Backplane Ethernet Low-Power Idle

Brad Booth, AMCC

May 2008

Contributors/Supporters

- D. Koenen – HP
- J. Chou - Realtek
- G. Zimmerman - Solarflare

Backplane PHYs

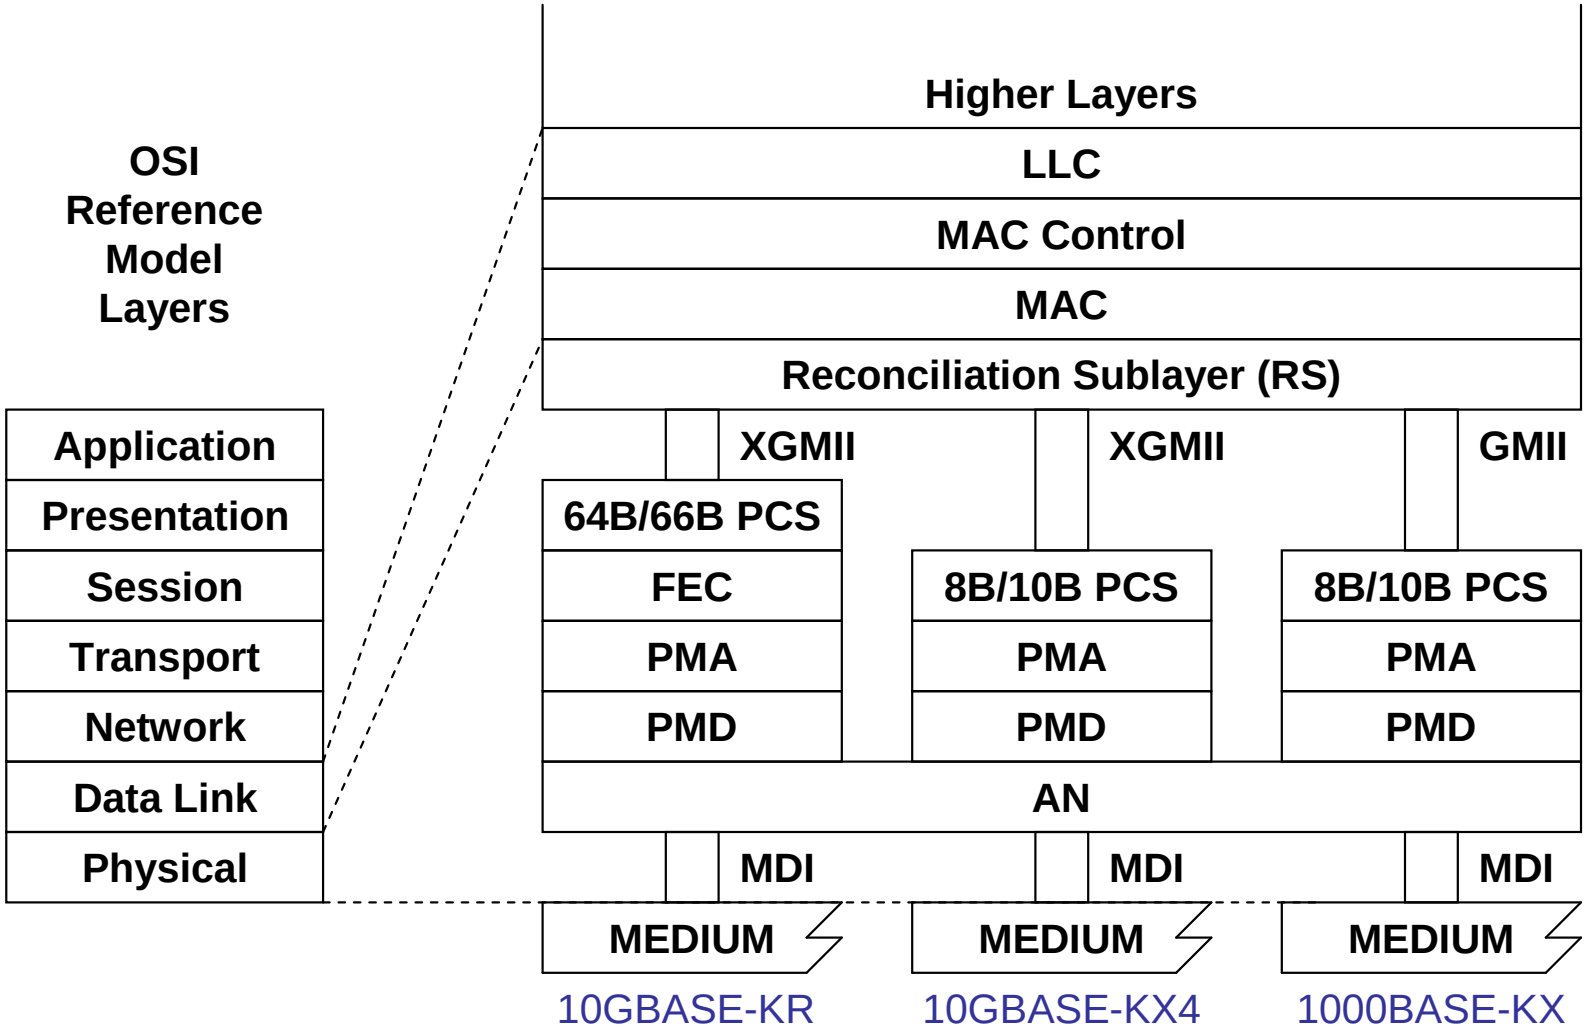

EEE should consider inclusion of 1000BASE-KX

1000BASE-KX

- At one point in time, it was assumed that 1000BASE-KX would be lower power PHY for 10GBASE-KR
- LPI changes that assumption
- P802.3az should consider energy-efficient operation for the 1000BASE-KX PHY
- Presentation also discusses 1000BASE-KX
  - Even though could be considered out of scope at this time

Backplane LPI

- Transmitter goes quiescent during quiet state
  - Similar concept to PCIe specification L2,L3 states
  - 10GBASE-KR has greater latency due to coding scheme & FEC
- Refresh operation
  - KR to use training mode during refresh cycle
  - KX4 and KX to use set idle pattern
  - FEC not refreshed
- Sleep transition
  - Fairly simple after the request is indicated
- Wake transition
  - Multi-staged to decrease synchronization requirements

LPI Operating States

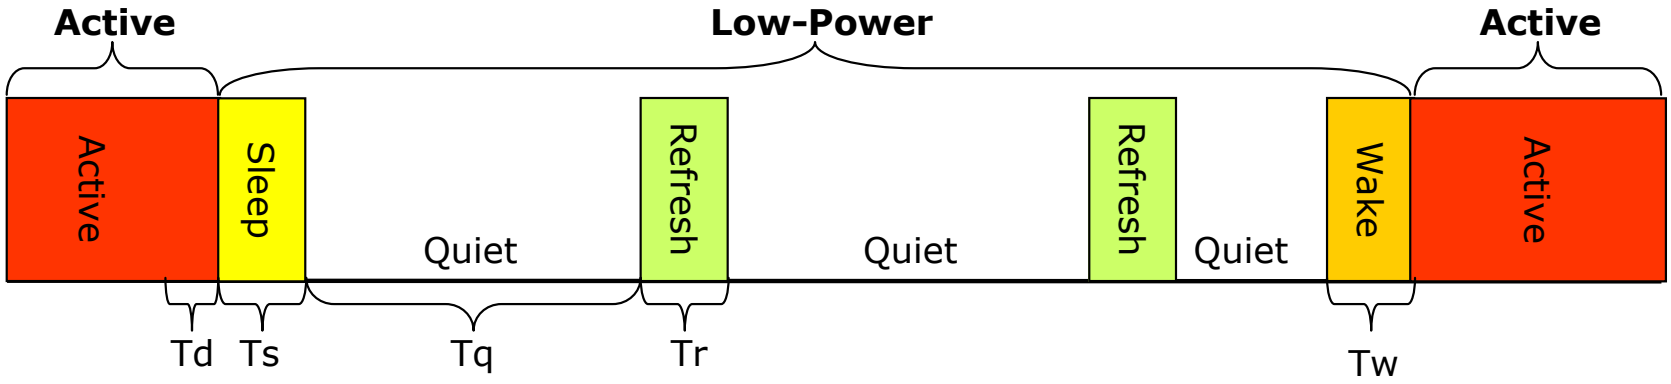

Term	Description
Active state	Existing state used for data transmission where either data packets or IPG/Idle symbols are transmitted.
Low-Power state	New state used during periods of no data transmission to allow system power reduction between data packet bursts.

- No change to this concept with backplane LPI

LPI New Line Signals

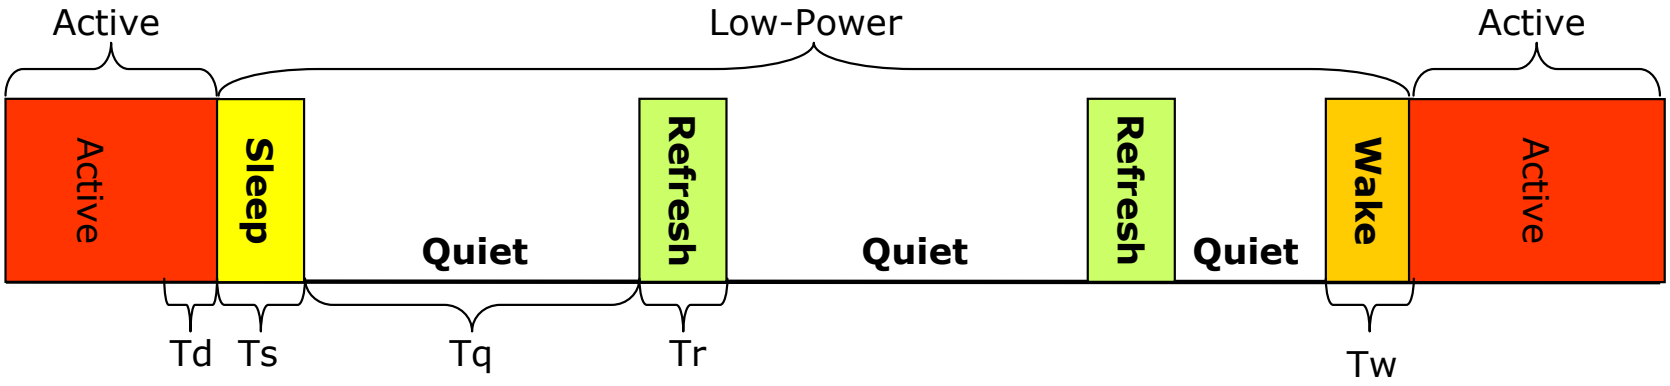

Term	Description
Sleep	Signal to inform link partner of entry into low power state
Quiet	Transmission of training frames only
Refresh	Periodic signal during low power state for PHY to maintain timing recovery and/or coefficients
Wake	Signal to inform link partner of entry back into active state

- Same concept as with BASE-T technologies

LPI Timing Parameters

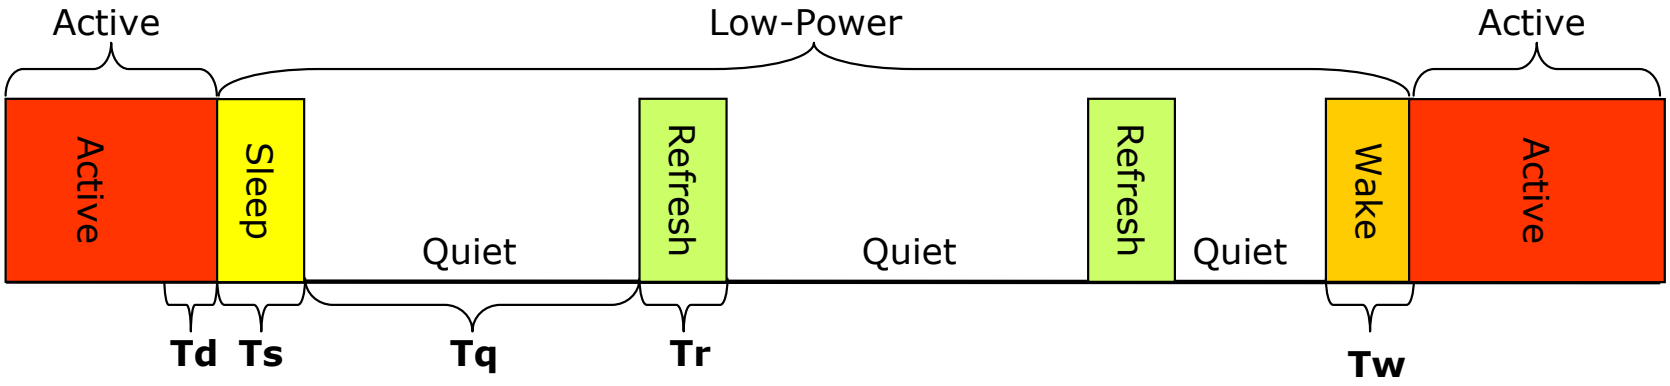

Term	Description
Decision Time (Td)	Higher-layer control policy timing; out of scope
Sleep Time (Ts)	Min. duration Sleep symbols sent before going to Quiet
Quiet Duration (Tq)	Max. duration PHY remains Quiet before Refresh
Refresh Duration (Tr)	Min. duration PHY sends Refresh symbols
Wake Time (Tw)	Max. period to permit the receiving system to wake up
Propagation Delay (Tp)	Max. transmission delay of the media

- No change to this concept with backplane LPI

Entering Quiet Mode

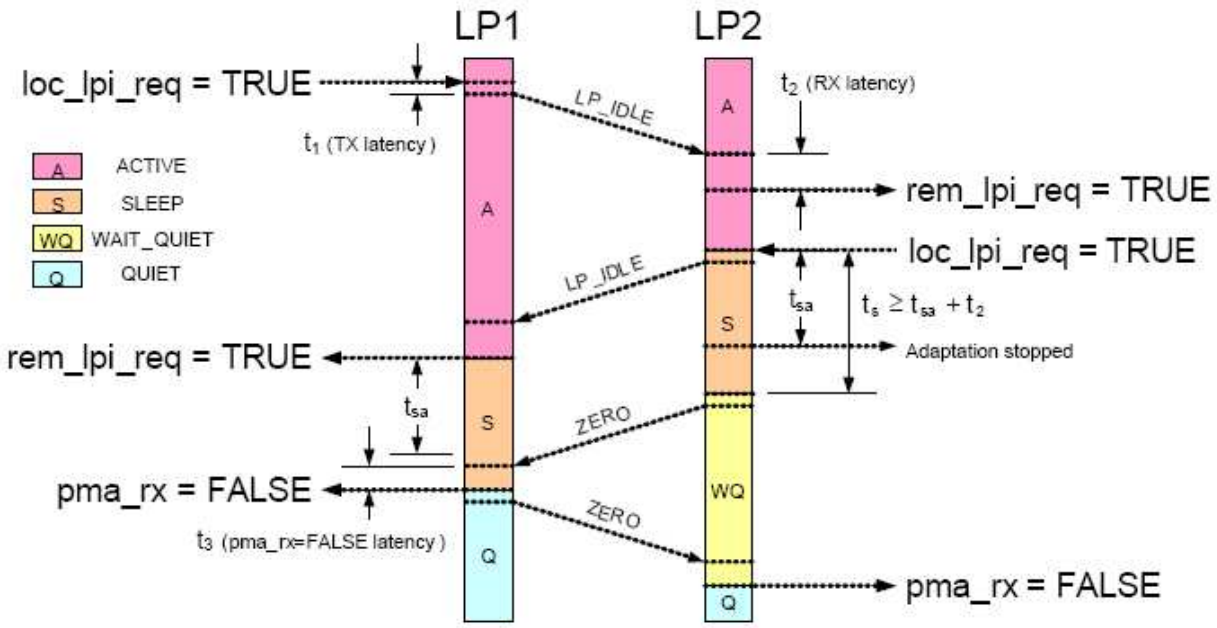

- Borrowed from Adam Healey
- Pass message to indicate transition to quiet
  - Ordered set with KR/KX4, special 8B/10B code with KX

Refresh Cycle

- KR
  - Use of training mode
  - Cell 14 of status report field changed from Reserved to be Refresh
- KX4
  - Use of KAR idle pattern only on lane 0
  - Lack of signals on lanes 1-3 indicates a refresh state
- KX
  - Per Dan Dove's presentation

/LPI/	Low Power Idle		Correcting /LI1/, Preserving /LI2/
/LI1/	Low Power Idle 1	2	/K28.5/D6.5/
/LI2/	Low Power Idle 2	2	/K28.5/D26.4/

KR Wake Cycle

- Enter Refresh Cycle first
  - LP2 indicates when receiver is ready
  - When LP1 is ready, transition to Wake state
- Wake Cycle
  - Add Wake to cell 13 of status report field
  - Transmits training pattern at line rate
  - Change to Active immediately after ready indication
- Known transition boundary
  - Simplifies FEC & coding sync

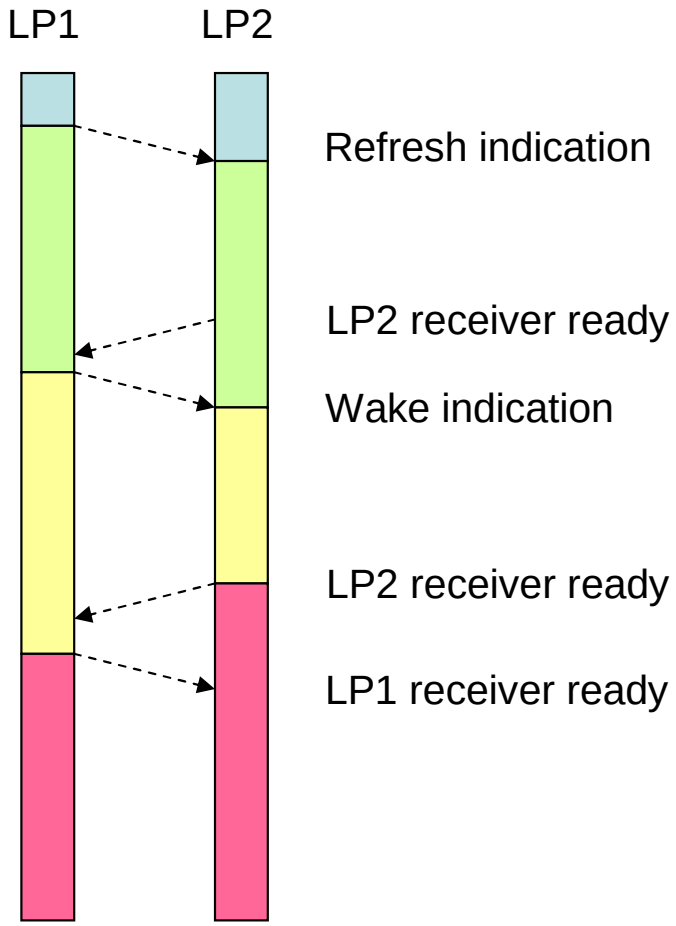

KX4/KX Wake Cycle

- Enter Wake Cycle directly
  - KX4 uses idle pattern on all 4 lanes
  - KX uses normal idle pattern
- Ordered sets
  - KX4 can use RF and LF to indicate when ready
  - Otherwise, just based upon time
- No FEC complications
  - Yippee!

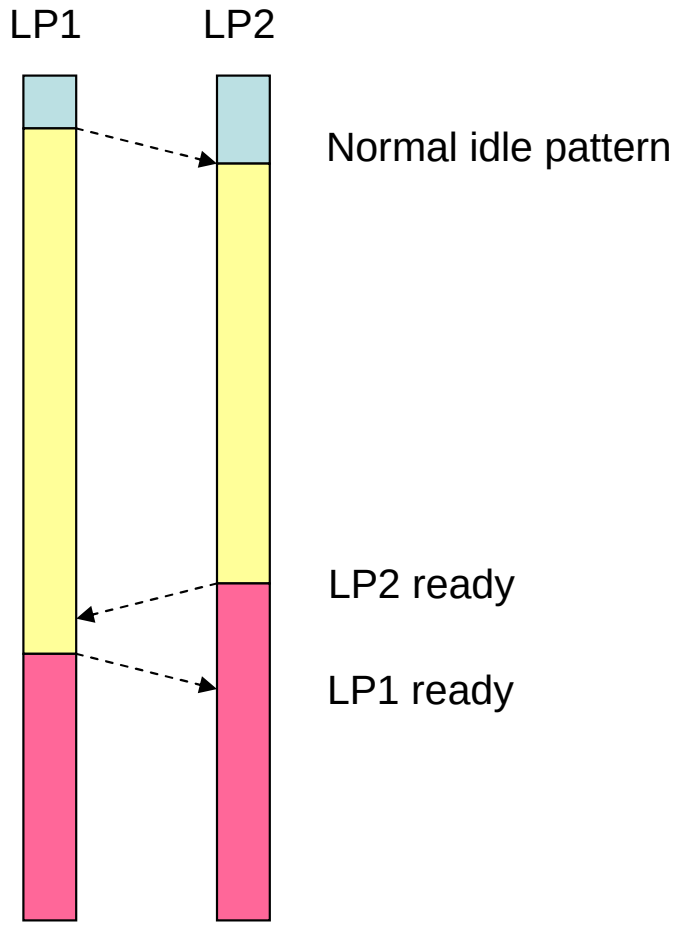

XAUI

- Not directly specified in the scope
- Expectation is this interface should exhibit energy-efficient properties
- Make use of the KX4 ability to switch to LPI
  - Very useful if the MAC and PHY are on different clock sources
  - Implementers could see greater savings if MAC and PHY on the same clock source
    - Transmit disable for XAUI could be used for this
    - No requirement to specify this behavior
  - Could also be applied to 10GBASE-CX4

Summary

- LPI can be applied to backplane Ethernet PHYs
  - Training mode for KR
  - Set idle patterns for other PHYs
- Include XAUI and 1000BASE-KX
  - Expectation that XAUI could exhibit power saving properties
  - 10GBASE-CX4 could also make use of this property
  - 1000BASE-KX should have energy saving properties
    - A slight oversight when we scoped the project
- Asymmetric support
  - Use of ordered sets in 10G (KX4 and KR)
  - Use of special idle codes for 1G (KX)