MMP in EPoC
how hard can it be ?

Marek Hajduczenia, PhD

marek.hajduczenia@zte.pt

MMP Proposal … in text

- There would be a maximum of 4 modulation rates in total, and a maximum of 2 modulation rates used by each CNU, which would not change dynamically. Each CNU would use one lowest common denominator (LDC) MCS for broadcast and multicast, and may use an additional MCS for unicast traffic if the transmission media (i.e., SNR) allowed it.
- For example, one CNU may have one LLID running the LCD at 256 QAM for broadcast and multicast and one LLID running at 1K QAM for unicast, while another CNU might be running the same LCD LLID at 256 QAM but the unicast LLID at 4K QAM, and a third CNU would run the LDC LLID and the unicast LLID at 256 QAM.

This text was retrieved from email from Jorge and edited for context only

… and in a picture

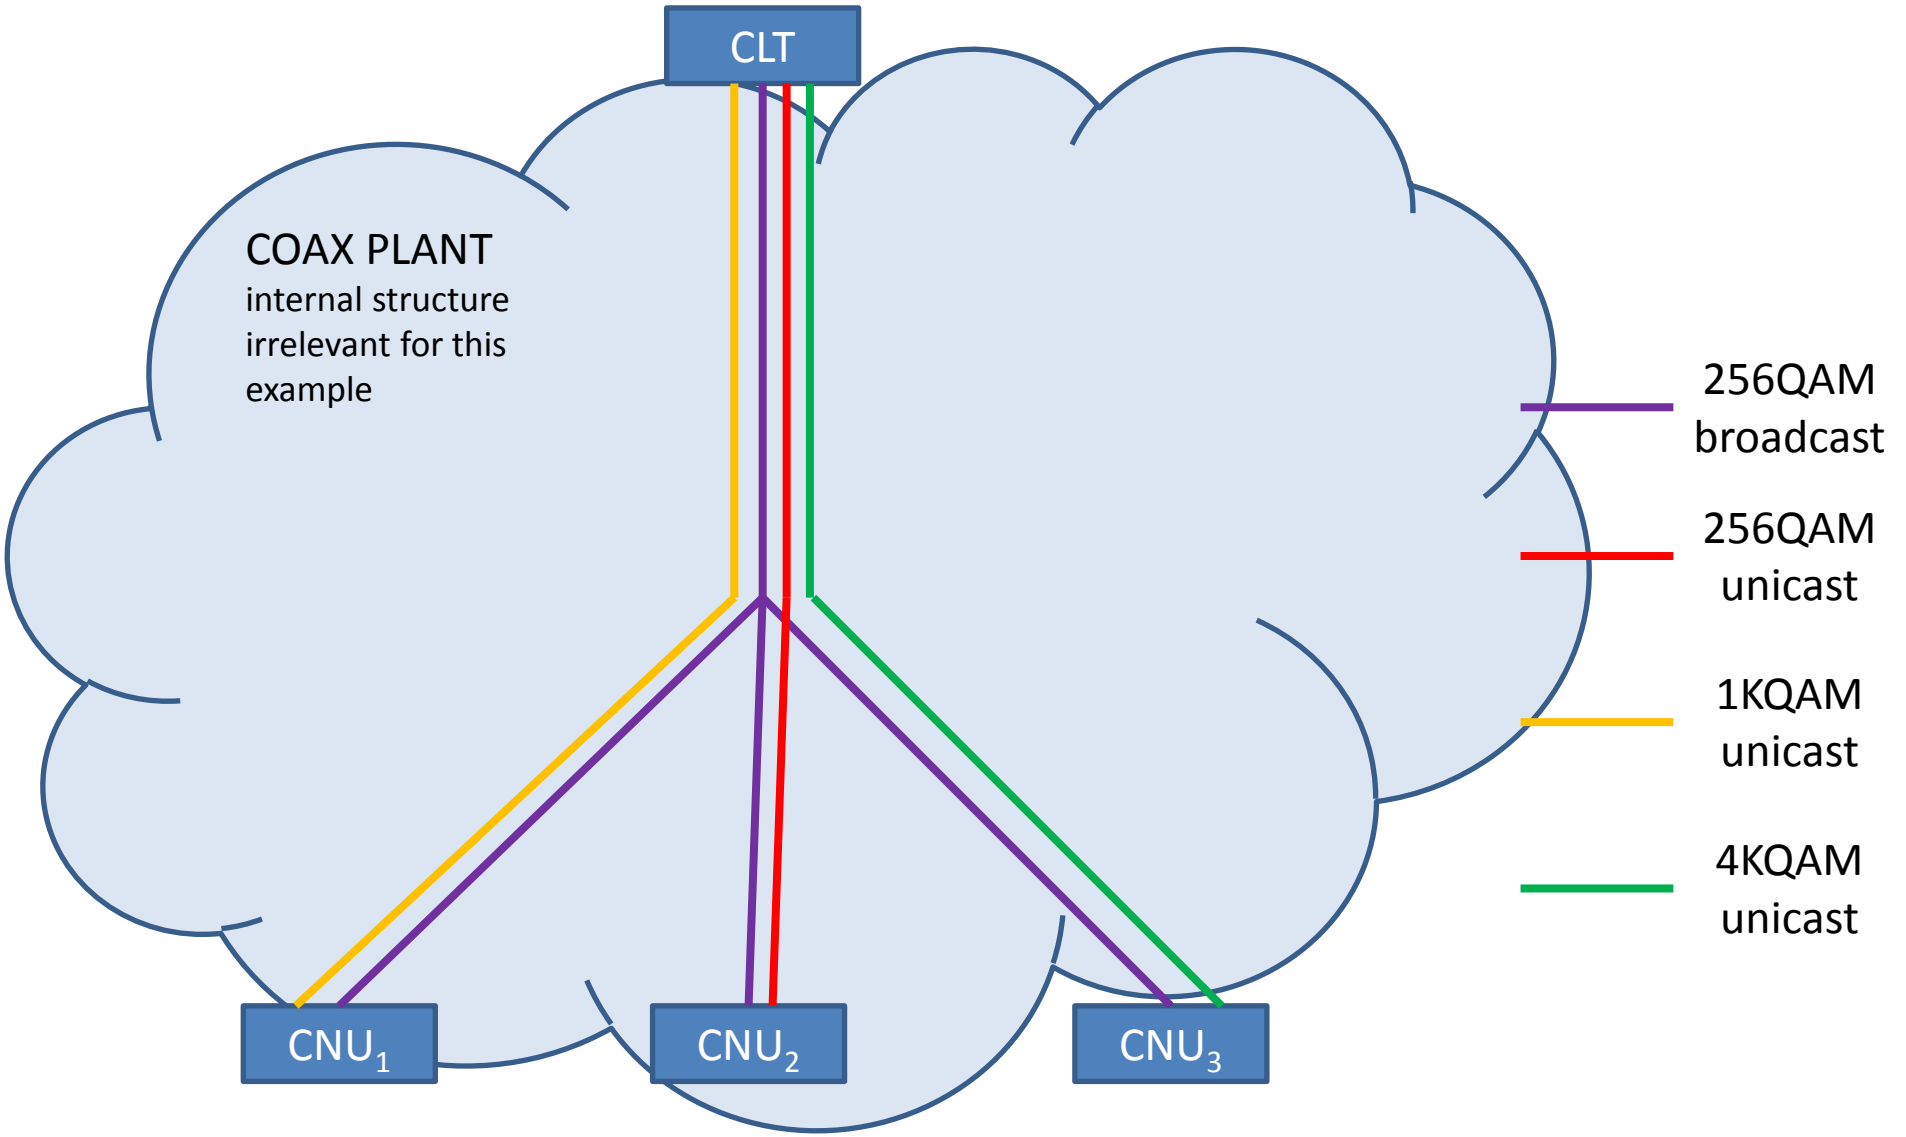

Reference slide

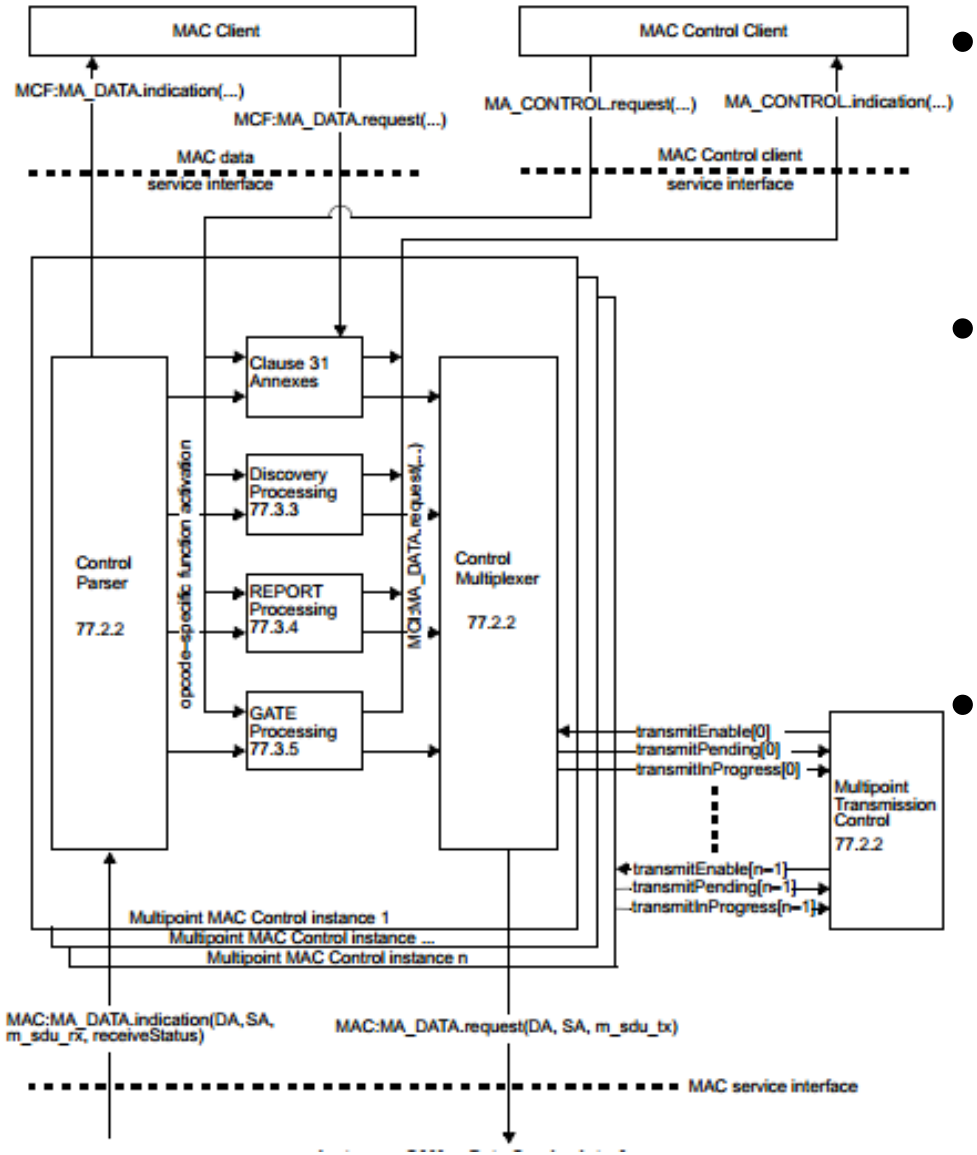

- See Figure 77-3: Multipoint MAC Control functional block diagram
- Similar architecture will be likely be used in EPoC (see law_01a_1112.pdf for more details)
- A single MAC instance transmits downstream at any time (function of Multipoint Transmission Control block)

At the logical layer (I)

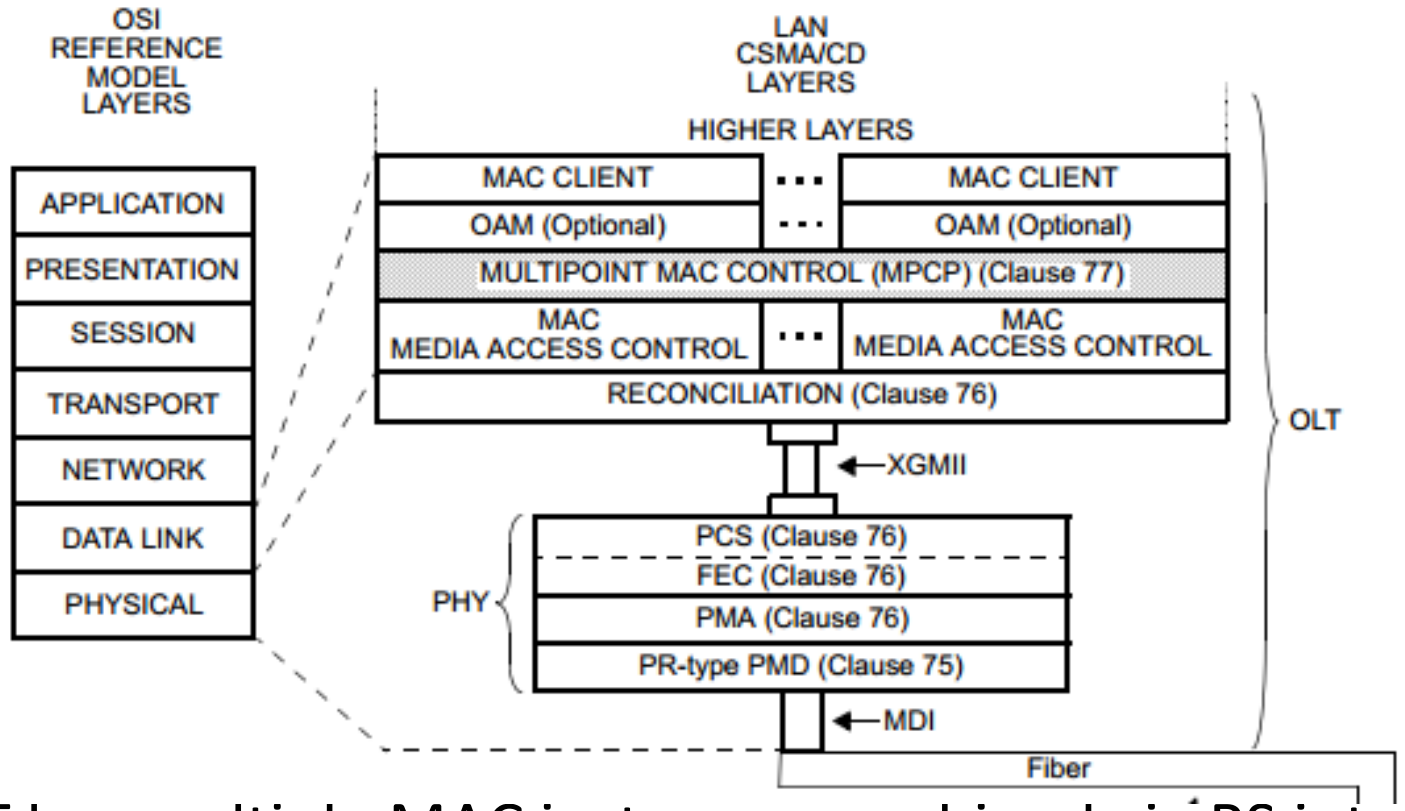

- OLT has multiple MAC instances combined via RS into a single XGMII and then sharing the same PHY
- Each MAC may operate at different effective data rate, but their line rates remain the same (10 Gb/s)

At the logical layer (II)

- Each OLT MAC is associated with a dedicated OLT Control Multiplexer (Figure 77-13 for reference) instance
- Each OLT Control Multiplexer instance operates at its own effective rate, depending on PHY rate, modulation profile, selected FEC, etc. How these parameters get configured is TBD at this time.
- Effectively, each MAC may transmit different number of IDLEs per data symbol to accommodate for de-rating at PHY layer.

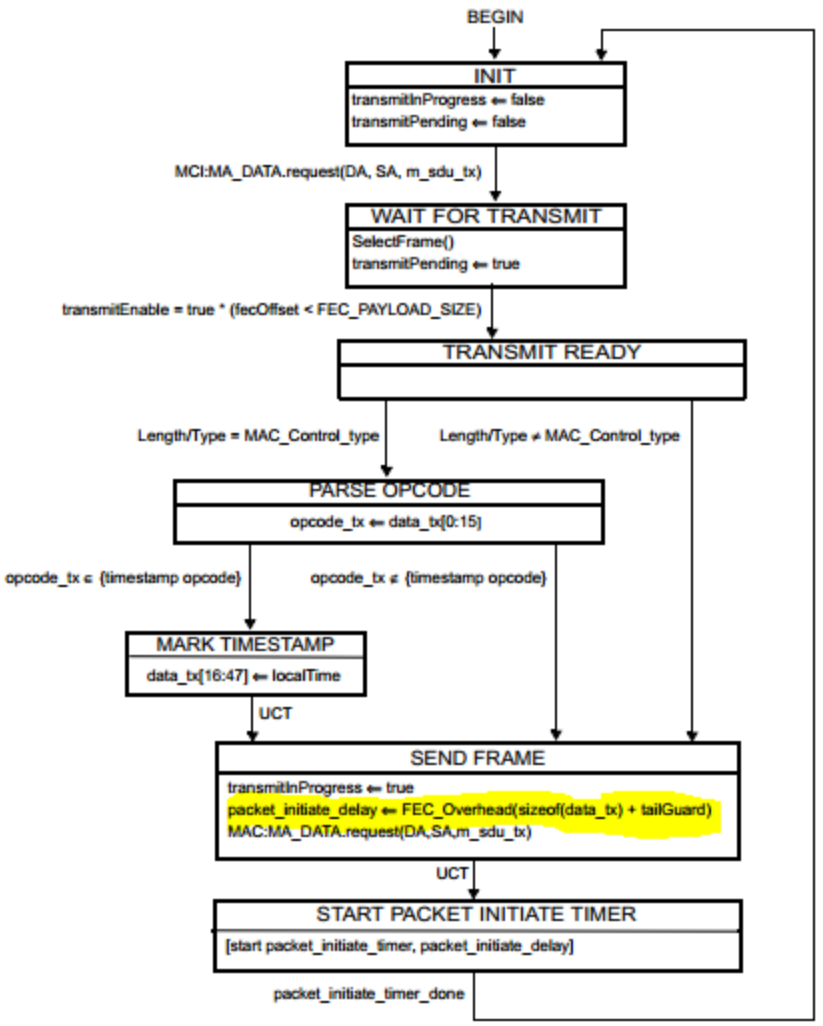

In a picture

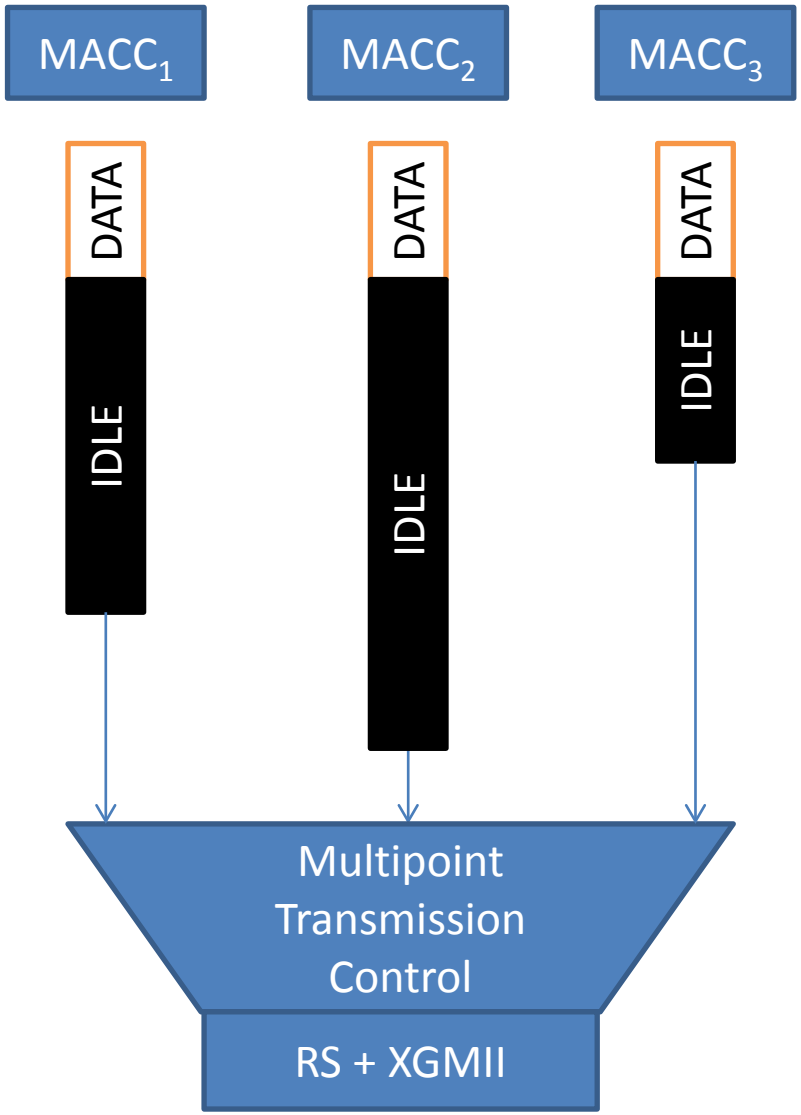

- $MACC_2$ has the lowest effective PHY rate and largest number of extra IDLEs need to be transmitted per data symbol to de-rate MAC
- Multipoint Transmission Control needs to figure out then which of the MAC instances can at the given time transmit data across RS and XGMII

MMP in PCS

- In EPON, PCS is designed around a single data stream transmitted between XGMII and PHY
- Data processing flow has been presented already many times. Figure 76-7 is shown for reference only (downstream).
- PCS has many moving pieces, inter-related to each other and working in tandem off the same data stream source (XGMII)

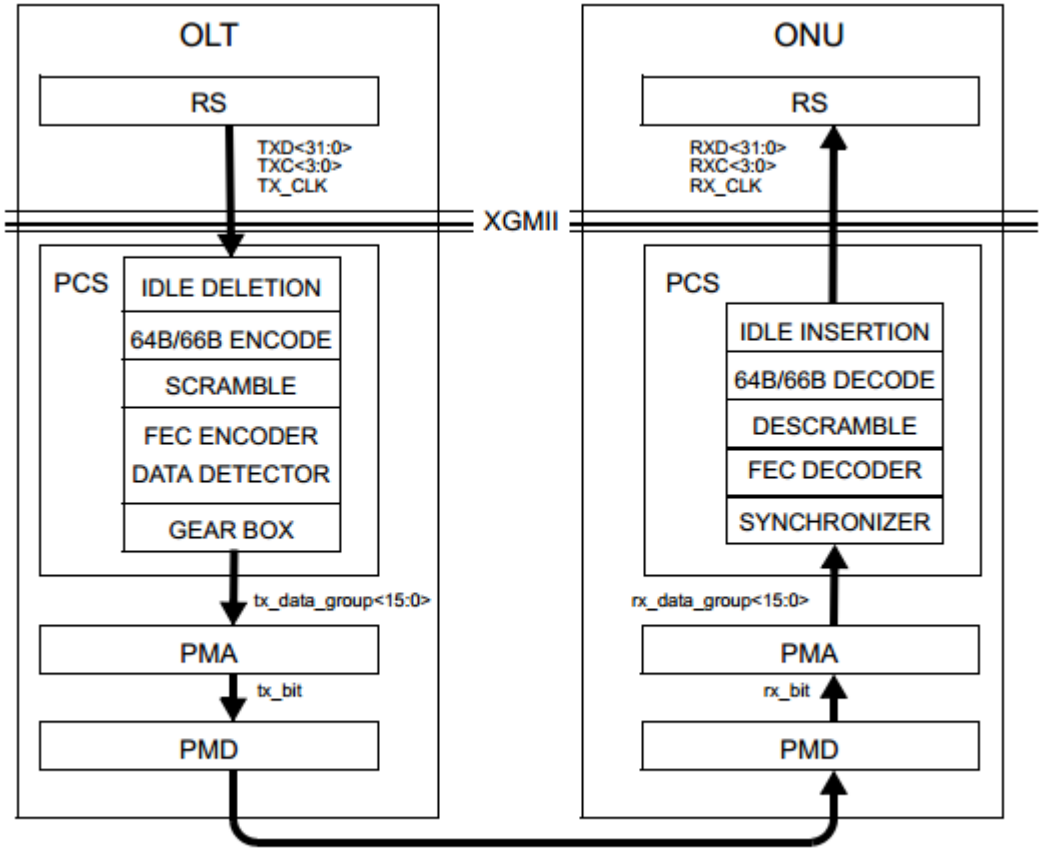

Figure 76–7—PCS functional block diagram, downstream path

PCS facts and figures

- A PCS does not have a notion of a frame – frame has been serialized, transmitted by XGMII and mixed with signaling by the time it reaches PCS
- All data within (10G-EPON) PCS is processed in vectors of 64 bits (prior to 64b/66b encoding) or 66 bits (after encoding)
- Typical 10G-EPON PCS implementations operate on XGMII clock (312.5 MHz). Recall XGMII operates at 312.5 Mtransfer/s, pushing 32 bit of data, 4 bits of control and 1 bit of clock towards PCS
- In 10G-EPON, data stream within PCS may be bursty (i.e., have gaps) after the Idle Deletion function. Gaps are filled in by FEC parity data prior to transmission over the medium.

MMP in PCS

- PCS does not have notion of frames. Data is processed in vectors (data sequences).
- To properly apply MMP, we need:
  - Idle Deletion function to be able to figure out how many excess IDLEs to delete after each data sequence. Each data sequence may belong to different profile.
  - Direct the resulting data sequence with IPG to proper FEC instance (assuming each profile gets its optimized FEC) for FEC encoding
  - Combine resulting data streams into a single MDI interface for presentation towards the medium
  - Assure that switching between individual profiles is done in sync with data transfer, i.e., there are no additional extra gaps in data stream due to profile switching
- Individual items will be examined next

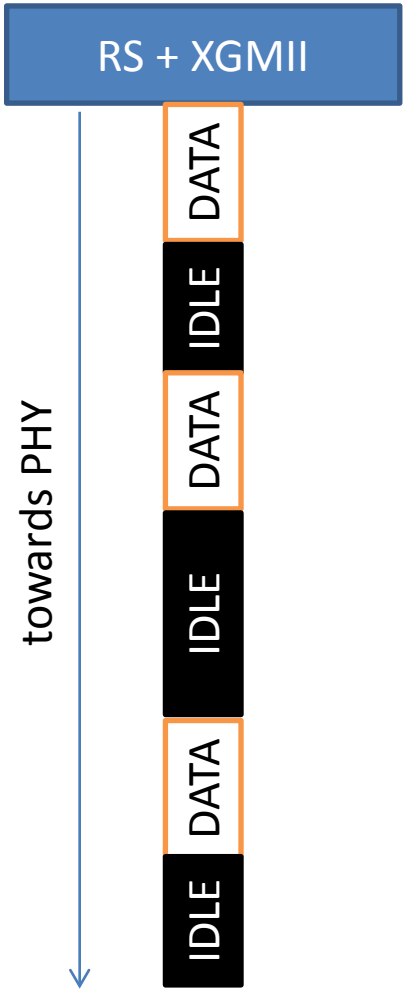

Idle Deletion & MMP (I)

- OLT Idle Deletion shown for reference only.
- Idle Deletion watches passing data vectors and figures out how many extra IDLEs need to be deleted to accommodate FEC parity and any PHY de-rating
- This function starts once and operates continuously as long as the station is powered on and data is transmitted by any of OLT MAC instances.

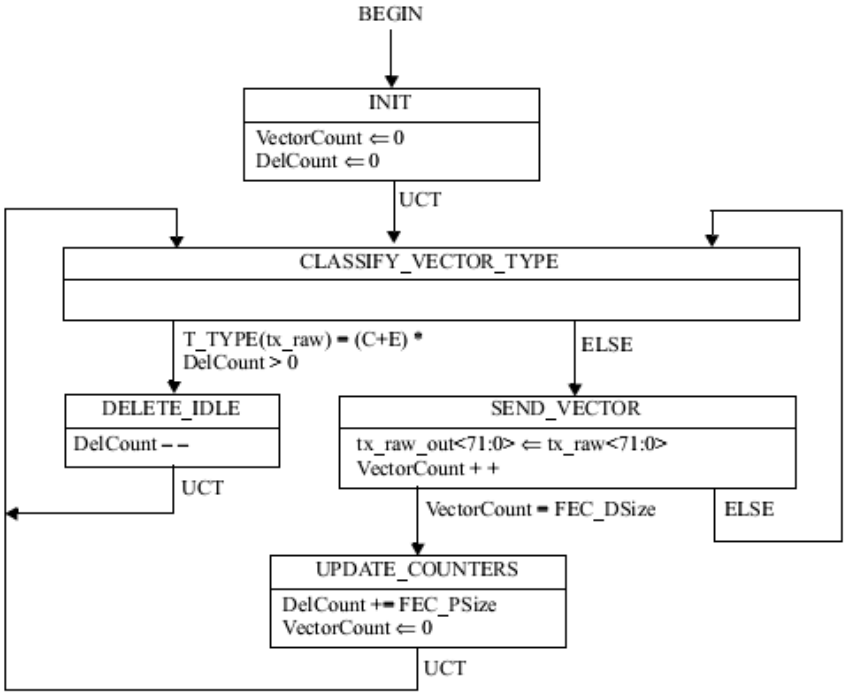

Figure 76–9—OLT Idle Deletion state diagram

Idle Deletion & MMP (II)

- MMP would require:
  - knowledge ahead of arrival of next frame as to what profile is used (and how many IDLEs need to be deleted)
  - some sort of real-time signaling between MPMC and PCS, across MAC and XGMII to indicate the next frame's profile
- Neither of these options are viable in 10G PCS
  - state diagram would need to be restarted for each data sequence (we would also need to have some SFD hunting function as well to know when to restart)
  - for each profile, proper set of parameters would need to be used (FEC overhead, PHY de-rating information, etc.) to delete appropriate number of IDLEs from the stream.

Smart Approach (I)

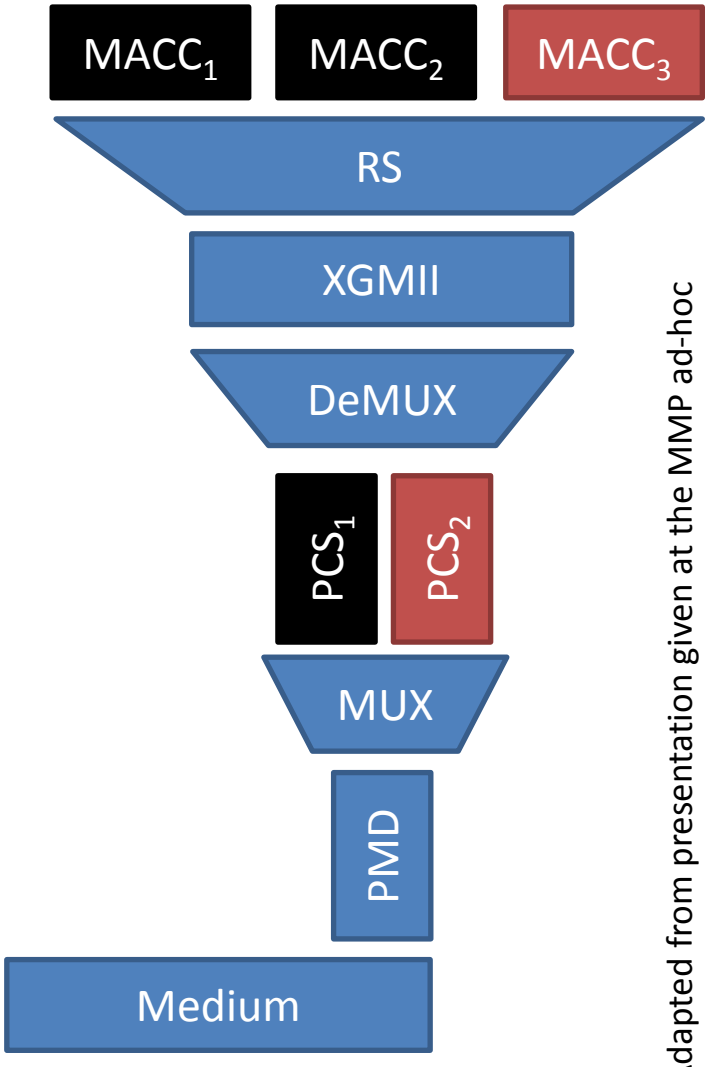

Adapted from presentation given at the MMP ad-hoc "Multiple_Profiles_for_EPoC_11Jan2013.pdf"

- It would work if the DeMUX function was not LLID based.
- LLID information requires sequence hunting for SFD in the incoming XGMII data stream
  - non-zero probability of false lock
- On false lock, frame fragment may be directed to incorrect PCS stack (see colored blocks).
- Such event will put MPMC view of PCS and actual PCS status out of sync for both correct and false-lock PCS instances. Only device reset will put things back in order.

Smart Approach (II)

- Mux function combines data streams from multiple PCS instances into one, to be presented to PHY
- Given different length FECs, it is possible that data becomes ready on more than one PCS instance (individual PCS instances run in parallel).
  In this case, MUX has to perform arbitration between different data sources. How selection is done, and controlled, is unclear at this time. It seems that some per-LLID decision would be needed, relying on information not available in PCS.
- It is also not clear, though, how profile information is propagated through PCS stack (new signaling interfaces ?) and correlated with individual data vectors.
- To a large extent, for this approach to work, Ethernet frame structure should be maintained within PCS, defeating the whole layering model used in 802.3.

Conclusions

- MMP at the logical layer seems technically feasible, and requires only minimum changes to definitions of the MPMC defined for 10G-EPON in Clause 77.
- MMP at the physical layer (PCS) is more complex and requires not only a brand new PCS, but also potentially require deeper changes in the layering architecture model. The scope of such changes would delay substantially the project and require changes to existing EPON MACs, defeating the very purpose of the project per its CFI.