

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 30 SC 30.5.1.1.17 P 39 L 46 # 19
 Ran, Adee Intel

Comment Type T Comment Status X

(comment is about text that has not changed from D2.1)

The maximum increment rates stated here seem to be incorrect.

The maximum occurs when every FEC codeword is corrected (which is close to the expectation with an uncorrelated BER close to 2e-4). For 200G/400G the codeword size is 5440 bits, and the durations are 2720/n UI, so 1360 UI = 51.2 ns for 200G and 680 UI = 25.6 ns for 400G.

Since there are two FEC instances, the maximum rate per instance corresponds to two codewords.

Accordingly the maximum rates are slightly below 10 million per second for 200G and 20 million per second for 400G.

Also applies to 30.5.1.1.18 (uncorrectable codewords) where the maximum rate is when all codewords are uncorrectable, e.g. when there is no link partner.

SuggestedRemedy

In 30.5.1.1.17, change "40 000 000" to "10 000 000" and "80 000 000" to "20 000 000".

Apply same changes in 30.5.1.1.18.

Proposed Response Response Status O

CI 30 SC 30.5.1.1.18 P 40 L 30 # 20
 Ran, Adee Intel

Comment Type T Comment Status X

(comment is about text that has not changed from D2.1)

"Each element of this array contains a count of corrected FEC blocks" seems to be a copy/paste error. aFECUncorrectableBlocks should count uncorrectable rather than corrected blocks

(The error appears in the base document, however the paragraph is amended so may be in scope of the project)

SuggestedRemedy

Change "corrected" to "uncorrectable".

Proposed Response Response Status O

CI 45 SC 45.2.3.13 P 72 L 13 # 21
 Ran, Adee Intel

Comment Type E Comment Status X

(comment is about text that has not changed from D2.1)

802.3bq has changed 10GBASE-T to MultiGBASE-T.

SuggestedRemedy

Change "10GBASE-T" to "MultiGBASE-T" in the following

- titles of 45.2.3.13, 45.2.3.13.1
- body of 45.2.3.13.1
- Table 119-5, first row (twice)

Proposed Response Response Status O

CI 78 SC 78.1 P 102 L 9 # 22
 Ran, Adee Intel

Comment Type T Comment Status X

(comment is about text that has not changed from D2.1)

The list of supported PHY types in should not include the new AUIs, since they are transparent to LPI (unlike 25GAUI, XLAUI and CAUI-n, which have special behavior in eep-sleep LPI). PMDs which are transparent to LPI (like all optical PMDs) are not listed.

However, the list should include the 200GXS and 400GXS, since they do have special requirements for relaying LPI signaling, which do apply in fast wake (like XGXS).

SuggestedRemedy

- Change "the 200GAUI-8 or 200GAUI-4" to "the 200GXS".
- Change "the 400GAUI-16 or 400GAUI-8" to "the 400GXS".

Proposed Response Response Status O

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 78 SC 78.5 P 103 L 4 # 24
 Ran, Adee Intel

Comment Type T Comment Status X

(comment is about text that has not changed from D2.1)

The LPI timing parameters for 200GXS and 400GXS are not listed.

Table 78-4 should include rows for 200GXS and 400GXS, similar to the row for XGXS in the base document.

Since these sublayers practically form a full 200GBASE-R/400GBASE-R link, it makes sense to assume that their timing parameters are the same as the corresponding PHYs. Any new PHY that includes 200GXS/400GXS would need to list its delay separately.

SuggestedRemedy

Add two rows for 200GXS and 400GXS with the same values as the existing rows for 200GBASE-R fast wake and 400GBASE-R fast wake.

Proposed Response Response Status O

Cl 78 SC 78.5.1 P 103 L 17 # 25
 Ran, Adee Intel

Comment Type T Comment Status X

(comment is about text that has not changed from D2.1)

78.5.1 (not included in the draft) is titled "10 Gb/s PHY extension using XGXS". Its content is relevant for 200GXS and 400GXS too.

SuggestedRemedy

Bring 78.5.1 into the draft.

Change its title from "10 Gb/s PHY extension using XGXS" to "PHY extension using extender sublayers".

Insert the following new paragraph at the end of 78.5.1:
 "The 200GXS/400GXS can be inserted between the RS and a 200 Gb/s or 400GXS PHY, respectively, to transparently extend the physical reach of the 200GMII/400GMII. The LPI signaling can operate through the 200GXS/400GXS with no change to the PHY timing parameters described in Table 78-4 or the operation of the Data Link Layer Capabilities negotiation described in 78.4."

Proposed Response Response Status O

Cl 78 SC 78.5.2 P 103 L 20 # 23
 Ran, Adee Intel

Comment Type T Comment Status X

(comment is about text that has not changed from D2.1)

There is no need to list the new AUIs here since they are transparent to LPI (unlike 25GAUI, XLAUI and CAUI-n). PMDs which are transparent to LPI (like all optical PMDs) are not listed.

SuggestedRemedy

Remove 78.5.2 and the editorial instructions to change it.

Proposed Response Response Status O

Cl 120 SC 120.5.11.2.3 P 200 L 54 # 43
 Dudek, Mike Cavium

Comment Type E Comment Status X

This paragraph duplicates the beginning of the next paragraph and is redundant.

SuggestedRemedy

Delete it.

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

Cl 120 SC 120.5.11.2.3 P 200 L 54 # 37
 Brown, Alan ADTRAN, Inc.

Comment Type E Comment Status X

A period is needed to close the sentence "Each section of PRBS31 is generated as if produced by the shift register implementation".

SuggestedRemedy

Add the period.

Proposed Response Response Status O

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 120 SC 120.5.11.2.3 P 200 L 54 # 1
 Ran, Adeee Intel

Comment Type E Comment Status X

The sentence "Each section of PRBS31 is generated as if produced by the shift register implementation" is cut prematurely, and is then repeated in the next paragraph.

Also, this is a list of rules, so it should be formatted accordingly.

SuggestedRemedy

Delete the quoted sentence.

Use dashed list format for the paragraphs from "Bit sequence A..." until "The repeating SSPRQ pattern..." (inclusive).

Proposed Response Response Status

CI 120 SC 120.5.11.2.4 P 201 L 46 # 4
 Ran, Adeee Intel

Comment Type T Comment Status X

The problem with square wave is related to the CDR functionality of the PMA at the receive side of the 200GAUI-4 or 400GAUI-8 (whether or not it is adjacent to the PMD).

There is nothing anywhere else in this clause that states that the PMA _receiver_ expects a CDR-friendly pattern and may not work well with a square wave (or, for that matter, with SSPR). This can occur even if there is no PMD in the system under test.

There is actually no specified receiver behavior for patterns other than PCS data and PRBS31/PRBS31Q. SSPR and square wave are used for transmitter testing but we do not expect receivers to work well with them. But as the text stands there is no special treatment for these patterns - the AUJ annexes requirements apply just the same, with their pater-agnostic BER statements. This is an overkill and probably not what we intend.

The text should state clearly that the receiver is not expected to cope with this kind of patterns.

This subclause deals with a transmitted test pattern, so it seems like a bad place to put such a statement. A better place to do that would be 120.5.1 which is titled "Per input-lane clock and data recovery". Alternatively it can be added to the BER subclause in each of the AUJ annexes.

SuggestedRemedy

Add a paragraph in 120.5.1:

"Clock and data recovery specifications apply for receiving PCS encoded data or PRBS31/PRBS31Q test patterns. Feeding other patterns (such as square wave or SSPR/SSPRQ) into a PMA through a physically instantiated interface may yield unexpected results".

Proposed Response Response Status

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 120 SC 120.5.11.2.4 P 201 L 46 # 2
 Ran, Adee Intel

Comment Type T Comment Status X

The "note that" sentence is a part of normative text (see style manual 16.1), but it is not clear how it specifies anything: "may" means "is allowed to", but this clause specifies the PMA and the PMA has no special "allowance" (in the current text; see another comment) for not forwarding data correctly when the data is a square wave.

Also this "note" is hard to find in the middle of this long run-on paragraph (inserting it made it even longer...)

It would be better to have this text stand out as an informative note (in a separate paragraph) after describing what the feature actually is.

SuggestedRemedy

Delete the sentence "Note that if a square wave is transmitted through a 200GAUI-4 or 400GAUI-8 it may not be correctly forwarded to the output of the PMD sublayer", and instead insert a paragraph break.

Add an informative note paragraph at the end of this subclause (after the "When enabled" paragraph):
 "NOTE--A square wave is not guaranteed to appear correctly on the output of a 200GAUI-4 or 400GAUI-8 PMD."

Proposed Response Response Status O

CI 120 SC 120.5.11.2.4 P 201 L 46 # 3
 Ran, Adee Intel

Comment Type T Comment Status X

Is the CDR problem specific to PAM4 only? Is a square wave guaranteed to appear correctly when the interface uses NRZ signaling (200GAUI-8 or 400GAUI-16)?

This is not a data signal anyway, so there is no need to assume it works like data in NRZ.

SuggestedRemedy

Change "200GAUI-4 or 400GAUI-8" to "200GAUI-n or 400GAUI-n".

Proposed Response Response Status O

CI 120D SC 120D.3.1.1 P 352 L 28 # 29
 Ran, Adee Intel

Comment Type TR Comment Status X

There seems to be a mismatch between the transmitter jitter specifications and the A_DD parameter.

Looking at the precedence in 83D:

- The maximum effective DJ allowance for the transmitter is 0.1 UI PtP (Table 83D-1)
- The channel is specified with COM parameter A_DD=0.05 (Table 83D-6), corresponding to 0.1 UI PtP.

In the current annex:

- Transmitter DJ is not specified directly, but using equations 120D-9 and 120D-10 with the maximum specified J4 (0.118 UI) and JRMS (0.019 UI) yields A_DD=0.015 and sigma_RJ=0.011
- The channel is specified with COM parameter A_DD=0.02 and sigma_RJ=0.01.

If the equations are correct, this means the channel specification assumes a significantly worse transmitter than what is actually allowed, and the transmitter specification may be relaxed.

SuggestedRemedy

Change specification to values that would yield the same values of A_DD and sigma_RJ from equations 120D-9 and 120D-10 as the values in table 120D-8. (I could not find such values)

Proposed Response Response Status O

CI 120D SC 120D.3.1.1 P 352 L 43 # 39
 Dawe, Piers Mellanox

Comment Type T Comment Status X

We don't need each of the 12 measurements to be within the J4 or Jrms limits; we just need the aggregate to do so because in COM we make all the edges have the jitter. Recognising this we can improve measurement time and cost 12-fold, which we need to do with multiple emphasis settings and up to over a hundred lanes on each IC.

SuggestedRemedy

After the first sentence, insert "Align the means of each histogram then add them together to obtain the the jitter probability density distribution."

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 120D SC 120D.3.1.1 P 352 L 43 # 38
 Dawe, Piers Mellanox

Comment Type T Comment Status X

Following 52.9.9.3 and 86.8.3.3.1, this says "Each histogram should include at least 10⁶ hits." Recommending such a detail (at least 10000 hits then) was OK for a single-lane stressed eye calibration in 52.9.9.3, and not right for a multi-lane yes/no product spec in 86.8.3.3.1, J2 Jitter, where the trade-off between margin and accuracy applies. But 10,000 hits x 4 or 10 lanes wasn't terrible, and we did not make the same mistake for J9. Here, we have a million hits, times multiple emphasis settings, times up to over a hundred lanes on each IC. It's far too much, and not necessary.

SuggestedRemedy

Delete "Each histogram should include at least 10⁶ hits". I considered adding words such as "to obtain an accurate measurement...", but a test engineer can work out what he needs for his own circumstances and should be free to do it.

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

CI 120D SC 120D.3.1.1 P 352 L 46 # 11
 Ran, Adeee Intel

Comment Type T Comment Status X

JRMS and J4 are defined twice in these two paragraphs. The second definition seems to be based on the first.

I assume that the first paragraph defines these values for a specific transition.

SuggestedRemedy

Change FROM

"J4 is defined as the time interval that includes all but 10–4 of the jitter probability density distribution, which is the time interval from the 0.005th to the 99.995th percentile of the jitter histogram. JRMS is defined as the RMS value of the jitter distribution."

TO

"J4(i) is defined as the time interval that includes all but 10–4 of the jitter probability density distribution for transition type i (i=1 to 12), which is the time interval from the 0.005th to the 99.995th percentile of the jitter histogram. JRMS(i) is defined as the RMS value of the jitter distribution for transition type i (i=1 to 12)."

change FROM

"J4 is the maximum of the 12 measurements. JRMS is the root mean square of the 12 measurements"

TO

"J4 is the maximum of the 12 measurements J4(i). JRMS is the root mean square of the 12 measurements JRMS(i)."

Proposed Response Response Status O

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 120D SC 120D.3.1.1 P 352 L 50 # 13
 Ran, Adee Intel

Comment Type T Comment Status X

JRMS is defined as the RMS or RMS measurements.

Assuming all measurements have the same number of samples, the RMS of RMS of equally-sized sets is the same as the std of the total population (which would include 12e6 samples).

RMS'ing 12 times the number of samples is expected to yield almost the same result, but take much longer time. If the measurement is separated to transitions, each measurement can be made shorter by a factor of 10 and we'll still have a higher confidence level than with the previous method. As suggested in another comment, this can be justified for J4 measurement too.

SuggestedRemedy

Change "Each histogram should include at least 10^6 hits" to "Each histogram should include at least 10^5 hits".

Proposed Response Response Status O

CI 120D SC 120D.3.1.1 P 352 L 50 # 12
 Ran, Adee Intel

Comment Type TR Comment Status X

J4 is defined as the maximum of the measurements. If one of the transitions has significantly higher jitter than the rest, this will cause exaggerated estimation of the effect of jitter on the detector error performance.

High jitter in a specific transition is practically diluted by the other transitions; the maximum-jitter transition only occurs in 1/12 of the transitions, or 1/16 of the UIs. With the new definition, the 1e-4 of a specific transition correspond to only 8.3e-6 of the total transition population or 6.25e-6 of the decisions.

Note that in COM the jitter is modeled as dual-dirac (using A_DD calculated from J4, see 120D-9). This may be quite far from the actual jitter distribution if most transitions have lower jitter. Using the average J4 across transitions (previous definition) as A_DD would be more accurate.

SuggestedRemedy

Proposed alternatives:

1. Specify that J4(i) of each transition is defined as the value that includes all but 12e-4 of the samples of transition i (so that it translates to 1e-4 of the total population).
2. Specify that J4 is the average of the highest n values of J4(i) (across the 12 transitions). This will reduce the effect of one transition. I suggest n=6.
3. Revert to the previous measurement method which does not measure each transition separately; that method inherently creates some averaging between transitions and prevents domination of the worst one.

Proposed Response Response Status O

CI 120D SC 120D.3.1.2 P 353 L 33 # 46
 Dudek, Mike Cavium

Comment Type E Comment Status X

The second sentence in the paragraph already says that the mean signal levels are defined in 120D.3.1.2.1. There is no need to repeat this.

SuggestedRemedy

Delete "The calculation of the mean signal levels is defined in 120D.3.1.2.1."

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 120D SC 120D.3.1.4 P 354 L 34 # 26
Ran, Adee Intel

Comment Type TR Comment Status X

The current steady-state voltage specification uses p(k), which is determined from the linear fit procedure, which is calculated separately for each equalizer setting. This specification reads as if it applies in all equalization settings.

It is not likely that the specified minimum steady-state voltage in Table 120D-1 (0.4 V) will be met in all equalization settings, and this is not the intent. Steady-state voltage should be specified only in unequalized state, to be consistent with precedent electrical clauses and AUI specifications.

SuggestedRemedy

Change FROM

"The linear fit pulse, p(k), is determined according to the linear fit procedure in 120D.3.1.3" TO

"The linear fit pulse, p(k), is determined according to the linear fit procedure in 120D.3.1.3 with Local_eq_cm1 and Local_eq_c1 set to 0".

Proposed Response Response Status O

Cl 120D SC 120D.3.1.4 P 354 L 34 # 27
Ran, Adee Intel

Comment Type E Comment Status X

Parentheses and numbers should not be italicised. Also, mutliplication should be denoted by a cross character.

SuggestedRemedy

Change numbers and parentheses to upright font.

Add cross character (0xD7) between "M" and "Nv".

Proposed Response Response Status O

Cl 120D SC 120D.3.1.5 P 354 L 44 # 28
Ran, Adee Intel

Comment Type E Comment Status X

(comment is about text that has not changed from D2.1)

Incorrect cross reference: 120D.3.1.2 describes transmitter linearity. The linear fit method is a different thing, and is described in 120D.3.1.3.

SuggestedRemedy

Change cross reference from 120D.3.1.2 to 120D.3.1.3.

Proposed Response Response Status O

Cl 120D SC 120D.3.1.7 P 356 L 23 # 47
Dudek, Mike Cavium

Comment Type T Comment Status X

Table 120D-7 is the jitter amplitudes and frequencies for the stressed receiver test and is not relevant.

SuggestedRemedy

Change "Table 120D-7" to Table 120D-8, on line 23 and on line 36.

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

Cl 120D SC 120D.3.1.7 P 356 L 23 # 33
Hidaka, Yasuo Fujitsu Lab. of Americ

Comment Type T Comment Status X

Optimization of two parameters of the second-order CTLE as described in 93A.1.4.3 with parameters in Table 120D-8 is not required for the loss of package and test fixture. The CTLE defined for chip-to-module interface in 120E.3.1.7 should be sufficient.

SuggestedRemedy

Change "SNR_ISI is defined by Equation (120D-8) computed from p_max and ISI_cursors after these have been re-calculated with the continuous time filter described in 93A.1.4.3 using the parameters in Table 120D-7 applied and optimized for maximum SNR_ISI." to

"SNR_ISI is defined by Equation (120D-8) computed from p_max and ISI_cursors after these have been re-calculated with the selectable continuous time linear equalizer (CTLE) which is described in 120E.3.1.7 by Equation (120E-2) with coefficients in Table 120E-2 and illustrated in Figure 120E-9 applied and optimized for maximum SNR_ISI."

Proposed Response Response Status O

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 120D SC 120D.3.1.7 P 356 L 23 # 32
 Hidaka, Yasuo Fujitsu Lab. of Americ

Comment Type E Comment Status X
 Table 120D-7 is referred for the parameters of the CTLE, but Table 120D-7 is a table of 200GAUI-4 and 400GAUI-8 receiver jitter tolerance parameters.

SuggestedRemedy
 Change "Table 120D-7" to "Table 120D-8".

Proposed Response Response Status O

Cl 120D SC 120D.3.1.7 P 356 L 38 # 14
 Ran, Adee Intel

Comment Type E Comment Status X
 Per the style manual (16.1), "Note" should be all-caps, followed by an em dash and use the note paragraph format.

SuggestedRemedy
 per comment

Proposed Response Response Status O

Cl 120D SC 120D.3.1.7 P 356 L 24 # 36
 Hidaka, Yasuo Fujitsu Lab. of Americ

Comment Type T Comment Status X
 The SNR_ISI specification is defined to be met for all transmit equalization settings. When the transmit equalization settings is stronger than required, the SNR_ISI includes not only ISI due to reflection, but also ISI due to over-equalization, because the CTLE in the COM parameter cannot suppress the high-frequency component.

SuggestedRemedy
 Change
 "The SNR_ISI specification shall be met for all transmit equalization settings."
 to
 "The SNR_ISI specification shall be met for all transmit equalization settings excepting those settings which makes the mean value of ISI_cursors always negative regardless of the continuous time filter settings."

Proposed Response Response Status O

Cl 120D SC 120D.3.1.7 P 356 L 38 # 31
 Ewen, John GlobalFoundries

Comment Type E Comment Status X
 Incorrect table reference for parameter Nb

SuggestedRemedy
 Replace Table 120D-7 with Table 120D-8

Proposed Response Response Status O

Cl 120D SC 120D.3.1.7 P 356 L 38 # 34
 Hidaka, Yasuo Fujitsu Lab. of Americ

Comment Type E Comment Status X
 M and N_p are not defined in 85.8.3.3.5.
 N_b is not found in Table 120D-7.

SuggestedRemedy
 Change "Note: M and N_p are defined in 85.8.3.3.5, and N_b is found in Table 120D-7."
 to
 "Note: M is defined in 85.8.3.3.4. N_p is defined in 120D.3.1.3. N_b is found in Table 120D-8."

Proposed Response Response Status O

Cl 120D SC 120D.3.1.8 P 356 L 40 # 16
 Ran, Adee Intel

Comment Type E Comment Status X
 The first three paragraphs of 120D.3.1.8, describing even-odd jitter signal, transitions, thresholds, filter, and what other lanes are transmitting, seem to repeat the corresponding text of "output jitter" in 120D.3.1.1. If there are any differences, they are difficult to identify.

It would help the readers to have the even-odd jitter definitions within the output jitter subclause, share definitions where it is possible, and note differences where they exist.

SuggestedRemedy
 Move the specific even-odd measurement text, p357 lines 1-25, to 120D.3.1.1, noting differences if there are any, with editorial license.

Delete 120D.3.1.8.

Proposed Response Response Status O

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 120D SC 120D.3.1.8 P 356 L 40 # 48
 Dudek, Mike Cavium

Comment Type E Comment Status X

It would read better if this Even-Odd Jitter section were placed next to the Output jitter section.

SuggestedRemedy

Make this a subsection 120D.3.1.1.2 . Also relabel the existing section 120D.3.1.1.as a sub-section 120D.3.1.1.1 called "J4 and Jrms"

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

CI 120D SC 120D.3.1.8 P 356 L 40 # 35
 Hidaka, Yasuo Fujitsu Lab. of Americ

Comment Type E Comment Status X

Specification of jitter is split to 120D.3.1.1 and 120D.3.1.8.

SuggestedRemedy

Reorganize 120D.3.1.1 and 120D.3.1.8 as follows:

- 120D.3.1.1 Output jitter
- 120D.3.1.1.1 J4 and J_RMS jitter
- 120D.3.1.1.2 Even-odd jitter

Change the references in Table 120D-1 as follows:

- J_RMS (max) 120D.3.1.1.1
- J4 (max) 120D.3.1.1.1
- Even-odd jitter (max) 120D.3.1.1.2

Proposed Response Response Status O

CI 120D SC 120D.3.1.8 P 356 L 50 # 17
 Ran, Adeee Intel

Comment Type T Comment Status X

"Even-odd jitter is measured with a single-pole high-pass filter with a 3 dB bandwidth of 4 MHz"

What is this filter applied to?

If this text stays here, it should refer to the CRU.

SuggestedRemedy

Unless this text is deleted as a result of another comment, change it to state that "Even-odd jitter is measured with a clock recovery unit (CRU) with a corner frequency of 4 MHz and a slope of 20 dB/decade".

Proposed Response Response Status O

CI 120D SC 120D.3.2.1 P 358 L 8 # 49
 Dudek, Mike Cavium

Comment Type T Comment Status X

This is a follow up to the un-satisfied comment #118 on draft 2.1. The change to Np from 13 to 200 while calibrating the Interference Tolerance test allows the test system to have bad reflections after 13UI that won't appear in the measurement of TxSNDR (and hence input to TxSNR for the COM calibration). This will overstress the receiver.

SuggestedRemedy

Either use Np =13 for the measurement of the TxSNDR of the test transmitter

Replace "The parameter SNR_{TX} is set to the measured value of SNDR" with "The parameter SNR_{TX} is set to the measured value of SNDR with Np=13,

or add an extra very tight specification of SNR_{isi} of 40dB for the test transmitter. (Variations in SNR_{isi} of the test transmitter will cause repeatability issues in the interference tolerance test if not calibrated out by the first solution).

Add an extra bullet after a) at line 53 page 357.

SNR_{isi} of the test transmitter shall be greater than 45dB.

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 120D SC 120D.3.2.1 P 358 L 14 # 30
 Ran, Adeee Intel

Comment Type TR Comment Status X

As a sanity check, I calculated what would happen with a purely dual-dirac jitter (no RJ) equal to the specified J4, and with purely random jitter (no DD) equal to the specified JRMS.

In the first case, J4 is 0.0118 and JRMS would be $\sqrt{0.0118}=0.109$; plugging these values to equations 120D-9 and 120D-10 yields A_DD=0.106 and sigma_RJ=0.192; instead of the expected A_DD=0.0118 and sigma_RJ=0.

In the second case, JRMS is 0.023 and J4 would be $0.023 \cdot Q(1e-4/2)=0.09$; plugging these values to equations 120D-9 and 120D-10 yields A_DD=0.022 and sigma_RJ=0.007; instead of the expected A_DD=0 and sigma_RJ=0.023.

The equations originated from comment #25 against D2.0 which has very little explanation. I have not found any further analysis and suspect that the equations may be incorrect.

SuggestedRemedy

Correct the equations. I will try to find a more detailed remedy for comment resolution.

Proposed Response Response Status W

[Editor's note: Category set to T]

Cl 120D SC 120D.3.2.2 P 359 L 21 # 15
 Ran, Adeee Intel

Comment Type TR Comment Status X

There seems to be a mismatch SJ in the jitter tolerance test and the A_DD parameter.

Looking at the precedence in 83D:

- The channel is specified with COM parameter A_DD=0.05 (Table 83D-6), corresponding to 0.1 UI PtP. The transmitter specification has the same value allowed for effective DJ.
- The SJ stress at high frequencies is 0.05 UI PtP (from Table 88-13).

This means the SJ stress is 50% lower than the maximum allowed for the transmitter; the test in 83D is understressed (unless the transmitter has intrinsic DJ of 0.05 UI PtP).

In the current annex

- The channel is specified with COM parameter A_DD=0.02 corresponding to 0.04 UI PtP (the transmitter specification may not match this value; as noted in another comment)
- The SJ stress at high frequencies is 0.05 UI PtP (Table 120D-7)

This means the SJ stress is 25% higher than the maximum allowed for the transmitter; the test is overstressed (even if the transmitter has no intrinsic DJ).

The SJ stress is supposedly based on the CRU bandwidth so all frequencies should be scaled similarly.

SuggestedRemedy

Change table 120D-7 so that the SJ is 0.04 UI PtP at high frequencies (cases C, D and E), 0.12 UI for case B, and 4 UI for case A.

Proposed Response Response Status O

Cl 120D SC 120D.5.4.1 P 364 L 51 # 50
 Dudek, Mike Cavium

Comment Type E Comment Status X

There are no pics for SNR_{risi} or TxSNDR.

SuggestedRemedy

Add Pics

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

CI 120E SC 120E.3.3.2.1 P 377 L 34 # 51
 Dudek, Mike Cavium

Comment Type T Comment Status X
 There is no mention of error counters in 119.2.5.3

SuggestedRemedy

Change "119.2.5.3" to "119.3.1"

Also on page 380 line 4

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

CI 121 SC 121.8.4 P 223 L 9 # 5
 Ran, Adee Intel

Comment Type T Comment Status X

The response to comment #49 on D2.1 had the unfortunate effect that the OMA specification is now stated as conditional: "if measured using a test pattern specified..." in all clauses.

Naturally the OMA has to be within the specified range regardless of the pattern (e.g. for data patterns too) and regardless of whether it is measured or not.

This applies to 121.8.4, 122.8.4, and 124.8.4.

SuggestedRemedy

Change in all three clauses

FROM:

"within the limits given in Table XXX if measured using a test pattern using specified for OMAouter in Table YYY"

TO:

"within the limits given in Table XXX. OMA_outer is measured using a test pattern specified in Table YYY"

(no change in the table numbers)

Proposed Response Response Status O

CI 121 SC 121.8.5.2 P 224 L 32 # 18
 Ran, Adee Intel

Comment Type E Comment Status X
 (comment is about text that has not changed from D2.1)

"DGD" appears in the row heading of table 121-11, and in the text line 51, without definition or explanation. It only appears expanded several pages later, in a footnote of table 121-13, which is the last occurrence of "DGD" in this clause.

Acronyms should be expanded and explained on the first usage.

Comment applies to clause 122 too.

SuggestedRemedy

Move acronym expansion and explanation from footnote of table 121-13 to footnote of table 121-11 or within the text.

Apply similar change in clause 122.

Proposed Response Response Status O

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 121 SC 121.8.5.3 P 225 L 6 # 44
 Dudek, Mike Cavium

Comment Type T Comment Status X

The change to use the equalized eye for measuring OMAouter creates significant potential confusion. The definition is for TDECQ but by inference it is for all OMAouter measurements as the same name is used. Because the DC gain of the equalizer depends on the tap weights this will effect all the link budgeting. On a dispersive channel Tx OMAouter minus Rx OMAouter will not equal the channel loss, because the tap weights will be different for the Tx signal versus the Rx signal. It also somewhat conflicts with the definition in 121.8.4.

SuggestedRemedy

Put the gain Cdc into the reference equalizer so that the reference equalizer has 0dB gain at dc.

Replace OMAouter*Cdc with OMAouter in equation 121-9.

Delete lines 1 and 2 on page 228.

add in 121.8.5.4 at line 13. "The reference equalizer contains a gain element with gain Cdc which ensures that the equalizer has unity DC gain for all equalizer settings." Move lines 4 to 9 on page 228 (including equation 121-10) immediately after this.

Alternatively clarify that OMAouter used in TDECQ is not the same as the OMAouter used in measuring the output of the Tx or calibrating the stressed input to the Rx. Change "OMAouter is measured according to 121.8.4 on the equalized signal" to "For this subsection only, OMAouter is measured on the equalized signal according to 121.8.4"

Make the equivalent changes in clauses 122.8.5.4 (or consider deleting this section and referencing clause 121.8.5.4 instead as the content is the same, (like 124.8.5 does))

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

Cl 121 SC 121.8.5.3 P 225 L 9 # 40
 Dawe, Piers Mellanox

Comment Type T Comment Status X

I didn't see a statement of whether averaging is used or not.

SuggestedRemedy

State that averaging is not used.

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

Cl 121 SC 121.8.5.3 P 225 L 12 # 6
 Ran, Adeed Intel

Comment Type T Comment Status X

Should "OMA" be "OMA_outer" which is defined above?

(if not, specify what "OMA" it is)

SuggestedRemedy

Change "OMA" to "OMA_Outer" across this subclause

Proposed Response Response Status O

Cl 121 SC 121.8.5.3 P 225 L 13 # 42
 Dawe, Piers Mellanox

Comment Type T Comment Status X

MMSE should be loaded with the amount of noise that could be added for a maximum-TDECQ signal, adjusted for scope noise already in the measurement

SuggestedRemedy

Add noise loading to the the mean square error calculation

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

Cl 121 SC 121.8.5.3 P 225 L 13 # 41
 Dawe, Piers Mellanox

Comment Type T Comment Status X

Window for equalizer tuning (the central 0.1 UI of the eye diagram) doesn't match the windows for TDECQ used later.

SuggestedRemedy

Do the tuning with the histogram windows used later.

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 121 SC 121.8.5.3 P 226 L 38 # 8
 Ran, Adee Intel

Comment Type T Comment Status X

The term "symbol error ratio" is used (along with the "unofficial" acronym) in several places, including within this draft, referring to the _FEC symbol_ error ratio, e.g. with 10-bit symbols. Here it seems to be used for _PAM4 symbol_ error ratio, but it is not stated that this is a different meaning than the usual one. This may be very confusing for the reader.

There is another term, detector error ratio (DER), that is used in several recent clauses when referring to physical receiver (PMD or AUI) decisions, regardless of the modulation. It is defined precisely in 93A.1.7, and it would be adequate to use it here too.

SuggestedRemedy

Change "symbol error ratio" to "detector error ratio" three times in this subclause. No need to introduce an acronym for this term.

Proposed Response Response Status

Cl 121 SC 121.8.5.3 P 226 L 47 # 10
 Ran, Adee Intel

Comment Type T Comment Status X

Equation (121-9) yields TDECQ in dB, but doesn't say that.

Since this value is used in a specification it is good to avoid confusion.

SuggestedRemedy

Add "(dB)" at the right of this equation.

Proposed Response Response Status

Cl 121 SC 121.8.5.3 P 226 L 53 # 9
 Ran, Adee Intel

Comment Type T Comment Status X

The "target PMD BER" in 121.1 is $2.4e-4$. With Gray coding it should be assumed that every detector error translates to a single bit error so the target detector error ratio should be $x2$ of the BER, so $4.8e-4$.

This seems to be inconsistent with the value of Q_t , since the Q-function yields $Q(3.414)=3.2e-4$.

It seems that the correct value should be $Q^{-1}(4.8e-4)=3.302$.

SuggestedRemedy

If there is another calculation that yields the current value, please clarify the text to prevent any suspicion.

Otherwise:

Change from "is 3.414 consistent with the BER and target symbol error ratio for Gray coded PAM4" to "is 3.302 consistent with the target BER (see 121.1) and using a single bit error for every PAM4 detector error, due to Gray coding".

Proposed Response Response Status

IEEE P802.3bs D2.2 200 Gb/s & 400 Gb/s Ethernet 2nd Working Group recirculation ballot comments

Cl 121 SC 121.8.5.3 P 227 L 2 # 7
 Ran, Adee Intel

Comment Type TR Comment Status X

The method of finding the "estimate of the partial symbol error ratio" is not clear; without a clear specification the definition of TDECQ is very ambiguous.

Reading the sentence "Each element of the cumulative probability function Cf1(yi) is multiplied by a value Gth1(yi), and then summed to calculate an approximation for the partial symbol error ratio (SER) for threshold 1."

The operation that should be performed is not clear from the text. Trying to guess what should be done I find some mathematical difficulties.

To find the SER estimate we should really find the value of the "bathtub function" at the threshold level. Cf1 definition makes it sort-of a bathtub function - it approaches 0.25 at the lowset values of yi, 0.75 at the highest values, and has a minimum at the threshold Pth1 (Cf1 is not a CDF since a CDF should start at 0 and rise monotonically to 1).

If each element of Cf1 at index yi is multiplied by a corresponding value of Gth1 at the same index yi (as the text suggests), then Gth2 is a weighting function operating on a bathtub function. What does that achieve?

"and then summed" suggests a convolution operation, but this is not obvious (I am not sure it is) and there is no equation that one can follow. Why should the bathtub function be summed? It is already cumulative; we only need the value at a specific point.

Assuming this is a convolution, this seems like incorrect math. A convolution between two PDFs of two independent variables yields the PDF of the sum of the variables; but here we have a PDF for one thing (approximated Gaussian noise Gth1(yi)) and a "bathtub curve" CF1(y1) for another thing (measured data). To add noise to the measured data, the convolution should be between Gth1(yi) and the normalized histogram f(yi); and then a bathtub function of the should can be calculated. From that bathtub function we can estimate the partial SER of that specific thredhold.

Note also that the total SER is the sum of partial SERs divided by 4 - not the sum as currently written - since each partial SER is a conditional probability (error rate given that the signal is within a specific eye); there is a probability of 1/4 to be at each of the 3 eyes plus 1/4 to be "outside of all eyes".

SuggestedRemedy

If the intent is to model adding Gaussian noise to the measured data: change the text so that the process is

1. f(y) is convolved with Gth(y) to yield fn1(y) (incldue equation)
2. fn1(y) is integrated to create bathtub function BF1(y) (include equation)
3. The value of BF1(y) at pth1 is the SER1, the partial SER for threshold 1
4. repeat for thresholds 2, 3
5. The total SER estimate is (SER1+SER2+SER3)/4.
6. adjust sigma_G so that the total SER from the previous steps becomes 4.8e-4.

TYPE: TR/technical required ER/editorial required GR/general required T/technical E/editorial G/general
 COMMENT STATUS: D/dispatched A/accepted R/rejected RESPONSE STATUS: O/open W/written C/closed U/unsatisfied Z/withdrawn
 SORT ORDER: Clause, Subclause, page, line

If there is another intent then please write clearly what is to be done here.

Proposed Response Response Status O

Cl 122 SC 122.7.3 P 255 L 32 # 45
 Dudek, Mike Cavium

Comment Type T Comment Status X

The footnote to the channel insertion loss is strange. Saying that it won't support operation at 10km isn't true if the channel insertion loss meets the 6.3dB specification. (which is a normative specification in table 122-17). It also isn't specific to 400GBASE-LR8 and would apply to 200GBASE-LR4 as well.

SuggestedRemedy

Delete the footnote here and add a footnote to the 6.3 in table 122-17 that says "To meet this specification with 10km of fiber using the 0.46dB/km at 1272.55nm attenuation for optical fiber cables derived from Appendix I of ITU-T G.695 the connection insertion loss must be less than 2dB." It might be better to amend 122.11.2.1 instead to use a lower allocation for connection and splice loss (1.6dB). Then the footnote would not be needed.

Proposed Response Response Status W

[Editor's note: This comment was sent after the close of the comment period]

Cl 122
 SC 122.7.3

Page 14 of 14
 16/12/2016 21:59:22