EPoC PHY and MAC proposal

Marek Hajduczenia, PhD
ZTE Corporation
marek.hajduczenia@zte.pt

Supporters

- Alan Brown, Aurora
- Ed Mallette, Bright House Networks

RF spectrum churn

- Chunks of RF spectrum may be allocated to EPoC downstream / upstream channels as they become available in the MSO network
  - RF spectrum allocation can differ among MSOs, from region to region or even within a single MSO plant, given differences in service plans, legacy equipment and plant conditions etc.
  - An adaptive solution would be needed to address RF spectrum allocation for majority of MSOs

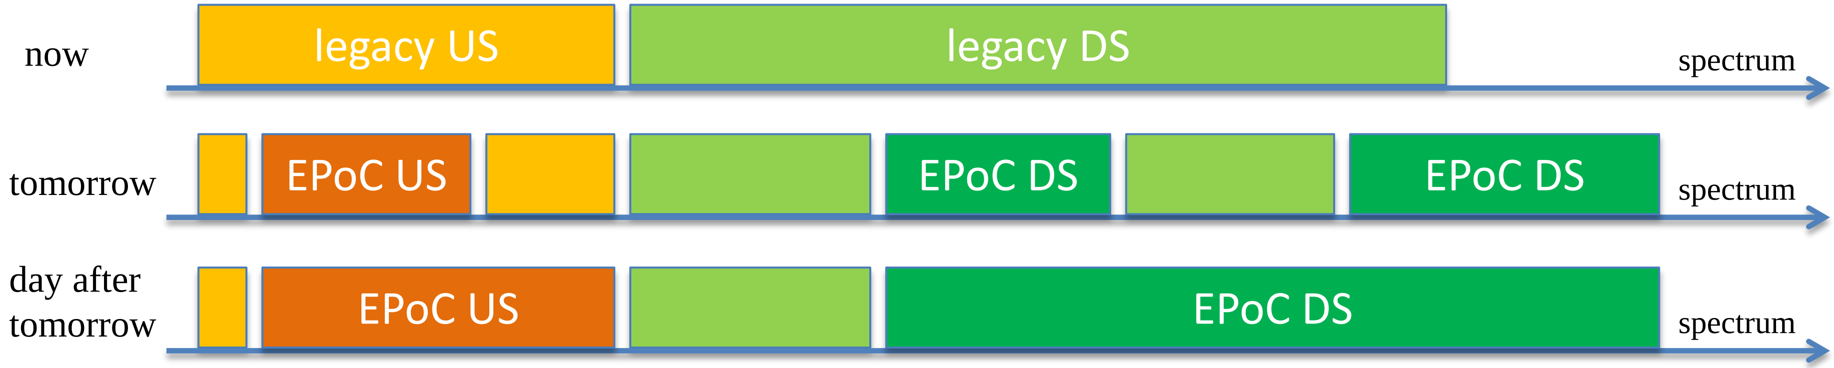

NOTE: images are illustrative only and DO NOT assume any specific RF allocation plan

Impact on P802.3bn TF work

- EPoC PMD must be (therefore) able to:
  - Switch selected RF spectrum sub-bands on or off, depending on selected configuration, availability and legacy services (baseline adopted in Geneva)
  - Report its current combined link capacity to MAC Control via MDIO for adjustment of the effective data rate (infrequent changes)
  - In the transmit direction, spread data stream coming from MAC across enabled RF spectrum bands, adapting to link conditions and its status.
  - In the receive direction, combine data stream received from enabled RF spectrum bands and recover single MAC data stream

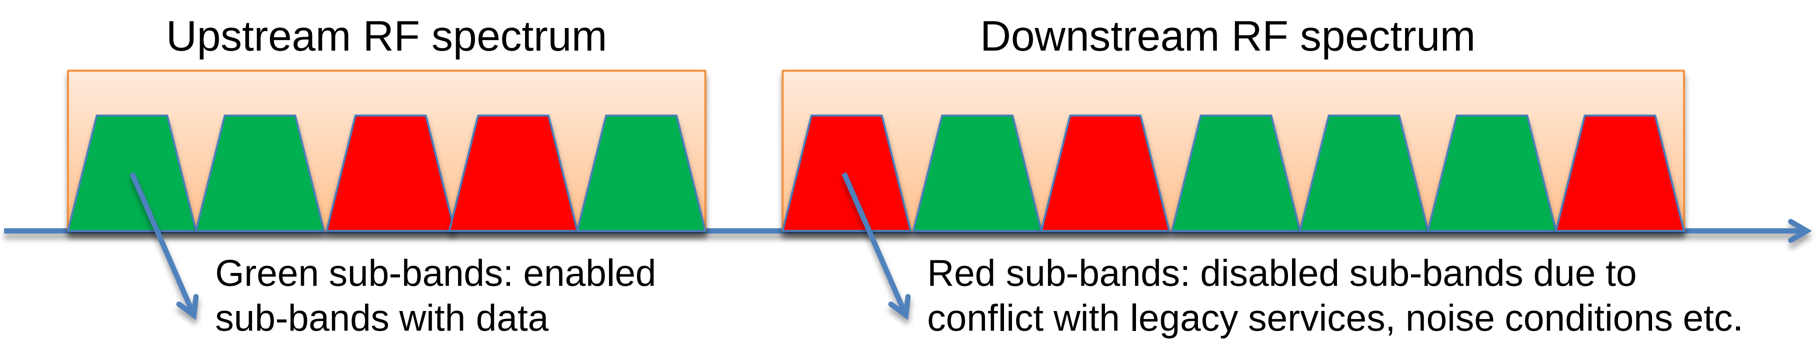

Sub-band specification

- P802.3bn TF will have to decide on:
  - Size, location and placement of upstream and downstream RF spectrum – all we have is size (192 MHz). This is where individual sub-carriers will be located.
  - Number of supported sub-carriers, their size, spacing and location within downstream / upstream RF spectrum bands (channelization)
  - Whether channelization is the same in downstream and upstream or not
  - And a whole bunch of parameters specific for the given multiple access technique and modulation format

Impact on MAC / PMD

- Sub-band specification also effects PCS / PMD design
  - MAC operates at 1G or 10G
    - there are no speeds like e.g. 150 Mbit/s
  - Need to decide which element in the EPoC stack enforces target data rate:
    - 10G-EPON-like data rate adaptation; or
    - Clause 61-like mechanism using Carrier Sense signaling
  - Need to decide how MAC data stream will be spread into individual sub-bands. Do we use:
    - 40G/100G P2P approach, or
    - Clause 61 architecture
- Individual items are discussed in the following slides

10G-EPON architecture

see hajduczenia_02_0712.pdf for more details

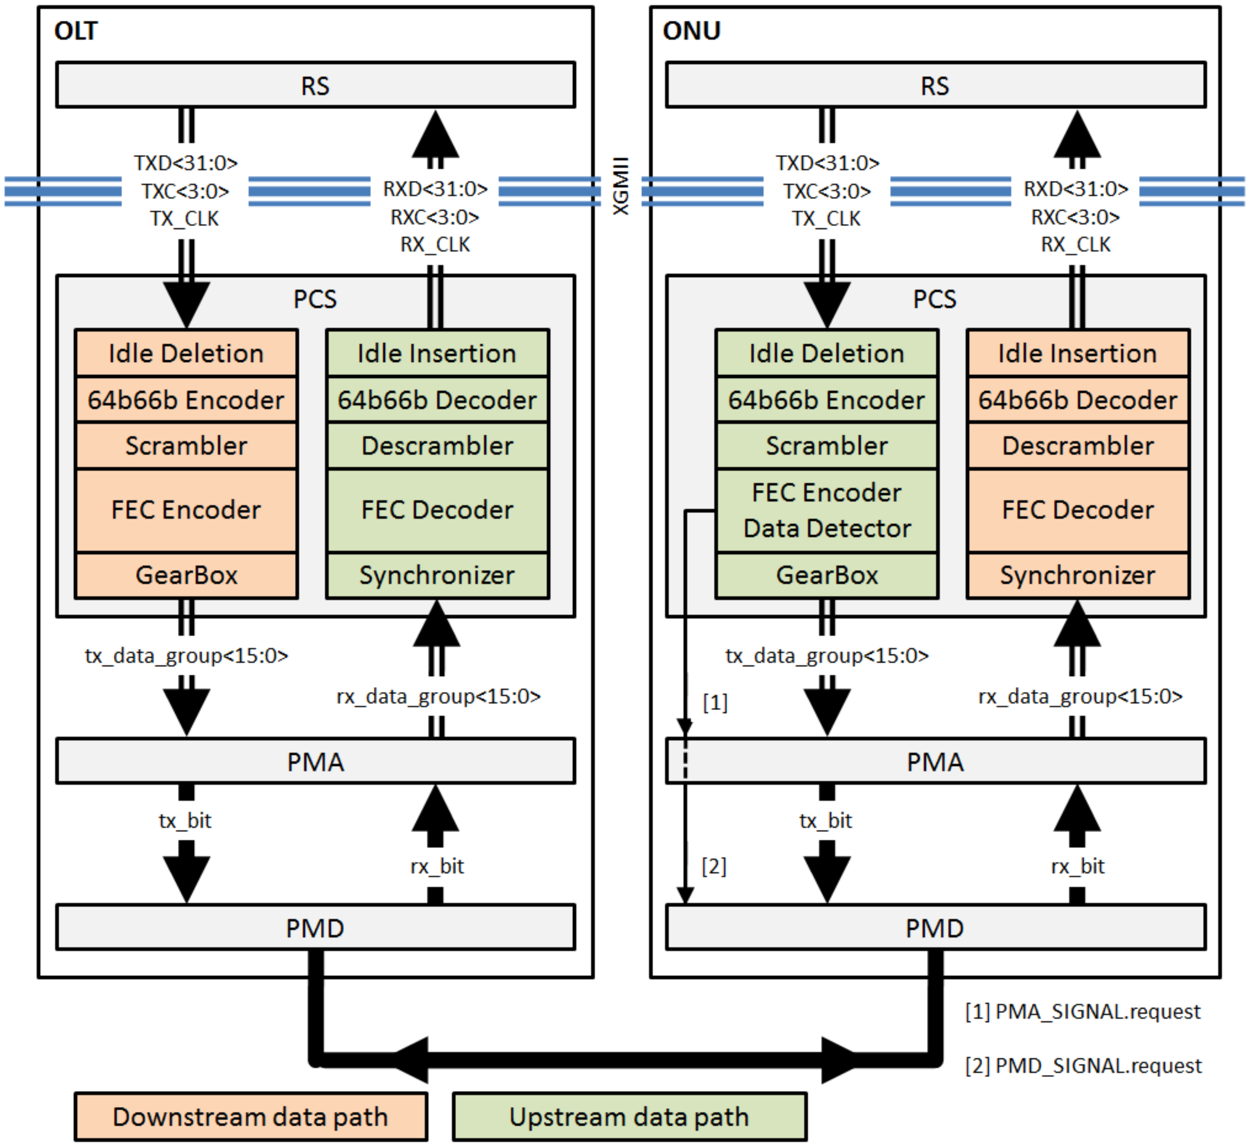

10G-EPON approach to data rate adaptation

see hajduczenia_02_0712.pdf for more details

- MAC transmits data at 10 Gbit/s across XGMII towards PMD
  - When no data is provided by upper layers, MAC transmits IDLE characters
- MAC Control function gates transmission of data frames to achieve certain predefined effective data rate
  - Effective PHY data rate is fixed at 223/255 × 10 Gbit/s (~8.74Gb/s) due to operation of mandatory, stream-based RS(255,223) FEC within PCS
- Idle Deletion function in transmit direction removes excess IDLE characters and prepares gaps for FEC parity insertion
  - Data detector combined with FEC encoder in transmit direction reconstructs continuous data stream at PMA, filling in gaps in data stream with FEC parity
- Idle Insertion function in receive direction fills in gaps from erased FEC parity with IDLE characters
  - MAC sees continuous 10 Gbit/s data stream again, with average of 32 IDLE characters per 255 received characters
- Why not use Carrier Sense in 10G-EPON?
  - See slide 9 in hajduczenia_02_0712.pdf presented in July 2012

Clause 61 architecture

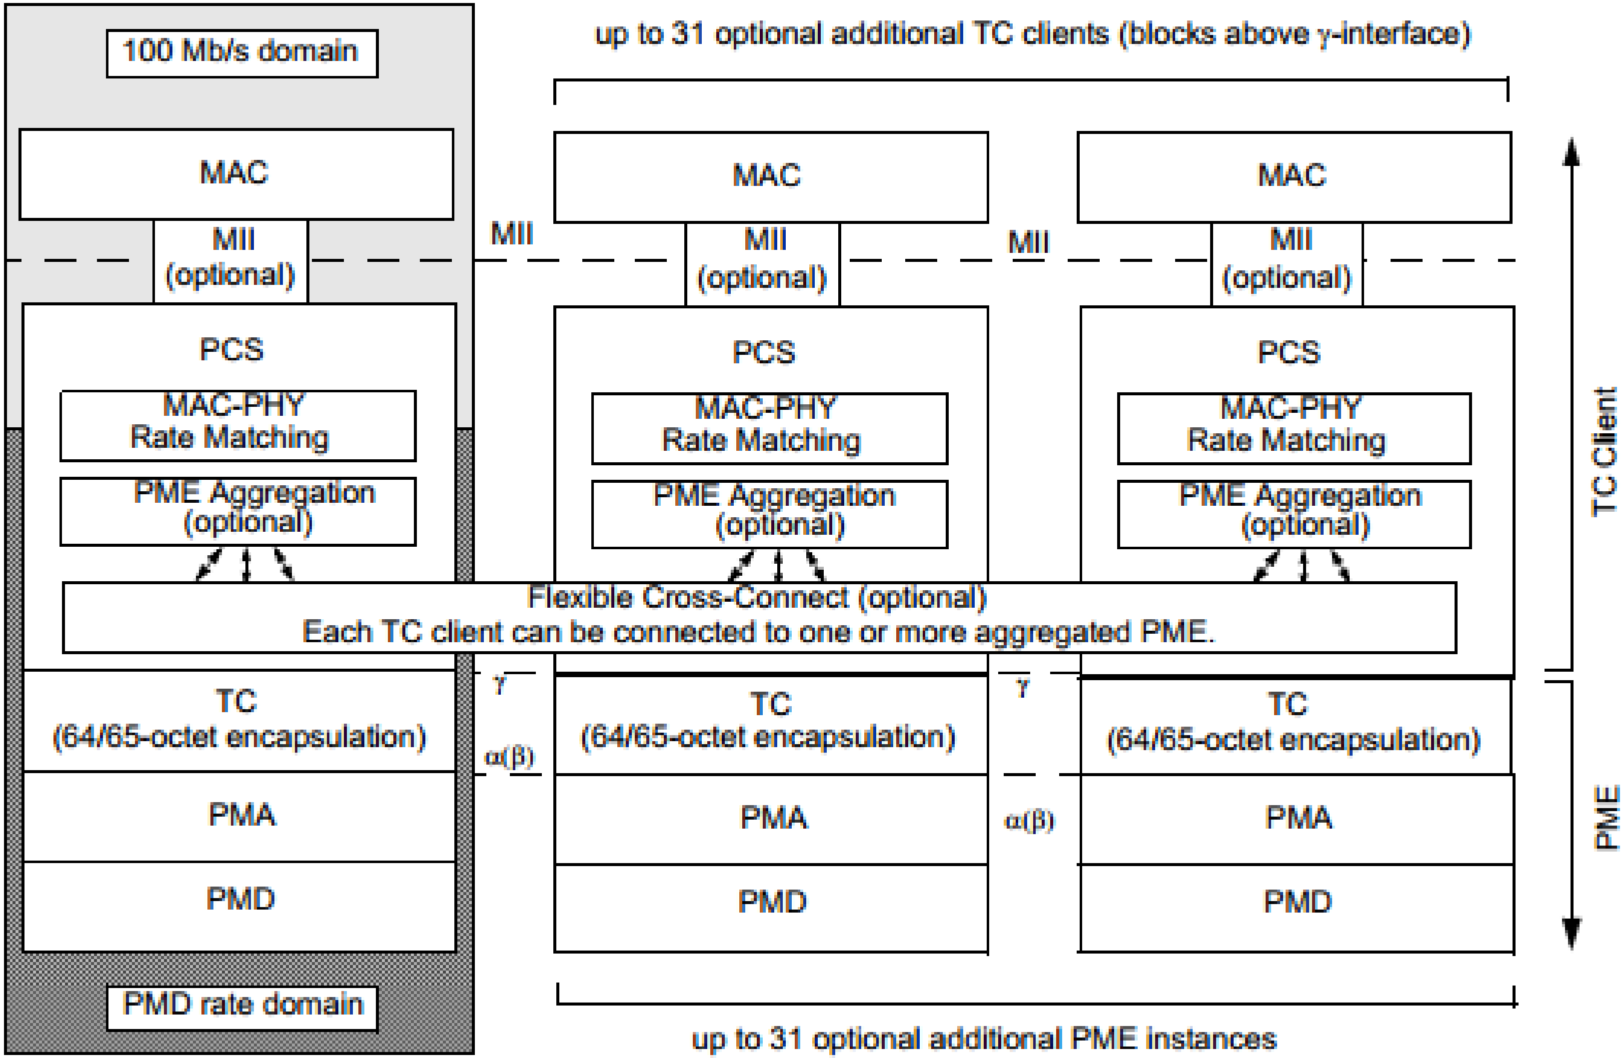

Figure 61–2—Overview of PCS functions

Clause 61 approach to multiple data lanes

- Clause 61 specifies PCS and TC layers for 10PASS-TS and 2BASE-TL PMDs using DSL-grade twisted pair media
- Data from a single MAC can be spread into a number of Physical Medium Entities (PMEs) through a PME Aggregation Function (PAF)
  - As a sort of free bonus, data from multiple MAC s can be aggregated as well using flexible cross connect function in PCS
- Each PME operates an independent PMD with separate medium attachment, data rate and link conditions
- PAF fragments incoming MAC frame into chunks and each chunk is then sent to separate PME, according to PME link capacity
- Each PAF-aggregated PMD can be managed and configured independently as part of a PMD bundle
  - Would a similar approach be welcome in EPoC ?

Clause 61 approach to data rate adaptation

- Data rate is adapted between PAF-aggregated set of PMD and MAC using the Carrier Sense (CRS) signal transmitted across MII:
  - When PCS buffers become full, CRS is raised, deferring transmission at MAC level until CRS is lowered.
  - MAC remains unaware of the actual PHY speed, operating at 100 Mbit/s and stopping / starting transmission based on CRS signal from PCS
- For this mechanism to work, MII has to be able to carry CRS signal between PCS and MAC
  - This works in MII (100 Mbit/s) and GMII (1 Gbit/s)
  - CRS signal is not supported in XGMII (10G), XLGMII (40G) and CGMII (100G) medium independent interfaces
- Conclusion:
  - Unless extensive work on design of extended XGMII is undertaken, Clause 61 approach to data rate adaptation is too limited for EPoC.
  - 10G-EPON solution seems to be more scalable and more future-proof

40GE/100GE architecture

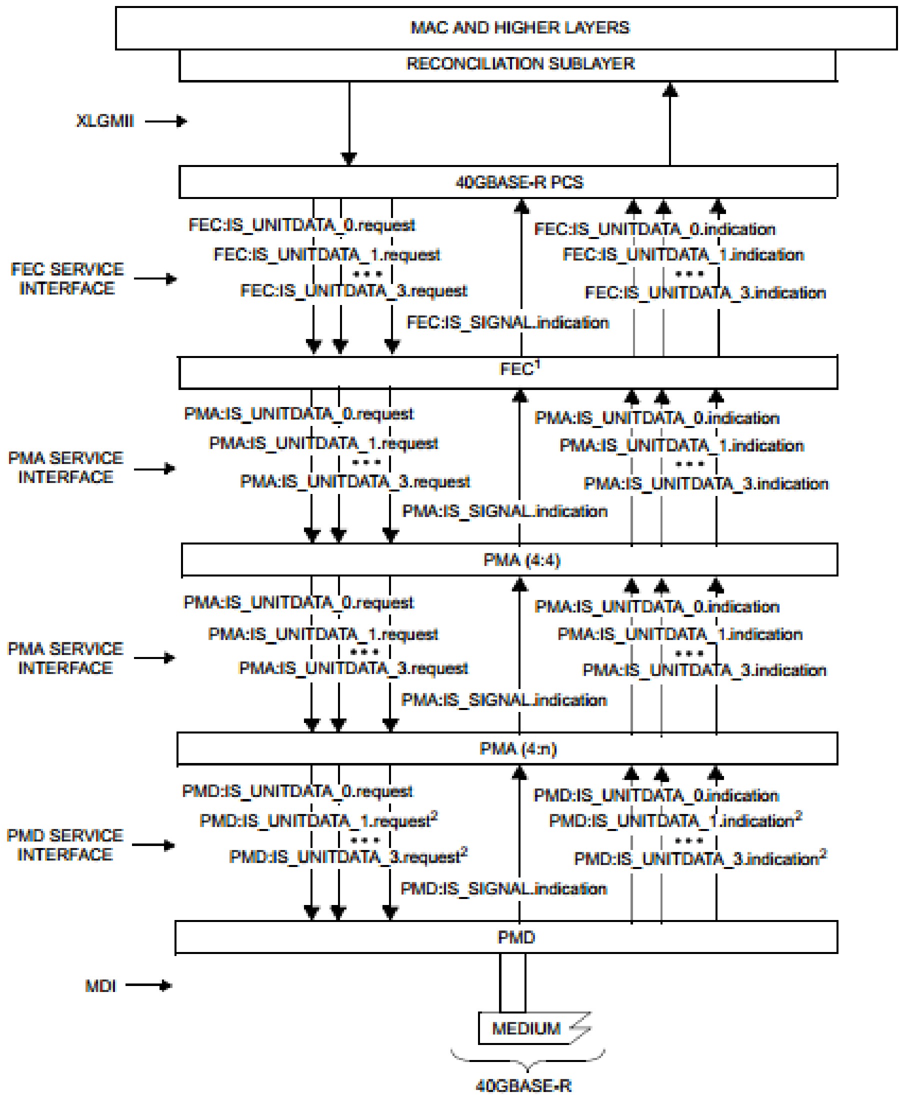

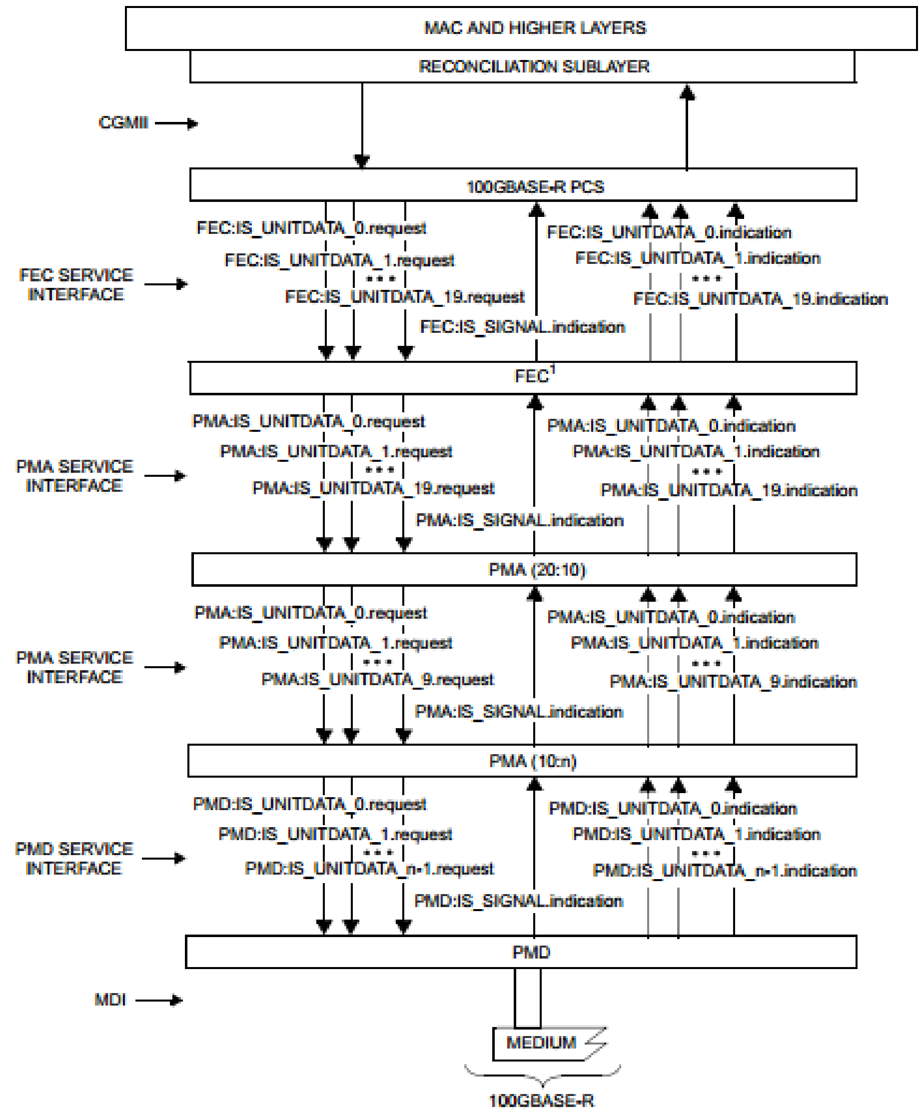

40G/100G approach to multiple lanes

- MAC transmits data at 40G or 100G across MII towards PMD
- There is no data rate adaptation function in P2P high-speed links
  - Data is always transmitted at the nominal MAC rate
  - PCS and PMD are prepared to handle such data stream from MAC
- When 40G/100G data stream hits PCS, the following happens:
  - Data stream is encoded in 64b/66b encoder
  - Resulting data stream is then scrambled and fed into the Block Distribution and Alignment Marker Insertion function
  - That function is responsible for spreading individual data blocks across N data lanes (N can be anywhere between 1 and 20 in existing PMDs)
  - FEC is applied to each and every individual spread data stream before it is handed off to PMA
- The process is reversed in the receive direction

How to put that all together?

- EPoC specification will define the following layers:
  - MAC Control (aka MPCP)
  - MAC (select only full-duplex or half-duplex mode)
  - RS
  - XGMII
  - PCS
  - PMA
  - PMD
- Individual layers, their functions and some of interfaces between them are described in the following slides

Proposed EPoC architecture

downstream example

Functions / Layers: MAC Control

- Functions:
  - Control MAC operation in P2MP environment of EPoC
  - Station discovery, registration and bandwidth allocation (MPCP)
  - Data rate adaptation to target PHY data rate based on PHY capabilities (data rate) retrieved via MDIO interface, accounting for overheads for 64b/66b encoding, FEC coding (parity), insertion of alignment markers etc.
- Specification blueprint:
  - Clause 77 (10G-EPON) for MPCP functions (first approximation)
  - Clause 77 (10G-EPON) for data rate adaptation, with changes to support variable data rate and different FEC type than in 10G-EPON
  - Clause 45 for MDIO registers (can be located in any EPoC layer)
- Challenges
  - Changes to data rate adaptation function from EPON, especially if changes in data rate are allowed during operation – CNU re-registration may be needed
  - Support for Multiple Coding Schemes (MCS): PCS and MAC Control have to become destination-aware, and modify the data rate in function of the destination station.

Functions / Layers: MAC and RS

- Functions:
  - In transmit direction, MAC serializes data packets received from MAC Control and fills in spaces between frames with IDLE characters; add preamble and FCS to frames received from MAC Control;.
  - In receive direction, MAC deserializes data, confirms that FCS is correct and prepares packet for transmission to higher layers.
  - In transmit direction, RS modifies content of preamble adding LLID information to reconcile P2MP plant with P2P Ethernet MAC
  - In receive direction, RS decides (based on LLID) which MAC instance to forward the given frame to.
- Specification blueprint:
  - Annex 4A for MAC
  - Clause 76 (10G-EPON) for RS
- Challenges
  - none identified right now

Interface: XGMII

- Functions:
  - XGMII acts as media independent layer connecting MAC (through RS) to the PHY.
  - XGMII is a bidirectional, 32-bit wide interface (4 data octets per transfer) in each direction, operating at the effective data rate of 10 Gbit/s per direction (312.5 Mtranfers / second).
  - XGMII also has 4 bit control interface (per direction) and a single clock lane (per direction)
- Specification blueprint:
  - Clause 46
- Challenges
  - none right now

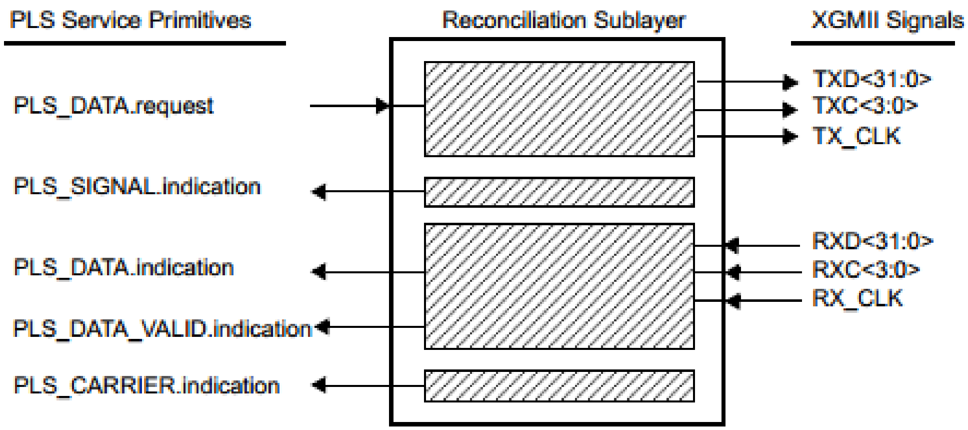

Figure 46–2—Reconciliation Sublayer (RS) inputs and outputs

Function / Layer: PCS

- Functions:
  - In the transmit direction, condition the serial bit stream from MAC to transmission over copper medium, encoding it accordingly, providing redundancy data (FEC) and preparing individual data streams for each individual transmitter
  - In the receive direction, decode data received from each receiver, perform data decoding and error recovery based on FEC parity, and recover a single bit stream suitable for MAC
- Specification blueprint:
  - Clause 76 (10G-EPON) for idle insertion / deletion functions
  - Clause 61 (10PASS-TS and 2BASE-TL) for PME aggregation function
- Challenges
  - Definition of a new FEC in downstream and upstream directions
  - Definition of PME aggregation function (larger number of PMDs is possible), which is more scalable than Clause 61 definition
  - Synchronization in transmit and receive directions, especially in upstream due to burst mode operation
  - Support for MCS: how to separate frames intended for different modulation depths from a serialized data stream without changing XGMII specifications and definition of new control characters for XGMII.

PCS: Idle Deletion / Idle Insertion

- Functions:
  - In the transmit direction, Idle Deletion function removes extra IDLE characters inserted by MAC between individual frames from MAC Control.
  - The number of deleted IDLE characters depends on the overhead resulting from 64b/66b encoding, FEC coding (parity), insertion of alignment markers and target PHY data rate.
  - In the receive direction, Idle Insertion function adds IDLE characters between individual recovered frames, preparing it for transmission across XGMII
- Specification blueprint:
  - Clause 76 (10G-EPON) for Idle Insertion / Deletion functions
- Challenges
  - None so far if MCS is not used. Substantial rewrite / redesign will be needed for MCS.

PCS: 64b/66b Encoder / Decoder

- Functions:
  - In the transmit direction, it encodes 64 data bits into 66 bit symbol, adding two control bits (Sync Header), indicating whether given symbol carries data or control character
  - In the receive direction, this function retrieve 64 bit blocks from 66 bit encoded stream after it has been processed by FEC decoder and de-interleaver
- Specification blueprint:
  - Clause 82 for 40G/100G 64b/66b encoder/decoder function with definition of alignment markers
- Challenges
  - Alignment markers might need modification. Spacing between alignment markers might need adjustment as well.

PCS: Configurable Interleaver / deInterleaver

- Functions:
  - In the transmit direction, the Configurable Interleaver performs the configurable interleaving operation on the received data stream. The Interleaver operates on bursty data.
  - In the receive direction, the Configurable deInterleaver performs the configurable deinterleaving operation on the received data stream. The deInterleaver operates on bursty data.
- Specification blueprint:
  - None so far, might reuse DOCSIS definitions (?)
- Challenges
  - None so far

PCS: Block Distribution + Alignment Marker Insertion

- Functions:
  - Block Distribution function distributes individual 66-bit symbols towards individual PMDs across associated PMAs (PMEs).
  - Additionally, every N symbols, a set of alignment markers is inserted into data stream, one marker per PME.
- Specification blueprint:
  - Clause 82 block distribution mechanism developed for multi-lane 40G and 100G P2P PHYs, including insertion of alignment markers for skew compensation.
- Challenges
  - Potentially large number of PMD/PMA (PME) entities connected to a single cross connect (larger than in any 40G/100G P2P PHY) – scalability may be an issue (needs better study)

PCS: PME Lane Deskew & Block Sync, Alignment Marker Removal

- Functions:
  - Performs lane deskew based on received alignment markers
  - It obtains lock to the 66-bit blocks in each lane using the sync headers and outputs 66-bit blocks.
  - Next, all alignment markers inserted into the data stream are removed prior to handing off data to deInterleaver function.
- Specification blueprint:
  - Clause 82 for PCS lane reordering, PCS lane deskew, Block synchronization and Alignment marker removal.
- Challenges
  - Potentially large number of PMD/PMA (PME) entities connected to a single cross connect (larger than in any 40G/100G P2P PHY) – scalability may be an issue (needs better study)

PCS: FEC / deFEC

- Functions:
  - In the transmit direction, performs FEC encoding the whole data stream transmitted by the given PME (data and IDLE characters). FEC encoder operates on bursty data from Block Distribution function and produces continuous data stream.
  - In the receive direction, performs FEC decoding on a continuous data stream, removes parity, leaving gaps in place of parity. FEC decoder produces bursty data, that is later fed to PME Lane Deskew function
- Specification blueprint:
  - May be derived from Clause 77 (10G-EPON) FEC encoder / decoder
- Challenges
  - Different data rates (clock synthesis will be needed) for operation in upstream and downstream directions; operation on bursty and continuous data streams

Functions / Layers: PMD + PMA (PME)

- Functions:
  - In the transmit direction, performs transmission of a sequence of 66 bit symbols with FEC parity across coaxial cable towards the remote station.
  - In the receive direction, receives bit stream comprising a sequence of 66 bit symbols, which are then aligned and presented across PMA to FEC decoder.
  - Each PMD in this model represents a single OFDM subcarrier used for transmission.
  - Naturally supports exclusion bands (control individual PMD status via management)
- Specification blueprint:
  - None, can reuse specification format included in Clause 75 (10G-EPON), but needs to be developed from scratch.
- Challenges
  - Depending on the number of subcarriers, PMD count can be high.
  - Specification format for large PMD number.
  - Might need a new entity, PMD bundle, to designate all PMDs combined together.
  - Clause 45 management for large PMD count may be complex.

Proposal

- Adopt the layering model shown in slide 14 as baseline for EPoC
- Start developing more rigorous description of individual functions / layers and interfaces, based on slides 15-26
  - Individual contributions could be geared towards draft D0.1
  - We need to start focusing discussions and showing more detailed proposals to keep with the timeline from Geneva
  - Development of TDD-specific definitions of sublayers can be left to the TDD track, but must keep common layers intact.
  - Support for MCS should be discussed separately.
    - Let's develop a single coding profile system first and then add MCS into it, if needed and when proven it can done.
    - We could also limit MCS to TDD mode only, where it could be easier to implement than in FDD system.
    - MCS support can be a distinguishing feature between FDD and TDD modes

Bringing you Closer

Thanks!