IEEE P802.3dj Editorial Team Report

Matt Brown (Alphawave, 802.3dj Chief Editor)
Gary Nicholl (Cisco)
Tom Huber (Nokia)
Adee Ran (Cisco)

IEEE P802.3dj Task Force

Introduction

❖ Some aspects of D1.1 implementation summarized:
  ➢ updates relating to new error ratio specifications
  ➢ updates adding link training and precoding to PAM4 optical PHYs
  ➢ revised C2M specification/test methodology
  ➢ Updates to Clause 176 SM-PMA symbol mux
  ➢ state of Clause 186, 800GBASE-ER1 PCS

Error ratios (data reliability)

IEEE P802.3dj Task Force

Error ratio introduction

❖ The resolution to Draft 1.0 comment #205 adopted new error ratio requirements.

❖ Comment response implemented in Draft 1.1 required some interpretation by the editorial team.

❖ Following slides point out results of the implementation in:
  ➢ Annex 174A – defining error ratio allocations to various portions of a network path
  ➢ Clause 118/171 – defining error ratio requirements for xMII extenders
  ➢ Clause/Annex 178/179/176D/176E – defining error ratio requirements for KR/CR/C2C/C2M
  ➢ Clause 180/181/182/183 – defining error ratio requirements for PAM4 optical
  ➢ Clause 185/187 – defining error ratio requirements for coherent optical

Scope and introduction

Data reliability ~~considerations~~

> New terminology as used in adopted proposal. More accurately, this clause defines an error ratio budget or allowance allocated to different portions of a network path.

174A.1 Scope

This clause defines the data reliability for Physical Layer implementations that include one or more physically instantiated inter-sublayer links (ISLs) with data rate of 200 Gb/s or higher per lane.

> Limits the scope to new PHY types introduced in this 802.3dj amendment.

174A.2 Introduction

Ethernet Physical Layers operating at a data rate of 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s are specified to enable a frame loss ratio (see 1.4.344) lower than 6.2×10^{-11} for 64-octet MAC frames with minimum interpacket gap. This frame loss ratio is equivalent to the frame loss ratio that would result from uncorrelated random bit errors on the MAC/PLS service interface with a bit error ratio (BER) of 10^{-13}. This defines the minimum expected data reliability for these physical layers.

> MAC frame loss ratio (FLR) has become the ultimate goal for an Ethernet network connection.

> Ultimate effective goal, but use of BER is obsolete due to bursty nature of errors at the detector and output of FEC decoder.

The use of forward error correction (FEC) allows higher BER at intermediate points within the physical layer, such as at the PMD, FEC, and PMA service interfaces. The acceptable BER specified at each interface depends on the data encoding as well as the distribution of errors over time.

This annex defines the error requirements at each interface that enables the expected minimum data reliability of the Physical Layer.

Network path

174A.3 Data reliability for an Ethernet network path

Data reliability of the path between two Ethernet DTEs is defined in terms of the frame loss ratio from the PLS service interface of the Reconciliation sublayer (RS) of the transmitting DTE to the PLS service interface of the RS of the receiving DTE.

The frame loss ratio for 64-octet MAC frames with minimum interpacket gap is expected to be less than 6.2×10^{-11}.

NOTE—The frame loss ratio is affected by multiple components along the network path and is not a normative requirement of a specific component.

The path from the MAC at one end to the MAC at the other end.

This is an "expectation", not a normative statement, as it is achieved by combination of Physical Layer implementations at each end and the channel between. The expectation should be achieved if all three elements are compliant to the specified requirements.

xMII Extender

174A.4 Data reliability for an xMII Extender

This subclause defines the data reliability for a 200GMII, 400GMII, 800GMII, or 1.6TMII Extender (see Clause 118 and Clause 171).

Data reliability of the xMII Extender is defined in terms of the frame loss ratio between the transmitting XS (DTE or PHY) and the receiving XS (PHY or DTE, respectively).

The frame loss ratio for 64-octet frames with minimum interpacket gap is expected to be less than 10^{-12}. This requirement is equivalent to a FEC codeword error ratio (see 174A.7) lower than 2.4×10^{-13}.

NOTE—The frame loss ratio is affected by multiple components within the xMII Extender and is not a normative requirement of a specific component.

174A.7 FEC codeword error ratio

FEC codeword error ratio is a ratio of the number of FEC codewords which are deemed to be uncorrectable (NE) by the FEC decoder to the number total number of FEC codewords processed by the FEC decoder (NT). The value of NT should be sufficiently large to reliably verify that block error ratio requirements are met, either by direct measurement or statistical projection.

An xMII Extender is bookended by an XS (equivalent to a PCS) with FEC at each end.

Again, frame loss ratio value is <u>expected</u> performance since attributed to set of multiple elements.

This is first time we specify the FEC codeword error ratio within the 802.3 standard. This new parameter is defined in 176A.7.

PHY-to-PHY link

174A.5 Data reliability for a PHY-to-PHY link

Data reliability of a PHY-to-PHY link is defined in terms of the frame loss ratio between the service interfaces of the transmitting PCS and the receiving PCS.

The frame loss ratio for 64-octet MAC frames with minimum interpacket gap is expected to be less than 6×10^{-11}.

For PHYs using the 200GBASE-R, 400GBASE-R, 800GBASE-R, or 1.6TBASE-R PCS, the frame loss ratio requirement is equivalent to an FEC codeword error ratio (see 174A.7) of less than 1.45×10^{-11}.

For PHYs using the 800GBASE-ER1 PCS, the frame loss ratio requirement is equivalent to a FEC codeword error ratio (see 174A.7) of less than TBD.

NOTE—The frame loss ratio is affected by multiple components within the PHYs and by the media, and is not a normative requirement of a specific component.

This link is bookended by a PCS (with FEC) at each end.

Two PCS types are recognized in 802.3dj.

Again, frame loss ratio value is **expected** performance since attributed to set of multiple elements.

Again, this is first time we specify the FEC codeword error ratio within the 802.3 standard.

Note that we have a TBD here for the ER1 PCS.

Inter-sublayer link (part 1)

174A.6 Data reliability for physically instantiated 200 Gb/s per lane inter-sublayer links

This subclause defines the data reliability for a physically instantiated inter-sublayer link with 200 Gb/s per lane signaling between a pair of PMAs including:

— a PMD at each end and a medium between

— an Inner FEC and PMD at each end and a medium between

— an xAUI-n between

176A.2 Conventions

For the purpose of this annex, the following definitions apply. Clause 1.4 should be referenced for terms not defined in this annex.

Inter-sublayer link (ISL)
> A physically instantiated link between a pair of adjacent sublayers. The inter-sublayer link may be an xAUI-n between a pair of PMA sublayers within the same physical layer implementation or a pair of PMDs and the medium between.

Although the ISL is defined literally as being between a pair of PMDs, the error measurement isn't meaningful until processed by the PMA and (in some cases) the Inner FEC. Thus the error ratio is measured at the PMA, rather than at the PMD.

Inter-sublayer link (ISL) is defined in 176A.2.

9

Inter-sublayer link (part 2)

The data reliability is measured on each lane using the following method:

a) At the transmitting device generate a PRBS31Q in the PMA (above the PMD).

b) At the output of the receiving PMA PAM4 decoder, insert random bit errors with a BER of BER_{added}.

c) Identify errored bits, either from the physical link or inserted per the previous step, using a pattern checker.

d) Divide the stream into a series of 10-bit symbols.

e) Within each block of 544 symbols count the number of errored symbols, where errored symbols are symbols with one or more bit errors.

f) Count the number of blocks with 16 or more errored symbols (NE).

g) Count the total number of blocks analyzed (NT). The value of NT should be sufficiently large to reliably verify that block error ratio requirements are met, either by direct measurement or statistical projection.

h) Calculate the block error ratio as NE / NT.

This procedure may alternately be achieved by using a PCS transmit function to generate a scrambled idle pattern in place of the PRBS31Q pattern and using the PCS receive function error counters to determine the FEC codeword error ratio (see 174A.7) in place of the block error ratio.

The measured block error ratio is expected be less than 1.45×10^{-11}.

BER_{added} represents the total random BER allocated to other physically instantiated inter-sublayer links in the PHY.

Adopted proposal defines this test using either PRBS31Q or scrambled idle, but hard to define together. Handled scrambled idle separately below.

Note that this could be done as part of the PMA error checking infrastructure or as a post-processing function in test equipment.

Alternate method using scrambled idle defined here.

BER_{added} is specified in the clause/annex that references this subclause.

Extender clauses

Defined simply by reference to 174A.4.

171.1a Data reliability

An 800GMII/1.6TMII Extender is expected to meet the frame loss ratio specifications in 174A.4.

Forgot to update Clause 118 for 200G/400G.

PMD (electrical) clauses and AUI annexes

178.2 Data reliability

A complete PHY is expected to meet the frame loss ratio specifications in 174A.5.

A PMD is expected to meet the block error ratio specifications in 174A.6 with BER_{added} equal to TBD.

178.9.3.3 Receiver interference tolerance

Receiver interference tolerance is defined by the procedure in Annex 93C. The receiver on each lane shall meet the expected block error ratio specified in 178.2 with channels matching the Channel Operating Margin (COM) and loss parameters for Test 1 and Test 2 in Table 178–10. The following additional considerations apply to the interference tolerance test.

Table 178–10—Receiver interference tolerance parameters

Parameter	Test 1 (low loss)			Test 2 (high loss)			Units
	Min	Max	Target	Min	Max	Target	
Block error ratio[a]			$< 1.45 \times 10^{-11}$				—
Insertion loss, *ILdd*, at 53.125 GHz[b]			—			—	dB
Receiver package class A	14.5	15.5		33.5	34.5		
Receiver package class B	14.5	15.5		30	31		
COM including effects of broadband noise[c]	—	—	3	—	—	3	dB

[a] The block error ratio (see 178.2) is measured instead of the FEC symbol error ratio in step 10) of the receiver interference tolerance method defined in 93C.2.

[b] ILdd measured between TPt and TP5 (see Figure 93C–4) minus ILdd of the specific package used by the test transmitter.

[c] The COM value is the target for the receiver noise level calibration defined in 93C.2 step 7). The channel noise voltage applied in 93C.2 step 8) should be as close as practical to the value needed to produce the target COM. If higher amplitude values are used, this would demonstrate margin to the specification but this is not required for compliance.

Defined by reference to 174A.5/6.
For KR/CR, it was not clear what allocation to give for the C2C AUIs that might be within this PHY type.

Changed symbol error ratio target to block error ratio targets.

rce

12

PMD (PAM4 optical) clauses

180.2 Data reliability

A complete PHY is expected to meet the frame loss ratio specifications in 174A.5.

A PMD is expected to meet the block error ratio specifications in 174A.6 with BER_{added} equal to 4×10^{-5}.

Defined by reference to 174A.5/6.

Table 180–8—200GBASE-DR1, 400GBASE-DR2, 800GBASE-DR4, and 1.6TBASE-DR8 receive characteristics

for 0.9 dB $\leq$ TECQ $\leq$ SECQ	-4.3 + TECQ	
Stressed receiver sensitivity (OMA$_{outer}$), each lane[c] (max)	-0.9	dBm
Conditions of stressed receiver sensitivity test[d]:		
Stressed eye closure for PAM4 (SECQ), lane under test	3.4	dB
OMA$_{outer}$ of each aggressor lane[e]	2.9	dBm

[a] The receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having this average power level. The receiver does not have to operate correctly at this input power.
[b] Average receive power, each lane (min) is 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 180.8) for the block error ratio specified in 180.2.
[d] These test conditions are for measuring stressed receiver sensitivity. They are not characteristics of the receiver.
[e] No aggressors needed for 200GBASE-DR1 in a single lane device.

Changed bit error ratio target to block error ratio target.

PMD (coherent optical) clauses

185.2 Data reliability

A complete PHY is expected to meet the frame loss ratio specifications in 174A.5.

A PMD is expected to meet <metric TBD>.

Editor's note: The block error ratio metric adopted per Draft 1.0 comment #205 and defined in 174A.5 does not address this PHY type. Contributions to address this are encouraged.

The adopted proposal did not address this 800 Gb/s PHY type.
Proposals for the error ratio metric encouraged.

187.2 Data reliability

A PHY is expected to meet the frame loss ratio specifications in 174A.4.

Defined only by frame loss ratio (and indirectly the FEC codeword error ratio).

Optical link training

IEEE P802.3dj Task Force

Optical link training

❖ At the May 2024 interim meeting we adopted a baseline for optical link training along with precoding for all 200 Gb/s per lane PAM4 PHY types in 802.3dj.

❖ The result of this is a set of updates to the following:
  ➢ Annex 176A – defining alternate ILT training frame structure and state diagram set
  ➢ Clause 120/177 – addition of precoding for these optical PHY types
  ➢ Clause 180/181/182/183 – updates to include ILT

Annex 176A (introduction)

Inter-sublayer link training for electrical and optical interfaces

Adding support for optical interfaces

176A.1 Overview

Inter-sublayer Link Training (ILT) for electrical and optical interfaces performs the path start-up protocol. ILT facilitates timing recovery and equalization in single inter-sublayer link (ISL) or multi-ISL paths (see 176A.2) while providing a mechanism through which the receiver may configure the transmitter to optimize performance. ILT supports these functions through the continuous exchange of fixed-length training frames between peers in an ISL and the transport of path end-to-end indications.

> *Editor's note: Two flavors of link training are defined.*
> *Type A1 addresses electrical interfaces including KR/CR PMDs, AUI-C2M, and AUI-C2C*
> *Type A2 addresses PAM4 optical interfaces including DR/DR-2/FR-500/FR PMDs.*

Define two types of training. A2 is for PAM4 optical

176A.4.2 Control and status fields

Two formats are defined for the control and status fields A1 and A2. Each interface using ILT shall identify which format is relevant for it.

> *Editor's note (to be removed in the next draft): Optical link training was adopted without specific nomenclature or guidance on how to distinguish it from the existing flavor. The terms "A1" and "A2" are placeholders chosen by the editor.*

rce

Annex 176A (training frame fields)

Table 176A–3—Control field structure for A2 interfaces

Bit(s)	Name	Description
15:10	Reserved	Transmit as 0, ignore on receipt
9:8	Modulation and precoding request	9 8 1 1 = PAM4 with precoding 1 0 = PAM4 0 1 = Reserved 0 0 = PAM2
7	Reserved	Transmit as 0, ignore on receipt
6:5	Test pattern request	6 5 1 1 = free -running PRBS31 1 0 = Reserved 0 1 = free-running PRBS13 0 0 = PRBS13
4:0	Reserved	Transmit as 0, ignore on receipt

Table 176A–5—Status field structure for A2 interfaces

Bit(s)	Name	Description
15	Receiver ready	1 = Training is complete and the receiver is ready for data 0 = Request for training to continue
14	ILT	Transmit as 1
13:12	Test pattern status	13 12 1 1 = free -running PRBS31 1 0 = Reserved 0 1 = free-running PRBS13 0 0 = PRBS13
11:10	Modulation and precoding status	11 10 1 1 = PAM4 with precoding 1 0 = PAM4 0 1 = Reserved 0 0 = PAM2
9	Receiver frame lock	1 = Frame boundaries identified 0 = Frame boundaries not identified
8	Reserved	Transmit as 0, ignore on receipt
7	Parity	Even parity bit
6	Extend training	1 = Continue training 0 = Switch to data when training is completed
5:0	Reserved	Transmit as 0, ignore on receipt

Annex 176A (state diagrams)

176A.11.3.5 State diagram figures

The training control state diagram (Figure 176A–7) defines the operation of ILT for AUIs and PMDs

The Training control state diagram is supported by the training frame lock, RTS update and coefficient update state diagrams. The training frame lock state diagram (Figure 176A–8) determines when ILT has detected training frame boundaries. The RTS update state diagram (Figure 176A–6) sets the value of the local_rts variable. The coefficient update state diagram (Figure 176A–9) defines the process for updating the local transmitter equalizer coefficients in response to requests from the path partner.

For A1 interfaces the interface control, frame lock and coefficient update state diagrams shall be implemented for each lane.

For A2 interfaces the interface control and frame lock state diagrams shall be implemented for each lane.

Uses a subset of the the electrical (A1) state diagrams.

NOTE — The frame lock state diagram is identical to the same state diagram in Figure 136–8.

Clause 180/181/182/183 PMD clauses (PHY tables)

Table 180–1—Physical Layer clauses associated with the 200GBASE-DR1 PMD

Associated clause	200GBASE-DR1
117—200 Gb/s RS	Required
117—200GMII[a]	Optional
118—200GMII Extender	Optional
119—200GBASE-R PCS	Required
120—200GBASE-R BM-PMA	Conditional[b]
120B—200GAUI-8 C2C	Optional[c]
120C—200GAUI-8 C2M	Optional[c]
120D—200GAUI-4 C2C	Optional[c]
120E—200GAUI-4 C2M	Optional[c]
120F—200GAUI-2 C2C	Optional[c]
120G—200GAUI-2 C2M	Optional[c]
176—200GBASE-R SM-PMA	Required[b]
176A—ILT	Required
176D—200GAUI-1 C2C	Optional[c]
176E—200GAUI-1 C2M	Optional[c]

Inter-sublayer link training added to PMD sublayer tables.

[a] The 200GMII is an optional interface. However, if the 200GMII is not implemented, a conforming implementation behaves functionally as though the RS and 200GMII were present.
[b] If one or two 200GAUI-n is implemented in a PHY, additional 200GBASE-R BM-PMA or SM-PMA sublayers are required according to the guidelines in 176B.4.1.
[c] One or two 200GAUI-n may be instantiated within a 200GBASE-DR1 PHY as described in 176B.4.1.

Clause 180/181/182/183 PMD clauses (service interface)

180.3 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the PMD. The service interface for these PMDs are described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data between the PMD and the PMD client. 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 116.3 for 200GBASE-R and 400GBASE-R, in 169.3 for 800GBASE-R, and 174.3 for 1.6TBASE-R. For PMD:IS_UNITDATA_i.request in the transmit direction and PMD:IS_UNITDATA_i.indication in the receive direction, where $i = 0$ to n–1, the tx_symbol and rx_symbol parameters are parallel PAM4 symbol streams with a nominal signaling rate of 106.25 GBd. The number of parallel streams, n, is 1 for 200GBASE-DR1, 2 for 400GBASE-DR2, 4 for 800GBASE-DR4, and 8 for 1.6TBASE-DR8.

The SIGNAL_OK parameter of the PMD:IS_SIGNAL.indication primitive corresponds to the variable training_status of the inter-sublayer training function, as defined in 176A.11.2.1. When SIGNAL_OK is either IN_PROGRESS or FAIL, the rx_symbol parameters of PMD:IS_UNITDATA_i.indication on all lanes are unspecified.

Editor's note: PMD:IS_SIGNAL.request(SIGNAL_OK) (transmit direction) as required by the ILT function needs to be defined. A proposal in this regard is encouraged.

Updated SIGNAL_OK (indication, RX direction) definition in line with the KR/CR clauses.

SIGNAL_OK (request, TX direction) needs to be defined.

Clause 180/181/182/183 PMD clauses (service interface)

180.5.4 PMD global signal detect function

The variable Global_PMD_signal_detect is a global indicator of the presence of optical signals on all n lanes. For the 200GBASE-DR1 PMD, the variable Global_PMD_signal_detect is set to PMD_signal_detect_0 (see 180.5.5). For the 400GBASE-DR2, 800GBASE-DR4, and 1.6TBASE-DR8

PMDs, the variable Global_PMD_signal_detect is set to the logical-and of PMD_signal_detect_i from all of the PMD lanes.

The state of the Global_PMD_signal_detect variable is conveyed to PMD client sublayers via the PMD service interface (see 180.3).

Editor's note: A means to control this variable in response to ILT is required. Proposals in this regard are required.

Global signal detect variable definition needs to be updated to reflect the training state, rather than per-lane signal detect states. (see 179.8.4)

179.8.4 PMD global signal detect function

The PMD global signal detect function is used by the PMD to indicate the successful completion of the start-up protocol by the PMD's inter-sublayer training (ILT) function (see 179.8.9). The variable Global_PMD_signal_detect is set to the value of remote_rts in the ILT function (see 176A.11.2.1).

Clause 180/181/182/183 PMD clauses (call out ILT)

180.5.12 Inter-sublayer link training function

A PMD shall provide the inter-sublayer link training (ILT) function for a Type A2 interface, specified in Annex 176A. When the variable mr_training_enable is true, the ILT function is used to request changes to the link partner transmitter state (modulation, training pattern, and precoder state), indicate the receiver state, and coordinate transition to DATA mode.

ILT called out here.

Clause 180/181/182/183 PMD clauses (MDI lane assignments)

180.6 Lane assignments

In order to support the ILT function as described in 180.5.12, optical lane assignments (within a group of transmit or receive lanes) for 200GBASE-DR1, 400GBASE-DR2, 800GBASE-DR4, or 1.6TBASE-DR8 are specified. The positioning and ordering of transmit and receive lanes at the MDI is specified in 180.8.3.1.
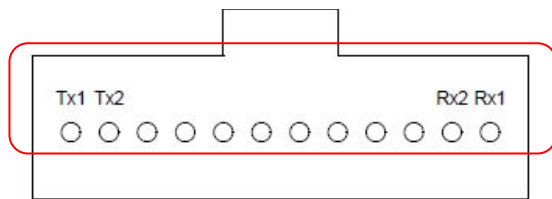
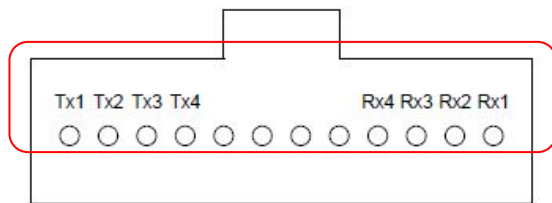Figure 180–7—400GBASE-DR2 optical lane assignments

MDI lanes are now explicitly numbered so that a TX lane can be associated with a corresponding RX lane.

Figure 180–8—800GBASE-DR4 optical lane assignments

Clause 120/173/176 PMA (precoding)

176.7.1.2 Precoding

The precoding specifications in this subclause apply to the input and output lanes of a PMA that are connected to the service interface of an xBASE-KRn, xBASE-CRn, xBASE-DRn, or xBASE-FRn-500 PMD, or are part of an xAUI-n C2C or C2M link.

The PMA shall provide $1/(1+D)$ mod 4 precoding capability on each output lane and may optionally provide $1/(1+D)$ mod 4 decoding capability on each input lane. Precoding is implemented as specified in 135.5.7.2.

The precoder is enabled independently on the Tx output, Rx input, Rx output, and Tx input on each lane. Precoding is enabled and disabled using variables precoder_tx_out_enable_i, precoder_rx_in_enable_i, precoder_rx_out_enable_i, and precoder_tx_in_enable_i, where i is in the range 0 to $(n-1)$.

If the PMA is connected to the service interface of an xBASE-KRn, xBASE-CRn, xBASE-DRn, or xBASE-FRn-500 PMD, or is part of an xAUI-n C2C or C2M link, and training is enabled by the management variable mr_training_enable (see 176A.12), then precoder_tx_out_enable_i and precoder_rx_in_enable_i shall be set as determined by the inter-sublayer link training (ILT) function in the LINK_READY state on lane i (see Figure 176A–7). The method by which the ILT function affects these variables is implementation dependent.

If the PMA is connected to the service interface of an xBASE-KRn, xBASE-CRn, xBASE-DRn, or xBASE-FRn-500 PMD, or is part of an xAUI-n C2C or C2M link, and ILT is disabled by the management variable mr_training_enable, then precoder_tx_out_enable_i, precoder_rx_in_enable_i, precoder_rx_out_enable_i, and precoder_tx_in_enable_i are set as required by the implementation.

Precoding is defined for DRn and FRn-500 and may be selected using ILT.

Clause 177 xBASE-R Inner FEC (precoding)

177.4.7.2 Precoding

The Inner FEC shall provide $1/(1+D) \bmod 4$ precoding capability on each transmit lane. Precoding is implemented as specified for output lanes in 135.5.7.2.

Precoding is enabled and disabled using variables precoder_tx_out_enable_i (where i is in the range 0 to 7). If training is enabled by the management variable mr_training_enable (see 176A.12), then precoder_tx_out_enable_i shall be set as determined by the inter-sublayer link training (ILT) function in the LINK_READY state on lane i (see Figure 176A–7). The method by which the ILT function affects this variable is implementation dependent. If ILT is disabled by the management variable mr_training_enable, precoder_tx_out_enable_i is set as required by the implementation.

177.5.1.1 Inverse precoding

The Inner FEC may optionally provide inverse $1/(1+D) \bmod 4$ precoding capability on each receive lane.

If inverse precoding is implemented, it is enabled or disabled as determined by the implementation.

If inverse precoding is enabled, the Inner FEC receive function processes the detected data equivalent to the process specified for input lanes in 135.5.7.2. If ILT function is enabled by the management variable mr_training_enable (see 176A.12), the precoding state on the link partner transmitter is requested using the ILT function. If ILT is disabled by the management variable mr_training_enable, the precoding state on the link partner transmitter is set by management.

Precoding is defined within the Inner FEC clause, thus implicitly for DRn-2, FR4, and LR4.

Precoding defined to be set using ILT.

C2M specification/test methodology

C2M input and output specifications

Following the acceptance of comment #186, the electrical characteristics and methodology subclauses (176E.4 and 176E.6) have been rewritten in D1.1.

The specification methodology is based on that of clause 179 with appropriate changes for the module side.

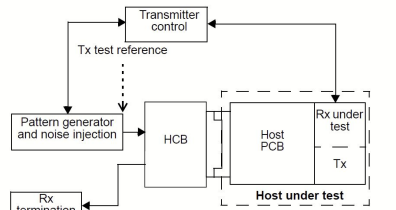The input test setup diagrams are shown below. For more details see D1.1.
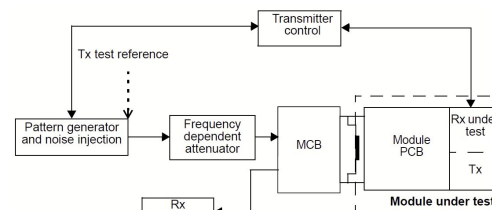
Figure 176E–7a—Host interference tolerance test setup

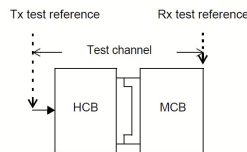**Figure 176E–8a—Module interference tolerance test setup**
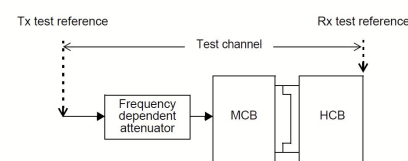
Figure 176E–7b—Host test channel calibration

Figure 176E–8b—Module test channel calibration

Clause 176 - Symbol-Mux PMA Restructuring

Introduction

100G lanes* > based on bit-muxing > use bit-muxing PMA (BM-PMA) > Cl 120, 173

200G lanes > based on RS-FEC symbol muxing > use symbol-muxing PMA (SM-PMA) > Cl 176

Converting between 100G lanes and 200G lanes:
- ❏ Use a BM-PMA to demux from 100G lanes to PCS lanes
- ❏ Use a SM-PMA to mux from PCS lanes to 200G lanes

176.1.4 Summary of PMA functions

An SM-PMA does not support bit multiplexing interfaces. To convert between bit multiplexing and symbol multiplexing interfaces, a BM-PMA and a SM-PMA are placed back-to-back with each PMA demultiplexing to PCSLs in between. For example, to convert from a 200GAUI-2 interface to a 200GAUI-1 interface, a 200GBASE-R 2:8 BM-PMA (see Clause 120) is placed next to a 200GAUI-1 8:1 SM-PMA.

40
41
42
43
44

* exception: 1.6TbE uses symbol-mutiplexing on 100G lanes (e.g. 1.6TAUI-16)

Three groupings (types) of symbol-multiplexing PMAs

m:n PMA — 1

    This term refers collectively to the 200GBASE-R 8:1 PMA, 400GBASE-R 16:2 PMA, — 2
    800GBASE-R 32:4 PMA, and 1.6TBASE-R 16:8 PMA. — 3

    — 4

n:m PMA — 5

    This term refers collectively to the 200GBASE-R 1:8 PMA, 400GBASE-R 2:16 PMA, — 6
    800GBASE-R 4:32 PMA, and 1.6TBASE-R 8:16 PMA. — 7

n:n PMA — 8

    This term refers collectively to the 200GBASE-R 1:1 PMA, 400GBASE-R 2:2 PMA, — 9
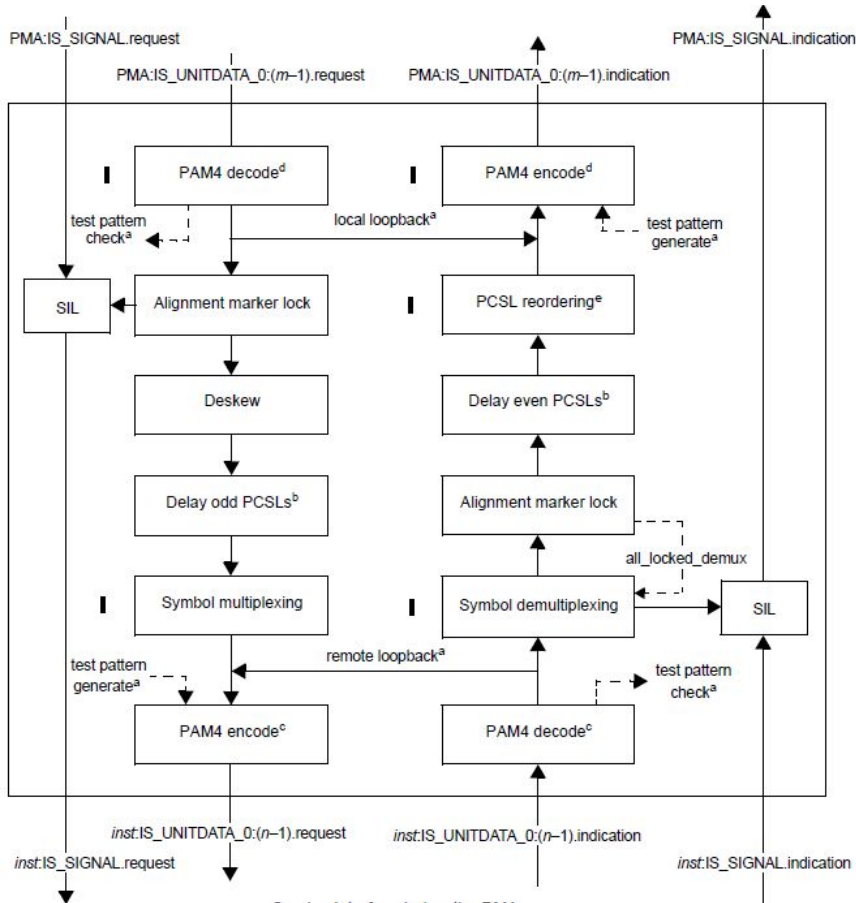   800GBASE-R 4:4 PMA, 1.6TBASE-R 8:8 PMA, and 1.6TBASE-R 16:16 PMA. — 10

— 11

PMA key truths :
- $m > n$
- m = number of PCS lanes
- n = number of 200G lanes (with one exception: 1.6T n:n with n=16)

Each type is defined in a single subclause with one block diagram and any differences explained.

m:n PMAs Functional Block Diagram

One block diagram applies to all 200G/400G/800G/1.6T m:n PMAs

Each step is described either as a common function or a function with common aspects, but some differences.

In the transmit flow:
- PAM4 decode/encode is the same for all Ethernet rates
- Deskew is different for 200G/400G vs. 800G vs.1.6T, but occurs at same place in the TX flow.
- Symbol mux occurs at same position in TX flow, but has slightly different rules for Ethernet rate
- Irrespective of Ethernet rate, each 200G lane has the exactly same format, i.e. an interleave of RS-FEC symbols from 4 different codewords

Clause 186 - 800GBASE-ER1 PCS/PMA

Clause 186 - overall status

- Subclauses 186.1 (overview), 186.2 (PCS), 186.3 (PMA) are mostly done, except as noted:
  - In the PCS, the functions related to the baseline for improving PTP accuracy are not complete (see other slide)
  - In the PMA, there is an open question about how to specify the pilot symbol sequence (see other slide)
- Several functions are specified by reference to documents from other SDOs; reviewers should consider whether this is sufficient

- 186.4-7 are not complete - content is copy/pasted from other clauses and needs to be updated

IEEE P802.3dj Task Force

Clause 186.3 - pilot symbols

- Below is a snippet of part of the table (inherited from P802.3cw)
- The intent is that we use the same pilot sequence as OIF 800ZR and ITU-T FlexO-8e-DO
- ITU-T and OIF specify the sequence in different formats
  - ITU uses bit values per polarization, OIF uses a format that shows each polarization as "I+Qj"
  - E.g., for index 0, X polarization, OIF would say "-3 + 3j", ITU would say "0010"

Is it reasonable to point to OIF, or should we express the values in the form 802.3cw was using?

Table 186–4—Pilot sequence

Index	X_I	X_Q	Y_I	Y_Q	Index	X_I	X_Q	Y_I	Y_Q	Index	X_I	X_Q	Y_I	Y_Q	Index	X_I	X_Q	Y_I	Y_Q
0	–3	3	–3	–3	29	3	–3	3	–3	58	3	–3	3	–3	87	3	–3	–3	3
1	3	3	–3	–3	30	–3	–3	–3	3	59	3	3	–3	3	88	–3	–3	–3	3
2	3	–3	3	–3	31	3	3	–3	–3	60	3	–3	3	3	89	3	–3	3	–3
3	–3	3	3	3	32	–3	3	3	–3	61	–3	–3	–3	–3	90	3	–3	3	3

IEEE P802.3dj Task Force

Clause 186 - PTP accuracy details to be worked through in D1.2

- The adopted baseline uses additional overhead in the 800GBASE-ER1 PCS to convey the position of the AMs from the transmitter to the receiver
- The removal and insertion of AMs is actually done in the PHY_XS, not the PCS
- From an implementation perspective this makes no difference; the PHY_XS and 800GBASE-ER1 PCS are in the same device
- Writing the specification is more complicated…
  - The PHY_XS doesn't know anything about the FEC frame that the PCS creates or the overhead that the PCS uses
  - The PHY_XS isn't exclusively used by the 800GBASE-ER1 PCS; not simple to add functions to PHY_XS that are specific to 800GBASE-ER1 PCS
  - In principle, 800GBASE-ER1 PCS isn't required to use the XS; not simple to add signals to the PCS service interface to convey the information to/from the PHY_XS