

IEEE P802.3cw D2.2 400 Gb/s over DWDM systems 2nd Working Group recirculation ballot comments

CI 156 SC 156.9.27 P 103 L 48 # 145

Dudek, Mike Marvell

Comment Type TR Comment Status A

The maximum ripple is specified as 2.5dB in table 156-8 but it is stated as being between 3dB points so with that definitions it must be at least 3dB.

SuggestedRemedy

Clarify the definition. Maybe it should be measured over a narrower wavelength range or maybe relative to a specific mask.

Response Response Status U

ACCEPT IN PRINCIPLE.

In 156.9.27 change "The ripple is the maximum peak-to-peak insertion loss variation between 3 dB points in the channel passband."

to

"The ripple is the maximum peak-to-peak insertion loss variation measured between the +/- 3 dB frequency points as defined by the maximum spectral excursion mask."

CI 155 SC 155 P 39 L 1 # 278

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.

IEEE P802.3cw D2.2 400 Gb/s over DWDM systems 2nd Working Group recirculation ballot comments

CI 116 SC 116.1.3 P 33 L 12 # 280

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.

CI 155 SC 155 P 39 L 1 # 281

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 156 SC 156.8 P 96 L 33 # 284

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 # 285

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.

IEEE P802.3cw D2.2 400 Gb/s over DWDM systems 2nd Working Group recirculation ballot comments

CI 156 SC 156.9.6 P 99 L 34 # 286

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 156.7 P 84 L 22 # 20334

Ghiasi, Ali Ghiasi Quantum/Marvell

Comment Type TR Comment Status R

The receiver must tolerate 26 dB OSNR and meet the required error rate, it is not clear what receive OSNR (min) of 29 dB provides

SuggestedRemedy

Need discussions on the intent

Response Response Status U

REJECT.

Receiver OSNR tolerance is measured without line impairments, see 156.9.24, which is different than Receiver OSNR which includes line impairments, see 156.9.23

CI 156 SC 156.10.1.1 P 93 L 44 # 20336

Ghiasi, Ali Ghiasi Quantum/Marvell

Comment Type TR Comment Status R

Assuming just 4 bits ENOB from 10 MHz to 29.9 MHz the reference receiver will have additional penalty than real receiver that has typically 6+ bits ENOB at low frequencies and about 4 bits at high frequency

SuggestedRemedy

If there is interest I can bring a frequency dependent ENOB mask

Response Response Status U

REJECT.

No suggested remedy provided

CI 156 SC 156.7.1 P 82 L 48 # 20337

Ghiasi, Ali Ghiasi Quantum/Marvell

Comment Type TR Comment Status R

For full interoperability using EVM may need additional constraints based on the data in rahn_3cw_01a_220223 and way_3cw_01a_220523

SuggestedRemedy

Need more data to prove that EVM will provide the IEEE level of interoperability

Response Response Status U

REJECT.

No suggested remedy provided

IEEE P802.3cw D2.2 400 Gb/s over DWDM systems 2nd Working Group recirculation ballot comments

CI 155 SC 155.3.3.5 P 58 L 45 # 20341

Zimmerman, George CME Consulting/APL Group, Cisco, Commscope, Ma

Comment Type TR Comment Status A rewrite bucket

"The signals are sampled by an ADC on each lane at a sampling rate." "The details of the ADC . are implementation specific". This is a description of an implementation, not appropriate for an interoperability specification. If someone could do the signal processing optically, analog, or by magic, it would still comply with the standard. The fact that an ADC is used, isn't a part of the interoperability standard, or even any of the characteristics of the ADC. Hence the mention is inappropriate and should be deleted. The sentence works just fine anyways and describes the processing without the "by an ADC".

SuggestedRemedy

Change header of 155.3.5 to Receive signal sampling.
On line 50, Delete "by an ADC"
Change line 54 to "The details of the sampling, including any quantization and the chosen sampling rate are implementation specific."
Replace "ADC" with "Sampler" in figure 155-10.

Response Response Status U

ACCEPT IN PRINCIPLE.

See response to comment #346.

CI 155 SC 155.3.3.1 P 52 L 28 # 20342

Zimmerman, George CME Consulting/APL Group, Cisco, Commscope, Ma

Comment Type TR Comment Status A rewrite bucket

"The received symbol signals are digitized into more than 4 discrete levels by the analog to digital converters (ADC) in the PMA sublayer and the number of bits for each signal is m/4 bits." This is a description of an implementation and is inappropriate for an interoperability standard. If some description is needed, one could rewrite this more generally, as is suggested in the remedy. Further, it appears that the "m/4 bits" is a detail that is unused in the draft (I searched). If it is used somewhere, please provide a pointer to where it is relevant. Otherwise delete the unnecessary detail which looks like a specification but isn't.

SuggestedRemedy

Preferably - delete the indicated sentence.
Alternatively, change the indicated sentence to read "The received symbol signals are sampled and quantized in the PMA sublayer."
If the m/4 bits is used somewhere, provide a reference.

Response Response Status U

ACCEPT IN PRINCIPLE.

See response to comment #346.

CI 155 SC 155.3.3.5 P 58 L 45 # 20343

Zimmerman, George CME Consulting/APL Group, Cisco, Commscope, Ma

Comment Type TR Comment Status A rewrite bucket

"The signals are sampled by an ADC on each lane at a sampling rate." "The details of the ADC . are implementation specific". This is a description of an implementation, not appropriate for an interoperability specification. If someone could do the signal processing optically, analog, or by magic, it would still comply with the standard. The fact that an ADC is used, isn't a part of the interoperability standard, or even any of the characteristics of the ADC. Hence the mention is inappropriate and should be deleted. The sentence works just fine anyways and describes the processing without the "by an ADC".

SuggestedRemedy

Change header of 155.3.5 to Receive signal sampling.
On line 50, Delete "by an ADC"
Change line 54 to "The details of the sampling, including any quantization and the chosen sampling rate are implementation specific."
Replace "ADC" with "Sampler" in figure 155-10.

Response Response Status U

ACCEPT IN PRINCIPLE.

See response to comment #346.

CI 155 SC 155.3.1.3 P 51 L 26 # 20345

Zimmerman, George CME Consulting/APL Group, Cisco, Commscope, Ma

Comment Type TR Comment Status A rewrite bucket

This figure is supposed to be a functional block diagram, not an implementation diagram. There are no characteristics for the DAC blocks defined in the specification. The closest thing in the text is 155.3.3.4 which are called the 16QAM encode and signal drivers. However, most other 802.3 PHY clauses leave out signal drivers, DACs and the like, and there are no specific requirements in 155.3.3.4, so deleting the blocks seems the right approach to making a functional block diagram.

SuggestedRemedy

Preferably, delete the "DAC" blocks from Figure 155-10 (going straight to the output is fine)
Alternatively, Relabel "16QAM Encoder and Signal Driver" (probably drawing as 2 blocks since you show I&Q paths)

Response Response Status U

ACCEPT IN PRINCIPLE.

See response to comment #346.

IEEE P802.3cw D2.2 400 Gb/s over DWDM systems 2nd Working Group recirculation ballot comments

Cl 155 SC 155.7.4.1 P 70 L 24 # 20346

Zimmerman, George CME Consulting/APL Group, Cisco, Commscope, Ma

Comment Type TR Comment Status A rewrite bucket

This is a general comment on the requirements. I am attaching it to these PICS because this is where it became apparent. The style of IEEE SA standards (and IEEE Std 802.3) is that requirements use the term "shall". Each PICS item should have an associated "shall" and each "shall" should have a PICS. However, 155.7.4.1 is a list of the subclauses for the most part. Further, looking at the subclauses, they are largely without "shalls". Most of the items in clause 155 are descriptive of an implementation, and do not use the term shall. They use "is" or other descriptive language. The PICS are a list of the functional blocks described, but most of those functional blocks are lacking actual requirements. Instead they often describe an implementation or, worse yet, sometimes try to require a particular implementation ("an implementation shall"). What needs to happen is that the clause needs to be rewritten carefully considering what requirements are needed for interoperability, and deleting the unnecessary implementation description. This is a big job, and, in my opinion, means the draft is not technically complete, and should not have begun initial working group ballot. I truly regret having to make a comment like this, but I believe this is a great example of why we have working group ballots in 802.

SuggestedRemedy

Unfortunately, the draft is so far from complete that I cannot propose a specific remedy for the systematic problem. I can suggest that the TF look at each subblock, determine what the observed behavior is, determine which parts matter to interoperability, and write "shall" statements in the subclauses. Then those shall statements can be made as PICS. Additionally, this will highlight where there is implementation description that can be deleted. When this is done, restart working group ballot.

Response Response Status U

ACCEPT IN PRINCIPLE.

With editorial license, restructure and clarify Clause 155 and 156 as appropriate: to identify interoperability requirements using "SHALL" statements, as needed. to address issues noted in https://www.ieee802.org/3/cw/public/22_10/dambrosia_3cw_01b_221018.pdf

Cl 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

IEEE P802.3cw D2.2 400 Gb/s over DWDM systems 2nd Working Group recirculation ballot comments

Cl 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

Cl 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.

Cl 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.