Cl **136** *SC* **136.9.3** *P* **225** *L* **23** # i-21

RAN, ADEE Intel Corporation

Comment Type **TR** *Comment Status* **R** *AC-coupling*

Scope connection through AC coupling is not specified in this clause. Transmitter tests should be done through AC coupling (except for common mode tests).

See http://www.ieee802.org/3/cd/public/adhoc/archive/ran_112717_3cd_adhoc.pdf

SuggestedRemedy

In the first paragraph:

"Unless specified otherwise, all transmitter measurements are made for each lane separately, at TP2, utilizing the test fixtures specified in Annex 136B, using a test system with a fourth-order Bessel-Thomson low-pass response with 33 GHz 3 dB bandwidth"

Append: "connected as shown in Figure 92-15".

Response *Response Status* **U**

REJECT.

[Editor changed CommentType from GR to TR]

A similar issue is being addressed in the comment resolution in the 802.3cj revision project. The task force prefers to close this issue based on the resolution of comments in 802.3cj.

Cl **138** *SC* **138.8.8** *P* **275** *L* **16** # i-58

RAN, ADEE Intel Corporation

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

The SRS methodology in 121.8.9.1 and 121.8.9.3 has several flaws that need to be addressed:

- Half of the SECQ should be obtained without noise or jitter, using the combination of low-pass filter and E/O converter (which is marked as "Tunable" in Figure 139-5, and also in Figure 122-5, but not in Figure 121-6). Different E/O converters that may be used in the test setup may have different characteristics (noise and BW), which will result in very different setting for the low-pass filter. This freedom enables very different test conditions, some of which may be favorable for some devices.

- The remaining SECQ is met by adjusting the Gaussian noise (with unspecified power), sinusoidal interferer amplitude (with unspecified amplitude and frequency), and low-pass filter (with no specified limits); the sinusoidal jitter stress (which is specified) also affects SECQ. There are too many degrees of freedom here, which again enable very different test conditions (as demonstrated in http://www.ieee802.org/3/cd/public/Nov17/chang_3cd_01_1117.pdf).

- The effect of sinusoidal jitter on SECQ measurement is difficult to predict, since the measurement is done with a CRU (which tracks all frequencies to some extent). Also, the pattern used for calibration is very short and the length captured is not specified (e.g. no requirement to measure at least a full cycle of the sinusoidal jitter, which may be much longer than the test pattern). This may result in repeatability problems.

The too many degrees of freedom need to be limited, ideally to one knob that has to be turned to reach the required SECQ. This is the motivation for the proposed change.

Also applies to 139.7.9 and 140.7.9.

SuggestedRemedy

Add exceptions or additions to the methods of 121.8.9.1 and 121.8.9.3 including the following:

1. Specify the combined bandwidth of the E/O and the low-pass filter (without equalization), e.g. -3 dB at 15 GHz (or an agreed upon value). This may be measured using a different transmitter (e.g. sinusoidal generator). This step is prior to any SECQ measurement

2. Specify that the target SECQ is achieved by addition of Gaussian noise only (without sinusoidal interference), this will be the knob to turn to achieve the SECQ.

3. Specify that SECQ is calibrated once before addition of sinusoidal jitter, and calibration is not repeated for every jitter frequency. (If necessary, reduce SECQ target to accommodate for expected jitter effect).

Implement the chosen solution (with different bandwidth and SECQ targets) also in 139.7.9 and 140.7.9.

Response *Response Status* **U**

REJECT.

It has not been demonstrated that there is a problem with the draft, nor has it been demonstrated that the proposed remedy fixes it.

The work presented in
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.

[Editor's note: Comments i-82, i-83 and i-84 address a similar issue.]

Cl **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.

Cl **135G** *SC* **135G.3.1** *P* **375** *L* **21** # i-61

RAN, ADEE Intel Corporation

Comment Type **TR** *Comment Status* **R** *jitter mismatch <cc>*

100GAUI-2 C2M host output is specified by reference to 120E.3.1. This means jitter is measured with a CRU with corner frequency of 4 MHz (per 120E.4.2).

Low-frequency jitter will be attenuated by the CRU - that means it is assumed to be tracked by the module's CDR.

This creates a problem if the module is a 100GBASE-DR PMD; the tracked jitter will be forwarded to the optical transmitter with the same time values, so doubled magnitude in UI terms.

This means that the link partner's optical receiver, with assumed CDR BW of 4 MHz too (per 140.7.9 and 121.8.9.4 SRS definitions), will see low frequency jitter that can be twice of what it is tested to tolerate.

The CDRs used in practice are second-order, so at very low frequencies this higher jitter level will likely be acceptable; but there is no specification for the integral gain of the CDR, so at medium frequencies the jitter tolerance is implementation dependent (even for fully compliant PMDs).

Having excessive untracked low-frequency jitter may be detrimental for BER even with FEC; the SNR will vary over time, and even if the average is good, uncorrectable codewords may be more frequent than what could be expected. This can cause unexpected deployment problems.

This issue was not resolved in 802.3bs although there have been comments about having the same CDR bandwidth for 50 and 100 Gb/s per lane interfaces. The least painful way to solve it at this point seems to be a recommendation for the host output jitter. This will leave all optical specs unmodified.

SuggestedRemedy

Add the following text after the single paragraph in 135G.3.1:

To limit the jitter at frequencies which a 100GBASE-DR PMD's optical receiver may not track well, it is recommended that in addition to the specifications in 120E.3.1, the Host output eye width and eye height specifications (120E.3.1.6) be met when measured using a clock recovery unit with a corner frequency of 2 MHz.

Response *Response Status* **U**

REJECT.

Reviewed http://www.ieee802.org/3/cd/public/Jan18/ghiasi_3cd_01_0118.pdf.

Straw poll #1 indicated lack of consensus to make any technical changes to the jitter specification.

Straw poll #1:
I would support making a technical change to the jitter specification.
Y: 4
N: 21

There is no support to make any changes to the jitter specifications.

Cl **138** *SC* **138.8.5** *P* **274** *L* **31** # i-79

Liu, Hai-Feng Intel Corporation

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

The sub-eye threshold levels in current TDECQ measurement are determined by the OMAouter and the average optical power of the PAM4 eye diagram (Pave) as defined in equations (121-1), (121-2) and (121-3). While this is good for perfectly linear PAM4 signals with 3 equal eye amplitudes, it would lead to pessimistic TDECQ values as compared to the link sensitivity penalty measurements where thresholds are adjusted by real receivers to achieve the lowest BER even if the signal is not perfectly linear.
Several vendors have contributed data (way_3bs_01a_0717, tamura_3bs_01a_0917, baveja_3cd_01_1117) showing many units that are able to close the link with good sensitivity/BER margin would fail to meet the maximum TDECQ specification, causing good transmitters to be failed.

SuggestedRemedy

Propose to adopt threshold optimization in TDECQ measurement as described in mazzini_120617_3cd_adhoc-v2 with the additional constraints on the allowable adjustment range.
Detailed presentation to be submitted for the January meeting with the summary of the proposal, measurement data to support the proposal, and suggested changes in details.

Response *Response Status* **U**

REJECT.

The presentation http://www.ieee802.org/3/cd/public/Jan18/liu_3cd_01a_0118.pdf was reviewed.
It does not provide sufficient details to implement.
It is not clear that the suggested remedy would be an improvement to the draft.
Also http://www.ieee802.org/3/cd/public/Jan18/king_3cd_01_0118.pdf was presented in support of the adequacy of the current specification.

Cl **139** *SC* **139.7.5** *P* **296** *L* **20** # i-80

Liu, Hai-Feng Intel Corporation

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

The sub-eye threshold levels in current TDECQ measurement are determined by the OMAouter and the average optical power of the PAM4 eye diagram (Pave) as defined in equations (121-1), (121-2) and (121-3). While this is good for perfectly linear PAM4 signals with 3 equal eye amplitudes, it would lead to pessimistic TDECQ values as compared to the link sensitivity penalty measurements where thresholds are adjusted by real receivers to achieve the lowest BER even if the signal is not perfectly linear.
Several vendors have contributed data (way_3bs_01a_0717, tamura_3bs_01a_0917, baveja_3cd_01_1117) showing many units that are able to close the link with good sensitivity/BER margin would fail to meet the maximum TDECQ specification, causing good transmitters to be failed.

SuggestedRemedy

Propose to adopt threshold optimization in TDECQ measurement as described in mazzini_120617_3cd_adhoc-v2 with the additional constraints on the allowable adjustment range.
Detailed presentation to be submitted for the January meeting with the summary of the proposal, measurement data to support the proposal, and suggested changes in details.

Response *Response Status* **U**

REJECT.

See resolution to comment i-79

[Editor's note added after comment resolution completed:

For reference, the response to comment i-79 is copied here:

REJECT.

The presentation http://www.ieee802.org/3/cd/public/Jan18/liu_3cd_01a_0118.pdf was reviewed.
It does not provide sufficient details to implement.
It is not clear that the suggested remedy would be an improvement to the draft.
Also http://www.ieee802.org/3/cd/public/Jan18/king_3cd_01_0118.pdf was presented in support of the adequacy of the current specification.

]

Cl **140** *SC* **140.7.5** *P* **319** *L* **19** # i-81

Liu, Hai-Feng Intel Corporation

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

The sub-eye threshold levels in current TDECQ measurement are determined by the OMAouter and the average optical power of the PAM4 eye diagram (Pave) as defined in equations (121-1), (121-2) and (121-3). While this is good for perfectly linear PAM4 signals with 3 equal eye amplitudes, it would lead to pessimistic TDECQ values as compared to the link sensitivity penalty measurements where thresholds are adjusted by real receivers to achieve the lowest BER even if the signal is not perfectly linear.
Several vendors have contributed data (way_3bs_01a_0717, tamura_3bs_01a_0917, baveja_3cd_01_1117) showing many units that are able to close the link with good sensitivity/BER margin would fail to meet the maximum TDECQ specification, causing good transmitters to be failed.

SuggestedRemedy

Propose to adopt threshold optimization in TDECQ measurement as described in mazzini_120617_3cd_adhoc-v2 with the additional constraints on the allowable adjustment range.
Detailed presentation to be submitted for the January meeting with the summary of the proposal, measurement data to support the proposal, and suggested changes in details.

Response *Response Status* **U**

REJECT.

See resolution to comment i-79

[Editor's note added after comment resolution completed:

For reference, the response to comment i-79 is copied here:

REJECT.

The presentation http://www.ieee802.org/3/cd/public/Jan18/liu_3cd_01a_0118.pdf was reviewed.
It does not provide sufficient details to implement.
It is not clear that the suggested remedy would be an improvement to the draft.
Also http://www.ieee802.org/3/cd/public/Jan18/king_3cd_01_0118.pdf was presented in support of the adequacy of the current specification.

]

Cl **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.]

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 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 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* **140.1** *P* **309** *L* **14** # i-85

Maki, Jeffery Juniper Networks, Inc.

Comment Type **TR** *Comment Status* **R** *jitter mismatch <cc>*

Table 140-1 lists a variety of AUI options (e.g., CAUI-4 C2M, 100GAUI-4 C2M, 100GAUI-2 C2M) to build a PHY using a 100GBASE-DR PMD with no explicit regard to the potential mismatch of the output jitter of the AUI and the compliant output jitter of the 100GBASE-DR PMD.

SuggestedRemedy

Add text stating, "The PMA between the AUI and the PMD is responsible for adapting the output jitter of the chosen AUI option to meet the compliant output jitter of the 100GBASE-DR PMD."

Response *Response Status* **U**

REJECT.

There is no consensus to make the proposed change.

[Editor's note added after comment resolution was complete.

The lack of consensus was based on opinions that the suggested new text was unecessary and the draft was sufficient as written.

]

Cl **135G** *SC* **135G.3.1** *P* **375** *L* **33** # i-87

Wertheim, Oded Mellanox Technologie

Comment Type **TR** *Comment Status* **R** *jitter mismatch <cc>*

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

SuggestedRemedy

Scale the corner frequency for 100GAUI-2 to 2MHz (half the corner frequency of 100GBASE-DR). The proposed resolution doesn't introduce constraints on future 100G per lane interfaces and provides simpler solution than alternative solutions that were investigated, with no change to the optical specs.

1. Add an exception to 135G.4 50GAUI-1 C2M and 100GAUI-2 C2M measurement methodology with an exception that:
a. The reference CRU for the Eye width and eye height measurement method has a corner frequency of 2MHz for the host output and module input tests.

2. Add an exception to 135G.3.4 50GAUI-1 C2M and 100GAUI-2 C2M module input characteristics:
With an exception that:
a. The reference CRU for the Module stressed input test has a corner frequency of 2MHz
b. The applied sinusoidal jitter values for 100GAUI-2 Module stressed input test shall be:
{Jitter frequency, Jitter amplitude}
Case A: {0.02, 5}
Case B: {0.66, 0.15}
Case C: {2, 0.05}
Case D: {6, 0.05}
Case E: {20, 0.05}

Response *Response Status* **U**

REJECT.

Resolve using the response to comment i-61.

[Editor's note added after comment resolution completed:

For reference, the response to comment i-61 is copied here:

REJECT.

Reviewed http://www.ieee802.org/3/cd/public/Jan18/ghiasi_3cd_01_0118.pdf.

Straw poll #1 indicated lack of consensus to make any technical changes to the jitter specification.

Straw poll #1:
I would support making a technical change to the jitter specification.
Y: 4
N: 21

There is no support to make any changes to the jitter specifications.
]

Cl **136** *SC* **136.9** *P* **225** *L* **39** # i-96

Rysin, Alexander Mellanox Technologie

Comment Type **TR** *Comment Status* **A** *ERL*

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

* Add annex describing ERL measurement and computation. See prior presentations for description.
* Remove the requirement for Differential return loss in Table 136-11.
* Add a requirement for Effective Return Loss (ERL) to be greater than 18.2 dB in Table 136-11.
* In 136.9.4 change "The receiver shall meet the return loss requirements specified in 92.8.4.2 and 92.8.4.3." to "The receiver shall meet the effective return loss requirement in 136.9.3."
* Add a paragraph in 137.9.2 and to 137.9.3 - "Effective Return Loss (ERL, min) is 16.2 dB. There is no frequency domain return loss mask."

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

Implement the changes according to http://www.ieee802.org/3/cd/public/Jan18/ran_3cd_01_0118.pdf.

Cl **136** *SC* **136.9** *P* **226** *L* **8** # i-97

Rysin, Alexander Mellanox Technologie

Comment Type **TR** *Comment Status* **A** *ERL*

Transmitter output residual ISI SNR_ISI (min) 36.8 dB (Clause 136) and 43 dB (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 136 and Clause 137 are even more stringent than in 120D. 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.

SuggestedRemedy

* Remove reference to SNR_ISI in Table 136-11 --Summary of transmitter specifications at TP2.
* Add a requirement for Effective Return Loss (ERL) to be greater than 18.2 dB in Table 136-11.
* 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 "Effective Return Loss (ERL) is calculated with Nb set to 12 (see Annex New). ERL shall be at least 16.2 dB. The Transmitter Output residual ISI SNR_ISI specification in Table in Table 120D-1 does not apply."

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

Implement the changes according to
http://www.ieee802.org/3/cd/public/Jan18/ran_3cd_01_0118.pdf.

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 **135G** *SC* **135G.3.1** *P* **375** *L* **21** # i-115

Dawe, Piers J G Mellanox Technologie

Comment Type **TR** *Comment Status* **R** *jitter mismatch <cc>*

As pointed out in both 802.3bs and this project, a host output with 50 Gb/s lanes is allowed to make twice as much low frequency jitter at very low frequencies as a receiver with 100 Gb/s lane(s) is required to receive. A jitter buffer does not fix this unless it is infinite. To assure interoperability, there must be industry-wide agreement that tightens 50G/lane host low frequency jitter generation, increases 100G/lane receiver low frequency jitter tolerance, or a combination. The proposed remedy is as simple as any of the options considered. Also it is likely to be compatible with 100G electrical lanes. This remedy must be applied to 100GAUI-2 C2M host outputs (unless another remedy is chosen), but may be applied to 50GAUI-1 host outputs and/or the corresponding module inputs for consistency. As any 50G/lane E/O conversions basically pass the low frequency jitter along for something else to tolerate, we can leave their specs alone.

SuggestedRemedy

Add to the end of the sentence "with the exception that the clock recovery unit's corner frequency (see 120E.4.2) is 2 MHz not 4 MHz".

If desired, change 135G.3.4: add "with the exceptions that the sinusoidal jitter (see 120E.3.4.1.1 and Table 120E-8) is defined by Table 135G-New, and that the reference CRU's corner frequency (see 120E.3.4.1.1of 4 MHz) is 2 MHz not 4 MHz".
Table 135G-New--Applied sinusoidal jitter
Parameter Case A Case B Case C Case D Case E Case F Units
Jitter frequency 0.02 0.667 2 6 20 60 MHz
Jitter amplitude 5 0.15 0.05 0.05 0.05 0.05 UI

Response *Response Status* **U**

REJECT.

Resolve using the response to comment i-61.

[Editor's note added after comment resolution completed:

For reference, the response to comment i-61 is copied here:

REJECT.

Reviewed http://www.ieee802.org/3/cd/public/Jan18/ghiasi_3cd_01_0118.pdf.

Straw poll #1 indicated lack of consensus to make any technical changes to the jitter specification.

Straw poll #1:
I would support making a technical change to the jitter specification.
Y: 4
N: 21

There is no support to make any changes to the jitter specifications.
]

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*log10(A_RMS/(s*3*Qt*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 additonal 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.

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.

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* **23** # i-136

Dawe, Piers J G Mellanox Technologie

Comment Type **TR** *Comment Status* **A** *ERL*

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

SuggestedRemedy

Either: Insert a new first item in the list of exceptions to Table 120D-1, create a new equation for Tx RL that is similar to the Cl.93 and the channel RL at low frequencies; 12 - 0.625f, 8.7-0.075f. Add figure to illustrate.
Or: change to an ERL spec or similar for the transmitter. Same Nb set to 12.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

Implement the changes according to
http://www.ieee802.org/3/cd/public/Jan18/ran_3cd_01_0118.pdf.

Cl **137** *SC* **137.9.2** *P* **251** *L* **28** # i-137

Dawe, Piers J G Mellanox Technologie

Comment Type **TR** *Comment Status* **A** *ERL*

Transmitter output residual ISI, SNR_ISI (min) 36.8 dB (Clause 136) and 43 dB (Clause 137) is still too high - can barely measure the IC through the test fixture. The warning NOTE in 120D.3.1.7 (where it's "only" 34.8 dB) shows the issue, but doesn't solve it. D2.0 comment 140, D21. comment 49.

SuggestedRemedy

Change to ERL spec or similar for the transmitter. Same Nb set to 12. Delete the SNR_ISI spec.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

Implement the changes according to
http://www.ieee802.org/3/cd/public/Jan18/ran_3cd_01_0118.pdf.

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.

Cl **137** *SC* **137.9.3** *P* **251** *L* **35** # i-141

Dawe, Piers J G Mellanox Technologie

Comment Type **TR** *Comment Status* **A** *ERL*

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

SuggestedRemedy

Either: Insert a new first item in the list of exceptions to Table 120D-5, create a new equation for Rx RL that is similar to the Cl.93 and the channel RL at low frequencies; 12 - 0.625f, 8.7-0.075f. Add figure to illustrate or pont to the figure for Tx RL (see another comment).
Or: change to an ERL spec or similar for the receiver. I think it can be more lenient than the transmitter spec because we have the receiver interference tolerance test.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

Implement the changes according to
http://www.ieee802.org/3/cd/public/Jan18/ran_3cd_01_0118.pdf.