

IEEE P802.3cw D2.4 400 Gb/s over DWDM systems 4th Working Group recirculation ballot comments

CI 116 SC 116.1.3 P 27 L 22 # 20419

Dawe, Piers Nvidia

Comment Type TR Comment Status R

The manipulations described in this draft don't describe a BASE-R "native Ethernet"; rather, they are like 10GBASE-W. An Ethernet signal is packed into a telecoms wrapper (then, based on SONET, here, based on OTN).

The combination is clumsy and messy. Starting from Ethernet building blocks, one would not engineer it like this. I understand that the rationale is because those designs were already there, and the cost of a clean design was thought to outweigh the inefficiencies of this scheme. But that calls "broad market potential" into question.

800G coherent will affect the market for this.

SuggestedRemedy

I can think of three options:

Redo Clause 155, leaving out GMP and FAW and simplifying the training sequence and pilot sequence to make an Ethernet PHY;

Cancel this project, and encourage those interested to feed their learnings into OIF's "400ZR" maintenance;

Rename this PHY to 400GBASE-ZW, which is more honest and leaves the "400GBASE-ZR" name available to any future native Ethernet PHY, should the broad market potential be found.

Response Response Status U

REJECT.

No consensus within the CRG to change the name of the 400GBASE-ZR PHY

CI 116 SC 116.1.3 P 33 L 12 # 21280

Dawe, Piers Nvidia

Comment Type TR Comment Status R

As is made clear by the non-BASE-R Table 116-5a and 116.4.3 and 116.4.4, "400GBASE-ZR" is not BASE-R. However, the "R in the name implies that it is, which causes confusion. Clause 155 describes a "WAN PHY" like 10GBASE-W: an Ethernet signal is carried in a telecoms wrapper (then, based on SONET, here, based on OTN). Also, misnaming this spec blocks the way for a future native BASE-R 400G Z class PHY. The name "400GBASE-ZW", while correct, doesn't flow very easily, but "400GBASE-Z" avoids the misrepresentation and provides a cleaner name.

SuggestedRemedy

Change "400GBASE-ZR" to "400GBASE-Z" throughout.

Response Response Status U

REJECT.

Changing the name from 400GBASE-ZR was previously considered in D2.0 comment #419 (https://www.ieee802.org/3/cw/comments/D2p0/8023cw_D2p0_comments_final_by_clause.pdf) and there was no consensus to make a change.

The comment does not provide sufficient justification to support the suggested remedy.

There was no consensus to make a change.

IEEE P802.3cw D2.4 400 Gb/s over DWDM systems 4th Working Group recirculation ballot comments

CI 155 SC 155 P 39 L 1 # 21278

Dawe, Piers Nvidia

Comment Type TR Comment Status R

This PCS/PMA is over-complicated and messy. We would not engineer it like this now (see nicholl_3dj_optx_01_230413 for a small step in the right direction, and maniloff_3dj_01a_2303 for an example of how to do coherent cleanly). OIF's so-called "400ZR" has had a draft since 2018, was issued in 2020 and revised last year. 800G coherent is coming in OIF and P802.3dj, which will take much of the market away. This P802.3cw project is on about its ninth draft and still the actual specifications are vague and incomplete, the previous draft was issued 8 months ago; not the usual two-monthly cadence we expect from an active project and an enthusiastic group. The moment for doing this spec in 802.3 has passed, it doesn't add significantly to 400ZR, and I observe there are not enough active participants in P802.3cw to justify it.

SuggestedRemedy

Cancel this project.
Encourage those interested to feed their learnings into OIF's "400ZR" maintenance.
Re-use relevant parts of the draft in P802.3dj when the time comes.

Response Response Status U

REJECT.

In the D2.0 review, 582 comments from 22 commentors were received which shows continued interest in the project.

In the D2.1 review, 290 comments from 13 commentors were received which shows continued interest in the project.

No consensus to cancel the project at this time.

CI 155 SC 155 P 39 L 1 # 21281

Dawe, Piers Nvidia

Comment Type TR Comment Status R

This PCS/PMA is way too complicated for just a "directive" specification, and much more complicated than the mainstream 256/257/RS-FEC. We need examples, as in Annex 91A, RS-FEC codeword examples, or Annex 76A, FEC Encoding example.
If no-one is willing to provide them, we don't have a quorum to complete the project.

SuggestedRemedy

Create examples of e.g. FEC and other blocks before and after coding. Smallish ones can go in the document, all can be uploaded to the directory that IEEE provides for these things.
Alternatively, cancel the project.

Response Response Status U

REJECT.

No data was provided for the editors to be able to implement this change. Contributions of such material would be welcomed.

Regarding the project cancel proposal see response to comment #278.

CI 155 SC 155 P 42 L 4 # 2318

Dawe, Piers Nvidia

Comment Type TR Comment Status R

D2.1 comment 281: this PCS/PMA is way too complicated for just a "directive" specification. We need examples, as in Annex 91A, RS-FEC codeword examples, or Annex 76A, FEC Encoding example, or the OIF test vectors for 400ZR.

SuggestedRemedy

Publish examples of e.g. FEC and other blocks before and after coding. Smallish ones can go in the document, all can be uploaded to the directory that IEEE provides for these things.
If no-one does the work needed, cancel the project.

Response Response Status U

REJECT.

As noted by commentator, this issue was previously raised in D2.1 comment #281 which was rejected with the response "No data was provided for the editors to be able to implement this change. Contributions of such material would be welcomed."

IEEE P802.3cw D2.4 400 Gb/s over DWDM systems 4th Working Group recirculation ballot comments

CI 155 SC 155 P 42 L 4 # 2317

Dawe, Piers Nvidia

Comment Type TR Comment Status R

D2.1 comment 278: this project is too slow, and has descended to only 25 comments from only four commenters when there is a lot to fix still. The moment for doing this spec in 802.3 has passed, it doesn't add significantly to 400ZR, it lacks momentum and there are not enough willing participants in P802.3cw to justify it.

SuggestedRemedy

Cancel this project.
Encourage those interested to feed their learnings into OIF's "400ZR" maintenance.
Re-use relevant parts of the draft in P802.3dj when the time comes.

Response Response Status U

REJECT.

As noted by commentor, this issue was previously raised in D2.1 comment #278 and there was no consensus to cancel the project.

https://www.ieee802.org/3/cw/comments/D2p1/8023cw_D2p1_comments_final_by_ID_230619.pdf.

Per Motion #1 from https://www.ieee802.org/3/cw/public/23_06/minutes_3cw_2306_approved.pdf the modified project timeline was approved. See https://www.ieee802.org/3/cw/proj_doc/timeline_3cw_230608.pdf

This plan of action was presented to the 802.3 WG at the July 2023 Plenary. See Slide #3 of https://www.ieee802.org/3/minutes/jul23/0723_3cw_open_report.pdf

There is no consensus to change this plan of action at this time.

CI 155 SC 155.1.5 P 35 L 1 # 20427

Dawe, Piers Nvidia

Comment Type TR Comment Status R

This PCS is too complicated for just a "directive" specification. We need examples.

SuggestedRemedy

Create examples of e.g. FEC and other blocks before and after coding. Smallish ones can go in the document, all can be uploaded to the directory that IEEE provides for these things. They might need to cover some of the PMA.

Response Response Status U

REJECT.

A detailed suggested remedy containing an editor's instruction on how to modify the draft was not provided.

The following straw poll was taken:

I would support rejecting comment #427
Yes - 10

N- 2

CI 155 SC 155.2.4.11 P 44 L 36 # 20463

Dawe, Piers Nvidia

Comment Type TR Comment Status R

generic operation ... in ITU-T G.709.3 Annex D: but that contains undefined symbols and terms.

SuggestedRemedy

As it seems it is not very long, write it out cleanly here

Response Response Status U

REJECT.

No consensus to make a change.

IEEE P802.3cw D2.4 400 Gb/s over DWDM systems 4th Working Group recirculation ballot comments

CI 155 SC 155.2.5.11 P 54 L 30 # 2338

Dawe, Piers Nvidia

Comment Type TR Comment Status R

D2.0 comment 463: generic operation ... in ITU-T G.709.3 Annex D: but that contains undefined symbols and terms. As it seems it is not very long, write it out cleanly here This is supposed to be a spec, we need a specific definition, not "generic". G.709.3 Annex D describes GMP (as referenced in 155.2.5.3), not the Hamming SD-FEC scheme. Also, G.709.3 is in revision. 400ZR 10.5, Inner Hamming Code, which is about one page long, specifically addresses a systematic (128, 119) double-extended Hamming code.

SuggestedRemedy

Copy the material from 400ZR 10.5, changing some of the b to m if appropriate to match the usual FEC notation in 802.3, and replacing the undefined symbols that look like ^ and V with the ones usually used in 802.3. Whatever symbols are used, say what they mean.

Response Response Status U

REJECT.

As noted by commentor, this issue was previously raised in D2.0 comment #463 which was rejected with the response "No consensus to make a change."

https://www.ieee802.org/3/cw/comments/D2p0/8023cw_D2p0_comments_final_by_ID.pdf.

ITU G.709.3 has been amended in November 2022, but there were no changes to Annex D.

CI 156 SC 156.8 P 96 L 33 # 21284

Dawe, Piers Nvidia

Comment Type TR Comment Status A Adjacent channel isolation

It is hard to grasp what this table is meant to say.

SuggestedRemedy

Provide a graph to illustrate it. Define the terms "frequency offset" and "isolation".

Response Response Status U

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #251.

Straw poll #1:

Do you support the addition of a graph as part of the resolution to this comment to further define adjacent channel isolation.

Yes: 5

No: 6

No consensus to add the graph to the draft.

CI 156 SC 156.9 P 97 L 12 # 21285

Dawe, Piers Nvidia

Comment Type TR Comment Status R

Multiple optical parameters are inadequately defined; some (or more) measurement methods are needed for some of them

SuggestedRemedy

Complete the definitions of the optical parameters, with measurement methods and references as necessary

Response Response Status U

REJECT.

Comment unclear and no suggested remedy provided.

CI 156 SC 156.9 P 102 L 13 # 2320

Dawe, Piers Nvidia

Comment Type TR Comment Status R

D2.1 comment 285, optical parameters are inadequately defined.

SuggestedRemedy

Review the 400ZR maintenance projects' activities for corrections and improvements and changes that would apply to this draft, including to EVM.

Response Response Status U

REJECT.

A detailed suggested remedy containing an editor's instruction on how to modify the draft was not provided.

IEEE P802.3cw D2.4 400 Gb/s over DWDM systems 4th Working Group recirculation ballot comments

CI 156 SC 156.9.1 P 102 L 45 # 2331

Dawe, Piers Nvidia
 Comment Type TR Comment Status R

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise. The header for this column is "Parameter" but "Laser frequency noise mask" is not an observable property of a signal, not even hypothetically. It's a mask, a property of the spec.

SuggestedRemedy

Change "Laser frequency noise mask" here, in Table 156-7 and in the title of 156.9.6. In 156.9.6, start by saying what frequency noise is before discussing the mask.

Response Response Status U

REJECT.

No consensus to make a change.

The CRG expressed interest in contributions related to laser frequency noise.

Contributions are encouraged.

CI 156 SC 156.9.5 P 105 L 46 # 9

Dawe, Piers Nvidia
 Comment Type TR Comment Status R

This says "Laser frequency noise is measured using an unmodulated laser as specified in Table 156-11" but frequency noise is not measured directly, it is derived from a measurement of something else. This doesn't say what is measured, or how, or how what

SuggestedRemedy

Change this spec to power spectrum or phase noise, or add the missing information so that "frequency noise" is defined.

Response Response Status U

REJECT. There was no consensus to make a change.

CI 156 SC 156.9.5 P 106 L 4 # 10

Dawe, Piers Nvidia
 Comment Type TR Comment Status R

The units of frequency noise are Hz²/Hz. No watts or dB involved. Frequency noise is not a power spectral density. D2.1 comments 285, optical parameters are inadequately defined, and other comments specifically on frequency noise.

SuggestedRemedy

Change this spec to power spectrum or phase noise, or change Table 156-13--Frequency vs spectral power density to 156-13--Frequency noise mask Change "One-sided frequency noise power spectral density (Hz²/Hz)" in the table and "One-sided frequency noi

Response Response Status U

REJECT. No consensus to make a change.

CI 156 SC 156.9.5 P 106 L 6 # 11

Dawe, Piers Nvidia
 Comment Type TR Comment Status R

"One-sided" is ambiguous and does not appear in the text. It might mean that only one side is shown, and the other is the same, or it might mean that both sides are to be summed (presumably in an RMS way). D2.1 comments 285, optical parameters are inadeq

SuggestedRemedy

In the text, say which is meant.

Response Response Status U

REJECT. No consensus to make a change.

IEEE P802.3cw D2.4 400 Gb/s over DWDM systems 4th Working Group recirculation ballot comments

CI 156 SC 156.9.6 P 99 L 34 # 21286

Dawe, Piers Nvidia

Comment Type TR Comment Status R

"Frequency noise" is extremely arcane, and not defined here. Phase noise is much more commonplace (but ambiguous, so that would need definition too). Also, it is not clear how the "frequency noise" is to be measured if the transmitter is transmitting Pattern 5; there needs to be a method that can tell unwanted "frequency noise" from the intended modulation.

SuggestedRemedy

If there is a well-known metric that does the job, use that instead. Either way, define the parameter with the relevant text, equation(s) and/or references, and write down how it may be measured.

Response Response Status U

REJECT.

No suitable definitions were found and a contribution to recommend a definition would be welcome.

No consensus to make a change at this time.

CI 156 SC 156.9.6 P 105 L 8 # 2325

Dawe, Piers Nvidia

Comment Type TR Comment Status R

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise and write down how it may be measured. For example, it is not stated what is measured in Hz². It is not stated adequately what to do with the two sidebands. The table column header says one-sided, but that's the wrong place to attempt a definition, and does it mean one folds both sidebands together, explicitly or as in a self-homodyne measurement, or takes the worst of the two, or what? It is not stated whether +ve and -ve frequencies are taken into account or just +ve. It seems that this extremely arcane term is more of a concept, or at most a laser modeller's input parameter, than an observable output, so it is not clear that it is the right thing to be specifying, as it may not be measurable.

SuggestedRemedy

Define and specify something relevant and measurable, clearly and completely, with an explanation of how it may be measured and what instrument may be used, and references as necessary. Probably an example is needed. Phase noise is a better-known parameter with some literature, although it needs careful definition to avoid ambiguity. See e.g. IEC 61280-1-3, Fibre optic communication subsystem test procedures--Part 1-3: General communication subsystems--Central wavelength and spectral width measurement for an example of a measurement spec that can be referred to in a definition.

Response Response Status U

REJECT.

No consensus to make a change.

The CRG expressed interest in contributions related to laser frequency noise.

Contributions are encouraged.

IEEE P802.3cw D2.4 400 Gb/s over DWDM systems 4th Working Group recirculation ballot comments

CI 156 SC 156.9.6 P 105 L 8 # 2336

Dawe, Piers Nvidia
 Comment Type TR Comment Status R

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise. The method of interpolation for the laser frequency noise mask is not specified. Figure 156-7 implies log-log interpolation but that is illustrative not normative.

SuggestedRemedy

State that log-log interpolation is used to build the mask is not specified.

Response Response Status U

REJECT.

No consensus to make a change.

The CRG expressed interest in contributions related to laser frequency noise.

Contributions are encouraged.

CI 156 SC 156.9.6 P 105 L 9 # 2328

Dawe, Piers Nvidia
 Comment Type TR Comment Status R

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise and write down how it may be measured. The laser frequency noise is supposed to be controlled down to less than 100 Hz. That's too vague for a spec. No indication is given of how it might be measured, but instruments that can measure GHz often don't measure kHz and below.

SuggestedRemedy

Either don't say anything about frequencies lower than the spec range, or use a separate recommendation (not expected to be testable). Review whether 100 Hz is feasible or necessary, change the limit if appropriate.

Response Response Status U

REJECT.

No consensus to make a change.

The CRG expressed interest in contributions related to laser frequency noise.

Contributions are encouraged.

CI 156 SC 156.9.6 P 105 L 9 # 2326

Dawe, Piers Nvidia
 Comment Type TR Comment Status R

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise. This text says "The mask frequencies are relative to the laser center frequency from *less than* 100 Hz to half the signaling rate", Table 156-13 has 10^2 to 10^9 Hz, and Figure 156-7 shows 10^2 to something indeterminate above 10^10.

SuggestedRemedy

Reconcile the frequency range for this spec, with clear and consistent lower and upper frequencies. For example, 100 Hz to 59.84375/2 = 29.921875 GHz, or 100 Hz to 30 GHz, or 100 Hz to 30.8 GHz to match the transmit spectrum.

Response Response Status U

REJECT.

No consensus to make a change.

The CRG expressed interest in contributions related to laser frequency noise.

Contributions are encouraged.

CI 156 SC 156.9.6 P 105 L 15 # 2337

Dawe, Piers Nvidia
 Comment Type TR Comment Status R

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise. This says "The definition of maximum laser linewidth is provided in ITU-T G.698.2." G.698.2, 7.2.8 Maximum laser linewidth, says "The laser linewidth is defined as: The level of the white noise component of the power spectrum density of the instantaneous laser frequency multiplied by pi." We need a definition of linewidth, not maximum laser linewidth. A power spectrum density would be in the dimensions of power per frequency, which is not inverse time, so this definition is not satisfactory as it stands.

SuggestedRemedy

Use another reference with a dimensionally correct definition, or write one for laser linewidth (not "maximum laser linewidth" here).

Response Response Status U

REJECT.

No consensus to make a change.

The CRG expressed interest in contributions related to laser frequency noise.

Contributions are encouraged.

IEEE P802.3cw D2.4 400 Gb/s over DWDM systems 4th Working Group recirculation ballot comments

CI 156 SC 156.9.6 P 105 L 21 # 2330

Dawe, Piers Nvidia
 Comment Type TR Comment Status R

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise and write down how it may be measured. This says "One-sided frequency noise power spectral density (Hz²/Hz)". I can see that a spectral density can be per hertz. Power has dimensions of energy per time, while Hz² is time⁻². These are incompatible.

SuggestedRemedy

If the units are not changed, delete "power" in the table row header and caption, and Figure 156-7, both y axis and caption.

Response Response Status U

REJECT.

No consensus to make a change.

The CRG expressed interest in contributions related to laser frequency noise.

Contributions are encouraged.

CI 156 SC 156.10.1.2.2 P 94 L 36 # 20564

Dawe, Piers Nvidia
 Comment Type TR Comment Status R

Need a bigger block size for at least one of these, to go with the jitter corner frequency

SuggestedRemedy

Response Response Status U

REJECT.

The CRG had no consensus to make a change at this, more study on a suitable solution is required.

CI 156 SC 156.10.1.2.4 P 112 L 21 # 4

Dawe, Piers Nvidia
 Comment Type TR Comment Status R

The measurement already has significant filtering: "The coherent receiver has a bandwidth of at least 30 GHz". Filtering it again without taking this into account would be too much.D2.1 comments 285, optical parameters are inadequately defined.

SuggestedRemedy

Say that the signal is further filtered so that the combined effect of the observation filter in 156.10.1.1 Calibrated coherent receiver and this filter is the RRC response.

Response Response Status U

REJECT. There is an understanding there are 2 stages of filtering. It is not clear if the RRC filter is adjusted based on the electrical bandwidth. There was no consensus to make a change at this time.