AQUANTIA
ACCELERATING CONNECTIVITY

EEE Improvements For Highly Sparse Traffic

Saied Benyamin

Ramin FarjadRad

George Zimmerman

Statement of problem

- The automotive infrastructure includes a large number of camera interfaces. These interfaces currently use MIPI or LVDS which are simplex connections and do not carry the burden of power consumption for the upstream direction. The upstream commands to the camera use low speed, low power consumption communication such as I2C.
- In order for ethernet to proliferate into this interface, it is highly desirable for us to have a solution with competing power.
- While the standard does not set power consumption specifications, it must allow for implementations to save power where possible.
- Given where we are in the lifecycle of the standard, the solution must be simple and cause minimal changes to the standard, and be compatible with a physical layer specification.

EEE operation

- EEE is controlled by the LPI client and allows the MAC & higher layers to go to sleep while the PHY sends special LPI signaling.
- The Physical layer only knows if it is in LPI or not.
- We get to specify the way the LPI is signaled on the line, but not up the stack.

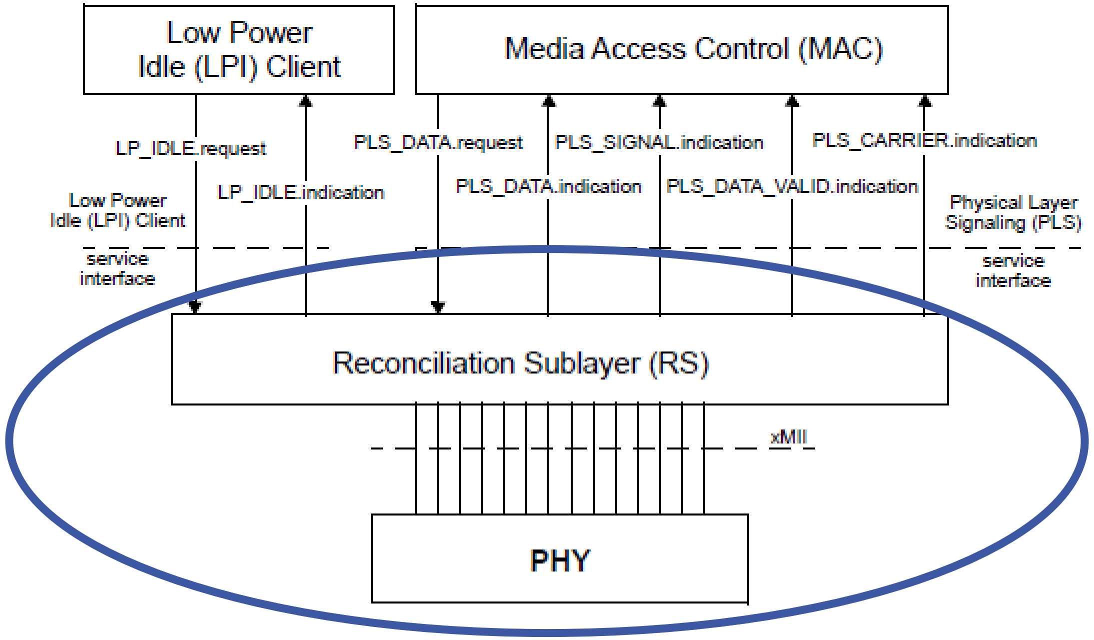

Figure 78–1—LPI Client and RS interlayer service interfaces

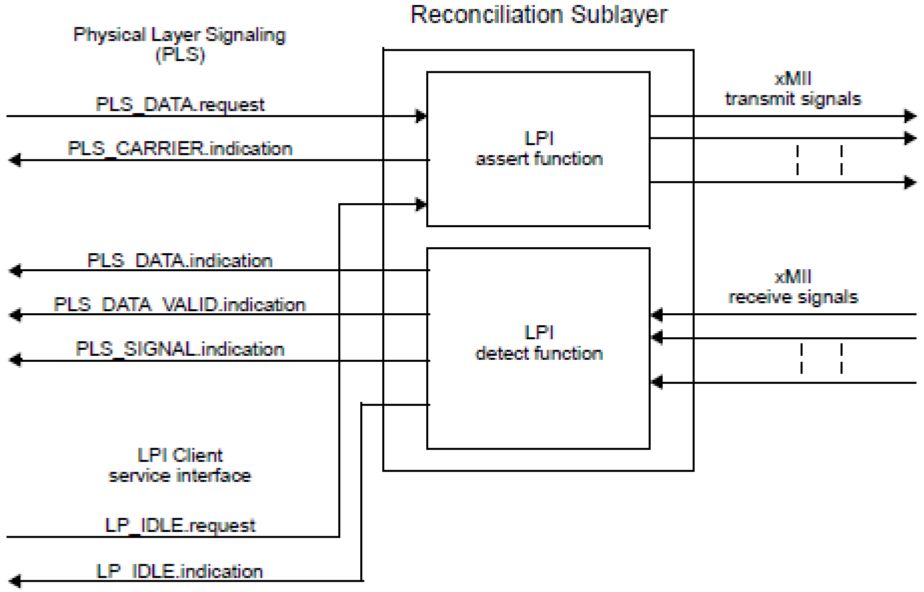

Figure 78–2—RS LPI assert and detect functions

EEE parameters

- We get to specify timing parameters associated with wake up
- These can be situational, if they are set via LLDP
- While Tw_phy is a minimum, Tw_sys_tx can be set longer, to accommodate longer PHY wake up.

Figure 78–5 illustrates the relationship between the LPI mode timing parameters and the minimum system wake time.

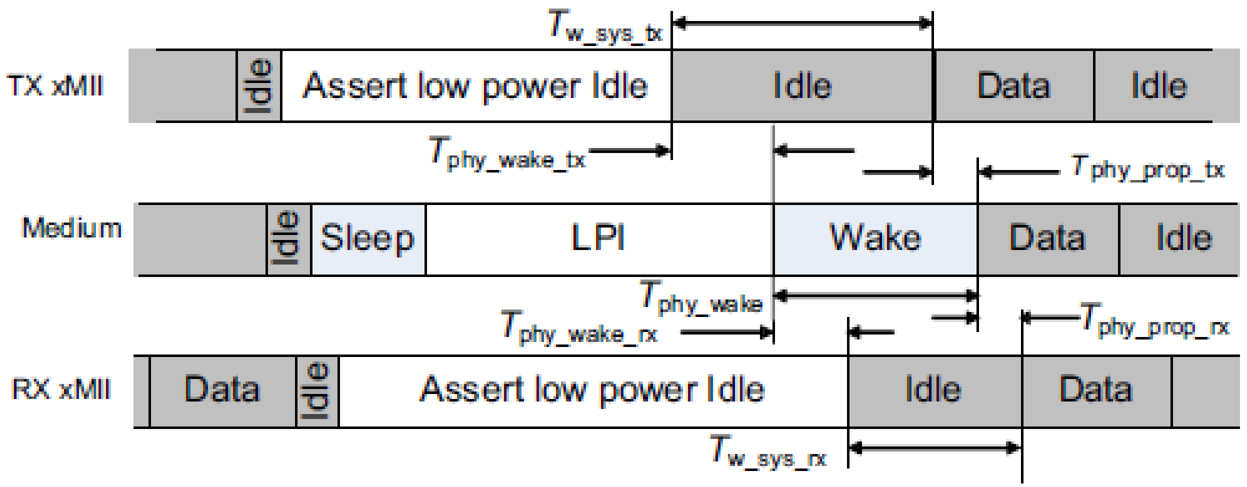

EEE power management

EEE was devised to bring power efficiency to low density traffic; these traditional low-power idle implementations have two components that affect power consumption:

- Power consumption of refresh
  - The quiet/refresh ratio is determined primarily by the needs of tracking clock drift. However, only the slave needs to do loop-timing, so this might be improved at least in one direction, saving power.
- Power consumption of alert detection
  - Since ALERT is a digital signal, using similar signaling to the data mode, large parts of the full receiver are generally activated to detect ALERT. A simpler ALERT, suitable for analog detection (or a simpler digital detection) would reduce this power consumption.
  - While the receiver can power down portions of its circuitry during quiet periods, it must still be constantly or very frequently listening for Alert. Extra power consumption needed because of demands for rapid ‘instant’ return to data mode. When traffic is not backing up, a longer time may be allowed to return to data transmission.
  - Many applications, such as camera devices, can tolerate a slow return to data traffic in the direction with sparse traffic. Taking advantage of this leads to possible simplifications in ALERT timing and recovery to full speed.

Proposed Improvements

- Improvements to LPI signaling in 802.3ch EEE to create opportunities for power savings.
  - Increasing the Quiet/Refresh ratio in the slave transmit to master receive direction
  - Enabling simple analog detection of the alert
- Optional improvements set by management based on use case (e.g., camera device)
  - Increasing Tw_sys_tx (set by management) and limiting when alert can be sent

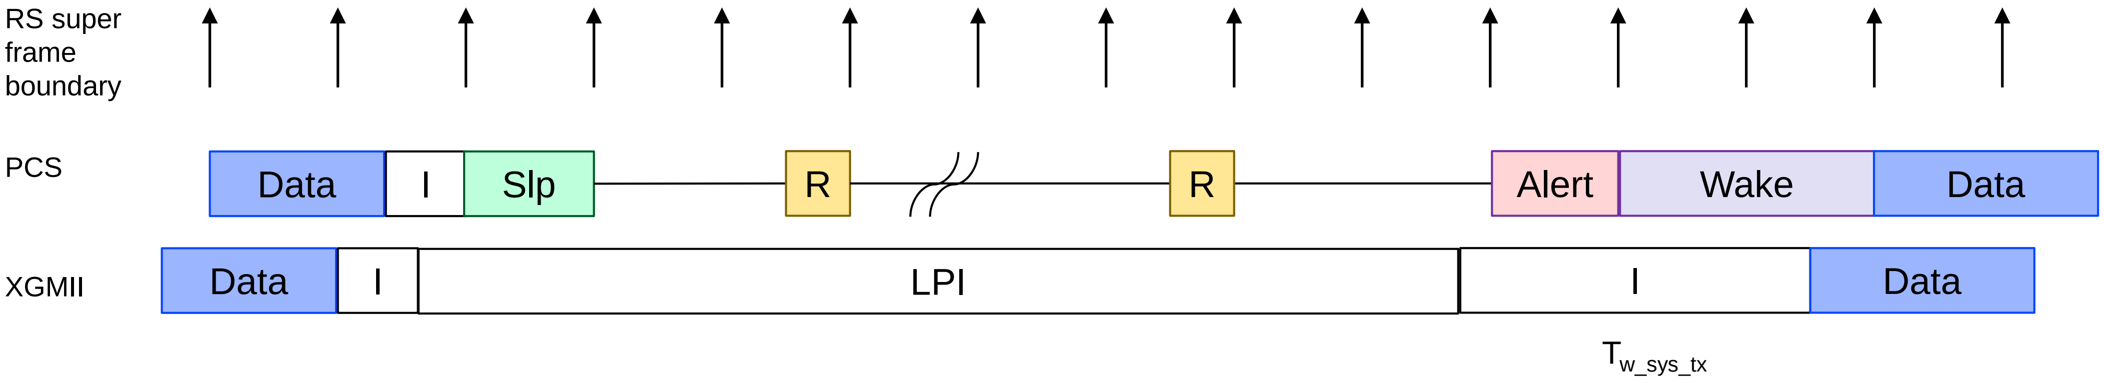

Change #1 – increase QR cycle

- The quiet/refresh ratio is determined primarily by the needs of tracking clock drift. However, only the slave needs to do loop-timing, so this might be improved at least in one direction, saving power.
- Reduce the refresh rate for slave transmitter since only the filters need to be refreshed on the master.
- Alternately, rather than reducing the refresh rate, the receiver (either master or slave) shall consider the SNR before determining the need for filter refresh.

Change #2. Change Alert to a sine wave

- Link Sync pattern used for Alert was not designed for power efficiency
- Using a sine wave avoids ISI, hence removing the need for filtering
- Using a low frequency signal avoids the need for echo cancellation
- Using two different frequencies for master/slave increase detection margin
- Propose a 180MHz/90MHz sine wave pattern of the same 1.28uS duration for alert for master/slave
- The pattern will be generated using PAM4 signaling
- A high tap count correlator is not required since a simple bandpass filter can detect it (in analog or using subsampling in digital domain)
- Emission impact from this pure frequency is negligible due to its short duration

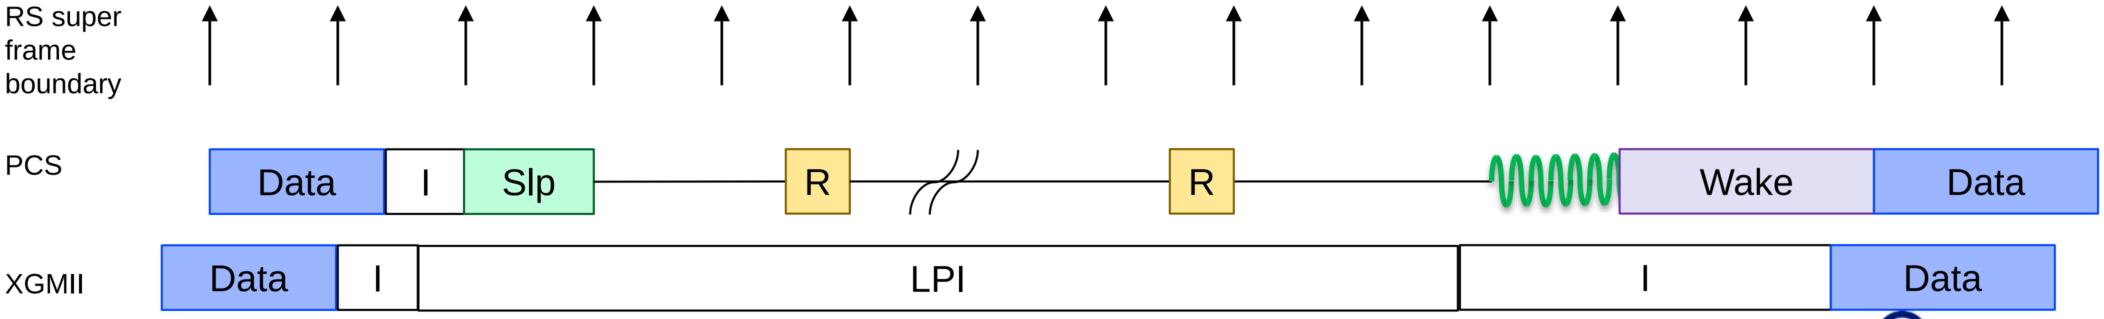

Use case based power savings

- The next proposal has minor negative effect on the system performance by extending the wake period, but is well suited for many use cases such as a camera interface where the longer wake period is inconsequential
- As such, we propose that an enable bit is employed called ‘Slow Wake’
- The system would enable this mode only when it knows that the traffic is light and bursty and the wake delay is not critical
- The exact method of setting of this bit is outside the scope of this standard, the LLDP can be used to control this bit on the link partner.
- Tw_sys_tx as discussed earlier will need to be adjusted on the MAC as well.

Longer wake-up can be accommodated via LLDP parameters

- LLDP exchanges allow a PHY to request the link partner to wait longer before sending data

79.3.5.1 Transmit T_w

Transmit $T_{w\_sys\_tx}$ (2 octets wide) shall be defined as the time (expressed in microseconds) that the transmitting link partner will wait before it starts transmitting data after leaving the Low Power Idle (LPI) mode. This is a function of the transmit system design and may be constrained, for example, by the transmit path buffering. The default value for Transmit $T_{w\_sys\_tx}$ is the T_{w_phy} defined for the PHY that is in use for the link. The Transmitting link partner expects that the Receiving link partner will be able to accept data after the time delay Transmit $T_{w\_sys\_tx}$ (expressed in microseconds).

79.3.5.2 Receive T_w

Receive $T_{w\_sys\_tx}$ (2 octets wide) shall be defined as the time (expressed in microseconds) that the receiving link partner is requesting the transmitting link partner to wait before starting the transmission data following the LPI. The default value for Receive $T_{w\_sys\_tx}$ is the T_{w_phy} defined for the PHY that is in use for the link. The Receive $T_{w\_sys\_tx}$ value can be larger but not smaller than the default. The extra wait time may be used by the receive link partner for power-saving mechanisms that require a longer wake-up time than the PHY-layer definitions.

Change #3. Less Frequent Alert, Slow Wake

- Alert is currently allowed to start on every 8th RS-frame
- We propose to limit the locations in the QR cycle at which alert can be transmitted (similar proposal by Lo_3ch_03a_0319.pdf - Low Datarate EEE)
- When SlowWake is set, alert can only start right after refresh
- This means that the Tw_sys_tx is increased by approximately 35uS in 10G.

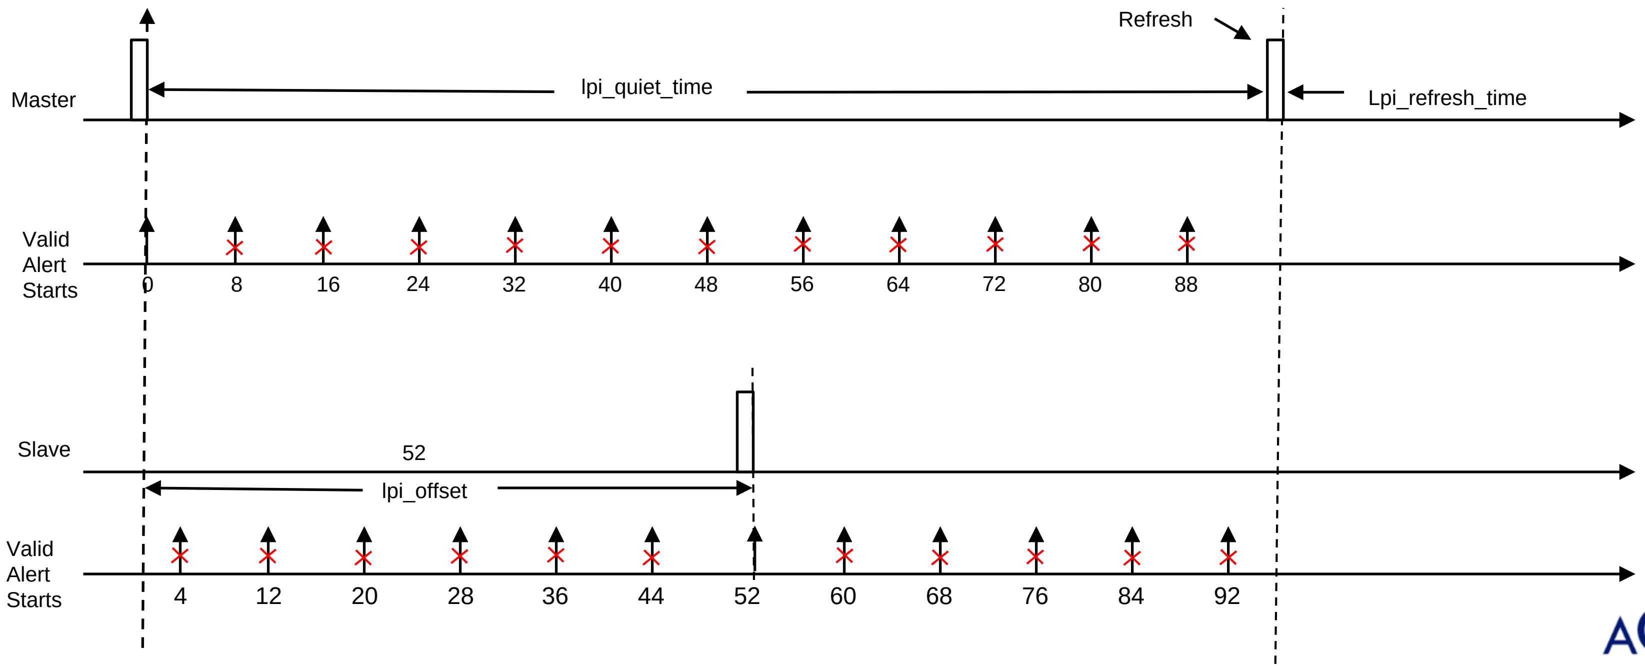

Change #3. Slow Wake

- Limit alert to the immediate RS frame location right after refresh.
- This allows the alert detect to remain powered down through most of the QR cycle.
- As you see in picture below, when XGMII idle starts, PCS has to wait for the next refresh to signal end of QR by inverting the refresh

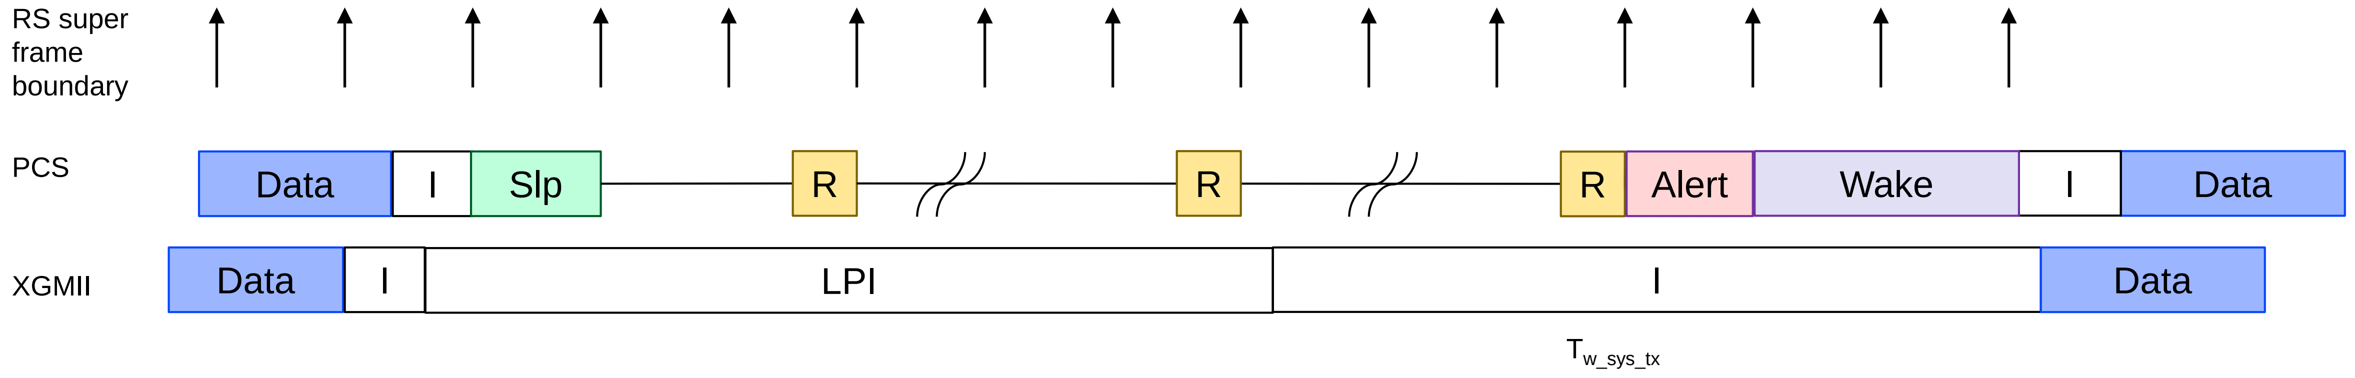

Conclusion

- Highly asymmetric traffic patterns require a more power efficient EEE, this is a crucial requirement for 802.3ch in order to proliferate into all the interfaces available in an automobile.
- We propose that we consider the following changes to the standard:

1. Replace alert signal with a low frequency sine wave
2. Reduce the refresh rate for master receive – propose that we leave this up to implementor on the receive side and bear the small extra power on the transmit for transmitting refresh – no change, implementor does this based on SNR
3. Allow a “Slow Wake” enable, which limits alert location

- Deeper study and simulation needs to be done on #1 to assure robust detection.

Thank you.

AQUANTIA®
ACCELERATING CONNECTIVITY