

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

CI 91 SC P118 L14 # 67  
Pillai, Velu Broadcom

Comment Type E Comment Status R

Fig 91-2 does not show the BER Monitor in the transmit path.

SuggestedRemedy

Add a block to show the BER Monitor attached to the Alignment lock and deskew.

Response Response Status C

REJECT.

The BER monitor is not required by the "Lane block synchronization" or "Alignment lock and deskew" functions. In the Clause 82 PCS, its function is to inhibit the operation of the PCS Receive state diagram when the BER is too large to reliably determine synchronization. It therefore has no function in the Clause 91 RS-FEC sublayer.

CI 91 SC 3 P116 L37 # 295  
Ofelt, David Juniper Networks

Comment Type TR Comment Status R

The current draft indicates that the RS FEC is only supported on services interfaces with width (p) of 4.

This is overly restrictive and ensures that when we develop 2 and 1 physical lane interfaces that we'll need to rework this part of the standard. It is possible to bit-interleave the four lanes into two or one, but the result does not handle burst errors well. An argument that comes up is that "we'll only support muxing for interfaces that are more unlikely to have burst errors (e.g. no DFE)". This is unsatisfying to me- we have an architecture from .3ba that handles a large variety of interface structures and then we follow it with the next rev of the PCS where we remove all that good flexibility or we can support it for a subset of the interface schemes.

SuggestedRemedy

Add text to 91.3 indicating something like:

"If a PMA wants to multiplex the four FEC lanes into two or one lanes, then the multiplexing shall be done at a Reed-Solomon codeword boundary"

I believe this is the necessary requirement to make FEC work properly once multiplexed.

With this change, we should have the features needed to implement all optics variety being discussed in .3bm.

Response Response Status C

REJECT.

1. It is not clear what it means to multiplex "at a Reed-Solomon codeword boundary."
2. The requirement is incomplete because it requires that the PMA also identify "codeword boundaries" to correctly demultiplex them for presentation to the RS-FEC sublayer. This is a non-trivial function, as can be seen by the mechanism Clause 91 uses for this purpose, but is omitted from the proposed requirement.
3. The proposed normative requirement applies to a PMA and such requirements should appear in the PMA clause.
4. There is no Physical Layer defined in P802.3bj that requires this feature.

While this feature could extend the applicability of the RS-FEC sublayer to a PHY, yet to be defined, based on less than 4 physical lanes, the suggested remedy is not complete and perhaps misplaced. It seems that the objective of the proposal is to add a new PMA that multiplexes 10-bit Reed-Solomon symbols rather than bits which could be done in the context of that new PHY.

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

CI 91 SC 91.5.2.5 P119 L19 # 88  
Sela, Oren Mellanox Technologies

Comment Type E Comment Status A

In bullet c) there is a redundant statement. In line 14 we establish that all synch header are valid so there is no need to state that both  $c_{<0>} = 1$  and  $c_{<1>} = 0$  it is enough to say that  $c_{<0>} = 1$

SuggestedRemedy

change:

Let  $c$  be the smallest value of  $j$  such that  $tx\_coded\_c_{<0>}=1$  and  $tx\_coded\_c_{<1>}=0$ . In other words,  $tx\_coded\_c$  is the first 66-bit control block that was received in the current group of four blocks.

To:

Let  $c$  be the smallest value of  $j$  such that  $tx\_coded\_c_{<0>}=1$ . In other words,  $tx\_coded\_c$  is the first 66-bit control block that was received in the current group of four blocks.

Response Response Status C

ACCEPT.

CI 91 SC 91.5.2.5 P119 L31 # 89  
Sela, Oren Mellanox Technologies

Comment Type E Comment Status R

bullet b) - change to  $tx\_xcoded_{<4:0>}=1111$

SuggestedRemedy

per comment

Response Response Status C

REJECT.

The text is correct as written.

CI 91 SC 91.5.2.6 P120 L28 # 69  
Pillai, Velu Broadcom

Comment Type ER Comment Status D bucket

payloads corresponding to PCS lanes 1, 5, 6, 13, and 17 are

is not correct

SuggestedRemedy

It needs to be

payloads corresponding to PCS lanes 1, 5, 9, 13, and 17 are

Proposed Response Response Status W

PROPOSED ACCEPT.

CI 91 SC 91.5.2.6 P122 L19 # 72  
Pillai, Velu Broadcom

Comment Type T Comment Status R

Text talks about bit error monitoring, but there are no counters attached to this statment. Either we should add error counters or remove this line.

SuggestedRemedy

Response Response Status C

REJECT.

BIP errors are monitored by the alignment marker removal function and the corresponding counters are cited there (see 91.5.2.4).

The paragraph in 91.5.2.6 is an advisory to the user that, while the BIP fields are preserved by the mapping function defined in that subclause, they should NOT be used to monitor errors over the FEC-protected link.

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

Cl 91 SC 91.5.2.6 P122 L28 # 110  
Sela, Oren Mellanox Technologies

Comment Type T Comment Status A

The tx\_lpi\_active reference to 82.2.7a is no longer correct and should be referenced to the new figure 91-10

SuggestedRemedy  
per comment

Response Response Status C

ACCEPT IN PRINCIPLE.

The reference to 82.2.7a should have been 82.2.8a and pertain to the definition of Rapid Alignment Markers.

tx\_lpi\_active is set by the Transmit LPI state diagram in Figure 91-10.

Correct the cross-reference to be 82.2.8a.

Cl 91 SC 91.5.2.7 P123 L34 # 374  
Cideciyan, Roy IBM

Comment Type ER Comment Status D bucket

Figure 91-5 states "symbol delay element, holds 1 10-bit symbol". The formulation can be improved.

SuggestedRemedy

Replace "symbol delay element, holds 1 10-bit symbol" by "symbol delay element, holds a 10-bit symbol"

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Reduces the risk the someone could interpret it read "holds 110-bit symbol".

Cl 91 SC 91.5.3.3 P126 L16 # 376  
Cideciyan, Roy IBM

Comment Type TR Comment Status D

MTTFPA computations in cideciyan\_01\_0512.pdf always assume that RS decoder reports (indicates) errors to PCS layer whenever there is an uncorrectable code word (error correction mode) or code word contains errors (error detection mode). Therefore, indication of errors to the PCS sublayer is not an option but a mandatory feature of the RS decoder in order to have satisfactory MTTFPA.

SuggestedRemedy

Replace "The Reed-Solomon decoder may optionally provide ..." by "The Reed-Solomon decoder shall provide ..."

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

See comment #369.

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

Cl 91 SC 91.5.3.3 P126 L16 # 369  
 Anslow, Pete Ciena

Comment Type TR Comment Status D

This says that the indication of uncorrected errors to the PCS is optional. But if uncorrected errors are not indicated, the MTTFPA will be poor because any FEC frame with uncorrected errors will contain at least 8 or 16 errored symbols.

Doing a simple minded calculation:

If the errors turn up in bursts of 8, then a BER of 1E-12 is a block of errors every 80 seconds. The only thing stopping this from being accepted as a good packet is the CRC. This fails with a probability of 2.3E-10 which is a false packet every 10,000 years.

If the BER falls to 1E-6, this is a false packet every 4 days.

I think Roy Cideciyan has shown that reporting errors with FEC enabled gives a MTTFPA of better than 10,000 years at 1E-6.

This is a huge improvement in performance, so marking uncorrected errors should be mandatory.

SuggestedRemedy

Make the indication of uncorrected errors mandatory in Clause 91. Make the appropriate changes to the other clauses e.g. Clause 45

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Straw poll (Track 2 ad hoc):  
 Make error indication mandatory.  
 Agree: 5  
 Disagree: 2

=== Response if error indication is mandatory ===

Change the fourth paragraph of 91.5.3.3 as follows and consolidate it with the last paragraph.  
 "The Reed-Solomon decoder shall indicate errors to the PCS sublayer by intentionally corrupting 66-bit block synchronization headers."

Change the first sentence of the last paragraph of 91.5.3.3 to:  
 "When the decoder determines."

Remove the "FEC error indication enable" variable from Table 91-2 as well as 91.6.2.

Remove the "FEC error indication ability" variable from Table 91-3 as well as 91.6.4.

Update Clause 45 management and the Clause 91 PICS accordingly.

=== Response if error indication is optional ===

Add a note to 91.5.3.3.  
 "NOTE 1 -- The ability to disable error indication is provided for applications that require the lowest possible latency. It should be understood that the mean time to false packet acceptance (MTTFPA) will be greatly reduced when correction is bypassed and error indication is disabled."

Cl 91 SC 91.5.3.3 P126 L17 # 377  
 Cideciyan, Roy IBM

Comment Type TR Comment Status D

MTTFPA computations in cideciyan\_01\_0512.pdf always assume that RS decoder reports (indicates) errors to PCS layer whenever there is an uncorrectable code word (error correction mode) or code word contains errors (error detection mode). Therefore, indication of errors to the PCS sublayer is not an option but a mandatory feature of the RS decoder in order to have satisfactory MTTFPA.

SuggestedRemedy

Omit the following two sentences: "The presence of this option is indicated by the assertion ... (see 91.6.4). When the option is provided, it is enabled ... (see 91.6.2).

Proposed Response Response Status W

PROPOSED ACCEPT.

See comment #369.

Cl 91 SC 91.5.3.3 P126 L21 # 378  
 Cideciyan, Roy IBM

Comment Type TR Comment Status D

MTTFPA computations in cideciyan\_01\_0512.pdf always assume that RS decoder reports (indicates) errors to PCS layer whenever there is an uncorrectable code word (error correction mode) or code word contains errors (error detection mode). Therefore, indication of errors to the PCS sublayer is not an option but a mandatory feature of the RS decoder in order to have satisfactory MTTFPA.

SuggestedRemedy

Replace "When the error indication function is enabled and the decoder determines that a code word ..." by "When the decoder determines that a code word ..."

Proposed Response Response Status W

PROPOSED ACCEPT.

See comment #369.

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

CI 91 SC 91.5.3.3 P126 L22 # 3  
 Szczepanek, Andre Inphi

Comment Type TR Comment Status A

"or is uncorrectable"  
 See previous comment related to line 9 on the same page.

SuggestedRemedy

Replace "or is uncorrectable"  
 with  
 "or contains errors and has not been corrected"

Response Response Status C

ACCEPT IN PRINCIPLE.

[changed Subclause to 91.5.3.3 for consistent sorting.]

Change the beginning of the first sentence of the last paragraph of 91.5.3.3 to:  
 "When the error indication function is enabled and the decoder determines that a codeword contains errors (when the bypass correction feature is enabled) or contains errors but was not corrected (when the bypass correction feature is not supported or not enabled)."

See also comment #375.

CI 91 SC 91.5.3.3 P126 L23 # 375  
 Cideciyan, Roy IBM

Comment Type T Comment Status D bucket

The formulation "... not supported or enabled" does not seem to be clear.

SuggestedRemedy

Replace "... not supported or enabled), ..." by "... not supported or not enabled), ..."

Proposed Response Response Status W

PROPOSED ACCEPT.

CI 91 SC 91.5.3.3 P126 L23 # 113  
 Sela, Oren Mellanox Technologies

Comment Type T Comment Status R

Should allow an implementation to nullify more than one 64/66 block in every other transcoding block - for example an implementation should be able to nullify all blocks

SuggestedRemedy

change to:  
 ...it shall ensure that, at least for every other 257-bit block within the codeword starting with the first (1st, 3rd, 5th, etc.), the synchronization header for the first 66-bit block at the output of the 256B/267B to 64B/66B transcoder, rx\_coded\_0<1:0>, is set to 11. In addition, it shall ensure rx\_coded\_3<1:0> corresponding to the last (20th) 257-bit block in the codeword is set to 11. This will cause the PCS to discard all frames 64 bytes and larger that are fully or partially within the codeword. The decoder may set rx\_coded\_j<1:0> to 11 and thus nullify more 66-bit blocks at the PCS.

Response Response Status C

REJECT.

If an implementation were to invalidate the synchronization headers of all 66-bit blocks included in a codeword, the PCS would lose block lock and this would result in an extended loss of data.

The synchronization header error pattern was chosen to ensure no packet could be incorrectly accepted while maintaining block lock.

CI 91 SC 91.5.3.3 P126 L25 # 68  
 Pillai, Velu Broadcom

Comment Type E Comment Status D bucket

256B/267B to 64B/66B transcoder, rx\_coded\_0<1:0>

SuggestedRemedy

Needs to be

256B/257B to 64B/66B transcoder, rx\_coded\_0<1:0>, is s

Proposed Response Response Status W

PROPOSED ACCEPT.

[Changed Subcl from 91.5.3.4 to 91.5.3.3.]

See comment #379.

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

Cl 91 SC 91.5.3.3 P126 L25 # 117  
Sela, Oren Mellanox Technologies

Comment Type T Comment Status D bucket  
typo - replace 256B/267B with 256B/257B

SuggestedRemedy  
per comment

Proposed Response Response Status W  
PROPOSED ACCEPT.

See comment #379.

Cl 91 SC 91.5.3.3 P126 L25 # 379  
Cideciyan, Roy IBM

Comment Type TR Comment Status D bucket  
Transcoder in the receiver is 256B/257B to 64B/66B transcoder.

SuggestedRemedy  
Replace "256B/267B to 64B/66B transcoder" by "256B/257B to 64B/66B transcoder"

Proposed Response Response Status W  
PROPOSED ACCEPT.

Cl 91 SC 91.5.3.3 P126 L9 # 112  
Sela, Oren Mellanox Technologies

Comment Type T Comment Status A  
The RS-FEC can't detect all the uncorrectable codewords

SuggestedRemedy  
change:  
The RS-FEC sublayer shall also be capable of detecting uncorrectable codewords  
To:  
The RS-FEC sublayer shall also be capable of detecting some of the uncorrectable codewords

Response Response Status C  
ACCEPT IN PRINCIPLE.

Change the last sentence of the second paragraph of 91.5.3.3 to:

"The RS-FEC sublayer shall also be capable of indicating when an errored codeword was not corrected."

Cl 91 SC 91.5.3.3 P126 L9 # 2  
Szczepanek, Andre Inphi

Comment Type TR Comment Status A  
"The RS-FEC sublayer shall also be capable of detecting uncorrectable codewords"  
It is not theoretically possible to detect all possible uncorrectable codewords as some error patterns can change one valid codeword into another valid codeword.  
The text in almost all of the rest of the clause has been altered to be consistent with clause 74 and use the terminology "corrected" and "uncorrected" codewords/blocks. This terminology was adopted for Clause 74 to avoid the issue of what is and isn't a correctable block and focus instead on what the sublayer actually does : correct, or fail to correct a block.

SuggestedRemedy  
Delete sentence "The RS-FEC sublayer shall also be capable of detecting uncorrectable codewords" as it includes a "shall" that isn't achievable or verifiable.

Response Response Status C  
ACCEPT IN PRINCIPLE.

[changed Sublause to 91.5.3.3 for consistent sorting.]

See comment #112.

Cl 91 SC 91.5.3.4 P126 L38 # 190  
Slavick, Jeff Avago Technologies

Comment Type E Comment Status D bucket  
If rx\_lpi\_active is asserted, then the Rx will see RAMs every other codeword.

SuggestedRemedy  
Change "The rx\_lpi\_active is true"  
to "When rx\_lpi\_active is true"

Proposed Response Response Status W  
PROPOSED ACCEPT.

[Changed Subcl to 91.5.3.4 for consistent sorting.]

In addition, change Page 126, Line 36 to:  
"...result in changes in the relative position."

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

CI 91 SC 91.5.3.5 P127 L31 # 73  
Pillai, Velu Broadcom

Comment Type TR Comment Status D bucket

If rx\_xcoded<0> is 0 and all rx\_coded<j+1>=1 is not correct.

SuggestedRemedy

It needs to be

If rx\_xcoded<0> is 0 and all rx\_xcoded<j+1>=1

Proposed Response Response Status W

PROPOSED ACCEPT.

CI 91 SC 91.5.3.5 P127 L34 # 71  
Pillai, Velu Broadcom

Comment Type T Comment Status D bucket

a)Set c = 1 and h<3:0> = 0000.  
The variable c is set to 1; On the transcoding side for the case of invalid sync header, c is set to 0

SuggestedRemedy

For consistency sake C should be set to 0

Proposed Response Response Status W

PROPOSED ACCEPT.

CI 91 SC 91.5.3.5 P127 L6 # 74  
Pillai, Velu Broadcom

Comment Type TR Comment Status D bucket

If rx\_xcoded<0> is 0 and any rx\_coded<j+1>=1 is not correct

SuggestedRemedy

It needs to be

If rx\_xcoded<0> is 0 and any rx\_xcoded<j+1>=0

Proposed Response Response Status W

PROPOSED ACCEPT.

CI 91 SC 91.5.4.2 P130 L36 # 115  
Sela, Oren Mellanox Technologies

Comment Type T Comment Status A

When EEE is supported lanes 16,17,18 and 19 should only be compared when rx\_lpi\_active is true - this is because in the next state the amp\_counter counts lower only when the rx\_lpi\_active is true. It is not broken as EEE capable device when rx\_lpi\_active false and first\_pcs1 is 16,17,18 or 19 then 4096 FEC code word later there should be lane 16, 17, 18 or 19 in the same possision but this was not the intent

SuggestedRemedy

change:  
For the optional EEE capability, each FEC lane also compares the candidate block to the alignment marker payload for PCS lanes 16, 17, 18, and 19  
To:  
For the optional EEE capability, when rx\_lpi\_active is true each FEC lane also compares the candidate block to the alignment marker payload for PCS lanes 16, 17, 18, and 19

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #207.

CI 91 SC 91.5.4.2.1 P130 L16 # 205  
Slavick, Jeff Avago Technologies

Comment Type T Comment Status A

With the inclusion of EEE into clause 82, Figure 82-12 now sets rx\_align\_status rather than align\_status. Other text in Clause 82 states that align\_status = rx\_align\_status when EEE is not supported. However, Clause 91 just references Figure 82-12.

SuggestedRemedy

Change align\_status variable name to be rx\_align\_status  
Change Figure 91-10 to use rx\_align\_status rather than align\_status  
Change tx\_quiet\_timer to refer to rx\_align\_status

Response Response Status C

ACCEPT.

[Changed Subcl to 91.5.4.2.1 for more consistent sorting.]

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

CI 91 SC 91.5.4.2.1 P130 L36 # 207  
 Slavick, Jeff Avago Technologies

Comment Type T Comment Status A

Setting amp\_valid true by comparing alignment markers to PCS lanes 16,17,18,19 is only valid when we're receiving RAMs.

SuggestedRemedy

Change "For the optional EEE capability, each FEC lane also compares the candidate block to the alignment marker payload for PCS lanes 16, 17, 18, and 19."

to:  
 "For the optional EEE capability, each FEC lane also compares the candidate block to the alignment marker payload for PCS lanes 16, 17, 18, and 19 when rx\_lpi\_active is true."

Response Response Status C

ACCEPT.

[Changed Subcl to 91.5.4.2.1 for more consistent sorting.]

CI 91 SC 91.5.4.2.1 P130 L39 # 212  
 Healey, Adam LSI Corporation

Comment Type T Comment Status A

Editor's note states the maximum distance of 3 nibbles may not be suitable for a 100GBASE-KP4 PHY.

However, the following argument has been suggested (by Zhongfeng Wang):

1. Estimates of the net coding gain imply about 0.4 dB additional coding gain for 100GBASE-KP4 FEC.
2. Therefore roughly assume the uncorrected error ratio for 100GBASE-KP4 could be 10x greater than for 100GBASE-KR4.
3. This implies, for the worst-case scenario, the mechanism would fail to lock with 6 RS-FEC codewords on an average of once every 1E7 years rather than 1E9 years for 100GBASE-KR4.

If this is the case, the likelihood of failure is very small and thus there is no compelling reason to modify the synchronization mechanism for 100GBASE-KP4.

SuggestedRemedy

Remove the editor's note.

Response Response Status C

ACCEPT.

CI 91 SC 91.5.4.2.1 P131 L50 # 206  
 Slavick, Jeff Avago Technologies

Comment Type T Comment Status A

ram\_valid and ramps\_valid are testing for valid Rapid Alignment Markers.

SuggestedRemedy

Change "valid alignment markers" to "valid Rapid Alignment Markers" for both ram\_valid and ramps\_valid variables.

Response Response Status C

ACCEPT IN PRINCIPLE.

[Changed Subcl to 91.5.4.2.1 for more consistent sorting.]

Strictly speaking, ramps\_valid tests for valid Rapid Alignment Marker payloads as the header bits are discarded in the mapping process.

Change the end of the definition of ram\_valid to:  
 "...are valid Rapid Alignment Markers and is set to false otherwise."

See #210 for the definition of ramps\_valid.

CI 91 SC 91.5.4.2.1 P131 L51 # 209  
 Healey, Adam LSI Corporation

Comment Type T Comment Status A

The bit error ratio of a CAUI that separates the PCS from the RS-FEC sublayer is expected to be low (less than 1E-12). Furthermore, it is unlikely (on the order of 1/2^50) to detect a valid alignment marker in random data.

Therefore, it is not necessary to check all PCS lanes for rapid alignment markers. The actual number to be checked is TBD.

SuggestedRemedy

For ram\_valid, set TBD to 2.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change the definition of ram\_valid to:  
 "Boolean variable that is set to true when the 66-bit blocks concurrently received on at least 2 PCS lanes are valid Rapid Alignment Markers with identical values in the Count Down fields and is set to false otherwise."



IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

CI 91 SC 91.5.4.2.1 P131 L8 # 70  
 Pillai, Velu Broadcom

Comment Type T Comment Status A  
 fec\_alignment\_valid variable description needs to indicate that each FEC lane needs to lock to a unique AM. This unique requirement is in the alignment\_valid variable description in CL82.2.18.2.2

SuggestedRemedy

Response Response Status C  
 ACCEPT IN PRINCIPLE.

Note that the lane mapping assignment is added by comment #183.

Change the definition of fec\_alignment\_valid to:  
 "Boolean variable that is set to true if all FEC lanes are aligned. FEC lanes are considered to be aligned when amps\_lock<x> is true for all x, each FEC lane is locked to a unique alignment marker payload sequence (see 91.5.2.6), and the FEC lanes are deskewed. Otherwise, this variable is set to false."

CI 91 SC 91.5.4.2.1 P132 L2 # 210  
 Healey, Adam LSI Corporation

Comment Type T Comment Status A  
 The variable ramps\_valid checks for "rapid" alignment marker payload sequences on the FEC lanes.

Since FEC codeword boundaries are known during this search, the corrected message could be used as the subject of the search (unless correction is bypassed).

If correction is not bypassed, it is unlikely that the RAM payload patterns would appear in random data. Therefore, it should be sufficient to check that a 64-bit block marker payload on any 2 FEC lanes corresponds to the first rapid alignment marker payload corresponding to that lane.

If the mechanism is intended to be operated with correction bypassed, a more complicated analysis of the appropriate distance between the reference pattern and the observed pattern must be performed.

SuggestedRemedy

Update the definition of ramps\_valid accordingly.

Response Response Status C  
 ACCEPT IN PRINCIPLE.

If correction is bypassed, it seems likely that the error probability is sufficiently low that an error in the Rapid Alignment Marker payload sequence would be very unlikely. If correction is not bypassed, the corrected Rapid Alignment Marker payload sequences are available to be examined with a low likelihood of error.

Given these assumptions, change the definition of ramps\_valid to:  
 "Boolean variable that is set to true if the received 64-bit blocks concurrently received on at least 2 FEC lanes are valid Rapid Alignment Marker payloads with identical values in the Count Down fields and is set to false otherwise."

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

Cl 91 SC 91.5.4.2.1 P133 L17 # 208  
 Slavick, Jeff Avago Technologies

Comment Type T Comment Status A  
 TBDs are in place for the quiet timers for Clause 91.

SuggestedRemedy  
 see slavick\_3bj\_01\_1112.pdf

Response Response Status C  
 ACCEPT IN PRINCIPLE.

[Changed Subcl to 91.5.4.2.1 for more consistent sorting.]

Specify the value of tx\_tq\_timer to be between 1.8 and 2 ms.  
 Specify the value of rx\_tq\_timer to be between 2 and 2.8 ms.

Cl 91 SC 91.5.4.2.3 P133 L17 # 211  
 Healey, Adam LSI Corporation

Comment Type T Comment Status A  
 The counters rx\_quiet\_timer and tx\_quiet\_timer are both TBD. Both timers should exceed the maximum value of the rx\_quiet\_timer at the PCS (currently set to 3 ms).

SuggestedRemedy  
 Set the range of both timers to 3.1 to 3.4 ms.

Response Response Status C  
 ACCEPT IN PRINCIPLE.

See comment #208.

Cl 91 SC 91.5.4.3 P136 L34 # 114  
 Sela, Oren Mellanox Technologies

Comment Type T Comment Status A  
 When only FW EEE is supported the arch from TX\_TEST\_NEXT to TX\_QUIET should not be taken

SuggestedRemedy  
 Add paramter called LPI\_FW - true in FW mode false in normal wake mode in Figure 91-10 - on the arch from TX\_TEST\_NEXT to TX\_QUIET add LPI\_FW\*(false!align\_status + !ram\_valid). And add an arch !LPI\_FW\*(false!align\_status + !ram\_valid) from TX\_TEST\_NEXT to TX\_FAULT

Response Response Status C  
 ACCEPT IN PRINCIPLE.

[Changed Subcl from 91-10 to 91.5.4.3 for consistent sorting. Added Line 34.]

It is true that a loss of alignment in the "fast wake" mode should be considered a fault and not a transition to a quiet line state.

Define new variable "fec\_lpi\_fw" as follows:  
 "Boolean variable that controls the behavior of the Transmit LPI and Receive LPI state diagrams. This variable is set to true when the local PCS is configured to use the Fast Wake mechanism and set to false otherwise."

Change the transition condition from TX\_TEST\_NEXT to TX\_QUIET to:  
 !fec\_lpi\_fw \* !rx\_align\_status

Add a transition from TX\_TEST\_NEXT to TX\_FAULT with the condition:  
 fec\_lpi\_fw \* !rx\_align\_status

Change the transition condition from RX\_TEST\_NEXT to RX\_QUIET to:  
 !fec\_lpi\_fw \* !fec\_align\_status

Add a transition from RX\_TEST\_NEXT to RX\_FAULT with the condition:  
 fec\_lpi\_fw \* !fec\_align\_status

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

Cl 91 SC 91.5.4.3 P136 L35 # 204  
 Slavick, Jeff Avago Technologies

Comment Type T Comment Status A

The last RAM down\_count value transmitted is 1 not 0. So figures 91-10 and 91-11 need to reflect that.

SuggestedRemedy

Change the test values on the exit of TX\_TEST\_NEXT and RX\_TEST\_NEXT to compare \*\_down\_count against 1.

Response Response Status C

ACCEPT IN PRINCIPLE.

[Changed Subcl to 91.5.4.3 for more consistent sorting.]

Define the following variables:

ram\_valid\_prev  
 Boolean variable that holds the value of ram\_valid from the previous expected Rapid Alignment Marker position.

ramps\_valid\_prev  
 Boolean variable that holds that value of ramps\_valid from the previous expected Rapid Alignment Marker payload position.

Add the following assignments:  
 In TX\_LPI, assign "ram\_valid\_prev <= ram\_valid"  
 In RX\_LPI, assign "ramps\_valid\_prev <= ramps\_valid"

Change the state transition conditions in Figure 91-10 and 91-11 as follows.

From TX\_TEST\_NEXT to TX\_LPI:  
 rx\_align\_status \* ((!ram\_valid \* ram\_valid\_prev) + (ram\_valid \* tx\_down\_count != 1))

From TX\_TEST\_NEXT to TX\_ACTIVE:  
 rx\_align\_status \* ((!ram\_valid \* !ram\_valid\_prev) + (ram\_valid \* tx\_down\_count=1))

From TX\_QUIET to TX\_FAULT:  
 tx\_quiet\_timer\_done

From RX\_TEST\_NEXT to RX\_LPI:  
 fec\_align\_status \* ((!ramps\_valid \* ramps\_valid\_prev) + (ramps\_valid \* rx\_down\_count != 1))

From RX\_TEST\_NEXT to RX\_ACTIVE:  
 fec\_align\_status \* ((!ramps\_valid \* !ramps\_valid\_prev) + (ramps\_valid \* rx\_down\_count=1))

From RX\_QUIET to RX\_FAULT:

rx\_quiet\_timer\_done

Cl 91 SC 91.6 P138 L26 # 183  
 Gustlin, Mark Xilinx

Comment Type T Comment Status A

Since a given FEC lane can be received on any of the four service interface lanes, add a register that captures which FEC lane is recieved at a given time on each service interface lane.

This is analogous to Lane x mapping register that is part of Clause 82 (Table 82-7).

SuggestedRemedy

Per the comment.

Response Response Status C

ACCEPT IN PRINCIPLE.

When the RS-FEC sublayer is connected to the PCS via CAUI, the PCS lane mapping for the RS-FEC transmit function would also be of interest.

Add PCS "Lane x mapping" registers similar to Clause 82, Table 82-7 to Table 91-3. The variables lane\_mapping<x> are assigned by Alignment marker lock state diagram (Figure 82-11) which is incorporated into Clause 91 by reference.

Add FEC "Lane x mapping" registers to Table 91-3. Add "fec\_lane\_mapping<x> <= fec\_lane" assignment to the "2\_GOOD" state of the FEC synchronization state diagram Figure 91-8. Define fec\_lane to be an fec\_lane number (0 to 3) that is derived from the values of first\_pcs1 and/or current\_pcs1 per the mapping defined in 91.5.2.6.

Add corresponding register space to Clause 45.

Cl 91 SC 91.6.2 P138 L35 # 380  
 Cideciyan, Roy IBM

Comment Type TR Comment Status D

MTTFPA computations in cideciyan\_01\_0512.pdf always assume that RS decoder reports (indicates) errors to PCS layer whenever there is an uncorrectable code word (error correction mode) or code word contains errors (error detection mode). Therefore, indication of errors to the PCS sublayer is not an option but a mandatory feature of the RS decoder in order to have satisfactory MTTFPA.

SuggestedRemedy

Omit subclause 91.6.2 as this variable is not needed.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

See comment #369.

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

CI 91 SC 91.6.3 P138 L47 # 191  
 Slavick, Jeff Avago Technologies

Comment Type E Comment Status D bucket

The FEC\_\*\_ability registers reference the wrong MDIO registers

SuggestedRemedy

Change FEC\_bypass\_correction\_ability to refer to 1.201.1  
 Change FEC\_error\_indication\_ability to refer to 1.201.2

Proposed Response Response Status W

PROPOSED ACCEPT.

[Changed Subcl to 91.6.3 for more consistent sorting.]

Note changes to Table 91-3 and 91.6.4 in addition to 91.6.3.

FEC\_error\_indication\_ability may be removed per comment #TBD which would overtake that portion of this response.

CI 91 SC 91.6.4 P138 L48 # 381  
 Cideciyan, Roy IBM

Comment Type TR Comment Status D

MTTFPA computations in cideciyan\_01\_0512.pdf always assume that RS decoder reports (indicates) errors to PCS layer whenever there is an uncorrectable code word (error correction mode) or code word contains errors (error detection mode). Therefore, indication of errors to the PCS sublayer is not an option but a mandatory feature of the RS decoder in order to have satisfactory MTTFPA.

SuggestedRemedy

Omit subclause 91.6.4 as this variable is not needed.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

See comment #369.

CI 91 SC 91.7.3 P141 L5 # 26  
 D'Ambrosia, John Dell

Comment Type TR Comment Status A

Item KR4 and KP4 have no corresponding shall statements. Also, both values are set to -KR4, which doesn't make sense.

SuggestedRemedy

delete  
 the determination of the KR4 and KP4 PHY is not done in the FEC sublayer

Response Response Status C

ACCEPT IN PRINCIPLE.

The RS-FEC sublayer implements a different Reed-Solomon code depending on whether it is used to form a complete 100GBASE-KR4 PHY or a complete 100GBASE-KP4 PHY. These options are defined in order to specify that conditional requirement (see TF9, TF10, RF3, and RF4).

Change Value/Comment for \*KP4 to be "Used to form a complete 100GBASE-KP4 PHY".

CI 91 SC 91.7.4.1 P142 L31 # 9  
 D'Ambrosia, John Dell

Comment Type E Comment Status A

TF9 is for 100GBASE-KR4 and 100GBASE-CR4

SuggestedRemedy

Add 100GBASE-CR4

Response Response Status C

ACCEPT IN PRINCIPLE.

In 91.7.3, change item \*KR4 as follows.  
 Feature: "100GBASE-CR4 or 100GBASE-KR4"  
 Value/Comment: "Used to form complete 100GBASE-CR4 or 100GBASE-KR4 PHY"

Change TF9 Feature to "Reed-Solomon encoder for 100GBASE-CR4 or 100GBASE-KR4"

Change RF3 Feature to "Reed-Solomon decoder for 100GBASE-CR4 or 100GBASE-KR4"

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

CI 91 SC 91.7.4.2 P143 L18 # 4

Szczepanek, Andre

Inphi

Comment Type TR Comment Status A

See previous comments related to the use of "uncorrectable" on page 126

SuggestedRemedy

Delete Item RF5

Response Response Status C

ACCEPT IN PRINCIPLE.

[Changed Clause from 19 to 91, changed Subclause to 91.7.4.2 for consistent sorting.]

Change RF5 Value/Comment to:  
"Capable of indicating when a codeword was not corrected."

CI 91 SC 91.7.4.2 P143 L21 # 5

Szczepanek, Andre

Inphi

Comment Type TR Comment Status A

See previous comments related to the use of "uncorrectable" on page 126

SuggestedRemedy

Replace "for uncorrectable codewords"  
with  
"for uncorrected errored codewords"

Response Response Status C

ACCEPT.

[Changed Clause from 19 to 91, changed Subclause to 91.7.4.2 for consistent sorting.]

Change RF6 Value/Comment to:  
"When enabled, corrupts 66-bit block synchronization headers for uncorrected errored codewords (or errored codewords when correction is bypassed)"

CI 91 SC 91.7.4.2 P143 L26 # 10

D'Ambrosia, John

Dell

Comment Type E Comment Status D bucket

subclause reference for RF7 wrong

SuggestedRemedy

change to 91.5.3.4

Proposed Response Response Status W

PROPOSED ACCEPT.

CI 91 SC 91.7.4.3 P143 L53 # 11

D'Ambrosia, John

Dell

Comment Type E Comment Status D bucket

Feature name for SD5 is incorrect

SuggestedRemedy

change to Rx LPI process

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Change to "Receive LPI process".

CI 91A SC 91A.1 P276 L1 # 66

Pillai, Velu

Broadcom

Comment Type E Comment Status R

The example RS-FEC blocks contains only Idle control characters. It will be better if we can have a block that has a mix of data and control codewords that addresses the different combinations. Basically a set that exercises the complex equations in subclause 91.5.2.5 and 91.5.3.5

SuggestedRemedy

Response Response Status C

REJECT.

This example is sufficient for the user to verify the correct bit order and implementation of the Reed-Solomon encoder.

Figure 91-3 was provided to illustrate the construction of 257-bit blocks for different mixtures of control and data words.

IEEE P802.3bj D1.2 100 Gb/s Backplane and Copper Cable 3rd Task Force review comments

CI 91A SC 91A.2 P277 L1 # 65  
Pillai, Velu Broadcom

Comment Type E Comment Status R

The CL91 text already clarifies in section 91.5.2.7 that when the transcoded data [0:256] is partitioned into 10-bit message symbols from left to right in the encoder, the resulting values are  $\{m_{<k-1>}[0:9], m_{<k-2>}[0:9], \dots, m_{<0>}[0:9]\}$ . An additional statement to section 91A.2 to indicate that when these values are used for parity symbol generation, the values must first be flipped end-to-end to become  $\{m_{<k-1>}[9:0], m_{<k-2>}[9:0], \dots, m_{<0>}[9:0]\}$  before being applied to the parity generation algorithm.

SuggestedRemedy

Response Response Status C

REJECT.

The annex clearly states the bit order for the contents of the tables and refers the reader to 91.5.2.7 which defines the how the bits are to be organized and ordered for processing by the Reed-Solomon encoder.

Correct implementation of the rules of 91.5.2.7 would yield the codewords included in Annex 91A.

No additional statements appear to be necessary.