PHY Frame Corrections

Ragnar Jonsson – Marvell
Mike Tu – Broadcom
George Zimmerman – CME consulting/Marvell

08/22/2022

Supporters

- Brett McClellan

Introduction

- Several issues were found in 802.3cy D2.0, related to frame alignment
- In particular
  - The definition and use of partial PHY frame was inconsistent
  - The concept of PHY frame was used, but never defined

What we actually intended

Align start of	PFC24	Offset	Comment
Partial (PHY) Frame	1	0	One PFC24 count every 1170 PMA_UNITDATA clocks
RS-FEC Frame	4	0	There are 4 PFC24 in each frame
Super-Frame L=1	4	0	One RS-FEC frame
Super-Frame L=2	8	0	Two RS-FEC frames
Super-Frame L=4	16	0	Four RS-FEC frames
Super-Frame L=8	32	0	Eight RS-FEC frames
OAM field	4	8450/2340	OAM field comes after first 8450 bits in RS-FEC frame
Training Frame	16	0	We don’t know L, so it cannot depend on L
Training allignment bit	1	1	Allignment bit at start of every partial frame
Training info field	16	15	Info Field is in last partial frame of Training Frame
QR Cycle	384	0	Each QR Cycle is 96 RS-Frames
Alert (Fast Wake)	32	16	Fourth frame of every 8 frame period
Alert (Slow Wake)	384	368	The 92nd frame of QR cycle
Refresh (Slave)	384	168	The 42nd frame of QR cycle
Refresh (Master)	384	360	The 90th frame of QR cycle
NOTE: Sleep and Wake signals have to be aligned on super-frame boundary			

Proposed Solution

- Instead of defining timing alignment in terms of poorly defined frames, use the PFC24 counter or frame counters to describe the signal alignment
- This can be done, either by referencing the table on previous slide, or by updating the text to be consistent with the table
- Remove all reference to “PHY Frame” and rename “partial PHY frame counter” to “partial frame counter”

Changes to page 69

When the receive channel is in training mode, the PCS Synchronization process continuously monitors PMA_RXSTATUS.indication(loc_rcvr_status). When loc_rcvr_status indicates OK, then the PCS Synchronization process accepts data-units via the PMA_UNITDATA.indication primitive. It attains frame and block synchronization based on the PMA training frames and conveys received blocks to the PCS Receive process. The PCS Synchronization process sets the block_lock flag to indicate whether the PCS has obtained synchronization. The PMA training frame includes an alignment bit every 1170 PAM2 symbols, which is aligned with the PCS partial ~~PHY~~ frame boundary, as well as an Infofield, which is inserted in the 16th PCS partial ~~PHY~~ frame. When the PCS Synchronization process is synchronized to this pattern, block_lock is asserted.

Changes on page 71

165.3.5 PMA training frame

During PMA training, the training frames are embedded with indicators to establish alignment to the RS-FEC superframe ~~comprised of 16 partial PHY frames that comprise the block~~. The last partial ~~PHY~~ frame, in each training frame, is embedded with an information field used to exchange messages between link partners. The timing relationship among training frame, partial frame, RS-FEC frame, superframe, and partial ~~PHY~~ frame count (PFC24) are shown in Figure 165–12.

[No change to Figure 165-12]

Figure 165–12—Timing relationship to PFC24

PMA training frame encoding is based on the generation, at time *n*, of the bit *Sn*. The first bit is inverted in the first 15 partial ~~PHY~~ frames of each training frame. ~~RS-FEC block~~. The first 96 bits of the 16th partial ~~PHY~~ frame are XORed with the contents of the Infofield. Each partial ~~PHY~~ frame is 1170 PMA_UNITDATA clocks ~~450 bits long, beginning at *Sn* where (*n* mod 450) = 0.~~ See Equation (165–7).

Changes on page 74

165.3.6.1 LPI synchronization

EEE-capable PHYs shall synchronize refresh intervals during the LPI mode. A PHY in SLAVE mode is responsible for synchronizing its partial ~~PHY~~ frame count (PFC24) to the MASTER's PFC24 during PAM2 training. For the requirements on the SLAVE and the MASTER frame alignment, see 165.4.2.4.10.

Refresh signaling is derived by tracking the RS-FEC frame count as shown in Figure 165–13, where:

RS-FEC frame count = floor (PFC24 / 4) mod 96.

Following the transition to PAM4, the PCS continues with the RS-FEC frame count and uses the count to generate refresh, alert, and wake control signals for the transmit functions.

Alert, a four RS-FEC frame long sequence (alert_length), shall start four frames after the beginning of any eighth ~~PHY~~ RS-FEC frame ~~boundary~~ counting from the start of the QR cycle. The MASTER and SLAVE shall derive the tx_refresh_active and tx_alert_start_next signals from the transmitted ~~PHY~~ RS-FEC frames as shown in Table 165–4 and Table 165–5. When Slow Wake is active, alert can be transmitted in only a single QR cycle location, starting ~~a~~at RS-FEC frame 92.

Changes on page 88

165.3.9.1 Definitions

The definitions for OAM are as defined in 149.3.9.1 for OAM frame, OAM symbol, OAM message, and OAM status.

OAM field: A 10-bit field in each ~~PHY~~ RS-FEC frame reserved for the OAM symbol as described in 165.3.2.2.14 or in each refresh cycle as described in 165.3.6.3.

Change on page 89

rx_boundary

This variable is set to TRUE whenever the receive data stream reaches the end of a ~~Reed-Solomon~~ RS-FEC frame and a dummy OAM symbol is not expected per 165.3.9.2.1 during normal power operation in the data mode, or at the end of a received refresh cycle during Low Power Idle operation. This variable is set to FALSE at other times.

….

tx_boundary

This variable is set to TRUE whenever the transmit data stream reaches the start of a ~~PHY~~ RS-FEC frame and a dummy OAM symbol is not transmitted per 165.3.9.2.1 during normal power operation in the data mode, or at the start of a transmit refresh cycle during Low Power Idle operation. This variable is set to FALSE at other times.

Changes on page 92

165.4.2.4 PHY Control function

PHY Control generates the control actions that are needed to bring the PHY into a mode of operation during which frames can be exchanged with the link partner. PHY Control shall comply with the state diagram in Figure 165–27.

During PMA training (TRAINING and COUNTDOWN states in Figure 165–27), PHY Control information is exchanged between link partners with a 12-octet Infofield, which is XORed with the first 96 bits of the 16th partial ~~PHY~~ frame (bits 17550 to 17645) of each training frame. The Infofield is also denoted IF. The link partner is not required to decode every IF transmitted but is required to decode IFs at a rate that enables the correct actions prior to the PAM2 to PAM4 transition.

The 12-octet Infofield shall include the fields in 165.4.2.4.2 through 165.4.2.4.8, also shown in Figure 165–22 and Figure 165–23. Infofield shall be transmitted at least 256 times with each change to octets 7 to 10.

Changes on page 93

165.4.2.4.3 Partial ~~PHY~~ frame count (PFC24)

The partial ~~PHY~~ frame count consists of 3 octets [Oct4<7:0>, Oct5<7:0>, Oct6<7:0>] and indicates the running count of partial ~~PHY~~ frames sent LSB first. The partial frame count increased by one for every 1170 PMA_UNITDATA clocks. There are 16 partial ~~PHY~~ frames per ~~PHY~~ training frame and the Infofield is embedded within the 16th partial ~~PHY~~ frame. The first partial ~~PHY~~ frame is zero, thus the first partial ~~PHY~~ frame count field after a reset is 15.

PFC24 continues to run uninterrupted for the duration of the link. The resolution of PFC24 is large enough that it does not rollover during the allotted training time. However, it will rollover if allowed to run indefinitely. PFC24 is defined to rollover to 0 after it reaches 16 776 959 to align with the EEE QR cycle.

Changes on page 94

165.4.2.4.6 Data switch partial ~~PHY~~ frame count

When PMA_state<7:6> = 01, then [Oct8<7:0>, Oct9<7:0>, Oct10<7:0>] contains the data switch partial ~~PHY~~ frame count (DataSwPFC24) sent LSB first. DataSwPFC24 indicates the partial ~~PHY~~ frame count when the transmitter switches from PAM2 to PAM4, which occurs at the start of an RS-FEC superframe. The last value of PFC24 prior to the transition is DataSwPFC24 – 1. DataSwPFC24 shall be set to an integer multiple of 32. ~~When the value of DataSwPFC24 is a multiple of 16 the switch from PAM2 to PAM4 occurs on a PHY frame boundary~~. DataSwPFC24 shall be a minimum of 4081 and a maximum of 4785 from the current PFC24 value.

Change “Reed-Solomon frame” to “RS-FEC frame”

Global change of “Reed-Solomon frame” to “RS-FEC frame” (or frame to frames) (and “a” to “an” where appropriate)

Page: 67

- L32: block of a **Reed-Solomon frame.**
- L34: composed of eight **Reed-Solomon frames** that contain

Page: 68

- L40: interleaved and the **Reed-Solomon frames** are decoded

Page: 75

- L44: Number of **Reed-Solomon frames** with uncorrectable
- L49: Number of **Reed-Solomon frames** received over bit error ratio interval.

Page: 76

- L4: when a new **Reed-Solomon frame** is available for testing
- L5: A new **Reed-Solomon frame** is available for testing when the Block Sync
- L7: the next **Reed-Solomon frame.**
- L39: TRUE if received **Reed-Solomon frame** is valid.
- L39: **Reed-Solomon frame** is valid if and only if all parity

Page: 80

- L27: number of invalid **Reed-Solomon frames** within the current
- L28: RFRX_CNT_LIMIT **Reed-Solomon frame** period
- L31: Count number **of** **Reed-Solomon frames** received during current period.
- L36: that a full **Reed-Solomon frame** has been decoded and the

Page: 89

- L6: end of a **Reed-Solomon frame** and a dummy OAM symbol is not expected

Page: 133

- L38: block of a **Reed-Solomon frame** is detected by the PCS Transmit function.

Thank You