Cl **162** *SC* **162.9.3** *P* **170** *L* **32** # 87

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *CR loss budget*

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

As in dawe_3ck_01a_0721.pdf:
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 the training control field to advertise A, B or C to the other end.
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: 9.5-6.875 = 2.625 dB higher (26.25 dB to 27.25 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.
Add MDIO registers to report local and remote host ability to station management, for inventory and diagnostics.

Response *Response Status* **U**

REJECT.
This comment is a restatement of comment #92 against D2.1, which was rejected by the task force. This new comment provides only minor changes to the suggested remedy. A related straw poll (#10) indicated strong opposition to adopting this proposal therefore there was no consensus to make the proposed changes.
July 2021 Straw Poll #10 is reproduced here for reference…
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* **189** *L* **38** # 89

Dawe, Piers Nvidia

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

As in previous comments: this common mode return loss spec RLcc becomes useless at the frequency when the MCB loss is 1.8/2 dB, which is only 8.5 GHz. We need a common mode return loss spec to stop large common-mode voltages building up through multiple low-loss reflections. The revised proposed remedy for D2.1 comment 79 seems OK: 1.8 dB 0.5<= f <= 4 GHz, 1.4+0.1*f dB 4< f <= 30 GHz. The 30 GHz fmax allows margin for real-world coax-PCB transitions (although the mated compliance boards are specified >=3 dB to 50 GHz); the cable itself should pass this comfortably because it is insulated from the test by the MCB loss.

SuggestedRemedy

Use a frequency-dependent mask 1.8 dB 0.5<= f <= 4 GHz, 1.4+0.1*f dB 4< f <= 30 GHz. f is in GHz. Similarly for Tx, Table 162-11, 162.9.3.6.

Response *Response Status* **U**

REJECT.
This comment is a restatement of D2.1 comment #79.
The suggested remedy does not provide sufficient additional justification to support the change to the draft.

Per straw poll #6, there was no consensus to make the proposed changes.

However, there was concern that the limits should be tightened. Further work and consensus is required.

Straw poll #6 (decision)
I support adopting the changes in comment #89 suggested remedy.
Yes: 11
No: 19

Cl **162** *SC* **162.11.7** *P* **191** *L* **39** # 90

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *COM DFE bgmax/min (CC)*

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. I know of only example channel with a tap like 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. In the case of Bch2_b2p5_7_t, reducing |bmaxg| from 0.05 to 0.03 increases COM by less than 0.1 dB, and the channel still passes comfortably. In this example, there were no taps that would be affected by reducing +ve bgmax from 0.05 to 0.03; one -ve tap was limited.

SuggestedRemedy

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

Response *Response Status* **U**

REJECT.
This is a restatement of comment #95 against D2.1 which was rejected by the task force due to insufficient supporting evidence. Some new information on the analysis of one channel is provided, but this is insufficient evidence to support the proposed changes.
[Editor's note: CC: 162, 163]

Cl **162** *SC* **162.11.7** *P* **191** *L* **38** # 91

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *COM DFE RSS (CC)*

The spec allows a cable 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. A limit of 0.045 works well with Bch2_b2p5_7_t. Similarly in 163.

Response *Response Status* **U**

REJECT.
This is a restatement of comment #96 against D2.1 which was rejected by the task force due to incomplete remedy and insufficient analysis. This new comment provides some new, but unsubtantiated information.
[Editor's note: CC: 162,163]

Cl **120G** *SC* **120G.3.2** *P* **264** *L* **11** # 93

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *MO 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 than FE for a well set up output and the spec should reflect that. Host designers know their own loss and medium-loss hosts can take advantage of a better signal that cost the module nothing.

SuggestedRemedy

Change the eye height, long near end, so that it is 3 dB above long far end, e.g. 15 mV (far) and 21 mV (near) if long far is not changed. 3 dB is about half the loss from long near end to long far end, so long far end remains the harder one to meet.

Response *Response Status* **U**

REJECT.

This comment is a restatement of D2.1 comment #98, for which there was no consensus to make the proposed changes.

The intent of specifications is to enforce what is necessary not what is possible. However, as this comment states, a long-mode host might be able to take advantage of the extra eye height.

There is insufficient evidence to make the proposed changes.

Cl **120G** *SC* **120G.5.2** *P* **279** *L* **43** # 95

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *EO mask*

The Gaussian weighting has the effect of destroying the histogram width, allowing bad fast eyes to pass, while giving the impression that the histogram width still applies. With a weighting standard deviation of 0.02 UI, the eye height is measured at around +/-0.03 UI rather than the +/-0.05 UI in the previous draft. Compare 120E with ESMW of 0.2 or 0.22 UI.

SuggestedRemedy

Remove the Gaussian weighting and set the eye height and VEC limits (which need revision anyway) appropriately.

Response *Response Status* **U**

REJECT.

The current method of determining eye height and VEC using a weighted window was introduced in D2.2 based on approved D2.1 comment #39. A final straw poll indicated acceptance of the response with a ratio (yes:no) of 21:11.

Per straw poll #9 and #10 there is no consensus to change the measurement method.

--- the following added 2021/10/4 ---

Straw poll #9 (pick one)
Straw poll #10 (chicago)
(direction)
I support the following method of determining eye height and VEC:
A: weighted window per Draft 2.2 (no change)
B: weighted window per Draft 2.2, except increase standard deviation
C: unweighted window per Draft 2.1 (perhaps with different width)
D: mask per D2.2 comment #101
#9: A: 17 B: 5 C: 6 D: 2
#10 A: 22 B: 12 C: 7 D: 3

Cl **120G** *SC* **120G.3.2.2.1** *P* **265** *L* **46** # 97

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *MO SI 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 more than far minus near for long. As real host channels are not exactly like the theoretical reference host channel and host makers hate avoidable precision, measurement and record-keeping, 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.

This comment is a restatement of D2.1 comment #102 for which
there was no consensus to make a change. However, the response notes that there may be some benefit to explore this further.

However, no further analysis or significant additional justification has been provided.

Further discussion indicated there are concerns with making the proposed changes.

There is no consensus to make the proposed changes.

Cl **120G** *SC* **120G.5.2** *P* **277** *L* **38** # 98

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **A** *EO RR gdc*

The limits for TP4 gDC, gDC2 should not be the same for short and long output modes. Obviously, different channels will need different CTLE settings. Obviously, CTLE settings that only signals outside what the spec is designed for use, should be excluded, to make implementers set up their product correctly.

SuggestedRemedy

Create separate limits for TP4 short and long output modes, so 4 sets for TP4+, in the style of TP1a. If you don't have any better numbers, create them anyway with the same numbers in each set - but see another comment.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

This comment is a restatement of D2.1 comment #103 and D2.0 comment #183, which were rejected on the basis of providing insufficient justification and detail.

This comment provides expanded justification.

Slides 7, 8, 11, 12 of the following presentation for a representation we reviewed by the task force.
https://www.ieee802.org/3/ck/public/21_09/kochuparambil_3ck_01b_0921.pdf

Slides 7, 8, and 11 of kochuparambil_01b provide a view the suggested remedy if implemented.

There was no consensus to provide separate gdc specifications for long and short modes.

However, some related editorial changes as follows are an improvement to the draft.

Update style of the TP4 gdc specifications in Table 120G-11 as shown in the referenced slide 12 of kochuparambil_01b. Include similar changes for g_dc2.

Implement with editorial license.

Cl **120G** *SC* **120G.5.2** *P* **277** *L* **46** # 99

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *EO 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.

SuggestedRemedy

For Continuous time filter, DC gain for TP4 far-end (gDC), change to sets of limits that depend on gDC2 in the same style as for TP1a. The allowed values should be subsets of those for TP1a. For TP4 long far end, use minimum gDC 1 dB higher than allowed for TP1a; for TP4 short far end, 3 dB higher than for TP1a.

Response *Response Status* **U**

REJECT.

This comment is a restatement of D2.1 comment #104 and D2.0 comment #178, which were rejected on the basis of providing insufficient justification and detail.

This comment provides no new justification, but does provide more details for implementation.

Cl **120G** *SC* **120G.5.2** *P* **277** *L* **32** # 100

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *EO RR bbmax*

My recent simulations don't use gDC as strong as the table allows, but occasionally, the first DFE tap hits the limit of 0.4

SuggestedRemedy

Increase bbmax(1) from 0.4 to 0.5, increase the minimum for gDC at TP1a and TP4 long far end.

Response *Response Status* **U**

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3ck D2.2 and D2.1 or the unsatisfied negative comments from previous drafts. Hence it is not within the scope of the recirculation ballot.

The comment provides only anecdotal evidence for the bbmax change.

For related changes to gdc see responses to comments 72 and 99.

There is no consensus to make the proposed changes to bb_max.

Cl **120G** *SC* **120G.5.2** *P* **279** *L* **6** # 101

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *EO mask*

This draft has a weighted rectangular eye mask spec with mask height = max(EHmin, EA/VECmax) and effective mask width ~2x0.03 UI, although it is described as a histogram 2x0.05 UI wide. Measuring a diamond eye with a rectangular mask provides weak and uncertain protection against too much jitter; de-weighting the sides of the histogram weakens it further; the effective BER criterion is hard to establish but seems to be around 1e-4, not 1e-5 as intended.
We need an eye mask that's more eye shaped, so that a higher proportion of the samples near the boundary are measured at full weight and contribute properly to the measurement. Eye mask measurement with a 10-sided mask has been pre-programmed into scopes for about 20 years, we should use established tools and methods where they work well.

SuggestedRemedy

Change from a 4-cornered weighted mask with corners at t = ts+/-0.05, V = y +/-H/2 to a 10-cornered unweighted mask with corners at t = ts+/-1/16, ts+/-0.05, ts+/-3/32, V = y +/-H/2, k +/-H*0.4, y. y is near VCmid, VCupp or VClow (vertically floating, as in D2.2).
H is max(EHmin, Eye Amplitude * 10^(-VECmax/20)). Eye Amplitude is AVupp, AVmid or AVlow, as in D2.2.
This simple scalable method can remain as the EH and VEC limits are revised.

Response *Response Status* **U**

REJECT.

This comment is a restatement of D2.1 comment #106 and D2.0 comment #180 for which there was no consensus to make the proposed changes. No new evidence or consensus has been provided.

Resolve using the response to comment #95.

Cl **162** *SC* **162.9.3.4** *P* **174** *L* **47** # 102

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *TX EOJ*

Having alternative normative patterns to measure one thing when the choice makes a difference, adds cost because the test has to be done both ways (if one way passes and the other fails). Also, the spec limit was relaxed from 0.019 UI to 0.025 to allow for PRBS13. We understand that the result would look better with PRBS9. There is no requirement to generate PRBS9.

SuggestedRemedy

Make PRBS13 normative, as usual. Use a different set of PRBS13Q pattern symbols used for jitter measurement vs. Table 120D-4 to reduce the pattern dependency issue.

Response *Response Status* **U**

REJECT.
This is a restatement of comment #109 against D2.1 which was rejected by the task force (insufficient remedy and lack of consensus to make the change). The comment does not provide new data or analysis to support it.

Cl **162** *SC* **162.9.3.4** *P* **174** *L* **49** # 103

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *TX EOJ*

We know that CRU corner frequency makes a difference to EOJ measurement. Allowing an unbounded "4 MHz or anything you like that's lower" is very bad: how many attempts must the tester try before he can fail a bad part?

SuggestedRemedy

Pick a single definitive CRU corner, e.g. 1 MHz or 2 MHz. Add informative NOTE saying that we expect that if it passes with the usual 4 MHz, it would also pass with the lower corner frequency.

Response *Response Status* **U**

REJECT.
This is a restatement of comment #109 against D2.1 which was rejected by the task force (insufficient remedy and lack of consensus to make the change). The comment does not provide new data or analysis to support it.

Cl **120G** *SC* **120G.3.3.5.2** *P* **270** *L* **22** # 148

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *HI SI method*

The host stressed input signal is emulating a module so must obey the same rules. VEC and eye height must be in spec for both near end and far end. The signal should be adjusted to minimise VEC for both, or possibly to minimise VEC for far end while keeping in spec at near end. The eye height should match the target at far end and be graeter at near end.

SuggestedRemedy

This procedure needs road-testing before the draft can be said to be "without technical issues". In the meantime, add text to the draft to explain more fully what the procedure is.

Response *Response Status* **U**

REJECT.

Item g) instructs that the eye height of the smallest eye match the target value in Table 120G-8. Table 120G-8 provides only one value to be used for both near-end and far-end measurements.
Item g) instructs that VEC is within the limits in Table 120G-8. Table 120G-8 provide only one range (with maximum and minimum) to be used for both near-end and far-end measurements.
The module output specifications for eye height and VEC are the same for near-end and far-end.

The comment does not provide sufficient evidence to support the proposed changes. The suggested remedy does not provide sufficient detail to implement.

Cl **162C** *SC* **162C.1** *P* **306** *L* **10** # 157

Ghiasi, Ali Ghiasi Quantum/Inphi

Comment Type **TR** *Comment Status* **R** *MDI pins table*

Per unsatisfied comment from D2.2.
Table 162C-3 needs to be better organized

SuggestedRemedy

An improved and beter organized table will be submited as ghiasi_3ck_01_0921.pdf

Response *Response Status* **U**

REJECT.

The following related presentation was considered by the task force:
https://www.ieee802.org/3/ck/public/21_09/ghiasi_3ck_01_0921.pdf

There is no consensus to make the proposed change.

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* **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 **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.

Cl **120G** *SC* **120G.3.2.1** *P* **254** *L* **51** # 21002

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **R** *MDIO 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.3.1** *P* **250** *L* **12** # 21046

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 **162** *SC* **162.9.3** *P* **163** *L* **18** # 21092

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** # 21094

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** # 21095

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** # 21096

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** # 21097

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** # 21098

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.5.2** *P* **265** *L* **16** # 21103

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** # 21104

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* **266** *L* **23** # 21106

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** # 21115

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.]