

IEEE P802.3 (IEEE 802.3cj) D3.0 Maintenance #12 (Revision) Initial Sponsor ballot comments

Cl 1 SC 1.4.281 P92 L4 # i-41
Nikolich, Paul INDEPENDENT

Comment Type TR Comment Status A

The current definition of 'lane' requires improvement.
Current definition: 1.4.281 lane: A bundle of signals that constitutes a logical subset of a point-to-point interconnect. A lane contains enough signals to communicate a quantum of data and/or control information between the two endpoints.

For example "bundle" is defined as a "group of signals", which is duplicated in "bundle of signals" above.
Per the definition of "bundle", it should be "A bundle that constitutes..."

Where is "quantum of data" defined? I couldn't find it.

Where is "endpoint" defined?

Unfortunately I don't have a good alternative definition.

SuggestedRemedy

Look through the draft and identify the various ways "lane" is used, then develop an appropriate single definition. If a single definition is not feasible, perhaps more than one definition is needed.

Response Response Status U

ACCEPT IN PRINCIPLE.
Replace the definition of "lane" with the following.

"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."

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

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

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.

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.

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.

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

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 quite 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.

Cl 121 SC 121.8.5.3 P132 L1 # 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 212-12. 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 212-12" 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

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

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."

]

Cl **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.

Cl **121** SC **121.8.5.4** P**135** L**18** # **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://ieeee802.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."

]

IEEE P802.3 (IEEE 802.3cj) D3.2 Maintenance #12 (Revision) 2nd Sponsor recirculation ballot comments

Cl 1 SC 1.4.300 P92 L 94 # r02-2
 Nikolich, Paul INDEPENDENT

Comment Type GR Comment Status R

Firstly, I disagree with the rejection of my earlier comments, as I believe the definition of "lane" should provide greater clarity, accuracy and precision.

Secondly, The term "lane" is used in the standard that is not consistent with the proposed definition. For example, later on in the definitions section the following definition is offered: 1.4.386 PCS lane (PCSL): In 40GBASE-R, 100GBASE-R, 200GBASE-R, and 400GBASE-R, the PCS distributes encoded data to multiple logical lanes, these logical lanes are called PCS lanes. One or more PCS lanes can be multiplexed and carried on a physical lane together at the PMA service interface. (See IEEE Std 802.3, Clause 83 and Clause 120.)

Note the use of the qualifiers "logical lane" and "physical lane". This implies there are at least two types of "lane", while the proposed definition appears to address "logical lane" and not "physical lane". At a minimum a definition for "physical lane" should be added to the standard.

SuggestedRemedy

- 1) Change the label on 1.4.300 to Logical Lane.
- 2) Add a definition for a Physical Lane.
- 3) Add illustrations to (1) and (2) above to improve the ability of a reader to correctly understand the definitions similar to what is used in 802.16-2017 definition of "protocol data unit" Figure 3-1

Response Response Status U

REJECT.

The definition of "lane" in 1.4.300 is correct for "logical lane", "physical lane", and "PCS lane". It is generic and addresses abstract/logical transfers of data "from one sublayer to an adjacent sublayer" and physical transfers of data across "the transmission medium (e.g., optical fiber, optical wavelength, wire pair)". The phrase "logical subset of the data and control information" does not limit the definition to "logical lanes" as physical lanes also convey "logical subsets" of the data.

As the definition of "PCS lane (PCSL)" states, it is a specific construct used "in 40GBASE-R, 100GBASE-R, 200GBASE-R, and 400GBASE-R" and it is fully consistent with the definition of lane in 1.4.300. It is a further qualification of the specific usage of lanes for those PHY families and introduces the term "physical lane" to distinguish bit-multiplexed PCS lanes from the PCS lanes themselves. The references to Clauses 83 and 120 can be followed for further details on these constructions.

Other PHYs specifications use "lanes" that are consistent with the definition 1.4.300 but are not handled (e.g., multiplexed) in the same way that "PCS lanes" as defined in 1.4.386 may be.

Therefore, it is too limiting to change the label of 1.4.300 from "lane" to "logical lane" and it

is not necessary to add a separate definition for "physical lane".

The comment provides no other indication as to where the current definition of "lane" lacks clarity, accuracy, or precision. The suggested remedy includes no other proposals other than to include a figure "similar to what is used in 802.16-2017". As stated in the response to comment r01-24 against P802.3/D3.1, "it is believed that the definition is clear as it is written and does not require a figure." In addition, it is unclear what relationship the Figure 3-1 from IEEE Std 802.16-2017 has to the definition of "lane". Therefore it is not clear what figure would satisfy the commenter.

IEEE P802.3 (IEEE 802.3cj) D3.2 Maintenance #12 (Revision) 2nd Sponsor recirculation ballot comments

Cl 121 SC 121.8.5.4 P136 L 20 # r02-7
 Dawe, Piers J G Mellanox Technologies

Comment Type TR Comment Status R

A much wider range of signals are allowed to be transmitted than are covered by SRS (required to be received).
 At present it is allowed to make a transmitter with a noisy or distorted signal, use heavy emphasis to get it to pass the TDECQ test, yet a compliant receiver that passes SRS would not need to receive it. The range needs to be bounded on the left hand side of the maps in daw_3cd_01a_0318 and daw_032118_3cd_adhoc so that the receiver design can be bounded in terms of having to "invert" heavily over-emphasised signals, and the gap between possible signals and SRS closed or narrowed.
 The remedy doesn't directly outlaw over-emphasised signals, but gives them worse TDECQ scores.
 D3.1 comment 35

SuggestedRemedy

This remedy lets the transmitter designer use reasonable amounts of emphasis, balancing his own transmitter bandwidth and the reference receiver front-end bandwidth.
 After saying where the largest magnitude tap coefficient is, add "The tap coefficients are constrained so that the sum of the other four tap coefficients is less than zero."
 Similarly in clauses 122, 124.

Response Response Status U

REJECT.
 This comment is a re-statement of unsatisfied negative comment r01-35 against D3.1 with a different suggested remedy.
 The need for additional restrictions on the equalizer tap coefficients in the TDECQ measurement for these approved SMF PMDs has not been established, and insufficient evidence has been provided that the proposed restriction fixes the claimed problem.
 To date no contribution has been made that demonstrates the problem described by unsatisfied negative comment r01-35 against D3.1 (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that restricting the sum of the four smallest magnitude tap coefficients to be less than zero prevents this issue from occurring.

The stressed receiver sensitivity (SRS) requirement is not intended to cover all possible transmitter waveforms and power levels. The argument used in the comment could be used to suggest that any transmitter with a waveform that does not match the SRS conformance test signal should be excluded. This would disallow a "good" transmitter with a much lower TDECQ than the maximum (and therefore with a lower minimum power).

Cl 121 SC 121.8.5.3 P136 L 14 # r02-8
 Dawe, Piers J G Mellanox Technologies

Comment Type TR Comment Status R

A much wider range of signals are allowed to be transmitted than are covered by SRS (required to be received).
 At present it is allowed to make a transmitter with a noisy or distorted signal and use emphasis to get a "noise enhancement credit" to pass the TDECQ test, yet the eye closure is more than the TDECQ limit and a compliant receiver that passes SRS would not need to receive it. The range needs to be bounded on the top side of the maps in daw_3cd_01a_0318 and daw_032118_3cd_adhoc so that the receiver design can be bounded in terms of resolution and patterning, and the gap between possible signals and SRS closed or narrowed.
 The first remedy has the disadvantage that errors in OMA measurement degrade its accuracy.
 D3.1 comment 35

SuggestedRemedy

Either:
 1. Limit TDECQ $-10 \cdot \log_{10}(C_{eq})$ to ≤ 2.8 dB.
 or:
 2. 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 (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).
 Limit 3 dB.
 Either remedy to apply to all PMDs that use TDECQ in Section 8, although it would not matter much for 400GBASE-FR8 if the over-emphasis limit (see another comment) is in force.

Response Response Status U

REJECT.
 This comment is a re-statement of unsatisfied negative comment r01-35 against D3.1 with changes to the options in the suggested remedy.

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

There is no consensus to make a change.

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, r03-27 against P802.3bs D3.3, and r01-35 against P802.3 (IEEE 802.3cj) D3.1 which were similarly rejected.

IEEE P802.3 (IEEE 802.3cj) D3.2 Maintenance #12 (Revision) 2nd Sponsor recirculation ballot comments

Cl 121 SC 121.8.5.4 P136 L 20 # r02-9
 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 the SRS test range: see right hand corner of the maps in daw_032118_3cd_adhoc. If this hole is not plugged, there could be interoperability issues, and/or some product receivers with more tap strength than is needed to receive the range of reasonable signals, degrading their cost/power/performance trade-off.

This issue is less severe than the lack of a limit on the left hand side, but should be considered nevertheless.

These remedies don't by themselves outlaw slower signals, but give them worse TDECQ scores.

D3.1 comment 36.

SuggestedRemedy

Either:

1. Set a maximum cursor strength limit, 1.59

or:

2. Set a maximum limit for $10 \cdot \log_{10}(\text{Ceq})$, 2.2 dB

Similarly in clauses 122, 124, although because the signalling rate for 124 is higher, the limit there might be higher or absent.

Response Response Status U

REJECT.

This comment is a re-statement of unsatisfied negative comment r01-36 against D3.1, which proposed to "Set a maximum cursor strength limit, which might be around 1.3".

The need for a limit to cursor strength or set a maximum limit for $10 \cdot \log_{10}(\text{Ceq})$ has not been established (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that the proposed limit of 1.59 for cursor weight or 2.2 dB for $10 \cdot \log_{10}(\text{Ceq})$ removes the demonstrated issue while not disallowing "reasonable" transmitters.

There was no consensus to make a change.

Cl 121 SC 121.8.5.4 P136 L 20 # r02-10
 Dawe, Piers J G Mellanox Technologies

Comment Type TR Comment Status R

dawe_3cd_01a_0318 showed that for the slowest, cleanest, most symmetrical allowed signal, putting the cursor at tap 3 has a negligible "benefit" vs. tap 2. This signal should probably not be allowed anyway (see another comment), and the reference receiver in TDECQ isn't meant to fully represent a real receiver. Rougher, noisier, faster, or less symmetric signals would see even less difference. Yet the option adds cost to real receivers (depending on implementation) and time to TDECQ measurements. In the last meeting, the effect of chromatic dispersion was mentioned. I have not yet found a chromatic dispersion effect that creates a slow leading edge, slower than trailing, for enough of the edges that it can be equalised. If it doesn't exist...
 D3.1 comment 37

SuggestedRemedy

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

Response Response Status U

REJECT.

This comment is a re-statement of part of unsatisfied negative comment r01-37 against D3.1.

It has not been demonstrated that disallowing tap 3 as having the largest magnitude tap coefficient is an improvement to the draft.

IEEE P802.3 (IEEE 802.3cj) D3.2 Maintenance #12 (Revision) 2nd Sponsor recirculation ballot comments

Cl 116 SC 116 P19 L1 # r02-11
 Dawe, Piers J G Mellanox Technologies

Comment Type TR Comment Status R

802.3cd has made and may make changes to material similar to clauses 116 to 124 and their annexes that should be applied here too. In particular, the 1% TDECQ threshold adjust should be common to all SMF clauses that use TDECQ, or absent from all.

SuggestedRemedy

Apply the changes as appropriate.

Response Response Status U

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3/D3.2 and IEEE P802.3/D3.1 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

Although changes have been made to material similar to that found in clauses 116 to 124, those changes are not "final" as IEEE P802.3cd is still in ballot. It is therefore not appropriate to make the same changes to this draft at this time. It is also unclear whether or not all of the changes made by IEEE P802.3cd are required to be made in this draft.

For the "the 1% TDECQ threshold adjust", making this change this would place an extra burden on 200 Gb/s and 400 Gb/s receivers in the field and this change in the P802.3cd draft is expected to lead to changes in other parameters (such as the maximum TDECQ value) in future versions of the draft P802.3cd specifications that would not be included here.

Cl 121 SC 121.8.5.4 P136 L19 # r02-12
 Dawe, Piers J G Mellanox Technologies

Comment Type TR Comment Status R

Two apparent causes of inaccuracy in TDECQ:

1. Somewhat arbitrary, pattern-dependent measurement of OMA directly affects TDECQ;
 2. The rule that the sum of the equalizer tap coefficients is equal to 1 seems to force the TDECQ algorithm to miss the optimum, at least sometimes. This appears to be not the same as the 1% threshold adjust issue.
- D3.1 comment 35.

SuggestedRemedy

Issue 1 is cancelled out in (OMA-TDECQ) but not in OMA, so the issue is controlling the signal quality (as opposed to its useful amplitude). Use of TDECQrms as in another comment partially addresses this.

For issue 2: could delete "The sum of the equalizer tap coefficients is equal to 1." The reference receiver could be described as having an offset so that the average power is mapped to zero at the FFE input. Then the thresholds are simply -OMAouter/3, 0, OMAouter/3.

Response Response Status U

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3/D3.2 and IEEE P802.3/D3.1 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

Unsatisfied negative comment r01-35 against D3.1 concerns "bad" transmitters that pass the TDECQ test but should be excluded because they "leave a realistic, compliant receiver with an unreasonable challenge". This is not related to claimed inaccuracy in the TDECQ measurement.

No evidence has been presented that supports the view that the measurement method specified for OMAouter causes inaccuracy in TDECQ.

There is no consensus to make a change.