Cl **116** *SC* **116.5** *P* **119** *L* **8** # i-104

Dawe, Piers J G — Mellanox Technologie

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

Table 116-7 has 80 ns for optical skew, and 100 ns for electrical (PCB), PMD and PMA skew. This is the same in ns as 802.3ba, but a total of 76,500 bits instead of 18,562.5, or 4.12 times as many bits to buffer. While this may not be as expensive as just a few bits in an optical module, some of this is an avoidable cost. The Skew limits need updating according to the principles used there (see http://ieee802.org/3/ba/public/may08/anslow_01_0508.pdf). The unit interval here is 38 (or 19) ps not 97 ps, and the number of lanes is 4 not 10.

SuggestedRemedy

Change SP1 from 29 ns, ~770 UI to 16 ns, ~425 UI.
Change SP2 from 43 ns, ~1142 UI to 24 ns, ~628 UI.
Change SP3 from 54 ns, ~1434 UI to 35 ns, ~930 UI.
Change SP4 from 134 ns, ~3559 UI to 115 ns, ~3055 UI.
Change SP5 from 145 ns, ~3852 UI to 126 ns, ~3347 UI.
Change SP6 from 160 ns, ~4250 UI to 134 ns, ~3559 UI.
Change "At PCS receive" from 180 ns, ~4781 UI to 145 ns, ~3852 UI.
Make the equivalent changes in the following clauses.

Response *Response Status* **U**

REJECT.
The initial Skew values were introduced into the P802.3ba draft by comment 240 against D1.0 with reference to:
http://www.ieee802.org/3/ba/public/nov08/giannakopoulos_01_1108.pdf
For example, the Skew at SP1 of 29 ns was justified by an analysis of an FPGA solution in:
http://www.ieee802.org/3/ba/public/may08/giannakopoulos_01_0508.pdf
The commenter has not provided equivalent analysis that shows that only 16 ns is required for 200 Gb/s and 400 Gb/s Ethernet.

Cl **116** *SC* **116.5** *P* **119** *L* **29** # i-105

Dawe, Piers J G — Mellanox Technologie

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

The Skew Variation limits need updating according to the principles in http://ieee802.org/3/ba/public/may08/anslow_01_0508.pdf as explained in http://ieee802.org/3/cd/public/Jan17/wertheim_3cd_01_0117.pdf The unit interval here is 38 (or 19) ps not 97 ps. The 8/4-lane module PMA is a completely different design to a host SerDes, and naturally, Tx and Rx sides are different designs. These relatively small FIFOs (just a few UI) are very expensive per UI in e.g. power, and consume some power even if never used.

SuggestedRemedy

Change SP1 from 0.2 ns, ~5 UI, N/A to 0.11 ns, ~3 UI, N/A.
Change SP2 from 0.4 ns, ~11 UI, N/A to 0.22 ns, ~6 UI, NA.
Change SP3 from 0.6 ns, ~16 UI, ~32 UI to 0.42 ns, ~11 UI, ~22 UI.
Change SP4 from 3.4 ns, ~90 UI, ~181 UI to 3.22 ns, ~86 UI, ~171 UI.
Change SP5 from 3.6 ns, ~96 UI, N/A to 3.42 ns, ~91 UI, N/A.
Change SP6 from 3.8 ns, ~101 UI, N/A to 3.53 ns, ~94 UI, N/A.
Change "At PCS receive" from 4 ns, ~106 UI, N/A to 3.73 ns, ~99 UI, N/A.
Make the equivalent changes in the following clauses.
It doesn't matter much if the SP4,5,6 and "At PCS receive" limits are changed or not.

Response *Response Status* **U**

REJECT.
The issue of whether to tighten the Skew Variation limits for PHYs using 25G lanes as proposed in
http://ieee802.org/3/cd/public/Jan17/wertheim_3cd_01_0117.pdf
was discussed in the P802.3cd Task Force in connection with comments #80 and #74 against P802.3cd D1.1 with the result that the same numbers as in the P802.3bs draft were adopted for 50 Gb/s Ethernet. See:
http://www.ieee802.org/3/cd/comments/8023cd_D11_final_comment_responses_by_clause.pdf

Cl **120** *SC* **120.5.1** *P* **190** *L* **20** # i-17

RAN, ADEE Intel

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

As noted in 120.5.11.2.4, a square wave may not be received correctly by the CDR of the PMA at the receive side of the 200GAUI-4 or 400GAUI-8 (whether or not it is adjacent to the PMD).

There is nothing in this clause that states that the PMA _receiver_ expects a CDR-friendly pattern and may not work well with a square wave (or, for that matter, with SSPR).

The PMA receiver behavior should only be specified for PCS data and for PRBS31/PRBS31Q. SSPR and square wave are used for transmitter testing, and we should not expect CDRs to operate with the same performance as with valid data. But as the text stands there is no special treatment for these patterns - the BER requirements in all AUI annexes are pattern-agnostic. This is an overkill.

This subclause seems to be the right place to state that the PMA receiver is not expected to cope with this kind of patterns.

SuggestedRemedy

Add a new paragraph at the end of 120.5.1:

"Clock and data recovery specifications apply for receiving PCS encoded data or PRBS31/PRBS31Q test patterns. Feeding other patterns (such as square wave or SSPR/SSPRQ) into a PMA through a physically instantiated interface may yield unexpected results".

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

As SSPRQ is used for optical tests, in principle, it could be generated from the last PMA on a host board when only a PMA providing an NRZ or PAM4 retimer function exists in the module. But for square wave, the concern is valid.

Add a paragraph at the end of 120.5.1:
"Test patterns that are intended for transmitter testing, such as a square wave, may not be correctly recovered by an adjacent PMA."

Cl **120** *SC* **120.5.11.2.3** *P* **201** *L* **5** # i-109

Dawe, Piers J G Mellanox Technologie

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

This SSPRQ is not suitable for use in TDECQ or stressed receiver calibration because measurements with this pattern do not give the correct penalty.

SuggestedRemedy

Change the first seed in Table 120-2 to one for which a minimally compliant transmitter with 0.4 dB baseline wander penalty (before and after FEC) with a random payload measures as minimally compliant (i.e. also 0.4 dB penalty) with SSPRQ.
It may be necessary to adjust another seed to get appropriate transition density characteristics.

Response *Response Status* **U**

REJECT.

The current SSPRQ pattern was adopted for use in the TDECQ test (after presentation of its baseline wander characteristics) by comment 50 against D1.3. A straw poll was taken in association with that comment: Do you support adopting the SSPRQ pattern for TDECQ and SRS calibration in Clauses 122 and 123? Yes 41 No 2 .

Comments i-130, i-133, and i-145 proposed to change the first seed in Table 120-2 but these comments were not accepted.

Cl **120D** *SC* **120D.3.1.7** *P* **356** *L* **23** # i-158

Hidaka, Yasuo — Fujitsu Laboratories of

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

Optimization of two parameters of the second-order CTLE as described in 93A.1.4.3 with parameters in Table 120D-8 is not required for the loss of package and test fixture.
The CTLE defined for chip-to-module interface in 120E.3.1.7 should be sufficient.

This is re-submission of comment #33 for D2.2.

SuggestedRemedy

Change
"SNR_ISI is defined by Equation (120D-8) computed from p_max and ISI_cursors after these have been re-calculated with the continuous time filter described in 93A.1.4.3 using the parameters in Table 120D-7 applied and optimized for maximum SNR_ISI."
to
"SNR_ISI is defined by Equation (120D-8) computed from p_max and ISI_cursors after these have been re-calculated with the selectable continuous time linear equalizer (CTLE) which is described in 120E.3.1.7 by Equation (120E-2) with coefficients in Table 120E-2 and illustrated in Figure 120E-9 applied and optimized for maximum SNR_ISI."

Response *Response Status* **U**

REJECT.
No consensus for a change at this time.

[Editor's note added after comment resolution completed. The consensus view was that the current measurement method is adequate and there is no need to simplfy it.]

Cl **120D** *SC* **120D.4** *P* **360** *L* **4** # i-73

Dudek, Michael — Cavium

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

Simulations presented in the 802.3cd task force have shown that the value of COM for 20dB channels varies significantly based on the values of Zc and Rd and that the presently used values do not provide the worst case result. No single set of values is the worst case for all channels. Some channels are showing 0.5dB less COM than the worst case package for that channel. (See http://grouper.ieee.org/groups/802/3/cd/public/adhoc/archive/hidaka_020117_3cd_adhoc.pdf and further as yet unpublished work)

SuggestedRemedy

Change the COM specification for the channel to 3.5dB here while leaving the COM calibration target for the receiver interference tolerance test at 3.0dB.

Response *Response Status* **U**

REJECT.
There was no consensus to make the equivalent change in P802.3cd

Straw Poll
Change the COM specification for the channel to 3.5dB 4
Make no change 9

Cl **120E** *SC* **120E** *P* **365** *L* **1** # i-118

Dawe, Piers J G — Mellanox Technologie

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

Are there discrepancies between CEI-56G-VSR-PAM4 and Annex 120E for which Annex 120E should change?

SuggestedRemedy

?

Response *Response Status* **U**

REJECT.
The comment identifies no issues, and proposes no remedies.

Cl **120E** *SC* **120E.3.1** *P* **369** *L* **19** # i-119

Dawe, Piers J G — Mellanox Technologie

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

The host is allowed to output a signal with large peak-to-peak amplitude but very small EH - in other words, a very bad signal. If the module is exactly like the reference receiver, that would work - but that's not a reasonable "if".

SuggestedRemedy

We may need some other spec to protect the module from unexpected signals.

Response *Response Status* **U**

REJECT.
No remedy provided. The commenter is encouraged to provide a presenation on this subject.

Cl **120E** *SC* **120E.3.1.6** *P* **370** *L* **41** # i-120

Dawe, Piers J G — Mellanox Technologie

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

There is no need for 31 UI offset between lanes. For PRBS13Q, only 1 UI offset is enough to give excellent decorrelation, better than 100-200 UI offset, and there is a spur at about 450 UI. PRBS31Q is believed to behave similarly (but it's such a long pattern I haven't checked). In some test setups, there is a master PRBS generator and an arrangement of splitters and cables; the cables must be kept short for good performance. 31 UI x 7 steps at 26.5625 GBd and 5 ns/m is 1.63 m - too long.

SuggestedRemedy

As the paths between the test points and the host PMA front-end circuitry are not likely to differ by more than 50 mm or about 10 UI, change 31 to 12. Also in 120E.3.3.2.1 Host stressed input test procedure.

Response *Response Status* **U**

REJECT.
31 UI was chosen as being large enough that it would not be removed by the 1 ns (about 27 UI) of Skew that is called out in footnote a to Table 116-7.

Cl **120E** *SC* **120E.3.2** *P* **374** *L* **10** # i-122

Dawe, Piers J G — Mellanox Technologie

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

The module output transition time min. spec is there to protect the module's input from too much crosstalk when connected to a host with more NEXT than the MCB. "Too much" doesn't depend on the module's output amplitude setting, so we should have an absolute spec here not a relative one.

SuggestedRemedy

This transition time spec should be replaced by a slew time spec, e.g. 3.5 ps between +/- 0.1 V. Definition of slew time similar to transition time but with fixed thresholds instead of the signal-dependent 20% and 80%.
There is less need to change the transition time spec for the host output because the connector is on the host board, so the NEXT is already in the measurement.

Response *Response Status* **U**

REJECT.
No consenus to make the change at this time.

[Editor's note added after comment resolution completed. The consensus view was that this is not a sufficiently significant issue to justify making this change.]

Cl **120E** *SC* **120E.3.2.1** *P* **374** *L* **26** # i-123

Dawe, Piers J G — Mellanox Technologie

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

There is no need for 31 UI offset between lanes. For PRBS13Q, only 1 UI offset is enough to give excellent decorrelation, better than 100-200 UI, and there is a spur at about 450 UI. PRBS31Q is believed to behave similarly (but it's such a long pattern I haven't checked). In some test setups, there is a master PRBS generator and an arrangement of splitters and cables; the cables must be kept short for good performance. 31 UI x 7 steps at 26.5625 GBd and 5 ns/m is 1.63 m - too long.

SuggestedRemedy

As the paths between the test points and the PMA front-end circuitry are not likely to differ by more than 20 mm or about 4 UI, change 31 to 6. Also in 120E.3.4.1.1 Module stressed input test procedure.

Response *Response Status* **U**

REJECT.
31 UI was chosen as being large enough that it would not be removed by the 1 ns (about 27 UI) of Skew that is called out in footnote a to Table 116-7.

Cl **121** *SC* **121.7.1** *P* **220** *L* **36** # i-127

Dawe, Piers J G — Mellanox Technologie

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

Requiring an extinction ratio of 4.5 dB restricts the range of transmitter technologies, pushing up the cost of this PMD, and 50GBASE-FR and 50GBASE-LR if they are aligned. Yet it does not benefit the link or the receiver significantly (they are protected by the TDECQ spec, and MPI penalty is a weak function of extinction ratio for PAM4 - very few 100th of dB difference). For an example of a modern direct-mod PMD spec and what a receiver can receive, 100GBASE-SR4 has a 2 dB limit. A transmitter optimized for PAM4 is likely to have a lower extinction ratio than one for NRZ, to reduce distortion.

SuggestedRemedy

Reduce the extinction ratio limit from 4.5 dB to 3 dB.

Response *Response Status* **U**

REJECT.
Insufficient justification for the proposed modification.
There is no agreement for 50GBASE-FR and 50GBASE-LR to make this modification quoting "While there was some support for the suggested remedy it may impact other parameters such as MPI. The impact should be evaluated before making the proposed change."
The commenter is invited to prepare a consensus presentation, including an analysis of the impact of the proposed modification.

Cl **121** *SC* **121.7.1** *P* **220** *L* **37** # i-128

Dawe, Piers J G — Mellanox Technologie

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

The purpose of the RIN spec has changed from something to ensure a good transmitter to something to ensure a good TDECQ measurement. The limit should be adjusted for the intended purpose.

SuggestedRemedy

When the way TDECQ handles measured noise and noise enhancement is clear, relax the RIN limits in 121, 122 and 124 according to what is necessary for successful TDECQ measurement

Response *Response Status* **U**

REJECT.
Insufficient justification and incomplete remedy.
The commenter is invited to prepare a consensus presentation with a complete proposal for a modification to the draft.

Cl **121** *SC* **121.8.1** *P* **222** *L* **39** # i-130

Dawe, Piers J G — Mellanox Technologie

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

This SSPRQ pattern will give misleading results when testing a range of transmitters - both product transmitters (line 39) and SRS signals (line 44). Same problem in clauses 122 and 124.

SuggestedRemedy

Change the first seed in Table 120-2 to one for which a minimally compliant transmitter with 0.4 dB baseline wander penalty (before and after FEC) with a random payload measures as minimally compliant (i.e. also 0.4 dB penalty) with SSPRQ.
It may be necessary to adjust another seed to get appropriate transition density characteristics.
Similarly in clauses 122, 124.

Response *Response Status* **U**

REJECT.
Insufficient evidence of the claimed problem and that the proposed remedy fixes the problem.
The current SSPRQ pattern was adopted for use in the TDECQ test (after presentation of its baseline wander characteristics) by comment 50 against D1.3. A straw poll was taken in association with that comment: Do you support adopting the SSPRQ pattern for TDECQ and SRS calibration in Clauses 122 and 123? Yes 41 No 2.
The commenter is invited to prepare a consensus presentation with a detailed analysis of the suggested problem.

Cl **121** *SC* **121.8.5.3** *P* **225** *L* **8** # i-133

Dawe, Piers J G — Mellanox Technologie

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

The draft says Pattern 6 (SSPRQ) should be used for TDECQ. But SSPRQ is a short, deliberately stressful pattern and therefore a TDECQ measurement does not give anything like the correct penalty for a range of reasonable compliant transmitters. Same problem in clauses 122 and 124.

SuggestedRemedy

Change the first seed in Table 120-2 to one for which a minimally compliant transmitter with 0.4 dB baseline wander penalty (before and after FEC) with a random payload measures as minimally compliant (i.e. also 0.4 dB penalty) with SSPRQ.
It may be necessary to adjust another seed to get appropriate transition density characteristics.

Response *Response Status* **U**

REJECT.
Insufficient evidence of the claimed problem and that the proposed remedy fixes the problem.

The current SSPRQ pattern was adopted for use in the TDECQ test (after presentation of its baseline wander characteristics) by comment 50 against D1.3. A straw poll was taken in association with that comment: Do you support adopting the SSPRQ pattern for TDECQ and SRS calibration in Clauses 122 and 123? Yes 41 No 2.

The commenter is invited to prepare a consensus presentation, with a detailed analysis of the implied problem.

Cl **121** *SC* **121.8.5.3** *P* **225** *L* **9** # i-134

Dawe, Piers J G — Mellanox Technologie

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

This says "...the oscilloscope is set up to capture samples from all symbols in the complete pattern". But with only 1 sample/UI, the record of the high frequency components of the signal would be made up by the instrument and test method, probably inaccurately. For comparison, 120E.4.2, Eye width and eye height measurement method, says "the capture includes a minimum of 3 samples per symbol, or equivalent", but an optical signal is likely to contain more high frequency components than 200GAUI-4, that could be good or bad.

SuggestedRemedy

Add "The capture includes a minimum of seven samples per symbol, or equivalent."

Response *Response Status* **U**

REJECT.
The optical signal is measured through a 0.75 x symbol rate BT4 low pass filter, so frequency content > the symbol rate is increasingly filtered out. The issue is being able to construct an eye diagram, which requires sampling of the signal waveform at many fractional UI through the signal waveform. Since the intent to construct an eye diagram is explicit in the description of the TDECQ measurement method, mandating 7 (or any other number of samples) per symbol just enforces a longer test, not a better one.
The minimum number of samples per UI would probably be different for the two types of scope allowed to be used.

Cl **121** *SC* **121.8.5.3** *P* **227** *L* **2** # i-23

RAN, ADEE Intel

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

The sentence "Each element of the cumulative probability function Cf1(yi) is multiplied by a value Gth1(yi), and then summed to calculate an approximation for the partial symbol error ratio (SER) for threshold 1" isn't quite clear.

What is "Each element of the cumulative probability function"? is it each term of the sum? What are the summation limits?

As a service to readers, please write the required calculation required to find the "approximation for the partial symbol error ratio (SER) for threshold 1" in equation form.

I assume the required calculation is

SER_1 = Sigma{y_i=-inf}{y_i=inf}C_f1(y_i)*G_th1(y_i)

SuggestedRemedy

Add a new equation (see comment, correct if necessary).

Replace the sentence "Each element of the cumulative probability function Cf1(yi) is multiplied by a value Gth1(yi), and then summed to calculate an approximation for the partial symbol error ratio (SER) for threshold 1" with a reference to the new equation.

Response *Response Status* **U**

REJECT.
The current text is in the context of an example of a linear vector, and the description of element by element multiplication was taken from a maths text book, and seems clear. A contribution with a clear equation describing the element by element multiplication would be helpful.

Cl **121** *SC* **121.8.5.3** *P* **227** *L* **22** # i-25

RAN, ADEE Intel

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

The noise definitions in the TDECQ calculation mix power and amplitude/RMS terms without clear indication which is which, and seem to include an error in the calculation of C_eq.

The noise R is an RMS value.

C_eq is a noise power enhancement compensation term.

N(w) is power spectral density; S_eq(w) is stated as frequency response, but this term is typically used for H_eq(w), the Fourier transform of the equalizer's continuous-time pulse response (T/2 pulse with energy 1). The noise transfer function is then the absolute square of the frequency response, |H_eq(w)|^2. It is not obvious that this is the intent.

C_dc is an "amplitude" correction term (unlike C_eq which is a power term).

This is very confusing and error prone. It would be useful to clarify which terms are RMS and which are power.

SuggestedRemedy

In line 22 change "The noise, R" to "The RMS value, R, of the noise".

In line 29 change "noise enhancement" to "noise power amplification".

In line 33, change "frequency response S_eq(w)" to "continuous frequency response H_eq(w)".

In equation 121-8, change "S_eq(w)" to "|H_eq(w)|^2".

Consider adding H_eq(w) to the equation definition list after N(w): "H_eq(w) is the Fourier transform of the equalizer's response to a T/2 pulse with energy 1".

Consider eliminating the term C_dc and using the coefficients A_i directly in equation 121-9, to minimize confusion with C_eq.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.
See response to comment i-59.

[Editor's note added after comment resolution completed.
The response to comment i-59 is:
Apply changes shown in http://www.ieee802.org/3/bs/public/17_03/king_3bs_01_0317.pdf with editorial license
]

Cl **121** *SC* **121.8.5.3** *P* **228** *L* **9** # i-140

Dawe, Piers J G — Mellanox Technologie

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

It may be 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.

SuggestedRemedy

Define TDECQrms = 10*log10(C_dc*A_RMS/(s*3*Qt*R)) where A_RMS is the standard deviation of the measured signal after the 19.34 GHz filter response and s is the standard deviation of a fast clean signal with OMA=0.5 and without emphasis, observed through the 19.34 GHz filter response (from memory I believe s is about 0.82). Require that TDECQrms shall not exceed the limit for TDECQ. If we think it's justified, we could allow a slightly higher limit for TDECQrms.

Response *Response Status* **U**

REJECT.
Insufficient evidence of the claimed problem and that the proposed remedy fixes the problem.
The commenter is invited to provide a contribution 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.

Cl **121** *SC* **121.8.7** *P* **228** *L* **19** # i-141

Dawe, Piers J G — Mellanox Technologie

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

In this draft (following 52.9.6), square wave is proposed for measuring the signal strength in a RIN measurement procedure. Clause 52 is 10GBASE-S/L/E, an NRZ clause. We should not use square wave here because it isn't PAM4; e.g. any transmitter linearity control circuits may fail because two of the expected PAM4 levels are missing. There is no need to use a special unnatural pattern for this. Using a mixed-frequency pattern is much more convenient and gives a slightly more relevant RIN, closer to SNR, anyway.

SuggestedRemedy

If a RIN spec is needed, define it based on PRBS13Q. All PAM4 optical clauses. Remove square wave for PAM4 from the draft.

Response *Response Status* **U**

REJECT.
This is a resubmit of comment #98 to D2.1 which was rejected with the following response: "The use of a square wave to measure RIN was discussed during the resolution of comment #152 against D2.0 with the consensus being to continue to use a square wave. The commenter is invited to provide the details of a measurement method for RIN which uses the PRBS13Q pattern."

Response to this comment is the same as to #98.

Cl **121** *SC* **121.8.9.2** *P* **230** *L* **41** # i-145

Dawe, Piers J G — Mellanox Technologie

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

Calibrating the signal for stressed receiver testing with this draft's SSPRQ then testing the receiver with PRBS31Q or scrambled idle won't work because the apparent penalty will be very different with the two patterns. This affects clauses 122 and 124 also.

SuggestedRemedy

Change the first seed in Table 120-2 to one for which a minimally compliant transmitter with 0.4 dB baseline wander penalty (before and after FEC) with a random payload measures as minimally compliant (i.e. also 0.4 dB penalty) with SSPRQ.
It may be necessary to adjust another seed to get appropriate transition density characteristics.

Response *Response Status* **U**

REJECT.
Insufficient evidence of the claimed problem and that the remedy fixes the problem.
The current SSPRQ pattern was adopted for use in the TDECQ test (after presentation of its baseline wander characteristics) by comment 50 against D1.3. A straw poll was taken in association with that comment: Do you support adopting the SSPRQ pattern for TDECQ and SRS calibration in Clauses 122 and 123? Yes 41 No 2 .

The commenter is invited to prepare a consensus presentation with a detailed analysis of the claimed problem.

Cl **122** *SC* **122.7.1** *P* **251** *L* **35** # i-148

Dawe, Piers J G — Mellanox Technologie

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

Requiring an extinction ratio of 4.5 dB restricts the range of transmitter technologies, pushing up the cost of this PMD and, unless they do better, 50GBASE-FR and 50GBASE-LR. Yet it does not benefit the link or the receiver significantly (they are protected by the TDECQ spec, and MPI penalty is a weak function of extinction ratio for PAM4 - very few 100th of dB difference). For an example of a modern direct-mod PMD spec and what a receiver can receive, 100GBASE-SR4 has a 2 dB limit. A transmitter optimized for PAM4 is likely to have a lower extinction ratio than one for NRZ, to reduce distortion.

SuggestedRemedy

Reduce the extinction ratio limit from 4.5 dB to 3 dB.

Response *Response Status* **U**

REJECT.
Insufficient justification for the requested modification.
There is no agreement for 50GBASE-FR and 50GBASE-LR to make this modification quoting "While there was some support for the suggested remedy it may impact other parameters such as MPI. The impact should be evaluated before making the proposed change."
The commenter is invited to prepare a consensus presentation, including an analysis of the impact of the required modification.

Cl **124** *SC* **124.7.1** *P* **297** *L* **31** # i-151

Dawe, Piers J G — Mellanox Technologie

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

Requiring an extinction ratio of 5 dB restricts the range of transmitter technologies, pushing up the cost of this PMD, and 100GBASE-DR if it is aligned. Yet it does not benefit the link or the receiver significantly (they are protected by the TDECQ spec, and MPI penalty is a weak function of extinction ratio for PAM4 - very few 100th of dB difference). Depending on technology, a transmitter optimized for PAM4 may need a lower extinction ratio than one for NRZ, to reduce distortion.

SuggestedRemedy

Reduce the extinction ratio limit from 5 dB to e.g. 3 dB.

Response *Response Status* **U**

REJECT.
Insufficient justification for the requested modification.
The reference to 100GBASE-DR is not appropriate, because there is no agreement to make this modification.
The commenter is invited to prepare a consensus presentation, including an analysis of the impact of the required modification.

Cl **124** *SC* **124.8.9** *P* **301** *L* **28** # i-153

Dawe, Piers J G — Mellanox Technologie

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

If the jitter corner frequency for 26.5625 GBd (NRZ and PAM4) is 4 MHz, shouldn't it be 8 MHz for 53 GBd PAM4? Or at least, the low frequency (sloping) part of the mask should scale with signalling rate, i.e. align if expressed in time vs. frequency. Compare 87.8.11.4 and 88.8.10: 4 MHz for 10.3125 GBd, 10 MHz for 25.78125 GBd.

SuggestedRemedy

Add another exception with a table like Table 121-12 but with the frequencies doubled.

Response *Response Status* **U**

REJECT.
The jitter corner frequency was extensively discussed within the Task Force with multiple presentations on the topic. The CRU corner frequency was chosen to be 4 MHz for all interfaces (including 400GBASE-DR4) in the March 2016 TF meeting as recorded in: http://www.ieee802.org/3/bs/public/16_03/anslow_3bs_04_0316.pdf.