Cl **163A** *SC* **163A.3.1.3** *P* **308** *L* **18** # 21

Hidaka, Yasuo — Credo Semiconductor, Inc.

Comment Type **TR** *Comment Status* **A** *measurement filter*

A measurement filter of BT filter is already included, because the step response is derived from the pulse response h(t) that uses the BT filter.

Figure 163A-3 is not correct, because the effect of BT filter is included.

SuggestedRemedy

Remove Editor's note in page 308.

Change Figure 163A-3 as follows:
Add H_BT(f) in the same way as Figure 163A-2.
Append a block of "Equation (163A-5)" followed by "Stepresponse u(t)" at the end after "Pulse response h(t)".

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

This subclause needs to be aligned with the interference tolerance test in 163 and 120F, but there is no consensus to make related changes at this time.

Add an editorial note that this method needs to be aligned with the interference tolerance test in 163 and 120F.

Cl **162** *SC* **162.9.3.1.1** *P* **165** *L* **5** # 29

Ran, Adee — Cisco systems

Comment Type **TR** *Comment Status* **R** *Np value*

Here it is stated that Np takes the value 29, but this value is only effective for calculation of SNDR. Other invocations of this procedure, for vf and vpeak, use Nv=200 instead. Nv appears several times and looks like a parameter, but it is not - it is a value that replaces Np; this is not stated anywhere.

In the remaining use of the linear fit, for calculation of the equalizer coefficients used in 162.9.3.1.3, 162.9.3.1.4, and 162.9.3.1.5, it does not matter whether 29 or 200 UI are used. So Np=29 is important only for SNDR, which is the exception.

Having two parameters instead of one parameter which takes two values is unnecessary and confusing.

SuggestedRemedy

In 162.9.3.1.1, change "Np=29" to "Np=200".

In 162.9.3.3 (Output SNDR) change "with the exception that the linear fit procedure in 162.9.3.1.1 is used" to "with the exception that the linear fit procedure in 162.9.3.1.1 is used with Np=29 instead of 200".

In 162.9.3.1.2 (Steady-state voltage and linear fit pulse peak) delete "using Nv=200".

In 163.9.2.3 (Difference steady state voltage) delete "with Nv = 200".

In 163A.3.1.1 (Steady-state voltage and pulse peak reference values) change "Nv" to "Np" (3 times).

In 163B.2 (Characteristics) delete "With Nv = 200".

With editorial license, change any remaining occurrence of Nv to Np.

Response *Response Status* **U**

REJECT.

The following presentation was reviewed by the task force at a previous ad hoc meeting. https://www.ieee802.org/3/ck/public/adhoc/july14_21/wu_3ck_adhoc_01a_071421.pdf.

There is no consensus to make the proposed changes at this time.

[Editor's note: CC: 162, 163, 163A, 163B]

Cl **120G** *SC* **120G.5.1** *P* **264** *L* **31** # 37

Ran, Adee Cisco systems

Comment Type **TR** *Comment Status* **R** *signal level (CC)*

This clause is referred to in Table 120G-1 and Table 120G-3 for the parameter differential PtP output voltage (max), among others.

The content is only a reference back to 120E.3.1.2: "The signal levels are as defined in 120E.3.1.2". 120E.3.1.2 does have a definition of differential signal but also states that "Unless otherwise noted, differential and common-mode signal voltages are measured with a PRBS13Q test pattern".

But PRBS13Q is not an appropriate signal for measurement of the PtP output voltage, because it has a maximum run length of 7 symbols and does not have any spectral content below 3 MHz. Much longer runs are possible in real data. Measurement with PRBS13Q over a lossy channel between the transmitter and the measurement point, without sufficient equalization, can thus yield peak-to-peak value lower than the value that real data would create.

Since there is no way to control the transmitter's swing or equalization, this may cause events of higher signal levels than the receiver expects, and cause periods of high BER, which can span many FEC symbols and cause uncorrectable codewords.

It is proposed to define the differential PtP explicitly as a requirement for any data pattern, and recommend to measure it using a pattern that contains low-frequency content, such as PRBS31Q or SSPRQ.

The definition of signal levels measurement using PRBS13Q also applies for CR/KR/C2C but in these cases the transmitter can be controlled to reduce the signal to an adequate level for the receiver, so it is less of an issue.

SuggestedRemedy

Replace the content of 120G.5.1 with the following:

"The definition of differential and common-mode signals can be found in 120E.3.1.2. The signal levels specifications for host and module outputs hold for any data pattern. It is recommended to measure differential peak to peak signal levels with PRBS31Q or SSPRQ test pattern."

Consider applying similar changes in 162, 163, and 120F, with editorial license.

Response *Response Status* **U**

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3ck D2.1 and D2.0 or the unsatisfied negative comments from the initial ballot.
Hence it is not within the scope of the recirculation ballot.

The proposal to refer "any data pattern" is rather broad.
SSPRQ has been previously used only for optical transmitter testing and has no advantages for this test.
It is not clear that similar changes are warranted for 162, 163, and 120F since the insertion loss to the test point is smaller.

There was some agreement that this specifications should be improved but there was no consensus on a resolution.

[Editor's note: CC: 120F, 120G, 162, 163]

Straw poll #13 (decision)
I support closing comment #37 updating 120G.5.1 as follows:
"The signal levels are as defined in 120E.3.1.2, with the exception that differential signal voltage is measured with a PRBS31Q (see 120.5.11.2.2) test pattern or a valid 100GBASE-R, 200GBASE-R, or 400GBASE-R signal."
Y: 10
N: 14

Cl **120G** *SC* **120G.3.1** *P* **250** *L* **12** # 46

Ran, Adee Cisco systems

Comment Type **TR** *Comment Status* **A** *AC CM noise*

"AC common-mode RMS output voltage (max)" specification of 17.5 mV is not feasible for high-volume, multi-port products. The common-mode output may include a component correlated to the differential output, e.g. from mode conversion on the host channel. A module receiver is expected to be quite tolerant to a correlated common-mode signal.

As suggested in ran_3ck_adhoc_20210630, there are two reasonable alternatives:
a) increase the allowed RMS voltage to 30 mV (as is allowed for the CR transmitter measured on an HCB - likely the same point - and where the common-mode concern is greater due to conversion in the cable assembly).
b) Keep the 17.5 mV specification but only for the component uncorrelated to the differential signal; use the linear fitted pulse response method (which is already referred to in 120G.5.2) to calculate the linear fitted pulse response characteristics of the common-mode output, and define the AC common-mode noise as the RSS of sigma_n and sigma_v.

Note: This comment is only about the host output; module output is more controlled and modules can be designed to have low mode conversion so the correlated component is expected to be small. Modules should not be allowed to generate 30 mV RMS, so if option a is chosen, the module output specification should not be changed.

SuggestedRemedy

Preferably implement option a in the comment.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE P802.3ck D2.1 and D2.0 or the unsatisfied negative comments from the initial ballot.
Hence it is not within the scope of the recirculation ballot.

Comment 121 proposes to increase the value to 25 mV.

This comment proposes to either:
(a) change the value to 30 mV
(b) change the parameter to relate to only the uncorrelated noise
There is not sufficient evidence that the correlated noise is indeed tolerable by the receiver (e.g., conversion from CM to DM in receiver might be non-linear or CM might have much larger channel transit time than DM)

The resolution to comment #123 indicates there is not consensus to make the change proposed in option (b), above.

Following straw polls #3 and #4, there was consensus to close this comment changing the value to 25 mV.

Change the AC common-mode RMS output voltage (max) for module output and host output to 25 mV.

Straw poll #3, pick one (direction)
Straw poll #4, Chicago rules (direction)
To address comments #46 and #121, for the module output and host output AC CM noise (max) I would support:
A: no change
B: change to 25 mV
C: change to 30 mV
Straw poll #3
A: 12 B: 13 C: 9
Straw poll #4
A: 15 B: 25 C: 21

Cl **120G** *SC* **120G.3.1** *P* **250** *L* **25** # 58

Ghiasi, Ali Ghiasi Quantum/Inphi

Comment Type **TR** *Comment Status* **A** *HO TT*

Transition time host requesting short mode or long mode is for TP4

SuggestedRemedy

Please revert to 10 ps in draft D2.0, please move this parameter to TP4 table 120G-3

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

This comment relates to the host output transition time specified in Table 120G-1.

Separate values for host long and short modes were added per D2.1 comment #188.

The justification was that the host input and host output PCB insertion loss will likely be similar, which is reflected in the transition times chosen for the host input crosstalk calibration. This must also be explicitly allowed and constrained at the hout output.

However, it would be helpful in Table 120G-1 to point to the subclause that defines long and short modes.

Add a footnote to the sub-rows for long and short modes in Table 120G-1 pointing to 120G.3.2.1.

Cl **120G** *SC* **120G.3.2** *P* **253** *L* **12** # 62

Ghiasi, Ali — Ghiasi Quantum/Inphi

Comment Type **TR** *Comment Status* **A** *MO VEC/EH*

TP4 VEC can be lowered from current 12 dB to 11 dB to allow additional penalty for real host channel and host ASIC

SuggestedRemedy

Reduce TP4 VEC=11 dB, see ghiasi_3ck_01_0721

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

This comment pertains to the module output VEC (max).

Slides 7 and 8 of the following presentation was reviewed by the task force:
https://www.ieee802.org/3/ck/public/21_07/ghiasi_3ck_01_0721.pdf

The slide shows that with the current g_dc constraints VEC fails for the long mode, near-end measurement. The comment suggests that g_dc max for TP4 far-end be increased from -3 dB to -2 dB. With this change to the g_DC limit there is no need to change VEC (max).

Cl **162** *SC* **162.9.3** *P* **163** *L* **18** # 92

Dawe, Piers — Nvidia

Comment Type **TR** *Comment Status* **R** *host/CA IL*

The draft CR loss budget wastes over 3 dB in nearly every case. The relative range of host losses, 6.875/2.3 = 3:1, is too small for switch layout yet not needed for NICs.
The recommendation for the host traces plus BGA footprint and host connector footprint, 6.875 dB, compares very poorly with C2M's host insertion loss up to 11.9 dB, making passive copper to this draft expensive and unattractive for a switch, yet a full range of NICs can be made with only 3.75 dB. Server-switch links are asymmetric in form factor (e.g. QSFP-DD to 2 x QSFP) and will get made with an asymmetric loss budget, so it would be better for the standard to regularise what will happen anyway. C2M already has short and long ports.
This change would also benefit CR switch-switch links because the shortest ports would get credit for their low loss.
The symmetric budget is used for some designs under way and may be useful in future for LOM, so it is kept here, and the better way added.

SuggestedRemedy

3 classes of CR ports, host loss allocations of A 10, B 6.875, C 3.75 dB. B is as D2.1.
A connects to C, B to B or C, C to A, B or C.
Use 2 bits in Clause 73 Auto-Negotiation Link codeword Base Page to advertise A, B or C to the other end. In the Priority Resolution function, an A port ignores a 100G/lane Technology Ability Field bit from an A or B port, a B port ignores a 100G/lane Technology Ability Field bit from an A port.
In Table 162-10, add limits A and C for linear fit pulse peak ratio (min). Change text in 162.9.3.1.2 to refer to the table.
In Table 162-14, add columns for Test 2 (high loss), A and C, with test channel insertion loss: A: 6.875-3.75 = 3.125 dB lower (20.5 dB to 21.5 dB), and C: 10-6.875 = 3.125 dB higher (26.75 dB to 27.75 dB). No change needed for Test 1.
In 162A.4, add equations for IL_PCBmax and ILHostMax A and B and show them in Fig 162A-1 and 2. In 162A.5, add Value columns A, C in Table 162A-1 (ILChmin and ILMaxHost differ). Adjust figures 162A-3 and 4.

Response *Response Status* **U**

REJECT.

D2.0 straw polls #6 and #7 indicated interest in exploring multiple CR port types. However, consensus is needed to make a change of this magnitude.

The following presentation was reviewed by the task force:
https://www.ieee802.org/3/ck/public/21_07/dawe_3ck_01a_0721.pdf

Based on straw poll #10, there is not sufficient consensus to implement the proposed changes in dawe_3ck_01a_0721.

Strawpoll #10 (direction)
I support P802.3ck specifying multiple CR host types such as in dawe_3ck_01_0721.
Y: 7

N: 24
A: 8

Cl **162** *SC* **162.11.6** *P* **181** *L* **38** # 94

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *CA RLcc*

Relaxing the already very loose CM RL spec from 2 dB to 1.8 dB at all frequencies isn't justified. This draft spec becomes useless at the frequency when the MCB loss is 1.8/2 dB, which is only 8.5 GHz.

SuggestedRemedy

Use a frequency-dependent mask e.g. 1.6 + 0.01f. Similarly for Tx, Table 162-11, 162.9.3.6.

Response *Response Status* **U**

REJECT.

The basis for the change to the cable assembly CM-to-CM RL spec from 2 dB to 1.8 dB was given in the following presentation.
Https://www.ieee802.org/3/ck/public/21_01/champion_3ck_01a_0121.pdf

The comment and suggested remedy does not provide sufficient information or justification to support a change to the draft.

Cl **162** *SC* **162.11.7** *P* **183** *L* **39** # 95

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *COM bbgmax*

The normalized DFE coefficient minimum limit bbmin for taps 3 to 12 is -0.03. It doesn't make sense that taps 13 to 40 could be worse, -0.05. If I have understood the data correctly, the example channels we have don't need this. (Remember, these are reference receiver limits not hard cable or channel limits anyway; a cable or channel can go beyond a tap limit if it makes up the COM another way, e.g. with acceptable crosstalk.)

SuggestedRemedy

Change bgmax 0.05 to bbgmax 0.05, bbgmax -0.03. Also in 163.

Response *Response Status* **U**

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3ck D2.1 and D2.0 or the unsatisfied negative comments from the initial ballot.
Hence it is not within the scope of the recirculation ballot.

The following presentation showed that some backplane channels had floating tap coefficient values of <-0.03.
https://www.ieee802.org/3/ck/public/19_09/heck_3ck_01_0919.pdf
The comment does not provide an assessment of the impact to those channels.
[Editor's note: CC: 162, 163]

Cl **162** *SC* **162.11.7** *P* **183** *L* **40** # 96

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *COM DFE RSS*

The spec allows a cable (not even the whole channel) to have its COM calculated with 9 taps in the range 13 to 24 clipped at +/-0.05 - which means that the channel's pulse response could be worse than +/-0.05 for all these 9 taps. That's a very bad cable! and not likely to get made: there won't be that many reflections in the same area. (Remember, these are reference receiver limits not hard cable limits anyway; a cable can go beyond a tap limit if it makes up the COM another way, e.g. with acceptable crosstalk.)
We don't need to provide all the receiver power and complexity to cope with unreasonably bad cables.

SuggestedRemedy

Use another DFE root-sum-of-squares limit for positions 13-24. Similarly in 163, but as 163 specifies the complete channel while 162 uses clean synthetic host traces, the limit should be higher.

Response *Response Status* **U**

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3ck D2.1 and D2.0 or the unsatisfied negative comments from the initial ballot.
Hence it is not within the scope of the recirculation ballot.

The suggested remedy is not complete nor has sufficient analysis been provided.

[Editor's note (added after the comment was addressed by the task force): The comment response incorrectly describes this comment as being out of scope as this comment is a restatement of unsatisfied D2.0 comment #235.]

Cl **120G** *SC* **120G.3.2** *P* **253** *L* **11** # 97

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *MO VEC/EH*

The driver swing has to be aggressively reduced from 600 mV pk-pk to deliver only 15 mV at near end, short mode. 120E has 70 mV, and D1.4 had 24 mV, ghiasi_3ck_adhoc_01a_042121 shows 35 mV (before Vpkpk was reduced). Yet a host can usefully optimise for e.g. different crosstalk or noise if given a reasonable signal strength. A NIC has no high-loss ports so it can do this even if a switch won't. There is room to increase this weak signal without overloading the receiver. Also, making the limits more like reality encourages more consistent module setup across the industry.

SuggestedRemedy

Increase the eye height, short mode near end, by 1.1 dB from 15 mV to 17 mV

Response *Response Status* **U**

REJECT.

This comment pertains to the module output eye height (min) for short mode, near end.

The comment does not provide sufficient evidence that the proposed change is necessary.

Cl **120G** *SC* **120G.3.2** *P* **253** *L* **11** # 98

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *MO VEC/EH*

If the eye height limit is the same at long near end as at long far end, there is huge margin at near end and the implementer is encouraged to optimise for far end or beyond, only limited by the NE VEC spec, while we want modules to be set up consistently, for the full range from near to far. EH is naturally larger at NE for a well set up output.

SuggestedRemedy

Increase the eye height, long mode near end, by 3 dB from 15 mV to 21 mV

Response *Response Status* **U**

REJECT.

This comment pertains to the module output eye height (min) for long mode, near end.

The comment does not provide sufficient evidence that the proposed change is necessary.

Cl **120G** *SC* **120G.3.2.2.1** *P* **254** *L* **51** # 102

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *MO SI host reference channel*

The near end and far end should be placed far enough apart so that the module implementer has little choice what emphasis to use, so that all modules are set up similarly. As short is easier than long, this means that far minus near (mm or dB) for short should be at least as much as far minus near for long. As real host channels are not exactly like the theoretical reference host channel, there should be a healthy overlap of short and long to give the host room for its implementation. D2.0's 160 mm delivered on both these criteria, D2.1's 133 mm doesn't.

SuggestedRemedy

Change 133 to 150, change 80 to 90

Response *Response Status* **U**

REJECT.

The comment does not provide sufficient justification for the proposed changes.

There may be some benefit to balancing the length range between short and long modes. Further analysis is encouraged.

Cl **120G** *SC* **120G.5.2** *P* **265** *L* **16** # 103

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *RR gdc*

The limits for TP4 gDC, gDC2 should not be the same for short and long output modes.

SuggestedRemedy

Create separate limits for TP4 short and long output modes, so 4 sets for TP4+, in the style of TP1a.

Response *Response Status* **U**

REJECT.

This comment is a restatement of D2.0 comment #179, which was rejected on the basis of insufficient justification and detail. It adds request to provide 4 sets of values in the style used for TP1a but does not provide specific values. No further justification is provided.

The comment does not provide sufficient justification for the proposed changes nor does the suggested remedy provide sufficient detail to implement.

Cl **120G** *SC* **120G.5.2** *P* **265** *L* **25** # 104

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *RR gdc*

As a lot of the channel for TP4 far-end is known exactly and the max loss to TP4 far end is less than to TP1a, the range of gDC, gDC2 combinations should be a subset of the TP1a ones. As for TP1a, I believe the strongest gDC and gDC2 should add to a constant.

SuggestedRemedy

For Continuous time filter, DC gain for TP4 far-end (gDC), change to a set of limits that depend on gDC2 in the same style as for TP1a, with the strongest gDC and gDC2 adding to a constant. The allowed values should be a subset of those for TP1a.

Response *Response Status* **U**

REJECT.
This comment is a restatement of D2.0 comment #178, which was rejected on the basis of insufficient justification and detail. No further justification or implementation detail is provided.
The comment does not provide sufficient justification for the proposed changes nor does the suggested remedy provide sufficient detail to implement.

Cl **120G** *SC* **120G.5.2** *P* **265** *L* **12** # 105

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *RR gdc*

When gDC2 is -2, we allow no more than -(-12-2) = 14 dB of peaking, yet when gDC2 is -3, we allow -(-13-3) = 16 dB, yet the channel loss should not be higher. This doesn't make sense.

SuggestedRemedy

For TP1a, change -12 -12 -13 to -12 -11 -10 or -12 -12 -11 (so the strongest CTLE peaking for the highest two gDC2 categories is the same).

Response *Response Status* **U**

REJECT.
The comment does not provide sufficient justification for the proposed changes.

Cl **120G** *SC* **120G.5.2** *P* **266** *L* **23** # 106

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **A** *EO method*

This draft has a primitive rectangular eye mask spec with mask height = max(EHmin, EA/VECmax) and mask width = 0.1 UI, although it is described as a histogram. Measuring a diamond eye with a rectangular mask is an inefficient, inaccurate way of measuring signal quality and provides weak and uncertain protection against too much jitter. Its effective width is less than its actual because of the 1e-5 probability criterion and the inefficient shape.
De-weighting the sides of the histogram/mask would make this worse, equivalent to increasing the target BER by 10x or so. A higher VEC / smaller EH limit with the rectangular mask would allow more jittered and more varied signals, particularly for very short host channels (see Mike Dudek's work) that can have faster edges than higher loss ones. The target BER is not going to change.
We need an eye mask that's more eye shaped, so that a higher proportion of the samples are near the boundary and contribute to the measurement.

SuggestedRemedy

Change from a 4-cornered mask with corners at t = ts+/-0.05, V = y +/-H/2 to a 10-cornered mask with corners at t = ts+/-0.05, ts+/-1/16, ts+/-3/32, V = y +/-H/2, k +/-H*0.4, y. y is near VCmid, VCupp or VClow (vertically floating, as in D2.1).
H is max(EHmin, Eye Amplitude * 10^(-VECmax/20)). Eye Amplitude is AVupp, AVmid or AVlow, as in D2.1.
This simple scalable method can remain as the EH and VEC limits are revised. Scopes have been measuring with 10-sided masks for many years, it's not more difficult than a rectangular mask and gives better results.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

This comment is a restatement of D2.0 comment #127, which was rejected on the basis of insufficient justification and insufficient analysis to show equivalent or better interoperability.

Straw polls 5, 6, and 7 indicate there is no consesus to make the proposed change. However, the resolution to comment #39 addresses the concern expressed in this comment.

Cl **162** *SC* **162.9.4.6** *P* **176** *L* **11** # 115

Dawe, Piers Nvidia

Comment Type **ER** *Comment Status* **R** *RLdc/RLcd graphs (bucket3)*

Don't waste the reader's time.

SuggestedRemedy

Combine the graphs for Transmitter common mode to differential return loss and Receiver differential to common-mode return loss.

Response *Response Status* **U**

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3ck D2.1 and D2.0 or the unsatisfied negative comments from the initial ballot.
Hence it is not within the scope of the recirculation ballot.

The two graphs represent requirements for different components, which happen in this case to have identical responses.

There is no consensus to make the proposed changes.

[Editor's note: Changed page from 175.]

[Editor's note (added after the comment was addressed by the task force): The comment response incorrectly describes this comment as being out of scope as the referenced figure was added in D2.1.]

Cl **163** *SC* **163.9.2.1.3** *P* **201** *L* **27** # 117

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **A** *TF RLcc (bucket2)*

Test fixture common-mode to common-mode return loss should be way better than the worst module connector! And needs to be significantly better than the spec for the IC+TF.

SuggestedRemedy

Change 2 to something sensible

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

This comment does not apply to the substantive changes between IEEE P802.3ck D2.1 and D2.0 or the unsatisfied negative comments from the initial ballot.
Hence it is not within the scope of the recirculation ballot.

This comment does not provide sufficient details for implementation.

The test fixture RLcc value is too small to permit measurement of a transmitter RLcc as specified. However, there is no consensus on an appropriate new specification. Further analysis and consensus is required.

Add an editor's note pointing out the issue as above calling for contributions to address this.

Cl **120G** SC **120G.3.1.5** P **252** L **13** # 119

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *pattern table*

As this annex uses several test patterns like an optical PMD, it should have a table of test patterns giving the pattern number, which this draft lacks, and description, and reference for definition.

SuggestedRemedy

Copy Table 167-10, Test patterns, leaving out the rows that don't apply. Refer to the table from elsewhere in the annex to reduce clutter end repetition.

Response *Response Status* **U**

REJECT.

Table 167-10 may be found in 802.3db.

It is not clear that the proposed table with pattern numbers will improve the draft all things considered.

It can indeed reduce some clutter for cases where multiple patterns are listed for a particular test step, but not in cases where a single pattern is referenced. It is more convenient to the reader to list the pattern names; the reader would otherwise have to memorize the relationship between pattern numbers and the pattern they represent. The test pattern names line up better with the test equipment controls.

Cl **162** SC **162.9.3** P **163** L **10** # 123

Mellitz, Richard Samtec

Comment Type **TR** *Comment Status* **R** *AC CM noise*

Table 162-10 specifies AC common-mode RMS voltage, vcmi (max) note b just changes to a PRBS13Q with method described in 93.8.1.3. The problem is that coherent CM signal are included in differential measurements like SNDR, Jitter, and Linear fit pulse peak ratio. That means it is the coherent part if AC CM is double counted.

SuggestedRemedy

Add note to line 10 (vcmi) indicating that the CM mode measurement is only for the non-coherent CM part of the measurement.

This applies to Tables 163-5, 120F-1, 120G-1, and 120G-3

Response *Response Status* **U**

REJECT.

[Editor's note: Changed clause/subclause from 163/163.9.3.]

This comment does not apply to the substantive changes between IEEE P802.3ck D2.1 and D2.0 or the unsatisfied negative comments from the initial ballot.
Hence it is not within the scope of the recirculation ballot.

The following presentation was reviewed by the task force:
https://www.ieee802.org/3/ck/public/21_07/mellitz_3ck_01a_0721.pdf.
Resolve in conjunction with comment #46.

Based on straw poll #2, there is not sufficient consensus to implement the proposed changes.

Straw poll #1 (direction)
I would support the AC CM voltage test methodology in Comment #123 and the related presentation mellitz_3ck_01_0721.
Yes: 18
No: 6
Need more information: 13
Abstain: 3

Straw poll #2 (decision)
For the resolution of comment #123, I support adopting the AC CM voltage test methodology in Comment #123 and the related presentation mellitz_3ck_01a_0721.
Yes: 15
No: 16

[Editor's note: CC: 163, 120F, 120G]

Cl **162** *SC* **162.9.3.4** *P* **158** *L* **39** # 20032

Ghiasi, Ali — Ghiasi Quantum/Inphi

Comment Type **TR** *Comment Status* **R** *EOJ CRU BW*

"Meeting even-odd jitter requriement with only one CRU bandwidth is sufficient" is not clear

SuggestedRemedy

What is the intention of only one CRU bandwidth, please make it clear.

Response *Response Status* **U**

REJECT.

The suggested remedy does not provide sufficient detail to implement.

There was some agreement that further clarification would be helpful. However, complete proposal is required.

Cl **120G** *SC* **120G.3.2** *P* **240** *L* **10** # 20034

Ghiasi, Ali — Ghiasi Quantum/Inphi

Comment Type **TR** *Comment Status* **R** *TP4 EH*

Given that now we have AUI-S/L far end eye would be AUI-S min eye opening

SuggestedRemedy

The eye opening with 50 mUI rectangular window for AUI-L is VEO=11 mV, see ghiasi_3ck_01_0121

Response *Response Status* **U**

REJECT.

Slide 9 of the following presentation was reviewed by the task force:
https://www.ieee802.org/3/ck/public/adhoc/apr21_21/ghiasi_3ck_adhoc_01a_042121.pdf

There was no consensus to make the proposed changes.

[Editor's note: Changed page/line from 164/13 to 240/10.]

Cl **163** *SC* **163.10.7** *P* **198** *L* **31** # 20037

Ghiasi, Ali — Ghiasi Quantum/Inphi

Comment Type **TR** *Comment Status* **R** *AC coupling*

Given that we have increased Baudrate it is logical to increase 3 dB cutoff by factor 2

SuggestedRemedy

Please increase 3 dB cutoff from 50 KHz to 100 KHz given that this standard is operating at 2x Baudrate of 802.3cd. It is well understood that if one needs to support 50G PAM4 then DC block corner frequency will be 50 KHz, but keeping 50 KHz for 100G PAM4 it just will force 200G gets force to 50 KHz assuming one generation support

Response *Response Status* **U**

REJECT.
There is insufficient justification that the suggested remedy does not degrade performance.
[Editor's note: CC: 162, 163]

Cl **162** *SC* **162.11** *P* **165** *L* **43** # 20038

Ghiasi, Ali — Ghiasi Quantum/Inphi

Comment Type **TR** *Comment Status* **R** *AC coupling*

Given that we have increased Baudrate it is logical to increase 3 dB cutoff by factor 2

SuggestedRemedy

Please increase 3 dB cutoff from 50 KHz to 100 KHz given that this standard is operating at 2x Baudrate of 802.3cd. It is well understood that if one needs to support 50G PAM4 then DC block corner frequency will be 50 KHz, but keeping 50 KHz for 100G PAM4 it just will force 200G gets force to 50 KHz assuming one generation support

Response *Response Status* **U**

REJECT.
The AC-coupling specification is used throughout 802.3ck and applied to predictive models as well as implemented in 802.3cd cable assemblies. The comment does not provide sufficient justification to support proposed change.
[Editor's note: CC: 162, 163]

Cl **120G** *SC* **120G.3.1** *P* **237** *L* **17** # 20039

Ghiasi, Ali — Ghiasi Quantum/Inphi

Comment Type **TR** *Comment Status* **R** *TP1 EH/VEC*

VEC limit of 12 dB and VEO limit of 10 mV results in well constructed host to fail, this was not the case prior to adding timing window of +/-50 mUI.

SuggestedRemedy

The agreement was not to shift the burden for host or module when we defined new values for VEC and VEO based on timing window ts=+/- 50 mUI. Unfortuntatly the VEC and VEO limits result in host that passed now will fail.
Propose new limits for VEO=8 mV and VEC=13.5 dB and see ghiasi_3ck_01_0421

Response *Response Status* **U**

REJECT.

Slide 3 to 9 of the following presentation were reviewed by the task force:
https://www.ieee802.org/3/ck/public/adhoc/apr21_21/ghiasi_3ck_adhoc_01a_042121.pdf

There is no consensus to change the VEC (max) or EH (min) values.

Cl **120G** *SC* **120G.3.4.1** *P* **247** *L* **17** # 20042

Ghiasi, Ali — Ghiasi Quantum/Inphi

Comment Type **TR** *Comment Status* **R** *TP4a SIT EH/VEC*

VEC limit of 12 dB and VEO limit of 10 mV results in well constructed host to fail, this was not the case prior to adding timing window of +/-50 mUI.

SuggestedRemedy

The agreement was not to shift the burden for host or module when we defined new values for VEC and VEO based on timing window ts=+/- 50 mUI. Unfortuntatly the VEC and VEO limits result in host that passed now will fail.
Propose new limits for VEO=8 mV and VEC=13.25 to 13.75 dB and see ghiasi_3ck_01_0421

Response *Response Status* **U**

REJECT.
[Editor's note: Changed page from 233 to 247 and subclause from 120G.3.1.5 to 120G.3.4.1]

Comment #39 proposed complementary changes to host output EH and VEC. However, the proposal in comment #39 was not adopted so no changes to the module input EH and VEC should be made.

See comment #39.

Cl **120G** *SC* **120G.3.4.1.1** *P* **248** *L* **1** # 20123

Ran, Adee — Cisco

Comment Type **TR** *Comment Status* **R** *TP2 additive noise*

In the module input stressed eye calibration procedure, "The stressed signal is generated by adding sinusoidal jitter, random jitter, and bounded uncorrelated jitter to a clean pattern, followed by frequency-dependent attenuation".

This signal does not necessarily represent a real host output, in which the EH and VEC can also be affected by additive noise (which is quite different from jitter in its effect on a receiver). Stressing the module with a high level of bounded uncorrelated jitter (which is not fully specified, and may create different stress for different DUTs) does not test its ability to operate with a noisy host.

Note that in a host transmitter it is often easier to control clock jitter than to reduce additive noise coupling from multiple sources in an ASIC.

Adjusting the VEC using additive noise, as done in the CR/KR/C2C tolerance tests, should at least be allowed instead of using "bounded uncorrelated jitter"; it may be preferable in some setups. For the time being, it is suggested as an alternative.

SuggestedRemedy

Add a wideband noise source to the diagram in Figure 120G–10, between the pattern generator and the frequency-dependent attenuator.

Add a description of the noise source to the text, with reference to 93C.1 (where noise source specification is defined) and setting f_NSD1 to 1 GHz, as in 163.9.3.4.

Add that calibrating the noise source level is an alternative method to adding BUJ for calibrating the EH and VEC.

Editorial license is suggested, but if necessary for accepting the comment I can provide candidate text before comment resolution.

Response *Response Status* **U**

REJECT.
Resolve using the response to comment #119.

Cl **162** *SC* **162.9.3** *P* **154** *L* **21** # 20166

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *CR port type*

The draft loss budget wastes over 3 dB in nearly every case.
The recommended maximum insertion loss allocation for the host traces plus BGA footprint and host connector footprint, of 6.875 dB, compares very poorly with C2M's host insertion loss up to 11.9 dB, making passive copper expensive and unattractive for a switch, while a full range of NICs can be made within only 3.75 dB. Server-switch links will get made with an asymmetric loss budget, so it would be better for the standard to regularise what will happen anyway. By the way, many server-switch links will be asymmetric anyway (different form factors at server and switch ends), and that's already allowed in this draft.
This change would also benefit CR switch-switch links because the shortest ports would get credit for their low loss.

SuggestedRemedy

As we have done for C2M, create two kinds of CR ports. Host loss allocations of 3.75 dB and 10 dB. Short can connect to short or long with same cable as today; long to long is not supported. Add entries in Clause 73 Auto-Negotiation to advertise short and long to the other end.
In Table 162-10, provide separate limits for Linear fit pulse peak (min).
In Table 162-14, provide separate rows for Test channel insertion loss: for testing the short host input the values for Test 2 are 10-6.875 = 3.125 dB higher (26.75 dB and 27.75 dB), while for the long host input the values for Test 2 are 6.875-3.75 = 3.125 dB lower (20.5 dB and 21.5 dB). No change needed for Test 1.
In 162A.4, provide two equations for each of IL_PCBmax and for ILHostMax and show them in Fig 162A-1 and 2. In 162A.5, provide two Value columns in Table 162A-1. Adjust figures 162A-3 and 4.

For discussion: should a "long" cable, 19.75+2*(6.875-3.75) = 19.75+6.25 = 26 dB max (maybe 3 m) be defined? A CR link could have no more than one of the three host, cable, and host being "long".

We could choose other names than "short" and "long" for the ports, possibly "short" and "medium" (as a C2M host can be "longer"), or A and B, somewhat like USB.

In 162.11.7.1.1, zp, representing the extra loss a host has above an MCB, could be made asymmetric but I believe that would not bring an improvement in accuracy.
There could be a third kind of CR port with 6.875 dB but this would not be useful for server-switch links, would be useful for only a subset of switch-switch links, for which passive copper is a subset anyway, so it doesn't seem worthwhile.

Response *Response Status* **U**

REJECT.

The following presentation was reviewed by the task force:
https://www.ieee802.org/3/ck/public/adhoc/apr28_21/dawe_3ck_adhoc_01_042821.pdf

The suggested remedy would require two or three different CR port types.

The assymetric-port approach was discussed early in this project.
Straw Poll #1 from the July 2018 Task Force meeting indicated strongest support for the current specification.
https://www.ieee802.org/3/ck/public/18_07/minutes_3ck_0718_approved.pdf

Based on discussion and straw poll 6 and 7, there is interest in exploring this proposal further. However, the proposal is not sufficiently complete at this time. A complete proposal and consensus is required.

Straw poll #6 (direction, chicago rule)
Straw poll #7 (direction, pick one)
I would support a new pair of CR port types with reduced host insertion loss limit on one end (e.g., NIC) and increased host loss limit on the other end (e.g., switch) similar to slide 7 of dawe_3ck_adhoc_01_042821.

Strawpoll #6
A: Yes 27
B: No 13
C: Need more information 29
D: Abstain 7

Straw poll #7
A: Yes 22
B: No 11
C: Need more information 11
D: Abstain 6

Cl **120G** *SC* **120G.3.2** *P* **240** *L* **9** # 20171

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *TP3 EH*

For a reasonably clean module (or test equipment in a host stressed eye test), the driver swing has to be aggressively reduced to deliver only 15 mV at near end, short mode. 120E has 70 mV, and the previous draft had 24 mV. Yet a host designer knows whether the host wants the short or long setting, and can usefully optimise for e.g. different crosstalk or noise or BER if given a reasonable signal strength. There is room to increase this weak signal without overloading the receiver.

SuggestedRemedy

Increase the eye height, short mode, from 15 mV to 18 mV

Response *Response Status* **U**

REJECT.

The resolution of comments #187 and #206 result in the differential peak-to-peak output voltage (max) value reduced from 900 mV to 600 mV for the short mode. There was no consensus to make the proposed change for this comment.

Cl **162** *SC* **162.11.6** *P* **169** *L* **27** # 20177

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *CA CM RL*

Relaxing the already very loose CM RL spec from 2 dB to 1.8 dB at all frequencies isn't justified. This spec becomes useless at the frequency when the MCB loss is 0.9 dB!

SuggestedRemedy

Restore it to 2 dB or use a frequency-dependent mask e.g. 1.8 + 0.01f

Response *Response Status* **U**

REJECT.

The basis for the change to the cable assmbly CM-to-CM RL spec from 2 dB to 1.8 dB was given in the following presentation.
https://www.ieee802.org/3/ck/public/21_01/champion_3ck_01a_0121.pdf

The commenter has not provided sufficient justification for the suggested remedy.

Cl **120G** *SC* **120G.5.2** *P* **252** *L* **25** # 20178

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *RR CTLE*

As a lot of the channel for TP4 far-end is known exactly, one would expect that a known subset of gDC, gDC2 combinations would be the only candidates to try. As for TP1a, I believe the strongest gDC and gDC2 should add to a constant.

SuggestedRemedy

For Continuous time filter, DC gain for TP4 far-end (gDC), change to a set of limits that depend on gDC2 in the same style as for TP1a, with the strongest gDC and gDC2 adding to a constant. The allowed values should be a subset of those for TP1a.

Response *Response Status* **U**

REJECT.

The comment does not provide sufficient justification to support any changes and the suggested remedy does not provide sufficient detail to implement.

Cl **120G** *SC* **120G.5.2** *P* **252** *L* **12** # 20179

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *RR CTLE*

By allowing stronger gDC with stronger gDC2, we can have up to 12 dB of peaking for gCD2 = -1 but up to 16 dB for gDC2 = -3 - yet we don't expect the maximum channel loss to vary like that.

SuggestedRemedy

For TP1a, change the second -12 to -11, and -13 to -10 (so the strongest "CTLE peaking" is 13).

Response *Response Status* **U**

REJECT.

The comment does not provide sufficient justification for the proposed change. It is not clear that the current specifications are harmful nor is there evidence that the proposed changes won't be harmful.

Cl **120G** *SC* **120G.5.2** *P* **253** *L* **23** # 20180

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *EH/VEC method*

This draft has a primitive rectangular eye mask (H = either EHmin or EA/VECmax), although it is described as a histogram. It's an inefficient/inaccurate way of measuring a signal quality vertically and provides weak and uncertain protection against too much jitter. This is worse with the higher VEC limit in the latest draft that allows worse and more varied signals, and is a particular concern for very short host channels (see Mike Dudek's work) that can have faster edges than higher loss ones.

SuggestedRemedy

Change from a 4-cornered mask with corners at t = ts+/-0.05, V = k +/-H/2 to a 10-cornered mask with corners at t = ts+/-0.05, ts+/-1/16, ts+/-3/32, V = k +/-H/2, k +/-H*0.4, k. k is VCmid, VCupp or VClow.
In case it's not clear, H is either EHmin or Eye Amplitude * 10^(-VECmax/20).
This simple scalable method can remain as the EH and VEC limits are revised. Scopes have been measuring with 10-sided masks for many years, it's not more difficult than a rectangular mask.

Response *Response Status* **U**

REJECT.
The currently methodology was chosen over an eye mask method like that being proposed in this comment.
See slide 3 of the following presentation was reviewed by the task force:
https://www.ieee802.org/3/ck/public/21_01/brown_3ck_04_0121.pdf
The comment does not provide sufficient justification to support the proposed changes.

Cl **120G** *SC* **120G.5.2** *P* **252** *L* **16** # 20183

Dawe, Piers
Nvidia

Comment Type **TR** *Comment Status* **R** *RR CTLE*

The limits for TP4 gDC, gDC2 should not be the same for short and long output modes.

SuggestedRemedy

Create separate limits for TP4 short and long output modes.

Response *Response Status* **U**

REJECT.

The comment does not provide sufficient justification to support any changes and the suggested remedy does not provide sufficient detail to implement.

Cl **162** *SC* **162.11.5** *P* **168** *L* **41** # 20201

Dudek, Mike
Marvell

Comment Type **TR** *Comment Status* **R** *CL-IL difference*

The differential to common mode conversion loss specification is very relaxed particularly at higher frequencies. As an example at 25GHz this specification is only approx 6dB more than the insertion loss. There is no specifiction for the common mode to common mode return loss of the Rx so all this common mode energy can be reflected back to the cable where through common mode to differential conversion it then becomes a differential signal interferer. Assuming this common mode to differential mode has approximately the same value as the differential to common mode conversion of approx 12.5dB this unwanted interferer is only 18.5dB below the wanted signal and will severely degrade the BER.

SuggestedRemedy

Add 10dB to this equation

Response *Response Status* **U**

REJECT.
The basis for a 10 dB tightening of the limit is not obvious in the stated comment and the correlation to the degradation of the BER is not provided.

Cl **120G** *SC* **120G.3.4.1.1** *P* **249** *L* **8** # 20224

Wu, Mau-Lin
MediaTek Inc.

Comment Type **TR** *Comment Status* **R** *module input SIT*

The frequency-dependent attenuation added from output of the pattern generator to TP1a is 18.2 dB, which is 16 dB channl loss with 2.2 dB for host transmitter package loss. However, 2.2 dB is too small a value for host transmitter package loss with 31 mm package trace length.

SuggestedRemedy

By leveraging what adopted in OIF CEI-112G-VSR-PAM4, propose to adopt the 19.5 dB value to replace 18.2 dB, where 3.5 dB representing host transmitter package loss is reasonable.

Response *Response Status* **U**

REJECT.

The comment does not provide sufficient evidence to make the proposed change.

Further work and a consensus proposal on this topic is encouraged.

Cl **120G** *SC* **120G.1** *P* **235** *L* **38** # 20234

Dawe, Piers Nvidia

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

Up to now, the optical PMD channels have not needed a very strong DFE, and the C2M loss (10 dB for C2M CAUI-4, 10.2 for 200GAUI-4 C2M, 16 for 400GAUI-4) is low enough that CR and KR PMDs don't need a very strong DFE when used as C2M. Therefore, we never have precoding on C2M at 50G/lane - simple. At 100G/lane, links such as active copper cables will benefit from a very strong DFE in the receiver in the cable end that's receiving from a higher loss in the cable. 802.3 enables such active cables via the C2M specs; up until now there was nothing more to say, so they don't get a mention in 802.3. Adding precoding after the signal has been serialised is best avoided, so it should be added in the host, so for the first time, there is something that 802.3 should do specifically about active cables.

SuggestedRemedy

Allow optional precoding abilities in 100G/lane C2M transmitters and receivers in the host. Add MDIO registers to advertise these abilities and to enable them.

Response *Response Status* **U**

REJECT.

Precoding if used is added and removed by the PMA at each end of a physical link as necessary. Similarly, an active cable can add precoding at the transmitter at one end and remove the precoding at the other end.
Precoding must be enabled (or disabled) on both Tx and Rx in the same direction; this is coordinated using training for CR/KR or by station management for C2C. Applying precoding internally within an active cable is still possible.

There is no consensus to implement the proposed.

Cl **162** *SC* **162.11.7** *P* **171** *L* **31** # 20235

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *CA COM DFE*

The spec allows a channel to have its COM calculated with 9 taps in the range 13 to 24 clipped at +/-0.05 - which means that the channel's pulse response could be a little worse than +/-0.05 for all these 9 taps. That's a very bad cable! and not likely to get made. We don't need to provide all the receiver power and complexity to cope with it.

SuggestedRemedy

Use another DFE root-sum-of-squares limit for positions 13-24. Similarly in 163, but as 163 specifies the complete channel while 162 uses clean synthetic host traces, the limit might differ.

Response *Response Status* **U**

REJECT.
The suggested remedy does not provide sufficient evidence that this is an issue and that the proposed change would not cause new issues.