

Multiple Profiles for EPoC

Andrea Garavaglia and Nicola Varanese - Qualcomm
IEEE 802.3bn Task Force Interim – Phoenix, AZ

Multiple Profile Architecture – Summary

- The function ***select()*** in the CLT Multipoint Transmission Control selects MAC Client grouping them over certain time interval – data or control packet can be selected, **control remain prioritized**
- Idle Insertion at CLT Control Multiplexer **matches the FEC of the profile** of the selected CNU, for correct parity insertion – de-rating can be applied according to the rate of that profile
- Idle deletion at CLT PCS transmitter **uses parameters of FEC and rate of the related profile**, for correct idle removal
- Idle insertion at the CNU PCS receiver is input by the FEC decoder of the corresponding profile, **inserting idle for FEC parity removal, un-decoded code words and XGMII rate match**

Why adaptive modulation and coding scheme?

- Two very important benefits are achievable by adapting the downstream transmission to the user channel conditions:
 - A. Increase of overall performance of the cable plant
 - B. Reduction of SNR margins and easier operations – the plant is automatically adjusted to the user condition
- These two benefits directly translates into CAPEX and OPEX savings/benefits for operators:
 - more data can be delivered in the plant due to optimized matching of channel conditions for groups of users and reduced SNR margins
 - automatic adjustments to slowly varying channel conditions achieved by assessing and grouping users, without human intervention

Capacity Analysis

SNR Estimates	Broadcast per plant		Broadcast per user group	
	Common MCS (reference)	Bit loading	Single MCS per user group	Bit loading per user group
Average spectral efficiency	7.2 bps/Hz	9.4 bps/Hz	9.5 bps/Hz	10.3 bps/Hz
Gain	--	30%	32%	43%
Peak User Spectral efficiency	7.2 bps/Hz	9.4 bps/Hz	10.8 bps/Hz	10.8 bps/Hz
Gain	--	30%	50%	50%

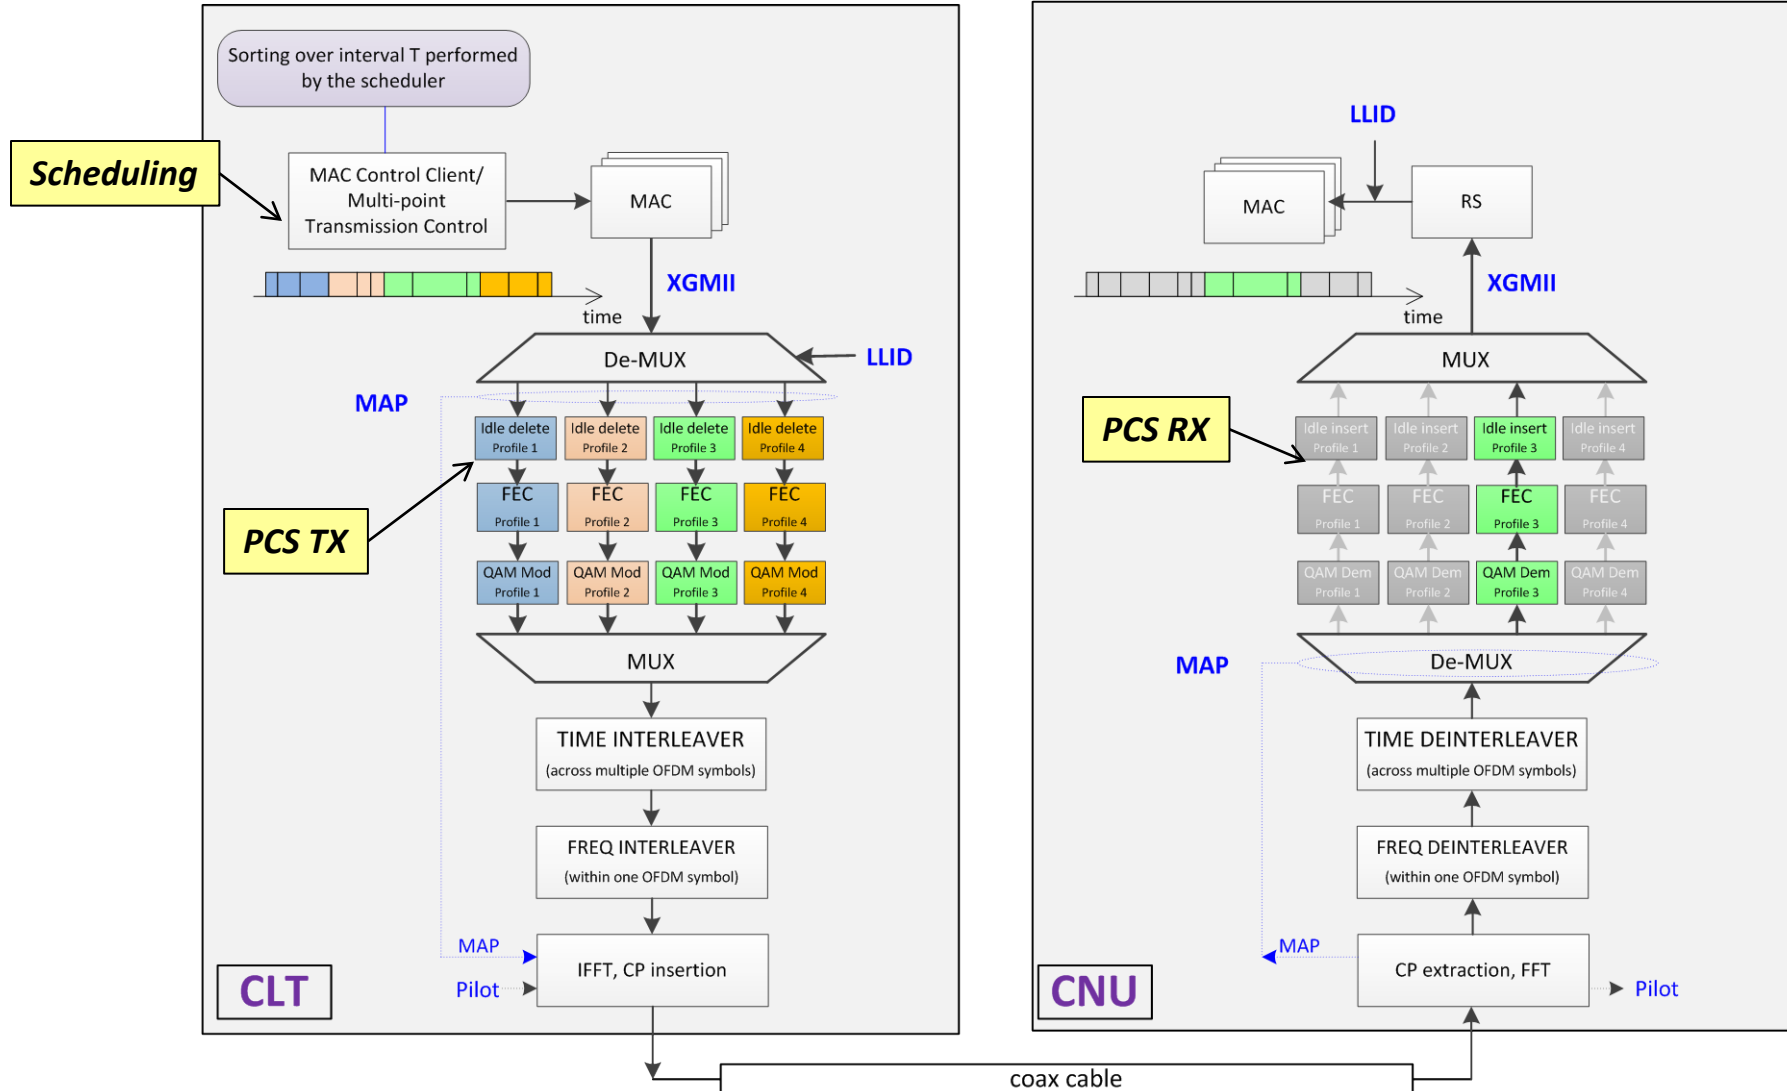
SNR Measurements	Broadcast per plant – Common MCS	Broadcast per user group – single MCS per user group
Average spectral efficiency	9.0 bps/Hz	10.5 bps/Hz
Gain	--	17%
Peak User Spectral efficiency	9.0 bps/Hz	10.8 bps/Hz
Gain	--	20%

OVERVIEW

Multiple Profiles - Overview

*Note: sorting only applies among LLID flows
Order of packets in one LLID flow does not change*

Note: CNU is expected to have one or two profiles, here represented a case with one profile (green)



Multiple Profile Design

- The downstream channel conditions are assessed and reported by each user during the PHY negotiation procedure, for the EPoC channel in use
 - The collected information is provided to upper layers via MDIO
- The CLT maintain a table in which users are grouped according to their supported Modulation and Coding Scheme (MCS), which is used by MAC Control to determine the supported rate and regulate idle insertion
- The PHY at the CLT processes the transmit packets on a per-group base, so that packets belonging to the same group can be encoded together
 - Stream based FEC applied to each group with large code words
 - MCS transitions optimized FEC efficiency by filling up the last code word as possible and eventually shortening it to free resources for the next group
 - A mapping of the OFDM symbol is included in the first part of the symbol (fixed robust MCS and reserved sub-carriers) to inform the receiver PHY in the CNU about how to decode incoming OFDM symbol and de-jitter packets
- The PHY at each CNU decodes the mapping of the incoming OFDM symbol and uses the information to process the related part in the PHY

MAC-Control Operations – key aspects

- For each packet to be transmitted, the MAC Control extracts from the MCS table the corresponding entry for the involved CNU
 - Depending on the MAC client which provides the packet, the CNU group is identified and the related MCS is read
 - Once the MCS is known, the corresponding data rate is calculated, which allows for the correct idle insertion after the packet (rate match)
- Note: This operation is the same if a single MCS is used in the plant, whereby here a table with more values is used (no complexity added)
 - Single MCS can still be supported as possible use case
- Number of MCS per plant: in practical deployments, a small number of groups (MCS) is sufficient to cover the range of typical SNRs
 - A range of e.g. 6-8 values can be included in the specification
 - In each plant, up to e.g. 4 groups can be configured at the time, choosing among those standardized values and a index/table is used to identify them

PHY Operations – key aspects

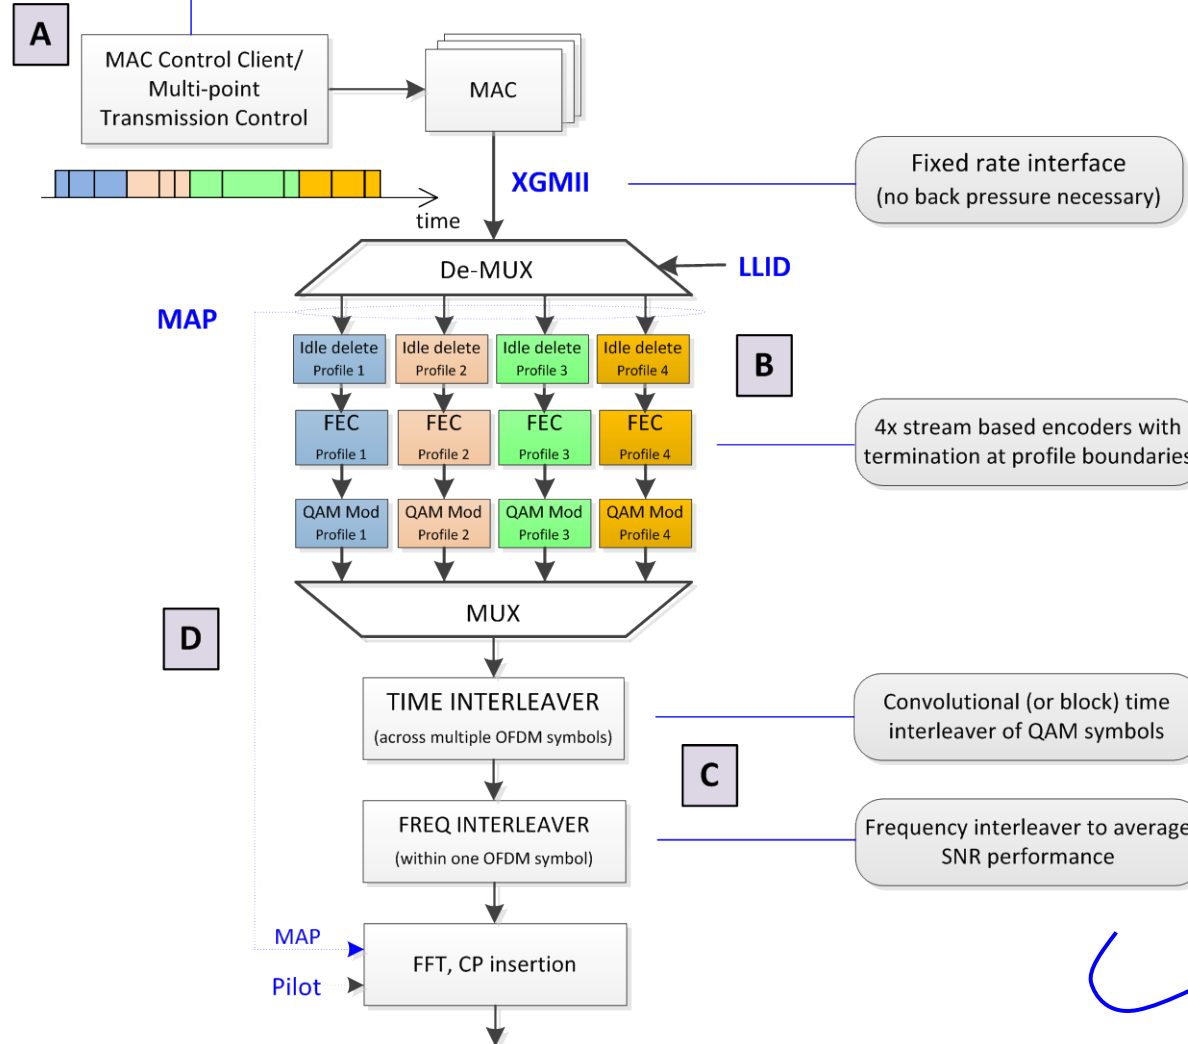
- To minimize FEC overhead, packets that are using the same MCS (i.e. for LLID pointing at CNU in the same group) are sent contiguously
 - Stream based FEC applied to each group, enable possible large code words
 - Group-to-group transition with optimized FEC efficiency by filling up the last code word as possible and eventually shortening it to free resources for the next group
 - MAC Control is aware of FEC formats to insert correct number of idles
 - The number of transitions is small (as the number of group in the plant is small)
- The grouping of packets can be performed
 - at MAC Control (e.g. via semi-static split of resources within a symbol, slowly adapting over time as traffic and groups change)
 - at PHY layer (e.g. distributing packets to the correct PHY processing chain based on their LLID – LLID maps uniquely into a MCS)
 - Both alternatives allows for single XGMII between MAC and PHY
- Broadcast traffic can be map into the group with smallest MCS, so that everybody can decode – similar concept could be used for multicast

DOWNSTREAM TRANSMISSION


Downstream operations – Grouping at MAC Control

Sorting over interval T performed by the scheduler

*Note: sorting only applies among LLID flows
Order of packets in one LLID flow does not change*



- Support of multiple groups via capability exchange
- Each group has fixed data rate (MCS), no backpressure needed
- Sorting applied at MAC Control over configurable interval – no impact on EPoC PHY delay budget

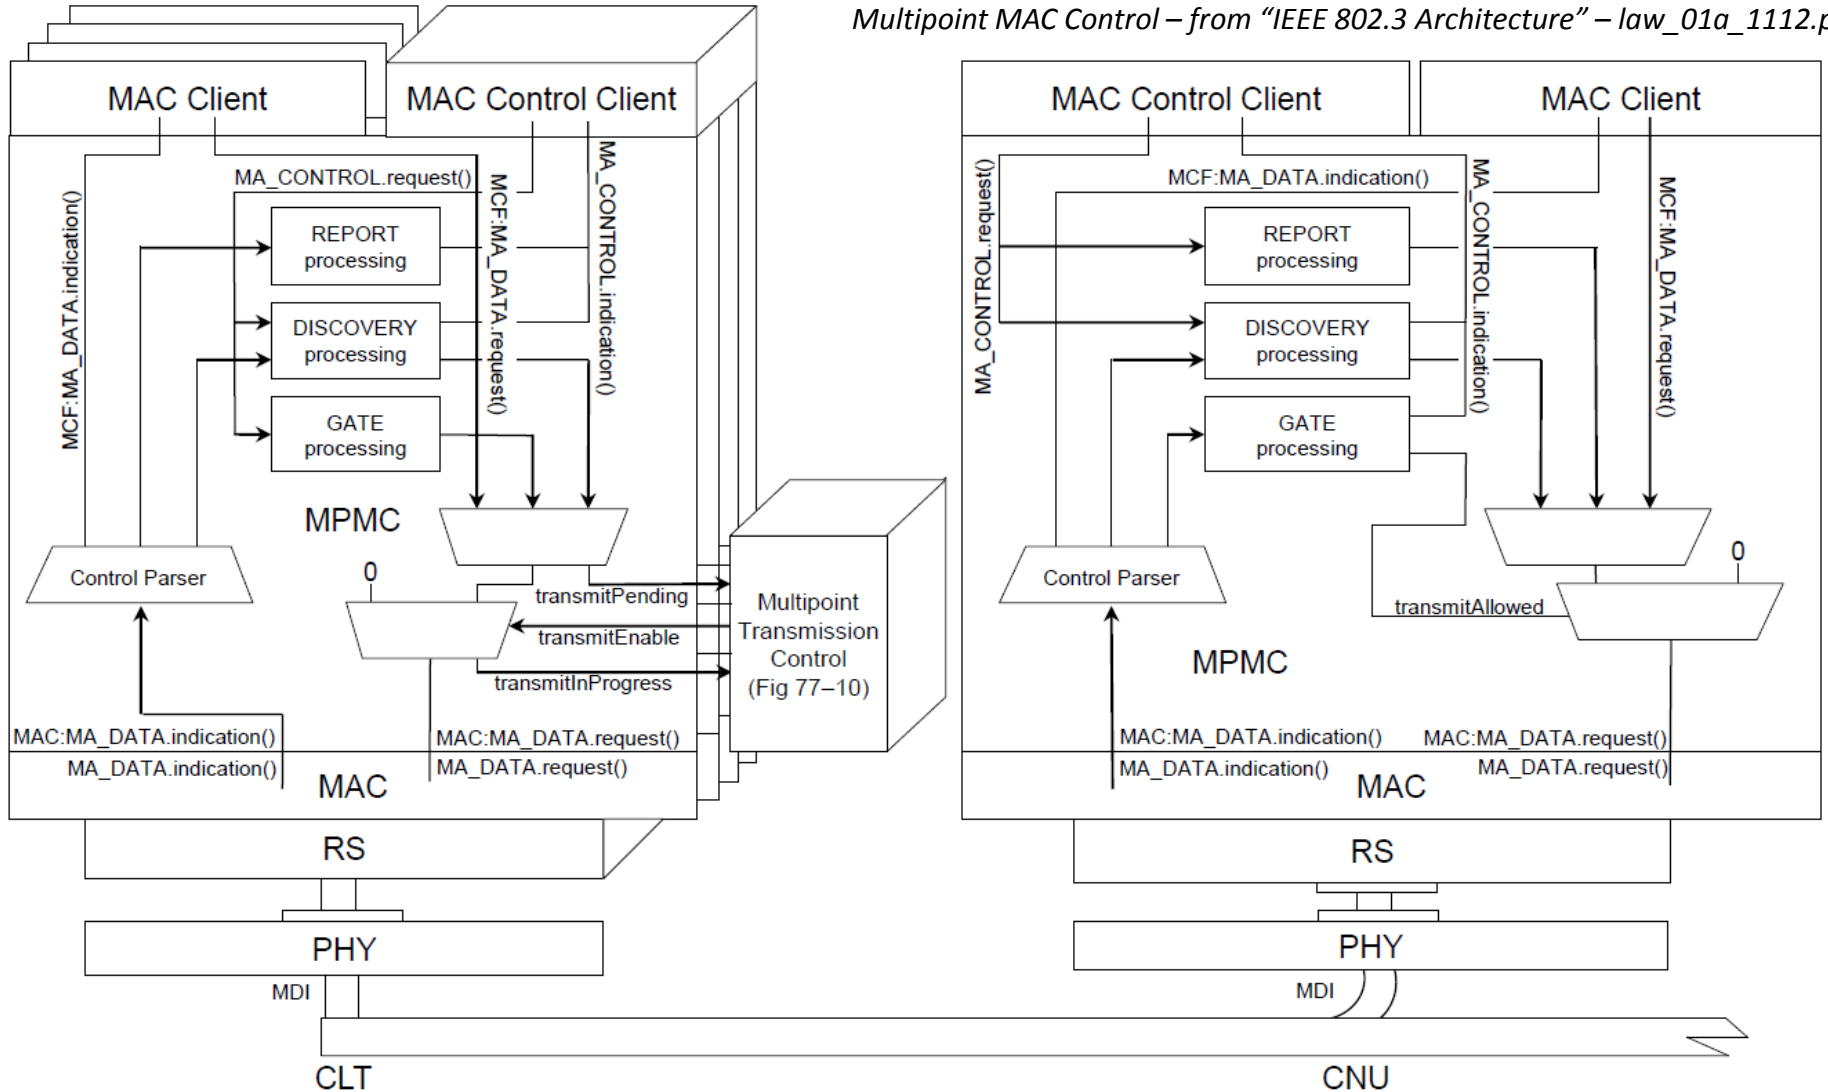


Downstream TX operations overview


- A. Concurrent downstream transmission requests from different MAC Client are coordinated by the Multi-Point Transmission Control entity in the CLT
 - packets are sorted by the scheduler based on their destination and grouped per profile, so they appear in a convenient order to be processed in PHY
 - idles are inserted after each packet for de-rating, upon need
- B. Within the baseband processing of each channel, packets are divided into N parallel flows (N is configurable, e.g. N = 1...4), one for each configured profile, and processed accordingly (idle deletion + FEC+ QAM modulation)
 - due to scheduler operations in A, packets are already grouped on a per profile base, so that the FEC encoder can be performed efficiently
- C. Processed data from each profile are multiplexed again together (same order is maintained to minimize possible time jitter) and the interleaving operations are performed (time and frequency) for each channel
- D. A MAP is created during PHY processing and inserted into the symbol to be conveyed to the receiver using a fixed robust modulation and coding scheme and reserving a small fixed number of OFDM subcarriers in a precise position

EPON – DS Multipoint MAC Control in CLT

Multipoint MAC Control – from “IEEE 802.3 Architecture” – law_01a_1112.pdf




EPON – DS Multipoint MAC Control in CLT

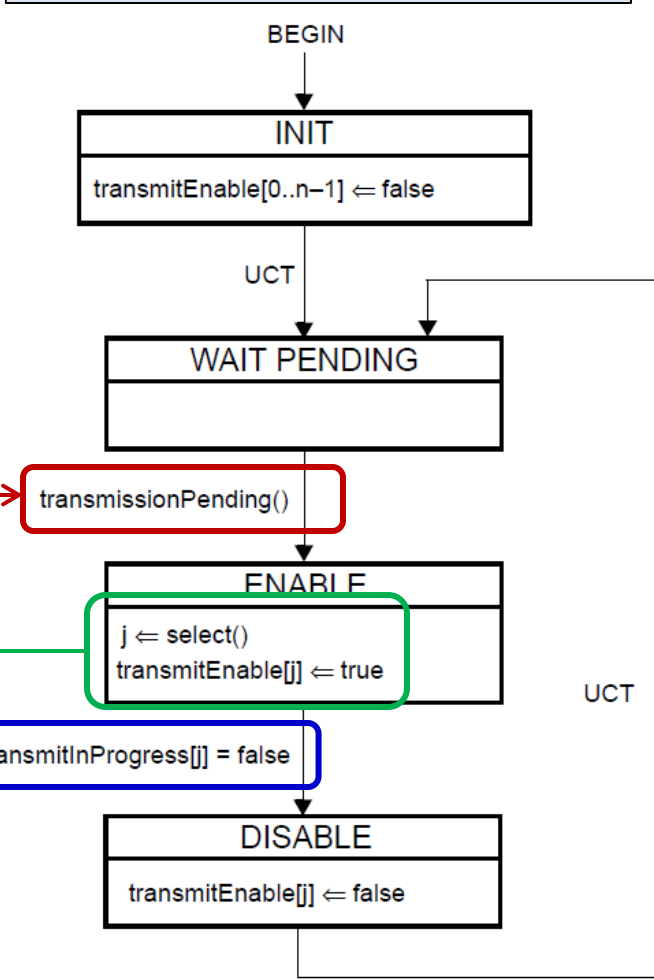


Scheduling – DS Multipoint CLT Control (77.2.2)

Control Multiplexer (fig. 77-13)



Multipoint Transmission Control (fig. 77-10)



Matching the FEC of the profile

Scheduling – DS Multipoint MAC Control in CLT

select(): This function selects the next Multipoint MAC Control instance allowed to initiate transmission of a frame. The function returns an index to the *transmitPending* array for which the value is not false. The selection criteria in the presence of multiple active elements in the list is implementation dependent.

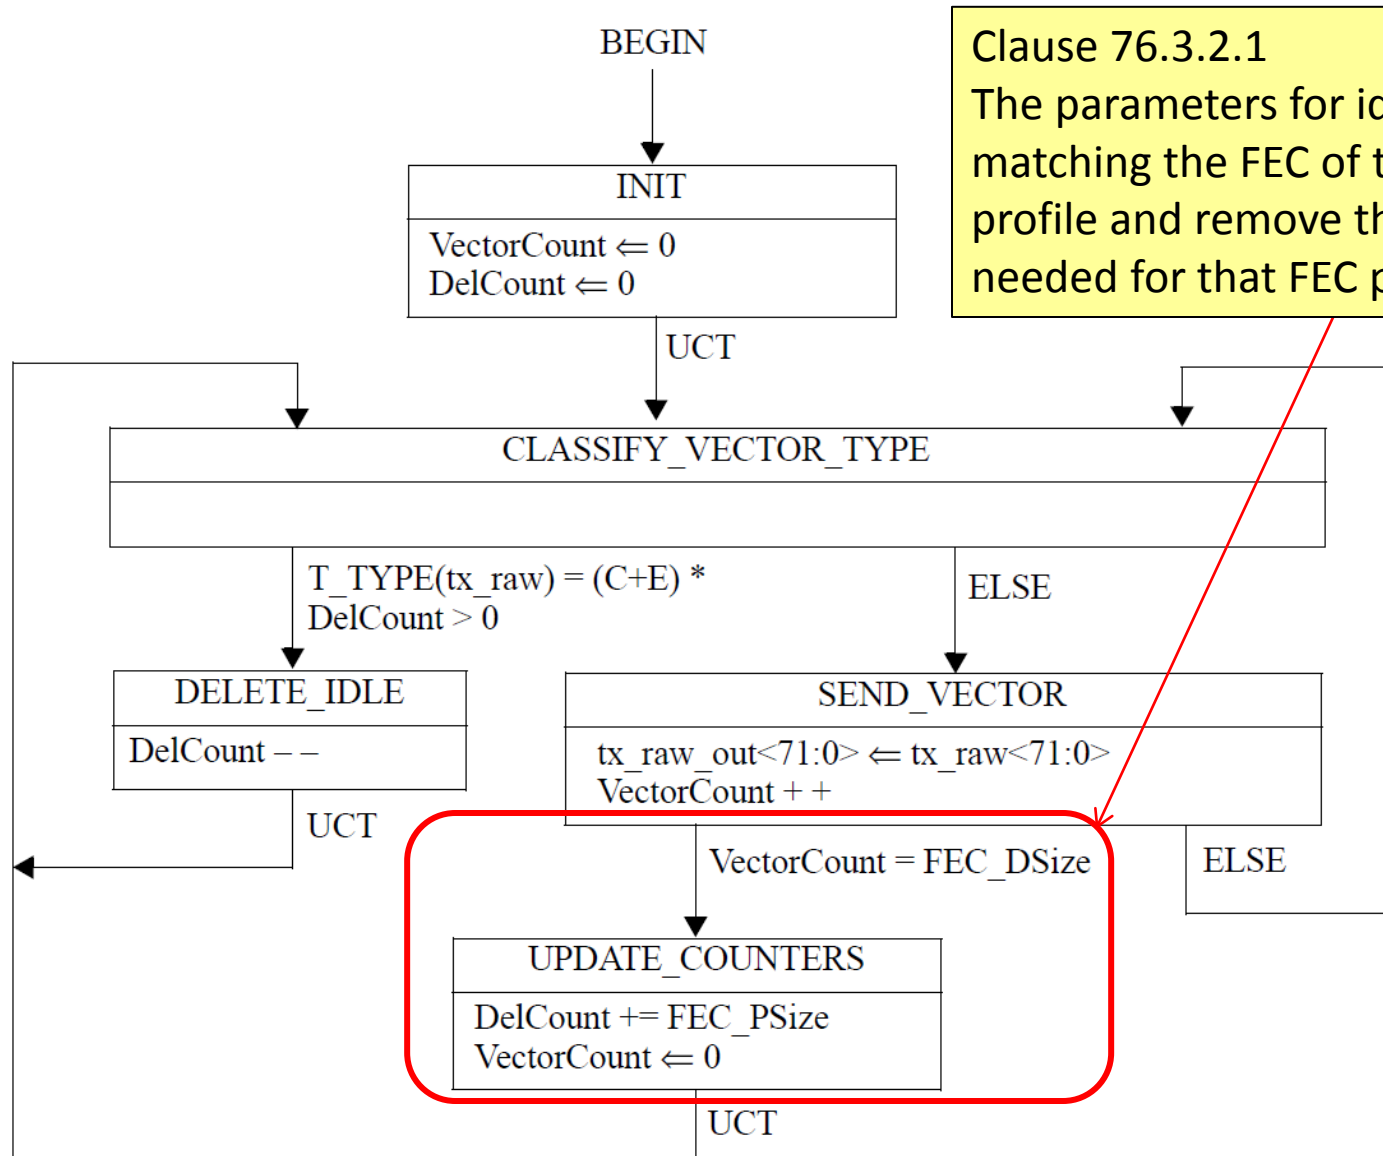
transmissionPending(): This function returns true if any of the Multipoint MAC Control instances has a frame waiting to be transmitted.

transmitEnable: This array contains one element per each Multipoint MAC Control instance. Elements of this array are used to control the transmit path in the Multipoint MAC Control instance at the CLT. Setting an element to TRUE indicates that the selected instance is permitted to transmit a frame. Setting it to FALSE inhibits the transmission of frames in the selected instance. Only one element of *transmitEnable* should be set to TRUE at a time.

transmitInProgress: This array contains one element per each Multipoint MAC Control instance. The element *j* of this array set to on indicates that the Multipoint MAC Control instance *j* is in the process of transmitting a frame.

- A → control frames of a MAC Instance are prioritized over data by *selectFrame()*
- B → the transmission token is given by the scheduler to user *j* ← *select()*
- C → at the end of the transmission, the token is given back to the MP TX Control

DS TX Idle deletion in CLT PCS



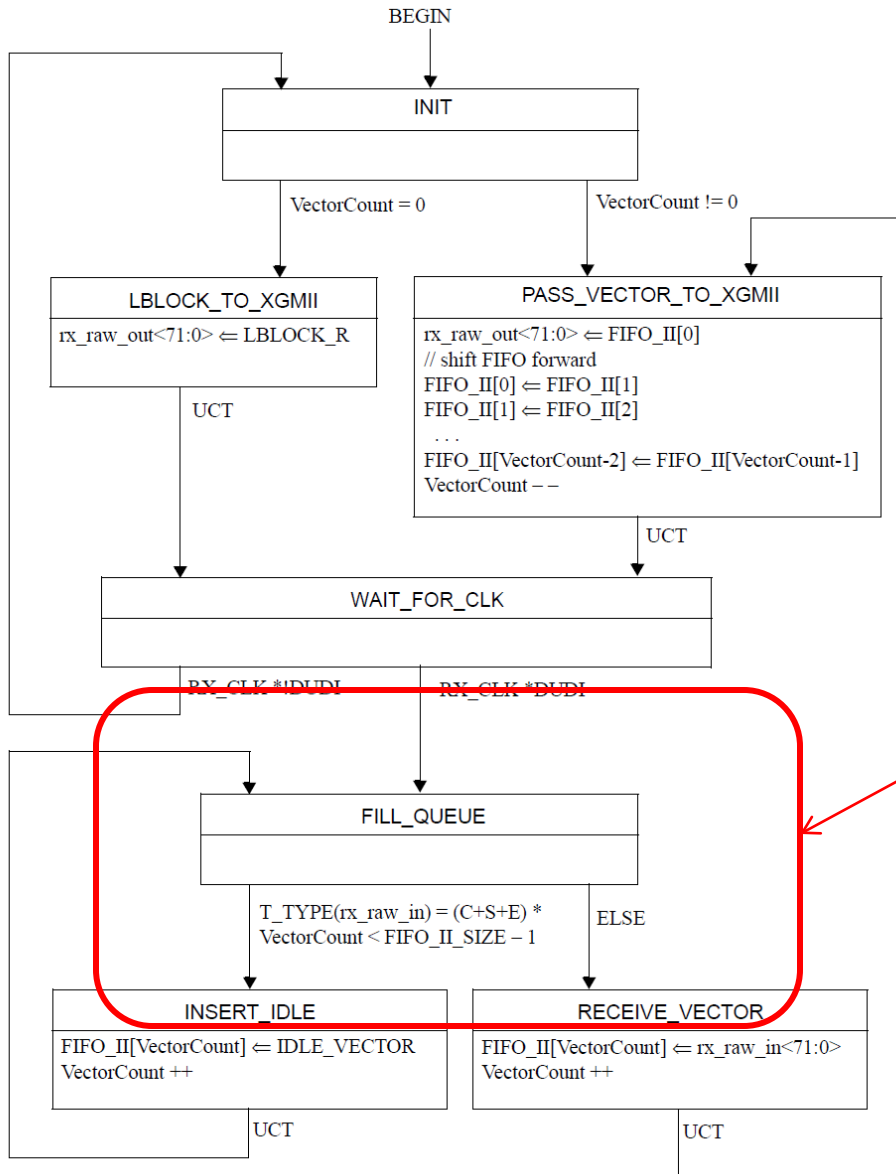
Clause 76.3.2.1
The parameters for idle deletion are matching the FEC of the corresponding profile and remove the number of idles needed for that FEC parity

DOWNSTREAM RECEPTION

CNU Reception

- The PHY at the CNU receiver first decodes the symbol map part in order to get instructions about what part of the OFDM symbol is of interest and where the corresponding code words are located
- All the code words for that group are decoded and all packets retrieved
 - PCS provides idle insertion (also) for packets which are not decoded
 - LLID is then used as usual to filter at RS what packets are for the CNU
 - No reordering of packet necessary within the same group
- Jitter handling:
 - In case the packet grouping is done at MAC Control, jitter is avoided

DS RX Idle insertion in CNU PCS



Clause 76.3.7
At the CNU receiver, idle insertion is performed at PCS

For each profile, the function of EPON could be in principle reused, input by the corresponding FEC decoder of the CNU

PERFORMANCE ASPECTS

Performance Evaluation

- 1) Capacity improvement with MMP (Slide 12).
 - Considered both bit-loading and MCS schemes.
 - The analysis provides an **upper bound** for Bit-loading as we employ Shannon's formula with a proper SNR margin.
 - The analysis provides a **lower-bound** for MCS as we are considering the worst SNR over the entire frequency range (e.g., we do not take frequency interleaving into account).

- 2) Achievable throughput with Bit loading and MCS (slide 13).
 - Single point-to-point link employing all the available bandwidth
 - More detailed prediction of the performance of Bit-loading Vs MCS (ReDeSign channel model).

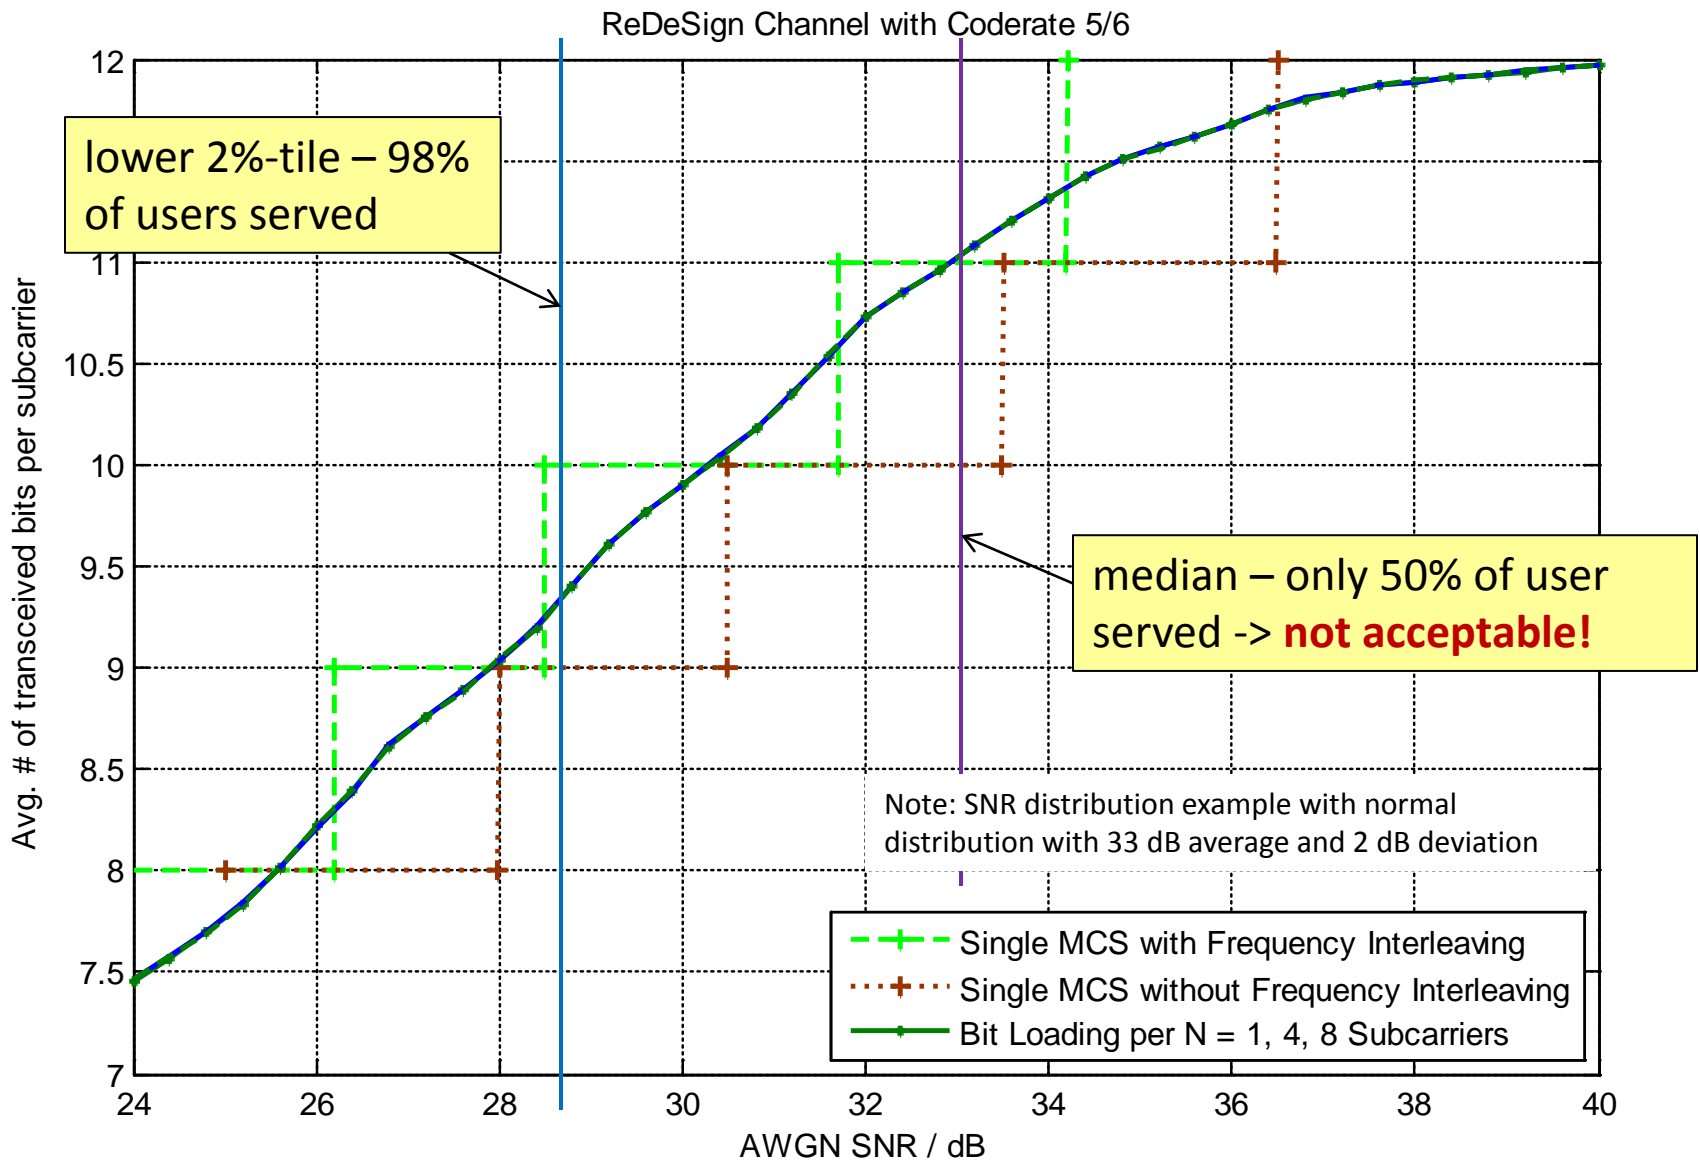
- 3) Throughput loss due to FEC codeword shortening (Slide 10).
 - Traffic model with a mixture of long and short frames.
 - A frame belongs to one of the available 4 profiles randomly with uniform probability.
 - Efficiency loss is recovered by grouping packet belonging to the same profile.
 - Performance can be improved by more advanced scheduling schemes.

1) Capacity Analysis

SNR Estimates	Broadcast per plant		Broadcast per user group	
	Common MCS (reference)	Bit loading	Single MCS per user group	Bit loading per user group
Average spectral efficiency	7.2 bps/Hz	9.4 bps/Hz	9.5 bps/Hz	10.3 bps/Hz
Gain	--	30%	32%	43%
Peak User Spectral efficiency	7.2 bps/Hz	9.4 bps/Hz	10.8 bps/Hz	10.8 bps/Hz
Gain	--	30%	50%	50%


SNR Measurements	Broadcast per plant – Common MCS	Broadcast per user group – single MCS per user group
Average spectral efficiency	9.0 bps/Hz	10.5 bps/Hz
Gain	--	17%
Peak User Spectral efficiency	9.0 bps/Hz	10.8 bps/Hz
Gain	--	20%

2) Throughput Comparison – Bit Loading vs. MCS



3) FEC Efficiency

FEC Efficiency Analysis - four profiles




- The potential FEC losses are limited to less than 3% and can be further optimized and tuned via sorting interval and code word size configuration


**ANNEX A –
CODE WORD SHORTENING**

Code word shortening

- In case of profile change, the last code word may need to be shortened to optimize the FEC overhead impact
- Compared to normal code word, a shortened code word has the same amount of parity bits as the regular code word, but less information bits. It is typically constructed by padding with zero the payload till reaching the regular size to compute parity, without transmitting those zeros
- Therefore a number of idles to account for full parity needs to be added at profile change in case the code word is not yet completed – **affects EPoC DS TX Idle Insertion and Idle Deletion state machine**

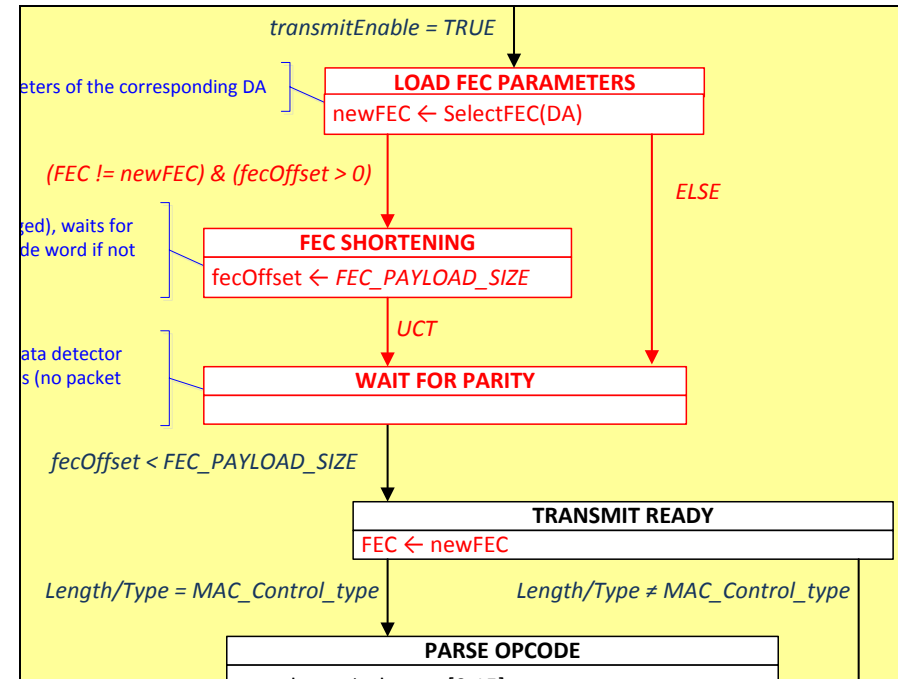
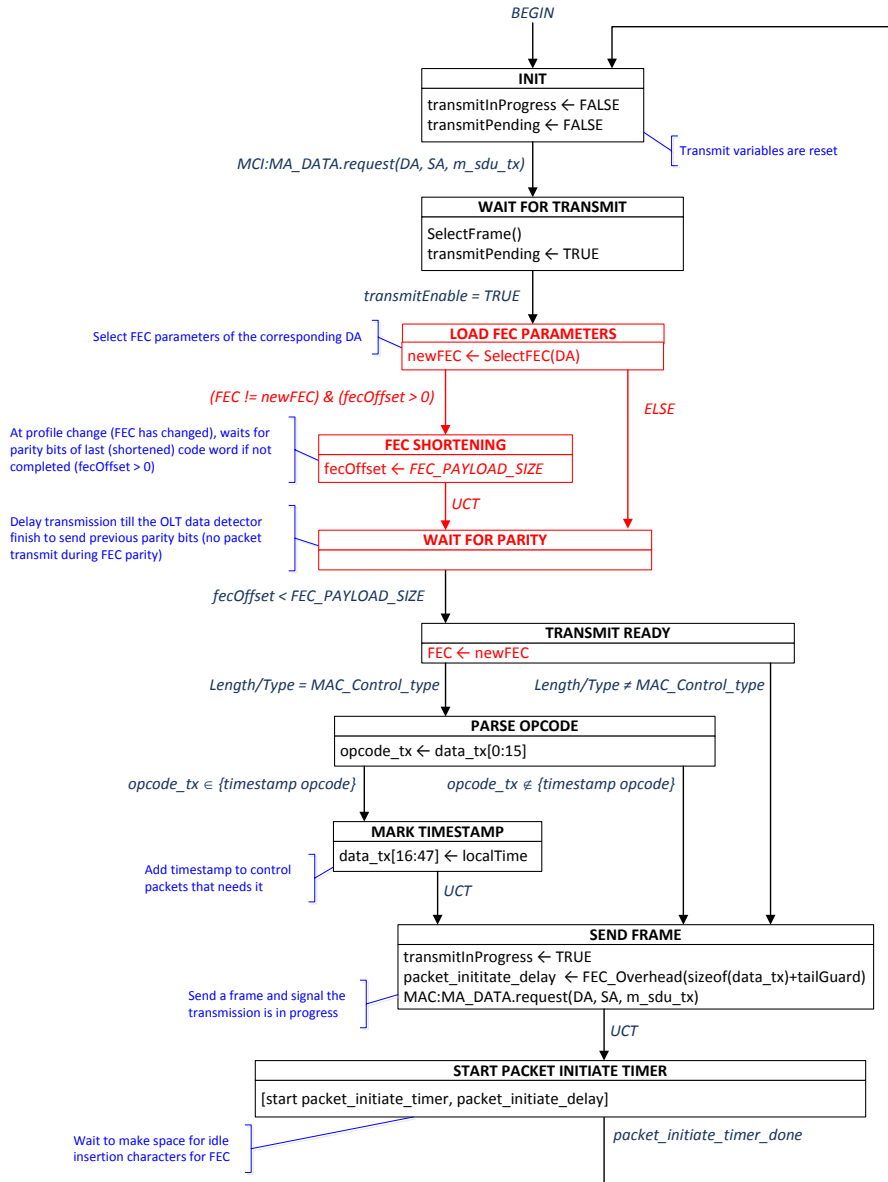


Regular Code Word



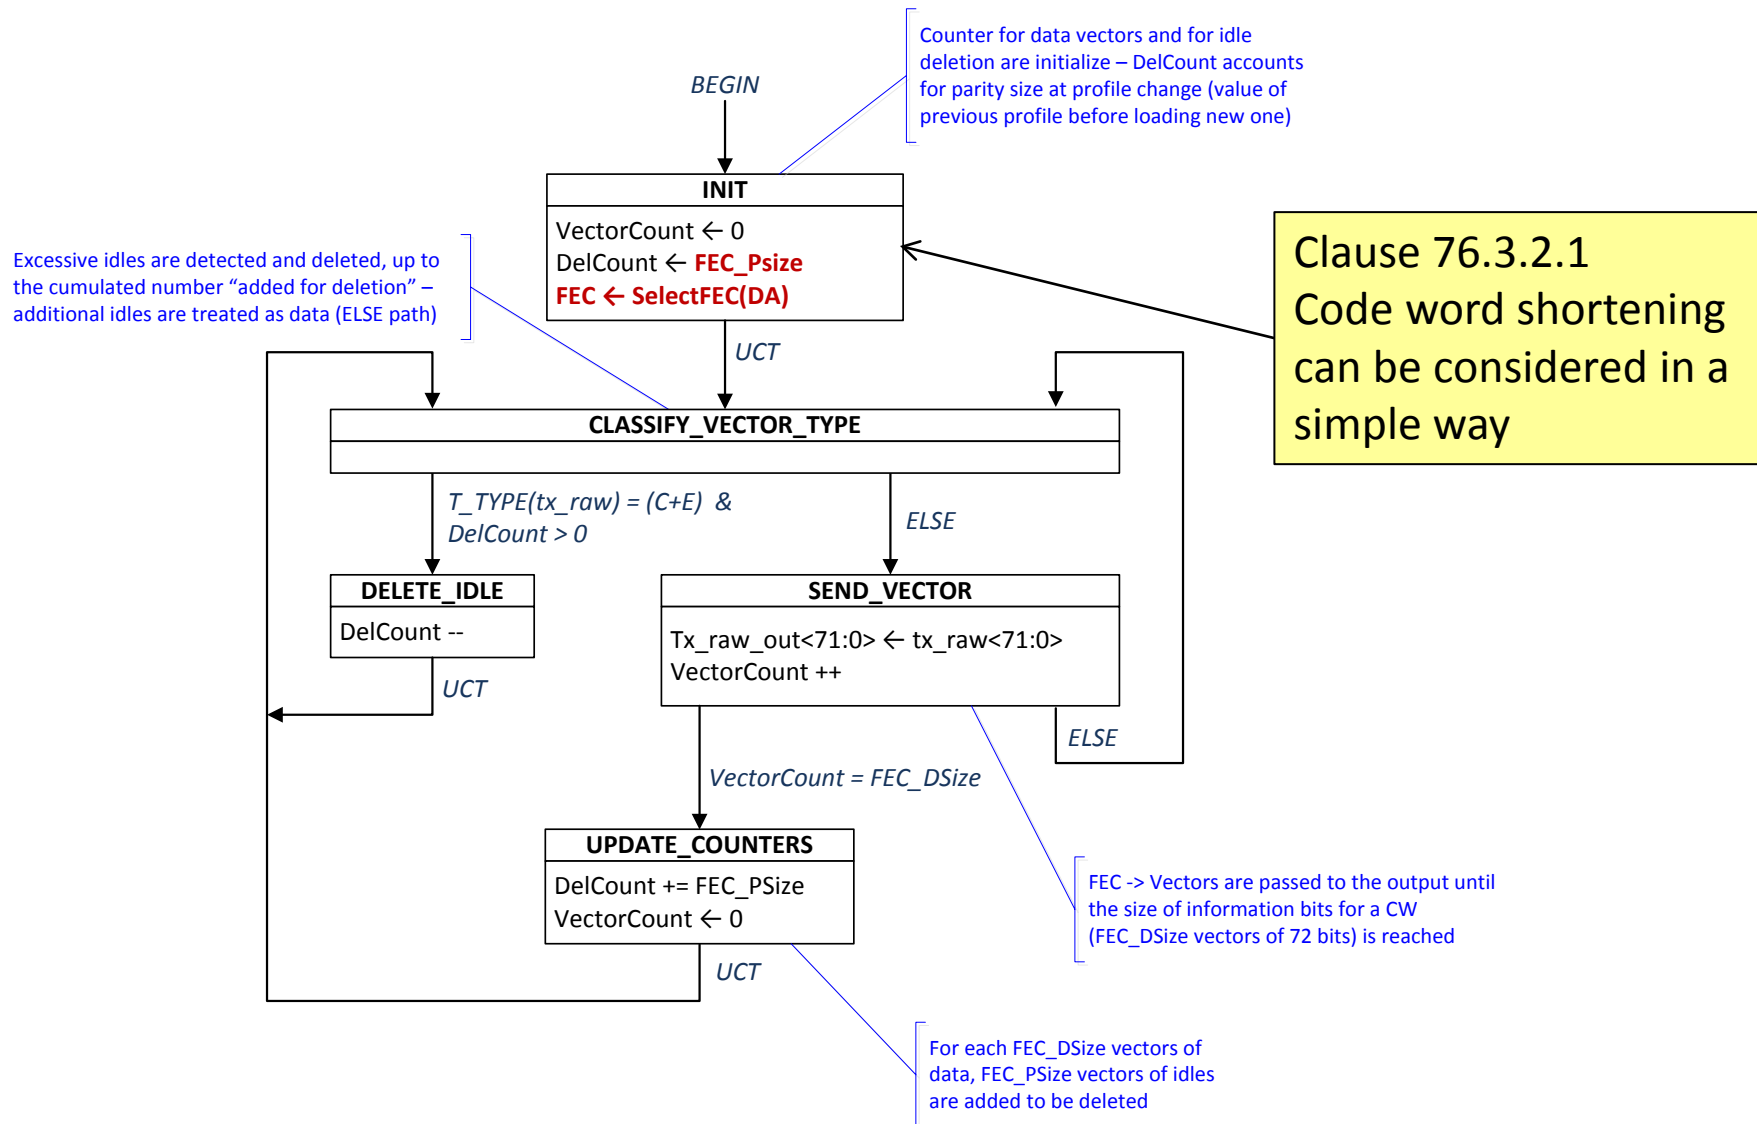
Shortened Code Word

CLT Multipoint MAC Control – state machine



Solution: set the variable *fecOffset* to the end of the payload in case of profile change: a space for parity is added automatically before being ready to transmit the first packet of new profile

DS TX Idle deletion in CLT PCS

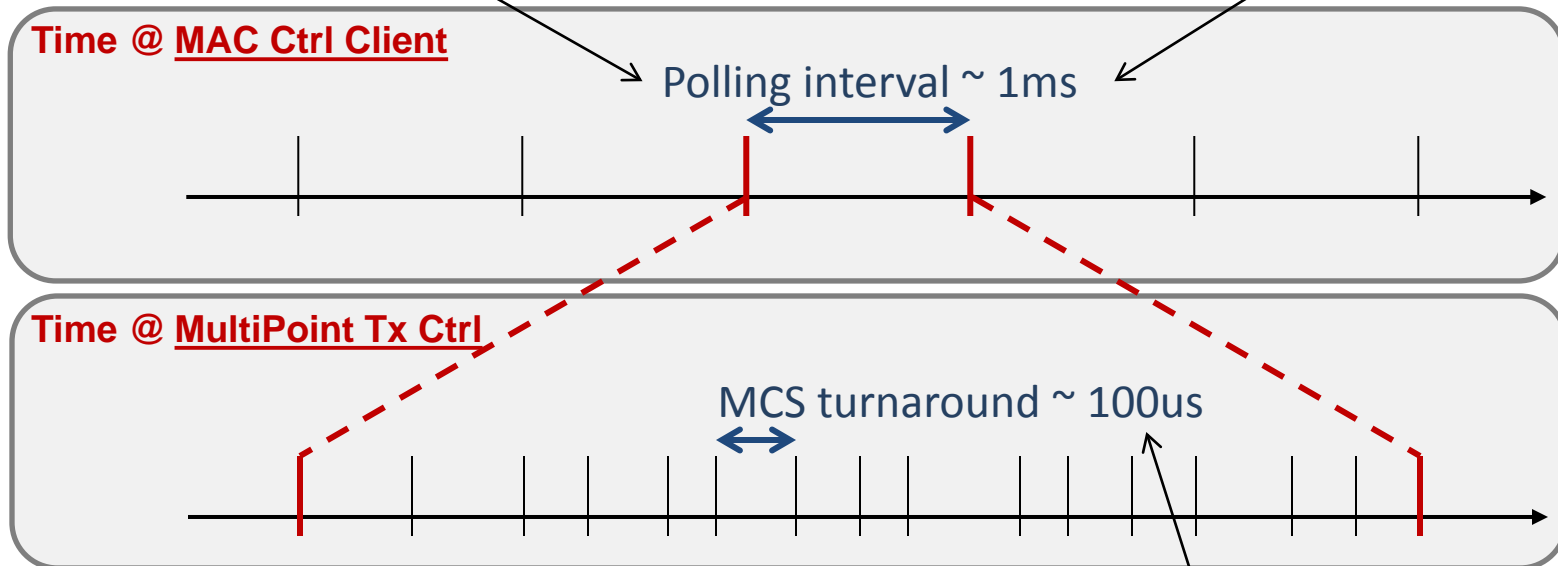


ANNEX B – GATE MESSAGE HANDLING

Scheduling – GATE message handling in CLT

The MAC Ctrl Client produces GATE messages in order to poll all CNUs regularly (polling interval).


The polling interval is typically in the order milliseconds (tens of OFDM symbols)



The scheduler within the Multipoint Tx Ctrl enables transmissions from MAC Clients according to the MCS they are associated with.

The turnaround time is the time the scheduler takes to cycle through all MCSs. The turnaround time may be variable (most general case), but in any case it will span only few OFDM symbols.

Scheduling – GATE/REPORT timeline



Reordering may cause 1st time delay of some GATE messages of some CNUs, but does not increase the GATE/REPORT cycle time afterwards (in steady state)