

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 00 SC 0 P0 L0 # 185

Brown, Matt Alphawave Semi

Comment Type T Comment Status D Machine Convention (bucket)

Many state diagrams in this draft as well as in the base standard use the operator "++" to indicate that the variable be incremented by 1. However, this operator is never defined.

SuggestedRemedy

Import Clause 21 andà

Amend 21.5 to include definition of "++".

Delete the following from state diagram conventions in multiple clauses. "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."

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Import Clause 21 and...

Amend 21.5 to include definition of "++".

Delete the following from state diagram conventions in 175.2.6.1, 176.5.1.6, 177.6.1, 184.6.1, 176A.10.1.

"The notation ++ after a counter or integer variable indicates that its value is to be incremented."

Implement with editorial license.

Cl 1 SC 1.3 P46 L33 # 506

Dawe, Piers Nvidia

Comment Type TR Comment Status D MDI references (bucket)

Add and update connector references as necessary. This is what is in 1.3:

SFF-8402, Rev 1.1, September 13, 2014, Specification for SFP+ 1X 28 Gb/s Pluggable Transceiver Solution (SFP28).

SFF-8432, Rev 5.1, August 8, 2012, Specification for SFP+ Module and Cage.

SFF-8436, Rev 4.8, October 31, 2013, Specification for QSFP+ 10 Gb/s 4X Pluggable Transceiver.

SFF-8665, Rev 1.9, June 29, 2015, Specification for QSFP+ 28 Gb/s 4X Pluggable Transceiver Solution (QSFP28).

SuggestedRemedy

Use these for now (most will be updated before this project is done):

OSFP Octal Small Form Factor Pluggable Module, Rev 5.0, October 2, 2022

QSFP-DD/QSFP-DD800/QSFP-DD1600 Hardware Specification for QSFP Double Density 8x Pluggable Transceivers, Rev 7.0, September 29, 2023

SFF-8665 Rev 1.9.4, 2022-04-01, QSFP+ 4X Pluggable Transceiver Solutions

SFF-TA-1011 Rev 1.1, 2024-04-19, Cross Reference to Select SFF Connectors and Modules

SFF-TA-1027, Rev 1.0, 2024-04-16, QSFP2 Connector, Cage, & Module Specification

SFF-TA-1031, Rev 1.0, 2023-06-11, SFP2 Cage, Connector, & Module Specification

<https://osfpmsa.org/specification.html>

<http://www.qsfp-dd.com/specification/>

Refer to these documents from 179C.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement suggested remedy with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 1 SC 1.4.184da P49 L43 # 309

D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei

Comment Type TR Comment Status D ER1 PHY (bucket)

800GBASE-ER1 is defined as using 800GBASE-R encoding, but per 802.3df-2024, 1.4.184e - "The term 800GBASE-R represents a family of Physical Layer devices using the Physical Coding Sublayer (PCS) defined in Clause 172 for 800 Gb/s operation." This PHY as noted in Table 169-3a,uses PCS encoding as defined in Clause 186.

SuggestedRemedy

Define new name for family / encoding based on Clause 186 encoding.
Modify definition of entry for 800GBASE-ER1 to reflect new family name.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

The comment correctly points out that the definition is not correct. However, it is not necessary to define a new family.

Change the definition of 800GBASE-ER1 and 800GBASE-ER1-20 to the following:
1.4.184da 800GBASE-ER1: IEEE 802.3 Physical Layer specification for 800 Gb/s PHY using 800GBASE-ER1 PCS and PMA encoding, dual polarization 16-state quadrature amplitude modulation (DP-16QAM) modulation, and coherent detection with reach up to at least 40 km. (See IEEE Std 802.3, Clause 186 and Clause 187).
1.4.184db 800GBASE-ER1-20: IEEE 802.3 Physical Layer specification for 800 Gb/s PHY using 800GBASE-ER1 PCS and PMA encoding, dual polarization 16-state quadrature amplitude modulation (DP-16QAM) modulation, and coherent detection with reach up to at least 20 km. (See IEEE Std 802.3, Clause 186 and Clause 187).
Implement with editorial license.

Cl 1 SC 1.4.184da P49 L44 # 111

Huber, Thomas Nokia

Comment Type T Comment Status D ER1 PHY (bucket)

Since 800GBASE-ER1 and -ER1-20 have a separate PCS, the definition for 800GBASE-ER1 and ER1-20 should refer to 800GBASE-ER1 encoding rather than 800GBASE-R encoding

SuggestedRemedy

Change 800GBASE-R to 800GBASE-ER1 for both the ER1 and ER1-20 definitions.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Resolve using the response to comment #309.

Cl 1 SC 1.4.184da P49 L47 # 310

D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei

Comment Type TR Comment Status D ER1 PHY (bucket)

800GBASE-ER1-20 is defined as using 800GBASE-R encoding, but per 802.3df-2024, 1.4.184e - "The term 800GBASE-R represents a family of Physical Layer devices using the Physical Coding Sublayer (PCS) defined in Clause 172 for 800 Gb/s operation." This PHY as noted in Table 169-3a,uses PCS encoding as defined in Clause 186.

SuggestedRemedy

Define new name for family / encoding based on Clause 186 encoding.
Modify definition of entry for 800GBASE-ER1 to reflect new family name.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Resolve using the response to comment #309.

Cl 1 SC 1.5 P51 L11 # 74

Lusted, Kent Intel Corporation

Comment Type TR Comment Status D (bucket)

The abbreviation "MLSD" is used numerous times in Annex 178A to reference Maximum Likelihood Sequence Detection and should be added to the abbreviations list.

SuggestedRemedy

Add MLSD | Maximum Likelihood Sequence Detection

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement suggested remedy with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 30 SC 30 P56 L33 # 369
 He, Xiang Huawei
 Comment Type TR Comment Status D timesync (bucket)
 Add TimeSync entity managed object classes for Inner FEC sublayers defined in Clause 177 and 184.
 SuggestedRemedy
 Add register set for Inner FEC sublayers in subclauses of 30.13.1: (30.13.1.1 - 30.13.1.14)
 (Presentation will be prepared for this comment.)
 Proposed Response Response Status W
 PROPOSED REJECT.
 The following related presentation was reviewed by the 802.3dj task force during the May Interim meeting:
https://www.ieee802.org/3/dj/public/24_05/he_3dj_01_2405.pdf
 This presentation does not provide sufficient detail to describe the requested change in Clause 30.

Cl 30 SC 30.3.2.1.2 P53 L11 # 112
 Huber, Thomas Nokia
 Comment Type T Comment Status D (bucket)
 There should also be an entry for 800GBASE-ER1 since it is a different PCS
 SuggestedRemedy
 Add a new editing instruction to insert 800GBASE-ER1 after 400GBASE-R.(or before the entry for 800GBASE-R).
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 30 SC 30.3.2.1.3 P53 L21 # 75
 Huber, Thomas Nokia
 Comment Type T Comment Status D (bucket)
 There should also be an entry for 800GBASE-ER1 since it is a different PCS
 SuggestedRemedy
 Add a new editing instruction to insert 800GBASE-ER1 after 400GBASE-R (or before the entry for 800GBASE-R).
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 45 SC 45 P57 L1 # 603
 de Koos, Andras Microchip Technology
 Comment Type T Comment Status D timesync(bucket)
 Inner FEC (Clause 177 or Clause 184) needs MDIO registers for TimeSync. They should look like the PMA/PMD clause registers.
 SuggestedRemedy
 Add the following MDIO registers for the Inner FEC, in the same style as the equivalent PMA/PMD MDIO registers
 - TimeSync capability
 - TimeSync transmit path data delay register
 - TimeSync receive path data delay register
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 The following related presentation was reviewed by the 802.3dj task force during the May Interim meeting:
https://www.ieee802.org/3/dj/public/24_05/he_3dj_01_2405.pdf

The register bits and names described on page 8 of the presentation will be used with the exception that the ability bits will be added to example register "TimeSync PMA/PMD capability (Register 1.1800)" and the new delay registers will be added to MMD 1 from location 1.1820 onwards.

Implement the register bits and names described on page 8 of the presentation and with the exception that the ability bits will be added to example register "TimeSync PMA/PMD capability (Register 1.1800)" and the new delay registers will be added to MMD 1 from location 1.1820 onwards.

Implement with editorial licence.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 45 SC 45 P81 L9 # 370

He, Xiang Huawei
 Comment Type TR Comment Status D timesync(bucket)

Add MDIO interface registers for Inner FEC sublayers defined in Clause 177 and 184.

SuggestedRemedy

Add definitions for the new register set defined for the Inner FEC sublayers in 30.3.1.1 - 30.1.1.14.

(Presentation will be prepared for this comment.)

Proposed Response Response Status W

PROPOSED REJECT.

The following related presentation was reviewed by the 802.3dj task force at the May Interim meeting:

https://www.ieee802.org/3/dj/public/24_05/he_3dj_01_2405.pdf

This presentation concerns TimeSync management and refers to the register set "30.13.1.1 – 30.13.1.14" rather than "30.3.1.1 – 30.1.1.14".

A different comment (#603) addresses adding registers for inner FEC TimeSync.

Another comment (#183) concerns adding additional status counters for the inner FEC which will require new registers.

There is insufficient detail given in this comment (#370) and comment #183 to make a change to Clause 45 for inner FEC register definitions at this time.

Cl 45 SC 45.2.1.60b P65 L17 # 507

Dawe, Piers Nvidia
 Comment Type T Comment Status D (bucket)

Shouldn't LR4 come before LR1 (same reach, narrower) and the order goes up the page, counting the bits forward

SuggestedRemedy

Swap 800GBASE-LR4 and 800GBASE-LR1

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 45 SC 45.2.1.60b P65 L24 # 508

Dawe, Piers Nvidia
 Comment Type T Comment Status D (bucket)

800GBASE-DR4-2 has longer reach than 800GBASE-FR4-500

SuggestedRemedy

Swap them

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 45 SC 45.2.1.60c P67 L21 # 509

Dawe, Piers Nvidia
 Comment Type T Comment Status D (bucket)

It's unfortunate that 800GBASE-ER1 and 800GBASE-ER1-20 are in different registers, and 800GBASE-ER1-20, having less reach, should come first

SuggestedRemedy

Move 800GBASE-ER1 from 1.73.14 to 1.74.0. 1.73.14 goes back to reserved - maybe it can be used for 800GBASE-LR20-1 ;)

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 73 SC 73 P83 L1 # 460

Slavick, Jeff Broadcom
 Comment Type T Comment Status D (bucket)

We are now using a Next Page to advertise IEEE defined PHYs. However the order of when Next Pages are introduced, defined and then used is a bit out of order. So re-arranging the order in which AN is specified would help readers to better understand what how Next Pages are defined, how to use them and when to use them.

SuggestedRemedy

Presentation will be provided.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

The following presentation was reviewed by the 802.3dj task force at the May Interim meeting.

https://www.ieee802.org/3/dj/public/24_05/slavick_3dj_01_2405.pdf

Implement the changes proposed in slavick_3dj_01_2405 with editorial licence and using appropriate editing instructions.

Cl 73 SC 73 P85 L9 # 149

Mi, Guangcan Huawei Technologies Co., Ltd
 Comment Type TR Comment Status D (bucket)

Table 73-5 is missing the indication of higherst priority.

SuggestedRemedy

change 1.6Tb/s 8lane in the capability column to 1.6Tb/s 8 lane, highest priority.

Proposed Response Response Status W

PROPOSED REJECT.

Table 73-5 already indicates "lowest priority" and 73.7.6 contains this text "priority as defined in Table 73-5 (listed from highest priority to lowest priority)". So adding "highest priority" in the Table 73-5 is redundant.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 90A SC 90A P519 L43 # 456
 Opsasnick, Eugene Broadcom
 Comment Type T Comment Status D (bucket)

In table 90A-1, the column titled "Alignment marker/ codeword marker insertion/removal" has a value of 2.56ns for 1.6T in the last row. This value should be the xMII time (at MAC data rate) of one Alignment marker block. The 1.6TE PCS lanes are now running at 100G vs 25G for slower speeds, so this number does not scale directly from the other entries. The value for the 1.6T row should be 1.28ns (a full AM group = 8 256b/257b blocks, so the MII time = 8 * 256 / 1600 = 1.28ns). Note that this column has correct values for 25G, 40G, 50G, and 100G. However, the value listed for 200G, 400G and 800G of 2.56ns should be 5.12ns and should also be fixed in maintenance.

SuggestedRemedy

Change the accuracy impairment value of 2.56 ns to 1.28 ns for the 1.6T Ethernet rate in Table 90A-1.

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 90A SC 90A.3 P519 L43 # 330
 de Koos, Andras Microchip Technology
 Comment Type T Comment Status D (bucket)

For the added row in Table 90A-1, the potential timestamp accuracy impairment due to alignment marker insertion/removal for 1.6T is incorrect. It should be 1.28ns, not 2.56ns. The values for 200G, 400G, and 800G are also erroneous (should all be 5.12ns). I've filed a maintenance request to correct these, too.

SuggestedRemedy

Change 2.56 to 1.28ns in the added row for Table 90A-1

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 93B SC 93B P520 L6710 # 55
 Mellitz, Richard Samtec
 Comment Type TR Comment Status D 93B (bucket)

We have been talking about "die-to-die" loss for while now. Add at test point reference to this and reference to section Annex 93B. One reference to this is in diminico_3dj_01_2307 slide 6 and 7.

SuggestedRemedy

Add TP0d and TP5d to figure 93B-1 and table 93B-1

Proposed Response Response Status W

PROPOSED REJECT.
 Annex 93B is not referenced anywhere in the draft, nor in previous backplane PMD clauses 163 and 137.
 There is no benefit in updating an annex that is not referenced.
 Figure 178-2 is used in this project instead.

Cl 116 SC 116 P94 L6 # 150
 Mi, Guangcan Huawei Technologies Co., Ltd
 Comment Type TR Comment Status D (bucket)

In table 116-3, the last two column, missusage of PMD names.

SuggestedRemedy

change PHY type of CL 178 and 179 in the table to the correct nomenclature, i.e., 200GBASE-KR1 and 200GBASE-CR1

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 116 SC 116 P95 L4 # 151
 Mi, Guangcan Huawei Technologies Co., Ltd
 Comment Type TR Comment Status D (bucket)

In table 116-3a, the last two column, missusage of PMD names.

SuggestedRemedy

change PHY type of CL 178 and 179 in the table to the correct nomenclature, i.e., 400GBASE-KR2 and 400GBASE-CR2

Proposed Response Response Status W

PROPOSED ACCEPT.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 116 SC 116 P102 L5 # 152

Mi, Guangcan Huawei Technologies Co., Ltd
 Comment Type TR Comment Status D (bucket)

200GBASE-R SM PMA delay constraint is missing

SuggestedRemedy

Proposed Response Response Status W

PROPOSED REJECT.
 A suggested remedy is not provided.
 200GBASE-R 8:1, 1:8, and 1:1 PMA types, all SM-PMA types are listed. Note that the term SM-PMA is used to reference any symbol multiplexing PMA, where it would otherwise be ambiguous. In the referenced text the multiplex ratio is unambiguous and the reference to Clause 176 in the notes column backs that up.

Cl 116 SC 116 P107 L4 # 153

Mi, Guangcan Huawei Technologies Co., Ltd
 Comment Type TR Comment Status D (bucket)

In Table 116-9, there should be no applicable SP1 and SP6 for 113.4375GBd PMD lane

SuggestedRemedy

change the content of row SP1 and SP6 in the column of 113.4375GBd PMD lane to N/A

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 116 SC 116.1.4 P94 L6 # 530

Rechtman, Zvi Nvidia
 Comment Type T Comment Status D Conditional PMA (bucket)

The comment refers to Table 116ü3.
 The SM_PMA and BM_PMA introduce a new case of optional PMA implementation. For instance 200GBASE-KR2 PHY cannot implement SM_PMA without implementing 200GAUI-1 C2C interface.
 It will be beneficial to add a note about the conditions which allow/require implementation of BM_PMA and SM_PMA
 Same apply to Table 116ü3a, Table 116ü4, Table 169ü2

SuggestedRemedy

Add a footnote labeled æbÆ next to the æOÆ marking for 200GBASE-R SM-PMA in the entries for 200GBASE-KR2, 200GBASE-KR4, 200GBASE-CR2, and 200GBASE-CR4. The footnote æbÆ should state: æApplicable only when 200GAUI-1 C2C interface is used within the PHY

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #312.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 116 SC 116.1.4 P94 L6 # 312

D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei
 Comment Type TR Comment Status D Conditional PMA (bucket)

200/400G BASE-R BM-PMA and 200/400G BASE-R-SM-PMA are noted as optional in Tables 116-3, 116-4, and 116-4a, but that is not quite correct. They are conditional dependent on the PHY type and on whether specific AUIs are implemented or not.

SuggestedRemedy

For 100Gb/s based PHYs the 200GBASE-R BM-PMA is mandatory, all AUIs are optional, and 200GBASE R SM PMA is "C" / conditional if either 200GAUI-1 is implemented.
 For 200Gb/s based PHYs the 200GBASE-R SM-PMA is mandatory, all AUIs are optional, and 200GBASE R BM PMA is "C" / conditional if either 200GAUI-2 is implemented.

For 100Gb/s based PHYs the 400GBASE-R BM-PMA is mandatory, all AUIs are optional, and 400GBASE R SM PMA is "C" / conditional if either 400GAUI-2 is implemented.
 For 200Gb/s based PHYs the 400GBASE-R SM-PMA is mandatory, all AUIs are optional, and 400GBASE R BM PMA is "C" / conditional if either 400GAUI-4 is implemented.

Change entries as described above in Tables 116-3, 116-4 and 116-4a for 800GBASE-R BM-PMA and 800GBASE-R-SM-PMA to C / with notes as stated above
 Modify entry in Table 178-1 to 200GBASE-R BM PMA to Conditional. Add note "c" A 200GBASE-R BM PMA must be implemented if a 200GAUI-2 C2C is implemented.
 Modify entry in Table 178-2 to 400GBASE-R BM PMA to Conditional. Add note "c" A 400GBASE-R BM PMA must be implemented if a 400GAUI-4 C2C is implemented.
 Modify entry in Table 179-1 to 200GBASE-R SM PMA to Conditional. Add note "c" A 200GBASE-R SM PMA must be implemented if a 200GAUI-1 C2C is implemented.
 Modify entry in Table 179-2 to 400GBASE-R SM PMA to Conditional. Add note "c" A 400GBASE-R SM PMA must be implemented if a 400GAUI-2 C2C is implemented.
 Modify entry in Table 181-1 to 200GBASE-R BM PMA to Conditional. Add note "c" A 200GBASE-R BM PMA must be implemented if a 200GAUI-2 C2C/C2M is implemented.
 Modify entry in Table 180-2 to 400GBASE-R BM PMA to Conditional. Add note "c" A 400GBASE-R BM PMA must be implemented if a 400GAUI-4 C2C/C2M is implemented.
 Modify entry in Table 182-1 to 200GBASE-R BM PMA to Conditional. Add note "c" A 200GBASE-R BM PMA must be implemented if a 200GAUI-2 C2C/C2M is implemented.
 Modify entry in Table 182-2 to 400GBASE-R BM PMA to Conditional. Add note "c" A 400GBASE-R BM PMA must be implemented if a 400GAUI-4 C2C/C2M is implemented.

Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #317.

Cl 116 SC 116.1.4 P98 L18 # 313

D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei
 Comment Type TR Comment Status D (bucket)

there is no PMD called 400GBASE-LR4

SuggestedRemedy

Change 400GBASE-LR4 to 400GBASE-LR4-6

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 116 SC 116.2.4 P99 L1 # 314

D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei
 Comment Type TR Comment Status D PMA introduction (bucket)

In support of 200 Gb/s per lane signaling - 200GBASE-R BM-PMA and 400GBASE-R PMA, Clause 176 was developed. No addition was made to 116.2 Summary of 200GbE and 400 GbE sublayers was made.

SuggestedRemedy

Modify last sentence of 116.2.4 and add additional text
 The 200GBASE-R and 400GBASE-R PMAs, which supports bit multiplexing, is specified in Clause 120.
 The 200GBASE-R and 400GBASE-R PMAs, which supports symbol multiplexing, is specified in Clause 176.
 Note that "PMA" is used as a general term to represent both types of PMAs for each speed.

Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 The comment appropriately proposes to add the new PMA types defined in Clause 176 and to differentiate the two based on multiplexing type. It is not necessary to point out that they may both be referred to as PMA and in fact this could be considered incorrect, since any PMA in the 802.3 standard might be called a PMA.
 Implement the following with editorial license:
 Replace the second sentence in 116.2.4 with appropriate editorial instructions to the following:
 200GBASE-R and 400GBASE-R PMAs that use bit multiplexing (BM-PMA) are specified in Clause 120.
 200GBASE-R and 400GBASE-R PMAs that use symbol multiplexing (SM-PMA) are specified in Clause 176.
 Implement with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 116 SC 116.5 P107 L46 # 510

Dawe, Piers Nvidia
 Comment Type T Comment Status D (bucket)

A new footnote has appeared "At the PCS receive input, 1 UI is equivalent to 1 bit." attached to an unchanged number. There is no equivalent footnote for Table 116-8. In 802.3, "bit" means MAC bit. I don't know what point the footnote is making - that PCS lanes use binary signalling not PAM4? Nor why it is here. If it were kept, it should say "1 bit on a PCS lane" or similar.

SuggestedRemedy

Delete footnote f

Proposed Response Response Status W

PROPOSED REJECT.
 The interface between the PMA and the PCS is an abstract interface. UI interval is the time span of a symbol. Since there there is no physical signal here, only bits are exchanged. The note clarifies that for this interface 1 UI is equivalent to 1 bit being transferred.

Cl 119 SC 119.2.4.1 P111 L26 # 333

de Koos, Andras Microchip Technology
 Comment Type T Comment Status D (bucket)

I understand why the use of the stateless encoder decoder is restricted to 200GBASE-R, and 400GBASE-R over 200Gbps lanes. Allowing it on other PMDs/AUIs would be out-of-scope for the 802.3dj project.
 HOWEVER, shouldn't common sense prevail, here?
 The stateless encoder/decoder was designed such that it is all-but-identical to the stateful encoder, only differing in their treatment of /E/ blocks. Since the 200GBASE-R and 400GBASE-R links are always protected by FEC, it is not as if /E/ blocks can occur at random causing divergent behaviour of the two encoder/decoder types.
 There is absolutely no danger of causing backward-compatibility issues, because the stateful encoder/decoder are still allowed for all PMDs
 The stateless encoder/decoder was added to the standard to allow greater implementation flexibility (removing long timing paths). But any new PCS implementation that may attach to either 100Gbps/lane or 200Gbps/lane PMDs would have to implement the stateful encoder/decoder! With the stateless encoder, the standard is offering more implementation flexibility that implemetors cannot actually use.

SuggestedRemedy

Consider removing the restriction on PMD type when using the stateless encoder and decoder in subclauses 119.2.4.1 and 119.2.5.8, respectively.

Proposed Response Response Status W

PROPOSED REJECT.
 As stated in the comment itself, adding an option to support stateless encoding/decoding for PHYs that are not part of the 802.3dj project is out-of-scope .

Cl 120 SC 120.1.1a P114 L30 # 66

Dudek, Mike Marvell
 Comment Type T Comment Status D PMA introduction (bucket)

Table 116-1 and Table 116-2 include the 200Gb/s per lane PMDs which require the symbol muxing PMA. This bit muxing PMA would only be used for lower speed AUIs. Saying it supports any of the PMDs in the tables is confusing.

SuggestedRemedy

Change to "The 200GBASE-R PMA(s) can support any of the two, or four lane 200Gb/s PMDs in Table 116-1 and the 400GBASE-R PMA(s) can support any of the four, or 8 lane 400Gb/s PMDs in Table 116-2". As a less preferred approach PMD's could be changed to PHYs in the original sentence and an additional sentence could be added saying "The single lane 200Gb/s PMDs in Table 116-1 and the two lane 400Gb/s in table 116-2 require the symbol-muxing PMAs described in clause 176."

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Indeed, the PMA defined in Clause 120 can support only PMDs with per-lane signaling rates of 100 Gb/s or less.
 The referenced paragraph should therefore be corrected.
 In Clause 116...
 Remove 200GBASE-KR1/CR1 from Table 116-3 and change table title to: "PHY type and clause correlation (200GBASE copper with 2 or 4 lanes)"
 Remove 400GBASE-KR2/CR2 from Table 116-3a and change table title to: "PHY type and clause correlation (200GBASE copper with 4 lanes)"
 Create new Table 116-3c with title "PHY type and clause correlation (200GBASE copper with 1 lanes)"
 Include 200GBASE-KR1/CR1 in this table.
 Create new Table 116-3d with title "PHY type and clause correlation (400GBASE copper with 2 lanes)"
 Include 400GBASE-KR2/CR2 in this table.
 In Clause 120...
 Change the referenced sentence to:
 "The 200GBASE-R PMA(s) can support any of the 200Gb/s PMDs in Table 116-3 and Table 116-4, and the 400GBASE-R PMA(s) can support any of the 400Gb/s PMDs in Table 116-3a and Table 116-5."
 Implement with editorial license.
 [Editor's note: CC 116, 120]

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 120F SC 120F.1 P522 L7 # 67
 Dudek, Mike Marvell
 Comment Type T Comment Status D Precoding (bucket)
 Clause 176 is for the symbol mux PMA it should not be used for Annex 120F
 SuggestedRemedy
 Remove the reference to 176.9.1.2
 Proposed Response Response Status W
 PROPOSED REJECT.
 Annex 120F is amended to include 1.6TAUI-16.
 176.8.4 defines the 1.6TBASE-R 16:16 PMA, which has a 16-lane interface that can use 1.6TAUI-16 as a physical interface.
 176.9.1.2 describes the precoding function for all symbol-muxing PMAs, which can also be used in the aforementioned PMA.

Cl 169 SC 169 P116 L15 # 155
 Mi, Guangcan Huawei Technologies Co., Ltd
 Comment Type TR Comment Status D PHY descriptions (bucket)
 same as the previous comment on 800GBASE-CR4
 SuggestedRemedy
 make the description consistent
 Proposed Response Response Status W
 PROPOSED REJECT.
 It is assumed that the referenced "previous comment" is Comment #154.
 The language used here is consistent with other similar PHY types in this table. There is similar differences between the PHYs described in this table and the definitions in 1.4.

Cl 169 SC 169 P116 L17 # 154
 Mi, Guangcan Huawei Technologies Co., Ltd
 Comment Type TR Comment Status D PHY descriptions (bucket)
 In Table 169-1, Row of 800GBASE-CR4 was described as 800Gb/s PHY using 800GBASE-R encoding over four lanes of twinaxial copper cable, which is inconsistent with the description in page 49, 1.4.184aa
 SuggestedRemedy
 make the language consistent.
 Proposed Response Response Status W
 PROPOSED REJECT.
 The language used here is consistent with other similar PHY types in this table. There are similar differences between the PHYs described in this table and the definitions in 1.4.

Cl 169 SC 169 P118 L4 # 156
 Mi, Guangcan Huawei Technologies Co., Ltd
 Comment Type TR Comment Status D (bucket)
 In table 169-3, Phy type and clause correlation was marked incorrectly for the columns of 800GBASE-DR8 PMD and 800GBASE-DR8-2 PMD
 SuggestedRemedy
 remove the unnecessary M in the following rows for 800GBASE-DR8 PMD: 800GBASE-DR4, 800GBASE-FR4-500. remove the unnecessary M in the following rows for 800GBASE-DR8-2 PMD: 800GBASE-DR4-2, 800GBASE-FR4, and 800GBASE-LR4.
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 169 SC 169 P123 L5 # 158
 Mi, Guangcan Huawei Technologies Co., Ltd
 Comment Type TR Comment Status D (bucket)
 In Table 169-4, the delay constraints on 800GBASE-R BM-PMA and 800GBASE-R SM-PMA are missing
 SuggestedRemedy
 add appropriate rows with TBD if no consensus has been built.
 Proposed Response Response Status W
 PROPOSED REJECT.
 800GBASE-R 32:4, 4:32, and 4;4, all SM-PMA types are listed in Table 169-4. Note that the term SM-PMA is used to reference any symbol multiplexing PMA, where it would otherwise be ambiguous. In the referenced text the multiplex ratio is unambiguous and the reference to Clause 176 in the notes column backs that up.

Cl 169 SC 169 P127 L4 # 157
 Mi, Guangcan Huawei Technologies Co., Ltd
 Comment Type TR Comment Status D (bucket)
 In Table 116-6, there should be no applicable SP1 and SP6 for 113.4375GBd PMD lane
 SuggestedRemedy
 change the content of row SP1 and SP6 in the column of 113.4375GBd PMD lane to N/A
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 It is assumed that the comment is referring to Table 169-6 rather than the referenced Table 116-6.
 Implement the suggested remedy with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 169 SC 169.1.3 P116 L42 # 315

D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei

Comment Type TR Comment Status D ER1 PHY (bucket)

800GBASE-ER1-20 and 800GBASE-ER1 are both defined as using 800GBASE-R encoding, but per 802.3df-2024, 1.4.184e - "The term 800GBASE-R represents a family of Physical Layer devices using the Physical Coding Sublayer (PCS) defined in Clause 172 for 800 Gb/s operation." These two PHYs as noted in Table 169-3a, they use PCS encoding as defined in Clause 186.

SuggestedRemedy

Define new name for family / encoding based on Clause 186 encoding.
 Eliminate table entries for ER1-20 and ER1 from Table 169-3a.
 Create new table for PHY type and clause correlation for new family based on Clause 186 encoding.
 Modify description of entry for 800GBASE-ER1-20 in Table 169-1 to reflect new family name.
 Modify description of entry for 800GBASE-ER1 in Table 169-1 to reflect new family name.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 This table lists ALL 800 Gb/s Ethernet PHY types (i.e., 800GBASE), not specifically 800GBASE-R PHY types. The description for 800GBASE-ER1 and 800GBASE-ER1-20 is deceiving and should be updated in line with the definitions in Clause 1. Table 169-3a, lists 800GBASE optical coherent PHY types (not specifically 800GBASE-R), so a separate nomenclature table is not required for 800GBASE-ER1/ER1-20.
 Note that comments 111, 310, and 311 propose changes to the definitions in Clause 1.
 In Table 169-1, change the definitions as follows:
 800GBASE-ER1-20 | 800 Gb/s PHY using 800GBASE-ER1 PCS and PMA encoding, dual polarization 16-state quadrature amplitude modulation (DP-16QAM) modulation, and coherent detection with reach up to at least 20 km (see Clause 187)
 800GBASE-ER1 | 800 Gb/s PHY using 800GBASE-ER1 PCS and PMA encoding, dual polarization 16-state quadrature amplitude modulation (DP-16QAM) modulation, and coherent detection with reach up to at least 40 km (see Clause 187)
 Implement with editorial license.

Cl 169 SC 169.1.3 P116 L43 # 76

Huber, Thomas Nokia

Comment Type T Comment Status D ER1 PHY (bucket)

The descriptions of 800GBASE-ER1-20 and 800GBASE-ER1 should refer to 800GBASE-ER1 encoding rather than 800GBASE-R encoding since the ER1[-20] PCS is distinct from the 800GBASE-R PCS

SuggestedRemedy

Change 800GBASE-R to 800GBASE-ER1 in the last two rows of the table.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #315.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 169 SC 169.1.4 P117 L12 # 317

D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei

Comment Type TR Comment Status D Conditional PMA (bucket)

800GBASE-R BM-PMA and 800GBASE-R-SM-PMA are noted as optional in Tables 169-2, 169-3, and Table 169-3a, but that is not quite correct. They are conditional dependent on the PHY type and on whether specific AUIs are implemented or not. .

SuggestedRemedy

For 100Gb/s based PHYs the 800GBASE-R BM-PMA is mandatory, all AUIs are optional, and 800GBASE R SM PMA is "C" / conditional if either 800GAUI-4 is implemented.
For 200Gb/s based PHYs the 800GBASE-R SM-PMA is mandatory, all AUIs are optional, and 800GBASE R BM PMA is "C" / conditional if either 800GAUI-8 is implemented.

Change entries as described above in Tables 169-2, 169-3 and 169-3a for 800GBASE-R BM-PMA and 800GBASE-R-SM-PMA to C / with notes as stated above.

Modify entry in Table 178-3 to 800GBASE-R BM PMA to Conditional. Add note "c" A 800GBASE-R BM PMA must be implemented if a 800GAUI-8 C2C is implemented.
Modify entry in Table 179-3 to 800GBASE-R SM PMA to Conditional. Add note "c" A 800GBASE-R SM PMA must be implemented if a 800GAUI-4 C2C is implemented.
Modify entry in Table 180-3 to 800GBASE-R BM PMA to Conditional. Add note "c" A 800GBASE-R BM PMA must be implemented if a 800GAUI-8 C2C/C2M is implemented.
Modify entry in Table 181-1 to 800GBASE-R BM PMA to Conditional. Add note "c" A 800GBASE-R BM PMA must be implemented if a 800GAUI-8 C2C/C2M is implemented.
Modify entry in Table 182-3 to 800GBASE-R BM PMA to Conditional. Add note "c" A 800GBASE-R BM PMA must be implemented if a 800GAUI-8 C2C/C2M is implemented.
Modify entry in Table 183-1 to 800GBASE-R BM PMA to Conditional. Add note "c" A 800GBASE-R BM PMA must be implemented if a 800GAUI-8 C2C/C2M is implemented.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Some guidance as to when the two PMA types are used would be helpful. However, it is not as simple as proposed in the suggested remedy. Guidance is required for all PMAs used within the various xAUIs. Annex 176B provides all of the necessary guidance. Each of the tables listing physical layer clauses associated with PMD types (e.g., Table 180-3 for 800GBASE-DR4) already include a reference to Annex 176B for the AUIs, but not for the two PMA types. Additional guidance in these tables would be helpful. In the nomenclature tables in Clause 169 it is not necessary to repeat all of these details nor is there any space in these already crowded tables; instead it would be sufficient, efficient, and future-proof to point back to the PMD clauses for guidance. For each new PMD (Clauses 178, 179, 180 to 183, 185, 186), update the PMD tables in the PMD clause and the associated nomenclature table in Clause 116, 169, and 174, similar to the following for the 800GBASE-DR4 defined in Clause 180. In Table 180-1, for the 800BASE-R BM-PMA row, change "Optional" to "Conditional" with the following footnote:
"If one or two 800GAUI-n is implemented in a PHY, additional 800GBASE-R BM-PMA or SM-PMA sublayers are required according to the guidelines in Annex 176B.6.1."

Attach the same footnote to "Required" in the row for 800GBASE-R SM-PMA.

In Table 169-3...

In the cell (800GBASE-DR4 row, 800GBASE-R BM-PMA column), change "O" to "C".

In footnote "a" add ", C = Conditional (refer to PMD clause for details)."

Implement with editorial license.

Cl 169 SC 169.1.4 P117 L12 # 316

D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei

Comment Type TR Comment Status D PMA introduction (bucket)

Table 169-2 introduces the 800GBASE-R BM-PMA and 800GBASE-R-SM-PMA in Table 169-2, but there is no real explanation to the use of the sub-layers - just the required PMA service interfaces, as noted in Items C&E. The clarification of these two sublayers is actually defined in 176.2 Conventions, which doesnt make sense.

SuggestedRemedy

Move definitions of 800GBASE-R BM-PMA and 800GBASE-R-SM-PMA from 176.2 to 169.1.3 Nomenclature

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

The terms BM-PMA and SM-PMA are defined in 120.1.1 and 176.1.1. The same terms are listed in 176.2, but the items in this larger list are terms for use only within Clause 176.

The definition of BM-PMA and SM-PMA should remain in the subclauses listed above. But they should also be introduced Clause 169.

Resolve using the response to comment #318.

Cl 169 SC 169.1.4 P118 L22 # 69

Dudek, Mike Marvell

Comment Type T Comment Status D (bucket)

There are errors in Table 169-3. 800GBASE-DR8-PMD is not needed for 800GBASE-DR4 or 800GBASE-FR4-500, 800GBASE-DR8-2 PMD is not needed for 800GBASE-DR4-2, 800GBASE-FR4, or 800GBASE-LR4,

SuggestedRemedy

Delete the offending "M"s

Proposed Response Response Status W

PROPOSED ACCEPT.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 169 SC 169.1.4 P118 L22 # 68
 Dudek, Mike Marvell
 Comment Type T Comment Status D (bucket)
 There are errors in Table 169-3. 800GBASE-DR8-PMD is not needed for 800GBASE-DR4 or 800GBASE-FR4-500, 800GBASE-DR8-2 PMD is not needed for 800GBASE-DR4-2, 800GBASE-FR4, or 800GBASE-LR4,
 SuggestedRemedy
 Delete the offending "M"s
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 169 SC 169.1.4 P119 L19 # 320
 D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei
 Comment Type TR Comment Status D Conditional PMA (bucket)
 For 800GBASE-LR1 in Table 169-3a
 800GBASE-R BM-PMA is conditional, pending implementation of 800GAUI-8 C2C/C2M
 800GBASE-R SM PMA is conditional, pending implementation of 800GAUI-4 C2C/C2M
 SuggestedRemedy
 Change entries for 800GBASE-LR1 to C for 800GBASE-R BM-PMA and 800GBASE-R SM-PMA
 Add note "C= Conditional, 800GBASE-R BM-PMA is conditional, pending implementation of 800GAUI-8 C2C/C2M
 800GBASE-R SM PMA is conditional, pending implementation of 800GAUI-4 C2C/C2M"
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #317.
 [Editor's note: Changed subclause from 169.1.3 to 169.1.4]

Cl 169 SC 169.1.4 P119 L20 # 77
 Huber, Thomas Nokia
 Comment Type T Comment Status D (bucket)
 The 800GXS can contain AUIs - so the C2C and C2M clauses should be marked as optional for the ER1 and ER1-20 PHYs, as should the associated PMAs.
 SuggestedRemedy
 Indicate that 800GBASE-R BM-PMA, 800GAUI-8 C2C, 800GAUI-8 C2M, 800GBASE-R SM-PMA, 800GAUI-4 C2C, and 800GAUI-4 C2M are optional for both ER1 and ER1-20 PHYs.
 Proposed Response Response Status W
 PROPOSED REJECT.
 The table references the optional 800GMII Extender which specifies the optional/condition AUIs and PMAs.

Cl 169 SC 169.2 P119 L28 # 319
 D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei
 Comment Type TR Comment Status D ER1 PHY (bucket)
 800GBASE-ER1 and 800GBASE-ER1-20 use the Clause 186 800GBASE-ER1 PCS/PMA.
 This layer is not described as part of 169.2.
 SuggestedRemedy
 Create 169.2.4c 800GBASE-ER1 PCS/PMA
 The 800GBASE-ER1 PCS performs encoding of data from the 800GMII, performs GMP mapping, applies FEC, and transfers the encoded data to the PMA. The 800GBASE-ER1 PMA sublayer 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.
 The 800GBASE-ER1 PCS is specified in Clause xxx.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Amend subclause 169.2.3 (from 802.3df) to the following with appropriate editorial instructions and mark-ups.
 The PCS performs encoding of data from the 800GMII data into a form compatible with the PMA and PMD.
 The 800GBASE-R PCS is specified in Clause 172.
 The 800GBASE-ER1 PCS is specified in Clause 186.
 Implement with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 169 SC 169.2 P119 L28 # 318

D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei
 Comment Type TR Comment Status D PMA introduction (bucket)

In support of 200 Gb/s per lane signaling - 800GBASE-R BM-PMA, Clause 176 was developed. No addition was made to 169.2 Summary of 800 GbE architecture

SuggestedRemedy

Modify 169.2.4 to read -
 The PMA sublayer provides a medium-independent means to support the use of a range of physical media.
 The 800GBASE-R PMA, which supports bit multiplexing, is specified in Clause 173.
 The 800GBASE-R PMA, which supports symbol multiplexing, is specified in Clause 176.
 Note that "PMA" is used as a general term to represent both types of PMAs.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 The comment appropriately proposes to add the new PMA types defined in Clause 176 and to differentiate the two based on multiplexing type. It is not necessary to point out that they may both be referred to as PMA and in fact this could be considered incorrect, since any PMA in the 802.3 standard might be called a PMA.
 Implement the following with editorial license:
 Replace the second sentence in 169.2.4 with appropriate editorial instructions to the following:
 The 800GBASE-R PMA that uses bit multiplexing (BM-PMA) is specified in Clause 173.
 The 800GBASE-R PMA that uses symbol multiplexing (SM-PMA) is specified in Clause 176.
 Implement with editorial license.

Cl 169 SC 169.2 P119 L31 # 193

Ran, Adeo Cisco
 Comment Type TR Comment Status D ER1 PHY (bucket)

A new 800GBASE-ER1 PCS is defined in clause 186. It should be mentioned in the introduction clause, 169.2.3 ("Physical Coding Sublayer (PCS)" in 802.3df) which currently only refers to the 800GBASE-R PCS.

SuggestedRemedy

Bring 169.2.3 into the draft and amend it to include the clause 186 PCS.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #319.

Cl 169 SC 169.3.2 P122 L14 # 322

D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei
 Comment Type TR Comment Status D (bucket)

There is no inter-sublayer interface for the PMA sublayer shown in the figure

SuggestedRemedy

Add placeholder text for future text.

Proposed Response Response Status W

PROPOSED REJECT.
 There is no PMA sublayer in the figure. The 800GBASE-LR1 is defined without a PMA sublayer, instead functions normally associated with a PMA are subsumed in the Inner FEC sublayer.

Cl 169 SC 169.3.2 P122 L35 # 78

Huber, Thomas Nokia
 Comment Type T Comment Status D ER1 PHY (bucket)

A similar diagram is needed for 800GBASE-ER1 and 800GBASE-ER1-20 PHYs.

SuggestedRemedy

Use figure 169-2b as a basis. Replace 800GBASE-R PCS with 800GBASE-ER1 PCS, 800GBASE-LR1 Inner FEC with 800GBASE-ER1 PMA, and 800GBASE-R PMD with 800GBASE-ER1 PMD (and of course renams all the service interfaces to align with that).

Proposed Response Response Status W

PROPOSED REJECT.
 A similar diagram for 800GBASE-ER1 and 800GBASE-ER1-20 is provided in Clause 185 which specifies both of these PMD types. No other PMD is of this form so it is not necessary to show a common diagram in Clause 169.

Cl 169 SC 169.3.2 P122 L54 # 321

D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei
 Comment Type TR Comment Status D ER1 PHY (bucket)

There is no figure describing 800GBASE-ER1/-20 describing inter-sublayer service interfaces including 800GBASE-ER1 PCS/PMA

SuggestedRemedy

Add placeholder text for future text.

Proposed Response Response Status W

PROPOSED REJECT.
 Resolve using the response to comment #78.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 169 SC 169.4 P123 L5 # 532
 Rechtman, Zvi Nvidia
 Comment Type TR Comment Status D (bucket)
 The comment refers to Table 169.4.
 The Inner-FEC delay appears to be missing from the table
 SuggestedRemedy
 add 800GBASE-R inner FEC (values are TBDs)
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 170 SC 170.1 P135 L12 # 461
 Slavick, Jeff Broadcom
 Comment Type T Comment Status D (bucket)
 The title of Clause 173 does include BM.
 SuggestedRemedy
 Remove the BM- from Table 171-1 for the Clause 173 entry and footnote A
 Proposed Response Response Status W
 PROPOSED REJECT.
 The term BM-PMA is used in Table 171-1, because this table includes reference to both BM and SM PMAs, and the convention we agreed on was in such cases to call out both PMAs explicitly. The same convention is used in tables 178-1, 179-1, 180-1, 181-1, 182-1 and 183-1.
 This is explained in 173.1.1 as follows:
 "When necessary for disambiguation, to differentiate the bit-multiplexing PMA (BM-PMA) types defined in this clause from the symbol-multiplexing PMA (SM-PMA) types defined in Clause 176, the term BM-PMA is used. Within this clause the term PMA refers specifically to the BM-PMA."

Cl 171 SC 171.3 P137 L41 # 386
 Nicholl, Gary Cisco
 Comment Type T Comment Status D (bucket)
 There is an issue with subclause 171.3.3 generated by 802.3df. There is an incorrect reference of "171.6.2" in the following bullets:
 ù An additional signal TXRD indicates the state of the rx_rm_degraded variable (see 171.6.2) as detected by the PHY 800GXS in the transmit direction
 ù An additional signal TXLD indicates the state of the FEC_degraded_SER variable (see 171.6.2) as detected by the PHY 800GXS in the transmit direction

SuggestedRemedy
 Import subclause 171.3.3 and correct the two bullets as follows:
 ù An additional signal TXRD indicates the state of the rx_rm_degraded variable (see 172.2.6.2.2) as detected by the PHY 800GXS in the transmit direction
 ù An additional signal TXLD is the logical OR of the FEC_degraded_SER and rx_local_degraded variables (see 172.2.6.2.2) as detected by the PHY 800GXS in the transmit direction.
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 171 SC 171.8 P144 L23 # 79
 Huber, Thomas Nokia
 Comment Type T Comment Status D (bucket)
 In tables 171-3 and 171-5, it is not clear what has changed in the rows that are shown.
 SuggestedRemedy
 Indicate the changes with revision marks
 Proposed Response Response Status W
 PROPOSED REJECT.

Although it may be hard to see, the draft is following 802.3 editing guidelines. The thing that changed in tables 171-3 and 171-5 is that an "_" was added between "FEC_symbol_error_counter" and "<0:31>" in the status variable column. Being added text, the "_" is underlined in keeping with 802.3 editing convention. The missing underscore was missed in the 802.3df draft, including during the final publication review.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 174 SC 174 P164 L20 # 159

Mi, Guangcan Huawei Technologies Co., Ltd

Comment Type TR Comment Status D (bucket)

In Table 174-4, the notes for 1.6TBASE-KR8 and 1.6TBASE-CR8 says includes the medium in one direction. No length of the medium was provided, nor any explicit delay due to the medium was provided. While In Table 169-4, a definitive of 14ns allocated for one direction through cable medium was provided for 800GBASE-CR4. One would assume 1.6TBASE-CR8 would be consistent with 800GBASE-CR4. The same problem applies to 1.6TBASE-KR8.

SuggestedRemedy

Put in explicit allocation of delay constraints for the medium used in 1.6T BASE-CR8 and 1.6TBASE-KR8. Align with that of 800GBASE-CR4 and 800GBASE-KR4, if technically feasibly.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Use the same text used for 800GBASE-KR8/CR8 in IEEE Std 802.3df-2024.

For the 800GBASE-KR4 row change the text in the note column to:

"Includes allocation of 14 ns for one direction through backplane medium. See 178.6."

For 800GBASE-CR4 row change the text in the note column to:

"Includes allocation of 14 ns for one direction through backplane medium. See 179.6."

Cl 175 SC 175 P169 L1 # 332

de Koos, Andras Microchip Technology

Comment Type T Comment Status D timesync (bucket)

Has any thought been given to how to calculate the latency through the 1.6TBASE-R PCS, i.e. the path data delay values for the purposes of TimeSync?

I do not see anything within the 1.6TBASE-R PCS that would prevent proper calculation of the path data delay values.

Clause 90.7.1 is instructive here, explaining that the path data delays should be "reported as if the DDMP is at the start of the FEC codeword". However, the existing language in 90.7.1 is awkward for PCSs with more than one FEC engine like the 1.6TBASE-R PCS, which has four FEC codewords in parallel.

SuggestedRemedy

No proposed change to Clause 175.

Clause 90.7.1 could be cleaned up to account for when there are multiple FEC codewords in parallel, but I assume that is out-of-scope for the 802.3dj project? I'll submit a maintenance request.

Proposed Response Response Status W

PROPOSED REJECT.

The suggested remedy does not propose an actionable (within the draft) remedy.

Cl 175 SC 175.2.1 P172 L26 # 376

Ofelt, David Juniper Networks

Comment Type T Comment Status D (bucket)

Text says to interleave two codewords from flow 0 and two from flow 1, but it isn't clear that those two should be from different FEC encoders.

SuggestedRemedy

After FEC encoding, a FEC codeword from each of the two encoders in flow 0 and a FEC codeword from each of the two encoders in flow 1 are then interleaved and distributed to individual PCS lanes.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the suggested remedy with editorial license.

Cl 175 SC 175.2.4.2 P173 L26 # 481

Slavick, Jeff Broadcom

Comment Type T Comment Status D timesync (bucket)

A note that modifying the data stream could affect TimeSync would be useful.

SuggestedRemedy

Add the following note:

"NOTE -- Insertion or removal of characters may affect protocols like times synchronization (see 90.4.1.2)"

Proposed Response Response Status W

PROPOSED REJECT.

It is not helpful to sprinkle notes related to time synchronization throughout the various sublayer clauses; this was not done in previous clauses/projects. Rather it would be preferable to add the necessary text into Clause 90/Annex 90A. A consensus presentation with a complete proposal is encouraged.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 175 SC 175.2.4.4 P173 L41 # 463

Slavick, Jeff Broadcom

Comment Type T Comment Status D (bucket)

The last sentence is giving the transcoded blocks sent to each flow a name. So it's not really make a flow of blocks. If anything it's making a series or stream of blocks.

SuggestedRemedy

Change the last sentence to read: "The transcoded blocks sent to flow 0 are referred to as tx_xcoded_f0<256:0> and the ones sent to flow 1 as tx_xcoded_f1<256:0>."

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the following with editorial license.

Change:

"This creates two flows of transcoded blocks, tx_xcoded_f0<256:0> to flow 0, and tx_xcoded_f1<256:0> to flow 1."

to:

"This creates two streams of transcoded blocks, tx_xcoded_f0<256:0> to flow 0, and tx_xcoded_f1<256:0> to flow 1."

Cl 175 SC 175.2.4.5 P173 L50 # 331

de Koos, Andras Microchip Technology

Comment Type T Comment Status D Scrambler seeds (bucket)

Different scrambler seeds for the two flows are NOT strictly necessary for the 1.6TBASE-R PCS. The output PCSs are never bit muxed, so having identical outputs from FEC A and FEC C, for example, should never have any adverse effect on "clock content" of the SerDes output.

It doesn't hurt to have the scramblers be seeded differently, however.

SuggestedRemedy

Consider changing the last sentence on page 173 from:
When reset is asserted, the two scramblers shall be initialized to a value other than zero and different from each other.

To:

When reset is asserted, the two scramblers shall be initialized to values other than zero.

(snuck in an editorial correction there, too!)

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Resolve using the response to comment #454.

Cl 175 SC 175.2.4.5 P174 L3 # 454

Opsasnick, Eugene Broadcom

Comment Type T Comment Status D Scrambler seeds (bucket)

The Editor's note at the end of subclause 175.2.4.5 "Scrambler" states that there are no requirements or restrictions in the 1.6TE PCS baselines for the scrambler seeds for each flow. The note also mentions that the corresponding sub-clause in 802.3df for 800GE PCS states that the two flows would have identical outputs if the seeds are identical and the data input is identical (such as after reset). The 1.6TE PCS does not have two separate sets of PCSs like 800GE PCS, but the PCSL formation could have back-to-back repeating RS-symbol values if identical seeds are used. Suggest to require different seeds after reset in the scramblers of each flow as written in the paragraph above the editor's note.

SuggestedRemedy

Remove the editor's note at the top of page 174, and leave the wording in 175.2.4.5 as-is with the requirement that the two scramblers are initialized with different seeds.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Comment #331 notes that the 1.6T PCS lanes are never bit-muxed so different seeds may not be necessary. While the effect of identical scrambler seeds is worse with bit-muxing than symbol-muxing, there may still be some detrimental effects with symbol muxing. If there are identical seeds and identical data, then the FEC-A and FEC-B codewords would be identical to the FEC-C and FEC-D codewords, respectively. With symbol muxing, the resulting data on a output lane would be symbols {A, B, C, D} where A=C and B=D. In general, it is safer to require different seeds to avoid any potential side-affect. As the comment #331 points out, it doesn't hurt to have the scramblers seeded differently.

Delete the editor's note near top of page 174.

Cl 175 SC 175.2.4.5 P174 L3 # 377

Ofelt, David Juniper Networks

Comment Type T Comment Status D Scrambler seeds (bucket)

Editor's Note asks if we should require different reset values for the scramblers.

SuggestedRemedy

Yes, we should!

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Resolve using the response to comment #454.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 175 SC 175.2.4.6 P174 L42 # 464

Slavick, Jeff Broadcom

Comment Type T Comment Status D (bucket)

tx_am_sf doesn't allow but provides a way to communicate the mandatory degrade status.

SuggestedRemedy

Change "allows the local PCS to communicate the status of the FEC degraded feature to the remote PCS" to "communicates the local PCS FEC degraded status to the remote PCS".

Proposed Response Response Status W

PROPOSED REJECT.

The draft is correct as written, and the proposed change does not improve clarity.

Cl 175 SC 175.2.4.6 P175 L22 # 453

Opsasnick, Eugene Broadcom

Comment Type T Comment Status D (bucket)

Sub-clause 172.2.4.6 has a reference to a text file containing the 800GBASE-R alignment marker values. CL 175 should add a similar note with a corresponding text file for the 1.6TBASE-R alignment markers.

SuggestedRemedy

Add text near line 22: "NOTE A text file containing the alignment marker patterns, as shown in Table 175 is available at <https://standards.ieee.org/downloads/802.3/>."

A presentation will be submitted with a corresponding text file containing the 1.6TBASE-R AM values.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Add note as suggested with additional reference to the text file from the May interim (https://www.ieee802.org/3/dj/public/24_05/opsasnick_3dj_02_2405.txt) as presented in https://www.ieee802.org/3/dj/public/24_05/opsasnick_3dj_01_2405.pdf Implement with editorial license.

Cl 175 SC 175.2.4.6 P176 L5 # 465

Slavick, Jeff Broadcom

Comment Type T Comment Status D (bucket)

am_mapped_f0 and am_mapped_f1 aren't solely based on the 10b-distribution and we never talk about how this two variables are us splitting the alignment marker group up.

SuggestedRemedy

Change:

The variables am_mapped_f0 and am_mapped_f1 are then derived from 10-bit interleaving the group of 16 alignment markers, am_x, using the following procedure:

To:
The alignment marker group is mapped into variables am_mapped_f0 and am_mapped_f1 as follows. First a 10-bit interleaving the group of 16 alignment markers, am_x, is done using the following procedure:

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the suggested remedy with editorial license.

Cl 175 SC 175.2.4.6 P176 L25 # 466

Slavick, Jeff Broadcom

Comment Type T Comment Status D (bucket)

am_mapped_f0 and am_mapped_f1 contain data that is sent into flow 0/1 and through codewords AB and CD.

SuggestedRemedy

Change:

Note that am_mapped_f0 contains the 10-bit symbols of FEC codewords A and B, and am_mapped_f1 contains the 10-bit symbols of FEC codewords C and D.

To:
Note that am_mapped_f0 is sent to flow 0 which produces FEC codewords A and B, and am_mapped_f1 is sent to flow 1 which produces FEC codewords C and D.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the suggested remedy with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 175 SC 175.2.4.6.2 P177 L 6 # 467

Slavick, Jeff Broadcom

Comment Type T Comment Status D (bucket)

Add a intro to what tx_scrambled is.

SuggestedRemedy

Change:

"The variables tx_scrambled_am_f0<10279:0> and tx_scrambled_am_f1<10279:0> are constructed in one of two ways."

To:

"In each flow a 10280-bit block of data is formed with two FEC codewords worth of message data, tx_scrambled_am_f0<10279:0> in flow 0 and tx_scrambled_am_f1<10279:0> in flow 1 and they are constructed in one of two ways. "

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
Implement the suggested remedy with editorial license.

Cl 175 SC 175.2.5.3 P182 L 9 # 469

Slavick, Jeff Broadcom

Comment Type T Comment Status D (bucket)

The Note about tracking statistics across all 4 decoders is missing from the bin counter.

SuggestedRemedy

Add this to the definition of the FEC_codeword_error_bin_i
"Note that this counter tracks codewords with errors across all four codewords."

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
Implement the suggested remedy with editorial license.

Cl 175A SC 175A P539 L 8 # 455

Opsasnick, Eugene Broadcom

Comment Type T Comment Status D (bucket)

Annex 175A contains tabular data for an example created by the 1.6TBASE-R PCS TX functions, including the scrambler output, RS-FEC codeword generation, and PCS lane interleaving. The editor's note on page 539 has a placeholder for a link to a text file that has the machine readable text data. That data file needs to be created.

SuggestedRemedy

A presentation is planned to submit a data file which corresponds to the Annex 176A example and can be referenced in the editor's note

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
Update the Editor's note with link to the text file
(https://www.ieee802.org/3/dj/public/24_05/opsasnick_3dj_03_2405.txt) as presented in
https://www.ieee802.org/3/dj/public/24_05/opsasnick_3dj_01_2405.pdf at the May interim.
Implement with editorial license.

Cl 176 SC 176 P195 L1 # 597
 de Koos, Andras Microchip Technology
 Comment Type T Comment Status D timesync (bucket)

Has any thought been put into how to calculate the path data delay values (MII-MDI latencies for timestamping) for the SM-PMAs? For bit-mux PMAs, it is very simple - i.e. it is all implementation delay, since the intrinsic delay from bit muxing/demultiplexing is negligible. But at first glance, determining the latency across the Clause 176 PMA looks like more of a challenge.

- a. I don't believe that the intrinsic (i.e. non-implementation) delay is deterministic, due to the partial deskew.
- b. But apart from the partial deskew, the latency across the SM-PMA should be deterministic using the principles in Annex 90A.7 (max latency value used for Tx path data delay, min latency value used for Rx path data delay).
- c. Traditionally, how to calculate the delays through the PHY layers has been an implementation concern, but this is because the calculation was straightforward at lower rates. At 200Gbps lanes, the standard does not have the luxury of being able to ignore this. If it is overly complicated or ambiguous, and opposite ends of a link do not implement it in the same fashion, the system Time Synchronization will be impaired.

SuggestedRemedy

Consider a note in Clause 176 (or next to the PMA path data delay MDIO registers - 45.2.1.176, 45.2.1.177) that the path data delay values for the SM-PMA should be calculated via the method in Annex 90A.7. I don't think it is necessary, but if a more detailed explanation is deemed useful, then a subclause could be added to Clause 90.7 spelling out explicitly how the path data delay values should be calculated for the SM-PMA.

Proposed Response Response Status W

PROPOSED REJECT.
 It is not helpful to sprinkle notes related to time synchronization throughout the various sublayer clauses; this was not done in previous clauses/projects. Rather it would be preferable to add the necessary text into Clause 90/Annex 90A. A consensus presentation with a complete proposal is encouraged.

Cl 176 SC 176.5.1.1 P200 L1 # 533
 Rechtman, Zvi Nvidia
 Comment Type TR Comment Status D DelayOddPCSLs (bucket)

The comment refers to Figure 176û2.
 The functions of "Delay odd PCSLs by 2 RS-FEC codewords" on Tx path and "Delay even PCSLs by 2 RS-FEC codewords" can be misleading, as they could be interpreted as a delay by 10,880 symbols. The intention is to delay the odd (Tx) and even (Rx) PCSLs by 136 symbols in order to get multiplex and demultiplex symbols from different 2 RS-FEC CWs. Same apply to Figure 176û9

SuggestedRemedy

Modify the description in the Tx path box from "Delay odd PCSLs by 2 RS-FEC codewords" to "Delay odd PCSLs by 136 symbols" and in the Rx path box from "Delay even PCSLs by 2 RS-FEC codewords" to "Delay even PCSLs by 136 symbols"

Proposed Response Response Status W

PROPOSED REJECT.
 The function in Fig 176-2 uses the words "2 RS-FEC codewords" as opposed to "136 RS-FEC symbols" because the function aims to align the 2 codewords on even lanes with 2 different codewords on odd lanes by delaying odd lanes by 2 codewords. This enables symbol multiplexing across 4 codewords. Same applies to Fig 176-9, 176-11 and 176-13. While it is not inaccurate to call it a "136 symbol delay", an advantage of using "2 RS-FEC codewords" as opposed to "136 symbols" is that the function name is equally applicable to both 200GE and 400GE SM-PMAs. Moreover, the first line of subclause 176.5.1.3.4 clearly specifies the delay as being 136 RS-FEC symbols, and the subsequent line shows this mathematically as "2 codewords x 544 symbols per codeword / 8 PCS lanes = 136 symbols." Similarly, subclause 176.6.1.2.4 (400GE 16:2 PMA) specifies the delay to be 68 symbols. Hence, the delay value is clearly specified and there is no room for misinterpretation.
 The comment proposes an alternate description which is technically correct but does not improve the accuracy or readability of the standard.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 176 SC 176.5.1.3.1 P201 L28 # 534

Rechtman, Zvi Nvidia
 Comment Type T Comment Status D (bucket)

There is reference in the text to lock process in Figure 119-12. However, there are exceptions to Figure 119-12 as outlined in 176.5.1.6. It can be beneficial to refer to 176.5.1.6 which include both the reference to Figure 119-12 and the list of exceptions list

SuggestedRemedy

Add a reference to 176.5.1.6 instead of Figure 119-12

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Add note in parenthesis "(see 176.5.1.6.4)" after Fig 119-12.

Implement with editorial license.

Cl 176 SC 176.5.1.3.1 P201 L29 # 475

Slavick, Jeff Broadcom
 Comment Type T Comment Status D (bucket)

There is more details to the AM lock function add a reference

SuggestedRemedy

add a "(see 175.5.1.6.4)" after Table 119-1

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Resolve using the response to comment #534.

[Editor's note: Changed clause, subclause from 175, 175.5.1.3.1 to 176, 176.5.1.3.1]

Cl 176 SC 176.5.1.3.3 P202 L45 # 535

Rechtman, Zvi Nvidia
 Comment Type T Comment Status D (bucket)

The comment refers to Figure 176-4
 The diagram represent a specific skew case between PCS lane, for instance in the absence of skew between the original PCS lanes, the "first" symbol A might be created by different A codeword which should be denote by A'.

SuggestedRemedy

Option1:

Modify only the first A symbol of the odd PCS lanes to be A'.

Option2:

Split the drawing into two: one for 200GBASE-R and another for 400GBASE-R. Then, add index numbers to the A, B symbols.

This could make it easier to understand the drawings and the roles of the symbols in each context.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Update the text referencing Fig 176-4 (in 176.5.1.3.3) and Fig 176-3 (in 176.5.1.3.2) to state that the RS-FEC symbols A and B belong to FEC-A and FEC-B. The "A" symbols could be from the same or different FEC-A codewords and the "B" symbols could be from the same or different FEC-B codewords.

Implement with editorial license.

Cl 176 SC 176.5.1.3.4 P202 L48 # 599

de Koos, Andras Microchip Technology
 Comment Type T Comment Status D (bucket)

The SM-PMA adds a lot of latency due to the 2x RS-FEC CW delay in the 8:1 and 16:2 SM-PMAs, as compared to the bit-mux PMAs

For setups with an MII-Extender it is actually worse, since the penalty would also exist between the DTE_XS and PHY_XS. If latency is a concern, it actually becomes preferable to use 100Gbps links for the DTE_XS <-> PHY_XS AUI interface, negating the advantages of 200Gbps links!

The latency penalty for the 8:1 and 16:2 PMAs should be noted in Clauses 176.5.1.3.4 and 176.6.1.2.4.

SuggestedRemedy

Add the following note to the 2xFEC CW delay sub-clauses (176.5.1.3.4 and 176.6.1.2.4):

Note that the delay added to the odd PCSs (and to the even PCSs at the far-end) causes an end-to-end latency increase of 51.4ns as compared to BM-PMAs.

Proposed Response Response Status W

PROPOSED REJECT.

The standard is not expected to note pros and cons of one PMA versus another (in this case the latency of SM-PMA versus a BM-PMA).

The comment proposes a change that does not improve the clarity or accuracy of the draft.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 176 SC 176.5.1.3.4 P202 L51 # 537

Rechtman, Zvi Nvidia
 Comment Type TR Comment Status D DelayOddPCSLs (bucket)

The sentence "This is equivalent to adding a delay of 2 RS-FEC codewords to the odd PCS lanes (2 codewords $\hat{=}$ 544 symbols per codeword / 8 PCS lanes = 136 symbols)." can be misinterpreted: 136 symbol delay x 4 odd PCS lanes = 544 symbols delay in total (not 2 RS-FEC codewords delay)

SuggestedRemedy

Remove "This is equivalent to adding a delay of 2 RS-FEC codewords to the odd PCS lanes (2 codewords $\hat{=}$ 544 symbols per codeword / 8 PCS lanes = 136 symbols)."

Modify: "Adding the two codeword delay to odd numbered lanes enables the multiplexing of four consecutive RSFEC symbols from four different codewords at the output of the 8:1 symbol multiplexer."

To: "Adding the 136 symbol delay to odd numbered lanes enables the multiplexing of four consecutive RSFEC symbols from four different codewords at the output of the 8:1 symbol multiplexer."

Proposed Response Response Status W

PROPOSED REJECT.
 The first line of subclause 176.5.1.3.4 clearly specifies that the odd lanes are delayed by 136 RS-FEC symbols, and the subsequent line describes mathematically that this (136 symbol delay) is equivalent to adding a delay of 2 codewords to the odd lanes by showing that "2 codewords x 544 symbols per codeword / 8 PCS lanes = 136 symbols". There is little room left for misinterpretation, since the delay in symbols is stated upfront.

Cl 176 SC 176.5.1.3.4 P203 L4 # 293

Galan, Jose Vicente Maxlinear Inc
 Comment Type T Comment Status D Figures (bucket)

For Figure 176-5, it has to be explained what A/E/B/E shall be.

SuggestedRemedy

Add an explanation for A/E/B/E, e. g. "A/E/B/E are the symbols from previous 2 CWs that are delayed"

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Update the text referencing Fig 176-5 (in 176.5.1.3.4) to state that RS-FEC symbols A and A' belong to different codewords from FEC-A, and B and B' belong to different codewords from FEC-B.
 Implement with editorial license.

Cl 176 SC 176.5.1.3.4 P203 L45 # 536

Rechtman, Zvi Nvidia
 Comment Type T Comment Status D Figures (bucket)

The comment refers to Figure 176-5
 The diagram represents a specific skew case between PCS lanes. For instance in the absence of skew between the PCS lanes in the PMA:IS_UNITDATA_0:7.request primitive, the first symbol of A' of the odd PCS lane should be marked as A" because of the additional one symbol delay prior to the 136 symbols delay

SuggestedRemedy

Option1:
 Modify only the first A' symbol of the odd PCS lanes to be A".

Option2:
 Split the drawing into two: one for 200GBASE-R and another for 400GBASE-R. Then, add index numbers to the A, B and A', B' symbols.
 This could make it easier to understand the drawings and the roles of the symbols in each context.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment # 293

Cl 176 SC 176.5.1.3.5 P204 L1 # 291

Galan, Jose Vicente Maxlinear Inc
 Comment Type T Comment Status D Figures (bucket)

In Figure 176-6, the output lane arrow is indicated in the opposite direction than the actual transmission order of the output PCSL symbols

SuggestedRemedy

Change the direction of the arrow to follow the actual transmission order.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Update Fig 176-6 to clarify the order of transmission on the output lane, with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 176 SC 176.5.1.4.2 P204 L42 # 595

de Koos, Andras Microchip Technology

Comment Type T Comment Status D Deskew (bucket)

Is there anything preventing an implementation from performing a full deskew at the Rx PMA? It is not technically required, but does not cause any adverse functional effects. A full deskew at the Rx SM-PMA would NOT change end-to-end latency, since the skew is all ultimately undone at the Rx PCS. A deskew upstream would simply offload the deskew from the Rx PCS.

Implementations with a SM-PMA attached to an RxPCS will undoubtedly perform the Alignment marker lock only once (not once in the PMA and again in the PCS). AM-lock plus deskew is a very natural coupling of functions.

SuggestedRemedy

Consider adding the following note to the Rx Alignment marker lock clauses (176.5.1.4.2, 176.6.1.3.2, 176.7.1.3.2, 176.8.1.3.2):

After the Alignment Marker lock, no deskew of the PCSLs is required. However, deskewing the PCSLs before the would not have and adverse functional effects.

Proposed Response Response Status W

PROPOSED REJECT.

An implementation of the PMA Rx could deskew the PCS lanes during alignment lock (as the comment suggests). However this is an implementation choice, and should not be called out in the standard.

Cl 176 SC 176.5.1.6.4 P206 L38 # 474

Slavick, Jeff Broadcom

Comment Type T Comment Status D (bucket)

Figure 119-12 uses functions and variables defined in CL119 but those aren't called out to be used, just that restart_lock_mux is used to replace restart_lock

SuggestedRemedy

add "using the state variables defined in 119.2.6.2" after Table 119-1 with editorial license

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the suggested remedy with editorial license.

Cl 176 SC 176.5.1.6.5 P206 L48 # 477

Slavick, Jeff Broadcom

Comment Type T Comment Status D (bucket)

Figure 119-12 uses functions and variables defined in CL119 but those aren't called out to be used, just that restart_lock_mux is used to replace restart_lock

SuggestedRemedy

add "using the state variables defined in 119.2.6.2" after Table 119-1 with editorial license

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement suggested remedy with editorial license.

Cl 176 SC 176.5.1.6.5 P208 L11 # 482

Slavick, Jeff Broadcom

Comment Type T Comment Status D (bucket)

Counter _done needs to be at the end of the counter name.

SuggestedRemedy

Change symbol_pair_lock_counter_done_demux to symbol_pair+lock_counter_demux_done

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

In Fig 176-8, change "symbol_pair_lock_counter_done_demux" to "symbol_pair_lock_counter_demux_done". Remove the definition of the variable "symbol_pair_lock_counter_done_demux" from 176.5.1.6.1. Implement with editorial license.

Cl 176 SC 176.5.1.6.6 P207 L6 # 378

Ofelt, David Juniper Networks

Comment Type T Comment Status D (bucket)

Should there be an arc from ALIGNMENT_FAIL to LOSS_OF_ALIGNMENT?

SuggestedRemedy

If so, add the arc

Proposed Response Response Status W

PROPOSED REJECT.

In the ALIGNMENT_FAIL state, restart_lock_mux is set to true which results in AM lock process of Fig 119-12 to be restarted on all lanes. This results in all_locked_mux to be set to false, which causes the state machine of 176-7 to go from ALIGNMENT_FAIL to LOSS_OF_ALIGNMENT state.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 176 SC 176.6.1 P214 L53 # 539

Rechtman, Zvi Nvidia
 Comment Type TR Comment Status D DelayOddPCSLs (bucket)

The comment refers to Figure 176-11.
 The functions of "Delay odd PCSLs by 2 RS-FEC codewords" on Tx path and "Delay even PCSLs by 2 RS-FEC codewords" can be misleading, as they could be interpreted as a delay by 10,880 symbols.
 The intention is to delay the odd (Tx) and even (Rx) PCSLs by 68 symbols in order to get multiplex and demultiplex symbols from different 2 RS-FEC CWs.
 Same apply to Figure 176-13

SuggestedRemedy

Modify the description in the Tx path box from "Delay odd PCSLs by 2 RS-FEC codewords" to "Delay odd PCSLs by 68 symbols" and in the Rx path box from "Delay even PCSLs by 2 RS-FEC codewords" to "Delay even PCSLs by 68 symbols"

Proposed Response Response Status W

PROPOSED REJECT.
 Resolve using the response to comment #533.

Cl 176 SC 176.6.1.2.5 P216 L1 # 290

Galan, Jose Vicente Maxlinear Inc
 Comment Type T Comment Status D Figures (bucket)

In Figure 176-12, the output lane arrow is indicated in the opposite direction than the actual transmission order of the output PCSL symbols

SuggestedRemedy

Change the direction of the arrow to follow the actual transmission order.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Update Fig 176-12 to clarify the order of transmission on the output lane, with editorial license.

Cl 176 SC 176.7.1.2.2 P223 L39 # 459

Opsasnick, Eugene Broadcom
 Comment Type T Comment Status D Figures (bucket)

In Figure 176-16 and Figure 176-17, on the following page, the symbol pattern of the even PCSLs in the upper half (PCSL 16-31) is not shown. It would be easier to see the RS symbol patterns if the figures included at least one even PCSL in the range of 16-31.

SuggestedRemedy

These two figures show PCSLs for lanes 0,1, and 31. Suggest to show the PCSL symbol pattern for lanes 0,1,15, 16, 17,31.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 176 SC 176.7.1.2.2 P223 L52 # 593

de Koos, Andras Microchip Technology
 Comment Type T Comment Status D Figures (bucket)

The 800GBASE-R PCS has 4 FEC engines, so figures 176-16, 176-17, 176-18 should use C,D to illustrate the symbols on PCSLs 16-31, rather than A',B'. The A',B' notation is used in 200GBASE-R and 400GBASE-R figures to denote CWs from engines A and B but with the 2CW delay.

SuggestedRemedy

Amend Figures 176-16, 176-17, 176-18 to avoid the A',B' notation

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Clause 176 avoids using "C" or "D" for 800GBASE-R PMAs because Clause 172 (800GBASE-R PCS) does not use FEC-C and FEC-D. Whereas, "C" and "D" are used in 1.6TBASE-R PMAs because Clause 175 (1.6TBASE-R PCS) uses FEC-C and FEC-D. However, the clarity of the draft will be improved by defining what A, B, A', B' are in the figures Fig 176-16, 176-17 and 176-18.
 Therefore, implement the following:
 Update the text referencing figures Fig 176-16, Fig 176-17 and 176-18 (in 176.7.1.2) to state the RS-FEC symbols A and B are from FEC-A and FEC-B in flow 0 of the 800GBASE-R PCS, while the RS-FEC symbols A' and B' are from flow 1 of the 800GBASE-R PCS.
 Implement with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 176 SC 176.7.1.2.2 P224 L38 # 294
 Galan, Jose Vicente Maxlinear Inc
 Comment Type T Comment Status D Figures (bucket)
 In all Figures in the 800G PMA section, it is referred to AÆ/BÆ symbols, although we have 4 RS CWs
 SuggestedRemedy
 Change to use A,B,C,D for the 4 RS CWs, instead of A, B, AÆ, BÆ
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment # 593

Cl 176 SC 176.7.1.2.4 P225 L1 # 289
 Galan, Jose Vicente Maxlinear Inc
 Comment Type T Comment Status D Figures (bucket)
 In Figure 176-18, the output lane arrow is indicated in the opposite direction than the actual transmission order of the output PCSL symbols
 SuggestedRemedy
 Change the direction of the arrow to follow the actual transmission order.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Update Figure 176-18 to clarify the order of transmission on the output lane, with editorial license.

Cl 176A SC 176A.2.1 P547 L3 # 563
 Law, David HPE
 Comment Type T Comment Status D ILT PICS (Bucket)
 The first 'shall' statement in Annex 176A (normative) 'Control function and start-up protocol for electrical interfaces' is in 176A.2.3.1 'PRBS13 function'. It seems, however, that there should be 'shall' statements in relation to the entire Training frame structure.
 SuggestedRemedy
 [1] In subclause 176A.2.1, change 'The training frame marker is a run ...' to read 'The training frame marker shall be a run ...'.
 [2] In subclause 176A.2.2, change 'The control field comprises ...' to read 'The control field shall be comprised of ...'.
 [3] In subclause 176A.2.2, change 'The status field comprises ...' to read 'The status field shall be comprised of ...'.
 [4] In subclause 176A.2.3, change 'The training pattern is the result of a ...' to read 'The training pattern shall be the result of a ...'.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement suggested remedy with editorial license.

Cl 176A SC 176A.2.3.2 P552 L14 # 199
 Ran, Adeo Cisco
 Comment Type TR Comment Status D ILT Pattern (Bucket)
 "The default identifier for each lane is its lane number (e.g., the default value for identifier_0 is 0 which selects polynomial_0)"
 Some interfaces have 8 lanes.
 The default mapping provided in Table 176A-1 can be used instead.
 SuggestedRemedy
 Change to "The default identifier for each lane is the same as that of the PRBS13 function, as shown in Table 176A-1".
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the following with editorial license.
 Change: "The default identifier for each lane is its lane number"
 To: "The default identifier for each lane is the same as that shown in Table 176A-1"

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 176A SC 176A.2.3.2 P552 L26 # 494
 Slavick, Jeff Broadcom
 Comment Type T Comment Status D ILT Pattern (Bucket)
 The PRBS gen should "stop" if training stops.
 SuggestedRemedy
 Add "while training is in progress while this mode is selected" after "is not stopped or reset".
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the following with editorial license.
 Add "while training is in progress and this mode is selected" after "is not stopped or reset".

Cl 176A SC 176A.2.3.3 P552 L43 # 495
 Slavick, Jeff Broadcom
 Comment Type T Comment Status D ILT Pattern (Bucket)
 The PRBS gen should "stop" if training stops.
 SuggestedRemedy
 Add "while training is in progress while this mode is selected" after "is not stopped or reset".
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the following with editorial license.
 Add "while training is in progress and this mode is selected" after "is not stopped or reset".

Cl 176A SC 176A.3.1 P553 L45 # 499
 Slavick, Jeff Broadcom
 Comment Type T Comment Status D ILT Coefficients (Bucket)
 Remove the specificity of how many presets there are.
 SuggestedRemedy
 Change:
 ôThe initial condition request bits are used to select one of the five predefined transmitter equalizer configurations (presets) specified in the AUI or PMD clauses. ô
 To:
 ôThe initial condition request bits are used to select a predefined transmitter equalizer configurations (presets) specified in the AUI or PMD clauses. ô
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the following with editorial license.
 Change: "The initial condition request bits are used to select one of the five predefined transmitter equalizer configurations (presets) specified in the AUI or PMD clauses." to:
 "The initial condition request bits are used to select one of the up to five predefined transmitter equalizer configurations (presets) specified in the AUI or PMD clauses."

Cl 176A SC 176A.4 P555 L10 # 549
 Rechtman, Zvi Nvidia
 Comment Type T Comment Status D ILT Frame (Bucket)
 The comment refers to Table 176A.3.3 Status field structure.
 The field in bit 14 - "One" require some explanation. It's unclear whether it refers to the support of the newly adopted test patterns, the support of multi-segment operation, or both.
 SuggestedRemedy
 Define the purpose of this bit
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.

Implement the following with editorial license.
 Add new section after the Receiver Ready section:
 "176A.4.2 One
 The one bit is set to 1 to signal the local receiver that the link partner supports the multi-segment control function."
 Note that comment #196 proposes to change "multi-segment control function" to "inter-sublayer link training", which may change the response text above.
 Address this comment have comment #196 is closed.

Cl 176A SC 176A.4 P555 L27 # 501
 Slavick, Jeff Broadcom
 Comment Type T Comment Status D ILT Frame (Bucket)
 You have self generated data you're sending but you don't have your self setup to send mission data yet.
 SuggestedRemedy
 Remove the "No data is available," from the option 1 of Extend training bit
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement suggested remedy with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 176A SC 176A.4.3 P556 L4 # 574

Law, David

HPE

Comment Type T Comment Status D ILT Frame (Bucket)

176A.4.3 'Receiver frame lock' says that 'When the receiver frame lock bit is set to 1, the receiver is indicating that it has identified training frame marker positions and is in a state where the response time requirements specified in 176A.10 are met.'. It then goes on to say 'Receiver frame lock ... is not set to 1 until training and local_tf_lock are both true.'

176A.10 is 'Variables, functions, timers, counters, and state diagrams', so I wonder if the reference should be to 176A.8 'Handshake timing'? In addition, I don't believe the variables training and local_tf_lock are conditioned on the response time requirements specified in 176A.10 being met, at least I didn't see it in their descriptions.

SuggestedRemedy

In 176A.4.3 change the text '... response time requirements specified in 176A.10 are met.' to read '... response time requirements specified in 176A.8 are met.' and the text '... and is not set to 1 until training and local_tf_lock are both true.' To read '... and is not set to 1 until training and local_tf_lock are both true and the response time requirements specified in 176A.10 can be met.'

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the following with editorial license.

Change: "... response time requirements specified in 176A.10 are met."

To: "... response time requirements specified in 176A.8 are met."

Change: "... and is not set to 1 until training and local_tf_lock are both true."

To: "... and is not set to 1 until training and local_tf_lock are both true and the response time requirements specified in 176A.8 can be met."

Cl 176A SC 176A.4.8 P556 L37 # 576

Law, David

HPE

Comment Type T Comment Status D ILT Frame (Bucket)

176A.4.8 'Coefficient status' says 'The acknowledge reflects the value of coef_sts resulting from the procedure described in 176A.6.3.'. While it is correct that the coef_sts variable is updated by the UPDATE_C(k) function in 176A.6.3, I believe the OUT_OF_SYNC, NEW_INDEX, and WAIT states of the Coefficient update state diagram also update the coef_sts variable. Further, 176A.10.3.2 says that the ENCODE_STS function 'Encodes portions of the status field of transmitted training frames.' and that '... coef_sts is mapped to the coefficient status bits ...'.

SuggestedRemedy

Since calls of the UPDATE_C(k) function and direct updates of the coef_sts variable all occur in the Coefficient update state diagram, suggest that 'The acknowledge reflects the value of coef_sts resulting from the procedure described in 176A.6.3.' in 176A.4.8 should be changed to just read 'The acknowledge reflects the value of coef_sts generated by the Coefficient update state diagram'.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

This comment appears to address the same concern expressed in comment #564.

Resolve using the response to comment #564.

Cl 176A SC 176A.4.8 P556 L37 # 564

Law, David

HPE

Comment Type T Comment Status D ILT Frame (Bucket)

176A.4.8 'Coefficient status' says that 'The acknowledge reflects the value of coef_sts resulting from the procedure described in 176A.6.3.'. I don't see a procedure that sets coef_sts in 176A.6.3, but there is one in 176A.6.4. With that said, is it correct that it is just this procedure that sets coef_sts? On review of Figure 176A.9 'Coefficient update state diagram', I see it directly sets coef_sts to 'not_upd' in the OUT_OF_SYNC state and indirectly sets coef_sts using the procedure described in 176A.6.4 through calls to the UPDATE_C(k) function in the NEW_REQUEST state. This seems to be confirmed by the first paragraph of 176A.6.4 which says 'The handling of incoming requests is specified by the coefficient update state diagram (Figure 176A.9). The behavior of the UPDATE_C(k) function shall be consistent with the following algorithm.'

SuggestedRemedy

Change 'The acknowledge reflects the value of coef_sts resulting from the procedure described in 176A.6.3.' to read 'The coefficient status bits reflect the value of coef_sts variable generated by the coefficient update state diagram (Figure 176A.9).'

Proposed Response Response Status W

PROPOSED ACCEPT.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 176A SC 176A.6 P557 L3 # 201

Ran, Adeo

Cisco

Comment Type TR Comment Status D ILT Coefficients (Bucket)

"When the interface control state diagram (Figure 176A.6) is in the TRAIN_LOCAL state, the device may request its link partner to..."

It is important to also note at which states requests from the link partner should be processed, and what happens in the other states - this may not be obvious.

SuggestedRemedy

Insert the following paragraphs after the first one:

When the interface control state diagram is in either the TRAIN_LOCAL or TRAIN_REMOTE state, the device shall respond to requests received from the link partner.

When the interface control state diagram is in any state other than TRAIN_LOCAL or TRAIN_REMOTE, the device shall not send any requests to the link partner and shall ignore requests from the link partner.

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 176A SC 176A.8 P559 L45 # 202

Ran, Adeo

Cisco

Comment Type TR Comment Status D ILT Coefficients (Bucket)

"When the receiver frame lock bit in the status field of transmitted training frames is set to 1, the time from the receipt of a new request to the acknowledgment of that request shall be less than 2 ms"

This requirement was defined in 802.3cd when training was limited in time (to 3 seconds) in order to prevent limiting the number of change requests due to delayed responses.

The new training scheme is not limited in time, and a receiver can use as many requests as it needs.

In some multi-tasking implementations, a hard 2 ms maximum may be challenging to meet. To avoid real-time requirements, it would be sufficient to have 2 ms as the average response time (and it does not need to be normative). The maximum response time can be relaxed without impact to the protocol.

SuggestedRemedy

Change to

"When the receiver frame lock bit in the status field of transmitted training frames is set to 1, the time from the receipt of a new request to the acknowledgment of that request shall be less than 20 ms. It is recommended that the average response time is less than 2 ms".

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 176A SC 176A.9.2 P562 L14 # 555

Law, David

HPE

Comment Type T Comment Status D ILT (Bucket)

Figure 176A.5 'Retimer reference model' shows the data multiplexor driven by the tx_mode value, with the multiplexor select set to 0 when tx_mode = training and set to 1 when tx_mode = data. Subclause 176A.10.2.1 'Variables', however, defines three values for tx_mode, training, local_pattern and data. Figure 176A.5, therefore, does not define the multiplexor select value for when tx_mode = local_pattern.

SuggestedRemedy

Update the figure to reflect the third value of tx_mode and the local pattern generator for each interface.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the following with editorial license.

Add the local_pattern option to the data selector.

Add a Local pattern box as an input to the data selector.

Cl 176A SC 176A.9.2 P562 L22 # 554

Law, David HPE
 Comment Type T Comment Status D ILT (Bucket)

The arrow pointing to the Interface A 'Driver' block and arrow pointing from the Interface B 'CDR' block both seem to be pointing in the wrong direction.

SuggestedRemedy

Reverse the direction of both arrows.

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 176A SC 176A.10.1 P562 L53 # 553

Law, David HPE
 Comment Type T Comment Status D ILT Diagrams (Bucket)

Subclause 176A.10.1 'State diagram conventions' says that 'The notation used in the state diagrams follows the conventions of 21.5.', however subclause 21.5 does not address the operation of timers.

SuggestedRemedy

Suggest that the text 'All timers operate in the manner described in 14.2.3.2.' be inserted as the new second sentence of the second paragraph of subclause 176A.10.1.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the following with editorial license.

Insert the text from clause 136.8.11.7.5: "State diagram timers follow the conventions of 14.2.3.2." as the new second sentence of the second paragraph of subclause 176A.10.1.

Cl 176A SC 176A.10.2.1 P563 L44 # 566

Law, David HPE
 Comment Type T Comment Status D ILT Diagrams (Bucket)

The last sentence of the tx_disable variable description says that the '... output on the lane is disabled.'. Is this correct, the first sentence says that tx_disable '... controls the transmitter's output on the interface.' and tx_disable is defined under subclause 176A.10.2 'Per-interface variables, functions and timers'. Suggest that the reference to 'lane' is changed to 'interface', or use 'all lanes of the interface' in the variable description to reflect the segment_ready variable description immediately above.

SuggestedRemedy

á

Either

á

[a] Change the text '... output on the lane is disabled.' in the last sentence of the tx_disable variable description to read '... output on the interface is disabled.'

á

or

á

[b] Change [1] the text '... the transmitter's output on the interface.' in the first sentence of both the tx_disable and tx_mode variable descriptions to read '... the transmitter output on all lanes of the interface.'; and [2] the text '... output on the lane is disabled.' in the last sentence of the tx_disable variable description to read '... output on all lanes of the interface is disabled.'

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

tx_disable is a per lane variable.

Implement the following with editorial license.

Move the definition of tx_disable to 176A.10.3.

Change the first sentence of the definition...

from: "Boolean variable that controls the transmitter's output on the interface."

to: "Boolean variable that controls the transmitter's output on the lane."

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 176A SC 176A.10.2.1 P563 L44 # 567

Law, David

HPE

Comment Type T Comment Status D ILT Diagrams (Bucket)

Suggest a description of what happens when the tx_disable variable is set to false is added to the variable description.

SuggestedRemedy

[1] Add 'When it is false, tx_mode controls the content of the transmitter's output on the interface.' or 'When it is false, tx_mode controls the content of the transmitter's output on all lanes of the interface.', depending on the response to my other comment, to the end of the tx_disable variable description.

[2] Change the text '... of the interface.' in the first sentence of the tx_mode variable description to read '... of the interface when tx_disable is false.'

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the following with editorial license.

Add the following sentence at the end of the tx_disable definition:
"When it is false, tx_mode controls the content of the transmitter's output on the lane."

Move the definition of tx_mode to 176A.10.3.1 and change the definition of tx_mode...
from: "Enumerated variable that controls the content of the transmitter's output of the interface."
to: "Enumerated variable that controls the content of the transmitter's output of the lane when tx_disable is false."

Cl 176A SC 176A.10.3.1 P565 L7 # 573

Law, David

HPE

Comment Type T Comment Status D ILT Diagrams (Bucket)

The description of the local_tf_lock variable in 176A.10.3.1 says that 'The value of this variable is encoded as the "training lock" bit in the status field of transmitted training frames.', however, there isn't a "training lock" bit defined for the training frames. Since 176A.4.3 'Receiver frame lock' says 'Receiver frame lock ... is not set to 1 until training and local_tf_lock are both true.' it seems that local_tf_lock is encoded in the 'Receiver frame lock' bit.

SuggestedRemedy

Change the text '... is encoded as the "training lock" bit ...' in the local_tf_lock variable description to read '.... is encoded in the "Receiver frame lock" bit ...'.

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 176A SC 176A.10.4 P566 L54 # 542

Rechtman, Zvi

Nvidia

Comment Type TR Comment Status D ILT Diagrams (Bucket)

The operation of precoding after the completion of the start-up protocol is missing

SuggestedRemedy

Add the following text:

"If the LINK_READY state is entered with local_tp_mode set to ôPAM4 with precoding, then the PMA shall transmit all subsequent data on the corresponding lane with precoding (see 176.9.1.2).

If the LINK_READY state is entered with remote_tp_mode set to ôPAM4 with precoding, then the PMA shall subsequently received data on the corresponding lane includes precoding (see 176.9.1.2)"

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the following with editorial license.

After the first paragraph of 176A.10, add the following text:

If the LINK_READY state in the Interface control state diagram (see Figure 176A-6) is entered with local_tp_mode set to "PAM4 with precoding", then the PMD or AUI shall cause the adjacent PMA to transmit all subsequent data on the corresponding lane with precoding (see 176.9.1.2).

If the LINK_READY state is entered with remote_tp_mode set to "PAM4 with precoding", then the PMD or AUI shall inform the adjacent PMA that all subsequently received data on the corresponding lane includes precoding (see 176.9.1.2).

Cl 176A SC 176A.10.4 P568 L20 # 551

Law, David

HPE

Comment Type T Comment Status D ILT Diagrams (Bucket)

There is a spurious '<' withing the transition condition from the state TRAIN_LOCAL to the state TRAIN_REMOTE.

SuggestedRemedy

Suggest that 'local_tf_lock<* local_rx_ready' should read 'local_tf_lock * local_rx_ready'.

Proposed Response Response Status W

PROPOSED ACCEPT.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 176A SC 176A.10.4 P568 L20 # 552
 Law, David HPE
 Comment Type T Comment Status D ILT Diagrams (Bucket)
 There should be an underscore between the timer name and 'done'.
 SuggestedRemedy
 Suggest that 'recovery_timer done' should be changed to read 'recovery_timer_done'.
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 176A SC 176A.10.4 P569 L17 # 556
 Law, David HPE
 Comment Type T Comment Status D ILT Diagrams (Bucket)
 The WAIT_ADJACENT to SWITCH_CLOCK transition condition uses the variable mr_training_enabled, however subclause 176A.10.2.1 'Variables' defines the variable mr_training_enable, not mr_training_enabled.
 SuggestedRemedy
 Change the transition condition '(!mr_training_enabled + segment_ready) * ...' to read '(!mr_training_enable + segment_ready) * ...'.
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 176D SC 176D.1 P595 L16 # 584
 Ghiasi, Ali Ghiasi Quantum/Marvell
 Comment Type T Comment Status D Channel ILdd (bucket)
 C2C loss is TBD
 SuggestedRemedy
 Assuming 28 dB budget and package A length ~300 mm and ~125 mm for package B
 Proposed Response Response Status W
 PROPOSED REJECT.
 The comment addresses an open TBD, but the suggested remedy is unclear.
 Also, the suggested remedy assumes the budget is 28 dB, but consensus on that has not been shown.

Cl 176D SC 176D.2 P596 L19 # 62
 Dudek, Mike Marvell
 Comment Type T Comment Status D (bucket)
 The note "The electrical specifications of C2C components are not equivalent to those of the corresponding PMD's isn't helpful. What does "not equivalent" mean?. Which corresponding PMD's?
 SuggestedRemedy
 Delete the note.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #64.

Cl 176D SC 176D.2 P596 L32 # 583
 Ghiasi, Ali Ghiasi Quantum/Marvell
 Comment Type T Comment Status D (bucket)
 Functional block diagram shown for C2C indicate ball-ball specifications
 SuggestedRemedy
 C2C component should be called C2C device and change the TP0 to TP0d and TP5 to TP5d
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 176D SC 176D.3.3 P597 L33 # 398
 Wu, Mau-Lin MediaTek
 Comment Type TR Comment Status D (bucket)
 The value of '106.255 +/- 50 ppm' is not correct.
 SuggestedRemedy
 Change '106.255' to '106.25'.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #361.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 176D SC 176D.3.3 P597 L33 # 423
 Li, Tobey MediaTek
 Comment Type TR Comment Status D (bucket)
 Signaling rate of 106.255 50 ppm in Table 176Dû1 is incorrect
 SuggestedRemedy
 Change "106.255 50 ppm" to "106.25 50 ppm"
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #361.

Cl 176D SC 176D.3.3 P597 L33 # 361
 Healey, Adam Broadcom Inc.
 Comment Type T Comment Status D (bucket)
 Typo.
 SuggestedRemedy
 Change "106.255" to "106.25".
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 176D SC 176D.3.4.4 P602 L47 # 424
 Li, Tobey MediaTek
 Comment Type TR Comment Status D ERL (bucket)
 Reference to ERL methodology is missing
 SuggestedRemedy
 Add reference to 176D.4.3.
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 176D SC 176D.3.4.4 P603 L30 # 426
 Li, Tobey MediaTek
 Comment Type TR Comment Status D (bucket)
 "Insertion loss at 26.5625 GHz"
 Nyquist frquency in Table 176Dû4 is incorrect
 SuggestedRemedy
 Change "26.5625 GHz" to "53.125 GHz"
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 176D SC 176D.3.4.4 P603 L31 # 451
 Simms, William NVIDIA
 Comment Type TR Comment Status D (bucket)
 Moot point maybe given table is all TBD, but the frequency should be 53.125GHz
 SuggestedRemedy
 change to 53.125GHz
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #426.

Cl 176D SC 176D.3.4.5 P604 L1 # 428
 Li, Tobey MediaTek
 Comment Type TR Comment Status D Editorial (bucket)
 Reference to test procedure is missing
 SuggestedRemedy
 Add reference to 176D.3.4.4
 Proposed Response Response Status W
 PROPOSED ACCEPT.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl **176D** SC **176D.4** P**604** L**27** # **429**
 Li, Tobey MediaTek
 Comment Type **TR** Comment Status **D** Editorial (bucket)
 Table reference is missing
 SuggestedRemedy
 Add reference of ERL to 176D.4.3.
 Add reference of differential-mode to common-mode return loss to 176D.4.4.
 Proposed Response Response Status **W**
 PROPOSED ACCEPT.

Cl **176D** SC **176D.4.1** P**605** L**16** # **122**
 Sakai, Toshiaki Socionext
 Comment Type **T** Comment Status **D** COM pkg tau (bucket)
 COM reference package parameter vlaue. (transmission line parameter tau)
 In "Table 176D06" class A package model Transmission line parameter t(tau) value is
 6.141e-4 ns/mm, but based on the adopted motion#10, Nov/2024, Ilim_3dj_01a_2311.pdf
 (page8-9), the value is 6.141e-3. The value should be 6.141e-3 ns/mm.
 SuggestedRemedy
 Change t(tau) value in Table 176D-6 (class A package) from 6.141e-4 ns/mm to 6.141e-3
 ns/mm.
 Or simply delete this row, as the t(tau) value in table 93A-3 is 6.141e-3 ns/mm.
 Proposed Response Response Status **W**
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #118.

Cl **176D** SC **176D.4.1** P**605** L**26** # **123**
 Sakai, Toshiaki Socionext
 Comment Type **T** Comment Status **D** COM pkg tau (bucket)
 COM reference package parameter vlaue. (transmission line parameter tau)
 In "Table 176D06" classB package model Transmission line parameter t(tau) value is
 6.141e-4 ns/mm, but based on the adopted motion#10, Nov/2024, Ilim_3dj_01a_2311.pdf
 (page8-9), the value is 6.141e-3. The value should be 6.141e-3 ns/mm.
 SuggestedRemedy
 Change t(tau) value in Table 176D-6 (class B package) from 6.141e-4 ns/mm to 6.141e-3
 ns/mm.
 Or simply delete this row, as the t(tau) value in table 93A-3 is 6.141e-3 ns/mm.
 Proposed Response Response Status **W**
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #118.

Cl **176D** SC **176D.4.2** P**607** L**31** # **63**
 Dudek, Mike Marvell
 Comment Type **T** Comment Status **D** Channel ILdd (bucket)
 An insertion loss of only 20dB is less than desirable and the equation is TBD. We
 shouldn't specify the loss at this time
 SuggestedRemedy
 Change 20dB to TBD.
 Proposed Response Response Status **W**
 PROPOSED ACCEPT IN PRINCIPLE.
 The value 20 dB was not adopted, and its appearance here is unintended.
 Slide 18 of https://www.ieee802.org/3/dj/public/24_01/ran_3dj_01a_2401.pdf states
 explicitly that the interconnect length is TBD.
 Implement suggested remedy with editorial license.

Cl **176E** SC **176E.2** P**615** L**20** # **64**
 Dudek, Mike Marvell
 Comment Type **T** Comment Status **D** (bucket)
 The note "The electrical specifications of C2C components are not equivalent to those of
 the corresponding PMD's. Specifically the test points at which module compliance is
 defined are different isn't helpful. What does "not equivalent" mean?. Which
 corresponding PMD's? Although the module test points are different those for the host are
 the same as Clause 179.
 SuggestedRemedy
 Delete the note.
 Proposed Response Response Status **W**
 PROPOSED ACCEPT IN PRINCIPLE.
 The corresponding PMDs are noted in the third paragraph of 176E.2, which states that a
 C2M component is functionally equivalent to a PMD.
 The note appears after the paragraph about the electrical characteristics, and highlights the
 essential difference between a C2M component and a PMD. It is specific about the test
 point difference for the module.
 The description of the C2M component's similarity to a PMD is new, and noting the
 differences is useful for readers.
 However, the term "corresponding PMDs" can be clarified.
 In 176E.2, change "the corresponding PMDs" to "the corresponding PMDs defined in
 Clause 179".
 In 176D.2, change "the corresponding PMDs" to "the corresponding PMDs defined in
 Clause 178".

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 176E SC 176E.2 P615 L23 # 129
 Ghiasi, Ali Ghiasi Quantum/Marvell
 Comment Type T Comment Status D Channel ILdd (bucket)
 Figure depicts loss should be bump-bump
SuggestedRemedy
 ...application and the associated ILdd bump-bump budget at 53.125 GHz
 To make it more clear Host C2M Component should be changed to Host C2M Device and Module C2M Device
Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 The C2M loss budget is currently TBD, but it is expected that it will be inclusive of packages.
 However, the suggested remedy does not significantly clarify this fact.
 It is preferable to align the diagram with Figure 179-2, where the paths between TP0d and TP1 and between TP4 and TP5d are shown to include the package.
 In figure 176E-2, change "Host ILdd" to "Host package and PCB ILdd", and "Module ILdd" to "Module package and PCB ILdd".

Cl 176E SC 176E.4.1 P632 L6 # 134
 Ghiasi, Ali Ghiasi Quantum/Marvell
 Comment Type T Comment Status D (bucket)
 Loss is TBD
SuggestedRemedy
 See Ghiasi C2M May-24 Contribution for background on the numbers
 Bump-bump Insertion loss at Nyquist frequency (53.125 GHz) is less than or equal to 28 dB
Proposed Response Response Status W
 PROPOSED REJECT.
 The loss in the text is TBD because equation 176E-3 has TBDs. When an equation is provided, the text can be changed accordingly, but the comment does not propose values for the equation.
 The following presentation was reviewed by the task force in the May 2024 interim meeting:
https://www.ieee802.org/3/dj/public/24_05/ghiasi_3dj_02_2405.pdf
 The presentation does not include a proposa equation 176E-3.
 [Editor's note: changed page from 621 to 632]

Cl 176E SC 176E.5.2 P633 L39 # 135
 Ghiasi, Ali Ghiasi Quantum/Marvell
 Comment Type T Comment Status D (bucket)
 Eye opening reference receiver parameters will be different between TP1d and TP4a measurement
SuggestedRemedy
 Given that number of module plug implementation will have COC or even if there is package it will be core-less ~8 mm so there is no need to add package after HCB given the loss of the HCB and plug boards are similar.
 At TP4a this is just the output of the module should be tested with synthetic
 - short trace
 - long trace
 recommendation is to measure at the ASIC ball otherwise we would need at least 2 test cases with Package A and 2 with Package B
Proposed Response Response Status W
 PROPOSED REJECT.
 The suggested remedy does not propose an actionable (within the draft) remedy.

Cl 176E SC 176E.5.2 P636 L49 # 523
 Dawe, Piers Nvidia
 Comment Type TR Comment Status D (bucket)
 "within the time interval $t_s \pm 0.05 U_I$ and with accumulated probability for each sample weighted by the function $w(t)$ defined by Equation (176E-4)": this makes the measurement too tolerant to jitter.
SuggestedRemedy
 Remove the Gaussian weighting function $w(t)$, increase ± 0.05 to ± 0.07 , same as TDECQ. This will make VEC look worse, but will be a better measurement to protect the link. Use this method for CR also, with "software channel" ("far end eye measurement") as appropriate.
Proposed Response Response Status W
 PROPOSED REJECT.
 The comment does not provide sufficient justification to support the suggested remedy.
 The suggested remedy does not provide sufficient detail to implement.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 177 SC 177.1.3 P249 L10 # 81
 Huber, Thomas Nokia
 Comment Type T Comment Status D (bucket)
 The second bullet could be written more clearly
 SuggestedRemedy
 Revise to read "Distributing (collecting) the convolutional interleaved data to (from) eight Inner FEC flows
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 177 SC 177.1.3 P249 L14 # 82
 Huber, Thomas Nokia
 Comment Type T Comment Status D (bucket)
 The fifth bullet could be written more clearly
 SuggestedRemedy
 Revise to read "8:1 interleaving (1:8 deinterleaving) the eight Inner FEC flows to (from) a single flow"
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 177 SC 177.1.4 P250 L25 # 83
 Huber, Thomas Nokia
 Comment Type T Comment Status D PAM4 decoding (bucket)
 Indicating PAM4 decoding as optional seems a bit misleading. The P{MD isn't doing soft-decoding in any case, so the FEC must do some sort of decoding to recover the bits from the PAM4 symbols.
 SuggestedRemedy
 Generalize the label in the box to "Decoding", and explain in the text in 177.5.x that there are multiple options for decoding.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Remove footnote in Figure 177-2.

Cl 177 SC 177.1.4 P250 L32 # 543
 Rechtman, Zvi Nvidia
 Comment Type T Comment Status D PAM4 decoding (bucket)
 The comment refers to Figure 177ü2.
 There is a footnote that PAM4 decoding is optional in case of soft decoding.
 However, the DataPath is defined using bit streams, also the FEC:IS_UNITDATA_i.indication primitives has two value of 0 or 1, therefore PAM4 decoding must to take place
 SuggestedRemedy
 Either remove the footnote, or elaborate on the intention of this footnote.

Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment # 83.

Cl 177 SC 177.4.1 P251 L36 # 605
 de Koos, Andras Microchip Technology
 Comment Type T Comment Status D timesync (bucket)
 Due primarily to the convolutional interleaver/deinterleaver, there is a large variation in the input-to-output latency of the Inner FEC sublayer. As such, there is concern that the method to properly calculate the path data delay for the Inner FEC sublayer should be explained in Clause 90, similarly to what is done for the variation from FEC codewords and PCS-lane distribution in clause 90.7.1.

SuggestedRemedy
 Do nothing.
 Using the general method in Clause 90A, allocating the maximum value of the intrinsic delay to the transmit PHY and the minimum value of the intrinsic delay to the receive PHY, there is no ambiguity.
 So it should not be necessary to add to Clause 90 for every new PHY type. The principles laid out in Annex 90A.7 should apply.
 If anything, a general note could be added in Clause 177 (or in Clause 45 with the MDIO registers for path data delay values) explaining that the Tx/Rx path data delay values should be calculated following the guidelines in Annex 90A.7, where the maximum latency value is used for the Tx path data delay, and the minimum latency value is used for Rx path data delay.

Proposed Response Response Status W
 PROPOSED REJECT.
 The suggested remedy does not propose an actionable (within the draft) remedy. It is not helpful to sprinkle notes related to time synchronization throughout the various sublayer clauses; this was not done in previous clauses/projects. Rather it would be preferable to add the necessary text into Clause 90/Annex 90A. A consensus presentation with a complete proposal is encouraged.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 177 SC 177.4.1 P251 L50 # 610

Huang, Kechao Huawei Technologies Co., Ltd.
 Comment Type T Comment Status D Cl (bucket)

"The convolutional interleaver is composed of 3 delay lines where the first delays the PHYs data by eight RS-FEC codewords, the second by four RS-FEC codewords and the last adds no delay" is correct only if the Q values are 544/272/136/68 for 200G/400G/800G/1.6T. However, the Q values should be 192/96/48/24 as shown in slides 6-11 of he_3dj_01_2307 for 200G/400G/800G/1.6TbE.

SuggestedRemedy

Suggest to modify Line 50-51 in page 251 as follows:
 The convolutional interleaver is composed of three parallel delay lines (numbered 0 to 2), as illustrated in Figure 177.3. Each delay operator $\delta D\delta$ represents a storage element of 40 bits. From one delay line to the next higher delay line, Q delay operators are deleted. Modify the Q values to 192/96/48/24 for 200G/400G/800G/1.6T

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #366.

Cl 177 SC 177.4.1 P251 L51 # 544

Rechtman, Zvi Nvidia
 Comment Type TR Comment Status D Cl (bucket)

The values of Q and the description of the Convolutional interleaver functionality doesn't match the adopted values in he_3dj_01_2307.pdf
 The values should be:
 200G BASE-R: Q = 192
 400G BASE-R: Q = 96
 800G BASE-R: Q = 48
 1.6T BASE-R: Q = 24

SuggestedRemedy

Modify the Q values to:
 200G BASE-R: Q = 192
 400G BASE-R: Q = 96
 800G BASE-R: Q = 48
 1.6T BASE-R: Q = 24

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #366.

Cl 177 SC 177.4.1 P252 L9 # 366

He, Xiang Huawei
 Comment Type TR Comment Status D Cl (bucket)

The Q values are not the same as the baseline adopted.

SuggestedRemedy

According to the adopted baseline, change the Q values as follows:
 — 200G BASE-R: Q = 192
 — 400G BASE-R: Q = 96
 — 800G BASE-R: Q = 48
 — 1.6T BASE-R: Q = 24

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 177 SC 177.4.1 P252 L9 # 292

Galan, Jose Vicente Maxlinear Inc
 Comment Type TR Comment Status D Cl (bucket)

The Q values of Convolutional interleaver are not in line with previous contributions, D0.1, D0.2, with the TP2 test vectors of Annex 177A and have to be corrected.

SuggestedRemedy

Q=24 for 1.6TBASE-R, Q=48 for 800GBASE-R, Q=96 for 400GBASE-R and Q=192 for 200GBASE-R

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #366.

Cl 177 SC 177.4.1 P252 L18 # 295

Galan, Jose Vicente Maxlinear Inc
 Comment Type T Comment Status D Cl (bucket)

Usually, a convolutional interleaver switches round-robin from low to high delay lines and the convolutional de-interleaver switches round-robin from high to low delay lines. Why in Figure 177-3 it is defined the other way round?

SuggestedRemedy

Change the convolutional interleaver order if that is the case.

Proposed Response Response Status W

PROPOSED REJECT.
 This is consistent with the adopted baseline. It is correct as documented.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

CI 177 SC 177.4.1 P252 L19 # 488
 Slavick, Jeff Broadcom
 Comment Type T Comment Status D CI (bucket)
 The delay line for CI177 starts with feeding data into the longest delay line while CI184 sends it to the delay line with the shortest delay.
 SuggestedRemedy
 Change CI177 to have the Delay Line 0 be the minimal delay and the Delay Line 2 to be the longest delay.
 Proposed Response Response Status W
 PROPOSED REJECT.
 This is consistent with the adopted baseline. It is correct as documented.

CI 177 SC 177.4.1 P256 L50 # 545
 Rechtman, Zvi Nvidia
 Comment Type TR Comment Status D CI - Editorial (bucket)
 The description in "The convolutional interleaver is composed of 3 delay lines where the first delays the PHYs data by eight RS-FEC codewords, the second by four RS-FEC codewords and the last adds no delay"
 Seems to represent block interleave and not convolutional interleave.
 SuggestedRemedy
 Modify to:
 "The convolutional interleaver is composed of 3 delay lines.
 For 200GBASE-R the first line (line0) delays the PHYs data by $4 \times 2 \times 192 = 1,536$ RS-FEC Symbols, the second line (line1) by $4 \times 1 \times 192 = 768$ RS-FEC symbols and the last line (line3) adds no delay.
 For 400GBASE-R the first line (line0) delays the PHYs data by $4 \times 2 \times 96 = 768$ RS-FEC Symbols, the second line (line1) by $4 \times 1 \times 96 = 384$ RS-FEC symbols and the last line (line3) adds no delay
 For 800GBASE-R the first line (line0) delays the PHYs data by $4 \times 2 \times 48 = 384$ RS-FEC Symbols, the second line (line1) by $4 \times 1 \times 48 = 192$ RS-FEC symbols and the last line (line3) adds no delay
 For 1.6TBASE-R the first line (line0) delays the PHYs data by $4 \times 2 \times 24 = 192$ RS-FEC Symbols, the second line (line1) by $4 \times 1 \times 24 = 96$ RS-FEC symbols and the last line (line3) adds no delay.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggest remedy with editorial license.

CI 177 SC 177.4.1 P256 L53 # 546
 Rechtman, Zvi Nvidia
 Comment Type T Comment Status D CI - Editorial (bucket)
 The input and output round-robin operation is defined relatively to the delay/buffering size of each lane. However, there are lines index that represent the delay and simplify the definition.
 SuggestedRemedy
 Change:
 "The input data round-robins between the three delay lines beginning with the eight RS-FEC delay line, then the four RS-FEC delay line and lastly the zero delay line. The output of the convolutional interleaver round-robins between the three delay lines receiving one RS-FEC symbol-quartet from each at a time beginning with the eight RS-FEC delay line, then four RS-FC delay line, and lastly the zero delay line"
 To:
 "The input data round-robins between the three delay lines beginning with the line0, then line1 delay line and lastly line2. The output of the convolutional interleaver round-robins between the three delay lines receiving one RS-FEC symbol-quartet (4 symbols) from each at a time beginning with line0, then line1, and lastly line2"
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggest remedy with editorial license.

CI 177 SC 177.4.3 P252 L37 # 607
 de Koos, Andras Microchip Technology
 Comment Type T Comment Status D Circular Shift (bucket)
 Was there not a proposal to make the circular shift optional, in order to minimize latency?
 SuggestedRemedy
 Consider removing the circular shift if it does offer not any worthwhile benefit.
 Proposed Response Response Status W
 PROPOSED REJECT.
 This is consistent with the baseline adopted. The comment does not provide sufficient justification to support the suggested remedy.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 177 SC 177.4.3 P252 L37 # 606

de Koos, Andras Microchip Technology
 Comment Type T Comment Status D Circular Shift (bucket)

I'm not convinced that the circular shift really adds any robustness. Yes, it distances bit-pairs belonging to the same RS-FEC codeword, butà Without the shift, the consecutive bit pairs (after 8:1 multiplexing) belonging to the same RS-FEC code words would each protected by different Inner FEC code words, would they not?
 So is the circular shift just protecting against uncorrected inner-FEC codewords that would all land on the same RS-FEC codeword? Seems overkill. Are there simulations/models showing the benefit of including circular shift?

SuggestedRemedy

Consider removing the circular shift if it does not offer any worthwhile benefit.

Proposed Response Response Status W

PROPOSED REJECT.
 This is consistent with the baseline adopted. The comment does not provide sufficient justification to support the suggested remedy.

Cl 177 SC 177.4.4 P253 L48 # 611

Huang, Kechao Huawei Technologies Co., Ltd.
 Comment Type T Comment Status D Inner FEC code(bucket)

The systematic Hamming code is most naturally defined in terms of its parity-check matrix, as pointed out in many textbooks and standard documents. One famous example is the systematic double-extended Hamming(128,119) code in OIF-400ZR and ITU-T G.709.3.

SuggestedRemedy

Suggest to include the construction process and parity-check matrix of the adopted Hamming(68,60) code to enhance the completeness of the document. A Supporting Presentation will be provided.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 The following presentation was reviewed by the 802.3dj task force at the May Interim meeting.
https://www.ieee802.org/3/dj/public/24_05/huang_3dj_01a_2405.pdf
 Implement the suggested remedy with editorial license.

Cl 177 SC 177.4.4 P253 L48 # 612

Huang, Kechao Huawei Technologies Co., Ltd.
 Comment Type T Comment Status D Inner FEC code (bucket)

"The generation matrix G(60,8) for the Hamming(68,60) encoder is given in Table 177—1" is not accurate. The generation matrix for the Hamming(68,60) should be with 60 rows and 68 columns, where the most-left 60 columns is the identity matrix.

SuggestedRemedy

Suggest to change the sentence to "The generator matrix of the Hamming(68,60) code is $G=[I_{60} ; G_{(60 \times 8)}]$, where I_{60} is the 60x60 identity matrix, and $G_{(60 \times 8)}$ is a 60x8 matrix used to generate the 8 parity bits given in Table 177-1."

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 The following presentation was reviewed by the 802.3dj task force at the May Interim meeting.
https://www.ieee802.org/3/dj/public/24_05/huang_3dj_01a_2405.pdf
 Implement the suggested remedy with editorial license.

Cl 177 SC 177.4.6 P254 L # 608

de Koos, Andras Microchip Technology
 Comment Type T Comment Status D pad insertion (bucket)

A figure illustrating the pad bits and their interval for each inner FEC flow would be useful. I always find myself referring to the equivalent RS-FEC Figures (Figure 119û6 and Figure 119û8)

SuggestedRemedy

Consider adding a figure illustrating the pad insertion and interval, in the same style as Figure 119-6

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggest remedy with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 177 SC 177.4.6 P254 L31 # 604

de Koos, Andras Microchip Technology
 Comment Type T Comment Status D timesync(bucket)

Phase of inner FEC pad bits vs outer FEC parity bits:
 - An inaccuracy in the path data delay of up to 12ps due to arbitrary phase between the output FEC parity bits and the inner FEC pad bits of the phase is not accounted for.
 - This arbitrary phase would affect the path data delay values.
 - Almost negligible, if my math is correct.

SuggestedRemedy

3 possible ways to address:
 a. Impose a phase relationship between the RS FEC code word boundaries and the inner FEC pad bits, which would mean large-scale changes to the draft.
 b. Specify (in clause 90, perhaps) that the path data delay contribution through the inner FEC sublayer shall be strictly additive to the path data delay contribution through the PCS and PMA layers.
 c. Ignore. Based on 90A.7, the effect here is small enough to not address specifically. "Whether the potential delay difference between the aggregated delay and the sum of the individual function delays is small enough to satisfy the timing requirements is up to the individual application."
 I prefer option (c). It should not be necessary to add specific text or impose new logical rules to the Inner FEC pad bits to address a potential 12ps path data delay impairment.

Proposed Response Response Status W

PROPOSED REJECT.
 The following related presentation was reviewed by the 802.3dj task force at the May Interim meeting.
https://www.ieee802.org/3/dj/public/24_05/he_3dj_01a_2405.pdf
 It appeared that there was no consensus to make any related changes to the draft.

Cl 177 SC 177.4.6 P254 L33 # 296

Galan, Jose Vicente Maxlinear Inc
 Comment Type T Comment Status D pad insertion (bucket)

It is not declared when the first pad insertion should happen.

SuggestedRemedy

Indicate in the text that the first pad insertion will happen right at the beginning of CWs, same as in the test vectors.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggest remedy with editorial license.

Cl 177 SC 177.4.6 P254 L44 # 489

Slavick, Jeff Broadcom
 Comment Type T Comment Status D pad insertion (bucket)

The last paragraph describing options for how the pad insertion could be done is unnecessary. The requirement that it occurs every 8704 CW and follows the Figure 177-6 is sufficient.

SuggestedRemedy

Remove the last paragraph of 177.4.6

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 177 SC 177.4.6 P254 L44 # 84

Huber, Thomas Nokia
 Comment Type T Comment Status D pad insertion (bucket)

The last paragraph on p254 is not necessary - implementations are always free to do things in different orders, as long as the end result matches the specified behavior.

SuggestedRemedy

Delete the paragraph.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggest remedy with editorial license.

Cl 177 SC 177.4.6.2 P255 L49 # 297

Galan, Jose Vicente Maxlinear Inc
 Comment Type T Comment Status D pad insertion (bucket)

The details of how to use the IBSF are beyond the scope of this standard. Does it mean this is vendor discretionary ? Or will it be defined in other standard ?

SuggestedRemedy

Clarify in the text where the use of the IBSF will be defined.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggest remedy with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 177 SC 177.5.1 P256 L25 # 86

Huber, Thomas

Nokia

Comment Type T Comment Status D Inner FEC Sync (bucket)

This subclause is confusing and seems to be prescribing a specific implementation. The goal of the process is to find codeword boundaries and remove the pad. If we simply reverse the processes of the tx, this process would (in a logical sense) be performed on the interleaved stream, and would search for the (interlaved) FS pattern

SuggestedRemedy

Rewrite the text to describe searching for the FS pattern and finding it at the expected interval

Proposed Response Response Status W

PROPOSED REJECT.

The comment does not provide sufficient justification to support the suggested remedy. The existing text is consistent with the adopted baseline.

Cl 177 SC 177.5.1 P256 L50 # 490

Slavick, Jeff

Broadcom

Comment Type T Comment Status D Inner FEC Sync (bucket)

Monitor and drop says you monitor on all flows. But Figure 177-7 is a per flow state diagram. So is each Flow checking for 140 bad out of 150? And 150 is not a multiple of 8 for it to span across all flows evenly.

SuggestedRemedy

Change:

"keeps monitoring 150 consecutive codewords on all flows, if at least 140 codewords are invalid, drop sync and restart from step a). "

To:

"each flow counts the number of invalid codewords seen in consecutive non-overlapping 150 codeword windows, if at least 140 codewords are invalid, drop sync and restart from step a). "

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the suggest remedy with editorial license.

Cl 177 SC 177.5.1 P257 L1 # 609

de Koos, Andras

Microchip Technology

Comment Type T Comment Status D Inner FEC Sync (bucket)

A figure illustrating the possible one bit-pair of skew and the relationship to the Inner FEC flows would be very helpful here. I only understand because I recall the Task Force presentations!

SuggestedRemedy

Consider adding a figure illustrating how the position of the 1 bit-pair of skew determines the Inner FEC flow number.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the suggest remedy with editorial license.

Cl 177 SC 177.5.3 P257 L29 # 183

Brown, Matt

Alphawave Semi

Comment Type T Comment Status D counters(bucket)

177.5.3 lists a few counter to be supported by the inner FEC. The defintion for some of these could be improved. Further, additional counters should be included provides bins of error counts to help estimate quality of the link.

SuggestedRemedy

A contribution with more details will be provided.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

The following presentation was reviewed by the 802.3dj task force at the May Interim meeting: https://www.ieee802.org/3/dj/public/24_05/brown_3dj_05a_2405.pdf.

Implement slides 7, 9 and 10 with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 177 SC 177.5.3.1 P257 L45 # 493

Slavick, Jeff Broadcom
 Comment Type T Comment Status D Inner FEC decode (bucket)

Defining how a miscorrected codeword can occur could be phrased more clearly.

SuggestedRemedy

Change:
 Note that for soft-decision decoded Inner FEC codewords, when there is more than one bit error in a codeword, there is always a non-zero chance that miscorrection could happen.
 To:
 Note that when there is more than one bit error in a codeword there is a chance that the soft decision decoder could miscorrect the codeword.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 177 SC 177.6.2.1 P258 L52 # 492

Slavick, Jeff Broadcom
 Comment Type T Comment Status D Inner FEC Sync (bucket)

Countes automatically have a _done variable created for them, so no need to define fc_cnt_done

SuggestedRemedy

Remove fc_cnt_done definition

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 177 SC 177.6.2.3 P260 L3 # 176

Ramesh, Sridhar Maxlinear Inc
 Comment Type TR Comment Status D counters(bucket)

Counters defined here do not seem consistent with those defined in Table 177-4.

SuggestedRemedy

Please make definitions of counters consistent with status variables shown on Table 177-4, page 263

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment # 183.

Cl 177 SC 177.6.2.3 P260 L3 # 175

Ramesh, Sridhar Maxlinear Inc
 Comment Type TR Comment Status D counters(bucket)

Add a counter for uncorrectable codewords (detected with additional one bit parity)

SuggestedRemedy

uncorr_cw_cnt
 Countes the number of inner FEC codewords considered uncorrectable by inner FEC decoder

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment # 183.

Cl 177A SC 177A P643 L5 # 306

Maki, Jeffery Juniper Networks
 Comment Type T Comment Status D (bucket)

Annex title unnecessarily uses the acronym IMDD. Not clear what purpose is achieved that cannot be achieved simply by omitting the use of the acronym IMDD.

SuggestedRemedy

Delete the acronym IMDD.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Change title to "Test vectors for 200GBASE-R, 400GBASE-R, 800GBASE-R, and 1.6TBASE-R Inner FEC"

Cl 178 SC 178.1 P268 L45 # 364

Healey, Adam Broadcom Inc.
 Comment Type T Comment Status D (bucket)

The Annex 176A control function is required and should be included in Table 178-1 (as is done in Table 179-1).

SuggestedRemedy

Add "176A - Control" as "Required" in Tables 178-1, 178-2, 178-3, and 178-4.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 178 SC 178.8.9 P275 L33 # 363

Healey, Adam Broadcom Inc.
 Comment Type T Comment Status D (bucket)

The reference to 179.8.9 seems inappropriate here since that subclause contains cross-references specific to the Clause 179.

SuggestedRemedy

Replicate the content of 179.8.9 here, replacing references to Clause 179 electrical requirements to the corresponding references in Clause 178.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 178 SC 178.9.2 P276 L18 # 452

Simms, William NVIDIA
 Comment Type T Comment Status D TX AC CM (bucket)

SCMR may need to be relaxed for 200Gb/s. Measure of 15dB full band at TP0v given full band Vcm noise of 80mVpp at TP2.

SuggestedRemedy

Likely need to tighten 80mV Vcm in table 179-7 for 200Gb/s

Proposed Response Response Status W

PROPOSED REJECT.
 The suggested remedy does not propose an actionable (within the draft) remedy. A question or call to action is not a valid request.

Cl 178 SC 178.9.3.4 P282 L45 # 401

Li, Tobey MediaTek
 Comment Type TR Comment Status D RX ITOL/JTOL (bucket)

"The test channel COM, calculated per items 3) through 7) in 93C.2, is at least 3 dB"

The reference to the test channel COM is wrong.

SuggestedRemedy

Change it to "The test channel COM, calculated peritem e) through h) in 178.9.3.3, is at least 3 dB" to be correct

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 178 SC 178.10 P284 L12 # 34

Mellitz, Richard Samtec
 Comment Type TR Comment Status D Channel lLdd (bucket)

reference is wrong and lLdd should reflect tp0d to tp05d.

SuggestedRemedy

change reference to 178.10.2
 and TBD to 40 dB
 or eliminate the reference to lLdd

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 The objective this clause is addressing is 40 dB die-to-die.
 Change the reference to 178.10.2 and the TBD to 40 dB with additional text to state that it is specified from TP0d to TP5d.
 Implement with editorial license.

Cl 178 SC 178.10.1 P285 L18 # 118

Sakai, Toshiaki Socionext
 Comment Type T Comment Status D COM pkg tau (bucket)

COM reference package parameter vlaue. (transmission line parameter tau)
 In "Table 178ú12" class A package model Transmission line parameter t(tau) value is 6.141e-4 ns/mm, but based on the adopted motion#10, Nov/2024, llim_3dj_01a_2311.pdf (page8-9), the value is 6.141e-3. The value should be 6.141e-3 ns/mm.

SuggestedRemedy

Change t(tau) value in Table 178-12 (class A package) from 6.141e-4 ns/mm to 6.141e-3 ns/mm.
 Or simply delete this row, as the t(tau) value in table 93A-3 is 6.141e-3 ns/mm.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 The value in D1.0 is a typo.
 Change 6.141e-4 to 6.141e-3 in Table 178–12, Table 179–15, and Table 176D–6 (twice in each table).

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 178 SC 178.10.1 P285 L19 # 356
 Healey, Adam Broadcom Inc.
 Comment Type T Comment Status D COM pkg tau (bucket)
 In Table 178-12, the transmission line parameter "tau" is set to 6.141e-4. In the adopted baseline proposal li_3dj_01a_2311 (slides 8 and 9), the value is specified to be 6.141e-3.
 SuggestedRemedy
 Replace the "tau" values in the Table 178-12 with the adopted value 6.141e-3 (2 instances). Similarly in Table 179-15 and Table 176D-6.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #118.

Cl 178 SC 178.10.1 P285 L28 # 119
 Sakai, Toshiaki Socionext
 Comment Type T Comment Status D COM pkg tau (bucket)
 COM reference package parameter vlaue.
 "Table 178-12" class B package model Transmission line parameter t(tau) value is 6.141e-4 ns/mm, but based on the adopted motion#10, Nov/2024, llim_3dj_01a_2311.pdf (page8-9), the value is 6.141e-3. The value should be 6.141e-3 ns/mm.
 SuggestedRemedy
 Change t(tau) value in Table 178-12 (class B package)from 6.141e-4 ns/mm to 6.141e-3 ns/mm.
 Or simply delete this row, as the t(tau) value in table 93A-3 is 6.141e-3 ns/mm.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #118.

Cl 178 SC 178.10.1 P285 L31 # 357
 Healey, Adam Broadcom Inc.
 Comment Type T Comment Status D COM ref pkg (bucket)
 In Table 178-12, the transmission line parameters for the "Class B package model" do not match the adopted baseline proposal li_3dj_01a_2311 slide 9.
 SuggestedRemedy
 Replace the characteristic impedance for stage 1 with 92 Ohms, and the length/characteristic impedances for stage 2 through 4 with 70 Ohms/1 mm, 80 Ohm/1 mm, and 100 Ohm/0.5 mm respectively. Similarly in Table 179-15 and Table 176D-6.
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 178 SC 178.10.2 P287 L37 # 40
 Mellitz, Richard Samtec
 Comment Type TR Comment Status D Channel lLdd (bucket)
 Define the channel insertion loss to include the package i.e TP0d to TP5d.
 SuggestedRemedy
 change TBD to 40 dB
 Proposed Response Response Status W
 PROPOSED REJECT.
 The comment addresses an open TBD, but the lLdd limit in this subclause is expected to be a frequency-dependent mask. The suggested remedy is a single number, which is inadequate.

Cl 178A SC 178A.1.5 P650 L7 # 228
 Noujeim, Leesa Google
 Comment Type T Comment Status D (bucket)
 The port labels on Figure 178A-6 are inconsistent with the cascade order implied in 178A-12 and with the text on line 1.
 SuggestedRemedy
 In Fig 178A-6 replace "Port 2" with "Port 1" and replace "Port 1" with "Port 2"
 Alternatively, replace Figure 178A-6 with a copy of Figure 178A-2 and reverse the arrow directions and swap Port 1 with Port 2.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 The comment correctly points out that port ordering conventions (1 is an input, 2 is an output) should be consistently applied.
 In Figure 178A-6, label the input to the "Host channel (optional)" as "Port 1" and label the output of the "Device termination" as "Port 2".
 Change the last sentence of 178A.1.5 to:
 "The port order of the resulting model is then reversed so that port 1 becomes the input to the optional host channel (or the device package when the host channel is not included) and port 2 becomes the output of the device termination."
 Implement with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 178A SC 178A.1.8 P654 L42 # 209
 Shakiba, Hossein Huawei Technologies Canada
 Comment Type T Comment Status D (bucket)
 Reference to the wrong section 178A.1.6.4
 SuggestedRemedy
 Change reference to section 178A.1.8.1
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 178A SC 178A.1.9 P657 L51 # 210
 Shakiba, Hossein Huawei Technologies Canada
 Comment Type T Comment Status D (bucket)
 h_ISI in equation (178A-29) should not include the main cursor (h_ISI(main) = 0)
 SuggestedRemedy
 Add a case to define h_ISI(n) = 0 for n = d+1
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 178A SC 178A.1.11.1 P660 L52 # 213
 Shakiba, Hossein Huawei Technologies Canada
 Comment Type T Comment Status D MLSD_PDF (bucket)
 Although clear, the result of the PDF convolution $\text{conv}[p(y), p(y/b1)]$ is a PDF and assumed to have been normalized to satisfy the PDF sum requirement.
 SuggestedRemedy
 Either mention that after convolution, the result should be normalized, or add a normalization coefficient of $1/b1$ in font of conv.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 On page 660, line 52, change " $\text{conv}[p(y), p(y/b1)]$ " to " $\text{conv}[p(y), p(y/b1)/|b1|]$ " where $|a|$ is the absolute value of a.
 In Equation (178A-39), change " $p(y/(1-b1))$ " to " $p(y/(1-b1))/|1-b1|$ ".
 Add a note that states that the operation $p(y/a)/|a|$ scales random variable Y by a factor of a, and that the scaled probability distribution function integrates to 1.
 Implement with editorial license.

Cl 178A SC 178A.1.11.1 P661 L1 # 214
 Shakiba, Hossein Huawei Technologies Canada
 Comment Type T Comment Status D MLSD_PDF (bucket)
 Although clear, the result of the PDF convolution of equation (178A-39) is a PDF and assumed to have been normalized to satisfy the PDF sum requirement.
 SuggestedRemedy
 Either mention that after convolution, the result should be normalized, or add a normalization coefficient of $1/(1-b1)$ in font of conv.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #213.

Cl 179 SC 179.9.4 P309 L44 # 511
 Dawe, Piers Nvidia
 Comment Type T Comment Status D TX AC CM (bucket)
 AC common-mode voltages are not as large as this in practice, even at 200G/lane
 SuggestedRemedy
 Reduce both AC common-mode voltage limits for CR, KR, C2C and C2M.
 Proposed Response Response Status W
 PROPOSED REJECT.
 The suggested remedy does not propose an actionable (within the draft) remedy. A question or call to action is not valid.

Cl 179 SC 179.9.4 P309 L46 # 512
 Dawe, Piers Nvidia
 Comment Type TR Comment Status D Tx swing (bucket)
 Supply voltages and voltage swing trend downwards over the years. This 1200 mV max has not changed since 10GBASE-KR, a long time ago. C2M has 750 mV.
 SuggestedRemedy
 Reduce 1200 mV to e.g. 1000 mV, here, in the receiver Table 179-10 and in the text in 179.9.5.2. Reduce the steady-state voltage v_f max from 0.6 V to 0.5 V. Similarly for KR and C2C.
 Proposed Response Response Status W
 PROPOSED REJECT.
 The comment does not provide sufficient justification to support the suggested remedy. Specifically, no issue was identified with allowing a device to have V_{dpp} of 1200 mV.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 179 SC 179.9.4 P310 L27 # 513
 Dawe, Piers Nvidia
 Comment Type TR Comment Status D Tx jitter, Tx SNDR (bucket)
 Our way of measuring jitter doesn't work well enough with the increased max host loss over 3ck. It is not clear that it can or should be fixed. Our way of defining SNDR doesn't work correctly over host loss either. This can be fixed, but "vertical and horizontal noise" act together to degrade BER: more of one goes with less of the other.
SuggestedRemedy
 Delete the SNDR and jitter specs. Add a VEC-like, TDECQ-like spec using this clause's COM reference receiver which can be implemented in a scope. Similarly for KR and C2C.
Proposed Response Response Status W
 PROPOSED REJECT.
 The suggested remedy does not propose an actionable (within the draft) remedy. A question or call to action is not valid.
 In addition, the comment includes a claim that measurements are not feasible, which is not substantiated and is contrasted by existing contributions, e.g. https://www.ieee802.org/3/dj/public/adhoc/electrical/24_0104/calvin_3dj_elec_01a_240104.pdf.

Cl 179 SC 179.9.4.6 P315 L15 # 514
 Dawe, Piers Nvidia
 Comment Type TR Comment Status D Tx jitter, Tx SNDR (bucket)
 As explained in other comments, up to 3ck the SNDR spec acted together with the jitter spec to protect the link performance - but we don't have a satisfactory way of measuring jitter at today's speeds and losses, and separating the two things out "leaves margin on the table".
SuggestedRemedy
 Delete the SNDR section. Add a VEC-like, TDECQ-like spec using this clause's COM reference receiver which can be implemented in a scope. Similarly for KR and C2C.
Proposed Response Response Status W
 PROPOSED REJECT.
 The suggested remedy does not propose an actionable (within the draft) remedy. A question or call to action is not valid.
 In addition, the comment includes a claim that measurements are not feasible, which is not substantiated and is contrasted by existing contributions, e.g. https://www.ieee802.org/3/dj/public/adhoc/electrical/24_0104/calvin_3dj_elec_01a_240104.pdf.

Cl 179 SC 179.9.4.7 P315 L24 # 515
 Dawe, Piers Nvidia
 Comment Type TR Comment Status D Tx jitter (bucket)
 Measuring jitter separately to other impairments relies on a better slew rate to noise ratio than we have at the observation point, and better than what is needed to make good links.
SuggestedRemedy
 Delete the jitter section. Add a VEC-like, TDECQ-like spec using this clause's COM reference receiver which can be implemented in a scope. Similarly for KR and C2C.
Proposed Response Response Status W
 PROPOSED REJECT.
 The suggested remedy does not propose an actionable (within the draft) remedy. A question or call to action is not valid.
 In addition, the comment includes a claim that measurements are not feasible, which is not substantiated and is contrasted by existing contributions, e.g. https://www.ieee802.org/3/dj/public/adhoc/electrical/24_0104/calvin_3dj_elec_01a_240104.pdf.
 Note that the importance of controlling jitter separately from other impairment has been addressed in https://www.ieee802.org/3/dj/public/24_05/ran_3dj_03_2405.pdf.

Cl 179 SC 179.11.1 P326 L27 # 389
 Kocsis, Sam Amphenol
 Comment Type T Comment Status D Nominal impedance (bucket)
 Nominal characteristic impedance of the cable assembly is "100-ohm"
SuggestedRemedy
 Contributions to the task force have demonstrated the nominal characteristic impedance of the cable assembly is ~92-ohm
Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 It is understood that the suggested remedy is to change the nominal impedance from 100 to 92 Ohm.
 However, as noted in comment #216, there is no need to specify a nominal impedance. Resolve with using the response to comment #216.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 179 SC 179.11.1 P326 L27 # 216

Noujeim, Leesa Google
 Comment Type T Comment Status D Nominal impedance (bucket)

There is no test method or definition for the nominal characteristic impedance of the cable assembly. The components (eg paddle card, twinax) within a cable assembly may have different nominal characteristic impedances. There is no need to specify the nominal characteristic impedance of the cable assembly, since the performance of the cable assembly is determined by cl 179.11.2-7.

SuggestedRemedy

Remove "The nominal characteristic impedance of the cable assembly is 100 ohms"

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 It is important to define the reference impedance for return loss specifications etc., but as the comment correctly suggests, there is no need to specify a nominal value.
 Implement the suggested remedy.

Cl 179 SC 179.11.1 P326 L27 # 516

Dawe, Piers Nvidia
 Comment Type T Comment Status D Nominal impedance (bucket)

"Nominal impedance" is something for a datasheet not a spec. If someone wants to build a cable assembly with 95 ohm bulk cable and it passes the spec - that's OK.

SuggestedRemedy

Delete "The nominal differential characteristic impedance of the cable assembly is 100 [ohm]". Move the one remaining sentence into 179.11.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #216.

Cl 179 SC 179.11.3 P327 L34 # 390

Kocsis, Sam Amphenol
 Comment Type T Comment Status D ERL (bucket)

ERL requirement for cable assemblie sthat have COM less than "4dB"

SuggestedRemedy

Change "4dB" to "TBD". Historical precedent may not be relevant for this specification

Proposed Response Response Status W

PROPOSED REJECT.
 The comment does not provide sufficient justification to support the suggested remedy.
 Note that any content of the draft can be changed if there is consensus, but changing from a number to TBD does not move us forward.

Cl 179 SC 179.11.7 P331 L18 # 120

Sakai, Toshiaki Socionext
 Comment Type T Comment Status D COM pkg tau (bucket)

COM reference package parameter vlaue. (transmission line parameter tau)
 In "Table 179û15" class A package model Transmission line parameter t(tau) value is 6.141e-4 ns/mm, but based on the adopted motion#10, Nov/2024, (Ilim_3dj_01a_2311.pdf (page8-9), the value is 6.141e-3. The value should be 6.141e-3 ns/mm.

SuggestedRemedy

Change t(tau) value in Table 179-15 (class A package) from 6.141e-4 ns/mm to 6.141e-3 ns/mm.
 Or simply delete this row, as the t(tau) value in table 93A-3 is 6.141e-3 ns/mm.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #118.

Cl 179 SC 179.11.7 P331 L28 # 121

Sakai, Toshiaki Socionext
 Comment Type T Comment Status D COM pkg tau (bucket)

COM reference package parameter vlaue. (transmission line parameter tau)
 In "Table 179û15" class B package model Transmission line parameter t(tau) value is 6.141e-4 ns/mm, but based on the adopted motion#10, Nov/2024, (Ilim_3dj_01a_2311.pdf (page8-9), the value is 6.141e-3. The value should be 6.141e-3 ns/mm.

SuggestedRemedy

Change t(tau) value in Table 179-15 (class B package) from 6.141e-4 ns/mm to 6.141e-3 ns/mm.
 Or simply delete this row, as the t(tau) value in table 93A-3 is 6.141e-3 ns/mm.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #118.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 179A SC 179A.2 P662 L 6710 # 56
 Mellitz, Richard Samtec
Comment Type TR Comment Status D 93B (bucket)
 Reference to a diagram with TP0d and TP5d is required
SuggestedRemedy
 Add TP0d and TP5d to figure 93B-1 and table 93B-1
Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Annex 93B is irrelevant for CR.
 Also, Annex 93B is not referenced anywhere in the draft, nor in previous backplane PMD clauses 163 and 137.
 A diagram with the new test points exists in Figure 179-2 and can be referenced instead.
 Add a reference in 179A.2 to Figure 179-2. Implement with editorial license.

Cl 179A SC 179A.5 P665 L 24 # 229
 Noujeim, Leesa Google
Comment Type T Comment Status D Channel ILdd (bucket)
 Doubling ILdd_(host+TFmax) implies both ends of the link have the same host designations.
SuggestedRemedy
 Replace "2*ILdd_(host+TFmax)" with "ILdd_(host+tFmax)_end1 + ILdd_(host+tFmax)_end2" or similar notation to accommodate asymmetric Link Configurations in Table 179A-3.
Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Replace "2*ILdd_(host+TFmax)" with "ILdd_(host+tFmax)_one end + ILdd_(host+tFmax)_other end" with editorial license to accommodate asymmetric Link Configurations in Table 179A-3.

Cl 179B SC 179B.1 P669 L 15 # 222
 Noujeim, Leesa Google
Comment Type T Comment Status D (bucket)
 Incorrect Annex reference 120G
SuggestedRemedy
 Replace 120G with 176E
Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 179B SC 179B.1 P669 L 17 # 223
 Noujeim, Leesa Google
Comment Type T Comment Status D HCB and MCB (bucket)
 Missing reference to Module compliance at TP1 and TP4
SuggestedRemedy
 Add "Module measurements for Modules specified in Annex 176E are made at TP1 and TP4 with test fixtures as specified in 179B.3. "
Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Insert the sentence:
 Module measurements for modules specified in Annex 176E are made at module compliance points TP1 and TP4 (see Figure 176E-4) with test fixtures as specified in 179B.3.

Cl 179B SC 179B.4.6 P676 L 26 # 224
 Noujeim, Leesa Google
Comment Type T Comment Status D HCB and MCB (bucket)
 SFPxxx is unclear
SuggestedRemedy
 Replace "The SFPxxx mated test fixture" with "The single-lane mated test fixture"
Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 In 179B replace SFPxxx with SFP112

Cl 179B SC 179B.4.26 P676 L 41 # 59
 Mellitz, Richard Samtec
Comment Type TR Comment Status D HCB and MCB (bucket)
 At least the symbol rate is known
SuggestedRemedy
 set fb to 106.25 GBd
Proposed Response Response Status W
 PROPOSED ACCEPT.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 179C SC 179C.1 P680 L15 # 525

Dawe, Piers Nvidia
 Comment Type T Comment Status D MDI references (bucket)

MDIs are mechanical entities. For 106.25 Gb/s operation, there are SFP2 (SFF-TA-1031) and QSFP2 (SFF-TA-1027). Any "SFP224" would be an SFP2 module or cable end with 200G-capable circuitry. But this annex is for the MDI, not the circuitry. Similarly for "QSFP224" and QSFP2.

SuggestedRemedy

Correct the names. Add references to SFF-TA-1011 which relates the names and specs for the SNIA-SFF modules, and SFF-8665, which defines the components of a QSFPx "solution".

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 There was broad consensus to use names of MDI types (part of baseline proposal) currently in the draft as follows: SFP224, SFP-DD224, QSFP224, QSFP-DD1600, QSFP1600.
 Resolve using the response to comment #506, which addresses the normative references.

Cl 179C SC 179C.1 P680 L17 # 526

Dawe, Piers Nvidia
 Comment Type TR Comment Status D MDI references (bucket)

Refer to the specification for each connector type where each is first mentioned. See another comment against 1.3 for the reference docs.

SuggestedRemedy

Per comment

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #506.

Cl 179C SC 179C.2.3 P688 L35 # 527

Dawe, Piers Nvidia
 Comment Type T Comment Status D MDI references (bucket)

This says "the mechanical interface". The mechanical spec is SFF-TA-1027, QSFP2. It is a standard, not an MSA.

SuggestedRemedy

Change " the TBD MSA" to "SFF-TA-1027".

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #506.

Cl 179C SC 179C.2.4 P689 L35 # 528

Dawe, Piers Nvidia
 Comment Type T Comment Status D MDI references (bucket)

There is no QSFP-DD1600 TBD MSA document. QSFP-DD1600 is defined in the singular QSFP-DD MSA document

SuggestedRemedy

Change "the QSFP-DD1600 TBD MSA" to "the QSFP-DD/QSFP-DD800/QSFP-DD1600 Hardware Specification".

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #506.

Cl 179C SC 179C.2.5 P690 L21 # 529

Dawe, Piers Nvidia
 Comment Type T Comment Status D MDI references (bucket)

There is no OSFP1600 TBD MSA document. OSFP1600 is defined in the singular OSFP MSA document, particularly section 4.

SuggestedRemedy

Change "the OSFP1600 TBD MSA" to "the OSFP Octal Small Form Factor Pluggable Module specification" or "section 4 of the OSFP Octal Small Form Factor Pluggable Module specification".

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #506.

Cl 180 SC 180.4.1 P350 L13 # 160

Yu, Rang-chen InnoLight
 Comment Type ER Comment Status D Editorial (bucket)

A typo of 'L3' in figure 180-2, right side, 3rd channel output label.

SuggestedRemedy

It should be 'L2'.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement with editorial license and discretion.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 180 SC 180.10 P368 L11 # 521
 Dawe, Piers Nvidia
 Comment Type T Comment Status D bit number (bucket)
 Bit number should match number of lanes
 SuggestedRemedy
 Change 1.9.4 to 1.9.n. Below, change 1.10.4 to 1.10.n. Similarly in other clauses.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 181 SC 181.1 P372 L16 # 4
 Johnson, John Broadcom
 Comment Type T Comment Status D Editorial (bucket)
 The PHY bracket in Figure 181-1 is shown encompassing the MDI layer, which isn't consistent with previous PMDs.
 SuggestedRemedy
 Shorten the PHY bracket to exclude the MDI layer.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 181 SC 181.8.5 P386 L41 # 2
 Johnson, John Broadcom
 Comment Type T Comment Status D Reference (bucket)
 The TDECQ methods reference channel requirements in 121.8.5.2 instead of the channel requirements in local clause 181.8.5.1.
 SuggestedRemedy
 Replace the reference to 121.8.5.2 with reference to 181.8.5.1.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 182 SC 182.1 P392 L44 # 301
 Maki, Jeffery Juniper Networks
 Comment Type TR Comment Status D IMDD acronym (bucket)
 Associated clause description is malformed. The acronym IMDD is used, which does not appear in the actual Clause 177 title. Why preclude that at some future point in time that Clause 177 is used for something other than IMDD? Also, there is no use of "Coherent" to describe Inner FECs used for coherent PMDs to setup the appropriate parallelism of terminology.
 SuggestedRemedy
 Delete the acronym IMDD.

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 182 SC 182.1 P393 L29 # 302
 Maki, Jeffery Juniper Networks
 Comment Type TR Comment Status D IMDD acronym (bucket)
 Associated clause description is malformed. The acronym IMDD is used, which does not appear in the actual Clause 177 title. Why preclude that at some future point in time that Clause 177 is used for something other than IMDD? Also, there is no use of "Coherent" to describe Inner FECs used for coherent PMDs to setup the appropriate parallelism of terminology.

SuggestedRemedy
 Delete the acronym IMDD.

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 182 SC 182.1 P394 L23 # 303
 Maki, Jeffery Juniper Networks
 Comment Type TR Comment Status D IMDD acronym (bucket)
 Associated clause description is malformed. The acronym IMDD is used, which does not appear in the actual Clause 177 title. Why preclude that at some future point in time that Clause 177 is used for something other than IMDD? Also, there is no use of "Coherent" to describe Inner FECs used for coherent PMDs to setup the appropriate parallelism of terminology.

SuggestedRemedy
 Delete the acronym IMDD.

Proposed Response Response Status W
 PROPOSED ACCEPT.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 182 SC 182.1 P394 L 50 # 304
 Maki, Jeffery Juniper Networks
 Comment Type TR Comment Status D IMDD acronym (bucket)
 Associated clause description is malformed. The acronym IMDD is used, which does not appear in the actual Clause 177 title. Why preclude that at some future point in time that Clause 177 is used for something other than IMDD? Also, there is no use of "Coherent" to describe Inner FECs used for coherent PMDs to setup the appropriate parallelism of terminology.
 SuggestedRemedy
 Delete the acronym IMDD.
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 182 SC 182.1 P395 L 21 # 5
 Johnson, John Broadcom
 Comment Type T Comment Status D Editorial (bucket)
 The PHY bracket in Figure 182-1 does not encompass the PMD layer, which isn't consistent with previous PMDs.
 SuggestedRemedy
 Lengthen the PHY bracket to include the PMD layer.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 183 SC 183.1 P418 L 39 # 305
 Maki, Jeffery Juniper Networks
 Comment Type TR Comment Status D IMDD acronym (bucket)
 Associated clause description is malformed. The acronym IMDD is used, which does not appear in the actual Clause 177 title. Why preclude that at some future point in time that Clause 177 is used for something other than IMDD? Also, there is no use of "Coherent" to describe Inner FECs used for coherent PMDs to setup the appropriate parallelism of terminology.
 SuggestedRemedy
 Delete the acronym IMDD.
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 184 SC 184.1.1 P441 L 8 # 308
 Bruckman, Leon Huawei
 Comment Type TR Comment Status D General (Bucket)
 The Inner FEC as defined, includes the PMA. Shall make this clear to the reader
 SuggestedRemedy
 Either add sentence: "This Inner FEC sublayer includes functionality often associated with the PMA sublayer", or split the PMA function
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the following with editorial license.
 Add sentence: "This Inner FEC sublayer includes functionality often associated with the PMA sublayer at the PMD service interface".
 Add similar text to the appropriate sub clause in clause 177
 [Editor's note: CC 184, 177]

Cl 184 SC 184.2 P443 L 7 # 87
 Huber, Thomas Nokia
 Comment Type T Comment Status D General (Bucket)
 Other diagrams of this type do not have dashed boxes around the transmit and received processes.
 SuggestedRemedy
 For consistency with the rest of the document, remove the dashed boxes
 Proposed Response Response Status W
 PROPOSED REJECT.
 The dashed boxes clearly denote the transmit and receive functions. Removing the dashed boxes does not improve clarity of the draft.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 184 SC 184.2 P444 L5 # 88

Huber, Thomas

Nokia

Comment Type T Comment Status D Functional (Bucket)

The second sentence of the paragraph (discussing the distribution to 32 lanes by the permutation function) seems to imply that the 32 lanes were interleaved into a serial stream after they were reordered and deskewed, but the text doesn't actually say that is done.

SuggestedRemedy

If the intent is that the 32 lanes are re-interleaved, and then the permutation function distributes the symbols back to 32 lanes (in something other than a round-robin manner), change the end of the first sentence to say "reordered, deskewed, and serialized". If the intent is that the permutation process just moves symbols around among the 32 lanes, change the second sentence to say "The RS-FEC symbols are then rearranged across the 32 lanes by a permutation function."

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
Implement the following with editorial license.
Change "The RS-FEC symbols are then distributed over the 32 lanes by a permutation function." to "The RS-FEC symbols are then rearranged across the 32 lanes by a permutation function."

Cl 184 SC 184.4 P445 L22 # 184

Brown, Matt

Alphawave Semi

Comment Type T Comment Status D Reorder (Bucket)

The Inner FEC transmit (184.4) and receive (184.5) functions provide a BCH encoder/decoder and other functions to be performed on each PCS lane. Although there is one per PCS lane, these should be called "flows" rather than "lanes" to be consistent with other FEC clauses and to differentiate between "lanes" that go between sublayers.

SuggestedRemedy

When describing the process applied to each PCS lane in each direction, use the word "flow" rather than "lane".

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
Implement the suggested remedy with editorial license.

Cl 184 SC 184.4.1 P445 L3 # 299

Loewenthal, Arnon

alphawave semi

Comment Type T Comment Status D Functional (bucket)

Need to further define the deskew requirement. For now it is defined as optional. In practice full deskew is optional, but doing 10b alignment of RS symbols is mandatory.

SuggestedRemedy

Replace lines 8-18 with the requirement of partial deskew, which means 10b RS symbols resolution deskew.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
Implement the following with editorial license.
In the first paragraph of clause 184.4.1 delete ", when implemented," and delete the second paragraph

Cl 184 SC 184.4.1 P445 L5 # 89

Huber, Thomas

Nokia

Comment Type T Comment Status D Functional (bucket)

There are always many implementation options, but we don't have to describe them in the document, we just have to describe the behavior that is required.

SuggestedRemedy

Delete "when implemented" from the first sentence, and delete the second paragraph.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
Comment #299 response implements suggested remedy.
Resolve using the response to comment #299

Cl 184 SC 184.4.1 P445 L12 # 90

Huber, Thomas

Nokia

Comment Type T Comment Status D Functional (Bucket)

What is the purpose of this mapping? There are 32 lanes being received; this process is simply aligning them based on the RS FEC frame, so it doesn't seem like a mapping is needed.

SuggestedRemedy

Either explain why this mapping process is needed, or delete it.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
Add text to explain the purpose of this mapping.
Implement with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 184 SC 184.4.1 P445 L12 # 178

Brown, Matt Alphawave Semi

Comment Type T Comment Status D Functional (Bucket)

The process provided in 184.4.1 "Alignment lock and deskew" merely maps bits on the FEC service interface to vectors; it does not include and RS-FEC symbol alignment. The process in 184.4.2 remaps the vectors such that there is alignment to the RS-FEC symbols and the lanes are properly ordered.

SuggestedRemedy

Either combine the two subclauses and process into one subclause or move the RS-FEC symbol alignment process in 184.4.2 to 184.4.1.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the following with editorial license.

Move the RS-FEC symbol alignment process in 184.4.2 to 184.4.1.

Cl 184 SC 184.4.2 P445 L19 # 300

Loewenthal, Arnon alphawave semi

Comment Type T Comment Status D Reorder. (Bucket)

Need to further define the lanes reorder requirement. For now it is defined as optional. In practice full lanes reorder is optional, but partial reorder, meaning having flow-0 on lanes 0-15 and flow-1 on lanes 16-31 is required. Not doing that would impact end to end FEC performance and margins.

SuggestedRemedy

Two options:

1. remove the word 'optional' from line 22.
2. Define the restriction of having flow-0 on lanes 0-15 and flow-1 on lanes 16-31.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement the following with editorial license.

Change: "If that is the case, the optional lane reorder function shall order the PCS lanes according to the PCS lane number." to: "The lane reorder function shall order the PCS lanes according to the PCS lane number."

Cl 184 SC 184.4.2 P445 L22 # 91

Huber, Thomas Nokia

Comment Type T Comment Status D Reorder (Bucket)

Lane reordering is not optional; the lanes have to be put in the correct order. If they happen to arrive in the correct order, it's a simple process.

SuggestedRemedy

Change the second sentence to say "The lane reorder process shall order the PCS lanes according to the PCS lane number."

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Resolve using the response to comment #300

Cl 184 SC 184.4.2 P445 L22 # 179

Brown, Matt Alphawave Semi

Comment Type T Comment Status D Reorder (Bucket)

The lane reorder process is stated as being optional, however, that is not the case. It is not required (or optional) if the lanes are already in order (e.g., connected to a PCS above) and mandatory if the lanes may not be in order (e.g., connected to an 8:32 PMA above), thus it is conditional, rather than optional.

SuggestedRemedy

Change the first 2 sentences in 184.4.2 to "If the sublayer above the Inner FEC does not provide the PCS lanes in order at the service interface, the lane reorder function shall reorder the PCS lanes according to the PCS lane number."

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 184 SC 184.4.2 P445 L26 # 92

Huber, Thomas Nokia

Comment Type T Comment Status D Reorder (Bucket)

It is not clear why this description is needed. Other clauses about reordering don't have this.

SuggestedRemedy

Delete the last paragraph

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Resolve using the response to comment #178

Cl 184 SC 184.4.3 P446 L1 # 93
 Huber, Thomas Nokia
 Comment Type T Comment Status D Reorder (Bucket)

This figure is not clear, nor is the relationship of the figure to the pseudocode beneath it. I think the columns 0-3 are just numbers that relate to the post-FEC distribution process. I have no idea why there are 32 sets of 4 symbols, as the algorithm doesn't do anything on a four-symbol basis. The function is simply reversing flow1 and flow0 every two columns, so that each lane has interleaved symbols from all four codewords. This could be described more simply by using blocks of 16 symbols in the figure (i.e., block 0 would be lanes 0-15 in column 0, block 1 would be lanes 16-31 in column 0, etc.).

SuggestedRemedy

Revise the figure as suggested. The input side would look like this (where each row here is corresponding to 16 PCS lanes in the figure):

```
0 2 4 6
1 3 5 7
and the output would be
0 2 5 7
1 3 4 6
```

This will remove any confusion about whether the 32 blocks are supposed to be somehow related to the 32 PCS lanes, and it will be easier to see what is changing between the figures.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 184 SC 184.4.3 P446 L45 # 94
 Huber, Thomas Nokia
 Comment Type T Comment Status D Algorithm (Bucket)

The algorithm is unnecessarily complex. There is no need for bit-level detail since the operation is performed on 10-bit symbols - though really it seems to be performed on 160-bit entities. Per figure 184-3, it's essentially receiving as input alternating sets of 160 bits from flow0 and flow1, and changing the order from 0, 1, 0, 1, 0, 1, 0, 1 to 0, 1, 0, 1, 1, 0, 1, 0.

SuggestedRemedy

A minimal change would be to state that the algorithm operates on 10-bit symbols, delete the for j loop and its terminator, and replace "10i+j" with "i" in the statement that describes the permutation..

Another option would be to rewrite the description around the 160-bit entities as described, and perhaps also change the figure to show those instead of 40-bit entities (which as noted in a previous comment seem to have no relevance to this process, or to the convolutional interleaver process that follows it).

Proposed Response Response Status W

PROPOSED REJECT.
 The algorithm is correct (and explicit) as written. This bit-wise mapping shows explicitly how the bits are mapped into the larger vector.
 Removing j does not seem to add clarity, better have the detailed function as described in the adopted baseline

Cl 184 SC 184.4.4 P447 L22 # 95
 Huber, Thomas Nokia
 Comment Type T Comment Status D Algorithm (Bucket)

The description of the convolutional interleaver process could be improved. The variable i is used in the first part of the subclause as an index for the delay lines and as an indication of time within a sequence. Then at the bottom of page 447 it's used as a symbol index.

SuggestedRemedy

Revise the list above the figure to read as follows, eliminating the overloading of the index i and improving the clarity a bit (and change the figure to label the lines as b=0, b=1, b=2)::

- The input and output switches are always aligned to the same row b, where b = 0 to 2
- a block of 40 bits is read from row b
- The contents of row b are shifted to the right by 40 bits
- A block of 40 bits is written to row b
- The switch position is updated to (b+1) mod 3

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement suggested remedy with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 184 SC 184.4.4 P447 L48 # 96
 Huber, Thomas Nokia
 Comment Type T Comment Status D Algorithm (Bucket)
 Since the convolutional interleaver operates separately on each PCS lane, there's no value in having an algorithm that includes the PCS lanes. Since it operates on 40-bit units, there's also no need to include bit-level description.
SuggestedRemedy
 State that the algorithm describes the operation on the 40 bit entities and is run on each PCS lane independently. This allows elimination of the p and j variables.
 Proposed Response Response Status W
 PROPOSED REJECT.
 This is correct as written.
 Removing the lanes and bits does not seem to add clarity, better have the detailed function as described in the adopted baseline.

Cl 184 SC 184.4.5 P448 L12 # 98
 Huber, Thomas Nokia
 Comment Type T Comment Status D Algorithm (Bucket)
 The first statement should not be a 'shall' (which indicates a PICS item of conformance). The second sentence is correct, in that there are 32 encoders, but what's actually required is that each lane has an encoder.
SuggestedRemedy
 Revise the paragraph to read: The BCH encoder works in conjunction with the RS(544,514) FEC to increase the FEC coding gain. There is a BCH encoder process for each PCS lane.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement the following with editorial license .
 Change: "The BCH encoder shall work in conjunction with the outer RS(544,514) FEC to provide a high-performance FEC for 800GBASE-LR1. There are 32 BCH encoder functions." to: "The BCH encoder works in conjunction with the outer RS(544,514) FEC to provide a high-performance FEC for 800GBASE-LR1. The Inner FEC shall implement 32 BCH encoder functions."

Cl 184 SC 184.4.5 P448 L40 # 99
 Huber, Thomas Nokia
 Comment Type T Comment Status D Algorithm (Bucket)
 The variable p is being overloaded - it is used at line 35 as a lane index, and at line 40 as the parity polynomial. Since the BCH encoding is done per lane, there is really no need to have a variable related to the lane number. The text can simply state that the algorithm is applied to each lane individually.
SuggestedRemedy
 Change the line above the dashed list to say "The BCH encoding is done separately on each lane. The encoding of of each BCH codeword u is defined as follows:
 At the top of page 449, remove the 'for pà' loop from the pseudocode.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Removing the lane does not seem to add clarity, better have the detailed function as described in the accepted baseline.
 Change the flow index from p to q to remove p overload.

Cl 184 SC 184.4.6 P449 L16 # 100
 Huber, Thomas Nokia
 Comment Type T Comment Status D Algorithm (Bucket)
 Clarify that the circular shift is applied per lane.
SuggestedRemedy
 Make similar changes to what was suggested in previous sections - remove the unnecessary variable p and associated for loop in the pseudocode, and add a sentence stating that the circular shift process is performed on each lane individually.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Removing the lane does not seem to add clarity, better have the detailed function as described in the accepted baseline.
 Add a sentence stating that the circular shift process is performed on each flow individually. Implement with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 184 SC 184.4.7.1 P450 L12 # 101
 Huber, Thomas Nokia
 Comment Type T Comment Status D Order (Bucket)
 The DSP frame should probably be a level 3 clause of its own, rather than a sub-clause under BCH interleaver.
 SuggestedRemedy
 Change to a level 3 heading
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 The "BCH interleaver" function includes the pilot insertion. Change clause 184.4.7 title to: BCH interleaver and pilot insertion"
 Implement with editorial license.

Cl 184 SC 184.4.7.1 P450 L14 # 371
 He, Xiang Huawei
 Comment Type TR Comment Status D DSP (Bucket)
 It is said " 4-bit pilot symbols (PS) are inserted every 64 4-bit blocks (one 4-bit PS, 63 4-bit message blocks)."
 But in Figure 184-5, message blocks m<0:63>, m<64-127>, between pilot symbols has 64 4-bit blocks.
 SuggestedRemedy
 Change Figure to match the text, i.e., change m<0:63> to m<0:62>, change m<64:127> to m<63:125>, etc.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Implement suggested remedy with editorial license.

Cl 184 SC 184.4.7.1 P450 L18 # 102
 Huber, Thomas Nokia
 Comment Type T Comment Status D DSP (Bucket)
 The first sentence of the second paragraph could be written more clearly.
 SuggestedRemedy
 Replace with "Two streams of DSP frames, one for each polarization, are generated by the inner FEC."
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 184 SC 184.4.7.2 P450 L45 # 103
 Huber, Thomas Nokia
 Comment Type T Comment Status D DSP (Bucket)
 It is not clear what "192 bits that are complemented with zeros" is intended to mean. Based on what is in Table 184-2, I think the intent is that a zero is inserted after each bit of the PRBS9 output to form the bit-pairs that become the PS symbols. Also, the text talks about 4-bit PS symbols, but Table 184-2 is showing bit-pairs for each component rather than 4-bit symbols without explaining that outputs 0 and 1 are for the X polarization (so the X PRBS is spread across outputs 0 and 1) and outputs 2 and 3 are for the Y polarization.
 SuggestedRemedy

Revise the two paragraphs above table 184-1 to read as follows:
 For both DSP frame_0 and DSP frame_1, the generator is initialized using the seed at the start of every DSP frame. The generator produces a sequence of 192 bits. A zero bit inserted after each bit to generate the bit-pairs that form the pilot symbols, which use the outer points of the 16QAM constellation.

The generator polynomial and seed values are shown in Figure 184-6 and listed in Table 184-1. The complete pilot sequence is shown in Table 184-2. The bit-pairs for the X polarization are distributed in a round-robin manner to outputs 0 and 1. The bit-pairs for the Y polarization are distributed in a round-robin manner to outputs 2 and 3.

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 184 SC 184.4.9 P452 L50 # 104
 Huber, Thomas Nokia
 Comment Type T Comment Status D Interface (Bucket)
 The editor's note suggesting that the mapping to analog signals probably belongs in the PMD clause seems to make sense, in which case this clause is really not "DP-16QAM mapping", it's really just mapping to 4-level signals, which the PMD will then turn into DP-16QAM.
 SuggestedRemedy
 Change the title to "4-level signal mapper", and make the corresponding change in 184.5.3.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 After the first sentence of subclause 184.4.9 add: "This four-level signals are used by the 800GBASE-LR1 PMD to generate a single optical DP-16QAM signal with orthogonal polarizations (see 185.4.2)."
 Implement with editorial license.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 184 SC 184.4.9 P452 L50 # 105

Huber, Thomas Nokia
 Comment Type T Comment Status D Order (Bucket)

The overall flow would be improved if it went BCH interleaver, 4-level signal mapping, DSP frame, with all the pilot symbol details then in the DSP frame clause.

SuggestedRemedy

Revise so the flow is like this:
 184.4.7 BCH interleaver
 184.4.8 Four-level signal mapping (current 184.4.9, without subclauses)
 184.4.9 DSP frame generation (current 184.4.7.1)
 184.4.9.1 Pilot sequence (current 184.4.7.2 and 184.4.9.1)

Proposed Response Response Status W

PROPOSED REJECT.
 The text is correct as written.
 The actual order is the right one. It describes the bit blocks generation and handling, then the mapping to four levels.

Cl 184 SC 184.5.1 P455 L42 # 106

Huber, Thomas Nokia
 Comment Type T Comment Status D Interface (Bucket)

The paragraph that begins with "the signals Rx_Xi, Rx_XQ, à" doesn't seem to make sense. The Tx and Rx signals are not guaranteed to be the same (i.e., Tx_XI can be received as any of the four components), but the contents of Tx_XI aren't distributed to all the Rx signals.

SuggestedRemedy

Revise to say: The signals Rx_XI, Rx_XQ, Rx_YI, and Rx_YQ each represent one of the corresponding Tx_XI, Tx_XQ, Tx_YI, Tx_YQ signals from the transmitting PMD. The association between Tx and Rx components is arbitrary (e.g., Rx_XI can be any of the 4 Tx components).

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 184 SC 184.5.8 P457 L45 # 107

Huber, Thomas Nokia
 Comment Type T Comment Status D Algorithm (Bucket)

Similar changes should be made in the convolutional de-interleaver as were requested for the convolutional interleaver in earlier comments

SuggestedRemedy

Revise the items in the lettered list and the algorithm to align with whatever changes are agreed for the convolutional interleaver.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 Implement suggested remedy with editorial license.

Cl 184 SC 184.6.5 P462 L3 # 307

Bruckman, Leon Huawei
 Comment Type TR Comment Status D Diagrams (Bucket)

Set TBD values of N and M

SuggestedRemedy

Set N=12, M=8. See contribution bruckman_3dj_01_241205

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 The following presentation (referenced in the suggested remedy) was reviewed by the 802.3dj task force at the May Interim meeting:
https://www.ieee802.org/3/dj/public/24_05/bruckman_3dj_01a_2405.pdf
 Implement the suggested remedy with editorial license.

Cl 184 SC 184.6.5 P462 L9 # 559

Law, David HPE
 Comment Type T Comment Status D Diagrams (Bucket)

The LOCK_INIT state in Figure 184ù9 'DSP lock state diagram' includes the action 'test_sym <= false', however the test_sym variable isn't defined in subclause 184.6.2 'Variables' and isn't used anywhere else in Figure 184ù9.

á
 It seems that this should have been 'test_ps <= false' as the test_ps variable isn't initialised during reset in the LOCK_INIT state but used to control the GET_SYMBOL to FIND_1ST transition below.

SuggestedRemedy

Change 'test_sym <= false' to read 'test_ps <= false'.

Proposed Response Response Status W

PROPOSED ACCEPT.

EEE P802.3dj D1.0 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet 1st Task Force review comment

Cl 185 SC 185.1 P468 L19 # 323

D'Ambrosia, John Futurewei, U.S. Subsidiary of Huawei

Comment Type TR Comment Status D Conditional PMA (bucket)

Table 185-1, Figure 185-1, Figure 185-2 does not reflect the PHY type and clause correlation in Table 169-3a. There is no mention of 800GBASE-R BM-PMA, 800GAUI-8 2C2, 800GAUI-8 C2M, 800GBASE SM-PMA, 800GAUI-4 C2C, and 800GAUI-4 C2M.

Baseline Proposal in https://www.ieee802.org/3/dj/public/23_07/kota_3dj_01a_2307.pdf shows support for 800GAUI's.

SuggestedRemedy

Clause 185 needs to be updated to reflect these layers.

Table185-1needs the following entries -

- 800GBASE-R BM-PMA - conditional
- 800GAUI-8 2C2 - optional
- 800GAUI-8 C2M - optional
- 800GBASE SM-PMA - conditional
- 800GAUI-4 C2C - optional
- 800GAUI-4 C2M - optional

Add note "C= Conditional, 800GBASE-R BM-PMA is conditional, pending implementation of 800GAUI-8 C2C/C2M

800GBASE-R SM PMA is conditional, pending implementation of 800GAUI-4 C2C/C2M"

Figure 185-1 should include a PMA sublayer in the diagram and be added to legend below
Figure 185-2 needs to be updated to show the 800GBASE-R PMA Sublayer and service interface between the PCS and Inner FEC

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Some optional and conditional sublayers are missing from Table 185-1 and the conditions for include the SM-PMA and BM-PMA should be included in this table.

Regarding Figure 185-1 and Figure 185-2, no PMA is shown because the 800GBASE-LR1 Inner FEC sublayer connects directly with the PCS; a PMA is not required between the PCS and the 800GBASE-LR1 Inner FEC. Note that the 800GBASE-LR1 Inner FEC subsumes some functions/services normally provided by a PMA for the PMD.

Add the following rows in Table 185-1:

- 800GBASE-R BM-PMA - conditional
- 800GAUI-8 C2C - optional
- 800GAUI-8 C2M - optional
- 800GBASE SM-PMA - conditional
- 800GAUI-4 C2C - optional
- 800GAUI-4 C2M - optional

Resolve the concern about conditional SM-PMA and BM-PMA related to Table 185-1 using the response to comment #317.

Implement with editorial license.

Cl 185 SC 185.7.1 P481 L21 # 375

He, Xiang Huawei

Comment Type TR Comment Status D test pattern (bucket)

The scrambled idle test pattern for 800GBASE-R PCS is defined in 172.2.4.11, not 175.2.4.11.

SuggestedRemedy

Change "175.2.4.11" to "172.2.4.11" and format as external reference.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Implement suggested remedy with editorial license.

Cl 186 SC 186 P491 L1 # 108

Huber, Thomas Nokia

Comment Type T Comment Status D (bucket)

The baseline for the 800GBASE-ER1[-20] PCS has issues with PTP accuracy when an extender sublayer is used.

SuggestedRemedy

Update the baseline per presentations in the May meeting proposing a mechanism to reduce the PTP inaccuracy.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Resolve using the proposal in https://www.ieee802.org/3/dj/public/24_05/sluski_3dj_01a_2405.pdf, which was presented in the May interim meeting. Impelemnt the suggested remedy in sluski_3dj_01a_2405 with editorial license.

Cl 186 SC 186 P491 L1 # 334

de Koos, Andras Microchip Technology

Comment Type T Comment Status D (bucket)

ER1 PCS: Planting the seed for when the PCS is ready to be properly reviewed.
How to calculate the path data delay across the ER1 PCS/PMA? Clause 90 and Annex 90A give general rules, like how to calculate the rx/tx path data delay when there are functions within the PHY that introduce cyclical delay.
But the path data delay in the ER1 PCS is very different from anything that has been imagined in Clause 90 - an Ethernet stream that floats within a GMP frame will present unique challenges; it is not immediately clear how to determine the min/max latency across such a PCS.
This might be worse than the Alignment marker issue!

SuggestedRemedy

Proposed Response Response Status W

PROPOSED REJECT.
The suggested remedy does not provide sufficient detail to implement.