IEEE Standard for Ethernet

SECTION SIX

This section includes Clause 78 through Clause 95 and Annex 83A through Annex 93C.

Contents

78. Energy-Efficient Ethernet (EEE)

78.1 Overview

The optional EEE capability combines the IEEE 802.3 Media Access Control (MAC) Sublayer with a family of Physical Layers defined to support operation in the Low Power Idle (LPI) mode. When the LPI mode is enabled, systems on both sides of the link can save power during periods of low link utilization.

EEE also provides a protocol to coordinate transitions to or from a lower level of power consumption and does this without changing the link status and without dropping or corrupting frames. The transition time in to and out of the lower level of power consumption is kept small enough to be transparent to upper layer protocols and applications.

EEE supports operation over twisted-pair cabling systems, twinax cable, electrical backplanes, optical fiber, the XGXS for 10 Gb/s PHYs, the XLAUI for 40 Gb/s PHYs, and the CAUI-10 or CAUI-4 for 100 Gb/s PHYs. Table 78–1 lists the supported PHYs and interfaces and their associated clauses.

In addition to the above, EEE defines a 10 Mb/s MAU (10BASE-Te) with reduced transmit amplitude requirements. The 10BASE-Te MAU is fully interoperable with 10BASE-T MAUs over 100 m of class D (Category 5) or better cabling as specified in ISO/IEC 11801:1995. These requirements can also be met by Category 5 cable and components as specified in ANSI/TIA/EIA-568-B-1995. The definition of 10BASE-Te allows a reduction in power consumption.

EEE also specifies means to exchange capabilities between link partners to determine whether EEE is supported and to select the best set of parameters common to both devices. Clause 78 provides an overview of EEE operation. PICS for the optional EEE capability for each specific PHY type are specified in the respective PHY clauses. Normative requirements for Data Link Layer capabilities are contained in 78.4.

78.1.1 LPI Signaling

LPI signaling allows the LPI Client to indicate to the PHY, and to the link partner, that a break in the data stream is expected, and the LPI Client can use this information to enter power-saving modes that require additional time to resume normal operation. LPI signaling also informs the LPI Client when the link partner has sent such an indication.

The definition of LPI signaling assumes 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 the LPI mode.

The LPI Client connects to the RS service interface. LPI signaling between the RS and PCS is performed by LPI encoding on the Media Independent Interface. The transmit PCS encodes LPI symbols, which are decoded by the link partner receive PCS. The receive and transmit PCS also generate service interface signals, which are passed down to the lower PHY sublayers and indicate when receive and transmit PHY functions may be powered down.

The EEE request signals from the PCS control transitions between quiescent and normal operation. The Clause 49 PCS and Clause 82 PCS also request transmit alert operation to assist the partner device PMD to detect the end of the quiescent state. Additionally the Clause 49 PCS and Clause 82 PCS generate the RX_LPI_ACTIVE signal, which indicates to the Clause 74 BASE-R FEC that it can use rapid block lock because the link partner PCS has bypassed scrambling.

Coding defined in Clause 83 also allows LPI transmit quiet and alert requests from the PCS to be signaled over the XLAUI and CAUI-n interfaces. The XLAUI and CAUI-n receive interfaces infer the quiet and alert

requests from the data received over the interface and use that to recreate the transmit or receive direction signaling. (See 83.5.11.1.)

The receive PCS checks that the link cycles out of the quiescent state at the correct time and that the received signals return to their expected state within the required time. The ENERGY_DETECT indicate signal is passed up from the PMA to the PCS to allow the PCS to monitor the waking process.

78.1.1.1 Reconciliation sublayer service interfaces

Figure 78–1 depicts the LPI Client and the RS interlayer service interfaces.

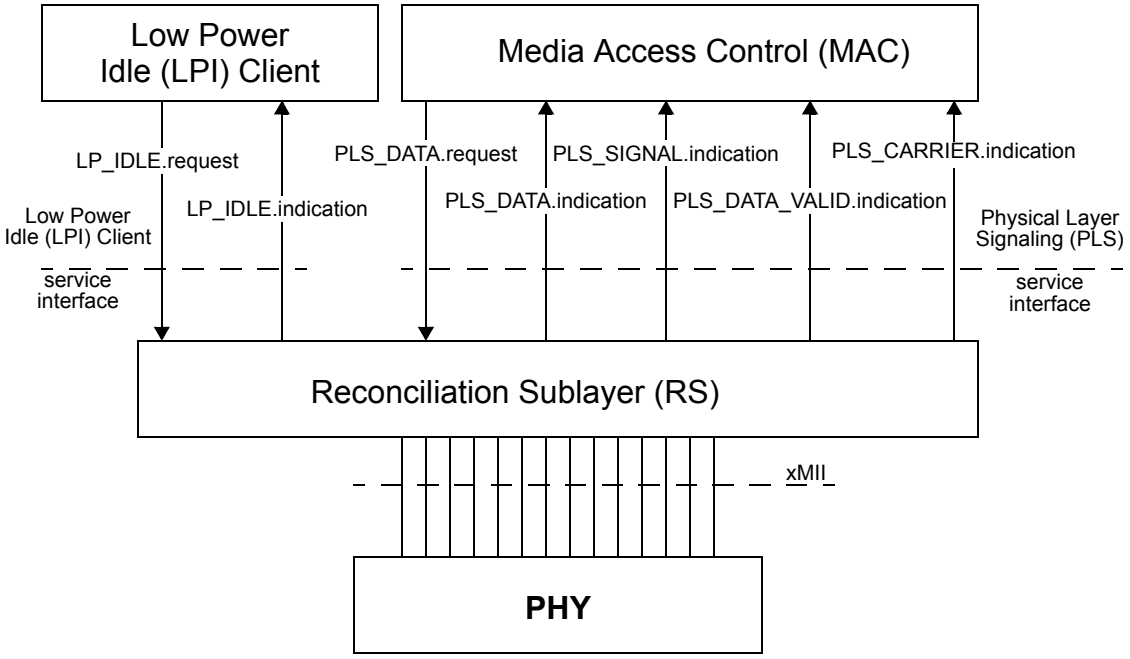

Figure 78–1—LPI Client and RS interlayer service interfaces

78.1.1.2 Responsibilities of LPI Client

The decision on when to signal LPI to the link partner is made by the LPI Client and communicated to the PHY through the RS. The LPI Client is also informed when the link partner is signaling LPI by the RS.

The conditions under which the LPI Client decides to send LPI, and what action are taken by the LPI Client when it receives LPI from the link partner, are implementation specific and beyond the scope of this standard.

78.1.2 LPI Client service interface

The following specifies the service interface provided by the RS to the LPI Client. These services are described in an abstract manner and do not imply any particular implementation.

The following primitives are defined:

LP_IDLE.request
LP_IDLE.indication

78.1.2.1 LP_IDLE.request

78.1.2.1.1 Function

A primitive used by the LPI Client to start or stop the signaling of LPI to the link partner.

78.1.2.1.2 Semantics of the service primitive

The semantics of the service primitive are as follows:

LP_IDLE.request (LPI_REQUEST)

The LPI_REQUEST parameter can take one of two values: ASSERT or DEASSERT. ASSERT initiates the signaling of LPI to the link partner. DEASSERT stops the signaling of LPI to the link partner. The effect of receipt of this primitive is undefined in any of the following cases:

a) link_status is not OK (see 28.2.6.1.1)
b) LPI_REQUEST=ASSERT within 1 s of the change of link_status to OK
c) The PHY is indicating LOCAL FAULT
d) The PHY is indicating REMOTE FAULT

78.1.2.1.3 When generated

Specification of the time when this primitive is generated by the LPI client is out of the scope of this standard.

78.1.2.1.4 Effect of receipt

The receipt of this primitive will cause the RS to start or stop signaling LPI to the link partner.

78.1.2.2 LP_IDLE.indication

78.1.2.2.1 Function

A primitive that is used to indicate to the LPI Client that the link partner has started or stopped signaling LPI.

78.1.2.2.2 Semantics of the service primitive

The semantics of the service primitive are as follows:

LP_IDLE.indication (LPI_INDICATION)

The LPI_INDICATION parameter can take one of two values: ASSERT or DEASSERT. ASSERT indicates that the link partner has started signaling LPI. DEASSERT indicates that the link partner has stopped signaling LPI.

78.1.2.2.3 When generated

This primitive is generated by the RS when it starts or stops receiving Assert LPI encoded on the receive xMII according to the rules defined in 78.1.3.2.

78.1.2.2.4 Effect of receipt

The effect of receipt of this primitive by the LPI client is unspecified.

78.1.3 Reconciliation sublayer operation

LPI assert and detect functions are contained in the Reconciliation Sublayer as shown in Figure 78–2. The xMII in this diagram represents any of the family of medium independent interfaces supported by EEE.

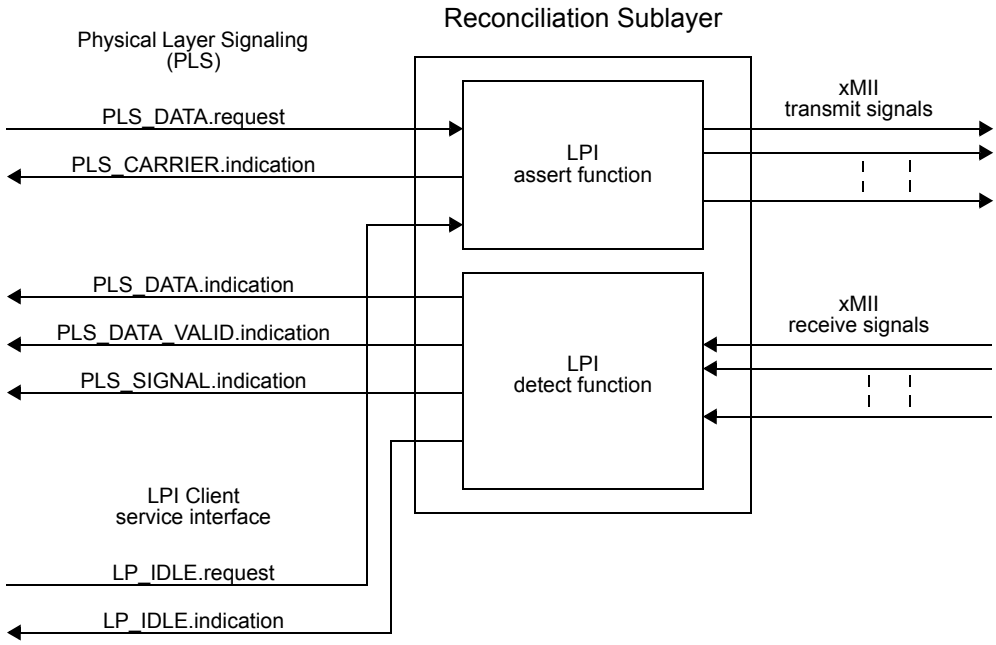

Figure 78–2—RS LPI assert and detect functions

The following provides an overview of RS LPI operation. The actual specification of RS LPI operation can be found in the respective RS clauses.

78.1.3.1 RS LPI assert function

In the absence of an LPI request, indicated by the LPI_REQUEST parameter set to DEASSERT in the LP_IDLE.request primitive of the LPI Client interface, the LPI assert function maps the PLS service interface to the transmit xMII signals as under normal conditions.

When an LPI request is asserted, indicated by the LPI_REQUEST parameter set to ASSERT in the LP_IDLE.request primitive of the LPI Client interface, the LPI assert function starts to transmit the "Assert LPI" encoding on the xMII. The LPI assert function also sets the CARRIER_STATUS parameter to CARRIER_ON in the PLS_CARRIER.indication primitive of the PLS service interface. This will prevent the MAC from transmitting.

When the LPI request is deasserted, indicated by the LPI_REQUEST parameter set to DEASSERT in the LP_IDLE.request primitive of the LPI Client interface, the LPI assert function starts to transmit the normal interframe encoding on the xMII. After a delay, the LPI assert function sets the CARRIER_STATUS parameter to CARRIER_OFF in the PLS_CARRIER.indication primitive of the PLS service interface,

allowing the MAC to start transmitting again. This delay is provided to allow the link partner to prepare for normal operation. This delay has a PHY dependent default value but this value can be adjusted using the Data Link Layer capabilities defined in 78.4.

78.1.3.2 LPI detect function

In the absence of LPI, indicated by an encoding other than "Assert LPI" on the receive xMII, the LPI detect function maps the receive xMII signals to the PLS service interface as under normal conditions.

At the start of LPI, indicated by the transition from normal interframe encoding to the "Assert LPI" encoding on the receive xMII, the LPI detect function continues to indicate idle on the PLS service interface, but sets LP_IDLE.indication(LPI_INDICATION) to ASSERT.

At the end of LPI, indicated by the transition from the "Assert LPI" encoding to any other encoding on the receive xMII, LP_IDLE.indication(LPI_INDICATION) is set to DEASSERT and the RS receive function resumes normal decode operation.

78.1.3.3 PHY LPI operation

The following provides an overview of PHY LPI operation. The specification of PHY LPI operation can be found in the respective PHY clauses (see Table 78–1).

78.1.3.3.1 PHY LPI transmit operation

When the start of "Assert LPI" encoding on the xMII is detected, the PHY signals sleep to its link partner to indicate that the local transmitter is entering LPI mode.

The EEE capability in most PHYs (for example, 100BASE-TX, 10GBASE-T, 1000BASE-KX, 10GBASE-KR, and 10GBASE-KX4) requires the local PHY transmitter to go quiet after sleep is signalled.

In the 1000BASE-T LPI mode, the local PHY transmitter goes quiet only after the local PHY signals sleep and receives a sleep signal from the remote PHY. If the remote PHY chooses not to signal LPI, then neither PHY can go into a low power mode; however, LPI requests are passed from one end of the link to the other regardless and system energy savings can be achieved even if the PHY link does not go into a low power mode.

The transmit function of the local PHY is enabled periodically to transmit refresh signals that are used by the link partner to update adaptive filters and timing circuits in order to maintain link integrity.

This quiet-refresh cycle continues until the reception of the normal interframe encoding on the xMII. The transmit function in the PHY communicates this to the link partner by sending a wake signal for a predefined period of time. The PHY then enters the normal operating state.

Figure 78–3 illustrates general principles of the EEE-capable transmitter operation.

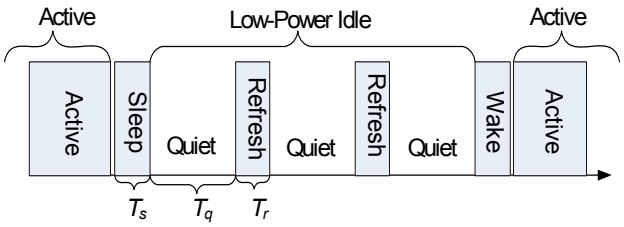

Figure 78–3—Overview of EEE LPI operation

No data frames are lost or corrupted during the transition to or from the LPI mode.

For PHYs with an operating speed of 40 Gb/s or greater that implement the optional EEE capability, two modes of LPI operation may be supported: deep sleep and fast wake. *Deep sleep* refers to the mode for which the transmitter ceases transmission during Low Power Idle (as shown in Figure 78–3) and is equivalent to the only mechanism defined for PHYs with an operating speed less than 40 Gb/s. Deep sleep support is optional for PHYs with an operating speed of 40 Gb/s or greater that implement EEE with the exception of the PHYs noted in Table 78–1 that do not support deep sleep. *Fast wake* refers to the mode for which the transmitter continues to transmit signals during Low Power Idle so that the receiver can resume operation with a shorter wake time (as shown in Figure 78–4). For transmit, other than the PCS encoding LPI, there is no difference between fast wake and normal operation. Fast wake support is mandatory for PHYs with an operating speed of 40 Gb/s or greater that implement EEE.

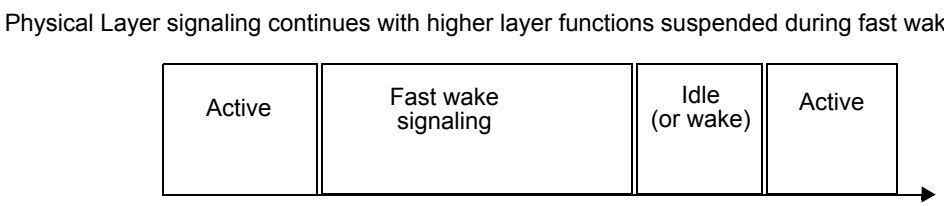

NOTE—Fast wake signaling continually indicates LPI in a normally constituted data stream.

Figure 78–4—Overview of fast wake operation

WARNING
The signaling in deep sleep operation precludes transparent mapping of the link over Optical Transport Networks. Only fast wake operation should be enabled for any link that is intended for transparent OTN mapping.

78.1.3.3.2 PHY LPI receive operation

In the receive direction, entering the LPI mode is triggered by the reception of a sleep signal from the link partner, which indicates that the link partner is about to enter the LPI mode. After sending the sleep signal, the link partner ceases transmission if not in fast wake mode. When the receiver detects the sleep signal, the local PHY indicates “Assert LPI” on the xMII and the local receiver can disable some functionality to reduce power consumption.

If not in fast wake mode the link partner periodically transmits refresh signals that are used by the local PHY to update adaptive coefficients and timing circuits. This quiet-refresh cycle continues until the link partner

initiates transition back to normal mode by transmitting the wake signal for a predetermined period of time controlled by the LPI assert function in the RS. This allows the local receiver to prepare for normal operation and transition from the "Assert LPI" encoding to the normal interframe encoding on the xMII. After a system specified recovery time, the link supports the nominal operational data rate.

78.1.4 PHY types optionally supporting EEE

EEE defines a low power mode of operation for the IEEE 802.3 PHYs and interfaces listed in Table 78–1. The table also lists the clauses associated with each PHY or sublayer. Normative requirements for the EEE capability for each PHY type and interface are in the associated clauses.

Table 78–1—Clauses associated with each PHY or interface type

PHY or interface type	Clause
10BASE-Te	14
100BASE-TX	24, 25
1000BASE-KX	70, 36
1000BASE-T	40
XGXS (XAUI)	47, 48
10GBASE-KX4	71, 48
10GBASE-KR	72, 51, 49, 74
10GBASE-T	55
XLAUI/CAUI-10[a]	83A
40GBASE-KR4	82, 83, 84, 74
40GBASE-CR4	82, 83, 85, 74
40GBASE-SR4[b]	82, 83, 86
40GBASE-FR[b]	82, 83, 89
40GBASE-LR4[b]	82, 83, 87
40GBASE-ER4[b]	82, 83, 87
CAUI-4[a]	83D
100GBASE-KP4	82, 91, 94
100GBASE-KR4	82, 83, 91, 93
100GBASE-CR10	82, 83, 85, 74
100GBASE-CR4	82, 83, 91, 92
100GBASE-SR10[b]	82, 83, 86
100GBASE-SR4[b]	82, 83, 91, 95
100GBASE-LR4[b]	82, 83, 88
100GBASE-ER4[b]	82, 83, 88

[a]XLAUI/CAUI-n shutdown is supported only when deep sleep is enabled for the associated PHY.
[b]The deep sleep mode of EEE is not supported for this PHY.

78.2 LPI mode timing parameters description

T_s The period of time that the PHY transmits the sleep signal before turning all transmitters off

T_q The period of time that the PHY remains quiet before sending the refresh signal

T_r Duration of the refresh signal

$T_{phy\_prop\_tx}$ The propagation delay of a given unit of data from the xMII to the MDI

$T_{phy\_prop\_rx}$ The propagation delay of a given unit of data from the MDI to the xMII

$T_{phy\_shrink\_tx}$ Transmitter shrinkage time, defined as the absolute time difference between the following two timing parameters:

—Delay between a transition from the "Assert LPI" to "Normal Idle" at the xMII and the corresponding start of the wake signal at the MDI

—$T_{phy\_prop\_tx}$

$T_{phy\_shrink\_rx}$ Receiver shrinkage time, defined as the absolute time difference between the following two timing parameters:

—Delay between start of the wake signal at the MDI and the corresponding transition from "Assert LPI" to "Normal Idle" at the xMII

—$T_{phy\_prop\_rx}$

T_{w_phy} Parameter employed by the system that corresponds to the behavior of the PHY. It is the period of time between reception of an IDLE signal on the xMII and when the first data codewords are permitted on the xMII. The wake time of a compliant PHY does not exceed T_{w_phy} (min).

$T_{w\_sys\_tx}$ Parameter employed by the system that corresponds to its requirements. It is the longest period of time the system has to wait between a request to transmit and its readiness to transmit.

$T_{w\_sys\_rx}$ Parameter employed by the system that corresponds to its requirements. It is the minimum time required by the system between a request to wake and its readiness to receive data.

Table 78–2 summarizes three key EEE parameters (T_s, T_q, and T_r) for supported PHYs.

Table 78–2—Summary of the key EEE parameters for supported PHYs or interfaces

PHY or interface type	T_s (μs)		T_q (μs)		T_r (μs)	
	Min	Max	Min	Max	Min	Max
100BASE-TX	200	220	20 000	22 000	200	220
1000BASE-KX	19.9	20.1	2 500	2 600	19.9	20.1
1000BASE-T	182	202	20 000	24 000	198	218.2
XGXS (XAUI)	19.9	20.1	2 500	2 600	19.9	20.1
10GBASE-KX4	19.9	20.1	2 500	2 600	19.9	20.1
10GBASE-KR	4.9	5.1	1 700	1 800	16.9	17.5

Table 78–2—Summary of the key EEE parameters for supported PHYs or interfaces *(continued)*

PHY or interface type	T_s (µs)		T_q (µs)		T_r (µs)	
	Min	Max	Min	Max	Min	Max
10GBASE-T	2.88	3.2	39.68	39.68	1.28	1.28
40GBASE-KR4	0.9	1.1	1 700	1 800	5.9	6.5
40GBASE-CR4	0.9	1.1	1 700	1 800	5.9	6.5
100GBASE-KP4	0.9	1.1	1 700	1 800	5.9	6.5
100GBASE-KR4	0.9	1.1	1 700	1 800	5.9	6.5
100GBASE-CR10	0.9	1.1	1 700	1 800	5.9	6.5
100GBASE-CR4	0.9	1.1	1 700	1 800	5.9	6.5

Figure 78–5 illustrates the relationship between the LPI mode timing parameters and the minimum system wake time.

78.3 Capabilities Negotiation

The EEE capability shall be advertised during the Auto-Negotiation stage, except for PHYs that only support fast wake operation. Auto-Negotiation provides a linked device with the capability to detect the abilities (modes of operation) supported by the device at the other end of the link, determine common abilities, and configure for joint operation. Auto-Negotiation is performed at power up, on command from management, due to link failure, or due to user intervention. Fast wake capability shall be advertised using L2 protocol frames as described in 78.4.

During Auto-Negotiation, both link partners indicate their EEE capabilities. EEE is supported only if during Auto-Negotiation both the local device and link partner advertise the EEE capability for the resolved PHY type. If EEE is not supported, all EEE functionality is disabled and the LPI client does not assert LPI. EEE deep sleep operation shall not be enabled unless both the local device and link partner advertise deep sleep capability during Auto-Negotiation for the resolved PHY type. If EEE is supported by both link partners for the negotiated PHY type, then the EEE function can be used independently in either direction.

Additional capabilities and settings using L2 protocol frames, including the adjustment of the $T_{w\_sys\_tx}$ parameter, are described in 78.4.

78.4 Data Link Layer Capabilities

Additional capabilities and settings are supported using frames based on the IEEE 802.3 Organizationally Specific TLVs are defined in Annex F of IEEE Std 802.1AB™-2009. Devices that require longer wake-up times prior to being able to accept data on their receive paths may use the Data Link Layer capabilities defined in this subclause to negotiate for extended system wake-up times from the transmitting link partner. This mechanism may allow for more or less aggressive energy saving modes.

The Data Link Layer capabilities shall be implemented for devices with an operating speed equal to or greater than 10 Gb/s and may be implemented for all other devices. The use of the EEE Fast Wake TLV shall be interpreted as an indication that the device supports EEE fast wake operation, regardless of the capability

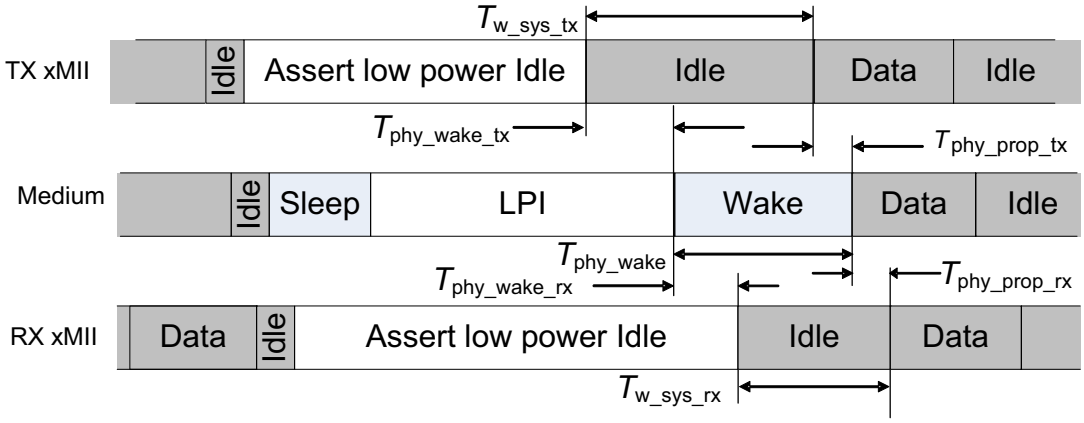

$T_{w\_sys\_tx}$ (min) = $T_{w\_sys\_rx}$ (min) + $T_{phy\_shrink\_tx}$ (max) + $T_{phy\_shrink\_rx}$ (max)
T_{w_phy} (min) = T_{phy_wake} (min) + $T_{phy\_shrink\_tx}$

$T_{w\_sys\_res}$ (min) is greater of $T_{w\_sys\_tx}$ (min) and T_{w_phy} (min)

$T_{phy\_shrink\_tx}$ (max) = ($T_{phy\_wake\_tx}$ (max) – $T_{phy\_prop\_tx}$ (min))
$T_{phy\_shrink\_rx}$ (max) = ($T_{phy\_wake\_rx}$(max) – $T_{phy\_prop\_rx}$ (min))

where

$T_{phy\_wake\_tx}$ = xMII start of wake to MDI start of wake delay
$T_{phy\_prop\_tx}$ = xMII to MDI data propagation delay
$T_{phy\_wake\_rx}$ = MDI start of wake to xMII start of wake delay
$T_{phy\_prop\_rx}$ = MDI to xMII data propagation delay
T_{phy_wake} = Minimum wake duration required by PHY

Figure 78–5—LPI mode timing parameters and their relationship to minimum system wake time

advertised during the Auto-Negotiation stage. A device shall not indicate deep sleep capability using the EEE Fast Wake TLV unless both the local device and link partner advertise deep sleep capability during Auto-Negotiation for the resolved PHY type.

Implementations that use the Data Link Layer capabilities shall comply with all mandatory parts of IEEE Std 802.1AB-2009; shall support the EEE Type, Length, Value (TLV) defined in 79.3.5; timing requirement in 78.4.1; and shall support the control state diagrams defined in 78.4.2. Devices with an operating speed equal to or greater than 40 Gb/s shall support EEE Fast Wake TLV as defined in 79.3.6.

The Data Link Layer capabilities are described from a unidirectional perspective on the link between transmitting and receiving link partners. For duplex EEE links that implement the Data Link Layer capabilities, each link partner shall implement the TLV, control and state diagrams for a transmitter as well as a receiver.

For purposes of Data Link Layer capabilities, all values that are negotiated and/or exchanged that have a fractional value shall be rounded up to the nearest integer number in microseconds.

78.4.1 Data Link Layer capabilities timing requirements

An EEE link partner shall send an LLDPDU containing an EEE TLV within 10 s of the Link Layer capability exchange being enabled when both the variables dll_enabled and dll_ready are asserted.

An LLDPDU containing an EEE TLV with an updated value for the "Echo Transmit $T_{w\_sys\_tx}$" field shall be sent within 10 s of receipt of an LLDPDU containing an EEE TLV where the value of "Transmit $T_{w\_sys\_tx}$" field is different from the previously communicated value.

An LLDPDU containing an EEE TLV with an updated value for the "Echo Receive $T_{w\_sys\_tx}$" field shall be sent within 10 s of receipt of an LLDPDU containing an EEE TLV where the value of "Receive $T_{w\_sys\_tx}$" field is different from the previously communicated value.

78.4.2 Control state diagrams

The control state diagrams for an EEE transmitting link partner and an EEE receiving link partner specify the externally observable behavior of an EEE transmitting link partner and an EEE receiving link partner implementing Data Link Layer capabilities respectively. EEE transmitting link partners implementing Data Link Layer capabilities shall provide the behavior of the state diagram as shown in Figure 78–6. EEE receiving link partners implementing Data Link Layer capabilities shall provide the behavior of the state diagram as shown in Figure 78–7.

78.4.2.1 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 state diagrams as described in 21.5.

78.4.2.2 Constants

PHY WAKE VALUE
: Integer (2 octets wide) representing the $T_{w\_sys\_tx}$ (min) defined for the PHY that is in use for the link. This parameter should be rounded up to the nearest integer number when it is calculated and examined according to 78.2 and Table 78–4.

78.4.2.3 Variables

Unless otherwise specified, all integers are assumed to be 2 octets wide.

LocTxSystemValue
: Integer that indicates the value of $T_{w\_sys\_tx}$ that the local system can support. This value is updated by the EEE DLL Transmitter state diagram. This variable maps into the aLldpXdot3LocTxTwSys attribute.

RemTxSystemValueEcho
: Integer that indicates the value Transmit $T_{w\_sys\_tx}$ echoed back by the remote system. This value maps from the aLldpXdot3RemTxTwSysEcho attribute.

LocRxSystemValue
: Integer that indicates the value of $T_{w\_sys\_tx}$ that the local system requests from the remote system. This value is updated by the EEE Receiver L2 state diagram. This variable maps into the aLldpXdot3LocRxTwSys attribute.

RemRxSystemValueEcho

Integer that indicates the value of Receive $T_{w\_sys\_tx}$ echoed back by the remote system. This value maps from the aLldpXdot3RemRxTwSysEcho attribute.

LocFbSystemValue

Integer that indicates the value of fallback $T_{w\_sys\_tx}$ that the local system requests from the remote system. This value is updated by the local system software.

RemTxSystemValue

Integer that indicates the value of $T_{w\_sys\_tx}$ that the remote system can support. This value maps from the aLldpXdot3RemTxTwSys attribute.

LocTxSystemValueEcho

Integer that indicates the remote system's Transmit $T_{w\_sys\_tx}$ that was used by the local system to compute the $T_{w\_sys\_tx}$ that it wants to request from the remote system. This value maps into the aLldpXdot3LocTxTwSysEcho attribute.

RemRxSystemValue

Integer that indicates the value of $T_{w\_sys\_tx}$ that the remote system requests from the local system. This value maps from the aLldpXdot3RemRxTwSys attribute.

LocRxSystemValueEcho

Integer that indicates the remote systems Receive $T_{w\_sys\_tx}$ that was used by the local system to compute the $T_{w\_sys\_tx}$ that it can support. This value maps into the aLldpXdot3LocRxTwSysEcho attribute.

LocResolvedTxSystemValue

Integer that indicates the current $T_{w\_sys\_tx}$ supported by the local system.

LocResolvedRxSystemValue

Integer that indicates the current $T_{w\_sys\_tx}$ supported by the remote system.

LPI_FW

Boolean variable controlling the wake mode for the LPI transmit and receive functions as defined in 82.2.19.2.2.

LocTxSystemFW

Boolean variable that indicates the state of LPI_FW that the local transmit system can support. This value is updated by the EEE DLL Transmit fast wake state diagram. This variable maps into the aLldpXdot3LocTxFw attribute.

RemTxSystemFWEcho

Boolean variable that indicates the state of transmit LPI_FW echoed back by the remote system. This value maps from the aLldpXdot3RemTxFwEcho attribute.

LocRxSystemFW

Boolean variable that indicates the state of LPI_FW that the local receive system requests from the remote system. This value is updated by the EEE DLL Receive fast wake state diagram. This variable maps into the aLldpXdot3LocRxFw attribute.

RemRxSystemFWEcho
: Boolean variable that indicates the state of receive LPI_FW echoed back by the remote system. This value maps from the aLldpXdot3RemRxFwEcho attribute.

RemTxSystemFW
: Boolean variable that indicates the LPI_FW that the remote transmit system requests from the local system. This value maps from the aLldpXdot3RemTxFw attribute.

LocTxSystemFWEcho
: Boolean variable that indicates the remote system's transmit LPI_FW that was used by the local system to decide the LPI_FW that it wants to request from the remote system. This value maps into the aLldpXdot3LocTxFwEcho attribute.

RemRxSystemFW
: Boolean variable that indicates the LPI_FW that the remote receive system requests from the local system. This value maps from the aLldpXdot3RemRxFw attribute.

LocRxSystemFWEcho
: Boolean variable that indicates the remote system's receive LPI_FW that was used by the local system to decide the LPI_FW that it can support. This value maps into the aLldpXdot3LocRxFwEcho attribute.

LocResolvedTxSystemFW
: Boolean that indicates the current LPI_FW supported by the local system.

LocResolvedRxSystemFW
: Boolean variable that indicates the current LPI_FW supported by the remote system.

TempTxFW
: Boolean variable used to store the value of LPI_FW.

TempRxFW
: Boolean variable used to store the value of LPI_FW.

local_system_FW_change
: An implementation-specific control variable that indicates that the local system wants to change either the Transmit LPI_FW or the Receive LPI_FW.

NEW_TX_FW
: Boolean variable that indicates the value of transmit LPI_FW that the local system can support.

NEW_RX_FW
: Boolean variable that indicates the value of receive LPI_FW that the local system wants the remote system to support.

TempTxVar
: Integer used to store the value of $T_{w\_sys\_tx}$.

TempRxVar
: Integer used to store the value of $T_{w\_sys\_tx}$.

local_system_change
: An implementation specific control variable that indicates that the local system wants to change either the Transmit $T_{w\_sys\_tx}$ or the Receive $T_{w\_sys\_tx}$.

tx_dll_ready
: Data Link Layer ready: This variable indicates that the tx system initialization is complete and is ready to update/receive LLDPDU containing EEE TLV. This variable is updated by the local system software.

rx_dll_ready
: Data Link Layer ready: This variable indicates that the rx system initialization is complete and is ready to update/receive LLDPDU containing EEE TLV. This variable is updated by the local system software.

NEW_TX_VALUE
: Integer that indicates the value of $T_{w\_sys\_tx}$ that the local system can support.

NEW_RX_VALUE
: Integer that indicates the value of $T_{w\_sys\_tx}$ that the local system wants the remote system to support.

A summary of cross-references between the EEE object class attributes and the transmit and receive control state diagrams, including the direction of the mapping, is provided in Table 78–3.

Table 78–3—Attribute to state diagram variable cross-reference

Entity	Object class	Attribute	Mapping	State diagram variable
TX	oLldpXdot3Loc-SystemsGroup	aLldpXdot3LocTxTwSys	⇐	LocTxSystemValue
		aLldpXdot3LocRxTwSysEcho	⇐	LocRxSystemValueEcho
		aLldpXdot3LocDllEnabled	⇒	tx_dll_enabled
		aLldpXdot3LocTxDllReady	⇐	tx_dll_ready
		aLldpXdot3LocTxFw	⇐	LocTxSystemFW
		aLldpXdot3LocRxFwEcho	⇐	LocRxSystemFWEcho
	oLldpXdot3Rem-SystemsGroup	aLldpXdot3RemRxTwSys	⇒	RemRxSystemValue
		aLldpXdot3RemTxTwSysEcho	⇒	RemTxSystemValueEcho
		aLldpXdot3RemRxFw	⇒	RemRxSystemFW
		aLldpXdot3RemTxFwEcho	⇒	RemTxSystemFWEcho

Table 78–3—Attribute to state diagram variable cross-reference *(continued)*

Entity	Object class	Attribute	Mapping	State diagram variable
RX	oLldpXdot3Loc-SystemsGroup	aLldpXdot3LocRxTwSys	⇐	LocRxSystemValue
		aLldpXdot3LocTxTwSysEcho	⇐	LocTxSystemValueEcho
		aLldpXdot3LocFbTwSys	⇐	LocFbSystemValue
		aLldpXdot3LocDllEnabled	⇒	rx_dll_enabled
		aLldpXdot3LocRxDllReady	⇐	rx_dll_ready
		aLldpXdot3LocRxFw	⇐	LocRxSystemFW
		aLldpXdot3LocTxFwEcho	⇐	LocTxSystemFWEcho
	oLldpXdot3Rem-SystemsGroup	aLldpXdot3RemTxTwSys	⇒	RemTxSystemValue
		aLldpXdot3RemRxTwSysEcho	⇐	RemRxSystemValueEcho
		aLldpXdot3RemTxFw	⇒	RemTxSystemFW
		aLldpXdot3RemRxFwEcho	⇐	RemRxSystemFWEcho

78.4.2.4 Functions

examine_Tx_change

This function computes the new value of $T_{w\_sys\_tx}$ that the local system can support when there is as updated request from the remote system or if local system conditions require a change in the value of the presently supported $T_{w\_sys\_tx}$.

examine_Rx_change

This function computes the new value of $T_{w\_sys\_tx}$ that the local system wants the remote system to support. This function is called when the remote system wants to change its presently allocated $T_{w\_sys\_tx}$ or if local system conditions require a change in the value of $T_{w\_sys\_tx}$ presently supported by the remote system.

examine_TxFW_change

This function decides if the new value of LPI_FW is acceptable by the local transmit system when there is an updated request from the remote system or if local system conditions require a change in the value of the presently supported LPI_FW.

examine_RxFW_change

This function decides if the new value of LPI_FW is acceptable by the local receive system when there is an updated request from the remote system or if local system conditions require a change in the value of the presently supported LPI_FW.

78.4.2.5 State diagrams

Control for placing data on the medium rests with the transmitting side, hence $T_{w\_sys\_tx}$ is enforced by the transmitter. For a given path between link partners (i.e., a transmitter and its associated receiver), the transmitting link partner shall wait for the time indicated by the Transmit $T_{w\_sys\_tx}$ after deasserting LPI (at the xMII) before sending data frames. The receiving link partner shall be ready to accept data based on its

echoed value of Transmit link partner's $T_{w\_sys\_tx}$. This ensures that the link partners transition out of LPI mode and receive frames without loss or corruption.

The general state change procedure for transmitter is shown in Figure 78–6.

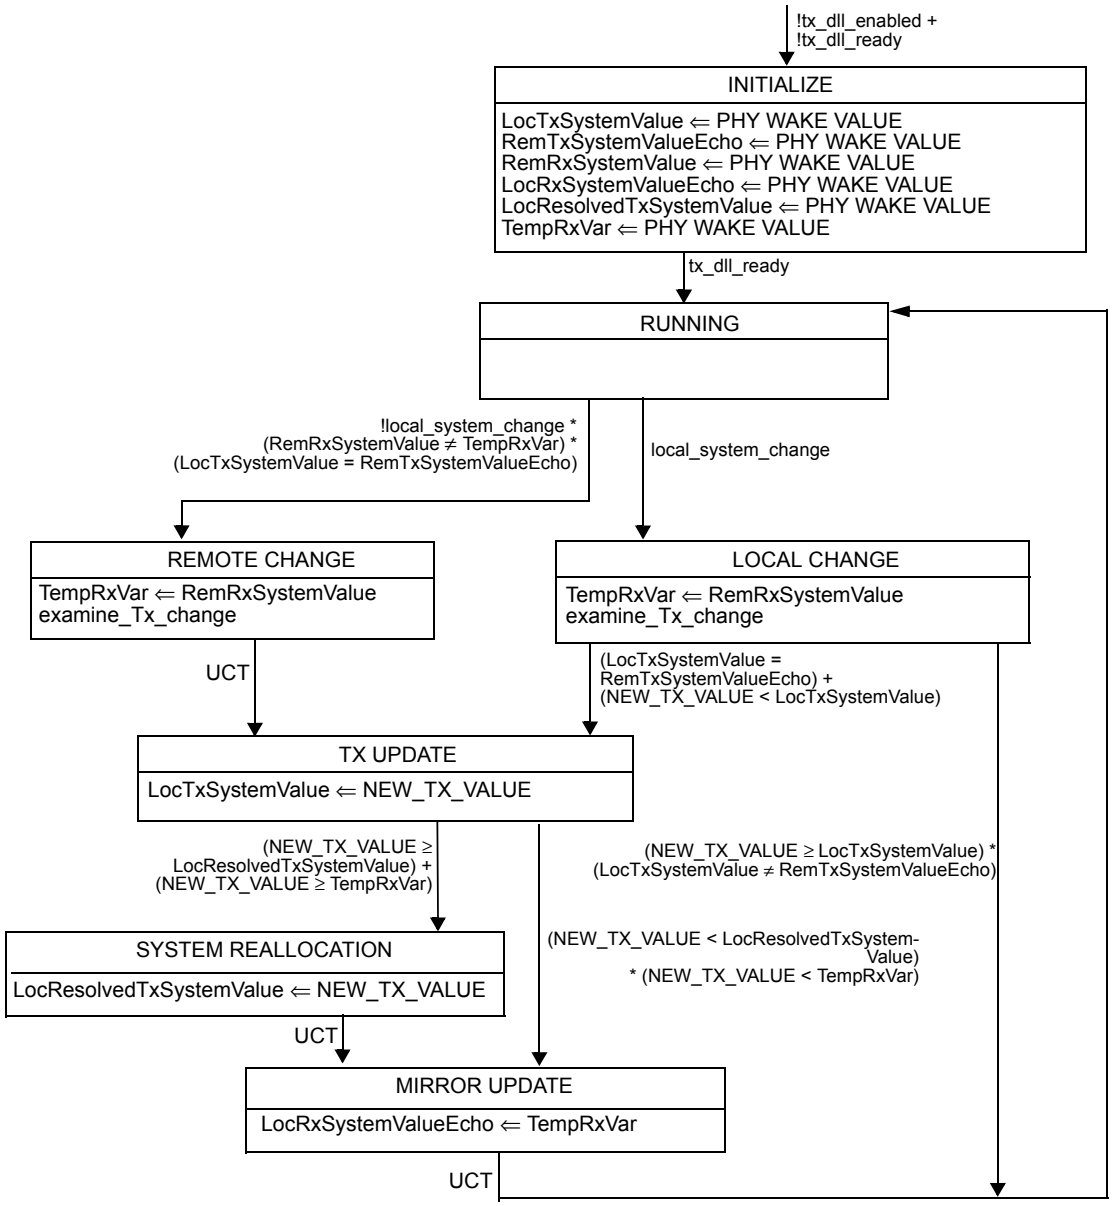

Figure 78–6—EEE DLL Transmitter state diagram

The general state change procedure for receiver is shown in Figure 78–7.

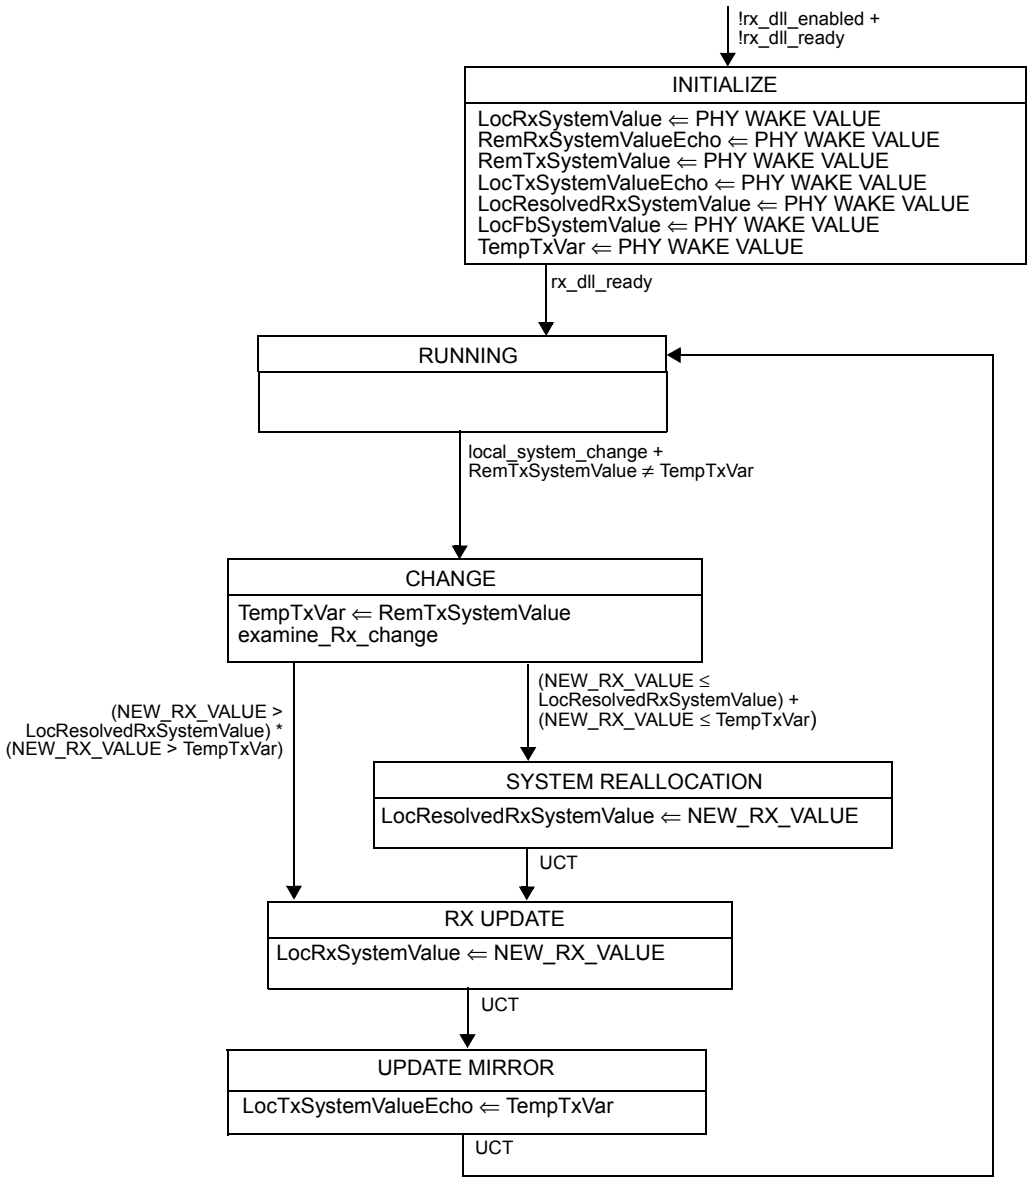

Figure 78–7—EEE DLL Receiver state diagram

The general state change procedure for transmitter fast wake is shown in Figure 78–8.

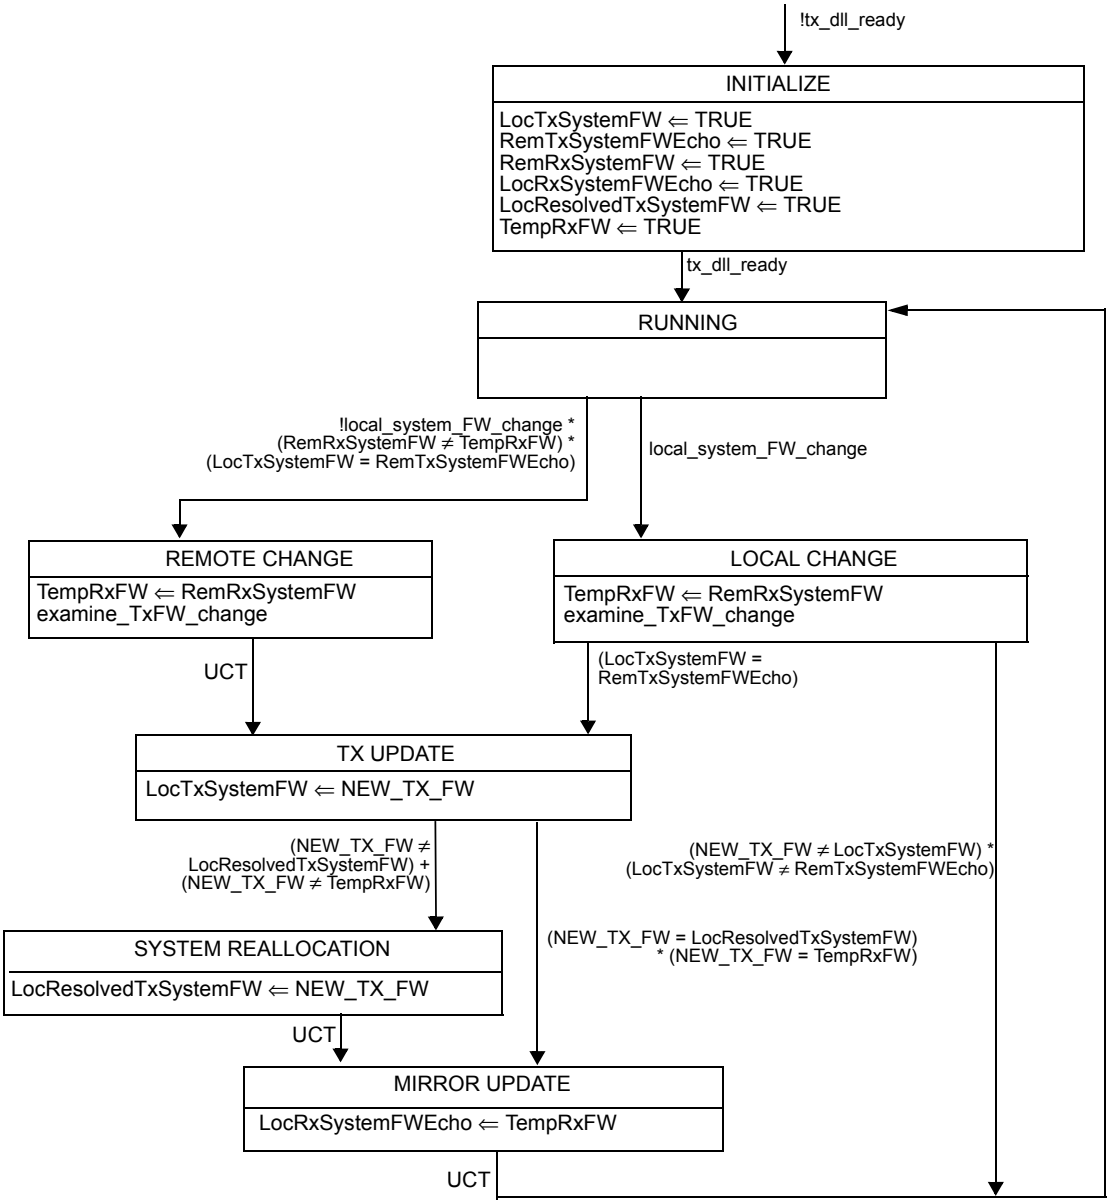

Figure 78–8—EEE DLL Transmitter fast wake state diagram

The general state change procedure for receiver fast wake is shown in Figure 78–9.

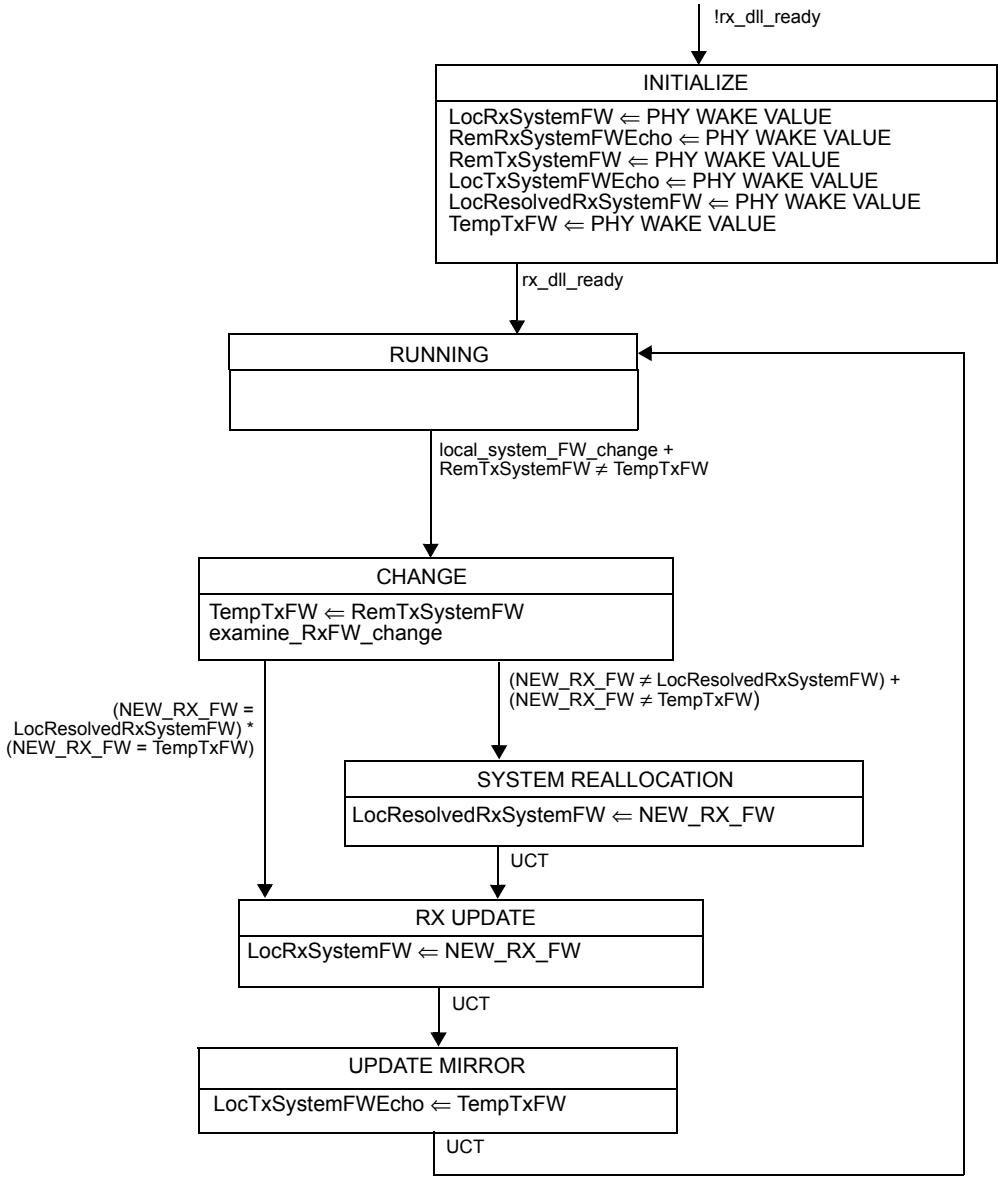

Figure 78–9—EEE DLL Receiver fast wake state diagram

78.4.3 State change procedure across a link

The transmitting and receiving link partners utilize the LLDP mechanism to advertise their various attributes to the other entity.

The initial $T_{w\_sys\_tx}$ defaults governing the EEE operation of the link default to the wake values required by the PHYs. This provides for EEE operation and functionality on initialization and prior to the exchange and processing of the TLVs.

The receiving link partner may request a new $T_{w\_sys\_tx}$ value through the aLldpXdot3LocRxTwSys (30.12.2.1.24) attribute in the LldpXdot3LocSystemsGroup managed object class (30.12.2). The request appears to the transmitting link partner as a change to the aLldpXdot3RemRxTwSys (30.12.3.1.21) attribute in the LldpXdot3RemSystemsGroup managed (30.12.3) object class. The transmitting link partner responds to its receiving partner's request through the aLldpXdot3LocTxTwSys (30.12.2.1.22) attribute in the LldpXdot3LocSystemsGroup managed object class (30.12.2). The transmitting link partner also copies the value of the aLldpXdot3RemRxTwSys (30.12.3.1.21) attribute in the LldpXdot3RemSystemsGroup managed (30.12.3) object class to the aLldpXdot3LocRxTwSysEcho (30.12.2.1.25) attribute in the LldpXdot3LocSystemsGroup managed object class (30.12.2).

The transmitting link partner may advertise new value of $T_{w\_sys\_tx}$ through the aLldpXdot3LocTxTwSys (30.12.2.1.22) attribute in the LldpXdot3LocSystemsGroup managed object class (30.12.2). This appears to the receiving link partner as a change to the aLldpXdot3RemTxTwSys (30.12.3.1.19) attribute in the LldpXdot3RemSystemsGroup managed (30.12.3) object class. The receiving link partner responds to a transmitter's request through the aLldpXdot3LocRxTwSys (30.12.2.1.24) attribute in the LldpXdot3LocSystemsGroup managed object class (30.12.2). The receiving link partner also copies the value of the aLldpXdot3RemTxTwSys (30.12.3.1.19) attribute in the LldpXdot3RemSystemsGroup managed (30.12.3) object class to the aLldpXdot3LocTxTwSysEcho (30.12.2.1.23) attribute in the LldpXdot3LocSystemsGroup managed object class (30.12.2). This appears to the transmitting link partner as a change to the aLldpXdot3RemTxTwSysEcho (30.12.3.1.20) attribute in the LldpXdot3RemSystemsGroup managed (30.12.3).

The state diagrams in Figure 78–6 and Figure 78–7 describe the preceding behavior.

The default state of Fast_Wake_Enable is TRUE for all PHYs that support the function. This provides for EEE operation and functionality on initialization and prior to the exchange and processing of the TLVs.

The receiving link partner may request a change of Fast_Wake_Enable through the aLldpXdot3LocRxFw (30.12.2.1.32) attribute in the LldpXdot3LocSystemsGroup managed object class (30.12.2). The request appears to the transmitting link partner as a change to the aLldpXdot3RemRxFw (30.12.3.1.26) attribute in the LldpXdot3RemSystemsGroup managed (30.12.3) object class. The transmitting link partner responds to its receiving partner's request through the aLldpXdot3LocTxFw (30.12.2.1.30) attribute in the LldpXdot3LocSystemsGroup managed object class (30.12.2). The transmitting link partner also copies the value of the aLldpXdot3RemRxFw (30.12.3.1.26) attribute in the LldpXdot3RemSystemsGroup managed (30.12.3) object class to the aLldpXdot3LocRxFwEcho (30.12.2.1.33) attribute in the LldpXdot3LocSystemsGroup managed object class (30.12.2).

The transmitting link partner may advertise a change of Fast_Wake_Enable through the aLldpXdot3LocTxFw (30.12.2.1.30) attribute in the LldpXdot3LocSystemsGroup managed object class (30.12.2). This appears to the receiving link partner as a change to the aLldpXdot3RemTxFw (30.12.3.1.24) attribute in the LldpXdot3RemSystemsGroup managed (30.12.3) object class. The receiving link partner responds to a transmitter's request through the aLldpXdot3LocRxFw (30.12.2.1.32) attribute in the LldpXdot3LocSystemsGroup managed object class (30.12.2). The receiving link partner also copies the value of the aLldpXdot3RemTxFw (30.12.3.1.24) attribute in the LldpXdot3RemSystemsGroup managed (30.12.3) object class to the aLldpXdot3LocTxFwEcho (30.12.2.1.31) attribute in the LldpXdot3LocSystemsGroup managed object class (30.12.2). This appears to the transmitting link partner as a change to the aLldpXdot3RemTxFwEcho (30.12.3.1.25) attribute in the LldpXdot3RemSystemsGroup managed (30.12.3).

The state diagrams in Figure 78–8 and Figure 78–9 describe the preceding behavior.

78.4.3.1 Transmitting link partner's state change procedure across a link

A transmitting link partner is said to be in sync with the receiving link partner if the presently advertised value of Transmit $T_{w\_sys\_tx}$ and the corresponding echoed value are equal.

During normal operation, the transmitting link partner is in the RUNNING state. If the transmitting link partner wants to initiate a change to the presently resolved value of $T_{w\_sys\_tx}$, the local_system_change is asserted and the transmitting link partner enters the LOCAL CHANGE state where NEW_TX_VALUE is computed. If the new value is smaller than the presently advertised value of $T_{w\_sys\_tx}$ or if the transmitting link partner is in sync with the receiving link partner, then it enters TX UPDATE state. Otherwise, it returns to the RUNNING state.

If the transmitting link partner sees a change in the $T_{w\_sys\_tx}$ requested by the receiving link partner, it recognizes the request only if it is in sync with the transmitting link partner. The transmitting link partner examines the request by entering the REMOTE CHANGE state where a NEW TX VALUE is computed and it then enters the TX UPDATE state.

Upon entering the TX UPDATE state, the transmitter updates the advertised value of Transmit $T_{w\_sys\_tx}$ with NEW_TX_VALUE. If the NEW_TX_VALUE is equal to or greater than either the resolved $T_{w\_sys\_tx}$ value or the value requested by the receiving link partner then it enters the SYSTEM REALLOCATION state where it updates the value of resolved $T_{w\_sys\_tx}$ with NEW_TX_VALUE. The transmitting link partner enters the MIRROR UPDATE state either from the SYSTEM REALLOCATION state or directly from the TX UPDATE state. The UPDATE MIRROR state then updates the echo for the Receive $T_{w\_sys\_tx}$ and returns to the RUNNING state.

A transmitting link partner is said to be in sync with the receiving link partner if the presently advertised value of Transmit Fast_Wake_Enable and the corresponding echoed value are equal.

During normal operation, the transmitting link partner is in the RUNNING state. If the transmitting link partner wants to initiate a change to the presently resolved value of Fast_Wake_Enable, the local_system_change is asserted and the transmitting link partner enters the LOCAL CHANGE state where NEW_TX_FW is computed. If the transmitting link partner is in sync with the receiving link partner, then it enters TX UPDATE state. Otherwise, it returns to the RUNNING state.

If the transmitting link partner sees a change in the Fast_Wake_Enable requested by the receiving link partner, it recognizes the request only if it is in sync with the transmitting link partner. The transmitting link partner examines the request by entering the REMOTE CHANGE state where a NEW_TX_FW is computed and it then enters the TX UPDATE state.

Upon entering the TX UPDATE state, the transmitter updates the advertised value of Transmit Fast_Wake_Enable with NEW_TX_FW. If the NEW_TX_FW is different to either the resolved Fast_Wake_Enable value or the value requested by the receiving link partner then it enters the SYSTEM REALLOCATION state where it updates the value of resolved Fast_Wake_Enable with NEW_TX_FW. The transmitting link partner enters the MIRROR UPDATE state either from the SYSTEM REALLOCATION state or directly from the TX UPDATE state. The UPDATE MIRROR state then updates the echo for the Receive Fast_Wake_Enable and returns to the RUNNING state.

78.4.3.2 Receiving link partner's state change procedure across a link

A receiving link partner is said to be in sync with the transmitting link partner if the presently requested value of Receive $T_{w\_sys\_tx}$ and the corresponding echoed value are equal.

During normal operation, the receiving link partner is in the RUNNING state. If the receiving link partner wants to request a change to the presently resolved value of $T_{w\_sys\_tx}$, the local_system_change is asserted. When local_system_change is asserted or when the receiving link partner sees a change in the $T_{w\_sys\_tx}$ advertised by the transmitting link partner, it enters the CHANGE state where NEW_RX_VALUE is computed. If NEW_RX_VALUE is less than either the presently resolved value of $T_{w\_sys\_tx}$ or the presently advertised value by the transmitting link partner, it enters the SYSTEM REALLOCATION state where it updates the resolved value of $T_{w\_sys\_tx}$ to NEW_RX_VALUE. The receiving link partner ultimately enters

the RX UPDATE state, either from the SYSTEM REALLOCATION state or directly from the CHANGE state.

In the RX UPDATE state, it updates the presently requested value to NEW_RX_VALUE, then it updates the echo for the Transmit $T_{w\_sys\_tx}$ in the UPDATE MIRROR state and finally goes back to the RUNNING state.

A receiving link partner is said to be in sync with the transmitting link partner if the presently requested value of Receive Fast_Wake_Enable and the corresponding echoed value are equal.

During normal operation, the receiving link partner is in the RUNNING state. If the receiving link partner wants to request a change to the presently resolved value of Fast_Wake_Enable, the local_system_change is asserted. When local_system_change is asserted or when the receiving link partner sees a change in the Fast_Wake_Enable advertised by the transmitting link partner, it enters the CHANGE state where NEW_RX_FW is computed. If NEW_RX_FW is different from either the presently resolved value of Fast_Wake_Enable or the presently advertised value by the transmitting link partner, it enters the SYSTEM REALLOCATION state where it updates the resolved value of Fast_Wake_Enable to NEW_RX_FW. The receiving link partner ultimately enters the RX UPDATE state, either from the SYSTEM REALLOCATION state or directly from the CHANGE state.

In the RX UPDATE state, it updates the presently requested value to NEW_RX_FW, then it updates the echo for the Transmit Fast_Wake_Enable in the UPDATE MIRROR state and finally goes back to the RUNNING state.

78.5 Communication link access latency

In the full duplex mode, predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delay 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 the concatenation of devices.

The EEE capability adds latency that has to be considered by the network designer. When in the LPI mode, the PHY link is not available immediately for transmission of data. The system has to wake it up by sending the normal IDLE code on the MAC interface. Following the reception of an IDLE code on the MAC interface, the PHY starts the wake-up process. The maximal PHY recovery time, T_{w_phy} is defined for each PHY.

Transmit and/or Receive wait time shrinkage can happen when $T_{phy\_shrink\_rx}$ or $T_{phy\_shrink\_tx}$ (as defined in 78.1.3) are not zero. This has to be taken into consideration in designing or configuring the network.

Table 78–4 summarizes critical timing parameters for supported PHYs. These are listed here to assist the system designer in assessing the impact of EEE on the operation of the link.

Case-1 of the 1000BASE-T PHY applies to PHYs in Master mode. Case-2 of the 1000BASE-T PHY applies to PHYs in Slave mode.

Case-1 of the 10GBASE-KR PHY applies to PHYs without FEC. Case-2 of the 10GBASE-KR PHY applies to PHYs with FEC.

Case-1 of the 10GBASE-T PHY applies when the PHY is requested to transmit the Wake signal before transmission of the Sleep signal to the Link Partner is complete. Case-2 of the 10GBASE-T PHY applies when the PHY is requested to transmit the Wake signal after transmission of the Sleep signal to the Link Partner is complete and if the PHY has not indicated LOCAL FAULT at any time during the previous 10 ms.

Case-1 of the 40GBASE-CR4, 40GBASE-KR4, and 100GBASE-CR10 PHYs applies to PHYs without FEC in deep sleep. Case-2 of these PHYs applies to PHYs with FEC in deep sleep.

Table 78–4—Summary of the LPI timing parameters for supported PHYs or interfaces

PHY or interface type	Case	$T_{w\_sys\_tx}$ (min) (μs)	T_{w_phy} (min) (μs)	$T_{phy\_shrink\_tx}$ (max) (μs)	$T_{phy\_shrink\_rx}$ (max) (μs)	$T_{w\_sys\_rx}$ (min) (μs)
100BASE-TX		30	20.5	5.0	15	10
1000BASE-T	Case-1	16.5	16.5	5.0	2.5	1.76
	Case-2	16.5	16.5	12.24	9.74	1.76
1000BASE-KX		13.26	11.25	5.0	6.5	1.76
XGXS (XAUI)		12.38	9.25	5.0	4.5	2.88
10GBASE-T	Case-1	7.36	7.36	4.48	0	2.88
	Case-2	4.48	4.48	1.6	0	2.88
10GBASE-KX4		12.38	9.25	5.0	4.5	2.88
10GBASE-KR	Case-1	15.38	12.25	5.0	7.5	2.88
	Case-2	17.38	14.25	5.0	9.5	2.88
40GBASE-R fast wake		0.34	0.3	0	0	0.25
40GBASE-CR4	Case-1	5.5	5.5	2	3	1.2
	Case-2	6.5	6.5	2	3	1.2
40GBASE-KR4	Case-1	5.5	5.5	2	3	1.2
	Case-2	6.5	6.5	2	3	1.2
100GBASE-R fast wake		0.34	0.3	0	0	0.25
100GBASE-CR10	Case-1	5.5	5.5	2	3	1
	Case-2	7.5	7.5	2	3	1
100GBASE-CR4		5.5	5.5	2	3	1
100GBASE-KR4		5.5	5.5	2	3	1
100GBASE-KP4		5.5	5.5	2	3	1
XLAUI/CAUI-n[a]		1				

[a] $T_{w\_sys\_tx}$ is increased by 1 μs for each instance of XLAUI/CAUI with shutdown enabled on the transmit path.The receiver should negotiate an increase for remote T_{w_sys} for the link partner of 1 μs for each instance of XLAUI/CAUI with shutdown enabled on the receive path.

78.5.1 10 Gb/s PHY extension using XGXS

The XGXS can be inserted between the RS and a 10 Gb/s PHY to transparently extend the physical reach of the XGMII. The LPI signaling can operate through the XGXS with no change to the PHY timing parameters described in Table 78–4 or the operation of the Data Link Layer Capabilities negotiation described in 78.4.

If the DTE XS XAUI stop enable bit (5.0.9) is asserted, the DTE XS may stop signaling on the XAUI in the transmit direction to conserve energy. If the DTE XS XAUI stop enable bit is asserted, the RS defers sending data following deassertion of LPI by an additional time equal to $T_{w\_sys\_tx} - T_{w\_sys\_rx}$ for the XGXS as shown in Table 78–4 (see 46.4.2.1).

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 negotiates an additional time 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.

78.5.2 40 Gb/s and 100 Gb/s PHY extension using XLAUI or CAUI-n

40 Gb/s PHYs may be extended using XLAUI and 100 Gb/s PHYs may be extended using CAUI-10 or CAUI-4 as a physical instantiation of the inter-sublayer service interface to separate functions between devices. The LPI signaling can operate across XLAUI/CAUI-n with no change to the PHY timing parameters described in Table 78–4 or the operation of the Data Link Layer Capabilities negotiation described in 78.4.

If PMA Egress AUI Stop Enable (PEASE, see 83.3; MDIO register bit 1.7.8) is asserted for any of the PMA sublayers, the PMA may stop signaling on the XLAUI/CAUI-n in the transmit direction to conserve energy. If PEASE is asserted, the RS defers sending data following deassertion of LPI by an additional time equal to $T_{w\_sys\_tx}-T_{w\_sys\_rx}$ as shown in Table 78–4 for each PMA with PEASE asserted (see 81.4.2).

If PMA Ingress AUI Stop Enable (PIASE, see 83.3; MDIO register bit 1.7.9) is asserted for any of the PMA sublayers, the PMA may stop signaling on the XLAUI/CAUI-n in the receive direction to conserve energy. The receiver should negotiate an additional time for the remote T_{w_sys} (equal to $T_{w\_sys\_tx}-T_{w\_sys\_rx}$ for the XLAUI/CAUI-n as shown in Table 78–4) for each PMA with PIASE to be asserted before setting the PIASE bits.

78.6 Protocol implementation conformance statement (PICS) proforma for EEE Data Link Layer Capabilities[1]

78.6.1 Introduction

The supplier of a protocol implementation that is claimed to conform to 78.4 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.

78.6.2 Identification

78.6.2.1 Implementation identification

Supplier[a]	
Contact point for inquiries about the PICS[a]	
Implementation Name(s) and Version(s)[a, c]	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)[b]	
[a]Required for all implementations [b]May be completed as appropriate in meeting the requirements for the identification. [c]The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

78.6.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, 78.4, EEE Data Link Layer Capabilities
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.

78.6.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
10G	Support 10G or higher operation	78.4	Support for 10 Gb/s or higher operation	O	Yes [ ] No [ ]
DLL1	DLL	78.4	DLL	10G:M	Yes [ ] N/A [ ]
DLL2	DLL	78.4	DLL	!10G:O	Yes [ ] No [ ]

In addition, the following predicate name is defined for use when different implementations from the preceding set have common parameters:

DLL = DLL1 OR DLL 2

78.6.4 DLL requirements

Item	Feature	Subclause	Value/Comment	Status	Support
DLR1	DLL Timing	78.4.1	Timing requirements	DLL:M	Yes [ ] N/A [ ]
DLR2	DLL Control state diagrams	78.4.2	State machines for TX and RX	DLL:M	Yes [ ] N/A [ ]

79. IEEE 802.3 Organizationally Specific Link Layer Discovery Protocol (LLDP) type, length, and value (TLV) information elements

79.1 Overview

The Link Layer Discovery Protocol (LLDP) specified in IEEE Std 802.1AB-2009 is a MAC Client protocol that allows stations attached to an IEEE 802 LAN to advertise to all other stations attached to the same IEEE 802 LAN: the major capabilities provided by the system incorporating that station, the management address or addresses of the entity or entities that provide management of those capabilities, and the identification of the station's point of attachment to the IEEE 802 LAN.

The information fields in each LLDP frame are contained in a Link Layer Discovery Protocol Data Unit (LLDPDU) as a sequence of short, variable length, information elements known as TLVs that each include a type, length, and value field.

Organizationally Specific TLVs can be defined by either the professional organizations or the individual vendors that are involved with the particular functionality being implemented within a system. The basic format and procedures for defining Organizationally Specific TLVs are provided in subclause 9.6 of IEEE Std 802.1AB-2009.

79.1.1 IEEE 802.3 LLDP frame format

The IEEE 802.3 LLDP frame format is illustrated in Figure 79–1.

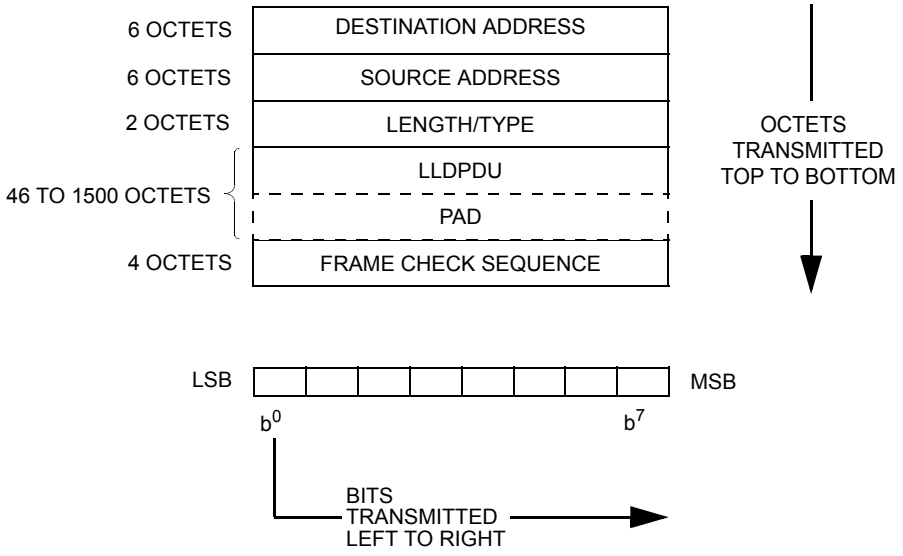

Figure 79–1—IEEE 802.3 LLDP frame format

NOTE—The illustration shows the simplest form of an IEEE 802.3 LLDP frame; i.e., where the frame has no IEEE Std 802.1Q™ tag header, or IEEE Std 802.1AE™ security tag, or any other form of encapsulation applied to it.[2]

[2]NOTES in text, tables, and figures are given for information only and do not contain requirements needed to implement the standard.

79.1.1.1 Destination Address field

The Destination Address field of an IEEE 802.3 LLDP frame contains a MAC address specified by 7.1 of IEEE Std 802.1AB-2009 (see 79.2).

79.1.1.2 Source Address field

The Source Address field of an IEEE 802.3 LLDP frame contains the 48-bit individual address of the station sending the frame.

79.1.1.3 Length/Type field

The Length/Type field of an IEEE 802.3 LLDP frame is a 2-octet field that contains the hexadecimal value: 88-CC. This value carries the Type interpretation (see 3.2.6), and has been universally assigned for LLDP.

79.1.1.4 LLDPDU field

The LLDPDU field of an IEEE 802.3 LLDP frame contains the LLDPDU which is a sequence of short, variable length, information elements known as TLVs that each include type, length, and value fields.

79.1.1.5 Pad field

A minimum MAC frame size is required for correct operation and, if necessary, a Pad field is appended after the LLDPDU field as defined in 3.2.8.

79.1.1.6 Frame Check Sequence field

The Frame Check Sequence (FCS) field contains the Frame Check Sequence, as defined in 3.2.9.

79.2 Requirements of the IEEE 802.3 Organizationally Specific TLV set

All IEEE 802.3 Organizationally Specific TLVs shall conform to the LLDPDU bit and octet ordering conventions of 8.1 of IEEE Std 802.1AB-2009.

79.3 IEEE 802.3 Organizationally Specific TLVs

The currently defined IEEE 802.3 Organizationally Specific TLVs are listed in Table 79–1. Any additions or changes to these TLVs will be included in this clause.

Table 79–1—IEEE 802.3 Organizationally Specific TLVs

IEEE 802.3 subtype	TLV name	Subclause reference
1	MAC/PHY Configuration/Status	79.3.1
2	Power Via Medium Dependent Interface (MDI)	79.3.2
3	Link Aggregation (deprecated)	79.3.3
4	Maximum Frame Size	79.3.4
5	Energy-Efficient Ethernet	79.3.5
6	EEE fast wake	79.3.6
7 to 255	Reserved	—

79.3.1 MAC/PHY Configuration/Status TLV

The MAC/PHY Configuration/Status TLV is an optional TLV that identifies the following:

a) The duplex and bit-rate capability of the sending IEEE 802.3 LAN node that is connected to the physical medium.
b) The current duplex and bit-rate settings of the sending IEEE 802.3 LAN node.
c) Whether the current duplex and bit-rate settings are the result of auto-negotiation during link initiation or of manual set override action.

Figure 79–2 shows the format of this TLV.

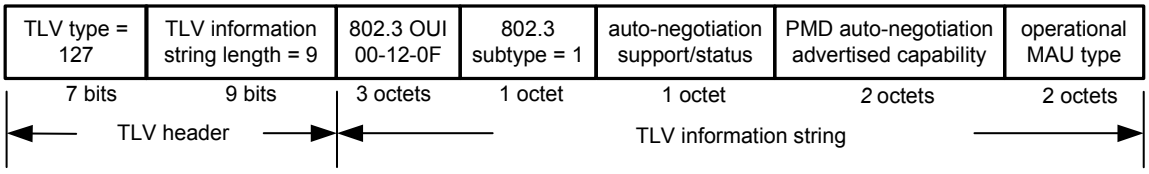

Figure 79–2—MAC/PHY configuration/status TLV format

79.3.1.1 Auto-negotiation support/status

The auto-negotiation support/status field shall contain a bitmap that identifies the auto-negotiation support and current status of the local IEEE 802.3 LAN station as defined in Table 79–2. If the auto-negotiation support bit (bit 0) is one and the auto-negotiation status bit (bit 1) is zero, the IEEE 802.3 physical media dependent sublayer (PMD) operating mode is determined by the operational Medium Attachment Unit (MAU) type field value rather than by auto-negotiation.

79.3.1.2 PMD auto-negotiation advertised capability field

The ‘PMD auto-negotiation capability’ field shall contain a 2-octet value that provides a bitmap of the ifMauAutoNegCapAdvertisedBits object, defined in IETF RFC 4836, of the sending device. Bit zero is the high order (left-most) bit in an octet string.

Table 79–2—IEEE 802.3 auto-negotiation support/status

Bit	Function	Value/meaning	IETF RFC 4836 reference
0	Auto-negotiation support	1 = supported 0 = not supported	ifMauAutoNegSupported
1	Auto-negotiation status	1 = enabled 0 = not enabled	ifMauAutoNegAdminStatus
7:2	—	Reserved for future standardization	—

79.3.1.3 Operational MAU type

The operational MAU type field contains an integer value indicating the MAU type of the sending device. This value shall be derived from the list position of the corresponding dot3MauType as listed in IETF RFC 4836 (or subsequent revisions) and is equal to the last number in the respective dot3MauType Object Identifier (OID). For example, if the ifMauType object is dot3MauType1000BaseTHD which corresponds to {dot3MauType 29}, the numerical value of this field is 29. For MAU types not listed in IETF RFC 4836 (or subsequent revisions), the value of this field shall be set to zero.

79.3.1.4 MAC/PHY Configuration/Status TLV usage rules

An LLDPDU should contain no more than one MAC/PHY Configuration/Status TLV.

79.3.2 Power Via MDI TLV

Clause 33 defines two option power entities: a Powered Device (PD) and Power Sourcing Equipment (PSE). These entities allow devices to draw/supply power over the sample generic cabling as used for data transmission. The Power Via MDI TLV allows network management to advertise and discover the MDI power support capabilities of the sending IEEE 802.3 LAN station. This TLV is also required to perform Data Link Layer classification as defined in 33.6. Figure 79–3 shows the format of this TLV.

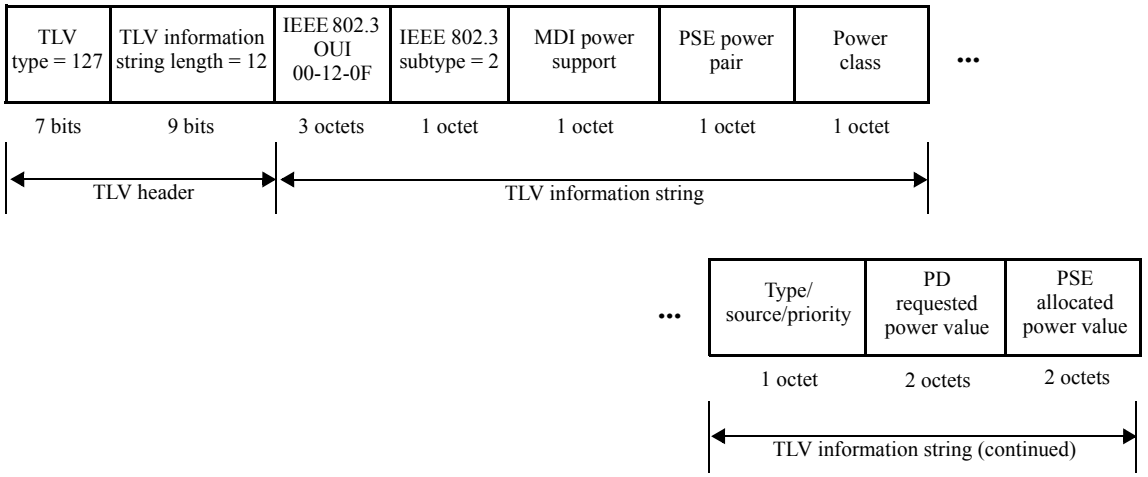

Figure 79–3—Power Via MDI TLV format

The TLV shown in Figure 79–3 is a revision of the legacy Power via MDI TLV originally defined in IEEE Std 802.1AB-2009 Annex F.3. The legacy TLV had only the first three fields of the TLV shown in the

figure. These three fields enable discovery and advertisement of MDI power support capabilities. The newly added fields provide Data Link Layer classification capabilities. The revised TLV can be used by the PSE only when it is supplying power to a PI encompassed within an MDI and by the PD only when it is drawing power from the PI. Power entities may continue to use the legacy TLV prior to supplying/drawing power to/from the PI. If the power entity implements Data Link Layer classification, it shall use the Power via MDI TLV shown in Figure 79–3 after the PI has been powered.

79.3.2.1 MDI power support

The MDI power support field shall contain a bitmap of the MDI power capabilities and status as defined in Table 79–3.

Table 79–3—MDI power capabilities/status

Bit	Function	Value/meaning	IETF RFC 3621 object reference
0	Port class	1 = PSE 0 = PD	See Note 1
1	Power Sourcing Equipment (PSE) MDI power support	1 = supported 0 = not supported	See Note 2 and Note 3
2	PSE MDI power state	1 = enabled 0 = disabled	pethPsePortAdminEnable
3	PSE pairs control ability	1 = pair selection can be controlled 0 = pair selection can not be controlled	pethPsePortPowerPairContolAbility
7:4	Reserved for future standardization	—	—

NOTE 1—Port class information is implied by the support of the PSE or PD groups.

NOTE 2—MDI power support information is implied by support of IETF RFC 3621.

NOTE 3—If bit 1 is zero, bit 2 has no meaning.

79.3.2.2 PSE power pair

The PSE power pair field shall contain an integer value as defined by the pethPsePortPowerPairs object in IETF RFC 3621.

79.3.2.3 Power class

The power class field shall contain an integer value as defined by the pethPsePortPowerClassifications object in IETF RFC 3621.

79.3.2.4 Requested power type/source/priority

The power type/source/priority field shall contain a bit-map of the power type, source and priority defined in Table 79–4 and is reported for the device generating the TLV.

79.3.2.4.1 Power type

This field shall be set according to Table 79–4.

Table 79–4—Power type/source/priority field

Bit	Function	Value/meaning
7:6	power type	7 6 1 1 = Type 1 PD 1 0 = Type 1 PSE 0 1 = Type 2 PD 0 0 = Type 2 PSE
5:4	power source	Where power type = PD 5 4 1 1 = PSE and local 1 0 = Reserved 0 1 = PSE 0 0 = Unknown Where power type = PSE 5 4 1 1 = Reserved 1 0 = Backup source 0 1 = Primary power source 0 0 = Unknown
3:2	Reserved	Transmit as zero, ignore on receive
1:0	power priority	1 0 1 1 = low 1 0 = high 0 1 = critical 0 0 = unknown (default)

79.3.2.4.2 Power source

When the power type is PD, this field shall be set to 01 when the PD is being powered only through the PI; to 11 when the PD is being powered from both; and to 00 when this information is not available.

When the power type is PSE, this field shall be set to 01 when the PSE is sourcing its power through the PI from its primary supply; to 10 when the PSE is sourcing its power through the PI from a backup source; and to 00 when this information is not available.

79.3.2.4.3 Power priority

When the power type is PD, this field shall be set to the power priority configured for the device. If a PD is unable to determine its power priority or it has not been configured, then this field shall be set to 00.

When the power type is PSE, this field reflects the PD priority that the PSE advertises to assign to the PD.

79.3.2.5 PD requested power value

The PD requested power value field shall contain the PD's requested power value defined in Table 79–5.

Table 79–5—PD requested power value field

Bit	Function	Value/meaning
15:0	PD requested power value	Power = 0.1 × (decimal value of bits) Watts. Valid values for these bits are decimal 1 through 255.

The PD requested power value is encoded according to Equation (79–1).

$$Power = \{0.1 \times X\}_{\mathrm{W}} \tag{79–1}$$

where

Power is the effective requested PD power value
X is the decimal value of the power value field, bits 15:0

"PD requested power value" is the maximum input average power (see 33.3.7.2) the PD wants to draw. "PD requested power value" is the power value at the input to the PD's PI.

79.3.2.6 PSE allocated power value

The PSE allocated power value field shall contain the PSE's allocated power value defined in Table 79–6.

Table 79–6—PSE allocated power value field

Bit	Function	Value/meaning
15:0	PSE allocated power value	Power = 0.1 × (decimal value of bits) Watts. Valid values for these bits are decimal 1 through 255.

The PSE allocated power value is encoded according to Equation (79–2).

$$Power = \{0.1 \times X\}_{\mathrm{W}} \tag{79–2}$$

where

Power is the effective allocated PSE power value
X is the decimal value of the power value field, bits 15:0

"PSE allocated power value" is the maximum input average power (see 33.3.7.2) the PSE expects the PD to draw. "PSE allocated power value" is the power at the input to the PD's PI. The PSE uses this value to compute P_{Class} as defined in 33.2.6.

79.3.2.7 Power Via MDI TLV usage rules

An LLDPDU should contain no more than one Power Via MDI TLV.

79.3.3 Link Aggregation TLV (deprecated)

The Link Aggregation TLV is an optional TLV that indicates whether the link is capable of being aggregated, whether the link is currently in an aggregation, and if in an aggregation, the port identification of the aggregation. Figure 79–4 shows the format for this TLV.

NOTE—As the Link Aggregation specification has now been removed from IEEE Std 802.3 and is now standardized as IEEE Std 802.1AX, new implementations of this standard are encouraged to make use of the Link Aggregation TLV that is now part of the IEEE 802.1 extension MIB specified in Annex E of IEEE Std 802.1AB-2009.

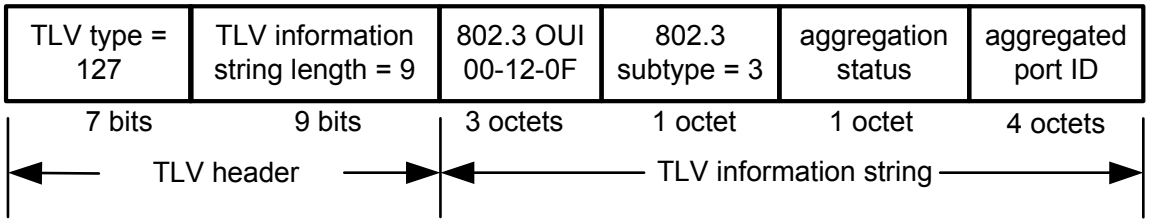

Figure 79–4—Link Aggregation TLV format

79.3.3.1 Aggregation status

The link aggregation status field shall contain a bitmap of the link aggregation capabilities and the current aggregation status of the link as defined in Table 79–7.

Table 79–7—Link aggregation capability/status

Bit	Function	Value/meaning
0	Aggregation capability	0 = not capable of being aggregated 1 = capable of being aggregated
1	Aggregation status	0 = not currently in aggregation 1 = currently in aggregation
7:2	reserved for future standardization	—

79.3.3.2 Aggregated port ID

The aggregated port ID field shall contain the IEEE 802.3 aggregated port identifier, aAggPortID, derived from the ifNumber in the ifIndex for the interface (see 30.7.2.1.1).

79.3.3.3 Link Aggregation TLV usage rules

An LLDPDU should contain no more than one Link Aggregation TLV.

79.3.4 Maximum Frame Size TLV

The Maximum Frame Size TLV is an optional TLV that indicates the maximum frame size capability of the implemented MAC and PHY. Figure 79–5 shows the format of this TLV.

NOTE—MAC and PHY support for a given frame size doesn't necessarily mean that the upper layers support that frame size.

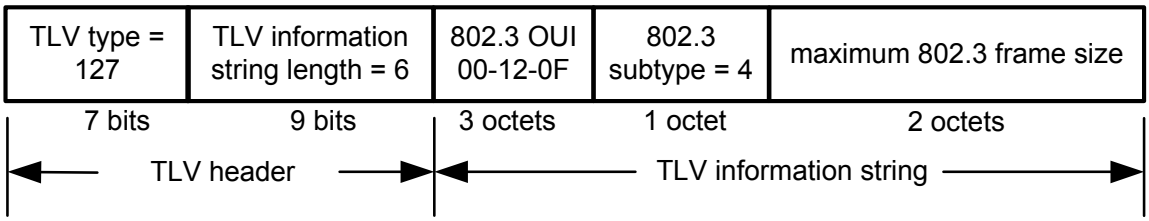

Figure 79–5—Maximum Frame Size TLV format

79.3.4.1 Maximum frame size

The maximum frame size field shall contain an integer value indicating the maximum supported frame size in octets as determined by the following:

a) If the MAC/PHY supports only basic frames (see 3.2.7) the maximum frame size field shall be set to 1518.
b) If the MAC/PHY supports Q-tagged frames (see 3.2.7) the maximum frame size field shall be set to 1522.
c) If the MAC/PHY supports envelope frames (see 3.2.7) the maximum frame size field shall be set to 2000.

79.3.4.2 Maximum Frame Size TLV usage rules

An LLDPDU should contain no more than one Maximum Frame Size TLV.

79.3.5 EEE TLV

The EEE TLV is used to exchange information about the EEE Data Link Layer capabilities. Figure 79–6 shows the format of this TLV.

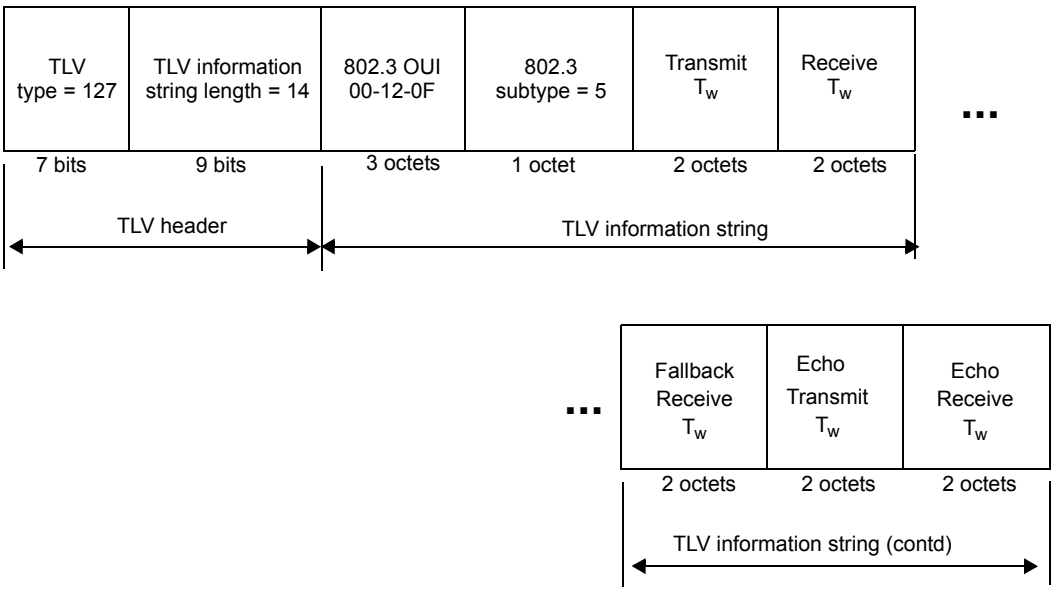

Figure 79–6—EEE TLV format

79.3.5.1 Transmit T_w

Transmit $T_{w\_sys\_tx}$ (2 octets wide) shall be defined as the time (expressed in microseconds) that the transmitting link partner will wait before it starts transmitting data after leaving the Low Power Idle (LPI) mode. This is a function of the transmit system design and may be constrained, for example, by the transmit path buffering. The default value for Transmit $T_{w\_sys\_tx}$ is the T_{w_phy} defined for the PHY that is in use for the link. The Transmitting link partner expects that the Receiving link partner will be able to accept data after the time delay Transmit $T_{w\_sys\_tx}$ (expressed in microseconds).

79.3.5.2 Receive T_w

Receive $T_{w\_sys\_tx}$ (2 octets wide) shall be defined as the time (expressed in microseconds) that the receiving link partner is requesting the transmitting link partner to wait before starting the transmission data following the LPI. The default value for Receive $T_{w\_sys\_tx}$ is the T_{w_phy} defined for the PHY that is in use for the link. The Receive $T_{w\_sys\_tx}$ value can be larger but not smaller than the default. The extra wait time may be used by the receive link partner for power-saving mechanisms that require a longer wake-up time than the PHY-layer definitions.

79.3.5.3 Fallback T_w

A receiving link partner may inform the transmitter of an alternate desired $T_{w\_sys\_tx}$. Since a receiving link partner is likely to have discrete levels for savings, this provides the transmitter with additional information that it may use for a more efficient allocation. As with the Receive $T_{w\_sys\_tx}$, this is 2 octets wide. Systems that do not implement this option default the value to be the same as that of the Receive $T_{w\_sys\_tx}$.

79.3.5.4 Echo Transmit and Receive T_W

The respective echo values shall be defined as the local link partner's reflection (echo) of the remote link partner's respective values. When a local link partner receives its echoed values from the remote link partner it can determine whether or not the remote link partner has received, registered, and processed its most recent values. For example, if the local link partner receives echoed parameters that do not match the values in its local MIB, then the local link partner infers that the remote link partner's request was based on stale information.

79.3.5.5 EEE TLV usage rules

An LLDPDU should contain no more than one EEE TLV.

79.3.6 EEE Fast Wake TLV

The EEE Fast Wake TLV is used to exchange information about the EEE fast wake capabilities. This TLV is only used by systems with links operating at speeds greater than 10 Gb/s. Figure 79–7 shows the format of this TLV.

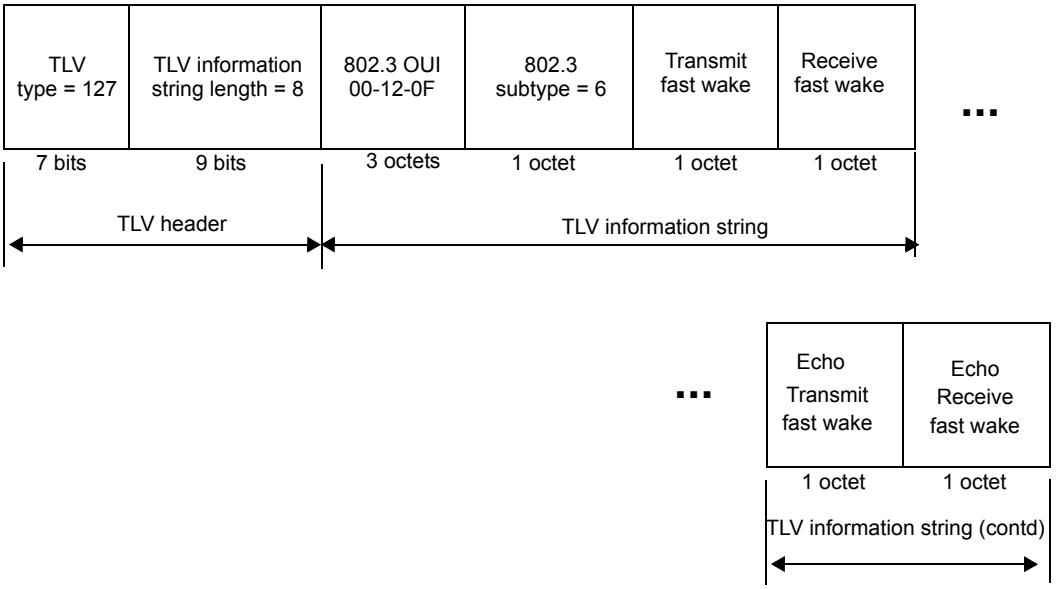

Figure 79–7—EEE Fast Wake TLV format

79.3.6.1 Transmit fast wake

Transmit fast wake (1 octet wide) is a logical indication that the transmit LPI state diagram intends to use the fast wake function (corresponding to the variable LPI_FW in 82.2.19.2.2). Transmit fast wake = 1 corresponds to LPI_FW being TRUE; Transmit fast wake = 0 corresponds to LPI_FW being FALSE. The default value for Transmit fast wake is 1 (TRUE). Transmit fast wake is set to TRUE unless the PHY is capable of deep sleep operation as determined by the PHY type and the results of auto-negotiation.

79.3.6.2 Receive fast wake

Receive fast wake (1 octet wide) is a logical indication that the receive LPI state diagram is expecting its link partner to use the fast wake function (corresponding to the variable LPI_FW in 82.2.19.2.2). Receive fast

wake = 1 corresponds to LPI_FW being TRUE; Receive fast wake = 0 corresponds to LPI_FW being FALSE. The default value for Receive fast wake is 1 (TRUE). Receive fast wake is set to TRUE unless the PHY is capable of deep sleep operation as determined by the PHY type and the results of auto-negotiation.

79.3.6.3 Echo of Transmit fast wake and Receive fast wake

The respective echo values are the local link partner's reflection (echo) of the remote link partner's respective values. When a local link partner receives its echoed values from the remote link partner, it can determine whether or not the remote link partner has received, registered, and processed its most recent values. For example, if the local link partner receives echoed parameters that do not match the values in its local MIB, then the local link partner infers that the remote link partner's request was based on stale information.

79.3.6.4 EEE Fast Wake TLV usage rules

An LLDPDU should contain no more than one EEE Fast Wake TLV.

79.4 IEEE 802.3 Organizationally Specific TLV selection management

TLV selection management consists of providing the network manager with the means to select which specific IEEE 802.3 Organizationally Specific TLVs are enabled for inclusion in an LLDPDU. The following LLDP variable cross references the LLDP local systems configuration MIB tables (see Clause 11 of IEEE Std 802.1AB-2009) to indicate which specific TLVs are enabled for the particular port(s) on the system. The specific port(s) through which each TLV is enabled for transmission may be set (or reset) by the network manager:

a) **mibXdot3TLVsTxEnable:** This variable lists the single-instance use IEEE 802.3 Organizationally Specific TLVs, each with a bitmap indicating the system ports through which the referenced TLV is enabled for transmission.

79.4.1 IEEE 802.3 Organizationally Specific TLV selection variable/LLDP Configuration managed object class cross reference

Table 79–8 lists the relationship both between IEEE 802.3 TLV selection variable and the corresponding LLDP Configuration managed object class (see 30.12.1) attribute.

Table 79–8—IEEE 802.3 Organizationally Specific TLV selection variable/LLDP MIB object cross reference

IEEE 802.3 TLV selection variable	LLDP Configuration managed object class attribute
mibXdot3TLVsTxEnable	aLldpXdot3PortConfigTLVsTxEnable

79.4.2 IEEE 802.3 Organizationally Specific TLV/LLDP Local and Remote System group managed object class cross references

The cross references between the IEEE 802.3 TLVs and the LLDP Local System Group managed object class (see 30.12.2) attributes are listed in Table 79–9. The cross references between the IEEE 802.3 TLVs and the LLDP Remote System Group managed object class (see 30.12.3) attributes are listed in Table 79–10.

The cross-references between the EEE TLV, the EEE Fast Wake TLV, and the EEE local (30.12.2) and remote (30.12.3) object class attributes are listed in Table 79–9 and Table 79–10.

Table 79–9—IEEE 802.3 Organizationally Specific TLV/LLDP Local System Group managed object class cross references

TLV name	TLV variable	LLDP Local System Group managed object class attribute
MAC/PHY Configuration/Status	Auto-negotiation support	aLldpXdot3LocPortAutoNegSupported
	Auto-negotiation status	aLldpXdot3LocPortAutoNegEnabled
	PMD auto-negotiation advertised capability	aLldpXdot3LocPortAutoNegAdvertisedCap
	Operational MAU type	aLldpXdot3LocPortOperMauType
Power via MDI	Port class	aLldpXdot3LocPowerPortClass
	PSE MDI power support	aLldpXdot3LocPowerMDISupported
	PSE MDI power state	aLldpXdot3LocPowerMDIEnabled
	PSE pairs control ability	aLldpXdot3LocPowerPairControlable
	PSE power pair	aLldpXdot3LocPowerPairs
	Power class	aLldpXdot3LocPowerClass
	Power type	aLldpXdot3LocPowerType
	Power source	aLldpXdot3LocPowerSource
	Power priority	aLldpXdot3LocPowerPriority
	PD requested power value	aLldpXdot3LocPDRequestedPowerValue
	PSE allocated power value	aLldpXdot3LocPSEAllocatedPowerValue
Link Aggregation (deprecated)	aggregation status	aLldpXdot3LocLinkAggStatus
	aggregated port ID	aLldpXdot3LocLinkAggPortId
Maximum Frame Size	maximum frame size	aLldpXdot3LocMaxFrameSize
EEE	Transmit $T_{w\_sys\_tx}$	aLldpXdot3LocTxTwSys
	Receive $T_{w\_sys\_tx}$	aLldpXdot3LocRxTwSys
	Echo Transmit $T_{w\_sys\_tx}$	aLldpXdot3LocTxTwSysEcho
	Echo Receive $T_{w\_sys\_tx}$	aLldpXdot3LocRxTwSysEcho
	Fallback $T_{w\_sys\_tx}$	aLldpXdot3LocFbTwSys
EEE Fast Wake	Transmit fast wake	aLldpXdot3LocTxFw
	Receive fast wake	aLldpXdot3LocRxFw
	Echo Transmit fast wake	aLldpXdot3LocTxFwEcho
	Echo Receive fast wake	aLldpXdot3LocRxFwEcho

Table 79–10—IEEE 802.3 Organizationally Specific TLV/LLDP Remote System Group managed object class cross references

TLV name	TLV variable	LLDP Remote System Group managed object class attribute
MAC/PHY Configuration/Status	Auto-negotiation support	aLldpXdot3RemPortAutoNegSupported
	Auto-negotiation status	aLldpXdot3RemPortAutoNegEnabled
	PMD auto-negotiation advertised capability	aLldpXdot3RemPortAutoNegAdvertisedCap
	Operational MAU type	aLldpXdot3RemPortOperMauType
Power via MDI	Port class	aLldpXdot3RemPowerPortClass
	PSE MDI power support	aLldpXdot3RemPowerMDISupported
	PSE MDI power state	aLldpXdot3RemPowerMDIEnabled
	PSE pairs control ability	aLldpXdot3RemPowerPairControlable
	PSE power pair	aLldpXdot3RemPowerPairs
	Power class	aLldpXdot3RemPowerClass
	Power type	aLldpXdot3RemPowerType
	Power source	aLldpXdot3RemPowerSource
	Power priority	aLldpXdot3RemPowerPriority
	PD requested power value	aLldpXdot3RemPDRequestedPowerValue
	PSE allocated power value	aLldpXdot3RemPSEAllocatedPowerValue
Link Aggregation (deprecated)	aggregation status	aLldpXdot3RemLinkAggStatus
	aggregated port ID	aLldpXdot3RemLinkAggPortId
Maximum Frame Size	maximum frame size	aLldpXdot3RemMaxFrameSize
EEE	Transmit $T_{w\_sys\_tx}$	aLldpXdot3RemTxTwSys
	Receive $T_{w\_sys\_tx}$	aLldpXdot3RemRxTwSys
	Echo Transmit $T_{w\_sys\_tx}$	aLldpXdot3RemTxTwSysEcho
	Echo Receive $T_{w\_sys\_tx}$	aLldpXdot3RemRxTwSysEcho
	Fallback $T_{w\_sys\_tx}$	aLldpXdot3RemFbTwSys
EEE Fast Wake	Transmit fast wake	aLldpXdot3RemTxFw
	Receive fast wake	aLldpXdot3RemRxFw
	Echo Transmit fast wake	aLldpXdot3RemTxFwEcho
	Echo Receive fast wake	aLldpXdot3RemRxFwEcho

79.5 Protocol implementation conformance statement (PICS) proforma for IEEE 802.3 Organizationally Specific Link Layer Discovery Protocol (LLDP) type, length, and value (TLV) information elements[3]

79.5.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 79, IEEE 802.3 Organizationally Specific Link Layer Discovery Protocol (LLDP) type, length, and value (TLV) information elements, 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.

79.5.2 Identification

79.5.2.1 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]	
NOTES 1—Required for all implementations. 2—May be completed as appropriate in meeting the requirements for the identification. 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

79.5.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 79, IEEE 802.3 Organizationally Specific Link Layer Discovery Protocol (LLDP) type, length, and value (TLV) information elements
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.

79.5.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*MP	MAC/PHY Configuration/Status TLV	79.3.1		O	Yes [ ] No [ ]
*PV	Power Via MDI TLV	79.3.2		O	Yes [ ] No [ ]
*LA	Link Aggregation TLV	79.3.3	TLV deprecated	O	Yes [ ] No [ ]
*FS	Maximum Frame Size TLV	79.3.4		O	Yes [ ] No [ ]
*EE	EEE TLV	79.3.5		O	Yes [ ] No [ ]
*EEFW	EEE Fast Wake TLV	79.3.6		O	Yes [ ] No [ ]

79.5.4 IEEE 802.3 Organizationally Specific TLV

Item	Feature	Subclause	Value/Comment	Status	Support
TLV1	Group MAC addresses	79.2	*Nearest device* group MAC addresses listed in Table 7-1 of IEEE Std 802.1AB-2009	M	Yes [ ]
TLV2	LLDPDU bit and octet ordering	79.2	Defined in subclause 8.1 of IEEE Std 802.1AB-2009	M	Yes [ ]

79.5.5 MAC/PHY Configuration/Status TLV

Item	Feature	Subclause	Value/Comment	Status	Support
MPT1	auto-negotiation support/status field	79.3.1.1	Identifies support and current status as defined in Table 79–2	MP:M	Yes [ ] N/A [ ]
MPT2	PMD auto-negotiation capability field	79.3.1.2	Bitmap of the ifMauAutoNeg-CapAdvertisedBits object defined in IETF RFC 4836	MP:M	Yes [ ] N/A [ ]
MPT3	operational MAU type field	79.3.1.3	Derived from the list position of the corresponding dot3Mau-Type as listed in IETF RFC 4836 (or subsequent revisions)	MP:M	Yes [ ] N/A [ ]
MPT4	operational MAU type field	79.3.1.3	Set to zero for MAU types not listed in IETF RFC 4836 (or subsequent revisions)	MP:M	Yes [ ] N/A [ ]
MPT5	Usage rules	79.3.1.4	LLDPDU contains no more than one MAC/PHY Configuration/Status TLV	MP:O	Yes [ ] No [ ] N/A [ ]

79.5.6 EEE TLV

Item	Feature	Subclause	Value/Comment	Status	Support
EET1	Transmit T_w field	79.3.5.1	2 octets representing time (expressed in microseconds) that the transmitting link partner will wait before it starts transmitting data after leaving the LPI mode	EE:M	Yes [ ] N/A [ ]
EET2	Receive T_w field	79.3.5.2	2 octets representing time (expressed in microseconds) that the receiving link partner is requesting the transmitting link partner to wait before it starts transmitting data following the LPI	EE:M	Yes [ ] N/A [ ]
EET3	Fallback field	79.3.5.3	2 octets representing time (expressed in microseconds)	EE:O	Yes [ ] N/A [ ]
EET4	Echo Transmit and Receive T_w fields	79.3.5.4	2 octets representing time (expressed in microseconds)	EE:M	Yes [ ] N/A [ ]
EET5	Usage rules	79.3.5.5	LLDPDU contains no more than one EEE TLV	EE:O	Yes [ ] No [ ] N/A [ ]

79.5.7 EEE Fast Wake TLV

Item	Feature	Subclause	Value/Comment	Status	Support
EFW1	Transmit fast wake field	79.3.6.1	1 octet representing fast wake option for transmit LPI function	EEFW: M	Yes [ ] N/A [ ]
EFW2	Receive fast wake field	79.3.6.2	1 octet representing fast wake option for receive LPI function	EEFW: M	Yes [ ] N/A [ ]
EFW3	Echo Transmit and Receive fast wake fields	79.3.6.3	2 octets representing received fast wake options	EEFW: M	Yes [ ] N/A [ ]

79.5.8 Power Via MDI TLV

Item	Feature	Subclause	Value/Comment	Status	Support
PVT1	MDI power support field	79.3.2.1	Bit map of the MDI power capabilities and status as defined in Table 79–3	PV:M	Yes [ ] N/A [ ]
PVT2	PSE power pair field	79.3.2.2	Integer value as defined by the pethPsePortPowerPairs object in IETF RFC 3621	PV:M	Yes [ ] N/A [ ]
PVT3	power class field	79.3.2.3	Integer value as defined by the pethPsePortPowerClassifications object in IETF RFC 3621	PV:M	Yes [ ] N/A [ ]
PVT4	Power type/source/priority field	79.3.2.4	Contains a bit-map of the power type, source, and priority defined in Table 79–4	PV:M	Yes [ ] N/A [ ]
PVT5	Power type field	79.3.2.4.1	Set according to Table 79–4	PV:M	Yes [ ] N/A [ ]
PVT6	Power source field when power type is PD	79.3.2.4.2	Set to ‘01’ when powered only through the PI; set to ‘11’ when powered from both; set to ‘00’ when information is not available	PV:M	Yes [ ] N/A [ ]
PVT7	Power source field when power type is PSE	79.3.2.4.2	When sourcing power through the PI, set to ‘01’ when using primary supply; set to ‘10’ when using backup source; set to ‘00’ when information is not available	PV:M	Yes [ ] N/A [ ]
PVT8	Power priority field when power type is PD	79.3.2.4.3	Set to the power priority configured for the device; set to ‘00’ if power priority is undetermined	PV:M	Yes [ ] N/A [ ]
PVT9	PD requested power value field	79.3.2.5	Contains the PD’s requested power value defined in Table 79–5	PV:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
PVT10	PSE allocated power value field	79.3.2.6	Contains the PSE's allocated power value defined in Table 79–6	PV:M	Yes [ ] N/A [ ]
PVT11	Usage rules	79.3.2.7	LLDPDU contains no more than one Power Via MDI TLV	PV:O	Yes [ ] No [ ] N/A [ ]

79.5.9 Link Aggregation TLV

Item	Feature	Subclause	Value/Comment	Status	Support
LAT1	link aggregation status field	79.3.3.1	Bitmap of the link aggregation capabilities and the current aggregation status as defined in Table 79–7	LA:M	Yes [ ] N/A [ ]
LAT2	aggregated port ID	79.3.3.2	IEEE 802.3 aggregated port identifier, aAggPortID	LA:M	Yes [ ] N/A [ ]
LAT3	Usage rules	79.3.3.3	LLDPDU contains no more than one Link Aggregation TLV	LA:O	Yes [ ] No [ ] N/A [ ]

79.5.10 Maximum Frame Size TLV

Item	Feature	Subclause	Value/Comment	Status	Support
FST1	maximum frame size field	79.3.4.1	Integer value indicating the maximum supported frame size	FS:M	Yes [ ] N/A [ ]
FST2	maximum frame size field	79.3.4.1	1518 for basic frames	FS:O/1	Yes [ ] No [ ] N/A [ ]
FST3	maximum frame size field	79.3.4.1	1522 for Q-tagged frames	FS:O/1	Yes [ ] No [ ] N/A [ ]
FST4	maximum frame size field	79.3.4.1	2000 for envelope frames	FS:O/1	Yes [ ] No [ ] N/A [ ]
FST5	Usage rules	79.3.4.2	LLDPDU contains no more than one Maximum Frame Size TLV	FS:O	Yes [ ] No [ ] N/A [ ]

80. Introduction to 40 Gb/s and 100 Gb/s networks

80.1 Overview

80.1.1 Scope

This clause describes the general requirements for 40 Gigabit and 100 Gigabit Ethernet. 40 Gigabit Ethernet uses the IEEE 802.3 MAC sublayer operating at a data rate of 40 Gb/s, coupled with any IEEE 802.3 40GBASE Physical Layer implementation. 100 Gigabit Ethernet uses the IEEE 802.3 MAC sublayer operating at a data rate of 100 Gb/s, coupled with any IEEE 802.3 100GBASE Physical Layer implementation. 40 Gb/s and 100 Gb/s Physical Layer entities, such as those specified in Table 80–1, provide a bit error ratio (BER) better than or equal to 10^{-12} at the MAC/PLS service interface.

40 Gigabit and 100 Gigabit Ethernet is defined for full duplex operation only.

80.1.2 Objectives

NOTE—The contents of this subclause have been deleted.

80.1.3 Relationship of 40 Gigabit and 100 Gigabit Ethernet to the ISO OSI reference model

40 Gigabit and 100 Gigabit Ethernet couples the IEEE 802.3 MAC to a family of 40 Gb/s and 100 Gb/s Physical Layers. The relationships among 40 Gigabit and 100 Gigabit Ethernet, the IEEE 802.3 MAC, and the ISO Open System Interconnection (OSI) reference model are shown in Figure 80–1.

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 XLGMII and CGMII, which, when implemented as a logical interconnection port between the MAC sublayer and the Physical Layer (PHY), uses a 64-bit wide data path as specified in Clause 81.
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 XLAUI (40 Gb/s Attachment Unit Interface) or CAUI-4 (100 Gb/s four-lane Attachment Unit Interface) at an observable interconnection port, uses a 4 lane data path as specified in Annex 83A, Annex 83B, Annex 83D, or Annex 83E.
d) The PMA service interface, which, when physically implemented as CAUI-10 (100 Gb/s ten-lane Attachment Unit Interface) at an observable interconnection port, uses a 10 lane data path as specified in Annex 83A or Annex 83B.
e) The PMD service interface, which, when physically implemented as XLPPI (40 Gb/s Parallel Physical Interface) at an observable interconnection port, uses a 4 lane data path as specified in Annex 86A.
f) The PMD service interface, which, when physically implemented as CPPI (100 Gb/s Parallel Physical Interface) at an observable interconnection port, uses a 10 lane data path as specified in Annex 86A.
g) The MDI as specified in Clause 89 for 40GBASE-FR uses a single lane data path.
h) The MDIs as specified in Clause 85 for 40GBASE-CR4, in Clause 86 for 40GBASE-SR4, in Clause 87 for 40GBASE-LR4 and 40GBASE-ER4, in Clause 88 for 100GBASE-LR4 and 100GBASE-ER4, in Clause 92 for 100GBASE-CR4, and in Clause 95 for 100GBASE-SR4 all use a 4 lane data path.
i) The MDIs as specified in Clause 85 for 100GBASE-CR10, and in Clause 86 for 100GBASE-SR10 use a 10 lane data path.

j) Although there is no electrical or mechanical specification of the MDI for backplane Physical Layers, the PMDs as specified in Clause 84 for 40GBASE-KR4, in Clause 93 for 100GBASE-KR4, and in Clause 94 for 100GBASE-KP4 all use a 4 lane data path.

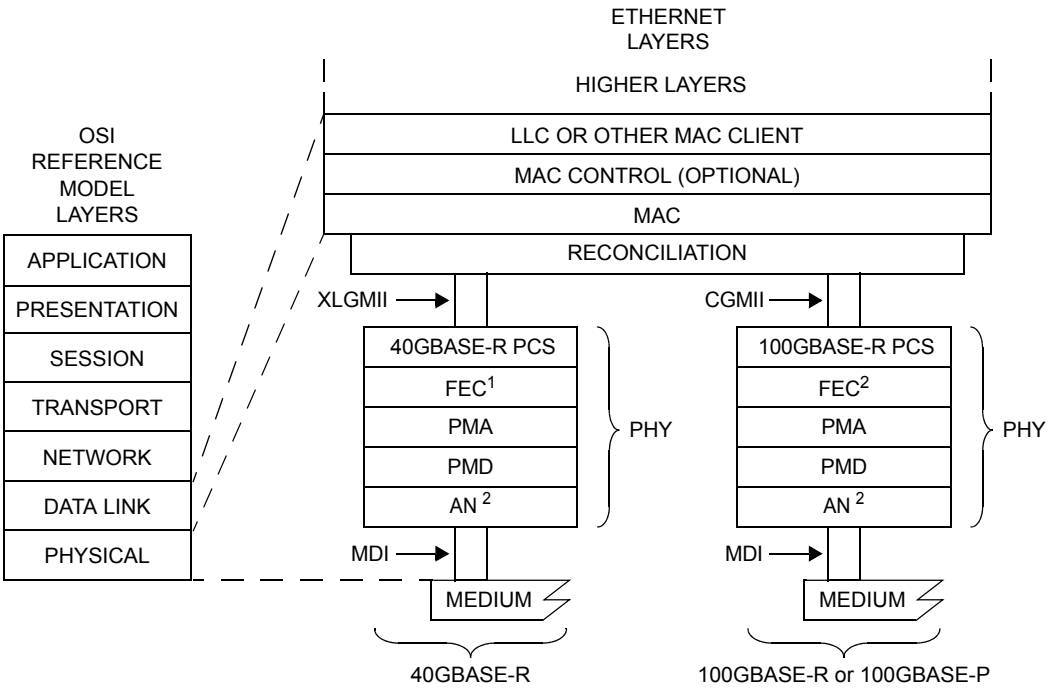

AN = AUTO-NEGOTIATION
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE
NOTE 2—CONDITIONAL BASED ON PHY TYPE

Figure 80–1—Architectural positioning of 40 Gigabit and 100 Gigabit Ethernet

80.1.4 Nomenclature

The nomenclature employed by the 40 Gb/s and 100 Gb/s Physical Layers is explained as follows.

The alpha-numeric prefix 40GBASE in the port type (e.g., 40GBASE-R) represents a family of Physical Layer devices operating at a speed of 40 Gb/s. The alpha-numeric prefix 100GBASE in the port type (e.g., 100GBASE-R) represents a family of Physical Layer devices operating at a speed of 100 Gb/s.

40GBASE-R represents a family of Physical Layer devices using the Clause 82 Physical Coding Sublayer for 40 Gb/s operation over multiple PCS lanes (see Clause 82). Some 40GBASE-R Physical Layer devices may also use the FEC of Clause 74.

100GBASE-R represents a family of Physical Layer devices using the Clause 82 Physical Coding Sublayer for 100 Gb/s operation over multiple PCS lanes (see Clause 82) and a PMD implementing 2-level pulse amplitude modulation (PAM). Some 100GBASE-R Physical Layer devices also use the transcoding and FEC of Clause 91 and some may also use the FEC of Clause 74.

100GBASE-P represents Physical Layer devices using the Clause 82 Physical Coding Sublayer for 100 Gb/s operation over multiple PCS lanes (see Clause 82) and a PMD implementing more than 2-level pulse amplitude modulation (PAM). Some 100GBASE-P Physical Layer devices also use the transcoding and FEC of Clause 91.

Physical Layer devices listed in Table 80–1 are defined for operation at 40 Gb/s and 100 Gb/s.

Table 80–1—40 Gb/s and 100 Gb/s PHYs

Name	Description
40GBASE-KR4	40 Gb/s PHY using 40GBASE-R encoding over four lanes of an electrical backplane, with reach up to at least 1 m (see Clause 84)
40GBASE-CR4	40 Gb/s PHY using 40GBASE-R encoding over four lanes of shielded balanced copper cabling, with reach up to at least 7 m (see Clause 85)
40GBASE-SR4	40 Gb/s PHY using 40GBASE-R encoding over four lanes of multimode fiber, with reach up to at least 100 m (see Clause 86)
40GBASE-FR	40 Gb/s PHY using 40GBASE-R encoding over one lane on single-mode fiber, with reach up to at least 2 km (see Clause 89)
40GBASE-LR4	40 Gb/s PHY using 40GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 10 km (see Clause 87)
40GBASE-ER4	40 Gb/s PHY using 40GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 40 km (see Clause 87)
100GBASE-KR4	100 Gb/s PHY using 100GBASE-R encoding, Clause 91 RS-FEC and 2-level pulse amplitude modulation over four lanes of an electrical backplane, with a total insertion loss up to 35 dB at 12.9 GHz (see Clause 93)
100GBASE-KP4	100 Gb/s PHY using 100GBASE-R encoding, Clause 91 RS-FEC and 4-level pulse amplitude modulation over four lanes of an electrical backplane, with a total insertion loss up to 33 dB at 7 GHz (see Clause 94)
100GBASE-CR4	100 Gb/s PHY using 100GBASE-R encoding and Clause 91 RS-FEC over four lanes of shielded balanced copper cabling, with reach up to at least 5 m (see Clause 92)
100GBASE-CR10	100 Gb/s PHY using 100GBASE-R encoding over ten lanes of shielded balanced copper cabling, with reach up to at least 7 m (see Clause 85)
100GBASE-SR10	100 Gb/s PHY using 100GBASE-R encoding over ten lanes of multimode fiber, with reach up to at least 100 m (see Clause 86)
100GBASE-SR4	100 Gb/s PHY using 100GBASE-R encoding over four lanes of multimode fiber, with reach up to at least 100 m (see Clause 95)
100GBASE-LR4	100 Gb/s PHY using 100GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 10 km (see Clause 88)
100GBASE-ER4	100 Gb/s PHY using 100GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 40 km (see Clause 88)

80.1.5 Physical Layer signaling systems

This standard specifies a family of Physical Layer implementations. Table 80–2, Table 80–3, and Table 80–4 specify the correlation between nomenclature and clauses. Implementations conforming to one or more nomenclatures must meet the requirements of the corresponding clauses.

Table 80–2—Nomenclature and clause correlation (40GBASE)

Nomenclature	Clause[a]														
	73	74	78	81		82	83	83A	83B	84	85	86	86A	87		89
	Auto-Negotiation	BASE-R FEC	EEE	RS	XLGMII	40GBASE-R PCS	40GBASE-R PMA	XLAUI	XLAUI	40GBASE-KR4 PMD	40GBASE-CR4 PMD	40GBASE-SR4 PMD	XLPPI	40GBASE-LR4 PMD	40GBASE-ER4 PMD	40GBASE-FR PMD
40GBASE-KR4	M	O	O	M	O	M	M	O		M						
40GBASE-CR4	M	O	O	M	O	M	M	O			M					
40GBASE-SR4			O	M	O	M	M	O	O			M	O			
40GBASE-FR			O	M	O	M	M	O	O							M
40GBASE-LR4			O	M	O	M	M	O	O				O	M		
40GBASE-ER4			O	M	O	M	M	O	O						M	

[a]O = Optional, M = Mandatory.

Table 80–3—Nomenclature and clause correlation (100GBASE copper)

Nomenclature	Clause[a]												
	73	74	78	81		82	83	83A	83D	85	91	92	93	94
	Auto-Negotiation	BASE-R FEC	EEE	RS	CGMII	100GBASE-R PCS	100GBASE-R PMA	CAUI-10	CAUI-4	100GBASE-CR10 PMD	RS-FEC	100GBASE-CR4 PMD	100GBASE-KR4 PMD	100GBASE-KP4 PMD
100GBASE-KR4	M		O	M	O	M	M	O	O		M		M	
100GBASE-KP4	M		O	M	O	M	O	O	O		M			M
100GBASE-CR4	M		O	M	O	M	M	O	O		M	M		
100GBASE-CR10	M	O	O	M	O	M	M	O	O	M				

[a]O = Optional, M = Mandatory.

Table 80–4—Nomenclature and clause correlation (100GBASE optical)

Nomenclature	Clause[a]														
	78	81		82	83	83A	83B	83D	83E	86	86A	88		91	95
	EEE	RS	CGMII	100GBASE-R PCS	100GBASE-R PMA	CAUI-10	CAUI-10	CAUI-4	CAUI-4	100GBASE-SR10 PMD	CPPI	100GBASE-LR4 PMD	100GBASE-ER4 PMD	RS-FEC	100GBASE-SR4 PMD
100GBASE-SR10	O	M	O	M	M	O	O	O	O	M	O				
100GBASE-SR4	O	M	O	M	M	O	O	O	O					M	M
100GBASE-LR4	O	M	O	M	M	O	O	O	O			M			
100GBASE-ER4	O	M	O	M	M	O	O	O	O				M		

[a]O = Optional, M = Mandatory.

80.2 Summary of 40 Gigabit and 100 Gigabit Ethernet sublayers

80.2.1 Reconciliation Sublayer (RS) and Media Independent Interface

The Media Independent Interface (Clause 81) provides a logical interconnection between the MAC sublayer and Physical Layer entities (PHY). The Media Independent Interface is not intended to be physically instantiated, rather it can logically connect layers within a device.

The XLGMII supports 40 Gb/s and CGMII supports 100 Gb/s operation through its 64-bit-wide transmit and receive data paths. The Reconciliation Sublayer (RS) provides a mapping between the signals provided at the Media Independent Interface (XLGMII and CGMII) and the MAC/PLS service definition.

While XLGMII and CGMII are optional interfaces, they are used extensively in this standard as a basis for functional specification and provides a common service interface for the physical coding sublayers defined in Clause 82.

80.2.2 Physical Coding Sublayer (PCS)

The terms 40GBASE-R, 100GBASE-R, and 100GBASE-P refer to a specific family of Physical Layer implementations based upon the 64B/66B data coding method specified in Clause 82 and the PMA specifications defined in Clause 83 or Clause 94. Clause 82 PCSs perform encoding (decoding) of data from (to) the XLGMII/CGMII to 64B/66B code blocks, distribute the data to multiple lanes, and transfer the encoded data to the PMA.

80.2.3 Forward Error Correction (FEC) sublayers

A Forward Error Correction sublayer is available for all 40GBASE-R and 100GBASE-R copper and backplane PHYs as well as 100GBASE-SR4. It is optional for 40GBASE-KR4, 40GBASE-CR4, and 100GBASE-CR10 PHYs and mandatory for 100GBASE-CR4, 100GBASE-KR4, 100GBASE-KP4, and

100GBASE-SR4 PHYs. The FEC sublayer can be placed in between the PCS and PMA sublayers or between two PMA sublayers.

The BASE-R FEC (see Clause 74) is instantiated for each PCS lane and operates autonomously on a per PCS lane basis. The Reed-Solomon FEC (see Clause 91) is instantiated once and requires 20 PCS lanes and 4 PMA lanes for operation.

80.2.4 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 40GBASE-R and 100GBASE-R PMAs perform the mapping of transmit and receive data streams between the PCS and PMA via the PMA service interface, and the mapping and multiplexing of transmit and receive data streams between the PMA and PMD via the PMD service interface. In addition, the PMAs perform retiming of the received data stream when appropriate, optionally provide data loopback at the PMA or PMD service interface, and optionally provide test pattern generation and checking.

The 40GBASE-R and 100GBASE-R PMAs are specified in Clause 83. The PMA specific to the 100GBASE-KP4 PHY is specified in Clause 94.

80.2.5 Physical Medium Dependent (PMD) sublayer

The Physical Medium Dependent sublayer is responsible for interfacing to the transmission medium. The PMD is located just above the Medium Dependent Interface (MDI). The MDI, logically subsumed within each PMD subclause, is the actual medium attachment for the various supported media.

The 40GBASE-R, 100GBASE-R, and 100GBASE-P PMDs and their corresponding media are specified in Clause 84 through Clause 89 and Clause 92 through Clause 95.

80.2.6 Auto-Negotiation

Auto-Negotiation provides a linked device with the capability to detect the abilities (modes of operation) supported by the device at the other end of the link, determine common abilities, and configure for joint operation.

Clause 73 Auto-Negotiation is used by the 40 Gb/s and 100 Gb/s backplane PHYs (40GBASE-KR4, 100GBASE-KP4, and 100GBASE-KR4) and the 40 Gb/s and 100 Gb/s copper PHYs (40GBASE-CR4, 100GBASE-CR10, and 100GBASE-CR4).

80.2.7 Management interface (MDIO/MDC)

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

80.2.8 Management

Managed objects, attributes, and actions are defined for all 40 Gigabit and 100 Gigabit Ethernet components. These items are defined in Clause 30.

80.3 Service interface specification method and notation

The service interface specification for 40GBASE-R, 100GBASE-R, and 100GBASE-P Physical Layers is as per the definition in 1.2.2. Note that the 40GBASE-R, 100GBASE-R, and 100GBASE-P inter-sublayer

service interfaces use multiple scalar REQUEST and INDICATION primitives, to indicate the transfer of multiple independent streams of data units, as explained in 80.3.1 through 80.3.3.

80.3.1 Inter-sublayer service interface

The inter-sublayer service interface is described in an abstract manner and does not imply any particular implementation. The inter-sublayer service interface primitives are defined as follows:

IS_UNITDATA_*i*.request
IS_UNITDATA_*i*.indication
IS_SIGNAL.indication

The IS_UNITDATA_*i*.request (where $i = 0$ to n – 1, and n is the number of streams of data units) primitive is used to define the transfer of multiple streams of data units from a sublayer N to the next lower sublayer N – 1. The IS_UNITDATA_*i*.indication (where $i = 0$ to n – 1, and n is the number of streams of data units) primitive is used to define the transfer of multiple streams of data units from a sublayer N – 1 to the next higher sublayer N. The IS_SIGNAL.indication primitive is used to define the transfer of signal status from a sublayer N – 1 to the next higher sublayer N.

If the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option is supported (see Clause 78, 78.1.3.3.1), then the inter-sublayer service interface includes five additional primitives defined as follows:

IS_TX_MODE.request
IS_RX_MODE.request
IS_ENERGY_DETECT.indication
IS_RX_LPI_ACTIVE.request
IS_RX_TX_MODE.indication

The IS_TX_MODE.request primitive is used to communicate the state of the PCS LPI transmit function to other sublayers in the PHY. The IS_RX_MODE.request primitive is used to communicate the state of the PCS LPI receive function to other sublayers. The IS_RX_TX_MODE.indication primitive is used to communicate the state of the rx_tx_mode parameter that reflects the inferred state of the link partner's tx_mode parameter from the PMA to other sublayers. The IS_RX_LPI_ACTIVE.request primitive is used to communicate to the BASE-R FEC (see Clause 74) that the PCS has detected LPI signaling. This allows the BASE-R FEC to use rapid block lock; the RS-FEC (see Clause 91) does not use this signal. The IS_ENERGY_DETECT.indication primitive is used to communicate that the PMD has detected the return of energy on the interface following a period of quiescence.

80.3.2 Instances of the Inter-sublayer service interface

The inter-sublayer interface can be instantiated between different sublayers, hence a prefix notation is defined to identify a specific instance of an inter-sublayer service interface. The following prefixes are defined:

a) PMD:—for primitives issued on the interface between the PMD sublayer and the PMA sublayer called the PMD service interface.
b) PMA:—for primitives issued on the interface between the PMA sublayer and the PCS (or the FEC) sublayer called the PMA service interface.
c) FEC:—for primitives issued on the interface between the FEC sublayer and the PCS (or the PMA) sublayer called the FEC service interface.

Examples of inter-sublayer service interfaces for 40GBASE-R, 100GBASE-R, and 100GBASE-P with their

corresponding instance names are illustrated in Figure 80–2, Figure 80–3, Figure 80–4, and Figure 80–5. For example, the primitives for one instance of the inter-sublayer service interface, named the PMD service interface, are identified as follows:

PMD:IS_UNITDATA_*i*.request
PMD:IS_UNITDATA_*i*.indication
PMD:IS_SIGNAL.indication.

Primitives for other instances, of inter-sublayer interfaces, are represented in a similar manner as described above.

80.3.3 Semantics of inter-sublayer service interface primitives

The semantics of the inter-sublayer service interface primitives for the 40GBASE-R and 100GBASE-R sublayers are described in 80.3.3.1 through 80.3.3.3.

80.3.3.1 IS_UNITDATA_*i*.request

The IS_UNITDATA_*i*.request (where i = 0 to n – 1) primitive is used to define the transfer of multiple streams of data units from a sublayer N to the next lower sublayer N – 1, where n is the number of parallel streams of data units.

80.3.3.1.1 Semantics of the service primitive

IS_UNITDATA_0.request(tx_bit)
IS_UNITDATA_1.request(tx_bit)
...
IS_UNITDATA_n–1.request(tx_bit)

The data conveyed by IS_UNITDATA_0.request to IS_UNITDATA_n–1.request consists of n parallel continuous streams of encoded bits, one stream for each lane. Each of the tx_bit parameters can take one of two values: one or zero.

80.3.3.1.2 When generated

The sublayer N continuously sends n parallel bit streams IS_UNITDATA_*i*.request(tx_bit) to the next lower sublayer N – 1, each at a nominal signaling rate defined by a specific instance of the inter-sublayer service interface.

80.3.3.1.3 Effect of receipt

The effect of receipt of this primitive is defined by the sublayer that receives this primitive.

80.3.3.2 IS_UNITDATA_*i*.indication

The IS_UNITDATA_*i*.indication (where i = 0 to n – 1) primitive is used to define the transfer of multiple streams of data units from the sublayer N – 1 to the next higher sublayer N, where n is the number of parallel streams of data units.

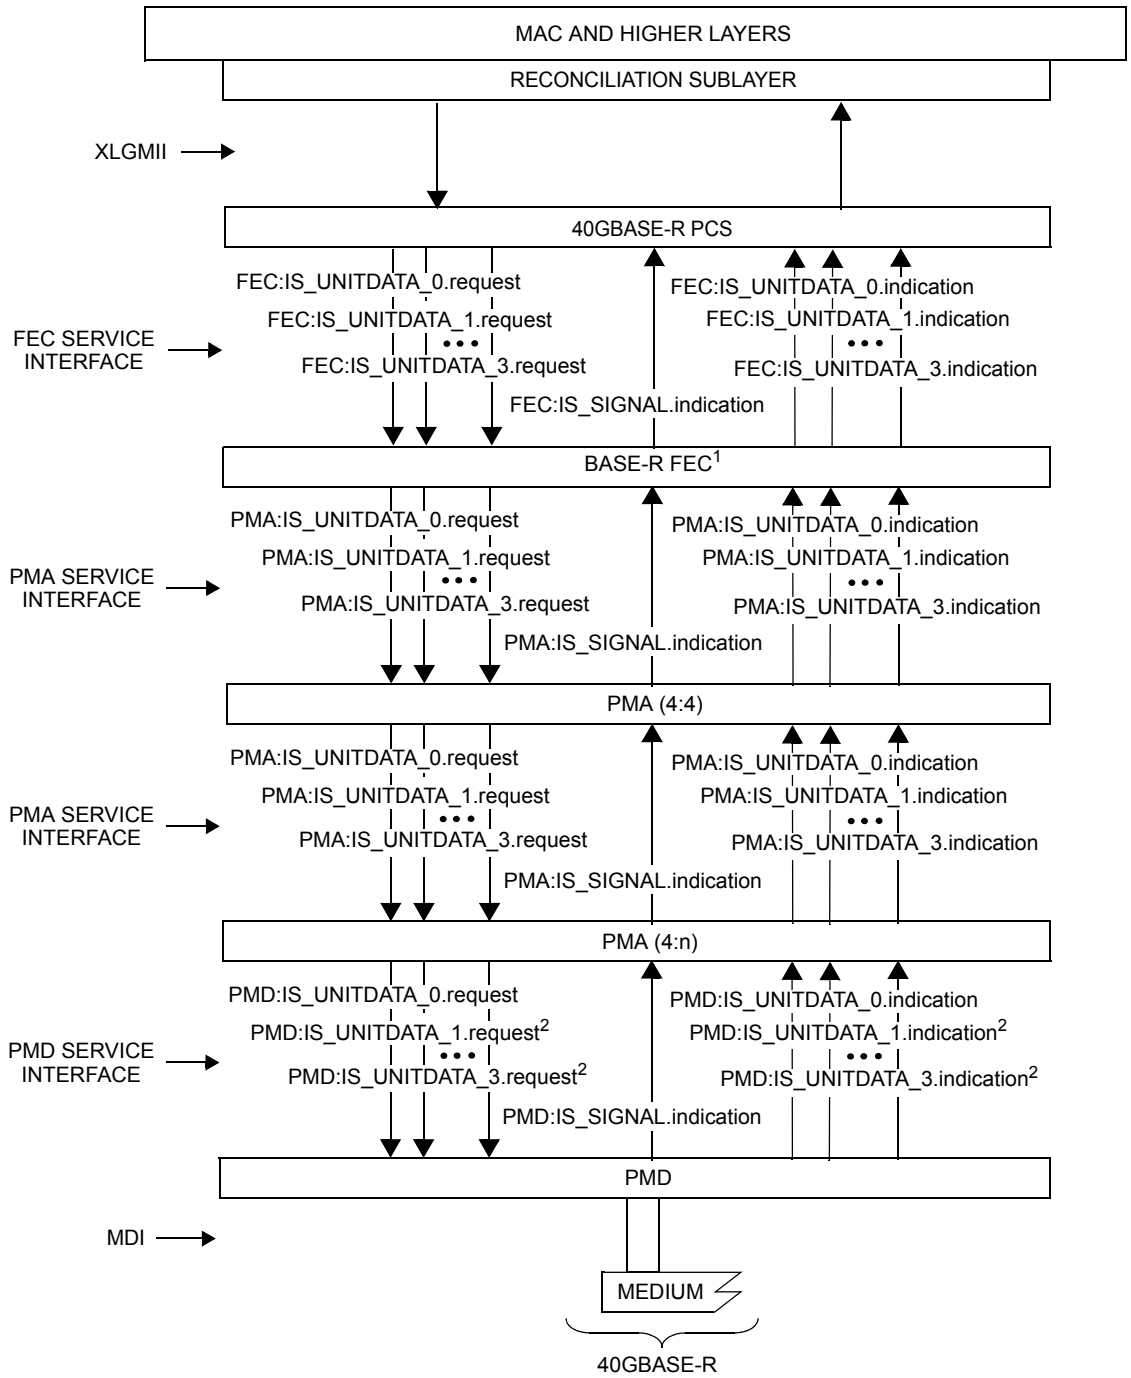

FEC = FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT

XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE
n = NUMBER OF PARALLEL STREAMS OF DATA UNITS

NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE
NOTE 2—DOES NOT EXIST FOR 40GBASE-FR PMD

Figure 80–2—40GBASE-R inter-sublayer service interfaces

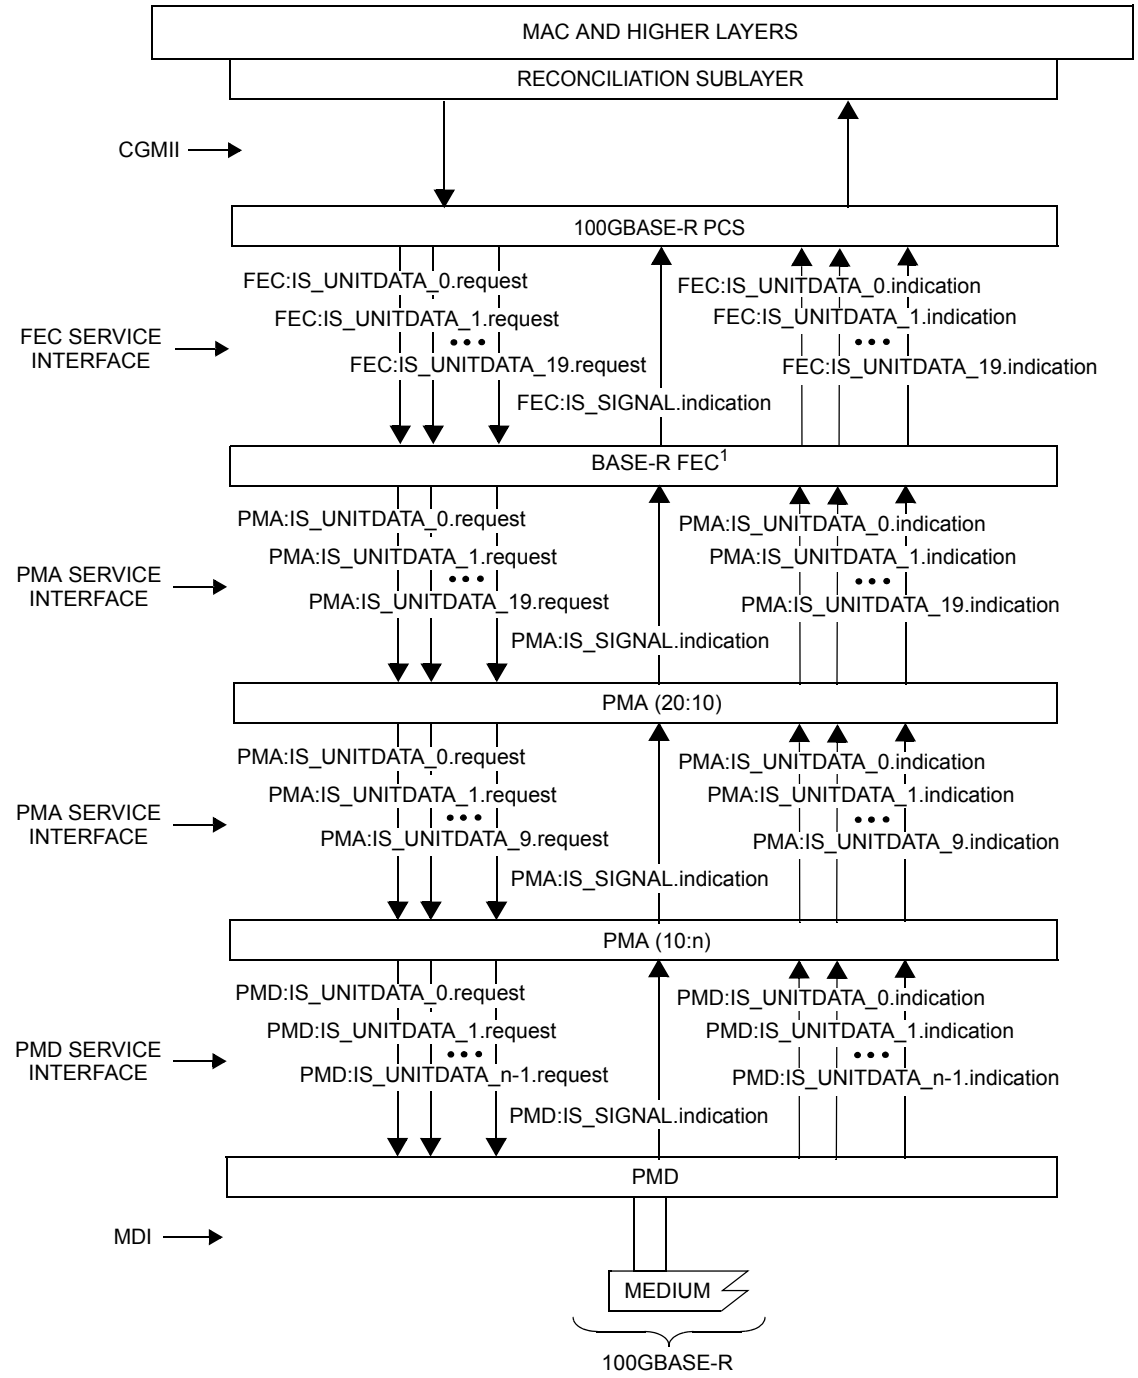

CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
n = NUMBER OF PARALLEL STREAMS OF DATA UNITS

NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE

Figure 80–3—100GBASE-R inter-sublayer service interfaces

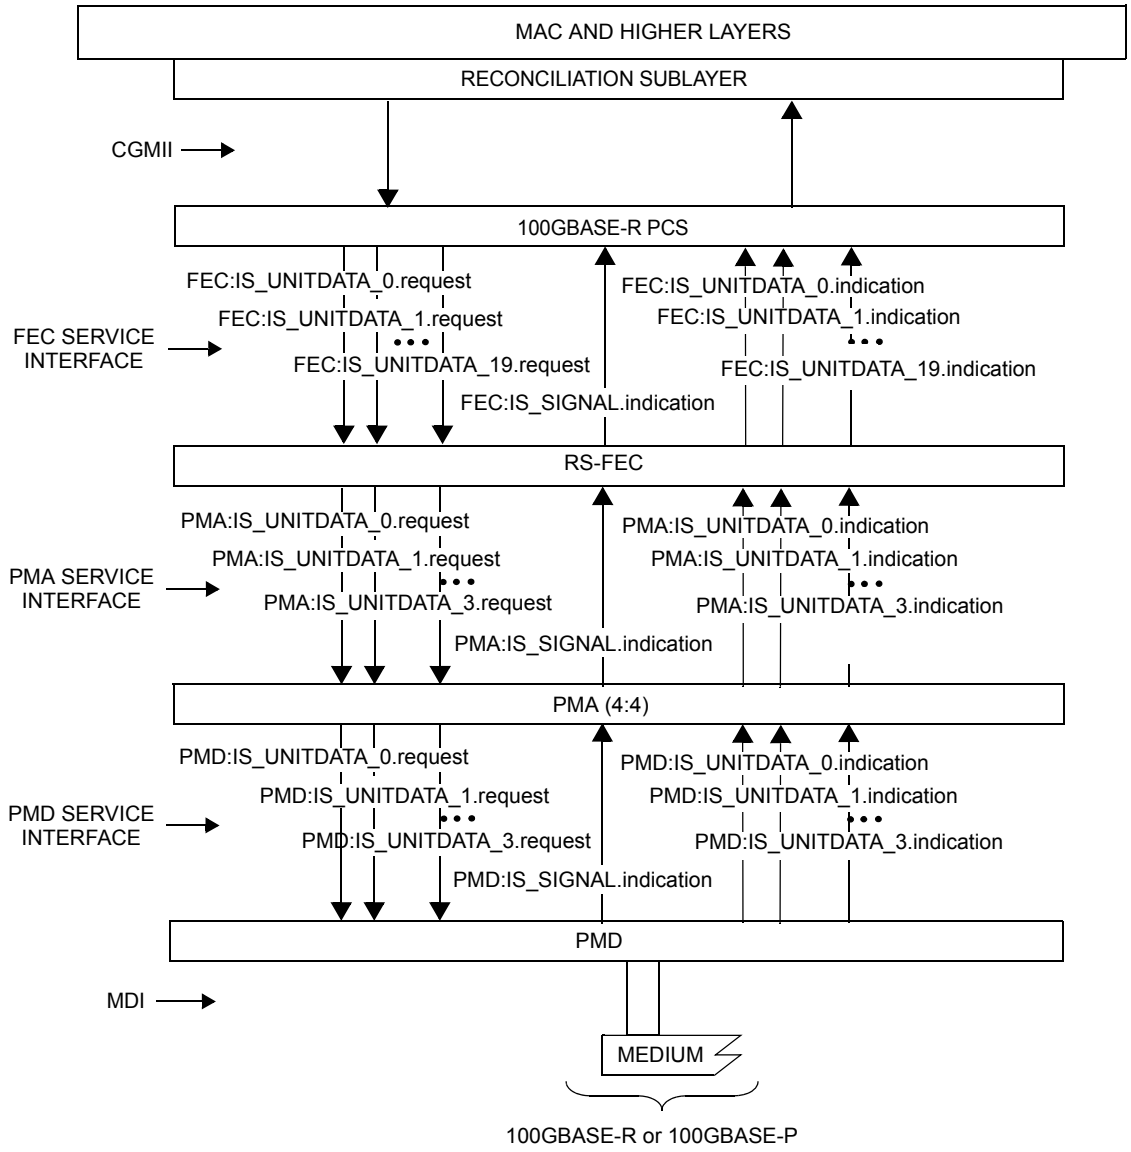

CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT

Figure 80–4—100GBASE-R and 100GBASE-P inter-sublayer service interfaces with RS-FEC

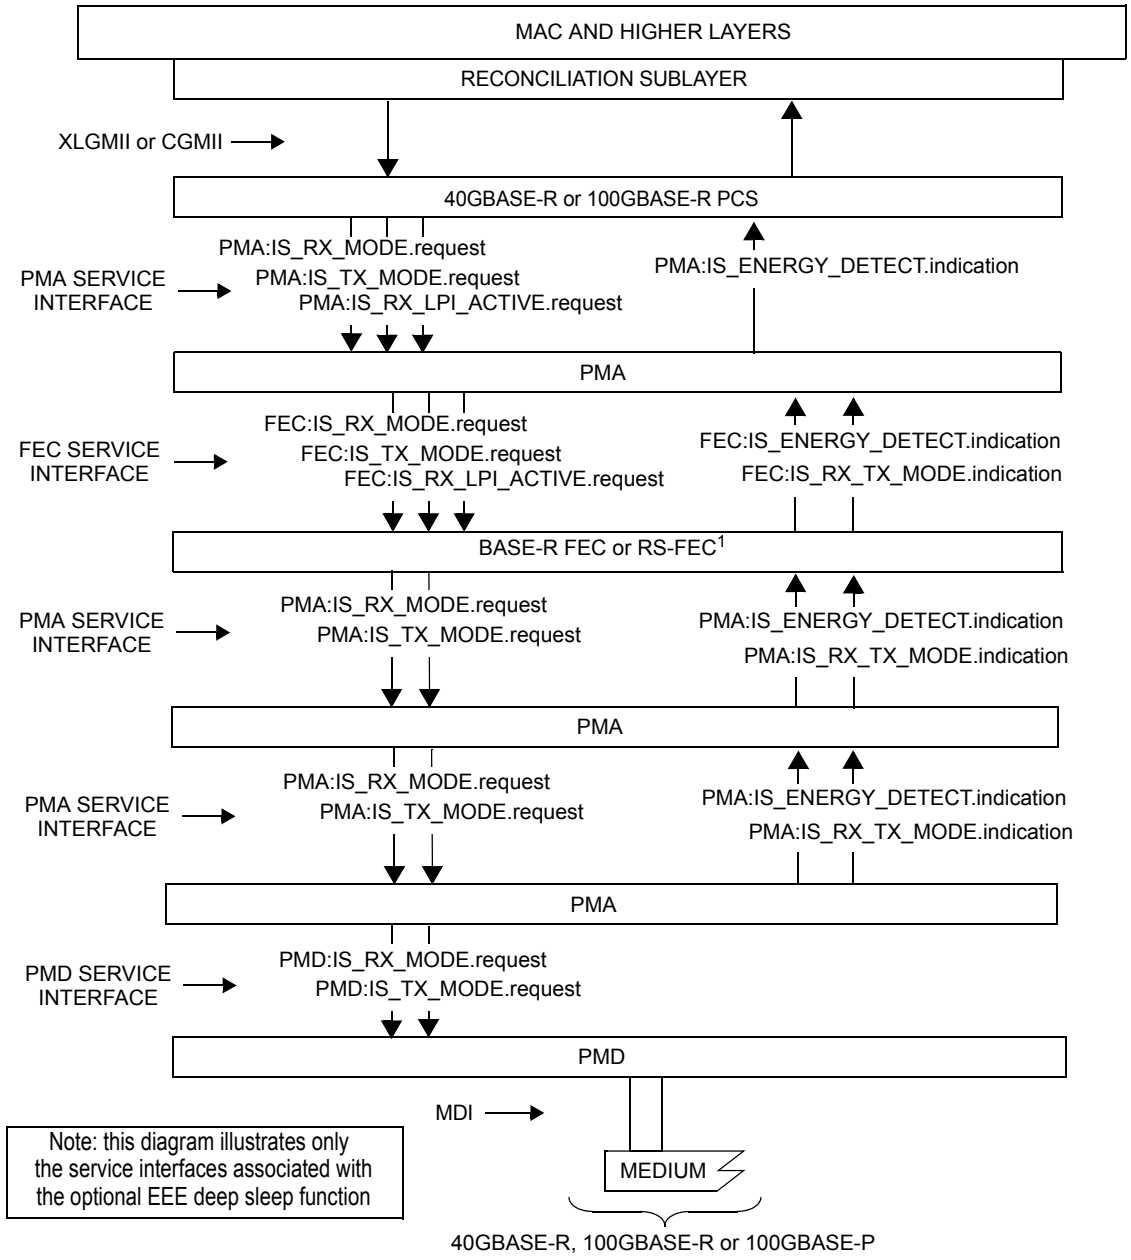

CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT

NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 80–5—Optional inter-sublayer service interfaces for EEE deep sleep support

80.3.3.2.1 Semantics of the service primitive

IS_UNITDATA_0.indication(rx_bit)
IS_UNITDATA_1.indication(rx_bit)
...
IS_UNITDATA_n–1.indication(rx_bit)

The data conveyed by IS_UNITDATA_0.indication to IS_UNITDATA_n–1.indication consists of n parallel continuous streams of encoded bits, one stream for each lane. Each of the rx_bit parameters can take one of two values: one or zero.

80.3.3.2.2 When generated

The sublayer N–1 continuously sends n parallel bit streams IS_UNITDATA_*i*.indication(rx_bit) to the next higher sublayer N, each at a nominal signaling rate defined by a specific instance of the inter-sublayer service interface.

80.3.3.2.3 Effect of receipt

The effect of receipt of this primitive is defined by the sublayer that receives this primitive.

80.3.3.3 IS_SIGNAL.indication

The IS_SIGNAL.indication primitive is generated by the sublayer N – 1 to the next higher sublayer N to indicate the status of the receive process. This primitive is generated by the receive process to propagate the detection of severe error conditions (e.g., no valid signal being received by the sublayer that generates this primitive) to the next higher sublayer N.

80.3.3.3.1 Semantics of the service primitive

IS_SIGNAL.indication(SIGNAL_OK)

The SIGNAL_OK parameter can take on one of two values: OK or FAIL. A value of FAIL denotes that invalid data is being presented (rx_bit parameters undefined) by the sublayer N – 1 to the next higher sublayer N. A value of OK does not guarantee valid data is being presented by the sublayer N – 1 to the next higher sublayer N.

80.3.3.3.2 When generated

The sublayer N – 1 generates the IS_SIGNAL.indication primitive to the next higher sublayer N whenever there is change in the value of the SIGNAL_OK parameter.

80.3.3.3.3 Effect of receipt

The effect of receipt of this primitive is defined by the sublayer that receives this primitive.

80.3.3.4 IS_TX_MODE.request

The IS_TX_MODE.request primitive communicates the tx_mode parameter generated by the PCS Transmit Process for EEE capability to invoke the appropriate PMA, FEC, and PMD transmit EEE states. Without EEE deep sleep mode capability, the primitive is never invoked and the sublayers behave as if tx_mode = DATA.

80.3.3.4.1 Semantics of the service primitive

IS_TX_MODE.request(tx_mode)

The tx_mode parameter takes on one of up to three values: DATA, QUIET, or ALERT.

80.3.3.4.2 When generated

This primitive is generated to indicate the low power mode of the transmit path.

80.3.3.4.3 Effect of receipt

The specific effect of receipt of this primitive is defined by the sublayer that receives this primitive. In general, when tx_mode is DATA the sublayer operates normally and when tx_mode is QUIET, the sublayer may go into a low power mode.

80.3.3.5 IS_RX_MODE.request

The IS_RX_MODE.request primitive communicates the rx_mode parameter generated by the PCS LPI receive function to other sublayers. Without EEE deep sleep mode capability, the primitive is never invoked and the sublayers behave as if rx_mode = DATA.

80.3.3.5.1 Semantics of the service primitive

IS_RX_MODE.request(rx_mode)

The rx_mode parameter takes on one of two values: DATA or QUIET.

80.3.3.5.2 When generated

This primitive is generated to indicate the state of the PCS LPI receive function.

80.3.3.5.3 Effect of receipt

The specific effect of receipt of this primitive is defined by the sublayer that receives this primitive. In general, when rx_mode is DATA the sublayer operates normally and when rx_mode is QUIET, the sublayer may go into a low power mode.

80.3.3.6 IS_RX_LPI_ACTIVE.request

The IS_RX_LPI_ACTIVE.request primitive communicates to the FEC that the PCS LPI receive function is active. This primitive may be passed through a PMA sublayer but has no effect on that sublayer. This primitive is only used for a PMA sublayer that is between the PCS and a Clause 74 FEC sublayer; in all other cases the primitive is never invoked and has no effect. Without EEE deep sleep mode capability, the primitive is never invoked and has no effect.

80.3.3.6.1 Semantics of the service primitive

IS_RX_LPI_ACTIVE.request(rx_lpi_active)

The parameter rx_lpi_active is Boolean.

80.3.3.6.2 When generated

This primitive is generated to indicate the state of the PCS LPI receive function. It is FALSE when in the RX_ACTIVE state and TRUE in all other states.

80.3.3.6.3 Effect of receipt

The specific effect of receipt of this primitive is defined by the FEC sublayer that receives this primitive. When rx_lpi_active is true the FEC sublayer uses rapid block lock to reestablish FEC operation following a period of quiescence.

80.3.3.7 IS_ENERGY_DETECT.indication

The IS_ENERGY_DETECT.indication primitive is used to communicate that the PMD has detected the return of energy on the interface following a period of quiescence. Without EEE deep sleep mode capability, the primitive is never invoked and has no effect.

80.3.3.7.1 Semantics of the service primitive

IS_ENERGY_DETECT.indication(energy_detect)

The parameter energy_detect is Boolean.

80.3.3.7.2 When generated

This primitive is generated by the PMA, reflecting the state of the signal_detect parameter received from the PMD.

80.3.3.7.3 Effect of receipt

The specific effect of receipt of this primitive is defined by the PCS sublayer that receives this primitive. This parameter is used to indicate that activity has returned on the interface following a period of quiescence.

80.3.3.8 IS_RX_TX_MODE.indication

The IS_RX_TX_MODE.indication primitive communicates the rx_tx_mode parameter. This parameter indicates the value of tx_mode that the PMA sublayer has inferred from the received signal. Without EEE deep sleep capability, the primitive is never generated and the sublayers behave as if rx_tx_mode=DATA.

80.3.3.8.1 Semantics of the service primitive

IS_RX_TX_MODE.indication(rx_tx_mode)

The parameter rx_tx_mode is assigned one of the following values: DATA, QUIET, or ALERT.

80.3.3.8.2 When generated

This primitive is generated whenever there is change in the value of the rx_tx_mode parameter.

80.3.3.8.3 Effect of receipt

The specific effect of receipt of this primitive is defined by the sublayer that receives it.

80.4 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. Table 80–5 contains the values of maximum sublayer delay (sum of transmit and receive delays at one end of the link) in bit times as specified in 1.4 and pause_quanta as specified in 31B.2. If a PHY contains an Auto-Negotiation sublayer, the delay of the Auto-Negotiation sublayer is included within the delay of the PMD and medium.

Equation (80–1) specifies the calculation of cable delay in nanoseconds 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, $c = 3 \times 10^8$ m/s.

$$\text{cable delay} = \frac{10^9}{nc} \text{ ns/m} \tag{80–1}$$

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 $n = 0.66$, which yields a default cable delay of 5 ns/m.

See 31B.3.7 for PAUSE reaction timing constraints for stations at operating speeds of 40 Gb/s and 100 Gb/s.

80.5 Skew constraints

Skew (or relative delay) can be introduced between lanes by both active and passive elements of a 40GBASE-R or 100GBASE-R link. Skew is defined as the difference between the times of the earliest PCS lane and latest PCS lane for the one to zero transition of the alignment marker sync bits. The PCS deskew function (see 82.2.13) compensates for all lane-to-lane Skew observed at the receiver. The Skew between the lanes must be kept within limits as shown in Table 80–6 so that the transmitted information on the lanes can be reassembled by the receive PCS.

Skew Variation may be introduced due to variations in electrical, thermal or environmental characteristics. Skew Variation is defined as the change in Skew between any PCS lane and any other PCS lane over the entire time that the link is in operation. From the time the link is brought up, Skew Variation must be limited to ensure that each PCS lane always traverses the same lane between any pair of adjacent sublayers while the link remains in operation.

Table 80–5—Sublayer delay constraints

Sublayer	Maximum (bit time)[a]	Maximum (pause_quanta)[b]	Maximum (ns)	Notes[c]
40G MAC, RS, and MAC Control	16384	32	409.6	See 81.1.4.
40GBASE-R PCS	11264	22	281.6	See 82.5.
40GBASE-R FEC	24576	48	614.4	See 74.6.
40GBASE-R PMA	4096	8	102.4	See 83.5.4.
40GBASE-KR4 PMD	2048	4	51.2	Includes delay of one direction through backplane medium. See 84.4.
40GBASE-CR4 PMD	4096	8	102.4	Does not include delay of cable medium. See 85.4.
40GBASE-SR4 PMD	1024	2	25.6	Includes 2 m of fiber. See 86.3.1.
40GBASE-FR PMD	1024	2	25.6	Includes 2 m of fiber. See 89.3.1.
40GBASE-LR4 PMD	1024	2	25.6	Includes 2 m of fiber. See 87.3.1.
40GBASE-ER4 PMD	1024	2	25.6	Includes 2 m of fiber. See 87.3.1.
100G MAC, RS, and MAC Control	24576	48	245.76	See 81.1.4.
100GBASE-R PCS	35328	69	353.28	See 82.5.
100GBASE-R FEC	122880	240	1228.8	See 74.6.
100GBASE-R RS-FEC	40960	80	409.60	See 91.4.
100GBASE-R PMA	9216	18	92.16	See 83.5.4.
100GBASE-KR4 PMD	2048	4	20.48	Includes delay of one direction through backplane medium. See 93.4.
100GBASE-KP4 PMA/PMD	8192	16	81.92	Includes delay of one direction through backplane medium. See 94.2.5.
100GBASE-CR4 PMD	2048	4	20.48	Does not include delay of cable medium. See 92.4.
100GBASE-CR10 PMD	9728	19	97.28	Does not include delay of cable medium. See 85.4.
100GBASE-SR4 PMD	2048	4	20.48	Includes 2 m of fiber. See 95.3.1.
100GBASE-SR10 PMD	2048	4	20.48	Includes 2 m of fiber. See 86.3.1.
100GBASE-LR4 PMD	2048	4	20.48	Includes 2 m of fiber. See 88.3.1.
100GBASE-ER4 PMD	2048	4	20.48	Includes 2 m of fiber. See 88.3.1.

[a] For 40GBASE-R, 1 bit time (BT) is equal to 25 ps and for 100GBASE-R, 1 bit time (BT) is equal to 10 ps. (See 1.4.117 for the definition of bit time.)
[b] For 40GBASE-R, 1 pause_quantum is equal to 12.8 ns and for 100GBASE-R, 1 pause_quantum is equal to 5.12 ns. (See 31B.2 for the definition of pause_quanta.)
[c] Should there be a discrepancy between this table and the delay requirements of the relevant sublayer clause, the sublayer clause prevails.

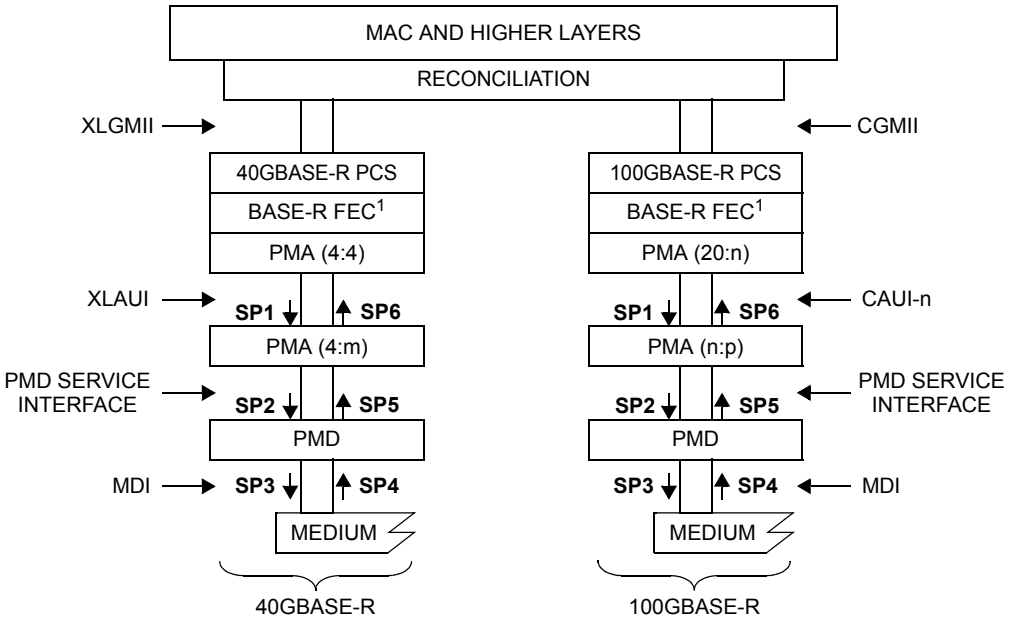

CAUI-n = 100 Gb/s ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLAUI = 40 Gb/s ATTACHMENT UNIT INTERFACE
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

m = 1 or 4
n = 4 or 10
p = 4 or 10
NOTE1—CONDITIONAL BASED ON PHY TYPE

Figure 80–6—40GBASE-R and 100GBASE-R Skew points for single XLAUI or CAUI-n

The maximum Skew and Skew Variation at physically instantiated interfaces is specified at Skew points SP1, SP2, and SP3 for the transmit direction and SP4, SP5, and SP6 for the receive direction as illustrated in Figure 80–6 (single XLAUI or CAUI-n interface) and Figure 80–7 (multiple XLAUI or CAUI-n interfaces).

In the transmit direction, the Skew points are defined in the following locations:

— SP1 on the XLAUI/CAUI-n interface, at the input of the PMA closest to the PMD;
— SP2 on the PMD service interface, at the input of the PMD;
— SP3 at the output of the PMD, at the MDI.

In the receive direction, the Skew points are defined in the following locations:

— SP4 at the MDI, at the input of the PMD;
— SP5 on the PMD service interface, at the output of the PMD;
— SP6 on the XLAUI/CAUI-n interface, at the output of the PMA closest to the PCS.

The allowable limits for Skew are shown in Table 80–6 and the allowable limits for Skew Variation are shown in Table 80–7.

The skew points are similarly illustrated for a PHY incorporating RS-FEC (see Clause 91) in Figure 80–8.

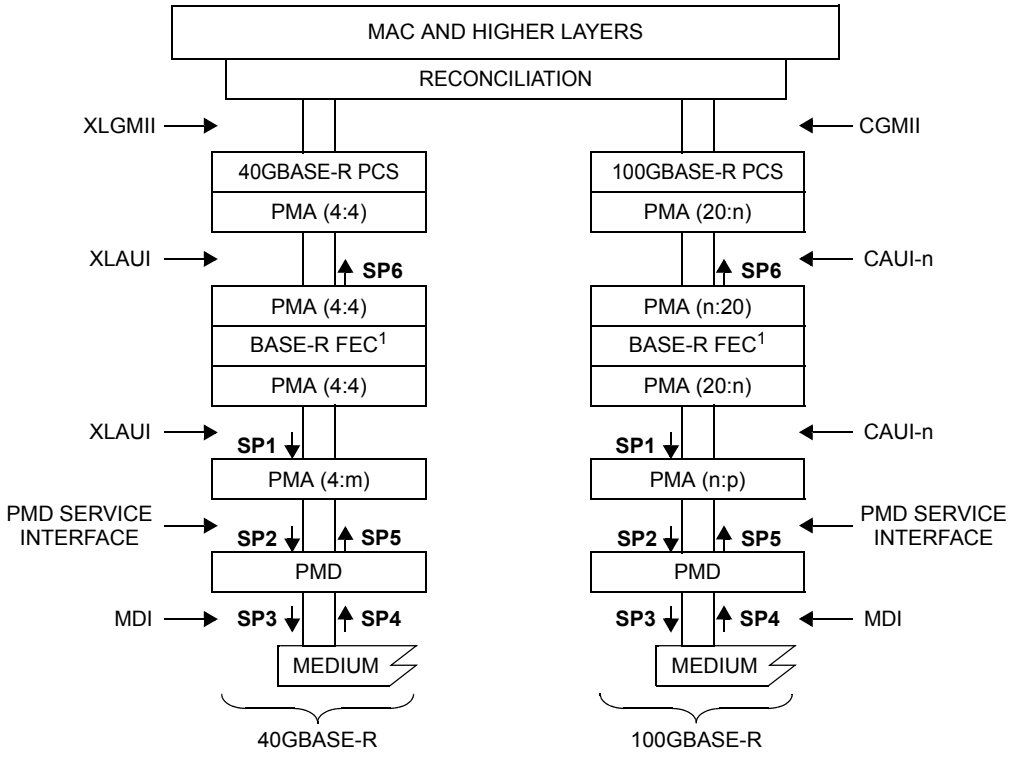

CAUI-n = 100 Gb/s ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLAUI = 40 Gb/s ATTACHMENT UNIT INTERFACE
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

m = 1 or 4
n = 4 or 10
p = 4 or 10
NOTE1—CONDITIONAL BASED ON PHY TYPE

Figure 80–7—40GBASE-R and 100GBASE-R Skew points for multiple XLAUI or CAUI-n

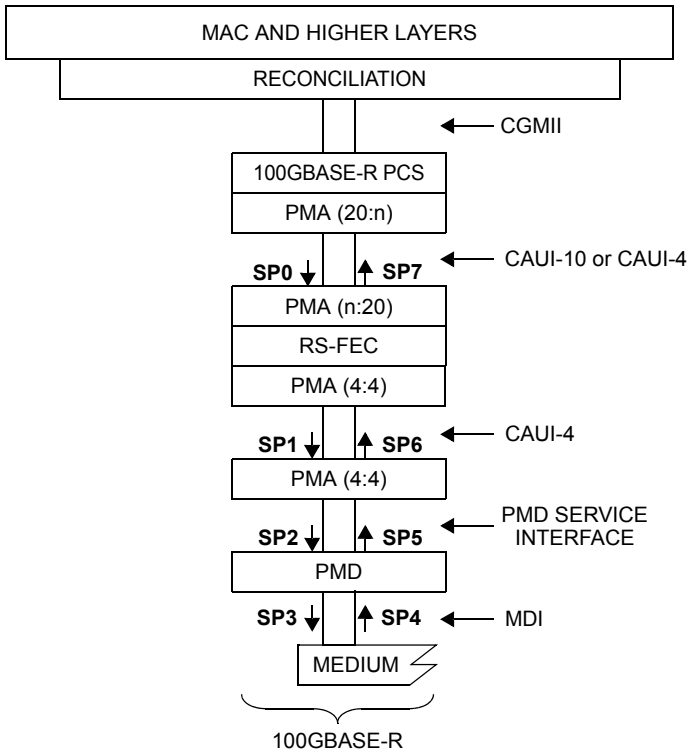

CAUI-4 = 100 Gb/s FOUR-LANE ATTACHMENT UNIT INTERFACE
CAUI-10 = 100 Gb/s TEN-LANE ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION
n = 4 or 10

Figure 80–8—100GBASE-R Skew points with RS-FEC and CAUI-n

The Skew requirements for the PCS, PMA and PMD sublayers are specified in the respective clauses as noted in Table 80–6 and Table 80–7.

Table 80–6—Summary of Skew constraints

Skew points	Maximum Skew (ns)[a]	Maximum Skew for 40GBASE-R PCS lane (UI)[b]	Maximum Skew for 100GBASE-R PCS lane (UI)[c]	Notes[d]
SP0	29	N/A	≈ 150	See 83.5.3.1
SP1	29	≈ 299	≈ 150	See 83.5.3.2
SP2	43	≈ 443	≈ 222	See 83.5.3.4, 84.5, 85.5, 86.3.2, 87.3.2, 88.3.2, 89.3.2, 92.5, 93.5, 94.3.4, or 95.3.2
SP3	54	≈ 557	≈ 278	See 84.5, 85.5, 86.3.2, 87.3.2, 88.3.2, 89.3.2, 92.5, 93.5, 94.3.4, or 95.3.2
SP4	134	≈ 1382	≈ 691	See 84.5, 85.5, 86.3.2, 87.3.2, 88.3.2, 89.3.2, 92.5, 93.5, 94.3.4, or 95.3.2
SP5	145	≈ 1495	≈ 748	See 84.5, 85.5, 86.3.2, 87.3.2, 88.3.2, 89.3.2, 92.5, 93.5, 94.3.4, or 95.3.2
SP6	160	≈ 1649	≈ 824	See 83.5.3.6
SP7	29	N/A	≈ 150	See 83.5.3.8
At PCS receive	180	≈ 1856	≈ 928	See 82.2.13
At RS-FEC transmit	49	N/A	≈ 253	See 91.5.2.2
At RS-FEC receive[e]	180	N/A	≈ 4641	See 91.5.3.1
At PCS receive (with RS-FEC)	49	N/A	≈ 253	See 82.2.13

[a]The Skew limit includes 1 ns allowance for PCB traces that are associated with the Skew points.
[b]The symbol ≈ indicates approximate equivalent of maximum Skew in UI for 40GBASE-R, based on 1 UI equals 96.969697 ps at PCS lane signaling rate of 10.3125 GBd.
[c]The symbol ≈ indicates approximate equivalent of maximum Skew in UI for 100GBASE-R, based on 1 UI equals 193.939394 ps at PCS lane signaling rate of 5.15625 GBd.
[d]Should there be a discrepancy between this table and the Skew requirements of the relevant sublayer clause, the sublayer clause prevails.
[e]The skew at the RS-FEC receive is the skew between RS-FEC lanes. The symbol ≈ indicates approximate equivalent of maximum Skew in UI for RS-FEC lanes with a signaling rate of 25.78125 GBd.

Table 80–7—Summary of Skew Variation constraints

Skew points	Maximum Skew Variation (ns)	Maximum Skew Variation for 10.3125 GBd PMD lane (UI)[a]	Maximum Skew Variation for 25.78125 GBd PMD lane (UI)[b]	Notes[c]
SP0	0.2	≈ 2	N/A	See 83.5.3.1
SP1	0.2	≈ 2	N/A	See 83.5.3.2
SP2	0.4	≈ 4	≈ 10	See 83.5.3.4, 84.5, 85.5, 86.3.2, 87.3.2, 88.3.2, 89.3.2, 92.5, 93.5, 94.3.4, or 95.3.2
SP3	0.6	≈ 6	≈ 15	See 84.5, 85.5, 86.3.2, 87.3.2, 88.3.2, 89.3.2, 92.5, 93.5, 94.3.4, or 95.3.2
SP4	3.4	≈ 35	≈ 88	See 84.5, 85.5, 86.3.2, 87.3.2, 88.3.2, 89.3.2, 92.5, 93.5, 94.3.4, or 95.3.2
SP5	3.6	≈ 37	≈ 93	See 84.5, 85.5, 86.3.2, 87.3.2, 88.3.2, 89.3.2, 92.5, 93.5, 94.3.4, or 95.3.2
SP6	3.8	≈ 39	≈ 98	See 83.5.3.6
SP7	0.2	≈ 2	N/A	See 83.5.3.8
At PCS receive	4	≈ 41	N/A	See 82.2.13
At RS-FEC transmit	0.4	N/A	≈ 10	See 91.5.2.2
At RS-FEC receive[d]	4	N/A	≈ 103	See 91.5.3.1
At PCS receive (with RS-FEC)	0.4	N/A	≈ 10	See 82.2.13

[a]The symbol ≈ indicates approximate equivalent of maximum Skew Variation in UI for 40GBASE-R, based on 1 UI equals 96.969697 ps at PMD lane signaling rate of 10.3125 GBd.
[b]The symbol ≈ indicates approximate equivalent of maximum Skew Variation in UI for 100GBASE-R, based on 1 UI equals 38.787879 ps at PMD lane signaling rate of 25.78125 GBd.
[c]Should there be a discrepancy between this table and the Skew requirements of the relevant sublayer clause, the sublayer clause prevails.
[d]The skew at the RS-FEC receive is the skew between RS-FEC lanes.

80.6 State diagrams

State diagrams take precedence over text.

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

Multiple states of a function that have a transition to a common state utilizing different qualifiers (for example, multiple exit conditions to an IDLE or WAIT state) may be indicated by a shared arrow. An exit transition arrow must connect to the shared arrow, and the qualifier must be met prior to termination of the transition arrow on the shared arrow. The shared arrow has no qualifier.

80.7 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, Clause 73, Clause 74, Clause 81 through Clause 89, Clause 91 through Clause 95, and related annexes 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 40 Gigabit and 100 Gigabit Ethernet PICS conforms to the same notation and conventions used in 21.6.

81. Reconciliation Sublayer (RS) and Media Independent Interface for 40 Gb/s and 100 Gb/s operation (XLGMII and CGMII)

81.1 Overview

This clause defines the characteristics of the Reconciliation Sublayer (RS) and the Media Independent Interface between Ethernet media access controllers and various PHYs. Figure 81–1 shows the relationship of the RS and Media Independent Interface to the ISO/IEC OSI reference model. Note that there are two variants of the Media Independent Interface in this clause, the 40 Gb/s Media Independent Interface (XLGMII) and the 100 Gb/s Media Independent Interface (CGMII).

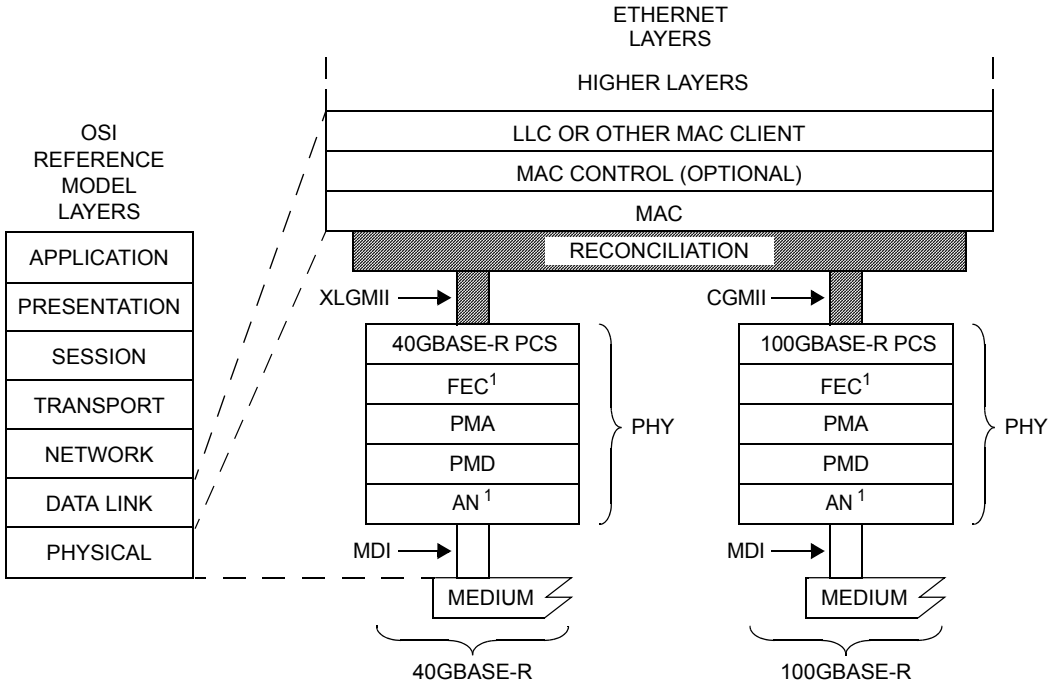

AN = AUTO-NEGOTIATION
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 81–1—RS and MII relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

The XLGMII and the CGMII are optional logical interfaces between the MAC sublayer and the Physical Layer (PHY).

The RS adapts the bit serial protocols of the MAC to the parallel format of the PCS service interface. Though the XLGMII/CGMII is an optional interface, it is used in this standard as a basis for specification. The Physical Coding Sublayer (PCS) is specified to the XLGMII/CGMII, so if not implemented, a conforming implementation shall behave functionally as if the RS and XLGMII/CGMII were implemented.

The XLGMII/CGMII has the following characteristics:

a) The XLGMII supports a speed of 40 Gb/s.
b) The CGMII supports a speed of 100 Gb/s.
c) Data and delimiters are synchronous to a clock reference.
d) It provides independent 64-bit wide transmit and receive data paths.
e) It supports full duplex operation only.

81.1.1 Summary of major concepts

The following are the major concepts of the XLGMII/CGMII:

a) The XLGMII/CGMII is functionally similar to other media independent interfaces that have been defined for lower speeds, 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 XLGMII/CGMII.
c) The RS maps the signal set provided at the XLGMII/CGMII 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) The XLGMII and CGMII may also support Low Power Idle (LPI) signaling for PHY types supporting Energy Efficient Ethernet (EEE) (see Clause 78).

81.1.2 Application

This clause applies to the interface between the MAC and PHY. This logical interface is used to provide media independence so that an identical media access controller may be used with supported PHY types.

81.1.3 Rate of operation

The XLGMII is specified to support 40 Gb/s operation and the CGMII is specified to support 100 Gb/s operation.

81.1.4 Delay constraints

The maximum cumulative MAC Control, MAC, and RS delay (sum of transmit and receive delays at one end of the link) shall meet the values specified in Table 81–1. A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

Table 81–1—Delay constraints

Sublayer	Maximum (bit time)	Maximum (pause_quanta)	Maximum (ns)
40 Gb/s MAC, RS, and MAC Control	16384	32	409.6
100 Gb/s MAC, RS, and MAC Control	24576	48	245.76

81.1.5 Allocation of functions

The allocation of functions at the XLGMII/CGMII balances the need for media independence with interface simplicity. The XLGMII and CGMII maximize media independence by separating the Data Link and Physical Layers of the OSI seven-layer reference model.

81.1.6 XLGMII/CGMII structure

The XLGMII/CGMII is composed of independent transmit and receive paths. Each direction uses 64 data signals (TXD<63:0> and RXD<63:0>), 8 control signals (TXC<7:0> and RXC<7:0>), and a clock (TX_CLK and RX_CLK). Figure 81–2 depicts a schematic view of the RS inputs and outputs.

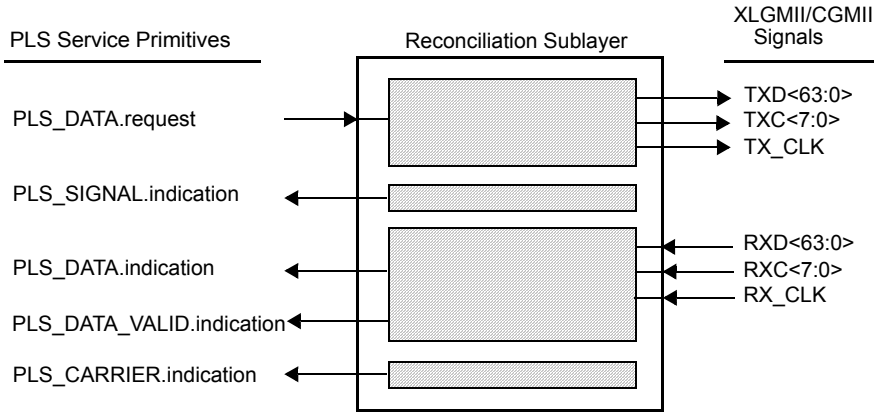

Figure 81–2—Reconciliation Sublayer (RS) inputs and outputs

The 64 TXD and 8 TXC signals shall be organized into eight data lanes, as shall the 64 RXD and 8 RXC signals (see Table 81–2). The eight lanes in each direction share a common clock, TX_CLK for transmit and RX_CLK for receive. The eight 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, the fifth to lane 4, the sixth to lane 5, the seventh to lane 6 and the eighth to lane 7, then repeating with the ninth 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 81–2—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
<39:32>	<4>	4
<47:40>	<5>	5
<55:48>	<6>	6
<63:56>	<7>	7

81.1.7 Mapping of XLGMII/CGMII signals to PLS service primitives

The Reconciliation Sublayer (RS) shall map the signals provided at the XLGMII/CGMII 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 40 Gb/s and 100 Gb/s; therefore, PLS service primitives supporting CSMA/CD operation are not mapped through the RS to the XLGMII/CGMII. This behavior and restrictions are the same as described in 22.7, with the details of the signaling described in 81.3. 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 40 Gb/s and 100 Gb/s operation:

PLS_DATA.request
PLS_DATA.indication
PLS_CARRIER.indication
PLS_SIGNAL.indication
PLS_DATA_VALID.indication

81.1.7.1 Mapping of PLS_DATA.request

81.1.7.1.1 Function

The RS maps the primitive PLS_DATA.request to the XLGMII/CGMII signals TXD<63:0>, TXC<7:0>, and TX_CLK.

81.1.7.1.2 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. One or zero represents a single data bit. The DATA_COMPLETE value signifies that the MAC sublayer has no more data to output.

81.1.7.1.3 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.

81.1.7.1.4 Effect of receipt

The OUTPUT_UNIT values are conveyed to the PHY by the signals TXD<63:0> and TXC<7:0> on each TX_CLK rising edge. Each PLS_DATA.request transaction shall be mapped to a TXD signal in sequence (TXD<0>, TXD<1>,... TXD<63>, TXD<0>) as described in 81.2, and for every eight transactions, a TXC signal in sequence (TXC<0>, TXC<1>,... TXC<7>, TXC<0>) is generated, as described in 81.3.1.2. After 64 PLS_DATA.request transactions from the MAC sublayer (eight octets of eight PLS_DATA.request transactions each), the RS requests transmission of 64 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<63:0> and TXC<7:0> shall be generated by the RS every 64 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 rising

edge. This may be on the same TX_CLK rising edge as the last data octet or the subsequent TX_CLK rising edge. When the Terminate control character is in lane 0, 1, 2, 3, 4, 5, or 6, the lanes following in sequence are encoded with an Idle control character.

81.1.7.2 Mapping of PLS_DATA.indication

81.1.7.2.1 Function

The RS maps the XLGMII/CGMII signals RXD<63:0>, RXC<7:0>, and RX_CLK to the primitive PLS_DATA.indication.

81.1.7.2.2 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.

81.1.7.2.3 When generated

The INPUT_UNIT values are derived from the signals RXC<7:0> and RXD<63:0> received from the PHY on each rising 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<63:0> during frame reception, the RS shall generate 64 PLS_DATA.indication transactions until the end of frame (Terminate control character), where 0, 8, 16, 24, 32, 40, 48, or 56 PLS_DATA.indication transactions will be generated from the RXD<63:0> 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<63>, RXD<0>) as described in 81.2.

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 XLGMII/CGMII error indications (see 81.3.3). Sequence ordered sets are not indicated to the MAC (see 81.3.4).

81.1.7.2.4 Effect of receipt

The effect of receipt of this primitive by the MAC sublayer is unspecified.

81.1.7.3 Mapping of PLS_CARRIER.indication

40 Gb/s and 100 Gb/s operation supports full duplex operation only. The RS never generates this primitive for PHYs that do not support EEE.

For PHYs that support EEE capability, CARRIER_STATUS is set in response to LPI_REQUEST as shown in Figure 81–13.

81.1.7.4 Mapping of PLS_SIGNAL.indication

40 Gb/s and 100 Gb/s operation supports full duplex operation only. The RS never generates this primitive.

81.1.7.5 Mapping of PLS_DATA_VALID.indication

81.1.7.5.1 Function

The RS maps the XLGMII/CGMII signals RXC<7:0> and RXD<63:0> to the primitive PLS_DATA_VALID.indication.

81.1.7.5.2 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 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.

81.1.7.5.3 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<7:0> and RXD<63:0> contained eight Idle characters, a Sequence ordered set, or a Terminate character. 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 81.3.3.1).

81.1.7.5.4 Effect of receipt

The effect of receipt of this primitive by the MAC sublayer is unspecified.

81.2 XLGMII/CGMII data stream

Packets transmitted through the XLGMII/CGMII shall be transferred within the XLGMII/CGMII data stream. A 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 81–3.

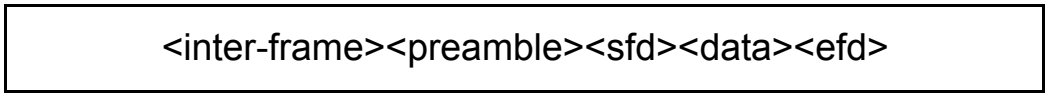

Figure 81–3—XLGMII/CGMII data stream

For the XLGMII/CGMII, transmission and reception of each bit and mapping of data octets to lanes shall be as shown in Figure 81–4.

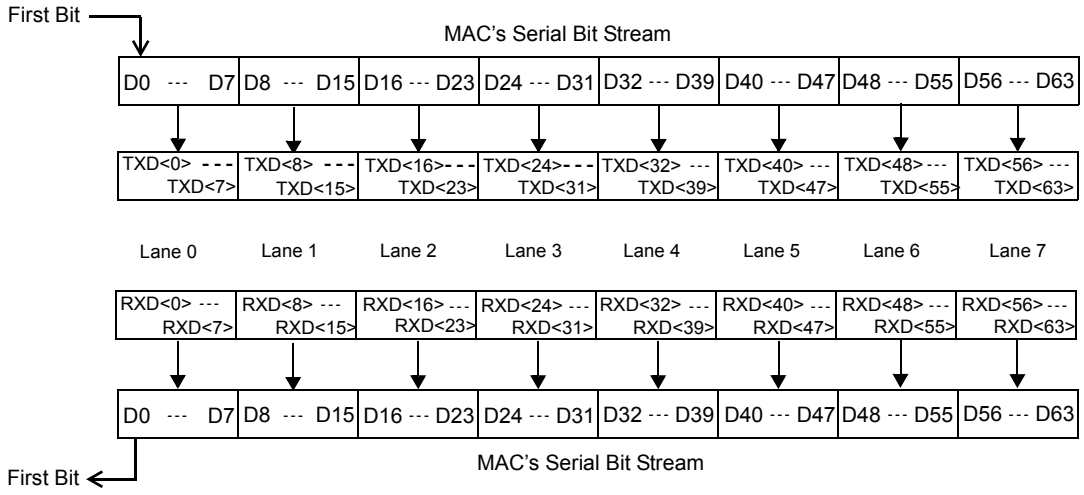

Figure 81–4—Relationship of data lanes to MAC serial bit stream

81.2.1 Inter-frame <inter-frame>

The inter-frame <inter-frame> period on an XLGMII/CGMII 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 or PHY clock rate compensation). The minimum IPG at the XLGMII/CGMII of the receiving RS is one octet.

The signaling of link status information logically occurs in the <inter-frame> period (see 81.3.4). Frame processing when signaling of link status information is initiated or terminated is described in 81.3.3.

81.2.2 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 XLGMII/CGMII. On transmit, the RS converts the first data octet of preamble transferred from the MAC into a Start control character. On receive, the RS converts the Start control character into a preamble data octet. The Start control character is aligned to lane 0 of the XLGMII/CGMII 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 XLGMII/CGMII is the same as the Start Frame Delimiter (SFD) specified in 4.2.6 and is equal to the following:

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 XLGMII/CGMII as octets sequentially ordered on the lanes of the XLGMII/CGMII. 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, the fourth on lane 3, the fifth on lane 4, the sixth on lane 5, the seventh on lane 6, and the SFD on lane 7, and the eight octets are transferred on the next rising edge of TX_CLK. The ninth octet is assigned to lane 0 with subsequent octets sequentially assigned to the lanes. The XLGMII/CGMII <preamble> and <sfd> are as follows:

Lane 0	**Lane 1**	**Lane 2**	**Lane 3**	**Lane 4**	**Lane 5**	**Lane 6**	**Lane 7**
Start	10101010	10101010	10101010	10101010	10101010	10101010	10101011

81.2.3 Data <data>

The data <data> in a frame shall consist of a set of data octets.

81.2.4 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 XLGMII/CGMII shall recognize the end of frame delimiter on any of the eight lanes of the XLGMII/CGMII.

81.2.5 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 81.1.7.5.2 and 30.3.2.1.5).

81.3 XLGMII/CGMII functional specifications

The XLGMII/CGMII 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 XLGMII/CGMII loopback is defined, but XLGMII/CGMII 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<63> to RXD<63>, TXC<0> to RXC<0> ... TXC<7> to RXC<7>, 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.

81.3.1 Transmit

81.3.1.1 TX_CLK

TX_CLK is a continuous clock used for operation at the appropriate frequency. TX_CLK provides the timing reference for the transfer of the TXC<7:0> and TXD<63:0> signals from the RS to the PHY. The values of TXC<7:0> and TXD<63:0> shall be sampled by the PHY on the rising edge of TX_CLK. TX_CLK is sourced by the RS.

The TX_CLK frequency shall be one-sixty-fourth of the MAC transmit data rate.

81.3.1.2 TXC<7:0> (transmit control)

TXC<7:0> indicate that the RS is presenting either data or control characters on the XLGMII/CGMII for transmission. The TXC signal for a lane shall be deasserted 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 deasserted by the RS for each octet of the preamble (except the first octet that is replaced with a Start control character) and remain deasserted while all octets to be transmitted are presented on the lanes of the XLGMII/CGMII. TXC<7:0> are driven by the RS and shall transition synchronously with respect to the rising edge of TX_CLK. Table 81–3 specifies the permissible encodings of TXD and TXC for an XLGMII/CGMII 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).

Table 81–3—Permissible encodings of TXC and TXD

TXC	TXD	Description	PLS_DATA.request parameter
0	0x00 through 0xFF	Normal data transmission	Zero, one (eight bits)
1	0x00 through 0x05	Reserved	—
1	0x06	Only valid on all 8 lanes simultaneously to request LPI	No applicable parameter (normal inter-frame)
1	0x07	Idle	No applicable parameter (normal inter-frame)
1	0x08 through 0x9B	Reserved	—
1	0x9C	Sequence (only valid in lane 0)	No applicable parameter (inter-frame status signal)
1	0x9D through 0xFA	Reserved	—
1	0xFB	Start (only valid in lane 0)	No applicable parameter, replaces first eight zero, one of a frame (preamble octet)
1	0xFC	Reserved	—
1	0xFD	Terminate	DATA_COMPLETE
1	0xFE	Error	No applicable parameter
1	0xFF	Reserved	—

A PHY with EEE capability shall interpret the combination of TXC and TXD as shown in Table 81–3 as an assertion of LPI. Transition into and out of the LPI state is shown in Figure 81–7.

81.3.1.3 TXD<63:0> (transmit data)

TXD is a bundle of 64 data signals organized into eight lanes of eight signals each (TXD<7:0>, TXD<15:8>, TXD<23:16>, TXD<31:24>, TXD<39:32>, TXD<47:40>, TXD<55:48>, and TXD<63:56>) that are driven by the RS. Each lane is associated with a TXC signal as shown in Table 81–2 and shall be encoded as shown in Table 81–3. TXD<63:0> shall transition synchronously with respect to the rising edge of TX_CLK. For each high TX_CLK transition, data and/or control are presented on TXD<63:0> 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, TXD<24> the least significant bit of lane 3, TXD<32> the least significant bit of lane 4, TXD<40> the least significant bit of lane 5, TXD<48> the least significant bit of lane 6, and TXD<56> the least significant bit of lane 7.

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 deasserted, 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 81–5.

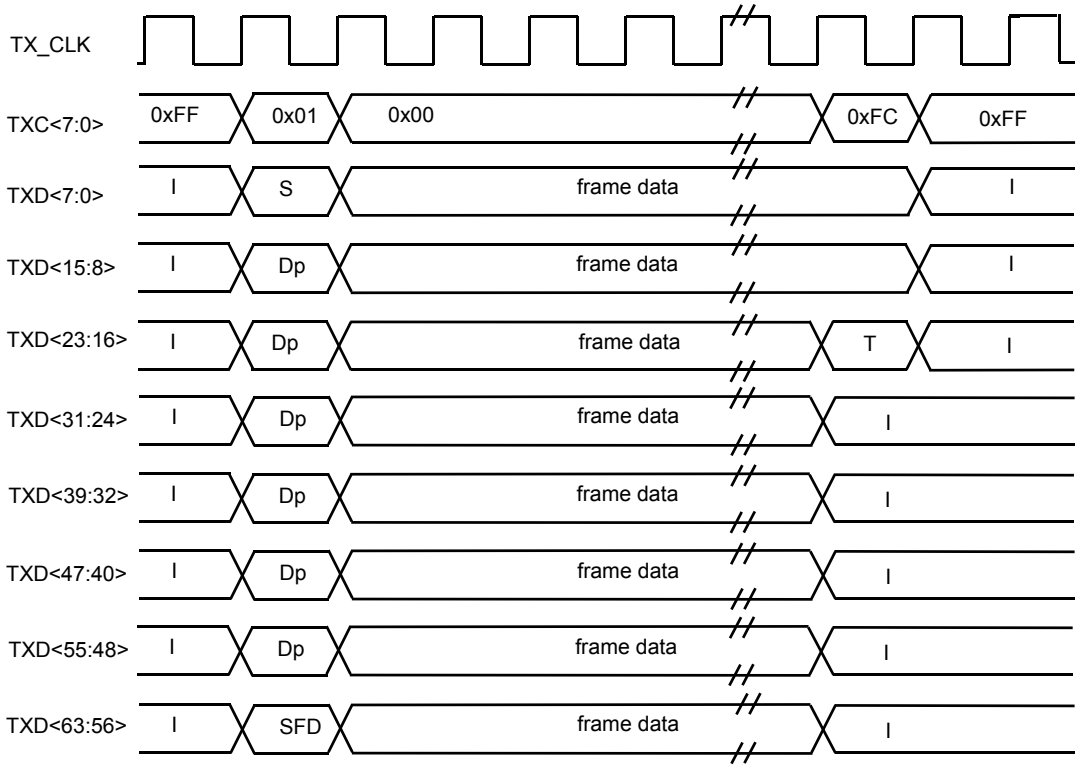

I: Idle control character, S: Start control character, Dp: preamble Data octet, T: Terminate control character, SFD: Start of Frame Delimiter

Figure 81–5—Normal frame transmission

Figure 81–6 shows the behavior of TXD and TXC during an example transmission of a frame propagating an error.

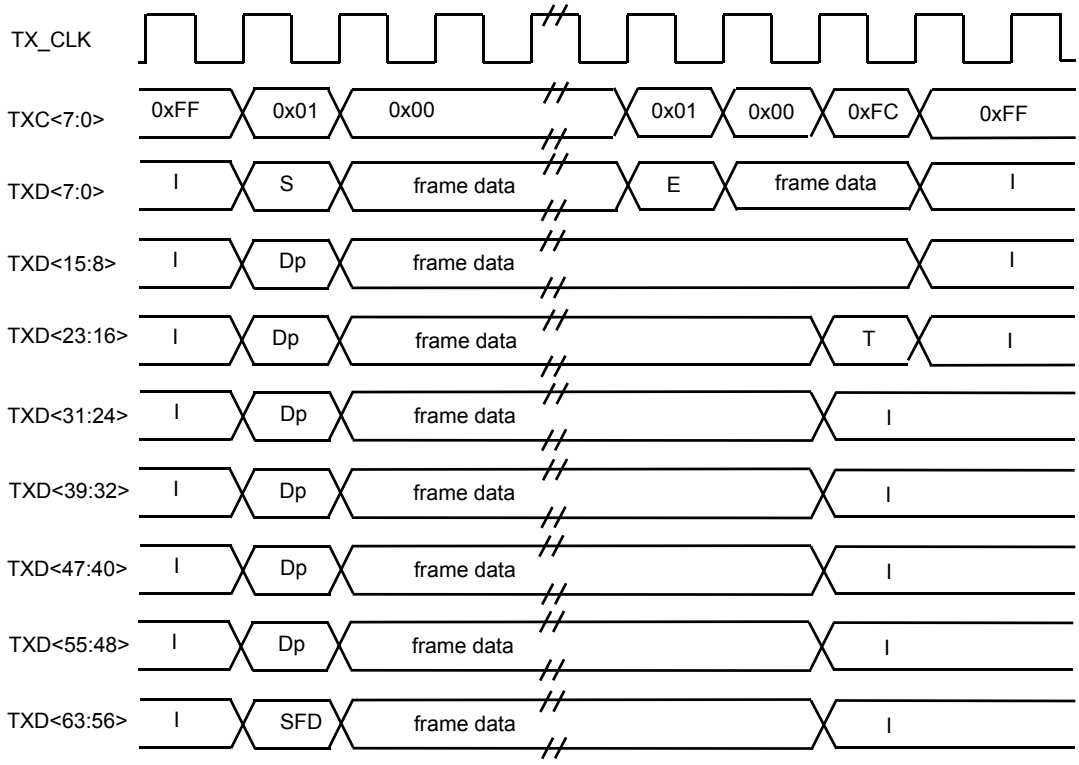

I: Idle control character, S: Start control character, Dp: preamble Data octet, T: Terminate control character, E: Error control character, SFD: Start of Frame Delimiter

Figure 81–6—Transmit Error Propagation

81.3.1.4 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:

a) A MAC implementation may incorporate this RS function into its design and always insert additional idle characters to align the start of preamble on an eight byte boundary. Note that this will reduce the effective data rate for certain packet sizes separated with minimum inter-frame spacing.

b) 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 seven. Note that this may result in inter-frame spacing observed on the transmit XLGMII/CGMII that is up to seven 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 the DIC as described.

81.3.1.5 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 asserts TXC and asserts IDLE on lanes 0 to 7 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 ($T_{w\_sys\_tx}$). The wake times are shown in Table 78–4

Figure 81–7 shows the behavior of TXC and TXD<7:0> during the transition into and out of the LPI state.

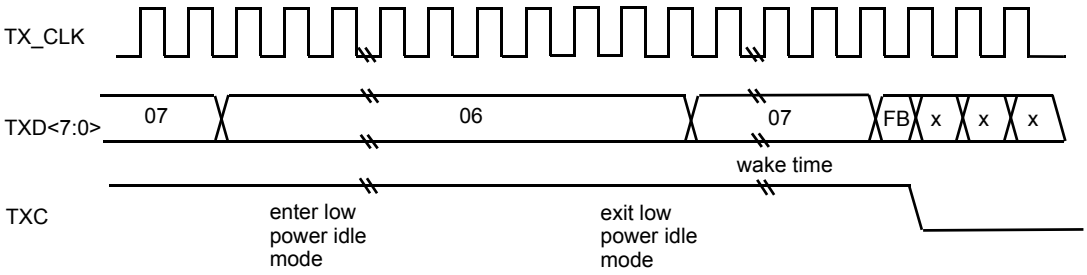

NOTE—TXC and TXD are shown for one lane, all 8 lanes behave identically during LPI.

Figure 81–7—LPI transition

Table 81–3 summarizes the permissible encodings of TXD<63:0>, TXC<7:0>.

81.3.2 Receive

81.3.2.1 RX_CLK (receive clock)

RX_CLK is a continuous clock that provides the timing reference for the transfer of the RXC<7:0> and RXD<63:0> signals from the PHY to the RS. RXC<7:0> and RXD<63:0> shall be sampled by the RS on the rising 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 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.

NOTE—This standard neither requires nor assumes a guaranteed phase relationship between the RX_CLK and TX_CLK signals.

81.3.2.2 RXC<7:0> (receive control)

RXC<7:0> indicate that the PHY is presenting either recovered and decoded data or control characters on the XLGMII/CGMII. The RXC signal for a lane shall be deasserted 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 deasserted by the PHY for each octet of the preamble (except the first octet that is replaced with a Start control character) and remain deasserted while all octets to be received are presented on the lanes of the XLGMII/CGMII. RXC<7:0> are driven by the PHY and shall transition synchronously with respect to the rising edge of RX_CLK. Table 81–4 specifies the permissible encodings of RXD and RXC for an XLGMII/CGMII 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).

Table 81–4—Permissible lane encodings of RXD and RXC

RXC	RXD	Description	PLS_DATA.indication parameter
0	0x00 through 0xFF	Normal data reception	Zero, one (eight bits)
1	0x00 through 0x05	Reserved	—
1	0x06	Only valid on all 8 lanes simultaneously to indicate LP_IDLE is asserted	No applicable parameter (normal inter-frame)
1	0x07	Idle	No applicable parameter (Normal inter-frame)
1	0x08 through 0x9B	Reserved	—
1	0x9C	Sequence (only valid in lane 0)	No applicable parameter (Inter-frame status signal)
1	0x9D through 0xFA	Reserved	—
1	0xFB	Start (only valid in lane 0)	No applicable parameter, first eight zero, one of a frame (a preamble octet)
1	0xFC	Reserved	—
1	0xFD	Terminate	No applicable parameter (start of inter-frame)
1	0xFE	Error	No applicable parameter
1	0xFF	Reserved	—

Figure 81–8 shows the behavior of RXC<7:0> during an example frame reception with no errors.

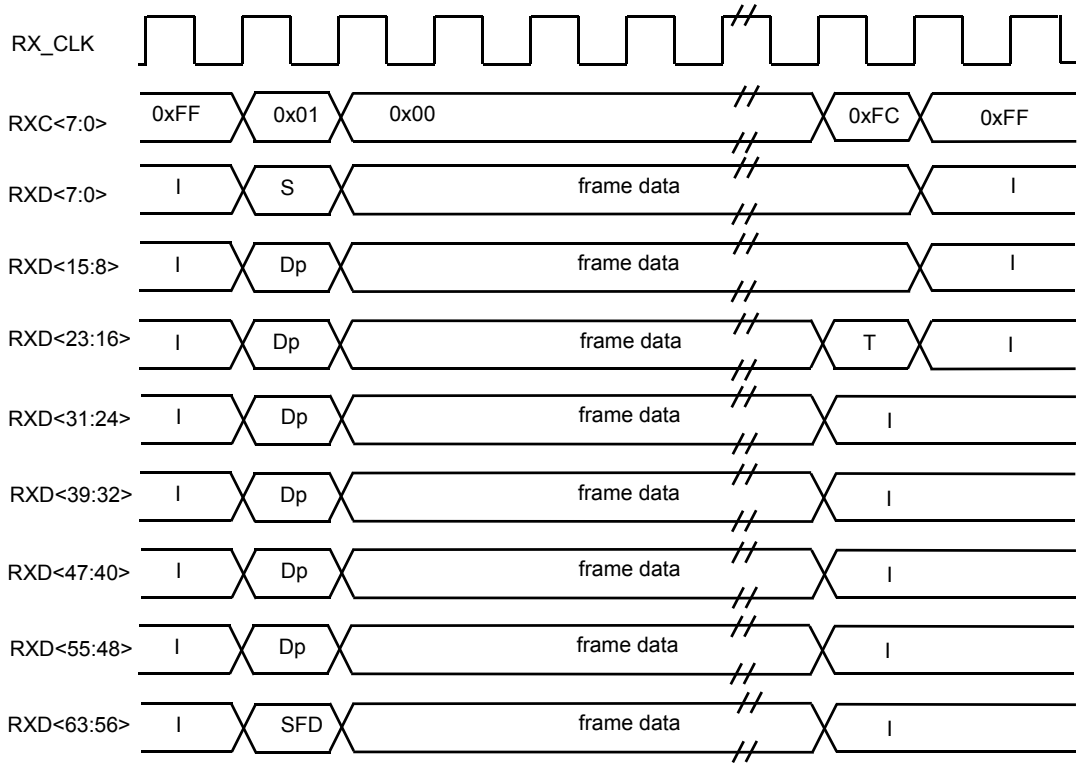

I: Idle control character, S: Start control character, Dp: preamble Data octet, T: Terminate control character, SFD: Start of Frame Delimiter

Figure 81–8—Frame reception without error

81.3.2.3 RXD (receive data)

RXD is a bundle of 64 data signals (RXD<63:0>) organized into eight lanes of eight signals each (RXD<7:0>, RXD<15:8>, RXD<23:16>, RXD<31:24>, RXD<39:32>, RXD<47:40>, RXD<55:48>, and RXD<63:56>) that are driven by the PHY. Each lane is associated with an RXC signal as shown in Table 81–2 and shall be decoded by the RS as shown in Table 81–4. RXD<63:0> shall transition synchronously with respect to the rising edge of RX_CLK. For each rising RX_CLK transition, received data and/or control are presented on RXD<63:0> 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, RXD<24> the least significant bit of lane 3, RXD<24> the least significant bit of lane 3, RXD<32> the least significant bit of lane 4, RXD<40> the least significant bit of lane 5, RXD<48> the least significant bit of lane 6, and RXD<56> the least significant bit of lane 7. Figure 81–8 shows the behavior of RXD<63:0> during frame reception.

While the RXC of a lane is deasserted, RXD of the lane is used by the RS to generate PLS_DATA.indications. 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 that 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 81.3.3.1. Figure 81–9 shows the behavior of RXC and RXD during the reception of an example frame with an error.

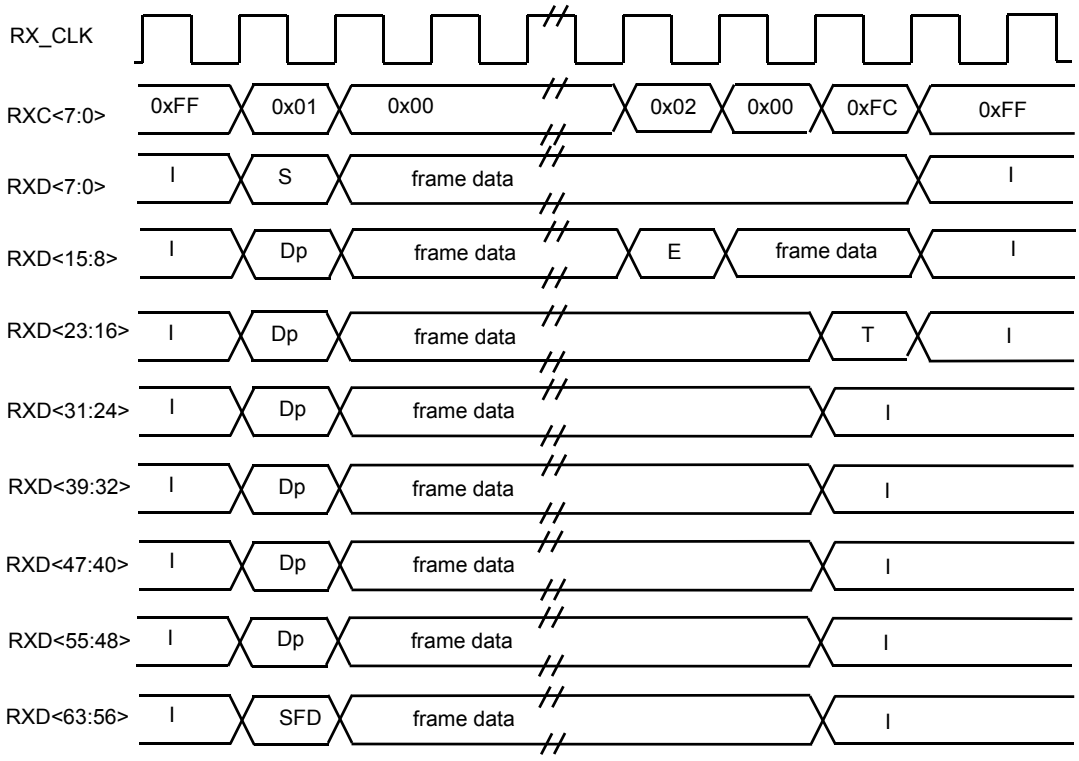

I: Idle control character, S: Start control character, Dp: preamble Data octet, T: Terminate control character, E: Error control character, SFD: Start of Frame Delimiter

Figure 81–9—Reception with error

81.3.2.4 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 all lanes 0 to 7 to return to a normal interframe state. The RS shall interpret the LPI coding as shown in Table 81–4.

Figure 81–10 shows the behavior of RXC and RXD<7:0> during LPI transitions.

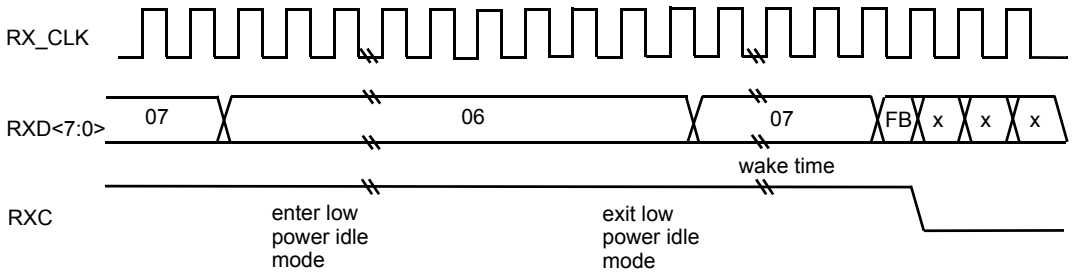

NOTE 1—RXC and RXD are shown for one lane, all 8 lanes behave identically during LPI.
NOTE 2—In some instances, LPI may be followed by characters other than IDLE during wake time.

Figure 81–10—LPI transition

81.3.3 Error and fault handling

81.3.3.1 Response to error indications by the XLGMII/CGMII

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.

81.3.3.2 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.

81.3.3.3 Response to received invalid frame sequences

The RS shall not indicate DATA_VALID to the MAC for a Start control character received on any lane other than lane 0. Error free operation will not change the SFD alignment in lane 7. A MAC/RS implementation is not required to process a packet that has an SFD in a position other than lane 7 in the XLGMII/CGMII transfer containing the Start control character.

81.3.4 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. The behavior of the fault signaling is the same as it is for Clause 46 with the exception that the ordered sets are aligned to eight byte boundaries, padding lanes 4 to 7 with 0x00.

Clause 46 uses the term column when describing data transfers on the XGMII. The eight lanes of data and control transferred per clock cycle on XLGMII/CGMII are equivalent to a column in the following description of link fault signaling.

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 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. Note that this behavior only supports bidirectional operation.

Status is signaled in an eight byte Sequence ordered set as shown in Table 81–5. The PHY indicates Local Fault with a Sequence control character in lane 0 and data characters of 0x00 in lanes 1, 2, 4, 5, 6, and 7 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, 2, 4, 5, 6, and 7 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. All other values in lanes 1 to 3 not shown in Table 81–5 are reserved. The link fault signaling state diagram allows future standardization of reserved Sequence ordered sets for functions other than link fault indications.

Table 81–5—Sequence ordered sets

Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Description
Sequence	0x00	0x00	0x00	0x00	0x00	0x00	0x00	Reserved
Sequence	0x00	0x00	0x01	0x00	0x00	0x00	0x00	Local Fault
Sequence	0x00	0x00	0x02	0x00	0x00	0x00	0x00	Remote Fault

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 81–11).

81.3.4.1 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 the rising clock transitions) unless reset.

fault_sequence
: A new column received on RXC<7:0> and RXD<63:0> comprising a Sequence ordered set of

eight bytes and consisting of a Sequence control character in lane 0 and a seq_type in lanes 1, 2, 3, 4, 5, 6, and 7 indicating either Local Fault or Remote Fault.

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, 0x00 in lane 4, 0x00 in lane 5, 0x00 in lane 6, 0x00 in lane 7.
Remote Fault; 0x00 in lane 1, 0x00 in lane 2, 0x02 in lane 3, 0x00 in lane 4, 0x00 in lane 5, 0x00 in lane 6, 0x00 in lane 7.

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.

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, 0x00 in lane 4, 0x00 in lane 5, 0x00 in lane 6, 0x00 in lane 7.
Remote Fault; 0x00 in lane 1, 0x00 in lane 2, 0x02 in lane 3, 0x00 in lane 4, 0x00 in lane 5, 0x00 in lane 6, 0x00 in lane 7.

81.3.4.2 State diagram

The Link Fault Signaling state diagram specifies the RS monitoring of RXC<7:0> and RXD<63:0> 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 fault sequence separated by less than 128 columns and no intervening fault_sequences of a different fault value. The state diagram is shown in Figure 81–11.

The variable link_fault is set to OK following any interval of 128 columns not containing a Remote Fault or Local Fault Sequence ordered set.

The RS output onto TXC<7:0> and TXD<63:0> 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
The RS shall continuously generate Idle control characters.

81.4 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

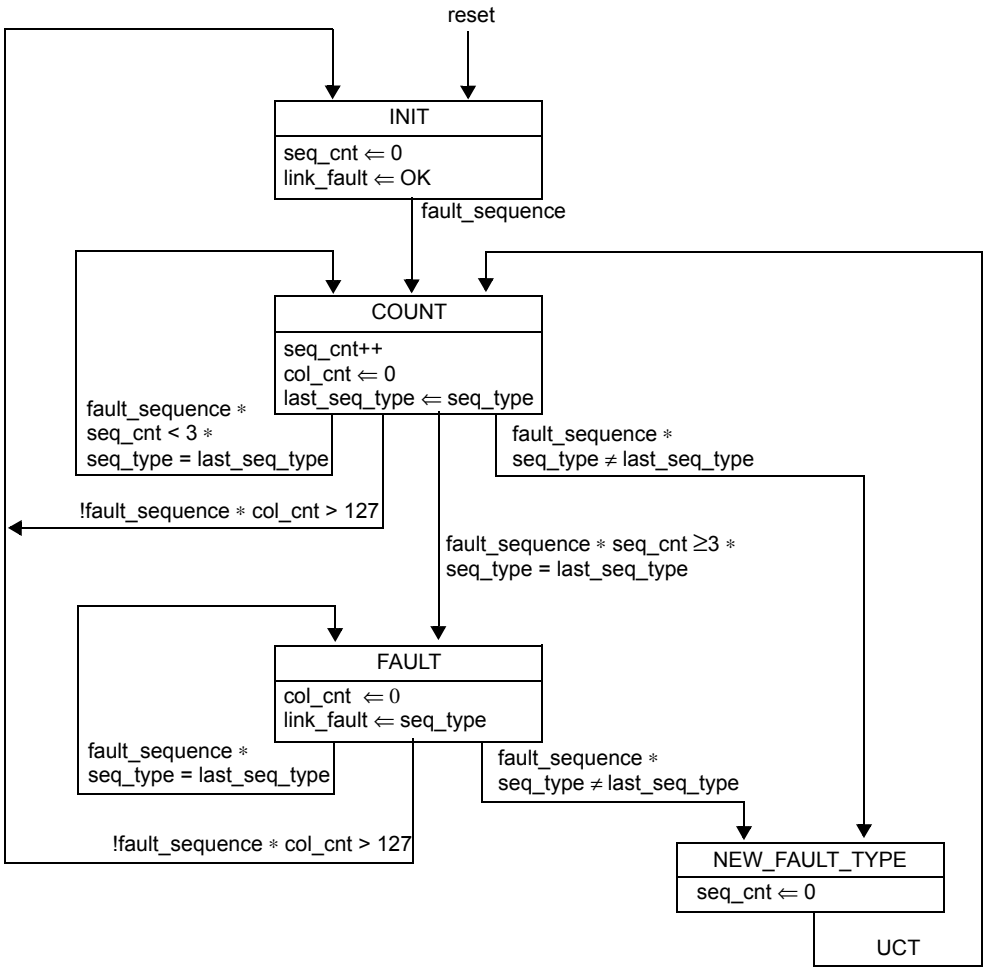

Figure 81–11—Link Fault Signaling state diagram

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 XLGMII and CGMII signals as shown in Figure 81–12.

The definition of TXC<7:0> and TXD<63:0> is derived from the state of PLS_DATA.request (81.1.7), except when it is overridden by an assertion of Remote Fault or LP_IDLE.request.

Similarly, RXC<7:0> and RXD<63:0> are mapped to PLS_DATA.indication except when LP_IDLE is detected.

PLS_CARRIER.indication(CARRIER_STATUS) is set to CARRIER_ON when the link is in LPI mode. See 81.1.7.3.

The timing of PLS_CARRIER.indication when used for the LPI function is controlled by the LPI transmit state diagram.

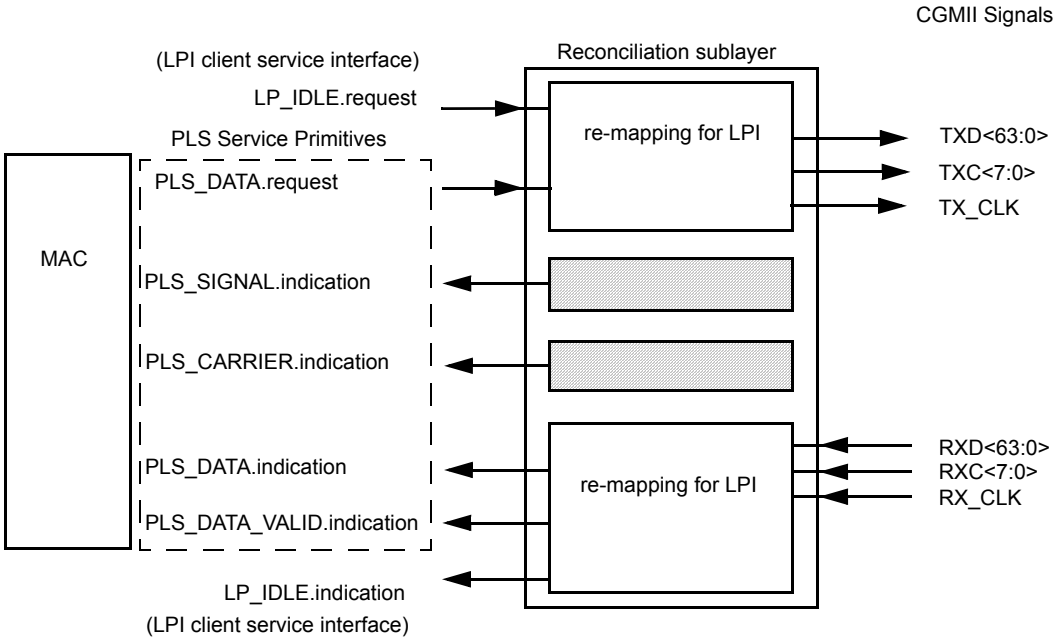

Figure 81–12—LPI assertion and detection mechanism

81.4.1 LPI messages

LP_IDLE.indication(LPI_INDICATION)

A primitive that indicates to the LPI client that the PHY has detected the assertion or deassertion of LPI from the link partner.

Values: DEASSERT: The link partner is operating with normal inter-frame 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 DEASSERT. ASSERT initiates the signaling of LPI to the link partner. DEASSERT 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 73.9.1.1) or within 1 s of the change of link_status to OK.

81.4.2 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 deasserted 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 PLS_CARRIER.indication(CARRIER_STATUS). The implementation shall conform to the behavior described by the transmit LPI state diagram shown in Figure 81–13.

81.4.2.1 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 81–13.

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.

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 the time since the deassertion of LPI. The terminal count of the timer is the value of the resolved $T_{w\_sys\_tx}$ as defined in 78.2. If PMA Ingress AUI Stop Enable (PIASE) bit (1.7.9) is asserted for any of the PMA sublayers, 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 XLAUI and CAUI-n as shown in Table 78–4 for each PMA with PIASE to be asserted. The signal tw_timer_done is asserted when tw_timer reaches its terminal count.

81.4.2.2 State Diagram

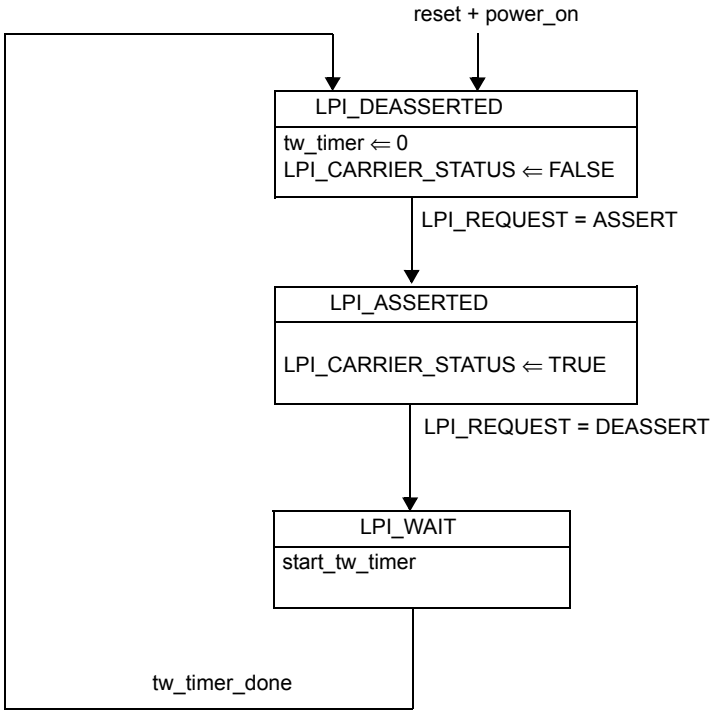

Figure 81–13—Transmit LPI State Diagram

81.4.3 Considerations for transmit system behavior

The transmit system should expect that egress data flow is halted for at least resolved $T_{w\_sys\_tx}$ (see 78.2) time after it requests the deassertion of LPI. Buffering and queue management should be designed to accommodate this behavior.

81.4.4 Considerations for receive system behavior

The mapping function of the Reconciliation Sublayer shall continue to signal DATA_NOT_VALID on PLS_DATA_VALID.indication while it is detecting LP_IDLE on the XLGMII and CGMII. The receive system should be aware that data frames may arrive at the XLGMII and CGMII following the deassertion of LPI_INDICATION with a delay corresponding to the link partner's resolved $T_{w\_sys\_rx}$ (as specified in 78.5) time.

If the PMA Ingress AUI Stop Enable (PIASE) bit (1.7.9) is asserted for any of the PMA sublayers, the PMA may stop signaling on the XLAUI and CAUI-n in the receive direction to conserve energy. The receiver should negotiate an additional time for the remote T_{w_sys} (equal to $T_{w\_sys\_tx}$–$T_{w\_sys\_rx}$ for the XLAUI and CAUI-n as shown in Table 78–4) for each PMA with PIASE to be asserted before setting the PIASE bits.

81.5 Protocol implementation conformance statement (PICS) proforma for Clause 81, Reconciliation Sublayer (RS) and Media Independent Interface for 40 Gb/s and 100 Gb/s operation[4]

81.5.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 81, Reconciliation Sublayer (RS) and Media Independent Interface for 40 Gb/s and 100 Gb/s operation, 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.

81.5.2 Identification

81.5.2.1 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).	

81.5.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 81, Reconciliation Sublayer (RS) and Media Independent Interface for 40 Gb/s and 100 Gb/s operation
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	

[4]*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.

81.5.2.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*PHY40	PHY support of XLGMII	81.2, 81.3		O	Yes [ ] No [ ]
*PHY100	PHY support of CGMII	81.2, 81.3		O	Yes [ ] No [ ]
*RS40	Reconciliation Sublayer support of XLGMII	81.2, 81.3		O	Yes [ ] No [ ]
*RS100	Reconciliation Sublayer support of CGMII	81.2, 81.3		O	Yes [ ] No [ ]
*LPI	Implementation of LPI	81.1.7		O	Yes [ ] No [ ]

81.5.3 PICS proforma tables for Reconciliation Sublayer and Media Independent Interface for 40 Gb/s and 100 Gb/s operation

81.5.3.1 General

Item	Feature	Subclause	Value/Comment	Status	Support
G1	PHY support of MAC data rate	81.1.3	Support MAC data rate of 40 Gb/s	PHY40:M	Yes [ ] N/A [ ]
G2	PHY support of MAC data rate	81.1.3	Support MAC data rate of 100 Gb/s	PHY100:M	Yes [ ] N/A [ ]
G3	Cumulative MAC Control, MAC and RS delay	81.1.4	Per Table 81–1 for 40 Gb/s	RS40:M	Yes [ ] N/A [ ]
G4	Cumulative MAC Control, MAC and RS delay	81.1.4	Per Table 81–1 for 100 Gb/s	RS100:M	Yes [ ] N/A [ ]
G5	Lane structure	81.1.6	Per Table 81–2	M	Yes [ ]

81.5.3.2 Mapping of PLS service primitives

Item	Feature	Subclause	Value/Comment	Status	Support
PL1	Mapping to Clause 6	81.1.7	RS implements mapping to Clause 6 PLS service primitives	RS:M	Yes [ ] N/A [ ]
PL2	Mapping of PLS_-DATA.requests	81.1.7.1.4	In sequence TXD<0> to TXD<63>	RS:M	Yes [ ] N/A [ ]
PL3	Start control character creation	81.1.7.1.4	First octet of preamble converted to Start control character	RS:M	Yes [ ] N/A [ ]
PL4	TXD and TXC generation	81.1.7.1.4	For each 64 PLS_-DATA.requests	RS:M	Yes [ ] N/A [ ]
PL5	Terminate control character creation	81.1.7.1.4	DATA_COMPLETE causes creation of Terminate control character in next lane in sequence	RS:M	Yes [ ] N/A [ ]
PL6	Mapping RXD to PLS_-DATA.indications	81.1.7.2.3	Create PLS_DATA.indications in sequence from RXD<0> to RXD<63>	RS:M	Yes [ ] N/A [ ]
PL7	PLS_DATA.indication generation	81.1.7.2.3	Generate 64 PLS_DATA.indications for each RXD<63:0> until Terminate then generating 0, 8, 16, 24, 32, 40, 48, or 56	RS:M	Yes [ ] N/A [ ]
PL8	Start control character conversion	81.1.7.2.3	Convert valid Start control character to preamble before generating PLS_-DATA.indications	RS:M	Yes [ ] N/A [ ]
PL9	Terminate control character	81.1.7.2.3	No PLS_DATA.indications generated	RS:M	Yes [ ] N/A [ ]
PL10	PLS_DATA_VALID.indication generation	81.1.7.5.3	On change of value of DATA_VALID_STATUS	RS:M	Yes [ ] N/A [ ]
PL11	DATA_VALID_STATUS	81.1.7.5.3	Value of DATA_VALID on a lane 0 Start control character preceded by eight idles, a Sequence ordered set, or a Terminate character	RS:M	Yes [ ] N/A [ ]
PL12	DATA_VALID_STATUS	81.1.7.5.3	Value of DATA_NOT_VALID on any control character but Error	RS:M	Yes [ ] N/A [ ]
PL13	Frame not ending with Terminate control character	81.1.7.5.3	Ensure MAC detects CRC error	RS:M	Yes [ ] N/A [ ]

81.5.3.3 Data stream structure

Item	Feature	Subclause	Value/Comment	Status	Support
DS1	Frame transfer	81.2	Within XLGMII/CGMII data stream	RS:M	Yes [ ] N/A [ ]
DS2	Bit mapping	81.2	Per Figure 81–4	RS:M	Yes [ ] N/A [ ]
DS3	Content of <data>	81.2.3	Consist of data octets	RS:M	Yes [ ] N/A [ ]
DS4	Recognition of <efd>	81.2.4	Terminate recognized in any lane	RS:M	Yes [ ] N/A [ ]

81.5.3.4 XLGMII/CGMII signal functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
FS1	TX_CLK active edge	81.3.1.1	TXD and TXC sampled on the rising edge of TX_CLK	XGE:M	Yes [ ] N/A [ ]
FS2	TX_CLK frequency	81.3.1.1	One-sixty-fourth of the MAC transmit data rate	XGE:M	Yes [ ] N/A [ ]
FS3	TXC assertion and deassertion	81.3.1.2	Deasserted for data, asserted for control character	RS:M	Yes [ ] N/A [ ]
FS4	TXC clock	81.3.1.2	Synchronous to TX_CLK	XGE:M	Yes [ ] N/A [ ]
FS5	TXD encoding	81.3.1.3	Per Table 81–3	RS:M	Yes [ ] N/A [ ]
FS6	TXD clock	81.3.1.3	Synchronous to TX_CLK	XGE:M	Yes [ ] N/A [ ]
FS7	Start alignment	81.3.1.4	Start control character aligned to lane 0	RS:M	Yes [ ] N/A [ ]
FS8	RX_CLK active edge	81.3.2.1	RXD and RXC sampled on the rising edge of RX_CLK	XGE:M	Yes [ ] N/A [ ]
FS9	RX_CLK frequency	81.3.2.1	One-sixty-fourth of the MAC receive data rate	XGE:M	Yes [ ] N/A [ ]
FS10	Loss of receive signal	81.3.2.1	Source RX_CLK from nominal clock	PHY:M	Yes [ ] N/A [ ]
FS11	RXC assertion and deassertion	81.3.2.2	Deasserted for data, asserted for control character	PHY:M	Yes [ ] N/A [ ]
FS12	RXC clock	81.3.2.2	Synchronous to RX_CLK	XGE:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
FS13	RXD decoding	81.3.2.3	Per Table 81–4	RS:M	Yes [ ] N/A [ ]
FS14	RXD clock	81.3.2.3	Synchronous to RX_CLK	XGE:M	Yes [ ] N/A [ ]
FS15	Received Error control character	81.3.3.1	RS cause MAC FrameCheck-Error	RS:M	Yes [ ] N/A [ ]
FS16	DATA_VALID assertion	81.3.3.3	RS not assert DATA_VALID unless Start control character in lane 0	RS:M	Yes [ ] N/A [ ]

81.5.3.5 Link fault signaling state diagram

Item	Feature	Subclause	Value/Comment	Status	Support
LF1	Link fault signaling state diagram	81.3.4	Implement per Figure 81–11	RS:M	Yes [ ] N/A [ ]
LF2	link_fault = OK and MAC frames	81.3.4.2	RS services MAC frame transmission requests	RS:M	Yes [ ] N/A [ ]
LF3	link_fault = OK and no MAC frames	81.3.4.2	In absence of MAC frames, RS transmits Idle control characters	RS:M	Yes [ ] N/A [ ]
LF4	link_fault = Local Fault	81.3.4.2	RS transmits continuous Remote Fault Sequence ordered sets	RS:M	Yes [ ] N/A [ ]
LF5	link_fault = Remote Fault	81.3.4.2	RS transmits continuous Idle control characters	RS:M	Yes [ ] N/A [ ]

81.5.3.6 LPI functions

Item	Feature	Subclause	Value/Comment	Status	Support
L1	Assertion of LPI in Tx direction	81.3.1.2	As defined in Table 81–3	LPI:M	Yes [ ] N/A [ ]
L2	Assertion of LPI wake time	81.4.2	As described by Figure 81–13	LPI:M	Yes [ ] N/A [ ]
L3	Assertion of LPI in Rx direction	81.3.2.2	As defined in Table 81–4	LPI:M	Yes [ ] N/A [ ]
L4	Signal DATA_NOT_VALID on PLS_DATA_VALID.indication	81.4.4	While detecting LP_IDLE on XLGMII or CGMII	LPI:M	Yes [ ] N/A [ ]

82. Physical Coding Sublayer (PCS) for 64B/66B, type 40GBASE-R and 100GBASE-R

82.1 Overview

82.1.1 Scope

This clause specifies the Physical Coding Sublayer (PCS) that is common to two families of (40 Gb/s and 100 Gb/s) Physical Layer implementations, known as 40GBASE-R and 100GBASE-R. The 40GBASE-R PCS is a sublayer of the 40 Gb/s PHYs listed in Table 80–1. The 100GBASE-R PCS is a sublayer of the 100 Gb/s PHYs listed in Table 80–1. The terms 40GBASE-R and 100GBASE-R are used when referring generally to Physical Layers using the PCS defined in this clause.

Both 40GBASE-R and 100GBASE-R are based on a 64B/66B code. The 64B/66B code supports transmission of data and control characters, while maintaining robust error detection. Data distribution is introduced to support multiple lanes in the Physical Layer. Part of the distribution includes the periodic insertion of an alignment marker, which allows the receive PCS to align data from multiple lanes.

82.1.2 Relationship of 40GBASE-R and 100GBASE-R to other standards

Figure 82–1 depicts the relationships among the 40GBASE-R and 100GBASE-R sublayers (shown shaded), the Ethernet MAC and Reconciliation Sublayers, and the higher layers.

This clause borrows heavily from Clause 49. 64B/66B encoding is reused with appropriate changes made to support 8 byte alignment versus 4 byte alignment in Clause 49. In addition to 64B/66B encoding, there is a methodology to add alignment markers and distribute data to multiple lanes.

82.1.3 Summary of 40GBASE-R and 100GBASE-R sublayers

Figure 82–1 shows the relationship of the 40GBASE-R PCS and 100GBASE-R PCS sublayers (shown shaded) with other sublayers to the ISO Open System Interconnection (OSI) reference model.

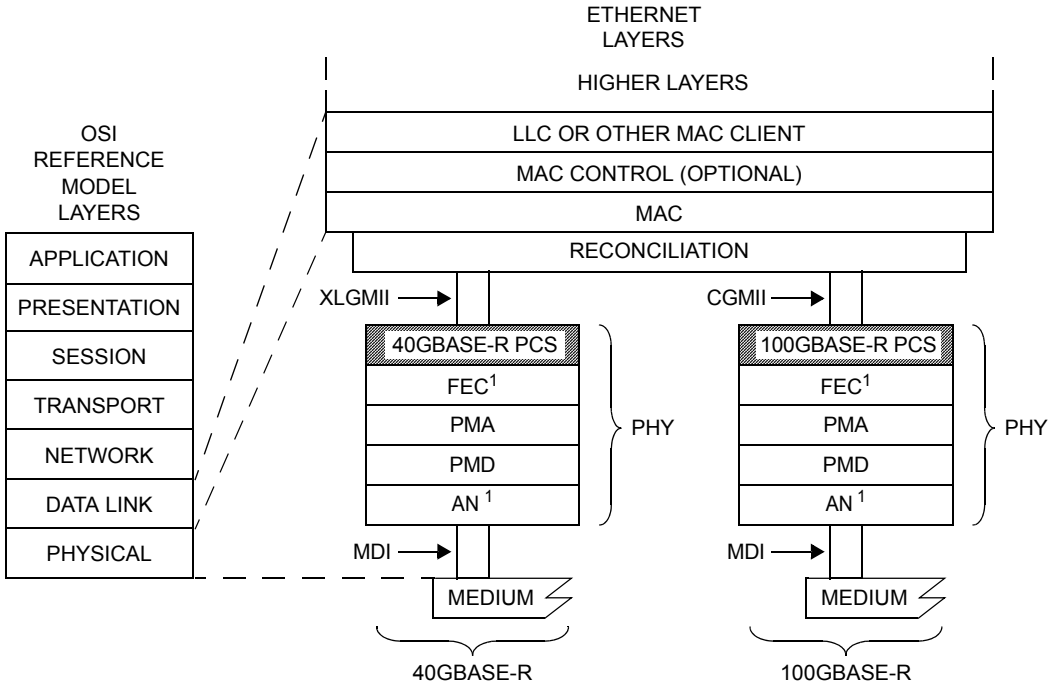

AN = AUTO-NEGOTIATION
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 82–1—40GBASE-R and 100GBASE-R PCS relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

82.1.3.1 Physical Coding Sublayer (PCS)

The PCS service interface is the Media Independent Interface (XLGMII/CGMII), which is defined in Clause 81. The 40 Gb/s variant of this is called the 40 Gb/s Media Independent Interface (XLGMII) and the 100 Gb/s variant of this interface is called the 100 Gb/s Media Independent Interface (CGMII). The XLGMII/CGMII provides a uniform interface to the Reconciliation Sublayer for all 40 Gb/s and 100 Gb/s PHY implementations.

The 40GBASE-R and 100GBASE-R PCSs provide all services required by the XLGMII/CGMII, including the following:

a) Encoding (decoding) of eight XLGMII/CGMII data octets to (from) 66-bit blocks (64B/66B).
b) Transferring encoded data to (from) the PMA.
c) Compensation for any rate differences caused by the insertion or deletion of alignment markers or due to any rate difference between the XLMII/CGMII and PMA through the insertion or deletion of idle control characters.
d) Determining when a functional link has been established and informing the management entity via the MDIO when the PHY is ready for use.

82.1.4 Inter-sublayer interfaces

The upper interface of the PCS may connect to the Reconciliation Sublayer through the XLGMII/CGMII. The lower interface of the PCS connects to the PMA sublayer to support a PMD. If the optional BASE-R FEC sublayer is implemented (see Clause 74) and an optional physical instantiation, i.e., XLAUI or CAUI-n, is not implemented directly below the PCS sublayer, then the lower interface connects to the BASE-R FEC sublayer. For Physical Layers that use Clause 91 RS-FEC, if an optional physical instantiation (i.e., CAUI-n) is not implemented directly below the PCS sublayer, then the lower interface connects to the RS-FEC sublayer. The 40GBASE-R PCS has a nominal rate at the PMA/FEC service interface of 10.3125 Gtransfers/s per PCS lane, which provides capacity for the MAC data rate of 40 Gb/s. The 100GBASE-R PCS has a nominal rate at the PMA/FEC service interface of 5.15625 Gtransfers/s per PCS lane, which provides capacity for the MAC data rate of 100 Gb/s.

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.

82.1.4.1 PCS service interface (XLGMII/CGMII)

The PCS service interface allows the 40GBASE-R or 100GBASE-R PCS to transfer information to and from a PCS client. The PCS client is the Reconciliation Sublayer. The PCS Service Interface is precisely defined as the Media Independent Interface (XLGMII/CGMII) in Clause 81.

82.1.4.2 Physical Medium Attachment (PMA) or Forward Error Correction (FEC) service interface

The PMA or FEC service interface for the PCS is described in an abstract manner and does not imply any particular implementation. The PMA/FEC Service Interface supports the exchange of encoded data between the PCS and PMA or FEC sublayer. The PMA or FEC service interface is defined in 83.3 or 94.2.1 and is an instance of the inter-sublayer service interface definition in 80.3 or 91.2.

82.1.5 Functional block diagram

Figure 82–2 provides a functional block diagram of the 40GBASE-R PCS and 100GBASE-R PCS.

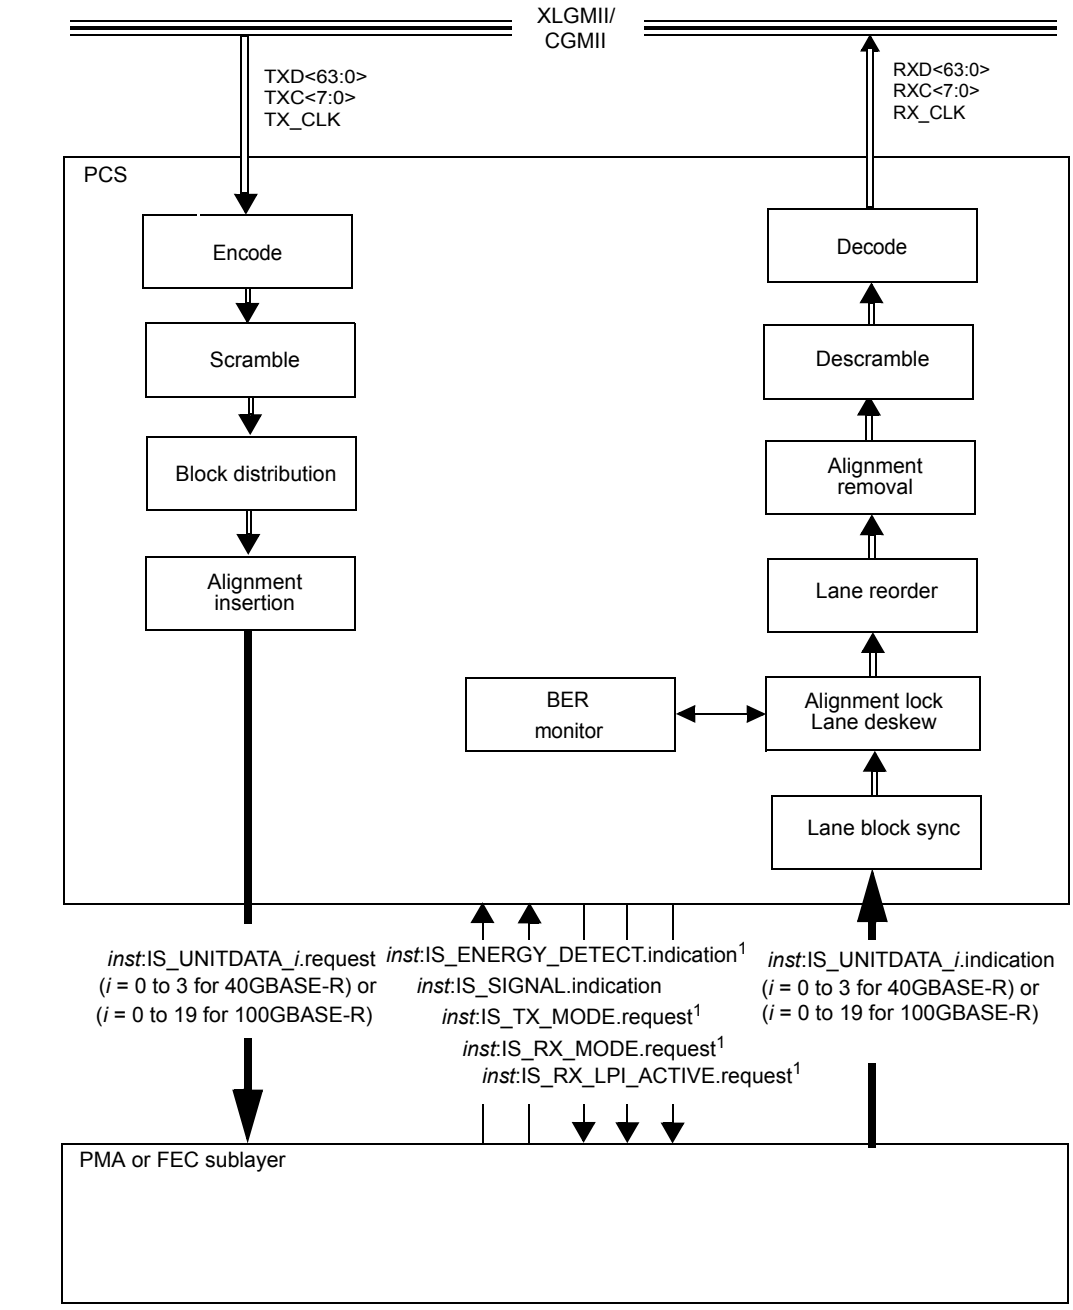

NOTE 1—FOR OPTIONAL EEE DEEP SLEEP CAPABILITY

Figure 82–2—Functional block diagram

82.2 Physical Coding Sublayer (PCS)

82.2.1 Functions within the PCS

The 40GBASE-R and 100GBASE-R PCSs comprise the PCS Transmit and PCS Receive processes for each rate of operation. 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 test-pattern mode.

When communicating with the XLGMII/CGMII, the PCS uses an eight octet-wide, synchronous data path, with packet delimiting being provided by transmit control signals (TXC<n> = 1) and receive control signals (RXC<n> = 1). When communicating with the PMA, the PCS uses multiple serial streams, 4 encoded bit streams for a 40GBASE-R PCS, or 20 encoded bit streams for a 100GBASE-R PCS. Alignment to 64B/66B blocks 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 XLGMII/CGMII format and the PMA/FEC service interface format.

Note that these serial streams originate from a common clock in each direction, but may vary in phase and Skew dynamically.

When the transmit channel is in normal mode, the PCS Transmit process continuously generates blocks based upon the TXD <63:0> and TXC <7:0> signals on the XLGMII/CGMII. The blocks are scrambled and then distributed to individual PCS lanes. After distribution, alignment marker blocks are periodically added to each PCS lane. Transmit data-units are sent to the service interface via the *inst*:IS_UNITDATA_*i*.request primitive. In these primitives, the *inst* can be either FEC or PMA depending on which sublayer is adjacent to the PCS sublayer.

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/FEC service interface via the *inst*:IS_UNITDATA_*i*.request primitive.

When the receive channel is in normal or test-pattern mode, the PCS Synchronization process continuously monitors *inst*:IS_SIGNAL.indication(SIGNAL_OK). When SIGNAL_OK indicates OK, then the PCS Synchronization process accepts data-units via the *inst*:IS_UNITDATA_*i*.indication primitive. It attains block synchronization based on the 2-bit synchronization headers on each one of the PCS lanes. Once block synchronization is found on a PCS lane, then alignment marker lock can be attained by searching for valid alignment markers. After alignment markers are found on all PCS lanes, the PCS lanes can be reordered and deskewed. Note that a particular transmit PCS lane can be received on any receive lane of the service interface due to the Skew and multiplexing that occurs in the path.

The PCS deskew process conveys received blocks to the PCS Receive process. The PCS deskew process deskews and aligns the individual PCS lanes, removes the alignment markers, forms a single stream, and sets the align_status flag to indicate whether the PCS has obtained alignment.

When the PCS deskew process has obtained alignment, the BER monitor process monitors the signal quality asserting hi_ber if excessive errors are detected. When align_status is asserted and hi_ber is deasserted, the PCS Receive process continuously accepts blocks and generates RXD <63:0> and RXC <7:0> on the MII.

When the receive channel is in test-pattern mode, the BER monitor process may be 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.

The PCS shall provide transmit test-pattern mode for the scrambled idle pattern (see 82.2.11), and shall provide receive test-pattern mode for the scrambled idle pattern (see 82.2.11). Test-pattern mode is activated

separately for transmit and receive. The PCS shall support transmit test-pattern mode and receive test-pattern mode operating simultaneously so as to support loopback testing.

82.2.2 Use of blocks

The PCS maps XLGMII/CGMII 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. The PCS functions ENCODE and DECODE generate, manipulate, and interpret blocks as defined in 82.2.3.

82.2.3 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 is a run-length-limited code.[5]

The relationship of block bit positions relative to XLGMII/CGMII, PMA, and other PCS functions is illustrated in Figure 82–3 for transmit and Figure 82–4 for receive. These figures illustrate the processing of a block containing 8 data octets. See 82.2.3.3 for information on how blocks containing control characters are mapped into 66-bit blocks. Note that the sync header is generated by the encoder and bypasses the scrambler.

82.2.3.1 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. The control characters /I/ and /E/ are labeled C_0 to C_7. The control characters, /Q/ and /Fsig/, for ordered sets are labeled as O_0 since they are only valid on the first octet of the XLGMII/CGMII. The control character for start is labeled as S_0 for the same reason. The control character for terminate is labeled as T_0 to T_7. The four trailing zero data octets in ordered sets are labeled as Z_4 to Z_7.

One XLGMII/CGMII transfer provides 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 XLGMII/CGMII transfer.

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.

[5]In 10GBASE-R the run length is limited to 66 bits, but in 40GBASE-R and 100GBASE-R, when multiplexing occurs in the PMA, this guaranteed run length limit increases. For example, if two PCS lanes are multiplexed in the PMA, then the possible run length would double.

82.2.3.2 Transmission order

Block bit transmission order shall be as illustrated in Figure 82–3 and Figure 82–4. Note that these figures show the mapping from XLGMII/CGMII to 64B/66B block for a block containing eight data characters.

82.2.3.3 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 is always guaranteed to contain 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, Terminate character, or ordered set, that character is implied by the block type field. Other control characters are encoded in a 7-bit control code. Each control block encodes eight characters.

The format of the blocks is as shown in Figure 82–5. 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 XLGMII/CGMII, 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 characters 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.

All unused values of block type field[6] are invalid; they shall not be transmitted and shall be considered an error if received.

[6]The block type field values have been chosen to have a 4-bit Hamming distance between them. There are five unused values that maintain this Hamming distance: 0x00, 0x2D, 0x33, 0x55, and 0x66.

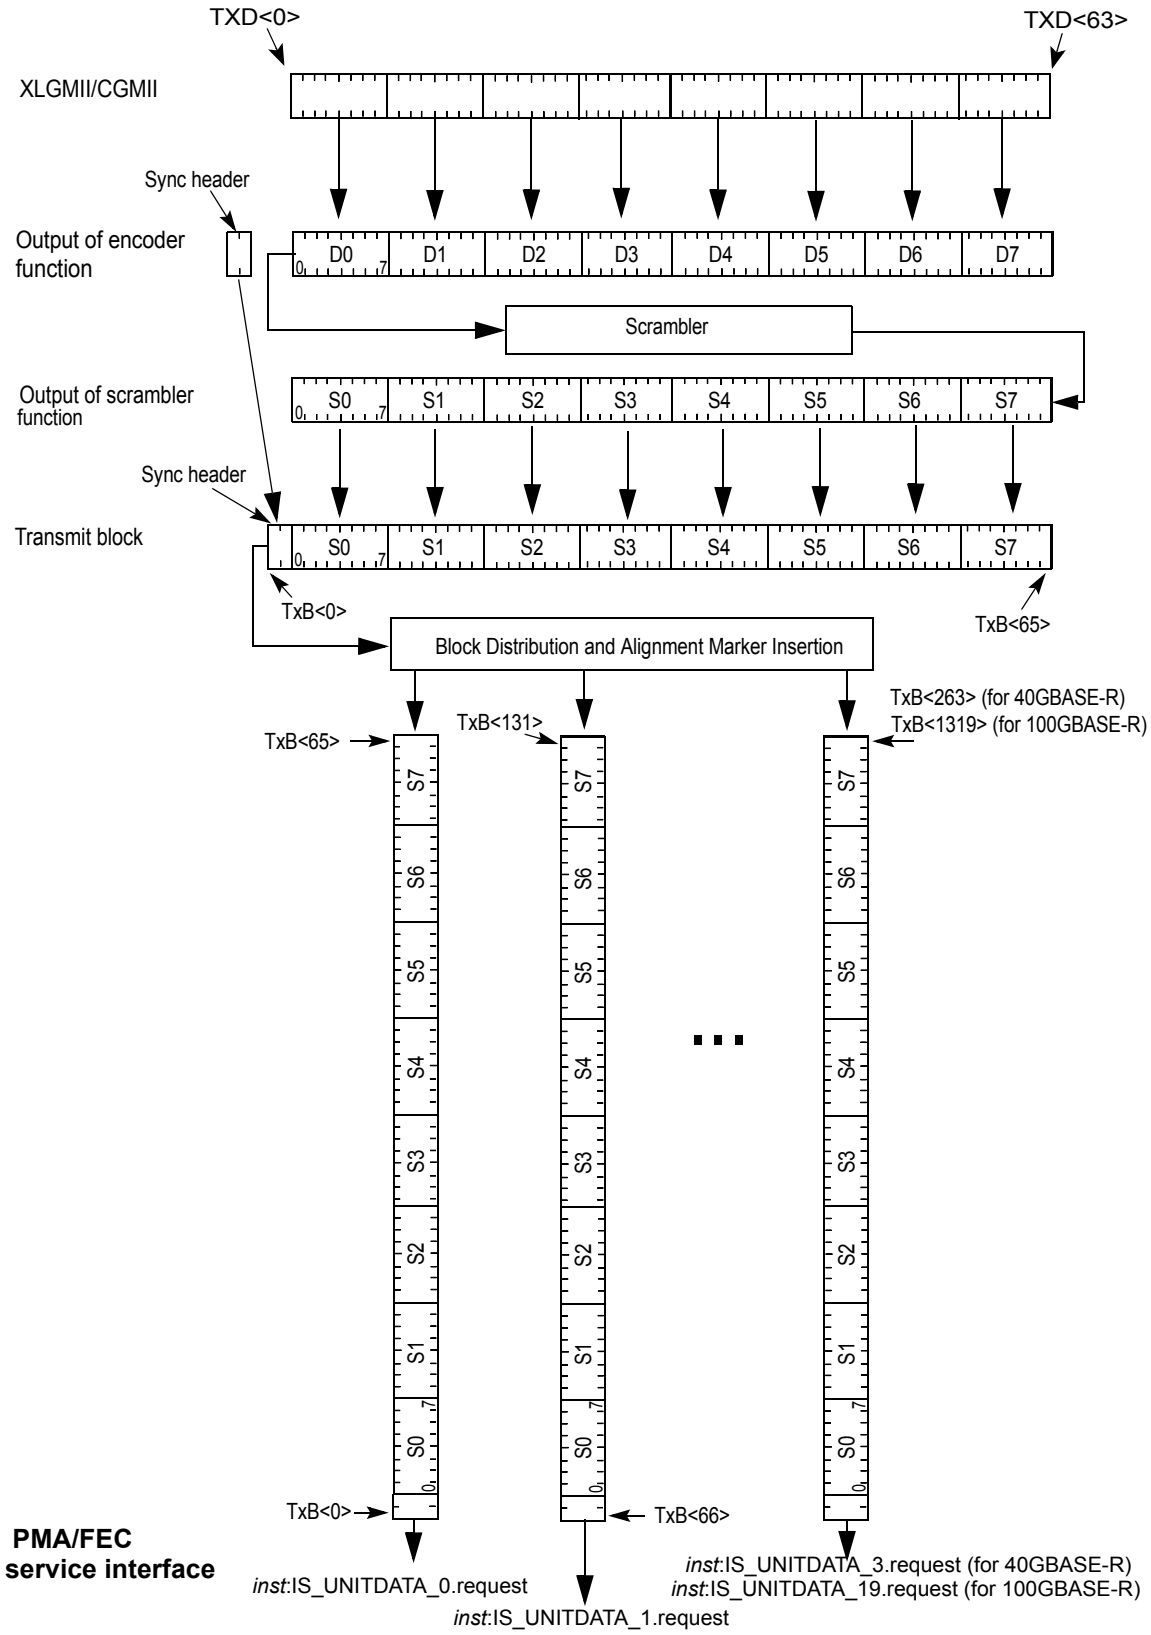

Figure 82–3—PCS Transmit bit ordering

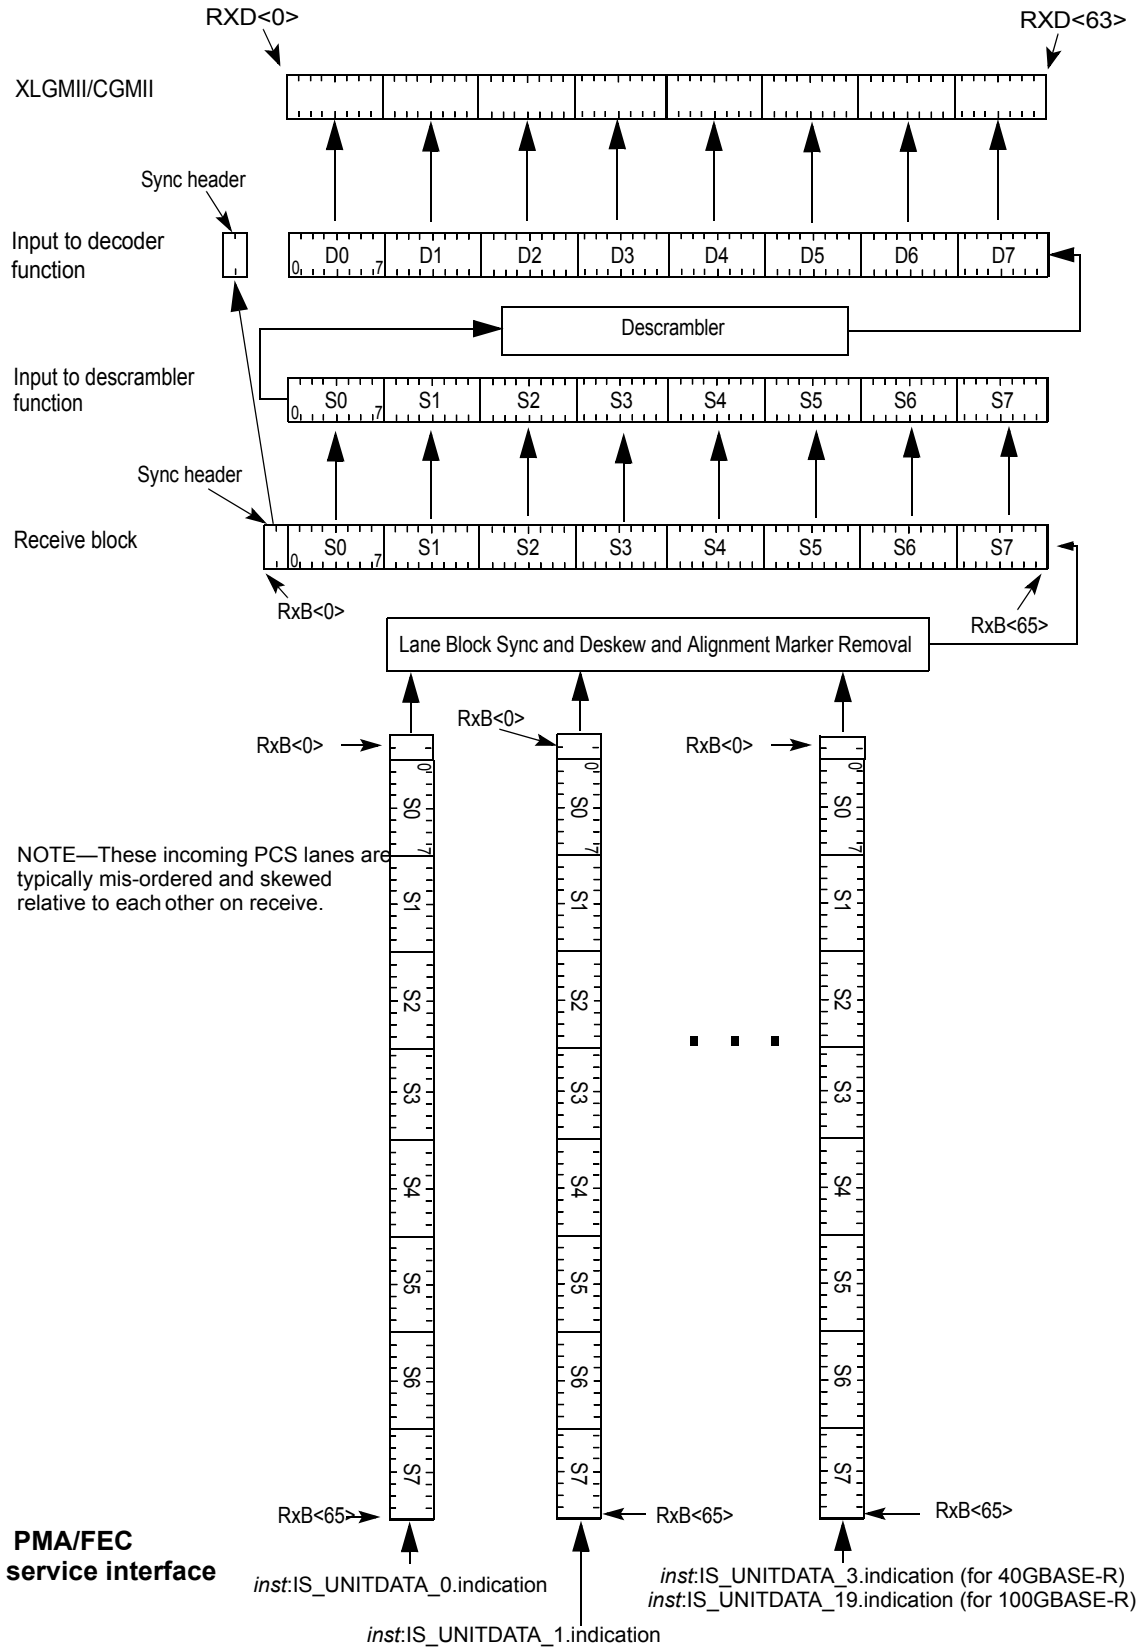

Figure 82–4—PCS Receive bit ordering

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
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 Z_4 Z_5 Z_6 Z_7	10	0x4B	D_1	D_2	D_3	O_0	0x000_0000		
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 82–5—64B/66B block formats

WARNING

The mapping of 40GBASE-R PCS blocks into OPU3 specified in ITU-T G.709 [B50] depends on the set of control block types shown in Figure 82–5. Any deviation from the coding specified in Figure 82–5 will break the mapping and may prevent 40GBASE-R PCS blocks from being mapped into OPU3 (see ITU-T G.709 [B50] for more details).

82.2.3.4 Control codes

The same set of control characters are supported by the XLGMII/CGMII and the PCS. The representations of the control characters are the control codes. XLGMII/CGMII encodes a control character into an octet (an eight bit value). The 40GBASE-R and 100GBASE-R PCS encode the start and terminate control characters implicitly by the block type field. The 40GBASE-R and 100GBASE-R PCS encode the ordered set control codes using the block type field. The 40GBASE-R and 100GBASE-R PCS encode each of the other control characters into a 7-bit control code.

The control characters and their mappings to 40GBASE-R and 100GBASE-R control codes and XLGMII/CGMII control codes are specified in Table 82–1. All XLGMII/CGMII, 40GBASE-R, and 100GBASE-R control code values that do not appear in the table shall not be transmitted and shall be considered 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 has not been negotiated, LPI shall not be transmitted and shall be treated as an error if received.

82.2.3.5 Valid and invalid blocks

A block is invalid if any of the following conditions exists:

a) The sync header has a value of 00 or 11.
b) The block type field contains an invalid value (one not included in Figure 82–5).
c) Any control character contains a value not included in Table 82–1.
d) The set of eight XLGMII/CGMII characters does not have a corresponding block format in Figure 82–5.

82.2.3.6 Idle (/I/)

Idle control characters (/I/) are transmitted when idle control characters are received from the XLGMII/CGMII. Idle control characters may be added or deleted by the PCS to adapt between clock rates. /I/ insertion and deletion shall occur in groups of 8. /I/s may be added following idle control characters or ordered sets. They shall not be added while data is being received.

Table 82–1—Control codes

Control character	Notation	XLGMII/ CGMII control code	40/100GBASE-R O code	40GBASE-R and 100GBASE-R control code
idle	/I/	0x07		0x00
LPI	/LI/	0x06		0x06
start	/S/	0xFB		Encoded by block type field
terminate	/T/	0xFD		Encoded by block type field
error	/E/	0xFE		0x1E
Sequence ordered set	/Q/	0x9C	0x0	Encoded by block type 0x4B plus O code, control codes are set to 0x00
Signal ordered set[a]	/Fsig/	0x5C	0xF	Encoded by block type 0x4B plus O code, control codes are set to 0x00

[a]Reserved for INCITS T11 Fibre Channel use.

To communicate LPI, the LPI control character /LI/ is sent continuously in place of /I/. LPI control characters are transmitted when LPI control characters are received from the XLGMII or CGMII. 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 8. /LI/s inserted for clock compensation may only be inserted following other LPI characters.

82.2.3.7 Start (/S/)

The start control character (/S/) indicates the start of a packet. This delimiter is only valid on the first octet of the XLGMII/CGMII (TXD<7:0> and RXD<7:0>). Receipt of an /S/ on any other octet of TXD indicates an error. Block type field values implicitly encode an /S/ as the first character of the block.

82.2.3.8 Terminate (/T/)

The terminate control character (/T/) indicates the end of a packet. Since packets vary in length, the /T/ can occur on any octet of the XLGMII/CGMII 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/ or an /E/.

82.2.3.9 ordered set (/O/)

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 followed by four zero data characters on the XLGMII/CGMII. Ordered sets always begin on the first octet of the XLGMII/CGMII. 40 Gigabit and 100 Gigabit Ethernet use one kind of ordered set: the sequence ordered set (see 81.3.4). An additional ordered set, the signal ordered set, has been reserved and it begins with another control code. The ordered set control characters (/Q/ and /Fsig/) indicate the start of an ordered set. The block type field plus the O code encode the specific control character for the ordered set. See Table 82–1 for the mappings.

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 only one of the two ordered sets may be deleted. Only idle control characters may be inserted for clock compensation. Signal ordered sets are not deleted for clock compensation.

82.2.3.10 Error (/E/)

In both the 64B/66B encoder and decoder, the /E/ is generated whenever an /E/ is detected. The /E/ is also generated when invalid blocks are detected. The /E/ allows the PCS to propagate detected errors. See R_TYPE and T_TYPE function definitions in 82.2.19.2.3 for further information.

82.2.4 Transmit process

The transmit process generates blocks based upon the TXD<63:0> and TXC<7:0> signals received from the XLGMII/CGMII. One XLGMII/CGMII data transfer is encoded into one 66-bit block. It takes 66 *inst*:IS_UNITDATA_*i* transfers to send a 66-bit block of data on each of the PCS lanes. The transmit process must delete idle control characters or sequence ordered sets to accommodate the transmission of alignment markers. If the PCS transmit process spans multiple clock domains, it may also perform clock rate compensation via the deletion of idle control characters or sequence ordered sets or the insertion of idle control characters.

There are sufficient idle control characters to delete in order to make room for alignment markers, in addition to handling clock compensation. Idle control characters or sequence ordered sets are removed, if necessary, to accommodate the insertion of the 66-bit alignment markers. See 82.2.3.6 for more details.

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.

82.2.5 Scrambler

The payload, tx_coded<65:2>, is scrambled with a self-synchronizing scrambler. The scrambler is identical to the scrambler used in Clause 49, see 49.2.6 for the definition of the scrambler. The sync bits, tx_coded<1:0>, are not scrambled.

82.2.6 Block distribution

Once the data is encoded and scrambled, it is distributed to multiple PCS lanes, 66-bit blocks at a time in a round robin distribution from the lowest to the highest numbered PCS lanes. This allows the PCS to support multiple physical lanes in the PMD and XLAUI or CAUI-n interfaces (see Annex 83A, Annex 83B, Annex 83D, and Annex 83E). The 40GBASE-R PCS distributes the 66-bit blocks to 4 PCS lanes, and the 100GBASE-R PCS distributes the blocks to 20 PCS lanes. The distribution process is shown in Figure 82–6.

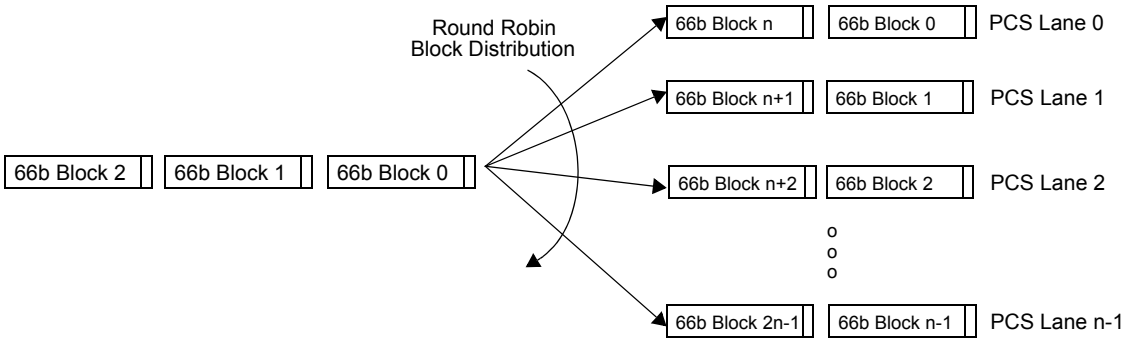

Figure 82–6—PCS Block distribution

82.2.7 Alignment marker insertion

In order to support deskew and reordering of individual PCS lanes at the receive PCS, alignment markers are added periodically to each PCS lane. The alignment marker has the form of a specially defined 66-bit block with a control block sync header. These markers interrupt any data transfer that is already in progress. This allows alignment markers to be inserted into all PCS lanes at the same time. Room for the alignment markers is created by periodically deleting IPG from the XLGMII/CGMII data stream. Other special properties of the alignment markers are that they are not scrambled and do not conform to the encoding rules as outlined in Figure 82–5. This is possible because the alignment markers are added after encoding is performed in the transmit PCS and the alignment markers are removed before 64B/66B decoding is performed in the receive PCS. The alignment markers are not scrambled in order to allow the receiver to find the alignment markers, deskew the PCS lanes, and reassemble the aggregate stream before descrambling is performed. The alignment markers themselves are formed from a known pattern that is defined to be balanced and with many transitions and therefore scrambling is not necessary for the alignment markers. The alignment markers shall be inserted after every 16383 66-bit blocks on each PCS lane. Alignment marker insertion is shown in Figure 82–7 and Figure 82–8.

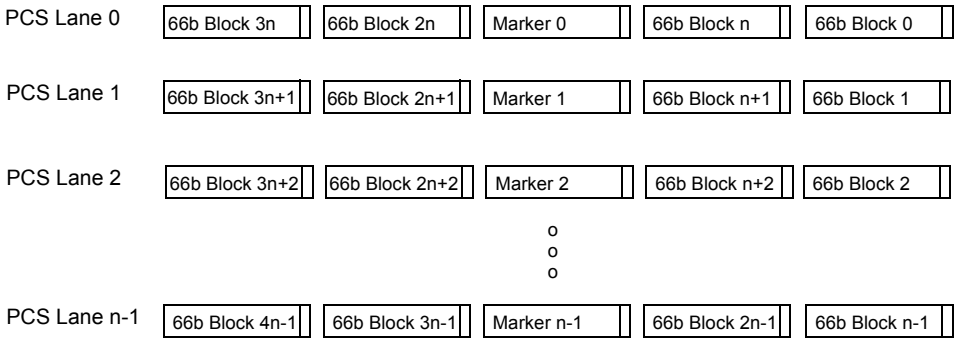

Figure 82–7—Alignment marker insertion

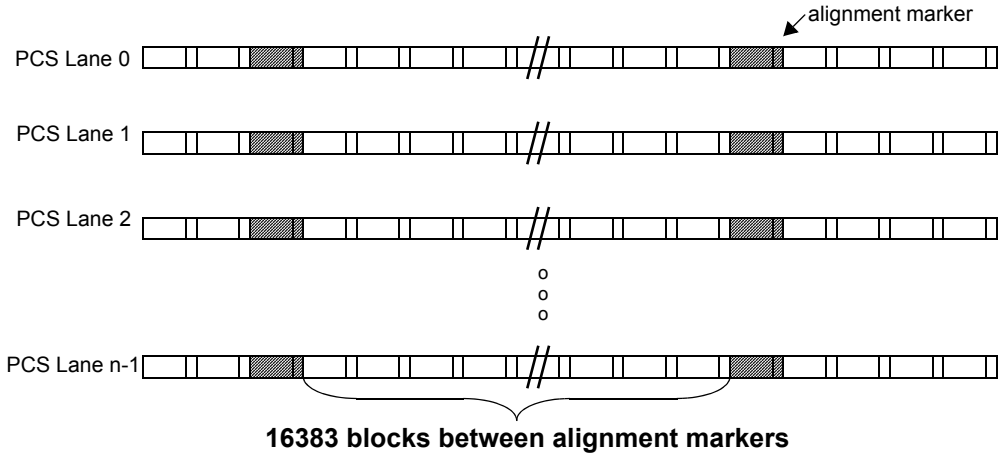

Figure 82–8—Alignment marker insertion period

The format of the alignment markers is shown in Figure 82–9.

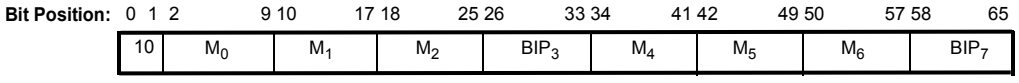

Figure 82–9—Alignment marker format

The content of the alignment markers shall be as shown in Table 82–2 for 100GBASE-R and in Table 82–3 for 40GBASE-R. The contents depend on the PCS lane number and the octet number. Note that M_4 through M_6 are the bit-wise inversion of M_0 through M_2, respectively. Also BIP_7 is the bit-wise inversion of BIP_3. This property allows the alignment markers to be DC-balanced. Lane markers 0 to 19 from Table 82–2 are used for the 100GBASE-R PCS and lane markers 0 to 3 from Table 82–3 are used for the 40GBASE-R PCS. As an example, the lane marker for 100GBASE-R lane number 0 is sent as (left most bit sent first):

10 10000011 00010110 10000100 BIP_3 01111100 11101001 01111011 BIP_7

After the alignment markers are inserted, data is sent to the PMA or FEC sublayer adjacent to the PCS.

Table 82–2—100GBASE-R Alignment marker encodings

PCS lane number	Encoding[a] {M_0, M_1, M_2, BIP_3, M_4, M_5, M_6, BIP_7}	PCS lane number	Encoding[a] {M_0, M_1, M_2, BIP_3, M_4, M_5, M_6, BIP_7}
0	0xC1, 0x68, 0x21, BIP_3, 0x3E, 0x97, 0xDE, BIP_7	10	0xFD, 0x6C, 0x99, BIP_3, 0x02, 0x93, 0x66, BIP_7
1	0x9D, 0x71, 0x8E, BIP_3, 0x62, 0x8E, 0x71, BIP_7	11	0xB9, 0x91, 0x55, BIP_3, 0x46, 0x6E, 0xAA, BIP_7
2	0x59, 0x4B, 0xE8, BIP_3, 0xA6, 0xB4, 0x17, BIP_7	12	0x5C, 0x B9, 0xB2, BIP_3, 0xA3, 0x46, 0x4D, BIP_7
3	0x4D, 0x95, 0x7B, BIP_3, 0xB2, 0x6A, 0x84, BIP_7	13	0x1A, 0xF8, 0xBD, BIP_3, 0xE5, 0x07, 0x42, BIP_7
4	0xF5, 0x07, 0x09, BIP_3, 0x0A, 0xF8, 0xF6, BIP_7	14	0x83, 0xC7, 0xCA, BIP_3, 0x7C, 0x38, 0x35, BIP_7
5	0xDD, 0x14, 0xC2, BIP_3, 0x22, 0xEB, 0x3D, BIP_7	15	0x35, 0x36, 0xCD, BIP_3, 0xCA, 0xC9, 0x32, BIP_7
6	0x9A, 0x4A, 0x26, BIP_3, 0x65, 0xB5, 0xD9, BIP_7	16	0xC4, 0x31, 0x4C, BIP_3, 0x3B, 0xCE, 0xB3, BIP_7
7	0x7B, 0x45, 0x66, BIP_3, 0x84, 0xBA, 0x99, BIP_7	17	0xAD, 0xD6, 0xB7, BIP_3, 0x52, 0x29, 0x48, BIP_7
8	0xA0, 0x24, 0x76, BIP_3, 0x5F, 0xDB, 0x89, BIP_7	18	0x5F, 0x66, 0x2A, BIP_3, 0xA0, 0x99, 0xD5, BIP_7
9	0x68, 0xC9, 0xFB, BIP_3, 0x97, 0x36, 0x04, BIP_7	19	0xC0, 0xF0, 0xE5, BIP_3, 0x3F, 0x0F, 0x1A, BIP_7

[a]Each octet is transmitted LSB to MSB.

Table 82–3—40GBASE-R Alignment marker encodings

PCS lane number	Encoding[a] {M_0, M_1, M_2, BIP_3, M_4, M_5, M_6, BIP_7}
0	0x90, 0x76, 0x47, BIP_3, 0x6F, 0x89, 0xB8, BIP_7
1	0xF0, 0xC4, 0xE6, BIP_3, 0x0F, 0x3B, 0x19, BIP_7
2	0xC5, 0x65, 0x9B, BIP_3, 0x3A, 0x9A, 0x64, BIP_7
3	0xA2, 0x79, 0x3D, BIP_3, 0x5D, 0x86, 0xC2, BIP_7

[a]Each octet is transmitted LSB to MSB.

82.2.8 BIP calculations

A PCS lane BIP field is carried in each PCS Lane alignment marker. This allows an accurate and fast measure of the bit error ratio of a given PCS Lane. This information is used to update error counters; no state machines use this information.

Each alignment marker has two Bit Interleaved Parity fields, BIP_3 and BIP_7. BIP_7 is a bit-wise inversion of BIP_3 in order to keep the alignment marker DC-balanced. The BIP_3 field contains the result of a bit interleaved parity calculation. Each bit in the BIP field is an even parity calculation over all of the previous specified bits of a given PCS Lane, from and including the previous alignment marker, but not including the current alignment marker. Using the bit definitions as shown in Figure 82–9, Table 82–4 has the bit assignments for each BIP_3 bit. As an example, BIP_3 bit 0 contains the result of XORing 131072 bits from 16384 66-bit words. BIP_3 bit 3 and bit 4 also include one sync header bit from each 66-bit word. Bit 3 and bit 4 each contain the result of XORing 147456 bits.

Table 82–4—BIP_3 bit assignments

BIP_3 bit number	Assigned 66-bit word bits
0	2, 10, 18, 26, 34, 42, 50, 58
1	3, 11, 19, 27, 35, 43, 51, 59
2	4, 12, 20, 28, 36, 44, 52, 60
3	0, 5, 13, 21, 29, 37, 45, 53, 61
4	1, 6, 14, 22, 30, 38, 46, 54, 62
5	7, 15, 23, 31, 39, 47, 55, 63
6	8, 16, 24, 32, 40, 48, 56, 64
7	9, 17, 25, 33, 41, 49, 57, 65

BIP_3 and BIP_7 are transmitted LSB to MSB. As an example, with BIP_3 = 0x0F, the PCS lane marker for 100GBASE-R lane number 0 is sent as (left most bit sent first):

10 10000011 00010110 10000100 11110000 01111100 11101001 01111011 00001111

If the EEE capability is supported, BIP statistics are only updated when the receiver is in the RX_ACTIVE state (see Figure 82–19). In all other states, the running parity is not calculated. The BIP statistics shall be first updated after transitioning from RAMs to normal AMs on the first received normal AM when LPI_FW is FALSE and on the second received AM after entering the RX_ACTIVE state when LPI_FW is TRUE.

82.2.9 Rapid alignment marker insertion

For the optional EEE capability, an alternate method of alignment is used when operating in the deep sleep low power state. Rapid Alignment Markers (RAMs) function in a similar manner to the alignment markers described in 82.2.7. Normal alignment markers are sent when the transmitter has an LPI transmit state of TX_ACTIVE; RAMs are sent in the TX_WAKE2 state until down_count_done is TRUE and in all the other LPI transmit states. Additionally, the BIP component defined for alignment markers is replaced by a count

down field (CD) so that the transition from RAMs to normal alignment markers can be indicated. The RAMs shall be inserted after every 7 66-bit blocks on each 100G PCS lane and every 15 66-bit blocks on each 40G PCS lane. RAM insertion is performed in the same manner as shown in Figure 82–7 and Figure 82–10. The transition from RAMs to normal alignment markers is shown in Figure 82–10. The count down field is also used to communicate some of the states of the tx_mode when it is not being used to coordinate the transition. After the LPI Transmit state diagram transitions from TX_ACTIVE to TX_SLEEP, the first RAM shall be inserted after at least one block of /LI/ has been transmitted on PCS lane 0. In order to force the RAMs to coincide with the start of an FEC block, the distance between the first RAM and preceding normal alignment marker shall be a multiple of 4 66-bit blocks.

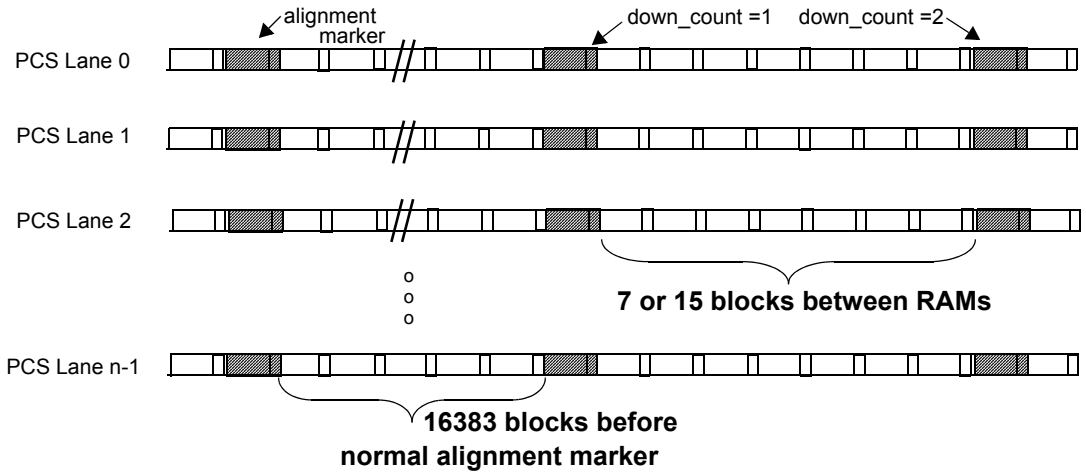

Figure 82–10—RAM transition

The format of the RAMs is shown in Figure 82–11.

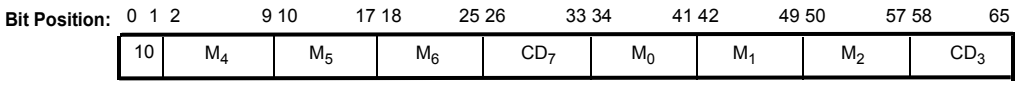

Figure 82–11—RAM format

The content of the RAMs shall be as shown in Table 82–5 for 100GBASE-R or Table 82–6 for 40GBASE-R. Note that these are similar to normal alignment markers, with CD_3 replacing BIP_3 and CD_7 replacing BIP_7, and also M_0 through M_2 and CD_3 swapped with M_4 through M_6 and CD_7, respectively. As an example, the lane marker for 100GBASE-R lane number 0 is sent as (left most bit sent first):

10 01111100 11101001 01111011 CD_7 10000011 00010110 10000100 CD_3

After the RAMs are inserted, data is sent to the PMA or FEC sublayer adjacent to the PCS.

The value of the CD_3 field is derived by the bit-wise XOR of the down_count variable with the M_0 value for the lane (therefore the last 5 RAMs sent by a 100GBASE-R PCS on PCS lane 0 would have CD_3 values: 0xC4, 0xC5, 0xC2, 0xC3, 0xC0; for PCS lane 1 these would be: 0x98, 0x99, 0x9E, 0x9F, 0x9C). The CD_7 field is the bit-wise inversion of CD_3. The CD field is used by the link partner to understand the expected

transition from RAMs to normal AMs. It may also be used by a device with a detached PMA or FEC sublayer to infer the state of the PCS.

Table 82–5—100GBASE-R RAM encodings

PCS lane number	Encoding[a] {M_4, M_5, M_6, CD_7, M_0, M_1, M_2, CD_3}	PCS lane number	Encoding[a] {M_4, M_5, M_6, CD_7, M_0, M_1, M_2, CD_3}
0	0x3E, 0x97, 0xDE, CD_7, 0xC1, 0x68, 0x21, CD_3	10	0x02, 0x93, 0x66, CD_7, 0xFD, 0x6C, 0x99, CD_3
1	0x62, 0x8E, 0x71, CD_7, 0x9D, 0x71, 0x8E, CD_3	11	0x46, 0x6E, 0xAA, CD_7, 0xB9, 0x91, 0x55, CD_3
2	0xA6, 0xB4, 0x17, CD_7, 0x59, 0x4B, 0xE8, CD_3	12	0xA3, 0x46, 0x4D, CD_7, 0x5C, 0xB9, 0xB2, CD_3
3	0xB2, 0x6A, 0x84, CD_7, 0x4D, 0x95, 0x7B, CD_3	13	0xE5, 0x07, 0x42, CD_7, 0x1A, 0xF8, 0xBD, CD_3
4	0x0A, 0xF8, 0xF6, CD_7, 0xF5, 0x07, 0x09, CD_3	14	0x7C, 0x38, 0x35, CD_7, 0x83, 0xC7, 0xCA, CD_3
5	0x22, 0xEB, 0x3D, CD_7, 0xDD, 0x14, 0xC2, CD_3	15	0xCA, 0xC9, 0x32, CD_7, 0x35, 0x36, 0xCD, CD_3
6	0x65, 0xB5, 0xD9, CD_7, 0x9A, 0x4A, 0x26, CD_3	16	0x3B, 0xCE, 0xB3, CD_7, 0xC4, 0x31, 0x4C, CD_3
7	0x84, 0xBA, 0x99, CD_7, 0x7B, 0x45, 0x66, CD_3	17	0x52, 0x29, 0x48, CD_7, 0xAD, 0xD6, 0xB7, CD_3
8	0x5F, 0xDB, 0x89, CD_7, 0xA0, 0x24, 0x76, CD_3	18	0xA0, 0x99, 0xD5, CD_7, 0x5F, 0x66, 0x2A, CD_3
9	0x97, 0x36, 0x04, CD_7, 0x68, 0xC9, 0xFB, CD_3	19	0x3F, 0x0F, 0x1A, CD_7, 0xC0, 0xF0, 0xE5, CD_3

[a]Each octet is transmitted LSB to MSB.

Table 82–6—40GBASE-R RAM encodings

PCS lane number	Encoding[a] {M_4, M_5, M_6, CD_7, M_0, M_1, M_2, CD_3}
0	0x6F, 0x89, 0xB8, CD_7, 0x90, 0x76, 0x47, CD_3
1	0x0F, 0x3B, 0x19, CD_7, 0xF0, 0xC4, 0xE6, CD_3
2	0x3A, 0x9A, 0x64, CD_7, 0xC5, 0x65, 0x9B, CD_3
3	0x5D, 0x86, 0xC2, CD_7, 0xA2, 0x79, 0x3D, CD_3

[a]Each octet is transmitted LSB to MSB.

82.2.10 PMA or FEC Interface

When the transmit channel is operating in normal mode, the 40GBASE-R PCS sends four data streams via *inst*:IS_UNITDATA_*i*.request primitives and the 100GBASE-R PCS sends twenty data streams via *inst*:IS_UNITDATA_*i*.request primitives.

The *inst*:IS_UNITDATA_*i*.request primitives are separate serial streams of bits. Since 66-bit blocks were distributed to each lane, that means for the 40GBASE-R PCS: bits 0 to 65 are sent on PCSL 0, bits 66 to 131

are sent on PCSL 1; bits 132 to 197 are sent on PCSL 2, bits 198 to 263 are sent on PCSL 3, then bits 264 to 329 are sent on PCSL 0 etc.

For 100GBASE-R it is: bits 0 to 65 on PCSL 0, bits 66 to 131 on PCSL 1, bits 132 to 197 on PCSL 2, bits 198 to 263 on PCSL 3, bits 264 to 329 on PCSL 4, bits 330 to 395 on PCSL 5, bits 396 to 461 on PCSL 6, bits 462 to 527 on PCSL 7, bits 528 to 593 on PCSL 8, bits 594 to 659 on PCSL 9, bits 660 to 725 on PCSL 10, bits 726 to 791 on PCSL 11, bits 792 to 857 on PCSL 12, bits 858 to 923 on PCSL 13, bits 924 to 989 on PCSL 14, bits 990 to 1055 on PCSL 15, bits 1056 to 1121 on PCSL 16, bits 1122 to 1187 on PCSL 17, bits 1188 to 1253 on PCSL 18, bits 1254 to 1319 on PCSL 19, then bits 1320 to 1385 on PCSL 0, etc.

82.2.11 Test-pattern generators

The PCS shall have the ability to generate and detect a scrambled idle test pattern. This test-pattern mode is suitable for receiver tests and for certain transmitter tests.

When a scrambled idle pattern is enabled, the test pattern is generated by the scrambler. No seeding of the scrambler is required during test-pattern operation. The input to the scrambler is a control block (block type=0x1E) with all idles as defined in Figure 82–5. Note that the sync headers and alignment markers are added to the stream so that the receive PCS can align and deskew the PCS lanes.

If a Clause 45 MDIO is implemented, then control of the test-pattern generation is from the BASE-R PCS test-pattern control register (bit 3.42.3).

When the transmit channel is operating in test-pattern mode, the encoded bit stream is distributed to the PCS Lanes as in normal operation (see 82.2.4).

82.2.12 Block synchronization

When the receive channel is operating in normal mode, the block synchronization function receives data via 4 (for 40GBASE-R) or 20 (for 100GBASE-R) IS_UNITDATA_*i*.indication primitives. The PCS forms 4 or 20 bit streams from the primitives by concatenating the bits from the indications of each primitive in order from each *inst*:IS_UNITDATA_0.indication to *inst*:IS_UNITDATA_3.indication or *inst*:IS_UNITDATA_0.indication to *inst*:IS_UNITDATA_19.indication. It obtains lock to the 66-bit blocks in each bit stream using the sync headers and outputs 66-bit blocks. Block lock is obtained as specified in the block lock state diagram shown in Figure 82–12.

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 82–19.

82.2.13 PCS lane deskew

Once the receiver achieves block lock on a lane, then it begins obtaining alignment marker lock as specified in the alignment marker lock state diagram shown in Figure 82–13. This process identifies the PCS lane number received on a particular lane of the service interface. After alignment marker lock is achieved on all lanes (4 or 20 lanes), then all inter-lane Skew is removed as shown in the PCS deskew state diagram in Figure 82–14. The Skew budget that the PCS receiver shall support is shown in Table 82–7.

Note that Skew is defined in 80.5.

Table 82–7—Skew tolerance requirements

PCS	Maximum Skew	Maximum Skew Variation
40GBASE-R	180 ns (~1856 bits)	4ns (~41 bits)
100GBASE-R	180 ns (~928 bits)	4ns (~21 bits)
100GBASE-R with RS-FEC	49 ns (~253 bits)	0.4 ns (~2 bits)

If EEE is not supported then align_status is identical to rx_align_status. Otherwise the relationship between align_status and rx_align_status is given by Figure 82–19.

82.2.14 PCS lane reorder

Transmit PCS lanes can be received on different lanes of the service interface from which they were originally transmitted due to Skew between lanes and multiplexing by the PMA. The receive PCS shall order the received PCS lanes according to the PCS lane number.

82.2.15 Alignment marker removal

After all PCS lanes are aligned and deskewed, the PCS lanes are multiplexed together in the proper order to reconstruct the original stream of blocks and the alignment markers are deleted from the data stream. The difference in rate from the deleted alignment markers is compensated for by inserting idle control characters by a function in the Receive process. Note that an alignment marker is always deleted when a given PCS Lane is in am_lock=true even if it does not match the expected alignment marker value (due to a bit error for example). Repeated alignment marker errors will result in am_lock being set to false for a given PCS Lane, but until that happens it is sufficient to delete the block in the alignment marker position.

As part of the alignment marker removal process, the BIP_3 field is compared to the calculated BIP value for each PCS lane. If a Clause 45 MDIO is implemented, then the appropriate BIP error counter register (registers 3.200 through 3.219) is incremented by one each time the calculated BIP value does not equal the value received in the BIP_3 field. The incoming bit error ratio can be estimated by dividing the BIP block error ratio by a factor of 1081344.

82.2.16 Descrambler

The descrambler is identical to that used in Clause 49, see 49.2.10 for the definition.

82.2.17 Receive process

The receive process decodes blocks to produce RXD<63:0> and RXC<7:0> for transmission to the XLGMII/CGMII. One XLGMII/CGMII data transfer is decoded from each block. The receive process must insert idle control characters to compensate for the removal of alignment markers. If the PCS receive process spans multiple clock domains, it may also perform clock rate compensation via the deletion of idle control characters or sequence ordered sets or the insertion of idle control characters.

The receive process decodes blocks as specified in the receive state diagram shown in Figure 82–17.

82.2.18 Test-pattern checker

When the receive channel is operating in scrambled idle test-pattern mode, the scrambled idle test-pattern checker checks the bits received via *inst*:IS_UNITDATA_*i*.indication primitives.

The scrambled idle test-pattern checker utilizes the block lock state diagram, the alignment marker state diagram, the PCS deskew state diagram, and the descrambler operating as they do during normal data reception. The BER monitor state diagram is disabled during receive test-pattern mode. When align_status is true and the scrambled idle receive test-pattern mode is active, the scrambled idle test-pattern checker observes the sync header and the output from the descrambler. When the sync header and the output of the descrambler is the all idle pattern, a match is detected. When operating in scrambled idle test pattern, the test-pattern error counter counts blocks with a mismatch. Any mismatch indicates an error and shall increment the test-pattern error counter. Because of the error multiplication characteristics of the descrambler, the incoming bit error ratio can be estimated by dividing the 66-bit block error ratio by a factor of 124.

If a Clause 45 MDIO is implemented, then control of the test-pattern reception is from the BASE-R PCS test-pattern control register (bit 3.42.2). In addition errors are counted in the BASE-R PCS test-pattern error counter register (3.43.15:0).

82.2.19 Detailed functions and state diagrams

82.2.19.1 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.

82.2.19.2 State variables

82.2.19.2.1 Constants

EBLOCK_R<71:0>
: 72 bit vector to be sent to the XLGMII/CGMII 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 XLGMII/CGMII containing one Local Fault ordered set. The Local Fault ordered set is defined in 81.3.4.

LBLOCK_T<65:0>
: 66 bit vector to be sent to the PMA containing one Local Fault ordered set.

82.2.19.2.2 Variables

align_status
: A variable set by the PCS deskew process to reflect the status of the PCS lane-to-lane alignment. Set true when all lanes are synchronized and aligned, set false when the deskew process is not complete.

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_align_status controlled according to Figure 82–13.

alignment_valid
: Boolean variable that is set true if all PCS lanes are aligned. It is valid when each lane is in am_lock, with each PCS lane locked to a unique alignment marker from Table 82–2 or Table 82–3, and when all PCS lanes are deskewed. Otherwise, alignment_valid is false.

am_lock<*x*>
: Boolean variable that is set true when receiver acquires alignment marker delineation for a given lane of the service interface, where $x = 0{:}3$ for 40GBASE-R and $x = 0{:}19$ for 100GBASE-R.

am_slip_done
: Boolean variable that is asserted true when the AM_SLIP requested by the alignment marker lock state diagram has been completed indicating that the next candidate 66-bit block position can be tested.

am_status
: A Boolean variable that is true when all lanes are in am_lock and false when at least one lane is not in am_lock.

am_counter_done
: Boolean variable that indicates that the alignment marker counter is done.

am_valid
: Boolean variable that is set true if received block rx_coded is a valid alignment marker. A valid alignment marker matches one of the encodings in Table 82–2, Table 82–3, Table 82–5, or Table 82–6, excluding the BIP_3 and BIP_7 fields.

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 to evaluate the header of the next block.

block_lock<*x*>
: Boolean variable that is set true when receiver acquires block delineation for a given lane of the service interface, where $x = 0{:}3$ for 40GBASE-R and $x = 0{:}19$ for 100GBASE-R.

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 according to Figure 82–12.

current_am
: This variable holds the lane number of the current alignment marker. This is compared to the variable first_am to determine if we have alignment marker lock and is always $n \times 16384$ 66-bit blocks away from the first_am.

enable_deskew
: A Boolean variable that indicates the enabling and disabling of the deskew process. Blocks may be discarded whenever deskew is enabled. True when deskew is enabled, false when deskew is disabled.

first_am
: A variable that holds the lane number of the first alignment marker that is recognized on a given lane. This is used later to compare to future alignment markers.

hi_ber
: Boolean variable which is asserted true when the ber_cnt equals or exceeds 97 indicating a bit error ratio $>10^{-4}$

lane_mapping<*x*>
: This variable indicates which PCS lane is received on lane *x* of the service interface when am_lock<*x*> = true, where *x* = 0:3 for 40GBASE-R and *x* = 0:19 for 100GBASE-R.

PCS_status
: A Boolean variable that is true when align_status is true and hi_ber is false.

r_block_type
: This variable contains the rx_coded<65:0> vector classification results, returned by the R_TYPE function. It can assume one of the following five values {C,S,T,D,E}, as defined by the R_TYPE function.

r_block_type_next
: This variable contains the rx_coded<65:0> vector classification results, returned by the R_TYPE_NEXT function. It can assume one of the following five values {C,S,T,D,E}, as defined by the R_TYPE_NEXT function.

r_test_mode
: Boolean variable that is asserted true when the receiver is in test-pattern mode.

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_align_status
: Variable used by the PCS lane deskew process to reflect the status of the PCS lane-to-lane alignment. This variable is set true when all lanes are synchronized and aligned, set false when the deskew process is not complete.

rx_block_lock<x>
: Variable used in Figure 82–12 to reflect the status of the code-group delineation for each lane. 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 82–5. The leftmost bit in the figure is rx_coded<0> and the rightmost bit is rx_coded<65>.

rx_down_count
: The value that results from the bit-wise exclusive-OR of the Count Down (CD3) byte and the M0 byte of the current received Rapid Alignment Marker (see 82.2.9).

rx_raw<71:0>
: Vector containing one XLGMII/CGMII transfer. RXC<0> through RXC<7> are from rx_raw<0> through rx_raw<7>, respectively. RXD<0> through RXD<63> are from rx_raw<8> through rx_raw<71>, respectively.

sh_valid
: Boolean variable 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 deasserted otherwise.

signal_ok
: Boolean variable that is set based on the most recently received value of *inst*:IS_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.

t_block_type
: This variable contains the tx_raw<71:0> vector classification results, returned by the T_TYPE function. It can assume one of the following five values {C,S,T,D,E}, as defined by the T_TYPE function.

test_am
: Boolean variable that is set true when a new block is available for testing and false when FIND_1ST state is entered. A new block is available for testing when the Block Sync process has accumulated enough bits from the PMA to evaluate the next block

test_sh
: Boolean variable that is set true when a new sync header is available for testing and false when TEST_SH and TEST_SH2 state are entered. A new sync header is available for testing when the Block Sync process has accumulated enough bits from the PMA 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 82–5. The leftmost bit in the figure is tx_coded<0> and the rightmost bit is tx_coded<65>.

tx_raw<71:0>
: Vector containing one XLGMII/CGMII transfer. TXC<0> through TXC<7> are placed in tx_raw<0> through tx_raw<7>, respectively. TXD<0> through TXD<63> are placed in tx_raw<8> through tx_raw<71>, respectively.

The following variables are used only for the EEE capability.

down_count_done
: Boolean variable that indicates that the down_count counter has reached zero.

down_count_enable
: Boolean variable controlling decrement of the counter down_count. This variable is set by the LPI transmit state diagram.

energy_detect
: A parameter generated by the PMA/PMD sublayer to reflect the state of the received signal. In the PMD this has the same definition as parameter signal_detect and is passed through without modification by the PMA (and FEC).

first_rx_lpi_active
: Boolean variable first_rx_lpi_active is set true when the receiver is in state RX_ACTIVE in the LPI receive state diagram (see Figure 82–19) and R_TYPE(rx_coded) = LI and is otherwise false.

LPI_FW
: Boolean variable controlling the wake mode for the LPI transmit and receive functions. This variable is set true when the link is to use the fast wake mechanism, and false when the link is to use the optional deep sleep mechanism for each direction. This variable defaults true and may only be set to false if the optional deep sleep mode is supported.

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 LPI indication:
: A Boolean variable indicating 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).

rx_mode
: A variable reflecting state of the LPI receive function as defined in Figure 82–19.The parameter has one of two values: DATA and QUIET.

scrambler_bypass
: This Boolean variable indicates whether the Tx PCS scrambler is to be bypassed in order to assist rapid synchronization following low power idle. When set to TRUE, the PCS passes the unscrambled data from the scrambler input rather than the scrambled data from the scrambler output. The scrambler continues to operate normally, shifting input data into the delay line. When scrambler_bypass is set to FALSE the PCS passes 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.

tx_mode

A variable reflecting state of the LPI transmit function as defined in Figure 82–18. When tx_mode is set to QUIET the sublayer may go into a low power state.

82.2.19.2.3 Functions

AM_SLIP

Causes the next candidate block position to be tested. The precise method for determining the next candidate block position is not specified and is implementation dependent. However, an implementation shall ensure that all possible blocks are evaluated.

DECODE(rx_coded<65:0>)

Decodes the 66-bit vector returning rx_raw<71:0>, which is sent to the XLGMII/CGMII. The DECODE function shall decode the block as specified in 82.2.3.

ENCODE(tx_raw<71:0>)

Encodes the 72-bit vector returning tx_coded<65:0> of which tx_coded<65:2> is sent to the scrambler. The two bits of the sync header bypass the scrambler. The ENCODE function shall encode the block as specified in 82.2.3.

R_TYPE(rx_coded<65:0>)

This function classifies the current rx_coded<65:0> vector as belonging to one of the following types, depending on its contents. The classification results are returned via the r_block_type variable.

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/ or /LI/;
b) A block type field of 0x4B.
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 the following:
a) 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.

Valid control characters are specified in Table 82–1.

NOTE—A PCS that does not support EEE classifies vectors containing one or more /LI/ control characters as type E.

R_TYPE_NEXT

This function classifies the 66-bit rx_coded vector that immediately follows the current rx_coded<65:0> vector as belonging to one of the five types defined in R_TYPE, depending on its contents. It is intended to perform a prescient end of packet check. The classification results are returned via the r_block_type_next variable.

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_TYPE = (tx_raw<71:0>)

This function classifies each 72-bit tx_raw vector as belonging to one of the following types depending on its contents. The classification results are returned via the t_block_type variable.

Values: C; The vector contains one of the following:
a) Eight valid control characters other than /O/, /S/, /T/, /LI/, and /E/;
b) One valid ordered set.
LI; For EEE capability, this vector contains eight /LI/ characters.
S; The vector contains an /S/ in its first character, 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 XLGMII/CGMII control code specified in Table 82–1. A valid ordered set consists of a valid /O/ character in the first character and data characters in the seven characters following the /O/. A valid /O/ is any character with a value for O code in Table 82–1.
NOTE—A PCS that does not support EEE classifies vectors containing one or more /LI/ control characters as type E.

82.2.19.2.4 Counters

am_counter
: This counter counts 66-bit blocks that separate two consecutive alignment markers. If the optional EEE deep sleep capability is supported, when rx_lpi_active is TRUE and LPI_FW is FALSE, the terminal count is 7 for 100GBASE-R PCS and 15 for 40GBASE-R PCS. If the optional EEE deep sleep capability is not supported, or rx_lpi_active is FALSE, the terminal count is 16383.

am_invld_cnt
: Count of the number of invalid alignment markers seen in a row. This counter is always reset to all zeros when a valid marker is detected.

ber_count
: A 22-bit counter that counts each time BER_BAD_SH state is entered. This counter is reflected in MDIO register bits 3.33.13:8 and 3.44.15:0.

ber_cnt
: Count up to a maximum of 97 of the number of invalid sync headers within the current 1.25 ms (40GBASE-R) or 500 μs (100GBASE-R) period.

errored_block_count:
: When the receiver is in normal mode, this 22-bit counter counts once for each time RX_E state is entered. This counter is reflected in MDIO register bits 3.33.7:0 and 3.45.13:0.

sh_cnt
: Count of the number of sync headers checked within the current 64 or 1024 block window.

sh_invld_cnt
: Count of the number of invalid sync headers within the current 64 or 1024 block window.

test_pattern_error_count:
: When the receiver is in test-pattern mode, this 16-bit counter counts errors as described in 82.2.18. This counter is reflected in MDIO register bits 3.43.15:0.

The following counters are used only for the EEE capability:

down_count
: A counter that is used in rapid alignment markers and is decremented each time a RAM is sent while variable down_count_enable = true. The counter initial value is set by the LPI transmit state diagram. When the down_count counter reaches zero, it sets the variable down_count_done = true.

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).

82.2.19.2.5 Timers

State diagram timers follow the conventions of 14.2.3.2.

*x*us_timer
: Timer that is triggered every 1.25 ms +1%, –25% for 40GBASE-R or 500 µs +1%, –25% for 100GBASE-R.

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 sets 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 sets 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 is set to a value no larger than the maximum value given for T_{WR} in Table 82–9. When the timer reaches terminal count, it sets 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 sets the rx_wf_timer_done = true.

scr_byp_timer
: This timer is started when the PCS transmitter enters the TX_SCR_BYPASS state. The timer terminal count is set to T_{BYP}. When the timer reaches terminal count, it sets the scr_byp_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 sets 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 sets 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 sets the tx_tw_timer_done = true.

tx_tw2_timer
: This timer is started when the PCS transmitter enters the TX_WAKE2 state. The timer terminal count is set to T_{WL2}. When the timer reaches terminal count, it sets the tx_tw2_timer_done = true.

82.2.19.3 State diagrams

The 40GBASE-R PCS shall implement four block lock processes as depicted in Figure 82–12. The 100GBASE-R PCS shall implement twenty block lock processes as depicted in Figure 82–12. A block lock process operates independently on each lane. Each block lock process looks for 64 valid sync headers in a row to declare lock. A valid sync header is either a 01 or a 10. Once in lock, the lock process looks for 65 invalid sync headers within a 1024 sync window to declare out of lock. An invalid sync header is a 11 or 00. Once block lock is achieved on a lane, then the alignment marker process starts.

The 40GBASE-R PCS shall implement four alignment marker lock processes as depicted in Figure 82–13. The 100GBASE-R PCS shall implement twenty alignment marker lock processes as depicted in Figure 82–13. An alignment marker lock process operates independently on each lane. The alignment marker lock state diagram shown in Figure 82–13 determines when the PCS has obtained alignment marker lock to the received bit stream for a given lane of the service interface. Each alignment marker lock process looks for two valid alignment markers 16384 66-bit blocks apart to gain alignment marker lock. On a given lane of the service interface, lane markers must match each other and an entry from Table 82–2 for 100GBASE-R or Table 82–3 for 40GBASE-R. Note that the BIP_3 and BIP_7 fields are excluded from the markers when making a match to each other or the tables. Once in lock, a lane will go out of alignment marker lock if four markers are received in a row that do not match the alignment marker that the lane is currently locked to. When the alignment marker lock process achieves lock for a lane, and if a Clause 45 MDIO is implemented, the PCS shall record the number of the PCS lane received on that lane of the service interface in the appropriate lane mapping register (3.400 to 3.419).

The PCS shall run the deskew process as depicted in Figure 82–14. The PCS deskew process is responsible for determining if the PCS is capable of presenting coherent data to the XLGMII/CGMII. The deskew process ensures that all PCS lanes have alignment marker lock, are locked to different alignment markers, and that the Skew is within the boundaries of what the PCS can deskew.

The BER Monitor state diagram shown in Figure 82–15 monitors the received aggregate signal for high bit error ratio. The high BER state shall be entered if 97 invalid 66-bit sync headers are detected within a 500 μs window for 100GBASE-R, or a 1.25 ms window for 40GBASE-R. The high BER state is exited once there are less than 97 invalid sync headers in the same window.

The Transmit state diagram shown in Figure 82–16 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 may not be operational during reset. Thus, the Local Fault ordered sets may not appear on the PMA/FEC service interface.

The Receive state diagram shown in Figure 82–17 controls the decoding of received blocks. It makes exactly one transition for each receive block processed.

The PCS shall perform the functions of block lock, alignment marker lock, PCS deskew, BER Monitor, Transmit, and Receive as specified in the respective state diagrams.

82.2.19.3.1 LPI state diagrams

A PCS that supports the EEE capability shall implement the LPI transmit and receive processes as shown in Figure 82–18 and Figure 82–19. 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.

Following a period of LPI, the receiver is required to achieve block synchronization within the wake-up time specified (see Figure 82–19). The implementation of the block synchronization state diagram should use

techniques to ensure that block lock is achieved with minimal numbers of slip attempts. If fast wake is selected then the receiver is expected to maintain sufficient state to allow much faster wake up.

The LPI functions shall use timer values for these state diagrams as shown in Table 82–8 for transmit and Table 82–9 for receive.

Table 82–8—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.	0.9	1.1	µs
T_{QL}	Local Quiet Time from when tx_mode is set to QUIET to entry into the TX_WAKE state.	1.7	1.8	ms
T_{WL}	Time spent in the TX_WAKE state.	1.5	1.6	µs
T_{WL2}	Time spent in the TX_WAKE2 state.	2.4	2.5	µs
T_{BYP}	Time spent in the TX_SCR_BYPASS state, 40 Gb/s operation.	0.9	1.1	µs
T_{BYP}	Time spent in the TX_SCR_BYPASS state, 100 Gb/s operation.	1.9	2.1	µs
T_{1U}	Time spent in the TX_ALERT state.	1.15	1.3	µs

Table 82–9—Receiver LPI timing parameters

Parameter	Description	Min	Max	Units
T_{QR}	Time the receiver waits for energy_detect to be set to true while in the RX_-QUIET state before asserting receive fault.	2	3	ms
T_{WR}	Time the receiver waits in the RX_WAKE state before indicating a wake time fault, scr_bypass_enable = FALSE.	—	4.5	µs
T_{WR}	Time the receiver waits in the RX_WAKE state before indicating a wake time fault, scr_bypass_enable = TRUE, 40 Gb/s.	—	5.5	µs
T_{WR}	Time the receiver waits in the RX_WAKE state before indicating a wake time fault, scr_bypass_enable = TRUE, 100 Gb/s.	—	6.5	µs
T_{WTF}	Wake time fault recovery time.	—	10	ms

82.3 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.

82.3.1 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PCS. Mapping of MDIO control variables to PCS control variables is

shown in Table 82–10. Mapping of MDIO status variables to PMD status variables is shown in Table 82–11.

Table 82–10—MDIO/PCS control variable mapping

MDIO control variable	PCS register name	Register/ bit number	PCS control variable
Reset	PCS control 1 register	3.0.15	reset
Loopback	PCS control 1 register	3.0.14	Loopback
Transmit test-pattern enable	BASE-R PCS test-pattern control register	3.42.3	tx_test_mode
Receive test-pattern enable	BASE-R PCS test-pattern control register	3.42.2	rx_test_mode
LPI_FW	LPI fast wake enable	3.20.0	LPI_FW

Table 82–11—MDIO/PMD status variable mapping

MDIO status variable	PCS register name	Register/ bit number	PCS status variable
BASE-R and 10GBASE-T receive link status	BASE-R and 10GBASE-T PCS status 1 register	3.32.12	PCS_status
BASE-R and 10GBASE-T PCS high BER	BASE-R and 10GBASE-T PCS status 1 register	3.32.1	hi_ber
Block *x* lock	Multi-lane BASE-R PCS alignment status 1 and 2 registers	3.50.7:0 3.51.11:0	block_lock<*x*>
Lane *x* aligned	Multi-lane BASE-R PCS alignment status 3 and 4 registers	3.52.7:0 3.53.11:0	am_lock<*x*>
PCS lane alignment status	Multi-lane BASE-R PCS alignment status 1 register	3.50.12	align_status
BER	BASE-R and 10GBASE-T PCS status 2 register BER high order counter register	3.33.13:8 3.44.15:0	ber_count
Errored blocks	BASE-R and 10GBASE-T PCS status 2 register Errored blocks high order counter register	3.33.7:0 3.45.13:0	errored_block_count
Test-pattern error counter	BASE-R PCS test-pattern error counter register	3.43.15:0	test_pattern_error_count
BIP error counter, lane *x*	BIP error counter, lane *x* register	3.200 through 3.219	bip_counter
Lane *x* mapping	Lane *x* mapping register	3.400 through 3.419	lane_mapping

Table 82–11—MDIO/PMD status variable mapping *(continued)*

MDIO status variable	PCS register name	Register/ bit number	PCS status variable
Tx LPI indication	Tx LPI indication	3.1.9	Tx LPI indication
Tx LPI received	Tx LPI received	3.1.11	Tx LPI received
Rx LPI indication	Rx LPI indication	3.1.8	Rx LPI indication
Rx LPI received	Rx LPI received	3.1.10	Rx LPI received
Wake_error_counter	Wake_error_counter	3.22	Wake_error_counter

82.4 Loopback

If a Clause 45 MDIO is implemented, then the PCS shall be placed in the loopback mode when the loopback bit from the PCS control 1 register (bit 3.0.14) is set to a one. In this mode, the PCS shall accept data on the transmit path from the XLGMII/CGMII and return it on the receive path to the XLGMII/CGMII. In addition, the PCS shall transmit what it receives from the XLGMII/CGMII to the PMA/FEC sublayer, and shall ignore all data presented to it by the PMA/FEC sublayer.

82.5 Delay constraints

The maximum delay contributed by the 40GBASE-R PCS (sum of transmit and receive delays at one end of the link) shall be no more than 11264 BT (22 pause_quanta or 281.6 ns). The maximum delay contributed by the 100GBASE-R PCS (sum of transmit and receive delays at one end of the link) shall be no more than 35328 BT (69 pause_quanta or 353.28 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

82.6 Auto-Negotiation

The following requirements apply to a PCS used with a 40GBASE-KR4 PMD, 40GBASE-CR4 PMD, 100GBASE-CR10, 100GBASE-CR4, 100GBASE-KR4, or 100GBASE-KP4 PMD where 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.

.

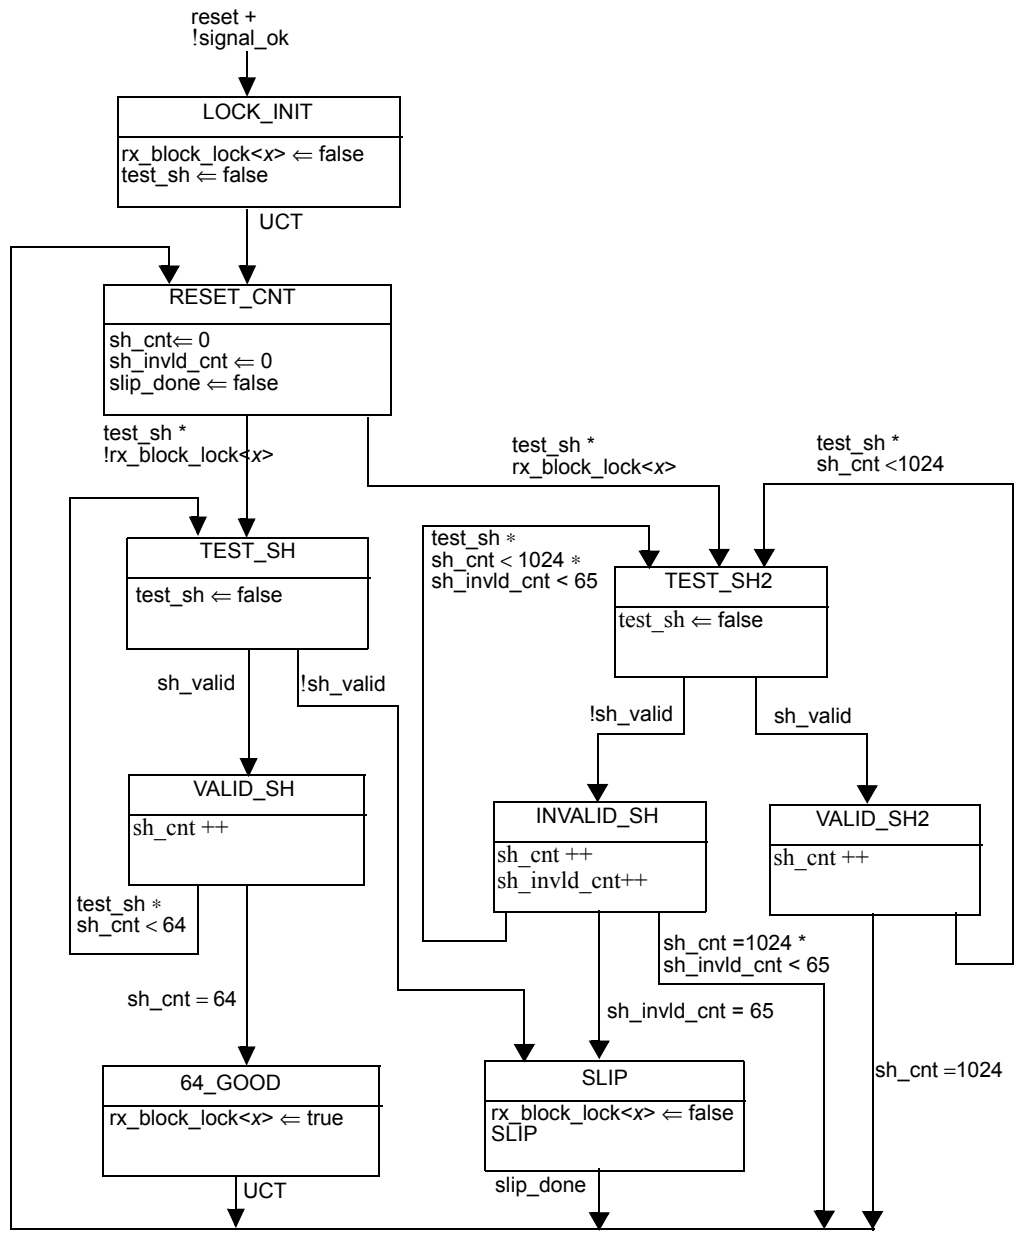

NOTE— rx_block_lock<*x*> refers to the received lane *x* of the service interface, where *x* = 0:3 (for 40GBASE-R) or 0:19 (for 100GBASE-R)

Figure 82–12—Block lock state diagram

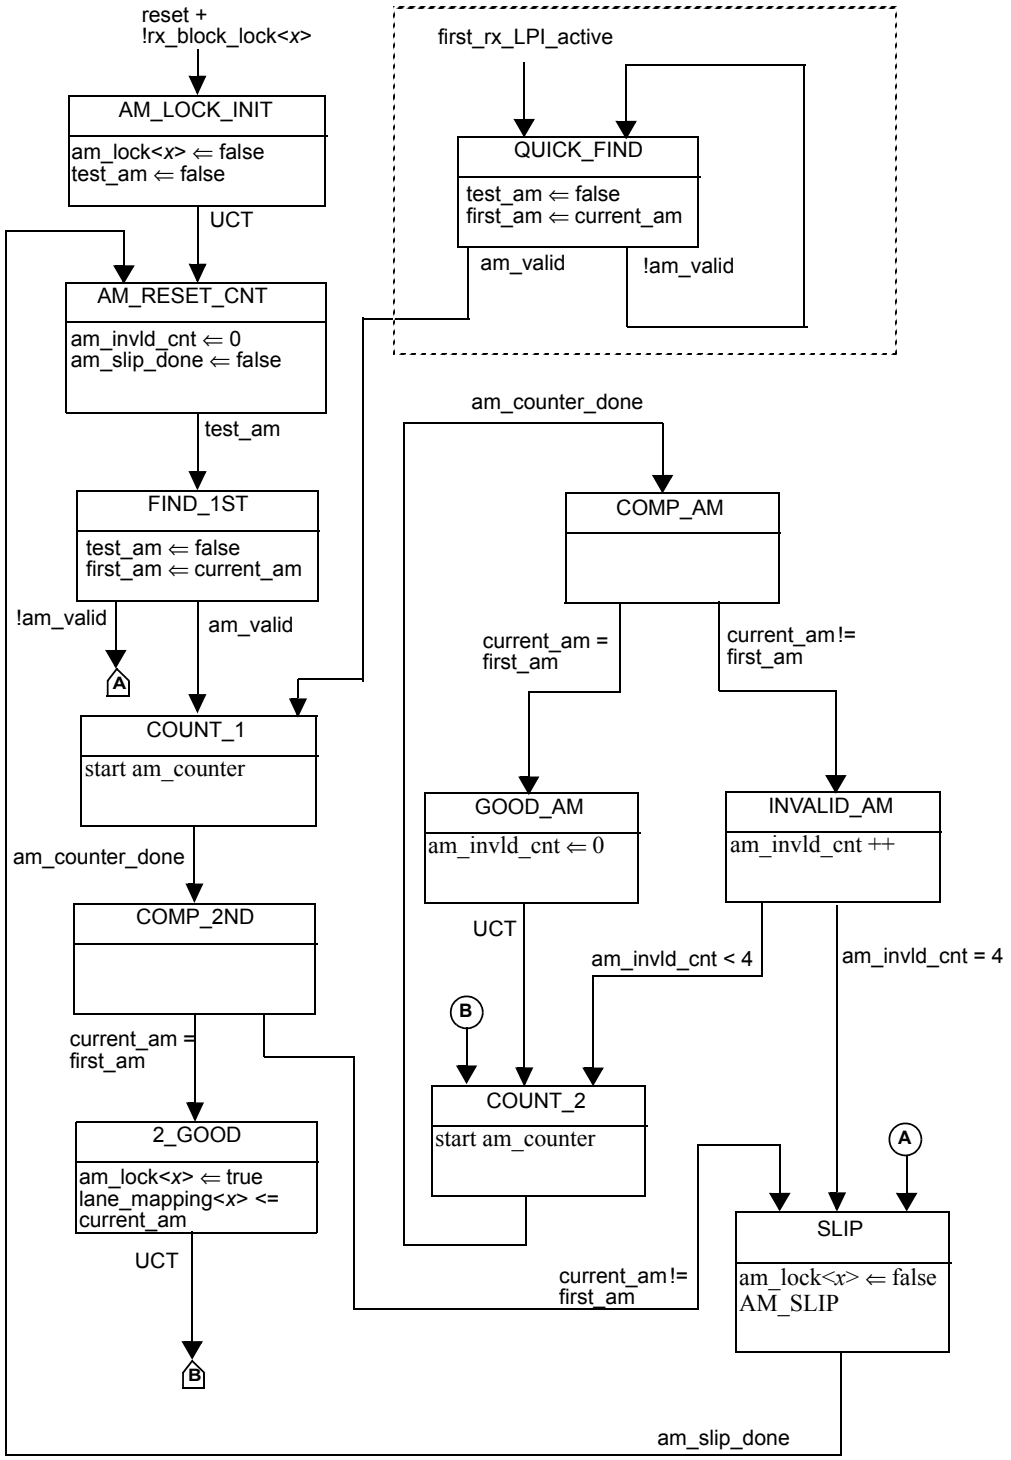

NOTE 1— am_lock<*x*> refers to the received lane *x* of the service interface, where *x* = 0:3 (for 40GBASE-R) or 0:19 (for 100GBASE-R)

NOTE 2—Optional state (inside the dotted box) is only required to support EEE capability.

Figure 82–13—Alignment marker lock state diagram

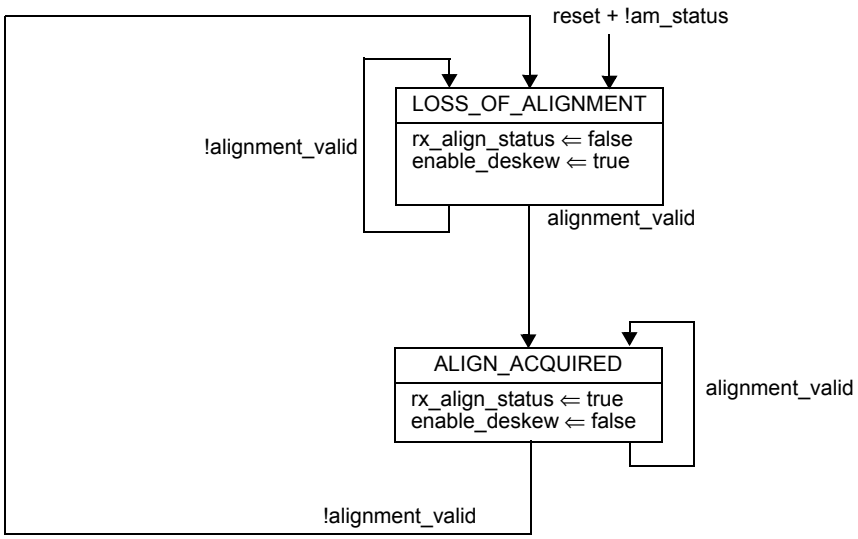

Figure 82–14—PCS deskew state diagram

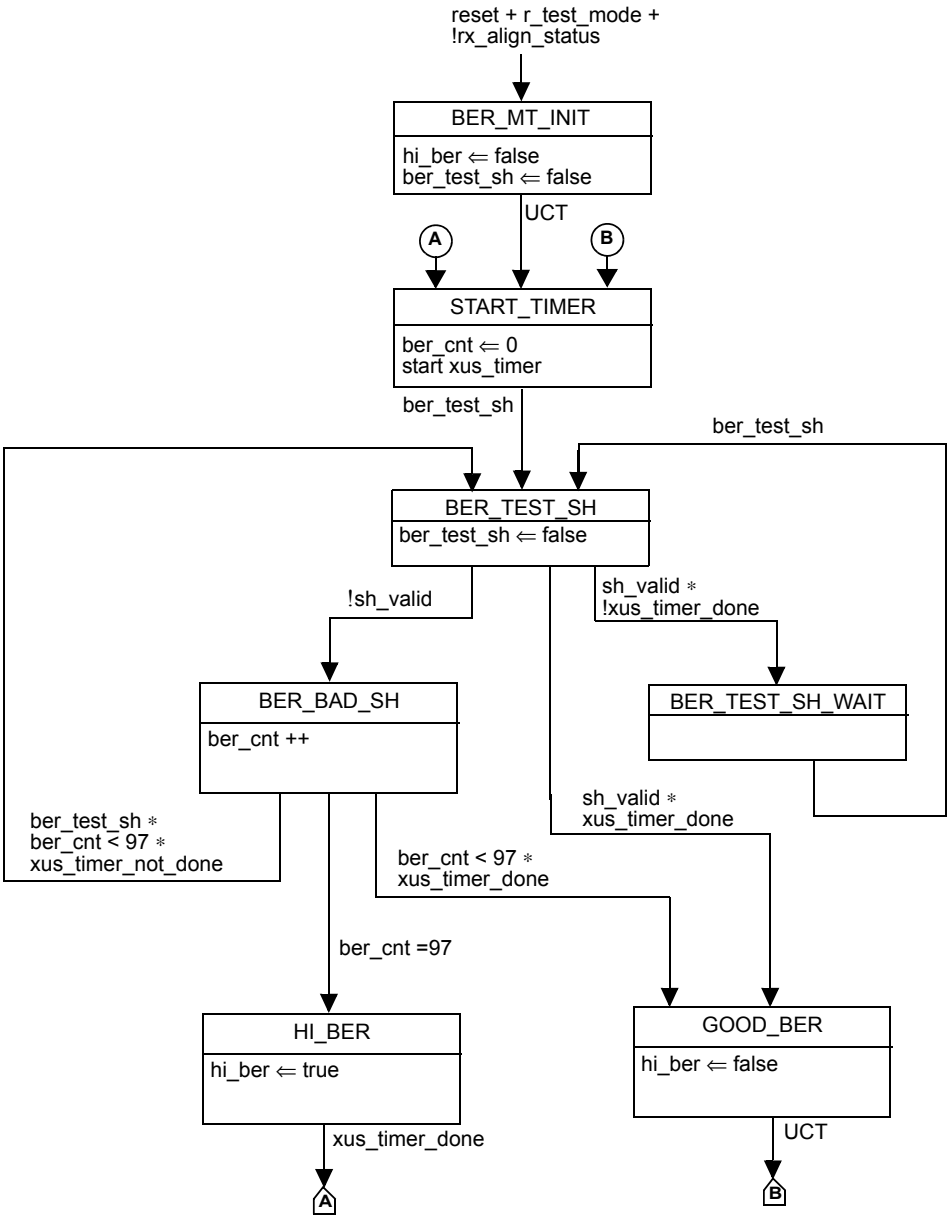

Figure 82–15—BER monitor state diagram

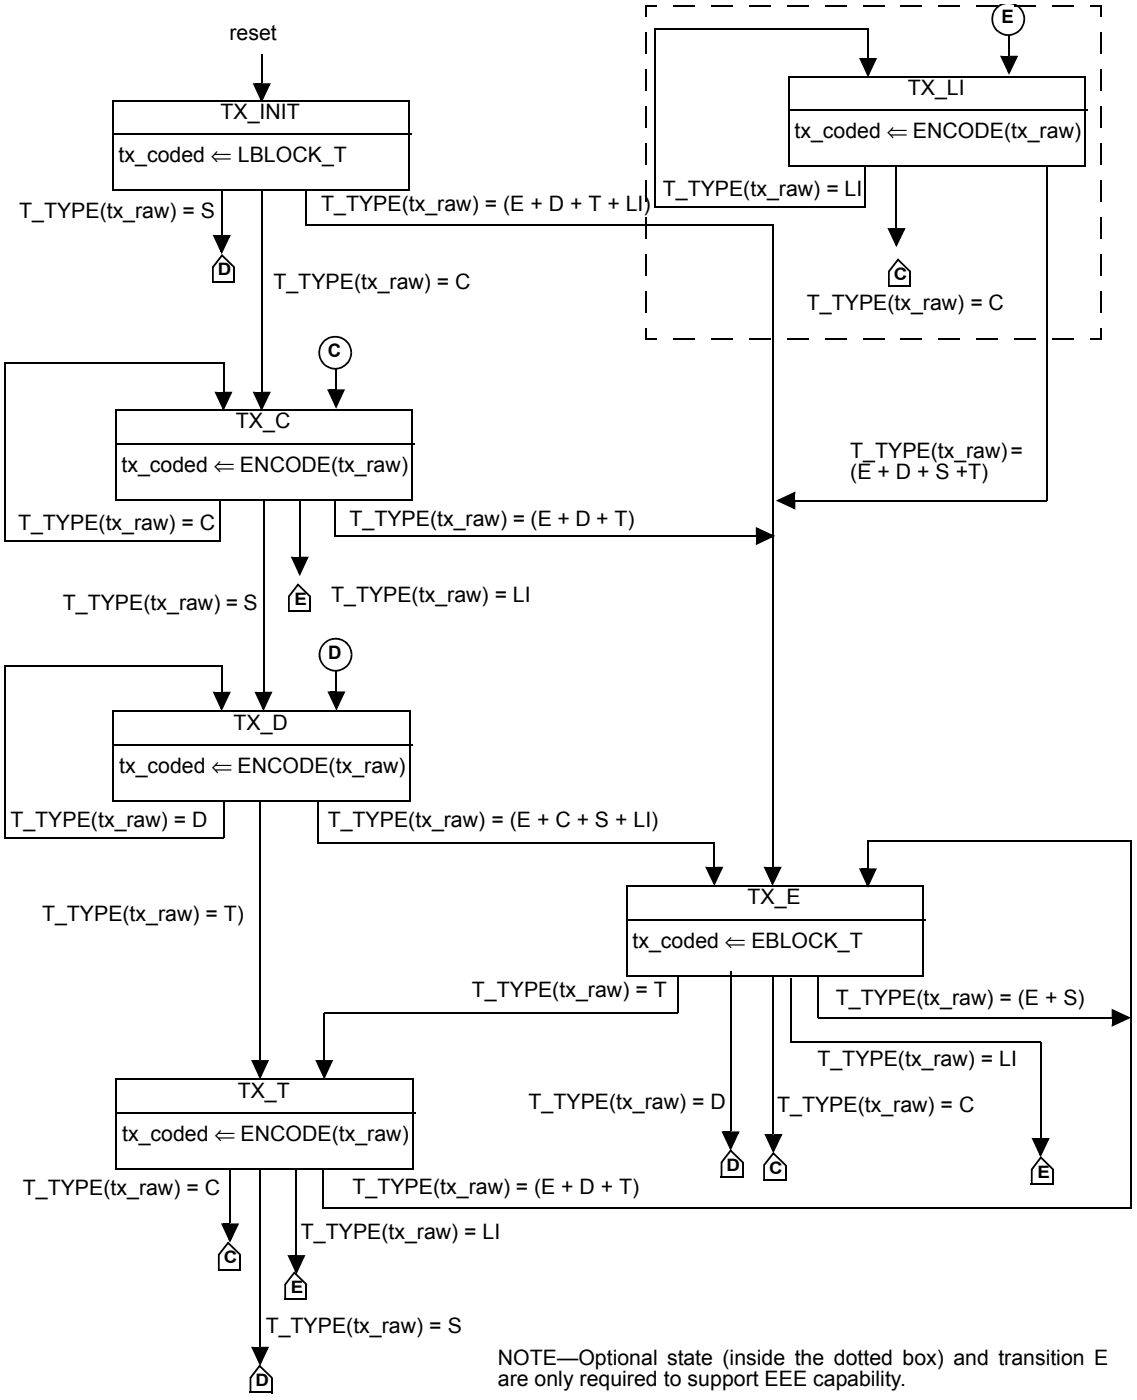

Figure 82–16—Transmit state diagram

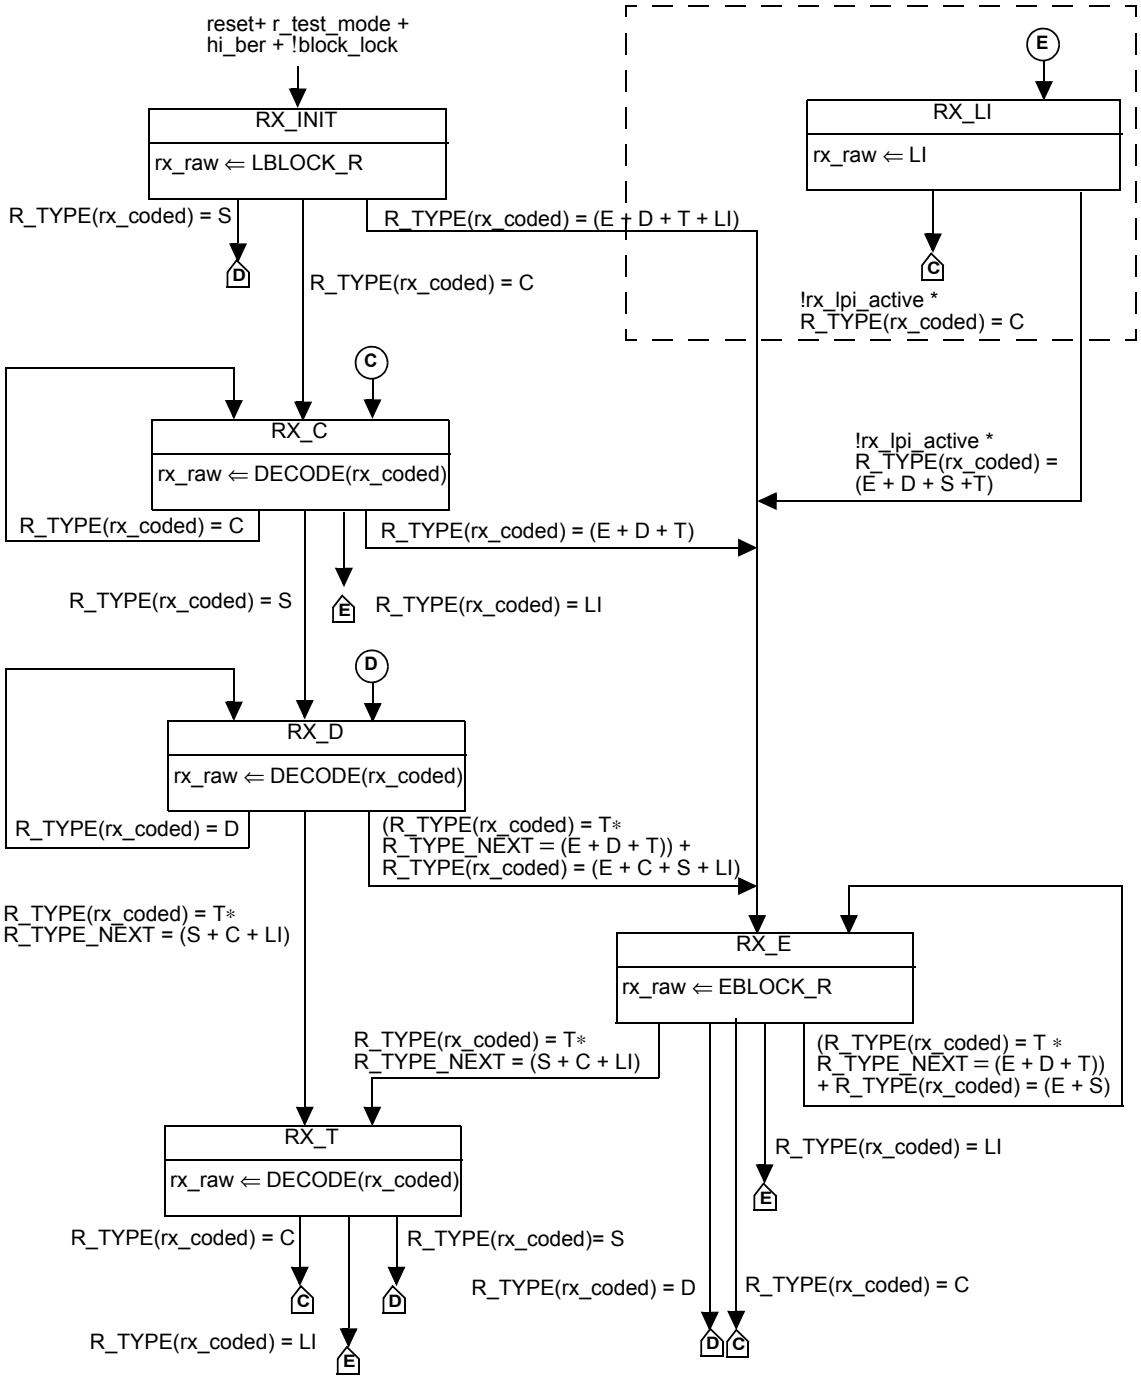

Figure 82–17—Receive state diagram

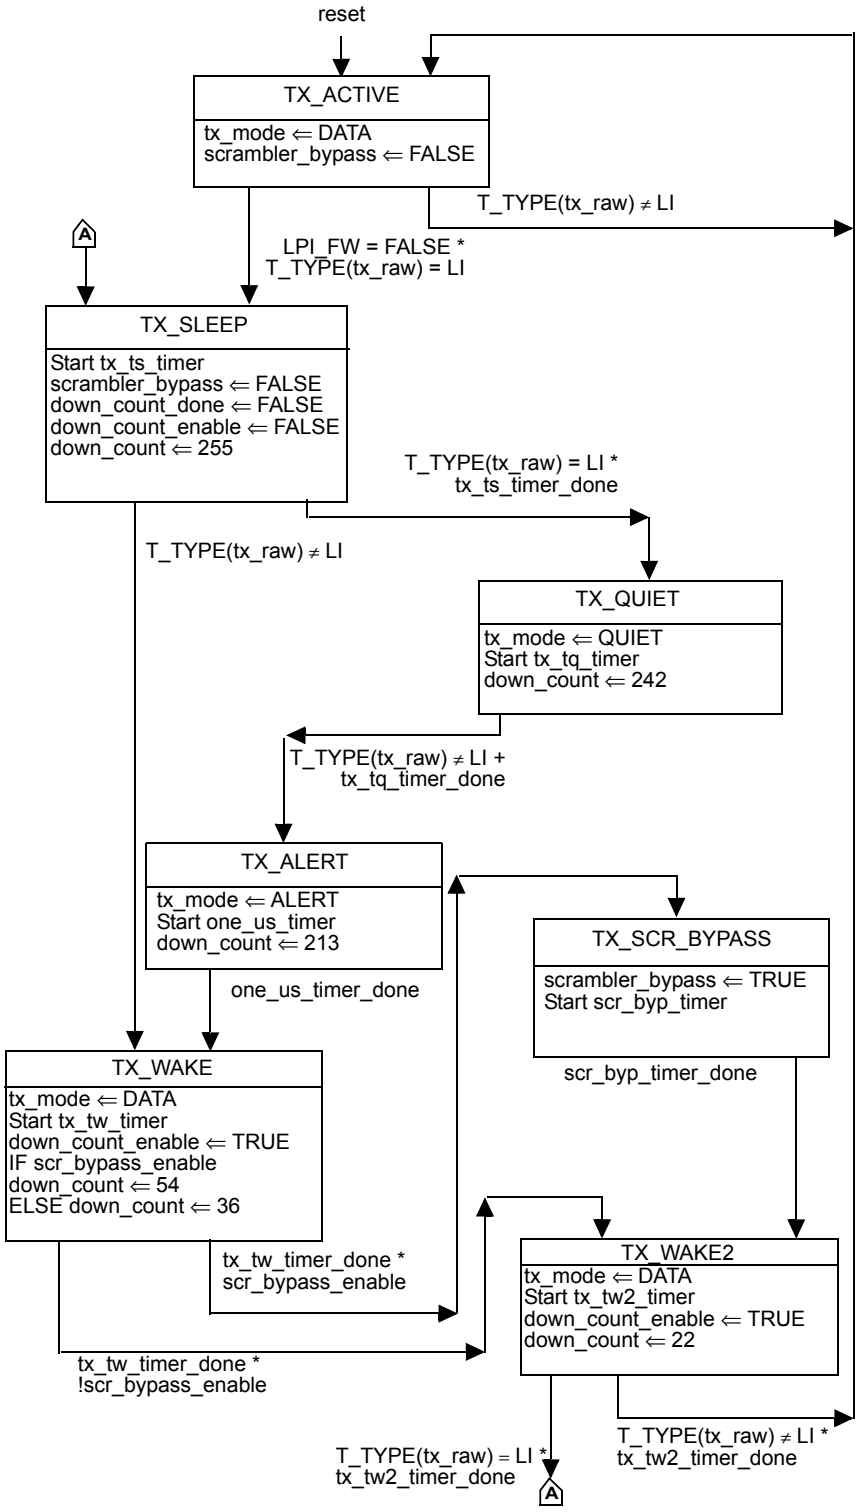

Figure 82–18—LPI Transmit state diagram

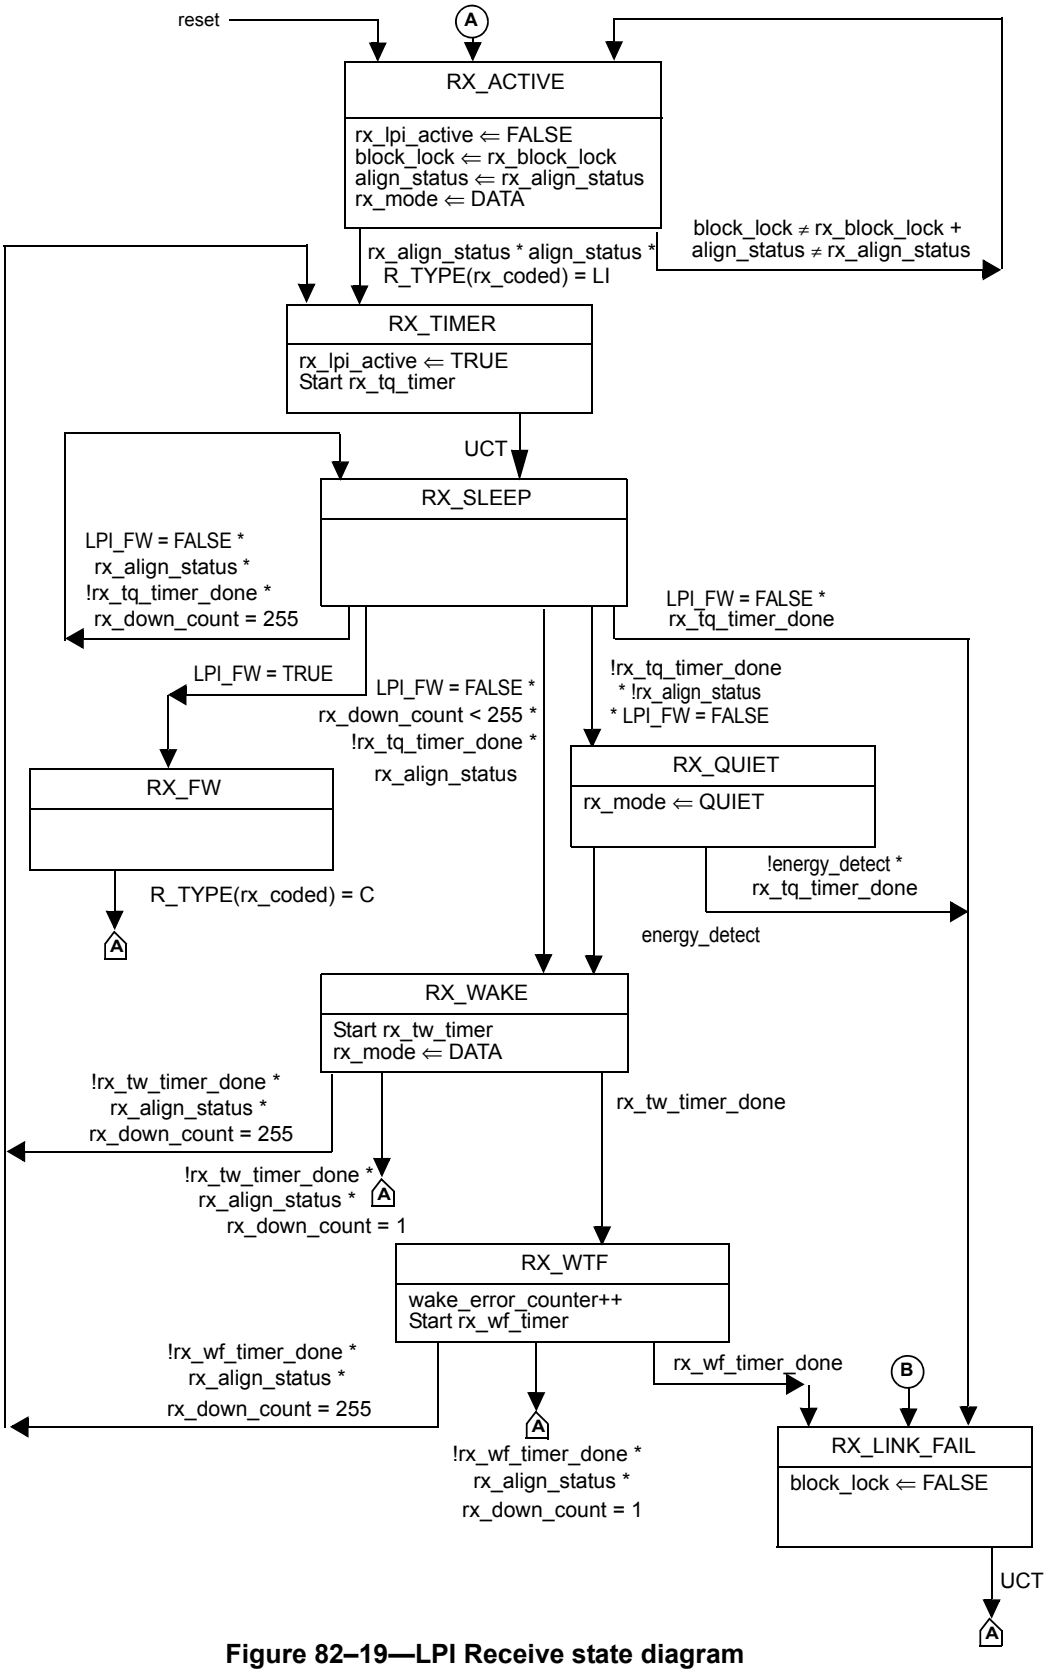

Figure 82–19—LPI Receive state diagram

82.7 Protocol implementation conformance statement (PICS) proforma for Clause 82, Physical Coding Sublayer (PCS) for 64B/66B, type 40GBASE-R and 100GBASE-R[7]

82.7.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 82, Physical Coding Sublayer (PCS) for 64B/66B, type 40GBASE-R and 100GBASE-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.

82.7.2 Identification

82.7.2.1 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).	

82.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 82, Physical Coding Sublayer (PCS) for 64B/66B, type 40GBASE-R and 100GBASE-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	

[7] *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.

82.7.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*PCS40	PCS for 40GBASE-R	82.1.1		O	Yes [ ] No [ ]
*PCS100	PCS for 100GBASE-R	82.1.1		O	Yes [ ] No [ ]
XGE40	XLGMII logical interface	81, 82.1.4	Logical interface is supported	O	Yes [ ] No [ ]
XGE100	XLGMII logical interface	81, 82.1.4	Logical interface is supported	O	Yes [ ] No [ ]
MD	MDIO	45, 82.3	Registers and interface supported	O	Yes [ ] No [ ]
PMA	Supports operation directly connected to a PMA	82.1.4.2		O.1	Yes [ ] No [ ]
FEC	Supports operation directly connected to a FEC sublayer	82.1.4.2		O.1	Yes [ ] No [ ]
*JTM	Supports test-pattern mode	82.2.1		PMA:M	Yes [ ] No [ ] N/A[ ]
*LPI	Implementation of LPI	82.2.3.4		O	Yes [ ] No [ ]

82.7.4 PICS Proforma Tables for PCS, type 40GBASE-R and 100GBASE-R

82.7.4.1 Coding rules

Item	Feature	Subclause	Value/Comment	Status	Support
C1	Encoder (and ENCODE function) implements the code as specified	82.2.3 and 82.2.19.2.3		M	Yes [ ] No [ ]
C2	Decoder (and DECODE function) implements the code as specified	82.2.3 and 82.2.19.2.3		M	Yes [ ] No [ ]
C3	Only valid block types are transmitted	82.2.3.3		M	Yes [ ] No [ ]
C4	Invalid block types are treated as an error	82.2.3.3		M	Yes [ ] No [ ]
C5	Only valid control characters are transmitted	82.2.3.4		M	Yes [ ] No [ ]
C6	Invalid control characters are treated as an error	82.2.3.4		M	Yes [ ] No [ ]
C7	Idles do not interrupt data	82.2.3.6		M	Yes [ ] No [ ]
C8	IDLE control code insertion and deletion	82.2.3.6	Insertion or Deletion in groups of 8 /I/s	M	Yes [ ] No [ ]
C9	Sequence ordered set deletion	82.2.3.9	Only one whole ordered set of two consecutive sequence ordered sets may be deleted	M	Yes [ ] No [ ]

82.7.4.2 Scrambler and Descrambler

Item	Feature	Subclause	Value/Comment	Status	Support
S1	Scrambler	82.2.5	Performs as shown in Figure 49–8	M	Yes [ ] No [ ]
S2	Descrambler	82.2.16	Performs as shown in Figure 49–10	M	Yes [ ] No [ ]

82.7.4.3 Deskew and Reordering

Item	Feature	Subclause	Value/Comment	Status	Support
DR1	Deskew	82.2.13	Able to deskew up to the value in Table 82–7	M	Yes [ ] No [ ]
DR2	Reordering	82.2.14	Performs reordering.	M	Yes [ ] No [ ]

82.7.4.4 Alignment Markers

Item	Feature	Subclause	Value/Comment	Status	Support
AM1	Alignment marker insertion	82.2.7	Alignment markers are inserted periodically as described in section 82.2.7	M	Yes [ ] No [ ]
AM2	Alignment marker form	82.2.7	Alignment markers are formed as described in section 82.2.7	M	Yes [ ] No [ ]
AM3	Lane mapping	82.2.19.3	PCS lane number is captured	MD:M	Yes [ ] No [ ]

82.7.5 Test-pattern modes

Item	Feature	Subclause	Value/Comment	Status	Support
JT1	Scrambled idle transmit test-pattern generator is implemented	82.2.11	Performs as in 82.2.11	JTM:M	Yes [ ] No [ ] N/A[ ]
JT2	Scrambled idle receive test-pattern checker is implemented	82.2.18	Performs as in 82.2.18	JTM:M	Yes [ ] No [ ] N/A[ ]
JT3	Transmit and receive test-pattern modes can operate simultaneously	82.2.1		JTM:M	Yes [ ] No [ ] N/A[ ]

82.7.5.1 Bit order

Item	Feature	Subclause	Value/Comment	Status	Support
B1	Transmit bit order	82.2.3.2	Placement of bits into the PCS lanes as shown in Figure 82–3	M	Yes [ ] No [ ]
B2	Receive bit order	82.2.3.2	Placement of bits into the XLGMII/CGMII as shown in Figure 82–4	M	Yes [ ] No [ ]

82.7.6 Management

Item	Feature	Subclause	Value/Comment	Status	Support
M1	Alternate access to PCS Management objects is provided	82.3		O	Yes [ ] No [ ]

82.7.6.1 State diagrams

Item	Feature	Subclause	Value/Comment	Status	Support
SM1	40GBASE-R Block Lock	82.2.19.3	Implements 4 block lock processes as depicted in Figure 82–12	PCS40:M	Yes [ ] No [ ]
SM2	100GBASE-R Block Lock	82.2.19.3	Implements 20 block lock processes as depicted in Figure 82–12	PCS100:M	Yes [ ] No [ ]
SM3	The SLIP function evaluates all possible bit positions	82.2.19.2.3		M	Yes [ ] No [ ]
SM4	40GBASE-R Alignment Marker Lock	82.2.19.3	Implements 4 alignment marker lock processes as depicted in Figure 82–13	PCS40:M	Yes [ ] No [ ]
SM5	100GBASE-R Alignment Marker Lock	82.2.19.3	Implements 20 alignment marker lock processes as depicted in Figure 82–13	PCS100:M	Yes [ ] No [ ]
SM6	The AM_SLIP functions evaluates all possible blocks	82.2.19.2.3		M	Yes [ ] No [ ]
SM7	40GBASE-R PCS deskew state diagram	82.2.19.3	Meets the requirements of Figure 82–14	PCS40:M	Yes [ ] No [ ]
SM8	100GBASE-R PCS deskew state diagram	82.2.19.3	Meets the requirements of Figure 82–14	PCS100:M	Yes [ ] No [ ]
SM9	40GBASE-R BER Monitor	82.2.19.3	Meets the requirements of Figure 82–15 with xus_timer_done set to 1.25 ms	PCS40:M	Yes [ ] No [ ]
SM10	100GBASE-R BER Monitor	82.2.19.3	Meets the requirements of Figure 82–15 with xus_timer_done set to 500 μs	PCS100:M	Yes [ ] No [ ]
SM11	40GBASE-R Transmit process	82.2.19.3	Meets the requirements of Figure 82–16	PCS40:M	Yes [ ] No [ ]
SM12	100GBASE-R Transmit process	82.2.19.3	Meets the requirements of Figure 82–16	PCS100:M	Yes [ ] No [ ]
SM13	40GBASE-R and 100GBASE-R Receive process	82.2.19.3	Meets the requirements of Figure 82–17	PCS40:M	Yes [ ] No [ ]
SM14	100GBASE-R Receive process	82.2.19.3	Meets the requirements of Figure 82–17	PCS100:M	Yes [ ] No [ ]

82.7.6.2 Loopback

Item	Feature	Subclause	Value/Comment	Status	Support
L1	Supports loopback	82.4	Performs as in 82.4	M	Yes [ ] No [ ] N/A[ ]
L2	When in loopback, transmits what it receives from the XLG-MII/CGMII	82.4	Performs as in 82.4	M	Yes [ ] No [ ]

82.7.6.3 Delay constraints

Item	Feature	Subclause	Value/Comment	Status	Support
TIM1	PCS Delay Constraint	82.5	No more than 11264 BT for sum of transmit and receive path delays for 40GBASE-R and 35328 BT for 100GBASE-R	M	Yes [ ] No [ ]

82.7.6.4 Auto-Negotiation for Backplane Ethernet functions

Item	Feature	Subclause	Value/Comment	Status	Support
*AN1	Support for use with a 40GBASE-KR4, 40GBASE-CR4, 100GBASE-CR10, 100GBASE-CR4, 100GBASE-KR4, or 100GBASE-KP4 PMD	82.6	AN technology dependent interface described in Clause 73	O	Yes [ ]
AN2	AN_LINK.indication primitive	82.6	Support of the primitive AN_LINK.indication(link_status), when the PCS is used with 40GBASE-KR4, 40GBASE-CR4, 100GBASE-CR10, 100GBASE-CR4, 100GBASE-KR4, or 100GBASE-KP4 PMD	AN1:M	Yes [ ]
AN3	link_status parameter	82.6	Takes the value OK or FAIL, as described in 82.6	AN1:M	Yes [ ]
AN4	Generation of AN_LINK.indication primitive	82.6	Generated when the value of link_status changes	AN1:M	Yes [ ]

82.7.6.5 LPI functions

Item	Feature	Subclause	Value/Comment	Status	Support
LP-01	Support for both wake modes	82.2.19.2.2	Variable LPI_FW may be true or false	LPI:O	Yes [ ] No [ ]
LP-02	Insertion and deletion of LPIs in groups of 8	82.2.3.6		LPI:M	Yes [ ] No [ ]
LP-03	BIP statistics during LPI	82.2.8	BIP statistics not updated during LPI	LPI:M	Yes [ ] No [ ]
LP-04	RAM insertion	82.2.9	Insertion of Rapid Alignment Markers meets the requirements of 82.2.9	LPI:M	Yes [ ] No [ ]
LP-05	Transmit state diagrams	82.2.19.3	Support LPI operation in Figure 82–16	LPI:M	Yes [ ] No [ ]
LP-06	Receive state diagrams	82.2.19.3	Support LPI operation in Figure 82–17	LPI:M	Yes [ ] No [ ]
LP-07	LPI transmit state diagrams	82.2.19.3.1	Meets the requirements of Figure 82–18	LPI:M	Yes [ ] No [ ]
LP-08	LPI receive state diagrams	82.2.19.3.1	Meets the requirements of Figure 82–19	LPI:M	Yes [ ] No [ ]
LP-09	LPI transmit timing	82.2.19.3.1	Meets the requirements of Table 82–8	LPI:M	Yes [ ] No [ ]
LP-10	LPI receive timing	82.2.19.3.1	Meets the requirements of Table 82–9	LPI:M	Yes [ ] No [ ]

83. Physical Medium Attachment (PMA) sublayer, type 40GBASE-R and 100GBASE-R

83.1 Overview

83.1.1 Scope

This clause specifies the Physical Medium Attachment sublayer (PMA) that is common to two families of (40 Gb/s and 100 Gb/s) Physical Layer implementations, known as 40GBASE-R and 100GBASE-R. The PMA allows the PCS (specified in Clause 82) to connect in a media-independent way with a range of physical media. The 40GBASE-R PMA(s) can support any of the 40 Gb/s PMDs in Table 80–2. The 100GBASE-R PMA(s) can support any of the 100 Gb/s PMDs in Table 80–3 or Table 80–4, but does not provide the PMD service interface for 100GBASE-KP4 (Clause 94). The terms 40GBASE-R and 100GBASE-R are used when referring generally to Physical Layers using the PMA defined in this clause.

40GBASE-R and 100GBASE-R can be extended to support any full duplex medium requiring only that the PMD be compliant with the appropriate PMA interface.

The interfaces for the inputs of the 40GBASE-R and 100GBASE-R PMAs are defined in an abstract manner and do not imply any particular implementation. The optional physical instantiation of the PMD service interfaces for 40GBASE-SR4, 40GBASE-LR4, and 100GBASE-SR10 PMDs, known as XLPPI and CPPI, are defined in Annex 86A. The PMD service interfaces for other PMDs are defined in an abstract manner according to 80.3.1. For 40GBASE-R PMAs, electrical interfaces connecting PMA sublayers, known as XLAUI, are defined in Annex 83A and Annex 83B. For 100GBASE-R PMAs, electrical interfaces connecting PMA sublayers, known as CAUI-n, are defined in Annex 83A, Annex 83B, Annex 83D, and Annex 83E.

83.1.2 Position of the PMA in the 40GBASE-R or 100GBASE-R sublayers

Figure 83–1 shows the relationship of the PMA sublayer (shown shaded) with other sublayers to the ISO Open System Interconnection (OSI) reference model.

83.1.3 Summary of functions

The following is a summary of the principal functions implemented (when required) by the PMA in both the transmit and receive directions:

a) Adapt the PCSL formatted signal to the appropriate number of abstract or physical lanes.
b) Provide per input-lane clock and data recovery.
c) Provide bit-level multiplexing.
d) Provide clock generation.
e) Provide signal drivers.
f) Optionally provide local loopback to/from the PMA service interface.
g) Optionally provide remote loopback to/from the PMD service interface.
h) Optionally provide test-pattern generation and detection.
i) Tolerate Skew Variation.

In addition, the PMA provides receive link status information in the receive direction.

83.1.4 PMA sublayer positioning

An implementation may use one or more PMA sublayers to adapt the number and rate of the PCS lanes to the number and rate of the PMD lanes. The number of PMA sublayers required depends on the partitioning

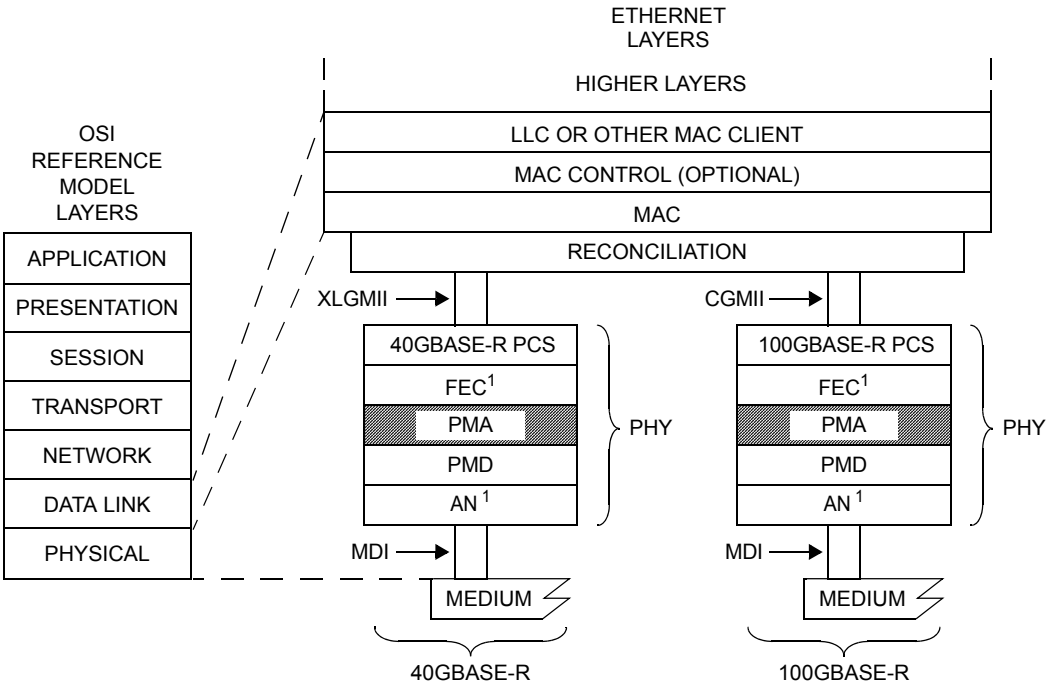

AN = AUTO-NEGOTIATION
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 83–1—40GBASE-R and 100GBASE-R PMA relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

of functionality for a particular implementation. An example is illustrated in Figure 83–2. This example illustrates the partitioning that might arise from use of a FEC device that is separate from the PCS. Additional examples are illustrated in Annex 83C. Each PMA maps the PCSLs from p PMA input lanes to q PMA output lanes in the Tx direction, and from q PMA input lanes to p PMA output lanes in the Rx direction.

Management Data Input/Output (MDIO) Manageable Device (MMD) addresses 1, 8, 9, 10, and 11 are available for addressing multiple instances of PMA sublayers (see Table 45–1 for MMD device addresses). If the PMA sublayer that is closest to the PMD is packaged with the PMD, it shares MMD 1 with the PMD. If the PMD service interface is physically instantiated as nPPI (see Annex 86A), the PMA sublayer that is closest to the PMD will be addressed as MMD 8. More addressable instances of PMA sublayers, each one separated from lower addressable instances by chip-to-chip interfaces, may be implemented and addressed allocating MMD addresses to PMAs in increasing numerical order going from the PMD toward the PCS. The example shown in Figure 83–2 could be implemented with four addressable instances: MMD 8 addressing the lowest PMA sublayer (note that this cannot share MMD 1 with the PMD as they are not packaged together in this example), MMD 9 addressing the PMA sublayer above the XLAUI/CAUI-4 below the FEC, MMD 10 addressing the PMA sublayer below the XLAUI/CAUI-10 above the FEC, and MMD 11 addressing the PMA sublayer closest to the PCS.

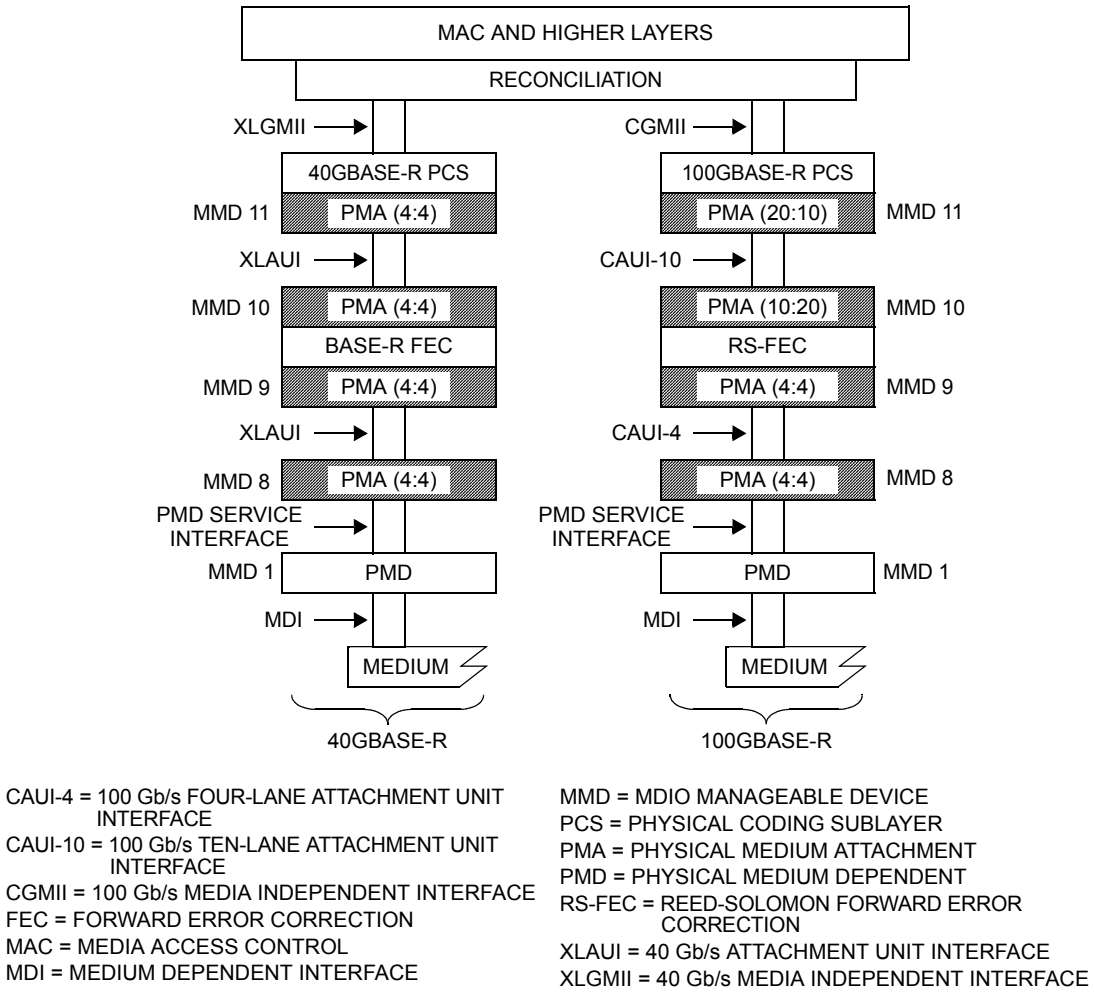

CAUI-4 = 100 Gb/s FOUR-LANE ATTACHMENT UNIT INTERFACE
CAUI-10 = 100 Gb/s TEN-LANE ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
MMD = MDIO MANAGEABLE DEVICE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION
XLAUI = 40 Gb/s ATTACHMENT UNIT INTERFACE
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

Figure 83–2—Example 40GBASE-R and 100GBASE-R PMA layering

The number of input lanes and the number of output lanes for a PMA are always divisors of the number of PCSLs. For PMA sublayers supporting 40GBASE-R PMDs, the number of PCSLs is 4, and for PMA sublayers supporting 100GBASE-R PMDs, the number of PCSLs is 20.

The following guidelines apply to the partitioning of PMAs:

a) The inter-sublayer service interface, defined in 80.3, is used for the PMA, FEC, and PMD service interfaces supporting a flexible architecture with optional FEC and multiple PMA sublayers.
   1) An instance of this interface can only connect service interfaces with the same number of lanes, where the lanes operate at the same rate.
b) XLAUI and CAUI-n are physical instantiations of the connection between two adjacent PMA sublayers.
   1) As a physical instantiation, it defines electrical and timing specification as well as requiring a receive re-timing function
   2) XLAUI is a 10.3125 GBd by 4 lane physical instantiation of the respective 40 Gb/s connection
   3) CAUI-10 is a 10.3125 GBd by 10 lane physical instantiation of the respective 100 Gb/s connection
   4) CAUI-4 is a 25.78125 GBd by 4 lane physical instantiation of the respective 100 Gb/s connection

c) The abstract inter-sublayer service interface can be physically instantiated as a XLAUI or CAUI-n, using associated PMAs to map to the appropriate number of lanes.
d) Opportunities for optional test-pattern generation, optional test-pattern detection, optional local loopback and optional remote loopback are dependent upon the location of the PMA sublayer in the implementation. See Figure 83–5.
e) A minimum of one PMA sublayer is required in a PHY.
f) A maximum of four PMA sublayers are addressable as MDIO MMDs.

83.2 PMA interfaces

All PMA variants for 40GBASE-R and 100GBASE-R signals are based on a generic specification of a bit mux function that applies to all input/output lane counts and each direction of transmission. Each direction of transmission employs one or more such bit muxes to adapt from the appropriate number of input lanes to the appropriate number of output lanes as illustrated in Figure 83–3.

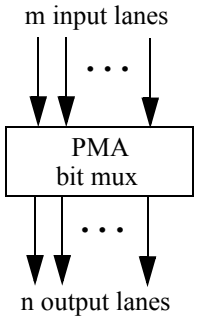

Figure 83–3—PMA bit mux used in both Tx and Rx directions

Conceptually, the PMA bit mux operates in one direction of transmission by demultiplexing PCSLs from m PMA input lanes and remultiplexing them into n PMA output lanes. The mapping of PCSLs from input to output lanes is not specified. See 83.5.2 and Figure 83–4 for details.

Figure 83–5 provides the functional block diagram of a PMA. The parameters of a PMA include the following:

— The aggregate rate supported (40GBASE-R or 100GBASE-R).
— The numbers of input and output lanes in each direction.
— Whether the PMA is adjacent to a physically instantiated interface (XLAUI/CAUI-n above or below, or PMD service interface below).
— Whether the PMA is adjacent to the PCS (or adjacent to FEC when FEC is adjacent to the PCS).
— Whether the PMA is adjacent to the PMD.

83.3 PMA service interface

The PMA service interface for 40GBASE-R and 100GBASE-R is an instance of the inter-sublayer service interface defined in 80.3. The PMA service interface primitives are summarized as follows:

PMA:IS_UNITDATA_*i*.request(tx_bit)
PMA:IS_UNITDATA_*i*.indication(rx_bit)
PMA:IS_SIGNAL.indication(SIGNAL_OK)

For a PMA with p lanes at the PMA service interface, the primitives are defined for $i = 0$ to $p - 1$.

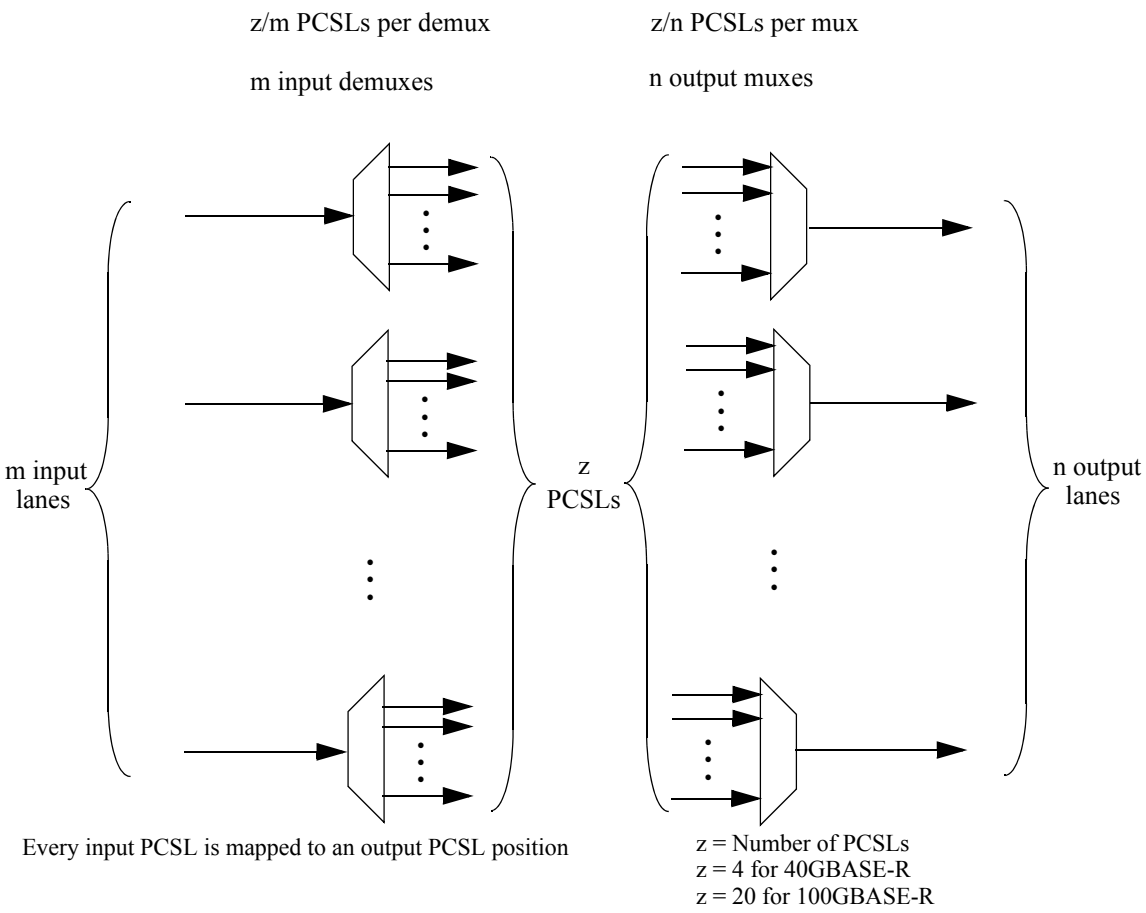

Figure 83–4—PMA bit mux operation used in both Tx and Rx directions

If the PMA client is the PCS or a BASE-R FEC sublayer (see Clause 74), the PMA (or PMA client) continuously sends four (for 40GBASE-R) or twenty (for 100GBASE-R) parallel bit streams to the PMA client (or PMA), each at the nominal signaling rate of the PCSL. If the PMA client is the 100GBASE-R RS-FEC sublayer (see Clause 91), the PMA continuously sends four parallel bit streams to the PMA client (or PMA), each at 25.78125 GBd.

If the PMA client is another PMA, for a PMA supporting a 40GBASE-R PMD, the number of PCSLs z = 4 and for a PMA supporting a 100GBASE-R PMD, the number of PCSLs z = 20. A PMA with p input lanes receives bits on each of its input lanes at z/p times the PCSL rate. Skew may exist between the bits received on each lane even though all lanes originate from the same synchronous source, so there is independence of arrival of bits on each lane.

In the Tx direction, if the bit from a PMA:IS_UNITDATA_*i*.request primitive is received over a physically instantiated interface (XLAUI/CAUI-n), clock and data are recovered on the lane receiving the bit. The bit is routed through the PMA to an output lane through a process that may demultiplex PCSLs from the input, perform any necessary buffering to tolerate Skew Variation across input lanes, and multiplex PCSLs to output lanes. The bit is sent on an output lane to the sublayer below using the *inst*:IS_UNITDATA_*k*.request (*k* not necessarily equal to *i*) primitive (see 83.4).

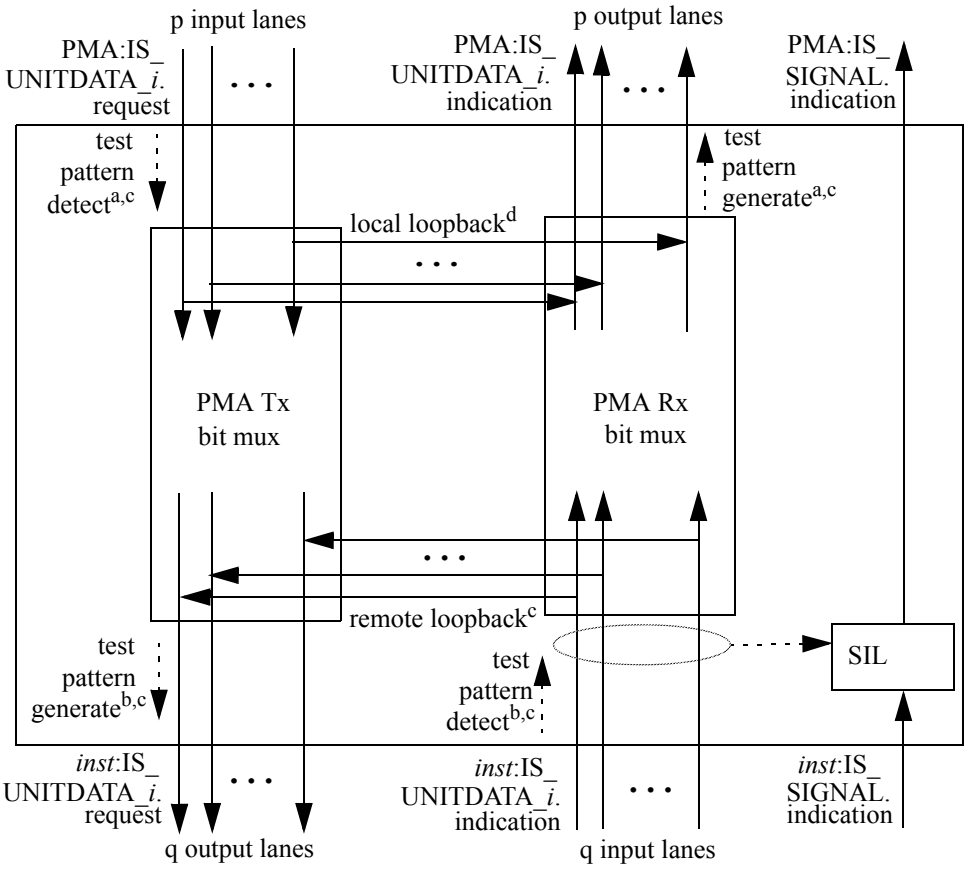

inst PMD, PMA, or FEC, depending on which sublayer is below this PMA
SIL Signal Indication Logic

[a] If physically instantiated interface (XLAUI/CAUI-n) immediately above this PMA.
[b] If physically instantiated interface (XLAUI/CAUI-n or PMD service interface) immediately below this PMA, or if this is the closest PMA to the PMD.
[c] Optional.
[d] Local loopback is required for PMAs adjacent to some PMDs, and optional for other PMAs. See 83.5.8.

Figure 83–5—PMA Functional Block Diagram

In the Rx direction, when data is being received from every input lane from the sublayer below the PMA that has a PCSL that is routed to a particular output lane at the PMA service interface, and (if necessary), buffers are filled to allow tolerating the Skew Variation that may appear between the input lanes, PCSLs are demultiplexed from the input lanes, remultiplexed to the output lanes, and bits are transferred over each output lane to the PMA client via the PMA:IS_UNITDATA_*i*.indication primitive.

The PMA:IS_SIGNAL.indication primitive is generated through a set of Signal Indication Logic (SIL) that reports signal health based on receipt of the *inst*:IS_SIGNAL.indication from the sublayer below, data being received on all of the input lanes from the sublayer below, buffers filled (if necessary) to accommodate Skew Variation, and bits being sent to the PMA client on all of the output lanes. When these conditions are met, the SIGNAL_OK parameter sent to the PMA client via the PMA:IS_SIGNAL.indication primitive will have the value OK. Otherwise, the SIGNAL_OK primitive will have the value FAIL.

If the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option is supported (see Clause 78, 78.1.3.3.1) then the inter-sublayer service interface includes four additional primitives defined as follows:

IS_TX_MODE.request
IS_RX_MODE.request
IS_ENERGY_DETECT.indication
IS_RX_TX_MODE.indication

The IS_TX_MODE.request primitive is used to communicate the state of the PCS LPI transmit function to other sublayers in the PHY. The IS_RX_MODE.request primitive is used to communicate the state of the PCS LPI receive function to other sublayers. The IS_RX_TX_MODE.indication primitive is used to communicate the state of the rx_tx_mode parameter, that reflects the inferred state of the link partner's tx_mode parameter, from the PMA to other sublayers. The IS_ENERGY_DETECT.indication primitive is used to communicate that the PMD has detected the return of energy on the interface following a period of quiescence.

A physically instantiated service interface with the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option (see 78.1.3.3.1) 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 in the ingress direction is indicated by register 1.1.9 (PMA Ingress AUI Stop Ability, PIASA) and in the egress direction by register 1.1.8 (PMA Egress AUI Stop Ability, PEASA). Transition to the low power state is enabled in the ingress direction by register 1.7.9 (PMA Ingress AUI Stop Enable, PIASE) and in the egress direction by register 1.7.8 (PMA Egress AUI Stop Enable, PEASE). The system shall not assert the enable bit for an interface unless the corresponding ability bit at the other side of the interface is also asserted. If the PIASE bit is TRUE, then the PMA may disable transmitters on the physical instantiation of the ingress AUI when aui_tx_mode is QUIET. If the PEASE bit is TRUE, then the PMA may disable transmitters on the physical instantiation of the egress AUI when tx_mode is QUIET.

83.4 Service interface below PMA

Since the architecture supports multiple PMA sublayers for various PMD lane counts and device partitioning, there are several different sublayers that may appear below a PMA, including FEC, the PMD, or another PMA. The variable *inst* represents whichever sublayer appears below the PMA (e.g., another PMA, FEC, or PMD).

The sublayer below the PMA utilizes the inter-sublayer service interface defined in 80.3. The service interface primitives provided to the PMA are summarized as follows:

inst:IS_UNITDATA_*i*.request(tx_bit)
inst:IS_UNITDATA_*i*.indication(rx_bit)
inst:IS_SIGNAL.indication(SIGNAL_OK)

The number of lanes q for the service interface matches the number of lanes expected by the PMA. The *inst*:IS_UNITDATA_*i* primitives are defined for each lane $i = 0$ to $q - 1$ of the service interface below the PMA. Note that electrical and timing specifications of the service interface are defined if the interface is physically instantiated (e.g., XLAUI/CAUI-n or nPPI), otherwise the service interface is specified only abstractly. The interface between the PMA and the sublayer below consists of q lanes for data transfer and a status indicating a good signal sent by the sublayer below the PMA (see Figure 83–5).

In the Tx direction, when data is being received via the PMA:IS_UNITDATA_*i*.request primitive from every input lane from the PMA client at the PMA service interface (see 83.3) that has a PCSL that is routed to this output lane, and (if necessary), buffers are filled to allow tolerating the Skew Variation that may appear between the input lanes from the PMA client, PCSLs are demultiplexed from the input lanes, remultiplexed to the output lanes, and bits are transferred over each output lane to the sublayer below the PMA.

In the Rx direction, if the bit is received over a physically instantiated interface (XLAUI/CAUI-n or nPPI), clock and data are recovered on the lane receiving the bit. The bit is routed through the PMA to an output

lane toward the PMA client through a process that may demultiplex PCSLs from the input, perform any necessary buffering to tolerate Skew Variation across input lanes, and multiplex PCSLs to output lanes, and finally sending the bit on an output lane to the PMA client using the PMA:IS_UNITDATA_*k*.indication (*k* not necessarily equal to *i*) primitive at the PMA service interface.

83.5 Functions within the PMA

The purpose of the PMA is to adapt the PCSL formatted signal to an appropriate number of abstract or physical lanes, to recover clock from the received signal (if appropriate), and optionally to provide test signals and loopback. Each input (Tx direction) or output (Rx direction) lane between the PMA and the PMA client carries one or more PCSLs that are bit-multiplexed. All input and output lanes between the PMA and the PMA client carry the same number of PCSLs and operate at the same nominal signaling rate. Likewise, each input (Rx direction) or output (Tx direction) lane between the PMA and the sublayer below the PMA carries one or more PCSLs that are bit-multiplexed. All input and output lanes between the PMA and the sublayer below the PMA carry the same number of PCSLs and operate at the same nominal signaling rate. As described in 83.1.4, the number of input lanes and the number of output lanes for a given PMA are divisors of the number of PCSLs for the interface type supported.

83.5.1 Per input-lane clock and data recovery

If the interface between the PMA client and the PMA is physically instantiated as XLAUI/CAUI-n, the PMA shall meet the electrical and timing specifications in Annex 83A, Annex 83B, Annex 83D, or Annex 83E as appropriate. If the interface between the sublayer below the PMA and the PMA is physically instantiated as XLAUI/CAUI-n or nPPI, the PMA shall meet the electrical and timing specifications at the service interface as specified in Annex 83A, Annex 83B, Annex 83D, Annex 83E, or Annex 86A as appropriate.

83.5.2 Bit-level multiplexing

The PMA provides bit-level multiplexing in both the Tx and Rx directions. In the Tx direction, the function is performed among the bits received from the PMA client via the PMA:IS_UNITDATA_*i*.request primitives (for PMA client lanes $i = 0$ to $p - 1$) with the result sent to the service interface below the PMA using the *inst*:IS_UNITDATA_*i*.request primitives (for service interface lanes $i = 0$ to $q - 1$), referencing the functional block diagram shown in Figure 83–5. The bit multiplexing behavior is illustrated in Figure 83–4.

The aggregate signal carried by the group of input lanes or the group of output lanes is arranged as a set of PCSLs. For PMA sublayers supporting 40GBASE-R interfaces, the number of PCSLs z is 4, and the nominal signaling rate R of each PCSL is 10.3125 GBd. For PMA sublayers supporting 100GBASE-R interfaces, the number of PCSLs z is 20, and the nominal signaling rate R of each PCSL is 5.15625 GBd.

For a PMA with m input lanes (Tx or Rx direction), each input lane carries, bit multiplexed, z/m PCSLs. Each input lane has a nominal signaling rate of $R \times z/m$. If bit x received on an input lane belongs to a particular PCSL, the next bit of that same PCSL is received on the same input lane at bit position $x+(z/m)$. The z/m PCSLs may arrive in any sequence on a given input lane.

For a PMA with n output lanes (Tx or Rx direction), each output lane carries, bit multiplexed, z/n PCSLs. Each output lane has a nominal signaling rate of $R \times z/n$. Each PCSL is mapped from a position in the sequence on one of the z/m input lanes to a position in the sequence on one of the z/n output lanes. If bit x sent on an output lane belongs to a particular PCSL, the next bit of that same PCSL is sent on the same output lane at bit position $x + (z/n)$. The PMA shall maintain the chosen sequence of PCSLs on all output lanes while it is receiving a valid stream of bits on all input lanes.

Each PCSL received in any temporal position on an input lane is transferred into a temporal position on an output lane. As the PCS (see Clause 82) has fully flexible receive logic, an implementation is free to perform the mapping of PCSLs from input lanes to output lanes without constraint. Figure 83–6 illustrates one possible bit ordering for a 10:4 PMA bit mux. Other bit orderings are also valid.

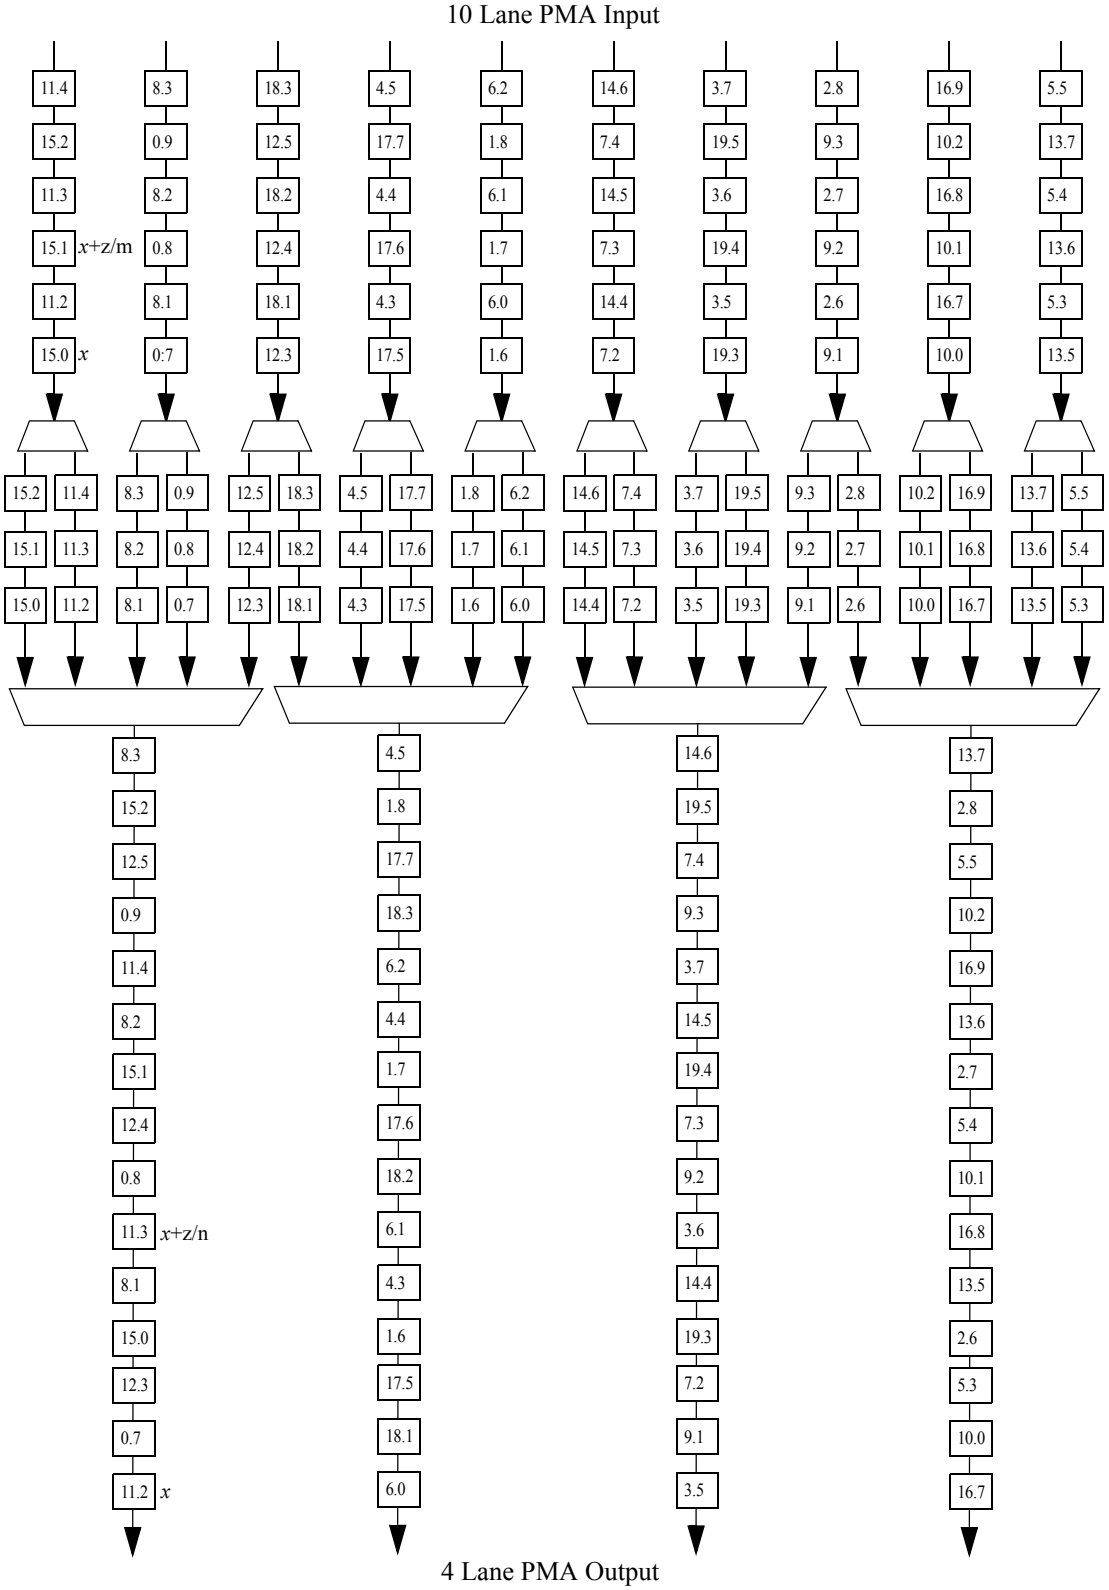

NOTE—i.k indicates bit k on PCSL i. Skew may exist between PCSLs.

Figure 83–6—Example 10: 4 PMA bit mux

83.5.3 Skew and Skew Variation

The Skew (relative delay) between the PCSLs must be kept within limits so that the information on the lanes can be reassembled by the PCS.

Any PMA that combines PCSLs from different input lanes onto the same output lane must tolerate Skew Variation between the input lanes without changing the PCSL positions on the output. Skew and Skew Variation are defined in 80.5. The limits for Skew and Skew Variation at physically instantiated interfaces are specified at Skew points SP0, SP1, and SP2 in the transmit direction and SP5, SP6, and SP7 in the receive direction as defined in 80.5 and illustrated in Figure 80–6, Figure 80–7, and Figure 80–8.

83.5.3.1 Skew generation toward SP0

In an implementation with one or more physically instantiated CAUI-n interfaces, the PMA that sends data in the transmit direction toward the CAUI-n that is closest to the RS-FEC (SP0 in Figure 80–8) shall produce no more than 29 ns of Skew between PCSLs toward the CAUI-n, and no more than 200 ps of Skew Variation.

83.5.3.2 Skew generation toward SP1

In an implementation with one or more physically instantiated XLAUI/CAUI-n interfaces, the PMA that sends data in the transmit direction toward the XLAUI/CAUI-n that is closest to the PMD (SP1 in Figure 80–6 and Figure 80–7) shall produce no more than 29 ns of Skew between PCSLs toward the XLAUI/CAUI-n, and no more than 200 ps of Skew Variation.

83.5.3.3 Skew tolerance at SP1

In an implementation with one or more physically instantiated XLAUI/CAUI-n interfaces, the PMA service interface that receives data in the transmit direction from the XLAUI/CAUI-n (SP1 in Figure 80–6 and Figure 80–7) shall tolerate the maximum amount of Skew Variation allowed at SP1 (200 ps) between input lanes while maintaining the bit ordering and position of each PCSL on each PMA lane in the transmit direction (toward the PMD).

83.5.3.4 Skew generation toward SP2

In an implementation with a physically instantiated PMD service interface, the PMA adjacent to the PMD service interface shall generate no more than 43 ns of Skew, and no more than 400 ps of Skew Variation between output lanes toward the PMD service interface (SP2 in Figure 80–6 and Figure 80–7). If there is a physically instantiated XLAUI/CAUI-n as well, then the Skew measured at SP1 is limited to no more than 29 ns of Skew and no more than 200 ps of Skew Variation.

83.5.3.5 Skew tolerance at SP5

In an implementation with a physically instantiated PMD service interface, the PMA adjacent to the PMD service interface (SP5) shall tolerate the maximum amount of Skew Variation allowed at SP5 (3.6 ns) between output lanes from the PMD service interface while maintaining the bit ordering and position of each PCSL on each PMA lane in the receive direction (toward the PCS).

83.5.3.6 Skew generation at SP6

In an implementation with one or more physically instantiated XLAUI/CAUI-n interfaces, at SP6 (the receive direction of the XLAUI/CAUI-n closest to the PCS), the PMA or group of PMAs between the PMD and the XLAUI/CAUI-n closest to the PCS shall deliver no more than 160 ns of Skew, and no more than 3.8 ns of Skew Variation between output lanes toward the XLAUI/CAUI-n in the Rx direction. If there is a

physically instantiated PMD service interface as well, the Skew measured at SP5 is limited to no more than 145 ns of Skew and no more than 3.6 ns of Skew Variation. If there is no physically instantiated PMD service interface, the Skew measured at SP4 is limited to no more than 134 ns of Skew, and no more than 3.4 ns of Skew Variation.

83.5.3.7 Skew tolerance at SP6

In an implementation with one or more physically instantiated XLAUI/CAUI-n interfaces, the PMA between the XLAUI/CAUI-n closest to the PCS and the PCS shall tolerate the maximum amount of Skew Variation allowed at SP6 (3.8 ns) between input lanes while maintaining the bit order and position of PCSLs on lanes sent in the receive direction towards the PCS.

83.5.3.8 Skew generation toward SP7

In an implementation with one or more physically instantiated CAUI-n interfaces and RS-FEC, at SP7 (the receive direction of the CAUI-n closest to the PCS), the PMA or group of PMAs between the RS-FEC and the CAUI-n closest to the PCS shall deliver no more than 160 ns of Skew, and no more than 3.8 ns of Skew Variation between output lanes toward the CAUI-n in the Rx direction.

83.5.4 Delay constraints

The maximum cumulative delay contributed by up to four PMA stages in a PHY (sum of transmit and receive delays at one end of the link) shall meet the values specified in Table 83–1. A description of overall system delay constraints and the definitions for bit-times and pause_quanta can be found in 80.4 and its references.

Table 83–1—Delay constraints

Sublayer	Maximum (bit time)	Maximum (pause_quanta)	Maximum (ns)
40GBASE-R PMA	4096	8	102.4
100GBASE-R PMA	9216	18	92.16

83.5.5 Clocking architecture

A PMA with m input lanes and n output lanes shall clock the output lanes at m/n times the rate of the input lanes. This applies in both the Tx and Rx directions of transmission. In the case where the interfaces between the PMA client and the PMA and/or the PMA and the sublayer below the PMA are physically instantiated, the PMA may derive its input clock(s) from the electrical interface on one or more of the input lanes, and generate the output clock(s) with an appropriate PLL multiplier/divider circuit.

There is no requirement that the PMA clock all output lanes in unison. Examples of independent clocking of output lanes include the following:

— The case where the number of input and output lanes are equal (the PMA is provided for retiming and regeneration of the signal). This may be implemented without any rearrangement of PCSLs between input lanes and output lanes (although rearrangements are allowed), and such a PMA may be implemented by driving each output lane using the clock recovered from the corresponding input lane.

— If the number of input and output lanes have a common factor, the PMA may be partitioned such that PCSLs from a subset of the input lanes are mapped only to a subset of the output lanes (for example, a 10:4 PMA could be implemented as two 5:2 PMAs). The output clock for one subset may be independent of the output clock for other subset(s).

83.5.6 Signal drivers

For cases where the interface between the PMA client and the PMA, or between the PMA and the sublayer below the PMA represent a physically instantiated interface, the PMA provides electrical signal drivers for that interface. The electrical and jitter/timing specifications for these interfaces appear in

— Annex 83A, which specifies the XLAUI/CAUI-10 interface for chip-to-chip applications.
— Annex 83B, which specifies the XLAUI/CAUI-10 interface for chip-to-module applications.
— Annex 83D, which specifies the CAUI-4 interface for chip-to-chip applications.
— Annex 83E, which specifies the CAUI-4 interface for chip-to-module applications.
— 86.2, which specifies the PMD service interface for 40GBASE-SR4 and 100GBASE-SR10 PMDs.
— 87.2, which specifies the PMD service interface for 40GBASE-LR4 and 40GBASE-ER4 PMDs.
— Annex 86A, which specifies the Parallel Physical Interface (XLPPI and CPPI), an optional physical instantiation of the PMD service interface for 40GBASE-SR4, 40GBASE-LR4, and 100GBASE-SR10 PMDs.

83.5.7 Link status

The PMA shall provide link status information to the PMA client using the PMA:IS_SIGNAL.indication primitive. The PMA continuously monitors the link status reported by the service interface below from the *inst*:IS_SIGNAL.indication primitive, and uses this as input to Signal Indication Logic (SIL) to determine the link status to report to the layer above. Other inputs to the SIL may include the status of clock and data recovery on the lanes from the service interface below the PMA and whether buffers/FIFOs have reached the required fill level to accommodate Skew Variation so that data is being sent on the output lanes.

83.5.8 PMA local loopback mode

PMA local loopback shall be provided by the PMA adjacent to the PMD for 40GBASE-KR4, 40GBASE-CR4, 100BASE-CR10, 100GBASE-KR4, and 100GBASE-CR4 PMDs. PMA local loopback mode is optional for other PMDs or for PMAs not adjacent to the PMD. If it is implemented, it shall be as described in this subclause (83.5.8).

The PMA sublayer may provide a local loopback function. The function involves looping back each input lane to the corresponding output lane. Each bit received from the PMA:IS_UNITDATA_*i*.request(tx_bit) primitive is looped back in the direction of the PCS using the PMA:IS_UNITDATA_*i*.indication(rx_bit) primitive.

During local loopback, the PMA performs normal bit muxing of PCSLs per 83.5.2 onto the lanes in the Tx direction toward the service interface below the PMA.

Ability to perform this function is indicated by the Local_loopback_ability status variable. If a Clause 45 MDIO is implemented, this variable is accessible through bit 1.8.0 (45.2.1.7.15). A device is placed in local loopback mode when the Local_loopback_enable control variable is set to one, and removed from local loopback mode when this variable is set to zero. If a Clause 45 MDIO is implemented, this variable is accessible through PMA/PMD control 1 register (bit 1.0.0, see 45.2.1.1.5).

83.5.9 PMA remote loopback mode (optional)

PMA remote loopback mode is optional. If implemented, it shall be as described in this subclause (83.5.9).

Remote loopback, if provided, should be implemented in a PMA sublayer close enough to the PMD to maintain the bit sequence on each individual PMD lane. When remote loopback is enabled, each bit received over a lane of the service interface below the PMA via *inst*:IS_UNITDATA_*i*.indication is looped back to the corresponding output lane toward the PMD via *inst*:IS_UNITDATA_*i*.request. Note that the service interface below the PMA can be provided by the FEC, PMD, or another PMA sublayer.

During remote loopback, the PMA performs normal bit muxing of PCSLs per 83.5.2 onto the lanes in the Rx direction towards the PMA client.

The ability to perform this function is indicated by the Remote_loopback_ability status variable. If a Clause 45 MDIO is implemented, this variable is accessible through bit 1.13.15 (45.2.1.12.1). A device is placed in remote loopback mode when the Remote_loopback_enable control variable is set to one, and removed from remote loopback mode when this variable is set to zero. If a Clause 45 MDIO is implemented, this variable is accessible through PMA/PMD Control register 1 (bit 1.0.1, see 45.2.1.1.4).

83.5.10 PMA test patterns (optional)

Where the output lanes of the PMA appear on a physically instantiated interface XLAUI/CAUI-n or the PMD service interface (whether or not it is physically instantiated), the PMA may optionally generate and detect test patterns. These test patterns are used to test adjacent layer interfaces for an individual PMA sublayer or to perform testing between a physically instantiated interface of a PMA sublayer and external testing equipment.

The ability to generate each of the respective test patterns in each direction of transmission are indicated by the PRBS9_Tx_generator_ability, PRBS9_Rx_generator_ability, PRBS31_Tx_generator_ability, and PRBS31_Rx_generator_ability status variables, which if a Clause 45 MDIO is implemented are accessible through bits 1.1500.5, 1.1500.4, 1.1500.3, and 1.1500.1, respectively (see 45.2.1.123).

The ability to check PRBS31 test patterns in each direction of transmission are indicated by the PRBS31_Tx_checker_ability and PRBS31_Rx_checker_ability status variables, which, if a Clause 45 MDIO is implemented, are accessible through bits 1.1500.2 and 1.1500.0, respectively (see 45.2.1.123).

If supported, when send Tx PRBS31 test pattern is enabled by the PRBS31_enable and PRBS_Tx_gen_enable control variables, the PMA shall generate a PRBS31 pattern (as defined in 49.2.8) on each of the lanes toward the service interface below the PMA via the *inst*:IS_UNITDATA_*i*.request primitive. To avoid correlated crosstalk, it is highly recommended that the PRBS31 patterns generated on each lane be generated from independent, random seeds or at a minimum offset of 20 000 UI between the PRBS31 sequence on any lane and any other lane. When send Tx PRBS31 test pattern is disabled, the PMA returns to normal operation performing bit multiplexing as described in 83.5.2. If a Clause 45 MDIO is implemented, the PRBS31_enable and PRBS_Tx_gen_enable control variables are accessible through bits 1.1501.7 and 1.1501.3 (see 45.2.1.124).

If supported, when send Rx PRBS31 test pattern is enabled by the PRBS31_enable and PRBS_Rx_gen_enable control variables, the PMA shall generate a PRBS31 pattern on each of the lanes toward the PMA client via the PMA:IS_UNITDATA_*i*.indication primitive. While this test pattern is enabled, the PMA also generates PMA:IS_SIGNAL.indication(SIGNAL_OK) toward the PMA client independent of the link status at the service interface below the PMA. When send Rx PRBS31 test pattern is disabled, the PMA returns to normal operation performing bit multiplexing as described in 83.5.2. If a Clause 45 MDIO is implemented, the PRBS31_enable and PRBS_Rx_gen_enable control variables are accessible through bits 1.1501.7 and 1.1501.1 (see 45.2.1.124).

If supported, when check Tx PRBS31 test pattern mode is enabled by the PRBS31_enable and PRBS_Tx_check_enable control variables, the PMA shall check for the PRBS31 pattern on each of the lanes received from the PMA client via the PMA:IS_UNITDATA_*i*.request primitive. The checker shall

increment the test-pattern error counter by one for each incoming bit error in the PRBS31 pattern for isolated single bit errors. Implementations should be capable of counting at least one error whenever one or more errors occur in a sliding 1000-bit window. If a Clause 45 MDIO is implemented, the PRBS31_enable and PRBS_Tx_check_enable control variables are accessible through bits 1.1501.7 and 1.1501.2 (see 45.2.1.124). The Tx test-pattern error counters Ln0_PRBS_Tx_test_err_counter through Ln9_PRBS_Tx_test_err_counter count, per lane, errors in detecting the PRBS31 pattern on the lanes from the PMA client. If a Clause 45 MDIO is implemented, these counters are accessible through registers 1.1600 through 1.1609 (see 45.2.1.126). When check Tx PRBS31 test pattern is disabled, the PMA expects normal traffic and test-pattern error counting does not continue. While in check Tx PRBS31 test-pattern mode, bit multiplexing continues as described in 83.5.2. Note that bit multiplexing of per-lane PRBS31 may produce a signal which is not meaningful for downstream sublayers.

If supported, when check Rx PRBS31 test-pattern mode is enabled by the PRBS31_enable and PRBS_Rx_check_enable control variables, the PMA shall check for the PRBS31 pattern on each of the lanes received from the service interface below the PMA via the *inst*:IS_UNITDATA_*i*.indication primitive. If a Clause 45 MDIO is implemented, the PRBS31_enable and PRBS_Rx_check_enable control variables are accessible through bits 1.1501.7 and 1.1501.0 (see 45.2.1.124). The Rx test-pattern error counters Ln0_PRBS_Rx_test_err_counter through Ln9_PRBS_Rx_test_error_counter count, per lane, errors in detecting the PRBS31 pattern on the lanes from the service interface below the PMA. If a Clause 45 MDIO is implemented, these counters are accessible through registers 1.1700 through 1.1709 (see 45.2.1.127). While in check Rx PRBS31 test-pattern mode, the PMA:IS_SIGNAL.indication primitive does not indicate a valid signal. When check Rx PRBS31 test pattern is disabled, the PMA returns to normal operation performing bit multiplexing as described in 83.5.2.

If supported, when send Tx PRBS9 test-pattern mode is enabled by the PRBS9_enable and PRBS_Tx_gen_enable control variables, the PMA shall generate a PRBS9 pattern (as defined in Table 68–6) on each lane toward the service interface below the PMA via the *inst*:IS_UNITDATA_*i*.request primitive. If a Clause 45 MDIO is implemented, the PRBS9_enable and PRBS_Tx_gen_enable control variables are accessible through bits 1.1501.6 and 1.1501.3 (see 45.2.1.124). When send Tx PRBS9 test-pattern mode is disabled, the PMA returns to normal operation performing bit multiplexing as described in 83.5.2.

If supported, when send Rx PRBS9 test-pattern mode is enabled by the PRBS9_enable and PRBS_Rx_gen_enable control variables, the PMA shall generate a PRBS9 pattern on each lane toward the PMA client via the PMA:IS_UNITDATA_*i*.indication primitive. The PMA will also generate PMA:IS_SIGNAL.indication(SIGNAL_OK) toward the PMA client independent of the link status at the service interface below the PMA. If a Clause 45 MDIO is implemented, the PRBS9_enable and PRBS_Rx_gen_enable control variables are accessible through bits 1.1501.6 and 1.1501.1 (see 45.2.1.124). When send Rx PRBS9 test-pattern mode is disabled, the PMA returns to normal operation performing bit multiplexing as described in 83.5.2.

Note that PRBS9 is intended to be checked by external test gear, and no PRBS9 checking function is provided within the PMA.

Transmit square wave test-pattern mode optionally applies to each lane of the Tx direction PMA towards a physically instantiated XLAUI/CAUI-n or towards the PMD service interface whether or not it is physically instantiated. The ability to perform this function is indicated by the Square_wave_ability status variable. If a Clause 45 MDIO is implemented, the Square_wave_ability status variable is accessible through the Square wave test ability bit 1.1500.12 (see 45.2.1.123). If implemented, the transmit square wave test-pattern mode is enabled by control variables Square_wave_enable_0 through Square_wave_enable_9. If a Clause 45 MDIO is implemented, these control variables are accessible through the square wave testing control and status register bits 1.1510.0 through 1.1510.9 (limited to the number of lanes of the service interface below the PMA, see 45.2.1.125). When enabled, the PMA shall generate a square wave test pattern (8 ones followed by 8 zeros) on the square wave enabled lanes toward the service interface below the PMA via the *inst*:IS_UNITDATA_*i*.request primitive. Lanes for which square wave is not enabled will transmit normal

data resulting from the bit multiplexing operations described in 83.5.2 or test patterns as determined by other registers. When transmit square wave test pattern is disabled for all lanes, the PMA will perform normal operation performing bit multiplexing as described in 83.5.2 or transmit test patterns as determined by other registers.

83.5.11 Energy Efficient Ethernet

When the optional Energy Efficient Ethernet (EEE) deep sleep capability is supported and the PMA service interface is physically instantiated as XLAUI or CAUI-n, the additional functions listed in this subclause are required. These functions enable the communication of service interface parameters that are essential to the operation of the EEE deep sleep capability. The timing parameters for EEE operation are shown in Table 83–2.

Table 83–2—EEE timing parameters

Timer	Symbol	Min.	Max.	Units
PMA quiet signal duration	T_{pq}	200	225	ns
Energy detect hold-off time	T_{ho}	750	800	ns
Time to assert PMA quiet detect	T_{dq}	25	50	ns
Time to assert PMA alert detect	T_{da}	—	25	ns
Time to hold rx_lpi_active = true	T_{ht}	4000	5500	ns

83.5.11.1 PMA quiet and alert signals

The PMA quiet and alert signals are generated on each lane with a self-synchronizing scrambler. The scrambler shall produce the same result as the implementation shown in Figure 83–7. This implements the scrambler polynomial defined by Equation (83–1).

$$G(x) = 1 + x^{28} + x^{31} \tag{83–1}$$

To generate the PMA quiet signal the input to the scrambler shall be 0. To generate the PMA alert signal the input to the scrambler shall be 1.

The initial state the scrambler of a given lane of PMA service interface is chosen to minimize the correlation between lanes.

83.5.11.2 Detection of PMA quiet and alert signals

Each lane detects the PMA quiet and alert signals at the output of a self-synchronizing descrambler that implements the polynomial defined in Equation (83–1). The descrambler shall produce the same result as the implementation shown in Figure 83–8.

The output of the descrambler is considered in consecutive, non-overlapping blocks of 256 bits. If the number of zeros detected in a given 256-bit block is greater than or equal 224, then the lane shall indicate that the PMA quiet signal is detected. If the number of ones detected in a given block is greater than or equal to 224, then the lane shall indicate that the PMA alert signal has been detected. Otherwise, the lane infers

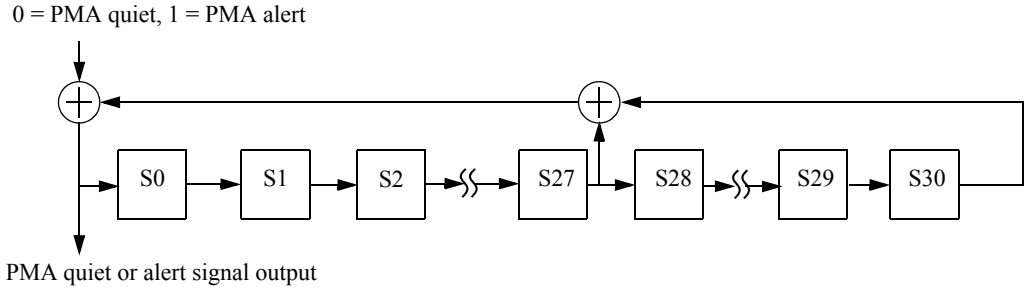

Figure 83–7—Scrambler for PMA quiet and alert signals

that normal data is being received and shall not indicate that either the PMA quiet or PMA alert signal has been detected.

The PMA shall indicate that the quiet signal is detected when all lanes of the PMA service interface have detected the quiet signal. The PMA shall indicate that the alert signal is detected when all lanes of the PMA service interface have detected the alert signal. Otherwise, the PMA infers that normal data is being received and shall not indicate that either the PMA quiet or PMA alert signal has been detected.

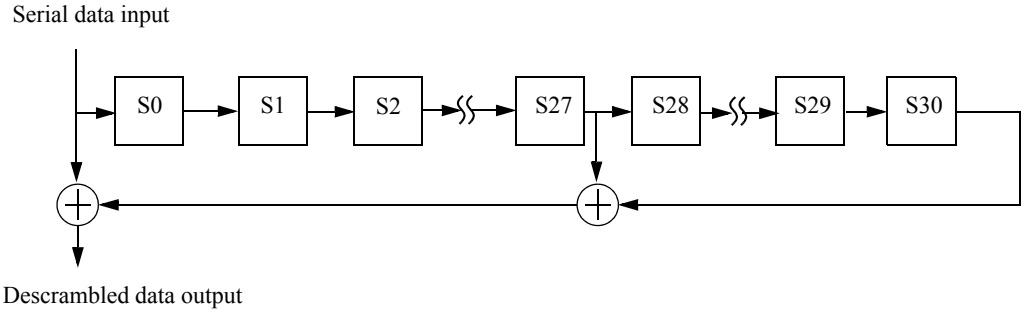

Figure 83–8—Descrambler for PMA quiet and alert signals

83.5.11.3 Additional transmit functions in the Tx direction

If the PMA client is the PCS, BASE-R FEC, or RS-FEC sublayer or is a PMA sublayer where the number of input lanes is not equal to the number of output lanes, then the PMA sublayer shall insert the PMA quiet and alert signals as follows: When the value of tx_mode is QUIET, the PMA inserts the PMA quiet signal defined in 83.5.11.1. When the value of tx_mode is ALERT, the PMA inserts the PMA alert signal defined in 83.5.11.1. For all other values of tx_mode, the PMA output is defined by the bit multiplexing function.

If XLAUI or CAUI-n is permitted to shut down (see 83.3), the variable aui_tx_mode shall be assigned the current value of tx_mode with the following exception. When tx_mode transitions from DATA to QUIET, the value of aui_tx_mode is held at DATA and the timer pma_quiet_timer (T_{pq}) is started. If tx_mode is QUIET when the timer expires, then aui_tx_mode is set to QUIET. If tx_mode is set to a value other than QUIET before the timer expires, then aui_tx_mode is set to DATA.

If XLAUI or CAUI-n is not permitted to shut down, aui_tx_mode shall be assigned the value DATA.

PMA functions in the Tx direction may be disabled in order to conserve energy while aui_tx_mode is QUIET.

83.5.11.4 Additional receive functions in the Tx direction

For a PMA that is separated from the PCS by XLAUI or CAUI-n, the value of tx_mode shall be assigned as follows: If the PMA quiet signal is detected, the value of tx_mode is set to QUIET. If the PMA alert signal is detected, the value of tx_mode is set to ALERT. Otherwise, the value of tx_mode is DATA.

If XLAUI or CAUI-n is permitted to shut down (see 83.3), then the variable aui_rx_mode shall be assigned as follows: The variable aui_rx_mode is initialized to DATA upon PMA power on or reset. When the PMA quiet signal is detected, the timer hold_off_timer (T_{ho}) is started. If the PMA alert signal is not detected before the timer expires, then aui_rx_mode is set to QUIET and SIGNAL_DETECT is set to FAIL. While aui_rx_mode is QUIET, it shall be set to DATA when SIGNAL_DETECT transitions from FAIL to OK. The value of tx_mode is inferred to be ALERT and the timer alert_timer (T_a) started upon a transition of aui_rx_mode from QUIET to DATA. The value of ALERT shall be held until alert_timer expires after which the value of tx_mode shall be set to DATA.

If XLAUI or CAUI-n is not permitted to shut down, aui_rx_mode shall be assigned the value DATA.

PMA functions in the Tx direction may be disabled in order to conserve energy while aui_rx_mode is QUIET.

83.5.11.5 Additional transmit functions in the Rx direction

For a PMA that is separated from the PCS by XLAUI or CAUI-n, the value of rx_mode shall be assigned as follows. The value of rx_mode is initialized to DATA upon PMA power on or reset. When the PMA quiet signal is detected, the timer hold_off_timer (T_{ho}) is started. If the PMA alert signal is not detected before the timer expires, then rx_mode is set to QUIET. While rx_mode is QUIET, it shall be set to DATA when the PMA alert signal is detected or energy_detect (or SIGNAL_OK) transitions from false to true.

The value of rx_tx_mode may be passed via the PMA:IS_RX_TX_MODE.indication primitive, otherwise it shall be assigned as follows. If the PMA quiet signal is detected, the value of rx_tx_mode is set to QUIET. The value of rx_tx_mode is set to be ALERT and the timer alert_timer (T_a) started upon a transition of the value of rx_mode from QUIET to DATA. The value of ALERT shall be held until alert_timer expires after which the value of rx_tx_mode shall be set to DATA.

If XLAUI or CAUI-n is permitted to shut down (see 83.3), the variable aui_tx_mode shall be assigned the current value of rx_tx_mode with the following exception. When rx_tx_mode transitions from DATA to QUIET, the value of aui_tx_mode is held at DATA and the timer pma_quiet_timer (T_{pq}) is started. If rx_tx_mode is QUIET when the timer expires, then aui_tx_mode is set to QUIET. If rx_tx_mode is set to a value other than QUIET before the timer expires, then aui_tx_mode is set to DATA.

If XLAUI or CAUI-n is not permitted to shut down, aui_tx_mode shall be assigned the value DATA.

PMA functions in the Rx direction may be disabled in order to conserve energy while aui_tx_mode is QUIET.

If the PMA is the client of the BASE-R FEC or RS-FEC sublayer or a PMA sublayer where the number of input lanes is not equal to the number of output lanes, then the PMA sublayer shall insert the PMA quiet and alert signals as follows. When the value of rx_tx_mode is QUIET, the PMA inserts the PMA quiet signal defined in 83.5.11.1. When the value of rx_tx_mode is ALERT, the PMA inserts the PMA alert signal defined in 83.5.11.1. For all other values of rx_tx_mode, the PMA output is defined by the bit multiplexing function.

83.5.11.6 Additional receive functions in the Rx direction

For a PMA that is separated from the PMD by XLAUI or CAUI, the value of energy_detect shall be assigned as follows. The value of energy_detect is initialized to true upon PMA power on or reset. When the value of rx_mode is set to QUIET, the value of energy_detect is set to false. The value of energy_detect is set to true when the PMA alert signal is detected or SIGNAL_DETECT transitions from FAIL to OK.

The value of rx_tx_mode shall be inferred as follows. If the PMA quiet signal is detected, the value of rx_tx_mode is set to QUIET. The value of rx_tx_mode is set to be ALERT and the timer alert_timer (T_a) started upon a transition of the value of rx_mode from QUIET to DATA. The value of ALERT shall be held until alert_timer expires after which the value of rx_tx_mode shall be set to DATA.

If XLAUI or CAUI is permitted to shut down (see 83.3), then the variable aui_rx_mode shall be assigned as follows. The variable aui_rx_mode is initialized to DATA upon PMA power on or reset. When the PMA quiet signal is detected, the timer hold_off_timer (T_{ho}) is started. If the PMA alert signal is not detected before the timer expires, then aui_rx_mode is set to QUIET. While aui_rx_mode is QUIET, it shall be set to DATA when SIGNAL_DETECT transitions from FAIL to OK. The value of tx_mode is assigned to be ALERT and the timer alert_timer (T_a) started upon a transition of aui_rx_mode from QUIET to DATA. The value of ALERT shall be held until alert_timer expires after which the value of tx_mode shall be set to DATA.

If XLAUI or CAUI is not permitted to shut down, aui_rx_mode shall be assigned the value DATA.

PMA functions in the Rx direction may be disabled in order to conserve energy while aui_rx_mode is QUIET.

83.5.11.7 Support for BASE-R FEC

When the PMA is a client of the BASE-R FEC sublayer, the rx_lpi_active parameter of the IS_RX_LPI_ACTIVE.request primitive shall be defined as follows. The value of rx_lpi_active is initialized to false upon PMA power on or reset. The value of rx_lpi_active is set to true and the timer rx_lpi_active_timer (T_{ht}) started upon a transition of the value of rx_mode from QUIET to DATA. When the timer expires, the value of rx_lpi_active is set to false.

83.6 PMA MDIO function mapping

The optional MDIO capability described in Clause 45 describes several variables that provide control and status information for and about the PMA. Since a given implementation may employ more than one PMA sublayer, the PMA control and status information is organized into multiple addressable instances, one for each possible PMA sublayer. See 45.2.1 and 83.1.4 for the allocation of MMD addresses to PMA sublayers. Control and status registers for MMD 8, 9, 10, and 11 will use the Extended PMA control and status registers at identical locations to those for MMD 1.

Mapping of MDIO control variables to PMA control variables is shown in Table 83–3. Mapping of MDIO status variables to PMA status variables is shown in Table 83–4. Mapping of MDIO counter to PMA counters is shown in Table 83–5. These tables provide the register and bit numbers for the PMA addressed as MMD 1. For implementations with multiple PMA sublayers, additional PMA sublayers use the corresponding register and bit numbers in MMDs 8, 9, 10, and 11 as necessary.

Table 83–3—MDIO/PMA control variable mapping

MDIO variable	PMA/PMD register name	Register/ bit number	PMA control variable
PMA remote loopback	PMA/PMD control 1	1.0.1	Remote_loopback_enable
PMA local loopback	PMA/PMD control 1	1.0.0	Local_loopback_enable
PRBS31 pattern enable	PRBS pattern testing control	1.1501.7	PRBS31_enable
PRBS9 pattern enable	PRBS pattern testing control	1.1501.6	PRBS9_enable
Tx generator enable	PRBS pattern testing control	1.1501.3	PRBS_Tx_gen_enable
Tx checker enable	PRBS pattern testing control	1.1501.2	PRBS_Tx_check_enable
Rx generator enable	PRBS pattern testing control	1.1501.1	PRBS_Rx_gen_enable
Rx checker enable	PRBS pattern testing control	1.1501.0	PRBS_Rx_check_enable
Lane 0 SW enable	Square wave testing control	1.1510.0	Square_wave_enable_0
Lane 1 SW enable	Square wave testing control	1.1510.1	Square_wave_enable_1
Lane 2 SW enable	Square wave testing control	1.1510.2	Square_wave_enable_2
Lane 3 SW enable	Square wave testing control	1.1510.3	Square_wave_enable_3
Lane 4 SW enable	Square wave testing control	1.1510.4	Square_wave_enable_4
Lane 5 SW enable	Square wave testing control	1.1510.5	Square_wave_enable_5
Lane 6 SW enable	Square wave testing control	1.1510.6	Square_wave_enable_6
Lane 7 SW enable	Square wave testing control	1.1510.7	Square_wave_enable_7
Lane 8 SW enable	Square wave testing control	1.1510.8	Square_wave_enable_8
Lane 9 SW enable	Square wave testing control	1.1510.9	Square_wave_enable_9
PIASE	PMA ingress AUI stop enable	1.7.9	PIASE
PEASE	PMA egress AUI stop enable	1.7.8	PEASE

Table 83–4—MDIO/PMA status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMA status variable
PMA remote loopback ability	40G/100G PMA/PMD extended ability register	1.13.15	Remote_loopback_ability
PMA local loopback ability	PMA/PMD status 2 register	1.8.0	Local_loopback_ability
PRBS9 Tx generator ability	Test-pattern ability register	1.1500.5	PRBS9_Tx_generator_ability
PRBS9 Rx generator ability	Test-pattern ability register	1.1500.4	PRBS9_Rx_generator_ability
PRBS31Tx generator ability	Test-pattern ability register	1.1500.3	PRBS31_Tx_generator_ability
PRBS31Tx checker ability	Test-pattern ability register	1.1500.2	PRBS31_Tx_checker_ability
PRBS31 Rx generator ability	Test-pattern ability register	1.1500.1	PRBS31_Rx_generator_ability
PRBS31 Rx checker ability	Test-pattern ability register	1.1500.0	PRBS31_Rx_checker_ability
Square wave test ability	Test-pattern ability register	1.1500.12	Square_wave_ability
PIASA	PMA ingress AUI stop ability	1.1.9	PIASA
PEASA	PMA egress AUI stop ability	1.1.8	PEASA

Table 83–5—MDIO/PMA counters mapping

MDIO variable	PMA/PMD register name	Register/bit number	PMA status variable
Error counter Tx, lane 0	PRBS Tx pattern testing error counter, lane 0	1.1600	Ln0_PRBS_Tx_test_err_counter
Error counter Tx, lane 1	PRBS Tx pattern testing error counter, lane 1	1.1601	Ln1_PRBS_Tx_test_err_counter
Error counter Tx, lane 2	PRBS Tx pattern testing error counter, lane 2	1.1602	Ln2_PRBS_Tx_test_err_counter
Error counter Tx, lane 3	PRBS Tx pattern testing error counter, lane 3	1.1603	Ln3_PRBS_Tx_test_err_counter
Error counter Tx, lane 4	PRBS Tx pattern testing error counter, lane 4	1.1604	Ln4_PRBS_Tx_test_err_counter
Error counter Tx, lane 5	PRBS Tx pattern testing error counter, lane 5	1.1605	Ln5_PRBS_Tx_test_err_counter
Error counter Tx, lane 6	PRBS Tx pattern testing error counter, lane 6	1.1606	Ln6_PRBS_Tx_test_err_counter
Error counter Tx, lane 7	PRBS Tx pattern testing error counter, lane 7	1.1607	Ln7_PRBS_Tx_test_err_counter

Table 83–5—MDIO/PMA counters mapping *(continued)*

MDIO variable	PMA/PMD register name	Register/bit number	PMA status variable
Error counter Tx, lane 8	PRBS Tx pattern testing error counter, lane 8	1.1608	Ln8_PRBS_Tx_test_err_counter
Error counter Tx, lane 9	PRBS Tx pattern testing error counter, lane 9	1.1609	Ln9_PRBS_Tx_test_err_counter
Error counter Rx, lane 0	PRBS Rx pattern testing error counter, lane 0	1.1700	Ln0_PRBS_Rx_test_err_counter
Error counter Rx, lane 1	PRBS Rx pattern testing error counter, lane 1	1.1701	Ln1_PRBS_Rx_test_err_counter
Error counter Rx, lane 2	PRBS Rx pattern testing error counter, lane 2	1.1702	Ln2_PRBS_Rx_test_err_counter
Error counter Rx, lane 3	PRBS Rx pattern testing error counter, lane 3	1.1703	Ln3_PRBS_Rx_test_err_counter
Error counter Rx, lane 4	PRBS Rx pattern testing error counter, lane 4	1.1704	Ln4_PRBS_Rx_test_err_counter
Error counter Rx, lane 5	PRBS Rx pattern testing error counter, lane 5	1.1705	Ln5_PRBS_Rx_test_err_counter
Error counter Rx, lane 6	PRBS Rx pattern testing error counter, lane 6	1.1706	Ln6_PRBS_Rx_test_err_counter
Error counter Rx, lane 7	PRBS Rx pattern testing error counter, lane 7	1.1707	Ln7_PRBS_Rx_test_err_counter
Error counter Rx, lane 8	PRBS Rx pattern testing error counter, lane 8	1.1708	Ln8_PRBS_Rx_test_err_counter
Error counter Rx, lane 9	PRBS Rx pattern testing error counter, lane 9	1.1709	Ln9_PRBS_Rx_test_err_counter

83.7 Protocol implementation conformance statement (PICS) proforma for Clause 83, Physical Medium Attachment (PMA) sublayer, type 40GBASE-R and 100GBASE-R[8]

83.7.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 83, Physical Medium Attachment (PMA) sublayer, type 40GBASE-R and 100GBASE-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 same, can be found in Clause 21.

83.7.2 Identification

83.7.2.1 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).	

83.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 83, Physical Medium Attachment (PMA) sublayer, type 40GBASE-R and 100GBASE-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	

[8]*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.

83.7.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*PMA40	PMA for 40GBASE-R	83.1.1		O.1	Yes [ ] No [ ]
*PMA100	PMA for 100GBASE-R	83.1.1		O.1	Yes [ ] No [ ]
LANES_ UPSTREAM	Number of lanes in direction of PCS	83.1.4	Divisor of number of PCS lanes	PMA40:M	4 [ ]
				PMA100: M	4 [ ] 10 [ ] 20 [ ]
LANES_ DOWNSTREAM	Number of lanes in direction of PMD	83.1.4	Divisor of number of PCS lanes	PMA40:M	1 [ ] 4 [ ]
				PMA100: M	4 [ ] 10 [ ] 20 [ ]
RX_CLOCK	Signaling rate of output lanes in Rx direction	83.5.5	LANES_DOWNSTREAM / LANES_UPSTREAM times signaling rate of input lanes in Rx direction	M	Yes [ ] No [ ]
TX_CLOCK	Signaling rate of output lanes in Tx direction	83.5.5	LANES_UPSTREAM / LANES _DOWNSTREAM times signaling rate of input lanes in Tx direction	M	Yes [ ] No [ ]
LANE_ MAPPING	Maintain lane mapping while link is in operation	83.5.2	Maintain sequence of PCSLs on all output lanes	M	Yes [ ] No [ ]
LNKS	PMA link status	83.5.7	Meets the requirements of 83.5.7	M	Yes [ ] No [ ]
JTP	Supports test-pattern mode	83.5.10		O	Yes [ ] No [ ] N/A [ ]
*KRCR	PMA adjacent to the PMD for 40GBASE-KR4, 40GBASE-CR4, 100GBASE-KR4, 100GBASE-CR4, or 100GBASE-CR10	83.5.8		O	Yes [ ] No [ ]
*LBL	PMA local loopback	83.5.8	Supports local loopback	KRCR:M !KRCR:O	Yes [ ] No [ ] N/A[ ]
*LBR	PMA remote loopback	83.5.9	Supports remote loopback	O	Yes [ ] No [ ] N/A[ ]

Item	Feature	Subclause	Value/Comment	Status	Support
MD	MDIO	83.6	Registers and interface supported	O	Yes [ ] No [ ]
*USP1SP6	Physically instantiated XLAUI or CAUI-n above (toward PCS)	83.5.3		O	Yes [ ] No [ ]
*DSP1SP6	Physically instantiated XLAUI or CAUI-n below (toward PMD)	83.5.3		O	Yes [ ] No [ ]
*SP2SP5	Physically instantiated PMD service interface	83.5.3		O	Yes [ ] No [ ]
*PPI	PMD service interface instantiated as nPPI	83.5.1, 83.5.5		O	Yes [ ] No [ ]
UNAUI	Electrical and timing requirements of Annex 83A or Annex 83B or Annex 83D or Annex 83E as appropriate met by upstream XLAUI/CAUI-n	83.5.1, 83.5.5		USP1SP6: M	Yes [ ] No [ ]
DNAUI	Electrical and timing requirements of Annex 83A or Annex 83B or Annex 83D or Annex 83E as appropriate met by downstream XLAUI/CAUI-n	83.5.1, 83.5.5		DSP1SP6: M	Yes [ ] No [ ]
PPIET	Electrical and timing requirements of Annex 86A met by PMD service interface	83.5.1, 83.5.5		PPI:M	Yes [ ] No [ ]
DELAY40	Roundtrip delay limit for 40GBASE-R	83.5.4	No more than 4096 BT or 8 pause_quanta	PMA40:M	Yes [ ] No [ ]
DELAY100	Roundtrip delay limit for 100GBASE-R	83.5.4	No more than 9216 BT or 18 pause_quanta	PMA100: M	Yes [ ] No [ ]
*LPI	Implementation of LPI with the deep sleep mode option	83.3		O	Yes [ ] No [ ]

83.7.4 Skew generation and tolerance

Item	Feature	Subclause	Value/Comment	Status	Support
S1	Skew generation toward SP1 in Tx direction	83.5.3.2	≤ 29 ns	DSP1SP6:M	Yes [ ] No [ ]
S2	Skew variation generation toward SP1 in Tx direction	83.5.3.2	≤ 200 ps	DSP1SP6:M	Yes [ ] No [ ]
S3	Skew variation tolerance at SP1	83.5.3.3	Minimum 200 ps	USP1SP6:M	Yes [ ] No [ ]
S4	Skew generation toward SP2 in Tx direction	83.5.3.4	≤ 43 ns	SP2SP5:M	Yes [ ] No [ ]
S5	Skew variation generation toward SP2 in Tx direction	83.5.3.4	≤ 400 ps	SP2SP5:M	Yes [ ] No [ ]
S6	Skew variation tolerance at SP5	83.5.3.5	Minimum 3.6 ns	SP2SP5:M	Yes [ ] No [ ]
S7	Skew generation toward SP6 in Rx direction	83.5.3.6	≤ 160 ns	USP1SP6:M	Yes [ ] No [ ]
S8	Skew variation generation toward SP6 in Rx direction	83.5.3.6	≤ 3.8 ns	USP1SP6:M	Yes [ ] No [ ]
S9	Skew variation tolerance at SP6	83.5.3.7	Minimum 3.8 ns	DSP1SP6:M	Yes [ ] No [ ]

83.7.5 Test patterns

Item	Feature	Subclause	Value/Comment	Status	Support
*JTP1	Physically Instantiated XLAUI/CAUI-n between PMA and PMA client	83.5.10		O	Yes [ ] No [ ]
*JTP2	Physically Instantiated XLAUI/CAUI-n between PMA and sublayer below the PMA, or adjacent to PMD whether or not PMD service interface is physically instantiated as nPPI	83.5.10		O	Yes [ ] No [ ]
J1	Send PRBS31 Tx	83.5.10		JTP2:O	Yes [ ] No [ ]
J2	Send PRBS31 Rx	83.5.10		JTP1:O	Yes [ ] No [ ]
J3	Check PRBS31 Tx	83.5.10		JTP1:O	Yes [ ] No [ ]
J4	Check PRBS31 Rx	83.5.10		JTP2:O	Yes [ ] No [ ]
J5	Send PRBS9 Tx	83.5.10		JTP2:O	Yes [ ] No [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
J6	Send PRBS9 Rx	83.5.10		JTP1:O	Yes [ ] No [ ]
J7	Send square wave Tx	83.5.10		JTP2:O	Yes [ ] No [ ]

83.7.6 Loopback modes

Item	Feature	Subclause	Value/Comment	Status	Support
LB1	PMA local loopback implemented	83.5.8	Meets the requirements of 83.5.8	LBL:M	Yes [ ] No [ ]
LB2	PMA remote loopback implemented	83.5.9	Meets the requirements of 83.5.9	LBR:M	Yes [ ] No [ ]

83.7.7 EEE deep sleep with XLAUI/CAUI

Item	Feature	Subclause	Value/Comment	Status	Support
RXDS	XLAUI/CAUI-n deep sleep Rx direction	83.5.11		LPI*USP1SP6:M or LPI*DSP1SP6:M	Yes [ ] No [ ]
TXDS	XLAUI/CAUI-n deep sleep Tx direction	83.5.11		LPI*USP1SP6:M or LPI*DSP1SP6:M	Yes [ ] No [ ]

84. Physical Medium Dependent sublayer and baseband medium, type 40GBASE-KR4

84.1 Overview

This clause specifies the 40GBASE-KR4 PMD. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 84–1, to the medium through the MDI and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 84–1—Physical Layer clauses associated with the 40GBASE-KR4 PMD

Associated clause	40GBASE-KR4
81—RS	Required
81—XLGMII[a]	Optional
82—PCS for 40GBASE-R	Required
74—BASE-R FEC	Optional
83—PMA for 40GBASE-R	Required
83A—XLAUI	Optional
73—Auto-Negotiation	Required
78—Energy Efficient Ethernet	Optional

[a]The XLGMII is an optional interface. However, if the XLGMII is not implemented, a conforming implementation must behave functionally as though the RS and XLGMII were present.

Figure 84–1 shows the relationship of the 40GBASE-KR4 PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model.

40GBASE-KR4 PHYs with the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option may optionally enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization (see Clause 78).

84.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 40GBASE-KR4 PMD. The service interface for this PMD is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data. The PMD translates the encoded data to and from signals suitable for the medium.

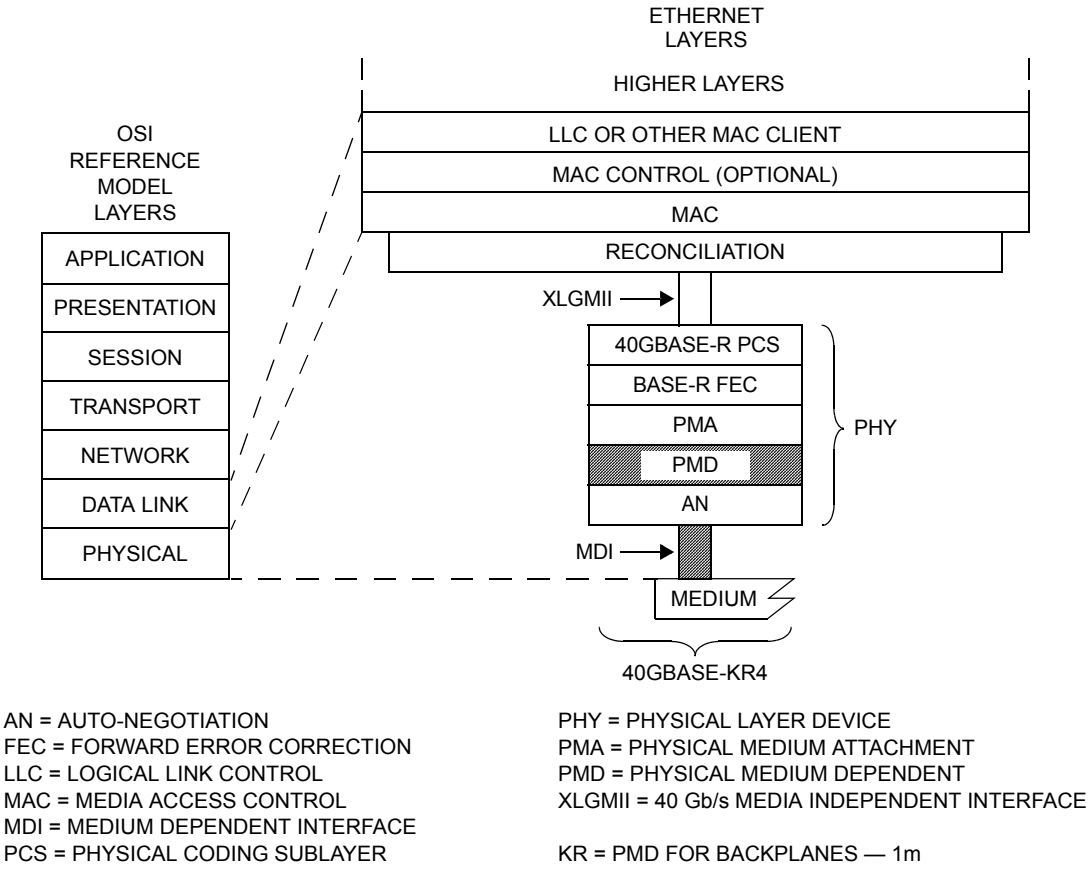

Figure 84–1—40GBASE-KR4 PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_*i*.request
PMD:IS_UNITDATA_*i*.indication
PMD:IS_SIGNAL.indication

The 40GBASE-KR4 PMD has four parallel bit streams, hence $i = 0$ to 3.

The PMA (or the PMD) continuously sends four parallel bit streams to the PMD (or the PMA), one per lane, each at a nominal signaling rate of 10.3125 GBd.

SIGNAL_DETECT in 40GBASE-KR4 indicates the successful completion of the start-up protocol on all four lanes.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the PMD:IS_UNITDATA_*i*.indication parameters are undefined.

The SIGNAL_DETECT parameter maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication primitive.

If the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option is supported (see Clause 78, 78.1.3.3.1), then the inter-sublayer service interface includes two additional primitives defined as follows:

PMD:IS_TX_MODE.request(tx_mode)
PMD:IS_RX_MODE.request(rx_mode)

The tx_mode parameter takes on one of up to three values: DATA, QUIET, or ALERT. When tx_mode = QUIET, transmission is disabled; when tx_mode = ALERT, the alert signal is transmitted (see 84.7.2).

The rx_mode parameter is used to communicate the state of the PCS LPI receive function and takes the value QUIET or DATA.

84.3 PCS requirements for Auto-Negotiation (AN) service interface

The PCS associated with this PMD is required to support the AN service interface primitive AN_LINK.indication defined in 73.9. (See 82.6.)

The 40GBASE-KR4 PHY may be extended using XLAUI as a physical instantiation of the inter-sublayer service interface between devices. If XLAUI is instantiated, the AN_LINK(link_status).indication is relayed from the device with the PCS sublayer to the device with the AN sublayer by means at the discretion of the implementer. As examples, the implementer may employ use of pervasive management or employ a dedicated electrical signal to relay the state of link_status as indicated by the PCS sublayer on one device to the AN sublayer on the other device.

84.4 Delay constraints

The sum of the transmit and the receive delays at one end of the link contributed by the 40GBASE-KR4 PMD, AN, and the medium in one direction shall be no more than 2048 bit times (4 pause_quanta or 51.2 ns). It is assumed that the one way delay through the medium is no more than 320 bit times (8 ns).

A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

84.5 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the PCS. The Skew Variation must also be limited to ensure that a given PCS lane always traverses the same physical lane. Skew and Skew Variation are defined in 80.5 and specified at the points SP1 to SP6 shown in Figure 80–6 and Figure 80–7.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 80.5.

84.6 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 84–2, and MDIO status variables to PMD status variables as shown in Table 84–3.

Table 84–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/ bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable register	1.9.0	Global_PMD_transmit_disable
PMD transmit disable 3	PMD transmit disable register	1.9.4	PMD_transmit_disable_3
PMD transmit disable 2	PMD transmit disable register	1.9.3	PMD_transmit_disable_2
PMD transmit disable 1	PMD transmit disable register	1.9.2	PMD_transmit_disable_1
PMD transmit disable 0	PMD transmit disable register	1.9.1	PMD_transmit_disable_0
Restart training	BASE-R PMD control register	1.150.0	mr_restart_training
Training enable	BASE-R PMD control register	1.150.1	mr_training_enable

Table 84–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	Global_PMD_signal_detect
PMD receive signal detect 3	PMD receive signal detect register	1.10.4	PMD_signal_detect_3
PMD receive signal detect 2	PMD receive signal detect register	1.10.3	PMD_signal_detect_2
PMD receive signal detect 1	PMD receive signal detect register	1.10.2	PMD_signal_detect_1
PMD receive signal detect 0	PMD receive signal detect register	1.10.1	PMD_signal_detect_0
40GBASE-KR4 deep sleep	EEE capability	1.16.0	—
Receiver status 3	BASE-R PMD status register	1.151.12	rx_trained_3
Frame lock 3	BASE-R PMD status register	1.151.13	frame_lock_3

Table 84–3—MDIO/PMD status variable mapping *(continued)*

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Start-up protocol status 3	BASE-R PMD status register	1.151.14	training_3
Training failure 3	BASE-R PMD status register	1.151.15	training_failure_3
Receiver status 2	BASE-R PMD status register	1.151.8	rx_trained_2
Frame lock 2	BASE-R PMD status register	1.151.9	frame_lock_2
Start-up protocol status 2	BASE-R PMD status register	1.151.10	training_2
Training failure 2	BASE-R PMD status register	1.151.11	training_failure_2
Receiver status 1	BASE-R PMD status register	1.151.4	rx_trained_1
Frame lock 1	BASE-R PMD status register	1.151.5	frame_lock_1
Start-up protocol status 1	BASE-R PMD status register	1.151.6	training_1
Training failure 1	BASE-R PMD status register	1.151.7	training_failure_1
Receiver status 0	BASE-R PMD status register	1.151.0	rx_trained_0
Frame lock 0	BASE-R PMD status register	1.151.1	frame_lock_0
Start-up protocol status 0	BASE-R PMD status register	1.151.2	training_0
Training failure 0	BASE-R PMD status register	1.151.3	training_failure_0

84.7 PMD functional specifications

84.7.1 Link block diagram

The 40GBASE-KR4 PMD uses the same Link block diagram as 10GBASE-KX4, as defined in 71.6.1.

84.7.2 PMD transmit function

The PMD Transmit function shall convert the four bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_3.request into four separate electrical streams. A positive output voltage of SL<p> minus SL<n> (differential voltage) shall correspond to tx_bit = one.

If the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option is supported (see Clause 78), then when tx_mode is set to ALERT, the PMD transmit function shall transmit a repeating 16-bit pattern, hexadecimal 0xFF00, on each lane. This sequence is transmitted regardless of the value of tx_bit presented by the PMD:IS_UNITDATA_i.request primitive. When tx_mode is ALERT, the transmitter equalizer taps shall be set to the preset state specified in 72.6.10.2.3.1. When tx_mode is QUIET, the transmitter shall be disabled as specified in 84.7.6. For all other states of tx_mode, the driver coefficients are restored to their states resolved during training.

84.7.3 PMD receive function

The PMD Receive function shall convert the four electrical streams from the MDI into four bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to

PMD:IS_UNITDATA_3.indication. A positive input voltage of DL<p> minus DL<n> (differential voltage) shall correspond to rx_bit = one.

84.7.4 Global PMD signal detect function

The Global PMD signal detect function shall continuously report the message PMD:IS_SIGNAL.indication (SIGNAL_DETECT) to the PMD service interface. SIGNAL_DETECT, while normally intended to be an indicator of signal presence, is used by 40GBASE-KR4 to indicate the successful completion of the start-up protocol on all lanes. When the PHY supports the optional EEE capability with the deep sleep mode, PMD_SIGNAL.indication is also used to indicate when the ALERT signal is detected, which corresponds to the beginning of a refresh or a wake.

When the PHY does not support the EEE capability or if the PHY supports the EEE capability and rx_mode is set to DATA, SIGNAL_DETECT shall be set to FAIL following system reset or the manual reset of the training state diagram. Upon successful completion of training on all lanes, SIGNAL_DETECT shall be set to OK.

When the PHY does not support the EEE capability or if the PHY supports the EEE capability and rx_mode is set to DATA, if training is disabled by management, SIGNAL_DETECT shall be set to OK. When the PHY supports the EEE capability with the deep sleep mode, SIGNAL_DETECT is set to FAIL following a transition from rx_mode = DATA to rx_mode = QUIET. When rx_mode = QUIET, SIGNAL_DETECT shall be set to OK within 500 ns following the application of a signal at the receiver input that corresponds to an ALERT transmission (see 84.7.2) from the link partner. While rx_mode = QUIET, SIGNAL_DETECT shall be held at FAIL as long as the signal at the receiver input corresponds to a QUIET tx_mode (see 84.7.6) of the link partner.

If the MDIO interface is implemented, then Global_PMD_signal_detect (1.10.0) shall be continuously set to the value of SIGNAL_DETECT as described in 45.2.1.9.7.

84.7.5 PMD lane-by-lane signal detect function

If the MDIO interface is implemented, then PMD_signal_detect_0 (1.10.1), PMD_signal_detect_1 (1.10.2), PMD_signal_detect_2 (1.10.3), and PMD_signal_detect_3 (1.10.4) shall be set to one or zero depending on whether a particular lane's signal_detect, as defined by the training state diagram in Figure 72–5, returns true or false (see 45.2.1.9).

84.7.6 Global PMD transmit disable function

The Global_PMD_transmit_disable function is mandatory if EEE with the deep sleep mode option is supported and is otherwise optional. When implemented, it allows all of the transmitters to be disabled with a single variable.

a) When the 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 with TX disabled as specified in Table 72–7.
b) If a PMD_fault (84.7.9) is detected, then the PMD may turn off the electrical transmitter in all lanes.
c) Loopback, as defined in 84.7.8, shall not be affected by Global_PMD_transmit_disable.
d) For EEE capability, the PMD_transmit_disable function shall turn off the transmitter after tx_mode is set to QUIET within a time and voltage level specified in 72.7.1.4. The PMD_transmit_disable function shall turn on the transmitter after tx_mode is set to DATA or ALERT within the time and voltage level specified in 72.7.1.4.

If the MDIO interface is implemented, then this function shall map to the Global_PMD_transmit_disable bit as specified in 45.2.1.8.7.

84.7.7 PMD lane-by-lane transmit disable function

The PMD_transmit_disable_*i* function (where *i* represents the lane number in the range 0:3) is optional and allows the electrical transmitter in each lane to be selectively disabled. When this function is supported, it shall meet the following requirements:

a) When a PMD_transmit_disable_*i* variable is set to one, this function shall turn off the transmitter associated with that variable such that it drives a constant level (i.e., no transitions) and does not exceed the maximum differential peak-to-peak output voltage with TX disabled as specified in Table 72–7.
b) If a PMD_fault (84.7.9) is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the electrical transmitter in each lane.
c) Loopback, as defined in 84.7.8, shall not be affected by PMD_transmit_disable_*i*.

84.7.8 Loopback mode

Local loopback shall be provided by the adjacent PMA (see 83.5.8) for the 40GBASE-KR4 PMD as a test function to the device. When loopback mode is enabled, transmission requests passed to each transmitter are sent directly to the corresponding receiver, overriding any signal detected by each receiver on its attached link. Note that loopback mode does not affect the state of the transmitter that continues to send data (unless disabled). 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 paths that are exercised in the loopback mode are implementation specific, but it is recommended that these signal paths 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.

84.7.9 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.

84.7.10 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 Global_PMD_transmit_disable function.

If a PMD_transmit_fault (optional) is detected, then the Global_PMD_transmit_disable function should also be asserted.

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.

84.7.11 PMD receive fault function

The PMD_receive_fault function is optional. The faults detected by this function are implementation specific.

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.

84.7.12 PMD control function

Each lane of the 40GBASE-KR4 PMD shall use the same control function as 10GBASE-KR, as defined in 72.6.10.

The random seed for the training pattern described in 72.6.10.2.6 shall be different for each of the lanes.

The variables rx_trained_*i*, frame_lock_*i*, training_*i,* and training_failure_*i* (where *i* goes from 0 to 3) report status for each lane and are equivalent to rx_trained, frame_lock, training, and training_failure as defined in 72.6.10.3.1.

If the MDIO interface is implemented, then this function shall map these variables to the appropriate bits in the BASE-R PMD status register (Register 1.151) as specified in 45.2.1.81.

84.8 40GBASE-KR4 electrical characteristics

84.8.1 Transmitter characteristics

Transmitter electrical characteristics at TP1 for 40GBASE-KR4 shall be the same as 10GBASE-KR, as detailed in 72.7.1.1 through 72.7.1.11.

84.8.1.1 Test fixture

The test fixture defined for 10GBASE-KR in 72.7.1.1 shall be used on all lanes.

84.8.2 Receiver characteristics

Receiver electrical characteristics at TP4 for 40GBASE-KR4 shall be the same as 10GBASE-KR, as detailed in 72.7.2.1 through 72.7.2.5.

84.8.2.1 Receiver interference tolerance

The receiver interference tolerance tests shall be the same as those described for 10GBASE-KR in 72.7.2.1 and Annex 69A.

For 40GBASE-KR4, each lane shall be tested individually using the methodology defined in Annex 69A with the transmitters of the unused lanes active and terminated by the reference impedance.

84.9 Interconnect characteristics

Informative interconnect characteristics for 40GBASE-KR4 are provided in Annex 69B.

84.10 Environmental specifications

84.10.1 General safety

All equipment subject to this clause shall conform to applicable sections (including isolation requirements) of IEC 60950-1.

84.10.2 Network safety

The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

84.10.3 Installation and maintenance guidelines

It is recommended that sound installation practice, as defined by applicable local codes and regulations, be followed in every instance in which such practice is applicable.

84.10.4 Electromagnetic compatibility

A system integrating the 40GBASE-KR4 PHY shall comply with applicable local and national codes for the limitation of electromagnetic interference.

84.10.5 Temperature and humidity

A system integrating the 40GBASE-KR4 PHY 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.

84.11 Protocol implementation conformance statement (PICS) proforma for Clause 84, Physical Medium Dependent sublayer and baseband medium, type 40GBASE-KR4[9]

84.11.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 84, Physical Medium Dependent sublayer and baseband medium, type 40GBASE-KR4, 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.

84.11.2 Identification

84.11.2.1 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).	

84.11.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 84, Physical Medium Dependent sublayer and baseband medium, type 40GBASE-KR4
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	

[9] *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.

84.11.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
XGE	XLGMII	84.1	Interface is supported	O	Yes [ ] No [ ]
PCS	Support of 40GBASE-R PCS	84.1		M	Yes [ ]
PMA	Support of 40GBASE-R PMA	84.1		M	Yes [ ]
XLAUI	XLAUI	84.1	Interface is supported	O	Yes [ ] No [ ]
AN	Auto-Negotiation	84.1	Device implements Auto-Negotiation	M	Yes [ ]
FEC	Forward Error Correction	84.1	Device implements BASE-R Forward Error Correction	O	Yes [ ]
DC	Delay constraints	84.4	Device conforms to delay constraints	M	Yes [ ]
DSC	Skew constraints	84.5	Device conforms to Skew and Skew Variation constraints	M	Yes [ ]
*MD	MDIO interface	84.6	Device implements MDIO	O	Yes [ ] No [ ]
*TD	Global_PMD_transmit_disable	84.7.6		O	Yes [ ] No [ ]
*LPI	Implementation of LPI with the deep sleep mode option	84.1		O	Yes [ ] No [ ]

84.11.4 PICS proforma tables for Clause 84, Physical Medium Dependent (PMD) sublayer and baseband medium, type 40GBASE-KR4

84.11.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
FS1	Transmit function	84.7.2	Conveys bits from PMD service interface to MDI	M	Yes [ ]
FS2	Transmitter signal	84.7.2	A positive differential voltage corresponds to tx_bit = one	M	Yes [ ]
FS3	Receive function	84.7.3	Conveys bits from MDI to PMD service interface	M	Yes [ ]
FS4	Receiver signal	84.7.3	A positive differential voltage corresponds to rx_bit = one	M	Yes [ ]
FS5	Signal detect	84.7.4	Report to PMD service interface	M	Yes [ ]
FS6	SIGNAL_DETECT value	84.7.4	Set to FAIL on reset	M	Yes [ ]
FS7	SIGNAL_DETECT value	84.7.4	Set to OK when training completes successfully	M	Yes [ ]
FS8	SIGNAL_DETECT value	84.7.4	Set to OK when training disabled	M	Yes [ ]
FS9	Transmit disable requirements	84.7.6	Requirements of 84.7.6 and 84.7.7	TD:M	Yes [ ] N/A[ ]
FS10	Loopback support	84.7.8	Provided for 40GBASE-KR4 PMD by transmitter and receiver	M	Yes [ ]
FS11	Control function	84.7.12	The same control function as 10GBASE-KR is used	M	Yes [ ]
FS12	Random seed	84.7.12	The random seed shall be different for each lane	M	Yes [ ]
FS13	Transmit function for EEE	84.7.2	Transmitter behavior during ALERT and QUIET	LPI:M	Yes [ ] No [ ]
FS14	Signal detect function for EEE	84.7.4		LPI:M	Yes [ ] No [ ]
FS15	Transmit disable during LPI	84.7.6	Disable transmitter during tx_mode = QUIET	LPI:M	Yes [ ] No [ ]

84.11.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
MF1	MDIO Variable Mapping	84.6	Per Table 84–2, and Table 84–3	MD:M	Yes [ ] N/A [ ]
MF2	Global signal detect	84.7.4	Value described in 45.2.1.9.7	MD:M	Yes [ ] N/A [ ]
MF3	Lane-by-lane signal detect	84.7.5	Value as described in 84.7.5	MD:M	Yes [ ] N/A [ ]
MF4	PMD_transmit_fault function	84.7.10	Sets PMD_transmit_fault as specified in 45.2.1.7.4	MD:M	Yes [ ] N/A [ ]
MF5	PMD_receive_fault function	84.7.11	Sets PMD_receive_fault as specified in 45.2.1.7.5	MD:M	Yes [ ] N/A [ ]
MF6	PMD training status	84.7.12	Sets training status as specified in 45.2.1.81	MD:M	Yes [ ] N/A [ ]

84.11.4.3 Transmitter electrical characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	Transmitter characteristics	84.8.1	The same transmitter characteristics as 10GBASE-KR are used	M	Yes [ ]
TC2	Test fixture	84.8.1.1	The same test fixture as 10GBASE-KR is used	M	Yes [ ]
TC3	Output Amplitude LPI voltage	84.7.6	Less than 30 mV within 500 ns of tx_quiet	LPI:M	Yes [ ] No [ ]
TC4	Output Amplitude ON voltage	84.7.6	Greater than 90% of previous level within 500 ns of tx_quiet deasserted	LPI:M	Yes [ ] No [ ]

84.11.4.4 Receiver electrical characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
RC1	Receiver characteristics	84.8.2	The same receiver characteristics as 10GBASE-KR are used	M	Yes [ ]
RC2	Receiver interference tolerance	84.8.2.1	The same receiver interference tolerance test as 10GBASE-KR is used	M	Yes [ ]
RC3	Receiver interference tolerance testing	84.8.2.1	Each lane shall be tested individually using the methodology defined in Annex 69A with the transmitters of the unused lanes active and terminated by the reference impedance	M	Yes [ ]

84.11.4.5 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
ES1	General safety	84.10.1	Complies with applicable section of IEC 60950-1	M	Yes [ ]
ES2	Electromagnetic interference	84.10.4	Complies with applicable local and national codes	M	Yes [ ]

85. Physical Medium Dependent sublayer and baseband medium, type 40GBASE-CR4 and 100GBASE-CR10

85.1 Overview

This clause specifies the 40GBASE-CR4 PMD and the 100GBASE-CR10 PMD (including MDI) and the baseband medium. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 85–1, to the medium through the MDI and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 85–1—Physical Layer clauses associated with the 40GBASE-CR4 and 100GBASE-CR10 PMDs

Associated clause	40GBASE-CR4	100GBASE-CR10
81—RS	Required	Required
81—XLGMII[a]	Optional	Not applicable
81—CGMII[b]	Not applicable	Optional
82—PCS for 40GBASE-R	Required	Not applicable
82—PCS for 100GBASE-R	Not applicable	Required
74—BASE-R FEC	Optional	Optional
83—PMA for 40GBASE-R	Required	Not applicable
83—PMA for 100GBASE-R	Not applicable	Required
83A—XLAUI	Optional	Not applicable
83A—CAUI-10	Not applicable	Optional
83D—CAUI-4	Not applicable	Optional
73—Auto-Negotiation	Required	Required
78—Energy Efficient Ethernet	Optional	Optional

[a]The XLGMII is an optional interface. However, if the XLGMII is not implemented, a conforming implementation must behave functionally as though the RS and XLGMII were present.
[b]The CGMII is an optional interface. However, if the CGMII is not implemented, a conforming implementation must behave functionally as though the RS and CGMII were present.

Figure 85–1 shows the relationship of the 40GBASE-CR4 and 100GBASE-CR10 PMD sublayers and MDI to the ISO/IEC Open System Interconnection (OSI) reference model.

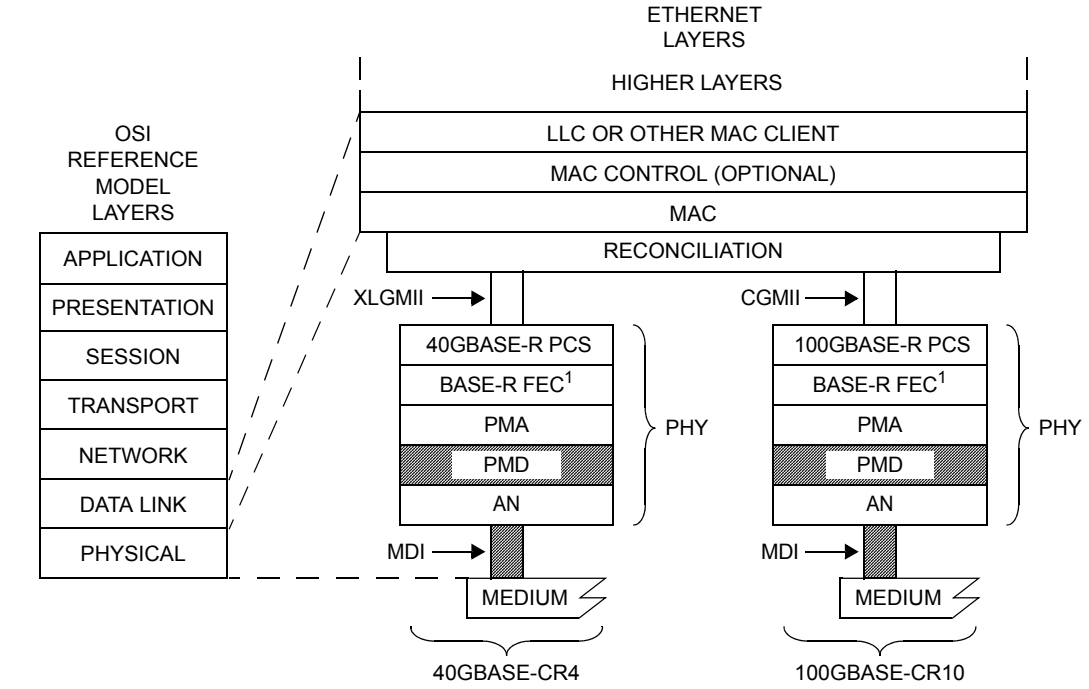

AN = AUTO-NEGOTIATION
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

CR = PMD FOR COPPER — 7 m
NOTE 1—OPTIONAL

Figure 85–1—40GBASE-CR4 and 100GBASE-CR10 PMDs relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

100GBASE-CR10 and 40GBASE-CR4 PHYs with the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option may optionally enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization (see Clause 78).

85.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 40GBASE-CR4 and 100GBASE-CR10 PMDs. The service interfaces for these PMDs are described in an abstract manner and do not imply any particular implementation. The PMD service interface supports the exchange of encoded data between the PMA and PMD entities. The PMD translates the encoded data to and from signals suitable for the specified medium.The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_*i*.request
PMD:IS_UNITDATA_*i*.indication
PMD:IS_SIGNAL.indication

The 40GBASE-CR4 PMD has four parallel bit streams, hence *i* = 0 to 3 for 40GBASE-CR4, and the 100GBASE-CR10 PMD has ten parallel bit streams, hence *i* = 0 to 9 for 100GBASE-CR10. The PMA (or the PMD) continuously sends four or ten parallel bit streams to the PMD (or the PMA), one per lane, each at a nominal signaling rate of 10.3125 GBd.

For 40GBASE-CR4, SIGNAL_DETECT indicates the successful completion of the start-up protocol on all four lanes. For 100GBASE-CR10, SIGNAL_DETECT indicates the successful completion of the start-up protocol on all ten lanes.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the PMD:IS_UNITDATA_*i*.indication parameters are undefined.

The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication primitive.

If the optional EEE capability with the deep sleep mode option is supported (see 78.1.3.3.1) then the inter-sublayer service interface includes two additional primitives defined as follows:

PMD:IS_TX_MODE.request(tx_mode)
PMD:IS_RX_MODE.request(rx_mode)

The tx_mode parameter takes on one of up to three values: DATA, QUIET, or ALERT. When tx_mode = QUIET, transmission is disabled; when tx_mode = ALERT, the alert signal is transmitted (see 85.7.2).

The rx_mode parameter is used to communicate the state of the PCS LPI receive function and takes the value QUIET or DATA.

85.3 PCS requirements for Auto-Negotiation (AN) service interface

The PCS associated with this PMD is required to support the AN service interface primitive AN_LINK.indication defined in 73.9. (See 82.6.)

The 40GBASE-CR4 PHY may be extended using XLAUI, a physical instantiation of the inter-sublayer service interface between devices. Similarly, the 100GBASE-CR10 PHY may be extended using CAUI-n. If XLAUI or CAUI-n is instantiated, the AN_LINK(link_status).indication is relayed from the device with the PCS sublayer to the device with the AN sublayer by means at the discretion of the implementer. As examples, the implementer may employ use of pervasive management or employ a dedicated electrical signal to relay the state of link_status as indicated by the PCS sublayer on one device to the AN sublayer on the other device.

85.4 Delay constraints

A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

The sum of the transmit and the receive delays at one end of the link contributed by the 40GBASE-CR4 PMD and AN shall be no more than 4096 bit times (8 pause_quanta or 102.4 ns). The delay through the medium is not included.

The sum of the transmit and the receive delays at one end of the link contributed by the 100GBASE-CR10 PMD and AN shall be no more than 9728 bit times (19 pause_quanta or 97.28 ns). The delay through the medium is not included.

85.5 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the PCS. The Skew Variation must also be limited to ensure that a given PCS lane always traverses the same physical lane. Skew and Skew Variation are defined in 80.5 and are specified at the points SP1 to SP6 shown in Figure 80–6 and Figure 80–7.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns. For more information on Skew and Skew Variation, see 80.5.

85.6 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 illustrated in Table 85–2 for 40GBASE-CR4 and 100GBASE-CR10. Mapping of MDIO status variables to PMD status variables is illustrated in Table 85–3 for 40GBASE-CR4 and 100GBASE-CR10.

Table 85–2—40GBASE-CR4 and 100GBASE-CR10 MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/ bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable register	1.9.0	Global_PMD_transmit_disable
PMD transmit disable 9 to PMD transmit disable 0	PMD transmit disable register	1.9.10 to 1.9.1	PMD_transmit_disable_9 to PMD_transmit_disable_0
Restart training	BASE-R PMD control register	1.150.0	mr_restart_training
Training enable	BASE-R PMD control register	1.150.1	mr_training_enable

Table 85–3—40GBASE-CR4 and 100GBASE-CR10 MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	Global_PMD_signal_detect
PMD receive signal detect 9 to PMD receive signal detect 0	PMD receive signal detect register	1.10.10 to 1.10.1	PMD_signal_detect_9 to PMD_signal_detect_0
40GBASE-CR4 deep sleep	EEE capability	1.16.1	—
100GBASE-CR10 deep sleep	EEE capability	1.16.8	—
Receiver status 9	BASE-R PMD status 3 register	1.157.4	rx_trained_9
Frame lock 9	BASE-R PMD status 3 register	1.157.5	frame_lock_9
Start-up protocol status 9	BASE-R PMD status 3 register	1.157.6	training_9
Training failure 9	BASE-R PMD status 3 register	1.157.7	training_failure_9
Receiver status 8	BASE-R PMD status 3 register	1.157.0	rx_trained_8
Frame lock 8	BASE-R PMD status 3 register	1.157.1	frame_lock_8
Start-up protocol status 8	BASE-R PMD status 3 register	1.157.2	training_8
Training failure 8	BASE-R PMD status 3 register	1.157.3	training_failure_8
Receiver status 7	BASE-R PMD status 2 register	1.156.12	rx_trained_7
Frame lock 7	BASE-R PMD status 2 register	1.156.13	frame_lock_7
Start-up protocol status 7	BASE-R PMD status 2 register	1.156.14	training_7
Training failure 7	BASE-R PMD status 2 register	1.156.15	training_failure_7
Receiver status 6	BASE-R PMD status 2 register	1.156.8	rx_trained_6
Frame lock 6	BASE-R PMD status 2 register	1.156.9	frame_lock_6
Start-up protocol status 6	BASE-R PMD status 2 register	1.156.10	training_6
Training failure 6	BASE-R PMD status 2 register	1.156.11	training_failure_6
Receiver status 5	BASE-R PMD status 2 register	1.156.4	rx_trained_5
Frame lock 5	BASE-R PMD status 2 register	1.156.5	frame_lock_5
Start-up protocol status 5	BASE-R PMD status 2 register	1.156.6	training_5
Training failure 5	BASE-R PMD status 2 register	1.156.7	training_failure_5
Receiver status 4	BASE-R PMD status 2 register	1.156.0	rx_trained_4
Frame lock 4	BASE-R PMD status 2 register	1.156.1	frame_lock_4
Start-up protocol status 4	BASE-R PMD status 2 register	1.156.2	training_4
Training failure 4	BASE-R PMD status 2 register	1.156.3	training_failure_4

Table 85–3—40GBASE-CR4 and 100GBASE-CR10 MDIO/PMD status variable mapping *(continued)*

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Receiver status 3	BASE-R PMD status register	1.151.12	rx_trained_3
Frame lock 3	BASE-R PMD status register	1.151.13	frame_lock_3
Start-up protocol status 3	BASE-R PMD status register	1.151.14	training_3
Training failure 3	BASE-R PMD status register	1.151.15	training_failure_3
Receiver status 2	BASE-R PMD status register	1.151.8	rx_trained_2
Frame lock 2	BASE-R PMD status register	1.151.9	frame_lock_2
Start-up protocol status 2	BASE-R PMD status register	1.151.10	training_2
Training failure 2	BASE-R PMD status register	1.151.11	training_failure_2
Receiver status 1	BASE-R PMD status register	1.151.4	rx_trained_1
Frame lock 1	BASE-R PMD status register	1.151.5	frame_lock_1
Start-up protocol status 1	BASE-R PMD status register	1.151.6	training_1
Training failure 1	BASE-R PMD status register	1.151.7	training_failure_1
Receiver status 0	BASE-R PMD status register	1.151.0	rx_trained_0
Frame lock 0	BASE-R PMD status register	1.151.1	frame_lock_0
Start-up protocol status 0	BASE-R PMD status register	1.151.2	training_0
Training failure 0	BASE-R PMD status register	1.151.3	training_failure_0

85.7 PMD functional specifications

85.7.1 Link block diagram

A 40GBASE-CR4 or 100GBASE-CR10 link in one direction is illustrated in Figure 85–2. For purposes of system conformance, the PMD sublayer is standardized at the test points described in this subclause. The electrical transmit signal is defined at TP2. Unless specified otherwise, all transmitter measurements and tests defined in Table 85–5 are made at TP2 utilizing the test fixture specified in 85.8.3.5. Unless specified otherwise, all receiver measurements and tests defined in 85.8.4 are made at TP3 utilizing the test fixture specified in 85.8.3.5. A mated connector pair has been included in both the transmitter and receiver specifications defined in 85.8.3 and 85.8.4. The maximum insertion loss from TP0 to TP2 or TP3 to TP5 including the test fixture is specified in 85.8.3.4.

The 40GBASE-CR4 and 100GBASE-CR10 channels are defined between the transmitter (TP0) and receiver blocks (TP5) to include the transmitter and receiver differential controlled impedance printed circuit board insertion loss and the cable assembly insertion loss as illustrated in Figure 85–2. Annex 85A provides information on parameters associated with test points TP0 and TP5 that may not be testable in an implemented system. All cable assembly measurements are to be made between TP1 and TP4 as illustrated in Figure 85–2. The cable assembly test fixture of Figure 85–13 or its functional equivalent, is required for measuring the cable assembly specifications in 85.10 at TP1 and TP4. Two mated connector pairs and the cable assembly test fixture have been included in the cable assembly specifications defined in 85.10.

Transmitter and receiver differential controlled impedance printed circuit board insertion losses defined between TP0–TP1 and TP4–TP5 respectively are provided informatively in Annex 85A.

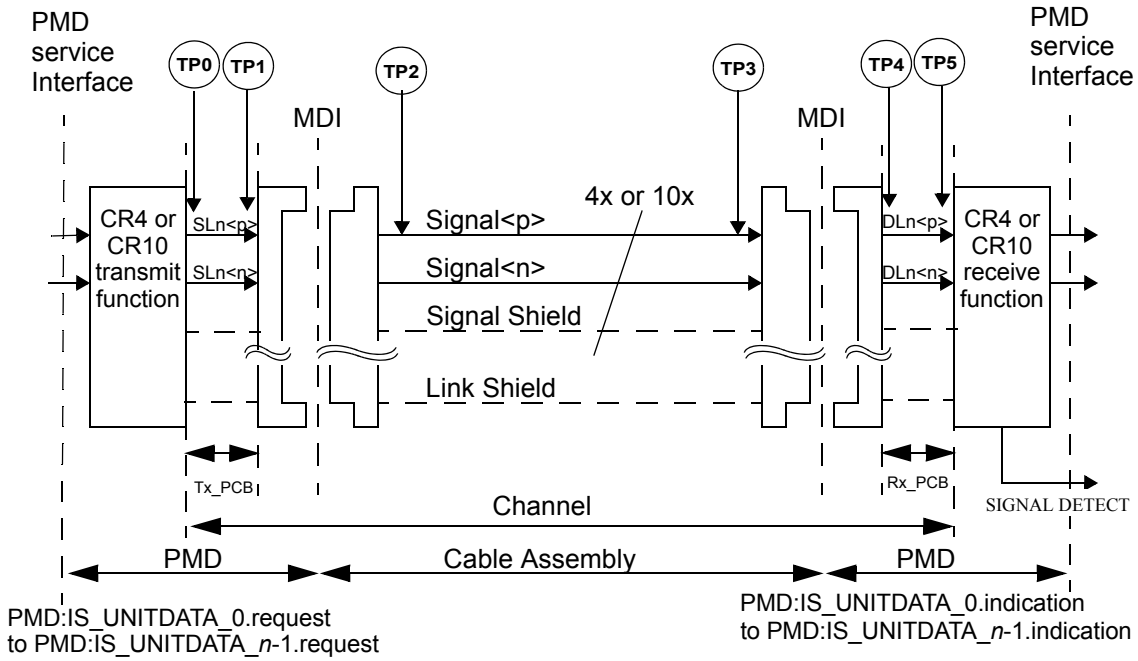

Figure 85–2—40GBASE-CR4 or 100GBASE-CR10 link (half link is illustrated)

Note that the source lanes (SL), signals SLn<p>, and SLn<n> are the positive and negative sides of the transmitters differential signal pairs and the destination lanes (DL) signals, DLn<p>, and DLn<n> are the positive and negative sides of the receivers differential signal pairs for lane n (n = 0, 1, 2, 3 or n = 0, 1, 2, 3, 4, 5, 6, 7, 8, 9).

Table 85–4 describes the defined test points illustrated in Figure 85–2.

Table 85–4—40GBASE-CR4 and 100GBASE-CR10 test points

Test points	Description
TP0 to TP5	The 40GBASE-CR4 and 100GBASE-CR10 channels including the transmitter and receiver differential controlled impedance printed circuit board insertion loss and the cable assembly insertion loss.
TP1 to TP4	All cable assembly measurements are to be made between TP1 and TP4 as illustrated in Figure 85–2. The cable assembly test fixture of Figure 85–13 or its functional equivalent, is required for measuring the cable assembly specifications in 85.10 at TP1 and TP4.
TP0 to TP1 TP4 to TP5	A mated connector pair has been included in both the transmitter and receiver specifications defined in 85.8.3 and 85.8.4. The maximum insertion loss from TP0 to TP2 or TP3 to TP5 including the test fixture is specified in 85.8.3.4.
TP2	Unless specified otherwise, all transmitter measurements and tests defined in Table 85–5 are made at TP2 utilizing the test fixture specified in 85.8.3.5.
TP3	Unless specified otherwise, all receiver measurements and tests defined in 85.8.4 are made at TP3 utilizing the test fixture specified in 85.8.3.5.

85.7.2 PMD Transmit function

The 40GBASE-CR4 PMD Transmit function shall convert the four bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_3.request into four separate electrical streams. The four electrical signal streams shall then be delivered to the MDI, all according to the transmit electrical specifications in 85.8.3. A positive output voltage of SL<p> minus SL<n> (differential voltage) shall correspond to tx_bit = one. The 100GBASE-CR10 PMD Transmit function shall convert the ten bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_9.request. The ten electrical signal streams shall then be delivered to the MDI, all according to the transmit electrical specifications in 85.8.3. A positive output voltage of SL<p> minus SL<n> (differential voltage) shall correspond to tx_bit = one.

If the EEE capability with the deep sleep mode option is supported (see Clause 78) then when tx_mode is set to ALERT, the PMD transmit function shall transmit a repeating 16-bit pattern, hexadecimal 0xFF00, on each lane. This sequence is transmitted regardless of the value of tx_bit presented by the PMD:IS_UNITDATA_i.request primitive. When tx_mode is ALERT, the transmitter equalizer taps are set to the preset state specified in 72.6.10.2.3.1. When tx_mode is QUIET, the transmitter is disabled as specified in 85.7.6. For all other states of tx_mode, the driver coefficients are restored to their states resolved during training.

85.7.3 PMD Receive function

The 40GBASE-CR4 PMD Receive function shall convert the four electrical streams from the MDI into four bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_3.indication. A positive input voltage of DL<p> minus DL<n> (differential voltage) shall correspond to rx_bit = one. The 100GBASE-CR10 PMD Receive function shall convert the ten electrical streams from the MDI into ten bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_9.indication. A positive input voltage of DL<p> minus DL<n> (differential voltage) shall correspond to rx_bit = one.

85.7.4 Global PMD signal detect function

The Global PMD signal detect function shall continuously report the message PMD:IS_SIGNAL.indication (SIGNAL_DETECT) to the PMD service interface. SIGNAL_DETECT, while normally intended to be an indicator of signal presence, is used by 40GBASE-CR4 and 100GBASE-CR10 to indicate the successful completion of the start-up protocol on all lanes. The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication primitive. When the PHY supports the optional EEE capability with deep sleep mode, PMD_SIGNAL.indication is also used to indicate when the ALERT signal is detected, which corresponds to the beginning of a refresh or a wake.

When the PHY does not support the EEE capability or if the PHY supports the EEE capability and rx_mode is set to DATA, SIGNAL_DETECT shall be set to FAIL following system reset or the manual reset of the training state diagram. Upon successful completion of training on all lanes, SIGNAL_DETECT shall be set to OK.

When the PHY does not support the EEE capability or if the PHY supports the EEE capability and rx_mode is set to DATA, if training is disabled by management, SIGNAL_DETECT shall be set to OK. When the PHY supports the EEE capability with deep sleep mode, SIGNAL_DETECT is set to FAIL following a transition from rx_mode = DATA to rx_mode = QUIET. When rx_mode = QUIET, SIGNAL_DETECT shall be set to OK within 500 ns following the application of a signal at the receiver input that corresponds to an ALERT transmission (see 85.7.2) from the link partner. While rx_mode = QUIET, SIGNAL_DETECT shall be held at FAIL as long as the signal at the receiver input corresponds to a QUIET tx_mode (see 85.7.6) of the link partner.

If the MDIO interface is implemented, then Global_PMD_signal_detect (1.10.0) shall be continuously set to the value of SIGNAL_DETECT as described in 45.2.1.9.7.

85.7.5 PMD lane-by-lane signal detect function

When the MDIO is implemented, each PMD_signal_detect_*i* value, where *i* represents the lane number in the range 0:3 for 40GBASE-CR4 and 0:9 for 100GBASE-CR10, shall be continuously updated as described in the following two paragraphs.

For 40GBASE-CR4 PMD_signal_detect_0 (1.10.1), PMD_signal_detect_1 (1.10.2), PMD_signal_detect_2 (1.10.3) and PMD_signal_detect_3 (1.10.4) shall be set to one or zero depending on whether a particular lane's signal_detect, as defined by the training state diagram in Figure 72–5, returns true or false.

For 100GBASE-CR10 PMD_signal_detect_0 (1.10.1), PMD_signal_detect_1 (1.10.2), PMD_signal_detect_2 (1.10.3), PMD_signal_detect_3 (1.10.4), PMD_signal_detect_4 (1.10.5), PMD_signal_detect_5 (1.10.6), PMD_signal_detect_6 (1.10.7), PMD_signal_detect_7 (1.10.8), PMD_signal_detect_8 (1.10.9), and PMD_signal_detect_9 (1.10.10) shall be set to one or zero depending on whether a particular lane's signal_detect, as defined by the training state diagram in Figure 72–5, returns true or false.

85.7.6 Global PMD transmit disable function

The Global_PMD_transmit_disable function is mandatory if EEE with the deep sleep mode option is supported and is otherwise optional. When implemented, it allows all of the transmitters to be disabled with a single variable.

a) When 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 85–5.
b) If a PMD_fault (85.7.9) is detected, then the PMD may turn off the electrical transmitter in all lanes.
c) Loopback, as defined in 85.7.8, shall not be affected by Global_PMD_transmit_disable.
d) For EEE capability, the PMD_transmit_disable function shall turn off the transmitter after tx_mode is set to QUIET within a time and voltage level specified in 72.7.1.4. The PMD_transmit_disable function shall turn on the transmitter after tx_mode is set to DATA or ALERT within the time and voltage level specified in 72.7.1.4.

85.7.7 PMD lane-by-lane transmit disable function

The PMD_transmit_disable_*i* function (where *i* represents the lane number in the range 0:3 or 0:9) is optional and allows the electrical transmitter in each lane to be selectively disabled. When this function is supported, it shall meet the following requirements:

a) When a PMD_transmit_disable_*i* variable is set to one, this function shall turn off the transmitter associated with that variable such that it drives a constant level (i.e., no transitions) and does not exceed the maximum differential peak-to-peak output voltage specified in Table 85–5.
b) If a PMD_fault (85.7.9) is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the electrical transmitter in each lane.
c) Loopback, as defined in 85.7.8, shall not be affected by PMD_transmit_disable_*i*.

85.7.8 Loopback mode

Local loopback mode shall be provided by the adjacent PMA (see 83.5.8) for the 40GBASE-CR4 and 100GBASE-CR10 PMDs as a test function to the device. When loopback mode is enabled, transmission requests passed to each transmitter are sent directly to the corresponding receiver, overriding any signal

detected by each receiver on its attached link. Note that loopback mode does not affect the state of the transmitter, which continues to send data (unless disabled).

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.

85.7.9 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. The PMD fault function shall be mapped to bit 1.1.7 as listed in Table 85–3.

85.7.10 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 Global_PMD_transmit_disable function.

If a PMD_transmit_fault (optional) is detected, then the Global_PMD_transmit_disable function should also be asserted.

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.

85.7.11 PMD receive fault function

The PMD_receive_fault function is optional. The faults detected by this function are implementation specific.

If the MDIO interface is implemented, then this function shall contribute to the Receive fault bit as specified in 45.2.1.7.5.

85.7.12 PMD control function

Each lane of the 40GBASE-CR4 or 100GBASE-CR10 PMD shall use the same control function as 10GBASE-KR, as defined in 72.6.10.

The random seed for the training pattern described in 72.6.10.2.6 shall be different for each of the lanes.

85.8 MDI Electrical specifications for 40GBASE-CR4 and 100GBASE-CR10

85.8.1 Signal levels

The 40GBASE-CR4 and 100GBASE-CR10 MDI is a low-swing AC-coupled differential interface. AC-coupling at the receiver, as defined in 85.8.4.5, allows for interoperability between components operating from different supply voltages. Low-swing differential signaling provides noise immunity and improved electromagnetic interference (EMI).

85.8.2 Signal paths

The 40GBASE-CR4 and 100GBASE-CR10 MDI signal paths are point-to-point connections. Each path corresponds to a 40GBASE-CR4 or 100GBASE-CR10 MDI lane and comprises two complementary signals, which form a balanced differential pair. For 40GBASE-CR4, there are four differential paths in each direction for a total of eight pairs, or sixteen connections. For 100GBASE-CR10, there are ten differential paths in each direction for a total of 20 pairs, or forty connections. The signal paths are intended to operate on twinaxial cable assemblies ranging from 0.5 m to 7 m in length, as described in 85.10.

85.8.3 Transmitter characteristics

Transmitter characteristics shall meet specifications summarized in Table 85–5 at TP2 unless otherwise noted. The transmitter specifications at TP0 are provided informatively in Annex 85A, Table 85A–1.

Table 85–5—Transmitter characteristics at TP2 summary

Parameter	Subclause reference	Value	Units
Signaling rate, per lane	85.8.3.9	10.3125 ± 100 ppm	GBd
Unit interval nominal	85.8.3.9	96.969697	ps
Differential peak-to-peak output voltage (max) with Tx disabled		30	mV
Common-mode voltage limits	72.7.1.4	0 to 1.9	V
Differential output return loss (min)	85.8.3.1	See Equation (85–1)	dB
AC common-mode output voltage (max., RMS)		30	mV
Common-mode voltage deviation (max) during LPI	72.7.1.4	150	mV
Amplitude peak-to-peak (max)	72.7.1.4	1200[a]	mV
Transmitter DC amplitude[b]	85.8.3.3	0.34 min, 0.6 max	V
Linear fit pulse (min)[c]	85.8.3.3	0.63 × Transmitter DC amplitude	V
Transmitted waveform max RMS normalized error (linear fit), "*e*" abs coefficient step size minimum precursor fullscale range minimum post cursor fullscale range	 85.8.3.3 85.8.3.3.2 85.8.3.3.3 85.8.3.3.3	 0.037 0.0083 min, 0.05 max 1.54 4	
Far-end transmit output noise (max) Low insertion loss channel High insertion loss channel	85.8.3.2	 2 See Equation (85–2) 1 See Equation (85–3)	mV
Max output jitter (peak-to-peak) Random jitter[d] Duty Cycle Distortion[e] Total jitter excluding data dependent jitter[f]		 0.15 0.035 0.25	 UI UI UI

[a]The 40GBASE-CR4 Style-1 connector may support 40GBASE-CR4 or XLPPI interfaces. For implementations that support both interfaces, the transmitter should not exceed the XLPPI voltage maximum until a 40GBASE-CR4 cable assembly has been identified. The 100GBASE-CR10 connector may support 100GBASE-CR10 or CPPI interfaces. For implementations that support both interfaces, the transmitter should not exceed the CPPI voltage maximum until a 100GBASE-CR10 cable assembly has been identified.
[b]The transmitter DC amplitude is the sum of linear fit pulse response $p(k)$ from step 3) divided by M from step 3).
[c]The peak of the linear fit pulse response $p(k)$ from step 3).

[d]Random jitter is specified at a BER of 10^{-12}
[e] See 72.7.1.9 for duty cycle distortion definition.
[f]Total jitter at a BER of 10^{-12} measured per 83A.5.1 excluding data dependent jitter (DDJ). DDJ is a jitter component where jitter that is not correlated to the data pattern has been removed. DDJ is measured with PRBS9 as specified in 83.5.10.

85.8.3.1 Transmitter differential output return loss

The differential output return loss, in dB, of the transmitter shall meet Equation (85–1). This output impedance requirement applies to all valid output levels. The reference impedance for differential return loss measurements shall be 100 Ω

$$Return\_loss(f) \geq \begin{cases} 12 - 2\sqrt{f} & 0.05 \leq f < 4.11 \\ 6.3 - 13\log_{10}(f/5.5) & 4.11 \leq f \leq 10 \end{cases} \quad \text{(dB)} \qquad (85\text{–}1)$$

where

f is the frequency in GHz
$Return\_loss(f)$ is the return loss at frequency f

85.8.3.2 Transmitter noise parameter measurements

The far-end transmitter output noise is an additional source of noise to the cable assembly's integrated crosstalk noise (ICN) specified in 85.10.7. The far-end transmitter output noise parameter is characterized using two reference channels; a "low-loss" cable assembly with insertion loss on the reference pair of 6 dB ± 1.5 dB at 5.15625 GHz and cable assembly integrated crosstalk noise (ICN) meeting the requirements of 85.10.7 and a "high-loss" cable assembly with insertion loss on the reference pair of 15.7 dB ± 1.5 dB at 5.15625 GHz and cable assembly ICN meeting the requirements of 85.10.7. The far-end transmitter output noise is characterized as a deviation from the cable assembly ICN using the following procedure:

1) Compute the far-end integrated crosstalk noise σ_{fx} into the reference lane of the cable assembly using the methodology of 85.10.7 and the parameters in Table 85–11.
2) Denote σ_l as the far-end ICN for the low-loss cable assembly.
3) Denote σ_h as the far-end ICN for the high-loss cable assembly.
4) The transmitter under test is connected to one end of the reference cable assembly and the other end is connected to the cable assembly test fixture specified in 85.10.8.
5) All lanes of the cable assembly test fixture are terminated in the reference impedance with the reference lane connected to the measuring instrument.
6) The reference lane of the transmitter under test sends a square wave test pattern as specified in 83.5.10 while all other adjacent transmitter lanes send either scrambled idle or PRBS31.
7) A fixed point on the square wave test pattern is chosen and the RMS deviation from the mean voltage at this observation point is measured. The histogram for RMS noise measurement is 1 UI wide.
8) The measurement should not include the measurement system noise.

For the low-loss cable assembly, the measured RMS deviation from the cable assembly ICN due to the far-end transmitter output noise shall meet the values determined using Equation (85–2).

$$RMSl_{dev} \leq \sqrt{\sigma_l^2 + 2^2} \text{ (mV)} \qquad (85\text{–}2)$$

For the high-loss cable assembly, the measured RMS deviation from the cable assembly ICN due to the far-end transmitter output noise shall meet the values determined using Equation (85–3).

$$RMSh_{dev} \leq \sqrt{\sigma_h^{\,2} + 1^2} \text{ (mV)} \quad (85–3)$$

85.8.3.3 Transmitter output waveform

The 40GBASE-CR4 and 100GBASE-CR10 transmit function includes programmable equalization to compensate for the frequency-dependent loss of the channel and facilitate data recovery at the receiver. The functional model for the transmit equalizer is the three tap transversal filter shown in Figure 85–3.

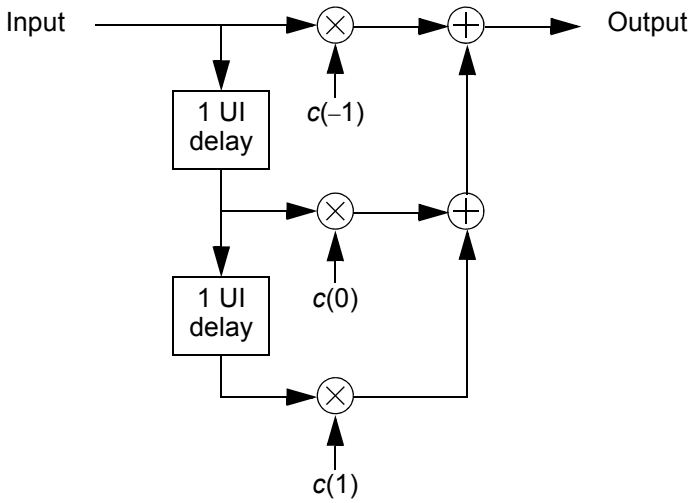

Figure 85–3—Transmit equalizer functional model

The state of the transmit equalizer and hence the transmitted output waveform may be manipulated via the PMD control function defined in 85.7.12 or via the management interface. The transmit function responds to a set of commands issued by the link partner's receive function and conveyed by a back-channel communications path.

This command set includes instructions to

a) Increment coefficient $c(i)$.
b) Decrement coefficient $c(i)$.
c) Hold coefficient $c(i)$ at its current value.
d) Set the coefficients to a pre-defined value (preset or initialize).

In response, the transmit function relays status information to the link partner's receive function. The status messages indicate that

a1) The requested update to coefficient $c(i)$ has completed (updated).
b1) Coefficient $c(i)$ is at its minimum value.
c1) Coefficient $c(i)$ is at its maximum value.
d1) Coefficient $c(i)$ is ready for the next update request (not_updated).

The requirements for the 40GBASE-CR4 and 100GBASE-CR10 transmit equalizer are intended to be similar to the requirements for 10GBASE-KR specified in 72.7.1.10. However, the signal path from the transmit function to TP2 introduces frequency-dependent loss and phase shift that distorts the signal and

makes it difficult to accurately characterize equalizer performance at TP2 using the methodology specified for 10GBASE-KR. The following process enables accurate characterization of the equalizer performance at TP2 by determining and correcting for the frequency-dependent loss and phase shift of the signal path from the transmit function to TP2.

1) The transmitter under test is preset as specified in 72.6.10.2.3.1 such that $c(-1)$ and $c(1)$ are zero and $c(0)$ is its maximum value.
2) Capture at least one complete cycle of the test pattern PRBS9 as specified in 83.5.10 at TP2 per 85.8.3.3.4.
3) Compute the linear fit to the captured waveform and the linear fit pulse response p(k) per 85.8.3.3.5.
4) Define t_x to be the time where the rising edge of the linear fit pulse, p, from step 3) crosses 50% of its peak amplitude.
5) Sample the linear fit pulse, p, at symbol-spaced intervals relative to the time $t_0 = t_x + 0.5$ UI, interpolating as necessary to yield the sampled pulse p_i.
6) Use p_i to compute the vector of coefficients, w, of a N_w-tap symbol-spaced transversal filter that equalizes for the transfer function from the transmit function to TP2 per 85.8.3.3.6.

The parameters of the pulse fit and the equalizing filter are given in Table 85–6. The DC amplitude, the sum of linear fit pulse response, p(k), from step 3) divided by M from step 3), shall be greater than 0.34 V and less than or equal to 0.6 V. The peak of the linear fit pulse response from step 3) shall be greater than $0.63 \times$ DC amplitude. The RMS value of the error between the linear fit from step 3) and the measured waveform, e, normalized to the peak value of the pulse, p, must be no greater than 0.037.

Table 85–6—Normalized transmit pulse template

Description	Symbol	Value	Units
Linear fit pulse length	N_p	8	UI
Linear fit pulse delay	D_p	2	UI
Equalizer length	N_w	7	UI
Equalizer delay	D_w	1	UI

For each configuration of the transmit equalizer

7) Configure the transmitter under test as required by the test.
8) Capture at least one complete cycle of the test pattern PRBS9 as specified in 83.5.10 at TP2 per 85.8.3.3.4.
9) Compute the linear fit to the captured waveform and the linear fit pulse response p(k) per 85.8.3.3.5.
10) Define t_x to be the time where the rising edge of the linear fit pulse response, p(k), from step 9 crosses 50% of its peak amplitude.
11) Sample the linear fit pulse response, p(k), at symbol-spaced intervals relative to the time $t_0 = t_x + 0.5$ UI, interpolating as necessary to yield the sampled pulse p_i.
12) Equalize the sampled pulse p_i using the coefficient vector, w, computed in step 6) per 85.8.3.3.6 to yield the equalized pulse q_i.

The RMS value of the error between the linear fit from step 9) and the measured waveform, e, normalized to the peak value of the pulse, p, must be no greater than 0.037.

The normalized amplitude of coefficient $c(-1)$ is the value of q_i at time $t_0 + (D_p - 1)$ UI. The normalized amplitude of coefficient $c(0)$ is the value of q_i at time $t_0 + D_p$ UI. The normalized amplitude of coefficient $c(1)$ is the value of q_i at time $t_0 + (D_p + 1)$ UI.

85.8.3.3.1 Coefficient initialization

When the PMD enters the INITIALIZE state of the Training state diagram (Figure 72–5) or receives a valid request to "initialize" from the link partner, the coefficients of the transmit equalizer shall be configured such that the ratio $(c(0)+c(1)-c(-1))/(c(0)+c(1)+c(-1))$ is 1.29 ±10% and the ratio $(c(0)-c(1)+c(-1))/(c(0)+c(1)+c(-1))$ is 2.57 ±10%. These requirements apply upon the assertion a coefficient status report of "updated" for all coefficients.

85.8.3.3.2 Coefficient step size

The change in the normalized amplitude of coefficient $c(i)$ corresponding to a request to "increment" that coefficient shall be between 0.0083 and 0.05. The change in the normalized amplitude of coefficient $c(i)$ corresponding to a request to "decrement" that coefficient shall be between –0.05 and –0.0083.

The change in the normalized amplitude of the coefficient is defined to be the difference in the value measured prior to the assertion of the "increment" or "decrement" request (e.g., the coefficient update request for all coefficients is "hold") and the value upon the assertion of a coefficient status report of "updated" for that coefficient.

85.8.3.3.3 Coefficient range

When sufficient "increment" or "decrement" requests have been received for a given coefficient, the coefficient will reach a lower or upper bound based on the coefficient range or restrictions placed on the minimum steady state differential output voltage or the maximum peak-to-peak differential output voltage.

With $c(-1)$ set to zero and both $c(0)$ and $c(1)$ having received sufficient "decrement" requests so that they are at their respective minimum values, the ratio $(c(0) - c(1))/(c(0) + c(1))$ shall be greater then or equal to 4.

With $c(1)$ set to zero and both $c(-1)$ and $c(0)$ having received sufficient "decrement" requests so that they are at their respective minimum values, the ratio $(c(0) - c(-1))/(c(0) + c(-1))$ shall be greater then or equal to 1.54.

Note that a coefficient may be set to zero by first asserting a coefficient preset request and then manipulating the other coefficients as required by the test.

85.8.3.3.4 Waveform acquisition

The transmitter under test repetitively transmits the specified test pattern. The waveform shall be captured with an effective sample rate that is M times the signaling rate of the transmitter under test. The value of M shall be an integer not less than 7. Averaging multiple waveform captures is recommended.

The captured waveform shall represent an integer number of repetitions of the test pattern totaling N bits. Hence the length of the captured waveform should be MN samples. The waveform should be aligned such that the first M samples of waveform correspond to the first bit of the test pattern, the second M samples to the second bit, and so on.

85.8.3.3.5 Linear fit to the waveform measurement at TP2

Given the captured waveform $y(k)$ and corresponding aligned symbols $x(n)$ derived from the procedure defined in 85.8.3.3.4, define the M-by-N waveform matrix Y as shown in Equation (85–4).

$$Y = \begin{bmatrix} y(1) & y(M+1) & \ldots & y(M(N-1)+1) \\ y(2) & y(M+2) & \ldots & y(M(N-1)+2) \\ \ldots & \ldots & \ldots & \ldots \\ y(M) & y(2M) & \ldots & y(MN) \end{bmatrix} \quad (85\text{–}4)$$

Rotate the symbols vector x by the specified pulse delay D_p to yield x_r as shown in Equation (85–5).

$$x_r = \left[x(D_p+1) \;\; x(D_p+2) \ldots x(N) \; x(1) \; \ldots \; x(D_p) \right] \quad (85\text{–}5)$$

Define the matrix X to be an N-by-N matrix derived from x_r as shown in Equation (85–6).

$$X = \begin{bmatrix} x_r(1) & x_r(2) & \ldots & x_r(N) \\ x_r(N) & x_r(1) & \ldots & x_r(N-1) \\ \ldots & \ldots & \ldots & \ldots \\ x_r(2) & x_r(3) & \ldots & x_r(1) \end{bmatrix} \quad (85\text{–}6)$$

Define the matrix X_1 to be the first N_p rows of X concatenated with a row vector of ones of length N. The M-by-$(N_p + 1)$ coefficient matrix, P, corresponding to the linear fit is then defined by Equation (85–7). The superscript "T" denotes the matrix transpose operator.

$$P = YX_1^T(X_1X_1^T)^{-1} \quad (85\text{–}7)$$

The error waveform, $e(k)$, is then read column-wise from the elements of E as shown in Equation (85–8).

$$E = PX_1 - Y = \begin{bmatrix} e(1) & e(M+1) & \ldots & e(M(N-1)+1) \\ e(2) & e(M+2) & \ldots & e(M(N-1)+2) \\ \ldots & \ldots & \ldots & \ldots \\ e(M) & e(2M) & \ldots & e(MN) \end{bmatrix} \quad (85\text{–}8)$$

Define P_1 to be a matrix consisting of the first N_p columns of the matrix P as shown in Equation (85–9). The linear fit pulse response, $p(k)$, is then read column-wise from the elements of P_1.

$$P_1 = \begin{bmatrix} p(1) & p(M+1) & \ldots & p(M(N_p-1)+1) \\ p(2) & p(M+2) & \ldots & p(M(N_p-1)+2) \\ \ldots & \ldots & \ldots & \ldots \\ p(M) & p(2M) & \ldots & p(MN_p) \end{bmatrix} \quad (85\text{–}9)$$

85.8.3.3.6 Transfer function between the transmit function and TP2

Rotate the sampled pulse response p_i by the specified equalizer delay D_w to yield p_r as shown in Equation (85–10).

$$p_r = \left[p_i(D_w+1) \;\; p_i(D_p+2) \ldots p_i(N_p) \; p_i(1) \; \ldots \; p_i(D_w) \right] \quad (85\text{–}10)$$

Define the matrix P_2 to be an N_p-by-N_p matrix derived from p_r as shown in Equation (85–11).

$$P_2 = \begin{bmatrix} p_r(1) & p_r(N_p) & \dots & p_r(2) \\ p_r(2) & p_r(1) & \dots & p_r(3) \\ \dots & \dots & \dots & \dots \\ p_r(N_p) & p_r(N_p - 1) & \dots & p_r(1) \end{bmatrix} \quad (85\text{–}11)$$

Define the matrix P_3 to be the first N_w columns of P_2. Define a unit pulse column vector x_p of length N_p. The value of element $x_p(D_p + 1)$ is 1 and all other elements have a value of 0. The vector of filter coefficients w that equalizes p_i is then defined by Equation (85–12).

$$w = (P_3^T P_3)^{-1} P_3^T x_p \quad (85\text{–}12)$$

Given the column vector of equalizer coefficients, w, the equalized pulse response q_i is determined by Equation (85–13).

$$q_i = P_3 w \quad (85\text{–}13)$$

85.8.3.4 Insertion loss TP0 to TP2 or TP3 to TP5

Transmitter measurements and tests defined in Table 85–5 are made at TP2 or TP3 using the test fixture of Figure 85–5, or its functional equivalent. The maximum insertion loss from TP0 to TP2 or TP3 to TP5 including the test fixture is determined using Equation (85–14).

$$Insertion\_loss(f) \le \left\{ \begin{array}{ll} 0.114 + 0.8914\sqrt{f} + 0.846f & 0.01 \le f < 7 \\ -35.91 + 6.3291f & 7 \le f < 8 \\ 14.72 & 8 \le f \le 10 \end{array} \right\} \text{ (dB)} \quad (85\text{–}14)$$

where

f is the frequency in GHz
$Insertion\_loss(f)$ is the insertion loss at frequency f

The maximum insertion loss of TP0 to TP2 or TP3 to TP5 is illustrated in Figure 85–4.

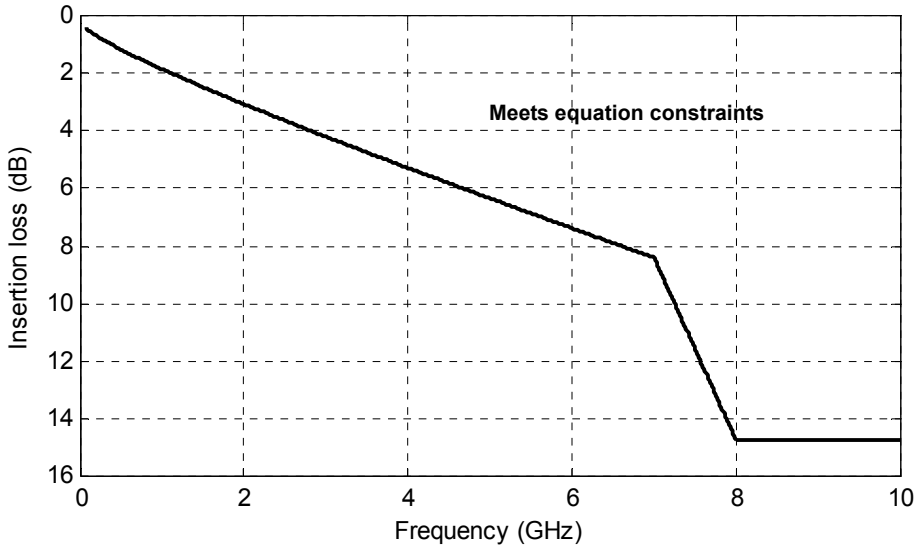

Figure 85–4—Maximum insertion loss TP0 to TP2 or TP3 to TP5

85.8.3.5 Test fixture

The test fixture of Figure 85–5, or its functional equivalent, is required for measuring the transmitter specifications in 85.8.3 at TP2 and the receiver return loss at TP3. TP2 and TP3 are illustrated in Figure 85–2. Figure 85–5 illustrates the test fixture attached to TP2 or TP3.

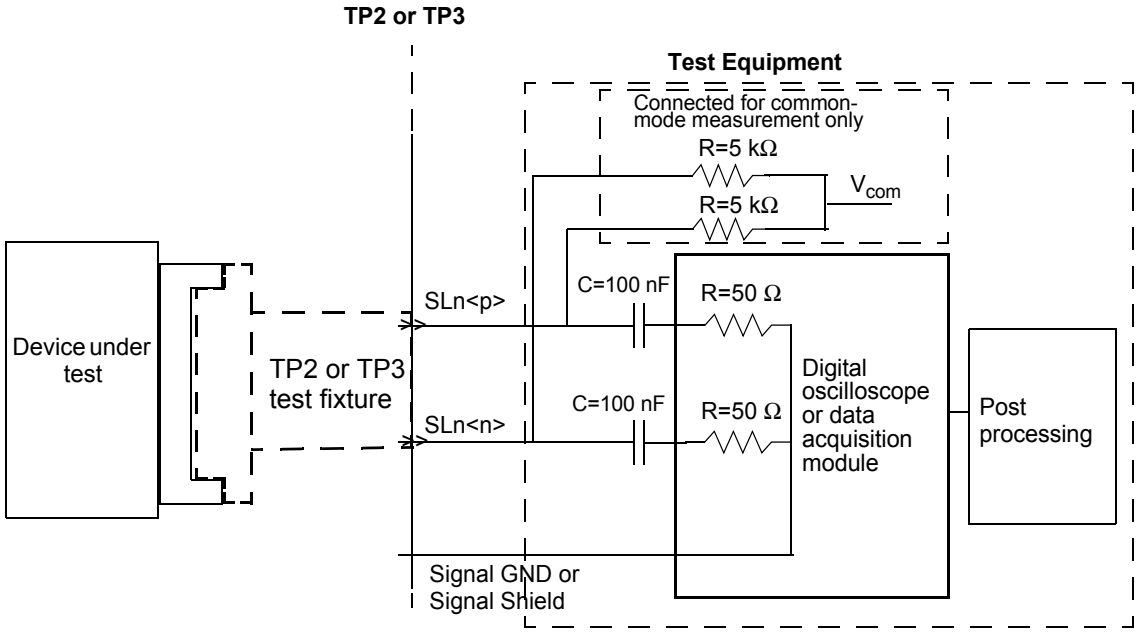

Figure 85–5—Transmitter test fixture

85.8.3.6 Test fixture impedance

The differential return loss, in dB, of the test fixture is specified in a mated state and shall meet the requirements of 85.10.9.2.

85.8.3.7 Test fixture insertion loss

The reference test fixture printed circuit board insertion loss shall meet the values determined using Equation (85–15). The effects of differences between the insertion loss of an actual test fixture and the reference insertion loss are to be accounted for in the measurements.

$$IL_{tfref}(f) = 0.01 + 0.3\sqrt{f} + 0.11f \text{ (dB)} \quad (85–15)$$

for 0.01 GHz ≤ f ≤ 10 GHz

where

f is the frequency in GHz
$IL_{tfref}(f)$ is the reference test fixture PCB insertion loss at frequency f

85.8.3.8 Data dependent jitter (DDJ)

An oscilloscope, time interval analyzer, or other instrument with equivalent capability may be used to measure DDJ. A repeating PRBS9 pseudo-random test pattern, 511 bits long, is used. For DDJ measurements, the measurement bandwidth should be at least 12 GHz. If the measurement bandwidth affects the result, it can be corrected for by post-processing.

Establish a crossing level equal to the average value of the entire waveform being measured. Synchronize the instrument to the pattern repetition frequency and average the waveforms or the crossing times sufficiently to remove the effects of random jitter and noise in the system. The PRBS9 pattern has 128 positive-going transitions and 128 negative-going transitions. The mean time of each crossing is then compared to the expected time of the crossing, and a set of 256 timing variations is determined. Crossings earlier than expected give a negative variation. Crossings later than expected give a positive variation. DDJ is the range (maximum minus minimum) of the timing variations. Note that it may be convenient to align the expected time of one of the crossings with the measured mean crossing.

Figure 85–6 illustrates the method. The vertical axis is in arbitrary units, and the horizontal axis is plotted in UI. The waveform is AC-coupled to an average value of 0, therefore 0 is the appropriate crossing level. The rectangular waveform shows the ideal crossing times, and the other is the waveform with jitter that is being measured. Only 16 UI are shown (out of 511). The waveforms have been arbitrarily aligned with ($\Delta t_2 = 0$) at 5 UI.

Data dependent jitter is defined as in Equation (85–16).

$$DDJ = max(\Delta t_1, \Delta t_2, \ldots \Delta t_{256}) - min(\Delta t_1, \Delta t_2, \ldots \Delta t_{256}) \quad (85–16)$$

85.8.3.9 Signaling rate range

The 40GBASE-CR4 and 100GBASE-CR10 MDI signaling rate shall be 10.3125 GBd ±100 ppm per lane. The corresponding unit interval is nominally 96.969697 ps.

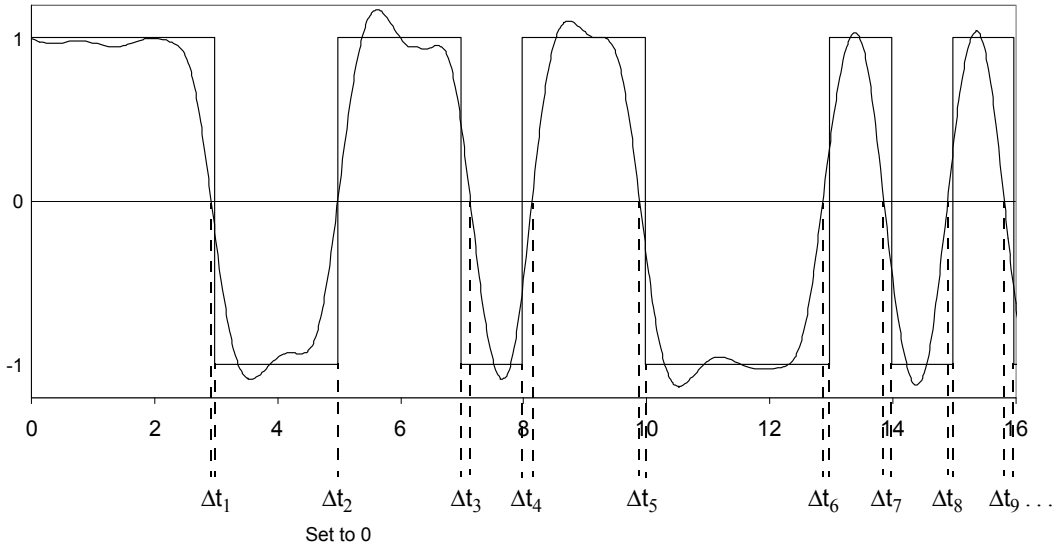

Figure 85–6—Data dependent jitter test method

85.8.4 Receiver characteristics at TP3 summary

The receiver characteristics shall meet the specifications summarized in Table 85–7 at TP3 unless otherwise noted.

Table 85–7—Receiver characteristics at TP3 summary

Parameter	Subclause reference	Value	Units
Bit error ratio	85.8.4.3	10^{-12} or better	
Signaling rate, per lane	85.8.4.4	10.3125 ± 100 ppm	GBd
Unit interval (UI) nominal	85.8.4.4	96.969697	ps
Differential peak-to-peak input amplitude tolerance (max)	72.7.2.4	1200	mV
Differential input return loss (min)[a]	85.8.4.1	Equation (85–17)	dB
Differential to common-mode input return loss		10 min from 10 MHz to 10 GHz	dB

[a]Relative to 100 Ω differential.

85.8.4.1 Receiver differential input return loss

The differential input return loss, in dB, of the receiver shall meet Equation (85–17). This return loss requirement applies at all valid input levels. The reference impedance for differential return loss measurements is 100 Ω.

$$Return\_loss(f) \geq \left\{ \begin{array}{ll} 12 - 2\sqrt{f} & 0.01 \leq f < 4.11 \\ 6.3 - 13\log_{10}(f/5.5) & 4.11 \leq f \leq 10 \end{array} \right\} \quad \text{(dB)} \qquad (85\text{–}17)$$

where

f is the frequency in GHz
$Return\_loss(f)$ is the return loss at frequency f

85.8.4.2 Receiver interference tolerance test

The receiver interference tolerance of each lane shall comply with both test 1 and test 2 using the parameters of Table 85–8 when measured according to the requirements of 85.8.4.2.1 to 85.8.4.2.5.

Table 85–8—40GBASE-CR4 and 100GBASE-CR10 interference tolerance parameters

Parameter	Test 1 values	Test 2 values	Units
Maximum BER	10^{-12}	10^{-12}	
Fitted insertion loss coefficients	$a_1 = 2.15$ $a_2 = 0.78$ $a_4 = 0.03$	$a_1 = 6.04$ $a_2 = 0.94$ $a_4 = 0.08$	dB/√GHz dB/GHz dB/GHz^2
Applied SJ[a] (peak-to-peak)	0.115	0.115	UI
Applied RJ[b] (peak-to-peak)	0.13	0.13	UI
Applied DCD (peak-to-peak)	0.035	0.035	UI
Calibrated far-end crosstalk (RMS)	6.3	2.2	mV
Calibrated ICN (RMS) – σ_{nx}	3.7	3.7	mV

[a]Applied SJ frequency >15 MHz, specified at TP0.
[b]Applied random jitter at TP0 is specified at 10^{-12}.

85.8.4.2.1 Test setup

The interference tolerance test is performed with the setup shown in Figure 85–7. The requirements of this subclause are verified at the pattern generator connection (PGC) or test references in Figure 85–7 and Figure 85–8. The lanes under test (LUT) are illustrated in Figure 85–7 and Figure 85–8. The cable assembly receive lanes are terminated in 100 Ω differentially.

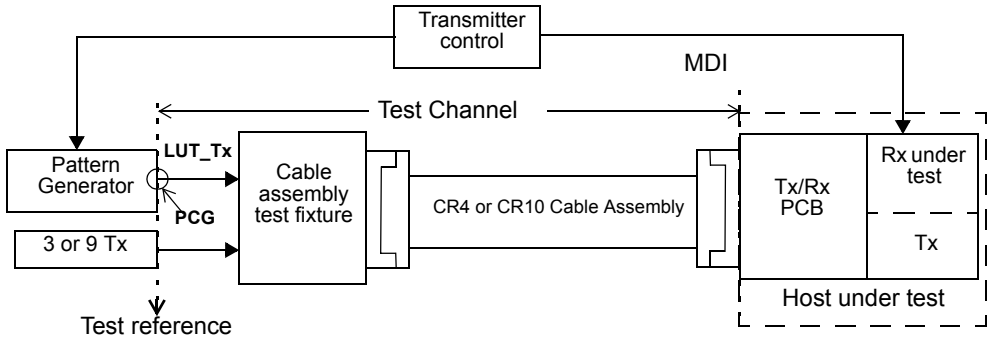

Figure 85–7—Interference tolerance test setup

85.8.4.2.2 Test channel

The test channel consists of the following:

a) A cable assembly
b) A cable assembly test fixture
c) A connecting path from the pattern generator to the cable assembly test fixture

85.8.4.2.3 Test channel calibration

The insertion loss, near-end integrated crosstalk noise, and far-end crosstalk of the test channels are characterized at the test references as illustrated in Figure 85–8 using the cable assembly test fixtures specified in 85.10.8.

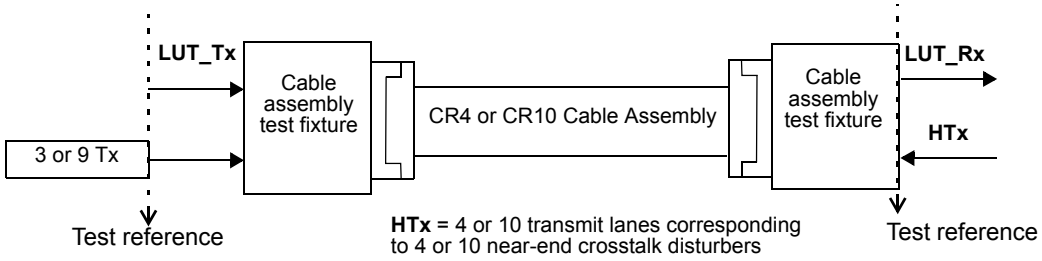

Figure 85–8—Test channel calibration

The fitted insertion loss coefficients of the lane under test (LUT), derived using the fitting procedure in 85.10.2, shall meet the test values in Table 85–8. It is recommended that the deviation between the insertion loss and the fitted insertion loss be as small as practical and that the fitting parameters be as close as practical to the values given in Table 85–8.

The MDNEXT is measured from points HTx to adjacent point LUT_Rx in Figure 85–8. HTx is the set of 4 or 10 transmit lanes of the device under test corresponding to the 4 or 10 near-end crosstalk disturbers with the parameters given in Table 85–11. The RMS value of the integrated MDNEXT crosstalk noise, determined using Equation (85–30) through Equation (85–34), shall meet the test values in Table 85–8.

The far-end crosstalk disturbers consist of 40GBASE-CR4 or 100GBASE-CR10 transmitters. It is recommended that the transition time, equalization setting, and path from the far-end crosstalk disturbers to the cable assembly test fixture emulate the pattern generator as much as practical. For 40GBASE-CR4 test channels, the crosstalk that is coupled into a receive lane is from three transmitters. For 100GBASE-CR10 test channels, the crosstalk that is coupled into a receiver lane is from nine transmitters. The disturber transmitters send either scrambled idle codes or PRBS31. The amplitudes of each of the disturbers should not deviate more than 3 dB from the mean of the disturber amplitudes. The amplitudes of the disturbers should be such that the calibrated far-end crosstalk in Table 85–8 is met in the calibration setup at the LUT point with no signal applied at the PGC, and HTx and PGC terminated in 100 Ω differentially.

85.8.4.2.4 Pattern generator

The pattern generator transmits data to the device under test. Its output amplitude shall be no more than 800 mV peak-to-peak differential when measured on an alternating one zero pattern. The rise and fall times of the pattern generator, as defined in 72.7.1.7, are 47 ps. If the rise and fall times of the pattern generator, T_r, are less than 47 ps the value of a_4 in Table 85–8 is increased by da_4 from Equation (85–18).

$$da_4 = 60.51 \times 10^{-6}(47^2 - T_r^2) \tag{85–18}$$

where T_r is the rise time in ps.

The pattern generator shall meet the jitter specification in Table 85–8. The output waveform of the pattern generator shall comply to 72.7.1.11.

85.8.4.2.5 Test procedure

For 40GBASE-CR4 or 100GBASE-CR10 testing, the pattern generator is first configured to transmit the training pattern defined in 72.6.10.2. During this initialization period, the device under test (DUT) configures the pattern generator equalizer, via transmitter control, to the coefficient settings it would select using the protocol described in 72.6.10 and the receiver will be tuned using its optimization method.

After the pattern generator equalizer has been configured and the receiver tuned, the pattern generator is set to transmit test pattern 3 as defined in 86.8.2. The receiver under test shall meet the target BER listed in Table 85–8. During the tests, the disturbers transmit at their calibrated level and all of the transmitters in the device under test transmit either scrambled idle characters or PRBS31, with the maximum compliant amplitude and equalization turned off (preset condition).

85.8.4.3 Bit error ratio

The receiver shall operate with a BER 10^{-12} or better when receiving a compliant transmit signal, as defined in 85.8.3, through a compliant cable assembly as defined in 85.10.

85.8.4.4 Signaling rate range

A 40GBASE-CR4 and 100GBASE-CR10 receiver shall comply with the requirements of 85.8.4.3 for any signaling rate in the range 10.3125 GBd ± 100 ppm. The corresponding unit interval is nominally 96.969697 ps.

85.8.4.5 AC-coupling

The 40GBASE-CR4 and 100GBASE-CR10 receivers are AC-coupled. AC-coupling shall be part of the receive function for Style-2 40GBASE-CR4 connectors. For Style-1 40GBASE-CR4 and 100GBASE-CR10 plug connectors, the receive lanes are AC-coupled; the coupling capacitors shall be within the plug

connectors. It should be noted that there may be various methods for AC-coupling in actual implementations. The low frequency 3 dB cutoff of the AC-coupling shall be less than 50 kHz.

It is recommended that the value of the coupling capacitors be 100 nF. This will limit the inrush currents and baseline wander.

85.9 Channel characteristics

The 40GBASE-CR4 and 100GBASE-CR10 channel is defined between TP0 and TP5 to include the transmitter and receiver differential controlled impedance printed circuit board and the cable assembly as illustrated in Figure 85–2. The channel parameters insertion loss, insertion loss deviation (ILD), insertion loss to crosstalk ratio, and the transmitter and receiver differential controlled impedance printed circuit boards for each differential lane are provided informatively in 85A.4 through 85A.7.

85.10 Cable assembly characteristics

The 40GBASE-CR4 and 100GBASE-CR10 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 7 m between network ports using controlled impedance cables. All cable assembly measurements are to be made between TP1 and TP4 with cable assembly test fixtures as specified in 85.10.8 and illustrated in Figure 85–13. These cable assembly specifications are based upon twinaxial cable characteristics, but other cable types are acceptable if the specifications of 85.10 are met.

Table 85–9 provides a summary of the cable assembly differential characteristics at 5.15625 GHz and references to the subclauses addressing each parameter.

Table 85–9—Cable assembly differential characteristics summary

Description	Reference	Value	Unit
Maximum insertion loss at 5.15625 GHz	85.10.2	17.04	dB
Minimum insertion loss at 5.15625 GHz		3	dB
Insertion loss deviation at 5.15625 GHz	85.10.3	max = 1.73 min = –1.73	dB
Minimum return loss at 5.15625 GHz	85.10.4	6.66	dB
MDNEXT loss	85.10.5	Equation (85–26)	dB
MDFEXT loss	85.10.6	Equation (85–27)	dB
Maximum integrated crosstalk noise	85.10.7	Equation (85–33)	mV

85.10.1 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.

85.10.2 Cable assembly insertion loss

The fitted cable assembly insertion loss $IL_{fitted}(f)$ as a function of frequency f is defined in Equation (85–19).

$$IL_{\text{fitted}}(f) = a_1\sqrt{f} + a_2 f + a_4 f^2 \quad \text{(dB)} \tag{85–19}$$

where

f is the frequency in GHz
$IL_{\text{fitted}}(f)$ is the fitted cable assembly insertion loss at frequency f

Given the cable assembly insertion loss measured between TP1 and TP4 is at N uniformly-spaced frequencies f_n spanning the frequency range 50 MHz to 7500 MHz with a maximum frequency spacing of 10 MHz, the coefficients of the fitted insertion loss are determined using Equation (85–20) and Equation (85–21).

Define the frequency matrix F as shown in Equation (85–20).

$$F = \begin{bmatrix} \sqrt{f_1} & f_1 & f_1^2 \\ \sqrt{f_2} & f_2 & f_2^2 \\ \cdots & \cdots & \cdots \\ \sqrt{f_N} & f_N & f_N^2 \end{bmatrix} \tag{85–20}$$

The polynomial coefficients a_1, a_2, and a_4 are determined using Equation (85–21). In Equation (85–21), T denotes the matrix transpose operator and IL is a column vector of the measured insertion loss values, IL_n at each frequency f_n.

$$\begin{bmatrix} a_1 \\ a_2 \\ a_4 \end{bmatrix} = (F^T F)^{-1} F^T IL \tag{85–21}$$

The maximum allowed values of the polynomial coefficients a_1, a_2, and a_4 of the fitted cable assembly insertion loss of each pair of the 40GBASE-CR4 and 100GBASE-CR10 in Equation (85–19) and the maximum insertion loss at 5.15625 GHz shall meet the specifications summarized in Table 85–10 unless otherwise noted. The fitted insertion loss corresponding to one example of the maximum insertion loss at 5.15625 GHz and the maximum allowed values of a_1, a_2, and a_4 is illustrated in Figure 85–9.

Table 85–10—Maximum cable assembly insertion loss characteristics

Description	Value	Unit
Maximum insertion loss at 5.15625 GHz	17.04[a]	dB
Maximum fitted insertion loss coefficient a_1	6	dB/√GHz
Maximum fitted insertion loss coefficient a_2	1	dB/GHz
Maximum fitted insertion loss coefficient a_4	0.08	dB/GHz^2

[a]The limit on the maximum insertion loss at 5.15625 GHz precludes the coefficients a_1, a_2, and a_4 from simultaneous maximum values.

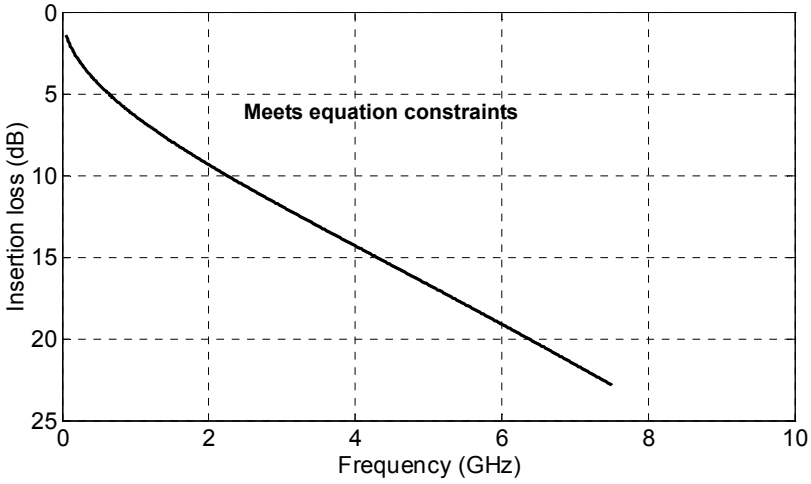

Figure 85–9—Example maximum cable assembly insertion loss

85.10.3 Cable assembly insertion loss deviation (ILD)

The cable assembly insertion loss deviation is the difference between the cable assembly insertion loss and the fitted cable assembly insertion loss determined using Equation (85–22).

$$ILD(f) = IL(f) - IL_{\text{fitted}}(f) \quad (85–22)$$

where

f is the frequency in MHz
$ILD(f)$ is the cable assembly insertion loss deviation at frequency f

The *ILD* shall be within the region defined by Equation (85–23) and Equation (85–24). This includes the insertion loss of the differential cabling pairs and the cable assembly connectors.

$$ILD(f) \geq ILD_{\text{min}}(f) = -0.7 - 0.2\times10^{-3} f \text{ (dB)} \quad (85–23)$$

$$ILD(f) \leq ILD_{\text{max}}(f) = 0.7 + 0.2\times10^{-3} f \text{ (dB)} \quad (85–24)$$

for 50 MHz $\leq f \leq$ 7500 MHz

where

f is the frequency in MHz

The insertion loss deviation limits are illustrated in Figure 85–10.

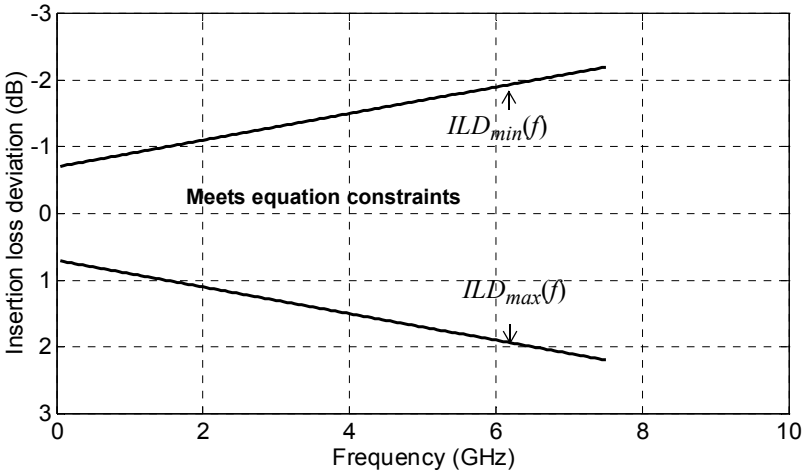

Figure 85–10—Maximum cable assembly insertion loss deviation

85.10.4 Cable assembly return loss

The return loss of each pair of the 40GBASE-CR4 and 100GBASE-CR10 cable assembly shall meet the values determined using Equation (85–25).

$$Return\_loss(f) \geq \begin{cases} 12 - 2\sqrt{f} & 0.05 \leq f < 4.1 \\ 6.3 - 13\log_{10}(f/5.5) & 4.1 \leq f \leq 10 \end{cases} \text{ (dB)} \qquad (85–25)$$

where

f is the frequency in GHz
$Return\_loss(f)$ is the return loss at frequency f

The minimum cable assembly return loss is illustrated in Figure 85–11.

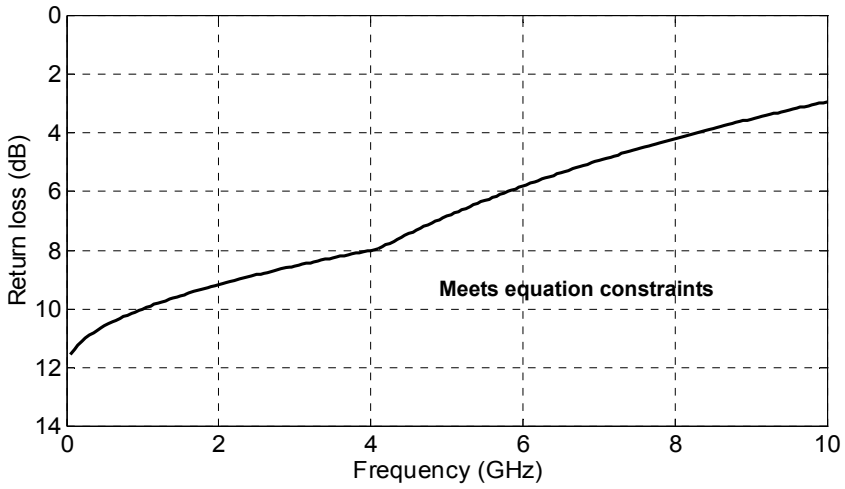

Figure 85–11—Minimum cable assembly return loss

85.10.5 Cable assembly multiple disturber near-end crosstalk (MDNEXT) loss

Since four lanes or ten lanes are used to transfer data between PMDs, the NEXT that is coupled into a receive lane will be from the four or ten transmit lanes. Multiple Disturber Near-End Crosstalk (MDNEXT) loss is determined using the individual NEXT losses.

MDNEXT loss is determined from the four or ten individual pair-to-pair differential NEXT loss values using Equation (85–26).

$$MDNEXT\_loss(f) = -10\log_{10}\left(\sum_{i=0}^{i=3\text{or }9} 10^{-NLi(f)/10}\right)\text{(dB)} \quad (85\text{–}26)$$

for 50 MHz ≤ f ≤ 10000 MHz

where

$MDNEXT_loss(f)$is the MDNEXT loss at frequency f
$NL_i(f)$is the NEXT loss at frequency f of pair combination i, in dB
f is the frequency in MHz
i is the 0 to 3 (pair-to-pair combination) or 0 to 9 (pair-to-pair combination)

85.10.6 Cable assembly multiple disturber far-end crosstalk (MDFEXT) loss

Since four lanes or ten 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 or nine other lanes in the same direction. MDFEXT loss is specified using the individual FEXT losses. MDFEXT loss is determined from the three or nine individual pair-to-pair differential FEXT loss values using Equation (85–27).

$$MDFEXT\_loss(f) = -10\log_{10}\left(\sum_{i=0}^{i=2\text{ or }8} 10^{-NLi(f)/10}\right)\text{(dB)} \quad (85\text{–}27)$$

for 50 MHz ≤ f ≤ 10000 MHz

where

MDFEXT_loss(f) is the MDFEXT loss at frequency f
$NL_i(f)$ is the FEXT loss at frequency f of pair combination i, in dB
f is the frequency in MHz
i is the 0 to 2 (pair-to-pair combination) or 0 to 8 (pair-to-pair combination)

85.10.7 Cable assembly integrated crosstalk noise (ICN)

In order to limit multiple disturber crosstalk noise at a receiver, the cable assembly integrated crosstalk noise (ICN) is specified in relationship to the measured insertion loss. ICN is calculated from the MDFEXT and MDNEXT. Given the multiple disturber near-end crosstalk loss *MDNEXT_loss*(f) and multiple disturber far-end crosstalk loss *MDFEXT_loss*(f) measured over N uniformly-spaced frequencies f_n spanning the frequency range 50 MHz to 10000 MHz with a maximum frequency spacing of 10 MHz, the RMS value of the integrated crosstalk noise shall be determined using Equation (85–28) through Equation (85–32). The RMS crosstalk noise is characterized at the output of a specified receive filter utilizing a specified transmitter waveform and the measured multiple disturber crosstalk transfer functions. The transmitter and receiver filters are defined in Equation (85–28) and Equation (85–29) as weighting functions to the multiple disturber crosstalk in Equation (85–30) and Equation (85–31). The sinc function is defined by $\mathrm{sinc}(x) = \sin(\pi x)/(\pi x)$.

Define the weight at each frequency f_n using Equation (85–28) and Equation (85–29).

$$W_{nt}(f_n) = (A_{nt}^2/f_b)\mathrm{sinc}(f_n/f_b)^2\left[\frac{1}{1+(f_n/f_{nt})^4}\right]\left[\frac{1}{1+(f_n/f_r)^8}\right] \quad (85–28)$$

$$W_{ft}(f_n) = (A_{ft}^2/f_b)\mathrm{sinc}(f_n/f_b)^2\left[\frac{1}{1+(f_n/f_{ft})^4}\right]\left[\frac{1}{1+(f_n/f_r)^8}\right] \quad (85–29)$$

where the equation parameters are given in Table 85–11.

Note that the 3 dB transmit filter bandwidths f_{nt} and f_{ft} are inversely proportional to the 20% to 80% rise and fall times T_{nt} and T_{ft} respectively. The constant of proportionality is 0.2365 (e.g., $T_{nt}f_{nt} = 0.2365$; with f_{nt} in hertz and T_{nt} in seconds). In addition, f_r is the 3 dB reference receiver bandwidth, which is set to 7.5 GHz.

The near-end integrated crosstalk noise σ_{nx} is calculated using Equation (85–30).

$$\sigma_{nx} = \left[2\Delta f\sum_n W_{nt}(f_n)10^{-MDNEXT_{loss}(f_n)/10}\right]^{1/2} \quad (85–30)$$

The far-end integrated crosstalk noise σ_{fx} is calculated using Equation (85–31).

$$\sigma_{fx} = \left[2\Delta f\sum_n W_{ft}(f_n)10^{-MDFEXT_{loss}(f_n)/10}\right]^{1/2} \quad (85–31)$$

where Δf is the uniform frequency step of f_n.

The total integrated crosstalk noise σ_x is calculated using Equation (85–32).

$$\sigma_x = \sqrt{\sigma_{nx}^2 + \sigma_{fx}^2} \quad (85–32)$$

The total integrated crosstalk noise for the cable assembly shall be computed using the parameters shown in Table 85–11.

Table 85–11—Cable assembly integrated crosstalk parameters

Description	Symbol	Value	Units
Symbol rate	f_b	10.3125	GBd
Near-end disturber peak differential output amplitude	A_{nt}	600	mV
Far-end disturber peak differential output amplitude	A_{ft}	600	mV
Near-end disturber 20% to 80% rise and fall times	T_{nt}	24	ps
Far-end disturber 20% to 80% rise and fall times	T_{ft}	24	ps

The total integrated crosstalk RMS noise voltage shall meet the values determined by Equation (85–33) illustrated in Figure 85–12.

$$\sigma_{x,\,\mathrm{ca}} \leq \left\{ \begin{array}{ll} 10 & 3 \leq IL \leq 5.3 \\ 12.4 - 0.45IL & 5.3 < IL \leq 17.04 \end{array} \right\} \quad (\mathrm{mV}) \quad (85–33)$$

where IL is the value of the cable assembly insertion loss in dB at 5.15625 GHz.

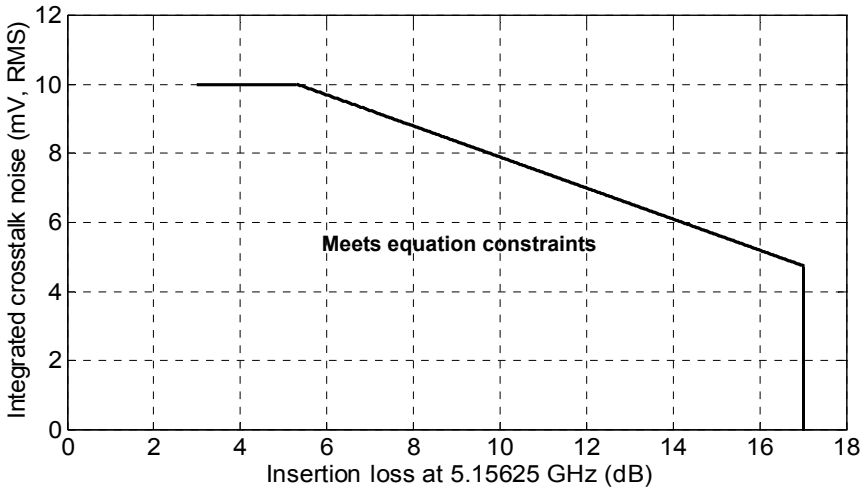

Figure 85–12—Integrated crosstalk noise limits

85.10.8 Cable assembly test fixture

The test fixture of Figure 85–13 or its functional equivalent, is required for measuring the cable assembly specifications in 85.10 at TP1 and TP4. TP1 and TP4 are illustrated in Figure 85–2 and Figure 85–13. The test fixture return loss is equivalent to the test fixture return loss specified in 85.8.3.6. The reference test fixture printed circuit board insertion loss is given in Equation (85–34). The effects of differences between the insertion loss of an actual test fixture and the reference insertion loss are to be accounted for in the measurements.

$$IL_{\text{catf}}(f) = 0.0006 + 0.16\sqrt{f} + 0.0587f \text{ (dB)} \quad (85\text{–}34)$$

for 0.01 GHz ≤ f ≤ 10 GHz

where

f is the frequency in GHz
$IL_{\text{catf}}(f)$ is the reference test fixture printed circuit board insertion loss at frequency f

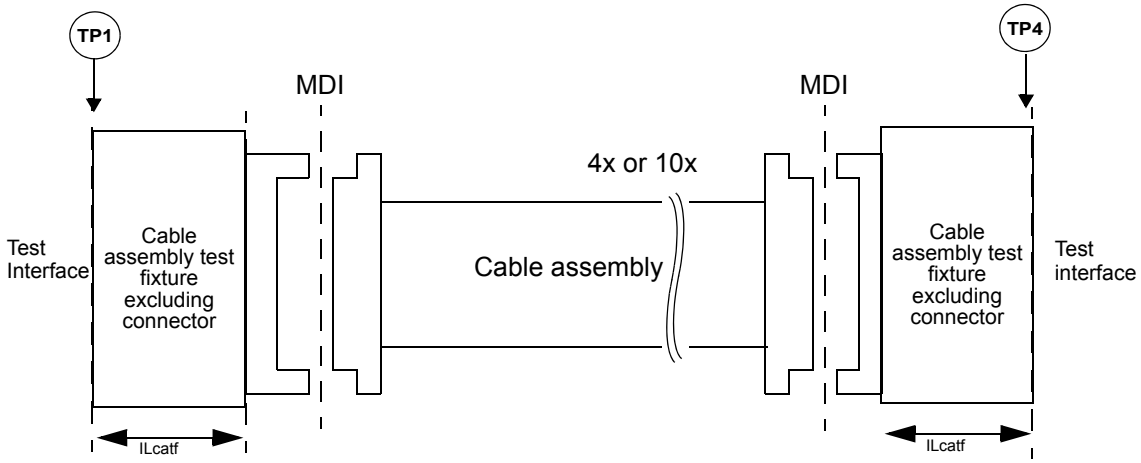

Figure 85–13—Cable assembly test fixtures

85.10.9 Mated test fixtures

The test fixtures of Figure 85–5 and Figure 85–13 are specified in a mated state illustrated in Figure 85–14.

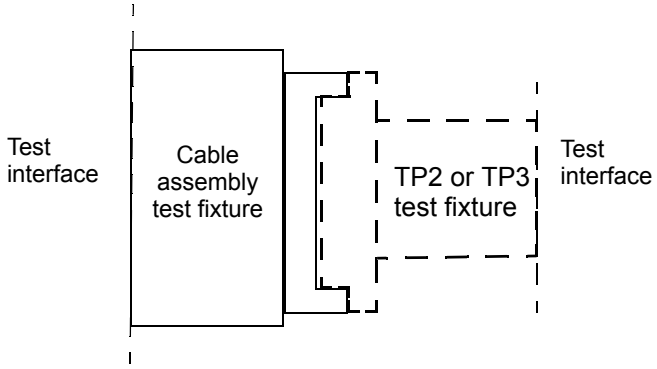

Figure 85–14—Mated test fixtures

85.10.9.1 Mated test fixtures insertion loss

The insertion loss of the mated test fixtures shall meet the values determined using Equation (85–35) and Equation (85–36).

$$IL(f) \geq IL_{MTFmin}(f) = -0.11 + 0.46\sqrt{f} + 0.16f \qquad 0.01 \leq f \leq 10 \qquad \text{(dB)} \qquad (85–35)$$

$$IL(f) \leq IL_{MTFmax}(f) = \left\{ \begin{array}{ll} 0.029 + 0.861\sqrt{f} + 0.158f & 0.01 \leq f < 5.5 \\ 0.2 + 0.65f & 5.5 \leq f \leq 10 \end{array} \right\} \text{(dB)} \qquad (85–36)$$

where

f is the frequency in GHz
$IL(f)$ is the mated test fixture insertion loss at frequency f

The mated test fixtures insertion loss limits are illustrated in Figure 85–15.

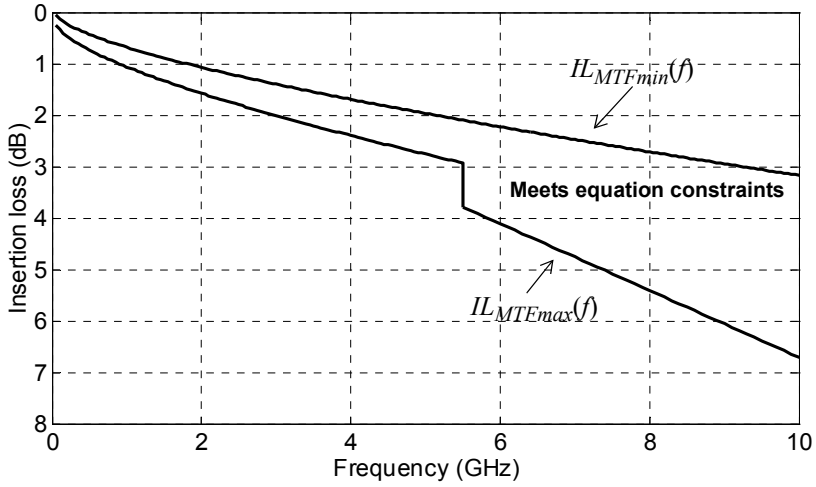

Figure 85–15—Mated test fixtures Insertion loss

85.10.9.2 Mated test fixtures return loss

The return loss of the mated test fixtures measured at either test fixture test interface shall meet the values determined using Equation (85–37).

$$Return\_loss(f) \geq \left\{ \begin{array}{ll} 20 - 2f & 0.01 \leq f < 2.5 \\ 15 & 2.5 \leq f < 5 \\ 13.8 - 28.85\log_{10}(f/5.5) & 5 \leq f \leq 10 \end{array} \right\} \text{(dB)} \qquad (85–37)$$

where

f is the frequency in GHz
$Return_loss(f)$is the return loss at frequency f

The mated test fixtures return loss is illustrated in Figure 85–16.

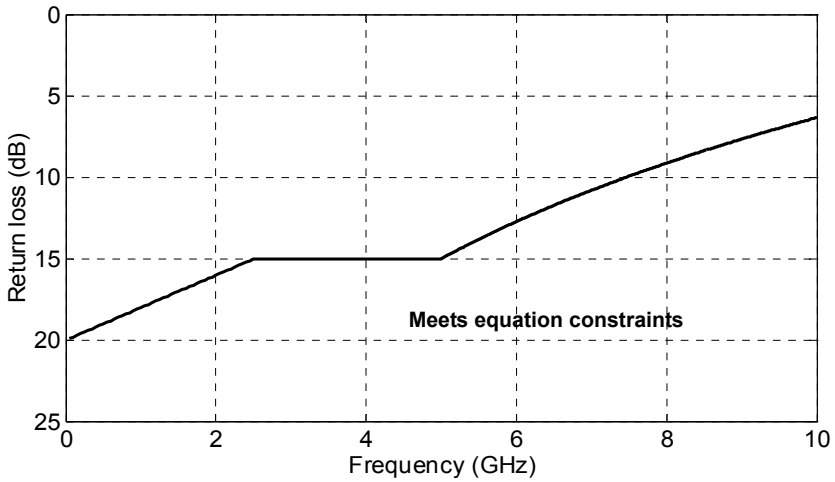

Figure 85–16—Mated test fixtures return loss

85.10.9.3 Mated test fixtures common-mode conversion loss

The common-mode conversion loss of the mated test fixtures measured at either test fixture test interface shall meet the values determined using Equation (85–38).

$$Conversion\_loss(f) \geq \left\{ \begin{array}{ll} 30 - 2.91f & 0.01 \leq f < 5.5 \\ 14 & 5.5 \leq f \leq 10 \end{array} \right\} \text{ (dB)} \qquad (85–38)$$

where

f is the frequency in GHz
$Conversion\_loss(f)$ is the return loss at frequency f

The mated test fixtures common-mode conversion loss is illustrated in Figure 85–17.

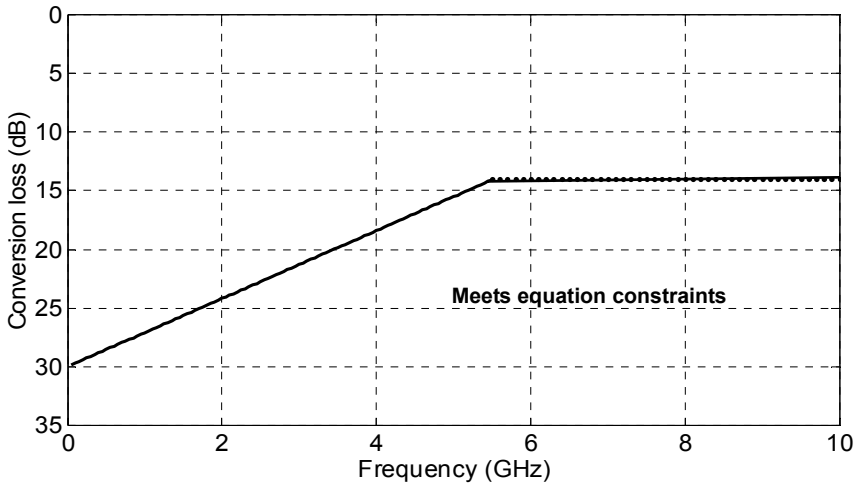

Figure 85–17—Common-mode conversion loss

85.10.9.4 Mated test fixtures integrated crosstalk noise

The mated test fixtures integrated crosstalk RMS noise voltages for the single-disturber near-end crosstalk loss and the single-disturber far-end crosstalk loss are determined using Equation (85–28) through Equation (85–32) by substituting the single disturber near-end for the multiple disturber near-end crosstalk loss and the single disturber far-end crosstalk loss for the multiple disturber far-end crosstalk loss. The values of the mated test fixtures integrated crosstalk RMS noise voltages determined using Equation (85–28) through Equation (85–32) for the single-disturber near-end crosstalk loss, the single-disturber far-end crosstalk loss, the multiple disturber near-end crosstalk loss, and the multiple disturber far-end crosstalk loss shall meet the specifications in Table 85–12.

Table 85–12—Mated test fixtures integrated crosstalk noise

Parameter	40GBASE-CR4	100GBASE-CR10	Units
Near-end integrated crosstalk noise voltage (max, RMS)	0.7	3	mV
Far-end integrated crosstalk noise voltage (max, RMS)	2.5	4	mV
MDNEXT integrated crosstalk noise voltage (max, RMS)	1	3	mV
MDFEXT integrated crosstalk noise voltage (max, RMS)	3.5	5	mV

85.10.10 Shielding

The cable assembly shall provide Class 2 or better shielding in accordance with IEC 61196-1.

85.10.11 Crossover function

The cable assembly shall be wired in a crossover fashion as illustrated in Figure 85–18, with each of the four or ten pairs being attached to the transmitter contacts at one end and the receiver contacts at the other end.

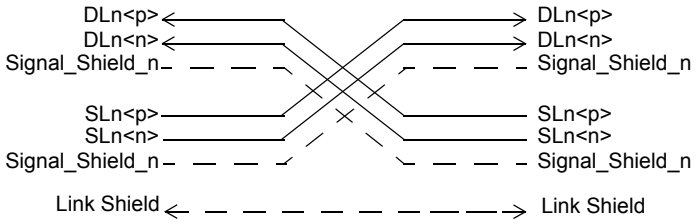

Figure 85–18—Cable assembly wiring

Note that the source lanes (SL) signals SLn<p> and SLn<n> are the positive and negative sides of the transmitters differential signal pairs and the destination lanes (DL) signals DLn<p> and DLn<n> are the positive and negative sides of the receivers differential signal pairs for lane n (n = 0, 1, 2, 3 or n = 0, 1, 2, 3, 4, 5, 6, 7, 8, 9). Signal_Shield_n is the signal shield of the differential signal pair for Lane n.

85.11 MDI specification

This subclause defines the Media Dependent Interface (MDI). The 40GBASE-CR4 and 100GBASE-CR10 PMD, as per 85.7, is coupled to the cable assembly, as per 85.10, by the MDI.

85.11.1 40GBASE-CR4 MDI connectors

Connectors meeting the requirements of 85.11.1.1 (Style-1) or 85.11.1.2 (Style-2) shall be used as the mechanical interface between the PMD of 85.7 and the cable assembly of 85.10. The plug connector shall be used on the cable assembly and the receptacle on the PHY. Style-1 or Style-2 connectors may be used as the MDI interface.

85.11.1.1 Style-1 40GBASE-CR4 MDI connectors

The connector for each end of the cable assembly shall be the quad small form factor pluggable (QSFP+) with the mechanical mating interface defined by SFF-8436 and illustrated in Figure 85–19. The MDI connector shall be the quad small form factor pluggable (QSFP+) receptacle with the mechanical mating interface defined by SFF-8436 and illustrated in Figure 85–20. These connectors have contact assignments matching that in Table 85–13, and electrical performance consistent with the signal quality and electrical requirements of 85.8 and 85.9.

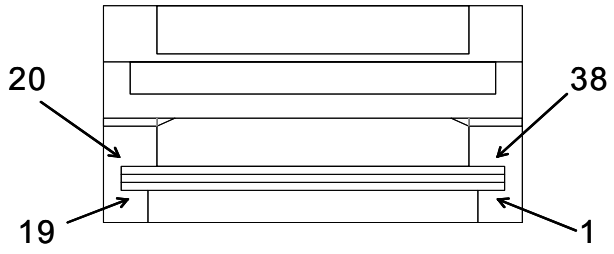

Figure 85–19—Example Style-1 cable assembly plug

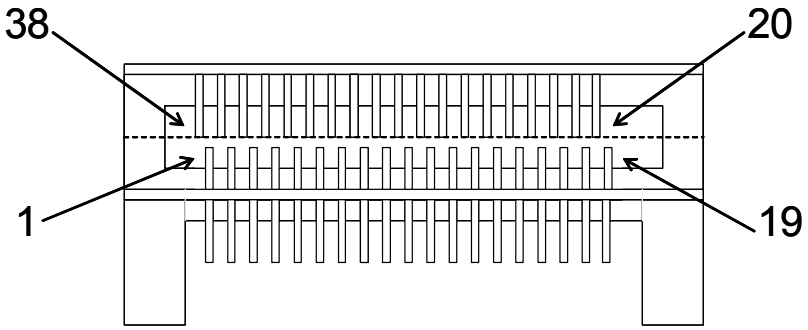

Figure 85–20—Example Style-1 MDI board receptacle

The Style-1 MDI connector of the 40GBASE-CR4 PMD comprises 38 signal connections. The Style-1 40GBASE-CR4 MDI connector contact assignments shall be as defined in Table 85–13.

Table 85–13—Style-1 40GBASE-CR4 lane to MDI connector contact mapping

Tx lane	MDI connector contact	Rx lane	MDI connector contact
signal gnd	S1	signal gnd	S13
SL1<n>	S2	DL2<p>	S14
SL1<p>	S3	DL2<n>	S15
signal gnd	S4	signal gnd	S16
SL3<n>	S5	DL0<p>	S17
SL3<p>	S6	DL0<n>	S18
signal gnd	S7	signal gnd	S19
SL2<p>	S33	DL1<n>	S21
SL2<n>	S34	DL1<p>	S22
signal gnd	S35	signal gnd	S23
SL0<p>	S36	DL3<n>	S24
SL0<n>	S37	DL3<p>	S25
signal gnd	S38	signal gnd	S26

NOTE—Although the 40GBASE-CR4 Style-1 MDI supports 38 connections only the transmitter and receiver contact assignments are specified.

85.11.1.1.1 Style-1 AC-coupling

For Style-1 40GBASE-CR4 plug connectors the receive lanes are AC-coupled; the coupling capacitors are contained within the plug connectors as specified in 85.8.4.5.

85.11.1.2 Style-2 40GBASE-CR4 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 85–21. The MDI connector shall be the latch-type receptacle with the mechanical mating interface defined by IEC 61076-3-113 and illustrated in Figure 85–22. These connectors have a pinout matching that in Table 85–14, and electrical performance consistent with the signal quality and electrical requirements of 85.8 and 85.9. Note that support of compatibility with 10GBASE-CX4 is at the Style-2 40GBASE-CR4 MDI.

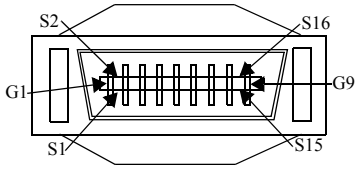

Figure 85–21—Example Style-2 cable assembly plug

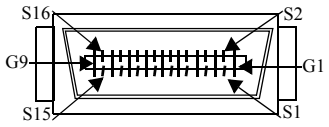

Figure 85–22—Example Style-2 MDI board receptacle

85.11.1.2.1 Style-2 40GBASE-CR4 Connector pin assignments

The MDI connector of the PMD comprises sixteen signal connections, eight signal shield connections, and one link shield connection. The 40GBASE-CR4 MDI connector pin assignments shall be as defined in Table 85–14.

Table 85–14—Style-2 40GBASE-CR4 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

85.11.2 100GBASE-CR10 MDI connectors

The connector for each end of the cable assembly shall be the plug with the mechanical mating interface defined in SFF-8642 and illustrated in Figure 85–23. The MDI connector shall be the receptacle with the mechanical mating interface defined by SFF-8642 and illustrated in Figure 85–24. These connectors have contact assignments matching that in Table 85–15, and electrical performance consistent with the signal quality and electrical requirements of 85.8 and 85.9.

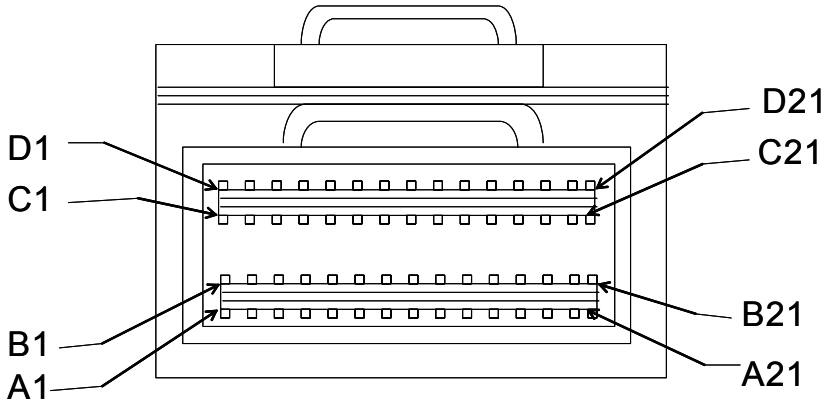

Figure 85–23—Example cable assembly plug

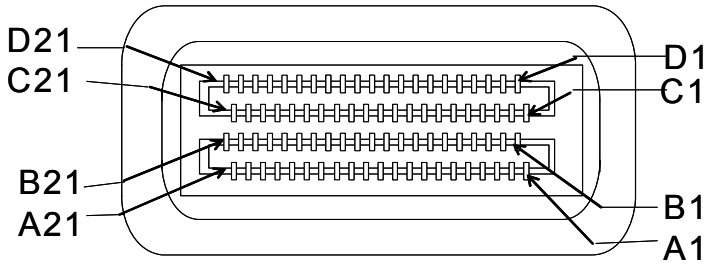

Figure 85–24—Example MDI board receptacle

The MDI connector of the PMD comprises 84 connections. The 100GBASE-CR10 PMD MDI connector contact assignments shall be as defined in Table 85–15.

Table 85–15—100GBASE-CR10 lane to MDI connector contact mapping

Tx lane	MDI connector contact	Tx lane	MDI connector contact	Rx lane	MDI connector contact	Rx lane	MDI connector contact
signal gnd	A1	signal gnd	B1	signal gnd	C1	signal gnd	D1
SL0<p>	A2	—	B2	DL0<p>	C2	—	D2
SL0<n>	A3	—	B3	DL0<n>	C3	—	D3
signal gnd	A4	signal gnd	B4	signal gnd	C4	signal gnd	D4
SL2<p>	A5	SL1<p>	B5	DL2<p>	C5	DL1<p>	D5
SL2<n>	A6	SL1<n>	B6	DL2<n>	C6	DL1<n>	D6
signal gnd	A7	signal gnd	B7	signal gnd	C7	signal gnd	D7
SL4<p>	A8	SL3<p>	B8	DL4<p>	C8	DL3<p>	D8
SL4<n>	A9	SL3<n>	B9	DL4<n>	C9	DL3<n>	D9
signal gnd	A10	signal gnd	B10	signal gnd	C10	signal gnd	D10
SL6<p>	A11	SL5<p>	B11	DL6<p>	C11	DL5<p>	D11
SL6<n>	A12	SL5<n>	B12	DL6<n>	C12	DL5<n>	D12
signal gnd	A13	signal gnd	B13	signal gnd	C13	signal gnd	D13
SL8<p>	A14	SL7<p>	B14	DL8<p>	C14	DL7<p>	D14
SL8<n>	A15	SL7<n>	B15	DL8<n>	C15	DL7<n>	D15
signal gnd	A16	signal gnd	B16	signal gnd	C16	signal gnd	D16
—	A17	SL9<p>	B17	—	C17	DL9<p>	D17
—	A18	SL9<n>	B18	—	C18	DL9<n>	D18
signal gnd	A19	signal gnd	B19	signal gnd	C19	signal gnd	D19

NOTE—Although the 100GBASE-CR10 MDI supports 84 connections only the transmitter and receiver contact assignments are specified.

85.11.2.1 100GBASE-CR10 MDI AC-coupling

For 100GBASE-CR10 plug connectors, the receive lanes are AC-coupled; the coupling capacitors are contained within the plug connectors as specified in 85.8.4.5.

85.11.3 Electronic keying

Electronic keying can be used to enable the detection of Style-1 40GBASE-CR4 MDI connectors or 100GBASE-CR10 MDI cable assembly plugs versus fiber modules or no modules present. Specifications of electronic keying are beyond the scope of this standard.

85.12 Environmental specifications

All equipment subject to this clause shall conform to the applicable requirements of 14.7.

85.13 Protocol implementation conformance statement (PICS) proforma for Clause 85, Physical Medium Dependent (PMD) sublayer and baseband medium, type 40GBASE-CR4 and 100GBASE-CR10[10]

85.13.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 85 Physical Medium Dependent (PMD) sublayer and baseband medium, type 40GBASE-CR4 and 100GBASE-CR10, 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.

85.13.2 Identification

85.13.2.1 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).	

85.13.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 85, Physical Medium Dependent (PMD) sublayer and baseband medium, type 40GBASE-CR4 and 100GBASE-CR10
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.

85.13.3 Major capabilities/options

Item[a]	Feature	Subclause	Value/Comment	Status	Support
XLGMII	XLGMII	85.1	Interface is supported	O	Yes [ ] No [ ]
CGMII	CGMII	85.1	Interface is supported	O	Yes [ ] No [ ]
XLAUI	XLAUI	85.1		O	Yes [ ] No [ ]
CAUI-10	CAUI-10	85.1		O	Yes [ ] No [ ]
CAUI-4	CAUI-4	85.1		O	Yes [ ] No [ ]
CR4	40GBASE-CR4 PMD	85.1	Can operate as 40GBASE-CR4 PMD	O.1	Yes [ ]
CR10	100GBASE-CR10 PMD	85.1	Can operate as 100GBASE-CR10 PMD	O.1	Yes [ ]
PCS	Support of 40GBASE-R PCS	85.1		CR4:M	Yes [ ]
PCS	Support of 100GBASE-R PCS	85.1		CR10:M	Yes [ ]
PMA	Support of 40GBASE-R PMA	85.1		CR4:M	Yes [ ]
PMA	Support of 100GBASE-R PMA	85.1		CR10:M	Yes [ ]
FEC	Forward error correction	85.1	Device implements BASE-R Forward Error Correction	O	Yes [ ] No [ ]
AN	Auto-negotiation	85.1	Device implements Auto-Negotiation	M	Yes [ ]
DC	Delay constraints	85.4	Device conforms to delay constraints specified in 85.4	M	Yes [ ]
DSC	Skew constraints	85.5	Device conforms to Skew and Skew Variation constraints specified in 85.5	M	Yes [ ]
*MD	MDIO capability	85.6	Registers and interface supported	O	Yes [ ] No [ ]
*CBL	Cable assembly	85.10	Items marked with CBL include cable assembly specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]
CA401	40GBASE-CR4 Style-1 cable assembly	85.10	Cable assembly supports 40GBASE-CR4 Style-1	CBL:O.3	Yes [ ] No [ ]
CA402	40GBASE-CR4 Style-2 cable assembly	85.10	Cable assembly supports 40GBASE-CR4 Style-2	CBL:O.3	Yes [ ] No [ ]
CA100	100GBASE-CR10 cable assembly	85.10	Cable assembly supports 100GBASE-CR10	CBL:O.3	Yes [ ] No [ ]

Item[a]	Feature	Subclause	Value/Comment	Status	Support
MDIST1	Style-1 MDI connector	85.11.1.1	40GBASE-CR4 device uses Style-1 MDI	O:2	Yes [ ] N/A [ ]
MDIST2	Style-2 MDI connector	85.11.1.2	40GBASE-CR4 device uses Style-2 MDI	O:2	Yes [ ] N/A [ ]
*LPI	Implementation of LPI with the deep sleep mode option	85.2		O	Yes [ ] No [ ]

[a]A "*" preceding an "Item" identifier indicates there are other PICS that depend on whether or not this item is supported.

85.13.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and baseband medium, type 40GBASE-CR4 and 100GBASE-CR10

85.13.4.1 PMD functional specifications

Item	Feature	Sub clause	Value/Comment	Status	Support
PF1	Transmit function	85.7.2	Converts four or ten logical bit streams into four or ten separate electrical streams	M	Yes [ ]
PF2	Transmit function	85.7.2	Conveys four or ten bits from PMD service interface to MDI lanes	M	Yes [ ]
PF3	Transmitter signal	85.7.2	A positive differential voltage corresponds to tx_bit = one	M	Yes [ ]
PF4	Receive function	85.7.3	Converts four or ten electrical signal streams from the MDI into four or ten logical bit streams	M	Yes [ ]
PF5	Receive function	85.7.3	Conveys four or ten logical bit streams from the MDI lanes to the PMD service interface	M	Yes [ ]
PF6	Receiver signal	85.7.3	A positive differential voltage corresponds to rx_bit = one	M	Yes [ ]
PF7	Global PMD Signal Detect function	85.7.4	Report state via PMD:IS_SIGNAL.indication (SIGNAL_DETECT)	M	Yes [ ]
PF8	Signal_Detect value	85.7.4	Set to FAIL following reset	M	Yes [ ]
PF9	Signal_Detect value	85.7.4	Set to OK on completion of training	M	Yes [ ]
PF10	Signal_Detect value	85.7.4	Set to OK if training disabled by management	M	Yes [ ]
PF11	Global_PMD_transmit_disable	85.7.6	Disables all transmitters by forcing a constant output state	O	Yes [ ] No [ ]
PF12	Global_PMD_transmit_disable	85.7.6	Loopback not affected	O	Yes [ ] No [ ]
PF13	Lane-by-lane_PMD_transmit_disable	85.7.7	Allows each lane transmitter to be selectively disabled	O	Yes [ ] No [ ]

Item	Feature	Sub clause	Value/Comment	Status	Support
PF14	Lane-by-lane_PMD_transmit_disable	85.7.7	Disables transmitters by forcing a constant output state	O	Yes [ ] No [ ]
PF15	Lane-by-lane_PMD_transmit_disable	85.7.7	Loopback not affected	O	Yes [ ] No [ ]
PF16	Loopback	85.7.8	Loopback function provided	M	Yes [ ]
PF17	PMD fault function	85.7.9	Mapped to bit 1.1.7 as listed in Table 85–3	M	Yes [ ]
PF18	PMD control function	85.7.12	Pattern described in 72.6.10 different for each of the lanes	M	Yes [ ]
PF19	Signal detect during LPI	85.7.4	Detect signal energy during LPI	LPI:M	Yes [ ] No [ ]
PF20	Signal detect for EEE	85.7.4	Transition timing to set SIGNAL_-DETECT	LPI:M	Yes [ ] No [ ]
PF21	Transmit disable during LPI	85.7.6	Disable transmitter during tx_mode = QUIET	LPI:M	Yes [ ] No [ ]

85.13.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
MF1	Global_PMD_signal_detect	85.7.4	Set to the value described in 45.2.1.9.7	O	Yes [ ]
MF2	Global_PMD_signal_detect	85.7.4	Set defined by the training state diagram in Figure 72–5	O	Yes [ ] No [ ]
MF3	Lane-by-Lane Signal Detect function	85.7.5	Sets PMD_signal_detect_n values on a lane-by-lane basis per requirements of 85.7.5	MD:M	Yes [ ] N/A [ ]
MF4	PMD_transmit_fault function	85.7.10	Mapped to the PMD_transmit_fault bit as specified in 45.2.1.7.4	MD:M	Yes [ ] N/A [ ]
MF5	PMD_receive_fault function	85.7.11	Contributes to the PMA/PMD receive fault bit as specified in 45.2.1.7.5	MD:M	Yes [ ] N/A [ ]

85.13.4.3 Transmitter specifications

Item	Feature	Subclause	Value/Comment	Status	Support
DS1	Meets specifications at TP2	85.8.3	Unless otherwise noted per Table 85–5	M	Yes [ ]
DS2	Test load	85.8.3.6	100 Ω differential load with return loss using Equation (85–1)	M	Yes [ ]
DS3	Test fixture return loss	85.8.3.6	Equation (85–37)	M	Yes [ ]
DS4	Test fixture insertion loss	85.8.3.7	Per Equation (85–15)	M	Yes [ ]
DS5	Signaling rate, per lane	85.8.3.9	10.3125 GBd ± 100 ppm	M	Yes [ ]
DS6	Output Amplitude LPI voltage	85.7.6	Less than 30 mV within 500 ns of tx_quiet	LPI:M	Yes [ ] No [ ]
DS7	Output Amplitude ON voltage	85.7.6	Greater than 90% of previous level within 500 ns of tx_quiet deasserted	LPI:M	Yes [ ] No [ ]

85.13.4.4 Receiver specifications

Item	Feature	Subclause	Value/Comment	Status	Support
RS1	Receiver tolerance	85.8.4.2	BER of better than 10^{-12}	M	Yes [ ]
RS2	Receiver tolerance	85.8.4.2	Maximum fitted insertion loss coefficients	M	Yes [ ]
RS3	Receiver tolerance	85.8.4.2	MDNEXT crosstalk noise	M	Yes [ ]
RS4	Receiver tolerance	85.8.4.2	Pattern generator output amplitude	M	Yes [ ]
RS5	Receiver tolerance	85.8.4.2	Pattern generator jitter specification	M	Yes [ ]
RS6	Bit Error Ratio	85.8.4.2.5	BER of better than 10^{-12}	M	Yes [ ]
RS7	Meets specifications at TP3	85.8.4	Unless otherwise noted per Table 85–7	M	Yes [ ]
RS8	Signaling rate, per lane	85.8.4.4	10.3125 GBd ± 100 ppm	M	Yes [ ]
RS9	AC-coupling	85.8.4.5	3 dB cutoff	M	Yes [ ]

85.13.4.5 Cable assembly specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CA1	Differential reference impedance	85.10.1	100 Ω	CBL:M	Yes [ ] N/A [ ]
CA2	Insertion loss	85.10.2	Per Equation (85–19) and Table 85–10	CBL:M	Yes [ ] N/A [ ]
CA3	Insertion loss deviation	85.10.3	Per Equation (85–23) and Equation (85–24)	CBL:M	Yes [ ] N/A [ ]
CA4	Return loss	85.10.4	Per Equation (85–25)	CBL:M	Yes [ ] N/A [ ]
CA5	Integrated crosstalk noise	85.10.7	Integrated crosstalk noise voltage per Equation (85–28) through Equation (85–33) and Table 85–11	CBL:M	Yes [ ] N/A [ ]
CA6	Cable assembly test fixture insertion loss	85.10.8	Per Equation (85–34)	CBL:M	Yes [ ] N/A [ ]
CA7	Mated test fixture insertion loss	85.10.9.1	Per Equation (85–35) and Equation (85–36)	CBL:M	Yes [ ] N/A [ ]
CA8	Mated test fixture return loss	85.10.9.2	Per Equation (85–38)	CBL:M	Yes [ ] N/A [ ]
CA9	Mated test fixtures integrated crosstalk noise	85.10.9.4	Per Equation (85–28), through Equation (85–32) and Table 85–12	CBL:M	Yes [ ] N/A [ ]
CA10	Shielding	85.10.10	Class 2 or better in accordance with IEC 61196-1	CBL:M	Yes [ ] N/A [ ]
CA11	Crossover function	85.10.11	Per Figure 85–18	CBL:M	Yes [ ] N/A [ ]
CA12	Cable assembly connector type	85.11.1	40GBASE-CR4 Style-1 plug (SFF-8436 plug)	CA401:M	Yes [ ] N/A [ ]
CA13	Pin assignments	85.11.1.1	Per Table 85–13		Yes [ ] N/A [ ]
CA14	Cable assembly connector type	85.11.1	IEC 61076-3-113 latch-type plug	CA402:M	Yes [ ] N/A [ ]
CA15	Pin assignments	85.11.1.2.1	Per Table 85–14		Yes [ ] N/A [ ]
CA16	Cable assembly connector type	85.11.2	100GBASE-CR10 plug (SFF-8642 plug)	CA100:M	Yes [ ] N/A [ ]
CA17	Pin assignments	85.11.2	Per Table 85–15		Yes [ ] N/A [ ]
CA18	AC-coupling	85.8.4.5	3 dB cutoff	M	Yes [ ]

85.13.4.6 MDI connector specifications

Item	Feature	Subclause	Value/Comment	Status	Support
MDC1	MDI connector type	85.11.1.1	40GBASE-CR4 Style-1 receptacle (SFF-8436 receptacle)	CR4*MDIST1:M	Yes [ ] N/A [ ]
MDC2	MDI connector type	85.11.1.2	IEC 61076-3-113 latch-type receptacle	CR4*MDIST2:M	Yes [ ]
MDC3	MDI connector type	85.11.2.1	100GBASE-CR10 receptacle (SFF-8642 receptacle)	CR10:M	Yes [ ] N/A [ ]

85.13.4.7 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	Environmental specifications	85.12		M	Yes [ ]

86. Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE–SR4 and 100GBASE–SR10

86.1 Overview

This clause specifies the 40GBASE–SR4 PMD and the 100GBASE–SR10 PMD together with the multimode fiber medium. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 86–1, to the medium through the MDI, and optionally to the management functions that are accessible through the management interface defined in Clause 45, or equivalent.

Table 86–1—Physical Layer clauses associated with the 40GBASE-SR4 and 100GBASE-SR10 PMDs

Associated clause	40GBASE–SR4	100GBASE–SR10
81—RS	Required	Required
81—XLGMII[a]	Optional	Not applicable
81—CGMII[a]	Not applicable	Optional
82—PCS for 40GBASE-R	Required	Not applicable
82—PCS for 100GBASE-R	Not applicable	Required
83—PMA for 40GBASE-R4	Required	Not applicable
83—PMA for 100GBASE-R10	Not applicable	Required
83A—XLAUI[b]	Optional	Not applicable
83A—CAUI-10[b]	Not applicable	Optional
83B—Chip to module XLAUI[b]	Optional	Not applicable
83B—Chip to module CAUI-10[b]	Not applicable	Optional
83D—CAUI-4	Not applicable	Optional
83E—Chip-to-module CAUI-4	Not applicable	Optional
86A—XLPPI	Optional	Not applicable
86A—CPPI	Not applicable	Optional
78—Energy Efficient Ethernet	Optional	Optional

[a] XLGMII and CGMII are optional interfaces. However, if the appropriate interface is not implemented, a conforming implementation must behave functionally as though the RS, and XLGMII or CGMII, were present.
[b] If XLAUI or CAUI-n is present, there is at least a PMA between the XLAUI or CAUI-n and the PMD.

Figure 86–1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 40 Gb/s and 100 Gb/s Ethernet is introduced

in Clause 80 and the purpose of each PHY sublayer is summarized in 80.2. Further relevant information may be found in Clause 1 (terminology and conventions, references, definitions and abbreviations) and Annex A (bibliography, referenced as [B1], [B2], etc.).

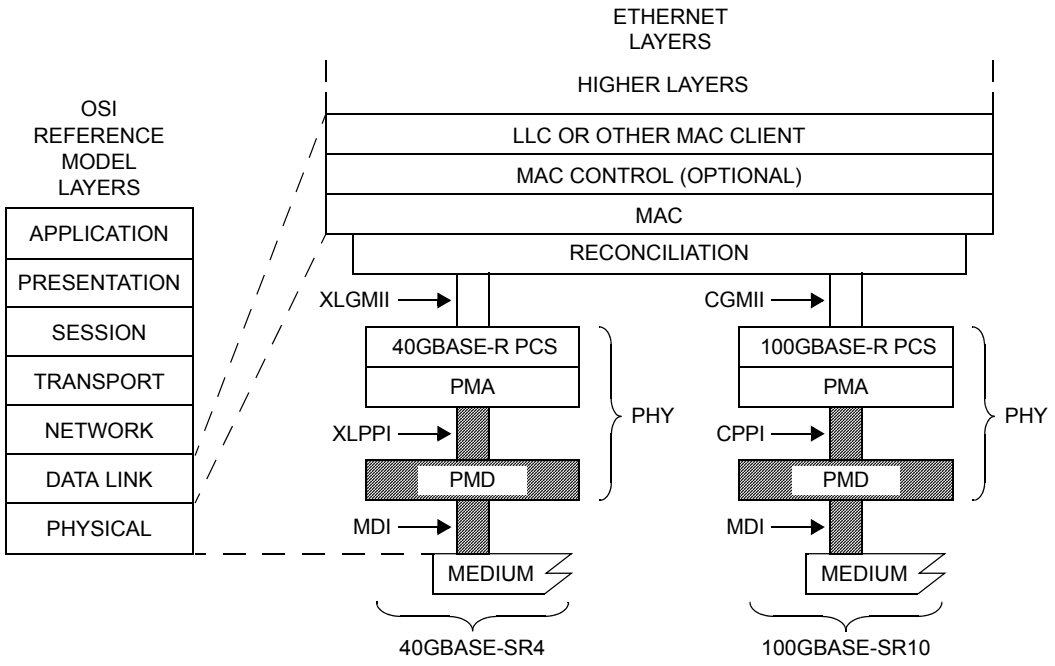

CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
CPPI = 100 Gb/s PARALLEL PHYSICAL INTERFACE
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE
XLPPI = 40 Gb/s PARALLEL PHYSICAL INTERFACE
SR = PMD FOR MULTIMODE FIBER

Figure 86–1—40GBASE-SR4 and 100GBASE-SR10 PMDs relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

The 40GBASE-SR4 and 100GBASE-SR10 PMD sublayers provide point-to-point 40 Gb/s and 100 Gb/s Ethernet links over four or ten pairs of multimode fiber, up to at least 100 m. Table 86–2 shows the primary attributes of each PMD type.

40GBASE–SR4 uses four identical lanes, while 100GBASE–SR10 uses ten of the same lanes. In this clause, where there are four or ten items (depending on PMD type) such as lanes, the items are numbered from 0 to $n - 1$, and an example item is numbered i. Thus n is 4 or 10.

The connection to the PMA may use the optional physical instantiation of the PMD service interface called XLPPI (four lanes, for 40GBASE-SR4) or CPPI (ten lanes, for 100GBASE-SR10). The term “nPPI” is used to denote either XLPPI or CPPI, or both.

This clause is arranged as follows: following the overview and an abstract description of the PMD service interface, delay and Skew specifications, control and status variables and registers, a block diagram and high-level specification of the PMD functions, and lane assignments, 86.7 (four parts) contains the optical specifications for 40GBASE-SR4 and 100GBASE-SR10. 86.8 defines test points and optical and dual-use parameters. 86.9 addresses safety, installation, environment and labeling, 86.10 defines the optical channel, and 86.11 contains the PICS. Annex 86A contains the electrical specifications for XLPPI and CPPI,

Table 86–2—Summary of 40GBASE–SR4 and 100GBASE–SR10

PMD Type	40GBASE–SR4	100GBASE–SR10	Unit
Fiber type	50/125 μm multimode, type A1a.2[a] (OM3) or A1a.3[b](OM4)		
Number of fiber pairs	4	10	
Nominal wavelength	850		nm
Required operating range	0.5 to 100 for OM3		m
	0.5 to 150 for OM4[c]		
Signaling rate, each lane	10.3125 ±100 ppm		GBd

[a] Type A1a.2 (OM3) specified in IEC 60793-2-10. See 86.10.2.1.
[b] Type A1a.3 (OM4) specified in IEC 60793-2-10. See 86.10.2.1.
[c] This is an engineered link with maximum 1 dB connection and splice loss.

electrical compliance boards, definitions of electrical parameters, and a recommended PCB response for the host (PMA).

40GBASE-SR4 and 100GBASE-SR10 PHYs with the optional Energy Efficient Ethernet (EEE) capability may enter the fast wake Low Power Idle (LPI) mode to conserve energy during periods of low link utilization (see Clause 78). The deep sleep mode of EEE is not supported.

86.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 40GBASE-SR4 and 100GBASE-SR10 PMDs. The service interfaces for these PMDs are described in an abstract manner and do not imply any particular implementation, although an optional implementation of the PMD service interface, the Parallel Physical Interface (nPPI), is specified in 86A.4.1 and 86A.4.2. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_*i*.request
PMD:IS_UNITDATA_*i*.indication
PMD:IS_SIGNAL.indication

The 40GBASE-SR4 PMD has four parallel bit streams, hence $i = 0$ to 3 for 40GBASE-SR4 and the 100GBASE-SR10 PMD has ten parallel bit streams, hence $i = 0$ to 9 for 100GBASE-SR10.

The PMA (or the PMD) continuously sends four or ten parallel bit streams to the PMD (or the PMA), one per lane, each at a nominal signaling rate of 10.3125 GBd (see 83.4).

Upon receipt of the PMD:IS_UNITDATA_*i*.request primitive, the PMD converts the specified streams of bits into the appropriate signals on the MDI.

The PMD:IS_UNITDATA_*i*.indication primitive corresponds to one of the signals received from the MDI. This primitive is received by the client (the PMA), as described in 83.4.

The PMD:IS_SIGNAL.indication(SIGNAL_DETECT) primitive is generated by the PMD to report the parameter SIGNAL_DETECT, which indicates the status of the signals being received from the MDI (see 86.5.4). There is one parameter and primitive, reporting on all the lanes as a group. This primitive is received by the client (the PMA) as PMD:IS_SIGNAL.indication(SIGNAL_OK), as described in 83.4.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_bit parameters are 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.

86.3 Delay and Skew

86.3.1 Delay constraints

The sum of the transmit and receive delays at one end of the link contributed by the 40GBASE-SR4 PMD including 2 m of fiber in one direction shall be no more than 1024 bit times (2 pause_quanta or 25.6 ns). The sum of the transmit and receive delays at one end of the link contributed by the 100GBASE-SR10 PMD including 2 m of fiber in one direction shall be no more than 2048 bit times (4 pause_quanta or 20.48 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

86.3.2 Skew and Skew Variation constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the PCS. The Skew Variation must also be limited to ensure that a given PCS lane always traverses the same physical lane. Skew and Skew Variation are defined in 80.5 and specified at the points SP1 to SP6 shown in Figure 80–6 and Figure 80–7. Skew points as they relate to the nPPI are shown in Figure 86–3.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 80.5. The measurements of Skew and Skew Variation are defined in 86.8.3.1.

86.4 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If MDIO is implemented, the mapping of MDIO control variables to PMD control variables shall be as shown in Table 86–3, and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 86–4.

Table 86–3—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable register	1.9.0	PMD_global_transmit_disable
PMD transmit disable 9[a]	PMD transmit disable register	1.9.10	PMD_transmit_disable_9
PMD transmit disable 8[a] to PMD transmit disable 0	PMD transmit disable register	1.9.9 to 1.9.1	PMD_transmit_disable_8 to PMD_transmit_disable_0

[a] For 40GBASE-SR4, the highest-numbered six of the ten lane-by-lane transmit disables do not apply.

Table 86–4—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	PMD_global_signal_detect
PMD receive signal detect 9[a]	PMD receive signal detect register	1.10.10	PMD_signal_detect_9
PMD receive signal detect 8[a] to PMD receive signal detect 0	PMD receive signal detect register	1.10.9 to 1.10.1	PMD_signal_detect_8 to PMD_signal_detect_0

[a] For 40GBASE-SR4, the highest-numbered six of the ten lane-by-lane signal detects do not apply.

86.5 PMD functional specifications

The 40GBASE–SR4 and 100GBASE–SR10 PMDs perform the Transmit and Receive functions, which convey data between the PMD service interface and the MDI.

86.5.1 PMD block diagram

The PMD block diagram is shown in Figure 86–2. Figure 86–3 shows the test points. It is not required that the PMD service interface be exposed or measurable (nPPI as defined in Annex 86A with compliance points TP1, TP1a, TP4, TP4a). However, if it is not, a conforming implementation must behave as though the interface were compliant.

For purposes of system conformance, the PMD sublayer is standardized at the test points described in 86.8.1. The transmit side electrical signal (PMA output and PMD electrical input) is defined at TP1 and TP1a; see Annex 86A. The optical transmit signal is defined at the output end of a 50 μm multimode fiber patch cord (TP2), between 2 m and 5 m in length. Unless specified otherwise, all optical transmitter measurements and tests defined in 86.8 are made at TP2. The optical receive signal is defined at the output

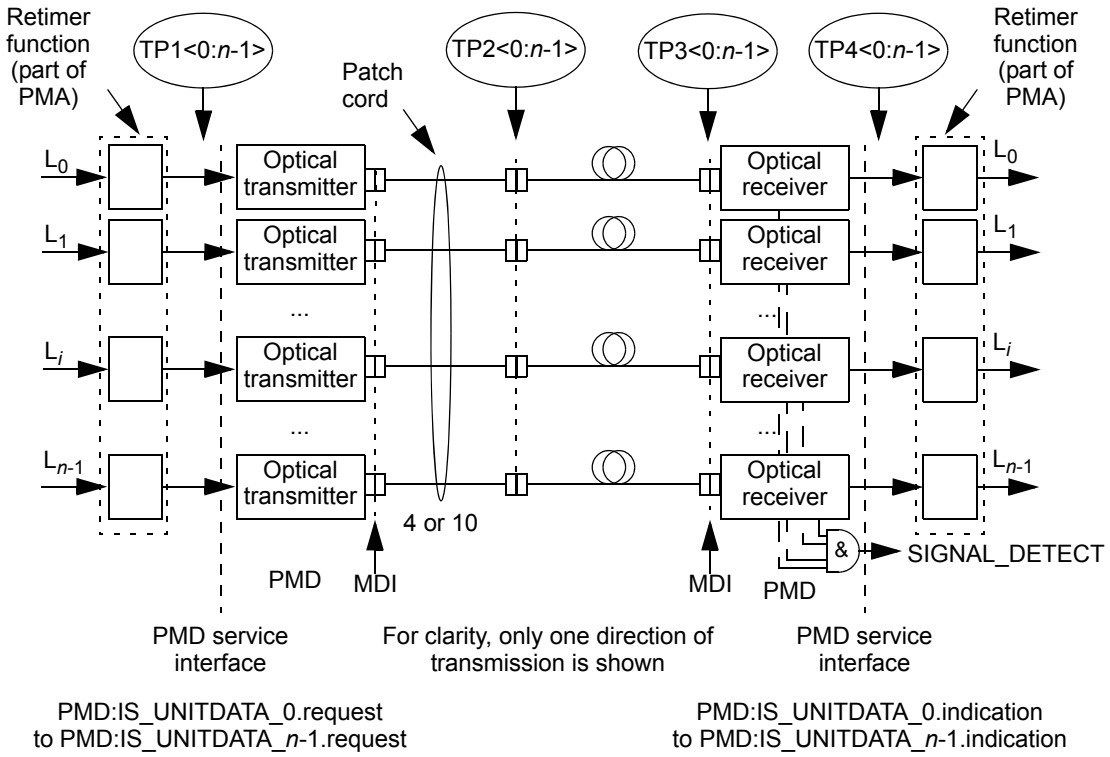

Figure 86–2—Block diagram for 40GBASE–SR4 and 100GBASE–SR10 transmit/receive paths

of the fiber optic cabling (TP3) at the MDI (see 86.10.3). Unless specified otherwise, all optical receiver measurements and tests defined in 86.8 are made at TP3. The receive side electrical signal (PMD electrical output and PMA input) is defined at TP4 and TP4a; see Annex 86A.

86.5.2 PMD transmit function

The PMD Transmit function shall convert the four or ten electronic bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_*n*-1.request into the same number of optical signal streams. The optical signal streams are delivered to the MDI, which contains four or ten parallel light paths for transmit, according to the transmit optical specifications in this clause. The higher optical power level in each signal stream shall correspond to tx_bit = one.

86.5.3 PMD receive function

The PMD Receive function shall convert the four or ten parallel optical signal streams received from the MDI into separate electronic bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_*n*-1.indication, all according to the receive optical specifications in this clause. The higher optical power level in each signal stream shall correspond to rx_bit = one.

86.5.4 PMD global signal detect function

The PMD global 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:IS_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 lanes. The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 86–5. The PMD receiver is not required to verify whether a compliant 40GBASE-SR4 or 100GBASE-SR10 signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

Table 86–5—SIGNAL_DETECT value definition

Receive conditions	SIGNAL_DETECT value
For any lane; Average optical power at TP3 ≤ –30 dBm	FAIL
For all lanes; [(Optical power at TP3 ≥ Minimum OMA, each lane, in Table 86–7) and (compliant 40GBASE–SR4 or 100GBASE–SR10 signal input as appropriate)]	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 including the effects of 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.

86.5.5 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_*i*, where *i* represents the lane number in the range 0:*n*-1, shall be continuously set in response to the optical signal on its associated lane, according to the requirements of Table 86–5.

86.5.6 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.

86.5.7 PMD global transmit disable function (optional)

The PMD_global_transmit_disable function is optional and allows all of the optical transmitters to be disabled.

a) When the PMD_global_transmit_disable variable is set to one, this function shall turn off all of the optical transmitters so that each transmitter meets the requirements of the average launch power of the OFF transmitter in Table 86–6.
b) If a PMD_fault is detected, then the PMD may set the PMD_global_transmit_disable to one, turning off the optical transmitter in each lane.

86.5.8 PMD lane-by-lane transmit disable function (optional)

The PMD_transmit_disable_*i* function (where *i* represents the lane number in the range 0:*n*-1) is optional and allows the optical transmitter in each lane to be selectively disabled.

a) When a PMD_transmit_disable_*i* 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 86–6.
b) If a PMD_fault is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the optical transmitter in each lane.

If the optional PMD_transmit_disable_*i* function is not implemented in MDIO, an alternative method may be provided to independently disable each transmit lane.

86.5.9 PMD fault function (optional)

If the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set PMD_fault to one. If the MDIO interface is implemented, PMD_fault shall be mapped to the PMA/PMD fault bit as specified in 45.2.1.2.3.

86.5.10 PMD transmit fault function (optional)

If the PMD has detected a local fault on any transmit lane, the PMD shall set the PMD_transmit_fault variable to one. If the MDIO interface is implemented, PMD_transmit_fault shall be mapped to the PMA/PMD transmit fault bit as specified in 45.2.1.7.4.

86.5.11 PMD receive fault function (optional)

If the PMD has detected a local fault on any receive lane, the PMD shall set the PMD_receive_fault variable to one. If the MDIO interface is implemented, PMD_receive_fault shall be mapped to the PMA/PMD receive fault bit as specified in 45.2.1.7.5.

86.6 Lane assignments

There are no lane assignments (within a group of transmit or receive lanes) for 40GBASE-SR4 or 100GBASE-SR10. While it is expected that a PMD will map electrical lane *i* to optical lane *i* and vice versa, there is no need to define the physical ordering of the lanes, as the PCS is capable of receiving the lanes in any arrangement. The positioning of transmit and receive lanes at the MDI is specified in 86.10.3.

86.7 PMD to MDI specifications for 40GBASE-SR4 or 100GBASE-SR10

The required operating range for the 40GBASE–SR4 and 100GBASE–SR10 PMD is defined in Table 86–2. A compliant PMD operates on 50/125 μm multimode fibers according to the specifications of Table 86–14. A PMD which exceeds the operating range requirement while meeting all other optical specifications is considered compliant (e.g., operating at 125 m meets the operating range requirement of 0.5 m to 100 m). The signaling rate for a lane of a 40GBASE–SR4 or 100GBASE–SR10 PMD shall be as defined in Table 86–2. The optical signal at the transmit and receive side of the MDI is specified in 86.7.1 and 86.7.3. The range of optical signals within the optical medium is defined in 86.7.2, and an illustrative link power budget is provided in 86.7.4. Test points are defined in 86.8.1.

86.7.1 Transmitter optical specifications

Each lane of a 40GBASE–SR4 or 100GBASE–SR10 optical transmitter shall meet the specifications of Table 86–6 per the definitions in 86.8.

Table 86–6—40GBASE–SR4 or 100GBASE–SR10 optical transmit characteristics

Description	Type	Value	Unit
Center wavelength	Range	840 to 860	nm
RMS spectral width[a]	Max	0.65	nm
Average launch power, each lane	Max	2.4	dBm
Average launch power, each lane	Min	–7.6	dBm
Optical Modulation Amplitude (OMA), each lane	Max	3	dBm
Optical Modulation Amplitude (OMA), each lane	Min	–5.6[b]	dBm
Difference in launch power between any two lanes (OMA)	Max	4	dB
Peak power, each lane	Max	4	dBm
Launch power in OMA minus TDP, each lane	Min	–6.5	dBm
Transmitter and dispersion penalty (TDP), each lane	Max	3.5	dB
Extinction ratio	Min	3	dB
Optical return loss tolerance	Max	12	dB
Encircled flux[c]		≥ 86% at 19 μm, ≤ 30% at 4.5 μm	
Transmitter eye mask definition {X1, X2, X3, Y1, Y2, Y3} Hit ratio 5×10^{-5} hits per sample	Spec values	0.23, 0.34, 0.43, 0.27, 0.35, 0.4	
Average launch power of OFF transmitter, each lane	Max	–30	dBm

[a] RMS spectral width is the standard deviation of the spectrum.
[b] Even if the TDP < 0.9 dB, the OMA (min) must exceed this value.
[c] If measured into type A1a.2 or type A1a.3 50 μm fiber in accordance with IEC 61280-1-4.

86.7.2 Characteristics of signal within, and at the receiving end of, a compliant optical channel

Table 86–7 gives the characteristics of a signal within, and at the receiving end of, a lane of a compliant 40GBASE–SR4 or 100GBASE–SR10 optical channel, and the aggregate signal. A signal with power in OMA or average power not within the ranges given cannot be compliant. However, a signal with power values within the ranges is not necessarily compliant.

Table 86–7—Characteristics of signal within, and at the receiving end of, a compliant optical channel

Description	Minimum		Maximum	Unit
Fiber type	OM3	OM4		
Total average power for 40GBASE–SR4	–3.5	–3.1	+8.4	dBm
Total average power for 100GBASE–SR10	+0.5	+0.9	+12.4	dBm
Average power, each lane	–9.5	–9.1	+2.4	dBm
Optical Modulation Amplitude (OMA), each lane	–7.5	–7.1	+3	dBm

NOTE—Table 86–7 provides information for diagnostic purposes that is needed by network operators in maintenance. There is no need to assure compliance to it in normal circumstances.

86.7.3 40GBASE–SR4 or 100GBASE–SR10 receiver optical specifications

Each lane of a 40GBASE–SR4 or 100GBASE–SR10 optical receiver shall meet the specifications defined in Table 86–8 per the definitions in 86.8.

Table 86–8—40GBASE–SR4 or 100GBASE–SR10 optical receiver characteristics

Description	Type	Value	Unit
Center wavelength, each lane	Range	840 to 860	nm
Damage threshold[a]	Min	+3.4	dBm
Average power at receiver input, each lane	Max	+2.4	dBm
	Min	–9.5	dBm
Receiver reflectance	Max	–12	dB
Optical Modulation Amplitude (OMA), each lane	Max	3	dBm
Stressed receiver sensitivity in OMA, each lane[b]	Max	–5.4	dBm
Peak power, each lane	Max	4	dBm
Conditions of stressed receiver sensitivity test:			
Vertical eye closure penalty (VECP)[c], each lane	—	1.9	dB
Stressed eye J2 Jitter[c], each lane	—	0.3	UI
Stressed eye J9 Jitter[c], each lane	—	0.47	UI
OMA of each aggressor lane	—	–0.4	dBm
Receiver jitter tolerance in OMA, each lane[d]	Max	–5.4	dBm
Conditions of receiver jitter tolerance test:			
Jitter frequency and peak-to-peak amplitude	—	(75, 5)	(kHz, UI)
Jitter frequency and peak-to-peak amplitude	—	(375, 1)	(kHz, UI)
OMA of each aggressor lane	—	–0.4	dBm

[a] The receiver shall be able to tolerate, without damage, continuous exposure to a modulated optical input signal having this power level on one lane. The receiver does not have to operate correctly at this input power.
[b]Measured with conformance test signal at TP3 (see 86.8.4.7).
[c]Vertical eye closure penalty and stressed eye jitter are test conditions for measuring stressed receiver sensitivity. They are not characteristics of the receiver. The apparent discrepancy between VECP and TDP is because VECP is defined at eye center while TDP is defined with ±0.15 UI offsets of the sampling instant.
[d]This is a test of the optical receiver's ability to track low-frequency jitter and is inappropriate for any subsystem that does not include a CRU.

86.7.4 40GBASE–SR4 or 100GBASE–SR10 illustrative link power budget

Illustrative power budgets and penalties for 40GBASE–SR4 or 100GBASE–SR10 optical channels are shown in Table 86–9.

Table 86–9—40GBASE–SR4 or 100GBASE–SR10 illustrative link power budgets

Parameter	OM3	OM4	Unit
Effective modal bandwidth at 850 nm[a]	2000	4700	MHz•km
Power budget (for maximum TDP)	8.3		dB
Operating distance	0.5 to 100	0.5 to 150	m
Channel insertion loss[b]	1.9	1.5	dB
Allocation for penalties (for maximum TDP)[c]	6.4	6.5	dB
Unallocated margin	0	0.3[d]	dB
Additional insertion loss allowed	0		dB

[a] Per IEC 60793-2-10.
[b] The channel insertion loss is calculated using the maximum distances specified in Table 86–2 and cabled optical fiber attenuation of 3.5 dB/km at 850 nm plus an allocation for connection and splice loss given in 86.10.2.2.1.
[c] Link penalties are used for link budget calculations. They are not requirements and are not meant to be tested.
[d] This unallocated margin is not available for use.

86.8 Definitions of optical and dual-use parameters and measurement methods

In this section, the test points and the parameters applicable to optical signals and for dual use (both optical and electrical) are defined. Test points are defined in 86.8.1, test patterns in 86.8.2, and parameters in 86.8.3 and 86.8.4. Further parameters for electrical use, and the use of electrical compliance boards, are defined in 86A.5.

86.8.1 Test points and compliance boards

Figure 86–3 shows the six test points for 40GBASE-SR4 and 100GBASE-SR10. These are TP1, TP1a, TP2, TP3, TP4, and TP4a; four of these are Skew points SP2, SP3, SP4, and SP5 as shown. Figure 86–3 also shows the substitution of compliance boards for PMD or PMA. This is explained in Annex 86A. Table 86–10 shows the parameters or signals measured at each point, including the electrical compliance points.

All transmitter optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified. A patch cord that connects the MDI transmit side to 4 or 10 individual connectors may be suitable.

86.8.2 Test patterns and related subclauses

While compliance is to be achieved in normal operation, specific test patterns are defined for measurement consistency and to enable measurement of some parameters. Table 86–11 lists the defined test patterns, and Table 86–12 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. As Pattern 3 is more demanding than

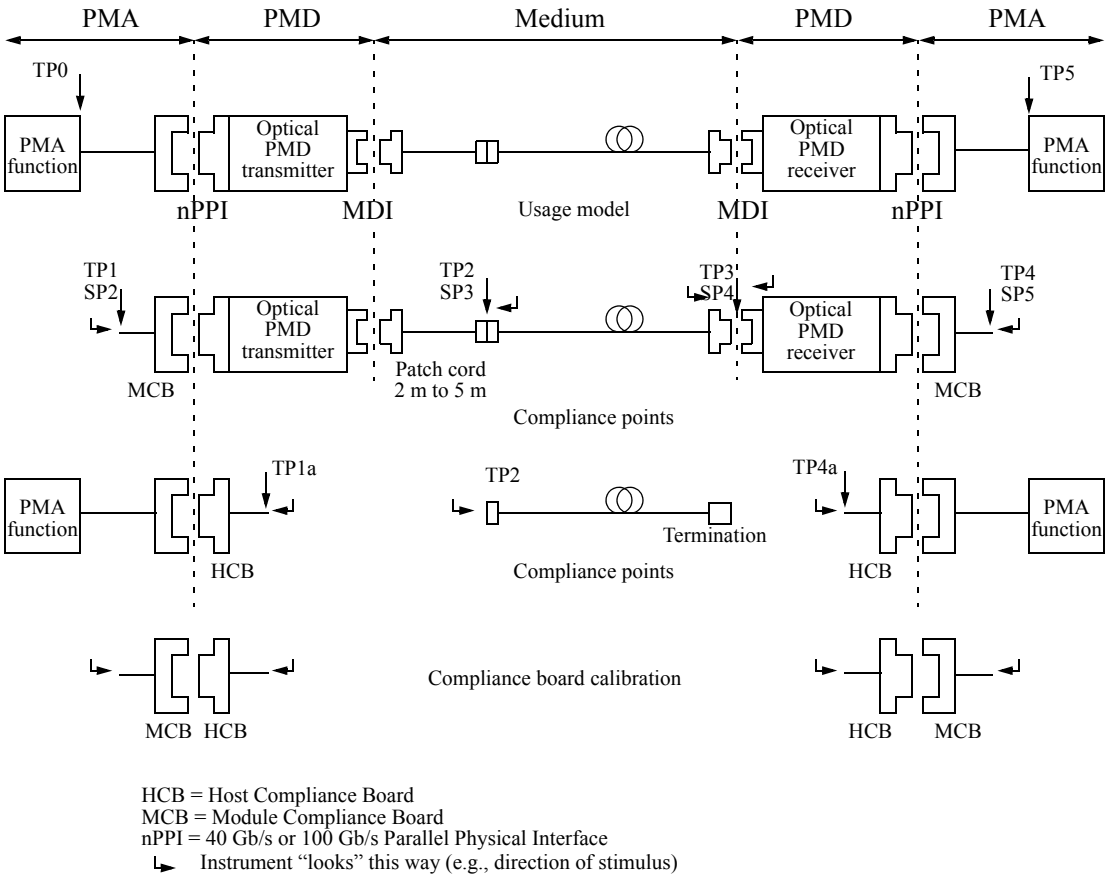

Figure 86–3—Test points for 40GBASE-SR4 and 100GBASE-SR10

Table 86–10—Parameters defined at each test point

Test point	Direction	Parameter
TP1	Looking downstream into PMD transmitter input	PMD transmitter input return loss
TP1a	Looking upstream into PMA transmitter output	PMA transmitter output signal and output return loss, PMD transmitter compliance signal calibration, PMA receiver compliance crosstalk signal calibration
TP2	Looking upstream into optical transmitter output	PMD transmitted signal, PMD transmitter reflectance
	Looking downstream into fiber	Optical return loss, connector reflections
TP3	Looking upstream into fiber	PMD receiver compliance signal
	Looking downstream into optical receiver input	PMD receiver reflectance
TP4	Looking upstream into PMD receiver output	PMD receiver output signal and output return loss, PMA receiver compliance signal calibration
TP4a	Looking downstream into PMA receiver input	PMA receiver input return loss

Table 86–11—Test patterns

Pattern no.	Pattern	Pattern defined in
Square wave	Square wave (8 ones, 8 zeros)	83.5.10
3	PRBS31	83.5.10
4	PRBS9	83.5.10
5	Scrambled idle	82.2.11

Pattern 5 (which itself is the same or more demanding than other 40GBASE-R or 100GBASE-R bit streams) an item that is compliant using Pattern 5 is considered compliant even if it does not meet the required limit using Pattern 3. Test patterns for further electrical parameters are given in Table 86A–6.

Table 86–12—Test patterns and related subclauses

Parameter	Pattern	Related subclause
Wavelength, spectral width	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.4.1
Average optical power	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.4.2
Transmitter OMA (modulated optical power)	Square wave or 4	86.8.4.3
Extinction ratio	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.4.5
Transmitted optical waveform (eye mask)	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.3.2, 86.8.4.6
TDP (transmitter and dispersion penalty)	3 or 5	86.8.4.4
Stressed receiver sensitivity	3 or 5	86.8.4.7
Receiver jitter tolerance	3 or 5	86.8.4.8
Calibration of OMA for receiver tests	Square wave or 4	52.9.9
Vertical eye closure penalty calibration	3 or 5	52.9.9
J2 Jitter	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.3.3.1
J9 Jitter	3 or 5	86.8.3.3.2
Data Dependent Pulse Width Shrinkage (DDPWS)	4	86A.5.3.4

Table 86–12—Test patterns and related subclauses *(continued)*

Parameter	Pattern	Related subclause
AC common-mode voltage	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86A.5.3.1
Transition time	Square wave or 4	86A.5.3.3
Electrical waveform (eye mask)	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.3.2, 86A.5.3.6

86.8.2.1 Multi-lane testing considerations

TDP is defined for each lane, at a BER of 10^{-12} on that lane. Stressed receiver sensitivity, receiver jitter tolerance and host input signal tolerance (in Annex 86A) are defined for an interface BER of 10^{-12}. The interface BER is the average of the four or ten BERs of the receive lanes when they are stressed.

Measurements with Pattern 3 (PRBS31) allow lane-by-lane BER measurements. Measurements with Pattern 5 (scrambled idle) give the interface BER if all lanes are stressed at the same time. If each lane is stressed in turn, the BER is diluted by the three or nine unstressed lanes, and the BER for that stressed lane alone must be found, e.g., by multiplying by 4 or 10 if the unstressed lanes have low BER. To allow TDP measurement with Pattern 5, unstressed lanes for the error detector may be created by setting the power at the reference receivers well above their sensitivities, or by copying the contents of the transmit lanes not under BER test to the error detector by other means. In stressed receiver sensitivity and receiver jitter tolerance measurements, unstressed lanes may be created by setting the power at the receiver under test well above its sensitivity and/or not stressing those lanes with ISI and jitter, or by other means. Each receive lane is stressed in turn while all are operated. All aggressor lanes are operated as specified. To find the interface BER, the BERs of all the lanes when stressed are averaged.

Where relevant, parameters are defined with all co-propagating and counter-propagating lanes operational so that crosstalk effects are included. Where not otherwise specified, the maximum amplitude (OMA or VMA) for a particular situation is used, and for counter-propagating lanes, the minimum transition time is used. Alternative test methods that generate equivalent results may be used. While the lanes in a particular direction may share a common clock, the Tx and Rx directions are not synchronous to each other. If Pattern 3 is used for the lanes not under test using a common clock, there is at least 31 UI delay between the PRBS31 patterns on one lane and any other lane.

86.8.3 Parameters applicable to both electrical and optical signals

86.8.3.1 Skew and Skew Variation

Skew and Skew Variation are defined in 80.5 and are required to remain within the limits given in 86.3.2 over the time that the link is in operation. Skew points as they relate to the nPPI are shown in Figure 86–3. The measurement of Skew and Skew Variation is made by acquiring the data on each lane using a clock and data recovery unit with a high-frequency corner bandwidth and slope as specified in 86.8.3.2. The arrival times of the one to zero transition of the alignment marker sync bits on each lane are then compared. This arrangement ensures that any high-frequency jitter that is present on the signals is not included in the Skew measurement.

86.8.3.2 Eye diagrams

Eye diagrams can be used to assess both electrical and optical signals. The measurement of the electrical eye is defined in 86A.5.3.6 and the measurement of the optical transmitter waveform is defined in 86.8.4.6. Whether electrical or optical eye diagrams, all co-propagating and counter-propagating lanes are active as crosstalk sources, using one of patterns 3, 5, or a valid 40GBASE-R or 100GBASE-R signal. The input lanes of the item under test are receiving signals that are asynchronous to those being output.

Normalized times of 0 and 1 on the unit interval scale are determined by the eye crossing means measured at the average value of the eye pattern. A clock recovery unit (CRU) is used to trigger the oscilloscope for mask measurements, as shown in Figure 52–9. It has a high-frequency corner bandwidth of 4 MHz and a slope of –20 dB/decade. The CRU tracks acceptable levels of low-frequency jitter and wander.

Consideration should be given as to whether a correction is needed for actual instrument properties.

86.8.3.2.1 Eye mask acceptable hit count examples

An example calculation relating hit count to hit ratio for an eye mask measurement using an oscilloscope is detailed below.

If an oscilloscope records 1350 samples/screen, and the time-base is set to 0.2 UI per division with 10 divisions across the screen, and the measurement is continued for 200 waveforms, then a signal that averages less than 6.75 hits is compliant to a hit ratio specification of 5×10^{-5}, i.e.,

$$\frac{5 \times 10^{-5} \times 200 \times 1350}{0.2 \times 10} = 6.75 \qquad (86–1)$$

Likewise, if a measurement is continued for 1000 waveforms, then an average of less than 33.75 hits is compliant. An extended measurement is expected to give a more accurate result, and a single reading of 6 hits in 200 waveforms would not give a statistically significant pass or fail.

The hit ratio limit has been chosen to avoid misleading results due to signal and oscilloscope noise.

86.8.3.3 Jitter

J2 Jitter and J9 Jitter are specified with all co-propagating and counter-propagating lanes active as crosstalk sources, using one of patterns 3, 5, or a valid 40GBASE-R or 100GBASE-R signal. The input lanes of the item under test are receiving signals that are asynchronous to those being output.

86.8.3.3.1 J2 Jitter

J2 Jitter is defined as the time interval that includes all but 10^{-2} of the jitter distribution, which is the time interval from the 0.5th to the 99.5th percentile of the jitter histogram. This may be measured using an oscilloscope, or if measured by plotting BER vs. decision time, J2 is the time interval between the two points with a BER of 2.5×10^{-3}. Oscilloscope histograms should include at least 10 000 hits, and should be taken over about 1% of the signal amplitude. Test patterns are given in Table 86–12.

86.8.3.3.2 J9 Jitter

J9 Jitter is defined as the time interval that includes all but 10^{-9} of the jitter distribution. If measured by plotting BER vs. decision time, it is the time interval between the two points with a BER of 2.5×10^{-10}. Test patterns are given in Table 86–12.

86.8.4 Optical parameter definitions

86.8.4.1 Wavelength and spectral width

The wavelength of each optical lane shall be within the range given in Table 86–6 if measured using the method given in TIA–455–127-A. The lane under test is modulated using the test pattern defined in Table 86–12.

86.8.4.2 Average optical power

Average optical power is defined by the methods given in IEC 61280-1-1.

86.8.4.3 Optical Modulation Amplitude (OMA)

OMA shall be as defined in 52.9.5 for measurement with a square wave (8 ones, 8 zeros) test pattern or 68.6.2 (from the variable MeasuredOMA in 68.6.6.2) for measurement with a PRBS9 test pattern, with the exception that each optical lane is tested individually. See 86.8.2 for test pattern information.

86.8.4.4 Transmitter and dispersion penalty (TDP)

Transmitter and dispersion penalty (TDP) shall be as defined for 10GBASE-S in 52.9.10 with the following exceptions:

a) Each optical lane is tested individually with all other lanes in operation.
b) The test pattern is as defined in Table 86–12.
c) The transmitter is tested using an optical channel with an optical return loss of 12 dB.
d) The reference receiver (including the effect of the decision circuit) has a fourth order Bessel-Thomson filter response with a bandwidth of 6.1 GHz. The transversal filter of 52.9.10.3 is not used.
e) The reference sensitivity S and the measurement P_DUT are both measured with the sampling instant displaced from the eye center by ± 0.15 UI. For each of the two cases (early and late), if P_DUT(i) is larger than S(i), the TDP(i) for the transmitter under test is the difference between P_DUT(i) and S(i), TDP(i) = P_DUT(i) – S(i). Otherwise, TDP(i) = 0. The TDP is the larger of the two TDP(i).
f) The test setup illustrated in Figure 52–12 shows the reference method. Other measurement implementations may be used with suitable calibration.
g) The BER of 10^{-12} is for the lane under test on its own. See 86.8.2.1 for multi-lane pattern considerations.

NOTE—Because practical receivers and decision circuits have noise and timing impairments, the sampling instant offsets have to be calibrated. One method of doing this is via a jitter bathtub method using a known low-jitter signal.

86.8.4.5 Extinction ratio

Extinction ratio is defined by the methods of IEC 61280–2–2 using the test pattern defined in Table 86–12.

NOTE—Extinction ratio and OMA are defined with different test patterns (see Table 86–12).

86.8.4.6 Transmitter optical waveform (transmit eye)

The required optical transmitter pulse shape characteristics are specified in the form of a mask of the transmitter eye diagram specified in Table 86–6 and shown in Figure 86–4. The transmitter optical waveform of a port transmitting the test pattern specified in Table 86–12 shall meet specifications according to the methods specified in 86.8.4.6.1 with the filter nominal reference frequency f_r of 7.5 GHz and filter tolerances as specified for STM-64 in ITU-T G.691.

86.8.4.6.1 Optical transmitter eye mask

The eye is measured with respect to the mask specified in Table 86–6 and shown in Figure 86–4 using a receiver with the fourth-order Bessel-Thomson response having a transfer function given by Equation (86–2) and Equation (86–3):

$$H(y) = \frac{105}{105 + 105y + 45y^2 + 10y^3 + y^4} \tag{86–2}$$

where:

$$y = 2.114p\ ; \quad p = \frac{j\omega}{\omega_r}\ ; \quad \omega_r = 2\pi f_r\ ; \quad f_r = \text{Reference frequency in GHz} \tag{86–3}$$

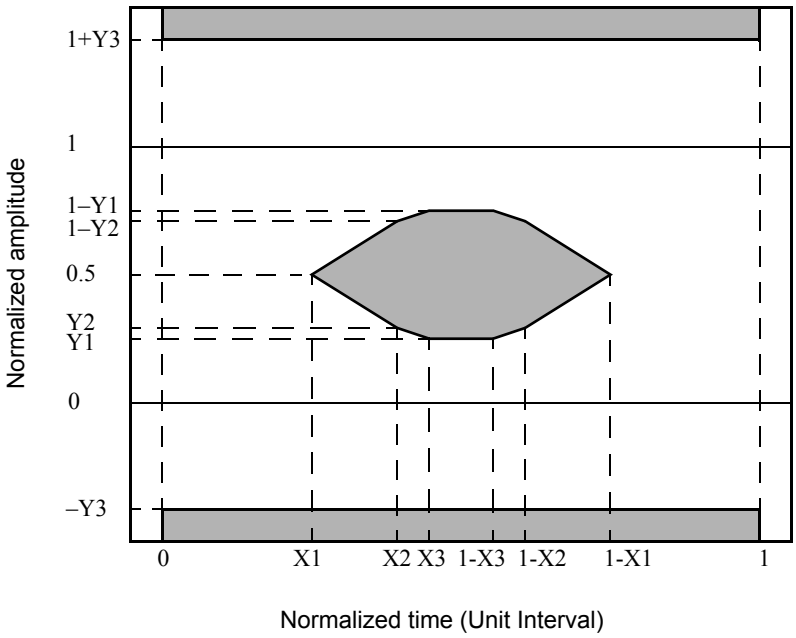

Figure 86–4—Transmitter eye mask definition

The 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. Compensation may be made for variation of the reference receiver filter response from an ideal fourth-order Bessel-Thomson response.

Normalized levels of 0 and 1 represent logic zero and one respectively. These are defined by the means of the lower and upper halves of the central 0.2 UI of the eye.

Further requirements are given in 86.8.3.2.

The transmitter shall achieve a hit ratio lower than the limit of hits per sample specified in the appropriate table or 5×10^{-5} hits per sample if not otherwise specified. “Hits” are the number of samples within the grey areas of Figure 86–4, and the sample count is the total number of samples from 0 UI to 1 UI. Some illustrative examples are provided in 86.8.3.2.1.

Further information on optical eye pattern measurement procedures may be found in IEC 61280-2-2.

86.8.4.7 Stressed receiver sensitivity

Stressed receiver sensitivity shall be within the limits given in Table 86–8 if measured using the method defined by 52.9.9 with the conformance test signal at TP3 and with the following exceptions:

a) The reference test procedure for a single lane is defined in 52.9.9. See 86.8.2.1 and below for multi-lane considerations.
b) The sinusoidal amplitude interferer is replaced by a Gaussian noise generator.
c) The fourth-order Bessel-Thomson filter is replaced by a low-pass filter followed by a limiter and a fourth-order Bessel-Thomson filter.
d) The sinusoidal jitter is at a fixed 80 MHz frequency and between 0 and 0.05 UI peak-to-peak amplitude.
e) The Gaussian noise generator, the amplitude of the sinusoidal jitter, and the Bessel-Thomson filter are adjusted so that the VECP, J2 Jitter and J9 Jitter specifications given in Table 86–8 are simultaneously met (the random noise effects such as RIN, random clock jitter do not need to be minimized).
f) The pattern for the received compliance signal is specified in Table 86–12.
g) The interface BER of the PMD receiver is the average of the BER of all receive lanes while stressed and at the specified receive OMA.
h) Where nPPI or XLAUI/CAUI-n is exposed, a PMD receiver is considered compliant if it meets the module electrical output specifications at TP4 given in Table 86A–3 for nPPI, or the requirements in Table 83B–3 for XLAUI/CAUI-10, or the requirements in Table 83E–3 for CAUI-4.

Stressed receiver sensitivity is defined with all transmit and receive lanes in operation. Each receive lane is tested in turn while all aggressor receive lanes are operated as specified in Table 86–8. Pattern 3 or Pattern 5, or a valid 40GBASE–R4 or 100GBASE–R10 signal is sent from the transmit section of the receiver under test. The signal being transmitted is asynchronous to the received signal. If Pattern 3 is used for the transmit and receive lanes not under test, there is at least 31 UI delay between the PRBS31 patterns generated on one lane and any other lane.

For 40GBASE-SR4 and 100GBASE-SR10, the relevant BER is the interface BER. The interface BER is the average of the four or ten BERs of the receive lanes when stressed: see 86.8.2.1.

86.8.4.8 Receiver jitter tolerance

Receiver jitter tolerance shall be as defined as in 68.6.11, with the following differences:

a) The reference test procedure for a single lane is defined in 68.6.11. See 86.8.2.1 for multi-lane considerations.
b) The pattern to be received is specified in Table 86–12.
c) The parameters of the signal are specified in Table 86–8 and the power in OMA at the receiver is set to the maximum for receiver jitter tolerance in OMA given in Table 86–8.
d) Each receive lane is tested in turn while all are operated. All aggressor lanes are operated as specified in Table 86–8.
e) The receive lanes not being tested are receiving Pattern 3, Pattern 5, or a valid 40GBASE-R4 or 100GBASE-R10 signal.
f) The transmitter is transmitting one of these signals using all lanes.
g) The transmitter and the receiver are not synchronous.
h) The interface BER of the PMD receiver is the average of the BER of all receive lanes when stressed.
i) The mode-conditioning patch cord suitable for 62.5/125 μm fiber is not used.

86.9 Safety, installation, environment, and labeling

86.9.1 General safety

All equipment subject to this clause shall conform to IEC 60950-1.

86.9.2 Laser safety

40GBASE–SR4 and 100GBASE–SR10 optical transceivers shall conform to Hazard Level 1M 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 explicitly defines requirements and usage restrictions on the host system necessary to meet these safety certifications.[11]

86.9.3 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.

86.9.4 Environment

The 40GBASE–SR4 and 100GBASE–SR10 operating environment specifications are as defined in 52.11, as defined in 52.11.1 for electromagnetic emission, and as defined in 52.11.2 for temperature, humidity, and handling.

86.9.5 PMD labeling

The 40GBASE–SR4 and 100GBASE–SR10 labeling recommendations and requirements are as defined in 52.12.

86.10 Optical channel

The fiber optic cabling (channel) contains 4 or 10 optical fibers for each direction to support 40GBASE-SR4 or 100GBASE-SR10, respectively. The fiber optic cabling interconnects the transmitters at the MDI on one end of the channel to the receivers at the MDI on the other end of the channel. As defined in 86.10.3, the optical lanes appear in defined locations at the MDI but the locations are not assigned specific lane numbers within this standard because any transmitter lane may be connected to any receiver lane.

86.10.1 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 86–5.

The channel shall comply with the specifications in Table 86–13.

[11]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.

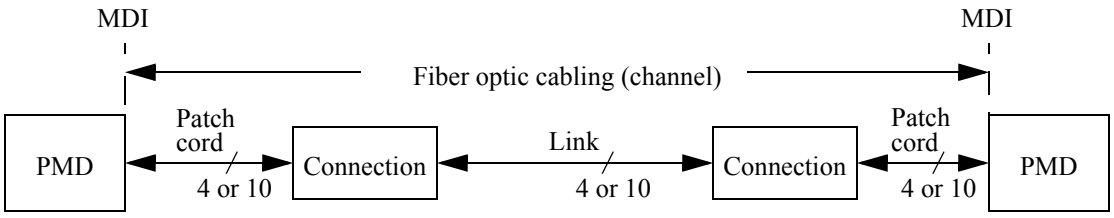

Figure 86–5—Fiber optic cabling model

Table 86–13—Fiber optic cabling (channel) characteristics at 850 nm

Description	Type	OM3	OM4	Unit
Operating distance	Max	100	150	m
Cabling Skew	Max	79		ns
Cabling Skew Variation[a]	Max	2.2		ns
Channel insertion loss	Min	0		dB
Channel insertion loss[b]	Max	1.9[c]	1.5[d]	dB

[a] An additional 0.6 ns of Skew Variation could be caused by wavelength changes, which are attributable to the transmitter not the channel.
[b] These channel insertion loss values include cable, connectors, and splices.
[c] 1.5 dB allocated for connection and splice loss.
[d] 1 dB allocated for connection and splice loss.

A channel may contain additional connectors as long as the optical characteristics of the channel, such as attenuation, modal dispersion, reflections and losses of all connectors and splices meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with IEC 61280-4-1:2009. As OM4 optical fiber meets the requirements for OM3, a channel compliant to the "OM3" column may use OM4 optical fiber, or a combination of OM3 and OM4.

86.10.2 Characteristics of the fiber optic cabling (channel)

86.10.2.1 Optical fiber cable

The fiber contained within the 40GBASE–SR4 or 100GBASE–SR10 fiber optic cabling shall comply with the specifications and parameters of Table 86–14. A variety of multimode cable types may satisfy these requirements, provided the resulting channel also meets the specifications of Table 86–13. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together.

86.10.2.2 Optical fiber connection

An optical fiber connection, as shown in Figure 86–5, consists of a mated pair of optical connectors.

Table 86–14—Optical fiber and cable characteristics

Description	OM3[a]	OM4[b]	Unit
Nominal core diameter	50		μm
Nominal fiber specification wavelength	850		nm
Effective modal bandwidth (min)[c]	2000	4700	MHz•km
Cabled optical fiber attenuation (max)	3.5		dB/km
Zero dispersion wavelength (λ_0)	$1295 \le \lambda_0 \le 1340$		nm
Chromatic dispersion slope (max) (S_0)	0.105 for $1295 \le \lambda_0 \le 1310$ and $0.000375 \times (1590 - \lambda_0)$ for $1310 \le \lambda_0 \le 1340$		ps/nm^2 km

[a] IEC 60793-2-10 type A1a.2
[b] IEC 60793-2-10 type A1a.3
[c] When measured with the launch conditions specified in Table 86–6.

86.10.2.2.1 Connection insertion loss

The maximum operating distances are based on an allocation of 1.5 dB or 1 dB total connection and splice loss. For example, these allocations support two connections, each with an insertion loss of 0.75 dB or 0.5 dB respectively. Connections with lower loss characteristics may be used provided the requirements of Table 86–13 are met. However, the loss of a single connection shall not exceed 0.75 dB.

86.10.2.2.2 Maximum discrete reflectance

The maximum discrete reflectance shall be less than –20 dB.

86.10.3 Medium Dependent Interface (MDI)

The 40GBASE–SR4 or 100GBASE–SR10 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 86–5). The 40GBASE–SR4 PMD is coupled to the fiber optic cabling through one connector plug into the MDI optical receptacle as shown in Figure 86–6. The 100GBASE–SR10 PMD is coupled to the fiber optic cabling through one or two connector plugs into the MDI optical receptacle(s), depending on choice of implementation, as shown in Figure 86–7. Example constructions of the MDI include the following:

a) PMD with a connectorized fiber pigtail plugged into an adapter;
b) PMD receptacle.

86.10.3.1 Optical lane assignments for 40GBASE-SR4

The four transmit and four receive optical lanes of 40GBASE-SR4 shall occupy the positions depicted in Figure 86–6 when looking into the MDI receptacle with the connector keyway feature on top. The interface contains eight active lanes within twelve total positions. The transmit optical lanes occupy the leftmost four positions. The receive optical lanes occupy the rightmost four positions. The four center positions are unused.

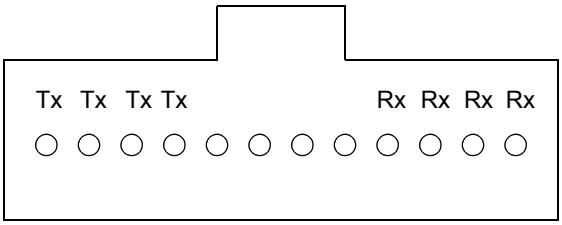

Figure 86–6—40GBASE-SR4 optical lane assignments

86.10.3.2 Optical lane assignments for 100GBASE-SR10

The ten transmit and ten receive optical lanes of 100GBASE-SR10 shall occupy the positions depicted in Figure 86–7 when looking into the MDI optical receptacle(s) with the connector keyway feature(s) on top. The single-receptacle Option A is recommended, the two-receptacle Option B and Option C are alternatives. The interface contains 20 active lanes within up to 24 total positions arranged in two rows of 10 or 12 positions. One row is dedicated to transmit optical lanes and the other row to receive optical lanes. For the depicted 12-position rows, the optical lanes occupy the center ten positions of each row with the outermost positions unused.

86.10.3.3 Medium Dependent Interface (MDI) requirements

The MDI adapter or receptacle shall meet the dimensional specifications of IEC 61754-7 interface 7-3, the MPO adapter interface. The plug terminating the optical fiber cabling shall meet the dimensional specifications of IEC 61754-7 interface 7-4, MPO female plug connector flat interface. The MDI shall optically mate with the plug on the optical fiber cabling. Figure 86–8 shows an MPO female plug connector with flat interface, and an MDI as a PMD receptacle using an MPO adapter interface.

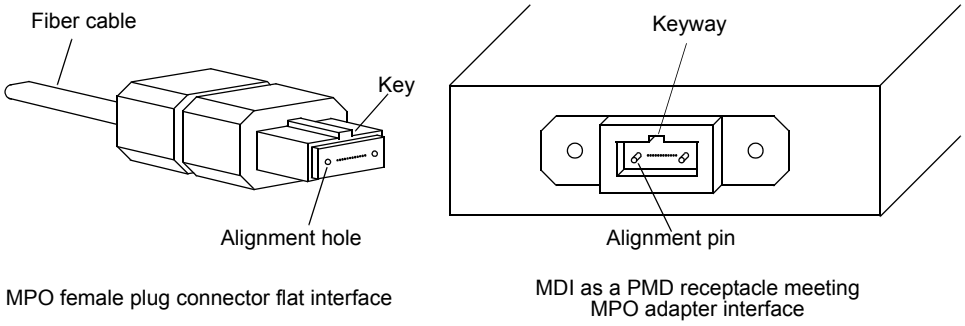

Figure 86–8—MPO female plug and MDI as a PMD receptacle using MPO adapter

The MDI shall meet the interface performance specifications of IEC 61753-1-1 and IEC 61753-022-2.

NOTE—Transmitter compliance testing is performed at TP2 as defined in 86.5.1, not at the MDI.

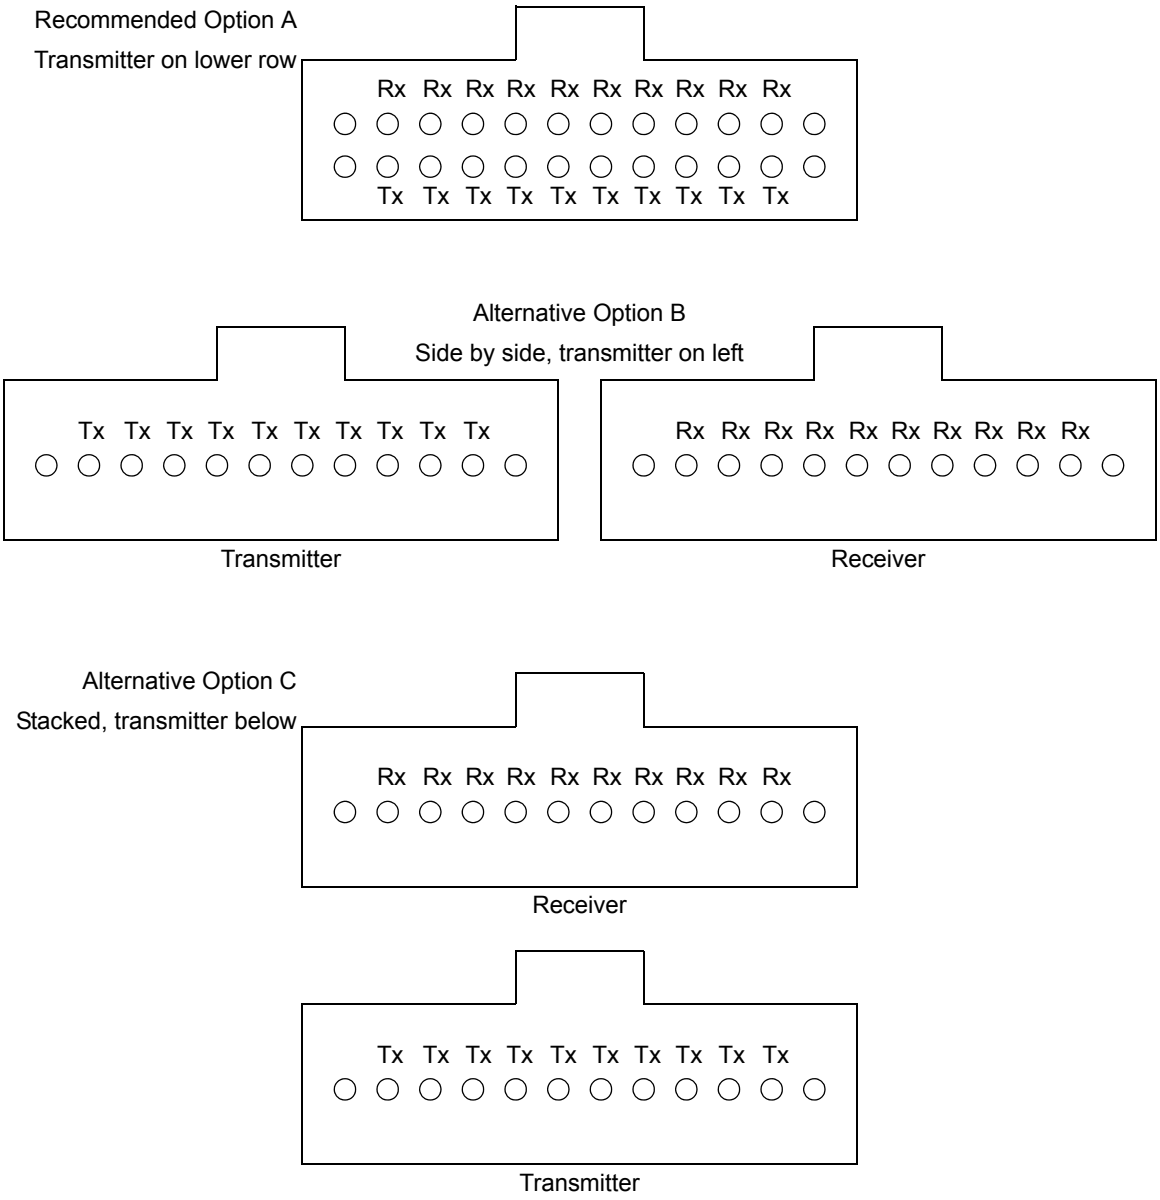

Figure 86–7—100GBASE-SR10 optical lane assignments

86.11 Protocol implementation conformance statement (PICS) proforma for Clause 86, Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE–SR4 and 100GBASE–SR10[12]

86.11.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 86, Physical Medium Dependent sublayer and medium, type 40GBASE–SR4 and 100GBASE–SR10, 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.

86.11.2 Identification

86.11.2.1 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).	

86.11.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 86, Physical Medium Dependent sublayer and medium, type 40GBASE–SR4 and 100GBASE–SR10
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.

86.11.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*SR4	40GBASE–SR4 PMD	86.7	Can operate as 40GBASE–SR4 PMD	O.1	Yes [ ] No [ ]
*SR10	100GBASE–SR10 PMD	86.7	Can operate as 100GBASE–SR10 PMD	O.1	Yes [ ] No [ ]
*INS	Installation / cable	86.10.1	Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]
*TP1	Compliance point TP1 exposed and available for testing	86.5.1	See text	O	Yes [ ] No [ ]
*TP4	Compliance point TP4 exposed and available for testing	86.5.1	See text	O	Yes [ ] No [ ]
*PIT	nPPI Tx interface	86.1	Uses XLPPI or CPPI host to module (see Annex 86A)	TP1:O	Yes [ ] No [ ]
*PIR	nPPI Rx interface	86.1	Uses XLPPI or CPPI module to host (see Annex 86A)	TP4:O	Yes [ ] No [ ]
*MD	MDIO capability	86.4	Registers and interface supported	O	Yes [ ] No [ ]

86.11.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, types 40GBASE–SR4 and 100GBASE–SR10

86.11.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
SF1	Compatible with 40GBASE–R or 100GBASE–R PCS and PMA	86.1	See text	M	Yes [ ]
SF2	Integration with management functions	86.1	See text	O	Yes [ ] No [ ]
D	Delay constraints	86.3.1	For SR4, max 25.6 ns. For SR10, max 20.48 ns.	M	Yes [ ]
SF3	Skew and Skew Variation at SP3 and SP4	86.3.2	At SP3, less than 54 ns, 600 ps. At SP4, less than 134 ns, 3.4 ns.	M	Yes [ ]
SF4	Skew and Skew Variation at SP5 (TP4)	86.3.2	If measurable, less than 145 ns, 3.6 ns.	TP4: M	Yes [ ] N/A [ ]
SF5	Transmit function	86.5.2	Conveys bits from PMD service interface to MDI	M	Yes [ ]
SF6	Delivery to the MDI	86.5.2	4 or 10 optical signal streams for delivery to the MDI	M	Yes [ ]
SF7	Mapping between optical signal and logical signal for transmitter	86.5.2	Higher optical power is a one	M	Yes [ ]
SF8	Receive function	86.5.3	Conveys bits from MDI to PMD service interface	M	Yes [ ]
SF9	Conversion of optical signals to electrical signals	86.5.3	For delivery to the PMD service interface	M	Yes [ ]
SF10	Mapping between optical signal and logical signal for receiver	86.5.3	Higher optical power is a one	M	Yes [ ]
SF11	Global Signal Detect function	86.5.4	Report to the PMD service interface the message PMD:IS_SIGNAL.indication(SIGNAL_DETECT)	M	Yes [ ]
SF12	Global Signal Detect behavior	86.5.4	Global indicator of the presence of optical signals on all lanes per Table 86–5	M	Yes [ ]
SF13	Lane-by-lane Signal Detect function	86.5.5	Sets PMD_signal_detect_*i* values on a lane-by-lane basis per Table 86–5	MD:O	Yes [ ] No [ ] N/A [ ]
SF14	PMD_reset function	86.5.6	Resets the PMD sublayer	MD:O	Yes [ ] No [ ] N/A [ ]

86.11.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
SM1	Management register set	86.4	See 86.4	MD:M	Yes [ ] N/A [ ]
SM2	Global transmit disable function	86.5.7	Disables all of the optical transmitters with Global_PMD_transmit_disable	MD:O	Yes [ ] No [ ] N/A [ ]
SM3	PMD_transmit_disable_*i* function	86.5.8	Disables the optical transmitter on the lane associated with PMD_transmit_disable_*i*	MD:O	Yes [ ] No [ ] N/A [ ]
SM4	PMD lane-by-lane transmit disable	86.5.8	Disables each optical transmitter independently if MD = No	!MD:O	Yes [ ] No [ ] N/A [ ]
SM5	PMD_fault function	86.5.9	Sets PMD_fault to one if any local fault is detected	O	Yes [ ] No [ ]
SM6	PMD_fault mapping to MDIO	86.5.9	Mapped to PMA/PMD fault bit	SM5*MD:M	Yes [ ] N/A [ ]
SM7	PMD_transmit_fault function	86.5.10	Sets PMD_transmit_fault to one if a local fault is detected on any transmit lane	O	Yes [ ] No [ ]
SM8	PMD_transmit_fault mapping to MDIO	86.5.10	Mapped to PMA/PMD transmit fault bit	SM7*MD:M	Yes [ ] N/A [ ]
SM9	PMD_receive_fault function	86.5.11	Sets PMD_receive_fault to one if a local fault is detected on any receive lane	O	Yes [ ] No [ ]
SM10	PMD_receive_fault mapping to MDIO	86.5.11	Mapped to PMA/PMD receive fault bit	SM9*MD:M	Yes [ ] N/A [ ]

86.11.4.3 Optical specifications for 40GBASE–SR4 or 100GBASE–SR10

Item	Feature	Subclause	Value/Comment	Status	Support
S1	Signaling rate per Table 86–2	86.7	10.3125 ±100 ppm	M	Yes [ ] N/A [ ]
S2	Optical transmitter meets specifications in Table 86–6	86.7.1	Per definitions in 86.8	M	Yes [ ] N/A [ ]
S3	Optical receiver meets specifications in Table 86–8	86.7.3	Per definitions in 86.8	M	Yes [ ] N/A [ ]

86.11.4.4 Definitions of parameters and measurement methods

Item	Feature	Subclause	Value/Comment	Status	Support
SOM1	Measurement cable	86.8.1	2 m to 5 m in length	M	Yes [ ]
SOM2	Center wavelength	86.8.4.1	Per TIA-455-127-A	M	Yes [ ]
SOM3	Average optical power	86.8.4.2	Per IEC 61280-1-1	M	Yes [ ]
SOM4	OMA measurements	86.8.4.3	Each lane, per methods of 52.9.5 or 68.6.2	M	Yes [ ]
SOM5	TDP	86.8.4.4	Each lane	M	Yes [ ]
SOM6	Extinction ratio	86.8.4.5	Per IEC 61280-2-2	M	Yes [ ]
SOM7	Transmit eye	86.8.4.6	Per 86.8.4.6.1 and 86.8.3.2, hit ratio lower than specified limit	M	Yes [ ]
SOM8	Stressed receiver conformance	86.8.4.7	See 86.8.4.7	M	Yes [ ]
SOM9	Receiver jitter tolerance	86.8.4.8	Per 68.6.11 as modified	M	Yes [ ]

86.11.4.5 Environmental and safety specifications

Item	Feature	Subclause	Value/Comment	Status	Support
SES1	General safety	86.9.1	Conforms to IEC 60950-1	M	Yes [ ]
SES2	Laser safety	86.9.2	Conforms to Hazard Level 1M laser requirements defined in IEC 60825-1 and IEC 60825-2	M	Yes [ ]
SES3	Electromagnetic interference	86.9.4	As 52.11. Complies with codes for the limitation of electromagnetic interference	M	Yes [ ]

86.11.4.6 Optical channel and MDI

Item	Feature	Subclause	Value/Comment	Status	Support
SOC1	Fiber optic cabling	86.10.1	Specified in Table 86–13	INS:M	Yes [ ] N/A [ ]
SOC2	Optical fiber characteristics	86.10.2.1	Per Table 86–14	INS:M	Yes [ ] N/A [ ]
SOC3	Maximum loss per connection	86.10.2.2.1	0.75 dB	INS:M	Yes [ ] N/A [ ]
SOC4	Maximum discrete reflectance	86.10.2.2.2	Less than –20 dB	INS:M	Yes [ ] N/A [ ]
SOC5	40G MDI layout	86.10.3.1	Optical lane assignments per Figure 86–6	SR4:M	Yes [ ] N/A [ ]
SOC6	100G MDI layout	86.10.3.2	Optical lane assignments per Figure 86–7	SR10:M	Yes [ ] N/A [ ]
SOC7	MDI dimensions	86.10.3.3	Per IEC 61754-7 interface 7-3	M	Yes [ ]
SOC8	Cabling connector dimensions	86.10.3.3	Per IEC 61754-7 interface 7-4	INS:M	Yes [ ] N/A [ ]
SOC9	MDI mating	86.10.3.3	MDI optically mates with plug on the cabling	M	Yes [ ]
SOC10	MDI requirements	86.10.3.3	Meets IEC 61753-1-1 and IEC 61753-022-2	M INS:M	Yes [ ] N/A [ ]

87. Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-LR4 and 40GBASE-ER4

87.1 Overview

This clause specifies the 40GBASE-LR4 PMD and the 40GBASE-ER4 PMD together with the single-mode fiber medium. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 87–1, to the medium through the MDI and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 87–1—Physical Layer clauses associated with the 40GBASE-LR4 and 40GBASE-ER4 PMDs

Associated clause	40GBASE-LR4	40GBASE-ER4
81—RS	Required	Required
81—XLGMII[a]	Optional	Optional
82—PCS for 40GBASE-R	Required	Required
83—PMA for 40GBASE-R	Required	Required
83A—XLAUI	Optional	Optional
83B—Chip-to-module XLAUI	Optional	Optional
86A—XLPPI	Optional	Not applicable
78—Energy Efficient Ethernet	Optional	Optional

[a]The XLGMII is an optional interface. However, if the XLGMII is not implemented, a conforming implementation must behave functionally as though the RS and XLGMII were present.

Figure 87–1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 40 Gb/s and 100 Gb/s Ethernet is introduced in Clause 80 and the purpose of each PHY sublayer is summarized in 80.2.

40GBASE-LR4 and 40GBASE-ER4 PHYs with the optional Energy Efficient Ethernet (EEE) fast wake capability may enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization (see Clause 78). The deep sleep mode of EEE is not supported.

87.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 40GBASE-LR4 and 40GBASE-ER4 PMDs. The service interface for these PMDs are described in an abstract manner and do not imply any particular implementation, although an optional implementation of the PMD service interface, the 40 Gb/s Parallel Physical Interface (XLPPI), is specified in Annex 86A. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

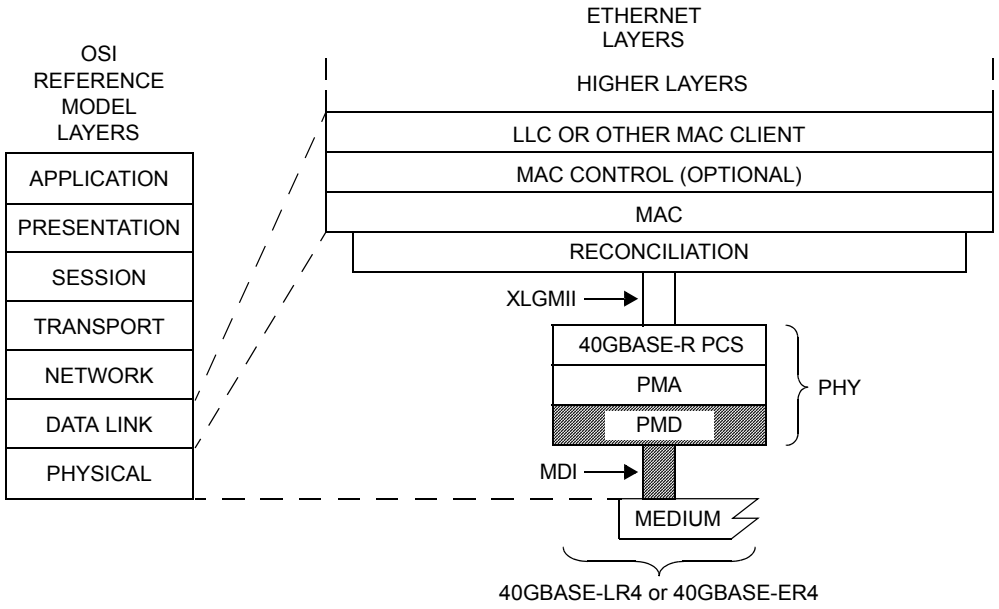

LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

LR = PMD FOR SINGLE-MODE FIBER — 10 km
ER = PMD FOR SINGLE-MODE FIBER — 40 km

Figure 87–1—40GBASE-LR4 and 40GBASE-ER4 PMDs relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_*i*.request
PMD:IS_UNITDATA_*i*.indication
PMD:IS_SIGNAL.indication

The 40GBASE-LR4 and 40GBASE-ER4 PMDs have four parallel bit streams, hence $i = 0$ to 3.

In the transmit direction, the PMA continuously sends four parallel bit streams to the PMD, one per lane, each at a nominal signaling rate of 10.3125 GBd. The PMD converts these streams of bits into appropriate signals on the MDI.

In the receive direction, the PMD continuously sends four parallel bit streams to the PMA corresponding to the signals received from the MDI, one per lane, each at a nominal signaling rate of 10.3125 GBd.

The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication(SIGNAL_OK) inter-sublayer service primitive defined in 80.3.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the rx_bit parameters are undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_bit parameters are 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.

87.3 Delay and Skew

87.3.1 Delay constraints

The sum of the transmit and the receive delays at one end of the link contributed by the 40GBASE-LR4 or 40GBASE-ER4 PMD including 2 m of fiber in one direction shall be no more than 1024 bit times (2 pause_quanta or 25.6 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

87.3.2 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the PCS. The Skew Variation must also be limited to ensure that a given PCS lane always traverses the same physical lane. Skew and Skew Variation are defined in 80.5 and specified at the points SP1 to SP6 shown in Figure 80–6 and Figure 80–7.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 80.5. The measurements of Skew and Skew Variation are defined in 87.8.2.

87.4 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If the MDIO interface is implemented, the mapping of MDIO control variables to PMD control variables shall be as shown in Table 87–2 and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 87–3.

87.5 PMD functional specifications

The 40GBASE-LR4 and 40GBASE-ER4 PMDs perform the Transmit and Receive functions, which convey data between the PMD service interface and the MDI.

87.5.1 PMD block diagram

The PMD block diagram is shown in Figure 87–2. 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 single-mode fiber patch cord (TP2), between 2 m and 5 m in length. Unless specified otherwise, all

Table 87–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/ bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable register	1.9.0	PMD_global_transmit_disable
PMD transmit disable 3	PMD transmit disable register	1.9.4	PMD_transmit_disable_3
PMD transmit disable 2	PMD transmit disable register	1.9.3	PMD_transmit_disable_2
PMD transmit disable 1	PMD transmit disable register	1.9.2	PMD_transmit_disable_1
PMD transmit disable 0	PMD transmit disable register	1.9.1	PMD_transmit_disable_0

Table 87–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	PMD_global_signal_detect
PMD receive signal detect 3	PMD receive signal detect register	1.10.4	PMD_signal_detect_3
PMD receive signal detect 2	PMD receive signal detect register	1.10.3	PMD_signal_detect_2
PMD receive signal detect 1	PMD receive signal detect register	1.10.2	PMD_signal_detect_1
PMD receive signal detect 0	PMD receive signal detect register	1.10.1	PMD_signal_detect_0

transmitter measurements and tests defined in 87.8 are made at TP2. The optical receive signal is defined at the output of the fiber optic cabling (TP3) at the MDI (see 87.11.3). Unless specified otherwise, all receiver measurements and tests defined in 87.8 are made at TP3.

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 accessible in an implemented system).

87.5.2 PMD transmit function

The PMD Transmit function shall convert the four bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_3.request 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 tx_bit = one.

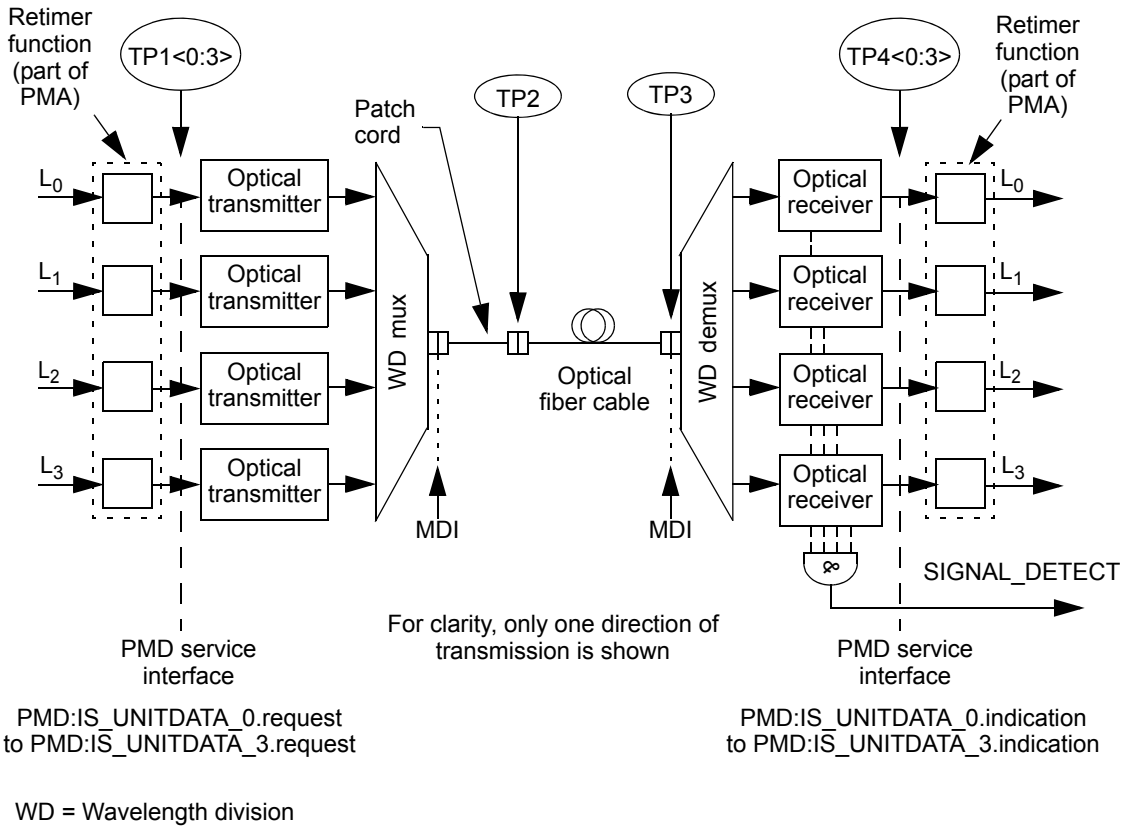

NOTE—Specification of the retimer function is beyond the scope of this standard.

Figure 87–2—Block diagram for 40GBASE-LR4 and 40GBASE-ER4 transmit/receive paths

87.5.3 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 bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_3.indication, all according to the receive optical specifications in this clause. The higher optical power level in each signal stream shall correspond to rx_bit = one.

87.5.4 PMD global signal detect function

The PMD global 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:IS_SIGNAL.indication message is generated when a change in the value of SIGNAL_DETECT occurs. The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the inter-sublayer service interface primitives defined in 80.3.

SIGNAL_DETECT shall be a global indicator of the presence of optical signals on all four lanes. The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 87–4. The PMD receiver is not required to verify whether a compliant 40GBASE-R signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

Table 87–4—SIGNAL_DETECT value definition

Receive conditions	SIGNAL_DETECT value
For any lane; Average optical power at TP3 ≤ –30 dBm	FAIL
For all lanes; [(Optical power at TP3 ≥ receiver sensitivity (max) in OMA in Table 87–8) AND (compliant 40GBASE–R 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 including the effects of 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.

87.5.5 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_*i*, where *i* represents the lane number in the range 0:3, shall be continuously set in response to the magnitude of the optical signal on its associated lane, according to the requirements of Table 87–4.

87.5.6 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.

87.5.7 PMD global transmit disable function (optional)

The PMD_global_transmit_disable function is optional and allows all of the optical transmitters to be disabled.

a) When the PMD_global_transmit_disable variable is set to one, this function shall turn off all of the optical transmitters so that each transmitter meets the requirements of the average launch power of the OFF transmitter in Table 87–7.
b) If a PMD_fault is detected, then the PMD may set the PMD_global_transmit_disable to one, turning off the optical transmitter in each lane.

87.5.8 PMD lane-by-lane transmit disable function

The PMD_transmit_disable_*i* (where *i* represents the lane number in the range 0:3) function is optional and allows the optical transmitters in each lane to be selectively disabled.

a) When a PMD_transmit_disable_*i* 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 87–7.
b) If a PMD_fault is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the optical transmitter in each lane.

If the optional PMD_transmit_disable_*i* function is not implemented in MDIO, an alternative method shall be provided to independently disable each transmit lane for testing purposes.

87.5.9 PMD fault function (optional)

If the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set PMD_fault to one. If the MDIO interface is implemented, PMD_fault shall be mapped to the fault bit as specified in 45.2.1.2.3.

87.5.10 PMD transmit fault function (optional)

If the PMD has detected a local fault on any transmit lane, the PMD shall set the PMD_transmit_fault variable to one. If the MDIO interface is implemented, PMD_transmit_fault shall be mapped to the transmit fault bit as specified in 45.2.1.7.4.

87.5.11 PMD receive fault function (optional)

If the PMD has detected a local fault on any receive lane, the PMD shall set the PMD_receive_fault variable to one. If the MDIO interface is implemented, PMD_receive_fault shall be mapped to the receive fault bit as specified in 45.2.1.7.5.

87.6 Wavelength-division-multiplexed lane assignments

The wavelength range for each lane of the 40GBASE-LR4 and 40GBASE-ER4 PMDs is defined in Table 87–5. The center wavelengths are members of the CWDM wavelength grid defined in ITU-T G.694.2 and are spaced at 20 nm.

Table 87–5—Wavelength-division-multiplexed lane assignments

Lane	Center wavelength	Wavelength range
L_0	1271 nm	1264.5 to 1277.5 nm
L_1	1291 nm	1284.5 to 1297.5 nm
L_2	1311 nm	1304.5 to 1317.5 nm
L_3	1331 nm	1324.5 to 1337.5 nm

NOTE—There is no requirement to associate a particular electrical lane with a particular optical lane, as the PCS is capable of receiving lanes in any arrangement.

87.7 PMD to MDI optical specifications for 40GBASE-LR4 and 40GBASE-ER4

The operating range for the 40GBASE-LR4 and 40GBASE-ER4 PMDs are defined in Table 87–6. A 40GBASE-LR4 or 40GBASE-ER4 compliant PMD operates on type B1.1, B1.3 or B6_a single-mode fibers according to the specifications defined in Table 87–14. A PMD that exceeds the required operating range while meeting all other optical specifications is considered compliant (e.g., a 40GBASE-LR4 PMD operating at 12.5 km meets the operating range requirement of 2 m to 10 km).

The 40GBASE-ER4 PMD interoperates with the 40GBASE-LR4 PMD provided that the channel requirements defined in 87.12 are met.

Table 87–6—40GBASE-LR4 and 40GBASE-ER4 operating ranges

PMD type	Required operating range
40GBASE-LR4	2 m to 10 km
40GBASE-ER4	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 worst case specified for B1.1, B1.3, or B6_a single-mode fiber.

87.7.1 40GBASE-LR4 and 40GBASE-ER4 transmitter optical specifications

The 40GBASE-LR4 transmitter shall meet the specifications defined in Table 87–7 per the definitions in 87.8. The 40GBASE-ER4 transmitter shall meet the specifications defined in Table 87–7 per the definitions in 87.8.

Table 87–7—40GBASE-LR4 and 40GBASE-ER4 transmit characteristics

Description	40GBASE-LR4	40GBASE-ER4	Unit
Signaling rate, each lane (range)	10.3125 ± 100 ppm		GBd
Lane wavelengths (range)	1264.5 to 1277.5 1284.5 to 1297.5 1304.5 to 1317.5 1324.5 to 1337.5		nm
Side-mode suppression ratio (SMSR), (min)	30		dB
Total average launch power (max)	8.3	10.5	dBm
Average launch power, each lane (max)	2.3	4.5	dBm
Average launch power, each lane[a] (min)	–7	–2.7	dBm
Optical Modulation Amplitude (OMA), each lane (max)	3.5	5	dBm
Optical Modulation Amplitude (OMA), each lane (min)[b]	–4	0.3	dBm
Difference in launch power between any two lanes (OMA) (max)	6.5	4.7	dB
Launch power in OMA minus TDP, each lane (min)	–4.8	–0.5	dBm
Transmitter and dispersion penalty (TDP), each lane (max)	2.6		dB
Average launch power of OFF transmitter, each lane (max)	–30		dBm
Extinction ratio (min)	3.5	5.5	dB

Table 87–7—40GBASE-LR4 and 40GBASE-ER4 transmit characteristics *(continued)*

Description	40GBASE-LR4	40GBASE-ER4	Unit
$RIN_{20}OMA$ (max)	–128		dB/Hz
Optical return loss tolerance (max)	20		dB
Transmitter reflectance[c] (max)	–12		dB
Transmitter eye mask definition {X1, X2, X3, Y1, Y2, Y3}	{0.25, 0.4, 0.45, 0.25, 0.28, 0.4}		

[a]Average launch power, each lane (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]Even if the TDP < 0.8 dB, the OMA (min) must exceed this value.
[c]Transmitter reflectance is defined looking into the transmitter.

87.7.2 40GBASE-LR4 and 40GBASE-ER4 receive optical specifications

The 40GBASE-LR4 receiver shall meet the specifications defined in Table 87–8 per the definitions in 87.8. The 40GBASE-ER4 receiver shall meet the specifications defined in Table 87–8 per the definitions in 87.8.

Table 87–8—40GBASE-LR4 and 40GBASE-ER4 receive characteristics

Description	40GBASE-LR4	40GBASE-ER4	Unit
Signaling rate, each lane (range)	10.3125 ± 100 ppm		GBd
Lane wavelengths (range)	1264.5 to 1277.5 1284.5 to 1297.5 1304.5 to 1317.5 1324.5 to 1337.5		nm
Damage threshold[a] (min)	3.3	3.8	dBm
Average receive power, each lane (max)	2.3	–4.5	dBm
Average receive power, each lane[b] (min)	–13.7	–21.2	dBm
Receive power, each lane (OMA) (max)	3.5	–4	dBm
Difference in receive power between any two lanes (OMA) (max)	7.5	7	dB
Receiver reflectance (max)	–26		dB
Receiver sensitivity (OMA), each lane[c] (max)	–11.5	–19	dBm
Receiver 3 dB electrical upper cutoff frequency, each lane (max)	12.3		GHz
Stressed receiver sensitivity (OMA), each lane[d] (max)	–9.6	–16.8	dBm

Table 87–8—40GBASE-LR4 and 40GBASE-ER4 receive characteristics *(continued)*

Description	40GBASE-LR4	40GBASE-ER4	Unit
Conditions of stressed receiver sensitivity test:			
Vertical eye closure penalty,[e] each lane	1.9	2.2	dB
Stressed eye J2 Jitter,[e] each lane	0.3		UI
Stressed eye J9 Jitter,[e] each lane	0.47		UI

[a]The receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having this average power level.
[b]Average receive power, each lane (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 (OMA), each lane (max) is informative.
[d]Measured with conformance test signal at TP3 (see 87.8.11) for BER = 10^{-12}.
[e]Vertical eye closure penalty, stressed eye J2 Jitter, and stressed eye J9 Jitter are test conditions for measuring stressed receiver sensitivity. They are not characteristics of the receiver.

87.7.3 40GBASE-LR4 and 40GBASE-ER4 illustrative link power budgets

Illustrative power budgets and penalties for 40GBASE-LR4 and 40GBASE-ER4 channels are shown in Table 87–9.

Table 87–9—40GBASE-LR4 and 40GBASE-ER4 illustrative link power budgets

Parameter	40GBASE-LR4	40GBASE-ER4		Unit
Power budget (for max TDP)	9.3	21.1		dB
Operating distance	10	30	40[a]	km
Channel insertion loss	6.7[b]	16.5[b]	18.5[a]	dB
Maximum discrete reflectance	–26	–26		dB
Allocation for penalties[c] (for max TDP)	2.6	2.6		dB
Additional insertion loss allowed	0	2	0	dB

[a]Links longer than 30 km are considered engineered links. Attenuation for such links needs to be less than the worst case for B1.1, B1.3, or B6_a single-mode cabled optical fiber.
[b]The channel insertion loss is calculated using the maximum distance specified in Table 87–6 and cabled optical fiber attenuation of 0.47 dB/km at 1264.5 nm plus an allocation for connection and splice loss given in 87.11.2.1.
[c]Link penalties are used for link budget calculations. They are not requirements and are not meant to be tested.

87.8 Definition of optical parameters and measurement methods

All transmitter optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified.

87.8.1 Test patterns for optical parameters

While compliance is to be achieved in normal operation, specific test patterns are defined for measurement consistency and to enable measurement of some parameters. Table 87–11 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. Any of the test patterns given for a particular test in Table 87–11 may be used to perform that test. The test patterns used in this clause are shown in Table 87–10.

Table 87–10—Test patterns

Pattern	Pattern description	Defined in
Square wave	Square wave (8 ones, 8 zeros)	83.5.10
3	PRBS31	83.5.10
4	PRBS9	83.5.10
5	Scrambled idle	82.2.11

Table 87–11—Test-pattern definitions and related subclauses

Parameter	Pattern	Related subclause
Wavelength	3, 5 or valid 40GBASE-LR4 or 40GBASE-ER4 signal	87.8.3
Side mode suppression ratio	3, 5 or valid 40GBASE-LR4 or 40GBASE-ER4 signal	—
Average optical power	3, 5 or valid 40GBASE-LR4 or 40GBASE-ER4 signal	87.8.4
Optical modulation amplitude (OMA)	Square wave or 4	87.8.5
Transmitter and dispersion penalty (TDP)	3 or 5	87.8.6
Extinction ratio	3, 5 or valid 40GBASE-LR4 or 40GBASE-ER4 signal	87.8.7
$RIN_{20}OMA$	Square wave or 4	87.8.8
Transmitter optical waveform	3, 5 or valid 40GBASE-LR4 or 40GBASE-ER4 signal	87.8.9
Stressed receiver sensitivity	3 or 5	87.8.11
Calibration of OMA for receiver tests	Square wave or 4	87.8.11
Vertical eye closure penalty calibration	3 or 5	87.8.11
Receiver 3 dB electrical upper cutoff frequency	3, 5 or valid 40GBASE-LR4 or 40GBASE-ER4 signal	87.8.12

87.8.2 Skew and Skew Variation

Skew and Skew Variation are defined in 80.5 and are required to remain within the limits given in 87.3.2 over the time that the link is in operation. The measurement of Skew and Skew Variation is made by separating optical lanes with an optical demultiplexer and then acquiring the data on each lane using clock and data recovery units with high-frequency corner bandwidths as specified in 86.8.3.2 and a slope of –20 dB/decade. The arrival times of the one to zero transition of the alignment marker sync bits on each lane are then compared. This arrangement ensures that any high frequency jitter that is present on the signals is not included in the Skew measurement.

87.8.3 Wavelength

The wavelength of each optical lane shall be within the ranges given in Table 87–5 if measured per TIA/EIA-455-127-A or IEC 61280-1-3. The lane under test is modulated using the test pattern defined in Table 87–11.

87.8.4 Average optical power

The average optical power of each lane shall be within the limits given in Table 87–7 if measured using the methods given in IEC 61280-1-1, with the sum of the optical power from all of the lanes not under test below –30 dBm, per the test setup in Figure 53–6.

87.8.5 Optical Modulation Amplitude (OMA)

OMA shall be as defined in 52.9.5 for measurement with a square wave (8 ones, 8 zeros) test pattern or 68.6.2 (from the variable MeasuredOMA in 68.6.6.2) for measurement with a PRBS9 test pattern, with the exception that each lane may be tested individually with the sum of the optical power from all of the lanes not under test being below –30 dBm, or if other lanes are operating, a suitable optical filter may be used to separate the lane under test.

87.8.6 Transmitter and dispersion penalty

Transmitter and dispersion penalty (TDP) shall be as defined in 52.9.10 with the exception that each optical lane is tested individually using an optical filter to separate the lane under test from the others. The measurement procedure for 40GBASE-LR4 and 40GBASE-ER4 is detailed in 87.8.6.1 to 87.8.6.4.

The optical filter passband ripple shall be limited to 0.5 dB peak-to-peak and the isolation is chosen such that the ratio of the power in the lane being measured to the sum of the powers of all of the other lanes is greater than 20 dB (see ITU-T G.959.1 Annex B). The lanes not under test shall be operating with PRBS31 or valid 40GBASE-R bit streams.

87.8.6.1 Reference transmitter requirements

The reference transmitter is a high-quality instrument-grade device that can be implemented by a CW laser modulated by a high-performance modulator. The basic requirements are as follows:

a) Rise/fall times of less than 25 ps at 20% to 80%.
b) The output optical eye is symmetric and passes the transmitter optical waveform test of 87.8.9.
c) In the center 20% region of the eye, the worst-case vertical eye closure penalty as defined in 52.9.9.2 is less than 0.5 dB.
d) Total Jitter less than 0.2 UI peak-to-peak.
e) RIN of less than –136 dB/Hz.

87.8.6.2 Channel requirements

The transmitter is tested using an optical channel that meets the requirements listed in Table 87–12.

Table 87–12—Transmitter compliance channel specifications

PMD type	Dispersion[a] (ps/nm)		Insertion loss[b]	Optical return loss[c]
	Minimum	Maximum		
40GBASE-LR4	$0.2325 \times \lambda \times [1 - (1324 / \lambda)^4]$	$0.2325 \times \lambda \times [1 - (1300 / \lambda)^4]$	Minimum	20 dB
40GBASE-ER4	$0.93 \times \lambda \times [1 - (1324 / \lambda)^4]$	$0.93 \times \lambda \times [1 - (1300 / \lambda)^4]$	Minimum	20 dB

[a]The dispersion is measured for the wavelength of the device under test (λ in nm). The coefficient assumes 10 km for 40GBASE-LR4 and 40 km for 40GBASE-ER4.
[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 40GBASE-LR4 or 40GBASE-ER4 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 87–12 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 IEC 60793-1-42 may be used. The measurement is made in the linear power regime of the fiber.

The channel provides an optical return loss specified in Table 87–12. The state of polarization of the back reflection is adjusted to create the greatest RIN.

87.8.6.3 Reference receiver requirements

The reference receiver is required to have the bandwidth given in 87.8.9. The sensitivity of the reference receiver is limited by Gaussian noise. The receiver has minimal threshold offset, deadband, hysteresis, baseline wander, deterministic jitter, or other distortions. Decision sampling has minimal uncertainty and setup/hold times.

The nominal sensitivity of the reference receiver, S, is measured in OMA using the setup 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 is 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}.

The clock recovery unit (CRU) used in the TDP measurement has a corner frequency of 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.

87.8.6.4 Test procedure

The test procedure is as defined in 52.9.10.4 with the exception that all lanes are operational in both directions (transmit and receive), each lane is tested individually using an optical filter to separate the lane under test from the others, and the BER of 1×10^{-12} is for the lane under test on its own.

87.8.7 Extinction ratio

The extinction ratio of each lane shall be within the limits given in Table 87–7 if measured using the methods specified in IEC 61280-2-2, with the sum of the optical power from all of the lanes not under test below –30 dBm. The extinction ratio is measured using the test pattern defined in Table 87–11.

NOTE—Extinction ratio and OMA are defined with different test patterns (see Table 87–11).

87.8.8 Relative Intensity Noise ($RIN_{20}OMA$)

RIN shall be as defined by the measurement methodology of 52.9.6 with the following exceptions:

a) The optical return loss is 20 dB.
b) Each lane may be tested individually with the sum of the optical power from all of the lanes not under test being below –30 dBm, or if other lanes are operating, a suitable optical filter may be used to separate the lane under test.
c) The upper –3 dB limit of the measurement apparatus is to be approximately equal to the signaling rate (i.e., 10.3 GHz).

87.8.9 Transmitter optical waveform (transmit eye)

The required optical transmitter pulse shape characteristics are specified in the form of a mask of the transmitter eye diagram as shown in Figure 86–4. The transmitter optical waveform of a port transmitting the test pattern specified in Table 87–11 shall meet specifications according to the methods specified in 86.8.4.6.1 with the filter nominal reference frequency f_r of 7.5 GHz and filter tolerances as specified for STM-64 in ITU-T G.691. Compensation may be made for variation of the reference receiver filter response from an ideal fourth-order Bessel-Thomson response.

87.8.10 Receiver sensitivity

Receiver sensitivity, which is defined for an ideal input signal, is informative and compliance 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.

87.8.11 Stressed receiver sensitivity

Stressed receiver sensitivity shall be within the limits given in Table 87–8 if measured using the method described in 87.8.11.1 and 87.8.11.5 with the conformance test signal at TP3 as described in 87.8.11.2. The BER is required to be met for the lane under test on its own.

For each lane, the stressed receiver sensitivity is defined with the transmit section in operation on all four lanes and with the receive lanes not under test also in operation. Pattern 3 or Pattern 5, or a valid 40GBASE-R signal, is sent from the transmit section of the PMD under test. The signal being transmitted is asynchronous to the received signal.

87.8.11.1 Stressed receiver conformance test block diagram

A block diagram for the receiver conformance test is shown in Figure 87–3. The patterns used for testing the receiver are specified in Table 87–11. The optical test signal is conditioned (stressed) using the stressed receiver methodology defined in 87.8.11.2, and has sinusoidal jitter applied as specified in 87.8.11.4. A suitable test set is needed to characterize and verify that the signal used to test the receiver has the appropriate characteristics.

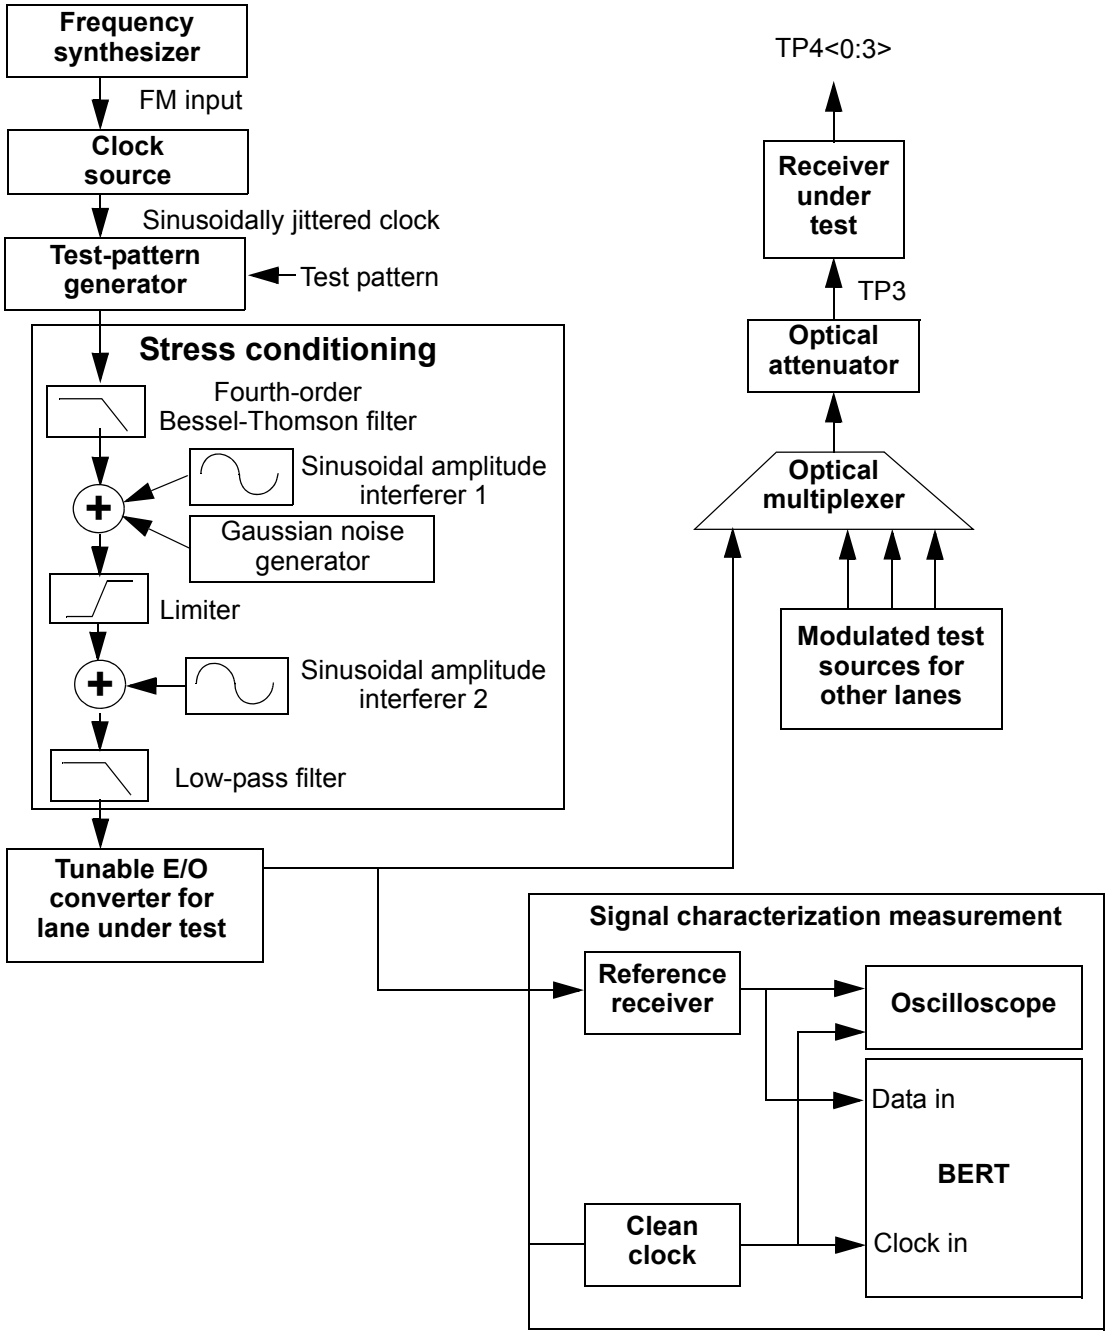

Figure 87–3—Stressed receiver conformance test block diagram

The low-pass filter is used to create ISI-induced vertical eye closure penalty (VECP). The low-pass filter, when combined with the E/O converter, should have a frequency response that results in the appropriate level of initial vertical eye closure before the sinusoidal terms are added.

The sinusoidal amplitude interferer 1 causes jitter that 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 amplitude interferer 2 causes additional eye closure, but in conjunction with the finite edge rates from the limiter, also causes some jitter. The sinusoidally jittered clock represents other forms of jitter and also verifies that the receiver under test can track low-frequency jitter. The sinusoidal amplitude interferers may be set at any frequency between 100 MHz and 2 GHz, although care should be taken to avoid harmonic relationships between the sinusoidal interferers, the sinusoidal jitter, the signaling rate, and the pattern repetition rate. The Gaussian noise generator, the amplitude of the sinusoidal interferers, and the low-pass filter are adjusted so that the VECP, stressed eye J2 Jitter, and stressed eye J9 Jitter specifications are met simultaneously.

For improved visibility for calibration, all elements in the signal path (cables, DC blocks, E/O converter, etc.) should have wide and smooth frequency response, and linear phase response, throughout the spectrum of interest. Baseline wander and overshoot and undershoot should be minimized.

The stressed receiver conformance signal verification is described in 87.8.11.3.

87.8.11.2 Stressed receiver conformance test signal characteristics and calibration

The conformance test signal is used to validate that the PMD receiver of the lane under test meets BER requirements with near worst-case waveforms at TP3.

The primary parameters of the conformance test signal are vertical eye closure penalty (VECP), stressed eye J2 Jitter, and stressed eye J9 Jitter. VECP 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). Stressed eye J2 Jitter is defined in 86.8.3.3.1 and stressed eye J9 Jitter is defined in 86.8.3.3.2. The values of these components are defined by their histogram results, as measured at the average optical power, which can be obtained with AC-coupling.

The vertical eye closure penalty is given by Equation (87–1).

$$\text{Vertical eye closure penalty} = 10\log_{10}\frac{OMA}{A_O} \quad \text{(dB)} \tag{87–1}$$

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

OMA is the optical modulation amplitude as defined in 87.8.5

An example stressed receiver conformance test setup is shown in Figure 87–3; however, any approach that modulates or creates the appropriate levels and frequencies of the VECP and jitter components is acceptable.

Residual low-probability noise and jitter should be minimized so that the outer slopes of the final amplitude histograms are as steep as possible.

The following steps describe a possible method for setting up and calibrating a stressed eye conformance signal when using a stressed receiver conformance test setup as shown in Figure 87–3:

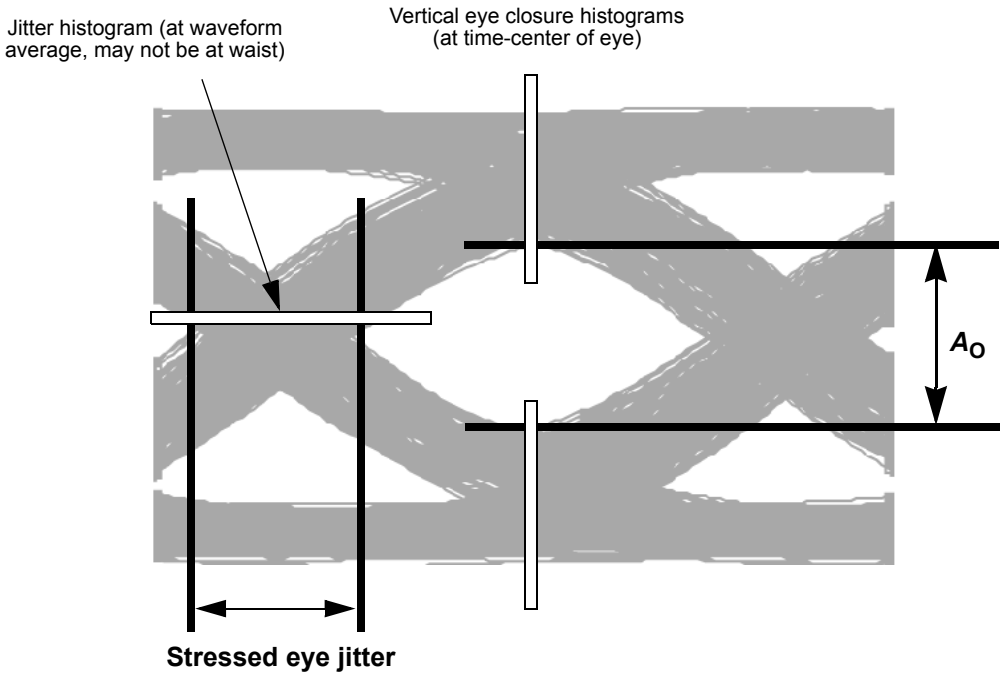

Figure 87–4—Required characteristics of the conformance test signal at TP3

1) Set the signaling rate of the test-pattern generator to meet the requirements in Table 87–8.
2) With sinusoidal interferers and sinusoidal jitter turned off, set the extinction ratio of the E/O to approximately the minimum specified in Table 87–7.
3) The required values of vertical eye closure penalty, stressed eye J2 Jitter, and stressed eye J9 Jitter of the stressed receiver conformance signal are given in Table 87–8.

   With the sinusoidal interferers and sinusoidal jitter turned off, greater than two thirds of the dB value of the VECP should be created by the selection of the appropriate bandwidth for the low-pass filter. Any remaining VECP must be created with sinusoidal interferer 2 or sinusoidal jitter.

   The sinusoidal amplitude interferers may be set at any frequency between 100 MHz and 2 GHz, although care should be taken to avoid harmonic relationships between the sinusoidal interferers, the sinusoidal jitter, the signaling rate and the pattern repetition rate.

   Sinusoidal jitter is added as specified in Table 87–13. When calibrating the conformance signal, the sinusoidal jitter frequency should be well within the 4 MHz to 10 times LB as defined in Table 87–13 and illustrated in Figure 87–5.

   Iterate the adjustments of sinusoidal interferers and Gaussian noise generator until the values of VECP, stressed eye J2 Jitter and stressed eye J9 Jitter meet the requirements in Table 87–8, and sinusoidal jitter above 4 MHz is as specified in Table 87–13. The resulting stressed eye conformance signal is required to have at least 0.05 UI of pulse width shrinkage.

Figure 87–3 shows the stress conditioned signal being applied to a tunable E/O converter. However, any optical source may be used that can meet the OMA and wavelength requirements for the lane under test as described in 87.8.11.5. Similarly, the other test sources that supply modulated signals to the other lanes may use any tunable or fixed sources that meet the OMA and wavelength requirements described in 87.8.11.5.

Each receiver lane is conformance tested in turn. The source for the lane under test is adjusted to supply a signal at the input to the receiver under test at the stressed receiver sensitivity OMA specified in Table 87–8, and the test sources for the other lanes are set to the required OMA and wavelength as described in 87.8.11.5.

87.8.11.3 Stressed receiver conformance test signal verification

The stressed receiver conformance test signal can be verified using an optical reference receiver with an ideal fourth-order Bessel-Thomson response with a reference frequency f_r of 7.5 GHz. Use of G.691 tolerance filters may significantly degrade this calibration. The clock output from the clock source in Figure 87–3 will be modulated with the sinusoidal jitter. To use an oscilloscope to calibrate the final stressed eye J2 Jitter and stressed eye J9 Jitter that includes the sinusoidal jitter component, a separate clock source (clean clock of Figure 87–3) is required that is synchronized to the source clock, but not modulated with the jitter source.

Care should be taken when characterizing the test signal because excessive noise/jitter in the measurement system will result in an input signal that does not fully stress the receiver under test. 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 standard, it is recommended that the implementer fully characterize the test equipment and apply appropriate guard bands to ensure that the stressed receiver conformance input signal meets the stress and sinusoidal jitter specified in 87.8.11.2 and 87.8.11.4.

87.8.11.4 Sinusoidal jitter for receiver conformance test

The sinusoidal jitter is used to test receiver jitter tolerance. The amplitude of the applied sinusoidal jitter is dependent on frequency as specified in Table 87–13 and is illustrated in Figure 87–5.

Table 87–13—Applied sinusoidal jitter

Frequency range	Sinusoidal jitter, peak-to-peak (UI)
$f < 40$ kHz	Not specified
40 kHz $< f \leq 4$ MHz	$2 \times 10^5 / f$
4 MHz $< f <$ 10 LB[a]	0.05

[a] LB = loop bandwidth; upper frequency bound for added sine jitter should be at least 10 times the loop bandwidth of the receiver being tested.

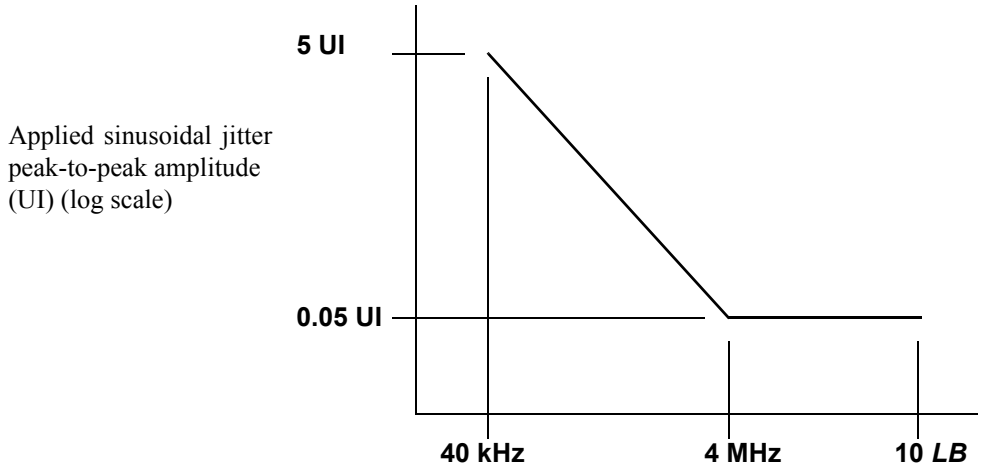

Figure 87–5—Illustration of the mask of the sinusoidal component of jitter tolerance

87.8.11.5 Stressed receiver conformance test procedure for WDM conformance testing

The receiver tests requiring the TP3 conformance test signal are performed on a per lane basis. For each lane, the stressed receiver sensitivity is defined with the transmit section in operation on all four lanes and with the receive lanes not under test also in operation. Pattern 3 or Pattern 5, or a valid 40GBASE-R signal is sent from the transmit section of the PMD under test. The signal being transmitted is asynchronous to the received signal. All test sources are modulated simultaneously, using valid 40GBASE-R signals. The transmitter of the transceiver under test is operating with test patterns as defined in Table 87–11.

A rigorous method for testing WDM conformance using tunable test sources is described as follows:

a) The OMAs of the test sources in the lanes other than the lane under test are set to the highest OMA relative to the test source in the lane under test allowed by the “difference in receive power between any two lanes (OMA) (max)” parameter in Table 87–8.
b) The test sources in the lanes other than the lane under test are tuned to wavelengths within their wavelength range corresponding to the worst-case crosstalk to the lane under test.
c) The test source for the lane under test is tuned to the wavelength within its wavelength range corresponding to the worst-case sensitivity for the receiver under test.

It is recognized that a test setup that uses multiple tunable sources, while allowing rigorous worst-case measurements to be made, is likely to be onerous in practice.

A more practical, alternative method for testing WDM conformance is described below. The alternative WDM conformance test avoids the need for tunable sources but may result in some over-stressing of the receiver under test:

1) The test sources in each lane can be at any wavelength within each lane’s wavelength range.
2) The OMAs of the test sources in the lanes other than the lane under test are set to the highest OMA relative to the test source in the lane under test allowed by the “difference in receive power between any two lanes (OMA) (max)” parameter in Table 87–8, plus an increment corresponding to insertion loss variation within the lane under test (NOTE 1), plus an increment corresponding to the isolation variation of the lane in question (NOTE 2).

NOTE 1—The increment corresponding to the variation of insertion loss with wavelength within the lane under test is equal to the dB variation of measured receiver sensitivity as the test source wavelength is swept across the wavelength range of the lane under test.

NOTE 2—For each of the test sources in the lanes not under test, an increment corresponding to the isolation variation is applied that is equal to the dB variation in optical crosstalk measured in the lane under test as the wavelength of each test source in the lanes not under test is swept across its respective wavelength range.

There are many ways to determine the size of the increments required, two example methods are given in the paragraphs that follow. Note that each lane will have one insertion loss variation value, and three values of optical crosstalk variation (one for each of the other lanes).

Example method 1: If a measure of the received signal strength is available for each lane, scan a tunable laser across the wavelength range of each lane while simultaneously recording the received signal in all lanes. As the laser sweeps across the wavelength range of a particular lane the dB variation in observed signal in that lane is equal to the insertion loss variation. The dB variation in signal observed in each of the other lanes is equal to the optical crosstalk variation.

Example method 2: If a discrete optical demultiplexer and discrete receivers are used, scan a tunable laser across the wavelength range of each channel while simultaneously recording the optical power of all four demultiplexer outputs. As the laser sweeps across the wavelength range of a particular channel, the dB variation in observed optical power in that lane's output is equal to the insertion loss variation. The dB variation in optical power observed in each of the other lanes outputs is equal to the optical crosstalk variation.

87.8.12 Receiver 3 dB electrical upper cutoff frequency

The receiver 3 dB electrical upper cutoff frequency shall be within the limits given in Table 87–8 if measured as described in 52.9.11. Each optical lane is measured using an optical signal or signals with its/their wavelength within the specified wavelength range of the lane to be tested. The test may use an electrical combiner and one optical source as in 53.9.13.

87.9 Safety, installation, environment, and labeling

87.9.1 General safety

All equipment subject to this clause shall conform to IEC 60950-1.

87.9.2 Laser safety

40GBASE-LR4 and 40GBASE-ER4 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 explicitly defines requirements and usage restrictions on the host system necessary to meet these safety certifications.[13]

[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.

87.9.3 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.

87.9.4 Environment

Normative specifications in this clause shall be met by a system integrating a 40GBASE-LR4 or 40GBASE-ER4 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.

87.9.4.1 Electromagnetic emission

A system integrating a 40GBASE-LR4 or 40GBASE-ER4 PMD shall comply with applicable local and national codes for the limitation of electromagnetic interference.

87.9.4.2 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.

87.9.5 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., 40GBASE-LR4).

Labeling requirements for Hazard Level 1 lasers are given in the laser safety standards referenced in 87.9.2.

87.10 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 87–6.

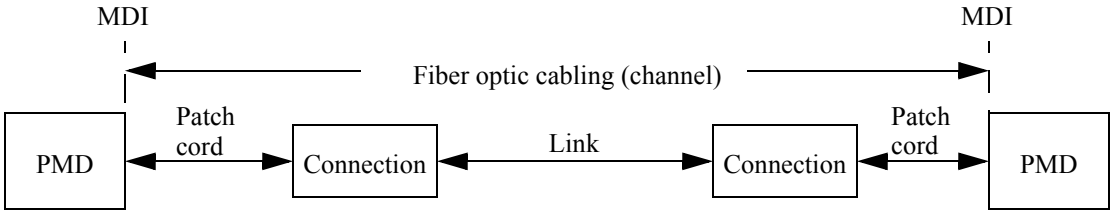

Figure 87–6—Fiber optic cabling model

The channel insertion loss is given in Table 87–14. A channel may contain additional connectors as long as the optical characteristics of the channel, such as attenuation, dispersion, reflections, and polarization mode

dispersion meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with 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 87–14—Fiber optic cabling (channel) characteristics

Description	40GBASE-LR4	40GBASE-ER4		Unit
Operating distance (max)	10	30	40	km
Channel insertion loss[a, b] (max)	6.7	18.5		dB
Channel insertion loss (min)	0	9		dB
Positive dispersion[b] (max)	33.5	100.5	134	ps/nm
Negative dispersion[b] (min)	–59.5	–178.5	–238	ps/nm
DGD_max[c]	10	12		ps
Optical return loss (min)	21	21		dB

[a]These channel insertion loss values include cable, connectors, and splices.
[b]Over the wavelength range 1264.5 nm to 1337.5 nm.
[c]Differential Group Delay (DGD) is the time difference at reception between the fractions of a pulse that were 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.

87.11 Characteristics of the fiber optic cabling (channel)

The 40GBASE-LR4 and 40GBASE-ER4 fiber optic cabling shall meet the specifications defined in Table 87–14. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together.

87.11.1 Optical fiber cable

The fiber optic cable requirements are satisfied by cables containing IEC 60793-2-50 type B1.1 (dispersion un-shifted single-mode), type B1.3 (low water peak single-mode), or type B6_a (bend insensitive) fibers or the requirements in Table 87–15 where they differ.

Table 87–15—Optical fiber and cable characteristics

Description	Value	Unit
Nominal fiber specification wavelength	1310	nm
Cabled optical fiber attenuation (max)	0.47[a] or 0.5[b]	dB/km
Zero dispersion wavelength (λ_0)	$1300 \leq \lambda_0 \leq 1324$	nm
Dispersion slope (max) (S_0)	0.093	ps/nm^2 km

[a]The 0.47 dB/km at 1264.5 nm attenuation for optical fiber cables is derived from Appendix I of ITU-T G.695.
[b]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 for 40GBASE-LR4 or 40 km for 40GBASE-ER4.

87.11.2 Optical fiber connection

An optical fiber connection, as shown in Figure 87–6, consists of a mated pair of optical connectors.

87.11.2.1 Connection insertion loss

The maximum link distances for single-mode fiber are calculated based on an allocation of 2 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 87–14 are met.

87.11.2.2 Maximum discrete reflectance

The maximum discrete reflectance shall be less than –26 dB.

87.11.3 Medium Dependent Interface (MDI) requirements

The 40GBASE-LR4 or 40GBASE-ER4 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 87–6). 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 IEC 61753-1-1 and IEC 61753-021-2.

NOTE—Transmitter compliance testing is performed at TP2 as defined in 87.5.1, not at the MDI.

87.12 Requirements for interoperation between 40GBASE-LR4 and 40GBASE-ER4

The 40GBASE-LR4 and 40GBASE-ER4 PMDs can interoperate with each other (over an engineered link) provided that the fiber optic cabling (channel) characteristics for 40GBASE-LR4 given in Fiber optic cabling (channel) characteristics are met, with the exception of the maximum and minimum channel insertion loss values, which are given in Channel insertion loss requirements for interoperation for the two link directions separately.

Table 87–16—Channel insertion loss requirements for interoperation

Direction	Min loss	Max loss	Unit
40GBASE-LR4 Transmitter to 40GBASE-ER4 Receiver	7.5	14.2	dB
40GBASE-ER4 Transmitter to 40GBASE-LR4 Receiver	2.2	11	dB

87.13 Protocol implementation conformance statement (PICS) proforma for Clause 87, Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-LR4 and 40GBASE-ER4[14]

87.13.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 87, Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-LR4 and 40GBASE-ER4, 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.

87.13.2 Identification

87.13.2.1 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).	

87.13.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 87, Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-LR4 and 40GBASE-ER4
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.

87.13.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*LR4	40GBASE-LR4 PMD	87.7	Device supports requirements for 40GBASE-LR4 PHY	O.1	Yes [ ] No [ ]
*ER4	40GBASE-ER4 PMD	87.7	Device supports requirements for 40GBASE-ER4 PHY	O.1	Yes [ ] No [ ]
*INS	Installation/cable	87.10	Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]
XLTP1	Reference point TP1 exposed and available for testing	87.5.1	This point may be made available for use by implementers to certify component conformance	O	Yes [ ] No [ ]
XLTP4	Reference point TP4 exposed and available for testing	87.5.1	This point may be made available for use by implementers to certify component conformance	O	Yes [ ] No [ ]
XLDC	Delay constraints	87.3.1	Device conforms to delay constraints	M	Yes [ ]
XLSC	Skew constraints	87.3.2	Device conforms to Skew and Skew Variation constraints	M	Yes [ ]
*MD	MDIO capability	87.4	Registers and interface supported	O	Yes [ ] No [ ]

87.13.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-LR4 and 40GBASE-ER4

87.13.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
XLF1	Compatible with 40GBASE-R PCS and PMA	87.1		M	Yes [ ]
XLF2	Integration with management functions	87.1		O	Yes [ ] No [ ]
XLF3	Transmit function	87.5.2	Conveys bits from PMD service interface to MDI	M	Yes [ ]
XLF4	Optical multiplexing and delivery to the MDI	87.5.2	Optically multiplexes the four optical signal streams for delivery to the MDI	M	Yes [ ]
XLF5	Mapping between optical signal and logical signal for transmitter	87.5.2	Higher optical power is a one	M	Yes [ ]
XLF6	Receive function	87.5.3	Conveys bits from MDI to PMD service interface	M	Yes [ ]
XLF7	Conversion of four optical signals to four electrical signals	87.5.3	For delivery to the PMD service interface	M	Yes [ ]
XLF8	Mapping between optical signal and logical signal for receiver	87.5.3	Higher optical power is a one	M	Yes [ ]
XLF9	Global Signal Detect function	87.5.4	Report to the PMD service interface the message PMD:IS_SIGNAL.indication(SIGNAL_DETECT)	M	Yes [ ]
XLF10	Global Signal Detect behavior	87.5.4	SIGNAL_DETECT is a global indicator of the presence of optical signals on all four lanes	M	Yes [ ]
XLF11	Lane-by-Lane Signal Detect function	87.5.5	Sets PMD_signal_detect_*i* values on a lane-by-lane basis per requirements of Table 87–4	MD:O	Yes [ ] No [ ] N/A [ ]
XLF12	PMD_reset function	87.5.6	Resets the PMD sublayer	MD:O	Yes [ ] No [ ] N/A [ ]

87.13.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
XLM1	Management register set	87.4		MD:M	Yes [ ] N/A [ ]
XLM2	Global transmit disable function	87.5.7	Disables all of the optical transmitters with the PMD_global_transmit_disable variable	MD:O	Yes [ ] No [ ] N/A [ ]
XLM3	PMD_lane_by_lane_transmit_disable function	87.5.8	Disables the optical transmitter on the lane associated with the PMD_transmit_disable_*i* variable	MD:O.2	Yes [ ] No [ ] N/A [ ]
XLM4	PMD_lane_by_lane_transmit_disable	87.5.8	Disables each optical transmitter independently if XLM3 = No	O.2	Yes [ ] No [ ]
XLM5	PMD_fault function	87.5.9	Sets PMD_fault to one if any local fault is detected	MD:O	Yes [ ] No [ ] N/A [ ]
XLM6	PMD_transmit_fault function	87.5.10	Sets PMD_transmit_fault to one if a local fault is detected on any transmit lane	MD:O	Yes [ ] No [ ] N/A [ ]
XLM7	PMD_receive_fault function	87.5.11	Sets PMD_receive_fault to one if a local fault is detected on any receive lane	MD:O	Yes [ ] No [ ] N/A [ ]

87.13.4.3 PMD to MDI optical specifications for 40GBASE-LR4

Item	Feature	Subclause	Value/Comment	Status	Support
XLLR1	Transmitter meets specifications in Table 87–7	87.7.1	Per definitions in 87.8	LR4:M	Yes [ ] N/A [ ]
XLLR2	Receiver meets specifications in Table 87–8	87.7.2	Per definitions in 87.8	LR4:M	Yes [ ] N/A [ ]

87.13.4.4 PMD to MDI optical specifications for 40GBASE-ER4

Item	Feature	Subclause	Value/Comment	Status	Support
XLER1	Transmitter meets specifications in Table 87–7	87.7.1	Per definitions in 87.8	ER4:M	Yes [ ] N/A [ ]
XLER2	Receiver meets specifications in Table 87–8	87.7.2	Per definitions in 87.8	ER4:M	Yes [ ] N/A [ ]

87.13.4.5 Optical measurement methods

Item	Feature	Subclause	Value/Comment	Status	Support
XLOM1	Measurement cable	87.8	2 m to 5 m in length	M	Yes [ ]
XLOM2	Center wavelength	87.8.3	Per TIA-455-127-A or IEC 61280-1-3 under modulated conditions	M	Yes [ ]
XLOM3	Average optical power	87.8.4	Per IEC 61280-1-1	M	Yes [ ]
XLOM4	Transmitter and dispersion penalty	87.8.6	Each lane	M	Yes [ ]
XLOM5	OMA	87.8.5	Each lane	M	Yes [ ]
XLOM6	Extinction ratio	87.8.7	Per IEC 61280-2-2	M	Yes [ ]
XLOM7	$RIN_{20}OMA$	87.8.8	Each lane	M	Yes [ ]
XLOM8	Transmit eye	87.8.9	Each lane	M	Yes [ ]
XLOM9	Stressed receiver sensitivity	87.8.11	Each lane	M	Yes [ ]
XLOM10	Receiver 3 dB electrical upper cutoff frequency	87.8.12	Each lane	M	Yes [ ]

87.13.4.6 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
XLES1	General safety	87.9.1	Conforms to IEC 60950-1	M	Yes [ ]
XLES2	Laser safety—IEC Hazard Level 1	87.9.2	Conforms to Hazard Level 1 laser requirements defined in IEC 60825-1 and IEC 60825-2	M	Yes [ ]
XLES3	Electromagnetic interference	87.9.4.1	Complies with applicable local and national codes for the limitation of electromagnetic interference	M	Yes [ ]

87.13.4.7 Characteristics of the fiber optic cabling and MDI

Item	Feature	Subclause	Value/Comment	Status	Support
XLOC1	Fiber optic cabling	87.11	Meets the requirements specified in Table 87–14	INS:M	Yes [ ] N/A [ ]
XLOC2	Maximum discrete reflectance	87.11.2.2	Less than –26 dB	INS:M	Yes [ ] N/A [ ]
XLOC3	MDI requirements	87.11.3	Meets the interface performance specifications listed	INS:M	Yes [ ] N/A [ ]

88. Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-LR4 and 100GBASE-ER4

88.1 Overview

This clause specifies the 100GBASE-LR4 PMD and the 100GBASE-ER4 PMD together with the single-mode fiber medium. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 88–1, to the medium through the MDI and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 88–1—Physical Layer clauses associated with the 100GBASE-LR4 and 100GBASE-ER4 PMDs

Associated clause	100GBASE-LR4	100GBASE-ER4
81—RS	Required	Required
81—CGMII[a]	Optional	Optional
82—PCS for 100GBASE-R	Required	Required
83—PMA for 100GBASE-R	Required	Required
83A—CAUI-10	Optional	Optional
83B—Chip to module CAUI-10	Optional	Optional
83D—CAUI-4	Optional	Optional
83E—Chip-to-module CAUI-4	Optional	Optional
78—Energy Efficient Ethernet	Optional	Optional

[a]The CGMII is an optional interface. However, if the CGMII is not implemented, a conforming implementation must behave functionally as though the RS and CGMII were present.

Figure 88–1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 40 Gb/s and 100 Gb/s Ethernet is introduced in Clause 80 and the purpose of each PHY sublayer is summarized in 80.2.

100GBASE-LR4 and 100GBASE-ER4 PHYs with the optional Energy Efficient Ethernet (EEE) fast wake capability may enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization (see Clause 78). The deep sleep mode of EEE is not supported.

88.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 100GBASE-LR4 and 100GBASE-ER4 PMDs. The service interfaces for these PMDs are described in an abstract manner and do not imply any particular implementation. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

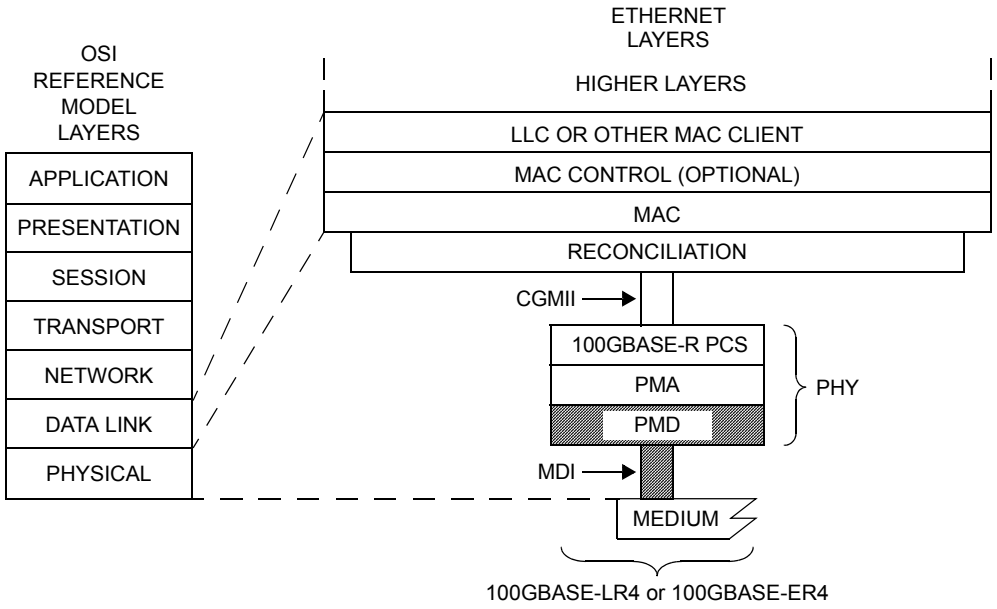

Figure 88–1—100GBASE-LR4 and 100GBASE-ER4 PMDs relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_*i*.request
PMD:IS_UNITDATA_*i*.indication
PMD:IS_SIGNAL.indication

The 100GBASE-LR4 and 100GBASE-ER4 PMDs have four parallel bit streams, hence $i = 0$ to 3.

In the transmit direction, the PMA continuously sends four parallel bit streams to the PMD, one per lane, each at a nominal signaling rate of 25.78125 GBd. The PMD then converts these streams of bits into the appropriate signals on the MDI.

In the receive direction, the PMD continuously sends four parallel bit streams to the PMA corresponding to the signals received from the MDI, one per lane, each at a nominal signaling rate of 25.78125 GBd.

The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication(SIGNAL_OK) inter-sublayer service interface primitive defined in 80.3.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the rx_bit parameters are undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_bit parameters are 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.

88.3 Delay and Skew

88.3.1 Delay constraints

The sum of the transmit and receive delays at one end of the link contributed by the 100GBASE-LR4 or 100GBASE-ER4 PMD including 2 m of fiber in one direction shall be no more than 2048 bit times (4 pause_quanta or 20.48 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

88.3.2 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the PCS. The Skew Variation must also be limited to ensure that a given PCS lane always traverses the same physical lane. Skew and Skew Variation are defined in 80.5 and specified at the points SP1 to SP6 shown in Figure 80–6 and Figure 80–7.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 80.5. The measurements of Skew and Skew Variation are defined in 87.8.2 with the exception that the clock and data recovery units' high-frequency corner bandwidths are 10 MHz.

88.4 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If the MDIO interface is implemented, the mapping of MDIO control variables to PMD control variables shall be as shown in Table 88–2 and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 88–3.

88.5 PMD functional specifications

The 100GBASE-LR4 and 100GBASE-ER4 PMDs perform the Transmit and Receive functions, which convey data between the PMD service interface and the MDI.

88.5.1 PMD block diagram

The PMD block diagram is shown in Figure 88–2. 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

Table 88–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable register	1.9.0	PMD_global_transmit_disable
PMD transmit disable 3	PMD transmit disable register	1.9.4	PMD_transmit_disable_3
PMD transmit disable 2	PMD transmit disable register	1.9.3	PMD_transmit_disable_2
PMD transmit disable 1	PMD transmit disable register	1.9.2	PMD_transmit_disable_1
PMD transmit disable 0	PMD transmit disable register	1.9.1	PMD_transmit_disable_0

Table 88–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	PMD_global_signal_detect
PMD receive signal detect 3	PMD receive signal detect register	1.10.4	PMD_signal_detect_3
PMD receive signal detect 2	PMD receive signal detect register	1.10.3	PMD_signal_detect_2
PMD receive signal detect 1	PMD receive signal detect register	1.10.2	PMD_signal_detect_1
PMD receive signal detect 0	PMD receive signal detect register	1.10.1	PMD_signal_detect_0

end of a single-mode fiber patch cord (TP2), between 2 m and 5 m in length. Unless specified otherwise, all transmitter measurements and tests defined in 88.8 are made at TP2. The optical receive signal is defined at the output of the fiber optic cabling (TP3) at the MDI (see 88.11.3). Unless specified otherwise, all receiver measurements and tests defined in 88.8 are made at TP3.

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 accessible in an implemented system).

88.5.2 PMD transmit function

The PMD Transmit function shall convert the four bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_3.request 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 tx_bit = one.

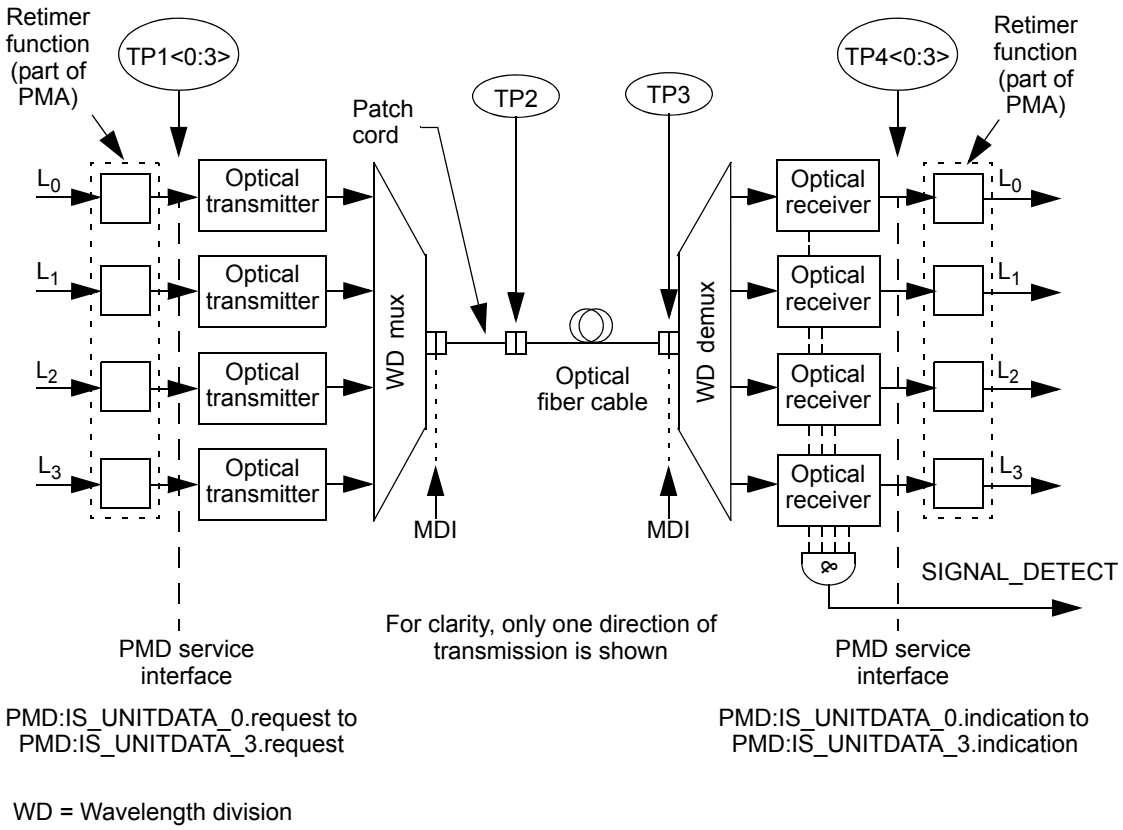

NOTE—Specification of the retimer function and the electrical implementation of the PMD service interface is beyond the scope of this standard.

Figure 88–2—Block diagram for 100GBASE-LR4 and 100GBASE-ER4 transmit/receive paths

88.5.3 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 bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_3.indication, all according to the receive optical specifications in this clause. The higher optical power level in each signal stream shall correspond to rx_bit = one.

88.5.4 PMD global signal detect function

The PMD global 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:IS_SIGNAL.indication message is generated when a change in the value of SIGNAL_DETECT occurs. The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the inter-sublayer service interface primitives defined in 80.3.

SIGNAL_DETECT shall be a global indicator of the presence of optical signals on all four lanes. The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 88–4. The PMD receiver is not required to verify whether a compliant 100GBASE-R signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

Table 88–4—SIGNAL_DETECT value definition

Receive conditions	SIGNAL_DETECT value
For any lane; Average optical power at TP3 ≤ –30 dBm	FAIL
For all lanes; [(Optical power at TP3 ≥ receiver sensitivity (max) in OMA in Table 88–8) AND (compliant 100GBASE–R 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 including the effects of 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.

88.5.5 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_*i*, where *i* represents the lane number in the range 0:3, shall be continuously set in response to the magnitude of the optical signal on its associated lane, according to the requirements of Table 88–4.

88.5.6 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.

88.5.7 PMD global transmit disable function (optional)

The PMD_global_transmit_disable function is optional and allows all of the optical transmitters to be disabled.

a) When the PMD_global_transmit_disable variable is set to one, this function shall turn off all of the optical transmitters so that each transmitter meets the requirements of the average launch power of the OFF transmitter in Table 88–7.
b) If a PMD_fault is detected, then the PMD may set the PMD_global_transmit_disable to one, turning off the optical transmitter in each lane.

88.5.8 PMD lane-by-lane transmit disable function

The PMD_transmit_disable_*i* (where *i* represents the lane number in the range 0:3) function is optional and allows the optical transmitters in each lane to be selectively disabled.

a) When a PMD_transmit_disable_*i* 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 88–7.
b) If a PMD_fault is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the optical transmitter in each lane.

If the optional PMD_transmit_disable_*i* function is not implemented in MDIO, an alternative method shall be provided to independently disable each transmit lane for testing purposes.

88.5.9 PMD fault function (optional)

If the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set PMD_fault to one.

If the MDIO interface is implemented, PMD_fault shall be mapped to the fault bit as specified in 45.2.1.2.3.

88.5.10 PMD transmit fault function (optional)

If the PMD has detected a local fault on any transmit lane, the PMD shall set the PMD_transmit_fault variable to one.

If the MDIO interface is implemented, PMD_transmit_fault shall be mapped to the transmit fault bit as specified in 45.2.1.7.4.

88.5.11 PMD receive fault function (optional)

If the PMD has detected a local fault on any receive lane, the PMD shall set the PMD_receive_fault variable to one.

If the MDIO interface is implemented, PMD_receive_fault shall be mapped to the receive fault bit as specified in 45.2.1.7.5.

88.6 Wavelength-division-multiplexed lane assignments

The wavelength range for each lane of the 100GBASE-LR4 and 100GBASE-ER4 PMDs is defined in Table 88–5. The center frequencies are members of the frequency grid for 100 GHz spacing and above defined in ITU-T G.694.1 and are spaced at 800 GHz.

Table 88–5—Wavelength-division-multiplexed lane assignments

Lane	Center frequency	Center wavelength	Wavelength range
L_0	231.4 THz	1295.56 nm	1294.53 to 1296.59 nm
L_1	230.6 THz	1300.05 nm	1299.02 to 1301.09 nm
L_2	229.8 THz	1304.58 nm	1303.54 to 1305.63 nm
L_3	229 THz	1309.14 nm	1308.09 to 1310.19 nm

NOTE—There is no requirement to associate a particular electrical lane with a particular optical lane, as the PCS is capable of receiving lanes in any arrangement.

88.7 PMD to MDI optical specifications for 100GBASE-LR4 and 100GBASE-ER4

The operating ranges for the 100GBASE-LR4 and 100GBASE-ER4 PMDs are defined in Table 88–6. A 100GBASE-LR4 or 100GBASE-ER4 compliant PMD operates on type B1.1, B1.3, or B6_a single-mode fibers according to the specifications defined in Table 88–14. A PMD that exceeds the operating range requirement while meeting all other optical specifications is considered compliant (e.g., a 100GBASE-LR4 PMD operating at 12.5 km meets the operating range requirement of 2 m to 10 km). The 100GBASE-ER4 PMD interoperates with the 100GBASE-LR4 PMD provided that the channel requirements for 100GBASE-LR4 are met.

Table 88–6—100GBASE-LR4 and 100GBASE-ER4 operating ranges

PMD type	Required operating range
100GBASE-LR4	2 m to 10 km
100GBASE-ER4	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 worst case specified for B1.1, B1.3, or B6_a single-mode fiber.

88.7.1 100GBASE-LR4 and 100GBASE-ER4 transmitter optical specifications

The 100GBASE-LR4 transmitter shall meet the specifications defined in Table 88–7 per the definitions in 88.8. The 100GBASE-ER4 transmitter shall meet the specifications defined in Table 88–7 per the definitions in 88.8.

Table 88–7—100GBASE-LR4 and 100GBASE-ER4 transmit characteristics

Description	100GBASE-LR4	100GBASE-ER4	Unit
Signaling rate, each lane (range)	25.78125 ± 100 ppm		GBd
Lane wavelengths (range)	1294.53 to 1296.59 1299.02 to 1301.09 1303.54 to 1305.63 1308.09 to 1310.19		nm
Side-mode suppression ratio (SMSR), (min)	30		dB
Total average launch power (max)	10.5	8.9	dBm
Average launch power, each lane (max)	4.5	2.9	dBm
Average launch power, each lane[a] (min)	–4.3	–2.9	dBm
Optical Modulation Amplitude (OMA), each lane (max)	4.5		dBm
Optical Modulation Amplitude (OMA), each lane (min)	–1.3[b]	0.1	dBm
Difference in launch power between any two lanes (OMA) (max)	5	—	dB
Difference in launch power between any two lanes (Average and OMA) (max)	—	3.6	
Launch power in OMA minus TDP, each lane (min)	–2.3	—	dBm
Transmitter and dispersion penalty (TDP), each lane (max)	2.2	2.5	dB
Average launch power of OFF transmitter, each lane (max)	–30		dBm
Extinction ratio (min)	4	8	dB
$RIN_{20}OMA$ (max)	–130		dB/Hz
Optical return loss tolerance (max)	20		dB
Transmitter reflectance[c] (max)	–12		dB
Transmitter eye mask definition {X1, X2, X3, Y1, Y2, Y3}	{0.25, 0.4, 0.45, 0.25, 0.28, 0.4}		

[a]Average launch power, each lane (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]Even if the TDP < 1 dB, the OMA (min) must exceed this value.
[c]Transmitter reflectance is defined looking into the transmitter.

88.7.2 100GBASE-LR4 and 100GBASE-ER4 receive optical specifications

The 100GBASE-LR4 receiver shall meet the specifications defined in Table 88–8 per the definitions in 88.8.
The 100GBASE-ER4 receiver shall meet the specifications defined in Table 88–8 per the definitions in 88.8.

Table 88–8—100GBASE-LR4 and 100GBASE-ER4 receive characteristics

Description	100GBASE-LR4	100GBASE-ER4	Unit
Signaling rate, each lane (range)	25.78125 ± 100 ppm		GBd
Lane wavelengths (range)	1294.53 to 1296.59 1299.02 to 1301.09 1303.54 to 1305.63 1308.09 to 1310.19		nm
Damage threshold[a]	5.5		dBm
Average receive power, each lane (max)	4.5[b]		dBm
Average receive power, each lane[c] (min)	–10.6	–20.9	dBm
Receive power, each lane (OMA) (max)	4.5		dBm
Difference in receive power between any two lanes (OMA) (max)	5.5	—	dB
Difference in receive power between any two lanes (Average and OMA) (max)	—	4.5	
Receiver reflectance (max)	–26		dB
Receiver sensitivity (OMA), each lane[d] (max)	–8.6	–21.4	dBm
Receiver 3 dB electrical upper cutoff frequency, each lane (max)	31		GHz
Stressed receiver sensitivity (OMA), each lane[e] (max)	–6.8	–17.9	dBm
Conditions of stressed receiver sensitivity test			
Vertical eye closure penalty,[f] each lane	1.8	3.5	dB
Stressed eye J2 Jitter,[f] each lane	0.3		UI
Stressed eye J9 Jitter,[f] each lane	0.47		UI

[a]The receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having this average power level.
[b]The average receive power, each lane (max) for 100GBASE-ER4 is larger than the 100GBASE-ER4 transmitter value to allow compatibility with 100GBASE-LR4 units at short distances.
[c]Average receive power, each lane (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.
[d]Receiver sensitivity (OMA), each lane (max) is informative.
[e]Measured with conformance test signal at TP3 (see 88.8.10) for BER = 10^{-12}.
[f]Vertical eye closure penalty, stressed eye J2 Jitter, and stressed eye J9 Jitter are test conditions for measuring stressed receiver sensitivity. They are not characteristics of the receiver.

88.7.3 100GBASE-LR4 and 100GBASE-ER4 illustrative link power budgets

Illustrative power budgets and penalties for 100GBASE-LR4 and 100GBASE-ER4 channels are shown in Table 88–9.

Table 88–9—100GBASE-LR4 and 100GBASE-ER4 illustrative link power budgets

Parameter	100GBASE-LR4	100GBASE-ER4		Unit
Power budget (for maximum TDP)	8.5	—		dB
Power budget	—	21.5		dB
Operating distance	10	30	40[a]	km
Channel insertion loss	6.3[b]	15	18	dB
Maximum discrete reflectance	–26	–26		dB
Allocation for penalties[c] (for maximum TDP)	2.2	—		dB
Allocation for penalties[c]	—	3.5		
Additional insertion loss allowed	0	3	0	dB

[a]Links longer than 30 km are considered engineered links. Attenuation for such links needs to be less than the worst case for B1.1, B1.3, or B6_a single-mode cabled optical fiber
[b]The channel insertion loss is calculated using the maximum distance specified in Table 88–6 for 100GBASE-LR4 and fiber attenuation of 0.43 dB/km at 1295 nm plus an allocation for connection and splice loss given in 88.11.2.1.
[c]Link penalties are used for link budget calculations. They are not requirements and are not meant to be tested.

88.8 Definition of optical parameters and measurement methods

All transmitter optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified.

88.8.1 Test patterns for optical parameters

While compliance is to be achieved in normal operation, specific test patterns are defined for measurement consistency and to enable measurement of some parameters. Table 88–11 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. Any of the test patterns given for a particular test in Table 88–11 may be used to perform that test. The test patterns used in this clause are shown in Table 88–10.

88.8.2 Wavelength

The wavelength of each optical lane shall be within the ranges given in Table 88–5 if measured per TIA/EIA-455-127-A or IEC 61280-1-3. The lane under test is modulated using the test pattern defined in Table 88–11.

Table 88–10—Test patterns

Pattern	Pattern description	Defined in
Square wave	Square wave (8 ones, 8 zeros)	83.5.10
3	PRBS31	83.5.10
4	PRBS9	83.5.10
5	Scrambled idle	82.2.11

Table 88–11—Test-pattern definitions and related subclauses

Parameter	Pattern	Related subclause
Wavelength	3, 5 or valid 100GBASE-R signal	88.8.2
Side mode suppression ratio	3, 5 or valid 100GBASE-R signal	—
Average optical power	3, 5 or valid 100GBASE-R signal	88.8.3
Optical modulation amplitude (OMA)	Square wave or 4	88.8.4
Transmitter and dispersion penalty (TDP)	3 or 5	88.8.5
Extinction ratio	3, 5 or valid 100GBASE-R signal	88.8.6
$RIN_{20}OMA$	Square wave or 4	88.8.7
Transmitter optical waveform	3, 5 or valid 100GBASE-R signal	88.8.8
Stressed receiver sensitivity	3 or 5	88.8.10
Calibration of OMA for receiver tests	Square wave or 4	87.8.11
Vertical eye closure penalty calibration	3 or 5	87.8.11
Receiver 3 dB electrical upper cutoff frequency	3, 5 or valid 100GBASE-R signal	88.8.11

88.8.3 Average optical power

The average optical power of each lane shall be within the limits given in Table 88–7 if measured using the methods given in IEC 61280-1-1, with the sum of the optical power from all of the lanes not under test below –30 dBm, per the test setup in Figure 53–6.

88.8.4 Optical Modulation Amplitude (OMA)

OMA shall be as defined in 52.9.5 for measurement with a square wave (8 ones, 8 zeros) test pattern or 68.6.2 (from the variable MeasuredOMA in 68.6.6.2) for measurement with a PRBS9 test pattern, with the exception that each lane may be tested individually with the sum of the optical power from all of the lanes not under test being below –30 dBm, or if other lanes are operating, a suitable optical filter may be used to separate the lane under test.

88.8.5 Transmitter and dispersion penalty (TDP)

Transmitter and dispersion penalty (TDP) shall be as defined in 52.9.10 with the exception that each optical lane is tested individually using an optical filter to separate the lane under test from the others. The measurement procedure for 100GBASE-LR4 and 100GBASE-ER4 is detailed in 88.8.5.1 to 88.8.5.4.

The optical filter passband ripple shall be limited to 0.5 dB peak-to-peak and the isolation is chosen such that the ratio of the power in the lane being measured to the sum of the powers of all of the other lanes is greater than 20 dB (see ITU-T G.959.1 Annex B). The lanes not under test shall be operating with PRBS31 or valid 100GBASE-R bit streams.

88.8.5.1 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. The basic requirements are as follows:

a) Rise/fall times of less than 12 ps at 20% to 80%.
b) The output optical eye is symmetric and passes the transmitter optical waveform test of 88.8.8.
c) In the center 20% region of the eye, the worst-case vertical eye closure penalty as defined in 87.8.11.2 is less than 0.5 dB.
d) Total Jitter less than 0.2 UI peak-to-peak.
e) RIN of less than –138 dB/Hz.

88.8.5.2 Channel requirements

The transmitter is tested using an optical channel that meets the requirements listed in Table 88–12.

Table 88–12—Transmitter compliance channel specifications

PMD type	Dispersion[a] (ps/nm)		Insertion loss[b]	Optical return loss[c]	Max mean DGD
	Minimum	Maximum			
100GBASE-LR4	$0.2325 \cdot \lambda \cdot [1 - (1324 / \lambda)^4]$	$0.2325 \cdot \lambda \cdot [1 - (1300 / \lambda)^4]$	Minimum	20 dB	0.8 ps
100GBASE-ER4	$0.93 \cdot \lambda \cdot [1 - (1324 / \lambda)^4]$	$0.93 \cdot \lambda \cdot [1 - (1300 / \lambda)^4]$	Minimum	20 dB	0.8 ps

[a]The dispersion is measured for the wavelength of the device under test (λ in nm). The coefficient assumes 10 km for 100GBASE-LR4 and 40 km for 100GBASE-ER4.
[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 100GBASE-LR4 or 100GBASE-ER4 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 88–12 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 IEC 60793-1-42 may be used. The measurement is made in the linear power regime of the fiber.

The channel provides an optical return loss specified in Table 88–12. The state of polarization of the back reflection is adjusted to create the greatest RIN.

The mean DGD of the channel is to be less than the value specified in Table 88–12.

88.8.5.3 Reference receiver requirements

The reference receiver is required to have the bandwidth given in 88.8.8. The sensitivity of the reference receiver is limited by Gaussian noise. The receiver has minimal threshold offset, deadband, hysteresis, baseline wander, deterministic jitter, or other distortions. Decision sampling has minimal uncertainty and setup/hold times.

The nominal sensitivity of the reference receiver, S, is measured in OMA using the setup 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 is 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}.

The clock recovery unit (CRU) used in the TDP measurement has a corner frequency of 10 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.

88.8.5.4 Test procedure

The test procedure is as defined in 52.9.10.4 with the exception that all lanes are operational in both directions (transmit and receive), each lane is tested individually using an optical filter to separate the lane under test from the others and the BER of 1×10^{-12} is for the lane under test on its own.

88.8.6 Extinction ratio

The extinction ratio of each lane shall be within the limits given in Table 88–7 if measured using the methods specified in IEC 61280-2-2, with the sum of the optical power from all of the lanes not under test below –30 dBm. The extinction ratio is measured using the test pattern defined in Table 88–11.

NOTE—Extinction ratio and OMA are defined with different test patterns (see Table 88–11).

88.8.7 Relative Intensity Noise ($RIN_{20}OMA$)

RIN shall be as defined by the measurement methodology of 52.9.6 with the following exceptions:

a) The optical return loss is 20 dB.
b) Each lane may be tested individually with the sum of the optical power from all of the lanes not under test being below –30 dBm, or if other lanes are operating, a suitable optical filter may be used to separate the lane under test.
c) The upper –3 dB limit of the measurement apparatus is to be approximately equal to the signaling rate (i.e., 25.8 GHz).

88.8.8 Transmitter optical waveform (transmit eye)

The required optical transmitter pulse shape characteristics are specified in the form of a mask of the transmitter eye diagram as shown in Figure 86–4. The transmitter optical waveform of a port transmitting the test pattern specified in Table 88–11 shall meet specifications according to the methods specified in 86.8.4.6.1 with the exception that the clock recovery unit's high-frequency corner bandwidth is 10 MHz. The filter nominal reference frequency f_r is 19.34 GHz and the filter tolerances are as specified for STM-64

in ITU-T G.691. Compensation may be made for variation of the reference receiver filter response from an ideal fourth-order Bessel-Thomson response.

88.8.9 Receiver sensitivity

Receiver sensitivity, which is defined for an ideal input signal, is informative and compliance 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.

88.8.10 Stressed receiver sensitivity

Stressed receiver sensitivity shall be within the limits given in Table 88–8 if measured using the method defined in 87.8.11 with the following exceptions:

a) Added sinusoidal jitter is as specified in Table 88–13.
b) The stressed eye J2 Jitter, stressed eye J9 Jitter, and vertical eye closure penalty are as given in Table 88–8.
c) The test pattern is as given in Table 88–11.
d) The reference receiver used to verify the conformance test signal is required to have the bandwidth given in 88.8.8.

Table 88–13—Applied sinusoidal jitter

Frequency range	Sinusoidal jitter, peak-to-peak (UI)
$f < 100$ kHz	Not specified
100 kHz $< f \leq$ 10 MHz	$5 \times 10^5 / f$
10 MHz $< f <$ 10 LB[a]	0.05

[a] LB = loop bandwidth; upper frequency bound for added sine jitter should be at least 10 times the loop bandwidth of the receiver being tested.

88.8.11 Receiver 3 dB electrical upper cutoff frequency

The receiver 3 dB electrical upper cutoff frequency shall be within the limits given in Table 88–8 if measured as described in 52.9.11. Each optical lane is measured using an optical signal or signals with its/their wavelength within the specified wavelength range of the lane to be tested. The test may use an electrical combiner and one optical source as in 53.9.13.

88.9 Safety, installation, environment, and labeling

88.9.1 General safety

All equipment subject to this clause shall conform to IEC 60950-1.

88.9.2 Laser safety

100GBASE-LR4 and 100GBASE-ER4 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 explicitly defines requirements and usage restrictions on the host system necessary to meet these safety certifications.[15]

88.9.3 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.

88.9.4 Environment

Normative specifications in this clause shall be met by a system integrating a 100GBASE-LR4 or 100GBASE-ER4 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.

88.9.5 Electromagnetic emission

A system integrating a 100GBASE-LR4 or 100GBASE-ER4 PMD shall comply with applicable local and national codes for the limitation of electromagnetic interference.

88.9.6 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.

88.9.7 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., 100GBASE-LR4).

Labeling requirements for Hazard Level 1 lasers are given in the laser safety standards referenced in 88.9.2.

[15]A host system that fails to meet the manufacturer's 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.

88.10 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 88–3.

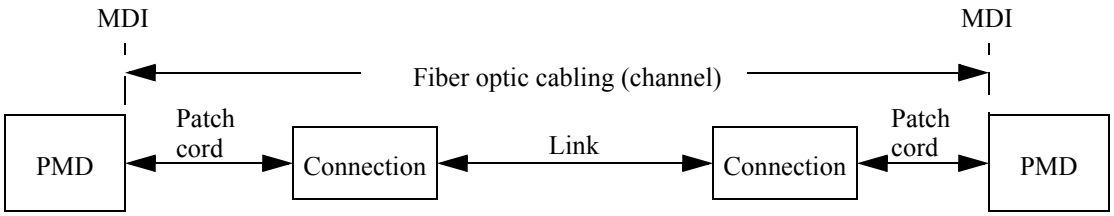

Figure 88–3—Fiber optic cabling model

The channel insertion loss is given in Table 88–14. A channel may contain additional connectors as long as the optical characteristics of the channel, such as attenuation, dispersion, reflections, and polarization mode dispersion meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with 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 88–14—Fiber optic cabling (channel) characteristics

Description	100GBASE-LR4	100GBASE-ER4		Unit
Operating distance (max)	10	30	40	km
Channel insertion loss[a, b] (max)	6.3	18	18	dB
Channel insertion loss (min)	0	0		dB
Positive dispersion[b] (max)	9.5	28	36	ps/nm
Negative dispersion[b] (min)	–28.5	–85	–114	ps/nm
DGD_max[c]	8	10.3	10.3	ps
Optical return loss (min)	21	21	21	dB

[a]These channel insertion loss values include cable, connectors, and splices.
[b]Over the wavelength range 1294.53 nm to 1310.19 nm
[c]Differential Group Delay (DGD) is the time difference at reception between the fractions of a pulse that were 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.

88.11 Characteristics of the fiber optic cabling (channel)

The 100GBASE-LR4 and 100GBASE-ER4 fiber optic cabling shall meet the specifications defined in Table 88–14. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together.

88.11.1 Optical fiber cable

The fiber optic cable requirements are satisfied by cables containing IEC 60793-2-50 type B1.1 (dispersion un-shifted single-mode), type B1.3 (low water peak single-mode), or type B6_a (bend insensitive) fibers or the requirements in Table 88–15 where they differ.

Table 88–15—Optical fiber and cable characteristics

Description	Value	Unit
Nominal fiber specification wavelength	1310	nm
Cabled optical fiber attenuation (max)	0.43[a] or 0.5[b]	dB/km
Zero dispersion wavelength (λ_0)	$1300 \leq \lambda_0 \leq 1324$	nm
Dispersion slope (max) (S_0)	0.093	ps/nm^2 km

[a]The 0.43 dB/km at 1295 nm attenuation for optical fiber cables is derived from Appendix I of ITU-T G.695.
[b]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 for 100GBASE-LR4 or 40 km for 100GBASE-ER4.

88.11.2 Optical fiber connection

An optical fiber connection, as shown in Figure 88–3, consists of a mated pair of optical connectors.

88.11.2.1 Connection insertion loss

The maximum link distance is based on an allocation of 2 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 88–14 are met.

88.11.2.2 Maximum discrete reflectance

The maximum discrete reflectance shall be less than –26 dB.

88.11.3 Medium Dependent Interface (MDI) requirements

The 100GBASE-LR4 or 100GBASE-ER4 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 88–3). 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 IEC 61753-1-1 and IEC 61753-021-2.

NOTE—Transmitter compliance testing is performed at TP2 as defined in 88.5.1, not at the MDI.

88.12 Protocol implementation conformance statement (PICS) proforma for Clause 88, Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-LR4 and 100GBASE-ER4[16]

88.12.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 88, Physical Medium Dependent sublayer and medium, type 100GBASE-LR4 and 100GBASE-ER4, 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.

88.12.2 Identification

88.12.2.1 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).	

88.12.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 88, Physical Medium Dependent sublayer and medium, type 100GBASE-LR4 and 100GBASE-ER4
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	

[16]*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.

88.12.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*LR4	100GBASE-LR4 PMD	88.7	Device supports requirements for 100GBASE-LR4 PHY	O.1	Yes [ ] No [ ]
*ER4	100GBASE-ER4 PMD	88.7	Device supports requirements for 100GBASE-ER4 PHY	O.1	Yes [ ] No [ ]
*INS	Installation / cable	88.10	Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]
CTP1	Reference point TP1 exposed and available for testing	88.5.1	This point may be made available for use by implementers to certify component conformance	O	Yes [ ] No [ ]
CTP4	Reference point TP4 exposed and available for testing	88.5.1	This point may be made available for use by implementers to certify component conformance	O	Yes [ ] No [ ]
CDC	Delay constraints	88.3.1	Device conforms to delay constraints	M	Yes [ ]
CSC	Skew constraints	88.3.2	Device conforms to Skew and Skew Variation constraints	M	Yes [ ]
*MD	MDIO capability	88.4	Registers and interface supported	O	Yes [ ] No [ ]

88.12.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, types 100GBASE-LR4 and 100GBASE-ER4

88.12.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CF1	Compatible with 100GBASE-R PCS and PMA	88.1		M	Yes [ ]
CF2	Integration with management functions	88.1		O	Yes [ ] No [ ]
CF3	Transmit function	88.5.2	Conveys bits from PMD service interface to MDI	M	Yes [ ]
CF4	Optical multiplexing and delivery to the MDI	88.5.2	Optically multiplexes the four optical signal streams for delivery to the MDI	M	Yes [ ]
CF5	Mapping between optical signal and logical signal for transmitter	88.5.2	Higher optical power is a one	M	Yes [ ]
CF6	Receive function	88.5.3	Conveys bits from MDI to PMD service interface	M	Yes [ ]
CF7	Conversion of four optical signals to four electrical signals	88.5.3	For delivery to the PMD service interface	M	Yes [ ]
CF8	Mapping between optical signal and logical signal for receiver	88.5.3	Higher optical power is a one	M	Yes [ ]
CF9	Global Signal Detect function	88.5.4	Report to the PMD service interface the message PMD:IS_SIGNAL.indication(SIGNAL_DETECT)	M	Yes [ ]
CF10	Global Signal Detect behavior	88.5.4	SIGNAL_DETECT is a global indicator of the presence of optical signals on all four lanes	M	Yes [ ]
CF11	Lane-by-lane Signal Detect function	88.5.5	Sets PMD_signal_detect_*i* values on a lane-by-lane basis per requirements of Table 88–4	MD:O	Yes [ ] No [ ] N/A [ ]
CF12	PMD reset function	88.5.6	Resets the PMD sublayer	MD:O	Yes [ ] No [ ] N/A [ ]

88.12.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
CM1	Management register set	88.4		MD:M	Yes [ ] N/A [ ]
CM2	Global transmit disable function	88.5.7	Disables all of the optical transmitters with the PMD_global_transmit_disable variable	MD:O	Yes [ ] No [ ] N/A [ ]
CM3	PMD_lane_by_lane_transmit_ disable function	88.5.8	Disables the optical transmitter on the lane associated with the PMD_transmit_disable_*i* variable	MD:O.2	Yes [ ] No [ ] N/A [ ]
CM4	PMD lane-by-lane transmit disable	88.5.8	Disables each optical transmitter independently if CM3 = No	O.2	Yes [ ] No [ ]
CM5	PMD_fault function	88.5.9	Sets PMD_fault to one if any local fault is detected	MD:O	Yes [ ] No [ ] N/A [ ]
CM6	PMD_transmit_fault function	88.5.10	Sets PMD_transmit_fault to one if a local fault is detected on any transmit lane	MD:O	Yes [ ] No [ ] N/A [ ]
CM7	PMD_receive_fault function	88.5.11	Sets PMD_receive_fault to one if a local fault is detected on any receive lane	MD:O	Yes [ ] No [ ] N/A [ ]

88.12.4.3 PMD to MDI optical specifications for 100GBASE-LR4

Item	Feature	Subclause	Value/Comment	Status	Support
CLR1	Transmitter meets specifications in Table 88–7	88.7.1	Per definitions in 88.8	LR4:M	Yes [ ] N/A [ ]
CLR2	Receiver meets specifications in Table 88–8	88.7.2	Per definitions in 88.8	LR4:M	Yes [ ] N/A [ ]

88.12.4.4 PMD to MDI optical specifications for 100GBASE-ER4

Item	Feature	Subclause	Value/Comment	Status	Support
CER1	Transmitter meets specifications in Table 88–7	88.7.1	Per definitions in 88.8	ER4:M	Yes [ ] N/A [ ]
CER2	Receiver meets specifications in Table 88–8	88.7.2	Per definitions in 88.8	ER4:M	Yes [ ] N/A [ ]

88.12.4.5 Optical measurement methods

Item	Feature	Subclause	Value/Comment	Status	Support
COM1	Measurement cable	88.8	2 m to 5 m in length	M	Yes [ ]
COM2	Center wavelength	88.8.2	Per TIA/EIA-455-127-A or IEC 61280–1–3 under modulated conditions	M	Yes [ ]
COM3	Average optical power	88.8.3	Per IEC 61280-1-1	M	Yes [ ]
COM4	OMA measurements	88.8.4	Each lane	M	Yes [ ]
COM5	Transmitter and dispersion penalty	88.8.5	Each lane	M	Yes [ ]
COM6	Extinction ratio	88.8.6	Per IEC 61280-2-2	M	Yes [ ]
COM7	RIN_{20}OMA measurement procedure	88.8.7	Each lane	M	Yes [ ]
COM8	Transmit eye	88.8.8	Each lane	M	Yes [ ]
COM9	Stressed receiver sensitivity	88.8.10	Each lane	M	Yes [ ]
COM10	Receiver 3 dB electrical upper cutoff frequency	88.8.11	Each lane	M	Yes [ ]

88.12.4.6 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CES1	General safety	88.9.1	Conforms to IEC 60950-1	M	Yes [ ]
CES2	Laser safety—IEC Hazard Level 1	88.9.2	Conforms to Hazard Level 1 laser requirements defined in IEC 60825-1 and IEC 60825-2	M	Yes [ ]
CES3	Electromagnetic interference	88.9.5	Complies with applicable local and national codes for the limitation of electromagnetic interference	M	Yes [ ]

88.12.4.7 Characteristics of the fiber optic cabling and MDI

Item	Feature	Subclause	Value/Comment	Status	Support
COC1	Fiber optic cabling	88.11	Meets requirements specified in Table 88–14	INS:M	Yes [ ] N/A [ ]
COC2	Maximum discrete reflectance	88.11.2.2	Less than –26 dB	INS:M	Yes [ ] N/A [ ]
COC3	MDI requirements	88.11.3	Meets IEC 61753-1-1 and IEC 61753-021-2	INS:M	Yes [ ] N/A [ ]

89. Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-FR

89.1 Overview

This clause specifies the 40GBASE-FR PMD together with the single-mode fiber medium. It is intended that devices compliant with this specification could be implemented to also be used for application VSR2000-3R2 as defined in ITU-T G.693 [B49] (but that is beyond the scope of this standard). The 40GBASE-FR PMD is defined using a specification and test methodology that is similar to that used in ITU-T G.693 [B49], which is different from the specification and test methodologies used for other optical PMDs at 1000 Mb/s and above in IEEE Std 802.3. For example, the transmitter specifications for 40GBASE-FR do not include OMA or TDP parameters but do include a dispersion penalty limit, and the receiver specifications for 40GBASE-FR do not include stressed receiver sensitivity parameters but do include a part-stressed sensitivity limit and a separate jitter tolerance limit. See 89.6.4.

When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 89–1, to the medium through the MDI and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 89–1—Physical Layer clauses associated with the 40GBASE-FR PMD

Associated clause	40GBASE-FR
81—RS	Required
81—XLGMII[a]	Optional
82—PCS for 40GBASE-R	Required
83—PMA for 40GBASE-R	Required
83A—XLAUI	Optional
83B—Chip to module XLAUI	Optional
78—Energy Efficient Ethernet	Optional

[a]The XLGMII is an optional interface. However, if the XLGMII is not implemented, a conforming implementation must behave functionally as though the RS and XLGMII were present.

Figure 89–1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 40 Gb/s and 100 Gb/s Ethernet is introduced in Clause 80 and the purpose of each PHY sublayer is summarized in 80.2.

40GBASE-FR PHYs with the optional Energy Efficient Ethernet (EEE) fast wake capability may enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization (see Clause 78). The deep sleep mode of EEE is not supported.

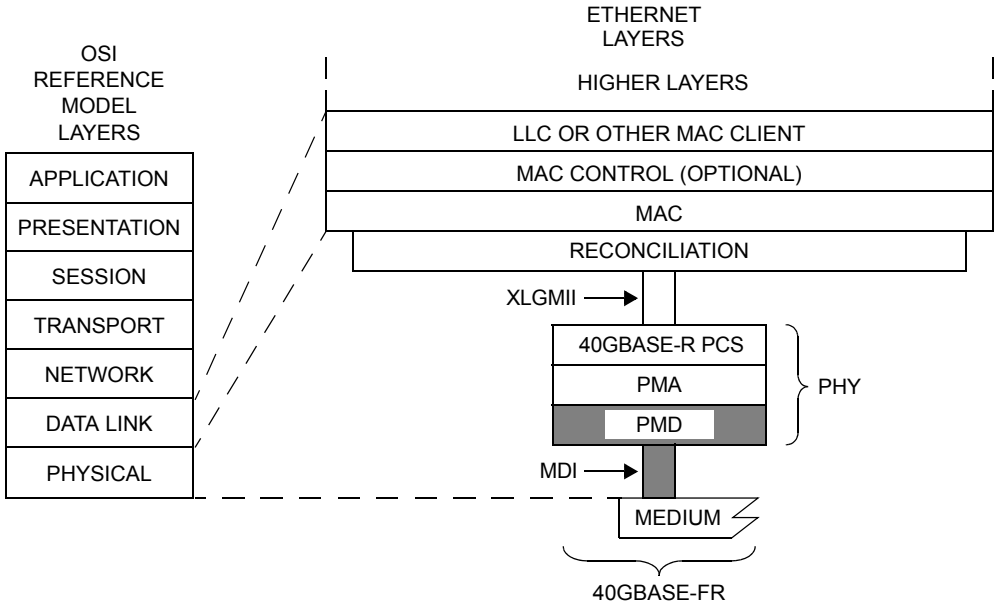

LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE
FR = PMD FOR SINGLE-MODE FIBER — 2 km

Figure 89–1—40GBASE-FR PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

89.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 40GBASE-FR PMD. The service interface for this PMD is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_0.request
PMD:IS_UNITDATA_0.indication
PMD:IS_SIGNAL.indication

In the transmit direction, the PMA continuously sends a single serial stream of bits to the PMD at a nominal signaling rate of 41.25 GBd. The PMD converts this stream of bits into an appropriate signal on the MDI.

In the receive direction, the PMD continuously sends a single serial stream of bits to the PMA corresponding to the signal received from the MDI at a nominal signaling rate of 41.25 GBd.

The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication(SIGNAL_OK) inter-sublayer service primitive defined in 80.3.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the rx_bit parameter is undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_bit parameter 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.

89.3 Delay and skew

89.3.1 Delay constraints

The sum of the transmit and the receive delays at one end of the link contributed by the 40GBASE-FR PMD including 2 m of fiber in one direction shall be no more than 1024 bit times (2 pause_quanta or 25.6 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

89.3.2 Skew constraints

The Skew (relative delay) between the PCS lanes must be kept within limits so that the information on the lanes can be reassembled by the PCS. The Skew Variation must also be limited to ensure that a given PCS lane always traverses the same physical lane. Skew and Skew Variation are defined in 80.5 and specified at the points SP1 to SP6 shown in Figure 80–6 and Figure 80–7.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns as defined by 83.5.3.4. Since the signal at the PMD service interface represents a serial bit stream, there is no Skew Variation at this point.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns. Since the signal at the MDI represents a serial bit stream, there is no Skew Variation at this point.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns. Since the signal at the MDI represents a serial bit stream, there is no Skew Variation at this point.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns. Since the signal at the PMD service interface represents a serial bit stream, there is no Skew Variation at this point.

For more information on Skew and Skew Variation, see 80.5. The measurements of Skew and Skew Variation are defined in 89.7.2.

89.4 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If the MDIO interface is implemented, the mapping of MDIO control variables to PMD control variables shall be as shown in Table 89–2, and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 89–3.

Table 89–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD Transmit disable register	1.9.0	PMD_global_transmit_disable

Table 89–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	PMD_global_signal_detect

89.5 PMD functional specifications

The 40GBASE-FR PMD performs the Transmit and Receive functions, which convey data between the PMD service interface and the MDI.

89.5.1 PMD block diagram

The PMD block diagram is shown in Figure 89–2. 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 single-mode fiber patch cord (TP2), between 2 m and 5 m in length. Unless specified otherwise, all transmitter measurements and tests defined in 89.7 are made at TP2. The optical receive signal is defined at the output of the fiber optic cabling (TP3) at the MDI (see 89.10.3). Unless specified otherwise, all receiver measurements and tests defined in 89.7 are made at TP3.

89.5.2 PMD transmit function

The PMD Transmit function shall convert the bit stream requested by the PMD service interface message PMD:IS_UNITDATA_0.request into an optical signal stream and deliver it to the MDI, according to the transmit optical specifications in this clause. The higher optical power level in the signal stream shall correspond to tx_bit = one.

89.5.3 PMD receive function

The optical signal stream received from the MDI shall be converted into a bit stream for delivery to the PMD service interface using the message PMD:IS_UNITDATA_0.indication, according to the receive optical specifications in this clause. The higher optical power level in the signal stream shall correspond to rx_bit = one.

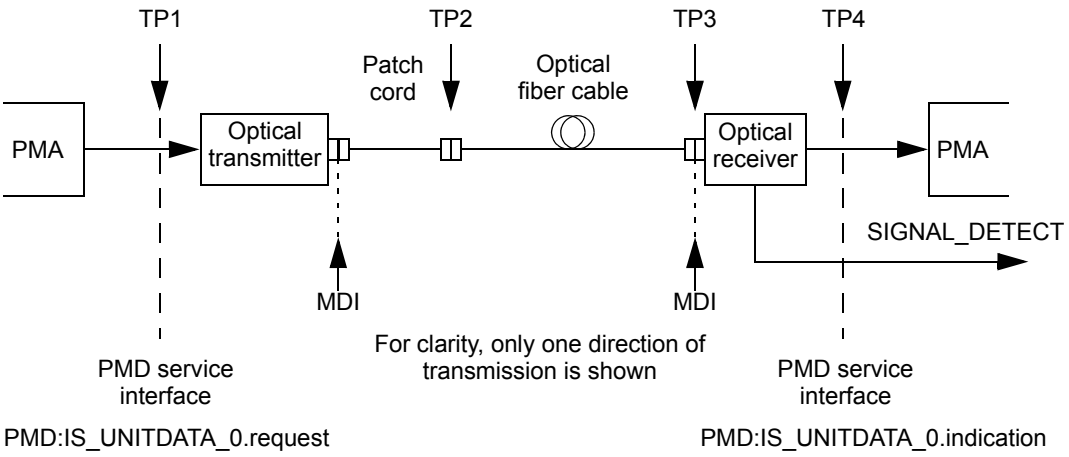

Figure 89–2—Block diagram for 40GBASE-FR transmit/receive paths

89.5.4 PMD global signal detect function

The PMD global 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:IS_SIGNAL.indication message is generated when a change in the value of SIGNAL_DETECT occurs. The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the inter-sublayer service interface primitives defined in 80.3.

SIGNAL_DETECT is an indicator of the presence of an optical signal. The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 89–4. The PMD receiver is not required to verify whether a compliant 40GBASE-R signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

Table 89–4—SIGNAL_DETECT value definition

Receive conditions	SIGNAL_DETECT value
Average optical power at TP3 ≤ –30 dBm	FAIL
(Optical power at TP3 ≥ receiver sensitivity (average power) (max) in Table 89–7) AND (compliant 40GBASE–R 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 including the effects of 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.

89.5.5 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.

89.5.6 PMD global transmit disable function (optional)

The PMD_global_transmit_disable function is optional and allows the optical transmitter to be disabled.

a) When the PMD_global_transmit_disable variable is set to one, this function shall turn off the optical transmitter so that it meets the requirements of the average launch power of the OFF transmitter in Table 89–6.
b) If a PMD_fault is detected, then the PMD may set the PMD_global_transmit_disable to one, turning off the optical transmitter.

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—The PMD lane-by-lane transmit disable function is not used for serial PMDs.

89.5.7 PMD fault function (optional)

If the PMD has detected a local fault on either the transmit or receive path, the PMD shall set PMD_fault to one. If the MDIO interface is implemented, PMD_fault shall be mapped to the fault bit as specified in 45.2.1.2.3.

89.5.8 PMD transmit fault function (optional)

If the PMD has detected a local fault on the transmitter, the PMD shall set the PMD_transmit_fault variable to one. If the MDIO interface is implemented, PMD_transmit_fault shall be mapped to the transmit fault bit as specified in 45.2.1.7.4.

89.5.9 PMD receive fault function (optional)

If the PMD has detected a local fault on the receiver, the PMD shall set the PMD_receive_fault variable to one. If the MDIO interface is implemented, PMD_receive_fault shall be mapped to the receive fault bit as specified in 45.2.1.7.5.

89.6 PMD to MDI optical specifications for 40GBASE-FR

The required operating range for the 40GBASE-FR PMD is defined in Table 89–5. A 40GBASE-FR compliant PMD operates on IEC 60793-2-50 type B1.1, B1.3 or B6_a single-mode fibers according to the specifications defined in Table 89–13. A PMD that exceeds the required operating range while meeting all other optical specifications is considered compliant (e.g., operating at 2.5 km meets the operating range requirement of 2 m to 2 km).

Table 89–5—40GBASE-FR required operating range

PMD type	Required operating range
40GBASE-FR	2 m to 2 km

89.6.1 40GBASE-FR transmitter optical specifications

The 40GBASE-FR transmitter shall meet the specifications defined in Table 89–6 per the definitions in 89.7.

Table 89–6—40GBASE-FR transmit characteristics

Description	Value	Unit
Signaling rate (range)	41.25 ± 100 ppm	GBd
Center wavelength (range)	1530 to 1565	nm
Side-mode suppression ratio (SMSR), (min)	35	dB
Average launch power (max)	3	dBm
Average launch power (min)	0	dBm
Dispersion penalty (max)	2	dB
Average launch power of OFF transmitter (max)	–30	dBm
Extinction ratio (min)	8.2	dB
$RIN_{20}OMA$ (max)	–132	dB/Hz
Optical return loss tolerance (max)	20	dB
Transmitter reflectance[a] (max)	–12	dB
Transmitter eye mask definition {X1, X2, X3, Y1, Y2, Y3}	{0.25, 0.4, 0.45, 0.22, 0.25, 0.4}	

[a]Transmitter reflectance is defined looking into the transmitter.

89.6.2 40GBASE-FR receive optical specifications

The 40GBASE-FR receiver shall meet the specifications defined in Table 89–7 per the definitions in 89.7.

Table 89–7—40GBASE-FR receive characteristics

Description	Value	Unit
Signaling rate (range)	41.25 ± 100 ppm	GBd
Center wavelength (range)	1290 to 1330 and 1530 to 1565	nm
Damage threshold[a] (min)	4	dBm
Average receive power (max)	3	dBm
Receiver reflectance (max)	–26	dB
Receiver sensitivity (average power)[b] (max)	–6	dBm
Receiver jitter tolerance (max)	1	dB
Receiver 3 dB electrical upper cutoff frequency (max)	49	GHz

[a]The receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having this average power level.

[b]Receiver sensitivity (average power) is defined in 89.7.9 and is to be met with a transmitter with worst-case transmit eye, extinction ratio, transmitter reflectance, and $RIN_{20}OMA$. This is a different definition of receiver sensitivity from that used in other clauses (e.g., that in Clause 38). See 89.6.4 for a comparison.

89.6.3 40GBASE-FR illustrative link power budget

An illustrative power budget and penalties for 40GBASE-FR channels are shown in Table 89–8.

Table 89–8—40GBASE-FR illustrative power budget

Parameter	Value	Unit
Power budget[a]	6	dB
Operating distance	2	km
Channel insertion loss[b]	4	dB
Maximum discrete reflectance	–26	dB
Path penalty[a, c]	2	dB
Additional insertion loss allowed	0	dB

[a]This budget does not include the effect of a non-ideal transmitter waveform. See 89.6.4.
[b]The channel insertion loss is calculated using the maximum distance specified in Table 89–5 and cabled optical fiber attenuation of 0.5 dB/km at 1530 nm plus an allocation for connection and splice loss given in 89.10.2.1.
[c]Path penalty is the combined penalty caused by chromatic dispersion and polarization mode dispersion.

89.6.4 Comparison of power budget methodology

This clause uses the budgeting methodology that is used for application VSR2000-3R2 in ITU-T G.693 [B49], which is different from the methodology used in other clauses of this standard (e.g., Clause 38, Clause 52, Clause 86, Clause 87, Clause 88). Figure 89–3 compares the terminology used in this clause with Clause 38. Receiver sensitivity in this clause is specified with a worst-case transmitter input whereas in the other clauses it is specified with a perfect signal without penalties. Stressed receiver sensitivity is not specified in this clause but is specified as the key requirement in the other clauses with a signal that includes both transmitter and link penalties.

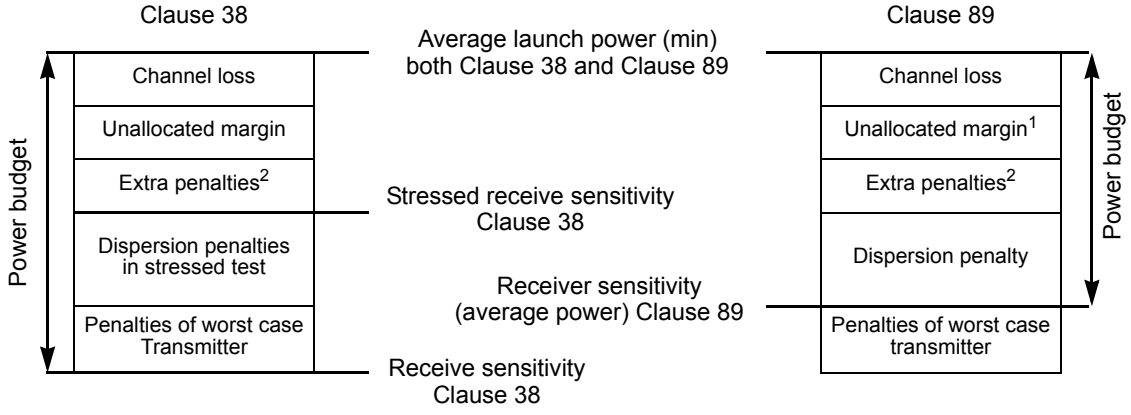

NOTE 1—For Clause 89 unallocated margin is zero and not mentioned in the rest of this clause.
NOTE 2—For Clause 38 extra penalties are any penalties that are not included in the stressed test while for Clause 89 extra penalties are any penalties that are not included in the dispersion penalty test.

Figure 89–3—Optical power budget methodology comparison with Clause 38

89.7 Definition of optical parameters and measurement methods

All transmitter optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified.

89.7.1 Test patterns for optical parameters

Compliance is to be achieved in normal operation. Table 89–10 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. Any of the test patterns given for a particular test in Table 89–10 may be used to perform that test. The test patterns used in this clause are shown in Table 89–9.

Table 89–9—Test patterns

Pattern	Pattern description	Defined in
Square wave	Square wave (8 ones, 8 zeros)	83.5.10
3	PRBS31	83.5.10
4	PRBS9	83.5.10
5	Scrambled idle	82.2.11

Table 89–10—Test-pattern definitions and related subclauses

Parameter	Pattern	Related subclause
Wavelength	3, 5 or valid 40GBASE-R signal	89.7.3
Side mode suppression ratio	3, 5 or valid 40GBASE-R signal	—
Average optical power	3, 5 or valid 40GBASE-R signal	89.7.4
Dispersion penalty	3 or 5	89.7.5
Extinction ratio	3, 5 or valid 40GBASE-R signal	89.7.6
$RIN_{20}OMA$	Square wave or 4	89.7.7
Transmitter optical waveform	3, 5 or valid 40GBASE-R signal	89.7.8
Receiver sensitivity	3 or 5	89.7.9
Receiver jitter tolerance	3 or 5	89.7.10
Receiver 3 dB electrical upper cutoff frequency	3, 5 or valid 40GBASE-R signal	89.7.11

89.7.2 Skew and Skew Variation

Skew and Skew Variation are defined in 80.5 and are required to remain within the limits given in 89.3.2 over the time that the link is in operation. Since the signal for this PMD is a serial bit stream at the PMD service interface and the MDI, there is no Skew Variation at skew points SP2, SP3, SP4, and SP5. The measurement of Skew is made by acquiring the data on the single lane using a clock and data recovery unit with high-frequency corner bandwidth of 16 MHz and a slope of –20 dB/decade. The arrival times of the one to zero transition of the alignment marker sync bits on each PCS lane are then compared. This arrangement ensures that any high-frequency jitter that is present on the signals is not included in the Skew measurement.

89.7.3 Wavelength

The center wavelength shall be within the range given in Table 89–6 if measured per TIA/EIA-455-127-A or IEC 61280-1-3 using the test pattern referenced in Table 89–10.

89.7.4 Average optical power

The average optical power shall be within the limits given in Table 89–6 and Table 89–7 if measured using the methods given in IEC 61280-1-1 using the test pattern referenced in Table 89–10.

89.7.5 Dispersion penalty

The dispersion penalty of a transmitter is defined as the difference in sensitivity (at a BER of 10^{-12}) of a reference receiver as defined in 89.7.5.2 when receiving the signal from that transmitter via the channel defined in 89.7.5.1 compared to the sensitivity obtained when receiving the signal via an attenuator only.

89.7.5.1 Channel requirements

The transmitter is tested using an optical channel that meets the requirements listed in Table 89–11.

Table 89–11—Transmitter compliance channel specifications

PMD type	Dispersion[a] (ps/nm)		Insertion loss[b]	Optical return loss[c]	Max mean DGD
	Minimum	Maximum			
40GBASE-FR	$0.0465 \times \lambda \times [1 - (1324 / \lambda)^4]$	$0.0465 \times \lambda \times [1 - (1300 / \lambda)^4]$	Minimum	20 dB	0.5 ps

[a]The dispersion is measured for the wavelength of the device under test (λ in nm). The coefficient assumes a maximum operating distance of 2 km.
[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 40GBASE-FR 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 89–11 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 IEC 60793-1-42 may be used. The measurement is made in the linear power regime of the fiber.

The channel provides an optical return loss specified in Table 89–11. The state of polarization of the back reflection is adjusted to create the greatest RIN.

The mean DGD of the channel is to be less than the value specified in Table 89–11.

89.7.5.2 Reference receiver requirements

The reference receiver is required to have the bandwidth given in 89.7.8. The sensitivity of the reference receiver is limited by Gaussian noise. The receiver has minimal threshold offset, deadband, hysteresis, baseline wander, deterministic jitter, or other distortions. Decision sampling has minimal uncertainty and setup/hold times.

The clock recovery unit (CRU) used in the dispersion penalty measurement has a corner frequency of 16 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 signal to the clock removes this low-frequency jitter from the measurement.

89.7.5.3 Test procedure

To measure the dispersion penalty, the following procedure shall be used:

a) Configure the test equipment as described above and illustrated in Figure 89–4.
b) With a short patch cable rather than the test fiber and with no reflection, adjust the attenuation of the optical attenuator to obtain a BER of 10^{-12}.
c) Record the optical power at the input to the reference receiver, P_ND, in dBm.
d) With the test fiber and reflection in place, adjust the attenuation of the optical attenuator to obtain a BER of 10^{-12}.
e) Record the optical power at the input to the reference receiver, P_D, in dBm.
f) If P_D is larger than P_ND, the dispersion penalty for the transmitter under test is P_D – P_ND. Otherwise, the dispersion penalty is zero.

It is to be ensured that the measurements are made in the linear power regime of the fiber.

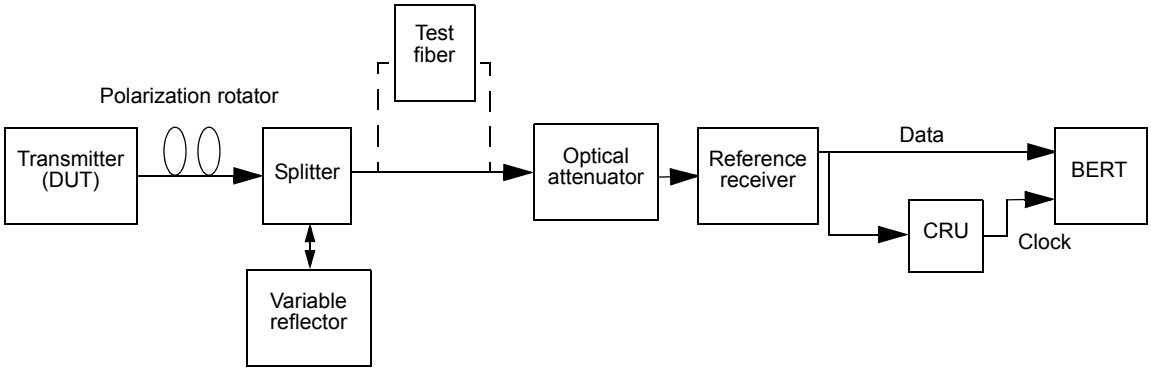

Figure 89–4—Test setup for measurement of dispersion penalty

89.7.6 Extinction ratio

The extinction ratio shall be within the limits given in Table 89–6 if measured using the methods specified in IEC 61280-2-2 using the test pattern referenced in Table 89–10.

NOTE—Extinction ratio and OMA are defined with different test patterns (see Table 89–10).

89.7.7 Relative Intensity Noise ($RIN_{20}OMA$)

The RIN measurement methodology shall be as defined in 52.9.6 with the following exceptions:

a) The optical return loss is 20 dB.
b) The upper –3 dB limit of the measurement apparatus is to be approximately equal to the signaling rate (i.e., 41.25 GHz).

89.7.8 Transmitter optical waveform (transmit eye)

The required optical transmitter pulse shape characteristics are specified in the form of a mask of the transmitter eye diagram as shown in Figure 86–4. The transmitter optical waveform of a port transmitting the test pattern specified in Table 89–10 shall meet specifications according to the methods specified in 86.8.4.6.1 with the filter nominal reference frequency f_r of 30.94 GHz and filter tolerances as specified for STM–64 in ITU–T G.691. Compensation may be made for variation of the reference receiver filter response from an ideal fourth-order Bessel-Thomson response.

89.7.9 Receiver sensitivity

Receiver sensitivity is defined as the minimum value of mean received power at TP3 to achieve a BER of 10^{-12}. This shall be met with a transmitter with worst-case transmit eye, extinction ratio, transmitter reflectance, and $RIN_{20}OMA$. The receiver sensitivity does not have to be met in the presence of dispersion, as this effect is specified separately as the dispersion penalty.

Receiver sensitivity is defined with the transmit section in operation. Pattern 3, or Pattern 5, or a valid 40GBASE–R signal is sent from the transmit section of the PMD under test. The signal being transmitted is asynchronous to the received signal.

89.7.10 Receiver jitter tolerance

Receiver jitter tolerance is defined as the largest increase in received optical power required to maintain a BER of 10^{-10} when the sinusoidal jitter defined in Table 89–12 is applied to the signal from a reference transmitter across the defined range of jitter frequencies.

The reference transmitter is a high-quality instrument-grade device, which can be implemented by a CW laser modulated by a high-performance modulator. The following basic requirements apply:

a) Rise/fall times are less than 10 ps at 20% to 80%.
b) The output optical eye is symmetric and passes the transmitter optical waveform test of 89.7.8.
c) In the center 20% region of the eye, the worst-case vertical eye closure penalty as defined in 52.9.9.2 is less than 0.5 dB.
d) Total Jitter is less than 0.2 UI peak-to-peak.

The receiver jitter tolerance shall be within the limits given in Table 89–7 if measured according to the above definition.

NOTE—Receiver jitter tolerance as defined here is made without additional amplitude stress.

89.7.11 Receiver 3 dB electrical upper cutoff frequency

The receiver 3dB electrical upper cutoff frequency shall be within the limits given in Table 89–7 if measured as described in 52.9.11 with the exception that the optical power level used in 52.9.11 b) is approximately equal to the receiver sensitivity level in Table 89–7.

Table 89–12—Applied sinusoidal jitter

Frequency range	Sinusoidal jitter, peak-to-peak (UI)
$f < 480$ kHz	Not specified
480 kHz $< f \leq$ 16 MHz	$2.88 \times 10^6 / f$
16 MHz $< f <$ 10 LB[a]	0.18

[a]LB = loop bandwidth; upper frequency bound for added sinusoidal jitter should be at least 10 times the loop bandwidth of the receiver being tested.

89.8 Safety, installation, environment, and labeling

89.8.1 General safety

All equipment subject to this clause shall conform to IEC 60950-1.

89.8.2 Laser safety

40GBASE-FR 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 explicitly defines requirements and usage restrictions on the host system necessary to meet these safety certifications.[17]

89.8.3 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.

89.8.4 Environment

Normative specifications in Clause 89 shall be met by a system integrating a 40GBASE-FR 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 Clause 89 will be met.

[17]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.

89.8.4.1 Electromagnetic emission

A system integrating a 40GBASE-FR PMD shall comply with applicable local and national codes for the limitation of electromagnetic interference.

89.8.4.2 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.

89.8.5 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., 40GBASE-FR).

Labeling requirements for Hazard Level 1 lasers are given in the laser safety standards referenced in 89.8.2.

89.9 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 89–5.

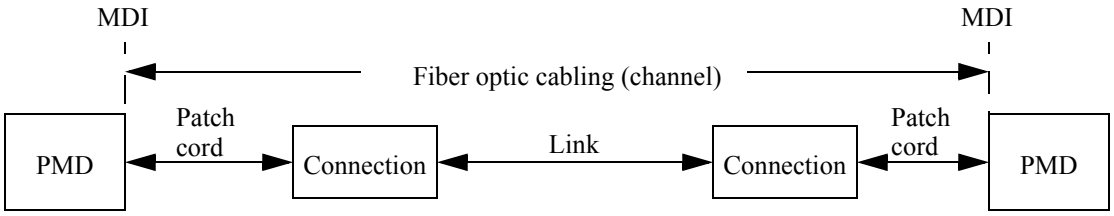

Figure 89–5—Fiber optic cabling model

The channel insertion loss is given in Table 89–13. A channel may contain additional connectors as long as the optical characteristics of the channel, such as attenuation, dispersion, reflections, and polarization mode dispersion, meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with ANSI/TIA/EIA-526-7 method A–1, chromatic dispersion is measured in accordance with IEC 60793-1-42, and polarization mode dispersion is measured in accordance with IEC 60793-1-48. 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.

89.10 Characteristics of the fiber optic cabling (channel)

The 40GBASE-FR fiber optic cabling shall meet the specifications defined in Table 89–13. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together.

89.10.1 Optical fiber cable

The fiber optic cable requirements are satisfied by cables containing IEC 60793-2-50 type B1.1 (dispersion un-shifted single-mode), type B1.3 (low water peak single-mode), or type B6_a (bend insensitive) fibers or the requirements in Table 89–14 where they differ.

Table 89–13—Fiber optic cabling (channel) characteristics for 40GBASE-FR

Description	Value	Unit
Operating distance (max)	2	km
Channel insertion loss[a, b] (max)	4	dB
Channel insertion loss (min)	0	dB
Dispersion[b] (max)	38	ps/nm
DGD_max[c]	3	ps
Optical return loss (min)	21	dB

[a]These channel insertion loss values include cable, connectors, and splices.
[b]Over the wavelength range 1530 nm to 1565 nm.
[c]Differential Group Delay (DGD) is the time difference at reception between the fractions of a pulse that were 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.

Table 89–14—Optical fiber and cable characteristics for 40GBASE-FR

Description	Value	Unit
Nominal fiber specification wavelength	1550	nm
Cabled optical fiber attenuation (max)	0.5[a]	dB/km
Zero dispersion wavelength (λ_0)	$1300 \leq \lambda_0 \leq 1324$	nm
Dispersion slope (max) (S_0)	0.093	ps/nm^2 km

[a]The 0.5 dB/km attenuation is provided for Outside Plant cable as defined in ANSI/TIA-568-C.3-2008.

89.10.2 Optical fiber connection

An optical fiber connection, as shown in Figure 89–5, consists of a mated pair of optical connectors.

89.10.2.1 Connection insertion loss

The maximum link distances for single-mode fiber are calculated based on an allocation of 3 dB total connection and splice loss. For example, this allocation supports six 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 89–13 are met.

89.10.2.2 Maximum discrete reflectance

The maximum discrete reflectance shall be less than –26 dB.

89.10.3 Medium Dependent Interface (MDI) requirements

The 40GBASE-FR 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 89–5). 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 IEC 61753-1-1 and IEC 61753-021-2.

NOTE—Transmitter compliance testing is performed at TP2 as defined in 89.5.1, not at the MDI.

89.11 Protocol implementation conformance statement (PICS) proforma for Clause 89, Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-FR[18]

89.11.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 89, Physical Medium Dependent sublayer and medium, type 40GBASE-FR, 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.

89.11.2 Identification

89.11.2.1 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).	

89.11.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 89, Physical Medium Dependent sublayer and medium, type 40GBASE-FR
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.

89.11.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*INS	Installation / cable	89.9	Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]
XLDC	Delay constraints	89.3.1	Device conforms to delay constraints	M	Yes [ ]
XLSC	Skew constraints	89.3.2	Device conforms to Skew and Skew Variation constraints	M	Yes [ ]
*MD	MDIO capability	89.4	Registers and interface supported	O	Yes [ ] No [ ]

89.11.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-FR

89.11.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
XLF1	Compatible with 40GBASE-R PCS and PMA	89.1		M	Yes [ ]
XLF2	Integration with management functions	89.1		O	Yes [ ] No [ ]
XLF3	Transmit function	89.5.2	Conveys bits from PMD service interface to MDI	M	Yes [ ]
XLF4	Mapping between optical signal and logical signal for transmitter	89.5.2	Higher optical power is a one	M	Yes [ ]
XLF5	Receive function	89.5.3	Conveys bits from MDI to PMD service interface	M	Yes [ ]
XLF6	Mapping between optical signal and logical signal for receiver	89.5.3	Higher optical power is a one	M	Yes [ ]
XLF7	Global Signal Detect function	89.5.4	Report to the PMD service interface the message PMD:IS_SIGNAL.indication(SIGNAL_DETECT)	M	Yes [ ]
XLF8	Global Signal Detect behavior	89.5.4	SIGNAL_DETECT is an indicator of the presence of an optical signal	M	Yes [ ]
XLF9	PMD_reset function	89.5.5	Resets the PMD sublayer	MD:O	Yes [ ] No [ ] N/A [ ]

89.11.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
XLM1	Management register set	89.4		MD:M	Yes [ ] N/A [ ]
XLM2	Global transmit disable function	89.5.6	Disables the optical transmitter with the PMD_global_transmit_disable variable	MD:O	Yes [ ] No [ ] N/A [ ]
XLM3	PMD_fault function	89.5.7	Sets PMD_fault to one if any local fault is detected	MD:O	Yes [ ] No [ ] N/A [ ]
XLM4	PMD_transmit_fault function	89.5.8	Sets PMD_transmit_fault to one if a local transmitter fault is detected	MD:O	Yes [ ] No [ ] N/A [ ]
XLM5	PMD_receive_fault function	89.5.9	Sets PMD_receive_fault to one if a local receiver fault is detected	MD:O	Yes [ ] No [ ] N/A [ ]

89.11.4.3 PMD to MDI optical specifications for 40GBASE-FR

Item	Feature	Subclause	Value/Comment	Status	Support
XLLR1	Transmitter meets specifications in Table 89–6	89.6.1	Per definitions in 89.7	M	Yes [ ] N/A [ ]
XLLR2	Receiver meets specifications in Table 89–7	89.6.2	Per definitions in 89.7	M	Yes [ ] N/A [ ]

89.11.4.4 Optical measurement methods

Item	Feature	Subclause	Value/Comment	Status	Support
XLOM1	Measurement cable	89.7	2 m to 5 m in length	M	Yes [ ]
XLOM2	Center wavelength	89.7.3	Per TIA-455-127-A or IEC 61280-1-3	M	Yes [ ]
XLOM3	Average optical power	89.7.4	Per IEC 61280-1-1	M	Yes [ ]
XLOM4	Dispersion penalty	89.7.5.3		M	Yes [ ]
XLOM5	Extinction ratio	89.7.6	Per IEC 61280-2-2	M	Yes [ ]
XLOM6	$RIN_{20}OMA$	89.7.7		M	Yes [ ]
XLOM7	Transmit eye	89.7.8		M	Yes [ ]
XLOM8	Receiver sensitivity	89.7.9		M	Yes [ ]
XLOM9	Receiver jitter tolerance	89.7.10		M	Yes [ ]
XLOM10	Receiver 3 dB electrical upper cutoff frequency	89.7.11		M	Yes [ ]

89.11.4.5 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
XLES1	General safety	89.8.1	Conforms to IEC 60950-1	M	Yes [ ]
XLES2	Laser safety—IEC Hazard Level 1	89.8.2	Conforms to Hazard Level 1 laser requirements defined in IEC 60825-1 and IEC 60825-2	M	Yes [ ]
XLES3	Electromagnetic interference	89.8.4.1	Complies with applicable codes for the limitation of electromagnetic interference	M	Yes [ ]

89.11.4.6 Characteristics of the fiber optic cabling and MDI

Item	Feature	Subclause	Value/Comment	Status	Support
XLOC1	Fiber optic cabling	89.10	Meets the requirements specified in Table 89–13	INS:M	Yes [ ] N/A [ ]
XLOC2	Maximum discrete reflectance	89.10.2.2	Less than –26 dB	INS:M	Yes [ ] N/A [ ]
XLOC3	MDI requirements	89.10.3	Meets the interface performance specifications listed	INS:M	Yes [ ] N/A [ ]

90. Ethernet support for time synchronization protocols

90.1 Introduction

This clause specifies the optional Time Synchronization Service Interface (TSSI). The TSSI can be used to support protocols that require knowledge of packet egress and ingress time.

The TSSI is defined for the full-duplex mode of operation only. It supports MAC operation at various data rates. The MII (Clause 22), GMII (Clause 35), XGMII (Clause 46), XLGMII (Clause 81) and CGMII (Clause 81) specifications are all compatible with the gRS sublayer defined in 90.5.

90.2 Overview

The goal of this clause is to provide an accurate indication of the transmission and reception initiation times of all packets, as required to support various time synchronization protocols, e.g., IEEE Std 1588-2008 [B46], and IEEE Std 802.1AS [B43].

The specific goals are to:

a) define the Time Synchronization Service Interface (TSSI), with associated service primitives;
b) add management registers to indicate the maximum and minimum data delays for estimation of link latency at the Time Synchronization Protocol (TimeSync) Client;

90.3 Relationship with other IEEE standards

Per 90.2, the TimeSync capability provides support for various time synchronization protocols, including e.g., IEEE Std 1588-2008 or IEEE Std 802.1AS. The definition of TimeSync Client, its capabilities and functions is outside the scope of IEEE Std 802.3.

90.4 Time Synchronization Service Interface (TSSI)

90.4.1 Introduction

This subclause specifies services provided by an extension to the Reconciliation Sublayers specified elsewhere in this standard.

90.4.1.1 Interlayer service interfaces

Figure 90–1 depicts the TimeSync Client and the RS interlayer service interfaces.

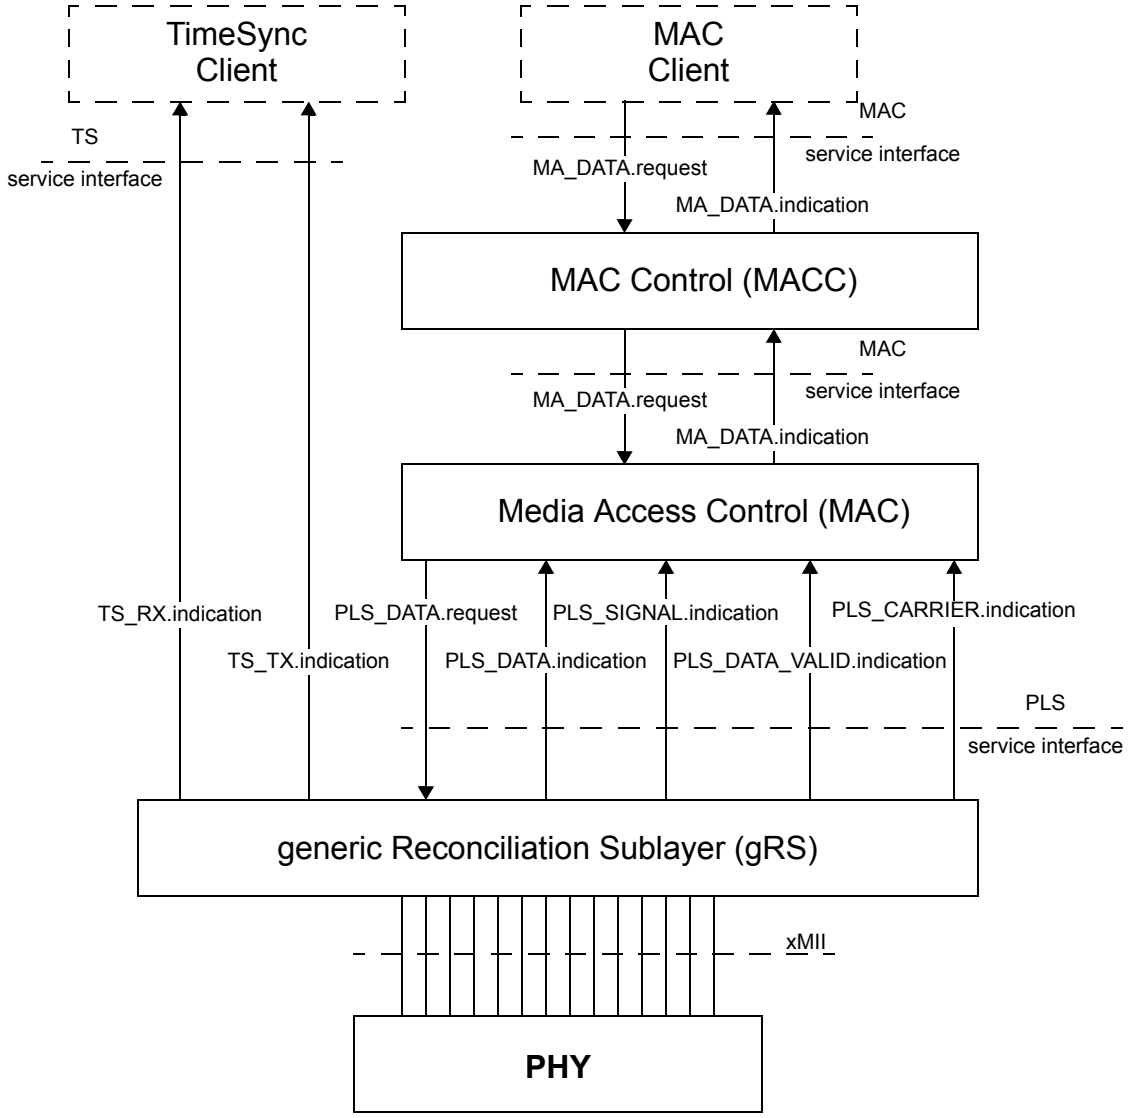

Figure 90–1—Relationship of the TimeSync Client, TSSI and gRS sublayer relative to MAC and MAC Client and associated interfaces

NOTE 1—In this figure, the xMII is used as a generic term for the Media Independent Interfaces for implementations of 100 Mb/s and above. For example: for 100 Mb/s implementations this interface is called MII; for 1 Gb/s implementations, it is called GMII; for 10 Gb/s implementations, it is called XGMII; etc.

NOTE 2—Optional Low Power Idle (LPI) Client service interface not shown.

90.4.1.2 Responsibilities of TimeSync Client

The TimeSync Client can use the indication of the event corresponding to the egress and ingress of packets provided by the TSSI, combined with knowledge of the protocol frames, to select the egress and ingress times relevant to the protocol. Which frames are of interest to any particular protocol is beyond the scope of this standard.

The TimeSync Client can use the indication of the event corresponding to the egress and ingress of packets at the xMII provided by the TSSI, together with the information provided by the TimeSync PHY transmit data delay and TimeSync PHY receive data delay (see 30.13.1).

90.4.2 TSSI

The following specifies the service interface provided by the RS to the TimeSync Client. These services are described in an abstract manner and do not imply any particular implementation. The model used in this service specification is identical to that used in 1.2.2.

The following primitives are defined:

— TS_TX.indication
— TS_RX.indication

90.4.3 Detailed service specification

90.4.3.1 TS_TX.indication primitive

This primitive defines the transfer of the indication of an event between gRS and the TimeSync Client.

90.4.3.1.1 Semantics

The semantics of the primitive are as follows:

TS_TX.indication(SFD)

The SFD parameter can take only one possible value, DETECTED. When asserted (SFD = DETECTED), the TimeSync Client is notified that a valid SFD was detected by the gRS sublayer TS_SFD_Detect_TX function (see 90.5.1) in the xMII transmit signals.

90.4.3.1.2 Condition for generation

This primitive is generated by the gRS sublayer in response to detection of a valid SFD in the data stream transmitted across the xMII transmit signals. Specific conditions for generation of this primitive are described in 90.5.1.

90.4.3.1.3 Effect of receipt

The receipt of this primitive by the TimeSync Client is undefined.

90.4.3.2 TS_RX.indication primitive

This primitive defines the transfer of the indication of an event between gRS and the TimeSync Client.

90.4.3.2.1 Semantics

The semantics of the primitive are as follows:

TS_RX.indication(SFD)

The SFD parameter can take only one possible value, DETECTED. When asserted (SFD = DETECTED), the TimeSync Client is notified that a valid SFD was detected by the gRS sublayer TS_SFD_Detect_RX function (see 90.5.2) in the xMII receive signals.

90.4.3.2.2 Condition for generation

This primitive is generated by the gRS sublayer in response to detection of a valid SFD in the data stream received across the xMII receive signals. Specific conditions for generation of this primitive are described in 90.5.2.

90.4.3.2.3 Effect of receipt

The receipt of this primitive by the TimeSync Client is undefined.

90.5 generic Reconciliation Sublayer (gRS)

Within the scope of this clause, the term generic Reconciliation Sublayer (gRS) is used to denote any IEEE 802.3 Reconciliation Sublayer (RS) used to interface a MAC with any PHY supporting the TimeSync capability through the xMII.

For the purpose of the TimeSync capabilities, two new functions are defined in this subclause, namely TS_SFD_Detect_TX (see 90.5.1) and TS_SFD_Detect_RX (see 90.5.2), which are responsible for generation of the TS_TX.indication and TS_RX.indication service primitives, as defined in 90.4. Figure 90–2 presents the TS_SFD_Detect_TX and TS_SFD_Detect_RX functions and their location within the RS sublayer.

90.5.1 TS_SFD_Detect_TX function

The TS_SFD_Detect_TX function observes the xMII transmit signals and detects the occurrence of the Start Frame Delimiter (SFD, see 3.1.1 and 3.2.2) in compliance with the specifications of the given type of the instantiated xMII.

The service primitive across the TSSI i.e. TS_TX.indication shall be generated only when the SFD is detected on the transmit signals of the xMII (SFD=DETECTED).

90.5.2 TS_SFD_Detect_RX function

The TS_SFD_Detect_RX function observes the xMII receive signals and detects the occurrence of the Start Frame (SFD, see 3.1.1 and 3.2.2) in compliance with the specifications of the given type of the instantiated xMII.

The service primitive across the TSSI i.e. TS_RX.indication shall be generated only when the SFD is detected on the receive signals of the xMII (SFD=DETECTED).

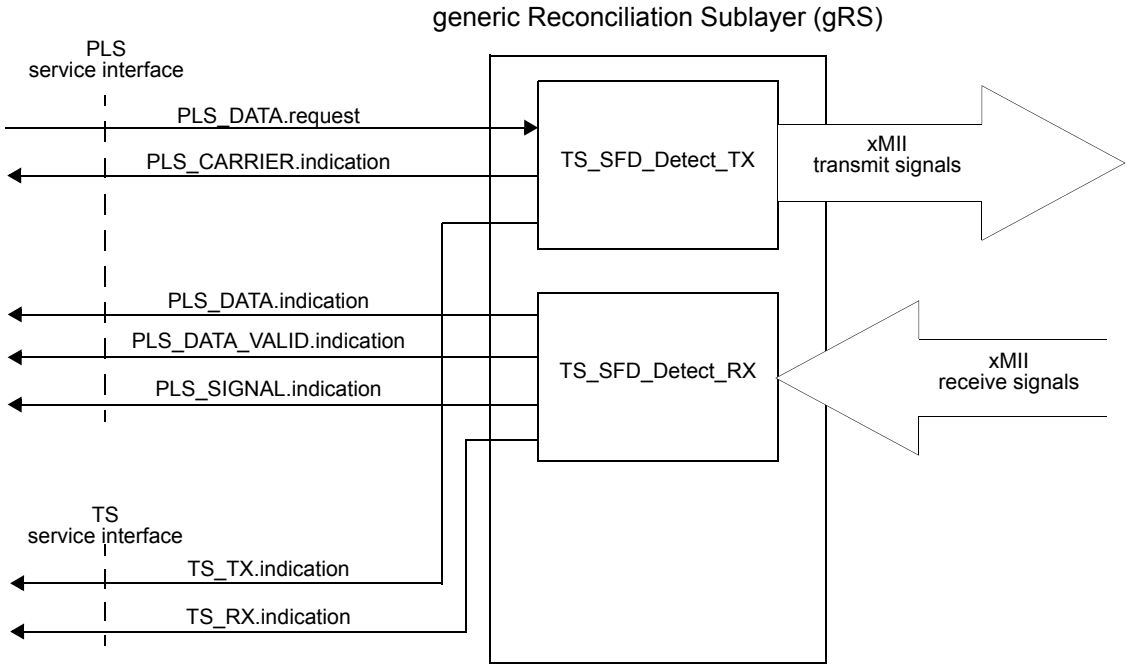

Figure 90–2—TS_SFD_Detect_TX and TS_SFD_Detect_RX functions within the generic Reconciliation Sublayer (gRS)

90.6 Overview of management features

Clause 30 describes the management functions for any of the IEEE Std 802.3 compliant PHYs. Objects defined in Clause 30 for the support of TimeSync capability are summarized below:

— oTimeSync managed object class, as defined in 30.13.1

— aTimeSyncCapabilityTX and aTimeSyncCapabilityRX managed objects, reflecting the status of a series of MDIO capability registers (1.1800, 2.1800, 3.1800, 4.1800, 5.1800, and 6.1800), as defined in 30.13.1.1 and 30.13.1.2, respectively

— aTimeSyncDelayTXmax and aTimeSyncDelayTXmin managed objects, representing the aggregate value of the series of transmit path data delay registers present in the instantiated MDIO registers (for each MMD, in case of multiple instances) 1.1801 through 1.1804, 2.1801 through 2.1804, 3.1801 through 3.1804, 4.1801 through 4.1804, 5.1801 through 5.1804, and 6.1801 through 6.1804, as defined in 30.13.1.3 and 30.13.1.4, respectively

— aTimeSyncDelayRXmax and aTimeSyncDelayRXmin managed objects, representing the aggregate value of the series of receive path data delay registers present in the instantiated MDIO registers (for each MMD, in case of multiple instances) 1.1805 through 1.1808, 2.1805 through 2.1808, 3.1805 through 3.1808, 4.1805 through 4.1808, 5.1805 through 5.1808, and 6.1805 through 6.1808, as defined in 30.13.1.5 and 30.13.1.6, respectively

The Management Data Input/Output (MDIO) capability described in Clause 45 defines several variables that provide TimeSync status information for the PMD, as shown in Table 90–1:

Table 90–1—Summary of TimeSync features in Clause 45

Register	Name	Reference
1.1800	TimeSync PMA/PMD capability register	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
2.1800	TimeSync WIS capability register	45.2.2.20
2.1801 through 2.1804	TimeSync WIS transmit path data delay	45.2.2.21
2.1805 through 2.1808	TimeSync WIS receive path data delay	45.2.2.22
3.1800	TimeSync PCS capability register	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
4.1800	TimeSync PHY XS capability register	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
5.1800	TimeSync DTE XS capability register	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
6.1800	TimeSync TC capability register	45.2.6.14
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

90.7 Data delay measurement

The TimeSync capability requires measurement of data delay in the transmit and receive paths, as shown in Figure 90–3. The transmit path data delay is measured from the input of the beginning of the SFD at the xMII to its presentation by the PHY to the MDI. The receive path data delay is measured from the input of the beginning of the SFD at the MDI to its presentation by the PHY to the xMII.

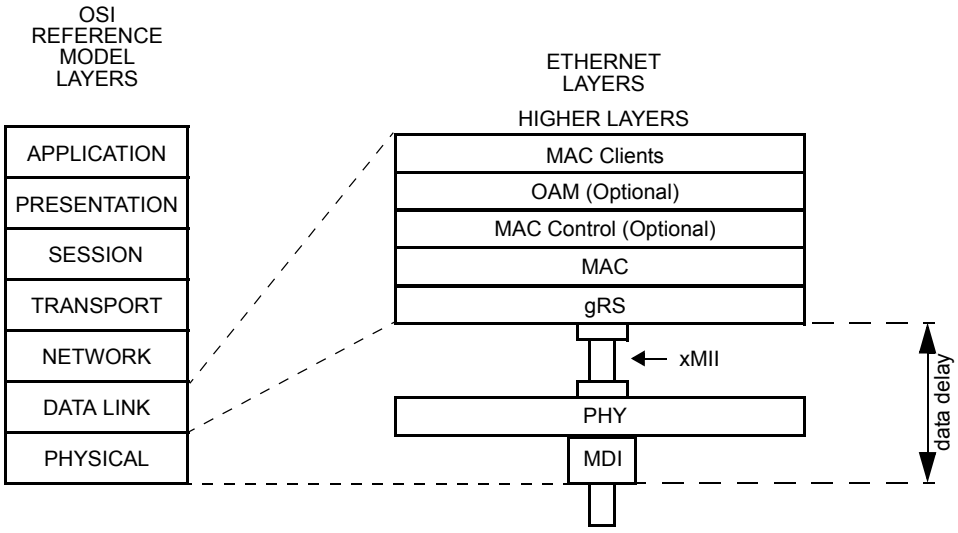

Figure 90–3—Data delay measurement

NOTE 1—The measured data delay in either the transmit or receive path should not include delay resulting from any length of the medium, regardless of the type of the medium used by the given PHY. It is observed that even if any medium-imposed delay is included in the measured data delay, it does not affect the operation of the TimeSync functions and only shifts the timing reference point farther from the MDI.

The obtained data delay measurement shall be reported in the form of a quartet of values; the maximum transmit data delay, the minimum transmit data delay, the maximum receive data delay, and the minimum receive data delay, as defined for the oTimeSync managed object class (30.13.1).

NOTE 2—The data delay values represent only the data delay in the PHY sublayers. The TimeSync Client may need to adjust for delays within the gRS. For example, the TimeSync Client may need to subtract the value of the delay associated with the TS_SFD_Detect_TX function from the sum of the minimum transmit data delay values reported by individual MMD(s). Likewise, the TimeSync Client may need to add the value of the delay associated with the TS_SFD_Detect_RX function to the sum of the maximum receive data delay values reported by individual MMD(s).

90.8 Protocol implementation conformance statement (PICS) proforma for Clause 90, Ethernet support for time synchronization protocols[19]

90.8.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 90 Ethernet support for time synchronization protocols, 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.

90.8.2 Identification

90.8.2.1 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).	

90.8.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 90, Ethernet support for time synchronization protocols
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	

[19] *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.

90.8.3 TSSI indication

Item	Feature	Subclause	Value/Comment	Status	Support
TS_TX	TS_TX.indication generation	90.5.1	TS_TX.indication is generated only when the SFD is detected on the transmit signals of the xMII.	M	Yes [ ]
TS_RX	TS_RX.indication generation	90.5.2	TS_RX.indication is generated only when the SFD is detected on the receive signals of the xMII.	M	Yes [ ]

90.8.4 Data delay reporting

Item	Feature	Subclause	Value/Comment	Status	Support
Report	Data delay reporting format	90.7	Data delay measurement is reported in the form of a quartet of values; the maximum transmit data delay, the minimum transmit data delay, the maximum receive data delay, and the minimum receive data delay, for each MMD instantiated in the device.	M	Yes [ ]

91. Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 100GBASE-R PHYs

91.1 Overview

91.1.1 Scope

This clause specifies a Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 100GBASE-R PHYs. Annex 91A provides examples of RS-FEC codewords constructed with the method specified in this clause.

91.1.2 Position of RS-FEC in the 100GBASE-R sublayers

Figure 91–1 shows the relationship of the RS-FEC sublayer to the ISO/IEC Open System Interconnection (OSI) reference model.

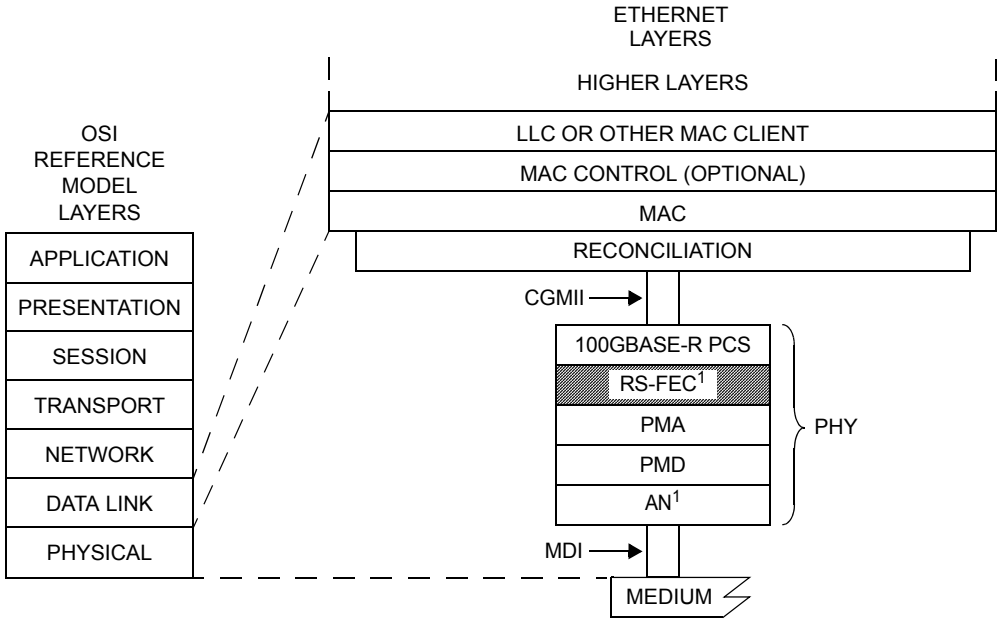

AN = AUTO-NEGOTATION
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION

NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 91–1—RS-FEC relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

91.2 FEC service interface

This subclause specifies the services provided by the RS-FEC sublayer. The service interface is described in an abstract manner and does not imply any particular implementation.

The FEC service interface is provided to allow the PCS to transfer information to and from the RS-FEC. The PCS may be connected to the RS-FEC using an optional instantiation of the PMA service interface (refer to Annex 83A or Annex 83D) in which case a PMA is the client of the FEC service interface.

The FEC service interface is an instance of the inter-sublayer service interface defined in 80.3. The FEC service interface primitives are summarized as follows:

FEC:IS_UNITDATA_*i*.request
FEC:IS_UNITDATA_*i*.indication
FEC:IS_SIGNAL.indication

The RS-FEC operates on 20 parallel bit streams, hence $i = 0$ to 19. The PCS (or PMA) continuously sends 20 parallel bit streams to the RS-FEC, one per lane, each at a nominal signaling rate of 5.15625 GBd. The RS-FEC continuously sends 20 parallel bit streams to the PCS (or PMA), one per lane, each at a nominal signaling rate of 5.15625 GBd

The SIGNAL_OK parameter of the FEC:IS_SIGNAL.indication primitive can take one of two values: OK or FAIL. The value is set to OK when the FEC receive function has identified codeword boundaries as indicated by fec_align_status equal to true. That value is set to FAIL when the FEC receive function is unable to reliably establish codeword boundaries as indicated by fec_align_status equal to false. When SIGNAL_OK is FAIL, the rx_bit parameters of the FEC:IS_UNITDATA_*i*.indication primitives are undefined.

If the optional EEE deep sleep capability is supported, then the FEC service interface includes four additional primitives as follows:

FEC:IS_TX_MODE.request
FEC:IS_RX_MODE.request
FEC:IS_RX_TX_MODE.indication
FEC:IS_ENERGY_DETECT.indication

When the tx_mode parameter of the FEC:IS_TX_MODE.request primitive is QUIET or ALERT, the RS-FEC sublayer may disable transmit functional blocks to conserve energy. Otherwise the RS-FEC transmit function operates normally. The value of tx_mode is passed to the client sublayer via the PMA:IS_TX_MODE.request primitive.

When the rx_mode parameter of the FEC:IS_RX_MODE.request primitive is QUIET, the RS-FEC sublayer may disable receive functional blocks to conserve energy. Otherwise the RS-FEC receive function operates normally. The value of rx_mode is passed to the client sublayer via the PMA:IS_RX_MODE.request primitive.

The rx_tx_mode parameter of the FEC:IS_RX_TX_MODE.indication primitive is used to communicate the link partner's value of tx_mode as inferred by the PMA. It is assigned the value that is received via the PMA:IS_RX_TX_MODE.indication primitive.

The energy_detect parameter of the FEC:IS_ENERGY_DETECT.indication primitive is used to communicate that the PMD has detected the return of energy on the interface following a period of quiescence. It is assigned the value that is received via the PMA:IS_ENERGY_DETECT.indication primitive.

91.3 PMA compatibility

The RS-FEC sublayer requires that the PMA service interface consist of exactly four upstream lanes and exactly four downstream lanes. Therefore, the RS-FEC sublayer may be a client of the PMA sublayer

defined in Clause 83 when the PMA service interface width, *p*, is set to 4. The RS-FEC sublayer may also be a client of the PMA sublayer defined in Clause 94.

In addition, all PMA service interfaces between the RS-FEC sublayer and the PMD sublayer are required to consist of four or fewer upstream lanes and four or fewer downstream lanes. A consequence of this constraint is that a physical instantiation of the ten-lane PMA service interface (CAUI-10) may not be used below the RS-FEC sublayer.

91.4 Delay constraints

The maximum delay contributed by the RS-FEC sublayer (sum of transmit and receive delays at one end of the link) shall be no more than 40960 bit times (80 pause_quanta or 409.6 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

91.5 Functions within the RS-FEC sublayer

91.5.1 Functional block diagram

A functional block diagram of the RS-FEC sublayer is shown in Figure 91–2.

91.5.2 Transmit function

91.5.2.1 Lane block synchronization

The RS-FEC transmit function forms 20 bit streams by concatenating the bits from each of the 20 FEC:IS_UNITDATA_*i*.request primitives in the order they are received. It obtains lock to the 66-bit blocks in each bit stream using the sync headers and outputs 66-bit blocks. Block lock is obtained as specified in the block lock state diagram shown in Figure 82–12.

91.5.2.2 Alignment lock and deskew

Once the RS-FEC transmit function achieves block lock on a PCS lane, it then begins obtaining alignment marker lock as specified by the alignment marker lock state diagram shown in Figure 82–13. This process identifies the PCS lane number received on a particular lane of the service interface. After alignment marker lock is achieved on all 20 lanes, all inter-lane Skew is removed as specified by the PCS deskew state diagram shown in Figure 82–14. The RS-FEC transmit function shall support a maximum Skew of 49 ns between PCS lanes and a maximum Skew Variation of 400 ps. Skew and Skew Variation are defined in 80.5.

91.5.2.3 Lane reorder

PCS lanes can be received on different lanes of the service interface from which they were originally transmitted due to Skew between lanes and multiplexing by the PMA. The RS-FEC transmit function shall order the PCS lanes according to the PCS lane number.

91.5.2.4 Alignment marker removal

After all PCS lanes are aligned and deskewed, the PCS lanes are multiplexed together in the proper order to reconstruct the original stream of blocks and the alignment markers are removed from the data stream. Note that an alignment marker is always removed when am_lock is true for a given PCS lane even if it does not match the expected alignment marker value (due to a bit error for example). Repeated alignment marker

errors result in am_lock being set to false for a given PCS lane, but until that happens it is sufficient to remove the block in the alignment marker position.

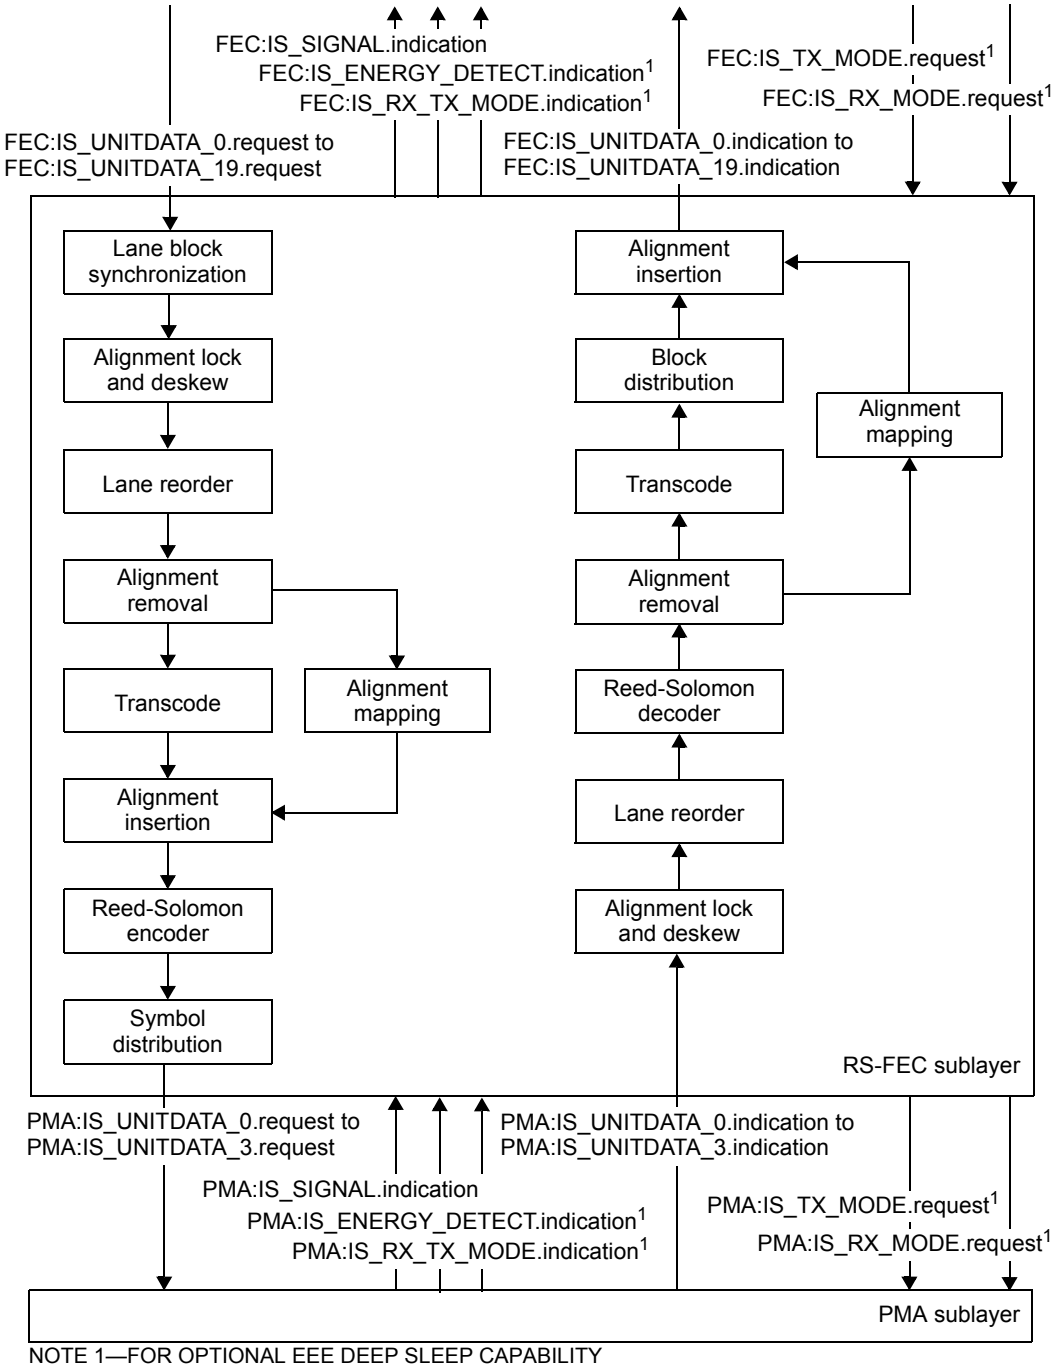

Figure 91–2—Functional block diagram

For the optional EEE deep sleep capability, transitions between normal alignment markers and Rapid Alignment Markers (RAMs) result in changes in relative position and frequency of alignment markers. These transitions are detected by the Transmit LPI state diagram (see Figure 91–10), and this information is used by the alignment marker removal function to determine which 66-bit blocks are to be removed.

As part of the alignment marker removal process, the BIP_3 field is compared to the calculated Bit Interleaved Parity (BIP) value (see 82.2.8) for each PCS lane. If a Clause 45 MDIO is implemented, then the appropriate BIP error counter register (registers 1.230 to 1.249) is incremented by one each time the calculated BIP value does not equal the value received in the BIP_3 field. The incoming bit error ratio can be estimated by dividing the BIP block error ratio by a factor of 1081344.

91.5.2.5 64B/66B to 256B/257B transcoder

The transcoder constructs a 257-bit block, tx_scrambled<256:0>, from a group of four 66-bit blocks, tx_coded_*j*<65:0> where *j*=0 to 3. For each group of four 66-bit blocks, *j*=3 corresponds to the most recently received block. Bit 0 in each 66-bit block is the first bit received and corresponds to the first bit of the synchronization header.

If for all *j*=0 to 3, tx_coded_*j*<0>=0 and tx_coded_*j*<1>=1, tx_xcoded<256:0> shall be constructed as follows:

a) tx_xcoded<0> = 1
b) tx_xcoded<(64*j*+64):(64*j*+1)> = tx_coded_*j*<65:2> for *j*=0 to 3

If for all *j*=0 to 3, tx_coded_*j*<0> ≠ tx_coded_*j*<1> (valid synchronization header) and for any *j*=0 to 3, tx_coded_*j*<0>=1 and tx_coded_*j*<1>=0, tx_xcoded<256:0> shall be constructed as follows:

a1) tx_xcoded<0> = 0
b1) tx_xcoded<*j*+1> = tx_coded_*j*<1> for *j*=0 to 3
c1) Let *c* be the smallest value of *j* such that tx_coded_*c*<0>=1. In other words, tx_coded_*c* is the first 66-bit control block that was received in the current group of four blocks.
d1) Let tx_payloads<(64*j*+63):64*j*> = tx_coded_*j*<65:2> for *j*=0 to 3
e1) Omit tx_coded_*c*<9:6>, which is the second nibble (based on transmission order) of the block type field for tx_coded_*c,* from tx_xcoded per the following expressions.
tx_xcoded<(64*c*+8):5> = tx_payloads<(64*c*+3):0>
tx_xcoded<256:(64*c*+9)> = tx_payloads<255:(64*c*+8)>

If for any *j*=0 to 3, tx_coded_*j*<0> = tx_coded_*j*<1> (invalid synchronization header), tx_xcoded<256:0> shall be constructed as follows:

a2) tx_xcoded<0> = 0
b2) tx_xcoded<*j*+1> = 1 for *j*=0 to 3
c2) Let tx_payloads<(64*j*+63):64*j*> = tx_coded_*j*<65:2> for *j*=0 to 3
d2) Omit the second nibble (based on transmission order) of tx_coded_0 per the following expressions.
tx_xcoded<8:5> = tx_payloads<3:0>
tx_xcoded<256:9> = tx_payloads<255:8>

Several examples of the construction of tx_xcoded<256:0> are shown in Figure 91–3. In Figure 91–3, d_*j* indicates the *j*th 66-bit block contains only data octets, c_*j* indicates the *j*th 66-bit block contains one or more control characters, f_*j* denotes the first nibble of the block type field for 66-bit block *j*, and s_*j* denotes the second nibble of the block type field for 66-bit block *j*.

Finally, scramble the first 5 bits, based on transmission order, of tx_xcoded<256:0> to yield tx_scrambled<256:0> as follows:

a3) Set tx_scrambled<4:0> to the result of the bit-wise exclusive-OR of the tx_xcoded<4:0> and tx_xcoded<12:8>.
b3) Set tx_scrambled<256:5> to tx_xcoded<256:5>

For each 257-bit block, bit 0 shall be the first bit transmitted.

91.5.2.6 Alignment marker mapping and insertion

The alignment markers that were removed per 91.5.2.4 are re-inserted after being processed by the alignment marker mapping function. The alignment marker mapping function compensates for the operation of the symbol distribution function defined in 91.5.2.8 and rearranges the alignment marker bits so that they appear on the FEC lanes intact and in the desired sequence. This preserves the properties of the alignment markers (e.g., DC balance, transition density) and provides a deterministic pattern for the purpose of synchronization. The RS-FEC receive function uses knowledge of this mapping to determine the FEC lane that is received on a given lane of the PMA service interface, to compensate for skew between FEC lanes, and to identify RS-FEC codeword boundaries.

The alignment marker mapping function operates on a group of 20 aligned and reordered alignment markers. Let am_tx_*x*<65:0> be the alignment marker for PCS lane *x*, *x*=0 to 19, where bit 0 is the first bit transmitted. The alignment markers shall be mapped to am_txmapped<1284:0> in a manner that yields the same result as the following process.

For *x*=0 to 19, amp_tx_*x*<63:0> is constructed as follows:

a) Set $y = 0$ when $x \leq 3$, set $y = 16$ when $x \geq 16$, otherwise set $y = x$.
b) amp_tx_*x*<23:0> is set to M_0, M_1, and M_2 as shown in Figure 82–9 (bits 25 to 2) using the values in Table 82–2 for PCS lane number *y*. If am_tx_*x* corresponds to a Rapid Alignment marker, then the M_4, M_5, and M_6 values are used instead (see Figure 82–11).
c) amp_tx_*x*<31:24> = am_tx_*x*<33:26>
d) amp_tx_*x*<55:32> is set to M_4, M_5, and M_6 as shown in Figure 82–9 (bits 57 to 34) using the values in Table 82–2 for PCS lane number *y*. If am_tx_*x* corresponds to a Rapid Alignment marker, then the M_0, M_1, and M_2 values are used instead (see Figure 82–11).
e) amp_tx_*x*<63:56> = am_tx_*x*<65:58>

This process replaces the fixed bytes of the alignment markers received, possibly with errors, with the values from Table 82–2. In addition it substitutes the fixed bytes of the alignment markers corresponding to PCS lanes 1, 2, and 3 with the fixed bytes for the alignment marker corresponding to PCS lane 0. Similarly, it substitutes the fixed bytes of the alignment markers corresponding to PCS lanes 17, 18, and 19 with the fixed bytes for the alignment marker corresponding to PCS lane 16. The variable bytes BIP or CD are unchanged. This process simplifies receiver synchronization since the receiver only needs to search for the fixed bytes corresponding to PCS lane 0 on each FEC lane. When the optional EEE deep sleep capability is supported, the receiver only needs to search for the fixed bytes corresponding to PCS lanes 0 and 16.

Construct a matrix of 4 rows and 320 columns, am_txpayloads, as shown in Figure 91–4. Given *i*=0 to 3, *j*=0 to 4, and *x*=*i*+4*j*, the matrix is derived per the following expression:

am_txpayloads<*i*, (64*j*+63):64*j*> = amp_tx_*x*<63:0>

Given *i*=0 to 3, *k*=0 to 31, and *y*=*i*+4*k*, am_txmapped may then be derived from am_txpayloads per the following expression:

am_txmapped<(10*y*+9):10*y*> = am_txpayloads<*i*, (10*k*+9):10*k*>

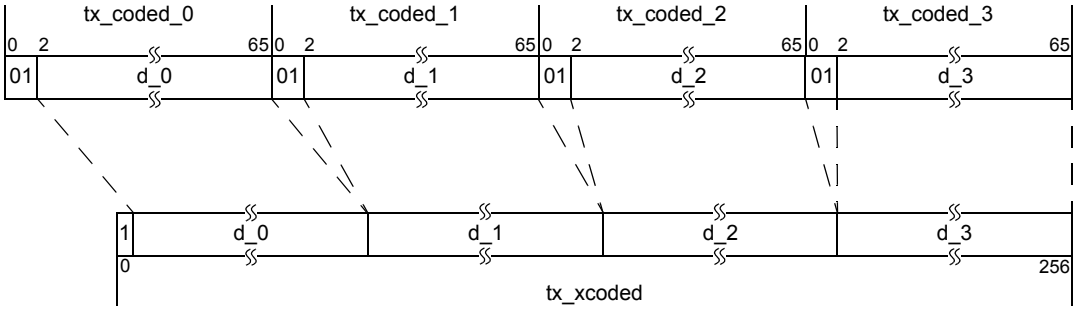

Example 1: All data blocks

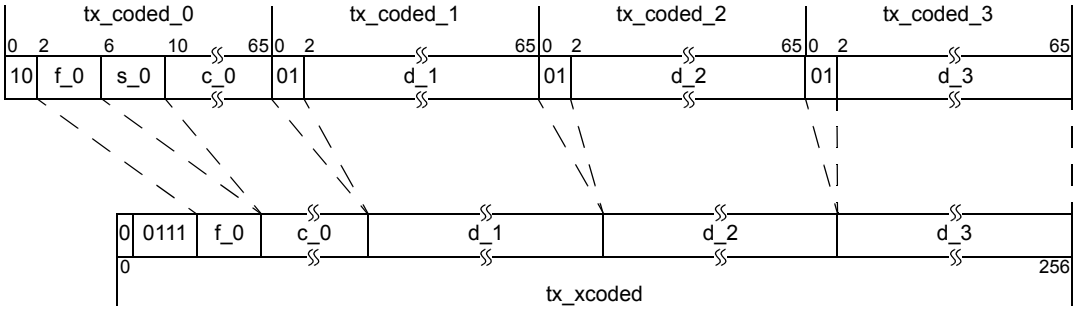

Example 2: Control block followed by three data blocks

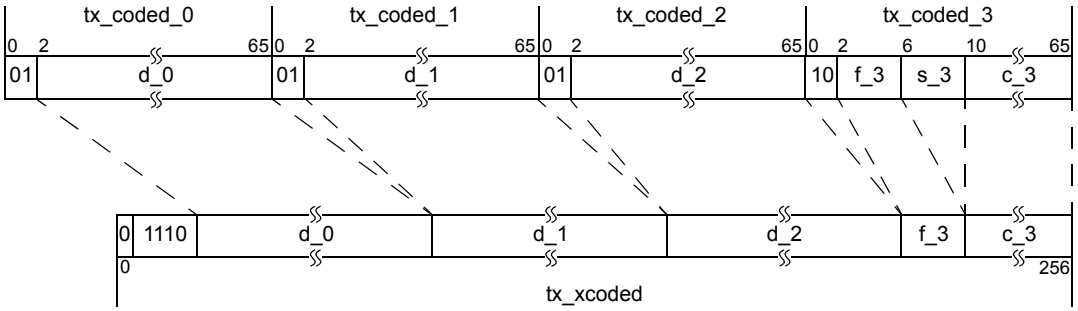

Example 3: Three data blocks followed by a control block

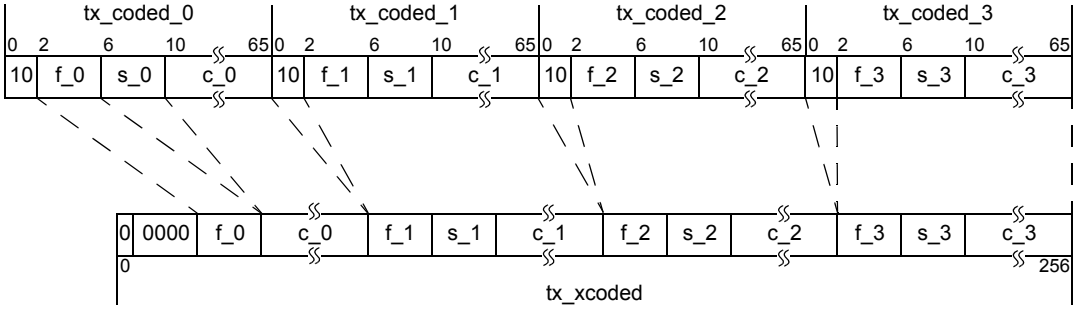

Example 4: All control blocks

Figure 91–3—Examples of the construction of tx_xcoded

A 5-bit pad is appended to the mapped alignment markers to yield the equivalent of five 257-bit blocks. The pad bits, am_txmapped<1284:1280>, shall be set to the binary values 00101 and 11010 (the leftmost bit is assigned to the highest bit index) in an alternating pattern. In other words, if a pad value of 00101 is used for the current iteration of the mapping function, a value of 11010 is used in the next iteration and vice versa.

The result of the alignment marker mapping function is a deterministic mapping between alignment marker payloads and FEC lanes. The alignment marker payloads corresponding to PCS lanes 0, 4, 8, 12, and 16 are transmitted on FEC lane 0, the alignment marker payloads corresponding to PCS lanes 0, 5, 9, 13, and 16 are transmitted on FEC lane 1, and so on (see Figure 91–4).

As a result of this process, the BIP_3 and BIP_7 fields from normal alignment markers are carried across the link protected by FEC. These fields cannot be used to monitor errors on the link protected by FEC as 64B/66B to 256B/257B transcoding and Reed-Solomon encoding alters the bit sequence. However, these fields may again be used to monitor errors after the original bit sequence is restored, i.e., following Reed-Solomon decoding and 256B/257B to 64B/66B transcoding.

One group of aligned and reordered alignment markers are mapped every 20 × 16384 66-bit blocks. This corresponds to 4096 Reed-Solomon codewords (refer to 91.5.2.7). The mapped alignment markers, am_tx-mapped<1284:0> shall be inserted as the first 1285 message bits to be transmitted from every 4096th codeword.

For the optional EEE deep sleep capability, when tx_lpi_active is true, one group of Rapid Alignment Markers (see 82.2.9) are mapped every 20 × 8 66-bit blocks. This corresponds to 2 Reed-Solomon codewords. The mapped Rapid Alignment Markers, am_txmapped<1284:0> shall be inserted as the first 1285 message bits to be transmitted from every other codeword.

The first 257-bit block inserted after am_txmapped shall correspond to the four 66-bit blocks received on PCS lanes 0, 1, 2, and 3 that immediately followed the alignment marker on each respective lane.

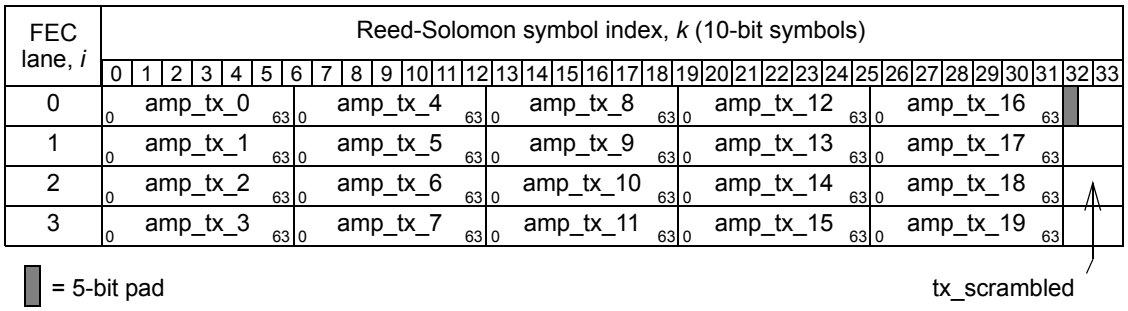

Figure 91–4—Alignment marker mapping to FEC lanes

91.5.2.7 Reed-Solomon encoder

The RS-FEC sublayer employs a Reed-Solomon code operating over the Galois Field GF(2^{10}) where the symbol size is 10 bits. The encoder processes k message symbols to generate $2t$ parity symbols, which are then appended to the message to produce a codeword of $n=k+2t$ symbols. For the purposes of this clause, a particular Reed-Solomon code is denoted RS(n,k).

When used to form a 100GBASE-CR4, 100GBASE-KR4, or 100GBASE-SR4 PHY, the RS-FEC sublayer shall implement RS(528,514). When used to form a 100GBASE-KP4 PHY, the RS-FEC sublayer shall implement RS(544,514). Each k-symbol message corresponds to 20 257-bit blocks produced by the transcoder. Each code is based on the generating polynomial given by Equation (91–1).

$$g(x) = \prod_{j=0}^{2t-1} (x - \alpha^j) = g_{2t}x^{2t} + g_{2t-1}x^{2t-1} + \ldots + g_1 x + g_0 \quad (91\text{–}1)$$

In Equation (91–1), α is a primitive element of the finite field defined by the polynomial $x^{10}+x^3+1$.

Equation (91–2) defines the message polynomial $m(x)$ whose coefficients are the message symbols m_{k-1} to m_0.

$$m(x) = m_{k-1}x^{n-1} + m_{k-2}x^{n-2} + \ldots + m_1 x^{2t+1} + m_0 x^{2t} \quad (91\text{–}2)$$

Each message symbol m_i is the bit vector $(m_{i,9}, m_{i,8}, \ldots, m_{i,1}, m_{i,0})$, which is identified with the element $m_{i,9}\alpha^9 + m_{i,8}\alpha^8 + \ldots + m_{i,1}\alpha + m_{i,0}$ of the finite field. The message symbols are composed of the bits of the transcoded blocks tx_scrambled (including a mapped group of alignment markers when appropriate) such that bit 0 of the first transcoded block in the message (or am_txmapped<0>) is bit 0 of m_{k-1} and bit 256 of the last transcoded block in the message is bit 9 of m_0. The first symbol input to the encoder is m_{k-1}.

Equation (91–3) defines the parity polynomial $p(x)$ whose coefficients are the parity symbols p_{2t-1} to p_0.

$$p(x) = p_{2t-1}x^{2t-1} + p_{2t-2}x^{2t-2} + \ldots + p_1 x + p_0 \quad (91\text{–}3)$$

The parity polynomial is the remainder from the division of $m(x)$ by $g(x)$. This may be computed using the shift register implementation illustrated in Figure 91–5. The outputs of the delay elements are initialized to zero prior to the computation of the parity for a given message. After the last message symbol, m_0, is processed by the encoder, the outputs of the delay elements are the parity symbols for that message.

The codeword polynomial $c(x)$ is then the sum of $m(x)$ and $p(x)$ where the coefficient of the highest power of x, $c_{n-1} = m_{k-1}$ is transmitted first and the coefficient of the lowest power of x, $c_0 = p_0$ is transmitted last. The first bit transmitted from each symbol is bit 0.

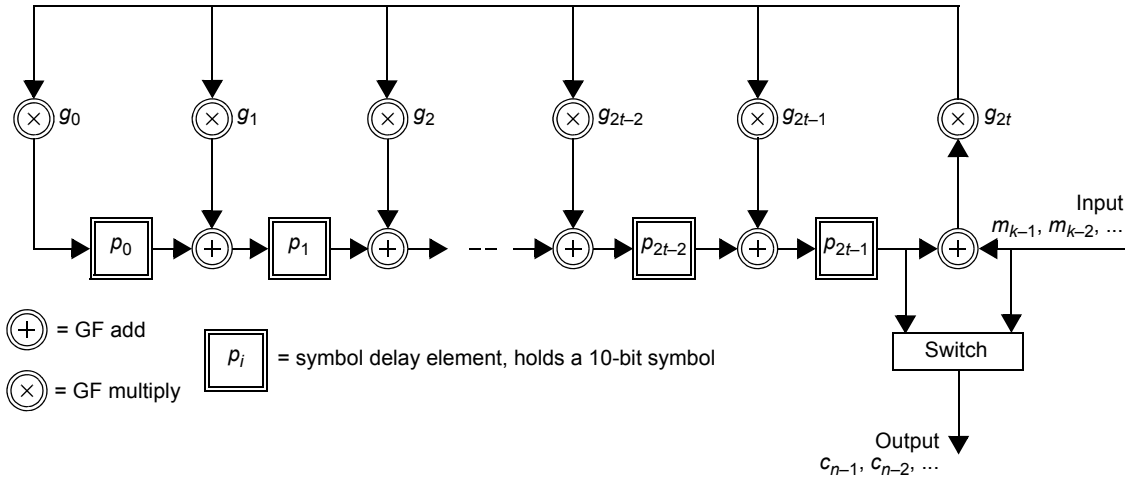

Figure 91–5—Reed-Solomon encoder functional model

The coefficients of the generator polynomial for each code are presented in Table 91–1. Example codewords for each code are provided in Annex 91A.

Table 91–1—Coefficients of the generator polynomial g_i (decimal)

i	RS(528,514)	RS(544,514)	*i*	RS(528,514)	RS(544,514)	*i*	RS(528,514)	RS(544,514)
0	432	523	11	701	883	22		565
1	290	834	12	6	503	23		108
2	945	128	13	904	942	24		1
3	265	158	14	1	385	25		552
4	592	185	15		495	26		230
5	391	127	16		720	27		187
6	614	392	17		94	28		552
7	900	193	18		132	29		575
8	925	610	19		593	30		1
9	656	788	20		249			
10	32	361	21		282			

91.5.2.8 Symbol distribution

Once the data has been Reed-Solomon encoded, it shall be distributed to 4 FEC lanes, one 10-bit symbol at a time in a round robin distribution from the lowest to the highest numbered FEC lane. The distribution process is shown in Figure 91–6.

When used to form a 100GBASE-KP4 PHY, the PMA:IS_UNITDATA_*i*.request primitive is defined to include an additional parameter (refer to 94.2.1.1.1). At the beginning of an FEC codeword, the parameter start=TRUE is asserted for the first bit of the first four symbols of the codeword transferred across the four primitives. Otherwise the parameter start is set to FALSE.

91.5.2.9 Transmit bit ordering

The transmit bit ordering is illustrated in Figure 91–6.

91.5.3 Receive function

91.5.3.1 Alignment lock and deskew

The RS-FEC receive function forms 4 bit streams by concatenating the bits from each of the 4 PMA:IS_UNITDATA_*i*.indication primitives in the order they are received. It obtains lock to the alignment markers as specified by the FEC synchronization state diagram shown in Figure 91–8.

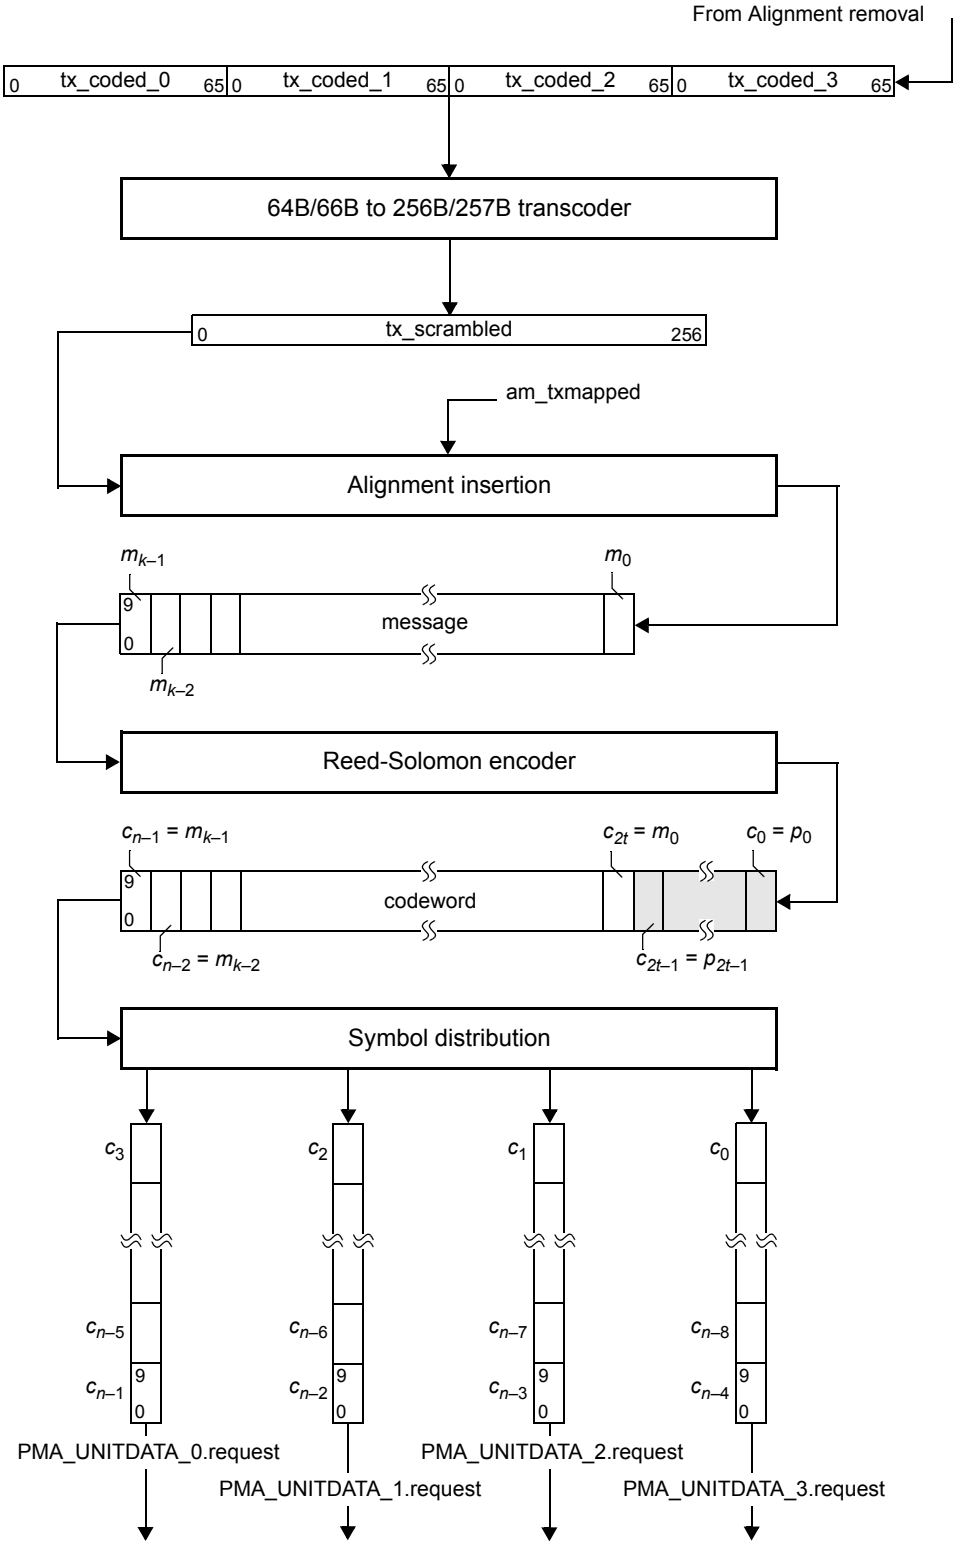

Figure 91–6—Transmit bit ordering

After alignment marker lock is achieved on all 4 lanes, all inter-lane Skew is removed as specified by the FEC alignment state diagram shown in Figure 91–9. The FEC receive function shall support a maximum Skew of 180 ns between FEC lanes and a maximum Skew Variation of 4 ns.

The 100GBASE-KP4 PMA transmit function (refer to 94.2.2) inserts PMA-specific overhead that is aligned with the start of a Reed-Solomon codeword. The 100GBASE-KP4 PMA receive function (refer to 94.2.3) synchronizes to this overhead and indicates the first bit of each of the first four symbols in a codeword by setting the PMA:IS_UNITDATA_*i*.indication parameter start=TRUE (see 94.2.1.2).

91.5.3.2 Lane reorder

FEC lanes can be received on different lanes of the service interface from which they were originally transmitted. The FEC receive function shall order the FEC lanes according to the FEC lane number (see 91.5.2.6). The FEC lane number is defined by the sequence of alignment markers that are mapped to each FEC lane.

After all FEC lanes are aligned, deskewed, and reordered, the FEC lanes are multiplexed together in the proper order to reconstruct the original stream of FEC codewords.

91.5.3.3 Reed-Solomon decoder

The Reed-Solomon decoder extracts the message symbols from the codeword, corrects them as necessary, and discards the parity symbols. The message symbols correspond to 20 transcoded blocks rx_scrambled.

When used to form a 100GBASE-CR4, 100GBASE-KR4, or 100GBASE-SR4 PHY, the RS-FEC sublayer shall be capable of correcting any combination of up to t=7 symbol errors in a codeword. When used to form a 100GBASE-KP4 PHY, the RS-FEC sublayer shall be capable of correcting any combination of up to t=15 symbol errors in a codeword. The RS-FEC sublayer shall also be capable of indicating when an errored codeword was not corrected. The probability that the decoder fails to indicate a codeword with t+1 errors as uncorrected is not expected to exceed 10^{-6}. This limit is also expected to apply for t+2 errors, t+3 errors, and so on.

The Reed-Solomon decoder may provide the option to perform error detection without error correction to reduce the delay contributed by the RS-FEC sublayer. The presence of this option is indicated by the assertion of the FEC_bypass_correction_ability variable (see 91.6.3). When the option is provided, it is enabled by the assertion of the FEC_bypass_correction_enable variable (see 91.6.1). This option shall not be used when the RS-FEC sublayer is used to form part of a 100GBASE-SR4 PHY.

NOTE—The PHY may rely on the error correction capability of the RS-FEC sublayer to achieve its performance objectives. It is recommended that acceptable performance of the underlying link is verified before error correction is bypassed.

The Reed-Solomon decoder indicates errors to the PCS sublayer by intentionally corrupting 66-bit block synchronization headers. When the decoder determines that a codeword contains errors (when the bypass correction feature is enabled) or contains errors that were not corrected (when the bypass correction feature is not supported or not enabled), it shall ensure that, for every other 257-bit block within the codeword starting with the first (1st, 3rd, 5th, etc.), the synchronization header for the first 66-bit block at the output of the 256B/257B to 64B/66B transcoder, rx_coded_0<1:0>, is set to 11. In addition, it shall ensure rx_coded_0<1:0> corresponding to the 6th 257-bit block and rx_coded_3<1:0> corresponding to the last (20th) 257-bit block in the codeword are set to 11. This causes the PCS to discard all frames 64 bytes and larger that are fully or partially within the codeword.

The Reed-Solomon decoder may optionally provide the ability to bypass the error indication feature to reduce the delay contributed by the RS-FEC sublayer. The presence of this option is indicated by the asser-

tion of the FEC_bypass_indication_ability variable (see 91.6.4). When the option is provided it is enabled by the assertion of the FEC_bypass_indication_enable variable (see 91.6.2).

When FEC_bypass_correction_enable is asserted, the decoder shall not bypass error indication and the value of FEC_bypass_indication_enable has no effect.

When FEC_bypass_indication_enable is asserted, additional error monitoring is performed by the RS-FEC sublayer to reduce the likelihood that errors in a packet are not detected. The Reed-Solomon decoder counts the number of symbol errors detected on all four FEC lanes in consecutive non-overlapping blocks of 8192 codewords. When the number of symbol errors in a block of 8192 codewords exceeds *K*, the Reed-Solomon decoder shall cause synchronization header rx_coded<1:0> of each subsequent 66-bit block that is delivered to the PCS to be assigned a value of 00 or 11 for a period of 60 ms to 75 ms. As a result, the PCS sets hi_ber=true, which inhibits the processing of received packets. When Auto-Negotiation is supported and enabled, assertion of hi_ber causes Auto-Negotiation to restart.

For the optional EEE deep sleep capability, the error monitor employed when FEC_bypass_indication_enable is asserted shall be disabled when rx_lpi_active=true. The next block of 8192 codewords considered by the error monitor shall begin on the codeword boundary following the transition of rx_lpi_active from true to false.

When the RS-FEC sublayer is used to form a 100GBASE-CR4, 100GBASE-KR4, or 100GBASE-SR4 PHY, the symbol error threshold shall be *K*=417. When the RS-FEC sublayer used to form a 100GBASE-KP4 PHY, the symbol error threshold shall be *K*=6380.

91.5.3.4 Alignment marker removal

The first 1285 message bits in every 4096th codeword is the vector am_rxmapped<1284:0> where bit 0 is the first bit received. The specific codewords that include this vector are indicated by the alignment lock and deskew function (refer to 91.5.3.1).

For the optional EEE deep sleep capability, transitions between normal alignment markers and Rapid Alignment Markers result in changes in the relative position and frequency of am_rxmapped<1284:0>. These transitions are detected by the Receive LPI state diagram (see Figure 91–11) and this information is used by the alignment marker removal function to determine which bits are to be removed. When rx_lpi_active is true, the first 1285 message bits in every other codeword is the vector am_rxmapped<1284:0>.

The vector am_rxmapped shall be removed prior to transcoding.

91.5.3.5 256B/257B to 64B/66B transcoder

The transcoder extracts a group of four 66-bit blocks, rx_coded_*j*<65:0> where *j*=0 to 3, from each 257-bit block rx_scrambled<256:0>. Bit 0 of the 257-bit block is the first bit received.

First, descramble the first 5 bits, based on reception order, of rx_scrambled<256:0> to yield rx_xcoded<256:0> as follows.

a) Set rx_xcoded<4:0> to the result of the bit-wise exclusive-OR of the rx_scrambled<4:0> and rx_scrambled<12:8>.
b) Set rx_xcoded<256:5> to rx_scrambled<256:5>.

If rx_xcoded<0> is 1, rx_coded_*j*<65:0> for *j*=0 to 3 shall be derived as follows.

a1) rx_coded_*j*<65:2> = rx_xcoded<(64*j*+64):(64*j*+1)> for *j*=0 to 3
b1) rx_coded_*j*<0>=0 and rx_coded_*j*<1>=1 for all *j*=0 to 3

If rx_xcoded<0> is 0 and any rx_xcoded<*j*+1>=0 for *j*=0 to 3, rx_coded_*j*<65:0> for *j*=0 to 3 shall be derived as follows.

a2) Let *c* be the smallest value of *j* such that rx_xcoded<*j*+1>=0. In other words, rx_coded_*c* is the first 66-bit control block in the resulting group of four blocks.
b2) Let rx_payloads be a vector representing the payloads of the four 66-bit blocks. It is derived using the following expressions:
rx_payloads<(64*c*+3):0> = rx_xcoded<(64*c*+8):5>
rx_payloads<(64*c*+7):(64*c*+4)> = 0000 (an arbitrary value that is later replaced by s_*c*)
rx_payloads<255:(64*c*+8)> = rx_xcoded<256:(64*c*+9)>
c2) rx_coded_*j*<65:2> = rx_payloads<(64*j*+63):64*j*> for *j*=0 to 3
d2) Let f_*c*<3:0> = rx_coded_*c*<5:2> be the scrambled first nibble (based on transmission order) of the block type field for rx_coded_*c*.
e2) Descramble f_*c*<3:0> to yield g<3:0> per the following expression where "^" denotes the exclusive-OR operation. When *c*=0, rx_coded_(*c*–1) corresponds to rx_coded_3 from the previous 257-bit block.
g<*i*> = f_*c*<*i*> ^ rx_coded_(*c*–1)<*i*+8> ^ rx_coded_(*c*–1)<*i*+27> for *i*=0 to 3
f2) The block type field may be uniquely identified by either its most or least significant nibble. Since g<3:0> is the least significant nibble of the block type field (per the transmission order), derive h<3:0> by cross-referencing to g<3:0> using Figure 82–5. For example, if g<3:0> is 0xE then h<3:0> is 0x1. If no match to g<3:0> is found, h<3:0> is set to 0000.
g2) If rx_xcoded<*j*+1>=0, rx_coded_*j*<0>=1 and rx_coded_*j*<1>=0 for *j*=0 to 3
h2) If rx_xcoded<*j*+1>=1, rx_coded_*j*<0>=0 and rx_coded_*j*<1>=1 for *j*=0 to 3
i2) If h<3:0> = 0000, rx_coded_*c*<1>=1 (invalidate synchronization header)

If rx_xcoded<0> is 0 and all rx_xcoded<*j*+1>=1 for *j*=0 to 3, rx_coded_*j*<65:0> for *j*=0 to 3 shall be derived as follows.

a3) Set *c* = 0 and h<3:0> = 0000.
b3) Let rx_payloads be a vector representing the payloads of the four 66-bit blocks. It is derived using the following expressions.
rx_payloads<(64*c*+3):0> = rx_xcoded<(64*c*+8):5>
rx_payloads<(64*c*+7):(64*c*+4)> = 0000 (an arbitrary value that is later replaced by s_*c*)
rx_payloads<255:(64*c*+8)> = rx_xcoded<256:(64*c*+9)>
c3) rx_coded_*j*<65:2> = rx_payloads<(64*j*+63):(64*j*)> for *j*=0 to 3
d3) rx_coded_*j*<0>=0 and rx_coded_*j*<1>=0 for *j*=0 and 2
e3) rx_coded_*j*<0>=1 and rx_coded_*j*<1>=1 for *j*=1 and 3

If rx_xcoded<0> is 0, scramble h<3:0> to yield s_*c*<3:0> and assign it to rx_coded_*c* per the following expressions.

a4) s_*c*<*i*> = h<*i*> ^ rx_coded_(*c*–1)<*i*+12> ^ rx_coded_(*c*–1)<*i*+31> for *i*=0 to 3
b4) rx_coded_*c*<9:6> = s_*c*<3:0>

The 66-bit blocks are transmitted in order from *j*=0 to 3. Bit 0 of each block is the first bit transmitted.

91.5.3.6 Block distribution

After the data has been transcoded, it shall be distributed to multiple PCS lanes, one 66-bit block at a time in a round robin distribution from the lowest to the highest numbered PCS lanes. The distribution process is shown in Figure 82–6.

91.5.3.7 Alignment marker mapping and insertion

The alignment marker mapping function compensates for operation of lane reorder function (refer to 91.5.3.2) to derive the PCS lane alignment markers, am_rx_*x*<65:0> for *x*=0 to 19, from am_rx-mapped<1284:0> (refer to 91.5.3.4).

The alignment markers shall be derived from am_rxmapped<1284:0> in a manner that yields the same result as the following process.

Given *i*=0 to 3, *k*=0 to 31, and *y*=*i*+4*k*, am_rxpayloads may be derived from am_rxmapped per the following expression:

am_rxpayloads<*i*, (10*k*+9):10*k*> = am_rxmapped<(10*y*+9):10*y*>

The 5-bit pad am_rxmapped<1284:1280> is ignored. Given *i*=0 to 3, *j*=0 to 4, and *x*=*i*+4*j*, amp_rx_*x* may be derived from am_rxpayloads by the following expression:

amp_rx_*x*<63:0> = am_rxpayloads<*i*, (64*j*+63):64*j*>

For *x*=0 to 19, am_rx_*x*<65:0> is constructed as follows:

a) am_rx_*x*<0>=1 and am_rx_*x*<1>=0.
b) am_rx_*x*<25:2> is set to M_0, M_1, and M_2 as shown in Figure 82–9 using the values in Table 82–2 for PCS lane number *x*. If amp_rx_*x* corresponds to a Rapid Alignment marker, then the M_4, M_5, and M_6 values are used instead (see Figure 82–11).
c) am_rx_*x*<33:26> = amp_rx_*x*<31:24>.
d) am_rx_*x*<57:34> is set to M_4, M_5, and M_6 as shown in Figure 82–9 using the values in Table 82–2 for PCS lane number *x*. If amp_rx_*x* corresponds to a Rapid Alignment marker, then the M_0, M_1, and M_2 values are used instead (see Figure 82–11).
e) am_rx_*x*<65:58> = amp_rx_*x*<63:56>.

One vector is mapped to 20 alignment markers every 4096 Reed-Solomon codewords (see 91.5.3.4). The alignment markers are simultaneously transmitted on the 20 PCS lanes after every 16383rd column of 20 66-bit blocks.

For the optional EEE deep sleep capability, when rx_lpi_active is true, one vector is mapped to 20 Rapid Alignment Markers every 2 Reed-Solomon codewords. The Rapid Alignment Markers are simultaneously transmitted on the 20 PCS lanes after every 7th column of 20 66-bit blocks.

The alignment markers am_rx_0 to am_rx_3 shall be inserted so that they are immediately followed by rx_coded_0 to rx_coded_3, respectively, as derived from the first 257-bit block following am_rxmapped. Similarly am_rx_4 to am_rx_7 are followed by the 66-bit blocks corresponding to the second 257-bit block following am_rxmapped, and so on.

91.5.3.8 Receive bit ordering

The receive bit ordering is illustrated in Figure 91–7. This illustration shows the case where the FEC lanes appear across the PMA:IS_UNITDATA_*i*.indication primitives in the correct order.

91.5.4 Detailed functions and state diagrams

91.5.4.1 State diagram conventions

The body of this subclause is comprised of state diagrams, including the associated definitions of variables, functions, and counters. 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. The notation ++ after a counter or integer variable indicates that its value is to be incremented.

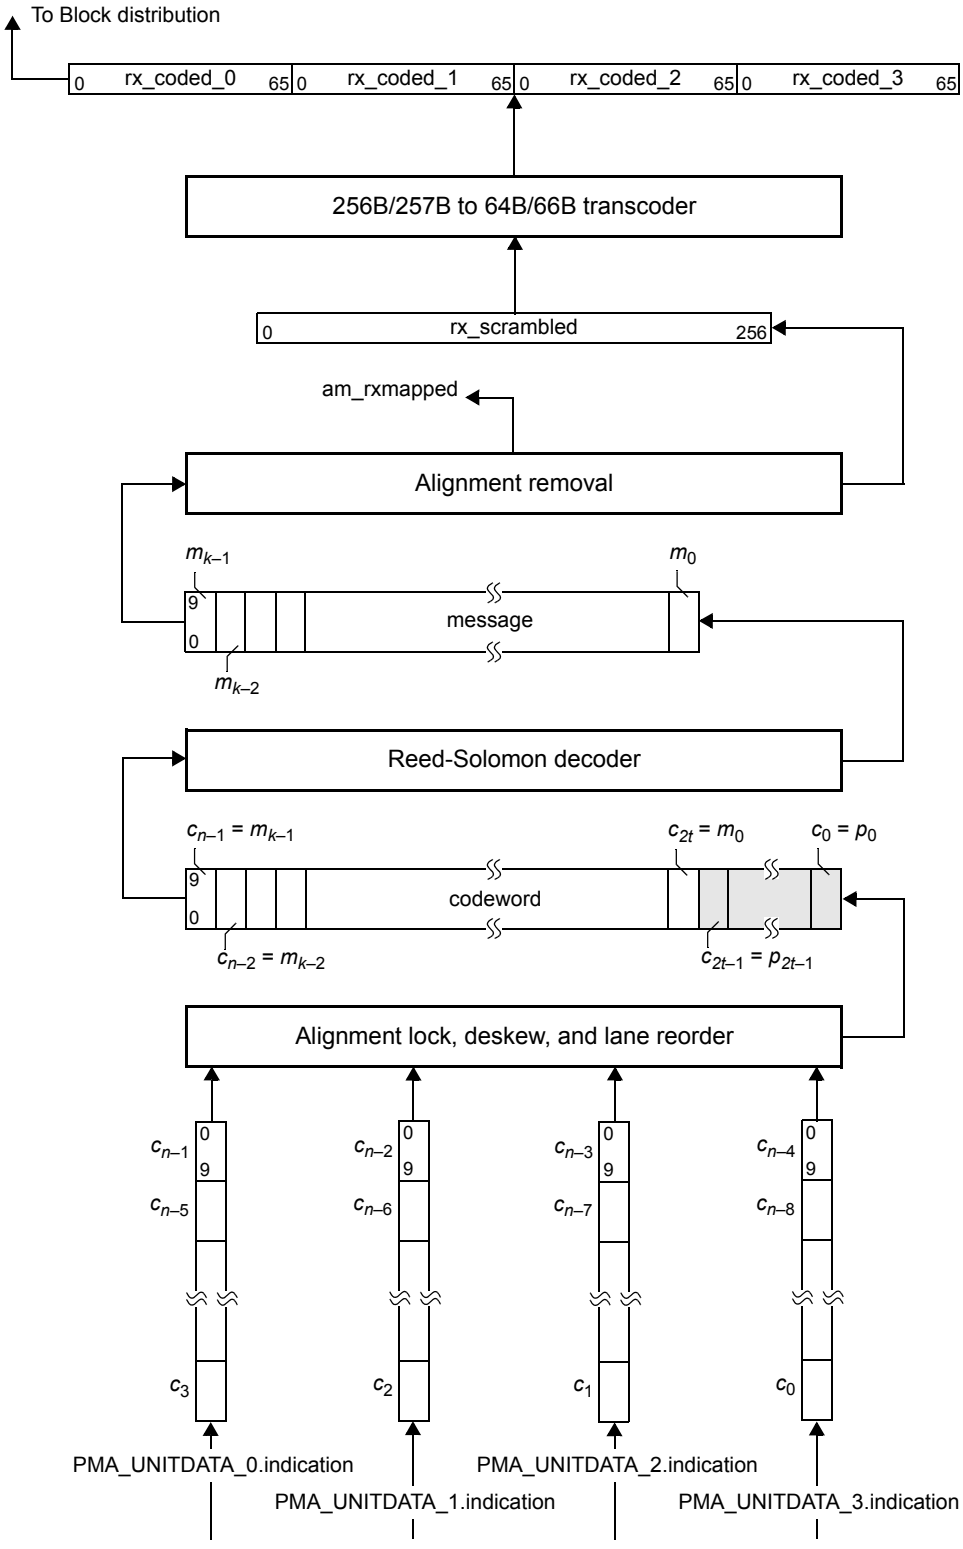

Figure 91–7—Receive bit ordering

91.5.4.2 State variables

91.5.4.2.1 Variables

all_locked
: A Boolean variable that is set to true when amps_lock<*x*> is true for all *x* and is set to false when amps_lock<*x*> is false for any *x*.

amp_counter_done
: Boolean variable that indicates that amp_counter has reached its terminal count.

amp_match
: Boolean variable that holds the output of the function AMP_COMPARE.

amp_valid
: Boolean variable that is set to true if the received 64-bit block is a valid alignment marker payload. The alignment marker payload, mapped to an FEC lane according to the process described in 91.5.2.6, consists of 48 known bits and 16 variable bits (the BIP_3 or CD_3 field and its complement BIP_7 or CD_7, see 82.2.7). The bits of the candidate block that are in the positions of the known bits in the alignment marker payload are compared on a nibble-wise basis (12 comparisons). If no more than 3 nibbles in the candidate block fail to match the corresponding known nibbles in the alignment marker payload, the candidate block is considered a valid alignment marker payload. For the normal mode of operation, each FEC lane compares the candidate block to the alignment marker payload for PCS lane 0. For the optional EEE deep sleep capability, each FEC lane also compares the candidate block to the alignment marker payload for PCS lane 16 when rx_lpi_active is true.

amps_lock<*x*>
: Boolean variable that is set to true when the receiver has detected the location of the alignment marker payload sequence for a given lane on the PMA service interface, where *x* = 0:3.

current_pcsl
: A variable that holds the PCS lane number corresponding to the current alignment marker payload that is recognized on a given lane of the PMA service interface. It is compared to the variable first_pcsl to confirm that the location of the alignment marker payload sequence has been detected.

cw_bad
: A Boolean variable that is set to true if the Reed-Solomon decoder (see 91.5.3.3) fails to correct the current FEC codeword and is set to false otherwise.

deskew_done
: A Boolean variable that is set to true when fec_enable_deskew is set to true and the deskew process is completed. Otherwise, this variable is set to false.

fec_align_status
: A variable set by the FEC alignment process to reflect the status of FEC lane-to-lane alignment. Set to true when all lanes are synchronized and aligned and set to false when the deskew process is not complete.

fec_alignment_valid
: Boolean variable that is set to true if all FEC lanes are aligned. FEC lanes are considered to be aligned when amps_lock<*x*> is true for all *x*, each FEC lane is locked to a unique alignment marker payload sequence (see 91.5.2.6), and the FEC lanes are deskewed. Otherwise, this variable is set to false.

fec_enable_deskew
: A Boolean variable that enables and disables the deskew process. Received bits may be discarded whenever deskew is enabled. It is set to true when deskew is enabled and set to false when deskew is disabled.

fec_lane
: A variable that holds the FEC lane number (0 to 3) received on lane *x* of the PMA service interface when amps_lock<*x*>=true. The FEC lane number is determined by the alignment marker payloads in the 2nd, 3rd, or 4th positions of the sequence based on the mapping defined in 91.5.2.6. The 48 bits that are in the positions of the known bits in the received alignment marker payload are

compared to the expected values for a given payload position and FEC lane on a nibble-wise basis (12 comparisons). If no more than 3 nibbles in the candidate block fail to match the corresponding known nibbles for any payload position on a given FEC lane, then the FEC lane number is assigned accordingly.

first_pcsl
: A variable that holds the PCS lane number that corresponds to the first alignment marker payload that is recognized on a given lane of the PMA service interface. It is compared to the PCS lane number corresponding to the second alignment marker payload that is tested.

reset
: Boolean variable that controls the resetting of the RS-FEC sublayer. 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 RS-FEC sublayer into low-power mode.

restart_lock
: Boolean variable that is set by the FEC alignment process to reset the synchronization process on all FEC lanes. It is set to true after 3 consecutive uncorrected codewords are received (3_BAD state) and set to false upon entry into the LOSS_OF_ALIGNMENT state.

rx_align_status
: Boolean variable that is set by the alignment lock and deskew function (see 91.5.2.2).

signal_ok
: Boolean variable that is set based on the most recently received value of PMA:IS_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 set to true when the SLIP requested by the synchronization state diagram has been completed indicating that the next candidate 64-bit block position can be tested.

test_amp
: Boolean variable that is set to true when a candidate block position is available for testing and false when the FIND_1ST state is entered.

test_cw
: Boolean variable that is set to true when a new FEC codeword is available for decoding and is set to false when the TEST_CW state is entered.

The following variables are only used for the optional EEE deep sleep capability. If this capability is not supported, the values of tx_lpi_active and rx_lpi_active are set to false.

1st_ram_counter_done
: Boolean variable that indicates that 1st_ram_counter has reached its terminal count.

1st_ramps_counter_done
: Boolean variable that indicates that 1st_ramps_counter has reached its terminal count.

fec_lpi_fw
: Boolean variable that controls the behavior of the Transmit LPI and Receive LPI state diagrams. This variable is set to true when the local PCS is configured to use the fast wake mechanism and set to false otherwise.

ram_counter_done
: Boolean variable that indicates that ram_counter has reached its terminal count.

ram_valid
: Boolean variable that is set to true when the 66-bit blocks concurrently received on at least 2 PCS lanes are valid Rapid Alignment Markers with identical values for rx_down_count and is set to false otherwise.

ram_valid_prev
: Boolean variable that holds the value of ram_valid from the previous expected Rapid Alignment Marker position.

ramps_counter_done
: Boolean variable that indicates that ramps_counter has reached its terminal count.

ramps_valid
: Boolean variable that is set to true when the 64-bit block payloads concurrently received on at least 2 FEC lanes are valid Rapid Alignment Marker payloads with identical values for rx_down_count and is set to false otherwise.

ramps_valid_prev
: Boolean variable that holds that value of ramps_valid from the previous expected Rapid Alignment Marker payload position.

rx_down_count
: The value that results from the bit-wise exclusive-OR of the Count Down (CD_3) byte and the M_0 byte of the current Rapid Alignment Marker payload (see 82.2.9).

rx_lpi_active
: A Boolean variable that is set to true when the RS-FEC sublayer infers that the Low Power Idle is being received from the link partner and is set to false otherwise.

rx_quiet_timer_done
: A Boolean variable that indicates that rx_quiet_timer has reached its terminal count.

tx_down_count
: The value that results from the bit-wise exclusive-OR of the Count Down (CD_3) byte and the M_0 byte of the current Rapid Alignment Marker (see 82.2.9).

tx_lpi_active
: A Boolean variable that is set to true when the RS-FEC sublayer infers that the local PCS is transmitting Low Power Idle and is set to false otherwise.

tx_quiet_timer_done
: Boolean variable that indicates that tx_quiet_timer has reached its terminal count.

91.5.4.2.2 Functions

AMP_COMPARE
: This function compares the values of first_pcsl and current_pcsl to determine if a valid alignment marker payload sequence has been detected and returns the result of the comparison using the variable amp_match. When rx_lpi_active is false, if current_pcsl and first_pcsl are 0, amp_match is set to true. When rx_lpi_active is true, the comparison is performed as follows. If first_pcsl is 0 then amp_match is set to true if current_pcsl is 16. If first_pcsl is 16 then amp_match is set to true if current_pcsl is 0. Otherwise, amp_match is set to false.

SLIP
: Causes the next candidate block position to be tested. The precise method for determining the next candidate block position is not specified and is implementation dependent. However, an implementation shall ensure that all possible block positions are evaluated.

91.5.4.2.3 Counters

amp_counter
: When rx_lpi_active is false, this counter counts the 4096 FEC codewords that separate the ends of two consecutive normal alignment marker payload sequences. An FEC codeword is 1320 bits per FEC lane for 100GBASE-KR4 and 1360 bits per FEC lane for 100GBASE-KP4. When rx_lpi_active is true, then amp_counter is defined as follows. If first_pcsl corresponds to PCS lane 0, it counts the 256 bits to the end of the expected location of the Rapid Alignment Marker payload corresponding to PCS lane 16. If first_pcsl corresponds to PCS lane 16, this counter counts the 2 FEC codewords minus 256 bits to the end of the expected location of the next Rapid Alignment Marker payload corresponding to PCS lane 0.

cw_bad_count
: Counts the number of consecutive uncorrected FEC codewords. This counter is set to zero when an FEC codeword is received and cw_bad is false for that codeword.

The following counters are only used for the optional EEE deep sleep capability.

1st_ram_counter
: This counter counts 4 66-bit blocks from the end of one candidate RAM position to the end of the next candidate RAM position. The first instance of the counter counts from the end of the last normal alignment marker received.

1st_ramps_counter
: This counter counts one FEC codeword from the end of one candidate RAM payload to the end of the next RAM payload position. An FEC codeword is 1320 bits per FEC lane for 100GBASE-KR4 and 1360 bits per FEC lane for 100GBASE-KP4

ram_counter
: This counter counts 8 66-bit blocks from the end of the current RAM to the end of the next expected RAM position.

ramps_counter
: This counter counts 2 FEC codewords from the end of the current RAM payload to the end of the next expected RAM payload position.

rx_quiet_timer
: This timer limits the maximum time fec_align_status may be deasserted before the Transmit LPI state diagram concludes that the link has failed. The value of this timer is between 2 ms and 2.8 ms.

tx_quiet_timer
: This timer limits the maximum time rx_align_status may be deasserted before the Transmit LPI state diagram concludes that the link has failed. The value of this timer is between 1.8 ms and 2 ms.

91.5.4.3 State diagrams

The FEC shall implement four synchronization processes as shown in Figure 91–8. The synchronization process operates independently on each lane. The synchronization state diagram determines when the FEC has detected the location of the alignment marker payload sequence in the received bit stream for a given lane of the service interface.

The FEC shall implement the alignment process as shown in Figure 91–9.

When the optional EEE deep sleep capability is supported, the FEC shall also implement the Transmit LPI process as shown in Figure 91–10 and the Receive LPI process as shown in Figure 91–11. The Transmit LPI state diagram infers when Low Power Idle is being transmitted by the local PCS by checking for the presence of Rapid Alignment Markers. The Receive LPI state diagram infers when Low Power Idle is being received from the link partner using a similar mechanism. Monitoring the position and frequency of alignment markers is also critical to the operation of the alignment marker removal function (see 91.5.2.4 and 91.5.3.4).

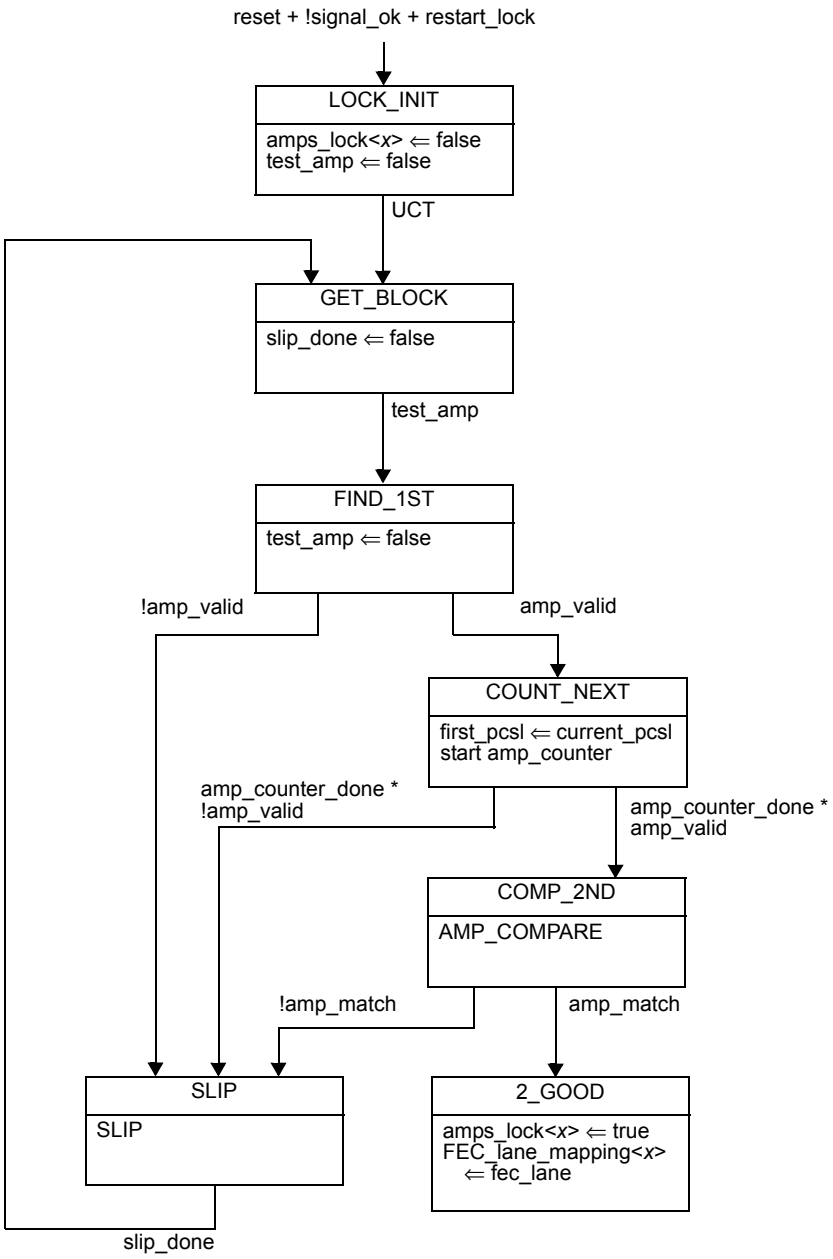

Figure 91–8—FEC synchronization state diagram

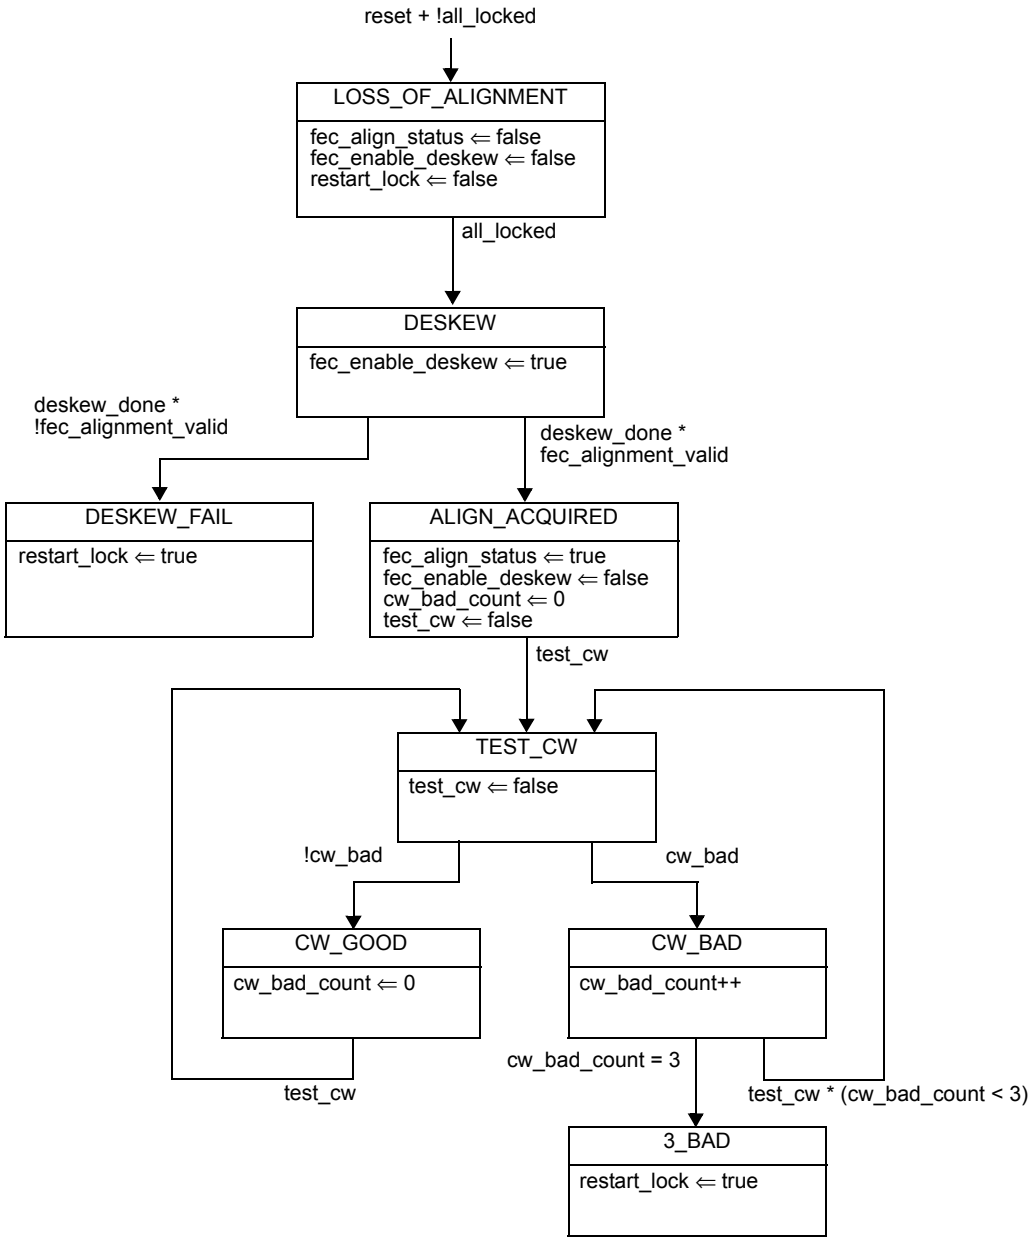

Figure 91–9—FEC alignment state diagram

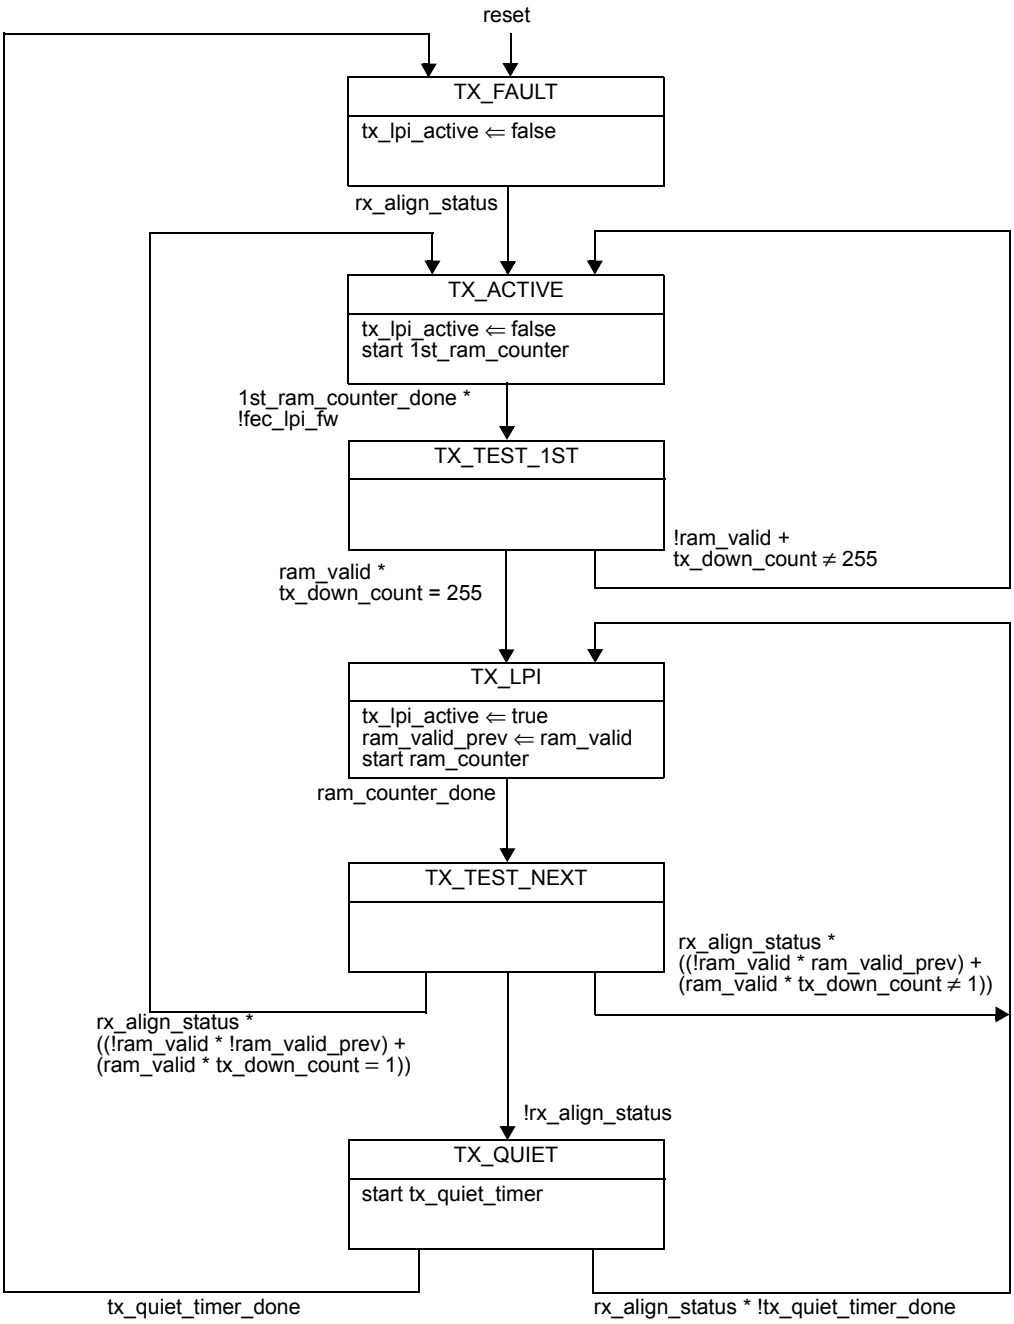

Figure 91–10—Transmit LPI state diagram

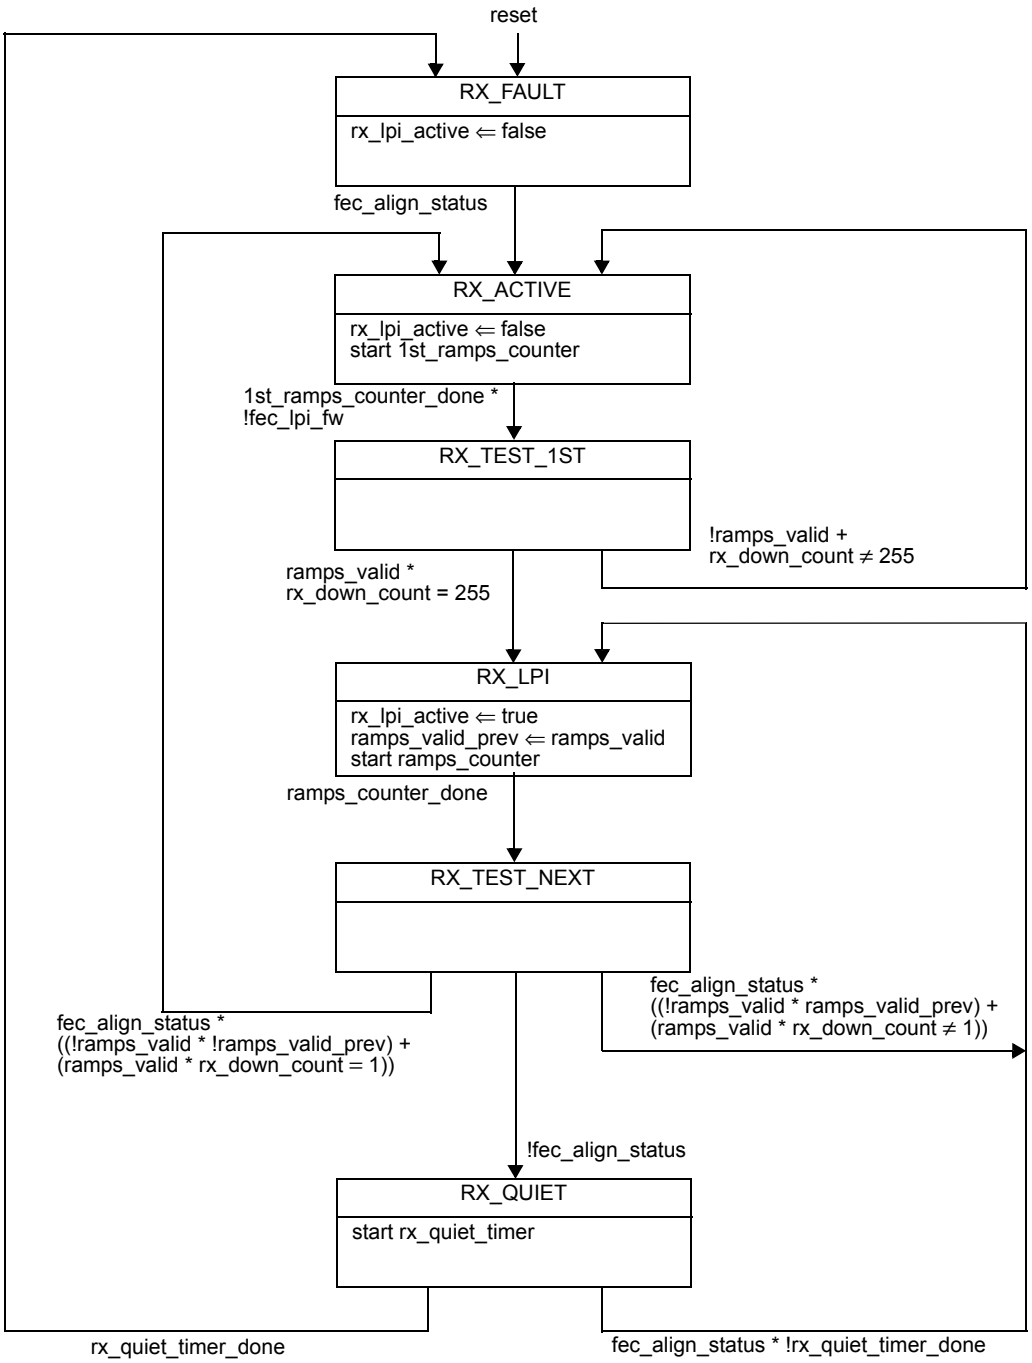

Figure 91–11—Receive LPI state diagram

91.6 RS-FEC MDIO function mapping

The optional MDIO capability described in Clause 45 defines several registers that provide control and status information for and about the RS-FEC. If MDIO is implemented, it shall map MDIO control bits to RS-FEC control variables as shown in Table 91–2, and MDIO status bits to RS-FEC status variables as

shown in Table 91–3, and if a separated PMA (see 45.2.1) is connected to the FEC service interface it shall map additional MDIO status bits to additional RS-FEC status variables as shown in Table 91–4.

Table 91–2—MDIO/RS-FEC control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	FEC variable
FEC bypass correction enable	RS-FEC control register	1.200.0	FEC_bypass_correction_enable
FEC bypass indication enable	RS-FEC control register	1.200.1	FEC_bypass_indication_enable

Table 91–3—MDIO/RS-FEC status variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	FEC variable
FEC bypass correction ability	RS-FEC status register	1.201.0	FEC_bypass_correction_ability
FEC bypass indication ability	RS-FEC status register	1.201.1	FEC_bypass_indication_ability
RS-FEC high SER	RS-FEC status register	1.201.2	hi_ser
FEC AM lock *x*, *x*=0 to 3	RS-FEC status register	1.201.8:11	amps_lock<*x*>
RS-FEC align status	RS-FEC status register	1.201.14	fec_align_status
FEC corrected codewords	RS-FEC corrected codewords counter register	1.202, 1.203	FEC_corrected_cw_counter
FEC uncorrected codewords	RS-FEC uncorrected codewords counter register	1.204, 1.205	FEC_uncorrected_cw_counter
FEC lane *x* mapping	RS-FEC lane mapping register	1.206	FEC_lane_mapping<*x*>
FEC symbol errors, FEC lanes 0 to 3	RS-FEC symbol error counter register, FEC lanes 0 to 3	1.210 to 1.217	FEC_symbol_error_counter_*i*

Table 91–4—MDIO/RS-FEC status variable mapping for separated PMA

MDIO control variable	PMA/PMD register name	Register/bit number	FEC variable
PCS align status	RS-FEC status register	1.201.15	align_status
BIP errors, PCS lanes 0 to 19	RS-FEC BIP error counter register, PCS lanes 0 to 19	1.230 to 1.249	BIP_error_counter_*i*
PCS lane *x* mapping	PCS lane *x* mapping register	1.250 to 1.269	lane_mapping<*x*>
Block *x* lock	RS-FEC PCS alignment status 1 and 2 registers	1.280 to 1.281	block_lock<*x*>
Lane *x* aligned	RS-FEC PCS alignment status 3 and 4 registers	1.282 to 1.283	am_lock<*x*>

The following subclauses define variables that are not otherwise defined, e.g., for use by state diagrams.

91.6.1 FEC_bypass_correction_enable

When this variable is set to one, the Reed-Solomon decoder performs error detection without error correction (see 91.5.3.3). When this variable is set to zero, the decoder also performs error correction. The default value of the variable is zero. This variable is mapped to the bit defined in 45.2.1.101 (1.200.0).

91.6.2 FEC_bypass_indication_enable

This variable is set to one to bypass the error indication function (see 91.5.3.3) when this ability is supported. When this variable is set to zero, the decoder indicates errors to the PCS sublayer. This variable has no effect (the decoder does not bypass error indication) if FEC bypass correction enable (1.200.0) is set to one. The default value of this variable is zero. This variable is mapped to the bit defined in 45.2.1.101 (1.200.1).

91.6.3 FEC_bypass_correction_ability

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 variable is set to one to indicate that the decoder has the ability to bypass error correction. The variable is set to zero if this ability is not supported. This variable is mapped to the bit defined in 45.2.1.102 (1.201.0).

91.6.4 FEC_bypass_indication_ability

The Reed-Solomon decoder may have the option to bypass the error indication function (see 91.5.3.3) to reduce the delay contributed by the RS-FEC sublayer. This variable is set to one to indicate that the decoder has the ability to bypass error indication. The variable is set to zero if this ability is not supported. This variable is mapped to the bit defined in 45.2.1.102 (1.201.1).

91.6.5 hi_ser

This variable is defined when the FEC_bypass_indication_ability variable is set to one. 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. This variable is mapped to the bit defined in 45.2.1.102 (1.201.2).

91.6.6 amps_lock<*x*>

These variables are assigned by the FEC alignment state diagram shown in Figure 91–9 (see 91.5.4.3). They are mapped to the bits defined in 45.2.1.102 (1.201.8 to 1.201.11 for FEC lanes 0 to 3, respectively).

91.6.7 fec_align_status

This variable assigned by the FEC alignment state diagram shown in Figure 91–9 (see 91.5.4.3). It is mapped to the bit defined in 45.2.1.102 (1.201.14).

91.6.8 FEC_corrected_cw_counter

A corrected FEC codeword is a codeword that contains errors and was corrected.

FEC_corrected_cw_counter is a 32-bit counter that counts once for each corrected FEC codeword processed when fec_align_status is true. This variable is mapped to the registers defined in 45.2.1.103 (1.202, 1.203).

91.6.9 FEC_uncorrected_cw_counter

An uncorrected FEC codeword is a codeword that contains errors (when the bypass correction feature is supported and enabled) or contains errors that were not corrected (when the bypass correction feature is not supported or not enabled).

FEC_uncorrected_cw_counter is a 32-bit counter that counts once for each uncorrected FEC codeword processed when fec_align_status is true. This variable is mapped to the registers defined in 45.2.1.104 (1.204, 1.205).

91.6.10 FEC_lane_mapping<*x*>

When the RS-FEC receive function detects and locks to an alignment marker payload on PMA service interface lane *x*, the FEC lane number corresponding to the detected alignment marker payload is assigned to the variable FEC_lane_mapping<*x*>. These variables are mapped to the register defined in 45.2.1.105 (1.206).

91.6.11 FEC_symbol_error_counter_*i*

FEC_symbol_error_counter_*i*, where *i*=0 to 3, is a 32-bit counter that counts once for each 10-bit symbol corrected on FEC lane *i* when fec_align_status is true. These variables are mapped to the registers defined in 45.2.1.106 and 45.2.1.107 (1.210 to 1.217).

91.6.12 align_status

This variable is assigned the value of rx_align_status as defined by the PCS deskew state diagram shown in Figure 82–14 (see 91.5.2.2). It is mapped to the bit defined in 45.2.1.102 (1.201.15).

91.6.13 BIP_error_counter_*i*

BIP_error_counter_*i*, where *i*=0 to 19, is a 16-bit counter that holds the BIP error count for PCS lane *i* as calculated by the RS-FEC transmit function (see 91.5.2.4). These variables are mapped to the registers defined in 45.2.1.108 and 45.2.1.109 (1.230 to 1.249).

91.6.14 lane_mapping<*x*>

When the RS-FEC transmit function detects and locks to an alignment marker on FEC service interface lane *x*, the PCS lane number corresponding to the detected alignment marker is assigned to the variable lane_mapping<*x*>. These variables are mapped to the registers defined in 45.2.1.110 and 45.2.1.111 (1.250 to 1.269).

91.6.15 block_lock<*x*>

These variables are assigned by the block lock state diagram shown in Figure 82–12 (see 91.5.2.1). They are mapped to the registers defined in 45.2.1.112 and 45.2.1.113 (1.280 to 1.281).

91.6.16 am_lock<*x*>

These variables are assigned by the alignment marker lock state diagram shown in Figure 82–13 (see 91.5.2.2). They are mapped to the registers defined in 45.2.1.114 and 45.2.1.115 (1.282 to 1.283).

91.7 Protocol implementation conformance statement (PICS) proforma for Clause 91, Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 100GBASE-R PHYs[20]

91.7.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 91, Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 100GBASE-R PHYs, 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.

91.7.2 Identification

91.7.2.1 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).	

91.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 91, Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 100GBASE-R PHYs
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	

[20] *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.

91.7.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*KR4	100GBASE-CR4 or 100GBASE-KR4		Used to form complete 100GBASE-CR4 or 100GBASE-KR4 PHY	O	Yes [ ] No [ ]
*KP4	100GBASE-KP4		Used to form complete 100GBASE-KP4 PHY	O	Yes [ ] No [ ]
*SR4	100GBASE-SR4		Used to form complete 100GBASE-SR4 PHY	O	Yes [ ] No [ ]
DC	Delay constraints	91.4	Conforms to delay constraints specified in 91.4	M	Yes [ ]
*MD	MDIO capability	91.6	Registers and interface supported	O	Yes [ ] No [ ]
*BEC	Bypass error correction	91.5.3.3	Capability is supported	O	Yes [ ] No [ ]
*BEI	Bypass error indication	91.5.3.3	Capability is supported	O	Yes [ ] No [ ]
*EEE	EEE capability	91.5.4.3	Capability is supported	O	Yes [ ] No [ ]

91.7.4 PICS proforma tables for Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 100GBASE-R PHYs

91.7.4.1 Transmit function

Item	Feature	Subclause	Value/Comment	Status	Support
TF1	Skew tolerance	91.5.2.2	Maximum Skew of 49 ns between PCS lanes and a maximum Skew Variation of 400 ps	M	Yes [ ]
TF2	Lane reorder	91.5.2.3	Order the PCS lanes according to the PCS lane number	M	Yes [ ]
TF3	64B/66B to 256B/257B transcoder	91.5.2.5	tx_xcoded<256:0> constructed per 91.5.2.5	M	Yes [ ]
TF4	257-bit block transmission order	91.5.2.5	First bit transmitted is bit 0	M	Yes [ ]
TF5	Alignment maker mapping	91.5.2.6	Map to am_tx-mapped<1284:0> per 91.5.2.6	M	Yes [ ]
TF6	Pad value	91.5.2.6	Binary values 00101 and 11010 (the leftmost bit is assigned to the highest bit index) in an alternating pattern	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
TF7	Alignment marker insertion	91.5.2.6	First 1285 message bits to be transmitted from every 4096th codeword	M	Yes [ ]
TF8	Alignment marker insertion when tx_lpi_active is true	91.5.2.6	First 1285 message bits to be transmitted from every other codeword	EEE:M	Yes [ ] N/A [ ]
TF9	Alignment marker insertion point	91.5.2.6	First 257-bit block inserted after am_txmapped corresponds to the four 66-bit blocks received on PCS lanes 0, 1, 2, and 3 that immediately followed the alignment marker on each respective lane	M	Yes [ ]
TF10	Reed-Solomon encoder for 100GBASE-CR4 or 100GBASE-KR4 or 100GBASE-SR4	91.5.2.7	RS(528,514)	KR4:M or SR4:M	Yes [ ] N/A [ ]
TF11	Reed-Solomon encoder for 100GBASE-KP4	91.5.2.7	RS(544,514)	KP4:M	Yes [ ] N/A [ ]
TF12	Symbol distribution	91.5.2.8	Distributed to 4 FEC lanes, one 10-bit symbol at a time in a round robin distribution from the lowest to the highest numbered FEC lane	M	Yes [ ]

91.7.4.2 Receive function

Item	Feature	Subclause	Value/Comment	Status	Support
RF1	Skew tolerance	91.5.3.1	Maximum Skew of 180 ns between FEC lanes and a maximum Skew Variation of 4 ns	M	Yes [ ]
RF2	Lane reorder	91.5.3.2	Order the FEC lanes according to the FEC lane number	M	Yes [ ]
RF3	Reed-Solomon decoder for 100GBASE-CR4 or 100GBASE-KR4	91.5.3.3	Corrects any combination of up to t=7 symbol errors in a codeword unless error correction bypassed	KR4:M	Yes [ ] N/A [ ]
RF4	Reed-Solomon decoder for 100GBASE-KP4	91.5.3.3	Corrects any combination of up to t=15 symbol errors in a codeword unless error correction bypassed	KP4:M	Yes [ ] N/A [ ]
RF5	Reed-Solomon decoder for 100GBASE-SR4	91.5.3.3	Corrects any combination of up to t=7 symbol errors in a codeword	SR4:M	Yes [ ] N/A [ ]
RF6	Error correction bypass for 100GBASE-SR4	91.5.3.3	Error correction is not bypassed	SR4:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
RF7	Reed-Solomon decoder	91.5.3.3	Capable of indicating when a codeword was not corrected.	M	Yes [ ]
RF8	Error indication function	91.5.3.3	Corrupts 66-bit block synchronization headers for uncorrected errored codewords (or errored codewords when correction is bypassed)	M	Yes [ ]
RF9	Error indication when error correction is bypassed	91.5.3.3	Error indication is not bypassed	BEI:M	Yes [ ] N/A [ ]
RF10	Error monitoring while error indication is bypassed	91.5.3.3	When the number of symbols errors in a block of 8192 codewords exceeds *K*, corrupt 66-bit block synchronization headers	BEI:M	Yes [ ] N/A [ ]
RF11	Symbol error threshold for 100GBASE-CR4, 100GBASE-KR4, and 100GBASE-SR4	91.5.3.3	*K*=417	BEI* KR4:M or BEI* SR4:M	Yes [ ] N/A [ ]
RF12	Symbol error threshold for 100GBASE-KP4	91.5.3.3	*K*=6380	BEI* KP4:M	Yes [ ] N/A [ ]
RF13	Error monitoring during LPI	91.5.3.3	Error monitor disabled when rx_lpi_active=true	BEI* EEE:M	Yes [ ] N/A [ ]
RF14	Start of error monitoring window	91.5.3.3	Begins on the codeword boundary following the transition of rx_lpi_active from true to false	BEI* EEE:M	Yes [ ] N/A [ ]
RF15	Alignment marker removal	91.5.3.4	am_rxmapped removed prior to transcoding	M	Yes [ ]
RF16	256B/257B to 64B/66B transcoder	91.5.3.5	rx_coded_*j*<65:0>, *j*=0 to 3 constructed per 91.5.3.5	M	Yes [ ]
RF17	Block distribution	91.5.3.6	One 66-bit block at a time in a round robin fashion from the lowest to the highest numbered PCS lane	M	Yes [ ]
RF18	Alignment marker mapping	91.5.3.7	Map to am_rx_*x*, *x*=0 to 19 per 91.5.3.7	M	Yes [ ]
RF19	Alignment marker insertion point	91.5.3.7	Alignment markers immediately followed by the 66-bit blocks derived from the 257-blocks immediately following am_rxmapped	M	Yes [ ]

91.7.4.3 State diagrams

Item	Feature	Subclause	Value/Comment	Status	Support
SD1	SLIP function	91.5.4.2.2	Ensure that all possible block positions are evaluated	M	Yes [ ]
SD2	Synchronization process	91.5.4.3	One instance per FEC lane per Figure 91–8	M	Yes [ ]
SD3	Alignment process	91.5.4.3	Per Figure 91–9	M	Yes [ ]
SD4	Transmit LPI process	91.5.4.3	Per Figure 91–10	EEE:M	Yes [ ] N/A [ ]
SD5	Receive LPI process	91.5.4.3	Per Figure 91–11	EEE:M	Yes [ ] N/A [ ]

92. Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-CR4

92.1 Overview

This clause specifies the 100GBASE-CR4 PMD and baseband medium. Annex 92A, an associated annex, provides information on parameters with test points that may not be testable in an implemented system.

When forming a complete Physical Layer, a PMD shall be connected as illustrated in Figure 92–1, to the appropriate PMA as shown in Table 92–1, to the medium through the MDI and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 92–1—Physical Layer clauses associated with the 100GBASE-CR4 PMD

Associated clause	100GBASE-CR4
81—RS	Required
81—CGMII[a]	Optional
82—PCS for 100GBASE-R	Required
91—RS-FEC	Required
83—PMA for 100GBASE-R[b]	Required
83A—CAUI-10	Optional
83D—CAUI-4	Optional
73—Auto-Negotiation	Required
78—Energy Efficient Ethernet	Optional

[a]The CGMII is an optional interface. However, if the CGMII is not implemented, a conforming implementation must behave functionally as though the RS and CGMII were present.
[b]There are limitations on the number of PMA lanes that may be used between sublayers, see 83.

When forming a complete 100GBASE-CR4 Physical Layer, the following guidelines apply.

Differential signals received at the MDI from a transmitter that meets the requirements of 92.8.3 and have passed through the cable assembly specified in 92.10 are received with a BER less than 10^{-5}.

For a complete Physical Layer, this specification is considered to be satisfied by a frame loss ratio (see 1.4.223) less than 6.2×10^{-10} for 64-octet frames with minimum interpacket gap.

A 100GBASE-CR4 PHY with the optional Energy-Efficient Ethernet (EEE) capability may optionally enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization.

Figure 92–1 shows the relationship of the 100GBASE-CR4 PMD sublayer and MDI to the ISO/IEC Open System Interconnection (OSI) reference model.

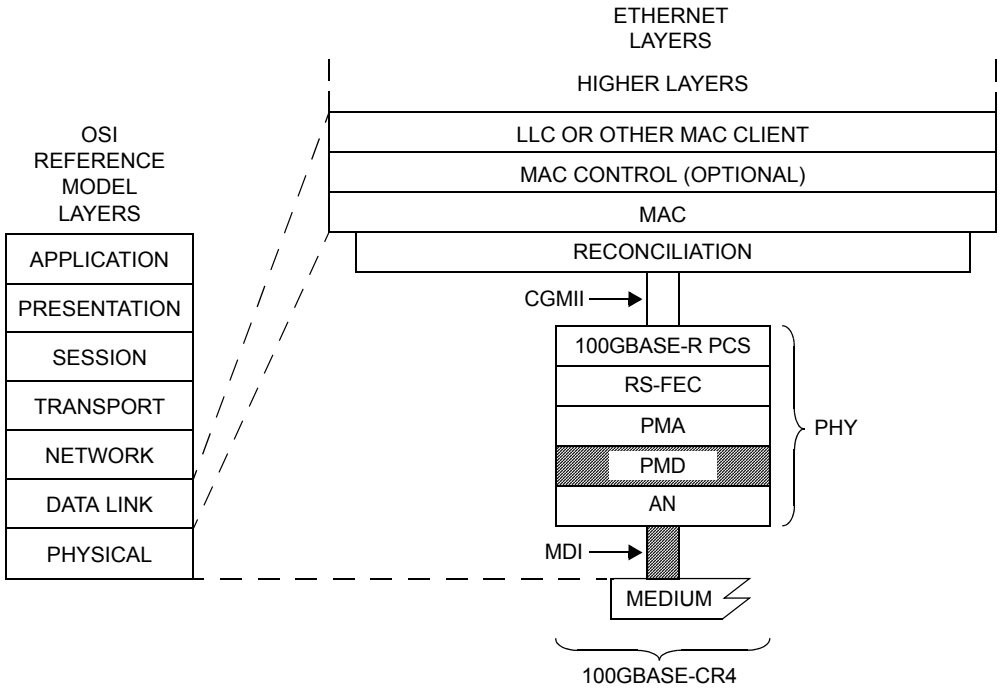

AN = AUTO-NEGOTATION
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION

Figure 92–1—100GBASE-CR4 PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

92.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 100GBASE-CR4 PMD. The service interface for this PMD is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data. The PMD translates the encoded data to and from signals suitable for the medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_*i*.request
PMD:IS_UNITDATA_*i*.indication
PMD:IS_SIGNAL.indication

The 100GBASE-CR4 PMD has four parallel bit streams, hence $i = 0$ to 3. The PMA (or the PMD) continuously sends four parallel bit streams to the PMD (or the PMA), one per lane, each at a nominal signaling rate of 25.78125 GBd.

The SIGNAL_OK parameter of the PMD:IS_SIGNAL.indication primitive corresponds to the variable Global_PMD_signal_detect as defined in 92.7.4. When Global_PMD_signal_detect is one, SIGNAL_OK shall be assigned the value OK. When Global_PMD_signal_detect is zero, SIGNAL_OK shall be assigned the value FAIL. When SIGNAL_OK is FAIL, the PMD:IS_UNITDATA_i.indication parameters are undefined.

If the optional EEE deep sleep capability is supported, then the PMD service interface includes two additional primitives as follows:

PMD:IS_TX_MODE.request
PMD:IS_RX_MODE.request

92.3 PCS requirements for Auto-Negotiation (AN) service interface

The PCS associated with this PMD is required to support the AN service interface primitive AN_LINK.indication defined in 73.9. (See 82.6.)

The 100GBASE-CR4 PHY may be extended using CAUI-n as a physical instantiation of the inter-sublayer service interface between devices. If CAUI-n is instantiated, the AN_LINK(link_status).indication is relayed from the device with the PCS sublayer to the device with the AN sublayer by means at the discretion of the implementer. As examples, the implementer may employ use of pervasive management or employ a dedicated electrical signal to relay the state of link_status as indicated by the PCS sublayer on one device to the AN sublayer on the other device.

92.4 Delay constraints

The sum of the transmit and the receive delays at one end of the link contributed by the 100GBASE-CR4 PMD and AN shall be no more than 2048 bit times (4 pause_quanta or 20.48 ns). It is assumed that the one way delay through the medium is no more than 6000 bit times (60 ns).

A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

92.5 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the RS-FEC sublayer. Skew and Skew Variation are defined in 80.5 and specified at the points SP1 to SP6 shown in Figure 80–8.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 80.5.

92.6 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several registers that provide control and status information for and about the PMD. If MDIO is implemented, it shall map MDIO control bits to PMD control variables as shown in Table 92–2, and MDIO status bits to PMD status variables as shown in Table 92–3.

Table 92–2—100GBASE-CR4 MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable	1.9.0	Global_PMD_transmit_disable
PMD transmit disable 3 to PMD transmit disable 0	PMD transmit disable	1.9.4 to 1.9.1	PMD_transmit_disable_3 to PMD_transmit_disable_0
Restart training	BASE-R PMD control	1.150.0	mr_restart_training
Training enable	BASE-R PMD control	1.150.1	mr_training_enable
Polynomial identifier 3	PMD training pattern 3	1.1453.12:11	identifier_3
Seed 3	PMD training pattern 3	1.1453.10:0	seed_3
Polynomial identifier 2	PMD training pattern 2	1.1452.12:11	identifier_2
Seed 2	PMD training pattern 2	1.1452.10:0	seed_2
Polynomial identifier 1	PMD training pattern 1	1.1451.12:11	identifier_1
Seed 1	PMD training pattern 1	1.1451.10:0	seed_1
Polynomial identifier 0	PMD training pattern 0	1.1450.12:11	identifier_0
Seed 0	PMD training pattern 0	1.1450.10:0	seed_0

Table 92–3—100GBASE-CR4 MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Fault	PMA/PMD status 1	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect	1.10.0	Global_PMD_signal_detect
PMD receive signal detect 3 to PMD receive signal detect 0	PMD receive signal detect	1.10.4 to 1.10.1	PMD_signal_detect_3 to PMD_signal_detect_0
100GBASE-CR4 deep sleep	EEE capability	1.16.11	—
Receiver status 3	BASE-R PMD status	1.151.12	rx_trained_3
Frame lock 3	BASE-R PMD status	1.151.13	frame_lock_3

Table 92–3—100GBASE-CR4 MDIO/PMD status variable mapping *(continued)*

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Start-up protocol status 3	BASE-R PMD status	1.151.14	training_3
Training failure 3	BASE-R PMD status	1.151.15	training_failure_3
Receiver status 2	BASE-R PMD status	1.151.8	rx_trained_2
Frame lock 2	BASE-R PMD status	1.151.9	frame_lock_2
Start-up protocol status 2	BASE-R PMD status	1.151.10	training_2
Training failure 2	BASE-R PMD status	1.151.11	training_failure_2
Receiver status 1	BASE-R PMD status	1.151.4	rx_trained_1
Frame lock 1	BASE-R PMD status	1.151.5	frame_lock_1
Start-up protocol status 1	BASE-R PMD status	1.151.6	training_1
Training failure 1	BASE-R PMD status	1.151.7	training_failure_1
Receiver status 0	BASE-R PMD status	1.151.0	rx_trained_0
Frame lock 0	BASE-R PMD status	1.151.1	frame_lock_0
Start-up protocol status 0	BASE-R PMD status	1.151.2	training_0
Training failure 0	BASE-R PMD status	1.151.3	training_failure_0

92.7 PMD functional specifications

92.7.1 Link block diagram

A 100GBASE-CR4 link in one direction is illustrated in Figure 92–2. For purposes of system conformance, the PMD sublayer is standardized at the test points described in this subclause. The electrical transmit signal is defined at TP2. Unless specified otherwise, all transmitter measurements and tests defined in 92.8.3 are made at TP2 utilizing the test fixture specified in 92.11.1. Unless specified otherwise, all receiver measurements and tests defined in 92.8.4 are performed at TP3 utilizing the test fixture specified in 92.11.1. A mated connector pair has been included in both the transmitter and receiver specifications defined in 92.8.3 and 92.8.4. The recommended maximum insertion loss from TP0 to TP2 or TP3 to TP5 including the test fixture is specified in 92.8.3.6.

The 100GBASE-CR4 channel is defined between the transmitter (TP0) and receiver (TP5) blocks to include the transmitter and receiver differential controlled impedance printed circuit board insertion loss and the cable assembly insertion loss, as illustrated in Figure 92–2. Annex 92A provides information on parameters associated with test points TP0 and TP5 that may not be testable in an implemented system. All cable assembly measurements are to be made between TP1 and TP4 as illustrated in Figure 92–2. The cable assembly test fixture of Figure 92–17, or its equivalent, is required for measuring the cable assembly specifications in 92.10 at TP1 and TP4. Two mated connector pairs and the cable assembly test fixture have been included in the cable assembly specifications defined in 92.10. Transmitter and receiver differential controlled impedance printed circuit board insertion losses defined between TP0–TP1 and TP4–TP5, respectively, are provided informatively in Annex 92A.

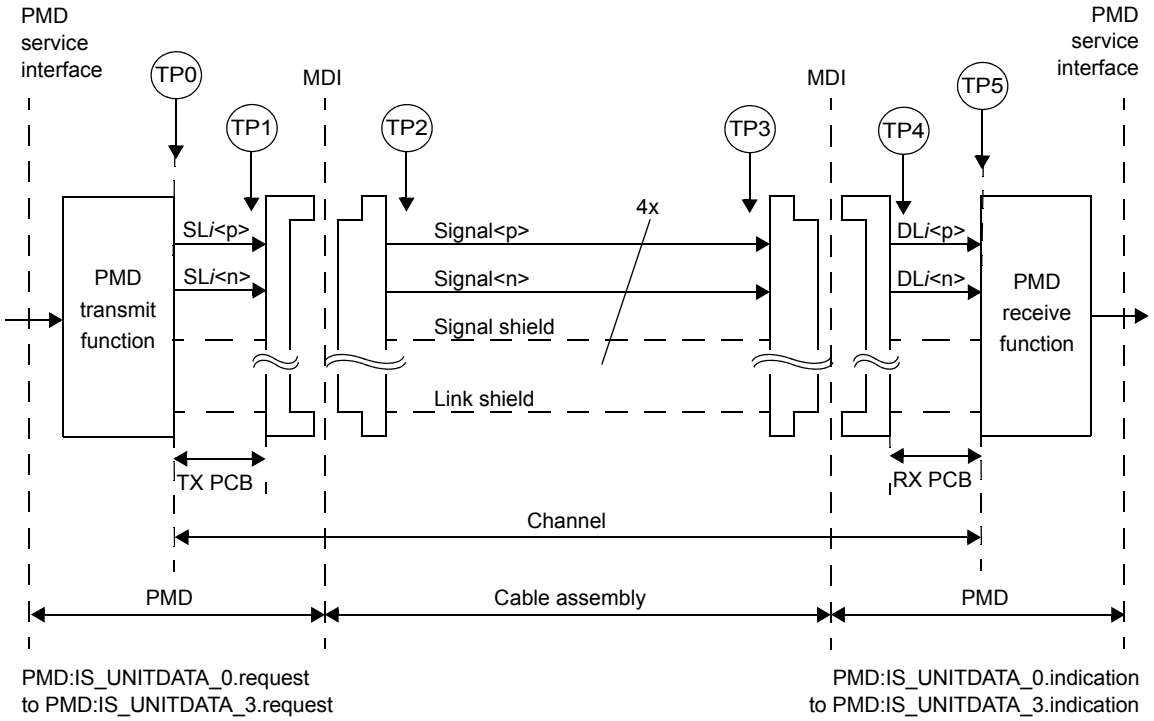

Figure 92–2—100GBASE-CR4 link (one direction is illustrated)

Note that the source lanes (SL), signals SL*i*<p>, and SL*i*<n> are the positive and negative sides of the transmitters differential signal pairs and the destination lanes (DL) signals, DL*i*<p>, and DL*i*<n> are the positive and negative sides of the receivers differential signal pairs for lane *i* (*i* = 0, 1, 2, 3).

Table 92–4 describes the defined test points illustrated in Figure 92–2.

Table 92–4—100GBASE-CR4 test points

Test points	Description
TP0 to TP5	The 100GBASE-CR4 channel including the transmitter and receiver differential controlled impedance printed circuit board insertion loss and the cable assembly insertion loss.
TP1 to TP4	All cable assembly measurements are to be made between TP1 and TP4 as illustrated in Figure 92–2. The cable assembly test fixture of Figure 92–17 or its equivalent, is required for measuring the cable assembly specifications in 92.10 at TP1 and TP4.
TP0 to TP2 TP3 to TP5	A mated connector pair has been included in both the transmitter and receiver specifications defined in 92.8.3 and 92.8.4. The recommended maximum insertion loss from TP0 to TP2 or TP3 to TP5 including the test fixture is specified in 92.8.3.6.
TP2	Unless specified otherwise, all transmitter measurements defined in Table 92–6 are made at TP2 utilizing the test fixture specified in 92.11.1.
TP3	Unless specified otherwise, all receiver measurements and tests defined in 92.8.4 are made at TP3 utilizing the test fixture specified in 92.11.1.

92.7.2 PMD Transmit function

The PMD transmit function shall convert the four bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_*i*.request (*i*=0 to 3) into four separate electrical signals. The four electrical signals shall then be delivered to the MDI, all according to the transmit electrical specifications in 92.8.3. A positive differential output voltage (SL*i*<p> minus SL*i*<n>) shall correspond to tx_bit = one.

If the optional EEE deep sleep capability is supported, the following requirements apply. When tx_mode is set to ALERT, the PMD transmit function shall transmit a periodic sequence, where each period of the sequence consists of 8 ones followed by 8 zeros, on each lane, with the transmit equalizer coefficients set to the preset values (see 72.6.10.2.3.1). This sequence is transmitted regardless of the value of tx_bit presented by the PMD:IS_UNITDATA_*i*.request primitive. When tx_mode is not set to ALERT, the transmit equalizer coefficients are set to the values determined via the start-up protocol (see 92.7.12).

92.7.3 PMD Receive function

The PMD receive function shall convert the four electrical signals from the MDI into four bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_*i*.indication (*i*=0 to 3). A positive differential input voltage (DL*i*<p> minus DL*i*<n>) shall correspond to rx_bit = one.

92.7.4 Global PMD signal detect function

The variable Global_PMD_signal_detect is the logical AND of the values of PMD_signal_detect_*i* for *i*=0 to 3.

When the MDIO is implemented, this function maps the variable Global_PMD_signal_detect to the register and bit defined in 92.6.

92.7.5 PMD lane-by-lane signal detect function

The PMD lane-by-lane signal detect function is used by the 100GBASE-CR4 PMD to indicate the successful completion of the start-up protocol by the PMD control function (see 92.7.12). PMD_signal_detect_*i* (where *i* represents the lane number in the range 0 to 3) is set to zero when the value of the variable signal_detect is set to false by the Training state diagram for lane *i* (see Figure 72–5). PMD_signal_detect_*i* is set to one when the value of signal_detect for lane *i* is set to true.

If training is disabled by the management variable mr_training_enable (see 92.6), PMD_signal_detect_*i* shall be set to one for *i*=0 to 3.

If the optional EEE deep sleep capability is supported, the following requirements apply. The value of PMD_signal_detect_*i* (for *i*=0 to 3) is set to zero when rx_mode is first set to QUIET. While rx_mode is set to QUIET, PMD_signal_detect_*i* shall be set to one within 500 ns of the application of the ALERT pattern defined in 92.7.2, with peak-to-peak differential voltage of 720 mV measured at TP2, to the differential pair at the input of the cable assembly that connects the transmitter to the receiver of lane *i*. While rx_mode is set to QUIET, PMD_signal_detect_*i* shall not be set to one when the voltage input to the differential pair of the cable assembly that connects the transmitter to the receiver of lane *i* is less than or equal to 70 mV peak-to-peak differential.

When the MDIO is implemented, this function maps the variables to registers and bits as defined in 92.6.

92.7.6 Global PMD transmit disable function

The Global PMD transmit disable function is mandatory if EEE deep sleep capability is supported and is otherwise optional. When implemented, it allows all of the transmitters to be disabled with a single variable.

a) When 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 92–6.
b) If a PMD fault (92.7.9) is detected, then the PMD may set Global_PMD_transmit_disable to one.
c) Loopback, as defined in 92.7.8, shall not be affected by Global_PMD_transmit_disable.
d) The following additional requirements apply when the optional EEE deep sleep capability is supported. The Global PMD transmit disable function shall turn off all of the transmitters as specified in 92.8.3.1 when tx_mode transitions to QUIET from any other value. The Global PMD transmit disable function shall turn on all of the transmitters as specified in 92.8.3.1 when tx_mode transitions from QUIET to any other value.

92.7.7 PMD lane-by-lane transmit disable function

The PMD lane-by-lane transmit disable function is optional and allows the electrical transmitter in each lane to be selectively disabled. When this function is supported, it shall meet the following requirements:

a) When a PMD_transmit_disable_*i* variable (where *i* represents the lane number in the range 0 to 3) is set to one, this function shall turn off the transmitter associated with that variable such that it drives a constant level (i.e., no transitions) and does not exceed the maximum differential peak-to-peak output voltage specified in Table 92–6.
b) If a PMD fault (92.7.9) is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the electrical transmitter in each lane.
c) Loopback, as defined in 92.7.8, shall not be affected by PMD_transmit_disable_*i*.

92.7.8 Loopback mode

Local loopback mode is provided by the adjacent PMA (see 83.5.8) as a test function. When loopback mode is enabled, transmission requests passed to each transmitter are sent directly to the corresponding receiver, overriding any signal detected by each receiver on its attached link. Note that loopback mode does not affect the state of the transmitter, which continues to send data (unless disabled).

Control of the loopback function is specified in 83.5.8.

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.

92.7.9 PMD fault function

If the MDIO is implemented, PMD_fault shall be mapped to the fault bit as specified in 45.2.1.2.3. PMD_fault is the logical OR of PMD_receive_fault, PMD_transmit_fault, and any other implementation specific fault.

92.7.10 PMD transmit fault function

The PMD transmit fault function is optional. The faults detected by this function are implementation specific, but the assertion of Global_PMD_transmit_disable is not considered a transmit fault.

If PMD_transmit_fault is set to one, then Global_PMD_transmit_disable should also be set to one.

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.

92.7.11 PMD receive fault function

The PMD receive fault function is optional. The faults detected by this function are implementation specific. A fault is indicated by setting the variable PMD_receive_fault to one.

If the MDIO interface is implemented, then PMD_receive_fault shall be mapped to the Receive fault bit as specified in 45.2.1.7.5.

92.7.12 PMD control function

Each lane of the 100GBASE-CR4 PMD shall use the same control function as 10GBASE-KR, as defined in 72.6.10, with the following differences:

a) The training frame structure used by the 100GBASE-CR4 PMD control function shall be as defined in 72.6.10 with the exception that 25.78125 GBd symbols replace 10.3125 GBd symbols and 100GBASE-CR4 UI replace 10GBASE-KR UI.
b) In addition to the coefficient update process specified in 72.6.10.2.5, the period from receiving a new request to responding to that request shall be less than 2 ms, except during the first 50 ms following the beginning of the start-up protocol. The beginning of the start-up protocol is defined to be entry into the AN_GOOD_CHECK state in Figure 73–10. The start of the period is the frame marker of the training frame with the new request and the end of the period is the frame marker of the training frame with the corresponding response. A new request occurs when the coefficient update field is different from the coefficient field in the preceding frame. The response occurs when the coefficient status report field is updated to indicate that the corresponding action is complete.
c) In addition, the training pattern defined in 72.6.10.2.6 is replaced with a set of training patterns designed to minimize the correlation between physical lanes. The training pattern for each lane shall consist of 4094 bits from the output of a pseudo-random bit sequence (PRBS) generator followed by two zeros. The PRBS generator for each lane shall implement each of the four generator polynomials (polynomial_*n* where *n* goes from 0 to 3) given in Table 92–5, selectable by identifier_*i* (where *i* is the lane number). By default identifier_*i* is set to lane number *i* (i.e., identifier_0 = 0; identifier_1 = 1, etc.). At the start of the training pattern, the state of the generator shall be set to the value in seed_*i* (where *i* is the lane number), with the default values given in Table 92–5. Note that a seed value of 0x000 is invalid. An example implementation of the PRBS generator for *n* = 0 with default settings is given in Figure 92–3. The first 32 bits of the training pattern for each polynomial is also provided in Table 92–5.

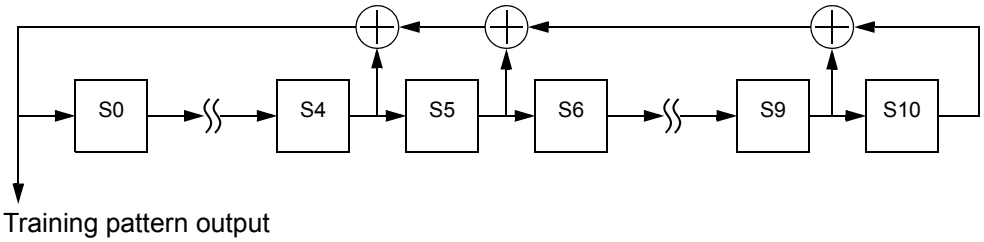

Figure 92–3—PRBS generator for polynomial_0

Table 92–5—PRBS parameters for each physical lane

n	Polynomial_*n*, *G*(*x*)	Default seed bits, S0 is the left most bit	Initial output[a]
0	$1 + x^5 + x^6 + x^{10} + x^{11}$	10101111110	fbf1cb3e
1	$1 + x^5 + x^6 + x^9 + x^{11}$	11001000101	fbb1e665
2	$1 + x^4 + x^6 + x^8 + x^{11}$	11100101101	f3fdae46
3	$1 + x^4 + x^6 + x^7 + x^{11}$	11110110110	f2ffa46b

[a]The first 32 bits of the training pattern are presented in a hexadecimal representation where the hex symbols are transmitted from left to right and the most significant bit of each hex symbol is transmitted first

The variables rx_trained_*i*, frame_lock_*i*, training_*i,* and training_failure_*i* (where *i* goes from 0 to 3) report status for each lane and are equivalent to rx_trained, frame_lock, training, and training_failure as defined in 72.6.10.3.1.

If the MDIO interface is implemented, then this function shall map the variables polynomial_*i*, seed_*i*, rx_trained_*i*, frame_lock_*i*, training_*i,* and training_failure_*i* to the registers and bits defined in 92.6.

92.8 100GBASE-CR4 electrical characteristics

92.8.1 Signal levels

The 100GBASE-CR4 MDI is a low-swing AC-coupled differential interface. AC-coupling within the plug connectors, as defined in 92.12.1, allows for interoperability between components operating from different supply voltages.

92.8.2 Signal paths

The 100GBASE-CR4 MDI signal paths are point-to-point connections. Each path corresponds to a 100GBASE-CR4 MDI lane and comprises two complementary signals, which form a balanced differential pair. For 100GBASE-CR4, 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 ranging from 0.5 m to 5 m in length, as described in 92.10.

92.8.3 Transmitter characteristics

Transmitter characteristics are summarized in Table 92–6. Unless specified otherwise, all transmitter measurements defined in Table 92–6 are made at TP2 utilizing the test fixtures specified in 92.11.1. A test system with a fourth-order Bessel-Thomson low-pass response with 33 GHz 3 dB bandwidth is to be used for all transmitter signal measurements, unless otherwise specified. The transmitter specifications at TP0 are provided informatively in Annex 92A.

Table 92–6—Transmitter characteristics at TP2 summary

Parameter	Subclause reference	Value	Units
Differential peak-to-peak output voltage (max.) with Tx disabled	92.8.3.1	35	mV
DC common-mode voltage (max.)	92.8.3.1	1.9	V
AC common-mode output voltage, v_{cmi} (max., RMS)	92.8.3.1	30	mV
Differential peak-to-peak voltage, v_{di} (max.)	92.8.3.1	1200	mV
Differential output return loss (min.)	92.8.3.2	See Equation (92–1)	dB
Common-mode to differential mode output return loss (min.)	92.8.3.3	See Equation (92–2)	dB
Common-mode to common-mode output return loss (min.)	92.8.3.4	See Equation (92–3)	dB
Transmitter steady-state voltage, v_f (min.) Transmitter steady-state voltage, v_f (max.)	92.8.3.5.2	0.34 0.6	V
Linear fit pulse peak (min.)	92.8.3.5.2	0.45 × v_f	V
Transmitted waveform abs coefficient step size (min.) abs coefficient step size (max.) minimum precursor full-scale ratio minimum post cursor full-scale ratio	 92.8.3.5.4 92.8.3.5.4 92.8.3.5.5 92.8.3.5.5	 0.0083 0.05 1.54 4	
Signal-to-noise-and-distortion ratio (min.)	92.8.3.7	26	dB
Output jitter (max.) Even-odd jitter, peak-to-peak Effective bounded uncorrelated jitter, peak-to-peak Effective total uncorrelated jitter, peak-to-peak	 92.8.3.8.1 92.8.3.8.2 92.8.3.8.2	 0.035 0.1 0.18	 UI UI UI
Signaling rate, per lane	92.8.3.9	25.78125±100 ppm	GBd
Unit interval nominal	92.8.3.9	38.787879	ps

92.8.3.1 Signal levels

The differential output voltage v_{di} is defined to be SL*i*<p> minus SL*i*<n>. The common-mode output voltage v_{cmi} is defined to be one half of the sum of SL*i*<p> and SL*i*<n>. These definitions are illustrated by Figure 92–4.

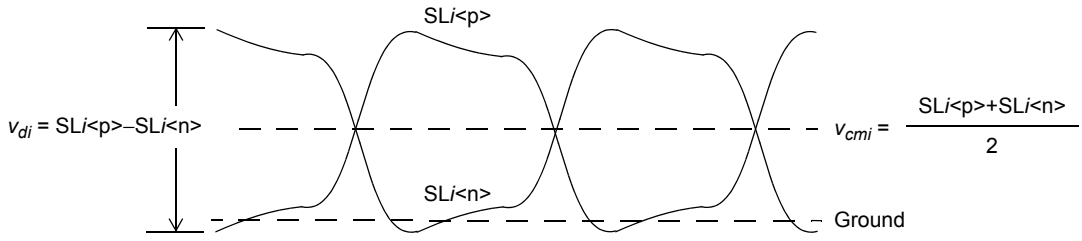

Figure 92–4—Transmitter output voltage definitions

The peak-to-peak differential output voltage shall be less than or equal to 1200 mV regardless of the transmit equalizer setting. The peak-to-peak differential output voltage shall be less than or equal to 35 mV while the transmitter is disabled (refer to 92.7.6 and 92.7.7).

The 100GBASE-CR4 Style-1 connector may support 100GBASE-CR4 or XLPPI interfaces. For implementations that support both interfaces, the transmitter should not exceed the XLPPI voltage maximum until a 100GBASE-CR4 cable assembly has been identified.

The DC common-mode output voltage shall be between 0 V and 1.9 V with respect to signal ground. The AC common-mode output voltage shall be less than or equal to 30 mV RMS with respect to signal ground. Common-mode output voltage requirements shall be met regardless of the transmit equalizer setting.

If the optional EEE deep sleep capability is supported the following requirements also apply. The peak-to-peak differential output voltage shall be less than 35 mV within 500 ns of the transmitter being disabled. The peak-to-peak differential output voltage shall be greater than 720 mV within 500 ns of the transmitter being enabled. The transmitter is enabled by the assertion of tx_mode=ALERT, and the preceding requirement applies when the transmitted symbols are the periodic pattern defined in 92.8.1 and the transmitter equalizer coefficients are assigned their preset values. The transmitter shall meet the requirements of 92.8.3 within 1 μs of the transmitter being enabled. While the transmitter is disabled, the DC common-mode output voltage shall be maintained to within ±150 mV of the value for the enabled transmitter.

Differential and common-mode signal levels are measured with a PRBS9 test pattern.

92.8.3.2 Transmitter differential output return loss

The differential output return loss, in dB, of the transmitter shall meet Equation (92–1). This output impedance requirement applies to all valid output levels. The reference impedance for differential return loss measurements shall be 100 Ω.

$$Return\_loss(f) \geq \begin{cases} 9.5 - 0.37f & 0.01 \leq f < 8 \\ 4.75 - 7.4\log_{10}(f/14) & 8 \leq f \leq 19 \end{cases} \quad \text{(dB)} \tag{92–1}$$

where

f is the frequency in GHz

$Return\_loss(f)$ is the differential output return loss at frequency f

The transmitter differential output return loss is illustrated in Figure 92–5.

92.8.3.3 Common-mode to differential mode output return loss

The common-mode to differential mode output return loss, in dB, of the transmitter shall meet Equation (92–2).

$$Return\_loss(f) \geq \begin{cases} 22 - (20/25.78)f & 0.01 \leq f < 12.89 \\ 15 - (6/25.78)f & 12.89 \leq f \leq 19 \end{cases} \quad \text{(dB)} \tag{92–2}$$

where

f is the frequency in GHz

$Return\_loss(f)$ is the common-mode to differential mode output return loss at frequency f

The common-mode to differential mode output return loss is illustrated in Figure 92–6.

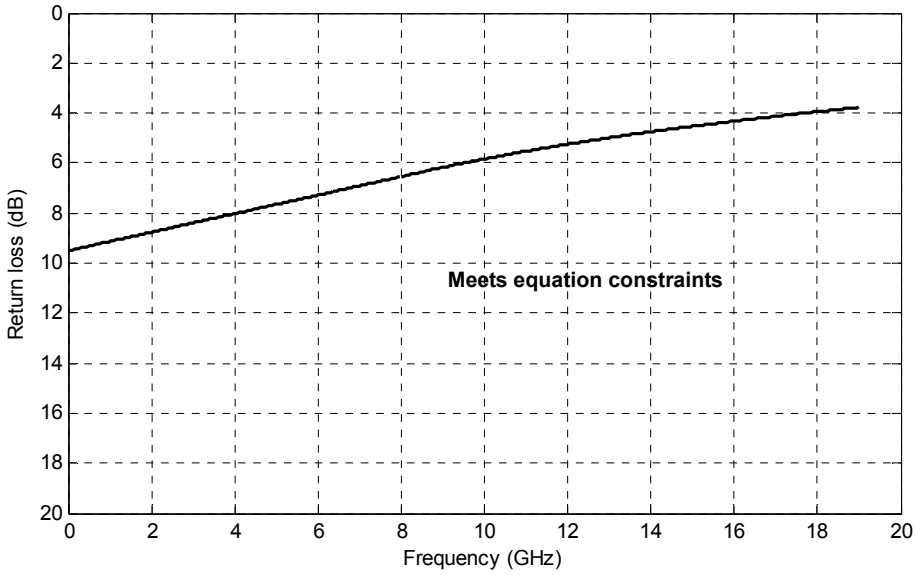

Figure 92–5—Transmitter differential output return loss

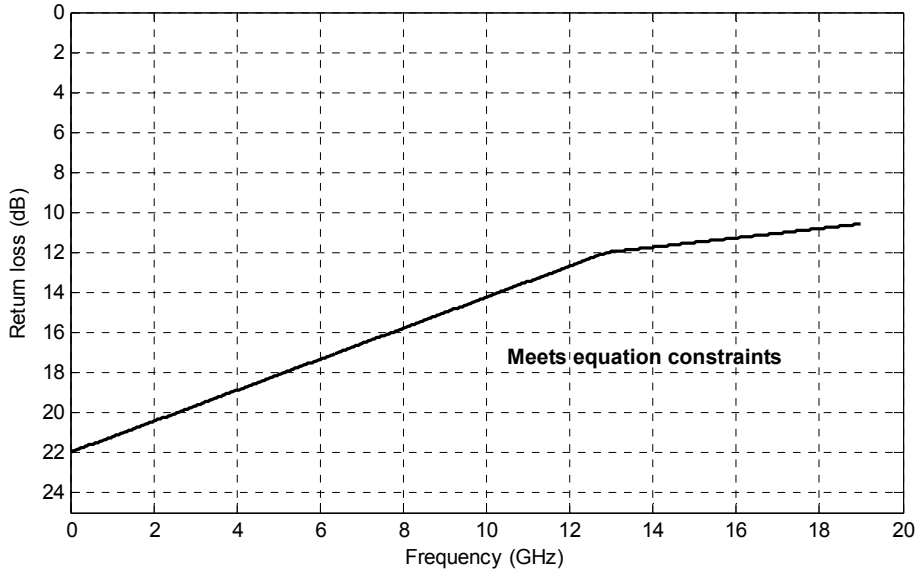

Figure 92–6—Common-mode to differential mode return loss

92.8.3.4 Common-mode to common-mode output return loss

The common-mode to common-mode output return loss, in dB, of the transmitter shall meet Equation (92–3).

$$Return\_loss(f) \geq 2 \quad \text{(dB)} \tag{92–3}$$

for $0.2 \leq f \leq 19$ GHz

where

f is the frequency in GHz

$Return\_loss(f)$ is the common-mode to common-mode return loss at frequency f

92.8.3.5 Transmitter output waveform

The 100GBASE-CR4 transmit function includes programmable equalization to compensate for the frequency-dependent loss of the channel and facilitate data recovery at the receiver. The functional model for the transmit equalizer is the three tap transversal filter shown in Figure 92–7.

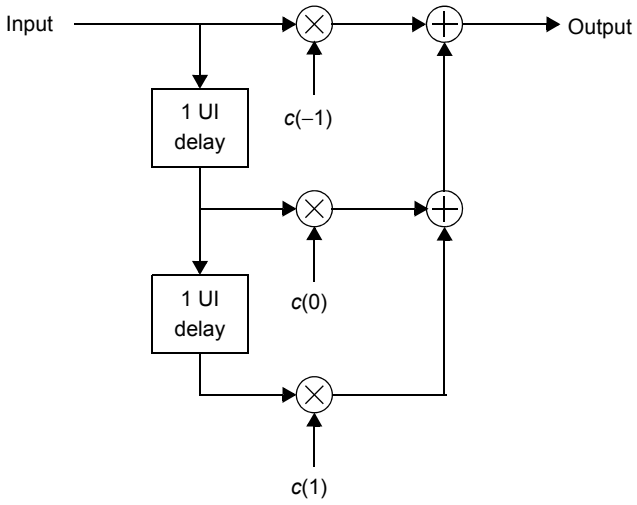

Figure 92–7—Transmit equalizer functional model

The state of the transmit equalizer and hence the transmitted output waveform may be manipulated via the PMD control function defined in 92.7.12 or via the management interface. The transmit function responds to a set of commands issued by the link partner's receive function and conveyed by a back-channel communications path.

This command set includes instructions to

a) Increment coefficient $c(i)$.
b) Decrement coefficient $c(i)$.
c) Hold coefficient $c(i)$ at its current value.
d) Set the coefficients to a predefined value (preset or initialize).

In response, the transmit function relays status information to the link partner's receive function. The status messages indicate that

a1) The requested update to coefficient $c(i)$ has completed (updated).
b1) Coefficient $c(i)$ is at its minimum value.
c1) Coefficient $c(i)$ is at its maximum value.
d1) Coefficient $c(i)$ is ready for the next update request (not_updated).

92.8.3.5.1 Linear fit to the measured waveform

For each configuration of the transmit equalizer, capture at least one complete cycle of the test pattern PRBS9 as specified in 83.5.10 at TP2 per 85.8.3.3.4. In the following calculation, M should be an integer not less than 32. Interpolation of the captured waveform may be used to achieve this. Compute the linear fit pulse response $p(k)$ from the captured waveform per 85.8.3.3.5 using $N_p = 14$ and $D_p = 2$. Define $r(k)$ to be the linear fit pulse response when transmit equalizer coefficients have been set to the "preset" values. The normalized coefficients for any configuration of the transmit equalizer are computed using the following method.

Define an MN_p-by-3 matrix R_m. The elements of R_m are assigned values per Equation (92–4) where $i = -1$ to 1, $j = 1$ to MN_p, and $m = -M/2$ to $M/2-1$ when M is even and $-(M-1)/2$ to $(M-1)/2$ when M is odd.

$$R_m(j, i+2) = \begin{cases} r(m+j-iM) & 1 \le m+j-iM \le MN_p \\ 0 & otherwise \end{cases} \quad (92\text{–}4)$$

The normalized coefficients of the transmit equalizer are computed using Equation (92–5).

$$\begin{bmatrix} c_m(-1) \\ c_m(0) \\ c_m(1) \end{bmatrix} = (R_m^T R_m)^{-1} R_m^T \begin{bmatrix} p(1) \\ \dots \\ p(MN_p) \end{bmatrix} \quad (92\text{–}5)$$

The linear fit pulse response is reconstructed from the matrix R_m and the normalized coefficients using Equation (92–6).

$$\begin{bmatrix} p_m(1) \\ \dots \\ p_m(MN_p) \end{bmatrix} = R_m \begin{bmatrix} c_m(-1) \\ c_m(0) \\ c_m(1) \end{bmatrix} \quad (92\text{–}6)$$

The sum of the squared error between $p(k)$ and $p_m(k)$ is computed using Equation (92–7). The normalized transmit equalizer coefficients $c(i)$ for a given linear fit pulse $p(k)$ are the values $c_m(i)$ for the value of m that minimizes $\varepsilon^2(m)$.

$$\varepsilon^2(m) = \sum_{k=1}^{MN_p} (p(k) - p_m(k))^2 \quad (92\text{–}7)$$

92.8.3.5.2 Steady-state voltage and linear fit pulse peak

The steady-state voltage v_f is defined to be the sum of the linear fit pulse $p(k)$ divided by M (refer to 85.8.3.3 step 3). The steady-state voltage shall be greater than or equal to 0.34 V and less than or equal to 0.6 V after the transmit equalizer coefficients have been set to the "preset" values.

The peak value of $p(k)$ shall be greater than $0.45 \times v_f$ after the transmit equalizer coefficients have been set to the "preset" values.

92.8.3.5.3 Coefficient initialization

When the PMD enters the INITIALIZE state of the Training state diagram (Figure 72–5) or receives a valid request to "initialize" from the link partner, the coefficients of the transmit equalizer shall be configured such that the ratio $(c(0)+c(1)–c(–1))/(c(0)+c(1)+c(–1))$ is $1.29 \pm 10\%$ and the ratio $(c(0)–c(1)+c(–1))/(c(0)+c(1)+c(–1))$ is $2.57 \pm 10\%$. These requirements apply upon the assertion of a coefficient status report of "updated" for all coefficients.

92.8.3.5.4 Coefficient step size

The change in the normalized amplitude of coefficient $c(i)$ corresponding to a request to "increment" that coefficient shall be between 0.0083 and 0.05. The change in the normalized amplitude of coefficient $c(i)$ corresponding to a request to "decrement" that coefficient shall be between –0.0083 and –0.05.

The change in the normalized amplitude of the coefficient is defined to be the difference in the value measured prior to the assertion of the "increment" or "decrement" request (e.g., the coefficient update request for all coefficients is "hold") and the value upon the assertion of a coefficient status report of "updated" for that coefficient.

92.8.3.5.5 Coefficient range

When sufficient "increment" or "decrement" requests have been received for a given coefficient, the coefficient reaches a lower or upper bound based on the coefficient range or restrictions placed on the minimum steady-state differential output voltage or the maximum peak-to-peak differential output voltage.

With $c(–1)$ set to zero and both $c(0)$ and $c(1)$ having received sufficient "decrement" requests so that they are at their respective minimum values, the ratio $(c(0) – c(1))/(c(0) + c(1))$ shall be greater than or equal to 4.

With $c(1)$ set to zero and both $c(–1)$ and $c(0)$ having received sufficient "decrement" requests so that they are at their respective minimum values, the ratio $(c(0) – c(–1))/(c(0) + c(–1))$ shall be greater than or equal to 1.54.

Note that a coefficient may be set to zero by first asserting a coefficient preset request and then manipulating the other coefficients as required by the test.

92.8.3.6 Insertion loss TP0 to TP2 or TP3 to TP5

The recommended maximum insertion loss from TP0 to TP2 or TP3 to TP5 including the test fixture is given by Equation (92–8). Note that the recommended maximum insertion loss from TP0 to TP2 or from TP3 to TP5 is 9.85 dB at 12.8906 GHz.

$$Insertion\_loss(f) \le \left\{ \begin{array}{ll} 0.08 + 0.57\sqrt{f} + 0.599f & 0.01 \le f < 14 \\ -19.067 + 2.119f & 14 \le f \le 19 \end{array} \right\} \text{ (dB)} \qquad (92–8)$$

where

f is the frequency in GHz

$Insertion\_loss(f)$ is the insertion loss at frequency f

The maximum insertion loss of TP0 to TP2 or TP3 to TP5 is illustrated in Figure 92–8.

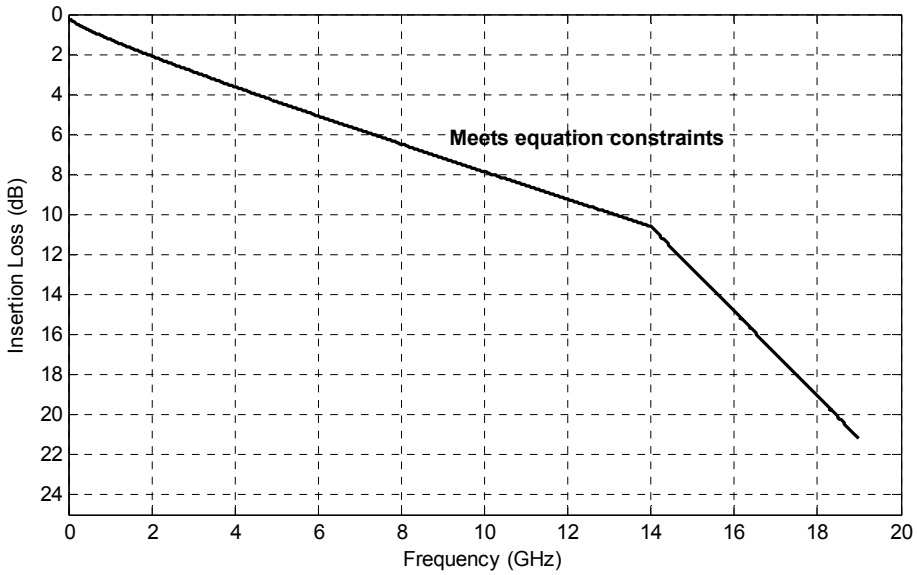

Figure 92–8—Maximum insertion loss TP0 to TP2 or TP3 to TP5

92.8.3.7 Transmitter output noise and distortion

Signal-to-noise-and-distortion ratio (SNDR) is measured at the transmitter output using the following method, with transmitters on all PMD lanes enabled and transmitting the same pattern with identical transmit equalizer settings.

Given a configuration of the transmit equalizer, capture at least one complete cycle of the test pattern PRBS9 as specified in 83.5.10 at TP0a per 85.8.3.3.4. Compute the linear fit pulse response $p(k)$ and the linear fit error waveform $e(k)$ from the captured waveform per 85.8.3.3.5 using $N_p = 14$ and $D_p = 2$. Denote the standard deviation of $e(k)$ as σ_e.

Given the same configuration of the transmit equalizer, measure the RMS deviation from the mean voltage at a fixed point in a run of at least 8 consecutive identical bits in a suitable pattern. PRBS9 is an example of a pattern that includes runs suitable to perform the measurement. It is recommended that the deviation is measured within the flattest portion of the waveform at a point where the slope is closest to zero. The RMS deviation is measured for a run of zeros and also a run of ones. The average of the two measurements is denoted as σ_n.

SNDR is defined by Equation (92–9) where $p_{\max}$ is the maximum value of $p(k)$.

$$SNDR = 10\log_{10}\left(\frac{p_{\max}^2}{\sigma_e^2 + \sigma_n^2}\right) \text{ dB} \tag{92–9}$$

SNDR shall be greater than 26 dB regardless of the transmit equalizer setting.

92.8.3.8 Transmitter output jitter

Three components of the transmitter output jitter are specified in this subclause: even-odd jitter, effective bounded uncorrelated jitter, and effective total uncorrelated jitter.

The effect of a single-pole high-pass filter with a 3 dB frequency of 10 MHz is applied to the jitter. The voltage threshold for the measurement of BER or crossing times is the mid-point (0 V) of the AC-coupled differential signal.

Jitter measurements are performed with transmitters on all PMD lanes enabled and transmitting the same pattern with identical transmit equalizer settings.

92.8.3.8.1 Even-odd jitter

Even-odd jitter is measured using two repetitions of a PRBS9 pattern. The deviation of the time of each transition from an ideal clock at the signaling rate is measured. Even-odd jitter is defined as the magnitude of the difference between the average deviation of all even-numbered transitions and the average deviation of all odd-numbered transitions, where determining if a transition is even or odd is based on possible transitions but only actual transitions are measured and averaged.

Even-odd jitter shall be less than or equal to 0.035 UI regardless of the transmit equalization setting.

NOTE—Even-odd jitter has been referred to as *duty cycle distortion* by other Physical Layer specifications for operation over electrical backplane or twinaxial copper cable assemblies (see 72.7.1.9). The term *even-odd jitter* is used here to distinguish it from the duty cycle distortion referred to by Physical Layer specifications for operation over fiber optic cabling.

92.8.3.8.2 Effective bounded uncorrelated jitter and effective random jitter

Effective bounded uncorrelated jitter and effective random jitter are measured on each of two specific transitions in a PRBS9 pattern (see 83.5.10). The two transitions occur in the sequence of five zeros and four ones and nine ones and five zeros, respectively. The sequences are located at bits 10 to 18 and 1 to 14, respectively, where bits 1 to 9 are the run of nine ones.

a) The jitter components are determined according to the following method. Acquire a horizontal histogram of a transition around the zero-crossing point. The number of acquired samples should be sufficiently large to yield consistent measurement results. Designate the total number of samples as *NS*, the number of bins as *NB*, the number of samples in each bin as N_i where i is the bin number from 1 to *NB*, and the sample time corresponding with the center of each bin as t_i.
b) Create two cumulative distribution curves $CDFL_i$ and $CDFR_i$ according to Equation (92–10) and Equation (92–11) and two corresponding curves QR_i and QL_i according to Equation (92–12) and Equation (92–13), where $\text{erfc}^{-1}(x)$ is the inverse of the complementary error function $\text{erfc}(x)$ defined by Equation (92–14).
c) Determine the parameters m_{left} and b_{left} of Equation (92–15) that best fit QL_i as a function of t_i for bins with $CDFL_i$ in the range of 10^{-3} to 2.5×10^{-2}. Similarly determine the parameters of m_{right} and b_{right} that best fit QR_i as a function of t_i for bins with $CDFR_i$ in the range of 10^{-3} to 2.5×10^{-2}.
d) Calculate the values of effective bounded uncorrelated jitter and effective total uncorrelated jitter according to Equation (92–17) and Equation (92–19), respectively. The peak-to-peak contribution of the effective random jitter in the effective total uncorrelated jitter is related to a bit error ratio of 10^{-5}.

$$CDFL_i = \sum_{k=1}^{i} \frac{N_k}{NS} \tag{92–10}$$

$$CDFR_i = \sum_{k=i}^{NB} \frac{N_k}{NS} \tag{92–11}$$

$$QL_i = \sqrt{2} \cdot \text{erfc}^{-1}(2 \cdot CDFL_i) \quad (92\text{–}12)$$

$$QR_i = \sqrt{2} \cdot \text{erfc}^{-1}(2 \cdot CDFR_i) \quad (92\text{–}13)$$

$$\text{erfc}(x) = \frac{2}{\sqrt{\pi}} \cdot \int_x^{\infty} e^{-t^2} dt \quad (92\text{–}14)$$

$$Q_{left} = m_{\text{left}} \cdot t + b_{\text{left}} \quad (92\text{–}15)$$

$$Q_{right} = m_{\text{right}} \cdot t + b_{\text{right}} \quad (92\text{–}16)$$

$$\text{effective bounded uncorrelated jitter} = EBUJ = b_{\text{left}}/m_{\text{left}} - b_{\text{right}}/m_{\text{right}} \quad (92\text{–}17)$$

$$\text{effective random jitter} = ERJ = \frac{m_{\text{left}} - m_{\text{right}}}{2 \cdot m_{\text{right}} \cdot m_{\text{left}}} \quad (92\text{–}18)$$

$$\text{effective total uncorrelated jitter} = 7.9 \cdot ERJ + EBUJ \quad (92\text{–}19)$$

Effective bounded uncorrelated jitter shall be less than or equal to 0.1 UI peak-to-peak regardless of the transmit equalization setting.

The effective total uncorrelated jitter shall be less than or equal to 0.18 UI peak-to-peak regardless of the transmit equalization setting.

92.8.3.9 Signaling rate range

The 100GBASE-CR4 MDI signaling rate shall be 25.78125 GBd ±100 ppm per lane. The corresponding unit interval is approximately 38.787879 ps.

92.8.4 Receiver characteristics

The receiver characteristics are summarized in Table 92–7. Unless specified otherwise, all receiver measurements defined in Table 92–7 are made at TP3 utilizing the test fixtures specified in 92.11.1. Unless otherwise specified, a test system with a fourth-order Bessel-Thomson low-pass response with 33 GHz 3 dB bandwidth is to be used for all receiver input signal measurements. The receiver specifications at TP5 are provided informatively in Annex 92A.

Table 92–7—Receiver characteristics at TP3 summary

Parameter	Subclause reference	Value	Units
Receiver input amplitude tolerance	92.8.4.1	1200 mV as measured at TP2	mV
Differential input return loss (min)	92.8.4.2	Equation (92–20)	dB
Differential to common-mode input return loss	92.8.4.3	Equation (92–21)	dB
Interference Tolerance	92.8.4.4	Table 92–8	—
Signaling rate, per lane	92.8.4.6	25.78125 ± 100 ppm	GBd
Unit interval (UI) nominal	92.8.4.6	38.787879	ps

92.8.4.1 Receiver input amplitude tolerance

100GBASE-CR4 receiver shall operate at a BER better than 10^{-5} when connected to a compliant transmitter whose peak-to-peak differential output voltage, as defined by 92.8.3.1 using preset equalizer coefficients, is 1200 mV using a compliant cable assembly with the minimum insertion loss defined in 92.10.2. The receiver is allowed to control the transmitter equalizer coefficients, using the protocol defined in 92.7.12 or an equivalent process, to meet this requirement.

92.8.4.2 Receiver differential input return loss

The differential input return loss, in dB, of the receiver shall meet Equation (92–20). This return loss requirement applies at all valid input levels. The reference impedance for differential return loss measurements shall be 100 Ω.

$$Return\_loss(f) \geq \begin{cases} 9.5 - 0.37f & 0.01 \leq f < 8 \\ 4.75 - 7.4\log_{10}(f/14) & 8 \leq f \leq 19 \end{cases} \quad \text{(dB)} \tag{92–20}$$

where

f is the frequency in GHz

$Return\_loss(f)$ is the differential input return loss at frequency f

92.8.4.3 Differential to common-mode input return loss

The differential to common-mode input return loss, in dB, of the receiver shall meet Equation (92–21).

$$Return\_loss(f) \geq \begin{cases} 22 - (20/25.78)f & 0.01 \leq f < 12.89 \\ 15 - (6/25.78)f & 12.89 \leq f \leq 19 \end{cases} \quad \text{(dB)} \tag{92–21}$$

where

f is the frequency in GHz

$Return\_loss(f)$ is the differential to common-mode input return loss at frequency f

92.8.4.4 Receiver interference tolerance test

The receiver interference tolerance of each lane shall comply with both test 1 and test 2 using the parameters of Table 92–8 when measured according to the requirements of 92.8.4.4.1 to 92.8.4.4.5. The cable assembly used in the test channel specified in 92.8.4.4.2 shall meet the cable assembly Channel Operating Margin (COM) specified in 92.10.7.

Table 92–8—100GBASE-CR4 interference tolerance parameters

Parameter	Test 1 values	Test 2 values	Units
RS-FEC symbol error ratio[a]	10^{-4}	10^{-4}	
Fitted insertion loss coefficients	$a_1 = 1.7$ $a_2 = 0.546$ $a_4 = 0.01$	$a_1 = 4.3$ $a_2 = 0.571$ $a_4 = 0.04$	dB/√GHz dB/GHz dB/GHz^2
Applied SJ[b] (peak-to-peak)	0.1	0.1	UI
Applied RJ (RMS)	0.01	0.01	UI
Even-odd jitter	0.035	0.035	UI
COM (max)	3	3	dB

[a]The FEC symbol error ratio is measured in step 11 of the receiver interference tolerance method defined in 93C.2.
[b]Applied SJ frequency >100 MHz, specified at TP0.

92.8.4.4.1 Test setup

The interference tolerance test is performed with the setup shown in Figure 92–9. The requirements of this subclause are verified at the pattern generator connection (PGC) or test references in Figure 92–9 and Figure 92–10. The lanes under test (LUT) are illustrated in Figure 92–9 and Figure 92–10. The cable assembly single-ended receive lanes are terminated in 50 Ω to provide 100 Ω differential termination.

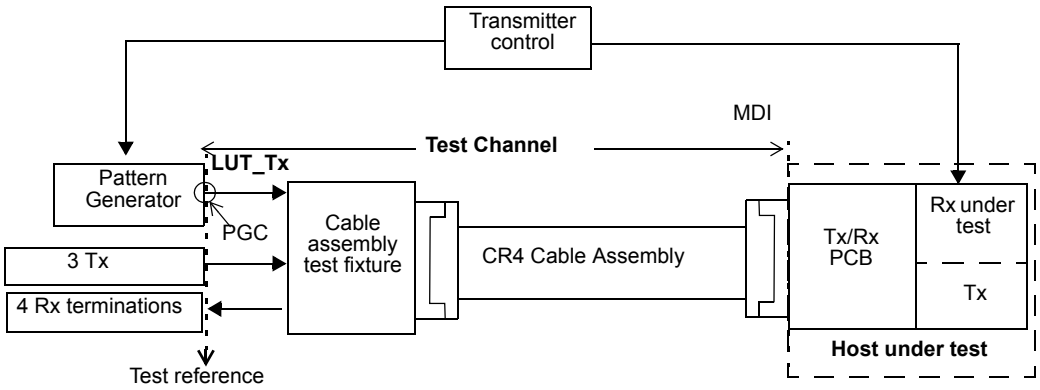

Figure 92–9—Interference tolerance test setup

92.8.4.4.2 Test channel

The test channel consists of the following:

a) A cable assembly that meets the cable assembly COM specified in 92.10.7.
b) A cable assembly test fixture
c) A connecting path from the pattern generator to the cable assembly test fixture

92.8.4.4.3 Test channel calibration

The scattering parameters of the test channels are characterized at the test references as illustrated in Figure 92–10 using the cable assembly test fixtures specified in 92.11.2.

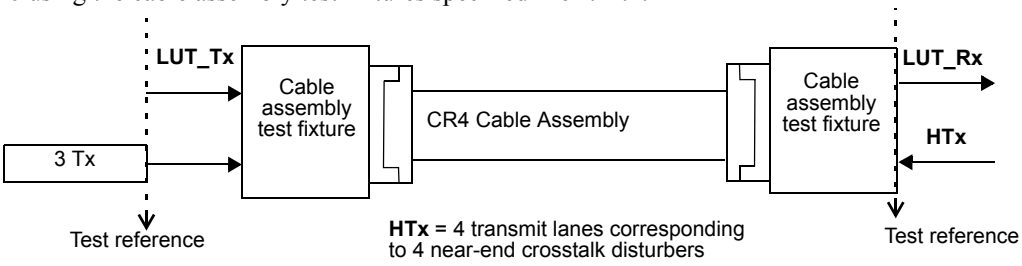

Figure 92–10—Test channel calibration

The fitted insertion loss coefficients of the lane under test (LUT), derived using the fitting procedure in 92.10.2, shall meet the test values in Table 92–8. It is recommended that the deviation between the insertion loss and the fitted insertion loss be as small as practical and that the fitting parameters be as close as practical to the values given in Table 92–8.

The far-end crosstalk disturbers consist of 100GBASE-CR4 transmitters. It is recommended that the transition time, equalization setting, and path from the far-end crosstalk disturbers to the cable assembly test fixture emulate the pattern generator as much as practical. For 100GBASE-CR4 test channels, the crosstalk that is coupled into a receive lane is from three transmitters. The disturber transmitters send scrambled idle encoded by RS-FEC. The amplitudes of each of the disturbers should be set to the value that results in the COM value given in Table 92–8 when calculated by the method given below.

The COM shall be calculated using the method and parameters of 92.10.7 with the following exceptions:

a) The channel signal path is $SCHS_p^{(k)} = \text{cascade}(\text{cascade}(S^{(CTSP)}, S^{(HOSP)}))$, where $S^{(CTSP)}$ is the measured channel between the test references for the LUT in Figure 92–10.

b) The channel far-end crosstalk path is $SCHNXT_p^{(k)} = \text{cascade}(\text{cascade}(S^{(CTFXTk)}, S^{(HOSP)}))$, where $S^{(CTFXTk)}$ is the measured FEXT channel between the test references [3 Tx] and LUT_Rx in Figure 92–10.

c) The value of the far-end aggressor amplitude A_{fe} is adjusted until the required COM is achieved. The far end aggressors ([3 Tx] in Figure 92–9) peak-to-peak amplitude is set to twice the resulting value for the test.

d) If the test transmitter presents a high-quality termination, e.g., it is a piece of test equipment, the transmitter device package model $S^{(tp)}$ is omitted from the calculation of $S_p^{(k)}$. Instead, the voltage transfer function is multiplied by the filter $H_t(f)$ defined by Equation (92–22) where T_r is the 20 to 80% transition time (see 86A.5.3.3) of the signal as measured at TP0a.

$$H_t(f) = \exp(-(\pi f T_r / 1.6832)^2) \tag{92–22}$$

92.8.4.4.4 Pattern generator

The pattern generator transmits data to the device under test. At the start of transmitter training, the pattern generator output amplitude shall be 800 mV peak-to-peak differential when measured on an alternating one-zero pattern. The output amplitude, measured on an alternating one zero pattern, is not permitted to exceed 800 mV peak-to-peak differential during transmitter training. The pattern generator shall be set to match the jitter specification in Table 92–8. The output waveform of the pattern generator shall comply to 93.8.1.

92.8.4.4.5 Test procedure

For 100GBASE-CR4 testing, the pattern generator is first configured to transmit the training pattern defined in 92.7.12. During this initialization period, the device under test (DUT) configures the pattern generator equalizer, via transmitter control, to the coefficient settings it would select using the protocol described in 72.6.10 and the receiver is tuned using its optimization method.

After the pattern generator equalizer has been configured and the receiver tuned, the pattern generator is set to generate scrambled idle encoded by RS-FEC. The receiver under test shall meet the target RS-FEC symbol error ratio listed in Table 92–8. During the tests, the disturbers transmit at their calibrated level and all of the transmitters in the device under test transmit scrambled idle encoded by RS-FEC, with the maximum compliant amplitude and equalization turned off (preset condition).

92.8.4.5 Receiver jitter tolerance

Receiver jitter tolerance is defined by the procedure in this subclause. When measured using the test setup shown in Figure 92–9, or its equivalent, the RS-FEC symbol error ratio for each lane of the receiver shall be less than or equal to 10^{-4} for each case listed in Table 92–9. The pattern generator meets the requirements of 92.8.4.4.4. The test channel meets the requirements of the interference tolerance test channel using Test 2 values listed in Table 92–8.

The test procedure is as described in 92.8.4.4.5 except that during the test the disturber transmitters are off and the pattern generator jitter is set to the frequency and peak-to-peak amplitude specified in Table 92–9.

Table 92–9—Receiver jitter tolerance parameters

Parameter	Case A values	Case B values	Units
Jitter frequency	190	940	kHz
Peak-to-peak jitter amplitude	5	1	UI

92.8.4.6 Signaling rate range

A 100GBASE-CR4 receiver shall comply with the requirements of 92.8.4.4 for any signaling rate in the range 25.78125 GBd ± 100 ppm. The corresponding unit interval is approximately 38.787879 ps.

92.9 Channel characteristics

The 100GBASE-CR4 channel is defined between TP0 and TP5 to include the transmitter and receiver differential controlled impedance printed circuit board and the cable assembly as illustrated in Figure 92–2. The channel parameters insertion loss, return loss, COM and the transmitter and receiver differential controlled impedance printed circuit board parameters for each differential lane are provided informatively in 92A.4 through 92A.7.

92.10 Cable assembly characteristics

The 100GBASE-CR4 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 5 m between network ports using controlled impedance cables. All cable assembly

measurements are to be made between TP1 and TP4 with cable assembly test fixtures as specified in 92.11.2 and illustrated in Figure 92–17. These cable assembly specifications are based upon twinaxial cable characteristics, but other cable types are acceptable if the specifications of 92.10 are met.

The cable assembly COM is specified in 92.10.7.

Table 92–10 provides a summary of the cable assembly characteristics and references to the subclauses addressing each parameter; reported values are at 12.8906 GHz.

Table 92–10—Cable assembly differential characteristics summary

Description	Reference	Value	Unit
Maximum insertion loss at 12.8906 GHz	92.10.2	22.48	dB
Minimum insertion loss at 12.8906 GHz	92.10.2	8	dB
Minimum return loss at 12.8906 GHz	92.10.3	6	dB
Differential to common-mode return loss	92.10.4	Equation (92–28)	dB
Differential to common-mode conversion loss	92.10.5	Equation (92–29)	dB
Common-mode to common-mode return loss	92.10.6	Equation (92–30)	dB

92.10.1 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.

92.10.2 Cable assembly insertion loss

The fitted cable assembly insertion loss $IL_{fitted}(f)$ as a function of frequency f is defined in Equation (92–23).

$$IL_{\text{fitted}}(f) = a_1\sqrt{f} + a_2 f + a_4 f^2 \quad \text{(dB)} \tag{92–23}$$

where

f is the frequency in GHz

$IL_{\text{fitted}}(f)$ is the fitted cable assembly insertion loss at frequency f

Given the cable assembly insertion loss measured between TP1 and TP4 is at N uniformly-spaced frequencies f_n spanning the frequency range 50 MHz to 19 000 MHz with a maximum frequency spacing of 10 MHz, the coefficients of the fitted insertion loss are determined using Equation (92–24) and Equation (92–25).

Define the frequency matrix F as shown in Equation (92–24).

$$F = \begin{bmatrix} \sqrt{f_1} & f_1 & f_1^2 \\ \sqrt{f_2} & f_2 & f_2^2 \\ \cdots & \cdots & \cdots \\ \sqrt{f_N} & f_N & f_N^2 \end{bmatrix} \tag{92–24}$$

The polynomial coefficients a_1, a_2, and a_4 are determined using Equation (92–25). In Equation (92–25), T denotes the matrix transpose operator and IL is a column vector of the measured insertion loss values, IL_n at each frequency f_n.

$$\begin{bmatrix} a_1 \\ a_2 \\ a_4 \end{bmatrix} = (F^T F)^{-1} F^T IL \tag{92–25}$$

The fitted insertion loss corresponding to one example of the maximum insertion loss at 12.8906 GHz is illustrated in Figure 92–11.

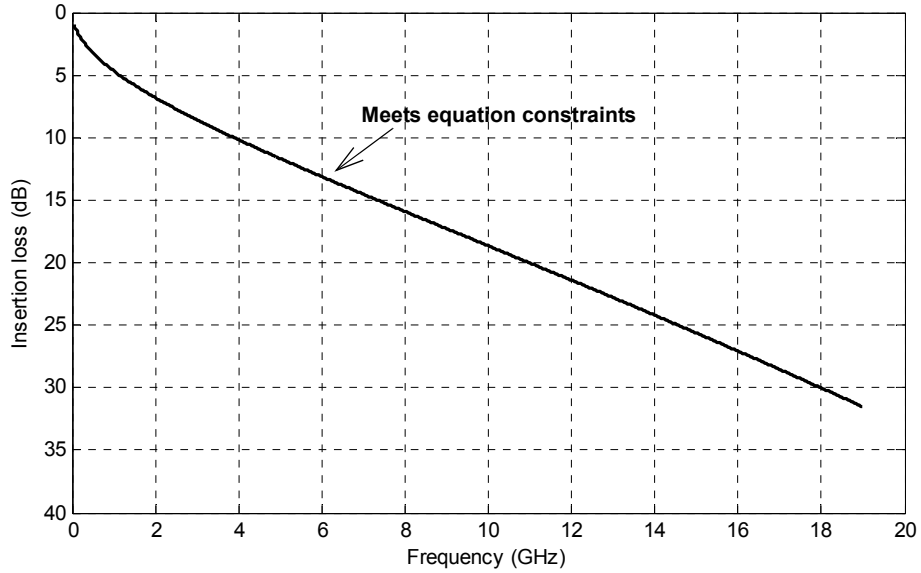

Figure 92–11—Example maximum cable assembly insertion loss

The measured insertion loss of the cable assembly shall be greater than or equal to the minimum cable assembly insertion loss given in Equation (92–26) and illustrated in Figure 92–12. The measured insertion loss of the cable assembly shall be less than or equal to the maximum cable insertion loss of 22.48 dB at 12.8906 GHz.

$$IL_{Cabmin}(f) = 0.7\sqrt{f} + 0.3f + 0.01f^2 \quad \text{(dB)} \tag{92–26}$$

where

f is the frequency in GHz

$IL_{\text{Cabmin}}(f)$ is the minimum cable assembly insertion loss at frequency f

Table 92–11—Maximum and minimum cable assembly insertion loss characteristics

Description	Value	Unit
Maximum insertion loss at 12.8906 GHz	22.48	dB
Minimum insertion loss at 12.8906 GHz	8	dB

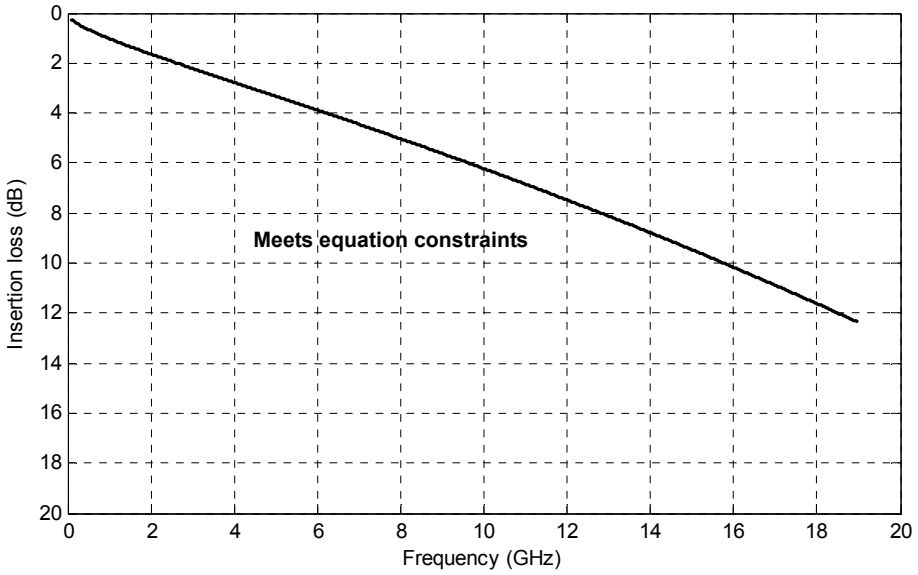

Figure 92–12—Minimum cable assembly insertion loss

92.10.3 Cable assembly differential return loss

The differential return loss of each pair of the 100GBASE-CR4 cable assembly shall meet the values determined using Equation (92–27).

$$Return\_loss(f) \geq \begin{cases} 16.5 - 2\sqrt{f} & 0.05 \leq f < 4.1 \\ 10.66 - 14\log_{10}(f/5.5) & 4.1 \leq f \leq 19 \end{cases} \quad \text{(dB)} \qquad (92\text{–}27)$$

where

f is the frequency in GHz

$Return\_loss(f)$ is the return loss at frequency f

The minimum cable assembly return loss is illustrated in Figure 92–13.

92.10.4 Differential to common-mode return loss

The differential to common-mode return loss, in dB, of the cable assembly shall meet Equation (92–28).

$$Return\_loss(f) \geq \begin{cases} 22 - (20/25.78)f & 0.01 \leq f < 12.89 \\ 15 - (6/25.78)f & 12.89 \leq f \leq 19 \end{cases} \quad \text{(dB)} \qquad (92\text{–}28)$$

where

f is the frequency in GHz

$Return\_loss(f)$ is the return loss at frequency f

The differential to common-mode cable assembly return loss is illustrated in Figure 92–14.

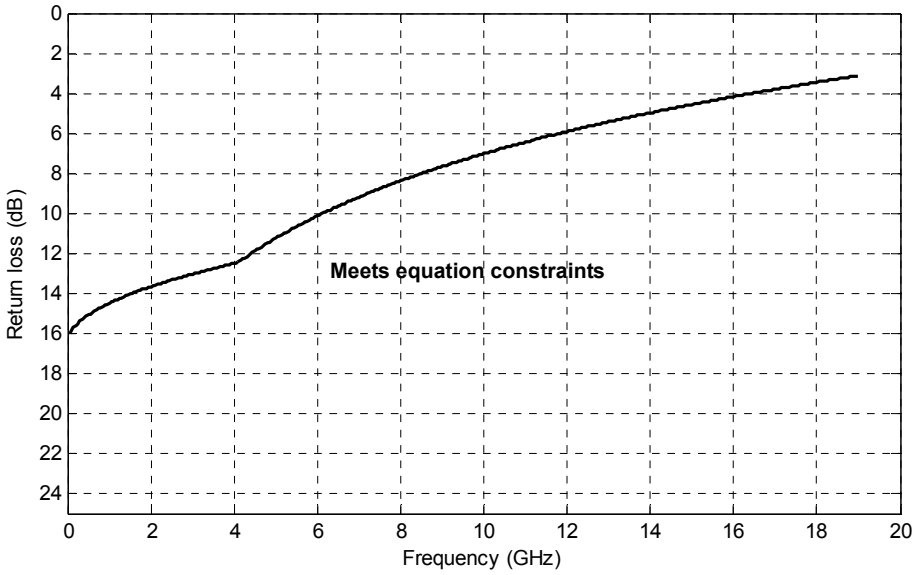

Figure 92–13—Minimum cable assembly return loss

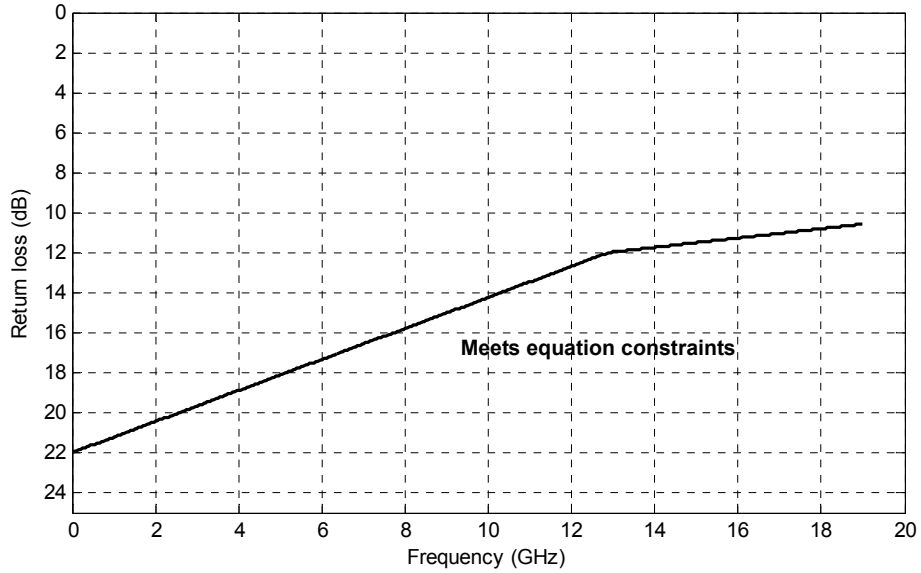

Figure 92–14—Differential to common-mode cable assembly return loss

92.10.5 Differential to common-mode conversion loss

The difference between the cable assembly differential to common-mode conversion loss and the cable assembly insertion loss shall meet Equation (92–29)

$$Conversion\_loss\ (f) - IL(f) \geq \left\{ \begin{array}{ll} 10 & 0.01 \leq f < 12.89 \\ 27 - (29/22)f & 12.89 \leq f < 15.7 \\ 6.3 & 15.7 \leq f \leq 19 \end{array} \right\} \quad (92–29)$$

where

f	is the frequency in GHz
$Conversion\_loss(f)$	is the cable assembly differential to common-mode conversion loss
$IL(f)$	is the cable assembly insertion loss

92.10.6 Common-mode to common-mode return loss

The common-mode to common-mode return loss, in dB, of the cable assembly shall meet Equation (92–30).

$$Return\_loss(f) \geq 2 \quad (\text{dB}) \quad (92–30)$$

for $0.2 \leq f \leq 19$ GHz

where

f	is the frequency in GHz
$Return\_loss(f)$	is the common-mode to common-mode return loss at frequency f

92.10.7 Cable assembly Channel Operating Margin

The cable assembly Channel Operating Margin (COM) for each victim signal path (receive lane) is derived from measurements of the cable assembly victim signal path, the four individual near-end crosstalk paths, and the three far-end crosstalk paths that can couple into a victim signal path. COM is computed using the procedure in 93A.1 with the Test 1 and Test 2 values in Table 93–8 and the signal paths defined in 92.10.7.1 and 92.10.7.2. Test 1 and Test 2 differ in the value of the device package model transmission line length z_p.

NOTE—For cable lengths greater than 4 m a frequency step (Δf) no larger than 5 MHz is recommended.

The cable assembly COM shall be greater than or equal to 3 dB for each test. This minimum value allocates margin for practical limitations on the receiver implementation as well as the largest step size allowed for transmitter equalizer coefficients.

92.10.7.1 Channel signal path

The channel signal path between TP0 and TP5 for the cable assembly COM consists of the measured cable assembly signal path (TP1 and TP4), a representative transmitter PCB signal path (TP0 to TP1), and a representative receiver PCB signal path (TP4 to TP5).

The channel signal path to be used in COM (93A.1.2) is the concatenation of the cable assembly signal path measurement, the transmitter PCB signal path, and the receiver PCB signal path using Equation (92–31) (see 93A.1.2.1). The transmitter and receiver PCB signal paths are calculated according to 92.10.7.1.1.

$$SCHS_p^{(k)} = \text{cascade}(\text{cascade}(S^{(HOSP)}, S^{(CASP)}), S^{(HOSP)}) \quad (92–31)$$

where

$SCHS_p^{(k)}$ is the channel signal path

$S^{(HOSP)}$ is the signal path calculated according to 92.10.7.1.1

$S^{(CASP)}$ is the cable assembly signal path

k is equal to zero

92.10.7.1.1 TP0 to TP1 and TP4 to TP5 signal paths

The additional transmitter and receiver PCB signal paths are calculated using the method defined in 93A.1.2.3. The scattering parameters for a PCB of length z_p are defined by Equation (93A–13), Equation (93A–14), and the parameter values given in Table 92–12.

For the channel signal path defined in 92.10.7.1 and calculated using Equation (92–31), the transmitter and receiver PCB model are each z_p = 151 mm in length representing an insertion loss of 6.26 dB at 12.89 GHz and are each denoted as $S^{(HOSP)}$.

For the channel crosstalk paths defined in 92.10.7.2 and calculated using Equation (92–32) and Equation (92–33), the receiver PCB model is $S^{(HOSP)}$. The transmitter PCB model is z_p = 72 mm in length representing an insertion loss of 3 dB at 12.89 GHz and is denoted as $S^{(HOTxSP)}$. The transmitter PCB insertion loss for the crosstalk channel is less than that for the signal channel to allow for a reasonable worst-case crosstalk in the COM calculation.

Table 92–12—Transmission line model parameters

Parameter	Value	Units
γ_0	0	1/mm
a_1	4.114×10^{-4}	$ns^{1/2}$/mm
a_2	2.547×10^{-4}	ns/mm
τ	6.191×10^{-3}	ns/mm
Z_c	109.8	Ω

92.10.7.2 Channel crosstalk paths

The channel structure includes three far-end and four near-end crosstalk paths. The MDI is the significant contributor to the channel crosstalk and is included in and characterized by the cable assembly crosstalk measurements. The cable assembly crosstalk signal paths to be used in COM are the four individual near-end crosstalk paths and the three far-end crosstalk paths that can couple into a victim signal path using the cascaded function defined in 93A.1.2.1 given in Equation (92–32) and Equation (92–33).

$$SCHNXT_p^{(k)} = \text{cascade}(\text{cascade}(S^{(HOTxSP)}, S^{(CANXTk)}), S^{(HOSP)}) \quad (92–32)$$

where

$SCHNXT_p^{(k)}$	is the channel near-end crosstalk path
$S^{(HOSP)}$	is the signal path calculated according to 92.10.7.1.1
$S^{(HOTxSP)}$	is the signal path calculated according to 92.10.7.1.1
$S^{(CANXTk)}$	is the cable assembly near-end crosstalk path k
k	is 1 to 4 near-end crosstalk paths

$$SCHFXT_p^{(k)} = \text{cascade}(\text{cascade}(S^{(HOTxSP)}, S^{(CASFXTk)}), S^{(HOSP)}) \quad (92–33)$$

where

$SCHFXT_p^{(k)}$	is the channel far-end crosstalk path
$S^{(HOSP)}$	is the signal path calculated according to 92.10.7.1.1
$S^{(HOTxSP)}$	is the signal path calculated according to 92.10.7.1.1
$S^{(CAFXTk)}$	is the cable assembly far-end crosstalk path k
k	is 1 to 3 far-end crosstalk paths

92.11 Test fixtures

Transmitter and receiver measurements are made at TP2 or TP3 utilizing the test fixture specified in 92.11.1. and illustrated in Figure 92–15. All cable assembly measurements are to be made between TP1 and TP4 with cable assembly test fixtures as specified in 92.11.2 and illustrated in Figure 92–17. The test fixtures of Figure 92–15 and Figure 92–17 are specified in a mated state, illustrated in Figure 92–18, to enable connections to measurement equipment. The requirements in this subclause are not MDI connector specifications for an implemented design.

92.11.1 TP2 or TP3 test fixture

The test fixture (also known as Host Compliance Board) of Figure 92–15, or its equivalent, is required for measuring the transmitter specifications in 92.8.3 at TP2 and the receiver return loss at TP3. The TP2 and TP3 test points are illustrated in Figure 92–2. Figure 92–15 illustrates the test fixture attached to TP2 or TP3.

92.11.1.1 Test fixture return loss

The differential return loss, in dB, of the test fixture is specified in a mated state and shall meet the requirements of 92.11.3.2.

92.11.1.2 Test fixture insertion loss

The test fixture printed circuit board insertion loss values determined using Equation (92–34) shall be used as the reference test fixture insertion loss. The effects of differences between the insertion loss of an actual test fixture and the reference insertion loss are to be accounted for in the measurements.

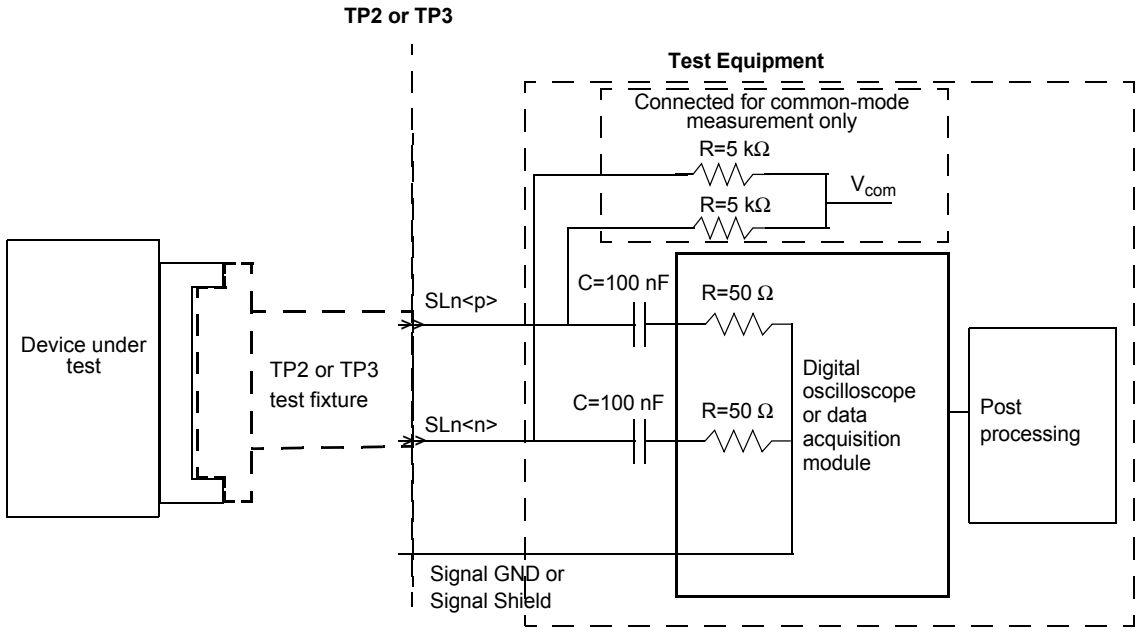

Figure 92–15—Transmitter and receiver test setup

$$IL_{tfref}(f) = -0.00144 + 0.13824\sqrt{f} + 0.06624\,f \quad \text{(dB)} \tag{92–34}$$

for $0.01 \le f \le 25$ GHz

where

f is the frequency in GHz

$IL_{tfref}(f)$ is the reference test fixture PCB insertion loss at frequency f

The reference test fixtures PCB insertion loss is illustrated in Figure 92–16.

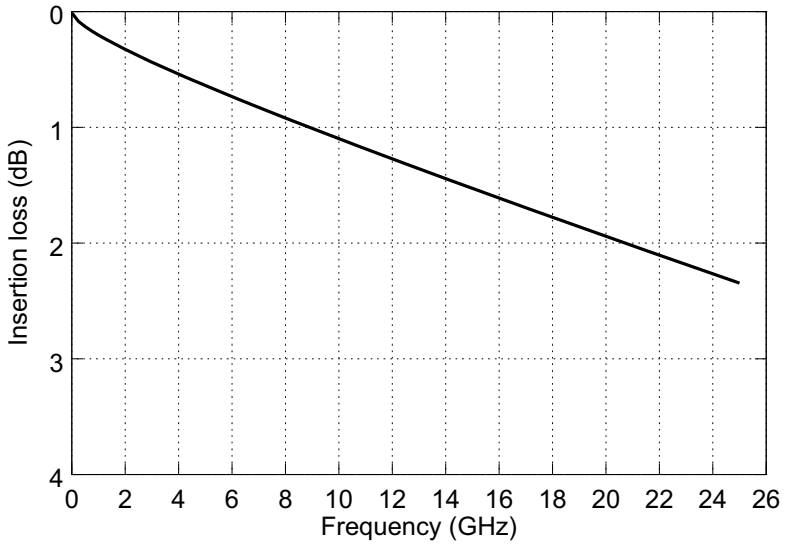

Figure 92–16—Reference test fixture insertion loss

92.11.2 Cable assembly test fixture

The test fixture of Figure 92–17 (also known as Module Compliance Board) or its equivalent, is required for measuring the cable assembly specifications in 92.10 at TP1 and TP4. The TP1 and TP4 test points are illustrated in Figure 92–2 and Figure 92–17. The test fixture return loss is equivalent to the test fixture return loss specified in 92.11.3.2. The test fixture printed circuit board insertion loss values determined using Equation (92–35) shall be used as the reference test fixture insertion loss. The effects of differences between the insertion loss of an actual test fixture and the reference insertion loss are to be accounted for in the measurements.

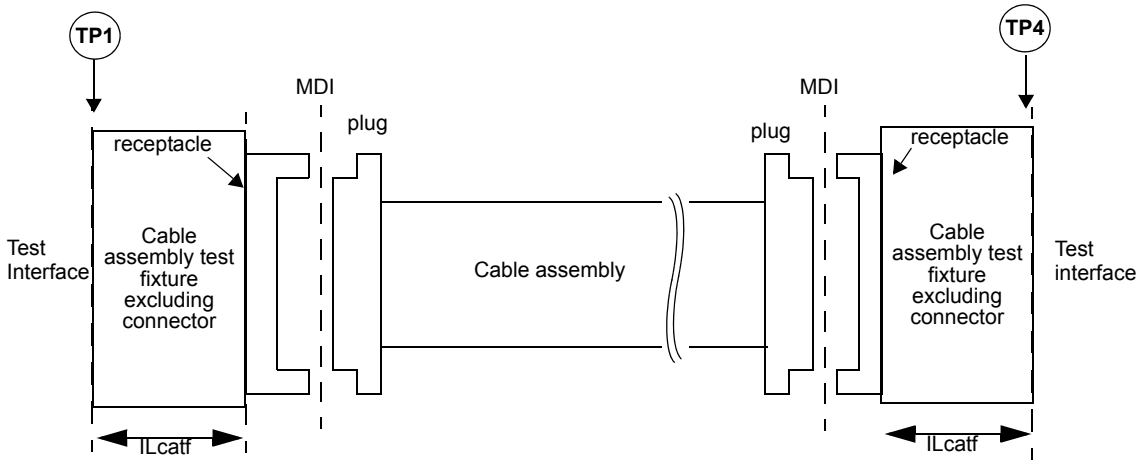

Figure 92–17—Cable assembly test fixtures

$$IL_{\mathrm{catf}}(f) = -0.00125 + 0.12\sqrt{f} + 0.0575f \quad \text{(dB)} \tag{92–35}$$

for $0.01\ \text{GHz} \le f \le 25\ \text{GHz}$

where

f is the frequency in GHz

$IL_{\mathrm{catf}}(f)$ is the reference test fixture printed circuit board insertion loss at frequency f

92.11.3 Mated test fixtures

The test fixtures of Figure 92–15 and Figure 92–17 are specified in a mated state illustrated in Figure 92–18. The mated test fixtures specifications shall be verified in both directions indicated by the arrows illustrated in Figure 92–18 except insertion loss, which shall be verified at either test interface illustrated in Figure 92–18.

92.11.3.1 Mated test fixtures insertion loss

The insertion loss of the mated test fixtures shall meet the values determined using Equation (92–36) and Equation (92–37).

$\mathrm{FOM_{ILD}}$ is calculated according to 93A.4 with f_b=25.78125 GHz, T_t=9.6 ps, and f_r=0.75 × f_b. The fitted insertion loss and insertion loss deviation are computed over the range f_{min}=0.01 GHz to f_{max}=25 GHz. $\mathrm{FOM_{ILD}}$ shall be less than 0.13 dB.

$$IL(f) \le IL_{\mathrm{MTFmax}}(f) = \left\{ \begin{array}{ll} 0.12 + 0.475\sqrt{f} + 0.221f & 0.01 \le f \le 14 \\ -4.25 + 0.66f & 14 < f \le 25 \end{array} \right\} \quad \text{(dB)} \tag{92–36}$$

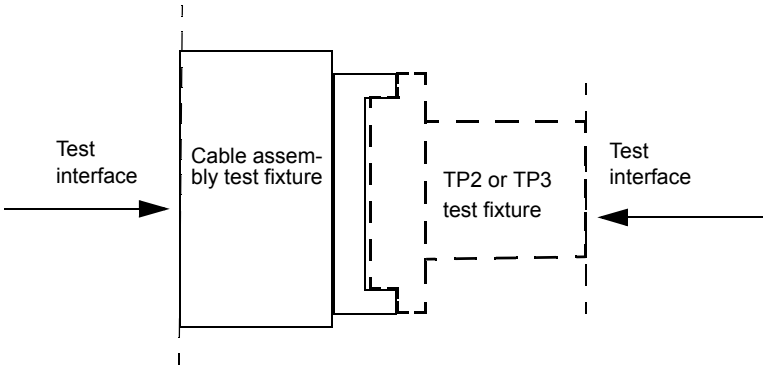

Figure 92–18—Mated test fixtures

$$IL(f) \geq IL_{\text{MTFmin}}(f) = 0.0656\sqrt{f} + 0.164f \quad 0.01 \leq f \leq 25 \quad \text{(dB)} \qquad (92\text{–}37)$$

where

f is the frequency in GHz

$IL(f)$ is the mated test fixture insertion loss at frequency f

The mated test fixtures insertion loss limits are illustrated in Figure 92–19.

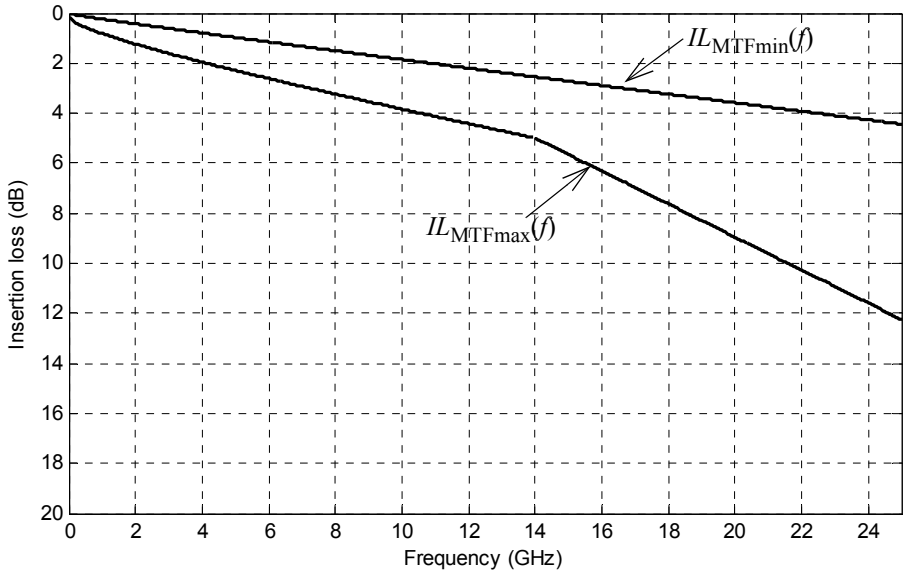

Figure 92–19—Mated test fixtures Insertion loss

92.11.3.2 Mated test fixtures return loss

The return loss of the mated test fixtures measured at each test fixture interface shall meet the values determined using Equation (92–38).

$$Return\_loss(f) \geq \left\{ \begin{array}{ll} 20 - f & 0.01 \leq f < 4 \\ 18 - 0.5f & 4 \leq f \leq 25 \end{array} \right\} \text{ (dB)} \quad (92\text{–}38)$$

where

f is the frequency in GHz

$Return\_loss(f)$ is the return loss at frequency f

The mated test fixtures return loss is illustrated in Figure 92–20.

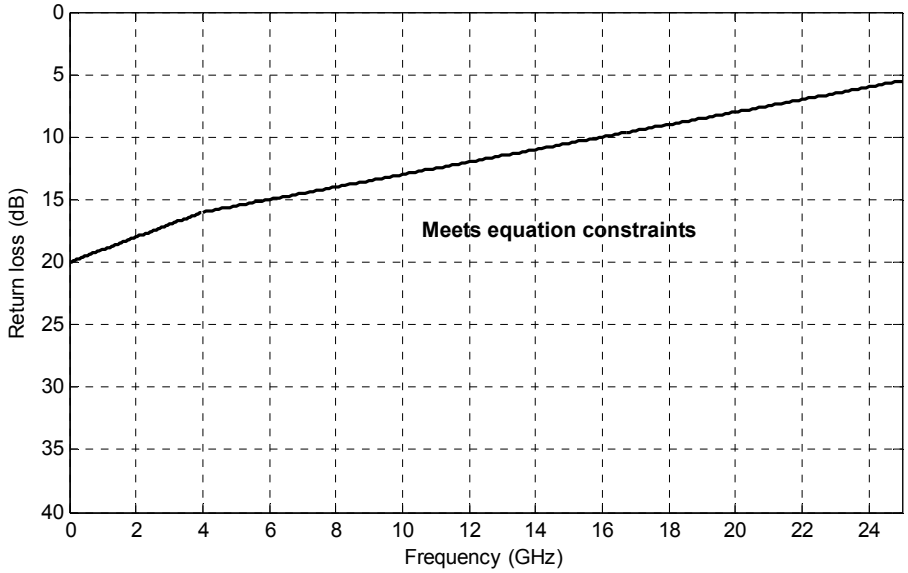

Figure 92–20—Mated test fixtures return loss

92.11.3.3 Mated test fixtures common-mode conversion insertion loss

The common-mode conversion insertion loss of the mated test fixtures measured at either test fixture test interface shall meet the values determined using Equation (92–39).

$$Conversion\_loss(f) \geq \left\{ \begin{array}{ll} 30 - (29/22)f & 0.01 \leq f < 16.5 \\ 8.25 & 16.5 \leq f \leq 25 \end{array} \right\} \text{ (dB)} \quad (92\text{–}39)$$

where

f is the frequency in GHz

$Conversion\_loss(f)$ is the conversion insertion loss at frequency f

The mated test fixtures common-mode conversion insertion loss is illustrated in Figure 92–21.

92.11.3.4 Mated test fixtures common-mode return loss

The common-mode return loss of the mated test fixtures measured at each test fixture test interface shall meet the values determined using Equation (92–40).

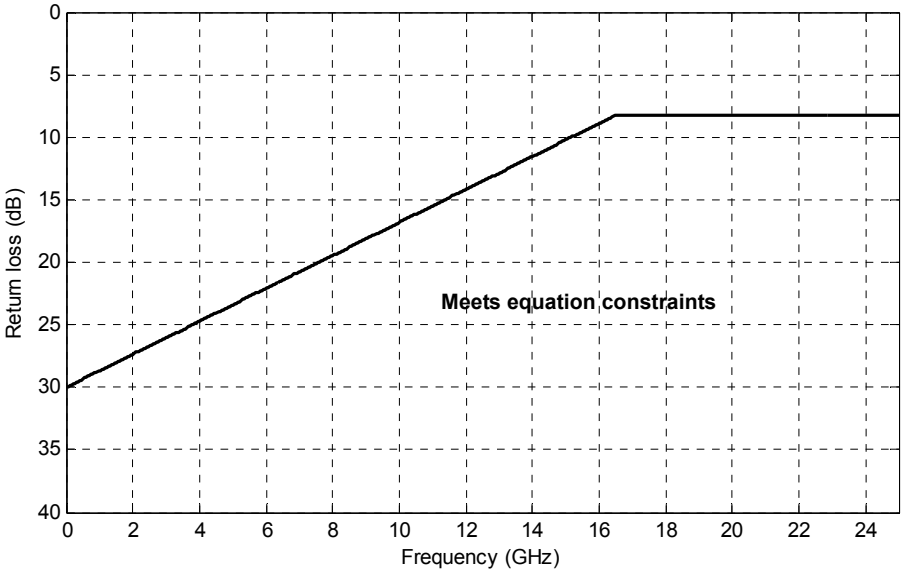

Figure 92–21—Common-mode conversion loss

$$Common\_mode\_return\_loss(f) \geq \begin{cases} 12 - 9f & 0.01 \leq f < 1 \\ 3 & 1 \leq f \leq 25 \end{cases} \quad \text{(dB)} \qquad (92\text{–}40)$$

where

f	is the frequency in GHz
$Common\_mode\_return\_loss(f)$	is the common-mode return loss at frequency f

The mated test fixtures common-mode return loss is illustrated in Figure 92–22.

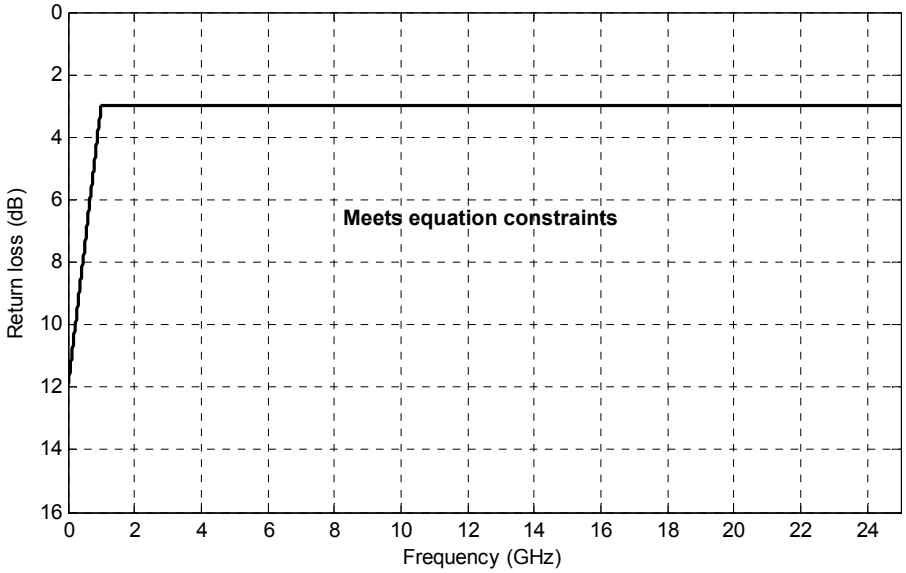

Figure 92–22—Common-mode return loss

92.11.3.5 Mated test fixtures common-mode to differential mode return loss

The common-mode to differential mode return loss of the mated test fixtures measured at each test fixture test interface shall meet the values determined using Equation (92–41).

$$Return\_loss(f) \geq \left\{ \begin{array}{ll} 30 - (30/25.78)f & 0.01 \leq f < 12.89 \\ 18 - (6/25.78)f & 12.89 \leq f \leq 25 \end{array} \right\} \text{ (dB)} \quad (92–41)$$

where

f is the frequency in GHz

$Return\_loss(f)$ is the common-mode to differential mode return loss at frequency f

The mated test fixtures common-mode to differential mode return loss is illustrated in Figure 92–23.

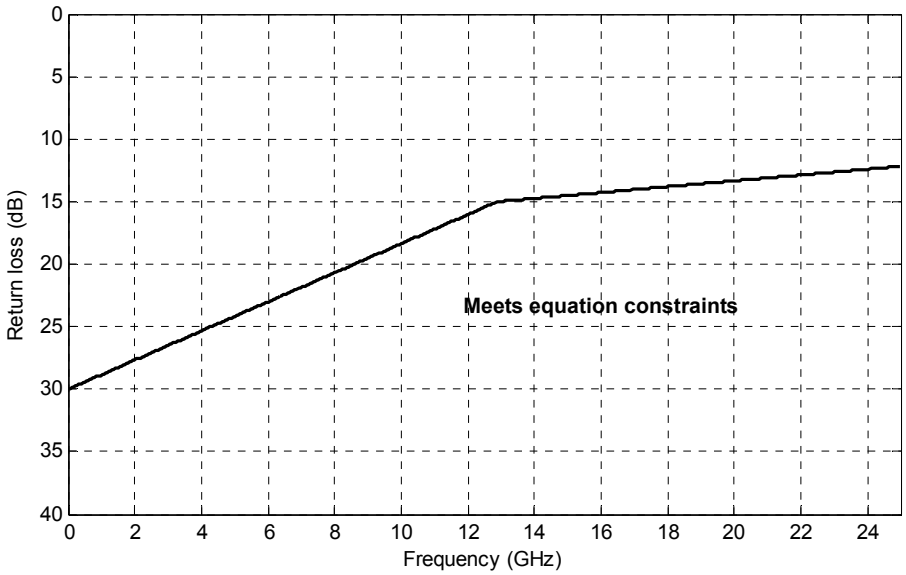

Figure 92–23—Common-mode to differential return loss

92.11.3.6 Mated test fixtures integrated crosstalk noise

The values of the mated test fixtures integrated crosstalk RMS noise voltages determined using Equation (92–44) through Equation (92–48) for the multiple disturber near-end crosstalk loss, and the multiple disturber far-end crosstalk loss shall meet the specifications in Table 92–13.

Table 92–13—Mated test fixtures integrated crosstalk noise

Parameter	100GBASE-CR4	Units
MDNEXT integrated crosstalk noise voltage	Less than 1.8	mV
MDFEXT integrated crosstalk noise voltage	Less than 4.8	mV

92.11.3.6.1 Mated test fixture multiple disturber near-end crosstalk (MDNEXT) loss

Since four lanes are used to transfer data between PMDs, the NEXT that is coupled into a receive lane is from the four transmit lanes. Multiple Disturber Near-End Crosstalk (MDNEXT) loss is determined using the individual NEXT losses.

MDNEXT loss is determined from the four individual pair-to-pair differential NEXT loss values using Equation (92–42).

$$MDNEXT\_loss(f) = -10\log_{10}\left(\sum_{i=0}^{i=3} 10^{-NLi(f)/10}\right) \quad \text{(dB)} \tag{92–42}$$

for 0.05 GHz ≤ f ≤ 19 GHz

where

$MDNEXT\_loss(f)$	is the MDNEXT loss at frequency f
$NL_i(f)$	is the NEXT loss at frequency f of pair combination i, in dB
f	is the frequency in GHz
i	is the 0 to 3 (pair-to-pair combination)

92.11.3.6.2 Mated test fixture multiple disturber far-end crosstalk (MDFEXT) loss

Since four lanes are used to transfer data between PMDs, the FEXT that is coupled into a data carrying lane is from the three other lanes in the same direction. MDFEXT loss is specified using the individual FEXT losses. MDFEXT loss is determined from the three individual pair-to-pair differential FEXT loss values using Equation (92–43).

$$MDFEXT\_loss(f) = -10\log_{10}\left(\sum_{i=0}^{i=2} 10^{-NLi(f)/10}\right) \quad \text{(dB)} \tag{92–43}$$

for 0.05 GHz ≤ f ≤ 19 GHz

where

$MDFEXT\_loss(f)$	is the MDFEXT loss at frequency f
$NL_i(f)$	is the FEXT loss at frequency f of pair combination i, in dB
f	is the frequency in GHz
i	is the 0 to 2 (pair-to-pair combination)

92.11.3.6.3 Mated test fixture integrated crosstalk noise (ICN)

ICN is calculated from the MDFEXT and MDNEXT. Given the multiple disturber near-end crosstalk loss $MDNEXT\_loss(f)$ and multiple disturber far-end crosstalk loss $MDFEXT\_loss(f)$ measured over N uniformly-spaced frequencies f_n spanning the frequency range 50 MHz to 19 000 MHz with a maximum frequency spacing of 10 MHz, the RMS value of the integrated crosstalk noise is determined using Equation (92–44) through Equation (92–48). The RMS crosstalk noise is characterized at the output of a specified receive filter utilizing a specified transmitter waveform and the measured multiple disturber crosstalk transfer functions. The transmitter and receiver filters are defined in Equation (92–44) and Equation (92–45) as weighting functions to the multiple disturber crosstalk in Equation (92–46) and Equation (92–47). The sinc function is defined by $\text{sinc}(x) = \sin(\pi x)/(\pi x)$.

Define the weight at each frequency f_n using Equation (92–44) and Equation (92–45).

$$W_{nt}(f_n) = (A_{nt}^2/f_b)\mathrm{sinc}^2(f_n/f_b)\left[\frac{1}{1+(f_n/f_{nt})^4}\right]\left[\frac{1}{1+(f_n/f_r)^8}\right] \quad (92\text{–}44)$$

$$W_{ft}(f_n) = (A_{ft}^2/f_b)\mathrm{sinc}^2(f_n/f_b)\left[\frac{1}{1+(f_n/f_{ft})^4}\right]\left[\frac{1}{1+(f_n/f_r)^8}\right] \quad (92\text{–}45)$$

where the equation parameters are given in Table 92–14.

Note that the 3 dB transmit filter bandwidths f_{nt} and f_{ft} are inversely proportional to the 20% to 80% rise and fall times T_{nt} and T_{ft} respectively. The constant of proportionality is 0.2365 (e.g., $T_{nt}f_{nt} = 0.2365$; with f_{nt} in hertz and T_{nt} in seconds). In addition, f_r is the 3 dB reference receiver bandwidth, which is set to 18.75 GHz.

The near-end integrated crosstalk noise σ_{nx} is calculated using Equation (92–46).

$$\sigma_{nx} = \left[2\Delta f\sum_n W_{nt}(f_n)10^{-MDNEXT_{loss}(f_n)/10}\right]^{1/2} \quad (92\text{–}46)$$

The far-end integrated crosstalk noise σ_{fx} is calculated using Equation (92–47).

$$\sigma_{fx} = \left[2\Delta f\sum_n W_{ft}(f_n)10^{-MDFEXT_{loss}(f_n)/10}\right]^{1/2} \quad (92\text{–}47)$$

where Δf is the uniform frequency step of f_n.

The total integrated crosstalk noise σ_x is calculated using Equation (92–48).

$$\sigma_x = \sqrt{\sigma_{nx}^2 + \sigma_{fx}^2} \quad (92\text{–}48)$$

The total integrated crosstalk noise for the mated test fixture is computed using the parameters shown in Table 92–14.

Table 92–14—Mated test fixture integrated crosstalk noise parameters

Description	Symbol	Value	Units
Symbol rate	f_b	25.78125	GBd
Near-end disturber peak differential output amplitude	A_{nt}	600	mV
Far-end disturber peak differential output amplitude	A_{ft}	600	mV
Near-end disturber 20% to 80% rise and fall times	T_{nt}	9.6	ps
Far-end disturber 20% to 80% rise and fall times	T_{ft}	9.6	ps

92.12 MDI specification

This subclause defines the Media Dependent Interface (MDI). The 100GBASE-CR4 PMD, as per 92.7, is coupled to the cable assembly, as per 92.10, by the MDI.

92.12.1 100GBASE-CR4 MDI connectors

Connectors meeting the requirements of 92.12.1.1 (Style-1) or 92.12.1.2 (Style-2) are used as the mechanical interface between the PMD of 92.7 and the cable assembly of 92.10. The plug connector is used on the cable assembly and the receptacle on the PHY. Style-1 or Style-2 connectors may be used as the MDI.

For Style-1 and Style-2 100GBASE-CR4 plug connectors, the receive lanes are AC-coupled; the AC-coupling shall be within the plug connectors. It should be noted that there may be various methods for AC-coupling in actual implementations. The low-frequency 3 dB cutoff of the AC-coupling shall be less than 50 kHz. It is recommended that the value of the coupling capacitors be 100 nF. The capacitor limits the inrush charge and baseline wander.

92.12.1.1 Style-1 100GBASE-CR4 MDI connectors

The plug connector for each end of the cable assembly shall be the QSFP+ 28 Gb/s 4X Pluggable (QSFP28) plugs defined in SFF-8665 and illustrated in Figure 92–24. The MDI connector shall be the QSFP+ 28 Gb/s 4X Pluggable (QSFP28) receptacle with the mechanical mating interface defined in SFF-8665 and illustrated in Figure 92–25. These connectors have contact assignments that are listed in Table 92–15 and electrical performance consistent with the signal quality and electrical requirements of 92.8 and 92.9.

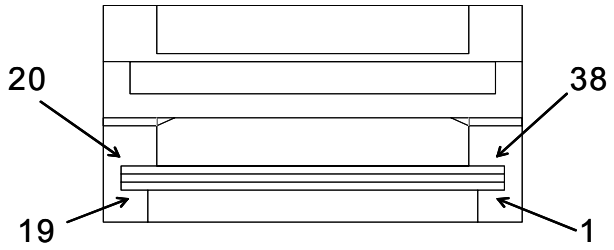

Figure 92–24—Style-1 example cable assembly

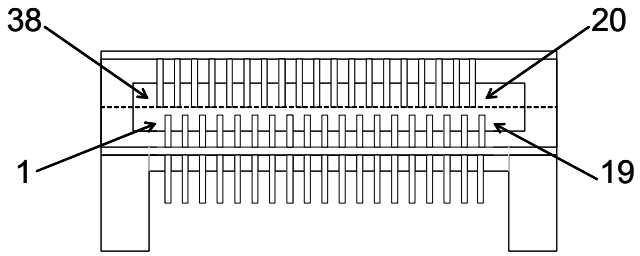

Figure 92–25—Style-1 example MDI board receptacle

The Style-1 MDI connector of the 100GBASE-CR4 PMD comprises 38 signal connections. The Style-1 100GBASE-CR4 MDI connector contact assignments shall be as defined in Table 92–15.

Table 92–15—100GBASE-CR4 lane to MDI connector contact mapping

Tx lane	MDI connector contact	Rx lane	MDI connector contact
signal gnd	S1	signal gnd	S13
SL1<n>	S2	DL2<p>	S14
SL1<p>	S3	DL2<n>	S15
signal gnd	S4	signal gnd	S16
SL3<n>	S5	DL0<p>	S17
SL3<p>	S6	DL0<n>	S18
signal gnd	S7	signal gnd	S19
signal gnd	S32	signal gnd	S20
SL2<p>	S33	DL1<n>	S21
SL2<n>	S34	DL1<p>	S22
signal gnd	S35	signal gnd	S23
SL0<p>	S36	DL3<n>	S24
SL0<n>	S37	DL3<p>	S25
signal gnd	S38	signal gnd	S26

92.12.1.2 Style-2 100GBASE-CR4 MDI connectors

The connector for each end of the cable assembly shall be the 100G Form Factor Pluggable 4 (CFP4) with the mechanical mating interface defined in CFP4 MSA HW Specification and illustrated in Figure 92–26. The MDI connector shall be the 100G Form Factor Pluggable (CFP4) receptacle with the mechanical mating interface defined by CFP4 MSA HW Specification and illustrated in Figure 92–27. These connectors have contact assignments that are listed in Table 92–16, and electrical performance consistent with the signal quality and electrical requirements of 92.8 and 92.9.

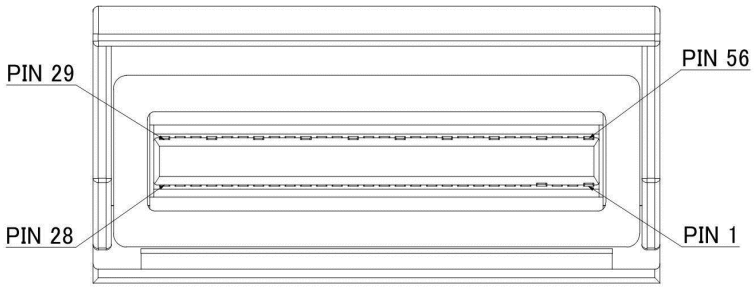

Figure 92–26—Style-2 example cable assembly

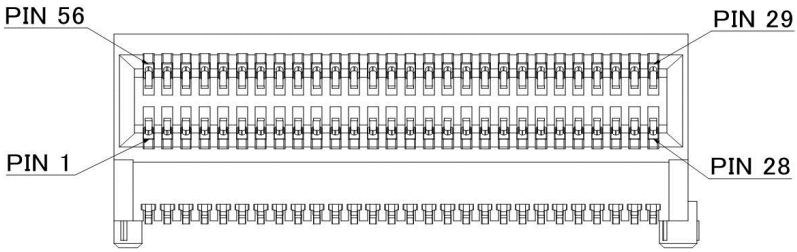

Figure 92–27—Style-2 example MDI board receptacle

The Style-2 MDI connector of the 100GBASE-CR4 PMD comprises 56 signal connections. The Style-2 100GBASE-CR4 MDI connector contact assignments shall be as defined in Table 92–16. Note that the source lanes (SL), signals SLi<p>, and SLi<n> are the positive and negative sides of the transmitters differential signal pairs and the destination lanes (DL) signals, DLi<p>, and DLi<n> are the positive and negative sides of the receivers differential signal pairs for lane i (i = 0, 1, 2, 3).

Table 92–16—100GBASE-CR4 lane to MDI connector contact mapping

Tx lanes	MDI connector contact	Rx lanes	MDI connector contact
signal gnd	44	signal gnd	29
SL0<p>	45	DL0<p>	30
SL0<n>	46	DL0<n>	31
signal gnd	47	signal gnd	32
SL1<p>	48	DL1<p>	33
SL1<n>	49	DL1<n>	34
signal gnd	50	signal gnd	35
SL2<p>	51	DL2<p>	36
SL2<n>	52	DL2<n>	37
signal gnd	53	signal gnd	38
SL3<p>	54	DL3<p>	39
SL3<n>	55	DL3<n>	40
signal gnd	56	signal gnd	41

92.13 Environmental specifications

All equipment subject to this clause shall conform to the applicable requirements of 14.7.

92.14 Protocol implementation conformance statement (PICS) proforma for Clause 92, Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-CR4[21]

92.14.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 92, Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-CR4, 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.

92.14.2 Identification

92.14.2.1 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).	

92.14.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 92, Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-CR4
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	

[21] *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.

92.14.3 Major capabilities/options

Item[a]	Feature	Subclause	Value/Comment	Status	Support
CGMII	CGMII	92.1	Interface is supported	O	Yes [ ] No [ ]
PCS	100GBASE-R PCS	92.1		M	Yes [ ]
RS-FEC	100GBASE-R RS-FEC	92.1		M	Yes [ ]
PMA	100GBASE-R PMA	92.1		M	Yes [ ]
CAUI-10	CAUI-10	92.1	Interface is supported	O	Yes [ ] No [ ]
CAUI-4	CAUI-4	92.1	Interface is supported	O	Yes [ ] No [ ]
CR4	100GBASE-CR4 PMD	92.1	Can operate as 100GBASE-CR4 PMD	M	Yes [ ]
AN	Auto-negotiation	92.1	Device implements Auto-Negotiation	M	Yes [ ]
DC	Delay constraints	92.4	Device conforms to delay constraints specified in 92.4	M	Yes [ ]
DSC	Skew constraints	92.5	Device conforms to Skew and Skew Variation constraints specified in 92.5	M	Yes [ ]
*MD	MDIO capability	92.6	Registers and interface supported	O	Yes [ ] No [ ]
*EEE	EEE deep sleep capability	92.1	Capability is supported	O	Yes [ ] No [ ]
*GTD	Global PMD transmit disable function	92.7.6	Function is supported	EEE:M	Yes [ ] No [ ]
*LTD	PMD lane-by-lane transmit disable function	92.7.7	Function is supported	O	Yes [ ] No [ ]
*CBL	Cable assembly	92.10	Items marked with CBL include cable assembly specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]
CAST1	100GBASE-CR4 Style-1 cable assembly	92.10	Cable assembly supports 100GBASE-CR4 Style-1	CBL:O.1	Yes [ ] No [ ]
CAST2	100GBASE-CR4 Style-2 cable assembly	92.10	Cable assembly supports 100GBASE-CR4 Style-2	CBL:O.1	Yes [ ] No [ ]
MDIST1	Style-1 MDI connector	92.12.1.1	100GBASE-CR4 device uses Style-1 MDI	O:2	Yes [ ] N/A [ ]
MDIST2	Style-2 MDI connector	92.12.1.2	100GBASE-CR4 device uses Style-2 MDI	O:2	Yes [ ] N/A [ ]

[a]A "*" preceding an "Item" identifier indicates there are other PICS that depend on whether or not this item is supported.

92.14.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-CR4

92.14.4.1 PMD functional specifications

Item	Feature	Sub-clause	Value/Comment	Status	Support
PF1	Transmit function	92.7.2	Converts four logical bit streams from the PMD service interface into four electrical signals and delivers them to the MDI	M	Yes [ ]
PF2	Transmitter signal	92.7.2	A positive differential voltage corresponds to tx_bit = one	M	Yes [ ]
PF3	ALERT signal	92.7.2	Transmit a periodic sequence, where each period of the sequence consists of 8 ones followed by 8 zeros, on each lane when tx_mode is set to ALERT	EEE:M	Yes [ ] N/A [ ]
PF4	Receive function	92.7.3	Converts four electrical signals from the MDI into four logical bit streams delivers them to the PMD service interface	M	Yes [ ]
PF5	Receiver signal	92.7.3	A positive differential voltage corresponds to rx_bit = one	M	Yes [ ]
PF6	Training disabled by management	92.7.5	PMD_signal_detect_*i* set to one for *i*=0 to 3	M	Yes [ ]
PF7	PMD_signal_detect_*i* asserted, rx_mode=QUIET	92.7.5	Set to one within 500 ns following the application of the signal defined in 92.7.5 to the input of the channel corresponding to the receiver of lane *i*	EEE:M	Yes [ ]
PF8	PMD_signal_detect_*i* not asserted, rx_mode=QUIET	92.7.5	Not set to one when the signal applied to the input of the channel corresponding to the receiver of lane *i* is less than or equal to 70 mV peak-to-peak differential	EEE:M	Yes [ ]
PF9	Global_PMD_transmit_disable	92.7.6	Disables all transmitters by forcing a constant output level	GTD:M	Yes [ ] N/A [ ]
PF10	Global_PMD_transmit_disable affect on loopback	92.7.6	No effect	GTD:M	Yes [ ] N/A [ ]
PF11	Global PMD transmit disable function, tx_mode transition to QUIET	92.7.6	Turn off all transmitters when tx_mode transitions to QUIET from any other value	EEE:M	Yes [ ] N/A [ ]
PF12	Global PMD transmit disable function, tx_mode transition from QUIET	92.7.6	Turn on all transmitters when tx_mode transitions from QUIET to any other value	EEE:M	Yes [ ] N/A [ ]
PF13	PMD_transmit_disable_*i* variable	92.7.7	When set to one, the transmitter for lane *i* satisfies the requirements of Table 92–6	LTD:M	Yes [ ] N/A [ ]

Item	Feature	Sub-clause	Value/Comment	Status	Support
PF14	PMD lane-by-lane transmit disable function affect on loopback	92.7.7	No effect	LTD:M	Yes [ ] N/A [ ]
PF15	PMD_fault variable mapping to MDIO	92.7.9	Mapped to the fault bit as specified in 45.2.1.2.3	MD:M	Yes [ ] N/A [ ]
PF16	PMD_transmit_fault variable mapping to MDIO	92.7.10	Mapped to Transmit fault bit as specified in 45.2.1.7.4	MD:M	Yes [ ] N/A [ ]
PF17	PMD_receive_fault variable mapping to MDIO	92.7.11	Mapped to Receive fault bit as specified in 45.2.1.7.5	MD:M	Yes [ ] N/A [ ]
PF18	PMD control function	92.7.12	Each lane uses the same control function as 10GBASE-KR, as defined in 72.6.10	M	Yes [ ]
PF19	PMD control response time	92.7.12	Response time less than 2 ms.	M	Yes [ ]
PF20	Training frame structure	92.7.12	Defined in 72.6.10 but adjusted for 100GBASE-CR4 signaling rate and use of uncorrelated training patterns	M	Yes [ ]
PF21	Training pattern	92.7.12	4094 bits from the output of a pseudo-random bit sequence (PRBS) generator defined in Table 92–5 followed by two zeros	M	Yes [ ]
PF22	Training pattern seed	92.7.12	Set to the value in Table 92–5 at the start of the training pattern	M	Yes [ ]
PF23	PMD control function variable mapping	92.7.12	Map variables to the appropriate bits as specified in 45.2.1.81	MD:M	Yes [ ] N/A [ ]

92.14.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
MF1	Global_PMD_signal_detect	92.7.4	Set to the value described in 45.2.1.9.7	O	Yes [ ]
MF2	Global_PMD_signal_detect	92.7.4	Set defined by the training state diagram in Figure 72–5	O	Yes [ ] No [ ]
MF3	Lane-by-Lane Signal Detect function	92.7.5	Sets PMD_signal_detect_n values on a lane-by-lane basis per requirements of 92.7.5	MD:M	Yes [ ] N/A [ ]
MF4	Lane-by-Lane Signal Detect function	92.7.5	If training is disabled by management, PMD_signal_-detect_*i* is set to one for i=0 to 3.	MD:M	Yes [ ] N/A [ ]
MF5	Lane-by-Lane Signal Detect function	92.7.5	Apply EEE requirements as specified in 92.7.5	EEE:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
MF6	Global_PMD_transmit_disable	92.7.6	Apply EEE requirements as specified in 92.7.6.	EEE:M	Yes [ ]
MF7	PMD_fault function	92.7.9	Mapped to the fault bit as specified in 45.2.1.2.3	MD:M	Yes [ ] N/A [ ]
MF8	PMD_transmit_fault function	92.7.10	Mapped to the PMD_transmit_fault bit as specified in 45.2.1.7.4	MD:M	Yes [ ] N/A [ ]
MF9	PMD_receive_fault function	92.7.11	Contributes to the PMA/PMD receive fault bit as specified in 45.2.1.7.5	MD:M	Yes [ ] N/A [ ]

92.14.4.3 Transmitter specifications

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	Test fixture return loss	92.11.1.1	Meets equation constraints	M	Yes [ ]
TC2	Test fixture insertion loss	92.11.1.2	Meets equation constraints	M	Yes [ ]
TC3	Signaling rate per lane	92.8.3.9	25.78125 GBd ± 100 ppm	M	Yes [ ]
TC4	Peak-to-peak differential output voltage	92.8.3.1	Less than or equal to 1200 mV regardless of transmit equalizer setting	M	Yes [ ]
TC5	Peak-to-peak differential output voltage, transmitter disabled	92.8.3.1	Less than or equal to 35 mV	M	Yes [ ]
TC6	DC common-mode output voltage	92.8.3.1	Between 0 V and 1.9 V with respect to signal ground	M	Yes [ ]
TC7	AC common-mode output voltage	92.8.3.1	Less than or equal to 30 mV RMS with respect to signal ground	M	Yes [ ]
TC8	The peak-to-peak differential output voltage	92.8.3.1	Less than 30 mV within 500 ns of the transmitter being disabled.	EEE:M	Yes [ ]
TC9	The peak-to-peak differential output voltage	92.8.3.1	Greater than 720 mV within 500 ns of the transmitter being enabled	EEE:M	Yes [ ]
TC10	The peak-to-peak differential output voltage	92.8.3.1	Meets the requirements of 92.8.3 within 1 μs of the transmitter being enabled.	EEE:M	Yes [ ]
TC11	DC common-mode output voltage while the transmitter is disabled.	92.8.3.1	Maintained to within ±150 mV of the value for the enabled transmitter.	EEE:M	Yes [ ]
TC12	Common-mode output voltage requirements	92.8.3.1	Met regardless of the transmit equalizer setting	M	Yes [ ]
TC13	Differential output return loss (min)	92.8.3.2	Meets equation constraints	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
TC14	Reference impedance for differential return loss measurements	92.8.3.2	100 Ω	M	Yes [ ]
TC15	Common-mode to differential mode output return loss	92.8.3.3	Meets equation constraints	M	Yes [ ]
TC16	Steady-state voltage, v_f	92.8.3.5.2	0.34 min, 0.6 max	M	Yes [ ]
TC17	Linear fit pulse peak (min)	92.8.3.5.2	$0.45 \times v_f$	M	Yes [ ]
TC18	Coefficient initialization	92.8.3.5.3	Satisfies the requirements of 92.8.3.5.3.	M	Yes [ ]
TC19	Normalized coefficient step size for “increment”	92.8.3.5.4	Between 0.0083 and 0.05	M	Yes [ ]
TC20	Normalized coefficient step size for “decrement”	92.8.3.5.4	Between –0.05 and –0.0083	M	Yes [ ]
TC21	Maximum post-cursor equalization ratio	92.8.3.5.5	Greater than or equal to 4	M	Yes [ ]
TC22	Maximum pre-cursor equalization ratio	92.8.3.5.5	Greater than or equal to 1.54	M	Yes [ ]
TC23	Transmitter output SNDR	92.8.3.7	Greater than or equal to 26 dB	M	Yes [ ]
TC24	Even-odd jitter	92.8.3.8.1	Less than or equal to 0.035 UI regardless of the transmit equalization setting	M	Yes [ ]
TC25	Effective bounded uncorrelated jitter	92.8.3.8.2	Less than or equal to 0.1 UI peak-to-peak regardless of the transmit equalization setting	M	Yes [ ]
TC26	Effective total uncorrelated jitter	92.8.3.8.2	Less than or equal to 0.18 UI RMS regardless of the transmit equalization setting	M	Yes [ ]

92.14.4.4 Receiver specifications

Item	Feature	Subclause	Value/Comment	Status	Support
RC1	Test fixture return loss	92.11.1.1	Meets the requirements of 92.11.3.2	M	Yes [ ]
RC2	Test fixture insertion loss	92.11.1.2	Meets equation constraints	M	Yes [ ]
RC3	Receiver input amplitude tolerance	92.8.4.1	1200 mV measured at TP2	M	Yes [ ]
RC4	Differential input return loss	92.8.4.2	Meets equation constraints	M	Yes [ ]
RC5	Reference impedance for differential return loss measurements	92.8.4.2	100 Ω	M	Yes [ ]
RC6	Common-mode input return loss	92.8.4.3	Meets equation constraints	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
RC7	Interference tolerance	92.8.4.4	Satisfy requirements summarized in Table 92–8	M	Yes [ ]
RC8	Interference tolerance	92.8.4.4.4	Pattern generator output amplitude	M	Yes [ ]
RC9	Interference tolerance	92.8.4.4.4	Pattern generator jitter specification	M	Yes [ ]
RC10	Receiver jitter tolerance for each lane	92.8.4.5	The RS-FEC symbol error ratio less than or equal to 10^{-4} for each test case	M	Yes [ ]
RC11	Signaling rate, per lane	92.8.4.6	25.78125 GBd ± 100 ppm	M	Yes [ ]

92.14.4.5 Cable assembly specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CA1	Differential reference impedance	92.10.1	100 Ω	CBL:M	Yes [ ] N/A [ ]
CA2	Minimum insertion loss	92.10.2	Per Equation (92–26)	CBL:M	Yes [ ] N/A [ ]
CA3	Maximum insertion loss	92.10.2	Less than or equal to 22.48 dB at 12.8906 GHz.	CBL:M	Yes [ ] N/A [ ]
CA4	Return loss	92.10.3	Per Equation (92–27)	CBL:M	Yes [ ] N/A [ ]
CA5	Differential to common-mode input and output return loss	92.10.4	Per Equation (92–28)	CBL:M	Yes [ ] N/A [ ]
CA6	Differential to common-mode conversion loss	92.10.5	Per Equation (92–29)	CBL:M	Yes [ ] N/A [ ]
CA7	Common-mode to common-mode return loss	92.10.6	Per Equation (92–30)	CBL:M	Yes [ ] N/A [ ]
CA8	Cable assembly Channel Operating Margin (COM)	92.10.7	Greater than 3 dB	CBL:M	Yes [ ] N/A [ ]
CA9	Test fixture reference printed circuit board insertion loss	92.11.2	Per Equation (92–35)	CBL:M	Yes [ ] N/A [ ]
CA10	Mated test fixture specifications	92.11.3	Verified in both directions illustrated in Figure 92–18 except insertion loss verified at either test interface in directions illustrated in Figure 92–18.	CBL:M	Yes [ ] N/A [ ]
CA11	Mated test fixture insertion loss	92.11.3.1	Per Equation (92–36) and Equation (92–37)	CBL:M	Yes [ ] N/A [ ]
CA12	Mated test fixture RMS insertion loss deviation	92.11.3.1	Less than 0.13 dB	CBL:M	Yes [ ] N/A [ ]
CA13	Mated test fixture return loss	92.11.3.2	Per Equation (92–38)	CBL:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
CA14	Mated test fixture common-mode conversion loss	92.11.3.3	Per Equation (92–39)	CBL:M	Yes [ ] N/A [ ]
CA15	Mated test fixture common-mode return loss	92.11.3.4	Per Equation (92–40)	CBL:M	Yes [ ] N/A [ ]
CA16	Mated test fixture common-mode to differential mode return loss	92.11.3.5	Per Equation (92–41)	CBL:M	Yes [ ] N/A [ ]
CA17	Mated test fixtures integrated crosstalk noise	92.11.3.6	Per Equation (92–44), through Equation (92–48) and Table 92–13	CBL:M	Yes [ ] N/A [ ]
CA18	Cable assembly connector type	92.12.1.1	100GBASE-CR4 Style-1 plug (SFF-8665 plug)	CAST1:M	Yes [ ] N/A [ ]
CA19	Pin assignments	92.12.1.1	Per Table 92–15	CAST1:M	Yes [ ] N/A [ ]
CA20	Cable assembly connector type	92.12.1.2	100GBASE-CR4 Style-2 plug (CFP4)	CAST2:M	Yes [ ] N/A [ ]
CA21	Pin assignments	92.12.1.2	Per Table 92–16	CAST2:M	Yes [ ] N/A [ ]
CA22	AC-coupling	92.12.1	3 dB cutoff	CBL:M	Yes [ ]

92.14.4.6 MDI connector specifications

Item	Feature	Subclause	Value/Comment	Status	Support
MDC1	MDI connector type	92.12.1.1	100GBASE-CR4 Style-1 receptacle (SFF-8665 receptacle)	MDIST1:M	Yes [ ] N/A [ ]
MDC2	MDI connector type	92.12.1.2	100GBASE-CR4 Style-2 receptacle (CFP4 receptacle)	MDIST2:M	Yes [ ]

92.14.4.7 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	Environmental specifications	92.13		M	Yes [ ]

93. Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-KR4

93.1 Overview

This clause specifies the 100GBASE-KR4 PMD and baseband medium. There are three associated annexes. Annex 93A defines characteristics of electrical backplanes, Annex 93B extends the electrical backplane reference model with additional informative test points, and Annex 93C defines the test method for receiver interference tolerance.

When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 93–1, to the medium through the MDI and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 93–1—Physical Layer clauses associated with the 100GBASE-KR4 PMD

Associated clause	100GBASE-KR4
81—RS	Required
81—CGMII[a]	Optional
82—PCS for 100GBASE-R	Required
91—RS-FEC	Required
83—PMA for 100GBASE-R[b]	Required
83A—CAUI-10	Optional
83D—CAUI-4	Optional
73—Auto-Negotiation	Required
78—Energy-Efficient Ethernet	Optional

[a]The CGMII is an optional interface. However, if the CGMII is not implemented, a conforming implementation must behave functionally as though the RS and CGMII were present.
[b]There are limitations on the number of PMA lanes that may be used between sublayers, see 83.3.

A 100GBASE-KR4 PHY with the optional Energy-Efficient Ethernet (EEE) capability may optionally enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization.

Figure 93–1 shows the relationship of the 100GBASE-KR4 PMD sublayer and MDI to the ISO/IEC Open System Interconnection (OSI) reference model.

The receive path of the RS-FEC sublayer may have the option to perform error detection without correction to reduce the data delay (see 91.5.3.3). When the receive path of the RS-FEC sublayer performs error correction, the link is required to operate with a BER of 10^{-5} or better. When the RS-FEC sublayer is configured to bypass error correction, the link is required to operate with a BER of 10^{-12} or better. In this context, a link consists of a compliant PMD transmitter, a compliant PMD receiver, and a channel meeting the requirements of 93.9.1.

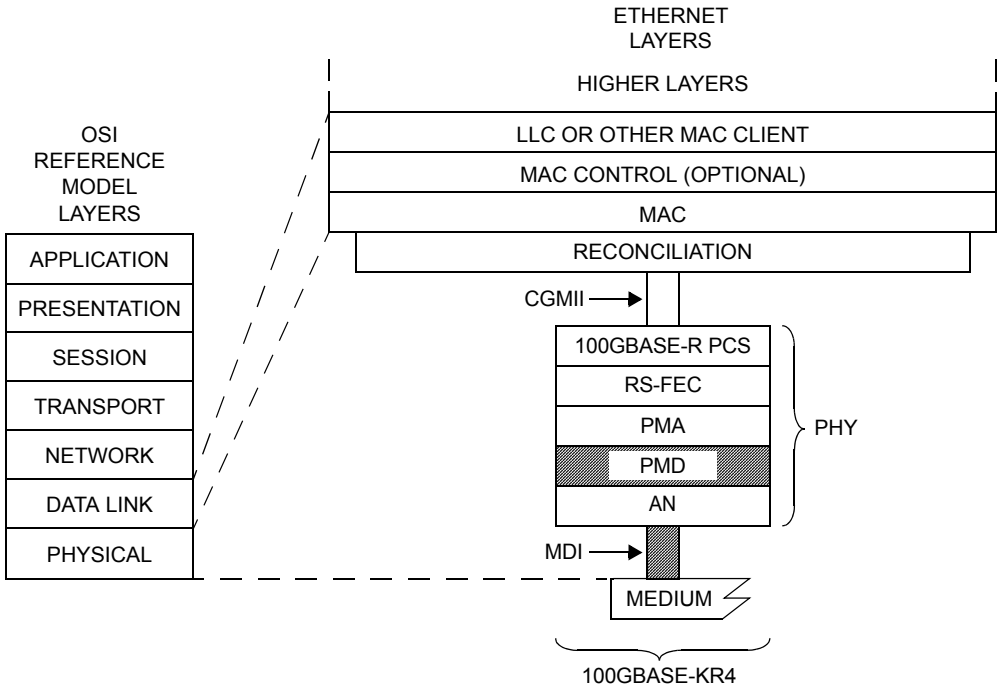

AN = AUTO-NEGOTATION
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION

Figure 93–1—100GBASE-KR4 PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

For a complete Physical Layer, this specification is considered to be satisfied by a frame loss ratio (see 1.4.223) of less than 6.2×10^{-10} for 64-octet frames with minimum interpacket gap.

93.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 100GBASE-KR4 PMD. The service interface for this PMD is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data. The PMD translates the encoded data to and from signals suitable for the medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_*i*.request
PMD:IS_UNITDATA_*i*.indication
PMD:IS_SIGNAL.indication

The 100GBASE-KR4 PMD has four parallel bit streams, hence *i* = 0 to 3. The PMA (or the PMD) continuously sends four parallel bit streams to the PMD (or the PMA), one per lane, each at a nominal signaling rate of 25.78125 GBd.

The SIGNAL_OK parameter of the PMD:IS_SIGNAL.indication primitive corresponds to the variable Global_PMD_signal_detect as defined in 93.7.4. When Global_PMD_signal_detect is one, SIGNAL_OK shall be assigned the value OK. When Global_PMD_signal_detect is zero, SIGNAL_OK shall be assigned the value FAIL. When SIGNAL_OK is FAIL, the PMD:IS_UNITDATA_*i*.indication parameters are undefined.

If the optional EEE deep sleep capability is supported, then the PMD service interface includes two additional primitives as follows:

PMD:IS_TX_MODE.request
PMD:IS_RX_MODE.request

93.3 PCS requirements for Auto-Negotiation (AN) service interface

The PCS associated with this PMD is required to support the AN service interface primitive AN_LINK.indication defined in 73.9. (See 82.6.)

The 100GBASE-KR4 PHY may be extended using CAUI-n as a physical instantiation of the inter-sublayer service interface between devices. If CAUI-n is instantiated, the AN_LINK(link_status).indication is relayed from the device with the PCS sublayer to the device with the AN sublayer by means at the discretion of the implementer. As examples, the implementer may employ use of pervasive management or employ a dedicated electrical signal to relay the state of link_status as indicated by the PCS sublayer on one device to the AN sublayer on the other device.

93.4 Delay constraints

The sum of the transmit and the receive delays at one end of the link contributed by the 100GBASE-KR4 PMD, AN, and the medium in one direction shall be no more than 2048 bit times (4 pause_quanta or 20.48 ns). It is assumed that the one way delay through the medium is no more than 800 bit times (8 ns).

A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

93.5 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the RS-FEC sublayer. Skew and Skew Variation are defined in 80.5 and specified at the points SP1 to SP6 shown in Figure 80–8.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 80.5.

93.6 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several registers that provide control and status information for and about the PMD. If MDIO is implemented, it shall map MDIO control bits to PMD control variables as shown in Table 93–2, and MDIO status bits to PMD status variables as shown in Table 93–3.

Table 93–2—100GBASE-KR4 MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable	1.9.0	Global_PMD_transmit_disable
PMD transmit disable 3 to PMD transmit disable 0	PMD transmit disable	1.9.4 to 1.9.1	PMD_transmit_disable_3 to PMD_transmit_disable_0
Restart training	BASE-R PMD control	1.150.0	mr_restart_training
Training enable	BASE-R PMD control	1.150.1	mr_training_enable
Polynomial identifier 3	PMD training pattern 3	1.1453.12:11	identifier_3
Seed 3	PMD training pattern 3	1.1453.10:0	seed_3
Polynomial identifier 2	PMD training pattern 2	1.1452.12:11	identifier_2
Seed 2	PMD training pattern 2	1.1452.10:0	seed_2
Polynomial identifier 1	PMD training pattern 1	1.1451.12:11	identifier_1
Seed 1	PMD training pattern 1	1.1451.10:0	seed_1
Polynomial identifier 0	PMD training pattern 0	1.1450.12:11	identifier_0
Seed 0	PMD training pattern 0	1.1450.10:0	seed_0

Table 93–3—100GBASE-KR4 MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Fault	PMA/PMD status 1	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect	1.10.0	Global_PMD_signal_detect
PMD receive signal detect 3 to PMD receive signal detect 0	PMD receive signal detect	1.10.4 to 1.10.1	PMD_signal_detect_3 to PMD_signal_detect_0
100GBASE-KR4 deep sleep	EEE capability	1.16.10	—
Receiver status 3	BASE-R PMD status	1.151.12	rx_trained_3
Frame lock 3	BASE-R PMD status	1.151.13	frame_lock_3
Start-up protocol status 3	BASE-R PMD status	1.151.14	training_3
Training failure 3	BASE-R PMD status	1.151.15	training_failure_3
Receiver status 2	BASE-R PMD status	1.151.8	rx_trained_2
Frame lock 2	BASE-R PMD status	1.151.9	frame_lock_2
Start-up protocol status 2	BASE-R PMD status	1.151.10	training_2
Training failure 2	BASE-R PMD status	1.151.11	training_failure_2
Receiver status 1	BASE-R PMD status	1.151.4	rx_trained_1
Frame lock 1	BASE-R PMD status	1.151.5	frame_lock_1
Start-up protocol status 1	BASE-R PMD status	1.151.6	training_1
Training failure 1	BASE-R PMD status	1.151.7	training_failure_1
Receiver status 0	BASE-R PMD status	1.151.0	rx_trained_0
Frame lock 0	BASE-R PMD status	1.151.1	frame_lock_0
Start-up protocol status 0	BASE-R PMD status	1.151.2	training_0
Training failure 0	BASE-R PMD status	1.151.3	training_failure_0

93.7 PMD functional specifications

93.7.1 Link block diagram

One direction for one lane of a 100GBASE-KR4 link is shown in Figure 93–2.

93.7.2 PMD Transmit function

The PMD transmit function shall convert the four bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_*i*.request (*i*=0 to 3) into four separate electrical signals. The four electrical signals shall then be delivered to the MDI, all according to the transmit electrical specifications in 93.8.1. A positive differential output voltage (SL*i*<p> minus SL*i*<n>) shall correspond to tx_bit = one.

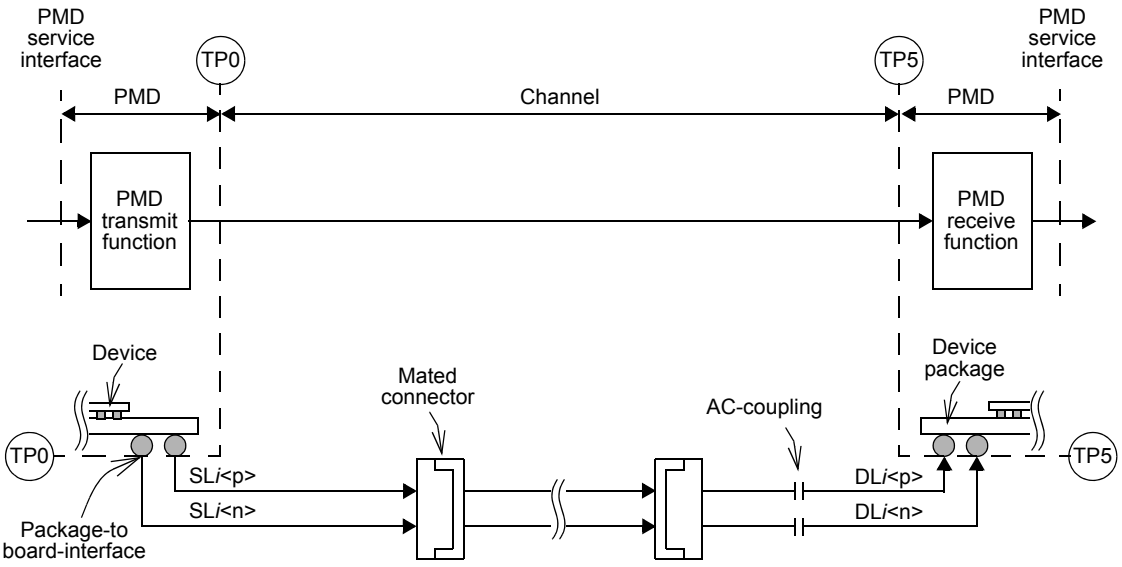

Figure 93–2—100GBASE-KR4 link (one direction for one lane is illustrated)

If the optional EEE deep sleep capability is supported, the following requirements apply. When tx_mode is set to ALERT, the PMD transmit function shall transmit a periodic sequence, where each period of the sequence consists of 8 ones followed by 8 zeros, on each lane, with the transmit equalizer coefficients set to the preset values (see 72.6.10.2.3.1). This sequence is transmitted regardless of the value of tx_bit presented by the PMD:IS_UNITDATA_*i*.request primitive.When tx_mode is not set to ALERT, the transmit equalizer coefficients are set to the values determined via the start-up protocol (see 93.7.12).

93.7.3 PMD Receive function

The PMD receive function shall convert the four electrical signals from the MDI into four bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_*i*.indication (*i*=0 to 3). A positive differential input voltage (DL*i*<p> minus DL*i*<n>) shall correspond to rx_bit = one.

93.7.4 Global PMD signal detect function

The variable Global_PMD_signal_detect is the logical AND of the values of PMD_signal_detect_*i* for *i*=0 to 3.

When the MDIO is implemented, this function maps the variable Global_PMD_signal_detect to the register and bit defined in 93.6.

93.7.5 PMD lane-by-lane signal detect function

The PMD lane-by-lane signal detect function is used by the 100GBASE-KR4 PMD to indicate the successful completion of the start-up protocol by the PMD control function (see 93.7.12). PMD_signal_detect_*i* (where *i* represents the lane number in the range 0 to 3) is set to zero when the value of the variable signal_detect is set to false by the Training state diagram for lane *i* (see Figure 72–5). PMD_signal_detect_*i* is set to one when the value of signal_detect for lane *i* is set to true.

If training is disabled by the management variable mr_training_enable (see 93.6), PMD_signal_detect_*i* shall be set to one for *i*=0 to 3.

If the optional EEE deep sleep capability is supported, the following requirements apply. The value of PMD_signal_detect_*i* (for *i*=0 to 3) is set to zero when rx_mode is first set to QUIET. While rx_mode is set to QUIET, PMD_signal_detect_*i* shall be set to one within 500 ns of the application of the ALERT pattern defined in 93.7.2, with peak-to-peak differential voltage of 720 mV as measured at TP0a, to the differential pair at the input of the channel that connects the transmitter to the receiver of lane *i*. While rx_mode is set to QUIET, PMD_signal_detect_*i* shall not be set to one when the voltage applied to the input of the differential pair of the channel that connects the transmitter to the receiver of lane *i* is less than or equal to 60 mV peak-to-peak differential.

When the MDIO is implemented, this function maps the variables to registers and bits as defined in 93.6.

93.7.6 Global PMD transmit disable function

The Global PMD transmit disable function is mandatory if the EEE deep sleep capability is supported and is otherwise optional. When implemented, it allows all of the transmitters to be disabled with a single variable.

a) When 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 93–4.
b) If a PMD fault (93.7.9) is detected, then the PMD may set Global_PMD_transmit_disable to one.
c) Loopback, as defined in 93.7.8, shall not be affected by Global_PMD_transmit_disable.
d) The following additional requirements apply when the optional EEE deep sleep capability is supported. The Global PMD transmit disable function shall turn off all of the transmitters as specified in 93.8.1.3 when tx_mode transitions to QUIET from any other value. The Global PMD transmit disable function shall turn on all of the transmitters as specified in 93.8.1.3 when tx_mode transitions from QUIET to any other value.

93.7.7 PMD lane-by-lane transmit disable function

The PMD transmit disable function is optional and allows the electrical transmitter in each lane to be selectively disabled. When this function is supported, it shall meet the following requirements:

a) When a PMD_transmit_disable_*i* variable (where *i* represents the lane number in the range 0 to 3) is set to one, this function shall turn off the transmitter associated with that variable such that it drives a constant level (i.e., no transitions) and does not exceed the maximum differential peak-to-peak output voltage specified in Table 93–4.
b) If a PMD fault (93.7.9) is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the electrical transmitter in each lane.
c) Loopback, as defined in 93.7.8, shall not be affected by PMD_transmit_disable_*i*.

93.7.8 Loopback mode

Local loopback mode is provided by the adjacent PMA (see 83.5.8) as a test function to the device. When loopback mode is enabled, transmission requests passed to each transmitter are sent directly to the corresponding receiver, overriding any signal detected by each receiver on its attached link. Note that loopback mode does not affect the state of the transmitter, which continues to send data (unless disabled).

Control of the loopback function is specified in 83.5.8.

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.

93.7.9 PMD fault function

PMD_fault is the logical OR of PMD_receive_fault, PMD_transmit_fault, and any other implementation specific fault. If the MDIO is implemented, PMD_fault shall be mapped to the fault bit as specified in 45.2.1.2.3.

93.7.10 PMD transmit fault function

The PMD transmit fault function is optional. The faults detected by this function are implementation specific, but the assertion of Global_PMD_transmit_disable is not considered a transmit fault.

If PMD_transmit_fault is set to one, then Global_PMD_transmit_disable should also be set to one.

If the MDIO interface is implemented, then PMD_transmit_fault shall be mapped to the Transmit fault bit as specified in 45.2.1.7.4.

93.7.11 PMD receive fault function

The PMD receive fault function is optional. The faults detected by this function are implementation specific. A fault is indicated by setting the variable PMD_receive_fault to one.

If the MDIO interface is implemented, then PMD_receive_fault shall be mapped to the Receive fault bit specified in 45.2.1.7.5.

93.7.12 PMD control function

Each lane of the 100GBASE-KR4 PMD shall use the same control function as 10GBASE-KR, as defined in 72.6.10, with the following differences:

a) The training frame structure used by the 100GBASE-KR4 PMD control function shall be as defined in 72.6.10 with the exception that 25.78125 GBd symbols replace 10.3125 GBd symbols and 100GBASE-KR4 UI replace 10GBASE-KR UI.
b) In addition to the coefficient update process specified in 72.6.10.2.5, the period from receiving a new request to responding to that request shall be less than 2 ms, except during the first 50 ms following the beginning of the start-up protocol. The beginning of the start-up protocol is defined to be entry into the AN_GOOD_CHECK state in Figure 73–10. The start of the period is the frame marker of the training frame with the new request and the end of the period is the frame marker of the training frame with the corresponding response. A new request occurs when the coefficient update field is different from the coefficient field in the preceding frame. The response occurs when the coefficient status report field is updated to indicate that the corresponding action is complete.
c) The training pattern defined in 72.6.10.2.6 shall be replaced with the set of training patterns defined in 92.7.12, which are designed to minimize the correlation between physical lanes.

The variables rx_trained_*i*, frame_lock_*i*, training_*i,* and training_failure_*i* (where *i* goes from 0 to 3) report status for each lane and are equivalent to rx_trained, frame_lock, training, and training_failure as defined in 72.6.10.3.1.

If the MDIO interface is implemented, then this function shall map these variables to the appropriate bits in the BASE-R PMD status register (Register 1.151) as specified in 45.2.1.81.

93.8 100GBASE-KR4 electrical characteristics

93.8.1 Transmitter characteristics

Transmitter characteristics measured at TP0a are summarized in Table 93–4.

Table 93–4—Summary of transmitter characteristics at TP0a

Parameter	Subclause reference	Value	Units
Signaling rate	93.8.1.2	25.78125±100 ppm	GBd
Differential peak-to-peak output voltage (max.) Transmitter disabled Transmitter enabled	93.8.1.3	 30 1200	 mV mV
DC common-mode output voltage (max.)	93.8.1.3	1.9	V
DC common-mode output voltage (min.)	93.8.1.3	0	V
AC common-mode output voltage (RMS, max.)	93.8.1.3	12	mV
Differential output return loss (min.)	93.8.1.4	Equation (93–3)	dB
Common-mode output return loss (min.)	93.8.1.4	Equation (93–4)	dB
Output waveform Steady-state voltage v_f (max.) Steady-state voltage v_f (min.) Linear fit pulse peak (min.) Normalized coefficient step size (min.) Normalized coefficient step size (max.) Pre-cursor full-scale range (min.) Post-cursor full-scale range (min.)	 93.8.1.5.2 93.8.1.5.2 93.8.1.5.2 93.8.1.5.4 93.8.1.5.4 93.8.1.5.5 93.8.1.5.5	 0.6 0.4 0.71 × v_f 0.0083 0.05 1.54 4	 V V V — — — —
Signal-to-noise-and-distortion ratio (min.)	93.8.1.6	27	dB
Output jitter (max.) Even-odd jitter Effective bounded uncorrelated jitter, peak-to-peak Effective total uncorrelated jitter, peak-to-peak	93.8.1.7	 0.035 0.1 0.18	 UI UI UI

93.8.1.1 Transmitter test fixture

Unless otherwise noted, measurements of the transmitter are made at the output of a test fixture (TP0a) as shown in Figure 93–5.

The insertion loss of the test fixture shall be between 1.2 dB and 1.6 dB at 12.89 GHz. The magnitude of the insertion loss deviation of the test fixture shall be less than or equal to 0.1 dB from 0.05 to 13 GHz.

The reference insertion loss of the test fixture is defined by Equation (93–1) where f is the frequency in GHz.

$$IL_{ref}(f) = -0.0015 + 0.144\sqrt{f} + 0.069f \text{ dB} \qquad 0.05 \le f \le 25 \tag{93–1}$$

The effects of differences between the insertion loss of an actual test fixture and the reference insertion loss are to be accounted for in the measurements. The reference insertion loss is illustrated in Figure 93–3.

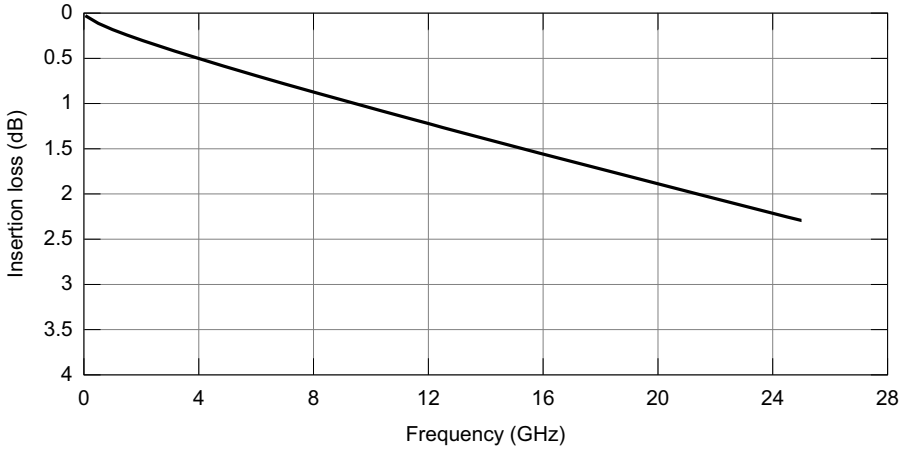

Figure 93–3—Test fixture reference insertion loss

The differential return loss of the test fixture, in dB, shall meet Equation (93–2) where f is the frequency in GHz.

$$RL_d(f) \geq \left\{ \begin{array}{ll} 20 - f & 0.05 \leq f \leq 5 \\ 15 & 5 < f \leq 13 \\ 20.57 - 0.4286f & 13 < f \leq 25 \end{array} \right\} \quad \text{dB} \tag{93–2}$$

The return loss limit is illustrated by Figure 93–4.

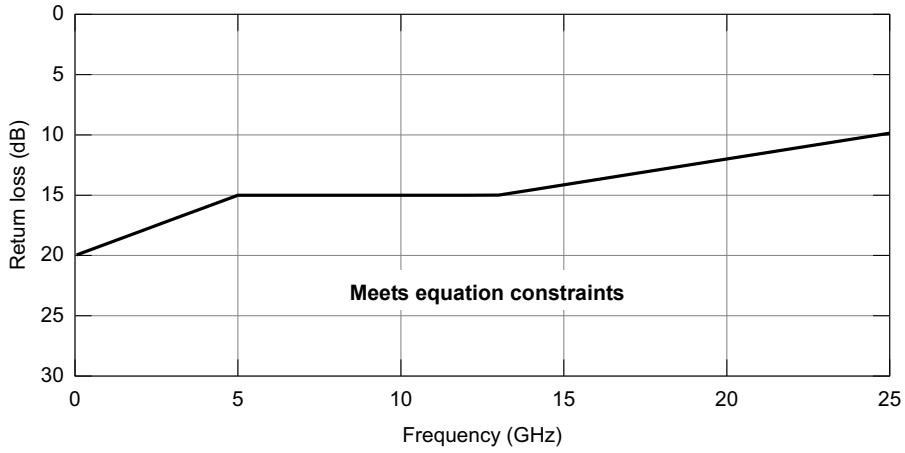

Figure 93–4—Test fixture differential return loss limit

The common-mode return loss of the test fixture shall be greater than or equal to 10 dB from 0.05 to 13 GHz.

A test system with a fourth-order Bessel-Thomson low-pass response with 33 GHz 3 dB bandwidth is to be used for all transmitter signal measurements, unless otherwise specified.

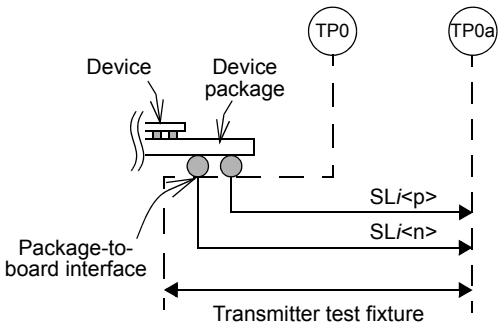

Figure 93–5—Transmitter test fixture and test points

93.8.1.2 Signaling rate and range

The 100GBASE-KR4 signaling rate shall be 25.78125 GBd ± 100 ppm per lane.

93.8.1.3 Signal levels

The differential output voltage v_{di} is defined to be SL*i*<p> minus SL*i*<n>. The common-mode output voltage v_{cmi} is defined to be one half of the sum of SL*i*<p> and SL*i*<n>. These definitions are illustrated by Figure 93–6.

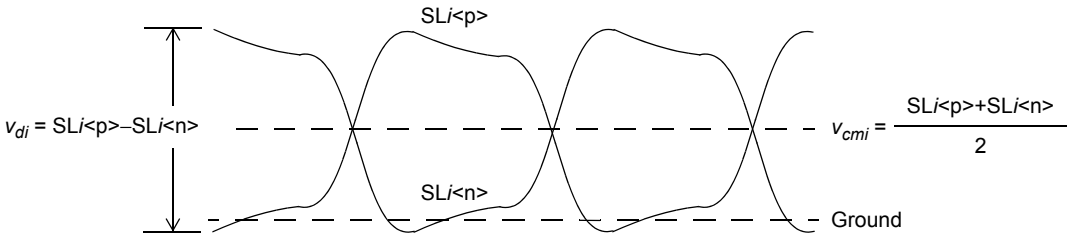

Figure 93–6—Transmitter output voltage definitions

The peak-to-peak differential output voltage shall be less than or equal to 1200 mV regardless of the transmit equalizer setting. The peak-to-peak differential output voltage shall be less than or equal to 30 mV while the transmitter is disabled (refer to 93.7.6 and 93.7.7).

The DC common-mode output voltage shall be between 0 V and 1.9 V with respect to signal ground. The AC common-mode output voltage shall be less than or equal to 12 mV RMS with respect to signal ground. Common-mode output voltage requirements shall be met regardless of the transmit equalizer setting.

If the optional EEE deep sleep capability is supported the following requirements also apply. The peak-to-peak differential output voltage shall be less than 30 mV within 500 ns of the transmitter being disabled. The peak-to-peak differential output voltage shall be greater than 720 mV within 500 ns of the transmitter being enabled. The transmitter is enabled by the assertion of tx_mode=ALERT and the preceding requirement applies when the transmitted symbols are the periodic pattern defined in 93.7.2 and the transmitter equalizer coefficients are assigned their preset values. The transmitter shall meet the requirements of 93.8.1 within 1 µs of the transmitter being enabled. While the transmitter is disabled, the DC common-mode output voltage shall be maintained to within ±150 mV of the value for the enabled transmitter.

Unless otherwise noted, differential and common-mode signal levels are measured with a PRBS9 test pattern.

93.8.1.4 Transmitter output return loss

The differential output return loss, in dB, of the transmitter shall meet Equation (93–3) where *f* is the frequency in GHz. This output impedance requirement applies to all valid output levels. The reference impedance for differential return loss measurements shall be 100 Ω.

$$RL_d(f) \geq \begin{Bmatrix} 12.05 - f & 0.05 \leq f \leq 6 \\ 6.5 - 0.075f & 6 < f \leq 19 \end{Bmatrix} \quad \text{dB} \tag{93–3}$$

The differential return loss limit is illustrated by Figure 93–7.

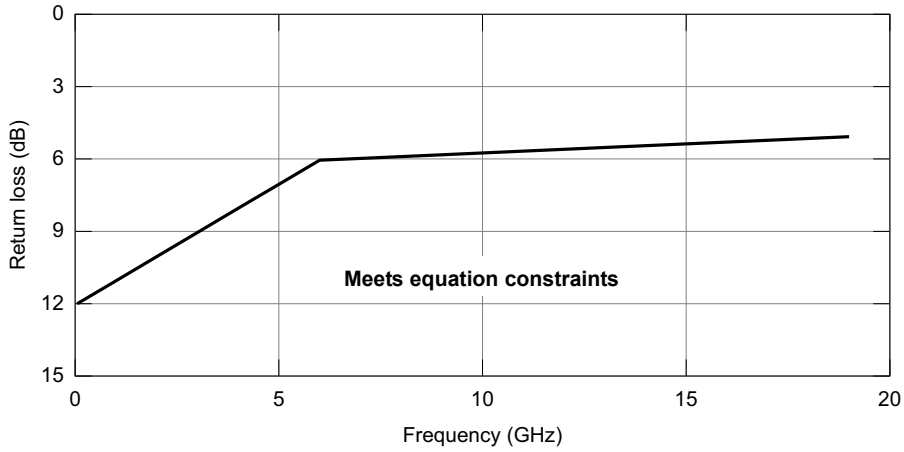

Figure 93–7—Differential return loss limit

The common-mode output return loss, in dB, of the transmitter shall meet Equation (93–4) where *f* is the frequency in GHz. This output impedance requirement applies to all valid output levels. The reference impedance for common-mode return loss measurements shall be 25 Ω.

$$RL_{cm}(f) \geq \begin{Bmatrix} 9.05 - f & 0.05 \leq f \leq 6 \\ 3.5 - 0.075f & 6 < f \leq 19 \end{Bmatrix} \quad \text{dB} \tag{93–4}$$

The common-mode return loss limit is illustrated by Figure 93–8.

93.8.1.5 Transmitter output waveform

The 100GBASE-KR4 transmit function includes programmable equalization to compensate for the frequency-dependent loss of the channel and facilitate data recovery at the receiver. The functional model for the transmit equalizer is the three tap transversal filter shown in Figure 93–9.

The state of the transmit equalizer and hence the transmitted output waveform may be manipulated via the PMD control function defined in 93.7.12 or via the management interface. The transmit function responds to a set of commands issued by the link partner's receive function and conveyed by a back-channel communications path.

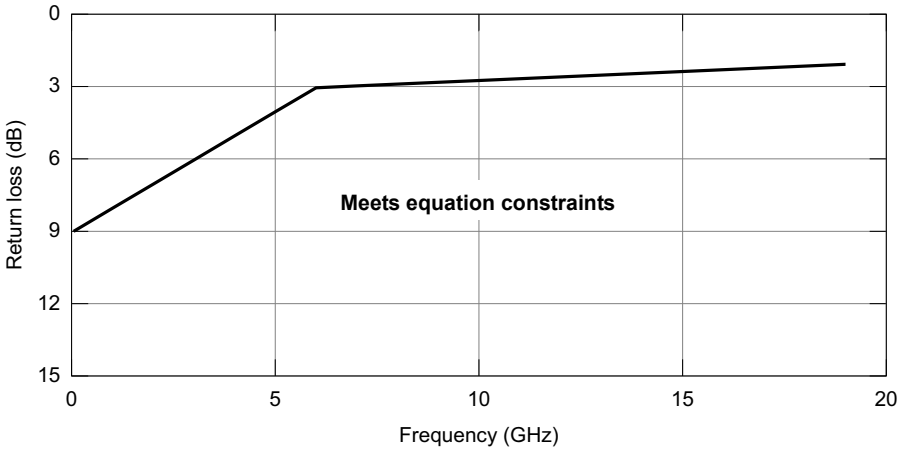

Figure 93–8—Common-mode return loss limit

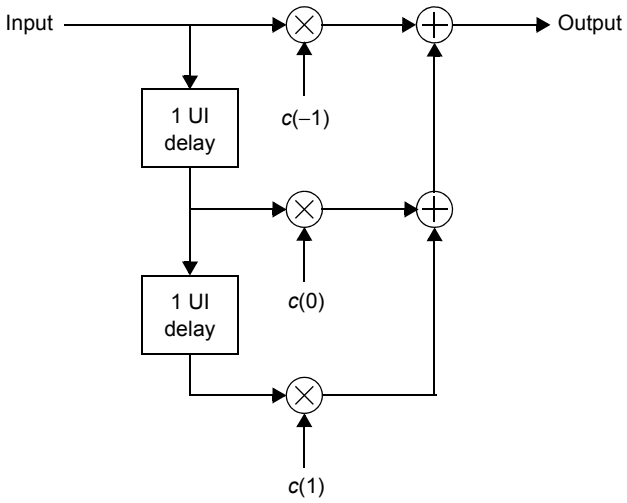

Figure 93–9—Transmit equalizer functional model

This command set includes instructions to:

Increment coefficient $c(i)$.
Decrement coefficient $c(i)$.
Hold coefficient $c(i)$ at its current value.
Set the coefficients to a predefined value (preset or initialize).

In response, the transmit function relays status information to the link partner's receive function. The status messages indicate that:

The requested update to coefficient $c(i)$ has completed (updated).
Coefficient $c(i)$ is at its minimum value.

Coefficient $c(i)$ is at its maximum value.
Coefficient $c(i)$ is ready for the next update request (not_updated).

93.8.1.5.1 Linear fit to the measured waveform

The transmitter output waveform is characterized using the procedure described in 92.8.3.5.1 with the exception that the measurement is performed at TP0a rather than TP2.

93.8.1.5.2 Steady-state voltage and linear fit pulse peak

The steady-state voltage v_f is defined to be the sum of the linear fit pulse $p(k)$ divided by M (refer to 85.8.3.3step 3). The steady-state voltage shall be greater than or equal to 0.4 V and less than or equal to 0.6 V after the transmit equalizer coefficients have been set to the "preset" values.

The peak value of $p(k)$ shall be greater than $0.71 \times v_f$ after the transmit equalizer coefficients have been set to the "preset" values.

93.8.1.5.3 Coefficient initialization

When the PMD enters the INITIALIZE state of the Training state diagram (Figure 72–5) or receives a valid request to "initialize" from the link partner, the coefficients of the transmit equalizer shall be configured such that the ratio $(c(0)+c(1)–c(–1))/(c(0)+c(1)+c(–1))$ is 1.29±10% and the ratio $(c(0)–c(1)+c(–1))/(c(0)+c(1)+c(–1))$ is 2.57±10%. These requirements apply upon the assertion a coefficient status report of "updated" for all coefficients.

93.8.1.5.4 Coefficient step size

The change in the normalized amplitude of coefficient $c(i)$ corresponding to a request to "increment" that coefficient shall be between 0.0083 and 0.05. The change in the normalized amplitude of coefficient $c(i)$ corresponding to a request to "decrement" that coefficient shall be between –0.05 and –0.0083.

The change in the normalized amplitude of the coefficient is defined to be the difference in the value measured prior to the assertion of the "increment" or "decrement" request (e.g., the coefficient update request for all coefficients is "hold") and the value upon the assertion of a coefficient status report of "updated" for that coefficient.

93.8.1.5.5 Coefficient range

When sufficient "increment" or "decrement" requests have been received for a given coefficient, the coefficient reaches a lower or upper bound based on the coefficient range or restrictions placed on the minimum steady-state differential output voltage or the maximum peak-to-peak differential output voltage.

With $c(–1)$ set to zero and both $c(0)$ and $c(1)$ having received sufficient "decrement" requests so that they are at their respective minimum values, the ratio $(c(0)–c(1))/(c(0)+c(1))$ shall be greater than or equal to 4.

With $c(1)$ set to zero and both $c(–1)$ and $c(0)$ having received sufficient "decrement" requests so that they are at their respective minimum values, the ratio $(c(0)–c(–1))/(c(0)+c(–1))$ shall be greater than or equal to 1.54.

Note that a coefficient may be set to zero by first asserting the preset control and then manipulating the other coefficients as required by the test.

93.8.1.6 Transmitter output noise and distortion

Signal-to-noise-and-distortion ratio (SNDR) measured at the transmitter output using the method described in 92.8.3.7 shall be greater than 27 dB regardless of the transmit equalizer setting.

93.8.1.7 Transmitter output jitter

The conditions for the measurement of transmitter output jitter (jitter filter, test pattern, etc.) are defined in 92.8.3.8.

Even-odd jitter is defined in 92.8.3.8.1. Even-odd jitter shall be less than or equal to 0.035 UI regardless of the transmit equalization setting.

Effective bounded uncorrelated jitter and effective total uncorrelated jitter are defined in 92.8.3.8.2. The effective bounded uncorrelated jitter shall be less than or equal to 0.1 UI peak-to-peak regardless of the transmit equalization setting. The effective total uncorrelated jitter shall be less than or equal to 0.18 UI peak-to-peak regardless of the transmit equalization setting.

93.8.2 Receiver characteristics

Receiver characteristics measured at TP5a are summarized in Table 93–5.

Table 93–5—Summary of receiver characteristics at TP5a

Parameter	Subclause reference	Value	Units
Differential input return loss (min.)	93.8.2.2	Equation (93–3)	dB
Differential to common-mode return loss (min.)	93.8.2.2	Equation (93–5)	dB
Interference tolerance	93.8.2.3	Table 93–6	—
Jitter tolerance	93.8.2.4	Table 93–7	—

93.8.2.1 Receiver test fixture

Unless otherwise noted, measurements of the receiver are made at the input to a test fixture as shown in Figure 93–10.

The insertion loss of the test fixture shall be between 1.2 dB and 1.6 dB at 12.89 GHz. The magnitude of the insertion loss deviation of the test fixture shall be less than or equal to 0.1 dB from 0.05 GHz to 13 GHz.

The reference insertion loss of the test fixture is defined by Equation (93–1) where f is the frequency in GHz. The effects of differences between the insertion loss of an actual test fixture and the reference insertion loss are to be accounted for in the measurements. The reference insertion loss is illustrated in Figure 93–3.

The differential return loss of the test fixture, in dB, shall meet Equation (93–2) where f is the frequency in GHz. The return loss limit is illustrated by Figure 93–4.

The common-mode return loss of the test fixture shall be greater than or equal to 10 dB from 0.05 GHz to 13 GHz.

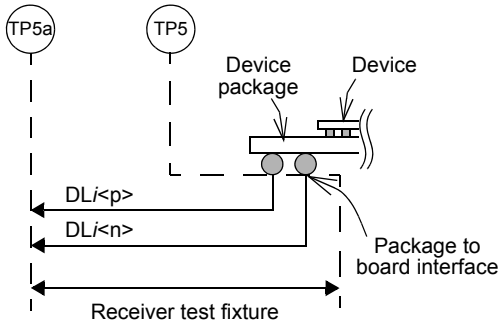

Figure 93–10—Receiver test fixture and test points

93.8.2.2 Receiver input return loss

The differential input return loss, in dB, of the receiver shall meet Equation (93–3) where *f* is the frequency in GHz. The reference impedance for differential return loss measurements shall be 100 Ω. The differential input return loss limit is illustrated by Figure 93–7.

The differential to common-mode return loss, in dB, of the receiver shall meet Equation (93–5).

$$RL_{cd}(f) = \left\{ \begin{array}{ll} 25 - 1.44f & 0.05 \le f \le 6.95 \\ 15 & 6.95 < f \le 19 \end{array} \right\} \quad \text{dB} \qquad (93\text{–}5)$$

The differential to common-mode return loss limit is illustrated by Figure 93–11.

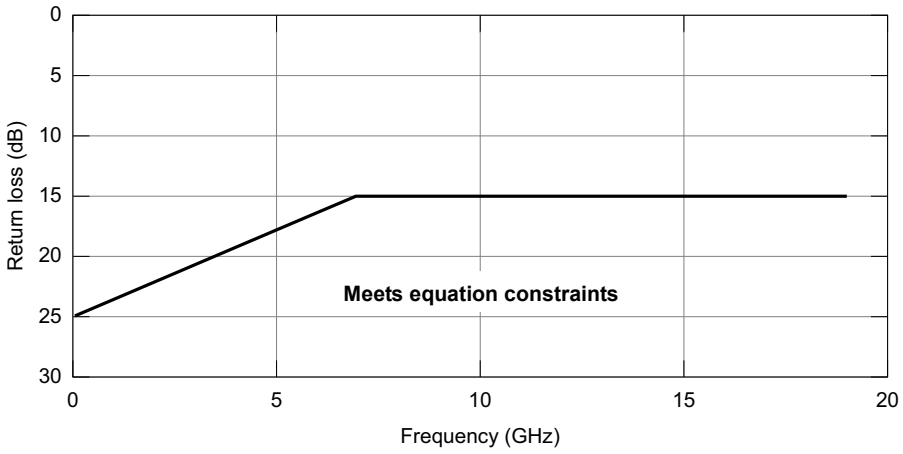

Figure 93–11—Receiver differential to common-mode return loss limit

93.8.2.3 Receiver interference tolerance

The receiver interference tolerance test setup and method are defined in Annex 93C. The receiver on each lane shall meet the RS-FEC symbol error ratio requirement with the channel defined for each test listed in Table 93–6. The parameter RSS_DFE4 is a figure of merit for the test channel that is defined in 93A.2.

The following considerations apply to the interference tolerance test. The test transmitter meets the specifications in 93.8.1 as measured at TP0a (see Figure 93C–3). The test transmitter is constrained such that for any transmit equalizer setting the differential peak-to-peak voltage (see 93.8.1.3) is less than or equal to 800 mV and the pre- and post-cursor equalization ratios (see 93.8.1.5.5) are less than or equal to 1.54 and 4, respectively. The lowest frequency f_{NSD1} for constraints on the noise spectral density is 1 GHz. The return loss of the test setup in Figure 93C–4 measured at TP5 replica meets the requirements of Equation (93–2).

The values of the parameters required for the calculation of Channel Operating Margin (COM) are given in Table 93–8 with the following exceptions. The COM parameter σ_{RJ} is set to the measured value of effective random jitter (see 92.8.3.8.2), the COM parameter A_{DD} is set to half the measured value of effective bounded uncorrelated jitter (see 93.8.1.7), and the COM parameter SNR_{TX} is set to the value of SNDR measured at TP0a (see 93.8.1.6). Tests 1 and 2 are for the case when error correction is bypassed in the RS-FEC sublayer (see 91.5.3.3) and for these cases COM is computed with a DER_0 value of 10^{-12}. The test pattern to be used is any valid PCS output (such as scrambled idle), which is subsequently encoded by the RS-FEC sublayer.

A test system with a fourth-order Bessel-Thomson low-pass response with 33 GHz 3 dB bandwidth is to be used for measurement of the signal applied by the pattern generator and for measurements of the broadband noise.

Table 93–6—Receiver interference tolerance parameters

Parameter	Test 1 values		Test 2 values		Test 3 values		Test 4 values		Units
	Min	Max	Min	Max	Min	Max	Min	Max	
RS-FEC symbol error ratio[a]	—	10^{-11}	—	10^{-11}	—	10^{-4}	—	10^{-4}	—
Insertion loss at 12.89 GHz[b]	—	16	30	—	—	30	35	—	dB
Coefficients of fitted insertion loss[c]									
a_0	–0.9	0.9	–0.9	0.9	–0.9	0.9	–0.9	0.9	dB
a_1	0	3.3	0	3.3	0	3.3	0	3.3	dB/GHz$^{1/2}$
a_2	0	—	0	—	0	—	0	—	dB/GHz
a_4	0	0.022	0	0.03	0	0.03	0	0.043	dB/GHz^2
RSS_DFE4	0.05	—	0.05	—	0.05	—	0.05	—	—
COM, including effects of broadband noise	—	3	—	3	—	3	—	3	dB

[a]The FEC symbol error ratio is measured in step 10 of the receiver interference tolerance method defined in 93C.2.
[b]Measured between TPt and TP5 (see Figure 93C–4).
[c]Coefficients are calculated from the insertion loss measured between TPt and TP5 (see Figure 93C–4) using the method in 93A.3 with f_{min} = 0.05 GHz, f_{max} = 25.78125 GHz, and maximum Δf = 0.01 GHz.

93.8.2.4 Receiver jitter tolerance

Receiver jitter tolerance is measured using the test setup shown in Figure 93–12, or its equivalent, for the receiver of each lane.

The test transmitter meets the specifications of 93.8.1. It is constrained so that its differential peak-to-peak output voltage does not exceed 800 mV at TP0a regardless of the transmitter equalizer setting (see 93.8.1.3). It is further constrained so that its maximum pre-cursor equalization ratio is 1.54 and its maximum post-cursor equalization ratio is 4 (see 93.8.1.5.5).

The test channel meets the requirements of the interference tolerance test channel using Test 4 values (see 93.8.2.3). No broadband noise is added for this test.

Receiver jitter tolerance is verified for each pair of jitter frequency and peak-to-peak amplitude values listed in Table 93–7. The synthesizer frequency is set to the specified jitter frequency and the synthesizer output amplitude is adjusted until the specified peak-to-peak jitter amplitude for that frequency is measured at TP0a. The output of the ISI channel is connected to the input of the receiver under test at TP5a. The link is initialized and the PMD start-up protocol is allowed to complete, thereby optimizing the test transmitter equalizer. The test transmitter then transmits any valid PCS output (such as scrambled idle), which is subsequently encoded by the RS-FEC sublayer. The RS-FEC symbol error ratio is measured using the errored symbol counter FEC_symbol_error_counter_*i* where *i* is the lane number of the receiver under test.

The RS-FEC symbol error ratio shall be less than or equal to 10^{-4} for each case listed in Table 93–7.

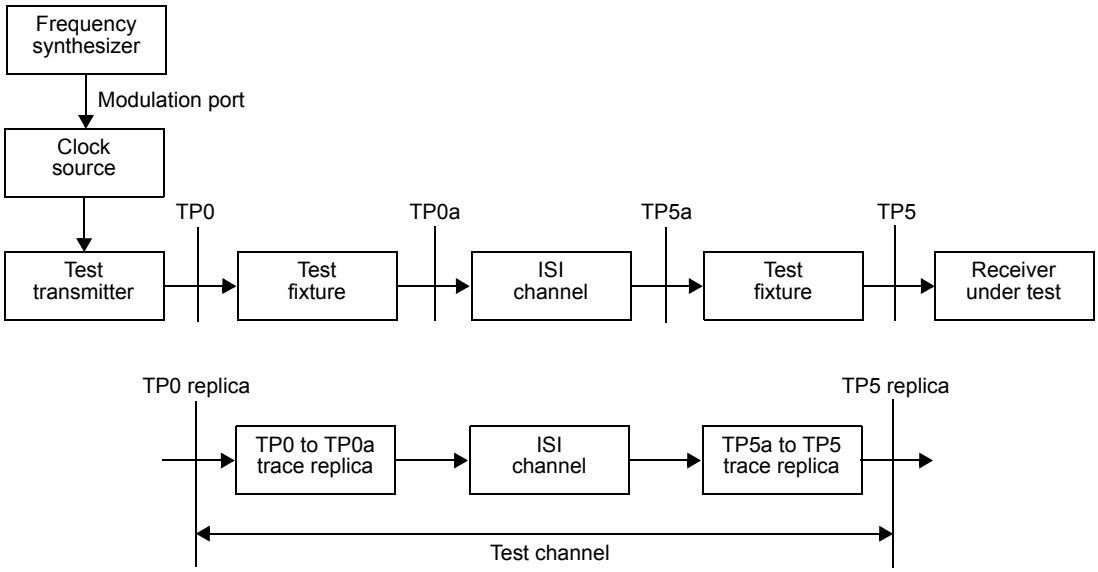

Figure 93–12—Jitter tolerance test setup

Table 93–7—Receiver jitter tolerance parameters

Parameter	Case A values	Case B values	Units
Jitter frequency	190	940	kHz
Peak-to-peak jitter amplitude	5	1	UI

93.9 Channel characteristics

93.9.1 Channel Operating Margin

The Channel Operating Margin (COM) is computed using the procedure in 93A.1 with the Test 1 and Test 2 values in Table 93–8. Test 1 and Test 2 differ in the value of the device package model transmission line length z_p.

COM shall be greater than or equal to 3 dB for each test. This minimum value allocates margin for practical limitations on the receiver implementation as well as the largest step size allowed for transmitter equalizer coefficients.

The receive path of the RS-FEC sublayer may have the option to perform error detection without correction to reduce the data delay (see 91.5.3.3). Channels that are compatible with this mode of operation shall meet this COM requirement with the value of DER_0 set to 10^{-12}.

93.9.2 Insertion loss

The insertion loss, in dB, of the channel is recommended to meet Equation (93–6).

$$IL(f) \le \begin{cases} 1.5 + 4.6\sqrt{f} + 1.318f & 0.05 \le f \le f_b/2 \\ -12.71 + 3.7f & f_b/2 < f \le f_b \end{cases} \text{ (dB)} \qquad (93\text{–}6)$$

where

f is the frequency in GHz
f_b is the signaling rate (25.78125) in GHz
$IL(f)$ is the insertion loss at frequency f

The insertion loss limit is illustrated by Figure 93–13.

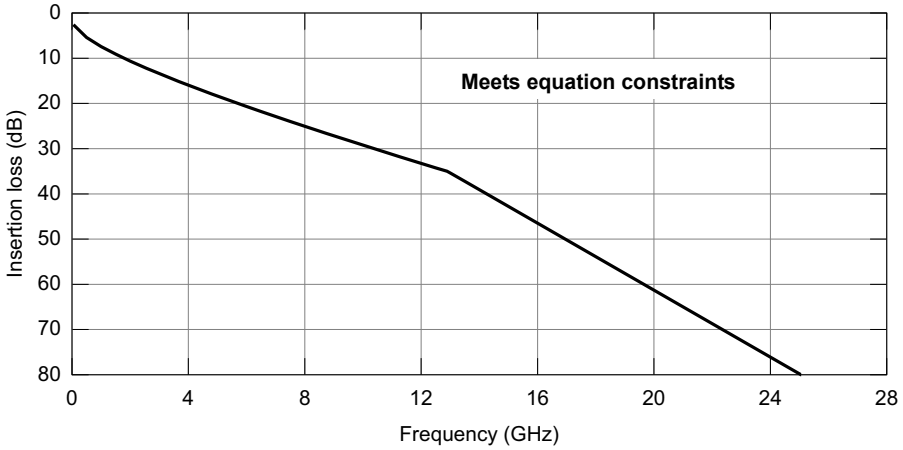

Figure 93–13—Insertion loss limit

93.9.3 Return loss

The return loss, in dB, of the channel is recommended to meet Equation (93–7).

Table 93–8—COM parameter values

Parameter	Symbol	Value	Units
Signaling rate	f_b	25.78125	GBd
Maximum start frequency	$f_{\min}$	0.05	GHz
Maximum frequency step	Δf	0.01	GHz
Device package model Single-ended device capacitance Transmission line length, Test 1 Transmission line length, Test 2 Single-ended package capacitance at package-to-board interface	 C_d z_p z_p C_p	 2.5×10^{-4} 12 30 1.8×10^{-4}	 nF mm mm nF
Single-ended reference resistance	R_0	50	Ω
Single-ended termination resistance	R_d	55	Ω
Receiver 3 dB bandwidth	f_r	$0.75 \times f_b$	GHz
Transmitter equalizer, minimum cursor coefficient	$c(0)$	0.62	—
Transmitter equalizer, pre-cursor coefficient Minimum value Maximum value Step size	$c(-1)$	 –0.18 0 0.02	 — — —
Transmitter equalizer, post-cursor coefficient Minimum value Maximum value Step size	$c(1)$	 –0.38 0 0.02	 — — —
Continuous time filter, DC gain Minimum value Maximum value Step size	g_{DC}	 –12 0 1	 dB dB dB
Continuous time filter, zero frequency	f_z	f_b / 4	GHz
Continuous time filter, pole frequencies	f_{p1} f_{p2}	f_b / 4 f_b	GHz
Transmitter differential peak output voltage Victim Far-end aggressor Near-end aggressor	 A_v A_{fe} A_{ne}	 0.4 0.4 0.6	 V V V
Number of signal levels	L	2	—
Level separation mismatch ratio	R_{LM}	1	—
Transmitter signal-to-noise ratio	SNR_{TX}	27	dB
Number of samples per unit interval	M	32	—
Decision feedback equalizer (DFE) length	N_b	14	UI
Normalized DFE coefficient magnitude limit, for n = 1 to N_b	$b_{\max}(n)$	1	—
Random jitter, RMS	σ_{RJ}	0.01	UI
Dual-Dirac jitter, peak	A_{DD}	0.05	UI
One-sided noise spectral density	η_0	5.2×10^{-8}	V^2/GHz
Target detector error ratio	DER_0	10^{-5}	—

$$RL_d(f) \geq \begin{cases} 12 & 0.05 \leq f \leq f_b/4 \\ 12 - 15\log_{10}(4f/f_b) & f_b/4 < f \leq f_b \end{cases} \quad \text{dB} \qquad (93\text{–}7)$$

where

f is the frequency in GHz
f_b is the signaling rate (25.78125) in GHz
$RL(f)$ is the return loss at frequency f

The differential return loss limit is illustrated by Figure 93–14.

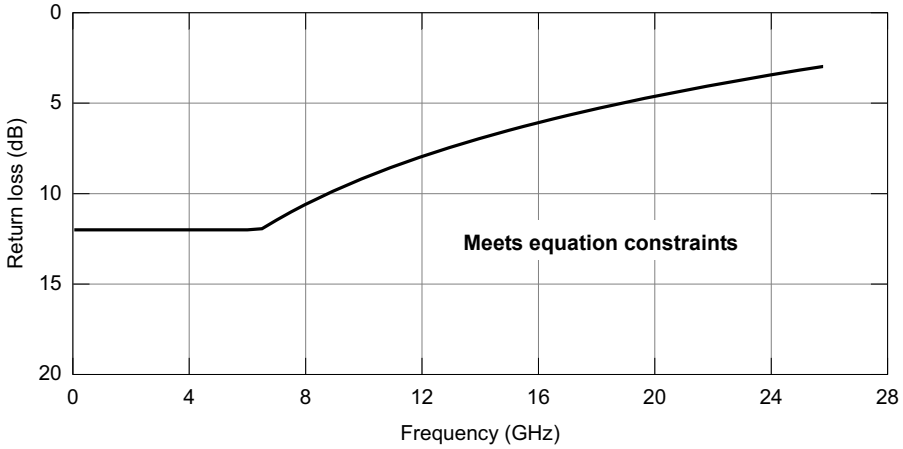

Figure 93–14—Differential return loss limit

93.9.4 AC-coupling

The 100GBASE-KR4 transmitter shall be AC-coupled to the receiver. Common-mode specifications are defined as if the DC-blocking capacitor is implemented between TP0 and TP5. Should the capacitor be implemented outside TP0 and TP5, the common-mode specifications in Table 93–4 may not be appropriate.

The impact of a DC-blocking capacitor implemented between TP0 and TP5 is accounted for within the channel specifications. Should the capacitor be implemented outside TP0 and TP5, it is the responsibility of implementers to consider any necessary modifications to common-mode and channel specifications required for interoperability as well as any impact on the verification of transmitter and receiver compliance.

The low-frequency 3 dB cutoff of the AC-coupling shall be less than 50 kHz.

93.10 Environmental specifications

93.10.1 General safety

All equipment subject to this clause shall conform to applicable sections (including isolation requirements) of IEC 60950-1.

93.10.2 Network safety

The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

93.10.3 Installation and maintenance guidelines

It is recommended that sound installation practice, as defined by applicable local codes and regulations, be followed in every instance in which such practice is applicable.

93.10.4 Electromagnetic compatibility

A system integrating the 100GBASE-KR4 PHY shall comply with applicable local and national codes for the limitation of electromagnetic interference.

93.10.5 Temperature and humidity

A system integrating the 100GBASE-KR4 PHY 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.

93.11 Protocol implementation conformance statement (PICS) proforma for Clause 93, Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-KR4[22]

93.11.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 93, Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-KR4, 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.

93.11.2 Identification

93.11.2.1 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).	

93.11.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 93, Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-KR4
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	

[22] *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.

93.11.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
CGMII	CGMII	93.1	Interface is supported	O	Yes [ ] No [ ]
PCS	100GBASE-R PCS	93.1		M	Yes [ ]
RS-FEC	100GBASE-R RS-FEC	93.1		M	Yes [ ]
PMA	100GBASE-R PMA	93.1		M	Yes [ ]
CAUI-10	CAUI-10	93.1	Interface is supported	O	Yes [ ] No [ ]
CAUI-4	CAUI-4	93.1	Interface is supported	O	Yes [ ] No [ ]
AN	Auto-negotiation	93.1		M	Yes [ ]
DC	Delay constraints	93.4	Conforms to delay constraints specified in 93.4	M	Yes [ ]
DSC	Skew constraints	93.5	Conforms to the Skew and Skew Variation constraints specified in 93.5	M	Yes [ ]
*MD	MDIO capability	93.6	Registers and interface supported	O	Yes [ ] No [ ]
*EEE	EEE capability	93.1	Capability is supported	O	Yes [ ] No [ ]
*GTD	Global PMD transmit disable function	93.7.6	Function is supported	EEE:M	Yes [ ] No [ ]
*LTD	PMD lane-by-lane transmit disable function	93.7.7	Function is supported	O	Yes [ ] No [ ]
*CHNL	Channel	93.9	Channel specifications not applicable to a PHY manufacturer.	O	Yes [ ] No [ ]

93.11.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-KR4

93.11.4.1 Functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
FS1	PMD transmit function	93.7.2	Converts four logical bit streams from the PMD service interface into four electrical signals and delivers them to the MDI	M	Yes [ ]
FS2	Mapping of logical signals to electrical signals	93.7.2	Positive differential output voltage corresponds to tx_bit=one	M	Yes [ ]
FS3	ALERT signal	93.7.2	Transmit a periodic sequence, where each period of the sequence consists of 8 ones followed by 8 zeros, on each lane when tx_mode is set to ALERT	EEE:M	Yes [ ] N/A [ ]
FS4	PMD receive function	93.7.3	Converts four electrical signals from the MDI into four logical bit streams delivers them to the PMD service interface	M	Yes [ ]
FS5	Mapping of electrical signals to logical signals	93.7.3	Positive differential input voltage corresponds to rx_bit=one	M	Yes [ ]
FS6	SIGNAL_OK mapping	93.2	Set to OK when Global_PMD_signal_detect is one and set to FAIL when Global_PMD_signal_detect is zero	M	Yes [ ]
FS7	Training disabled by variable mr_training_enable	93.7.5	PMD_signal_detect_*i* set to one for *i*=0 to 3	M	Yes [ ]
FS8	PMD_signal_detect_*i* asserted, rx_mode=QUIET	93.7.5	Set to one within 500 ns following the application of the signal defined in 93.7.5 to the input of the channel corresponding to the receiver of lane *i*	EEE:M	Yes [ ] N/A [ ]
FS9	PMD_signal_detect_*i* not asserted, rx_mode=QUIET	93.7.5	Not set to one when the signal applied to the input of the channel corresponding to the receiver of lane *i* is less than or equal to 60 mV peak-to-peak differential	EEE:M	Yes [ ] N/A [ ]
FS10	Global_PMD_transmit_disable variable	93.7.6	When set to one, all transmitters satisfy the requirements of 93.7.6	GTD:M	Yes [ ] N/A [ ]
FS11	Global PMD transmit disable function affect on loopback	93.7.6	No effect	GTD:M	Yes [ ] N/A [ ]
FS12	Global PMD transmit disable function, tx_mode transition to QUIET	93.7.6	Turn off all transmitters when tx_mode transitions to QUIET from any other value	EEE:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
FS13	Global PMD transmit disable function, tx_mode transition from QUIET	93.7.6	Turn on all transmitters when tx_mode transitions from QUIET to any other value	EEE:M	Yes [ ] N/A [ ]
FS14	PMD_transmit_disable_*i* variable	93.7.7	When set to one, the transmitter for lane *i* satisfies the requirements of 93.8.1.3	LTD:M	Yes [ ] N/A [ ]
FS15	PMD lane-by-lane transmit disable function affect on loopback	93.7.7	No effect	LTD:M	Yes [ ] N/A [ ]
FS16	PMD_fault variable mapping to MDIO	93.7.9	Mapped to the fault bit as specified in 45.2.1.2.3	MD:M	Yes [ ] N/A [ ]
FS17	PMD_transmit_fault variable mapping to MDIO	93.7.10	Mapped to Transmit fault bit as specified in 45.2.1.7.4	MD:M	Yes [ ] N/A [ ]
FS18	PMD_receive_fault variable mapping to MDIO	93.7.11	Mapped to Receive fault bit as specified in 45.2.1.7.5	MD:M	Yes [ ] N/A [ ]
FS19	PMD control function	93.7.12	Defined in 72.6.10	M	Yes [ ]
FS20	Training frame structure	93.7.12	Defined in 72.6.10 but adjusted for 100GBASE-KR4 signaling rate	M	Yes [ ]
FS21	Training patterns	93.7.12	Training patterns defined in 92.7.12.	M	Yes [ ]
FS22	PMD control function variable mapping to MDIO	93.7.12	Map variables as specified in 93.7.12	MD:M	Yes [ ] N/A [ ]
FS23	PMD control response time	93.7.12	Response time less than 2 ms	M	Yes [ ]

93.11.4.2 Transmitter characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	Test fixture insertion loss	93.8.1.1	Between 1.2 dB and 1.6 dB at 12.89 GHz	M	Yes [ ]
TC2	Test fixture insertion loss deviation	93.8.1.1	Magnitude less than 0.1 dB	M	Yes [ ]
TC3	Test fixture differential return loss	93.8.1.1	Meets equation constraints	M	Yes [ ]
TC4	Test fixture common-mode return loss	93.8.1.1	Greater than or equal to 10 dB from 0.05 to 13 GHz	M	Yes [ ]
TC5	Signaling rate per lane	93.8.1.2	25.78125 GBd ± 100 ppm	M	Yes [ ]
TC6	Peak-to-peak differential output voltage	93.8.1.3	Less than or equal to 1200 mV regardless of transmit equalizer setting	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
TC7	Peak-to-peak differential output voltage, transmitter disabled	93.8.1.3	Less than or equal to 30 mV	M	Yes [ ]
TC8	DC common-mode output voltage	93.8.1.3	Between 0 V and 1.9 V with respect to signal ground	M	Yes [ ]
TC9	AC common-mode output voltage	93.8.1.3	Less than or equal to 12 mV RMS with respect to signal ground	M	Yes [ ]
TC10	Common-mode output voltage requirements	93.8.1.3	Met regardless of the transmit equalizer setting	M	Yes [ ]
TC11	Transmitter disable timing	93.8.1.3	Peak-to-peak differential output voltage less than 30 mV within 500 ns of the transmitter being disabled	EEE:M	Yes [ ] N/A [ ]
TC12	Transmitter enable timing	93.8.1.3	Peak-to-peak differential output voltage greater than 720 mV within 500 ns of the transmitter being enabled and meet all requirements of 93.8.1 within 1 μs	EEE:M	Yes [ ] N/A [ ]
TC13	Common-mode output voltage, transmitter disabled	93.8.1.3	Maintained to within ±150 mV of the value for the enabled transmitter	EEE:M	Yes [ ] N/A [ ]
TC14	Differential input return loss	93.8.1.4	Meets equation constraints	M	Yes [ ]
TC15	Reference impedance for differential return loss measurements	93.8.1.4	100 Ω	M	Yes [ ]
TC16	Common-mode output return loss	93.8.1.4	Meets equation constraints	M	Yes [ ]
TC17	Reference impedance for common-mode return loss measurements	93.8.1.4	25 Ω	M	Yes [ ]
TC18	Steady-state voltage, v_f	93.8.1.5.2	Greater than or equal to 0.4 V and less than or equal to 0.6 V after the transmit equalizer coefficients have been set to the "preset" values	M	Yes [ ]
TC19	Linear fit pulse peak	93.8.1.5.2	Greater than $0.71 \times v_f$ after the transmit equalizer coefficients have been set to the "preset" values	M	Yes [ ]
TC20	Coefficient initialization	93.8.1.5.3	Satisfies the requirements of 93.8.1.5.3.	M	Yes [ ]
TC21	Normalized coefficient step size for "increment"	93.8.1.5.4	Between 0.0083 and 0.05	M	Yes [ ]
TC22	Normalized coefficient step size for "decrement"	93.8.1.5.4	Between –0.05 and –0.0083	M	Yes [ ]
TC23	Maximum post-cursor equalization ratio	93.8.1.5.5	Greater than or equal to 4	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
TC24	Maximum pre-cursor equalization ratio	93.8.1.5.5	Greater than or equal to 1.54	M	Yes [ ]
TC25	Transmitter output noise and distortion	93.8.1.6	SNDR greater than or equal to 27 dB	M	Yes [ ]
TC26	Even-odd jitter	93.8.1.7	Less than or equal to 0.035 UI regardless of the transmit equalization setting	M	Yes [ ]
TC27	Effective bounded uncorrelated jitter	93.8.1.7	Less than or equal to 0.1 UI peak-to-peak regardless of the transmit equalization setting	M	Yes [ ]
TC28	Effective total uncorrelated jitter	93.8.1.7	Less than or equal to 0.18 UI peak-to-peak regardless of the transmit equalization setting	M	Yes [ ]

93.11.4.3 Receiver characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
RC1	Test fixture insertion loss	93.8.2.1	Between 1.2 dB and 1.6 dB at 12.89 GHz	M	Yes [ ]
RC2	Test fixture insertion loss deviation	93.8.2.1	Magnitude less than 0.1 dB	M	Yes [ ]
RC3	Test fixture differential return loss	93.8.2.1	Meets equation constraints	M	Yes [ ]
RC4	Test fixture common-mode return loss	93.8.2.1	Greater than or equal to 10 dB from 0.05 GHz to 13 GHz	M	Yes [ ]
RC5	Differential input return loss	93.8.2.2	Meets equation constraints	M	Yes [ ]
RC6	Reference impedance for differential return loss measurements	93.8.2.2	100 Ω	M	Yes [ ]
RC7	Differential to common-mode return loss	93.8.2.2	Meets equation constraints	M	Yes [ ]
RC8	Receiver interference tolerance	93.8.2.3	Satisfy requirements summarized in Table 93–6	M	Yes [ ]
RC9	Receiver jitter tolerance	93.8.2.4	RS-FEC symbol error ratio less than or equal to 10^{-4} for each case listed in Table 93–7	M	Yes [ ]

93.11.4.4 Channel characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	Channel Operating Margin (COM)	93.9.1	Greater than or equal to 3 dB for each test listed in Table 93–8	CHNL:M	Yes [ ] N/A [ ]
CC2	COM when error correction is bypassed by the RS-FEC sublayer	93.9.1	Greater than or equal to 3 dB for each test listed in Table 93–8 but with $DER_0 = 10^{-12}$	CHNL:M	Yes [ ] N/A [ ]
CC3	AC-coupling	93.9.4	Channel AC-couples the transmitter to the receiver	CHNL:M	Yes [ ] N/A [ ]
CC4	AC-coupling 3 dB cut-off frequency	93.9.4	Less than 50 kHz	CHNL:M	Yes [ ] N/A [ ]

93.11.4.5 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
ES1	General safety	93.10.1	Conform to applicable sections of IEC 60950-1	M	Yes [ ]
ES2	Electromagnetic compatibility	93.10.4	Comply with applicable local and national codes	M	Yes [ ]

94. Physical Medium Attachment (PMA) sublayer, Physical Medium Dependent (PMD) sublayer, and baseband medium, type 100GBASE-KP4

94.1 Overview

This clause specifies the Physical Medium Attachment (PMA) sublayers, Physical Medium Dependent (PMD) sublayer, and medium for the 100GBASE-KP4 PHY.

When forming a complete Physical Layer, the PMA shall be connected to the RS-FEC, the PMD shall be connected to the medium through the MDI as shown in Figure 94–1, and the PMA and PMD shall be connected to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 94–1—Physical Layer clauses associated with the 100GBASE-KP4 PMD

Associated clause	100GBASE-KP4
81—RS	Required
81—CGMII[a]	Optional
82—PCS for 100GBASE-R	Required
83—PMA for 100GBASE-R	Optional
83A—CAUI-10	Optional
83D—CAUI-4	Optional
91—RS-FEC	Required
73—Auto-Negotiation	Required
78—Energy-Efficient Ethernet	Optional

[a]The CGMII is an optional interface. However, if the CGMII is not implemented, a conforming implementation must behave functionally as though the RS and CGMII were present.

Figure 94–1 shows the relationship of the 100GBASE-KP4 PMA and PMD sublayers and MDI to the ISO/IEC Open System Interconnection (OSI) reference model.

Differential signals received at the MDI from a transmitter that meets the requirements of 94.2.2 and 94.3.12 and have passed through the channel specified in 94.4 are received with a BER less than 10^{-5} as measured at the PMA service interface.

For a complete Physical Layer, this specification is considered to be satisfied by a frame loss ratio (see 1.4.223) of less than 6.2×10^{-10} for 64-octet frames with minimum interpacket gap.

A 100GBASE-KP4 PHY with the optional Energy-Efficient Ethernet (EEE) capability may optionally enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization.

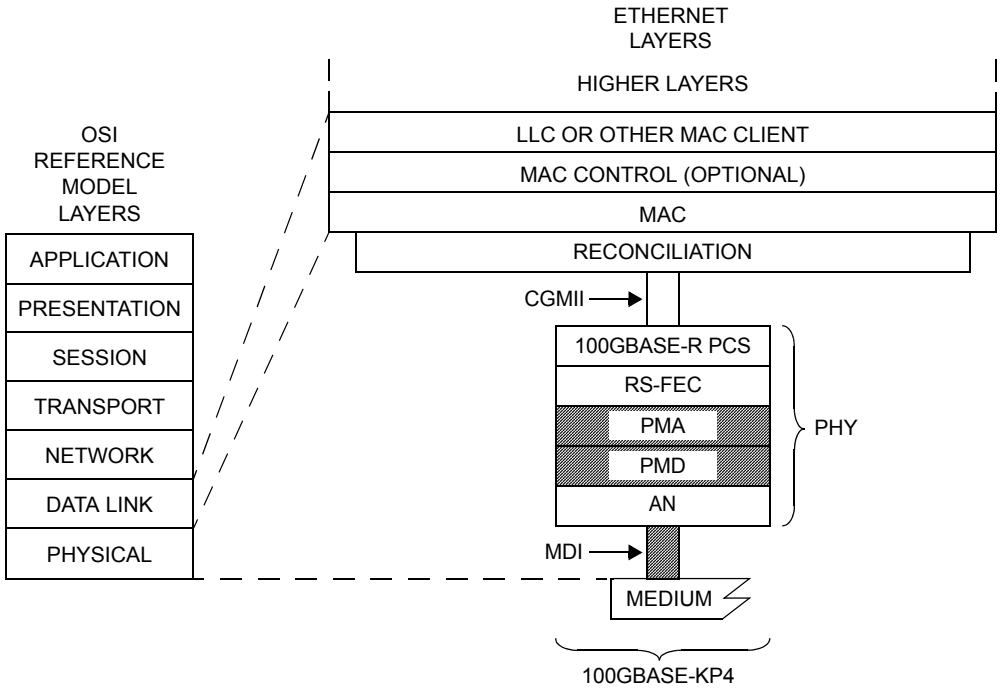

AN = AUTO-NEGOTATION
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION

Figure 94–1—100GBASE-KP4 PMA and PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

94.2 Physical Medium Attachment (PMA) Sublayer

94.2.1 PMA Service Interface

The PMA service interface for 100GBASE-KP4 PMA is based on the inter-sublayer service interface defined in 80.3. This interface is defined in an abstract manner and does not imply any particular implementation.

The PMA service interface primitives are summarized as follows:

PMA:IS_UNITDATA_*i*.request
PMA:IS_UNITDATA_*i*.indication
PMA:IS_SIGNAL.indication

If the optional EEE deep sleep capability is supported (see Clause 78, 78.1.3.3.1), then the PMA service interface includes four additional primitives as follows:

PMA:IS_TX_MODE.request
PMA:IS_RX_MODE.request
PMA:IS_ENERGY_DETECT.indication
PMA:IS_RX_TX_MODE.indication

94.2.1.1 PMA:IS_UNITDATA_i.request

The PMA:IS_UNITDATA_*i*.request (where *i*=0 to 3) primitive is used to define the transfer of four streams of data units from the PMA client to the PMA.

94.2.1.1.1 Semantics of the service primitive

PMA:IS_UNITDATA_0.request(tx_bit,start)
PMA:IS_UNITDATA_1.request(tx_bit,start)
PMA:IS_UNITDATA_2.request(tx_bit,start)
PMA:IS_UNITDATA_3.request(tx_bit,start)

The data conveyed by PMA:IS_UNITDATA_0.request to IS_UNITDATA_3.request consists of four parallel continuous streams of encoded bits, one stream for each lane. Each of the tx_bit parameters can take one of two values: one or zero. The start parameter is TRUE to indicate that the concurrent tx_bit is the first bit of the first, second, third, or fourth FEC symbol in a FEC codeword and is otherwise FALSE.

94.2.1.1.2 When generated

The PMA client continuously sends four parallel bit streams PMA:IS_UNITDATA_*i*.request(tx_bit,start) to the PMA, each at a nominal signaling rate of 26.5625 Gb/s.

94.2.1.1.3 Effect of receipt

Upon receiving each instance of PMA:IS_UNITDATA_*i*.request, the tx_bit and start parameters are passed to the PMA framing process corresponding to each stream.

94.2.1.2 PMA:IS_UNITDATA_i.indication

The PMA:IS_UNITDATA_*i*.indication (where *i* = 0 to 3) primitive is used to define the transfer of four streams of data units from the PMA to the PMA client.

94.2.1.2.1 Semantics of the service primitive

PMA:IS_UNITDATA_0.indication(rx_bit, start)
PMA:IS_UNITDATA_1.indication(rx_bit, start)
PMA:IS_UNITDATA_2.indication(rx_bit, start)
PMA:IS_UNITDATA_3.indication(rx_bit, start)

The data conveyed by PMA:IS_UNITDATA_0.indication to PMA:IS_UNITDATA_3.indication consists of four parallel continuous streams of encoded bits, one stream for each lane. Each of the rx_bit parameters can take one of two values: one or zero. The start parameter is TRUE to indicate that the concurrent rx_bit is the first bit of the first, second, third, or fourth FEC symbol in the FEC codeword and is otherwise FALSE.

94.2.1.2.2 When generated

The PMA continuously sends four parallel bit streams PMA:IS_UNITDATA_*i*.indication(rx_bit,start) to the PMA client, each at a nominal signaling rate of 26.5625 Gb/s.

94.2.1.2.3 Effect of receipt

The effect of receipt of this primitive is defined by the PMA client.

94.2.1.3 PMA:IS_SIGNAL.indication

The PMA:IS_SIGNAL.indication primitive is generated by the PMA to the PMA client to indicate the status of the receive process. This primitive is generated by the receive process to propagate the detection of severe error conditions (e.g., loss of synchronization) to the PMA client.

94.2.1.3.1 Semantics of the service primitive

PMA:IS_SIGNAL.indication(SIGNAL_OK)

The SIGNAL_OK parameter can take on one of two values: OK or FAIL. A value of FAIL denotes that invalid data is being presented (rx_bit parameters undefined) by the PMA to the PMA client. A value of OK does not guarantee valid data is being presented by the PMA to the PMA client.

94.2.1.3.2 When generated

The PMA generates the PMA:IS_SIGNAL.indication primitive to the PMA client whenever there is a change in the value of the SIGNAL_OK parameter.

94.2.1.3.3 Effect of receipt

The effect of receipt of this primitive is defined by the PMA client.

94.2.1.4 PMA:IS_TX_MODE.request

The PMA:IS_TX_MODE.request primitive communicates the tx_mode parameter generated by the PCS LPI transmit process to invoke the appropriate PMA, FEC, and PMD transmit EEE states. Without EEE deep sleep capability, this primitive is never invoked and the sublayers behave as if tx_mode = DATA.

94.2.1.4.1 Semantics of the service primitive

PMA:IS_TX_MODE.request(tx_mode)

The tx_mode parameter takes on one of up to three values: DATA, QUIET, or ALERT.

94.2.1.4.2 When generated

This primitive is generated to indicate the state of the PCS LPI transmit function.

94.2.1.4.3 Effect of receipt

When this primitive is received, PMD:IS_TX_MODE.request(tx_mode) is generated with the value received in PMA:IS_TX_MODE.request(tx_mode).

If the value is DATA or ALERT, the PMA operates normally.

If the value is QUIET, the PMA may go into a low power mode.

94.2.1.5 PMA:IS_RX_MODE.request

The PMA:IS_RX_MODE.request primitive communicates the rx_mode parameter generated by the PCS LPI receive process. Without EEE deep sleep capability, this primitive is never invoked and the sublayers behave as if rx_mode = DATA.

94.2.1.5.1 Semantics of the service primitive

PMA:IS_RX_MODE.request(rx_mode)

The rx_mode parameter takes on one of two values: DATA or QUIET.

94.2.1.5.2 When generated

This primitive is generated to indicate the state of the PCS LPI receive function.

94.2.1.5.3 Effect of receipt

When this primitive is received, PMD:IS_RX_MODE.request(rx_mode) is generated with the value received in PMA:IS_RX_MODE.request(rx_mode).

If the value is DATA, the PMA operates normally.

If the value is QUIET, the PMA may be go into a low power mode.

94.2.1.6 PMA:IS_ENERGY_DETECT.indication

The PMA:IS_ENERGY_DETECT.indication primitive is used to communicate that the PMD has detected the presence of energy on the interface following a period of quiescence. Without EEE deep sleep capability, this primitive is never invoked and has no effect.

94.2.1.6.1 Semantics of the service primitive

PMA:IS_ENERGY_DETECT.indication(energy_detect)

The parameter energy_detect is Boolean.

94.2.1.6.2 When generated

This primitive is generated by the PMA, reflecting the state of PMD:IS_SIGNAL.indication(SIGNAL_OK) received from the PMD (see 94.3.1.3). When SIGNAL_OK indicates OK, energy detect indicates TRUE. When SIGNAL_OK indicates FAIL, energy_detect indicates FALSE.

94.2.1.6.3 Effect of receipt

The effect of receipt of this primitive is defined by the PMA client sublayers that receive it.

94.2.1.7 PMA:IS_RX_TX_MODE.indication

The PMA:IS_RX_TX_MODE.indication primitive communicates the value of the rx_tx_mode parameter. This parameter indicates the value of tx_mode that the PMA sublayer has inferred from the received signal.

Without EEE deep sleep capability, the primitive is never generated and the sublayers behave as if rx_tx_mode=DATA.

94.2.1.7.1 Semantics of the service primitive

PMA:IS_RX_TX_MODE.indication(rx_tx_mode)

The parameter rx_tx_mode is assigned one of the following values: DATA, QUIET, or ALERT. DATA is assigned when the PMA is reset or when PMA frames are being received. QUIET is assigned if PMA frame reception ceases. ALERT is assigned if rx_tx_mode = QUIET and PMD:IS_SIGNAL.indication(SIGNAL_OK) transitions from FAIL to OK.

94.2.1.7.2 When generated

This primitive is generated whenever there is change in the value of the rx_tx_mode parameter.

94.2.1.7.3 Effect of receipt

The RS-FEC sublayer passes this primitive through to the PMA sublayer that may exist above.

94.2.2 PMA Transmit Functional Specifications

In the transmit direction, the role of the 100GBASE-KP4 PMA is to adapt the signal from the FEC (the PMA client) to a PAM4 encoded signal to be passed to the PMD for transfer over the attached medium. The adaptation processes shown in Figure 94–2 include insert overhead, insert termination bits, apply Gray coding, apply 1/(1+D) mod 4 precoding, and apply PAM4 encoding.

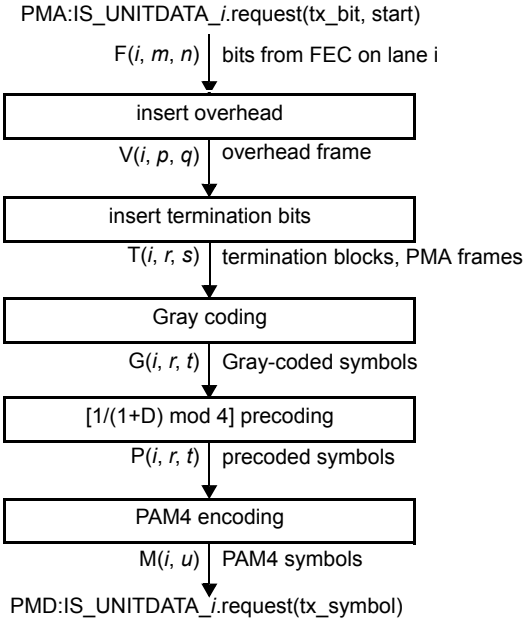

Figure 94–2—Transmit adaptation process diagram

94.2.2.1 FEC Interface

The PMA transmit process receives FEC bits via the PMA:IS_UNITDATA_*i*(tx_bit, start) primitive (see 94.2.1.1). The index *i* indicates the PMA lane number: 0, 1, 2, or 3.

On each transaction, tx_bit is assigned to F(*i*, *m*, *n*), where

i is the lane number
m is an index indicating the FEC codeword number and increments at the start of each codeword
n is an index indicating the bit number within a codeword with a range 1 to 1360

The start of a codeword is determined by the start parameter associated with the tx_bit parameter being equal to TRUE.

94.2.2.2 Overhead Frame

The PMA transmit process shall create a sequence of overhead frames by inserting 40 overhead bits for every 31280 FEC bits as specified in this subclause.

The FEC bits, F(*i*, *m*, *n*), are mapped into a continuous sequence of overhead frames. The overhead frame is 31320 bits in length.

Each bit in the overhead frame is denoted V(*i*, *p*, *q*), where:

i is the lane number
p is an index that indicates the frame number and increments at the start of each frame
q is an index that indicates the bit number within a frame with a range 1 to 31320

The first 40 bits of the frame, V(*i*, *p*, 1) to V(*i*, *p*, 40), are the overhead bits (see 94.2.2.3). The next 31280 bits, V(*i*, *p*, 41) to V(*i*, *p*, 31320), are composed of the bits from 23 consecutive FEC codewords.

The overhead bits are inserted in the frame as follows:

V(*i*, *p*, 1) = H(*i*, *p*, 1)
V(*i*, *p*, 2) = H(*i*, *p*, 2)
V(*i*, *p*, ...) = H(*i*, *p*, ...)
V(*i*, *p*, 40) = H(*i*, *p*, 40)

The FEC codeword bits are aligned such that V(*i*, *p*, 41) is the first bit of a codeword, e.g., V(*i*, *p*, 41) = F(*i*, *m*, 1). The FEC bits are inserted into the frame in the order in which they were received from the FEC, e.g., V(*i*, *p*, 42) = F(*i*, *m*, 2), V(*i*, *p*, 43) = F(*i*, *m*, 3), and so on. The method for aligning the FEC codeword with the start of the overhead frame is outside the scope of this standard.

94.2.2.3 Overhead

The PMA transmit process shall form the overhead bits in each overhead frame as specified in this subclause.

The overhead bits are denoted H(*i*, *p*, *k*), where

i is the lane number
p is an index that indicates the frame number and increments at the start of each frame
k is an index that indicates the header bit number with a range 1 to 40

Bits are mapped to the overhead in a sequence of five groups of 8 bits each. Each 8-bit group takes on the value A or An. The values of the 8 bits in A are such that A(7:0) = TX_OH_pattern(7:0). The values of the 8 bits in An are such that each bit is the inverse of the corresponding bit in A. For each lane i, the sequence

of 8-bit groups is according to the bits in TX_OH_sequence_i(4:0) such that H(*i*,*p*,((*a*+1)*8):(*a**8+1)) is equal to A(7:0) or An(7:0) if TX_OH_sequence_i(*a*) is equal to 0 or 1, respectively, where a ∈ {0,1,2,3,4}.

If the optional Clause 45 MDIO is implemented, the overhead function maps the TX_OH_pattern and TX_OH_sequence_i status variables to the registers and bits defined in 94.2.10. The default values for each of the variables are summarized in Table 94–2.

Table 94–2—Default overhead configuration values

Parameter	Values (binary)
TX_OH_pattern(7:0)	01100110
TX_OH_sequence_0(4:0)	00110
TX_OH_sequence_1(4:0)	01010
TX_OH_sequence_2(4:0)	10101
TX_OH_sequence_3(4:0)	11001

94.2.2.4 Termination Blocks

The PMA transmit process shall create a sequence of termination blocks by inserting two termination bits for every 90 overhead frame bits as specified in this subclause. The termination block is 92 bits in length. The overhead frame mapped into 348 consecutive termination blocks forms a PMA frame.

Each bit in a termination block is denoted T(*i*, *r*, *s*), where:

i is the lane number
r is an index indicating block number and increments at the start of each block
s is an index indicating the bit number within a termination block with a range 1 to 92

The first two bits in each termination block, T(*i*, *r*, 1) and T(*i*, *r*, 2), are populated with the output of a PRBS13 generator of the form specified in 94.3.10.8. For each termination block, the PRBS13 generator generates a block of 92 pseudo-random bits, R(*i*,1:92). The first two bits are used for the termination bits such that T(*i, r,* 1)=R(*i*, 1) and T(*i, r*, 2)=R(*i*, 2). The PRBS13 generator is initialized during training (94.3.10.8). Upon the transition from the last training frame to the first PMA frame the PRBS13 generator used during training advances without reseeding (see 94.3.10.7.2) and without inversion, and the output is used to generate the termination bits. The PRBS13 generator continues to advance without reseeding and without inversion.

The remaining 90 bits of each termination block, T(*i*, *r*, 3) to T(*i*, *r*, 92), are overhead frame bits (see 94.2.2.2). The overhead frame bits are aligned with the termination blocks such that the first overhead bit, V(*i*, *p*, 1), corresponds to the third bit of a termination block, T(*i*, *r*, 3).

Overhead frame bits are mapped to the termination blocks in order of location within the overhead frame, e.g., T(*i*, *r*, 4) = V(*i*, *p*, 2), T(*i*, *r*, 5) = V(*i*, *p*, 3), and so on.

The termination bit PRBS13 generator is initialized during the training process. When training is complete the state of the termination bit PRBS13 generator is retained and the resulting output is used for the termination bits in the PMA frame.

94.2.2.5 Gray Mapping

The PMA transmit process shall map consecutive pairs of bits to one of four Gray-coded symbols as specified in this subclause.

Each pair of bits, {A, B}, of each termination block are converted to a Gray-coded symbol with one of the four Gray-coded levels as follows:

{0, 0} maps to 0,
{0, 1} maps to 1,
{1, 1} maps to 2, and
{1, 0} maps to 3.

Gray-coded symbols corresponding to each termination block are denoted G(i, r, t), where:

i is the lane number
r is an index indicating the termination block number
t is an index indicating the symbol number within a termination block with a range 1 to 46

Pairing of bits is such that the first two bits of each termination block, T(i, r, 1) and T(i, r, 2), form a pair. Each bit pair {T(i, r, 2t–1), T(i, r, 2t)} maps to {A, B} and the Gray-coded result is assigned to G(i, r, t). The Gray-coded symbol G(i, r, 1) is formed from the first two bits of a termination block, the termination bits, thus forming a termination symbol.

94.2.2.6 Precoding

The PMA transmit process shall precode the Gray-coded symbols as specified in this subclause.

The precoder output symbols are denoted, P(i, r, t), where:

i is the lane number
r is an index indicating the termination block number
t is an index indicating the symbol number within a termination block with a range 1 to 46.

For each Gray-coded symbol G(i, r, t), a precoded symbol, P(i, r, t) is determined by the following algorithm:

If t = 1 then
    P(i, r, t) = G(i, r, t)
Else
    P(i, r, t) = (G(i, r, t) – P(i, r, t–1)) mod 4
End If

The bits contributing to the Gray-coded termination symbol, G(i, r, 1), are the termination bits. The precoding algorithm applies this symbol directly to the output rather than combining it with the previous non-deterministic symbols and thus this termination symbol is always deterministic.

94.2.2.7 PAM4 encoding

The PMA transmit process shall encode each precoder output symbol to one of four PAM4 levels as specified in this subclause.

The PAM4 encoded symbols are denoted M(i, u), where

i is the lane number
u is an index indicating the symbol number

Each consecutive precoder output symbol, P(*i*, *r*, *t*), is mapped to one of four PAM4 levels and assigned to the PAM4 encoder output M(*i*, u).

Mapping from the precoder output symbol P(*i*, *r*, *t*) to a PAM4 encoded symbol M(*i*, *u*) is as follows:

0 maps to –1,
1 maps to –1/3,
2 maps to +1/3, and
3 maps to +1.

94.2.2.8 PMD Interface

The PMA transmit process shall transmit each PAM4 encoded symbol, M(*i*, *u*) to the PMD via the PMD:IS_UNITDATA_*i*(tx_symbol) primitive at a symbol transfer rate of 13.59375 GBd.

94.2.3 PMA Receive Functional Specifications

The receive process shall recover the data encoded by the transmit process meeting the performance requirements specified in 94.1 after the overhead and termination bits have been removed.

The process by which the receiver recovers the data to meet this requirement is outside the scope of this standard. The signal structure encoded by the transmitter process including the overhead bits, Gray coding, and termination symbols may be leveraged by the receiver implementation at the discretion of the implementer. The remainder of this subclause specifies the receiver processes to reverse the transmitter encoding and report status.

In the received direction, the role of the 100GBASE-KP4 PMA is to adapt the PAM4 encoded signal from the PMD to a FEC encoded signal to be passed to the FEC for further processing. The adaptation processes shown in Figure 94–3 include PAM4 decoding, (1+D) mod 4 decoding, inverse Gray coding, remove termination bits, and remove overhead.

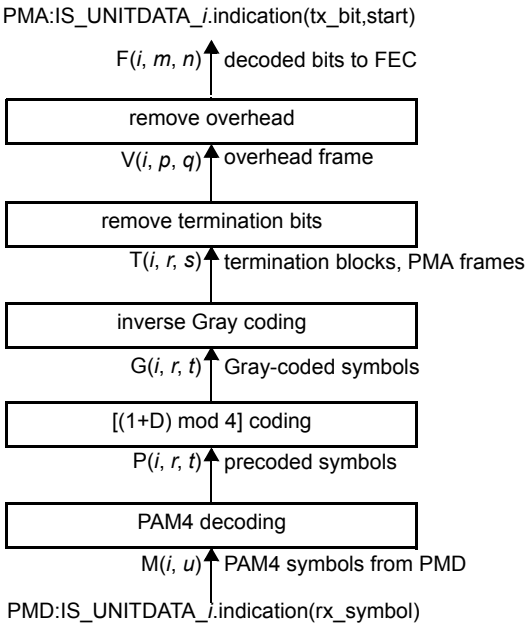

Figure 94–3—Receive adaptation process diagram

94.2.3.1 Overhead

The PMA receive process shall decode the overhead bits in each overhead frame as specified in this subclause. The format of the overhead bits is specified in 94.2.2.3.

The receive process decodes the received overhead bits in a sequence of five groups of 8 bits each into a sequence of 5 bits. Each group of 8 bits is compared with A and An, where A(7:0) = RX_OH_pattern(7:0) and the values of the 8 bits in An are such that each bit is the inverse of the corresponding bit in A. A match of each 8-bit group with A or An results in a decoded value of 0 or 1, respectively. The decoded value when the 8-bit group matches neither A or An is not specified. The decoded values from H(*i*,*p*,((*a*+1)*8):(*a**8+1)) are assigned to RX_OH_sequence_i(*a*), where $a \in \{0,1,2,3,4\}$. If all of the 8-bit groups match either A or An, the RX_OH_sequence_i(5) is set to 1 and is otherwise set to 0.

If the optional Clause 45 MDIO is implemented, the PMA receive process maps the RX_OH_pattern and RX_OH_sequence_i variables to the registers and bits defined in 94.2.10. The default values for each of the control variables are summarized in Table 94–3.

Table 94–3—Default overhead control values

Parameter	Values (binary)
RX_OH_pattern(7:0)	01100110

94.2.4 Skew constraints

Skew considerations for the 100GBASE-KP4 PMA, PMD, and AN are specified in 94.3.4.

94.2.5 Delay constraints

Delay considerations for the 100GBASE-KP4 PMA, PMD, AN, and medium are specified in 94.3.3.

94.2.6 Link status

The PMA shall provide link status information to the PMA client using the PMA:IS_SIGNAL.indication primitive (see 94.2.1.3). The PMA continuously monitors the link status reported by the PMD from the PMD:IS_SIGNAL.indication primitive, and uses this as input to Signal Indication Logic (SIL) to determine the link status to report to the PMA client. Other inputs to the SIL may include status of clock and data recovery on the lanes from the PMD and frame synchronization.

94.2.7 PMA local loopback mode

PMA local loopback shall be provided. This function involves looping back each input lane from the PMA service interface to the corresponding output lane on the PMA service interface. Each received instance of the PMA:IS_UNITDATA_*i*.request(tx_bit,start) primitive is looped back in the direction of the PMA client using the PMA:IS_UNITDATA_*i*.indication(rx_bit,start) primitive.

During local loopback, the PMA performs normal framing and precoding onto the lanes in the Tx direction toward the PMD service interface.

Ability to perform this function is indicated by the Local_loopback_ability status variable. The Local_loopback_ability status variable is always set to 1. If a Clause 45 MDIO is implemented, this variable is accessible through bit 1.8.0 (45.2.1.7.15). A device is placed in local loopback mode when the

Local_loopback_enable control variable is set to one, and removed from local loopback mode when this variable is set to zero. If a Clause 45 MDIO is implemented, this variable is accessible through PMA/PMD control 1 register (bit 1.0.0, see 45.2.1.1.5).

94.2.8 PMA remote loopback mode (optional)

PMA remote loopback mode is optional. If implemented, it shall be as described in this subclause.

Remote loopback, if provided, should be implemented close enough to the PMD to maintain the bit sequence on each individual PMD lane. When remote loopback is enabled, each bit received over a lane of the service interface below the PMA via PMD:IS_UNITDATA_*i*.indication is looped back to the corresponding output lane toward the PMD via PMD:IS_UNITDATA_*i*.request.

During remote loopback, the PMA performs normal bit processing operation in the Rx direction towards the PMA client.

The ability to perform this function is indicated by the Remote_loopback_ability status variable. If a Clause 45 MDIO is implemented, this variable is accessible through bit 1.13.15 (Clause 45.2.1.2.3). A device is placed in remote loopback mode when the Remote_loopback_enable control variable is set to one, and removed from remote loopback mode when this variable is set to zero. If a Clause 45 MDIO is implemented, this variable is accessible through PMA/PMD Control register 1 (bit 1.0.1, see Clause 45.2.1.1.4).

94.2.9 PMA test patterns

94.2.9.1 JP03A test pattern

A 100GBASE-KP4 PMA shall include a JP03A test pattern generator as specified in this subclause.

The JP03A test pattern is generated prior to PAM4 encoding. When the JP03A test pattern is enabled, it replaces the signal from the precoder. The JP03A test pattern is a repeating {0,3} sequence.

The JP03A test pattern is enabled by the test_pattern_enable and JP03A_enable control variables. If the optional Clause 45 MDIO is implemented, the control variables map to the registers and bits defined in 94.2.10.

94.2.9.2 JP03B test pattern

A 100GBASE-KP4 PMA shall include a JP03B test pattern generator as specified in this subclause.

The JP03B test pattern is generated prior to PAM4 encoding. When the JP03B test pattern is enabled, it replaces the signal from the precoder. The JP03B test pattern is a repeating sequence of {0,3} repeated 15 times followed by {3,0} repeated 16 times. The entire 62-symbol pattern is shown in Equation (94–1).

$$03030303030303030303030303033030303030303030303030303030303030 \quad (94–1)$$

The JP03B test pattern is enabled by the test_pattern_enable and JP03B_enable control variables. If the optional Clause 45 MDIO is implemented, the control variables map to the registers and bits defined in 94.2.10.

94.2.9.3 Quaternary PRBS13 test pattern

A 100GBASE-KP4 PMA shall include a quaternary PRBS13 (QPRBS13) pattern generator as specified in this subclause.

The QPRBS13 test pattern is a repeating 15548-symbol (338 training frame words) sequence equivalent to the training pattern specified in 94.3.10.8.

The PRBS13 pattern generator is re-initialized for each repetition of QPRBS13 with the same seeds specified in Table 94–11.

The QPRBS13 test pattern is enabled by the test_pattern_enable and QPRBS13_enable control variables. If the optional Clause 45 MDIO is implemented, the control variables map to the registers and bits defined in 94.2.10.

94.2.9.4 Transmitter linearity test pattern

A 100GBASE-KP4 PMA shall include a transmitter linearity test pattern generator as specified in this subclause.

The transmitter linearity test pattern is a repeating 160-symbol pattern with a sequence of 10 symbol values each 16 UI in duration. The 10 values correspond to the set of PAM4 symbols shown in Equation (94–2).

$$\{-1,-1/3,+1/3,+1,-1,+1,-1,+1,+1/3,-1/3\} \quad (94–2)$$

The transmitter linearity test pattern is enabled by the test_pattern_enable and TX_linearity_enable control variables. If the optional Clause 45 MDIO is implemented, the control variables map to the registers and bits defined in 94.2.10.

94.2.10 PMA MDIO function mapping

Clause 45 specifies the optional MDIO capability that describes several registers that provide control and status information for and about the PMA. 45.2.1 describes the Management Data Input/Output (MDIO) Manageable Device (MMD) addresses. If MDIO is implemented, it shall map MDIO control bits to PMA control variables as shown in Table 94–4, and MDIO status bits to PMA status variables as shown in Table 94–5.

Table 94–4—100GBASE-KP4 MDIO/PMA control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMA control variable
PMA local loopback	PMA/PMD control 1	1.0.0	Local_loopback_enable
PMA remote loopback	PMA/PMD status 2	1.0.1	Remote_loopback_enable
PMA Tx generator enable	PRBS pattern testing control	1.1501.3	test_pattern_enable
JP03A pattern enable	PRBS pattern testing control	1.1501.8	JP03A_enable
JP03B pattern enable	PRBS pattern testing control	1.1501.9	JP03B_enable
QPRBS13 pattern enable	PRBS pattern testing control	1.1501.10	QPRBS13_enable
TX linearity pattern enable	PRBS pattern testing control	1.1501.11	TX_linearity_enable
PMA transmit overhead pattern	PMA overhead control 1	1.162.7:0	TX_OH_pattern
PMA transmit overhead sequence 0	PMA overhead control 1	1.162.12:8	TX_OH_sequence_0
PMA transmit overhead sequence 1	PMA overhead control 2	1.163.4:0	TX_OH_sequence_1

Table 94–4—100GBASE-KP4 MDIO/PMA control variable mapping *(continued)*

MDIO control variable	PMA/PMD register name	Register/bit number	PMA control variable
PMA transmit overhead sequence 2	PMA overhead control 2	1.163.9:5	TX_OH_sequence_2
PMA transmit overhead sequence 3	PMA overhead control 2	1.163.14:10	TX_OH_sequence_3
PMA receive overhead pattern	PMA overhead control 3	1.164.7:0	RX_OH_pattern

Table 94–5—100GBASE-KP4 MDIO/PMA status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMA status variable
PMA local loopback ability	PMA/PMD status 2	1.8.0	Local_loopback_ability
PMA remote loopback ability	40G/100G PMA/PMD extended ability	1.13.15	Remote_loopback_ability
PMA receive overhead sequence 0	PMA overhead status 1	1.165.5:0	RX_OH_sequence_0
PMA receive overhead sequence 1	PMA overhead status 1	1.165.11:6	RX_OH_sequence_1
PMA receive overhead sequence 2	PMA overhead status 2	1.166.5:0	RX_OH_sequence_2
PMA receive overhead sequence 3	PMA overhead status 2	1.166.11:6	RX_OH_sequence_3

94.3 Physical Medium Dependent (PMD) Sublayer

94.3.1 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 100GBASE-KP4 PMD. The service interface for this PMD is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data. The PMD translates the encoded data to and from signals suitable for the medium.

The PMD service interface is based on the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_*i*.request
PMD:IS_UNITDATA_*i*.indication
PMD:IS_SIGNAL.indication

If the optional EEE deep sleep capability is supported, then the PMD service interface includes two additional primitives as follows:

PMD:IS_TX_MODE.request
PMD:IS_RX_MODE.request

94.3.1.1 PMD:IS_UNITDATA_i.request

The PMD:IS_UNITDATA_*i*.request (where *i*=0 to 3) primitive is used to define the transfer of four streams of data units from the PMA to the PMD.

94.3.1.1.1 Semantics of the service primitive

PMD:IS_UNITDATA_0.request(tx_symbol)
PMD:IS_UNITDATA_1.request(tx_symbol)
PMD:IS_UNITDATA_2.request(tx_symbol)
PMD:IS_UNITDATA_3.request(tx_symbol)

The data conveyed by PMD:IS_UNITDATA_i.request consists of four parallel continuous streams of encoded symbols, tx_symbol, one stream for each lane. Each of the tx_symbol parameters can take one of four values: –1, –1/3, +1/3, or +1.

94.3.1.1.2 When generated

The PMA continuously sends four parallel symbol streams PMD:IS_UNITDATA_*i*.request(tx_symbol) to the PMD, each at a nominal signaling rate of 13.59375 GBd.

94.3.1.1.3 Effect of receipt

Upon receiving each instance of PMD:IS_UNITDATA_*i*.request, the tx_symbol parameter is passed to the PMD transmit process corresponding to each stream.

94.3.1.2 PMD:IS_UNITDATA_i.indication

The PMD:IS_UNITDATA_*i*.indication (where *i*=0 to 3) primitive is used to define the transfer of four streams of data units from the PMD to the PMA.

94.3.1.2.1 Semantics of the service primitive

PMD:IS_UNITDATA_0.indication(rx_symbol)
PMD:IS_UNITDATA_1.indication(rx_symbol)
PMD:IS_UNITDATA_2.indication(rx_symbol)
PMD:IS_UNITDATA_3.indication(rx_symbol)

The data conveyed by PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_3.indication consists of four parallel continuous streams of encoded symbols, one stream for each lane. Each of the rx_symbol parameters can take one of four values: –1, –1/3, +1/3, or +1.

94.3.1.2.2 When generated

The PMD continuously sends four parallel encoded symbol streams PMD:IS_UNITDATA_*i*.indication(rx_symbol) to the PMD client, each at a nominal signaling rate of 13.59375 GBd.

94.3.1.2.3 Effect of receipt

The effect of receipt of this primitive is undefined by the PMD.

94.3.1.3 PMD:IS_SIGNAL.indication

The PMD:IS_SIGNAL.indication primitive is generated by the PMD to the PMA to indicate the status of the PMD receive process. This primitive is generated by the PMD receive process to propagate the detection of severe error conditions (e.g., loss of synchronization) to the PMA.

94.3.1.3.1 Semantics of the service primitive

PMD:IS_SIGNAL.indication(SIGNAL_OK).

The SIGNAL_OK parameter corresponds to the variable Global_PMD_signal_detect as defined in 94.3.6.4. When Global_PMD_signal_detect is one, SIGNAL_OK shall be assigned the value OK. When Global_PMD_signal_detect is zero, SIGNAL_OK shall be assigned the value FAIL. When SIGNAL_DETECT = FAIL, the PMD:IS_UNITDATA_*i*.indication parameters are undefined.

94.3.1.3.2 When generated

The PMD generates the PMD:IS_SIGNAL.indication primitive to the PMD client whenever there is change in the value of the Global_PMD_signal_detect parameter.

94.3.1.3.3 Effect of receipt

The effect of receipt of this primitive is undefined by the PMD.

94.3.2 PCS requirements for Auto-Negotiation (AN) service interface

The PCS associated with this PMD is required to support the AN service interface primitive AN_LINK.indication defined in 73.9. (See 82.6.)

The 100GBASE-KP4 PHY may be extended using CAUI-n as a physical instantiation of the inter-sublayer service interface between devices. If CAUI-n is instantiated, the AN_LINK(link_status).indication is relayed from the device with the PCS sublayer to the device with the AN sublayer by means at the discretion of the implementer. As examples, the implementer may employ use of pervasive management or employ a dedicated electrical signal to relay the state of link_status as indicated by the PCS sublayer on one device to the AN sublayer on the other device.

94.3.3 Delay constraints

The sum of the transmit and the receive delays contributed by the 100GBASE-KP4 PMA, PMD, AN, and the medium in one direction shall be no more than 8192 bit times (16 pause_quanta or 81.92 ns). It is assumed that the one way delay through the medium is no more than 800 bit times (8 ns).

A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

94.3.4 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the RS-FEC sublayer. Skew and Skew Variation are defined in 80.5 and specified at the points SP1 to SP6 shown in Figure 80–8.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 0.4 ns.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 0.6 ns.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 80.5.

94.3.5 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several registers that provide control and status information for and about the PMD. If MDIO is implemented, it shall map MDIO control bits to PMD control variables as shown in Table 94–6, and MDIO status bits to PMD status variables as shown in Table 94–7.

Table 94–6—100GBASE-KP4 MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable	1.9.0	Global_PMD_transmit_disable
PMD transmit disable 3 to PMD transmit disable 0	PMD transmit disable	1.9.4 to 1.9.1	PMD_transmit_disable_3 to PMD_transmit_disable_0
Restart training	BASE-R PMD control	1.150.0	mr_restart_training
Training enable	BASE-R PMD control	1.150.1	mr_training_enable

Table 94–7—100GBASE-KP4 MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Fault	PMA/PMD status 1	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect	1.10.0	Global_PMD_signal_detect
PMD receive signal detect 3 to PMD receive signal detect 0	PMD receive signal detect	1.10.4 to 1.10.1	PMD_signal_detect_3 to PMD_signal_detect_0
100GBASE-KP4 deep sleep	EEE capability	1.16.9	—
Receiver status 3	BASE-R PMD status	1.151.12	rx_trained_3
Frame lock 3	BASE-R PMD status	1.151.13	frame_lock_3

Table 94–7—100GBASE-KP4 MDIO/PMD status variable mapping *(continued)*

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Start-up protocol status 3	BASE-R PMD status	1.151.14	training_3
Training failure 3	BASE-R PMD status	1.151.15	training_failure_3
Receiver status 2	BASE-R PMD status	1.151.8	rx_trained_2
Frame lock 2	BASE-R PMD status	1.151.9	frame_lock_2
Start-up protocol status 2	BASE-R PMD status	1.151.10	training_2
Training failure 2	BASE-R PMD status	1.151.11	training_failure_2
Receiver status 1	BASE-R PMD status	1.151.4	rx_trained_1
Frame lock 1	BASE-R PMD status	1.151.5	frame_lock_1
Start-up protocol status 1	BASE-R PMD status	1.151.6	training_1
Training failure 1	BASE-R PMD status	1.151.7	training_failure_1
Receiver status 0	BASE-R PMD status	1.151.0	rx_trained_0
Frame lock 0	BASE-R PMD status	1.151.1	frame_lock_0
Start-up protocol status 0	BASE-R PMD status	1.151.2	training_0
Training failure 0	BASE-R PMD status	1.151.3	training_failure_0

94.3.6 PMD functional specifications

94.3.6.1 Link block diagram

One direction for one lane of a 100GBASE-KP4 link is shown in Figure 94–4.

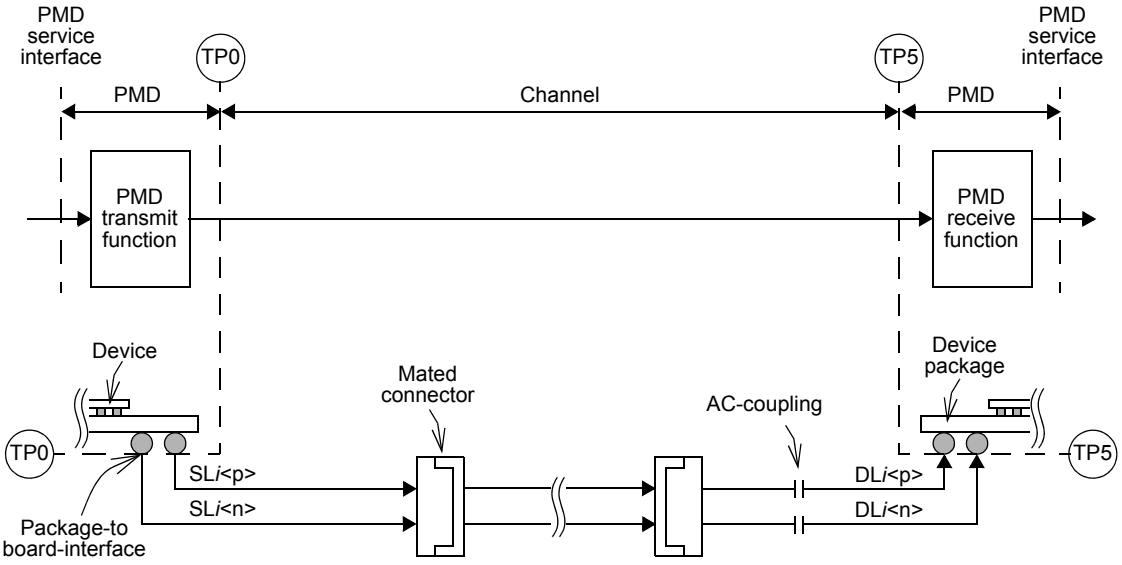

Figure 94–4—100GBASE-KP4 link (one direction for one lane is illustrated)

94.3.6.2 PMD Transmit function

The PMD transmit function shall convert the four encoded symbol streams requested by the PMD service interface messages PMD:IS_UNITDATA_i.request (i=0 to 3) into four separate electrical signals. The four electrical signals shall then be delivered to the MDI, all according to the transmit electrical specifications in 94.3.12. A positive differential output voltage (SL*i*<p> minus SL*i*<n>) shall correspond to a positive tx_symbol value.

If the optional EEE deep sleep capability is supported, the PMD transmit function shall transmit a periodic sequence, where each period of the sequence is an ALERT frame (see 94.3.11.1) when tx_mode is set to ALERT. Regardless of tx_mode, the transmit equalizer coefficients shall be set to the values determined via the start-up protocol (see 94.3.10).

94.3.6.3 PMD Receive function

The PMD receive function shall convert the four electrical signals from the MDI into four encoded symbol streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_i.indication (i=0 to 3). A positive differential input voltage of (DL*i*<p> minus DL*i*<n>) shall correspond to a positive rx_symbol value.

94.3.6.4 Global PMD signal detect function

The variable Global_PMD_signal_detect is the logical AND of the values of PMD_signal_detect_i for i=0 to 3.

When the MDIO is implemented, this function maps the variable Global_PMD_signal_detect to the register and bit specified in 94.3.5.

94.3.6.5 PMD lane-by-lane signal detect function

The PMD lane-by-lane signal detect function is used by the 100GBASE-KP4 PMD to indicate the successful completion of the start-up protocol by the PMD control function (see 94.3.10). PMD_signal_detect_i (where i represents the lane number in the range 0 to 3) is set to zero when the value of the variable signal_detect is set to false by the Training state diagram for lane i (see Figure 72–5). PMD_signal_detect_i is set to one when the value of signal_detect for lane i is set to true.

If training is disabled by the management variable mr_training_enable (see 94.3.5), PMD_signal_detect_i shall be set to one for i=0 to 3.

If the optional EEE deep sleep capability is supported, the following requirements apply. The value of PMD_signal_detect_i (for i=0 to 3) is set to zero when rx_mode is first set to QUIET. While rx_mode is set to QUIET, PMD_signal_detect_*i* shall be set to one within 500 ns of the application of the ALERT pattern defined in 94.3.6.2 and meeting the EEE transmit-enabled amplitude requirement of 94.3.12.3. While rx_mode is set to QUIET, PMD_signal_detect_*i* shall not be set to one when the output of the transmitter on the same lane meets the EEE transmit-disabled amplitude requirement of 94.3.12.3.

When the MDIO is implemented, this function maps the variables to registers and bits as defined in 94.3.5.

94.3.6.6 Global PMD transmit disable function

The Global PMD transmit disable function is mandatory if EEE deep sleep capability is supported and is otherwise optional. When implemented, it allows all of the transmitters to be disabled with a single variable.

a) When 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 94–13.

b) If a PMD fault (94.3.7) is detected, then the PMD may set Global_PMD_transmit_disable to one.

c) Loopback, as defined in 94.3.6.8, shall not be affected by Global_PMD_transmit_disable.

d) The following additional requirements apply when the optional EEE deep sleep capability is supported. The Global PMD transmit disable function shall turn off all of the transmitters as specified in 94.3.12.3 when tx_mode transitions to QUIET from any other value. The Global PMD transmit disable function shall turn on all of the transmitters as specified in 94.3.12.3 when tx_mode transitions from QUIET to any other value.

94.3.6.7 PMD lane-by-lane transmit disable function

The PMD lane-by-lane transmit disable function is optional and allows the electrical transmitter in each lane to be selectively disabled. When this function is supported, it shall meet the following requirements:

a) When a PMD_transmit_disable_*i* variable (where i represents the lane number in the range 0 to 3) is set to one, this function shall turn off the transmitter associated with that variable such that it drives a constant level (i.e., no transitions) and does not exceed the maximum differential peak-to-peak output voltage specified in Table 94–13.

b) If a PMD_fault (94.3.7) is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the electrical transmitter in each lane.

c) Loopback, as defined in 94.3.6.8, shall not be affected by PMD_transmit_disable_*i*.

94.3.6.8 Loopback mode

Local loopback mode is provided by the PMA (94.2.7). Loopback shall not affect the state of the transmitter, which continues to send data unless disabled (94.3.6.7).

NOTE—Placing a network port into loopback mode can be disruptive to a network.

94.3.7 PMD fault function

PMD_fault is the logical OR of PMD_receive_fault, PMD_transmit_fault, and any other implementation specific fault.

If the MDIO is implemented, PMD_fault shall be mapped to the fault bit as specified in 45.2.1.2.3.

94.3.8 PMD transmit fault function

The PMD transmit fault function is optional. The faults detected by this function are implementation specific, but the assertion of Global_PMD_transmit_disable is not considered a transmit fault. A fault is indicated by setting the variable PMD_transmit_fault to one.

If PMD_transmit_fault is asserted, then Global_PMD_transmit_disable should also be asserted.

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.

94.3.9 PMD receive fault function

The PMD receive fault function is optional. The faults detected by this function are implementation specific. A fault is indicated by setting the variable PMD_receive_fault to one.

If the MDIO interface is implemented, then PMD_receive_fault shall contribute to the Receive fault bit as specified in 45.2.1.7.5.

94.3.10 PMD control function

94.3.10.1 Overview

The PMD control function generates the control actions required to bring the PMD from initialization to a mode in which data may be exchanged with the link partner.

The PMD control function is based upon the 10GBASE-KR start-up protocol. This protocol facilitates timing recovery and equalization while also providing a mechanism through which the receiver can tune the transmit equalizer to optimize performance over the backplane interconnect. The protocol supports these mechanisms through the continuous exchange of fixed-length training frames.

Each lane of the 100GBASE-KP4 PMD shall have an independent control function as defined in this subclause.

The variables rx_trained_*i*, frame_lock_*i*, training_*i,* and training_failure_*i* (where *i* goes from 0 to 3) report status for each lane and are equivalent to rx_trained, frame_lock, training, and training_failure as defined in 72.6.10.3.1. If the MDIO interface is implemented, then this function shall map these variables to the appropriate bits in the BASE-R PMD status register (Register 1.151) as specified in 45.2.1.81.

94.3.10.2 Training frame structure

The training frame is a fixed length structure that is sent continuously during training. The training frame, shown in Figure 94–5, is 348 training frame words (94.3.10.3) in length and contains a frame marker, a control channel, and training pattern. The frame marker delimits the beginning of a training frame. The control channel provides a means for the each receiver to control the taps on the link partner transmitter and communicate status. The training pattern provides content rich pattern for receiver convergence.

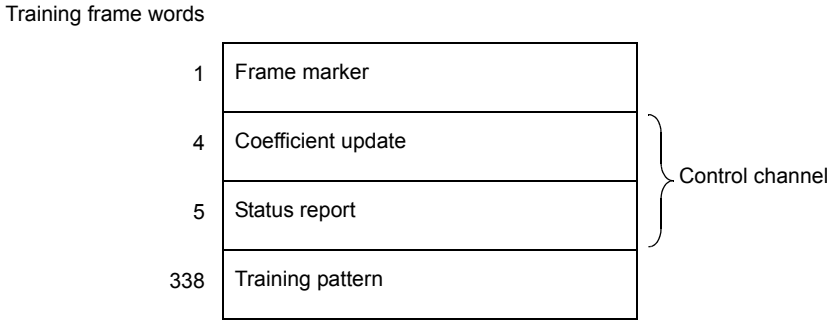

Figure 94–5—Training frame structure

94.3.10.3 Training frame words

Each training frame is composed of a series of 348 training frame words. Each training frame word is 46 symbols in length, equivalent in size to a termination block described in 94.2.2.4.

94.3.10.4 Frame marker

Each training frame shall be delimited by a frame marker as described in this subclause. The frame marker is a training frame word composed of a 46-symbol pattern of 23 +1 symbols followed by 23 –1 symbols. This pattern does not appear in the control channel or the training pattern and therefore serves as a unique indicator of the start of a training frame.

94.3.10.5 Control channel encoding

94.3.10.5.1 Differential Manchester encoding

The control channel shall be transmitted using differential Manchester encoding (DME). DME guarantees transition density and DC balance while the reduced rate of transmission facilitates reception over non-optimally equalized channels.

DME cells shall be encoded using the following rules:

a) Each DME cell represents one bit of information.
b) The upper value is represented by a series of PAM4 +1 symbols.
c) The lower value is represented by a series of PAM4 –1 symbols.
d) A data transition occurs at each cell boundary.
e) A mid-cell data transition is used to signal a logical one.
f) The absence of a mid-cell data transition is used to signal a logical zero.

If a coding violation is detected within the bounds of the control channel in a given training frame, the contents of the control channel for that frame shall be ignored.

94.3.10.5.2 Control channel structure

The control channel shall be constructed of a series of DME cells as described in this subclause.

The control channel is composed of a series of 9 training frame words. Each training frame word is composed of four 10-symbol control channel DME cells and a 6-symbol control overhead DME cell.

The control overhead cell is always transmitted as a one following the DME rules. In other words, the control overhead cell is transmitted as either three +1 symbols followed by three –1 symbols or vice versa depending on the previously transmitted control channel cell.

The coefficient update field is transmitted in the first 16 control channel DME cells. The status report field is transmitted in the next 24 control channel DME cells. The structure of the frame marker and control channel are shown in Table 94–8.

94.3.10.6 Coefficient update field

The coefficient update field carries correction information from the local receiver to the link partner transmit equalizer. The field consists of preset controls, initialization controls, coefficient updates for three transmit equalizer taps, and parity. The coefficient update field is mapped into the first 16 control channel DME

Table 94–8—Frame marker and control channel structure

Training frame word	Symbols 1:10	Symbols 11:20	Symbols 21:30	Symbols 31:40	Symbols 41:46	Training frame fields
1	cell 15	cell 14	cell 13	cell 12	overhead	coefficient update
2	cell 11	cell 10	cell 9	cell 8	overhead	
3	cell 7	cell 6	cell 5	cell 4	overhead	
4	cell 3	cell 2	cell 1	cell 0	overhead	
5	cell 19	cell 18	cell 17	cell 16	overhead	status report
6	cell 15	cell 14	cell 13	cell 12	overhead	
7	cell 11	cell 10	cell 9	cell 8	overhead	
8	cell 7	cell 6	cell 5	cell 4	overhead	
9	cell 3	cell 2	cell 1	cell 0	overhead	

The format of the coefficient update field shall be as shown in Table 94–9. Cell 15 of the coefficient update field shall be transmitted first. The preset, initialize, and coefficient update fields are set by the receiver adaptation process. The algorithm employed by the receiver adaptation process is beyond the scope of this standard.

Table 94–9—Coefficient update field

Cell(s)	Name	Description
15:14	Reserved	Transmitted as 0, ignored on reception.
13	Preset	1 = Preset coefficients 0 = Normal operation
12	Initialize	1 = Initialize coefficients 0 = Normal operation
11:7	Reserved	Transmitted as 0, ignored on reception.
6	Parity	Even parity of all other coefficient update cells.
5:4	Coefficient (+1) update	5 4 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold
3:2	Coefficient (0) update	3 2 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold
1:0	Coefficient (–1) update	1 0 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold

94.3.10.6.1 Preset

The behavior in response to the preset field shall be as specified in 72.6.10.2.3.1.

94.3.10.6.2 Initialize

The behavior in response to the initialize field shall be as specified in 72.6.10.2.3.2, except that the conditions for the INITIALIZE state are specified in 94.3.12.5.4 instead of 72.6.10.4.2.

94.3.10.6.3 Parity

The parity bit shall be set so that all bits in the coefficient update field including the parity bit exhibit even parity. The parity bit protects against acceptance of errored control messages and preserves DC balance. If a parity violation is detected within a received coefficient update field, that field shall not be used to update the transmitter coefficients.

94.3.10.6.4 Coefficient (*k*) update

The coefficient update fields shall be encoded as described in this subclause.

Each coefficient, identified by the index k, is assigned a 2-bit field describing a requested update, where $k \in \{-1, 0, 1\}$ denoting the pre-cursor, main, and post-cursor taps, respectively. The encoding of the coefficient update fields are as shown in Table 94–9.

Three request encodings are defined: increment, decrement, and hold. The default state of each tap is hold, which corresponds to no change in the coefficient. The increment or decrement encodings are transmitted to request that the corresponding coefficient be increased or decreased. The amount of change implemented by the transmitter in response to the coefficient update request meets the requirements of 94.3.12.5.5. An increment or decrement request is transmitted continuously until the update status (94.3.10.7.4) for that tap indicates updated, maximum, or minimum. At that point, the outgoing requests for that tap may be set to hold. The hold setting must be maintained until the incoming status message for that tap reverts to not_updated. A new request to increment or decrement a tap may be sent only when the incoming status message for that tap is not_updated.

Coefficient increment and decrement update requests must not be sent in combination with initialize or preset.

94.3.10.7 Status report field

The status report field is used to signal state information from the local PMD to the link partner. The format of the status report field of training frames shall be as shown in Table 94–10. Cell 19 of the status report field shall be transmitted first.

Table 94–10—Status report field for training and alert frames

Cell(s)	Name	Description
19	Parity	Set to achieve even parity for status report field.
18	Mode	0: Training 1: EEE
17:16	Frame countdown	Number of frames remaining before transition to data mode.

Table 94–10—Status report field for training and alert frames *(continued)*

Cell(s)	Name	Description
15:13	PMA alignment offset	Relative location of the next alert frame within the PMA frame (set to zero for training frames).
12:7	Reserved	Transmitted as zeros.
6	Receiver ready	1 = The local receiver has determined that training is complete and is prepared to receive data. 0 = The local receiver is requesting that training continue.
5:4	Coefficient (+1) status	5 4 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated (and for EEE alert frames)
3:2	Coefficient (0) status	3 2 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated (and for EEE alert frames)
1:0	Coefficient (–1) status	1 0 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated (and for EEE alert frames)

94.3.10.7.1 Parity

The parity cell shall be set so that all bits in the status report field including the parity bit exhibit even parity. The parity bit protects against acceptance of errored status messages and preserves DC balance. If a parity violation is detected within a received status field, that field shall not be used to determine the link partner status.

94.3.10.7.2 Training frame countdown

The training frame countdown cell shall signal the transition from training to data mode as described in this subclause. When training begins, countdown is set to the value 3 and remains so until all receivers have completed training. When the received status report receiver ready is 1 in all four received lanes and the transmitted status report receiver ready is 1 in all four transmitted lanes, the transmitter on each transmitted lane decrements the countdown in three successive frames. The countdown values are equal in all four lanes. In other words, in the last three training frames countdown contains 2, 1, and 0, respectively. Immediately after the last training frame word of the last training frame is sent, transmission of the PMA frame begins starting with the termination block containing the PMA overhead (see 94.3.10.9).

94.3.10.7.3 Receiver ready

The receiver ready bit shall signal the local receiver state to the link partner as described in this subclause. When training begins the receiver ready bit is deasserted and remains so until the receiver has concluded training. The receiver ready bit is asserted to indicate that the local receiver has concluded training and is prepared to receive data. The encoding of the receiver ready bit is as shown in Table 94–10.

94.3.10.7.4 Coefficient (*k*) status

The behavior of the coefficient (*k*) status fields shall be as specified in 72.6.10.2.4.5.

94.3.10.7.5 Coefficient update process

The coefficient update process shall behave as specified in 72.6.10.2.5.

In addition, the period from receiving a new request to responding to that request shall be less than 2 ms, except during the first 50 ms following the beginning of the start-up protocol. The beginning of the start-up protocol is defined to be entry into the AN_GOOD_CHECK state in Figure 73–10. The start of the period is the frame marker of the training frame with the new request and the end of the period is the frame marker of the training frame with the corresponding response. A new request occurs when the coefficient update field is different from the coefficient field in the preceding frame. The response occurs when the coefficient status report field is updated to indicate that the corresponding action is complete.

94.3.10.8 Training pattern

The training pattern shall be encoded as specified in this subclause.

The training pattern is mapped into a series of 338 training frame words. Each training frame words is encoded as a PMA signal as specified in 94.2.2 with the exception that the input is from a PRBS13 generator rather than from the PMA service interface and no PMA overhead (94.2.2.2) is inserted.

For each training frame, the PRBS13 generator is used to produce 31096 bits. Three full cycles of 8191 bits and one truncated cycle of 6523 bits are concatenated to form the 31096 bit sequence, R(1:31096). Bits in the first and third cycle, R(1:8191) and R(16383:24573), are not inverted and bits in the second and fourth cycles, R(8192:16382) and R(24574:31096), are inverted.

The PRBS13 pattern generator produces the same result as the implementation shown in Figure 94–6, which implements the generator polynomial shown in Equation (94–3). The PRBS13 pattern generator is initialized for each frame using a unique seed for each lane. The 13-bit seed and the initial 16 bits for each lane are annotated in Table 94–11.

$$G(x) = 1 + x + x^2 + x^{12} + x^{13} \tag{94–3}$$

The PRBS13 pattern is mapped into the 92 bits of each training frame word. The first 2 bits of each training frame word form the termination bits (94.2.2.4) and each training frame word in the training pattern is equivalent to a termination block (94.2.2.4). The resulting termination blocks are gray-mapped (94.2.2.5), precoded (94.2.2.6), and mapped to PAM4 levels (94.2.2.7).

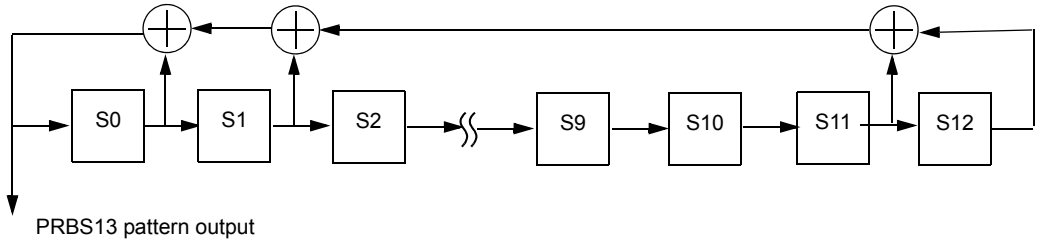

Figure 94–6—PRBS13 pattern generator

The outputs of PRBS13 generator, gray mapper, and precoder for the first two training frame words are provided in Table 94–12.

Table 94–11—PRBS13 seeds and initial output

PMD Lane	Seed bits (leftmost bit in S0, rightmost in S12)	Initial 16 bits (in order of transmission)
0	0000010101011	0100100110110011
1	0011101000001	1101111101010100
2	1001000101100	1100101111000011
3	0100010000010	0110111101000111

Table 94–12—Training pattern initial sequences

PMD Lane	Output of	Contents of first (top) and second (bottom) training frame words transmitted left to right
0	PRBS13	0100100110110011110001010101100001001001110111100111010000011 1010011011101001110011001010111 0001111111010111011011111101000101101111101001111011001010110 0111001001110000111100001101011
	Gray code	1031320220111130103121231210012102121023131112 0122211213222101132233123203320231023012301332
	Precoder	1301200200101031003201123322233220110021032320 0111101103333223211121021130331123112233001211
1	PRBS13	1101111101010100000010010011011001111000101010110000100100111 0111100111010000011101001101110 1001110011001010111000111111101011101101111110100010110111110 1001111011001010110011100100111
	Gray code	2122111000310213123033320031023220233002331323 3120203323022233232122330321221022131113120312
	Precoder	2333232222100230112212113123112022030002123021 3200221203111121120111213023332202301012331233
2	PRBS13	1100101111000011111011101110110011001100111000111000110000110 0001110111000001100110000001110 0011011000011000101011000110010011101010100011001001000011100 1111011101010110011001001010101
	Gray code	2032200223232320202023023020020023230020200023 0213013033201310233330203100231232333202031111
	Precoder	2211131112033022002203112200022203300022000021 0230012212001231121213312313301120303311301010
3	PRBS13	0110111101000111101111101011001101111110001111011010111011110 0100001011000101100101011111000 1001011010111100101001011000100111110000101011010110010011111 1000101011011101001000101111100
	Gray code	1322101232233202122302213323220301130320332230 3113322033113031220033211310222011132331011220
	Precoder	1202310211121133202133321203331223213022120213 3230333121012210200030232100202232302123101113

94.3.10.9 Transition from training to data

The transmitted signal shall transition from the training signal to normal data as described in this subclause.

The transition from the training signal to normal data occurs when the training countdown is complete, as indicated by the training frame countdown being equal to 0 (94.3.10.7.2). Immediately after the last bit of the last training frame, transmission of the first PMA frame (94.2.2.4) begins with the termination block containing the overhead, T(i,1,1:92). The PRBS13 generator used during training to generate the training pattern is used to generate the termination bits in data mode. The state of the training PRBS13 generator is retained and 92 new bits are generated without reseeding or inverting. Termination bits are assigned and the PRBS13 generator continues to operate as specified in 94.2.2.4. The transition from training to data mode and mapping of the PRBS13 to training frame and termination bits is depicted in Figure 94–7.

94.3.10.10 Frame lock state diagram

The 10GBASE-KP4 PMD shall implement the frame lock state diagram as depicted in Figure 72–4 including compliance with the associated state variables, timers, counters, and functions specified in 72.6.10.3. The frame lock state diagram determines when the PMD control function has detected the frame boundaries in the received data stream.

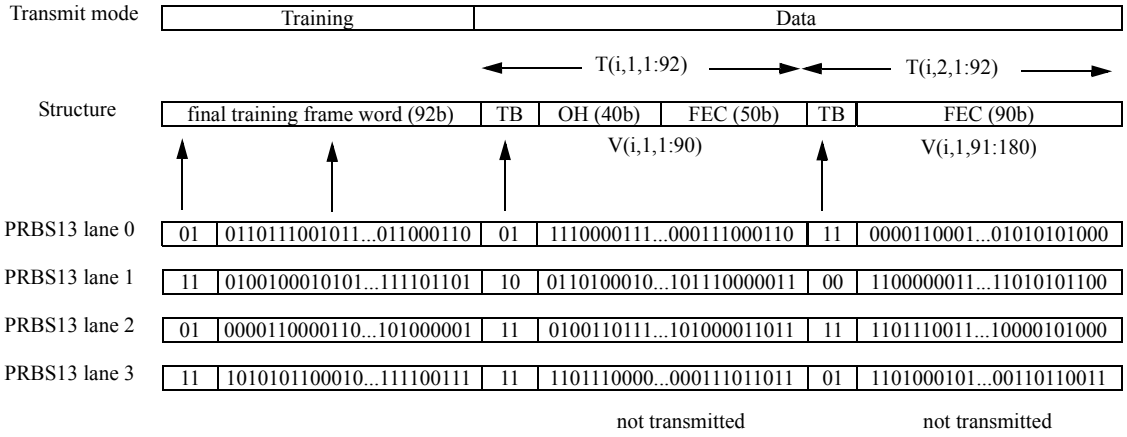

Figure 94–7—Transition from training to data mode

94.3.10.11 Training state diagram

The 10GBASE-KP4 PMD shall implement the training state diagram as depicted in Figure 72–5 including compliance with the associated state variables specified in 72.6.10.3. The training state diagram defines the operation of the 100GBASE-KP4 start-up protocol.

When the training state diagram enters the INITIALIZE state, the transmitter equalizer shall be configured such that the output meets the requirements of 94.3.12.5.4.

94.3.10.12 Coefficient update state diagram

For each tap, the 100GBASE-KP4 PMD shall implement an instance of the coefficient update state diagram in Figure 72–6 including compliance with the associated state variables as specified in 72.6.10.3. The coefficient update state diagram defines the process for updating transmit equalizer coefficients in response to requests from the link partner and also defines the coefficient update status to be reported in outgoing training frames.

94.3.11 PMD LPI function

The PMD LPI function responds to the transitions between Active, Sleep, Quiet, Refresh, and Wake states via the PMD:IS_TX_MODE.request and PMD:IS_RX_MODE.request. Implementation of the function is optional. EEE capabilities and parameters are advertised during the Backplane Auto-negotiation, as described in 45.2.7.13. The transmitter on the local device informs the link partner's receiver when to sleep, refresh, and wake. The local receiver transitions are controlled by the link partner's transmitter and can change independent of the local transmitter states and transitions.

94.3.11.1 Alert Signal

During refresh and wake, to enable effective detection and quick receiver synchronization, an alert frame is sent prior to sending normal PMA frames. The alert signal is a series of repeating alert frames.

The alert frame shall be composed of a frame marker, control channel, and training pattern as depicted in Figure 94–8. The alert frame is based on the training frame specified in 94.3.10.2. The distinguishing differences are that the training pattern is truncated to 48 training frame words (4416 bits) and the coefficient update and status report fields are encoded differently. The alert frame is a total of 58 training frame words in length.

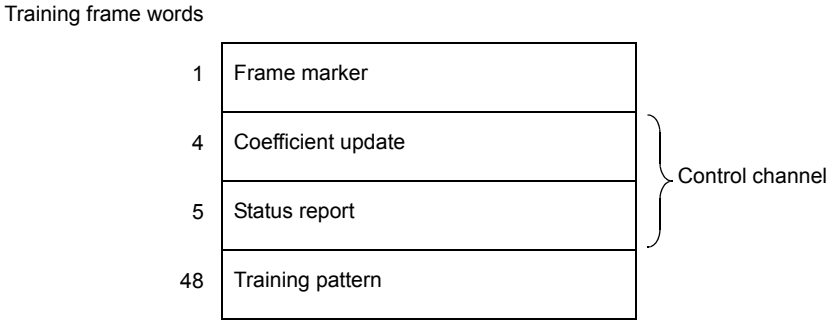

Figure 94–8—Alert frame structure

94.3.11.1.1 Frame marker

The frame marker shall be implemented as specified in 94.3.10.4.

94.3.11.1.2 Coefficient update field

The coefficient update field is unused in the alert frame. All bits in the coefficient update field are reserved and shall be transmitted as zeros.

94.3.11.1.3 Status report field

The status report field is used to signal state information from the local PMD to the link partner. The format of the status report field of alert frames shall be as shown in Table 94–10. Cell 19 of the status report field shall be transmitted first.

94.3.11.1.4 Parity

The parity field shall have the same behavior and purpose as specified for the training frame in 94.3.10.7.1.

94.3.11.1.5 Mode

The mode field indicates whether the frame is a training frame (mode = 0) or an alert frame (mode = 1). This field differentiates the alert frame from a training frame. The mode field in the alert frame shall always indicate 1.

94.3.11.1.6 Alert frame countdown

The alert frame countdown field shall be updated as specified for the training frame in 94.3.10.7.2. The alert frame countdown may be used by the receiver to determine when the signal transitions from the alert frame to the PMA frame (see 94.3.11.1.9).

94.3.11.1.7 PMA alignment offset

The PMA alignment offset (PAO) shall indicate the relative position in the PMA frame in relation to the beginning of the next alert frame as described in this subclause. The PMA alignment offset may be used by the receiver to synchronize to the PMA frame without an additional frame synchronization process (see 94.3.11.1.9). The beginning of the PMA frame is defined as the termination block containing the PMA overhead (94.2.2.2).

The PMA frame length is exactly 6 times the alert frame length. The PMA alignment offset field indicates one of six offsets for the next alert frame within the PMA frame. The offset in number of training frame words of the next alert frame is determined by multiplying PMA alignment offset by 32. As a reference point, a PMA alignment offset of zero indicates that the start of the next alert frame is aligned with the start of a PMA frame. Valid values for the PMA alignment offset are {0,1,2,3,4,5}. The values {6,7} are not valid.

94.3.11.1.8 Receiver ready

The receiver ready cell shall always be set to 1 to indicate training is complete and the link is up.

94.3.11.1.9 Transition from alert to data

The transmitted signal shall transition from the alert signal to normal data as described in this subclause.

The transition from the alert signal to normal data occurs when the alert countdown is complete, as indicated by the alert frame countdown being equal to 0 (94.3.11.1.6). Immediately after the last bit of the last alert frame, transmission of the first PMA frame (94.2.2.4) begins with the termination block indicated by the PMA alignment offset, T(i,32*PAO,1:92). The PRBS13 generator used during alert to generate the training pattern is used to generate the termination bits in data mode. The state of the training PRBS13 generator is retained and 92 new bits are generated without reseeding or inverting. Termination bits are assigned and the PRBS13 generator continues to operate as specified in 94.2.2.4. The transition from alert to data mode and mapping of the PRBS13 to alert frame and termination bits is depicted in Figure 94–9. The values for the PRBS13 in Figure 94–9 are specific to a transition with a PAO of zero. The values are different for other PAO values.

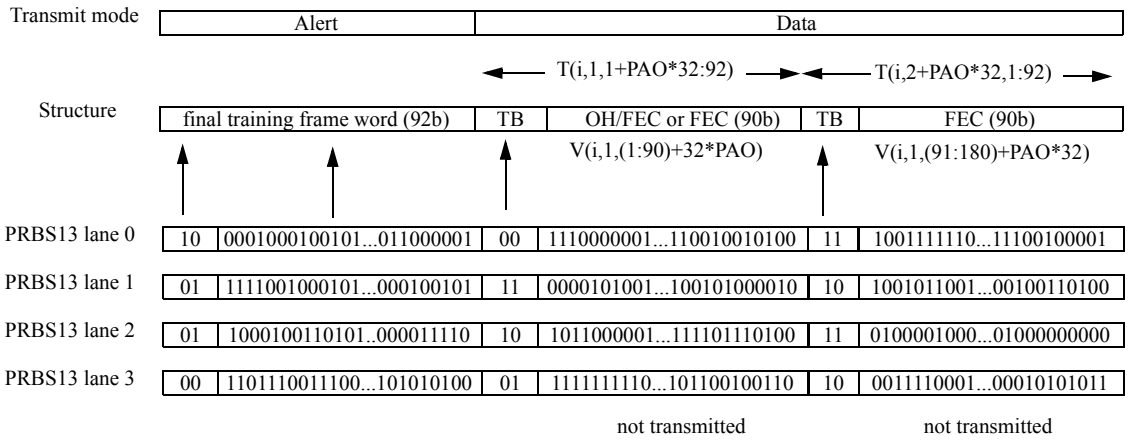

Figure 94–9—Transition from alert to data mode

94.3.12 PMD Transmitter electrical characteristics

Transmitter characteristics measured at TP0a are summarized in Table 94–13.

94.3.12.1 Test fixture

The test fixture of Figure 94–10 or its equivalent is required for measuring the transmitter specifications described in 94.3.12.

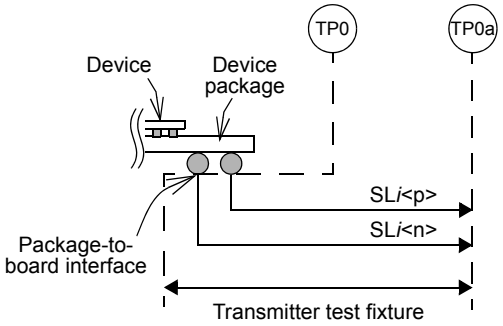

Figure 94–10—Transmitter test fixture and test points

94.3.12.1.1 Test fixture impedance

The differential load impedance applied to the transmitter output by the test fixture depicted in Figure 94–10 shall be 100 Ω.

The differential return loss, in dB with f in GHz, of the test fixture shall meet the requirement of Equation (94–4). The return loss limit $RL_{\min}(f)$ is shown Figure 94–11.

Table 94–13—Summary of transmitter characteristics at TP0a

Parameter	Subclause reference	Value	Units
Signaling rate	94.3.12.2	13.59375 ± 100 ppm	GBd
Differential peak-to-peak output voltage (max.) Transmitter disabled Transmitter enabled	94.3.12.3	 30 1200	 mV mV
DC common-mode output voltage (max.)	94.3.12.3	1.9	V
DC common-mode output voltage (min.)	94.3.12.3	0	V
AC common-mode output voltage (RMS, max.)	94.3.12.3	30	mV
Differential output return loss (min.)	94.3.12.4	Equation (94–7)	dB
Common-mode output return loss (min.)	94.3.12.4	Equation (94–8)	dB
Output waveform Level separation mismatch ratio, R_{LM} (min) Steady-state voltage v_f (max.) Steady-state voltage v_f (min.) Linear fit pulse peak (min.) Normalized coefficient step size (min.) Normalized coefficient step size (max.) Pre-cursor full-scale range (min.) Post-cursor full-scale range (min.)	 94.3.12.5.1 94.3.12.5.3 94.3.12.5.3 94.3.12.5.3 94.3.12.5.5 94.3.12.5.5 94.3.12.5.6 94.3.12.5.6	 0.92 0.6 0.4 0.85 × v_f 0.0083 0.05 1.54 4	 — V V V — — — —
Output jitter and linearity Clock random jitter, RMS (max.) Clock deterministic jitter, peak-to-peak (max.) Even-odd jitter (max.) Signal-to-noise-and-distortion ratio (min.)	 94.3.12.6.1 94.3.12.6.1 94.3.12.6.2 94.3.12.7	 0.005 0.05 0.019 31	 UI UI UI dB

$$RL(f) \geq RL_{\min}(f) = \left\{ \begin{array}{ll} 20 - f & 0.05 \leq f \leq 5 \\ 15 & 5 < f \leq 13 \\ 20.57 - 0.4286f & 13 < f \leq 14 \end{array} \right\} \text{ (dB)} \qquad (94\text{–}4)$$

The common-mode return loss, in dB, with f in GHz, of the test fixture shall meet the requirement of Equation (94–5). The return loss limit $RL_{\min}(f)$ is shown Figure 94–12.

$$RL(f) \geq RL_{\min}(f) = 10 \text{ (dB)} \qquad 0.05 \leq f \leq 14 \qquad (94\text{–}5)$$

94.3.12.1.2 Test fixture insertion loss

The insertion loss of the test fixture measured at 12.89 GHz shall be between 1.2 dB and 1.6 dB.

The insertion loss deviation of the test fixture from 0.05 GHz to 10 GHz shall be less than 0.1 dB.

The reference insertion loss of the test fixture is defined by Equation (94–6), where f is the frequency in GHz, and is shown in Figure 94–13.

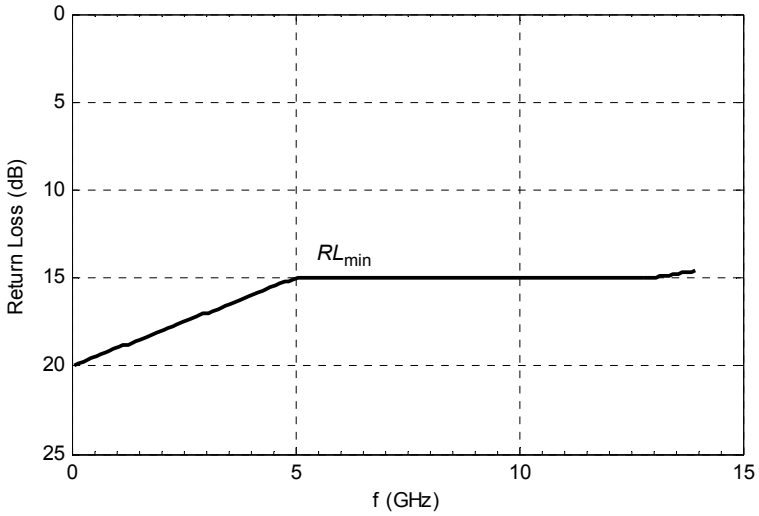

Figure 94–11—Test fixture differential return loss limit

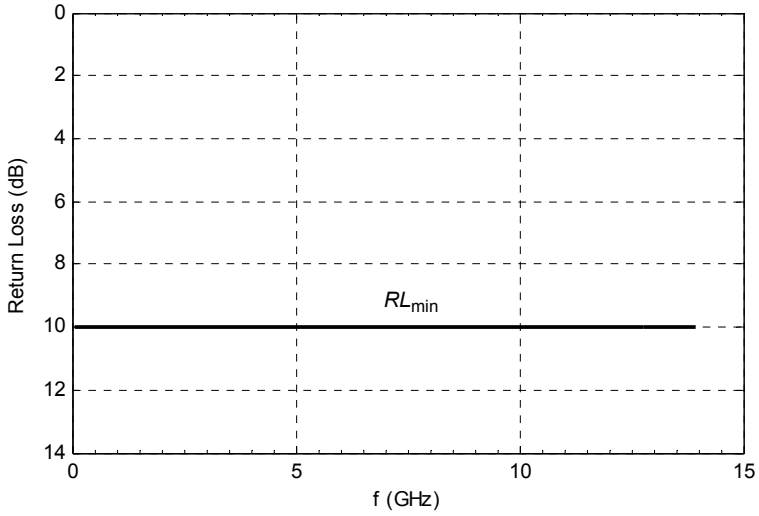

Figure 94–12—Test fixture common-mode return loss limit

$$IL_{ref}(f) = -0.0015 + 0.144\sqrt{f} + 0.069f \ \text{(dB)} \qquad 0.05 \le f \le 14 \tag{94–6}$$

The effects of differences between the insertion loss of an actual test fixture and the reference insertion loss are to be accounted for in the measurements.

A test system with a fourth-order Bessel-Thomson low-pass response with 17 GHz 3 dB bandwidth is to be used for all transmitter signal measurements, unless otherwise specified.

94.3.12.2 Signaling rate and range

The 100GBASE-KP4 signaling rate shall be 13.59375 GBd ± 100 ppm per lane.

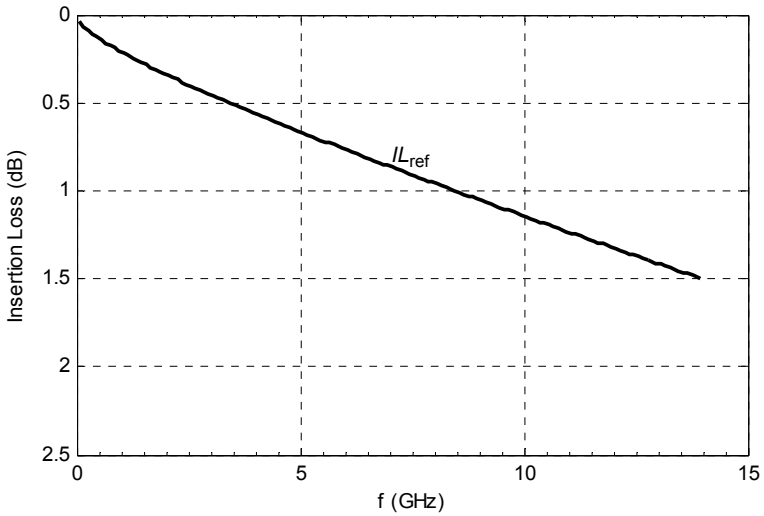

Figure 94–13—Test fixture reference insertion loss

94.3.12.3 Signal levels

The differential output voltage v_{di} is defined to be SL*i*<p> minus SL*i*<n>. The common-mode output voltage v_{cmi} is defined to be one half of the sum of SL*i*<p> and SL*i*<n>. These definitions are illustrated by Figure 94–14.

For a QPRBS13 test pattern (94.2.9.3), the peak-to-peak differential output voltage shall be less than or equal to 1200 mV regardless of the transmit equalizer setting. The peak-to-peak differential output voltage shall be less than or equal to 30 mV while the transmitter is disabled (refer to 94.3.6.6 and 94.3.6.7).

The DC common-mode output voltage shall be between 0 V and 1.9 V with respect to signal ground. The AC common-mode output voltage shall be less than or equal to 30 mV RMS with respect to signal ground. Common-mode output voltage requirements shall be met regardless of the transmit equalizer setting.

If the optional EEE deep sleep capability is supported, the following requirements also apply. The peak-to-peak differential output voltage shall be less than 30 mV within 500 ns of the transmitter being disabled.

When the transmitter transitions from disabled to enabled: (a) The amplitude of the frame marker of the third complete alert frame (see 94.3.11.1) after the transmitter is enabled shall be greater than 90% of the steady-state value (see 94.3.12.5.3), and (b) the transmitter output shall meet the requirements of 94.3.12 within 1 μs of the transmitter being enabled.

While the transmitter is disabled, the DC common-mode output voltage shall be maintained to within ±150 mV of the value for the enabled transmitter.

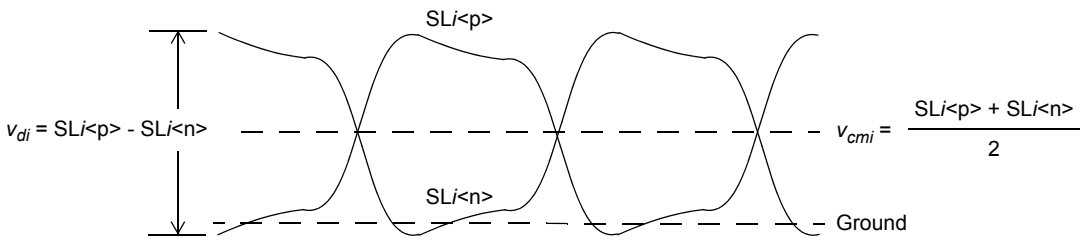

Figure 94–14—Transmitter output voltage definitions

94.3.12.4 Transmitter output return loss

The differential output return loss, in dB, of the transmitter shall meet Equation (94–7), where *f* is the frequency in GHz. The return loss limit $RL_{\min}(f)$ is shown Figure 94–15. This output impedance requirement applies to all valid output levels. The reference impedance for differential return loss measurements shall be 100 Ω

$$RL(f) \geq RL_{\min}(f) = \left\{ \begin{array}{cc} 12.05 - f & 0.05 \leq f \leq 6 \\ 6.5 - 0.075f & 6 < f \leq 10 \end{array} \right\} \quad \text{(dB)} \tag{94–7}$$

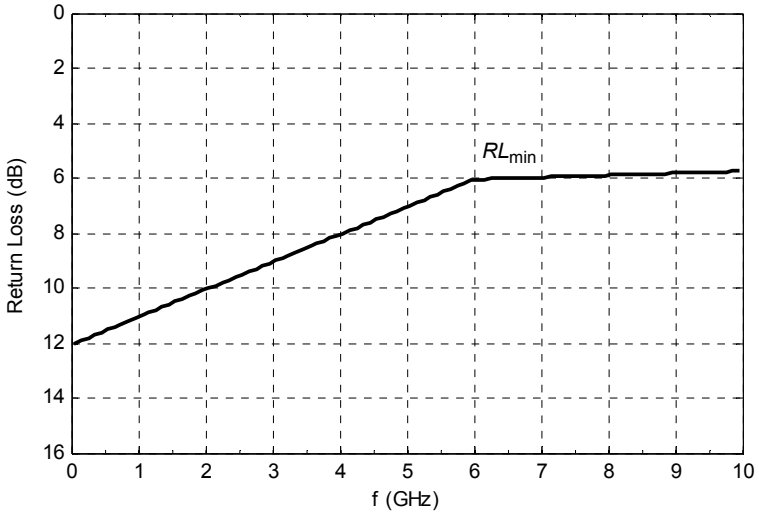

Figure 94–15—Transmitter differential return loss limit

The common-mode output return loss, in dB, of the transmitter shall meet Equation (94–8), where *f* is the frequency in GHz. The return loss limit $RL_{\min}(f)$ is shown Figure 94–16. This output impedance requirement applies to all valid output levels. The reference impedance for common-mode return loss measurements shall be 25 Ω.

$$RL(f) \geq RL_{\min}(f) = \left\{ \begin{array}{cc} 9.05 - f & 0.05 \leq f \leq 6 \\ 3.5 - 0.075f & 6 < f \leq 10 \end{array} \right\} \quad \text{(dB)} \tag{94–8}$$

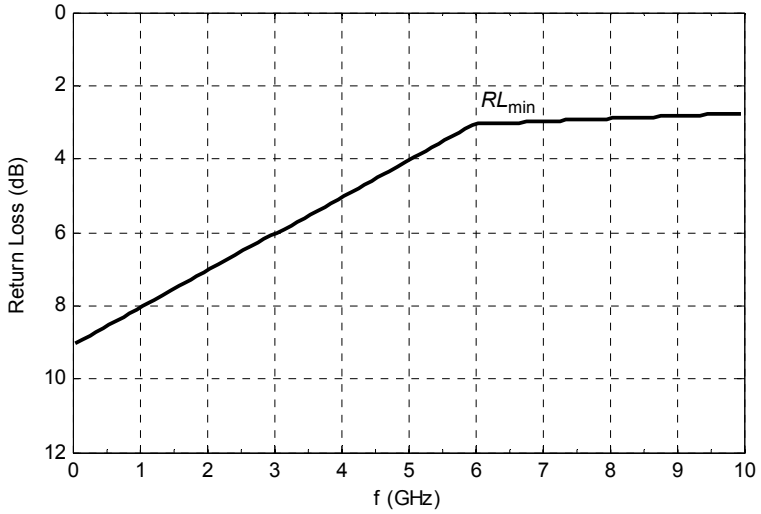

Figure 94–16—Transmitter common-mode return loss limit

94.3.12.5 Transmitter output waveform

The 100GBASE-KP4 transmit function includes programmable equalization to compensate for the frequency-dependent loss of the channel and facilitate data recovery at the receiver. The functional model for the transmit equalizer is the three tap transversal filter shown in Figure 94–17.

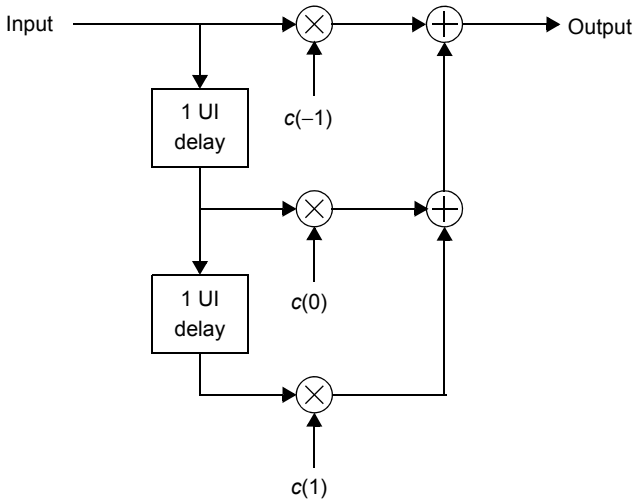

Figure 94–17—Transmit equalizer functional model

The state of the transmit equalizer and hence the transmitted output waveform may be manipulated via the PMD control function defined in 94.3.10 or via the management interface. The transmit function responds to a set of commands issued by the link partner's receive function and conveyed by a back-channel communications path.

This command set includes instructions to:

Increment coefficient $c(i)$.
Decrement coefficient $c(i)$.
Hold coefficient $c(i)$ at its current value.
Set the coefficients to a predefined value (preset or initialize).

In response, the transmit function relays status information to the link partner's receive function. The status messages indicate that:

The requested update to coefficient $c(i)$ has completed (updated).
Coefficient $c(i)$ is at its minimum value.
Coefficient $c(i)$ is at its maximum value.
Coefficient $c(i)$ is ready for the next update request (not_updated).

94.3.12.5.1 Transmitter linearity

Transmitter linearity is measured using the transmitter linearity test pattern (see 94.2.9.4).

The resulting waveform is shown in Figure 94–18. Each measured level, V_A, V_B, V_C, and V_D, is measured within a 2 UI period starting 7 UI after the last level transition time. The minimum signal level, S_{min}, effective symbol levels ES_1 and ES_2, and level separation mismatch ratio, R_{LM}, are calculated based on Equation (94–9), Equation (94–10), Equation (94–11), Equation (94–12), and Equation (94–13), respectively.

The level separation mismatch ratio shall be greater than 0.92.

$$S_{\text{min}} = \frac{\min(V_D - V_C, V_C - V_B, V_B - V_A)}{2} \tag{94–9}$$

$$V_{\text{avg}} = \frac{V_A + V_B + V_C + V_D}{4} \tag{94–10}$$

$$ES_1 = \frac{V_B - V_{\text{avg}}}{V_A - V_{\text{avg}}} \tag{94–11}$$

$$ES_2 = \frac{V_C - V_{\text{avg}}}{V_D - V_{\text{avg}}} \tag{94–12}$$

$$R_{\text{LM}} = \frac{6 \cdot S_{\text{min}}}{V_D - V_A} \tag{94–13}$$

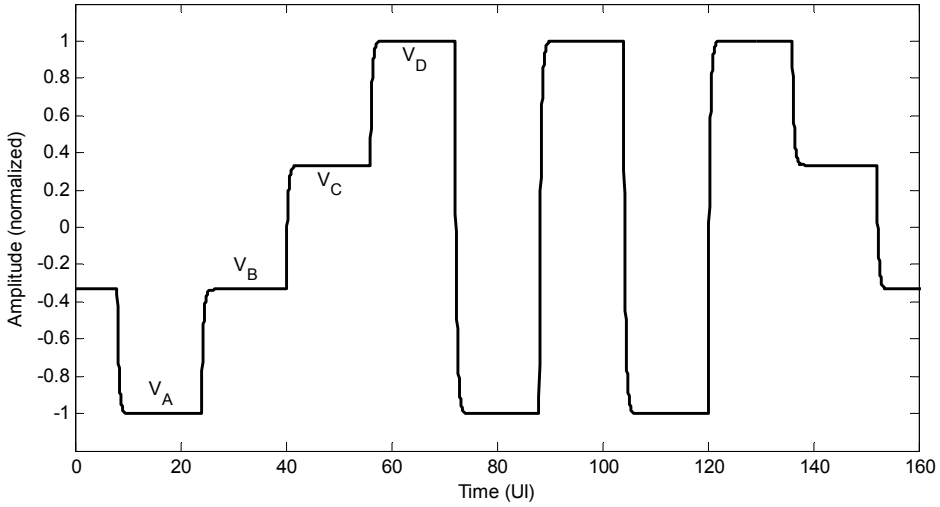

Figure 94–18—Transmitter linearity test pattern

94.3.12.5.2 Linear fit to the measured waveform

The following test procedure shall be followed to determine the linear fit pulse response, linear fit error, and normalized transmitter coefficient values.

For each configuration of the transmit equalizer, capture at least one complete cycle of the QPRBS13 test pattern (94.2.9.3) at TP0a per 85.8.3.3.4.

Compute the linear fit pulse response $p(k)$ from the captured waveform per 85.8.3.3.5 using $N_p = 16$ and $D_p = 2$. For aligned symbol values x(n) use –1, $-ES_1$, ES_2, and 1 to represent symbol values of 0, 1, 2, and 3, respectively, and where ES_1 and ES_2 are the effective symbol levels determined in 94.3.12.5.1.

Define $r(k)$ to be the linear fit pulse response when transmit equalizer coefficients have been set to the "preset" values (72.6.10.2.3.1).

For each configuration of the transmit equalizer, compute the normalized transmit equalizer coefficients, *c(i)*, according to 92.8.3.5.1.

94.3.12.5.3 Steady-state voltage and linear fit pulse peak

The linear fit pulse, $p(k)$, is determined according to the linear fit procedure in 94.3.12.5.2. The steady-state voltage v_f is defined to be the sum of the linear fit pulse $p(k)$ divided by M, determined in step 3 of the linear fit procedure.

The steady-state voltage shall be greater than or equal to 0.4 V and less than or equal to 0.6 V.

The peak value of $p(k)$ shall be greater than $0.85 \times v_f$.

94.3.12.5.4 Coefficient initialization

When the PMD enters the INITIALIZE state of the Training state diagram (Figure 72–5) or receives a valid request to "initialize" from the link partner, the coefficients of the transmit equalizer shall be configured such that the ratio (c(0)+c(1)–c(–1))/(c(0)+c(1)+c(–1)) is 1.29 ±10%, the ratio

(c(0)–c(1)+c(–1))/(c(0)+c(1)+c(–1)) is 2.57 ±10%, and the steady-state voltage, v_f, (see 94.3.12.5) is greater than or equal to 140 mV. These requirements apply upon the assertion of a coefficient status report of "updated" for all coefficients.

94.3.12.5.5 Coefficient step size

The normalized amplitude of each coefficient $c(i)$ is determined according to the linear fit procedure in 94.3.12.5.2.

The change in the normalized amplitude of coefficient $c(i)$ corresponding to a request to "increment" that coefficient shall be between 0.0083 and 0.05. The change in the normalized amplitude of coefficient $c(i)$ corresponding to a request to "decrement" that coefficient shall be between −0.05 and −0.0083.

The change in the normalized amplitude of the coefficient is defined to be the difference in the value measured prior to the assertion of the "increment" or "decrement" request (e.g., the coefficient update request for all coefficients is "hold") and the value upon the assertion of a coefficient status report of "updated" for that coefficient.

94.3.12.5.6 Coefficient range

When sufficient "increment" or "decrement" requests have been received for a given coefficient, the coefficient reaches a lower or upper bound based on the coefficient range or restrictions placed on the minimum steady-state differential output voltage or the maximum peak-to-peak differential output voltage.

With $c(-1)$ set to zero and both $c(0)$ and $c(1)$ having received sufficient "decrement" requests so that they are at their respective minimum values, the ratio $(c(0)–c(1))/(c(0)+c(1))$ shall be greater than or equal to 4.

With $c(1)$ set to zero and both $c(-1)$ and $c(0)$ having received sufficient "decrement" requests so that they are at their respective minimum values, the ratio $(c(0)–c(-1))/(c(0)+c(-1))$ shall be greater than or equal to 1.54.

Note that a coefficient may be set to zero by first asserting the preset control and then manipulating the other coefficients as required by the test.

94.3.12.6 Transmitter output jitter

Jitter measurements in this subclause are performed with transmitters on all PMD lanes enabled and transmitting the same pattern with identical transmit equalizer settings.

94.3.12.6.1 Clock random jitter and clock deterministic jitter

Clock random jitter (CRJrms) measured at the transmitter output using the methodology described in this subclause shall be less than 0.005 UI RMS regardless of transmit equalization setting.

Clock deterministic jitter (CDJ) measured at the transmitter output using the methodology described in this subclause shall be less than 0.05 UI peak-to-peak regardless of transmit equalization setting.

CRJrms and CDJ are determined using the following procedure:

1) CRJrms and CDJ are measured using the JP03A test pattern (94.2.9.1).
2) Using appropriate test equipment and procedure, capture the zero-crossing times, $T_{ZC}(i)$, of a pattern of length, N, of 10^7 symbols or greater.
3) Determine the average pulse width ΔT_{Avg} using Equation (94–14).
4) Determine the jitter series, $\tau(k)$, using Equation (94–15).
5) Apply the effect of a high-pass filter with the response given by Equation (94–16) to the jitter samples to obtain $\tau_{HPF}(k)$, where f is the frequency in MHz, f_n is 2.12 MHz, T is 0.0286 μs, and $j = \sqrt{-1}$.
6) Create a CDF as a function of $\tau_{HPF}(k)$.
7) From the CDF, determine J_5 as the difference between τ_{HPF} at the $(1–0.5 \times 10^{-5})$ and 0.5×10^{-5} probabilities, respectively, and J_6 as the difference between τ_{HPF} at the $(1–0.5 \times 10^{-6})$ and 0.5×10^{-6} probabilities, respectively.
8) Calculate CRJrms and CDJ using the relationship in Equation (94–17).

$$\Delta T_{Avg} = \frac{T_{ZC}(N) - T_{ZC}(1)}{N-1} \tag{94–14}$$

$$\tau(k) = T_{ZC}(k) - (k-1)\cdot\Delta T_{Avg} - T_{ZC}(1) \qquad k = 2, 3, \ldots N \tag{94–15}$$

$$G(f) = \frac{f}{f - j \times f_n e^{(-j2\pi fT)}} \tag{94–16}$$

$$\begin{bmatrix} CRJrms \\ CDJ \end{bmatrix} = \begin{bmatrix} 1.0538 & -1.0538 \\ -9.3098 & 10.3098 \end{bmatrix} \begin{bmatrix} J_6 \\ J_5 \end{bmatrix} \tag{94–17}$$

94.3.12.6.2 Even-odd jitter

Even-odd jitter (EOJ) measured at the transmitter output using the methodology described in this subclause shall be less than 0.019 UI peak-to-peak regardless of transmit equalization setting.

EOJ is determined using the following procedure:

1) EOJ is measured using the JP03B test pattern (94.2.9.2).
2) Using appropriate test equipment and procedure, capture the time for each of the 60 transitions.
3) Averaging of the vertical waveform or of each zero-crossing time is recommended to mitigate the contribution of uncorrelated noise and jitter.
4) Denote the averaged zero-crossing times as $T_{ZC}(i)$, where $i = \{1,2,...60\}$ and where $i = 1$ designates the transition from 3 to 0 after the consecutive pair of symbols {3,3}.
5) The set of 40 pulse widths, $\Delta T(j)$, isolated from the double-width pulses are determined using the relationship in Equation (94–18).
6) EOJ is calculated using the relationship in Equation (94–19).

$$\Delta T(j) = \left\{ \begin{array}{ll} T_{ZC}(j+10) - T_{ZC}(j+9) & 1 \le j \le 20 \\ T_{ZC}(j+19) - T_{ZC}(j+18) & 21 \le j \le 40 \end{array} \right\} \tag{94–18}$$

$$EOJ = \frac{\left| \sum_{j=1}^{20} \Delta T(2 \cdot j) - \sum_{j=1}^{20} \Delta T(2 \cdot j - 1) \right|}{40} \quad (94\text{–}19)$$

94.3.12.7 Transmitter output noise and distortion

Signal-to-noise-and-distortion ratio (SNDR) is measured at the transmitter output using the following method, with transmitters on all PMD lanes enabled and transmitting the same pattern with identical transmit equalizer settings.

Compute the linear fit to the captured waveform and the linear fit pulse response, $p(k)$, and error, $e(k)$, according to 94.3.12.5.2. Denote the standard deviation of $e(k)$ as σ_e.

Given the same configuration of the transmit equalizer, measure the RMS deviation from the mean voltage at a fixed point in a run of at least 8 consecutive identical levels. The transmitter linearity test pattern as specified in 94.3.12.5.1 is an example of a pattern that includes runs suitable to perform the measurement. The RMS deviation is measured for a run of each of the four levels. The average of the four measurements is denoted as σ_n.

SNDR is defined by Equation (94–20) where $p_{\max}$ is the maximum value of $p(k)$.

$$SNDR = 10\log_{10}\left(\frac{p_{\max}^2}{\sigma_e^2 + \sigma_n^2}\right) \text{ (dB)} \quad (94\text{–}20)$$

SNDR shall be greater than 31 dB for any allowable transmit equalizer setting.

94.3.13 PMD Receiver electrical characteristics

Receiver characteristics measured at TP5a are summarized in Table 94–14.

Table 94–14—Summary of receiver characteristics at TP5a

Parameter	Subclause reference	Value	Units
Differential input return loss (min.)	94.3.13.2	Equation (94–7)	dB
Differential to common-mode return loss (min.)	94.3.13.2	Equation (94–21)	dB
Interference tolerance	94.3.13.3	Table 94–15	—
Jitter tolerance	94.3.13.4	Table 94–16	—

94.3.13.1 Test fixture

The test fixture of Figure 94–19 or its equivalent is required for measuring the receiver specifications described in 94.3.13. The test fixture shall meet the requirements for insertion loss, insertion loss deviation, differential return loss, and common-mode return loss in 94.3.12.1.

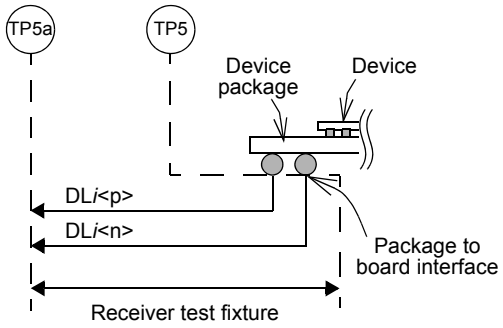

Figure 94–19—Receiver test fixture and test points

94.3.13.2 Receiver input return loss

The differential input return loss, in dB, of the receiver shall meet Equation (94–7). The reference impedance for differential return loss measurements shall be 100 Ω

The differential to common-mode return loss, in dB, of the receiver shall meet Equation (94–21). The return loss limit $RL_{\min}(f)$ is shown Figure 94–20.

$$RL(f) \geq RL_{\min}(f) = \left\{ \begin{array}{cc} 25 - 1.44f & 0.05 \leq f \leq 6.95\text{GHz} \\ 15 & 6.95 \leq f \leq 10\text{GHz} \end{array} \right\} \text{ (dB)} \quad (94\text{–}21)$$

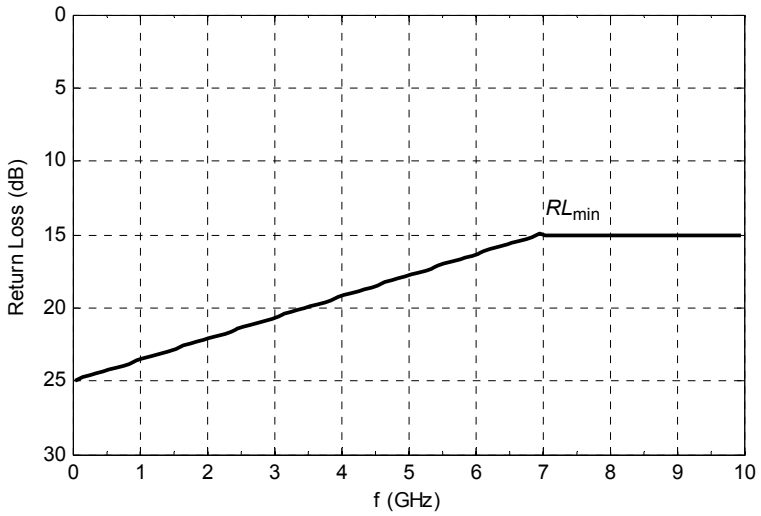

Figure 94–20—Receiver differential to common-mode return loss limit

94.3.13.3 Receiver interference tolerance

Receiver interference tolerance is defined by the procedure in Annex 93C. The receiver on each lane shall meet the FEC symbol error ratio requirement with channels matching the Channel Operating Margin (COM) and loss parameters for Test 1 and Test 2 in Table 94–15. Example fitted-insertion-loss curves for Test 1 and Test 2, as well as bounds resulting from the constraints on the fitted-insertion-loss coefficients with insertion loss at the limit specified for each test, are shown in Figure 94–21 and Figure 94–22, respectively. The

parameter RSS_DFE4 in Table 94–15 is a figure of merit for the test channel that is defined by Equation (93A–50) (see 93A.2).

The following considerations apply to the interference tolerance test. The transmitter noise parameter is SNDR (see 94.3.12.7). The transmitter output levels are set such that R_{LM} is equal to 0.92. The test transmitter meets the specifications in 94.3.12. The test transmitter is constrained such that for any transmitter equalizer setting the differential peak-to-peak voltage (see 94.3.12.3) is less than 800 mV, the pre-cursor peaking ratio (see 94.3.12.5.6) is less than 1.54, and the post-cursor peaking ratio (see 94.3.12.5.6) is less than 4. The lower frequency bound for the noise spectral density constraints, f_{NSD1}, is 1 GHz. The jitter parameters to be measured are CRJrms and CDJ (see 94.3.12.6.1). The return loss of the test setup in Figure 93C–4 measured at TP5 replica meets the requirements of Equation (94–4). The COM parameter σ_{RJ} is set to the measured value of CRJrms and the COM parameter A_{DD} is set to half the measured value of CDJ. Other COM parameters are set according to the values in Table 94–17. The test pattern to be used is the scrambled idles test pattern. A test system with a fourth-order Bessel-Thomson low-pass response with 17 GHz 3 dB bandwidth is to be used for measurement of the signal applied by the pattern generator and for measurements of the broadband noise.

Table 94–15—Receiver interference tolerance parameters

Parameter	Test 1 values		Test 2 values		Units
	Min	Max	Min	Max	
FEC symbol error ratio[a]		3.3×10^{-3}		3.3×10^{-3}	
Test channel parameters:					
COM, including effects of broadband noise		3		3	dB
Insertion loss at 7 GHz[b]		14.4	33		dB
a_0[c]	–1.5	1	–1.5	2	dB
a_1	0	1.6	0	3.8	dB/√GHz
a_2	0	1.6	0	4.2	dB/GHz
a_4	0	0.03	0	0.065	dB/GHz^2
RSS_DFE4	0.05	—	0.05	—	—

[a]The FEC symbol error ratio is measured in step 10 of the interference tolerance test method in 93C.2.
[b]Measured between TPt and TP5 (see Figure 93C–4).
[c]Coefficients are determined from insertion loss measured between TPt and TP5 (see Figure 93C–4) using the methodology in 93A.3 with f_{min} of 0.05 GHz, f_{max} of 13.59375 GHz, and maximum Δf of 0.01 GHz.

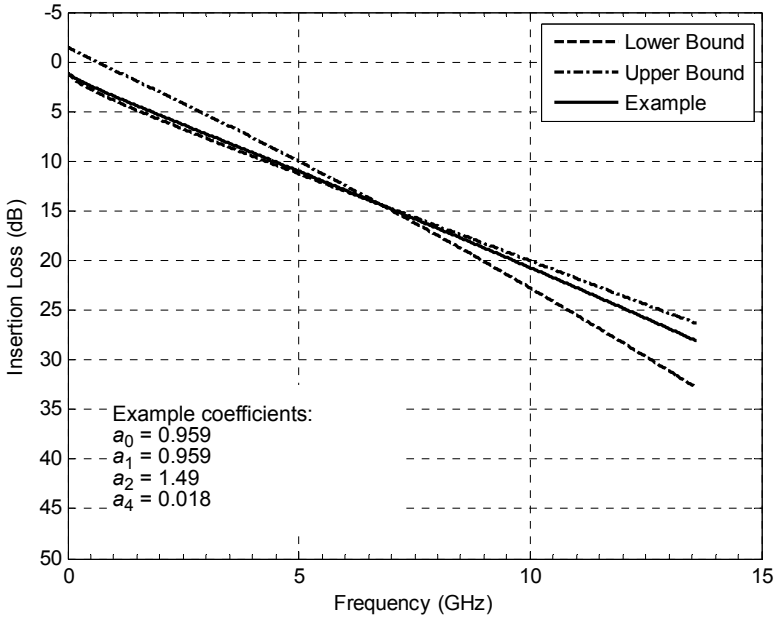

Figure 94–21—Insertion loss example and bounds for Test 1 channel

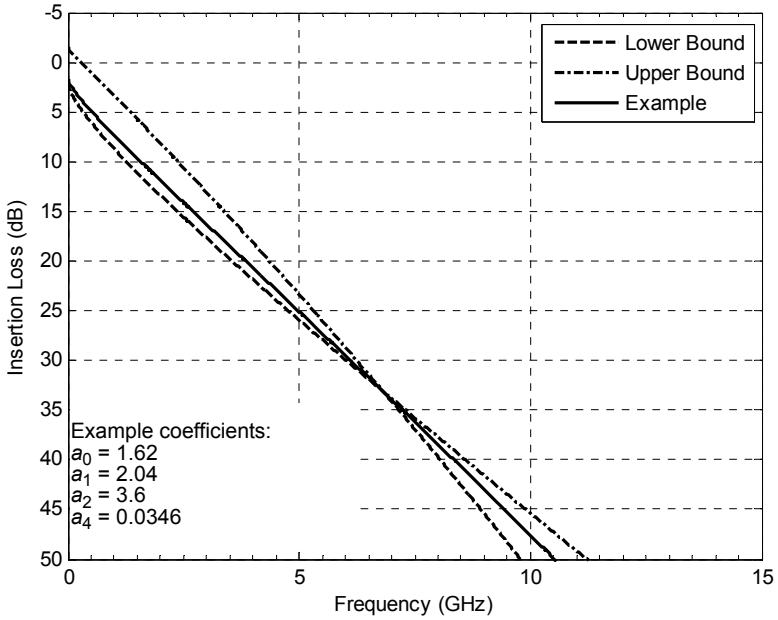

Figure 94–22—Insertion loss example and bounds for Test 2 channel

94.3.13.4 Receiver jitter tolerance

Receiver jitter tolerance is defined by the procedure defined in 94.3.13.4.2. The receiver FEC symbol error ratio shall be less than the maximum value for each pair of jitter frequency and peak-to-peak amplitude values listed in Table 94–16.

Table 94–16—Receiver jitter tolerance parameters

Parameter	Case A values	Case B values	Units
Maximum FEC symbol error ratio[a]	3.3×10^{-3}	3.3×10^{-3}	
Jitter frequency	16	160	kHz
Jitter amplitude	5	0.5	UI

[a]The FEC symbol error ratio is measured in step 3 of the jitter tolerance test method in 94.3.13.4.2.

94.3.13.4.1 Test setup

Jitter tolerance is measured using the test setup in Figure 93C–2 on each lane. The transmitter output is constrained as described in 93C.1. The Tx and channel noise sources are disabled. The test channel (TPt to TP5 replica) meets the requirements for the channel used for Test 2 in 94.3.13.3. The low-frequency jitter specified in Table 94–16 is applied to the output of the transmitter and is measured at TP0a.

94.3.13.4.2 Test method

The following jitter tolerance test method is repeated for each pair of jitter frequency and peak-to-peak amplitude values listed in Table 94–16.

1) Configure the transmitter with the corresponding jitter frequency and amplitude from Table 94–16.
2) Initiate training on the receiver under test and allow training to complete.
3) Measure the FEC symbol error ratio on each lane using the errored symbol counter, FEC_symbol-_error_counter_*i*, where *i* is the lane under test.

94.4 Channel characteristics

94.4.1 Channel Operating Margin

The Channel Operating Margin (COM) is computed using the procedure in 93A.1 with the Test 1 and Test 2 values in Table 94–17. Test 1 and Test 2 differ in the value of the device package model transmission line length z_p.

COM shall be greater than or equal to 3 dB for each test. This minimum value allocates margin for practical limitations on the receiver implementation as well as the largest step size allowed for transmitter equalizer coefficients.

94.4.2 Channel insertion loss

The insertion loss, in dB, of the channel is recommended to meet Equation (94–22). The insertion loss limit is shown Figure 94–23.

Table 94–17—COM parameter values

Parameter	Symbol	Value	Units
Signaling rate	f_b	13.59375	GBd
Maximum start frequency	$f_{\min}$	0.05	GHz
Maximum frequency step	Δf	0.01	GHz
Device package model Single-ended device capacitance Transmission line length, Test 1 Transmission line length, Test 2 Single-ended package capacitance at package-to-board interface	 C_d z_p z_p C_p	 2.5×10^{-4} 12 30 1.8×10^{-4}	 nF mm mm nF
Single-ended reference resistance	R_0	50	Ω
Single-ended termination resistance	R_d	55	Ω
Receiver 3 dB bandwidth	f_r	$0.75 \times f_b$	
Transmitter equalizer, minimum cursor coefficient	$c(0)$	0.62	—
Transmitter equalizer, pre-cursor coefficient Minimum value Maximum value Step size	$c(-1)$	 –0.18 0 0.02	 — — —
Transmitter equalizer, post-cursor coefficient Minimum value Maximum value Step size	$c(1)$	 –0.38 0 0.02	 — — —
Continuous time filter, DC gain Minimum value Maximum value Step size	g_{DC}	 –12 0 1	 dB dB dB
Continuous time filter, zero frequency	f_z	$f_b / 4$	GHz
Continuous time filter, pole frequencies	f_{p1} f_{p2}	$f_b / 4$ f_b	GHz
Transmitter differential peak output voltage Victim Far-end aggressor Near-end aggressor	 A_v A_{fe} A_{ne}	 0.4 0.4 0.6	 V V V
Number of signal levels	L	4	—
Level separation mismatch ratio	R_{LM}	0.92	—
Transmitter signal-to-noise ratio	SNR_{TX}	31	dB
Number of samples per unit interval	M	32	—
Decision feedback equalizer (DFE) length	N_b	16	UI
Normalized DFE coefficient magnitude limit for n = 1 for n = 2 to N_b	$b_{max}(n)$	 1 0.2	—
Random jitter, RMS	σ_{RJ}	0.005	UI
Dual-Dirac jitter, peak	A_{DD}	0.025	UI
One-sided noise spectral density	η_0	5.2×10^{-8}	V^2/GHz
Target detector error ratio	DER_0	3×10^{-4}	—

$$IL(f) \le IL_{\max}(f) = \left\{ \begin{array}{cc} a_0 + a_1 \cdot \sqrt{f} + a_2 \cdot f + a_3 \cdot f^2 + a_4 \cdot f^3 & f_{\min} \le f \le f_2 \\ a_5 + a_6 \cdot (f - f_2) & f_2 \le f \le f_{\max} \end{array} \right\} \text{ (dB)} \quad (94\text{–}22)$$

where

$IL(f)$ is the insertion loss in dB at frequency f
$IL_{\max}(f)$ is the maximum allowable insertion loss at frequency f
f is the measurement frequency in Hz
$f_{\min} = 0.05$ GHz
$f_2 = 7$ GHz
$f_{\max} = 15$ GHz
$a_0 = 0.8$
$a_1 = 1.7372 \times 10^{-4}$
$a_2 = 1.1554 \times 10^{-9}$
$a_3 = 2.7795 \times 10^{-19}$
$a_4 = -1.0423 \times 10^{-29}$
$a_5 = 33.467$
$a_6 = 1 \times 10^{-8}$

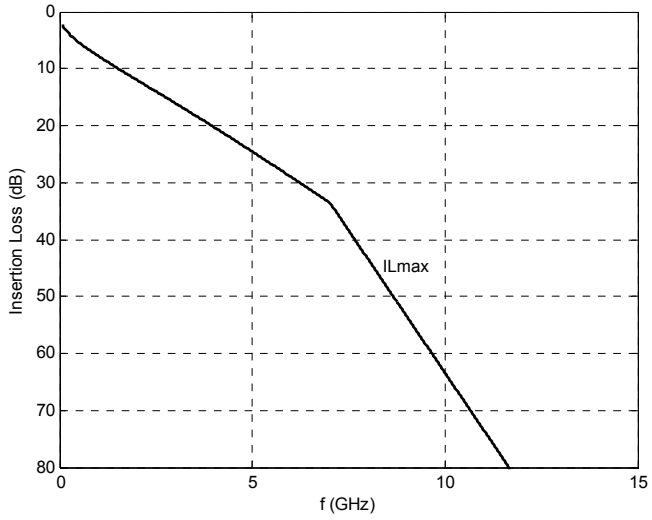

Figure 94–23—Channel insertion loss limit

94.4.3 Channel return loss

The return loss, in dB, of the channel is recommended to meet Equation (94–23). The return loss limit $RL_{\min}(f)$ is shown Figure 94–24.

$$RL(f) \ge RL_{\min}(f) = \left\{ \begin{array}{cc} 12 & 0.05 \le f \le f_b/4 \\ 12 - 15\log_{10}(4f/f_b) & f_b/4 < f \le f_b \end{array} \right\} \text{ (dB)} \quad (94\text{–}23)$$

where

$RL(f)$ is the return loss at frequency f in dB
$RL_{min}(f)$ is the minimum allowable return loss in dB
f is the measurement frequency in GHz
f_b is the signaling rate (13.59375) in GHz

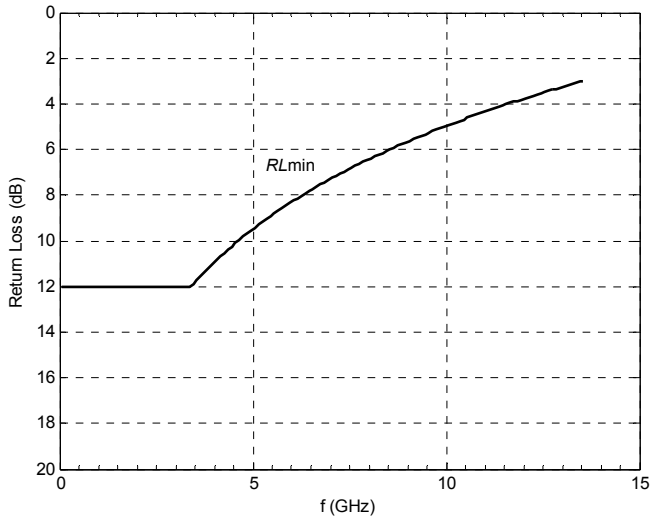

Figure 94–24—Channel return loss limit

94.4.4 Channel AC-coupling

The 100GBASE-KP4 transmitter shall be AC-coupled to the receiver. Common-mode specifications are defined as if the DC-blocking capacitor is implemented between TP0 and TP5. Should the capacitor be implemented outside TP0 and TP5, the common-mode specifications in Table 94–13 may not be appropriate.

The impact of a DC-blocking capacitor implemented between TP0 and TP5 is accounted for within the channel specifications. Should the capacitor be implemented outside TP0 and TP5, it is the responsibility of implementers to consider any necessary modifications to common-mode and channel specifications required for interoperability as well as any impact on the verification of transmitter and receiver compliance.

The low-frequency 3 dB cutoff of the AC-coupling shall be less than 50 kHz.

94.5 Environmental specifications

94.5.1 General safety

All equipment subject to this clause shall conform to applicable sections (including isolation requirements) of IEC 60950-1.

94.5.2 Network safety

The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

94.5.3 Installation and maintenance guidelines

It is recommended that sound installation practice, as defined by applicable local codes and regulations, be followed in every instance in which such practice is applicable.

94.5.4 Electromagnetic compatibility

A system integrating the 100GBASE-KP4 PHY shall comply with applicable local and national codes for the limitation of electromagnetic interference.

94.5.5 Temperature and humidity

A system integrating the 100GBASE-KP4 PHY 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.

94.6 Protocol implementation conformance statement (PICS) proforma for Clause 94, Physical Medium Attachment (PMA) sublayer, Physical Medium Dependent (PMD) sublayer, and baseband medium, type 100GBASE-KP4[23]

94.6.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 94, Physical Medium Attachment (PMA) sublayer, Physical Medium Dependent (PMD) sublayer, and baseband medium, type 100GBASE-KP4, 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.

94.6.2 Identification

94.6.2.1 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).	

94.6.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 94, Physical Medium Attachment (PMA) sublayer, Physical Medium Dependent (PMD) sublayer, and baseband medium, type 100GBASE-KP4
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	

[23] *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.

94.6.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
CGMII	CGMII	94.1	Interface is supported	O	Yes [ ] No [ ]
PCS	100GBASE-R PCS	94.1		M	Yes [ ]
RS-FEC	100GBASE-R RS-FEC	94.1		M	Yes [ ]
PMA	100GBASE-R PMA	94.1		O	Yes [ ] No [ ]
CAUI-10	CAUI-10	94.1	Interface is supported	O	Yes [ ] No [ ]
CAUI-4	CAUI-4	94.1	Interface is supported	O	Yes [ ] No [ ]
AN	Auto-negotiation	94.1		M	Yes [ ]
*MD	MDIO capability	94.3.5	Registers and interface supported	O	Yes [ ] No [ ]
*EEE	EEE capability	94.1	Capability is supported	O	Yes [ ] No [ ]
*GTD	Global PMD transmit disable function	94.3.6.6	Function is supported	O	Yes [ ] No [ ]
*LTD	PMD lane-by-lane transmit disable function	94.3.6.7	Function is supported	O	Yes [ ] No [ ]
*PDI	Physically instantiated PMD service interface	94.3.4	Interface is supported	O	Yes [ ] No [ ]
*CHNL	Channel	94.4	Channel specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]

94.6.4 PICS proforma tables for Physical Medium Attachment (PMA) sublayer, Physical Medium Dependent (PMD) sublayer, and baseband medium, type 100GBASE-KP4

94.6.4.1 PMA functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
AFS1	Overhead frame	94.2.2.2	Transmitter maps FEC bits to overhead frame	M	Yes [ ]
AFS2	Overhead	94.2.2.3	Transmitter maps sequence overhead bits	M	Yes []
AFS3	Termination blocks	94.2.2.4	Transmitter maps overhead frame bits to termination blocks	M	Yes [ ]
AFS4	Gray mapping	94.2.2.5	Transmitter maps each pair of termination block bits to Gray-mapped symbols	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
AFS5	Precoder	94.2.2.6	Transmitter precodes each Gray-mapped symbol	M	Yes [ ]
AFS6	PAM4 encoder	94.2.2.7	Transmitter maps each pre-coded symbol to PAM4 levels	M	Yes [ ]
AFS7	Transmit symbols	94.2.2.8	Transmitter sends each PAM4 symbol to the PMD	M	Yes [ ]
AFS8	Recover data	94.2.3	Receiver recovers data, meets performance requirements, and removes termination bits and overhead	M	Yes [ ]
AFS9	Overhead	94.2.3.1	Recover overhead sequence	M	Yes [ ]
AFS10	Link status	94.2.6	Provide link status to PMA client via PMA:IS_SIG-NAL.indication primitive	M	Yes [ ]
AFS11	PMA local loopback	94.2.7	Provide loopback from PMA SI input to PMA SI output	M	Yes [ ]
AFS12	PMA remote loopback	94.2.8	Provide loopback from PMA SI output to PMA SI input	O	Yes [ ] No [ ]
AFS13	JP03A pattern	94.2.9.1	Provide JP03A test pattern	M	Yes [ ]
AFS14	JP03B pattern	94.2.9.2	Provide JP03B test pattern	M	Yes [ ]
AFS15	QPRBS13 test pattern	94.2.9.3	Provide QPRBS13 test pattern	M	Yes [ ]
AFS16	Transmitter linearity test pattern	94.2.9.4	Provide transmitter linearity test pattern	M	Yes [ ]
AFS17	PMA control variables	94.2.10	Map PMA control variables to MDIO per Table 94–4	MD:M	Yes [ ] N/A [ ]
AFS18	PMA status variables	94.2.10	Map PMA status variables to MDIO per Table 94–5	MD:M	Yes [ ] N/A [ ]

94.6.4.2 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
DFS1	SIGNAL_OK assignment	94.3.1.3.1	Set SIGNAL_OK based on Global_PMD_signal_detect	M	Yes[]
DFS2	Sum of receive and transmit delays in one direction for PMA, PMD, and AN	94.3.3	Less than 8192 bit times	M	Yes [ ]
DFS3	Skew at SP3	94.3.4	Less than 54 ns	M	Yes [ ]
DFS4	Skew variation at SP3	94.3.4	Less than 0.6 ns	M	Yes [ ]
DFS5	Skew at SP4	94.3.4	Less than 134 ns	M	Yes [ ]
DFS6	Skew variation at SP4	94.3.4	Less than 3.4 ns	M	Yes [ ]
DFS7	Skew at SP5	94.3.4	Less than 145 ns	PDI:M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
DFS8	Skew variation at SP5	94.3.4	Less than 3.6 ns	PDI:M	Yes [ ]
DFS9	PMD control variables	94.3.5	Map PMD control variables to MDIO per Table 94–6	MD:M	Yes [ ] N/A [ ]
DFS10	PMD status variables	94.3.5	Map PMD status variables to MDIO per Table 94–7	MD:M	Yes [ ] N/A [ ]
DFS11	Transmit function	94.3.6.2	Convert 4 encoded symbol streams from PMD SI to 4 electrical signals at the MDI	M	Yes [ ]
DFS12	Transmit symbol value	94.3.6.2	A positive value for SLi<p> minus SLi<n> corresponds to a positive symbol value	M	Yes [ ]
DFS13	Transmit EEE alert signal	94.3.6.2	Send ALERT frame	EEE:M	Yes [ ] N/A [ ]
DFS14	Transmit EEE alert transmitter setting	94.3.6.2	Using coefficients determined during start-up	EEE:M	Yes [ ] N/A [ ]
DFS15	Receive function	94.3.6.3	Convert 4 electrical signals from MDI to 4 encoded symbol streams at the PMD SI	M	Yes [ ]
DFS16	Receive symbol value	94.3.6.3	A positive value for DLi<p> minus DLi<n> corresponds to a positive symbol value	M	Yes [ ]
DFS17	Signal detect parameter	94.3.6.4	Continuously send SIGNAL_DETECT to PMD SI.	M	Yes [ ]
DFS18	PMD_signal_detect_i when training is disabled by management	94.3.6.5	Set to one for all lanes.	M	Yes []
DFS19	PMD_signal_detect_i assertion time	94.3.6.5	Within 500 ns of compliant signal	EEE:M	Yes [ ] N/A [ ]
DFS20	PMD_signal_detect_i when transmitter is disabled	94.3.6.5	Not asserted when transmitter output meets requirements for disabled state	EEE:M	Yes [ ] N/A [ ]
DFS22	Global_PMD_transmit_ disable variable	94.3.6.6	When set to one, all transmitters satisfy the requirements of 94.3.12.3	GTD:M	Yes [ ] N/A [ ]
DFS21	Loopback when transmitter disabled	94.3.6.6	Loopback not affected	GTD:M	Yes [ ] N/A [ ]
DFS22	Transmitter output on transition to QUIET	94.3.6.6	Turn off and meet requirements in 94.3.12.3	EEE:M	Yes [ ] N/A [ ]
DFS23	Transmitter output on transition from QUIET	94.3.6.6	Turn on and meet requirements in 94.3.12.3	EEE:M	Yes [ ] N/A [ ]
DFS24	PMD_transmit_disable_*i* variable	94.3.6.7	When set to one, the transmitter for lane *i* satisfies the requirements of 94.3.12.3	LTD:M	Yes [ ] N/A [ ]
DFS25	PMD lane-by-lane transmit disable function affect on loopback	94.3.6.7	No effect	LTD:M	Yes [ ] N/A [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
DFS26	Loopback mode	94.3.6.8	Provided in adjacent PMA	M	Yes [ ]
DFS27	Loopback effect	94.3.6.8	Does not affect transmitter	M	Yes [ ]
DFS28	PMD_fault variable mapping to MDIO	94.3.7	Mapped to the fault bit as specified in45.2.1.2.3	MD:M	Yes [ ] N/A [ ]
DFS29	PMD_transmit_fault variable mapping to MDIO	94.3.8	Mapped to Transmit fault bit as specified in 45.2.1.7.4	MD:M	Yes [ ] N/A [ ]
DFS30	PMD_receive_fault variable mapping to MDIO	94.3.9	Mapped to Receive fault bit as specified in 45.2.1.7.5	MD:M	Yes [ ] N/A [ ]
DFS31	Control function	94.3.10.1	Independent per lane	M	Yes [ ]
DFS32	Training variables	94.3.10.1	Map to MDIO bits per 45.2.1.76	MD:M	Yes [ ] N/A [ ]
DFS33	Training frame marker	94.3.10.4	Frame marker encoded per 94.3.10.4	M	Yes [ ]
DFS34	Training frame control channel encoding	94.3.10.5.1	Differential manchester encoding	M	Yes [ ]
DFS35	Training frame control channel DME	94.3.10.5.1	Encoded based on rules in 94.3.10.5.1	M	Yes [ ]
DFS36	Training frame coding violation	94.3.10.5.1	Discard control channel if there is any coding violation	M	Yes [ ]
DFS37	Training frame control channel structure	94.3.10.5.2	Series of DME cells specified in 94.3.10.5.2	M	Yes [ ]
DFS38	Training frame coefficient update field	94.3.10.6	Format per Table 94–9	M	Yes [ ]
DFS39	Training frame coefficient transmission order.	94.3.10.6	Cell 15 first	M	Yes [ ]
DFS40	Training frame preset	94.3.10.6.1	Response per 72.6.10.2.3.1	M	Yes [ ]
DFS41	Training frame initialize	94.3.10.6.2	Response per 72.6.10.2.3.2 and conditions per 94.3.12.5.4	M	Yes [ ]
DFS42	Training frame coefficient update parity	94.3.10.6.3	Set for even parity in coefficient update field.	M	Yes [ ]
DFS43	Training frame coefficient update parity violation	94.3.10.6.3	Discard control channel if parity violation	M	Yes [ ]
DFS44	Training frame coefficient (*k*) update	94.3.10.6.4	Encoded per 94.3.10.6.4	M	Yes [ ]
DFS45	Training and alert frame status report field	94.3.10.7	Format per Table 94–10	M	Yes [ ]
DFS46	Training frame status transmission order	94.3.10.7	Cell 19 first	M	Yes [ ]
DFS47	Training frame status report parity	94.3.10.7.1	Set for even parity in status report field	M	Yes [ ]
DFS48	Training frame status report parity violation	94.3.10.7.1	Discard control channel if parity violation.	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
DFS49	Training frame countdown	94.3.10.7.2	Indicate transition per 94.3.10.7.2	M	Yes [ ]
DFS50	Training frame receiver ready	94.3.10.7.3	Signal local receiver state per 94.3.10.7.3	M	Yes [ ]
DFS51	Training frame coefficient status	94.3.10.7.4	Behavior per 72.6.10.2.4.5.	M	Yes [ ]
DFS52	Training frame coefficient update process	94.3.10.7.5	Update coefficients per 72.6.10.2.5	M	Yes [ ]
DFS53	PMD control response time	94.3.10.7.5	Response time less than 2 ms	M	Yes [ ]
DFS54	Training frame training pattern	94.3.10.8	Encode per 94.3.10.8	M	Yes [ ]
DFS55	Signal on transition from training mode to data mode	94.3.10.9	Signal according to 94.3.10.9	M	Yes [ ]
DFS56	Training frame lock state diagram	94.3.10.10	Implement per Figure 72–4 and 72.6.10.3.	M	Yes [ ]
DFS57	Training state diagram	94.3.10.11	Implement per Figure 72–5 and 72.6.10.3	M	Yes [ ]
DFS58	Transmitter setting in INITIALIZE state	94.3.10.11	Transmitter configured according to 94.3.12.5.4	M	Yes [ ]
DFS59	Coefficient update state diagram	94.3.10.12	Implement per Figure 72–6 and 72.6.10.3	M	Yes [ ]
DFS60	Alert frame	94.3.11.1	Structure per Figure 94–8	EEE:M	Yes [ ] N/A [ ]
DFS61	Alert frame marker	94.3.11.1.1	Implement per 94.3.10.4	EEE:M	Yes [ ] N/A [ ]
DFS62	Alert frame coefficient update field	94.3.11.1.2	Transmitted as all zeros	EEE:M	Yes [ ] N/A [ ]
DFS63	Alert frame status report field	94.3.11.1.3	Format per Table 94–10	EEE:M	Yes [ ] N/A [ ]
DFS64	Alert frame status report field order of transmission	94.3.11.1.3	Transmit cell 19 first	EEE:M	Yes [ ] N/A [ ]
DFS65	Alert frame status report parity	94.3.11.1.4	Behavior per 94.3.10.7.1	EEE:M	Yes [ ] N/A [ ]
DFS66	Alert frame mode	94.3.11.1.5	Always set to 1	EEE:M	Yes [ ] N/A [ ]
DFS67	Alert frame countdown	94.3.11.1.6	Indicate transition per 94.3.10.7.2	EEE:M	Yes [ ] N/A [ ]
DFS68	Alert frame PMA alignment offset	94.3.11.1.7	Indicate relative position of PMA frame and encoded per 94.3.11.1.7	EEE:M	Yes [ ] N/A [ ]
DFS69	Alert frame receiver ready	94.3.11.1.8	Always set to 1	EEE:M	Yes [ ] N/A [ ]
DFS70	Signal on transition from alert to data mode	94.3.11.1.9	Signal according to 94.3.11.1.9	M	Yes [ ]

94.6.4.3 PMD transmitter characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	Test fixture impedance	94.3.12.1.1	100 Ω	M	Yes [ ]
TC2	Test fixture differential return loss	94.3.12.1.1	Equation (94–4)	M	Yes [ ]
TC3	Test fixture common-mode return loss	94.3.12.1.1	Equation (94–5)	M	Yes [ ]
TC4	Test fixture differential insertion loss	94.3.12.1.2	≥ 1.2 dB and ≤ 1.6 dB at 12.89 GHz.	M	Yes [ ]
TC5	Test fixture insertion loss deviation	94.3.12.1.2	Magnitude < 0.1 dB	M	Yes [ ]
TC6	Signaling rate per lane	94.3.12.2	13.59375 GBd ± 100 ppm	M	Yes [ ]
TC7	Peak-to-peak differential output voltage	94.3.12.3	≤ 1200 mV regardless of transmit equalizer setting	M	Yes [ ]
TC8	Peak-to-peak differential output voltage, transmitter disabled	94.3.12.3	≤ 30 mV	M	Yes [ ]
TC9	DC common-mode output voltage	94.3.12.3	Between 0 V and 1.9 V with respect to signal ground	M	Yes [ ]
TC10	AC common-mode output voltage	94.3.12.3	≤ 30 mV RMS regardless of transmit equalizer setting	M	Yes [ ]
TC11	EEE transmitter output level when disabled	94.3.12.3	< 35 mV peak-to-peak differential within 500 ns	EEE:M	Yes [ ] N/A [ ]
TC12	EEE transmitter output level when enabled	94.3.12.3	90% of steady-state voltage by third alert frame marker and meet requirements of 94.3.12 with 1 µs	EEE:M	Yes [ ] N/A [ ]
TC13	Transmitter output DC common-mode voltage when disabled	94.3.12.3	Within ±150 mV of value when transmitter is enabled	M	Yes [ ]
TC14	Differential output return loss	94.3.12.4	Equation (94–7) with 100 Ω reference impedance	M	Yes [ ]
TC15	Common-mode output return loss	94.3.12.4	Equation (94–8) with 25 Ω reference impedance	M	Yes [ ]
TC16	Transition times	94.3.12.5	≥ 18 ps when transmit equalization is disabled	M	Yes [ ]
TC17	Level separation mismatch ratio, R_{LM}	94.3.12.5.1	> 0.92	M	Yes []
TC18	Steady-state voltage, v_f	94.3.12.5.3	≥ 0.4 V and ≤ 0.6 V after the transmit equalizer coefficients have been set to the “preset” values	M	Yes [ ]
TC19	Linear fit pulse peak	94.3.12.5.3	> 0.85 × v_f after the transmit equalizer coefficients have been set to the “preset” values	M	Yes [ ]

Item	Feature	Subclause	Value/Comment	Status	Support
TC20	Coefficient initialization	94.3.12.5.4	Satisfies the requirements of 94.3.12.5.4	M	Yes [ ]
TC21	Normalized coefficient step size for “increment”	94.3.12.5.5	Between 0.0083 and 0.05	M	Yes [ ]
TC22	Normalized coefficient step size for “decrement”	94.3.12.5.5	Between –0.05 and –0.0083	M	Yes [ ]
TC23	Maximum post-cursor equalization ratio	94.3.12.5.6	≥ 4	M	Yes [ ]
TC24	Maximum pre-cursor equalization ratio	94.3.12.5.6	≥ 1.54	M	Yes [ ]
TC25	Clock random jitter RMS	94.3.12.6.1	≤ 0.005 UI RMS regardless of the transmit equalization setting	M	Yes [ ]
TC26	Clock deterministic jitter	94.3.12.6.1	≤ 0.05 UI regardless of the transmit equalization setting	M	Yes [ ]
TC27	Even-odd jitter	94.3.12.6.2	≤ 0.03 UI regardless of the transmit equalization setting	M	Yes [ ]
TC28	SNDR	94.3.12.7	≥ 31 dB	M	Yes [ ]

94.6.4.4 PMD receiver characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
RC1	Test fixture insertion loss	94.3.13.1	Meet requirements in 94.3.12.1	M	Yes [ ]
RC2	Differential input return loss	94.3.13.2	Meets Equation (94–7) measured with a reference impedance of 100 Ω	M	Yes [ ]
RC3	Differential to common-mode return loss	94.3.13.2	Meets Equation (94–21).	M	Yes [ ]
RC4	Interference tolerance	94.3.13.3	Satisfy requirements in Table 94–15	M	Yes [ ]
RC5	Jitter tolerance	94.3.13.4	Satisfy requirements in Table 94–16	M	Yes [ ]

94.6.4.5 Channel characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	Channel Operating Margin (COM)	94.4.1	Greater than or equal to 3 dB	CHNL:M	Yes [ ] N/A [ ]
CC2	AC-coupling	94.4.4	Channel AC-couples the transmitter to the receiver	CHNL:M	Yes [ ] N/A [ ]
CC3	AC-coupling 3 dB cut-off frequency	94.4.4	Less than 50 kHz	CHNL:M	Yes [ ] N/A [ ]

94.6.4.6 Environment specifications

Item	Feature	Subclause	Value/Comment	Status	Support
ES1	General safety	94.5.1	Complies with application section of IEC 60950-1	M	Yes [ ]
ES2	Electromagnetic interference	94.5.4	Complies with applicable local and national codes	M	Yes [ ]

95. Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-SR4

95.1 Overview

This clause specifies the 100GBASE-SR4 PMD together with the multimode fiber medium. The PMD sublayer provides a point-to-point 100 Gb/s Ethernet link over four pairs of multimode fiber, up to at least 100 m. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 95–1, to the medium through the MDI and optionally with the management functions that may be accessible through the management interface defined in Clause 45, or equivalent.

Table 95–1—Physical Layer clauses associated with the 100GBASE-SR4 PMD

Associated clause	100GBASE-SR4
81—RS	Required
81—CGMII[a]	Optional
82—PCS for 100GBASE-R	Required
83—PMA for 100GBASE-R	Required
91—RS-FEC[b]	Required
83A—CAUI-10	Optional
83B—Chip-to-module CAUI-10[c]	Optional
83D—CAUI-4	Optional
83E—Chip-to-module CAUI-4	Optional
78—Energy Efficient Ethernet	Optional

[a]The CGMII is an optional interface. However, if the CGMII is not implemented, a conforming implementation must behave functionally as though the RS and CGMII were present.
[b]The option to bypass the Clause 91 RS-FEC correction function is not supported.
[c]This option requires the RS-FEC sublayer to be within the module. See 91.3.

Figure 95–1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 40 Gb/s and 100 Gb/s Ethernet is introduced in Clause 80 and the purpose of each PHY sublayer is summarized in 80.2.

100GBASE-SR4 PHYs with the optional Energy Efficient Ethernet (EEE) fast wake capability may enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization (see Clause 78). The deep sleep mode of EEE is not supported.

Further relevant information may be found in Clause 1 (terminology and conventions, references, definitions and abbreviations) and Annex A (bibliography, referenced as [B1], [B2], etc.).

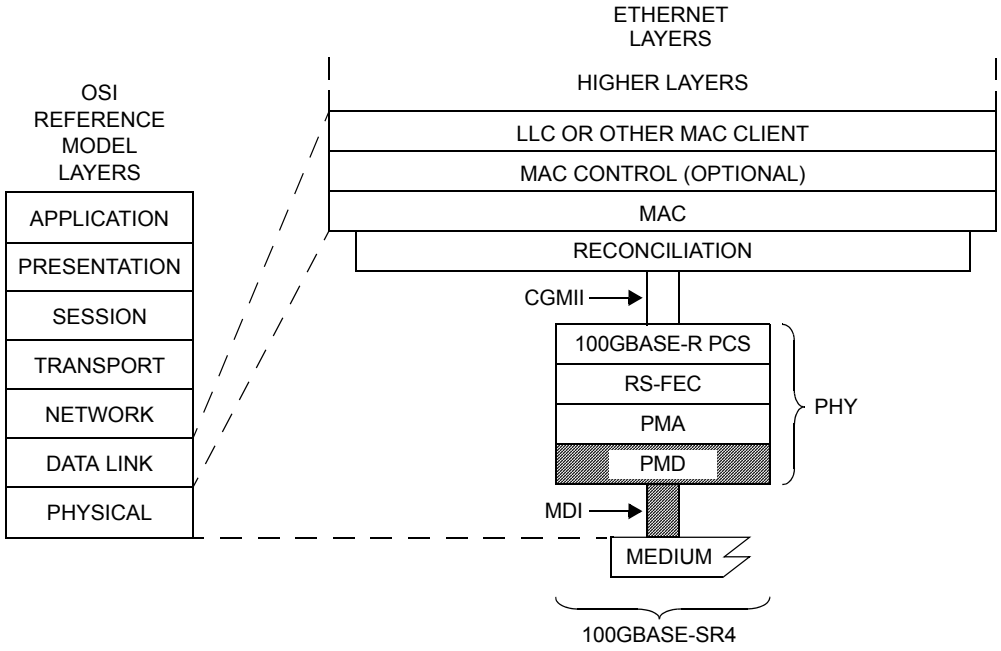

CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION
SR = PMD FOR MULTIMODE FIBER

Figure 95–1—100GBASE-SR4 PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

95.1.1 Bit error ratio

The bit error ratio (BER) shall be less than 5×10^{-5} provided that the error statistics are sufficiently random that this results in a frame loss ratio (see 1.4.223) of less than 6.2×10^{-10} for 64-octet frames with minimum interpacket gap when processed according to Clause 91.

If the error statistics are not sufficiently random to meet this requirement, then the BER shall be less than that required to give a frame loss ratio of less than 6.2×10^{-10} for 64-octet frames with minimum interpacket gap when processed according to Clause 91.

95.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 100GBASE-SR4 PMD. The service interface for this PMD is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_*i*.request
PMD:IS_UNITDATA_*i*.indication
PMD:IS_SIGNAL.indication

The 100GBASE-SR4 PMD has four parallel bit streams, hence $i = 0$ to 3.

In the transmit direction, the PMA continuously sends four parallel bit streams to the PMD, one per lane, each at a nominal signaling rate of 25.78125 GBd. The PMD converts these streams of bits into appropriate signals on the MDI.

In the receive direction, the PMD continuously sends four parallel bit streams to the PMA corresponding to the signals received from the MDI, one per lane, each at a nominal signaling rate of 25.78125 GBd.

The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication(SIGNAL_OK) inter-sublayer service primitive defined in 80.3.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the rx_bit parameters are undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_bit parameters are 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 BER defined in 95.1.1.

95.3 Delay and Skew

95.3.1 Delay constraints

The sum of the transmit and receive delays at one end of the link contributed by the 100GBASE-SR4 PMD including 2 m of fiber in one direction shall be no more than 2048 bit times (4 pause_quanta or 20.48 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

95.3.2 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the RS-FEC sublayer. Skew and Skew Variation are defined in 80.5 and specified at the points SP0 to SP7 shown in Figure 80–8.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 80.5. The measurements of Skew and Skew Variation are defined in 86.8.3.1 with the exception that the clock and data recovery units' high-frequency corner bandwidths are 10 MHz.

95.4 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If the MDIO interface is implemented, the mapping of MDIO control variables to PMD control variables shall be as shown in Table 95–2, and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 95–3.

Table 95–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable register	1.9.0	PMD_global_transmit_disable
PMD transmit disable 3	PMD transmit disable register	1.9.4	PMD_transmit_disable_3
PMD transmit disable 2	PMD transmit disable register	1.9.3	PMD_transmit_disable_2
PMD transmit disable 1	PMD transmit disable register	1.9.2	PMD_transmit_disable_1
PMD transmit disable 0	PMD transmit disable register	1.9.1	PMD_transmit_disable_0

Table 95–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	PMD_global_signal_detect
PMD receive signal detect 3	PMD receive signal detect register	1.10.4	PMD_signal_detect_3
PMD receive signal detect 2	PMD receive signal detect register	1.10.3	PMD_signal_detect_2
PMD receive signal detect 1	PMD receive signal detect register	1.10.2	PMD_signal_detect_1
PMD receive signal detect 0	PMD receive signal detect register	1.10.1	PMD_signal_detect_0

95.5 PMD functional specifications

The 100GBASE-SR4 PMD performs the Transmit and Receive functions, which convey data between the PMD service interface and the MDI.

95.5.1 PMD block diagram

The PMD block diagram is shown in Figure 95–2. 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 multimode fiber patch cord (TP2), between 2 m and 5 m in length. Unless specified otherwise, all transmitter measurements and tests defined in 95.8 are made at TP2. The optical receive signal is defined at the output of the fiber optic cabling (TP3) at the MDI (see 95.11.3). Unless specified otherwise, all receiver measurements and tests defined in 95.8 are made at TP3.

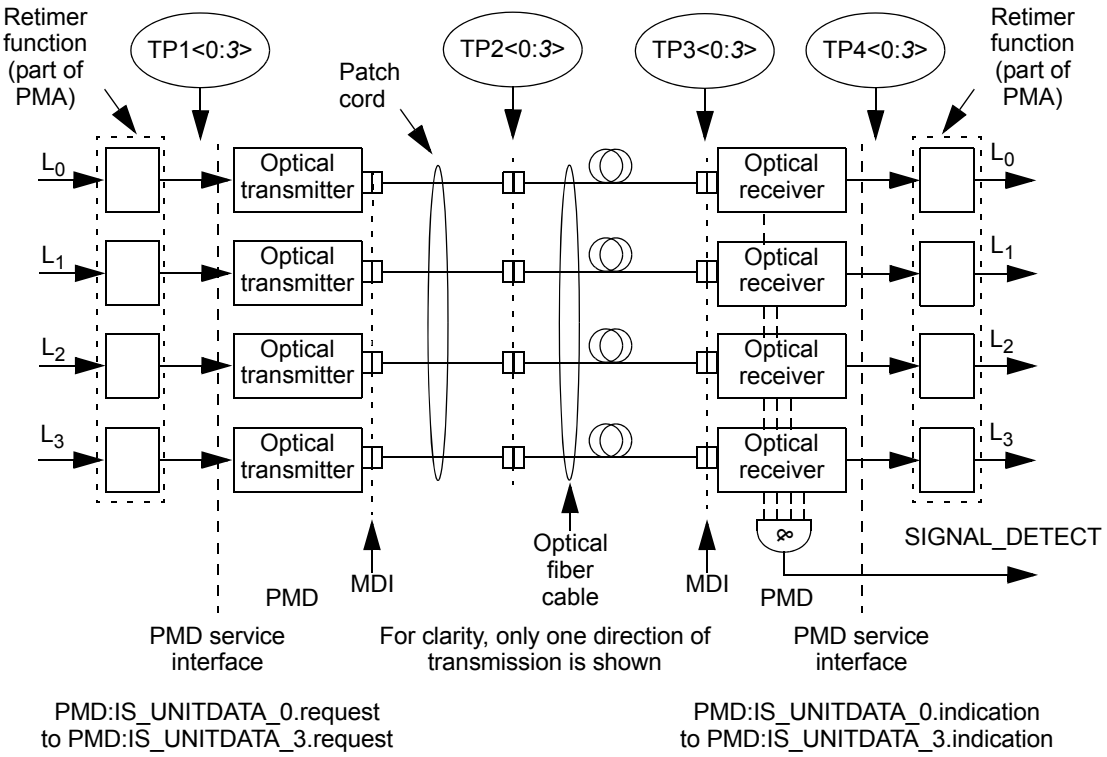

Figure 95–2—Block diagram for 100GBASE-SR4 transmit/receive paths

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 accessible in an implemented system).

95.5.2 PMD transmit function

The PMD Transmit function shall convert the four bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_3.request into four separate optical signal streams. The four optical signal streams shall then be delivered to the MDI, which contains four parallel light paths for transmit, according to the transmit optical specifications in this clause. The higher optical power level in each signal stream shall correspond to tx_bit = one.

95.5.3 PMD receive function

The PMD Receive function shall convert the four parallel optical signal streams received from the MDI into separate bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_3.indication, all according to the receive optical specifications in this clause. The higher optical power level in each signal stream shall correspond to rx_bit = one.

95.5.4 PMD global signal detect function

The PMD global 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:IS_SIGNAL.indication message is generated when a change in the value of SIGNAL_DETECT occurs. The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the inter-sublayer service interface primitives defined in 80.3.

SIGNAL_DETECT shall be a global indicator of the presence of optical signals on all four lanes. The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 95–4. The PMD receiver is not required to verify whether a compliant 100GBASE-SR4 signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

Table 95–4—SIGNAL_DETECT value definition

Receive conditions	SIGNAL_DETECT value
For any lane; Average optical power at TP3 ≤ –30 dBm	FAIL
For all lanes; [(Optical power at TP3 ≥ average receive power, each lane (min) in Table 95–7) AND (compliant 100GBASE-SR4 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 including the effects of 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.

95.5.5 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_*i*, where *i* represents the lane number in the range 0:3, shall be continuously set in response to the magnitude of the optical signal on its associated lane, according to the requirements of Table 95–4.

95.5.6 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.

95.5.7 PMD global transmit disable function (optional)

The PMD_global_transmit_disable function is optional and allows all of the optical transmitters to be disabled.

a) When the PMD_global_transmit_disable variable is set to one, this function shall turn off all of the optical transmitters so that each transmitter meets the requirements of the average launch power of the OFF transmitter in Table 95–6.
b) If a PMD_fault is detected, then the PMD may set the PMD_global_transmit_disable to one, turning off the optical transmitter in each lane.

95.5.8 PMD lane-by-lane transmit disable function (optional)

The PMD_transmit_disable_*i* (where *i* represents the lane number in the range 0:3) function is optional and allows the optical transmitter in each lane to be selectively disabled.

a) When a PMD_transmit_disable_*i* 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 95–6.
b) If a PMD_fault is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the optical transmitter in each lane.

If the optional PMD_transmit_disable_*i* function is not implemented in MDIO, an alternative method may be provided to independently disable each transmit lane.

95.5.9 PMD fault function (optional)

If the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set PMD_fault to one.

If the MDIO interface is implemented, PMD_fault shall be mapped to the fault bit as specified in 45.2.1.2.3.

95.5.10 PMD transmit fault function (optional)

If the PMD has detected a local fault on any transmit lane, the PMD shall set PMD_transmit_fault to one.

If the MDIO interface is implemented, PMD_transmit_fault shall be mapped to the transmit fault bit as specified in 45.2.1.7.4.

95.5.11 PMD receive fault function (optional)

If the PMD has detected a local fault on any receive lane, the PMD shall set the PMD_receive_fault variable to one.

If the MDIO interface is implemented, PMD_receive_fault shall be mapped to the receive fault bit as specified in 45.2.1.7.5.

95.6 Lane assignments

There are no lane assignments (within a group of transmit or receive lanes) for 100GBASE-SR4. While it is expected that a PMD will map electrical lane *i* to optical lane *i* and vice versa, there is no need to define the physical ordering of the lanes, as the RS-FEC sublayer is capable of receiving the lanes in any arrangement. The positioning of transmit and receive lanes at the MDI is specified in 95.11.3.1.

95.7 PMD to MDI optical specifications for 100GBASE-SR4

The operating range for the 100GBASE-SR4 PMD is defined in Table 95–5. A 100GBASE-SR4 compliant PMD operates on 50/125 μm multimode fibers, type A1a.2 (OM3) or type A1a.3 (OM4), according to the specifications defined in Table 95–12. A PMD that exceeds the operating range requirement while meeting all other optical specifications is considered compliant (e.g., a 100GBASE-SR4 PMD operating at 120 m meets the operating range requirement of 0.5 m to 100 m).

Table 95–5—100GBASE-SR4 operating range

PMD type	Required operating range[a]
100GBASE-SR4	0.5 m to 70 m for OM3
	0.5 m to 100 m for OM4

[a]The RS-FEC correction function may not be bypassed for any operating distance.

95.7.1 100GBASE-SR4 transmitter optical specifications

Each lane of a 100GBASE-SR4 transmitter shall meet the specifications in Table 95–6 per the definitions in 95.8.

Table 95–6—100GBASE-SR4 transmit characteristics

Description	Value	Unit
Signaling rate, each lane (range)	25.78125 ± 100 ppm	GBd
Center wavelength (range)	840 to 860	nm
RMS spectral width[a] (max)	0.6	nm
Average launch power, each lane (max)	2.4	dBm
Average launch power, each lane (min)	–8.4	dBm
Optical Modulation Amplitude (OMA), each lane (max)	3	dBm
Optical Modulation Amplitude (OMA), each lane (min)[b]	–6.4	dBm
Launch power in OMA minus TDEC (min)	–7.3	dBm
Transmitter and dispersion eye closure (TDEC), each lane (max)	4.3	dB
Average launch power of OFF transmitter, each lane (max)	–30	dBm
Extinction ratio (min)	2	dB
Optical return loss tolerance (max)	12	dB
Encircled flux[c]	≥ 86% at 19 μm ≤ 30% at 4.5 μm	
Transmitter eye mask definition {X1, X2, X3, Y1, Y2, Y3} Hit ratio 1.5×10^{-3} hits per sample	{0.3, 0.38, 0.45, 0.35, 0.41, 0.5}	

[a]RMS spectral width is the standard deviation of the spectrum.
[b]Even if the TDEC < 0.9 dB, the OMA (min) must exceed this value.
[c]If measured into type A1a.2 or type A1a.3 50 μm fiber in accordance with IEC 61280-1-4.

95.7.2 100GBASE-SR4 receive optical specifications

Each lane of a 100GBASE-SR4 receiver shall meet the specifications in Table 95–7 per the definitions in 95.8.

Table 95–7—100GBASE-SR4 receive characteristics

Description	Value	Unit
Signaling rate, each lane (range)	25.78125 ± 100 ppm	GBd
Center wavelength (range)	840 to 860	nm
Damage threshold[a] (min)	3.4	dBm
Average receive power, each lane (max)	2.4	dBm
Average receive power, each lane[b] (min)	–10.3	dBm
Receive power, each lane (OMA) (max)	3	dBm
Receiver reflectance (max)	–12	dB
Stressed receiver sensitivity (OMA), each lane[c] (max)	–5.2	dBm
Conditions of stressed receiver sensitivity test:[d]		
Stressed eye closure (SEC), lane under test	4.3	dB
Stressed eye J2 Jitter, lane under test	0.39	UI
Stressed eye J4 Jitter, lane under test (max)	0.53	UI
OMA of each aggressor lane	3	dBm
Stressed receiver eye mask definition {X1, X2, X3, Y1, Y2, Y3} Hit ratio 5×10^{-5} hits per sample	{0.28, 0.5, 0.5, 0.33, 0.33, 0.4}	

[a]The receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having this average power level on one lane. The receiver does not have to operate correctly at this input power.
[b]Average receive power, each lane (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]Measured with conformance test signal at TP3 (see 95.8.8) for the BER specified in 95.1.1.
[d]These test conditions are for measuring stressed receiver sensitivity. They are not characteristics of the receiver.

95.7.3 100GBASE-SR4 illustrative link power budget

An illustrative power budget and penalties for 100GBASE-SR4 channels are shown in Table 95–8.

95.8 Definition of optical parameters and measurement methods

All transmitter optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified.

Table 95–8—100GBASE-SR4 illustrative link power budget

Parameter	OM3	OM4	Unit
Effective modal bandwidth at 850 nm[a]	2000	4700	MHz.km
Power budget (for max TDEC)	8.2		dB
Operating distance	0.5 to 70	0.5 to 100	m
Channel insertion loss[b]	1.8	1.9	dB
Allocation for penalties[c] (for max TDEC)	6.3		dB
Additional insertion loss allowed	0.1	0	dB

[a]Per IEC 60793-2-10.
[b]The channel insertion loss is calculated using the maximum distance specified in Table 95–5 and cabled optical fiber attenuation of 3.5 dB/km at 850 nm plus an allocation for connection and splice loss given in 95.11.2.1.
[c]Link penalties are used for link budget calculations. They are not requirements and are not meant to be tested.

95.8.1 Test patterns for optical parameters

While compliance is to be achieved in normal operation, specific test patterns are defined for measurement consistency and to enable measurement of some parameters. Table 95–10 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. Any of the test patterns given for a particular test in Table 95–10 may be used to perform that test. As Pattern 3 is more demanding than Pattern 5 (which itself is the same or more demanding than other 100GBASE-R bit streams) an item that is compliant using Pattern 5 is considered compliant even if it does not meet the required limit using Pattern 3. The test patterns used in this clause are shown in Table 95–9.

Table 95–9—Test patterns

Pattern	Pattern description	Defined in
Square wave	Square wave (8 ones, 8 zeros)	83.5.10
3	PRBS31	83.5.10
4	PRBS9	83.5.10
5[a]	RS-FEC encoded scrambled idle	82.2.11[a]

[a]The pattern defined in 82.2.11 as encoded by Clause 91 RS-FEC for 100GBASE-SR4.

95.8.1.1 Multi-lane testing considerations

Stressed receiver sensitivity is defined for an interface at the BER specified in 95.1.1. The interface BER is the average of the four BERs of the receive lanes when they are stressed.

Measurements with Pattern 3 (PRBS31) allow lane-by-lane BER measurements. Measurements with Pattern 5 (RS-FEC encoded scrambled idle) give the interface BER if all lanes are stressed at the same time. If each lane is stressed in turn, the BER is diluted by the three unstressed lanes, and the BER for that stressed

Table 95–10—Test-pattern definitions and related subclauses

Parameter	Pattern	Related subclause
Wavelength, spectral width	3, 5 or valid 100GBASE-SR4 signal	95.8.2
Average optical power	3, 5 or valid 100GBASE-SR4 signal	95.8.3
Optical modulation amplitude (OMA)	Square wave or 4	95.8.4
Transmitter and dispersion eye closure (TDEC)	3, 5 or valid 100GBASE-SR4 signal	95.8.5
Extinction ratio	3, 5 or valid 100GBASE-SR4 signal	95.8.6
Transmitter optical waveform	3, 5 or valid 100GBASE-SR4 signal	95.8.7
Stressed receiver sensitivity	3, 5 or valid 100GBASE-SR4 signal	95.8.8
Stressed eye closure (SEC), calibration	3, 5 or valid 100GBASE-SR4 signal	95.8.8

lane alone must be found, e.g., by multiplying by four if the unstressed lanes have low BER. In stressed receiver sensitivity measurements, unstressed lanes may be created by setting the power at the receiver under test well above its sensitivity and/or not stressing those lanes with ISI and jitter, or by other means. Each receive lane is stressed in turn while all are operated. All aggressor lanes are operated as specified. To find the interface BER, the BERs of all the lanes when stressed are averaged.

Where relevant, parameters are defined with all co-propagating and counter-propagating lanes operational so that crosstalk effects are included. Where not otherwise specified, the maximum amplitude (OMA or VMA) for a particular situation is used, and for counter-propagating lanes, the minimum transition time is used. Alternative test methods that generate equivalent results may be used. While the lanes in a particular direction may share a common clock, the Tx and Rx directions are not synchronous to each other. If Pattern 3 is used for the lanes not under test using a common clock, there is at least 31 UI delay between the PRBS31 patterns on one lane and any other lane.

95.8.2 Center wavelength and spectral width

The center wavelength and RMS spectral width of each optical lane shall be within the range given in Table 95–6 if measured per TIA/EIA-455-127-A or IEC 61280-1-3. The lane under test is modulated using one of the test patterns specified in Table 95–10.

95.8.3 Average optical power

The average optical power of each lane shall be within the limits given in Table 95–6 if measured using the methods given in IEC 61280-1-1. The average optical power is measured using the test pattern defined in Table 95–10.

95.8.4 Optical Modulation Amplitude (OMA)

OMA shall be within the limits given in Table 95–6 if measured as defined in 52.9.5 for measurement with a square wave (8 ones, 8 zeros) test pattern or as defined in 68.6.2 (from the variable MeasuredOMA in 68.6.6.2) for measurement with a PRBS9 test pattern, with the exception that each optical lane is tested individually. See 95.8.1 for test pattern information.

95.8.5 Transmitter and dispersion eye closure (TDEC)

TDEC of each lane shall be within the limits given in Table 95–6 if measured using the methods specified in 95.8.5.1 and 95.8.5.2.

TDEC is a measure of each optical transmitter's vertical eye closure; it is based upon vertical histogram data from an eye diagram measured through an optical to electrical converter (O/E) with a bandwidth equivalent to a combined reference receiver and worst case optical channel. Table 95–10 specifies the test patterns to be used for measurement of TDEC.

95.8.5.1 TDEC conformance test setup

A block diagram for the TDEC conformance test is shown in Figure 95–3. Other measurement implementations may be used with suitable calibration.

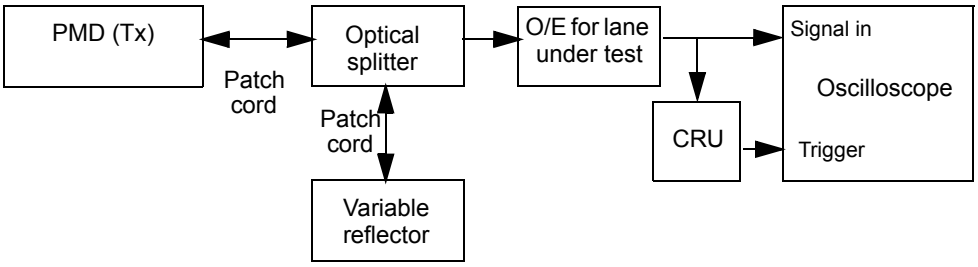

Figure 95–3—TDEC conformance test block diagram

Each optical lane is tested individually with all other lanes in operation. The optical splitter and variable reflector are adjusted so that each transmitter is tested with an optical return loss of 12 dB.

The combination of the O/E and the oscilloscope used to measure the optical waveform has a fourth-order Bessel-Thomson filter response with a bandwidth of 12.6 GHz. Compensation may be made for any deviation from an ideal fourth-order Bessel-Thomson response.

The clock recovery unit (CRU) has a corner frequency of 10 MHz and a slope of 20 dB/decade.

95.8.5.2 TDEC measurement method

The oscilloscope is set up to accumulate samples of the optical eye diagram for the transmitter under test, as illustrated in Figure 95–4.

OMA is measured according to 95.8.4.

The standard deviation of the noise of the O/E and oscilloscope combination, S, is determined with no optical input signal and the same settings as used to capture the histograms described below.

The average optical power (P_{ave}), the crossing points of the eye diagram, and the four vertical histograms used to calculate TDEC, are all measured using the same test pattern selected from those identified for TDEC in Table 95–10. The 0 UI and 1 UI crossing points are determined by the average of the eye diagram crossing times, as measured at P_{ave}, as illustrated in Figure 95–4.

Four vertical histograms are measured through the eye diagram, centered at 0.4 UI and 0.6 UI, and above and below P_{ave}, as illustrated in Figure 95–4.

Each histogram window has a width of 0.04 UI. Each histogram window has an inner height boundary, which is set close to P_{ave} (so that no further samples would be captured by moving it closer to P_{ave}), and an outer height boundary, which is set beyond the outer-most samples of the eye diagram (so that no further samples would be captured by increasing the outer boundary of the histogram window).

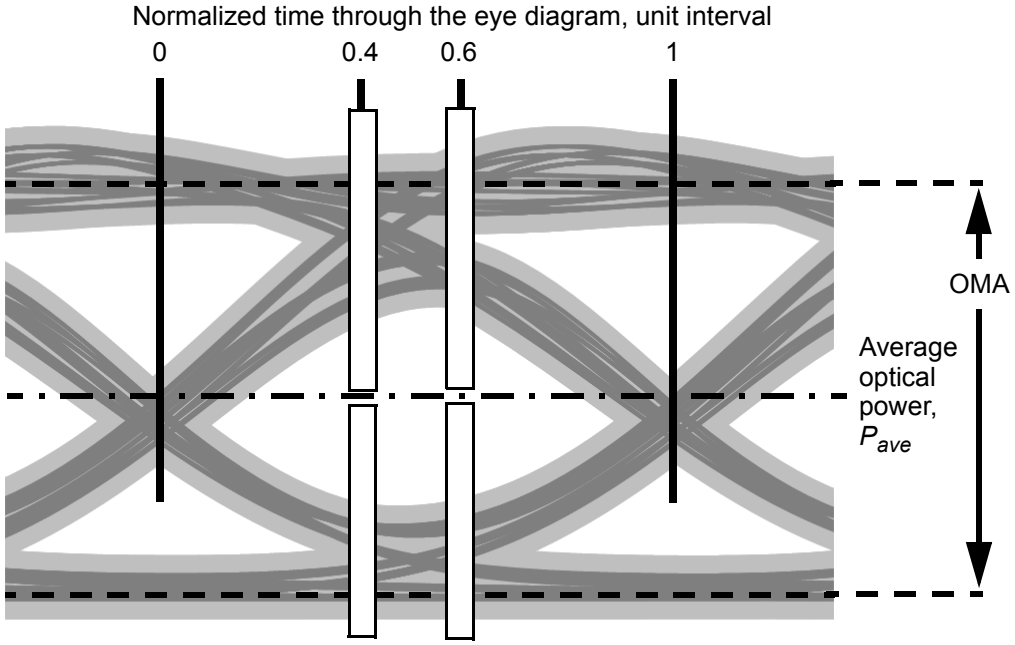

Figure 95–4—Illustration of the TDEC measurement

The distributions of the two histograms on the left are each multiplied by Q functions, which represent an estimate of the probability of errors caused by each part of the distribution for the greatest tolerable noise that could be added by an optical channel and a receiver. The resulting distributions are integrated and each integral is divided by the integral of the distribution it was derived from, giving two bit error probabilities. The Q function uses a standard deviation, σ_L, chosen so that the average of these two bit error probabilities is 5×10^{-5}. Similarly, for the two histograms on the right, a standard deviation, σ_R, is found.

Q(x) is the area under a Normal curve for values larger than *x* (the tail probability, related to the "complementary error function"), as shown in Equation (95–1):

$$Q(x) = \int_{x}^{\infty} \frac{e^{(-z^2/2)}}{\sqrt{2\pi}} dz \qquad (95\text{–}1)$$

where

x is $(y–P_{ave})/\sigma_G$ or $(P_{ave}–y)/\sigma_G$, as in Equation (95–2)

This procedure finds a value of σ_G such that Equation (95–2) is satisfied:

$$\frac{1}{2}\left(\frac{\int fu(y)Q\left(\frac{y-P_{ave}}{\sigma_G}\right)dy}{\int fu(y)dy}\right)+\frac{1}{2}\left(\frac{\int fl(y)Q\left(\frac{P_{ave}-y}{\sigma_G}\right)dy}{\int fl(y)dy}\right) = 5\times10^{-5} \quad (95–2)$$

where

fu(y), fl(y)	are the upper and lower distributions
σ_G	is the left or right standard deviation, σ_L or σ_R

The lesser of σ_L and σ_R is *N*.

The noise, *R*, that could be added by a receiver is given by Equation (95–3).

$$R = (1-M_1)\sqrt{N^2+S^2-M_2^2} \quad (95–3)$$

where

M_1, M_2	defined in Equation (95–4) and Equation (95–5), account for mode partition noise and modal noise that could be added by the optical channel
S	is the standard deviation of the noise of the O/E and oscilloscope combination

$$M_1 = 0.04 \quad (95–4)$$

$$M_2 = 0.0175P_{ave} \quad (95–5)$$

where

P_{ave}	is the average optical power of the eye diagram.

TDEC is given by Equation (95–6).

$$TDEC = 10\log_{10}\left(\frac{OMA}{2}\times\frac{1}{3.8906R}\right) \quad (95–6)$$

where

OMA	is the optical modulation amplitude as defined in 95.8.4

The factor 3.8906 is chosen for consistency with the BER of 5×10^{-5} given in 95.1.1.

The method described in 95.8.5.1 and 95.8.5.2 is the reference measurement method. Other equivalent measurement methods may be used with suitable calibration.

95.8.6 Extinction ratio

The extinction ratio of each lane shall be within the limits given in Table 95–6 if measured using the methods specified in IEC 61280-2-2. The extinction ratio is measured using one of the test patterns specified for extinction ratio in Table 95–10.

NOTE—Extinction ratio and OMA are defined with different test patterns (see Table 95–10).

95.8.7 Transmitter optical waveform (transmit eye)

The required optical transmitter pulse shape characteristics are specified in the form of a mask of the transmitter eye diagram as shown in Figure 86–4 with the transmitter eye mask coordinates and hit ratio in Table 95–6. The transmitter optical waveform of a port transmitting the test pattern specified in Table 95–10 shall meet specifications according to the methods specified in 86.8.4.6.1 with the following exceptions:

— The clock recovery unit's high-frequency corner bandwidth is 10 MHz.
— The filter nominal reference frequency f_r is 19.34 GHz and the filter tolerances are as specified for STM-64 in ITU-T G.691.

Compensation may be made for variation of the reference receiver filter response from an ideal fourth-order Bessel-Thomson response and for any excess reference receiver noise.

95.8.8 Stressed receiver sensitivity

Stressed receiver sensitivity shall be within the limits given in Table 95–7 if measured using the method defined by 95.8.8.1 and 95.8.8.5, with the conformance test signal at TP3 as described in 95.8.8.2.

Stressed receiver sensitivity is defined with all transmit and receive lanes in operation. Pattern 3 or Pattern 5, or a valid 100GBASE-SR4 signal is sent from the transmit section of the PMD under test. The signal being transmitted is asynchronous to the received signal. The interface BER of the PMD receiver is the average of the BER of all receive lanes while stressed and at the specified receive OMA.

95.8.8.1 Stressed receiver conformance test block diagram

A block diagram for the receiver conformance test is shown in Figure 95–5. The patterns used for the received conformance signal are specified in Table 95–10. The optical test signal is conditioned (stressed) using the stressed receiver methodology defined in 95.8.8.2 and has sinusoidal jitter applied as specified in 95.8.8.5. A suitable test set is needed to characterize and verify that the signal used to test the receiver has the appropriate characteristics. The fourth-order Bessel-Thomson filter has a 3 dB bandwidth of approximately 19 GHz. The low-pass filter is used to create ISI. The combination of the low-pass filter and the E/O converter should have a frequency response that results in the level of stressed eye closure (SEC) before the sinusoidal and Gaussian noise terms are added, as described in 95.8.8.2.

The sinusoidal amplitude interferer 1 causes jitter that 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 amplitude interferer 2 causes additional eye closure, but in conjunction with the finite edge rates from the limiter, also causes some jitter. The sinusoidally jittered clock represents other forms of jitter and also verifies that the receiver under test can track low-frequency jitter. The sinusoidal amplitude interferers may be set at any frequency between 100 MHz and 2 GHz, although care should be taken to avoid harmonic relationships between the sinusoidal interferers, the sinusoidal jitter, the signaling rate, and the pattern repetition rate. The Gaussian noise generator, the amplitude of the sinusoidal interferers, and the low-pass filter are adjusted so that the SEC and stressed eye J2 Jitter specifications given in Table 95–7 are met simultaneously, while the maximum stressed eye J4 Jitter specified in Table 95–7 is not exceeded, and while

also passing the stressed receiver eye mask in Table 95–7 according to the methods specified in 95.8.7 (the random noise effects such as RIN, or random clock jitter, do not need to be minimized).

For improved visibility for calibration, all elements in the signal path (cables, DC blocks, E/O converter, etc.) should have wide and smooth frequency response, and linear phase response, throughout the spectrum of interest. Baseline wander and overshoot and undershoot should be minimized.

The stressed receiver conformance test signal verification is described in 95.8.8.4.

Stressed receiver sensitivity is defined with all transmit and receive lanes in operation. Each receive lane is tested in turn while all aggressor receive lanes are operated as specified in Table 95–7. Pattern 3, Pattern 5, or a valid 100GBASE–SR4 signal is sent from the transmit section of the receiver under test. The signal being transmitted is asynchronous to the received signal. If Pattern 3 is used with a common clock for the transmit or receive lanes not under test, there is at least 31 UI delay between the PRBS31 patterns generated on one lane and any other lane.

For 100GBASE-SR4 the relevant BER is the interface BER at the PMD service interface. The interface BER is the average of the four BER of the receive lanes when stressed: see 95.8.1.1. If present, the RS-FEC sublayer can measure the lane symbol error ratio at its input. The lane BER can be assumed to be one tenth of the lane symbol error ratio. If each lane is stressed in turn, the PMD interface BER is the average of the BERs of all the lanes when stressed: see 95.8.1.1.

95.8.8.2 Stressed receiver conformance test signal characteristics and calibration

The conformance test signal is used to validate that each lane of the PMD receiver meets BER requirements with near worst-case waveforms at TP3.

The primary parameters of the stressed receiver conformance test signal are its stressed eye closure (SEC), stressed eye J2 Jitter, and stressed eye J4 Jitter. The SEC of the stressed receiver conformance test signal is measured according to 95.8.5, except that the combination of the O/E and the oscilloscope used to measure the waveform has a fourth-order Bessel-Thomson filter response with a bandwidth of 19.34 GHz, and the values of M_1 and M_2 in Equation (95–3) are set to zero. Stressed eye J2 Jitter and stressed eye J4 Jitter are defined in 95.8.8.3.

An example stressed receiver conformance test setup is shown in Figure 95–5; however, alternative test setups that generate equivalent stress conditions may be used.

The following steps describe a possible method for setting up and calibrating a stressed receiver conformance test signal when using a stressed receiver conformance test setup as shown in Figure 95–5:

1) Set the signaling rate of the test pattern generator to meet the requirements in Table 95–7.
2) With the sinusoidal jitter, sinusoidal interferers, and the Gaussian noise generator turned off, set the extinction ratio of the E/O to approximately the minimum specified in Table 95–6.
3) The required values of SEC and J2 Jitter, and the maximum value of J4 Jitter of the stressed receiver conformance test signal are given in Table 95–7. A valid stressed receiver conformance test signal may have a lower value of J4 jitter than the maximum specified in Table 95–7, provided it meets the required values of SEC and J2 Jitter.

   With the sinusoidal jitter, sinusoidal interferer 1, sinusoidal interferer 2, and the Gaussian noise generator turned off, at least 2.5 dB of SEC should be created by the selection of the appropriate bandwidth for the combination of the low-pass filter and the E/O converter. Any remaining SEC must be created with a combination of sinusoidal jitter, sinusoidal interference, and Gaussian noise.

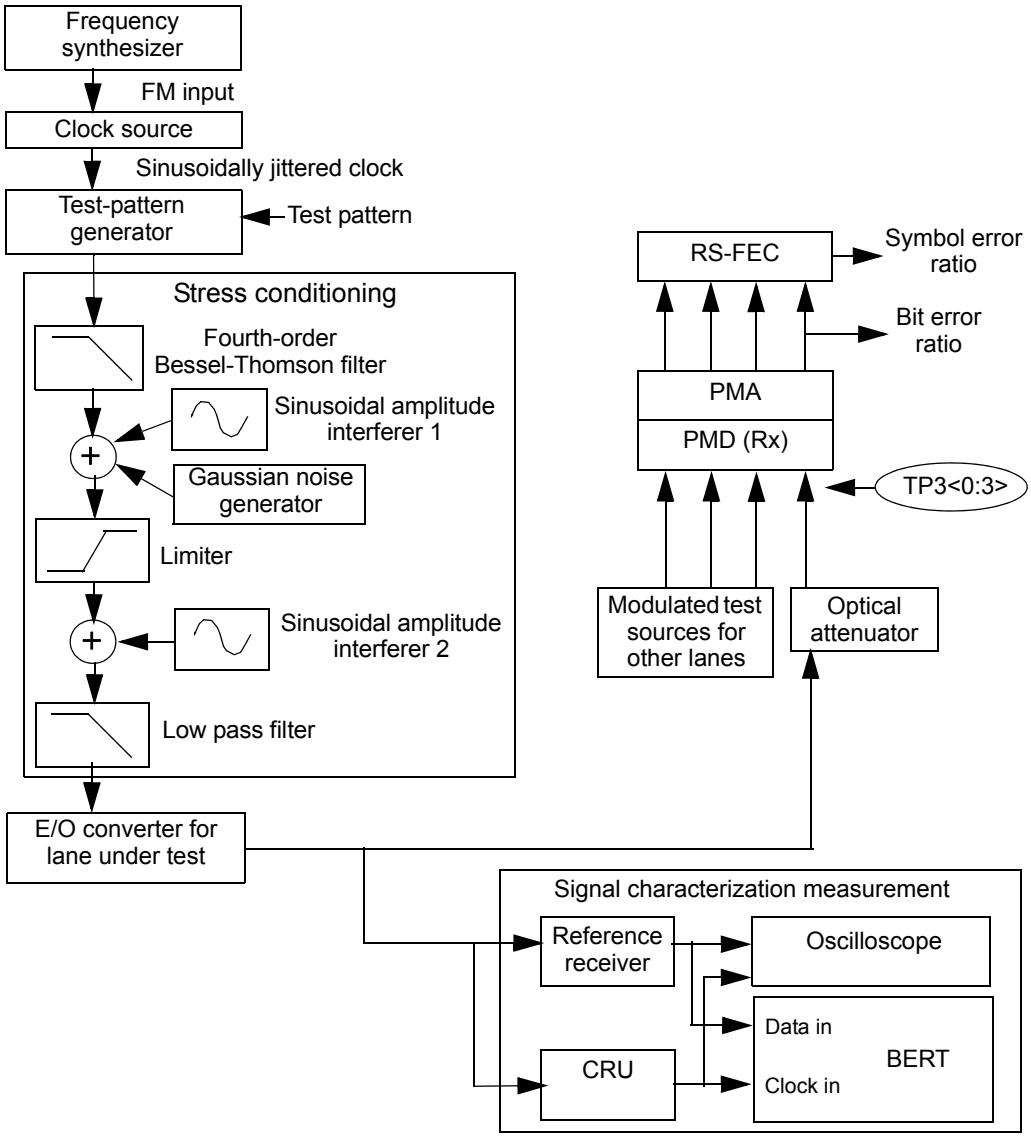

Figure 95–5—Stressed receiver conformance test block diagram

Sinusoidal jitter is added as specified in Table 95–11. When calibrating the conformance signal, the sinusoidal jitter frequency should be between 50 MHz and 10 times LB as defined in Table 95–11. Sinusoidal jitter amplitude may be calibrated by measuring the jitter on the oscilloscope, while transmitting the square wave pattern, and using a clean clock in place of the CRU to trigger the oscilloscope.

Iterate the adjustments of sinusoidal interferers and Gaussian noise generator and extinction ratio until the values of SEC and stressed eye J2 Jitter are met, while also meeting the following conditions: the maximum value of stressed eye J4 Jitter, as specified in Table 95–7, is not exceeded; the extinction ratio is approximately the minimum specified in Table 95–6; and sinusoidal jitter is as specified in Table 95–11.

Each receiver lane is conformance tested in turn. The source for the lane under test is adjusted to supply a signal at the input to the receiver under test at the “Stressed receiver sensitivity (OMA), each lane (max)”

specified in Table 95–7, and the test sources for the other lanes are set to the "OMA of each aggressor lane" specified in Table 95–7.

95.8.8.3 J2 and J4 Jitter

J2 Jitter is defined as the time interval at the average optical power level that includes all but 10^{-2} of the jitter distribution, which is the time interval from the 0.5th to the 99.5th percentile of the jitter histogram. J2 Jitter is defined using a clock recovery unit as in 95.8.7. If measured using an oscilloscope, the histogram should include at least 10 000 hits and should be taken over about 1% of the signal amplitude. If measured by plotting BER vs. decision time, J2 is the time interval between the two points with a BER of 2.5×10^{-3}.

J4 Jitter is defined as the time interval at the average optical power level that includes all but 10^{-4} of the jitter distribution. J4 Jitter is defined using a clock recovery unit as in 95.8.7. If measured using an oscilloscope, the histogram should include at least 1 000 000 hits and should be taken over about 1% of the signal amplitude. If measured by plotting BER vs. decision time, J4 is the time interval between the two points with a BER of 2.5×10^{-5}.

95.8.8.4 Stressed receiver conformance test signal verification

The stressed receiver conformance test signal can be verified using an optical reference receiver with an ideal fourth-order Bessel-Thomson response with a reference frequency f_r of 19.34 GHz. Use of G.691 tolerance filters may significantly degrade this calibration. The clock output from the clock source in Figure 95–5 is modulated with the sinusoidal jitter. To use an oscilloscope to calibrate the final stressed eye J2 Jitter and stressed eye J4 Jitter that includes the sinusoidal jitter component, a clock recovery unit (CRU of Figure 95–5) is required.

Care should be taken when characterizing the test signal because excessive noise/jitter in the measurement system would result in an input signal that does not fully stress the receiver under test. 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 standard, it is recommended that the implementer fully characterize the test equipment and apply appropriate guard bands to ensure that the stressed receiver conformance input signal meets the stress and sinusoidal jitter specified in 95.8.8.2 and 95.8.8.5.

95.8.8.5 Sinusoidal jitter for receiver conformance test

The sinusoidal jitter is used to test receiver jitter tolerance. The amplitude of the applied sinusoidal jitter is dependent on frequency as specified in Table 95–11.

Table 95–11—Applied sinusoidal jitter

Frequency range	Sinusoidal jitter peak-to-peak (UI)
$f < 100$ kHz	Not specified
100 kHz $< f \leq$ 10 MHz	$5 \times 10^5 / f$
10 MHz $< f \leq 10\ LB$[a]	0.05

[a] LB = loop bandwidth; upper frequency bound for added sinusoidal jitter should be at least 10 times the loop bandwidth of the receiver being tested.

95.9 Safety, installation, environment, and labeling

95.9.1 General safety

All equipment subject to this clause shall conform to IEC 60950-1.

95.9.2 Laser safety

100GBASE-SR4 optical transceivers shall conform to Hazard Level 1M 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 explicitly defines requirements and usage restrictions on the host system necessary to meet these safety certifications.[24]

95.9.3 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.

95.9.4 Environment

Normative specifications in this clause shall be met by a system integrating a 100GBASE-SR4 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.

95.9.5 Electromagnetic emission

A system integrating a 100GBASE-SR4 PMD shall comply with applicable local and national codes for the limitation of electromagnetic interference.

95.9.6 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.

[24]A host system that fails to meet the manufacturer's 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.

95.9.7 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., 100GBASE-SR4).

Labeling requirements for Hazard Level 1M lasers are given in the laser safety standards referenced in 95.9.2.

95.10 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 95–6.

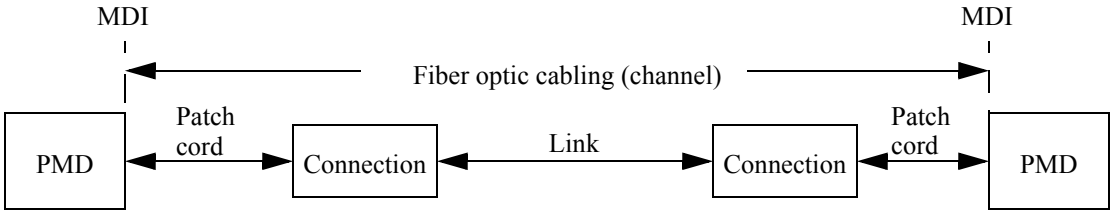

Figure 95–6—Fiber optic cabling model

The channel insertion loss is given in Table 95–12. A channel may contain additional connectors as long as the optical characteristics of the channel (such as attenuation, modal dispersion, reflections and losses of all connectors and splices) meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with IEC 61280-4-1:2009. As OM4 optical fiber meets the requirements for OM3, a channel compliant to the "OM3" column may use OM4 optical fiber, or a combination of OM3 and OM4. The term *channel* is used here for consistency with generic cabling standards.

Table 95–12—Fiber optic cabling (channel) characteristics for 100GBASE-SR4

Description	OM3	OM4	Unit
Operating distance (max)	70	100	m
Cabling Skew (max)	79		ns
Cabling Skew Variation[a] (max)	2.4		ns
Channel insertion loss[b] (max)	1.8	1.9	dB
Channel insertion loss (min)	0		dB

[a]An additional 400 ps of Skew Variation could be caused by wavelength changes, which are attributable to the transmitter not the channel.
[b]These channel insertion loss values include cable loss plus 1.5 dB allocated for connection and splice loss, over the wavelength range 840 nm to 860 nm.

95.11 Characteristics of the fiber optic cabling (channel)

The 100GBASE-SR4 fiber optic cabling shall meet the specifications defined in Table 95–12. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together.

95.11.1 Optical fiber cable

The fiber contained within the 100GBASE-SR4 fiber optic cabling shall comply with the specifications and parameters of Table 95–13. A variety of multimode cable types may satisfy these requirements, provided the resulting channel also meets the specifications of Table 95–12.

Table 95–13—Optical fiber and cable characteristics

Description	OM3[a]	OM4[b]	Unit
Nominal core diameter	50		μm
Nominal fiber specification wavelength	850		nm
Effective modal bandwidth (min)[c]	2000	4700	MHz.km
Cabled optical fiber attenuation (max)	3.5		dB/km
Zero dispersion wavelength (λ_0)	$1295 \le \lambda_0 \le 1340$		nm
Chromatic dispersion slope (max) (S_0)	0.105 for $1295 \le \lambda_0 \le 1310$ and $0.000375 \times (1590 - \lambda_0)$ for $1310 \le \lambda_0 \le 1340$		ps/nm^2 km

[a]IEC 60793-2-10 type A1a.2
[b]IEC 60793-2-10 type A1a.3
[c]When measured with the launch conditions specified in Table 95–6

95.11.2 Optical fiber connection

An optical fiber connection, as shown in Figure 95–6, consists of a mated pair of optical connectors.

95.11.2.1 Connection insertion loss

The maximum link distance is 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 per connection of 0.5 dB. Connections with lower loss characteristics may be used provided the requirements of Table 95–12 are met. However, the loss of a single connection shall not exceed 0.75 dB.

95.11.2.2 Maximum discrete reflectance

The maximum discrete reflectance shall be less than –20 dB.

95.11.3 Medium Dependent Interface (MDI)

The 100GBASE-SR4 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 95–6). The 100GBASE-SR4 PMD is

coupled to the fiber optic cabling through one connector plug into the MDI optical receptacle as shown in Figure 95–7. Example constructions of the MDI include the following:

a) PMD with a connectorized fiber pigtail plugged into an adapter;
b) PMD receptacle.

95.11.3.1 Optical lane assignments

The four transmit and four receive optical lanes of 100GBASE-SR4 shall occupy the positions depicted in Figure 95–7 when looking into the MDI receptacle with the connector keyway feature on top. The interface contains eight active lanes within twelve total positions. The transmit optical lanes occupy the left-most four positions. The receive optical lanes occupy the right-most four positions. The four center positions are unused.

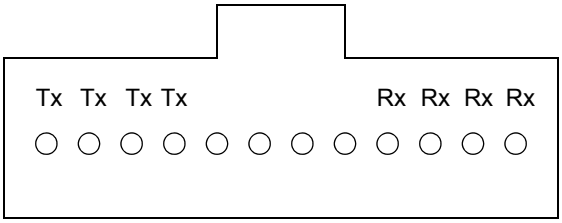

Figure 95–7—100GBASE-SR4 optical lane assignments

95.11.3.2 Medium Dependent Interface (MDI) requirements

The MDI adapter or receptacle shall meet the dimensional specifications for interface 7-1-3: *MPO adapter interface – opposed keyway configuration*, or interface 7-1-10: *MPO active device receptacle, flat interface*, as defined in IEC 61754-7-1. The plug terminating the optical fiber cabling shall meet the dimensional specifications of interface 7-1-4: *MPO female plug connector, flat interface for 2 to 12 fibres*, as defined in IEC 61754-7-1. The MDI shall optically mate with the plug on the optical fiber cabling. Figure 95–8 shows an MPO female plug connector with flat interface, and an MDI.

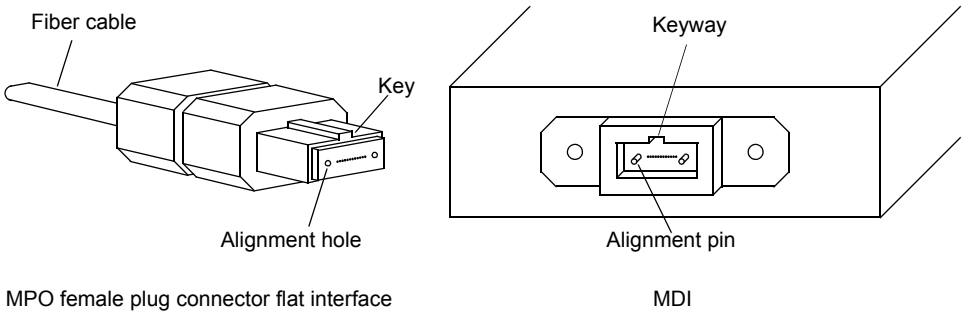

Figure 95–8—MPO female plug with flat interface and MDI

The MDI shall meet the interface performance specifications of IEC 61753-1 and IEC 61753-022-2.

NOTE—Transmitter compliance testing is performed at TP2 as defined in 95.5.1, not at the MDI.

95.12 Protocol implementation conformance statement (PICS) proforma for Clause 95, Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-SR4[25]

95.12.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 95, Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-SR4, 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.

95.12.2 Identification

95.12.2.1 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).	

95.12.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Clause 95, Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-SR4
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	

[25] *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.

95.12.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
SR4	100GBASE-SR4 PMD	95.7	Device supports requirements for 100GBASE-SR4 PHY	O	Yes [ ] No [ ]
*INS	Installation / cable	95.11	Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]
CTP1	Reference point TP1 exposed and available for testing	95.5.1	This point may be made available for use by implementers to certify component conformance	O	Yes [ ] No [ ]
CTP4	Reference point TP4 exposed and available for testing	95.5.1	This point may be made available for use by implementers to certify component conformance	O	Yes [ ] No [ ]
CDC	Delay constraints	95.3.1	Device conforms to delay constraints	M	Yes [ ]
CSC	Skew constraints	95.3.2	Device conforms to Skew and Skew Variation constraints	M	Yes [ ]
*MD	MDIO capability	95.4	Registers and interface supported	O	Yes [ ] No [ ]

95.12.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-SR4

95.12.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CF1	Compatible with 100GBASE-R RS-FEC, PCS, and PMA	95.1		M	Yes [ ]
CF2	Integration with management functions	95.1		O	Yes [ ] No [ ]
CF3	Bit error ratio	95.1.1	Meets the BER specified in 95.1.1	M	Yes [ ]
CF4	Transmit function	95.5.2	Conveys bits from PMD service interface to MDI	M	Yes [ ]
CF5	Mapping between optical signal and logical signal for transmitter	95.5.2	Higher optical power is a one	M	Yes [ ]
CF6	Receive function	95.5.3	Conveys bits from MDI to PMD service interface	M	Yes [ ]
CF7	Conversion of four optical signals to four electrical signals	95.5.3	For delivery to the PMD service interface	M	Yes [ ]
CF8	Mapping between optical signal and logical signal for receiver	95.5.3	Higher optical power is a one	M	Yes [ ]
CF9	Global Signal Detect function	95.5.4	Report to the PMD service interface the message PMD:IS_SIGNAL.indication(SIGNAL_DETECT)	M	Yes [ ]
CF10	Global Signal Detect behavior	95.5.4	SIGNAL_DETECT is a global indicator of the presence of optical signals on all four lanes	M	Yes [ ]
CF11	Lane-by-lane Signal Detect function	95.5.5	Sets PMD_signal_detect_*i* values on a lane-by-lane basis per requirements of Table 95–4	MD:O	Yes [ ] No [ ] N/A [ ]
CF12	PMD reset function	95.5.6	Resets the PMD sublayer	MD:O	Yes [ ] No [ ] N/A [ ]

95.12.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
CM1	Management register set	95.4		MD:M	Yes [ ] N/A [ ]
CM2	Global transmit disable function	95.5.7	Disables all of the optical transmitters with the PMD_global_transmit_disable variable	MD:O	Yes [ ] No [ ] N/A [ ]
CM3	PMD_lane_by_lane_transmit_ disable function	95.5.8	Disables the optical transmitter on the lane associated with the PMD_transmit_disable_*i* variable	MD:O	Yes [ ] No [ ] N/A [ ]
CM4	PMD lane-by-lane transmit disable	95.5.8	Disables each optical transmitter independently if CM3 = No	!MD:O	Yes [ ] No [ ]
CM5	PMD_fault function	95.5.9	Sets PMD_fault to one if any local fault is detected	MD:O	Yes [ ] No [ ] N/A [ ]
CM6	PMD_transmit_fault function	95.5.10	Sets PMD_transmit_fault to one if a local fault is detected on any transmit lane	MD:O	Yes [ ] No [ ] N/A [ ]
CM7	PMD_receive_fault function	95.5.11	Sets PMD_receive_fault to one if a local fault is detected on any receive lane	MD:O	Yes [ ] No [ ] N/A [ ]

95.12.4.3 PMD to MDI optical specifications for 100GBASE-SR4

Item	Feature	Subclause	Value/Comment	Status	Support
CSR1	Transmitter meets specifications in Table 95–6	95.7.1	Per definitions in 95.8	M	Yes [ ] N/A [ ]
CSR2	Receiver meets specifications in Table 95–7	95.7.2	Per definitions in 95.8	M	Yes [ ] N/A [ ]

95.12.4.4 Optical measurement methods

Item	Feature	Subclause	Value/Comment	Status	Support
COM1	Measurement cable	95.8	2 m to 5 m in length	M	Yes [ ]
COM2	Center wavelength and spectral width	95.8.2	Per TIA/EIA-455-127-A or IEC 61280-1-3 under modulated conditions	M	Yes [ ]
COM3	Average optical power	95.8.3	Per IEC 61280-1-1	M	Yes [ ]
COM4	OMA measurements	95.8.4	Each lane	M	Yes [ ]
COM5	Transmitter and dispersion eye closure (TDEC)	95.8.5	Each lane	M	Yes [ ]
COM6	Extinction ratio	95.8.6	Per IEC 61280-2-2	M	Yes [ ]
COM7	Transmit eye	95.8.7	Each lane	M	Yes [ ]
COM8	Stressed receiver sensitivity	95.8.8	See 95.8.8	M	Yes [ ]

95.12.4.5 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CES1	General safety	95.9.1	Conforms to IEC 60950-1	M	Yes [ ]
CES2	Laser safety—IEC Hazard Level 1M	95.9.2	Conforms to Hazard Level 1M laser requirements defined in IEC 60825-1 and IEC 60825-2	M	Yes [ ]
CES3	Electromagnetic interference	95.9.5	Complies with applicable local and national codes for the limitation of electromagnetic interference	M	Yes [ ]

95.12.4.6 Characteristics of the fiber optic cabling and MDI

Item	Feature	Subclause	Value/Comment	Status	Support
COC1	Fiber optic cabling	95.11	Meets requirements specified in Table 95–12	INS:M	Yes [ ] N/A [ ]
COC2	Optical fiber characteristics	95.11.1	Per Table 95–13	INS:M	Yes [ ] N/A [ ]
COC3	Maximum discrete reflectance	95.11.2.2	Less than –20 dB	INS:M	Yes [ ] N/A [ ]
COC4	MDI layout	95.11.3.1	Optical lane assignments per Figure 95–7	M	Yes [ ]
COC5	MDI dimensions	95.11.3.2	Per IEC 61754-7-1 interface 7-1-3 or interface 7-1-10	M	Yes [ ]
COC6	Cabling connector dimensions	95.11.3.2	Per IEC 61754-7-1 interface 7-1-4	INS:M	Yes [ ] N/A [ ]
COC7	MDI mating	95.11.3.2	MDI optically mates with plug on the cabling	M	Yes [ ]
COC8	MDI requirements	95.11.3.2	Meets IEC 61753-1 and IEC 61753-022-2	INS:M	Yes [ ] N/A [ ]

Annex 83A

(normative)

40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s ten-lane Attachment Unit Interface (CAUI-10)

83A.1 Overview

This annex defines the functional and electrical characteristics for the optional 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s ten-lane Attachment Unit Interface (CAUI-10). Figure 83A–1 shows the relationships of the XLGMII, PMA, XLAUI, and PMD for 40 Gb/s and CGMII, PMA, CAUI-10, and PMD for 100 Gb/s.

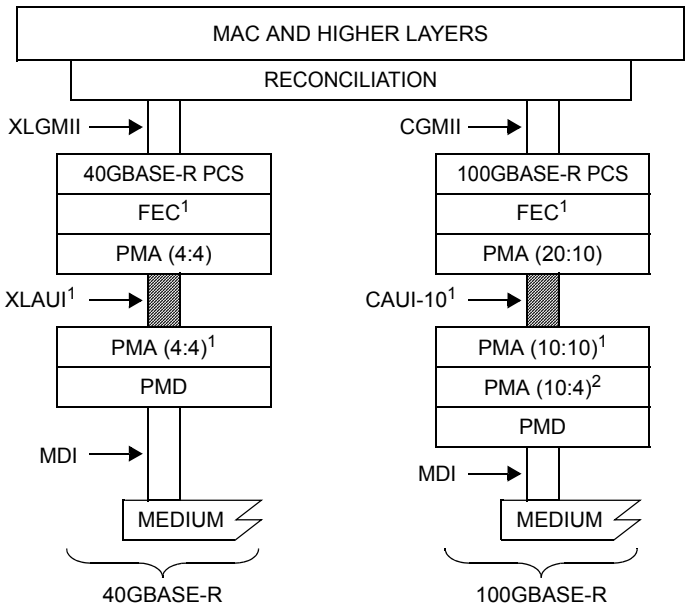

CAUI-10 = 100 Gb/s TEN-LANE ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLAUI = 40 Gb/s ATTACHMENT UNIT INTERFACE
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE
NOTE 2—CONDITIONAL BASED ON PMD TYPE

Figure 83A–1—Example relationship of XLAUI and CAUI-10 to IEEE 802.3 Ethernet model

The purpose of the optional XLAUI or CAUI-10 is to provide a flexible chip-to-chip and chip-to-module interconnect for 40 Gb/s or 100 Gb/s components. Annex 83A provides compliance requirements for

XLAUI/CAUI-10 transmitters and receivers while Annex 83B specifies the electrical requirements for the chip-to-module interconnection.

The XLAUI/CAUI-10 allows interconnect distances of approximately 25 cm over printed circuit board including one connector, see 83A.4.

An example application of CAUI-10 includes providing a physical connection between a ten-lane 100 Gb/s PMA and a 10:4 PMA mapping element. An example application of XLAUI is to provide lane extension for interfacing MAC and PHY components in a 40 Gb/s Ethernet system distributed across a circuit board.

The optional XLAUI/CAUI-10 interface has the following characteristics:

a) Independent transmit and receive data paths
b) Differential AC-coupled signaling with low voltage swing
c) Self-timed interface
d) Shared technology with other 40 Gb/s or 100 Gb/s interfaces
e) Utilization of 64B/66B coding

83A.1.1 Summary of major concepts

The following is a list of the major concepts of XLAUI and CAUI-10:

a) The optional XLAUI/CAUI-10 interface can be inserted between PMA layers in the IEEE 802.3 Ethernet model to transparently enable chip-to-chip communication
b) The XLAUI is organized into four lanes, the CAUI-10 is organized into ten lanes
c) The XLAUI/CAUI-10 interface is a parallel electrical interface with each lane running at a nominal rate of 10.3125 Gb/s

83A.1.2 Rate of operation

The XLAUI interface supports the 40 Gb/s data rate and the CAUI-10 interface supports the 100 Gb/s data rate. For 40 Gb/s applications, the data stream shall be presented in four lanes as described in Clause 83. For 100 Gb/s applications, the data stream shall be presented in ten lanes as described in Clause 83. The data is 64B/66B coded. The nominal signaling rate for each lane in both 40 Gb/s and 100 Gb/s applications shall be 10.3125 Gb/s.

83A.2 XLAUI/CAUI-10 link block diagram

XLAUI/CAUI-10 link is illustrated in Figure 83A–2. XLAUI/CAUI-10 channel is defined from the transmit pad to the receive pad including any AC-coupling in the path.

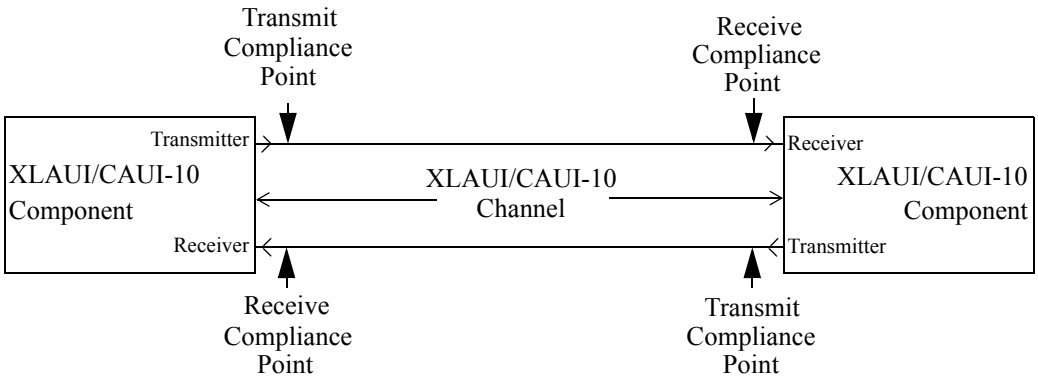

Figure 83A–2—Definition of transmit and receive compliance points

83A.2.1 Transmitter compliance points

The reference differential insertion loss, expressed in decibels, between the transmitter and the transmit compliance point is defined in Equation (83A–1) and illustrated in Figure 83A–3. The effects of differences between the actual insertion loss and the reference insertion loss are to be accounted in the measurements.

$$Insertion\_loss(f) = -0.00086 + 0.2286\sqrt{f} + 0.08386f \quad \text{dB} \qquad 0.01 \le f \le 11.1 \qquad \text{(83A–1)}$$

where

Insertion_loss(*f*) is the differential insertion loss at frequency *f*
f is the frequency in GHz

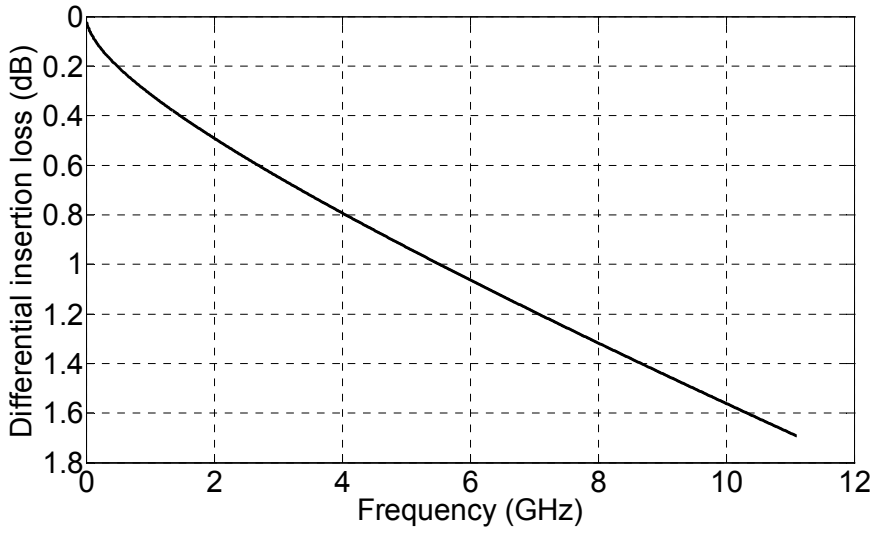

Figure 83A–3—Insertion loss between transmitter and transmit compliance point

83A.2.2 Receiver compliance points

The reference differential insertion loss, expressed in decibels, between the receiver compliance point and the receiver is defined in Equation (83A–2) and illustrated in Figure 83A–4. The effects of differences between the actual insertion loss and the reference insertion loss are to be accounted in the measurements.

$$Insertion\_loss(f) = -0.00086 + 0.2286\sqrt{f} + 0.08386f \quad \text{dB} \qquad 0.01 \le f \le 11.1 \qquad \text{(83A–2)}$$

where

Insertion_loss(*f*) is the differential insertion loss at frequency *f*
f is the frequency in GHz

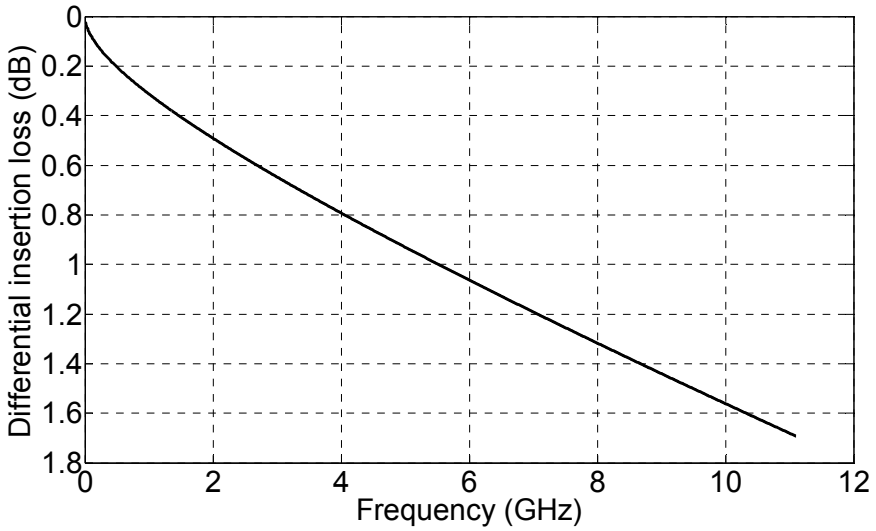

Figure 83A–4—Insertion loss between receive compliance point and receiver

83A.3 XLAUI/CAUI-10 electrical characteristics

The electrical characteristics of the XLAUI/CAUI-10 interface are specified such that they can be applied within a variety of 40 Gb/s Ethernet or 100 Gb/s Ethernet equipment types. The electrical characteristics for XLAUI/CAUI-10 shall meet the specifications defined in 83A.3.1, 83A.3.2, 83A.3.4, and 83A.3.5.

83A.3.1 Signal levels

The XLAUI/CAUI-10 is a low-swing AC-coupled differential interface. AC-coupling allows for interoperability between components operating from different supply voltages. Differential signal swings are defined in the following subclauses, and depend on several factors such as transmitter de-emphasis and transmission line losses.

83A.3.2 Signal paths

The XLAUI/CAUI-10 signal paths are point-to-point connections. Each path corresponds to a XLAUI/CAUI-10 lane, and is comprised of two complementary signals making a balanced differential pair. For XLAUI, there are four differential paths in each direction for a total of eight pairs, or sixteen connections. For CAUI-10, there are ten differential paths in each direction for a total of twenty pairs, or 40 connections.

83A.3.3 EEE operation

If the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option is supported (see Clause 78, 78.1.3.3.1), then the inter-sublayer service interface includes four additional primitives as described in 83.3 and may also support CAUI-10 shutdown.

If the EEE capability includes XLAUI/CAUI-10 shutdown (see 78.5.2), then when aui_tx_mode (see 83.5.11.3) is set to ALERT, the transmit direction sublayer sends a repeating 16-bit pattern, hexadecimal 0xFF00, which is transmitted across the XLAUI/CAUI-10. This sequence is transmitted regardless of the value of tx_bit presented by the PMA:IS_UNITDATA_i.request primitive or the rx_bit presented by the PMA:IS_UNITDATA_i.indication primitive. When aui_tx_mode is QUIET, the transmit direction XLAUI/CAUI-10 transmitter is disabled as specified in 83A.3.4.1.1. Similarly when the received aui_tx_mode is set to ALERT, the receive direction sublayer sends a repeating 16-bit pattern, hexadecimal 0xFF00, which is transmitted across the XLAUI/CAUI-10. This sequence is transmitted regardless of the value of tx_bit presented by the PMA:IS_UNITDATA_i.request primitive or the rx_bit presented by the PMA:IS_UNITDATA_i.indication primitive. When the received aui_tx_mode is QUIET, the receive direction XLAUI/CAUI-10 transmitter is disabled as specified in 83A.3.4.1.1.

83A.3.4 Transmitter characteristics

The XLAUI/CAUI-10 transmitter characteristics measured at the transmitter compliance point are specified in Table 83A–1. The XLAUI/CAUI-10 signaling rate shall be the signaling rate defined in Table 83A–1.

Table 83A–1—Transmitter characteristics

Parameter	Subclause reference	Value	Units
Signaling rate per lane (range)	—	10.3125 ± 100 ppm	GBd
Single-ended output voltage maximum minimum	83A.3.4.1	 4 –0.4	 V V
Maximum differential output voltage, peak-to-peak	83A.3.4.1	760	mV
Minimum de-emphasis	83A.3.4.1	4.4	dB
Maximum de-emphasis	83A.3.4.1	7	dB
Minimum VMA	83A.3.4.1	See Equation (83A–4)	mV
Maximum termination mismatch at 1 MHz	86A.5.3.2	5	%
Maximum output AC common-mode voltage, RMS	86A.5.3.1	15	mV
Minimum output rise and fall time (20% to 80%)	83A.3.4.2	24	ps
Differential output return loss	83A.3.4.3	See Equation (83A–5)	dB

Table 83A–1—Transmitter characteristics *(continued)*

Parameter	Subclause reference	Value	Units
Common-mode output return loss	83A.3.4.4	See Equation (83A–6)	dB
Maximum Total Jitter	83A.3.4.5	0.32	UI
Maximum Deterministic Jitter	83A.3.4.5	0.17	UI
Transmitter eye mask definition X1	83A.3.4.5	0.16	UI
Transmitter eye mask definition X2	83A.3.4.5	0.38	UI
Transmitter eye mask definition Y1	83A.3.4.5	200	mV
Transmitter eye mask definition Y2	83A.3.4.5	380	mV

83A.3.4.1 Output amplitude

Driver differential output amplitude shall be less than the maximum differential output voltage defined in Table 83A–1 including any transmit de-emphasis. DC-referenced logic levels are not defined since the receiver is AC-coupled. Single-ended output voltage shall be within the range specified in Table 83A–1 with respect to ground.

De-emphasis shall be greater than the minimum de-emphasis and less than the maximum de-emphasis defined in Table 83A–1. De-emphasis is defined as the ratio between the amplitude following a transition and the amplitude during a non-transition bit as seen in Equation (83A–3). VMA is defined in 86A.5.3.5 using the square wave or PRBS9 (Pattern 4) defined in Table 86–11.

See Figure 83A–5 for an illustration of absolute driver output voltage limits, definition of differential peak-to-peak amplitude, and definition of the parameters used to calculate de-emphasis. SLi<P> and SLi<N> are positive and negative sides of a differential signal pair for lane i (i = 0, 1, 2, 3 for XLAUI. For CAUI-10 i = 0:9).

$$\text{De-emphasis (dB)} = 20\log_{10}\left(\frac{\text{Differential peak-to-peak amplitude}}{\text{VMA}}\right) \quad (83A–3)$$

$$\text{Minimum VMA (mV)} = (234.64 - 2.13x + 0.18x^2) \times 1.32(10^{-y/20}) \quad (83A–4)$$

where

x is the rise or fall time (whichever is larger) in ps
y is de-emphasis value in dB

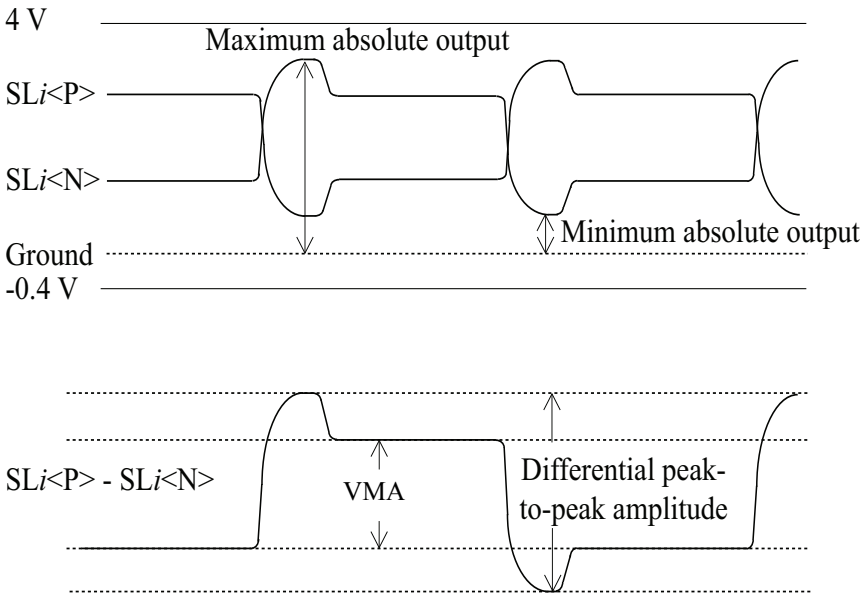

Figure 83A–5—Driver output voltage limits and definitions

83A.3.4.1.1 Amplitude and swing

For EEE capability with XLAUI/CAUI-10 shutdown, the XLAUI/CAUI-10 transmitter lane's differential peak-to-peak output voltage shall be less than 30 mV within 500 ns of aui_tx_mode changing to QUIET in the relevant direction. Furthermore, the XLAUI/CAUI-10 transmitter lane's differential peak-to-peak output voltage shall be greater than 720 mV within 500 ns of aui_tx_mode ceasing to be QUIET in the relevant direction.

83A.3.4.2 Rise/fall time

Differential rise/fall times shall be greater than the minimum output rise and fall time defined in Table 83A–1, as measured from the 20% to the 80% levels. Shorter transitions may result in excessive high-frequency components and increase EMI and crosstalk. The upper limit is defined by the transmit eye mask shown in Figure 83A–8. Rise/fall time is measured with de-emphasis off as defined in 83A.5.1. Rise/fall time measurements are taken using a square wave test pattern as defined in 83.5.10.

83A.3.4.3 Differential output return loss

Differential output return loss shall meet the requirements defined in Table 83A–1. Differential output return loss is given in Equation (83A–5) and is illustrated Figure 83A–6. Differential output 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. The maximum termination mismatch at 1 MHz shall be less than the requirement defined in Table 83A–1.

$$Return\_loss(f) \geq \begin{cases} 12 & 0.01 \leq f < 2.125 \\ 6.5 - 13.33\log_{10}\left(\frac{f}{5.5}\right) & 2.125 \leq f \leq 11.1 \end{cases} \text{ (dB)} \quad \text{(83A–5)}$$

where

Return_loss(*f*) is the differential output return loss at frequency *f*
f is the frequency in GHz

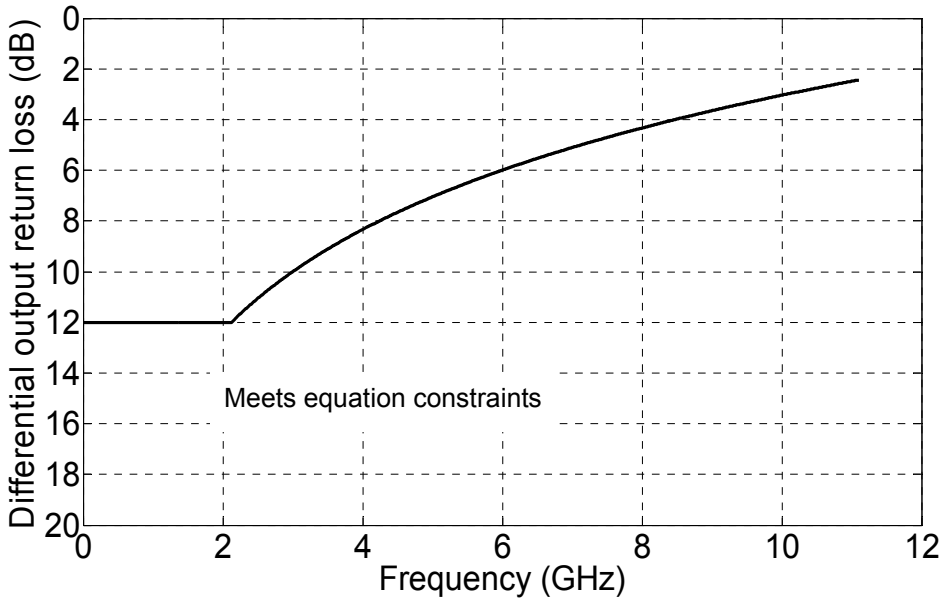

Figure 83A–6—Differential output return loss

83A.3.4.4 Common-mode output return loss

Common-mode output return loss shall meet the requirements defined in Table 83A–1. Common-mode output return loss is given in Equation (83A–6) and is illustrated in Figure 83A–7. Common-mode output 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 common-mode return loss measurements is 25 Ω.

$$Return\_loss(f) \geq \begin{cases} 9 & 0.01 \leq f < 2.125 \\ 3.5 - 13.33\log_{10}\left(\frac{f}{5.5}\right) & 2.125 \leq f < 7.1 \\ 2 & 7.1 \leq f \leq 11.1 \end{cases} \text{ (dB)} \quad \text{(83A–6)}$$

where

Return_loss(*f*) is the common-mode output return loss at frequency *f*
f is the frequency in GHz

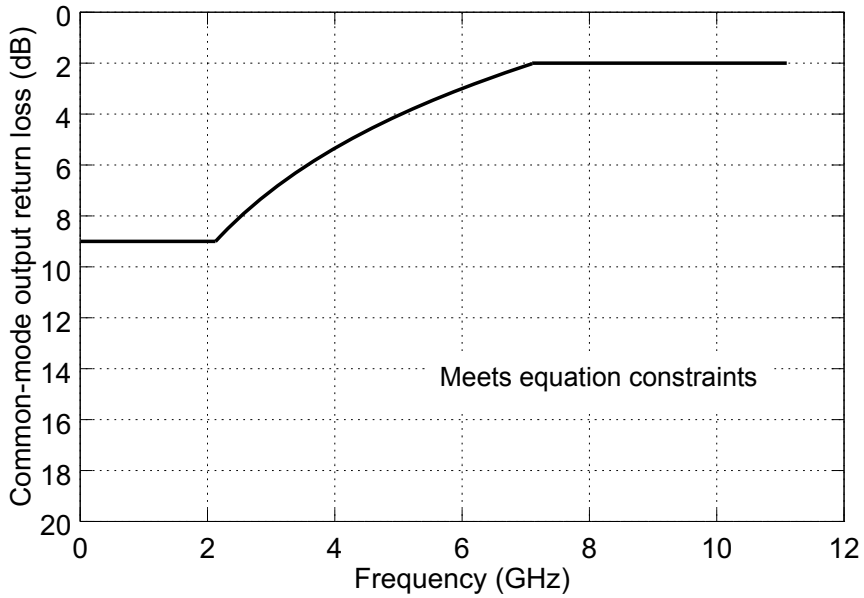

Figure 83A–7—Common-mode output return loss

83A.3.4.5 Transmitter eye mask and transmitter jitter definition

The measured transmit signal at the transmit compliance point shall meet the eye template specified in Figure 83A–8 and Table 83A–1. The template measurement requirements are specified in 83A.5.1. The measured jitter at the transmit compliance point shall be less than the maximum Total Jitter as defined in Table 83A–1 and a maximum Deterministic Jitter as defined in Table 83A–1. Jitter and eye mask measurement requirements are described in 83A.5.1 and are conducted with de-emphasis off.

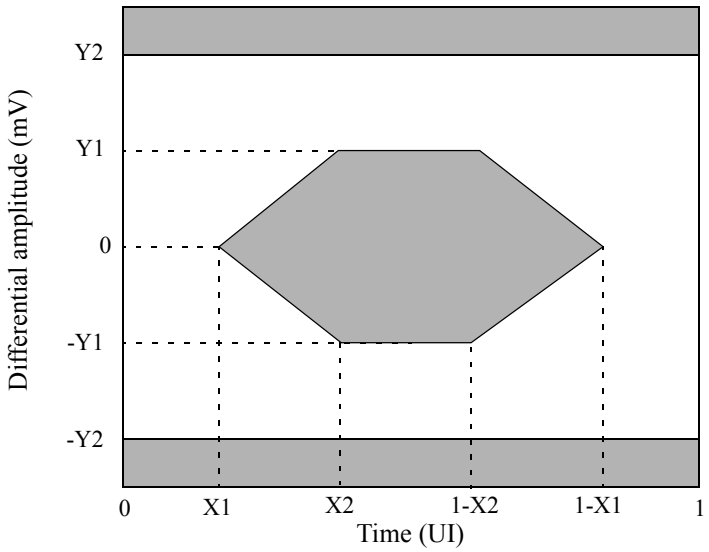

Figure 83A–8—Transmitter eye mask

83A.3.4.6 Global transmit disable function

Global transmit disable is optional for EEE capability. The transmit disable function shall turn off all transmitter lanes for a physically instantiated AUI in either the ingress or the egress direction. In the egress direction, the PMA may turn off all the transmitter lanes for the egress direction XLAUI/CAUI-10 if PEASE is asserted and aui_tx_mode is QUIET. In the ingress direction, the PMA may turn off all the transmitter lanes for the ingress direction XLAUI/CAUI-10 if PIASE is asserted and the received aui_tx_mode is QUIET. In both directions, the transmit disable function shall turn on all transmitter lanes after the appropriate direction aui_tx_mode changes to any state other than QUIET within a time and voltage level specified in 83A.3.4.1.1.

83A.3.5 Receiver characteristics

Receiver characteristics at the receiver compliance point are specified in Table 83A–2 and detailed in 83A.3.5.1 through 83A.3.5.6.

83A.3.5.1 Bit error ratio

The receiver shall operate with a BER of better than 10^{-12} in the presence of a compliant input signal as defined in 83A.3.5.2.

NOTE—A transmitter capable of operating at TJ = 0.3 UI and DJ = 0.16 UI and receiver capable of operating at stress jitter tolerance of TJ = 0.64 UI and DJ= 0.41 UI would have sufficient margin for operation at approximately BER 10^{-15}.

83A.3.5.2 Input signal definition

A compliant input signal to a XLAUI/CAUI-10 receiver has characteristics determined by a compliant XLAUI/CAUI-10 driver and channel. The input signal definition satisfies the eye mask defined in Table

Table 83A–2—Receiver characteristics

Parameter	Subclause reference	Value	Units
Signaling rate per lane (range)	—	10.3125 ± 100 ppm	GBd
Minimum input AC common-mode voltage tolerance, RMS	86A.5.3.1	20	mV
Minimum input rise and fall time tolerance (20% to 80%)	83A.3.4.2	24	ps
Differential input return loss	83A.3.5.3	See Equation (83A–7)	dB
Differential to common-mode input return loss	83A.3.5.4	See Equation (83A–8)	dB
Stressed receiver tolerance			
Minimum Total Input Jitter Tolerance	83A.3.5.2	0.62	UI
Minimum Deterministic Input Jitter Tolerance	83A.3.5.2	0.42	UI
Receiver eye mask definition X1	83A.3.5.2	0.31	UI
Receiver eye mask definition X2	83A.3.5.2	0.5	UI
Receiver eye mask definition Y1	83A.3.5.2	42.5	mV
Receiver eye mask definition Y2	83A.3.5.2	425	mV

83A–2 and Figure 83A–9. Input signal jitter does not exceed the jitter tolerance requirements specified in 83A.3.5.6. Stressed receiver measurement requirements are specified in 83A.5.2.

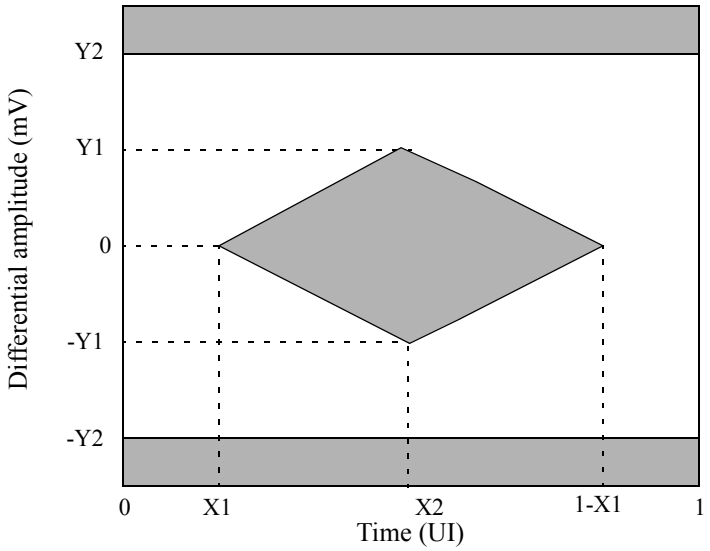

Figure 83A–9—Receiver template

83A.3.5.3 Differential input return loss

Differential input return loss shall meet the requirements defined in Table 83A–2. Differential input return loss is given in Equation (83A–7) and is illustrated in Figure 83A–10. Differential input return loss includes contributions from on-chip circuitry, chip packaging, and any off-chip components related to the receiver. The reference impedance for differential return loss measurements is 100 Ω.

$$Return\_loss(f) \geq \begin{cases} 12 & 0.01 \leq f < 2.125 \\ 6.5 - 13.33\log_{10}\left(\frac{f}{5.5}\right) & 2.125 \leq f \leq 11.1 \end{cases} \text{ (dB)} \quad \text{(83A–7)}$$

where

Return_loss(*f*) is the differential input return loss at frequency *f*
f is the frequency in GHz

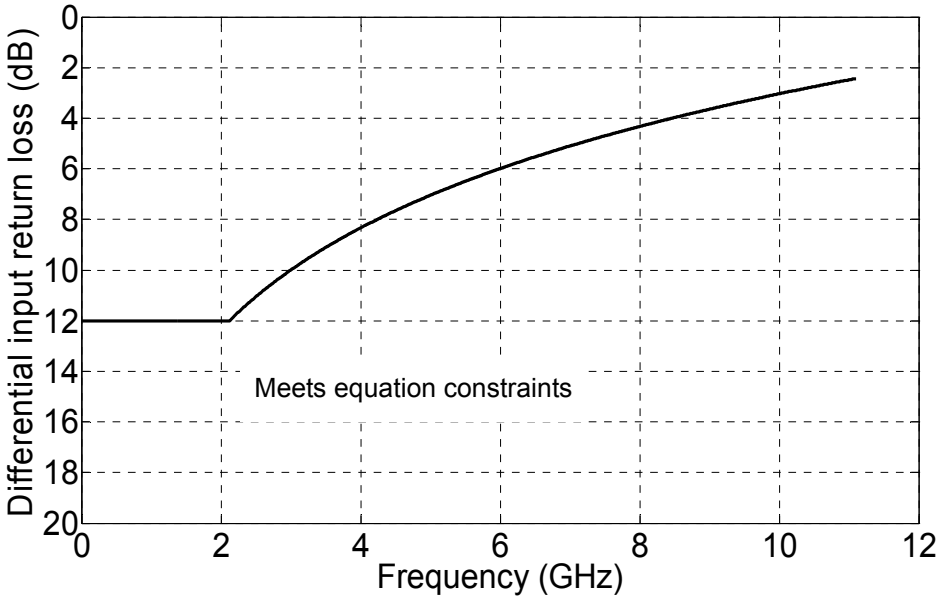

Figure 83A–10—Differential input return loss

83A.3.5.4 Differential to common-mode input return loss

Differential to common-mode input return loss shall meet the requirements defined in Table 83A–2. Differential to common-mode input return loss is given in Equation (83A–8) and is illustrated in Figure 83A–11. Differential to common-mode input return loss includes contributions from on-chip circuitry, chip packaging, and any off-chip components related to the receiver. The reference impedance for differential return loss measurements is 100 Ω. The reference impedance for common-mode return loss measurements is 25 Ω.

$$Return\_loss(f) \geq 15 \qquad 0.01 \leq f < 11.1 \qquad \text{(dB)} \quad \text{(83A–8)}$$

where

Return_loss(*f*) is the differential to common-mode input return loss at frequency *f*
f is the frequency in GHz

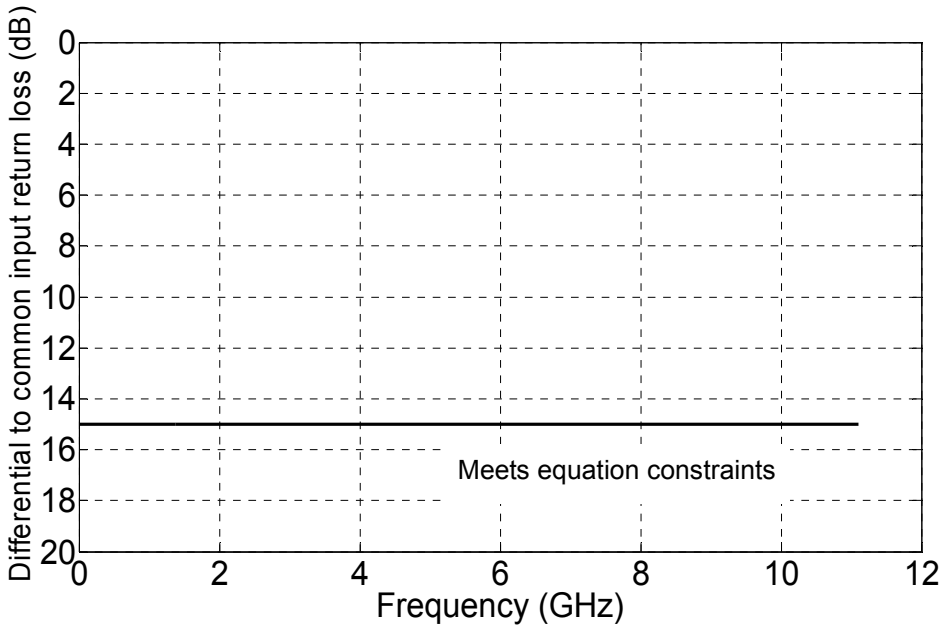

Figure 83A–11—Differential to common-mode input return loss

83A.3.5.5 AC-coupling

The XLAUI/CAUI-10 receiver shall be AC-coupled to the XLAUI/CAUI-10 transmitter 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 except when interfacing with modules defined in Annex 83B or explicitly stated otherwise. It should be noted that there may be various methods for AC-coupling in actual implementations.

83A.3.5.6 Jitter tolerance

The XLAUI/CAU-10 receiver shall have a peak-to-peak total jitter amplitude tolerance of at least the minimum total input jitter tolerance defined in Table 83A–2. This total jitter is composed of two components: deterministic jitter and random jitter. Deterministic jitter tolerance shall be at least the minimum deterministic input jitter tolerance defined in Table 83A–2. The XLAUI/CAUI-10 receiver shall tolerate sinusoidal jitter with any frequency and amplitude defined by the mask of Figure 83A–12. This sub-component of deterministic jitter is intended to ensure margin for low-frequency jitter, wander, noise, crosstalk, and other variable system effects.

83A.3.5.7 Global energy detect function

The global energy detect function is mandatory for EEE capability with the deep sleep mode option and XLAUI/CAUI-10 shutdown. The global energy detect function indicates whether or not signaling energy is being received on the physical instantiation of the inter sublayer interface (in each direction as appropriate). The energy detection function may be considered a subset of the signal indication logic. If no energy is being received on the XLAUI/CAUI-10 for the ingress direction, SIGNAL_DETECT is set to FAIL following a transition from aui_rx_mode = DATA to aui_rx_mode = QUIET. When aui_rx_mode = QUIET, SIGNAL_DETECT shall be set to OK within 500 ns following the application of a signal at the receiver input that corresponds to an ALERT signal driven from the XLAUI/CAUI-10 link partner. While aui_rx_mode = QUIET, SIGNAL_DETECT changes from FAIL to OK only after the valid ALERT signal is received.

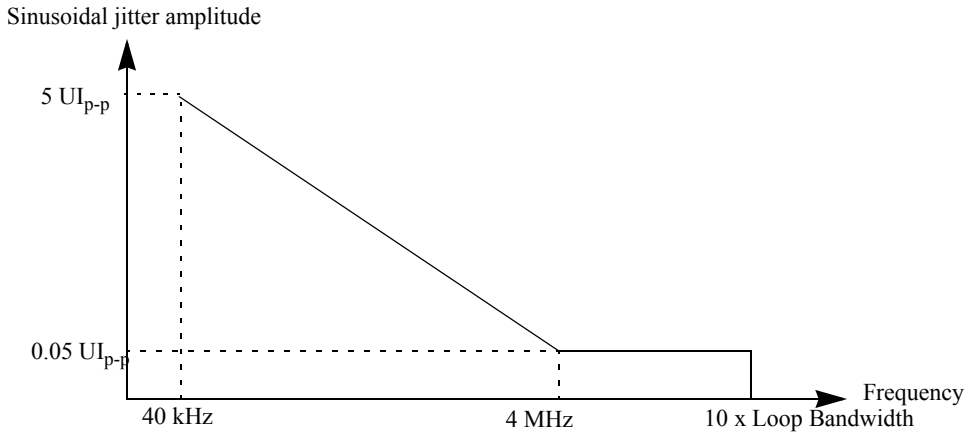

Figure 83A–12—Single-tone sinusoidal jitter mask

83A.4 Interconnect characteristics

This subclause describes recommended characteristics which are used to characterize a XLAUI/CAUI-10 channel as shown in Figure 83A–2. The value for differential insertion loss is summarized in Equation (83A–9) and illustrated in Figure 83A–13. The value for minimum return loss is summarized in Equation (83A–10) and illustrated in Figure 83A–14. The channel is terminated with 100 Ω differential impedance. Other impairments such as crosstalk can have a material impact on the link performance and should be minimized.

$$Insertion\_loss(f) < \left\{ \begin{array}{ll} 0.15 + 1.39\sqrt{f} + 1.4f & 0.01 \le f < 7 \\ -15.86 + 4.2f & 7 \le f \le 11.1 \end{array} \right\} \text{ (dB)} \quad \text{(83A–9)}$$

where

Insertion_loss(*f*) is the differential insertion loss at frequency *f*
f is the frequency in GHz

$$Return\_loss(f) \ge \left\{ \begin{array}{ll} 12.5 & 0.01 \le f < 5 \\ 12.5 - 27.5\log_{10}\left(\frac{f}{5}\right) & 5 \le f \le 11.1 \end{array} \right\} \text{ (dB)} \quad \text{(83A–10)}$$

where

Return_loss(*f*) is the differential input return loss at frequency *f*
f is the frequency in GHz

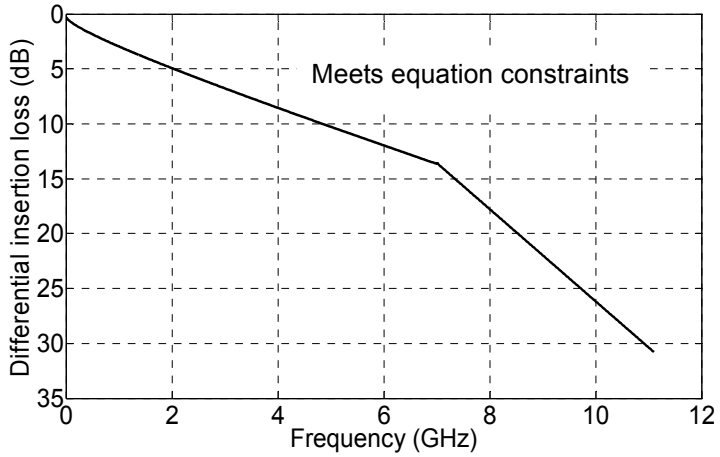

Figure 83A–13—Differential insertion loss

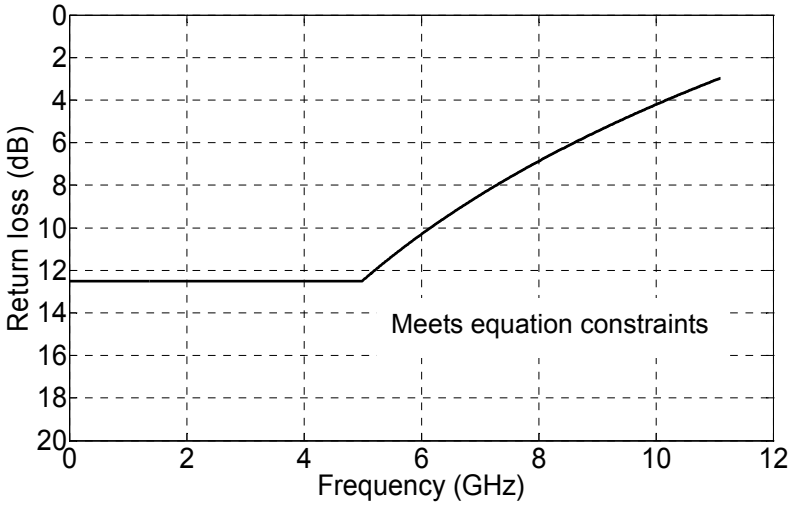

Figure 83A–14—Differential input return loss

The XLAUI/CAUI-10 is primarily intended as a point-to-point interface of up to approximately 25 cm between integrated circuits using controlled impedance traces on low-cost printed circuit boards (PCBs). Longer reaches for the XLAUI/CAUI-10 may be achieved by the use of better PCB materials, as the performance of an actual XLAUI/CAUI-10 interconnect is highly dependent on the implementation.

83A.4.1 Characteristic impedance

The recommended differential characteristic impedance of circuit board trace pairs is 100 Ω ± 10%.

83A.5 Electrical parameter measurement methods

This subclause describes the measurement methodology that is to be used to verify XLAUI/CAUI-10 compliance. Eye templates are measured with AC-coupling and centered at 0 V differential. The signal waveform, eye, and jitter may be measured using a receiver with at least an equivalent 12 GHz low-pass filter response. Jitter values and eye masks are specified for BER 10^{-12}.

83A.5.1 Transmit jitter

Transmit jitter is defined with respect to a test procedure resulting in a BER bathtub curve such as that described in Annex 48B.3. For the purpose of jitter measurement, the effect of a single-pole high-pass filter with a 3 dB point at 4 MHz is applied to the jitter. The test pattern for jitter measurements shall be PRBS31 test pattern in 83.5.10 or scrambled idle in 82.2.11. Crossing times are defined with respect to the mid-point (0 V) of the AC-coupled differential signal. De-emphasis shall be off during jitter testing. Transmit de-emphasis off state is defined by any setting that gives optimal performance for transmitter jitter and eye mask evaluation. All XLAUI/CAUI-10 transmitter lanes shall be active and all XLAUI/CAUI-10 receive lanes shall be receiving maximum amplitude and fastest rise time (as defined in Table 83A–1) during transmit jitter testing to ensure maximum lane-lane crosstalk is included in the jitter evaluation.

83A.5.2 Receiver tolerance

The XLAUI/CAUI-10 jitter tolerance test setup in Figure 83A–15 or its equivalent shall meet the receiver eye mask defined in Table 83A–2. Applied jitter is measured using the methodology described in Annex 48B.3. Deterministic jitter is added to a clean test pattern by adding sinusoidal jitter as defined in 83A.3.5.6, along with low-pass filter stress, followed by a limiting function, and frequency-dependent attenuation. The low-pass filter stress is added until the 0.34 UI Deterministic Jitter is achieved. Frequency-dependent attenuation is then added using PCB trace or frequency-dependent attenuation which emulates PCB loss. Frequency-dependent attenuation is added until 0.42 UI Deterministic Jitter is achieved. Random jitter is added to the test signal using an interference generator which is a broadband noise source capable of producing white Gaussian noise with adjustable amplitude. The power spectral density shall be flat to ±3 dB from 50 MHz to 6 GHz with a crest factor of no less than 5. The amplitude and output jitter of the filter stress plus limiter and random jitter injection shall meet the receiver eye mask defined in Table 83A–2. All XLAUI/CAUI-10 lanes shall be active during jitter tolerance testing. The PRBS31 pattern defined in 83.5.10 or scrambled idle defined in 82.2.11 is used for evaluating XLAUI/CAUI-10 jitter tolerance.

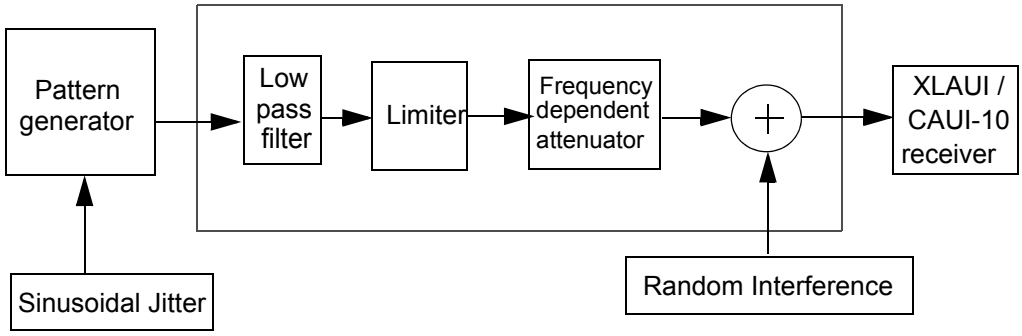

Figure 83A–15—Stressed-eye and jitter tolerance test setup

83A.6 Environmental specifications

83A.6.1 General safety

All equipment subject to this clause shall conform to applicable sections (including isolation requirements) of IEC 60950-1.

83A.6.2 Network safety

The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

83A.6.3 Installation and maintenance guidelines

It is recommended that sound installation practice, as defined by applicable local codes and regulations, be followed in every instance in which such practice is applicable.

83A.6.4 Electromagnetic compatibility

A system integrating the XLAUI/CAUI-10 shall comply with applicable local and national codes for the limitation of electromagnetic interference.

83A.6.5 Temperature and humidity

A system integrating the XLAUI/CAUI-10 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.

83A.7 Protocol implementation conformance statement (PICS) proforma for Annex 83A, 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s ten-lane Attachment Unit Interface (CAUI-10)[26]

83A.7.1 Introduction

The supplier of a XLAUI/CAUI-10 implementation that is claimed to conform to Annex 83A, 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s ten-lane Attachment Unit Interface (CAUI-10), 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.

83A.7.2 Identification

83A.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).	

83A.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Annex 83A, 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s ten-lane Attachment Unit Interface (CAUI-10)
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	

[26]*Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this annex so that it can be used for its intended purpose and may further publish the completed PICS.

83A.7.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
NOL	XLAUI is organized into four lanes, CAUI-10 is organized into ten lanes	83A.1.2	See Clause 83	M	Yes [ ]
RATE	Each XLAUI/CAUI-10 lane operates at 10.3125 Gb/s	83A.1.2	10.3125 Gb/s (nominal)	M	Yes [ ]
IO	Meets XLAUI/CAUI-10 Electrical Characteristics	83A.3	Supports transmit and receive compliance points	M	Yes [ ]
*LPI	Support for CAUI-10 shutdown	83A.3.3		O	Yes [ ] No [ ]

83A.7.4 XLAUI/CAUI-10 transmitter requirements

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	Signaling rate	83A.3.4	10.3125 GBd ±100 ppm	M	Yes [ ]
TC2	Single ended output voltage range	83A.3.4.1	–0.4 V to 4 V	M	Yes [ ]
TC3	Minimum de-emphasis	83A.3.4.1	4.4 dB	M	Yes [ ]
TC4	Maximum de-emphasis	83A.3.4.1	7 dB	M	Yes [ ]
TC5	Maximum termination mismatch	83A.3.4.3	5%	M	Yes [ ]
TC6	Maximum differential output voltage	83A.3.4.1	760 mV	M	Yes [ ]
TC7	Minimum output rise and fall time (20% to 80%)	83A.3.4.2	24 ps	M	Yes [ ]
TC8	Differential output return loss	83A.3.4.3	See Equation (83A–5)	M	Yes [ ]
TC9	Common-mode output return loss	83A.3.4.4	See Equation (83A–6)	M	Yes [ ]
TC10	Transmitter eye mask	83A.3.4.5	See Figure 83A–8	M	Yes [ ]
TC11	Amplitude and swing for XLAUI/CAUI-10 shutdown	83A.3.4.1.1		LPI:M	Yes [ ]
TC12	Transmit disable for XLAUI/CAUI-10 shutdown	83A.3.4.6		LPI:M	Yes [ ]

83A.7.5 XLAUI/CAUI-10 receiver requirements

Item	Feature	Subclause	Value/Comment	Status	Support
RC1	Receiver BER	83A.3.5.1	Better than 10^{-12}	M	Yes [ ]
RC2	Differential input return loss	83A.3.5.3	See Equation (83A–7)	M	Yes [ ]
RC3	Differential to common-mode input return loss	83A.3.5.4	See Equation (83A–8)	M	Yes [ ]
RC4	Total Jitter Tolerance	83A.3.5.6	0.62 UI	M	Yes [ ]
RC5	Deterministic Jitter Tolerance	83A.3.5.6	0.42 UI	M	Yes [ ]
RC6	Sinusoidal Jitter Tolerance	83A.3.5.6	See Figure 83A–12	M	Yes [ ]
RC7	Receiver AC-coupling	83A.3.5.5	Present	M	Yes [ ]
RC8	Signal detect for XLAUI/CAUI-10 shutdown	83A.3.5.7		LPI:M	Yes [ ]

83A.7.6 Electrical measurement methods

Item	Feature	Subclause	Value/Comment	Status	Support
EM1	De-emphasis setting during jitter measurements	83A.5.1	Off	M	Yes [ ]
EM2	Jitter Tolerance stressed input	83A.5.2	0.42 UI DJ, 0.2 UI RJ	M	Yes [ ]
EM3	Random jitter interference generator	83A.5.2	±3 dB from 50 MHz to 6 GHz with a crest factor of no less than 5	M	Yes [ ]
EM4	Meet jitter tolerance requirement with all XLAUI/CAUI-10 channels active	83A.5.2	Yes	M	Yes [ ]
EM5	XLAUI/CAUI-10 Jitter Tolerance Test Pattern	83A.5.2	PRBS31 or scrambled idle	M	Yes [ ]

83A.7.7 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
ES1	Conformance to applicable sections of IEC 60950-1	83A.6.1	Yes	M	Yes [ ]
ES2	Compliance with applicable local and national codes for the limitation of electromagnetic interference	83A.6.4	Yes	M	Yes [ ]

Annex 83B

(normative)

Chip-to-module 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s ten-lane Attachment Unit Interface (CAUI-10)

83B.1 Overview

This annex defines the functional and electrical characteristics for the optional chip-to-module 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s ten-lane Attachment Unit Interface (CAUI-10). The purpose of this annex is to provide electrical characteristics and associated compliance points for pluggable module applications that use the XLAUI/CAUI-10 interface and shall use the same number of lanes and signaling rate defined in Annex 83A. Figure 83B–3 and Table 83B–1 summarize an example differential insertion loss budget associated with the chip-to-module application. The insertion loss of Equation (83A–9), excluding a 0.5 dB connector loss at 5.15625 GHz, is linearly scaled to 7.9 dB loss at 5.15625 GHz for the host XLAUI/CAUI-10 component, and 2.1 dB loss at 5.15625 GHz for the module as per Table 83B–1 and Equation (83B–1) for the host and Equation (83B–2) for the module. Tradeoffs can be made between the host PCB insertion loss and the connector loss. Equation (83B–1) is illustrated in Figure 83B–1 and Equation (83B–2) is illustrated in Figure 83B–2.

$$Insertion\_loss(f) \le \left\{ \begin{array}{ll} 0.111 + 1.046\sqrt{f} + 1.05f & 0.01 \le f < 7 \\ -11.82 + 3.15f & 7 \le f \le 11.1 \end{array} \right\} \text{ (dB)} \qquad \text{(83B–1)}$$

where

Insertion_loss(*f*) is the differential insertion loss at frequency *f*
f is the frequency in GHz

$$Insertion\_loss(f) \le \left\{ \begin{array}{ll} 0.03 + 0.278\sqrt{f} + 0.28f & 0.01 \le f < 7 \\ -3.155 + 0.84f & 7 \le f \le 11.1 \end{array} \right\} \text{ (dB)} \qquad \text{(83B–2)}$$

where

Insertion_loss(*f*) is the differential insertion loss at frequency *f*
f is the frequency in GHz

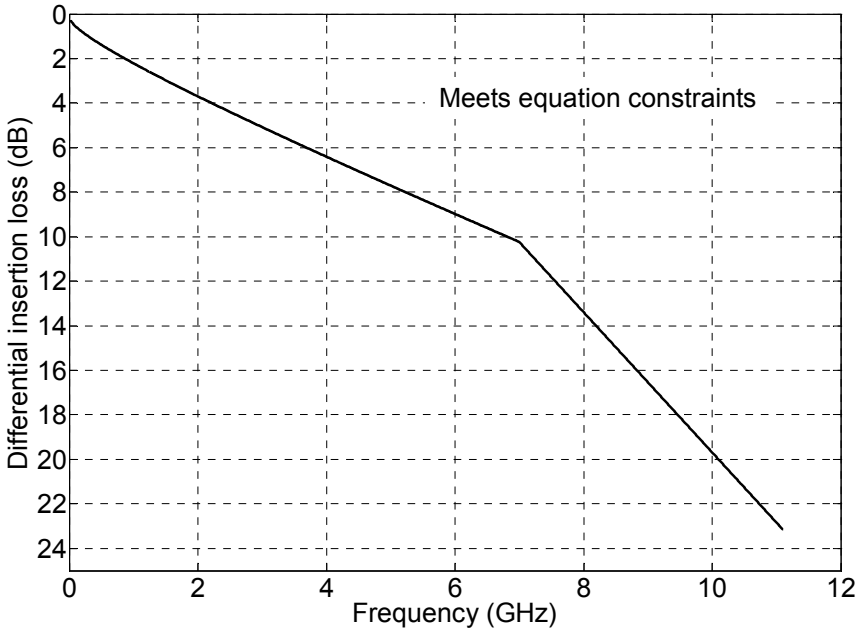

Figure 83B–1—Host insertion loss

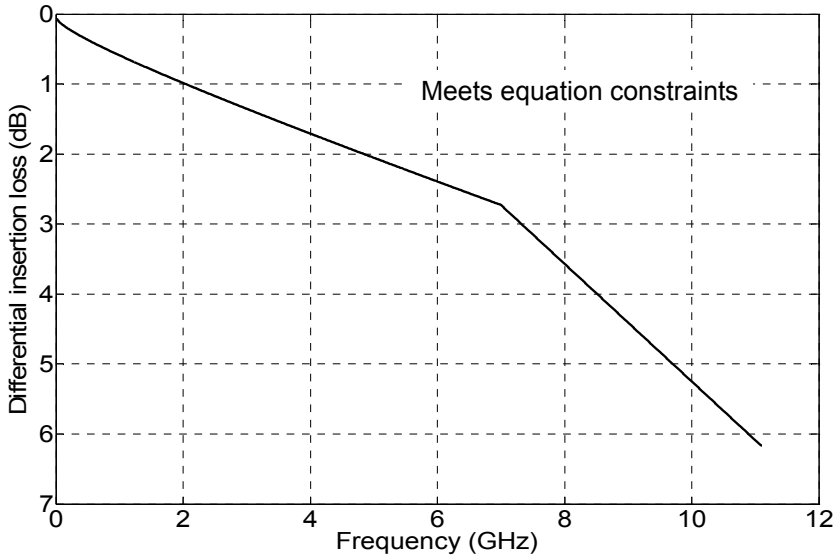

Figure 83B–2—Module insertion loss

.

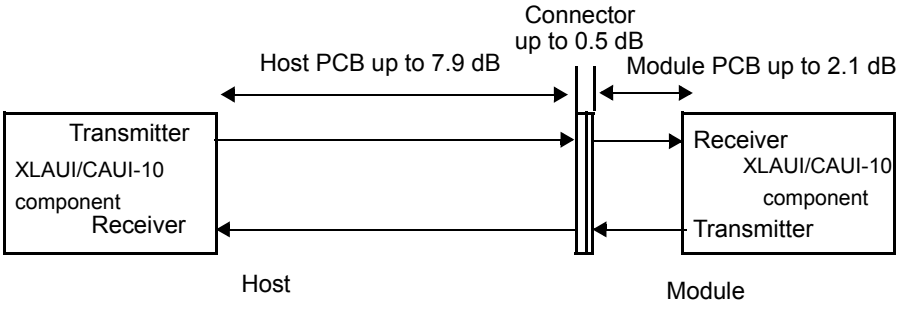

Figure 83B–3—Chip-to-module loss budget at 5.15625 GHz

Table 83B–1—Chip-to-module loss budget

Section	Differential insertion loss max. (at 5.15625 GHz)
Host XLAUI/CAUI-10 component to connector	7.9 dB
Connector loss	0.5 dB
Connector to module XLAUI/CAUI-10 component	2.1 dB

83B.2 Compliance point specifications for chip-to-module XLAUI/CAUI-10

The chip-to-module XLAUI/CAUI-10 interface specifies compliance points around the module connector as depicted in Figure 83B–5 and Figure 83B–7. Chip-to-module devices shall meet the electrical characteristics defined in 83B.2, 83B.2.1, 83B.2.2, and 83B.2.3. Compliance points are defined to ensure interoperability between hosts and modules. A Module Compliance Board (MCB) is used to isolate the characteristics of the module, and a Host Compliance Board (HCB) is used to isolate the characteristics of the host. Figure 83B–5 and Figure 83B–7 include the loss associated with the HCB and MCB at 5.15625 GHz.

The reference differential insertion loss of the HCB PCB is given in Equation (83B–3) and illustrated in Figure 83B–4. The effects of differences between the insertion loss of an actual HCB and the reference insertion loss are to be accounted in the measurements.

$$Insertion\_loss(f) = 0.017 + 0.5\sqrt{f} + 0.1836f \qquad 0.01 \le f \le 11.1 \qquad \text{(83B–3)}$$

where

Insertion_loss(*f*) is the differential insertion loss at frequency *f*
f is the frequency in GHz

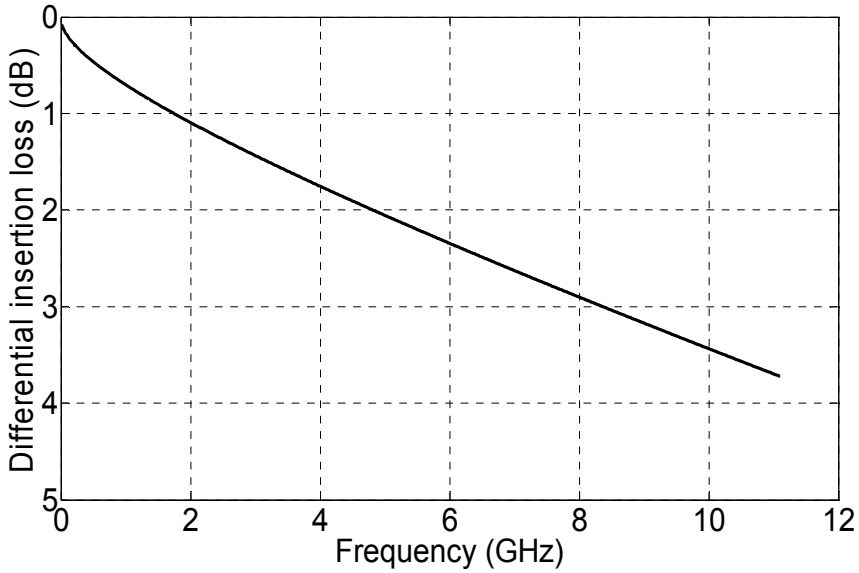

Figure 83B–4—HCB insertion loss

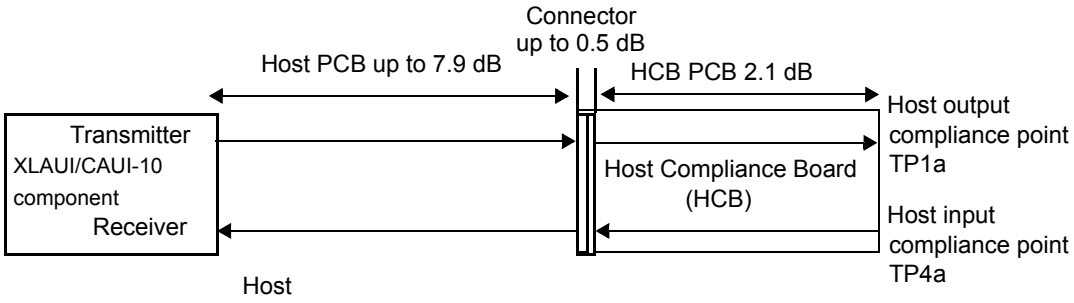

Figure 83B–5—Chip-to-module HCB insertion loss budget at 5.15625 GHz

The reference differential insertion loss of the MCB PCB is given in Equation (83B–4) and illustrated in Figure 83B–6. The effects of differences between the insertion loss of an actual MCB and the reference insertion loss are to be accounted in the measurements.

$$Insertion\_loss(f) = -0.00086 + 0.2286\sqrt{f} + 0.08386f \quad \text{dB} \qquad 0.01 \le f \le 11.1 \qquad (83B–4)$$

where

Insertion_loss(f) is the differential insertion loss at frequency *f*
f is the frequency in GHz

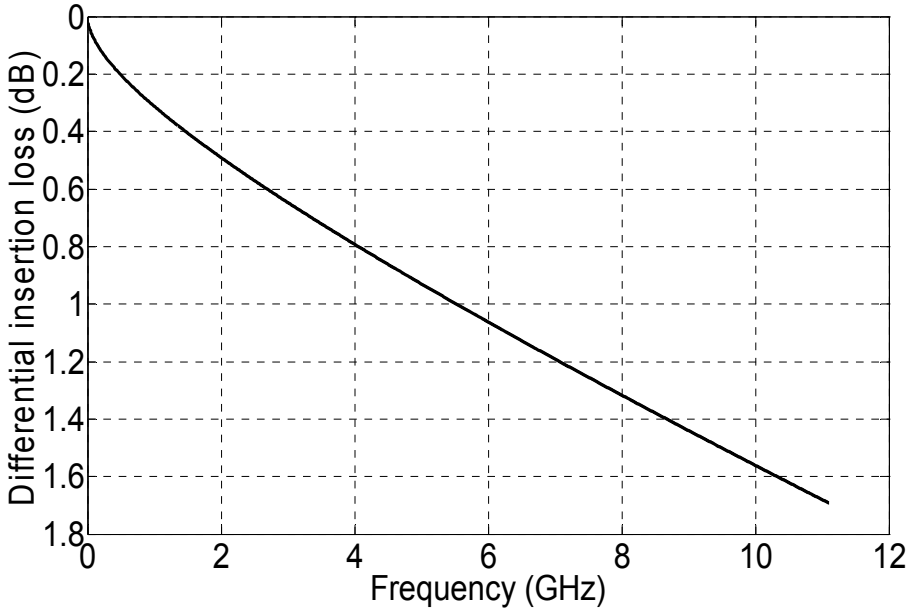

Figure 83B–6—MCB insertion loss

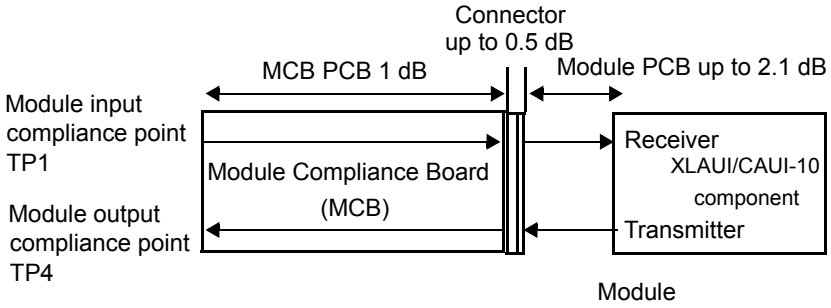

Figure 83B–7—Chip-to-module with MCB insertion loss budget at 5.15625 GHz

83B.2.1 Module specifications

A module that uses XLAUI/CAUI-10 to interface with a host shall meet the characteristics outlined in Table 83B–2 and Table 83B–3 when measured using the MCB and HCB (where the HCB is used to calibrate inputs to the module). Table 83B–2 also lists the equivalent test points for the XLPPI/CPPI (see Figure 86–3).

Modules may support additional de-emphasis states, but the specification of additional states is outside the scope of this standard. De-emphasis shall be off during eye mask and jitter testing. Module electrical output de-emphasis off state is the optimal setting for module electrical output jitter and eye mask evaluation. AC-coupling for both Tx and Rx paths shall be located in the module.

Table 83B–2—Specifications at module compliance points

Reference	Compliance point		Value	Unit
Minimum module differential input return loss	MCB input	TP1	See Equation (83B–5)	dB
Module input tolerance signal	HCB output	TP1a	See 83A.5.2	
Module output signal	MCB output	TP4	See Table 83B–3	
Minimum module differential output return loss	MCB output	TP4	See Equation (83B–5)	dB

$$Return\_loss(f) \geq \begin{cases} 12 - 2\sqrt{f} & 0.01 \leq f < 2.19 \\ 5.56 - 8.7\log_{10}\left(\frac{f}{5.5}\right) & 2.19 \leq f \leq 11.1 \end{cases} \text{ (dB)} \qquad (83B–5)$$

where

Return_loss(*f*) is the differential input return loss at frequency *f*
f is the frequency in GHz

Minimum module differential input/output return loss is illustrated in Figure 83B–8.

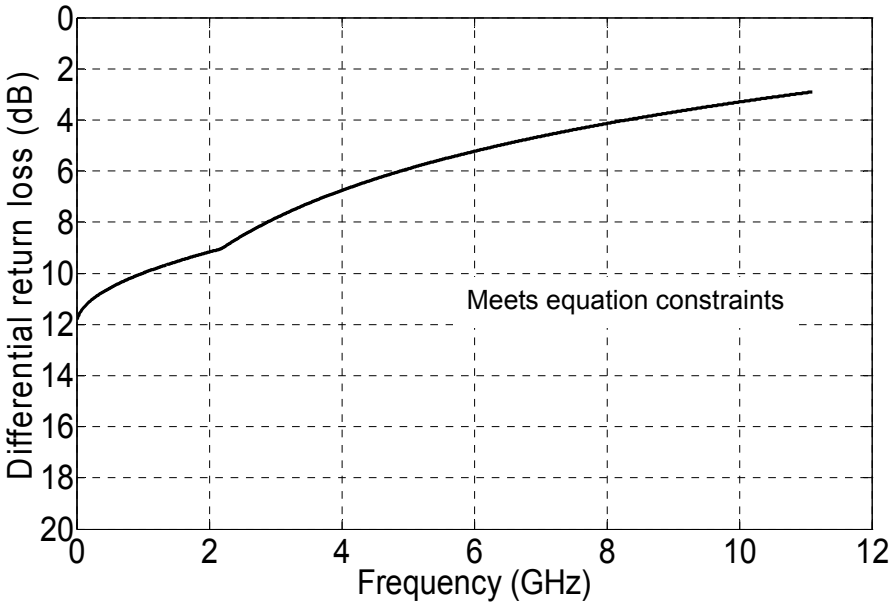

Figure 83B–8—Module input/output return loss

Table 83B–3—Module electrical output

Parameter	Subclause reference	Value	Unit
Maximum differential output voltage, peak-to-peak	83A.3.4.1	760	mV
Minimum de-emphasis	83A.3.4.1	3.5	dB
Maximum de-emphasis	83A.3.4.1	6	dB
Minimum VMA	83A.3.4.1	See Equation (83B–6)	mV
Maximum termination mismatch at 1 MHz	86A.5.3.2	5	%
Maximum output AC common-mode voltage, RMS	86A.5.3.1	15	mV
Minimum output rise and fall time (20% to 80%)	83A.3.4.2	24	ps
Maximum Total Jitter	83A.3.4.5	0.4	UI
Maximum Deterministic Jitter	83A.3.4.5	0.25	UI
Module electrical output eye mask definition X1	83A.3.4.5	0.2	UI
Module electrical output eye mask definition X2	83A.3.4.5	0.5	UI
Module electrical output eye mask definition Y1	83A.3.4.5	136	mV
Module electrical output eye mask definition Y2	83A.3.4.5	380	mV

$$\text{Minimum VMA (mV)} = (-110 - 2.13x + 0.32x^2) \times (10^{-y/20}) \quad \text{(83B–6)}$$

where

x is the rise or fall time, or 38, (whichever is larger) in ps
y is de-emphasis value in dB

83B.2.2 Host specifications

A host that uses XLAUI/CAUI-10 to interface with a module shall meet the characteristics outlined in Table 83B–4 and Table 83B–5 when measured using the HCB and MCB (where MCB is used to calibrate inputs to the host). Table 83B–4 also lists the equivalent test points for the XLPPI/CPPI.

Table 83B–4—Specifications at host compliance points

Reference	Compliance point		Value	Unit
Host output signal specifications	HCB output	TP1a	See Table 83B–5	
Minimum host differential output return loss	HCB output	TP1a	See Equation (83B–7)	dB
Minimum host differential input return loss	HCB input	TP4a	See Equation (83B–7)	dB
Host input tolerance signal	MCB output	TP4	See 83B.2.3	

$$Return\_loss(f) \geq \begin{cases} 12 - 2\sqrt{f} & 0.01 \leq f < 2.19 \\ 5.56 - 8.7\log_{10}\left(\frac{f}{5.5}\right) & 2.19 \leq f \leq 11.1 \end{cases} \text{ (dB)} \quad \text{(83B–7)}$$

where

$Return\_loss(f)$ is the differential input/output return loss at frequency f
f is the frequency in GHz

Minimum host differential input/output return loss is illustrated in Figure 83B–9.

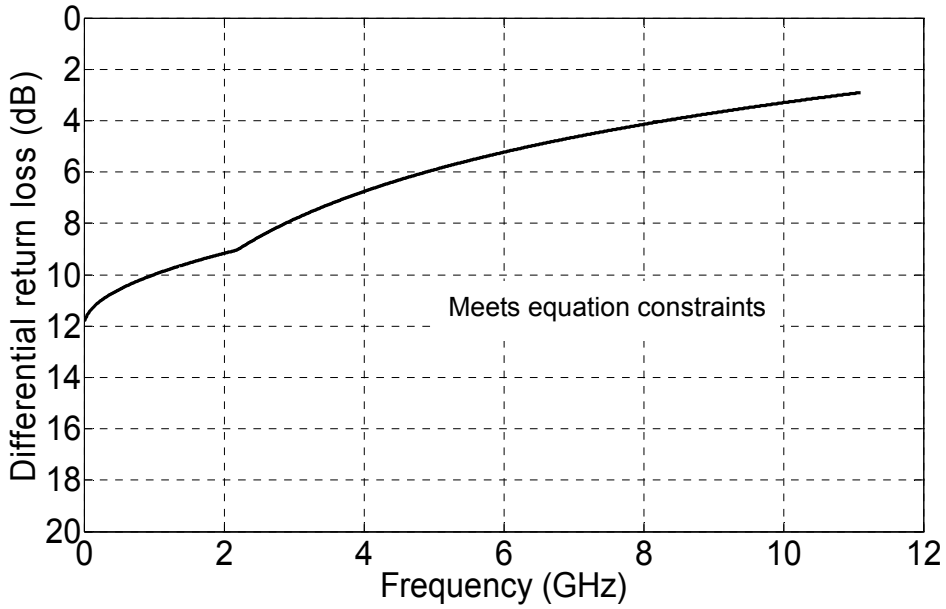

Figure 83B–9—Host input/output return loss

Table 83B–5—Host electrical output

Parameter	Subclause reference	Value	Units
Maximum output AC common-mode voltage, RMS	86A.5.3.1	20	mV
Minimum output rise and fall time (20% to 80%)	83A.3.4.2	24	ps
Maximum Total Jitter	83A.3.4.5	0.62	UI
Maximum Deterministic Jitter	83A.3.4.5	0.42	UI
Host electrical output eye mask definition X1	83A.3.4.5	0.31	UI
Host electrical output eye mask definition X2	83A.3.4.5	0.5	UI
Host electrical output eye mask definition Y1	83A.3.4.5	42.5	mV
Host electrical output eye mask definition Y2	83A.3.4.5	425	mV

83B.2.3 Host input signal tolerance

Host XLAUI/CAUI-10 jitter tolerance evaluation shall be defined by a stressed input signal that comprises 0.25 UI Deterministic Jitter, and 0.15 UI random jitter for BER 10^{-12}. Deterministic Jitter is added to a clean test pattern as sinusoidal jitter defined in 83A.3.5.6. The limited low-pass filter stress is added until the 0.22 UI Deterministic Jitter is achieved. Frequency-dependent attenuation is then added using PCB trace or frequency-dependent attenuation that emulates PCB loss. Frequency-dependent attenuation is added until 0.25 UI Deterministic Jitter is achieved. Random jitter is added to the test signal using an interference generator, which is a broadband noise source capable of producing white Gaussian noise with adjustable

amplitude. The power spectral density shall be flat to ±3 dB from 50 MHz to 6 GHz with a crest factor of no less than 5. Figure 83B–10 depicts a XLAUI/CAUI-10 jitter tolerance test setup. The amplitude and output jitter of the filter stress plus limiter and random jitter injection shall meet the receiver eye mask illustrated in Figure 83A–9 with the values for X1, X2, Y1, Y2 given in Table 83B–3. All XLAUI/CAUI-10 lanes shall be active during jitter tolerance testing. The PRBS31 pattern defined in 83.5.10 or scrambled idle defined in 82.2.11 is used for evaluating XLAUI/CAUI-10 jitter tolerance.

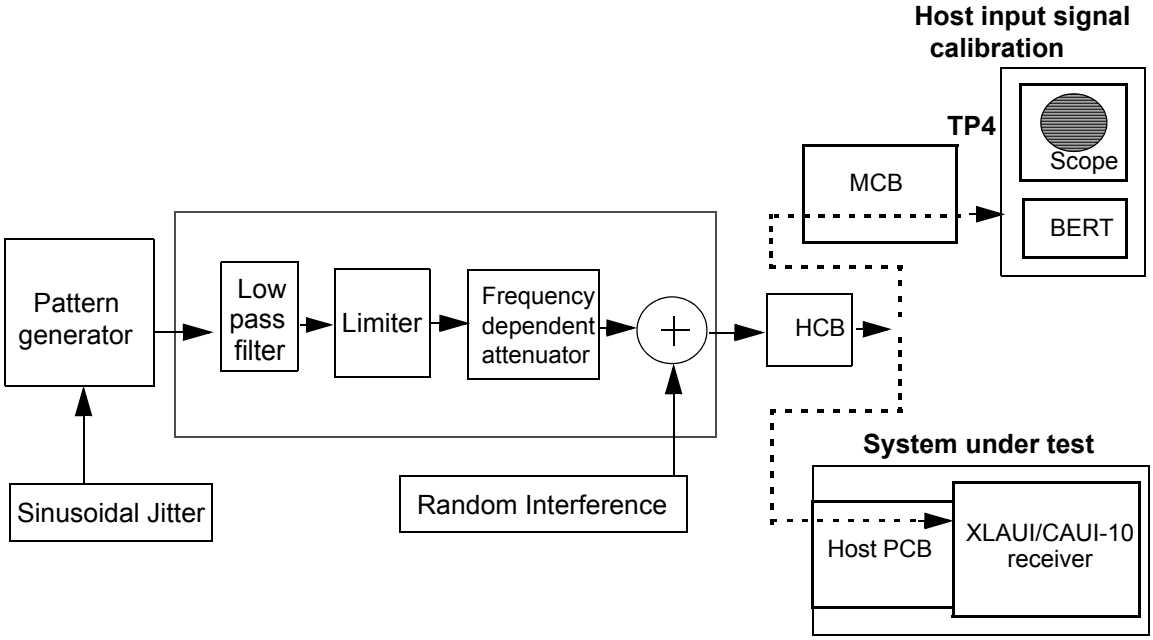

Figure 83B–10—Stressed-eye and jitter tolerance test setup

83B.3 Environmental specifications

83B.3.1 General safety

All equipment subject to this clause shall conform to applicable sections (including isolation requirements) of IEC 60950-1.

83B.3.2 Network safety

The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

83B.3.3 Installation and maintenance guidelines

It is recommended that sound installation practice, as defined by applicable local codes and regulations, be followed in every instance in which such practice is applicable.

83B.3.4 Electromagnetic compatibility

A system integrating the XLAUI/CAUI-10 shall comply with applicable local and national codes for the limitation of electromagnetic interference.

83B.3.5 Temperature and humidity

A system integrating the XLAUI/CAUI-10 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.

83B.4 Protocol implementation conformance statement (PICS) proforma for Annex 83B, Chip-to-module 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s ten-lane Attachment Unit Interface (CAUI-10)[27]

83B.4.1 Introduction

The supplier of a chip to module XLAUI/CAUI-10 implementation that is claimed to conform to Annex 83B, Chip-to-module 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s ten-lane Attachment Unit Interface (CAUI-10), 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.

83B.4.2 Identification

83B.4.2.1 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).	

83B.4.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Annex 83B, Chip-to-module 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s ten-lane Attachment Unit Interface (CAUI-10)
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	

[27] *Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this annex so that it can be used for its intended purpose and may further publish the completed PICS.

83B.4.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
NOL	XLAUI is organized into four lanes, CAUI-10 is organized into ten lanes	83A.1.2	See Clause 83	M	Yes [ ]
RATE	Each XLAUI/CAUI-10 lane operates at 10.3125 Gb/s	83A.1.2	10.3125 Gb/s (nominal)	M	Yes [ ]
IO	Meets chip-module XLAUI/CAUI-10 Electrical Characteristics	83B.2	Supports host/module compliance points	M	Yes [ ]

83B.4.4 Module requirements

Item	Feature	Subclause	Value/Comment	Status	Support
MC1	XLAUI/CAUI-10 compliant module	83B.2.1	Meets requirements defined in Table 83B–2 and Table 83B–3	M	Yes [ ]
MC2	De-emphasis setting during module jitter evaluation	83B.2.1	Off	M	Yes [ ]
MC3	AC-coupling for Tx and Rx	83B.2.1	Present in module	M	Yes [ ]

83B.4.5 Host requirements

Item	Feature	Subclause	Value/Comment	Status	Support
HC1	XLAUI/CAUI-10 compliant host	83B.2.2	Meets requirements defined in Table 83B–4 and Table 83B–5	M	Yes [ ]
HC2	All XLAUI/CAUI-10 lanes active during jitter evaluation	83B.2.3	Yes	M	Yes [ ]
HC3	Applied jitter for host jitter tolerance evaluation	83B.2.3	0.25UI Deterministic Jitter, 0.15UI random jitter	M	Yes [ ]

83B.4.6 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
ES1	Conformance to applicable sections of IEC 60950-1	83B.3.1	Yes	M	Yes [ ]
ES2	Compliance with applicable local and national codes for the limitation of electromagnetic interference	83B.3.4	Yes	M	Yes [ ]

Annex 83C

(informative)

PMA sublayer partitioning examples

The following subclauses provide various partitioning examples. Partitioning guidelines and MMD numbering conventions are described in 83.1.4.

83C.1 Partitioning examples with FEC

The example of BASE-R FEC (see Clause 74) implemented in a separate device from either the PCS or the PMD is illustrated in Figure 83–2.

83C.1.1 FEC implemented with PCS

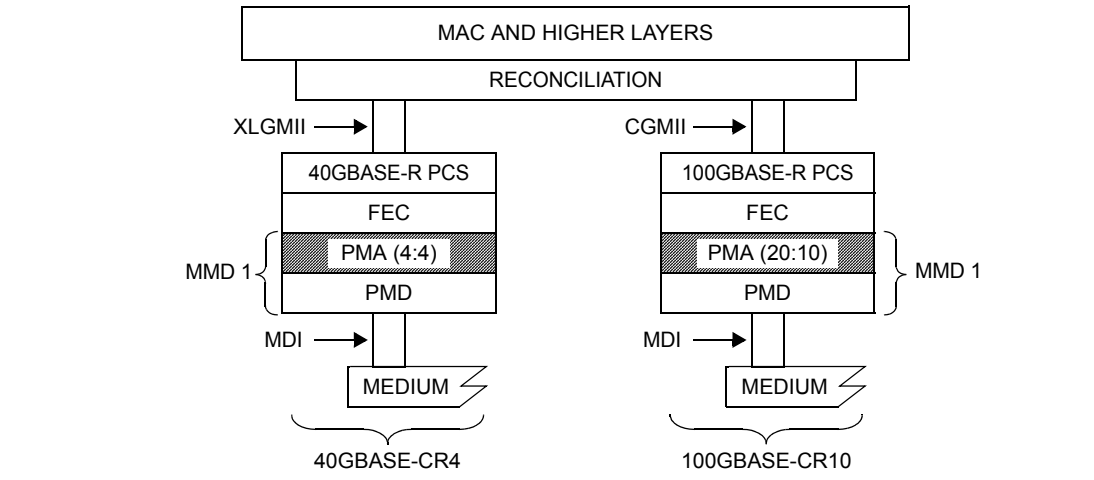

CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
MMD = MDIO MANAGEABLE DEVICE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

Figure 83C–1—Example FEC implemented with PCS

83C.1.2 FEC implemented with PMD

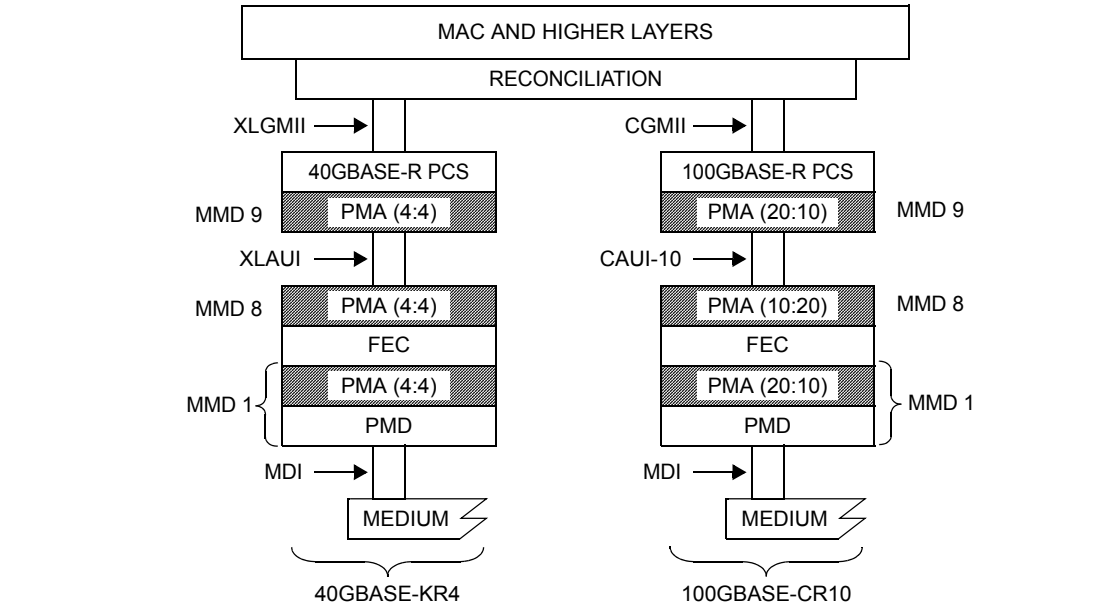

CAUI-10 = 100 Gb/s TEN-LANE ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLAUI = 40 Gb/s ATTACHMENT UNIT INTERFACE
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

Figure 83C–2—Example FEC implemented with PMD

83C.2 Partitioning examples with RS-FEC

83C.2.1 Single PMA sublayer with RS-FEC

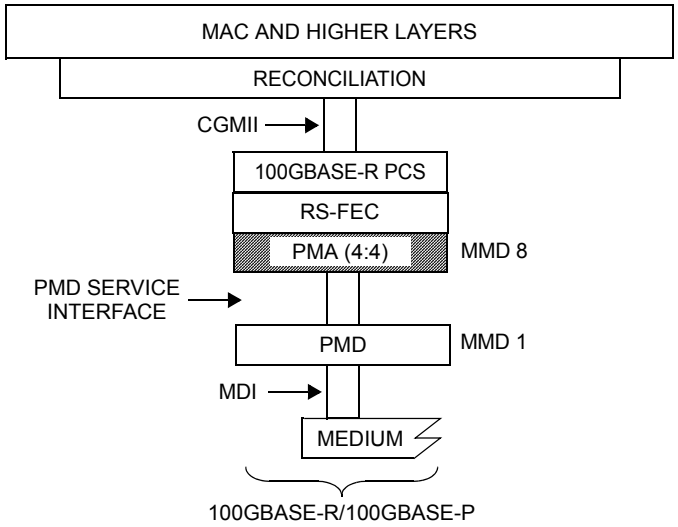

CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
MMD = MDIO MANAGEABLE DEVICE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION

Figure 83C–3—Example single PMA sublayer with RS-FEC

83C.2.2 Single CAUI-10 with RS-FEC

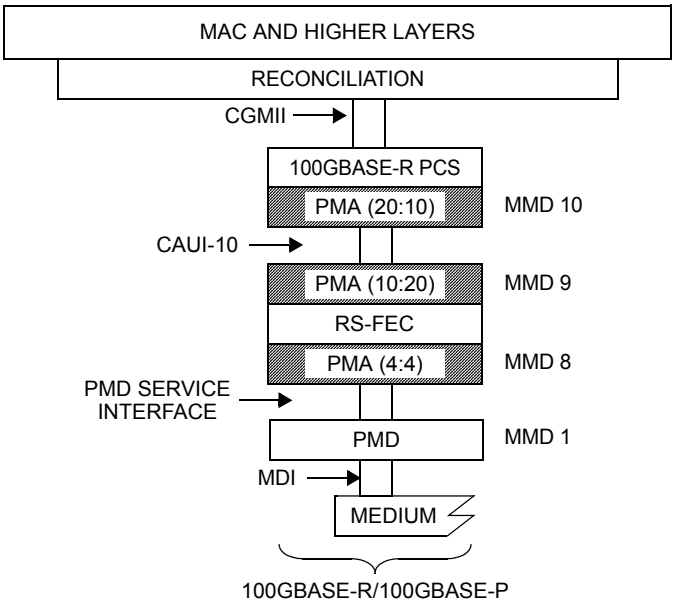

CAUI-10 = 100 Gb/s TEN-LANE ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
MMD = MDIO MANAGEABLE DEVICE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION

Figure 83C–4—Example single CAUI-10 with RS-FEC

83C.3 Partitioning examples without FEC

83C.3.1 Single PMA sublayer without FEC

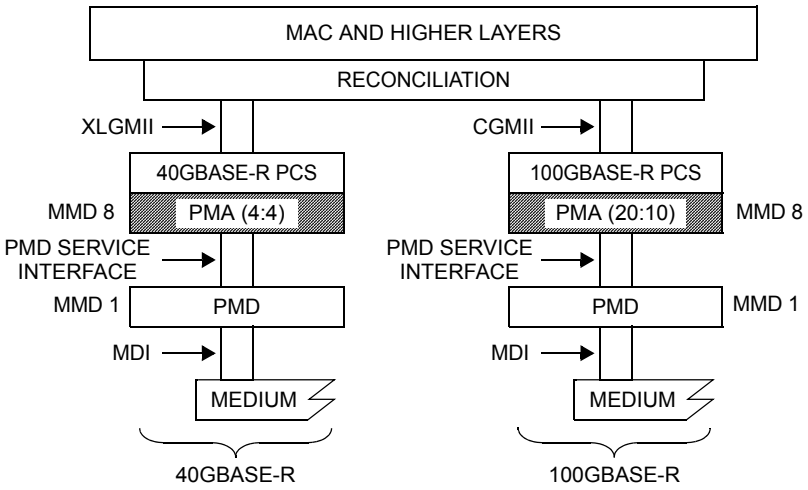

CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
MMD = MDIO MANAGEABLE DEVICE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

Figure 83C–5—Example single PMA sublayer without FEC

83C.3.2 Single XLAUI/CAUI-4 without FEC

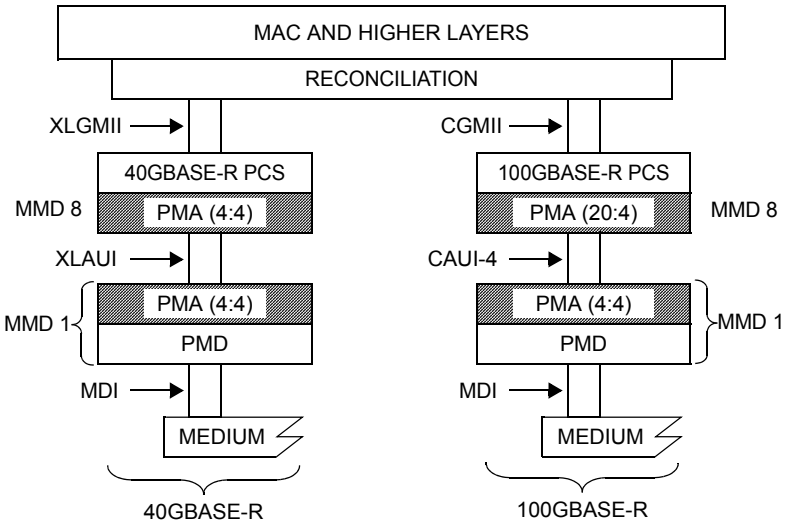

CAUI-4 = 100 Gb/s FOUR-LANE ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
MMD = MDIO MANAGEABLE DEVICE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLAUI = 40 Gb/s ATTACHMENT UNIT INTERFACE
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

Figure 83C–6—Example single XLAUI/CAUI-4 without FEC

83C.3.3 Separate SERDES for optical module interface

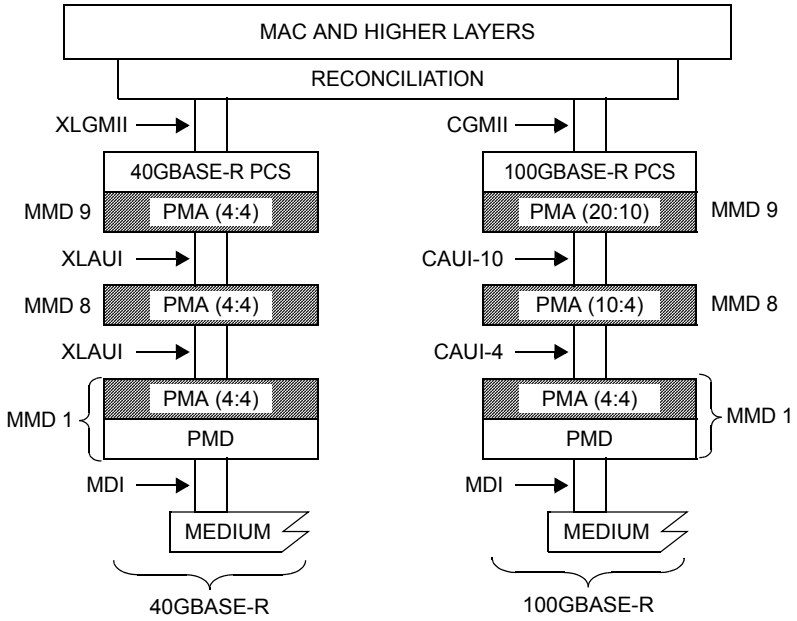

CAUI-4 = 100 Gb/s FOUR-LANE ATTACHMENT UNIT INTERFACE
CAUI-10 = 100 Gb/s TEN-LANE ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
XLAUI = 40 Gb/s ATTACHMENT UNIT INTERFACE
XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

Figure 83C–7—Separate SERDES for optical module interface

Annex 83D

(normative)

Chip-to-chip 100 Gb/s four-lane Attachment Unit Interface (CAUI-4)

83D.1 Overview

This annex defines the functional and electrical characteristics for the optional chip-to-chip 100 Gb/s four-lane Attachment Unit Interface (CAUI-4). Figure 83D–1 shows an example relationship of the CAUI-4 chip-to-chip interface to the ISO/IEC Open System Interconnection (OSI) reference model. The chip-to-chip interface provides electrical characteristics and associated compliance points, which can optionally be used when designing systems with electrical interconnect of approximately 25 cm in length.

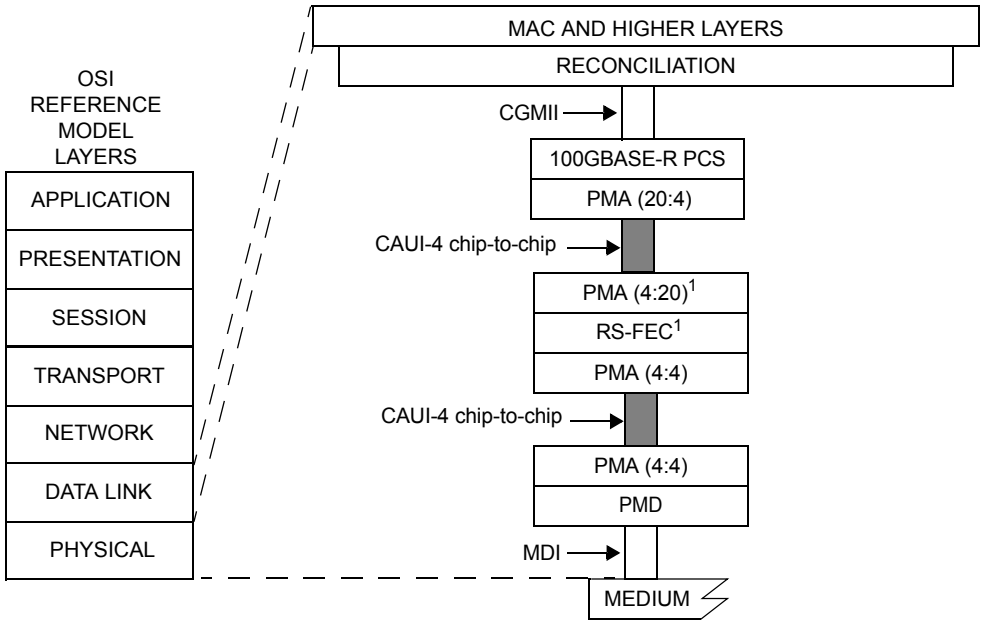

Figure 83D–1—Example CAUI-4 chip-to-chip relationship to the ISO/IEC Open System Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

The CAUI-4 bidirectional link is described in terms of a CAUI-4 transmitter, a CAUI-4 channel, and a CAUI-4 receiver. Figure 83D–2 depicts a typical CAUI-4 application, and Equation (83D–1) (illustrated in Figure 83D–3) summarizes the informative differential insertion loss budget associated with the chip-to-chip application. The CAUI-4 chip-to-chip interface comprises independent data paths in each direction. Each data path contains four differential lanes, which are AC-coupled. The nominal signaling rate for each lane is 25.78125 GBd. The CAUI-4 transmitter on each end of the link is adjusted to an appropriate setting

based on channel knowledge. If implemented, the transmitter equalization feedback mechanism described in 83D.3.3.2 may be used to identify an appropriate setting. The adaptive or adjustable receiver performs the remainder of the equalization.

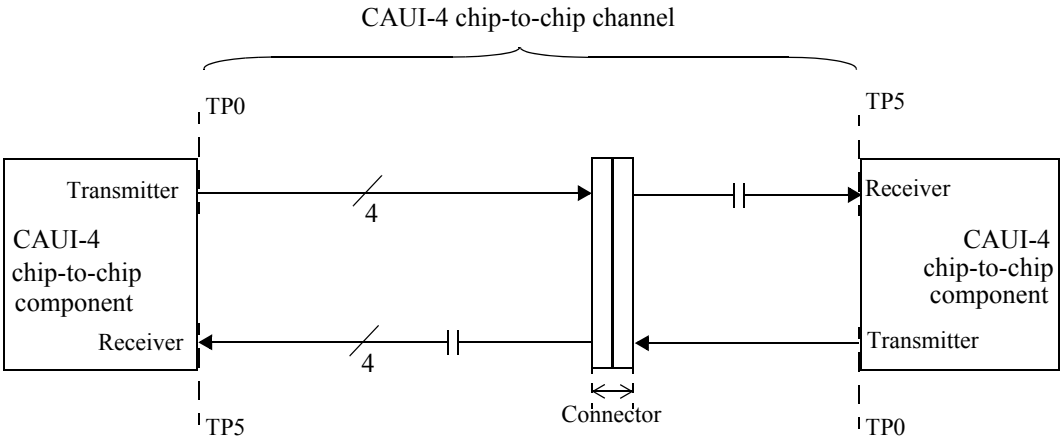

Figure 83D–2—Typical CAUI-4 chip-to-chip application

The normative channel compliance is through CAUI-4 COM as described in 83D.4. Actual channel loss could be higher or lower than that given by Equation (83D–1) due to the channel ILD, return loss, and cross-talk.

$$Insertion\_loss(f) \le \left\{ \begin{array}{cc} 1.083 + 2.543\sqrt{f} + 0.761f & 0.01 \le f < 12.89 \\ -17.851 + 2.936f & 12.89 \le f < 25.78 \end{array} \right\} \text{ (dB)} \qquad (83D–1)$$

where

f	is the frequency in GHz
Insertion_loss(f)	is the informative CAUI-4 chip-to-chip insertion loss

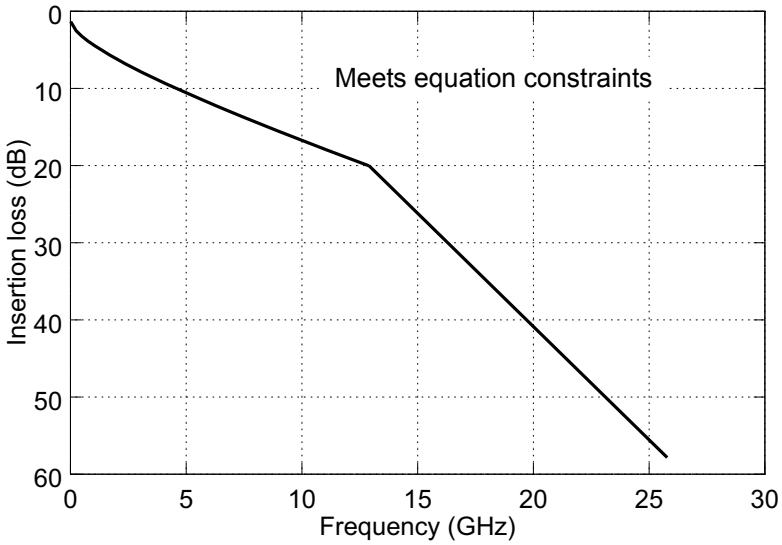

Figure 83D–3—CAUI-4 chip-to-chip channel insertion loss

83D.2 CAUI-4 chip-to-chip compliance point definition

The electrical characteristics for the CAUI-4 chip-to-chip interface are defined at compliance points for the transmitter (TP0a) and receiver (TP5a), respectively. The location of TP0a and electrical characteristics of the test fixture used to measure transmitter characteristics are defined in Figure 93–5 and 93.8.1.1, respectively. The location of TP5a and electrical characteristics of the test fixture used to measure the receiver are defined in Figure 93–10 and 93.8.2.1, respectively.

83D.3 CAUI-4 chip-to-chip electrical characteristics

83D.3.1 CAUI-4 transmitter characteristics

A CAUI-4 chip-to-chip transmitter shall meet the specifications defined in Table 83D–1 if measured at TP0a. While the CAUI-4 chip-to-chip transmitter requirements are similar to those in Clause 93, they differ in that they do not assume transmitter training or a back-channel communications path. Also, the transmit output waveform is not manipulated via the PMD control function described in 93.7.12, but may optionally be manipulated via the feedback mechanism described in 83D.3.3.2.

A test system with a fourth-order Bessel-Thomson low-pass response with 33 GHz 3 dB bandwidth is to be used for all transmitter signal measurements, unless otherwise specified.

Table 83D–1—CAUI-4 transmitter characteristics at TP0a

Parameter	Reference	Value	Units
Signaling rate per lane (range)	93.8.1.2	25.78125 ± 100 ppm	GBd
Differential peak-to-peak output voltage (max) Transmitter disabled Transmitter enabled	93.8.1.3	 30 1200	 mV mV
Common-mode voltage (max)	93.8.1.3	1.9	V
Common-mode voltage (min)	93.8.1.3	0	V
AC common-mode output voltage (max, RMS)	93.8.1.3	12	mV
Differential output return loss (min)	93.8.1.4	Equation (93–3)	dB
Common-mode output return loss (min)	93.8.1.4	Equation (93–4)	dB
Output waveform[a] Steady state voltage v_f(max) Steady state voltage v_f(min) Linear fit pulse peak (min) Pre-cursor equalization Post-cursor equalization	 93.8.1.5.2[b] 93.8.1.5.2[b] 93.8.1.5.2[b] 83D.3.1.1 83D.3.1.1	 0.6 0.4 0.71 x v_f Table 83D–2 Table 83D–3	 V V V — —
Signal-to-noise-and-distortion ratio (min)	93.8.1.6[b]	27	dB
Output Jitter (max) Even-odd jitter Effective bounded uncorrelated jitter, peak-to-peak[c] Effective total uncorrelated jitter, peak-to-peak[cd]	92.8.3.8	 0.035 0.1 0.26	 UI UI UI

[a]The state of the transmit equalizer is controlled by management interface.
[b]The values of the parameters are measured as defined in the referenced subclause except that the values of N_p and N_w are 5.
[c]Effective bounded uncorrelated jitter and effective total uncorrelated jitter are measured as defined in 92.8.3.8.2 except that the range for fitting $CDFL_i$ and $CDFR_i$, as defined in 92.8.3.8.2 c), shall be from 10^{-6} to 10^{-4}.
[d]Effective total uncorrelated jitter, peak-to-peak is specified to a 10^{-15} probability.

83D.3.1.1 Transmitter equalization settings

The CAUI-4 chip-to-chip transmitter includes programmable equalization to compensate for the frequency-dependent loss of the channel and to facilitate data recovery at the receiver. The functional model for the transmit equalizer is the three tap transversal filter shown in Figure 83D–4. The transmitter output equalization is characterized using the linear fit method described in 93.8.1.5.1 where the state of the CAUI-4 transmit output is manipulated via management.

The variable *Local_eq_cm1* controls the weight of the pre-cursor tap c(–1), by changing the ratio c(–1)/(|c(–1)|+|c(0)|+|c(1)|). The valid values of *Local_eq_cm1* and the corresponding ratios are specified in Table 83D–2. The variable *Local_eq_c1* controls the weight of the post-cursor tap c(1), by changing the ratio c(1)/(|c(–1)|+|c(0)|+|c(1)|). The valid values of *Local_eq_c1* and the corresponding ratios are specified in Table 83D–3. *Local_eq_cm1* and *Local_eq_c1* are independent of each other and independent on each lane. Each successive step in *Local_eq_cm1* and *Local_eq_c1* value shall result in a monotonic change in transmitter equalization.

If a Clause 45 MDIO is implemented, *Local_eq_cm1* and *Local_eq_c1* for each lane (0 through 3) and direction (transmit and receive) are accessible through registers 1.180 through 1.187 (see 45.2.1.97 through 45.2.1.100).

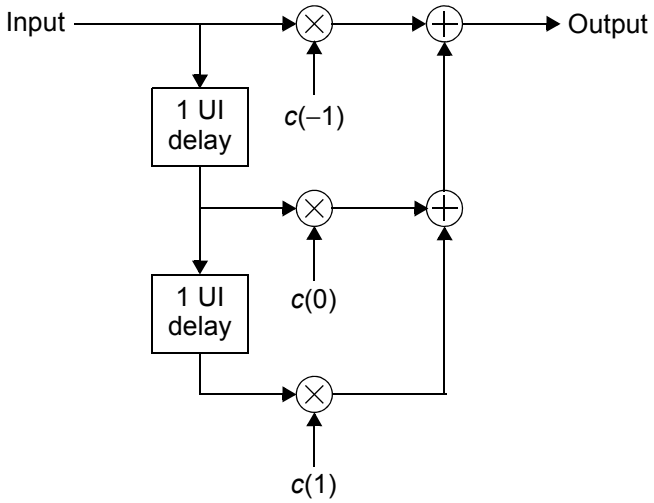

Figure 83D–4—Transmit equalizer functional model

.

Table 83D–2—Pre-cursor equalization

***Local_eq_cm1* value**	***c*(–1) ratio** $\left(\frac{c(-1)}{\lvert c(-1)\rvert + \lvert c(0)\rvert + \lvert c(1)\rvert}\right)$
0	0 ±0.04
1	–0.05 ±0.04
2	–0.1 ±0.04
3	–0.15 ±0.04

83D.3.2 Optional EEE operation

If the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option is supported (see Clause 78 and 78.3), then the inter-sublayer service interface includes four additional primitives as described in 83.3 and may also support CAUI-4 shutdown.

If the EEE capability includes CAUI-4 shutdown (see 78.5.2), then when aui_tx_mode (see 83.5.11.3) is set to ALERT, the transmit direction sublayer sends a repeating 16-bit pattern, hexadecimal 0xFF00, which is transmitted across the CAUI-4. This sequence is transmitted regardless of the value of tx_bit presented by the PMA:IS_UNITDATA_i.request primitive or the rx_bit presented by the PMA:IS_UNITDATA_i.indication primitive. When aui_tx_mode is QUIET, the transmit direction CAUI-4 transmitter is disabled as specified below. Similarly when the received aui_tx_mode is set to ALERT, the receive direction sublayer sends a repeating 16-bit pattern, hexadecimal 0xFF00, which is transmitted across the CAUI-4. This sequence is

Table 83D–3—Post-cursor equalization

Local_eq_c1 value	*c*(1) ratio $\left(\frac{c(1)}{\lvert c(-1)\rvert + \lvert c(0)\rvert + \lvert c(1)\rvert}\right)$
0	0 ±0.04
1	–0.05 ±0.04
2	–0.1 ±0.04
3	–0.15 ±0.04
4	–0.2 ±0.04
5	–0.25 ±0.04

transmitted regardless of the value of tx_bit presented by the PMA:IS_UNITDATA_i.request primitive or the rx_bit presented by the PMA:IS_UNITDATA_i.indication primitive. When the received aui_tx_mode is QUIET, the receive direction CAUI-4 transmitter is disabled as specified below.

For EEE capability with CAUI-4 shutdown, the CAUI-4 transmitter lane's differential peak-to-peak output voltage shall be less than 30 mV within 500 ns of aui_tx_mode changing to QUIET in the relevant direction. Furthermore, the CAUI-4 transmitter lane's differential peak-to-peak output voltage shall be greater than 720 mV within 500 ns of aui_tx_mode ceasing to be QUIET in the relevant direction.

Global transmit disable is optional for EEE capability. The transmit disable function shall turn off all transmitter lanes for a physically instantiated AUI in either the ingress or the egress direction. In the egress direction, the PMA may turn off all the transmitter lanes for the egress direction CAUI-4 if PEASE is asserted and aui_tx_mode is QUIET. In the ingress direction, the PMA may turn off all the transmitter lanes for the ingress direction CAUI-4 if PIASE is asserted and the received aui_tx_mode is QUIET. In both directions, the transmit disable function shall turn on all transmitter lanes after the appropriate direction aui_tx_mode changes to any state other than QUIET within a time and voltage level specified in this subclause.

83D.3.3 CAUI-4 receiver characteristics

A CAUI-4 chip-to-chip receiver shall meet the specifications defined in Table 83D–4 if measured at TP5a.

Table 83D–4—CAUI-4 receiver characteristics at TP5a

Parameter	Subclause reference	Value	Units
Differential input return loss (min)	93.8.2.2	Equation (93–3)	dB
Differential to common-mode input return loss	93.8.2.2	Equation (93–5)	dB
Interference tolerance	83D.3.3.1	Table 83D–5	—

83D.3.3.1 Receiver interference tolerance

The receiver shall satisfy the requirements for interference tolerance defined in Table 83D–5. The interference tolerance test uses the method described in Annex 93C as specified by 93.8.2.3, with the following exceptions:

a) The parameters in Table 83D–5 replace the parameters in Table 93–6.

b) The transmitter taps are set via management (see 83D.3.1.1) to the settings that provide the lowest BER.

c) Sinusoidal jitter is added to the test transmitter by modulating the clock source.

Table 83D–5—Receiver interference tolerance parameters

Parameter	Test 1 values			Test 2 values			Units
	Min	Max	Target	Min	Max	Target	
Bit error ratio[ab]	—	10^{-15}		—	10^{-15}		—
Applied pk-pk sinusoidal jitter			Table 88–13			Table 88–13	
Insertion loss at 12.89 GHz[c]	19.5	20.5		9.5	10.5		dB
Coefficients of fitted insertion loss[d] a_0 a_1 a_2 a_4	 –1 0 0 0	 2 2.937 1.599 0.03		 –1 0 0 0	 1 0.817 0.801 0.01		 dB dB/GHz$^{1/2}$ dB/GHz dB/GHz^2
RSS_DFE2[e]	0.05	—		0.025	—		—
COM including effects of broad-band noise	—	—	2	—	—	2	dB

[a]Bit error ratio replaces the RS symbol error ratio measurement in 93.8.2.3.
[b]Maximum BER assumes errors are not correlated to ensure sufficiently high mean time to false packet acceptance (MTTFPA) assuming 64B/66B coding. Actual implementation of the receiver is beyond the scope of this standard.
[c]Measured between TPt and TP5 (see Figure 93C–4).
[d]Coefficients are calculated from the insertion loss measured between TPt and TP5 (see Figure 93C–4) using the method in 93A.3 with f_{min} = 0.05 GHz, and f_{max} = 25.78125 GHz, and maximum Δf = 0.01 GHz.
[e]RSS_DFE2 is equivalent to RSS_DFE4 described in 93A.2 except that n_1=2 and n_2=5.

83D.3.3.2 Transmitter equalization feedback (optional)

Transmitter equalization feedback is an optional capability for a CAUI-4 chip-to-chip receiver. If implemented, it shall operate as described in this subclause.

Transmitter equalization feedback is generated for each lane (0 through 3) and direction (transmit and receive) independently. The variables that control transmitter equalization feedback are specific for each lane and direction.

A CAUI-4 chip-to-chip receiver may generate a request to change the transmit equalization coefficients of the remote transmitter to new values by setting the *Request_flag* variable to 1. The variables *Request_eq_cm1* and *Request_eq_c1* indicate the request values of *Local_eq_cm1* and *Local_eq_c1*, respectively, in the remote transmitter (see Table 83D–2 and Table 83D–3). The requested setting may be generated from the remote CAUI-4 chip-to-chip transmitter's equalization setting, which is stored in variables *Remote_eq_cm1* and *Remote_eq_c1*, and from information internal to the receiver, in an implementation specific manner.

When a CAUI-4 chip-to-chip receiver does not request a change of the remote transmitter's transmit equalization setting, it sets the *Request_flag* variable to 0. A CAUI-4 chip-to-chip receiver that does not implement transmitter equalization feedback always sets *Requests_flag* to 0.

If a Clause 45 MDIO is implemented, the variables *Request_flag, Requested_eq_cm1, Requested_eq_c1, Remote_eq_cm1,* and *Remote_eq_c1* for each lane and direction are accessible through registers 1.180 through 1.187 (see 45.2.1.97 through 45.2.1.100).

83D.3.4 Global energy detect function for optional EEE operation

The global energy detect function is mandatory for EEE capability with the deep sleep mode option and CAUI-4 shutdown. The global energy detect function indicates whether or not signaling energy is being received on the physical instantiation of the inter sublayer interface (in each direction as appropriate). The energy detection function may be considered a subset of the signal indication logic. If no energy is being received on the CAUI-4 for the ingress direction, SIGNAL_DETECT is set to FAIL following a transition from aui_rx_mode = DATA to aui_rx_mode = QUIET. When aui_rx_mode = QUIET, SIGNAL_DETECT shall be set to OK within 500 ns following the application of a signal at the receiver input that corresponds to an ALERT signal driven from the CAUI-4 link partner. While aui_rx_mode = QUIET, SIGNAL_DETECT changes from FAIL to OK only after the valid ALERT signal is received.

83D.4 CAUI-4 chip-to-chip channel characteristics

The Channel Operating Margin (COM), computed using the procedure in Annex 93A and the parameters in Table 83D–6, shall be greater than or equal to 2 dB. This minimum value allocates margin for practical limitations on the receiver implementation as well as the allowed transmitter equalization coefficients.

Table 83D–6—Channel Operating Margin parameters

Parameter	Symbol	Value	Units
Signaling rate	f_b	25.78125	GBd
Maximum start frequency	f_{min}	0.05	GHz
Maximum frequency step	Δf	0.01	GHz
Device package model Single-ended device capacitance Transmission line length, Test 1 Transmission line length, Test 2 Single-ended board capacitance	 C_d z_p z_p C_b	 2.5×10^{-4} 12 30 1.8×10^{-4}	 nF mm mm nF
Single-ended reference resistance	R_o	50	ohms
Single-ended termination resistance	R_d	55	ohms
Receiver 3 dB bandwidth	f_r	$0.75 \times f_b$	GHz
Transmitter equalizer, minimum cursor coefficient	$c(0)$	0.6	—
Transmitter equalizer, pre-cursor coefficient Minimum value Maximum value Step size	$c(-1)$	 –0.15 0 0.05	—

Table 83D–6—Channel Operating Margin parameters *(continued)*

Parameter	Symbol	Value	Units
Transmitter equalizer, post-cursor coefficient Minimum value Maximum value Step size	$c(1)$	 –0.25 0 0.05	—
Continuous time filter, DC gain Minimum value Maximum value Step size	g_{DC}	 –12 0 1	dB
Continuous time filter, zero frequency	f_z	$f_b/4$	GHz
Continuous time filter, pole frequencies	f_{p1} f_{p2}	$f_b/4$ f_b	GHz
Transmitter differential peak output voltage Victim Far-end aggressor Near-end aggressor	 Av A_{fe} A_{ne}	 0.4 0.4 0.6	 V V V
Number of signal levels	L	2	—
Level separation mismatch ratio	R_{LM}	1	
Transmitter signal-to-noise ratio	SNR_{TX}	27	dB
Number of samples per unit interval	M	32	—
Decision feedback equalizer (DFE) length	N_b	5	UI
Normalized DFE coefficient magnitude limit, for $n = 1$ to N_b	$b_{max}(n)$	0.3	—
Random jitter, RMS	σ_{RJ}	0.01	UI
Dual-Dirac jitter, peak	A_{DD}	0.05	UI
One-sided noise spectral density	η_o	5.2×10^{-8}	V^2/GHz
Target detector error ratio	DER_0	10^{-15}	—

83D.5 Example usage of the optional transmitter equalization feedback

83D.5.1 Overview

If implemented, transmitter equalization feedback from a CAUI-4 chip-to-chip receiver may be used to tune the equalization settings of the transmitter at the other end of the CAUI-4 chip-to-chip link to the values requested by the receiver. An example of a possible transmitter equalization tuning process using transmitter equalization feedback is provided in this subclause.

In this example, two components, A and B, are connected by a CAUI-4 chip-to-chip link, such that A is closest to the PCS and B is closest to the PMD. Clause 45 MDIO is implemented by both components, with component A at device address 11 and component B at device address 10. Transmitter equalization feedback is implemented by either component A, component B, or both. One Station Management (STA) controls both components.

Figure 83D–5 depicts the components of the CAUI-4 chip-to-chip link and the registers used during the tuning procedure.

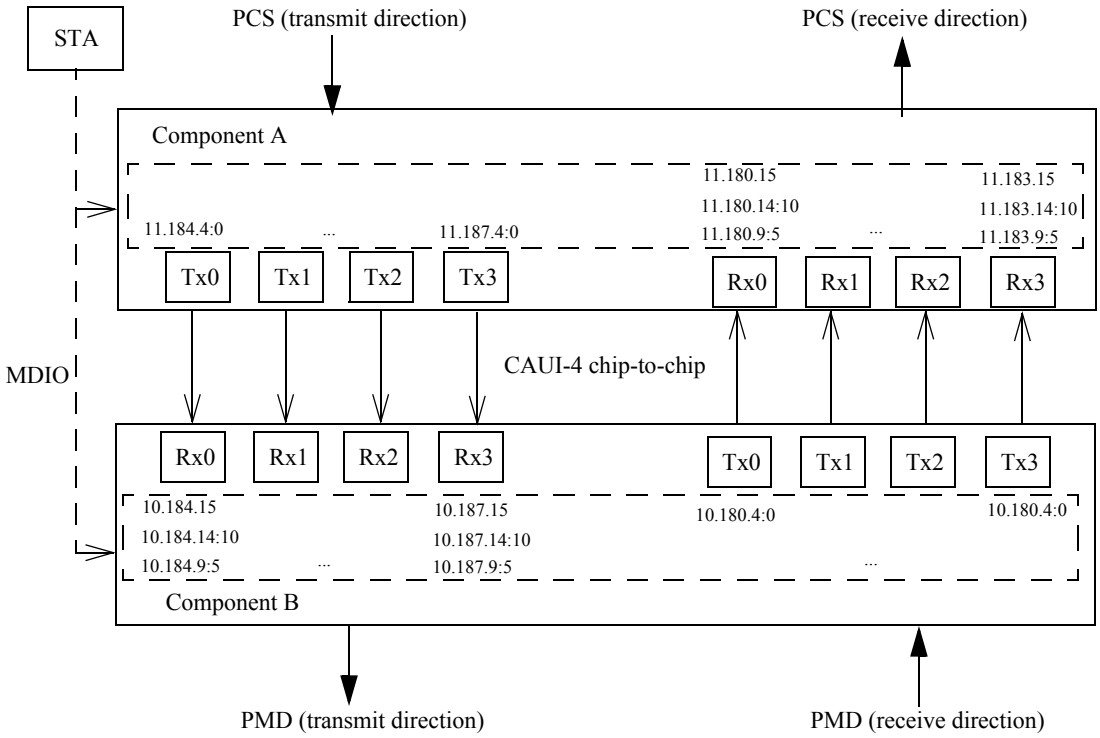

Figure 83D–5—Example transmitter equalization feedback components and registers

The STA performs the procedures described in 83D.5.2 and 83D.5.3 to tune lane 0 equalization settings in both sides of the CAUI-4 chip-to-chip link. When these procedures are completed, the STA uses similar procedures to tune equalization settings in lanes 1 through 3. When all lanes are tuned, the STA may repeat the process with another pair of components connected by CAUI-4 chip-to-chip.

NOTE—Using non-optimal transmitter equalization settings (or changing them) during the tuning procedure may interrupt data communication. The CAUI-4 bit error ratio is assumed to meet the requirements of 83D.3.3.1 upon completion of the tuning process.

83D.5.2 Tuning equalization settings on lane 0 in the transmit direction

1) Read *Local_eq_cm1* (11.184.1:0) and *Local_eq_c1* (11.184.4:2) from component A.
2) Write *Local_eq_cm1* and *Local_eq_c1* read from component A to *Remote_eq_cm1* (10.184.6:5) and *Remote_eq_c1* (10.184.9:7), respectively, in component B.
3) Read *Request_flag* (10.184.15), *Requested_eq_cm1* (10.184.11:10), and *Requested_eq_c1* (10.184.14:12) from component B.
4) If *Request_flag* is 0, go to tuning equalization settings on lane 0 in the Receive direction (83D.5.3).
5) If *Request_flag* is 1, write *Requested_eq_cm1* and R*equested_eq_c1* read from component B to *Local_eq_cm1* (11.184.1:0) and *Local_eq_c1* (11.184.4:2), respectively, in component A.
6) Go to step 1).

83D.5.3 Tuning equalization settings on lane 0 in the receive direction

1) Read *Local_eq_cm1* (10.180.1:0) and *Local_eq_c1* (10.180.4:2) from component B.
2) Write *Local_eq_cm1* and *Local_eq_c1* read from component B to *Remote_eq_cm1* (11.180.6:5) and *Remote_eq_c1* (11.180.9:7), respectively, in component A.
3) Read *Request_flag* (11.180.15), *Requested_eq_cm1* (11.180.11:10), and *Requested_eq_c1* (11.180.14:12) from component A.
4) If *Request_flag* is 0, proceed to tuning lane 1.
5) If *Request_flag* is 1, write *Requested_eq_cm1* and *Requested_eq_c1* read from component A to *Local_eq_cm1* (10.180.1:0) and *Local_eq_c1* (10.180.4:2), respectively, in component B.
6) Go to step 1).

83D.6 Protocol implementation conformance statement (PICS) proforma for Annex 83D, Chip-to-chip 100 Gb/s four-lane Attachment Unit Interface (CAUI-4)[28]

83D.6.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Annex 83D, Chip-to-chip 100 Gb/s four-lane Attachment Unit Interface (CAUI-4), 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.

83D.6.2 Identification

83D.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).	

83D.6.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Annex 83D, Chip-to-chip 100 Gb/s four-lane Attachment Unit Interface (CAUI-4)
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	

[28] *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.

83D.6.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
NOL	Number of differential AC-coupled lanes	83D.1	Four independent data paths in each direction	M	Yes [ ]
*CHAN	Channel	83D.4	Items marked with CHAN include channel specifications not applicable to a PHY manufacturer	O	Yes [ ] No [ ]
*LPI	Support for CAUI-4 shutdown	83D.3.2		O	Yes [ ] No [ ]

83D.6.4 PICS proforma tables for chip-to-chip 100 Gb/s four-lane Attachment Unit Interface (CAUI-4)

83D.6.4.1 Transmitter

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	Signaling rate	83D.3.1	25.78125 GBd ± 100 ppm per lane	M	Yes [ ]
TC2	Peak-to-peak differential output voltage	83D.3.1	1200 mV (max)	M	Yes [ ]
TC3	Peak-to-peak differential output voltage, transmitter disabled	83D.3.1	Less than or equal to 30 mV	M	Yes [ ]
TC4	DC common-mode voltage	83D.3.1	Between 0 V and 1.9 V with respect to signal ground	M	Yes [ ]
TC5	AC common-mode output voltage	83D.3.1	12 mV RMS with respect to signal ground	M	Yes [ ]
TC6	Differential output return loss	83D.3.1	Meets Equation (93–3) constraints	M	Yes [ ]
TC7	Common-mode output return loss	83D.3.1	Meets Equation (93–4) constraints	M	Yes [ ]
TC8	Output waveform	83D.3.1	Meets Table 83D–1 constraints	M	Yes [ ]
TC9	Output jitter	83D.3.1	Meets Table 83D–1 constraints	M	Yes [ ]
TC10	Amplitude and swing for CAUI-4 shutdown	83D.3.2		LPI:M	Yes [ ]
TC11	Transmit disable for CAUI-4 shutdown	83D.3.2		LPI:M	Yes [ ]
TC12	Transmit equalization	83D.3.1.1	Each successive step results in a monotonic change	M	Yes [ ]

83D.6.4.2 Receiver

Item	Feature	Subclause	Value/Comment	Status	Support
RC1	Differential input return loss	93.8.2.2	Meets Equation (93–3) constraints	M	Yes [ ]
RC2	Differential to common-mode input return loss	93.8.2.2	Meets Equation (93–5) constraints	M	Yes [ ]
RC3	Interference tolerance	83D.3.3.1	Satisfy requirements in Table 83D–5	M	Yes [ ]
RC4	Jitter tolerance	93.8.2.4	Satisfy requirements in Table 93–7	M	Yes [ ]
RC5	Signal detect for CAUI-4 shutdown	83D.3.4		LPI:M	Yes [ ]

83D.6.4.3 Channel

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	Channel Operating Margin (COM)	83D.4	Greater than or equal to 2 dB	CHAN:M	Yes [ ]

Annex 83E

(normative)

Chip-to-module 100 Gb/s four-lane Attachment Unit Interface (CAUI-4)

83E.1 Overview

This annex defines the functional and electrical characteristics for the optional chip-to-module 100 Gb/s four-lane Attachment Unit Interface (CAUI-4). Figure 83E–1 shows the relationship of the CAUI-4 chip-to-module interface to the ISO/IEC Open System Interconnection (OSI) reference model. The chip-to-module interface provides electrical characteristics and associated compliance points, which can optionally be used when designing systems with pluggable module interfaces.

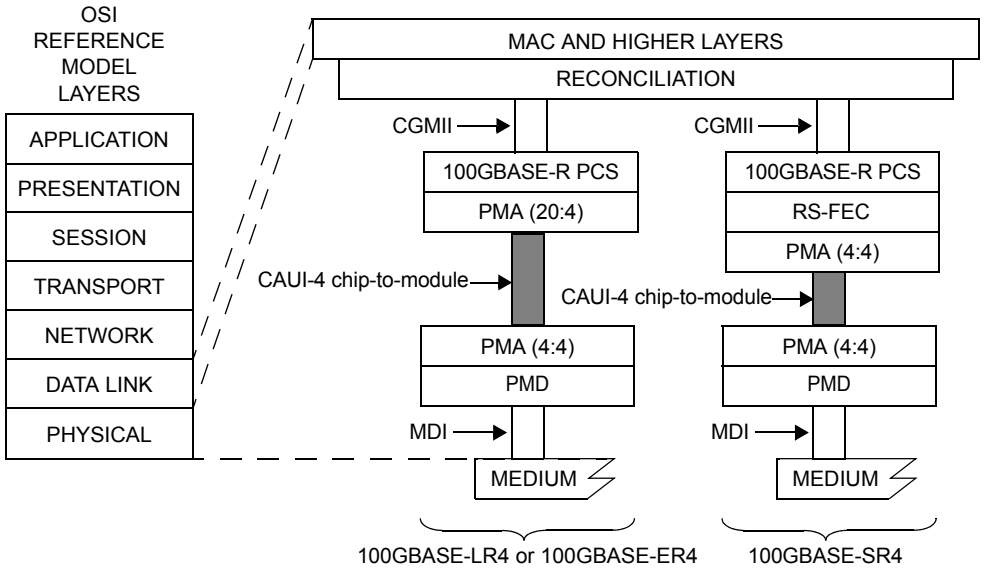

CAUI-4 = 100 Gb/s FOUR-LANE ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION

Figure 83E–1—Example CAUI-4 chip-to-module relationship to the ISO/IEC Open System Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

The CAUI-4 link is described in terms of a host CAUI-4 component, a CAUI-4 channel with associated insertion loss, and a module CAUI-4 component. Figure 83E–2 and Equation (83E–1) depict a typical CAUI-4 application and summarize the differential insertion loss budget associated with the chip-to-module application, which is shown in Figure 83E–3. The CAUI-4 chip-to-module interface comprises independent data paths in each direction. Each data path contains four differential lanes, which are AC-coupled within the module. The nominal signaling rate for each lane is 25.78125 GBd. The chip-to-module interface is

defined using a specification and test methodology that is similar to that used for CEI-28G-VSR defined in OIF-CEI-03.1 [B55].

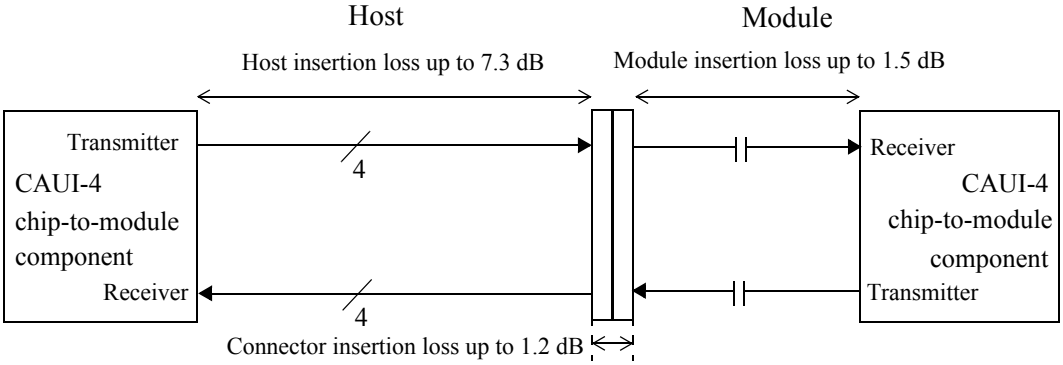

Figure 83E–2—Chip-to-module insertion loss budget at 12.89 GHz

$$Insertion\_loss(f) \le \left\{ \begin{array}{cc} 1.076(0.075 + 0.537\sqrt{f} + 0.566f) & 0.01 \le f < 14 \\ 1.076(-18 + 2f) & 14 \le f < 18.75 \end{array} \right\} \text{ (dB)} \qquad (83E\text{–}1)$$

where

f is the frequency in GHz
Insertion_loss(f) is the CAUI-4 chip-to-module insertion loss

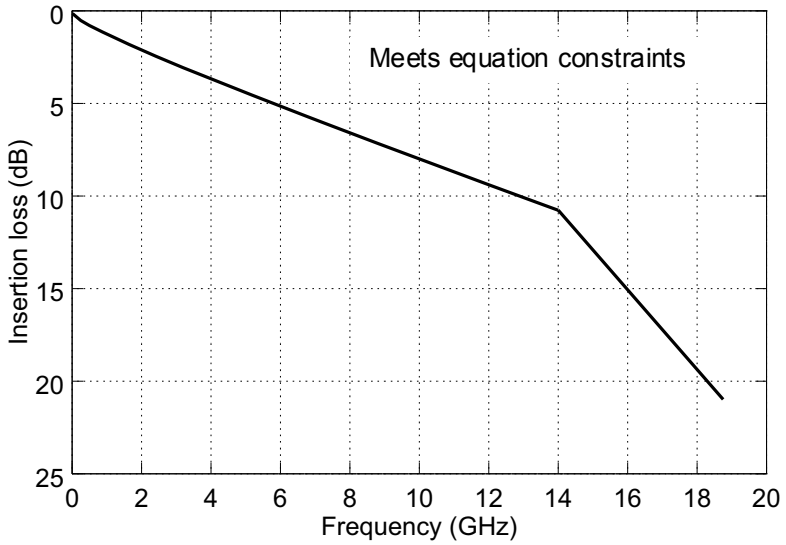

Figure 83E–3—CAUI-4 chip-to-module channel insertion loss

83E.1.1 Bit error ratio

The bit error ratio (BER) shall be less than 10^{-15} with any errors sufficiently uncorrelated to ensure an acceptably high mean time to false packet acceptance (MTTFPA) assuming 64B/66B coding.

83E.2 CAUI-4 chip-to-module compliance point definitions

The electrical characteristics for the CAUI-4 chip-to-module interface are defined at compliance points for the host and module, respectively. Reference test fixtures, called compliance boards, are used to access the electrical specification parameters. Figure 83E–4 depicts the location of compliance points when measuring host CAUI-4 compliance. The output of the Host Compliance Board (HCB) is used to verify the host electrical output signal at TP1a. Similarly, the input of the HCB at TP4a is used to verify the host input compliance.

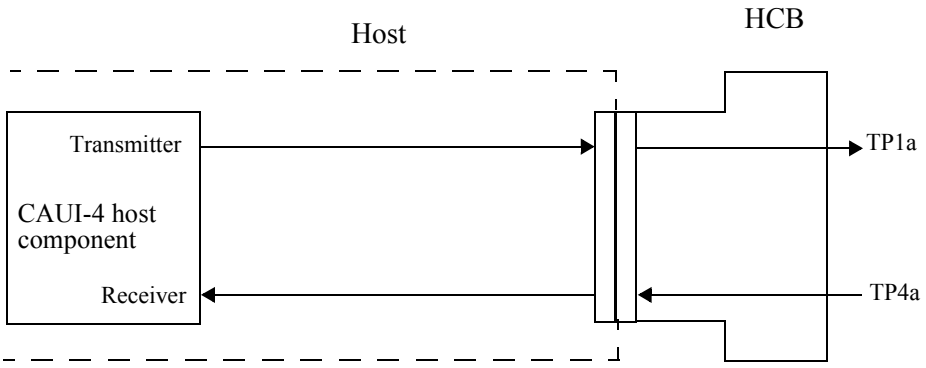

Figure 83E–4—Host CAUI-4 compliance points

Figure 83E–5 depicts the location of compliance points when measuring module CAUI-4 compliance. The output of the Module Compliance Board (MCB) is used to verify the module electrical output signal at TP4. Similarly, the input of the MCB at TP1 is used to verify the module input compliance. Additional details on the requirements for the MCB and HCB are given in 83E.4.1.

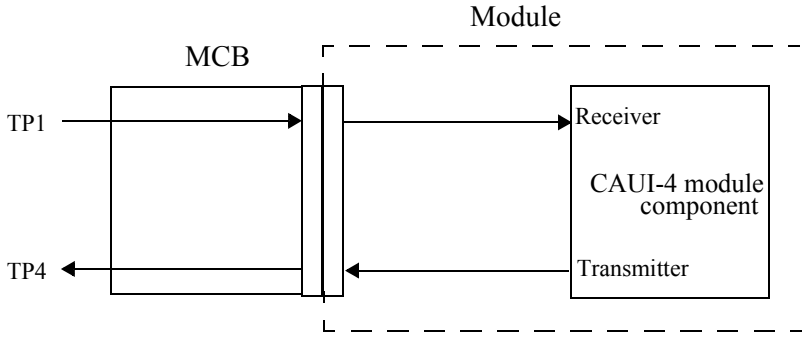

Figure 83E–5—Module CAUI-4 compliance points

83E.3 CAUI-4 chip-to-module electrical characteristics

83E.3.1 CAUI-4 host output characteristics

A CAUI-4 host output shall meet the specifications defined in Table 83E–1 if measured at TP1a.

Table 83E–1—CAUI-4 host output characteristics (at TP1a)

Parameter	Reference	Value	Units
Signaling rate per lane (range)	83E.3.1.1	25.78125 ± 100 ppm	GBd
DC common-mode output voltage (max)	83E.3.1.2	2.8	V
DC common-mode output voltage (min)	83E.3.1.2	–0.3	V
Single-ended output voltage (max)	83E.3.1.2	3.3	V
Single-ended output voltage (min)	83E.3.1.2	–0.4	V
AC common-mode output voltage (max, RMS)	83E.3.1.2	17.5	mV
Differential peak-to-peak output voltage (max) Transmitter disabled Transmitter enabled	83E.3.1.2	 35 900	mV
Eye width (min)	83E.3.1.6	0.46	UI
Eye height A, differential (min)	83E.3.1.6	95	mV
Eye height B, differential (min)	83E.3.1.6	80	mV
Differential output return loss (min)	83E.3.1.3	Equation (83E–2)	dB
Common to differential mode conversion return loss (min)	83E.3.1.3	Equation (83E–3)	dB
Differential termination mismatch (max)	83E.3.1.4	10	%
Transition time (min, 20% to 80%)	83E.3.1.5	10	ps

A test system with a fourth-order Bessel-Thomson low-pass response with 33 GHz 3 dB bandwidth is to be used for all output signal measurements, unless otherwise specified.

83E.3.1.1 Signaling rate and range

The CAUI-4 signaling rate is 25.78125 GBd ± 100 ppm per lane. This translates to a nominal unit interval of 38.787879 ps.

83E.3.1.2 Signal levels

The differential output voltage v_{di} is defined to be the difference between the single-ended output voltages, SL*i*<p> minus SL*i*<n>. The common-mode voltage v_{cmi} is defined to be one half of the sum of SL*i*<p> and SL*i*<n>. These definitions are illustrated by Figure 83E–6.

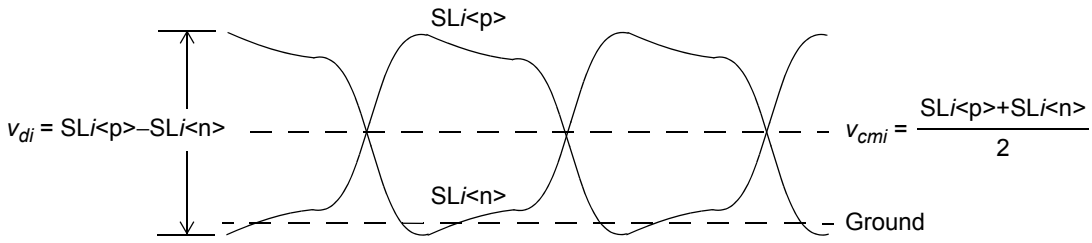

Figure 83E–6—Voltage definitions

The peak-to-peak differential output voltage is less than or equal to 900 mV. The peak-to-peak differential output voltage is less than or equal to 35 mV when the transmitter is disabled.

The DC common-mode output voltage and AC common-mode output voltage are defined with respect to signal ground.

83E.3.1.3 Output return loss

The differential output return loss, in dB, of the output is shown in Equation (83E–2) and illustrated in Figure 83E–7. This output requirement applies to all valid output levels. The reference impedance for differential return loss measurements is 100 Ω.

$$RLd(f) \geq \begin{cases} 9.5 - 0.37f & 0.01 \leq f < 8 \\ 4.75 - 7.4\log_{10}\left(\frac{f}{14}\right) & 8 \leq f < 19 \end{cases} \text{ (dB)} \qquad \text{(83E–2)}$$

where

f	is the frequency in GHz
RLd	is the CAUI-4 chip-to-module host output differential return loss

Common to differential output conversion return loss, in dB, of the output is shown in Equation (83E–3) and illustrated in Figure 83E–8.

$$RLdc(f) \geq \begin{cases} 22 - 20\left(\frac{f}{25.78}\right) & 0.01 \leq f < 12.89 \\ 15 - 6\left(\frac{f}{25.78}\right) & 12.89 \leq f < 19 \end{cases} \text{ (dB)} \qquad \text{(83E–3)}$$

where

f	is the frequency in GHz
$RLdc$	is the CAUI-4 chip-to-module output common to differential mode conversion return loss

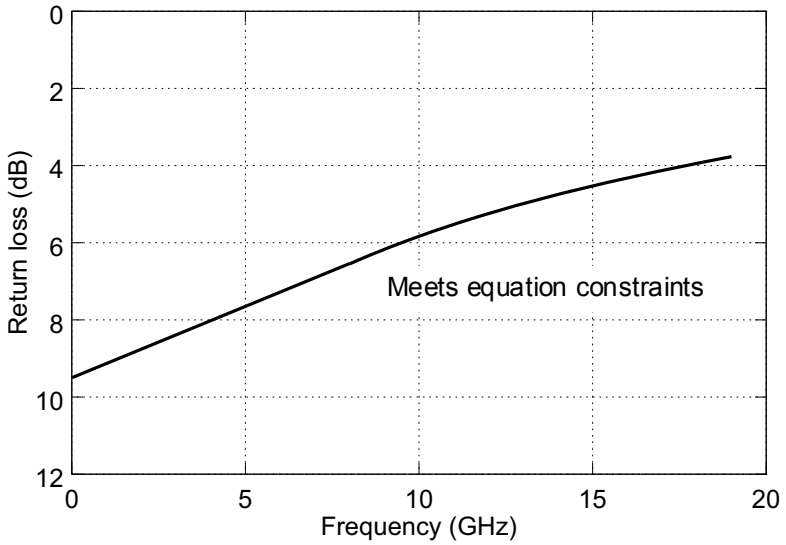

Figure 83E–7—Output differential return loss

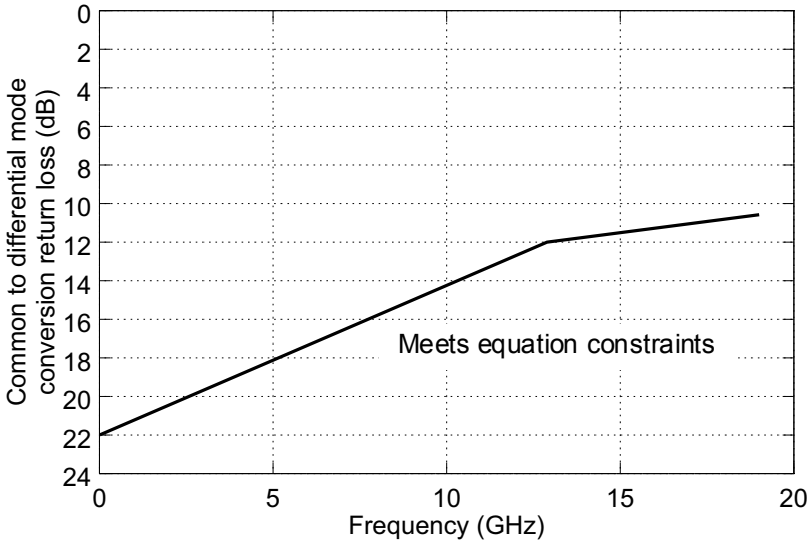

Figure 83E–8—Output common to differential mode conversion return loss

83E.3.1.4 Differential termination mismatch

Differential termination mismatch is defined in 86A.5.3.2.

83E.3.1.5 Transition time

The transition times (rise and fall times) are defined in 86A.5.3.3 with the exception that the observation is through a 33 GHz low-pass filter response.

83E.3.1.6 Host output eye width and eye height

Figure 83E–9 depicts an example host output eye width and eye height test configuration. Host output eye width and eye height are measured at TP1a using compliance boards defined in 83E.2. The host output eye is measured using a reference receiver with a continuous time linear equalizer (CTLE) defined in 83E.3.1.6.1. The recommended CTLE peaking value is used for host output eye measurements. In addition, it is provided to the module via the variable *Recommended_CTLE_value.* If a Clause 45 MDIO is implemented, this variable is accessible in the module through register 1.179 (see 45.2.1.96). Eye width and eye height measurement methodology is described in 83E.4.2. All counter-propagating signals shall be asynchronous to the co-propagating signals using Pattern 5 (with or without FEC encoding), Pattern 3, or a valid 100GBASE-R signal. Patterns 3 and 5 are described in Table 86–11. For the case where Pattern 3 is used with a common clock, there is at least 31 UI delay between the PRBS31 patterns on one lane and any other lane. The crosstalk generator is calibrated at TP4 with target differential peak-to-peak amplitude of 900 mV and target transition time of 12 ps.

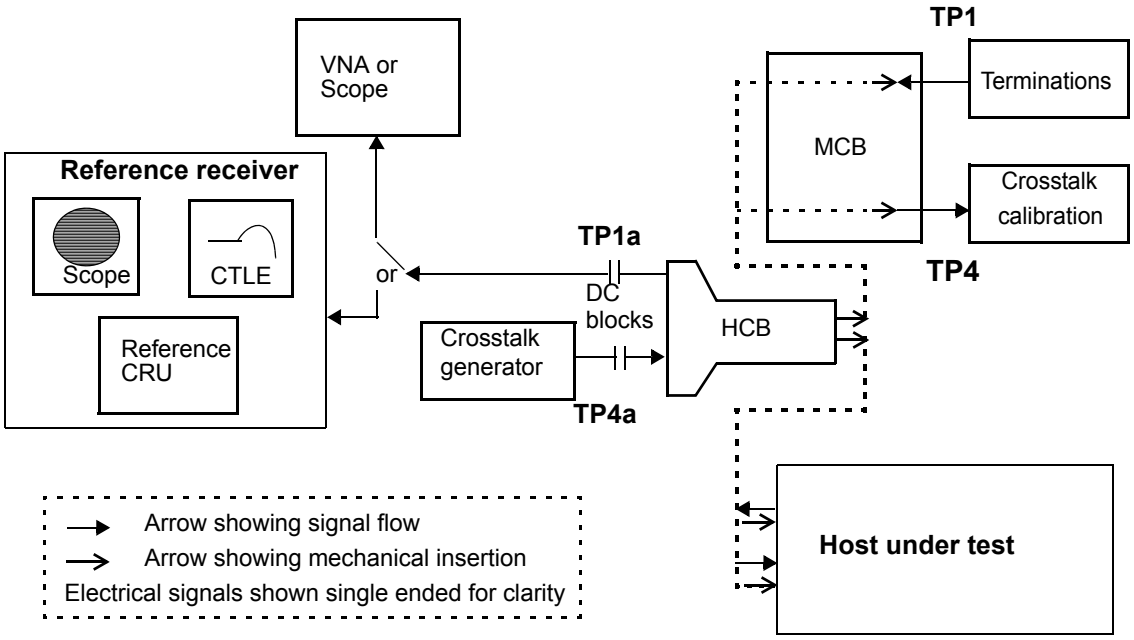

Figure 83E–9—Example host output test configuration

83E.3.1.6.1 Reference receiver for host output eye width and eye height evaluation

The reference receiver is used to measure host eye width and eye height. The reference receiver includes a selectable continuous time linear equalizer (CTLE), which is described by Equation (83E–4) with coefficients given in Table 83E–2 and illustrated in Figure 83E–10. The equalizer may be implemented in software; however, the measured signal is not averaged.

$$H(f) = \frac{GP_1P_2}{Z_1} \times \frac{\mathrm{j}2\pi f + Z_1}{(\mathrm{j}2\pi f + P_1)(\mathrm{j}2\pi f + P_2)} \quad (83E–4)$$

where

$H(f)$ is the CTLE transfer function
G is the CTLE gain

P_1, P_2 are the CTLE poles in Grad/s
Z_1 is the CTLE zero in Grad/s
j is the square root of –1
f is the frequency in GHz

Table 83E–2—Reference CTLE coefficients

Peaking (dB)	G	$\frac{P_1}{2\pi}$	$\frac{P_2}{2\pi}$	$\frac{Z_1}{2\pi}$
1	0.89125	18.6	14.1	8.364
2	0.79433	18.6	14.1	7.099
3	0.70795	15.6	14.1	5.676
4	0.63096	15.6	14.1	4.9601
5	0.56234	15.6	14.1	4.358
6	0.50119	15.6	14.1	3.844
7	0.44668	15.6	14.1	3.399
8	0.39811	15.6	14.1	3.012
9	0.35481	15.6	14.1	2.672

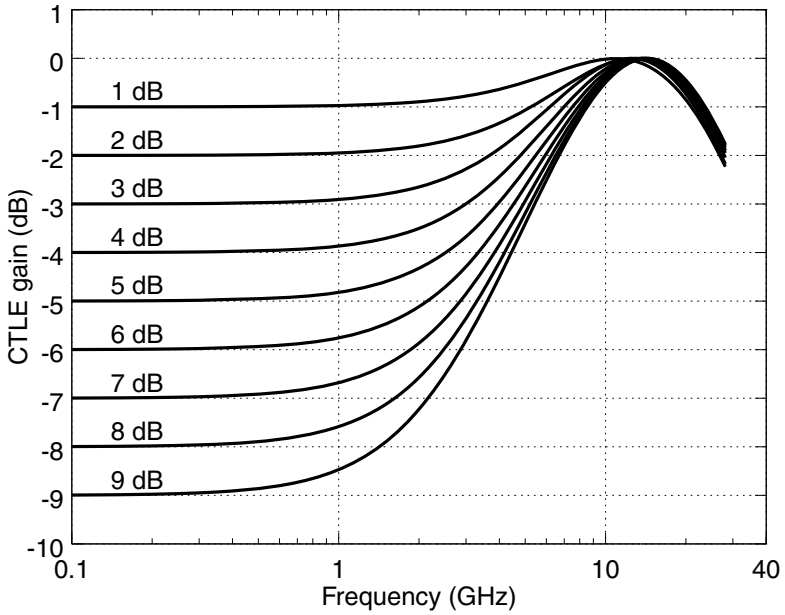

Figure 83E–10—Selectable continuous time linear equalizer (CTLE) characteristic

83E.3.2 CAUI-4 module output characteristics

A CAUI-4 module output shall meet the specifications defined in Table 83E–3 if measured at TP4. A test system with a fourth-order Bessel-Thomson low-pass response with 33 GHz 3 dB bandwidth is to be used for all output signal measurements, unless otherwise specified.

Table 83E–3—CAUI-4 module output characteristics (at TP4)

Parameter	Reference	Value	Units
Signaling rate per lane (range)	83E.3.1.1	25.78125 ± 100 ppm	GBd
AC common-mode output voltage (max, RMS)	83E.3.1.2	17.5	mV
Differential output voltage (max)	83E.3.1.2	900	mV
Eye width (min)	83E.3.2.1	0.57	UI
Eye height, differential (min)	83E.3.2.1	228	mV
Vertical eye closure (max)	83E.4.2.1	5.5	dB
Differential output return loss (min)	83E.3.1.3	Equation (83E–2)	dB
Common to differential mode conversion return loss (min)	83E.3.1.3	Equation (83E–3)	dB
Differential termination mismatch (max)	83E.3.1.4	10	%
Transition time (min, 20% to 80%)	83E.3.1.5	12	ps
DC common mode voltage (min)[a]	83E.3.1.2	–350	mV
DC common mode voltage (max)[a]	83E.3.1.2	2850	mV

[a]DC common mode voltage is generated by the host. Specification includes effects of ground offset voltage.

83E.3.2.1 Module output eye width and eye height

Module output eye width is greater than 0.57 UI. Module output eye height is greater than 228 mV. Figure 83E–11 depicts an example module output eye width and eye height test configuration. Module output eye width and eye height are measured at TP4 using compliance boards defined in 83E.2. The module output eye is measured using a reference receiver with a continuous time linear equalizer (CTLE) defined in 83E.3.2.1.1. Eye width and eye height measurement methodology is described in 83E.4.2. All counter-propagating signals shall be asynchronous to the co-propagating signals using Pattern 5 (with or without FEC encoding), Pattern 3, or a valid 100GBASE-R signal. Patterns 3 and 5 are described in Table 86–11. For the case where Pattern 3 is used with a common clock, there is at least 31 UI delay between the PRBS31 patterns on one lane and any other lane. The crosstalk generator is calibrated at TP1a with target differential peak-to-peak amplitude of 900 mV and target transition time of 19 ps.

83E.3.2.1.1 Reference receiver for module output eye width and eye height evaluation

A reference receiver is used to measure module eye width and eye height. The reference receiver includes a selectable continuous time linear equalizer (CTLE), which is described by Equation (83E–4) with coefficients given in the first two rows of Table 83E–2. The equalizer may be implemented in software; however, the measured signal is not averaged. Either of the two equalizer settings may be used to meet the output eye width and eye height requirement.

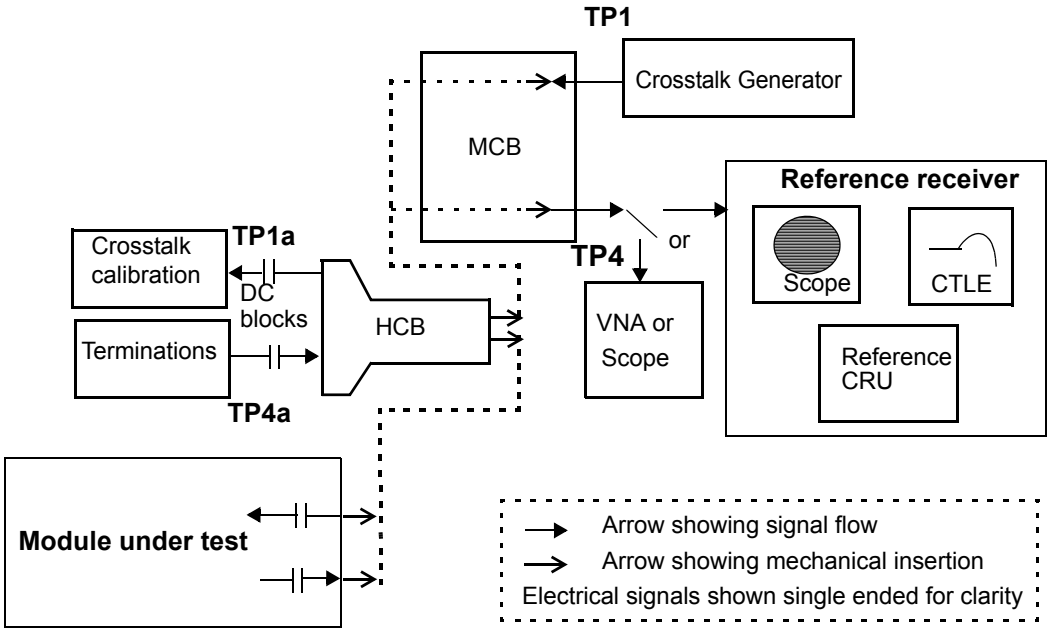

Figure 83E–11—Example module output test configuration

83E.3.3 CAUI-4 host input characteristics

A CAUI-4 host input shall meet the specifications defined in Table 83E–4 if measured at the appropriate test point.

Table 83E–4—CAUI-4 host input characteristics

Parameter	Reference	Test point	Value	Units
Signaling rate, per lane (range)	83E.3.1.1	TP4a	25.78125 ± 100 ppm	GBd
Differential pk-pk input voltage tolerance (min)	83E.3.1.2	TP4	900	mV
Differential input return loss (min)	83E.3.3.1	TP4a	Equation (83E–5)	dB
Differential to common-mode input return loss (min)	83E.3.3.1	TP4a	Equation (83E–6)	dB
Host stressed input test[a]	83E.3.3.2	TP4	See 83E.3.3.2	
Differential termination mismatch (max)	83E.3.1.4	TP4a	10	%
Common-mode voltage[b] Min Max	83E.3.1.2	TP4a	 –0.3 2.8	V

[a]Meets BER specified in 83E.1.1.
[b]Generated by host, referred to host ground.

83E.3.3.1 Input return loss

The differential input return loss, in dB, of the input is shown in Equation (83E–5) and illustrated in Figure 83E–12. The reference impedance for differential return loss measurements is 100 Ω.

$$RLd(f) \geq \left\{ \begin{array}{ll} 9.5 - 0.37f & 0.01 \leq f < 8 \\ 4.75 - 7.4\log_{10}\left(\frac{f}{14}\right) & 8 \leq f < 19 \end{array} \right\} \text{ (dB)} \qquad \text{(83E–5)}$$

where

f is the frequency in GHz
RLd is the CAUI-4 chip-to-module input differential return loss

Differential to common-mode input return loss, in dB, of the input is shown in Equation (83E–6) and illustrated in Figure 83E–13.

$$RLcd(f) \geq \left\{ \begin{array}{ll} 22 - 20\left(\frac{f}{25.78}\right) & 0.01 \leq f < 12.89 \\ 15 - 6\left(\frac{f}{25.78}\right) & 12.89 \leq f < 19 \end{array} \right\} \text{ (dB)} \qquad \text{(83E–6)}$$

where

f is the frequency in GHz
$RLcd$ is the CAUI-4 chip-to-module input differential to common-mode input return loss

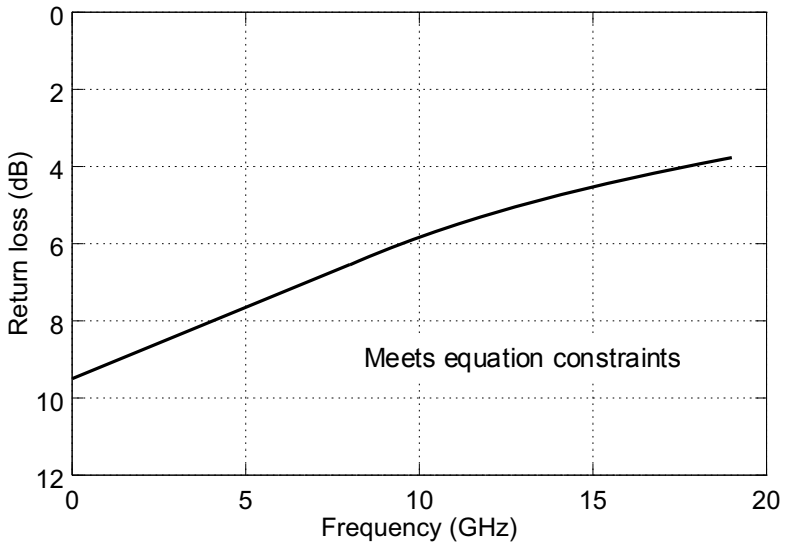

Figure 83E–12—Differential input return loss

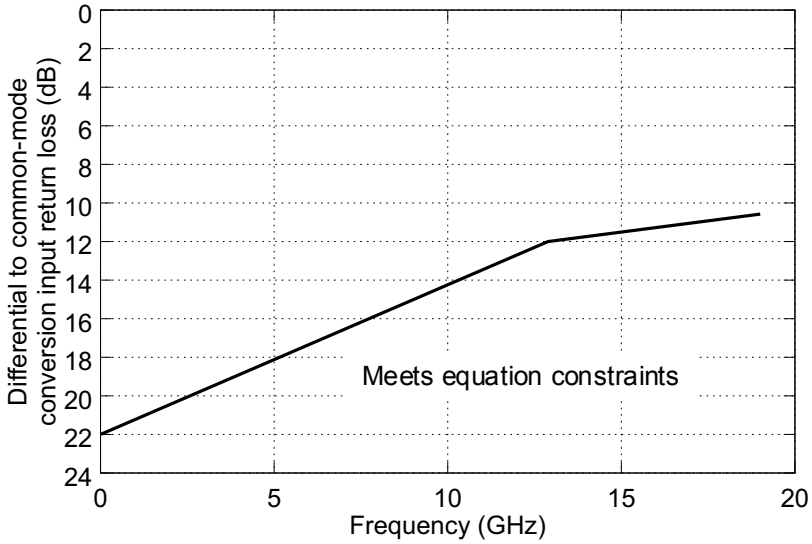

Figure 83E–13—Differential to common-mode conversion input return loss

83E.3.3.2 Host stressed input test

The host stressed input tolerance is measured using the procedure defined in 83E.3.3.2.1. The input shall satisfy the input tolerance defined in Table 83E–5.

Table 83E–5—Host stressed input parameters

Parameter	Value
Eye width	0.57 UI
Applied pk-pk sinusoidal jitter	Table 88–13
Eye height	228 mV

83E.3.3.2.1 Host stressed input test procedure

The host stressed input test is summarized in Figure 83E–14. The stressed signal is applied at TP4a and is calibrated at TP4. A reference CRU with a corner frequency of 10 MHz and slope of 20 dB/decade is used to calibrate the stressed signal using Pattern 4 (PRBS9, see Table 86–11 and Table 68–6). The reference receiver includes a selectable CTLE given by Equation (83E–4) and the first two rows of Table 83E–2. The stressed signal is generated by adding sinusoidal jitter, random jitter, and bounded uncorrelated jitter to a clean signal. The amount of applied peak-to-peak sinusoidal jitter used for the host stressed input test is given in Table 83E–5. Bounded uncorrelated jitter provides a source of bounded high probability jitter uncorrelated with the signal stream. This jitter stress source may not be present in all stressed pattern generators or bit error ratio testers. It can be generated by driving the pattern generator external jitter modulation input with a filtered PRBS pattern. The PRBS pattern length should be between PRBS7 and PRBS9. The PRBS signaling rate should be approximately 1/10 of the stressed signal's signaling rate (i.e., approximately

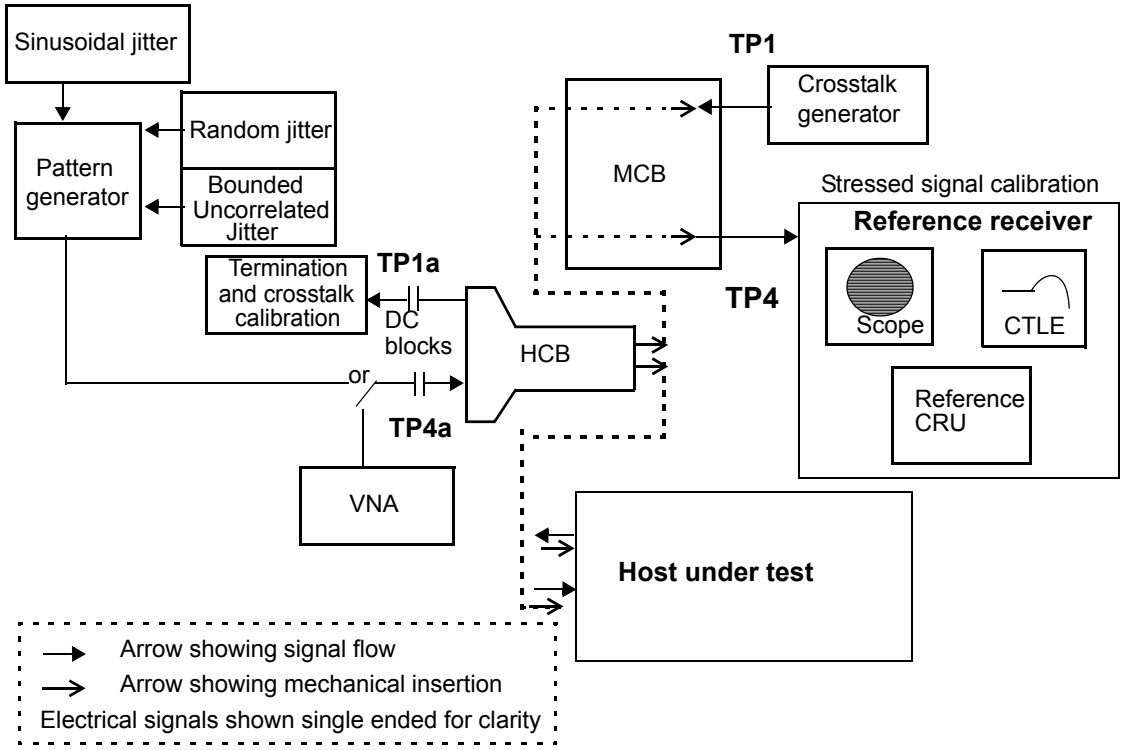

Figure 83E–14—Example host stressed input test

2.578 GBd). The clock source for the PRBS generator is asynchronous to the pattern generator clock source to assure non-correlation of the jitter. The low pass filter that operates on the PRBS pattern to generate the bounded uncorrelated jitter should exhibit 20 dB/decade roll-off with a –3 dB corner frequency between 150 MHz and 300 MHz. This value must also be below the upper frequency limit of the pattern generator external modulator input. Random jitter and bounded uncorrelated jitter are added such that the output of the pattern generator approximates a jitter profile given in Table 83E–6.

Table 83E–6—Pattern generator jitter characteristics

Parameter	Value
Total Jitter (pk-pk)[a]	0.28 UI
Random Jitter (pk-pk)[b]	0.15 UI
Max even-odd jitter (pk-pk)[c]	0.035 UI

[a]Total Jitter at BER of 10^{-15}
[b]Random Jitter at BER of 10^{-15}
[c]As defined in 92.8.3.8.1

The counter propagating crosstalk channels during calibration of the stressed signal are asynchronous with target amplitude of 900 mV peak-to-peak differential and 20% to 80% target transition time of 19 ps as measured at TP1a. The crosstalk signal transition time is calibrated with Pattern 4. The pattern is changed to Pattern 5 (with or without FEC encoding), Pattern 3, or a valid 100GBASE-R signal for amplitude calibration and the stressed input test. Patterns 3, 4, and 5 are described in Table 86–11. For the case where Pattern 3 is used with a common clock, there is at least 31 UI delay between the PRBS31 patterns on one lane and any

other lane. Any one of these patterns is sufficient as a crosstalk aggressor with all lanes active during the stressed input test.

Eye height and eye width, extrapolated to a probability of 10^{-15}, are then measured at TP4 based on the eye measurement methodology given in 83E.4.2. Random jitter and the pattern generator output amplitude are adjusted (without exceeding the differential pk-pk input voltage tolerance specification as shown in Table 83E–4) to result in the eye height and eye width given in Table 83E–5 using the reference receiver with the setting of the CTLE that maximizes the product of eye height and eye width.

A host input stressed signal should have a vertical eye closure in the range of 4.5 dB to 5.5 dB with a target value of 5 dB.

The pattern is then changed to Pattern 5 (with or without FEC encoding), Pattern 3, or a valid 100GBASE-R signal for the input test, which is conducted by inserting the HCB into the host under test.

83E.3.4 CAUI-4 module input characteristics

A CAUI-4 module input shall meet the specifications defined in Table 83E–7 if measured at the appropriate test point.

Table 83E–7—CAUI-4 module input characteristics

Parameter	Reference	Test point	Value	Units
Signaling rate per lane (range)	83E.3.1.1	TP1	25.78125 ± 100 ppm	GBd
Differential pk-pk input voltage tolerance (min)	83E.3.1.2	TP1a	900	mV
Differential input return loss (min)	83E.3.3.1	TP1	Equation (83E–5)	dB
Differential to common-mode input return loss (min)	83E.3.3.1	TP1	Equation (83E–6)	dB
Differential termination mismatch (max)	83E.3.1.4	TP1	10	%
Module stressed input test[a]	83E.3.4.1	TP1a	See 83E.3.4.1	
Single-ended voltage tolerance range (min)	83E.3.1.2	TP1a	–0.4 to 3.3	V
DC common mode voltage (min)[b]	83E.3.1.2	TP1	–350	mV
DC common mode voltage (max)[b]	83E.3.1.2	TP1	2850	mV

[a]Meets BER specified in 83E.1.1.
[b]DC common mode voltage generated by the host. Specification includes effects of ground offset voltage.

83E.3.4.1 Module stressed input test

The module stressed input tolerance is measured using the procedure defined in 83E.3.4.1.1. The input shall satisfy the input tolerance defined in Table 83E–8.

83E.3.4.1.1 Module stressed input test procedure

The module stressed input test is summarized in Figure 83E–15. The stressed signal is applied at TP1 and is calibrated at TP1a. A reference CRU with a corner frequency of 10 MHz and slope of 20 dB/decade is used

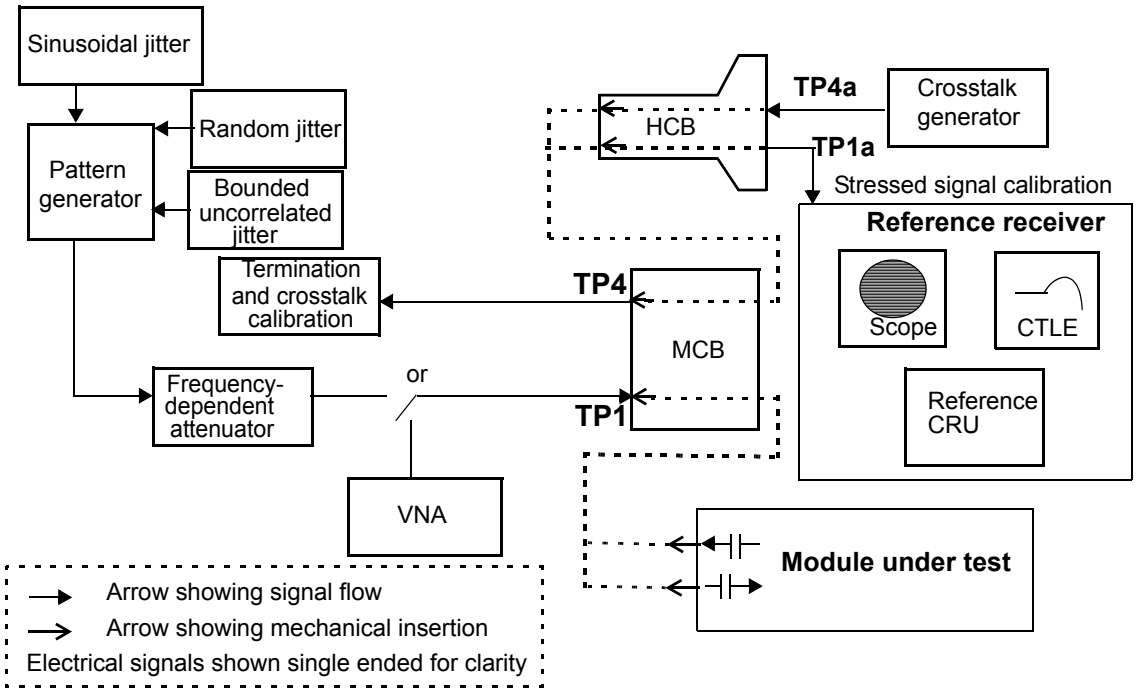

Figure 83E–15—Example module stressed input test

Table 83E–8—Module stressed input parameters

Parameter	Value
Eye width	0.46 UI
Applied pk-pk sinusoidal jitter	Table 88–13
Eye height	95 mV

to calibrate the stressed signal using Pattern 4. The reference receiver includes a selectable CTLE given by Equation (83E–4) and Table 83E–2. The stressed signal is generated by adding sinusoidal jitter, random jitter, and bounded uncorrelated jitter to a clean signal, followed by frequency-dependent attenuation. The frequency-dependent attenuator represents the host channel and may be implemented with PCB traces. The amount of applied peak-to-peak sinusoidal jitter used for the module stressed input test is given in Table 83E–8. Bounded uncorrelated jitter provides a source of bounded high probability jitter uncorrelated with the signal stream. This jitter stress source may not be present in all stressed pattern generators or bit error ratio testers. It can be generated by driving the pattern generator external jitter modulation input with a filtered PRBS pattern. The PRBS pattern length should be between PRBS7 and PRBS9. The PRBS signaling rate should be approximately 1/10 of the stressed signal's signaling rate (i.e., approximately 2.578 GBd). The clock source for the PRBS generator is asynchronous to the pattern generator clock source to assure non-correlation of the jitter. The low pass filter that operates on the PRBS pattern to generate the bounded uncorrelated jitter should exhibit 20 dB/decade roll-off with a –3 dB corner frequency between 150 MHz and 300 MHz. This value must also be below the upper frequency limit of the pattern generator external modulator input. Random jitter and bounded uncorrelated jitter are added such that the output of the pattern generator approximates a jitter profile given in Table 83E–9. The target pattern generator 20% to 80%

transition time in the module stressed input test is 9.5 ps. The return loss of the test system as measured at TP1 meets the specification given in Equation (83E–2).

Table 83E–9—Pattern generator jitter characteristics

Parameter	Value
Total Jitter (pk-pk)[a]	0.28 UI
Random Jitter (pk-pk)[b]	0.15 UI
Max even-odd jitter (pk-pk)[c]	0.035 UI

[a]Total Jitter at BER of 10^{-15}.
[b]Random Jitter at BER of 10^{-15}.
[c]As defined in 92.8.3.8.1.

The counter propagating crosstalk channels during calibration of the stressed signal are asynchronous with target amplitude of 900 mV peak-to-peak differential and 20% to 80% target transition time of 12 ps as measured at TP4. The crosstalk signal transition time is calibrated with Pattern 4. The pattern is changed to Pattern 5 (with or without FEC encoding), Pattern 3, or a valid 100GBASE-R signal for amplitude calibration and the stressed input test. Patterns 3, 4, and 5 are described in Table 86–11. For the case where Pattern 3 is used with a common clock, there is at least 31 UI delay between the PRBS31 patterns on one lane and any other lane. Any one of these patterns is sufficient as a crosstalk aggressor with all lanes being active during the stressed input test.

Two levels of frequency-dependent attenuation are used for the module stressed input test: high loss and low loss. For the high loss case, frequency-dependent attenuation is added such that the loss at 12.89 GHz from the output of the pattern generator to TP1a is 13.8 dB. The 13.8 dB loss represents 10.25 dB channel loss with an additional allowance for host transmitter package loss. Eye height and eye width, extrapolated to a probability of 10^{-15}, are then measured at TP1a based on the eye measurement methodology given in 83E.4.2. Random jitter and the pattern generator output amplitude are adjusted (without exceeding the differential pk-pk input voltage tolerance specification as shown in Table 83E–7) to result in the eye height and eye width given in Table 83E–8 using the reference receiver with the setting of the CTLE that maximizes the product of eye height and eye width. For the low loss case, discrete frequency-dependent attenuation is removed such that the connection from the output of the pattern generator to TP1a comprises the mated HCB/MCB pair as described in 83E.4.1. Eye height and eye width at TP1a are then adjusted in the same way as described for the high loss case. In both the low loss and high loss cases, the module under test is provided with the reference CTLE setting used to meet eye width and eye height requirements via the variable *Recommended_CTLE_value*. If a Clause 45 MDIO is implemented, this variable is accessible through register 1.179 (see 45.2.1.96). The pattern is then changed to Pattern 5 (with or without FEC encoding), Pattern 3, or a valid 100GBASE-R signal for the input test, which is conducted by inserting the module into the MCB. The module CAUI-4 receiver under test shall meet the BER requirement as described in 83E.1.1 using three *Recommended_CTLE_value* values for both the high loss test and low loss test:

a) The CTLE setting used to meet eye width and eye height requirements
b) The value 1 dB higher if present in Table 83E–2
c) The value 1 dB lower if present in Table 83E–2

Modules may optionally elect not to use the *Recommended_CTLE_value.*

83E.4 CAUI-4 measurement methodology

This subclause describes common measurement tools and methodologies to be used for the CAUI-4 chip-to-module interface. Details of HCB and MCB characteristics are given in 83E.4.1 and details of the eye diagram measurement methodology are given in 83E.4.2.

83E.4.1 HCB/MCB characteristics

HCB characteristics are described in 92.11.1 where the HCB performs the equivalent function as the TP2 or TP3 test fixture. The MCB characteristics are described in 92.11.2 where the MCB performs the equivalent functionality as the cable assembly test fixture.

83E.4.2 Eye width and eye height measurement method

Eye diagrams in CAUI-4 chip-to-module are measured using a reference receiver. The reference receiver includes a fourth-order Bessel-Thomson low-pass filter response with 33 GHz 3 dB bandwidth, and a selectable continuous time linear equalizer (CTLE) to measure eye height and width. The pattern used for output eye diagram measurements is Pattern 4. The following procedure should be used to obtain eye height and eye width parameters:

1) Capture Pattern 4 using a clock recovery unit with a corner frequency of 10 MHz and slope of 20 dB/decade and a minimum sampling rate of 3 samples per bit. Collect sufficient samples equivalent to at least 4 million bits to allow for construction of a normalized cumulative distribution function (CDF) to a probability of 10^{-6} without extrapolation.
2) Apply the reference receiver including the appropriate CTLE to the captured signal. For modules, any single CTLE setting as described in 83E.3.2.1.1 that meets both eye width and eye height requirements is acceptable. For host compliance, the CTLE peaking in the reference receiver shall be set to three values:
   a) The recommended CTLE peaking value provided by the host
   b) The value 1 dB higher if present in Table 83E–2
   c) The value 1 dB lower if present in Table 83E–2

   A compliant host passes both the eye width and eye height A limit specified in Table 83E–1 using at least one of the settings and passes eye height B in Table 83E–1 at all of the two or three settings.
3) Use the differential equalized signal from step 2) to construct the CDF of the jitter zero crossing for both the left edge (CDFL) and right edge (CDFR), as a distance from the center of the eye. Calculate the eye width (EW6) as the difference in time between CDFR and CDFL with a value of 10^{-6}. CDFL and CDFR are calculated as the cumulative sum of histograms of the zero crossing samples at the left and right edges of the eye normalized by the total number of sampled bits. For a pattern with 50% transition density the maximum value for the CDFL and CDFR would be 0.5. The CDFL and CDFR are equivalent to bath tub curves where the BER is plotted versus sampling time.
4) Leveraging the Dual-Dirac jitter model described in 48B.1.1, estimate the random jitter. Calculate the best linear fit in Q-scale over the range of probabilities of 10^{-4} to 10^{-6} of the CDFL and CDFR to yield the random jitter on the left edge (RJL) and the random jitter on the right edge (RJR), respectively. The eye width is then given by Equation (83E–7).

$$EW15 = EW6 - 3.19 \times (RJR + RJL) \quad \text{(83E–7)}$$

where

$EW15$	is the eye width extrapolated to 10^{-15} probability
$EW6$	is the eye width at 10^{-6} probability
RJL	is the RMS value of the jitter estimated from CDFL
RJR	is the RMS value of the jitter estimated from the CDFR

5) Use the differential equalized signal from step 2) to construct the CDF of the signal voltage in the central 5% of the eye, for both logic 1 (CDF1) and logic 0 (CDF0), as a distance from the center of the eye. Calculate the eye height (EH6) as the difference in voltage between CDF1 and CDF0 with a value of 10^{-6}. CDF0 and CDF1 are calculated as the cumulative sum of histograms of the voltage at the top and bottom of the eye normalized by the total number of sampled bits. For a well-balanced number of ones and zeros, the maximum value for CDF0 and CDF1 will be 0.5.

6) Apply the Dual-Dirac and tail fitting techniques to CDF1 and CDF0 to estimate the noise at the middle of the eye. Calculate the best linear fit in Q-scale over the range of probabilities of 10^{-4} and 10^{-6} of CDF1 and CDF0 to yield relative noise one (RN1) and relative noise zero (RN0). The eye height is then given by Equation (83E–8)

$$EH15 = EH6 - 3.19 \times (RN0 + RN1) \quad \text{(83E–8)}$$

where

$EH15$	is the eye height extrapolated to 10^{-15} probability
$EH6$	is the eye height at 10^{-6} probability
$RN1$	is the RMS value of the noise estimated from CDF1
$RN0$	is the RMS value of the noise estimated from CDF0

83E.4.2.1 Vertical eye closure

Vertical eye closure is calculated using Equation (83E–9).

$$VEC = 20\log\left(\frac{AV}{EH15}\right) \quad \text{(83E–9)}$$

where

VEC	is vertical eye closure in dB
AV	is the eye amplitude of the equalized waveform. Eye amplitude is defined as the mean value of logic one minus the mean value of logic zero in the central 5% of the eye
$EH15$	is given in Equation (83E–8)

83E.5 Protocol implementation conformance statement (PICS) proforma for Annex 83E, Chip-to-module 100 Gb/s four-lane Attachment Unit Interface (CAUI-4)[29]

83E.5.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Annex 83E, Chip-to-module 100 Gb/s four-lane Attachment Unit Interface (CAUI-4), 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.

83E.5.2 Identification

83E.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).	

83E.5.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Annex 83E, Chip-to-module 100 Gb/s four-lane Attachment Unit Interface (CAUI-4)
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	

[29] *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.

83E.5.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
NOL	Number of differential AC-coupled lanes	83E.1	Four independent data paths in each direction	M	Yes [ ]
BER	Meets CAUI-4 BER requirement	83E.1.1	See 83E.1.1	M	Yes [ ]
ADR	Adaptive receiver	83E.3.4.1.1	Module CAUI-4 receiver does not use *Recommended_CTLE_value*	O	Yes [ ] No [ ]

83E.5.4 PICS proforma tables for chip-to-module 100 Gb/s four-lane Attachment Unit Interface (CAUI-4)

83E.5.4.1 Host output

Item	Feature	Subclause	Value/Comment	Status	Support
TH1	Signaling rate	83E.3.1.1	25.78125 GBd ± 100 ppm per lane	M	Yes [ ]
TH2	Peak-to-peak differential output voltage	83E.3.1.2	900 mV (max)	M	Yes [ ]
TH3	Peak-to-peak differential output voltage, transmitter disabled	83E.3.1.2	Less than or equal to 35 mV	M	Yes [ ]
TH4	DC common-mode output voltage	83E.3.1	Between –0.3 V and 2.8 V with respect to signal ground	M	Yes [ ]
TH5	AC common-mode output voltage	83E.3.1	17.5 mV RMS with respect to signal ground	M	Yes [ ]
TH6	Differential output return loss	83E.3.1.3	Meets Equation (83E–2) constraints	M	Yes [ ]
TH7	Reference impedance for output return loss	83E.3.1.3	100 Ω	M	Yes [ ]
TH8	Common to differential mode conversion	83E.3.1.3	Meets Equation (83E–3) constraints	M	Yes [ ]
TH9	Differential termination mismatch	83E.3.1	Less than 10%	M	Yes [ ]
TH10	Transition time	83E.3.1.5	Greater than or equal to 10 ps	M	Yes [ ]
TH11	Eye width	83E.3.1	0.46 UI	M	Yes [ ]
TH12	Eye height A	83E.3.1	95 mV	M	Yes [ ]
TH13	Eye height B	83E.3.1	80 mV	M	Yes [ ]
TH14	Crosstalk source	83E.3.1.6	Asynchronous crosstalk source using Pattern 5, Pattern 3, or valid 100GBASE-R signal	M	Yes [ ]

83E.5.4.2 Module output

Item	Feature	Subclause	Value/Comment	Status	Support
TM1	Signal rate	83E.3.1.1	25.78125 GBd ± 100 ppm per lane	M	Yes [ ]
TM2	Peak-to-peak differential output voltage	83E.3.1.2	900 mV (max)	M	Yes [ ]
TM3	AC common-mode output voltage	83E.3.1.2	17.5 mV RMS with respect to signal ground	M	Yes [ ]
TM4	Differential output return loss	83E.3.1.3	Meets Equation (83E–2) constraints	M	Yes [ ]
TM5	Reference impedance for output return loss	83E.3.1.3	100 Ω	M	Yes [ ]
TM6	Common to differential mode conversion	83E.3.1.3	Meets Equation (83E–3) constraints	M	Yes [ ]
TM7	Differential termination mismatch	83E.3.1.4	Less than 10%	M	Yes [ ]
TM8	Transition time	83E.3.1.5	Greater than or equal to 12 ps	M	Yes [ ]
TM9	Eye width	83E.3.2.1	0.57 UI	M	Yes [ ]
TM10	Eye height	83E.3.2.1	228 mV	M	Yes [ ]
TM11	Crosstalk source	83E.3.2.1	Asynchronous crosstalk source using Pattern 5, Pattern 3, or valid 100GBASE-R signal	M	Yes [ ]
TM12	Vertical eye closure	83E.4.2.1	5.5 dB (max)	M	Yes [ ]

83E.5.4.3 Host input

Item	Feature	Subclause	Value/Comment	Status	Support
RH1	CAUI-4 host input characteristics	83E.3.3	Table 83E–4	M	Yes [ ]

83E.5.4.4 Module input

Item	Feature	Subclause	Value/Comment	Status	Support
RM1	CAUI-4 module input characteristics	83E.3.4	Table 83E–7	M	Yes [ ]
RM2	BER requirement	83E.3.4.1.1	As 83E.1.1 with settings associated with *Recommended_CTLE_value*	M	Yes [ ]

Annex 85A

(informative)

40GBASE-CR4 and 100GBASE-CR10 TP0 and TP5 test point parameters

85A.1 Overview

Annex 85A provides information on parameters associated with test points TP0 and TP5 that may not be testable in an implemented system. TP0 and TP5 test points are illustrated in the 40GBASE-CR4 and 100GBASE-CR10 link block diagram of Figure 85–2.

85A.2 Transmitter characteristics at TP0

The specifications at TP0 are summarized in Table 85A–1.

Table 85A–1—Transmitter characteristics at TP0 summary

Parameter	Subclause reference	Value	Units
Signaling rate, per lane	85.8.3.9	10.3125 ± 100 ppm	GBd
Unit interval nominal	85.8.3.9	96.969697	ps
Differential peak-to-peak output voltage (max) with Tx disabled		30	mV
Common-mode voltage limits	72.7.1.4	0 to 1.9	V
Differential output return loss (min)	72.7.1.5	See Equation (72–4) and Equation (72–5)	dB
Common-mode output return loss (min)	72.7.1.6	See Equation (72–6)and Equation (72–7)	dB
AC common-mode output voltage (max, RMS)		30	mV
Transition time (20%–80%)	72.7.1.7	24 to 47	ps
Max output jitter (peak-to-peak)[a] Random jitter[b] Deterministic jitter Duty Cycle Distortion[c] Total jitter	72.7.1.9	 0.15 0.15 0.035 0.28	 UI UI UI UI

[a]Jitter is measured with emphasis off.
[b]Jitter is specified at BER 10^{-12}.
[c]Duty Cycle Distortion is considered part of the deterministic jitter distribution.

85A.3 Receiver characteristics at TP5

The receiver characteristics at TP5 are summarized in Table 85A–2.

Table 85A–2—Receiver characteristics at TP5 summary

Parameter	Subclause reference	Value	Units
Bit error ratio	85.8.4.3	10^{-12}	
Signaling rate, per lane	85.8.4.4	10.3125 ± 100 ppm	GBd
Unit interval (UI) nominal	85.8.4.4	96.969697	ps
Receiver coupling	85.8.4.5	AC	
Differential input peak-to-peak amplitude (max)	72.7.2.4	1200[a]	mV
Differential input return loss (min)[b]	72.7.2.5	See Equation (72–4) and Equation (72–5)	dB
Differential to common-mode input return loss	85.8.4	10 min from 10 MHz to 10 GHz	dB

[a]The receiver shall tolerate amplitudes up to 1600 mV without permanent damage.
[b]Relative to 100 Ω differential.

85A.4 Transmitter and receiver differential printed circuit board trace loss

With the insertion loss TP0 to TP2 or TP3 to TP5 given in 85.8.3.4 and an assumed mated connector loss of 1.74 dB, the maximum insertion loss allocation for the transmitter and receiver differential controlled impedance printed circuit boards for each differential lane (i.e., the maximum value of the sum of the insertion losses from TP0 to the MDI host receptacle and from TP5 to the MDI host receptacle) are determined using Equation (85A–1). The maximum insertion loss allocation for the transmitter and receiver differential controlled impedance printed circuit boards is 7 dB at 5.15625 GHz. The maximum insertion loss for the transmitter or the receiver differential controlled impedance printed circuit board is one half of the maximum insertion loss $IL_{PCBmax}(f)$.

$$IL_{\mathrm{PCB}}(f) \le IL_{\mathrm{PCBmax}}(f) = (0.3)[20\log_{10}(\mathrm{e})(b_1\sqrt{f} + b_2 f + b_3 f^2 + b_4 f^3)] \ \text{(dB)} \qquad \text{(85A–1)}$$

for 10 MHz ≤ f ≤ 7500 MHz.

where

f is the frequency in Hz
$IL_{\mathrm{PCB}}(f)$ is the insertion loss for the transmitter and receiver PCB
$IL_{\mathrm{PCBmax}}(f)$ is the maximum insertion loss for the transmitter and receiver PCB
b_1 2×10^{-5}
b_2 1.1×10^{-10}
b_3 3.2×10^{-20}
b_4 -1.2×10^{-30}
e ≈ 2.71828

The minimum insertion loss allocation for the transmitter and receiver differential controlled impedance printed circuit boards for each differential lane (i.e., the minimum value of the sum of the insertion losses from TP0 to MDI receptacle and TP5 to MDI receptacle) are determined using Equation (85A–2) and the coefficients b_1 through b_4 are given in Equation (85A–1). The minimum insertion loss for the transmitter or the receiver differential controlled impedance printed circuit board is one half of the minimum insertion loss $IL_{\mathrm{PCBmin}}(f)$.

$$IL_{\mathrm{PCB}}(f) \geq IL_{\mathrm{PCBmin}}(f) = (0.0574)[20\log_{10}(\mathrm{e})(b_1\sqrt{f} + b_2 f + b_3 f^2 + b_4 f^3)] \text{ (dB)} \quad (85A–2)$$

for 10 MHz ≤ f ≤ 7500 MHz.

where

f is the frequency in Hz
$IL_{\mathrm{PCB}}(f)$ is the insertion loss for the transmitter and receiver PCB
$IL_{\mathrm{PCBmin}}(f)$ is the minimum insertion loss for the transmitter and receiver PCB

85A.5 Channel insertion loss

This subclause provides information on channel insertion losses for intended topologies ranging from 0.5 m to 7 m in length. The maximum channel insertion loss associated with the 7 m topology is determined using Equation (85A–3). The channel insertion loss associated with the 0.5 m topology and a maximum host channel is determined by Equation (85A–4). The channel insertion loss budget at 5.15625 GHz is illustrated in Figure 85A–1.

The maximum channel insertion loss is determined using Equation (85A–3). The maximum channel insertion loss is 24.44 dB at 5.15625 GHz.

$$IL_{\mathrm{Chmax}}(f) = IL_{\mathrm{Camax}}(f) + 2IL_{\mathrm{Host}}(f) - 2IL_{\mathrm{MatedTF}}(f) \text{ (dB)} \quad (85A–3)$$

for 50 MHz ≤ f ≤ 7500 MHz.

where

f is the frequency in MHz
$IL_{\mathrm{Ch}max}(f)$ is the maximum channel insertion loss between TP0 and TP5
$IL_{\mathrm{Camax}}(f)$ is the maximum cable assembly insertion loss using Equation (85–19)
$IL_{\mathrm{Host}}(f)$ is the maximum insertion loss from TP0 to TP2 or TP3 to TP5 using Equation (85–14)
$IL_{\mathrm{MatedTF}}(f)$ is the maximum insertion loss of the mated test fixture using Equation (85–36)

The channel insertion loss between TP0 and TP5 representative of a 0.5 m cable assembly and a maximum host channel is determined using Equation (85A–4).

$$(IL_{\mathrm{Ch0.5}m}(f) = 0.275IL_{\mathrm{Camax}}(f) + 2IL_{\mathrm{Host}}(f) - 2IL_{\mathrm{MatedTF}}(f) \text{ (dB)} \quad (85A–4)$$

for 50 MHz ≤ f ≤ 7500 MHz.

where

f is the frequency in MHz

$IL_{Ch0.5m}(f)$ is the channel insertion loss between TP0 and TP5 representative of a 0.5 m cable assembly and a maximum host channel

$IL_{Camax}(f)$ is the maximum cable assembly insertion loss from TP1 to TP4 using Equation (85–19)

$IL_{Host}(f)$ is the maximum insertion loss from TP0 to TP2 or TP3 to TP5 using Equation (85–14)

$IL_{MatedTF}(f)$ is the maximum insertion loss of the mated test fixture using Equation (85–36)

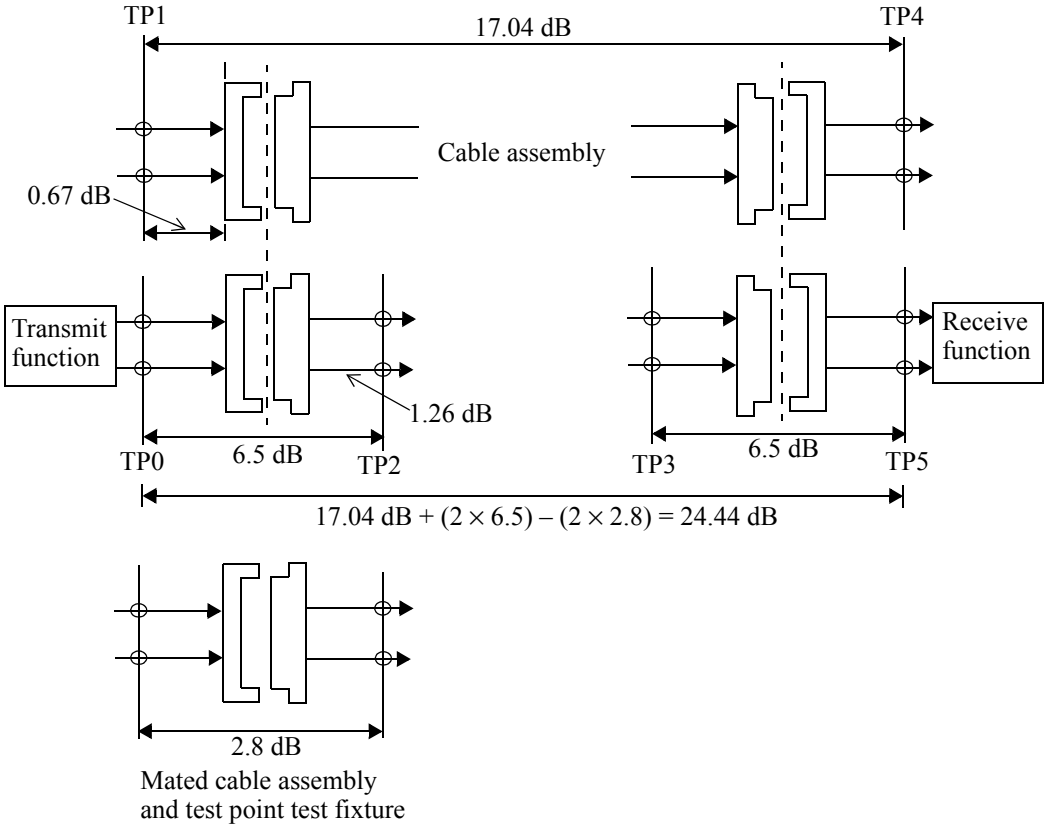

Figure 85A–1—Illustration of channel insertion loss budget at 5.15625 GHz

85A.6 Channel return loss

The return loss of each lane of the 40GBASE-CR4 or 100GBASE-CR10 channel is recommended to meet the values determined using Equation (85–25).

85A.7 Channel insertion loss deviation (ILD)

The channel insertion loss deviation is the difference between the channel insertion loss and the fitted channel insertion loss is determined using Equation (85A–5).

$$ILD(f) = IL_{Ch}(f) - IL_{fitted}(f) \text{ (dB)} \quad (85A–5)$$

where

f is the frequency in MHz

$ILD(f)$ is the channel insertion loss deviation at frequency f

$IL_{\text{fitted}}(f)$ is defined in Equation (85–19)

Given the channel insertion loss is at N uniformly-spaced frequencies f_n spanning the frequency range 50 MHz to 7500 MHz with a maximum frequency spacing of 10 MHz, the coefficients of the fitted channel insertion loss are determined using Equation (85–20) and Equation (85–21).

The channel insertion loss deviation is recommended to be within the region defined by Equation (85–23) and Equation (85–24) for all frequencies from 50 MHz to 7500 MHz.

85A.8 Channel integrated crosstalk noise (ICN)

Since four lanes or ten lanes are used to transfer data between PMDs, the near-end crosstalk (NEXT) that is coupled into a victim receiver will be from the four or ten adjacent transmitters. The channel multiple disturber NEXT loss, $MDNEXT\_loss_1(f)$, is specified using the individual NEXT losses as shown in Equation (85–26).

In addition, the far-end Crosstalk (FEXT) that is coupled into a receiver will be from the three or nine other transmitters adjacent to the victim transmitter. The channel multiple disturber FEXT loss, $MDFEXT\_loss(f)$, is specified using the individual FEXT losses as shown in Equation (85–27).

Given the channel $MDNEXT\_loss(f)$ and $MDFEXT\_loss(f)$ measured over N uniformly-spaced frequencies f_n spanning the frequency range 50 MHz to 10000 MHz with a maximum frequency spacing of 10 MHz, the RMS value of the integrated crosstalk noise is determined using Equation (85–28) through Equation (85–32) and the parameters shown in Table 85–11.

The total integrated crosstalk RMS noise voltage of the channel is recommended to meet the values determined using Equation (85A–6) illustrated in Figure 85A–2.

$$\sigma_{x,\,ch} \leq \left\{ \begin{array}{ll} 10 & 3 \leq IL \leq 7.5 \\ 13.4 - 0.45IL & 7.5 < IL \leq 24.44 \end{array} \right\} \quad \text{(mV)} \qquad \text{(85A–6)}$$

where IL is the value of the channel insertion loss in dB at 5.15625 GHz.

Figure 85A–2—Channel integrated crosstalk noise

Annex 86A

(normative)

Parallel Physical Interface (nPPI) for 40GBASE-SR4 and 40GBASE-LR4 (XLPPI) and 100GBASE-SR10 (CPPI)

86A.1 Overview

The Parallel Physical Interface (nPPI) is an optional instantiation of the PMD service interface for the PMDs in Clause 86, which is described in 86.2 and 80.3, or for the 40GBASE-LR4 PMD service interface described in 87.2 and 80.3. It allows the construction of compact optical transceiver modules for 40GBASE-SR4, 40GBASE-LR4, or 100GBASE-SR10 with no clock and data recovery circuits inside. The 40 Gb/s Parallel Physical Interface (XLPPI, four lanes) is used with 40GBASE-SR4 or 40GBASE-LR4, and the 100 Gb/s Parallel Physical Interface (CPPI, ten lanes) with 100GBASE-SR10. The term "nPPI" denotes either or both.

The PMD and PMA attached to the nPPI are required to comply with the delay, Skew, and Skew Variation requirements in 86.3, 83.5.3, and 83.5.4 as appropriate. The PMD MDIO function mapping given in 86.4 may apply. The PMD functional specifications are as given in 86.5.

This annex is arranged as follows: following the overview, a reminder of the block diagram, a brief introduction to the test points, and lane assignments, 86A.4 contains the electrical specifications for nPPI from host to module (Tx side) and then module to host (Rx side). Test points, compliance boards, and electrical parameters are defined in 86A.5. The host PCB ("channel") response recommendation is provided in 86A.6. Safety, installation, environment, and labeling is addressed in 86A.7 and PICS is provided in 86A.8. This annex is very closely related to Clause 86.

86A.2 Block diagram and test points

The PMD block diagram is shown in Figure 86–2. Figure 86–3 shows the test points.

The nPPI is standardized at the test points described in 86.8.1. The transmit side electrical signal (host electrical output and module electrical input) is defined at the output of the Host Compliance Board (TP1a), and other specifications of the module electrical input port are defined at the input of the Module Compliance Board (TP1). The receive side electrical signal (module electrical output and host electrical input) is defined at the output of the Module Compliance Board (TP4), and other specifications of the host electrical input port are defined at the input of the Host Compliance Board (TP4a). Test points and compliance boards are defined more thoroughly in 86A.5.1.

86A.3 Lane assignments

There are no lane assignments (within a group of transmit or receive lanes) for XLPPI or CPPI. While it is expected that a PMD will map electrical lane *i* to optical lane *i* and vice versa, there is no need to define the physical ordering of the lanes, as the PCS is capable of receiving the lanes in any arrangement.

86A.4 Electrical specifications for nPPI

The signaling rate for a lane of an XLPPI or CPPI interface shall be as defined in Table 86–2. 86A.4.1 and 86A.4.2 specify the host to module (Tx side) and module to host (Rx side) respectively of the nPPI. Parameters are defined in 86A.5 and 86.8. A recommended PCB ("channel") response for the host (PMA) is provided in 86A.6. Test points are defined in 86A.5.1.

86A.4.1 nPPI host to module electrical specifications

Each output lane and signal of the nPPI host (PMA), if measured at TP1a (see 86A.5.1) with the specified crosstalk signals applied on all input lanes, shall meet the specifications of Table 86A–1 per the definitions in 86A.5. Each lane of the nPPI module (PMD) electrical input, if measured at TP1 and TP1a with all Rx lanes (module output) operating, shall meet the specifications of Table 86A–2 per the definitions in 86A.5. The module electrical input shall be AC-coupled, i.e., it shall present a high DC common-mode impedance at TP1. There may be various methods for AC-coupling in actual implementations.

Table 86A–1—nPPI host electrical output specifications at TP1a

Parameter description	Min	Max	Units	Conditions
Single ended output voltage	–0.3	4	V	Referred to signal common
AC common-mode output voltage	—	15	mV	RMS
Termination mismatch at 1 MHz	—	5	%	
Differential output return loss	See 86A.4.1.1	—	dB	
Output transition time, 20% to 80%	28	—	ps	
J2 Jitter output	—	0.17	UI	
J9 Jitter output	—	0.29	UI	
Data Dependent Pulse Width Shrinkage (DDPWS)	—	0.07	UI	
Q_{sq} for XLPPI	45	—	V/V	
Q_{sq} for CPPI	43	—	V/V	
	Specification values			
Eye mask coordinates: X1, X2 Y1, Y2	0.11, 0.31 95, 350		UI mV	Hit ratio = 5×10^{-5}
Crosstalk source VMA, each input lane	700		mV	At TP4
Crosstalk source transition times, 20% to 80%	34		ps	At TP4

Table 86A–2—nPPI module electrical input specifications at TP1 and TP1a

Parameter description	Test point	Min	Max	Units	Conditions
Single ended input voltage tolerance[a]	TP1a	–0.3	4	V	Referred to TP1 signal common
AC common-mode input voltage tolerance	TP1a	15	—	mV	RMS
Differential input return loss	TP1	See 86A.4.1.1	—	dB	10 MHz to 11.1 GHz
Differential to common-mode input return loss	TP1	10	—	dB	10 MHz to 11.1 GHz
J2 Jitter tolerance	TP1a	0.17	—	UI	
J9 Jitter tolerance	TP1a	0.29	—	UI	
Data Dependent Pulse Width Shrinkage (DDPWS) tolerance	TP1a	0.07	—	UI	
		Specification values			
Eye mask coordinates: X1, X2 Y1, Y2	TP1a	0.11, 0.31 95, 350		UI mV	Hit ratio = 5×10^{-5}
Crosstalk calibration signal VMA	TP4	850		mV	While calibrating compliance signal[b]
Crosstalk calibration signal transition times, 20% to 80%	TP4	34		ps	

[a] The single ended input voltage tolerance is the allowable range of the instantaneous input signals.
[b] The crosstalk calibration signals are applied to the mated HCB-MCB at TP4a and measured at TP4 following the same principles as the host electrical input calibration (see 86A.5.3.8.5). They are removed before testing.

86A.4.1.1 Differential return losses at TP1 and TP1a

From 10 MHz to 11.1 GHz, the magnitude in decibels of the module differential input return loss at TP1 and the host differential output return loss at TP1a (see 86A.5.1) shall not exceed the limit given in Equation (86A–1) and illustrated in Figure 86A–1.

$$Return\_loss(f) \geq \begin{cases} 12 - 2\sqrt{f} & 0.01 \leq f < 4.11 \\ 6.3 - 13\log_{10}\left(\frac{f}{5.5}\right) & 4.11 \leq f \leq 11.1 \end{cases} \text{ dB} \qquad (86A–1)$$

where
Return_loss(*f*)is the return loss at frequency *f*
f is the frequency in GHz

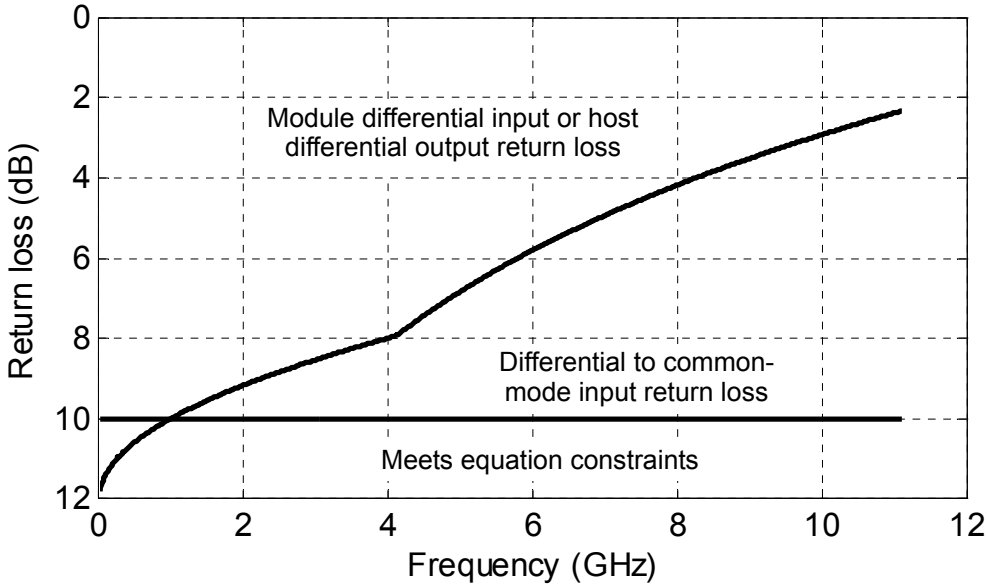

Figure 86A–1—Return loss specifications

86A.4.2 nPPI module to host electrical specifications

Each electrical output lane and signal of the nPPI module (PMD), if measured at TP4 shall meet the specifications of Table 86A–3 per the definitions in 86A.5 while the specified crosstalk sources are applied to the module's electrical input Each lane of the nPPI host (PMA) input shall meet the specifications of Table 86A–4 at TP4 and/or TP4a per the definitions in 86A.5. The module electrical output shall be AC-coupled, i.e., it shall present a high DC common-mode impedance at TP4. There may be various methods for AC-coupling in actual implementations..

Table 86A–3—nPPI module electrical output specifications at TP4

Parameter description	Min	Max	Units	Conditions
Single ended output voltage tolerance	–0.3	4	V	Referred to signal common
AC common-mode output voltage (RMS)	—	7.5	mV	
Termination mismatch at 1 MHz	—	5	%	
Differential output return loss	See 86A.4.2.1	—	dB	10 MHz to 11.1 GHz
Output transition time, 20% to 80%	28	—	ps	
J2 Jitter output	—	0.42	UI	
J9 Jitter output	—	0.65	UI	
	Specification values			
Eye mask coordinates: X1, X2 Y1, Y2	0.29, 0.5 150, 425		UI mV	Hit ratio = 5×10^{-5}
Crosstalk source VMA, each lane	700		mV	At TP1a
Crosstalk source transition times, 20% to 80%	37		ps	At TP1a

86A.4.2.1 Differential return losses at TP4 and TP4a

From 10 MHz to 11.1 GHz, the magnitude in decibels of the module differential output return loss at TP4 and the host differential input return loss at TP4a shall not exceed the limit given in Equation (86A–2) and illustrated in Figure 86A–1.

$$Return\_loss(f) \geq \begin{cases} 12 - 2\sqrt{f} & 0.01 \leq f < 4.11 \\ 6.3 - 13\log_{10}\left(\frac{f}{5.5}\right) & 4.11 \leq f \leq 11.1 \end{cases} \text{ dB} \qquad (86A–2)$$

where

Return_loss(*f*)is the return loss at frequency *f*

f is the frequency in GHz

Table 86A–4—nPPI host electrical input specifications at TP4 and TP4a

Parameter description	Test point	Min	Max	Units	Conditions
Single ended input voltage[a]	TP4	–0.3	4	V	Referred to signal common
AC common-mode input voltage tolerance	TP4	7.5	—	mV	RMS
Differential input return loss	TP4a	See 86A.4.2.1	—	dB	
Differential to common-mode input return loss	TP4a	10	—	dB	10 MHz to 11.1 GHz
Host input signal tolerance, interface BER limit		—	10^{-12}	—	
Conditions of host electrical receiver signal tolerance test:[b]					
		Specification values			
Eye mask coordinates: X1, X2 Y1, Y2	TP4	0.29, 0.5 150, 425		UI mV	Hit ratio = 5×10^{-5}
Transition time, 20% to 80%		34		ps	
J2 Jitter	TP4	0.42		UI	
J9 Jitter	TP4	0.65		UI	
Data Dependent Pulse Width Shrinkage (DDPWS)		0.34		UI	
VMA of aggressor lanes	TP4	850		mV	
Crosstalk calibration signal VMA	TP1a	700		mV	
Crosstalk calibration signal transition times, 20% to 80%	TP1a	37		ps	

[a] The host is required to tolerate (work correctly with) input signals with instantaneous voltages anywhere in the specified range.
[b] The specification values are test conditions for measuring signal tolerance and are not characteristics of the host (see 86A.5.3.8).

86A.5 Definitions of electrical parameters and measurement methods

Test points are defined in 86A.5.1, compliance boards in 86A.5.1.1, test patterns in 86A.5.2 and parameters in 86A.5.3 and 86.8. Multi-lane testing considerations are given in 86.8.2.1.

86A.5.1 Test points and compliance boards

Figure 86–3 shows the six test points for 40GBASE-SR4 and 100GBASE-SR10. These are TP1, TP1a, TP2, TP3, TP4, and TP4a; four of these are Skew points SP2, SP3, SP4, and SP5 as shown. For 40GBASE-LR4, points TP1, TP1a, TP4, TP4a, SP2, and SP5 are in equivalent positions. Figure 86–3 also shows the substitution of compliance boards for module (PMD) or host (PMA). These compliance boards are defined to connect generic test equipment to the module and host for test purposes. The module can be plugged into

the Module Compliance Board (MCB), which has specified electrical parameters. The Host Compliance Board (HCB), which also has specified electrical parameters, can plug into the host. The MCB and the HCB can be plugged together for calibration of compliance signals and to check the electrical parameters of the boards. Table 86A–5 shows the parameters or signals measured at each point. Also, TP0 and TP5 define the host ends of the electrical channel, at the PMA IC.

All electrical measurements and parameters are defined as through HCB and/or MCB as appropriate, with corrections if necessary.

CAUTION

A PMD with an nPPI interface is AC-coupled; however, an HCB is not. The user should take care that the test equipment does not improperly load or damage a host under test.

Table 86A–5—Parameters defined at each test point

Test point	Direction	Parameter
TP1	Looking downstream into module transmitter input	Module transmitter input return loss
TP1a	Looking upstream into host transmitter output	Host transmitter output signal and output return loss, module transmitter compliance signal calibration, host receiver compliance crosstalk signal calibration
TP4	Looking upstream into module receiver output	Module receiver output signal and output return loss, host receiver compliance signal calibration
TP4a	Looking downstream into host receiver input	Host receiver input return loss

86A.5.1.1 Compliance board parameters

The electrical characteristics of the HCB and MCB are given in 86A.5.1.1.1 and 86A.5.1.1.2. If boards are used that do not match the specifications given, the measurement results for nPPI shall be corrected for the differences between the actual HCB or MCB's properties and the reference differential insertion losses given in 86A.5.1.1.1. As it may be impractical to correct eye measurements for a board with differential insertion loss outside the limits given in 86A.5.1.1.2, such boards shall not be used. Boards that do not meet the specifications for mated HCB-MCB in 86A.5.1.1.2 shall not be used.

86A.5.1.1.1 Reference insertion losses of HCB and MCB

The reference differential insertion loss in decibels of the HCB and of the MCB, excluding the module connector, are given in Equation (86A–3) and Equation (86A–4) respectively, and illustrated in Figure 86A–2.

For the HCB,

$$Insertion\_loss(f) = 0.01 + 0.3\sqrt{f} + 0.11f \quad \text{dB} \qquad 0.01 \le f \le 11.1 \qquad \text{(86A–3)}$$

where

Insertion_loss(*f*)is the insertion loss at frequency *f*

f is the frequency in GHz

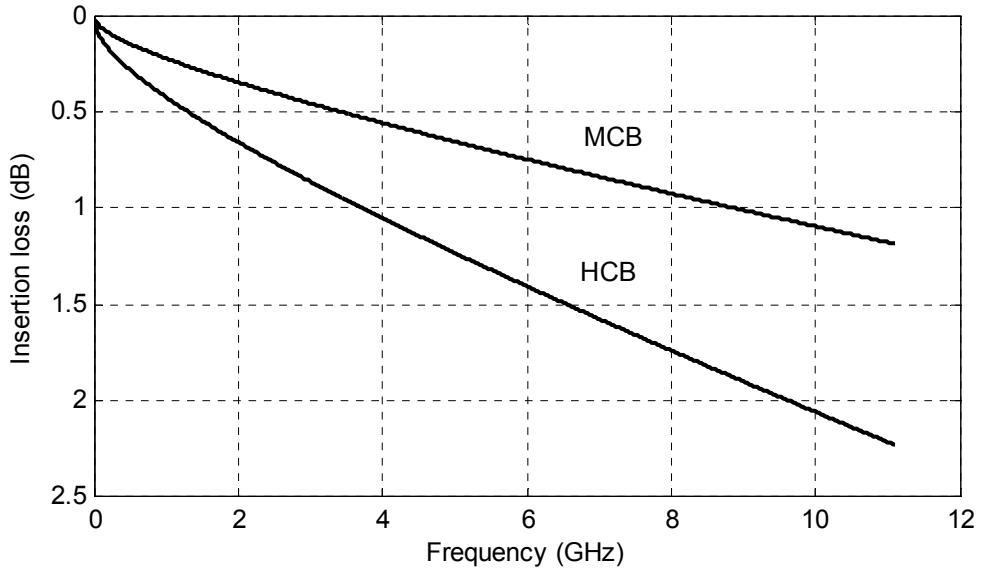

Figure 86A–2—Reference differential insertion losses of HCB, MCB excluding connector

For the MCB,

$$Insertion\_loss(f) = 0.0006 + 0.16\sqrt{f} + 0.0587f \quad \text{dB} \qquad 0.01 \le f \le 11.1 \tag{86A–4}$$

where

Insertion_loss(*f*)is the insertion loss at frequency *f*
f is the frequency in GHz

86A.5.1.1.2 Electrical specifications of mated HCB and MCB

The limits on the differential insertion loss in decibels of the mated HCB and MCB (in either direction) are given in Equation (86A–5) (which defines the minimum insertion loss) and Equation (86A–6) (which defines the maximum insertion loss). These limits are illustrated in Figure 86A–3.

$$Insertion\_loss(f) \ge -0.11 + 0.46\sqrt{f} + 0.16f \quad \text{dB} \qquad 0.01 \le f \le 11.1 \tag{86A–5}$$

$$Insertion\_loss(f) \le \left\{ \begin{array}{ll} 0.029 + 0.861\sqrt{f} + 0.158f & 0.01 \le f < 5.5 \\ 0.2 + 0.65f & 5.5 \le f \le 11.1 \end{array} \right\} \quad \text{dB} \tag{86A–6}$$

where

Insertion_loss(*f*)is the insertion loss at frequency *f*
f is the frequency in GHz

The limit on the differential return loss in decibels of the mated HCB and MCB, looking into the HCB or looking into the MCB, is given in Equation (86A–7), and illustrated in Figure 86A–4.

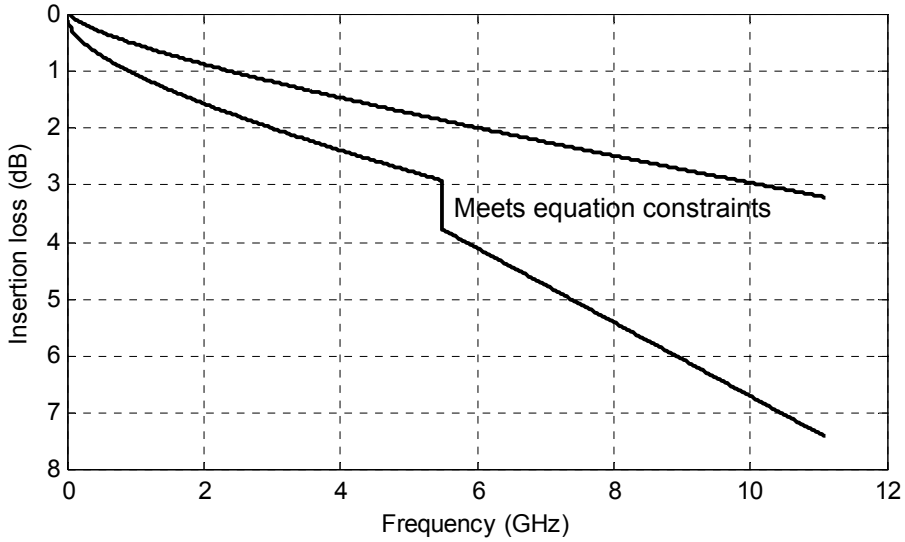

Figure 86A–3—Differential insertion loss limits of mated HCB-MCB

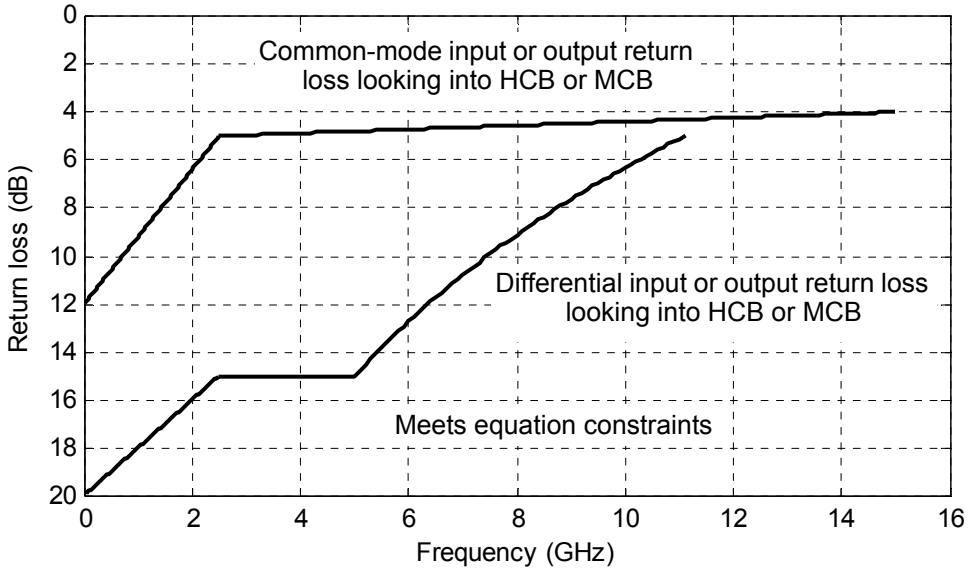

Figure 86A–4—Return loss limits of mated HCB-MCB

$$Return\_loss(f) \geq \begin{cases} 20 - 2f & 0.01 \leq f < 2.5 \\ 15 & 2.5 \leq f < 5 \\ 13.8 - 28.85\log_{10}\left(\frac{f}{5.5}\right) & 5 \leq f \leq 11.1 \end{cases} \text{ dB} \qquad (86A–7)$$

where

Return_loss(*f*)is the return loss at frequency *f*
f is the frequency in GHz

The limit on the common-mode return loss in decibels of the mated HCB and MCB is given in Equation (86A–8) and illustrated in Figure 86A–4.

$$Return\_loss(f) \geq \left\{ \begin{array}{ll} 12 - 2.8f & 0.01 \leq f < 2.5 \\ 5.2 - 0.08f & 2.5 \leq f \leq 15 \end{array} \right\} \text{ dB} \quad (86A–8)$$

where

Return_loss(*f*)is the return loss at frequency *f*
f is the frequency in GHz

The limit on the differential to common-mode conversion loss in decibels of the mated HCB and MCB, for input to HCB and output from MCB, or input to MCB and output from HCB, is given in Equation (86A–9) and illustrated in Figure 86A–5.

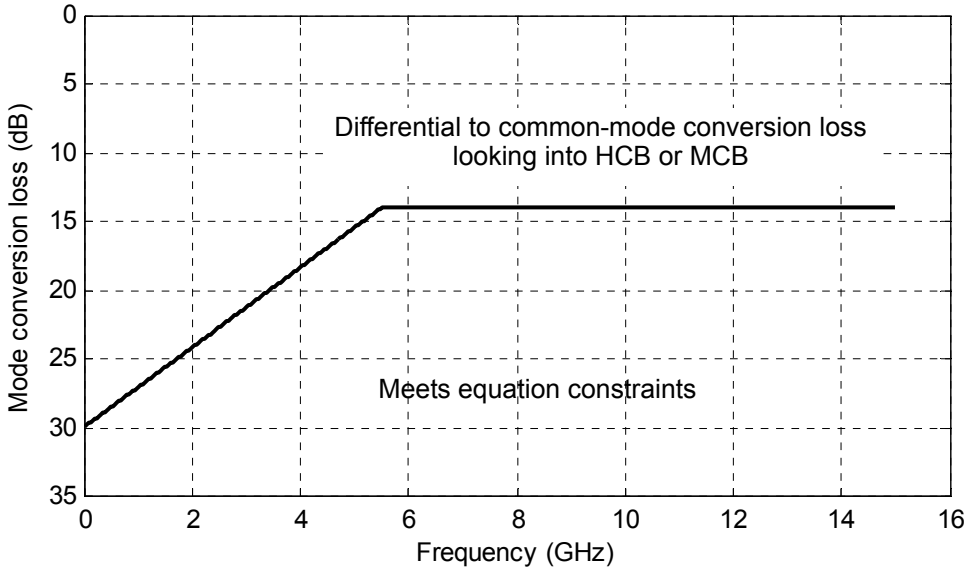

Figure 86A–5—Differential to common-mode conversion loss limit of mated HCB-MCB

$$Mode\_conversion\_loss(f) \geq \left\{ \begin{array}{ll} 30 - 2.91f & 0.01 \leq f < 5.5 \\ 14 & 5.5 \leq f \leq 15 \end{array} \right\} \text{ dB} \quad (86A–9)$$

where

Mode_conversion_loss(*f*)is the mode conversion loss at frequency *f*
f is the frequency in GHz

The limits on integrated crosstalk noise of the mated HCB and MCB for XLPPI and CPPI are as specified in 85.10.9.4 for 40GBASE-CR4 and 100GBASE-CR10 respectively, with the exception that the frequency range is 0.01 GHz to 12 GHz.

86A.5.2 Test patterns and related subclauses

While compliance is to be achieved in normal operation, specific test patterns are defined for measurement consistency and to enable measurement of some parameters. Table 86–11 lists the defined test patterns, and Table 86A–6 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. Multi-lane testing considerations are

Table 86A–6—Test patterns and related subclauses

Parameter	Pattern	Related subclause
J2 Jitter	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.3.3.1
J9 Jitter	3 or 5	86.8.3.3.2
Data Dependent Pulse Width Shrinkage (DDPWS)	4	86A.5.3.4
AC common-mode voltage	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86A.5.3.1
Transition time	Square wave or 4	86A.5.3.3
Electrical waveform (eye mask)	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.3.2, 86A.5.3.6
Q_{sq}	Square wave or 4	86A.5.3.5
Host electrical receiver signal tolerance	3 or 5	86A.5.3.8

given in 86.8.2.1. As Pattern 3 is more demanding than Pattern 5 (which itself is the same or more demanding than other 40GBASE-R or 100GBASE-R bit streams) an item which is compliant using Pattern 5 is considered compliant even if it does not meet the required limit using Pattern 3.

86A.5.3 Parameter definitions

In addition to the parameter definitions in the following subclause, some definitions with dual use (both optical and electrical) are given in 86.8.3.

86A.5.3.1 AC common-mode voltage

The common-mode voltage of a differential signal at any time is the average of signal+ and signal– at that time. RMS AC common-mode voltage may be calculated by applying the histogram function over 1 UI to the common-mode signal. As AC common-mode generation is very sensitive to the cable or oscilloscope delay mismatch, it is recommended to delay match the oscilloscope inputs for any measurements.

86A.5.3.2 Termination mismatch

Termination mismatch is the percentage difference between the two low-frequency impedances to common of a differential electrical port. Termination mismatch is defined as shown in Equation (86A–10).

$$\Delta Z_M = 2 \times \frac{|Z_p - Z_n|}{Z_p + Z_n} \times 100\ \% \quad (86A–10)$$

Termination mismatch can be measured by applying a low-frequency test tone to the differential inputs as shown in Figure 86A–6. The test frequency must be high enough to overcome the high pass effects of any AC-coupling capacitors. The measured differential output or input impedance is designated by Z_{diff}.

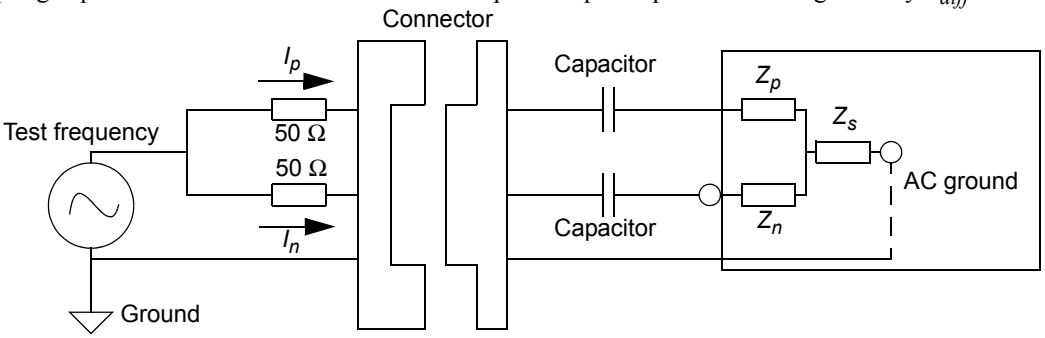

Figure 86A–6—Measurement of AC termination mismatch

For a 100 Ω port, low-frequency termination mismatch is then given by Equation (86A–11).

$$\Delta Z_M = 2 \times \frac{|I_p - I_n|}{I_p + I_n} \times \frac{Z_{diff} + 100}{Z_{diff}} \times 100\ \% \quad (86A–11)$$

where I_p and I_n are the currents flowing into the port as shown in Figure 86A–6.

Z_s is the effective series impedance between the terminations Z_p and Z_n and the AC ground.

86A.5.3.3 Transition time

In this annex, transition times (rise and fall times) are defined as the time between the 20% and 80% times, or 80% and 20% times, respectively, of isolated edges.

If the test pattern is the square wave with eight ones and eight zeros, the 0% level and the 100% level are as defined by the OMA measurement procedure (see 68.6.2).

If the test pattern is PRBS9, the transitions within sequences of five zeros and four ones, and nine ones and five zeros, respectively, are measured. These are bits 10 to 18 and 1 to 14, respectively, where bits 1 to 9 are the run of nine ones. In this case, the 0% level and the 100% level may be estimated as ZeroLevel and ZeroLevel + MeasuredOMA in the TWDP code (see 68.6.6.2), or by the average signal within windows from –3 UI to –2 UI and from 2 UI to 3 UI relative to the edge.

For electrical signals, the waveform is observed through a 12 GHz low-pass filter response (such as a Bessel-Thomson response).

NOTE—This definition is not the same as the rise and fall times typically reported by an oscilloscope from an eye diagram, which take all the edges into account.

86A.5.3.4 Data Dependent Pulse Width Shrinkage (DDPWS)

An oscilloscope, time interval analyzer, or other instrument with equivalent capability may be used to measure DDPWS. A clock recovery unit (CRU) as defined in 86.8.3.2 is used to trigger the oscilloscope. A repeating PRBS9 pseudo-random test pattern, 511 bits long, is used. For electrical jitter measurements, the

measurement bandwidth is 12 GHz (such as a Bessel-Thomson response). If the measurement bandwidth affects the result, it can be corrected for by post-processing. However, a bandwidth above 12 GHz is expected to have little effect on the results.

The crossing level is the average value of the entire waveform being measured. The instrument is synchronized to the pattern repetition frequency and the waveforms or the crossing times are averaged sufficiently to remove the effects of random jitter and noise in the system. The PRBS9 pattern has 128 positive-going transitions and 128 negative-going transitions. The crossing times t_1 to t_{256} of each transition of the averaged waveform (when the averaged waveform crosses the crossing level) are found as shown in Figure 86A–7.

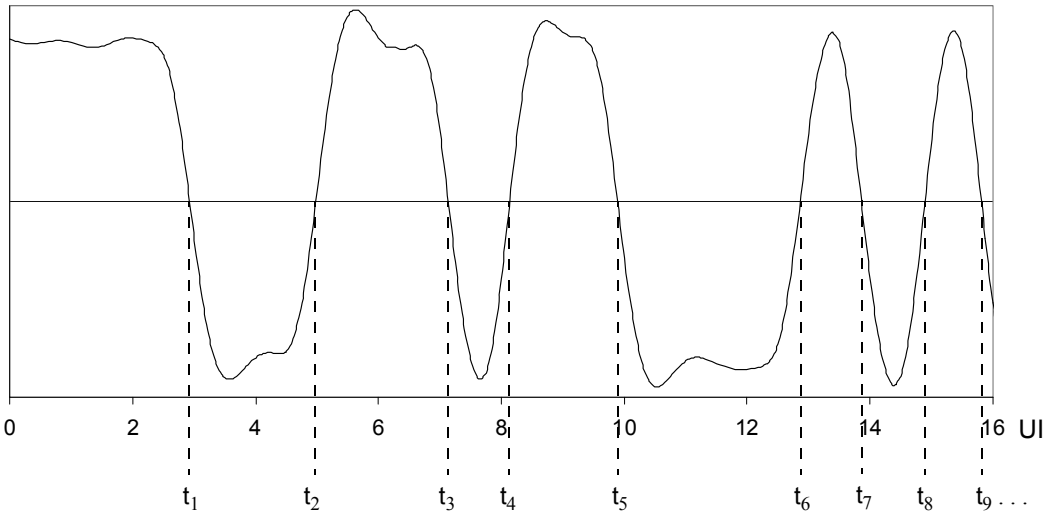

Figure 86A–7—Data Dependent Pulse Width Shrinkage test method

The DDPWS is the difference between one symbol period and the minimum of all the differences between pairs of adjacent transitions as follows in Equation (86A–12).

$$DDPWS = \mathrm{T} - \min(t_2 - t_1, t_3 - t_2, \ldots\ldots t_{256+1} - t_{256}) \quad (86A–12)$$

where T is one symbol period.

Note that the difference from the next edge in the repeating sequence, t_{256+1}, is also considered.

86A.5.3.5 Signal to noise ratio Q_{sq}

Q_{sq} is a measure of signal to noise ratio. For an electrical signal, it is analogous to Q_{sq} defined for an optical signal in 68.6.7, and it relates the low-frequency signal amplitude to the noise in an electrical –3 dB bandwidth of 12 GHz. It is defined with all co-propagating and counter-propagating crosstalk sources active, using one of patterns 3, 5, or a valid 40GBASE-SR4 or 100GBASE-SR10 signal. The input lanes of the item under test are receiving signals that are asynchronous to those being output.

Q_{sq} may be measured using an oscilloscope as follows:

a) The Voltage Modulation Amplitude (VMA) of the output lane is measured, using a square wave (8 ones, 8 zeros) or PRBS9 (Pattern 4). VMA is the difference between the 0% level and the 100% level defined in 86A.5.3.3; it is defined by analogy to OMA (see 68.6.2) but with a 12 GHz observation bandwidth.

b) Using the same pattern, the RMS noise over flat regions of the logic one and logic zero portions of the signal (see Figure 68–4 for the analogous optical measurement), is measured, compensating for noise in the measurement system. If possible, means should be used to prevent noise of frequency less than 1 MHz from affecting the result.
c) Q_{sq} is given by Equation (86A–13).

$$Q_{sq} = VMA / (n1 + n0) \quad (86A–13)$$

where

n1 is the RMS noise of logic one
n0 is the RMS noise of logic zero

86A.5.3.6 Eye mask for TP1a and TP4

The eye mask is defined by parameters X1, X2, Y1, and Y2. Unlike the optical eye mask, the vertical dimensions are fixed rather than scaled to the signal. Figure 83A–8 (an example of a hexagonal eye mask such as at TP1a) and Figure 83A–9 (a diamond mask such as at TP4) show the meaning of the parameters X1, X2, Y1 and Y2. The eye is defined as measured using a receiver with an electrical –3 dB bandwidth of 12 GHz (such as a Bessel-Thomson response). Further requirements are given in 86.8.3.2.

86A.5.3.7 Reference impedances for electrical measurements

The reference impedance for differential electrical measurements is 100 Ω and the reference impedance for common-mode electrical measurements is 25 Ω.

86A.5.3.8 Host input signal tolerance

To be compliant the host input signal tolerance shall satisfy the requirements of 86A.5.3.8.1 to 86A.5.3.8.6.

86A.5.3.8.1 Introduction

This subclause provides guidance for jitter tolerance testing at the host input (PMA) compliance point TP4/TP4a. Compliance is required with input jitter, vertical eye closure (Y1), and vertical peak level (Y2) as specified in Table 86A–4. Compliance is defined at an interface BER (the average of the BERs of all the lanes when stressed) of 10^{-12}. There are two test conditions: one each for the sensitivity and overload vertical eye parameters conditions. The reference test procedure is described in detail for a single stressed lane. Each Rx lane is tested in turn while all lanes are operated. Aggressor lanes are operated with the VMA specified in Table 86A–4.

86A.5.3.8.2 Test equipment and setup

A jitter tolerance test configuration is shown in Figure 86A–8. A test source is used to continuously generate the test signal. The test signal is conditioned within the guidelines outlined in 86A.5.3.8.3 to exhibit the appropriate jitter stress.

An RF attenuator or other output amplitude control of the test source may be required to set the vertical eye opening of the stressed eye.

The test equipment measured at TP4a looking towards the test signal source has better than 20 dB output return loss up to 12 GHz. The output return loss of the test system when measured at TP4 with the Module Compliance Board (looking towards the test signal source) is 2 dB better than the specifications of Table 86A–3 up to 8 GHz and 1 dB better up to 11 GHz.

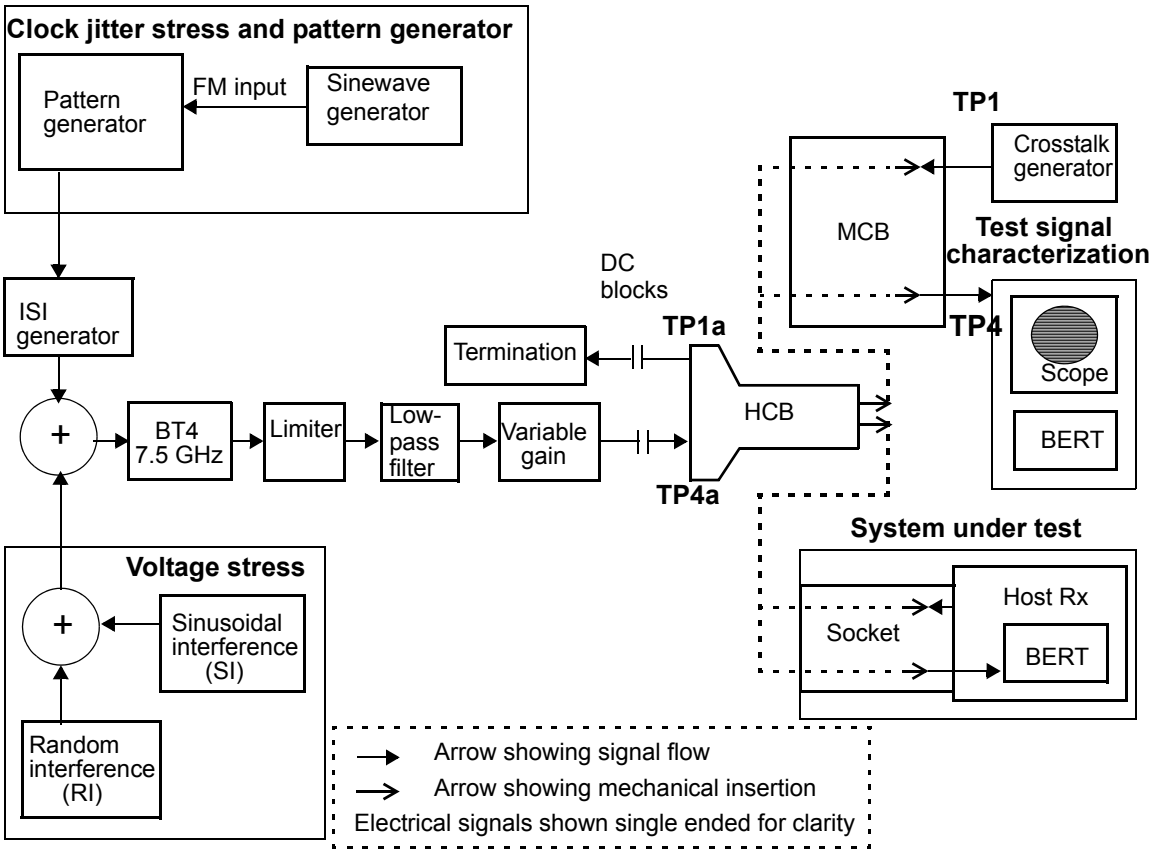

Figure 86A–8—Example jitter tolerance test configuration

86A.5.3.8.3 Stressed eye jitter characteristics

This subclause describes required test signal characteristics along with considerations and suggested approaches for test signal generation. The test signal is generated by the functions shown in Figure 86A–8 or by equivalent means. Figure 86A–9 illustrates how the jitter parameters in Table 86A–4 map to the jitter components in the stressed-eye test signal.

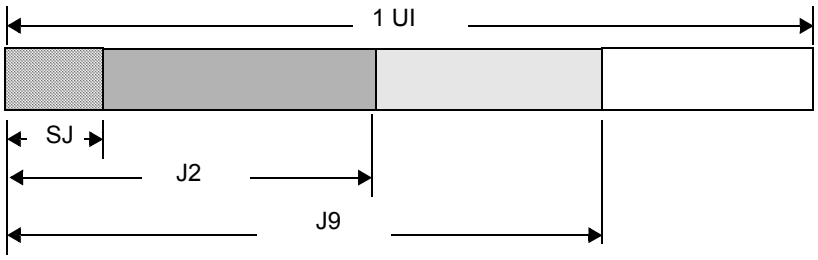

Figure 86A–9—Stressed eye jitter components

The 0.05 UI Sinusoidal Jitter (SJ) component of J2 Jitter is defined for frequencies much higher than the CDR bandwidth (e.g., ~20 MHz). At lower frequencies, the CDR must track additional applied SJ as detailed in the relevant specifications (see Figure 86A–10 and 52.8.1).

The balance of the J2 Jitter is created by the following mechanisms: an ISI generator, sinusoidal interference (SI), and random interference (RI), all passed through a limiting function.

The test signal at TP4 has DDPWS as defined by Table 86A–4. Any duty cycle distortion (DCD) in the test signal does not exceed 0.02 UI. DCD is the difference between the mean position of all falling edges and the mean position of all rising edges of a waveform. It is measured at the average level of the waveform in a 12 GHz bandwidth, using Pattern 3 (PRBS31), Pattern 4 (PRBS9), Pattern 5, or a valid 40GBASE-SR4 or 100GBASE-SR10 signal.

The ISI generator may be implemented by a low-pass filter, length of FR4 trace, length of coax cable, or other equivalent method. It is required that the resulting signal be passed through a limiter function. A suitable limiter function may be implemented using a discrete limiting amplifier followed by a low-pass filter and an attenuator. The low-pass filter emulates the bandwidth and/or slew rate of a practical limiter. The attenuator is used to set the output amplitude to minimum and maximum values allowed by the eye mask coordinates of Table 86A–4.

A voltage stress is to be applied before the limiter function. This stress is composed of a single tone sinusoidal interferer (SI) in the frequency range 100 MHz to 2 GHz and a broadband noise source (RI) with a minimum power spectrum –3 dB point of 6 GHz and minimum crest factor of 7. It is the intent that this combination of voltage stress and limiting function introduce pulse-shrinkage jitter behavior. However, no more than 20% of the J2 Jitter is created by the sinusoidal interferer.

Jitter generation mechanisms for the pattern generator are typically based on phase modulation of the clock source, edge modulation of a variable delay line or a combination thereof.

Any approach that modulates or creates the appropriate levels and frequencies of the jitter components is acceptable.

86A.5.3.8.4 Calibration

Calibration of the test signal is performed using the guidelines for test setup in 86A.5.3.8.2 and illustrated in Figure 86A–8. The aim of the calibration is to achieve a test signal exhibiting jitter stress in accordance with Table 86A–4.

The test signal is calibrated differentially into standard instrumentation loads. If complementary single-ended signals are used, they are carefully matched in both amplitude and phase.

For accurate calibration, it is imperative that all elements in the signal path (cables, DC blocks, etc.) have wide and smooth frequency response as well as linear phase response throughout the spectrum of interest. Baseline wander and overshoot/undershoot are minimized.

An AC-coupling –3 dB corner frequency of 20 kHz is expected to be adequate to eliminate baseline wander effects; however, high-frequency performance is critical and must not be sacrificed by the AC-coupling.

Given random jitter and the nature of the long test patterns, low probability jitter events will likely be present. It is recommended for jitter calibration that a technique that can accurately measure low probability events be used to avoid overly stressful test conditions.

It is recommended that the actual compliance test pattern be used during calibration. For jitter stress calibration it is permissible, however, to use any appropriate test pattern that still results in the creation of a compliance test pattern with the appropriate jitter stress.

86A.5.3.8.5 Calibration procedure

The vertical eye opening and peak level are set approximately to the levels specified in Table 86A–4.

With an applied calibration test pattern and no additional jitter stress applied, the intrinsic J2 Jitter and J9 Jitter of the test source due to intrinsic noise and finite bandwidth effects are measured. At this stage, J2 Jitter is less than 0.15 UI and J9 Jitter less than 0.25 UI.

SJ is added until the J2 Jitter increases by 0.05 UI above this intrinsic J2 Jitter level. The SJ jitter frequency is well above the CDR bandwidth and asynchronous to the characteristic frequencies of the signal.

Next, additional high probability jitter as specified in 86A.5.3.8.3 is added by the ISI generator until at least 80% of the J2 Jitter has been created. The sinusoidal interferer amplitude is then turned on and adjusted until the required level of J2 Jitter is achieved. The frequency of any sinusoidal interferer is asynchronous to the characteristic frequencies of the signal.

A compliant test signal exhibits Data Dependent Pulse Width Shrinkage (defined in 86A.5.3.4) as specified in Table 86A–4. This is measured with noise and clock-jitter sources turned off.

Once the required level of J2 Jitter has been achieved, turn on the crosstalk source that is set such that at the output of the Host Compliance Board, the amplitude and the transition time are as given in Table 86A–4. The crosstalk pattern is Pattern 3 (PRBS31), Pattern 5, or a valid 40GBASE-SR4 or 100GBASE-SR10 signal, and is asynchronous with the test signal. Then the RI (random interference) voltage stress is added until the specified value of J9 Jitter is achieved.

If necessary, the sine interferer is readjusted to obtain the required level of J2 Jitter and if the sinusoidal interferer is changed then the random interferer is readjusted to obtain the required level of J9 Jitter. Iterative adjustments of the sinusoidal interferer and random interferer are made until the required values of both J2 Jitter and J9 Jitter are achieved.

If necessary, the vertical eye opening is readjusted to required levels.

The vertical eye opening and peak level specifications are verified.

Care must be taken when characterizing the signal used to make receiver tolerance measurements. The intrinsic noise and jitter introduced by the calibration measurement equipment (e.g., filters, oscilloscope and BERT) must be accounted for and controlled. If equipment imperfections affect the results materially, corrections such as root-sum-of-squares deconvolution of Gaussian noise and jitter are used.

86A.5.3.8.6 Test procedure

Testing is performed differentially through a Host Compliance Board (see 86A.5.1).

Using a test signal arranged according to 86A.5.3.8.2 and calibrated according to 86A.5.3.8.5, operate the system with the test pattern specified in Table 86A–6. Each lane is tested in turn while all are operated. Aggressor lanes are operated with the VMA specified in Table 86A–4. The BERs of all the lanes when stressed are averaged to form the interface BER. See 86.8.2.1.

All signals and reference clocks that operate during normal operation are active during the test including all the other host lanes in both directions. The test signal and the host's transmitted signals are asynchronous. The host transmits Pattern 3 (PRBS31), Pattern 5, or a valid 40GBASE-SR4 or 100GBASE-SR10 signal. The sinusoidal jitter is stepped across the frequency and amplitude range according to Table 86A–7 and illustrated in Figure 86A–10, while monitoring the BER of the lane(s). The interface BER of a compliant host receiver remains below 10^{-12}.

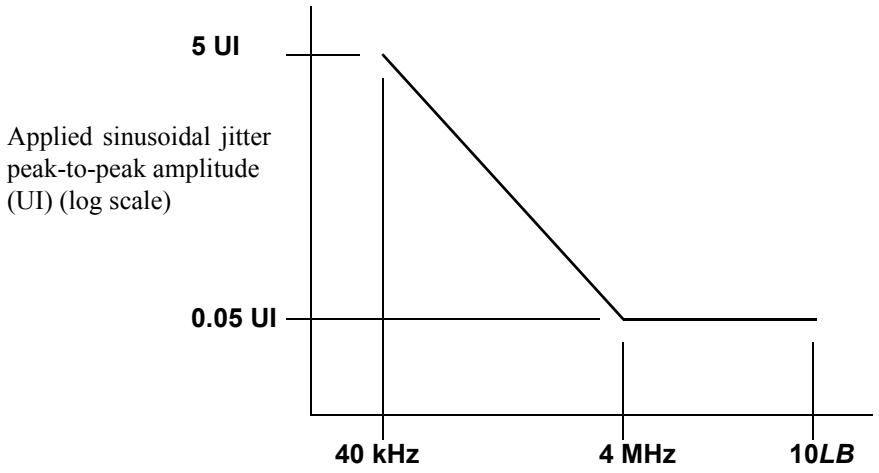

Figure 86A–10—Mask of the sinusoidal component of jitter tolerance

Table 86A–7—Applied sinusoidal jitter

Frequency range	Sinusoidal jitter, peak-to-peak, (UI)
$f < 40$ kHz	Not specified
40 kHz $\le f <$ 4 MHz	$2 \times 10^5 / f$
4 MHz $\le f \le 10\ LB$[a]	0.05

[a]*LB* = loop bandwidth; upper frequency bound for added sine jitter should be at least 10 times the loop bandwidth of the receiver being tested.

86A.6 Recommended electrical channel

The recommended limits for the differential insertion loss in decibels of the host PCB and connector mated to the HCB, between the PMA IC (TP0 or TP5) and TP1a or TP4a, are given in Equation (86A–14) and Equation (86A–15), and illustrated in Figure 86A–11. It is recommended that

$$Insertion\_loss(f) \le \left\{ \begin{array}{ll} 0.5 & 0.01 \le f < 0.11 \\ 0.114 + 0.8914\sqrt{f} + 0.846f & 0.11 \le f < 7 \\ -35.91 + 6.3291f & 7 \le f < 8 \\ 14.72 & 8 \le f \le 11.1 \end{array} \right\} \text{ dB} \qquad (86A–14)$$

and

$$Insertion\_loss(f) \ge \left\{ \begin{array}{ll} -0.22 + 0.46f & 0.01 \le f < 7 \\ 3 & 7 \le f \le 11.1 \end{array} \right\} \text{ dB} \qquad (86A–15)$$

where

Insertion_loss(*f*)is the insertion loss at frequency *f*

f is the frequency in GHz

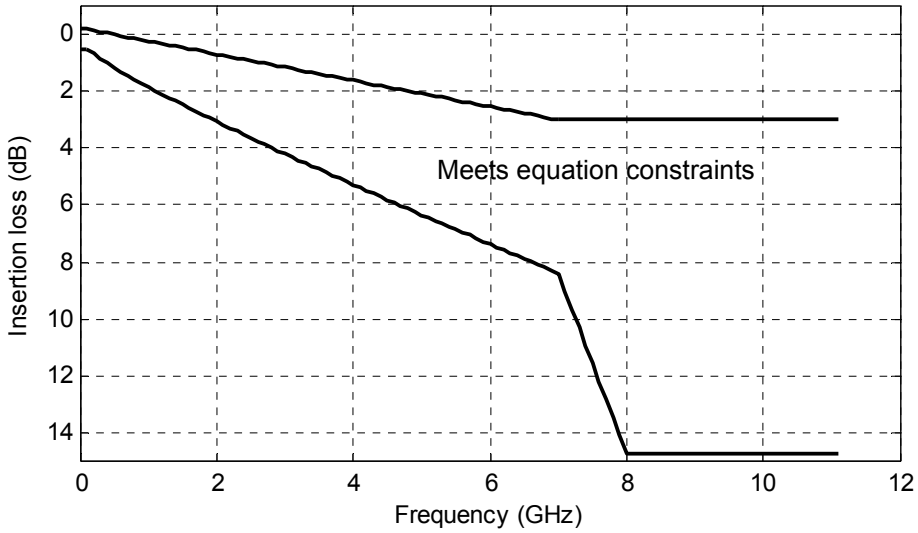

Figure 86A–11—Recommended insertion loss limits of host PCB, connector and HCB

The recommended maximum loss of the host PCB only (without connector or HCB) at 5.15625 GHz is 4.4 dB.

86A.7 Safety, installation, environment, and labeling

86A.7.1 General safety

All equipment subject to this annex shall conform to IEC 60950-1.

86A.7.2 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.

86A.7.3 Environment

The 40GBASE–SR4 and 100GBASE–SR10 operating environment specifications are as defined in 52.11, as defined in 52.11.1 for electromagnetic emission, and as defined in 52.11.2 for temperature, humidity, and handling.

86A.7.4 PMD labeling

The 40GBASE–SR4 and 100GBASE–SR10 labeling recommendations and requirements are as defined in 52.12.

86A.8 Protocol implementation conformance statement (PICS) proforma for Annex 86A, Parallel Physical Interface (nPPI) for 40GBASE-SR4 and 40GBASE-LR4 (XLPPI) and 100GBASE-SR10 (CPPI)[30]

86A.8.1 Introduction

The supplier of an XLPPI/CPPI implementation that is claimed to conform to Annex 86A, Parallel Physical Interface (nPPI) for 40GBASE-SR4 and 40GBASE-LR4 (XLPPI) and 100GBASE-SR10 (CPPI), 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.

86A.8.2 Identification

86A.8.2.1 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).	

86A.8.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2015, Annex 86A, Parallel Physical Interface (nPPI) for 40GBASE-SR4 and 40GBASE-LR4 (XLPPI) and 100GBASE-SR10 (CPPI)
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	

[30] *Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this annex so that it can be used for its intended purpose and may further publish the completed PICS.

86A.8.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*R4	40GBASE–R4	86A.4	Can operate as part of 40GBASE–SR4 or 40GBASE-LR4 PMA or PMD	O.1	Yes [ ] No [ ]
*R10	100GBASE–SR10	86A.4	Can operate as part of 100GBASE–SR10 PMA or PMD	O.1	Yes [ ] No [ ]
*MO	Module		Items marked with MO are applicable to a module (PMD)	O.2	Yes [ ] No [ ]
*HO	Host		Items marked with HO are applicable to a host (PMA)	O.2	Yes [ ] No [ ]
*MD	MDIO capability	86A.1	Registers and interface supported	O	Yes [ ] No [ ]

86A.8.4 PICS proforma tables for Parallel Physical Interface (nPPI) for 40GBASE-SR4 and 40GBASE-LR4 (XLPPI) and 100GBASE-SR10 (CPPI)

86A.8.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
SF1	Compatible with 40GBASE–R or 100GBASE–R PCS and PMA and 40GBASE–SR4, 40GBASE–LR4 or 100GBASE–SR10 PMD	86A.1		M	Yes [ ]

86A.8.4.2 Electrical specifications for nPPI

Item	Feature	Subclause	Value/Comment	Status	Support
S1	Signaling rate per Table 86–2	86A.4	10.3125 GBd ±100 ppm	M	Yes [ ]
S2	Host output signal per Table 86A–1	86A.4.1	Per definitions in 86A.5	HO:M	Yes [ ] N/A [ ]
S3	Module electrical input per Table 86A–2	86A.4.1	Per definitions in 86A.5	MO:M	Yes [ ] N/A [ ]
S4	Module electrical input AC-coupled	86A.4.1		MO:M	Yes [ ] N/A [ ]
S5	Module electrical differential input return loss at TP1	86A.4.1.1	Per Equation (86A–1)	MO:M	Yes [ ] N/A [ ]
S6	Host differential output return loss at TP1a	86A.4.1.1	Per Equation (86A–1)	HO:M	Yes [ ] N/A [ ]
S7	Module electrical output per Table 86A–3	86A.4.2	Per definitions in 86A.5	MO:M	Yes [ ] N/A [ ]
S8	Host input per Table 86A–4	86A.4.2	Per definitions in 86A.5	HO:M	Yes [ ] N/A [ ]
S9	Module electrical output AC-coupled	86A.4.2		MO:M	Yes [ ] N/A [ ]
S10	Module electrical output return loss at TP4	86A.4.2.1	Per Equation (86A–2)	MO:M	Yes [ ] N/A [ ]
S11	Host differential input return loss at TP4a	86A.4.2.1	Per Equation (86A–2)	HO:M	Yes [ ] N/A [ ]

86A.8.4.3 Definitions of parameters and measurement methods

Item	Feature	Subclause	Value/Comment	Status	Support
SEM1	Compliance boards	86A.5.1	Use compliance boards, correct as necessary	M	Yes [ ]
SEM2	Compliance boards	86A.5.1.1	Results corrected	M	Yes [ ]
SEM3	Compliance boards	86A.5.1.1	Individual insertion losses per 86A.5.1.1.2	M	Yes [ ]
SEM4	Compliance boards	86A.5.1.1	Mated HCB-MCB per 86A.5.1.1.2	M	Yes [ ]
SEM5	Host input signal tolerance	86A.5.3.8	As specified	HO:M	Yes [ ] N/A [ ]

86A.8.4.4 Environmental and safety specifications

Item	Feature	Subclause	Value/Comment	Status	Support
SES1	General safety	86A.7.1	Conforms to IEC 60950-1	M	Yes [ ]
SES3	Electromagnetic interference	86A.7.3	Complies with applicable local and national codes for the limitation of electromagnetic interference	M	Yes [ ]

Annex 91A

(informative)

RS-FEC codeword examples

This annex provides example RS-FEC codewords produced by the 64B/66B to 256B/257B transcoding and Reed-Solomon encoding defined in Clause 91. This annex presents data in a tabular form. The contents of the tables are transmitted from left to right within each row starting from the top row and ending at the bottom row. The tables contain both binary and hexadecimal representations of the data. For the hexadecimal representation, the most significant bit of each hex symbol is transmitted first.

91A.1 Input to the 64B/66B to 256B/257B transcoder

Table 91A–1 contains a sequence of 80 66-bit blocks corresponding to the PCS transmission of Idle control characters. The initial value of the scrambler was set to bits 6 to 63 of the first 64-bit payload in the first row of Table 74A–2. Bit 6 is assigned to S57 and bit 63 is assigned to S0 (see 49.2.6).

Table 91A–1—64B/66B to 256B/257B transcoder input

Sync <0:1>	64-bit payload, hex <2:65>	Sync <0:1>	64-bit payload, hex <2:65>	Sync <0:1>	64-bit payload, hex <2:65>	Sync <0:1>	64-bit payload, hex <2:65>
10	ad5a3bf86d9acf5c	10	de55cb85df0f7ca0	10	e6ccff8e8212b1c6	10	d63bc6c309000638
10	70e3b0ce30e0497d	10	dc8df31ec3ab4491	10	66fb9139c81cd37b	10	b57477d4f05e3602
10	8cfd495012947a31	10	e7777cf0c6d06280	10	44529cf4b4900528	10	85ce1d27750ad61b
10	456d5c71743f5c69	10	c1bf62e5dc5464b5	10	dc6011be7ea1ed54	10	1cf92c450042a75f
10	cc4b940eaf3140db	10	77bb612a7abf401f	10	c22d341e90545d98	10	ce6daf1f248bbd6d
10	dd22d0b3f9551ed6	10	574686c3f9e93898	10	2e52628f4a1282ce	10	f20c86d71944aab1
10	55133c9333808a2c	10	1aa825d8b817db4d	10	637959989f3021eb	10	976806641b26aae9
10	6a37d4531b7ed5f2	10	53c3e96d3b12fb46	10	528c7eb8481bc969	10	ab8f9980d5a54559
10	9a4d2abfda65cc33	10	94fe646efe5af02d	10	9a65ae5fcd88c03a	10	5ef08673168def9b
10	220c871a953fffc6	10	ce0bb95ac263e6c1	10	4f6a917d1a676571	10	5890918c7b687d75
10	44d2b3e43096f836	10	84cdd4fc48b79608	10	b3e4503e3c824a8c	10	fd6d0b1a39687929
10	1730167c08302a69	10	4c15ff56de92b1ad	10	d0c2f0d4ff0dee95	10	e1422ee2e8b92125
10	ed5acaf86592fcee	10	de799be0b903c880	10	2714ffbf40bc09f6	10	c3be97c3c285009f
10	1020faf19f606631	10	93007cabbb3f8c9d	10	ef6955f7f43df5d0	10	4dbd0616afe60e1f
10	3a1e49b7c7f7bb5d	10	901d828746ceec61	10	71ed3c097158c224	10	11adb3d81e13d263
10	a350d1a343b2394b	10	eab30ca27b5b34e3	10	90359ef711ed53d9	10	9b446763c8627ea8
10	6e891c0f4842b823	10	c4d786a25727a7fc	10	094fe7da31fb60cd	10	9f9a004de5e70767
10	054bdd77b7cb4e7b	10	c598cb710558af67	10	fc386d1f99d3a925	10	4928e0b43e781893
10	5a44dd3eb8b2ad6c	10	94462af4f583d770	10	8061ba9381f51f55	10	476d4eded7c90fcc
10	1efc25aa6a7e0b4c	10	93dd968c06a56809	10	9768e9d1ba74d3b6	10	014e9dc9f13670bb

91A.2 Output of the RS(528,514) encoder

Table 91A–2 contains a RS(528,514) codeword. Each row of Table 91A–1 is a set of four 66-bit blocks that is converted to one 257-bit block using the procedure defined in 91.5.2.5. The resulting set of 20 257-bit blocks constitute the message portion of the codeword. The parity is computed using the encoder defined in 91.5.2.7 and is appended to the message to complete the codeword.

Table 91A–2—RS(528,514) codeword

Header <0:4>	Payload, hex <5:64>	Payload, hex <65:128>	Payload, hex <129:192>	Payload, hex <193:256>
00101	a5a3bf86d9acf5c	de55cb85df0f7ca0	e6ccff8e8212b1c6	d63bc6c309000638
11110	7e3b0ce30e0497d	dc8df31ec3ab4491	66fb9139c81cd37b	b57477d4f05e3602
01111	8fd495012947a31	e7777cf0c6d06280	44529cf4b4900528	85ce1d27750ad61b
00110	46d5c71743f5c69	c1bf62e5dc5464b5	dc6011be7ea1ed54	1cf92c450042a75f
00100	c4b940eaf3140db	77bb612a7abf401f	c22d341e90545d98	ce6daf1f248bbd6d
10010	d22d0b3f9551ed6	574686c3f9e93898	2e52628f4a1282ce	f20c86d71944aab1
10001	5133c9333808a2c	1aa825d8b817db4d	637959989f3021eb	976806641b26aae9
00011	637d4531b7ed5f2	53c3e96d3b12fb46	528c7eb8481bc969	ab8f9980d5a54559
10100	94d2abfda65cc33	94fe646efe5af02d	9a65ae5fcd88c03a	5ef08673168def9b
00000	20c871a953fffc6	ce0bb95ac263e6c1	4f6a917d1a676571	5890918c7b687d75
01101	4d2b3e43096f836	84cdd4fc48b79608	b3e4503e3c824a8c	fd6d0b1a39687929
10011	130167c08302a69	4c15ff56de92b1ad	d0c2f0d4ff0dee95	e1422ee2e8b92125
00101	e5acaf86592fcee	de799be0b903c880	2714ffbf40bc09f6	c3be97c3c285009f
10010	120faf19f606631	93007cabbb3f8c9d	ef6955f7f43df5d0	4dbd0616afe60e1f
10001	31e49b7c7f7bb5d	901d828746ceec61	71ed3c097158c224	11adb3d81e13d263
00101	a50d1a343b2394b	eab30ca27b5b34e3	90359ef711ed53d9	9b446763c8627ea8
01000	6891c0f4842b823	c4d786a25727a7fc	094fe7da31fb60cd	9f9a004de5e70767
00100	04bdd77b7cb4e7b	c598cb710558af67	fc386d1f99d3a925	4928e0b43e781893
10100	544dd3eb8b2ad6c	94462af4f583d770	8061ba9381f51f55	476d4eded7c90fcc
11111	1fc25aa6a7e0b4c	93dd968c06a56809	9768e9d1ba74d3b6	014e9dc9f13670bb
Parity, hex <0:63>	**Parity, hex <64:127>**	**Parity, hex <128:139>**		
ed0e78f1734bc808	a38c0c417bd68f36	825		

91A.3 Output of the RS(544,514) encoder

Table 91A–3 contains a RS(544,514) codeword. Each row of Table 91A–1 is a set of four 66-bit blocks that is converted to one 257-bit block using the procedure defined in 91.5.2.5. The resulting set of 20 257-bit blocks constitute the message portion of the codeword. The parity is computed using the encoder defined in 91.5.2.7 and is appended to the message to complete the codeword.

Table 91A–3—RS(544,514) codeword

Header <0:4>	Payload, hex <5:64>	Payload, hex <65:128>	Payload, hex <129:192>	Payload, hex <193:256>
00101	a5a3bf86d9acf5c	de55cb85df0f7ca0	e6ccff8e8212b1c6	d63bc6c309000638
11110	7e3b0ce30e0497d	dc8df31ec3ab4491	66fb9139c81cd37b	b57477d4f05e3602
01111	8fd495012947a31	e7777cf0c6d06280	44529cf4b4900528	85ce1d27750ad61b
00110	46d5c71743f5c69	c1bf62e5dc5464b5	dc6011be7ea1ed54	1cf92c450042a75f
00100	c4b940eaf3140db	77bb612a7abf401f	c22d341e90545d98	ce6daf1f248bbd6d
10010	d22d0b3f9551ed6	574686c3f9e93898	2e52628f4a1282ce	f20c86d71944aab1
10001	5133c9333808a2c	1aa825d8b817db4d	637959989f3021eb	976806641b26aae9
00011	637d4531b7ed5f2	53c3e96d3b12fb46	528c7eb8481bc969	ab8f9980d5a54559
10100	94d2abfda65cc33	94fe646efe5af02d	9a65ae5fcd88c03a	5ef08673168def9b
00000	20c871a953fffc6	ce0bb95ac263e6c1	4f6a917d1a676571	5890918c7b687d75
01101	4d2b3e43096f836	84cdd4fc48b79608	b3e4503e3c824a8c	fd6d0b1a39687929
10011	130167c08302a69	4c15ff56de92b1ad	d0c2f0d4ff0dee95	e1422ee2e8b92125
00101	e5acaf86592fcee	de799be0b903c880	2714ffbf40bc09f6	c3be97c3c285009f
10010	120faf19f606631	93007cabbb3f8c9d	ef6955f7f43df5d0	4dbd0616afe60e1f
10001	31e49b7c7f7bb5d	901d828746ceec61	71ed3c097158c224	11adb3d81e13d263
00101	a50d1a343b2394b	eab30ca27b5b34e3	90359ef711ed53d9	9b446763c8627ea8
01000	6891c0f4842b823	c4d786a25727a7fc	094fe7da31fb60cd	9f9a004de5e70767
00100	04bdd77b7cb4e7b	c598cb710558af67	fc386d1f99d3a925	4928e0b43e781893
10100	544dd3eb8b2ad6c	94462af4f583d770	8061ba9381f51f55	476d4eded7c90fcc
11111	1fc25aa6a7e0b4c	93dd968c06a56809	9768e9d1ba74d3b6	014e9dc9f13670bb
Parity, hex <0:63>	**Parity, hex <64:127>**	**Parity, hex <128:191>**	**Parity, hex <192:255>**	**Parity, hex <256:299>**
d6983839edc3e5ac	c3cb45691ddba6cb	c26d756ea6f5b73d	249e30f415aa60b1	5743dc81c21

91A.4 Reed-Solomon encoder model

This annex also includes a model of the Reed-Solomon encoder, defined in 91.5.2.7, written in the C programming language. To emulate the RS(528,514) encoder, declare global variables per 91A.4.1. To emulate the RS(544,514) encoder, declare global variables per 91A.4.2. The generic components of the model are defined in 91A.4.3 to 91A.4.6.

91A.4.1 Global variable declarations for RS(528,514)

These global variables define the codeword size (in symbols) and generator polynomial coefficients (see Table 91–1) for the RS(528,514) code. Elements of GF(2^{10}) are presented as decimal values.

```
long n_symbols = 528;
unsigned long generator_polynomial[1024] =
   {904,6,701,32,656,925,900,614,391,592,265,945,290,432};
```

91A.4.2 Global variable declarations for RS(544,514)

These global variables define the codeword size and generator polynomial coefficients for the RS(544,514) code.

```
long n_symbols = 544;
unsigned long generator_polynomial[1024] =
   {575,552,187,230,552,1,108,565,282,249,593,132,94,720,495,385,942,503,883,36
   1,788,610,193,392,127,185,158,128,834,523};
```

91A.4.3 Other global variable declarations

The following global variables are declared for both RS(528,514) and RS(544,514). The field polynomial is assigned its decimal representation (1033 corresponds to $x^{10}+x^3+1$).

```
long polynomial = 1033;
long k_symbols = 514;
long check_symbols;
unsigned long codeword[1024];
```

91A.4.4 GF(2^{10}) multiplier function

This function implements multiplication over GF(2^{10}) using the expansion and reduction algorithm.

```
unsigned long multiply(long aa, long bb)
{
   unsigned long expand = 0;
   long k;

   for (k = 0; k < 10; k++)
   {
      if (bb & (1 << k))
         expand = expand ^ (aa << k);
   }

   for (k = 0; k < 9; k++)
   {
      if ((expand >> (18-k)) & 1)
         expand = expand ^ (polynomial << (8-k));
   }

   return expand;
}
```

91A.4.5 Reed-Solomon encoder function

This function implements the Reed-Solomon encoder. It uses the multiply() function.

```
void encode()
{
   long k, j;
   unsigned long multiplier;
   unsigned long generator_vector[1024];
   unsigned long encoder_divide[1024];

   for (k = 0; k < check_symbols; k++)
      encoder_divide[k] = 0;

   for (k = 0; k < k_symbols; k++)
   {
      multiplier = codeword[k] ^ encoder_divide[0];

      for (j = 0; j < check_symbols; j++)
         generator_vector[j] = multiply(multiplier, generator_polynomial[j]);

      for (j = 0; j < check_symbols-1; j++)
         encoder_divide[j] = generator_vector[j] ^ encoder_divide[j+1];

      encoder_divide[check_symbols-1] = generator_vector[check_symbols-1];

      for (j = 0; j < check_symbols; j++)
         codeword[j+k_symbols] = encoder_divide[j];
   }
}
```

91A.4.6 Main function

This sample main function defines a hypothetical message consisting of a countdown from 1023 to 510 (514 Reed-Solomon symbols). It then computes the parity and produces a codeword using the encode() function. The resulting codeword is printed to the console.

```
void main()
{
   long k;
   check_symbols = n_symbols-k_symbols;

   //*** Generate simple message symbols ***
   for (k = 0; k < k_symbols; k++)
      codeword[k] = 1023-k;

   encode();

   for (k = 0; k < n_symbols; k++)
      printf("%ld ", codeword[k]);
}
```

Annex 92A

(informative)

100GBASE-CR4 TP0 and TP5 test point parameters and channel characteristics

92A.1 Overview

Annex 92A provides information on parameters associated with test points TP0 and TP5 that may not be testable in an implemented system. TP0 and TP5 test points are illustrated in the 100GBASE-CR4 link block diagram of Figure 92–2. It also provides information on channel characteristics.

92A.2 Transmitter characteristics at TP0

The transmitter characteristics at TP0 are constrained at TP0a by 93.8.1.

92A.3 Receiver characteristics at TP5

The receiver characteristics at TP5 are constrained at TP5a by 93.8.2.

92A.4 Transmitter and receiver differential printed circuit board trace loss

The recommended maximum insertion loss allocation for the transmitter or receiver differential controlled impedance printed circuit boards is determined using Equation (92A–1) and illustrated in Figure 92A–1. Note that the recommended maximum insertion loss allocation for the transmitter or receiver differential controlled impedance printed circuit boards is 6.81 dB at 12.8906 GHz. The recommended maximum insertion loss allocation for the transmitter or receiver differential controlled impedance printed circuit boards is consistent with the insertion loss TP0 to TP2 or TP3 to TP5 given in 92.8.3.6 and an assumed mated connector loss of 1.69 dB.

$$IL_{\mathrm{PCB}}(f) \le IL_{\mathrm{PCBmax}}(f) = 0.5(0.0694 + 0.4248\sqrt{f} + 0.9322f) \quad \text{(dB)} \qquad \text{(92A–1)}$$

for 0.01 GHz $\le f \le$ 19 GHz.

where

f	is the frequency in GHz
$IL_{\mathrm{PCB}}(f)$	is the insertion loss for the transmitter and receiver PCB
$IL_{\mathrm{PCBmax}}(f)$	is the recommended maximum insertion loss for the transmitter and receiver PCB

The minimum insertion loss allocation for the transmitter and receiver differential controlled impedance printed circuit boards for each differential lane (i.e., the minimum value of the sum of the insertion losses from TP0 to MDI receptacle or TP5 to MDI receptacle) is determined using Equation (92A–2) and illustrated in Figure 92A–1.

$$IL_{PCB}(f) \geq IL_{PCBmin}(f) = 0.086(0.0694 + 0.4248\sqrt{f} + 0.9322f) \quad \text{(dB)} \qquad \text{(92A–2)}$$

for 0.01 GHz ≤ f ≤ 19 GHz.

where

f	is the frequency in GHz
$IL_{PCB}(f)$	is the insertion loss for the transmitter and receiver PCB
$IL_{PCBmin}(f)$	is the minimum insertion loss for the transmitter and receiver PCB

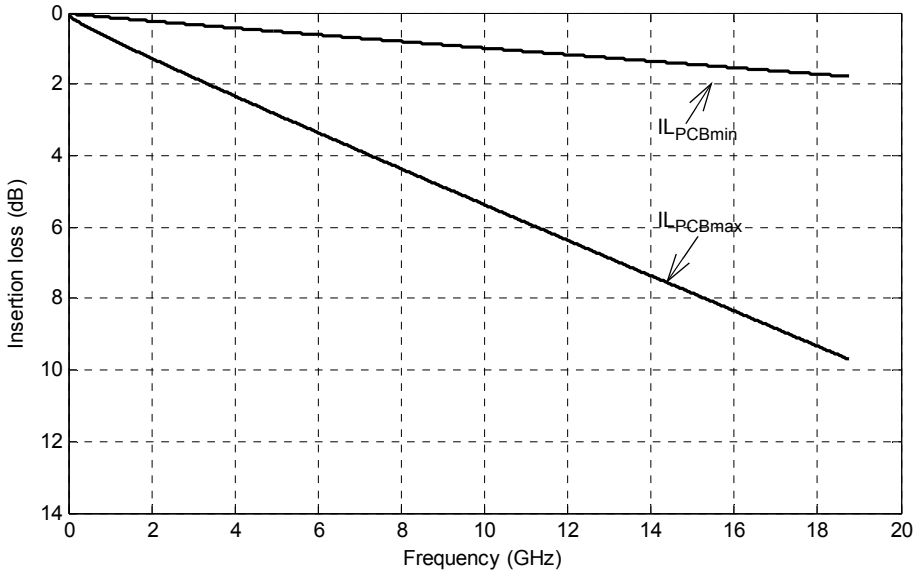

Figure 92A–1—Insertion Loss Tx or Rx PCB max and min

92A.5 Channel insertion loss

This subclause provides information on channel insertion losses for intended topologies ranging from 0.5 m to 5 m in length. The maximum channel insertion loss associated with the 5 m topology is determined using Equation (92A–3). The channel insertion loss associated with the 0.5 m topology and a maximum host channel is determined by Equation (92A–5). The channel insertion loss budget at 12.8906 GHz for the 5 m topology is illustrated in Figure 92A–2.

The maximum channel insertion loss for the 5 m topology is determined using Equation (92A–3). The maximum channel insertion loss is 35 dB at 12.8906 GHz.

$$IL_{Chmax35dB}(f) = IL_{Camax5m}(f) + 2IL_{Host}(f) - 2IL_{MatedTF}(f) \quad \text{(dB)} \tag{92A–3}$$

for 0.05 GHz ≤ f ≤ 19 GHz.

where

f	is the frequency in GHz
$IL_{Chmax35dB}(f)$	is the maximum channel insertion loss between TP0 and TP5 representative of a 5 m cable assembly and a maximum host channel
$IL_{Camax5m}(f)$	is the maximum 5 m cable assembly insertion loss.
$IL_{Host}(f)$	is the maximum insertion loss from TP0 to TP2 or TP3 to TP5 using Equation (92–8)
$IL_{MatedTF}(f)$	is the nominal insertion loss of the mated test fixture using Equation (92A–4)

The nominal insertion loss of the mated test fixture is determined using Equation (92A–4).

$$IL_{MatedTF}(f) = 0.1148\sqrt{f} + 0.287f \quad \text{(dB)} \tag{92A–4}$$

for 0.01 GHz ≤ f ≤ 25 GHz.

where

f	is the frequency in GHz
$IL_{MatedTF}(f)$	is the nominal insertion loss of the mated test fixture.

The channel insertion loss between TP0 and TP5 representative of a 0.5 m cable assembly and a maximum host channel is determined using Equation (92A–5).

$$IL_{Ch0.5m}(f) = IL_{Camin0.5m}(f) + 2IL_{Host}(f) - 2IL_{MatedTF}(f) \quad \text{(dB)} \tag{92A–5}$$

for 0.05 GHz ≤ f ≤ 19 GHz.

where

f	is the frequency in GHz
$IL_{Ch0.5m}(f)$	is the channel insertion loss between TP0 and TP5 representative of a 0.5 m cable assembly and a maximum host channel
$IL_{Camin0.5m}(f)$	is the minimum 0.5 m cable assembly insertion loss given in Equation (92–26) and illustrated in Figure 92–12.
$IL_{Host}(f)$	is the maximum insertion loss from TP0 to TP2 or TP3 to TP5 using Equation (92–8)
$IL_{MatedTF}(f)$	is the nominal insertion loss of the mated test fixture using Equation (92A–4)

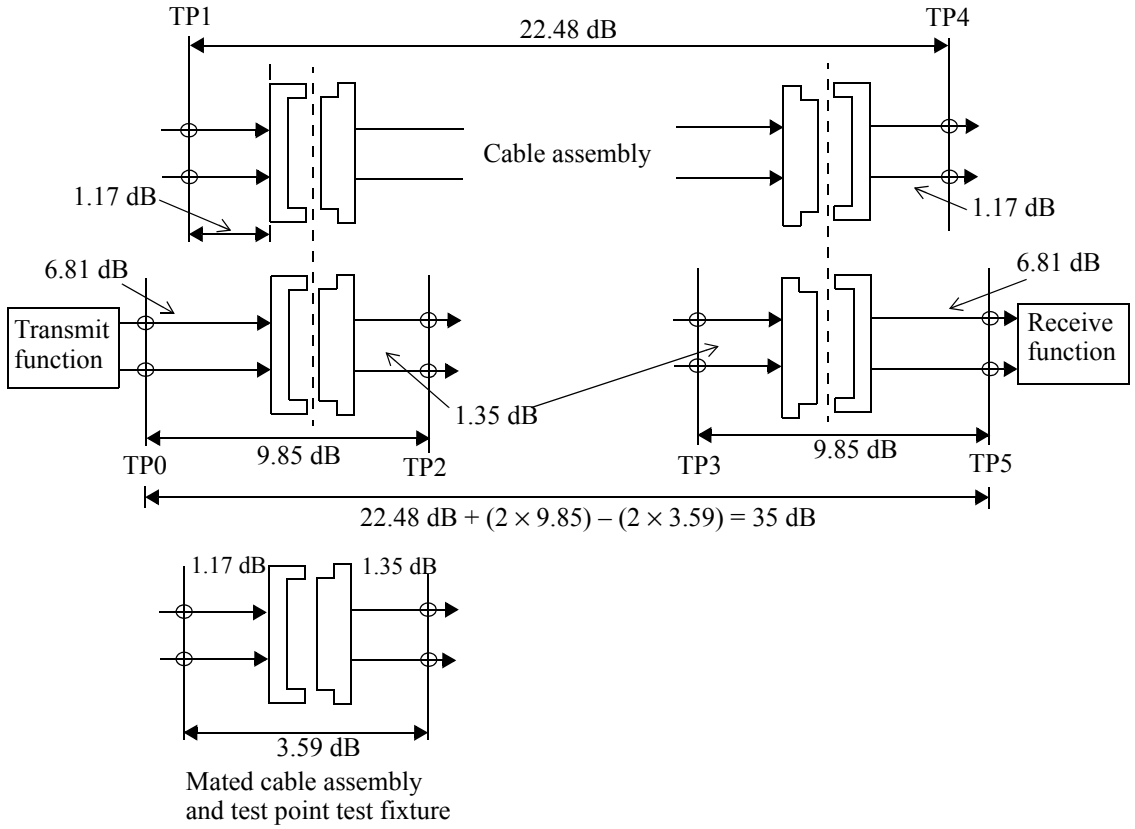

NOTE—The connector insertion loss is 1.07 dB for the mated test fixture. The host connector is allocated 0.62 dB of additional margin.

Figure 92A–2—35 dB channel insertion loss budget at 12.8906 GHz

92A.6 Channel return loss

The return loss of each lane of the 100GBASE-CR4 channel is recommended to meet the values determined using Equation (92–27).

92A.7 Channel Operating Margin (COM)

The Channel Operating Margin (COM) for the channel between TP0 and TP5, computed using the procedure in 93A.1 and the parameters in Table 93–8, is recommended to be greater than or equal to 3 dB.

NOTE—For cable lengths greater than 4 m, a frequency step (Δf) no larger than 5 MHz is recommended.

Annex 93A

(normative)

Specification methods for electrical channels

93A.1 Channel Operating Margin

The Channel Operating Margin (COM) is a figure of merit for a channel derived from a measurement of its scattering parameters. COM is related to the ratio of a calculated signal amplitude to a calculated noise amplitude as defined by Equation (93A–1).

$$COM = 20\log_{10}(A_s/A_{ni}) \quad \text{(93A–1)}$$

COM shall be calculated using the method described in this annex. The signal amplitude A_s is defined in 93A.1.6 and the noise amplitude A_{ni} is defined in 93A.1.7.

Figure 93A–1 illustrates the reference model that is the basis for the calculation for COM. The parameters used to calculate COM are listed in Table 93A–1. The values assigned to these parameters are defined by the Physical Layer specification that invokes the method (see Table 93A–2).

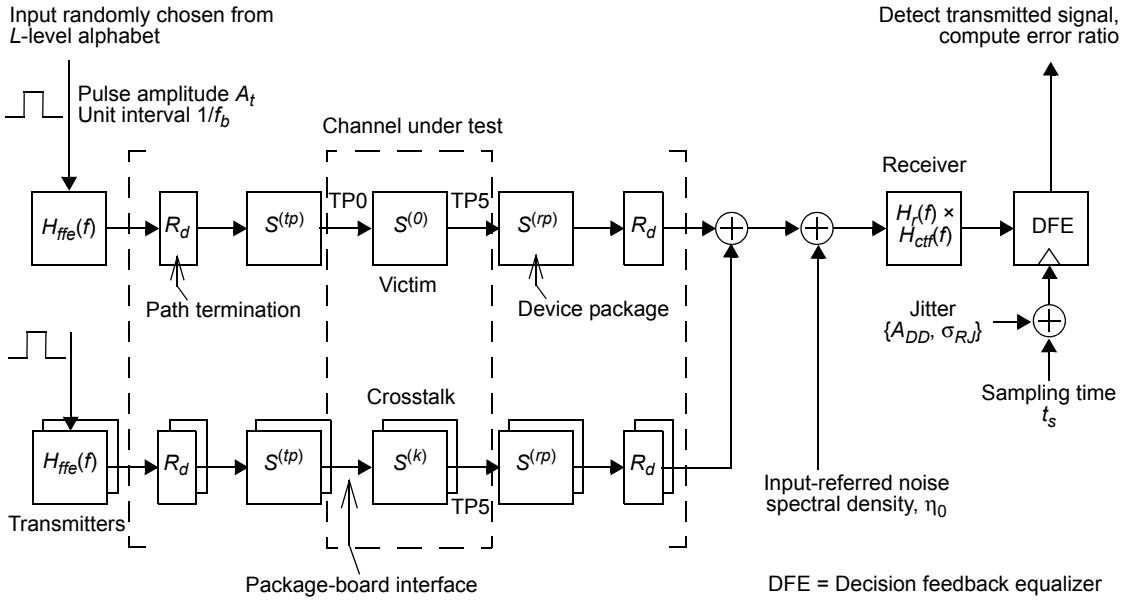

Figure 93A–1—COM reference model

Table 93A–1—COM parameters

Parameter	Reference	Symbol	Units
Signaling rate	93A.1.1	f_b	GBd
Maximum start frequency	93A.1.1	$f_{\min}$	GHz
Maximum frequency step	93A.1.1	Δf	GHz
Device package model Single-ended device capacitance Transmission line length Single-ended package capacitance at package-to-board interface	93A.1.2	 C_d z_p C_p	 nF mm nF
Single-ended reference resistance	93A.1.2	R_0	Ω
Single-ended termination resistance	93A.1.3	R_d	Ω
Receiver 3 dB bandwidth	93A.1.4.1	f_r	GHz
Transmitter equalizer, minimum cursor coefficient	93A.1.4.2	$c(0)$	—
Transmitter equalizer, pre-cursor coefficient Minimum value Maximum value Step size	93A.1.4.2	$c(-1)$	 — — —
Transmitter equalizer, post-cursor coefficient Minimum value Maximum value Step size	93A.1.4.2	$c(1)$	 — — —
Continuous time filter, DC gain Minimum value Maximum value Step size	93A.1.4.3	g_{DC}	 dB dB dB
Continuous time filter, zero frequency	93A.1.4.3	f_z	GHz
Continuous time filter, pole frequencies	93A.1.4.3	f_{p1} f_{p2}	GHz
Transmitter differential peak output voltage Victim Far-end aggressor Near-end aggressor	93A.1.5	 A_v A_{fe} A_{ne}	 V V V
Number of signal levels	93A.1.6	L	—
Level separation mismatch ratio	93A.1.6	R_{LM}	—
Transmitter signal-to-noise ratio	93A.1.6	SNR_{TX}	dB
Number of samples per unit interval	93A.1.6	M	—
Decision feedback equalizer (DFE) length	93A.1.6	N_b	UI
Normalized DFE coefficient magnitude limit	93A.1.6	$b_{\max}(n)$	—
Random jitter, RMS	93A.1.6	σ_{RJ}	UI
Dual-Dirac jitter, peak	93A.1.6	A_{DD}	UI
One-sided noise spectral density	93A.1.6	η_0	V^2/GHz
Target detector error ratio	93A.1.7	DER_0	—

Table 93A–2—Physical Layer specifications that employ COM

Physical Layer	Parameter values
100GBASE-CR4 (Clause 92)	Table 93–8
100GBASE-KR4 (Clause 93)	Table 93–8
100GBASE-KP4 (Clause 94)	Table 94–17
CAUI-4 (Annex 83D)	Table 83D–6

93A.1.1 Measurement of the channel

The channel consists of a victim signal path plus some number of far-end and near-end crosstalk paths. The total number of paths for a given channel is denoted as K and, by convention, the path index k=0 corresponds to the victim path. The number of crosstalk paths is a function of the structure of the system. All significant contributors to the channel crosstalk should be included in the calculation of COM.

Each signal path is represented by a set of frequency-dependent scattering parameters. For the purpose of the calculation of COM, references to scattering parameters correspond to the differential-mode scattering parameters. The scattering parameters measured at frequency f are presented as the 2 x 2 matrix $S(f)$ as defined by Equation (93A–2).

$$S(f) = \begin{bmatrix} s_{11}(f) & s_{12}(f) \\ s_{21}(f) & s_{22}(f) \end{bmatrix} \quad \text{(93A–2)}$$

The relationship between $S(f)$ and other commonly cited characteristics is as follows. The insertion loss is the magnitude in dB of either $1/s_{12}(f)$ or $1/s_{21}(f)$. The input and output return loss are the magnitude in dB of $1/s_{11}(f)$ and $1/s_{22}(f)$, respectively.

The scattering parameters for the victim signal path are measured from TP0 to TP5. The scattering parameters for each crosstalk path are measured from the package-to-board interface of the aggressor transmitter to TP5. The frequency-dependent scattering matrix for signal path k is denoted as $S^{(k)}(f)$. The reference impedance for scattering parameter measurements is 100 Ω.

It is recommended that the scattering parameters be measured with uniform frequency step no larger than Δf from a start frequency no larger than $f_{\min}$ to a stop frequency of at least the signaling rate f_b.

93A.1.2 Transmitter and receiver device package models

Each signal path in the channel is augmented to reflect the likely influence of transmitter and receiver device packages. The device package models are two-port networks defined by their scattering parameters. The scattering parameters are calculated using the method defined in 93A.1.2.1 through 93A.1.2.4.

Each signal path in the channel is represented by the scattering matrix $S^{(k)}$. The augmented signal path is denoted as $S_p^{(k)}$ and is defined by Equation (93A–3).

$$S_p^{(k)} = \text{cascade}(\text{cascade}(S^{(tp)}, S^{(k)}), S^{(rp)}) \quad \text{(93A–3)}$$

The function cascade() is defined in 93A.1.2.1. $S^{(tp)}$ and $S^{(rp)}$ are defined in 93A.1.2.4. If k corresponds to a near-end crosstalk path, $S^{(tp)}$ is calculated with the smallest value of z_p specified by the clause that invokes this method.

93A.1.2.1 Cascade connection of two-port networks

The connection of a pair of two-port networks x and y such that port 2 of network x is connected to port 1 of network y may be represented by an equivalent two-port network z. Port 1 of network z corresponds to port 1 of network x and port 2 network z corresponds to port 2 of network y. The scattering parameters of network z are given in terms of the scattering parameters of networks x and y by Equation (93A–4) through Equation (93A–7).

$$s_{11}^{(z)} = s_{11}^{(x)} + \frac{s_{21}^{(x)} s_{11}^{(y)} s_{12}^{(x)}}{1 - s_{22}^{(x)} s_{11}^{(y)}} \tag{93A–4}$$

$$s_{12}^{(z)} = \frac{s_{12}^{(x)} s_{12}^{(y)}}{1 - s_{22}^{(x)} s_{11}^{(y)}} \tag{93A–5}$$

$$s_{21}^{(z)} = \frac{s_{21}^{(x)} s_{21}^{(y)}}{1 - s_{22}^{(x)} s_{11}^{(y)}} \tag{93A–6}$$

$$s_{22}^{(z)} = s_{22}^{(y)} + \frac{s_{12}^{(y)} s_{22}^{(x)} s_{21}^{(y)}}{1 - s_{22}^{(x)} s_{11}^{(y)}} \tag{93A–7}$$

For the purpose of this annex, this set of operations is referred to using the shorthand notation $S^{(z)} = \text{cascade}(S^{(x)}, S^{(y)})$.

93A.1.2.2 Two-port network for a shunt capacitance

The scattering parameters for a shunt capacitance with value C are defined by Equation (93A–8) where $j = \sqrt{-1}$ and $\omega = 2\pi f$.

$$S(C) = \frac{1}{2 + j\omega C R_0} \begin{bmatrix} -j\omega C R_0 & 2 \\ 2 & -j\omega C R_0 \end{bmatrix} \tag{93A–8}$$

The scattering parameters for the device capacitance C_d are denoted as $S^{(d)} = S(C_d)$ and the scattering parameters for the board capacitance C_p are denoted as $S^{(p)} = S(C_p)$.

93A.1.2.3 Two-port network for the package transmission line

The scattering parameters for the package transmission line model are a function of the complex propagation coefficient defined by Equation (93A–9), Equation (93A–10), and Equation (93A–11) and the reflection coefficient defined by Equation (93A–12). The values of the parameters that appear in these equations are defined in Table 93A–3. The units of f are GHz.

$$\gamma(f) = \left\{ \begin{array}{ll} \gamma_0 & f = 0 \\ \gamma_0 + \gamma_1\sqrt{f} + \gamma_2(f)f & f > 0 \end{array} \right\} \quad \text{(93A–9)}$$

$$\gamma_1 = a_1(1 + j) \quad \text{(93A–10)}$$

$$\gamma_2(f) = a_2(1 - j(2/\pi)\log_e(f/1\ \text{GHz})) + j2\pi\tau \quad \text{(93A–11)}$$

$$\rho = \frac{Z_c - 2R_0}{Z_c + 2R_0} \quad \text{(93A–12)}$$

Table 93A–3—Transmission line model parameters and values

Parameter	Value	Units
γ_0	0	1/mm
a_1	1.734×10^{-3}	$ns^{1/2}$/mm
a_2	1.455×10^{-4}	ns/mm
τ	6.141×10^{-3}	ns/mm
Z_c	78.2	Ω

The scattering parameters for a package transmission line of length z_p are defined by Equation (93A–13) and Equation (93A–14). The units of z_p are mm.

$$s_{11}^{(l)}(f) = s_{22}^{(l)}(f) = \frac{\rho(1 - \exp(-\gamma(f)2z_p))}{1 - \rho^2\exp(-\gamma(f)2z_p)} \quad \text{(93A–13)}$$

$$s_{21}^{(l)}(f) = s_{12}^{(l)}(f) = \frac{(1 - \rho^2)\exp(-\gamma(f)z_p)}{1 - \rho^2\exp(-\gamma(f)2z_p)} \quad \text{(93A–14)}$$

The transmission line scattering parameter matrix is then denoted as $S^{(l)}$.

93A.1.2.4 Assembly of transmitter and receiver device package models

The scattering parameters for the transmitter device package model $S^{(tp)}$ are the result of the cascade connection of the device capacitance, package transmission line, and board capacitance as defined by Equation (93A–15).

$$S^{(tp)} = \text{cascade}(\text{cascade}(S^{(d)}, S^{(l)}), S^{(p)}) \quad \text{(93A–15)}$$

Similarly, the scattering parameters for the receiver device package model $S^{(rp)}$ are the result of the cascade connection of the board capacitance, package transmission line, and device capacitance as defined by Equation (93A–16).

$$S^{(rp)} = \text{cascade}(\text{cascade}(S^{(p)}, S^{(l)}), S^{(d)}) \quad (93A–16)$$

93A.1.3 Path terminations

The input to each signal path is terminated by an impedance defined by the reflection coefficient Γ_1. The output of each signal path is terminated by an impedance defined by the reflection coefficient Γ_2.

The reflection coefficients Γ_1 and Γ_2 are defined by Equation (93A–17).

$$\Gamma_1 = \Gamma_2 = \frac{R_d - R_0}{R_d + R_0} \quad (93A–17)$$

The voltage transfer function of the terminated signal path is defined by Equation (93A–18) where $\Delta S(f) = s_{11}(f)s_{22}(f) - s_{12}(f)s_{21}(f)$.

$$H_{21}(f) = \frac{s_{21}(f)(1 - \Gamma_1)(1 + \Gamma_2)}{1 - s_{11}(f)\Gamma_1 - s_{22}(f)\Gamma_2 + \Gamma_1\Gamma_2\Delta S(f)} \quad (93A–18)$$

The voltage transfer function for the signal path represented by $S_p^{(k)}(f)$ is denoted $H_{21}^{(k)}(f)$.

93A.1.4 Filters

The voltage transfer function for each signal path $H_{21}^{(k)}(f)$ (see 93A.1.3) is multiplied by a set of filter transfer functions to yield $H^{(k)}(f)$ as shown in Equation (93A–19).

$$H^{(k)}(f) = H_{ffe}(f)H_{21}^{(k)}(f)H_r(f)H_{ctf}(f) \quad (93A–19)$$

The receiver noise filter $H_r(f)$ is defined in 93A.1.4.1, the transmitter equalizer $H_{ffe}(f)$ is defined in 93A.1.4.2, and the receiver equalizer $H_{ctf}(f)$ is defined in 93A.1.4.3.

The filtered voltage transfer function $H^{(k)}(f)$ is used to compute the pulse response (see 93A.1.5).

93A.1.4.1 Receiver noise filter

$H_r(f)$ is a noise filter defined by Equation (93A–20).

$$H_r(f) = \frac{1}{1 - 3.414214(f/f_r)^2 + (f/f_r)^4 + j2.613126(f/f_r - (f/f_r)^3)} \quad (93A–20)$$

93A.1.4.2 Transmitter equalizer

$H_{ffe}(f)$ is defined by Equation (93A–21) and is intended to represent the transmitter equalizer. If k corresponds to a near-end crosstalk path, then $c(-1)$ and $c(1)$ are zero regardless of the values used for the other paths. The value of the "cursor" coefficient $c(0)$ is set to $1 - |c(-1)| - |c(1)|$ for any value of $c(-1)$ and $c(1)$. If the value of $c(0)$ is less than the specified minimum value, the corresponding combination of $c(-1)$ and $c(1)$ is considered invalid and is not used to calculate COM.

$$H_{ffe}(f) = \sum_{i=-1}^{1} c(i)\exp(-j2\pi(i+1)(f/f_b)) \quad (93A–21)$$

93A.1.4.3 Receiver equalizer

$H_{ctf}(f)$ is defined by Equation (93A–22).

$$H_{ctf}(f) = \frac{10^{g_{DC}/20} + jf/f_z}{(1 + jf/f_{p1})(1 + jf/f_{p2})} \quad (93A–22)$$

93A.1.5 Pulse response

The pulse response of a signal path is defined to be the output of the path following the application of a rectangular pulse one unit interval in duration at its input. First define the function $X(f)$ per Equation (93A–23) where $\mathrm{sinc}(x) = \sin(\pi x)/(\pi x)$ and $T_b = 1/f_b$ is the unit interval.

$$X(f) = A_t T_b \mathrm{sinc}(fT_b) \quad (93A–23)$$

$X(f)$ is a function of A_t, which in turn is based on the path index k. If k=0, i.e., the victim path, then $A_t = A_v$. If k corresponds to a far-end crosstalk path, then $A_t = A_{fe}$. If k corresponds to a near-end crosstalk path, then $A_t = A_{ne}$.

The pulse response $h^{(k)}(t)$ is derived from the voltage transfer function $H^{(k)}(f)$ (see 93A.1.4) using Equation (93A–24).

$$h^{(k)}(t) = \int_{-\infty}^{\infty} X(f) H^{(k)}(f) \exp(j2\pi ft) df \quad (93A–24)$$

NOTE 1—COM is expected to be computed from measurements at discrete frequencies that cover a limited span (see 93A.1.1). The inverse Fourier transform depicted in Equation (93A–24) is likely to be implemented as a discrete Fourier transform and the filtered voltage transfer function may need to be extrapolated (both to DC and to one half of the sampling frequency) for this computation. The extrapolation method and sampling frequency must be chosen carefully to limit the error in the COM computation.

NOTE 2—The time span of the pulse response in unit intervals, N, is limited in practice by frequency step Δf ($N = f_b/\Delta f$) but in general should be set to include all significant components of the pulse response.

93A.1.6 Determination of variable equalizer parameters

COM is a function of the variables $c(-1)$, $c(1)$, and g_{DC}. The following procedure is used to determine the values of these variables that are used to calculate COM.

a) Compute the pulse response $h^{(k)}(t)$ of each signal path k for a given $c(-1)$, $c(1)$, and g_{DC} using the procedure defined in 93A.1.5.
b) Define t_s to be the time that satisfies Equation (93A–25). If there are multiple values of t_s that satisfy the equation, then the first value prior to the peak of $h^{(0)}(t)$ is selected. The coefficients of the decision feedback equalizer $b(n)$ are computed as shown in Equation (93A–26). If N_b is 0, then the $b(n)$ is considered to be zero for all n.
c) Define A_s to be $R_{LM} h^{(0)}(t_s)/(L–1)$.
d) Compute σ_{TX}^2 per Equation (93A–30) and Equation (93A–29). This represents the noise output from the transmitter.
e) Compute $h_{ISI}(n)$ per Equation (93A–27). This represents the residual intersymbol interference (ISI) after decision feedback equalization. The corresponding ISI amplitude variance σ_{ISI}^2 is computed per Equation (93A–31) and Equation (93A–29).

f) Compute the slope of the pulse response of the victim path $h_J(n)$ as shown in Equation (93A–28). The variance of the amplitude error due to timing jitter σ_J^2 is computed per Equation (93A–32) and Equation (93A–29).

g) The variance of the amplitude for path k is given by Equation (93A–33) where the phase index m can assume any integer value from 0 to M–1. Denote the value of m that maximizes the variance for path k as i. The variance of the amplitude for the combination of all crosstalk paths σ_{XT}^2 is then computed using Equation (93A–34), which is the sum of the maximum variances for the individual paths k=1 to K–1.

h) Compute the variance of the noise at the output of the receive equalizer σ_N^2 based on the one-sided spectral density η_0 referred to the receiver noise filter input per Equation (93A–35).

i) Compute the figure of merit (FOM) per Equation (93A–36).

$$h^{(0)}(t_s - T_b) = h^{(0)}(t_s + T_b) - h^{(0)}(t_s)b(1) \tag{93A–25}$$

$$b(n) = \begin{Bmatrix} -b_{\max}(n) & h^{(0)}(t_s + nT_b)/h^{(0)}(t_s) < -b_{\max}(n) \\ b_{\max}(n) & h^{(0)}(t_s + nT_b)/h^{(0)}(t_s) > b_{\max}(n) \\ h^{(0)}(t_s + nT_b)/h^{(0)}(t_s) & \text{otherwise} \end{Bmatrix} \tag{93A–26}$$

$$h_{ISI}(n) = \begin{Bmatrix} 0 & n = 0 \\ h^{(0)}(t_s + nT_b) - h^{(0)}(t_s)b(n) & 1 \le n \le N_b \\ h^{(0)}(t_s + nT_b) & \text{otherwise} \end{Bmatrix} \tag{93A–27}$$

$$h_J(n) = \frac{h^{(0)}(t_s + (n + 1/M)T_b) - h^{(0)}(t_s + (n - 1/M)T_b)}{2/M} \tag{93A–28}$$

$$\sigma_X^2 = \frac{L^2 - 1}{3(L - 1)^2} \tag{93A–29}$$

$$\sigma_{TX}^2 = [h^{(0)}(t_s)]^2 10^{-SNR_{TX}/10} \tag{93A–30}$$

$$\sigma_{ISI}^2 = \sigma_X^2 \sum_n h_{ISI}^2(n) \tag{93A–31}$$

$$\sigma_J^2 = (A_{DD}^2 + \sigma_{RJ}^2)\sigma_X^2 \sum_n h_J^2(n) \tag{93A–32}$$

$$[\sigma_m^{(k)}]^2 = \sigma_X^2 \sum_n [h^{(k)}((m/M + n)T_b)]^2 \tag{93A–33}$$

$$\sigma_{XT}^2 = \sum_{k=1}^{K-1} [\sigma_i^{(k)}]^2 \tag{93A–34}$$

$$\sigma_N^2 = \eta_0 \int_0^{\infty} \left| H_r(f) H_{ctf}(f) \right|^2 df \tag{93A–35}$$

$$FOM = 10\log_{10}\left(\frac{A_s^2}{\sigma_{TX}^2 + \sigma_{ISI}^2 + \sigma_J^2 + \sigma_{XT}^2 + \sigma_N^2} \right) \tag{93A–36}$$

The FOM is calculated for each permitted combination of $c(-1)$, $c(1)$, and g_{DC} values per Table 93A–1. The combination of values that maximizes the FOM, including the corresponding value of t_s, is used for the calculation of the interference and noise amplitude in 93A.1.7 and the calculation of COM in 93A.1.

93A.1.7 Interference and noise amplitude

Given the values of $c(-1)$, $c(1)$, g_{DC}, and t_s derived in 93A.1.6, compute the combined interference and noise distribution $p(y)$ per 93A.1.7.3. The corresponding cumulative distribution function is $P(y)$ as defined by Equation (93A–37).

$$P(y) = \int_{-\infty}^{y} p(y)dy \tag{93A–37}$$

The noise amplitude, A_{ni}, is the magnitude of the value of y_0 that satisfies the relationship $P(y_0) = DER_0$ where DER_0 is the target detector error ratio. The detector error ratio is the probability that the detector fails to identify the signal level that was transmitted.

In 93A.1.7.1 through 93A.1.7.3, "*" denotes convolution, which is defined by Equation (93A–38).

$$f(t) * g(t) = \int_{-\infty}^{\infty} f(\tau)g(t-\tau)d\tau \tag{93A–38}$$

93A.1.7.1 Interference amplitude distribution

The interference amplitude distribution is computed from the sampled pulse response $h(n)$ with the assumption that the transmitted symbols are independent, identically distributed random variables and that the symbols are uniformly distributed across the set of L possible values. For the purpose of this subclause, $h(n)$ is a general notation that corresponds to $A_{DD}h_f(n)$ (see 93A.1.7.2), $h_{ISI}(n)$, or $h^{(k)}((i / M + n)T_b)$ (see 93A.1.7.3).

Equation (93A–39) defines the nth component of the interference amplitude distribution function where $\delta(y)$ is the Dirac delta function.

$$p_n(y) = \frac{1}{L}\sum_{l=0}^{L-1} \delta\left(y - \left(\frac{2l}{L-1} - 1 \right) h(n) \right) \tag{93A–39}$$

The set of N such components are combined via convolution to obtain the complete interference amplitude distribution. Initialize $p(y)$ to $\delta(y)$ and then evaluate Equation (93A–40) sequentially for n=0 to N–1.

$$p(y) = p(y) * p_n(y) \tag{93A–40}$$

NOTE 1—COM is expected to be numerically computed using a quantized amplitude axis y. The amplitude step Δy introduces quantization error in the calculated distribution function that is compounded by subsequent convolutions with other quantized distribution functions. It is recommended that Δy be no larger than 0.1% of A_s or 0.01 mV, whichever is smaller.

NOTE 2—It is recommended that components of the pulse response whose amplitude is less than 0.1% of A_s be ignored as they likely correspond to measurement noise or numerical artifacts.

93A.1.7.2 Noise amplitude distribution

The calculation of COM includes two noise terms that are described in terms of their distribution function. The first term has a Gaussian amplitude distribution function with zero mean and variance σ_G^2. The variance is defined by Equation (93A–41) where $H_r(f)$ is defined in 93A.1.4.1, $H_{ctf}(f)$ is defined in 93A.1.4.3, and σ_X^2 and $h_J(n)$ are defined in 93A.1.6.

$$\sigma_G^2 = \sigma_{TX}^2 + \sigma_{RJ}^2\sigma_X^2\sum_n h_J^2(n) + \eta_0\int_0^\infty \left|H_r(f)H_{ctf}(f)\right|^2 df \quad \text{(93A–41)}$$

The amplitude distribution of the Gaussian noise term is defined by Equation (93A–42).

$$p_G(y) = \frac{\exp(-y^2/(2\sigma_G^2))}{\sqrt{2\pi\sigma_G^2}} \quad \text{(93A–42)}$$

The second term is denoted as p_{DD} and is related to the amplitude noise resulting from dual-Dirac jitter. It is computed using the procedure defined in 93A.1.7.1 with $h(n) = A_{DD}h_J(n)$.

The components are combined using convolution to yield the overall noise amplitude distribution function as defined in Equation (93A–43).

$$p_n(y) = p_G(y) * p_{DD}(y) \quad \text{(93A–43)}$$

93A.1.7.3 Combination of interference and noise distributions

Compute the intersymbol interference amplitude distribution using the procedure defined in 93A.1.7.1 with $h(n) = h_{ISI}(n)$ as defined by Equation (93A–27) and denote the result as $p(y)$.

The contributions of the K–1 crosstalk paths to the total interference are included as follows. Determine the phase index $m = i$ that maximizes the variance of the amplitude for path k as defined by Equation (93A–33). Compute the interference amplitude distribution using the procedure defined in 93A.1.7.1 with $h(n) = h^{(k)}((i / M + n)T_b)$ and denote the result as $p^{(k)}(y)$.

Compute $p^{(k)}(y)$ and evaluate Equation (93A–44) sequentially for integer values k=1 to K–1.

$$p(y) = p(y) * p^{(k)}(y) \quad \text{(93A–44)}$$

The noise distribution $p_n(y)$ defined in 93A.1.7.2 is then included to yield the combined interference and noise amplitude distribution as shown in Equation (93A–45).

$$p(y) = p(y) * p_n(y) \quad \text{(93A–45)}$$

93A.2 Test channel calibration using COM

A generalized block diagram of the interference tolerance test channel is shown in Figure 93A–2.

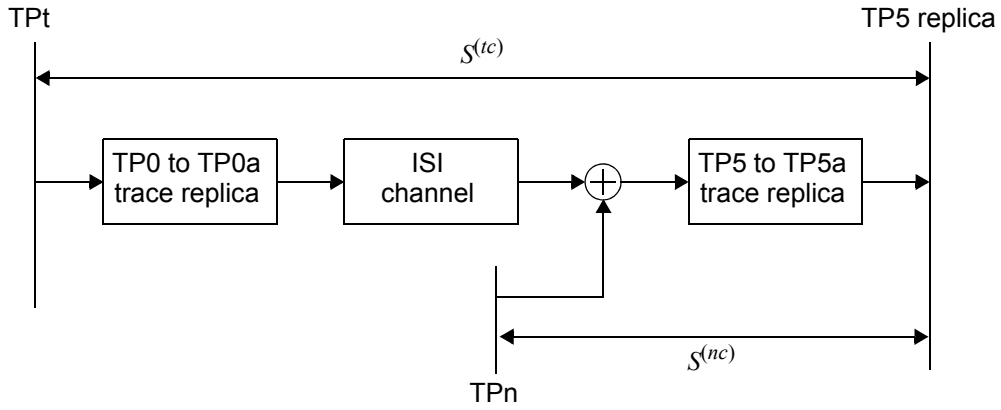

Figure 93A–2—Generalized interference tolerance test channel

The signal path from the test transmitter connected at TPt to the receiver under test connected at TP5 consists of replicas of the test fixture traces, a controlled ISI channel, and the means by which additive broadband noise is coupled into the path. This path is represented by the scattering parameters $S^{(tc)}$ measured from TPt to the TP5 replica. The signal path from the broadband noise source connected at TPn to the receiver consists of the means of broadband noise coupling and a replica of TP5 to TP5a test fixture trace. This path is represented by scattering parameters $S^{(nc)}$ measured from TPn to the TP5 replica.

COM is used to calibrate the interference tolerance test channel. The values assigned to the parameters listed in Table 93A–1 are defined by the Physical Layer specification that invokes this method.

The calculations defined in 93A.1 are evaluated for $S^{(tc)}$ with the following exceptions.

If the test transmitter presents a high-quality termination, e.g., it is a piece of test equipment, the transmitter device package model $S^{(tp)}$ is omitted from the calculation of $S_p^{(k)}$. Instead, the voltage transfer function is multiplied by the filter $H_t(f)$ defined by Equation (93A–46) where T_r is the 20 to 80% transition time (see 86A.5.3.3) of the signal as measured at TP0a.

$$H_t(f) = \exp(-(\pi f T_r / 1.6832)^2) \tag{93A–46}$$

The approximate voltage transfer function for the path from TPn to the output of the receiver equalizer is defined by Equation (93A–47).

$$H^{(ne)}(f) = s_{21}^{(nc)}(f)\, s_{21}^{(rp)}(f)\, H_r(f)\, H_{ctf}(f) \tag{93A–47}$$

The broadband noise source applies noise at TPn that has a Gaussian amplitude distribution with zero mean and standard deviation σ_{bn}. The power spectral density of the noise is flat from $-f_b / 2$ to $f_b / 2$ and is zero elsewhere. The standard deviation of the noise at the receiver equalizer output σ_{ne} is defined by Equation (93A–48).

$$\sigma_{ne}^2 = \frac{2\sigma_{bn}^2}{f_b} \int_0^{f_b/2} \left|H^{(ne)}(f)\right|^2 df \tag{93A–48}$$

Equation (93A–41) defines the standard deviation of the Gaussian noise amplitude distribution function. When COM is used to calibrate the interference tolerance test channel, this definition is replaced by Equation (93A–49). The value of σ_{bn} is adjusted until the target COM value is achieved.

$$\sigma_G^2 = \sigma_{TX}^2 + \sigma_{RJ}^2\sigma_X^2\sum_n h_J^2(n) + \eta_0\int_0^\infty \left|H_r(f)H_{ctf}(f)\right|^2 df + \sigma_{ne}^2 \tag{93A–49}$$

An additional figure of merit for the test channel is the root-sum-square of the magnitude terms n_1 to n_2 of the equalized pulse response where n_2 is less than or equal to N_b. This measure of the relative usage of the decision feedback equalizer is defined by Equation (93A–50).

$$u_b(n_1, n_2) = \frac{1}{A_s}\sqrt{\sum_{n=n_1}^{n_2}\left(h^{(0)}(n)\right)^2} \tag{93A–50}$$

The shorthand notation RSS_DFE4 is used to represent $u_b(4, N_b)$.

93A.3 Fitted insertion loss

The fitted insertion loss as a function of frequency is given by Equation (93A–51).

$$IL_{fitted}(f) = a_0 + a_1\sqrt{f} + a_2 f + a_4 f^2 \tag{93A–51}$$

Denote the insertion loss, in dB, measured at frequency f_n as $IL(f_n)$. Given the insertion loss measured at N uniformly-spaced frequencies from start frequency $f_{\min}$ to stop frequency $f_{\max}$ with step no larger than Δf, the coefficients for the fitted insertion loss shall be calculated as follows.

Define the weighted frequency matrix F using Equation (93A–52).

$$F = \begin{bmatrix} 10^{-IL(f_1)/20} & \sqrt{f_1}10^{-IL(f_1)/20} & f_1 10^{-IL(f_1)/20} & f_1^2 10^{-IL(f_1)/20} \\ 10^{-IL(f_2)/20} & \sqrt{f_2}10^{-IL(f_2)/20} & f_2 10^{-IL(f_2)/20} & f_2^2 10^{-IL(f_2)/20} \\ \cdots & \cdots & \cdots & \cdots \\ 10^{-IL(f_N)/20} & \sqrt{f_N}10^{-IL(f_N)/20} & f_N 10^{-IL(f_N)/20} & f_N^2 10^{-IL(f_N)/20} \end{bmatrix} \tag{93A–52}$$

Define the weighted insertion loss vector L using Equation (93A–53).

$$L = \begin{bmatrix} IL(f_1)10^{-IL(f_1)/20} \\ IL(f_2)10^{-IL(f_2)/20} \\ \cdots \\ IL(f_N)10^{-IL(f_N)/20} \end{bmatrix} \tag{93A–53}$$

The fitted insertion loss coefficients are then given by Equation (93A–54).

$$\begin{bmatrix} a_0 \\ a_1 \\ a_2 \\ a_4 \end{bmatrix} = (F^T F)^{-1} F^T L \quad \text{(93A–54)}$$

The values assigned to f_{min}, f_{max}, and Δf are defined by the Physical Layer specification that invokes this method.

93A.4 Insertion loss deviation

The insertion loss deviation $ILD(f)$ is the difference between the measured insertion loss $IL(f)$ and the fitted insertion loss $IL_{fitted}(f)$ (see 93A.3) as shown in Equation (93A–55).

$$ILD(f) = IL(f) - IL_{fitted}(f) \quad \text{(93A–55)}$$

A figure of merit for a channel that is based on $ILD(f)$ is given by Equation (93A–56). In Equation (93A–56), f_n are the frequencies considered in the computation of the fitted insertion loss and $W(f_n)$ is the weight at each frequency as defined by Equation (93A–57).

$$FOM_{ILD} = \left[\frac{1}{N} \sum_n W(f_n) ILD^2(f_n) \right]^{1/2} \quad \text{(93A–56)}$$

$$W(f_n) = \text{sinc}^2(f_n / f_b) \left[\frac{1}{1 + (f_n / f_t)^4} \right] \left[\frac{1}{1 + (f_n / f_r)^8} \right] \quad \text{(93A–57)}$$

The variable f_b is the signaling rate. The 3 dB transmit filter bandwidth f_t is inversely proportional to the 20% to 80% rise and fall time T_t. The constant of proportionality is 0.2365 (e.g., $T_t f_t = 0.2365$; with f_t in Hertz and T_t in seconds). The variable f_r is the 3 dB reference receiver bandwidth.

The values assigned to f_b, T_t, and f_r are defined by the Physical Layer specification that invokes this method.

Annex 93B

(informative)

Electrical backplane reference model

This annex describes additional informative test points that may used to partition the electrical backplane channel. See Figure 93B–1 and Table 93B–1

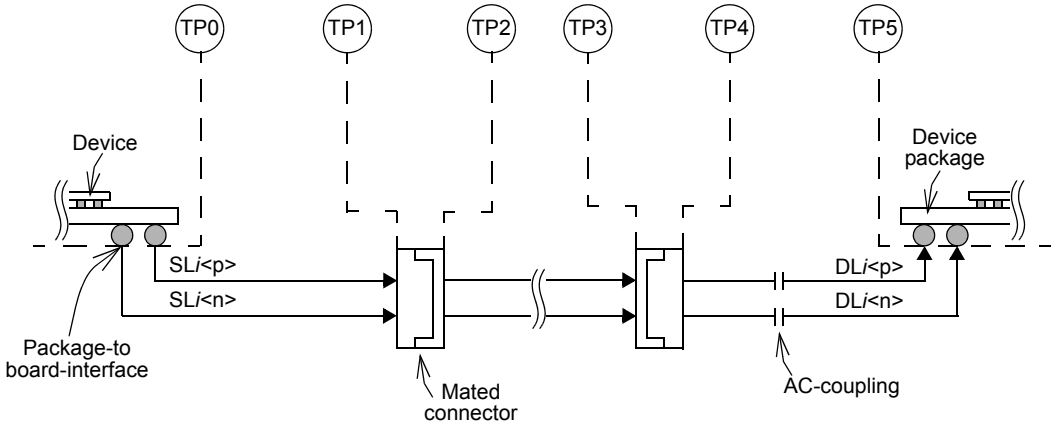

Figure 93B–1—Reference model (one direction from one lane is illustrated)

Table 93B–1—Description of channel components

Test points	Description
TP0 to TP1	The printed circuit board between the transmitter and the separable connector closest to the transmitter. TP1 is defined to be the interface between the board and connector plug.
TP2 to TP3	The electrical path from the separable connector closest to the transmitter to the separable connector closest to the receiver. TP2 and TP3 are defined to be the interface between connector receptacle and the printed circuit board.
TP4 to TP5	The printed circuit board between the receiver and the separable connector closest to the receiver. TP4 is defined to be the interface between the board and connector plug. It is recommended that the AC-coupling capacitors are implemented between TP4 and TP5
TP0 to TP5	The electrical backplane channel as defined in 93.9 and 94.4. TP0 and TP5 are defined to be the interface between the device package and the printed circuit board.

Annex 93C

(normative)

Receiver interference tolerance

This annex defines a test setup (see 93C.1) and method (see 93C.2) for testing receiver interference tolerance. The PMD clause that invokes this method specifies the following items:

a) Constraint limit values for peak-to-peak voltage, the pre-cursor peaking ratio, and the post-cursor peaking ratio for test setup,
b) Lower frequency bound for the noise spectral density constraints (f_{NSD1}),
c) Jitter parameters to be measured in test method step 3,
d) Target Channel Operating Margin (COM) and RSS_DFE4 values for the test system in test method step 7,
e) COM parameter table in test method step 7,
f) Jitter transformation method in test method step 7,
g) Test pattern in test method step 9, and
h) Test system frequency response.

NOTE—The intent of the interference tolerance test is to ensure that the PHY receiver operates correctly with transmitter parameters anywhere within the specified limits including the case where all parameters are at the specified limits. Testing of the receiver with transmitter parameters beyond the specified limits may be helpful to determine margin or to provide comparative metrics, but failure of the receiver to operate correctly under these conditions is not to be interpreted as non-compliance.

93C.1 Test setup

The interference tolerance test is performed with the setup shown in Figure 93C–2 or its equivalent. Calibration and characterization of the various elements in the test setup is accomplished using the test configurations in Figure 93C–3, Figure 93C–4, Figure 93C–5, and Figure 93C–6.

The transmitter is functionally and parametrically compliant to the requirements of the invoking PMD clause. The ISI channel emulates the frequency dependent loss of a backplane channel. The channel noise source emulates crosstalk, transmitter noise, and unequalizable signal distortions introduced by a channel.

The transmitter output, as measured at TP0a, meets all transmitter specifications as indicated by the invoking PMD clause. In addition, the transmitter output, as measured at TP0a, is constrained such that for any transmitter equalizer setting the maximum differential peak-to-peak voltage, the pre-cursor peaking ratio, and the post-cursor peaking ratio are constrained as indicated by the PMD clause that invokes this method.

The channel noise source has an adjustable output such that the level may be set according to the test procedure. The noise produced by the channel noise source is measured directly at the output of the noise source (see Figure 93C–6). The noise is Gaussian with a crest factor of at least 5. The noise spectral density, $NSD(f)$, is normalized and constrained according to the relations in Equation (93C–1), where f_b is the symbol rate and f_{NSD1} is specified by the PMD clause that invokes this method. $NSD(f)$ is in units of V^2/Hz. The average noise spectral density, NSD_{average}, is determined according Equation (93C–2). An example constraint template with f_{NSD1} equal to $0.08f_b$ is illustrated in Figure 93C–1.

$$\left.\begin{aligned} 10\log_{10}\left(\frac{NSD(f)}{NSD_{average}}\right) &< 3 \\ 10\log_{10}\left(\frac{NSD(f)}{NSD_{average}}\right) &> -3(1 - 1.2f/f_b) \end{aligned}\right\} f_{NSD1} \le f \le f_b/2 \qquad \text{(93C–1)}$$

$$NSD_{average} = \frac{\int_{f_{NSD1}}^{f_b/2} NSD(f)df}{f_b/2 - f_{NSD1}} \qquad \text{(93C–2)}$$

The receiver on one lane at a time is tested for compliance. The input to the receiver on each of the other lanes is generated by a transmitter with similar levels and equalization settings and transmitted through a similar channel, such that the input signals are similar to the input signal on the lane under test.

NOTE—FEXT and NEXT in the test setup are not accounted for in the test channel calibration (see 93A.2). It is recommended that the test setup be designed to minimize these effects.

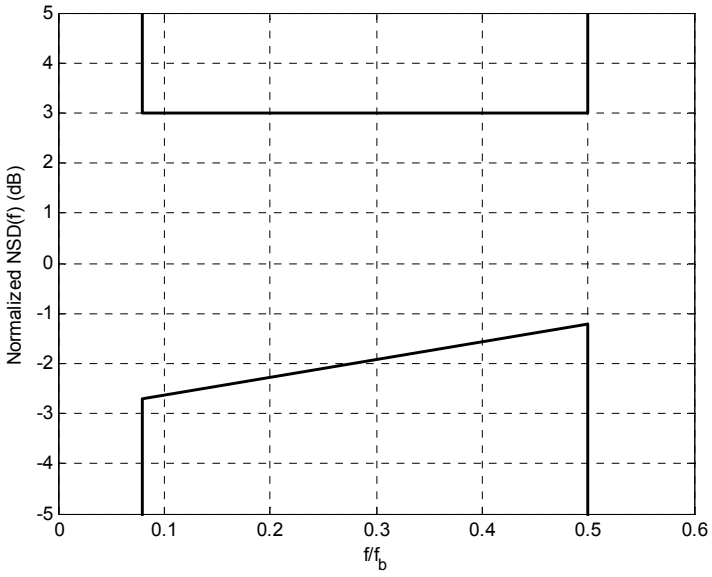

Figure 93C–1—Example NSD(f) constraint template

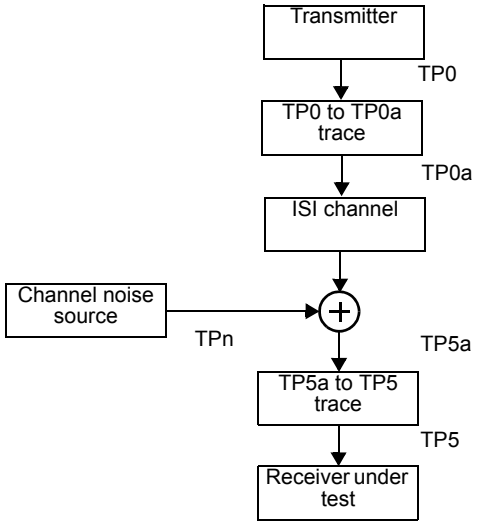

Figure 93C–2—Interference tolerance test setup

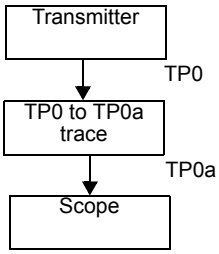

Figure 93C–3—Interference tolerance transmitter test setup

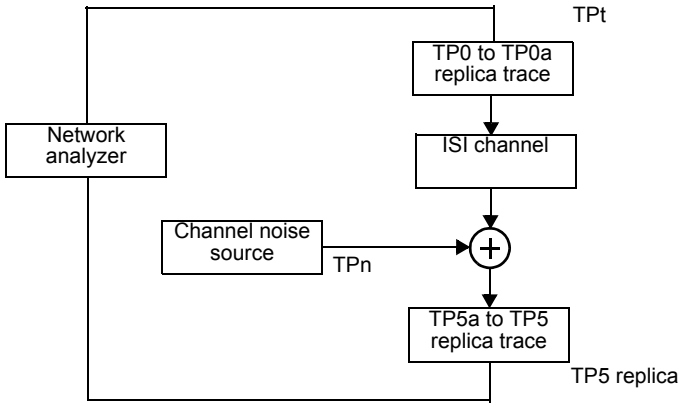

Figure 93C–4—Interference tolerance channel s-parameter test setup

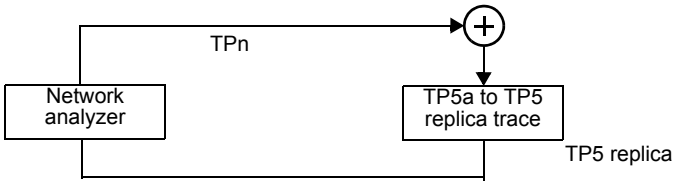

Figure 93C–5—Interference tolerance channel noise path test setup

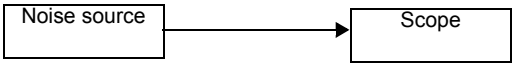

Figure 93C–6—Interference tolerance channel noise level test setup

93C.2 Test method

The interference tolerance test is performed using the following method:

1) Set the channel noise source to zero.
2) Using the test setup in Figure 93C–2, initiate the training sequence, allow the training sequence to complete, and retain the resulting transmitter tap coefficients.
3) Measure the jitter parameters relevant to the PMD clause that invokes this method that are to be used to set the value of σ_{RJ} and A_{DD} in step 7.
4) Measure the noise parameters relevant to the PMD clause that invokes this method that are to be used to set the value SNR_{TX}.
5) Using the test setup in Figure 93C–4 (also see Figure 93A–2), measure the scattering parameters, $S^{(tc)}$, of the test channel (TPt to TP5 replica).
6) Using the test setup in Figure 93C–5 (also see Figure 93A–2), measure the scattering parameters, $S^{(nc)}$, of the noise addition network (TPn to TP5 replica).
7) Using the procedure defined in 93A.2: (a) determine the receiver noise level, σ_{bn}, required to achieve the COM value specified in the PMD clause that invokes this method, and (b) verify that RSS_DFE4 is greater than or equal to the value specified in the PMD clause that invokes this method. The procedure is based on the calculation of COM, which uses the parameters defined in the COM parameter table in the PMD clause that invokes this method with the following exceptions. The value of σ_{RJ} and A_{DD} are set based on a transformation of measured parameters as specified in the PMD clause that invokes this method. The value of SNR_{TX} is set based on a transformation of the measured parameters specified in the PMD clause that invokes this method. In the COM computation the transmitter package model is included only if a compliant transmitter with a similar termination is used. If a transmitter with high quality termination is used, in the COM calculation, the termination is modeled as ideal and a Gaussian low pass filter is added to Equation (93A–19), which has the same 20% to 80% transition time as the transmitter measured at TP0a.
8) Using the test setup in Figure 93C–6, measure the channel noise voltage σ_{bnm} and adjust it so that it equals σ_{bn} determined in step 7. The channel noise voltage is determined from the measured $NSD(f)$ according to Equation (93C–3).
9) Using the test setup in Figure 93C–2, the transmitter taps as determined in step 2, and the channel noise as determined in step 7, configure the transmitter to transmit the test pattern specified in the PMD clause that invokes this method. Also configure the transmitters of the PMD under test to transmit the same test pattern, with their transmitters in the preset condition.

10) Measure the FEC symbol error ratio on the receiver under test using the errored symbol counter, FEC_symbol_error_counter_*i*, where *i* is the lane number of the receiver under test.

$$\sigma_{bnm} = \sqrt{\int_0^{f_b/2} NSD(f)df} \qquad (93C–3)$$

A test system with frequency response specified in the PMD clause that invokes this method is to be used for measurement of the signal applied by the pattern generator and for measurements of the broadband noise.

Consensus

WE BUILD IT.

Connect with us on:

Facebook: https://www.facebook.com/ieeesa

Twitter: @ieeesa

LinkedIn: http://www.linkedin.com/groups/IEEESA-Official-IEEE-Standards-Association-1791118

IEEE-SA Standards Insight blog: http://standardsinsight.com

YouTube: IEEE-SA Channel

IEEE
standards.ieee.org
Phone: +1 732 981 0060 Fax: +1 732 562 1571