Cl **136** *SC* **136.11.7** *P* **225** *L* **9** # 71

Hidaka, Yasuo — Fujitsu Lab. of Americ

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

As explained in hidaka_061417_3cd_01_adhoc.pdf, the limit of variation of compliant channels will grow, if we use a single reference value for the COM impedance parameters, and the single reference value is different from the nominal value. In order to minimize the variation of compliant channels, we should use the nominal value as the single reference value, or we should use multiple reference values. Reduction of variation helps to improve margin for interoperability, which is not guaranteed in the current specification. When we change the COM impedance parameters, we should also consistently change A_v, A_fe, A_ne to get the same signal amplitude at TP0a from reference Tx in COM, and we should also change the COM value to avoid changing the pass / fail status of existing channels. The consistent changes required to A_v, A_fe, and A_ne were reported in hidaka_060717_3cd_adhoc-v2.pdf slide 12. The consistent change required to COM value was reported in hidaka_061417_3cd_01_adhoc.pdf slide 14-18.

SuggestedRemedy

Change the following COM parameter values in Table 136-15:

Package Z_c from 90 ohm to 95 ohm
R_d from 55 ohm to 50 ohm
A_v from 0.44 V to 0.415 V
A_fe from 0.44 V to 0.415 V
A_ne from 0.63 V to 0.604 V

In the second paragraph of 136.11.7.1, P226, L31, change
"the parameter values given in Table 92-12"
to
"the parameter values given in Table 92-12 excepting that Z_c is 100.0 ohm".

For clarification of the intention of the value, in the parameter column of Table 136-15, change
"Package transmission line characteristic impedance"
to
"Package transmission line nominal characteristic impedance".

In Table 136-14, change the value of Minimum COM from 3 dB to 3.3 dB.

In the third paragraph of 136.11.7, P224, L40, change
"shall be greater than or equal to 3 dB"
to
"shall be greater than or equal to 3.3 dB".

Response *Response Status* **U**

REJECT.

hidaka_3cd_01_0717 was reviewed. There is no consensus to make the proposed changes.

The commenter is encourged to build consensus on a proposal in the ad hoc meetings.

Cl **136** *SC* **136.9.4.2.2** *P* **220** *L* **28** # 72

Hidaka, Yasuo — Fujitsu Lab. of Americ

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

Test channel of receiver interference tolerance test is specified as the cable assembly meets the requirements of 136.11 and the cable assembly test fixture meets the requirements of Annex 136B. However, as explained in hidaka_3cd_01a_0517.pdf and hidaka_060717_3cd_adhoc-v2.pdf, the cable assembly just meeting the requirements of 136.11 allows use of a cable assembly with the worst return loss, which will cause interoperability problems between compliant channel and compliant Rx. As explained in hidaka_3cd_02_adhoc-v2.pdf, the return of of the test channel for Rx ITT is important to improve margin for interoperability. We should specify the Rx-side return loss of the test channel tighter than the return loss of the compliant channel so that a good test channel is always used for Rx ITT.

SuggestedRemedy

Change the sentence of 136.9.4.2.2
"The test channel is the same as the one defined in 110.8.4.2.2, except that the cable assembly meets the requirements of 136.11 and the cable assembly test fixture meets the requirements of Annex 136B."

to

"The test channel is the same as the one defined in 110.8.4.2.2, except that the cable assembly meets the requirements of 136.11, the cable assembly test fixture meets the requirements of Annex 136B, and the differential return loss of the test channel measured at Rx test reference including the cable assembly meets Equation (92-38)."

Response *Response Status* **U**

REJECT.

There is no consensus to implement the suggested remedy.

Cl **136** *SC* **136.11** *P* **223** *L* **42** # 113

Dudek, Mike — Cavium

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

Equation 92-27 for the differential return loss gives 5.3dB return loss at 13.28GHz. This is not the 6dB listed and is a relatively poor value and could lead to significant differences between system performance with a real host and the COM calculated with the single 110 Ohm host board trace equivalent. Work on backplanes and C2C (e.g. Hidaka_3cd_01a_0317, Dudek_3bs_02_0517) has shown that this affect is significant and it would be better to test COM with nominal impedances and have a guard band between the channel COM and the Interference tolerance COM.

SuggestedRemedy

Change 6 to 5.3 Change the COM value to 3.5dB. In table 136-15 change the value of Rd to 50 Ohm, the value of Zc to 95 Ohm, On page 224 line 40 change the value of COM to 3.5dB. Change the impedance of the test trace from TP0 to TP1 and TP4 to TP5 to 100 Ohm by changing on page 226 line 41 from "using zp = 151 mm in length, representing an insertion loss of 6.42 dB at 13.28 GHz on each PCB." to "using Zc = 100 Ohm and zp = 151 mm in length, representing an insertion loss of 6.42 dB at 13.28 GHz on each PCB." Also change to 3.5dB in PICS CA8.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

In Table 136-14, Change "Minimum differential return loss at 13.28 GHz" from 6 dB to 5.3 dB.

The rest of the suggested remedy requires more consensus building.

See also #71

Cl **139** *SC* **139.7.1** *P* **286** *L* **19** # 126

Dawe, Piers — Mellanox

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

For SRS testing, while Table 138-12 following 802.3by Table 95-10 allows PRBS31Q, scrambled idle (with FEC) or valid 50GBASE-SR, 100GBASE-SR2, or 200GBASE-SR4 signal, but this Table 139-10 (following the older 802.3ba?) allows only PRBS31Q and scrambled idle. The 58-bit scrambler is so long that we can't tell the statistics of RS-FEC encoded scrambled idle from any other valid 50GBASE-R signal. RF, which is a valid 50GBASE-R signal, is often more convenient than scrambled idle. Table 89-10 (40GBASE-FR) also allows PRBS31, scrambled idle or valid 40GBASE-R signal.

SuggestedRemedy

Change "3 or 5" to "3, 5, 6 or valid 50GBASE-R signal". Also in Table 140-10.

Response *Response Status* **U**

REJECT.

The recommended test patterns 3 (PRBS31Q) or 5 (scrambled idles) are more than adequate for SRS testing. The current approach is used in in-force SMF Clauses 87 and 88 and in progress (for P802.3bs) Clauses 121, 122 and 124. For consistency with corresponding Clauses in P802.3bs the pattern set should stay as it is.

SSPRQ (pattern 6) is intended only for transmitter testing. Therefore it is not relevant for this test and may overstress the receiver.

Cl **138** *SC* **138.7.1** *P* **262** *L* **18** # 127

Dawe, Piers — Mellanox

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 (up to 4/2 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.
This is an issue for all the PAM4 optical PMDs, although it may be worse for MMF because of the high TDECQ limit.

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 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 - can be calculated when the filter bandwidth is stable).
Set limit for TDECQrms according to what level of dirty-but-emphasised signal we decide is acceptable, add max TDECQrms row to the table. Alternatively, if the same relative limit is acceptable for all PAM4 optical PMDs, the limit could be in the TDECQ procedure 121.8.5.3 as proposed in P802.3bs D3.2 comment r02-35.
Similarly in clauses 139, 140.

Response *Response Status* **U**

REJECT.

A similar comment was made to P802.3bs D3.2 via comment r02-35, which was rejected.

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

Cl **140** *SC* **140.6.1** *P* **306** *L* **33** # 128

Dawe, Piers — Mellanox

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 100GBASE-DR by 0.5 or 1 dB (with other adjustments for other reasons). Review the TDECQ limit.

Response *Response Status* **U**

REJECT.

The suggested remedy does not propose a specific change to the draft.

Cl **139** *SC* **139.7.7** *P* **289** *L* **15** # 133

Dawe, Piers Mellanox

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, so a T-spaced reference equalizer and a T/2-spaced product equalizer are compatible from this point of view, I think. As 52.9.6 says, this RIN method is intended for components (TOSAs) not a "system level test" suitable for a complete optical module.
This is much the same as P802.3bs D3.2 comment r02-39.

SuggestedRemedy

Review; reduce the bandwidth and simplify RIN measurement to a Qsq measurement (see 68.6.7) or eliminate as appropriate. Remove 135.5.10.2.4 Square wave (quaternary) test pattern and any associated registers.
Similarly in 140.7.9.

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 of P802.3bs 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 **140** *SC* **140.7.9** *P* **310** *L* **28** # 134

Dawe, Piers Mellanox

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

The lack of consistency between the low frequency jitter specs in 802.3bs affects 802.3cd also. Here is P802.3bs D3.2 comment r02-40 for those who have not been following this issue. Depending how this inconsistency is fixed, there may be little or no explicit change in the P802.3cd draft.
Following up on P802.3bs 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 [fourth] sentence have been considered or solved by the [P802.3bs] 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.

One option in 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 **137** *SC* **137.9** *P* **241** *L* **1** # 136
Dawe, Piers Mellanox

Comment Type **TR** *Comment Status* **R** *Electrical <NSR>*

We don't yet know how to write a spec for 30 dB channels that isn't bleeding edge for ICs and/or channels. This isn't Ethernet "broad market" today, it's a specialist niche.

SuggestedRemedy

Keep working on it in Working Group ballot and if things don't improve, reduce the 30 dB objective and reduce the high loss RITT loss. It might be OK to leave the channel recommended insertion loss limit if the COM spec protects the Tx and Rx.

Response *Response Status* **U**

REJECT.

No specific change to the draft is suggested.

Cl **137** *SC* **137.9.2** *P* **241** *L* **22** # 139
Dawe, Piers Mellanox

Comment Type **TR** *Comment Status* **R** *Electrical <NSR>*

Signal-to-noise-and-distortion ratio (min) 32.5 dB is too high (even worse than 120D) - 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.

dawe_3cd_02_0717 was presented.

The comment highlights some issues in the current draft, but there was no consensus for adopting any of the proposed solutions.

The commenter is encourged to build consensus and bring a new proposal.

Cl **137** *SC* **137.9.2** *P* **241** *L* **21** # 140
Dawe, Piers Mellanox

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

Output residual ISI SNR_ISI (min) 43 dB is way 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 notes the issue (for 34.8 dB), but doesn't solve it.

SuggestedRemedy

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

Response *Response Status* **U**

REJECT.

dawe_3cd_02_0717 was presented.

The comment highlights an issue in the current draft, but there was no consensus for adopting any of the proposed solutions.

The commenter is encourged to build consensus and bring a new proposal.

See #139.

Cl **137** *SC* **137.9.3.1** *P* **241** *L* **46** # 141
Dawe, Piers Mellanox

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

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-LR 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**

ACCEPT IN PRINCIPLE.

This issue was discussed in 802.3bs and resulted in a change to the similar specification (Comment #r02-41).

In 137.9.3.1 (Receiver input return loss), append the following text to the first paragraph: "The test fixture return loss may be de-embedded from the return loss measurements."

No need to add this in 137.9.2 (Transmitter characteristics) since it refers to 120D.3.1.1, where a similar change was applied by 802.3bs (indirectly through Table 120D-1). Update exceptions if necessary.

Implement with editorial license.

Cl **136** *SC* **136.9.3** *P* **216** *L* **11** # 143

Dawe, Piers Mellanox

Comment Type **TR** *Comment Status* **R** *Electrical <NSR>*

J4, now called J4u (all but 1e-4 of the edges, or 1e-4*0.75 of the number of UI, divided between early and late, so 3.75e-5 per UI or 1.875e-5 per bit) is overkill for the spec BER of 2.4e-4, and J3u (1.875e-4 per bit) is a good match to the spec BER - just as J4u is a good match to the BER of 1e-5 for 120D. Also, not all edges cause errors. We can make the spec better (more accurate, less performance left on the table) and reduce test time. Futher, the jitter at TP2 won't be the same as at TP0a in 137.9.2 (expected to be more).

SuggestedRemedy

Change J4 to J3u. Choose the limit at TP2 considering jitter limit at TP0a and the mated compliance board crosstalk specs, among other factors.

Response *Response Status* **U**

REJECT.

The suggested remedy lacks sufficient detail required for implementation - the limits for TP2 are not included.

The commenter is encouraged to suggest and build consensus for specific limits at TP2, as well as the suggestion to change J4u to J3u.

Cl **137** *SC* **137.9.2** *P* **241** *L* **24** # 144

Dawe, Piers Mellanox

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

J4u in 120D (all but 1e-4 of the edges, or 1e-4*0.75 of the number of UI, divided between early and late, so 3.75e-5 per UI or 1.875e-5 per bit) is overkill for the spec BER of 2.4e-4, and J3u (1.875e-4 per bit) is a good match to the spec BER - just as J4u is a good match to the BER of 1e-5 for 120D. Also, not all edges cause errors. We can make the spec better (more accurate, less performance left on the table) and reduce test time.

SuggestedRemedy

Change J4 to J3u, max 0.106 UI (from eq 136-6 and 7). In Eq 136-6 and 136-7 and the NOTE, change Q4=3.8906 to Q3=3.2905, Q(Q3) = 5 x10^-4.

Response *Response Status* **U**

REJECT.

Note that the suggested change (J4u to J3u) seems to enable a shorter measurement while keeping the same sigma_RJ and A_DD for COM, by changing the conversion equations (136-6 and 136-7).

The task force discussed the suggested remedy. Since currently both clauses 136 and 137 use the same equations, there is preference to make changes to both clauses together. There is no consensus for changing just this clause.

See comment #143.

Cl **138** *SC* **138.7.1** *P* **262** *L* **17** # 147

Dawe, Piers Mellanox

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

This PMD needs more study, and knowing what TDECQ is feasible is probably the key.

SuggestedRemedy

While in WG ballot, show evidence of technical feasibility for the numbers in the spec: eyes, receiver waterfall plots, TDECQ measurements and so on. Adjust the draft as appropriate. TR because this could take a few meeting cycles.

Response *Response Status* **U**

REJECT.

[Editors note: This comment is a repeat of comment 42 against draft 1.3]

No specific changes to the draft suggested.

Task force participants are encouraged to prepare consensus presentations with proposals for specific changes to the draft if necessary.

Cl **139** *SC* **139.6.1** *P* **283** *L* **36** # 152

Dawe, Piers
Mellanox

Comment Type **TR** *Comment Status* **R** *power budget*

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. However, stassar_061417_3cd_adhoc-v2 shows plenty of receiver sensitivity margin (although not yet shown with SSPRQ). As more measurements with with new receiver designs and the new TDECQ method become available, it appears the optical power levels can be reduced and the spec as in this draft will be uneconomic (particularly 50GBASE-FR which should be low cost, low power, convenient for quad or octal packaging).

SuggestedRemedy

Bring more evidence for what optical power levels and TDECQ limits are right, including TDECQ measurements with SSPRQ, and correlation to actual receiver performance.
Based on evidence, reduce all the optical power levels for 50GBASE-FR and 50GBASE-LR by 0.5, 1 or 1.5 dB (with other adjustments for other reasons). Review the TDECQ limit.

Response *Response Status* **U**

REJECT.

The suggested remedy does not propose a specific change to the draft.