

IEEE P802.3cm D1.1 400 Gb/s over Multimode Fiber 2nd Task Force review comments

CI 138 SC 138.1 P 28 L 20 # 11
Dawe, Piers Mellanox

Comment Type E Comment Status A

"a complete Physical Layer ... as shown in Table 138-1, Table 138-2, Table 138-3, or Table 138-3a": too many tables showing almost the same information makes it hard for the reader to see what is common and what is different.

SuggestedRemedy

Combine to one table with columns for sublayer or interface, clause by PHY type, and status (required or optional). 6 columns, 16 rows including header row. Mock-up available.

Response Response Status C

ACCEPT IN PRINCIPLE.

Add a column to Table 138-3 for 400GBASE-SR8 instead of adding Table 138-3a.

CI 138 SC 138.7.1 P 270 L 28 # 9
Dawe, Piers Mellanox

Comment Type TR Comment Status R

The transition time spec is intended to protect the receiver from unreasonably slow signals, and it should be possible to use a common equalizer IC across all 50G/lane PAM4 optical PMDs without having to carry an extra burden for just one or a few PMD types.

802.3cd chose 34 ps as the slowest after a slow channel (SMF clauses) but also used 34 ps for the slowest MMF signal after a fast channel, equivalent to 36 ps (observed in 13.28125 GHz) after a slow channel - but still used 34 ps for the slowest signal in SRS.

This is inconsistent. The survey results for MMF (dawe_3cd_01b_0518 slide 8 green and slide 11 brown) show that actual transition times are significantly faster than these numbers, so there is room to correct the spec and still allow plenty of margin for measurement.

Also, it is more convenient to use the same bandwidth for transition time as for TDECQ. If someone prefers to use a different bandwidth, he can read the results across, similar to the second alternative in the remedy.

Someone using emphasis to make a slow transmitter look faster will find that it makes the transition time shorter too. If his transmitter is slow enough to worry about the transition time spec, he won't have a problem with tightening the cursor tap strength limit.

SuggestedRemedy

Either: in 138.8.7, Transmitter transition time, for 400GBASE-SR8, change 13.28125 GHz to 11.2 GHz and 26.5625 GHz to 22.4 GHz (twice) (same as 138.8.5, TDECQ).

Or:

For 400GBASE-SR8, change 34 ps to 32 ps, and:

In 138.8.10 Stressed receiver sensitivity, change "the transition time is no greater than the value specified in Table 138-8" to "the transition time is no greater than 34 ps", or add a limit could of 34 ps to Table 138-9, Receive characteristics, in the section for Conditions of stressed receiver sensitivity test.

Response Response Status C

REJECT.

This comment is similar to comment #44 against D1.0, which was rejected.

It is highly desirable to keep the per lane specifications for 400GBASE-SR8 as close as possible to the other PMDs in Clause 138 and changing the transition time for 50GBASE-SR, 100GBASE-SR2, and 200GBASE-SR4 is out of scope for this project.

IEEE P802.3cm D1.1 400 Gb/s over Multimode Fiber 2nd Task Force review comments

CI 138 SC 138.8.5 P37 L 31 # 4
Dawe, Piers Mellanox

Comment Type TR Comment Status R

The 0.1 dB allocation for modal noise and mode partition noise is too little. See dawes_3cm_adhoc_01_101118, castro_3cm_01_1118.pdf and pepeljugoski_1_1104: we need 0.1 to 0.2 dB for MN and another 0.1 dB for MPN if we keep the total penalties to 4.6 dB (which is unreasonably high). This should be done with a formula, as for 100GBASE-SR4, so as not to penalise good transmitters.

SuggestedRemedy

Insert:
For 400GBASE-SR16, Equation (138-1) is used in place of Equation (121-11).
 $R = \sqrt{\sigma^2 + \sigma^2 M^2 - M^2}$ (138-1)
where $M = 0.0065P_{ave}$
In 138.8.10 Stressed receiver sensitivity, either refer to the new Eq. 138-1 (as above) and say that:
the values of M in Equation (138-1) is set to zero
or, leave this section referring to Eq. 121-11 but to avoid confusion, add:
NOTE--The parameter M of Equation (138-1) is not used.

Response Response Status C

REJECT.
This comment is similar to comment #39 against D1.0, which was rejected.
It is highly desirable to keep the per lane specifications for 400GBASE-SR8 as close as possible to the other PMDs in Clause 138 and changing the TDECQ definition for 50GBASE-SR, 100GBASE-SR2, and 200GBASE-SR4 is out of scope for this project.
Insufficient evidence has been provided to show that the penalty is large enough to warrant a change to the link budget at this time. See the following for previous analysis:
http://www.ieee802.org/3/aq/public/nov04/pepeljugoski_1_1104.pdf.
castro_3cm_01_0119 was reviewed.
Further experimental data to support a possible change is welcomed.

CI 138 SC 138.8.5 P273 L 40 # 19
Dawe, Piers Mellanox

Comment Type E Comment Status R

Punctuation per style manual

SuggestedRemedy

Add full stops to these two sentences.

Response Response Status C

REJECT.
The text that is the subject of this comment is not present in the P802.3cm draft, so this appears to be a comment on D3.5 of P802.3cd and is therefore out of scope. If this issue is not addressed in the publication process for IEEE Std 802.3cd-2018, the commenter is recommended to pursue this matter via IEEE 802.3 Maintenance.

CI 138 SC 138.8.5 P273 L 41 # 18
Dawe, Piers Mellanox

Comment Type T Comment Status A

There is no figure showing the back reflection, nor any mention of it in the text, except "The optical channel requirements in 121.8.5.2 do not apply", and 121.8.5.2 says "The channel provides an optical return loss specified in Table 121-11. The state of polarization of the back reflection is adjusted to create the greatest RIN." I believe deleting the back reflection was not the intention.

SuggestedRemedy

After "The optical channel requirements in 121.8.5.2 do not apply", add: "Instead, the optical splitter and variable reflector are adjusted so that each transmitter is tested with the optical return loss tolerance (max) given in Table 138-8."
or add: "For 400GBASE-SR8, the optical splitter and variable reflector are adjusted so that each transmitter is tested with the optical return loss tolerance (max) given in Table 138-8."

Response Response Status C

ACCEPT IN PRINCIPLE.
Add: "For 400GBASE-SR8, the optical splitter and variable reflector are adjusted so that each transmitter is tested with an optical return loss equal to the "Optical return loss tolerance (max)" given in Table 138-8".
Changing the TDECQ measurement conditions for 50GBASE-SR, 100GBASE-SR2, and 200GBASE-SR4 is out of scope for this project. For these three PMDs, the commenter is recommended to pursue this matter via IEEE 802.3 Maintenance.

IEEE P802.3cm D1.1 400 Gb/s over Multimode Fiber 2nd Task Force review comments

CI 138 SC 138.8.5.1 P 274 L 5 # 7
Dawe, Piers Mellanox

Comment Type TR Comment Status R

Equalizing a signal after an 11.2 GHz BT4 filter with a 5-tap FFE needs at least one precursor unless the signal is carefully pre-distorted. If it is, and a fourth post-cursor is needed, the same transmitter seen after a fast channel, e.g. a short fibre, will be difficult to receive because the 5-tap FFE can't correct the fourth post-cursor and the (now -ve) first precursor at the same time. This is equivalent to trying to receive a neutral but bad signal after an 11.2 GHz filter with no precursor in the equaliser.

In practice, it seems that TDECQ uses at least one precursor for real transmitters. There is an alternative remedy: defining MMF TDECQ with fast and slow channels, in the same spirit as SMF with high and low dispersion.

SuggestedRemedy

To ensure that the 400GBASE-SR8 transmitter is not gaming the spec like this, change:

"Tap 1, tap 2, or tap 3, has the largest magnitude tap coefficient..." to "For 50GBASE-SR, 100GBASE-SR2, and 200GBASE-SR4, tap 1, tap 2, or tap 3, has the largest magnitude tap coefficient... For 400GBASE-SR8, tap 2, or tap 3, has the largest magnitude tap coefficient..."

Response Response Status C

REJECT.

This comment is similar to comment #42 against D1.0, which was rejected.

It is highly desirable to keep the per lane specifications for 400GBASE-SR8 as close as possible to the other PMDs and changing the constraint on which tap can have the largest magnitude for 50GBASE-SR, 100GBASE-SR2, and 200GBASE-SR4 is out of scope for this project.

Limiting to at most three post-cursors in the reference EQ means that the transmitted signal, when propagated through the TDECQ reference response, cannot have a significant amount of fourth post-cursor response at the receiver without suffering higher TDECQ penalty.

Insufficient evidence has been provided to justify a change.

CI 138 SC 138.8.5.1 P 274 L 5 # 12
Dawe, Piers Mellanox

Comment Type T Comment Status D

The receiver is assessed with a stressed eye generator that "should have wide and smooth frequency response, and linear phase response". So it won't need an unusual second precursor. A real transmitter, being more "causal" than neutral unless pre-distorted, will need weaker precursors than the SRS signal. Yet a transmitter is allowed to use pre-distortion to need stronger precursors, maybe of the opposite sign, than the SRS signal, and we should ensure that the transmitter combined with the range of channels can't be significantly worse than the SRS signal. For some low power equalizer architectures, precursors are much more expensive than post-cursors (sun_3cd_042518_adhoc), yet we expect MMF to be low power.

A straightforward transmitter won't benefit much from a second precursor. A clever transmitter can be set up to avoid a second precursor.

Note there is a separate comment that explains why at least one precursor is needed. Accepting both comments ("Tap 2 has") has an additional benefit of simplifying and speeding up TDECQ measurement.

SuggestedRemedy

Change:

"Tap 1, tap 2, or tap 3, has the largest magnitude tap coefficient..." to

"For 50GBASE-SR, 100GBASE-SR2, and 200GBASE-SR4, tap 1, tap 2, or tap 3, has the largest magnitude tap coefficient... For 400GBASE-SR8, tap 1 or tap 2 has the largest magnitude tap coefficient..."

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

This comment is similar to comment #47 against D1.0, which was rejected.

Allowing just one pre-cursor in the reference EQ means the transmitted signal, when propagated through a worst-case channel, cannot have a significant amount of pre-cursor response at the receiver without suffering higher TDECQ penalty. An electrical channel typically can guarantee that, however, the chromatic and modal dispersion effects of the optical channel in combination with laser performance may require the extra tap. Insufficient evidence has been provided to justify a change.

IEEE P802.3cm D1.1 400 Gb/s over Multimode Fiber 2nd Task Force review comments

CI 138 SC 138.8.5.1 P 274 L 6 # 6
Dawe, Piers Mellanox

Comment Type TR Comment Status R

All the PAM4 specs should allow the same range of over-emphasis so that a common equalizer IC can be used for all, without all SMF equalizers carrying a burden because of the MMF spec. 802.3cd chose a largest magnitude tap coefficient of at least 0.8 as a way of protecting the receiver from excessively peaky signals that abuse the receiver's dynamic range, resolution or sensitivity but don't benefit the transmitter implementer - however they did not implement it fully.

While SMF TDECQ is measured for both extremes of channel, MMF TDECQ is measured for the slow channel only. We could measure MMF TDECQ for the fast channel too. If not, we can read across to the other case we don't measure, but recognise that a signal after a slow channel will look less emphasised than what the receiver has to tolerate after a fast channel.

The reference equalizer's largest magnitude tap coefficient (0.8 for a fast channel) should be set consistently (as from the same transmitter) for the slow channel. The survey results for MMF (green points, slide 8, daw_3cd_01b_0518) are all to the right of +0.5 dB (or tap strength about 1.1); with the slower filter for 400GBASE-SR4.2 they will be further to the right (bigger again). Anyone using emphasis to make a slow transmitter look faster will start well to the right (large tap strength) and will not be concerned by this limit. This proposal is consistent with the SMF specs and still allows a strongly over-emphasised transmitter.

SuggestedRemedy

After "the largest magnitude tap coefficient, which is constrained to be at least 0.8", add "for 50GBASE-SR, 100GBASE-SR2, and 200GBASE-SR4, and at least 0.85 for 400GBASE-SR8".

Response Response Status C

REJECT.

This comment is similar to comment #41 against D1.0, which was rejected.

It is highly desirable to keep the per lane specifications for 400GBASE-SR8 as close as possible to the other PMDs and changing the constraint on the largest magnitude tap coefficient for 50GBASE-SR, 100GBASE-SR2, and 200GBASE-SR4 is out of scope for this project.

In addition, VCSEL measurements to date have shown slightly higher TDECQ penalties than SMF transmitters due to low bandwidth, and the use of peaking can help to improve yield and reduce cost especially at process, temperature, and voltage corners. Increasing the minimum coefficient of the largest magnitude tap would reduce the flexibility for the transmitter design.

CI 138 SC 138.11.4.6 P 44 L 22 # 16
Dawe, Piers Mellanox

Comment Type E Comment Status R

OC11 and OC12 Value/Comment should refer to the correct IEC 61754-7-1 interface(s), not 7-1-1

SuggestedRemedy

Compare 95.12.4.6.

Response Response Status C

REJECT.

Changing the PICS for 100GBASE-SR2 and 200GBASE-SR4 is out of scope for this project. If this issue is not addressed in the publication process for IEEE Std 802.3cd-2018, the commenter is recommended to pursue this matter via IEEE 802.3 Maintenance.

CI 150 SC 150.6 P 52 L 12 # 23
Dawe, Piers Mellanox

Comment Type E Comment Status A

Make the text easier to use by improving the layout. The example lists in the style manual start each item of a list on a new line.

Also, we don't use (i) and (ii) anywhere else and the order of entries isn't significant technically, so an unordered (e.g. dashed) list might be a better choice.

SuggestedRemedy

-- TxRx pair type TR comprises a transmitter that uses the wavelength range 844 nm to 863 nm and a receiver that uses the wavelength range 900 nm to 918 nm.

-- TxRx pair type RT comprises a transmitter that uses the wavelength range 900 nm to 918 nm and a receiver that uses the wavelength range 844 nm to 863 nm.

Response Response Status C

ACCEPT.

CI 150 SC 150.6 P 52 L 12 # 22
Dawe, Piers Mellanox

Comment Type T Comment Status A

"From Table 150-5, it is evident that" is unnecessary, and "TxRx pair" isn't defined as promised by Table 150-4 note a.

SuggestedRemedy

Change "From Table 150-5, it is evident that there are two types of TxRx pair:" to "Two TxRx pair types (combinations of Tx and Rx types that connect to a single fiber) are defined:".

Response Response Status C

ACCEPT.

IEEE P802.3cm D1.1 400 Gb/s over Multimode Fiber 2nd Task Force review comments

CI **150** SC **150.7.1** P **53** L **29** # **8**
Dawe, Piers Mellanox

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

The transition time spec is intended to protect the receiver from unreasonably slow signals, and it should be possible to use a common equalizer IC across all 50G/lane PAM4 optical PMDs without having to carry an extra burden for the bidi spec. 802.3cd chose 34 ps as the slowest after a slow channel (SMF clauses). Here, we have 34 ps for the slowest MMF signal after a fast channel, equivalent to 38 ps (observed in 13.28125 GHz) after a slow channel - but 34 ps is used for the slowest signal in SRS. This is inconsistent. The survey results for MMF (dawe_3cd_01b_0518 slide 8 green and slide 11 brown) show that actual transition times are significantly faster than these numbers, and transmitters for 150 m have to be better than those for 100 m, so there is room to correct this spec and still allow plenty of margin for measurement.

Also, it is more convenient to use the same bandwidth for transition time as for TDECQ. If someone prefers to use a different bandwidth, he can read the results across, similar to the second alternative in the remedy.

Someone using emphasis to make a slow transmitter look faster will find that it makes the transition time shorter too. If his transmitter is slow enough to worry about the transition time spec, he won't have a problem with tightening the cursor tap strength limit.

SuggestedRemedy

Either: in 150.8.7, Transmitter transition time, change 13.28125 GHz to 9 GHz and 26.5625 GHz to 18 GHz (twice) (same as 150.8.5, TDECQ).

Or:

Change 34 ps to 30 ps, and:

In 150.8.10 Stressed receiver sensitivity, change "the transition time is no greater than the value specified in Table 150-7" to "the transition time is no greater than 34 ps", or add a limit could of 34 ps to Table 150-8, Receive characteristics, in the section for Conditions of stressed receiver sensitivity test.

Response Response Status **C**

REJECT.

This comment is similar to comment #43 against D1.0, which was rejected.

VCSEL measurements shown in

http://www.ieee802.org/3/cd/public/May18/king_3cd_03_0518.pdf and

http://www.ieee802.org/3/cd/public/July18/king_3cd_02a_0718.pdf had transition times as high as 33 ps, with no receiver problems seen.

CI **150** SC **150.7.3** P **54** L **38** # **21**
Dawe, Piers Mellanox

Comment Type **E** Comment Status **A**

The reader will realise that there are two budgets, one for each wavelength range.

SuggestedRemedy

Write down which wavelength range(s) this example applies to.

Response Response Status **C**

ACCEPT IN PRINCIPLE.

Table 150-9 is applicable to both wavelength ranges.

Insert ", which apply to both wavelength ranges," after "An illustrative power budget and penalties for 400GBASE-SR4.2 channels".

CI **150** SC **150.7.3** P **54** L **45** # **17**
Dawe, Piers Mellanox

Comment Type **E** Comment Status **A**

Use the right multiplication dot, not a full stop (unless IEC do). See tables 86-9 and 86-14 (tables 59-16, 60-19, 68-9 and 75-14 use what may be a decimal point or middle dot).

SuggestedRemedy

Twice in Table 150-9, three places in Table 150-14.

Response Response Status **C**

ACCEPT IN PRINCIPLE.

Replace the five instances with a multidot.

IEEE P802.3cm D1.1 400 Gb/s over Multimode Fiber 2nd Task Force review comments

CI 150 SC 150.7.1, Table 150-7 P 53 L 15 # 24
Murty, Ramana Broadcom Inc.

Comment Type T Comment Status A

The maximum RMS spectral width (Uw) is specified as 0.6 nm for both the 844 - 863 and 900 - 918 nm channels. The lower fiber chromatic dispersion at 900 - 918 nm (compared to 844 - 863 nm) can support a larger Uw without exceeding the MPN penalty of the 844 - 863 nm channel, thereby helping reduce the manufacturing/test costs of 900 - 918 nm VCSEL arrays.

The impact of a max Uw change from 0.60 to 0.65 nm on the 900 - 918 nm channel is presented in murty_3cm_01_1118. No change to the TDECQ filter BW of 9 GHz is required.

SuggestedRemedy

Set the maximum RMS spectral width as 0.65 nm for the 900 - 918 nm channel.

Response Response Status C

ACCEPT IN PRINCIPLE.

This comment was discussed in detail during the ad-hoc teleconference held on 10 January 2019 (see murty_3cm_adhoc_01_010919).

Change the RMS spectral width (max) for the 900 nm to 918 nm wavelength range to 0.65 nm and change the TDECQ reference response bandwidth to 8.96 GHz.

CI 150 SC 150.8.5 P 56 L 41 # 20
Dawe, Piers Mellanox

Comment Type T Comment Status A

There is no figure showing the back reflection, nor any mention of it in the text, except "The optical channel requirements in 121.8.5.2 do not apply", and 121.8.5.2 says "The channel provides an optical return loss specified in Table 121-11. The state of polarization of the back reflection is adjusted to create the greatest RIN." I believe deleting the back reflection is not the intention.

SuggestedRemedy

After "The optical channel requirements in 121.8.5.2 do not apply", add: "Instead, the optical splitter and variable reflector are adjusted so that each transmitter is tested with the optical return loss tolerance (max) given in Table 150-7."

Response Response Status C

ACCEPT IN PRINCIPLE.

After "The optical channel requirements in 121.8.5.2 do not apply", add: "Instead, the optical splitter and variable reflector are adjusted so that each transmitter is tested with an optical return loss equal to the "Optical return loss tolerance (max)" given in Table 150-7."

CI 150 SC 150.8.5 P 56 L 44 # 1
Abbott, John Corning Incorporated

Comment Type TR Comment Status R

The "3dB bandwidth of 9GHz" should be specified as "not more than 9.0GHz" or something similar -- rather than designing the O/E converter and oscilloscope to be roughly 9GHz it should be designed so that it is no more than 9GHz.

SuggestedRemedy

Change "has a 3dB bandwidth of 9GHz" to "has a 3dB bandwidth of not more than 9.0GHz"

Response Response Status C

REJECT.

The current wording is consistent with other clauses. See, for example, 121.8.5.1, 122.8.5.1, 124.8.5, 138.8.5 and 140.7.5. The suggested change would mean that any -3 dB bandwidth below 9 GHz would be acceptable. The test engineer should make reasonable effort to be as near to 9 GHz as possible, which is the implication of the current wording.

Editor's note: The value of 9 GHz was modified in comment #24 to 8.96 GHz.

IEEE P802.3cm D1.1 400 Gb/s over Multimode Fiber 2nd Task Force review comments

CI **150** SC **150.8.5** P **56** L **51** # **10**
Dawe, Piers Mellanox

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

The 0.1 dB allocation for modal noise and mode partition noise is too little. See dawes_3cm_adhoc_01_101118, castro_3cm_01_1118.pdf and pepeljugoski_1_1104: we need 0.1 to 0.2 dB for MN and another 0.2 to 0.4 dB for MPN if we keep the total penalties to 4.6 dB (which is unreasonably high). This should be done with a formula, as for 100GBASE-SR4, so as not to penalise good transmitters.

This remedy keeps the 150 m reach for OM5, but the 100 m links are paying a penalty for support of 150 m. Alternatively, the reach for OM5 could be reduced to keep the same budget for OM5 as for OM4.

SuggestedRemedy

Insert:

Equation (150-1) is used in place of Equation (121-11).

$R = \sqrt{\sigma_G^2 + \sigma_S^2 - M^2}$ (150-1)

where $M = 0.0065P_{\text{ave}}$

In 150.8.10 Stressed receiver sensitivity, either refer to the new Eq. 150-1 (as above) and say that:

the values of M in Equation (150-1) is set to zero

or, leave this section referring to Eq. 121-11 but to avoid confusion, add:

NOTE--The parameter M of Equation (150-1) is not used.

Reduce the limits for TDECQ and TDECQ-10log10(Ceq), from 4.5 dB to 4.3 dB (0.2 dB lower than the SECQ values, allowing for 0.3 dB MPN penalty with associated Pcross, including the 0.1 dB already in the draft budget).

In the budget table 150-9, the power budget and allocation for penalties don't change, but the additional insertion losses for 70 m and 100 m increase by 0.1 dB each.

Response Response Status **C**

REJECT.

This comment is similar to comment #45 against D1.0, which was rejected.

Insufficient evidence has been provided to show that the penalty is large enough to warrant a change to the link budget at this time. See the following for previous analysis:

http://www.ieee802.org/3/aq/public/nov04/pepeljugoski_1_1104.pdf.

castro_3cm_01_0119 was reviewed.

Further experimental data to support a possible change is welcomed.

CI **150** SC **150.8.5.1** P **57** L **2** # **13**
Dawe, Piers Mellanox

Comment Type **T** Comment Status **D**

The receiver is assessed with a stressed eye generator that "should have wide and smooth frequency response, and linear phase response". So it won't need an unusual second precursor. A real transmitter, being more "causal" than neutral unless pre-distorted, will need weaker precursors than the SRS signal. Yet a transmitter is allowed to use pre-distortion to need stronger precursors, maybe of the opposite sign, than the SRS signal, and we should ensure that the transmitter combined with the range of channels can't be significantly worse than the SRS signal. For some low power equalizer architectures, precursors are much more expensive than post-cursors (sun_3cd_042518_adhoc), yet we expect MMF to be low power.

A straightforward transmitter won't benefit much from a second precursor. A clever transmitter can be set up to avoid a second precursor.

Note there is a separate comment that explains why at least one precursor is needed.

Accepting both comments ("Tap 2 has") has an additional benefit of simplifying and speeding up TDECQ measurement.

SuggestedRemedy

Change "Tap 1, tap 2, or tap 3, has" to "Tap 1 or tap 2 has".

Proposed Response Response Status **Z**

REJECT.

This comment was WITHDRAWN by the commenter.

This comment is similar to comment #47 against D1.0, which was rejected.

Allowing just one pre-cursor in the reference EQ means the transmitted signal, when propagated through a worst-case channel, cannot have a significant amount of pre-cursor response at the receiver without suffering higher TDECQ penalty. An electrical channel typically can guarantee that, however, the chromatic and modal dispersion effects of the optical channel in combination with laser performance may require the extra tap.

Insufficient evidence has been provided to justify a change.

IEEE P802.3cm D1.1 400 Gb/s over Multimode Fiber 2nd Task Force review comments

CI 150 SC 150.8.5.1 P57 L2 # 14
Dawe, Piers Mellanox

Comment Type TR Comment Status R

Equalizing a signal after an 9 GHz BT4 filter with a 5-tap FFE needs at least one precursor unless the signal is carefully pre-distorted. If it is, and a fourth post-cursor is needed, the same transmitter seen after a fast channel, e.g. a short fibre, will be difficult to receive because the 5-tap FFE can't correct the fourth post-cursor and the (now -ve) first precursor at the same time. This is equivalent to trying to receive a neutral but bad signal after an 9 GHz filter with no precursor in the equaliser.

In practice, it seems that TDECQ uses at least one precursor for real transmitters.

There is an alternative remedy: defining MMF TDECQ with fast and slow channels, in the same spirit as SMF with high and low dispersion.

SuggestedRemedy

To ensure that the transmitter is good enough without having to rely on a particular channel bandwidth and a fourth post-cursor, change "Tap 1, tap 2, or tap 3, has" to "Tap 2 or tap 3 has".

Response Response Status C

REJECT.

This comment is similar to comment #48 against D1.0, which was rejected.

Limiting to at most three post-cursors in the reference EQ means that the transmitted signal, when propagated through the TDECQ reference response, cannot have a significant amount of fourth post-cursor response at the receiver without suffering higher TDECQ penalty.

Insufficient evidence has been provided to justify a change.

CI 150 SC 150.8.5.1 P57 L3 # 5
Dawe, Piers Mellanox

Comment Type TR Comment Status R

All the PAM4 specs should allow the same range of over-emphasis so that a common equalizer IC can be used for all, without all their equalizers carrying a burden because of the bidi spec. 802.3cd chose a largest magnitude tap coefficient of at least 0.8 as a way of protecting the receiver from excessively peaky signals that abuse the receiver's dynamic range, resolution or sensitivity but don't benefit the transmitter implementer.

While SMF TDECQ is measured for both extremes of channel, MMF TDECQ is measured for the slow channel only. We could measure MMF TDECQ for the fast channel too. If not, we can read across to the other case we don't measure, but recognise that a signal after a slow channel will look less emphasised than what the receiver has to tolerate after a fast channel.

The reference equalizer's largest magnitude tap coefficient (0.8 for a fast channel) should be set consistently (as from the same transmitter) for the slow channel. The survey results for MMF (green points, slide 8, daw_3cd_01b_0518) are all to the right of +0.5 dB (or tap strength about 1.1); with the slower filter for 400GBASE-SR4.2 they will be further to the right (bigger again). Anyone using emphasis to make a slow transmitter look faster will start well to the right (large tap strength) and will not be concerned by this limit. This proposal is consistent with the SMF specs and still allows a strongly over-emphasised transmitter.

SuggestedRemedy

In "the largest magnitude tap coefficient, which is constrained to be at least 0.8", change 0.8 to 0.9.

Response Response Status C

REJECT.

This comment is similar to comment #40 against D1.0, which was rejected.

VCSEL measurements to date have shown slightly higher TDECQ penalties than SMF transmitters due to low bandwidth, and the use of peaking can help to improve yield and reduce cost, especially at process, temperature, and voltage corners. Increasing the minimum coefficient of the largest magnitude tap would reduce the flexibility for the transmitter design.

IEEE P802.3cm D1.1 400 Gb/s over Multimode Fiber 2nd Task Force review comments

CI 150 SC 150.10.2 P 62 L 21 # 25
Swanson, Steven Corning Incorporated

Comment Type E Comment Status A

In Table 150-14, there is a footnote "e" associated with the Effective modal bandwidth at 953 nm (min) noting that Effective modal bandwidth guidance is provided at all wavelengths in the 840 nm to 953 nm range in IEC 60793-2-10. A value of 2470 is included in the table for OM5 because 953 nm is in fact specified at 953 nm for OM5. For OM3 and OM4, it is not specified. However, it is believed that it would be useful to the reader to know what that guidance is for OM3 and OM4 so it is proposed to add the calculated guidance at 953 nm of for those fibers.

Suggested Remedy

Replace "Effective modal bandwidth guidance is provided at all wavelengths in the 840 nm to 953 nm range in IEC 60793-2-10." with "Effective modal bandwidth guidance is provided at all wavelengths in the 840 nm to 953 nm range in IEC 60793-2-10. For OM3, the guidance is 1033 MHz.km at 953 nm and for OM4, the guidance is 1459 MHz.km at 953 nm."

Response Response Status C

ACCEPT IN PRINCIPLE.
See response to comment #2.

CI 150 SC 150.10.2.1 P 62 L 8 # 2
Abbott, John Corning Incorporated

Comment Type TR Comment Status A

In Table 150-14 we should include an extra row with the informative IEC guidance for EMB at 910nm consistent with table 150-9 because 910nm is the wavelength of interest for clause 150 (not 953nm), and rather than 'not specified (e)' for OM3,OM4 EMB at 953nm we should use the informative IEC guidance along with footnote e. The reason for including the 910nm information is that this goes into the link model and TDECQ analysis.

Suggested Remedy

To table 150-14

1. add a row like the 850nm and 953nm rows for "Effective modal bandwidth at 910nm (min) (e) " [for 910nm row use (e) not (d)] .

Use the values in 150-9 in the form

OM3 1260 (informative)(e)

OM4 1980 (informative)(e)

OM5 3100 (informative) (e)

2. In the 953nm row do some thing similar for the guidance values for OM3 and OM4:

OM3 1033 (informative) (e)

OM4 1459 (informative) (e)

Response Response Status C

ACCEPT IN PRINCIPLE.

In table 150-14:

Add a row for "Effective modal bandwidth at 910 nm (min)" with "Not specified (e)" for all fiber types.

Replace footnote (e) "Effective modal bandwidth guidance is provided at all wavelengths in the 840 nm to 953 nm range in IEC 60793-2-10." with "Effective modal bandwidth guidance is provided at all wavelengths in the 840 nm to 953 nm range in IEC 60793-2-10. For OM3, the guidance is 1260 MHz.km at 910 nm and 1033 MHz.km at 953 nm. For OM4, the guidance is 1980 MHz.km at 910 nm and 1459 MHz.km at 953 nm. For OM5, the guidance is 3100 MHz.km at 910 nm."

IEEE P802.3cm D1.1 400 Gb/s over Multimode Fiber 2nd Task Force review comments

CI **150** SC **150.10.2.1** P **62** L **11** # **26**
Swanson, Steven Corning Incorporated

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

In Table 150-14, the cabled optical fiber attenuation is specified at 3.5 dB/km for OM3, OM4 and OM5. This does not agree with the IEC cable specification, IEC 60794-2-11, nor the structured cabling standards in TIA 568.3-D and ISO/IEC 11801-1. Part of the problem is that none of the standards were harmonized when OM5 was added to each.

The current draft of IEC 60794-2-11 (86A/1858/CDV) specifies 3.5 dB/km for OM3 and OM4 but 3.0 dB/km for OM5. The stated reason for the revision of the cable specification is to align 60794-2-11 with the new revision of ISO/IEC 11801-1, consistent with the Scope of -2-11. There has been discussion that the values of OM3 and OM4 should be set at the same value as OM5.

The current edition of TIA-568.3, which defines the bulk of OM5 application space, uses values of 3.0 dB/km for OM3, OM4 and OM5.

11801-1 has published with the 3.5 dB/km for OM3, 3.5 dB/km for OM4 and 3.0 dB/km for OM5.,The US National Committee to IEC recommended that ISO/IEC JTC1 SC25/WG3 initiate an amendment to IS-11801-1 modifying the cable attenuation of OM3 and OM4 from 3.5 dB/km to 3.0 dB/km consistent with that specified for OM5 which has been agreed to by IEC SC86A. There is no technical reason for OM5 to be specified at a lower attenuation than OM3 or OM4.

It is my strong desire to harmonize ISO 11801-1 with TIA-568.3-D, IEC 60794-2-11 and the IEEE 802.3cm standard. To ignore these discrepancies continues to foster misunderstanding in the market as to what specifications are correct.

SuggestedRemedy

Replace "Cabled optical fiber attenuation (max) 3.5 dB/km" with "Cabled optical fiber attenuation (max) 3.0 dB/km" in Table 150-14.

Draft a liaison letter to IEC, ISO/IEC and TIA noting this change and asking IEC and ISO/IEC to harmonize with TIA 568.3-D by specifying 3.0 dB/km for OM3, OM4 and OM5. Again, there is no technical reason for OM5 to be specified at a lower attenuation than OM3 or OM4.

Response Response Status **C**

ACCEPT IN PRINCIPLE.

To align with IEC 60794-2-11, replace "Cabled optical fiber attenuation (max) 3.5 dB/km" with "Cabled optical fiber attenuation (max) 3 dB/km" in Table 150-14.

CI **150** SC **150.10.2.1** P **62** L **14** # **27**
Swanson, Steven Corning Incorporated

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

In Table 150-14, the chromatic dispersion slope (max) (S0) is specified with one value for OM3 and OM4 and another value for OM5.

There is no technical reason for OM5 to be specified with lower chromatic dispersion than OM3 or OM4.

SuggestedRemedy

Delete the current chromatic dispersion requirement for OM3 and OM4 and specify the chromatic dispersion currently specified for OM5 for all three cabling categories, OM3, OM4 and OM5. Again, there is no technical reason for OM5 to be specified with lower chromatic dispersion than OM3 or OM4.

Draft a liaison letter to IEC, ISO/IEC and TIA noting this change and asking that the chromatic dispersion specifications be harmonized for OM3, OM4 and OM5.

Response Response Status **C**

REJECT.

The current draft correctly cites the relevant IEC specifications for the chromatic dispersion slope of OM3, OM4 and OM5 cable types. Liaisons with other standards bodies cannot be considered during comment resolution. The commenter is invited to raise the topic of liaisons once the comment resolution session is complete.

Straw poll:

Do you support the proposed change in the commenter's suggested remedy?

Y: 3

N: 12

IEEE P802.3cm D1.1 400 Gb/s over Multimode Fiber 2nd Task Force review comments

CI 150 SC 150.10.12.1 P 62 L 14 # 3
Abbott, John Corning Incorporated

Comment Type TR Comment Status R

The link model/link budget and TDECQ models for OM3 and OM4 should be done using the chromatic dispersion formula for OM5, which was determined during the TIA OM5 project by testing OM3 and OM4 fibers. The revision represents improved estimates of OM3, OM4, and OM5 fibers rather than a specific characteristic of OM5. This might also apply to clause 138 (at 850nm) but we get get started.

SuggestedRemedy

- (a) For the OM3-OM4 chromatic dispersion, but a footnote (f) which states "round-robin testing during TIA OM5 development indicates OM3, OM4, OM5 fibers all meet the OM5 specification; this guidance is informative for OM3, OM4"
- (b) generate a liaison letter to TIA or IEC as appropriate so that either informative guidance can be included in IEC60793-2-10 or the specification can be updated.

Response Response Status C

REJECT.
The current draft correctly cites the relevant IEC specifications for the chromatic dispersion of OM3, OM4 and OM5 cable types. It is not appropriate for an IEEE standard to cite results from "round-robin testing" performed in another standards body. Liaisons with other standards bodies cannot be considered during comment resolution. The commenter is invited to raise the topic of liaisons once the comment resolution session is complete.

CI 150 SC 150.11.4.6 P 68 L 18 # 15
Dawe, Piers Mellanox

Comment Type E Comment Status A

OC6 and OC7 Value/Comment should refer to the correct IEC 61754-7-1 interface(s), not 7-1-1.

SuggestedRemedy

Compare 95.12.4.6.

Response Response Status C

ACCEPT IN PRINCIPLE.
Change OC6 Value/Comment from "Per IEC 61754-7-1 interface 7-1-1" to "Per IEC 61754-7-1 interface 7-1-3 or interface 7-1-10".
Change OC7 Value/Comment from "Per IEC 61754-7-1 interface 7-1-1" to "Per IEC 61754-7-1 interface 7-1-4".