

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 001 SC 1.4 P 39 L 3 # 1
 Robert Grow RMG Consulting

Comment Type E Comment Status R <OOS>

The inserted definitions do not follow the dictionary sort order to be used for Std 802.3. When consulting with IEEE editors on a previous project, they consider changes to subclass numbering and related editing instructions to be non-substantive. (Specifically, they would be willing to make the changes during publication preparation without need for Sponsor ballot recirculation.) Applying this principal, there is ample precedent for making similar changes during WG ballot or in preparation for Sponsor Ballot. Based on current plan, this project will be an amendment to the now in process revision of IEEE Std 802.3-2015.

Because the sort order for 1.4 both in 802.3-2015 and its approved and anticipated amendments has no consistency, it does not make sense to spend much time on working for this amendment until the inserted definitions are included and the complete set of definitions are resorted to a consistent dictionary sort order.

SuggestedRemedy

When the content of P802.3 has all amendments that will be included in the revision merged into the P802.3 draft, update the editing instructions and subclass numbers in P802.3cd to be based on that revision draft. (Hopefully before the initial Sponsor ballot on P802.3cd.)

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

The commenter correctly points out that P802.3cd must be aligned with the revision project P802.3cj at some point in time.

P802.3cj is still undergoing significant changes especially assuming that P802.3bs will be merged into an upcoming draft.

P802.3cd is likely going to go to sponsor ballot after the November plenary so changes in the next draft should be minimized.

The commenter is encouraged to resubmit this comment at Sponsor ballot.

Cl 136 SC 136.9.3 P 225 L 36 # 2
 Hidaka, Yasuo Fujitsu Lab. of Americ

Comment Type T Comment Status R <OOS>

AC common-mode RMS output voltage (max) is specified as 30mV at TP2 normatively in 136.9.3. It is also specified as 30mV at TP0a in 137.9.2 by a reference to Table 120D-1 normatively for clause 137 and informatively for clause 136 by a reference from 136A.2. There should be some difference in these numbers in order to take account of the mode conversion from differential mode to common mode in signal propagation from TP0 to TP2. In the past clauses, the difference was often 18mV (12mV at TP0a and 30mV at TP2).

SuggestedRemedy

Change AC common-mode RMS output voltage (max) at TP2 in Table 136-11 from 30mV to 48mV.

Or, add the following exception in 137.9.2 as an exception to Table 120D-1: The AC common-mode output voltage (max, RMS) is 12mV.

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

The loss on the host PCB and connector is likely to reduce the AC common mode noise from the host transmitter. But the host PCB may also introduce more AC common mode.

The effect of relaxing this specification on system performance has not been shown.

The commenter may provide supporting data, gain consensus, and resubmit the comment at Sponsor ballot.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 136 SC 136.8.11.1.3 P 212 L 18 # 3
 Hidaka, Yasuo Fujitsu Lab. of Americ

Comment Type E Comment Status A <bucket><oos>
 120.5.11.2.3 describes SSPRQ test pattern.

SuggestedRemedy
 Change 120.5.11.2.3 to 120.5.11.2.1.

Response Response Status C
 ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

However, the suggested change corrects an apparent error and would be an improvement to the draft.

Implement suggested remedy.

CI 136 SC 136.8.11.1.3 P 212 L 1 # 4
 Hidaka, Yasuo Fujitsu Lab. of Americ

Comment Type T Comment Status R
 Although the variable "n" was changed to "p", the statement starting "The polynomials for each identifier value p and" is still difficult to read, because "i" is used as the variable for the identifier in page 211 at line 30, and it is written as "p=i" in page 211 at line 46. It seems that the variable "p" is not necessary, and if we can avoid the variable "p", we can simplify the description. Or, if the variable "p" is necessary, I recommend to rephrase the statement at line 1, page 212.

SuggestedRemedy
 At line 46, page 211, change "for p = i" to "for identifier i".

At line 1, page 212, change "The polynomials for each identifier value p and the default seeds" to "The polynomial and the default seed for each identifier value i".

In Table 136-8, change "p" to "i" in the first column, and change "Polynomial_p" to "Polynomial" in the second column.

Response Response Status C
 REJECT.

The letter "i" is used specifically for the lane number. The letter "p" is used for the identifier value for the polynomial, which is configurable and not necessarily equal to the lane number.

The text in P211 L46 specifies that the default value for the seed for each lane i is given in the table for p=i. However, seed_i may be configured to a different value.

CI 140 SC 140.7.9 P 320 L 26 # 5
 Maki, Jeffery Juniper Networks

Comment Type T Comment Status R jitter corner frequency
 The applied sinusoidal jitter requirements in 121.8.9 are not correct in this case. An additional exception is needed.

SuggestedRemedy
 Add the exception:
 -- 80 kHz is to be used instead of 40 kHz and 8 MHz instead of 4 MHz.

Note that this proposed remedy places the burden solely on this PMD to have the correct requirements and there is no need to change any of the AUJ specifications. This approach is similar to what was done originally for 100GBASE-LR4 with CAUI-10.

Response Response Status C
 REJECT.

See response for comment #7.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl **FM** SC **0** P **1** L **35** # **6**
 Geoff Thompson GraCaSI S.A.

Comment Type **E** Comment Status **R** <00s>

The text on the front page does not describe where this draft is in the process. Instead of: "Draft D2.2 is prepared for Working Group Ballot." it would have been more appropriate to say: "Draft D2.2 is prepared for Working Group 2nd recirculation ballot."

SuggestedRemedy

Use the expanded description in the front page description on all future drafts.

Response Response Status **C**

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

The referenced text is consistent with draft documents for other concurrent IEEE 802.3 amendment projects. The relationship between the draft and the recirculation is clearly communicated in the ballot invitation email.

Cl **135G** SC **135G.3.4** P **375** L **35** # **7**
 Wertheim, Oded Mellanox Technologie

Comment Type **T** Comment Status **R** jitter corner frequency

The jitter specification for the 100G per lane 100GBASE-DR1 receiver uses the same frequency corner as the 50G per lane 100GAUI-2 with the same jitter but with half the peak-to-peak jitter as the jitter mask is defined in UIs. This requires the 100GBASE-DR transceiver PMA to implement a de-jitterizer, which requires to add a PLL to handle the low frequency jitter and a jitter buffer. This adds unnecessary complexity, cost and power to the transceiver.

SuggestedRemedy

Scale the corner frequency for 100GAUI-2 to 2MHz (half the corner frequency of 100GBASE-DR).

The proposed resolution doesn't introduce constraints on future 100G per lane interfaces and provides simpler solution than previous options that were investigated.

1.Add an exception to 135G.3.4 50GAUI-1 C2M and 100GAUI-2 C2M module input characteristics:

With an exception that:

a.The reference CRU for the Module stressed input test and Host stressed input test has a corner frequency of 2MHz

b.The applied sinusoidal jitter values for 100GAUI-2 Module stressed input test and Host stressed input test shall be:

{Jitter frequency, Jitter amplitude}

Case A: {0.02, 5}

Case B: {0.66, 0.15}

Case C: {2, 0.05}

Case D: {6, 0.05}

Case E: {20, 0.05}

2.Add an exception to 135G.4 50GAUI-1 C2M and 100GAUI-2 C2M measurement methodology

With an exception that:

a.The reference CRU for the Eye width and eye height measurement method has a corner frequency of 2MHz

See presentation to follow with additional details.

Response Response Status **C**

REJECT.

The specifications for 100GBASE-DR in P802.3cd are intentionally the same as for 400GBASE-DR4 in P802.3bs.

The potential problem identified in comments #5 and #7 was discussed during the cd Ad Hoc on 25 October 2017 in association with

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

http://www.ieee802.org/3/cd/public/adhoc/archive/dietrich_102517_3cd_adhoc.pdf, explaining the potential problem and proposing some solutions.

Based on that discussion, it is not clear if any changes are warranted or if the proposed changes may result in new problems. A more thorough analysis of the highlighted problem and the impact of the proposed solutions is required.

Further analysis and building of consensus supporting both the highlighted issue and a proposed solution is encouraged to happen.

The comments may be resubmitted in sponsor ballot with any updated information.

Cl 134 **SC 134.5.3.1** **P 158** **L 1** # **8**

Dawe, Piers Mellanox

Comment Type **E** **Comment Status** **A** <bucket><ooos>

After alignment marker lock is achieved the two FEC lanes

SuggestedRemedy
on both FEC lanes (as in 134.5.3.2)? on each FEC lane? on the two FEC lanes?

Response *Response Status* **C**
ACCEPT IN PRINCIPLE.

[The editor changed the page/line from 156/12 to 158/1.]

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

Change:
"After alignment marker lock is achieved the two FEC lanes,"
To:
"After alignment marker lock is achieved on the two FEC lanes,"

Cl 135 **SC 135.1** **P 172** **L 5** # **9**

Dawe, Piers Mellanox

Comment Type **E** **Comment Status** **R** <ooos>

Missing text.

SuggestedRemedy
Add some text in for the overview explaining what this clause is about. Mention all the annexes briefly, in the style of 136.1.

Response *Response Status* **C**
REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

No specific solution is provided.

All of the related normative Annexes are referenced in the final paragraph in 135.1.1.

The organization of this clause is consistent with Clause 120 (PMA for 200GE and 400GE) and 109 (PMA for 25GE).

Cl 135 **SC 135.1** **P 172** **L 20** # **10**

Dawe, Piers Mellanox

Comment Type **E** **Comment Status** **A** <bucket><ooos>

defined in Clause 135B through Clause 135G. ... in Clause 135D through Clause 135G.

SuggestedRemedy
defined in Annex 135B through Annex 135G. ... in Annex 135D through Annex 135G.

Response *Response Status* **C**
ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

However, this change is an improvement to the draft and is editorial in nature.

Implement the suggested remedy.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 135 SC 135.1.3 P 172 L 20 # 11
 Dawe, Piers Mellanox

Comment Type E Comment Status R <oos>

We have added another function, precoding. This isn't the same as Gray mapping, which is part of PAM4 coding - a PMA might do precoding but not PAM4 coding.

SuggestedRemedy

add item k, In some circumstances, perform precoding for PAM4.
 Add full stop to item j.

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

Precoding is part of the processing required for PAM4 modulated signals. It is never required for NRZ modulated signals.

Note that under the subclause 135.5.7 "PAM4 Encoding" includes subclauses 135.5.7.1 "Gray mapping ..." and 135.5.7.2 "Precoding ...".

Cl 135 SC 135.5.7.2 P 184 L 12 # 12
 Dawe, Piers Mellanox

Comment Type T Comment Status R

Because a lane can run through PMAs or PMDs, this text is ambiguous: does an indirect connection count? In the first paragraph we have "PMA lanes connected to" and in the last two paragraphs we have "PMA lanes adjacent to".
 Also, per 120D.1, "The... C2C link is described in terms of a ... C2C transmitter, a ... C2C channel, and a ... C2C receiver." So a PMA lane connected to a C2C link (not part of the link) might be further up or down the chain.

SuggestedRemedy

Change "For PMA lanes connected to a 50GAUI-1 C2C or 100GAUI-2 C2C link, or to the PMD service interface of a 50GBASE-CR, 50GBASE-KR, 100GBASE-CR2, or 100GBASE-KR2 PMD, the PMA shall provide 1/(1+D) mod 4 precoding capability on each output lane and may optionally provide 1/(1+D) mod 4 decoding capability on each input lane."
 to "A PMA shall provide 1/(1+D) mod 4 precoding capability on each output lane that is part of a 50GAUI-1 C2C or 100GAUI-2 C2C transmitter, or is adjacent to the PMD service interface of a 50GBASE-CR, 50GBASE-KR, 100GBASE-CR2, or 100GBASE-KR2 PMD. A PMA may optionally provide 1/(1+D) mod 4 decoding capability on each input lane that is part of a 50GAUI-1 C2C or 100GAUI-2 C2C receiver, or is adjacent to the PMD service interface of a 50GBASE-CR, 50GBASE-KR, 100GBASE-CR2, or 100GBASE-KR2 PMD."
 In the penultimate paragraph, change "For PMA inputs and outputs adjacent to a 50GBASE-CR PMD" to "For PMA lanes adjacent to a 50GBASE-CR PMD".
 In the last paragraphs, change "For PMA lanes adjacent to a 50GAUI-1 C2C" to "For PMA inputs and outputs that are part of a 50GAUI-1 C2C".

Response Response Status C

REJECT.

The text is accurate and is sufficiently clear as written.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 136 SC 136.1 P 197 L 11 # 13
 Dawe, Piers Mellanox

Comment Type E Comment Status A <bucket><oos>

"There are three associated annexes." No, there are four.

SuggestedRemedy

Change three to four. Add sentence for 136C.

Response Response Status C

ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

However, the suggested change corrects an apparent error and would be an improvement to the draft.

Change "three" to "four".

Insert after "Annex 136B specifies test fixtures.":

"Annex 136C specifies MDIs."

Cl 136 SC 136.9.3 P 226 L 10 # 14
 Dawe, Piers Mellanox

Comment Type TR Comment Status R <OOS>

As noted in D2.0 comment 143 and 144, and D2.1 comment 43, these TP2 Jrms and J4u limits, which are copies of the ones in Table 120D-1 (different BER, different test point) should be replaced with Jrms and J3u limits that are consistent (not the same) as the TP0a limits. Crosstalk at the connector combined with the slower edges increases J3u from TP0a to TP2.

SuggestedRemedy

Change J4u to J3u, here and in 137. Choose the limits at TP2 considering the jitter limit at TP0a, the mated compliance board crosstalk specs, and the slower edges at TP2. In 136.9.4.2.3 step e, change J4u to J3u (3 places).

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

Note that the similar comment #43 against D2.1 was rejected with the following response: "REJECT.

The suggested remedy is not specific and cannot be used to apply a change in the draft. More consensus around a specific remedy is required."

The suggested remedy in this new comment is still not specific and cannot be used to apply a change in the draft. More consensus around a specific remedy is required.

See comment #24.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 136 SC 136.9.4.2 P 230 L 26 # 15
 Dawe, Piers Mellanox

Comment Type TR Comment Status R <OOS>

The COM value in the receiver interference tolerance isn't a maximum, it's the reference value that defines what we mean by receiver interference tolerance, and it is used as a target when adjusting the injected noise. See maintenance D2.0 comments 135 and 136.

SuggestedRemedy

In Table 136-13, straddle the "Min" and "Max" columns for the "COM" row and place the contents of the "Max" column into the straddled column. Add the following table footnote to the "COM" parameter label.

"The COM value is the target value for the SNR_TX calibration defined in 136.9.4.2.3 item f). The SNR_TX value measured at the Tx test reference should be as close as practical to the value needed to produce the target COM. If lower SNR_TX values are used, this would demonstrate margin to the specification but this is not required for compliance."

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

However, this change is potentially an improvement, so the commenter is encouraged to resubmit at Sponsor ballot.

Cl 136 SC 136.9.4.2.2 P 230 L 42 # 16
 Dawe, Piers Mellanox

Comment Type T Comment Status R <OOS>

As pointed out in hidaka_3cd_01a_0517.pdf and hidaka_060717_3cd_adhoc-v2.pdf, and D2.0 comment 72, we need a spec for the test channel RL (Rx end) that's better than the regular cable RL spec given by 92.10.3, eq 92-27: 16.5-2rt.f to 4.1 GHz then 10.66-14log10(f/5.5). The comment proposed the mated test fixtures return loss limit, eq 92-38, 20-f to 4 GHz then 18-0.5f. Adopting a limit about half way between these two would be much better than doing nothing. See hidaka_3cd_01a_0517 slides 17/18 to end.

SuggestedRemedy

Insert new requirement into 136.9.4.2.2:

The test channel is the same as the one defined in 110.8.4.2.2, except that the cable assembly meets the requirements of 136.11, the differential return loss of the test channel measured at the Rx test reference (see Figure 110-3b) meets Equation (136-new)."

Eq 136-new to be about half way between eq 92-27 and eq 92-38. 18-f to 4 GHz then 16-0.5f

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

However, there is potential for improvement in this area, so it is encouraged to gain consensus and resubmit at Sponsor ballot.

Cl 136 SC 136.9.4.2.3 P 231 L 3 # 17
 Dawe, Piers Mellanox

Comment Type T Comment Status D <OOS>

It is not likely that the frequency dependent attenuator would have 109.8 ohm impedance.

SuggestedRemedy

If the PCB impedance in 136.11.7.1 (referring to Table 92-12) is not changed (see another comment), add an exception here that the PCB impedance is 100 ohm.

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 136 SC 136.11.7 P 235 L 18 # 18
 Dawe, Piers Mellanox

Comment Type TR Comment Status R <OOS>

The COM impedances should be moved towards neutral, as explained in D2.0 comment 71 and 113.

SuggestedRemedy

Make changes similar to D2.0 comment 71 and hidaka_3cd_01_0717

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

Comment #71 against D2.0 was rejected with the following response:

"REJECT.
 hidaka_3cd_01_0717 was reviewed. There is no consensus to make the proposed changes."

The effect of the proposed changes on system performance has not been shown.

Further supporting data and consensus building are encouraged.

The commenter may resubmit the comment at Sponsor ballot.

It would be helpful for further cable channel data to be provided for use by the task force.

Cl 136 SC 136.11.7.1 P 236 L 39 # 19
 Dawe, Piers Mellanox

Comment Type T Comment Status R <OOS>

Using 109.8 ohm PCB impedance in COM could provide an incentive to build cables to that (wrong) impedance, which seems unhelpful.

SuggestedRemedy

Add another exception to Table 92-12: Zc = 100. In 136.11.7.1.1 and 136.11.7.1.2, delete "and the parameter values given in Table 92-12" (because that is already stated in 136.11.7.1).

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

The text marked for removal in the suggested remedy is intended to draw attention to the different values and prevent misinterpretation (as explained in the editor's notes).

The effect of the proposed changes on system performance has not been shown.

Further supporting data and consensus is encouraged.

The commenter may resubmit the comment at Sponsor ballot.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 136 SC 136.11 P 234 L 5 # 20
 Dawe, Piers Mellanox

Comment Type TR Comment Status D <withdrawn>

This cable loss limit is based on bad reasoning (copying a number from a backplane spec, which is something that doesn't apply to this cable spec). It will be a benefit to the market if a 50GBASE-CR cable can also be a CA-25G-S cable, so the limit should be consistent with 16.48 dB, adjusted for Nyquist frequency. Setting it too high is objective creep (the objective is "copper twin-axial cables with lengths up to at least 3m"), and creates a class of 50GBASE-CR cables that aren't CA-25G-S compliant. I made an estimate of the adjustment and got 16.93 dB. This can be rounded off to 17, which is still significantly more than the 16.06 dB in D1.3.

This comment is a refinement of D2.0 comment 44.

SuggestedRemedy

Change the max loss from 17.16 to 17, to be consistent with CA-25G-S, in Table 136-14, 136.11.2, Table 136A-1 and Figure 136A-1 (two places). Change the RITT losses in Table 136-13 from 15.16 and 17.16 to 15 and 17. In Table 136A-1 and Figure 136A-1, change ILChmax from 30 to 29.84.

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

CI 137 SC 137.9.2 P 251 L 23 # 21
 Dawe, Piers Mellanox

Comment Type T Comment Status R <OOS>

Now that COM is defined with a near-neutral termination and package impedance, we don't expect transmitter return loss to align to the COM model any more. This RL is much tighter than CEI-56G-LR-PAM4 at low (and high) frequency (although apparently looser between 4 and 9 GHz). At low frequencies it is tighter than the channel RL. The effect of (good) RL at low frequency is much less than the less good RL at higher frequencies anyway, and there is less concern about end-to-end reflections than in C2C because the loss is higher when the receiver is challenged. So we can go back to what we had a few drafts ago.

SuggestedRemedy

Insert a new first item in the list of exceptions to Table 120D-1, create a new equation for Tx RL that is similar to the CI.93 and the channel RL at low frequencies; $12 - 0.625f$, $8.7 - 0.075f$. Add figure to illustrate.

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

However, there is potential for improvement in this area, so further analysis and consensus building are encouraged.

The commenter may resubmit the comment at Sponsor ballot.

Note that the similar comment #52 against D2.1 was rejected with the following response: "REJECT.

The presentation dawe_3cd_01a_0917 was reviewed. Further information was requested by the task force on the system implications of the proposed return loss relaxation. There was no consensus to implement the proposed changes."

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 137 SC 137.9.2 P 251 L 28 # 22
 Dawe, Piers Mellanox

Comment Type TR Comment Status R <OOS>

Transmitter output residual ISI SNR_ISI (min) 36.8 dB (Clause 136) and 43 dB (Clause 137) is still 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. D2.0 comment 140, D2.1 comment 49.

SuggestedRemedy

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

There is no suggested remedy however a presentation was submitted.

Viewed and discussed presentation daw_3cd_02_1117.

There is no consensus to make any changes at this time.

Further analysis and consensus building is encouraged.

Cl 137 SC 137.9.2 P 251 L 29 # 23
 Dawe, Piers Mellanox

Comment Type TR Comment Status R <OOS>

Signal-to-noise-and-distortion ratio (min), increased to 33.3 dB (Clause 136) and to 32.5 dB (Clause 137) for all Tx emphasis settings, is still too high. D2.0 comment 139, D2.1 comment 50.

SuggestedRemedy

Response Response Status W

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

There is no suggested remedy however a presentation was submitted.

Viewed and discussed presentation daw_3cd_02_1117. The SNDR was not directly addressed in this presentation. However, SNDR was addressed in a previous presentation http://www.ieee802.org/3/cd/public/Sept17/rysin_3cd_01_0917.pdf.

There is no consensus to make any changes at this time.

Further analysis and consensus building is encouraged.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 137 SC 137.9.2 P 251 L 30 # 24
 Dawe, Piers Mellanox

Comment Type TR Comment Status R <OOS>

This clause with a BER of 2.4e-4 needs a J3u spec, just as 120D with a BER of 1e-5 uses J4u. Using J3u enables a shorter measurement as well as a more relevant, accurate one.

SuggestedRemedy

Add exception 5: the J4u limit in Table 120E-1 does not apply but the maximum J3u is 0.106 UI.

In Eq 136-7 and 136-8 and the NOTE, change J4u to J3u, Q4=3.8906 to Q3=3.2905, Q(Q3) = 5 x10^-4.

Jrms and its value don't change.

If wished, add an informative NOTE in 137.9.2 saying that the J3u limit here is consistent with the J4u limit in Table 120D-1.

Add a new subclause:

136.9.3.n J3u Jitter

J3u is defined similarly to J4u (see 120D.3.1.8). J3u is defined as the time interval that includes all but 10-3 of fJ(t), from the 0.05th to the 99.95th percentile of fJ(t).

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

However, this change is potentially an improvement, so the commenter is encouraged to resubmit at Sponsor ballot.

CI 137 SC 137.9.3 P 251 L 35 # 25
 Dawe, Piers Mellanox

Comment Type TR Comment Status R <OOS>

Now that COM is defined with a near-neutral termination and package impedance, receiver mismatch is the receiver designer's concern, not the standard's, unless it is very extreme, because the receiver interference tolerance test finds its effect combined with other receiver attributes. And we don't expect receiver return loss to align to the COM model any more. This RL is much tighter than CEI-56G-LR-PAM4 at low (and high) frequency (although apparently looser between 4 and 9 GHz). At low frequencies it is tighter than the channel RL. The effect of (good) RL at low frequency is much less than the less good RL at higher frequencies anyway. So we can go back to what we had a few drafts ago.

SuggestedRemedy

Insert a new first item in the list of exceptions to Table 120D-5, create a new equation for Rx RL that is similar to the CI.93 and the channel RL at low frequencies; 12 -0.625f, 8.7-0.075f. Add figure to illustrate or pont to the figure for Tx RL (see another comment).

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

However, there is potential for improvement in this area, so further analysis and consensus building is encouraged.

The commenter may resubmit this comment at Sponsor ballot.

Note that the similar comment #37 against D2.1 was rejected with the following response: "REJECT.

The presentation dawe_3cd_01a_0917 was reviewed. Further information was requested by the task force on the system implications of the proposed return loss relaxation. There was no consensus to implement the proposed changes."

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 138 SC 138.5.7 P 269 L 20 # 26
 Dawe, Piers Mellanox

Comment Type E Comment Status A <bucket><oos>

Function names don't have underscores like this, although functional variable names do. See maintenance D2.0 comments 139, 142, compare 136.8.6.

SuggestedRemedy

Change "PMD_global_transmit_disable function" to "PMD global transmit disable function". Similarly in 139, 140.

Response Response Status C

ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

However, the commenter points out a editorial improvement.

Implement suggested remedy.

CI 138 SC 138.5.7 P 269 L 26 # 27
 Dawe, Piers Mellanox

Comment Type E Comment Status A <bucket><oos>

Improving the language. See maintenance D2.0 comment 140.

SuggestedRemedy

Change "set the PMD_global_transmit_disable to one" to "set the PMD_global_transmit_disable variable to one" Similarly in 139, 140.

Response Response Status C

ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

However, commenter points out an editorial improvement.

Implement suggested remedy.

CI 138 SC 138.5.8 P 269 L 30 # 28
 Dawe, Piers Mellanox

Comment Type E Comment Status A <bucket><oos>

Function names don't have underscores, don't need lane numbers. See maintenance D2.0 comment 141, compare 136.8.7.

SuggestedRemedy

Change: The PMD_transmit_disable_i (where i represents the lane number in the range 0:3) function is optional...

a) When a PMD_transmit_disable_i variable is set...
 If the optional PMD_transmit_disable_i function is not...
 to:

The PMD lane-by-lane transmit disable function is optional...

a) When a PMD_transmit_disable_i variable (where i represents the lane number in the range 0:3) is set...
 If the optional PMD lane-by-lane transmit disable function is not...

Response Response Status C

ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

However, commenter points out an editorial improvement.

Implement suggested remedy.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 135F SC 135F.1 P 367 L 7 # 29
 Dawe, Piers Mellanox

Comment Type T Comment Status R <OOS>

This annex does not refer to Clause 135 at all, nor does it mention precoding for the data path.

SuggestedRemedy

Make reference to 135.
 Here, add sentence saying that a receiver may request precoding and a transmitter should? shall? follow the request.
 In 135F.3.1, say that in addition the C2C transmitter provides a precoding function that can be switched on and off.
 In 135F.3.2, say that in addition the C2C receiver may provide an inverse precoding function.

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

The commenter correctly points out that except for the discussion of transmitter precoding request specification in 135F.3.2.1 it is not stated explicitly that precoding is a configurable option.

The commenter is encouraged to resubmit this comment at Sponsor ballot.

Cl 135G SC 135G.3.1 P 375 L 22 # 30
 Dawe, Piers Mellanox

Comment Type TR Comment Status R <OOS>

As shown in http://ieee802.org/3/bs/public/adhoc/elect/05Oct_17/dawe_01b_100517_elect.pdf there is a need for an additional spec to protect the module from e.g. very noisy hosts, and a max VEC spec provides worthwhile protection.

SuggestedRemedy

Here, add a requirement for VEC, max 12 dB. In 135G.4, add definition of VEC, which was in P802.3bs D2.0 120E.4.2.1 (the AVs were illustrated in Figure 120E-13, although they could be on Fig 120E-14 and the text under what was equation 120E-3 is clear enough so we don't have to add them to the figure).
 Add PICS to 135G.5.4.1.

Response Response Status C

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3cd D2.1 and D2.2 or the unsatisfied negative comments from the initial ballot. Hence it is not within the scope of the recirculation ballot.

Viewed dawe_3cd_01_1117. There was general agreement with the proposal.

A presentation providing detailed implementation is encouraged.

The commenter is encouraged to resubmit this comment at Sponsor ballot.

Cl 136C SC 136C.1 P 387 L 41 # 31
 Dawe, Piers Mellanox

Comment Type T Comment Status R

The paragraph about AC coupling, which should be a property of and requirement on the cable not the MDI, is in the wrong place. The subclause reference in PICS CA9 is wrong.

SuggestedRemedy

Move this paragraph to 136.11 just before 136.11.1 (older clauses have it in the equivalent of 136.12, which is not really correct but at least it's in the clause). Update the subclause reference in PICS CA9.

Response Response Status C

REJECT.

The referenced text is correct as written.

However, it may be helpful for the reader if the text is moved to a more appropriate location.

The commenter is invited to resubmit the comment in sponsor ballot.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 136 SC 136.11 P 223 L 42 # 20113
 Dudek, Mike Cavium

Comment Type TR Comment Status A Cable

Equation 92-27 for the differential return loss gives 5.3dB return loss at 13.28GHz. This is not the 6dB listed and is a relatively poor value and could lead to significant differences between system performance with a real host and the COM calculated with the single 110 Ohm host board trace equivalent. Work on backplanes and C2C (e.g. Hidaka_3cd_01a_0317, Dudek_3bs_02_0517) has shown that this affect is significant and it would be better to test COM with nominal impedances and have a guard band between the channel COM and the Interference tolerance COM.

SuggestedRemedy

Change 6 to 5.3 Change the COM value to 3.5dB. In table 136-15 change the value of Rd to 50 Ohm, the value of Zc to 95 Ohm, On page 224 line 40 change the value of COM to 3.5dB. Change the impedance of the test trace from TP0 to TP1 and TP4 to TP5 to 100 Ohm by changing on page 226 line 41 from "using zp = 151 mm in length, representing an insertion loss of 6.42 dB at 13.28 GHz on each PCB." to "using Zc = 100 Ohm and zp = 151 mm in length, representing an insertion loss of 6.42 dB at 13.28 GHz on each PCB." Also change to 3.5dB in PICS CA8.

Response Response Status C

ACCEPT IN PRINCIPLE.

[Editor's note: This D2.0 comment was unsatisfied.]

In Table 136-14, Change "Minimum differential return loss at 13.28 GHz" from 6 dB to 5.3 dB.

The rest of the suggested remedy requires more consensus building.

See also #71

Cl 138 SC 138.7.1 P 262 L 18 # 20127
 Dawe, Piers Mellanox

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 (up to 4/2 dB worse than the SRS test?) With some of the changed low-bandwidth TDECQ being used to equalize the reference receiver's own bandwidth, this issue becomes more apparent. This is an issue for all the PAM4 optical PMDs, although it may be worse for MMF because of the high TDECQ limit.

SuggestedRemedy

Define TDECQrms = $10 \cdot \log_{10}(A_{RMS}/(s^3 \cdot Q_t \cdot R))$ where A_RMS is the standard deviation of the measured signal after the 13.28125 GHz filter response. s is the standard deviation of a fast clean signal with OMA=2 and without emphasis, observed through the 13.28125 GHz filter response (around 0.7 - can be calculated when the filter bandwidth is stable). Set limit for TDECQrms according to what level of dirty-but-emphasised signal we decide is acceptable, add max TDECQrms row to the table. Alternatively, if the same relative limit is acceptable for all PAM4 optical PMDs, the limit could be in the TDECQ procedure 121.8.5.3 as proposed in P802.3bs D3.2 comment r02-35. Similarly in clauses 139, 140.

Response Response Status C

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

A similar comment was made to P802.3bs D3.2 via comment r02-35, which was rejected.

Insufficient evidence of the claimed problem and that the proposed remedy fixes the problem.

A contribution is invited that demonstrates the problem (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that the proposed additional requirement prevents this issue from occurring.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 140 SC 140.6.1 P 306 L 33 # 20128
 Dawe, Piers Mellanox

Comment Type TR Comment Status R

PAM4 optics is still new and raw, we are still debugging the specification methodology, and we have seen too little experimental information showing technical and economic feasibility. As measurements with the new TDECQ method and with new receiver designs become available, it may be that optical power levels can be reduced and the spec as in this draft would be uneconomic.

SuggestedRemedy

Bring more evidence for what optical power levels and TDECQ limits are right; in particular, TDECQ measurements with SSPRQ, and correlation to actual receiver performance. Based on evidence, reduce all the optical power levels for 100GBASE-DR by 0.5 or 1 dB (with other adjustments for other reasons). Review the TDECQ limit.

Response Response Status C

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

The suggested remedy does not propose a specific change to the draft.

Cl 139 SC 139.7.7 P 289 L 15 # 20133
 Dawe, Piers Mellanox

Comment Type TR Comment Status R

With the lower receiver bandwidth, measuring RIN in approximately the signaling rate (twice as much) seems too much; 1/2 to 3/4 would be better. A T-spaced equalizer cannot independently adjust for good ISI and RIN filtering, so can an adequate estimate of RIN can be obtained as a by-product of the TDECQ procedure? While a T/2-spaced equalizer could enhance the RIN, it would not choose to do so if RIN were a problem, so a T-spaced reference equalizer and a T/2-spaced product equalizer are compatible from this point of view, I think. As 52.9.6 says, this RIN method is intended for components (TOSAs) not a "system level test" suitable for a complete optical module. This is much the same as P802.3bs D3.2 comment r02-39.

SuggestedRemedy

Review; reduce the bandwidth and simplify RIN measurement to a Qsq measurement (see 68.6.7) or eliminate as appropriate. Remove 135.5.10.2.4 Square wave (quaternary) test pattern and any associated registers. Similarly in 140.7.9.

Response Response Status C

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

The suggested remedy suggests 2 different approaches to change the draft. Changing the RIN measurement to a Qsq measurement has not been demonstrated to provide the same safeguards that are expected from the RIN requirement. Eliminating the RIN measurement was discussed in the response to comment #130 against D2.0 of P802.3bs on the basis that "The transmitter RINxOMA spec is intended to screen out potentially bad transmitters even if the noise correction required by the TDECQ test is not very accurate."

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 140 SC 140.7.9 P 310 L 28 # 20134
 Dawe, Piers Mellanox

Comment Type TR Comment Status R jitter

The lack of consistency between the low frequency jitter specs in 802.3bs affects 802.3cd also. Here is P802.3bs D3.2 comment r02-40 for those who have not been following this issue. Depending how this inconsistency is fixed, there may be little or no explicit change in the P802.3cd draft.

Following up on P802.3bs D3.0 comment 153 and D3.1 comment 55: if the jitter corner frequency for 26.5625 GBd (NRZ and PAM4) is 4 MHz, the low frequency ends of the jitter masks must align or be in the right order if expressed in time vs. frequency, i.e. should scale with signalling rate if in UI. If this is not done, the required depth of the LF jitter buffer in the 2:1 muxes in a 400GBASE-DR4 module is unbounded and the low frequency jitter generation requirements on the module become unreasonable. Compare 87.8.11.4 and 88.8.10: 4 MHz for 10.3125 GBd, 10 MHz for 25.78125 GBd. History: anslow_3bs_04_0316 does not contain reasoning, refers to ghiasi_3bs_01_0316 which does not address wander and buffering. ghiasi_3bs_01a_0116.pdf#page=15 shows FIFOs but does not establish a workable spec. Slide 14 shows they can be avoided: this is what we have for 400GAUI-8 or 400GAUI-16 with 400GBASE-xR8. I have no evidence that the problems described in the [fourth] sentence have been considered or solved by the [P802.3bs] committee.

SuggestedRemedy

Add another exception for the SRS procedure, with a table like Table 121-12 replacing second row after the header row:

80 kHz < f <= 250 kHz 4e5/f
 250 kHz < f <= 500 kHz 1e11/f^2
 1 MHz < f <= 4 MHz 2e5/f

Or, with the UIs doubled vs. Table 121-12:

f < 40 kHz Not specified
 40 kHz < f <= 4 MHz 4e5/f
 4 MHz < f <= 10 LB 0.1

Increase the TDECQ limit to share the burden appropriately between transmitter and receiver.

This option means the 100G/lane receiver has to tolerate no more timing slew rate (in ps/us) than that agreed for 50G/lanes.

Or, increase jitter by 50% and corner frequency by 33%:

f < 40 kHz Not specified
 40 kHz < f <= 6 MHz 4e5/f
 5.333 MHz < f <= 10 LB 0.075

and add an exception in 124.8.5 that the CRU corner frequency is 5.333 MHz. Increase the TDECQ limit to share the burden between transmitter and receiver.

To do the job properly with the first option, in 124.8.5 we should add another exception to the CRU with a corner frequency of 4 MHz and a slope of 20 dB/decade (in 121.8.5.1): add a pole at 250 kHz and a zero at 500 kHz. I am advised that this can be done in hardware (in software, anything is possible).

Response Response Status C

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

One option in the suggested remedy is proposing to place an extra burden on the receiver by allowing transmitters with a higher level of TDECQ which may be due to ISI and also by requiring a higher level of jitter tolerance.

The commenter has not demonstrated that this extra burden is less onerous than putting a buffer in the PMA.

For the second option in the suggested remedy the commenter is invited to build consensus for an increase of the corner frequency to be above 4 MHz.

CI 137 SC 137.9 P 241 L 1 # 20136
 Dawe, Piers Mellanox

Comment Type TR Comment Status R Electrical <NSR>

We don't yet know how to write a spec for 30 dB channels that isn't bleeding edge for ICs and/or channels. This isn't Ethernet "broad market" today, it's a specialist niche.

SuggestedRemedy

Keep working on it in Working Group ballot and if things don't improve, reduce the 30 dB objective and reduce the high loss RITT loss. It might be OK to leave the channel recommended insertion loss limit if the COM spec protects the Tx and Rx.

Response Response Status C

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

No specific change to the draft is suggested.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 137 SC 137.9.2 P 241 L 22 # 20139
 Dawe, Piers Mellanox

Comment Type **TR** Comment Status **R** Electrical <NSR>

Signal-to-noise-and-distortion ratio (min) 32.5 dB is too high (even worse than 120D) - probably can't measure the IC through the test fixture and cables. I suspect there is double counting of jitter in SNDR and as jitter, in COM.

SuggestedRemedy

Remove the double counting. Reduce the SNDR limit to something that can reasonably be measured, or change the measurement method.

Response Response Status **C**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

dawe_3cd_02_0717 was presented.

The comment highlights some issues in the current draft, but there was no consensus for adopting any of the proposed solutions.

The commenter is encouraged to build consensus and bring a new proposal.

CI 137 SC 137.9.2 P 241 L 21 # 20140
 Dawe, Piers Mellanox

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

Output residual ISI SNR_ISI (min) 43 dB is way too high - probably can't measure the IC through the test fixture and cables, even test equipment fails this limit. The warning NOTE in 120D.3.1.7 notes the issue (for 34.8 dB), but doesn't solve it.

SuggestedRemedy

It may be necessary to move away from the SNR_ISI method.

Response Response Status **C**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

dawe_3cd_02_0717 was presented.

The comment highlights an issue in the current draft, but there was no consensus for adopting any of the proposed solutions.

The commenter is encouraged to build consensus and bring a new proposal.

See #139.

CI 137 SC 137.9.3.1 P 241 L 46 # 20141
 Dawe, Piers Mellanox

Comment Type **TR** Comment Status **A** Electrical

The low frequency RL at 14.25 dB is insignificant for signal integrity compared with the 8.7 dB at 6 GHz. This RL is much tighter than CEI-56G-LR at low (and high) frequency (although apparently looser between 4 and 9 GHz).

SuggestedRemedy

Change 14.25 - f to 12 -0.625f

Response Response Status **C**

ACCEPT IN PRINCIPLE.

[Editor's note: This D2.0 comment was unsatisfied.]

This issue was discussed in 802.3bs and resulted in a change to the similar specification (Comment #r02-41).

In 137.9.3.1 (Receiver input return loss), append the following text to the first paragraph: "The test fixture return loss may be de-embedded from the return loss measurements."

No need to add this in 137.9.2 (Transmitter characteristics) since it refers to 120D.3.1.1, where a similar change was applied by 802.3bs (indirectly through Table 120D-1). Update exceptions if necessary.

Implement with editorial license.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 136 SC 136.9.3 P 216 L 11 # 20143
 Dawe, Piers Mellanox

Comment Type TR Comment Status R Electrical <NSR>

J4, now called J4u (all but 1e-4 of the edges, or 1e-4*0.75 of the number of UI, divided between early and late, so 3.75e-5 per UI or 1.875e-5 per bit) is overkill for the spec BER of 2.4e-4, and J3u (1.875e-4 per bit) is a good match to the spec BER - just as J4u is a good match to the BER of 1e-5 for 120D. Also, not all edges cause errors. We can make the spec better (more accurate, less performance left on the table) and reduce test time. Futher, the jitter at TP2 won't be the same as at TP0a in 137.9.2 (expected to be more).

SuggestedRemedy

Change J4 to J3u. Choose the limit at TP2 considering jitter limit at TP0a and the mated compliance board crosstalk specs, among other factors.

Response Response Status C

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

The suggested remedy lacks sufficient detail required for implementation - the limits for TP2 are not included.

The commenter is encouraged to suggest and build consensus for specific limits at TP2, as well as the suggestion to change J4u to J3u.

Cl 137 SC 137.9.2 P 241 L 24 # 20144
 Dawe, Piers Mellanox

Comment Type TR Comment Status R Electrical

J4u in 120D (all but 1e-4 of the edges, or 1e-4*0.75 of the number of UI, divided between early and late, so 3.75e-5 per UI or 1.875e-5 per bit) is overkill for the spec BER of 2.4e-4, and J3u (1.875e-4 per bit) is a good match to the spec BER - just as J4u is a good match to the BER of 1e-5 for 120D. Also, not all edges cause errors. We can make the spec better (more accurate, less performance left on the table) and reduce test time.

SuggestedRemedy

Change J4 to J3u, max 0.106 UI (from eq 136-6 and 7). In Eq 136-6 and 136-7 and the NOTE, change Q4=3.8906 to Q3=3.2905, Q(Q3) = 5 x10^-4.

Response Response Status C

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

Note that the suggested change (J4u to J3u) seems to enable a shorter measurement while keeping the same sigma_RJ and A_DD for COM, by changing the conversion equations (136-6 and 136-7).

The task force discussed the suggested remedy. Since currently both clauses 136 and 137 use the same equations, there is preference to make changes to both clauses together. There is no consensus for changing just this clause.

See comment #143.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 138 SC 138.7.1 P 262 L 17 # 20147
 Dawe, Piers Mellanox

Comment Type TR Comment Status R

This PMD needs more study, and knowing what TDECQ is feasible is probably the key.

SuggestedRemedy

While in WG ballot, show evidence of technical feasibility for the numbers in the spec: eyes, receiver waterfall plots, TDECQ measurements and so on. Adjust the draft as appropriate. TR because this could take a few meeting cycles.

Response Response Status U

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

[Editors note: This comment is a repeat of comment 42 against draft 1.3]

No specific changes to the draft suggested.

Task force participants are encouraged to prepare consensus presentations with proposals for specific changes to the draft if necessary.

CI 139 SC 139.6.1 P 283 L 36 # 20152
 Dawe, Piers Mellanox

Comment Type TR Comment Status R power budget

PAM4 optics is still new and raw, we are still debugging the specification methodology, and we have seen too little experimental information showing technical and economic feasibility. However, stassar_061417_3cd_adhoc-v2 shows plenty of receiver sensitivity margin (although not yet shown with SSPRQ). As more measurements with new receiver designs and the new TDECQ method become available, it appears the optical power levels can be reduced and the spec as in this draft will be uneconomic (particularly 50GBASE-FR which should be low cost, low power, convenient for quad or octal packaging).

SuggestedRemedy

Bring more evidence for what optical power levels and TDECQ limits are right, including TDECQ measurements with SSPRQ, and correlation to actual receiver performance. Based on evidence, reduce all the optical power levels for 50GBASE-FR and 50GBASE-LR by 0.5, 1 or 1.5 dB (with other adjustments for other reasons). Review the TDECQ limit.

Response Response Status C

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

The suggested remedy does not propose a specific change to the draft.

CI 137 SC 137.9.3.1 P 250 L 1 # 21037
 Dawe, Piers Mellanox

Comment Type TR Comment Status R Return loss

Now that COM is defined with a near-neutral termination and package impedance, receiver mismatch is the receiver designer's concern, not the standard's, unless it is very extreme, because the receiver interference tolerance test finds its effect combined with other receiver attributes. And we don't expect transmitter return loss to align to the COM model any more. This RL is much tighter than CEI-56G-LR at low (and high) frequency (although apparently looser between 4 and 9 GHz). At low frequencies it is tighter than the channel RL. The effect of (good) RL at low frequency is much less than the less good RL at higher frequencies anyway. So we can go back to what we had a few drafts ago.

SuggestedRemedy

Change "shall meet Equation (137-1)" to "shall meet Equation (93-3)" and delete Eq 137-1 and Fig 137-3. Or, change 14.25 - f to 12 -0.625f, revise the figure.

Response Response Status C

REJECT.

[Editor's note: This D2.1 comment was unsatisfied.]

The presentation daw_3cd_01a_0917 was reviewed. Further information was requested by the task force on the system implications of the proposed return loss relaxation.

There was no consensus to implement the proposed changes.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 138 SC 138.7.1 P 270 L 10 # 21038
 Dawe, Piers Mellanox

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. With some of the changed low-bandwidth TDECQ being used to equalize the reference receiver's own bandwidth, this issue becomes more apparent. 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.

SuggestedRemedy

1. To screen for noisy or distorted signals with heavy emphasis
 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 filter response, Q_t and R are as already in Eq 212-12. s is the standard deviation of a fast clean signal with OMA=2 and without emphasis, observed through the 13.28125 GHz filter response (around 0.7 - can be calculated when the filter bandwidth is stable). Set limit for $TDECQ_{rms}$ according to what level of dirty-but-emphasised signal we decide is acceptable, add max $TDECQ_{rms}$ row to the table. Alternatively, if the same relative limit is acceptable for all PAM4 optical PMDs, the limit could be in the TDECQ procedure 121.8.5.3 as proposed in bs comment(s). Similarly in clauses 139, 140.
2. To protect the TIA input, consider a peak power spec as in Clause 86.
3. To protect the TIA and any AGC and TIA from unreasonable signals, consider a crest factor spec.
4. To protect the equalizer from having to support unnecessary settings, require that the cursor is one of the first three taps.
5. To protect the receiver from having to "invert" heavily over-emphasised signals, set a minimum cursor weight.

Response Response Status U

REJECT.

[Editor's note: This D2.1 comment was unsatisfied.]

This comment is related to unsatisfied comments i-140 and r02-35 against 802.3bs draft 3.2.

The resolution to P802.3bs comment r02-35 was:

"REJECT

Insufficient evidence of the claimed problem and that the proposed remedy fixes the problem. The commenter is invited to provide a contribution that demonstrates the problem (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that the proposed additional requirement prevents this issue from occurring."

Insufficient evidence was provided of the claimed problem and that the suggested remedy

fixes the problem. A contribution is invited that demonstrates the problem (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that the proposed additional requirements prevent this issue from occurring.

CI 139 SC 139.6.1 P 291 L 36 # 21040
 Dawe, Piers Mellanox

Comment Type TR Comment Status R

The discussion around D2.0 comment 152 implied that there is receiver margin to spare in 50GBASE-FR.

SuggestedRemedy

reduce all the optical power levels for 50GBASE-FR (except Rx damage) by 1 dB. Bring more evidence for what optical power levels and TDECQ limits are right, including TDECQ measurements with SSPRQ, and correlation to actual receiver performance. Review the TDECQ limit.

Response Response Status U

REJECT.

[Editor's note: This D2.1 comment was unsatisfied.]

This comment is a follow up comment to comment #152 to D2.0.

The current values are based on the adoption of a baseline proposal in http://www.ieee802.org/3/cd/public/May16/cole_3cd_01_0516.pdf during the May 2016 meeting in Whistler by a motion with the following results. Y: 54 N: 0 A: 25.

It is known that there are margins in both transmitter and receiver specifications when the baseline proposal was adopted.

No analysis has been provided that changing the current values by 1 dB would enable lower cost solutions and/or better performance.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 140 SC 140.6.1 P 314 L 33 # 21042
 Dawe, Piers Mellanox

Comment Type TR Comment Status R

D2.0 comment 128: PAM4 optics is still new and raw, we are still debugging the specification methodology, and we have seen too little experimental information showing technical and economic feasibility. As measurements with the new TDECQ method and with new receiver designs become available, it may be that optical power levels can be reduced and the spec as in this draft would be uneconomic.

SuggestedRemedy

Reduce all the optical power levels for 100GBASE-DR by 0.5 dB.
 Bring more evidence for what optical power levels and TDECQ limits are right; in particular, TDECQ measurements with SSPRQ, and correlation to actual receiver performance.
 Review the TDECQ limit.

Response Response Status U

REJECT.

[Editor's note: This D2.1 comment was unsatisfied.]

No analysis has been provided that changing the current values by 0.5 dB would enable lower cost solutions and/or better performance.

Furthermore the existing values for 100GBASE-DR are intentionally consistent with the values for one lane in 400GBASE-DR4 in P802.3bs.

A presentation (dawe_3bs_03_0917) containing similar proposals pertaining to 400GBASE-DR4 in P802.3bs D3.3 was not accepted.

CI 136 SC 136.9.3 P 224 L 10 # 21043
 Dawe, Piers Mellanox

Comment Type TR Comment Status R <NSR>

As explained before, J4u should be changed to J3u. The equivalent J3u is known (D2.0 comment 144) but we need an estimate of the difference in jitter between TP0a and TP2 so that we can choose more appropriate limits for the two test points (D2.0 comment 143).

SuggestedRemedy

Change J4u to J3u, here and in 137. Choose the limit at TP2 considering jitter limit at TP0a, the mated compliance board crosstalk specs, and the slower edges at TP2.

Response Response Status C

REJECT.

[Editor's note: This D2.1 comment was unsatisfied.]

The suggested remedy is not specific and cannot be used to apply a change in the draft.

More consensus around a specific remedy is required.

CI 136 SC 136.11.2 P 232 L 28 # 21044
 Dawe, Piers Mellanox

Comment Type TR Comment Status R Cable assembly

Where did 17.16 dB come from? the limit should be consistent with other 3 m cables: 16.48 or 15.5 dB (CA-25G-S CA-25G-N), adjusted for Nyquist frequency. Setting it too high is objective creep.

SuggestedRemedy

Set the max loss to be no more than consistent with CA-25G-S. Set the RITT losses accordingly.

Response Response Status C

REJECT.

[Editor's note: This D2.1 comment was unsatisfied.]

The value 17.16 dB is included in the resolution of comment #124 against D2.0, based on palkert_3cd_01b_0717 and the task force discussion following the presentation. This number makes the channel IL the same as for Clause 137.

No further changes are required to close the budget.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 136 SC 136.11.7 P 233 L 18 # 21047
 Dawe, Piers Mellanox

Comment Type TR Comment Status R Cable assembly

The COM impedances should be moved towards neutral, as explained in D2.0 comment 71 and 113.

SuggestedRemedy

Make changes proposed in D2.0 comment 71 and hidaka_3cd_01_0717 - except don't change the parameter name unless it is coordinated with the name used in Annex 93A.

Response Response Status C

REJECT.

[Editor's note: This D2.1 comment was unsatisfied.]

Comment #71 against D2.0 suggested changing COM parameters to use well-matched impedances: terminations of 50 Ohm, package impedance of 95 Ohm and board impedance of 100 Ohms.

D2.0 comment #71 was rejected due to lack of consensus.

The related changes suggested in D2.0 comment #113 were also not in consensus.

The comment does not provide any new information, nor address any concerns that prevented the prior comments from being adopted.

CI 137 SC 137.9.2 P 249 L 28 # 21049
 Dawe, Piers Mellanox

Comment Type TR Comment Status R Tx specs

Transmitter output residual ISI SNR_ISI (min) 36.8 dB (Clause 136) and 43 dB (Clause 137) is still too high see daw_3bs_04_0717 and daw_3cd_02a_0717 - 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. D2.0 comment 140

SuggestedRemedy

See presentation.

Response Response Status C

REJECT.

[Editor's note: This D2.1 comment was unsatisfied.]

The task force reviewed rysin_3cd_02_0917.

There was no consensus for implementing the proposed changes.

CI 137 SC 137.9.2 P 249 L 29 # 21050
 Dawe, Piers Mellanox

Comment Type TR Comment Status A Tx specs

Signal-to-noise-and-distortion ratio (min), increased to 33.3 dB (Clause 136) and to 32.5 dB (Clause 137) for all Tx emphasis settings, is too high: see daw_3bs_04_0717 and daw_3cd_02a_0717 - can barely measure the IC through the test fixture. It seems SNDR depends on emphasis, while COM assumes the spec limit at all emphasis settings which is pessimistic and not realistic. Also I suspect there is double counting of jitter in SNDR and as jitter, in COM. D2.0 comment 139.

SuggestedRemedy

Apply a SNDR limit that accounts for the way Pmax varies with emphasis: SNDR0+20log10(Pmax_equalized/Pmax_unequalized), or apply the SNDR spec for no emphasis only.

Response Response Status C

ACCEPT IN PRINCIPLE.

[Editor's note: This D2.1 comment was unsatisfied.]

The response to comment #139 against D2.0 was:

"REJECT.

daw_3cd_02_0717 was presented.

The comment highlights some issues in the current draft, but there was no consensus for adopting any of the proposed solutions.

The commenter is encouraged to build consensus and bring a new proposal."

The suggested remedy is a new proposal.

The commenter points out an issue and proposed solution that need further investigation.

There is no consensus to make the proposed changes.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 137 SC 137.9.2 P 249 L 29 # 21051
 Dawe, Piers Mellanox

Comment Type TR Comment Status R Tx specs

COM SNR_TX is defined at the TX output. SNDR is measured thru package and TF by real (imperfect) test equipment therefore is lower than SNR_TX, causing some double counting in COM. D2.0 comment 139.

SuggestedRemedy

Reduce the SNDR specification to 29 dB for both Clause 136 and 137 to account for the degradation caused by the package and test fixture as well as by the measurement impairments.

Response Response Status C

REJECT.

[Editor's note: This D2.1 comment was unsatisfied.]

The task force reviewed rysin_3cd_01_0917.

The package and test fixture effects are linear, so are effectively de-embedded in the linear fit procedure.

The claim that measured SNDR is lower than "real" SNDR is not substantiated.

Creating a difference of 3.5 dB between the COM parameter (SNR_TX) and the corresponding TX parameter (SNDR) would break the budget. Bad transmitters may pass the Tx specs but cause their partner's receiver to fail.

There is no consensus to make the proposed changes.

[Editor's note: The following response was added at the November 2017 Task Force meeting.]

It may be possible to improve the SNDR measurement technique to improve correlation with system performance.

Cl 137 SC 137.9.2 P 249 L 30 # 21052
 Dawe, Piers Mellanox

Comment Type TR Comment Status R Return loss

Now that COM is defined with a near-neutral termination and package impedance, we don't expect transmitter return loss to align to the COM model any more. This RL is much tighter than CEI-56G-LR at low (and high) frequency (although apparently looser between 4 and 9 GHz). At low frequencies it is tighter than the channel RL. The effect of (good) RL at low frequency is much less than the less good RL at higher frequencies anyway, and there is less concern about end-to-end reflections than in C2C because the loss is higher when the receiver is challenged. So we can go back to what we had a few drafts ago.

SuggestedRemedy

If bs doesn't fix this, add another exception and create new equation for Tx RL that is similar to the Cl.93 and the channel RL at low frequencies; 12 -0.625f, 8.7-0.075f. Add figure to illustrate. Refer to new equation instead of existing 137-1. If 137-1 is revised as above for the receiver, can continue to point to it.

Response Response Status C

REJECT.

[Editor's note: This D2.1 comment was unsatisfied.]

The presentation daw_3cd_01a_0917 was reviewed. Further information was requested by the task force on the system implications of the proposed return loss relaxation.

There was no consensus to implement the proposed changes.