



AQUANTIA

ACCELERATING CONNECTIVITY

EEE for 802.3ch
Saied Benyamin
Dr. Paul Langner

AQUANTIA

EEE


- Previous presentations have discussed EEE implementation similar to 1000BASE-T1.
- This presentation proposes EEE based on 1000BASE-T1 (Clause 97) with the following functional modifications:
 - Optional support for OAM transmission during LPI
 - Add low-frequency Alert signal before wake, similar to 10GBASE-T
 - Take out limitation on Wake frame start (i.e. start anywhere)
 - Master/Slave refresh are exactly half QR cycle apart

Timing Synchronization – Partial Frames

- Similar to 1000BASE-T1, we will have a frame counter synchronized between master and slave so that we can align refresh properly. This counter can be a maximum of one count different between master and slave. The counter must be set to zero align with the start of a super-frame when interleaving is used.
 - All proposed full RS frames are quite long in 2.5G/5G/10GBASE-T1
 - Concept of partial frame counter was introduced in 1GBASE-T1
- We define a partial frame as 1/8 the time period of a PHY FEC frame. The following timings are based on FEC of RS1024(576,514).

RS1024(576,514)	2.5G	5G	10G
PHY frame (ns)	2048	1024	512
Partial PHY frame (ns)	256	128	64
PAM4 symbols per partial Frame	360	360	360

Signal Timing



	PFrame	2.5G	5G	10G
Quiet/Refresh Cycle	128	32.768us	16.384us	8.192us
Refresh	4	1.024us	512ns	256ns

- $lpi_offset = lpi_qr_time / 2$

Alert

- Similar to 10GBASE-T, in order to reduce complexity, a low frequency Alert signal is introduced. This signal can be detected without the need for use of the full DSP path. Alert and sleep can start at the beginning of any full RS super-frame. This means that if we use 4x interleaving, Alert may have to wait up to 3 additional frames to start. Receive side can align start of detection with start of super-frame to save additional power.
- Alert sequence is a unique PAM2 64-bit sequence repeated 7 times, followed by sixty four zero's for a total of 512 bits. We use unique sequences for master and slave.
- The sequence is directly copied from 10GBASE-T. Whereas the sequence is bit- replicated 2x in 10GBASE-T, we replicate it 16x here.
 - 512 bits x 16 bit-repeat = 8192 symbols
 - This sequence is padded by 448 zero's to fill three RS times (2880 symbols per RS frame)
- Alert is always followed by one full WAKE superframe of idles
 - Need to send interleaved superframe to ensure reception within noise environment



Asymmetric Channel Options

1. LVDS SERDES one direction with I²C FDM'ed onto cable in DC null
 - Pros: Simple, and low power
 - Cons: No upstream OAM, fixed low upstream BW, and PoDL potentially an issue
 2. Standard SERDES one direction, with synchronous, lower BW signal upstream
 - Pros: AC coupled upstream
 - Cons: Non-standard system I/F, fixed low BW upstream with no additional on demand bandwidth increase
- The issues with both of these approaches are:
1. No standard system interface
 2. Custom SERDES required to support asymmetric rates
 3. Fixed upstream BW
 4. Long PAUSE response times due to channel asymmetry
 5. No EEE

Asymmetric Channel Options with EEE

1. Regular EEE
 - Pros: On-demand BW by dropping out of EEE, standard system I/F
 - Cons: No background OAM
2. Regular EEE with upstream OAM as in 1GBASE-T1
 - Pros: On-demand BW by dropping out of EEE, standard system I/F
 - Cons: ?
3. Low-upstream BW during EEE
 - Pros: low bandwidth upstream plus On-demand BW by dropping out of EEE
 - Cons: Complicated non-standard system I/F, potentially requiring a modified USXGMII scheme to allow asymmetric rates and switching in and out of full-rate


➤ **Recommendation is #2 if OAM is required, and #1 if not required**

EEE Examples

1. 400 kb/s SMBus with write word command (48 bits) every second
 - Map into 72 byte packet (57.6 ns @ 10G)
 - With 3072ns turn-on and 512ns turn-off and 4x interleaved RS frames minimum payload size, this gives: $3072 + 2048 + 512 = 5632$ ns on every second = **0.00055% duty cycle**
2. 400 kb/s SMBus with back-to-back write word commands (48 bits)
 - Map into 72 byte packets (57.6 ns @ 10G) sent every 120 us
 - With 3072ns turn-on and 512ns turn-off and 4x interleaved RS frames minimum payload size, this gives: $3072 + 2048 + 512 = 5632$ ns on every 120us = **4.7% duty cycle**
3. 40 Mb/s upstream channel mapped into 1500 byte packets (1.2us @ 10G which requires six 4x interleaved RS frames) sent every 287.2us
 - This gives: $3072 + 6 \times 2048 + 512 = 15.872$ us every 287.2us = **5.5% duty cycle**

Proposed Refresh signaling 10GBASE-T1

- PCS will be sending out zero's during LPI, except when the RS frame coincides with a refresh. In this case, the OAM alone is sent out as the last 10 bits of the frame.
- We can use repetition coding to make sure OAM is not corrupted since it is not protected by RS or interleaving as in data mode.




Backup



1GBASE-T1 EEE refresh

- PCS sends out zero's during LPI in place of data. For the RS frame that coincides with refresh, OAM and RS parity are transmitted.



Thank you.

AQUANTIA®
ACCELERATING CONNECTIVITY