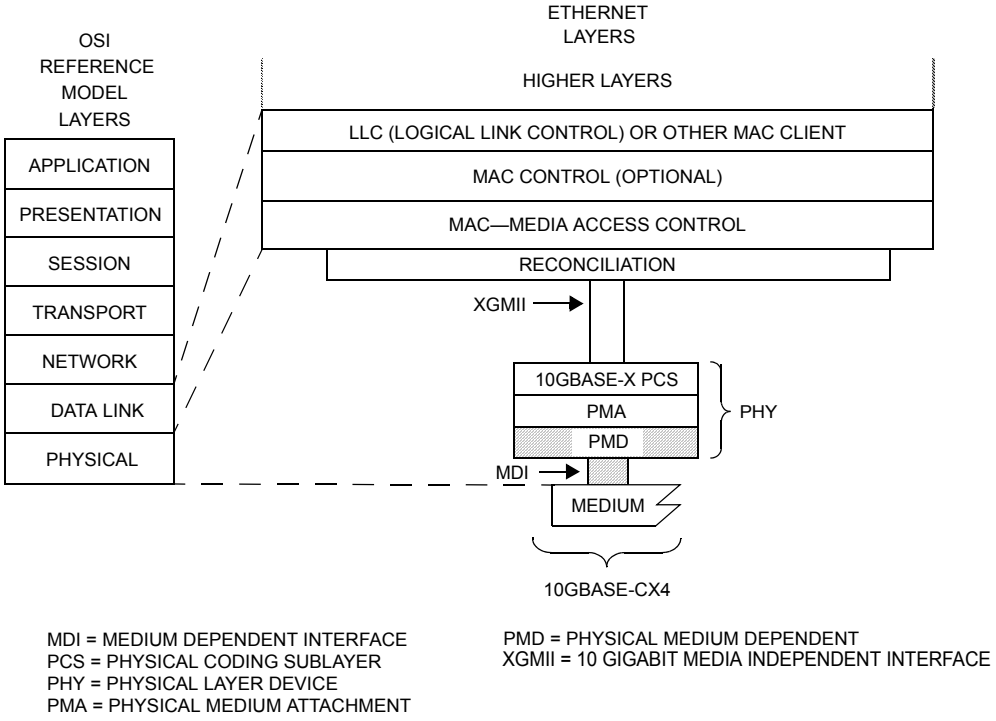

Figure 54–1—10GBASE-CX4 PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

54.2 Physical Medium Dependent (PMD) service interface

The 10GBASE-CX4 PMD utilizes the PMD service interface defined in 53.1.1. The PMD service interface is summarized as follows:

PMD_UNITDATA.request
PMD_UNITDATA.indication
PMD_SIGNAL.indication

54.3 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that implementers of MAC, MAC Control, and PHY must consider the delay maxima, and that network planners and administrators consider the delay constraints regarding the cable topology and concatenation of devices. A description of overall system delay constraints and the definitions for bit-times and pause_quanta can be found in 44.3.

The sum of the transmit and the receive delays contributed by the 10GBASE-CX4 PMD shall be no more than 512 BT or 1 pause_quantum.

54.4 PMD MDIO function mapping

The 10GBASE-CX4 PMD uses the same MDIO function mapping as 10GBASE-LX4, as defined in 53.3.

54.5 PMD functional specifications

The 10GBASE-CX4 PMD performs the transmit and receive functions (which convey data between the PMD service interface and the MDI), and provides various management functions if the optional MDIO is implemented.

54.5.1 Link block diagram

A 10GBASE-CX4 link is shown in Figure 54–2. For purposes of system conformance, the PMD sublayer is standardized at the points described in this subclause. The electrical transmit signal is defined at the output end of the mated connector (TP2). Unless specified otherwise, all transmitter measurements and tests defined in 54.6.3 are made at TP2. Unless specified otherwise, all receiver measurements and tests defined in 54.6.4 are made at the input end of the mated connector (TP3). A mated connector pair has been included in both the transmitter and receiver specifications defined in 54.6.3 and 54.6.4. Two mated connector pairs have been included in the cable assembly specifications defined in 54.7.

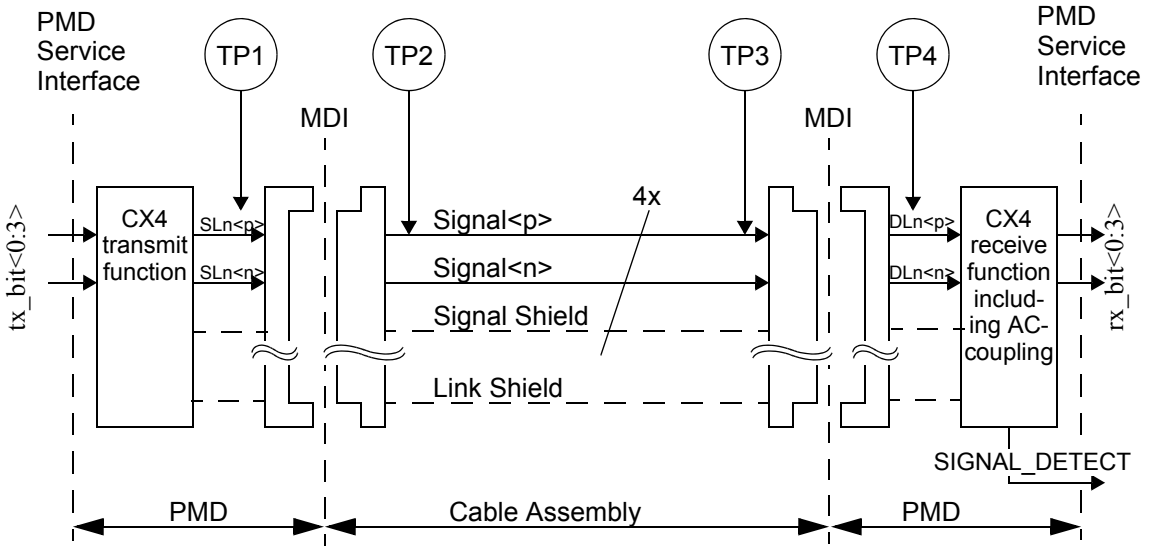

Figure 54–2—10GBASE-CX4 link (half link is shown)

NOTE—SLn<p> and SLn<n> are the positive and negative sides of the transmit differential signal pair and DLn<p> and DLn<n> are the positive and negative sides of the receive differential signal pair for lane n (n = 0, 1, 2, 3).

54.5.2 PMD Transmit function

The PMD Transmit function shall convert the four logical bit streams requested by the PMD service interface message PMD_UNITDATA.request (tx_bit<0:3>) into four separate electrical signal streams. The four electrical signal streams shall then be delivered to the MDI, all according to the transmit electrical specifications in 54.6.3. A positive output voltage of SLn<p> minus SLn<n> (differential voltage) shall correspond to tx_bit = ONE.

The PMD shall convey the bits received from the PMD service interface using the message PMD_UNITDATA.request(tx_bit<0:3>) to the MDI lanes, where (SL0<p>/<n>, SL1<p>/<n>, SL2<p>/<n>, SL3<p>/<n>) = tx_bit<0:3>.

54.5.3 PMD Receive function

The PMD Receive function shall convert the four electrical signal streams from the MDI into four logical bit streams for delivery to the PMD service interface using the message PMD_UNITDATA.indication (rx_bit<0:3>), all according to the receive electrical specifications in 54.6.4. A positive input voltage level in each signal stream of DLn<p> minus DLn<n> (differential voltage) shall correspond to a rx_bit = ONE.

The PMD shall convey the bits received from the MDI lanes to the PMD service interface using the message PMD_UNITDATA.indication(rx_bit<0:3>), where rx_bit<0:3> = (DL0<p>/<n>, DL1<p>/<n>, DL2<p>/<n>, DL3<p>/<n>).

54.5.4 Global PMD signal detect function

The Global_PMD_signal_detect function shall report the state of SIGNAL_DETECT via the PMD service interface. The SIGNAL_DETECT parameter is signaled continuously, while the PMD_SIGNAL.indication message is generated when a change in the value of SIGNAL_DETECT occurs.

SIGNAL_DETECT is a global indicator of the presence of electrical signals on all four lanes. The PMD receiver is not required to verify whether a compliant 10GBASE-CX4 signal is being received; however, it shall assert SIGNAL_DETECT = OK within 100 μs after the absolute differential peak-to-peak input voltage on each of the four lanes at the MDI has exceeded 175 mV for at least 1 UI (unit interval).

After any such assertion of SIGNAL_DETECT = OK, SIGNAL_DETECT = FAIL shall not be asserted for at least 250 μs. The PMD shall have asserted SIGNAL_DETECT = FAIL when the absolute differential peak-to-peak input voltage on any of the four lanes at the MDI has dropped below 50 mV and has remained below 50 mV for longer than 500 μs.

Table 54–2—SIGNAL_DETECT summary (informative)

Parameter	Value	Units
SIGNAL_DETECT = OK level (maximum differential peak-to-peak amplitude)	175	mV
SIGNAL_DETECT = OK width (minimum)	1	UI
SIGNAL_DETECT = OK assertion time (maximum)	100	μs
SIGNAL_DETECT = FAIL level (minimum differential peak-to-peak amplitude)	50	mV
SIGNAL_DETECT = FAIL de-assertion time maximum minimum	 500 250	 μs μs

NOTE—SIGNAL_DETECT may not activate with a continuous 1010… pattern, such as the high-frequency pattern of 48A.1, but it will be activated by an interpacket gap (IPG).

54.5.5 PMD lane-by-lane signal detect function

When the MDIO is implemented, each PMD_signal_detect_n value, where n represents the lane number in the range 0:3, shall be continuously updated in response to the amplitude of the receive signal on its associated lane, according to the requirements of 54.5.4.

54.5.6 Global PMD transmit disable function

The Global_PMD_transmit_disable function is optional. When implemented, it allows all of the transmitters to be disabled with a single variable.

a) When a Global_PMD_transmit_disable variable is set to ONE, this function shall turn off all of the transmitters such that each transmitter drives a constant level (i.e., no transitions) and does not exceed the maximum differential peak-to-peak output voltage in Table 54–3.
b) If a PMD_fault (54.5.9) is detected, then the PMD may turn off the electrical transmitter in all lanes.
c) Loopback, as defined in 54.5.8, shall not be affected by Global_PMD_transmit_disable.

54.5.7 PMD lane-by-lane transmit disable function

The PMD_transmit_disable_n function is optional. It allows the electrical transmitters in each lane to be selectively disabled.

a) When a PMD_transmit_disable_n variable is set to ONE, this function shall turn off the transmitter associated with that variable such that the corresponding transmitter drives a constant level (i.e., no

transitions) and does not exceed the maximum differential peak-to-peak output voltage in Table 54–3.

b) If a PMD_fault (54.5.9) is detected, then the PMD may turn off the electrical transmitter in all lanes.
c) Loopback, as defined in 54.5.8, shall not be affected by PMD_transmit_disable_n.

NOTE—Turning off a transmitter can be disruptive to a network.

54.5.8 Loopback mode

Loopback mode shall be provided for the 10GBASE-CX4 PMD by the transmitter and receiver of a device as a test function to the device. When loopback mode is selected, transmission requests passed to the transmitter are shunted directly to the receiver, overriding any signal detected by the receiver on its attached link. The transmitters shall not be disabled when loopback mode is enabled. A device must be explicitly placed in loopback mode because loopback mode is not the normal mode of operation of a device. Loopback applies to all lanes as a group (the lane 0 transmitter is directly connected to the lane 0 receiver, the lane 1 transmitter is directly connected to the lane 1 receiver, etc.). The method of implementing loopback mode is not defined by this standard.

Control of the loopback function is specified in 45.2.1.1.5.

NOTE 1—The signal path that is exercised in the loopback mode is implementation specific, but it is recommended that this signal path encompass as much of the circuitry as is practical. The intention of providing this loopback mode of operation is to permit diagnostic or self-test functions to test the transmit and receive data paths using actual data. Other loopback signal paths may also be enabled independently using loopback controls within other devices or sublayers.

NOTE 2—Placing a network port into loopback mode can be disruptive to a network.

54.5.9 PMD fault function

If the MDIO is implemented, and the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set PMD_fault to ONE; otherwise, the PMD shall set PMD_fault to ZERO.

54.5.10 PMD transmit fault function

If the MDIO is implemented, and the PMD has detected a local fault on any transmit lane, the PMD shall set the PMD_transmit_fault variable to ONE; otherwise, the PMD shall set PMD_transmit_fault to ZERO.

54.5.11 PMD receive fault function

If the MDIO is implemented, and the PMD has detected a local fault on any receive lane, the PMD shall set the PMD_receive_fault variable to ONE; otherwise, the PMD shall set PMD_receive_fault to ZERO.

54.6 MDI Electrical specifications for 10GBASE-CX4

54.6.1 Signal levels

The 10GBASE-CX4 MDI is a low-swing AC-coupled differential interface. Transmitter to receiver path AC-coupling, as defined in 54.6.4.3, allows for interoperability between components operating from different supply voltages. Low-swing differential signaling provides noise immunity and improved electromagnetic interference (EMI).

54.6.2 Signal paths

The 10GBASE-CX4 MDI signal paths are point-to-point connections. Each path corresponds to a 10GBASE-CX4 MDI lane and comprises two complementary signals, which form a balanced differential pair. There are four differential paths in each direction for a total of eight pairs, or sixteen connections. The signal paths are intended to operate on twinaxial cable assemblies up to 15 m in length, as described in 54.7.

54.6.3 Transmitter characteristics

Transmitter characteristics shall meet specifications at TP2, unless otherwise noted. The specifications are summarized in Table 54–3 and detailed in 54.6.3.1 through 54.6.3.9.

Table 54–3—Transmitter characteristics' summary (informative)

Parameter	Subclause reference	Value	Units
Signaling speed, per lane	54.6.3.3	3.125 ± 100 ppm	GBd
Unit interval nominal	54.6.3.3	320	ps
Differential peak-to-peak output voltage maximum minimum	54.6.3.4	 1200 800	 mV mV
Differential peak-to-peak output voltage difference (maximum)	54.6.3.4	150	mV
Common-mode voltage limits maximum minimum	54.6.3.4	 1.9 –0.4	 V V
Differential output return loss minimum	54.6.3.5	[See Equation (54–1) and Equation (54–2)]	dB
Differential output template	54.6.3.6	(See Figure 54–6 and Table 54–4)	V
Transition time maximum minimum	54.6.3.7	 130 60	 ps ps
Output jitter (peak-to-peak) Random jitter Deterministic jitter[a] Total jitter	54.6.3.8	 0.27 0.17 0.35	 UI UI UI

[a]Deterministic jitter is already incorporated into the differential output template.

54.6.3.1 Test fixtures

The test fixture of Figure 54–3, or its functional equivalent, is required for measuring the transmitter specifications described in 54.6.3.

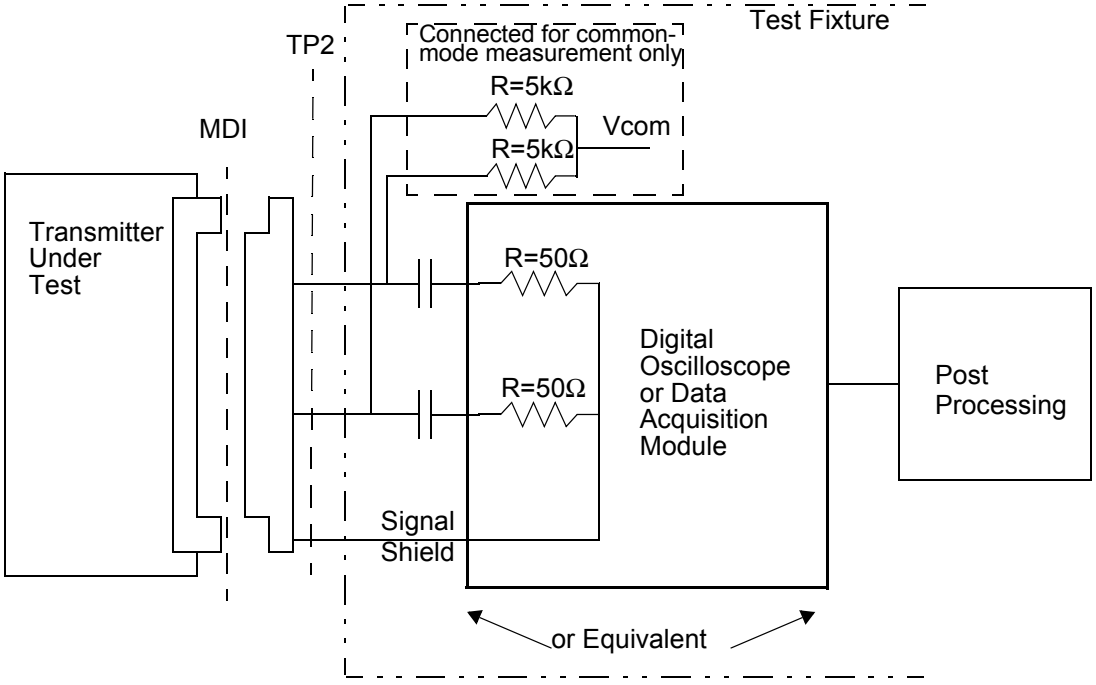

Figure 54–3—Transmit test fixture

54.6.3.2 Test-fixture impedance

The differential load impedance applied to the transmitter output by the test fixture depicted in Figure 54–3 shall have a return loss greater than 20 dB from 100 MHz to 2000 MHz. The reference impedance for differential return loss measurements shall be 100 Ω.

54.6.3.3 Signaling speed range

The 10GBASE-CX4 MDI signaling speed shall be 3.125 GBd ±100 ppm. The corresponding unit interval is nominally 320 ps.

54.6.3.4 Output amplitude

While transmitting the test pattern specified in 48A.2:

a) The transmitter maximum differential peak-to-peak output voltage shall be less than 1200 mV.
b) The minimum differential peak-to-peak output voltage shall be greater than 800 mV.
c) The maximum difference between any two lanes' differential peak-to-peak output voltage shall be less than or equal to 150 mV.

See Figure 54–4 for an illustration of the definition of differential peak-to-peak output voltage.

DC-referenced logic levels are not defined since the receiver is AC-coupled. The common-mode voltage of SLn<p> and SLn<n> shall be between –0.4 V and 1.9 V with respect to Signal Shield as measured at Vcom in Figure 54–3.

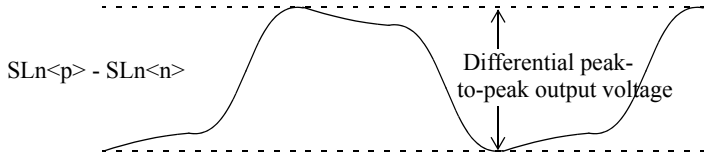

Figure 54–4—Transmitter differential peak-to-peak output voltage definition

NOTE—SLn<p> and SLn<n> are the positive and negative sides of the differential signal pair for Lane n (n = 0,1,2,3).

54.6.3.5 Output return loss

For frequencies from 100 MHz to 2000 MHz, the differential return loss, in dB with f in MHz, of the transmitter shall meet Equation (54–1) and Equation (54–2). This output impedance requirement applies to all valid output levels. The reference impedance for differential return loss measurements shall be 100 Ω.

$$Return\ Loss\ (f) \geq 10 \tag{54–1}$$

for 100 MHz ≤ f < 625 MHz and

$$Return\ Loss\ (f) \geq 10 - 10 \times \log\left(\frac{f}{625}\right) \tag{54–2}$$

for 625 MHz ≤ f ≤ 2000 MHz.

The minimum transmit differential output return loss is shown in Figure 54–5.

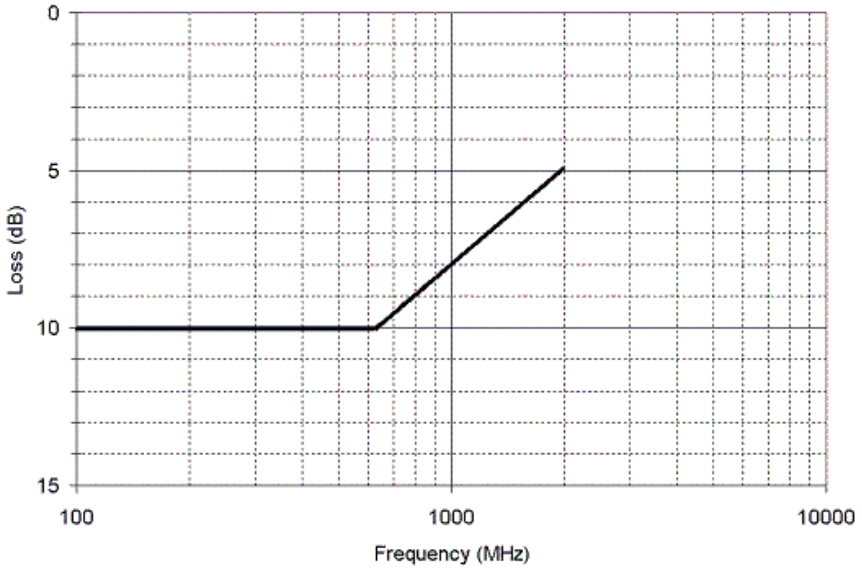

Figure 54–5—Minimum transmit differential output return loss (informative)

54.6.3.6 Differential output template

The transmitter differential output signal is defined at TP2, as shown in Figure 54–2. The transmitter shall provide equalization such that the output waveform falls within the template shown in Figure 54–6 for the test pattern specified in 48A.2, with all other transmitters active. Voltage and time coordinates for inflection points on Figure 54–6 are given in Table 54–4. The signals on each pair at TP2 shall meet the transmit template specifications when connected to the transmitter test fixture shown in Figure 54–3, with all other transmitters active. The waveform under test shall be normalized by using the following procedure:

1) Align the output waveform under test, to achieve the best fit along the horizontal time axis.
2) Calculate the +1 low frequency level as V_{lowp} = average of any 2 successive unit intervals (2UI) between 2.5 UI and 5.5 UI.
3) Calculate the 0 low frequency level as V_{lowm} = average of any 2 successive unit intervals (2UI) between 7.5 UI and 10.5 UI.
4) Calculate the vertical offset to be subtracted from the waveform as $V_{off} = (V_{lowp} + V_{lowm}) / 2$.
5) Calculate the vertical normalization factor for the waveform as $V_{norm} = (V_{lowp} – V_{lowm}) / 2$.
6) Calculate the normalized waveform as Normalized_Waveform=(Original_Waveform – V_{off}) × (0.69/V_{norm}).
7) Align the Normalized_Waveform under test, to achieve the best fit along the horizontal time axis.

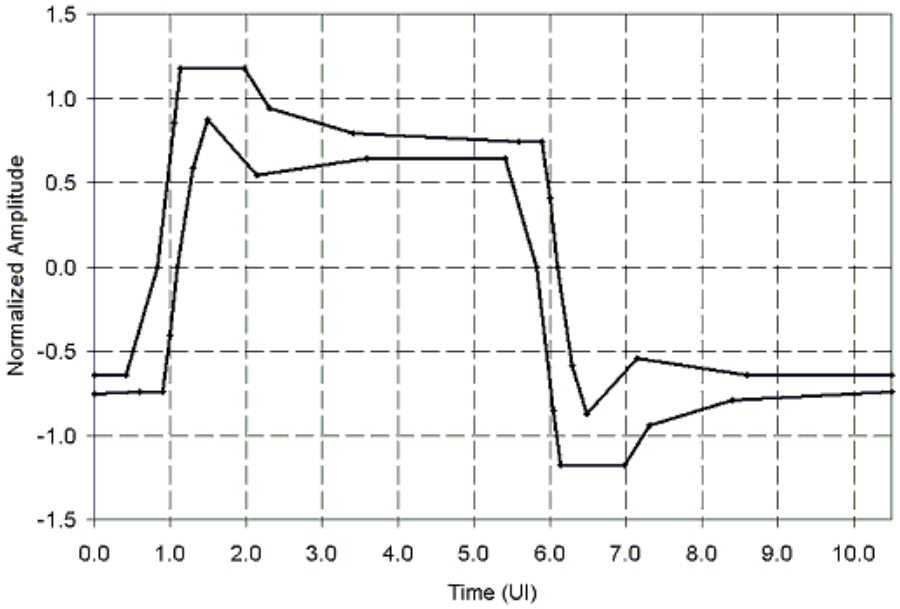

Figure 54–6—Normalized transmit template

Table 54–4—Normalized transmit time domain template

Upper limit				Lower limit			
Time (UI)	**Amplitude**	**Time (UI)**	**Amplitude**	**Time (UI)**	**Amplitude**	**Time (UI)**	**Amplitude**
0.000	–0.640	5.897	0.740	0.000	–0.754	5.409	0.640
0.409	–0.640	5.997	0.406	0.591	–0.740	5.828	0.000
0.828	0.000	6.094	0.000	0.897	–0.740	6.050	–0.856
1.050	0.856	6.294	–0.586	0.997	–0.406	6.134	–1.175
1.134	1.175	6.491	–0.870	1.094	0.000	6.975	–1.175
1.975	1.175	7.141	–0.546	1.294	0.586	7.309	–0.940
2.309	0.940	8.591	–0.640	1.491	0.870	8.500	–0.790
3.409	0.790	10.500	–0.640	2.141	0.546	10.500	–0.742
5.591	0.740	—	—	3.591	0.640	—	—

54.6.3.7 Transition time

The rising-edge transition time shall be between 60 ps and 130 ps as measured at the 20% and 80% levels of the peak-to-peak differential value of the waveform using the high-frequency test pattern of 48A.1. The falling edge transition time shall be between 60 ps and 130 ps as measured at the 80% and 20% levels of the peak-to-peak differential value of the waveform using the high-frequency test pattern of 48A.1.

54.6.3.8 Transmit jitter

The transmitter shall satisfy the jitter requirements of 54.6.3.9 with a maximum total jitter of 0.350 UI peak-to-peak, a maximum deterministic component of 0.170 UI peak-to-peak and a maximum random component of 0.270 UI peak-to-peak. Jitter specifications include all but 10^{-12} of the jitter population. Transmit jitter test requirements are specified in 54.6.3.9.

54.6.3.9 Transmit jitter test requirements

Transmit jitter is defined with respect to a test procedure resulting in a BER bathtub curve such as that described in Annex 48B. For the purpose of jitter measurement, the effect of a single-pole, high-pass filter with a 3 dB point at 1.875 MHz is applied to the jitter. The data pattern for jitter measurements shall be the CJPAT pattern defined in 48A.5. All four lanes of the 10GBASE-CX4 transceiver are active in both directions, and opposite ends of the link use asynchronous clocks. Crossing times are defined with respect to the mid-point (0 V) of the AC-coupled differential signal.

54.6.4 Receiver characteristics

Receiver characteristics are summarized in Table 54–5 and detailed in 54.6.4.1 through 54.6.4.5.

Table 54–5—Receiver characteristics' summary (informative)

Parameter	Subclause reference	Value	Units
Bit error ratio	54.6.4.1	10^{-12}	
Signaling speed, per lane	54.6.4.2	3.125 ± 100 ppm	GBd
Unit interval (UI) nominal	54.6.4.2	320	ps
Receiver coupling	54.6.4.3	AC	—
Differential input peak-to-peak amplitude (maximum)	54.6.4.4	1200	mV
Return loss[a] differential (minimum)	54.6.4.5	[See Equation (54–1) and Equation (54–2)]	dB

[a]Relative to 100 Ω differential.

54.6.4.1 Bit error ratio

The receiver shall operate with a BER of better than 10^{-12} when receiving a compliant transmit signal, as defined in 54.6.3, through a compliant cable assembly as defined in 54.7.

NOTE—The BER should be met with a worst-case insertion loss, long cable, as well as a low-loss, short cable. The low-loss cable may be a more stringent requirement on the system due to higher reflections and crosstalk than with long cables.

54.6.4.2 Signaling speed range

A 10GBASE-CX4 receiver shall comply with the requirements of 54.6.4.1 for any signaling speed in the range 3.125 GBd ± 100 ppm. The corresponding unit interval is nominally 320 ps.

54.6.4.3 AC-coupling

The 10GBASE-CX4 receiver shall be AC-coupled to the cable assembly to allow for maximum interoperability between various 10 Gb/s components. AC-coupling is considered to be part of the receiver for the purposes of this standard unless explicitly stated otherwise. It should be noted that there may be various methods for AC-coupling in actual implementations.

NOTE—It is recommended that the maximum value of the coupling capacitors be limited to 470 pF. This will limit the inrush currents to the receiver that could damage the receiver circuits when repeatedly connected to transmit modules with a higher voltage level.

54.6.4.4 Input signal amplitude

10GBASE-CX4 receivers shall accept differential input signal peak-to-peak amplitudes produced by compliant transmitters connected without attenuation to the receiver, and still meet the BER requirement specified in 54.6.4.1. Note that this may be larger than the 1200 mV differential maximum of 54.6.3.4 due to the actual transmitter output and receiver input impedances. The input impedance of a receiver can cause the minimum signal into a receiver to differ from that measured when the receiver is replaced with a 100 Ω test

load. Since the 10GBASE-CX4 receiver is AC-coupled, the absolute voltage levels with respect to the receiver ground are dependent on the receiver implementation.

54.6.4.5 Input return loss

For frequencies from 100 MHz to 2000 MHz, the differential return loss (in dB with *f* in MHz) of the receiver shall be greater than or equal to Equation (54–1) and Equation (54–2). This input impedance requirement applies to all valid input levels. The reference impedance for differential return loss measurements is 100 Ω.

54.7 Cable assembly characteristics

The 10GBASE-CX4 cable assembly contains insulated conductors terminated in a connector at each end for use as a link segment between MDIs. This cable assembly is primarily intended as a point-to-point interface of up to 15 m between network ports using controlled impedance cables. All cable assembly measurements are to be made between TP1 and TP4 as shown in Figure 54–2. These cable assembly specifications are based upon twinaxial cable characteristics, but other cable types are acceptable if the specifications are met.

Table 54–6—Cable assembly differential characteristics' summary (informative)

Description	Reference	Value	Unit
Maximum Insertion loss at 1.5625 GHz	54.7.2 and 54.7.3	16	dB
Minimum Return loss at 1.5625 GHz	54.7.3	12.0	dB
Minimum NEXT loss at 1.5625 GHz	54.7.4.1	31.8	dB
Minimum MDNEXT loss at 1.5625 GHz	54.7.4.2	29.8	dB
Minimum ELFEXT loss at 1.5625 GHz	54.7.5.1	23.3	dB
Minimum MDELFEXT loss at 1.5625 GHz	54.7.5.2	21.1	dB

54.7.1 Characteristic impedance and reference impedance

The nominal differential characteristic impedance of the cable assembly is 100 Ω. The differential reference impedance for cable assembly specifications shall be 100 Ω.

54.7.2 Cable assembly insertion loss

The insertion loss (in dB with f in MHz) of each pair of the 10GBASE-CX4 cable assembly shall be:

$$Insertion\ Loss\ (f) \leq (0.2629 \times \sqrt{f}) + (0.0034 \times f) + \left(\frac{12.76}{\sqrt{f}}\right) \quad (54\text{–}3)$$

for all frequencies from 100 MHz to 2000 MHz. This includes the attenuation of the differential cabling pairs and the assembly connectors.

The maximum cable assembly insertion loss is shown in Figure 54–7.

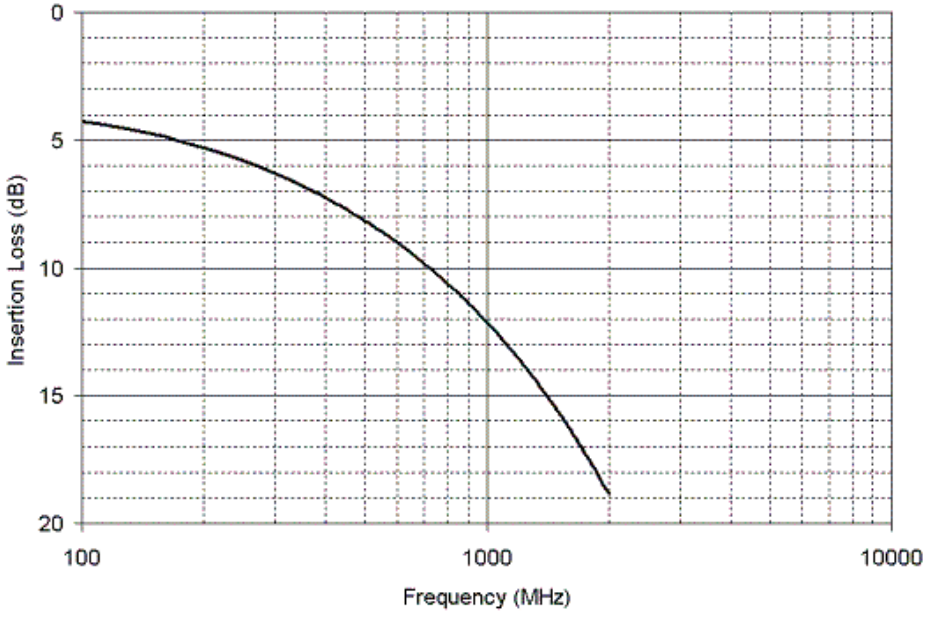

Figure 54–7—Maximum cable assembly insertion loss (informative)

54.7.3 Cable assembly return loss

The return loss (in dB with *f* in MHz) of each pair of the 10GBASE-CX4 cable assembly shall be:

$$Return\ Loss(f) \geq 22.35 - 17.19 \times \log\left(\frac{f}{100}\right) \quad (54\text{–}4)$$

for 100 MHz ≤ *f* < 400 MHz.

$$Return\ Loss(f) \geq 12 \quad (54\text{–}5)$$

for 400 MHz ≤ *f* ≤ 2000 MHz.

The minimum cable assembly return loss is shown in Figure 54–8.

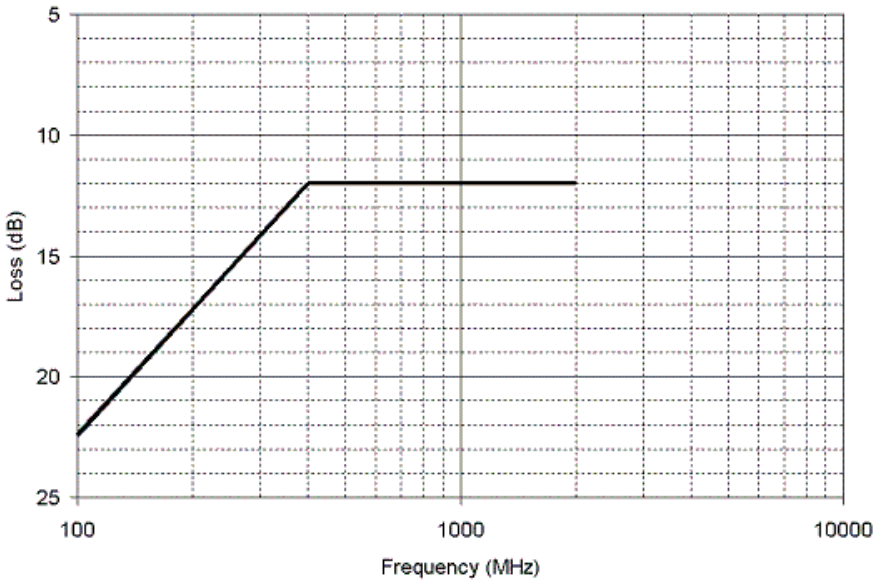

Figure 54–8—Minimum cable assembly return loss (informative)

54.7.4 Near-End Crosstalk (NEXT)

54.7.4.1 Differential Near-End Crosstalk

In order to limit the crosstalk at the near end of a link segment, the differential pair-to-pair Near-End Crosstalk (NEXT) loss between any of the four transmit lanes and any of the four receive lanes is specified to meet the BER objective specified in 54.6.4.1. The NEXT loss between any transmit and receive lane of a link segment (in dB with *f* in MHz) shall be at least:

$$NEXT(f) \geq 30 - 17 \times \log\left(\frac{f}{2000}\right) \quad (54\text{–}6)$$

for all frequencies from 100 MHz to 2000 MHz.

54.7.4.2 Multiple Disturber Near-End Crosstalk (MDNEXT)

Since four transmit and four receive lanes are used to transfer data between PMDs, the NEXT that is coupled into a receive lane will be from the four transmit lanes. To ensure the total NEXT coupled into a receive lane is limited, multiple disturber NEXT loss is specified as the power sum of the individual NEXT losses.

The Power Sum loss between a receive lane and the four transmit lanes (in dB with f in MHz) shall be at least:

$$MDNEXT(f) \geq 28 - 17 \times \log\left(\frac{f}{2000}\right) \quad (54–7)$$

for all frequencies from 100 MHz to 2000 MHz.

MDNEXT loss is determined by summing the power of the four individual pair-to-pair differential NEXT loss values over the frequency range 100 MHz to 2000 MHz as follows:

$$MDNEXT_{\text{loss}}(f) = -10 \times \log\left(\sum_{i=0}^{i=3} 10^{-NL(f)_i/10}\right) \quad (54–8)$$

where

$MDNEXT_{\text{loss}}(f)$ is the MDNEXT loss at frequency f,
$NL(f)_i$ is the power of the NEXT loss at frequency f of pair combination i, in dB,
f is frequency ranging from 100 MHz to 2000 MHz,
i is the 0, 1, 2, or 3 (pair-to-pair combination).

The minimum cable assembly NEXT / MDNEXT loss is shown in Figure 54–9.

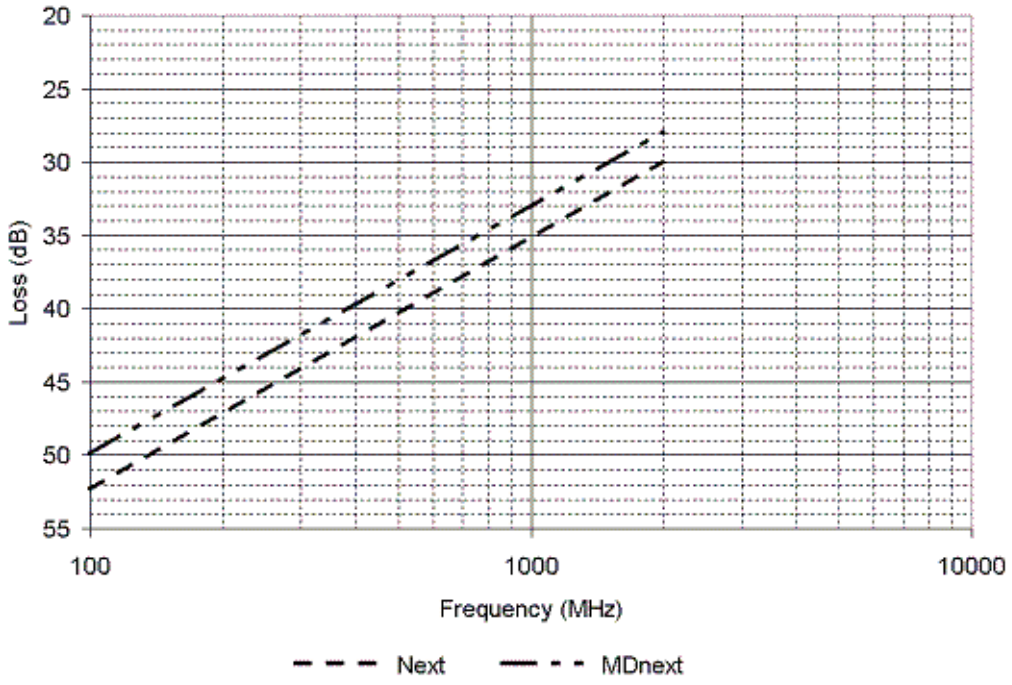

Figure 54–9—Minimum cable assembly NEXT / MDNEXT loss (informative)

54.7.5 Far-End Crosstalk (FEXT)

54.7.5.1 Equal Level Far-End Crosstalk (ELFEXT) loss

Equal Level Far-End Crosstalk (ELFEXT) loss is specified in order to limit the crosstalk at the far end of each link segment and meet the BER objective specified in 54.6.4.1. Far-End Crosstalk (FEXT) is crosstalk that appears at the far end of a lane (disturbed lane), which is coupled from another lane (disturbing lane) with the noise source (transmitters) at the near end. FEXT loss is defined as

$$FEXT\_Loss(f) = 20 \times \log(V\text{pds}(f)/V\text{pcn}(f))$$

and ELFEXT Loss is defined as

$$ELFEXT\_Loss(f) = 20 \times \log(V\text{pds}(f)/V\text{pcn}(f)) - SLS\_Loss(f)$$

where

FEXT_Loss(*f*)is the FEXT loss at frequency *f,*
ELFEXT_Loss(*f*)is the ELFEXT loss at frequency *f,*
*V*pds is the peak voltage of the disturbing signal (near-end transmitter),
*V*pcn is the peak crosstalk noise at the far end of the disturbed lane,
SLS_Loss(*f*)is the insertion loss of the disturbed lane in dB,
f is frequency ranging from 100 MHz to 2000 MHz.

The worst pair ELFEXT loss between any two lanes shall be at least:

$$ELFEXT(f) \geq 21 - 20 \times \log\left(\frac{f}{2000}\right) \quad (54\text{–}9)$$

for all frequencies from 100 MHz to 2000 MHz.

54.7.5.2 Multiple Disturber Equal Level Far-End Crosstalk (MDELFEXT) loss

Since four lanes are used to transfer data between PMDs, the FEXT that is coupled into a data carrying lane will be from the three other lanes in the same direction. To ensure the total FEXT coupled into a lane is limited, multiple disturber ELFEXT loss is specified as the power sum of the individual ELFEXT losses.

The Power Sum loss (labeled as MDELFEXT) between a lane and the three adjacent disturbers shall be at least:

$$MDELFEXT(f) \geq 19 - 20 \times \log\left(\frac{f}{2000}\right) \quad (54\text{–}10)$$

for all frequencies from 100 MHz to 2000 MHz.

MDELFEXT loss is determined by summing the power of the three individual pair-to-pair differential ELFEXT loss values over the frequency range 100 MHz to 2000 MHz as follows:

$$MDELFEXT_{\text{loss}}(f) = -10 \times \log\left(\sum_{i=0}^{i=3} 10^{-NL(f)_i/10}\right) \quad (54\text{–}11)$$

where

$MDELFEXT_{\text{loss}}(f)$ is the MDELFEXT loss at frequency f,
$NL(f)_i$ is the power of ELFEXT loss at frequency f of pair combination i, in dB,
f is frequency ranging from 100 MHz to 2000 MHz,
i is the 1, 2, or 3 (pair-to-pair combination).

The minimum cable assembly ELFEXT / MDELFEXT loss is shown in Figure 54–10.

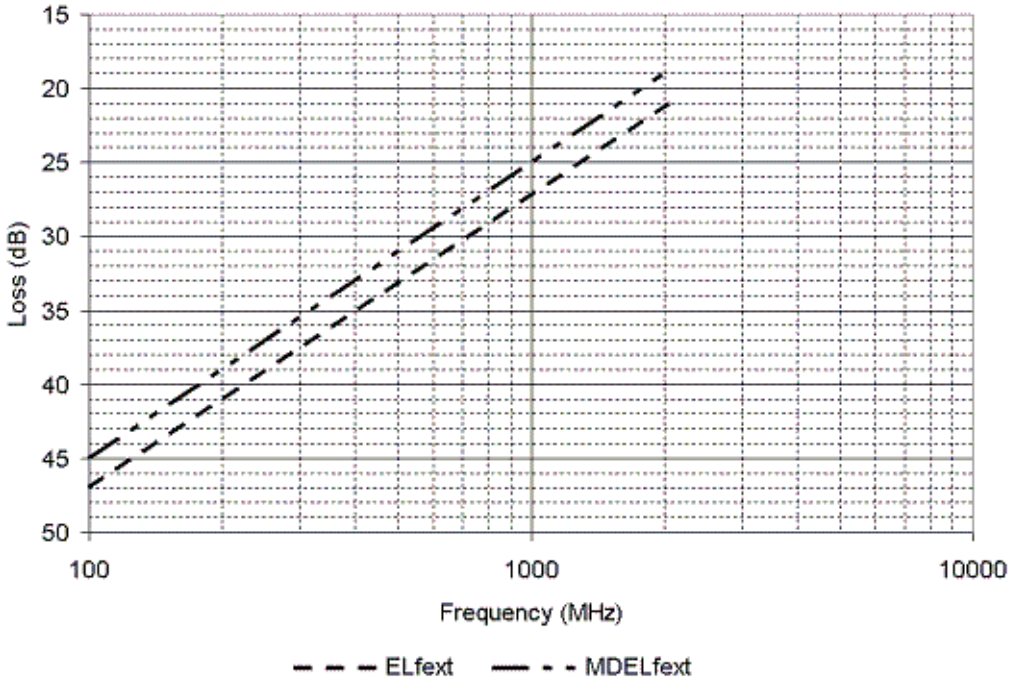

Figure 54–10—Minimum cable assembly ELFEXT / MDELFEXT loss (informative)

54.7.6 Shielding

The cable assembly shall provide Class 2 or better shielding in accordance with IEC 61196-1.

54.7.7 Crossover function

The cable assembly shall be wired in a crossover fashion as shown in Figure 54–11, with each of the four pairs being attached to the transmitter contacts at one end and the receiver contacts at the other end.

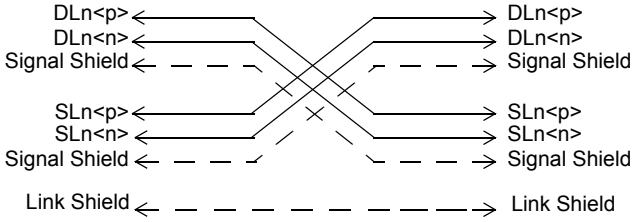

Figure 54–11—Cable wiring

NOTE—SLn<p> and SLn<n> are the positive and negative sides of the differential signal pair for Lane n (n = 0,1,2,3)

54.8 MDI specification

This subclause defines the Media Dependent Interface (MDI). The 10GBASE-CX4 PMD, as per 54.6, is coupled to the cable assembly, as per 54.7, by the MDI.

54.8.1 MDI connectors

The connector for each end of the cable assembly shall be the latch-type plug with the mechanical mating interface defined by IEC 61076-3-113 and illustrated in Figure 54–12. The MDI connector shall be the latch-type receptacle with the mechanical mating interface defined by IEC 61076-3-113 and illustrated in Figure 54–13. These connectors have a pinout matching that in Table 54–7, and electrical performance consistent with the signal quality and electrical requirements of 54.6 and 54.7.

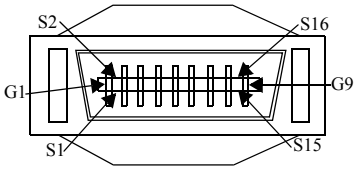

Figure 54–12—Example cable assembly plug (informative)

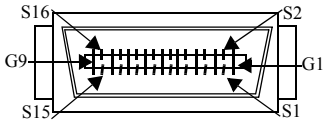

Figure 54–13—Example MDI board receptacle (informative)

54.8.2 Connector pin assignments

The MDI connector of the PMD comprises 16 signal connections, eight signal shield connections, and one link shield connection. The 10GBASE-CX4 PMD MDI connector pin assignments shall be as defined in Table 54–7.

Table 54–7—CX4 lane to MDI connector pin mapping

Rx lane	MDI Connector pin	Tx lane	MDI Connector pin
DL0<p>	S1	SL0<p>	S16
DL0<n>	S2	SL0<n>	S15
DL1<p>	S3	SL1<p>	S14
DL1<n>	S4	SL1<n>	S13
DL2<p>	S5	SL2<p>	S12
DL2<n>	S6	SL2<n>	S11
DL3<p>	S7	SL3<p>	S10
DL3<n>	S8	SL3<n>	S9
Signal Shield	G1	Signal Shield	G5
Signal Shield	G2	Signal Shield	G6
Signal Shield	G3	Signal Shield	G7
Signal Shield	G4	Signal Shield	G8
—	—	Link Shield	G9

54.9 Environmental specifications

All equipment subject to this clause shall conform to the applicable requirements of 14.7.

54.10 Protocol implementation conformance statement (PICS) proforma for Clause 54, Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-CX4[18]

54.10.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 54 Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-CX4, shall complete the following protocol implementation conformance statement (PICS) proforma. A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

54.10.2 Identification

54.10.2.1 Implementation identification

Supplier[1]	
Contact point for inquiries about the PICS[1]	
Implementation Name(s) and Version(s)[1,3]	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)[2]	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

54.10.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 54., Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-CX4
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [ ] Yes [ ] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2015.)	

Date of Statement	

[18] *Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

54.10.3 PICS proforma tables for 10GBASE-CX4 and baseband medium

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	Environmental specifications	54.9		M	Yes [ ]

54.10.4 Major capabilities/options

Item[a]	Feature	Subclause	Value/Comment	Status	Support
XGE	XGMII interface	46, 54.1	Interface is supported	O	Yes [ ] No [ ]
XGXS	XGXS and XAUI	47, 54.1		O	Yes [ ] No [ ]
PCS	Support of 10GBASE-X PCS / PMA	48, 54.1		O	Yes [ ] No [ ]
DC	Delay constraints	54.3	Delay no more than 512 BT or 1 pause_quantum	M	Yes [ ]
*MD	MDIO capability	54.4	Registers and interface supported	O	Yes [ ] No [ ]
*CBL	Cable assembly	54.7	Items marked with CBL include cable assembly specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]

[a]A "*" preceding an "Item" identifier indicates there are other PICS that depend on whether or not this item is supported.

54.10.4.1 PMD Functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
PF1	Transmit function	54.5.2	Convert the four logical signals requested by PMD_UNIT-DATA.request (tx_bit<0:3>) to 4 electrical signals	M	Yes [ ]
PF2	Delivery to the MDI	54.5.2	Supplies four electrical signal streams for delivery to the MDI	M	Yes [ ]
PF3	Mapping between logical signal and electrical signal for transmitter	54.5.2	A positive differential voltage is a one	M	Yes [ ]
PF4	Transmit Signal order	54.5.2	PMD_UNITDATA.request (tx_bit<0:3>) = (SL0<p>/<n>, SL1<p>/<n>, SL2<p>/<n>, SL3<p>/<n>)	M	Yes [ ]
PF5	Receive function	54.5.3	Convert the four electrical signals received from the MDI to 4 logical signals PMD_UNITDATA.indication(rx_bit<0:3>) per 54.6.4	M	Yes [ ]
PF6	Mapping between electrical signal and logical signal for receiver	54.5.3	A positive differential voltage is a one	M	Yes [ ]
PF7	Receive Signal order	54.5.3	(DL0<p>/<n>, DL1<p>/<n>, DL2<p>/<n>, DL3<p>/<n>) = PMD_UNITDATA.indication (rx_bit<0:3>)	M	Yes [ ]
PF8	Global PMD Signal Detect function	54.5.4	Report state via PMD_SIGNAL.indication (SIGNAL_DETECT)	M	Yes [ ]
PF9	Global PMD Signal Detect OK threshold	54.5.4	SIGNAL_DETECT = OK for signal value ≥ 175 mV for at least 1 UI on each of the four lanes	M	Yes [ ]
PF10	Global PMD Signal Detect FAIL threshold, minimum	54.5.4	SIGNAL_DETECT = FAIL not asserted for signal level < 50 mV for < 250 µs on all four lanes	M	Yes [ ]
PF11	Global PMD Signal Detect FAIL threshold, maximum	54.5.4	SIGNAL_DETECT = FAIL asserted for signal level < 50 mV for > 500 µs on any of the four lanes	M	Yes [ ]
PF12	Global_PMD_transmit_disable	54.5.6	Disables all transmitters by forcing a constant output state	O	Yes [ ] No [ ]
PF13	PMD_fault disables transmitter	54.5.6	Disables all transmitters by forcing a constant output state when a fault is detected	O	Yes [ ] No [ ]
PF14	Effect on loopback of Global_PMD_transmit_disable	54.5.6	Global_PMD_transmit_disable does not affect loopback function	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
PF15	PMD_transmit_disable_n	54.5.7	Disables transmitter n (n=0:3) by forcing a constant output state	O	Yes [ ]
PF16	PMD_fault disables transmitter n	54.5.7	Disables transmitter n (n=0:3) by forcing a constant output state when a fault is detected	O	Yes [ ] No [ ]
PF17	Effect on loopback of PMD_transmit_disable_n	54.5.7	PMD_transmit_disable_n does not affect loopback function	M	Yes [ ]
PF18	Loop Back	54.5.8	Loopback function provided	M	Yes [ ]
PF19	Transmitters on during loopback	54.5.8	Loopback function does not disable the transmitters	M	Yes [ ]

54.10.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
MF1	Lane-by-Lane Signal Detect function	54.5.5	Sets PMD_signal_detect_n values on a lane-by-lane basis per requirements of 54.5.4	MD:M	Yes [ ] N/A [ ]
MF2	PMD_fault function	54.5.9	Sets PMD_fault to a logical 1 if any local fault is detected otherwise set to 0	MD:M	Yes [ ] N/A [ ]
MF3	PMD_transmit_fault function	54.5.10	Sets PMD_transmit_fault to a logical 1 if a local fault is detected on the transmit path otherwise set to 0	MD:M	Yes [ ] N/A [ ]
MF4	PMD_receive_fault function	54.5.11	Sets PMD_receive_fault to a logical 1 if a local fault is detected on the receive path otherwise set to 0	MD:M	Yes [ ] N/A [ ]

54.10.4.3 Transmitter specifications

Item	Feature	Subclause	Value/Comment	Status	Support
DS1	Meets specifications at TP2	54.6.3		M	Yes [ ]
DS2	Test load	54.6.3.2	100 Ω differential load with return loss > 20 dB	M	Yes [ ]
DS3	Signaling speed	54.6.3.3	3.125 GBd ± 100 ppm	M	Yes [ ]
DS4	Maximum transmitter differential peak-to-peak output amplitude	54.6.3.4	Less than 1200 mV	M	Yes [ ]
DS5	Minimum transmitter differential peak-to-peak output amplitude	54.6.3.4	Greater than 800 mV	M	Yes [ ]
DS6	Maximum transmitter differential peak-to-peak amplitude difference	54.6.3.4	Less than 150 mV	M	Yes [ ]
DS7	Common-mode output voltage	54.6.3.4	Between –0.4 V and +1.9 V	M	Yes [ ]
DS8	Transmitter output return loss	54.6.3.5	Per Equation (54–1) and Equation (54–2)	M	Yes [ ]
DS9	Transmitter output return loss reference impedance	54.6.3.5	100 Ω	M	Yes [ ]
DS10	Transmitter output template test pattern	54.6.3.6	Per 48A.2	M	Yes [ ]
DS11	Transmitter output template compliance	54.6.3.6	Met while connected to test fixture shown in Figure 54–3, with all outputs active	M	Yes [ ]
DS12	Transmitter output normalization	54.6.3.6	Per process defined in 54.6.3.6	M	Yes [ ]
DS13	Transmitter output template	54.6.3.6	Lies within template of Figure 54–6 and Table 54–4	M	Yes [ ]
DS14	Rising edge transition time	54.6.3.7	Between 60 ps and 130 ps as measured per 54.6.3.7	M	Yes [ ]
DS15	Falling edge transition time	54.6.3.7	Between 60 ps and 130 ps as measured per 54.6.3.7	M	Yes [ ]
DS16	Jitter requirements	54.6.3.8	Meet BER bathtub curve, See Annex 48B	M	Yes [ ]
DS17	Transmit jitter, peak-to-peak	54.6.3.8	Meet BER bathtub curve, See Annex 48B, with Total jitter < 0.35 UI Deterministic jitter < 0.17 UI Random jitter < 0.27 UI	M	Yes [ ]
DS18	Jitter test patterns	54.6.3.9	As per Annex 48A.5	M	Yes [ ]

54.10.4.4 Receiver specifications

Item	Feature	Subclause	Value/Comment	Status	Support
RS1	Bit Error Ratio	54.6.4.1	BER of better than 10^{-12}	M	Yes [ ]
RS2	Signaling speed	54.6.4.2	3.125 GBd ± 100 ppm	M	Yes [ ]
RS3	AC-coupling	54.6.4.3	—	M	Yes [ ]
RS4	Input peak-to-peak amplitude tolerance	54.6.4.4	Accepts signals compliant with 54.6.3, may be larger than 1200 mV	M	Yes [ ]
RS5	Receiver input return loss	54.6.4.5	Per Equation (54–1) and Equation (54–2)	M	Yes [ ]

54.10.4.5 Cable assembly specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CA1	Differential reference impedance	54.7.1	100 Ω	CBL:M	Yes [ ] N/A [ ]
CA2	Insertion loss	54.7.2	Per Equation (54–3)	CBL:M	Yes [ ] N/A [ ]
CA3	Return loss	54.7.2	Per Equation (54–4), Equation (54–5), and Equation (54–6)	CBL:M	Yes [ ] N/A [ ]
CA4	NEXT	54.7.4.1	Per Equation (54–6)	CBL:M	Yes [ ] N/A [ ]
CA5	MDNEXT	54.7.4.2	Per Equation (54–7)	CBL:M	Yes [ ] N/A [ ]
CA6	ELFEXT	54.7.5.1	Per Equation (54–9)	CBL:M	Yes [ ] N/A [ ]
CA7	MDELFEXT	54.7.5.2	Per Equation (54–10)	CBL:M	Yes [ ] N/A [ ]
CA8	Shielding	54.7.6	Class 2 or better in accordance with IEC 61196-1	CBL:M	Yes [ ] N/A [ ]
CA9	Crossover function	54.7.7	Per Figure 54–11	CBL:M	Yes [ ] N/A [ ]
CA10	Cable assembly connector type	54.8.1	IEC 61076-3-113 latch-type plug	CBL:M	Yes [ ] N/A [ ]
CA11	Pin assignments	54.8.2	Per Table 54–7	CBL:M	Yes [ ] N/A [ ]

54.10.4.6 MDI connector specifications

Item	Feature	Subclause	Value/Comment	Status	Support
MDC1	MDI connector type	54.8.1	IEC 61076-3-113 latch-type receptacle	M	Yes [ ]

55. Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer and baseband medium, type 10GBASE-T

55.1 Overview

The 10GBASE-T PHY is one of the 10 Gigabit Ethernet family of high-speed network specifications. The 10GBASE-T Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA), and baseband medium specifications are intended for users who want 10 Gb/s performance over balanced twisted-pair structured cabling systems. 10GBASE-T signaling requires four pairs of balanced cabling, as specified in ISO/IEC 11801:2002 with appropriate augmentation as specified in 55.7. It is recommended that the guidelines in TIA TSB-155-A, ISO/IEC TR 24750, ANSI/TIA-568-C.2, and ISO/IEC 11801:2002/ Amendment 1 be considered before the installation of 10GBASE-T equipment for any cabling system.

This clause defines the type 10GBASE-T PCS, type 10GBASE-T PMA sublayer, and type 10GBASE-T Medium Dependent Interface (MDI). Together, the PCS and the PMA sublayer define a 10GBASE-T Physical Layer (PHY). Functional, electrical, and mechanical specifications for the type 10GBASE-T PCS, PMA, and MDI are provided in this document. This clause also specifies the baseband medium used with 10GBASE-T. Management is specified in Clause 30.

This clause also specifies 10GBASE-T Low Power Idle (LPI) as part of Energy-Efficient Ethernet (EEE). This allows the PHY to enter a low power mode of operation during periods of low link utilization as described in Clause 78. 10GBASE-T PHYs may optionally support a fast retrain mechanism.

55.1.1 Objectives

The objectives of 10GBASE-T are as follows:

a) Support full duplex operation only
b) Support star-wired local area networks using point-to-point links and structured cabling topologies
c) Support a speed of 10 Gb/s at the MAC/PLS service interface
d) Support copper medium from ISO/IEC 11801:2002, with appropriate augmentation as specified in 55.7
e) Support operation over 4-connector structured 4-pair, twisted copper cabling for all supported distances and Classes
f) Define a single 10 Gb/s PHY that would support links of up to 100 m on 4-pair balanced copper cabling as specified in 55.7
g) Preserve the IEEE 802.3/Ethernet frame format at the MAC client service interface
h) Preserve minimum and maximum frame size of the current IEEE 802.3 standard
i) Support Auto-Negotiation (Clause 28)
j) Meet CISPR/FCC Class A EMC requirements
k) Support a BER of less than or equal to 10^{-12} on all supported distances and Classes
l) Support a EEE capability as part of Energy-Efficient Ethernet (Clause 78)

55.1.2 Relationship of 10GBASE-T to other standards

Relations between the 10GBASE-T PHY, the ISO Open Systems Interconnection (OSI) reference model, and the IEEE 802.3 Ethernet model are shown in Figure 55–1. The PHY sublayers (shown shaded) in Figure 55–1 connect the IEEE 802.3 MAC to the medium.

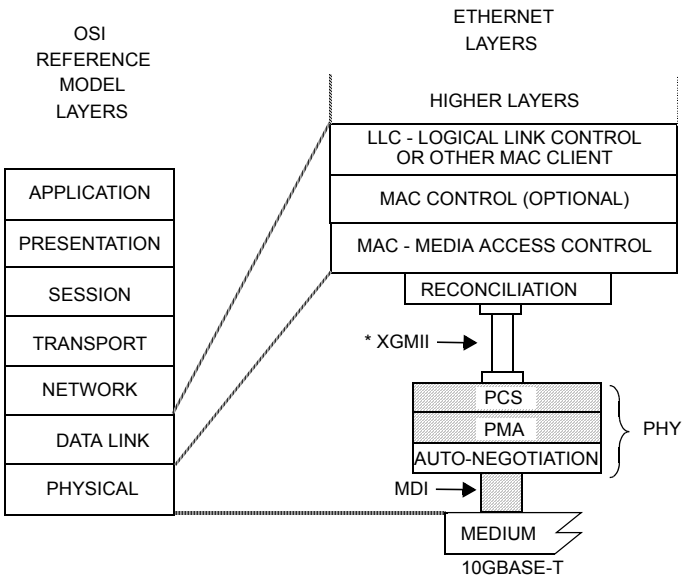

MDI = MEDIUM DEPENDENT INTERFACE
XGMII = TEN GIGABIT MEDIA INDEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PHY = PHYSICAL LAYER DEVICE

*XGMII is optional.

Figure 55–1—Type 10GBASE-T PHY relationship to the ISO Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

55.1.3 Operation of 10GBASE-T

The 10GBASE-T PHY employs full duplex baseband transmission over four pairs of balanced cabling. The aggregate data rate of 10 Gb/s is achieved by transmitting 2500 Mb/s in each direction simultaneously on each wire pair, as shown in Figure 55–2. Baseband 16-level PAM signaling with a modulation rate of 800 Megasymbol per second is used on each of the wire pairs. Ethernet data and control characters are encoded at a rate of 3.125 information bits per PAM16 symbol, along with auxiliary channel bits. Two consecutively transmitted PAM16 symbols are considered as one two-dimensional (2D) symbol. The 2D symbols are selected from a constrained constellation of 128 maximally spaced 2D symbols, called DSQ128[19] (double square 128). After link startup, PHY frames consisting of 512 DSQ128 symbols are continuously transmitted. The DSQ128 symbols are determined by 7-bit labels, each comprising 3 uncoded bits and 4 LDPC-encoded bits. The 512 DSQ128 symbols of one PHY frame are transmitted as 4×256 PAM16 symbols over the four wire pairs. Data and Control symbols are embedded in a framing scheme that runs continuously after startup of the link. The modulation symbol rate of 800 Msymbols/s results in a symbol period of 1.25 ns.

[19]The DSQ128 symbols are obtained by concatenating two time-adjacent 1D PAM16 symbols and retaining among the 256 possible Cartesian product combinations, 128 maximally spaced 2D symbols. The resulting checkerboard constellation is based on a lattice called RZ^2 in the literature (see Forney [B31]). DSQ constellations have previously been introduced under the name "AMPM" (see [B32] for examples of 8 point and 32 point AMPM/DSQ constellations).

A 10GBASE-T PHY can be configured either as a MASTER PHY or as a SLAVE PHY. The MASTER-SLAVE relationship between two stations sharing a link segment is established during Auto-Negotiation (see Clause 28, 55.6, Annex 28B, Annex 28C, and Annex 28D). The MASTER PHY uses a local clock to determine the timing of transmitter operations. The MASTER-SLAVE relationship may include loop timing. If loop timing is implemented, the SLAVE PHY recovers the clock from the received signal and uses it to determine the timing of transmitter operations, i.e., it performs loop timing, as illustrated in Figure 55–3. If loop timing is not implemented, the SLAVE PHY transmit clocking is identical to the MASTER PHY transmit clocking.

10GBASE-T PHYs optionally provide support for LPI as part of EEE (see Clause 78). This extension allows PHYs to enter an LPI mode when either the local or link partner system requests low power operation. The transmit and receive functions may enter and leave the LPI mode independently so that both symmetric and asymmetric operation is supported. While the PHY is in the LPI mode, the PHY periodically transmits a refresh signal to allow the remote PHY to refresh its receiver state (e.g., timing recovery, adaptive filter coefficients) and thereby track long-term variation in the timing of the link or the underlying channel characteristics. An easily detectable alert signal is transmitted to signal an end to the LPI mode. The alert signal is followed by a wake signal to enable a rapid transition back to the normal operational mode.

10GBASE-T PHYs may optionally support a fast retrain mechanism. This function allows PHYs to quickly recover from link degradation without a normal two second retrain.

The PCS and PMA subclauses of this document are summarized in 55.1.3.1 and 55.1.3.2. The EEE capability is summarized in 55.1.3.3. Figure 55–3 shows the functional block diagram.

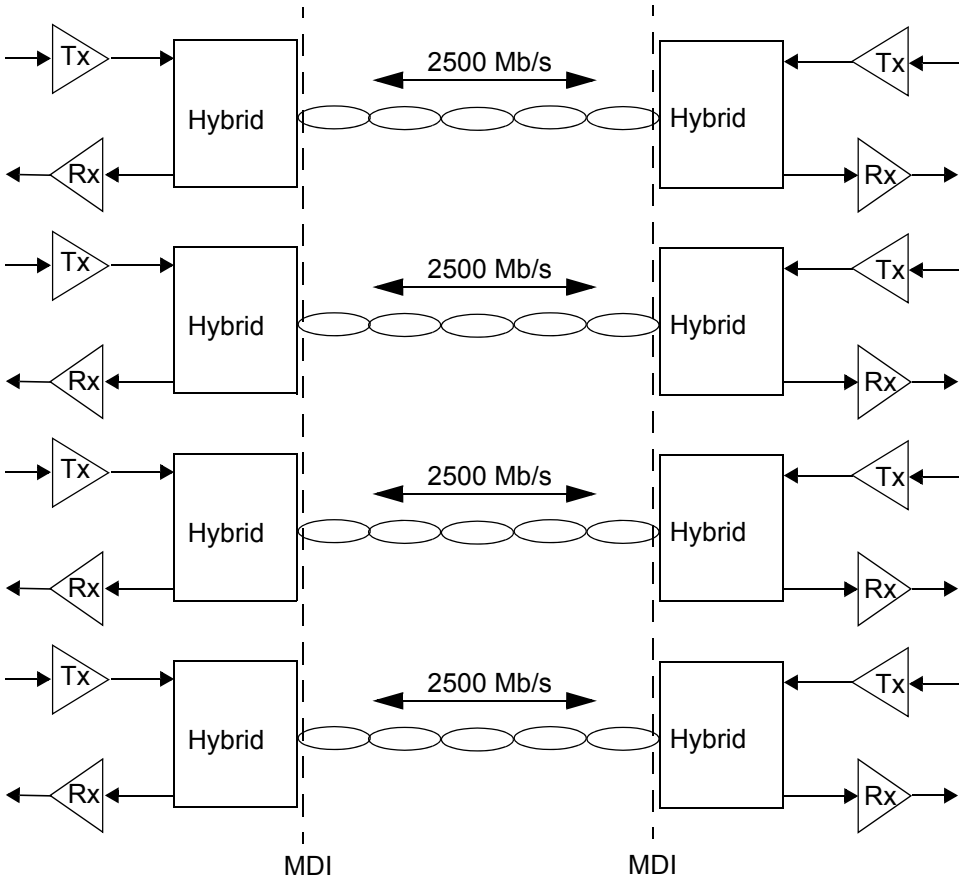

Figure 55–2—10GBASE-T topology

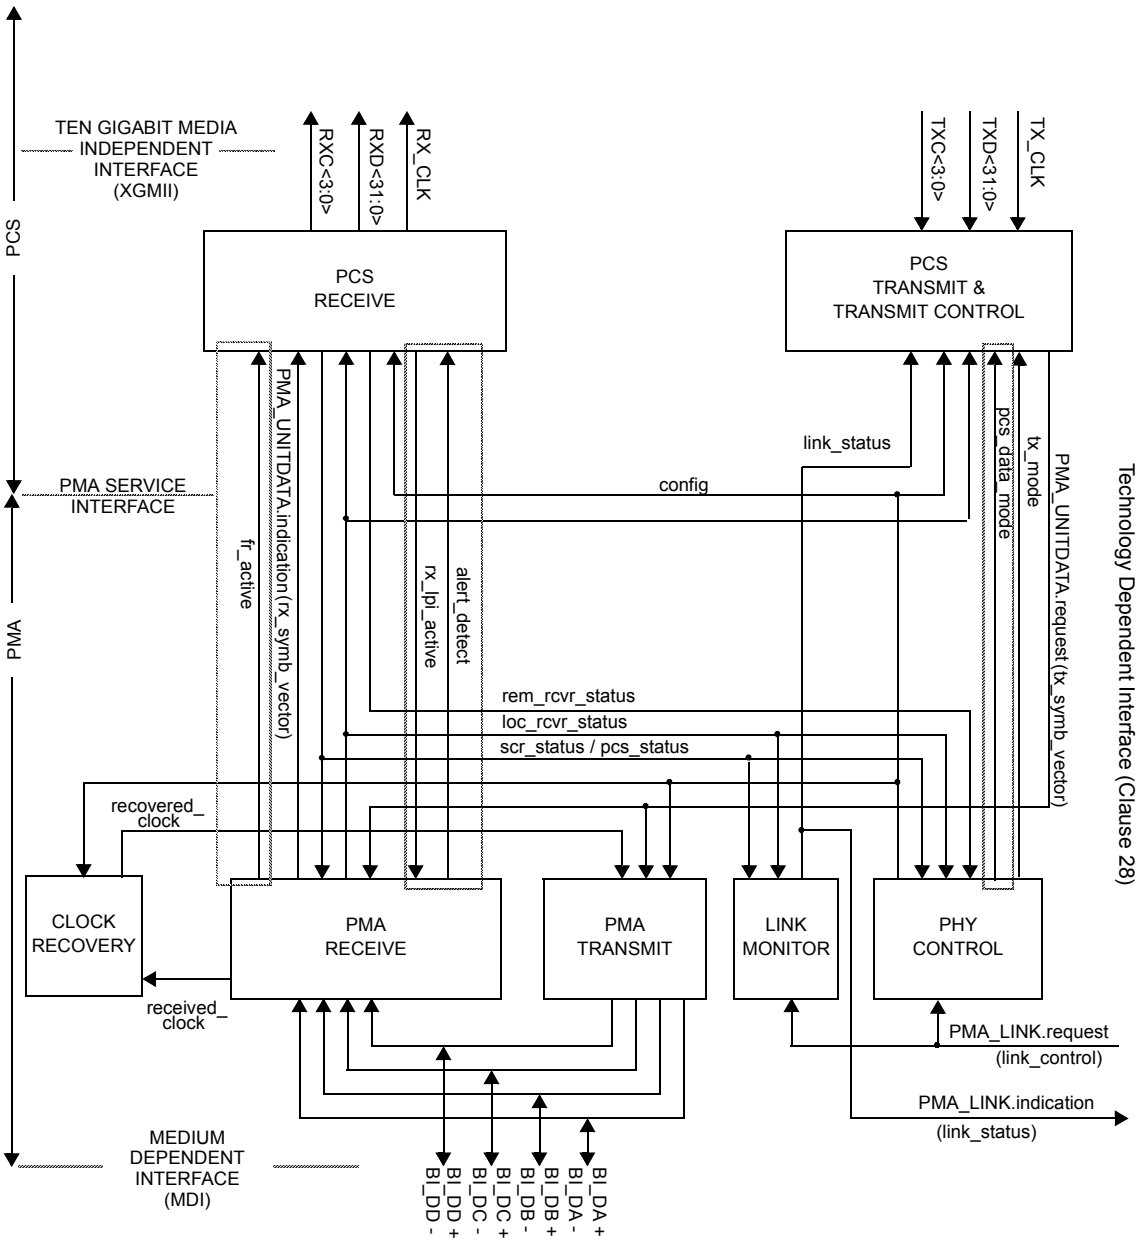

NOTE 1—The recovered_clock arc is shown to indicate delivery of the received clock signal back to PMA TRANSMIT for loop timing.

NOTE 2—pcs_data_mode is required only for the EEE or fast retrain capabilities; alert_detect and rx_lpi_active are only required for the EEE capability; fr_active is only required for the fast retrain capability.

Figure 55–3—Functional block diagram

55.1.3.1 Physical Coding Sublayer (PCS)

The 10GBASE-T PCS couples a Ten Gigabit Media Independent Interface (XGMII), as described in Clause 46, to the 10GBASE-T Physical Medium Attachment (PMA) sublayer.

In addition to the normal mode of operation, the PCS supports a training mode. Furthermore, the PCS contains a management interface.

In the transmit direction, in normal mode, the PCS receives eight XGMII data octets provided by two consecutive transfers on the XGMII service interface on TXD<31:0> and groups them into 64-bit blocks with the 64-bit block boundaries aligned with the boundary of the two XGMII transfers. Each group of eight octets along with the data/control indications is transcoded into a 65-bit block. The resulting 65-bit blocks are scrambled and assembled in a group of 50 blocks. Adding CRC8 check bits yields a CRC-checked Ethernet payload of $50 \times 65 + 8 = 3258$ bits. An auxiliary channel bit is added to obtain a block of 3259 bits.

The 3259 bits are divided into 3×512 bits and 1723 bits. The 3×512 bits, among them the auxiliary channel bit, remain uncoded. The 1723 bits are encoded by a systematic LDPC(1723,2048) encoder, which adds 325 LDPC check bits to form an LDPC codeword of 2048 coded bits. The 3×512 uncoded bits and the $2048 = 4 \times 512$ coded bits are arranged in a frame of 512 7-bit labels. Each 7-bit label comprises 3 uncoded bits and 4 coded bits.

The 512 7-bit labels are mapped into 512 2D modulation symbols selected from a DSQ128 constellation. The DSQ128 symbols are obtained by concatenating two time-adjacent 1D PAM16 symbols and retaining among the 256 possible Cartesian product combinations, 128 maximally spaced 2D symbols. The resulting checkerboard constellation is based on a lattice called RZ^2 in the literature (see Forney [B31]).

The DSQ128 constellation is partitioned into 16 subsets, each subset containing eight maximally spaced 2D symbols. The four coded bits of each 7-bit label select one DSQ128 subset, and the three uncoded bits of the label select one 2D symbol in this subset.

The obtained PHY frame of 512 DSQ128 symbols is passed on to the PMA as PMA_UNITDATA.request. The PMA transmits the DSQ128 symbols over the four wire pairs in the form of 256 constituent PAM16 symbols per pair. Details of the PCS function are covered in 55.3.

In the receive direction, in normal mode, the PCS processes code-groups received from the remote PHY via the PMA in 256 4D symbol blocks and maps them to the XGMII service interface in the receive path. In this receive processing scheme, symbol clock synchronization is done by the PMA Receive function.

The PCS functions and state diagrams are specified in 55.3. The signals provided by the PCS at the XGMII conform to the interface requirements of Clause 46. The interface to the PMA is an abstract message-passing interface specified in 55.2.

55.1.3.2 Physical Medium Attachment (PMA) sublayer

The PMA couples messages from the PCS service interface onto the balanced cabling physical medium via the Medium Dependent Interface (MDI) and provides the link management and PHY Control functions. The PMA provides full duplex communications at 800 Msymbols/s over four pairs of balanced cabling up to 100 m in length.

The PMA Transmit function comprises four transmitters to generate continuous time analog signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD, as described in 55.4.3.1. In normal mode, each four-dimensional (4D) symbol received from the PCS Transmit function undergoes multiple stages of processing. First the symbol goes through a Tomlinson-Harashima precoder (THP), which maps the PAM16 input (as described in 55.3.2.2.19) in each dimension of the four-dimensional symbol into a quasi-continuous discrete

time value in the range $-16 \le x < 16$. This THP-processed four-dimensional symbol stream may be further processed by a digital transmit filter and is then passed on to four digital to analog converters (DACs). The DAC outputs may be further processed with continuous time filters to roll off the high-frequency spectral response to limit high-frequency emissions and are then applied to each of the four balanced pairs via the MDI port.

The PMA Receive function comprises four independent receivers for pulse-amplitude modulated signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD, as described in 55.4.3.2. The receivers are responsible for acquiring symbol timing and, when operating in normal mode, for cancelling echo, near-end crosstalk, far-end crosstalk, and equalizing the signal. The 4D symbols are provided to the PCS Receive function via the PMA_UNITDATA.indication message. The PMA also contains functions for Link Monitor.

The PMA PHY Control function generates signals that control the PCS and PMA sublayer operations. PHY Control begins following the completion of Auto-Negotiation and provides the startup functions required for successful 10GBASE-T operation. It determines whether the PHY operates in a normal mode, enabling data transmission over the link segment, or whether the PHY sends special PAM2 code-groups that are used in the training mode. The latter occurs when either one or both of the PHYs that share a link segment are not operating reliably.

PMA functions and state diagrams are specified in 55.4. PMA electrical specifications are given in 55.5.

The PMA sublayer may also support a fast retrain function. The fast retrain function is specified in 55.4.2.5.15.

55.1.3.3 EEE capability

A 10GBASE-T PHY may optionally support the EEE capability, as described in 78.3. The EEE capability is a mechanism by which 10GBASE-T PHYs are able to reduce power consumption during periods of low link utilization. PHYs can enter this mode of operation after reaching PCS data mode. Each direction of the full duplex link is able to enter and exit the LPI mode independently, supporting symmetric and asymmetric LPI operation. This allows power savings when only one side of the full duplex link is in a period of low utilization. No data frames are lost or corrupted during the transition to or from the LPI mode.

In the transmit direction the transition to the LPI transmit mode begins when the PCS transmit function detects an LPI control character in all four lanes of two consecutive transfers of TXD[31:0] that will be mapped into a single 64B/65B block. Following this event a sleep signal is transmitted by the PMA. The sleep signal is composed of LDPC frames that contain only LP_IDLE 64B/65B blocks. The sleep signal indicates to the link partner that the transmit function of the PHY is entering the LPI transmit mode. Immediately after the transmission of the sleep frames, the transmit function of the local PHY enters the LPI transmit mode. While the transmit function is in the LPI mode the PHY may disable data path and control logic to save additional power. Periodically the transmit function of the local PHY transmits refresh frames that are used by the link partner to update adaptive filters and timing circuits in order to maintain link integrity. The LPI mode begins with quiet signaling or with a full refresh period. Partial refreshes (defined as a refresh signal shorter than 4 LDPC frames) that immediately follow the transition to the LPI mode are replaced with quiet signaling. The quiet-refresh cycle continues until the PCS function detects IDLE characters on the XGMII. These characters signal to the PHY that the LPI transmit mode should end. The PMA Transmit function in the PHY then sends an alert message to the link partner. The alert signal begins on a LDPC frame boundary, but has no fixed relationship to the quiet-refresh cycle. The alert signal wakes the link partner from sleep. The alert signal is followed by a wake signal, composed of LDPC frames containing only IDLE 64B/65B blocks. After a short recovery time the normal operational mode is resumed.

In the receive direction the transition to the LPI mode is triggered when the PCS Receive function detects LPI control characters within received LDPC frames. This indicates that the link partner is about to enter the LPI transmit mode. Following these frames the link partner ceases transmission and begins quiet-refresh

signaling. During the quiet time it is highly recommended that the local receiver power off circuits to reduce power consumption. Periodically the link partner transmits refresh frames that are used by the receiver to update adaptive coefficients and timing circuits. This quiet-refresh cycle continues until the link partner transmits the alert signal, initiating a transition back to the normal operational mode. The alert signal is detected in the PMA and signals that normal data frames will follow. The alert signal is followed by a wake signal that allows the local receiver time to prepare for the normal operational mode. The wake signal is composed of repeated IDLE 64B/65B blocks. After a short recovery time the normal operational mode is resumed.

Support for the EEE capability is advertised during Auto-Negotiation. Transitions to and from the LPI transmit mode are controlled via XGMII signaling. Transitions to and from the LPI receive mode are controlled by the link partner using sleep, alert, and wake signaling.

The PCS 64B/65B Transmit state diagram in Figure 55–16 and Figure 55–17 includes additional states for EEE. The PCS 64B/65B Receive state diagram in Figure 55–18 and Figure 55–19 includes additional states for EEE. The EEE Transmit state diagram is contained in the PCS Transmit function and is specified in Figure 55–20.

55.1.4 Signaling

10GBASE-T signaling is performed by the PCS generating continuous code-group sequences that the PMA transmits over each wire pair. The signaling scheme achieves a number of objectives including:

a) Forward error correction (FEC) coded symbol mapping for data.
b) Algorithmic mapping from TXD<31:0> and TXC<3:0> to four-dimensional symbols in the transmit path.
c) Algorithmic mapping from the received four-dimensional signals on the MDI port to RXD<31:0> and RXC<3:0> on the XGMII interface.
d) Uncorrelated symbols in the transmitted symbol stream.
e) No correlation between symbol streams traveling both directions on any pair combination.
f) No correlation between symbol streams on pairs BI_DA, BI_DB, BI_DC, and BI_DD.
g) Block framing and other control signals.
h) Ability to signal the status of the local receiver to the remote PHY to indicate that the local receiver is not operating reliably and requires retraining.
i) Ability to automatically detect and correct for pair swapping and crossover connections.
j) Ability to automatically detect and correct for incorrect polarity in the connections.
k) Ability to automatically correct for differential delay variations across the wire-pairs.
l) Ability to support refresh, quiet and alert signaling during LPI operation.

The PHY operates in two modes—normal mode or training mode. In normal mode, PCS generates a continuous stream of four-dimensional symbols that are transmitted via the PMA at one of eight power levels. In training mode, the PCS is directed to generate only PAM2 symbols for transmission by the PMA, which enable the receiver at the other end to train until it is ready to operate in normal mode. (See the PCS reference diagram in 55.2.)

PHYs may also support the EEE capability as described in 55.1.3.3. Transitions to the LPI mode are supported after reaching normal mode.

55.1.5 Interfaces

All 10GBASE-T PHY implementations are compatible at the MDI and at a physically exposed XGMII, if made available. Physical implementation of the XGMII is optional. Designers are free to implement circuitry within the PCS and PMA in an application-dependent manner provided that the MDI and XGMII

(if the XGMII is implemented) specifications are met. System operation from the perspective of signals at the MDI and management objects are identical whether the XGMII is implemented or not.

55.1.6 Conventions in this clause

The body of this clause contains state diagrams, including definitions of variables, constants, and functions. Should there be a discrepancy between a state diagram and descriptive text, the state diagram prevails.

The notation used in the state diagrams follows the conventions of 21.5.

Default initializations, unless specifically specified, are left to the implementer.

55.2 10GBASE-T service primitives and interfaces

10GBASE-T transfers data and control information across the following four service interfaces:

a) Ten Gigabit Media Independent Interface (XGMII)
b) Technology Dependent Interface
c) PMA service interface
d) Medium dependent interface (MDI)

The XGMII is specified in Clause 46; the Technology Dependent Interface is specified in Clause 28. The PMA service interface is defined in 55.2.2 and the MDI is defined in 55.8.

55.2.1 Technology Dependent Interface

10GBASE-T uses the following service primitives to exchange status indications and control signals across the Technology Dependent Interface as specified in Clause 28:

PMA_LINK.request (link_control)

PMA_LINK.indication (link_status)

55.2.1.1 PMA_LINK.request

This primitive allows the Auto-Negotiation algorithm to enable and disable operation of the PMA as specified in 28.2.6.2.

55.2.1.1.1 Semantics of the primitive

PMA_LINK.request (link_control)

The link_control parameter can take on one of three values: SCAN_FOR_CARRIER, DISABLE, or ENABLE.

SCAN_FOR_CARRIER	Used by the Auto-Negotiation algorithm prior to receiving any fast link pulses. During this mode the PMA reports link_status=FAIL. PHY processes are disabled.
DISABLE	Set by the Auto-Negotiation algorithm in the event fast link pulses are detected. PHY processes are disabled. This allows the Auto-Negotiation algorithm to determine how to configure the link.

ENABLE	Used by Auto-Negotiation to turn control over to the PHY for data processing functions.

55.2.1.1.2 When generated

Auto-Negotiation generates this primitive to indicate a change in link_control as described in Clause 28.

55.2.1.1.3 Effect of receipt

This primitive affects operation of the PMA Link Monitor function as defined in 55.4.2.6.

55.2.1.2 PMA_LINK.indication

This primitive is generated by the PMA to indicate the status of the underlying medium as specified in 28.2.6.1. This primitive informs the PCS, PMA PHY Control function, and the Auto-Negotiation algorithm about the status of the underlying link.

55.2.1.2.1 Semantics of the primitive

PMA_LINK.indication (link_status)

The link_status parameter can take on one of three values: FAIL, READY, or OK.

FAIL	No valid link established.
READY	For 10GBASE-T link_status does not take the value READY.
OK	The Link Monitor function indicates that a valid 10GBASE-T link is established. Reliable reception of signals transmitted from the remote PHY is possible.

55.2.1.2.2 When generated

The PMA generates this primitive to indicate a change in link_status in compliance with the state diagram given in Figure 55–31.

55.2.1.2.3 Effect of receipt

The effect of receipt of this primitive is specified in 55.3.6.2.

55.2.2 PMA service interface

10GBASE-T uses the following service primitives to exchange symbol vectors, status indications, and control signals across the service interfaces:

PMA_TXMODE.indication (tx_mode)

PMA_CONFIG.indication (config)

PMA_UNITDATA.request (tx_symb_vector)

PMA_UNITDATA.indication (rx_symb_vector)

PMA_SCRSTATUS.request (scr_status)

PMA_PCSSTATUS.request (pcs_status)

PMA_RXSTATUS.indication (loc_rcvr_status)

PMA_REMRXSTATUS.request (rem_rcvr_status)

EEE-capable PHYs additionally support the following service primitives:

PMA_ALERTDETECT.indication (alert_detect)

PCS_RX_LPI_STATUS.request (rx_lpi_active)

PMA_PCSDATAMODE.indication (pcs_data_mode)

Fast retrain capable PHYs additionally support the following service primitive:

PMA_FR_ACTIVE.indication (fr_active)

The use of these primitives is illustrated in Figure 55–4. Connections from the management interface (signals MDC and MDIO) to the sublayers are pervasive and are not shown in Figure 55–4.

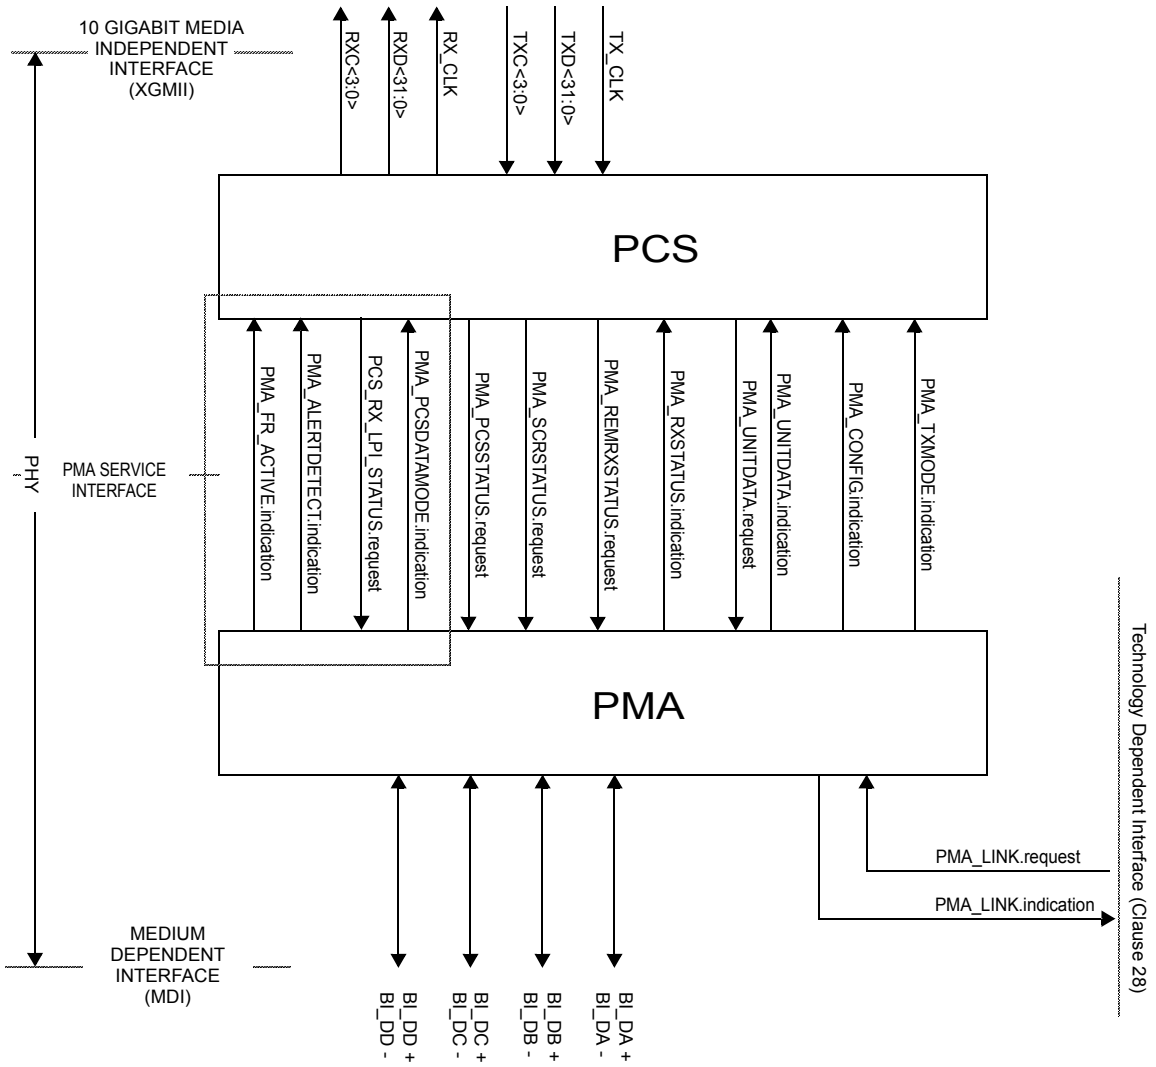

NOTE—
PMA_PCSDATAMODE.indication is required only for the EEE or fast retrain capabilities.
PMA_ALERTDETECT.indication and PCS_RX_LPI_STATUS.request are only required for the EEE capability.
PMA_FR_ACTIVE.indication is only required for the fast retrain capability.

Figure 55–4—10GBASE-T service interfaces

55.2.2.1 PMA_TXMODE.indication

The transmitter in a 10GBASE-T link normally sends over the four pairs, four-dimensional symbols that represent an XGMII data stream with framing, scrambling and encoding of data, control information, or idles.

55.2.2.1.1 Semantics of the primitive

PMA_TXMODE.indication (tx_mode)

PMA_TXMODE.indication specifies to PCS Transmit via the parameter tx_mode what sequence of code-groups the PCS should be transmitting. The parameter tx_mode can take on one of the following three values of the form:

SEND_N	This value is continuously asserted when transmission of sequences of four-dimensional symbols representing an XGMII data stream in normal mode.
SEND_T	This value is continuously asserted in case transmission of sequences of code-groups representing the training mode is to take place.
SEND_Z	This value is continuously asserted in case transmission of zeros is required.

55.2.2.1.2 When generated

The PMA PHY Control function generates PMA_TXMODE.indication messages to indicate a change in tx_mode.

55.2.2.1.3 Effect of receipt

Upon receipt of this primitive, the PCS performs its transmit function as described in 55.3.2.2.

55.2.2.2 PMA_CONFIG.indication

Each PHY in a 10GBASE-T link is capable of operating as a MASTER PHY and as a SLAVE PHY. MASTER-SLAVE configuration is determined during Auto-Negotiation (55.6.1). The result of this negotiation is provided to the PMA.

55.2.2.2.1 Semantics of the primitive

PMA_CONFIG.indication (config)

PMA_CONFIG.indication specifies to PCS and PMA Transmit via the parameter config whether the PHY operates as a MASTER PHY or as a SLAVE PHY. The parameter config can take on one of the following two values of the form:

MASTER	This value is continuously asserted when the PHY operates as a MASTER PHY.
SLAVE	This value is continuously asserted when the PHY operates as a SLAVE PHY.

55.2.2.2.2 When generated

PMA generates PMA_CONFIG.indication messages to indicate a change in config.

55.2.2.2.3 Effect of receipt

PCS and PMA Clock Recovery perform their functions in MASTER or SLAVE configuration according to the value assumed by the parameter config.

55.2.2.3 PMA_UNITDATA.request

This primitive defines the transfer of code-groups in the form of the tx_symb_vector parameter from the PCS to the PMA. The code-groups are obtained in the PCS Transmit function using the encoding rules defined in 55.3.2.2 to represent XGMII data and control streams or other sequences.

55.2.2.3.1 Semantics of the primitive

PMA_UNITDATA.request (tx_symb_vector)

During transmission, the PMA_UNITDATA.request simultaneously conveys to the PMA via the parameter tx_symb_vector the value of the symbols to be sent over each of the four transmit pairs BI_DA, BI_DB, BI_DC, and BI_DD. For EEE-capable PHYs, the vector also requests the PMA to send the ALERT signal during LPI. The tx_symb_vector parameter takes on the form:

SYMB_4D	A vector of four multi-level symbols, one for each of the four transmit pairs BI_DA, BI_DB, BI_DC, and BI_DD. In normal operation, each symbol may take on one of the values in the set {–15, –13, –11, –9, –7, –5, –3, –1, 1, 3, 5, 7, 9, 11, 13, 15}. The symbols may additionally take the value 0 when zeros are to be transmitted in the following two cases: 1) when PMA_TXMODE.indication is SEND_Z during PMA training, and 2) after data mode is reached, the transmit function is in the LPI transmit mode and lpi_tx_mode is QUIET
ALERT	A vector used to indicate that the PMA should transmit the alert sequence. ALERT will be asserted for a time equal to 4 LDPC frames.

The symbols that are elements of tx_symb_vector are called, according to the pair on which each is transmitted, tx_symb_vector[BI_DA], tx_symb_vector[BI_DB], tx_symb_vector[BI_DC], and tx_symb_vector[BI_DD].

55.2.2.3.2 When generated

The PCS generates PMA_UNITDATA.request (SYMB_4D) synchronously with every transmit clock cycle.

55.2.2.3.3 Effect of receipt

Upon receipt of this primitive the PMA transmits on the MDI the signals corresponding to the indicated symbols after processing with the THP, the transmit filter and other specified PMA Transmit processing. The parameter tx_symb_vector is also used by the PMA Receive function to process the signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DD for cancelling the echo and near-end crosstalk (NEXT).

55.2.2.4 PMA_UNITDATA.indication

This primitive defines the transfer of code-groups in the form of the rx_symb_vector parameter from the PMA to the PCS.

55.2.2.4.1 Semantics of the primitive

PMA_UNITDATA.indication (rx_symb_vector)

During reception the PMA_UNITDATA.indication simultaneously conveys to the PCS via the parameter rx_symb_vector the values of the symbols detected on each of the four receive pairs BI_DA, BI_DB, BI_DC, and BI_DD. The rx_symb_vector parameter takes on the form:

SYMB_4D	A vector of the four 1D symbols that is the receiver's best estimate of the symbols that were sent by the remote transmitter across the four pairs with reliability measures.

55.2.2.4.2 When generated

The PMA generates PMA_UNITDATA.indication (SYMB_4D) messages synchronously every four symbols received at the MDI. The nominal rate of the PMA_UNITDATA.indication primitive is 800 MHz, as governed by the recovered clock.

55.2.2.4.3 Effect of receipt

The effect of receipt of this primitive is unspecified.

55.2.2.5 PMA_SCRSTATUS.request

This primitive is generated by PCS Receive to communicate the status of the descrambler for the local PHY. The parameter scr_status conveys to the PMA Receive function the information that the training mode descrambler has achieved synchronization.

55.2.2.5.1 Semantics of the primitive

PMA_SCRSTATUS.request (scr_status)

The scr_status parameter can take on one of two values of the form:

OK	The training mode descrambler has achieved synchronization.
NOT_OK	The training mode descrambler is not synchronized.

55.2.2.5.2 When generated

PCS Receive generates PMA_SCRSTATUS.request messages to indicate a change in scr_status.

55.2.2.5.3 Effect of receipt

The effect of receipt of this primitive is specified in 55.4.2.4, 55.4.2.5, and 55.4.6.1.

55.2.2.6 PMA_PCSSTATUS.request

This primitive is generated by PCS Receive to indicate the fully operational state of the PCS for the local PHY. The parameter pcs_status conveys to the PMA Receive function the information that the PCS is operating reliably in data mode.

55.2.2.6.1 Semantics of the primitive

PMA_PCSSTATUS.request (pcs_status)

The pcs_status parameter can take on one of two values of the form:

OK	The PCS is operating reliably in data mode.
NOT_OK	The PCS is not operating reliably in data mode.

55.2.2.6.2 When generated

PCS Receive generates PMA_PCSSTATUS.request messages to indicate a change in pcs_status.

55.2.2.6.3 Effect of receipt

The effect of receipt of this primitive is specified in 55.4.6.

55.2.2.7 PMA_RXSTATUS.indication

This primitive is generated by PMA Receive to indicate the status of the receive link at the local PHY. The parameter loc_rcvr_status conveys to the PCS Transmit, PCS Receive, PMA PHY Control function, and Link Monitor the information on whether the status of the overall receive link is satisfactory or not. Note that loc_rcvr_status is used by the PCS Receive decoding functions. The criterion for setting the parameter loc_rcvr_status is left to the implementer. It can be based, for example, on observing the mean-square error at the decision point of the receiver and detecting errors during reception of symbol streams.

55.2.2.7.1 Semantics of the primitive

PMA_RXSTATUS.indication (loc_rcvr_status)

The loc_rcvr_status parameter can take on one of two values of the form:

OK	This value is asserted and remains true during reliable operation of the receive link for the local PHY.
NOT_OK	This value is asserted whenever operation of the link for the local PHY is unreliable.

55.2.2.7.2 When generated

PMA Receive generates PMA_RXSTATUS.indication messages to indicate a change in loc_rcvr_status on the basis of signals received at the MDI.

55.2.2.7.3 Effect of receipt

The effect of receipt of this primitive is specified in Figure 55–28 and in subclauses 55.2 and 55.4.6.3.

55.2.2.8 PMA_REMRXSTATUS.request

This primitive is generated by PCS Receive to indicate the status of the receive link at the remote PHY as communicated by the remote PHY via its encoding of its loc_rcvr_status parameter. The parameter rem_rcvr_status conveys to the PMA PHY Control function the information on whether reliable operation of the remote PHY is detected or not. The criterion for setting the parameter rem_rcvr_status is left to the implementer. It can be based, for example, on asserting rem_rcvr_status is NOT_OK until loc_rcvr_status is OK and then asserting the detected value of rem_rcvr_status after proper PCS Receive decoding is achieved.

55.2.2.8.1 Semantics of the primitive

PMA_REMRXSTATUS.request (rem_rcvr_status)

The rem_rcvr_status parameter can take on one of two values of the form:

OK	The receive link for the remote PHY is operating reliably.
NOT_OK	Reliable operation of the receive link for the remote PHY is not detected.

55.2.2.8.2 When generated

The PCS generates PMA_REMRXSTATUS.request messages to indicate a change in rem_rcvr_status on the basis of signals received at the MDI.

55.2.2.8.3 Effect of receipt

The effect of receipt of this primitive is specified in Figure 55–28.

55.2.2.9 PMA_ALERTDETECT.indication

This primitive is generated by PMA Receive to indicate the status of the receive link at the local PHY when rx_lpi_active is TRUE. The parameter alert_detect conveys to the PCS receive function information regarding the detection of the LPI alert signal by the PMA receive function. The criterion for setting the parameter alert_detect is left to the implementer.

55.2.2.9.1 Semantics of the primitive

PMA_ALERTDETECT.indication (alert_detect)

The alert_detect parameter can take on one of two values of the form:

TRUE	The alert signal has been reliably detected at the local receiver.
FALSE	The alert signal at the local receiver has not been detected.

55.2.2.9.2 When generated

The PMA generates PMA_ALERTDETECT.indication messages to indicate a change in the alert_detect status.

55.2.2.9.3 Effect of receipt

The effect of receipt of this primitive is specified in 55.3.2.3, Figure 55–18, and Figure 55–19.

55.2.2.10 PCS_RX_LPI_STATUS.request

When the PHY supports the EEE capability this primitive is generated by the PCS receive function to indicate the status of the receive link at the local PHY. The parameter PCS_RX_LPI_STATUS.request conveys to the PCS transmit and PMA receive functions information regarding whether the receive function is in the LPI receive mode. The parameter is generated by the Receive 64B/65B state diagram in Figure 55–18.

55.2.2.10.1 Semantics of the primitive

PCS_RX_LPI_STATUS.request (rx_lpi_active)

The rx_lpi_active parameter can take on one of two values of the form:

TRUE	The receive function is in the LPI receive mode.
FALSE	The receive function is not in the LPI receive mode.

55.2.2.10.2 When generated

The PCS generates PCS_RX_LPI_STATUS.request messages to indicate a change in the rx_lpi_active variable as determined by the receive state diagram in Figure 55–18.

55.2.2.10.3 Effect of receipt

The effect of receipt of this primitive is specified in 55.3.2.3 and Figure 55–32.

55.2.2.11 PMA_PCSDATAMODE.indication

This primitive indicates whether or not the PCS state diagrams are able to transition from their initialization states. The pcs_data_mode variable is generated by the PMA PHY Control function. It is passed to the PCS Control function via the PMA_PCSDATAMODE.indication primitive.

55.2.2.11.1 Semantics of the primitive

PMA_PCSDATAMODE.indication (pcs_data_mode)

55.2.2.11.2 When generated

The PMA PHY Control function generates PMA_PCSDATAMODE.indication messages continuously.

55.2.2.11.3 Effect of receipt

Upon receipt of this primitive, the PCS performs its transmit function as described in 55.3.2.2.

55.2.2.12 PMA_FR_ACTIVE.indication

This primitive indicates whether or not the PMA is currently performing a fast retrain. The fr_active variable is generated by the PMA PHY Control function. It is passed to the PCS Receive Control function via the PMA_FR_ACTIVE.indication primitive. This primitive is only supported by PHYs with the fast retrain capability.

55.2.2.12.1 Semantics of the primitive

PMA_FR_ACTIVE.indication (fr_active)

55.2.2.12.2 When generated

The PMA PHY Control function generates PMA_FR_ACTIVE.indication messages continuously.

55.2.2.12.3 Effect of receipt

The effect of receipt of this primitive is specified in Figure 55–18.

55.3 Physical Coding Sublayer (PCS)

55.3.1 PCS service interface (XGMII)

The PCS service interface allows the 10GBASE-T PCS to transfer information to and from a PCS client. The PCS Interface is precisely defined as the 10 Gigabit Media Independent Interface (XGMII) in Clause 46.

55.3.2 PCS functions

The PCS comprises one PCS Reset function and two simultaneous and asynchronous operating functions. The PCS operating functions are: PCS Transmit and PCS Receive. All operating functions start immediately after the successful completion of the PCS Reset function.

The PCS reference diagram, Figure 55–5, shows how the two operating functions relate to the messages of the PCS-PMA interface. Connections from the management interface (signals MDC and MDIO) to other layers are pervasive and are not shown in Figure 55–5.

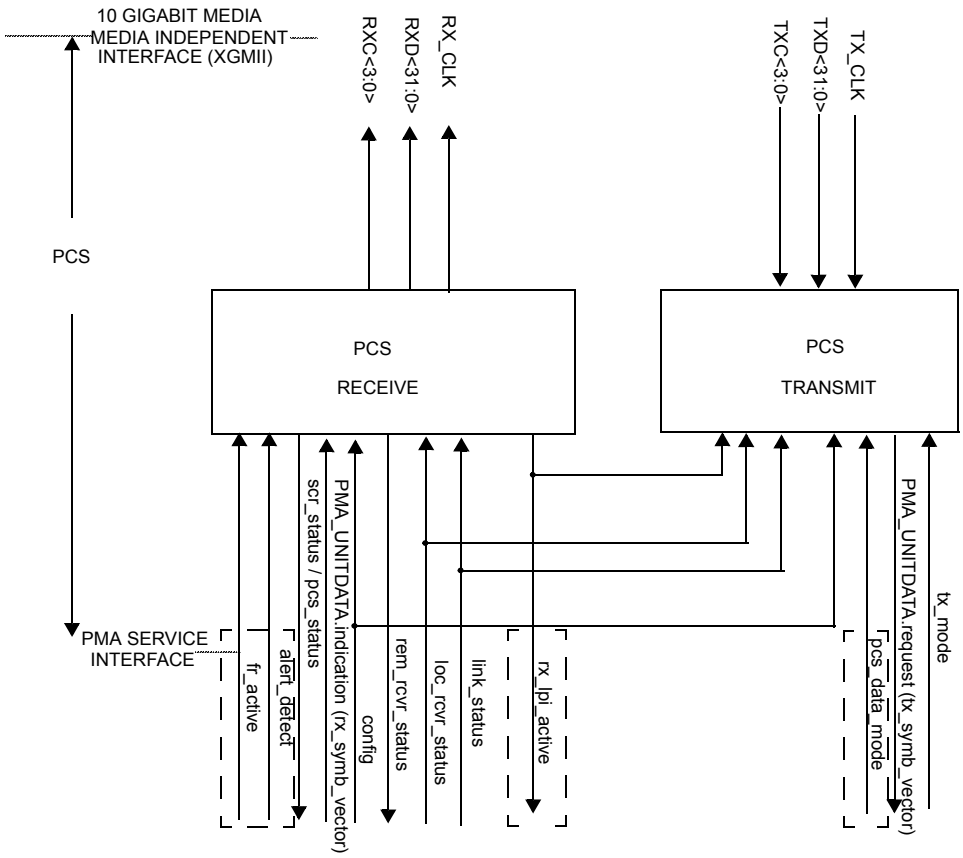

NOTE—
pcs_data_mode is required only for the EEE or fast retrain capabilities.
alert_detect and rx_lpi_active are only required for the EEE capability.
fr_active is only required for the fast retrain capability.

Figure 55–5—PCS reference diagram

55.3.2.1 PCS Reset function

PCS Reset initializes all PCS functions. The PCS Reset function shall be executed whenever one of the following conditions occur:

a) Power on (see 55.3.6.2.2).
b) The receipt of a request for reset from the management entity.

PCS Reset sets pcs_reset=ON while any of the above reset conditions hold true. All state diagrams take the open-ended pcs_reset branch upon execution of PCS Reset. The reference diagrams do not explicitly show the PCS Reset function.

55.3.2.2 PCS Transmit function

The PCS Transmit function shall conform to the PCS 64B/65B Transmit state diagram in Figure 55–16 and the PCS Transmit bit ordering in Figure 55–6 and Figure 55–8.

Dashed rectangles in Figure 55–16 and Figure 55–17 are used to indicate states and state transitions in the transmit process state diagram that shall be supported by PHYs with the EEE capability. PHYs without the EEE capability do not support these transitions.

When communicating with the XGMII, the PCS uses a four octet-wide, synchronous data path, with packet delimiting being provided by transmit control signals and receive control signals. Alignment to 64B/65B is performed in the PCS. The PMA sublayer operates independently of block and packet boundaries. The PCS provides the functions necessary to map packets between the XGMII format and the PMA service interface format.

When the transmit channel is in normal mode, the PCS Transmit process continuously generates 65B blocks based upon the TXD <31:0> and TXC <3:0> signals on the XGMII. The subsequent functions of the PCS Transmit process then pack the resulting blocks, and split the bits into the uncoded set and the coded set, which is processed by a low density parity check (LDPC) encoder and then joint mapped into a transmit LDPC frame of DSQ128 symbols. Transmit data-units are sent to the PMA service interface via the PMA_UNITDATA.request primitive.

In each symbol period, when communicating with the PMA, the PCS Transmit generates a code-group (A_n, B_n, C_n, D_n) that is transferred to the PMA via the PMA_UNITDATA.request primitive. The PMA transmits symbols A_n, B_n, C_n, D_n over wire-pairs BI_DA, BI_DB, BI_DC, and BI_DD respectively. The integer, n, is a time index that is introduced to establish a temporal relationship between different symbol periods. A symbol period, T, is 1.25 ns.

If a PMA_TXMODE.indication message has the value SEND_Z, PCS Transmit passes a vector of zeros at each symbol period to the PMA via the PMA_UNITDATA.request primitive.

If a PMA_TXMODE.indication message has the value SEND_T, PCS Transmit generates sequences of code-groups (TA_n, TB_n, TC_n, TD_n) defined in 55.3.4.2 to the PMA via the PMA_UNITDATA.request primitive. These code-groups are used for training mode and only transmit the values {–9, 9} to keep the transmit power in the training mode the same as the transmit power in normal mode.

During training mode an InfoField is transmitted at regular intervals containing messages for startup operation. By this mechanism, a PHY indicates the status of its own receiver to the link partner and makes requests for remote transmitter settings. (See 55.4.2.5.)

In the normal mode of operation, the PMA_TXMODE.indication message has the value SEND_N, and the PCS Transmit function uses a 65B-LDPC coding technique to generate at each symbol period code-groups that represent data or control. During transmission, the 65B bits are scrambled by the PCS using a PCS scrambler, 8 CRC check bits and an auxiliary channel bit are added, then frames are encoded into a code-group of four-dimensional symbols and transferred to the PMA. During data encoding, PCS Transmit utilizes a LDPC frame encoder.

After reaching the normal mode of operation, EEE-capable PHYs may enter the LPI transmit mode under the control of the MAC via the XGMII. The EEE Transmit state diagram is contained within the PCS Transmit function. The EEE capability is described in 55.3.2.2.22.

55.3.2.2.1 Use of blocks

The PCS maps XGMII signals into 65-bit blocks inserted into an LDPC frame, and vice versa, using a 65B-LDPC coding scheme. The PAM2 PMA training frame synchronization allow establishment of LDPC frame and 65B boundaries by the PCS Synchronization process. Blocks and frames are unobservable and have no meaning outside the PCS. During the LPI mode, LDPC frame boundaries delimit sleep, wake, refresh, quiet, and alert cycles. The PCS functions ENCODE and DECODE generate, manipulate, and interpret blocks and frames as provided by the rules in 55.3.2.2.2.

55.3.2.2.2 65B-LDPC transmission code

The PCS uses a transmission code to improve the transmission characteristics of information to be transferred across the link and to support transmission of control and data characters. The encoding defined by the transmission code ensure that sufficient information is present in the PHY bit stream to make clock recovery possible at the receiver. The encoding also preserves the likelihood of detecting any LDPC frame errors that may occur during transmission and reception of information. In addition, the code enables the receiver to achieve PCS synchronization alignment on the incoming PHY bit stream.

The relationship of block bit positions to XGMII, PMA, and other PCS constructs is illustrated in Figure 55–6 for transmit and Figure 55–7 for receive. These figures illustrate the processing of a multiplicity of blocks containing 8 data octets. See 55.3.2.2.5 for information on how blocks containing control characters are mapped.

55.3.2.2.3 Notation conventions

For values shown as binary, the leftmost bit is the first transmitted bit.

64B/65B encodes 8 data octets or control characters into a block. Blocks containing control characters also contain a block type field. Data octets are labeled D_0 to D_7. Control characters other than /O/, /S/ and /T/ are labeled C_0 to C_7. The control character for ordered set is labeled as O_0 or O_4 since it is only valid on the first octet of the XGMII. The control character for start is labeled as S_0 or S_4 for the same reason. The control character for terminate is labeled as T_0 to T_7.

Two consecutive XGMII transfers provide eight characters that are encoded into one 65-bit transmission block. The subscript in the above labels indicates the position of the character in the eight characters from the XGMII transfers.

Contents of block type fields, data octets and control characters are shown as hexadecimal values. The LSB of the hexadecimal value represents the first transmitted bit. For instance, the block type field 0x1E is sent from left to right as 01111000. The bits of a transmitted or received block are labeled TxB<64:0> and RxB<64:0> respectively where TxB<0> and RxB<0> represent the first transmitted bit. The value of the data/ctrl header is shown as a binary value. Binary values are shown with the first transmitted bit (the LSB) on the left.

55.3.2.2.4 Transmission order

The PCS Transmit bit ordering shall conform to Figure 55–6 and Figure 55–8. Note that these figures show the mapping from XGMII to 64B/65B block for a block containing eight data characters.

55.3.2.2.5 Block structure

Blocks consist of 65 bits. The first bit of a block is the data/ctrl header. Blocks are either data blocks or control blocks. The data/ctrl header is 0 for data blocks and 1 for control blocks. The remainder of the block contains the payload.

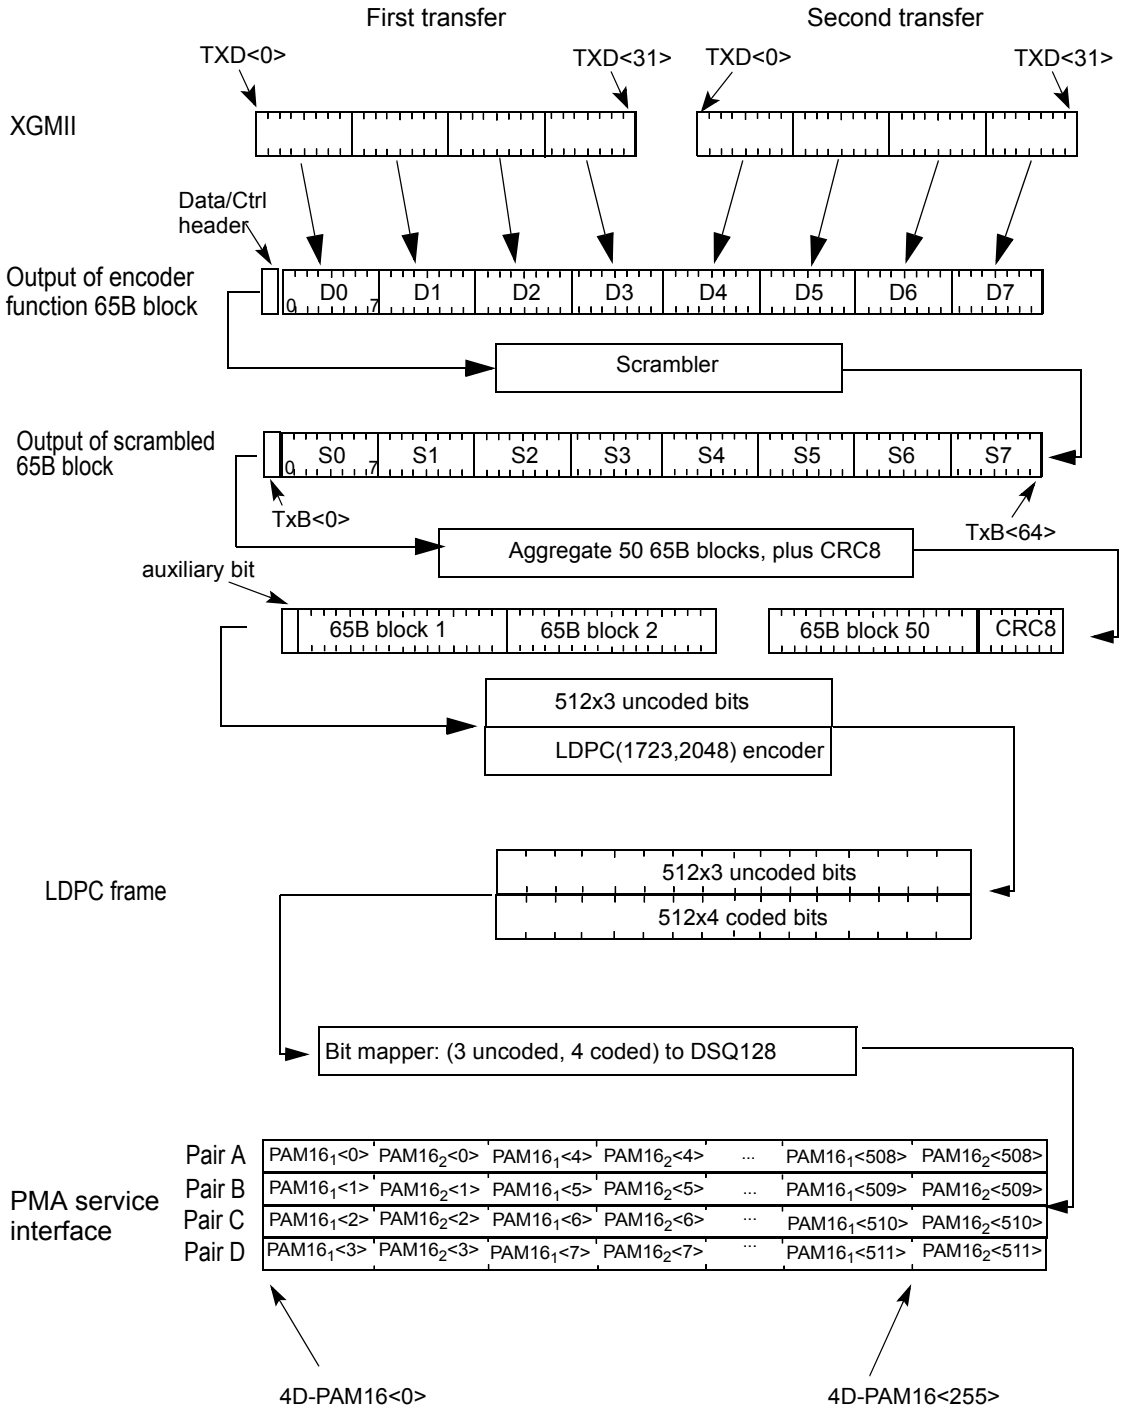

Figure 55–6—PCS Transmit bit ordering

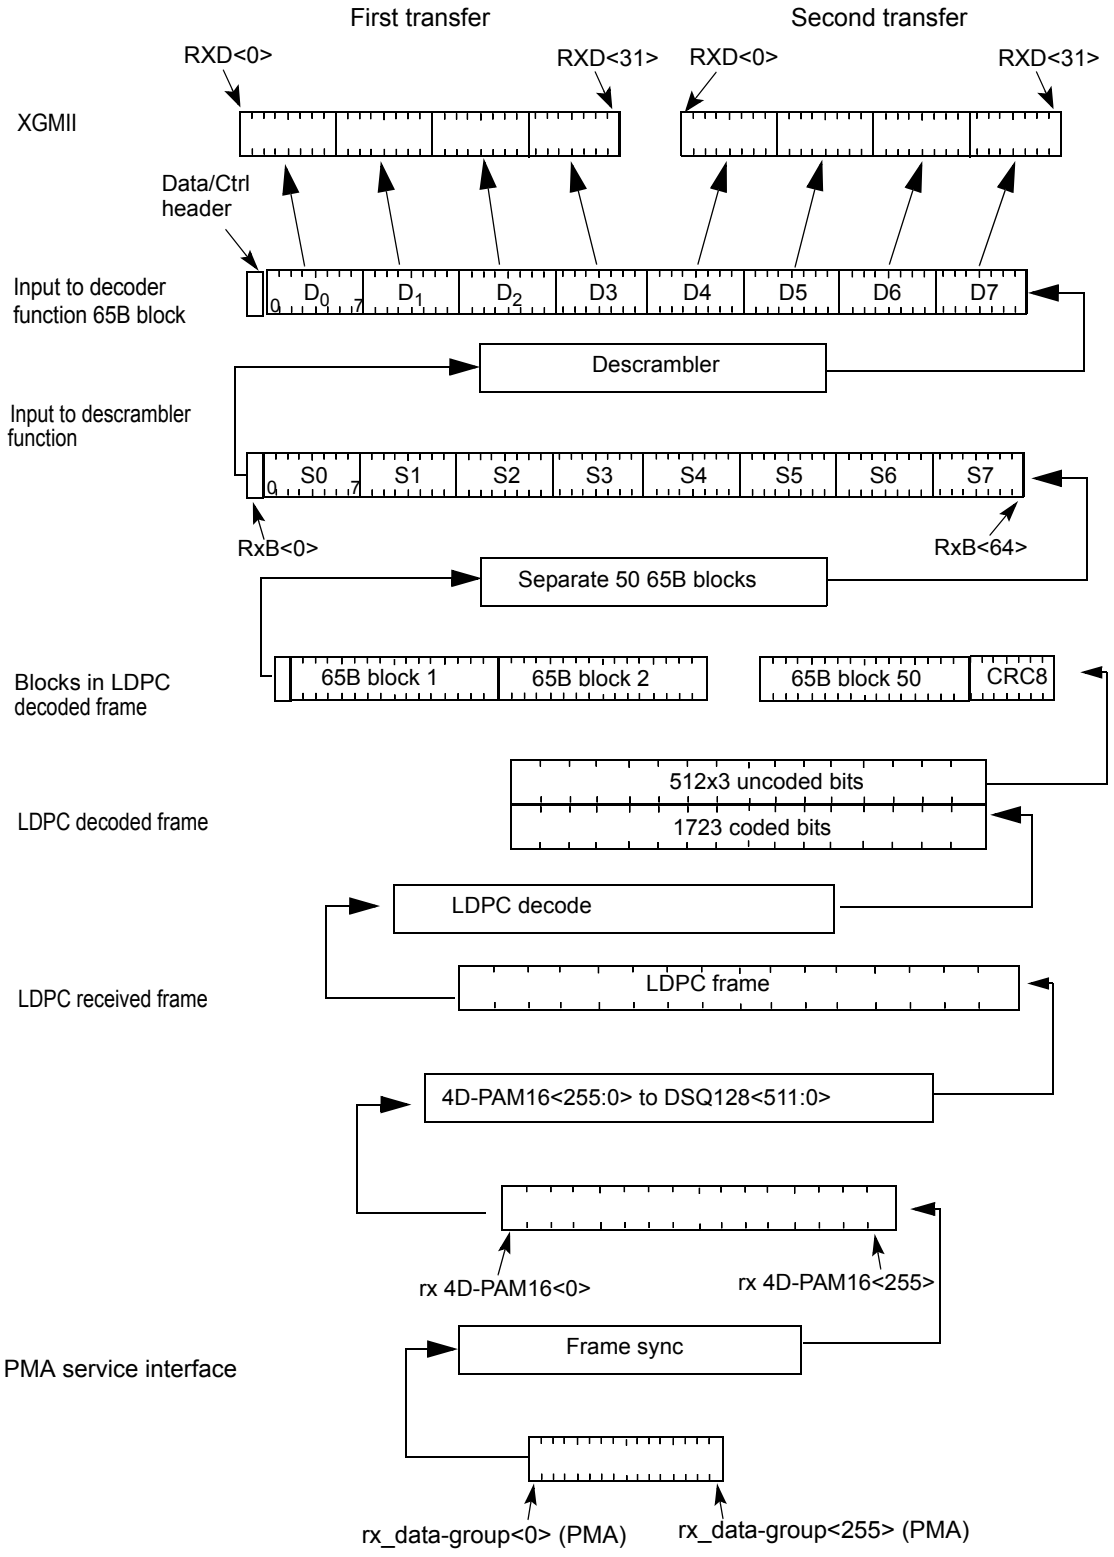

Figure 55–7—PCS Receive bit ordering

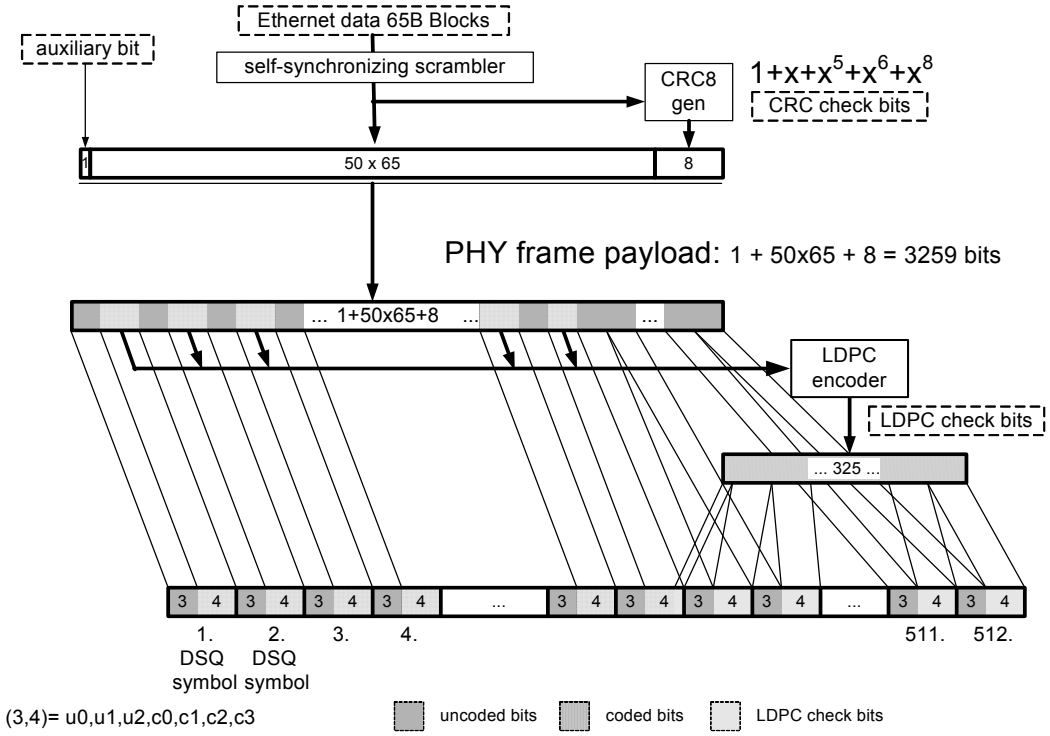

Figure 55–8—PCS detailed transmit bit ordering

Data blocks contain eight data characters. Control blocks begin with an 8-bit block type field that indicates the format of the remainder of the block. For control blocks containing a Start or Terminate character, that character is implied by the block type field. Other control characters are encoded in a 7-bit control code or a 4-bit O Code. Each control block contains eight characters.

The format of the blocks is as shown in Figure 55–9. In the figure, the column labeled Input Data shows, in abbreviated form, the eight characters used to create the 65-bit block. These characters are either data characters or control characters and, when transferred across the XGMII interface, the corresponding TXC or RXC bit is set accordingly. Within the Input Data column, D_0 through D_7 are data octets and are transferred with the corresponding TXC or RXC bit set to zero. All other characters are control octets and are transferred with the corresponding TXC or RXC bit set to one. The single bit fields (thin rectangles with no label in the figure) are sent as zero and ignored upon receipt.

Bits and field positions are shown with the least significant bit on the left. Hexadecimal numbers are shown prepended with '0x', and with the least significant digit on the right. For example the block type field 0x1E is sent as 01111000 representing bits 1 through 8 of the 65-bit block. The least significant bit for each field is placed in the lowest numbered position of the field.

All unused values of block type field[20] are reserved.

[20]The block type field values have been chosen to have a 4-bit Hamming distance between them. The only unused value that maintains the Hamming distance is 0x00.

55.3.2.2.6 Control codes

The same set of control characters are supported by the XGMII and the 10GBASE-T PCS. The representations of the control characters are the control codes. XGMII encodes a control character into an octet (an 8-bit value). The 10GBASE-T PCS encodes the start and terminate control characters implicitly by the block type field. The 10GBASE-T PCS encodes the ordered set control codes using a combination of the block type field and a 4-bit O code for each ordered set. The 10GBASE-T PCS encodes each of the other control characters into a 7-bit C code).

The control characters and their mappings to 10GBASE-T control codes and XGMII control codes are specified in Table 55–1. All XGMII and 10GBASE-T control code values that do not appear in the table shall not be transmitted and shall be treated as an error if received.

Input Data	data ctrl header	**Block Payload**							
Bit Position: / **Data Block Format:**	0	1							64
$D_0\ D_1\ D_2\ D_3/D_4\ D_5\ D_6\ D_7$	0	D_0	D_1	D_2	D_3	D_4	D_5	D_6	D_7
Control Block Formats:		Block							
$C_0\ C_1\ C_2\ C_3/C_4\ C_5\ C_6\ C_7$	1	0x1E	C_0	C_1	C_2	C_3	C_4	C_5	C_6	C_7
$C_0\ C_1\ C_2\ C_3/O_4\ D_5\ D_6\ D_7$	1	0x2D	C_0	C_1	C_2	C_3	O_4	D_5	D_6	D_7
$C_0\ C_1\ C_2\ C_3/S_4\ D_5\ D_6\ D_7$	1	0x33	C_0	C_1	C_2	C_3		D_5	D_6	D_7
$O_0\ D_1\ D_2\ D_3/S_4\ D_5\ D_6\ D_7$	1	0x66	D_1	D_2	D_3	O_0		D_5	D_6	D_7
$O_0\ D_1\ D_2\ D_3/O_4\ D_5\ D_6\ D_7$	1	0x55	D_1	D_2	D_3	O_0	O_4	D_5	D_6	D_7
$S_0\ D_1\ D_2\ D_3/D_4\ D_5\ D_6\ D_7$	1	0x78	D_1	D_2	D_3	D_4		D_5	D_6	D_7
$O_0\ D_1\ D_2\ D_3/C_4\ C_5\ C_6\ C_7$	1	0x4B	D_1	D_2	D_3	O_0	C_4	C_5	C_6	C_7
$T_0\ C_1\ C_2\ C_3/C_4\ C_5\ C_6\ C_7$	1	0x87		C_1	C_2	C_3	C_4	C_5	C_6	C_7
$D_0\ T_1\ C_2\ C_3/C_4\ C_5\ C_6\ C_7$	1	0x99	D_0		C_2	C_3	C_4	C_5	C_6	C_7
$D_0\ D_1\ T_2\ C_3/C_4\ C_5\ C_6\ C_7$	1	0xAA	D_0	D_1		C_3	C_4	C_5	C_6	C_7
$D_0\ D_1\ D_2\ T_3/C_4\ C_5\ C_6\ C_7$	1	0xB4	D_0	D_1	D_2		C_4	C_5	C_6	C_7
$D_0\ D_1\ D_2\ D_3/T_4\ C_5\ C_6\ C_7$	1	0xCC	D_0	D_1	D_2	D_3		C_5	C_6	C_7
$D_0\ D_1\ D_2\ D_3/D_4\ T_5\ C_6\ C_7$	1	0xD2	D_0	D_1	D_2	D_3	D_4		C_6	C_7
$D_0\ D_1\ D_2\ D_3/D_4\ D_5\ T_6\ C_7$	1	0xE1	D_0	D_1	D_2	D_3	D_4	D_5		C_7
$D_0\ D_1\ D_2\ D_3/D_4\ D_5\ D_6\ T_7$	1	0xFF	D_0	D_1	D_2	D_3	D_4	D_5	D_6	

Figure 55–9—64B/65B block formats

55.3.2.2.7 Ordered sets

Ordered sets are used to extend the ability to send control and status information over the link such as remote fault and local fault status. Ordered sets consist of a control character followed by three data characters. Ordered sets always begin on the first octet of the XGMII. 10 Gigabit Ethernet uses one kind of ordered set: the sequence ordered set (see 46.3.4). The sequence ordered set control character is denoted /Q/. An additional ordered set, the signal ordered set, has been reserved and it begins with another control code. The 4-bit O field encodes the control code. See Table 55–1 for the mappings.

55.3.2.2.8 Valid and invalid blocks

A block is invalid if any of the following conditions exists:

a) The block type field contains a reserved value.
b) Any control character contains a value not in Table 55–1.
c) Any O code contains a value not in Table 55–1.
d) The set of eight XGMII characters does not have a corresponding block format in Figure 55–9.
e) The block contains the payload of an invalid received PHY frame or the first 64B/65B block following an invalid received PHY frame to account for self-synchronizing scrambler error propagation.

55.3.2.2.9 Idle (/I/)

Idle control characters (/I/) are transmitted when idle control characters are received from the XGMII. Idle characters may be added or deleted by the PCS to adapt between clock rates. /I/ insertion and deletion shall occur in groups of 4. /I/s may be added following idle or ordered sets. They shall not be added while data is being received. When deleting /I/s, the first four characters after a /T/ shall not be deleted.

.

Table 55–1—Control codes

Control character	Notation	XGMII control codes	10GBASE-T control codes	10GBASE-T O code	8B/10B code[a]
idle	/I/	0x07	0x00		K28.0 or K28.3 or K28.5 without D20.5[b]
LPI	/LI/	0x06	0x06		K28.0 or K28.3 or K28.5 with D20.5[b]
start	/S/	0xFB	Encoded by block type field		K27.7
terminate	/T/	0xFD	Encoded by block type field		K29.7
error	/E/	0xFE	0x1E		K30.7
Sequence ordered set	/Q/	0x9C	Encoded by block type field plus O code	0x0	K28.4
reserved0	/R/[c]	0x1C	0x2D		K28.0
reserved1		0x3C	0x33		K28.1
reserved2	/A/	0x7C	0x4B		K28.3
reserved3	/K/	0xBC	0x55		K28.5
reserved4		0xDC	0x66		K28.6
reserved5		0xF7	0x78		K23.7
Signal ordered set[d]	/Fsig/	0x5C	Encoded by block type field plus O code	0xF	K28.2

[a]For information only. The 8B/10B code is specified in Clause 36. Usage of the 8B/10B code for 10 Gb/s operation is specified in Clause 48.
[b]Use of idle and LPI ordered sets per 48.2.4.2.
[c]The codes for /A/, /K/, and /R/ are used on the XAUI interface to signal idle. They are not present on the XGMII when no errors have occurred, but certain bit errors cause the XGXS to send them on the XGMII.
[d]Reserved for INCITS T11 Fibre Channel use.

55.3.2.2.10 LPI (/LI/)

Low power idle (LPI) control characters (/LI/) on the XGMII indicate that the LPI client is requesting operation in the LPI transmit mode. A continuous stream of LPI control characters (/LI/) is used to maintain a link in the LPI transmit mode. Idle control characters (/I/) are used to transition from the LPI transmit mode to the normal mode. EEE compliant PHYs respond to the LPI XGMII control characters using the

procedure outlined in 55.1.3.3. LPI characters may be added or deleted by the PCS to adapt between clock rates. /LI/ insertion and deletion shall occur in groups of four. /LI/s may be added following low power idle characters. They shall not be added while data is being received.

If EEE is not supported, then /LI/ is not a valid control character.

55.3.2.2.11 Start (/S/)

The start control character (/S/) indicates the start of a packet. This delimiter is only valid on the first octet of the XGMII (TXD<0:7> and RXD<0:7>). Receipt of an /S/ on any other octet of TxD indicates an error. Block type field values implicitly encode an /S/ as the fifth or first character of the block. These are the only characters of a block on which a start can occur.

55.3.2.2.12 Terminate (/T/)

The terminate control character (/T/) indicates the end of a packet. Since packets may be any length, the /T/ can occur on any octet of the XGMII interface and within any character of the block. The location of the /T/ in the block is implicitly encoded in the block type field. A valid end of packet occurs when a block containing a /T/ is followed by a control block that does not contain a /T/.

55.3.2.2.13 ordered set (/O/)

The ordered set control characters (/O/) indicate the start of an ordered set. There are two kinds of ordered sets: the sequence ordered set and the signal ordered set (which is reserved). When it is necessary to designate the control character for the sequence ordered set specifically, /Q/ is used. /O/ is only valid on the first octet of the XGMII. Receipt of an /O/ on any other octet of TXD indicates an error. Block type field values implicitly encode an /O/ as the first or fifth character of the block. The 4-bit O code encodes the specific /O/ character for the ordered set.

Sequence ordered sets may be deleted by the PCS to adapt between clock rates. Such deletion shall only occur when two consecutive sequence ordered sets have been received and shall delete only one of the two. Only Idles may be inserted for clock compensation. Signal ordered sets are not deleted for clock compensation.

55.3.2.2.14 Error (/E/)

The /E/ is sent whenever an /E/ is received. It is also sent when invalid blocks are received. The /E/ allows physical sublayers such as the XGXS and PCS to propagate received errors. See R_BLOCK_TYPE and T_BLOCK_TYPE function definitions in 55.3.6.2.4 for further information.

55.3.2.2.15 Transmit process

The transmit process generates blocks based upon the TXD<31:0> and TXC<3:0> signals received from the XGMII. Two XGMII data transfers are encoded into each block. It takes 256 PMA_UNITDATA transfers to send an LDPC frame of data. Therefore, if the PCS is connected to an XGMII and PMA sublayer where the ratio of their transfer rates is exactly 25:64, then the transmit process does not need to perform rate adaptation. Where the XGMII and PMA sublayer data rates are not synchronized to that ratio, the transmit process needs to insert idles, delete idles, or delete sequence ordered sets to adapt between the rates.

The transmit process generates blocks as specified in the transmit process state diagram. The contents of each block are contained in a vector tx_coded<64:0>, which is passed to the scrambler. tx_coded<0> contains the data/ctrl header and the remainder of the bits contain the block payload.

55.3.2.2.16 PCS scrambler

The payload of the PCS PHY frame is scrambled with a self-synchronizing scrambler. The scrambler for the MASTER shall produce the same result as the implementation shown in Figure 55–10. This implements the scrambler polynomial:[21]

$$G(x) = 1 + x^{39} + x^{58} \tag{55–1}$$

The scrambler for the SLAVE shall produce the same result as the implementation shown in Figure 55–10. This implements the scrambler polynomial:

$$G(x) = 1 + x^{19} + x^{58} \tag{55–2}$$

The initial seed values for the MASTER and SLAVE are left to the implementer. The scrambler is run continuously on all payload bits.

PCS scrambler employed by the MASTER

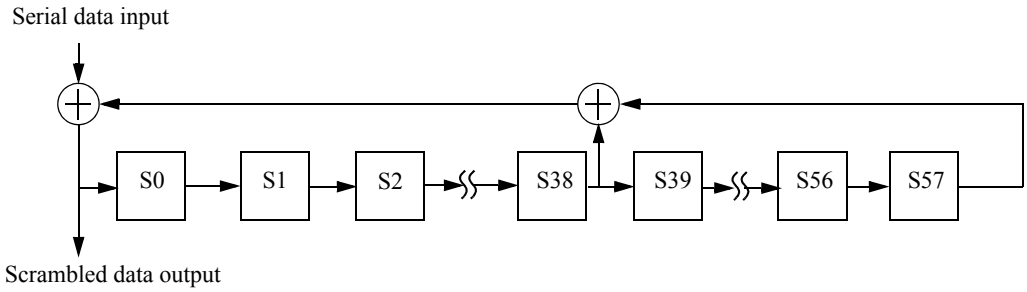

PCS scrambler employed by the SLAVE

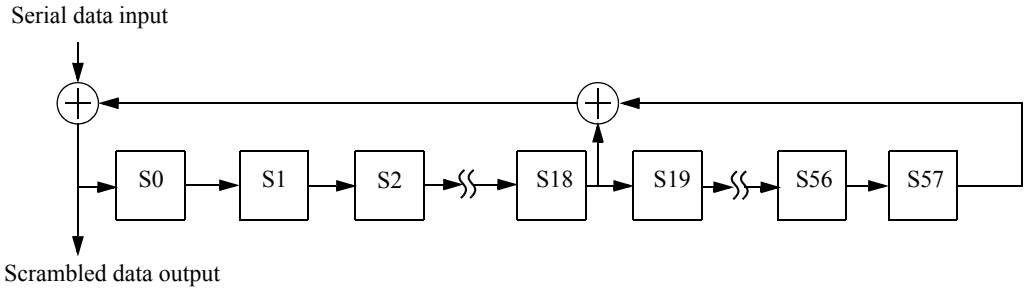

Figure 55–10—MASTER and SLAVE PCS scramblers

[21]The convention here, which considers the most recent bit into the scrambler to be the lowest order term, is consistent with most references and with other scramblers shown in this standard. Some references consider the most recent bit into the scrambler to be the highest order term and would therefore identify this as the inverse of the polynomial in Equation (55–1). In case of doubt, note that the conformance requirement is based on the representation of the scrambler in the figure rather than the polynomial equation.

55.3.2.2.17 CRC8

The aggregated 50 65B blocks shall be used to calculate the cyclic redundancy check parity bits for each PHY frame. The parity bits are generated by the following CRC8 cyclic generator polynomial:

$$C(x) = 1 + x + x^5 + x^6 + x^8 \quad (55–3)$$

The CRC8 shall produce the same result as the implementation shown in Figure 55–11. In Figure 55–11 the 8 delay elements S0,..., S7, shall be initialized to zero. Afterwards the 50 65B blocks of serial data input are used to compute the CRC8 with the switch connected, which is setting CRCgen in Figure 55–11. After all the 50 65B blocks have been processed, the switch is disconnected (setting CRCout) and the 8 values stored in the delay elements are transmitted in the order illustrated, first S7, followed by S6, and so on until the final value S0.

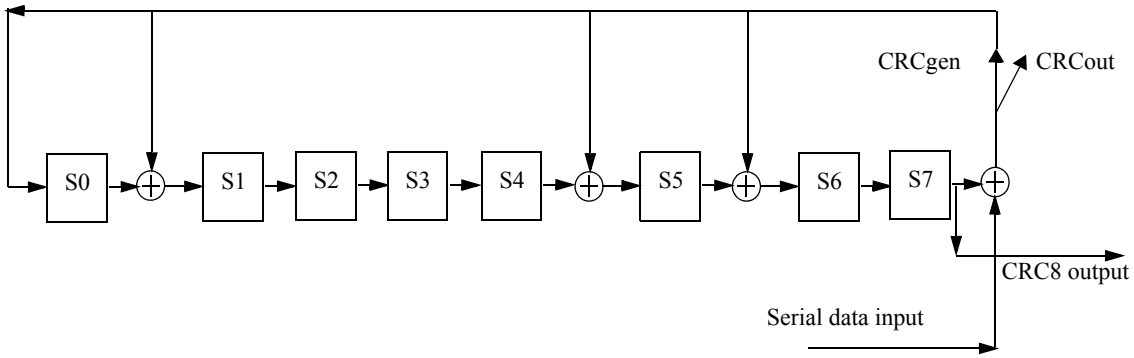

Figure 55–11—CRC8

55.3.2.2.18 LDPC encoder

The resulting payload of scrambled 50 65B blocks, followed by the corresponding 8 bits from the CRC8 and preceded by 1 auxiliary bit results in a total payload of 50 × 65 + 8 +1 = 3259 bits. The use of this bit for vendor-specific communication is outside the scope of this document. For the purposes of this standard it is ignored by the link partner. From the total payload of 3259 bits, 1536 (3 bits for each of the 512 DSQ128 symbols) are uncoded bits and the remaining 1723 shall be encoded by the LDPC(1723, 2048) generator matrix G. G is described in Annex 55A.

The LPDC encoding takes the 1723 bit input code vector $x = [x_0\ x_1\ x_2\ ...\ x_{1722}]$, and shall generate the 2048 bit codeword c represented by the matrix multiplication $c = x \times G$. For both x and c the encoder shall follow the notation described in 55.3.2.2.3 where the LSB (leftmost element of the vectors x and c) is the first bit into the LDPC encoder and the first transmitted bit.

55.3.2.2.19 DSQ128 bit mapping

DSQ128 refers to a two dimensional constellation with 128 possible 2D values, such that the combined 2D symbol carries $\log_2(128)$ or 7 bits. The PHY frame of 1536 uncoded bits and 2048 (LDPC output) coded bits described in Figure 55–6 shall be partitioned into 7-bit groups of (3 uncoded, 4 coded) as described in Figure 55–8. The bit partition is as follows, the first 1290 (3 × 430) uncoded bits are paired with the first 1720 (4 × 430) LDPC input bits. The following 3 uncoded bits are paired with the last 3 LDPC input bits and the first LDPC parity bit. Finally the remaining 243 (3 × 81) uncoded bits are paired with the remaining 324 (4 × 81) LDPC parity bits.

Each 2D-DSQ128 value has two PAM16 components denoted $PAM16_1$ and $PAM16_2$ respectively. The DSQ128 can be constructed by pruning the 256 values of a 2D-PAM16 where every other point in 2D is discarded (like the black or white squares in a checkerboard). The PAM16 components $PAM16_1$ and $PAM16_2$ can each take any of the values from the set {–15, –13, –11, –9, –7, –5, –3, –1, 1, 3, 5, 7, 9, 11, 13, 15}. The mapping from 7 bits where $u_0\ u_1\ u_2$ denote the 3 uncoded bits and $c_0\ c_1\ c_2\ c_3$ denote the 4 coded bits to the DSQ128 is described by the following four steps (the bits from the scrambler output shall be read LSB first):

Step 1:

$x_{13} = (!u_0)\ \&\ u_2$

$x_{12} = u_0\ \text{XOR}\ u_2$

$x_{11} = c_0$

$x_{10} = c_0\ \text{XOR}\ c_1$

$x_{23} = (u_1\ \&\ u_2)\ \text{OR}\ (u_0\ \&\ !u_1)$

$x_{22} = u_1\ \text{XOR}\ u_2$

$x_{21} = c_2$

$x_{20} = c_2\ \text{XOR}\ c_3$

Step 2:

$x_1 = 8x_{13} + 4x_{12} + 2x_{11} + x_{10}$

$x_2 = 8x_{23} + 4x_{22} + 2x_{21} + x_{20}$

Step 3:

$y_1 = (x_1 + x_2)\ \text{mod16}$

$y_2 = (–x_1 + x_2)\ \text{mod16}$

Step 4:

$PAM16_1 = 2y_1 – 15$

$PAM16_2 = 2y_2 – 15$

where '*n* mod16' for an integer *n*, is defined as the integer value *p* in the range 0 to 15 (both inclusive) such that '$p = n + 16m$', for some integer *m*.

55.3.2.2.20 DSQ128 to 4D-PAM16

The DSQ mapper generates 512 2D-DSQ128 symbols per LDPC frame that are mapped onto 256 4D-PAM16 symbols prior to sending to the PMA via PMA_UNITDATA.request. The mapping of DSQ128 to 4D-PAM16 is illustrated in Figure 55–6. As shown in Figure 55–6, the two PAM16 components of each DSQ128 symbol are mapped onto two consecutive time periods on the same wire pair.

55.3.2.2.21 65B-LDPC framer

The 65B-LDPC framer adapts between the 65-bit width of the 65B blocks and the 4D-PAM16 width of the PMA. When the transmit channel is operating in normal mode, the 65B-LDPC sends four PAM16 of transmit data at a time via PMA_UNITDATA.request primitives. The PMA_UNITDATA.request primitives are fully packed with bits.

55.3.2.2.22 EEE capability

The optional 10GBASE-T EEE capability allows compliant PHYs to transition to an LPI mode of operation when link utilization is low.

EEE compliant PHYs shall implement the EEE transmit state diagram, shown in Figure 55–20, within the PCS.

When PCS_Reset is asserted or pcs_data_mode is not asserted, the state diagram enters the TX_NORMAL state.

When a complete 64B/65B block of LPI characters is generated by the PCS transmit function, the PHY transmits the sleep signal to indicate to the link partner that it is transitioning to the LPI transmit mode. If the sleep signal begins on an LDPC frame boundary, then it contains 9 full LDPC frames each composed entirely of 64B/65B LDPC-encoded LP_IDLE blocks. If the sleep signal does not begin on an LDPC frame boundary, then it contains one LDPC frame partially composed of LP_IDLE blocks followed by 9 LDPC frames fully composed of LP_IDLE blocks.

Following the transmission of the sleep signal, quiet-refresh signaling begins, as described in 55.3.5.

After the sleep signal is transmitted LPI control characters shall be input to the PCS scrambler continuously until the PCS Transmit Function exits the LPI transmit mode.

While the PMA asserts SEND_N, the lpi_tx_mode variable shall control the transmit signal through the PMA_UNITDATA.request primitive described as follows:

> When the PHY is not in the PCS_Data state, the lpi_tx_mode variable is ignored.
>
> When the lpi_tx_mode variable takes the value NORMAL and the PMA asserts SEND_N, the PCS passes coded data to the PMA via the PMA_UNITDATA.request primitive as described in 55.3.2.2.
>
> When the lpi_tx_mode variable takes the value QUIET and the PMA asserts SEND_N, the PCS passes zeros to the PMA through the PMA_UNITDATA.request primitive.
>
> When the lpi_tx_mode variable takes the value REFRESH_A and the PMA asserts SEND_N, the PCS passes the PMA training signal to the PMA on pair A, to allow both the local and remote PHY to refresh adaptive filters and timing loops. The PCS passes zeros to all other pairs in this condition. REFRESH_B, REFRESH_C and REFRESH_D operate in a analogous manner for the other pairs.
>
> When the lpi_tx_mode variable takes the value ALERT and the PMA asserts SEND_N, the PCS passes the ALERT vector to the PMA.

The quiet-refresh cycle is repeated until codewords other than LP_IDLE are detected at the XGMII. These codewords indicate that the local system is requesting a transition back to the normal operational mode. Following this event, the PMA_UNITDATA.request message is set to the value ALERT. The alert signal is not synchronized with respect to the quiet-refresh cycle but shall be synchronized so that the alert signal from the PMA begins on a LDPC frame boundary.

The PHY will also transition back to the normal operation mode if an error condition occurs. This error condition is defined as the detection of any characters other than LPI or IDLE at the XGMII.

After the alert signal the PCS completes the transition from LPI mode to normal mode by sending a wake signal containing lpi_wake_time LDPC frames composed of IDLE 64B/65B blocks.

lpi_wake_time is a fixed parameter that is defined as 9 LDPC frames as shown in Table 55–2. The maximum PHY wake time when wake is requested before sleep has been completely transmitted is 7.36 μs (lpi_wake_timer=T_{w_phy} as defined by Clause 78). The maximum PHY wake time when wake is requested after sleep has been completely transmitted is 4.48 μs.

Table 55–2—LPI wake time

lpi_wake_time	lpi_wake_timer when wake starts before sleep signal is complete		lpi_wake_timer when wake starts after sleep signal is complete	
(frames)	(frames)	(μs)	(frames)	(μs)
9	23	7.36	14	4.48

55.3.2.3 PCS Receive function

The PCS Receive function shall conform to the PCS 64B/65B receive state diagram in Figure 55–18 and the PCS Receive bit ordering in Figure 55–7 including compliance with the associated state variables as specified in 55.3.6.

The PCS Receive function accepts received code-groups provided by the PMA Receive function via the parameter rx_symb_vector. The PCS receiver uses knowledge of the encoding rules to correctly align the 65B-LDPC frames. The received 65B-LDPC frames are decoded with error correction; the CRC8 and framing is checked; and the 64B/65B ordered sets are converted to 64-bit data blocks to obtain the signals RXD<31:0> and RXC<3:0> for transmission to the XGMII. Two XGMII data transfers are decoded from each block. Where the XGMII and PMA sublayer data rates are not synchronized to a 25:64 ratio, the receive process inserts idles, delete idles, or delete sequence ordered sets to adapt between rates.

During PMA training mode, PCS Receive checks the received PAM2 framing and signals the reliable acquisition of the descrambler state by setting the parameter scr_status to OK.

When the PCS Synchronization process has obtained synchronization, the LDPC frame error ratio (LFER) monitor process monitors the signal quality asserting hi_lfer if excessive LDPC frame errors are detected (LDPC parity error or CRC8 error). If 40 consecutive LDPC frame errors are detected, the block_lock flag is de-asserted. When block_lock is asserted and hi_lfer is de-asserted, the PCS Receive process continuously accepts blocks. The PCS Receive process monitors these blocks and generates RXD <31:0> and RXC <3:0> on the XGMII.

When the receive channel is in training mode, the PCS Synchronization process continuously monitors PMA_RXSTATUS.indication (loc_rcvr_status). When loc_rcvr_status indicates OK, then the PCS Synchronization process accepts data-units via the PMA_UNITDATA.request primitive. It attains frame and block synchronization based on the PMA training frames and conveys received blocks to the PCS Receive process. The PCS Synchronization process sets the block_lock flag to indicate whether the PCS has obtained synchronization. The PMA training sequence includes 1 bit pattern on pair A every 256 PAM2 symbols, which is aligned with the PCS PHY frame boundary. When the PCS Synchronization process is synchronized to this pattern, block_lock is asserted.

PHYs with the EEE capability support transition to the LPI mode when the PHY has successfully completed training and pcs_data_mode is TRUE. Transitions to and from the LPI mode are allowed to occur independently in the transmit and receive functions. The PCS receive function is responsible for detecting transitions to and from the LPI receive mode and indicating these transitions using signals defined in 55.2.2.

The link partner signals a transition to the LPI mode of operation by transmitting 9 LDPC frames composed entirely of 64B/65B blocks of /LI/. When blocks of /LI/ are detected at the output of the 64B/65B decoder, rx_lpi_active is asserted by the PCS receive function and the /LI/ character is continuously asserted at the receive XGMII. These frames may be preceded by a frame composed partially of /LI/ characters. After these frames the link partner begins transmitting zeros, and it is recommended that the receiver power down receive circuits to reduce power consumption. The receive function uses LDPC frame counters to maintain synchronization with the remote PHY and receives periodic refresh signals that are used to update coefficients, so that the integrity of adaptive filters and timing loops in the PMA is maintained. LPI signaling is defined in 55.3.5. The quiet-refresh cycle continues until the PMA asserts alert_detect to indicate that the alert signal has been reliably detected. After the alert signal the link partner transmits repeated /I/ characters, representing a wake signal. The PHY receive function sends /I/ to the XGMII for 9 LDPC frame periods and then resumes normal operation.

55.3.2.3.1 Frame and block synchronization

When the receive channel is operating in normal mode, the frame and block synchronization function receives data via 4D-PAM16 PMA_UNITDATA.request primitives. It shall form a 4D-PAM16 stream from the primitives by concatenating requests with the PAM16s of each primitive in order from rx_data-group<0> to rx_data-group<255> (see Figure 55–7). It obtains block_lock to the LDPC frames during the PAM2 training pattern using synchronization bits provided on pair A. The 65-bit blocks are extracted based on their location in the LDPC frame.

55.3.2.3.2 PCS descrambler

The descrambler processes the payload to reverse the effect of the scrambler using the same polynomial. It shall produce the same result as the implementations shown in Figure 55–12 for the MASTER and the SLAVE.

55.3.2.3.3 CRC8 receive function

The PCS Receive function shall check the integrity of the CRC8 parity bits defined in 55.3.2.2.17. If the CRC check fails the PHY frame is invalid.

The PCS Receive function may decode and check the CRC8 parity bits simultaneous to resolving the LDPC error correction function. It is recommended that the PCS receiver not use the CRC8 parity check code to assist the LDPC convergence.

55.3.3 Test-pattern generators

The test-pattern generator mode is provided for enabling joint testing of the local transmitter, the channel and remote receiver. When the transmit PCS is operating in test-pattern mode it shall transmit continuously as illustrated in Figure 55–6, with the input to the scrambler set to zero and the initial condition of the scrambler set to any non-zero value. When the receiver PCS is operating in test-pattern mode it shall receive continuously as illustrated in Figure 55–7. After acquiring the self-synchronizing scrambler state, the output of the received scrambled values should ideally be zero. Any nonzero values correspond to receiver bit errors. This mode is further described as test mode 7 in 55.5.2

PCS descrambler employed by the MASTER

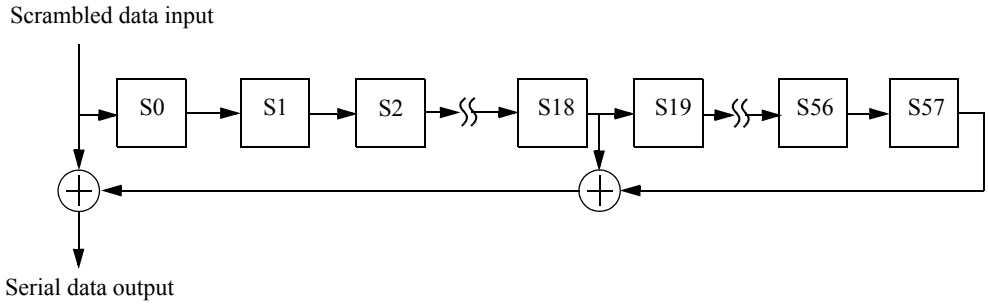

PCS descrambler employed by the SLAVE

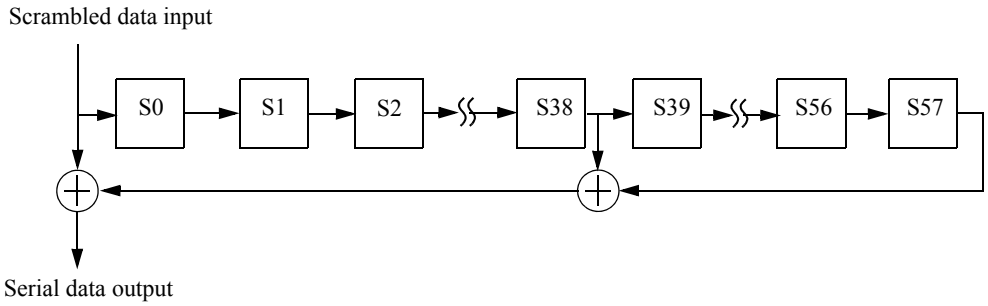

Figure 55–12—MASTER and SLAVE PCS descramblers

55.3.4 PMA training side-stream scrambler polynomials

The PCS Transmit function employs side-stream scrambling for generating 2-level PAM PMA training sequences as shown in Figure 55–13. An implementation of MASTER and SLAVE PHY side-stream scramblers is shown in the "Main PN sequence" box. The bits stored in the shift register delay line at time *n* are denoted by $Scr_n[32:0]$. At each symbol period, the shift register is advanced by one bit, and one new bit represented by $Scr_n[0]$ is generated. The transmitter side-stream scrambler is reset upon execution of the PCS Reset function. If PCS Reset is executed, all bits of the 33-bit vector representing the side-stream scrambler state are arbitrarily set. The initialization of the scrambler state is left to the implementer. In no case shall the scrambler state be initialized to all zeros.

Moreover during Auto-Negotiation each transceiver may request the remote transceiver to reinitialize the values of its scrambler state after every 16384 symbol periods, to generate a periodically repeating pattern with repetition period 16384. The initial 33-bit values of the scrambler state shall be generated by combining 0x39A422 for the 22 MSBs and random value SB10-SB0 from Table 55–15 generated by the local device for the 11 LSBs as shown in Figure 55–13.

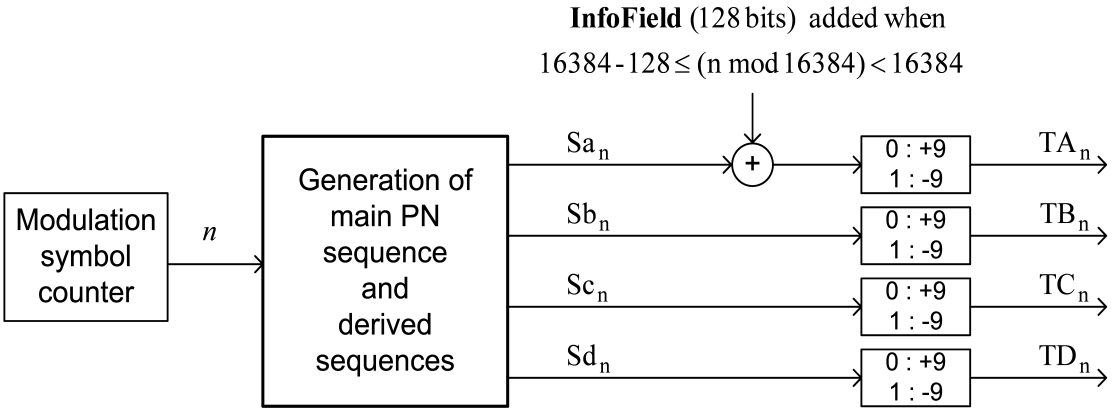

Main PN sequence

n mod 16384 = 0 and
periodic initialization: $Scr_n[32:11] = 0x39A422$

$Scr_n[10:0] =$ SB10 to SB0 of Table 55-15

else: $Scr_n[32:1] = Scr_{n-1}[31:0]$

$$Scr_n[0] = \begin{cases} \mathrm{Scr}_{n-1}[12] \oplus \mathrm{Scr}_{n-1}[32] & \text{if PMA\_CONFIG} = \text{MASTER} \\ \mathrm{Scr}_{n-1}[19] \oplus \mathrm{Scr}_{n-1}[32] & \text{if PMA\_CONFIG} = \text{SLAVE} \end{cases}$$

Derived sequences

$$Sa_n = \begin{cases} Scr_n[0] \oplus 1 & \text{if } n \bmod 256 = 0 \\ Scr_n[0] & otherwise \end{cases}$$

$$Sb_n = Scr_n[3] \oplus Scr_n[8]$$

$$Sc_n = Scr_n[6] \oplus Scr_n[16]$$

$$Sd_n = Scr_n[9] \oplus Scr_n[14] \oplus Scr_n[19] \oplus Scr_n[24]$$

Figure 55–13—A realization of PMA training PAM2 sequences

55.3.4.1 Generation of bits Sa_n, Sb_n, Sc_n, Sd_n

PMA training signal encoding rules are based on the generation, at time *n*, of the four bits Sa_n, Sb_n, Sc_n, Sd_n. These four bits are generated in a systematic fashion using the bits in $Scr_n[32:0]$, and an auxiliary generating polynomial. For both MASTER and SLAVE PHYs, they are obtained by the same linear combinations of bits stored in the transmit scrambler shift register delay line. These four bits are derived from elements of the same maximum-length shift register sequence of length $2^{33}–1$ as $Scr_n[0]$, but shifted in time. The associated delays are all large and different so that there is no short-term correlation among the bits Sa_n, Sb_n, Sc_n, Sd_n. The four bits are generated using the bit $Scr_n[0]$ and the equations in Figure 55–13 in the "Derived sequences" box.

55.3.4.2 Generation of 4D symbols TA_n, TB_n, TC_n, TD_n

The four bits Sa_n, Sb_n, Sc_n, Sd_n are mapped to a 4D symbol (TA_n, TB_n, TC_n, TD_n) as shown in Figure 55–13.

The inversion on pair A at 256 intervals ($n = k \times 256, k = 0, 1, 2, \ldots$) defines the LDPC boundary during data mode. If requested by the link partner, the PCS shall reset the training mode scrambler every 16384 periods aligned with the 256 symbol period inversion on pair A, which corresponds to the time instants $n = m \times 16384, m = 0, 1, 2, \ldots$

Notice that over the repeating time intervals of 16384 and of length 128, $m \times 16384 - 128 \le n < m \times 16384, m = 1, 2, 3, \ldots$, the PMA training pattern in pair A is XOR'ed with the InfoField. Thus, pair A transmits the InfoField, which communicates to the remote transceiver settings of THP and power backoff and other control information.

55.3.4.3 PMA training mode descrambler polynomials

The PHY shall acquire descrambler state synchronization to the PAM2 training sequence and report success through scr_status. For side-stream descrambling, the MASTER PHY shall employ the receiver descrambler generator polynomial $g'_M(x) = 1 + x^{20} + x^{33}$ and the SLAVE PHY shall employ the receiver descrambler generator polynomial $g'_S(x) = 1 + x^{13} + x^{33}$.

55.3.5 LPI signaling

PHYs with EEE capability have transmit and receive functions that can enter and leave the LPI mode independently. The PHY can transition to the LPI mode when the PHY has successfully completed training and pcs_data_mode is TRUE. The transmit function of the PHY initiates a transition to the LPI transmit mode when it generates 64B/65B blocks composed entirely of LPI control characters, as described in 55.3.2.2.22. The transmit function of the link partner signals the transition using the sleep signal. When the transmitter begins to send the sleep signal, it asserts tx_lpi_active and the transmit function enters the LPI transmit mode.

Within the LPI mode PHYs use a repeating quiet-refresh cycle (see Figure 55–14). The first part of this cycle is known as the quiet period and lasts for a time lpi_quiet_time equal to 124 LDPC frame periods. The quiet period is defined in 55.3.5.2. The second part of this cycle is known as the refresh period and lasts for a time lpi_refresh_time equal to 4 LDPC frame periods. The refresh period is defined in 55.3.5.3. A cycle composed of one quiet period and one refresh period is known as a single pair LPI cycle and lasts for a time lpi_qr_time equal to 128 LDPC frame periods. The time taken to complete a quiet-refresh cycle for all four pairs is known as a complete LPI cycle.

lpi_offset, lpi_quiet_time, lpi_refresh_time, lpi_qr_time, and lpi_allpairs_qr_time are timing parameters that are integer multiples of the LDPC frame period. lpi_offset is a fixed value equal to lpi_qr_time/2 that is used to ensure refresh signals are appropriately offset by the link partners.

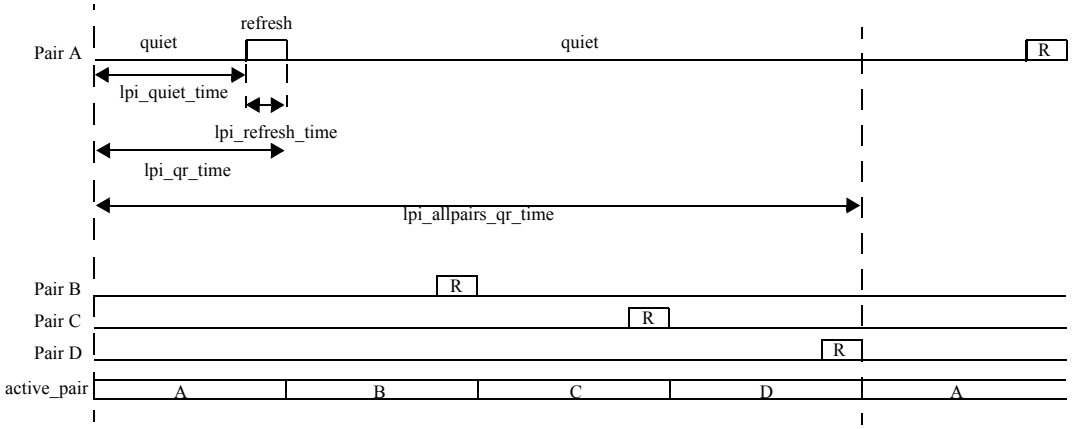

Figure 55–14—Timing periods for LPI signals

PHYs begin the transition from the LPI receive mode when the alert signal is detected by the PMA as defined in 55.4.2.4.

55.3.5.1 LPI Synchronization

To maximize power savings, maintain link integrity, and ensure interoperability, EEE-capable PHYs must synchronize refresh intervals during the LPI mode. The transition to PCS_Test is used as a fixed timing reference for the link partners. Refresh signaling is derived by counting LDPC frames from the transition to PCS_Test. An EEE-capable PHY shall support loop timing and loop timing shall be enabled on the slave PHY.

In initial training, normal retraining, and fast retraining, with or without the EEE capability being supported, the master and slave signal when they will transition to PCS_Test using the transition counter following the procedure described in 55.4.2.5.14.

A EEE-capable PHY in slave mode is responsible for synchronizing its PMA training frame to the master's PMA training frame during the transition to PMA_Training_Init_S. The slave shall ensure that its PMA training frames are synchronized to the master's PMA training frames within 1 LDPC frame, measured at the slave MDI on pair A. In addition, the slave shall initialize its transition counter so that it transitions to PCS_Test within 1 LDPC frame of the master PHY's transition to PCS_Test, measured at the slave PHY's MDI on pair A. This mechanism ensures that the refresh offset is bounded to a small value at both MDI interfaces, thus ensuring there is no overlap of master and slave signals when both transmit and receive are in the LPI mode.

Following the transition to PCS_Test, the PCS counts transmitted and received LDPC frames, and uses these counters to generate refresh and pair control signals for the transmit and receive functions. The transmitted LDPC frame count is named tx_ldpc_frame_cnt. The received LDPC frame count is named rx_ldpc_frame_cnt.

The master and slave shall derive the active pair and refresh_active signals from the LDPC frame counters as shown in Table 55–3 and Table 55–4.

Table 55–3—Synchronization logic derived from slave signal LDPC frame count

Slave-side variable	Master-side variable	for master u=rx_ldpc_frame_cnt for slave u=tx_ldpc_frame_cnt
tx_refresh_active=true	rx_refresh_active=true	lpi_offset – lpi_refresh_time ≤ mod(u,lpi_qr_time) < lpi_offset
tx_lpi_full_refresh=true	N/A	lpi_offset – lpi_refresh_time = mod(u,lpi_qr_time)
tx_active_pair=PAIR_A	rx_active_pair=PAIR_A	lpi_offset + lpi_qr_time ≤ u < lpi_offset + 2 x lpi_qr_time
tx_active_pair=PAIR_B	rx_active_pair=PAIR_B	lpi_offset + 2 x lpi_qr_time ≤ u < lpi_offset + 3 x lpi_qr_time
tx_active_pair=PAIR_C	rx_active_pair=PAIR_C	lpi_offset + 3 x lpi_qr_time ≤ u < 4 x lpi_qr_time OR 0 ≤ u < lpi_offset
tx_active_pair=PAIR_D	rx_active_pair=PAIR_D	lpi_offset ≤ u < lpi_offset + lpi_qr_time

Table 55–4—Synchronization logic derived from master signal LDPC frame count

Slave-side variable	Master-side variable	for master v=tx_ldpc_frame_cnt for slave v=rx_ldpc_frame_cnt
rx_refresh_active=true	tx_refresh_active=true	lpi_quiet_time ≤ mod(v,lpi_qr_-time)
N/A	tx_lpi_full_refresh=true	lpi_quiet_time = mod(v,lpi_qr_-time)
rx_active_pair=PAIR_A	tx_active_pair=PAIR_A	0 ≤ v < lpi_qr_time
rx_active_pair=PAIR_B	tx_active_pair=PAIR_B	lpi_qr_time ≤ v < 2 x lpi_qr_time
rx_active_pair=PAIR_C	tx_active_pair=PAIR_C	2 x lpi_qr_time ≤ v < 3 x lpi_qr_-time
rx_active_pair=PAIR_D	tx_active_pair=PAIR_D	3 x lpi_qr_time ≤ v < 4 x lpi_qr_-time

55.3.5.2 Quiet period signaling

During the quiet period the transmitters on all four pairs should be turned off. Average launch power (as measured from 28 LDPC frames after a refresh period to 28 LDPC frames before the next refresh period on the same lane) for each Transmitter shall be less than –41dBm. This requirement does not apply to the periods when the alert signal is transmitted as defined in 55.4.2.2.1.

55.3.5.3 Refresh period signaling

During the LPI mode 10GBASE-T PHYs use staggered, out-of-phase refresh signaling to maximize power savings. Two-level PAM refresh symbols are generated using the PMA side-stream scrambler polynomials described in 55.3.4 and exactly as is shown in Figure 55–13 with the exception that the InfoField consists of a sequence of 128 zeros. The training sequence without periodic reinitialization described in 55.3.4 shall be used during the LPI mode, with the scramblers free-running from PCS Reset. If scrambler reinitialization is used for normal training, it shall be disabled and the scramblers shall begin free-running when the PHY Control state diagram enters the PCS_Test state.

Refresh signals shall be sent using the THP filter as described in 55.4.3.1. At the start of each refresh signal the THP feedback delay line shall be initialized with zeros.

While a transmit function is in the LPI transmit mode only one of the transmit pairs will be active during a refresh period. tx_symb_vector for all transmit pairs that are not active shall be set to zero.

When tx_symb_vector has the value ALERT and the PHY is master, the transmitter on pair A shall be active and all other pairs shall be quiet. When tx_symb_vector has the value ALERT and the PHY is slave, the transmitter on pair C shall be active and all other pairs shall be quiet. If lpi_tx_mode=REFRESH_A on a MASTER PHY or lpi_tx_mode=REFRESH_C on a SLAVE PHY, and tx_symb_vector has the value ALERT, then the alert signaling shall be transmitted in place of the refresh signaling where the signals overlap.

55.3.6 Detailed functions and state diagrams

55.3.6.1 State diagram conventions

The body of this subclause is comprised of state diagrams, including the associated definitions of variables, constants, and functions. Should there be a discrepancy between a state diagram and descriptive text, the state diagram prevails.

The notation used in the state diagrams follows the conventions of 21.5. State diagram timers follow the conventions of 14.2.3.2. The notation ++ after a counter or integer variable indicates that its value is to be incremented.

55.3.6.2 State diagram parameters

55.3.6.2.1 Constants

EBLOCK_R<71:0>
: 72 bit vector to be sent to the XGMII interface containing /E/ in all the eight character locations.

EBLOCK_T<64:0>
: 65 bit vector to be sent to the LDPC encoder containing /E/ in all the eight character locations.

LBLOCK_R<71:0>
: 72 bit vector to be sent to the XGMII interface containing two Local Fault ordered sets. The Local Fault ordered set is defined in 46.3.4.

LBLOCK_T<64:0>
: 65 bit vector to be sent to the LDPC encoder containing two Local Fault ordered sets.

LPBLOCK_R<71:0>
: 72 bit vector to be sent to the XGMII containing /LI/ in all the eight character locations.

LPBLOCK_T<64:0>
: 65 bit vector to be sent to the LDPC encoder containing /LI/ in all the eight character locations.

IBLOCK_R<71:0>
: 72 bit vector to be sent to the XGMII containing /I/ in all the eight character locations.

IBLOCK_T<64:0>
: 65 bit vector to be sent to the LDPC encoder containing /I/ in all the eight character locations.

UBLOCK_R<71:0>
: 72 bit vector to be sent to the XGMII containing two Link Interruption ordered sets. The Link Interruption ordered set is defined in 46.3.4.

55.3.6.2.2 Variables

lfer_test_lf
: Boolean variable that is set true when a new LDPC frame is available for testing and false when LFER_TEST_LF state is entered. A new LDPC frame is available for testing when the Block Sync process has accumulated enough symbols from the PMA to evaluate the next LDPC frame.

block_lock
: Boolean variable that is set true when receiver acquires block delineation.

hi_lfer
: Boolean variable which is asserted true when the lfer_cnt exceeds 16 indicating a bit error ratio > 4×10^{-4}.

pcs_reset
: Boolean variable that controls the resetting of the PCS. It is true whenever a reset is necessary

including when reset is initiated from the MDIO, during power on, and when the MDIO has put the PCS into low-power mode.

rx_coded<64:0>

Vector containing the input to the 64B/65B decoder. The format for this vector is shown in Figure 55–9. The leftmost bit in the figure is rx_coded<0> and the rightmost bit is rx_coded<64>.

rx_raw<71:0>

Vector containing two successive XGMII output transfers. RXC<0> through RXC<3> for the first transfer are taken from rx_raw<0> through rx_raw<3>, respectively. RXC<0> through RXC<3> for the second transfer are taken from rx_raw<4> through rx_raw<7>, respectively. RXD<0> through RXD<31> for the first transfer are taken from rx_raw<8> through rx_raw<39>, respectively. RXD<0> through RXD<31> for the second transfer are taken from rx_raw<40> through rx_raw<71>, respectively.

lf_valid

Boolean indication that is set true if received LDPC frame is valid. LDPC frame is valid if:

a. All parity checks of the coded bits are satisfied
b. CRC8 check is satisfied

tx_coded<64:0>

Vector containing the output from the 64B/65B encoder. The format for this vector is shown in Figure 55–9. The leftmost bit in the figure is tx_coded<0> and the rightmost bit is tx_coded<64>.

tx_raw<71:0>

Vector containing two successive XGMII transfers. TXC<0> through TXC<3> for the first transfer are placed in tx_raw<0> through tx_raw<3>, respectively. TXC<0> through TXC<3> for the second transfer are placed in tx_raw<4> through tx_raw<7>, respectively. TXD<0> through TXD<31> for the first transfer are placed in tx_raw<8> through tx_raw<39>, respectively. TXD<0> through TXD<31> for the second transfer are placed in tx_raw<40> through tx_raw<71>, respectively.

The following variables are required for PHYs that support the EEE capability:

tx_lpi_active

A Boolean variable that is set true when the PHY transmit function is operating in the LPI transmit mode and during transitions to and from the LPI transmit mode (i.e., at any time when the PHY is transmitting sleep, alert, wake, or quiet-refresh signaling). It is set false otherwise.

tx_lpi_qr_active

A Boolean variable that is set true during the LPI transmit mode, when the PHY is transmitting quiet-refresh signaling. Set false otherwise.

rx_lpi_active

A Boolean variable that is set true when the PHY receive function is operating in the LPI receive mode and set false otherwise. The LPI receive mode begins when the sleep signal is detected and lasts until the alert signal is detected. When the EEE capability is not supported, rx_lpi_active is set false.

tx_lpi_req

A Boolean variable that is set true when the LPI client indicates that it is requesting operation in the LPI transmit mode via the XGMII and set false otherwise.

alert_detect

Indicates that an alert signal from the link partner has been received at the MDI as indicated by PMA_ALERTDETECT.indication(alert_detect).

tx_lpi_alert_active

A Boolean variable that is set true when the PHY is transmitting ALERT signaling. Set false otherwise.

rx_lpi_wake

A Boolean variable that is set true when the PHY receiver is in the WAKE state and sending IDLE to the XGMII. Set false otherwise. When the EEE capability is not supported, rx_lpi_wake is set false.

tx_active_pair
: A variable indicating the transmit active pair during the LPI transmit mode. The variable may take the values PAIR_A, PAIR_B, PAIR_C, PAIR_D. This variable is defined in 55.3.5.1.

lpi_tx_mode
: A variable indicating the signaling to be used from the PCS to the PMA across the PMA_UNITDATA.request (tx_symb_vector) interface.
lpi_tx_mode controls tx_symb_vector only when tx_mode is set to SEND_N.
The variable is set to NORMAL when (!tx_lpi_qr_active * !tx_lpi_alert_active), indicating that the PCS is in the normal mode of operation and will encode code-groups as described in Figure 55–16 and Figure 55–17.
The variable is set to REFRESH_A when (tx_lpi_qr_active * (tx_active_pair==PAIR_A) * tx_refresh active).
The variable is set to REFRESH_B when (tx_lpi_qr_active * (tx_active_pair==PAIR_B) * tx_refresh active).
The variable is set to REFRESH_C when (tx_lpi_qr_active * (tx_active_pair==PAIR_C) * tx_refresh active).
The variable is set to REFRESH_D when (tx_lpi_qr_active * (tx_active_pair==PAIR_D) * tx_refresh active).
The variable is set to QUIET when (tx_lpi_qr_active * (!tx _refresh_active + tx_lpi_initial_quiet))
The variable is set to ALERT when (tx_lpi_alert_active)

tx_refresh_active
: A Boolean value. This variable is set true following the logic described in 55.3.5.1.

tx_lpi_full_refresh
: A Boolean value. This variable is set true following the logic described in 55.3.5.1.

tx_lpi_initial_quiet
: A Boolean value. This variable is set true when the transmit function enters the LPI transmit mode and a partial refresh will be replaced by quiet signaling.

ldpc_frame_done
: A Boolean value. This variable is set true when the final symbol of each LDPC frame is transmitted and is set false otherwise.

The following variable is only required for PHYs that support the fast retrain capability:

fr_sigtype
: If fast retrain is supported, this variable is set based on the value in 1.147.2:1 as follows:
00 IBLOCK_R
01 LBLOCK_R
10 UBLOCK_R
11 Reserved

55.3.6.2.3 Timers

State diagram timers follow the conventions described in 14.2.3.2.

125_ustimer:
: Timer that is triggered every 125 μs +1%, –25%.

The following timers are required for PHYs that support the EEE capability:

lpi_tx_sleep_timer
: This timer defines the time the local transmitter sends the sleep signal to the link partner.

Values: The condition lpi_tx_sleep_timer_done becomes true upon timer expiration.
Duration: This timer shall have a period equal to 9 LDPC frame periods.

lpi_tx_alert_timer
This timer defines the time the local transmitter transmits the alert signal.
Values: The condition lpi_tx_alert_timer_done becomes true upon timer expiration.
Duration: This timer shall have a period equal to 4 LDPC frame periods.

lpi_tx_wake_timer:
This timer defines the time the local transmitter transmits the wake signal.
Values: The condition lpi_tx_wake_timer_done becomes true upon timer expiration.
Duration: This timer shall have a period equal to lpi_wake_time LDPC frame periods.

lpi_rx_wake_timer:
This timer defines the time the receiver sends IDLE blocks to the XGMII after the alert signal is detected.
Values: The condition lpi_rx_wake_timer_done becomes true upon timer expiration.
Duration: This timer shall have a period equal to lpi_wake_time LDPC frame periods.

55.3.6.2.4 Functions

DECODE(rx_symb_vector<64:0>)
In the PCS Receive process, this function takes as its argument 65-bit rx_coded<64:0> from the LDPC decoder and decodes the 65B-LDPC bit vector returning a vector rx_raw<71:0>, which is sent to the XGMII. The DECODE function shall decode the block based on code specified in 55.3.2.2.2.

ENCODE(tx_raw<71:0>)
Encodes the 72-bit vector received from the XGMII, returning 65 bit vector tx_coded. The ENCODE function shall encode the block as specified in 55.3.2.2.2.

R_BLOCK_TYPE = {C, S, T, D, E, I, LI, LII}
When the EEE capability is not supported, this function classifies each 65 bit rx_coded vector as belonging to one of the five types {C, S, T, D, E} depending on its contents.
When the EEE capability is supported, this function classifies each 65 bit rx_coded vector as belonging to the eight types depending on its contents. A vector may simultaneously belong to the C and I types when it contains eight valid control characters that are all /I/, but in every other case the vector belongs to only one type.
Values: C; The vector contains a data/ctrl header of 1 and one of the following:
a) A block type field of 0x1E and eight valid control characters other than /E/ and /LI/;
b) A block type field of 0x2D or 0x4B, a valid O code, and four valid control characters;
c) A block type field of 0x55 and two valid O codes.

S; The vector contains a data/ctrl header of 1 and one of the following:
a) A block type field of 0x33 and four valid control characters;
b) A block type field of 0x66 and a valid O code;
c) A block type field of 0x78.

T; The vector contains a data/ctrl header of 1, a block type field of 0x87, 0x99, 0xAA, 0xB4, 0xCC, 0xD2, 0xE1 or 0xFF and all control characters are valid.

D; The vector contains a data/ctrl header of 0.

I; If the optional EEE capability is supported, then the I type is a special case of the C type where the vector contains a data/ctrl header of 1, a block type field of 0x1e, and eight control characters of /I/.

LI: If the optional EEE capability is supported, then the LI type occurs when the vector contains a data/ctrl header of 1, a block type field of 0x1e, and eight control characters of /LI/.

LII: If the optional EEE capability is supported, then the LII type occurs when the vector contains a data/ctrl header of 1, a block type field of 0x1E, and one of the following:
a) four control characters of /LI/ followed by four control characters of /I/;
b) four control characters of /I/ followed by four control characters of /LI/
E; The vector does not meet the criteria for any other value.

A valid control character is one containing a 10GBASE-T control code specified in Table 55–1. A valid O code is one containing an O code specified in Table 55–1.

R_TYPE(rx_coded<64:0>)

Returns the R_BLOCK_TYPE of the rx_coded<64:0> bit vector.

R_TYPE_NEXT

Prescient end of packet check function. It returns the R_BLOCK_TYPE of the rx_coded vector immediately following the current rx_coded vector.

T_BLOCK_TYPE = {C, S, T, D, E, I, LI, LII}

When the EEE capability is not supported, this function classifies each 72-bit tx_raw vector as belonging to one of the five types {C, S, T, D, E} depending on its contents.

When the EEE capability is supported, this function classifies each 72-bit tx_raw vector as belonging to the eight types depending on its contents. A vector may simultaneously belong to the C and I types when it contains eight valid control characters that are all /I/, but in every other case the vector belongs to only one type.

Values: C; The vector contains one of the following:
a) eight valid control characters other than /O/, /S/, /T/, /E/, and /LI/;
b) one valid ordered set and four valid control characters other than /O/, /S/ and /T/;
c) two valid ordered sets.
S; The vector contains an /S/ in its first or fifth character, any characters before the S character are valid control characters other than /O/, /S/ and /T/ or form a valid ordered set, and all characters following the /S/ are data characters.
T; The vector contains a /T/ in one of its characters, all characters before the /T/ are data characters, and all characters following the /T/ are valid control characters other than /O/, /S/ and /T/.
D; The vector contains eight data characters.
I; If the optional EEE capability is supported, then the I type is a special case of the C type where the vector contains eight control characters of /I/.
LI: If the optional EEE capability is supported, then the LI type occurs when the vector contains eight control characters of /LI/.
LII: If the optional EEE capability is supported, then the LII type occurs when the vector contains one of the following:
a) four control characters of /LI/ followed by four control characters of /I/;
b) four control characters of /I/ followed by four control characters of /LI/.
E; The vector does not meet the criteria for any other value.

A tx_raw character is a control character if its associated TXC bit is asserted. A valid control character is one containing an XGMII control code specified in Table 55–1. A valid ordered set consists of a valid /O/ character in the first or fifth characters and data characters in the three characters following the /O/. A valid /O/ is any character with a value for O code in Table 55–1.

T_TYPE(tx_raw<71:0>)

Returns the T_BLOCK_TYPE of the tx_raw<71:0> bit vector.

T_TYPE_NEXT

Prescient end of packet check function. It returns the FRAME_TYPE of the tx_raw vector immediately following the current tx_raw vector.

55.3.6.2.5 Counters

lfer_cnt
: Count up to a maximum of 16 of the number of invalid LDPC frames within the current 125 μs period.

The following counters are required for PHYs that support the EEE capability:

tx_ldpc_frame_cnt
: An integer value that counts transmit LDPC frame periods. The counter is reset when the first symbol of the first LDPC frame crosses the MDI on pair A in the transmit direction after normal training or fast retraining. It is incremented after the last symbol of each transmitted LDPC frame. tx_ldpc_frame_cnt is reset to 0 when tx_ldpc_frame_cnt = lpi_qr_time x 4.

rx_ldpc_frame_cnt
: An integer value that counts receive LDPC frame periods. The counter is reset when the first symbol of the first LDPC frame crosses the MDI on pair A in the receive direction after normal training or fast retraining. It is incremented after the last symbol of each received LDPC frame. rx_ldpc_frame_cnt is reset to 0 when rx_ldpc_frame_cnt = lpi_qr_time x 4.

lpi_rxw_err_cnt
: An integer value that counts the number of receive wake on error conditions. lpi_rxw_err_cnt is reset to zero during PCS_Test. The counter is reflected in register 3.22 (see 45.2.3.10).

55.3.6.3 Messages

PMA_UNITDATA.indication (rx_symb_vector)
: A signal sent by PMA Receive indicating that a vector of four symbols is available in rx_symb_vector. (See 55.2.2.4.)

PMA_UNITDATA.request (tx_symb_vector)
: A signal sent to PMA Transmit indicating that a vector of four PAM16 symbols is available tx_symb_vector. (See 55.2.2.3.)

PCS_status
: Indicates whether the PCS is in a fully operational state. (See 55.3.7.1.)

55.3.6.4 State diagrams

The LFER Monitor state diagram shown in Figure 55–15 monitors the received signal for high LDPC frame error ratio.

The 64B/65B Transmit state diagram shown in Figure 55–16 controls the encoding of 65B transmitted blocks. It makes exactly one transition for each 65B transmit block processed. Though the Transmit state diagram sends Local Fault ordered sets when reset is asserted, the scrambler and 65B-LDPC may not be operational during reset. Thus, the Local Fault ordered sets may not appear on the PMA service interface.

The 64B/65B Receive state diagram shown in Figure 55–18 controls the decoding of 65B received blocks. It makes exactly one transition for each receive block processed except for the transition from RX_WE to RX_E, which occurs immediately after the RX_WE processes are complete.

The PCS shall perform the functions of LFER Monitor, Transmit, and Receive as specified in these state diagrams. The PCS shall not perform the LFER Monitor function during LPI receive operation from the time that the PCS 64B/65B Receiver enters the state RX_L, until the state RX_W is exited.

Transitions surrounded by dashed rectangles indicate requirements for 10GBASE-T EEE-capable implementations.

55.3.7 PCS management

The following objects apply to PCS management. If an MDIO Interface is provided (see Clause 45), they are accessed via that interface. If not, it is recommended that an equivalent access be provided.

55.3.7.1 Status

PCS_status:
: Indicates whether the PCS is in a fully operational state. It is only true if block_lock is true and hi_lfer is false. This status is reflected in MDIO register 3.32.12. A latch low view of this status is reflected in MDIO register 3.1.2 and a latch high of the inverse of this status, Receive fault, is reflected in MDIO register 3.8.10.

block_lock:
: Indicates the state of the block_lock variable. This status is reflected in MDIO register 3.32.0. A latch low view of this status is reflected in MDIO register 3.33.15.

hi_lfer:
: Indicates the state of the hi_lfer variable. This status is reflected in MDIO register 3.32.1. A latch high view of this status is reflected in MDIO register 3.33.14.

Rx LPI indication:
: For EEE capability, this variable indicates the current state of the receive LPI function. This flag is set to TRUE (register bit set to one) when the PCS 64B/65B Receive state diagram (Figure 55–19) is in the RX_L or RX_W states. This status is reflected in MDIO register 3.1.8. A latch high view of this status is reflected in MDIO register 3.1.10 (Rx LPI received).

Tx LPI indication
: For EEE capability, this variable indicates the current state of the transmit LPI function. This flag is set to TRUE (register bit set to one) when the PCS 64B/65B Transmit state diagram (Figure 55–17) is in the TX_L or TX_W states. This status is reflected in MDIO register 3.1.9. A latch high view of this status is reflected in MDIO register 3.1.11 (Tx LPI received).

55.3.7.2 Counters

The following counters are reset to zero upon read and upon reset of the PCS. When they reach all ones, they stop counting. Their purpose is to help monitor the quality of the link.

lfer_count:
: 6-bit counter that counts each time LFER_BAD_LF state is entered. This counter is reflected in MDIO register bits 3.33.13:8. The counter is reset when register 3.33 is read by management. Note that this counter counts a maximum of 16 counts per 125 μs since the LFER_BAD_LF can be entered a maximum of 16 times per 125 μs window.

errored_block_count:
8-bit counter. When the receiver is in normal mode, errored_block_count counts once for each time RX_E state is entered. This counter is reflected in MDIO register bits 3.33.7:0

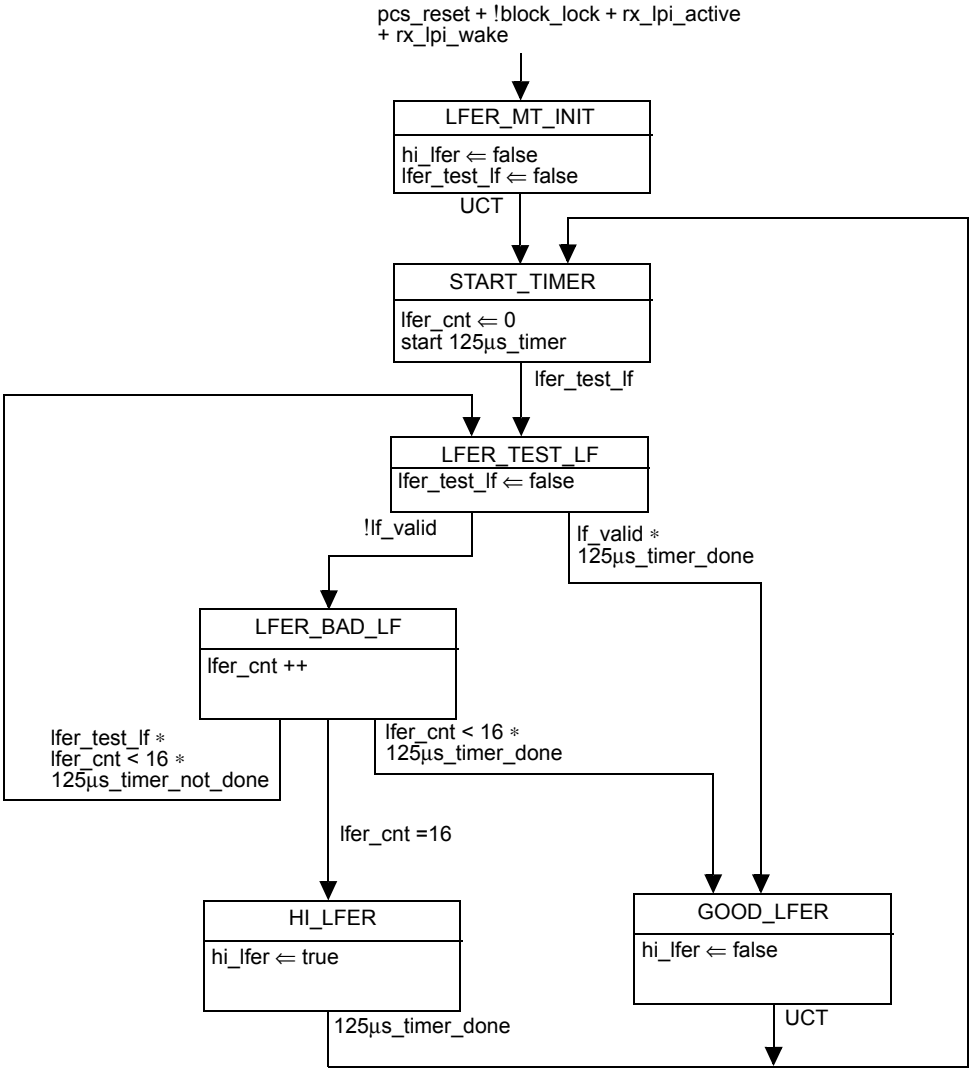

Figure 55–15—LFER monitor state diagram

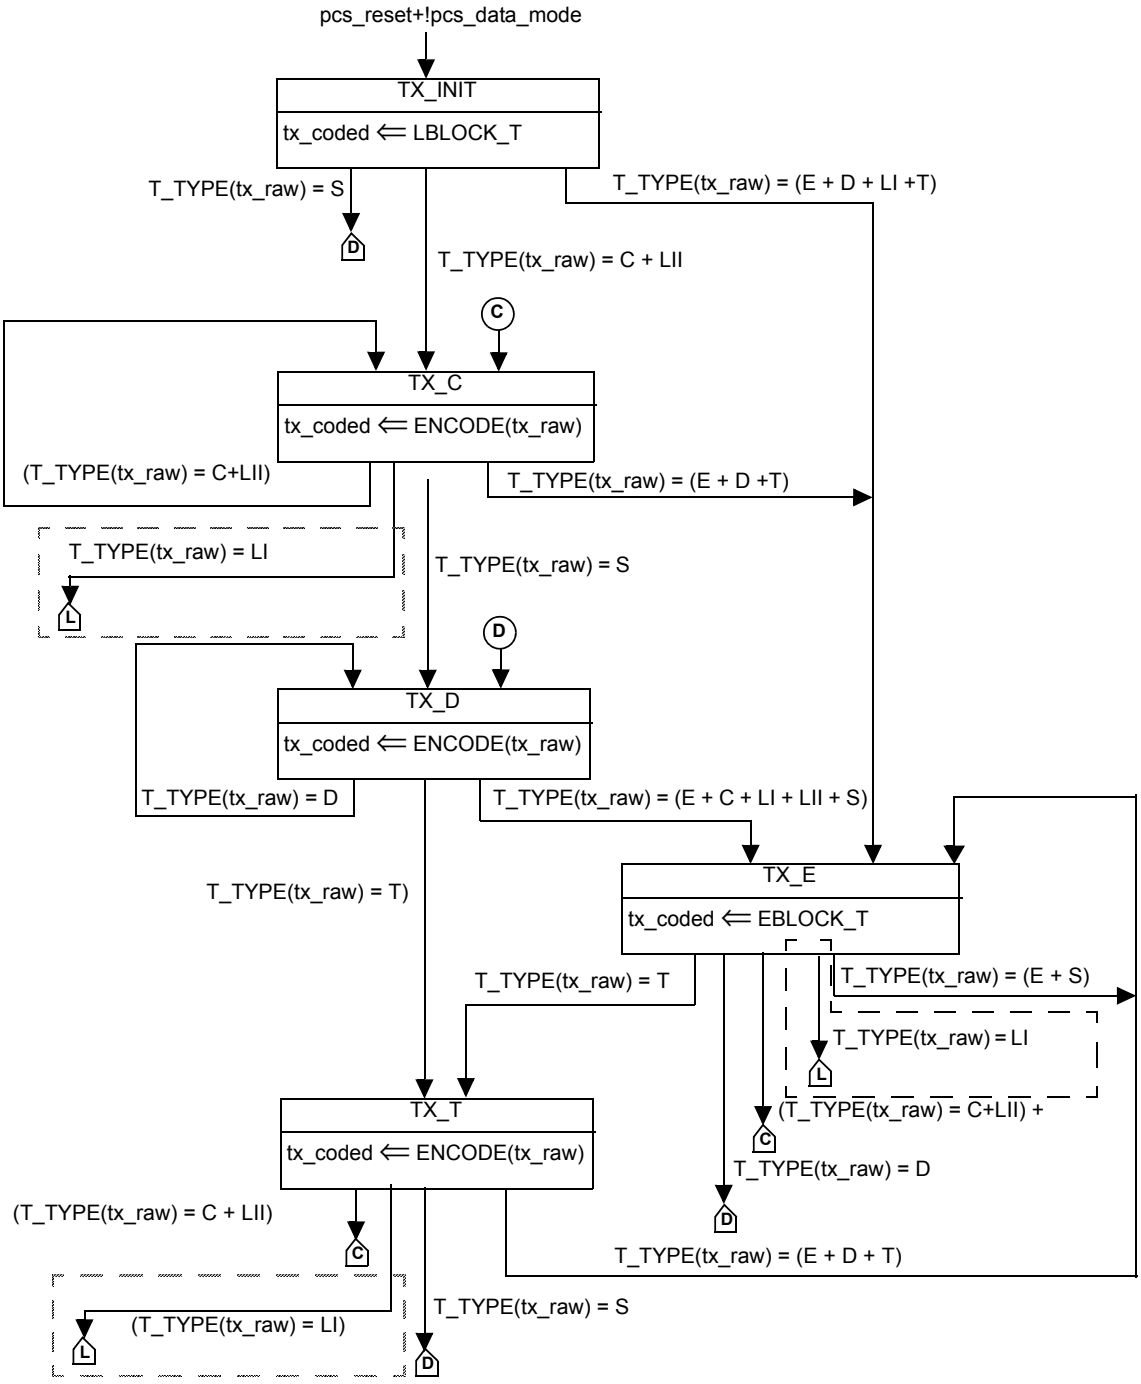

NOTE—Transitions inside dashed boxes are only required for the EEE capability.

Figure 55–16—PCS 64B/65B Transmit state diagram, part a

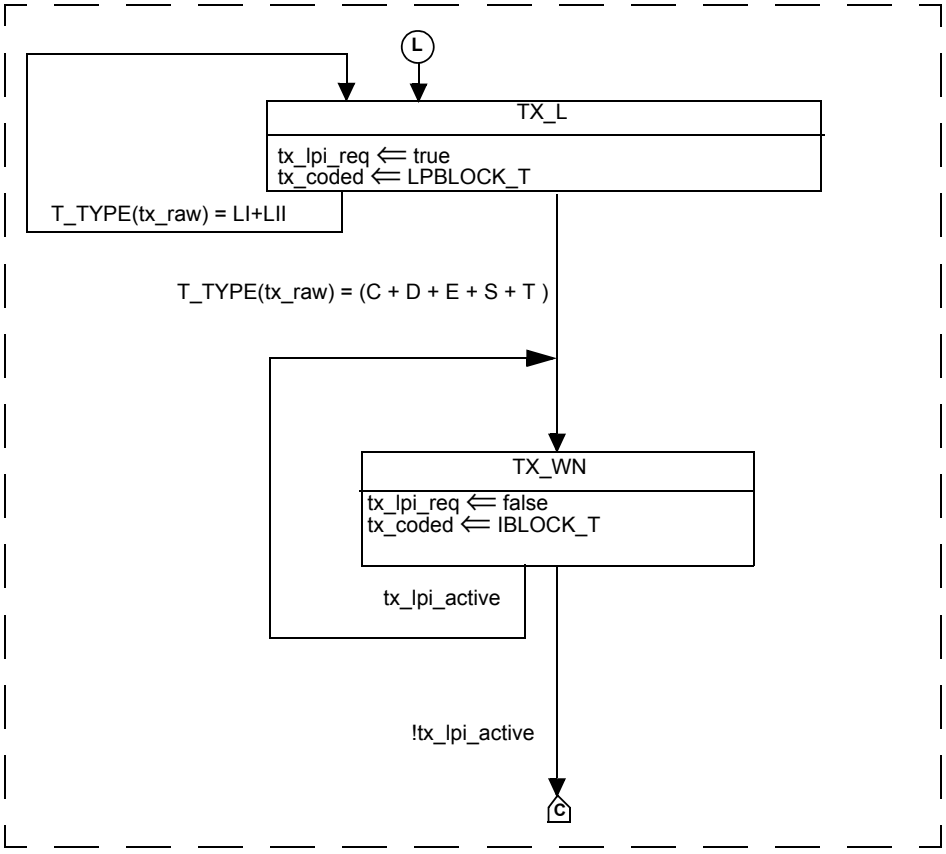

NOTE—This figure is mandatory for PHYs with the EEE capability.

Figure 55–17—PCS 64B/65B Transmit state diagram, part b

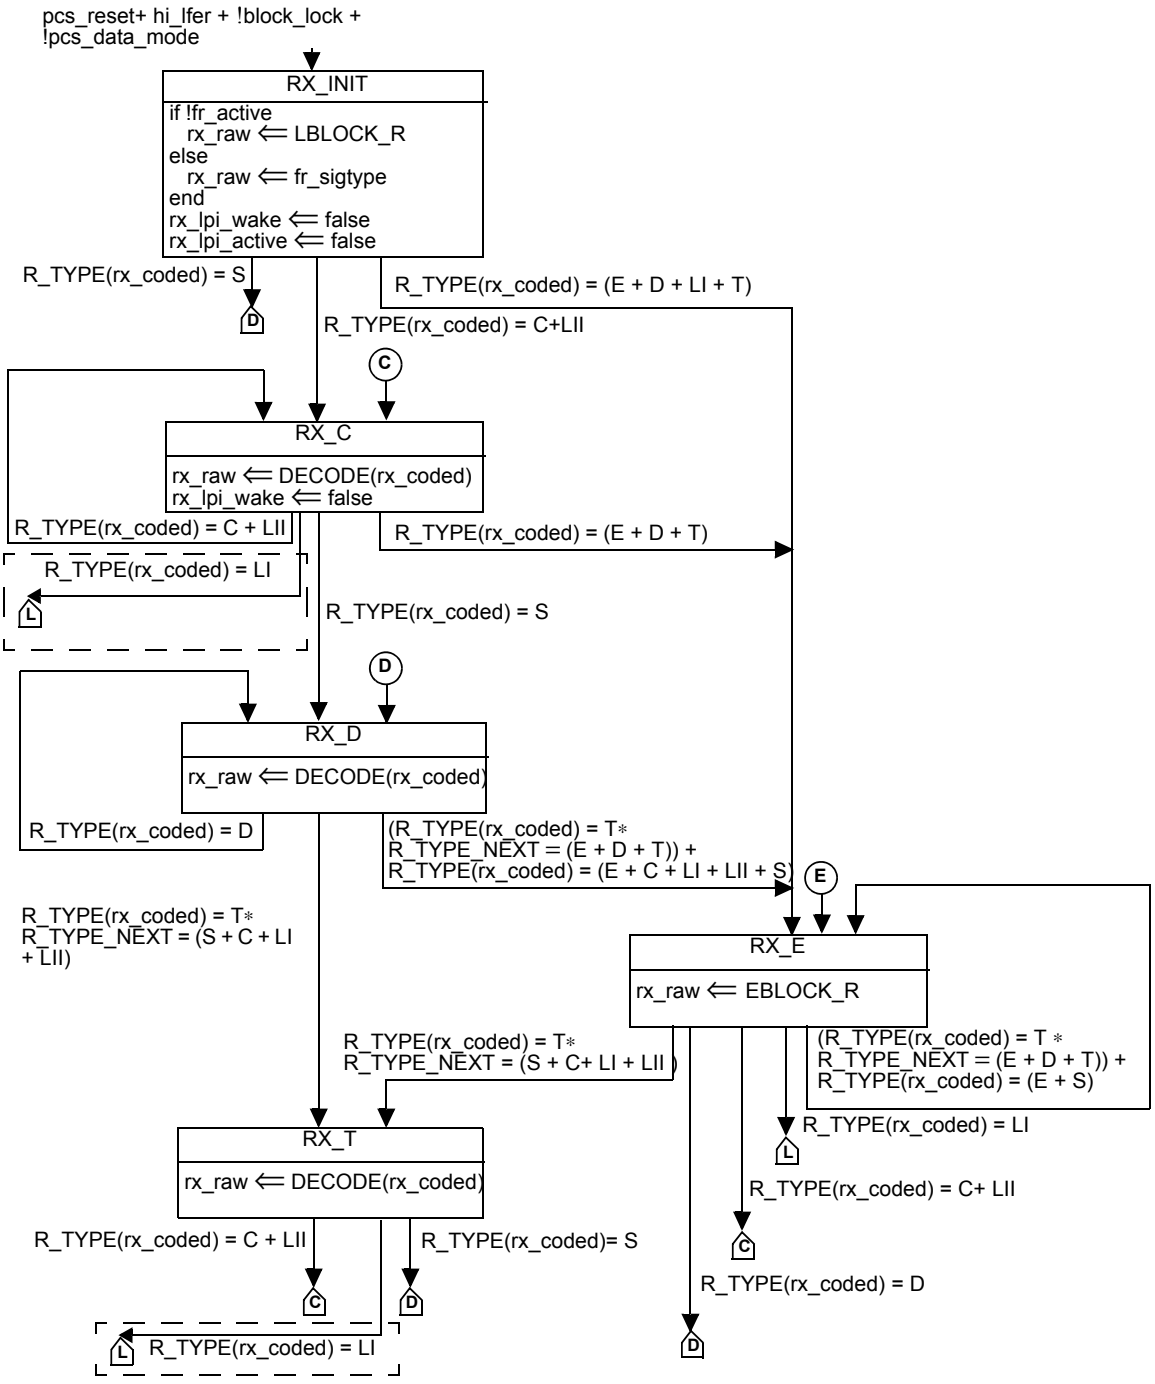

NOTE—Signals and functions shown with dashed lines are only required for the EEE capability.

Figure 55–18—PCS 64B/65B Receive state diagram, part a

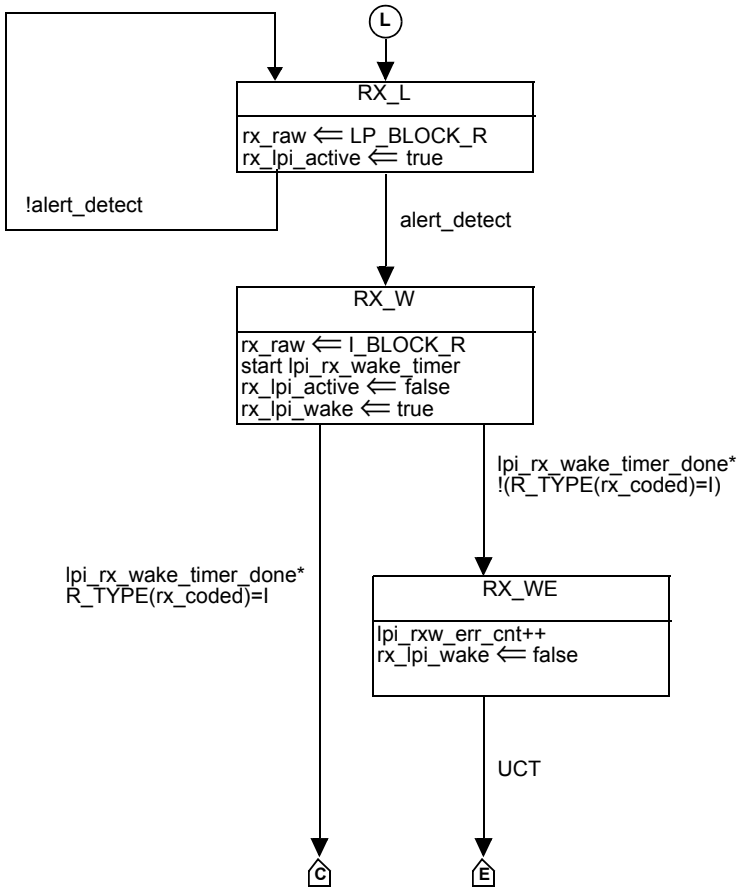

NOTE—This figure is mandatory for PHYs with the EEE capability.

Figure 55–19—PCS 64B/65B Receive state diagram, part b

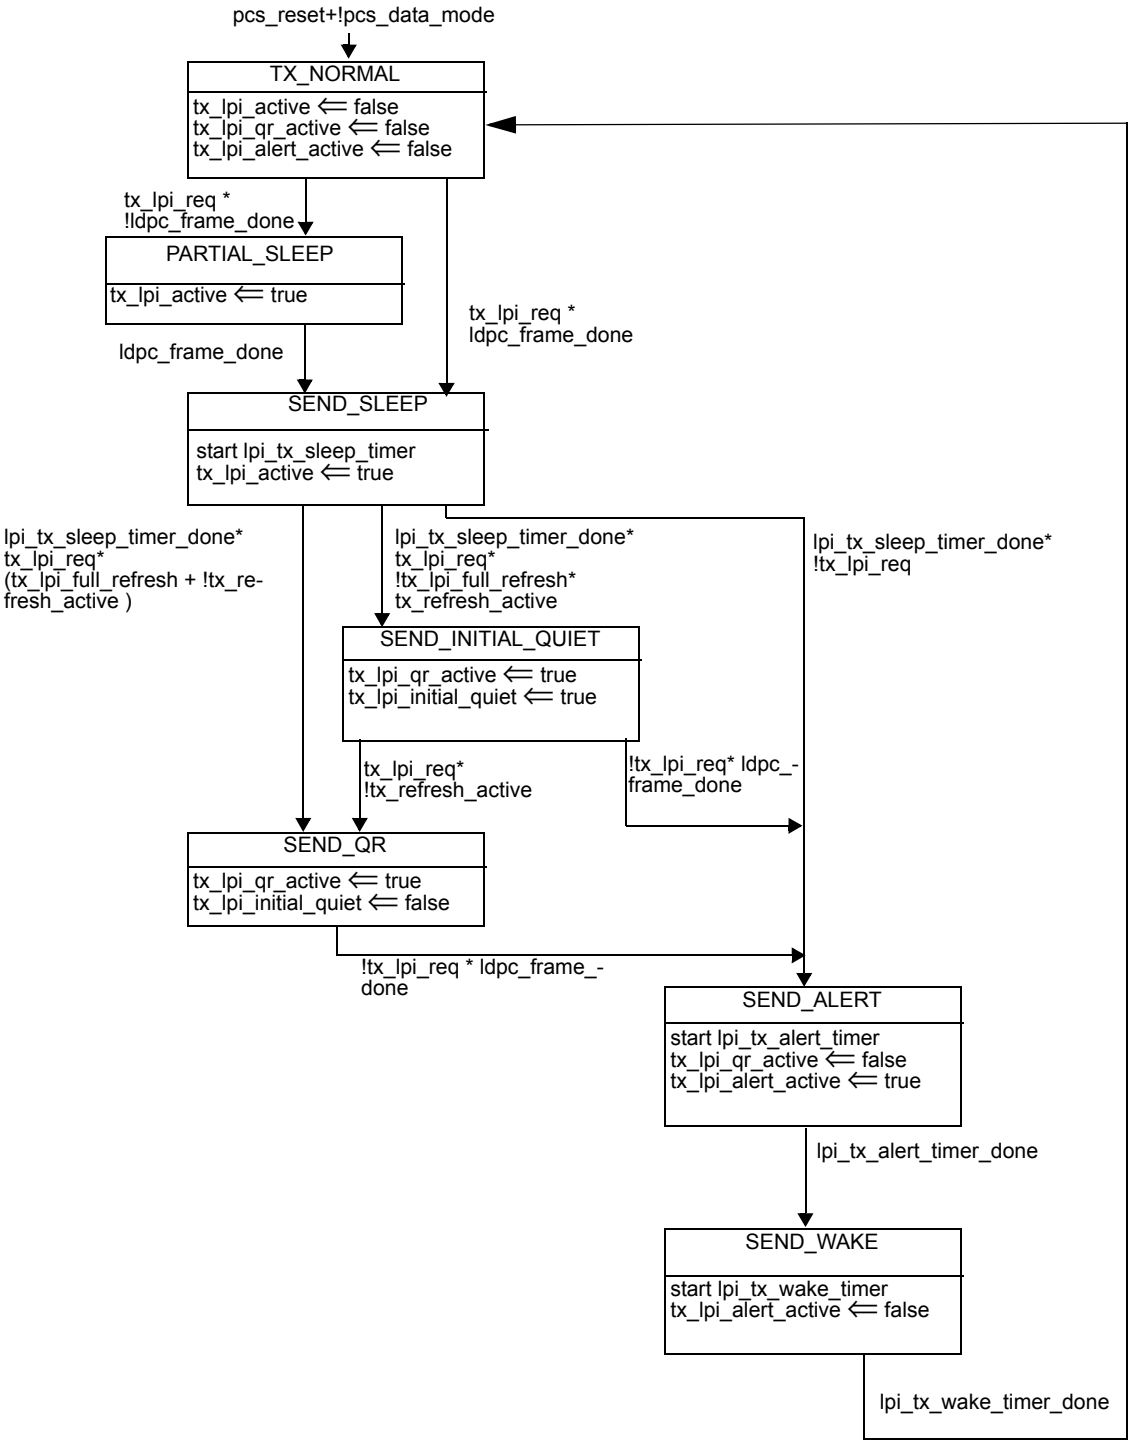

NOTE—This figure is mandatory for PHYs with the EEE capability.

Figure 55–20—EEE transmit state diagram

55.3.7.3 Loopback

The PCS shall be placed in loopback mode when the loopback bit in MDIO register 3.0.14 is set to a one. In this mode, the PCS shall accept data on the transmit path from the XGMII and return it on the receive path to the XGMII. In addition, the PCS shall transmit a continuous stream of 65B-LDPC encoded 4D-PAM16 symbols to the PMA sublayer, and shall ignore all data presented to it by the PMA sublayer.

55.4 Physical Medium Attachment (PMA) sublayer

55.4.1 PMA functional specifications

The PMA couples messages from a PMA service interface specified in 55.2.2 to the 10GBASE-T baseband medium, specified in 55.7.

The interface between PMA and the baseband medium is the Medium Dependent Interface (MDI), which is specified in 55.8.

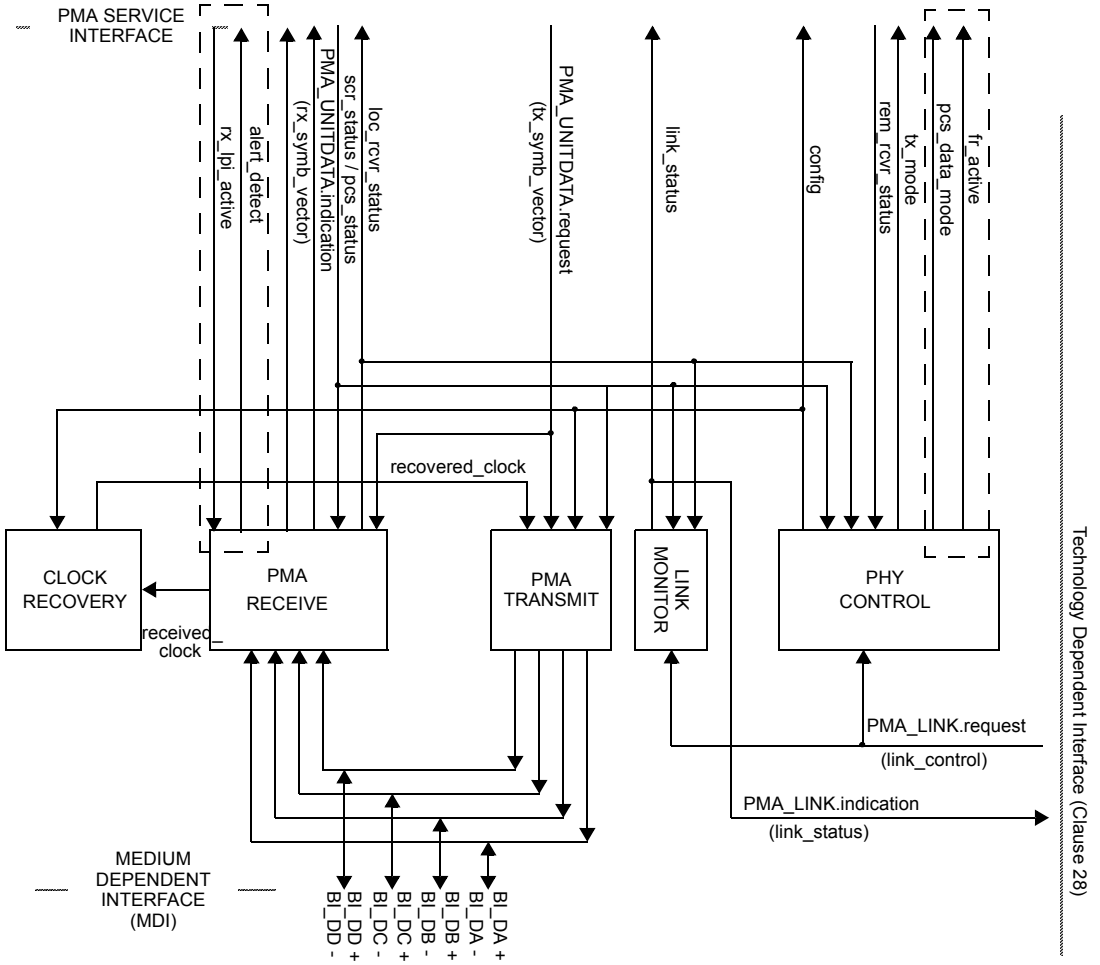

NOTE 1—The recovered_clock arc is shown to indicate delivery of the recovered clock signal back to PMA TRANSMIT for loop timing.

NOTE 2—pcs_data_mode is required only for the EEE or fast retrain capabilities alert_detect and rx_lpi_active are only required for the EEE capability fr_active is only required for the fast retrain capability

Figure 55–21— PMA reference diagram

55.4.2 PMA functions

The PMA sublayer comprises one PMA Reset function and five simultaneous and asynchronous operating functions. The PMA operating functions are PHY Control, PMA Transmit, PMA Receive, Link Monitor, and Clock Recovery. All operating functions are started immediately after the successful completion of the PMA Reset function.

The PMA reference diagram, Figure 55–21, shows how the operating functions relate to the messages of the PMA Service interface and the signals of the MDI. Connections from the management interface, comprising the signals MDC and MDIO, to other layers are pervasive and are not shown in Figure 55–21.

55.4.2.1 PMA Reset function

The PMA Reset function shall be executed whenever one of the two following conditions occur:

a) Power on (see 55.3.6.2.2)
b) The receipt of a request for reset from the management entity

All state diagrams take the open-ended pma_reset branch upon execution of PMA Reset. The reference diagrams do not explicitly show the PMA Reset function.

55.4.2.2 PMA Transmit function

The PMA Transmit function comprises four synchronous transmitters to generate four pulse-amplitude modulated signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD. While send_fail is FALSE and ALERT is not indicated by tx_symb_vector, PMA Transmit shall continuously transmit onto the MDI pulses modulated by the symbols given by tx_symb_vector[BI_DA], tx_symb_vector[BI_DB], tx_symb_vector[BI_DC], and tx_symb_vector[BI_DD], respectively after processing with the THP, optional transmit filtering, digital to analog conversion (DAC) and subsequent analog filtering. When ALERT is indicated by tx_symb_vector, the alert signal is transmitted as specified in 55.4.2.2.1. When send_fail is TRUE, the link failure signal is transmitted as specified in 55.4.2.2.2. The four transmitters shall be driven by the same transmit clock, TX_TCLK. The signals generated by PMA Transmit shall follow the mathematical description given in 55.4.3.1, and shall comply with the electrical specifications given in 55.5.

When the PMA_CONFIG.indication parameter config is MASTER, for both normal and LPI operation, the PMA Transmit function shall source TX_TCLK from a local clock source while meeting the transmit jitter requirements of 55.5.3.3. The MASTER/SLAVE relationship may include loop timing. If loop timing is implemented and the PMA_CONFIG.indication parameter config is SLAVE, the PMA Transmit function shall source TX_TCLK from the recovered clock of 55.4.2.8 while meeting the jitter requirements of 55.5.3.3. If loop timing is not implemented, the SLAVE PHY transmit clocking is identical to the MASTER PHY transmit clocking. An EEE-capable PHY shall operate with loop timing when configured as SLAVE.

The PMA Transmit fault function is optional. The faults detected by this function are implementation specific. If the MDIO interface is implemented, then this function shall be mapped to the transmit fault bit as specified in 45.2.1.7.4.

EEE-capable PHYs shall implement a PMA Transmit function that generates the alert signal as defined in 55.4.2.2.1. PHYs that support the fast retrain capability shall implement a PMA Transmit function that generates the link failure signal as defined in 55.4.2.2.2. If ALERT is indicated by tx_symb_vector at the same time as send_fail is TRUE, then link failure signaling is transmitted.

55.4.2.2.1 Alert signal

PHYs that support the optional EEE capability will transmit the following PAM2 sequence when the PMA_UNITDATA.request parameter is set to ALERT. The alert signal is sent for a total of 4 LDPC frame periods and begins on a LDPC frame boundary. The alert signal is transmitted without THP filtering. The alert signal is transmitted on pair A when the PHY operates as a MASTER. The alert signal is transmitted on pair C when the PHY operates as a SLAVE. All other pairs transmit quiet as described in 55.3.5.

When the PMA_CONFIG.indication(config) is MASTER the alert signal is composed of 7 repetitions of the following 128 symbol PAM2 sequence, followed by 128 zero symbols.

xpr_master =

9 9 -9 -9 -9 -9 -9 -9 9 9 -9 -9 9 9 9 9

9 9 9 9 -9 -9 9 9 9 9 -9 -9 9 9 -9 -9

-9 -9 -9 -9 -9 -9 9 9 -9 -9 -9 -9 -9 -9 9 9

-9 -9 -9 -9 -9 -9 -9 -9 9 9 -9 -9 9 9 -9 -9

-9 -9 9 9 9 9 9 9 9 9 9 9 -9 -9 -9 -9

9 9 -9 -9 -9 -9 9 9 9 9 -9 -9 9 9 -9 -9

-9 -9 -9 -9 -9 -9 -9 -9 9 9 9 9 -9 -9 9 9

9 9 -9 -9 9 9 -9 -9 9 9 9 9 -9 -9 -9 -9

When the PMA_CONFIG.indication(config) is SLAVE the alert signal is composed of 7 repetitions of the following 128 symbol PAM2 sequence, followed by 128 zero symbols.

xpr_slave =

-9 -9 -9 -9 9 9 9 9 -9 -9 9 9 -9 -9 9 9

9 9 -9 -9 9 9 9 9 -9 -9 -9 -9 -9 -9 -9 -9

-9 -9 9 9 -9 -9 9 9 9 9 -9 -9 -9 -9 9 9

-9 -9 -9 -9 9 9 9 9 9 9 9 9 9 9 -9 -9

-9 -9 9 9 -9 -9 9 9 -9 -9 -9 -9 -9 -9 -9 -9

9 9 -9 -9 -9 -9 -9 -9 9 9 -9 -9 -9 -9 -9 -9

-9 -9 9 9 -9 -9 9 9 9 9 -9 -9 9 9 9 9

9 9 9 9 -9 -9 9 9 -9 -9 -9 -9 -9 -9 9 9

The alert signal is followed by a wake signal composed of repeated IDLE characters encoded using the 64B/65B encoding technique. At the start of the wake signal all THP feedback delay lines are initialized with zeros.

55.4.2.2.2 Link failure signal

PHYs that support the fast retrain capability transmit the link failure signal under the control of the Fast Retrain state diagram. The link failure signal indicates to the link partner that a link failure has been detected and that the link partners should begin the fast retrain procedure.

The link failure signal is sent for 4 LDPC frames and begins on a LDPC frame boundary. The link failure signal is transmitted without THP filtering. The link failure signal is transmitted on pair A when the PHY operates as a MASTER. The link failure signal is transmitted on pair C when the PHY operates as a SLAVE. All other pairs transmit quiet as described in 55.3.5.

When the PMA_CONFIG.indication(config) is MASTER, the link failure signal is composed of 7 repetitions of the following 128 symbol PAM2 sequence, followed by 128 zero symbols.

$xfr\_master = xpr\_master \times (-1)$

When the PMA_CONFIG.indication(config) is SLAVE the link failure signal is composed of 7 repetitions of the following 128 symbol PAM2 sequence, followed by 128 zero symbols.

$xfr\_slave = xpr\_slave \times (-1)$

55.4.2.3 PMA transmit disable function

55.4.2.3.1 Global PMA transmit disable function

The Global_PMA_transmit_disable function allows all of the transmitters to be disabled, when either.

a) When a Global_PMA_transmit_disable variable is set to TRUE, this function shall turn off all of the transmitters so that the each transmitter Average Launch Power of the OFF Transmitter is less than –53 dBm.
b) If a PMA_transmit_fault is detected, then the PMA may set the Global_PMA_transmit_disable to TRUE, turning off the transmitter on each pair.

55.4.2.3.2 PMA pair by pair transmit disable function

The PMA_transmit_disable function allows the transmitters on each pair to be selectively disabled.

When a PMA_transmit_disable_N variable is set to TRUE, this function shall turn off the transmitter associated with that variable so that the transmitter Average Launch Power of the OFF Transmitter is less than –53 dBm.

55.4.2.3.3 PMA MDIO function mapping

The MDIO capability described in Clause 45 defines several variables that provide control and status information for and about the PMA. Mapping of MDIO control variables to PMA control variables is shown in Table 55–5. Mapping of MDIO status variables to PMA status variables is shown in Table 55–6.

Table 55–5—MDIO/PMA control variable mapping

MDIO control variable	PMA register name	Register/bit number	PMA control variable
Reset	Control register 1	1.0.15	PMA_reset
Global transmit disable	Transmit disable register	1.9.0	Global_PMA_transmit_disable
Transmit disable pair D	Transmit disable register	1.9.4	PMA_transmit_disable_D
Transmit disable pair C	Transmit disable register	1.9.3	PMA_transmit_disable_C
Transmit disable pair B	Transmit disable register	1.9.2	PMA_transmit_disable_B
Transmit disable pair A	Transmit disable register	1.9.1	PMA_transmit_disable_A

Table 55–6—MDIO/PMA status variable mapping

MDIO status variable	PMA register name	Register/bit number	PMA status variable
Fault	Status register 1	1.1.7	PMA_fault
Transmit fault	Status register 2	1.8.11	PMA_transmit_fault
Receive fault	Status register 2	1.8.10	PMA_receive_fault

55.4.2.4 PMA Receive function

The PMA Receive function comprises four independent receivers for pulse-amplitude modulated signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD. PMA Receive contains the circuits necessary to both detect symbol sequences from the signals received at the MDI over receive pairs BI_DA, BI_DB, BI_DC, and BI_DD and to present these sequences to the PCS Receive function. The signals received at the MDI are described mathematically in 55.4.3.2. The PMA translates the signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DB into the PMA_UNITDATA.indication parameter rx_symb_vector. The quality of these symbols shall allow LFER of less than 3.2×10^{-9} after LDPC decoding, over a channel meeting the requirements of 55.7. The receiver shall correct for differential delay variations of up to 50 ns across the wire-pairs. The delay skew is removed by computing the relative received delay of the four known transmit patterns described in 55.3.4.

To achieve the indicated performance, it is highly recommended that PMA Receive include the functions of signal equalization, echo and crosstalk cancellation. The sequence of code-groups assigned to tx_symb_vector is needed to perform echo and self near-end crosstalk cancellation.

The PMA Receive function uses the scr_status parameter and the state of the equalization, cancellation, estimation, and LPI functions to determine the quality of the receiver performance, and generates the loc_rcvr_status variable accordingly. The precise algorithm for generation of loc_rcvr_status is implementation dependent.

The receiver uses the sequence of symbols during the training sequence to detect and correct for pair swaps and crossovers. The receiver pairs BI_DA, BI_DB, BI_DC, and BI_DD may be connected in any manner described in 55.4.4 to the corresponding transmit pairs. The receiver also detects and corrects for polarity mismatches on any pairs and corrects for differential delay variations across the wire-pairs.

The PMA Receive fault function is optional. The PMA Receive fault function is the logical OR of the link_status = FAIL and any implementation specific fault. If the MDIO interface is implemented, then this function shall contribute to the receive fault bit specified in 45.2.1.7.5.

PMA receive functions that support the optional EEE capability shall generate alert_detect when the alert signal is detected at the receiver. The PMA receive function asserts alert_detect after the entire alert signal (3.5 LDPC frame periods of the xpr_master or xpr_slave sequence and 0.5 frames of silence) has been detected. The alert signal is specified in 55.4.2.2.1. The criterion used to generate alert_detect is left to the implementer.

PHYs that support the fast retrain capability shall set link_fail_detect to TRUE when the link failure signal is reliably detected at the receiver. The PMA receive function asserts link_fail_detect after the entire link failure signal (3.5 LDPC frame periods of the xfr_master or xfr_slave sequence and 0.5 frames of silence) has been detected. The link failure signal is specified in 55.4.2.2.2. The criterion used to generate link_fail_detect is left to the implementer. It is highly recommended that the generation of link_fail_detect is qualified with repeated errored frames at the LDPC decoder output.

55.4.2.5 PHY Control function

PHY Control generates the control actions that are needed to bring the PHY into a mode of operation during which frames can be exchanged with the link partner. PHY Control shall comply with the state diagram description given in Figure 55–28.

During PMA training (includes PMA_Training_Init_M, PMA_Training_Init_S, PMA_PBO_Exch, PMA_Coeff_Exch, and PMA_Fine_Adjust states in Figure 55–28), PHY Control information is exchanged between link partners with a 16 octet InfoField, which is XOR'ed with the last 128 bits of the PMA 16384 PAM2 frame on pair A (see Figure 55–13). The InfoField is also denoted IF. The link partner is not required to decode every IF transmitted but is required to decode IFs at a rate that enables the correct actions to timer expiration times, transition counter values, etc. described in Figure 55–28, Figure 55–29, and Figure 55–30.

The 16 octet InfoField shall include the fields in 55.4.2.5.2 through 55.4.2.5.13, also shown in the overview Figure 55–22, and the more detailed Figure 55–23, Figure 55–24, and Figure 55–25.

Start of Frame Delimiter 0xBBA70000	3 Transmitter Settings	Message Field	SNR Margin	Message Field Dependent	Message Field Dependent	CRC16
4 octets	3 octets	1 octet	4 bits	1.5 octets	4 octets	2 octets

Figure 55–22—InfoField format

Start of Frame Delimiter 0xBBA70000	3 Transmitter Settings	Message Field	SNR Margin	Reser-ved	Transition Counter	Reser-ved	Vendor Specific	CRC16
4 octets	3 octets	1 octet	4 bits	2 bits	10 bits	2 octets	2 octets	2 octets

Figure 55–23—InfoField transition counter format

Start of Frame Delimiter 0xBBA70000	3 Transmitter Settings	Message Field	SNR Margin	Coefficient Exchange	Coefficient Field	CRC16
4 octets	3 octets	1 octet	4 bits	1.5 octets	4 octets	2 octets

Figure 55–24—InfoField coefficient exchange format

Start of Frame Delimiter 0xBBA70000	3 Transmitter Settings	Message Field	SNR Margin	Reserved	Reserved	Vendor Specific	CRC16
4 octets	3 octets	1 octet	4 bits	1.5 octets	2 octets	2 octets	2 octets

Figure 55–25—InfoField not transition counter and not coefficient exchange format

55.4.2.5.1 Infofield notation

For all the InfoField notation below, Reserved<bit location> represents any unused values and shall be set to zero and ignored by the link partner. For all PBO InfoField values below, the PBO<6:4> are unsigned 3-bit values 000, 001, 010, 011, 100, 101, 110, and 111 shall indicate power backoffs of 0 dB, 2 dB, 4 dB, 6 dB, 8 dB, 10 dB, 12 dB, and 14 dB respectively. The InfoField is transmitted following the notation described in 55.3.2.2.3 where the LSB of each octet is sent first and the octets are sent in increasing number order (that is, the LSB of Oct1 is sent first).

55.4.2.5.2 Start of Frame Delimiter

The start of Frame Delimiter consist of 4 octets [Oct1<7:0>, Oct2<7:0>, Oct3<7:0>, Oct4<7:0>] and shall use the hexadecimal value 0xBBA70000. 0xBB corresponds to Oct1<7:0> and so forth.

55.4.2.5.3 Current transmitter settings

Current transmitter setting (1 octet). Represented by the octet Oct5{Valid<7>, PBO<6:4>, Reserved<3:0>} and shown in Figure 55–26. Used to announce the current fixed PBO setting during PMA_Training_Init_M, PMA_Training_Init_S and PMA_PBO_Exch, and the current programmable PBO setting during PMA_Coeff_Exch. For every other state this octet is set to zero and ignored by the link partner. The bit Valid shall be set to one if the corresponding octet information is valid and shall be set to zero if it the octet information is not valid. If Valid is set to zero, the octet is ignored by the link partner.

Single transmitter setting detail (one for current, next or requested)

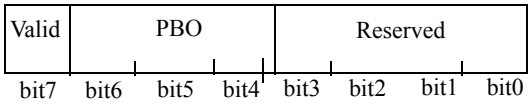

Figure 55–26—InfoField transmitter setting format

55.4.2.5.4 Next transmitter settings

Next transmitter setting (1 octet). Represented by the octet Oct6{Valid<7>, PBO<6:4>, Reserved<3:0>} and shown in Figure 55–26. Used to announce the next programmable PBO setting during PMA_PBO_Exch that takes effect upon entering PMA_Coeff_Exch state. For every other state, this octet is set to zero and ignored by the link partner. The bit Valid shall be set to one if the corresponding octet information is valid and shall be set to zero if it the octet information is not valid. If Valid is set to zero, the octet is ignored by the link partner.

55.4.2.5.5 Requested transmitter settings

Requested remote transmitter setting (1 octet). Represented by the octet Oct7{Valid<7>, PBO<6:4>, Reserved<3:0>} and shown in Figure 55–26. Used to request the remote transmitter programmable PBO setting during PMA_PBO_Exch that takes effect upon entering PMA_Coeff_Exch state. For every other state, this octet is set to zero and ignored by the link partner. The bit Valid shall be set to one if the corresponding octet information is valid and shall be set to zero if it the octet information is not valid. If Valid is set to zero, the octet is ignored by the link partner.

55.4.2.5.6 Message Field

Message Field (1 octet). For the MASTER, this field is represented by the octet Oct8{PMA_state<7:6>, loc_rcvr_status<5>, en_slave_tx<4>, trans_to_Coeff_Exch<3>, Coeff_exchange<2>, trans_to_Fine_Adjust<1>, trans_to_PCS_Test<0>}. For the SLAVE, this field is represented by the octet Oct8{PMA_state<7:6>, loc_rcvr_status<5>, timing_lock_OK<4>, trans_to_Coeff_Exch<3>, Coeff_exchange<2>, trans_to_Fine_Adjust <1>, trans_to_PCS_Test<0>}.

The two state-indicator bits PMA_state<7:6> shall communicate the state of the transmitting transceiver to the link partner. PMA_state<7:6>=00 indicates PMA_Training_Init_M or PMA_Training_Init_S, PMA_state<7:6>=01 indicates PMA_PBO_Exch, PMA_state<7:6>=10 indicates PMA_Coeff_Exch, and PMA_state<7:6>=11 indicates PMA_Fine_Adjust.

All possible Message Field settings are listed in Table 55–7 for the MASTER and Table 55–8 for the SLAVE. Any other value shall not be transmitted and shall be ignored at the receiver. The Message Field setting for the first transmitted PMA frame shall be the first row of Table 55–7 for the MASTER and the first row of Table 55–8 for the SLAVE. Moreover, for a given Message Field setting, the following Message Field setting shall be the same Message Field setting or the Message Field setting corresponding to a row below the current setting. When loc_rcvr_status=OK the InfoField variable is set to loc_rcvr_status<5>=1 and set to 0 otherwise.

Table 55–7—InfoField message field valid MASTER settings

PMA_state<7:6>	loc_rcvr_ status	en_slave_tx	trans_to_ Coeff_Exch	Coeff_ exchange	trans_to_ Fine_Adjust	trans_to_ PCS_test
00	0	0	0	0	0	0
00	0	1	0	0	0	0
01	0	1	0	0	0	0
01	0	1	1	0	0	0
10	0	1	0	0	0	0
10	0	1	0	1	0	0
10	0	1	0	0	0	0
10	0	1	0	0	1	0
11	0/1	1	0	0	0	0
11	1	1	0	0	0	1

Table 55–8—InfoField message field valid SLAVE settings

PMA_state<7:6>	loc_rcvr_ status	timing_lock _OK	trans_to_ Coeff_Exch	Coeff_ exchange	trans_to_ Fine_Adjust	trans_to_ PCS_test
00	0	0	0	0	0	0
00	0	0/1	0	0	0	0
01	0	1	0	0	0	0
01	0	1	1	0	0	0
10	0	0/1	0	0	0	0
10	0	1	0	1	0	0
10	0	1	0	0	0	0
10	0	1	0	0	1	0
11	0	0/1	0	0	0	0
11	0/1	1	0	0	0	0
11	1	1	0	0	0	1

55.4.2.5.7 SNR_margin

SNR_margin (4 bits). Represented by the half octet Oct9<7:4>, which reports received decision point SNR margin in 1/2 dB steps. SNR_margin is relative to the SNR required for reception of LDPC-coded DSQ128 at an LDPC frame error ratio of less than 3.2×10^{-9}. The SNR_margin<7:4> 4-bit values, 0010, 0011, 0100, 0101, 0110, 0111, 1000, 1001, 1010, 1011, 1100, 1101, 1110 shall indicate the decision point SNR margin

values of –1.5, –1, –0.5, 0, 0.5, 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5 dB respectively. The value 0001 shall indicate a margin of –2 dB or less, and the value 1111 shall indicate 5 dB or more. Finally the value 0000 shall indicate that the SNR margin value is unknown.

55.4.2.5.8 Transition counter

Transition counter (10 bits). Represented by the 1.25 octets [Oct9<1:0>, Oct10<7:0>]. When configured as Transition counter (Coeff_exchange<2>=0 and a transition is announced to PMA_Coeff_Exch, PMA_Fine_Adjust or PCS_Test) this field is used as a 10 bit counter that counts the number of remaining frames until the next transition (PMA_Coeff_Exch, PMA_Fine_Adjust, PCS_Test).

55.4.2.5.9 Coefficient exchange handshake

Coefficient exchange handshake (12 bits). Represented by the 1.5 octets [Oct9<3:0>, Oct10<7:0>]. If Coeff_exchange<2>=1, this field is configured as a Coefficient exchange handshake and is used as a handshake control channel during programmable THP coefficient exchange. The details of the coefficient exchange are described in 55.4.2.5.14.

55.4.2.5.10 Reserved Fields

All InfoField fields denoted Reserved in Figure 55–23, Figure 55–24, and Figure 55–25 are reserved for future use. This includes octets Oct11 and Oct12 when Coeff_exchange<2>=0, Oct9<3:2> when transition counter is announced and [Oct9<3:0>, Oct10<7:0>] when no transition is announced and no coefficients are exchanged.

55.4.2.5.11 Vendor-specific field

If Coeff_exchange<2>=0 octets, Oct13 and Oct14 are vendor-specific fields. If during Auto-Negotiation both transceivers agree on the use of the two vendor-specific octets, they may be used as a PHY communication channel; otherwise they are set to zero and ignored by the link partner. Represented by the 2 octets [Oct13<7:0>, Oct14<7:0>].

55.4.2.5.12 Coefficient Field

Coefficient Field (4 octets). Represented by the octets [Oct11<7:0>, Oct12<7:0>, Oct13<7:0>, Oct14<7:0>]. When Coeff_exchange<2>=1, this field is used to exchange programmable THP coefficients. It transmits four 8-bit THP coefficients out of the total of 64 (16 coefficients over each of the 4 pairs). The order is pair A, coefficients 0:3, followed by coefficients 4:7, followed by 8:11 and 12:15. For all cases the first coefficient (indices 0, 4, 8 and 12) is mapped to Oct11, the second coefficient (indices 1, 5, 9, 13) is mapped to Oct12 and so on. The same coefficient order is followed to transmit the coefficients for pair B, followed by pair C, and finally pair D. The details of the coefficient exchange are described in 55.4.2.5.14.

55.4.2.5.13 CRC16

CRC16 (2 octets). Shall implement the CRC16 polynomial $(x+1)(x^{15}+x+1)$ of the previous 10 octets, Oct5<7:0>, Oct6<7:0>, Oct7<7:0>, Oct8<7:0>, Oct9<7:0>, Oct10<7:0>, Oct11<7:0>, Oct12<7:0>, Oct13<7:0>, and Oct14<7:0>. The CRC16 shall produce the same result as the implementation shown in Figure 55–27. In Figure 55–27 the 16 delay elements S0,..., S15, shall be initialized to zero. Afterwards Oct5 through Oct14 are used to compute the CRC16 with the switch connected, which is setting CRCgen in Figure 55–27. After all the 10 octets have been processed, the switch is disconnected (setting CRCout) and the 16 values stored in the delay elements are transmitted in the order illustrated, first S15, followed by S14, and so on, until the final value S0.

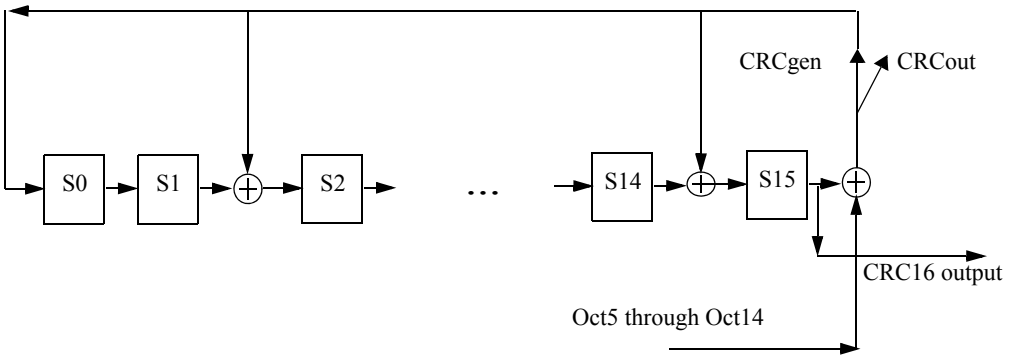

Figure 55–27—CRC16

55.4.2.5.14 Startup sequence

The startup sequence shall comply with the state diagram description given in Figure 55–28 and the transition counter state diagrams Figure 55–29 and Figure 55–30.

During Auto-Negotiation, PHY Control is in the DISABLE_10GBASE-T_TRANSMITTER state and the transmitters are disabled. During normal training, prior to enabling the transmitter, the THP coefficients are set to zero.

When the Auto-Negotiation process asserts link_control=ENABLE, PHY Control enters the INIT_MAXWAIT_TIMER state. Upon entering this state, the maxwait_timer is started and PHY Control enters the SILENT state, which starts the minwait_timer and forces transmission of zeros by setting tx_mode=SEND_Z.

In MASTER mode, after expiration of the minwait_timer, PHY Control transitions to the PMA_Training_Init_M state.

Upon entering the PMA_Training_Init_M and PMA_Training_Init_S states, the PHY Control forces transmission into the training mode by asserting tx_mode=SEND_T, which includes the transmission of InfoFields.

Upon entering state PMA_Training_Init_M, the MASTER starts transmission with a fixed transmit power level, PBO=4 (corresponding to a power backoff of 8 dB). The PBO variable is communicated to the link partner via the current transmitter octet of the InfoField.

Initially the MASTER is not ready for the SLAVE to respond and sets en_slave_tx=0, which is communicated to the link partner via the InfoField. After the MASTER has sufficiently converged the necessary circuitry, the MASTER must set en_slave_tx=1 to allow the SLAVE to transition to PMA_Training_Init_S.

In SLAVE mode, PHY Control transitions to the PMA_Training_Init_S state only after the SLAVE PHY acquires timing, converges its equalizers, acquires its descrambler state and sets loc_SNR_margin=OK. The SLAVE shall respond using the fixed PBO transmit power level, PBO=4 (corresponding to a power backoff of 8 dB). For PHYs with the EEE capability, further requirements for this transition are described in 55.3.5.1.

While in states PMA_Training_Init_S, PMA_PBO_Exch, or PMA_Coeff_Exch, whenever a SLAVE operating in loop timing mode loses the MASTER timing reference (for example, after transmit power level transitions) it sets timing_lock_OK=0, which is communicated to the link partner via the InfoField. Otherwise, timing_lock_OK is set to one.

In MASTER mode, PHY Control enters the PMA_PBO_Exch state after loc_SNR_margin=OK and in SLAVE mode PHY Control enters the PMA_PBO_Exch state after the loc_SNR_margin=OK and minwait_timer expires. In the PMA_PBO_Exch state, after the MASTER has computed the final desired programmable PBO level, it shall request a PBO change using the requested transmitter setting in the InfoField (octet 7). In SLAVE mode, after the MASTER has requested the desired PBO level, the SLAVE shall request a desired PBO level that is within two levels (within 4 dB) of the requested MASTER PBO level.

Following PBO exchange for both transceivers, each PHY shall announce the next PBO setting using the next transmitter setting (octet 6). Afterwards, each PHY announces a transition to the PMA_Coeff_Exch state using the trans_to_Coeff_Exch=1 and transition_count as described in 55.4.5.1. MASTER initiates the transition to PMA_Coeff_Exch count with the trans_to_Coeff_Exch=1 flag and a transition counter value of 2^9 (10 ms). The SLAVE responds prior to the MASTER transition counter reaching 2^6 (1 ms) by setting trans_to_Coeff_Exch=1 flag and a transition counter value matching the MASTER. The PMA frame after each transceiver transition_count reaches zero, the PHYs shall enter the PMA_Coeff_Exch state and enable the requested PBO. Therefore, both PHYs will enter the PMA_Coeff_Exch state within one PMA frame.

While both MASTER and SLAVE are in state PMA_Coeff_Exch, when either end has computed the programmable THP settings, the programmable THP coefficient exchange process can begin, using the 1.5 octet Coefficient exchange handshake and the 4 octet Coefficient Field as follows:

a) During PMA_Coeff_Exch each PHY begins a coefficient exchange by setting the Coeff_Exchange flag to 1 in the Message Field.
b) During coefficient exchange, the transition counter bits are used as the Coefficient Exchange Handshake
    1) Oct9{Reserved<3:0>}: unused
    2) Coefficient Pair Received, Oct10<7:6>: 01 for local transmitter pair A, 10 for B, 11 for C and 00 for D (default). This is the handshake to tell the remote unit the last coefficients received.
    3) Coefficient Group Received, Oct10<5:4>: 01 for coefficients 0:3, 10 for 4:7, 11 for 8:11 and 00 for 12:15 (default). This is the handshake to tell the remote unit the last coefficients received.
    4) Coefficient Pair Sent, Oct10<3:2>: 01 for remote transmitter pair A, 10 for B, 11 for C and 00 for D (default). This is the handshake to tell the remote unit the current coefficients being sent.
    5) Coefficient Group Sent, Oct10<1:0>: 01 for 0:3, 10 for 4:7, 11 for 8:11 and 00 for 12:15 (default). This is the handshake to tell the remote unit the current coefficients being sent.
c) The Coefficient Field is used to send 4 8-bit coefficients in each frame designated by the Coefficient Pair Sent and Coefficient Group Sent bits. The coefficient format is:
    1) 8 bits per coefficient. Use one octet per coefficient in twos complement notation
    2) Coefficient range is –2.0 to 1.984375 in steps of 0.015625
    3) The sign of the coefficients shall be consistent with Equation (55–4)
d) Each PHY begins the exchange by sending pair A coefficients 0:3 with Coefficient Pair Sent=01 and Coefficient Group Sent=01.
e) The remote unit acknowledges by setting Coefficient Pair Received=01 and Coefficient Group Received=01.
f) Following each acknowledgement, the PHY increments through the Coefficient Group and then Coefficient Pair settings until Coefficient Pair Sent=00 and Coefficient Group Sent=00 and Coefficient Pair Received=00 and Coefficient Group Received=00. At this time, coefficient exchange is done and both PHYs set Coeff_Exchange=0.

Following coefficient exchange for both transceivers, each PHY announces a transition to the PMA_Fine_Adjust state (trans_to_Fine_Adjust=1) and starts the transition_count as described in 55.4.5.1. During the first PMA frame after the transition_count reaches zero, the PHYs enter the PMA_Fine_Adjust state and enable the THP precoders with the requested coefficients. At the closure of the THP feedback loop, the initial state of the THP feedback filters shall be the last 16 symbols from the state PMA_Coeff_Exch.

The THP coefficients and PBO setting may not be changed during PMA_Fine_Adjust. The final convergence of the adaptive filter parameters is completed in the PMA_Fine_Adjust state.

After the PHY completes successful training and establishes proper receiver operations, PCS Transmit conveys this information to the link partner via transmission of the parameter InfoField value loc_rcvr_status. The link partner's value for loc_rcvr_status is stored in the local device parameter rem_rcvr status. When the condition loc_rcvr_status=OK and rem_rcvr_status=OK is satisfied, each PHY announces a transition to the PCS_Test state (trans_to_PCS_Test=1) and start the transition counter as described in 55.4.5.1. For PHYs with the EEE capability, further requirements for this transition are described in 55.3.5.1.

The normal mode of operation corresponds to the PCS_Data state, where PHY Control asserts tx_mode=SEND_N and transmission of data over the link can take place.

PHY Control may force the transmit scrambler state to be initialized to an arbitrary value by requesting the execution of the PCS Reset function defined in 55.3.2.1.

The operation of the maxwait_timer requires that the PHY complete the startup sequence from state SILENT to PMA_Fine_Adjust in the PHY Control state diagram (Figure 55–28) in less than 2000 ms to avoid link_status being changed to FAIL by the Link Monitor state diagram (Figure 55–31). To ensure interoperability the timing in Table 55–9 should be observed.

After reaching the PCS_Data state PHYs with the EEE capability can transition to the LPI receive mode under the control of the link partner and to the LPI transmit mode under control of the local LPI client.

Table 55–9—Recommended startup sequence timing

Master	Recommended maximum time (ms)	Recommended average time (ms)	Slave
SILENT plus (PMA_Training_Init_M state AND en_slave_tx = 0)	350	315	SILENT
(PMA_Training_Init_M state AND en_slave_tx = 1) plus PMA_PBO_Exch state	480	432	PMA_Training_Init_S state plus PMA_PBO_Exch state
PMA_Coeff_Exch state	100	90	PMA_Coeff_Exch state with timing_lock_OK=0
	520	468	Total for PMA Coeff Exch state
PMA_Fine_Adjust state	650	585	PMA_Fine_Adjust state
Total	2000	1800	

55.4.2.5.15 Fast retrain function

PHYs that support the fast retrain capability shall implement the fast retrain state diagram shown in Figure 55–33. PHYs may request a fast retrain by setting the variable loc_fr_req to TRUE. This causes the transmission of an easily-detected link failure signal specified in 55.4.2.2.2. After completing the link failure signal the PHY shall transition to the PMA_Coeff_Exch state, keep its THP turned on with its previously exchanged coefficients, and send PAM2 signaling within a time period equivalent to 9 LDPC frame periods.

After the detection of the link failure signal, a PHY shall transition to the PMA_Coeff_Exch state and respond with PAM2 signaling within a time period equivalent to 9 LDPC frame periods after receiving the link failure signal.

The PAM2 symbols are generated using the PMA sidestream scrambler polynomials shown in Figure 55–13. The training sequence without periodic re-initialization described in 55.3.4 shall be used during fast retraining, with the scramblers free-running from PCS Reset. If scrambler re-initialization is used for normal training, it shall be disabled and the scramblers shall begin free-running when the PHY Control state diagram enters the PCS_Test state and the variable fr_active is FALSE.

Note that reliable traffic on the transmitter may be interrupted when the local receiver requests a fast retrain.

Following the link failure signal, the two link partners transition back to the PMA_Coeff_Exch state and follow the training procedure described in 55.4.2.5.14, with the exception that the initial infofield countdown values are reduced as indicated in Figure 55–29 and Figure 55–30.

To ensure interoperability the training times in Table 55–10 should be observed during the fast retrain.

Table 55–10—Recommended fast retrain sequence timing

State	Recommended maximum time (ms)
PMA_Coeff_Exch state	20
PMA_Fine_Adjust state	10

55.4.2.6 Link Monitor function

Link Monitor determines the status of the underlying receive channel and communicates it via the variable link_status. Failure of the underlying receive channel typically causes the PMA's clients to suspend normal operation.

The Link Monitor function shall comply with the state diagram of Figure 55–31.

Upon power on, reset, or release from power down, the Auto-Negotiation algorithm sets link_control=SCAN_FOR_CARRIER and, during this period, sends fast link pulses to signal its presence to a remote station. If the presence of a remote station is sensed through reception of fast link pulses, the Auto-Negotiation algorithm sets link_control=DISABLE and exchanges Auto-Negotiation information with the remote station. During this period, link_status=FAIL is asserted. If the presence of a remote 10GBASE-T station is established, the Auto-Negotiation algorithm permits full operation by setting link_control=ENABLE. As soon as reliable transmission is achieved, the variable link_status=OK is asserted, upon which further PHY operations can take place.

55.4.2.7 Refresh Monitor function

The Refresh monitor is required for PHYs that support the EEE capability. The Refresh monitor operates when the PHY is in the LPI receive mode. The Refresh monitor shall comply with the state diagram of Figure 55–19. The function forces a link retrain if a refresh signal is not reliably detected within a moving time window equivalent to 50 complete LPI cycles (nominally equal to 8.192 ms), when the PHY is in the lower power receive mode.

55.4.2.8 Clock Recovery function

The Clock Recovery function couples to all four receive pairs. It may provide independent clock phases for sampling the signals on each of the four pairs.

The Clock Recovery function shall provide clocks suitable for signal sampling on each line so that the LDPC FER indicated in 55.4.2.4 is achieved. The received clock signal should be stable and ready for use when training has been completed (loc_rcvr_status=OK). The received clock signal is supplied to the PMA Transmit function by received_clock.

55.4.3 MDI

Communication through the MDI is summarized in 55.4.3.1 and 55.4.3.2.

55.4.3.1 MDI signals transmitted by the PHY

The symbols to be transmitted by the PMA on the four pairs BI_DA, BI_DB, BI_DC, and BI_DD are denoted by tx_symb_vector[BI_DA], tx_symb_vector[BI_DB], tx_symb_vector[BI_DC], and tx_symb_vector[BI_DD], respectively. The modulation scheme used over each pair is PAM16. PMA Transmit generates a pulse-amplitude modulated signal on each pair in the following form:

$$x_n = M(a_n - \sum_{k=1}^{16} x_{n-k}c_k) = a_n + 32m_n - \sum_{k=1}^{16} x_{n-k}c_k \quad (55\text{–}4)$$

$$s(t) = \sum_{n=0}^{\infty} x_n h_T(t - nT) \quad (55\text{–}5)$$

In Equation (55–4), a_n is the PAM16 modulation symbol from the set {–15, –13, –11, –9, –7, –5, –3, –1, 1, 3, 5, 7, 9, 11, 13, 15} to be transmitted at time nT. Each of the 16 THP coefficients c_1, c_2,..., c_{16} per wire pair is represented in two's complement form by 8 bits described in 55.4.2.5. The nonlinear THP operation given by $M(\alpha) = (\alpha + 16) mod 32 - 16$ corresponds to changing the modulation symbol a_n to an augmented modulation symbol $\tilde{a}_n = a_n + 32m_n$ with the integer m_n chosen such that the THP output lies in the interval $-16 \le x_n < 16$. Equation (55–5)describes the convolution of the THP output signals with the transmitter symbol response $h_T(t)$ to obtain the transmit signal $s(t)$ at the MDI. The values of the programmable THP coefficients are exchanged in the InfoField during PMA_Coeff_Exch. The THP filter coefficients shall be fixed after startup.

The nominal power (denoted Ptx) and the symbol response of the PMA transmitted signal $s(t)$, shall comply with the electrical specifications given in 55.5. When the link segment does not experience the maximum insertion loss (IL), each transceiver indicates to the link partner that the link partner PMA Transmit signal shall be reduced in increments of 2 dB. The minimum power backoff level requested shall comply with the power backoff schedule in Table 55–11. If a given receiver has sufficient decision point SNR margin, it may choose to request from the link partner larger power backoff (up to 14 dB) than shown in Table 55–11. Additionally, the Slave shall select a PBO level as described in the PMA_PBO_Exch state of 55.4.2.5.14. The PMA Transmit shall be capable of eight power backoff settings in approximately 2 dB

steps. The difference between each consecutive power setting shall be 2 ± 0.25 dB, and each step shall be centered at 2 × *n* dB (*n* = 0 to 7) reduction from nominal, with a maximum error of ± 1 dB.

The received signal power at the MDI, *P* (dBm), in Table 55–11, should be the estimate of the average received power across all four pairs from the remote transmitter when the link partner PMA Transmit is at nominal power (after accounting for local transmitter power). If the remote transmitter is not at nominal power during the measurement, the estimate of the received power should be incremented by the amount of power backoff of the link partner transmitter during the measurement. Nominal power refers to the transmit power without any power backoff and is specified in 55.5.3.4. The estimate of the received signal power is stored in registers 1.141 to 1.144 as described in 45.2.1. The values in the length, *L* (m), column in Table 55–11 are for reference only (not required for power backoff evaluation) and have been computed using the scaled insertion loss equation in 55.7.

Table 55–11—Power backoff schedule

Received signal power at MDI, *P* (dBm)	Length *L*(m) (reference)	Minimum power backoff (dB)
$-1.1 < P$	$0 \le L < 35$	10
$-2.3 < P \le -1.1$	$35 \le L < 45$	8
$-3.3 < P \le -2.3$	$45 \le L < 55$	6
$-4.2 < P \le -3.3$	$55 \le L < 65$	4
$-5.0 < P \le -4.2$	$65 \le L < 75$	2
$P \le -5.0$	$75 \le L$	0

55.4.3.2 Signals received at the MDI

Signals received at the MDI can be expressed for each pair as pulse-amplitude modulated signals that are corrupted by noise as follows:

$$r(t) = \sum_{n=0}^{\infty} \tilde{a}_n h_R(t - nT) + w(t) \quad (55\text{–}6)$$

In Equation (55–6), $\tilde{a}_n$ are the augmented PAM16 modulation symbols described in 55.4.3.1, $h_R(t)$ denotes the symbol response of the overall channel from the THP precoder to the MDI at the receiver and $w(t)$ represents the contribution of various noise sources including uncancelled crosstalk. The four signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DD are processed within the PMA Receive function to yield the received symbols rx_symb_vector.

55.4.4 Automatic MDI/MDI-X configuration

Automatic MDI/MDI-X configuration is intended to eliminate the need for crossover cables between similar devices. Automatic MDI/MDI-X configuration is required for 10GBASE-T devices and shall comply with 40.4.4.1 and 40.4.4.2.

Having established MDI/MDI-X configuration, the receiver shall detect and correct for several configurations of pair swaps and crossovers and arbitrary polarity swaps. The receiver pairs BI_DA, BI_DB,

BI_DC, and BI_DD might be connected to the corresponding transmit pairs in any of the following ways with arbitrary polarity:

a) No crossover
b) A/B crossover only
c) C/D crossover only
d) A/B crossover and C/D crossover

For EEE-capable PHYs, the MDI/MDIX function configuration shall apply to refresh and alert signaling. For PHYs with the fast retrain capability, the MDI/MDIX function configuration shall apply to link failure signaling.

55.4.5 State variables

55.4.5.1 State diagram variables

coeff_exchange_done
: This variable reports that both transceivers have received the corresponding coefficients from the link partner.
Values: TRUE: The coefficient exchange has completed.
FALSE: The coefficient exchange has not completed.

config
: The PMA shall generate this variable continuously and pass it to the PCS via the PMA_CONFIG.indication primitive.
Values: MASTER or SLAVE

link_control
: This variable is defined in 28.2.6.2.

link_status
: The link_status parameter set by PMA Link Monitor and passed to the PCS via the PMA_LINK.indication primitive.
Values: OK or FAIL

loc_rcvr_status
: Variable set by the PMA Receive function to indicate correct or incorrect operation of the receive link for the local PHY.
Values: OK: The receive link for the local PHY is operating reliably.
NOT_OK: Operation of the receive link for the local PHY is unreliable.

loc_SNR_margin
: This variable reports whether the local device has sufficient SNR margin to continue to the next state. The criterion for setting the parameter loc_SNR_margin is left to the implementer.
Values: OK: The local device has sufficient SNR margin.
NOT_OK: The local device does not have sufficient SNR margin.

master_transition_counter
: This variable reports the current value of the MASTER's transition counter reported in the InfoField defined in 55.4.2.5.
Values: 0 to 2^9

MessageField_IF
: This variable reports that a receiver has successfully received and decoded the InfoField

from the remote device. This variable takes on the value contained in the Message Field. If the Message Field cannot be decoded or no explicit action is outstanding the value Null is returned.
Values: trans_to_Coeff_Exch, trans_to_Fine_Adjust, trans_to_PCS_Test or Null

PBO

PBO is a variable that can take any integer value from 0 to 7 and indicates the power backoff level. Denoting Ptx as the maximum nominal power, the PBO values are:
Values: 0, 1, 2, 3, 4, 5, 6, 7, which correspond to transmit power levels of *Ptx*, *Ptx*–2 dB, *Ptx*–4 dB, *Ptx*–6 dB, *Ptx*–8 dB, *Ptx*–10 dB, *Ptx*–12 dB, *Ptx*–14 dB respectively

PBO_next

PBO_next is a variable that can take any integer value from 0 to 7 and indicates the next power backoff level to be used at the local transmitter. The value is taken from the fixed set of values during PMA_Training_Init_M and PMA_Training_Init_S as described in 55.4.2.5. The value is taken from the decoded value of the link partner InfoField during PMA_PBO_Exch
Values: 0, 1, 2, 3, 4, 5, 6, 7, which correspond to transmit power levels of *Ptx*, *Ptx*–2 dB, *Ptx*–4 dB, *Ptx*–6 dB, *Ptx*–8 dB, *Ptx*–10 dB, *Ptx*–12 dB, *Ptx*–14 dB respectively

PBO_tx

PBO_tx is a variable that can take any integer value from 0 to 7 and indicates the power backoff level currently used at the local transmitter.
Values: 0, 1, 2, 3, 4, 5, 6, 7, which correspond to transmit power levels of *Ptx*, *Ptx*–2 dB, *Ptx*–4 dB, *Ptx*–6 dB, *Ptx*–8 dB, *Ptx*–10 dB, *Ptx*–12 dB, *Ptx*–14 dB respectively

PBO_exchange_done

This variable reports that both transceivers have received the corresponding PBO levels from the link partner.
Values: TRUE: The PBO exchange has completed.
FALSE: The PBO exchange has not completed.

pma_reset

Allows reset of the PHY Control and Link Monitor state diagrams.
Values: ON or OFF

rem_rcvr_status

Variable set by the PCS Receive function to indicate whether correct operation of the receive link for the remote PHY is detected or not.
Values: OK: The receive link for the remote PHY is operating reliably.
NOT_OK: Reliable operation of the receive link for the remote PHY is not detected.

THP_next

THP is a variable that contains sixteen 8-bit values and describes the next transmitter setting of the THP coefficients. It refers to the programmable THP coefficients selected during coefficient exchange described in 55.4.2.5.
Values: 16 coefficients of 8-bit values each. Range is –2.0 to 1.984375 in steps of 0.015625

THP_tx

THP is a variable that contains sixteen 8-bit values and describes the current transmitter setting of the THP coefficients. It refers to the programmable THP coefficients selected during

coefficient exchange described in 55.4.2.5.
Values: 16 coefficients of 8-bit values each. Range is –2.0 to 1.984375 in steps of 0.015625

trans_to_Coeff_Exch
Message field variable defined in 55.4.2.5 that flags a transition by the local device to the PMA_Coeff_Exch state.
Values: 1: The local device transitions to the PMA_Coeff_Exch state on expiration of the transition counter.
0: The local device does not transition to the PMA_Coeff_Exch state.

trans_to_Fine_Adjust
Message field variable defined in 55.4.2.5 that flags a transition by the local device to the PMA_Fine_Adjust state.
Values: 1: The local device transitions to the PMA_Fine_Adjust state on expiration of the transition counter.
0: The local device does not transition to the PMA_Fine_Adjust state.

trans_to_PCS_Test
Message field variable defined in 55.4.2.5 that flags a transition by the local device to the PCS_test state.
Values: 1: The local device transitions to the PCS_test state on expiration of the transition counter.
0: The local device does not transition to the PCS_test state.

transition_count
This variable reports the value of the transition counter contained in the InfoField sent to the remote device. Transition_count must comply with the state diagram description given in 55.4.6.2. When the message field contains a flag for a state transition, the transition counter denotes the remaining number of InfoField until the next state transition. MASTER initiates the transition to PMA_Coeff_Exch count with the trans_to_Coeff_Exch=1 flag and a counter value of 2^9 (10 ms). The SLAVE responds prior to the counter reaching 2^6 (1 ms) with the same flag and a count value matching the MASTER. Then both PHY's transition to PMA_Coeff_Exch within one PMA frame. The same sequence is performed in the transition to PMA_Fine_Adjust state and PCS_Test state using the trans_to_Fine_Adjust=1 and trans_to_PCS_Test=1 flags respectively. In EEE-capable PHYs, synchronization of the PMA frames is tightly controlled as described in 55.3.5.1. When the message field does not contain a flag for a state transition, the transition counter is set to zero and ignored by the receiver.
Values: 0 to 2^9

tx_mode
PCS Transmit sends code-groups according to the value assumed by this variable.
Values: SEND_N: This value is continuously asserted when transmission of sequences of code-groups representing a XGMII data stream take place.

SEND_T: This value is continuously asserted when transmission of sequences of code-groups representing the sequences of code-groups (TA_n, TB_n, TC_n, TD_n) defined in 55.3.4.2 is to take place.

SEND_Z: This value is asserted when transmission of zero code-groups is to take place.

The following variables are required only for PHYs that support the EEE capability:

lpi_refresh_detect
: Set TRUE when the receiver has reliably detected refresh signaling and FALSE otherwise. The exact criteria left to the implementer.

pcs_data_mode
: Generated by the PMA PHY Control function and indicates whether or not the local PHY may transition its PCS state diagrams out of their initialization states. The current value of the pcs_data_mode is passed to the PCS via the PMA_PCSDATAMODE.indicate primitive. In the absence of the optional EEE and fast retrain capabilities, the PHY operates as if the value of this variable is TRUE.

mtc
: mtc is the transition count for a MASTER PHY during normal training and fast retraining. mtc shall be equal to 2^9 for normal training and 2^5 for fast retrain.

stc
: stc is the transition count for a SLAVE PHY during normal training and fast retraining. stc shall be equal to 2^6 for normal training and 2^4 for fast retrain.

The following six variables are required only for PHYs that support the fast retrain capability:

fr_enable
: This variable is set to TRUE if 1.147.0 is set to 1 and fast retrain is supported. The variable is set to FALSE otherwise.

loc_fr_req
: Set TRUE when the receiver has detected a link failure condition and is requesting a fast retrain; set FALSE otherwise.

loc_fr_detect
: Set TRUE when the receiver has reliably detected the link failure signal. It is highly recommended that loc_fr_detect is qualified with the reception of errored blocks at the LDPC decoder output. Set FALSE when the link failure signal is not detected.

send_link_fail
: When TRUE indicates that the PMA should send the link failure signal. When FALSE the variable has no effect.

fr_active
: Set TRUE when the PHY is performing a fast retrain and set FALSE otherwise.

fast_retrain_flag
: Set TRUE after the PHY generates or detects a link failure signal and set FALSE otherwise.

55.4.5.2 Timers

All timers operate in the manner described in 14.2.3.2.

maxwait_timer
: A timer used to limit the amount of time during which a receiver dwells in the SILENT and TRAINING states. The timer shall expire 2000 ms ± 10 ms after being started. This timer is used jointly in the PHY Control and Link Monitor state diagrams. The

maxwait_timer is tested by the Link Monitor to force link_status to be set to FAIL if the timer expires and loc_rcvr_status is NOT_OK. See Figure 55–28 and Figure 55–31.

minwait_timer
: A timer used to determine the minimum amount of time the PHY Control stays in the SILENT, PMA_Training_Init_S, PCS_Test and PCS_Data states. The timer shall expire 1 ms ± 0.1 ms after being started.

The following timer is required only for PHYs that support the EEE capability:

lpi_refresh_rx_timer
: This timer is used to monitor link quality during the LPI receive mode. If the PHY does not reliably detect reliable refresh signaling before this timer expires then a full retrain is performed.
Values: The condition lpi_refresh_rx_timer_done becomes true upon timer expiration.
Duration: This timer shall have a period equal to 50 complete quiet-refresh signal periods, equivalent to 8.192 ms.

The following two timers are required only for PHYs that support the fast retrain capability:

link_fail_sig_timer
: Determines the period of time the PHY sends the link failure signal.
Values: The condition link_fail_sig_timer_done becomes true upon timer expiration.
Duration: This timer shall have a period equal to 4 LDPC frame periods.

fr_maxwait_timer
: Determines the period of time the PHY has to transition its PCS Control State to PCS_Test following a fast retrain before the fast retrain is aborted and a full retrain performed.
Values: The condition fr_maxwait_timer_done becomes true upon timer expiration.
Duration: This timer shall have a period equal to 30 ms.

55.4.5.3 Functions

Exchange_Final_PBO
: This function transmits and receives the final PBO settings using the InfoField as described in 55.4.2.5.

Exchange_THP_Coefficients
: This function compiles and sends to the link partner and receives from the link partner the desired programmable THP coefficients using the InfoField as described in 55.4.2.5.

55.4.5.4 Counters

The following two counters are required only for PHYs that support the fast retrain capability:

fr_tx_counter
: Counts the number of times the PHY initiates a fast link retrain by transmitting the link failure signal. This counter is reflected in MDIO register 1.147.10:6 specified in 45.2.1.79.2.

fr_rx_counter
: Counts the number of times the PHY begins a fast link retrain in response to the detection of link failure signaling from the link partner. This counter is reflected in MDIO register 1.147.15:11 specified in 45.2.1.79.1.

55.4.6 State diagrams

55.4.6.1 PHY Control state diagram

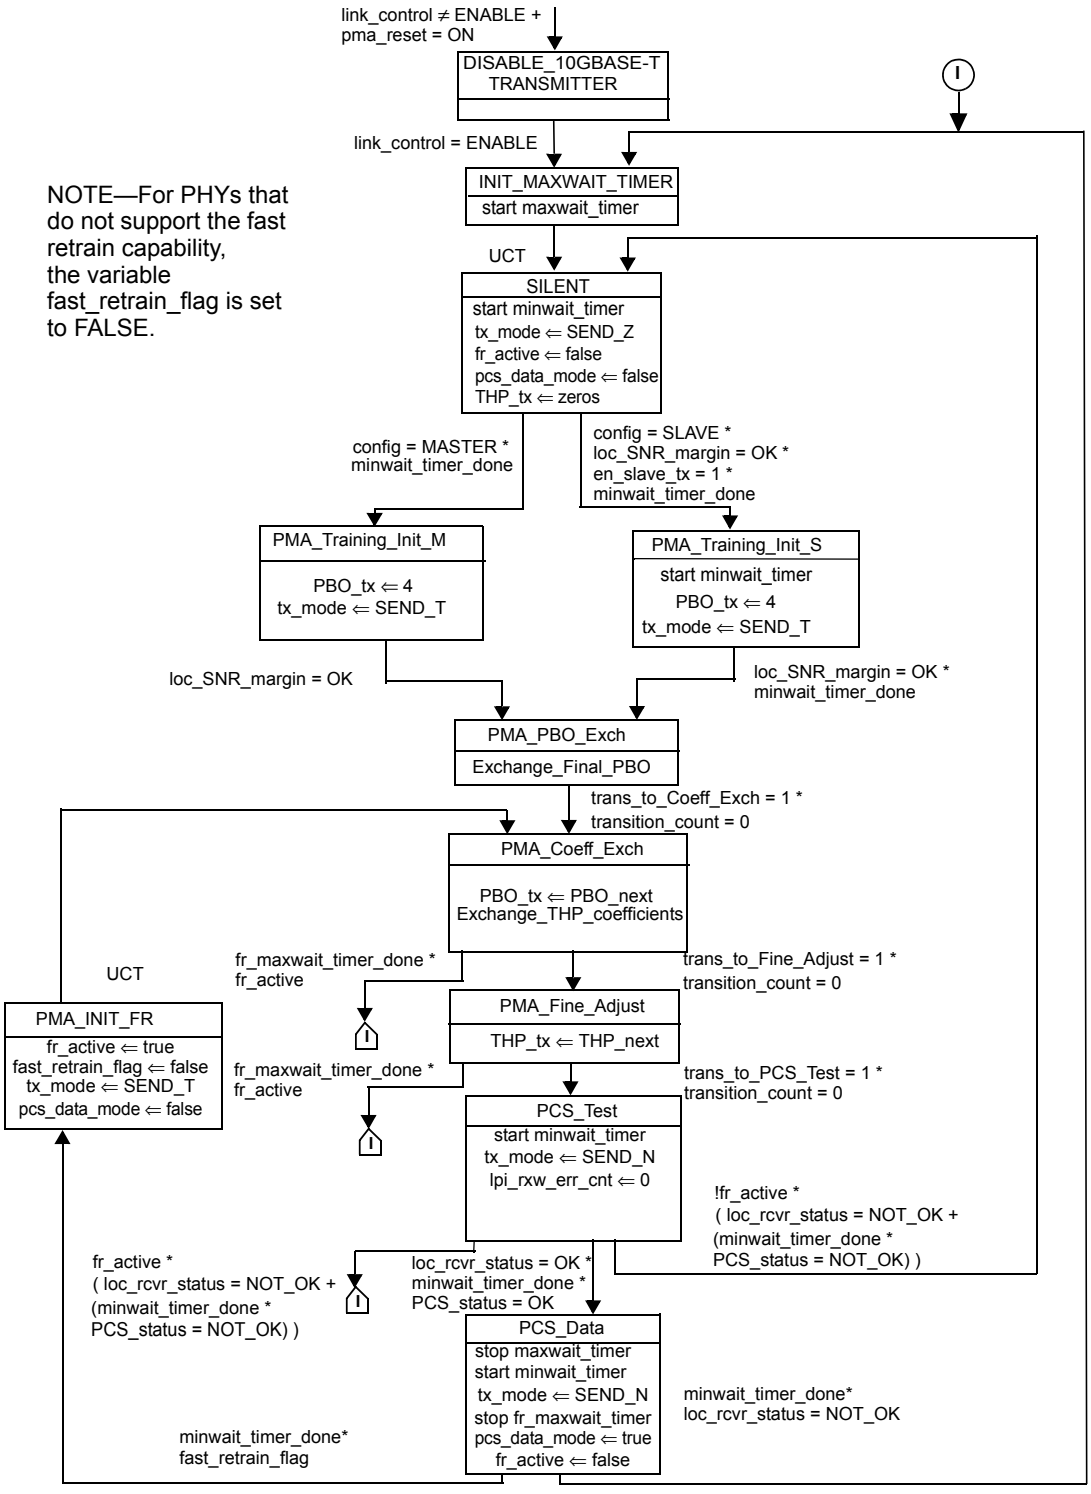

Figure 55–28—PHY Control state diagram

55.4.6.2 Transition counter state diagrams

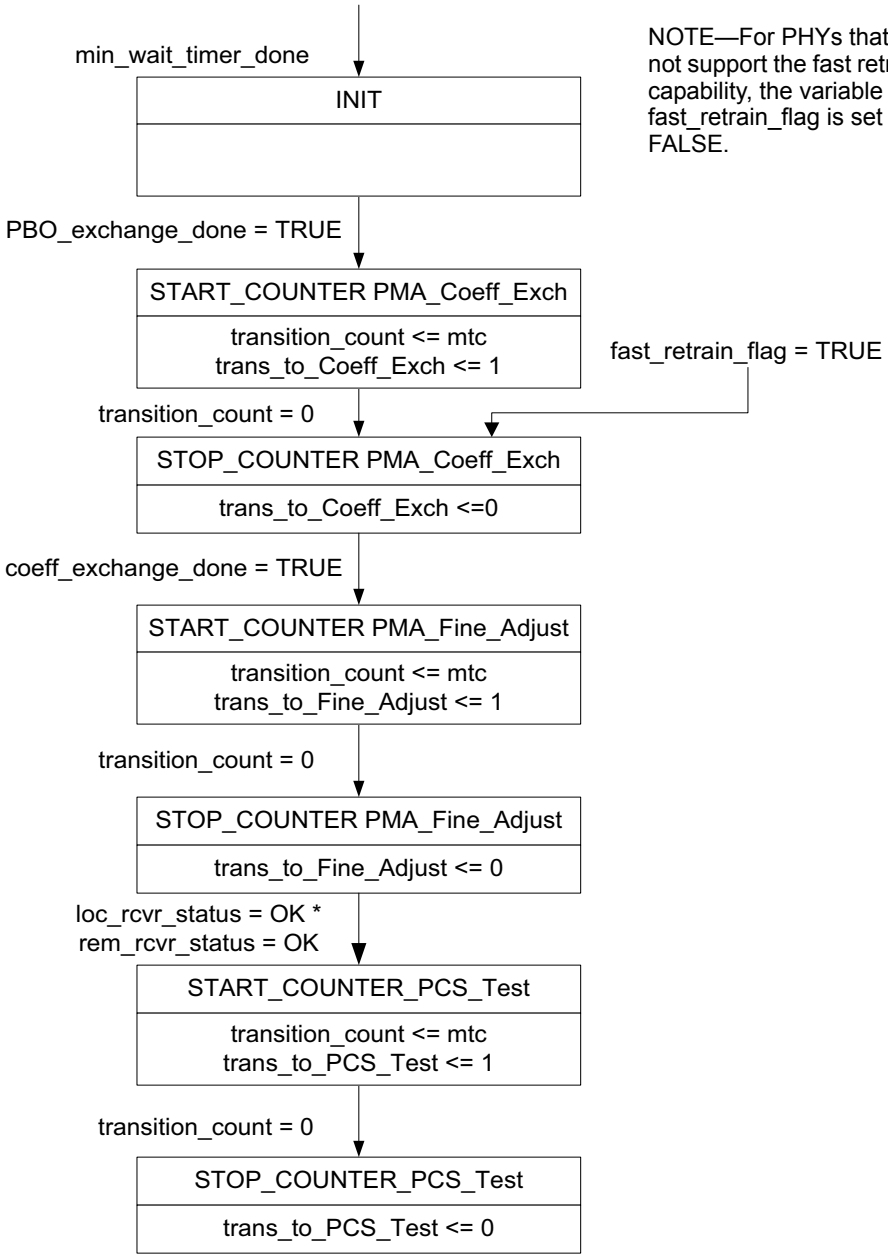

Figure 55–29—MASTER transition counter state diagram

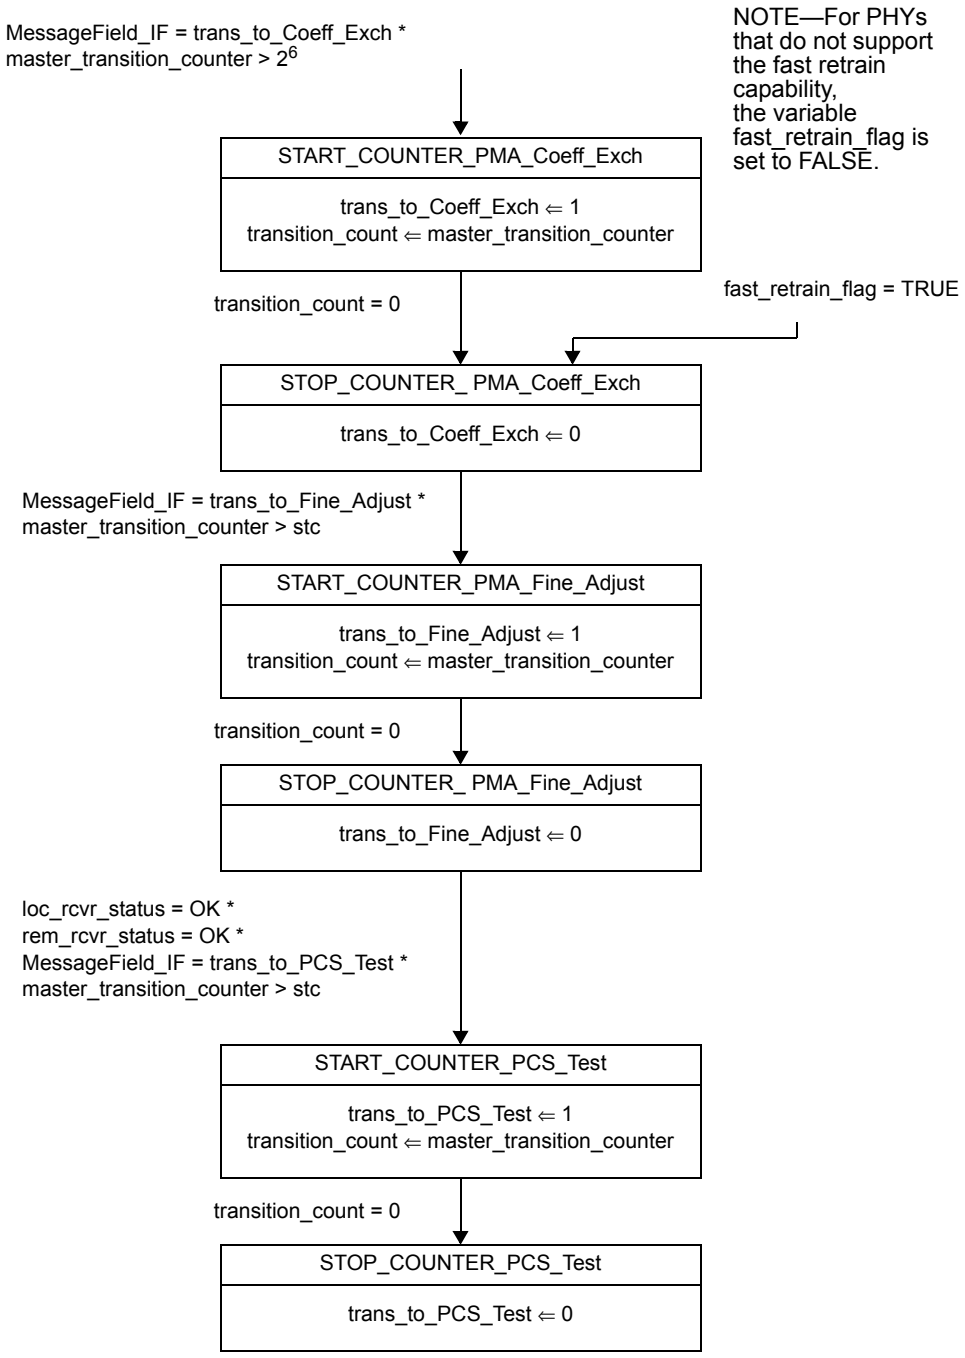

Figure 55–30—SLAVE transition counter state diagram

55.4.6.3 Link Monitor state diagram

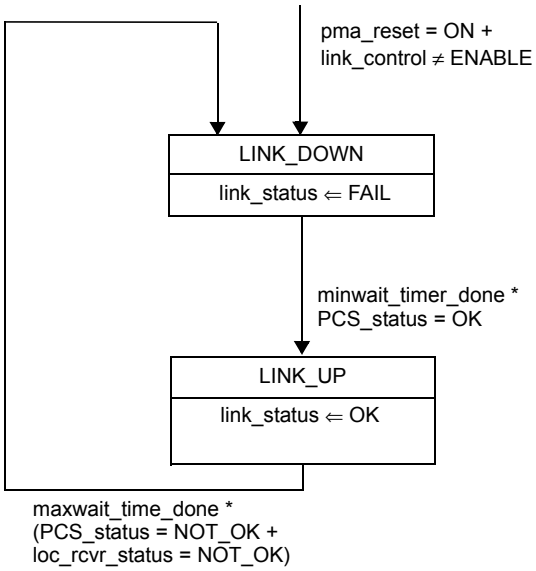

NOTE 1—maxwait_timer is started in PHY Control state diagram (see Figure 55–28).
NOTE 2—The variables link_control and link_status are designated as link_control_10GigT and link_status_10GigT, respectively, by the Auto-Negotiation Arbitration state diagram (Figure 28–16).

Figure 55–31—Link Monitor state diagram

55.4.6.4 EEE Refresh monitor state diagram

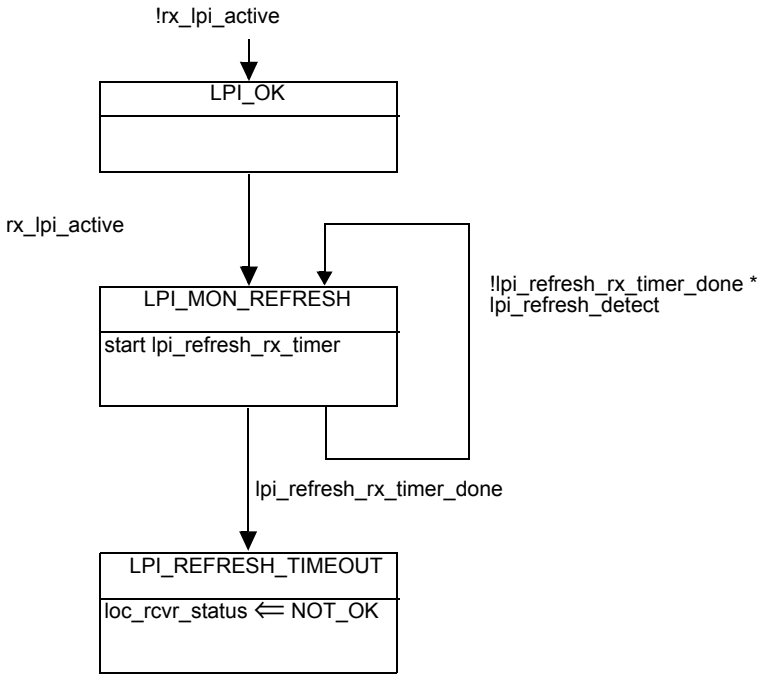

NOTE—This state diagram is only required when the PHY supports the EEE capability.

Figure 55–32—EEE Refresh monitor state diagram

55.4.6.5 Fast retrain state diagram

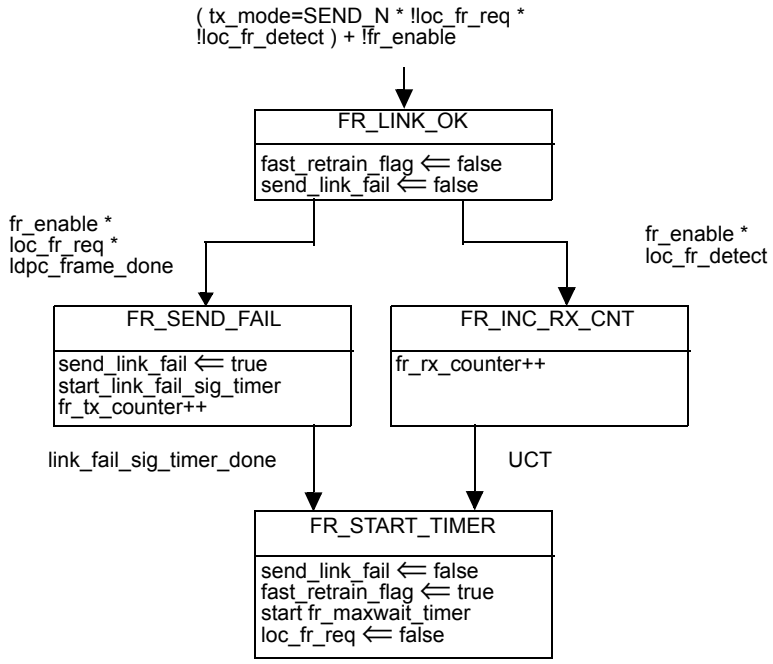

NOTE—This state diagram is only required when the PHY supports the fast retrain capability.

Figure 55–33—Fast retrain control state diagram

55.5 PMA electrical specifications

This subclause defines the electrical characteristics of the PMA and specifies PMA-to-MDI interface tests.

Common-mode tests use the common-mode return point as a reference.

55.5.1 Isolation requirement

The PHY shall provide electrical isolation between the port device circuits, including frame ground (if any) and all MDI leads. This electrical isolation shall withstand at least one of the following electrical strength tests:

a) 1500 V rms at 50 Hz to 60 Hz for 60 s, applied as specified in Section 5.2.2 of IEC 60950-1:2001.
b) 2250 V dc for 60 s, applied as specified in Section 5.2.2 of IEC 60950-1:2001.
c) A sequence of ten 2400 V impulses of alternating polarity, applied at intervals of not less than 1 s. The shape of the impulses is 1.2/50 μs (1.2 μs virtual front time, 50 μs virtual time or half value), as defined in Annex N of IEC 60950-1:2001.

There shall be no insulation breakdown, as defined in Section 5.2.2 of IEC 60950-1:2001, during the test. The resistance after the test shall be at least 2 MΩ, measured at 500 V dc.

55.5.2 Test modes

The test modes described below shall be provided to allow for testing of the transmitter waveform, transmitter distortion, transmitted jitter, transmitter droop and BER testing.

For a PHY with an MDIO management interface, these modes shall be enabled by setting bits 1.132.15:13 (10GBASE-T test mode register) of the MDIO Management register set as shown in Table 55–12. These test modes shall only change the data symbols provided to the transmitter circuitry and shall not alter the electrical and jitter characteristics of the transmitter and receiver from those of normal (non-test mode) operation. PHYs without a MDIO shall provide a means to enable these modes for conformance testing.

Table 55–12—MDIO management register settings for test modes

1.132.15	1.132.14	1.132.13	Mode
0	0	0	Normal operation.
0	0	1	Test mode 1—Setting of MASTER transmitter required by SLAVE for transmit jitter test in SLAVE mode.
0	1	0	Test mode 2—Transmit jitter test in MASTER mode.
0	1	1	Test mode 3—Transmit jitter test in SLAVE mode (if loop timing is supported).
1	0	0	Test mode 4—Transmit distortion test.
1	0	1	Test mode 5—Normal operation with no power backoff. This is for the PSD mask and power level test.
1	1	0	Test mode 6—Transmitter droop test mode.
1	1	1	Test mode 7—Pseudo random test mode for BER Monitor.

Test mode 1 is a mode provided for enabling testing of timing jitter on a SLAVE transmitter. When test mode 1 is enabled, the PHY shall transmit the PMA training pattern (PRBS 33) continually from all four transmitters with the THP turned off, with no power backoff and with the transmitted symbols timed from its local clock source.

Test mode 2 is for transmitter jitter testing when transmitter is in MASTER timing mode. When test mode 2 is enabled, the PHY shall transmit {Two +16 symbols followed by two –16 symbols} continually from all four transmitters with the THP turned off, with no power backoff and with the transmitted symbols timed from its local clock source. The transmitter output is a 200 MHz signal.

Test mode 3 is optional for a PHY that does not support loop timing. When test mode 3 is enabled on a PHY that supports loop timing, the PHY shall transmit, with THP turned off, the data symbol sequence {two +16 symbols followed by two –16 symbols} repeatedly on pair D with the symbols timed from its recovered receive data clock in SLAVE timing mode. A PHY that supports loop timing operates in test mode 3 when there is no input signal on pair D. The transmitter output is a 200 MHz signal on pair D and shall be silence on pairs A, B, and C.

Test mode 4 is for transmitter linearity testing. When test mode 4 is enabled, the PHY shall transmit, with the THP turned off, transmitted symbols, timed from a transmit clock (as specified in 55.5.3.5) in the MASTER timing mode, defined by the bits1.132.12:10 and Table 55–13.

Table 55–13—MDIO management register settings for transmit frequencies in Test mode 4

1.132.12	1.132.11	1.132.10	Output waveform frequencies in MHz
			Two tone frequency pairs
0	0	0	Reserved
0	1	1	Reserved
1	1	1	Reserved
0	0	1	(800/1024) × 47, (800/1024) × 53
0	1	0	(800/1024) × 101, (800/1024) × 103
1	0	0	(800/1024) × 179, (800/1024) × 181
1	0	1	(800/1024) × 277, (800/1024) × 281
1	1	0	(800/1024) × 397, (800/1024) × 401

The peak-to-peak levels used in this test shall correspond to the ± 16 symbol levels and the relative amplitudes of the tones in a two-tone pair shall be within 0.5 dB of each other.

Test mode 5 is for checking whether the transmitter is compliant with the transmit PSD mask and the transmit power level. When test mode 5 is enabled, the PHY shall transmit as in normal operation but with the power backoff disabled.

Test mode 6 is for testing transmitter droop. When test mode 6 is enabled, the PHY shall transmit the following sequence of data symbols A_n, B_n, C_n, D_n, of 55.4.3.1 continually from all four transmitters, with the THP turned off:

{One hundred twenty eight +16 followed by one hundred twenty eight –16 symbols}.

Test mode 7 is for enabling measurement of the bit error ratio of the link including the LDPC encoder/decoder, the transmit and receive analog front ends of the PHY and a cable connecting two PHYs. This mode reuses the 10GBASE-T PCS scrambler and is defined in detail in 55.3.3.

55.5.2.1 Test fixtures

The following fixtures (illustrated by Figure 55–34, Figure 55–35, and Figure 55–36), or their functional equivalents, can be used for measuring the transmitter specifications described in 55.5.3.

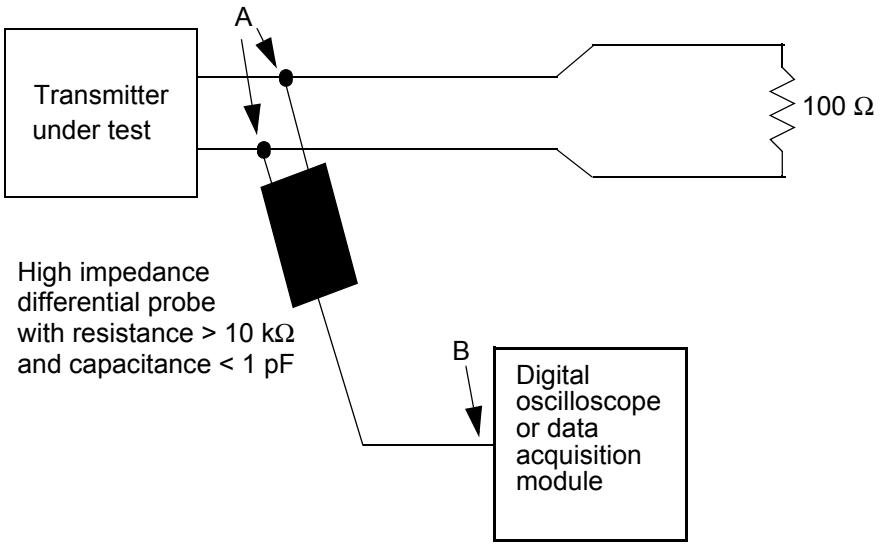

Figure 55–34—Transmitter test fixture 1 for transmitter droop measurement

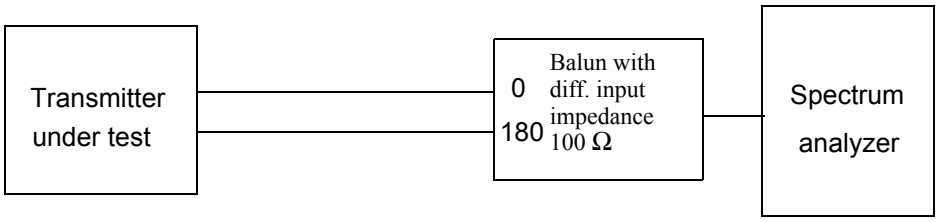

Figure 55–35—Transmitter test fixture 2 for linearity measurement, power spectral density measurement and transmit power level measurement

The high impedance probe shown in Figure 55–34 in transmitter test fixture 1 has resistance > 10 kΩ and capacitance < 1 pF over the frequency range of 1 MHz to 400 MHz. Figure 55–35 includes a power summer or balun device to couple the 100 Ω differential output of the transmitter to the 50 Ω single-ended input typically found in a spectrum analyzer input. The center frequency (f_c) of the band pass filter show in Figure 55–36 is 200 MHz ± 200 kHz and the band pass filter noise bandwidth (B_n) is 2 MHz ± 200 kHz.

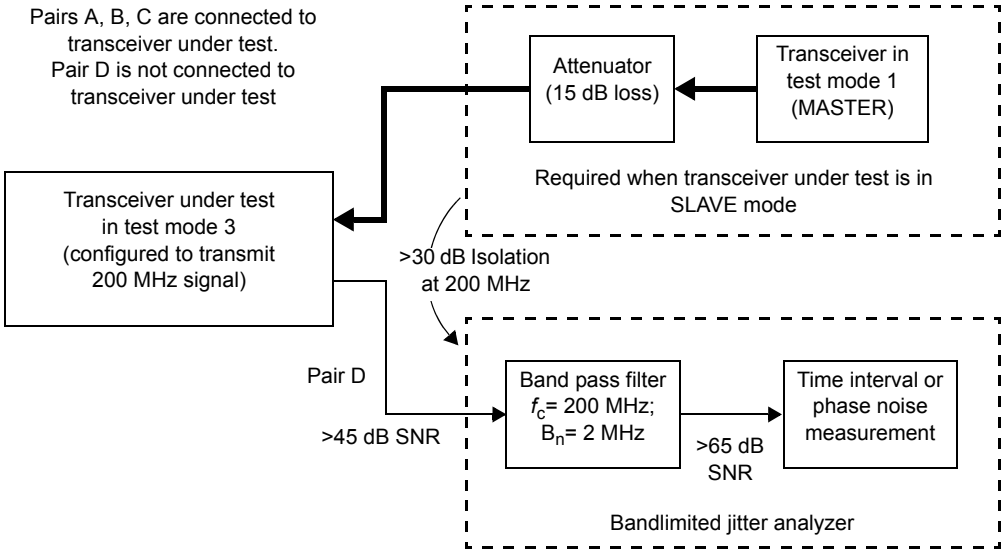

Figure 55–36—Transmitter test fixture 3 for transmitter jitter measurement

55.5.3 Transmitter electrical specifications

The PMA provides the Transmit function specified in 55.4.2.2 in accordance with the electrical specifications of this clause. The PMA shall operate with AC-coupling to the MDI.

Where a load is not specified, the transmitter shall meet the requirements of this clause with a 100 Ω resistive differential load connected to each transmitter output.

55.5.3.1 Maximum output droop

With the transmitter in test mode 6 and using the transmitter test fixture 1, the magnitude of both the positive and negative droop shall be less than 10%, measured with respect to an initial value at 10 ns after the zero crossing and a final value at 90 ns after the zero crossing.

55.5.3.2 Transmitter linearity.

When in test mode 4 and observing the spectrum of the differential signal output at the MDI using transmitter test fixture 2, for each pair, with no intervening cable, the transmitter linearity mask is defined as follows:

The SFDR of the transmitter, with dual tone inputs as specified in test mode 4, shall meet the requirement that:

$$\mathrm{SFDR} \geq \quad 2.5 + \ \min\{52,\ 58 - 20 \times \log_{10}(f/25)\} \tag{55–7}$$

where f is the maximum frequency of the two test tones in MHz and SFDR is the ratio in dB of the minimum RMS value of either input tone to the RMS value of the worst intermodulation product in the frequency range of 1 MHz to 400 MHz.

This specification on transmit linearity is derived from the requirement for interoperability with the far-end device.

55.5.3.3 Transmitter timing jitter

When in test mode 2, the PHY transmits {two +16 symbols followed by two –16 symbols} continually with the THP turned off and with no power backoff. In this mode, the transmitter output should be a 200 MHz signal and the RMS period jitter measured at the PHY MDI output shall be less than 5.5 ps. The RMS period jitter is measured as per the test configuration shown in Figure 55–36 over an integration time interval of 1 ms ± 10%.

For a PHY that can operate in loop timing mode, the SLAVE mode RMS period jitter test shall be run using the test configuration shown in Figure 55–36. For this test, the MASTER PHY is in test mode 1 and the SLAVE PHY is in test mode 3. The MASTER is transmitting the PMA training pattern (PRBS 33) to the SLAVE PHY on pairs A, B, and C. The SLAVE PHY is in loop timing mode, synchronizing its transmit clock to the signals received from the MASTER PHY on pairs A, B, and C. In this configuration, the transmitter output on pair D should be a 200 MHz signal and the RMS period jitter measured at the SLAVE PHY MDI output shall be less than 5.5 ps. The RMS period jitter is measured over an integration time interval of 1 ms ± 10%.

RMS period jitter over an integration time interval of 1 ms ± 10% is defined as the root mean square period difference from the average period ($T - T_{avg}$), accumulated over a sample size of 200 000 ± 20 000:

$$\text{RMS period jitter} = \sqrt{\frac{\sum[(T - T_{avg})^2]}{\text{Sample size}}} \tag{55–8}$$

55.5.3.4 Transmitter power spectral density (PSD) and power level

In test mode 5 (normal operation with no power backoff), the transmit power shall be in the range 3.2 dBm to 5.2 dBm and the power spectral density of the transmitter, measured into a 100 Ω, load using the test fixture shown in Figure 55–35 shall be between the upper and lower masks specified in Equation (55–9) and Equation (55–10). The masks are shown graphically in Figure 55–37.

$$\text{Upper PSD}\,(f) \le \begin{cases} -78.5 \text{ dBm/Hz} & 0 < f \le 70 \\ -78.5 - \left(\dfrac{f-70}{80}\right) \text{dBm/Hz} & 70 < f \le 150 \\ -79.5 - \left(\dfrac{f-150}{58}\right) \text{dBm/Hz} & 150 < f \le 730 \\ -79.5 - \left(\dfrac{f-330}{40}\right) \text{dBm/Hz} & 730 < f \le 1790 \\ -116 \text{ dBm/Hz} & 1790 < f \le 3000 \end{cases} \tag{55–9}$$

and

$$\text{Lower PSD}\ (f) \geq \begin{cases} -83\ \text{dBm/Hz} & 5 \leq f \leq 50 \\ -83 - \left(\dfrac{f-50}{50}\right)\text{dBm/Hz} & 50 < f \leq 200 \\ -86 - \left(\dfrac{f-200}{25}\right)\text{dBm/Hz} & 200 < f \leq 400 \end{cases} \quad (55\text{–}10)$$

where f is in MHz.

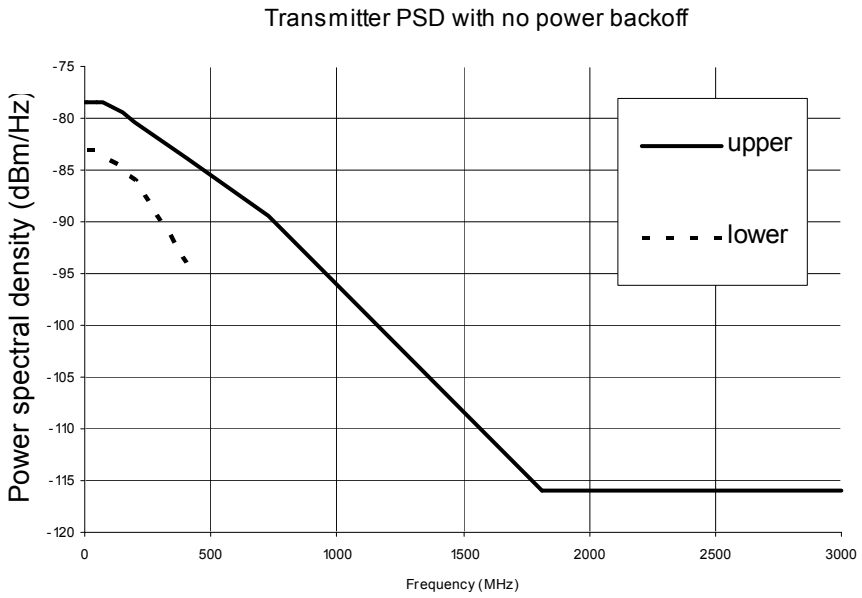

Figure 55–37—Transmitter power spectral density mask

55.5.3.5 Transmit clock frequency

The symbol transmission rate on each pair of the MASTER PHY shall be within the range 800 MHz ± 50 ppm.

When the transmitter is in the LPI transmit mode or when the receiver is in the LPI receive mode the transmitter clock short-term rate of frequency variation shall be less than 0.1 ppm/second. The short-term frequency variation limit shall also apply when switching to and from the LPI mode.

55.5.4 Receiver electrical specifications

The PMA provides the Receive function specified in 55.4.2.4 in accordance with the electrical specifications of this clause using patch cabling and interconnecting hardware that is within the limits specified in 55.7.

55.5.4.1 Receiver differential input signals

Differential signals received at the MDI that were transmitted from a remote transmitter within the specifications of 55.5.3 and have passed through a link specified in 55.7 are received with a BER less than

10^{-12} and sent to the PCS after link reset completion. This specification shall be satisfied by a frame error ratio less than 9.6×10^{-9} for 800 octet frames with minimum IPG or greater than 799 octet IPG.

55.5.4.2 Receiver frequency tolerance

The receive feature shall properly receive incoming data, per the requirements of 55.5.4.1, with a symbol rate within the range 800 MHz ± 50 ppm.

55.5.4.3 Common-mode noise rejection

This specification is provided to limit the sensitivity of the PMA receiver to common-mode noise from the cabling system. Common-mode noise generally results when the cabling system is subjected to electromagnetic fields.

The common-mode noise can be simulated using the cable clamp test defined in 40.6.1.3.3. A 6 dBm sine wave signal from 80 MHz to 1000 MHz can be used to simulate an external electromagnetic field. Operational requirements of the transceiver during the test are determined by the manufacturer. A system integrating a 10GBASE-T PHY may perform this test.

55.5.4.4 Alien crosstalk noise rejection

While receiving data from a transmitter compliant with specifications in 55.5.3, through a 100 m link segment compliant with the specifications in 55.7, a receiver shall operate with an Ethernet frame error ratio less than 9.6×10^{-9} for 800 octet frames with either a minimum IPG or greater than 799 octet IPG with four noise sources at the specified levels representing alien crosstalk, one connected to each of the four pairs. Independent noise sources should be injected into each MDI inputs using couplers that do not significantly alter the link segment characteristics. Each noise source shall have a flat noise spectrum, with 3 dB bandwidth at least 10 MHz to 400 MHz and a power spectral density such that at the MDI port of the device under test the power spectral density of the injected noise is –141.9 dBm/Hz. A flat noise source is chosen to model the sum of all alien noise sources. See Figure 55–38.

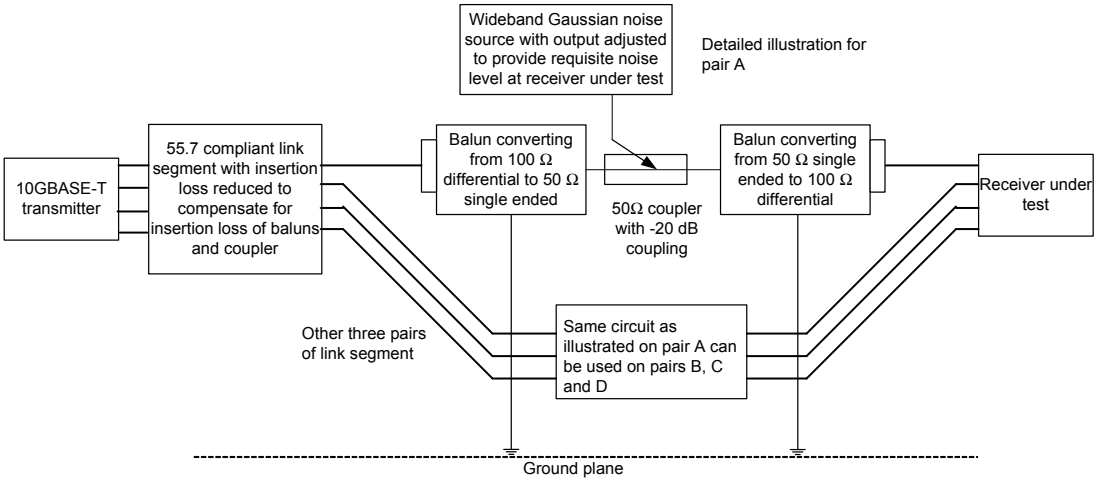

Figure 55–38—Alien crosstalk noise rejection test

The structure shown for injecting the noise in Figure 55–38 is illustrative and alternative approaches are possible. The loss of the coupling structure shown in Figure 55–38, which consists of two baluns and a coupler, is approximately 2.5 dB. The overall insertion loss of the link segment together with the insertion loss of the coupling structure should be adjusted to match the insertion loss specified in 55.7.2.1 to within ± 0.5 dB. This can be achieved by shortening the cable by approximately 5 m. The balun-coupler-balun

structure shown in Figure 55–38 can be replaced by resistively coupling a balanced noise source to the twisted pair using 500 Ω resistors. In either case, calibration of the test setup is required to confirm the overall insertion loss and the injected noise power at the MDI of the receiver under test.

55.5.4.5 Short reach mode link test

In short reach mode (indicating operation over a short reach channel) while receiving data from a transmitter compliant with specifications in 55.5.3 (whether or not in short reach mode), through both short reach test channels, a receiver shall operate with a frame error ratio less than 9.6×10^{-9} for 800 octet frames with minimum IPG or greater than 799 octet IPG (e.g., operate with a BER less than 10^{-12}).

The PHY short reach register setting 1.131.0 indicates whether the PHY is operating in the short reach mode.

55.5.4.5.1 Short reach test channels

One short reach test channel specification is consistent with the transmission characteristics of a cabling link of 30 m consisting of two connectors, 10 m of cable cord, and a horizontal cable distance of 20 m, that meets the ISO/IEC 11801 Category 7 component specifications.

One short reach test channel specification is consistent with the transmission characteristics of a cabling link of 30 m consisting of two connectors, 10 m of cable cord, and a horizontal cable distance of 20 m that meets the Category 6A component specifications.

55.6 Management interfaces

10GBASE-T makes extensive use of the management functions that may be provided by the MDIO (Clause 45), and the communication and self-configuration functions provided by Auto-Negotiation (Clause 28). Additional Auto-Negotiation requirements are set forth within this subclause.

55.6.1 Support for Auto-Negotiation

All 10GBASE-T PHYs shall provide support for Auto-Negotiation (Clause 28) and shall be capable of operating as MASTER or SLAVE. All 10GBASE-T PHYs shall provide support for Extended Next Pages as defined in 28.2.3.4.2 and shall support and use optimized FLP Burst to FLP burst timing as defined in 28.2.1.1.1, and nlp_link_test_min_timer and link_fail_inhibit_timer as defined in 28.3.2.

Auto-Negotiation is performed as part of the initial set-up of the link, and allows the PHYs at each end to advertise their capabilities (speed, PHY type, half or full duplex) and to automatically select the operating mode for communication on the link. Auto-Negotiation signaling is used for the following primary purposes for 10GBASE-T:

a) To negotiate that the PHY is capable of supporting 10GBASE-T transmission.
b) To determine the MASTER-SLAVE relationship between the PHYs at each end of the link.
c) To determine whether the local PHY performs PMA training pattern reset.
d) To determine whether the local PHY supports the EEE capability.
e) To determine whether the local PHY supports the fast retrain capability.

55.6.1.1 10GBASE-T use of registers during Auto-Negotiation

A 10GBASE-T PHY shall use the management register definitions and values specified in Table 55–14.

Table 55–14—10GBASE-T registers

Register	Bit	Name	Description	Type [a]
7.0	7.0.15:0	AN control register	Defined in 45.2.7.1	R/W
7.1	7.1.15:0	AN status register	Defined in 45.2.7.2	RO
7.2, 7.3	7.2.15:0, 7.3.15:0	AN device identifier registers	Defined in 45.2.7.3	R/W
7.5, 7.6	7.5.15:0, 7.6.15:0	AN devices in package registers	Defined in 45.2.7.4	R/W
7.14, 7.15	7.14.15:0, 7.15.15:0	AN package identifier registers	Defined in 45.2.7.5	R/W
7.16	7.16.15:0	AN advertisement register	Defined in 45.2.7.6	R/W
7.19	7.19.15:0	AN LP Base Page ability register	Defined in 45.2.7.7	RO
7.22, 7.23, 7.24	7.22.15:0, 7.23.15:0, 7.24.15:0	AN XNP transmit register	Defined in 45.2.7.8	R/W
7.25, 7.26, 7.27	7.25.15:0, 7.26.15:0, 7.27.15:0	AN LP XNP ability register	Defined in 45.2.7.9	RO
7.32	7.32.15:0	10GBASE-T AN control register	Defined in 45.2.7.10	R/W
7.33	7.33.15:0	10GBASE-T AN status register	Defined in 45.2.7.11	RO

[a] R/W = Read/Write, RO = Read only

55.6.1.2 10GBASE-T Auto-Negotiation page use

10GBASE-T PHYs shall exchange a 10GBASE-T and 1000BASE-T formatted Extended Next Page, as specified in Table 55–15, immediately following the exchange of the Base Page.

Note that the Acknowledge 2 bit is not utilized and has no meaning when used for the 10GBASE-T message page exchange.

Table 55–15—10GBASE-T Base and Next Pages bit assignments

Bit	Name	Description
Base Page		
D15	Next Page	Defined in 28.2.1.2.6
D14	Acknowledge	Defined in 28.2.1.2.5
D13	Remote Fault	Defined in 28.2.1.2.4
D12	Extended Next Page	Defined in 28.2.1.2.3
D11:D5	Technology Ability Field	Defined in 28.2.1.2.2
D4:D0	Selector Field	Defined in 28.2.1.2.1
Extended Next Page (Message Code Field and Flags Field)		
M10:M0	Next Page message code	Defined in Annex 28C
T	Toggle	Defined in 28.2.3.4.7

Table 55–15—10GBASE-T Base and Next Pages bit assignments *(continued)*

Bit	Name	Description
Ack2	Acknowledge 2	Defined in 28.2.3.4.6
MP	Message Page	Defined in 28.2.3.4.5
Ack	Acknowledge	Defined in 28.2.3.4.4
NP	Next Page	Defined in 28.2.3.4
Extended Next Page (Unformatted Message Code Field)		
U31:U25	Reserved, transmit as 0	
U24	10GBASE-T EEE (1 = Advertise EEE capability for 10GBASE-T 0 = Do not advertise EEE capability for 10GBASE-T)	Defined in 45.2.7.13.10
U23	1000BASE-T EEE (1 = Advertise EEE capability for 1000BASE-T 0 = Do not advertise EEE capability for 1000BASE-T)	Defined in 45.2.7.13.11
U22	100BASE-TX EEE (1 = Advertise EEE capability for100BASE-TX 0 = Do not advertise EEE capability for 100BASE-TX)	Defined in 45.2.7.13.12
U21	Reserved	
U20	LD PMA training reset request (1 = Local Device requests that Link Partner reset PMA training PRBS every frame 0 = Local Device requests that Link Partner run PMA training PRBS continuously)	Defined in 45.2.7.10.5
U19	Fast retrain ability (1 = Advertise PHY as supporting fast retrain, 0 = Advertise PHY as not supporting fast retrain)	Defined in 45.2.7.10.6
U18	PHY short reach mode (1 = PHY of Local Device is operating in short reach mode 0 = PHY of Local Device is operating in normal mode)	Defined in 45.2.1.64.2
U17	LD loop timing ability (1 = Advertise PHY as capable of loop timing and 0 = do not advertise PHY as capable of loop timing)	Defined in 45.2.7.10.7
U16	10GBASE-T ability (1 = support of 10GBASE-T and 0 = no support)	Defined in 45.2.7.10.4
U15	1000BASE-T half duplex (1 = half duplex and 0 = no half duplex)	
U14	1000BASE-T full duplex (1 = full duplex and 0 = no full duplex)	
U13	Port type bit (1 = multiport device and 0 = single-port device)	Defined in 45.2.7.10.3
U12	10GBASE-T MASTER-SLAVE config value (1 = MASTER and 0 = SLAVE) This bit is ignored if 7.32.15 =0.	Defined in 45.2.7.10.1
U11	10GBASE-T MASTER-SLAVE manual config enable (1 = manual configuration enable) This bit is intended to be used for manual selection in a particular MASTER-SLAVE mode and is to be used in conjunction with bit 7.32.14.	Defined in 45.2.7.10.2
U10	MASTER-SLAVE seed Bit 10 (SB10) (MSB)	
U9	MASTER-SLAVE seed Bit 9 (SB9)	

Table 55–15—10GBASE-T Base and Next Pages bit assignments *(continued)*

Bit	Name	Description
U8	MASTER-SLAVE seed Bit 8 (SB8)	
U7	MASTER-SLAVE seed Bit 7 (SB7)	
U6	MASTER-SLAVE seed Bit 6 (SB6)	
U5	MASTER-SLAVE seed Bit 5 (SB5)	
U4	MASTER-SLAVE seed Bit 4 (SB4)	
U3	MASTER-SLAVE seed Bit 3 (SB3)	
U2	MASTER-SLAVE seed Bit 2 (SB2)	
U1	MASTER-SLAVE seed Bit 1 (SB1)	
U0	MASTER-SLAVE seed Bit 0 (SB0)	

55.6.1.3 Sending Next Pages

Implementers who do not wish to send additional Next Pages (i.e., Next Pages in addition to those required to perform PHY configuration as defined in this clause) can use Auto-Negotiation as defined in Clause 28.

55.6.2 MASTER-SLAVE configuration resolution

Since both PHYs that share a link segment are capable of being MASTER or SLAVE, a prioritization scheme exists to ensure that the correct mode is chosen. The MASTER-SLAVE relationship shall be determined during Auto-Negotiation using Table 55–16 with the 10GBASE-T Technology Ability Next Page bit values specified in Table 55–15 and information received from the link partner. This process is conducted at the entrance to the FLP LINK GOOD CHECK state shown in the Arbitration state diagram (Figure 28–18.)

The following four equations are used to determine these relationships:

manual_MASTER = U11 * U12
manual_SLAVE = U11 * !U12
single-port device = !U11 * !U13
multiport device = !U11 * U13

where

U11 is bit 11 of 10GBASE-T/1000BASE-T Technology message code,
U12 is bit 12 of 10GBASE-T/1000BASE-T Technology message code,
U13 is bit 13 of 10GBASE-T/1000BASE-T Technology message code (see Table 55–15).

A 10GBASE-T PHY is capable of operating either as the MASTER or SLAVE. In the scenario of a link between a single-port device and a multiport device, the preferred relationship is for the multiport device to be the MASTER PHY and the single-port device to be the SLAVE. However, other topologies may result in contention. The resolution function of Table 55–16 is defined to handle any relationship conflicts.

In the instance when both link partners support the optional loop timing mode, or both link partners do not support the loop timing mode, as designated by bit U17, the resolution shown in Table 55–16 shall be used.

When only one link partner supports loop timing, the device that supports loop timing shall be forced to SLAVE and the other device shall be forced to MASTER.

Table 55–16—10GBASE-T MASTER-SLAVE configuration resolution table

Local device type	Remote device type	Local device resolution	Remote device resolution
single-port device	multiport device	SLAVE	MASTER
single-port device	manual_MASTER	SLAVE	MASTER
manual_SLAVE	manual_MASTER	SLAVE	MASTER
manual_SLAVE	multiport device	SLAVE	MASTER
multiport device	manual_MASTER	SLAVE	MASTER
manual_SLAVE	single-port device	SLAVE	MASTER
multiport device	single-port device	MASTER	SLAVE
multiport device	manual_SLAVE	MASTER	SLAVE
manual_MASTER	manual_SLAVE	MASTER	SLAVE
manual_MASTER	single-port device	MASTER	SLAVE
single-port device	manual_SLAVE	MASTER	SLAVE
manual_MASTER	multiport device	MASTER	SLAVE
multiport device	multiport device	The device with the higher seed value is configured as MASTER, otherwise SLAVE.	The device with the higher seed value is configured as MASTER, otherwise SLAVE.
single-port device	single-port device	The device with the higher seed value is configured as MASTER, otherwise SLAVE	The device with the higher seed value is configured as MASTER, otherwise SLAVE.
manual_SLAVE	manual_SLAVE	MASTER-SLAVE configuration fault	MASTER-SLAVE configuration fault
manual_MASTER	manual_MASTER	MASTER-SLAVE configuration fault	MASTER-SLAVE configuration fault

A 10GBASE-T multiport device has higher priority than a single-port device to become the MASTER. In the case where both devices are of the same type, e.g., both devices are multiport devices and both devices have the same loop timing support, the device with the higher MASTER-SLAVE seed bits (SB0...SB10), where SB10 is the MSB, shall become the MASTER and the device with the lower seed value shall become the SLAVE. In case both devices have the same seed value, both should assert link_status_10GigT=FAIL (as defined in 28.3.1) to force a new cycle through Auto-Negotiation. Successful completion of the MASTER-SLAVE resolution shall be treated as MASTER-SLAVE configuration resolution complete.

The method of generating a random or pseudorandom seed is left to the implementer. The generated random seeds should belong to a sequence of independent, identically distributed integer numbers with a uniform distribution in the range of 0 to 2^{11}– 2. The algorithm used to generate the integer should be designed to minimize the correlation between the number generated by any two devices at any given time. A seed counter shall be provided to track the number of seed attempts. The seed counter shall be set to zero at start-up and shall be incremented each time a seed is generated. When MASTER-SLAVE resolution is complete,

the seed counter shall be reset to 0 and bit 7.33.15 shall be set to zero. A MASTER-SLAVE resolution fault shall be declared if resolution is not reached after the generation of seven seeds.

The MASTER-SLAVE manual configuration enable bit (control register bit 7.32.15) and the MASTER-SLAVE config value bit (control register bit 7.32.14) are used to manually set a device to become the MASTER or the SLAVE. In case both devices are manually set to become the MASTER or the SLAVE, this condition shall be flagged as a MASTER-SLAVE configuration fault condition; thus the MASTER-SLAVE configuration fault bit (status register bit 7.33.15) shall be set to logical one. In the situation where one link partner supports loop timing and the other does not, a MASTER-SLAVE configuration fault bit (status register bit 7.33.15) shall be set to one if the link partner with loop timing support is manually set to MASTER or the link partner with no loop timing support is manually set to SLAVE. The MASTER-SLAVE configuration fault condition shall be treated as MASTER-SLAVE configuration resolution complete and link_status_10GigT shall be set to FAIL, because the MASTER-SLAVE relationship was not resolved. This forces a new cycle through Auto-Negotiation after the link_fail_inhibit_timer has expired. Determination of MASTER-SLAVE values occur on the entrance to the FLP LINK GOOD CHECK state (Figure 28–18) when the highest common denominator (HCD) technology is 10GBASE-T. The resulting MASTER-SLAVE value is used by the 10GBASE-T PHY Control (55.4.2.5).

If MASTER-SLAVE manual configuration is disabled (bit 7.32.15 is set to 0) and the local device detects that both the local device and the remote device are of the same type (either multiport device or single-port device and identical loop timing support) and that both have generated the same random seed, it generates and transmits a new random seed for MASTER-SLAVE negotiation by setting link_status to FAIL and cycling through the Auto-Negotiation process again.

The MASTER-SLAVE configuration process returns one of the three following outcomes:

a) *Successful:* Bit 7.33.15 of the 10GBASE-T status register is set to logical zero and bit 7.33.14 is set to logical one for MASTER resolution or for logical zero for SLAVE resolution. 10GBASE-T returns control to Auto-Negotiation (at the entrance to the FLP LINK GOOD CHECK state in Figure 28–18) and passes the value MASTER or SLAVE to PMA_CONFIG.indicate (see 55.2.2.2).
b) *Unsuccessful:* link_status_10GigT is set to FAIL and Auto-Negotiation restarts (see Figure 28–18).
c) *Fault detected:* Bit 7.33.15 of the 10GBASE-T status register is set to logical one to indicate that a configuration fault has been detected. This bit also is set when seven attempts to configure the MASTER-SLAVE relationship via the seed method have failed. When a fault is detected, link_status_10GigT is set to FAIL, causing Auto-Negotiation to cycle through again.

NOTE—MASTER-SLAVE arbitration only occurs if 10GBASE-T or 1000BASE-T is selected as the highest common denominator.

55.7 Link segment characteristics

10GBASE-T is designed to operate over ISO/IEC 11801 Class E or Class F 4-pair balanced cabling that meets the additional requirements specified in this subclause. Each of the four pairs supports an effective data rate of 2500 Mb/s in each direction simultaneously. The term “link segment” used in this clause refers to four twisted-pairs operating in full duplex. Specifications for a link segment apply equally to each of the four twisted-pairs. All implementations of the balanced cabling link segment specification shall be compatible at the MDI. It is recommended that the guidelines in TIA TSB-155-A, ISO/IEC TR 24750, ANSI/TIA-568-C.2, and ISO/IEC 11801:2002/Amendment 1 be considered before the installation of 10GBASE-T equipment for any cabling system.

55.7.1 Cabling system characteristics

The cabling system used to support 10GBASE-T requires 4 pairs of ISO/IEC 11801 Class E or Class F balanced cabling with a nominal impedance of 100 Ω. Operation on other classes of cabling may be supported if the link segment meets the requirements of 55.7.

Additionally:

a) 10GBASE-T uses a star topology with Class E or Class F balanced cabling used to connect PHY entities.
b) 10GBASE-T is an ISO/IEC 11801 Class E and Class F application with the additional transmission requirements specified in this subclause. The ISO/IEC 11801 cabling limit calculation minimums apply to the link segment specifications. Mitigation practices may be required—see Annex 55B.

55.7.2 Link segment transmission parameters

A link segment consisting of up to 100 m of Class E or up to 100 m of Class F that meets the transmission parameters of this subclause provides a reliable medium. The transmission parameters of the link segment include insertion loss, delay parameters, nominal impedance, NEXT loss, ELFEXT, and return loss. In addition, the requirements for the alien crosstalk coupled "between" link segments is specified.

Table 55–17 lists the supported cabling types and distances.

Table 55–17— Cabling types and distances

Cabling	Supported link segment distances	Cabling references
Class E / Category 6	55 m to 100 m[a]	ISO/IEC TR 24750 / TIA TSB-155-A
Class E / Category 6: unscreened	55 m	ISO/IEC TR 24750 / TIA TSB-155-A
Class E / Category 6: screened	100 m	ISO/IEC TR 24750 / TIA TSB-155-A
Class F	100 m	ISO/IEC TR 24750
Class E_A / Category 6A	100 m	ISO/IEC 11801:2002/Amendment 1 / ANSI/TIA-568-C.2
Class F_A	100 m	ISO/IEC 11801:2002 Amendment 1

[a]Supported link segments up to 100 m shall meet the alien crosstalk to insertion loss requirements specified in 55.7.3.1.2 and 55.7.3.2.2.

The link segment transmission parameters of insertion loss and ELFEXT specified are ISO/IEC 11801 Class E specifications extended by extrapolating the formulas to a frequency up to 500 MHz with appropriate adjustments for length when applicable as specified in ISO/IEC TR 24750 and TIA TSB-155-A. The link segment transmission parameters of NEXT loss, MDNEXT loss, and return loss specified are ISO/IEC 11801 Class E specifications extended beyond 250 MHz by utilizing the equations referenced in TIA TSB-155.

55.7.2.1 Insertion loss

The insertion loss of each duplex channel shall meet the values determined using Equation (55–11).

$$\text{Insertion loss}(f) \leq 1.05\left(1.82 \times \sqrt{f} + 0.0169 \times f + \frac{0.25}{\sqrt{f}}\right) + 4 \times 0.02 \times \sqrt{f} \quad \text{(dB)} \qquad (55\text{–}11)$$

where

f is the frequency in MHz; $1 \leq f \leq 500$

This includes the insertion loss of the balanced cabling pairs, including work area and equipment cables plus connector losses within each duplex channel.

For the purpose of calculating the link segment insertion loss for cabling less than 100 m the cable insertion loss is assumed to scale linearly with length as defined in Equation (55–27).

55.7.2.2 Differential characteristic impedance

The nominal differential characteristic impedance of each link segment duplex channel, which includes cable cords and connecting hardware, is 100 Ω for all frequencies between 1 MHz and 500 MHz.

55.7.2.3 Return loss

In order to limit the noise at the receiver due to impedance mismatches in the cabling system, each link segment duplex channel shall meet the values determined using Equation (55–12) at all frequencies from 1 MHz to 500 MHz. The reference impedance for the return loss specification is 100 Ω.

$$\text{Return loss} \geq \left\{ \begin{array}{ll} 19 & 1 \leq f < 10 \\ 24 - 5\log_{10}(f) & 10 \leq f < 40 \\ 32 - 10\log_{10}(f) & 40 \leq f < 400 \\ 6 & 400 \leq f \leq 500 \end{array} \right\} \quad \text{(dB)} \qquad (55\text{–}12)$$

where

f is the frequency in MHz.

55.7.2.4 Coupling parameters between duplex channels comprising one link segment

In order to limit the noise coupled into a duplex channel from adjacent duplex channels, near-end crosstalk (NEXT) loss and equal level far-end crosstalk (ELFEXT) are specified for each link segment. In addition, each duplex channel can be disturbed by more than one duplex channel. To ensure the total NEXT loss and FEXT loss coupled into a duplex channel is limited, multiple disturber near-end crosstalk (MDNEXT) and multiple disturber ELFEXT (MDELFEXT) is specified.

55.7.2.4.1 Differential near-end crosstalk

In order to limit the crosstalk at the near end of a link segment, the differential pair-to-pair near-end crosstalk (NEXT) loss between a duplex channel and the other three duplex channels is specified to meet the bit error ratio objective specified in 55.1.1. The NEXT loss between any two duplex channels of a link segment shall meet the values determined using Equation (55–13)

$$\text{NEXT loss}(f) \geq \begin{cases} -20\log_{10}\left(10^{\frac{74.3-15\log_{10}(f)}{-20}} + 2\times 10^{\frac{94-20\log_{10}(f)}{-20}}\right) \text{ (dB)} & 1 \leq f < 330 \\ 31 - 50\log_{10}\left(\frac{f}{330}\right) \text{ (dB)} & 330 \leq f \leq 500 \end{cases} \quad (55\text{–}13)$$

where

f is the frequency in MHz.

Calculations that result in NEXT loss values greater than 65 dB shall revert to a requirement of 65 dB minimum.

55.7.2.4.2 Multiple disturber near-end crosstalk (MDNEXT) loss

Since four duplex channels are used to transfer data between PMDs, the NEXT that is coupled into a data carrying channel is from the three adjacent disturbing duplex channels.

To ensure the total NEXT coupled into a duplex channel is limited, multiple disturber NEXT loss is specified as the power sum of the individual NEXT losses. The power sum loss between a duplex channel and the three adjacent disturbers shall meet the values determined using Equation (55–14).

$$\text{MDNEXT loss}(f) \geq \begin{cases} -20\log_{10}\left(10^{\frac{72.3-15\log_{10}(f)}{-20}} + 2\times 10^{\frac{90-20\log_{10}(f)}{-20}}\right) \text{ (dB)} & 1 \leq f < 330 \\ 28 - 50\log_{10}\left(\frac{f}{330}\right) \text{ (dB)} & 330 \leq f \leq 500 \end{cases} \quad (55\text{–}14)$$

where

f is the frequency in MHz.

Calculations that result in PS NEXT loss values greater than 62 dB shall revert to a requirement of 62 dB minimum.

55.7.2.4.3 Multiple disturber power sum near-end crosstalk (PSNEXT) loss

PS NEXT loss is determined by summing the power of the three individual pair-to-pair differential NEXT loss values over the frequency range 1 MHz to 500 MHz as follows in Equation (55–15).

$$\text{PSNEXT loss}(f) = \left(-10\log_{10}\sum_{i=1}^{n} 10^{\frac{-NL(f)i}{10}}\right) \text{ (dB)} \quad (55\text{–}15)$$

where

$NL(f)i$ is the magnitude in dB of NEXT loss at frequency f (in MHz) of pair combination i

i is the pair-to-pair combination (1 to *n*)

n is the number of pair-to-pair combinations (*n*=3)

55.7.2.4.4 Equal level far-end crosstalk (ELFEXT)

Equal level far-end crosstalk (ELFEXT) is specified in order to limit the crosstalk at the far end of each link segment duplex channel and meet the bit error ratio objective specified in 55.1.1. Far-end crosstalk (FEXT) is crosstalk that appears at the far end of a duplex channel (disturbed channel), which is coupled from another duplex channel (disturbing channel) with the noise source (transmitters) at the near end.

FEXT loss is defined in Equation (55–16) as

$$\text{FEXT loss}(f) = 20\log_{10}\left(\frac{V_{pds}(f)}{V_{pcn}(f)}\right) \quad \text{(dB)} \tag{55–16}$$

and ELFEXT is defined in Equation (55–17) as

$$\text{ELFEXT}\ (f) = 20\log_{10}\left(\frac{V_{pds}(f)}{V_{pcn}(f)}\right) - \text{SLS\_Loss(f)}\ \text{(dB)} \tag{55–17}$$

where

Vpds is the peak voltage of disturbing signal (near-end transmitter),
Vpcn is the peak crosstalk noise at far end of disturbed channel,
SLS_Loss is the insertion loss of disturbed channel in dB.

The worst pair ELFEXT between any two duplex channels shall meet the values determined using Equation (55–18).

$$\text{ELFEXT} \geq -20\log_{10}\left(10^{\frac{67.8 - 20\log_{10}(f)}{-20}} + 4 \times 10^{\frac{83.1 - 20\log_{10}(f)}{-20}}\right) \quad \text{(dB)} \tag{55–18}$$

where

f is the frequency in MHz and $1 \leq f \leq 500$.

The numerator of the first term “raised to the power of 10” in Equation (55–18) is the cable portion of the duplex channel of 100 m. The ELFEXT of the cable improves as the cable is reduced in length as defined in Equation (55–19). The equation assumes coupling over 100 m of cable including horizontal cable and cable cords

$$\text{ELFEXT}(f, L) = (67.8 - 20\log_{10}(f)) - 10\log_{10}\left(\frac{L}{100}\right) \quad \text{(dB)} \tag{55–19}$$

where

L is the length in meters of the link segment,
f is the frequency in MHz.

55.7.2.4.5 Multiple disturber equal level far-end crosstalk (MDELFEXT)

Since four duplex channels are used to transfer data between PMDs, the FEXT that is coupled into a data carrying channel is from the three adjacent disturbing duplex channels. To ensure the total FEXT coupled into a duplex channel is limited, multiple disturber ELFEXT is specified as the power sum of the individual ELFEXT disturbers. The power sum loss between a duplex channel and the three adjacent disturbers shall meet the values determined using Equation (55–20).

$$\mathrm{MDELFEXT}(f) \geq -20\log_{10}\left(10^{\frac{64.8 - 20\log_{10}(f)}{-20}} + 4 \times 10^{\frac{80.1 - 20\log_{10}(f)}{-20}}\right) \quad \text{(dB)} \qquad \text{(dB)} \qquad (55\text{–}20)$$

where

f is the frequency $1 \leq f \leq 500$.

55.7.2.4.6 Multiple disturber power sum equal level far-end crosstalk (PS ELFEXT)

PS ELFEXT is determined by summing the power of the three individual pair-to-pair differential ELFEXT values over the frequency range 1 MHz to 500 MHz as follows in Equation (55–21):

$$\mathrm{PSELFEXT}\ (f) = -10\log_{10}\sum_{i=1}^{i=n} 10^{\frac{-EL(f)i}{10}} \quad \text{(dB)} \qquad (55\text{–}21)$$

where

EL(f)i is the magnitude of ELFEXT at frequency *f* of pair combination *i*,
i is the pair-to-pair combination (1 to *n*),
f is the frequency in MHz,
n is the number of pair-to-pair combinations (*n*=3).

55.7.2.5 Maximum link delay

The propagation delay of a link segment shall not exceed 570 ns at all frequencies between 2 MHz and 500 MHz.

55.7.2.6 Link delay skew

The difference in propagation delay, or skew, between all duplex channel pair combinations of a link segment, under all conditions, shall not exceed 50 ns at all frequencies from 2 MHz to 500 MHz. It is a further functional requirement that, once installed, the skew between any two of the four duplex channels due to environmental conditions shall not vary more than 10 ns within the above requirement.

55.7.3 Coupling parameters between link segments

Noise coupled between the disturbed duplex channel in a link segment and the disturbing duplex channels in other link segments is referred to as alien crosstalk noise. To ensure the total alien NEXT loss and alien FEXT loss coupled between link segments is limited, multiple disturber alien near-end crosstalk (MD ANEXT) loss and multiple disturber alien FEXT (MDAFEXT) loss is specified. In addition, to ensure the reliable operation of the link segment, a minimum alien crosstalk to insertion loss ratio are specified.

55.7.3.1 Multiple disturber alien near-end crosstalk (MDANEXT) loss

In order to limit the alien crosstalk at the near end of a link segment, the differential pair-to-pair near-end crosstalk (NEXT) loss between the disturbed duplex channel in a link segment and the disturbing duplex channels in other link segments is specified to meet the bit error ratio objective specified in 55.1.1. To ensure the total alien NEXT coupled into a duplex channel is limited, multiple disturber alien NEXT loss is specified as the power sum of the individual alien NEXT disturbers.

55.7.3.1.1 Multiple disturber power sum alien near-end crosstalk (PSANEXT) loss

PS ANEXT loss is determined by summing the power of the individual pair-to-pair differential alien NEXT loss values over the frequency range 1 MHz to 500 MHz as follows in Equation (55–22):

$$\mathrm{PSANEXT}_N(f) = -10\log_{10}\sum_{j=1}^{m}\sum_{i=1}^{4} 10^{\frac{-AN(f)_{i,j,N}}{10}} \quad (\mathrm{dB}) \tag{55–22}$$

where

$AN(f)_{i,j,N}$ is the magnitude in dB of the alien NEXT loss at frequency f of the individual pair combination i (1 to 4) of the disturbing link j (1 to m) for each disturbed pair N.

Annex 55B provides additional information on identifying the number of adjacent link segments to consider in the PSANEXT calculation.

The power sum ANEXT loss between a disturbed duplex channel in a link segment and the disturbing duplex channels in other link segments shall meet the values determined using Equation (55–23) independent of length.

$$\mathrm{PSANEXT}(f) \geq \begin{cases} X1 - 10\log_{10}\left(\frac{f}{100}\right) \ (\mathrm{dB}) & 1 \leq f \leq 100 \\ X1 - 15\log_{10}\left(\frac{f}{100}\right) \ (\mathrm{dB}) & 100 < f \leq 500 \end{cases} \tag{55–23}$$

where

f is the frequency in MHz,
$X1$ is the value at f=100 MHz. $X1$ is referred to as the PSANEXT constant.

When the computed PSANEXT value at a certain frequency exceeds 67 dB, the PSANEXT result at that frequency is for information only.

The PSANEXT constant is determined in 55.7.3.1.2 constrained by the ratio of the PSANEXT to the insertion loss.

The average PSANEXT between a disturbed duplex channel in a link segment and the disturbing duplex channels in other link segments is the average of the individual-pair PSANEXT “across the 4-pairs” determined using Equation (55–24):

$$\text{PSANEXT}avg(f) = \frac{\left(\sum_{i=1}^{n} \text{PSANEXT}(f)\_\text{pair\_i (dB)} \right)}{4} \quad (55\text{–}24)$$

where

PSANEXT*(f)*_pair_i is the magnitude in dB of PSANEXT loss at frequency *f* of pair_i,
n is the number of individual-pairs (*n* = 4).

The average PSANEXT loss of the 4-pairs shall meet the values determined using Equation (55–25) independent of length.

$$\text{PSANEXT}avg(f) \geq \begin{cases} (X1 + 2.25) - 10\log_{10}(\frac{f}{100}) \text{ (dB)} & 1 \leq f \leq 100 \\ (X1 + 2.25) - 15\log_{10}(\frac{f}{100}) \text{ (dB)} & 100 < f \leq 500 \end{cases} \quad (55\text{–}25)$$

where

f is the frequency in MHz,
X1 is the value at *f* =100 MHz. *X1* is referred to as the PSANEXT constant.

55.7.3.1.2 PSANEXT loss to insertion loss ratio requirements

To ensure reliable operation, a minimum insertion loss to alien crosstalk ratio shall be maintained. The PSANEXT loss requirement of 55.7.3.1.1 can be relaxed based on a reduction in the maximum insertion loss specified in 55.7.2.1. The insertion loss reduction can be achieved by utilizing link segments less than 100 m defined by (55–27) or the insertion loss at 250 MHz of the supported cabling types in Table 55–17.

The adjusted PSANEXT loss requirement is determined by first calculating the PSANEXT constant and utilizing the constant in the PSANEXT limit line Equation (55–23) and Equation (55–25).

The PSANEXT_constant shall be specified by the following Equation (55–26):

$$\text{PSANEXT\_constant} = 62 - \frac{10GBTIL(250\text{ MHz}) - \text{IL}(250\text{ MHz})}{1.04} \text{ (dB)} \quad (55\text{–}26)$$

where

10*GBTIL(*250*MHz)* is the link segment insertion loss at 250 MHz [Equation (55–11)],
IL(250*MHz*) is the insertion loss of a link segment less than 100 m [Equation (55–27)] or the insertion loss at 250 MHz of the supported cabling types in Table 55–17. For measurement-based calculations, IL(250 MHz) shall be the average of the insertion loss of the 4-pairs at 250 MHz.

The calculated PSANEXT_constant values less than 33.5 dB shall revert to a value of 33.5 dB.

For the purpose of calculating the link segment insertion loss for cabling less than 100 m the cable insertion loss is assumed to scale linearly with length.

The scaled link segment is defined by the following Equation (55–27):

$$\text{Scaled\_10GBTIL} = \frac{\text{Length\_m}}{100} \times 1.05\left(1.82 \times \sqrt{f} + 0.0169 \times f + \frac{0.25}{\sqrt{f}}\right) + 4 \times 0.02 \times \sqrt{f} \quad \text{(dB)} \qquad (55\text{–}27)$$

where

Length_m is in meters,
f is the frequency in MHz.

Table 55–18 provides the calculated PSANEXT_constants for link segments of 100 m and 55 m.

Table 55–18— Calculated PS ANEXT constants

Link segment distance	PSANEXT_constant (dB)	PSANEXT_constant_avg average of the 4-pairs (dB)	Insertion loss at 250 MHz (dB)
100 m	62	64.25	35.9
55 m	47	49.25	20.3

Table 55–19 provides the calculated PSANEXT_constants for the supported cabling types.

Table 55–19— Calculated PSANEXT constants for supported cabling types

Cabling	Insertion loss at 250 MHz (dB)	PSANEXT_constant (dB)	PSANEXT_constant_avg average of the 4-pairs (dB)
Category 6/ Class E	35.9[a], 20.3[b]	62[a], 47[b]	64.25[a], 49.25[b]
Class E_A / Category 6A	33.8	60	62.25
Class F	33.8	60	62.25

[a]Link segment length of 100 m.
[b]Link segment length of 55 m.

Table 55–20 provides the calculated PSANEXT to insertion loss ratio for the supported cabling types.

Table 55–20— Calculated PSANEXT to insertion loss ratio at 250 MHz

Cabling	Insertion loss at 250 MHz (dB)	PSANEXT at 250 MHz (dB)	PSANEXT to insertion loss ratio at 250 MHz (dB)
Category 6/ Class E	35.9[a], 20.3[b]	56.0[a], 41.0[b]	20.1[a], 20.7[b]
Class E_A / Category 6A	33.8	54.0	20.2
Class F	33.8	54.0	20.2

[a]Link segment length of 100 m.
[b]Link segment length of 55 m.

For simulating PHY performance to estimate system margin, the PSANEXT_constant is increased by 3.5 dB to account for an averaging of the PSANEXT over frequency and averaging the PSANEXT *across the 4-pairs*.

55.7.3.2 Multiple disturber alien far-end crosstalk (MDAFEXT) loss

In order to limit the alien crosstalk at the far-end of a link segment, the differential pair-to-pair alien far-end crosstalk loss between the disturbed duplex channel in a link segment and the disturbing duplex channels in other link segments is specified to meet the bit error ratio objective specified in 55.1.1. To ensure the total alien FEXT coupled into a duplex channel is limited, multiple disturber ELFEXT is specified as the power sum of the individual alien ELFEXT disturbers.

55.7.3.2.1 Multiple disturber power sum alien equal level far-end crosstalk (PSAELFEXT)

Power sum alien ELFEXT is determined by summing the power of the individual pair-to-pair differential alien ELFEXT values over the frequency range 1 MHz to 500 MHz as follows in Equation (55–28):

$$\mathrm{PSAELFEXT}_N(f) = -10\log_{10}\sum_{j=1}^{m}\sum_{i=1}^{4}10^{\frac{-EL(f)_{i,j,N}}{10}} \quad \text{(dB)} \tag{55–28}$$

where

$EL(f)_{i,j,N}$ is determined using Equation (55–29) as the magnitude in dB of the alien ELFEXT of the coupled length at frequency f of the individual pair combination i (1 to 4) of the disturbing link j (1 to m) for each disturbed pair N corrected by subtracting the $10\log_{10}$ ratio of the disturbed length insertion loss to the coupled length insertion loss. The coupled length is the length of cabling over which the crosstalk coupling can occur.

$$EL(f)_{i,j,N} = AlienELFEXT(f)_{i,j,N} - 10\log_{10}\left(\frac{DisturbedIL_N}{CoupledlengthIL_{i,j,N}}\right) \tag{55–29}$$

where

$AlienELFEXT(f)_{i,j,N}$ is determined using Equation (55–30) as the difference of the magnitude in dB of the Alien FEXT of the coupled length at frequency f of the individual pair combination i (1 to 4) of the disturbing link j (1 to m) for each disturbed pair N and the insertion loss of the coupled length.

$$AlienELFEXT(f)_{i,j,N} = AFEXT(f)_{i,j,N} - CoupledlengthIL(f)_{i,j,N} \tag{55–30}$$

where

$CoupledlengthIL(f)_{i,j,N}$ is determined using Equation (55–31) as the minimum of the insertion loss of the disturbed pair N and the disturbing individual pair i (1 to 4) of the disturbing link j (1 to m).

$$CoupledlengthIL(f)_{i,j,N} = min(DisturbedIL_N, DisturbingIL_{i,j}) \tag{55–31}$$

NOTE—The computation of PSAELFEXT in Equation (55–28) is consistent with the computation of power sum alien attenuation to crosstalk ratio far end (PS AACR-F). The term PS AACR-F is used in ISO/IEC TR 24750 and in the 1st amendment to the second edition of ISO/IEC 11801.

The PSAELFEXT between a disturbed duplex channel in a link segment and the disturbing duplex channels in other link segments shall meet the values determined using Equation (55–32).

$$\text{PSAELFEXT}(f) \geq X2 - 20\log_{10}\left(\frac{f}{100}\right) \quad (dB) \qquad 1 \leq f \leq 500 \tag{55–32}$$

where

f is the frequency in MHz,
X2 is the value at *f* =100 MHz. *X2* is referred to as the PSAELFEXT constant.

When the measured PSAFEXT limit value at a certain frequency exceeds 72-15log$_{10}$(*f*/100) dB, 67 dB max, the PSAELFEXT result at that frequency is for information only.

The average PSAELFEXT between a disturbed duplex channel in a link segment and the disturbing duplex channels in other link segments is the average of the individual-pair PS AELFEXT *across the 4-pairs* determined using Equation (55–33).

$$\text{PSAELFEXT}avg(f) = \frac{\left(\sum_{i=1}^{n} \text{PSAELFEXT}(f)\_\text{pair\_i (dB)}\right)}{4} \tag{55–33}$$

where

PSAELFEXT*(f)*_pair_i is the magnitude in dB of PSAELFEXT loss at frequency *f* of pair *i*,
n is the number of individual-pairs (*n* = 4).

The average PSAELFEXT of the 4-pairs shall meet the values determined using Equation (55–34).

$$\text{PSAELFEXT}avg(f) \geq (X2 + 4) - 20\log_{10}\left(\frac{f}{100}\right) \quad (\text{dB}) \qquad 1 \leq f \leq 500 \tag{55–34}$$

where

f is the frequency in MHz,
X2 is the value at *f* =100 MHz. The value at *f* =100 MHz is referred to as the PSAELFEXT constant.

The PSAELFEXT constant is determined in 55.7.3.2.2 by the ratio of the PS AELFEXT to the insertion loss and a length correction term. The insertion loss dependent term maintains the required SNR ratio and the IL independent term enables a relaxation of the PS AELFEXT due the reduction in crosstalk coupling for lengths less than 100 m.

55.7.3.2.2 PSAELFEXT to insertion loss ratio requirements

To ensure reliable operation, a minimum insertion loss to alien crosstalk ratio shall be maintained. The PS AELFEXT requirement of 55.7.3.2.1 can be relaxed based on a reduction in the maximum insertion loss specified in 55.7.2.1 and a reduction in the AELFEXT coupling for cabling lengths less than 100 m.

The adjusted PSAELFEXT requirement is determined by first calculating the PSAELFEXT constant and utilizing the constant in the PSAELFEXT limit line model of Equation (55–35).

The PSAELFEXT_constant shall be specified by the following equation:

$$\text{PSAELFEXT\_constant} = 37.9 - \left(\frac{10GBTIL(250\ \text{MHz}) - \overline{\overline{IL(250\ \text{MHz})}}}{2.29} \right) - 10\log_{10}\left(\frac{L}{100}\right) \quad (dB) \qquad (55\text{–}35)$$

where

$10GBTIL($250 MHz) is the link segment insertion loss at 250 MHz [Equation (55–11)],
IL(250 MHz) is the insertion loss of a link segment less than 100 m [Equation (55–27)] or the insertion loss of the supported cabling types in Table 55–17 at 250 MHz. For measurement based calculations, IL(250 MHz) shall be the average of the insertion loss of the 4-pairs at 250 MHz.
L is the length in meters of the link segment. The equation assumes coupling over 100 m of horizontal cable and cable cords. For measurement based calculations (e.g., field testing), L can derived from the measured insertion loss and substituted in Equation (55–35) as defined in Equation (55–36).

$$L = 2.77 \times (\text{measured insertion loss at 250 MHz}) \qquad (55\text{–}36)$$

For insertion loss values less than 10 dB at 250 MHz the calculated PSAELFEXT_constant values that exceed 32.5 dB shall revert to a value of 32.5 dB.

The field testing of length and insertion loss are addressed in TIA TSB-155-A and ISO/IEC TR 24750.

Table 55–21 provides the calculated PSAELFEXT_constants for link segments of 100 m and 55 m.

Table 55–21—Calculated PS AELFEXT constants

Link segment distance	PS AELFEXT_constant (dB)	PS AELFEXT_constant_avg average of the 4-pairs (dB)	Insertion loss at 250 MHz (dB)
100 m	37.9	41.9	35.9
55 m	33.7	37.7	20.3

Table 55–22 provides the calculated PSAELFEXT_constants for the supported cabling types.

Table 55–22—Calculated PSAELFEXT constants for supported cabling types

Cabling	Insertion loss at 250 MHz (dB)	PSAELFEXT_constant (dB)	PSAELFEXT_constant_avg average of the 4-pairs (dB)
Category 6/ Class E	35.9[a], 20.3[b]	37.9[a], 33.7[b]	41.9[a], 37.7[a]
Class E_A / Category 6A	33.8	37	41
Class F	33.8	37	41

[a]Link segment length of 100 m.
[b]Link segment length of 55 m.

Table 55–23 provides the calculated PS AELFEXT at 250 MHz for the supported cabling types.

Table 55–23—Calculated PS AELFEXT at 250 MHz for supported cabling types

Cabling	Insertion loss at 250 MHz (dB)	PS AELFEXT at 250 MHz (dB)
Category 6/ Class E	35.9[a], 20.3[b]	29.9[a], 25.7[b]
Class E_A / Category 6A	33.8	29.0
Class F	33.8	29.0

[a]Link segment length of 100 m.
[b]Link segment length of 55 m.

55.7.3.3 Alien crosstalk margin computation

The objective of alien crosstalk margin computation is to further characterize the alien crosstalk coupling "between" link segments. The alien crosstalk margin computation ensures the total combined PSAFEXT and PSANEXT coupled into a duplex channel is limited in order to maintain the minimum signal-to-noise ratio. The alien crosstalk margin computation can be applied in the event that the PSANEXT limits specified in 55.7.3.1.1 [Equation (55–23) and Equation (55–25)] or the PSAELFEXT limits specified in 55.7.3.2. [Equation (55–32) and Equation (55–34)] are not met. The alien crosstalk margin is specified for each of the individual 4-pairs as well as the average "across the 4-pairs".

The alien crosstalk margin is determined by the following algorithm:

Step 1. Determine the length of the disturbed link segment and the disturbing link segments using Equation (55–37).

$$L = 277 \times \mathrm{IL\_avg} \tag{55–37}$$

where

IL_*avg*is the average measured insertion loss at 250 MHz "across the 4-pairs" of each disturbed and disturbing link segment.
L is the length in meters for each disturbed and disturbing link segment derived from the measured insertion loss.

Step 2. Determine the minimum power backoff (dB) for each disturbed and disturbing link segment from Table 55–11. Power backoff schedule utilizing the calculated link segment length from Step 1 (e.g., from Table 55–11; a length of 30 m has a minimum power backoff of 10 dB).

Step 3. Determine the insertion loss backoff factor using Equation (55–38).

$$\mathrm{IL\_bof}_j = \mathrm{min\_PBO\_disturbing\ link}_j - \mathrm{min\_PBO\_disturbed\ link} \tag{55–38}$$

where

IL_*bof*$_j$is calculated as the difference between the minimum power backoff (dB) of the disturbing link j and the disturbed link segments determined in Step 2.

Step 4. Determine the PSANEXT and the PSAFEXT from the measured ANEXT and AFEXT and the insertion loss backoff factor for each disturbed pair N = 1, 2, 3, 4 of a link segment using Equation (55–39) and Equation (55–40) respectively.

$$\text{PSANEXT}_N(f) = -10\log_{10}\sum_{j=1}^{m}\sum_{i=1}^{4} 10^{\frac{-(AN\_\text{pr(f)}_{i,j,N} + IL\_\text{bof}_j)}{10}} \quad \text{(dB)} \qquad (55\text{–}39)$$

$$\text{PSAFEXT}_N(f) = -10\log_{10}\sum_{j=1}^{m}\sum_{i=1}^{4} 10^{\frac{-(AF\_\text{pr(f)}_{i,j,N} + IL\_\text{bof}_j)}{10}} \quad \text{(dB)} \qquad (55\text{–}40)$$

where

$AN\_pr(f)_{i,j,N}$ is the measured ANEXT of the individual pair combination i (1 to 4) of the disturbing link j (1 to m) for each disturbed pair N.
$AF\_pr(f)_{i,j,N}$ is the measured AFEXT of the individual pair combination i (1 to 4) of the disturbing link j (1 to m) for each disturbed pair N.
IL_bof_j is the insertion loss backoff factor determined in Equation (55–38).

Step 5. Determine the individual-pair margin for each of the 4-pairs using Equation (55–41).

$$XWn(f) = -10\log_{10}\left(10^{\frac{\text{PSANEXT}_N(f)}{-10}} + 10^{\frac{\text{PSAFEXT}_N(f)}{-10}}\right) + 10\log_{10}\left(10^{\frac{AN\_ipl(f)+2.5}{-10}} + 10^{\frac{AF\_ipl(f)}{-10}}\right) \quad \text{(dB)} \qquad (55\text{–}41)$$

where: f is the frequency in MHz,

$XWn(f)$ is the individual-pair margin for each of the 4-pairs,

$\text{PSANEXT}_N(f)$ is the measured PSANEXT loss in dB in the frequency range $10 \le f \le 400$ determined in Equation (55–39) adjusted for the insertion loss power backoff,

$\text{PSAFEXT}_N(f)$ is the measured PSAFEXT loss in dB in the frequency range $10 \le f \le 400$ determined in Equation (55–40) adjusted for the insertion loss power backoff,

$AN\_ipl(f)$ is the individual-pair limit line for PSANEXT calculated using Equation (55–42). $AN\_ipl(f)$ is derived utilizing the average of the insertion loss of the 4-pairs at 250 MHz determined in Step 1 to calculate the PSANEXT constant using Equation (55–26).

$$AN\_\text{ipl}(f) = \begin{cases} \text{PSANEXTconstant} - 10\log_{10}\left(\frac{f}{100}\right) & \text{(dB)} \quad 1 \le f \le 100 \\ \text{PSANEXTconstant} - 15\log_{10}\left(\frac{f}{100}\right) & \text{(dB)} \quad 100 < f \le 500 \end{cases} \qquad (55\text{–}42)$$

$AF\_ipl(f)$ is the individual-pair line for PSAFEXT calculated using Equation (55–43) $AF\_ipl(f)$ is derived by adding the average of the measured insertion loss of the 4-pairs at 250 MHz determined in Step 1 to the PSAELFEXT limit line. The PSAELFEXT constant is derived using Equation (55–32) utilizing the average of the insertion loss of the 4-pairs at 250 MHz determined in Step 1.

$$AF\_\text{ipl(f)} = (\text{ PSAEFEXT constant }) - 20\log_{10}(\frac{f}{100}) + IL\_\text{avg}(f)\ \ (dB)\ \ \ 1 \le f \le 500 \qquad (55\text{–}43)$$

NOTE—The 2.5 dB is the PSANEXT allowance for the peak-to-average difference across frequency.

Step 6. Determine the average value "across frequency" of XWn(f) from 10 MHz to 400 MHz, for each individual-pair of the 4-pair cabling using Equation (55–44).

$$XAWn = \frac{\int_{f=10}^{400} XWn(f)df}{390}(\text{dB}) \qquad (55\text{–}44)$$

where

f is the frequency in MHz,

XAWn is the average value "across frequency" for the individual-pair number n (n = 1,2,3,4).

Step 7. Determine the individual-pair margin as the minimum of the average value "across frequency" of each of the individual pairs of the 4-pair cabling determined from Step 2 using Equation (55–45).

$$Y_{\text{inp}} = \min(XAW1, XAW2, XAW3, XAW4)(\text{dB}) \qquad (55\text{–}45)$$

where

Y_{inp} is the individual-pair margin.

Step 8. Determine the power sum of the PSANEXT and PSAFEXT for each of the 4-pairs using Equation (55–46).

$$XXn(f) = 10\log_{10}\left(10^{\frac{AN(f)}{-10}} + 10^{\frac{AF(f)}{-10}}\right)(\text{dB}) \qquad (55\text{–}46)$$

where

f is the frequency in MHz,

XXn(f) is the power sum of the PSANEXT and PSAFEXT for each of the 4-pairs,

AN(f) is the measured PSANEXT loss in dB in the frequency range $10 \le f \le 400$ determined in Equation (55–39) adjusted for the insertion loss power backoff,

AF(f) is the measured PSAFEXT loss in dB in the frequency range $10 \le f \le 400$ determined in Equation (55–40) adjusted for the insertion loss power backoff.

Step 9. Determine the average of the power sum of the PSANEXT and PSAFEXT of the 4-pairs using Equation (55–47).

$$XXavg(f) = \frac{\left(\sum_{i=1}^{n} XXi(f)\right)}{4} (\text{dB}) \quad (55\text{–}47)$$

where

XXavg(f) is the average of the power sum of the PSANEXT and PSAFEXT of the 4-pairs,

n is the number of individual pairs; $n = 4$.

Step 10. Determine the average margin using Equation (55–48).

$$XA(f) = -XXavg + 10\log_{10}\left(10^{\frac{AN\_avgl(f)}{-10}} + 10^{\frac{AF\_avgl(f)}{-10}}\right) \quad (\text{dB}) \quad (55\text{–}48)$$

where

XA(f) is the average margin,

AN_avgl(f) is the average limit line for PSANEXT calculated using Equation (55–49). AN_avgl(f) is derived using the PSANEXT constant determined in Step 5.

$$AN\_\text{avgl(f)} = \begin{cases} PSANEXT_{\text{constant}} + 3.5 - 10\log_{10}\left(\frac{f}{100}\right) & (\text{dB}) \quad 1 \le f \le 100 \\ PSANEXT_{\text{constant}} + 3.5 - 15\log_{10}\left(\frac{f}{100}\right) & (\text{dB}) \quad 100 < f \le 500 \end{cases} \quad (55\text{–}49)$$

where

f is the frequency in MHz,

AF_avgl(f) is the average limit line for PSAFEXT calculated using Equation (55–50). *AF_avgl(f)* is derived by adding the average of the measured insertion loss of the 4-pairs at 250 MHz determined in Step 1 to the PSAELFEXT limit line using the PSAELFEXT constant determined in Step 5.

$$AF\_avg1(f) = (\text{PSAELFEXT constant} + 4) - 20\log_{10}\left(\frac{f}{100}\right) + IL\_\text{avg}(f) \quad (\text{dB}) \quad 1 \le f \le 500 \quad (55\text{–}50)$$

Step 11. Determine the average margin as the average value across frequency of *XA(f)* using Equation (55–51).

$$Y_{\text{avg}} = \frac{\int_{f=10}^{400} XAfdf}{390} (\text{dB}) \quad (55\text{–}51)$$

where

f is the frequency in MHz,
Y_{avg} is the average value "across frequency" of $XA(f)$ from 10 to 400 MHz.

Step 12. Determine the alien crosstalk margin as the minimum value of the individual pair margin [Equation (55–45)] and the average margin [Equation (55–51)] using Equation (55–52).

$$YL = min(Y_{inp}, Y_{avg}) \text{ (dB)} \quad (55–52)$$

The alien crosstalk margin YL shall be greater than zero.

55.7.4 Noise environment

The 10GBASE-T noise environment consists of noise from many sources. The primary noise sources that impact the objective BER are the crosstalk and echo interference of a link segment, which are reduced to a small residual noise, and the noise coupled between the link segments referred to as alien crosstalk noise. The remaining noise sources, which are secondary sources, are discussed in the following. The 10GBASE-T noise environment consists of the following:

a) Echo from the local transmitter on the same duplex channel (cable pair). Echo is caused by the hybrid function used to achieve simultaneous bi-directional transmission of data and by impedance mismatches in the link segment. It is impractical to achieve the objective BER without using echo cancellation. Since the symbols transmitted by the local disturbing transmitter are available to the cancellation processor, echo interference can be reduced to a small residual noise using echo cancellation methods.

b) Near-end crosstalk (NEXT) interference from the local transmitters on the duplex channels (cable pairs) of the link segment. Each receiver experiences NEXT interference from three adjacent transmitters. NEXT cancellers are used to reduce the interference from each of the three disturbing transmitters to a small residual noise. NEXT cancellation is possible since the symbols transmitted by the three disturbing local transmitters are available to the cancellation processor.

c) Far-end crosstalk (FEXT) noise at a receiver is from three disturbing transmitters at the far end of the duplex channel (cable pairs) of the link segment. FEXT noise can be reduced through cross coupled equalizers although the symbols from the remote transmitters are not immediately available.

d) Intersymbol interference (ISI). ISI is the extraneous energy from one signaling symbol that interferes with the reception of another symbol on the same channel. 10GBASE-T supports the use of Tomlinson-Harashima Precoding as a mechanism to reduce the effects of ISI.

e) Noise from non-idealities in the duplex channel, transmitters, and receivers; for example, DAC/ADC non-linearity, electrical noise (shot and thermal), and non-linear channel characteristics. 10GBASE-T limits the effects of some of these non-idealities by a variety of PMA electrical specifications.

f) Noise coupled between link segments. Noise coupled between the disturbed duplex channel in a link segment and the disturbing duplex channels in other link segments is referred to as alien crosstalk noise. Since the transmitted symbols from the alien crosstalk noise sources are not available to the cancellation processor (they are in another cable), it is very difficult to cancel the alien crosstalk noise. To ensure robust operation the alien crosstalk is specified in 55.7.3.

g) The background noise for 10GBASE-T is expected not to exceed –150 dBm/Hz. A background noise limit of –150 dBm/Hz was assumed for determining the minimum signal-to-noise ratio.

55.8 MDI specification

This subclause defines the MDI. The link topology requires a crossover function in a DTE-to-DTE connection. See 55.4.4 for a description of the automatic MDI/MDI-X configuration.

55.8.1 MDI connectors

Eight-pin connectors meeting the requirements of IEC 60603-7-4 (unscreened) or IEC 60603-7-5 (screened) shall be used as the mechanical interface to the balanced cabling. The plug connector shall be used on the balanced cabling and the jack on the PHY. These connectors are depicted (for informational use only) in Figure 55–39 and Figure 55–40. The assignment of PMA signals to connector contacts for PHYs is shown in Table 55–24.

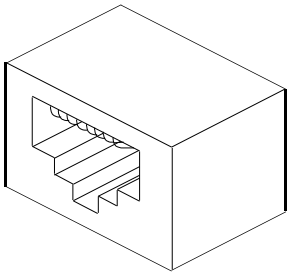

Figure 55–39—MDI connector

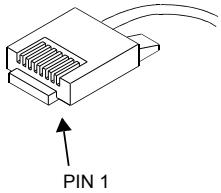

Figure 55–40—Balanced cabling connector

Table 55–24—Assignment of PMA signal to MDI and MDI-X pin-outs

Contact	MDI	MDI-X
1	BI_DA+	BI_DB+
2	BI_DA–	BI_DB–
3	BI_DB+	BI_DA+
4	BI_DC+	BI_DD+
5	BI_DC–	BI_DD–
6	BI_DB–	BI_DA–
7	BI_DD+	BI_DC+
8	BI_DD–	BI_DC–

55.8.2 MDI electrical specifications

The MDI connector (jack) when mated with a specified balanced cabling connector (plug) shall meet the FEXT requirements as specified in Equation (55–53):

$$\text{FEXT loss} \geq 43.1 - 20\log_{10}(f/100) \quad (\text{dB}) \tag{55–53}$$

where f is the frequency over the range 1 MHz to 500 MHz between all contact pair combinations shown in Table 55–24.

No spurious signals shall be emitted onto the MDI when the PHY is held in power-down mode (as defined in 22.2.4.1.5), when released from power-down mode, or when external power is first applied to the PHY.

55.8.2.1 MDI return loss

The differential impedance at the MDI for each transmit/receive channel shall be such that any reflection due to differential signals incident upon the MDI from a balanced cabling having a nominal differential characteristic impedance of 100 Ω is attenuated, relative to the incident signal as per the relationship:

$$\text{Return loss} \geq \begin{cases} 16 & 1 \leq f \leq 40 \quad (\text{dB}) \\ 16 - 10\log_{10}(f/40) & 40 < f \leq 400 \quad (\text{dB}) \\ 6 - 30\log_{10}(f/400) & 400 < f \leq 500 \quad (\text{dB}) \end{cases} \tag{55–54}$$

where f is in MHz.

55.8.2.2 MDI impedance balance

Impedance balance is a measure of the impedance-to-ground difference between the two MDI contacts used by a duplex link channel and is referred to as common-mode-to-differential-mode impedance balance. The common-mode-to-differential-mode impedance balance, *Bal*(*f*), of each channel of the MDI shall meet the relationship:

$$Bal(f) \geq \begin{cases} 48 & 1 \leq f < 30 \quad (\text{dB}) \\ 44 - 19.2\log_{10}\left(\frac{f}{50}\right) & 30 \leq f \leq 500 \quad (\text{dB}) \end{cases} \tag{55–55}$$

where f is the frequency in MHz when the transmitter is transmitting random or pseudo random data. Test-mode 4 may be used to generate an appropriate transmitter output. The impedance balance (in dB) is defined as:

$$Bal(f) = 20\log_{10}\left(\frac{E_{cm}(f)}{E_{dif}(f)}\right) - IL_{cal}(f) \quad (\text{dB}) \tag{55–56}$$

where E_{cm} is an externally applied common-mode sinusoidal voltage as shown in Figure 55–41, E_{dif} is the resulting differential waveform due only to the applied sine wave E_{cm} and not the transmitted data, and $IL_{cal}(f)$ is the coupling circuit common-mode insertion loss (in dB). The coupling circuit common-mode insertion loss is defined as:

$$IL_{\text{cal}}(f) = 20\log_{10}\left(\frac{E_{cin}(f)}{E_{cout}(f)}\right) \quad (\text{dB}) \tag{55–57}$$

where E_{cin} and E_{cout} are test signals used to measure the coupling circuit common-mode insertion loss as shown in Figure 55–42.

The impedance balance is measured with a common-mode coupling circuit presenting a differential termination of 100 Ω and a common-mode termination of 50 Ω to the MDI pair under test as shown in Figure 55–41. The return loss for both differential and common terminations in the common-mode coupling circuit should be greater than 20 dB from 10 MHz to 500 MHz. The insertion loss IL_{cal} (in dB) of the common-mode coupling circuit can be measured with the test circuit shown in Figure 55–42.

NOTE 1—Triggered averaging can be used to separate the component due to the applied common-mode sine wave from the transmitted data component.
NOTE 2—The imbalance of the test equipment should be insignificant relative to the balance requirements.

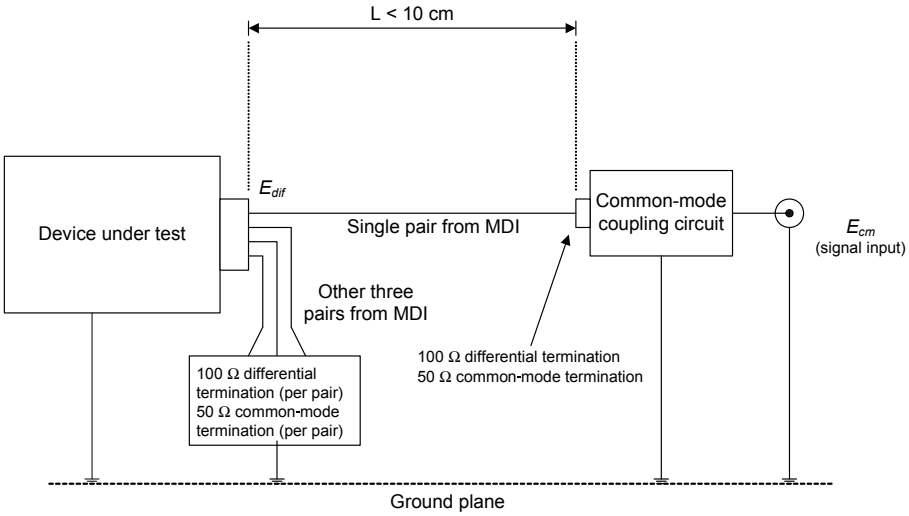

Figure 55–41—MDI impedance balance test circuit

The impedance balance may also be measured with a mixed mode four port network analyzer capable of measuring the common-mode voltage and differential mode voltage of a balanced port. Impedance balance is the S parameter measurement of Scd11 in dB at the MDI where two ports of the four port network analyzer are connected between two MDI contacts used by a duplex link channel and these two ports are configured as a single balanced port. For this test the PHY ground is connected to the network analyzer ground. The other two ports of the network analyzer are unconnected. The network analyzer should be capable of measuring Scd11 to at least –60 dB and should use a differential input impedance of 100 Ω and a common-mode impedance of 75 Ω.

During the test the PHY is connected to the MDI as in normal operation, but with the transmitter output disabled.

55.8.2.3 MDI fault tolerance

Each wire pair of the MDI shall, under all operating conditions, withstand without damage the application of short circuits of any wire to any other wire within the 4-pair cable for an indefinite period of time and shall resume normal operation after the short circuit(s) are removed. The magnitude of the current through such a short circuit shall not exceed 300 mA.

A 10GBASE-T PHY shall be able to sustain, without damage, connection to a PSE and shall not cause damage to the PSE as defined in 33.2.

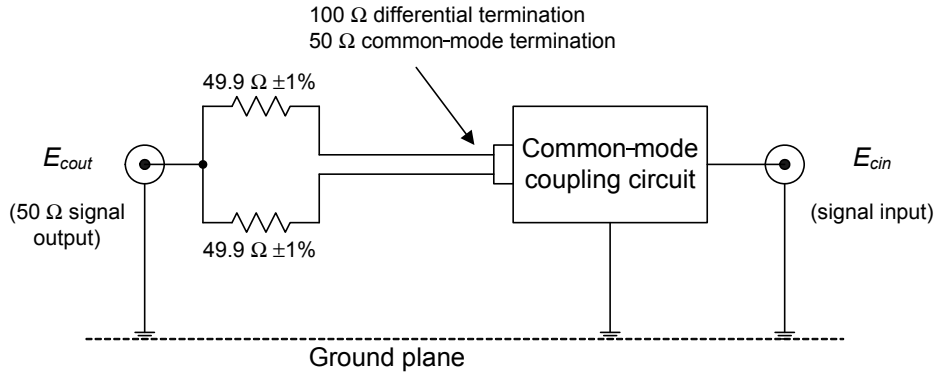

Figure 55–42— Impedance balance calibration circuit

Each wire pair shall withstand without damage a 1000 V common-mode impulse of either polarity. The shape of the impulse is 0.3/50 μs (300 ns virtual front time, 50 μs virtual time of half value), as defined in IEC 60060. See Figure 40–33 in 40.8.3.4.

55.9 Environmental specifications

55.9.1 General safety

All equipment meeting this standard shall conform to IEC 60950-1.

55.9.2 Network safety

This subclause sets forth a number of recommendations and guidelines related to safety concerns; the list is neither complete nor does it address all possible safety issues. The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

LAN cabling systems described in this subclause are subject to at least four direct electrical safety hazards during their installation and use. These hazards are as follows:

a) Direct contact between LAN components and power, lighting, or communications circuits.
b) Static charge buildup on LAN cabling and components.
c) High-energy transients coupled onto the LAN cabling system.
d) Voltage potential differences between safety grounds to which various LAN components are connected.

Such electrical safety hazards should be avoided or appropriately protected against for proper network installation and performance. In addition to provisions for proper handling of these conditions in an operational system, special measures should be taken to ensure that the intended safety features are not negated during installation of a new network or during modification or maintenance of an existing network.

55.9.3 Installation and maintenance guidelines

It is recommended that sound installation practice, including screen management, as defined by applicable local codes and regulations, be followed in every instance in which such practice is applicable. In addition, Annex 55B provides additional cabling guidelines for 10GBASE-T deployment on balanced copper cabling systems.

It is a mandatory requirement that, during installation of the cabling plant, care be taken to ensure that non-insulated network cabling conductors do not make electrical contact with unintended conductors or ground.

55.9.4 Telephone voltages

The use of building wiring brings with it the possibility of wiring errors that may connect telephony voltages to 10GBASE-T. Other than voice signals (which are low voltage), the primary voltages that may be encountered are the "battery" and ringing voltages. Although there is no universal standard, the following maximums generally apply.

a) Battery voltage to a telephone line is generally 56 V dc applied to the line through a balanced 400 Ω source impedance.
b) Ringing voltage is a composite signal consisting of an ac component and a dc component. The ac component is up to 175 V peak at 20 Hz to 60 Hz with a 100 Ω source resistance. The dc component is 56 V dc with a 300 Ω to 600 Ω source resistance. Large reactive transients can occur at the start and end of each ring interval.

Although 10GBASE-T equipment is not required to survive such wiring hazards without damage, application of any of the above voltages shall not result in any safety hazard.

55.9.5 Electromagnetic compatibility

A system integrating the 10GBASE-T shall comply with applicable local and national codes for the limitation of electromagnetic interference.

55.9.6 Temperature and humidity

A system integrating the 10GBASE-T is expected to operate over a reasonable range of environmental conditions related to temperature, humidity, and physical handling (such as shock and vibration). Specific requirements and values for these parameters are considered to be beyond the scope of this standard.

55.10 PHY labeling

It is recommended that each PHY (and supporting documentation) be labeled in a manner visible to the user with at least the following parameters:

a) Data rate capability and units thereof
b) Power level in terms of maximum current drain (for external PHYs)
c) Port type (i.e., 10GBASE-T)
d) Any applicable safety warnings
e) EEE support
f) Fast retrain support

55.11 Delay constraints

In full duplex mode, predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) also demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementers conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices.

The sum of the transmit and receive data delays for an implementation of a 10GBASE-T PHY shall not exceed 25 600 BT. Transmit data delay is measured from the input of a given unit of data at the XGMII to the presentation of the same unit of data by the PHY to the MDI. Receive data delay is measured from the

input of a given unit of data at the MDI to the presentation of the same unit of data by the PHY to the XGMII.

NOTE—The physical medium interconnecting two PHYs introduces additional delay in a link. Equation (44–1) specifies the calculation of bit time per meter of electrical cable and Table 44–3 can also be used to convert electrical cable delay values specified relative to the speed of light or in nanoseconds per meter.

55.12 Protocol implementation conformance statement (PICS) proforma for Clause 55—Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer and baseband medium, type 10GBASE-T[22]

The supplier of a protocol implementation that is claimed to conform to this clause shall complete the protocol implementation conformance statement (PICS) proforma listed in the following subclauses.

Instructions for interpreting and filling out the PICS proforma may be found in Clause 21.

55.12.1 Identification

55.12.1.1 Implementation identification

Supplier	
Contact point for inquiries about the PICS	
Implementation Name(s) and Version(s)	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)	
NOTE 1—Only the first three items are required for all implementations; other information may be completed as appropriate in meeting the requirements for the identification. NOTE 2—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

55.12.1.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 55, Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer, and baseband medium, type 10GBASE-T
Identification of amendments and corrigenda to this PICS proforma which have been completed as part of this PICS	
Have any Exceptions items been required? No [ ] Yes [ ] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2015.)	
Date of Statement	

[22]*Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so it can be used for its intended purpose and may further publish the completed PICS.

55.12.2 Major capabilities/options

Item	Feature	Subclause	Status	Support	Value/Comment
XGE	XGMII compatibility interface	46, 55.1.5	O	Yes [ ] No [ ]	Compatibility interface is supported
*LT	Support of loop timing		O	Yes [ ] No [ ]	
*EEE	Support of EEE capability		O	Yes [ ] No [ ]	55.1.3.3
*FR	Support of Fast Retrain capability		O	Yes [ ] No [ ]	55.4.2.5.15

55.12.3 Physical Coding Sublayer (PCS)

Item	Feature	Subclause	Status	Support	Value/Comment
PCT1	PCS Transmit function state diagram	55.3.2.2	M	Yes [ ]	See Figure 55–16
PCT2	PCS Transmit function state diagram with EEE states	55.3.2.2	EEE:M	Yes [ ]	See Figure 55–16 and Figure 55–17
PCT3	PCS Transmit bit ordering	55.3.2.2.4	M	Yes [ ]	See Figure 55–6 and Figure 55–8
PCT4	Invalid control code handling	55.3.2.2.6	M	Yes [ ]	
PCT5	/I/ insertion and deletion	55.3.2.2.9	M	Yes [ ]	
PCT6	/LI/ insertion and deletion	55.3.2.2.10	EEE:M	Yes [ ]	
PCT7	/O/ deletion	55.3.2.2.13	M	Yes [ ]	
PCT8	Scrambler as MASTER	55.3.2.2.16	M	Yes [ ]	
PCT9	Scrambler as SLAVE	55.3.2.2.16	M	Yes [ ]	
PCT10	CRC8	55.3.2.2.17	M	Yes [ ]	See Figure 55–11
PCT11	LDPC encoding	55.3.2.2.18	M	Yes [ ]	Generator matrix is described in Annex 55A
PCT12	DSQ128 mapping	55.3.2.2.19	M	Yes [ ]	
PCT13	EEE Transmit function state diagram	55.3.2.2.22	EEE:M	Yes [ ]	See Figure 55–20
PCT14	LP_IDLE input to scrambler during LPI mode	55.3.2.2.22	EEE:M	Yes [ ]	
PCT15	lpi_tx_mode control	55.3.2.2.22	EEE:M	Yes [ ]	
PCT16	PCS test pattern mode	55.3.3	M	Yes [ ]	See Figure 55–6
PCT17	PMA training - MASTER scrambler	55.3.4	M	Yes [ ]	
PCT18	PMA training - SLAVE scrambler	55.3.4	M	Yes [ ]	

Item	Feature	Subclause	Status	Support	Value/Comment
PCT19	PMA training scrambler reset	55.3.4	M	Yes [ ]	If requested by Link Partner during Auto Negotiation
PCT20	PMA training scrambler initial state	55.3.4	M	Yes [ ]	In no case shall the scrambler state be initialized to all zeros
PCT21	LPI active pair and refresh_active signals	55.3.5.1	EEE:M	Yes [ ]	
PCT22	Alert signaling in place of refresh signaling	55.3.5.3	EEE:M	Yes [ ]	
PCT23	Slave synchronization	55.3.5.1	EEE:M	Yes [ ]	
PCT24	Quiet launch power	55.3.5.2	EEE:M	Yes [ ]	
PCT25	LPI sleep timer	55.3.6.2.3	EEE:M	Yes [ ]	
PCT26	LPI alert timer	55.3.6.2.3	EEE:M	Yes [ ]	
PCT27	LPI wake timer	55.3.6.2.3	EEE:M	Yes [ ]	
PCT28	LPI rx wake timer	55.3.6.2.3	EEE:M	Yes [ ]	
PCT29	LPI tx wake timer	55.3.6.2.3	EEE:M	Yes [ ]	
PCT30	LPI scrambler	55.3.5.3	EEE:M	Yes [ ]	The training sequence without periodic re-initialization described in 55.3.5 shall be used
PCT31	Disable scrambler reinitialization	55.3.5.3	EEE:M	Yes [ ]	
PCT32	Refresh using THP	55.3.5.3	EEE:M	Yes [ ]	
PCT33	Reset THP at the start of refresh	55.3.5.3	EEE:M	Yes [ ]	
PCT34	Master alert on pair A, other pairs silent	55.3.5.3	EEE:M	Yes [ ]	
PCT35	Slave alert on pair C, other pairs silent	55.3.5.3	EEE:M	Yes [ ]	
PCT36	Inactive pairs transmit zeros	55.3.5.3	EEE:M	Yes [ ]	
PCT37	ENCODE function	55.3.6.2.4	M	Yes [ ]	Encode the block as specified in 55.3.2.2.2
PCT38	PCS loopback setup	55.3.7.3	M	Yes [ ]	

55.12.3.1 PCS Receive functions

Item	Feature	Subclause	Status	Support	Value/Comment
PCR1	PCS Receive function state diagram	55.3.2.3	M	Yes [ ]	See Figure 55–18 and state variables as specified in 55.3.6.2.
PCR2	PCS Receive bit ordering	55.3.2.3.1	M	Yes [ ]	See Figure 55–7
PCR3	Descrambling as MASTER	55.3.2.3.2	M	Yes [ ]	
PCR4	Descrambling as SLAVE	55.3.2.3.2	M	Yes [ ]	
PCR5	PMA training descrambler - MASTER	55.3.4.3	M	Yes [ ]	
PCR6	PMA training descrambler - SLAVE	55.3.4.3	M	Yes [ ]	
PCR7	DECODE operation	55.3.6.2.4	M	Yes [ ]	Decode the block as specified in 55.3.4
PCR8	LFER monitor	55.3.6.4	M	Yes [ ]	See state diagrams in Figure 55–15, Figure 55–16 and Figure 55–18.

55.12.3.2 Other PCS functions

Item	Feature	Subclause	Status	Support	Value/Comment
PCO1	PCS Reset function	55.3.2.1	M	Yes [ ]	

55.12.4 Physical Medium Attachment (PMA)

Item	Feature	Subclause	Status	Support	Value/Comment
PMF1	PMA Reset function	55.4.2.1	M	Yes [ ]	
PMF2	PMA transmission	55.4.2.2	M	Yes [ ]	
PMF3	Transmitter clocking	55.4.2.2	M	Yes [ ]	All driven by TX_TCLK
PMF4	PMA Transmit mapping	55.4.2.2	M	Yes [ ]	Per mathematical description given in 55.4.3.1
PMF5	PMA Transmit electrical compliance	55.4.2.2	M	Yes [ ]	See PMA electrical specifications given in 55.5
PMF6	Clocking as MASTER	55.4.2.2	M	Yes [ ]	Source the transmit clock TX_TCLK from a local clock
PMF7	Clocking as SLAVE in loop timed mode	55.4.2.2	M	Yes [ ]	Source the transmit clock TX_TCLK from the recovered clock

Item	Feature	Subclause	Status	Support	Value/Comment
PMF8	Transmit fault mapping	55.4.2.2	O	Yes [ ]	Contribute to the transmit fault bit as specified in 45.2.1.7.4
PMF9	Generates alert signal	55.4.2.2	EEE:M	Yes [ ]	Generates the alert signal defined in 55.4.2.2.1
PMF10	Generates link failure signaling	55.4.2.2	FR:M	Yes [ ] No [ ]	Generates the link failure signal defined in 55.4.2.2.2
PMF11	PMA Receive function performance	55.4.2.4	M	Yes [ ]	LDPC frame error ratio of less than 3.2×10^{-9}
PMF12	Receive fault mapping	55.4.2.4	O	Yes [ ]	Contribute to the receive fault bit as specified in 45.2.1.7.5
PMF13	Implement alert_detect	55.4.2.4	EEE:M	Yes [ ]	Generates alert_detect when the alert signal is detected at the receiver
PMF14	Detect link failure signaling	55.4.2.4	FR:M	Yes [ ] No [ ]	Sets link_fail_detect to true when the link failure signal is detected
PMF15	PHY Control function	55.4.2.5	M	Yes [ ]	See state diagram in Figure 55–28 and state diagrams in Figure 55–29 and Figure 55–30
PMF16	InfoField	55.4.2.5	M	Yes [ ]	All subclauses from 55.4.2.5.1 to 55.4.2.5.13
PMF17	THP initialization	55.4.2.5.14	M	Yes [ ]	Last 16 symbols of PMA_Coeff_Exch
PMF18	PBO exchange	55.4.2.5.14	M	Yes [ ]	
PMF19	Slave PBO setting	55.4.2.5.14	M	Yes [ ]	Slave's PBO final setting within two levels (4 dB) of the MASTER's PBO level
PMF20	THP coefficient exchange	55.4.2.5.14	M	Yes [ ]	
PMF21	Recommended startup sequence timing	55.4.2.5.14	O	Yes [ ] No [ ]	See Table 55–9
PMF22	Implements fast retrain state diagram	55.4.2.5.15	FR:M	Yes [ ] No [ ]	
PMF23	Behavior after fast retrain request	55.4.2.5.15	FR:M	Yes [ ] No [ ]	Transmit PAM2 within 9 LDPC frame periods following link failure request
PMF24	Behavior after fast retrain signal detection	55.4.2.5.15	FR:M	Yes [ ] No [ ]	Transmit PAM2 within 9 LDPC frame periods following link failure signal detection
PMF25	Link Monitor	55.4.2.6	M	Yes [ ]	See state diagram in Figure 55–31
PMF26	Refresh monitor state diagram	55.4.2.7	EEE:M	Yes [ ] No [ ]	Implements state diagram of Figure 55–19
PMF27	Recommended fast retrain sequence timing	55.4.2.5.15	FR:O	Yes [ ] No [ ]	See Table 55–10

Item	Feature	Subclause	Status	Support	Value/Comment
PMF28	Clock Recovery function	55.4.2.8	M	Yes [ ]	
PMF29	MDIX for EEE refreshes and alert	55.4.4	EEE:M	Yes [ ]	
PMF30	Symbol response	55.4.3.1	M	Yes [ ]	Per electrical specifications given in 55.5
PMF31	THP filter coefficient setting	55.4.3.1	M	Yes [ ]	Fixed after startup
PMF32	PMA Transmit power backoff settings	55.4.3.1	M	Yes [ ]	
PMF33	Minimum power backoff requested	55.4.3.1	M	Yes [ ]	Per Table 55–11
PMF34	Automatic configuration	55.4.4	M	Yes [ ]	Comply with the specifications of 40.4.4.1 and 40.4.4.2
PMF35	Pair/Polarity swap detection and correction	55.4.4	M	Yes [ ]	
PMF36	PMA_CONFIG.indicate generation	55.4.5.1	M	Yes [ ]	
PMF37	Maxwait_timer expiration	55.4.5.2	M	Yes [ ]	
PMF38	Minwait_timer expiration	55.4.5.2	M	Yes [ ]	

55.12.5 Management interface

Item	Feature	Subclause	Status	Support	Value/Comment
MF1	Support for Auto-Negotiation	55.6.1	M	Yes [ ]	See Clause 28
MF2	MASTER and SLAVE operation	55.6.1	M	Yes [ ]	Capable of operating either as MASTER or SLAVE
MF3	Extended Next Page support	55.6.1	M	Yes [ ]	
MF4	Optimized FLP timing	55.6.1	M	Yes [ ]	
MF5	Management registers	55.6.1.1	M	Yes [ ]	As defined in Table 55–14
MF6	10GBASE-T Extended Next Page bit assignments	55.6.1.2	M	Yes [ ]	As defined in Table 55–15
MF7	EEE advertisement	55.6.1.2	EEE:M	Yes [ ]	As defined in Table 55–15
MF8	Fast retrain ability advertisement	55.6.1.2	FR:M	Yes [ ]	As defined in Table 55–15
MF9	MASTER-SLAVE resolution with both or neither devices supporting loop timing	55.6.2	M	Yes [ ]	As defined in Table 55–16
MF10	MASTER-SLAVE resolution with one device supporting loop timing	55.6.2	M	Yes [ ]	Device supporting loop timing forced to SLAVE

Item	Feature	Subclause	Status	Support	Value/Comment
MF11	Resolution completion	55.6.2	M	Yes [ ]	Successful completion of resolution treated as MASTER-SLAVE configuration resolution complete.
MF12	Seed counter	55.6.2	M	Yes [ ]	Counter provided to track number of seed attempts
MF13	Counter set to zero at startup	55.6.2	M	Yes [ ]	
MF14	Counter increment	55.6.2	M	Yes [ ]	
MF15	Counter reset	55.6.2	M	Yes [ ]	After resolution is complete
MF16	Bit 7.33.15 set to zero after resolution is complete	55.6.2	M	Yes [ ]	After resolution is complete
MF17	Resolution fault declared	55.6.2	M	Yes [ ]	After generation of seven seeds
MF18	MASTER-SLAVE fault condition	55.6.2	M	Yes [ ]	Condition occurs when both devices manually select MASTER or SLAVE
MF19	MASTER-SLAVE fault condition bit	55.6.2	M	Yes [ ]	Set to one upon fault condition
MF20	MASTER-SLAVE fault resolution	55.6.2	M	Yes [ ]	Fault condition treated as MASTER-SLAVE resolution complete
MF21	MASTER-SLAVE fault condition indication	55.6.2	M	Yes [ ]	link_status_10GigT set to FAIL

55.12.6 PMA Electrical Specifications

Item	Feature	Subclause	Status	Support	Value/Comment
PME1	Electrical isolation	55.5.1	M	Yes [ ]	One of three electrical strength tests listed in 55.5.1
PME2	Insulation breakdown after test	55.5.1	M	Yes [ ]	>2 MΩ, measured at 500 V dc
PME3	Test modes supported	55.5.2	M	Yes [ ]	
PME4	Test mode enablement	55.5.2	M	Yes [ ]	Per management register settings shown in Table 55–12
PME5	The test modes only change the data symbols	55.5.2	M	Yes [ ]	
PME6	Alternate way to enable the test modes	55.5.2	O	Yes [ ]	Mandatory for PHYs without MDIO
PME7	Test mode 1 operation	55.5.2	M	Yes [ ]	
PME8	Test mode 2 operation	55.5.2	M	Yes [ ]	
PME9	Test mode 3 operation, pair D	55.5.2	LT:M	Yes [ ] N/A [ ]	Mandatory for PHY that supports loop timing

Item	Feature	Subclause	Status	Support	Value/Comment
PME10	Test mode 3 operation, pairs A, B and C	55.5.2	M	Yes [ ]	Transmit silence
PME11	Test mode 4 waveform	55.5.2	M	Yes [ ]	Tones per Table 55–13
PME12	Test mode 4 levels	55.5.2	M	Yes [ ]	
PME13	Test mode 5 operation	55.5.2	M	Yes [ ]	
PME14	Test mode 6 operation	55.5.2	M	Yes [ ]	
PME15	Test mode 7 operation	55.5.2	M	Yes [ ]	
PME16	Text fixture 3 isolation	55.5.2	M	Yes [ ]	>30 dB between signals on any of {pairs A, B, C} and pair D
PME17	Transmitter nominal load	55.5.3	M	Yes [ ]	
PME18	AC-coupling to the MDI	55.5.3	M	Yes [ ]	
PME19	Droop test	55.5.3.1	M	Yes [ ]	
PME20	SFDR of transmitter	55.5.3.2	M	Yes [ ]	
PME21	Transmitter jitter as MASTER	55.5.3.3	M	Yes [ ]	
PME22	Transmitter jitter as loop-timed SLAVE	55.5.3.3	LT: M	N/A [ ] Yes [ ]	Applicable only if loop timing is supported
PME23	Transmit power level	55.5.3.4	M	Yes [ ]	
PME24	Transmitter PSD	55.5.3.4	M	Yes [ ]	
PME25	MASTER symbol rate	55.5.3.5	M	Yes [ ]	
PME26	Maximum short term rate of frequency variation during LPI	55.5.3.5	EEE:M	Yes [ ]	Less than 0.1 ppm/s
PME27	Maximum short term rate of frequency variation when switching to and from LPI	55.5.3.5	EEE:M	Yes [ ]	Less than 0.1 ppm/s
PME28	BER over a 55.7 compliant link	55.5.4.1	M	Yes [ ]	
PME29	Receiver frequency tolerance	55.5.4.2	M	Yes [ ]	
PME30	Alien noise tolerance	55.5.4.4	M	Yes [ ]	

55.12.7 Characteristics of the link segment

Item	Feature	Subclause	Status	Support	Value/Comment
LKS1	MDI compatibility	55.7	M	Yes [ ]	
LKS2	Insertion loss of each duplex channel	55.7.2.1	M	Yes [ ]	See Equation (55–11)
LKS3	Return loss	55.7.2.3	M	Yes [ ]	See (55–12)
LKS4	NEXT loss within pairs in a link segment	55.7.2.4.1	M	Yes [ ]	See (55–13)

Item	Feature	Subclause	Status	Support	Value/Comment
LKS5	Power sum NEXT loss	55.7.2.4.2	M	Yes [ ]	See (55–14)
LKS6	Worst pair ELFEXT loss	55.7.2.4.4	M	Yes [ ]	See (55–18)
LKS7	Power sum ELFEXT loss	55.7.2.4.5	M	Yes [ ]	See (55–20)
LKS8	Propagation delay	55.7.2.5	M	Yes [ ]	
LKS9	Delay skew	55.7.2.6	M	Yes [ ]	
LKS10	Delay skew change	55.7.2.6	M	Yes [ ]	
LKS11	Power sum ANEXT loss	55.7.3.1.1	M	Yes [ ]	See (55–23)
LKS12	Average PS ANEXT loss	55.7.3.1.1	M	Yes [ ]	See (55–25)
LKS13	PSANEXT_constant	55.7.3.1.2	M	Yes [ ]	See (55–26)
LKS14	IL(250) used in PSANEXT_ constant equation	55.7.3.1.2	M	Yes [ ]	Actual measured insertion loss at 250 MHz
LKS15	PS AELFEXT loss	55.7.3.2.1	M	Yes [ ]	See (55–32)
LKS16	Average PS AELFEXT loss	55.7.3.2.1	M	Yes [ ]	See (55–34)
LKS17	PSAELFEXT_constant	55.7.3.2.2	M	Yes [ ]	See (55–35)
LKS18	IL(250) used in PSAELFEXT_ constant equation	55.7.3.2.2	M	Yes [ ]	Actual measured insertion loss at 250 MHz
LKS19	Alien crosstalk margin *YL*	55.7.3.3	M	Yes [ ]	Greater than zero

55.12.8 MDI requirements

Item	Feature	Subclause	Status	Support	Value/Comment
MDI1	MDI connector	55.8.1	M	Yes [ ]	8-Way connector per IEC 60603-7:1996
MDI2	Cabling connector	55.8.1	M	Yes [ ]	Plug
MDI3	PHY connector	55.8.1	M	Yes [ ]	Jack (as opposed to plug)
MDI4	MDI connector jack plus plug performance	55.8.2	M	Yes [ ]	
MDI5	Mated MDI connector FEXT loss	55.8.2	M	Yes [ ]	Per (55–53)
MDI6	MDI power down	55.8.2	M	Yes [ ]	
MDI7	MDI return loss	55.8.2.1	M	Yes [ ]	
MDI8	MDI impedance balance	55.8.2.2	M	Yes [ ]	
MDI9	MDI fault tolerance to short circuits	55.8.2.3	M	Yes [ ]	No damage with an indefinite short
MDI10	Recovery from short	55.8.2.3	M	Yes [ ]	
MDI11	Short circuit current	55.8.2.3	M	Yes [ ]	Less than 300 mA

Item	Feature	Subclause	Status	Support	Value/Comment
MDI12	Connection to PSE	55.8.2.3	M	Yes [ ]	No damage to PHY
MDI13	Connection to PSE	55.8.2.3	M	Yes [ ]	No damage to PSE
MDI14	Common-mode impulse tolerance	55.8.2.3	M	Yes [ ]	1000 V common-mode impulse of either polarity

55.12.9 General safety and environmental requirements

Item	Feature	Subclause	Status	Support	Value/Comment
ENV1	Conformance to safety specifications	55.9.1	M	Yes [ ]	IEC 60950-1
ENV2	Installation isolation integrity	55.9.3	INS:M	N/A [ ] Yes [ ]	No electrical contact with unintended conductors or ground
ENV3	Phone voltage immunity	55.9.4	M	Yes [ ]	
ENV4	Electromagnetic compatibility	55.9.5	INS:M	N/A [ ] Yes [ ]	With local and national codes

55.12.10 Timing requirements

Item	Feature	Subclause	Status	Support	Value/Comment
TR1	Delay	55.11	M	Yes [ ]	

Annex 44A

(informative)

Diagram of Data Flow

This annex contains diagrams of the data flow from the MAC to the MDI in the transmit direction and from the MDI to the MAC in the receive direction. The diagrams are provided as a reference, and the corresponding clause should be reviewed for the complete information.

44A.1 10GBASE-R bit ordering

Figure 44A–1 shows the bit ordering on the LAN serial transmit data path.

Figure 44A–2 shows the bit ordering on the LAN serial receive data path.

The 10GBASE-R diagrams reference the following clauses: 46 (XGMII and RS), 47 (XGXS and XAUI), 49 (64B/66B PCS), and 51 (serial PMA).

44A.2 10GBASE-W serial bit ordering

Figure 44A–3 shows the bit ordering on the WAN serial transmit data path.

Figure 44A–4 shows the bit ordering on the WAN serial receive data path.

The 10GBASE-W serial diagrams reference the following clauses: 46 (XGMII and RS), 47 (XGXS and XAUI), 49 (64B/66B PCS), 50 (WIS), and 51 (serial PMA).

44A.3 10GBASE-LX4 bit ordering

Figure 44A–5 shows the bit ordering on the LAN WDM transmit path.

Figure 44A–6 shows the bit ordering on the LAN WDM receive path.

The 10GBASE-LX4 diagrams reference the following clauses: 46 (XGMII and RS) and 48 (8B/10B PCS and PMA).

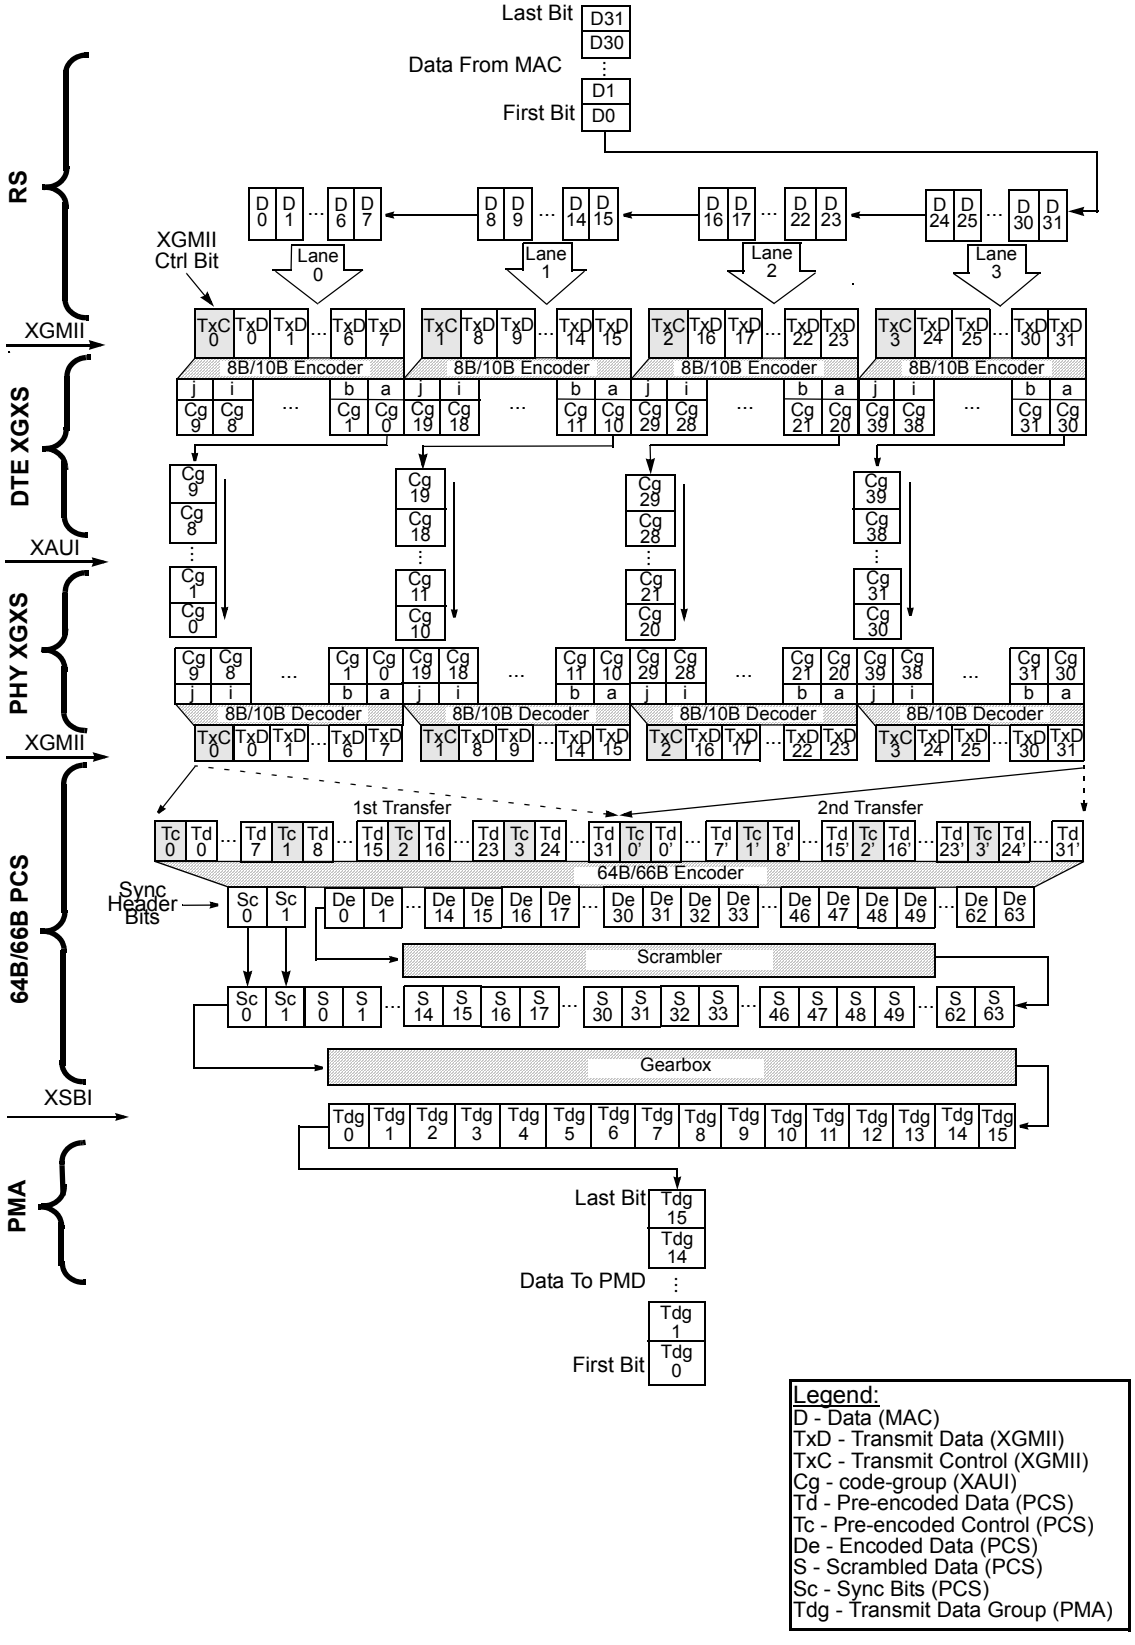

Figure 44A–1—10GBASE-R transmit data path bit ordering

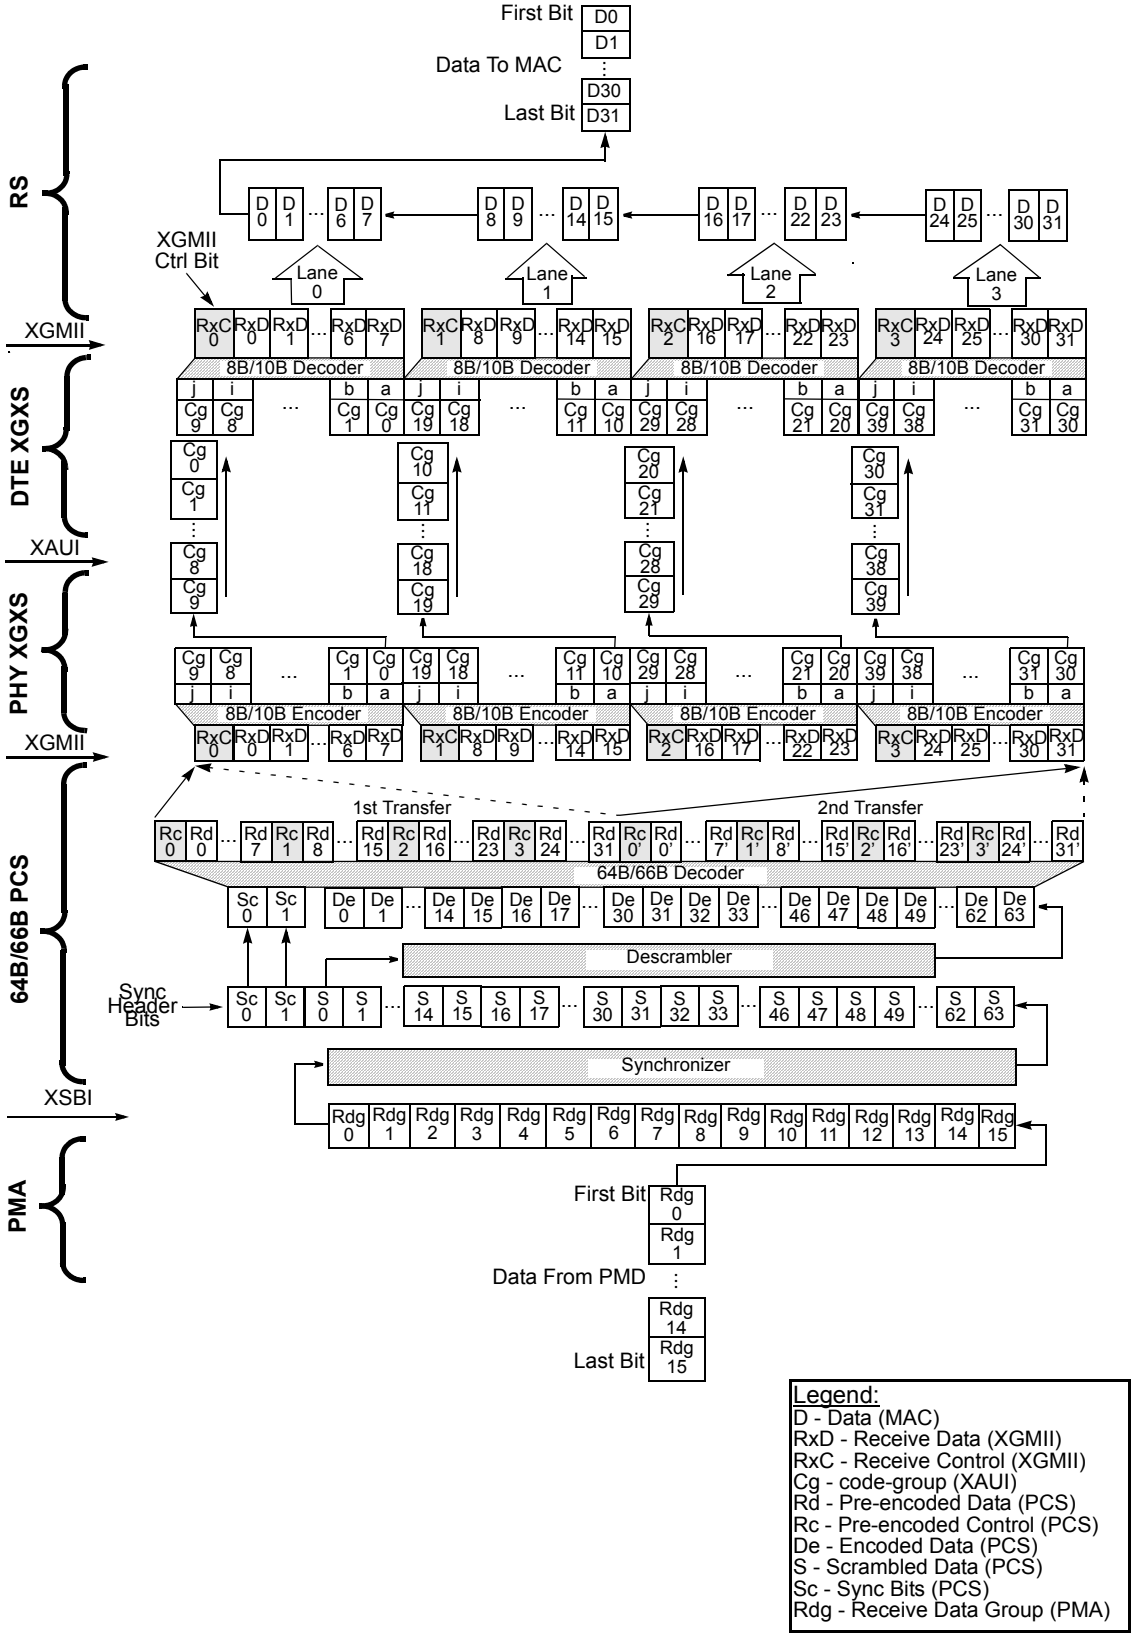

Figure 44A–2—10GBASE-R receive data path bit ordering

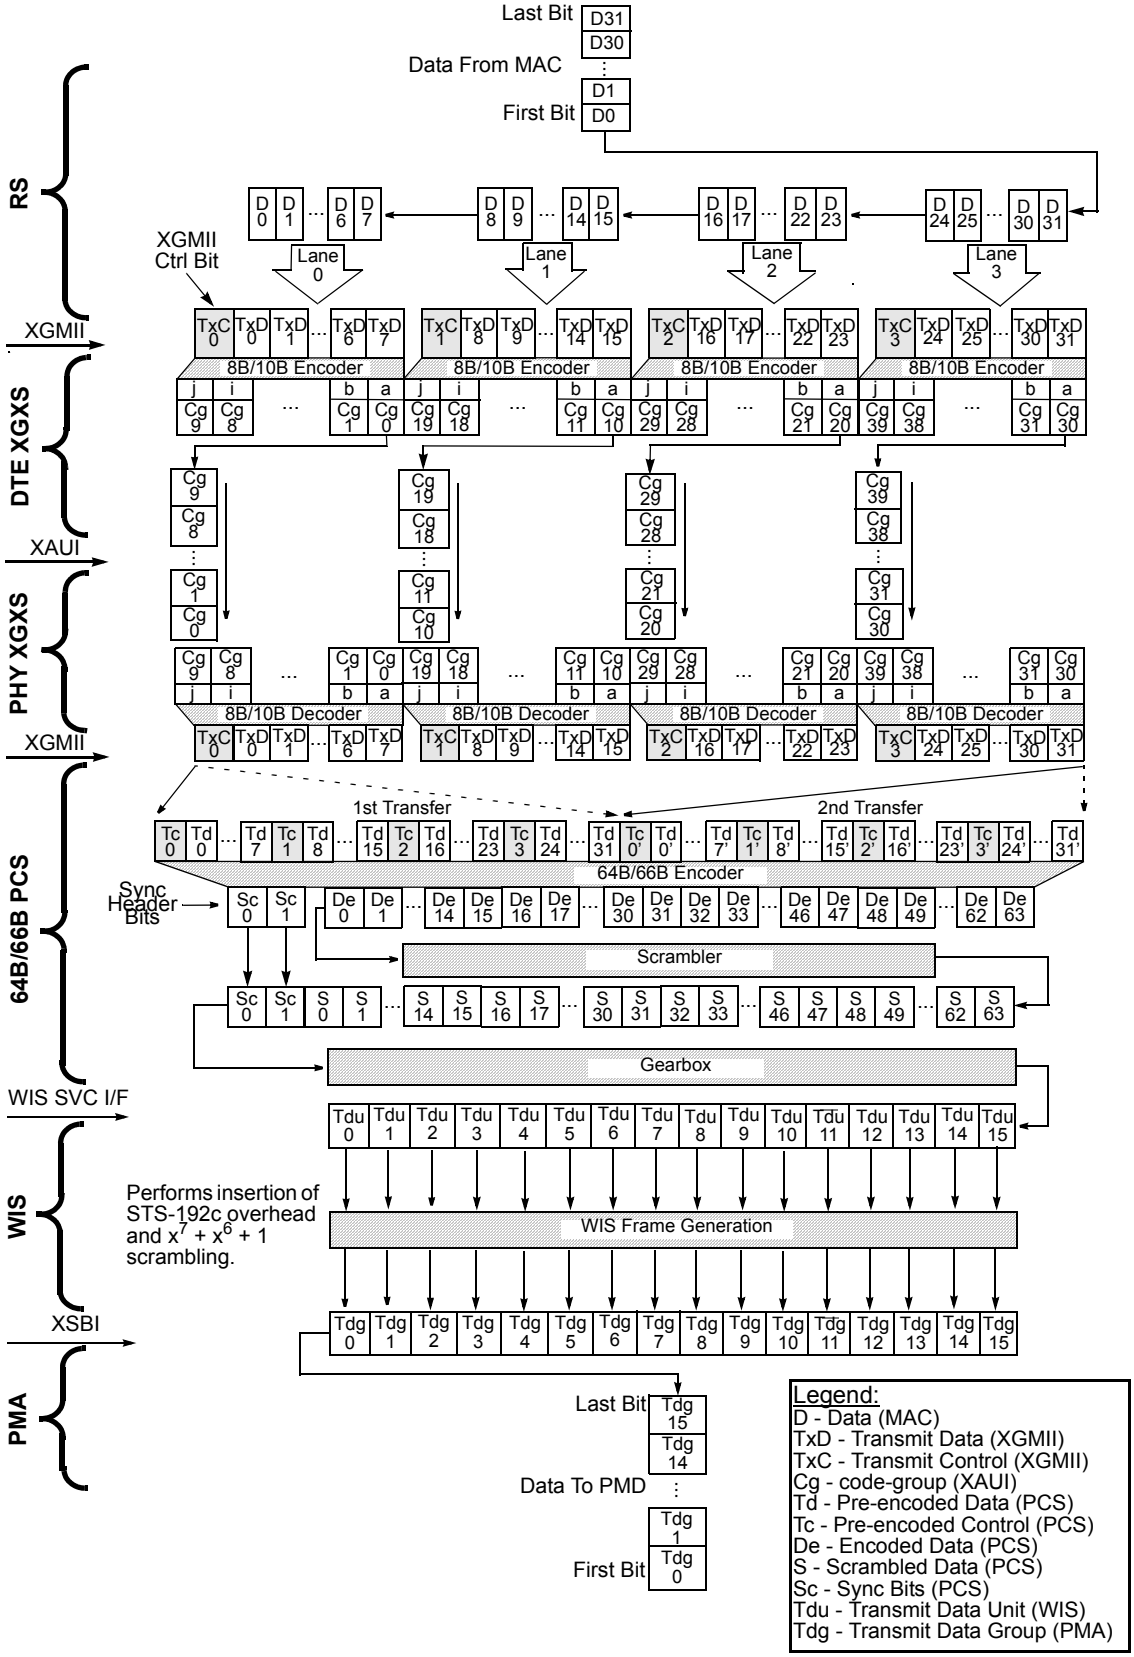

Figure 44A–3—10GBASE-W serial transmit data path bit ordering

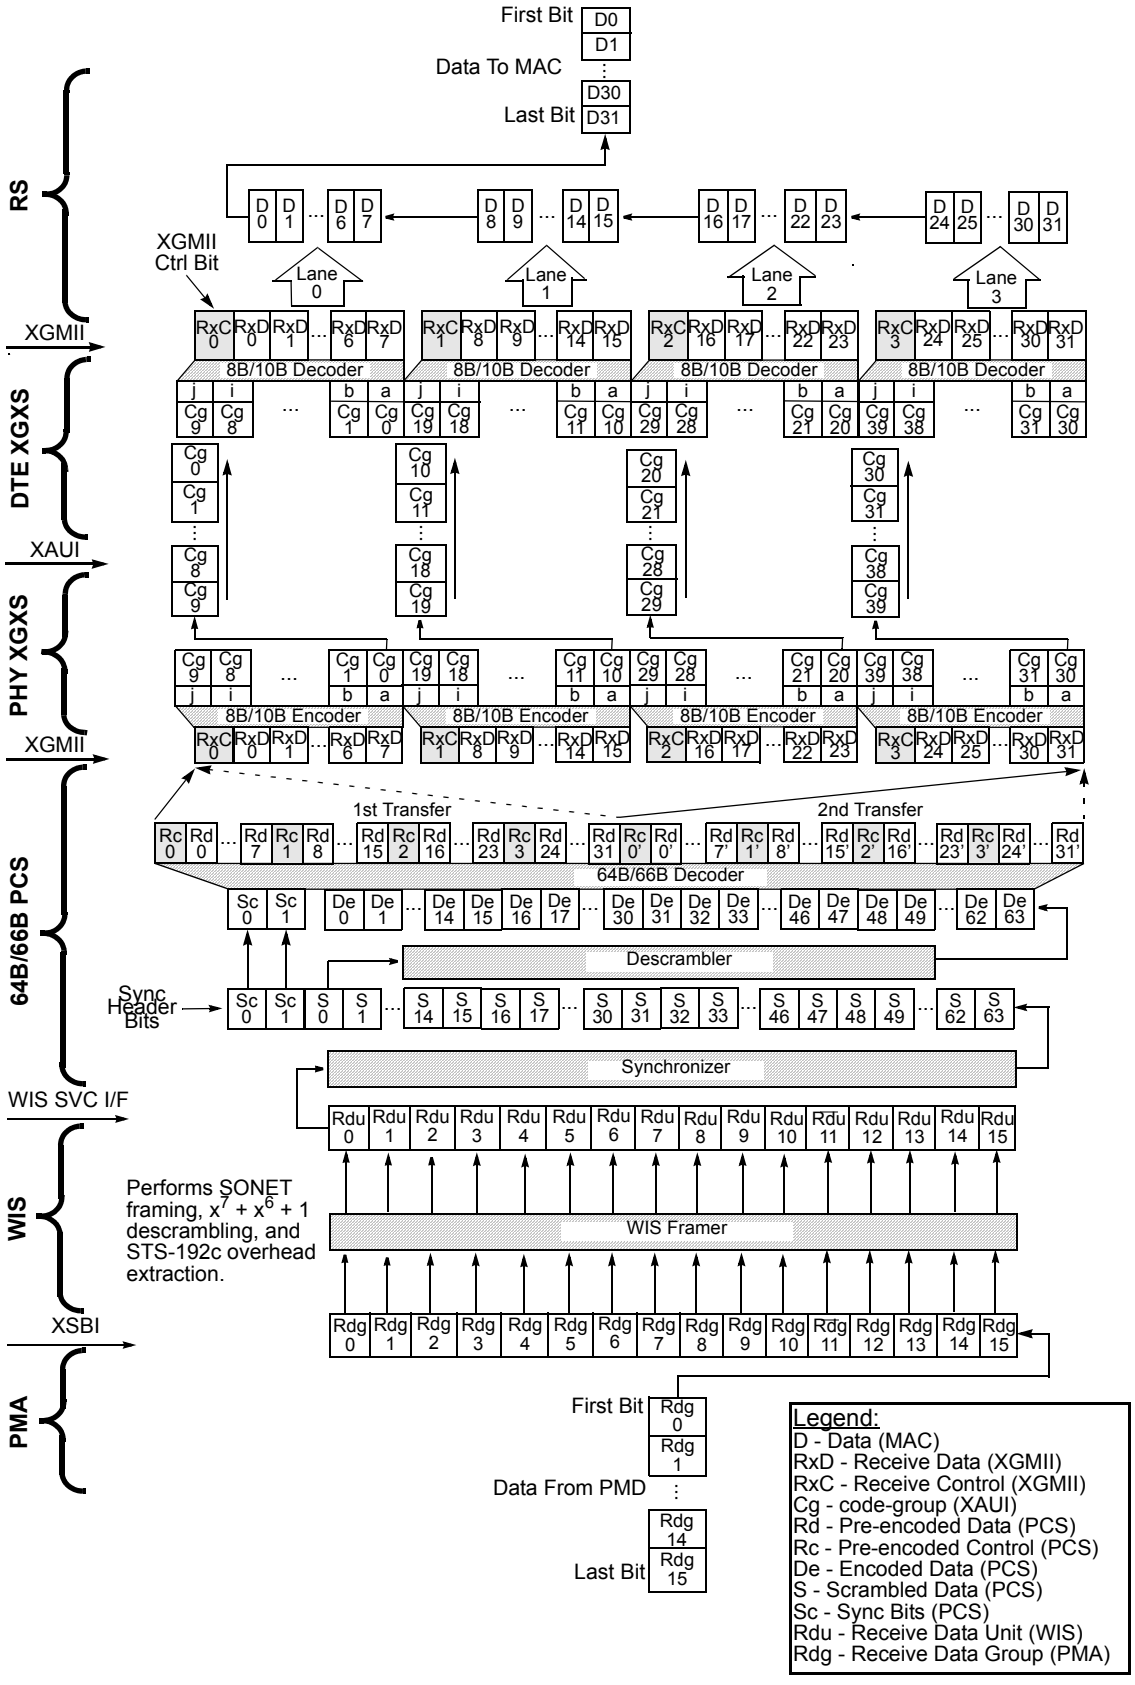

Figure 44A–4—10GBASE-W serial receive data path bit ordering

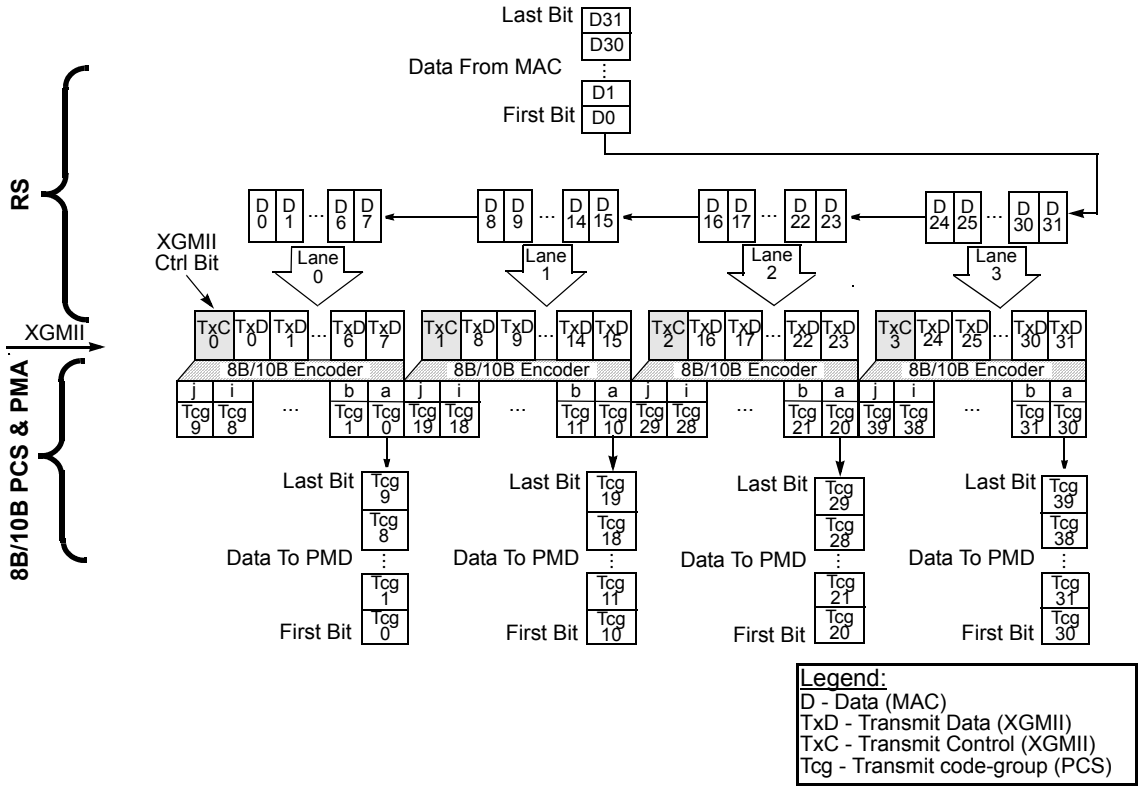

Figure 44A–5—10GBASE-LX4 transmit data path bit ordering

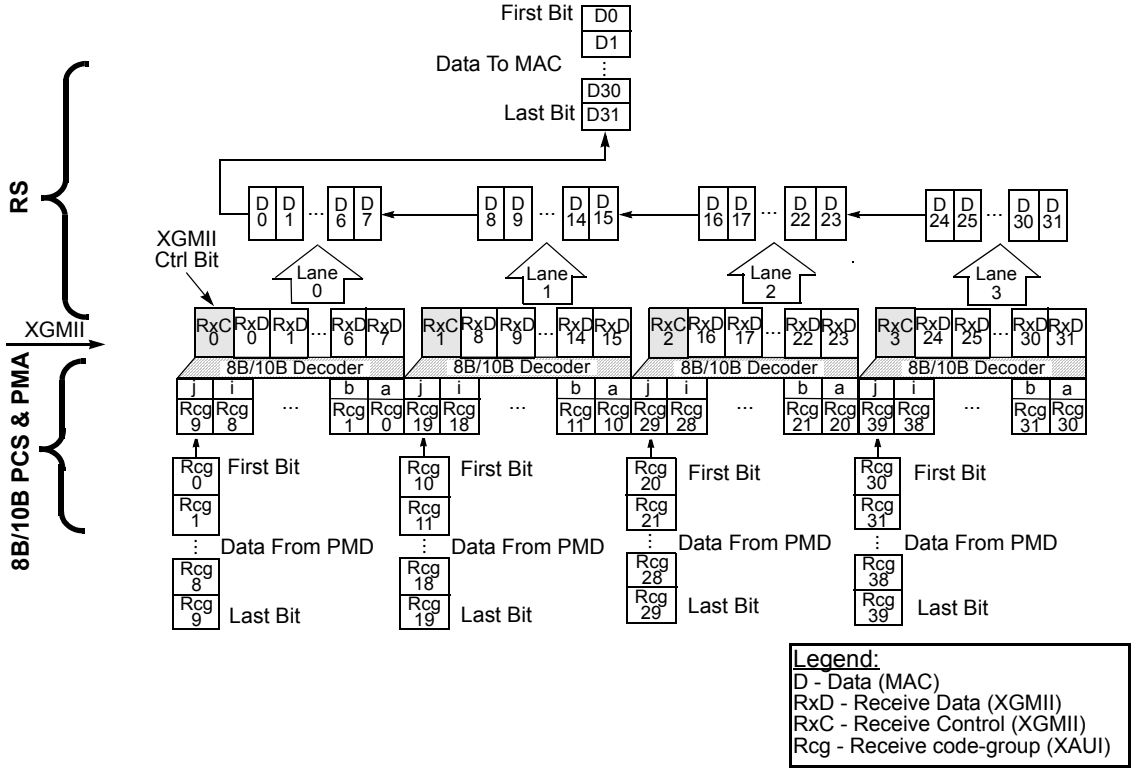

Figure 44A–6—10GBASE-LX4 receive data path bit ordering

44A.4 Loopback locations

Figure 44A–7 diagrams the impact of the various loopback signals on the generation of signal detect and local fault. Clause 46 describes the handling of local fault via the XGMII. If the WIS is not present, the PMA_SIGNAL.indication passes directly from the PMA sublayer to the PCS as is the case in 10GBASE-X and 10GBASE-R devices.

Figure 44A–7 represents an implementation of loopback where loopback is implemented in accordance with the recommendation that it exercises as much of the circuitry in the sublayer as possible. If the loopback is implemented above part of the circuitry in a sublayer, then any data valid signals produced below the loopback point should be ignored during loopback as are any data valid signals from below the sublayer. For instance, if loopback in a WIS implementation was done above the framer, then data valid signals from the framer should be logically OR'ed with loopback.

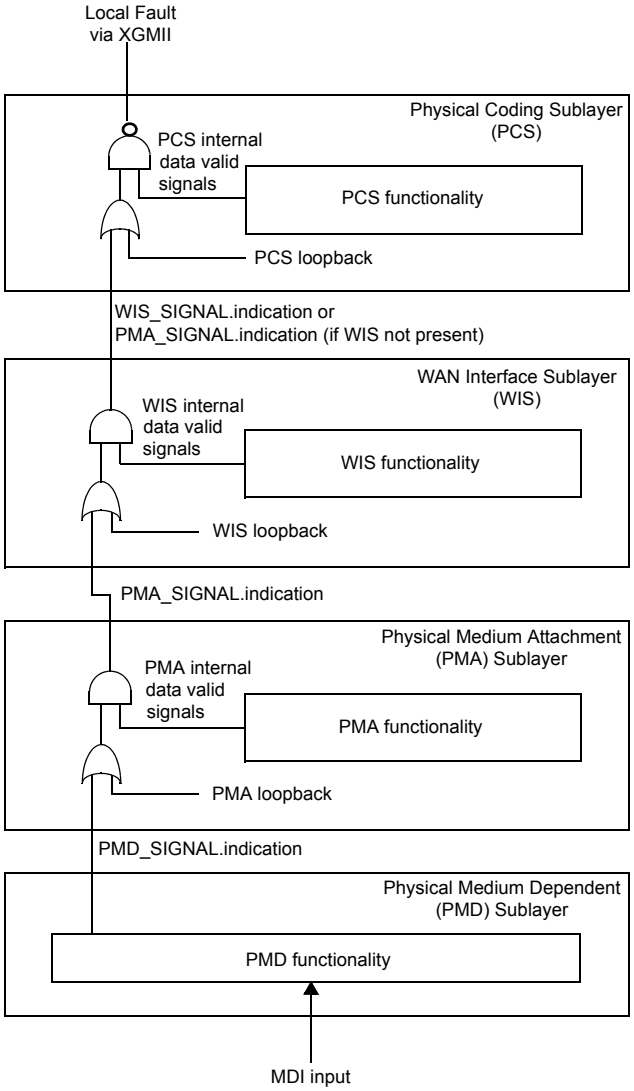

Figure 44A–7—Signal Detect handling across sublayers

Annex 45A

(informative)

Clause 45 MDIO electrical interface

45A.1 MDIO driver

It is possible to implement the MDIO electrical driver using an open drain driver, as the MDIO signal is pulled up to 1.2 V using a resistor. This is shown in Figure 45A–1.

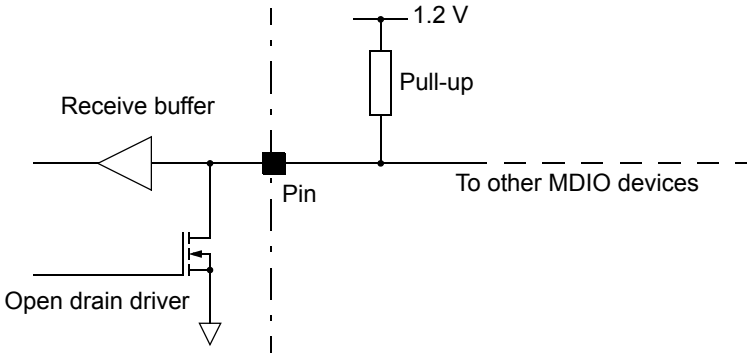

Figure 45A–1—Use of open drain driver for MDIO electrical interface

45A.2 Single Clause 45 electrical interface

As the Clause 45 MDIO electrical interface is significantly different from the Clause 22 MII Management electrical interface, it is not possible to directly attach a Clause 45 MMD to a Clause 22 MII Management Interface. Clause 45 MMDs that include Clause 22 registers may have their Clause 22 registers accessed via the Clause 45 electrical interface. This obviates the requirement to provide a Clause 22 compliant electrical interface in such systems. This is shown in Figure 45A–2.

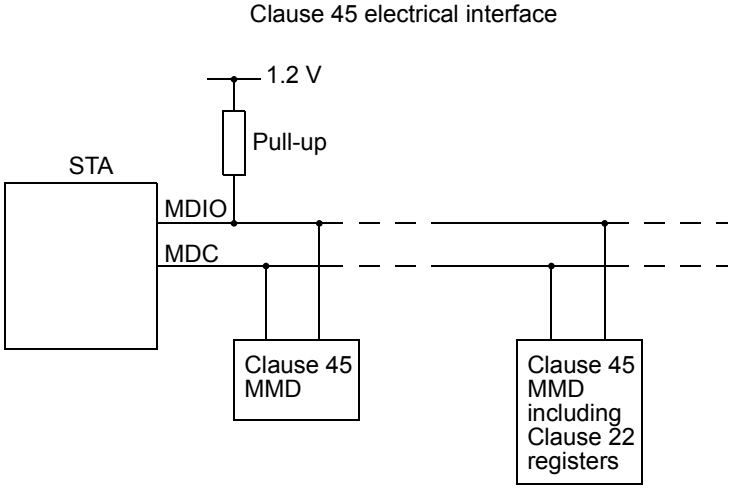

Figure 45A–2—Support of Clause 22 registers via the Clause 45 MDIO interface

45A.3 Clause 45 electrical interface for STA with Clause 22 electrical interface to PHYs

If Clause 22 PHYs are to be attached to a Clause 45 MDIO interface, then a voltage translation device is required. One possible solution is a protocol aware voltage translation device. This arrangement is shown in Figure 45A–3. Such a translation device must correctly drive both the Clause 45 MDIO interface and the Clause 22 MII Management interface using the voltages and timings specified in the appropriate clause. In order to drive the interfaces correctly, the device must have knowledge of the Clause 22 PHY addresses present on the Clause 22 Management interface. When any PHY attached to the Clause 22 MII Management interface is driving the MDIO, the translation device must drive the Clause 45 MDIO interface with the same logical state that is being driven on to the Clause 22 Management interface by the PHY. Whenever the STA drives the Clause 45 MDIO interface, the translation device must drive the Clause 22 Management interface with the same logical state that is being driven on to the Clause 45 MDIO interface. In addition, the translation device must drive the Clause 45 MDC on to the Clause 22 Management interface. Even though the Clause 45 MDIO frames using the ST=00 frame code will also be driven on to the Clause 22 MII Management interface, the Clause 22 PHYs will ignore the frames.

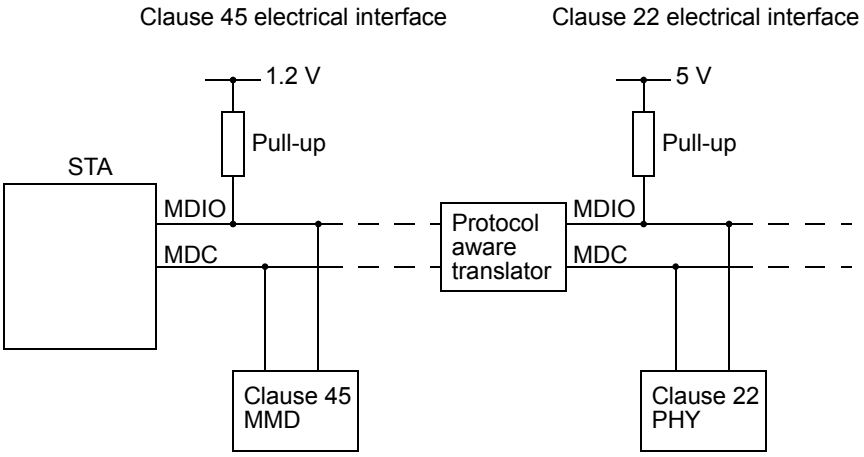

Figure 45A–3—Protocol aware voltage translator between Clause 45 and Clause 22

45A.4 Clause 22 electrical interface for STA with Clause 45 electrical interface to MMDs

If Clause 45 MMDs are to be attached to a Clause 22 MII management interface, then a voltage translation device is required. One possible solution is a protocol aware voltage translation device. This arrangement is shown in Figure 45A–4. Such a translation device must correctly drive both the Clause 22 MII Management interface and the Clause 45 MDIO interface using the voltages and timings specified in the appropriate clause. In order to drive the interfaces correctly, the device must have knowledge of any Clause 22 PHY addresses present within MMDs attached to the Clause 45 MDIO Management interface. When any Clause 45 MMD, including embedded Clause 22 PHYs, attached to the Clause 45 MDIO Management interface is driving the MDIO, the translation device must drive the Clause 22 MII Management interface with the same logical state that is being driven on to the Clause 45 MDIO interface by the PHY. Whenever the STA drives the Clause 22 MII Management interface, the translation device must drive the Clause 45 MDIO interface with the same logical state that is being driven on to the Clause 22 MII Management interface. In addition, the translation device must drive the Clause 22 MII Management interface MDC on to the Clause 45 MDIO interface. The Clause 22 MII Management interface must be capable of generating the MDIO frame structure specified in Clause 45. Even though the Clause 45 MDIO frames using the ST=00

frame code will be driven on to the Clause 22 MII Management interface, the Clause 22 PHYs will ignore such frames.

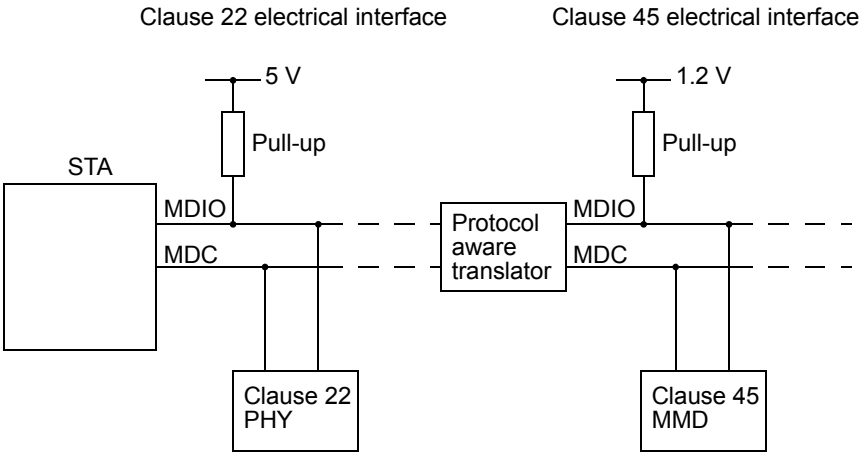

Figure 45A–4—Protocol aware voltage translator between Clause 22 and Clause 45

Annex 48A

(normative)

Jitter test patterns

This annex defines test patterns that allow the 10GBASE-X PHY described in Clause 48 to test either its attached PMD described in Clause 53 or its XAUI interface described in Clause 47 for compliance in a system environment, or for unspecified diagnostic purposes. The patterns may be implemented at a bit, code-group or frame level and may be used for transmitter testing. The receiver may not have the capability to accept these diagnostic sequences; however, system debug can be improved if a receiver is able to test for one or more of these patterns and report bit errors (e.g., 8B/10B decoder errors) back to the user.

48A.1 High-frequency test pattern

The intent of this test pattern is to test random jitter (RJ) at a BER of 10^{-12}, and also to test the asymmetry of transition times. This pattern is not intended for jitter compliance testing. This high-frequency test pattern generates a one, or light on, for a duration of 1 unit interval (UI), followed by a zero, or light off, for a duration of 1 UI. This pattern repeats continuously on each lane for the duration of the test. For example, 10101010101010101010101010101010101010...

NOTE—This pattern can be generated by the repeated transmission of the D21.5 code-group on each lane. Disparity rules are followed.

48A.2 Low-frequency test pattern

The intent of this test pattern is to test low-frequency RJ and also to test PLL tracking error. This pattern is not intended for jitter compliance testing. This low-frequency test pattern generates a one, or light on, for a duration of 5 UI, followed by a zero, or light off, for a duration of 5 UI. This pattern repeats continuously on each lane for the duration of the test. For example, 11111000001111100000111110000011111100000...

NOTE—This pattern can be generated by the repeated transmission of the K28.7 code-group on each lane. Disparity rules are followed.

48A.3 Mixed-frequency test pattern

The intent of this test pattern is to test the combination of RJ and deterministic jitter (DJ) due to high-frequency ISI. This pattern is not intended for jitter compliance testing. This mixed-frequency test pattern generates a one, or light on, for a duration of 5 UI, followed by a zero, or light off, for a duration of 1 UI, followed by a one for 1 UI followed by a zero for 1 UI followed by a one for 2 UI followed by a zero for 5 UI followed by a one for 1 UI followed by a zero for 1 UI followed by a one for 1 UI followed by a zero for 2 UI. This pattern repeats continuously on each lane for the duration of the test. For example, 11111010110000010100111110101100000101 00...

NOTE—This pattern can be generated by the repeated transmission of the K28.5 code-group on each lane. Disparity rules are followed.

48A.4 Continuous random test pattern (CRPAT)

The continuous random test pattern is intended to provide broad spectral content and minimal peaking that can be used for the measurement of jitter at either a component or system level. This pattern is not intended for jitter compliance testing.

NOTE—The basis of this pattern may be found in NCITS TR-25:1999 [B47], “Fibre Channel - Methodologies for Jitter Specification”. This annex uses similar modifications to fit the RPAT test pattern into an IEEE 802.3 frame.

The continuous random test pattern consists of a continuous stream of identical packets, separated by a minimum IPG. Each packet is encapsulated within delimiters as specified in Clause 48 in the ordinary way. The contents of each packet is composed of the following octet sequences, as observed at the XGMII, before 8B/10B coding.

Each packet in the continuous random test pattern consists of 8 octets of PREAMBLE/SFD, followed by 1488 data octets (124 repetitions of the 12-octet modified RPAT sequence), plus 4 CRC octets, followed by a minimum IPG of 12 octets of IDLE.

START/PREAMBLE/SFD:

FB 55 55 55 55 55 55 D5

MODIFIED RPAT SEQUENCE (REPEAT ENTIRE SEQUENCE 31 TIMES CONSECUTIVELY):

BE for 4 octets;

D7 for 4 octets;

23 for 4 octets;

47 for 4 octets;

6B for 4 octets;

8F for 4 octets;

B3 for 4 octets;

14 for 4 octets;

5E for 4 octets;

FB for 4 octets;

35 for 4 octets;

59 for 4 octets.

CRC

F8 79 05 59

IPG

FD 07 07 07 07 07 07 07 07 07 07 07

END

NOTE—The data values for this pattern are in hexadecimal.

48A.5 Continuous jitter test pattern (CJPAT)

The continuous jitter test pattern is intended to expose a receiver's CDR to large instantaneous phase jumps. The pattern alternates repeating low-transition density patterns with repeating high-transition density patterns. The repeating code-group durations should be longer than the time constants in the receiver clock recovery circuit. This assures that the clock phase has followed the systematic pattern jitter and the data sampling circuitry is exposed to large systematic phase jumps. This stresses the timing margins in the received eye. The following pattern is intended for receive jitter compliance testing.

NOTE—The basis of this pattern may be found in NCITS TR-25:1999 [B47]. This annex uses similar modifications to fit the JTPAT test pattern into an IEEE 802.3 frame.

The continuous jitter test pattern consists of a continuous stream of identical packets, separated by a minimum IPG. Each packet is encapsulated within delimiters, as specified in Clause 48, in the ordinary way. The contents of each packet is composed of the following octet sequences, as observed at the XGMII, before 8B/10B coding.

Each packet in the continuous jitter test pattern consists of 8 octets of PREAMBLE/SFD, followed by 1504 data octets, plus 4 CRC octets, followed by a minimum IPG of 12 octets of IDLE.

START/PREAMBLE/SFD:

FB 55 55 55 55 55 55 D5

MODIFIED JTPAT SEQUENCE

0B for 1 octet (lane 0);

7E for 3 octets (lanes 1, 2, 3);

7E for 524 octets—Low-density transition pattern;

F4 for 4 octets—Phase jump;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

AB for 4 octets—Phase jump;

B5 for 160 octets—High-density transition pattern;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

7E for 528 octets—Low-density transition pattern;

F4 for 4 octets—Phase jump;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

AB for 4 octets—Phase jump;

B5 for 160 octets—High-density transition pattern;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump;

EB for 4 octets—Phase jump;

F4 for 4 octets—Phase jump.

CRC

BD 9F 1E AB

IPG

FD 07 07 07 07 07 07 07 07 07 07 07

END

NOTE—The data values for this pattern are in hexadecimal.

48A.5.1 Continuous jitter test pattern (CJPAT) 10 bit values

The following tables depict a 10-bit encoding for CJPAT. This is the encoding which will occur when each lane has negative disparity before the Start column. The actual 10-bit encoding sent when CJPAT is transmitted will be one of 16 encodings depending upon the disparity of each lane at the beginning of the packet. CJPAT has been designed to produce the same disparity on each lane after the Terminate column as that lane had before the Start column. When CJPAT is sent continuously, two of the sixteen possible encodings will occur because Idle will either leave the disparity the same or flip the disparity on all lanes.

Table 48A–1—Start/Preamble/SFD/First 8 Data Octets

Lane0 abcdei fghj		Lane1 abcdei fghj		Lane2 abcdei fghj		Lane3 abcdei fghj	
FB	1101101000	55	1010100101	55	1010100101	55	1010100101
55	1010100101	55	1010100101	55	1010100101	D5	1010100110
0B	1101001011	7E	0111100011	7E	0111100011	7E	0111100011
7E	1000011100	7E	1000011100	7E	1000011100	7E	1000011100

Table 48A–2—Low-density transition pattern (repeat 65 times)

Lane0 abcdei fghj		Lane1 abcdei fghj		Lane2 abcdei fghj		Lane3 abcdei fghj	
7E	0111100011	7E	0111100011	7E	0111100011	7E	0111100011
7E	1000011100	7E	1000011100	7E	1000011100	7E	1000011100

Table 48A–3—Phase jump

Lane0 abcdei fghj		Lane1 abcdei fghj		Lane2 abcdei fghj		Lane3 abcdei fghj	
F4	0010110111	F4	0010110111	F4	0010110111	F4	0010110111
EB	1101001000	EB	1101001000	EB	1101001000	EB	1101001000
F4	0010110111	F4	0010110111	F4	0010110111	F4	0010110111
EB	1101001000	EB	1101001000	EB	1101001000	EB	1101001000
F4	0010110111	F4	0010110111	F4	0010110111	F4	0010110111
EB	1101001000	EB	1101001000	EB	1101001000	EB	1101001000
F4	0010110111	F4	0010110111	F4	0010110111	F4	0010110111
AB	1101001010	AB	1101001010	AB	1101001010	AB	1101001010

Table 48A–4—High-density transition pattern (repeat 20 times)

Lane0 abcdei fghj		Lane1 abcdei fghj		Lane2 abcdei fghj		Lane3 abcdei fghj	
B5	1010101010	B5	1010101010	B5	1010101010	B5	1010101010
B5	1010101010	B5	1010101010	B5	1010101010	B5	1010101010

Table 48A–5—Phase jump (repeat 4 times)

Lane0 abcdei fghj		Lane1 abcdei fghj		Lane2 abcdei fghj		Lane3 abcdei fghj	
EB	1101001000	EB	1101001000	EB	1101001000	EB	1101001000
F4	0010110111	F4	0010110111	F4	0010110111	F4	0010110111

Table 48A–6—Low-density transition pattern (repeat 66 times)

Lane0 abcdei fghj		Lane1 abcdei fghj		Lane2 abcdei fghj		Lane3 abcdei fghj	
7E	1000011100	7E	1000011100	7E	1000011100	7E	1000011100
7E	0111100011	7E	0111100011	7E	0111100011	7E	0111100011

Table 48A–7—Phase jump

Lane0 abcdei fghj		Lane1 abcdei fghj		Lane2 abcdei fghj		Lane3 abcdei fghj	
F4	0010110001	F4	0010110001	F4	0010110001	F4	0010110001
EB	1101001110	EB	1101001110	EB	1101001110	EB	1101001110
F4	0010110001	F4	0010110001	F4	0010110001	F4	0010110001
EB	1101001110	EB	1101001110	EB	1101001110	EB	1101001110
F4	0010110001	F4	0010110001	F4	0010110001	F4	0010110001
EB	1101001110	EB	1101001110	EB	1101001110	EB	1101001110
F4	0010110001	F4	0010110001	F4	0010110001	F4	0010110001
AB	1101001010	AB	1101001010	AB	1101001010	AB	1101001010

Table 48A–8—High-density transition pattern (repeat 20 times)

Lane0 abcdei fghj		Lane1 abcdei fghj		Lane2 abcdei fghj		Lane3 abcdei fghj	
B5	1010101010	B5	1010101010	B5	1010101010	B5	1010101010
B5	1010101010	B5	1010101010	B5	1010101010	B5	1010101010

Table 48A–9—Phase jump (repeat 4 times)

Lane0 abcdei fghj		Lane1 abcdei fghj		Lane2 abcdei fghj		Lane3 abcdei fghj	
EB	1101001110	EB	1101001110	EB	1101001110	EB	1101001110
F4	0010110001	F4	0010110001	F4	0010110001	F4	0010110001

Table 48A–10—CRC/Terminate/IPG

Lane0 abcdei fghj		Lane1 abcdei fghj		Lane2 abcdei fghj		Lane3 abcdei fghj	
BD	1011101010	9F	1010110010	1E	0111100100	AB	1101001010
FD	0100010111	07	0011111010	07	0011111010	07	0011111010
07	1100001100	07	1100001100	07	1100001100	07	1100001100
07	0011110100	07	0011110100	07	0011110100	07	0011110100

Annex 48B

(informative)

Jitter test methods

This annex specifies the definitions and measurement requirements for the jitter specification of the XGXS and XAUI described in Clause 47, the 10GBASE-LX4 PMD described in Clause 53 and the 10GBASE-CX4 PMD described in Clause 54. These measurement methods and specifications are intended to be used for jitter and wander compliance testing, but are not definitive.

While jitter specifications are required to be met for all lanes during operation, the methods described in this annex are written for testing of a single lane. Each lane can be tested with these methods individually or concurrently.

48B.1 BER and jitter model

Measurement of BER within a data eye is not only the fundamental indicator of signal quality, it is a valuable tool to infer jitter properties. Insight into the relationship between jitter, eye opening and error ratio can be gained through mathematical modeling.

48B.1.1 Description of dual Dirac mathematical model

Figure 48B–1 considers a typical eye diagram that may be seen on an oscilloscope. For the purpose of discussion, assume that an ideal trigger source is used so that the eye is accurately depicted. Jitter is indicated by distributed transitions (crossings) of the threshold as the data toggles between logic states. Histograms of transition regions can be taken at the threshold level. The width of the histograms can then be estimated, including standard deviation, etc.

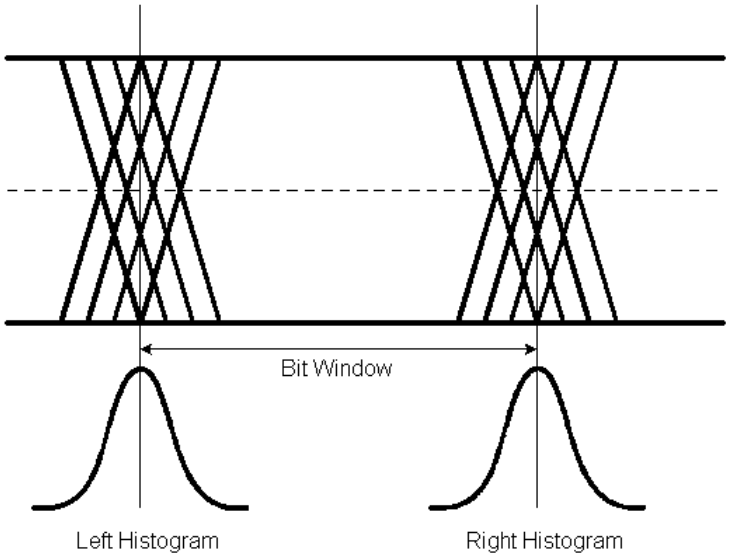

Figure 48B–1—Eye Diagram sketch

The histograms represent probability density functions (PDFs) of the jitter and statistically describe the locations in time of the transitions. The PDFs are placed with their means at the ideal transition times of the

logic states. To simplify matters, the time scale is given in terms of unit intervals (UI) with 0.5 located at the exact center of the eye. The means of the two PDFs are at 0 and +1 UI.

Ideally, the receiver samples the eye at the center where the tails of the transition histograms are small, as shown in Figure 48B–2. To calculate the probability of either transition causing an error due to jitter, the area under its PDF tail on the errored side of the sample point (time) must be calculated. This is the complementary cumulative distribution function, or CDF. For the left hand PDF, the tail is integrated from to the sample point to $+\infty$; the right hand PDF's tail is integrated from $-\infty$ to the sample point. The overall probability of transition error is the sum of the two CDFs. It is assumed that the tails of the neighboring bits do not contribute to the probability of error.

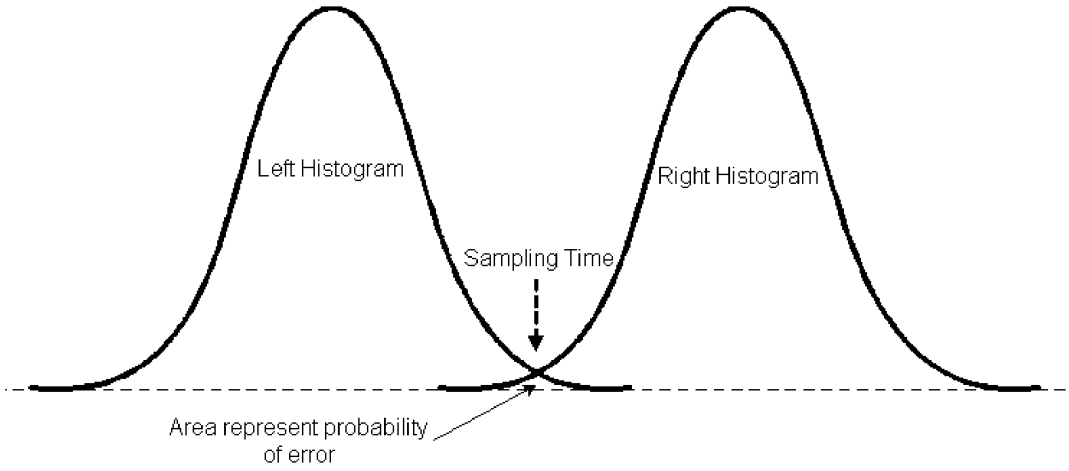

Figure 48B–2—Eye sampling and probability of error

To determine the BER, the probability of a transition-caused error must be multiplied by the probability of a transition occurring. Nominally, the latter may be seen as the average transition density. This model assumes typical data streams have a transition density of 50%.

To demonstrate these concepts, define a general jitter PDF, $JT(\tau, W, \sigma)$, centered at 0, where τ is time, W is the peak-to-peak value of deterministic jitter, and σ is the rms value of random jitter. The left PDF histogram (centered at 0) causes bit errors as:

$$BER_{left}(\tau_{sample}, W, \sigma) = \Gamma_{transistion} \cdot \int_{\infty}^{\tau_{sample}} JT(\tau, W, \sigma)\delta\tau \quad (48B\text{-}1)$$

where τ_{sample} is the sampling instant in time, and $\Gamma_{transistion}$ is the transition density. Similarly, the right PDF histogram (shifted and centered at 1 UI) causes bit errors as

$$BER_{right}(\tau_{sample}, W, \sigma) = \Gamma_{transistion} \cdot \int_{-\infty}^{\tau_{sample}} JT(\tau - UI, W, \sigma)\delta\tau \quad (48B\text{-}2)$$

The sum of these two functions provides the total (BER) due to jitter,

$$BER_{total}(\tau_{sample}, W, \sigma) = BER_{left}(\tau_{sample}, W, \sigma) + BER_{right}(\tau_{sample}, W, \sigma) \quad (48B\text{-}3)$$

In the BERT scan, BER is measured as the sample point, τ_{sample}, is swept between the two eye crossings. The probability of error as a function of the sample point is commonly known as a BER bathtub curve.

48B.1.2 Random Jitter

For random or Gaussian jitter (RJ) only the standard deviation, σ, is necessary to define the PDF

$$RT(\tau, \sigma) = \frac{1}{\sqrt{2\pi}} \cdot \frac{1}{\sigma} \cdot e^{-\left(\frac{\tau^2}{2\sigma^2}\right)} \tag{48B-4}$$

The CDF for this Gaussian PDF function is the complementary error function (erfc).

48B.1.3 Addition of Deterministic Jitter

Total jitter is comprised of both random and deterministic components. The DJ component can be defined, as for RJ, by a PDF, where the combined total jitter PDF is a convolution of the DJ and RJ PDFs.

Simple peak-to-peak is insufficient as the measurement of DJ. An overall weighting function that captures not only the peak-to-peak but also the shape of the density function is required. This may be known as "effective DJ." For purposes of simplicity, effective DJ in this annex is based on the dual Dirac function, where it is assumed that the DJ PDF is comprised only of a pair of delta functions. Other PDFs are certainly possible, however, in the prediction of low BERs, this simplification is sufficient. All references to DJ in Clauses 47 and 53 should be understood as effective DJ.

An example of effective DJ is duty cycle distortion (DCD), which when convolved with RJ results in two Gaussian functions, one centered at each of the two delta functions.

The magnitude of the DJ, is given as peak-to-peak amplitude, *W*. Therefore, each delta function is offset from the mean crossing position by the peak value of DJ, *W*/2. The PDF for deterministic jitter, centered at 0, is therefore defined as:

$$JT(\tau, W, \sigma) = \frac{\delta\left(\tau, -\frac{W}{2}\right)}{2} + \frac{\delta\left(\tau, \frac{W}{2}\right)}{2} \tag{48B-5}$$

When convolved with random jitter, the PDF, centered at 0, becomes:

$$JT(\tau, W, \sigma) = \frac{1}{2\sqrt{2\pi}} \cdot \frac{1}{\sigma} \cdot \left(e^{-\left(\frac{\left(\tau - \frac{W}{2}\right)^2}{2\sigma^2}\right)} + e^{-\left(\frac{\left(\tau + \frac{W}{2}\right)^2}{2\sigma^2}\right)} \right) \tag{48B-6}$$

48B.1.4 Effects of jitter high-pass filtering and CJPAT on deterministic jitter

It is understood that CDRs track low-frequency jitter, and that including this effect in the specifications could ease requirements on clock oscillators (lower cost designs tend to exhibit low-frequency RJ), serializer (SERDES, same advantage) designs and switching power supplies, layouts, bypassing, etc. All jitter specifications include the effects of a high-pass filter (to suppress the significance of low-frequency jitter) to emulate CDR tracking.

It is also realized that, due to frequency content, long complex patterns cause phenomena that are not observed with short patterns—data dependent jitter (DDJ, a form of DJ) can have extreme ranges of frequency content from well below to well above the CDR corner frequencies. Effects are usually seen in both transmitters and receivers. Jitter test patterns are specified in Annex 48A.

48B.2 Jitter tolerance test methodologies

An important measurement in determining link integrity is the characterization of a receiver's (i.e., CDR's) ability to tolerate jittered inputs yet recover error-free data. This is accomplished by inputting a well-controlled, jittered signal to a CDR while measuring the BER at the output of the CDR. As the source signal is modified in amplitude and spectral content, the change in BER is measured. This subclause describes some useful test methodologies for testing a receiver's jitter tolerance.

48B.2.1 Calibration of a signal source using the BERT scan technique

The jitter model described in 48B.1 can be used to calibrate the signal source for tolerance measurements as well as to provide a method for extrapolating lower BER. In this approach, curve fitting of the dual Dirac model onto bathtub curves derived from BERT Scans give the jitter content.

Figure 48B–3 shows an example of how to generate controlled amounts of jitter in a signal to be used for tolerance measurement.[23] Three sources of jitter are provided: Deterministic (DDJ and sinusoidal (SJ)) and Gaussian (RJ). Each of these is added to the serial signal generated by a BERT Pattern Generator. A sinewave generator modulates the clock to the Pattern Generator to provide Sinusoidal jitter modulation from 10 Hz to 20 MHz. A filter or cable adds deterministic jitter. A white noise source[24] with a bandwidth greater than 500 MHz[25] is added to the signal to provide random jitter.

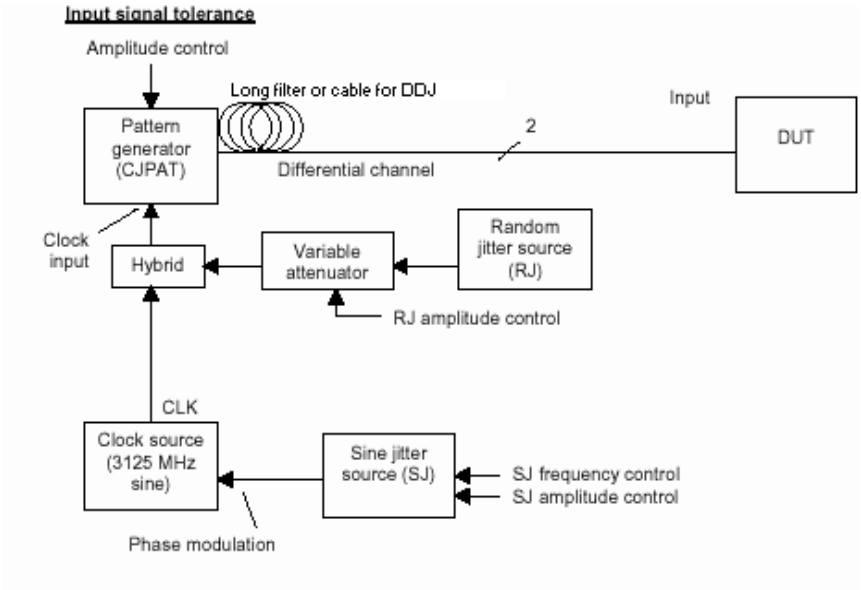

Figure 48B–3—BERT Jitter Tolerance Source

[23]Beware of instrument noise floors and data dependencies, as these may lead to test signals with less stress than intended.

[24]RJ must be Gaussian and not clipped out to >BER being tested (spec is 1E-12). Simply measuring the rms value does not guarantee that the tails exist and that the Rx will be sufficiently stressed as specified.

[25]The spectrum of the portion of RJ that is calibrated must be from greater than 20MHz. The sketch as shown assumes the pattern generator passes clock input jitter through to the data with >>20 MHz BW. This may not be the case. The RJ spectrum need not be white above 20 MHz, but a wider spectrum is preferred.

Amplitude calibration is not possible on a scope due to the low probability of the peak-to-peak RJ effects. For better visibility for amplitude calibration, the RJ portion may be temporarily replaced with an equivalent amount of peak-to-peak SJ to bring TJ up to the same specified output jitter level, allowing the desired amplitude value to be met.[26,27]

If the clock data recovery circuit of the system under test has a corner frequency above the maximum sine wave generator frequency, be sure to increase the amount of the applied deterministic jitter by the amount of the specified high-frequency value sinusoidal jitter.

When calibrating the jitter tolerance test setup, the effects of high pass filtering in the time/jitter domain must be understood and included per the specifications. Per 48B.1.4, high pass filtering can have profound effects on pattern dependent jitter (it will increase measured DJ with patterns such as CJPAT) and also track out jitter below the specified corner frequency. [28]

After complete, calibration must be verified using the methods described in 48B.3 and iterated as needed until the required specifications are met.

It should be noted that the XAUI is a differential input and calibration of the test signal must be performed using the differential signal. If test equipment only provides single-ended possibilities, care should be taken to calibrate and test with a balun to convert differential signals to single-ended signals. Single-ended signals should not be directly used, due to the possibility of asymmetric characteristics.

48B.3 Jitter output test methodologies

Three viable methodologies for measuring jitter output are described in this subclause. They are as follows:

a) Time domain measurement using an oscilloscope to characterize the data eye.
b) Time domain measurement using BERT scan by moving of the data sampling point within the data eye.
c) Time interval analysis based on accurate measurement of the time interval between threshold crossings of the transmitter waveform;

48B.3.1 Time domain measurement—Scope and BERT scan

It is easy to grasp that if the eye is larger than the eye mask, it must be ok. This is not necessarily true. Given the probabilistic nature of random jitter, it is necessary either to test for an extended amount of time to reach a high confidence level for achieving 1×10^{-12} BER or to perform some kind of statistical extrapolation.

48B.3.1.1 Jitter high pass filtering (using Golden PLL)

Due to the issues mentioned in 48B.1.4, when testing either at the component level or the system level, the test clock should be derived from the data with a given high-pass filter function of the jitter using, for example, a golden PLL.

Given a typical test setup, the serial bit pattern is transmitted to the input of a BERT and a golden PLL. The golden PLL extracts a reference clock to trigger the sampling scope or clock the BERT.

[26]Amplitude calibration should be done before addition of SJ, as SJ is designed to add margin into the design. NOTE—This footnote is not referring to the suggestion of temporarily substituting SJ and RJ during amplitude calibration, but to the SJ that is required as part of the specifications.

[27]It may be required to add amplification to achieve the required test signal amplitude.

[28]Beware of autocorrelation effects.

48B.3.1.2 Time domain scope measurement

Time domain measurement uses the high-speed sampling scopes to view the jitter output data eye. Most high-speed sampling scopes today provide features to collect and present data on the output jitter. Some oscilloscopes provide a feature to compare the measured data to an "eye mask." An eye mask is a specification for allowed eye opening. The advantage of an eye mask is that it tests for amplitude as well as timing compliance. The general physical media transmitter pulse shape characteristics are specified in the form of a mask of the transmitter eye diagram at any of the compliance measurement points. These characteristics include rise time, fall time, pulse overshoot, pulse undershoot, and ringing, all of which prevent excessive degradation of the receiver sensitivity. For the purpose of an assessment of the transmit signal, it is important to not only consider the eye opening, but also the overshoot and undershoot limitations. The parameters specifying the mask of the transmitter diagram (eye mask) can be found in the clause of the applicable Physical Layer specification whether it be copper or fiber physical media. The eye mask, through its use of a specified time range in which the transmit signal can change state from the logic low to logic high levels, is also specifying a measure of the allowed jitter.

It is strongly advised not to use the eye mask to verify that total jitter is within specification because of the nature of the test. The eye mask test is generally a short test, and thus the eye diagram is not captured for a sufficient time to capture the full extent of the random jitter's peak-to-peak value. It is necessary to capture the eye pattern for a sufficient period to insure the full extent of the deterministic jitter is captured by the test instrument. Since the peak-to-peak of the random jitter for a BER of 1×10^{-12} is 14 times the sigma, it is not possible to capture the full extent of random jitter's peak-to-peak value due to the low sampling rate.

48B.3.1.3 BERT Scan

As implied by the jitter model in 48B.1, a BERT scan approach can provide random and deterministic jitter components and provides a mechanism to extrapolate to lower BER (less than 1×10^{-12}) without impossibly long test times due to the inherent high sampling rate.

48B.3.1.3.1 Approximate curve-fitting for BERT scan

48B.3.1.3 describes a technique for using a BERT scan to determine eye opening and jitter. For highest accuracy, the bathtub curve should be measured over a high number points at low BER and curve-fitted with a least-squares method to estimate equivalent DJ, RJ, and TJ values. However, a simple and fast method for estimating these values may be applied using only two measurement points.

The following steps describe a process for estimating equivalent RJ, DJ, and TJ values from a two point BERT scan measurement:

a) Measure the eye opening at 2 different BER, BER_0 @τ_0 and BER_1 @τ_1 (e.g., 10^{-9} and 10^{-5}).

b) For each BER_n, determine the associated Q_n from the inverse normal cumulative probability distribution, adjusted for transition density, e.g., $Q = 3.94$ for BER = 10^{-5}, and $Q = 5.77$ for BER = 10^{-9}, where transition density is assumed to be 0.5.

c) Calculate the individual jitter components using the equations for Random Jitter and Deterministic Jitter, and then calculate the total jitter for the given BER = 10^{-12}, using the following three equations:

Random Jitter—

$$RJ_{RMS} = 0.5\left|\frac{t_1 - t_0}{Q_1 - Q_0}\right| \quad \text{(48B-7)}$$

Deterministic Jitter—

$$DJ = UI - t_0 - (2Q_0 RJ_{RMS}) \quad \text{(48B-8)}$$

Total Jitter—

$$TJ = DJ + 13.8 \cdot RJ_{RMS} \quad \text{(48B-9)}$$

The minimum value for measured BER is constrained by test time (10 errors are suggested as an absolute minimum to get reasonable statistical confidence); the maximum value is constrained by potential departure of actual results from the assumed curve fit shape (BER = 10^{-5} should be the maximum value used).

48B.3.2 Time Interval Analysis

Time Interval Analysis (TIA) uses many accurate, single-shot, edge-to-edge time measurements. The statistics of these measurements can be used to perform jitter calculations. True random jitter (RJ) and true deterministic jitter (DJ) can be measured using this technique on repeated patterns as well as random data streams. From the true RJ and DJ measurements, an error probability density function, or bathtub curve can be generated. Using the bathtub curve, effective DJ and effective RJ can be calculated. There are two measurement methods defined for executing this test:

a) TIA with Golden PLL
b) TIA armed on pattern trigger

Each method uses the basic TIA methodology, but with different schemes for acquiring the data. In the first method, the TIA is used with a Golden PLL, as described in 48B.3.1.1, and time measurements are made from the recovered reference clock to the data transitions. The low-frequency content that would be tracked by the receiver is presented on the reference clock, and thus eliminated from the jitter data before being presented to the TIA. The second method is used to acquire data when a Golden PLL is not available, or, when jitter on just the data is desired. This method requires a pattern trigger to synchronize the jitter measurements relative to a specific location in the pattern. The TIA can extract the pattern trigger from the data based on a unique bit sequence or pattern length.

48B.3.2.1 TIA with Golden PLL

In this option, a Golden PLL is required to provide a reference clock derived from the data stream to trigger the TIA. This scheme effectively removes low-frequency content from the measured data in a manner similar to that done in a receiver. The configuration used for this option is shown in Figure 48B–4 (note that only a single channel is shown). The Golden PLL in Figure 48B–4 is shown with a differential input; other approaches are possible, but it is important that the balance of the data signals is not disturbed and that both phases are included in clock recovery. A Golden PLL is used for each data signal so as to ensure proper CDR modeling on each data signal. Each Golden PLL is intended to track low-frequency jitter content occurring on the data signal applied to the input stage of that PLL. It cannot be assumed that the clock extracted from one data signal can be used for the other signals since the low-frequency jitter content may be different. Some TIA models have this algorithm integrated internally in the instrument, and some TIA models have integrated Golden PLL or high-pass filtering algorithms.

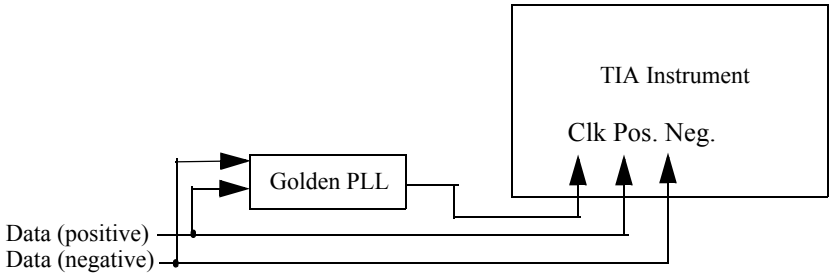

Figure 48B–4—TIA with Golden PLL configured for jitter measurement

48B.3.2.1.1 Test method

The TIA is capable of measuring the time from the data transitions to the reference clock as extracted with the Golden PLL. This results in a clock to data histogram that describes the jitter occurring on the data transitions relative to the reference clock (as is the case with a bit clock triggered BERT measurement). RJ can be estimated by fitting a Gaussian curve to the tail region of the histogram. DJ can be estimated by measuring the separation of the means of the fitted Gaussian curves. Some sampling TIA instruments have built in algorithms specifically designed for this process, and can automatically estimate RJ, DJ, and TJ in this configuration. Figure 48B–5 shows a typical clock to data histogram for a communication signal. Note the highlighted curves that represent the fitted Gaussian curves. Also note that the rising edges are treated separately from the falling edges so as to accurately group the jitter contributors. This is critical in cases where the cross point of the data signal is not at the measured voltage. In Figure 48B–5, the rising edge histogram is shifted to the right of the falling edge histogram, indicating that the cross point was below the test voltage. The outer most histogram on a given side is used for the RJ and DJ calculations for that side of the histogram. Each side is treated independently.

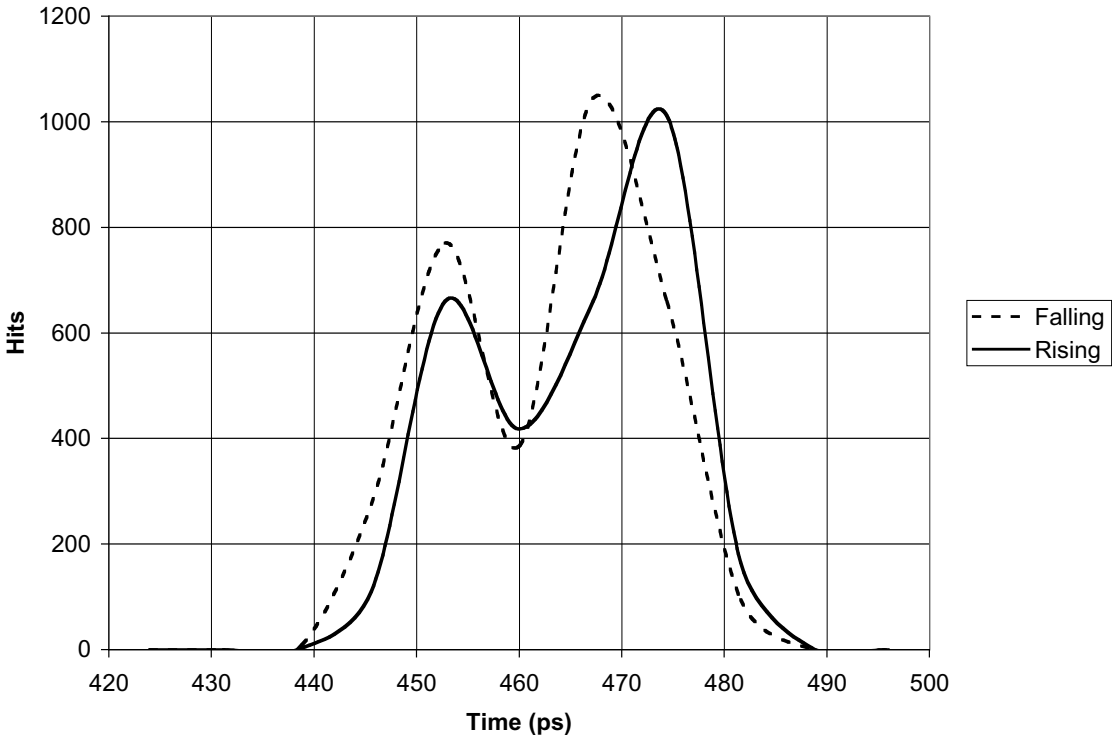

Figure 48B–5—Typical clock to data histogram with fitted Gaussian curves

Jitter values such as DJ, RJ, and TJ can also be displayed on the plot corresponding to 1E-12 BER as shown in Figure 48B–5. Effective Jitter is calculated from the bathtub curve using the same calculations used by the BERT scan method described in 48B.3.1.3.1. To quantify the quality of the measurement, a goodness of fit calculation can be made which compares the quality of the fit of the Gaussian to the actual histogram. Acceptable goodness-of-fit values (χ) for probability confidence level (approximately 1E-4) should be in the range of $0.8 < \chi < 1.2$. The measured DJ and TJ values using this method can be compared against the maximum jitter as specified.

48B.3.2.2 TIA with pattern trigger

The basic setup for this option is shown in Figure 48B–6 (note that only one channel is shown). The pattern trigger can be generated internally to the TIA based on the identification of a unique bit sequence within the data stream or pattern length of the repeated pattern. No measurements are made from the pattern trigger. All measurements are made within the data signal thereby eliminating any error attributed to the bit clock. The high-pass filter characteristic of the Golden PLL is applied using a DSP algorithm thereby eliminating the need for external hardware. Some TIA instruments have this algorithm fully integrated.

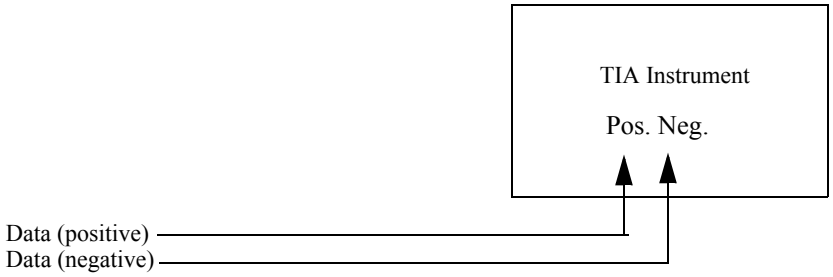

Figure 48B–6—TIA with pattern trigger configured for jitter measurement

48B.3.2.2.1 Test Method

All time measurement are made from the same reference edge within the data pattern. This edge is identified by the pattern trigger. Thousands of measurements are made of each transition within the data pattern relative to the reference edge. Histograms for each transition are captured from this data. An FFT of the autocorrelation algorithm of the variance information from these histograms is used to extract SJ. Using the tail fit algorithm described in 48B.3.2.1.1, RJ is measured from each of these histograms and the results plotted in the FFT display. The RJ, bounded by the high-pass filter function is calculated from the FFT of the autocorrelation function. Data dependent jitter is measured from the histograms of each transition based on the displacement of each histogram mean relative to integer multiples of the measured bit period. TJ is calculated and the bathtub curve generated based on the convolution of the RJ PDF, DDJ PDF, and the SJ PDF. This method directly measures the DJ PDF and thereby does not need to assume that the DJ PDF is a dual Dirac function. Effective DJ and Effective RJ is calculated from the bathtub curve using the same formulas used for the BERT scan method, thereby ensuring correlation. Some TIA instruments have this algorithm fully integrated, simplifying the measurement methodology.

48B.3.2.3 Approximate curve fitting for TIA bathtub curve

Since the TIA is capable of measuring true DJ and its probability density function (PDF), the TIA can accurately display total jitter (TJ) as a function of BER in the form of an error probability density function, or bathtub curve. From the Error PDF, the effective DJ can be estimated using the same curve fit technique described in 48B.3.1.3.1.

Annex 50A

(informative)

Thresholds for Severely Errored Second calculations

This annex provides tables of threshold values that may be used to determine the generation of Section, Line and Path Severely Errored Second (SES) events for various error ratios. An SES event is deemed to have occurred when the number of bit errors, detected via the corresponding Bit Interleaved Parity (BIP) check, exceeds some predefined threshold value when accumulated over a 1 s interval; reporting of an SES event indicates that the actual bit error ratio on the medium has increased to a point where Station Management must be notified. These threshold values are referred to as *x* in the Layer Management definitions in Clause 30. This annex also defines the terms "Path Block Error" and "Far End Path Block Error."

The values in these tables are derived according to the algorithms given in Annex C of ANSI T1.231-1997.

50A.1 Section SES threshold

The Section SES threshold is referenced as *x* in 30.8.1.1.3.

Table 50A–1 gives the values to be used for the Section SES threshold in order to report Section SES events at various limiting bit error ratios. A BER of 10^{-6} corresponds to the default Section SES threshold of 8554 in 30.8.1.1.3.

Table 50A–1—Section SES

BER	Threshold, *x*
10^{-6}	8554
10^{-7}	980
10^{-8}	99
10^{-9}	10
10^{-10}	1

50A.2 Line SES threshold

The Line SES threshold is referenced as *x* in 30.8.1.1.11.

Table 50A–2 gives the values to be used for the Line SES threshold in order to report Line SES events at various limiting bit error ratios. A BER of 10^{-6} corresponds to the default Line SES threshold of 9835 in 30.8.1.1.11.

50A.3 Path SES threshold

The Path SES threshold is referenced as *x* in 30.8.1.1.19.

Table 50A–2—Line SES

BER	Threshold, *x*
10^{-6}	9835
10^{-7}	984
10^{-8}	98
10^{-9}	10
10^{-10}	1

Table 50A–3 gives the values to be used for the Path SES threshold in order to report Path SES events at various limiting bit error ratios. The default Path SES threshold of 2400 in 30.8.1.1.19 corresponds to a situation where 30% of all the SPEs being received have parity errors detected via the Path BIP checks.

Table 50A–3—Path SES

Fraction of SPEs with parity errors	Threshold, *x*
30%	2400
25%	2000
20%	1600
15%	1200
10%	800

50A.4 Definition of Path Block Error

A Path Block Error (see Section 8.1.1.1.2 of ANSI T1.231-1997) is declared when one or more of the 8 BIP-8 parity bits that are computed over the received Synchronous Payload Envelope, and compared to the B3 octet of the Path Overhead, are found to be in error. At most one Path Block Error can be declared per WIS frame (i.e., a maximum rate of 8000 counts per second), regardless of the actual number of parity errors detected using a particular B3 octet.

50A.5 Definition of Far End Path Block Error

A Far End Path Block Error is declared when the G1 octet extracted from a received WIS frame indicates that the far end WIS has detected one or more BIP-8 parity errors in the SPE of a previously transmitted WIS frame. The far end WIS reports these parity errors via the REI-P field of the G1 octet. At most one Far End Path Block Error can be declared per received WIS frame (i.e., a maximum rate of 8000 counts per second), regardless of the actual number of parity errors reported in the REI-P field of a particular G1 octet.

Annex 55A

(normative)

LDPC details

55A.1 The generator matrix

The generator matrix, G, uniquely specifies the mapping of information bits to codewords. G.txt contains a representation of G. G.txt contains 1723 rows, one for each row of G. Each row has numbers ranging from 0 to 2047 separated by spaces. Each number represents the column index of the "1" entries in the specific row. All other entries of G are "0" and are not referenced in G.txt. G.txt is available online in the file matrices.zip.[29]

55A.2 The sparse parity check matrix H

Codewords that are generated using G are in the null space of a sparse parity check matrix H. The H matrix is available in machine readable format in file H.txt. This H matrix was obtained by permuting the columns of another sparse matrix using the permutations specified in col_swap.txt. The file row_swap.txt describes row permutations of H that can put it in pseudo-lower triangular form. H.txt, col_swap.txt and row_swap.txt are informative and are also available online in the file matrices.zip.[29]

For further information on LDPC codes, see Gallagher [B32].

[29]matrices.zip is available at: http://standards.ieee.org/downloads/802.3/.

Annex 55B

(informative)

Additional cabling design guidelines for 10GBASE-T

This annex provides additional cabling guidelines for 10GBASE-T deployment on balanced copper cabling systems as specified in 55.7. These guidelines are intended to supplement those in Clause 55.

55B.1 Alien crosstalk considerations

The noise coupled between the adjacent link segments illustrated in Figure 55B–3 and Figure 55B–4 is referred to as alien crosstalk (Figure 55B–1).

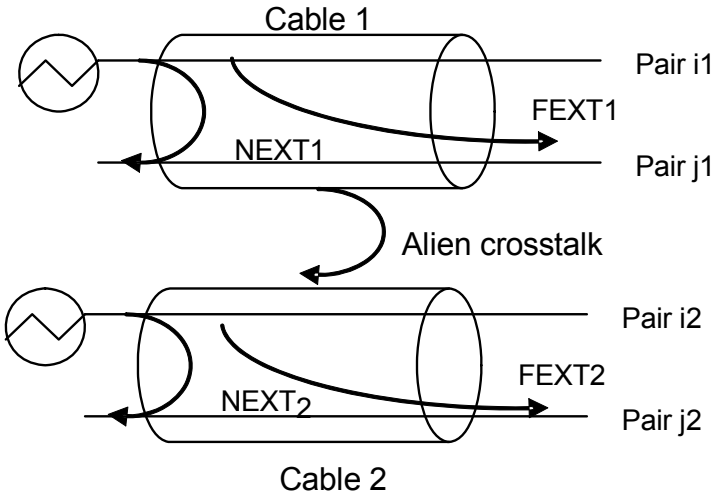

Figure 55B–1—Alien crosstalk coupling between link segments

The alien crosstalk noise levels are dependent on the number and proximity of adjacent cables and connectors. The alien crosstalk link segment specifications of PS ANEXT and PS AELFEXT given in 55.7 are based on the alien crosstalk measurements between cabling channels configured in a six-around-one orientation where the disturbed cable is the central cable and adjacent to all of the other disturbing cables. The alien crosstalk coupling between a cable and six adjacent cables tightly bound in a six-around-one configuration exhibits a "worse case".

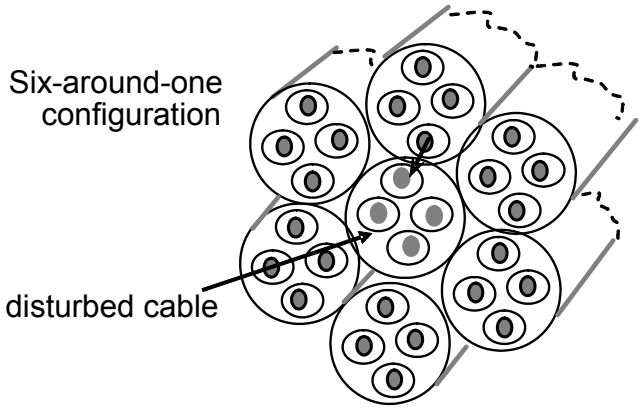

Figure 55B–2—Six-around-one cable configuration

The alien crosstalk coupling is a function of the distance between the cabling components. Removing or loosening cable bindings, separating equipment and patch cordage, and utilizing non-adjacent connectors all have a significant impact on reducing the levels of alien crosstalk noise.

In many topologies, the cables are only bundled together for relatively short distances and laid in cable trays. The star wiring topology, where the cables are distributed radially from a centralized telecommunication closet to each work area, reduces the distances over which link segments are in close proximity.

55B.1.1 Alien crosstalk mitigation

10GBASE-T is designed to operate over a channel that meets the specifications of 55.7 and the channel configuration shown Figure 55B–3. However, if the channel specifications of 55.7 cannot be met when using the channel configuration shown in Figure 55B–3, the configuration shown in Figure 55B–4 is recommended. This optimized channel for a 10GBASE-T link segment deletes the consolidation point or optional transition point and runs an equipment patch cord directly between the LAN equipment and the connector termination of the permanent link. This reduces the number of connectors and cordage and their associated alien crosstalk.

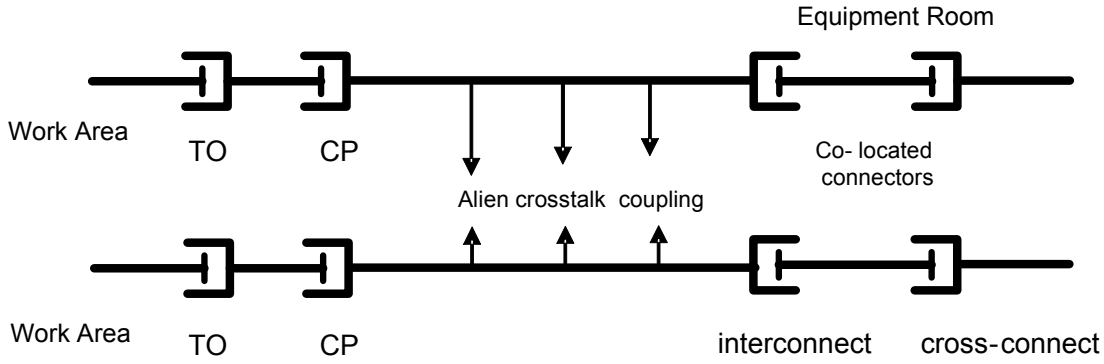

CP = consolidation point connector
TO = telecommunication outlet connector

Figure 55B–3—Maximum horizontal cabling configuration

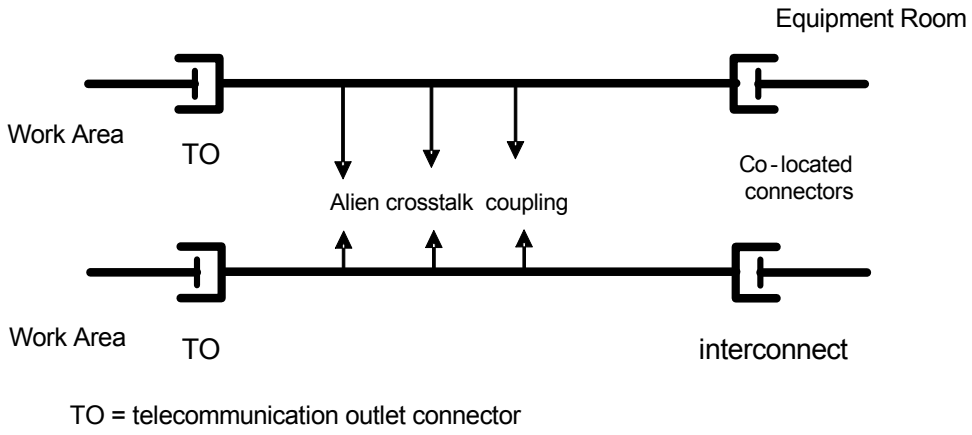

Figure 55B–4—Minimum horizontal cabling configuration

55B.1.2 Alien crosstalk mitigation procedure

This annex provides procedures and cabling guidelines designed to mitigate the alien crosstalk in the event that the alien crosstalk transmission parameters given in 55.7 are not met. For more information on mitigation techniques, see TIA TSB-155-A and ISO/IEC TR 24750.

The mitigation actions outlined below are based on four connector channels. In the majority of initially non-compliant cases, fewer than all corrective actions are required. Select the option(s) that is most appropriate for your situation.

a) When selective deployment of 10GBASE-T is possible, non-adjacent patch panel positions in the equipment room should be used. The adjacent positions may be used for other applications.
b) Reduce the number of co-located connectors by implementing an interconnect configuration to attach equipment to the horizontal cabling rather than a cross-connect. In general, connectors and cordage in the work area are not co-located.
c) Reduce the alien crosstalk coupling in the first 5 m to 20 m of the horizontal cabling by separating the equipment cords and the patch cords and un-bundling the horizontal cabling: in the case of a telecommunications room un-bundle the cabling to the point it exits the telecommunications room. A significant portion of the ANEXT coupling occurs in less than the first 20 m of cabling.
d) An alternative to separating equipment cords is to utilize equipment cords sufficiently specified to mitigate the alien crosstalk coupling.
e) Replace connectors with Category 6A/Class E_A.