Thoughts on 802.3 Asymmetric Rates

IEEE 802.3

MultiGigabit Automotive Ethernet PHY Study Group Ad Hoc

George Zimmerman

CME Consulting, Inc.

7 February 2017

Asymmetry is a Systems Issue

- Dynamic Asymmetry of rates:
  - Capable of varying degrees of asymmetry (including symmetric) operation
  - Flexible for varying demands and use cases
  - May or may not go above the PHY layer
  - Variation may be engineered into PHY silicon
- Fixed Asymmetry of rates
  - Fixed at system design time
  - Locked in to traffic assumptions and use cases
  - May be engineered into PHY and MAC silicon
- This is not about “full duplex” or “half duplex”

Full Duplex vs. Half Duplex

- See Clause 4 (the MAC):
  - "In half duplex mode, stations contend for the use of the physical medium, using the CSMA/CD algorithms specified. Bidirectional communication is accomplished by rapid exchange of frames, rather than full duplex operation."
  - "The full duplex mode of operation can be used when all of the following are true:

    1) The physical medium is capable of supporting simultaneous transmission and reception without interference (e.g., 10BASE-T, 10BASE-FL, and 100BASE-TX/FX).

    2) There are exactly two stations on the LAN. This allows the physical medium to be treated as a full duplex point-to-point link between the stations. Since there is no contention for use of a shared medium, the multiple access (i.e., CSMA/CD) algorithms are unnecessary.

    3) Both stations on the LAN are capable of and have been configured to use full duplex operation."
- Has variable throughput and network-level considerations
  - No MACs at rates greater than 1Gbps support half duplex
- Requires planning of link segment propagation delay and packet length
- Does NOT require the rate is the same – it is about contention for use of the media

Examples of Dynamic Asymmetry in 802.3

- EEE low power idle (Clause 78, see 78.1.1)
  - Defined for many PHYs, including BASE-Ts and 1000BASE-T1, but not 100BASE-T1
  - Includes signaling at the MAC interface for systems-level power savings
- PHY-only rate adaptation with MAC-PHY rate matching (Clause 61.2.1)
  - Defined for 10PASS-TS and 2PASS-TL Ethernet in the First Mile PHYs
    - Rates subject to spectral planning, interference, loop conditions, etc.
  - PHY-only rate changes, MAC runs at full rate (100Mb/s in Clause 61)
- Half duplex transmission
  - Only one side transmits at a time, based on traffic
  - Media contention/delay planning gets harder as rate increases
    - This relationship broke at Gigabit Ethernet
  - 2.5G/5G/10G/25G/40G/100G Ethernet are defined for full duplex only
    - (Clauses 44, 105, 125)
    - NO MACs (switches) at these rates support half duplex!

EEE Basics

Where EEE Fits

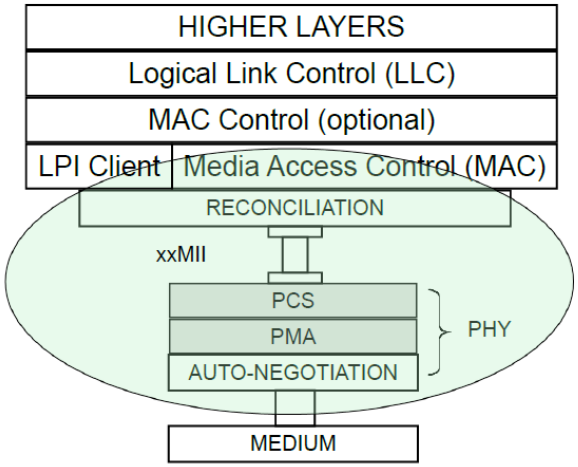

IEEE 802 Plenary March 2011 Page 4

What is Low Power Idle?

- Concept: Transmit data as fast as possible, return to Low-Power Idle
- Saves energy by cycling between Active and Low Power Idle
  - –Power reduced by turning off unused circuits during LPI
  - –Energy use scales with bandwidth utilization

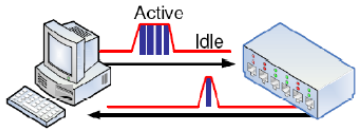

IEEE 802 Plenary March 2011 Page 5

Source: EEE Overview, M. Bennett, http://www.ieee802.org/3/100GCU/public/mar11/bennett_01_0311.pdf

- Like normal IDLE signaling in Ethernet, control of LPI mode is asymmetrical
- Packets may be buffered and burst transmitted over the link in either direction

LPI Overview

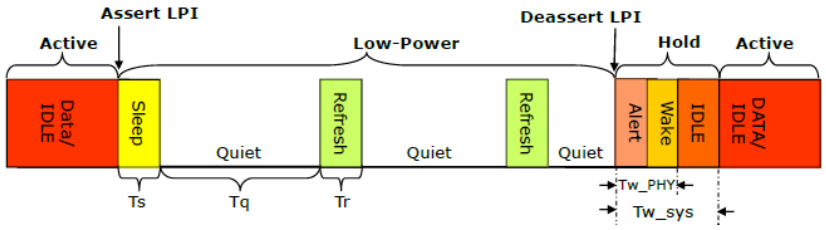

Examples of Fixed Asymmetry in 802.3

- Modes of 10PASS-TS (ADSL)
  - Fixed asymmetry in the phy, dynamic variation
- Various EPON PHYs
  - 10G/1G EPON: 10Gb/s downstream, 1Gb/s upstream, Cl 75
  - 10PASS-XR (802.3br): up to 10Gb/s down, 1.6Gb/s up, Cl 100
  - NG EPON (in development)
- EPON is different from most Ethernet
  - A dual MAC-interface model
  - Extensions to Reconciliation Sublayer, PHYs for unidirectional transport (Cl 66)
  - Point to Multipoint protocols
    - Need to be careful that our solution works in the point-to-point context
  - Also allows dynamic bandwidth allocation on request
    - See http://www.ieee802.org/3/10SPE/public/adhoc/Multiaccess in Ethernet Passive Optical Networks.pdf for detail on EPON P2MP protocols

EPON dual-MAC Model

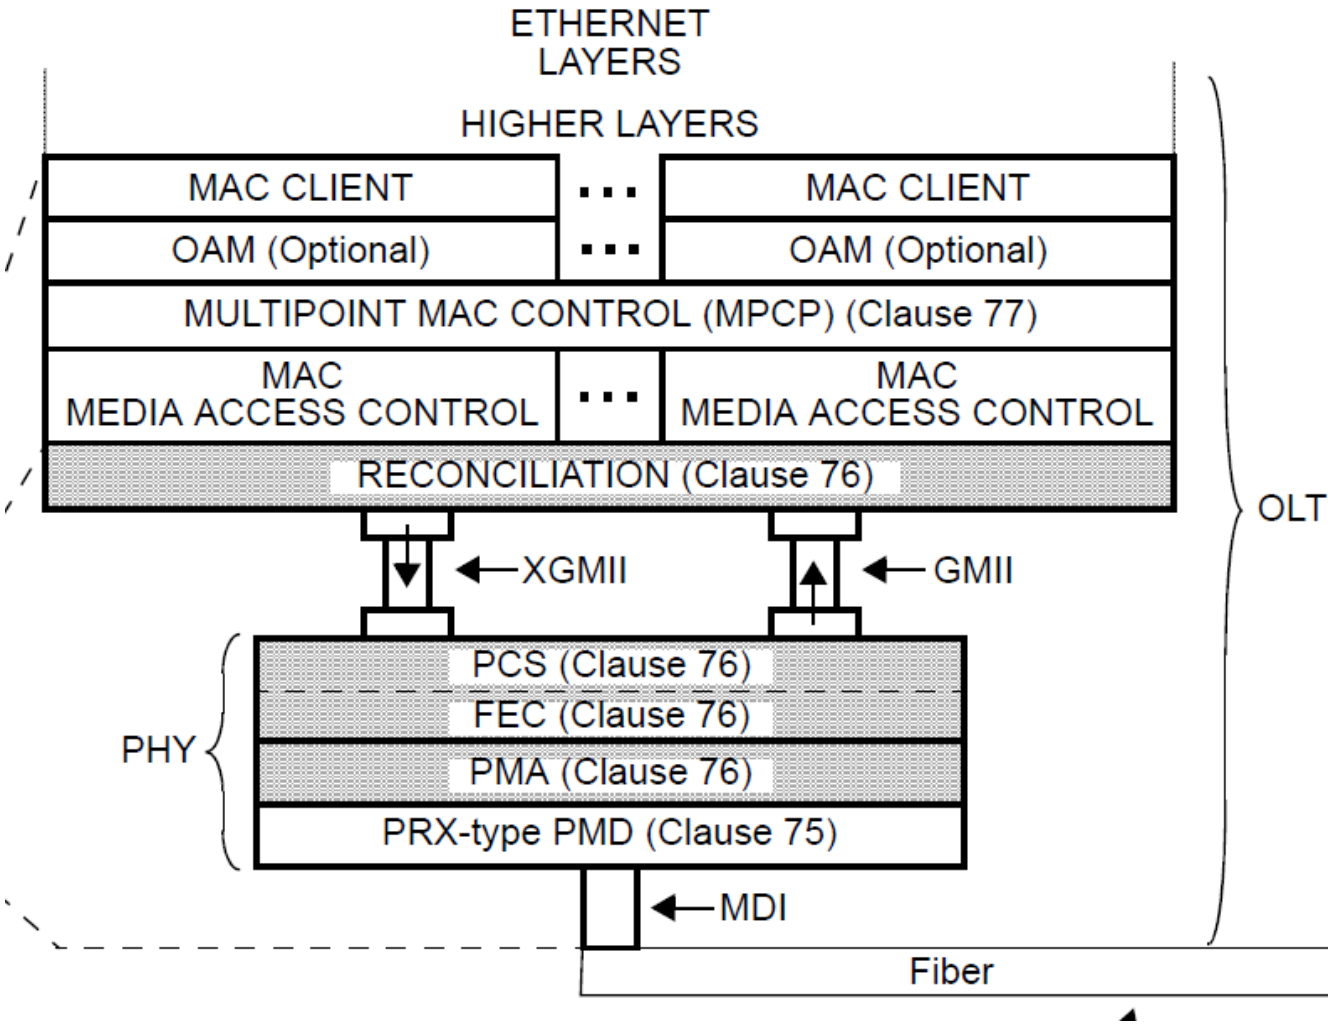

Changing data rate in PHYs (easy)

- Scale the duty cycle of the transmission
  - EEE and half duplex transmission do this dynamically
  - Power savings related to what circuitry can be shut down
- Scale the symbol rate (baud) to use more/less frequency
  - 2.5G/5GBASE-T and 25G/40GBASE-T do this
  - Same PCS & PMA signaling, at different speed
  - Enables savings on the link segment (cabling) parameters (lower bandwidth in the cabling)
  - Saves analog & digital power in the PHY design
  - Saves area through parallelization/bandwidth and architecture
- Scale the modulation/information density using same frequency span
  - For example, 100BASE-TX to 1000BASE-T to 2.5GBASE-T
    - PAM-3 to PAM-5 to PAM-16
    - 100M -> 250M -> 390.625 M bps/pair at 125Mbaud or equivalent (625 Mbps @ 200 Mbaud for 2.5G)
  - Different PCS & PMA signaling, at same or different speeds
  - Minimal savings on link segment (cabling)
  - Saves some power (mostly digital) in PHY design, depends on architecture
  - Saves area (mostly digital) in PHY design

Asymmetry in implementation (1)

- Half duplex
  - Pros:
    - Already defined
    - No need for echo-cancelled receiver signal processing
  - Cons:
    - Not defined in IEEE 802.3 at greater than 1 Gb/s
    - Variable network throughput
- Asymmetric EEE
  - Pros:
    - Already defined
    - Power savings – PHY circuitry shut down when not in use
    - Systems power savings – MAC, switch and other circuitry know to sleep too
    - Flexibility – can burst to full rate when needed
    - Network and PHY proven – can leverage experience and work
    - Could be used with asymmetric transmission, but not defined
  - Cons:
    - Wakeup delay limits savings when traffic gets close to symmetry
    - If low latency wakeup is needed, analog savings may suffer
    - Clock stability may be an issue if fast wakeup is desired after long periods of inactivity

Asymmetry in implementation (2)

- Fixed-asymmetry PHY implementations
  - Pros:
    - Transmit and receive circuitry optimized for rate
      - Power savings in the PHY
      - Potential for silicon area savings at the PHY
        - » Not guaranteed - depends on the relationships of clocks, bandwidths, etc.
  - Cons:
    - Twice as hard to define the PHY (need to design 2 PHYs)
    - Multiple PHY types (can be avoided by dropping savings)
    - MAC level complexity
    - No direct system level savings
    - MAC/system level power savings reduced
      - Circuits are more often active than in the “EEE buffer and burst” mode

Conclusion

- No need for an ‘asymmetric’ objective at this time
  - Unless it effects the MAC interface, but nonstandard rates would be out of scope
- Stick with full duplex operation
  - Don’t break > 1Gbps Ethernet!
- Asymmetric rates can be handled by mechanisms like EEE or PHY modifications, already in the objectives
  - Additional PHY optimization can be done in Task Force
  - Many ways to do this
- Consider systems-level implications

Thank You!