Cl **162** *SC* **162.11.7** *P* **188** *L* **46** # R2-16

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *COM parameter*

93A.1.1 says "It is recommended that the scattering parameters be measured with uniform frequency step no larger than Delta f from a start frequency no larger than fmin to a stop frequency of at least the signaling rate fb". But the test fixtures are defined to 50 GHz, and other specs such as RLdc are defined to 40 GHz. 93A.1.5 says "the filtered voltage transfer function may need to be extrapolated (both to DC and to one half of the sampling frequency) for this computation. The extrapolation method and sampling frequency should be chosen carefully to limit the error in the COM computation."

For cable COM, there is the sinc function for NRZ signalling + driver Gaussian filter Tr + minimum ~16 dB cable loss even at 40 GHz + PCBs + packages + Butterworth filter + extra pole of the CTLE. The result is quite tolerant to the extrapolation.
For ERL, there is sinc function, Tr, Butterworth filter, and Tukey filter (17.7 dB at 50 GHz), and twice the test fixture trace loss. There can be very little energy between 50 GHz and 53.125 GHz where the Tukey filter cuts off.
Extrapolating RL (as opposed to IL) is not reliable anyway.

SuggestedRemedy

To ensure consistency between measurements, define the maximum measurement frequency for COM as 50 GHz, then COM is calculated with careful extrapolation as mentioned.
Define the maximum frequency for ERL as 50 GHz, with no extrapolation.
Both these could be achieved by inserting a row for fmax, 50 GHz, in the tables for COM parameter values.
Apply to 162 and 120G which rely on test fixtures with connectors that are defined to 50 GHz.
Apply to 163 and 120F ERL also because 50 GHz is a natural break point for network analysers.
Unless we find that doing so opens a hole in the spec, apply to 163 and 120F COM also.

Response *Response Status* **U**

REJECT.

This comment is a restatement of Draft 3.1 comment R1-52 and of Draft 3.0 comment I-186.
The resolution to these comments is provided in the following files:
https://www.ieee802.org/3/ck/comments/draft3p0/8023ck_D3p0_final_closedcomments_sortedByNumber.pdf
https://www.ieee802.org/3/ck/comments/draft3p1/8023ck_D3p1_final_closedcomments_sortedByNumber.pdf

This comment provides no new evidence to support the proposed changes.

There is no consensus to make the proposed changes.

Cl **120G** *SC* **120G.5.2** *P* **275** *L* **50** # R2-17

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *HO/MO EW*

As we know, this Gaussian "weighting" function de-weights the sides of the histogram, allowing worse eye width (jitter) than otherwise. As healey_3ck_01a_1020 shows, for the same VEC, ESMW varies across channels by at least 130 mUI, plus some more for driver output edge rate. As e.g. dudek_3ck_01_0921 slide 7 shows, there can be a great variety of eyes for only slightly different channels. It turns out that unsymmetric eyes are possible (significantly different to left and right) - see presentation. The draft spec skews the spec to passing signals with relatively bad eye width, which endanger the link BER, while failing signals with usable VEC and eye height and better eye width.

We need better control of eye width, as has been pointed out in D3.0 comments I-107, I-108, I-115, I-116, I-211, I-212 and R1-55, with two clear alternative remedies proposed: the 10-sided mask or explicit ESMW limits.

SuggestedRemedy

Add ESMW spec limits:
Host output and module stressed input >=120 mUI;
Module output and host stressed input >= 130 mUI.
ESMW is defined around ts in the same way that ESMW is defined around Tcmid in 120E.

The reason for host spec being less than module is that almost all the bad stuff is in the host measurement, but not all the host channel and package impairments are in the module measurement, even "far end".
The limits in 120E are host 0.22 UI, module near 0.265 UI, module far 0.2 UI (with a less capable equaliser), so these specs are allowing much worse eyes than 120E, but not totally out of control.

Response *Response Status* **U**

REJECT.

This comment is a restatement of Draft 3.0 comments I-107, I-108, I-115, I-116, I-211 and I-212, and Draft 3.1 comment R1-55. The resolution to these comments is provided in the following files:
https://www.ieee802.org/3/ck/comments/draft3p0/8023ck_D3p0_final_closedcomments_sortedByNumber.pdf
https://www.ieee802.org/3/ck/comments/draft3p1/8023ck_D3p1_final_closedcomments_sortedByNumber.pdf

These comments were closed on the basis of no consensus to make the related changes. The result of straw poll #11 recorded in the response to comment I-211 (see https://www.ieee802.org/3/ck/comments/draft3p0/8023ck_D3p0_final_closedcomments_sortedByNumber.pdf) indicated consensus to not make these proposed changes.

The following related presentation was reviewed by the task force:
https://www.ieee802.org/3/ck/public/22_06/dawe_3ck_01a_0622.pdf

This new comment provides an alternative suggested remedy and the presentation provides new evidence.

Per straw poll #7, there is no consensus to make the proposed changes.

Straw poll #7 (direction)
I support adding an ESMW specification for C2M.
Yes: 8
No: 16

Cl **120G** *SC* **120G.3.2** *P* **261** *L* **11** # I-187

Dawe, Piers J G NVIDIA

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

On one hand: the eye height measurement method is very inaccurate, host receivers that implement CR can cope with much smaller eye height than this, VEC is much more important. Receiver noise is already in the measurement, C2M drivers are traditionally 900/1200 as strong as CR/KR drivers, and the end-to-end loss is lower by a much larger ratio. So a small EH is acceptable.
On the other hand: if the eye height limit is the same at near end as at far end, there is huge margin at near end and the implementer can optimise beyond far end, only limited by the NE VEC spec, while we want modules to be set up consistently, for the full range from near to far. NE and FE EH naturally differ, and the spec should reflect that. Also, host designers know their own loss and low-loss hosts (NICs) can take advantage of a naturally larger signal that cost the module nothing. This applies to both the short and long modes.

SuggestedRemedy

Change the far end eye height so that it is 2 dB below near end: if near can remain at 15 mV, far becomes 12 mV. Far end remains the one with less margin, that the implementer should tune the module for.

Response *Response Status* **U**

REJECT.

The comment makes reference to the capabilities of a CR SERDES. Annex 120G is specifying C2M recievers and transmitters. Although it is true that the host might have a CR-capable SERDES that may not be universally the case. Note that there are different host channel budgets for CR and C2M.

The comment does not provide sufficient justification for the proposed changes. Analysis is required to demonstrate the need.

There is no consensus to make the proposed changes.

Cl **120G** *SC* **120G.3.2** *P* **261** *L* **11** # I-188

Dawe, Piers J G NVIDIA

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

The module output eye height and VEC have to comply at both near end and far end, and depending on the cleanliness of its signal, a module can be tuned to either end or somewhere in the middle, or even somewhere outside the range. The host stressed input signal is tuned to far end, only, so the host isn't required to receive those other tuning choices. This is inconsistent and a serious flaw in the spec. Yet we would rather not have multiple host stress tests, nor require the host to receive unnecessary and sub-optimal signal tunings, so we need to make sure that modules are tuned correctly.

SuggestedRemedy

Tighten the equaliser limits for module output so that modules are tuned consistently across the industry. Because the channel losses in short and long mode testing are significantly different, in Table 20G-11 use separate gDC limits for short and long mode (see other comments). To discourage module implementers from mis-tuning modules so they are optimised significantly beyond the far end, in Table 120G-3, ensure that each near end VEC is 0.5 dB less (better) than its corresponding far end VEC, and the far end EHs are 2 dB less than the corresponding near end EHs. Note other comments that address what these values should be.

Response *Response Status* **U**

REJECT.

The comment provides insufficient evidence evidence that the proposed changes are necessary or improve the interoperability.

Cl **120G** *SC* **120G.5.2** *P* **275** *L* **27** # I-206

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *MO gDC values*

The limits for TP4 gDC, gDC2 should not be the same for short and long output modes. The range of losses in a module is much less than the range of losses of the four reference host channels. So, obviously, different channels will need different CTLE settings. Obviously, CTLE settings that represent signals outside what the spec makes a host capable of receiving in a particular mode, should be excluded, to make module 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. See other comments.

Response *Response Status* **U**

REJECT.

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* **275** *L* **34** # I-208

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *MI gDC values*

The weakest (max, least -ve) gDC + gDC2 is -2 for TP1a, -2 for TP4 near end, -3 for TP4 far end and -10.5 for module stressed input high loss. There is about 10 dB loss difference between short near end and long far end, but 1 dB difference in max gDC + gDC2 which is far too little. It looks like TP4 far end (-9 to -2 in the draft) is out of step, with a much wider range than TP4 near end. TP4 LONG far end should never use this wide range as most of the channel loss is fixed. We should not be encouraging modules to try to do a job the host receiver does better, and we want modules to be set up consistently so that the short/long mode choice means something.
Also, if we include an allowance for host transmitter package loss for the host stressed input test, it would make sense to include the same allowance for far-end module output specs.

SuggestedRemedy

Impose a max gDC + gDC2 limit of -5 for TP4 long far end, e.g. with gDC, gDC2 ranges in the same style as TP1a:
Range for gDC2 = 0 -9 to -5
Range for -1 <= gDC2 < 0 -9 to -4
Range for -2 <= gDC2 < -1 -9 to -3
Range for -3 <= gDC2 < -2 -9 to -2

Response *Response Status* **U**

REJECT.

There is some agreement with the direction of the proposal but further analysis is required to determine appropriate values.

Cl **120G** *SC* **120G.5.2** *P* **275** *L* **34** # I-209

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *MO gDC values*

As a most 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 short 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.
See another comment for TP4 long far end.
For TP4 short far end, change from -9 to -2, to:
Range for gDC2 = 0 -7 to -3
Range for -1 <= gDC2 < 0 -7 to -2
Range for -2 <= gDC2 < -1 -7 to -2
Range for -3 <= gDC2 < -2 -7 to -2

Response *Response Status* **U**

REJECT.

There is some agreement with the direction of the proposal but further analysis is required to determine appropriate values.

Cl **120G** *SC* **120G.5.2** *P* **277** *L* **6** # I-211

Dawe, Piers J G NVIDIA

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

This draft has a (de-)weighted rectangular eye mask spec with mask height = max(EHmin, EA/VECmax) and effective mask width ~2x0.03 to 2x0.035 UI, although it is described as a histogram 2x0.05 UI wide. This is too narrow; compare 120E with ESMW of 0.2 or 0.22 UI. It's half as wide as TDECQ with histograms extending to +/-0.07 UI.

This de-weighted histogram might have worked if there had been a guarantee that no host or module would ever produce a fast, highly jittered eye, but we don't have that guarantee. Work needs to be done to repair the hole in the spec.

See healey_3ck_01a_1020 slide 6, orange dots for +/-0.025 UI which is the closest to the current draft. For VEC of 10 dB, EW can be anywhere in the range 160 to 290 mUI: an almost 2:1 range. Driver risetime is not reported; if it is always the COM default slowest-reasonable 7.5 ps, then even worse EW is possible with faster or peaked drivers. This is too much worse than 120E. As the plot shows, a wide range of eye widths are possible, so we don't need to allow the worst ones by an oversight.

De-weighting the sides of the histogram with flat top and bottom, rather than chamfering the corners, means that infringing the corners by a mile is counted the same as infringing by an inch, which is bad.
Most of the weight of samples is in the middle of the eye which is a waste of measurement time; we know the corners will fail first so we should measure them, not the middle Hence the 2-offsets approach of TDEC and healey_3ck_01a_1020.
The effective BER criterion of the (de-)weighted mask seems to be around 1e-4, not 1e-5 as before.

The distribution of repeated measurements is very skewed.

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.

The 10-sided mask controls the eye on the diagonal more strongly than the rectangular uniform histogram/mask because hits are collected over the time of the chamfer, rather than just in corners. The de-weighted rectangular histogram controls the eye on the diagonal more weakly than the rectangular uniform histogram/mask because hits are collected just in corners, and de-weighted.

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, y +/-H*0.4, y. y is near VCmid, VCupp or VClow (vertically floating, as in D3.0).
H is max(EHmin, Eye Amplitude * 10^(-VECmax/20)). Eye Amplitude is AVupp, AVmid or AVlow, as today.

This simple scalable method gives VEC results 0.5 to 1 dB more optimistic than the unweighted rectangular mask. It can remain as the EH and VEC limits are revised in the light of experience.

Response *Response Status* **U**

REJECT.

Straw polls #8 and #9 indicate strong consensus to continue with a weighted window approach. Straw polls #10 and #11 indicate strong consensus to continue with the currently specified weighting function.

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

Straw poll #8 (chicago rules)
Straw poll #9 (choose one)
I support the following direction of the eye opening specification method:
A. weighted window per Draft 3.0 (as is or with some improvements)
B. revert to uniform weighted window per D2.1 (D3.0 comment #212)
C. 10pt mask per D3.0 comment #211
#8 A: 31 B: 12 C: 6
#9 A: 27 B: 5 C: 1

Note: Straw poll #8 and #9 are the same question and answers except #8 is chicago rules (pick any) and #9 is choose one.

Straw poll #10 (chicago rules)
Straw poll #11 (choose one)
To address eye width issues expressed, I support the following method to modify the weighted window:
A. no change
B. “wider” weighting mask (e.g., larger sigma, alternate distribution shape)
C. add jitter specification
D. add eye width specification (i.e., per D3.0 comments 107, 108, 115, 116)
#10 A: 26 B: 15 C: 9 D:9
#11 A: 19 B: 5 C: 3 D: 4

Note: Straw poll #10 and #11 are the same question and answers except #10 is chicago rules (pick any) and #11 is choose one.

Cl **120G** *SC* **120G.5.2** *P* **277** *L* **6** # I-212

Dawe, Piers J G NVIDIA

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

The Gaussian weighting has the effect of destroying the histogram width, allowing bad fast eyes to pass, while failing less bad slow eyes. It gives the false impression that the histogram width still applies. With a weighting standard deviation of 0.02 UI, the eye height is measured at around +/-0.035 UI rather than the +/-0.05 UI with the unweighted histogram - depending on eye shape. Compare 120E with ESMW of 0.2 or 0.22 UI, and TDECQ with histograms extending twice as wide, to +/-0.07 UI.
This weighting is equivalent to relaxing the VEC spec by 1.5 to 2 dB - but it depends on the eye shape, it weakens the spec most for the worst-shaped eyes, which is bad. It applies a worse BER criterion than the 1e-5 intended.

SuggestedRemedy

Remove the Gaussian weighting and set the eye height and VEC limits (which need revision anyway) appropriately. ghiasi_3ck_01_0721, which was not given the presentation time it deserved, says that the minimum eye height in particular needs to be reduced for TP1 and TP4 far end.

Response *Response Status* **U**

REJECT.

There is no consensus to make the proposed changes.

For details, see the reponse to comment i-211.

Cl **120G** *SC* **120G.5.2** *P* **275** *L* **50** # R1-55

Dawe, Piers J G NVIDIA

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

As noted, this weighting function skews the spec to passing signals with relatively bad eye width, whether from jitter or other cause, which endanger the link BER, while failing signals with usable VEC and eye height and better eye width.

SuggestedRemedy

Pick one of the proposed solutions and fix the problem. Notice that the apparent VEC and EH numbers are likely to change in step.

Response *Response Status* **U**

REJECT.

This comment is a restatement of D3.0 comments i-211 and i-212 recorded in the following comment report:
https://www.ieee802.org/3/ck/comments/draft3p0/8023ck_D3p0_final_closedcomments_sortedByNumber.pdf

No further evidence nor any alternate remedies are provided.

Straw poll #11 (recorded in the response to comment i-211) indicated consensus to make no changes to the measurement method.