

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 30 SC 30.5.1.1.4 P439 L26 # r01-1
Marris, Arthur Cadence Design Syste

Comment Type ER Comment Status A Bucket

Full stop/period should come after the quotation marks.

SuggestedRemedy

Change "remote fault." to "remote fault" on lines 9, 21 and 26.

Make similar change on line 12, 20 and on page 441 line 9, page 454 line 42, page 457 etc.

Response Response Status C

ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE P802.3/D3.1 and IEEE P802.3/D3.0 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

However, this is an editorial issue that should be addressed. It is not limited to "remote fault" (e.g., p439/l11 "offline," and l12 "auto neg error."). In Clause 30, ensure that all enumerations cited in quotation marks are as listed in "APPROPRIATE SYNTAX" and that any punctuation that is not part of the enumeration is outside of the quotation marks.

CI 82 SC 82.6 P171 L44 # r01-2
Marris, Arthur Cadence Design Syste

Comment Type E Comment Status A Bucket

There is something not right about the cross reference link to Clause 73

SuggestedRemedy

Make "Clause" part of the cross reference link for Clause 73

Response Response Status C

ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

However, the changes suggested are an improvement to the draft.
Implement the suggested remedy.

CI 119 SC 119.2.6.2.2 P76 L48 # r01-3
Gustlin, Mark Xilinx

Comment Type T Comment Status A Bucket

The restart_lock definition was not updated when the PCS synchronization state diagram was updated with comment r02-6 against D3.2 of 802.3bs:

http://www.ieee802.org/3/bs/comments/P802d3bs_D3p2_comments_final_ID.pdf#page=3

SuggestedRemedy

Change:

"restart_lock: Boolean variable that is set by the PCS synchronization process to restart the alignment marker lock process on all PCS lanes. It is set to true after 3 consecutive uncorrected codewords are received (3_BAD state) or when 5 Alignment Markers in a row fail to match (5_BAD state) and set to false upon entry into the LOSS_OF_ALIGNMENT state.

To:

"restart_lock: Boolean variable that is set by the PCS synchronization process to restart the alignment marker lock process on all PCS lanes. It is set to true after 3 consecutive uncorrected codewords are received (3_BAD state) and set to false upon entry into the LOSS_OF_ALIGNMENT state."

Response Response Status C

ACCEPT.

CI 0 SC 0 P L # r01-4
Berger, Catherine

Comment Type G Comment Status A Bucket

This draft meets all editorial requirements.

SuggestedRemedy

Response Response Status C

ACCEPT.

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

Cl 45 SC 45.2.3.16.4 P252 L2 # r01-5
Marris, Arthur Cadence Design Syst

Comment Type T Comment Status A Bucket

The first sentence of 45.2.3.16.4 needs improving.

SuggestedRemedy

Change:

The errored blocks counter is an eight bit count defined by the errored_block_count counter specified in 49.2.14.2 for 10/25GBASE-R, in 82.3.1 for 40/100GBASE-R and defined by counter errored_block_count in 126.3.7.2 in 2.5GBASE-T and 5GBASE-T, 55.3.7.2 for 10GBASE-T, and in 113.3.7.2 for 25GBASE-T and 40GBASE-T.

To:

The errored blocks counter is an eight bit count defined by the counter errored_block_count specified in 49.2.14.2 for 10/25GBASE-R, in 82.3.1 for 40/100GBASE-R, in 126.3.7.2 for 2.5GBASE-T and 5GBASE-T, in 55.3.7.2 for 10GBASE-T, and in 113.3.7.2 for 25GBASE-T and 40GBASE-T.

Also in 45.2.3.16.3 change "in 126.3.7.2 in 2.5GBASE-T and 5GBASE-T," to "in 126.3.7.2 for 2.5GBASE-T and 5GBASE-T,"

Response Response Status C

ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

However, the changes suggested are an improvement to the draft.

Implement the suggested remedy, with the exception that in 45.2.3.16.3 change "in 126.3.7.2 in 2.5GBASE-T and 5GBASE-T, 55.3.7.2 for 10GBASE-T" to "in 126.3.7.2 for 2.5GBASE-T and 5GBASE-T, in 55.3.7.2 for 10GBASE-T"

Cl 30 SC 30.5.1.1.33 P448 L40 # r01-6
Anslow, Peter Ciena Corporation

Comment Type E Comment Status A Bucket

This subclause is about PCS FEC not RS-FEC. In the last sentence "the RS-FEC control register (see 45.2.3.59)" should be "the PCS FEC control register (see 45.2.3.59)"

SuggestedRemedy

Change "the RS-FEC control register" to "the PCS FEC control register"

Response Response Status C

ACCEPT.

Cl 114 SC 114.5.6 P820 L31 # r01-7
Anslow, Peter Ciena Corporation

Comment Type E Comment Status A Bucket

Here is one remaining example of "When PMD_global_transmit_disable is set ..." which should be changed to "When the PMD_global_transmit_disable variable is set ..."

SuggestedRemedy

Change "When PMD_global_transmit_disable is set ..." to "When the PMD_global_transmit_disable variable is set ..."

Response Response Status C

ACCEPT.

Cl 120B SC 120B.3.2 P350 L31 # r01-8
Anslow, Peter Ciena Corporation

Comment Type E Comment Status A Bucket

Comment i-67 against D3.0:
<http://www.ieee802.org/3/cj/comments/P8023-D3p0-Comments-Final-byID.pdf#page=22>
was not implemented correctly. The word "where" should not have been removed.

SuggestedRemedy

Change:

"... defined by Equation (93A-46), Tr is calculated ..." to:

"... defined by Equation (93A-46), where Tr is calculated ..."

Response Response Status C

ACCEPT.

Cl 120D SC 120D.3.2.1 P373 L1 # r01-9
Anslow, Peter Ciena Corporation

Comment Type E Comment Status A Bucket

Space missing in "Equation (93A-3)in"

SuggestedRemedy

Add the space

Response Response Status C

ACCEPT.

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 120D SC 120D.3.2.1 P373 L3 # r01-10
Anslow, Peter Ciena Corporation

Comment Type E Comment Status A Bucket

Comment i-67 against D3.0:
<http://www.ieee802.org/3/cj/comments/P8023-D3p0-Comments-Final-byID.pdf#page=22>
was not implemented correctly. The word "where" should not have been removed.

SuggestedRemedy

Change:
"... defined by Equation (93A-46), Tr is calculated ..." to:
"... defined by Equation (93A-46), where Tr is calculated ..."

Response Response Status C
ACCEPT.

CI 30 SC 30.5.1.1.15 P443 L4 # r01-11
Ran, Adeo Intel Corporation

Comment Type T Comment Status A Bucket

The aFECAbility attribute seems inadequate for clause 119; the clause 119 PCS has forward error correction internally, and unconditionally has FEC ability.

All the other clauses listed here are optional for some PCSs (a PHY with the same PCS can either have or not have FEC ability).

Other PHYs that have such unconditional FEC functionality as part of the PCS do not have this attribute. It makes sense since in these PHYs there is no need for this attribute; FEC is always supported. The same is true in clause 119.

Examples of PHYs which have FEC unconditionally are clauses 55, 113, 115. If they are not listed here, why should clause 119 be listed?

The change in the second paragraph is a result of the inclusion of clause 119 and results in awkward text.

SuggestedRemedy

Revert the changes in this subclause to the text in D3.0, removing the reference to clause 119 and the changed text in the second paragraph.

Response Response Status C

ACCEPT IN PRINCIPLE.
Change the text below "BEHAVIOUR DEFINED AS:" to the following.
"A read-only value that indicates if the PHY supports an optional FEC sublayer for forward error correction (see 65.2, Clause 74, Clause 91, and Clause 108).
If a Clause 45 MDIO Interface is present, then this attribute maps to the FEC capability register (see 45.2.10.2 or 45.2.1.97).;"

CI 30 SC 30.5.1.1.32 P448 L11 # r01-12
Ran, Adeo Intel Corporation

Comment Type T Comment Status A Bucket

The attribute name aPCSFECIndicationAbility is misleading: it is the ability to bypass indication, not to indicate.

If one reads a value of "not supported", it could be wrongly interpreted as if indication ability is not supported.

Similarly in 30.5.1.1.33 the name is "aPCSFECIndicationEnable" but it is the enable for bypassing indication, not for indication.

These names also contradict the way indication ability/enable are defined in the BASE-R FEC, see 45.2.1.101.2 and 45.2.1.102.2.

I noticed that the same problem exists in the existing attributes aRSFECIndicationAbility (30.5.1.1.29) and aRSFECIndicationEnable (30.5.1.1.31); These should be corrected too.

SuggestedRemedy

Change the attribute name in 30.5.1.1.29 to "aRSFECBypassIndicationAbility".
Change the attribute name in 30.5.1.1.31 to "aRSFECBypassIndicationEnable".
Change the attribute name in 30.5.1.1.32 to "aPCSFECBypassIndicationAbility".
Change the attribute name in 30.5.1.1.33 to "aPCSFECBypassIndicationEnable".

Update Table 30-1e accordingly.

Response Response Status C
ACCEPT.

CI 30 SC 30.5.1.1.33 P448 L40 # r01-13
Ran, Adeo Intel Corporation

Comment Type T Comment Status A Bucket

"If a Clause 45 MDIO Interface is present, then this attribute maps to the RS-FEC control register (see 45.2.3.59)."

This is the PCS FEC control register.

SuggestedRemedy

Change "RS-FEC" to "PCS FEC".

Response Response Status C
ACCEPT.

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 45 SC 45.2.1.6.3 P72 L51 # r01-14
 Ran, Adeo Intel Corporation
 Comment Type E Comment Status A Bucket
 Serial comma should be a serial semicolon to match the other semicolons in this list.
 SuggestedRemedy
 Change "ability register, and the 400G" to "ability register; and the 400G"
 Response Response Status C
 ACCEPT.

CI 45 SC 45.2.1.8 P80 L5 # r01-15
 Ran, Adeo Intel Corporation
 Comment Type E Comment Status A Bucket
 The PMD transmit disable register is now extended by the PMD transmit disable extension register (Register 1.27) that includes the bit for lane 15 (which does not fit here).
 It may be helpful for readers to mention this and provide cross references.
 Similarly for the PMD receive signal detect register in 45.2.1.9.
 SuggestedRemedy
 Insert another NOTE before the current NOTE, with the text
 "This register is extended by the PMD transmit disable extension register (See 45.2.1.22)."
 Insert a NOTE similarly in 45.2.1.22 with the text
 "This register is an extension of the PMD transmit disable register (See 45.2.1.8)."
 Insert a NOTE after the first paragraph of 45.2.1.9 with the text
 "This register is extended by the PMD receive signal detect extension register (See 45.2.1.23)."
 Add a NOTE after the first paragraph of 45.2.1.23 with the text
 "This register is an extension of the PMD receive signal detect register (See 45.2.1.9)."
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Change the existing note in 45.2.1.8 to be "NOTE 2"
 Insert NOTE 1 before NOTE 2 in 45.2.1.8, with the text:
 "NOTE 1-This register is extended by the PMD transmit disable extension register (see 45.2.1.22)."
 Change the existing note in 45.2.1.22 to be "NOTE 2"
 Insert NOTE 1 before NOTE 2 in 45.2.1.22, with the text:
 "NOTE 1-This register is an extension of the PMD transmit disable register (see 45.2.1.8)."
 Insert a NOTE after the first paragraph of 45.2.1.9 with the text:
 "NOTE-This register is extended by the PMD receive signal detect extension register (see 45.2.1.23)."
 Insert a NOTE after the first paragraph of 45.2.1.23 with the text:
 "NOTE-This register is an extension of the PMD receive signal detect register (see 45.2.1.9)."

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 45 SC 45.2.3.61 P285 L3 # r01-16
Ran, Adee Intel Corporation

Comment Type TR Comment Status A

The PCS FEC corrected codewords counter and the PCS FEC uncorrected codewords counter registers do not have explicit definitions. It isn't specified anywhere when these registers should be incremented and by how much.

This may seem obvious for the "corrected" counter, but the "uncorrected" counter can be interpreted as "no need for correction" rather than "uncorrectable" (which is proper).

Compare to the corresponding counter registers for RS-FEC (45.2.1.112 and 45.2.1.113) that point to explicit definitions for the register, which clarify the meaning; there are no corresponding definitions in clause 119.

Applies similarly to 45.2.4.23, 45.2.4.24, 45.2.5.23, 45.2.5.24 (same counters in the XS).

This may also apply to the symbol error counters in 45.2.3.57 and the corresponding XS registers 45.2.4.19 and 45.2.5.19 (compare to RS-FEC registers in 45.2.1.115).

This may also apply to other registers without explicit definition in clause 119. Compare to 91.6.

Suggested Remedy

Add new subclauses under 119.3 with definitions of FEC_corrected_cw_counter and FEC_uncorrected_cw_counter, using 91.6.9 and 91.6.10 respectively as examples.

Add cross-references to these definitions in 45.2.3.61, 45.2.3.62, 45.2.4.23, 45.2.4.24, 45.2.5.23, and 45.2.5.24.

Alternatively refer to the existing definitions in 91.6.9 and 91.6.10 without adding anything in 119.3 (though this may be awkward).

Consider adding similar definitions for symbol error counters, and referring to them in 45.2.3.57, 45.2.4.19, and 45.2.5.19.

Consider applying to other registers without explicit definition in clause 119.

Response Response Status C

ACCEPT IN PRINCIPLE.

Add new subclauses under 119.3:

119.3.2 FEC_corrected_cw_counter

A corrected FEC codeword is a codeword that contains errors and was corrected.

FEC_corrected_cw_counter is a 32-bit counter that counts once for each corrected FEC codeword processed when align_status is true. This variable is mapped to the registers defined in 45.2.3.61 (3.802, 3.803).

119.3.3 FEC_uncorrected_cw_counter

An uncorrected FEC codeword is a codeword that contains errors that were not corrected. FEC_uncorrected_cw_counter is a 32-bit counter that counts once for each uncorrected FEC codeword processed when align_status is true. This variable is mapped to the registers defined in 45.2.3.62 (3.804, 3.805).

119.3.4 FEC_symbol_error_counter_i

FEC_symbol_error_counter_i (where i=0 to 7 for the 200GBASE-R PCS and i=0 to 15 for the 400GBASE-R PCS) are 32-bit counters that count once for each 10-bit symbol corrected on PCS lane l when align_status is true. These variables are mapped to the registers defined in 45.2.3.57 and 45.2.3.58 (3.600 to 3.631).

In 45.2.3.61, 45.2.4.23, and 45.2.5.23 insert a new second sentence:
"See 119.3.2 for a definition of this counter."

In 45.2.3.62, 45.2.4.24, and 45.2.5.24, insert a new second sentence:
"See 119.3.3 for a definition of this counter."

In 45.2.3.57, 45.2.4.19, and 45.2.5.19, insert a new second sentence:
"See 119.3.4 for a definition of this counter."

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 121 SC 121.8.5.4 P135 L18 # r01-17
King, Jonathan Finisar Corporation

Comment Type T Comment Status A

"Comment i-107 against P802.3cd D3.0:
http://www.ieee802.org/3/cd/comments/8023cd_D30_final_comment_responses_by_clause.pdf#page=57
added a constraint on the main tap location for the equalisers specified in:
138.8.5.1 for 50GBASE-SR, 100GBASE-SR2, and 200GBASE-SR4
139.7.5.4 for 50GBASE-FR and 50GBASE-LR
140.7.5.1 for 100GBASE-DR
as detailed in:
http://www.ieee802.org/3/cd/public/Jan18/king_3cd_03_0118.pdf
and as justified in:
http://www.ieee802.org/3/cd/public/Jan18/sun_3cd_01a_0118.pdf

Since 50GBASE-FR, 50GBASE-LR, and 100GBASE-DR are expected to use the same technology as one lane of the corresponding multi-lane PMDs now included in the revision (200GBASE-FR4, 400GBASE-FR8, 200GBASE-LR4, 400GBASE-LR8, and 400GBASE-DR4), this has introduced an inconsistency between the two sets of specifications that should be removed."

SuggestedRemedy

"Make changes to:
121.8.5.4 for 200GBASE-DR4 (and by reference 400GBASE-DR4)
122.8.5.4 for 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8
equivalent to the changes made in P802.3cd 139.7.5.4 between D3.0 and D3.1:
Add the text:
""A functional model of the reference equalizer is shown in Figure 12x-y.""
""Tap 1, tap 2, or tap 3, has the largest magnitude tap coefficient.""
and a figure in each case."
"

Response Response Status C
ACCEPT.

CI 121 SC 121.7.2 P128 L17 # r01-18
King, Jonathan Finisar Corporation

Comment Type T Comment Status A

"Comment i-78 against P802.3cd D3.0:
http://www.ieee802.org/3/cd/comments/8023cd_D30_final_comment_responses_by_clause.pdf#page=61
Changed the informative receiver sensitivity specified in:
138.8.7 for 50GBASE-SR, 100GBASE-SR2, and 200GBASE-SR4
139.7.8 for 50GBASE-FR and 50GBASE-LR
140.7.8 for 100GBASE-DR
as detailed in:
http://www.ieee802.org/3/cd/public/Jan18/king_3cd_04_0118.pdf
also see related proposal in:
http://www.ieee802.org/3/cd/public/Jan18/lewis_3cd_01_0118.pdf

Since 50GBASE-FR, 50GBASE-LR, and 100GBASE-DR are expected to use the same technology as one lane of the corresponding multi-lane PMDs now included in the revision (200GBASE-FR4, 400GBASE-FR8, 200GBASE-LR4, 400GBASE-LR8, and 400GBASE-DR4), this has introduced an inconsistency between the two sets of specifications that should be removed."

SuggestedRemedy

"Make changes to:
Table 121-7 and 121.8.8 for 200GBASE-DR4
Table 122-11, Table 122-12, and 122.8.8 for 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8
Table 124-7 and 124.8.8 for 400GBASE-DR4
equivalent to the changes made in P802.3cd Table 139-7 and 139.7.8 between D3.0 and D3.1:
In the tables, replace the Receiver sensitivity value with a cross-reference to a new equation and modify the table footnotes
In the Receiver sensitivity subclauses replace the existing text with equations and illustrative figures"

Response Response Status C
ACCEPT IN PRINCIPLE.
Implement the changes in <http://www.ieee802.org/3/maint/public/anslow_1_0318.pdf>.

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 116 SC 116.6 P32 L9 # r01-19
 Slavick, Jeff Broadcom Limited

Comment Type TR Comment Status A Bucket

The FEC degrade feature is partially optional. The detection of a degraded signal is optional. However, it's mandatory that the signalling of a degraded condition is done by all.

SuggestedRemedy

Change "FEC degrade is an optional feature allowing for the detection of a non-service affecting link degradation condition based on exceeding a threshold for FEC corrected errors."

To "FEC degrade is a feature allowing for the optional detection of a non-service affecting link degradation condition based on exceeding a threshold for FEC corrected errors."

Response Response Status C

ACCEPT.

CI 121 SC 121.8.9.2 P137 L46 # r01-20
 Ran, Adeo Intel Corporation

Comment Type T Comment Status A

Several comments against P802.3cd D3.0 noted that the SRS test conditions can be calibrated in multiple ways.

(Note that although for the reference receiver the SRS result is independent of the choice of stress conditions, this may not be true for specific implementations. For example, a receiver with better equalization capabilities than the reference receiver but with more internal noise may pass the test if the stress is mostly ISI, but fail if the stress is mostly uncorrelated noise.)

The response to comment i-58 against P802.3cd D3.0 indicated that there is deliberate freedom in setting up the SRS test source.

Discussions following presentations related to that comment (e.g. schube_011718_3cd_adhoc) indicated that this freedom is desirable, since different PMD transmitters with different characteristics can be used by link partners (for example, high bandwidth with large noise, or low bandwidth with low noise). Narrowing down the test parameters may exclude conditions caused by some compliant transmitters.

This implies that in order to interoperate with any compliant transmitter, a receiver should pass the SRS test regardless of how the stress signal is calibrated.

This may seem obvious for people with deep understanding of the standard, but test engineers may have different interpretations, and may decide based on only one test condition that happens to make the DUT pass. This approach also enables "gaming the test" by choosing particular test conditions that are favorable for a device.

It is suggested to clarify the intent of the freedom of choice of stress conditions with an informative note.

Note that a similar comment is submitted against P802.3cj D3.1 for several subclauses. Since the definitions in 121.8.9.2 are inherited by all SRS test subclauses in clauses 122, 123, 124, 138, 139, and 140, adding a single note in 121.8.9.2 may be sufficient.

SuggestedRemedy

Add the following note at the end of 121.8.9.2:

NOTE--The stress conditions in the SRS test can be calibrated in several ways. A compliant PMD receiver is expected to meet the sensitivity requirements with a calibrated conformance test signal regardless of the choice of stress components.

Consider adding similar notes in 122.8.9.2, 123.8.8, and 124.8.9.

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

Response **Response Status** **C**

ACCEPT IN PRINCIPLE.
Add the following note at the end of 121.8.9.2:
"NOTE-A compliant PMD receiver is expected to meet the stressed receiver sensitivity requirements with a calibrated conformance test signal regardless of the proportion (as long as it is above half) of the dB value of the SECQ that is due to the frequency response of the combination of the low-pass filter and the E/O converter."

There is no need to add this text to 122.8.9.2, 123.8.8, or 124.8.9 as all of these refer back to 121.8.9.2 for the details of the procedure.

CI 120D	SC 120D.3.1	P366	L27	# r01-21
Ran, Ade		Intel Corporation		
Comment Type E	Comment Status A		Bucket	

"J_RMS" seems to be in larger point size than the rest of the text in this table.

SuggestedRemedy

Unify text size as appropriate.

Response **Response Status** **C**

ACCEPT IN PRINCIPLE.
Change the font size for "J_RMS" to be consistent with other text in the table.

CI 45	SC 45.2.1.139	P181	L1	# r01-22
Marris, Arthur		Cadence Design Syste		
Comment Type E	Comment Status A		Bucket	

The hex character fields don't begin with 0x

SuggestedRemedy

Change "for lane 0, fbf1cb3e; for lane 1, fbb1e665; for lane 2, f3fdae46; for lane 3, f2ffa46b" to be "for lane 0, 0xfb1cb3e; for lane 1, 0xfbb1e665; for lane 2, 0xf3fdae46; for lane 3, 0xf2ffa46b"

Response **Response Status** **C**

ACCEPT IN PRINCIPLE.
This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

However, the changes suggested are an improvement to the draft.
Implement the suggested remedy.

CI 120D	SC 120D.3.1	P366	L14	# r01-23
Rysin, Alexander		Mellanox Technologies		

Comment Type **TR** **Comment Status** **R**

Transmitter output residual ISI SNR_ISI (min) 34.8 dB (Clause 120D) is too high - can barely measure the IC through the test fixture. The warning NOTE in 120D.3.1.7 shows the issue, but doesn't solve it. 802.3cd D2.0 comment 140, D2.1 comment 49, D2.2 comment 22. Since both SNR_ISI and Effective Return Loss (ERL) represent uncompensated reflections from the transmitter and the test fixtures, measurements of ERL can replace SNR_ISI.
Also, frequency domain return loss mask does not truly represent digital signaling at a given bit error ratio. There is no real proof that violating return loss masks is directly tied to failures and a number of false negatives have been shown. 802.3cd D2.0 comment 141, D2.1 comments 26, 27 and 28, D3.0 comment 98.

SuggestedRemedy

- * Add an Annex describing ERL computation method and parameters. The Annex can be copied from 93A-5 in 802.3cd D3.1.
- * Add a parameter Table, copying Table 137-5 for 802.3cd D3.1.
- * Add a description of the ERL computation and parameters as follows:

Effective return loss (ERL) of the transmitter at TP0a is computed using the procedure in Annex (new) with the values in Table TBD. Parameters that do not appear in Table TBD take values from Table 120D-8. The value of Tfx is twice the delay from TP0 to TP0a. Nbx is set to the value of Nb in Table 120D-8. ERL shall be at least 16.1 dB.

- * Add a reference in 120D.3.2 to Annex (new) and to Table TBD for a description of the ERL computation and parameters as follows:

Effective return loss (ERL) of the receiver computed using the procedure in Annex (new) with the values in Table TBD. Parameters that do not appear in Table TBD take values from Table 120D-8. The value of Tfx is twice the delay from TP5a to TP5. Nbx is set to the value of Nb in Table 120D-8. ERL shall be at least 16.1 dB.

- * Remove the requirement for Differential return loss in Table 120D-1.
- * Add a requirement for Effective Return Loss (ERL) to be greater than 16.1 dB in Table 120D-1
- * Remove the requirement for Differential input return loss in Table 120D-5
- * Add a requirement for Effective Return Loss (ERL) to be greater than 16.1 dB in Table 120D-5
- * Remove reference to Transmitter Output residual ISI SNR_ISI(min) in Table 120D-1.

Response **Response Status** **U**

REJECT.
Annex 93A.5 and Effective Return Loss (ERL) specifications were first introduced in IEEE P802.3cd/D3.1 (January 2018). During the March 2018 meeting of IEEE P802.3cd Task Force, numerous changes were proposed to the ERL parameters and ERL requirements (see <http://www.ieee802.org/3/cd/public/Mar18/dudek_3cd_02_0318.pdf>) and many

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

changes were adopted. Many of these new values are to include notes that state values are "to be confirmed". Based on this, it appears that the new specification is not mature enough to incorporate into the draft at this time.

The specifications in Annex 120D (CDAUI-8 chip-to-chip) were approved as part of IEEE Std 802.3bs-2017. The concept of ERL was introduced in the IEEE P802.3cd amendment for the backplane and copper cable interfaces. In these cases the link budget margins are considerably lower making the imprecision of return loss masks more impactful and SNR_ISI requirements more demanding (34.8 dB for Annex 120D vs. the a placeholder value of 43 dB for IEEE P802.3cd/D3.1 Clause 137). It has not been established that the use of ERL for chip-to-chip (or chip-to-module interfaces, which are not mentioned in the comment despite their use of return loss masks) provide benefits that outweigh the risk of imposing new requirements on devices compliant to the original standard. It has been shown that devices that pass current return loss requirements do not necessarily pass the proposed ERL requirements.

Cl 1	SC 1.4.281	P92	L4	# r01-24
Nikolich, Paul		INDEPENDENT		

Comment Type GR Comment Status R

*** Comment submitted with the file 96131200003-20180124_163855.jpg attached ***

The proposed resolution is an improvement, but unacceptable:

"A logical subset of the data and control information transmitted from one sublayer (e.g., PCS, PMA) to an adjacent sublayer across the inter-sublayer interface or from one PHY to another across the transmission medium (e.g. optical fiber, optical wavelength, wire pair). Lanes are transmitted in parallel and combine to deliver the full set of data and control information across the interface."

My comments:

a) The proposed text doesn't quiet capture the concept of arbitrary recombination of the smallest subsets into larger subsets (which are not identical to the originating superset. Perhaps adding the word 'superset' will help as follows:

"A logical subset of a superset of data and control information transmitted from one sublayer (e.g.,PCS, PMA)..."

b) The text should be accompanied by an illustrative figure similar to the one you drew for me in Geneva. See attached file.

SuggestedRemedy

See suggestion in above comment.

Response Response Status U

REJECT.

The definition is specific to the transmission of control and data information from "one sublayer (e.g., PCS, PMA) to an adjacent sublayer across the inter-sublayer interface or from one PHY to another across the transmission medium." While the number of output lanes may be changed from the number of input lanes by a sublayer (e.g., it may aggregate subsets into larger subsets or divide subsets into smaller subsets), this is a function of the sublayer and not inherent to the definition of a lane. The definition of lane applies to the input of the sublayer and the output of the sublayer while the functions within the sublayer are beyond the scope of this definition. The proposed addition of the term "superset" does not appear to improve the definition in this context.

The inclusion of a figure with a definition is unprecedented in IEEE Std 802.3 (although it is acknowledged there is an example of this in IEEE Std 802.16-2017 and other standards under IEEE-SA). Regardless, it is believed that the definition is clear as it is written and does not require a figure.

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 82 SC 82.2.19.2.2 P162 L3 # r01-25
Marris, Arthur Cadence Design Syste

Comment Type T Comment Status A

The QUICK_FIND state is only used when in Deep Sleep mode. Consider gating "first_rx_LPI_active" with "LPI_FW = FALSE"

SuggestedRemedy

on page 156 line 29 change:

first_rx_lpi_active

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

To:

first_rx_lpi_active

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

Response Response Status C

ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

However, the changes suggested are an improvement to the draft. Implement the suggested remedy.

CI 120E SC 120E.3.3.2.1 P393 L21 # r01-26
Dawe, Piers J G Mellanox Technologies

Comment Type E Comment Status A Bucket

"The counter propagating crosstalk channels... are asynchronous" Not channels, should be lanes or signals. One can think in terms of signals that run on or in lanes. In 83E.3.1.6 we changed to signals.

SuggestedRemedy

Change channels to signals. Also p395 line 46.

Response Response Status C

ACCEPT.

CI 45 SC 45.2.1 P63 L11 # r01-27
Dawe, Piers J G Mellanox Technologies

Comment Type E Comment Status A Bucket

Although this table is far too long, it should not be made longer than necessary.

SuggestedRemedy

Make the middle column wider, so that the table is full width.

Response Response Status C

ACCEPT.

CI 93 SC 93.8.1.3 P470 L37 # r01-28
Dawe, Piers J G Mellanox Technologies

Comment Type T Comment Status A Bucket

"Measurement of the DC common-mode voltage is made with a high-impedance connection to TP0a where TP0a is AC-coupled to a 100 ohm differential termination." Compare Fig 85-6 or 92-15, where there are two 50 ohm resistors to GND - not the same as 100 ohm differential. Which is right?

SuggestedRemedy

If the 2 x 50 ohm method is correct, change to "Measurement of the DC common-mode voltage is made with a high-impedance connections to TP0a where TP0a is AC-coupled to 50 ohm terminations." ?

Response Response Status C

ACCEPT IN PRINCIPLE.

Change to:

"Measurement of the DC common-mode voltage is made with a high-impedance connection to TP0a where each conductor of the differential pair is AC-coupled to a 50 Ohm termination."

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 114 SC 114.6.1 P821 L54 # r01-29
Dawe, Piers J G Mellanox Technologies

Comment Type E Comment Status A Bucket

Table is broken over a page break (and it doesn't have a thin bottom border before the break), but the next page is mostly empty.

SuggestedRemedy

Might as well start 114.6.1 at the top of the next page

Response Response Status C

ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

However, the changes suggested are an improvement to the draft.
Implement the suggested remedy.

CI 45 SC 45.2.1.4 P70 L26 # r01-30
Dawe, Piers J G Mellanox Technologies

Comment Type E Comment Status A Bucket

Layout

SuggestedRemedy

Make the middle column wider, so that the table is full width. Make other columns narrower if needed to make this text fit on one line.

Response Response Status C

ACCEPT.

CI 120E SC 120E.3.1 P385 L19 # r01-31
Dawe, Piers J G Mellanox Technologies

Comment Type TR Comment Status A

Please add the host output VEC spec to 120E that P802.3cd has adopted in 135G. See http://ieee802.org/3/cd/public/Jan18/dawe_3cd_01_0118.pdf for more information. It seems better to put the module stressed input VEC limit in the text than in Table 120E-8, Module stressed input parameters, because the items there are calibration targets and VEC for module stressed input isn't, it's a maximum.

SuggestedRemedy

Add a 12 dB max Vertical eye closure (VEC) spec in Table 120E-1, 200GAUI-4 and 400GAUI-8 C2M host output characteristics (at TP1a).

Add a PICS item.

Copy 802.3cd 135G.4.1, Vertical Eye Closure, to the end of 120E.4.2, or after step 7 if it is preferred to complete the vertical specs before addressing the horizontal specs.

In 120E.3.4.1.1, change as follows:

Eye height and eye width are then measured at TP1a... to

Eye height, eye width, and vertical eye closure are then measured at TP1a...

After "restriction that the CTLE setting has to be greater than or equal to 7 dB does not apply.", add:

In both cases, the input vertical eye closure is less than 12 dB.

It would be good to insert a paragraph break before "The pattern is then changed to Pattern 5, Pattern 3, or a valid 200GBASE-R/400GBASE-R signal for the input test..." as it's about doing the test (as are the next two paragraphs) rather than setup and calibration.

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement the changes in http://www.ieee802.org/3/maint/public/anslow_2_0318.pdf.

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI **83E** SC **83E.3.2** P**638** L**19** # **r01-32**
Dawe, Piers J G Mellanox Technologies

Comment Type **E** Comment Status **R** Bucket

Vertical Eye Closure or vertical eye closure? Mostly it's in lower case, but it is used several ways:
vertical eye closure penalty as defined in 52.9.9.2 (a defined parameter)
vertical eye closure (a different defined parameter)
TDECQ is a measure of each optical transmitter's vertical eye closure when...
Vertical eye closure histograms

SuggestedRemedy

It would be as well to make both vertical eye closure and vertical eye closure penalty into proper nouns, to distinguish them from each other and the ordinary English meaning of the words in those other phrases.

Response Response Status **C**

REJECT.
This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

The phrase "vertical eye closure" is consistently presented in lower case in this draft (except where it is grammatically correct to capitalize the first word of the phrase). Making "vertical eye closure" a proper noun does not add clarity and may have unintended consequences in some cases.

CI **120** SC **120.5.5** P**102** L**53** # **r01-33**
Dawe, Piers J G Mellanox Technologies

Comment Type **T** Comment Status **A** LF_Jitter

This might be a suitable, though obscure, place to add hints that the implementer may have to pay attention to the low frequency jitter issue.
This comment is similar to 802.3cd comment 54.

SuggestedRemedy

Add text e.g. "The PMA output attached to an AUI or PMD conditions the output clock such that the AUI output or PMD transmitter meets its requirements." At page 103 line 11, add NOTE--Excessive low-frequency jitter might prevent the PMA from providing adequate clock quality, particularly when multiple input lanes are mapped to a single output lane.

Response Response Status **C**

ACCEPT IN PRINCIPLE.
Add the following note after the first paragraph in 120.5.5:
"NOTE-For a PMA where the signaling rate on each output lane is higher than the signaling rate on each input lane, any low-frequency jitter on the input lanes may result in more jitter relative to the UI on the output lanes."

CI **83** SC **83.5.3** P**188** L**8** # **r01-34**
Dawe, Piers J G Mellanox Technologies

Comment Type **T** Comment Status **R** Bucket

Correct this text to acknowledge that not all PMA interfaces are multi-lane, so not all have Skew Variation, and some Skew values are not as given.

SuggestedRemedy

Change: The limits for Skew and Skew Variation at physically instantiated interfaces are specified at Skew points SP0, SP1, and SP2 in the transmit direction and SP5, SP6, and SP7 in the receive direction as defined in 80.5 and illustrated in Figure 80-6, Figure 80-7, and Figure 80-8. to:

For 40GBASE-FR, the limits for Skew at physically instantiated interfaces are specified at Skew points SP0, SP1, and SP2 in the transmit direction and SP5, SP6, and SP7 in the receive direction as defined in 80.5 and illustrated in Figure 80-6, Figure 80-7. For 40GBASE-FR, the limits for Skew Variation at physically instantiated interfaces are specified at Skew points SP0 and SP1 in the transmit direction, and SP6 and SP7 in the receive direction, as defined in 80.5 and illustrated in Figure 80-6 and Figure 80-7. Except for 40GBASE-FR, the limits for Skew and Skew Variation at physically instantiated interfaces are specified at Skew points SP0, SP1, and SP2 in the transmit direction and SP5, SP6, and SP7 in the receive direction as defined in 80.5 and illustrated in Figure 80-6, Figure 80-7, and Figure 80-8.

Response Response Status **C**

REJECT.
This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

Also, this text is correct as it is.

It is true that for 40GBASE-FR the Skew at SP5 (on the PMD service interface, at the output of the PMD) is the same as the Skew at SP2 (on the PMD service interface, at the input of the PMD). However, the IEEE P802.3bg project chose to maintain the same limit for Skew at SP5 for this PMD as for all of the other 40GBASE-R PMDs on the basis that the 40GBASE-FR PMD would have no difficulty in meeting this limit and there was no advantage in specifying a lower value for this PMD than for all of the others. Similarly, although there is no Skew Variation for 40GBASE-FR at SP2, SP3, SP4, or SP5, the IEEE P802.3bg project chose to maintain the same limit for Skew Variation at these skew points for all of the PMDs including 40GBASE-FR for the same reasons. The text in 89.3.2 states that: "there is no Skew Variation at this point" for SP2, SP3, SP4, and SP5 so it is obvious that the limit is met.

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI **121** SC **121.8.5.3** P**132** L**1** # **r01-35**
Dawe, Piers J G Mellanox Technologies

Comment Type **TR** Comment Status **R**

It seems that it is possible to make a bad transmitter (e.g. with a noisy or distorted signal), use emphasis to get it to pass the TDECQ test, yet leave a realistic, compliant receiver with an unreasonable challenge, such as high peak power, high crest factor, or a need to remove emphasis from the signal, contrary to what equalizers are primarily intended to do. Note the receiver is tested for a very slow signal only, not for any of these abusive signals. This is an issue for all the PAM4 optical PMDs, although it may be worse for MMF because of the high TDECQ limit and because the signal is measured in a particularly low bandwidth. This comment updates 802.3cd D3.1 comment 71. With luck it will be possible to follow 802.3cd's action on this topic.

SuggestedRemedy

1. To screen for noisy or distorted signals with heavy emphasis:
 - 1a. Define a metric similar to TDECQ but with C_{eq} held at 1, that measures how closed the eye after the reference equalizer is. Set a limit for it.
 - or:
 - 1b. Define $TDECQ_{rms} = 10 \cdot \log_{10}(A_{RMS}/(s^3 \cdot Q_t \cdot R))$ where A_{RMS} is the standard deviation of the measured signal after the 13.28125 GHz or 11.2 GHz filter response (before the FFE), Q_t and R are as already in Eq 121-12. s is the standard deviation of a fast clean signal with OMA=2 and without emphasis, observed through the filter response (0.6254 for 13.28125 GHz, 0.6006 for 11.2 GHz).
- Either, set limit for $TDECQ_{rms}$ according to what level of dirty-but-emphasised signal we decide is acceptable, add max $TDECQ_{rms}$ row to each transmitter table.
- Or, if the same relative limit is acceptable for all PAM4 optical PMDs, the limit could be here in the TDECQ procedure. E.g. make the $TDECQ_{rms}$ limit the same as the TDECQ limit, say here that both TDECQ and $TDECQ_{rms}$ must meet the TDECQ spec.
2. To protect the receiver from having to "invert" heavily over-emphasised signals, set a minimum cursor weight, 0.9. Similarly in clauses 122, 124.
- To protect the equalizer from having to support unnecessary settings for waveforms that can't or shouldn't ever happen, constrain the cursor position - see other comments.

Response Response Status **U**

REJECT.
There are no PAM4 optical PMDs (that would use the TDECQ test) over MMF in the draft. "Eq 121-12" in the suggested remedy should be "Eq 121-12".

The need for additional transmitter specs for the SMF PMDs has not been established, and insufficient evidence has been provided that the proposed alternative remedies fix the claimed problem.

To date no contribution has been made that demonstrates the problem (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that one of the proposed additional requirements prevents this issue from occurring.

A similar proposal to create a $TDECQ_{rms}$ spec was suggested in comments i-140 against P802.3bs D3.0, r02-35 against P802.3bs D3.2 and r03-27 against P802.3bs D3.3 which were similarly rejected.

A peak power spec has not been shown to be necessary, and a definition and value has

not been provided.

A crest factor limit has not been shown to be necessary, and a definition and value has not been provided.

The need for a limit to cursor weight has not been established.

Constraints have been placed on the cursor position due to the changes made in response to comment r01-17.

[Editor's note added after comment resolution completed.

The suggested remedy for comment r01-17 was accepted. The suggested remedy is:

"Make changes to:

121.8.5.4 for 200GBASE-DR4 (and by reference 400GBASE-DR4)

122.8.5.4 for 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8

equivalent to the changes made in P802.3cd 139.7.5.4 between D3.0 and D3.1:

Add the text:

""A functional model of the reference equalizer is shown in Figure 12x-y.""

""Tap 1, tap 2, or tap 3, has the largest magnitude tap coefficient.""

and a figure in each case."

]

CI **121** SC **121.8.5.3** P**134** L**45** # **r01-36**
Dawe, Piers J G Mellanox Technologies

Comment Type **TR** Comment Status **R**

The TDECQ method allows signals that are slower than 100GBASE-LR4, probably slower than the original T/2-spaced TDECQ allowed, and slower than anticipated. If this hole is not plugged, product receivers will have to provide more tap strength than is needed to receive the range of reasonable signals, degrading their cost/power/performance trade-off. This issue became more clear after the 802.3cd comments were written, but with luck, 802.3cd will consider the matter as part of their TDECQ comment resolution anyway.

SuggestedRemedy

Set a maximum cursor strength limit, which might be around 1.3.

Similarly in clauses 122, 124.

Response Response Status **U**

REJECT.

The need for a limit to cursor weight has not been established (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that the proposed limit of 1.3 removes the demonstrated issue while not disallowing "reasonable" transmitters.

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 121 SC 121.8.5.4 P135 L18 # r01-37
Dawe, Piers J G Mellanox Technologies

Comment Type TR Comment Status A

802.3cd has adopted cursor position rules that should apply here too. Further, the rules should be tightened (see http://ieee802.org/3/cd/public/Mar18/dawe_3cd_01_0318.pdf).

SuggestedRemedy

Copy the new material from 138.8.5.1, including Figure 138-3, TDECQ reference equalizer functional model. However, (802.3cd comment 76, instead of "Tap 1, tap 2, or tap 3, has the largest magnitude tap coefficient", use "Tap 1 or tap 2 has the largest magnitude tap coefficient".

Specifications work at different levels: functional, logic/digital, analog (electrical or optical), and "Functional" is the highest/most abstract, while this FFE diagram is part of the specification of an analog quantity (more at 802.3cd comment 72). So instead of "symbol period. A functional model of the reference equalizer is shown in Figure 138-3" use "symbol period, as shown in Figure 138-3", and in the figure title, instead of "TDECQ reference equalizer functional model" use "TDECQ reference equalizer".

Response Response Status U

ACCEPT IN PRINCIPLE.

See response to comment r01-17 which applies the restriction that the main tap has to be tap1, tap2, or tap3.

It has not been demonstrated that disallowing tap 3 as having the largest magnitude tap coefficient is an improvement to the draft. (Indeed, several of the contributed measurements have shown tap3 as the largest magnitude tap coefficient for the optimum tap setting.)

Regarding the "functional model" description, the text and figure follow the precedent set in IEEE Std 802.3bs-2017 Annex 120D for an equivalent type of equalizer.

[Editor's note added after comment resolution completed.

The suggested remedy for comment r01-17 was accepted. The suggested remedy is:

"Make changes to:

121.8.5.4 for 200GBASE-DR4 (and by reference 400GBASE-DR4)

122.8.5.4 for 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 equivalent to the changes made in P802.3cd 139.7.5.4 between D3.0 and D3.1:

Add the text:

""A functional model of the reference equalizer is shown in Figure 12x-y.""

""Tap 1, tap 2, or tap 3, has the largest magnitude tap coefficient.""

and a figure in each case."

]

CI 120E SC 120E.3.3.2.1 P393 L7 # r01-38
Dawe, Piers J G Mellanox Technologies

Comment Type E Comment Status A Bucket

Remove the ambiguity in the table and be clear like Table 121-12 and similar.

SuggestedRemedy

Change "Jitter amplitude" to "Jitter amplitude (pk-pk)".

Response Response Status C

ACCEPT.

CI 116 SC 116 P19 L1 # r01-39
Dawe, Piers J G Mellanox Technologies

Comment Type T Comment Status R

802.3cd may make changes to material similar to clauses 116 to 124 and their annexes that should be applied here too.

SuggestedRemedy

Apply them as appropriate.

Response Response Status C

REJECT.

The proposed change in the comment does not contain sufficient detail to enable the specific changes that satisfy the commenter to be understood.

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 85 SC 85.7.1 P224 L51 # r01-40
Dawe, Piers J G Mellanox Technologies

Comment Type E Comment Status R Bucket

Specifications work at different levels: functional, logic/digital, analog (electrical or optical). "Functional" is the highest/most abstract, while this FFE diagram is part of the specification of an analog quantity. Examples "A functional block diagram of the RS-FEC sublayer is shown in Figure 134-2", "if the 50GMII is not implemented, a conforming implementation must behave functionally as though the RS and 50GMII were present", "PMD functional specifications". I know that several copper clauses say "functional model for the transmit equalizer", and it may be too much effort to correct them all, but this isn't a "transmit equalizer", it's a test fixture, a piece of test equipment. Analog, with specs. Compare 92.7.1, Table 92-4, 92.11.1, 92.11.2.

SuggestedRemedy

Change "The cable assembly test fixture of Figure 85-14 or its functional equivalent, is required" to "The cable assembly test fixture of Figure 85-14 or its equivalent, is required". Similarly, delete "functional" in Table 85-4, 85.8.3.4, 85.8.3.5, 85.10.8

Change the figure title from "TDECQ reference equalizer functional model" to "TDECQ reference equalizer". Similarly in 139.7.5.4 and 140.7.5.1.

Response Response Status C

REJECT.

The last paragraph of the suggested remedy does not pertain to Clause 85 or IEEE P802.3/D3.1 (Clauses 139 and 140 are defined in a separate amendment project). It appears to refer to text in Clause 121 which is [partly] the subject of comment r01-37 and this part of the suggested remedy is considered there.

This remainder of this comment does not apply to the substantive changes between IEEE P802.3/D3.1 and IEEE P802.3/D3.0 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

While the term "functional" may not be necessary, removing it is not a clear improvement to the draft. The term "functional" is not universally used throughout IEEE 802.3 to express the notion of "highest" or "most abstract". Further, this text was approved in 2010 and was unchanged during revisions in 2012 and 2015. There is no clear motivation to make the proposed changes at this point.

[Editor's note added after comment resolution completed.

The response to comment r01-37 includes the following with respect to the use of "functional" model in Clause 121:

"Regarding the "functional model" description, the text and figure follow the precedent set in IEEE Std 802.3bs-2017 Annex 120D for an equivalent type of equalizer."

]

CI 124 SC 124.8.5 P204 L34 # r01-41
Dawe, Piers J G Mellanox Technologies

Comment Type T Comment Status A Bucket

I don't think the reference equalizer as described in 121.8.5.4 is suitable because there, T the symbol period is twice what we need here.

SuggestedRemedy

Add text explaining that the symbol period T is not the same as in 121.8.5.4.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change:

"using a reference equalizer as described in 121.8.5.4, with the following exceptions:" to:

"using a reference equalizer as described in 121.8.5.4 where T is the symbol period for 400GBASE-DR4, with the following exceptions:"

CI 120A SC 120A.3 P345 L27 # r01-42
Dawe, Piers J G Mellanox Technologies

Comment Type T Comment Status D LF_Jitter

As pointed out in both 802.3bs and 802.3cd, a host output with 50 Gb/s lanes is allowed to make twice as much low frequency jitter at very low frequencies as a receiver with 100 Gb/s lane(s) is required to receive. If we don't fix the specs we must warn implementers.

SuggestedRemedy

Add text: e.g. NOTE--The sinusoidal jitter in the 400GAUI-8 module stressed input test represents twice as much, in time or bits, as the sinusoidal jitter in the stressed receiver sensitivity test for the 400GBASE-DR4 PMD.

Also for Figure 120A-6, Example 400GBASE-DR4 PMA layering with single 400GAUI-8 chip-to-module interface

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 120E SC 120E.4.2 P397 L17 # r01-43
Dawe, Piers J G Mellanox Technologies

Comment Type TR Comment Status R LF_Jitter

As pointed out in both 802.3bs and 802.3cd, a host output with 50 Gb/s lanes is allowed to make twice as much low frequency jitter at very low frequencies as a receiver with 100 Gb/s lane(s) is required to receive. A jitter buffer does not fix this unless it is infinite. To assure interoperability, there must be industry-wide agreement that tightens 50G/lane host low frequency jitter generation, increases 100G/lane receiver low frequency jitter tolerance, or a combination: see http://ieee802.org/3/cd/public/Jan18/dawe_3cd_02a_0118.pdf slide 8. The proposed remedy is as simple as any of the options considered. Also it is likely to be compatible with 100G electrical lanes. This remedy should be applied to 400GAUI-8 C2M host outputs (unless another remedy is chosen). It could be applied to 400GAUI-8 host outputs, if it is anticipated that they will ever be connected to 400GBASE-DR4 modules. As any 50G/lane E/O conversions basically pass the low frequency jitter along for something else to tolerate, we can leave their specs alone. 802.3cd may find an alternative solution which could be used instead.

SuggestedRemedy

Add text: To limit the jitter at frequencies which a 400GBASE-DR4 PMD's optical receiver may not track well, it is recommended that for 400GAUI-8, the host output eye width and eye height specifications, and the vertical eye closure specification, be met when measured using a clock recovery unit with a corner frequency of 2 MHz.

Response Response Status C

REJECT.

The same change to a similar interface was proposed in IEEE P802.3cd/D3.0 comment i-115. The response to this comment referred to the response to IEEE P802.3cd/D3.0 comment i-61. The response to this comment was the following.

"REJECT.

Reviewed http://www.ieee802.org/3/cd/public/Jan18/ghiasi_3cd_01_0118.pdf.

Straw poll #1 indicated lack of consensus to make any technical changes to the jitter specification.

Straw poll #1:

I would support making a technical change to the jitter specification.

Y: 4

N: 21

There is no support to make any changes to the jitter specifications."

The poll showed no support to make this change to the IEEE P802.3cd draft. As the same principles apply to this subclause, a similar result is expected here. At this time, no additional information has been provided to support the change.

CI 83 SC 83.5.3.6 P188 L54 # r01-44
Dawe, Piers J G Mellanox Technologies

Comment Type T Comment Status R Bucket

Correct this text to acknowledge that not all PMA interfaces are multi-lane, so not all have Skew Variation, and some Skew values are not as given.

SuggestedRemedy

Change:

If there is a physically instantiated PMD service interface as well, the Skew measured at SP5 is limited to no more than 145 ns of Skew and no more than 3.6 ns of Skew Variation. If there is no physically instantiated PMD service interface, the Skew measured at SP4 is limited to no more than 134 ns of Skew, and no more than 3.4 ns of Skew Variation. to: If there is a physically instantiated PMD service interface that allows the Skew to be measured, the Skew measured at SP5 is limited to no more than 43 ns of Skew for 40GBASE-FR or 145 ns of Skew for a 4-lane PMD, and to no more than 3.6 ns of Skew Variation for a 4-lane PMD. If there is no physically instantiated PMD service interface, the Skew measured at SP4 is limited to no more than 43 ns of Skew for 40GBASE-FR or 134 ns of Skew for a 4-lane PMD, and to no more than 3.4 ns of Skew Variation for a 4-lane PMD.

Response

Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

See response to comment r01-34.

[Editor's note added after comment resolution completed.

The response to comment r01-34 is:

"This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

Also, this text is correct as it is.

It is true that for 40GBASE-FR the Skew at SP5 (on the PMD service interface, at the output of the PMD) is the same as the Skew at SP2 (on the PMD service interface, at the input of the PMD). However, the IEEE P802.3bg project chose to maintain the same limit for Skew at SP5 for this PMD as for all of the other 40GBASE-R PMDs on the basis that the 40GBASE-FR PMD would have no difficulty in meeting this limit and there was no advantage in specifying a lower value for this PMD than for all of the others.

Similarly, although there is no Skew Variation for 40GBASE-FR at SP2, SP3, SP4, or SP5, the IEEE P802.3bg project chose to maintain the same limit for Skew Variation at these skew points for all of the PMDs including 40GBASE-FR for the same reasons. The text in 89.3.2 states that: "there is no Skew Variation at this point" for SP2, SP3, SP4, and SP5 so it is obvious that the limit is met."

]

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 120	SC 120.5.3.5	P102	L8	# r01-45
Dawe, Piers J G		Mellanox Technologies		
Comment Type	E	Comment Status	A	Bucket
Correct the subclause title to reflect the contents (like 120.5.3.3)				
SuggestedRemedy				
Change:				
120.5.3.5 Skew generation at SP6 to:				
120.5.3.5 Skew generation toward SP6				
Response	Response Status C			
ACCEPT IN PRINCIPLE.				
The title of 120.5.3.5 is correct as it is.				
As shown in Figure 116-4 and detailed in the text of 116.5, SP2 is "on the PMD service interface, at the input of the PMD" while SP6 is "on the 200GAUI-n/400GAUI-n interface, at the output of the PMA closest to the 200GBASE-R/400GBASE-R PCS or DTE 200GXS/400GXS".				
Consequently, the title of 120.5.3.3 "Skew generation toward SP2" reflects the fact that SP2 is not at the output of the PMA closest to the PMD, but at the other end of the PMD service interface at the input of the PMD, hence "toward SP2".				
SP6, however, is at the output of the PMA closest to the PCS, so the title of 120.5.3.5 says "at SP6".				
However, it was noted that the text in 120.5.3.5 would be easier to understand if broken into two paragraphs. Start a new paragraph with "If there is a physically instantiated PMD service interface...".				

CI 83	SC 83.5.3.6	P188	L49	# r01-46
Dawe, Piers J G		Mellanox Technologies		
Comment Type	E	Comment Status	A	Bucket
Correct the subclause title to reflect the contents (like 80.5.3.4)				
SuggestedRemedy				
Change:				
80.5.3.6 Skew generation at SP6 to:				
80.5.3.6 Skew generation toward SP6				
Response	Response Status C			
ACCEPT IN PRINCIPLE.				
[Editor's note: in the comment "80.5.3.4" should be "83.5.3.4" and in the suggested remedy, "80.5.3.6" should be "83.5.3.6" (2 instances)]				
This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.				
Also, the title of 83.5.3.6 is correct as it is.				
As shown in Figure 80-6 and detailed in the text of 80.5, SP2 is "on the PMD service interface, at the input of the PMD" while SP6 is "on the XLAUI/CAUI-n interface, at the output of the PMA closest to the PCS".				
Consequently, the title of 83.5.3.4 "Skew generation toward SP2" reflects the fact that SP2 is not at the output of the PMA closest to the PMD, but at the other end of the PMD service interface at the input of the PMD, hence "toward SP2".				
SP6, however, is at the output of the PMA closest to the PCS, so the title of 83.5.3.6 says "at SP6".				
However, it was noted that the text in 83.5.3.6 would be easier to understand if broken into two paragraphs. Start a new paragraph with "If there is a physically instantiated PMD service interface...".				

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 89 SC 89.3.2 P350 L28 # r01-47
Dawe, Piers J G Mellanox Technologies

Comment Type T Comment Status R

The Skew at SP3 (the output of the PMD), SP4 (the receiver MDI) and at SP5 (PMD service interface, output) has to be the same as at SP2 (PMD service interface, input of the PMD) for 40GBASE-FR, a serial PMD. As the receiver can't do anything about it, the "shall"s for SP4 and SP5 are not appropriate. What 802.3ba (all multilane) did can't all be applied to a serial PMD. It's the SP6 spec that can be common to serial and non-serial PMDs, not SP3-5.

SuggestedRemedy

Change:

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns. Since the signal at the MDI represents a serial bit stream, there is no Skew Variation at this point.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns. Since the signal at the MDI represents a serial bit stream, there is no Skew Variation at this point.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns. Since the signal at the PMD service interface represents a serial bit stream, there is no Skew Variation at this point. to: The Skew at SP3 (the transmitter MDI) shall also be less than 43 ns. Since the signal at the MDI represents a serial bit stream, there is no Skew Variation at this point.

The Skew at SP4 (the receiver MDI) and SP5 (the output of the PMD at the PMD service interface) is the same as at SP2, and there is no Skew Variation at these points.

Correct Table 80-6, Summary of Skew constraints, e.g. by inserting columns for 40GBASE-FR, or adding notes to the entries for SP3 SP4 SP5 saying that for 40GBASE-FR, the maximum Skew is as for SP2. Or simply saying that the entries for SP3 to SP5 don't apply to 40GBASE-FR.

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

It is true that for 40GBASE-FR the Skew at SP5 (on the PMD service interface, at the output of the PMD) is the same as the Skew at SP2 (on the PMD service interface, at the input of the PMD). However, the IEEE P802.3bg project chose to maintain the same limit for Skew at SP3 through SP5 for this PMD as for all of the other 40GBASE-R PMDs on the basis that the 40GBASE-FR PMD would have no difficulty in meeting this limit and there was no advantage in specifying a lower value for this PMD than for all of the others.

CI 80 SC 80.5 P102 L4 # r01-48
Dawe, Piers J G Mellanox Technologies

Comment Type E Comment Status R Bucket

Are these column headings correct? They don't seem to apply to a 10-lane 100G PMD.

SuggestedRemedy

?

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

Also, the column headings for Table 80-6 are correct as written.

The approximate skew requirements for 100GBASE-R are given in terms of the unit interval of the PCS lanes, not the the unit interval of the physical lanes. If the table was not done this way, then the numbers would have to be different at each skew point depending on how many PCS lanes are multiplexed together at each point.

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

CI 85 SC 85.8.3.5 P236 L48 # r01-49
Dawe, Piers J G Mellanox Technologies

no clear motivation to make the proposed changes at this point.

Comment Type E Comment Status R Bucket

The test fixture with no name should be given its names, consistent with Clause 92.
Compare:
85.8.3.5 Test fixture
The test setup illustrated in Figure 85-6, or its functional equivalent, is required...
85.10.8 Cable assembly test fixture
The test fixture of Figure 85-14 or its functional equivalent, is required...
92.11.1 TP2 or TP3 test fixture
The test fixture (also known as Host Compliance Board) of Figure 92-15, or its equivalent, is required...
92.11.2 Cable assembly test fixture
The test fixture of Figure 92-17 (also known as Module Compliance Board) or its equivalent, is required...

SuggestedRemedy

Insert "(also known as Module Compliance Board)" in 85.10.8 and "(also known as Host Compliance Board)" in 85.8.3.5. E.g. change
85.8.3.5 Test fixture
The test setup illustrated in Figure 85-6, or its functional equivalent, is required for measuring the transmitter specifications in 85.8.3 at TP2 and the receiver return loss at TP3. TP2 and TP3 are illustrated in Figure 85-2. Figure 85-6 illustrates the test fixture attached to TP2 or TP3.
to
85.8.3.5 TP2 or TP3 test fixture
The test setup illustrated in Figure 85-6, or its functional equivalent, is required for measuring the transmitter specifications in 85.8.3 at TP2 and the receiver return loss at TP3. TP2 and TP3 are illustrated in Figure 85-2. Figure 85-6 illustrates the TP2 or TP3 test fixture (also known as Host Compliance Board) attached to TP2 or TP3.
See another comment for removal of "functional".

Response Response Status C

REJECT.
This comment does not apply to the substantive changes between IEEE 802.3cj/D3.0 and IEEE 802.3cj/D3.1 or the unsatisfied negative comments from the previous ballot. Hence it is not within the scope of the recirculation ballot.

The position of 85.8.3.5 in the subclause heirarchy, as well as the first sentence of the subclause, clarify that it is the "TP2 or TP3 test fixture". The parenthetical references to "host compliance board" and "module compliance board" imply an equivalence but the test fixtures in Clause 85 are defined independently from the compliance boards in Annex 86A (although it is recognized there were efforts to make them compatible). This is in contrast to 83E.4.1 (CAUI-4 chip-to-module) which defines the characteristics of the compliance boards by reference to Clause 92 (and hence a clear equivalence).

Further, the headings and contents of these subclauses were approved in 2010 and were unchanged during revisions in 2012 and 2015 (after the introduction of Clause 92). There is