EPoC Delay

Ed Boyd, Broadcom

Overview

- The goal for the EPoC system is a fiber alternative over coax cable.
- To achieve this goal, the system should efficiently carry Ethernet traffic with similar delay and equipment costs as EPON.
- This presentation attempts to breakdown the delay components in an EPON/EPoC system to start discussion on a delay budget.
- This presentation gives some ball park numbers for the delays so there is “give and take” when selecting solutions for pieces of the design.
- This is not a baseline proposal. It is an evaluation.
- Task Force evaluation should include impact to the delay budget.
- TDD options are considered with cost and delay impacts.

Evaluation Boundaries

- The 1G or 10G EPON systems are assumed to be based on the current EPON standards.
- Low cost, high density 1G OLTs are currently deployed widely.
- 10G OLTs are being deployed now and will be deployed in high volume by the time EPoC systems are standardized.
- If a new generation of OLT/ONU chips on new OLT systems are required to achieve a high performing EPoC system, there will be significant impacts to cost and availability.
- EPoC can’t require drastic changes to the Ethernet layering diagram or standards. It should be just a PHY.

Service Group Size

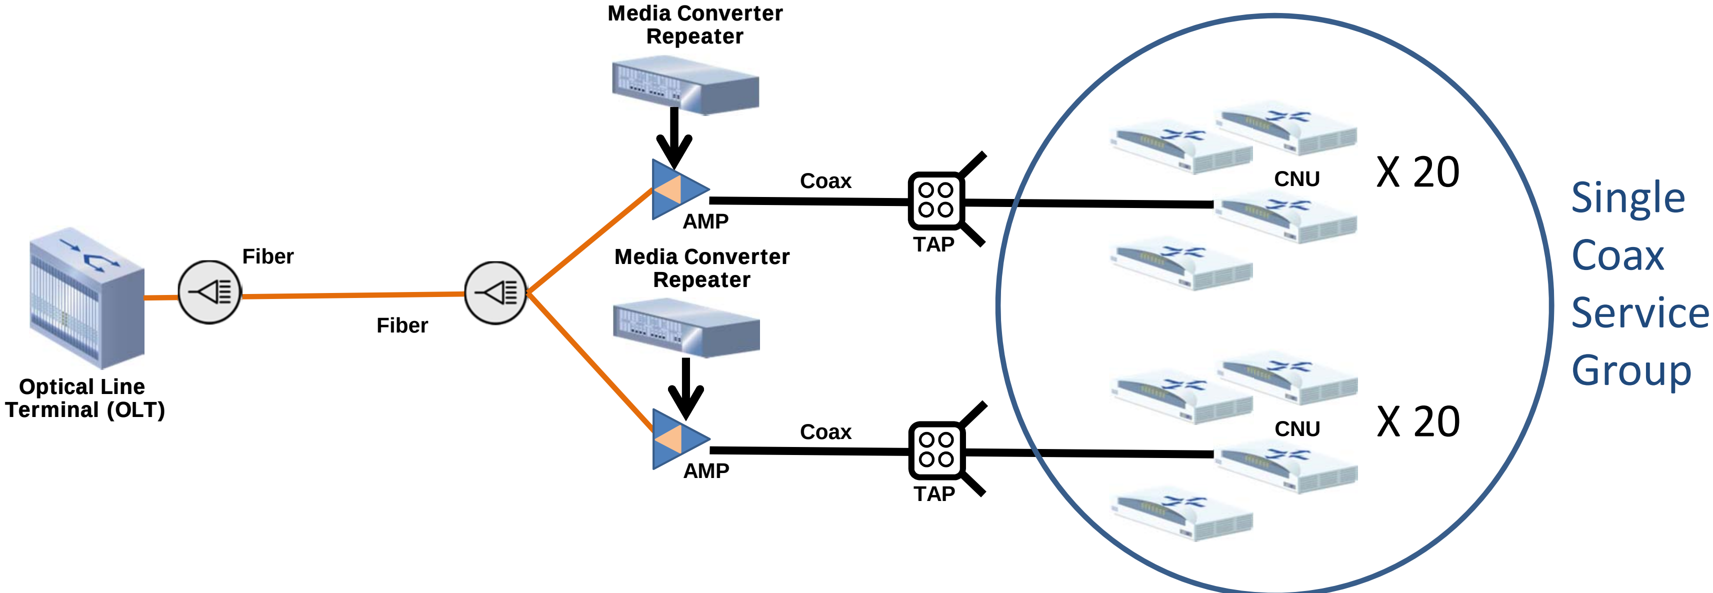

- Amplified Coax or Multiple Repeaters on Node + 0 can be used to create large service groups.
- Service groups of 32/64 (EPON size) and 256 (operator requested) should be considered.
- Smaller serving groups require more OLT ports, more wavelengths, and deeper fiber so larger groups must be evaluated.
- Bridges allow for fewer OLT ports but a more complicated and expensive outside plant device with longer delays so repeaters will be the focus. If a repeater meets the delay budget, a bridge can be considered later. (See May IEEE presentation)

Sources of EPoC Delay

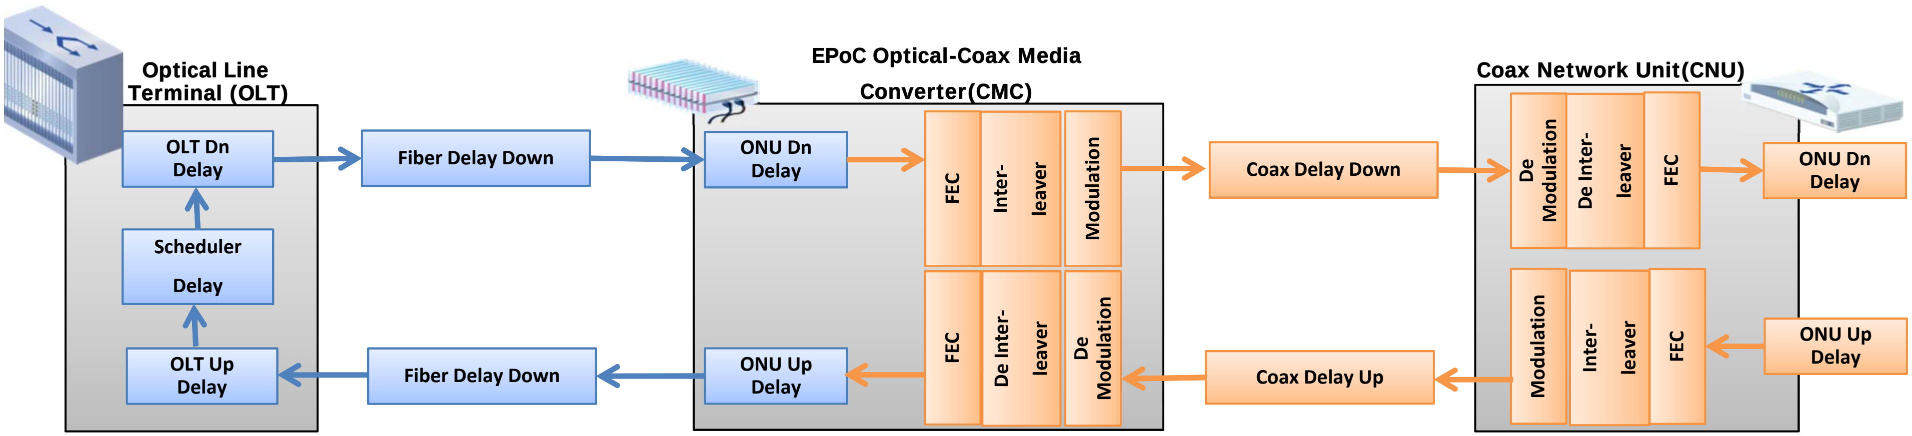

- EPON System Delay (Fiber Only)
  - Fiber Propagation Delay (100us for 20Km)
  - PHY/MAC fixed delay (roughly 25us)
  - 125us of fixed delay in each direction.
- Scheduler Delay (Network size, service, and implementation specific)
  - Number of active stations in service group times maximum longest burst size.
  - If scheduler has any processing cycles, they would be additional delay.
- Polling Delay
  - Solicited Granting requires a grant looking for a non-zero queue status.
  - Shorter polling cycles have higher bandwidth overhead but decrease delay.
- EPoC PHY Delay
  - Symbol blocks, FEC blocks, Interleaving, duplexing delays (if TDD)

EPoC Continuous Downstream

- FEC
  - Longer code words for better efficiency can be used.
  - No need for shortened code words.
  - 90% LDPC could be used for starting point for analysis.
  - 4.5K bits (4.5us @ 1Gbps) code word size.

- Interleaver
  - Continuous downstream can work with convolutional.
  - Convolutional is half the delay and memory of a block interleaver for the same burst error protection. No efficiency impact.
  - DOCSIS like J.83 Convolutional interleaver will be used as a starting point for analysis.

- Modulation Blocks/Symbols
  - Long symbols can be used to reduce cyclic prefix overhead.
  - For now, go with a simple 32us symbol and 1us CP overhead (96.87%)

- **Total PHY Layer Downstream Efficiency = (90%*96.87%) = 87.2%**

EPoC Burst Upstream

- FEC
  - Shorter code words for better efficiency can be used.
  - Option of short or long code words is challenging without short/long indication added into REPORT and GATE MPCP frames. More study needed.
  - Shortened Blocks are needed to improve efficiency
  - 80% Efficient FEC as a starting point.

- Interleaver
  - Burst upstream can't have convolutional since packets must finish before burst boundaries.
  - Block interleaver is needed so no packets span the burst boundaries.
  - Double the delay and size of convolutional but no efficiency impact.

- Modulation Blocks/Symbols
  - Short symbols and blocks for lower delay.
  - 16 symbol blocks with 1 symbol for burst preamble (93.75%)
  - 1us CP overhead (93.75%)

- **Total PHY Layer Upstream Efficiency = 70.3125% (NOTE: Need effect of shortened FEC)**

EPON+EPoC Fixed (PHY) Delay Budget

- Upstream
  - FEC/Interleaver
    - 10us burst protection
    - 256us TX + 256us Rx
  - 16us symbols with 16 symbol blocks (256us)
    - Overlaps with interleaving (2x16us symbol delay).
  - 100us propagation delay (sum of coax + fiber)
  - 25us of MAC up delay
  - 25us of ONU up delay
  - Total 512us+100us+32us+50us = 694us.

- Downstream
  - FEC/Interleaver
    - 10us burst protection
    - 128us TX + 128us RX
    - 5us FEC
  - 32us symbols
    - 2x32us symbol delay
  - 100us propagation delay (sum of coax + fiber)
  - Total 261us+64us+100us = 425us.

Total Bidirectional Fixed Delay = 694us + 425us = 1119us

EPoC Fixed (PHY) Delay Implications

- Up+Down Compared to Fiber Only
  - 250us vs 1.1ms
- REPORT Frame Implication
  - MPCP REPORT maximum queue size is 2^16*16ns or 1.05ms.
  - If delay exceeds 1.05ms, full line rate can't be achieved.
  - Example:
    - 1.1ms: 1.05ms/1.1ms*1 Gbps = 954 Mbps max BW per CNU.
    - 2ms=525 Mbps max BW per CNU.
- Solicited Fixed Upstream Delay Implication
  - 2xUp+1xDn=2x694us+438us=1.8ms (vs 725us Fiber)
  - 134K Bytes more buffering on CNU - OR -
  - 33% increase in polling BW for MEF23H or
  - 12% increase in polling BW for MEF23M/L
  - See NCTA paper for more details.

1 millisecond is the practical limit for EPoC PHY Delay without significant ONU buffer increases or low performance

EPoC Small Burst Overhead

- Maximum Number of transmitters adds penalty to small upstream bursts.

- Simple Example
  - 256us Symbol Block at 1 Gbps = 32K Bytes
  - 32 Transmitter Limit = 1K Byte Min Burst Size
  - 64 Transmitter Limit = 512 Byte Min Burst Size
  - 276 Byte Min Burst Size in EPON (From May IEEE presentation)
  - EPoC could take double the polling bandwidth or total bandwidth for short

- Based on 64 station system polling and 50% small burst distribution: EPoC efficiency could be 90% of EPON due to small burst. More analysis needed.

This is a significant issue for EPoC and should be explored in the task force.

EPoC FDD Summary

- Delay is at the limit and should be reduced if possible.
  - If we reduce delay, longer symbols or FEC codes with better efficiency are possible.
  - Delay can not increase beyond 1ms without new MPCP frames and DBA interface to 802.1.

- Downstream efficiency at 87.2% is a good start.

- Upstream efficiency at 70% without small burst penalty and 63% with small burst penalty is low. Future contributions should be focused in this area.

EPoC FDD can work but upstream performance is key item for study.

EPoC TDD Options

Ed Boyd, Broadcom

Ethernet Layer Diagram

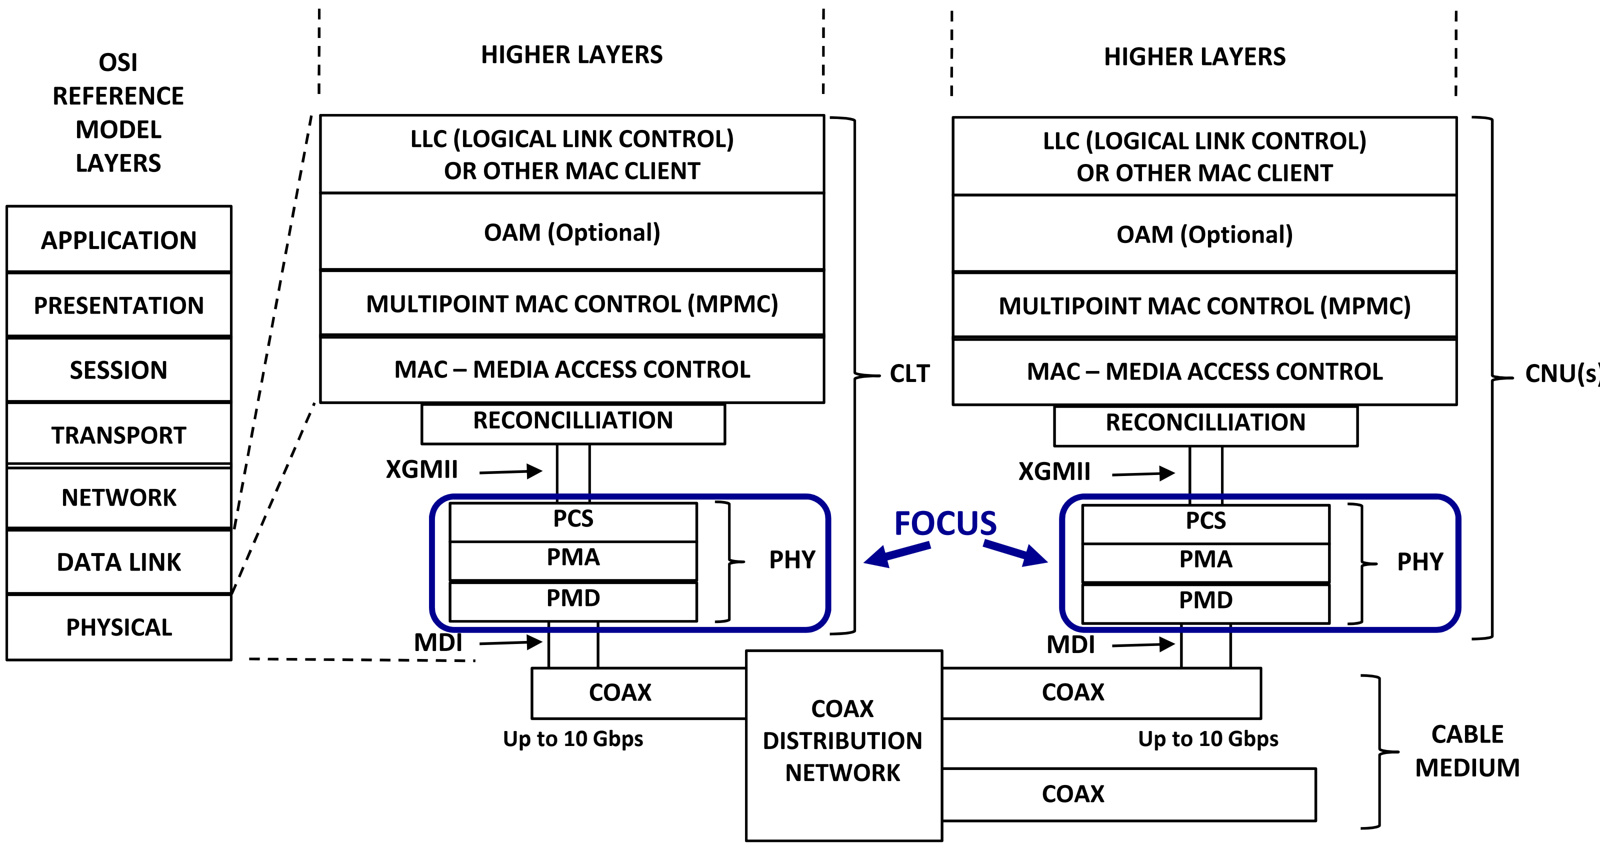

CLT – COAX LINE TERMINAL
CNU – COAX NETWORK UNIT
MDI – MEDIUM DEPENDENT INTERFACE
OAM – OPERATIONS, ADMINISTRATION, & MAINTENANCE

PCS – PHYSICAL CODING SUBLAYER
PHY – PHYSICAL LAYER DEVICE
PMA – PHYSICAL MEDIUM ATTACHMENT
PMD – PHYSICAL MEDIUM DEPENDENT
XGMII – GIGABIT MEDIA INDEPENDENT INTERFACE

Layer Diagram Requirements for Burst Mode

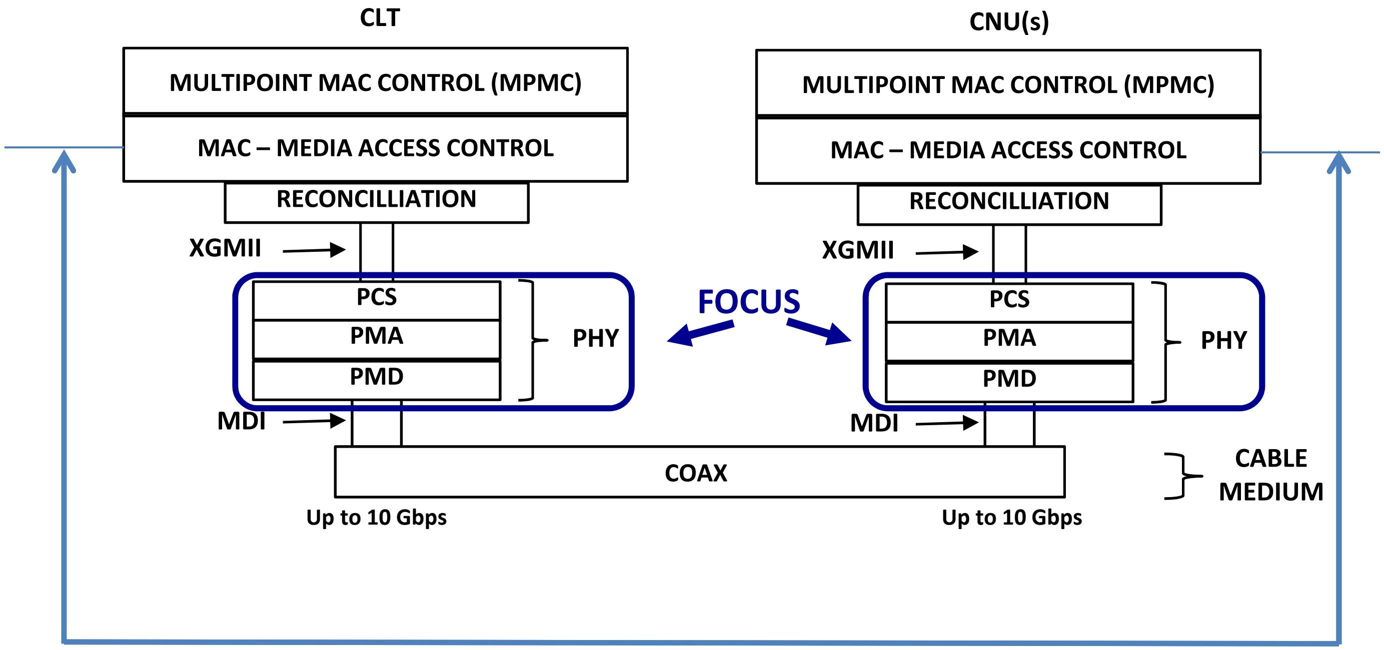

Fixed Delay for Full Packets without Fragmentation

- Packets from MAC Control in either direction must have a fixed delay (8 TQ jitter) for MPCP discovery to accurately time the loop and align slots.
- 802.3ah (EPON) and 802.3bf are two examples of broken standards if excessive jitter is introduced.
- Packets can't span bursts. There is no ability for fragmentation to MAC and jitter would be over the limit.

TDD: TX/RX Logic Sharing

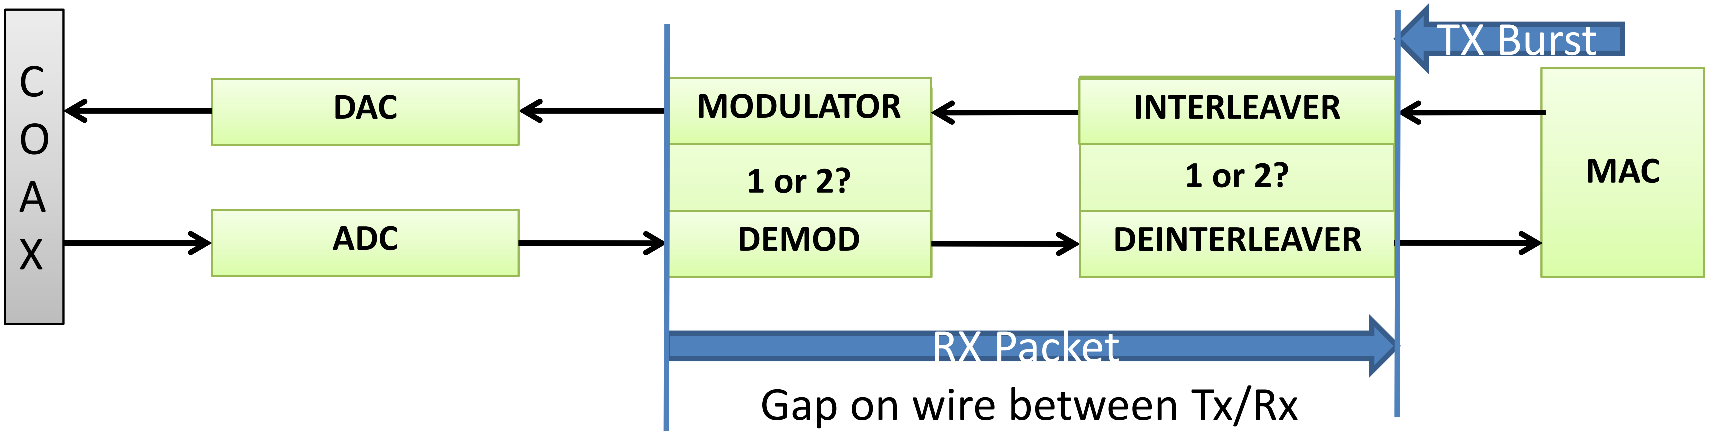

- **TDD needs wider data path for CNU transmit than FDD.**
  - **1Gbps Up/5 Gbps Down would have upstream of 1Gbps for FDD and 6 Gbps for TDD.**
  - **To avoid the cost of 2x6Gbps channel on TDD CNU. Logic should be shared for TX and RX.**
- **AFE**
  - **Amplifier, ADC, DAC can't be shared. TDD requires more expensive front end.**
- **Modulator/Demodulator**
  - **Receive must finish processing before transmit starts to process.**
  - **2 symbol gap between upstream and downstream needed.**
- **Interleaver/Deinterleaver**
  - **Block Interleaver so memory is empty between bursts. Not Convolutional.**
  - **Dual load and unload could be possible to avoid delay between TX and RX.**

TDD: MAC Control or PHY Control

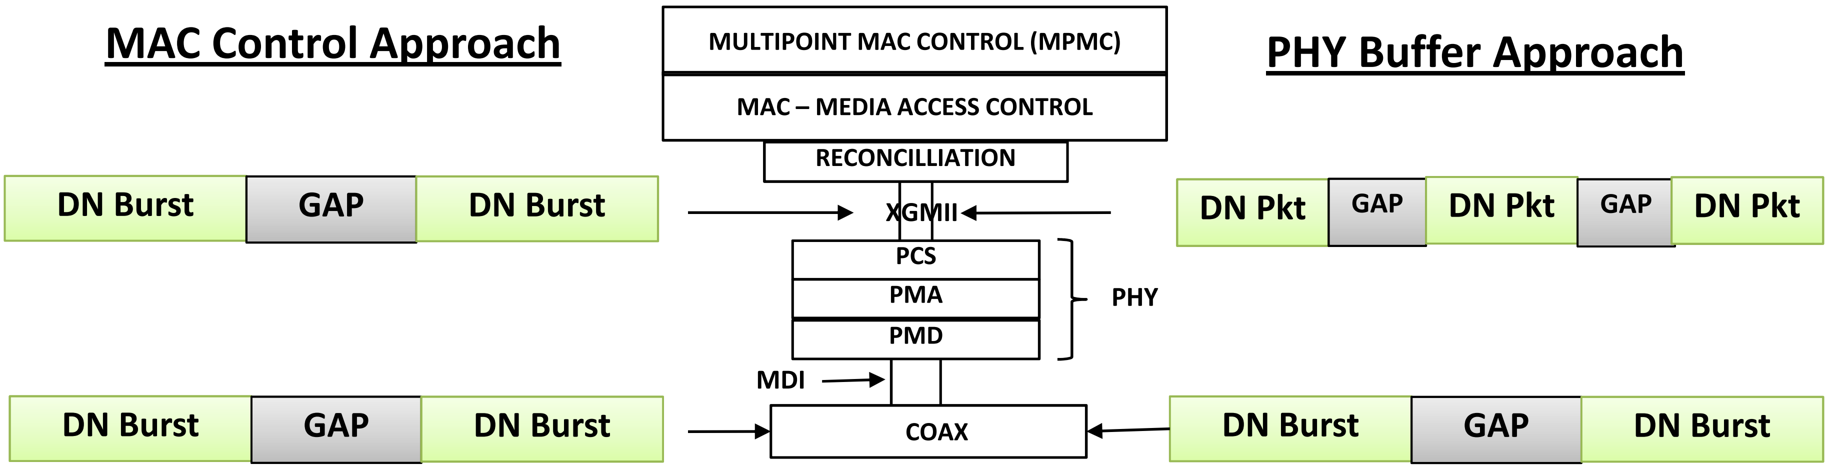

- MAC Control stops downstream packets to create gap on coax for upstream.
- MAC Control schedules upstream bursts into downstream GAP.
- MPCP timestamps and lengths match at XGMII and PHY output.
- Data Detector in current EPON upstream PHY turns ON and OFF downstream PHY.
- Packets can’t span the gap without violating fixed delay rule.

- Additional GAP between packets is added together to make large GAP in PHY.
- PHY has buffer to hold burst data and stream out.
- MAC control has no awareness of GAP.
- Schedules continuous upstream and PHY shift packets to GAPs.
- MPCP timestamps are not correct on wire in upstream and downstream on cable.

TDD: MAC Control Fixed or Flexible

- MAC Control Fixed
  - CLT MAC Control block could have a fixed size for the upstream.
  - CNU would discover the GAP size during registration and it would remain relatively fixed.
  - Burst sizes in the upstream must be split to fit into slot. (frame alignment lost with REPORT frame)
  - Downstream frames must avoid GAP since fragmentation isn't allowed
- MAC Control Flexible
  - CLT MAC Control block would schedule variable size upstream and downstream size.
  - New Downstream MPCP GATE frame needed to announce downstream size to CNUs.
  - Additional downstream delay for MAC Control to get frame size ahead of Downstream MPCP GATE frame.
  - Flexible size could allow for variable upstream & downstream split.
  - Avoids frame alignment issues.

TDD Choices

- PHY Buffer
  - PRO: Same MAC interface as FDD
  - PRO: Full 10Gbps possible in both directions
  - CON: Different PHY than FDD upstream. More expensive. Violates single PHY objective.
  - CON: Large buffer required in PHY: 100's of kilobytes.
  - CON: Timestamps are aligned MPCP upstream slots and downstream.
  - CON: Upstream frames and bursts are fragmented into multiple bursts.

- MAC Control Fixed Size
  - PRO: Lower cost PHY, same as FDD, no buffer added to PHY, MPCP time is correct.
  - CON: Downstream Frame Boundary Alignment will lower efficiency
  - CON: Upstream Burst Alignment will lower efficiency
  - CON: XGMII limits bandwidth to 10Gbps for sum of upstream and downstream

- MAC Control Flexible Size
  - PRO: Lower cost PHY, same as FDD, no buffer added to PHY, MPCP time is correct.
  - PRO: Dynamic Change of upstream and downstream bandwidth.
  - PRO: No Frame alignment issues.
  - CON: New Downstream MPCP frame to announce the burst size
  - CON: Additional delay of downstream data to determine MPCP frame
  - CON: XGMII limits bandwidth to 10Gbps for sum of upstream and downstream

TDD: Performance Metrics

- If frame alignment to burst boundary isn't addressed, a 2K Byte waste is possible on any upstream or downstream slot end. Assume 1K Byte error on average – 8us @ 1Gbps.
- Sharing transmit and receive logic adds 2 symbols (16us each) of GAP plus propagation delay difference. Estimate GAP at 40us.
- Total direction switch burst overhead is estimated at 48us
- 480us Burst Size with 48us overhead (10% turnaround penalty)
- Total Cycle of 528us for each direction.
- Block Interleaver is used for upstream and downstream.

EPoC TDD Performance

- Downstream
  - Efficiency
    - 10% penalty for direction change so 78.5% overall
  - Delay
    - Block Interleaver adds 250us
    - Disruption for change of Direction adds 528us.
    - 425us + 250us + 528us = 1.203ms

- Upstream
  - Efficiency
    - 10% penalty for direction change overhead
    - 5% additional penalty based on small burst penalty and doubling of maximum packet size. (This is very rough based on 1Gbps bidirectional channel)
    - 60.1% upstream overall
  - Delay
    - Disruption for change of Direction adds 528us.
    - 694us + 528us = 1.222ms

Total Fixed Delay = 1.2ms + 1.2ms = 2.4ms

EPoC TDD Delay Implications

- Up+Down Compared to Fiber Only
  - 250us vs 2.4ms

- REPORT Frame Implication
  - MPCP REPORT maximum queue size is 2^16*16ns or 1.05ms.
  - 1.05ms/2.4ms * 1 Gbps = 437 Mbps

- Solicited Fixed Upstream Delay Implication
  - 2xUp+1xDn=2x1.2ms+1.2ms=3.6ms (vs 725us Fiber)
  - 450K Bytes more buffering on CNU –OR-
  - 36% increase in polling BW for MEF23M/L
  - No increase in MEF23H polling can recover 3ms

TDD requires a new MPCP REPORT frame, new 802.1 interface, and larger CNU buffers.

TDD Conclusions

- Do it right or not at all.
  - Flexible MAC Control provides the best solution in performance, cost, and standards compatibility.
  - Significant additions are needed in MAC Control to make EPOC TDD perform well compared to other TDD solutions.
- Two Projects?
  - Same PHY could be used for TDD and FDD but MAC Control must be different.
  - TDD EPoC won't match fiber performance.
  - A study should be focused on comparing TDD EPoC with proper MAC Control to other TDD solutions.
  - Applications could be expanded beyond cable access.
  - A second 802.3 project for EPoC MAC Control changes for TDD should be considered if market requires it.

Thank You!