Cl **180** *SC* **180.9.6.1** *P* **480** *L* **53** # R1-154

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

Comment Type **TR** *Comment Status* **R** *TDECQ/TFT AUI (O)*

802.3dj has spend last 1.5 year how to best capture jitter in TDECQ and FRx, unless xAUI-n (C2M) interface operate with jitter tolerance condition in 176D.8.14 TDECQ measurement can be too optimistic.

SuggestedRemedy

Add sentence: For those cases where there is an xAUI-n chip-to-chip (C2C) or chip-to-module (C2M) interface (see Table 180û1 through Table 180û4), the clock source for the tes pattern is derived from the clock recovered from the xAUI-n input signal. The xAUI-n input signal is that described for the Jitter tolerance test in 176D.8.14. The pattern of the xAUI-n input signal may be PRBS31Q, scrambled idle, or a valid xBASE-R signal.
Suggest above text and a fiigure to be provided to move into a section called "Transmitter Under Test Requirements"
See ghiasi_3dj_02_2605

Response *Response Status* **U**

REJECT.

This comment is resubmission of comments I-258 in D3.0 in the following report
https://www.ieee802.org/3/dj/comments/D3p0/8023dj_D3p0_comments_final_clause.pdf

A straw poll related to comment I-258 is shown here:

Straw poll TF-17 (directional) from May 2026
I support adoption of the proposed change (blue text) in the top paragraph on slide 9 of ghiasi_3dj_02_2605.
Yes: 10
No: 23

There was clear consensus from the May 2026 straw poll to keep the draft as is, not adding the jitter tolerance condition to the AUI interface when measuring the TDECQ of the PMD transmitter.

The following presentation was reviewed by the CRG
https://www.ieee802.org/3/dj/public/26_07/ghiasi_3dj_02_2607.pdf

No additional new information supporting the suggested remedy or the proposed implementation.

After CRG discussion there was no consensus to make any changes.

Straw poll TF-7 (directional)
I support the proposed changes on slide 9 of ghiasi_3dj_02_2607
Yes: 10
No: 17

Cl **181** *SC* **181.9.6** *P* **519** *L* **53** # R1-155

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

Comment Type **TR** *Comment Status* **R** *TDECQ/TFT AUI (O)*

802.3dj has spend last 1.5 year how to best capture jitter in TDECQ and FRx, unless xAUI-n (C2M) interface operate with jitter tolerance condition in 176D.8.14 TDECQ measurement can be too optimistic.

SuggestedRemedy

Add sentence: For those cases where there is an xAUI-n chip-to-chip (C2C) or chip-to-module (C2M) interface (see Table 180û1 through Table 180û4), the clock source for the test pattern is derived from the clock recovered from the xAUI-n input signal. The xAUI-n input signal is that described for the Jitter tolerance test in 176D.8.14. The pattern of the xAUI-n input signal may be PRBS31Q, scrambled idle, or a valid xBASE-R signal.
Suggest above text and a fiigure to be provided to move into a section called "Transmitter Under Test Requirements" Either locally create the section or just reference the section created in 180.9.6.x
See ghiasi_3dj_02_2605

Response *Response Status* **U**

REJECT.

Resolved using the response to comment R1-154.

Cl **182** *SC* **182.9.6** *P* **551** *L* **36** # R1-156

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

Comment Type **TR** *Comment Status* **R** *TDECQ/TFT AUI (O)*

802.3dj has spend last 1.5 year to best capture jitter in TDECQ and FRx, unless xAUI-n (C2M) interface operate with jitter tolerance condition in 176D.8.14 TDECQ measurement can be too optimistic
see unsatisfied D3.0 comment 260

SuggestedRemedy

Add sentence: For those cases where there is an xAUI-n chip-to-chip (C2C) or chip-to-module (C2M) interface (see Table 180û1 through Table 180û4), the clock source for the test pattern is derived from the clock recovered from the xAUI-n input signal. The xAUI-n input signal is that described for the Jitter tolerance test in 176D.8.14. The pattern of the xAUI-n input signal may be PRBS31Q, scrambled idle, or a valid xBASE-R signal.
Suggest above text and a fiigure to be provided to move into a section called "Transmitter Under Test Requirements" Either locally create the section or just reference the section created in 180.9.6.x
See ghiasi_3dj_02_2605

Response *Response Status* **U**

REJECT.

Resolved using the response to comment R1-154.

Cl **183** *SC* **183.9.6** *P* **583** *L* **49** # R1-157

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

Comment Type **TR** *Comment Status* **R** *TDECQ/TFT AUI (O)*

802.3dj has spend last 1.5 year to best capture jitter in TDECQ and FRx, unless xAUI-n (C2M) interface operate with jitter tolerance condition in 176D.8.14 TDECQ measurement can be too optimistic
see unsatisfied D3.0 comment 261

SuggestedRemedy

Add sentence: For those cases where there is an xAUI-n chip-to-chip (C2C) or chip-to-module (C2M) interface (see Table 180û1 through Table 180û4), the clock source for the tes pattern is derived from the clock recovered from the xAUI-n input signal. The xAUI-n input signal is that described for the Jitter tolerance test in 176D.8.14. The pattern of the xAUI-n input signal may be PRBS31Q, scrambled idle, or a valid xBASE-R signal.
Suggest above text and a fiigure to be provided to move into a section called "Transmitter Under Test Requirements" Either locally create the section or just reference the section created in 180.9.6.x
See ghiasi_3dj_02_2605

Response *Response Status* **U**

REJECT.

Resolved using the response to comment R1-154

Cl **180** *SC* **180.9.14** *P* **492** *L* **30** # R1-208

Dudek, Michael — Marvell

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

The receiver sensitivity test is intended to ensure inter-operability with worst case transmitters that have the given TDECQ/TECQ. It is expected that a low extinction ratio will be more worst case for a receiver than a high extinction ratio, therefore the min extinction ratio should be used for the LRS test.

SuggestedRemedy

Change the min extinction ratio to 3.5dB and the max to 4.5dB in table 180-18.

Response *Response Status* **U**

REJECT.

The CRG reviewed slides 21-23 of
https://www.ieee802.org/3/dj/public/26_07/cole_3dj_01a_2607.pdf

The ER specification of LRS compliance signal is correctly set.
The idea of LRS is to create a very good Tx signal as low stress to find out the receiver's best performance (OMA-TDECQ) so shifting ER to lower side is not low stress.

Cl **180** *SC* **180.9.1** *P* **478** *L* **45** # R1-262

Swenson, Norman — Nokia Corporation,Norman Swenson Consulting,Point

Comment Type **ER** *Comment Status* **R** *'TECQ nomenclature (Bp) (O)*

TDECQ and TECQ have changed significantly from previous versions with the conversion from an FFE to a DFE for the reference equalizer. To avoid confusion with TDECQ and TECQ for 100G, the name should be changed to something else.

SuggestedRemedy

Change "Transmitter and Dispersion Eye Closure for PAM4 (TDECQ)" to "Transmitter and Dispersion Eye Closure for PAM4 - DFE (TDECQ-D)". Make similar changes for TECQ. Make the changes where appropriate throughout the document.

Response *Response Status* **U**

REJECT.

This change is unnecessary, in effect could result in confusion. Using a different name of the same parameter leads to confusion to reader.
DFE is part of the reference equalizer, and the reference equalizer has been evolving since TDECQ was adopted as a transmitter parameter. We used different ref. equalizers for different speeds and MMF vs SMF. TDECQ's purpose and meaning is not changed by the reference equalizer, just like implementation DSPs using different equalizer schemes are all DSPs, and all form 800GBASE-DR4 module.

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* **180.1** *P* **463** *L* **6** # R1-272

Dawe, Piers J G NVIDIA

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

D3.0 comment 418 which was misunderstood: Table 180-3 and 4, Physical Layer clauses associated with... are not quite 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 and makes it hard to spot the differences (PCS and PMA). Also it is inaccurate: there is no "120F--800GAUI-8 C2C" for example; Annex 120F is for many MAC rates. 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, 800GBASE-DR4 (with the abbreviated ~titles from Table 180-3), 1.6TBASE-DR8 (with the abbreviated titles from Table 180-4), Status. There will be one more row than in Table 180-3.
Combine footnote d.
Follow Table 167-2, not Table 122-1 which inefficiently interleaves rows from two tables.

Response *Response Status* **U**

REJECT.
Resolved using the response to comment R1-271.

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* **180.2** *P* **465** *L* **48** # R1-275

Dawe, Piers J G NVIDIA

Comment Type **ER** *Comment Status* **R** *editorial (Bp) (CB)*

D3.0 comment 420 (once in the bucket): the logic of this subclause is obscured. The idea is that we have an FLR expectation, which is translated to a PHY receiver BLER expectation, which is further translated to a PMD receiver BLER expectation using the BERadded method. But the second and third paragraphs in this subclause (and similar) are out of order.

SuggestedRemedy

Swap the second and third paragraphs so that the logic behind these requirements is more apparent, here and in similar subclauses.

Response *Response Status* **U**

REJECT.
This comment is a resubmission of Draft 3.0 comment I-420 in the following report:
https://www.ieee802.org/3/dj/comments/D3p0/8023dj_D3p0_comments_final_clause.pdf
There was no consensus to make the proposed change.
No new substantive evidence or rationale is provided in the comment.
No new substantive information on the desired outcome is provided in the suggested remedy.

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* **180.9.5** *P* **479** *L* **53** # R1-278

Dawe, Piers J G NVIDIA

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

The improved OMA definition leaves some loose ends:
If the waveform is an inverted SSPRQ, does one use 7+ threes and 6+ zeros or 6+ threes and 7+ zeros?
Why 6 and 7?
The statistics of 7 vs. 7 or more could be better or worse; it depends what's in SSPRQ.
Other things being equal, excluding "or more" makes the analysis simpler (although no-one would do it by hand anyway). Similarly, does "central 2 UI" give better or worse consistency than 2.5 to 4.5 / 2 to 4? Starting at 2 is probably undesirable.
We changed the definition because a measurement on just two runs is open to error caused by ISI, but measuring with PRBS13Q still has that problem. PRBS13Q is the reason why we say 7 and 6 not 7 and 7.

SuggestedRemedy

Change the text so that:
The definition is with SSPRQ;
Runs of 7, whether zeros or threes, except when using PRBS13Q;
Users are warned that PRBS13Q (7 and 6 because that's what's available) is less accurate;
in case of a conflict, SSPRQ takes precedence;
Choose "runs of" vs. "or more" and central vs. 2.5 to 4.5 according to the evidence.

Response *Response Status* **U**

REJECT.

The comment does not provide sufficient evidence that the proposed changes are warranted.

The draft as it stands is correct and allows both SSPRQ and PRBS 13Q as test pattern, as what has been done in the past for IMDD optical PMDs. Both test patterns are widely adopted in the industry.

Using all patterns that are 6 zeros or more and 7 threes or more is consistent with the industry implementations, therefore it provides better consistency.

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* **483** *L* **19** # R1-281

Dawe, Piers J G NVIDIA

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

D3.0 comment 468: The tap weights were created by drawing limit lines that passed a population of transmitters, some time ago. The main tap maximum and the other FFE tap limits need review, taking the DFE into account and considering what is cost-effective for receivers. In particular, the we will see less of the characteristic alternating FFE tap weights because the DFE can do that job better, so taps 3+ should be tightened up.

SuggestedRemedy

Review the tap weight limits and make changes as appropriate.

Response *Response Status* **U**

REJECT.

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

The suggested remedy does not provide sufficient detail to implement.

Cl **180** *SC* **180.9.6.3** *P* **483** *L* **25** # R1-282

Dawe, Piers J G NVIDIA

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

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.

SuggestedRemedy

Remove the absolute bars | |. If necessary, add a minimum limit, e.g. -0.55, remembering that there are acceptable signals off the bottom of the plot in dawe_3ds_01_260527.

Response *Response Status* **U**

REJECT.

This is a resubmission of Draft 3.0 Comment I-468 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

Note that the contribution cited in the comment (dawe_3ds_01_260527) was not presented to the IEEE 802.3dj task force or CRG, but rather was presented to another task force (802.3ds).

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

Cl **180** *SC* **180.9.8.1** *P* **488** *L* **50** # R1-283

Dawe, Piers J G NVIDIA

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

This says "the best estimate of the actual MPI and DGD penalties of the test fiber", yet the fibre alone doesn't make much MPI. The reflections between transmitter and connectors, and receiver and connectors, are more important. It is not clear which reflections are taken as baseline and which are attributed to the fibre, particularly as both transmitter and receiver reflectances are not necessarily at the spec limits. MPI varies with polarisation yet there is no polarisation control in Figure 180-12.

SuggestedRemedy

If the TFT section is kept;
if MPI penalty variations are important, explain how the MPI penalty is to be estimated and add polarisation control. If MPI penalty variations are not so important or not feasible to estimate with confidence, use a fixed allocation

Response *Response Status* **U**

REJECT.

The CRG reviewed slides 6 and 7 of
https://www.ieee802.org/3/dj/public/26_07/issenhuth_3dj_01_2607.pdf
and
slides 26 and 27 of
https://www.ieee802.org/3/dj/public/26_07/cole_3dj_01a_2607.pdf

After discussion there was no consensus to make related changes to the draft.

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* **180.9.8.2** *P* **489** *L* **43** # R1-286

Dawe, Piers J G NVIDIA

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

"Minimum measurement time is 60 seconds", "a value greater than 0 in counter".
1. 802.3 is not a test specification. There was an 802.3 test specification; it was a separate document which was withdrawn a long time ago. We specify observable behaviours at interfaces.
2. With not maximum measurement time and no indication of a statistical criterion, this is down to chance. It's not a specification.
I said I would TR such material.

SuggestedRemedy

If the TFT section is kept, define a proper, observable, repeatable criterion, for bins 9+. No limit for some of them would be proper.

Response *Response Status* **U**

REJECT.

The suggested remedy does not provide sufficient detail to implement.

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 180ı16, 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* **180.7.1** *P* **472** *L* **50** # R1-288

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Tx overshoot (O)*

27% overshoot is a lot. As it is on each side, the envelope is 156%of the OMA.
As SSPRQ contains 65,000 UI we are basing the measurement on the 6th-worst UI (roughly), which is on the edge of too few for a good measurement. If a signal's overshoot involves ISI but not on those UI, we would miss that. Another SSPRQ-like pattern could give a different result.

SuggestedRemedy

I wonder if 1e-4 is too ambitious and we should set the spec at e.g. 5e-3, with different limits.

Response *Response Status* **U**

REJECT.

The change on overshoot was adopted in D3.0. CRG reached consensus after reviewing the following supporting contribution with data shown in P5:
https://www.ieee802.org/3/dj/public/26_05/rodes_3dj_01b_2605.pdf

The commentor noted that in the comment the statement " 6th-worst UI" should have been " 2nd-worst UI" and in the suggested remedy that spec should have been "5e-4" not "5e-3".

The comment does not provide sufficient justification to support the suggested remedy.

The suggested remedy does not provide sufficient detail to implement.

Cl **180** *SC* **180.5.1** *P* **467** *L* **44** # R1-289

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *signal detect (CG)*

D3.0 comment 433: This shows the signal detect after the usual AND gate at the bottom of the diagram, but an unconnected SIGNAL_OK at the top of the diagram. But these are connected; the ILT/RTS functions will be asleep until signal detect reports that there is light coming in and it is worth waking up. A consideration of squelch shows that getting this right can be important. Optical signal detect protects the DSP from processing crosstalk (successfully or not) when there is no input, and it definitively keeps the receiver from getting confused if the light has disappeared and confusing those debugging the network. For these purposes it does not matter that it is slow, nor that it doesn't identify a good signal; it unambiguously identifies a too-weak or no signal. Every optical clause has had this for a long time; 187 for example still does, in the traditional way. But there might be implementations where delivering the signal detect information to the right place in RTS is impractical.

SuggestedRemedy

Show signal detect as a recommended input to "ILT/RTS functions". Several clauses.

Response *Response Status* **U**

REJECT.
This comment is a submission of comment I-436 against Draft 3.0.
Resolved using the response to comment R1-290.

Cl **180** *SC* **180.5.4** *P* **468** *L* **48** # R1-290

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *signal detect (CG)*

D3.0 comment 436: global signal detect should gate RTS. It would be advisable for the lane-by-lane signal detect to gate ILT training (no light = nothing to train on): see another comment.

SuggestedRemedy

Add sentence that Global_PMD_signal_detect should (recommended) be used as an input to the RTS function of 178A. Several optical clauses.

Response *Response Status* **U**

REJECT.
This comment is a restatement of comment I-436 against Draft 3.0.
This comment refers to another comment ("see another comment"), likely R1-289.
IEEE 802.3 is an interoperability specification, not an implementations specification.
Although use of the signal detect variable might be helpful in determining the state of the input signal it is not necessary to mandate or recommend the use of that variable for the purposes listed in the comment.
This clause defines Global_PMD_signal_detect only as a management variable and consequently bits in a management register. Use beyond that is not defined. The quality of a received optical signal is determined by the receiver. The means by which a receiver determines the quality of the optical signal is beyond the scope of this standard, though it might include use of the Global_PMD_signal_detect variable.

Cl **178B** *SC* **178B.6** *P* **878** *L* **46** # R1-291

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Signal detect (CA)*

D3.0 comment 433 and others: signal detect should gate RTS. RTS should not move to DATA mode if no light is coming in, whether or not a retimer has locked its clock to something.

SuggestedRemedy

Add text that for an optical PMD or PMD/PMA, local_rts should not be set to true if signal detect indicates that there is no optical input.

Response *Response Status* **U**

REJECT.

The resolution to comments R1-289 and R1-290 determined that signal detect would not be explicitly specified as a recommended input to the ILT and RTS functions.

There is no consensus to make the proposed changes.

Cl **178B** *SC* **178B.7.2** *P* **880** *L* **1** # R1-292

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Signal detect (CA)*

D3.0 comment 433 and 436: a receiving PMD or PMA should not claim all is well if no light is coming in.

SuggestedRemedy

Add text here and possibly in 178B.8.3 saying that if the lane-by-lane signal detect is 0, per-lane ISL should (recommended) not move to mission mode or stay there (although it may transmit onto the line).

Response *Response Status* **U**

REJECT.

The resolution to comments R1-289 and R1-290 determined that signal detect would not be explicitly specified as a recommended input to the ILT and RTS functions.

There is no consensus to make the proposed changes.

Cl **178B** *SC* **178B.7** *P* **879** *L* **41** # R1-293

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Signal detect (CA)*

This says "Recovery from the ISL_FAIL state in the training control state diagram (Figure 178Bû10b) requires management to assert mr_restart_training or reset; the timing of this action is implementation dependent and beyond the scope of this standard" which sounds permanent or at least aggravating.

SuggestedRemedy

I wonder if this should say that if the link was up and signal_detect transitions from OK to not, the state diagram should be restarted, or whether it should stay down until the return of the optical signal causes something else to restart it.

Response *Response Status* **U**

REJECT.

The resolution to comments R1-289 and R1-290 determined that signal detect would not be explicitly specified as a recommended input to the ILT and RTS functions.

There is no consensus to make the proposed changes.

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* **471** *L* **44** # R1-295

Dawe, Piers J G NVIDIA

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

D3.0 comment 466: With the current histogram spacing and width, good links can be made with non-compliant TDECQ and TECQ - because it is not a useful thing to limit, as explained before. It is a penalty, which the receiver can tolerate with no specific limit, as it just one component of required sensitivity. Optimising ECQ at the expense of performance is bad. The things that matter (at 4.56e-4) are the net useful modulation amplitude, OMA-TDECQ, and the residual penalty, K = ECQ -dBo(Ceq), and the overshoot. 800GBASE-LR4 has a limit of 3.9 dB, showing what is possible.

SuggestedRemedy

If the histogram spacing and width remains as is, increase the TDECQ and TECQ limits to 3.8 dB. The net useful modulation amplitude, OMA-TDECQ, is unchanged.
Do this for all IMDD PMDs except 800GBASE-LR4.

Response *Response Status* **U**

REJECT.
This comment is a resubmission of Draft 3.0 Comment I-439 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.

Cl **180** *SC* **180.9.6.3** *P* **482** *L* **32** # R1-297

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TDECQ editorial (Bp) (O)*

The vector w is not a set, it is a vector or "tuple" where the order matters (but credit for putting w in bold).

SuggestedRemedy

Remove the curly brackets

Response *Response Status* **U**

REJECT.
This represents a set of elements and use of "curly brackets" is a conventional way of representing a set of elements.

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* **180.9.6.3** *P* **482** *L* **18** # R1-301

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **A** *TDECQ variable (B) (O)*

The notation seems inconsistent. We have z_n and x(n); both seem to be series, or an element in a series, but they are written differently.

SuggestedRemedy

Remove the unnecessary algebra and clean up the remaining notation.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.
Resolved using the response to comment R1-94 and R1-114.

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* **180.9.8.1** *P* **489** *L* **19** # R1-306

Dawe, Piers J G NVIDIA

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

This says that an example estimate of the transmitter under test actual TDECQ value over the
test fiber is the measured DUT_TECQ value plus the best estimate of the actual transmitter under test CD penalty over the test fiber. Clearly this is wrong. The best estimate of TDECQ is TDECQ. We know that CD penalty is not exactly the same as TDECQ-TECQ.

SuggestedRemedy

If the TFT section is kept, decide which is meant and use it.

Response *Response Status* **U**

REJECT.

The suggested remedy does not provide sufficient detail to implement.

The referenced text in the draft was discussed and there was agreement that the text was helpful to the reader.

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.

Cl **176D** *SC* **176D.6.4** *P* **820** *L* **36** # I-242

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

Comment Type **TR** *Comment Status* **R** *C2M methodology (E)*

Comments submitted against D2.x requested demonstration that using transmitter out test parameters with SNDR are sufficient for interoperability. These explicit jitter value either will are too strengthen or too relaxed from the point of what receiver cares!

SuggestedRemedy

The established method in CK VEO and VEC has direct correlation to what receiver can tolerate, the EECQ is another test developed in the OIF Linear and RTLR is a composite test that determines goodness/badness from the perspective of the receiver.
TF either need to demonstrate that the current jitter limits is sufficient for interoperability otherwise consider adding VEO/VEC or EECQ.
Other standards that typically follow IEEE in this case plan to continue with VEC/VEO and some major customers also not convinced what we have for C2M is sufficient and likely will create their own specifications

Response *Response Status* **U**

REJECT.
The comment does not indicate any problem with the existing specification. Rather, it requests a demonstration that it is "sufficient for interoperability". This is a call for action, and is not implementable in the draft.

The existing methodology has been in place since D1.1 (based on the response to comment #186 against D1.0). It is aligned with the KR/CR methodology due to the addition of ILT to C2M interfaces. The methods mentioned in the suggested remedy are specified for fixed-waveform signals and do not take into account the effect of variable transmitter equalization. Measurement data of the currently specified parameters with existing test equipment has been provided in several contributions. No such data for the suggested alternative methods has been contributed.
A methodology change will have a widespread impact on the draft, since input specifications (receiver tolerance tests) are aligned with the transmitter specified paramters for calculation of COM. It would risk the project schedule.

The suggested remedy does not provide sufficient detail for the CRG to understand the proposed changes and for the editors to implement.

Cl **180** *SC* **180.9.6.1** *P* **477** *L* **44** # I-258

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

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

802.3dj has spend last 1 year to best capture jitter in TDECQ and FRx, unless xAUI-n (C2M) interface operate with jitter tolerance condition in 176D.8.14 TDECQ measurement maybe be too optimistic see also unsatisfied comment 147 D2.2

SuggestedRemedy

Add sentence: For those cases where there is an xAUI-n chip-to-chip (C2C) or chip-to-module (C2M) interface (see Table 180û1 through Table 180û4), the clock source for the test pattern is derived from the clock recovered from the xAUI-n input signal. The xAUI-n input signal is that described for the Jitter tolerance test in 176D.8.14. The pattern of the xAUI-n input signal may be PRBS31Q, scrambled idle, or a valid xBASE-R signal.
See ghiasi_3dj_02_2605

Response *Response Status* **U**

REJECT.

The following contribution was reviewed by the CRG:
https://www.ieee802.org/3/dj/public/26_05/ghiasi_3dj_02_2605.pdf

Based on straw poll TF-17 there is no consensus to adopt the proposed changes.

Straw poll TF-17 (directional)
I support adoption of the proposed change (blue text) in the top paragraph on slide 9 of ghiasi_3dj_02_2605.
Yes: 10
No: 23

Cl **181** *SC* **181.9.6** *P* **517** *L* **53** # I-259

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

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

802.3dj has spend last 1 year to best capture jitter in TDECQ and FRx, unless xAUI-n (C2M) interface operate with jitter tolerance condition in 176D.8.14 TDECQ measurement maybe be too optimistic see also unsatisfied comments aginst D2.1 and D2.2

SuggestedRemedy

Add sentence: For those cases where there is an xAUI-n chip-to-chip (C2C) or chip-to-module (C2M) interface (see Table 180û1 through Table 180û4), the clock source for the test pattern is derived from the clock recovered from the xAUI-n input signal. The xAUI-n input signal is that described for the Jitter tolerance test in 176D.8.14. The pattern of the xAUI-n input signal may be PRBS31Q, scrambled idle, or a valid xBASE-R signal.
See ghiasi_3dj_02_2605

Response *Response Status* **U**

REJECT.
Resolved using the response to comment I-258

Cl **182** *SC* **182.9.6** *P* **549** *L* **38** # I-260

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

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

802.3dj has spend last 1 year to best capture jitter in TDECQ and FRx, unless xAUI-n (C2M) interface operate with jitter tolerance condition in 176D.8.14 TDECQ measurement maybe be too optimistic see also unsatisfied comments aginst D2.1 and D2.2

SuggestedRemedy

Add sentence: For those cases where there is an xAUI-n chip-to-chip (C2C) or chip-to-module (C2M) interface (see Table 180û1 through Table 180û4), the clock source for the tes pattern is derived from the clock recovered from the xAUI-n input signal. The xAUI-n input signal is that described for the Jitter tolerance test in 176D.8.14. The pattern of the xAUI-n input signal may be PRBS31Q, scrambled idle, or a valid xBASE-R signal.
See ghiasi_3dj_02_2605

Response *Response Status* **U**

REJECT.
Resolved using the response to comment I-258

Cl **183** *SC* **183.9.6** *P* **582** *L* **46** # I-261

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

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

802.3dj has spend last 1 year to best capture jitter in TDECQ and FRx, unless xAUI-n (C2M) interface operate with jitter tolerance condition in 176D.8.14 TDECQ measurement maybe be too optimistic see also unsatisfied comments aginst D2.1 and D2.2

SuggestedRemedy

Add sentence: For those cases where there is an xAUI-n chip-to-chip (C2C) or chip-to-module (C2M) interface (see Table 180û1 through Table 180û4), the clock source for the tes pattern is derived from the clock recovered from the xAUI-n input signal. The xAUI-n input signal is that described for the Jitter tolerance test in 176D.8.14. The pattern of the xAUI-n input signal may be PRBS31Q, scrambled idle, or a valid xBASE-R signal.
See ghiasi_3dj_02_2605

Response *Response Status* **U**

REJECT.
Resolved using the response to comment I-258

Cl **180** *SC* **180.9.9** *P* **487** *L* **40** # I-262

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

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

802.3dj has spend last 1 year to best capture jitter in TDECQ and FRx, unless xAUI-n (C2M) interface operate with jitter tolerance condition in 176D.8.14 TDECQ measurement maybe be too optimistic see also unsatisfied comments aginst D2.1 and D2.2

SuggestedRemedy

Add sentence: For those cases where there is an xAUI-n chip-to-chip (C2C) or chip-to-module (C2M) interface (see Table 180û1 through Table 180û4), the clock source for the test pattern is derived from the clock recovered from the xAUI-n input signal. The xAUI-n input signal is that described for the Jitter tolerance test in 176D.8.14. The pattern of the xAUI-n input signal may be PRBS31Q, scrambled idle, or a valid xBASE-R signal.
See ghiasi_3dj_02_2605

Response *Response Status* **U**

REJECT.
Resolved using the response to comment I-258.

Cl **176D** *SC* **176D.7.2** *P* **824** *L* **41** # I-269

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

Comment Type **TR** *Comment Status* **R** *C2M reference channel (E)*

COM reference model provided is module-chip and only used for interference tolerance, but current transmitter relies on jitter output which is incomplete not having a reference COM model

SuggestedRemedy

Create an ASIC-Module COM model and also a Module-ASIC COM model. If there is interest in TF, I can provide it COM models.

Response *Response Status* **U**

REJECT.
The COM reference model applies to both host-to-module and module-to-host channels. The direction may require different inputs to the analysis (s-parameter files, port order, and possibly 802.3 COM tool configuration sheets) but these are not part of the standard.
The COM reference model is not used in jitter measurement.

The suggested remedy does not include sufficient detail to implement.

Cl **179** *SC* **179.9.4.7.2** *P* **431** *L* **38** # I-305

Dudek, Michael Marvell

Comment Type **TR** *Comment Status* **R** *Jitter (E)*

For calculating JH4u from Djdd and RJrms Q4 is used with a value of 3.891 approximating Q(Q4) of 5x10^-5 (equation 179-13). However for the calculation of ADD and sigmaRj from JH4u and Jrms in the interference tolerance test in 179.9.5.3.3 Q4d is used with a value of 3.719. approximating Q(Q4d)-1x10^-4 . That means that Add is not equal to Djdd. and sigmaRJ is not equal to RJRMS. This seems strange. (Note the same Q4d is also used in 178.9.3.4.2)

SuggestedRemedy

Reconcile this difference. Change "Q4" to "Q4d" with value 3.719 in 4 places in 179.9.4.7.2 and change the Q value to 1x10^-4.

Response *Response Status* **U**

REJECT.
Equations (179-17) and (179-18) are solutions of standard deviation from a measured quantile, and use Q4d.
The term Q4d is consistent with Q3d defined in 162.9.5.3.3, which was changed from Q3 by the resolution of comment #209 against 802.3ck D2.0 (see <www.ieee802.org/3/ck/comments/draft2p0/8023ck_D2p0_final_closedcomments.pdf#page=45> and related contributions <https://www.ieee802.org/3/ck/public/adhoc/apr14_21/hidaka_3ck_adhoc_01_041421.pdf> and <https://www.ieee802.org/3/ck/public/21_05/li_3ck_02c_0521.pdf>).

The calculation in Equation (179-13) uses Q4 to convert the parameters of the Dual Dirac model to the 5e-5 quantile on each side, to calculate the "all but 1e-4" range. This is a different calculation and it can be tested to yield a sufficiently accurate result. Using Q4d would yield the 1e-4 quantile, which would result in "all but 2e-4" range instead.

There was no consensus to make the suggested change.

Cl **176D** *SC* **176D.6.4** *P* **821** *L* **33** # I-312

Dudek, Michael Marvell

Comment Type **TR** *Comment Status* **R** *SNDR (E)*

The SNDR for the host output is measured with the worst case next agressor after the major loss of the channel. (i.e. close to the end of the channel). It is unrealistic and unnecessary to meet the values at the CR test point

SuggestedRemedy

Change the requirement (just for the host output) from "Table 179-9" to 27.5dB. On line 17 Delete "The minimum SNDR requirements in Table 179û9 apply for both module output and host output."

Response *Response Status* **U**

REJECT.

The CRG discussion indicated concerns about the proposed change, and lack of data to support it.

There was no consensus to make a change.

Cl **180** *SC* **180.1** *P* **458** *L* **6** # I-417

Dawe, Piers J G NVIDIA

Comment Type **ER** *Comment Status* **R** *editorial (Bx) (OI)*

Table 180-1 and 2, Physical Layer clauses associated with... are clones of each other. Each really has 3 columns. The status (optional, required or conditional) is the same in each. But the presentation obscures the very high level of commonality.

SuggestedRemedy

Present these 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.

Response *Response Status* **U**

REJECT.

Combining the two tables together does not improve the clarity of the draft but rather results in a table that is more difficult to read since it would need to include a two PCS rows (one for each Ethernet rate), two xMII rows (one for each Ethernet rate), etc., where only one of those rows would be relevant to any one PHY type. An example with with two Ethernet rates can be found in Table 122-1 in the base standard.

Cl **180** *SC* **180.1** *P* **460** *L* **6** # I-418

Dawe, Piers J G NVIDIA

Comment Type **ER** *Comment Status* **R** *editorial (Bx) (OI)*

Table 180-3 and 4, Physical Layer clauses associated with... are not quite clones of each other. Each really has 3 columns. The status (optional, required or conditional) is the same in each. But the presentation obscures the very high level of commonality and makes it hard to spot the differences (PCS and PMA).

SuggestedRemedy

Present these 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. There will be one more row than in Table 180-3.
Combine footnote d.

Response *Response Status* **U**

REJECT.
Resolved using the response to comment #I-417.

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 **180** *SC* **180.2** *P* **461** *L* **41** # I-420

Dawe, Piers J G NVIDIA

Comment Type **ER** *Comment Status* **R** *editorial (Bx) (OI)*

The paragraphs in this subclause are ordered top - bottom - middle.

SuggestedRemedy

Swap the second and third paragraphs so that the logic behind these requirements is more apparent. Several clauses.

Response *Response Status* **U**

REJECT.

There is not 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 **175** *SC* **175.2.4.6.1** *P* **291** *L* **10** # I-423

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *syntax (L)*

802.3 is a specification, not an example, lecture or tutorial. It is definitive and should be written in proper formal English. am_x<119:0> is not an example; it is *the* name for this thing. Other sections use better language, for example "The variable tx_am_sf is set as follows".

SuggestedRemedy

Change:
Let am_x<119:0> be the alignment marker for PCS lane x, x=0 to 15, where bit 0 is the first bit transmitted.
am_x<119:0> is constructed as shown in Figure 175-3 using the values in Table 175-4 for each PCS lane number x.
to:
The alignment marker for PCS lane x is am_x<119:0>, where x=0 to 15 and bit 0 is the first bit transmitted. It is constructed as shown in Figure 175-3 using the values in Table 175-4 for each PCS lane number x.
Scrub the draft for similar misuse of "let".

Response *Response Status* **U**

REJECT.
The sentence is grammatically correct and is a well-structured definition commonly used in technical specifications to define a variable (am_x<119:0>), its context (PCS lane x), its range (x=0 to 15), and the interpretation of its bits ("bit 0 is the first bit transmitted"). This precise phrasing is used to ensure the transmitter and receiver agree on how the alignment marker is constructed and serialized.

The CRG reviewed the slide 29-33 of https://www.ieee802.org/3/dj/public/26_05/opsasnick_3dj_01a_2605.pdf.

There are many examples of this language in the base standard, see slide 32 of opsasnick_3dj_01a_2605.pdf.

The is no consensus to make a change.

Cl **178** *SC* **178.10.1** *P* **401** *L* **16** # I-424

Dawe, Piers J G NVIDIA

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

This table contains "C_d^(1)", "C_d^(2)" and so on. A superscript denotes either exponentiation or some sort of footnote or reference, yet the former would not be consistent with the units and the latter are missing. Computer programs don't have variable names containing superscript, and such names are difficult for humans too.

SuggestedRemedy

Change to C_d1", "C_d2" and so on. Similarly R_d and so on
Scrub the draft for unnecessary superscript notation, so that it does not appear outside 178A.1, Channel Operating Margin.

Response *Response Status* **U**

REJECT.

The format of the die capacitance symbol is consistent with the convention specified in 178A and used in 179, 176C, and 176D.
It is also consistent with similar symbols defined in Annex 93A.

There is no consensus to make the suggested changes.

[CC: 178, 179, 176C, 176D]

Cl **178** *SC* **178.10.1** *P* **403** *L* **25** # I-425

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Jitter (E)*

This says "Random jitter, RMS", "sigma_RJ" yet it is clear by inspection that what is meant is Gaussian jitter, and the its cause (random, known cause, or unknown cause) is not the point. A new COM definition is our chance to follow OIF and break with this old name from 8B/10B times which in the age of scrambled coding is inaccurate and encourages sloppy thinking.

SuggestedRemedy

Change "Random jitter, RMS", "sigma_RJ" to "Gaussian jitter, RMS", "sigma_GJ" throughout. Add notes as necessary pointing out that similar quantities in older clauses (Annex 93A for example) use the old names.

Response *Response Status* **U**

REJECT.
Resolved using the response to comment I-426.

Cl **179** *SC* **179.9.4.7** *P* **448** *L* **8** # I-426

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Jitter (E)*

This says "Random jitter, RMS", "sigma_RJ" yet it is clear by inspection that what is meant is Gaussian jitter, and the its cause (random, known cause, or unknown cause) is not the point. A new COM definition is our chance to follow OIF and break with this old name from 8B/10B times which in the age of scrambled coding is inaccurate and encourages sloppy thinking.

SuggestedRemedy

Change "Random jitter, RMS", "sigma_RJ" to "Gaussian jitter, RMS", "sigma_GJ" throughout. Add notes as necessary pointing out that similar quantities in older clauses (Annex 93A for example) use the old names.

Response *Response Status* **U**

REJECT.
The term "random jitter" and associated symbols has been used in several previous places in 802.3.
The suggested change is a matter of preference.

There was no consensus to make the suggested changes.

[Editor's note: changed Page/Line from 403/25 to 448/8]
[CC: 178, 179, 176C, 176D, 178A]

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* **180.5.1** *P* **464** *L* **7** # I-433

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *signal_detect (Bx) (OI)*

This shows the signal detect after the usual AND gate at the bottom of the diagram, but an unconnected SIGNAL_OK at the top of the diagram. But these are connected; the ILT/RTS functions will be asleep until signal detect reports that there is light coming in and it is worth waking up. A consideration of squelch shows that getting this right can be important.

SuggestedRemedy

Show signal detect as an input to "ILT/RTS functions"

Response *Response Status* **U**

REJECT.

802.3dj now uses a different mechanism local_rx_ready that defines the perceived quality of the signal from the perspective of ILT.

Signal detect as defined may or may not be helpful at the discretion of the implementor.

Cl **180** *SC* **180.5.4** *P* **465** *L* **29** # I-436

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *signal_detect (Bx) (OI)*

In existing optical PMDs, the relation between signal detect function and SIGNAL_OK is simple. In this draft it is not stated in these sections where one would expect it. There should be some text explaining how the global and/or lane-by-lane signal detect function relates to SIGNAL_OK.

SuggestedRemedy

Add text explaining how the global and/or lane-by-lane signal detect function relates to SIGNAL_OK.

Response *Response Status* **U**

REJECT.

For PMDs that uses RTS and ILT function, SIGNAL_OK is dependent on the rts_status and not related to the signal detect functions. The local_rx_ready variable is used to define the perceived quality of the signal in ILT function.
Signal detect as defined may or may not be helpful at the discretion of the implementor.

Cl **180** *SC* **180.5.5** *P* **466** *L* **1** # I-437

Dawe, Piers J G NVIDIA

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

This text, which has been modified from text for 1000BASE-LX and 1000BASE-SX, says: "Implementations need to provide adequate margin between the input optical power level at which the PMD_signal_detect_i variable is set to 1 and the inherent noise level of the PMD". This could be read as a pseudo-requirement that (for decreasing signal strength), the SD should trip before the signal is too noisy to use, or after - the direction of the margin is not indicated. Actually, too noisy to use is not the requirement. The inherent noise level of the PMD is not on the same scale as the input optical power level, and for SD based on average optical power, it is barely relevant.
The paragraph would be more use in a copper clause.
"need" should be avoided.

SuggestedRemedy

Say what we actually mean: NOTE--Implementors should ensure that the PMD_signal_detect_i variable is not incorrectly set to 1 by the action of the inherent noise level of the PMD, including the effects of crosstalk, power supply noise, etc.
Or better, delete the paragraph. It is only an attempt at guidance to designers, and confusing.
Six clauses.

Response *Response Status* **U**

REJECT.

Strawpoll O-2 (directional) Chicago rules
I support
1: Leave text as is
2: Change text per suggested remedy per I-437
3: Delete the paragraph

1: 6
2: 5
3: 7

Strawpoll O-3 (directional) Pick one
I support
1: Leave text as is
2: Change text per suggested remedy per I-437
3: Delete the paragraph

1: 10
2: 1
3: 5

No consensus to make a change.

Cl **180** *SC* **180.7.1** *P* **468** *L* **28** # I-439

Dawe, Piers J G NVIDIA

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

There is a strong relation between ECQ, and histogram spacing + width. Histogram spacing is very important to protect receivers from narrow eyes caused by transmitter jitter.

SuggestedRemedy

If the histogram spacing + width remains as is, increase the TDECQ and TECQ limits to 3.8 dB. The useful signal strength, OMA-TDECQ, is unchanged.
Several clauses.

Response *Response Status* **U**

REJECT.

The comment proposes to increase the limit of TDECQ and TECQ. However, no data or simulation analysis has been provided indicating that 3.4dB is failing properly working transmitter.

Changing the TDECQ limit would require a corresponding change in the stressed receiver sensitivity limit.

Insufficient evidence was provided to justify to proposed changes.

There is no consensus to make the proposed changes.

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* **36** # I-441

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **A** *Overshoot (OI)*

The transmitter overshoot and undershoot limit is far too high (D2.2 comments 223 and 252). A signal with overshoot like this has been badly set up. It is outside the range of signal shapes a receiver is expected to see, as well as wasting transmitter headroom and receiver ADC resolution.

SuggestedRemedy

For the current 1e-2 hit ratio, the limit should be reduced to from 22% to 13%, and transmitter power excursion should be reduced from 2.3 dBm to 1.8 dBm. See another comment for the hit ratio.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

The resolution to comment I-248 changed the hit ratio for overshoot/undershoot measurement to 1E-4 and adopted a new value for overshoot/undershoot (max).

Further evidence and data is required to make a change to power excursion (max) value. There was no consensus to change either the measurement hit ratio or the value for power excursion (max).

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* **180.7.1** *P* **468** *L* **47** # I-445

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *RINxxOMA (Bx) (OI)*

Transmitter noise is included in TDECQ and TECQ. We introduced a RIN limit too, -132 dB/Hz at 50G/lane, because we don't attempt to find the shape of the RIN spectrum, which determines how the RIN is amplified by the equalizer, so we just keep fairly small. At 100G/lane, the limit is -136 (more noise bandwidth). At 200G/lane, we reduced the limit by another 3 dB, keeping things in proportion. Then we added the DFE (in the reference receiver and in real DSP), so there is much less noise enhancement. TDECQ and TECQ pick up electrical noises such as from power supplies as well as true RIN - these are not at high enough frequencies to undergo noise enhancement.

SuggestedRemedy

Change the RIN limit to -137 dB/Hz.
Four clauses.

Response *Response Status* **U**

REJECT.

Noise enhancement dependent on DSP implementation, it is not necessarily that DFE is implemented in a commercial DSP. Modern CW laser or EML could support -139dB/Hz which helps in providing clean transmitter signal. Data or evidence is needed to reduce the RIN limit.

There is no consensus to make the proposed change.

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* **180.7.2** *P* **470** *L* **9** # I-447

Dawe, Piers J G NVIDIA

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

Tweaking the stressed receiver sensitivity for a change to max TDECQ and TECQ (see another comment).

SuggestedRemedy

Change the stressed sensitivity from -0.9 to -0.5 and change SECQ from 3.4 to 3.8 dB.
Similarly in other IMDD clauses.

Response *Response Status* **U**

REJECT.

Resolved using the response to comment #I-439.

Cl **180** *SC* **180.7.3** *P* **471** *L* **7** # I-448

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Power budget (OI)*

Tweaking the budget for changes to max TDECQ and TECQ and MPI (see other comments).

SuggestedRemedy

Change the power budget from 6.7 to 7.2 dB.
Change the allocation for penalties from 3.7 to 4.2 dB
In table footnotes b and c, change 0.1 dB to 0.2 dB.
Similarly in other IMDD clauses.

Response *Response Status* **U**

REJECT.

This comment is partially resolved with the response to comments I-439 and I-442.

The comment does not provide sufficient evidence to support the proposed changes.

There is no consensus to make the proposed changes.

Cl **180** *SC* **180.9.2** *P* **476** *L* **5** # I-452

Dawe, Piers J G NVIDIA

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

The choice of fourth-order Bessel-Thomson response is a historical accident. Fifth order rolls off a bit steeper at high frequencies, which is more realistic, and has no other significant difference. The better roll-off leads to better measurements. This is also true for measurements where reflections at an instrument's electrical connectors are a concern. This is our chance to make the change, as the industry transitions to a new speed and new test equipment.

SuggestedRemedy

Change fourth-order to fifth-order throughout the draft.

Response *Response Status* **U**

REJECT.

The BT filter response is used to bound the implementation of a reference reciever, which typically are implemented in test equipement. Fourth order BT response has been used in previous signaling rates, changing it to a fifth order BT response may require re-validation of existing implementations.

No evidence has been provided in the comment showing that the current spec as is will cause false result of transmitter characteristic measurements.

Cl **180** *SC* **180.9.6.1** *P* **477** *L* **46** # I-457

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *test pattern (Bx) (CO)*

SSPRQ contains a section for challenging CDRs. There is no point challenging the CRU in the scope; that only degrades the measurement. SSPRQ is not used for testing receivers. So it must be there to challenge the PMAs between the PCS and the Tx PMD. Yet this says "The pattern of the xAUI-n input signal may be PRBS31Q, scrambled idle, or a valid xBASE-R signal" - which I believe is for practical considerations in ICs, some years ago.

SuggestedRemedy

Change "The pattern of the xAUI-n input signal may be PRBS31Q, scrambled idle, or a valid xBASE-R signal."
to "The pattern of the xAUI-n input signal should be SSPRQ on the lane under test, and test patterns 3, 5 or 6 on the other lanes. However, test patterns 3, 5 or 6 may be used on the lane under test.

Response *Response Status* **U**

REJECT.

The SSPRQ test pattern was adopted in IEEE 802.3bs specifically for testing of TDECQ (for optical transmitters) and SECQ (for optical receiver test calibration). See comment #50 here:

http://www.ieee802.org/3/bs/comments/P802d3bs_D1p3_comments_final_ID.pdf

The high stress and low-frequency content was intended to challenge the "low-frequency" control loops that might be implemented within an optical transmitter. It was not intended for receiver testing. There is no assurance that AUI or PMD receiver will be able to effectively track SSPRQ without penalty compared to other test patterns. Note that the analyses referenced in comment #50 shows that the phase stress is considerably higher than for the patterns specified for RX testing.

There is no consensus to make to proposed changes.

Cl **180** *SC* **180.9.6.3** *P* **479** *L* **36** # I-465

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *Reference equalizer (OI)*

A floating main tap is not attractive, whether in a reference equalizer (software) or a real product equalizer - it is more trouble than small floating taps well after the main tap.

SuggestedRemedy

Fix the main tap position e.g. at position 4, and add up to 3 more FFE taps as necessary to preserve the FFE's effectiveness. Remove the variable a in Figure 180-10.

Response *Response Status* **U**

REJECT.

Based on straw poll TF-9 and TF-10 there is no consensus to make any of the changes proposed in the suggested remedy.

Straw poll TF-9 (choose 1) and TF-10 (chicago) (directional)
I support changing the FFE length and number of equalizer pre-cursor taps for the reference equalizer taps as follows:
A: no change
B: set number of pre-cursor taps to fixed number 3 and leave FFE length at 15
C: set number of pre-cursor taps to fixed number 3 and change FFE length to 18

TF-9 (choose 1) A: 39 B: 1 C: 6
TF-10 (chicago) A: 39 B: 5 C: 16

Cl **180** *SC* **180.9.6.3** *P* **479** *L* **37** # I-467

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *tap weight (OI)*

The tap weights were created by drawing limit lines that passed a population of transmitters, some time ago. The main tap maximum and the other FFE tap limits need review, taking the DFE into account and considering what is cost-effective for receivers. In particular, the we will see less of the characteristic alternating FFE tap weights because the DFE can do that job better, so taps 4+ should be tightened up.

SuggestedRemedy

Review the tap weight limits and make changes as appropriate.

Response *Response Status* **U**

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

Cl **180** *SC* **180.9.6.3** *P* **480** *L* **17** # I-468

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *tap weight-equation (OI)*

Causal vs. anticausal pulse shapes
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/dk/public/2511/3dk_dawe_2511_1.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, but the opposite is not.

SuggestedRemedy

Remove the absolute bars | |

Response *Response Status* **U**

REJECT.

If the absolute bars are removed then the associated limit likely needs to be adjusted. Further work to justify the change and the associated limit is required.

There is no consensus to make the proposed change.

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* **180.9.6.4** *P* **481** *L* **27** # I-472

Dawe, Piers J G NVIDIA

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

The apparent TDECQ or TECQ is a very strong function of histogram window spacing. Reducing the spacing does not make the signal any better, but it lowers our guard against signals with a jitter problem. This should be treated very carefully.

SuggestedRemedy

As in another comment, increase the headline TDECQ, TECQ limit from 3.4 to 3.8 dB. Reduce the histogram window spacing from 0.1 UI between centres to 0.09 UI (if sampling at 50 s/UI is feasible) or 3/32 = 0.09375 UI (if sampling at 32/UI is preferred).

Response *Response Status* **U**

REJECT.

First, since there is another comment #I-439 specifically on TDECQ limit, response to this comment would focus on the histogram window spacing.
In theory, the two histogram window spacing can be arbitrary. Reducing the histogram spacing could reduce the reported TDECQ value. But the effect of changing the spacing by less than 0.01 UI is questionable. Note that, previous contribution https://www.ieee802.org/3/dj/public/25_11/rodes_3dj_01a_2511.pdf provides data of TDECQ comparing impact of histogram spacing and SER target. It showed for a transmitter of 2.5dB TDECQ, changing spacing from 0.1UI to 0.08UI would reduce TDECQ to 2.44dB. The result of using 0.09UI or 0.09375UI will be closer to 2.5dB, and the difference could be within measurement variation.

However, there is value in reconsidering the factors that could affect TDECQ value, such as histogram window spacing and SER target. Contributions are encouraged to support further consensus building on this topic.

Cl **180** *SC* **180.9.6.4** *P* **481** *L* **29** # I-473

Dawe, Piers J G NVIDIA

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

The histogram window width should be chosen carefully. A wide width allows points in the windows to be collected more quickly, but the samples towards the time-centre of the eye are benign samples that just dilute the measurement, making it less discriminating. The evidence in rodes_3dj_01_2601.pdf slide 4 (which is for TDECQcer) shows that a narrower window improves repeatability. Common practice is to collect very close to 32, 64 or 128 samples/UI, for which 0.04 is not optimal.

SuggestedRemedy

Make the histogram window a little narrower according to what sampling ratio is appropriate: 1/32 = 0.03125 UI for 32 s/UI or 0.02 UI for 50 s/UI.

Response *Response Status* **U**

REJECT.

There is some agreement that changes in the direction of the suggested remedy may be warranted. However, further analysis and a complete proposal are required.

There is no consensus to make the proposed change.

Cl **180** *SC* **180.9.6.4** *P* **483** *L* **23** # I-476

Dawe, Piers J G NVIDIA

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

The appropriate probability for ECQ, for the right correlation, should be considered. It won't be the same as the target pre-FEC SER because of geometrical factors. The target pre-FEC SER should not be 4.56e-4 anyway; we know that is not adequate for the FLR objective.

SuggestedRemedy

Choose a more appropriate probability. Adjust Qt consistent with this.

Response *Response Status* **U**

REJECT.

The suggested remedy does not provide sufficient detail to implement.

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-483

Dawe, Piers J G NVIDIA

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

If this mask were useful, there should be a figure showing it. But actually, any transmitter-receiver pair will have the least margin at bin 8 so none of the others matter.

SuggestedRemedy

If there will be a thought-through mask, illustrate it.

Response *Response Status* **U**

REJECT.

The commenter indicated that the intent of the suggested remedy was to add a figure illustrating the mask.

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

There is no consensus to make the proposed changes.

Straw poll O-4 (directional)
I support adding a figure illustrating the transmitter symbol error mask.
Y: 12
N: 22

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-486

Dawe, Piers J G NVIDIA

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

AGCs have levels but attenuators have attenuations

SuggestedRemedy

Change "VOA_level " to "OA" for optical attenuation

Response *Response Status* **U**

REJECT.

There is no consensus to make the proposed change.

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.