Cl **180** *SC* **180.7.1** *P* **473** *L* **9** # R2-195

christensen, steen NVIDIA Corporation

Comment Type **TR** *Comment Status* **R** *Tx transition time (OI)*

Risetime of 8ps is overly restrictive for modern receivers. With DFE and standard FFE design a risetime of e.g. 10ps is absolutely acceptable and will allow for more types - and cheaper transmitters.

SuggestedRemedy

Response *Response Status* **U**

REJECT.

Based on the comment, it appears that the intention of the comment is to change the transition time (max) from 8ps to 10ps.

Resolved using the response to comment R2-200.

Cl **180** *SC* **180.9.15** *P* **494** *L* **4** # R2-196

christensen, steen NVIDIA Corporation

Comment Type **TR** *Comment Status* **R** *Rx sensitivity (OI)*

2GHz tone for making the transmitter used in compliance testing is too restictive for many receivers. Noise is expected to be spread across the full spectrum. It should not be expected to have such transmitters in systems.

SuggestedRemedy

Response *Response Status* **U**

REJECT.

The following related contributions were reviewed by the CRG:
https://www.ieee802.org/3/dj/public/26_09/dawe_3dj_02a_2609.pdf

The related contribution proposed three options as follows:
Option 1: Slide 14
Option 2: Slide 15
Option 3: Slide 18

There is no consensus to make the proposed changes.

Cl **180** *SC* **180.9.14** *P* **492** *L* **16** # R2-198

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Rx sensitivity (OI)*

802.3 specifies equipments by their behaviour at external interfaces. It does not require a particular implementation; all implementations should be held to the same standard so that they are interoperable and the spec does not exhibit favouritism. Additionally it provides example specifications for some internal interfaces such as AUI C2M, but these are not required.
D3.2 erroneously disallows methods 174A.9.7 (BLER method for all lanes using PMA measurements) and 174A.11 (using PCS measurements). The former is incorrect because errors are shared by a common FEC), and also, measuring an individual lane and extrapolating as far as 6.9e-28 is subject to errors that can be reduced by pooling the lane measurements as described in 174A.9.7. The latter is incorrect because it's at the PCS where the errors matter, and there may be no AUI in a product.
Compare 174A.10, Error ratio tests for 800GBASE-LR1... where "the number of lanes p = 1" although there are four signals rx_signal_xi, rx_signal_xq, rx_signal_yi, rx_signal_yq whose errors could have been counted separately. That might be interesting for diagnostics but it is not appropriate as a spec limit.

In general (for the receive side), we should consider errors at or near the PCS first, and fall back to counters nearer to the PMD when necessary. Because this is an equipment spec not a component spec.

As explained before, receiver implementers who don't like the all-lane methods can conservatively use 174A.9.5 or 174A.9.6, if they allow for extra scatter. Qualifying a module for 8 x DR1 is a different matter anyway (different BLER block time on a lane).
The all-lanes calculation is still conservative because each receiver lane is tested with minimum victim and maximum optical crosstalk aggressor, and they can't all be like that, and because each transmitter lane is specified separately so they are all at least as good as minimum that the receiver is tested against.

SuggestedRemedy

Revert the changes that eliminated 174A.9.7 and 174A.11 in the IMDD clauses.

Response *Response Status* **U**

REJECT.

The CRG agreed on the resolution of comment R1-195 against 3.1 to fix the ambiguity in the receiver sensitivity measurement method and make it appropriate for a per-lane metric. Reverting the change as proposed in the Suggested Remedy would bring us back to an ambiguous definition of receiver sensitivity.

The resolution of comment R1-195 is in the following comment report:
https://www.ieee802.org/3/dj/comments/D3p1/8023dj_D3p1_comments_final_clause.pdf

The following related contribution was reviewed by the CRG:
https://www.ieee802.org/3/dj/public/26_09/dawe_3dj_04_2609.pdf

There is no consensus to make the proposed changes.

Straw Poll TF-4 (directional)
I support making the changes proposed in the suggested remedy of comment R2-198.
Yes: 11 No: 40

Cl **180** *SC* **180.7.1** *P* **473** *L* **8** # R2-199

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *PMD spec (OI)*

D3.0 comment 442 and D3.1 comment R1-298: as line rates are increased, maintaining a particular extinction ratio becomes increasingly challenging (this will be even worse at 400G). However, MPI can get worse at low extinction ratio, when it happens (rarely - fading) although for PAM4 this is a fairly weak function. On the other hand ECQ improves with low extinction ratio (linearity), differently for different technologies but it always happens (not fading). Also, RF power and supply voltage can be reduced at lower extinction ratio, and/or modulators can be made smaller or faster. We have addressed this before, in 100GBASE-DR where we reduced the minimum from 5 dB to 3.5 dB for a 0.3 increase in OMA-TDECQ (see https://ieee802.org/3/cd/public/July17/dawe_3cd_01_0717.pdf), then in 100GBASE-FR1, 100GBASE-LR1 and later PMDs when we specified 3.5 dB across the board.
It is time to take the next step in this long-running secular trend.
Notice D3.0 comment 72, proposing reducing the OMA for DR by 0.6 dB, had some support. So it is reasonable to take the simple approach of increasing the budget by 0.2 dB by improving the sensitivity by only 0.2 dB, not changing the max OMA and average power limits.
By the way, if an additional 0.2 dB seems large when Table 190-9 says "includes an allocation of 0.1 dB for MPI penalty" some transmitter-related MPI penalty is measured in TDECQ, and some is needed to preserve the insertion loss vs. discrete reflectances trade-off in e.g. Table 180-12 and 182-12. The 0.1 dB in the draft for no additional reflectances is barely changed, but I am not proposing changing Table 180-12 or 182-12. In 2017 we estimated 0.3 dB extra for going from 5 dB to 3.5 dB extinction ratio, consistent with this proposal.
https://ieee802.org/3/dj/public/26_07/dawe_3dj_01_2607.pdf provides background.

SuggestedRemedy

In Table 180-7 (transmitter), change Extinction ratio, each lane (min) from 3.5 dB to 3 dB;
In Table 180-8 (receiver), change Receiver sensitivity (OMAouter), each lane (max) from -3.4 to -3.6, and from -4.3 + TECQ to -4.5 + TECQ
and change Stressed receiver sensitivity (OMAouter), each lane (max) from -0.9 dBm to -1.1 dBm;
Revise Figure 180-4, Receiver sensitivity (OMAouter), each lane (max), and Figure 180-5 (Tx and Rx vs. ECQ), to reflect this change;
In Table 180-9 (budget), change, change Power budget (for max TDECQ) from 6.7 dB to 6.9 dB, change Allocation for penalties (for max TDECQ) from 3.7 dB to 3.9 dB;
In table footnote b, change "The channel insertion loss of 3 dB shown in this illustrative link power budget is valid for a maximum MPI penalty of 0.1 dB, which depends on the number and value of discrete reflectances present in the link." to "The channel insertion loss of 3 dB shown in this illustrative link power budget is valid for a maximum MPI penalty of 0.3 dB. MPI penalty depends on the number and value of discrete reflectances present in the link.";
In table footnote c, change "an allocation of 0.1 dB for MPI penalty" to "an allocation of 0.1 dB to 0.3 dB for MPI penalty" and change "6.7 - channel insertion loss" to "6.9 - channel insertion loss".
Apply this to clauses 180 for DR and 182 for DR-2.

Response *Response Status* **U**

REJECT.

This comment is a resubmission of Draft 3.1 Comment R1-298 in the following comment report:
https://www.ieee802.org/3/dj/comments/D3p1/8023dj_D3p1_comments_final_clause.pdf

The response to comment R1-298 includes the following straw poll, indicating strong consensus of leaving the spec as is written in D3.1.

Straw Poll TF-10 (decision) of Draft 3.1 comment resolution at the 2026 July Plenary meeting is :
I support adopting the first option proposed in the suggested remedy of comment R1-298.
Yes: 12
No: 44

The following related contribution was reviewed by the CRG:
https://www.ieee802.org/3/dj/public/26_09/dawe_3dj_01_2609.pdf

Slides 9 and 10 propose the desired changes rather than the suggested remedy.

There is no consensus to make the proposed changes.

Cl **180** *SC* **180.7.1** *P* **473** *L* **9** # R2-200

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Tx transition time (OI)*

D3.0 comments 443, 319-321 and D3.1 comment R1-299: The transmitter transition time limit was scaled from 34 ps at 50G where a linear receiver e.g. 5-tap FFE was considered, and somehow rounded down from 8.5 to 8 ps for DR, then not adjusted when the DFE was added. However, for 800GBASE-LR4, the limit is 13 ps for a 7% faster line rate. Today at 200G, transmitter bandwidth relative to the line rate is more challenging, and receivers have DSP, much longer than 5 taps, with design learnings from copper where the signals are *much* slower, so they tolerate slow signals very well. 34/2/2 = 8.5 ps max is no longer appropriate; 34/2/2/(1-DFEmax) = 12.7 ps would be suitable scaling (as for 800GBASE-LR4). As to a relationship between (bandwidth and shape) -> group delay response -> difficult to equalise and error floor: kimber_3dj_01a_2603 swept a BT filter BW, without Tx FIR, with a fixed extinction ratio. Other filters and transmitters would have given other results. As this is a specification; we specify outcomes, not possible causes. A signal that the ECQ equaliser finds difficult to equalise, because of its group delay (phase) response for example, gets a bad rating anyway.
Transmitter transition time is not very accurate because it is based on single edges subject to ISI and relies on the definition of OMA, which is less reliable than we hoped and based on other parts of the pattern.
The relation between bandwidth and transition time depends on the filter type. A transmitter implementer with poor bandwidth or filter shape can get a pass on transition time by cranking up the emphasis, which makes the signal worse but compliant. This spec item encourages bad behaviour. The ECQ limit, main tap max, and overshoot limit, catch undesirable signals but in a more reasonable, balanced way. We don't need a much tighter limit for transition time than 800GBASE-LR4 has, as well.

SuggestedRemedy

Change 8 ps to 10 ps or 12.7 ps for all IMDD except 800GBASE-LR4.

Response *Response Status* **U**

REJECT.

This comment is a resubmission of Draft 3.1 Comment R1-299 in the following comment report:
https://www.ieee802.org/3/dj/comments/D3p1/8023dj_D3p1_comments_final_clause.pdf
and Draft 3.0 Comment I-319 in the following comment report
https://www.ieee802.org/3/dj/comments/D3p0/8023dj_D3p0_comments_final_clause.pdf
A straw poll was conducted when addressing comment I-319.
Previous evidence indicated a preference for a smaller transition time, there is no evidence for a longer transition time.

The following related contribution was reviewed by the CRG:
https://www.ieee802.org/3/dj/public/26_09/simms_3dj_01a_2609.pdf

The contribution proposes changing the transition time (max) from 8 ps to 10 ps for all PMDs defined in Clause 180, Clause 181, and Clause 182, and for 800GBASE-FR4 PMD defined in Clause 183.

There is no consensus to make the proposed changes.

Straw poll TF-5 (directional)
I support changing the transition time (max) from 8 ps to 10 ps for all PMDs defined in Clause 180, Clause 181, and Clause 182, and for 800GBASE-FR4 PMD defined in Clause 183.
Yes: 16 No: 48

Cl **180** *SC* **180.9.8.1** *P* **489** *L* **49** # R2-201

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TFT (OI)*

This says "A failure of the TFT accurately identifies a transmitter that is not generally interoperable". But the TFT method can fail *any* transmitter, because it requires 1.75e-12 in bin 8 (Table 180-17) while a receiver is required to deliver < about 3.9e-5 in bin 8 (Table 180-20), and there is no requirement to be any particular amount better than that.

SuggestedRemedy

Revise the method to use a calibrated receiver, or recognise that this is a screen for transmitter-receiver pairs that play nicely with each other (a useful thing to have but not a traditional spec item), rename "transmitter functional test" to "transmitter-receiver functional test", and remove the row from Table 180-7, transmit characteristics, because the transmitter because no transmitter has the ability to ensure a pass.

Response *Response Status* **U**

REJECT.

The following related contributions were reviewed by the CRG:
https://www.ieee802.org/3/dj/public/26_09/dawe_3dj_07_2609.pdf
https://www.ieee802.org/3/dj/public/26_09/cole_3dj_01b_2609.pdf (slide 4 is related to this comment)

The suggested remedy proposes two options to consider.

The suggested remedy and related contribution do not provide sufficient detail to implement a calibrated receiver (the first of two options proposed).

There is no consensus to adopt either of the options proposed in the suggested remedy.

Cl **180** *SC* **180.9.15** *P* **510** *L* **28** # R2-203

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Rx SRS (OI)*

The frequency of the sinusoidal interferer was changed from minimum 100 MHz to minimum 200 MHz, but it's still a single frequency (comment R1-304). The minimum remains at 2 GHz, same as it was at 10 Gb/s. Back then, for PAM2, without equalisers or FEC, that was a range from 4% to 39% of Nyquist; now it's 0.4% to 1.9%. Most unequalisable distortion, and reflections and crosstalk have spectra like the victim or weighted to higher frequencies, not lower. Also a stressed signal with significant sinusoidal interferer can have bimodal levels (which didn't really matter for PAM2 without equaliser). So, both the spectrum and the distribution are unrealistic and not suitable to subject real receivers to. The sampling scope aliases the real SI frequency to some other frequency which can be much higher than Nyquist and the software FFE will amplify it according to the apparent frequency, not the real one, causing a calibration error which can be surprisingly large. A PAM2 interferer has a better spectrum than the sinusoid but the same bimodal issue.

SuggestedRemedy

1. Change from a sine tone to a QPRBS generator (not synchronous to the victim) with a signalling rate similar to the victim and passed through the same or similar filtering to the victim.
2. Apply a maximum limit to the bounded interference (now QPRBS), or apply a minimum limit from filtering, e.g. minimum 1.9 dB from filtering alone.
3. In 121.8.10.1, we say "care should be taken to avoid harmonic relationships between the sinusoidal interferer, the sinusoidal jitter, the signaling rate, and the pattern repetition rate". Modify this as necessary for the different interferer and after taking advice from scope experts, describe how the frequency of the interferer and the sampling rate of the scope should be related.

Response *Response Status* **U**

REJECT.

This comment is a resubmission of Draft 3.1 Comment R1-304 in the following comment report:
https://www.ieee802.org/3/dj/comments/D3p1/8023dj_D3p1_comments_final_clause.pdf

The following related contributions were reviewed by the CRG:
https://www.ieee802.org/3/dj/public/26_09/dawe_3dj_02a_2609.pdf

The related contribution proposed three options as follows:
Option 1: Slide 14
Option 2: Slide 15
Option 3: Slide 18

There is no consensus to make the proposed changes.

Cl **180** *SC* **180.9.6.3** *P* **484** *L* **25** # R2-204

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TDECQ ref EQ coeff (OI)*

D3.0 comment 468: causal vs. anticausal pulse shapes are not equivalent.
This pre-post equalizer coefficient difference limit |w(1)/w(0) - b(1) - w(-1)/w(0)| seeks to keep the attack and decay of the signal pulse similar. Yet the evidence from https://ieee802.org/3/dj/public/25_05/chayeb_3dj_01_2505.pdf#page=8 and 9, plotted in https://ieee802.org/3/ds/public/Interim_May_27-28-2026/dawe_3ds_01_260527.pdf , shows that for a real receiver, the situation is very far from symmetrical; relatively fast attack (c(-1) close to zero) and slow decay (c(+1) strongly -ve) is tolerated well (naturally), but the opposite is not. A constrained optimiser can find a solution by choosing a bad sampling phase - but the reported ECQ would be bad.

SuggestedRemedy

Remove the absolute bars | |. Add a minimum limit, -0.4, remembering that there are acceptable signals off the bottom of the plot in dawe_3ds_01_260527.

Response *Response Status* **U**

REJECT.

This comment is a resubmission of Draft 3.1 Comment R1-282 in the following comment report:
https://www.ieee802.org/3/dj/comments/D3p1/8023dj_D3p1_comments_final_clause.pdf

The following related contribution was reviewed by the CRG:
https://www.ieee802.org/3/dj/public/26_09/dawe_3dj_05a_2609.pdf

The proposed change to the draft is to remove the absolute bars | | in the equation, add a minimum limit, -0.4, and leave the maximum as is.

There is no consensus to make the proposed changes.

Cl **180** *SC* **180.9.6.3** *P* **483** *L* **13** # R2-205

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TDECQ (OI)*

802.3 is not a test specification (the 802.3 companion test specification was withdrawn).

We are not supposed to be writing a model of a DSP. We are supposed to be defining TECQ and TDECQ, which are measured with a scope. Different. The scope does not do T-spaced sampling, for example.

Previously, and for ECQ in other clauses, we had " One way of doing this, using estimated PAM4 symbol error ratio as the figure of merit for the equalized signal, is described below." and "Alternative optimization methods such as minimum mean squared error (MMSE) may be used to determine equalizer tap weights to reduce test time, and are expected to report equal or higher values of TDECQ. These alternative methods should not be used for receiver sensitivity and stressed receiver sensitivity calibration." So it was clear that the normative value is the best value (which does not take long to find, in spite of assertions that it must). We said that a short cut that reads high may be used, where reading high does not weaken the spec. Which would have been the case even if not stated.

But this says: "For a given sampling phase phi0 and assumed RMS value sigmaG for Gaussian noise eta(t), determine the equalizer coefficients w and b(1), subject to the limits in Table 180–16, that result in the minimum mean squared error (MMSE) between yn and x(n). Alternative optimization methods may be used to determine equalizer tap weights if they report comparable values of TDECQ." This tries to make the inaccurate high-reading over-prescriptive method normative.

There are other flaws in this but the whole concept is misguided; they should be removed when the damage in this draft is undone.

SuggestedRemedy

What a mess. Replace the whole of 180.9.6 with the normal stable ECQ definition, following 168.7.5 (the latest sound version, I believe) with the few known relevant modifications for signalling rate and number FFE taps and the addition of the DFE.

Response *Response Status* **U**

REJECT.

This comment is a resubmission of Draft 3.1 Comment R1-287 in the following comment report:
https://www.ieee802.org/3/dj/comments/D3p1/8023dj_D3p1_comments_final_clause.pdf

The following related contribution was reviewed by the CRG:
https://www.ieee802.org/3/dj/public/26_09/dawe_3dj_06a_2609.pdf

The following straw poll was taken after contributions were hear by the CRG:
Straw Poll #O-1
I support adopting the changes to TDECQ in the direction proposed in dawe_3dj_06a_2609 (“ECQ method”, comment R2-205).

Yes: 9 , No: 22

Straw poll TF-6 was taken when this comment was being addressed by the CRG.

There is no consensus to make the proposed changes.

Straw poll TF-6 (directional)
I support adopting the changes to TDECQ in the direction proposed in dawe_3dj_06a_2609 ("ECQ method", comment R2-205).
Yes: 8 No: 54

Cl **180** *SC* **180.2** *P* **461** *L* **37** # I-419

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *error ratio allocation (CG)*

Based on the FLR specification and FEC CER of 174A.6 and the sentence at line 41, the sentence at line 37 is applicable to pluggable modules where the decisions are made in the module, errors can be counted at its PMA, and some error budget is reserved for the following AUIs in the receive path. But it is not correct for RTLR, LPO, NPO and CPO where there is no C2M AUI.

SuggestedRemedy

Change the first sentence to:
With a compliant input signal, a PMD receiver is expected to meet the block error ratio (BLER) of 1.45 x 10^-11 (see 174A.6), measured at the PMA adjacent to the PMD using the method described in 174A.9, with the appropriate BER_added. For a receiver in a module with an AUI C2M, BER_added is 6.4 x 10^-5. For a receiver where the PMA adjacent to the PMD is the final PMA in the receiver (e.g. in a host ASIC), BER_added is 3.2 x 10^-5. For a receiver where the PMA adjacent to the PMD is followed by an AUI C2C then the final PMA in the receiver, BER_added is 5 x 10^-5.
Similarly in other optical clauses.

Response *Response Status* **U**

REJECT.

The comment points out that some implementations may not align with sublayer partitioning as specified by IEEE 802.3.
How such a device determines its compliance is left to the implementer.
Note the last paragraph of 180.2 describes the behavior the PHY is expected to meet for cases when the PMD can not be evaluated standalone. Also, 174A.9 explains how BER_added is determined and 174A.13 summarizes the allocation of error ratios to each physical interface.

There is no consensus to make the proposed changes.

Cl **174A** *SC* **174A.9** *P* **745** *L* **22** # I-421

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *(Bx) (CG)*

The FEC bin histogram, BER_added and masks methodology is very complicated. It needs more accessible explanation.

SuggestedRemedy

Add a more accessible explanation, with graphs and examples to illustrate what is going on. Show how the method can be simplified when assessing error rates around PMDs, where BER_added is relatively small.

Response *Response Status* **U**

REJECT.

The suggested remedy does not provide sufficient detail to implement.

Cl **180** *SC* **180.3** *P* **461** *L* **49** # I-429

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *service interface (CA)*

"The service interface" language was imposed from the world of software onto 802.3 a very long time ago. For years, it was boilerplate that we did not need to take seriously. Now, with ILT, APSU, RTLR, LPO, CPO, the distinction becomes interesting and we should get this right so as to prepare solid foundations for future projects.
As is abundantly clear from the base standard and the definitions 1.4.469, Physical Medium Attachment (PMA) sublayer, and 1.4.470 Physical Medium Dependent (PMD) sublayer, a PMA provides functions for transmission, reception, clock recovery and skew alignment while the PMD interfaces to the transmission medium concerned (coax, twinax, MMF, SMF). The PMD is untimed, it's a media convertor. In a traditional retimed optical module, there are formally, a PMA which provides CDR, pattern generation and maybe pattern checking, above a PMD that converts untimed electrical signals between the internal electrical format of the module or an IC within it, and the optical format - and vice versa.
Even for a spec where there is a DFE in a reference receiver for a PMD, the spec has no concept of bits or errors at the DFE output; the digital, countable symbol steam is inside the PMA or at the interface above it, and the interface below the PMA (AUI has an analog specification.

SuggestedRemedy

Here, change "The PMD translates the encoded data to and from signals suitable for the specified medium" to "The PMD translates the signals from and to the adjacent PMA to and from signals suitable for the specified medium".
Replace the next paragraph, which describes a digital interface, with a definition for an interface carrying signals (rather than bits or a symbol stream) on the datapath. The auxiliary primitives PMD:IS_SIGNAL.request and PMD:IS_SIGNAL.indication can remain "digital"; they aren't associated with timing-critical pattern generators and such.
Make appropriate changes to 178A (APSU). Items such as the retimer reference model in Figure 178B-13 will be formally in a PMA, always. This will allow a consistent definition framework for copper, IMDD and coherent optical PHYs.

Response *Response Status* **U**

REJECT.
As defined explicitly in the draft (see 180.3), the output of the PMD is the detected PAM4 symbols (i.e., after the DSP); the PMA or Inner FEC above processes those detected symbols. The draft does not explicitly define how the the PMD is implemented. Based on that definition, for a linear optical interface, the PMD would encompass the linear optical front end and associated electronics on the module, the electrical channel between the host device and the module, and the DSP on the host device.
The definition of PMA in 1.4.469 notes some functions are only for some PHY types: "(depending on the PHY) collision detection, clock recovery and skew alignment".
The comment is proposing to change this definition to a different reference point. As currently specified the interface between the PMD and sublayer above provides a clean and consistent interface. Changing the PMD interface as proposed would required the PMD service interface to be fully redefined as an analog interface.

Cl **180** *SC* **180.7.1** *P* **468** *L* **36** # I-440

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TFSEH (OI)*

The transmitter functional symbol error histogram method tests transmitter-receiver pairs, not just transmitters. Depending on the receiver used, a transmitter could be falsely failed or falsely passed. The FEC bin 8 criterion is not connected to the FLR objective of this project, nor to a target such as OIF's 1e-15 effective BER. Extrapolation from bin 8 to bins 14-16 which are the ones that matter, is lacking. Although the idea of counting errors can be useful for deploying links in volume where the implementation of transmitter and receiver (technology, supplier, revision) really are paired, 802.3 is an interoperability standard.

SuggestedRemedy

Remove this row from this table. Remove row from Table 180-14. Move 180.9.9 to an annex.
Four clauses.

Response *Response Status* **U**

REJECT.

This comment is proposing to remove the "Transmitter functional symbol error histogram" parameter as a mandatory requirement for all PMDs defined in Clauses 180 through 183. Instead, make it informative specification.

A previous poll to determine support for removing this feature was taken by the CRG at the November 2025 Plenary meeting during Working Group ballot. The poll showed high support for retaining the TFSEH feature at the time.
See the response to comment I-138 in the following document:
https://www.ieee802.org/3/dj/comments/D2p2/8023dj_D2p2_comments_final_id.pdf

Based on straw poll TF-4, there is no consensus to make the proposed change.

Straw poll TF-3 (directional):
I support adopting the suggested remedy to comment I-440.
Y: 33
N: 16

Concern was expressed over the language and meaning of this straw poll so the poll was repeated with better clarifying language in TF-4.

Straw poll TF-4 (directional):
I support adopting the suggested remedy (as provided by commenter) to comment I-440.
Y: 20 (change to informative specification)
N: 41 (retain as mandatory specification)

Cl **180** *SC* **180.7.1** *P* **468** *L* **42** # I-442

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *ER and MPI (OI)*

The extinction ratio limit is there to protect the receiver from MPI. MPI can get worse at low extinction ratio, when it happens (rarely - fading) but for PAM4 this is a fairly weak function. On the other hand ECQ improves with low extinction ratio (linearity), differently for different technologies but it always happens (not fading). Also, RF power and supply voltage can be reduced at lower extinction ratio.

SuggestedRemedy

Reduce the extinction ratio limit from 3.5 dB to 3 dB.
The definition of TDECQ, which is already channel-dependent, can be modified to include MPI, similar to Eq. 95-3; if so, increase the TDECQ limit by 0.1 dB for the reference case in tables 180-9 and 180-12. The definition of TECQ and SECQ, where the optical channel is negligible, are not changed. The receiver sensitivity is lowered by 0.1 dB.
Or, TDECQ can remain without MPI: receiver sensitivity and stressed receiver sensitivity are lowered by 0.1 dB.
Apply this to all IMDD clauses. I can imagine that it would this would be particularly helpful for 800GBASE-LR4 where a low extinction ratio will enable a good chromatic dispersion penalty.

Response *Response Status* **U**

REJECT.
The comment proposes to reduce the limit of extinction ratio to 3dB. However, no data or simulation analysis has been provided indicating that 3.5dB can not be supported by current modulator technology, nor any data showing 3dB is sufficient for worst case of the link , particularly in high loss and penalty cases such as LR4. Thererfore, the limit of ER should remain the same until further evidence be presented to the task force.

The power budget is built with dedicated penalty allocation for MPI and DGD for the past two PAM4 signaling rates. Changing to include penalty in TDECQ as the proposed remedy could cause unecessary confusion to the industry.
TDECQ measurement setup (Figure 180-9) could not emulate the MPI effect, i.e. the channel that the optical signal passes follows the requriement specified in Table 180-15 with minimum insertion loss, and not necessarily the mulltiple reflection from the multiple connenctions or splices as in the case of real application.

Regarding the alternate approach, the current power budget and (receiver sensitiviy-OMA_min) already takes into account of MPI penalty. -0.1-(-3.4) = 3.3dB, which includes 3dB(max. IL) and 0.1dB MPI and 0.2dB DGD penalty.

There was interest expressed in this topic. Further contributions on this topic are encouraged.

Cl **180** *SC* **180.7.1** *P* **468** *L* **43** # I-443

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Transition time (OI)*

The transmitter transition time limit was scaled from 100G where a retiming (pluggable) receiver with FFE-based or other linear equalization was considered. Today at 200G, receivers are DSP, with design learnings from copper where the signals are *much* slower. Generating high bandwidth signals is far from free, and with today's receivers that tolerate slow signals well, 8 ps max is not necessary.

SuggestedRemedy

Change 8 ps to 12 ps, or remove the item completely and rely on well-chosen ECQ, tap weight and overshoot limits to protect us from bad signals on the far right of the ECQ map.

Response *Response Status* **U**

REJECT.

Resolved using the response to comment #I-319.

Cl **180** *SC* **180.7.2** *P* **470** *L* **7** # I-446

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *ER and MPI (OI)*

Tweaking the receiver sensitivities for MPI accounting (see another comment).

SuggestedRemedy

Change the receiver sensitivity from -3.4, -4.3 + TECQ to -3.5, -4.4 + TECQ. Change the stressed receiver sensitivity from -0.9 to -1.
Similarly in other IMDD clauses.

Response *Response Status* **U**

REJECT.

Resolved using the response to comment I-442.

Cl **180** *SC* **180.9.6.3** *P* **480** *L* **22** # I-469

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **A** *tap weight-DFE (OI)*

The DFE tap limit seems too weak. Copper clauses use a limit of 0.85, but with a different definition.

SuggestedRemedy

Set the limit to at least 0.35 on the old scale, or the equivalent of 0.55 on the new scale or as defined in copper clauses and COM. This leaves another 0.3 for losses between TP2 and TP5.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.
Resolved using the response to comment #I-221.

Cl **180** *SC* **180.9.9** *P* **487** *L* **13** # I-480

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TFSEH (OI)*

This functional symbol error histogram method tests transmitter-receiver pairs. As the receiver is not calibrated, and subject to much more possible variation than a CRU/scope, it tests transmitter-receiver pairs, not transmitters.

SuggestedRemedy

Remove spec row from Table 180-7. Remove row from Table 180-14. Move 180.9.9 to an annex.
Four clauses.

Response *Response Status* **U**

REJECT.
Resolved using the response to comment #I-440.

Cl **180** *SC* **180.9.9** *P* **487** *L* **13** # I-481

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TFSEH (OI)*

A real compliant receiver will have better jitter tolerance and other properties than a spec-worst receiver, so "passing" transmitter-receiver pairs can contain unsatisfactory transmitters because of the lack of receiver calibration.

SuggestedRemedy

Calibrate the receiver or delete the subclause.

Response *Response Status* **U**

REJECT.

Resolved using the response to comment #I-25.

Cl **180** *SC* **180.9.9** *P* **488** *L* **18** # I-482

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TFSEH (OI)*

A receiver is not required to meet 1.75e-12 in bin 8 under any circumstances. A receiver that exceeds this will make any transmitter-receiver pair "fail" this test.

SuggestedRemedy

Provide a thought-through mask or delete the subclause.

Response *Response Status* **U**

REJECT.

The suggested remedy does not provide sufficient detail on an alternate mask to implement.

There is no consensus to change the mask or delete the subclause.

Cl **180** *SC* **180.9.9** *P* **488** *L* **8** # I-484

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TFSEH (OI)*

A transmitter-receiver pair that is marginal to this mask (i.e. close to the limit at bin 8) would probably have a FLR in the 1e-20 to 1e-25 region, which is far removed from the project objective (FLR 6.2e-11 = effective BER 1e-13) or the OIF's target (effective BER 1e-15) or reasonable market expectations. The difference between bin 8 and bins 14-16 is very large. Just drawing a line sideways from bin 9 is not adequate.

SuggestedRemedy

Decide what the performance objective of this test should be. The obvious candidate is the OIF's target (effective BER 1e-15). Use FEC bin extrapolation as we do for receiver testing. Or delete the subclause.

Response *Response Status* **U**

REJECT.

The suggested remedy does not provide sufficient detail to implement.

There is no consensus to make the proposed changes.

Cl **180** *SC* **180.9.9.1** *P* **488** *L* **40** # I-487

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TFSEH (OI)*

This is 802.3. We specify observable signals at compliance points. In the context of an optical spec, "DUT will be interpreted as device = optical module, while what is intended to be tested is the transmitter under test as it says on line 44, and what is actually tested is a transmitter-receiver pair.

SuggestedRemedy

Change "DUT" to "PUT", and when the receiver calibration is sorted out, change to "TUT".

Response *Response Status* **U**

REJECT.

There is no consensus to make the proposed change.

Cl **180** *SC* **180.9.15** *P* **491** *L* **3** # I-493

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Rx sensitivity (OI)*

Receiver sensitivity is a property of the lanes together ("interface BER" concept), not each lane separately. We ensure that each transmit lane is good enough, and each channel lane. We test a receiver by applying the same optical power on each lane, but the receiver can share its error allocation across its lanes how it chooses, just as it can allocate errors to LSB and MSB within a lane as it chooses. This is what happens naturally if the errors are checked at the PCS. 174A.9.5 says "If this test passes for each lane, then the PHY or xMII Extender will meet the expected codeword error ratio. If this test fails, then the performance may be further verified using the method in 174A.9.6 or 174A.9.7."

SuggestedRemedy

Delete "of each lane" here and in 180.9.16. In Table 180-8, the table entries are optical powers which are by lane, not error ratios which are not, so maybe "each lane" can stay there, although it is misleading.

Response *Response Status* **U**

REJECT.

CRG discussion on this comment highlighted some ambiguity regarding the definition of receiver sensitivity and the associated measurement method, specifically with respect to per-lane and all-lane measurements.

Further consensus building and contribution is encouraged to address this matter.

There was no consensus on adopting changes.

Cl **180** *SC* **180.1** *P* **462** *L* **6** # R1-271

Dawe, Piers J G NVIDIA

Comment Type **ER** *Comment Status* **R** *PHY tables (B) (CG)*

D3.0 comment 417 which was misunderstood: Table 180-1 and 2, Physical Layer clauses associated with... are clones of each other, and they don't fit on a single page. Each really has three columns. The status (optional, required or conditional) is the same in each. But the presentation obscures the very high level of commonality. Also it is inaccurate: there is no "Clause 117--400 Gb/s RS" for example; Clause 117 is for both 200 Gb/s and 400 Gb/s operation. Really, we are indicating what is relevant in the clause, not which clause we mean, so it should be in a separate column from the clause number.

SuggestedRemedy

Present these two tables as a single table with four columns: Clause number, 200GBASE-DR1 (with the abbreviated titles from Table 180-1), 400GBASE-DR2 (with the abbreviated titles from Table 180-2), Status.
Combine footnotes c and e.
Follow Table 167-2, not Table 122-1 which inefficiently interleaves rows from two tables.
Make similar improvements in other clauses

Response *Response Status* **U**

REJECT.
This comment is re-submission of comments I-417 and I-418 against D3.0. The responses to the comments may be found in the following report:
https://www.ieee802.org/3/dj/comments/D3p0/8023dj_D3p0_comments_final_id.pdf
This new comment points out a different format than considered in the response to those comments.
Nevertheless, the separate tables as currently in the draft provide clear definition for each of the PMD types defined in the clause without creating a rather dense table combining all four PHY types into a single table. Note also that for many sublayer types multiple rows will be required, e.g., for PCS 119 for 200G/400G, 172 for 800G, 175 for 1.6T. This is also helpful on concise wording in the footnotes. Clause number and "abbreviated" titles are always combined in the same column for previous PMD clauses with rare exceptions; separating the two into separate columns would not be inappropriate but at the same time is not necessary.

Cl **180** *SC* **180.2** *P* **464** *L* **48** # R1-274

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *BLER (CB)*

D3.0 comment 419: 174A, particularly Table 174A-1, lays out the standard error allocations for the sublayers in a PHY. This paragraph is applicable to pluggable modules where the decisions are made in the module, errors can be counted at its PMA, and some error budget is reserved for the following AUIs in the receive path. But it is not correct for RTLR, LPO, NPO and CPO where there is no C2M AUI. For some of these, a different BERadded should be used, and for others, the PCS method at line 53 is more relevant.

SuggestedRemedy

Either:
Change the first sentence to:
With a compliant input signal, a PMD receiver is expected to meet the block error ratio (BLER) of 1.45 x 10^-11 (see 174A.6), measured at the PMA adjacent to the PMD using the method described in 174A.9, with the appropriate BERadded. For a receiver in a module with an AUI C2M, BERadded is 6.4 x 10^-5. For a receiver where the PMA adjacent to the PMD is the final PMA in the receiver (e.g. in a host ASIC), BERadded is 3.2 x 10^-5. For a receiver where the PMA adjacent to the PMD is followed by an AUI C2C then the final PMA in the receiver, BERadded is 5 x 10^-5.
Or,
Add this information in 180.9.14 (see another comment) and refer to it here.
Similarly in other optical clauses.

Response *Response Status* **U**

REJECT.
This comment is a resubmission of comment I-419 against Draft 3.0 recorded in the following report.
https://www.ieee802.org/3/dj/comments/D3p0/8023dj_D3p0_comments_final_id.pdf
The approved response to I-419 is repeated here:
<start>
REJECT.
The comment points out that some implementations may not align with sublayer partitioning as specified by IEEE 802.3.
How such a device determines its compliance is left to the implementer.
Note the last paragraph of 180.2 describes the behavior the PHY is expected to meet for cases when the PMD can not be evaluated standalone. Also, 174A.9 explains how BER_added is determined and 174A.13 summarizes the allocation of error ratios to each physical interface.
There is no consensus to make the proposed changes.
<end>

The error allocation provided in 180.2 accounts for the possibility of there being one or two electrical interfaces (e.g., one C2M AUI and one C2C AUI) within a PHY.

Based on the result of straw poll TF-9 there is no consensus to make the proposed changes.

Straw poll TF-9 (directional)
I support adopting the changes proposed in the suggested remedy of comment R1-274.
Yes: 22
No: 34

Cl **180** *SC* **180.9.14** *P* **492** *L* **5** # R1-277

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **A** *Rx Sensitivity (O)*

The correct BERadded is a normative fact; it should not be in an informative NOTE.
Table 180-19 provides BERadded for only one scenario, while there are several (D3.0 comment 419).
We should do a thorough job in the table, and either refer to it in 180.2 which is explaining why the spec is like it is, or repeat the info there.

SuggestedRemedy

At line 2, change "measured at the PMA using the test method in either 174A.9.5, 174A.9.6, or 174A.9.7" to "measured at the PMA using the test method in either 174A.9.5, 174A.9.6, or 174A.9.7, or at the PCS using the method of 174A.11". Delete the first informative note.
Change "The error mask Hmax(k) to be used in the method of 174A.9.5 is provided in Table 180-20." to "The error mask Hmax(k) to be used in the method of 174A.9.5 for a detachable PMD/PMA is provided in Table 180-20."
In Table 180-19, provide each BERadded as a separate row or sub-row:
Errors counted at a detachable PMA that could be followed by AUI C2M and AUI C2C 6.4e-5
Errors counted at a detachable PMD/PMA that is followed by an AUI C2C 4 x 10^-5
Errors counted at a non-detachable PMA in the host 3.2 x 10^-5
Errors counted at the PCS in the host 3.2 x 10^-5
In 180.2, after "BERadded equal to 6.4 x 10^-5.", add "See Table 180-19 for other scenarios."
Similarly in other optical clauses.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

Several of the suggested remedy points were solved in the response to comment R1-195 however after CRG discussion there was no consensus to adopt the other proposed changes in the suggested remedy that were not addressed by the resolution of R1-195.

Cl **180** *SC* **180.9.6.3** *P* **483** *L* **29** # R1-280

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TDECQ tap limit (O)*

D3.0 comment 469: the feedback tap coefficient limit, 0.33, is too weak, and even worse than in D3.0 (the definition has changed). The low limit encourages bad transmitter setup. Copper clauses use a limit of 0.85.

SuggestedRemedy

At least restore the status quo (0.43) which would leave a good allocation for losses between TP2 and TP5. There is very little traditional frequency-dependent loss between TP2 and TP3, and a moderate amount between TP3 and TP4, so that's a good allocation for losses between TP4 and TP5. Preferably, change the limit to 0.5.

Response *Response Status* **U**

REJECT.

The comment didn't provide sufficient justification for the suggested remedy.

The suggested remedy is suggesting to change b(1) limit to 0.43 to account for allocation between losses (assuming it is referred to the passive channel) between TP2 and TP5.

After CRG discussion there was no consensus to the make the proposed change.

Cl **180** *SC* **180.9.6.3** *P* **482** *L* **30** # R1-284

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TDECQ sampling (O)*

This makes out that the ECQ measurement apparatus takes T-spaced samples at phi, then phi is swept. Neither a sampling scope nor a real-time scope do that; they collect samples all across the waveform / pattern; certain samples fall in the histogram windows, others don't.

SuggestedRemedy

Delete Eq. 180-2. Delete the "sampler" on the left of Figure 180-10, or at least make it clear that it samples the whole waveform at uniform density.

Response *Response Status* **U**

REJECT.

The method is consistent itself. Implementations using non-integer number of sample/UI is upon vendor's choice.

During CRG discussion it was noted that the sampler shown is part of the reference equalizer and not intended to show the operation of the scope.

There was no consensus to make the proposed change.

Cl **180** *SC* **180.9.6.4** *P* **484** *L* **1** # R1-287

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TDECQ method (O)*

802.3 is not a test specification (the 802.3 companion test specification was withdrawn).

Previously, and for ECQ in other clauses, we had " One way of doing this, using estimated PAM4 symbol error ratio as the figure of merit for the equalized signal, is described below." and "Alternative optimization methods such as minimum mean squared error (MMSE) may be used to determine equalizer tap weights to reduce test time, and are expected to report equal or higher values of TDECQ. These alternative methods should not be used for receiver sensitivity and stressed receiver sensitivity calibration." So it was clear that the normative value is the best value (which does not take long to find, in spite of assertions that it must). We said that a short cut that reads high may be used, where reading high does not weaken the spec. Which would have been the case even if not stated.

But this says: "For a given sampling phase phi0 and assumed RMS value sigmaG for Gaussian noise eta(t), determine the equalizer coefficients w and b(1), subject to the limits in Table 180ů16, that result in the minimum mean squared error (MMSE) between yn and x(n). Alternative optimization methods may be used to determine equalizer tap weights if they report equal or lower mean values of TDECQ." This tries to make the inaccurate high-reading over-prescriptive method normative, but allows better methods, but doesn't say when.

There are other flaws in this but the whole concept is misguided they should be removed when the damage in this draft is undone.

SuggestedRemedy

What a mess. Replace the whole of 180.9.6 with the normal stable ECQ definition, following 167.8.6 (the latest sound version, I believe) with the few known relevant modifications for SMF vs. MMF and the addition of the DFE.

Response *Response Status* **U**

REJECT.

Straw poll TF-4 (directional):

I support the replacement of 180.9.6 (TDECQ definition with ECQ definition) as proposed in comment R1-287.
Yes: 16
No: 28
Abstain: 21

The suggested remedy does not provide sufficient detail to implement. Based on straw poll TF-4 there was not consensus to make the proposed change.

Cl **180** *SC* **180.9.14** *P* **491** *L* **45** # R1-294

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Rx Sensitivity (O)*

D3.0 comment 493: as is explained at length in 174A.9 (referenced at line 53), receiver sensitivity is a property of the lanes together ("interface BER" concept, transferred to BLER), not each lane separately. We ensure that each transmit lane is good enough individually, and each channel lane individually. We test a receiver by applying the same optical power on each lane, but the receiver can share its error allocation across its lanes how it chooses, just as it can allocate errors to LSB and MSB, or I and Q, within a lane, as it chooses. This is what happens naturally if the errors are checked at the PCS, and it's fine: we put symbol distribution in the FEC to prepare for it. 174A.9.5 says "If this test passes for each lane, then the PHY or xMII Extender will meet the expected codeword error ratio. If this test fails, then the performance may be further verified using the method in 174A.9.6 or 174A.9.7." Those who wish to qualify a multilane receiver (e.g., 800GBASE-DR4) as multiple single-lane PMDs (e.g., 4 x 200GBASE-DR1) will continue to do so; those who don't like 174A.9.6 (evaluate the performance of all physical lanes in an AUI component or PMD) because they don't want to convolve the per-lane BLER histograms together, can continue to use 174A.9.5 or 174A.9.7 (evaluate the performance of each physical lane in an AUI component or PMD).

SuggestedRemedy

Here, remove "of each lane" from line 45 (because the sensitivity of a single lane in a multi-lane PMD is not a spec item) and add "on each lane" on page 492 line 1 after "OMAouter" because this really is the OMA to be applied to each victim lane. Similarly in 180.9.15 for SRS. Similarly in other IMDD clauses.

I hope that the electrical clauses and AIs don't have this defect to be corrected.

Response *Response Status* **U**

REJECT.

Based on the resolution of R1-195 there was no consensus to make the proposed change.

Cl **180** *SC* **180.7.1** *P* **472** *L* **8** # R1-298

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *OMA (O)*

D3.0 comment 442: as line rates are increased, maintaining a particular extinction ratio becomes increasingly challenging (this will be even worse at 400G). However, MPI can get worse at low extinction ratio, when it happens (rarely - fading) although for PAM4 this is a fairly weak function. On the other hand ECQ improves with low extinction ratio (linearity), differently for different technologies but it always happens (not fading). Also, RF power and supply voltage can be reduced at lower extinction ratio, and/or modulators can be made smaller or faster. We have addressed this before, in 100GBASE-DR where we reduced the minimum from 5 dB to 3.5 dB for a 0.3 increase in OMA-TDECQ (see https://ieee802.org/3/cd/public/July17/dawe_3cd_01_0717.pdf), then in 100GBASE-FR1, 100GBASE-LR1 and later PMDs when we specified 3.5 dB across the board.
It is time to take the next step in this long-running secular trend.
Notice D3.0 comment 72, proposing reducing the OMA for DR by 0.6 dB, had some support. So it might be reasonable to take the simpler approach of increasing the budget by 0.2 dB by improving the sensitivity by 0.2 dB.
By the way, if an additional 0.2 dB seems large when Table 190-9 says "includes an allocation of 0.1 dB for MPI penalty" some transmitter-related MPI penalty is measured in TDECQ, and in 2017 we estimated 0.3 dB extra for going from 5 dB to 3.5 dB extinction ratio, consistent with this proposal.
We could use a budgeting formula that relates predicted MPI penalty to extinction ratio (not changing TDECQ) if the numbers were thought large enough to make it worthwhile.

SuggestedRemedy

Reduce the extinction ratio limit from 3.5 dB to 3 dB.
Similarly to what we did for 100GBASE-DR: change OMA min from
-0.1, -1+max(TECQ, TDECQ) for all extinction ratios, to
-0.1, -1+max(TECQ, TDECQ) for extinction ratio >= 3.5 dB
+0.1, -0.8+max(TECQ, TDECQ) for extinction ratio <3.5 dB.
In Table 180-9, Illustrative link power budget, change "Power budget (for max TDECQ)" to Power budget (for max TDECQ and 3.5 dB extinction ratio).
Apply this to DR (180) and DR-2 (182) (the other IMDD clauses have higher reflections).

Or, more simply:
Reduce the extinction ratio limit from 3.5 dB to 3 dB.
Reduce the receiver sensitivity and stressed receiver sensitivity by 0.2 dB
Increase the budget and the allocation for MPI penalty by 0.2 dB.

Apply this to clauses 180 and 182.

Response *Response Status* **U**

REJECT.

This comment is a resubmission of Draft 3.0 Comment I-442 in the following comment report

https://www.ieee802.org/3/dj/comments/D3p0/8023dj_D3p0_comments_final_clause.pdf

No new substantive information on the desired outcome is provided in the suggested remedy

Based on the result of straw poll TF-10 there is no consensus to adopt the proposed changes.

Straw Poll TF-10 (decision)
I support adopting the first option proposed in the suggested remedy of comment R1-298.
Yes: 12
No: 44

Cl **180** *SC* **180.7.1** *P* **472** *L* **9** # R1-299

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Tx transition time (O)*

D3.0 comments 443, 319-321: The transmitter transition time limit was scaled from 100G where a product receiver with FFE-based or other linear equalization was considered. Today at 200G, transmitter bandwidth relative to the line rate is more challenging, and receivers have DSP, with design learnings from copper where the signals are *much* slower, so they tolerate slow signals well. 34/2/2 ~8 ps max is no longer necessary. There may be a relationship between (bandwidth and shape) -> group delay response -> difficult to equalise and error floor, which could be investigated further. However, this is a specification; we specify outcomes, not possible causes. A signal that the ECQ equaliser finds difficult to equalise, gets a bad rating anyway. Transmitter transition time is not very accurate because it relies on the definition of OMA, which is less reliable than we hoped. The relation between bandwidth and transition time is not 1:1. A transmitter implementer with poor (bandwidth and shape) can get a pass on transition time by cranking up the emphasis, which makes the signal worse but compliant. This spec item encourages bad behaviour. The ECQ main tap max, and the overshoot limit, also catch undesirable signals but in a more reasonable, balanced way. We don't need this as well.
Notice that for 800GBASE-LR4, the limit is 13 ps for a 7% faster line rate, and it's probably the same receiver DSP that tolerates this.

SuggestedRemedy

Change 8 ps to 10 ps for all IMDD except 800GBASE-LR4, or remove the item completely and rely on well-chosen ECQ, tap weight and overshoot limits to protect us from bad signals on the far right of the ECQ map.

Response *Response Status* **U**

REJECT.

This is a resubmission of Draft 3.0 Comment I-443 in the following comment report:
https://www.ieee802.org/3/dj/comments/D3p0/8023dj_D3p0_comments_final_clause.pdf

Also, I-443 was resolved in conjunction with comments I-319, I-320, I-321, and I-322. The straw poll taken by the CRG related to this set of comments indicated a clear preference for not increasing the transition time.

This new comments provides no new substantive information to support the proposed changes.

Cl **180** *SC* **180.9.15** *P* **493** *L* **45** # R1-304

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *SRS interference (O)*

The frequency of the sinusoidal interferer was changed from minimum 100 MHz to minimum 200 MHz, but it's still a single frequency. The sampling scope aliases it to some other frequency and the software FFE will amplify it according to the apparent frequency, not the real one, causing error.

SuggestedRemedy

1. Change from a sine tone to a PRBS generator (not synchronous to the victim) with a high enough signalling rate that wherever the aliasing takes the apparent frequencies to (there are many real frequencies in a PRBS), the overall enhancement will be reasonably consistent. A PAM2 PRBS is more "bounded-like" than a PAM4 one as well as having a uniform spectrum.
2. Limit the amount of bounded interference.
3. In 121.8.10.1, we say "care should be taken to avoid harmonic relationships between the sinusoidal interferer, the sinusoidal jitter, the signaling rate, and the pattern repetition rate". After taking advice from scope experts: add the sampling rate of the scope to that list.

Response *Response Status* **U**

REJECT.

The following contribution was reviewed by the CRG:
https://www.ieee802.org/3/dj/public/26_07/dawe_3dj_02_2607.pdf

The contribution provided good material to understand the limitation of one aspect of the current SRS specification and provided a better understanding of the proposal.

However, the proposal is not sufficiently well developed to adopt in its current form.

There is no consensus to adopt the proposed changes.

Cl **180** *SC* **180.9.8.1** *P* **488** *L* **40** # R1-305

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **A** *TFT wording (O) (B)*

This says that Test_margin equals 1.5 dB, which decreases the required ORx operating BER to
2.4e5. But, 1.5 dB and 2.4e5 are not related, and there is no *required* ORx operating BER, although an ORx operating BER for calculating the bin mask is chosen in 180.9.8.2. It doesn't "equal" 1.5 dB as the output of some formula or measurement; it was chosen.

SuggestedRemedy

If the TFT section is kept, change to just "It is set to 1.5 dB."

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.
Resolved using the response to comment R1-118.

Cl **180** *SC* **180.9.8** *P* **487** *L* **32** # R1-307

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TFT (O)*

The TFT section has a blow-by-blow recipe but no introductory statement of what it is trying to do.

SuggestedRemedy

If the TFT section is kept,
Add a sentence explaining what it is for;
Address the relation between the numbers in this section and the other clauses and annexes that have very different numbers;
Adjust the numbers in this section so they relate reasonably to the rest of the draft.

Response *Response Status* **U**

REJECT.

The suggested remedy does not provide sufficient detail to implement.