Pre-FEC BER Monitoring: Comments against P802.3bs D1.2

David Ofelt

2016-03

Background

- Previous presentations:
  - maki_3bs_01a_1115.pdf
  - ofelt_3b_01a_0116.pdf
- Proposed to add pre-FEC BER tracking to:
  - Proactively fail the link in the face of a configurable high BER condition
  - Signal a configurable degraded pre-FEC BER situation to the local node and, using in-band communication, to the transmitter.
- Straw poll held for the degrade feature in January 2016 meeting in Atlanta:
  - 24 Yes, 0 No, 14 Abstain
- Fault feature overlapped with the FEC indication bypass – which was accepted with a separate straw poll.

Basic Approach

- Pre-FEC Fault Feature
  - Constructed by extending the FEC bypass indication feature
  - Separates out High SER function from the FEC bypass indication function
  - Makes the interval and threshold values configurable from the current fixed values of 8120 and 5560.
- Pre-FEC Degrade Feature
  - Clones the High SER function to make Degrade SER function
  - Separates the assert and de-assert interval and threshold values to provide hysteresis
  - Local and Remote Degrade bits placed into and extracted from the Alignment Markers

Changes from the initial proposal

- Original proposal in ofelt_3bs_01a_0116
  - http://www.ieee802.org/3/bs/public/16_01/ofelt_3bs_01a_0116.pdf
- This version:
  - Changes “BER” terminology to “SER” to be consistent with the current draft
  - Adds separate activate and deactivate controls for the Local Degrade feature
  - Adds voting to set the received LD/RD bits from the alignment marker to avoid glitches in the face of uncorrectable errors when FEC bypass indication mode is enabled.

Notes

- This set of changes does not try to assign exact bit positions, register numbers, or sub clause numbers.
- These changes try to make cross references make sense, but the exact numbering will depend on the final edits and will be up to the cause editors

Clause 119 – 1

- 119.2.4.4 : p97 : l39 : Alignment Marker Insertion
- New text - bit positions to be picked once the new AM structure is settled but the last two bits of the AM pad in the last arriving 257-bit alignment marker block were considered plausible choices on the logic calls. This text assumes this location to make things concrete.

  - *The last two bits of the final 257-bit alignment marker block provide the Local Degrade (tx_am_ld) and Remote Degrade (tx_am_rd) bits. If the FEC_degraded_SER_enable bit (see 119.3) is zero then both tx_am_ld and tx_am_rd are set to zero. If the FEC_degraded_SER_enable bit is one, then tx_am_ld is set to one if this PCS is in an extender sublayer and either the FEC_degraded_SER (see 119.3) bit is set or the incoming am_ld is set (see 119.2.5.4). If this PCS is not an extender sublayer then tx_am_ld is set to zero. The tx_am_rd bit is set if either the incoming am_ld bit is set or the local FEC_degraded_SER is set.*

Clause 119 – 2

- 119.2.5.3 : p105 : l4 : Reed-Solomon Decoder
  - Move the 2nd to last paragraph in subsection 119.2.5.3 (“The Reed-Solomon decoder may … variable (see 119.3) should be moved to the end of the subsection (after the current last paragraph) and modified to read:
  - …Assertion of the FEC_bypass_indication_enable variable (see 119.3). *In order to have safe operation in the face of a large SER, the high SER function should be enabled to protect against false packet acceptance. The appropriate minimum SER is set by having a FEC_high_SER_interval set to 8192 codewords and a FEC_high_SER_threshold set to 5560.*
- 119.2.5.3. : p105 : l9
  - Modify the current last paragraph (“When FEC_bypass_indication … 75ms) to:
  - *When FEC_high_SER_enable is asserted, additional error monitoring is performed by the RS_FEC sublayer to bring down the link in the face of too high a pre-FEC symbol error rate. The Reed-Solomon decoder counts the number of symbol errors detected on all PCS lanes in consecutive non-overlapping blocks of FEC_high_SER_interval (see 119.3) codewords. When the number of symbol errors in a block of FEC_high_SER_interval codewords exceeds the threshold set in FEC_high_SER_threshold (see 119.3), the Reed-Solomon decoder shall cause the synchronization header rx_coded<1:0> of each subsequent 66-bit block that is delivered to the PCS decoder to be assigned a value of 11 for a period of 60 ms to 75 ms.*

Clause 119 – 3

- 119.2.5.3 : p105 : l15 : Reed-Solomon Decoder
- New Paragraph
  - *When FEC_degraded_SER_enable (see 119.3) is asserted, additional error monitoring is performed by the RS_FEC sublayer to indicate that the interface may be experiencing a significant error rate. The Reed-Solomon decoder counts the number of symbol errors detected on all PCS lanes in consecutive non-overlapping blocks of FEC_degraded_SER_assert_interval (see 119.3) codewords and in consecutive non-overlapping blocks of FEC_degraded_SER_deassert_interval (see 119.3) codewords. When the number of symbol errors in a block of FEC_degraded_SER_assert_interval codewords exceeds the threshold set in FEC_degreaded_SER_assert_threshold (see 119.3), the Reed-Solomon decoder will assert the FEC_degraded_SER bit (see 119.3). If the FEC_degraded_SER bit is set and there are fewer than FEC_degraded_SER_deassert_threshold (see 119.3) symbol errors in a block of FEC_degraded_SER_deassert_interval (see 119.3) codewords, then the FEC_degraded_SER bit is cleared.*

Clause 119 – 4

- 119.2.5.5 : p105 : l28 : Alignment Marker Removal
- Add a new paragraph after line 28.
  - *If the FEC_degraded_SER_enable bit is set, then the last 2 bits of the final 257-bit block in the alignment marker section provide the values for the received local degrade (rx_am_ld) and remote degrade (rx_am_rd) bits. If the FEC_degraded_SER_enable bit is zero, then rx_am_ld and rx_am_rd are both set to zero. If two consecutive alignment markers have the same values for each of rx_am_ld and rx_am_rd then am_ld (see 119.3) and am_rd (see 119.3) are set to the respective values.*

Clause 119 – 4

- 119.3 : p115 : l10 : Table 119-3 MDIO/PCS control variable mappings

MDIO Control Variable	PCS Register Name	Register / bit number	PCS control variable
High SER Enable	PCS FEC control register	3.800.x	FEC_high_SER_enable
Degraded SER Enable	PCS FEC control register	3.800.y	FEC_degraded_SER_enable
PCS FEC High SER Threshold	PCS FEC High SER Threshold register	3.80a, 3.80b	FEC_high_SER_threshold
PCS FEC High SER Interval	PCS FEC High SER Interval register	3.80c, 3.80d	FEC_high_SER_interval
PCS FEC Degraded SER Activate Threshold	PCS FEC Degraded SER Activate Threshold register	3.80e, 3.80f	FEC_degraded_SER_activate_threshold
PCS FEC Degraded SER Activate Interval	PCS FEC Degraded SER Activate Interval	3.80g, 3.80h	FEC_degraded_SER_activate_interval
PCS FEC Degraded SER Dectivate Threshold	PCS FEC Degraded SER Dectivate Threshold register	3.80e, 3.80f	FEC_degraded_SER_deactivate_threshold
PCS FEC Degraded SER Dectivate Interval	PCS FEC Degraded SER Dectivate Interval	3.80g, 3.80h	FEC_degraded_SER_deactivate_interval

Clause 119 – 5

- 119.3 : p115 : l29 : Table 119-4 MDIO/PCS status variable mapping

MDIO Control Variable	PCS Register Name	Register / bit number	PCS control variable
FEC High SER Ability	PCS FEC status register	3.801.a	FEC_high_SER_ability
FEC Degraded SER Ability	PCS FEC status register	3.801.b	FEC_degraded_SER_ability
FEC Degraded SER	PCS FEC status register	3.800.c	FEC_degraded_SER
Alignment Marker Local Degrade	PCS FEC status register	3.800.d	am_ld
Alignment Marker Remote Degrade	PCS FEC status register	3.800.e	am_rd

Clause 45 – PCS FEC Control bits

- 45.2.3.47c : p61 : l31 – Table 45-160b “PCS FEC control register bit definitions”
  - Add 2 new bits

Bit(s)	Name	Description	R/W
3.800.x	High SER Enable	1 = High SER detection enabled 0 = High SER detection disabled	RW
3.800.y	Degraded SER Enable	1 = Degraded SER detection enabled 0 = Degraded SER detection disabled	RW

Clause 45 – PCS FEC Control bits 2

- 45.2.3.47c.x : p61 : l47
  - Add description for “**PCS High SER Enable**”
  - This bit enables the PCS FEC high SER feature. When set to a one, the FEC High SER bit (3.801.2) will be set when the number of FEC symbol errors crosses a threshold as described in 45.2.3.47d.1. When set to a zero, the FEC High SER bit will not be set. Writes to this bit are ignored and reads return a zero if the PCS FEC does not have the ability to detect high SER (bit 3.801.a)
- 45.2.3.47c.y : p61 : l47
  - Add description for “**PCS Degraded SER Enable**”
  - This bit enables the PCS FEC degraded SER feature. When set to a one, the FEC Degraded SER bit (3.801.c) will be set when the number of FEC symbol errors crosses a threshold as described in 45.2.3.47d.X. When set to a zero, the FEC Degraded SER bit will not be set. Writes to this bit are ignored and reads return a zero if the PCS FEC does not have the ability to detect degraded SER (bit 3.801.b)

Clause 45 – PCS FEC Status bits 1

- 45.2.3.47d : p62 : l3 – Table 45-160c “PCS FEC status register bit definitions”
  - Add 5 new bits

Bit(s)	Name	Description	R/W
3.801.a	PCS FEC High SER Ability	1 = FEC Decoder has the ability to do high SER detection 0 = FEC Decoder does not have the ability to do high SER detection	RO
3.801.b	PCS FEC Degraded SER Ability	1 = FEC Decoder has the ability to do degraded SER detection 0 = FEC Decoder does not have the ability to do degraded SER detection	RO
3.801.c	PCS FEC Degraded SER	1 = FEC errors have exceeded the degrade threshold 0 = FEC errors have not exceeded the degrade threshold	RO
3.801.d	Alignment Marker Local Degrade Set	1 = Local Degrade bit set in the received alignment marker 0 = Local Degrade bit not set in the received alignment marker	RO
3.801.e	Alignment Marker Remote Degrade Set	1 = Remote Degrade bit set in the received alignment marker 0 = Remote Degrade bit not set in the received alignment marker	RO

Clause 45 – PCS FEC Status bits 2

- 45.2.3.47d.1: p62 : l16 : **PCS FEC high SER** (3.801.2)
  - Change text to:
  - *When the PCS FEC High SER enable (bit 3.800.x) is set to one, this bit is set to one if the number of FEC symbol errors in a window of PCS FEC High SER Interval (registers 3.80c and 3.80d) is above the PCS FEC High SER Threshold (registers 3.80a and 3.80b) (see 119.2.5.3) and is set to zero otherwise. This bit is set to zero if the PCS FEC High SER enable (bit 3.800.x) is set to zero. This bit shall be set with latching high behavior.*
- 45.2.3.47d.a: p62 : l28
  - Add description for "**PCS FEC High SER Ability**"
  - *The PCS FEC decoder may have the option to detect a high SER condition (see 119.2.5.3). This bit is set to one to indicate that the decoder has the ability to detect high SER conditions. The bit is set to zero if the ability is not supported.*

Clause 45 – PCS FEC Status bits 3

- 45.2.3.47d.b: p62 : l28
  - Add description for “**PCS FEC Degraded SER Ability**”
  - *The PCS FEC decoder may have the option to detect a degraded SER condition (see 119.2.5.3). This bit is set to one to indicate that the decoder has the ability to detect degraded SER conditions. The bit is set to zero if the ability is not supported.*
- 45.2.3.47d.c: p62 : l28
  - Add description for “**PCS FEC Degraded SER**”
  - *When the PCS FEC Degraded SER enable (bit 3.800.y) is set to one, this bit is set to one if the number of FEC symbol errors in a window of PCS FEC Degraded SER Activate Interval (registers 3.80g and 3.80h) is above the PCS FEC Degraded SER Activate Threshold (registers 3.80e and 3.80f) and is cleared if it was set and there were fewer than PCS FEC Degraded Deactivate Threshold (registers 3.80i and 3.80j) FEC symbol errors at the end of the PCS FEC Degraded Deactivate Interval (registers 3.80k and 3.80l) (see 119.2.5.3). This bit is set to zero if the PCS FEC Degraded SER enable (bit 3.800.y) is set to zero. This bit shall be set and cleared with latching behavior.*

Clause 45 – PCS FEC Status bits 3

- 45.2.3.47d.d: p62 : l28
  - Add description for “**Alignment Marker Local Degrade Set**”
  - *The Alignment Marker Local Degrade Set bit is set to one if two successive alignment markers have the Local Degrade bit set. The bit is set to zero if two successive alignment markers have the Local Degrade bit clear (see 119.2.5.5).*
- 45.2.3.47d.e: p62 : l28
  - Add description for “**Alignment Marker Remote Degrade Set**”
  - *The Alignment Marker Remote Degrade Set bit is set to one if two successive alignment markers have the Remote Degrade bit set. The bit is set to zero if two successive alignment markers have the Remote Degrade bit clear (see 119.2.5.5).*

Clause 45 – New PCS Registers

- After 45.2.3.47f : p63 : l14
- Add the following new registers – details for each on the following pages
  - PCS FEC High SER Threshold
  - PCS FEC High SER Interval
  - PCS FEC Degraded SER Activate Threshold
  - PCS FEC Degraded SER Activate Interval
  - PCS FEC Degraded SER Deactivate Threshold
  - PCS FEC Degraded SER Deactivate Interval

Clause 45 – New PCS Registers 2

- **45.2.3.47x PCS FEC High SER Threshold**
  - *The assignment of bits in the PCS FEC High SER Threshold register is shown in table 45-160x. The value controls the threshold used to set the PCS FEC high SER status bit (3.801.2). This value should be set to 5560 if just the PCS FEC bypass indication feature is being used.*
  - Table 45-160x
    - 3.80a.15:0 | High SER Threshold lower | FEC_high_ser_threshold[15:0] | R/W
    - 3.80b.15:0 | High SER Threshold upper | FEC_high_ser_threshold[31:16] | RW
- **45.2.3.47x PCS FEC High SER Interval**
  - *The assignment of bits in the PCS FEC High SER Interval register is shown in table 45-160y. The value controls the interval used to set the PCS FEC high SER status bit (3.801.2). This value should be set to 8192 if just the PCS FEC bypass indication feature is being used.*
  - Table 45-160y
    - 3.80c.15:0 | High SER Interval lower | FEC_high_ser_interval[15:0] | R/W
    - 3.80d.15:0 | High SER Interval upper | FEC_high_ser_interval[31:16] | RW

Clause 45 – New PCS Registers 3

- **45.2.3.47x PCS FEC Degraded SER Activate Threshold**
  - *The assignment of bits in the PCS FEC Degraded SER Activate Threshold register is shown in table 45-160x. The value controls the threshold used to set the PCS FEC degraded SER status bit (3.801.x).*
  - Table 45-160x
    - 3.80e.15:0 | Degraded SER Activate Threshold lower | FEC_degraded_ser_activate_threshold[15:0] | R/W
    - 3.80f.15:0 | Degraded SER Activate Threshold lower | FEC_degraded_ser_activate_threshold[31:16] | R/W
- **45.2.3.47x PCS FEC Degraded SER Activate Interval**
  - *The assignment of bits in the PCS FEC Degraded SER Activate Interval register is shown in table 45-160y. The value controls the interval used to set the PCS FEC degraded SER status bit (3.801.x).*
  - Table 45-160y
    - 3.80g.15:0 | Degraded SER Activate Interval lower | FEC_degraded_ser_activate_interval[15:0] | R/W
    - 3.80h.15:0 | Degraded SER Activate Interval upper | FEC_degraded_ser_activate_interval[31:16] | RW

Clause 45 – New PCS Registers 4

- **45.2.3.47x PCS FEC Degraded SER Deactivate Threshold**
  - *The assignment of bits in the PCS FEC Degraded SER Deactivate Threshold register is shown in table 45-160x. The value controls the threshold used to clear the PCS FEC degraded SER status bit (3.801.x).*
  - Table 45-160x
    - 3.80i.15:0 | Degraded SER Deactivate Threshold lower | FEC_degraded_ser_deactivate_threshold[15:0] | R/W
    - 3.80j.15:0 | Degraded SER Deactivate Threshold lower | FEC_degraded_ser_deactivate_threshold[31:16] | R/W
- **45.2.3.47x PCS FEC Degraded SER Deactivate Interval**
  - *The assignment of bits in the PCS FEC Degraded SER Deactivate Interval register is shown in table 45-160y. The value controls the interval used to clear the PCS FEC degraded SER status bit (3.801.x).*
  - Table 45-160y
    - 3.80k.15:0 | Degraded SER Dectivate Interval lower | FEC_degraded_ser_deactivate_interval[15:0] | R/W
    - 3.80l.15:0 | Degraded SER Dectivate Interval upper | FEC_degraded_ser_deactivate_interval[31:16] | RW