

IEEE P802.3 (IEEE 802.3cj) D3.0 Maintenance #12 (Revision) Initial Sponsor ballot comments

Cl 1 SC 1.4.281 P92 L4 # i-41
Nikolich, Paul INDEPENDENT

Comment Type TR Comment Status A

The current definition of 'lane' requires improvement.
Current definition: 1.4.281 lane: A bundle of signals that constitutes a logical subset of a point-to-point interconnect. A lane contains enough signals to communicate a quantum of data and/or control information between the two endpoints.

For example "bundle" is defined as a "group of signals", which is duplicated in "bundle of signals" above.
Per the definition of "bundle", it should be "A bundle that constitutes..."

Where is "quantum of data" defined? I couldn't find it.

Where is "endpoint" defined?

Unfortunately I don't have a good alternative definition.

SuggestedRemedy

Look through the draft and identify the various ways "lane" is used, then develop an appropriate single definition. If a single definition is not feasible, perhaps more than one definition is needed.

Response Response Status U

ACCEPT IN PRINCIPLE.
Replace the definition of "lane" with the following.

"A logical subset of the data and control information transmitted from one sublayer (e.g., PCS, PMA) to an adjacent sublayer across the inter-sublayer interface or from one PHY to another across the transmission medium (e.g. optical fiber, optical wavelength, wire pair). Lanes are transmitted in parallel and combine to deliver the full set of data and control information across the interface."

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

Cl 120D SC 120D.3.1 P366 L14 # r01-23
 Rysin, Alexander Mellanox Technologies

changes were adopted. Many of these new values are to include notes that state values are "to be confirmed". Based on this, it appears that the new specification is not mature enough to incorporate into the draft at this time.

Comment Type TR Comment Status R

Transmitter output residual ISI SNR_ISI (min) 34.8 dB (Clause 120D) is too high - can barely measure the IC through the test fixture. The warning NOTE in 120D.3.1.7 shows the issue, but doesn't solve it. 802.3cd D2.0 comment 140, D2.1 comment 49, D2.2 comment 22. Since both SNR_ISI and Effective Return Loss (ERL) represent uncompensated reflections from the transmitter and the test fixtures, measurements of ERL can replace SNR_ISI.

Also, frequency domain return loss mask does not truly represent digital signaling at a given bit error ratio. There is no real proof that violating return loss masks is directly tied to failures and a number of false negatives have been shown. 802.3cd D2.0 comment 141, D2.1 comments 26, 27 and 28, D3.0 comment 98.

The specifications in Annex 120D (CDAUI-8 chip-to-chip) were approved as part of IEEE Std 802.3bs-2017. The concept of ERL was introduced in the IEEE P802.3cd amendment for the backplane and copper cable interfaces. In these cases the link budget margins are considerably lower making the imprecision of return loss masks more impactful and SNR_ISI requirements more demanding (34.8 dB for Annex 120D vs. the a placeholder value of 43 dB for IEEE P802.3cd/D3.1 Clause 137). It has not been established that the use of ERL for chip-to-chip (or chip-to-module interfaces, which are not mentioned in the comment despite their use of return loss masks) provide benefits that outweigh the risk of imposing new requirements on devices compliant to the original standard. It has been shown that devices that pass current return loss requirements do not necessarily pass the proposed ERL requirements.

SuggestedRemedy

- * Add an Annex describing ERL computation method and parameters. The Annex can be copied from 93A-5 in 802.3cd D3.1.
- * Add a parameter Table, copying Table 137-5 for 802.3cd D3.1.
- * Add a description of the ERL computation and parameters as follows:

Effective return loss (ERL) of the transmitter at TP0a is computed using the procedure in Annex (new) with the values in Table TBD. Parameters that do not appear in Table TBD take values from Table 120D-8. The value of Tfx is twice the delay from TP0 to TP0a. Nbx is set to the value of Nb in Table 120D-8. ERL shall be at least 16.1 dB.

- * Add a reference in 120D.3.2 to Annex (new) and to Table TBD for a description of the ERL computation and parameters as follows:

Effective return loss (ERL) of the receiver computed using the procedure in Annex (new) with the values in Table TBD. Parameters that do not appear in Table TBD take values from Table 120D-8. The value of Tfx is twice the delay from TP5a to TP5. Nbx is set to the value of Nb in Table 120D-8. ERL shall be at least 16.1 dB.

- * Remove the requirement for Differential return loss in Table 120D-1.
- * Add a requirement for Effective Return Loss (ERL) to be greater than 16.1 dB in Table 120D-1
- * Remove the requirement for Differential input return loss in Table 120D-5
- * Add a requirement for Effective Return Loss (ERL) to be greater than 16.1 dB in Table 120D-5
- * Remove reference to Transmitter Output residual ISI SNR_ISI(min) in Table 120D-1.

Response Response Status U

REJECT.
 Annex 93A.5 and Effective Return Loss (ERL) specifications were first introduced in IEEE P802.3cd/D3.1 (January 2018). During the March 2018 meeting of IEEE P802.3cd Task Force, numerous changes were proposed to the ERL parameters and ERL requirements (see <http://www.ieee802.org/3/cd/public/Mar18/dudek_3cd_02_0318.pdf>) and many

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

Cl 1 SC 1.4.281 P92 L4 # r01-24
 Nikolich, Paul INDEPENDENT

Comment Type GR Comment Status R

*** Comment submitted with the file 96131200003-20180124_163855.jpg attached ***

The proposed resolution is an improvement, but unacceptable:
 "A logical subset of the data and control information transmitted from one sublayer (e.g., PCS, PMA) to an adjacent sublayer across the inter-sublayer interface or from one PHY to another across the transmission medium (e.g. optical fiber, optical wavelength, wire pair). Lanes are transmitted in parallel and combine to deliver the full set of data and control information across the interface."

My comments:

a) The proposed text doesn't quite capture the concept of arbitrary recombination of the smallest subsets into larger subsets (which are not identical to the originating superset. Perhaps adding the word 'superset' will help as follows:
 "A logical subset of a superset of data and control information transmitted from one sublayer (e.g., PCS, PMA)..."

b) The text should be accompanied by an illustrative figure similar to the one you drew for me in Geneva. See attached file.

SuggestedRemedy

See suggestion in above comment.

Response Response Status U

REJECT.
 The definition is specific to the transmission of control and data information from "one sublayer (e.g., PCS, PMA) to an adjacent sublayer across the inter-sublayer interface or from one PHY to another across the transmission medium." While the number of output lanes may be changed from the number of input lanes by a sublayer (e.g., it may aggregate subsets into larger subsets or divide subsets into smaller subsets), this is a function of the sublayer and not inherent to the definition of a lane. The definition of lane applies to the input of the sublayer and the output of the sublayer while the functions within the sublayer are beyond the scope of this definition. The proposed addition of the term "superset" does not appear to improve the definition in this context.

The inclusion of a figure with a definition is unprecedented in IEEE Std 802.3 (although it is acknowledged there is an example of this in IEEE Std 802.16-2017 and other standards under IEEE-SA). Regardless, it is believed that the definition is clear as it is written and does not require a figure.

Cl 121 SC 121.8.5.3 P132 L1 # r01-35
 Dawe, Piers J G Mellanox Technologies

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, such as high peak power, high crest factor, or a need to remove emphasis from the signal, contrary to what equalizers are primarily intended to do. Note the receiver is tested for a very slow signal only, not for any of these abusive signals. This is an issue for all the PAM4 optical PMDs, although it may be worse for MMF because of the high TDECQ limit and because the signal is measured in a particularly low bandwidth. This comment updates 802.3cd D3.1 comment 71. With luck it will be possible to follow 802.3cd's action on this topic.

SuggestedRemedy

1. To screen for noisy or distorted signals with heavy emphasis:
 - 1a. Define a metric similar to TDECQ but with C_{eq} held at 1, that measures how closed the eye after the reference equalizer is. Set a limit for it.
 - or:
 - 1b. Define $TDECQ_{rms} = 10 \cdot \log_{10}(A_{RMS}/(s^3 \cdot Q_t \cdot R))$ where A_{RMS} is the standard deviation of the measured signal after the 13.28125 GHz or 11.2 GHz filter response (before the FFE), Q_t and R are as already in Eq 212-12. s is the standard deviation of a fast clean signal with $OMA=2$ and without emphasis, observed through the filter response (0.6254 for 13.28125 GHz, 0.6006 for 11.2 GHz). Either, set limit for $TDECQ_{rms}$ according to what level of dirty-but-emphasised signal we decide is acceptable, add max $TDECQ_{rms}$ row to each transmitter table. Or, if the same relative limit is acceptable for all PAM4 optical PMDs, the limit could be here in the TDECQ procedure. E.g. make the $TDECQ_{rms}$ limit the same as the TDECQ limit, say here that both TDECQ and $TDECQ_{rms}$ must meet the TDECQ spec.
2. To protect the receiver from having to "invert" heavily over-emphasised signals, set a minimum cursor weight, 0.9. Similarly in clauses 122, 124. To protect the equalizer from having to support unnecessary settings for waveforms that can't or shouldn't ever happen, constrain the cursor position - see other comments.

Response Response Status U

REJECT.
 There are no PAM4 optical PMDs (that would use the TDECQ test) over MMF in the draft. "Eq 212-12" in the suggested remedy should be "Eq 121-12".

The need for additional transmitter specs for the SMF PMDs has not been established, and insufficient evidence has been provided that the proposed alternative remedies fix the claimed problem.

To date no contribution has been made that demonstrates the problem (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that one of the proposed additional requirements prevents this issue from occurring.

A similar proposal to create a $TDECQ_{rms}$ spec was suggested in comments i-140 against P802.3bs D3.0, r02-35 against P802.3bs D3.2 and r03-27 against P802.3bs D3.3 which were similarly rejected.

A peak power spec has not been shown to be necessary, and a definition and value has

IEEE P802.3 (IEEE 802.3cj) D3.1 Maintenance #12 (Revision) 1st Sponsor recirculation ballot comments

not been provided.
 A crest factor limit has not been shown to be necessary, and a definition and value has not been provided.
 The need for a limit to cursor weight has not been established.
 Constraints have been placed on the cursor position due to the changes made in response to comment r01-17.

[Editor's note added after comment resolution completed.
 The suggested remedy for comment r01-17 was accepted. The suggested remedy is:
 "Make changes to:
 121.8.5.4 for 200GBASE-DR4 (and by reference 400GBASE-DR4)
 122.8.5.4 for 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8
 equivalent to the changes made in P802.3cd 139.7.5.4 between D3.0 and D3.1:
 Add the text:
 ""A functional model of the reference equalizer is shown in Figure 12x-y.""
 ""Tap 1, tap 2, or tap 3, has the largest magnitude tap coefficient.""
 and a figure in each case."
]

Cl 121 SC 121.8.5.3 P134 L45 # r01-36
 Dawe, Piers J G Mellanox Technologies

Comment Type TR Comment Status R

The TDECQ method allows signals that are slower than 100GBASE-LR4, probably slower than the original T/2-spaced TDECQ allowed, and slower than anticipated. If this hole is not plugged, product receivers will have to provide more tap strength than is needed to receive the range of reasonable signals, degrading their cost/power/performance trade-off. This issue became more clear after the 802.3cd comments were written, but with luck, 802.3cd will consider the matter as part of their TDECQ comment resolution anyway.

SuggestedRemedy

Set a maximum cursor strength limit, which might be around 1.3.
 Similarly in clauses 122, 124.

Response Response Status U

REJECT.
 The need for a limit to cursor weight has not been established (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that the proposed limit of 1.3 removes the demonstrated issue while not disallowing "reasonable" transmitters.

Cl 121 SC 121.8.5.4 P135 L18 # r01-37
 Dawe, Piers J G Mellanox Technologies

Comment Type TR Comment Status A

802.3cd has adopted cursor position rules that should apply here too. Further, the rules should be tightened (see http://ieeee802.org/3/cd/public/Mar18/dawe_3cd_01_0318.pdf).

SuggestedRemedy

Copy the new material from 138.8.5.1, including Figure 138-3, TDECQ reference equalizer functional model. However, (802.3cd comment 76, instead of "Tap 1, tap 2, or tap 3, has the largest magnitude tap coefficient", use "Tap 1 or tap 2 has the largest magnitude tap coefficient".
 Specifications work at different levels: functional, logic/digital, analog (electrical or optical), and "Functional" is the highest/most abstract, while this FFE diagram is part of the specification of an analog quantity (more at 802.3cd comment 72). So instead of "symbol period. A functional model of the reference equalizer is shown in Figure 138-3" use "symbol period, as shown in Figure 138-3", and in the figure title, instead of "TDECQ reference equalizer functional model" use "TDECQ reference equalizer".

Response Response Status U

ACCEPT IN PRINCIPLE.
 See response to comment r01-17 which applies the restriction that the main tap has to be tap1, tap2, or tap3.
 It has not been demonstrated that disallowing tap 3 as having the largest magnitude tap coefficient is an improvement to the draft. (Indeed, several of the contributed measurements have shown tap3 as the largest magnitude tap coefficient for the optimum tap setting.)
 Regarding the "functional model" description, the text and figure follow the precedent set in IEEE Std 802.3bs-2017 Annex 120D for an equivalent type of equalizer.

[Editor's note added after comment resolution completed.
 The suggested remedy for comment r01-17 was accepted. The suggested remedy is:
 "Make changes to:
 121.8.5.4 for 200GBASE-DR4 (and by reference 400GBASE-DR4)
 122.8.5.4 for 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8
 equivalent to the changes made in P802.3cd 139.7.5.4 between D3.0 and D3.1:
 Add the text:
 ""A functional model of the reference equalizer is shown in Figure 12x-y.""
 ""Tap 1, tap 2, or tap 3, has the largest magnitude tap coefficient.""
 and a figure in each case."
]