

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 00 SC 0 P L # 172
 Anslow, Pete Ciena

Comment Type E Comment Status D bucket

Now that IEEE Std 802.3-2012 has been approved, update all references in the draft to reflect 2012 and remove the reference to "Draft 3.1" in the frontmatter.

SuggestedRemedy

Update all 802.3 references in the draft to be "IEEE Std 802.3-2012" and remove the reference to "Draft 3.1" in the frontmatter.

Proposed Response Response Status W

PROPOSED ACCEPT.

The frontmatter will be updated under the guidance of the Working Group chair.

In addition, replace all references to the base document with IEEE Std 802.3-2012.

Cl 00 SC 0 P L # 180
 Anslow, Pete Ciena

Comment Type T Comment Status D bucket

The content of the P802.3bj draft seems to be sufficiently stable that the content of Clause 45, Clause 30 Annex 91A and the various PICS proforma should now be populated.

SuggestedRemedy

Complete the content of Clause 45, Clause 30 Annex 91A and the various PICS proforma.

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 69 SC 69.1.2 P 28 L 29 # 31
 Anslow, Pete Ciena

Comment Type E Comment Status A

The editing instruction says "Delete 69.1.2."
 When applied to the base document, this will have the effect of renumbering 69.1.3 to be 69.1.2.
 The modification to what was formerly 69.1.3 just below should reflect this change.

Note, the same issue for 80.1.2 is the subject of a separate comment.

SuggestedRemedy

Change the editing instruction to "Delete 69.1.2 and renumber 69.1.3 to 69.1.2 accordingly."
 For 69.1.3, move the editing instruction above the title, renumber to 69.1.2 and amend the editing instruction to be:
 "Change the first paragraph of 69.1.3 (now renumbered to 69.1.2) as shown:"

Response Response Status C

ACCEPT IN PRINCIPLE.

This subclause will be handled in a manner consistent with the treatment of 80.1.2 (see comment #6).

Cl 69 SC 69.1.2 P 28 L 32 # 106
 Barrass, Hugh Cisco

Comment Type E Comment Status A

For consistency - and also so that commenters can see what is changing - show the deleted text.

SuggestedRemedy

Show the deleted text.

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #31.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 69 SC 69.1.3 P 28 L 51 # 2
 Anslow, Pete Ciena

Comment Type E Comment Status D bucket

The editing instruction says "Change Figure 69-1 and insert Figure 69-1a as shown:" but Figure 69-1 does not show any changes, it is a replacement figure.

SuggestedRemedy

Change the editing instruction to:
 "Replace Figure 69-1 and insert Figure 69-1a as shown:"

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 69 SC 69.1.3 P 29 L 16 # 423
 Dawe, Piers IPtronics

Comment Type E Comment Status D bucket

For consistency with Fig 80-1,

SuggestedRemedy

Mark the FEC for 10GBASE-KR, and 40GBASE-KR4 (Fig 69-1a), as optional.

Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.

Also change FEC to RS-FEC for 100GBASE-KR4 and 100GBASE-KP4 (Figure 69-1a).

Cl 69 SC 69.1.3 P 30 L 45 # 436
 Dawe, Piers IPtronics

Comment Type T Comment Status A

Not so fast! It's still the case that a 2-lane 10GBASE-KX4 wouldn't be compliant, and so on. As the channel or medium isn't normative for older BPE, and MDI is shown in other places, it may be convenient to attach this requirement to the MDI.

SuggestedRemedy

Reinstate item f but change "as specified in" to "of". Add the new PMD types. Rework to say MDIs for types A, B, C have one pair/differential electrical path in each direction while X, Y, Z have four. No need for clause numbers:
 f) The MDI for 1000BASE-KX and 10GBASE-KR uses one pair of electrical connections for each direction, while 10GBASE-KX4, 40GBASE-KR4 and ... have four pairs.

Response Response Status C

ACCEPT IN PRINCIPLE.

During consideration of this comment, it was observed that XLAUI is not included in the list for 40GBASE-KR4.

Replace item f):

"f) The PMA service interface, which, when physically implemented as XLAUI (40 Gigabit Attachment Unit Interface) at an observable interconnection port, uses a 4 lane data path as specified in Annex 83A or Annex 83B."

Add item g):

"g) The MDI for 1000BASE-KX and 10GBASE-KR use a serial data path while the MDI for 10GBASE-KX4, 40GBASE-KR4, 100GBASE-KR4, and 100GBASE-KP4 use a four lane data path."

Cl 69 SC 69.2.4 P 32 L 6 # 3
 Anslow, Pete Ciena

Comment Type E Comment Status D bucket

The cell borders for Table 69-1a in the Nomenclature row are not consistent for clauses 91, 93 and 94

SuggestedRemedy

Change the left and right borders in the Nomenclature row for clauses 91 and 93 to be "very thin"

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 69 SC 69.5 P 32 L 47 # 20
 Anslow, Pete Ciena

Comment Type T Comment Status D bucket

The text:
 "The supplier of a protocol implementation that is claimed to conform to any part of IEEE Std 802.3, Clause 70 through Clause 74, demonstrates compliance by completing a protocol implementation conformance statement (PICS) proforma." has been changed to:
 "The supplier of a protocol implementation that is claimed to conform to any part of IEEE Std 802.3 demonstrates compliance by completing a protocol implementation conformance statement (PICS) proforma."
 But this is not a true statement. There are many clauses in 802.3 that do not have an accompanying PICS proforma.
 Same issue for 80.7

SuggestedRemedy

Remove the deletion of ", Clause 70 through Clause 74," in 69.5 and also remove the deletion of ", Clause 45, Clause 73, Clause 74, Clause 81 through Clause 89, and related annexes" from 80.7.
 Augment these two statements as required to reflect the new clauses added by the amendment.

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 73 SC 73.10.7 P 35 L 12 # 424
 Dawe, Piers IPtronics

Comment Type E Comment Status D bucket

Make the document easier to use with consistent ordering.

SuggestedRemedy

Put the PMAs and PMDs in the reverse order to Table 73-5 Priority Resolution. Also the list for single_link_ready.

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 73 SC 73.7.2 P 34 L 30 # 437
 Dawe, Piers IPtronics

Comment Type T Comment Status A

Wordsmithing:
 "... the Receive Switch function shall connect the MDI to ... and to the receive path of the 1000BASE-KX ... and 100GBASE-CR4 if the PHY is present."

SuggestedRemedy

"... the Receive Switch function shall connect the MDI to ... and to the receive path of each PMD that is present and has Auto-Negotiation enabled."

Response Response Status C

ACCEPT IN PRINCIPLE.

Considering 73.6.10 and 73.7.2 from the base document, it appears that the Transmit/Receive switch functions connect the HCD PHY to the medium once Auto-Negotiation has completed.

This is reinforced by the requirement in 73.6.10 that only "DME page generator" is connected to the MDI during Auto-Negotiation.

To be consistent with 73.6.10, 73.7.2 should state that, during Auto-Negotiation, the DME page receiver and the receive path of the 1000BASE-KX and 10GBASE-KX4 (if present) to support parallel detection.

It would also be valuable to quantify what it means to be "in Auto-Negotiation."

[Change these two subclauses as shown.]

73.6.10 Transmit Switch function

Prior to entry into the AN_GOOD_CHECK state, the Transmit Switch function shall connect only the DME page generator controlled by the Transmit State Diagram to the MDI.

Upon entry into the AN_GOOD_CHECK state, the Transmit Switch function shall connect the transmit path from a single technology-dependent (highest common denominator) PHY to the MDI.

When a PHY is connected to the MDI through the Transmit Switch function, the signals at the MDI shall conform to all of the PHY's specifications.

73.7.2 Receive Switch function

Prior to entry into the AN_GOOD_CHECK state, the Receive Switch function shall connect the DME page receiver to the MDI. For the Parallel Detection function, the Receive Switch function shall also connect the receive path of the 1000BASE-KX and 10GBASE-KX4 PHY to the MDI when those PHYs are present.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Upon entry into the AN_GOOD_CHECK state, the Receive Switch function shall connect the receive path from a single technology-dependent (highest comment denominator) PHY to the MDI.

Cl 89 SC 1 P 30 L 10 # 298
Ghiasi, Ali Broadcom

Comment Type TR Comment Status D bucket
A more detailed disclaimer needs to be added including the fact VSR2000-3R2 does not have the same level of interoperability or BER objective

Suggested Remedy

The specifications in this clause therefore use a similar methodology to that used in ITU-T G.693 [Bx1] and not recommended for reuse as it does not provide the same level of interoperability or BER other 40GBASE-R PMDs provide.

Proposed Response Response Status W

PROPOSED REJECT.

This comment appears to have been submitted in error. Clause 89 is beyond the scope of P802.3bj.

Cl 89 SC 5.1 P 34 L 33 # 299
Ghiasi, Ali Broadcom

Comment Type TR Comment Status D bucket
PMD service interface TP1 and TP4 are not applicable as they are not currently defined

Suggested Remedy

Remove TP1 and TP4
Add XLAUI interface to the PMA

Proposed Response Response Status W

PROPOSED REJECT.

This comment appears to have been submitted in error. Clause 89 is beyond the scope of P802.3bj.

Cl 89 SC 6.3 P 37 L 36 # 300
Ghiasi, Ali Broadcom

Comment Type TR Comment Status D bucket

With the transmitter center wavelength at 1550 nm compatible with VSR3, there is no need to require FR receiver be dual wavelength. If the reason to add 1310 nm band for some future 1310 nm targeted for lower power and cost but we already declared at the beginning SONET VSR methodology is not recommended for reuse for not having same level of interoperability as IEEE specifications.

Suggested Remedy

Remove the 1310 nm window

Proposed Response Response Status W

PROPOSED REJECT.

This comment appears to have been submitted in error. Clause 89 is beyond the scope of P802.3bj.

Cl 89 SC 6.3 P 37 L 46 # 301
Ghiasi, Ali Broadcom

Comment Type TR Comment Status D bucket
Receiver jitter tolerance test method missing

Suggested Remedy

Add receiver jitter tolerance

Proposed Response Response Status W

PROPOSED REJECT.

This comment appears to have been submitted in error. Clause 89 is beyond the scope of P802.3bj.

Cl 89 SC 7.10 P 42 L 4 # 302
Ghiasi, Ali Broadcom

Comment Type TR Comment Status D bucket

The receiver jitter tolerance here is unstress which is different than 802.3 and note should be added to clarify

Suggested Remedy

Add note receiver jitter tolerance is unstress

Proposed Response Response Status W

PROPOSED REJECT.

This comment appears to have been submitted in error. Clause 89 is beyond the scope of P802.3bj.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 89 SC 9 P 4 L 17 # 303
 Ghiasi, Ali Broadcom
 Comment Type TR Comment Status D bucket
 Definition and test method for dispersion is missing
 SuggestedRemedy
 Add definition and test method
 Proposed Response Response Status W
 PROPOSED REJECT.
 This comment appears to have been submitted in error. Clause 89 is beyond the scope of P802.3bj.

Cl 89 SC 9 P 4 L 19 # 304
 Ghiasi, Ali Broadcom
 Comment Type TR Comment Status D bucket
 Test method for DGD is missing
 SuggestedRemedy
 Add test method
 Proposed Response Response Status W
 PROPOSED REJECT.
 This comment appears to have been submitted in error. Clause 89 is beyond the scope of P802.3bj.

Cl 91 SC 91 P 104 L 0 # 196
 Slavick, Jeff Avago Technologies
 Comment Type T Comment Status A
 No definitions for counter to track the following have been added to the RS-FEC.
 Corrected_block_count
 Uncorrected_block_count
 Symbol_error_count_0
 Symbol_error_count_1
 Symbol_error_count_2
 Symbol_error_count_3
 SuggestedRemedy
 Add a new section named RS-FEC Error monitoring capability which defines the following counters and create MDIO access methods for these as well.
 Corrected_block_count - 32b counter which increments each time a codeword is successfully corrected when fec_bypass_correction is true.
 Uncorrected_block_count - 32b counter which increments each time a codeword is uncorrectable when fec_bypass_correction is false and when the local parity and received parity's don't match when fec_bypass_correction is true.
 Symbol_error_count_0..3 - 32b counter, one for each PMD lane, which increments each time a symbol for the given lane is corrected when fec_bypass_correction is true.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Add a summary of management variables per healey_3bj_02_0912 and define the corresponding register and bits to MMD 1 in Clause 45. Give the editor license to assign registers and bit number, but begin a new contiguous address space starting at 1.200.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.1.2 P91 L 29 # 470
 Cideciyan, Roy IBM

Comment Type TR Comment Status D

RS encoding is mandatory, i.e., not conditional based on PHY type.

SuggestedRemedy

Delete "NOTE 1-CONDITIONAL BASED ON PHY TYPE" and omit superscript "1" in sublayers RS-FEC and AN.

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

The 100GBASE-R family is not limited to 100GBASE-CR4, 100GBASE-KR4, and 100GBASE-KP4. For example, 100GBASE-LR4 and 100GBASE-ER4 do not include the RS-FEC sublayer. Therefore, inclusion of the RS-FEC sublayer is "conditional based on PHY type."

Cl 91 SC 91.2 P92 L 21 # 239
 Healey, Adam LSI Corporation

Comment Type T Comment Status A

Now that the FEC synchronization state diagram has been included in the draft, the assignment of the SIGNAL_OK parameter of the FEC:IS_UNITDATA.indication primitive can be defined.

SuggestedRemedy

Specify that SIGNAL_OK=OK when align_status=TRUE and SIGNAL_OK=FAIL when align_status=FALSE. Also define the value of the rx_bit parameter for the FEC:IS_UNITDATA_i.indication primitives when SIGNAL_OK=FAIL.

Response Response Status C

ACCEPT IN PRINCIPLE.

Define SIGNAL_OK per the comment (note the variable name has changes to fec_align_status).

Specify that when SIGNAL_OK=FAIL, the value of rx_bit is undefined.

Cl 91 SC 91.2 P92 L 33 # 95
 Barrass, Hugh Cisco

Comment Type T Comment Status A

For change of LPI Rx function

rx_mode needs to change direction, also energy_detect and rx_lpi_active need to be added.

SuggestedRemedy

Change:

IS_RX_MODE.indication

To:

IS_RX_MODE.request
 IS_ENERGY_DETECT.indication
 IS_RX_LPI_ACTIVE

Response Response Status C

ACCEPT IN PRINCIPLE.

Clause 91 does not require the IS_RX_LPI_ACTIVE primitive.

Add IS_ENERGY_DETECT and change the direction of IS_RX_MODE per the comment.

Cl 91 SC 91.3 P92 L 44 # 161
 Ran, Adeo Intel

Comment Type TR Comment Status D bucket

RS-FEC is defined only to be a client of the 100GBASE-R PCS where the number of upstream lanes is 20.

Also: the terms p and q only appear in one paragraph in subclause 83.1.4 in a descriptive manner, and are not used or officially defined anywhere else. It would be easier to search for the more unique terms LANES_UPSTREAM and LANES_DOWNSTREAM that appear in 83.7.3. Perhaps a maintenance change in 83.1.4 is also due.

SuggestedRemedy

Change "four upstream lanes" to "20 upstream lanes".
 Change "PMA service interface width, p, is set to 4" to "PMA service interface widths LANES_UPSTREAM and LANES_DOWNSTREAM are set to 20 and 4 respectively".

Proposed Response Response Status Z

PROPOSED REJECT.

This comment was WITHDRAWN by the commenter.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.4 P 92 L 52 # 245
 Healey, Adam LSI Corporation
 Comment Type T Comment Status A
 The Clause 91 architecture has stabilized to the point where a delay constraint can be provided.
 SuggestedRemedy
 Specify the maximum delay contributed by the RS-FEC sublayer.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 See comment #190.

Cl 91 SC 91.4 P 92 L 53 # 190
 Slavick, Jeff Avago Technologies
 Comment Type T Comment Status A
 Need to replace TBDs with values for maximum delay contributed by the RS-FEC. Clause 74 was set to ~3x FEC frame size.
 SuggestedRemedy
 Change TBDs to be 4096 BT, 158.3ns, 8 pause_quanta
 That's ~3.01 RS-FEC frames for KP4 and 3.1 for KR4/CR4
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 It should be noted that the purpose of this Delay specification is to bound the delay through a link for MAC Control PAUSE operation. Low latency implementations are certainly possible.
 Set TBD to 80 pause_quanta (derive equivalent for other units). This enables a wide range of implementations.
 In addition, comment #241 requests more information on the impact of error marking on FEC latency. The specified value is inclusive of error marking and for the stated purpose of this requirement, a limit without error marking does not need to be specified.

Cl 91 SC 91.5.1 P 94 L 4 # 99
 Barrass, Hugh Cisco
 Comment Type T Comment Status A
 For change of LPI Rx function
 Fix the block diagram in Fig 91-2
 SuggestedRemedy
 Change the direction FEC:IS_RX_MODE.request
 Add FEC:IS_ENERGY_DETECT.indication
 Add FEC:IS_RX_LPI_ACTIVE.request
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Clause 91 does not use the IS_RX_LPI_ACTIVE primitive. Implement the other changes in the suggested remedy.

Cl 91 SC 91.5.1 P 94 L 40 # 100
 Barrass, Hugh Cisco
 Comment Type T Comment Status A
 For change of LPI Rx function
 Fix the block diagram in Fig 91-2
 SuggestedRemedy
 Change the direction FEC:IS_RX_MODE.request
 Add FEC:IS_ENERGY_DETECT.indication
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Change the direction of PMA:IS_RX_MODE.request and add PMA:IS_ENERGY_DETECT.indication

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.2.2 P 93 L 27 # 222
 Gustlin, Mark Xilinx
 Comment Type T Comment Status D
 The skew variation of 0.2ns is discussed, but it would be good to also refer to SP1 in this sentence, similar to how it is referenced in 83.5.3.3.
 SuggestedRemedy
 Per the comment.
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 See healey_02_0912.

Cl 91 SC 91.5.2.4 P 93 L 46 # 197
 Slavick, Jeff Avago Technologies
 Comment Type T Comment Status A
 Replace TBD with the BIP error counter register that already exist in MDIO.
 SuggestedRemedy
 Change TBD with 3.200 to 3.219
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 3.200 and 3.129 are PCS bits/registers. As the BIP check is done by the RS-FEC sublayer, new counters should be defined in MMD 1 (Clause 74 FEC register space resides in this MMD, so it is proposed that the Clause 91 register space also reside here).
 See comment #196.

Cl 91 SC 91.5.2.5 P 95 L 1 # 53
 Szczepanek, Andre Inphi
 Comment Type TR Comment Status A
 The output of the transcoder for invalid sync headers is not defined.
 If for any j=0 to 3, tx_coded_j<1> == tx_coded_j<0> what is tx_xcoded ?
 SuggestedRemedy
 for any j=0 to 3, tx_coded_j<1> == tx_coded_j<0>
 then the transcoded output should be equivalent to the transcode of four Local_fault input words
 Response Response Status C
 ACCEPT IN PRINCIPLE.

[Added Clause (91) to Sbcl field for consistent sorting.]

Specify that, if any of the four 66-bit blocks tx_coded_j has an invalid sync. header, then tx_xcoded<0> is set to 0 and tx_xcoded<4:1> is set to 1111. The second nibble from the first 64-bit block payload is deleted.

Specify that, when rx_xcoded is received with rx_xcoded<0>=0 and rx_xcoded<4:1>=1111, then the sync. headers of the blocks rx_coded_j are set to be invalid: 00, 11, 00, 11. The second nibble from the first 66-bit block payload is set to zero and scrambled per the current procedure.

An error not considered by the commenter is the case where an invalid first nibble of the block type field is received by the 256B/257B to 64B/66B transcoder. Specify that this case is handled per healey_3bj_02_0912 slide 20.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.2.5 P95 L 12 # 15
 Anslow, Pete Ciena

Comment Type E Comment Status D bucket

This says "such that tx_coded_c<1:0>=01."
 The usual arrangement for the sync bits is to show them with the first bit transmitted on the left (i.e. for control, sync = 10).
 Consequently, it would be clearer to show each bit separately.

Also, it would keep the sync bits in the usual order if the <0> index was shown first.

Similar issues in 91.5.3.5 and 91.5.3.7

SuggestedRemedy

On line 1, change:
 "tx_coded_j<1>=1 and tx_coded_j<0>=0," to:
 "tx_coded_j<0>=0 and tx_coded_j<1>=1,"

On line 7 change:
 "tx_coded_j<1>=0 and tx_coded_j<0>=1," to:
 "tx_coded_j<0>=1 and tx_coded_j<1>=0,"

On line 12 change:
 "such that tx_coded_c<1:0>=01." to:
 "such that tx_coded_c<0>=1 and tx_coded_c<1>=0

On page 101, line 30 change:
 "rx_coded_j<1> = 1 and rx_coded_j<0> = 0" to:
 "rx_coded_j<0> = 0 and rx_coded_j<1> = 1"

On page 101, line 35 change:
 "rx_coded_j<1> = 0 and rx_coded_j<0> = 1" to:
 "rx_coded_j<0> = 1 and rx_coded_j<1> = 0"

On page 101, line 36 change:
 "rx_coded_j<1> = 1 and rx_coded_j<0> = 0" to:
 "rx_coded_j<0> = 0 and rx_coded_j<1> = 1"

On page 102, line 32 change:
 "Finally, am_x<1:0> = 01" to:
 "Finally, am_x<0> = 1 and am_x<1> = 0"

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 91 SC 91.5.2.5 P95 L 12 # 240
 Healey, Adam LSI Corporation

Comment Type T Comment Status D bucket

Clarify the assignment of tx_coded_c<1:0>.

SuggestedRemedy

Change to tx_coded_c<1:0>=01 to tx_coded_c<1>=0 and tx_coded_c<0>=1.

Proposed Response Response Status W
 PROPOSED ACCEPT.

See comment #15.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.2.5 P 95 L 15 # 56
 Szczepanek, Andre Inphi

Comment Type ER Comment Status R

The function for omission of the first codeword "s" nibble is unnecessarily terse and makes it difficult to understand what is required. As c only has 4 possible values, why not just state all 4 possible bit muxes.

SuggestedRemedy

Replace :

e)Omit tx_coded_c<9:6>, which is the second nibble (based on transmission order) of the block type field for tx_coded_c, from tx_xcoded per the following expressions.
 tx_xcoded<(64c+8):5> = tx_payloads<(64c+3):0>
 tx_xcoded<256:(64c+9)> = tx_payloads<255:(64c+8)>

With :

e)Omit tx_coded_c<9:6>, which is the second nibble (based on transmission order) of the block type field for tx_coded_c, from tx_xcoded per the following :
 if (c==0) tx_coded <256:5> = tx_payloads<255:8> :: tx_payloads<3:0>
 if (c==1) tx_coded <256:5> = tx_payloads<255:72> :: tx_payloads<67:0>
 if (c==2) tx_coded <256:5> = tx_payloads<255:136> :: tx_payloads<131:0>
 if (c==3) tx_coded <256:5> = tx_payloads<255:200> :: tx_payloads<195:0>

Response Response Status C

REJECT.

[Added Clause (91) to Sbcl field for consistent sorting.]

The text is correct as written. Illustrations have been added (see Figure 91-3) to help the reader understand the process.

The suggested remedy includes notation for array concatenation ":" that is not used elsewhere in IEEE 802.3. The existing definition does not require new array concatenation notation.

While the mathematical description is precise, it requires the user to do a number of index computations to understand the construction of the codeword. It is not clear that the calculations involving the variable c are more onerous than the others.

See also comment #52.

Cl 91 SC 91.5.2.5 P 95 L 20 # 155
 Ran, Adee Intel

Comment Type ER Comment Status A

It is not absolutely clear from the text whether the XOR occurs only for the case where at least one 66-bit block is a control block, or for all cases including all-data blocks. I assume the latter is correct, but it is preferable to avoid possible confusion.

The examples in figure 91-3 fail to depict this operation - bits 4:0 are shown as in the original assignment.

Also: the second sentence in this paragraph should be in a separate paragraph.

SuggestedRemedy

Use a temporary variable tx_xcoded_header<4:0> for all the assignments to tx_xcoded<4:0> that occur before this paragraph.

Update figure 91-3 to include both tx_xcoded_header<4:0> and tx_xcoded<4:0>. (May require restructuring the figure).

Change the paragraph in lines 20-22 to the following:

"
 Set tx_coded<4:0> to the result of the bit-wise exclusive-OR of tx_xcoded_header<4:0> and tx_xcoded<12:8>."

Several examples that illustrate the transcoding process are shown in Figure 91-3.

Response Response Status C

ACCEPT IN PRINCIPLE.

In the first paragraph of 91.5.2.5, change reference to tx_xcoded<256:0> to tx_scrambled<256:0>.

Replace the last paragraph of 91.5.2.5 with following definition of tx_scrambled.

"Several examples of the construction of tx_xcoded<256:0> are shown in Figure 91-3.

Finally, scramble tx_xcoded<256:0> to yield tx_scrambled<256:0> as follows.

- a) Set tx_scrambled<4:0> to the result of the bit-wise exclusive-OR of the tx_xcoded<4:0> and tx_xcoded<12:8>.
- b) Set tx_scrambled<256:5> to tx_xcoded<256:5>."

Re-name Figure 91-3 to be "Examples of the construction of tx_xcoded".

Change 91.5.2.7, page 98, line 8 to "The message symbols are composed of the bits of the transcoded blocks tx_scrambled (including a mapped group of alignment markers when appropriate) such that bit 0 of the first transcoded block in the message (or am_txmapped<0>)."

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

In Figure 91-6, replace tx_xcoded with tx_scrambled.

CI 91 SC 91.5.2.5 P 95 L 20 # 198
 Slavick, Jeff Avago Technologies

Comment Type T Comment Status A
 Figure 91-3 doesn't incorporate the XOR function in it's illustration of the transcoding process

SuggestedRemedy
 Change
 "Several examples that illustrate the transcoding process are shown in Figure 91-3."
 to
 "Several examples that illustrate the transcoding process steps a-e are shown in Figure 91-3."

Response Response Status C
 ACCEPT IN PRINCIPLE.
 See comment #155.

CI 91 SC 91.5.2.5 P 95 L 21 # 471
 Cideciyan, Roy IBM

Comment Type TR Comment Status A
 Figure 91-3 does not show the final change of tx_xcoded<4:0> by using bitwise XOR which is part of the transcoder description.

SuggestedRemedy
 Replace sentence "Several examples that illustrate ... in Figure 91-3." by "Several examples that illustrate the transcoding process without the final modification of tx_xcoded<4:0> are shown in Figure 91-3."

Response Response Status C
 ACCEPT IN PRINCIPLE.
 See comment #155.

CI 91 SC 91.5.2.5 P 95 L 7 # 162
 Ran, Adee Intel

Comment Type TR Comment Status A
 The transcoding procedure does not handle all possible values of tx_coded_j<1:0>. The values 00 and 11 are indeed invalid, but can still occur (e.g. due to errors in reception from upper layers). This is likely to happen more often than once in MTFFPA.

Since the header must be compressed, the reasonable behavior in such cases would be to mark the 66-bit block in question as a control block with /E/ on transmission, to make sure they are discarded by the receiving PCS.

SuggestedRemedy
 Change the condition in line 7 to:
 "If for all j=0 to 3, tx_coded_j<1>!=tx_coded_j<0>, and for at least one value of j, tx_coded_j<1>=0 and tx_coded_j<0>=1"

Add text based on the following paragraph after line 19 (expand the text inside braces to be technically accurate according to comment):

"
 If for any j=0 to 3, tx_coded_j<1>=tx_coded_j<0>, tx_xcoded<256:0> shall be constructed as follows:

- a) tx_coded<0>=0
- b) tx_xcoded<k+1> = tx_coded_k<1> for k=0 to 3 except for k=j
- [c) and on: specify that any blocks where invalid header was found be replaced by control blocks containing /E/]

Add a suitable example to figure 91-3.

Response Response Status C
 ACCEPT IN PRINCIPLE.
 See comment #53.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.2.5 P 96 L 47 # 473
 Cideciyan, Roy IBM
 Comment Type TR Comment Status D bucket
 Header bit (first bit) of transcoded block that contains 4 control blocks not correct.
 SuggestedRemedy
 Replace header bit (first bit) of transcoded block by 0.
 Proposed Response Response Status W
 PROPOSED ACCEPT.
 Comment is against Figure 91-3.

Cl 91 SC 91.5.2.6 P L # 464
 Cideciyan, Roy IBM
 Comment Type ER Comment Status D bucket
 Title of subclause is "Alignment mapping and insertion" whereas title of subclause 91.5.3.7 is "Alignment marker mapping and insertion"
 SuggestedRemedy
 Both subclauses should have the same title, i.e., either "Alignment mapping and insertion" or "Alignment marker mapping and insertion". My preference is that both subclauses have the more descriptive title "Alignment marker mapping and insertion".
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 Change heading of 91.5.2.6 to "Alignment marker mapping and insertion".

Cl 91 SC 91.5.2.6 P 113 L 38 # 206
 Zhong, Qiwen Huawei
 Comment Type E Comment Status A
 "Figure 91 - 64B/66B to 256B/257B transcoding example" Especially "Example 3: Alternating data and control blocks" might misguide readers as the Ethernet Packet with min length of 64 bytes and 8 bytes Preamble+SFD, and with min 12 bytes Interframe GAPS. It means that the example of Alternating data and control blocks in an 256/257 Block would not appeared!
 SuggestedRemedy
 Remove or modify the example!
 Response Response Status C
 ACCEPT IN PRINCIPLE.

Alternating control and data blocks can appear when errors are enforced during packet transmission. Refer to the possible transition between TX_D and TX_E states in Figure 82-14
 However, it would be better to an example that reflects a more common mapping. Change example three to be three data blocks followed by a control block.

Cl 91 SC 91.5.2.6 P 95 L 26 # 156
 Ran, Adeo Intel
 Comment Type ER Comment Status A
 This subclause describes the mapping operation but it is unclear how the mapped markers are re-inserted into the normal stream, paired with their removal in clause 91.5.2.4.
 SuggestedRemedy
 A figure showing the input and output of these two operations is required. Unfortunately I do not understand the proposed procedure enough to provide it.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Figure 91-4 was intended to be the requested illustration.
 See comment #150.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.2.6 P 95 L 40 # 54
 Szczepanek, Andre Inphi
 Comment Type TR Comment Status D bucket
 The upper limit of the range of variable "j" is wrong.
 The range of j should be 0 to 4 consistent with the 5 AMs per row shown in Figure 91-4
 SuggestedRemedy
 Replace "j=0 to 5" with "j=0 to 4"
 Proposed Response Response Status W
 PROPOSED ACCEPT.
 [Added Clause (91) to Sbcl field for consistent sorting.]
 See comment #472.

Cl 91 SC 91.5.2.6 P 95 L 40 # 472
 Cideciyan, Roy IBM
 Comment Type TR Comment Status D bucket
 j should run from 0 to 4
 SuggestedRemedy
 Given i=0, j=0 to 4, and x=i+4j, ...
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 91 SC 91.5.2.6 P 95 L 40 # 163
 Ran, Adee Intel
 Comment Type TR Comment Status D bucket
 x should takes PCS lane values (0..19), but if j=0..5 and i=0..3, x=i+4j can take values from 0 to 23. Seems that j should be only within 0..4.
 SuggestedRemedy
 Change "j=0 to 5" to "j=0 to 4".
 Proposed Response Response Status W
 PROPOSED ACCEPT.
 [Comment was entered against Subcl 91.5.2.5, but is actually against 91.5.2.6.]
 See comment #472.

Cl 91 SC 91.5.2.6 P 95 L 45 # 57
 Szczepanek, Andre Inphi
 Comment Type ER Comment Status A
 This mapping process really needs a diagram to show what is going on.
 A mapping equation though succinct is not descriptive.
 A diagram was provided in gustlin_01_0312, why not use it.
 SuggestedRemedy
 Add mapping diagram based on slide 15 of gustlin_01_0312.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 [Added Clause (91) to Sbcl field for consistent sorting.]
 Figure 91-4 was included for this purpose.
 See comment #150.

Cl 91 SC 91.5.2.6 P 95 L 50 # 150
 Ran, Adee Intel
 Comment Type E Comment Status A
 The 5-bit pad should better be depicted in figure 91-4 or elsewhere to show the five 257-bit blocks structure.
 SuggestedRemedy
 Preferably, update figure 91-4.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Augment Figure 91-4 to show the inclusion of the 5-bit pad and the transition into the next 257-bit block.
 Also clarify the assignment of pad bits in the text.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.2.6 P 95 L 51 # 463
 Cideciyan, Roy IBM
 Comment Type T Comment Status A
 am_txmapped<1284:1280> contains 5 bits whereas 0x05 and 0x1A contain 8 bits. Therefore, the notation is not very clear.
 SuggestedRemedy
 Replace 0x05 by 00101 and 0x1A by 11010
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 [Commenter did not specify CommentType. Set to T.]
 Given previous comments on the ambiguity of assignment of elements of binary array to a vector variable x<i:j>, the assignment needs to be further clarified.
 See comment #150.

Cl 91 SC 91.5.2.7 P 97 L 33 # 48
 Szczepanek, Andre Inphi
 Comment Type ER Comment Status D bucket
 Why do we refer to w-bit symbols rather than 10bit symbols. The rest of this clause has been written on the basis of 10bit symbols, So "w" is not a variable.
 SuggestedRemedy
 Replace "GF(2^w) where w=10 is the symbol size in bits" with "GF(2^10) where the symbol size is 10 bits"
 Proposed Response Response Status W
 PROPOSED ACCEPT.
 [Added Clause (91) to Sbcl field for consistent sorting.]
 Substitute the value 10 for all instances of w in Clause 91.

Cl 91 SC 91.5.2.6 P 96 L 48 # 182
 Slavick, Jeff Avago Technologies
 Comment Type E Comment Status D bucket
 Figure 91-3. Header bit for a All Control blocks TC block is 0, not 1.
 SuggestedRemedy
 Change the 1 in the 0 bit location of tx_xcoded to a 0 for example 4.
 Proposed Response Response Status W
 PROPOSED ACCEPT.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.2.7 P97 L 41 # 443
 Dawe, Piers IPtronics

Comment Type T Comment Status A

As well as telling us the error correction capability, please tell us the error detection capability of these codes. Also, while a code may be capable of something, the spec needs to say what an implementation must do.

SuggestedRemedy

Add text giving the error detection capability of these codes, and the expected/required error correction and detection capability of implementations.

Response Response Status C

ACCEPT IN PRINCIPLE.

The error detection capability of a bounded distance decoder is $(n-k) = 2^t$ symbols. For $(n-k+1)$ or more symbol errors, there is a chance that the decoder will incorrectly recognize the input as a different codeword. In these cases, it is only possible to bound the probability that errors will be detected (see [1]). Methods that achieve this require one additional codeword of decoding latency. However, there are other methods of error detection that offer reduced latency but are not guaranteed to detect all uncorrectable errors. There is no intention to preclude such methods.

The statement of error correcting capability was intended to establish the relevance of the parameter t . Since 91.5.2.7 specifies the operation of the encoder, decoder requirements should not be added here.

76.3.3.3 states the following:

"Implementations shall be capable of correcting up to 16 symbols in a codeword and detecting uncorrectable codewords."

Using this as a model, add the following paragraph after the first paragraph of 91.5.3.3.

"When used to form a 100GBASE-CR4 or 100GBASE-KR4 PHY, the RS-FEC sublayer shall be capable of correcting any combination of up to $t=7$ symbol errors in a codeword. When used to form a 100GBASE-KP4 PHY, the RS-FEC sublayer shall be capable of correcting any combination of up to $t=15$ symbol errors in a codeword. The RS-FEC sublayer shall also be capable of detecting uncorrectable codewords."

In 91.5.2.7, remove "This code has the capability to correct any combination of $t=?$ symbols errors in a codeword." These two sentences are redundant with the information proposed to be added to 91.5.3.3.

[1] R. J. McEliece and L. Swanson, "On the decoder error probability for Reed-Solomon codes," IEEE Trans. Inform. Theory, vol. 32, pp. 701-703, Sep. 1986.

Cl 91 SC 91.5.2.7 P98 L 1 # 465
 Cideciyan, Roy IBM

Comment Type ER Comment Status D bucket

Typographical error

SuggestedRemedy

Replace "polynomial" by "polynomial"

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 91 SC 91.5.2.7 P98 L 12 # 466
 Cideciyan, Roy IBM

Comment Type ER Comment Status D bucket

Typographical error

SuggestedRemedy

Replace "whose the coefficients" by "whose coefficients"

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 91 SC 91.5.2.7 P98 L 23 # 467
 Cideciyan, Roy IBM

Comment Type ER Comment Status D bucket

Missing blank

SuggestedRemedy

Insert blank between "... is transmitted last." and "The first bit ..."

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.2.7 P98 L47 # 59
 Szczepanek, Andre Inphi

Comment Type ER Comment Status A

Why are the generator polynomial coefficients relegated to a (presumably informative) annex ?.
 Although they can be derived from field polynomial and number of check symbols this requires a good bit of maths. So why not state them here. The coefficients are normative after all, there is no discretion in their values.

SuggestedRemedy

Add list of generator polynomial coefficients for the two FEC codes, in a format consistent with Figure 91-5.

Response Response Status C

ACCEPT.

[Added Clause (91) to Sbcl field for consistent sorting.]

See comment #234.

Cl 91 SC 91.5.2.7 P99 L1 # 234
 Healey, Adam LSI Corporation

Comment Type T Comment Status A

The RS-FEC encoding is sufficiently stable to define the generator polynomial coefficients and example codewords to assist users of the standard.

SuggestedRemedy

Add Annex 91A with FEC codeword examples in the style of Annex 74A. Include coefficients of the generator polynomial, gi, in Clause 91 or in the proposed annex.

Response Response Status C

ACCEPT IN PRINCIPLE.

Remove the editor's note. Add a table to the end of 91.5.2.7 that defines the coefficients of the generator polynomials for 100GBASE-KR4 and 100GBASE-KP4.

Add Annex 91A which includes an example of an FEC codeword (input, transcoded output, FEC encoded output).

Refer to langhammer_3bj_01_0912 for a C model of the encoders. These will also be included in the Annex.

Cl 91 SC 91.5.2.8 P99 L13 # 151
 Ran, Adeee Intel

Comment Type E Comment Status D bucket

A cross-reference to the relevant place in clause 94 could be useful.

SuggestedRemedy

After "When used to form a 100GBASE-KP4 PHY" add " (refer to 94.2.1.1.1)".

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Add "(refer to 94.2.1.1)" to the end of the first sentence.

In 91.5.3.1, add "(refer to 94.2.1.2)" to the end of the last sentence of the last paragraph.

Cl 91 SC 91.5.2.8 P99 L9 # 474
 Cideciyan, Roy IBM

Comment Type TR Comment Status D bucket

There is no scrambler at Tx of RS-FEC.

SuggestedRemedy

Replace "Once the data is scrambled and encoded, ..." by "Once the data is transcoded and encoded, ..."

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

See comment #183.

Cl 91 SC 91.5.2.8 P99 L9 # 183
 Slavick, Jeff Avago Technologies

Comment Type E Comment Status D bucket

We no longer are scrambling the data within the RS-FEC

SuggestedRemedy

Remove the words "scrambled and" along with the comma after encoded. In the first sentence of 91.5.2.8

Remove the words "descrambling and" from the last sentence in 91.5.3.4

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Change to:

"Once the data has been Reed-Solomon encoded, it shall..."

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.2.8 P99 L9 # 498
 Dawe, Piers IPtronics

Comment Type T Comment Status D bucket

This says "Once the data is scrambled and encoded" yet I can't see any mention of scrambling on the Tx side, nor de-scrambling the 58-bit scrambler in Clause 82. On the receive side, I can see that three bits in 257 are sometimes descrambled and three are scrambled. Also that the received first nibble is scrambled (where were they scrambled?). In 91.5.3.6 receive block distribution, "Once the data is encoded and scrambled" - I wouldn't say the data is scrambled. First, I would not call it data because it should consist of data blocks and also control blocks. Second, if only three block type bits in 66? are scrambled, it would be misleading to imply the whole stream is scrambled.

SuggestedRemedy

Does the Tx process scramble or not? Make the next draft clearer.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Scrambling/descrambling was removed from the RS-FEC sublayer. The paragraph must be updated to reflect this.

See comment #183.

Cl 91 SC 91.5.3.1 P99 L31 # 49
 Szczepanek, Andre Inphi

Comment Type ER Comment Status A

"FEC Deskew state diagram" is a misnomer. The SM shown in Figure 91-9 has very little to do with deskew (despite inheriting the functions of Figure 82-12), instead it is all about verifying FEC block Lock.

The functions of FEC lane deskew and testing for FEC block lock are functionally independent and will be implemented at quite different positions in the datapath and possibly in different clock regimes.

I see no real need to combine these two functions into one SM. Why not just re-use Figure 82-12 as is for FEC lane deskew, and provide a separate FEC block Lock SM.

SuggestedRemedy

Replace Figure 91-1 with a copy of Figure 82-12.
 Edit existing Figure 91-1 to use the "align_status" output from the deskew lock SM.

Response Response Status C

ACCEPT IN PRINCIPLE.

[Added Clause (91) to Sbcl field for consistent sorting.]

It is true that the actual "deskew" operation is a small portion of the state diagram and the majority of the functionality pertains to monitoring whether or not proper FEC codewords are being received.

A stand-alone FEC deskew state diagram would be trivial. Relative placement of deskew and FEC decode blocks, clock domains, etc. are implementation-specific considerations that should have little bearing on this generalized description of the required behavior.

From a behavioral point of view, defining operations for each FEC lane (Figure 91-8) and operations for the aggregate (deskew or "lane alignment", error monitoring) is a reasonable way to partition the problem. Both aspects are required to establish and monitor FEC codeword lock.

To avoid giving undue weight to the deskew operation, rename Figure 91-9 to be the "FEC alignment state diagram".

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.3.1 P99 L 32 # 26
 Anslow, Pete Ciena

Comment Type T Comment Status D

This says "The FEC receive function shall support a maximum Skew of 134 ns between FEC lanes and a maximum Skew Variation of 3.4 ns."

These are the skew and skew variation requirements at SP4 which is the input of the PMD sublayer, but they should be the values at SP5 which is at the output of the PMD sublayer as per the new Figure 80-5a

SuggestedRemedy

Change to:
 "The FEC receive function shall support a maximum Skew of 145 ns between FEC lanes and a maximum Skew Variation of 3.6 ns."

Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.

See healey_02_0912.pdf.

Cl 91 SC 91.5.3.2 P99 L 42 # 453
 Dawe, Piers IPtronics

Comment Type TR Comment Status D bucket

The medium is allowed to mix the lanes up, that's no error. See 86.6 Lane assignments

SuggestedRemedy

Delete "due to connection errors in the underlying medium".

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 91 SC 91.5.3.2 P99 L 42 # 152
 Ran, Adee Intel

Comment Type E Comment Status D bucket

If lane reordering is mandatory then physical lane swapping should not be considered an error. For some media this may happen intentionally and consistently.

Compare to 82.2.13 where the reason for possible re-ordering is stated as "due to Skew between lanes and multiplexing by the PMA". No "error" is mentioned.

SuggestedRemedy

Change "due to connection errors in the underlying medium" to "due to possible swapping in the underlying medium".

Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.

See comment #453.

Cl 91 SC 91.5.3.2 P99 L 43 # 50
 Szczepanek, Andre Inphi

Comment Type ER Comment Status A

Where is the FEC lane number defined ?. Stating "The FEC lane number is defined by the sequence of alignment markers mapped to each FEC lane" only tells half the story.

SuggestedRemedy

Explicitly state that FEC lane number zero is the lane that carries AM_0, lane 1 AM_1, lane 2 AM_2, and lane 2 AM_3.

Response Response Status C
 ACCEPT IN PRINCIPLE.

[Committer submitted the comment against Clause 99. Changed to Clause 91. Added Clause to Sbcl field for consistent sorting.]

The other half of the story is in 91.5.2.6 and Figure 91-4.

In 91.5.3.2, add a cross-reference to 91.5.2.6 at the end of the last sentence of the first paragraph.

In 91.5.2.6, state that alignment marker payloads corresponding to PCS lanes 0, 4, 8, 12, and 16 correspond to FEC lane 0, alignment marker payloads corresponding to PCS lanes 1, 5, 9, 13, and 17 correspond to FEC lane 1, and so on see Figure 91-4).

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.3.3 P 101 L 10 # 475
 Cideciyan, Roy IBM
 Comment Type TR Comment Status D bucket
 Suggestion to increase clarity and change from passive form to active form. Minimum frame size is 64 bytes. Minimum packet size, I believe, is 64+8=72 bytes.
 SuggestedRemedy
 "This will cause the PCS to discard all frames 64 bytes and larger that are fully or partially within the uncorrectable codeword."
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 91 SC 91.5.3.3 P 101 L 10 # 468
 Cideciyan, Roy IBM
 Comment Type ER Comment Status D bucket
 64-bytes should not be one word. It is not used as an adjective in this sentence.
 SuggestedRemedy
 Replace "64-bytes" by "64 bytes".
 Proposed Response Response Status W
 PROPOSED ACCEPT.
 See comment #475.

Cl 91 SC 91.5.3.3 P 101 L 11 # 186
 Slavick, Jeff Avago Technologies
 Comment Type T Comment Status A
 Ability to bypass the FEC correction function is not defined.
 SuggestedRemedy
 Add the following text to 91.5.3.3
 When fec_bypass_correction is set true and the incoming parity of the codeword does not match the received parity the decoder shall corrupt the codeword in the same manner as if an uncorrectable codeword was received.
 Added an MDIO register bit to control fec_bypass_correction
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 While gustlin_01a_0712 discusses the possibility that an implementation may choose to disable error correction to reduce latency when the operating conditions allow it, it was not proposed that implementations are required to do so or to expose this feature via a management variable.
 However, after discussion, it was decided that this feature should be an option and an ability bit will be added in addition to the proposed enable bit.
 The management variables are described in healey_3bj_02_0912. Add corresponding text to 91.5.3.3 describing the option.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.3.3 P 101 L 6 # 55
 Szczepanek, Andre Inphi

Comment Type TR Comment Status R

"If the decoder determines that a codeword is uncorrectable, it shall"
 What is the definition of uncorrectable ?
 This is important as it has a "shall" tied to it.
 Without a definition of "uncorrectable" how can we determine compliance

SuggestedRemedy

Add the following definition of an uncorrectable 802.3bj codeword.
 An uncorrectable codeword is a codeword whose error locator polynomial has a degree greater than 7 (t), or where the error locator or error evaluator polynomials cannot be determined (The key equation cannot be solved).

This definition provides a definitive minimum requirement for codeword marking.

Response Response Status C

REJECT.

[Added Clause (91) to Sbcl field for consistent sorting.]

The commenter defines the term "uncorrectable codeword" while introducing three new undefined terms ("error locator polynomial", "error evaluator polynomial", and "key equation"). This is not an equitable trade.

After discussion, it was clear that there were multiple approaches to error detection that offer trade-offs between coverage or latency. This is an implementation specific issue that should not be constrained by the draft.

See comment #443.

Cl 91 SC 91.5.3.3 P 101 L 6 # 241
 Healey, Adam LSI Corporation

Comment Type T Comment Status A

Clause 74 error marking is optional presumably due to its impact on latency. What is the latency impact of the error marking specified in this subclause?

If the increase is significant, consider optional error marking for Clause 91.

SuggestedRemedy

Evaluate the impact of error marking on latency and determine whether or not the feature should be optional.

Response Response Status C

ACCEPT IN PRINCIPLE.

Make error marking optional. Modify text in 91.5.3.3 to indicate this. Add "error indication" ability and enable bits to management per healey_3bj_02_0912.

It should be noted that deactivating error marking would have an adverse impact on MTTFPA

As stated in the comment, the other consideration for error marking is any added latency which is discussed in the context of comment #190.

Cl 91 SC 91.5.3.4 P 101 L 17 # 476
 Cideciyan, Roy IBM

Comment Type TR Comment Status D bucket

Data is not descrambled prior to transcoding at Rx.

SuggestedRemedy

Replace "... prior to descrambling and transcoding." by "... prior to transcoding."

Proposed Response Response Status W

PROPOSED ACCEPT.

See comment #51.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.3.4 P 101 L 17 # 51
 Szczepanek, Andre Inphi
 Comment Type ER Comment Status D bucket
 Descrambling no longer forms part of the receive datapath.
 SuggestedRemedy
 Remove "descrambling and"
 Proposed Response Response Status W
 PROPOSED ACCEPT.
 [Added Clause (91) to Sbcl field for consistent sorting.]

Cl 91 SC 91.5.3.4 P 101 L 18 # 242
 Healey, Adam LSI Corporation
 Comment Type T Comment Status A
 This subclause does not address the case where rapid alignment markers are being received
 SuggestedRemedy
 Modify the subclause to address both normal and rapid alignment markers.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Grant editorial license to craft to text to be consistent with changes to EEE functionality suggested by other comments.
 See comment #243.

Cl 91 SC 91.5.3.5 P 101 L 25 # 157
 Ran, Adeee Intel
 Comment Type ER Comment Status A
 Assuming rx_rxcoded<4:0> in this line is a typo, then rx_xcoded<4:0> is assigned twice. This can be confusing.
 It would be preferred to define another variable rx_xcoded_header and use it as in my comment on subclause 91.5.2.5.
 SuggestedRemedy
 Change this paragraph to:
 "Set rx_xcoded_header<4:0> to the result of the bit-wise exclusive-OR of rx_xcoded<4:0> and rx_xcoded<12:8>".
 Use rx_xcoded_header<0> instead of rx_xcoded<0>, and rx_xcoded_header<j+1> instead of rx_xcoded<j+1> in the following steps.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Add the following sentence to the end of the first paragraph of 91.5.3.3.
 "The message symbols correspond to 20 transcoded blocks rx_scrambled."
 In the first paragraph of 91.5.3.5, change reference to rx_xcoded<256:0> to rx_scrambled<256:0>.
 Replace the second paragraph of 91.5.2.5 with following.
 "First, descramble rx_scrambled<256:0> to yield rx_xcoded<256:0> as follows.
 a) Set rx_xcoded<4:0> to the result of the bit-wise exclusive-OR of the rx_scrambled<4:0> and rx_scrambled<12:8>.
 b) Set rx_xcoded<256:5> to rx_scrambled<256:5>."
 In Figure 91-6, replace rx_xcoded with rx_scrambled.

Cl 91 SC 91.5.3.5 P 101 L 25 # 477
 Cideciyan, Roy IBM
 Comment Type TR Comment Status D bucket
 Notation not correct
 SuggestedRemedy
 Replace "rx_rxcoded<4:0>" by "rx_xcoded<4:0>".
 Proposed Response Response Status W
 PROPOSED ACCEPT.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.3.5 P 101 L 39 # 52
 Szczepanek, Andre Inphi

Comment Type ER Comment Status R

The function for re-insertion of the first codeword "s" nibble is unnecessarily terse and makes it difficult to understand what is required. As c only has 4 possible values, why not just state all 4 possible bit muxes. In order to understand what is going the reader will have to calculate these four bit muxes - so why not do it for them.

SuggestedRemedy

Replace :

d)let rx_payloads be a vectorrepresenting the payloads of the four 66-bit blocks. It is derived using the following expressions:

rx_payloads<(64c+3):0> = rx_xcoded<(64c+8):5>
 rx_payloads<(64c+7):(64c+4)> = 0000 (an arbitrary value that is later replaced, see step j)
 rx_payloads<255:(64c+8)> = rx_xcoded<256:(64c+9)>

With :

d)let rx_payloads be a vectorrepresenting the payloads of the four 66-bit blocks. It is derived using the following expressions:

if (c==0) rx_payloads <255:0> = rx_xcoded<256:9> :: 4'b000 :: rx_xcoded <8:5>
 if (c==1) rx_payloads <255:0> = rx_xcoded<256:73> :: 4'b000 :: rx_xcoded <72:5>
 if (c==2) rx_payloads <255:0> = rx_xcoded<256:137> :: 4'b000 :: rx_xcoded <136:5>
 if (c==3) rx_payloads <255:0> = rx_xcoded<256:201> :: 4'b000 :: rx_xcoded <200:5>
 where 4'b000 is an arbitrary value that will be replaced later in step j

Response Response Status C

REJECT.

[Commenter submitted this comment against Clause 00. Changed to Clause 91, Subcl 91.5.3.5, Page 101, Line 39.]

The text is correct as written. Illustrations have been added (see Figure 91-3) to help the reader understand the process.

The suggested remedy includes notation for array concatenation "::" and definition of binary vectors 4b'xxxx, that is not used elsewhere in IEEE 802.3. The existing definition does not require new array concatenation notation.

While the mathematical description is precise, it requires the user to do a number of index computations to understand the construction of the codeword. It is not clear why the calculations involving the variable c are more onerous than the others.

Cl 91 SC 91.5.3.5 P 101 L 45 # 164
 Ran, Adee Intel

Comment Type TR Comment Status R

According to accepted change in transcoding (gustlin_02_0712) there is no additional scrambling following transcoding. Unscrambling described in step g does not seem to have a counterpart in the original 64B/66B to 256B/257B transcoding procedure in 91.5.2.5.

SuggestedRemedy

Delete steps f and g?
 Make sure this clause describes exactly the inverse operation of 91.5.2.5.

Response Response Status C

REJECT.

The 64B/66B to 256B/257B transcoder (see 91.5.2.5) removes 4 scrambled bits from the input 66-bit blocks (if any of the blocks are control blocks). The 256B/257B to 64B/66B transcoder must restore these bits, scrambled in a manner consistent with the surrounding bits, to produce valid 66B blocks.

To restore the bits, the decoder must first descramble the first nibble in order to determine what the second nibble should be (step f). It must then scramble the second nibble based on the learned scrambler state (step g).

The steps are integral to the processing defined in gustlin_02_0712 and adopted via Draft 1.0 comment #70. They will not be deleted.

Cl 91 SC 91.5.3.6 P 102 L 9 # 478
 Cideciyan, Roy IBM

Comment Type TR Comment Status D bucket

Encoding and scrambling is not performed at Rx.

SuggestedRemedy

Replace "Once the data is encoded and scrambled, it shall ..." by "Once the data is decoded and transcoded, it shall ..."

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Change to:

"After the data has been transcoded, it shall be distributed to multiple PCS lanes, one 66-bit block at a time..."

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.3.7 P 102 L 16 # 480
 Cideciyan, Roy IBM

Comment Type TR Comment Status D bucket

There may be errors at the RS decoder output. Therefore, am_x and am_payloads in Section 91.5.2.6 does not have to be the same as am_x and am_payloads in Section 91.5.3.7

SuggestedRemedy

In Section 91.5.2.6 replace am_x and am_payloads by am_tx and am_txpayloads
 In Section 91.5.3.7 replace am_x and am_payloads by am_rx and am_rxpayloads

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

In 91.5.2.6, change am_x to am_tx_x and am_payloads to am_txpayloads.

In 91.5.2.6, change am_x to am_rx_x and am_payloads to am_rxpayloads.

The notation is changed from the suggested remedy to clearly separate "tx" and "rx" from the variable "x" (PCS lane number).

Cl 91 SC 91.5.3.7 P 102 L 27 # 479
 Cideciyan, Roy IBM

Comment Type TR Comment Status D bucket

j runs from 0 to 4

SuggestedRemedy

Given i=0 to 3, j=0 to 4, and x=i+4j, the ...

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 91 SC 91.5.4.2.1 P 104 L # 211
 Sela, Oren Mellanox Technologies

Comment Type E Comment Status A

There are many variables that have the same name in CL82 and may cause unnecessary confusion.

SuggestedRemedy

Change the naming:
 align_status --> RS_FEC_align_status
 alignment_valid --> RS_FEC_alignment_valid
 all_locked --> amps_all_locked
 enable_deskew --> RS_FEC_enable_deskew

Response Response Status C

ACCEPT IN PRINCIPLE.

Some variable names clash with those incorporated by reference (see 91.5.2.1 and 91.5.2.2).

Change the following variable names:

align_status to fec_align_status
 alignment_valid to fec_alignment_valid
 enable_deskew to fec_enable_deskew

all_locked is not a variable name in Clause 82 and does not require change.

Cl 91 SC 91.5.4.2.1 P 104 L # 209
 Sela, Oren Mellanox Technologies

Comment Type T Comment Status A

restart_lock variable is not defined in the varabile section

SuggestedRemedy

add restart_lock definition

Response Response Status C

ACCEPT IN PRINCIPLE.

Define restart_lock as follows (do not include text in <>):

Boolean variable that is set by the FEC alignment <see comment #49> process to reset the synchronization process on all FEC lanes. It is set to true after 3 consecutive uncorrectable codewords are received (3_BAD state) and set to false upon entry into the LOSS_OF_ALIGNMENT state.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.4.2.1 P 104 L 16 # 495
 Dawe, Piers IPtronics
 Comment Type E Comment Status R late
 I can't see the difference between align_status (true when all lanes are synchronized and aligned) and alignment_valid. I think they can be the same.
 SuggestedRemedy
 Combine them into one variable, or if not, add text to explain why there are two/what the difference is.
 Response Response Status C
 REJECT.
 This portion of the state diagram (and corresponding variables) is similar to what is used in the PCS deskew state diagram (refer to 82-12). There is no clear incentive to deviate from this familiar form.

Cl 91 SC 91.5.4.2.1 P 104 L 26 # 213
 Sela, Oren Mellanox Technologies
 Comment Type ER Comment Status D bucket
 typo - am_lock<x> should be amps_lock<x>
 SuggestedRemedy
 Change:
 "A Boolean variable that is set to true when amps_lock<x> is true for all x and is set to false when am_lock<x> is false for any x."
 "
 To:
 "A Boolean variable that is set to true when amps_lock<x> is true for all x and is set to false when amps_lock<x> is false for any x."
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 91 SC 91.5.4.2.1 P 104 L 39 # 243
 Healey, Adam LSI Corporation
 Comment Type T Comment Status A
 How does the RS-FEC sublayer discriminate between normal operation and the optional EEE capability? The intent of this statement is to specify that the state diagram behaves one way when normal alignment markers are expected but behaves a different way when rapid alignment markers are expected.
 The RS-FEC sublayer should use the EEE service interface primitives defined in 91.2 to determine if normal or rapid alignment markers are expected.
 SuggestedRemedy
 Tie the behavior of the state diagram to the EEE service interface primitives defined in 91.2.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Implement changes for the optional EEE capability per healey_3bj_02_0912.

Cl 91 SC 91.5.4.2.1 P 104 L 46 # 225
 Gustlin, Mark Xilinx
 Comment Type T Comment Status D
 This editor's note can be removed, Zhongfeng Wang has looked at this and the current SM is sufficiently robust for KP4 also.
 SuggestedRemedy
 Per the comment.
 Proposed Response Response Status Z
 REJECT.
 This comment was WITHDRAWN by the commenter.

Cl 91 SC 91.5.4.2.1 P 105 L 3 # 469
 Cideciyan, Roy IBM
 Comment Type ER Comment Status D bucket
 typographical error
 SuggestedRemedy
 Replace "maker" by "marker"
 Proposed Response Response Status W
 PROPOSED ACCEPT.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91.5.4.2.1 P 105 L 54 # 208
Sela, Oren Mellanox Technologies

Comment Type T Comment Status A

Also for the optional EEE capability, if first_amp corresponds to PCS lane 16, 17, 18, or 19, this counter counts the 4096 FEC codewords minus 256 bits to the end of the expected location of the next alignment marker payload corresponding to PCS lanes 0, 1, 2, or 3

This means that for waking in up from EEE the 4096 FEC block time is longer than the RAMs - meaning that it will also take longer for the PCS to lock

SuggestedRemedy

Option 1 -
Change amp_valid to look for lanes 0,1,2 or 3 only in FIND_1ST state for both EEE and normal mode, and to look for 16, 17,18 or 19 in COMP_2ND state for EEE.

Option 2-
Have the same behavior for normal and EEE mode for the amp_valid and amp_counter should be 4096 FEC codewords when rx_mode = data and 8 FEC codewords when rx_mode != data.

If option 1 is chosen then the AMP_COMPARE should be changed so that for EEE amp_match should be set to true if current_pcs1 = first_pcs1+16 only
If option 2 is chosen then AMP_COMPARE should change so that - if current_pcs1 equals first_pcs1, amp_match is set to true - is applicable for both EEE and normal mode

Response Response Status C

ACCEPT IN PRINCIPLE.

The definition of amp_counter is incorrect. During low power idle, if first_amp corresponds to PCS lanes 16, 17, 18, or 19, amp_counter should count 2 FEC codewords minus 256 bits to the end of the expected location of the next alignment marker payload corresponding to PCS lanes 0, 1, 2, or 3.

See also comment #243.

Cl 91 SC 91.5.4.2.1 P 107 L 3 # 199
Slavick, Jeff Avago Technologies

Comment Type T Comment Status A

Figure 91-8. The variable restart_lock is not defined in the State Variables section.

SuggestedRemedy

Add a definition for restart_lock to 91.5.4.2.1

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #209.

Cl 91 SC 91.5.4.2.3 P 106 L 3 # 204
Slavick, Jeff Avago Technologies

Comment Type T Comment Status D bucket

The term first_amp is used but the variable name is first_psc1

SuggestedRemedy

Change all first_amp references to first_psc1 in the amp_counter definition.

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 91 SC 91.5.4.3 P 107 L 3 # 226
Gustlin, Mark Xilinx

Comment Type T Comment Status A

The signal restart_lock is not a defined variable. Add it to the list of variables.

SuggestedRemedy

Per the comment.

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #209.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

CI 91 SC 91.5.4.3 P 108 L 37 # 205
 Slavick, Jeff Avago Technologies

Comment Type T Comment Status A

Figure 91-9. The transition out of TEST_CW should be gated by a new codeword being available instead of gating the exit from a cw_bad_count adjustment state being gated.

SuggestedRemedy

Change the following state transitions to be:
 TEST_CW -> CW_GOOD: test_cw & !cw_bad
 TEST_CW -> CW_BAD: test_cw & cw_bad
 CW_GOOD -> TEST_CW: UCT
 CW_BAD -> TEST_CW: cw_bad_count < 3

Response Response Status C

ACCEPT IN PRINCIPLE.

[Added Subcl 91.5.4.3 for consistent sorting.]

The Suggested Remedy would cause the first codeword received after ALIGN_REQUIRED to not be considered in cw_bad_count. Otherwise, there is no difference between the existing state diagram and proposed modifications.

The problem with existing state diagram is not made clear. This is the form used in clauses 41 and 82. There is no obvious advantage to the suggested remedy.

However, in the course of considering this comment, two errors were found. In Figure 91-8, test_amp should be assigned the value FALSE in the LOCK_INIT state. In Figure 91-9, test_cw should be assigned the value FALSE in the ALIGN_ACQUIRED state. Add the assignments to the corresponding state diagrams.

CI 91 SC 91.6 P 108 L 52 # 244
 Healey, Adam LSI Corporation

Comment Type T Comment Status A

The RS-FEC architecture has stabilized to the point where MDIO status and control variables can be defined.

SuggestedRemedy

Include tables defining RS-FEC status and control variables and amend Clause 45 accordingly.

Response Response Status C

ACCEPT IN PRINCIPLE.

Refer to comment #196.

CI 91 SC 91-2 P 94 L # 207
 Sela, Oren Mellanox Technologies

Comment Type T Comment Status R

In the receive path should merge the alignment lock and deskew block with the Lane reorder block - all 3 action are done by acquiring FEC block lock based on the alignment markers. Also this will make is consistent with Figure 91-7

SuggestedRemedy

Create one block "alignment lock, deskew and lane reorder" to replace the 2 blocks in the receive path in figure 91-2

Response Response Status C

REJECT.

Figure 91-7 is intended to describe bit order and for that purpose there was no advantage to showing "lane reorder" as a separate block.

Figure 91-2 is partitioned to correspond with the organization of subclauses.

Lane reordering is not needed to obtain alignment lock. Lane reordering is needed to verify that valid codewords are being received after alignment lock which requires information from the Reed-Solomon decoder. Therefore, even with the proposed consolidation, the functions are still not self-contained.

For these reasons the partition will remain as is.

CI 91 SC 91-8 P 107 L # 210
 Sela, Oren Mellanox Technologies

Comment Type T Comment Status A

The FEC synchronization state diagram doesn't take into account the fast lock needed for EEE wakeup from LPI QUIET - need to specify that amp_count should count 4096 FEC codeword when rx_mode is DATA and 8 FEC codeword when rx_mode is not DATA.

SuggestedRemedy

per comment

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #243.

IEEE P802.3bj D1.0 100 Gb/s Backplane and Copper Cable 1st Task Force review comments

Cl 91 SC 91-9 P 108 L # 212
 Sela, Oren Mellanox Technologies

Comment Type E Comment Status A

The name: "FEC deskew" is not the right name for that diagram. This diagram doesn't only enable/disable deskew but also monitors the FEC block lock

SuggestedRemedy

Change the name of the Figure to: "FEC block lock state diagram" or "FEC block lock and deskew state diagram"

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #49.

Cl 91 SC Figure 91-4 P 97 L 4 # 58
 Szczepanek, Andre Inphi

Comment Type ER Comment Status A

This figure describes the mapping process specified on line 43 page 95, but the column heading description "Reed Solomon Symbol Index, k" does not relate to this mapping process

SuggestedRemedy

The columns should be labelled either by alignment marker column index "j" or by column (0 to 319). Better still with both as it makes the mapping easire to understand.

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #150.

Figure 91-4 illustrates the am_payloads matrix and "k" does indeed relate to the mapping per page 95, lines 45 to 48.

Cl 91 SC Figure 91-5 P 98 L 39 # 47
 Szczepanek, Andre Inphi

Comment Type ER Comment Status D bucket

Why do we refer to w-bit symbols rather than 10bit symbols.
 The rest of this clause has been written on the basis of 10bit symbols,
 So "w" is not a variable.

SuggestedRemedy

Replace "symbol delay element, holds 1 w-bit symbol"
 with "symbol delay element, holds 1 10-bit symbol"

Proposed Response Response Status W

PROPOSED ACCEPT.

See comment #48.

Cl 99 SC P 5 L 11 # 29
 Anslow, Pete Ciena

Comment Type E Comment Status D bucket

It is usual for amendments to 802.3 to include a short summary of their content immediately after the text that describes the sections of IEEE Std 802.3.
 This is missing from this draft.
 For example IEEE Std 802.3ap-2007 contained:
 IEEE Std 802.3ap-2007
 This amendment includes changes to IEEE Std 802.3-2005 and adds Clause 69 through Clause 74 and Annex 69A, Annex 69B, Annex 73A and Annex 74A. This amendment adds new Physical Layers that support the exchange of IEEE Std 802.3 format frames over electrical backplanes at 1 Gb/s and 10 Gb/s.

This paragraph will then also appear in the frontmatter of other amendments being developed such as 802.3bk

SuggestedRemedy

Add a paragraph describing 802.3bj

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

The frontmatter will be updated under the guidance of the Working Group chair.