Energy Efficient 100G Ethernet with Modular Low Power Idle

September 2011

Universidad Antonio de Nebrija

100G over copper or backplane

- Intended for datacenter connectivity
- 4x25G using four lanes in each direction
- Receiver must perform
  - Timing recovery
  - Equalization (FFE+DFE)
  - FEXT cancellation (No NEXT due to lane placement??)
  - FEC decoding (FEC blocks span over the four lanes)

100G and Low Power Idle

- Assume transition times are reduced from 10GBASE-T (Tw = 4.5µs and Tw = 2.9µs) to a combined Tw+Ts = 2µs
- Packet transmission time for a 1500 byte packet is only 0.12µs
- This means a large transition overhead
- Consumption during transitions will not be negligible
- For Poisson traffic and 1500 byte frames we get the following energy versus load curve assuming consumption is the same in active mode and transitions. This provides an indication of performance

100G and Low Power Idle

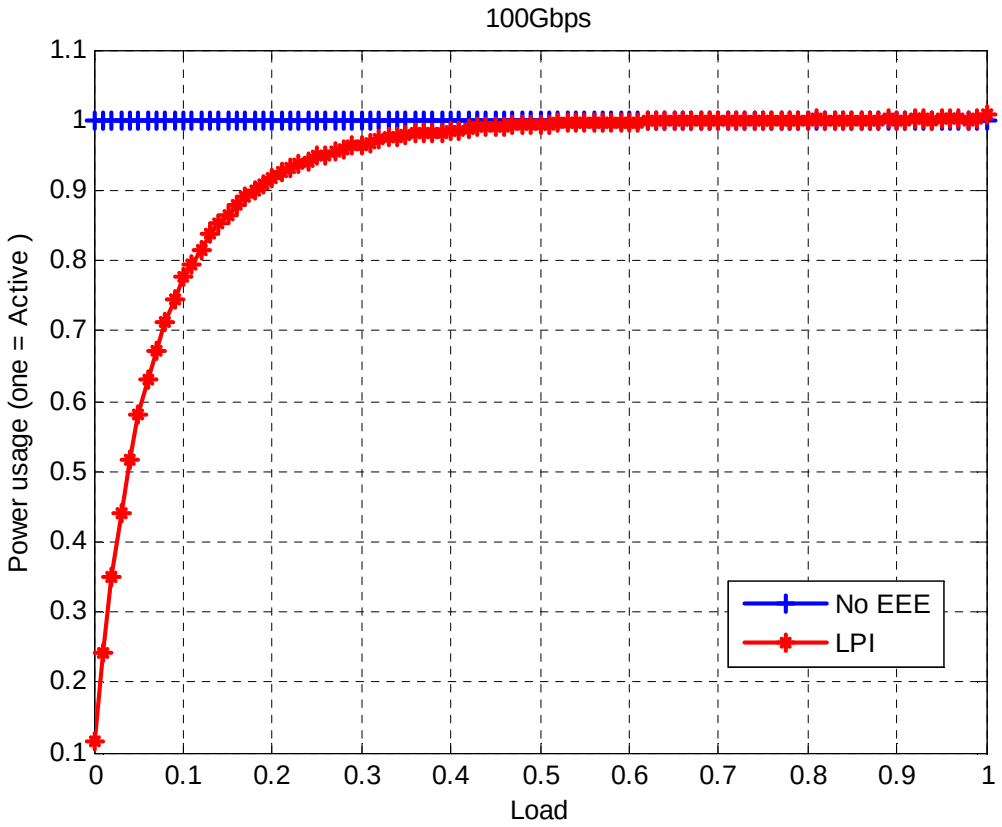

100G and Low Power Idle

- Load in 100G applications (datacenter) will be significant and much larger than for BASE-T PHYs (end-user/offices)
- Therefore the performance of LPI will be poor
- Colaescing can be used to improve savings at the expense of latency but
  - Latency is critical in many of the applications for which 100G is intended
  - Colaescing is outside the scope of the standard

EE 100G with Modular LPI

- Make LPI modular per lane such that it can be graded
  - 1 lane active -> 25G
  - 2 lanes active -> 50G
  - 3 lanes active -> 75G
  - 4 lanes active -> 100G

EE 100G with Modular LPI

Advantages

- More energy savings
- Transition time to add a lane once a lane is active will be much lower (timing is acquired on the first lane)
- FEXT will not consume power (estimates for savings ??) when the rest of the lanes are off
  - Four lanes active -> 12 FEXT cancellers
  - Three lanes active -> 6 FEXT cancellers
  - Two lanes active -> 2 FEXT cancellers

EE 100G with Modular LPI

- Assume transition times to add/remove lanes when at least one is active is half that of activating the first lane.
- Trigger an additional lane activation when queue reaches 8 frames
- Poisson traffic with 1500 byte frames
- This provides an indication of performance that greatly depends on the algorithm to add/remove active lanes to trade energy savings for latency

EE 100G with Modular LPI

- Results

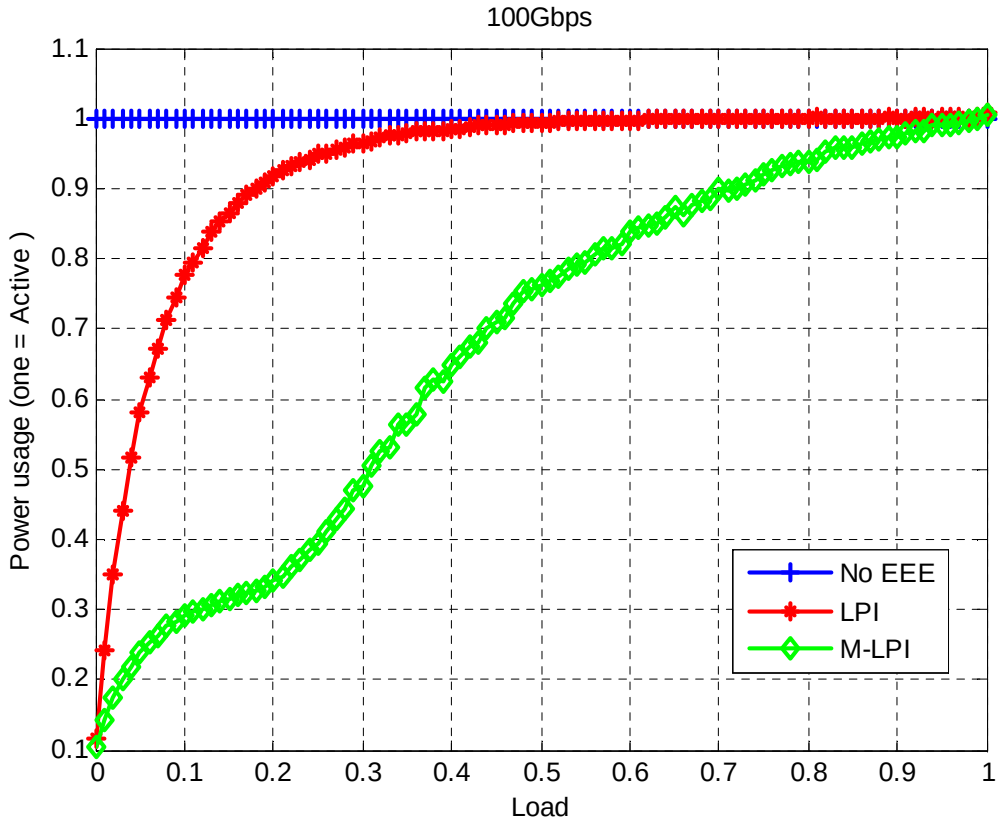

EE 100G with Modular LPI

- Challenges
  - Start/stop lanes without disrupting traffic (FEXT)
  - FEC and frame realignments
  - Algorithms to control the number of lanes (left to vendors to differentiate)

Modular LPI signaling on the CGMII

- CGMII is also structured in lanes (8 lanes with 8 data bits each)
- The modular approach may be extended to the CGMII activating groups of two lanes
  - 2 lanes active -> 25G
  - 4 lanes active -> 50G
  - 6 lanes active -> 75G
  - 8 lanes active -> 100G
- Wake/Sleep Signaling can be done in each lane in a similar way to EEE

Conclusions

- LPI performance may be poor at 100G
- Modular LPI can enable larger energy savings at the expense of a more complex implementation