EEE ALERT signal for 100GBASE-KP4

Matt Brown, AppliedMicro
Bart Zeydel, AppliedMicro
Adee Ran, Intel
Kent Lusted, Intel

(Regarding Comments 39 and 10234)

Supporters

- Brad Booth, Dell
- Rich Mellitz, Intel
- Andre Szczepanek , Inphi

Overview

- Addresses comments 39 and 10234 for 802.3bj Draft 1.1.
- Proposal for 100GBASE-KP4 EEE ALERT signal used during REFRESH and WAKE.
- Base the ALERT signal on the training frame proposed in lusted_01_0912.

EEE overview

- EEE LPI transitions to quiet state with occasional refresh states.
- Tim to wake up and transition to normal data mode (tw_phy) is targeted at 5 us for 100G PHYs.

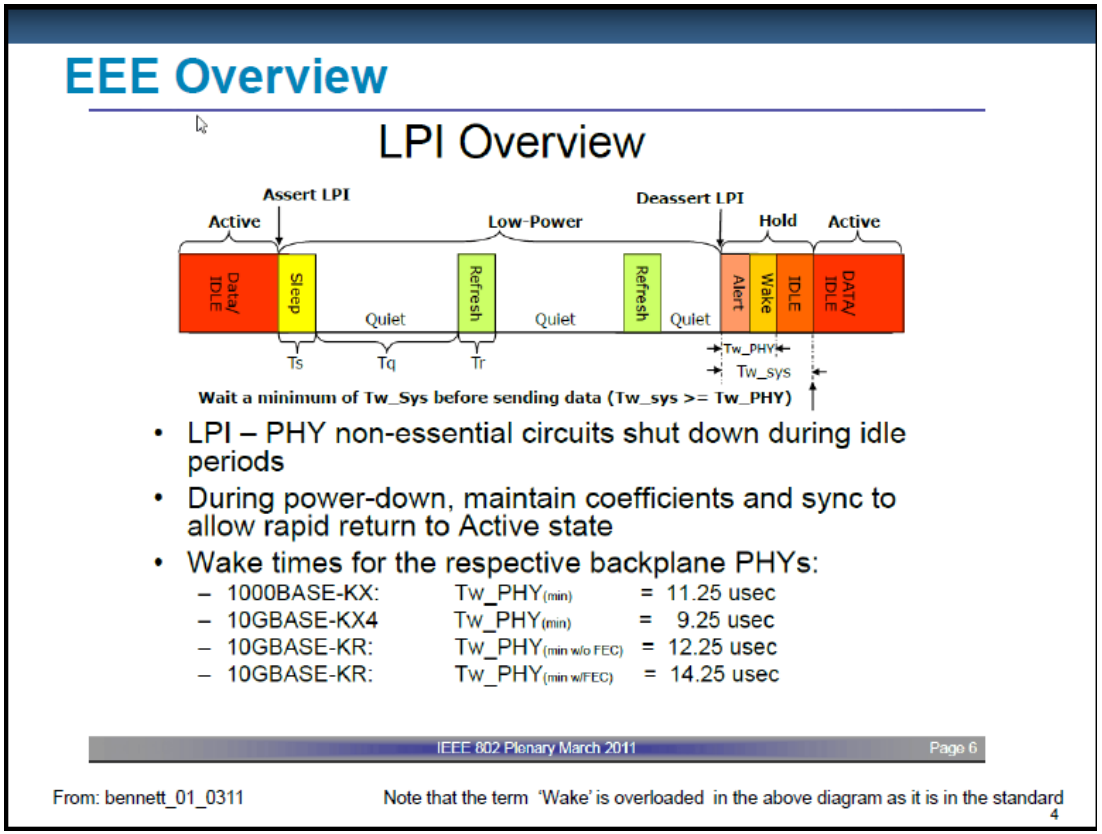

From gustlin_02_0112

EEE state machine

- Draft 1.1 Figure 82-16 (right) shows the transmit low power idle (LPI) state diagram.
- When transitioning from the QUIET state (TX_QUIET) toward either WAKE (TX_WAKE) or REFRESH (TX_RF_WAKE) an alert signal is transmitted.
- ALERT signal provides:
  - Strong signal to detect and initiate wake up.
  - Frame alignment signal for fast alignment to training frame and line coding.
  - Control channel to indicate EEE state and to handoff from ALERT frame to PMA frame.

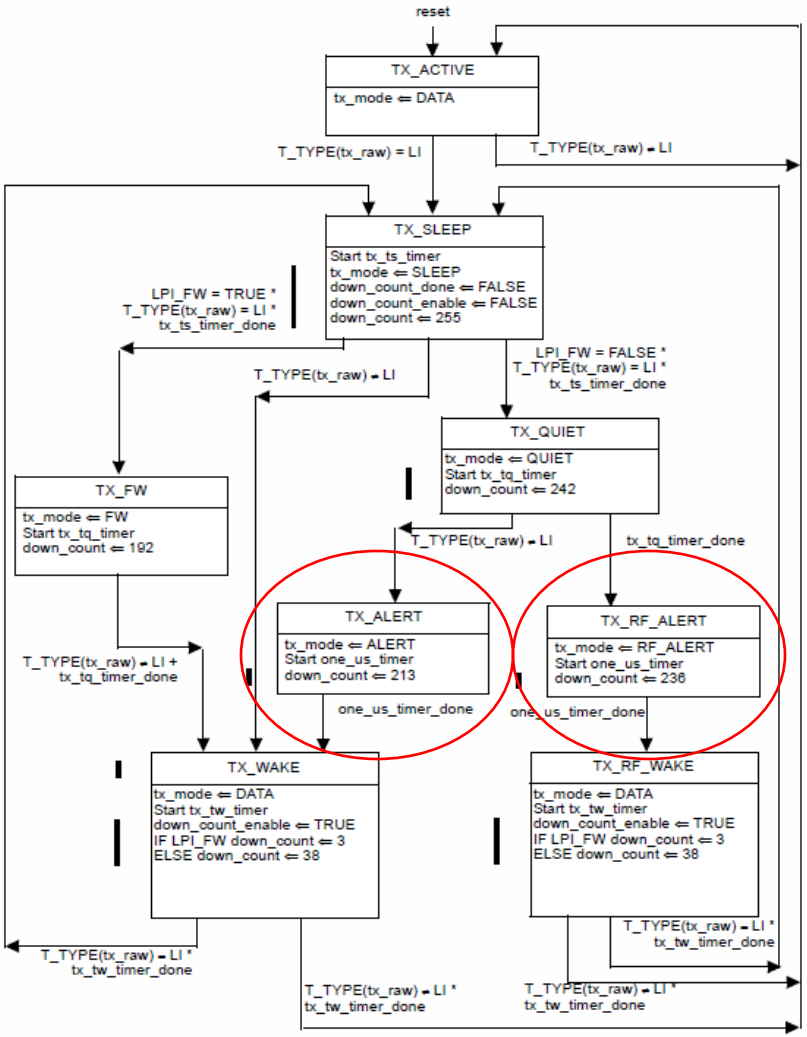

Figure 82–16—LPI Transmit state diagram

EEE Refresh and Wake Synchronization

- For EEE, it is necessary to synchronize very quickly to the signal on transitions from QUIET to WAKE or REFRESH.
- The PMA must synchronize in 3-4 us after receiving the ALERT signal.
  - Orders of magnitude faster than for initial synchronization.
  - The total targeted budget is 5 us, but this must be allocated among transmitter, power up, equalization settling/convergence, etc.
- Once the PMA/PMD is synchronized the RS-FEC layer must also synchronize.

Challenges

- The wake or refresh signal must be reliably discernible from noise to prevent missing or falsely detecting the WAKE/REFRESH signal.
- The PAM4 PHY (unlike PAM2) is not able to easily make use of the PCS alignment markers for synchronization due to the encoding and 4-level (rather than 2-level) signaling.
- The PHY receiver may not be able to effectively decode until synchronization is achieved.
- Even with effective equalization, without FEC synchronization the BER will be very high (~1E-5).

ALERT Signal

- The ALERT signal is a repeating ALERT frame.
- The ALERT frame is based on the training frame in lusted_01_0912 except:
  - The training pattern is truncated to 4320 bits.
    - Use the same seeds.
- The majority of the specified control channel fields are not used for the ALERT signal.
  - Unless specified otherwise in subsequent slides, the control channel is ignored.

100GBASE-KP4 EEE Alert Frame

Field	TB46 #	Size (TB46)	TFW #	Size (TFW)
Frame marker	0:1	2	0	1
Control channel	2:17	16	1:8	8
Training pattern	18:115	98	9:58	49
Total		116		58

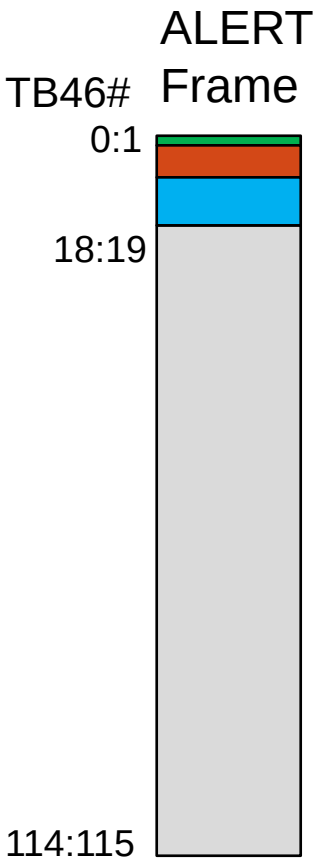

ALERT Frame Attributes

- The short ALERT frame enables completing the countdown sequence in less than 1 us
  - 116/2 = 58 TFW * 46 PAM4 symbols * 73ps = ~194 nsec
- The ALERT frame contains 116 TB46
  - Exactly 6 ALERT frames fit into a PMA frame.
  - Each ALERT frame is aligned to one of the 6 offsets in the PMA frame.
- The training patterns are slightly DC imbalanced, but the effect in LPI mode is insignificant (see backup slide).

Coefficient Update Field

Same as lusted_01_0912.pdf

11:7 Reserved
6 Parity Check

Cell ordering not finalized

In EEE mode, all these cells are transmitted as 0, ignored on reception.

Table 72–4—Coefficient update field

Cell(s)	Name	Description
15:14	Reserved	Transmitted as 0, ignored on reception.
13	Preset	1 = Preset coefficients 0 = Normal operation
12	Initialize	1 = Initialize coefficients 0 = Normal operation
11:6	Reserved	Transmitted as 0, ignored on reception.
5:4	Coefficient (+1) update	5 4 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold
3:2	Coefficient (0) update	3 2 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold
1:0	Coefficient (–1) update	1 0 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold

Status Report Field

In EEE mode…

- Cell 6 = 1 because training already completed
- Cells 5:0 are transmitted as 0, ignored on reception.
- Calculate parity field as in training frame
- Cell ordering not finalized

Cell(s)	Name	Description
19	Parity Check	Parity calculation for Status Report Field
18:14	EEE State	Current EEE state of local transmitter, if EEE is implemented.
13:12	Training Frame Countdow	Number of training frames remaining before link training process transitions to data mode
11:7	PMA Alignment Offset	Relative location of the next training frame within the PMA frame
6	Receiver ready	1 = The local receiver has determined that training is complete and is prepared to receive data. 0 = The local receiver is requesting that training continue.
5:4	coefficient (+1) status	5 4 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated
3:2	coefficient (0) status	3 2 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated
1:0	Coefficient (-1) status	1 0 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated

Status Report Cells Used for EEE - 1

- EEE State (Cells 18:14)
  - Cell 18 indicates mode
    - 0 = training (the link is in start-up training mode)
    - 1 = EEE (the link is in LPI mode)
  - Cells 17:16 indicate EEE state (see 802.3bj draft 1.1 80.3.3.4.1)
    - 00 = Wake, 01 = Refresh, other values reserved
  - Cells 15:14, reserved and set to 0.
- Cells 13:12 Countdown counter
  - Same as lusted_01_0912.pdf

Status Report Cells Used for EEE - 2

- Cells 11:6 PMA Alignment Offset (PAO)
  - Same function as lusted_01_0912.pdf
  - In EEE mode, PAO encodes the relative location of the TB46 after the end of the Alert frame as a 3-bit integer in the range 0 to 5.
  - The start of the next Alert frame is 116 * PAO offset from the start of the PMA frame
    - 0: marker aligned with 40-bit overhead
    - 1: marker is at offset of 1*116=116 termination blocks from 40-bit overhead
    - …
    - 5: marker is at offset 5*116 =580 termination blocks from 40-bit overhead
    - 6 to 7: invalid, never transmitted, ignored on reception

ALERT Frame Summary

- Long period frame marker and control channel permit efficient detection and alignment.
- Pseudo-random pattern allows fine phase alignment and receiver convergence.
- Small frame size and the countdown/PAO control fields allow precise cutover from the ALERT signal to the PMA frame.
- PMA frame and thus FEC are aligned and error-free data detection and alignment marker synchronization may begin immediately.
- EEE state control field gives early indication of the EEE state (e.g., wake vs refresh).

Conclusion

- Specify 100GBASE-KP4 ALERT frame as proposed.

BACKUP

PMA Alignment Offset Example

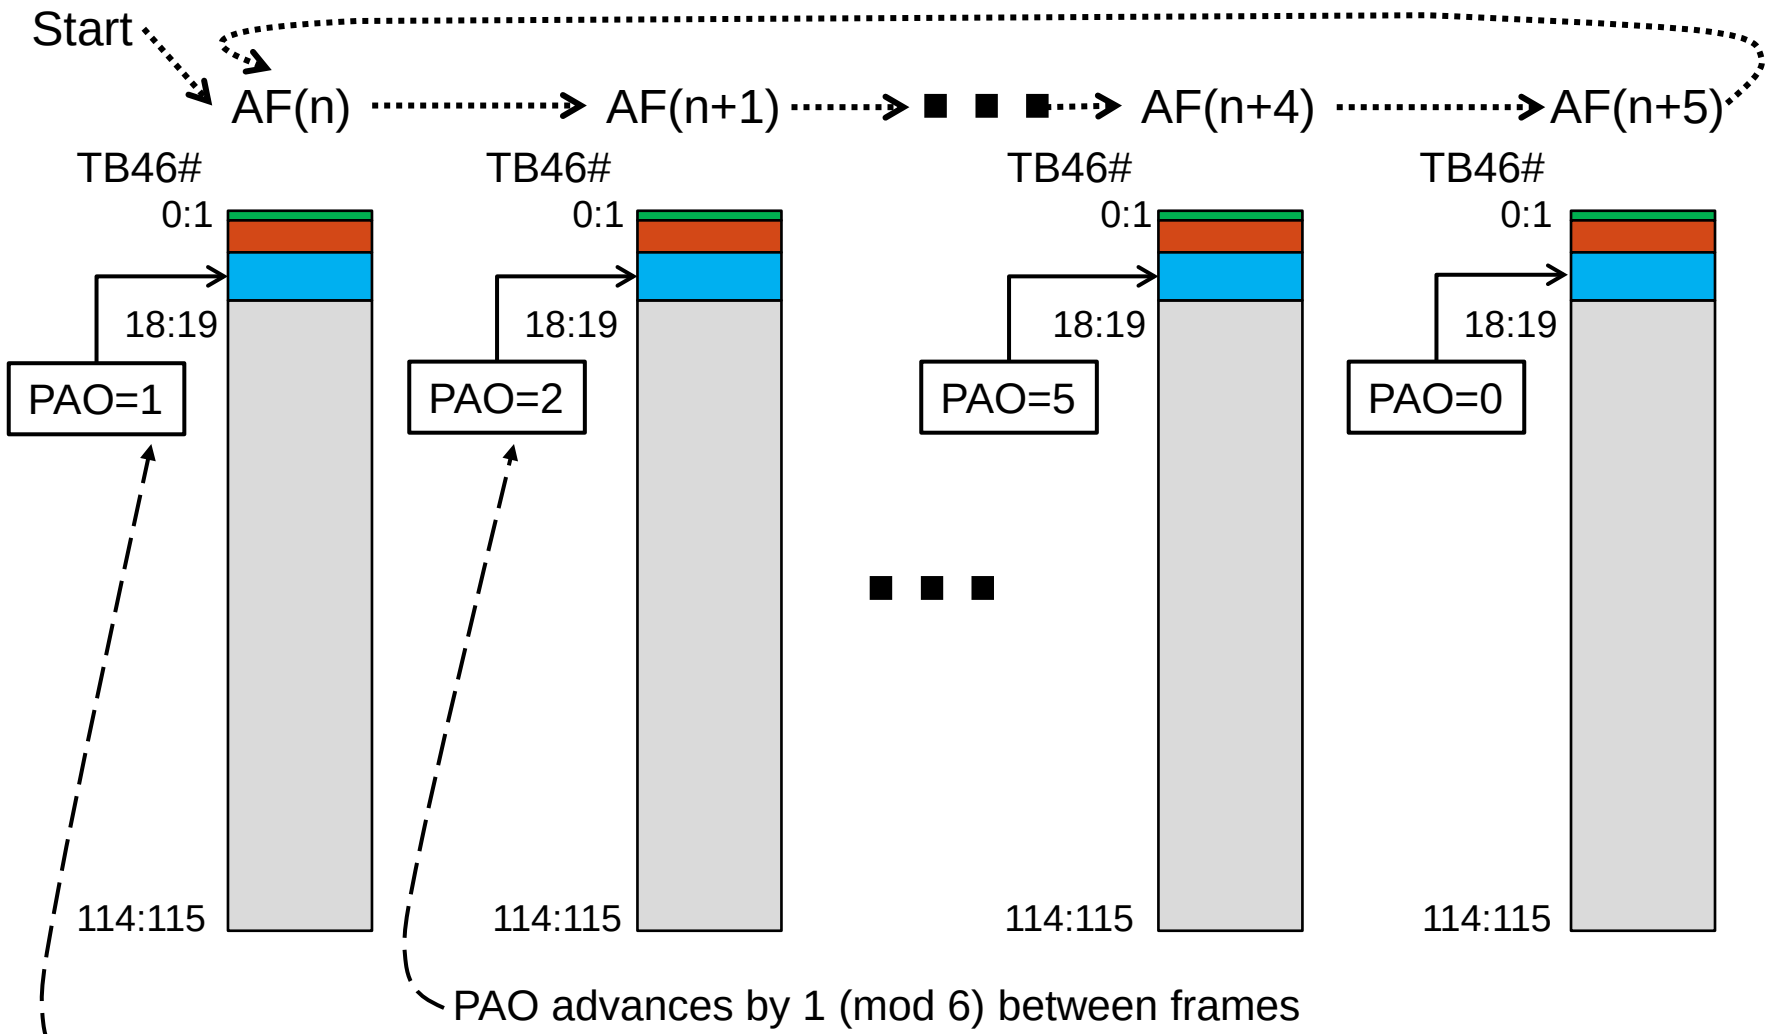

1 ALERT Frame = 116 Termination Blocks (TB46)

DC Balance

- Worst case running disparity = 32 of the +1 PAM4 symbols
  - Up to 1.5% of the truncated training pattern shift in DC balance