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 **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 (See http://grouper.ieee.org/groups/802/3/cd/public/adhoc/archive/hidaka_020117_3cd_adhoc.pdf and further as yet unpublished work)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 **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* **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* **121.8.5.1** *P* **227** *L* **52** # r01-13

RAN, ADEE
Intel

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

(page 224 according to footer in CMP document)

This is a follow-up on i-131 due to changes in 121.8.5.a and 121.8.5.3 which make it more relevant.

The 31-UI offset is now required "so that the symbols on each lane are not correlated within the PMD". But that is incorrect; the symbols are fully correlated, with a constant offset.

The rebuttal of comment i-131 claimed that having crosstalk "locked to the pattern under test" enables it to be "correctly processed by the equalizer". But this makes the crosstalk strongly correlated with the measured signal (even with 31 UI offset) and appear as a high-probablity noise component (due to the short SSPRQ length); where in real life, crosstalk will be totally uncorrelated with the transmitter signal, and likely closer to Gaussian. This results in overly pessimistic accounting of crosstalk.

With TDECQ being tested without averaging (as now added in 121.8.5.3), there seems to be no need for requiring the SSPRQ pattern on all lanes. The statistics of uncorrelated crosstalk will be represented better if the measurement is done with adjacent lanes transmitting a signal with a different period, such as PRBS31Q or PRBS13Q. Since the measurement is not averaged, the statistics can be captured correctly.

In addition for making it a more representative test, controlling SSPRQ per lane and not requiring a 31-UI offset (which does not really help anyway) may reduce complexity in the PMA design.

SuggestedRemedy

Require TDECQ measurement to be performed with SSPRQ transmitted only on the lane under test, with other lanes transmitting PRBS31Q or a valid PCS pattern.

Change SSPRQ generator control to be per-lane (in 120.5.11.2.3 and 45.2.1.124).

Delete the requirement to have at least a 31 UI delay between lanes in 120.5.11.2.3 and in 121.8.5.1, and delete the words "so that the symbols on each lane are not correlated within the PMD" (they are incorrect).

Apply corresponding changes in the TDECQ subclauses of other PMD clauses.

Grant license to the editors to implement the changes correctly across the multiple clauses involved.

Response *Response Status* **U**

REJECT.
This comment makes a similar proposal to comment i-131, which was rejected with the response:

"The TDECQ test (and SECQ test) are based on capturing the complete SSPRQ pattern and passing it through a reference equalizer. The measurement is allowed to be made using an equivalent-time sampling oscilloscope. By requiring that all lanes are receiving the SSPRQ pattern, any crosstalk from the other lanes is locked to the pattern under test, captured by the oscilloscope as a distortion of the waveform and correctly processed by the equalizer. Because of the offset between the lanes, the crosstalk will be different for the various occurrences of each symbol type. If the draft is changed to allow PRBS13Q or PRBS31Q on the other lanes, then the crosstalk will no longer be locked to the pattern under test and will appear as noise when captured using an equivalent-time sampling oscilloscope and will not be processed correctly by the reference equalizer since the frequency profile of the crosstalk is lost."

The advantage of retaining the frequency content of the crosstalk when using an equivalent time oscilloscope outweighs any advantage of improved randomness when using a different pattern on the other lanes.

Cl **120D** *SC* **120D.3.1.1** *P* **357** *L* **29** # r01-22

RAN, ADEE
Intel

Comment Type **GR** *Comment Status* **A**

(page 353 according to footer in CMP document)

Current SNR_ISI value of 38 dB is too high to be the minimum requirement (although stated as maximum - this is the subject of another comment).

In measurements performed with state-of-the-art scope and an instrument-grade pattern generator, connected by a short instrument-grade cable, the best SNR_ISI achieved was 39.3 dB, and that was with equalization off. This is only 1.3 dB better than the current minimum. This may be an "ISI floor" of the scope, cables, etc., or actual ISI in the transmitter.

Using a packaged transmitter with a supplied evaluation board, high-performance connectors, with short cables to the same scope, resulted in only 36.9 dB at room temperature and without equalization.

With maximum equalization, the pulse peak will be 60% of the unequalized peak, while the ISI can be assumed to be roughly the same. This will result in a degradation of 4.4 dB in SNR_ISI, so the instrument-grade transmitter will actually have SNR_ISI of only 34.9 dB.

For the channels targeted by the C2C specification, and with a CTLE+DFE equivalent assumed in the receiver, operating at the maximum Tx equalization state is unlikely (as this would reduce the signal and exacerbate the effects of TX ISI, crosstalk and other noises). The COM analysis of contributed channels resulted in Tx equalization much lower than the maximum. Therefore, it is reasonable not to judge the transmitter by this state. More likely, the Tx equalization will reduce the peak by up to 2 dB relative to the unequalized pulse.

To achieve technical feasibility with a broad market potential, the standard should allow some margin for manufacturing variability and temperature dependence. The specification should be such that an instrument-grade transmitter will have a margin of ~2 dB.

At the bottom line, the proposal is to specify minimum SNR_ISI as 4 dB below the best measured value with an instrument-grade unequalized transmitter, or 35.3 dB.

The current value was set by comment i-69 which states: "the RSS sum of the SNDR and SNRisi should equal the RSS sum of the TxSNR used in COM plus the SNRisi produced by the COM package". The normalized RSS of the current values of SNDR and SNR_ISI is 0.03, or 30.2 dB below the signal; to keep it the same with SNR_ISI of 35.3 dB, the required SNDR should be slightly increased to 31.8 dB.

SuggestedRemedy

Change the minimum SNR_ISI value from 38 to 35.3 dB.

Change the minimum SNDR from 31 to 31.8 dB.

In 120D.3.1.7, change "The SNR_ISI specification shall be met for all transmit equalization settings" to "The SNR_ISI is measured with Local_eq_cm1 and Local_eq_c1 set to zero".

Add another NOTE at the end of 120D.3.1.7:
NOTE 2--The observed SNR_ISI can be significantly influenced by the measurement setup, e.g. reflections in cables and connectors. High-precision measurement and careful calibration of the setup are recommended.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

In Table 120D-1:
Change the minimum SNR_ISI value from 38 to 34.8 dB.
Change the minimum SNDR from 31 to 31.5 dB.
Change Linear fit pulse peak (min) from 0.736*Vf to 0.76*Vf

In Table 120D-8:
Change Av and Afe values from 0.45 to 0.44

Add another NOTE at the end of 120D.3.1.7:
NOTE 2--The observed SNR_ISI can be significantly influenced by the measurement setup, e.g. reflections in cables and connectors. High-precision measurement and careful calibration of the setup are recommended.

Cl **120** *SC* **120.5.11.2.3** *P* **202** *L* **18** # r01-32

Dawe, Piers J G — Mellanox Technologie

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

Following up D3.0 comment 109: this SSPRQ is not suitable for use in TDECQ or stressed receiver calibration because measurements with this pattern do not give the correct (post FEC) penalty. Neither dawe_3bs_01a_0317 nor anslow_01_0417_smf show a suitable pattern. See associated comment against 121.8.5.3, 122, 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 after FEC with a random payload measures as minimally compliant (i.e. also 0.4 dB baseline wander penalty) on a pre-FEC BER basis with SSPRQ. This will be a pattern between the red and light brown curves in dawe_3bs_01a_0317 slide 6.

Response *Response Status* **U**

REJECT.

A similar proposal was made in i-109 which was rejected. No consensus has been reached on changes to this pattern in the ad hoc calls.

After further discussion there is still no consensus for a change to the draft.

[Editor's note added after comment resolution completed.
The response to comment i-109 is:
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.3.1.1** *P* **353** *L* **24** # r01-36

Dawe, Piers J G — Mellanox Technologie

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

Transmitter Output residual ISI SNR_ISI (max) 38 dB is too high - probably can't measure the IC through the test fixture and cables.

SuggestedRemedy

Start by checking whether Gaussian assumptions are tripping us up.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

See response to comment #r01-22

[Editor's note added after comment resolution completed.
The response to comment r01-22 is:
In Table 120D-1:
Change the minimum SNR_ISI value from 38 to 34.8 dB.
Change the minimum SNDR from 31 to 31.5 dB.
Change Linear fit pulse peak (min) from 0.736*Vf to 0.76*Vf

In Table 120D-8:
Change Av and Afe values from 0.45 to 0.44

Add another NOTE at the end of 120D.3.1.7:
NOTE 2--The observed SNR_ISI can be significantly influenced by the measurement setup, e.g. reflections in cables and connectors. High-precision measurement and careful calibration of the setup are recommended.
]

Cl **120D** *SC* **120D.3.1.8** *P* **358** *L* **46** # r01-41

Dawe, Piers J G — Mellanox Technologie

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

I doubt that the low frequency RL at 14.25 dB is significant for signal integrity compared with the 8.7 dB at 6 GHz. This RL is much tighter than CEI-56G-MR at low (and high) frequency but looser between 4 and 9 GHz.

SuggestedRemedy

Change 14.25 - f to 12 -0.625f

Response *Response Status* **U**

REJECT.
No consensus to make a change at this time, but further investigation is encouraged.

[Editor's note added after comment resolution completed. The consensus view was that further investigation of the effect of Return Loss at low frequencies should take place, but no change to the equation can be justified at this time.]

Cl **120E** *SC* **120E.3.2** *P* **376** *L* **5** # r01-42

Dawe, Piers J G — Mellanox Technologie

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

Far-end pre-cursor ratio doesn't seem like the right tool to solve the issue raised in healey_3bs_01a_0317, which seeks to outlaw "transmitter A1" that gives more than 4 dB COM anyway, so the limit for far-end pre-cursor ratio seems too restrictive. The complaint seems to be that even if the eye is open after the software channel, some receivers might struggle after their own package loss.

SuggestedRemedy

If there is an issue, consider increasing the loss in the software channel to moving the "far end" to after a reasonable package loss, and making a small adjustment the FE eye height and width to compensate. Anyway, relax the far-end pre-cursor ratio limit. If a limit remains, consider if there needs to be a minimum as well as a maximum limit.
Review the way this works for a reasonable variety of channels.

Response *Response Status* **U**

REJECT.
The commenter has not provided any evidence to support his assertion that the limit for far-end pre-cursor ratio is more restrictive than necessary.

Cl **121** *SC* **121.8.5.3** *P* **226** *L* **8** # r01-48

Dawe, Piers J G — Mellanox Technologie

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

Following up on D3.0 comment 133: the draft says Pattern 6 (SSPRQ) should be used for TDECQ. Today's SSPRQ is more stressful in pre-FEC measurements than the service pattern (long scrambler) with FEC, so today's TDECQ measurement does not give the correct penalty for a range of reasonable and compliant transmitters. Same problem in clauses 122 and 124. See associated comment against 120.5.11.2.3.

SuggestedRemedy

Change the first seed in Table 120-2 to one for which a minimally compliant transmitter with 0.4 dB baseline wander penalty after FEC with a random payload measures as minimally compliant (i.e. also 0.4 dB baseline wander penalty) on a pre-FEC BER basis with SSPRQ. This will be a pattern between the red and light brown curves in dawe_3bs_01a_0317 slide 6.

Response *Response Status* **U**

REJECT.
This topic has been discussed at the SMF Ad Hoc with no consensus being reached for a change.
After further discussion there is still no consensus for a change to the draft.

[Editor's note added after comment resolution completed.
Evidence that no change is needed was given in:
http://www.ieee802.org/3/bs/public/17_05/anslow_3bs_03_0517.pdf
]

Cl **124** *SC* **124.8.9** *P* **302** *L* **31** # r01-55

Dawe, Piers J G — Mellanox Technologie

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

Following up on D3.0 comment 153: if the jitter corner frequency for 26.5625 GBd (NRZ and PAM4) is 4 MHz, the low frequency (sloping) part of the jitter mask should scale with signalling rate, i.e. align if expressed in time vs. frequency, to avoid a need for a poorly specified wander buffer in the 2:1 muxes in a 400GBASE-DR4 module. Compare 87.8.11.4 and 88.8.10: 4 MHz for 10.3125 GBd, 10 MHz for 25.78125 GBd. History: anslow_3bs_04_0316 does not contain reasoning, refers to ghiasi_3bs_01_0316 which does not address wander and buffering.

SuggestedRemedy

Add another exception for the SRS procedure, with a table like Table 121-12 but with the frequencies doubled.
Or, replacing second row after the header row:
80 kHz < f <= 500 kHz 4e5/f
500 kHz < f <= 1 MHz 2e11/f^2
1 MHz < f <= 4 MHz 2e5/f

Response *Response Status* **U**

REJECT.
This issue was already discussed in response to comment i-153 to D3.0 which was:
"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."

The possible need for a buffer was discussed in presentations made leading up to this decision. For example, see:
http://www.ieee802.org/3/bs/public/16_01/ghiasi_3bs_01a_0116.pdf#page=15

There was no consensus to make a change to the draft.

Cl **121** *SC* **121.7.1** *P* **221** *L* **25** # r02-28

Dawe, Piers J G — Mellanox Technologie

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

PAM4 optics is still new and raw, we are still debugging the specification methodology, and we have seen far too little experimental information showing technical and economic feasibility. It looks like this PMD can be made to work but as measurements with the new TDECQ method and with new receiver designs become available, we expect the optical power levels can be reduced and the spec as in this draft will be uneconomic.

SuggestedRemedy

Bring more evidence for what optical power levels and TDECQ limits are right; in particular, TDECQ measurements with SSPRQ, and correlation to actual receiver performance.
Based on evidence, reduce all the optical power levels for 200GBASE-DR4 by 0.5, 1 or 1.5 dB (with other adjustments for other reasons). Review the TDECQ limit.

Response *Response Status* **U**

REJECT.
This comment does not apply to the substantive changes between IEEE P802.3bs/D3.2 and IEEE P802.3bs/D3.1 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

The suggested remedy does not propose any changes to the draft.

Cl **121** *SC* **121.8.5.1** *P* **226** *L* **49** # r02-31

Dawe, Piers J G — Mellanox Technologie

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

Using the same pattern on the aggressor lanes (correlated crosstalk) is very unusual. Does what we gain in correctly handling the spectrum of the deterministic part of the crosstalk outweigh what we lose in inconsistency vs. UI- and sub-UI phasing? As D3.1 comment 13 points out, using the conventional uncorrelated crosstalk can simplify the PMA. It should be possible to calculate the relative measurement accuracy of the two approaches.

SuggestedRemedy

Work out which is better; change the crosstalk patterns here and the related pattern generator options in Clause 120 as appropriate.

Response *Response Status* **U**

REJECT.
The suggested remedy does not propose any changes to the draft.

The commenter is invited to perform the calculation suggested in the comment and prepare a consensus presentation with proposed changes to the draft.

Cl **121** *SC* **121.8.5.3** *P* **229** *L* **34** # r02-34

Dawe, Piers J G — Mellanox Technologie

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

The change of the reference bandwidth from 19.34 GHz to 13.28125 means that an ideal signal (fast, no noise or jitter, no emphasis) has a TDECQ that is far from zero. We could live with this and change many other numbers including "results in at least half of the dB value of the stressed eye closure (SECQ)" but doing so makes the budget hard to understand. In the remedy I assume the offset is 0.5 dB; this should be checked.

SuggestedRemedy

In Eq. 121-12, change 1 to 0.891, which is 0.5 dB less. Add a NOTE to explain that this number represents the TDECQ of an ideal signal (fast edges, no noise or jitter, no emphasis).
Or, change 1 to a new parameter, value 0.891, add to the "where" list.
Or, modify equation to TDECQ = 10 log10(...) - TDECQ0 where TDECQ0 is 0.5 ...

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.
See response to comment r02-2

[Editor's note added after comment resolution completed.
The response to comment r02-2 is:

Implement the changes shown in
http://www.ieee802.org/3/bs/public/adhoc/smf/17_06_27/anslow_02_0617_smf.pdf
with the following exceptions:
In Tables 121-7, 122-11, 122-12, and 124-7:
leave the Receiver sensitivity (OMAouter), each lane (max) unchanged
In footnote c of each table change the addition to "and is defined for a transmitter with SECQ of 0.9 dB"
]

Cl **121** *SC* **121.8.5.3** *P* **229** *L* **42** # r02-35

Dawe, Piers J G — Mellanox Technologie

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

Updating D3.0 comment 140:
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 (up to 2.5/2 dB worse than the SRS test?) With some of the changed low-bandwidth TDECQ being used to equalize the reference receiver's own bandwidth, this issue becomes more apparent.

SuggestedRemedy

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. s is close to the standard deviation of a fast clean signal with OMA=0.5 and without emphasis, observed through the 13.28125 GHz filter response, according to what level of dirty-but-emphasised signal we decide is acceptable. Require that TDECQrms shall not exceed the limit for TDECQ.

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 **122** *SC* **122.7.1** *P* **252** *L* **14** # r02-36

Dawe, Piers J G — Mellanox Technologie

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

PAM4 optics is still new and raw, we are still debugging the specification methodology, and we have seen far too little experimental information showing technical and economic feasibility. As measurements with the new TDECQ method and with new receiver designs become available, it may be that optical power levels can be reduced and the spec as in this draft would be uneconomic.

SuggestedRemedy

Bring more evidence for what optical power levels and TDECQ limits are right; in particular, TDECQ measurements with SSPRQ, and correlation to actual receiver performance.
Based on evidence, consider reducing all the optical power levels in this clause except the -30 dBm signal detect limit by 0.5 or 1 dB (with other adjustments for other reasons).
Review the TDECQ limits.

Response *Response Status* **U**

REJECT.
This comment does not apply to the substantive changes between IEEE P802.3bs/D3.2 and IEEE P802.3bs/D3.1 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

The suggested remedy does not propose any changes to the draft.

Cl **124** *SC* **124.7.1** *P* **298** *L* **4** # r02-37

Dawe, Piers J G — Mellanox Technologie

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

PAM4 optics is still new and raw, we are still debugging the specification methodology, and we have seen too little experimental information showing technical and economic feasibility. As measurements with the new TDECQ method and with new receiver designs become available, it may be that optical power levels can be reduced and the spec as in this draft would be uneconomic.

SuggestedRemedy

Bring more evidence for what optical power levels and TDECQ limits are right; in particular, TDECQ measurements with SSPRQ, and correlation to actual receiver performance.
Based on evidence, reduce all the optical power levels for 400GBASE-DR4 by 0.5 or 1 dB (with other adjustments for other reasons). Review the TDECQ limit.

Response *Response Status* **U**

REJECT.
This comment does not apply to the substantive changes between IEEE P802.3bs/D3.2 and IEEE P802.3bs/D3.1 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

The suggested remedy does not propose any changes to the draft.

Cl **121** *SC* **121.8.7** *P* **302** *L* **20** # r02-39

Dawe, Piers J G — Mellanox Technologie

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

With the lower receiver bandwidth, measuring RIN in approximately the signaling rate (twice as much) seems too much; 1/2 to 3/4 would be better. A T-spaced equalizer cannot independently adjust for good ISI and RIN filtering, so can an adequate estimate of RIN can be obtained as a by-product of the TDECQ procedure? While a T/2-spaced equalizer could enhance the RIN, it would not choose to do so if RIN were a problem.

SuggestedRemedy

Review; simplify RIN measurement to a Qsq measurement (see 68.6.7) or eliminate as appropriate. Remove 120.5.11.2.4 Square wave (quaternary) test pattern, and associated registers.

Response *Response Status* **U**

REJECT.

The suggested remedy suggests 2 different approaches to change the draft.
Changing the RIN measurement to a Qsq measurement has not been demonstrated to provide the same safeguards that are expected from the RIN requirement.

Eliminating the RIN measurement was discussed in the response to comment #130 against D2.0 on the basis that "The transmitter RINxOMA spec is intended to screen out potentially bad transmitters even if the noise correction required by the TDECQ test is not very accurate."

Cl **124** *SC* **124.8.9** *P* **302** *L* **46** # r02-40

Dawe, Piers J G — Mellanox Technologie

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

Following up on D3.0 comment 153 and D3.1 comment 55: if the jitter corner frequency for 26.5625 GBd (NRZ and PAM4) is 4 MHz, the low frequency ends of the jitter masks must align or be in the right order if expressed in time vs. frequency, i.e. should scale with signalling rate if in UI. If this is not done, the required depth of the LF jitter buffer in the 2:1 muxes in a 400GBASE-DR4 module is unbounded and the low frequency jitter generation requirements on the module become unreasonable. Compare 87.8.11.4 and 88.8.10: 4 MHz for 10.3125 GBd, 10 MHz for 25.78125 GBd. History: anslow_3bs_04_0316 does not contain reasoning, refers to ghiasi_3bs_01_0316 which does not address wander and buffering. ghiasi_3bs_01a_0116.pdf#page=15 shows FIFOs but does not establish a workable spec. Slide 14 shows they can be avoided: this is what we have for 400GAUI-8 or 400GAUI-16 with 400GBASE-xR8. I have no evidence that the problems described in the second sentence have been considered or solved by the committee.

SuggestedRemedy

Add another exception for the SRS procedure, with a table like Table 121-12 replacing second row after the header row:
80 kHz < f <= 250 kHz 4e5/f
250 kHz < f <= 500 kHz 1e11/f^2
1 MHz < f <= 4 MHz 2e5/f
Or, with the UIs doubled vs. Table 121-12:
f < 40 kHz Not specified
40 kHz < f <= 4 MHz 4e5/f
4 MHz < f <= 10 LB 0.1
Increase the TDECQ limit to share the burden appropriately between transmitter and receiver.
This option means the 100G/lane receiver has to tolerate no more timing slew rate (in ps/us) than that agreed for 50G/lanes.
Or, increase jitter by 50% and corner frequency by 33%:
f < 40 kHz Not specified
40 kHz < f <= 6 MHz 4e5/f
5.333 MHz < f <= 10 LB 0.075
and add an exception in 124.8.5 that the CRU corner frequency is 5.333 MHz. Increase the TDECQ limit to share the burden between transmitter and receiver.
To do the job properly with the first option, in 124.8.5 we should add another exception to the CRU with a corner frequency of 4 MHz and a slope of 20 dB/decade (in 121.8.5.1): add a pole at 250 kHz and a zero at 500 kHz. I am advised that this can be done in hardware (in software, anything is possible).

Response *Response Status* **U**

REJECT.
The suggested remedy is proposing to place an extra burden on the receiver by allowing transmitters with a higher level of TDECQ which may be due to ISI and also by requiring a higher level of jitter tolerance.
The commenter has not demonstrated that this extra burden is less onerous than putting a buffer in the PMA.

For the second option in the suggested remedy the commenter is invited to build consensus for an increase of the corner frequency to be above 4 MHz.

Cl **120D** *SC* **120D.3.1.1** *P* **353** *L* **24** # r02-42

Dawe, Piers J G — Mellanox Technologie

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

Signal-to-noise-and-distortion ratio (min) 31.5 dB is too high (increased by D3.1 comment 22, so even worse than before) - probably can't measure the IC through the test fixture and cables. I suspect there is double counting of jitter in SNDR and as jitter, in COM.

SuggestedRemedy

Remove the double counting. Reduce the SNDR limit to something that can reasonably be measured, or change the measurement method.

Response *Response Status* **U**

REJECT.
The presentation:
http://www.ieee802.org/3/bs/public/17_07/dawe_3bs_04_0717.pdf was reviewed.
Changing the SNDR limit to 28.5 dB is considered to be placing too great a burden on the receiver and it has not been demonstrated that implementations cannot meet the current specification.

Cl **120D** *SC* **120D.3.1.1** *P* **353** *L* **26** # r02-43

Dawe, Piers J G — Mellanox Technologie

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

Following D3.1 comments 22 and 36: transmitter Output residual ISI SNR_ISI (min) 34.8 dB is still too high - probably can't measure the IC through the test fixture and cables, even test equipment fails this limit. The warning NOTE in 120D.3.1.7 shows the issue, but doesn't solve it.

SuggestedRemedy

It may be necessary to move away from the SNR_ISI method.

Response *Response Status* **U**

REJECT.
No remedy provided

Cl **120D** *SC* **120D.3.1.1** *P* **354** *L* **36** # r02-44

Dawe, Piers J G — Mellanox Technologie

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

Following D3.1 comment 41: the low frequency RL at 14.25 dB is insignificant for signal integrity compared with the 8.7 dB at 6 GHz. This RL is much tighter than CEI-56G-MR at low (and high) frequency (although apparently looser between 4 and 9 GHz).

SuggestedRemedy

Change 14.25 - f to 12 -0.625f

Response *Response Status* **U**

REJECT.
Re-statement of comment r01-41 which was rejected with the response:
No consensus to make a change at this time, but further investigation is encouraged.
[Editor's note added after comment resolution completed. The consensus view was that further investigation of the effect of Return Loss at low frequencies should take place, but no change to the equation can be justified at this time.]

While additional work has been done on this topic, there is still no consensus to make a change.

Cl **120E** *SC* **120E.3.1** *P* **371** *L* **20** # r02-46

Dawe, Piers J G — Mellanox Technologie

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

Building on D3.0 comment 119: The host is allowed to output a signal with 900 mV peak-to-peak amplitude but only 32 mV eye height - a very bad signal. If the module is exactly like the reference receiver, that would work, but with a good but slightly different receiver the eye will collapse.

SuggestedRemedy

We need some other spec to protect the module from such unexpected signals. A vertical eye closure spec will probably work. I'll try to bring a presenttaion.

Response *Response Status* **U**

REJECT.
No presentation providing a suggested remedy for this comment was submitted.
While a vertical eye closure specification was considered worth further investigation, no consensus was reached to make a change to the draft.

Cl **120E** *SC* **120E.3.2** *P* **376** *L* **7** # r02-47

Dawe, Piers J G — Mellanox Technologie

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

It turns out that meeting the five module output specs simultaneously with good tolerances is not feasible (near and far end eye height and width, far-end pre-cursor ratio). And, according to my understanding of healey_3bs_01a_0317, a far-end pre-cursor ratio of 1%, 2% or 9% provides a healthy COM for a C2C receiver but a C2M receiver after a COM package with a now obsolete Cd has a problem with 9%, so the 2.5% limit in the draft seems arbitrary.
This is a follow-up to D3.1 comment 42.

SuggestedRemedy

Decrease the limit for far-end eye height from 70 mV to 45 mV.
Widen the pre-cursor ratio limit from +/-2.5% to +/-3.5%.
Consider increasing the loss in the software channel (moving the "far end" to after a reasonable package loss), and making a small adjustment to the far-end eye height and width to compensate.
If the loss is not increased, consider if an asymmetrical pre-cursor ratio limit would be more effective.
Review the way this works for a reasonable variety of channels.
Review what range of CTLE peaking is consistent with the insertion loss budget.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

This issue of changing the near end eye height and pre-cursor ratio was discussed at the 6th July electrical ad hoc, but no consenus was reached on how to address it.

A Straw poll was taken:
A) Change the near end eye height from 70 mV to 60 mV
B) Make no change to the draft
A 9
B 16

Change
". The setting of the reference CTLE is the same used to measure eye width and height."
To
". Any setting of the reference CTLE for which the eye width and height satisfy the limits in Table 120E-3, may be used."

Cl **120D** *SC* **120D.4** *P* **362** *L* **9** # r02-56

Dudek, Michael — Cavium

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

Variations in package impedance and die impedance while still meeting the Tx and Rx specifications (including return loss) cause worse COM for some channels than is obtained with the values used in the COM test for the channel resulting in a "hole" in the budget. (See e.g. Hidaka_3cd_01a_0317, Dudek_3bs_02_0517). This hole is around 0.5dB.

SuggestedRemedy

Change the required value of COM for the channel from 3.0dB to 3.5dB while leaving the calibration of the interference tolerance test at 3.0dB COM. As an alternative the burden to close the budget could be shifted from the channel to the Rx by using 3.0dB as the channel COM and 2.5dB COM for the interference tolerance test calibration or could be shared as long as there is 0.5dB difference between them.. Change PICS CC1 to this revised value.

Response *Response Status* **U**

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3bs/D3.2 and IEEE P802.3bs/D3.1 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

A straw poll was taken:
I support the following option (choose one):
A) Change the required value of COM for the channel from 3 dB to 3.1 dB and change the calibration of the interference tolerance test COM from 3 dB to 2.9 dB.
B) Change the required value of COM for the channel from 3 dB to 3.2 dB while leaving the calibration of the interference tolerance test COM at 3 dB.
C) No change (i.e., both COM for the channel and calibration of the RX ITT remain at 3 dB).
A 2
B 0
C 24