Thoughts on Energy Efficient Multidrop Systems

G. Zimmerman 6/30/21

CME Consulting

EEE in point-to-point Ethernet

- Locally Controlled
  - Controlled asymmetrically at the transmitter
- Receiver follows transmitter's direction
  - Gets extra time to come back on-line when given ALERT signal
- MII codes to assert LPI, and receive LPI, these are interpreted as 'go to sleep' and 'wake up' by system

- All control is local – derived from LPI Client

Multidrop and EEE are Synergistic

- Multidrop is the right solution when nodes do NOT regularly consume the bandwidth
- Managing energy at the system level is the right solution when activity is not always required
- Both are consistent with long periods of inactivity at each node

How EEE Works

- EEE is controlled locally by the Low Power Idle (LPI) client
  - Allows the MAC & higher layers to go to sleep
  - Physical layer only knows it is in LPI state and is quiet most of the time
- LPI client signals and controls the transition in/out of Low Power Idle state
- Signalling to LPI client via primitives is the KEY to system power savings
- LPI Client operation is unspecified – can be controlled over network

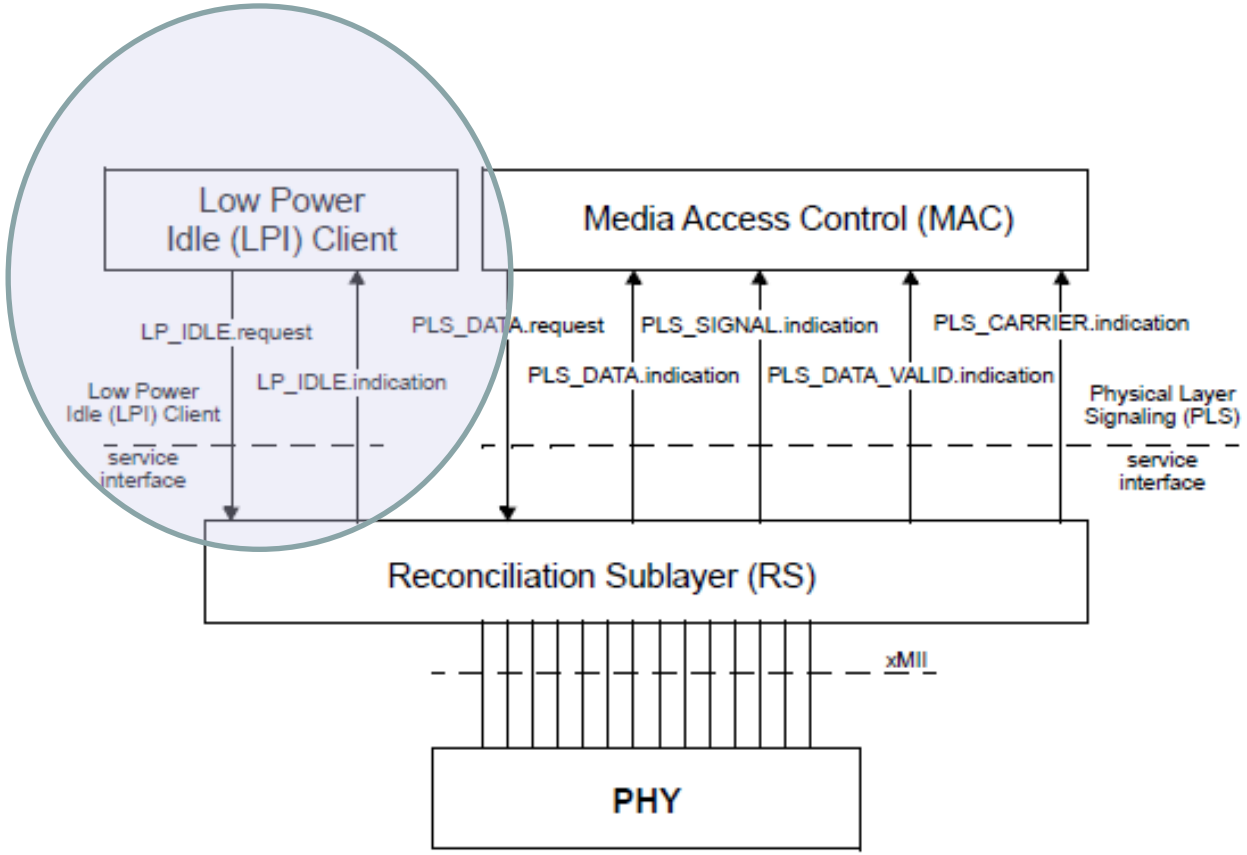

Figure 78–1—LPI Client and RS interlayer service interfaces

EEE Primitives / Mechanisms

- Across xMII:
  - LP_IDLE.request (ASSERT / DEASSERT)
    - A primitive used by the LPI Client to start or stop the signaling of LPI to the link partner.
  - LP_IDLE.indication (ASSERT / DEASSERT)
    - A primitive that is used to indicate to the LPI Client that the link partner has started or stopped signaling LPI.
- Within PHY
  - WAKE, SLEEP, ALERT signals or signal sequences defined on the wire
    - Either PCS symbols or by state diagram / sequences of PCS control codes

What is ALERT?

- ALERT is a special PHY signal used to begin wakeup process
  - Usually a PMA primitive – begins initializing receiver
    - Transmitted when LP_IDLE.request is DEASSERTED
    - Reception can start system wake up by deasserting LP_IDLE.indication
    - Can be out-of-band
  - Can’t be addressed (layer violation)

EEE in multidrop Ethernet

- MII codes can be the same - assert LPI, and receive LPI, interpreted as ‘go to sleep’ and ‘wake up’ by system
  - Requires PCS (4B5B) mapping for these codes, or state diagram for multiple code sequence
  - Allows node to turn completely off
- BUT: receiver could receive from any of many transmitters
  - Scheme to get extra time to come back on line when given ALERT doesn’t work unless all transmitters are asleep
  - Can individual nodes go to sleep & wake, or does it have to be the entire segment?

Putting Nodes to Sleep/Wake Up

- Sleep can be locally controlled per PHY
  - Control via local input or via management entity (network or LLDP)
  - LPI Client asserts LP_IDLE.request
    - 802.3da only has to specify the PHYs response (MII codes already exist)
  - Can be coordinated through data channel – above physical layer
- Local wake-up is easy if you wake everyone else.
  - LPI Client signals (global)ALERT on mixing segment – everybody responds – may be energy inefficient
- Non-global wake-up is hard - requires addressability
  - Ethernet is peer-to-peer – any node must be able to wake up or put any other node to sleep
  - LLDP/Frame (MAC) addressing can’t be used without total PHY wakeup
- Suggests non-Ethernet data, out-of-band control channel, e.g., ALERT!

Encoded-ALERT: Physical Layer Remote Wake Up

- Basic Principle
  - Encode ALERT
    - ALERT is a sequence which can use a simple, robust correlation receiver
  - LPI Client can maintain MAC address client mapping
    - Requires augmentation of LP_IDLE.request DEASSERT with a phy ID
    - Requires maximum node count to be identified
    - Requires global command and individual command

Discussion?