

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

CI 136 SC 136.11.7 P 235 L 18 # i-60
 RAN, ADEE Intel Corporation

Comment Type TR Comment Status R

Package transmission line characteristic impedance is set at 90 Ohm. This is an increase from the default value in Annex 93A which is 78.2 Ohm.

The reason for the relatively low value 78.2 Ohm was that to typical packages (especially large ones with many lanes) have lower impedance to improve their matching to silicon and ball impedances, and to reduce the trace insertion loss. This is not expected to change; most practical packages will not have impedance close to 100 Ohm.

In practice, termination can be adjusted and board design can be optimized to match lower impedance package and improve performance (even if cables are 100 Ohm)

It is suggested to acknowledge the expected lower impedance of practical devices in the reference package and termination parameters: assume packages are 80 Ohm while termination and board are 90 Ohm (imperfect matching).

Also applies in 137.10 (Table 137-5).

SuggestedRemedy

In both Table 136-15, and Table 137-5, change the value of Zc to 80 Ohm and Rd to 45 Ohm.

In 136.11.7.1, add an exception to the parameter values from Table 92-12: Z_c is set to 90 Ohm.

Consider changing the reference impedance for channels from 100 Ohm to 85 Ohm (136.11.1 and 137.10, and COM tables).

Response Response Status U

REJECT.

The response to comment i-161 resulted in different changes than the ones in the suggested remedy.

CI 139 SC 139.7.9.1 P 298 L 45 # i-82
 Liu, Hai-Feng Intel Corporation

Comment Type TR Comment Status R

PAM4 test results have shown (see chang_3cd_01_1117, particularly p. 20) that the composition and ratio of the stressors in the stressed receiver sensitivity test has a strong impact on link performance. In particular, the same SECQ can generate widely varying BER performance from the same receiver depending on whether the dominant stressor added to the bandwidth filtering was Gaussian noise or sinusoidal interferer. To address this we propose to more specifically prescribe the stressor ratio used to create the stressed Rx sensitivity conformance test input, to avoid understressing the receiver and causing interoperability issues.

SuggestedRemedy

In the second paragraph of section 139.7.9.1, after the existing sentence "The combination of the low-pass filter and the E/O converter should...", add the sentence "Of the remaining dB value of stressed eye closure (SECQ), at least half should be from the Gaussian noise stressor."

Response Response Status U

REJECT.

http://www.ieee802.org/3/cd/public/Nov17/chang_3cd_01_1117.pdf showed good correlation between SECQ and Rx sensitivity and the freedom to set up the SRS stress was explored quite thoroughly.

The freedom to set up the SRS test source is a balance between pragmatism and precision; the SECQ test metric ensures that the penalty (for the reference equalizer) of the induced stresses for different test source set-ups, is identical.

A late presentation http://www.ieee802.org/3/cd/public/Jan18/schube_3cd_01a_0118.pdf was reviewed also addressing the claimed problem. There was no consensus to make a change to the draft and further work was necessary to investigate the problem and provide a complete proposed remedy.

[Editor's note: Comment i-58 addresses a similar issue.]

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

Cl 139 SC 139.7.9.2 P 299 L 54 # i-83
Liu, Hai-Feng Intel Corporation

Comment Type TR Comment Status R

[note that a comment is needed in this section in addition to the comment above to avoid any confusion with the less clear instructions in the referenced 802.3bs section 121.8.9.2] PAM4 test results have shown (see chang_3cd_01_1117, particularly p. 20) that the composition and ratio of the stressors in the stressed receiver sensitivity test has a strong impact on link performance. In particular, the same SECQ can generate widely varying BER performance from the same receiver depending on whether the dominant stressor added to the bandwidth filtering was Gaussian noise or sinusoidal interferer. To address this we propose to more specifically prescribe the stressor ratio used, to avoid understressing the receiver and causing interoperability issues.

SuggestedRemedy

Add the following sentence to the end of section 139.7.9.2: "As outlined in section 139.7.9.1 above, half of the dB value of stressed eye closure (SECQ) should be from bandwidth limitations from the low-pass filter and E/O converter, while of the remaining dB value of stressed eye closure (SECQ), at least half should be from the Gaussian noise stressor."

Response Response Status U

REJECT.

See response to comment i-82

[Editor's note added after comment resolution completed:

For reference, the response to comment i-82 is copied here:

REJECT.

[Http://www.ieee802.org/3/cd/public/Nov17/chang_3cd_01_1117.pdf](http://www.ieee802.org/3/cd/public/Nov17/chang_3cd_01_1117.pdf) showed good correlation between SECQ and Rx sensitivity and the freedom to set up the SRS stress was explored quite thoroughly.

The freedom to set up the SRS test source is a balance between pragmatism and precision; the SECQ test metric ensures that the penalty (for the reference equalizer) of the induced stresses for different test source set-ups, is identical.

A late presentation http://www.ieee802.org/3/cd/public/Jan18/schube_3cd_01a_0118.pdf was reviewed also addressing the claimed problem. There was no consensus to make a change to the draft and further work was necessary to investigate the problem and provide a complete proposed remedy.

[Editor's note: Comment i-58 addresses a similar issue.]

]

Cl 140 SC 140.7.9 P 320 L 15 # i-84
Liu, Hai-Feng Intel Corporation

Comment Type TR Comment Status R

PAM4 test results have shown (see chang_3cd_01_1117, particularly p. 20) that the composition and ratio of the stressors in the stressed receiver sensitivity test has a strong impact on link performance. In particular, the same SECQ can generate widely varying BER performance from the same receiver depending on whether the dominant stressor added to the bandwidth filtering was Gaussian noise or sinusoidal interferer. To address this we propose to more specifically prescribe the stressor ratio used to create the stressed Rx sensitivity conformance test input, to avoid understressing the receiver and causing interoperability issues.

SuggestedRemedy

Add the following bullet to the end of section 140.7.9, "Of the remaining half of stressed eye closure (SECQ) that is not generated by bandwidth limitations from the low-pass filter and E/O converter, at least half of the remaining stress (in dB of SECQ) should be from the Gaussian noise stressor."

Response Response Status U

REJECT.

See resolution to comment i-82

[Editor's note added after comment resolution completed:

For reference, the response to comment i-82 is copied here:

REJECT.

[Http://www.ieee802.org/3/cd/public/Nov17/chang_3cd_01_1117.pdf](http://www.ieee802.org/3/cd/public/Nov17/chang_3cd_01_1117.pdf) showed good correlation between SECQ and Rx sensitivity and the freedom to set up the SRS stress was explored quite thoroughly.

The freedom to set up the SRS test source is a balance between pragmatism and precision; the SECQ test metric ensures that the penalty (for the reference equalizer) of the induced stresses for different test source set-ups, is identical.

A late presentation http://www.ieee802.org/3/cd/public/Jan18/schube_3cd_01a_0118.pdf was reviewed also addressing the claimed problem. There was no consensus to make a change to the draft and further work was necessary to investigate the problem and provide a complete proposed remedy.

[Editor's note: Comment i-58 addresses a similar issue.]

]

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

Cl 135F SC 135F.3 P 367 L 18 # i-98
 Rysin, Alexander Mellanox Technologie

Comment Type TR Comment Status R ERL

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

SuggestedRemedy

Change 135F.3.1 from "A 50GAUI-1 C2C or a 100GAUI-2 C2C transmitter shall meet all specifications in 120D.3.1" to "A 50GAUI-1 C2C or a 100GAUI-2 C2C transmitter shall meet all specifications in 120D.3.1 with the following exceptions: Effective Return Loss (ERL) is calculated with Nb set to 10 (see Annex New). ERL shall be at least 16.2 dB. The Transmitter Output residual ISI SNR_ISI and the return loss specifications in Table in Table 120D-1 do not apply."

Change 135F.3.2 from "A 50GAUI-1 C2C or a 100GAUI-2 C2C receiver shall meet all specifications in 120D.3.1" to "A 50GAUI-1 C2C or a 100GAUI-2 C2C transmitter shall meet all specifications in 120D.3.2 with the following exceptions: Effective Return Loss (ERL) is calculated with Nb set to 10 (see Annex New). ERL shall be at least 16.2 dB. There is no frequency domain return loss mask."

Response Response Status U

REJECT.

Although ERL was adopted for clauses 137 and 136, it is not clear whether it should be adopted for Annex 135F, since its electrical characteristics were intended to be essentially identical to 120D.

There is no consensus to implement the suggested remedy.

Cl 138 SC 138.8.5 P 274 L 39 # i-116
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

It seems that it is possible to make a bad transmitter (e.g. with a noisy or distorted signal), use emphasis to get it to pass the TDECQ test, yet leave a realistic, compliant receiver with an unreasonable challenge, such as high peak power, high crest factor, or a need to remove emphasis from the signal, contrary to what equalizers are primarily intended to do. Note the receiver is tested for a very slow signal only, not for any of these abusive signals. This is an issue for all the PAM4 optical PMDs, although it may be worse for MMF because of the high TDECQ limit.

SuggestedRemedy

1. To screen for noisy or distorted signals with heavy emphasis Define TDECQrms = $10 \cdot \log_{10}(A_RMS/(s^3 \cdot Qt \cdot R))$ where A_RMS is the standard deviation of the measured signal after the 13.28125 GHz filter response, Qt and R are as already in Eq 212-12. s is the standard deviation of a fast clean signal with OMA=2 and without emphasis, observed through the 13.28125 GHz filter response (around 0.7). Set limit for TDECQrms according to what level of dirty-but-emphasised signal we decide is acceptable, add max TDECQrms row to each transmitter table. Alternatively, if the same relative limit is acceptable for all PAM4 optical PMDs, the limit could be here in the TDECQ procedure. Similarly in clauses 139, 140.
2. To protect the TIA input, consider a peak power spec as in Clause 86.
3. To protect the TIA and any AGC and TIA from unreasonable signals, consider a crest factor spec.
4. To protect the receiver from having to "invert" heavily over-emphasised signals, set a minimum cursor weight. To protect the equalizer from having to support unnecessary settings for waveforms that can't or shouldn't ever happen, constrain the cursor position - see other comments .

Response Response Status U

REJECT.

The need for additional transmitter specs has not been established, and insufficient evidence has been provided that the proposed remedy fixes the claimed problem.

A contribution is invited that demonstrates the problem (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that the proposed additional requirement prevents this issue from occurring. A similar proposal to create a TDECQrms spec was suggested in comment #r02-35 against 802.3bs D3.2, which was similarly rejected.

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

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

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

Cl 138 SC 138.7.1 P 272 L 17 # i-119
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

A TDECQ limit of 4.9 seems very high, given that the same fibres and transmitter and receiver front-ends that should not be worse can do 100GBASE-SR4 (PAM2, almost the same signalling rate) without the FFE.

SuggestedRemedy

This needs more study. We should be able to use information from 802.3bm.

Response Response Status U

REJECT.

No change to document suggested.

The issue caused by a TDECQ limit of 4.9 dB has not been clarified. There is precedence for this kind of transmitter quality metric to be higher in MMF specifications than in SMF specifications.

Cl 138 SC 138 P 261 L 1 # i-122
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

This clause has received next to no attention - it's still the baseline. It needs more (some) study.

SuggestedRemedy

Do the work. Show technical feasibility for the draft spec (after improvements).
 The alternative is to withdraw the clause, which would be a pity.

Response Response Status U

REJECT.

No change to document suggested.

The presentation http://www.ieee802.org/3/cd/public/Jan18/king_3cd_02_0118.pdf was reviewed and provides supporting evidence for the specification in Clause 138.

Cl 136 SC 136.6.1 P 202 L 19 # i-123
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R skew <cc>

The Skew at SP4 (the receiver MDI) has to be the same as the Skew at SP3 (the transmitter MDI) for these serial PMDs.

SuggestedRemedy

Correct the numbers at SP4 and SP5. Correct Table 131-5, Summary of Skew constraints - all 50GBASE-R PMDs are serial so it's simple to do. Also 137.6.1 138.3.2.1 139.3.2.

Response Response Status U

REJECT.

The skew constraints for 100G in Table 80-5 and for 50G in Table 131-5 are consistent with the budget and methodology adopted by 802.3ba and 802.3bg and used in subsequent projects (e.g., 802.3bm, 802.3bs).

The skew constraints are established to ensure that the FEC/PCS skew tolerance is sufficient to support the worst case skew for any currently specified or potential (within reason) future PHY (e.g., 2-lane PMD for reach longer than 40 km). This is accomplished by having the same skew constraint at SP5 regardless of the PMD type.

The skew constraint at SP5 includes allocation for skew accumulated through the TX PMD (SP2 to SP3), the medium (SP3 to SP4), and the RX PMD (SP4 to SP5). Rather than specifying unique values for SP3, SP4, and SP5 based on PMD type, the adopted approach was to use the same numbers for all PMD types for consistency.

The approach described above is consistent for all PHY types defined by 802.3ba and subsequent projects. For instance, the medium skew accumulation (SP3 to SP4) of 80 ns was based on an 80 km multi-lane optical PMD. Nevertheless, the same value is used for other PMDs where the skew would be considerably lower (e.g., 100GBASE-SR4, 100GBASE-KR4, 100GBASE-CR4, etc.).

This specification methodology does not preclude an engineered implementation that optimizes the FEC/PCS skew buffering based on assumed lower PMD and medium skew accumulation. However, it should be noted that this implementation would not be compliant to 802.3cd.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

Cl 140 SC 140.3.2 P 311 L 49 # i-125
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R Skew <cc>

The Skew at SP4 (the receiver MDI) has to be the same as the Skew at SP3 (the transmitter MDI) for this serial PMD.

SuggestedRemedy

Correct the numbers at SP4 and SP5. Correct Table 80-5, Summary of Skew constraints, at least for SP2-6, e.g. by using Table 131-5 (corrected) for 100G serial.

Response Response Status U

REJECT.

Resolve with the response to comment i-123.

[Editor's note: For reference, the response to comment i-123 is copied here:

REJECT.

The skew constraints for 100G in Table 80-5 and for 50G in Table 131-5 are consistent with the budget and methodology adopted by 802.3ba and 802.3bg and used in subsequent projects (e.g., 802.3bm, 802.3bs).

The skew constraints are established to ensure that the FEC/PCS skew tolerance is sufficient to support the worst case skew for any currently specified or potential (within reason) future PHY (e.g., 2-lane PMD for reach longer than 40 km). This is accomplished by having the same skew constraint at SP5 regardless of the PMD type.

The skew constraint at SP5 includes allocation for skew accumulated through the TX PMD (SP2 to SP3), the medium (SP3 to SP4), and the RX PMD (SP4 to SP5). Rather than specifying unique values for SP3, SP4, and SP5 based on PMD type, the adopted approach was to use the same numbers for all PMD types for consistency.

The approach described above is consistent for all PHY types defined by 802.3ba and subsequent projects. For instance, the medium skew accumulation (SP3 to SP4) of 80 ns was based on an 80 km multi-lane optical PMD. Nevertheless, the same value is used for other PMDs where the skew would be considerably lower (e.g., 100GBASE-SR4, 100GBASE-KR4, 100GBASE-CR4, etc.).

This specification methodology does not preclude an engineered implementation that optimizes the FEC/PCS skew buffering based on assumed lower PMD and medium skew accumulation. However, it should be noted that this implementation would not be compliant to 802.3cd.

]

Cl 137 SC 137.9.2 P 251 L 29 # i-138
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R Tx electrical

Signal-to-noise-and-distortion ratio (min), increased to 33.3 dB (Clause 136) and to 32.5 dB (Clause 137) for all Tx emphasis settings, is still too high. D2.0 comment 139, D2.1 comment 50. It turns out that the SNDR method captures sort of "high frequency distortion" that is filtered out by a real channel and receiver 3fb/4 bandwidth (see 93A.1.4.1), partly un-filtered by the equalizer. So it should be measured in something less than ~19 GHz.

SuggestedRemedy

Add ", when sigma_e and sigma_n are found from signals observed with a fourth-order Bessel-Thomson low-pass response with 19.34 GHz 3 dB bandwidth.

NOTE--pmax is found from a signal observed with a fourth-order Bessel-Thomson low-pass response with 33 GHz 3 dB bandwidth."

If we wish, we can tweak the limit for pmax and measure it in the same 19.34 GHz, which would more correctly remove the harmonics from the measurement.

Response Response Status U

REJECT.

The sigma_TX term in COM is calculated under the assumption that the spectrum of the noise and the distortion is identical to the spectrum of the ideal signal at the transmitter output (sinc shaped per Eq. 93A-23). If that is the case, the signal, noise and distortion all go through the same transfer function, which includes the transmitter, receiver, and channel (Eq. 93A-19).

The actual effect on the receiver depends on the Tx noise and distortion spectrum (if high frequencies dominate, sigma_tx is too high because they will be more attenuated by channel and Rx than the signal; if low frequencies dominate, sigma_tx is too low since they will be less attenuated).

The suggested remedy includes a specific new filter for noise and distortion measurement but there is insufficient evidence that this filter is more suitable than the current filter.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

Cl 135F SC 135F.3 P 408 L 27 # r01-27
 Rysin, Alexander Mellanox Technologie

Comment Type TR Comment Status R ERL AUI

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

SuggestedRemedy

Change 135F.3.1 from "A 50GAUI-1 C2C or a 100GAUI-2 C2C transmitter shall meet all specifications in 120D.3.1" to: "A 50GAUI-1 C2C or a 100GAUI-2 C2C transmitter shall meet all specifications in 120D.3.1 with the following exceptions:
 * Effective return loss (ERL) of the transmitter at TP0a is computed using the procedure in 93A.5 with the values in Table 137-5. Parameters that do not appear in Table 137-5 take values from Table 120D-8. The value of Tfx is twice the delay from TP0 to TP0a. Nbx is set to the value of Nb in Table 120D-8. ERL shall be at least 16.1 dB. The Transmitter Output residual ISI SNR_ISI and the return loss specifications in Table in Table 120D-1 do not apply."

Change 135F.3.2 from "A 50GAUI-1 C2C or a 100GAUI-2 C2C receiver shall meet all specifications in 120D.3.1" to: "A 50GAUI-1 C2C or a 100GAUI-2 C2C receiver shall meet all specifications in 120D.3.2 with the following exceptions:
 * Effective return loss (ERL) of the receiver computed using the procedure in 93A.5 with the values in Table 137-5. Parameters that do not appear in Table 137-5 take values from Table 120D-8. The value of Tfx is twice the delay from TP5a to TP5. Nbx is set to the value of Nb in Table 120D-8. ERL shall be at least 16.1 dB.

Response Response Status U

REJECT.

This comment is similar to the unsatisfied comment i-98. The response to that comment was:
 "Although ERL was adopted for clauses 137 and 136, it is not clear whether it should be adopted for Annex 135F, since its electrical characteristics were intended to be essentially identical to 120D.
 There is no consensus to implement the suggested remedy."

Straw Poll ET-4

I support adding the ERL specifications to Annexes 135D, 135E, 135F, or 135G.

A. Yes: 3
 B. No: 17

There is no consensus to make the proposed change.

Cl 137 SC 137.9.2 P 281 L 28 # r01-28
 Rysin, Alexander Mellanox Technologie

Comment Type TR Comment Status A SNR_ISI

Requirements for Transmitter output residual ISI SNR_ISI (min) of 43 dB and SNDR (min) of 32.5 dB in Clause 137 is too high - can barely measure the IC through the test fixture. The warning NOTE in 120D.3.1.7 shows the issue, but doesn't solve it. The limits for SNR_ISI in Clause 137 are even more stringent than in 120D. COM packages were shown to generate worse SNDR and SNR_ISI for the target SNR_TX. D2.0 comment 140, D2.1 comment 49, D2.2 comment 22, D3.0 comments 71, 74, 97. Previous comments, suggesting ERL should replace SNR_ISI suggest a partial remedy.

SuggestedRemedy

* Change paragraph 3 in 137.9.2 from "SNR_ISI is computed with Nb set to 12 and Dp set to 3. The value of SNR_ISI (min) is 43 dB." to "SNR_ISI is computed with Nb set to 12 and Dp set to 3. The value of SNR_ISI (min) is 30.5 dB"
 * Change paragraph 4 in 137.9.2 from "The value of SNDR (min) is 32.5 dB." to "The value of SNDR (min) is 32 dB".

See presentation.

Response Response Status U

ACCEPT IN PRINCIPLE.

To address SNR_ISI, implement the changes proposed in http://www.ieee802.org/3/cd/public/Mar18/dudek_3cd_02_0318.pdf with editorial license.

Relative to SNDR, http://www.ieee802.org/3/cd/public/Mar18/rysin_3cd_01_0318.pdf was reviewed.

There was no consensus to make a change to SNDR.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

Cl 131 SC 131.5 P 131 L 12 # r01-51
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R skew variation <cc>

This Table 131-6 (Skew Variation) still does not agree with e.g. 138.3.2.1, which says "Since the signal at XX represents a serial bit stream, there is no Skew Variation at this point". All 50GBASE-R PMDs are serial.

SuggestedRemedy

Either:
 Delete the rows for SP2 to SP5, adding a table note to explain that there is no SV at those points; or:
 For SP2, delete the reference to 135.5.3.5, which is not relevant for a serial PMA/PMD interface,
 For SP5, delete the reference to 135.5.3.6, which is not relevant for a serial PM/PMA interface, and
 for SP2 to SP5, change the numbers to N/A.

Response Response Status U

REJECT.

The specifications at SP2 and SP5 ensure that the PMA is compatible with any current or future PHY. The specifications at SP3 and SP4 provide the skew variation limits for the net budget that are assumed for any future PHY that might have a 50GAUI with more than one lane and to be consistent with the budget methodology used for 40G, 100G, 200G, and 400G in base standard.

The references to 135.5.3.5/6 are retained as they would be relevant to any future multi-lane PMD.

This specification methodology does not preclude an engineered implementation that optimizes the FEC/PCS skew buffering based on assumed lower PMA, PMD, and medium skew variation. However, it should be noted that such an implementation would not be compliant to 802.3cd.

[Editor's note: Comments r01-51, r01-53, r01-55, r01-56, and r01-58 from the same commenter relate to a similar topic.]

Cl 135 SC 135.5.3 P 177 L 49 # r01-53
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R skew variation <cc>

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

SuggestedRemedy

Change:
 The limits for Skew and Skew Variation at physically instantiated interfaces are specified at Skew points SP0, SP1, and SP2 in the transmit direction and SP5, SP6, and SP7 in the receive direction as defined in 131.5 and illustrated in Figure 131-3 for 50GBASE-R and as defined in 80.5 and illustrated in Figure 80-8 for 100GBASE-P. to:
 For 50GBASE-R, the limits for Skew at physically instantiated interfaces are specified at Skew points SP0, SP1, and SP2 in the transmit direction and SP5, SP6, and SP7 in the receive direction as defined in 131.5 and illustrated in Figure 131-3. For 50GBASE-R, the limits for Skew Variation at physically instantiated interfaces are specified at Skew points SP0 and SP1 in the transmit direction, and SP6 and SP7 in the receive direction, as defined in 131.5 and illustrated in Figure 131-3. For 100GBASE-P, the limits for Skew and Skew Variation at physically instantiated interfaces are specified at Skew points SP0, SP1, and SP2 in the transmit direction and SP5, SP6, and SP7 in the receive direction as defined in 80.5 and illustrated in Figure 80-8 for 100GBASE-P.

Response Response Status U

REJECT.

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

The description in the referenced paragraph relates to currently specified as well as any future PMD which may have more than one lane. The specific requirements for each PMD are specified in the PMD clause.

The beginning of the paragraph points out that skew variation only applies to cases with multiple lanes. "Any PMA that combines PCSLs/FECLs from different input lanes onto the same output lane must tolerate Skew Variation between the input lanes without changing the PCSL/FECL positions on the output."

[Editor's note: Comments r01-51, r01-53, r01-55, r01-56, and r01-58 from the same commenter relate to a similar topic.]

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

Cl 135 SC 135.5.3.5 P 179 L 12 # r01-55

Dawe, Piers J G

Mellanox Technologie

Comment Type TR Comment Status R skew variation <cc>

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

SuggestedRemedy

Change:

... 43 ns of Skew, and no more than 0.4 ns of Skew Variation ... to:

... 43 ns of Skew, and, for 100GBASE-P, no more than 0.4 ns of Skew Variation ...

Response Response Status U

REJECT.

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

The description in the referenced paragraph relates to currently specified as well as any future PMD which may have more than one lane. The specific requirements for each PMD are specified in the PMD clause.

[Editor's note: Comments r01-51, r01-53, r01-55, r01-56, and r01-58 from the same commenter relate to a similar topic.]

Cl 135 SC 135.5.3.6 P 179 L 17 # r01-56

Dawe, Piers J G

Mellanox Technologie

Comment Type TR Comment Status R skew variation <cc>

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

SuggestedRemedy

Change:

135.5.3.6 Skew tolerance at SP5

If the PMD service interface... to:

135.5.3.6 Skew tolerance at SP5 for 100GBASE-P

If a 100GBASE-P PMD service interface...

Response Response Status U

REJECT.

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

The description in the referenced paragraph relates to currently specified as well as any future PMD which may have more than one lane. The specific requirements for each PMD are specified in the PMD clause.

[Editor's note: Comments r01-51, r01-53, r01-55, r01-56, and r01-58 from the same commenter relate to a similar topic.]

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

Cl 135 SC 135.5.3.7 P 179 L 30 # r01-58
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R skew variation <cc>

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

SuggestedRemedy

Change:

If there is a physically instantiated PMD service interface that allows the Skew to be measured, the Skew measured at SP5 is limited to no more than 145 ns of Skew and no more than 3.6 ns of Skew Variation. If there is no physically instantiated PMD service interface, the Skew measured at SP4 is limited to no more than 134 ns of Skew, and no more than 3.4 ns of Skew Variation. to:

If there is a physically instantiated PMD service interface that allows the Skew to be measured, the Skew measured at SP5 is limited to no more than 43 ns of Skew for 50GBASE-R or 145 ns of Skew for 100GBASE-P, and to no more than 3.6 ns of Skew Variation for 100GBASE-P. If there is no physically instantiated PMD service interface, the Skew measured at SP4 is limited to no more than 43 ns of Skew for 50GBASE-R or 134 ns of Skew for 100GBASE-P, and to no more than 3.4 ns of Skew Variation for 100GBASE-P.

Response Response Status U

REJECT.

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

The description in the referenced paragraph relates to currently specified as well as any future PMD which may have more than one lane. The specific requirements for each PMD are specified in the PMD clause.

[Editor's note: Comments r01-51, r01-53, r01-55, r01-56, and r01-58 from the same commenter relate to a similar topic.]

Cl 136 SC 136.6.1 P 200 L 16 # r01-59
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R skew <cc>

The Skew at SP3 (the output of the PMD), SP4 (the receiver MDI) and at SP5 (PMD service interface, output) have to be the same as at SP2 (PMD service interface, input of the PMD) for 50GBASE-CR, a serial PMD. As the receiver can't do anything about it, the "shall"s for SP4 and SP5 are not appropriate. What 802.3ba (all multilane) or 802.3bg (not a good precedent) did is not binding, nor a good choice for a family of serial PMDs. Any KR4-based 2-lane PMD can have its own independent Skew budget. Any future KP4-based 2-lane PMD can also have its own Skew budget, that could be like the 802.3bs one. What we write for a 1-lane PMA input cannot bind any 2-lane PMA. It's the SP6 spec that determines what future non-serial PMDs could be like, not SP3-5. D3.0 comment 123.

SuggestedRemedy

Change:

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns. Since the signal at the MDI represents a serial bit stream, there is no Skew Variation at this point. The Skew at SP4 (the receiver MDI) shall be less than 134 ns. Since the signal at the MDI represents a serial bit stream, there is no Skew Variation at this point. If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns. Since the signal at the PMD service interface represents a serial bit stream, there is no Skew Variation at this point. to: The Skew at SP3 (the transmitter MDI) shall also be less than 43 ns. Since the signal at the MDI represents a serial bit stream, there is no Skew Variation at this point. The Skew at SP4 (the receiver MDI) and SP5 (the output of the PMD at the PMD service interface) is the same as at SP2, and there is no Skew Variation at these points.

Correct Table 131-5, Summary of Skew constraints - as 50GBASE-R PMDs are serial it's simple to do. Change 54 134 145 to 43, 1434 3559 and 3852 to 1142. For SP2, remove the reference to 135.5.3.5. For SP5, remove the reference to 135.5.3.6. Also 137.6.1 138.3.2.1 139.3.2.

Response Response Status U

REJECT.

Comments on this same topic with a similar suggested remedy have been addressed at previous task force meetings. Examples include comments #147, #148, #220, #221 against Draft 1.2, comments #40, #41 against Draft 1.3, and comments i-123 and i-125 against Draft 3.0.

The common response to the Draft 1.2 comments was as follows: "REJECT.

Based on discussion and comment resolution at the January 2017 task force meeting WRT to the skew specifications for single-lane PMDs the consensus was to implement the specifications consistent with 40G, 100G, and 200G PHYs already specified in IEEE Std 802.3-2015 and P802.3bs.

See the final response for P802.3cd Draft 1.1 Comment #10"

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

The common response to the Draft 1.3 comments pointed back D1.2 comment #120 adding the note:

"There is no new information in the comment to support the suggested change."

The common response to the Draft 3.0 comments upheld and elaborated upon the previous responses as follows:

"REJECT."

The skew constraints for 100G in Table 80-5 and for 50G in Table 131-5 are consistent with the budget and methodology adopted by 802.3ba and 802.3bg and used in subsequent projects (e.g., 802.3bm, 802.3bs).

The skew constraints are established to ensure that the FEC/PCS skew tolerance is sufficient to support the worst case skew for any currently specified or potential (within reason) future PHY (e.g., 2-lane PMD for reach longer than 40 km). This is accomplished by having the same skew constraint at SP5 regardless of the PMD type.

The skew constraint at SP5 includes allocation for skew accumulated through the TX PMD (SP2 to SP3), the medium (SP3 to SP4), and the RX PMD (SP4 to SP5). Rather than specifying unique values for SP3, SP4, and SP5 based on PMD type, the adopted approach was to use the same numbers for all PMD types for consistency.

The approach described above is consistent for all PHY types defined by 802.3ba and subsequent projects. For instance, the medium skew accumulation (SP3 to SP4) of 80 ns was based on an 80 km multi-lane optical PMD. Nevertheless, the same value is used for other PMDs where the skew would be considerably lower (e.g., 100GBASE-SR4, 100GBASE-KR4, 100GBASE-CR4, etc.).

This specification methodology does not preclude an engineered implementation that optimizes the FEC/PCS skew buffering based on assumed lower PMD and medium skew accumulation. However, it should be noted that this implementation would not be compliant to 802.3cd."

In the suggested remedy for this comment, the commenter is proposing essentially the same changes as in these previously addressed comments and the commenter is providing no new evidence to support the proposed changes.

As noted in the response above, the task force has consistently exhibited consensus to retain the specification methodology for Skew and Skew Variation used for 40G, 100G, and 200G PHYs specified in the base standard. The specifications for Skew and Skew Variation in this draft are consistent with those in the base standard.

<i>Cl</i> 137	<i>SC</i> 137.9.2	<i>P</i> 251	<i>L</i> 29	# r01-64
Dawe, Piers J G		Mellanox Technologie		

<i>Comment Type</i>	TR	<i>Comment Status</i>	R	<i>SNDR <scope></i>
---------------------	-----------	-----------------------	----------	---------------------------

SNDR is measured in 33 GHz while the effect of SNR_TX is calculated (Annex 93A) in a different, lower bandwidth. This seems to lead to an error - probably because sigma_e and sigma_n are affected by bandwidth more strongly than pmax is. SNDR should be measured in something less than ~19 GHz.
D3.0 comment 138.

SuggestedRemedy

Add ", when sigma_e and sigma_n are found from signals observed with a fourth-order Bessel-Thomson low-pass response with 19.34 GHz 3 dB bandwidth.

NOTE--pmax is found from a signal observed with a fourth-order Bessel-Thomson low-pass response with 33 GHz 3 dB bandwidth."

<i>Response</i>	<i>Response Status</i>	U
-----------------	------------------------	----------

REJECT.

This is essentially a resubmit of comment i-138.

Comment i-138 was rejected with the following response:

"REJECT.

The sigma_TX term in COM is calculated under the assumption that the spectrum of the noise and the distortion is identical to the spectrum of the ideal signal at the transmitter output (sinc shaped per Eq. 93A-23). If that is the case, the signal, noise and distortion all go through the same transfer function, which includes the transmitter, receiver, and channel (Eq. 93A-19).

The actual effect on the receiver depends on the Tx noise and distortion spectrum (if high frequencies dominate, sigma_tx is too high because they will be more attenuated by channel and Rx than the signal; if low frequencies dominate, sigma_tx is too low since they will be less attenuated).

The suggested remedy includes a specific new filter for noise and distortion measurement but there is insufficient evidence that this filter is more suitable than the current filter."

There is no new information that would justify accepting this comment now.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

Cl 138 SC 138 P 263 L 1 # r01-69
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

This clause has received next to no attention - it's still the baseline, with some TDECQ changes inherited from other clauses. It needs more study. D3.0 comment 122.

SuggestedRemedy

Do the work. Show technical feasibility for the draft spec (after improvements).

The alternatives are:

withdraw the clause, which would be a pity; or
 delay the project until the work gets done.

Response Response Status U

REJECT.

No specific change to document suggested.

Measured data has been presented to the task force supporting the current specifications.

See:

http://www.ieee802.org/3/cd/public/Jan18/king_3cd_02_0118.pdf

http://www.ieee802.org/3/cd/public/adhoc/archive/chang_011018_3cd_02_adhoc-v2.pdf

Cl 138 SC 138.7.1 P 273 L 22 # r01-70
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

A TDECQ limit of 4.9 seems very high, given that the same fibres and transmitter, and receiver front-ends that should not be worse, can do 100GBASE-SR4 (PAM2, almost the same signalling rate) without the FFE. D.30 comment 119.

Also, it seems that the TDECQ spec limit can be "gamed" (D3.0 comment 116).

SuggestedRemedy

Compare a minimally compliant 100GBASE-SR4 transmitter and set the TDECQ limit accordingly. Provide a signal quality spec that cannot be "gamed".

Response Response Status U

REJECT.

No specific change to document suggested.

The issue that might be caused by a TDECQ limit of 4.9 dB has not been clarified. There is precedence for this kind of transmitter quality metric to be higher in MMF specifications than in SMF specifications.

To date no contribution has been made that demonstrates the problem, for example, a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation.

Measured data has been presented to the task force supporting the current specifications.

See:

http://www.ieee802.org/3/cd/public/Jan18/king_3cd_02_0118.pdf

http://www.ieee802.org/3/cd/public/adhoc/archive/chang_011018_3cd_02_adhoc-v2.pdf

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

CI 138 SC 138.8.5 P 276 L 33 # r01-71
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

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

D3.0 comment 116.

SuggestedRemedy

- To screen for noisy or distorted signals with heavy emphasis:
 Define $TDECQ_{rms} = 10 \cdot \log_{10}(A_{RMS}/(s^3 \cdot Qt \cdot R))$ where A_{RMS} is the standard deviation of the measured signal after the 13.28125 GHz or 11.2 GHz filter response (before the FFE), Qt and R are as already in Eq 212-12. s is the standard deviation of a fast clean signal with OMA=2 and without emphasis, observed through the filter response (0.6254 for 13.28125 GHz, 0.6006 for 11.2 GHz).
 Either, set limit for $TDECQ_{rms}$ according to what level of dirty-but-emphasised signal we decide is acceptable, add max $TDECQ_{rms}$ row to each transmitter table.
 Or, if the same relative limit is acceptable for all PAM4 optical PMDs, the limit could be here in the TDECQ procedure. e.g. make the $TDECQ_{rms}$ limit the same as the TDECQ limit, say here that both TDECQ and $TDECQ_{rms}$ must meet the TDECQ spec.
- To protect the receiver from having to "invert" heavily over-emphasised signals, set a minimum cursor weight, 0.9.
 Similarly in clauses 139, 140.
 To protect the equalizer from having to support unnecessary settings for waveforms that can't or shouldn't ever happen, constrain the cursor position - see other comments .

Response Response Status U

REJECT.

The need for additional transmitter specs has not been established, and insufficient evidence has been provided that the proposed remedy fixes the claimed problem.

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

A similar proposal to create a $TDECQ_{rms}$ spec was suggested in comment #r02-35 against 802.3cd D3.0, which was similarly rejected.

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

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

been provided.

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

CI 138 SC 138.8.5.1 P 276 L 38 # r01-73
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

Further investigation of possible minimally compliant MMF signals and their associated TDECQ FFE settings indicates that 2 pre, 2 post (making the cursor the third tap) is never significantly better than 1 pre, 3 post (making it the second tap), for compliant signals. Further refining the TDECQ search rules will avoid inefficiency both in product receiver design, testing and operation, and in TDECQ testing.

SuggestedRemedy

Change "Tap 1, tap 2, or tap 3, has" to "Tap 1 or tap 2 has". There is a separate comment for SMF because the different TDECQ limit there could lead to a different conclusion.

Response Response Status U

REJECT.

A similar proposal was made against draft 3.0 (comments i-107 i-117 and i120) which was reviewed by the Task Force.

The agreed resolution was to limit the main tap to tap 1, tap 2, or tap 3.
http://www.ieee802.org/3/cd/public/Mar18/dawe_3cd_01a_0318.pdf was reviewed by the Task Force.

There was no consensus to make the proposed change.

The resolution to i-117 was:

ACCEPT IN PRINCIPLE.

Implement the changes proposed in
http://www.ieee802.org/3/cd/public/Jan18/king_3cd_03_0118.pdf with editorial license

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

Cl 139 SC 139.7.5.4 P 301 L 1 # r01-76

Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

Further investigation of possible minimally compliant SMF signals and their associated TDECQ FFE settings indicates that 2 pre, 2 post (making the cursor the third tap) is never significantly better than 1 pre, 3 post (making it the second tap), for compliant signals. Further refining the TDECQ search rules will avoid inefficiency both in product receiver design, testing and operation, and in TDECQ testing.

SuggestedRemedy

Change "Tap 1, tap 2, or tap 3, has" to "Tap 1 or tap 2 has". Do the same in 140.7.5.1 because the TDECQ limit is similar. There is a separate comment for MMF because the different TDECQ limit there could lead to a different conclusion.

Response Response Status U

REJECT.

See response to comment r01-73.

[Editor's note added after comment resolution completed:

For reference, the response to comment r01-73 is copied here:

REJECT.

A similar proposal was made against draft 3.0 (comments i-107 i-117 and i120) which was reviewed by the Task Force.

The agreed resolution was to limit the main tap to tap 1, tap 2, or tap 3. http://www.ieee802.org/3/cd/public/Mar18/dawe_3cd_01a_0318.pdf was reviewed by the Task Force.

There was no consensus to make the proposed change.

The resolution to i-117 was:

ACCEPT IN PRINCIPLE.

Implement the changes proposed in http://www.ieee802.org/3/cd/public/Jan18/king_3cd_03_0118.pdf with editorial license

]

Cl 140 SC 140.3.2 P 315 L 46 # r01-77

Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

The Skew at SP3 (the output of the PMD), SP4 (the receiver MDI) and at SP5 (PMD service interface, output) have to be the same as at SP2 (PMD service interface, input of the PMD) for 100GBASE-DR, a serial PMD. As the receiver can't do anything about it, the "shall"s for SP4 and SP5 are not appropriate. What we write for a 1-lane PMD and PMA input doesn't affect the multi-lane PMA interfaces and PMDs: the point that is common to different PMDs is SP6, not SP3-5. D3.0 comment 125.

SuggestedRemedy

Change:

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns. Since the signal at the MDI represents a serial bit stream, there is no Skew Variation at this point. The Skew at SP4 (the receiver MDI) shall be less than 134 ns. Since the signal at the MDI represents a serial bit stream, there is no Skew Variation at this point. If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns. Since the signal at the PMD service interface represents a serial bit stream, there is no Skew Variation at this point. The Skew at SP3 (the transmitter MDI) shall also be less than 43 ns. Since the signal at the MDI represents a serial bit stream, there is no Skew Variation at this point. The Skew at SP4 (the receiver MDI) and SP5 (the output of the PMD at the PMD service interface) is the same as at SP2, and there is no Skew Variation at these points.

Correct Table 80-6, Summary of Skew constraints - add notes to the entries for SP3 SP4 SP4 saying that for 100GBASE-DR, the maximum Skew is as for SP2.

Response Response Status U

REJECT.

See response to comment r01-59.

[Editor's note added after comment resolution completed:

For reference, the response to comment r01-59 is copied here:

REJECT.

Comments on this same topic with a similar suggested remedy have been addressed at previous task force meetings. Examples include comments #147, #148, #220, #221 against Draft 1.2, comments #40, #41 against Draft 1.3, and comments i-123 and i-125 against Draft 3.0.

The common response to the Draft 1.2 comments was as follows:

"REJECT.

Based on discussion and comment resolution at the January 2017 task force meeting WRT to the skew specifications for single-lane PMDs the consensus was to implement the

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

specifications consistent with 40G, 100G, and 200G PHYs already specified in IEEE Std 802.3-2015 and P802.3bs.

See the final response for P802.3cd Draft 1.1 Comment #10"

The common response to the Draft 1.3 comments pointed back D1.2 comment #120 adding the note:

"There is no new information in the comment to support the suggested change."

The common response to the Draft 3.0 comments upheld and elaborated upon the previous responses as follows:

"REJECT.

The skew constraints for 100G in Table 80-5 and for 50G in Table 131-5 are consistent with the budget and methodology adopted by 802.3ba and 802.3bg and used in subsequent projects (e.g., 802.3bm, 802.3bs).

The skew constraints are established to ensure that the FEC/PCS skew tolerance is sufficient to support the worst case skew for any currently specified or potential (within reason) future PHY (e.g., 2-lane PMD for reach longer than 40 km). This is accomplished by having the same skew constraint at SP5 regardless of the PMD type.

The skew constraint at SP5 includes allocation for skew accumulated through the TX PMD (SP2 to SP3), the medium (SP3 to SP4), and the RX PMD (SP4 to SP5). Rather than specifying unique values for SP3, SP4, and SP5 based on PMD type, the adopted approach was to use the same numbers for all PMD types for consistency.

The approach described above is consistent for all PHY types defined by 802.3ba and subsequent projects. For instance, the medium skew accumulation (SP3 to SP4) of 80 ns was based on an 80 km multi-lane optical PMD. Nevertheless, the same value is used for other PMDs where the skew would be considerably lower (e.g., 100GBASE-SR4, 100GBASE-KR4, 100GBASE-CR4, etc.).

This specification methodology does not preclude an engineered implementation that optimizes the FEC/PCS skew buffering based on assumed lower PMD and medium skew accumulation. However, it should be noted that this implementation would not be compliant to 802.3cd."

In the suggested remedy for this comment, the commenter is proposing essentially the same changes as in these previously addressed comments and the commenter is providing no new evidence to support the proposed changes.

As noted in the response above, the task force has consistently exhibited consensus to retain the specification methodology for Skew and Skew Variation used for 40G, 100G, and 200G PHYs specified in the base standard. The specifications for Skew and Skew Variation in this draft are consistent with those in the base standard.

]

<i>Cl</i> 139	<i>SC</i> 139.7.10.2	<i>P</i> 299	<i>L</i> 54	# r01-100
Liu, Hai-Feng		Intel Corporation		

<i>Comment Type</i>	TR	<i>Comment Status</i>	R
---------------------	-----------	-----------------------	----------

[note that a comment is needed in this section in addition to the comment above to avoid any confusion with the less clear instructions in the referenced 802.3bs section 121.8.9.2] PAM4 link analysis has shown (see schube_3cd_02_0118) that the composition and ratio of the stressors in the stressed receiver sensitivity test can have a strong impact on link performance. In particular, the same SECQ can generate widely varying BER performance from the same receiver depending on the amount of ISI/bandwidth limitation as a portion of the overall SECQ stress. To address this we propose to clarify the current language describing the stressor ratio to be used to create the stressed Rx sensitivity conformance test input, to avoid understressing the receiver and causing interoperability issues.

Suggested Remedy

Add the following sentence to the end of section 139.7.10.2: "Note that regardless of calibration method, and regardless of the characteristics of the reference/test transmitter before stressors are added, at least half of the total dB value of stressed eye closure (SECQ) should be from bandwidth limitations / ISI, as outlined in section 139.7.9.1 above."

<i>Response</i>	<i>Response Status</i>	U
-----------------	------------------------	----------

REJECT.

Subclause 139.7.10.2 does not exist. This should be 139.7.9.2 starting on page 303 of the draft.

The requirement that "The combination of the low-pass filter and the E/O converter should have a frequency response that results in at least half of the dB value of the stressed eye closure (SECQ) specified in Table 139-7 for 50GBASE-FR and 50GBASE-LR before the sinusoidal and Gaussian noise terms are added, according to the methods specified in 139.7.9.2." is already present in 139.7.9.1, so it is not necessary to repeat the requirement in 139.7.9.2.

[Editor's note: Comment r01-19 deals with a related topic]

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 1st Sponsor recirculation ballot comments

Cl 140 SC 140.7.10 P 320 L 15 # r01-101
Liu, Hai-Feng Intel Corporation

Comment Type TR Comment Status R

PAM4 link analysis has shown (see schube_3cd_02_0118) that the composition and ratio of the stressors in the stressed receiver sensitivity test can have a strong impact on link performance. In particular, the same SECQ can generate widely varying BER performance from the same receiver depending on the amount of ISI/bandwidth limitation as a portion of the overall SECQ stress. To address this we propose to clarify the current language describing the stressor ratio to be used to create the stressed Rx sensitivity conformance test input, to avoid understressing the receiver and causing interoperability issues.

SuggestedRemedy

Add the following sentence to the end of section 140.7.10: "Note that regardless of calibration method, and regardless of the characteristics of the reference/test transmitter before stressors are added, at least half of the total dB value of stressed eye closure (SECQ) should be from bandwidth limitations / ISI."

Response Response Status U

REJECT.

See response to comment r01-100.

[Editor's note added after comment resolution completed:

For reference, the response to comment r01-100 is copied here:

REJECT.

Subclause 139.7.10.2 does not exist. This should be 139.7.9.2 starting on page 303 of the draft.

The requirement that "The combination of the low-pass filter and the E/O converter should have a frequency response that results in at least half of the dB value of the stressed eye closure (SECQ) specified in Table 139-7 for 50GBASE-FR and 50GBASE-LR before the sinusoidal and Gaussian noise terms are added, according to the methods specified in 139.7.9.2." is already present in 139.7.9.1, so it is not necessary to repeat the requirement in 139.7.9.2.

[Editor's note: Comment r01-19 deals with a related topic]

]