

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

CI 138 SC 138.8.5 P38 L 44 # 1
Dawe, Piers Mellanox

Comment Type TR Comment Status R

The 0.1 dB allocation for both modal noise and mode partition noise is too little. See daw_3cm_adhoc_01_101118, castro_3cm_01_1118, pepeljugoski_1_1104 and castro_3cm_01_0119: we need 0.1 to 0.2 dB for MN (castro_3cm_01_0119 says 0.23 to 0.45 dB) as well as 0.1 dB for MPN. The total penalties should be kept below 4.6 dB, which is unreasonably high already. This should be done with a formula, as for 100GBASE-SR4, so as not to penalise good transmitters.

SuggestedRemedy

Insert:
For 400GBASE-SR8, Equation (138-1) is used in place of Equation (121-11).
 $R = \sqrt{\sigma_G^2 + \sigma_S^2 - M^2}$ (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 comments #39 against D1.0 and #4 against D1.1, which were rejected.
It is highly desirable to keep the per lane specifications for 400GBASE-SR8 identical 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/ae/public/nov04/pepeljugoski_1_1104.pdf,
http://www.ieee802.org/3/cd/public/Oct18/king_3cd_01_1018.pdf
and see the following for updated analysis:
http://www.ieee802.org/3/cm/public/March19/king_3cm_01_0319.pdf.
The experimental results in
http://www.ieee802.org/3/cm/public/adhoc/sun_3cm_adhoc_01_022819.pdf are consistent with the current budget.

CI 150 SC 150.8.5.1 P58 L 29 # 2
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, recognising that a signal after the slow measurement channel looks 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 comments #40 against D1.0 and #5 against D1.1, which were 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 D1.2 400 Gb/s over Multimode Fiber 3rd Task Force review comments

CI 138 SC 138.8.5.1 P38 L45 # 3
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, recognising that a signal after the slow measurement channel looks 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, dawe_3cd_01b_0518) are all to the right of +0.5 dB (or tap strength about 1.1). 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

Change the fourth sentence in 138.8.5.1 as follows: change "Tap 1, tap 2, or tap 3, has the largest magnitude tap coefficient, which is constrained to be at least 0.8." to "...constrained to be at least 0.8 for 50GBASE-SR, 100GBASE-SR2, and 200GBASE-SR4, and at least 0.85 for 400GBASE-SR8".

Note another comment relates to the same sentence.

Response Response Status C

REJECT.

This comment is similar to comments #41 against D1.0 and #6 against D1.1, which were rejected.

It is highly desirable to keep the per lane specifications for 400GBASE-SR8 identical 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 138.8.5.1 P38 L45 # 4
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, can be difficult to receive (outside the TDECQ spec limit) because the 5-tap FFE can't correct the fourth post-cursor and the (now -ve) first precursor at the same time.

In practice, it seems that TDECQ uses at least one precursor for real MMF 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 the fourth sentence in 138.8.5.1 as follows: 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..."

Note another comment relates to the same sentence.

Response Response Status C

REJECT.

This comment is similar to comments #42 against D1.0 and #7 against D1.1, which were rejected.

It is highly desirable to keep the per lane specifications for 400GBASE-SR8 identical 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 equalizer 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.

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

CI 150 SC 150.7.1 P 54 L 30 # 5
Dawe, Piers Mellanox

Comment Type TR Comment Status R

The transition time spec is not consistent for transmit and SRS specs, and too slow for this 400GBASE-SR4.2 channel which needs faster transmitters. See slides 6 and 7 of daw_3cm_01a_0119.
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 (daw_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 comments #43 against D1.0 and #8 against D1.1, which were 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, and no evidence has been shown demonstrating a problem with the current draft.

CI 138 SC 138.7.1 P 36 L 53 # 6
Dawe, Piers Mellanox

Comment Type TR Comment Status R

The transition time spec is not consistent for transmit and SRS specs.
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 (daw_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:
In Table 138-8, Transmit characteristics, add a second Transmitter transition time row for 400GBASE-SR8, max 32 ps (not 34), 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 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 comments #44 against D1.0 and #9 against D1.1, which were rejected.
It is highly desirable to keep the per lane specifications for 400GBASE-SR8 identical 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 D1.2 400 Gb/s over Multimode Fiber 3rd Task Force review comments

CI 150 SC 150.8.5 P 58 L 22 # 7
Dawe, Piers Mellanox

Comment Type TR Comment Status R

The 0.1 dB allocation for both modal noise and mode partition noise is too little. See dawes_3cm_adhoc_01_101118, castro_3cm_01_1118, pepeljugoski_1_1104 and castro_3cm_01_0119: we need 0.1 to 0.2 dB for MN (castro_3cm_01_0119 says 0.23 to 0.45 dB) as well as 0.2 to 0.4 dB for MPN. The total penalties should be kept below 4.6 dB, which is unreasonably high already. 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, although the 100 m transmitters have to be slightly better than needed for 100 m on 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 value 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 doesn't change, the allocation for penalties for 70 m and 100 m decrease from 4.6 to 4.5 dB and the additional insertion losses for 70 m and 100 m increase by 0.1 dB to 0.3, 0.2 dB.

Response Response Status C

REJECT.
This comment is similar to comments #45 against D1.0 and #10 against D1.1, which were 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/ae/public/nov04/pepeljugoski_1_1104.pdf,
http://www.ieee802.org/3/cd/public/Oct18/king_3cd_01_1018.pdf
and see the following for updated analysis:
http://www.ieee802.org/3/cm/public/March19/king_3cm_01_0319.pdf.
The experimental results in
http://www.ieee802.org/3/cm/public/March19/sun_3cm_01a_0319.pdf
are consistent with the current budget.
Further contributions on this topic are welcomed.

CI 138 SC 138.1 P 28 L 30 # 8
Dawe, Piers Mellanox

Comment Type E Comment Status R

This table can be presented more simply, as is already done for the first and last rows, 117--RS and 78--Energy-Efficient Ethernet.

SuggestedRemedy

Combine the rows:
117--200GMII Optional Not applicable
117--400GMII Not applicable Optional
Into a single row:
117--200GMII or 400GMII Optional Optional
and so on. Notice that the columns for 200GBASE-SR4 and 400GBASE-SR8 become identical, and can be combined too.

Response Response Status C

REJECT.
The current implementation retains the base format of Clause 138, which is helpful to the reader and follows the same format as Clause 86, also a clause with multiple PMDs with identical lane specifications. The proposed change makes it less obvious which MII applies to which PMD and does not represent an improvement to the draft.

CI 150 SC 150.8.5.1 P 58 L 28 # 9
Dawe, Piers Mellanox

Comment Type TR Comment Status R

Equalizing a signal after a 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, can be difficult to receive (outside the TDECQ spec limit) because the 5-tap FFE can't correct the fourth post-cursor and the (now -ve) first precursor at the same time.
In practice, it seems that TDECQ uses at least one precursor for real MMF 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 for the intended range of channel bandwidths, 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 comments #48 against D1.0 and #14 against D1.1, which were rejected.
Limiting to at most three post-cursors in the reference equalizer 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.

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

CI 138 SC 138.10.1 P 39 L 44 # 10
Swanson, Steven Corning Incorporated

Comment Type E Comment Status A

Missing text.

SuggestedRemedy

Add text after "Change note 'a' of Table 138–14 as follows:"

Response Response Status C

ACCEPT IN PRINCIPLE.

Insert:

Replace "Only applies to 100GBASE-SR2 and 200GBASE-SR4" with "Only applies to 100GBASE-SR2, 200GBASE-SR4, and 400GBASE-SR8".

CI 150 SC 150.7 P 53 L 39 # 11
Swanson, Steven Corning Incorporated

Comment Type ER Comment Status A

Since our objective is to support 100m, the example used here should be based on 100m consistent with Clause 138.7 and not 150m.

SuggestedRemedy

Replace "A PMD that exceeds the operating range requirement while meeting all other optical specifications is considered compliant (e.g., a 400GBASE-SR4.2 PMD operating at 170 m meets the operating range requirement of 0.5 m to 150 m)." with Replace "A PMD that exceeds the operating range requirement while meeting all other optical specifications is considered compliant (e.g., a 400GBASE-SR4.2 PMD operating at 120 m meets the operating range requirement of 0.5 m to 100 m)."

Response Response Status C

ACCEPT IN PRINCIPLE.

Straw poll:

Do you support (Chicago rules):

A The commenter's Suggested Remedy?

B The commenter's Suggested Remedy modified to include the fiber type of OM4?

C No change?

Results:

A 3

B 9

C 7

Replace "A PMD that exceeds the operating range requirement while meeting all other optical specifications is considered compliant (e.g., a 400GBASE-SR4.2 PMD operating at 170 m meets the operating range requirement of 0.5 m to 150 m)." with "A PMD that exceeds the operating range requirement while meeting all other optical specifications is considered compliant (e.g., a 400GBASE-SR4.2 PMD operating at 120 m on OM4 meets the operating range requirement of 0.5 m to 100 m on OM4)."

CI 150 SC 150.7 P 56 L 22 # 12
Swanson, Steven Corning Incorporated

Comment Type ER Comment Status A

The cabled optical fiber attenuation in footnote b is incorrectly stated.

SuggestedRemedy

Replace "bThe channel insertion loss is calculated using the maximum distance specified in Table 150–6 and cabled optical fiber attenuation of 3.5 dB/km at 850 nm plus an allocation for connection and splice loss given in 150.10.2.2.1." with "bThe channel insertion loss is calculated using the maximum distance specified in Table 150–6 and cabled optical fiber attenuation of 3 dB/km at 850 nm plus an allocation for connection and splice loss given in 150.10.2.2.1."

Response Response Status C

ACCEPT.

CI 150 SC 150.7 P 56 L 17 # 13
Swanson, Steven Corning Incorporated

Comment Type T Comment Status A

In Table 150-9, the channel insertion loss values must change to reflect the new cabled optical fiber attenuation value of 3 dB/km.

SuggestedRemedy

For OM3, replace "1.8" with "1.7."

For OM4, replace "1.9" with "1.8."

Response Response Status C

ACCEPT IN PRINCIPLE.

In Table 150-9:

In the row for parameter "Channel insertion loss", replace "1.8" with "1.7" and replace "1.9" with "1.8".

In the row for parameter "Additional insertion loss allowed", replace "0.2" with "0.3" and replace "0.1" with "0.2".

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

<i>CI</i> 150	<i>SC</i> 150.10.1	<i>P</i> 63	<i>L</i> 33	<i>#</i> 14
Swanson, Steven		Corning Incorporated		

Comment Type **T** *Comment Status* **A**

In Table 150-13, the channel insertion loss values must change to reflect the new cabled optical fiber attenuation value of 3 dB/km.

SuggestedRemedy

or OM3, replace "1.8" with "1.7."

For OM4, replace "1.9" with "1.8."

Response *Response Status* **C**

ACCEPT.