(major) challenges ahead of EPoC

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

Motivation

- P802.3bn TF is at the beginning of its PMD development path and a number of challenges have been identified (see next slide), requiring deeper analysis and contributions from TF membership
- In this contribution I examine two of these challenges with the focus on:
  - Project timeline (we need to have EPoC done ASAP)
  - Project objectives (as approved)
  - 802.3 history and existing specifications
- A few recommendations for tackling individual challenges are made. Based on TF feedback, these may be converted into baseline proposals

(major) challenges identified so far

Based on online and offline discussions, a few technical challenges can be identified:

1. Discontinuous RF spectrum allocation in upstream and downstream directions
2. Support for future EPoC evolution and coexistence on the same CCDN
3. Selection of multiple-access (MA) scheme for EPoC
4. Bandwidth allocation in EPoC
5. Retaining existing MPCP definitions in EPoC
6. RF frontend design and its technical parameters
7. Channel / link model
8. EPoC performance (delay / jitter / bandwidth availability)
9. …

Of these list, the next slides will address items 1, 2, and 5 only at this time. Items 3, 4, 6, 7, and 8 are discussed in separate contributions to this meeting.

The remaining aspects will need more discussion at the following meetings and require more technical details on EPoC PHY design

Supporters

- Ed Boyd, Broadcom

CHALLENGE 1

Discontinuous RF spectrum allocation in upstream and downstream directions

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, etc.
  - There is no way to force all operators to use the same RF spectrum allocation plan given local and global differences in service plans, legacy equipment and plant conditions
  - An adaptive solution would be needed to address RF spectrum allocation for majority of MSOs

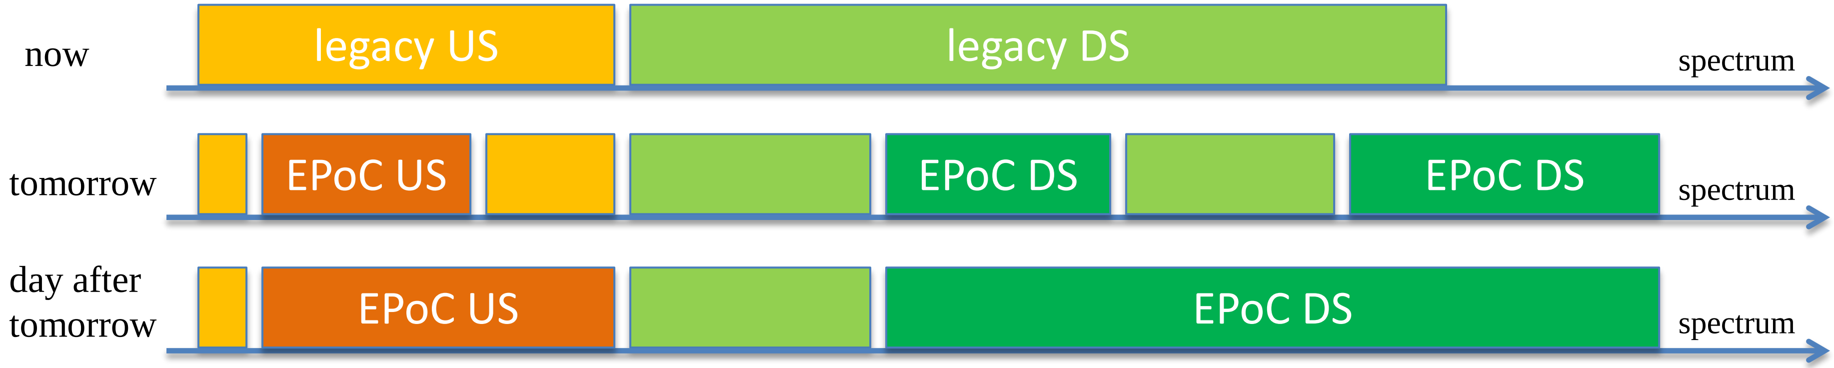

NOTE: images are illustrative only and DO NOT assume any specific RF allocation plan

Impact on P802.3bn TF work

- RF spectrum churn influences PMD design, putting the following requirements. PMD must be able to:
  - Switch selected RF spectrum sub-bands on or off, depending on selected configuration, availability and legacy services
  - Report its current combined link capacity to MAC Control via MDIO for adjustment of the effective data rate
  - Spread data stream coming from MAC across available RF spectrum bands, adapting to number of available sub-bands, their link conditions and current status.
  - Combine data stream received from available RF spectrum bands and recover single MAC data stream

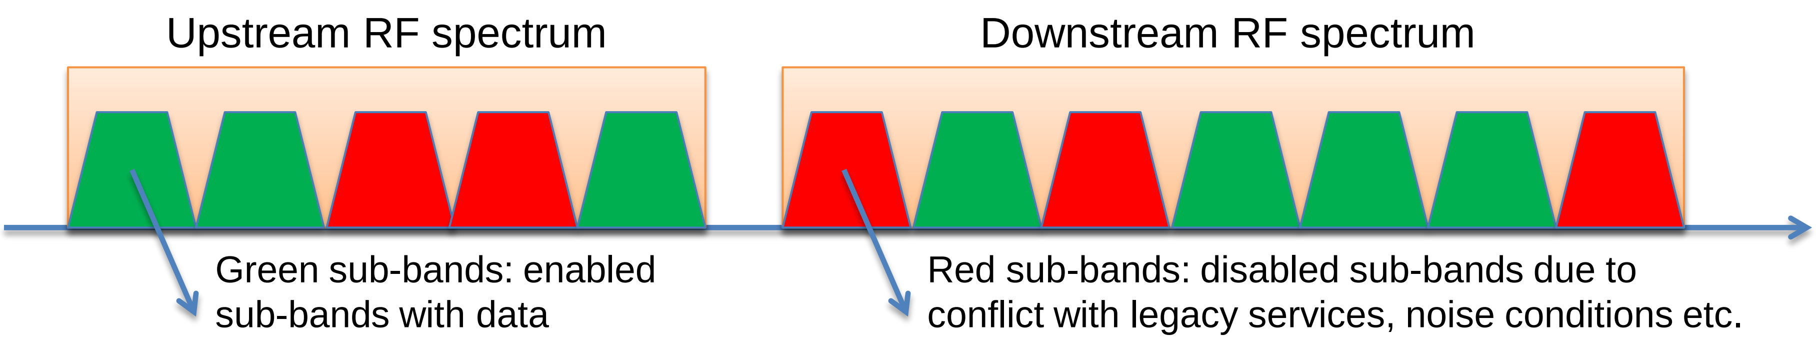

Sub-band specification

- P802.3bf TF at some point of time will have to decide on:
  - Size, location and placement of upstream and downstream RF spectrum bands within which individual RF spectrum sub-bands will then be defined
  - Number of supported RF spectrum sub-bands, their size, spacing and location within downstream / upstream RF spectrum bands
  - Whether channelization is the same in downstream and upstream or not
  - And a whole bunch of parameters specific for the given MA and modulation format selected by TF
- Sub-band specification has also effect on 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 rate adaptation or something similar to Clause 61 mechanisms using Carrier Sense signalling?)
  - 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?

10G-EPON architecture

see hajduczenia_02_0712.pdf for more details

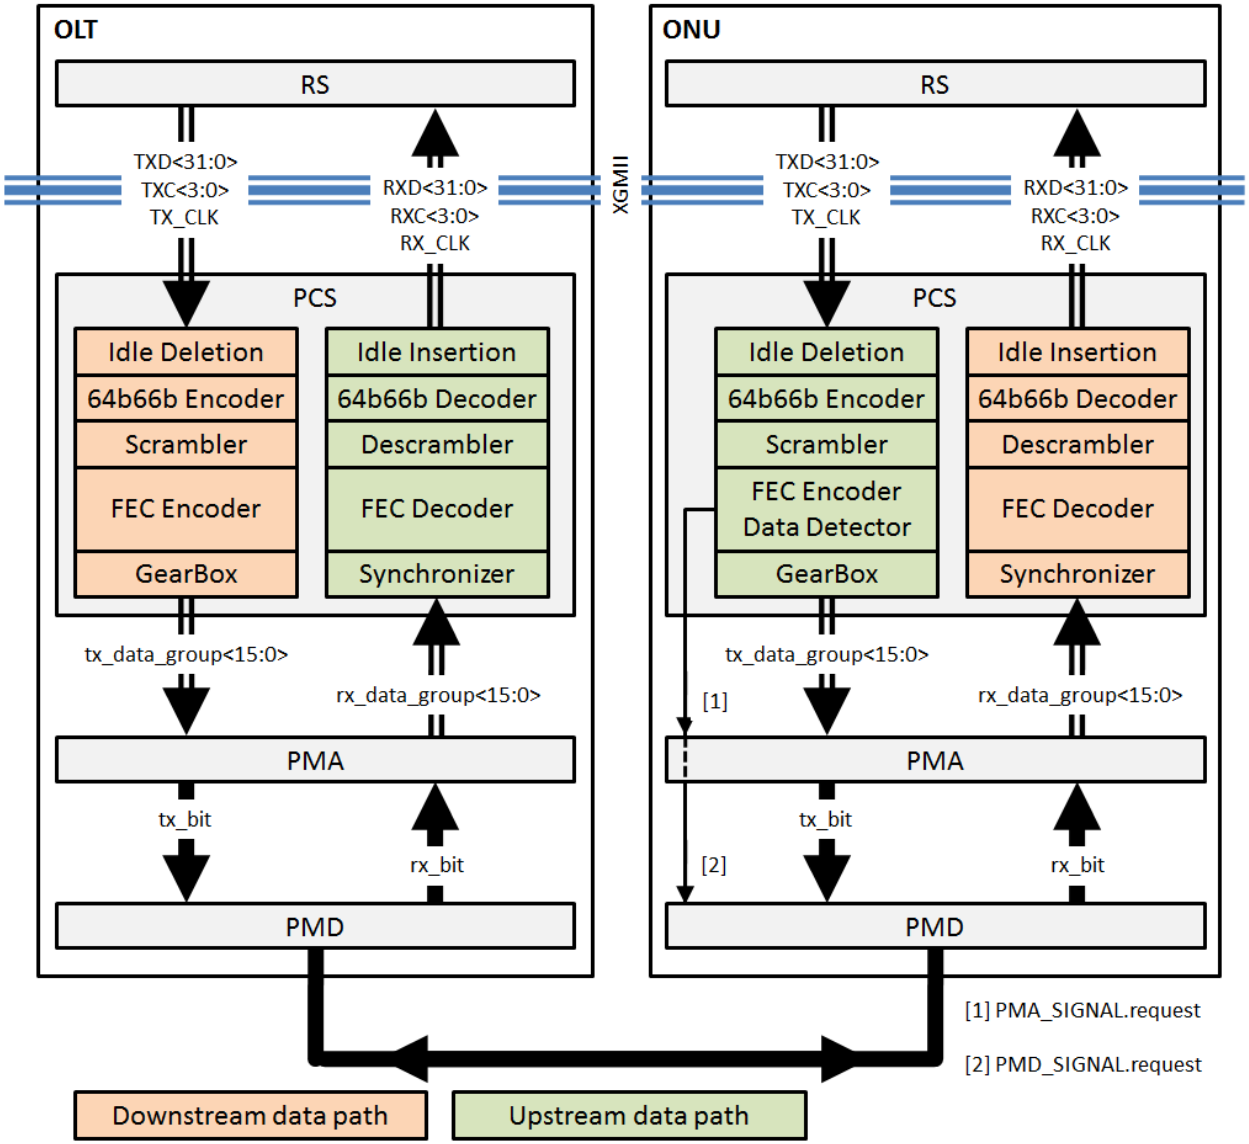

10G-EPON approach to data rate adaptation

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 idles
- MAC Control function gates transmission of data frames to achieve certain predefined effective data rate
  - In 10G-EPON, this data rate is fixed at 223/255 × 10 Gbit/s due to operation of stream-based RS(255,223) FEC
- Idle deletion function in transmit direction removes excess idles 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 idles characters per 255 received characters
- Why not use Carrier Sense in 10G-EPON?
  - See slide 9 in hajduczenia_02_0712.pdf presented in July

Clause 61 architecture

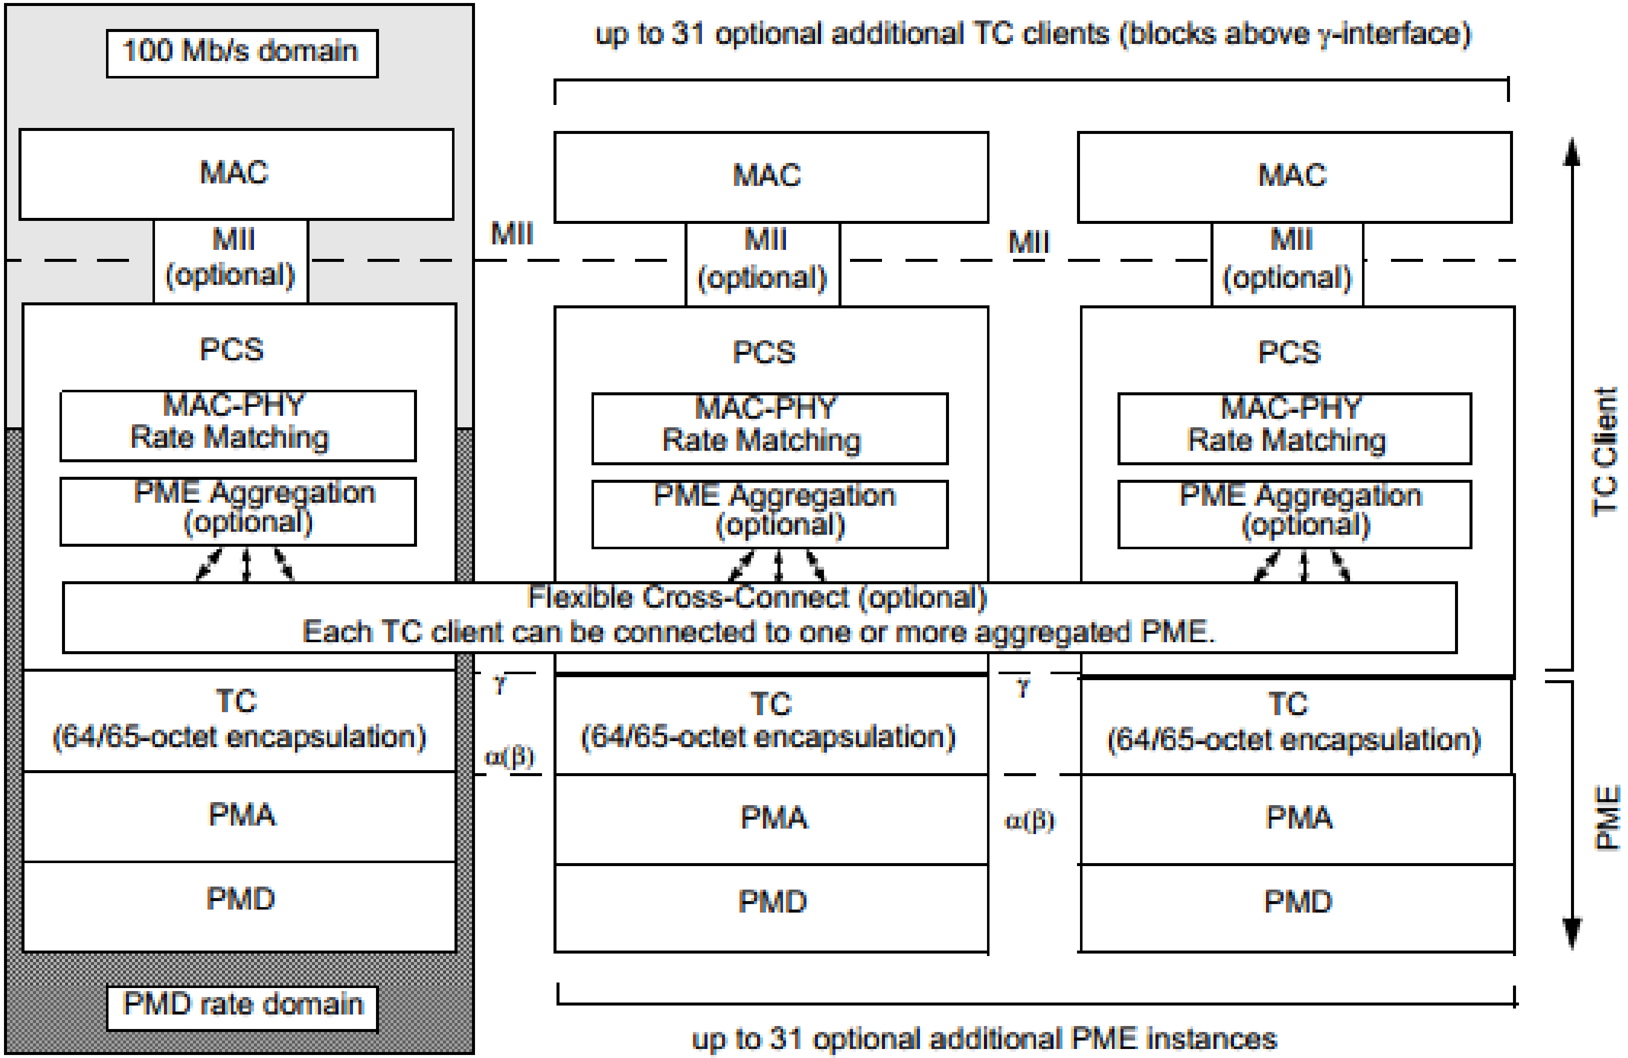

Figure 61–2—Overview of PCS functions

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 unaware, 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

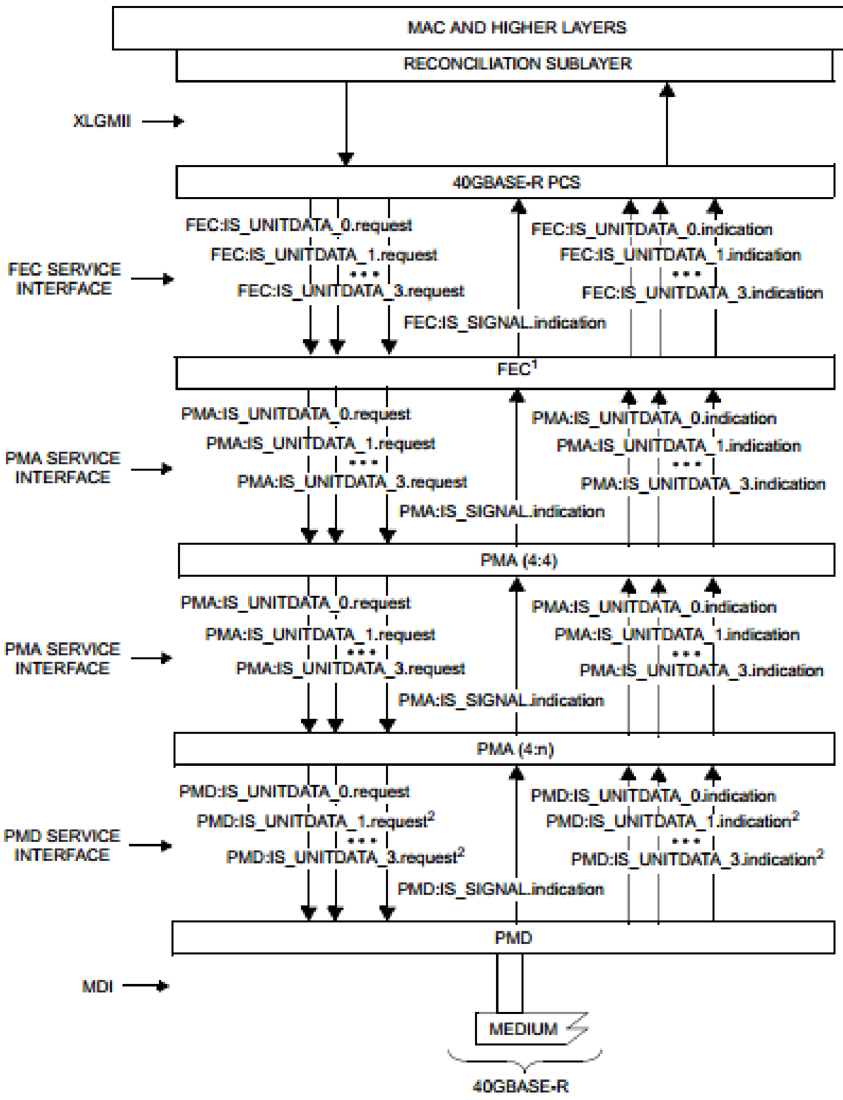

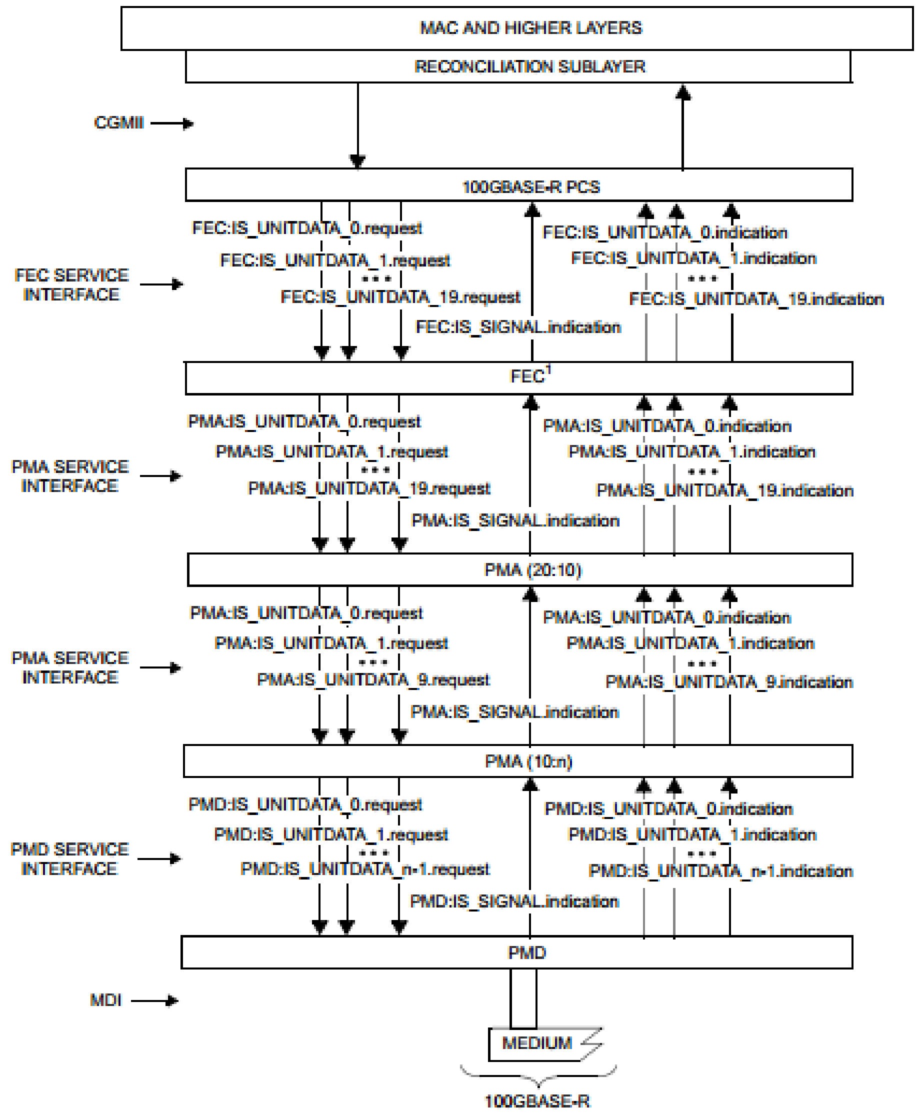

40G/100G approach to multiple lanes

- 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

Proposed approach for EPoC

- A high level proposal to solve EPoC challenge 1:
  - Adopt 10G-EPON-like data rate adaptation mechanism, where MAC Control gates transmission towards MAC in a way to achieve average data rate ≤ than data rate reported by PMD
  - Use Clause 61-like multi-lane approach with PAF-like aggregation function for multiple PMD instances.
  - Each PMD instance is to report its data rate capabilities across MDIO.
  - Whole PAF-aggregation PME bundle also reports aggregate bandwidth capabilities via MDIO
  - Individual PMEs can be selectively enabled / disabled by management to adapt to RF spectrum availability on the given CCDN
- Points to consider:
  - Support for MAC aggregation using Clause 61-like mechanism may be hard to achieve in P2MP architecture of EPoC

CHALLENGE 2

Support for future EPoC evolution and coexistence on the same CCDN

Support for multiple generations

- 802.3 standards evolve over time, adding higher data rate support and/or new features in each next generation
- P802.3bn "EPoC" should be no different in this aspect
  - We start with 1st generation EPoC based on technically and economically feasible technical solutions available today
  - But at this time, it is not possible to determine with 100% certainty what 2nd generation of EPoC will look like …
- Support for multiple equipment generations is something that P802.3bn TF needs to address sooner rather than later:
  - This affects technical choices we will have to take going forward
  - Wrong choices or no discussion right now is likely to influence negatively extensibility of this technology in the future

How can we support evolution of EPoC?

- Solution 1: forget about EPoC evolution. Fiber is the end-game and EPoC is single-generation only.
  - Not an ideal response to MSOs with new coax plants, brown-field scenarios and retrofitting applications
- Solution 2: use an incremental evolution model, similar to DOCSIS 3.0
  - As time goes by, new sub-bands are added to RF spectrum band available for EPoC.
  - More bandwidth can be transmitted up to the MAC data rate
  - Some sort of channel (PHY) bonding will be then needed
- Solution 3: use a revolutionary model, similar to EPON
  - When needed, new RF spectrum band will be appropriated by next-gen EPoC and used exclusively for higher-speed users

Solution 2

- Advantages:
  - pay-as-you-grow approach, where higher-speed devices can reuse previously allocated spectrum
  - older devices still get their data on allocated sub-bands
  - newer devices get their data on all allocated sub-bands
- Disadvantages:
  - Technical challenges associated with selective spreading of data frames: some frames are transmitted only on legacy channels and some are transmitted across whole RF spectrum
  - Such a channel bonding technology is already available – we're not developing anything better than DOCSIS 3.0 in that respect. All challenges of DOCSIS 3.0 channel bonding would be still present in EPoC

Solution 3

- Advantages:
  - Clear delineation between legacy and next-gen systems
  - No complex technical solutions required for achieving coexistence. Data for legacy and next-gen systems is mixed on CCDN in different RF spectrum bands
  - Multiple generations can be stacked on top of each other as along as there is enough RF spectrum available
- Disadvantages:
  - Data replication is needed for some of broadcast and multicast data. EPON seems to handle that gracefully
  - Spectrum expensive approach where RF spectrum occupied by legacy cannot be reused for next-gen systems.

Solution X ?

- What if we combine Solution 3 and Solution 2:
  - Assuming EPoC is designed with flexible approach to RF spectrum bundling (per slide 15):
  - Gen1 is allocated band F1 – F2
  - Gen2 is allocated band F1 – F3

  - As long as Gen1 devices are present on CCDN, Gen2 devices are configured to operate on F2-F3 band only

Solution X ? (cnt'd)

– Over time, as Gen1 devices are replaced on CCDN with Gen2 devices, F2 can be moved gradually to the left, through reconfiguration of Gen1 and Gen2 devices

– When the last Gen1 device is removed, Gen2 devices can occupy the whole F1 – F3 RF spectrum band

– This approach can be scaled into multiple equipment generations and use contiguous / non-contiguous RF spectrum sub-bands

Now, what do we need to do with it?

- Nothing at this point of time, as long as requirements from slide 15 can be addressed
  - We can come back to discuss this topic a few years from now, when and if next-gen EPoC systems need to be designed
- Coexistence and evolution can be then guaranteed by inherent EPoC PMD flexibility
  - No need to spend too much time on it right now
  - First gen EPoC devices could deliver up to XXX G bandwidth – XXX is TBD at this time, subject to detailed technical discussions. 5G seems within reach, 10G may be possible.
  - Higher data rates on coax may be challenging to achieve, if possible at all

MSO awareness is the key

- MSOs need to be then aware of the evolution path in this RF spectrum-restricted environment:
  - Next-gen devices can only reach full bandwidth potential when legacy devices disappear from the network
  - In coexistence mode, only non-overlapping spectrum can be used to their fullest extent by next-gen devices.
    - This is similar to the use of RF spectrum in DOCSIS - some channels end up being used for DOCSIS 1, 2, and 3 simultaneously, with, for example, differing modulation profiles or channel bonding capabilities.
  - Legacy devices always have access to their bandwidth, as configured by the operator

CHALLENGE 5

Retaining existing MPCP definitions in EPoC

“fixing” the MPCP – good or bad?

- MPCP is the primary target for optimization attempts, “fixing” and “improvements” in multiple directions
  - Something always can be done better, nicer or using more flexible message format
- However, in many such change attempts, the main purpose of MPCP is lost
  - It controls Multi-Point transmission in PON, as the name suggests, allowing stations to register and access services
  - It also facilitates data exchange, giving the central node a mechanism to grant bandwidth and poll stations for queue status
- MPCP is part of MAC Control
  - should be exposed to the minimum level of details about the underlying physical channel properties
  - In 10G-EPON, it is only aware of the upstream data rate supported by the device to allow DBA to calculate grants correctly

Some facts on MPCP

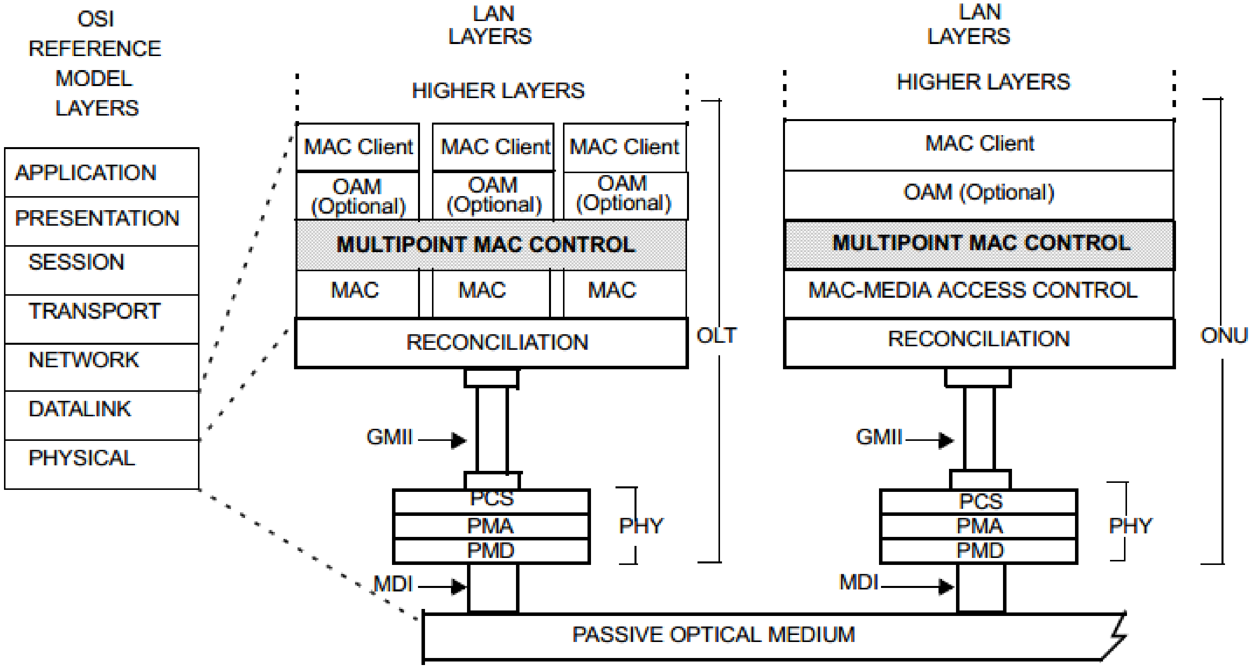

- MPCP is part of MAC Control
  - Its role is to control MAC behaviour, by gating transmissions and handling MAC Control messages
  - MAC Control has no role in controlling PHY layer parameters – these should not be exposed to MAC Control

Changing MPCP role

- MPCP changes seem fairly easy
  - Just add new field here and there, fix codes, and we're done
- Think about the bigger picture, though:
  - Objectives state clearly to limit changes to MPCP to "minimum augmentation" only
  - What "minimum" is, can be obviously argued in both directions
  - However, the function of MPCP and its role in EPoC stack should remain unchanged from EPON and Ethernet at large
- MPCP kicks in only when link has been established
  - It should have no part in initial channel set up, link configuration, selection of carriers etc. – this is the role of PHY
  - The role of MPCP should remain the station discovery and bandwidth allocation once the link is established and operational

Changing MPCP messages and SDs

- Again, fairly easy target to shoot at
  - We have reserved fields, padding – plenty of space for innovation
  - A number of reserved OPCODEs available for MAC Control as well
- Impact of every such change
  - We have a very aggressive timeline for this project
  - Every change in MPCP SDs from EPON requires months, if not longer, of debugging, design and testing
  - That alone can put project timeline through the roof, not to mention so desired product availability
- Conclusions
  - We should avoid changes to MPCP message format and MPCP SD if we can live without them, even at the cost of decreased efficiency
  - Overall system efficiency is not dictated only by MPCP but combination of physical link, PHY definitions and ability to grant bandwidth. Such impact can be only assessed through simulation.

EPoC operation steps

- PHY-only operation
  - CNU PHY goes online, initial (local) configuration of PHY parameters is performed to adapt PHY to target RF spectrum
  - PHY receiver is enabled, pilot tone hunting
  - PHY link established, auto-negotiation with CLT PHY begins. CNU MAC is not yet active.
  - Link parameters between CNU and CLT were established, CNU switches to target configuration
- MAC and PHY operation stage
  - Only when PHY layer link is established at target data rate, MAC is enabled on CNU and CLT
  - MPCP, OAM, extended OAM discovery and registration processes are started

5 commandments for MPCP changes

- Don't:
  - Mix PHY link establishment with MPCP discovery – these are separate processes, executed at distinct times and with different purposes
  - Modify the role of MPCP in P2MP environment: station discovery on working PHY link and bandwidth allocation only
  - Solve PHY layer problems at MAC and MAC Control layers. MAC should be PHY unaware as much as possible.
  - Bring in proposals for MPCP changes without clear motivation for them.
    - We could have a separate project on such "improvements" but these are outside the scope of this particular project

Changes to MPCP – path forward ?

- Challenges for proposals for MPCP changes:
  - Demonstrate what is broken in current MPCP station discovery that it cannot work in EPoC. Do we need to change anything?
  - Demonstrate added value of ultra-precise bandwidth allocation in EPoC when 1G+ link capacity is available upstream
    - In 10G-EPON, certain inefficiencies were allowed for to simplify specification, design and limit changes to 1G-EPON MPCP
    - Consider similar approach in EPoC. 90% efficient link is not worth it, if it costs 10x more to build than 80% efficient link

Changes to MPCP – path forward ? (cont’d)

- Proposals to change MPCP should be valued based on:
  - impact on efficiency (how much more bandwidth we get if we do it);
  - design complexity (how much more expensive it gets); and
  - timeline impact (how much longer it will take to get the project done).
- Backward compatibility with EPON is strongly desired:
  - Simpler equipment design, fewer bugs to root, quicker time to market
  - Reuse of existing platforms, e.g. DPoE, to manage ONUs and CNUs

Thanks!