Practical Lessons from D3.1 Deployments and a Profile Management Application (PMA)

A Technical Paper prepared for SCTE•ISBE by

Karthik Sundaresan
Distinguished Technologist
CableLabs
858 Coal Creek Circle, Louisville, CO 80027
3036613895
k.sundaresan@cablelabs.com

Jay Zhu
Senior Engineer
CableLabs
858 Coal Creek Circle, Louisville, CO 80027
3036613312
j.zhu@cablelabs.com

Mayank Mishra
R&D Wired Intern
CableLabs
858 Coal Creek Circle, Louisville, CO 80027
3036619100
M.Mishra@cablelabs.com

James Lin
Lead Engineer
Kyrio/CableLabs
858 Coal Creek Circle, Louisville, CO 80027
3036613871
j.lin@kyrio.com

Table of Contents

Title **Page Number**

List of Figures

List of Tables

Title **Page Number**

Introduction

DOCSIS 3.1 OFDM/A Profiles provide a wide range of modulation choices that can be used to fine-tune the CMTS downstream and CM Upstream transmissions to get the best performance from the current network conditions. A well-designed, optimized set of modulation profiles allows a channel to operate more robustly against ingress noise and also enables an overall higher user throughput.

This paper will discuss the D3.1 Profile Management Application and how it is used by operators in deployment with D3.1 CMTSs and CMs to create Downstream profiles and Upstream IUCs/ Profiles. The paper will share our experiences from creating such an application. It will discuss the effect of noise on the choice of a profile assigned to a CM. We share results on how the number of FEC codewords corrected vary with the noise, and at different modulation orders. We also talk about the gain in the network capacity seen by smart Profile creation algorithms versus using a flat profile and simply assigning CMs to the least common denominator profile which fits the CM.

The paper will also describe MAC layer state machine and interactions between the CM & CMTS when a profile fails. The interaction between CM-STATUS messages from a CM (for flagging failures and recovery on a profile) and the CMTS response to that message in changing the profile used, results in “Profile Flapping”. This paper recommends on how to design CM and CMTS MAC Layer settings to make the system be robust when these interactions take place in the D3.1 system.

D3.1 OFDM/A channels can have interference in parts of the channel and different modems experience this differently. Deploying well designed profiles for each channel will decrease the number of errors seen on the channel, reducing trouble calls. It could also unlock a solid 200~400 of Mbps of capacity gain on each channel.

1. OFDM and OFDMA Profiles

DOCSIS 3.1 specifications introduced features that leverage the OFDM-based PHY layer, including variable bit loading, and the ability to define multiple modulation profiles on downstream and upstream channels. Other new features include the ability to measure the quality of a downstream channel and test out the profiles in use, and features like upstream probes to measure the quality of the upstream OFDMA signal. The new MAC management messages support this on a per CM basis. There are also extensive additions to important operational items within proactive network maintenance (PNM) which enables measurement of various physical layer metrics and exposes that data to the operators.

The configuration, initiation logic and compute processing needed to optimize some of these functions Downstream Profile creation, Upstream IUC/Profile creation, are not defined in the DOCSIS 3.1 MAC and PHY specifications. This allows such functionality to be moved out of a CMTS and implemented as a Profile Management Application (PMA) running outside the CCAP. Here the idea is to move the profile creation process as an application external to the

CMTS. The PMA can communicate with a data lake and the CCAP to gather the needed information, process the data, and make intelligent decisions to set up the CCAP as needed.

To leverage the new OFDM/A physical layer to its maximum benefit, different subcarriers use different modulation orders. Optimizing the downstream/upstream profiles allows a downstream/upstream channel to be able to operate with lower Signal-to-Noise Ratio (SNR) margin, potentially allowing a channel to operate at an overall higher throughput. The logic to achieve this can be external to a CCAP and enable innovation. For a cable operator, it allows uniform operation of such algorithms across different CCAP platforms.

[D31PMA-INTX16] describes methods for designing OFDM/A profiles and choosing the appropriate modulation orders for a profile. It answers the questions around which profile is appropriate for a CM and what is the optimal set of profiles to use across the an OFDM/A channel for a given set of CMs.

A D3.1 CM supports two or more OFDM channels, each occupying a spectrum of up to 192 MHz in the downstream. The OFDM signal is composed of: Data subcarriers, Scattered pilots, Continuous pilots and PLC subcarriers. A modulation profile consists of a vector of bit-loading values, an integer value for each active subcarrier in the downstream channel. The modulation orders range from 16-QAM to 16384-QAM, the range of bit-loading values is from 4 to 14 (skipping 5); however, it is expected that very low bit-loading values, 7 or less, will be used very infrequently since most plants support 256 QAM today, but those will likely be in use in the roll-off regions.

Each CM will support and can be assigned up to four modulation profiles, including Profile A (used for broadcast frames), an optimized profile for the CM's unicast traffic, and possibly two additional profiles that could be used for multicast traffic.

A CMTS on each OFDM channel, per the DOCSIS specifications, needs to support up to 16 profiles. In the short term, CMTSs support 3~4 profiles per channel and many CMTS equipment assign all the profiles they support to all the modems. Today these profiles are flat, i.e. the same modulation order across the whole channel. In the long term as CMTS vendors support more than 4 profiles per channel, they will assign one profile A and one optimized profile to a CM. There may be a third profile assigned as a fall back profile between the optimized profile and profile A.

DOCSIS 3.1 Downstream

2. Noise Characteristics on an OFDM Channel

Several different types of noise have been identified in the field creating challenges with profile flapping and partial service flapping with the OFDM channel. These includes LTE ingress, Sweep Generator, Suckouts, Channel Roll-offs, Tilt etc. Examples of these are shown in the figures below. Any of these noise ingress can cause uncorrectable codewords on the channel, depending on the severity. All the below graphs show the RxMER level measured at the CM along the y-axis, and the frequency of the OFDM channel in MHz, along the x-axis.

2.1. CMs with clean OFDM channels

Depending on the plant, many CMs have quite good and clean RxMER levels across the whole D3.1 OFDM channel. This example is a 96 MHz channel starting at 696 MHz (in 850 MHz plant).

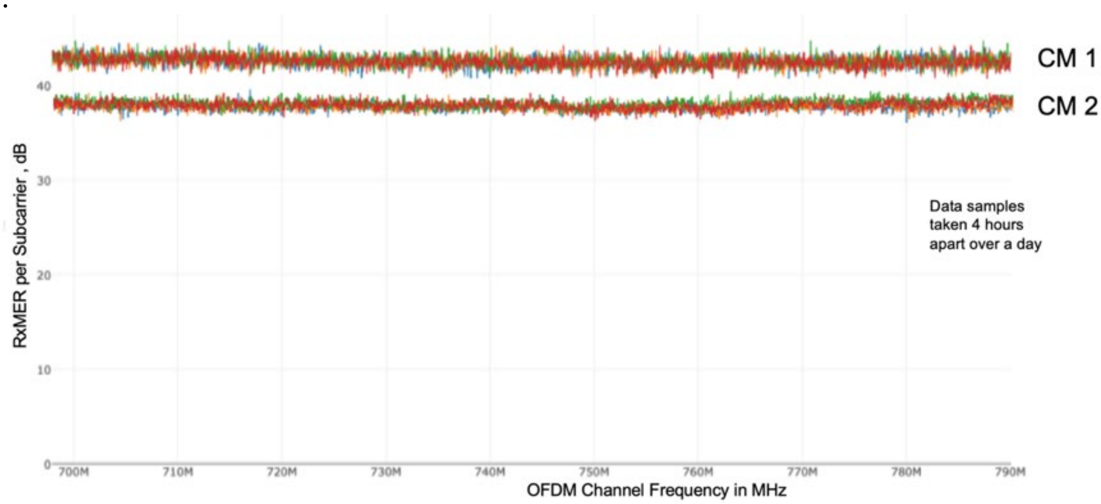

Figure 1 – Examples of Clean RxMER from CMs

2.2. LTE Ingress

The below picture shows two channels, a 96 MHz OFDM channel from 660-756 MHz, and second one from 700-796 MHz. Both of these channels clearly show LTE ingress noise in the CM's DS RxMER, as detailed later in this paper. As seen the LTE ingress is at a specific location in frequency but the ingress level is highly time variant. The bottom graph in this figure shows samples every 30 seconds, and the ingress level is always changing.

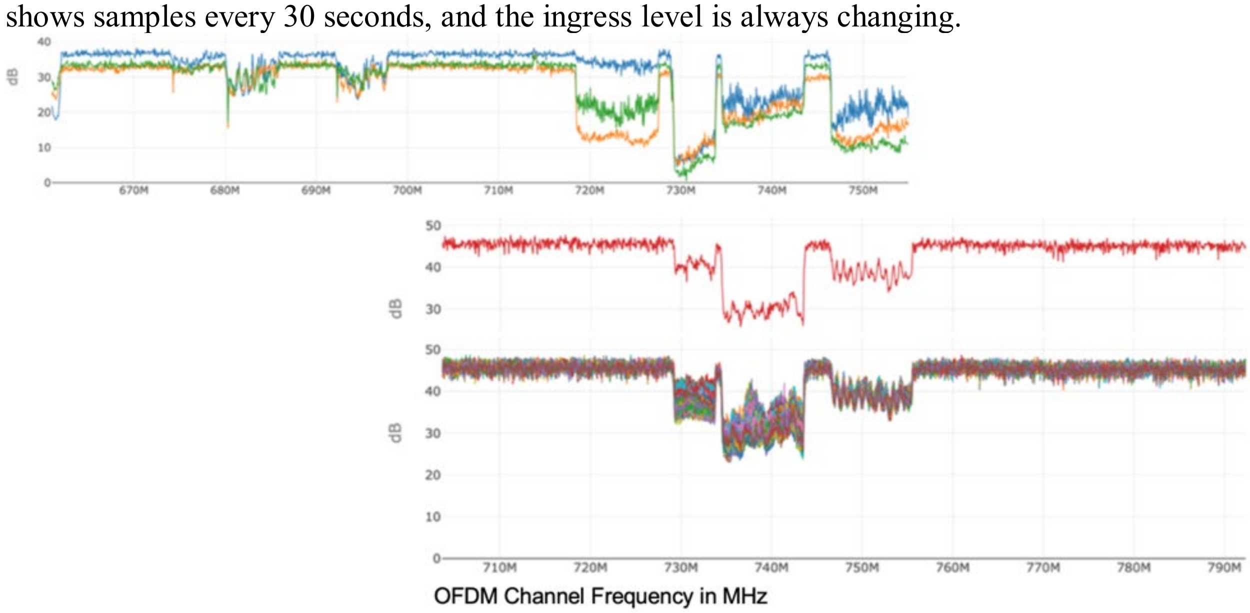

Figure 2 –LTE Interference

2.3. Interference from a Sweep Generator

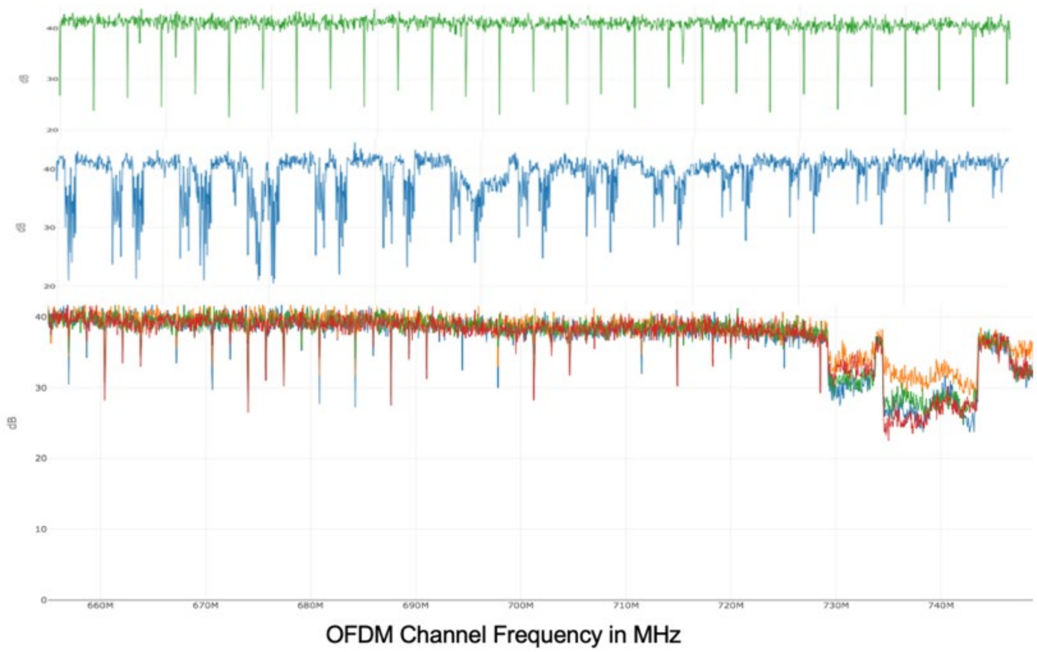

Figure 3 –Sweep Generator Ingress

2.4. DS Roll off

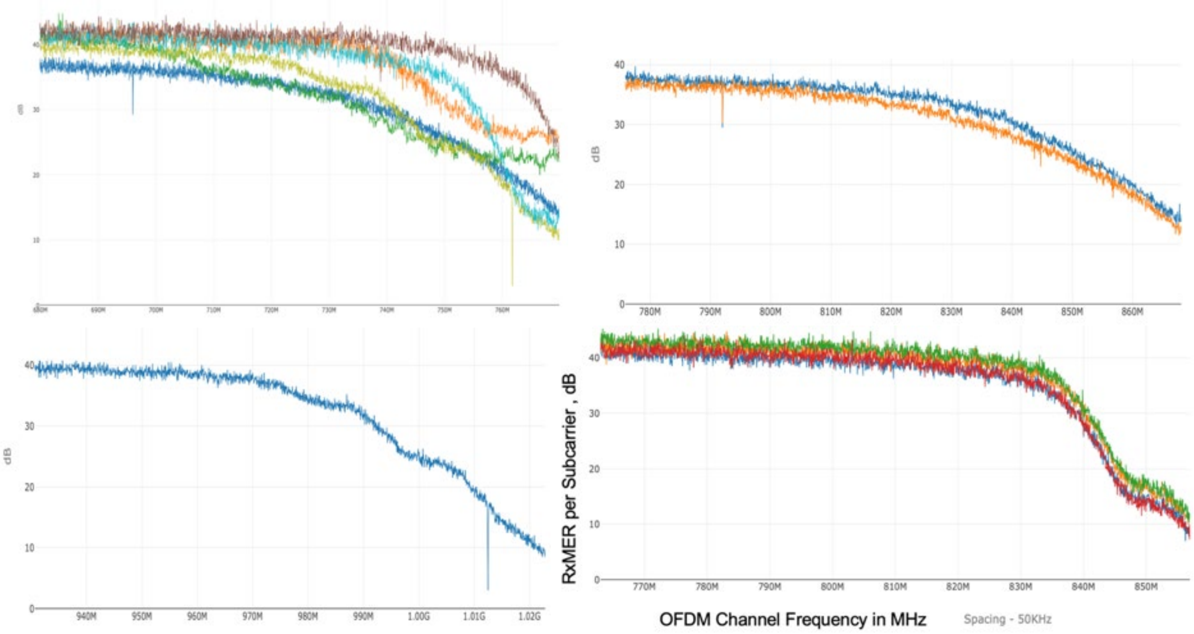

Figure 4 – OFDM Channels in the Roll-off

2.5. Suckouts

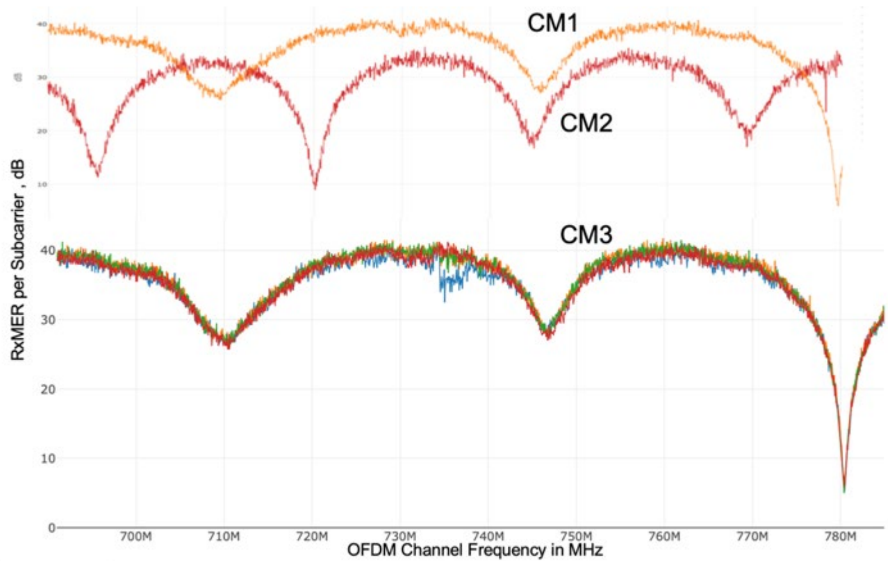

Figure 5 – Examples of Suckout

2.6. Standing Waves

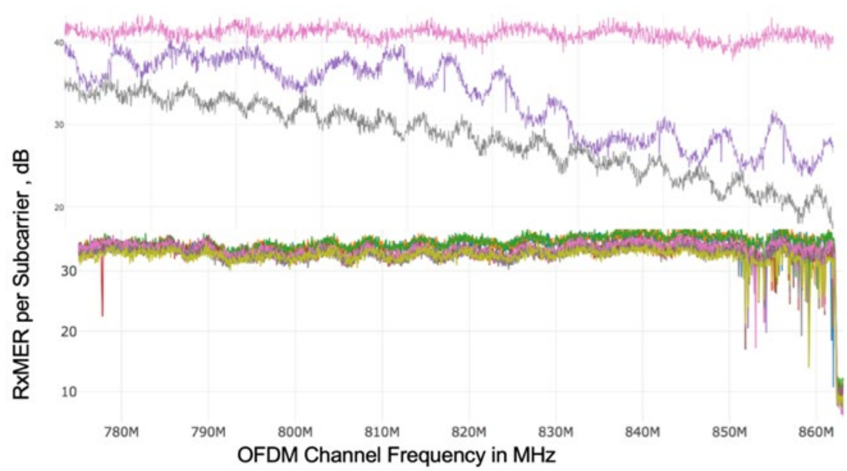

Figure 6 – Examples of Standingwave

2.7. Downstream Tilt

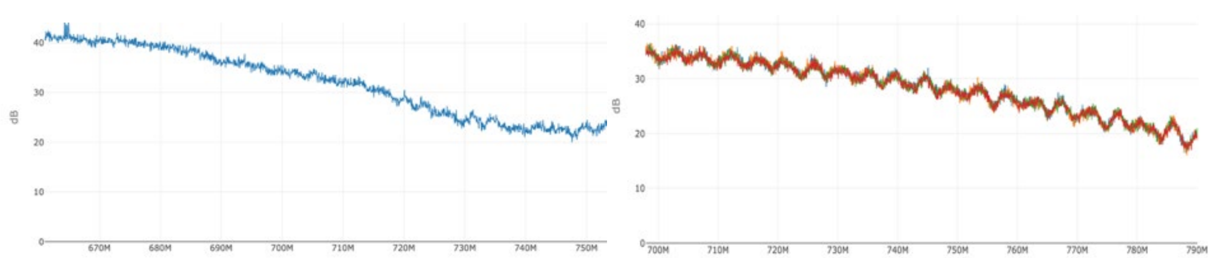

Figure 7 – Examples of DS Tilt

3. Downstream Forward Error Correction behavior

Forward error correction (FEC) is a way of adding redundancy to messages so that the receiver can both detect and correct common errors. The D3.1 system uses an outer BCH coding and an inner LDPC coding. The devices support the 8/9 code rate for the codeword (16,200 bits). Codeword shortening is used when there is insufficient data to fill complete codewords and to achieve strong burst noise protection.
A PMA needs to be able to map RxMER values in the Downstream, to an appropriate QAM level, when creating a profile. These mappings are defined in [PHYv3.1] and summarized below.

Table 1 - DS RxMER to QAM Level mapping

Downstream Constellation/ Bit Loading	DS MER (dB)
16 QAM	15.0
64 QAM	21.0
128 QAM	24.0
256 QAM	27.0
512 QAM	30.5
1024 QAM	34.0
2048 QAM	37.0
4096 QAM	41.0
8192 QAM	*46.0*
16384 QAM	*52.0*

3.1. Lab testing of DS FEC behavior on D3.1 equipment

We wanted to understand the performance of the DS FEC on real D3.1 equipment, in an effort to understand at what points will an operator start seeing failures in the system. We are testing these with 2 different CMTSs and 2 different CMs. We hope to expand the testing to include other devices. The test was run on a 96 MHz OFDM channel at 300 MHz and 702 MHz

3.2. Baseline test (no noise)

This test checks RxMER at the CM and runs downstream traffic from the CMTS to the CM:

1. Send traffic to modem starting at 200 Mbps for 30 seconds.
2. Increase traffic in steps of 100 Mbps until Traffic rate of 500 Mbps.
3. Repeat steps 1 and 2 for the profile configured at each QAM level between 256QAM – 512QAM-1024QAM-2048QAM-4096QAM.

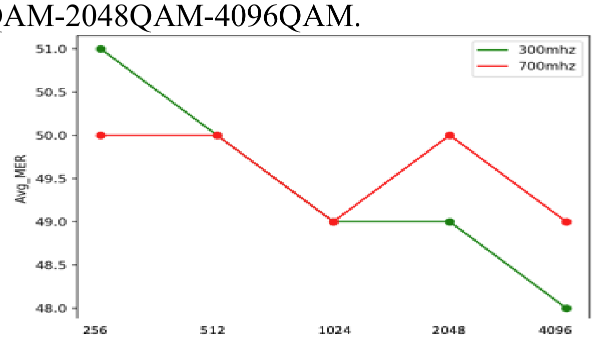

Figure 8 – Average D3.1 RxMER for a clean channel when carrying traffic

For a particular flat profile (or modulation order), there was little significant difference between the average RxMER values across the channel at different traffic rates both at 300 or 700MHz. However, the difference between the RxMER values measured when using different profiles at different modulation orders varies a bit. This was something we couldn't readily explain as to why the RxMER measurements could be different. In this baseline test there were no corrected FEC codewords for any of the 5 modulation profiles.

3.3. Test to discover failure points, noise across entire channel

The next step was to determine at what points would CMs start seeing DS codeword errors. The idea here is to increase the noise floor on the channel (AWGN) and see how the system performs. For each modulation order: QAM 256, 512, 1024, 2048, 4096, the goal was to identify the average RxMER of the channel at which

- the first corrected codeword is seen
- 100% corrected Codewords are seen
- the first uncorrected Codeword is seen

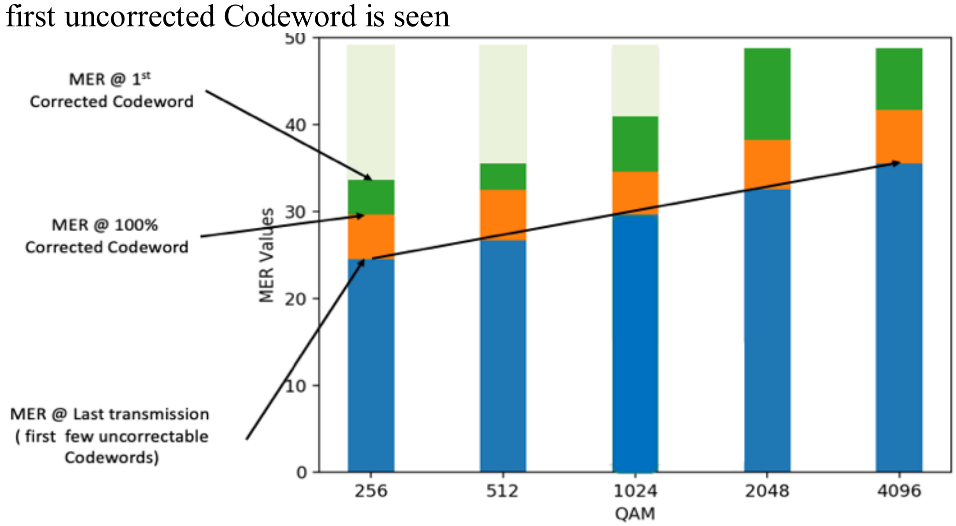

Figure 9 – Average D3.1 RxMER levels for Correctable and Uncorrectables

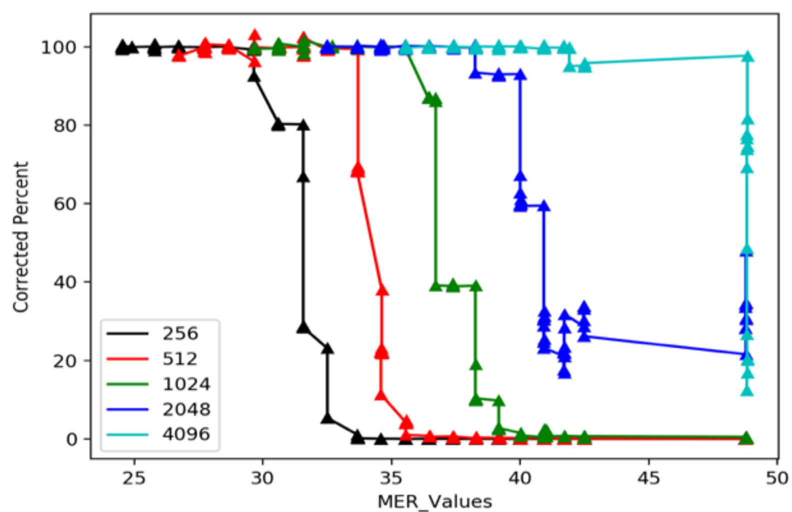

Figure 10 – Alternate view Corrected Codeword percentage vs RxMER

We identified the different code word corrected codeword failure levels at the CMs, and the corresponding MER values tied to those failure point. As seen in the figure there is an inverted S-curve growth in corrected codewords as noise increases. As expected the lower the modulation order, the more noise is needed to get to the first corrected, 100% corrected and first uncorrectable codewords.
Most CMs stopped receiving data on the profile as soon as the first few uncorrectable codewords are seen. This is because the CM sends out a CM-STATUS message and the CMTS reacts by downgrading the profile or putting the CM in partial service.

3.4. Test QAM Level tolerance to noise

This test was to determine if higher modulation orders are progressively more susceptible to noise. The test methodology was to send traffic to the CM at 300 Mbps. We then insert AWGN noise with a MER of 50db over entire channel. The noise was increased in steps of 1db, each noise level was maintained for 15 seconds. The test continued until the modem reported 100% corrected codewords, for each QAM level between 256– 4096.

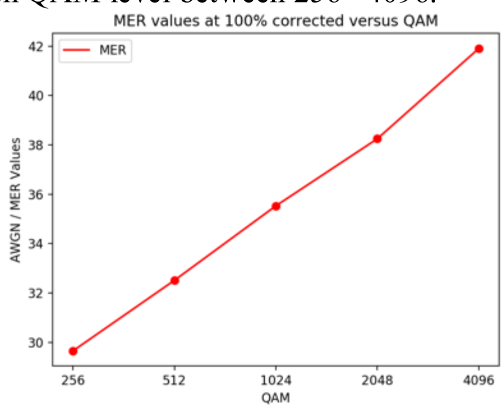

Figure 11 –D3.1 RxMER levels for 100 % Corrected vs QAM level

The RxMER values that yield 100% corrected codewords grows linearly with higher QAM.

3.5. Test tolerance to narrowband noise (AWGN)

The test is to determine the performance of downstream FEC when injecting narrow band noise to the OFDM channel. In this test, the downstream traffic is set to a rate of 500 Mbps. Different noise widths including 1MHz, 12MHz, and 25MHz are injected into the channel. The following figure shows the OFDM downstream RxMER with different noise output attenuations when injecting 1MHz noise:

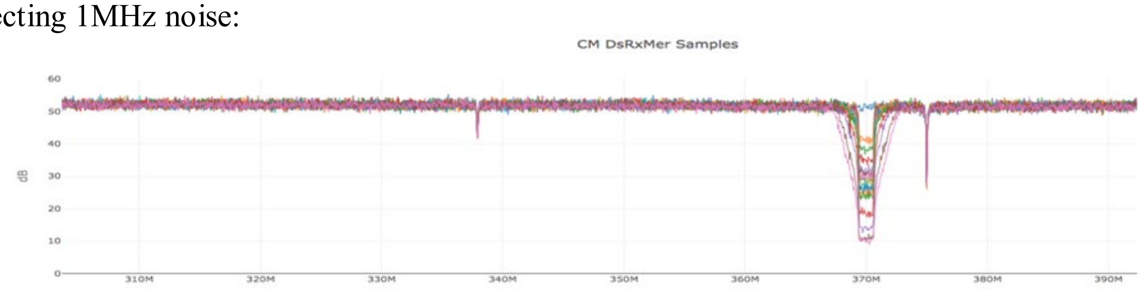

Figure 12 –D3.1 Noise Injection test

The noise injection test results are put together with the baseline test result (full-band AWGN noise) for comparison. The following figure shows the results with 256-QAM as the data profile modulation order:

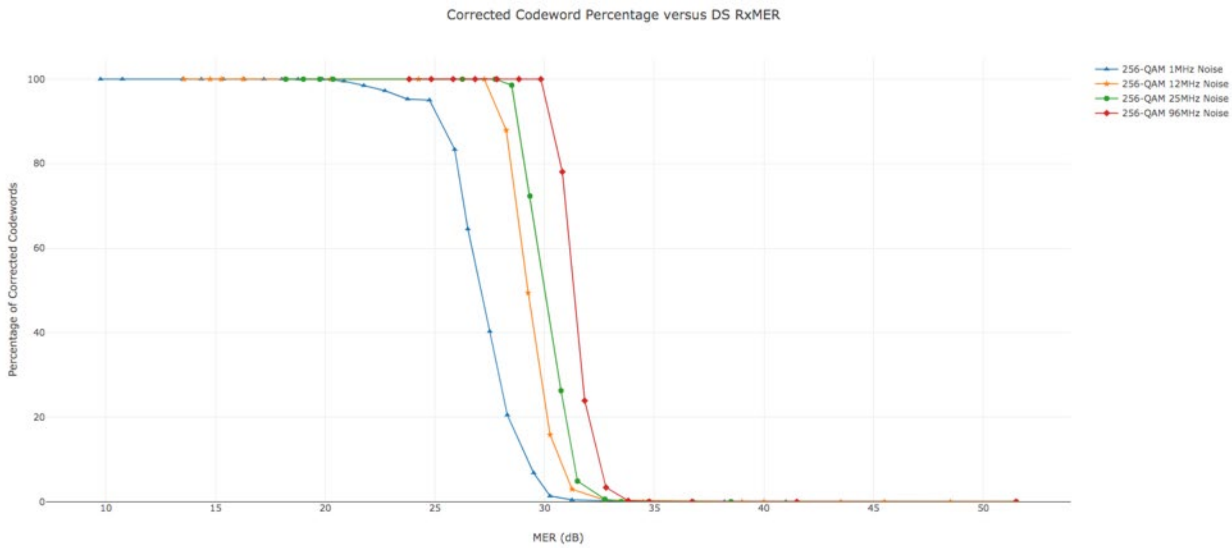

Figure 13 –D3.1 Noise Injection test results

As the noise width increases, the LDPC decoder has to correct more codewords. Wider noise injection causes the LDPC decoder to reach 100% corrected codewords at higher RxMER levels (lower noise power) than narrower noise injection. Also, a wider noise width creates uncorrectable codewords at higher RxMER levels and the CM goes into partial channel state on the OFDM channel. This can cause the CMTS to move the CM to a lower modulation order profile or remove the channel from the CM' receive channel set to ensure service quality.

3.6. Test narrow band noise (LTE ingress)

We recorded an LTE signal, a single channel at ~10 MHz bandwidth. We injected this into an OFDM channel and gradually decreased the output attenuation. The plot below shows CM downstream RxMER with different noise output attenuation levels that we tested:

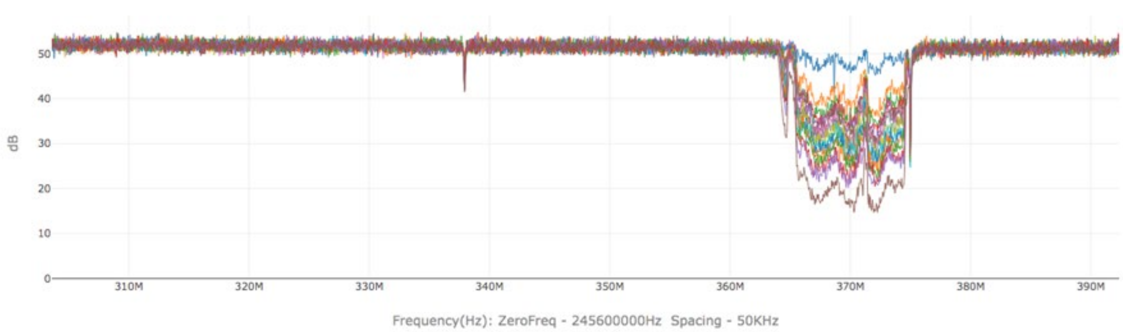

Figure 14 –D3.1 DS RxMER plots w Noise Injection LTE signal (1 channel)

The following plot shows the test results with 2 different flat profiles with modulation order at 256-QAM and 1024-QAM. The corresponding average RxMER values shown are calculated from CM downstream RxMER values only within the ingress frequency range.

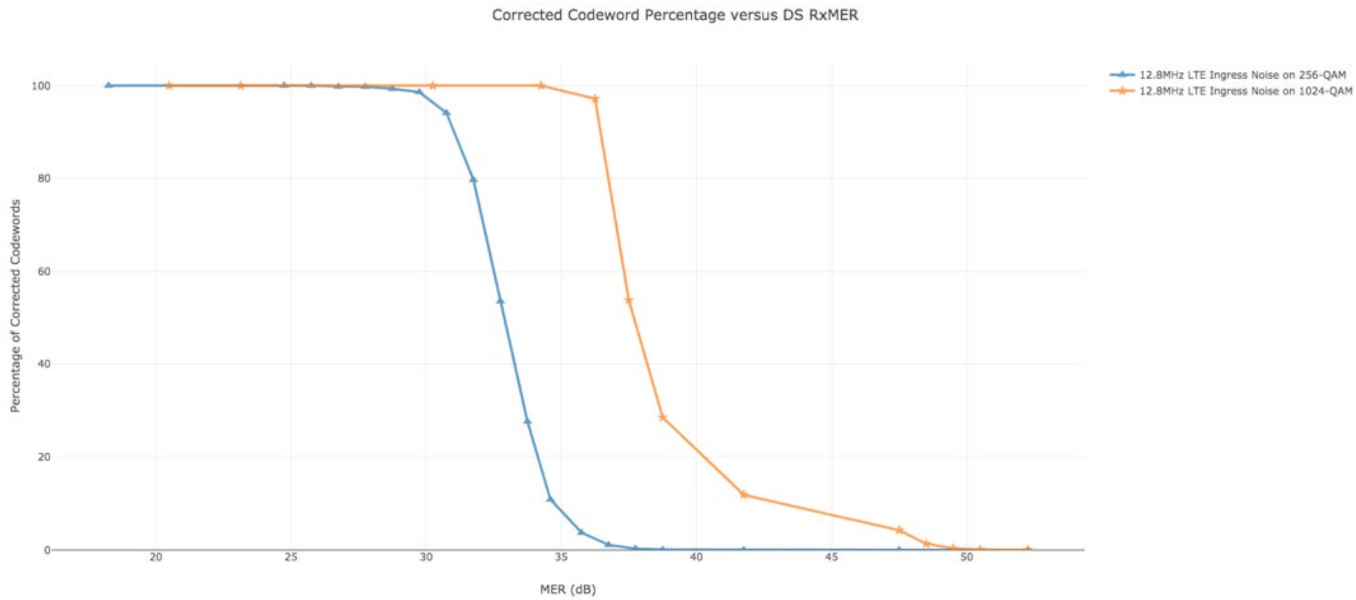

Figure 15 –FEC behavior w noise injection of LTE signal (1 channel)

Compare to the previous test result of injecting 12MHz AWGN noise (256-QAM), the LTE ingress noise adds slightly more pressure on the LDPC decoder. The CM starts to have corrected codewords at 38.75dB, which is 6dB higher than where we observed corrected codewords with 12MHz AWGN noise. However, we observed 100% corrected codewords around 27dB with both noise sources at 256-QAM.

We also recorded an LTE signal with 3 adjacent LTE channels. We injected this to an OFDM channel and gradually decreased the output attenuation of the noise signal. The plot below shows CM downstream RxMER with different noise output attenuation levels that we tested:

The following plot shows the test results with 2 different flat profiles with the modulation orders at 256-QAM and 1024-QAM. The corresponding average RxMER:

Figure 16 – D3.1 DS RxMER plots w Noise Injection LTE signal (3 channels)

The following plot shows the test results with 2 different flat profiles with the modulation orders at 256-QAM and 1024-QAM. The corresponding average RxMER values shown are calculated from CM downstream RxMER values only within the most powerful LTE channel's ingress frequency range.

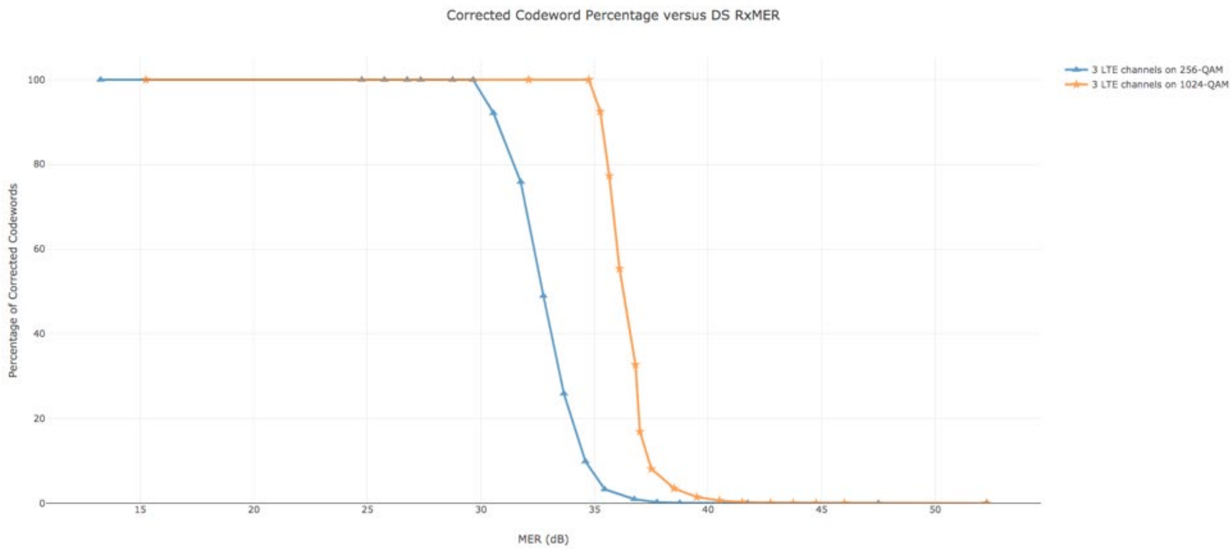

Figure 17 – FEC behavior w noise injection of LTE signal (3 channels)

When the noise ingress level is large enough, uncorrectable codewords can be observed at the moment the noise injection into the OFDM channel happens. The test injects LTE noise 3 different times. As the ingress noise power increases, an increasing number uncorrectable codeword spikes can be observed (0.18% uncorrectable codewords at the highest spike for one CM, and 0.5% for another CM).

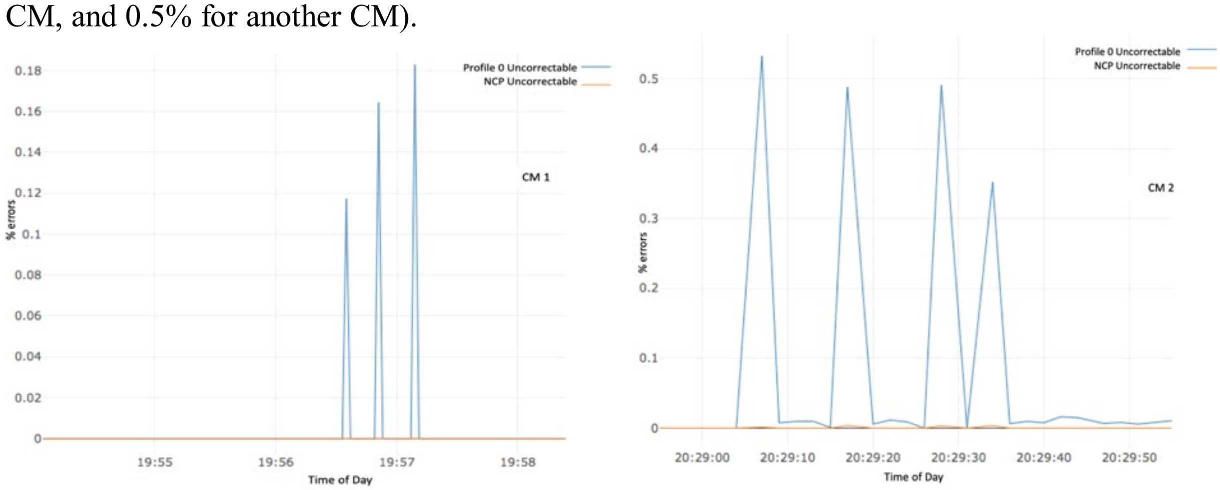

Figure 18 –D3.1 DS Uncorrectables at the instant of Ingress (2 different CMs)

The CM also sends CM-STATUS (16-Profile Failure) messages in order to inform the CMTS about its codeword failures on a profile. In the field, ingress/burst noise from the environment can be powerful enough to cause the CMs to have uncorrectable codewords and send CM-STATUS messages to the CMTS. As discussed in detail in Chapter 5, the CMTS acts based on

CM-STATUS messages and changes the CMs' profile, and over time this happens back and forth. The probability of capturing & observing ingress/burst noise in the cable plant is low, since MSOs are capturing RxMER measurements a few times a day, on a CM is low. This can make it hard to diagnose the root cause of data rate issues in the system. Different CMTS/CM vendors can use different algorithms/thresholds to generate & handle CM-STATUS messages, which leaves room for system improvements on this issue.

3.7. Understand profile switchover behaviour

This test determines the behavior of the system when using multiple flat profiles and when noise is injected. The goal was to identify the RxMER values at which the system switched profiles (to a lower modulation order profile or back up to a higher modulation profile). Traffic was sent to the CM at 500 Mbps. The CM was ensured to be on the highest profile (Profile 2, Profile 0 is the default). The test introduces noise until the system changes the profile for the traffic to profile 1 and then to Profile 0. The test also removes the noise slowly and notes when the CM returned to profile 1 and then back to Profile 2. When injecting noise, on a downward move from a higher profile to a lower profile, the movement is immediately after the CM sends a CM-STATUS message. On the way back, from a lower profile to a higher profile, there is a built-in hysteresis in the CMTS settings for profile recovery and the time to recover depends on the settings supported by the CMTS.

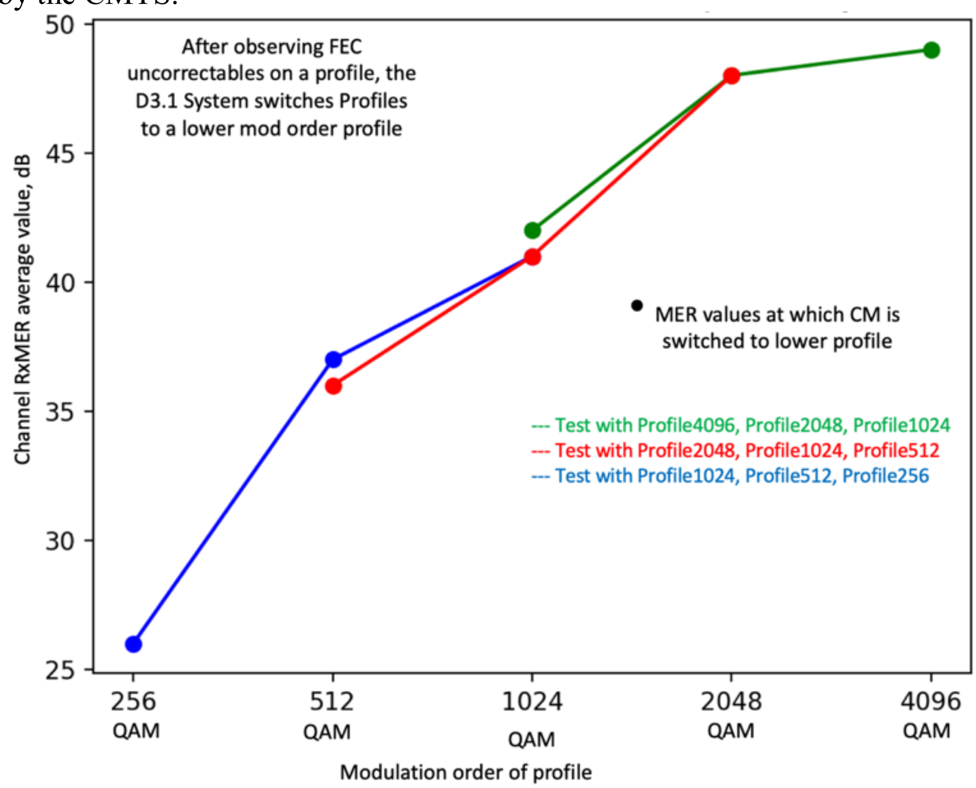

Figure 19 –CM-CMTS DS Profile Switchover at specific

4. OFDM Channel Configuration

There are various items to consider for an operator when Deploying OFDM channels and configuring them on the plant. These include the location of the OFDM channel itself and the location of the PLC within the channel. Care needs to be taken to work around interference/ingress noise which can affect system performance. There are also various physical layer settings like cyclic prefix, Roll-off, Interleaver depth, etc.). There are also many MAC layer settings like CM-STATUS messaging which can affect system stability. Care needs to be taken to avoid problems like profile flapping and ultimately there is a need to build a solution like PMA.

4.1. OFDM Channel Location

The CM supports a two independently configurable OFDM channels each occupying a spectrum of up to 192 MHz in the downstream. As the operator opens up spectrum to locate OFDM channels, some operators need to put the OFDM channels in the roll-off region in their plant. Per the results seen above, CMs which are in the roll-off will suffer from FEC codeword errors in the roll off region. At a minimum the operator will need to design profiles which lower the modulation orders used in the roll-off region, to ensure the CMs have a reliable data connection.

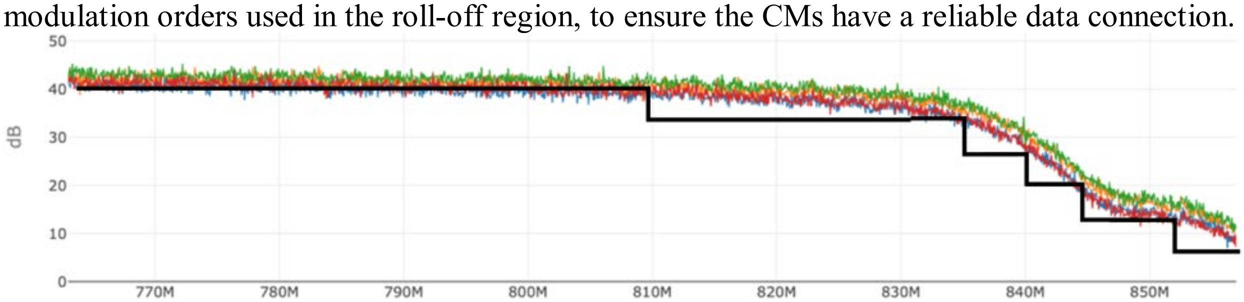

Figure 20 – Profile for an OFDM Channel in the Roll Off

4.2. PLC Location within OFDM Channel

The aim of the PLC is for the CMTS to convey to the CM the physical properties of the OFDM channel. When acquiring the OFDM channel, the CM needs the physical parameters of the channel. The CM first acquires the PLC, and from this extracts the parameters needed to acquire the complete OFDM channel.

The CMTS places the PLC at the center of a 6 MHz encompassed spectrum with no excluded subcarriers and places the 6 MHz encompassed spectrum containing the PLC on a 1 MHz grid. For 4K FFT OFDM, this 6 MHz will contain 56 subcarriers followed by the 8 PLC subcarriers followed by another 56 subcarriers. For 8K FFT OFDM, this 6 MHz will contain 112 subcarriers followed by the 16 PLC subcarriers followed by another 112 subcarriers. The PLC consists of 8 symbols of preamble followed by 120 symbols of data subcarriers, 16-QAM is used for the PLC subcarriers. The CMTS interleaves scattered pilots and data subcarriers, but does not interleave continuous pilots, the PLC, and subcarriers belonging to excluded regions. The insertion of continuous pilots, PLC and excluded regions happens after both time and frequency interleaving. There is no interleaving of the 400 kHz of PLC subcarriers (8 or 16).

So essentially any interference in the same spectrum as the PLC means that the CM could lose the PLC and hence the OFDM channel itself. It is very important that the operator and the CMTS place the PLC in a part of the spectrum that is less susceptible to noise and interference. The operator with an overriding configuration is expected to place the PLC as appropriate within the channel.

Many operators have built a geographical map of the LTE and other wireless channels in their footprint. This map is then correlated to the location of the OFDM channels in the cable plant in those areas. Specifically, the idea is to find the spectrum with the least amount of overlap with wireless carriers in that region. The best 6 MHz chunk of spectrum within the OFDM channel is then chosen to locate the PLC.

As an example, LTE Band 12,13, 17 are some of the more common bands used by cellular providers in USA [LTEFreq]. These correspond to the following frequencies:

- Band 12: Uplink of 699 – 716 MHz, Downlink of 729 – 746MHz,
- Band 13: Uplink of 777 – 787 MHz, Downlink of 746 – 756 MHz,
- Band 17: Uplink of 704 – 716 MHz, Downlink of 734 – 746 MHz.

A DS OFDM Channel which overlaps an LTE band, say from 648-744MHz or 702-798MHz, could see LTE ingress from these bands. See figure below for an example of how LTE ingress into the cable plant can affect the RxMER of a CM in that region.

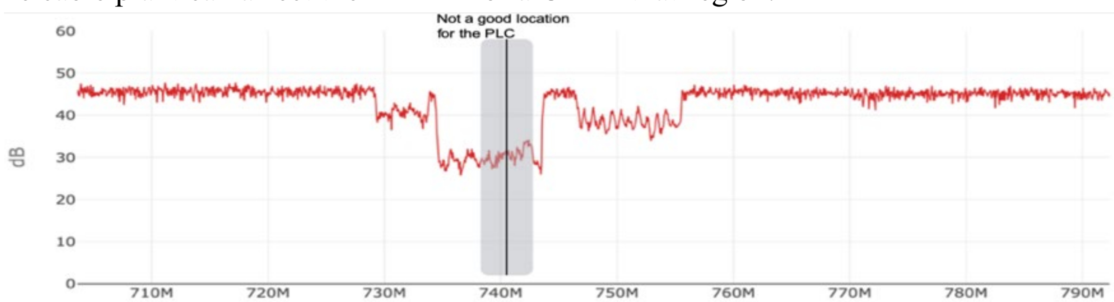

Figure 21 – LTE Ingress in cable plant, seen in CM RxMER (729-756 MHz)

Similarly, for European deployments in addition to LTE bands, Digital Audio Broadcast (DAB+) uses a wide-bandwidth broadcast technology and typically spectra have been allocated for it in Band III (174–240 MHz). The recommendation is to be aware and keep the PLC out of the DAB+ Region and any other potential source of ingress into the cable plant.

4.3. OFDM Channel Parameters

From many of the initial deployments, the operators have quickly realized that there are many CMTS defaults that need to be revisited, to optimize the channel operation.

4.3.1. Cyclic Prefix and Roll-Off

The addition of a cyclic prefix (CP) enables the receiver to overcome the effects of inter-symbol-interference caused by micro-reflections in the channel. The time duration of the CP should be chosen to be longer than the time of the longest significant reflection. Windowing maximizes channel capacity by sharpening the edges of the spectrum of the OFDM signal. Spectral edges occur at the two ends of the spectrum of the OFDM symbol, as well as at the ends of internal exclusion bands. The roll off must be integrated within the duration of the CP. The choice of a larger CP from 0.9375µs to 5µs, will affect the efficiency and the capacity accordingly from 1.8 Gbps to 1.5 Gbps, for a 192 MHz OFDM channel. A CP of 1.25µs and a Roll-off period of 0.625µs have given a few operators the stability and performance they are looking for.

4.3.2. Interleaver

The OFDM symbols, comprised of data subcarriers, scattered pilot placeholders, and NCPs, are subject to time and frequency interleaving. Time interleaving mitigates the impact of burst noise, while frequency interleaving mitigates the effect of ingress.

Time interleaving disperses the subcarriers of an input symbol over a set of output symbols, based on the depth of interleaving. Therefore, if an OFDM symbol is corrupted by a noise burst, this burst is spread over the symbols when it is de-interleaved, thereby reducing the error correction burden on the decoder.

Frequency interleaving occurs after time interleaving. Frequency interleaving disperses subcarriers of the symbol along the frequency axis; therefore, OFDM subcarriers impacted by narrowband ingress are distributed between several codewords, reducing the number of errors in each codeword. For time interleaving, the CMTS supports a maximum value of M equal to 32 for 20 µs symbol duration (50 kHz subcarrier spacing) and 16 for 40 µs symbol duration (25 kHz subcarrier spacing).

Initially running the system with M=16, for 50 kHz spacing looks to protect against many noise sources. Care also needs to be taken that frequency interleaving is enabled, one of the early CMTS implementations turned off frequency interleaving in some cases.

4.3.3. NCP modulation

When the data codewords are mapped to subcarriers within a symbol, a pointer is needed to identify where a data codeword starts, this is the Next Codeword Pointer (NCP). There are a variable number of NCP message blocks (MBs) on each OFDM symbol each pointing to the next codeword. Each FEC encoded NCP MB is 48 bits. The NCP QAM constellation can be QPSK, QAM-16, QAM-64.

The CMTS places the NCP subcarriers beginning from the frequency location of the highest frequency active data subcarrier of the OFDM symbol, and going downwards along active data subcarriers of the OFDM symbols before they are time and frequency interleaved. The CMTS time and frequency interleaves the NCP subcarriers using the algorithm applied to data subcarriers. This essentially allows the NCP modulation order to be QAM-16 or QAM-64 in practice. The QAM-16 looks to be sufficient to handle all variations in the plant noise across CMs in majority of deployments and many deployments are running with QAM-64 as well.

4.3.4. Pilot multiplier

Downstream pilots are subcarriers modulated by the CMTS with a defined modulation pattern that is known to all the CMs in the system to allow interoperability. There are two types of pilots: continuous and scattered. Continuous pilots occur at fixed frequencies in every symbol. Scattered pilots occur at different frequency locations in different symbols.

The Number of Continuous Pilots = $min\ (max\ (8,\ ceil\ (M * (Fmax - Fmin)/190e6))),\ 120)$.

The value of M in equation is as a parameter between ($120 \geq M \geq 48$) that can be adjusted by the operator or the CMTS. The typical value seen in deployments for M is 48, which will give 56 pilots (48 + 8 PLC pilots) for a 190 MHz channel. There is no observed change in behavior/stability by increasing the pilot scale.

5. DS CM-STATUS Interactions and Settings

CM-STATUS messages are needed in cases where the CM detects a failure that the CMTS cannot detect directly, or where the CM can send valuable information to the CMTS when an error or a recovery event occurs (for example, the CM can report a T3 timeout to the CMTS). Upon receiving an error indication, the CMTS is expected to act in order to correct the error. A CM transmits a CM-STATUS message on any available channel when it detects an event condition and the reporting of the event type is enabled on the CM. These events are summarized in table CM-STATUS with the Event Type Codes and Status Events and are also defined in detail in the DOCSIS specifications [MULPIv3.1]. Some event types are for a particular downstream channel, a particular upstream channel, or a DSID. For each such event, the CM maintains a separate state variable as to whether the event condition is considered "on" or "off" for each channel or DSID. The CM-STATUS message includes the Event Type Code and a unique Transaction ID for each occurrence of the event (IDs start from 1 and go up to 65535 before wrapping around to 1).

Fault detection and recovery occurs at multiple levels. At the physical level, FEC is used to correct errors where possible. The MAC protocol protects against errors through the use of checksum fields across both the MAC Header and the data portions of the packet. All MAC management messages are protected with a CRC covering the entire message, At the network layer and above, the MAC Sublayer considers messages to be data packets protected by the CRC field of the data packet; any packets with bad CRCs are discarded. Recovery from these lost packets is in accordance with the upper layer protocol.

When loss of lock is detected on an OFDM channel, it is possible that the channel is still partially functional and receives data. So, in case of high FEC errors detected in PLC, NCP or Profile A, the CM attempts to continue using the channel and enters a Partial Channel Mode. If the upstream communication is available, the CM sends out a CM-STATUS message to inform the CMTS of the Lost Lock event. If CMTS becomes aware of an interruption of a CM's Primary Downstream Channel (via a CM-STATUS message from the affected CM or from another CM), the CMTS can take an appropriate action. It can potentially send a DBC-REQ to the CM to reprogram the downstream channel or the profile as appropriate.

Table 2 – CM-STATUS Events related to DS Channel Failure

Event Code	Event Description	Channel applicable	Introduced in
2	QAM/FEC Failure	SC-QAM	D3.0
5	QAM/FEC Recovery	SC-QAM	D3.0
16	OFDM Profile Failure	OFDM	D3.1
24	OFDM Profile Recovery	OFDM	D3.1
20	NCP Profile Failure	OFDM	D3.1
22	NCP Profile Recovery	OFDM	D3.1
21	PLC Failure	OFDM	D3.1
23	PLC Recovery	OFDM	D3.1

5.1. CM CM-STATUS State Machine

When the CM detects an event, it creates a CM-STATUS message and it starts a timer for a value between 0 and the max-hold-off with 20ms resolution, before sending the message. When the timer expires the CM-STATUS message is placed sent to the CMTS.
The CM will send the CM status message again if it has not reached the max-reports number of times, or it has not received a CM-STATUS-ACK from the CMTS (with the same transaction number for the same event ID), or the event type has not been disabled. After the first message it will send the message again if applicable on the period of the full max-hold-off timer. The CM transitions back to the IDLE state after the max-reports or a CM-STATUS-ACK is received.
The recovery and failure events are treated as separate events, and as we have seen in the field a recovery message can quickly follow a failure message based on the CM algorithms. If the status transitions from "on" to "off" to "on" while a hold-off is in place in the SENDING state, the max reports is reset and additional CM-STATUS messages are sent as a new transaction-Id.

5.2. CMTS Management of CM-STATUS Messages

When the CMTS receives a CM-STATUS from a CM reporting a Partial Channel Mode, the CMTS can either stop sending data for that modem to the reported profile or move the service flows for the CM to another working profile on that channel. The CMTS attempts to resolve partial channel situations, such as by shifting the service flow to other profiles or other channels. Once the CMTS receives the CM-STATUS message it may react based on any algorithm it chooses. The CMTS may use a DBC message to ask the CM to make a change to its RCS (add or delete a DS channel), or ask the CM to change the DS profile in use. The DBC transaction includes the following: A Request (DBC-REQ) from CMTS to CM, the CM sending back a DBC Response (DBC-RSP) to the CMTS, and the CMTS acknowledging with a DBC-ACK. These message exchanges will require a request and grant cycle and may be on the order of 100s of ms (maybe 300-800 ms). The CMTS may choose to immediately quit forwarding data on the profile or the channel if it believes the CM can receive the data on another channel or profile which the CMTS sends the data.
Some CMTS implementations immediately trigger a profile update upon receiving a CM-STATUS message if the CM is not in profile 0, or an RCS change if the CM is in profile 0. The only way to recover the profile is based on CM-STATUS messaging. This can be used to add stability from the CMTS perspective. If no other CM-STATUS failures occur for that profile during the timer period, then the CMTS will move the traffic back to the higher capacity profile as soon as the OFDM recovery hysteresis time expires.
The CMTS picks the profile order by taking the bps/Hz for each subcarrier and summing all active subcarriers ranking from highest capacity least robust to lowest capacity most robust.

5.3. CMTS Polling CMs on RxMER and FEC

There are a couple of different ways the CMTS can check how the CM is performing on the Downstream. The OPT-REQ MAC Management message is used by the CMTS to cause a CM to test various aspects of an OFDM downstream channel. A single OPT-REQ message can be used to test the CM's ability to receive the specified downstream OFDM profile by checking for FEC statistics, alternatively it can be used to query the CM's RxMER statistics.
CMTSs today follow one or both of the above approaches, they periodically check for RxMER of the CM or the FEC statistics for a particular profile. A CMTS does this periodically by

issuing OPT-REQ messages to every CM on the OFDM channel. The rate at which the CMTS does this can be configured in a vendor proprietary manner.
The CMTS compares the numbers from these measurements with a certain threshold, before it decides it needs to downgrade or upgrade the profile on which the traffic reaches the CM.

5.4. Recommended configurations

The configuration settings for the CMTS-CM CM-STATUS state machines have the following goals. First is to avoid any unnecessary profile flapping in the presence of transient noise or noise near modulation profile error thresholds. The next is to ensure that data is forwarded to the CM on the correct profile when a more robust profile is chosen. The goal is to avoid unnecessary partial service flapping of the OFDM channel in the presence of noise that impacts profile 0. An operator wants to ensure that when the CM enters the partial service state with the OFDM channel no longer available, data can still be sent over the D3.0 channels.

5.4.1. CM Event Thresholds for CM-STATUS Messaging

CMs implement different methods to identify the failure of a profile on which it is receiving data. A CM also can implement different failure/recovery thresholds, for each event e.g. failure/recovery of PLC, NCP, Profile 0, Profile 1,2,3 etc. There are different ways of detecting profile codeword errors: using a raw FEC CW error count vs. using a count of time over which errors have occurred.
The first method can be thought of as a count of the Number of Errored Codewords. e.g. if 50 out of 1000 codewords have errors, the CM could declare the profile to have failed and raise the event by sending a CM-STATUS message. Similarly, if 990 codewords out of the last 1000 have no errors, then the CM can declare that the profile is working correctly again and clear the event with another CM-STATUS message.
The second method can be thought of as a count of the time there are errored Codewords on the channel. e.g. if 2 out of last 20 seconds have codeword errors, the CM could declare the profile to have failed and raise the event by sending a CM-STATUS message. Similarly, if 18 out of the last 20 seconds have no errors, then the CM can declare that the profile is working correctly again and clear the event with another CM-STATUS message.
Another method which a failure conditions can be detected would be a combination of the codeword error count and time for which errors are present.
Both the failure threshold and the recovery thresholds can be adjusted to bring out the desired behavior. This then leads to the question of what is the desired behavior, which leads to the question which kind of noise sources are we trying to recover from quickly. If an operator can characterize the noise sources in the plant, then these settings can be fine-tuned for each kind of plant. e.g. a short bursty noise which is present for some amount of time, or a noise which may be present for a longer time durations or strong noise which affects large number of subcarriers and hence codewords, vs a low-grade noise which only causes a small number of errored codewords.

The philosophy which yields the most optimal results is as follows: *Fail fast and recover slowly.*

This means we want the CMs to detect failures in profiles, as soon as possible. This also needs to happen without being unreasonable, i.e. we don't want to declare a profile failure with say just a few codeword errors. So once the errors in a profile are confidently detected then it is better to raise the failure as soon as possible. The idea is to minimize the impact to the customer data

traffic. The earlier the profile failure occurs the sooner the CMTS can move the data to a more robust profile.
Also, on the recovery side the CM needs to make sure the profile is very stable before declaring the profile is good to use. Prematurely moving the traffic to a higher modulation profile will also lead to traffic loss and the profile flapping behavior.

5.4.2. CMTS Thresholds for CM-STATUS

The CM-STATUS reporting mechanism includes a random holdoff prior to transmission of status report messages. This value is set on the CMTS and makes the CM its dampen CM-STATUS messages. The Maximum Event Holdoff Timer indicates the value of that random holdoff timer to be used by the CM when determining when/whether to transmit a CM-STATUS message. This TLV associates a separate hold-off timer value with each CM-STATUS event type code managed by the CMTS.

A Maximum Reports Count Timer value controls how often repeated CM-STATUS messages for the same Transaction Identifier are sent by the CM. It controls how many CM-STATUS messages for the same Transaction Identifier are transmitted by the CM. A Maximum Reports Count of zero signals that the CM continues sending CM-STATUS messages as long as the event condition is "on" and is enabled for reporting. If the CMTS receives a CM-STATUS message from the CM, the CMTS transmits a CM-STATUS-ACK message with the received event type and transaction ID.

When the CMTS receives a CM-STATUS message, it needs to act to remedy the situation. The CMTS action is based on the received event in the CM-STATUS message.
As an example, on Profile Failures: the CMTS could change the profile for traffic, this will be a combination of DBC, OPT, (DPD) etc. If the PLC or Profile 0 failure occurs, the CMTS could reassign the CM receive channel set by removing the affected channel and replacing it with another OFDM channel if available.

5.5. DS Profile Flapping

When the CM experiences any kind of intermittent noise on the OFDM channel, e.g. LTE, sweep generator, or other ingress noise, or if the CM is simply at the margin for a profile, the RxMER for the channel falls below the required level for the CM to decode a certain modulation order. The interleaving and FEC are unable to correct for noise and the CM will experience code word errors on the channel on a particular profile assigned to it. Let's say the CM is assigned 3 profiles: Profile A (256QAM), Profile-1024QAM, and Profile-4096QAM. All of these profiles are flat, i.e. every subcarrier has the same modulation order.
If the CM is currently operating on Profile-4096QAM, and it experiences code word errors, then as described above the CM will send a CM-STATUS message to the CMTS, informing the CMTS that it has a profile failure on Profile-4096QAM. Now the CMTS will start sending traffic on the next lower profile, in this case that is Profile-1024QAM. Now let's assume the noise interference is just enough that it can operate on Profile-1024QAM without code word errors. This switchover of profiles happens quite seamlessly on the CMTS. Now the CMTS is still sending traffic on Profile-4096QAM, to other CMs on the profile. So, this CM will continue to try and decode codewords on that profile, though the traffic is not destined to it.

Now if the noise ingress disappears quickly, say within 10 seconds or so, then the CM will be able to start successfully decoding packets on the Profile-4096QAM. If the CM is currently experiencing no more code word errors, (thresholds as described before), the CM will now send a CM-STATUS message to the CMTS, informing the CMTS that it has a profile recovery on Profile-4096QAM. Now the CMTS will move the traffic destined to this CM from Profile-1024QAM back to Profile-4096QAM.

If the noise ingress is strong enough, a CM could drop through multiple profiles all the way down to Profile A, and then back up all the way when the noise disappears.

Now this whole process of receiving traffic on a high profile and then to a lower profile and back is known as profile flapping.

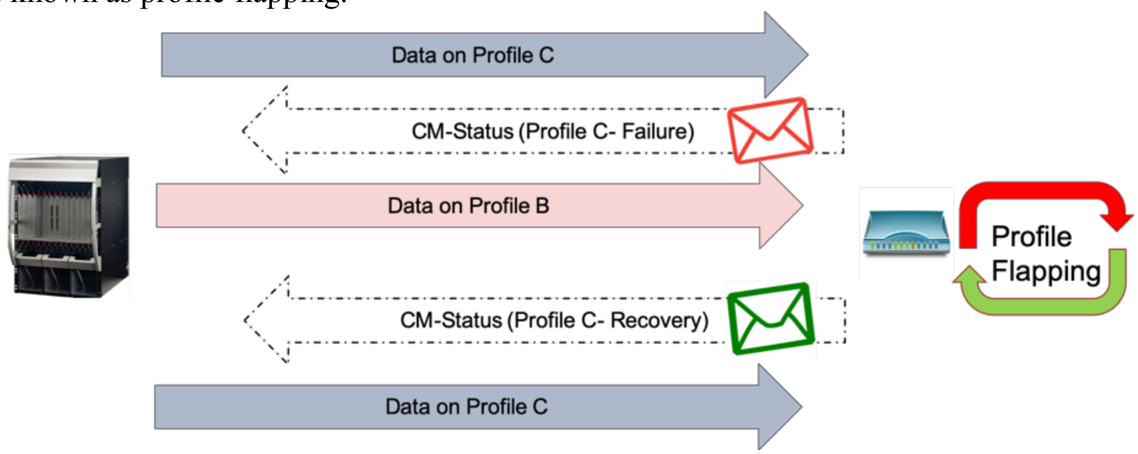

Figure 22 – D3.1 Profile Flapping Behavior

In practice depending on the settings on the CMTS and the CM, of the CM-status messaging and of the hysteresis time which the CMTS applies before promoting the CM back to the highest profile, the CM could see significantly lower capacity. The least common denominator profile, Profile A could be 64QAM to accommodate the worst-case user and if others CMs also spent a lot of time on profile A due to the above profile flapping issue, the CM and the network see a much-reduced capacity.

the channel maybe intermittently impaired for some CMs leading to a cycle of profile failure and recovery messages. Each time the CM is forced to a lower profile, the CMTS as per the configured hysteresis settings does not recover the profile fast enough. Profile Downgrade is immediate while the profile upgrade is slow due to the hysteresis in the system reacting to the recovery.

If errors are seen on Profile A, then the consequences are more catastrophic as this means the channel itself is unusable, and the CMTS now considers the OFDM channel inoperable and the CM from partial channel, has moved to partial service as it has lost the use of the whole OFDM channel. Even though the CMTS reacts to both the profile failure and recovery messages, this profile downgrade - profile flapping/ partial service scenarios results in user traffic being dropped by the CMTS-CM. This starts manifesting itself to the user as intermittent connectivity issues and slow speeds in general.

The figure below explains the behavior visually using code word counters over time, as data moves across profiles, along with the timelines for CM-STATUS messages (not drawn to scale).

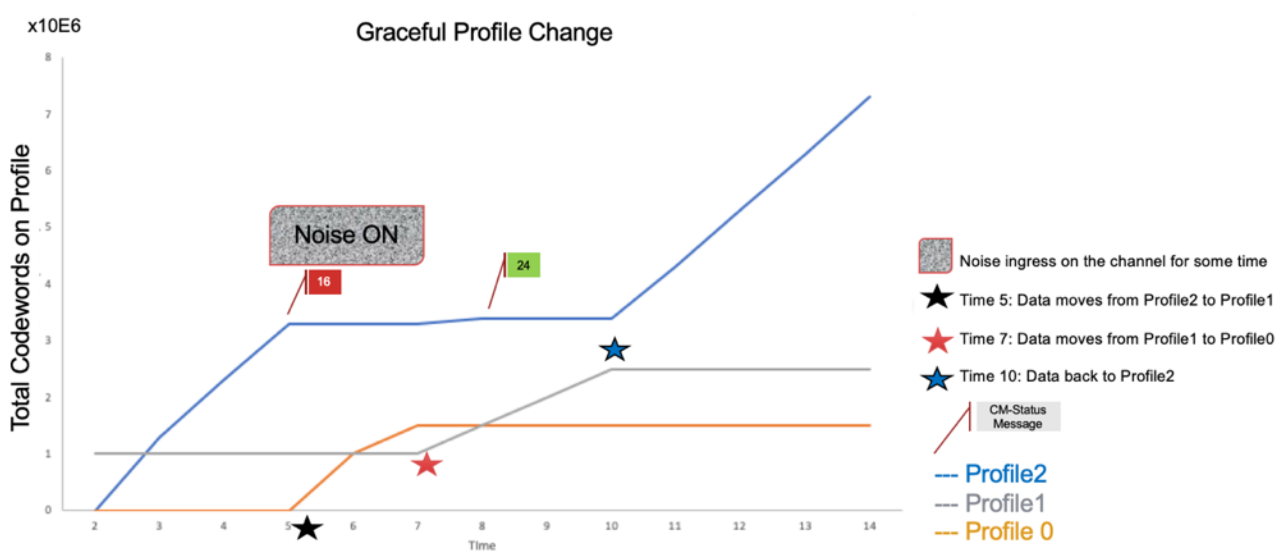

Figure 23 – D3.1 Profile Flapping Example

5.6. Suggested CM-STATUS Settings

On the CMTS the CM-STATUS Holdoff Timers need to be configured different from default as below for better performance.

Table 3 – CMTS CM-STATUS Settings

CM-STATUS Event	CM-STATUS Description	Setting (multiple of 20 ms)	Time (seconds)
2	QAM/FEC Failure	50	1
5	QAM/FEC Recovery	**500**	10
16	OFDM Profile Failure	**50**	1
24	OFDM Profile Recovery	**500**	10
20	NCP Profile Failure	100	2
22	NCP Profile Recovery	100	2
21	PLC Failure	100	2
23	PLC Recovery	100	2

CMTS Automatic Profile Recovery behavior is not specified in the DOCSIS standards and these are non-Standard CMTS Configurations, which need to be customized per CMTS.

Table 4 – CMTS Maximum reports and ACK

CMTS setting	Value
CM-STATUS Maximum reports	5
CM-STATUS-ACK	ON
Vendor Proprietary Mechanism (Hysteresis for profile upgrades). Maybe known as profile guard time, Unfit time, recovery time etc.	5 mins or lower (2 Mins) if supported

6. Downstream Profile Mangement Application

DOCSIS 3.1 introduces the concept of modulation profiles or bit loading characteristics for each subcarrier within the OFDM/A channels. A modulation profile is a list of modulation orders or bit loading configurations, defined for each subcarrier within an OFDM channel, or for each minislot in an OFDMA channel. A CMTS can define multiple modulation profiles/IUCs for use on a channel, where the profiles differ in the modulation orders assigned to each subcarrier or minislot. A CMTS can assign different downstream and upstream modulation profiles for different groups of CMs.

As seen in Chapter 2, the interference patterns/RxMER on the different CMs on the same channel/plant are quite different. The best way to prevent profile flapping and enable robust operation on a D3.1 OFDM channel is to custom design the profiles for the CMs and to the state of the plant. Determining the best modulation profile to use on a channel is difficult, given the number of CMs and the differences in signal quality that they experience. PMA helps operators design the best modulation profiles are for each channel, given the channel characteristics seen by each CM on the network.

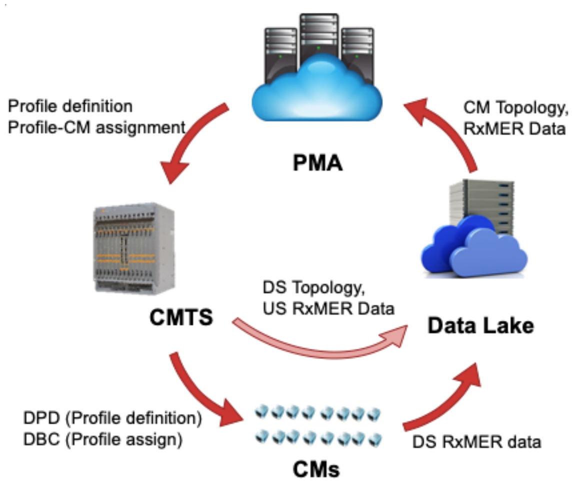

Figure 24 – Profile Management Appplication Deployment Architecture

PMA Goal: The goal of designing profiles is to increase reliable operation and throughput per CM. PMA essentially consists of intelligent clustering algorithms to group CMs which share similar noise characteristics together: Groups of CMs get assigned a unique custom designed profile, which works around specific ingress issues etc.

The tasks an external PMA performs for both downstream and upstream profiles are as follows:

1. Create a set of optimized modulation profiles for use on an OFDM or OFDMA channel by selecting the best modulation order for each subcarrier based on the channel quality measured at the CMs/CMTS using the channel profile test or probes. (For all CMs)
2. For a new CM joining the network and periodically for all active CMs, find the best fit among existing modulation profiles and recommend modulation profile usage. (Per CM)
3. Create backup profiles or downgrade a CM based on errors on a certain profile. E.g. based on CM performance and SNR margin, provide a better modulation profile for a CM. (Per CM)

PMA Benefits: Well-designed profiles minimize codeword errors which lead to data loss and sub-optimal throughput. Designing profiles around noisy areas in the plant makes the system operation more robust. CMs can get downgraded/ upgraded on profiles when the noise is intermittent, PMA can prevent this profile flapping. PMA can maximize network capacity by optimizing the bit loading of every subcarrier. The bandwidth gains in running a well-designed set of profiles can be anywhere from 20% to 40% capacity increase on a channel, compared with running the whole channel at 256-QAM. This can translate to a solid 200 to 400 Mbps extra capacity on each OFDM channel. This enables an operator to match growing bandwidth demands and defer potential node-splits and new equipment costs. PMA is in full scale deployment with one large operator, and in field trial with another. 10 other operators are trialing out PMA in their labs (for Downstream and upstream) as they ramp up their D3.1 deployments.

Since CMs can be assigned to modulation profiles that are optimized for their channel conditions, there is no longer a fixed value for channel capacity. The cleaner the channel is to any CM, the higher the modulation order its traffic is carried on, raising the overall average efficiency, and hence the overall capacity of the channel. Capitalizing on this capability requires that the PMA can determine the channel conditions to the set of CMs in the Service Group, and that from this information determine the optimal set of (up to 16) modulation profiles. CMTSs today support ~4 DS profiles, with plans to increase the number of profiles supported.

6.1. DS PMA Software System architecture

Different systems/solutions can be built for downstream PMA deployment. A reference architecture is show in the figure below.

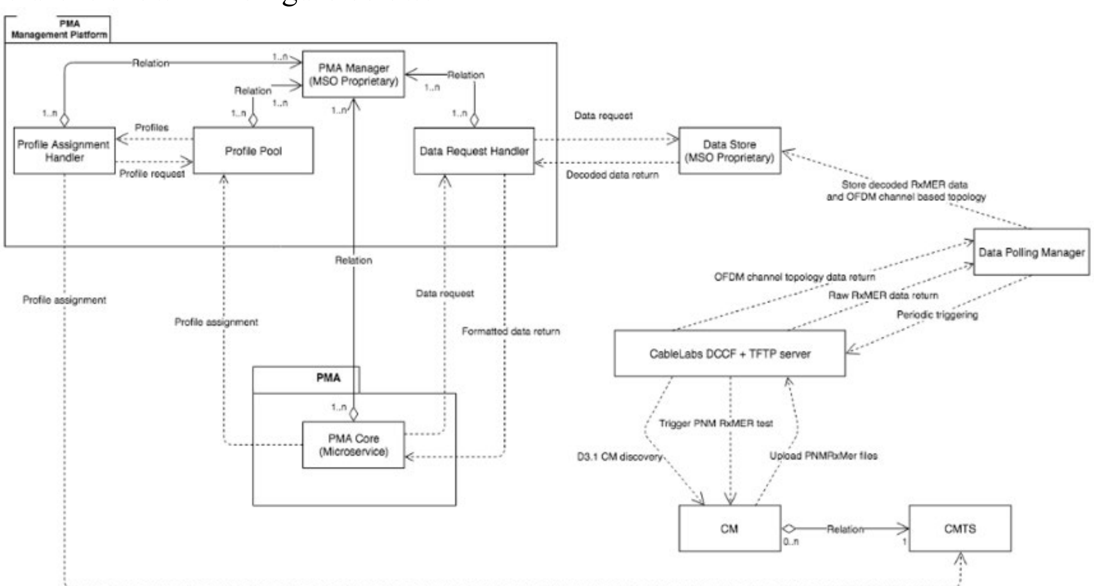

Figure 25 – Profile Management Appplication Software Architecture

A DOCSIS data collection framework is necessary to collect RxMER data from the CMs so that a PMA can create the appropriate profiles for the channel. DCCF, the DOCSIS common collection framework, developed by CableLabs along with industry partners is being used by

many operators and vendors. Here DCCF is used as an example of a data collector that can handle CM discovery, DOCSIS 3.1 PNM (& SNMP) data polling, and preprocessing for OFDM/A channel-based topology discovery, basic data storing and retrieving functionalities, and polling job management, etc. While DCCF is one of the frameworks for data collection from the cable plant, alternatives of DCCF can be built and chosen by vendors and operators.
In the architecture diagram, the following components perform the tasks described below.

- PMA Manager is the PMA Management Platform that handles profile assignment scheduling, policy input management, CMTS management, MER data redirection, and profile output management.
- PMA Core is the micro-service of the PMA core algorithm. It's a light-weight server that can handle multiple profile calculation requests efficiently.
- DCCF (or XCCF) with TFTP server (or MSO's own remote TFTP server) is a tool that will help MSOs setup and collect data. It attaches to one or multiple CMTSs and automatically discovers existing DOCSIS 3.1 CMs in operation. It's also able to discover the OFDM/A channel-based topology for all these DOCSIS 3.1 CMs, which is an essential function for PMA, which calculates profiles for each OFDM channel and needs the CM MAC address list to be hosted in CMTS_IP/slot/port/channel fashion.
- Data Polling Manager periodically triggers DCCF to start a new round of OFDM DS RxMER data polling, and also triggers OFDM channel-based topology discovery on the CMTS. It could be an MSO specific implementation with an ability to pre-process the raw data from the data collector (DCCF in the system).
- Data Store is where the PMA gets the RxMER data and MAC addresses of CMs that are on each OFDM channel. It consists of two parts, a Data Lake and a thin service layer that serves data with restful APIs.
- Profile Translator is a shim layer that translates PMA's profile output to actual assignment information and automation process to CMTSs through available interfaces such as RESTCONF or CMTS CLI.
- Data Request Handler is the service that handles the PMA Core's data request and reply with data from the Data Store.
- Profile Pool is where original profile assignments are stored.
- Profile Assignment Handler handles actual profile assignment to the CMTS by requesting profile data from the Profile Pool.

The recommended functions of a Data Polling Manager are:

- PNM RxMER test triggering functions, including periodic triggering scheduler functions
- Manage DCCF API calls and translate DS RxMER data to a PMA understandable format
- Understand and translate returned OFDM channel-based topology from DCCF
- Perform highly efficient Data Store API calls for storing pre-processed data in place

The recommended functions of a Data Store are:

- The restful API service of the Data Store must comply the APIs defined by the PMA and must serve the data in formats that PMA can understand
- The Data Lake of the Data Store could be a database or an HDFS, per actual needs

PMA interacts with the Data Store on the link (data request and decoded data return) below using APIs standardized by CableLabs in the PMA YANG model [PMA-TR-CL]. The advantages of this model are that it is highly scalable. Request don't hit the network (CMTSs/CMs) directly when there's demand, and it can also serve plant data for other applications without querying the devices multiple times.

6.2. DS PMA Practical gains

We have collected and analyzed field RxMER data from D3.1 CMs from a few different cities from 3 different operators. This data includes approximately 142 unique OFDM channels with about 25000 D3.1 CMs spread across these D3.1 channels. In most cases, each channel had anywhere from 100 to 300 CMs (average of 174 CMs per channel), and in a few of the cases the channel had 40~80 CMs. We ignored channels which had less than 20 D3.1 CMs.
We design optimal profiles for the set of CMs on each channel, and then calculate the capacity gain for each of these channels, when using a set of profiles, as described in the paper on profile management algorithms [D31PMA-INTX16].
This capacity gain metric is essentially the capacity gain of a channel when using multiple profiles for unique sets of CMs and comparing that to all CMs using the single 256-QAM profile. The gain in capacity (J) increases as the number of profiles increases. For the sample channel below, one can see that once we get past ~7 or 8 profiles there is only minor incremental benefit to the overall capacity of the channel

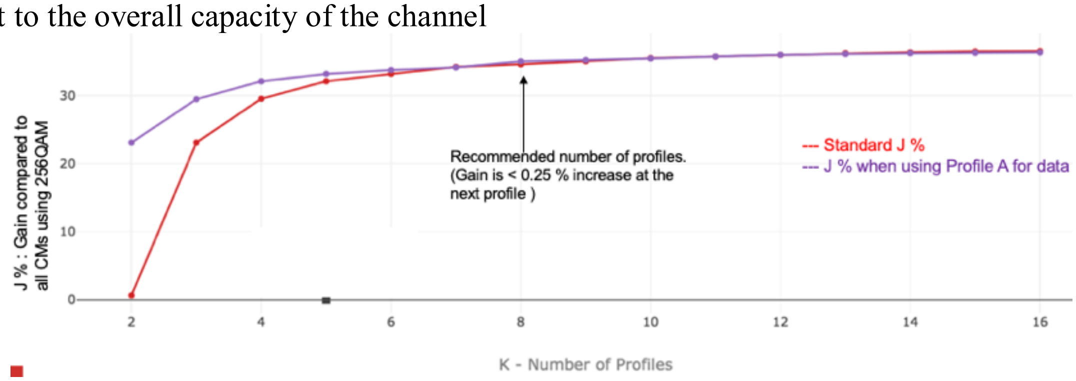

Figure 26 – Typical PMA Gain vs number of profiles (J-K Correlation)

Using this method, we calculated the capacity gain for each channel in our field data set of 142 channels and for a range of 3-8 profiles. Below are the plots of the gains seen across this sample population of OFDM channels, shown as a histogram. Each histogram is for a certain number of profiles. As the number of profiles supported on the D3.1 CMTS increases, the gain of each channel and hence the gain of the collective of the channels increases.

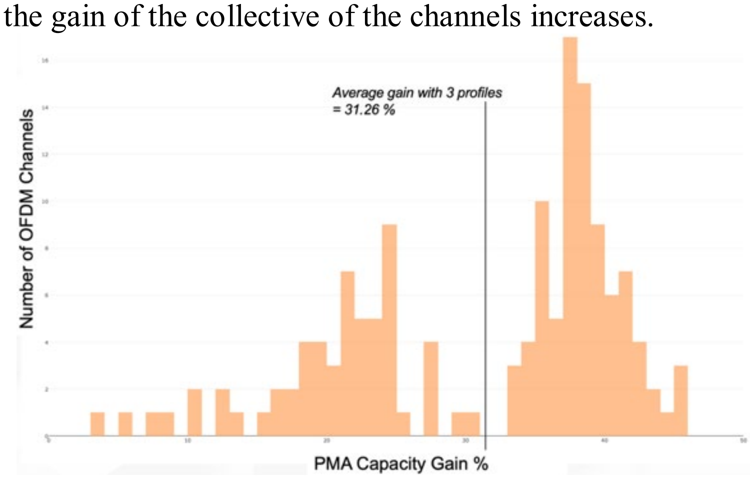

Figure 27 -PMA capacity gain histogram when using 3 profiles

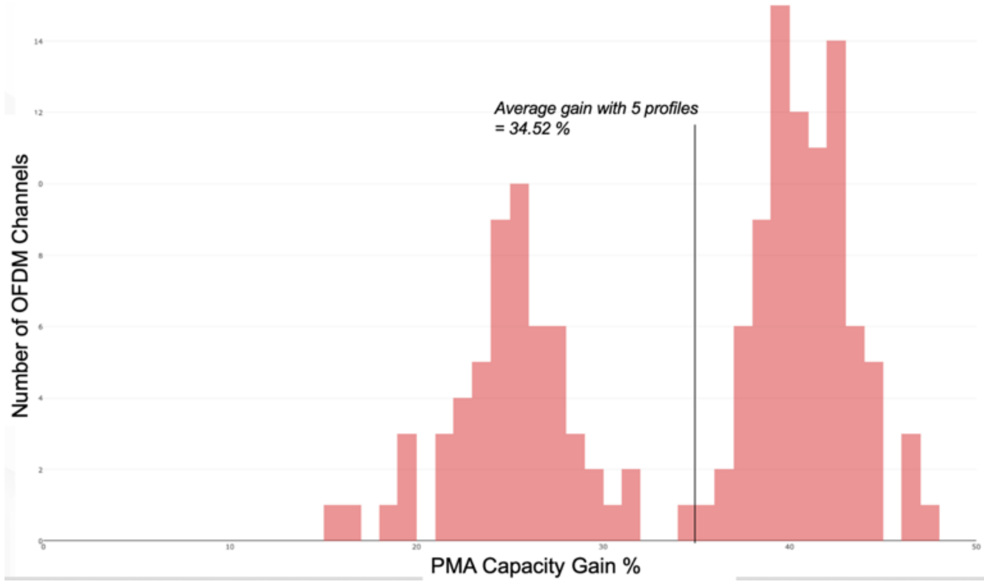

Figure 28 -PMA capacity gain histogram when using 5 profiles

Again, the PMA algorithm here has created robust profiles which work around the noise seen by each CM, and at the same time it also groups CMs to profiles so that the overall capacity of the channel is maximized.

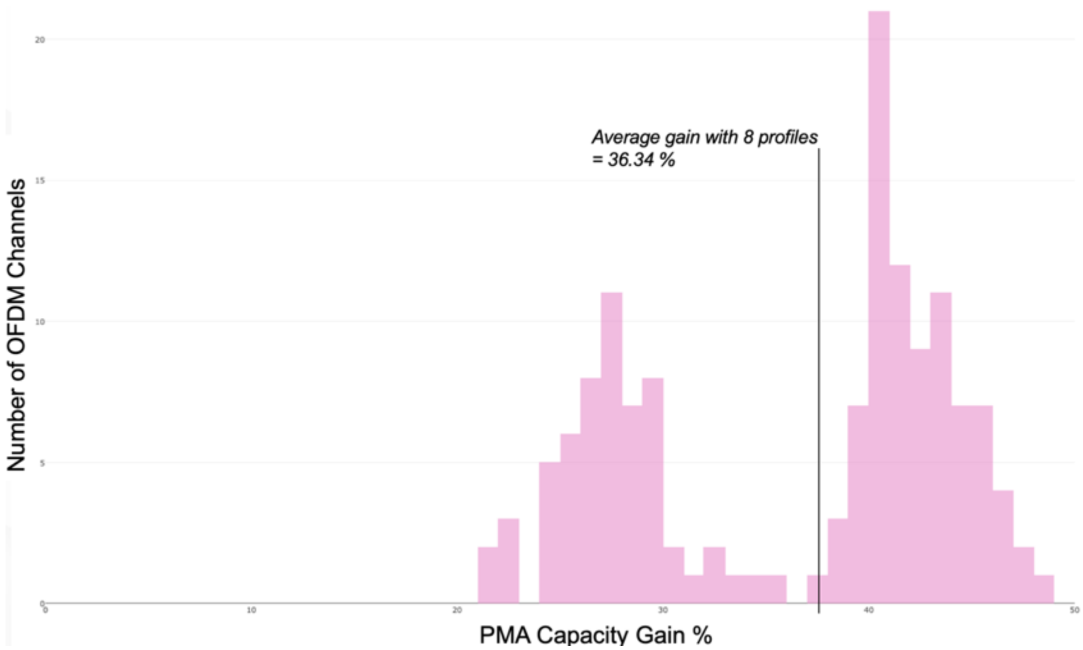

Figure 29 -PMA capacity gain histogram when using 8 profiles

The next histogram below shows a comparison of the capacity gains when using 2 customized profiles on a channel vs 10 customized profiles on a channel. This comparison gives a visual of how the capacity gains for 2 profiles (in blue) shift to the right for 10 profiles (in light green). A higher number of well-designed profiles unlocks significant capacity gains across the plant.

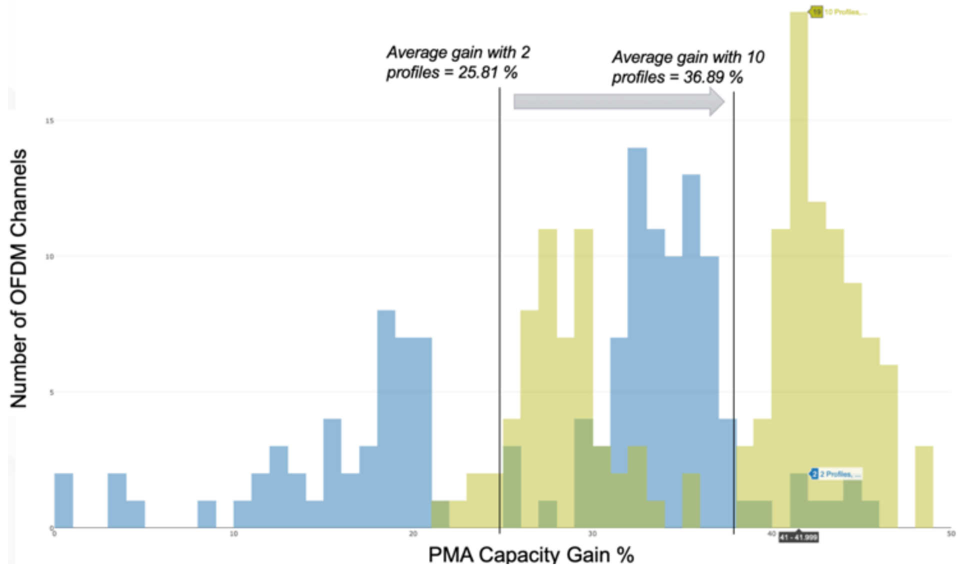

Figure 30 -PMA capacity gains, 2 profiles vs 10 profiles on a channel

The next histogram below shows a comparison of the recommended number of profiles. When you add another profile to the system, the capacity gain(J) will increase. The recommended number of profiles for that channel is when that improvement in capacity is less than an incremental thresholds value. The below plot shows 3 histograms each with the recommended number of profiles with 3 different threshold values.

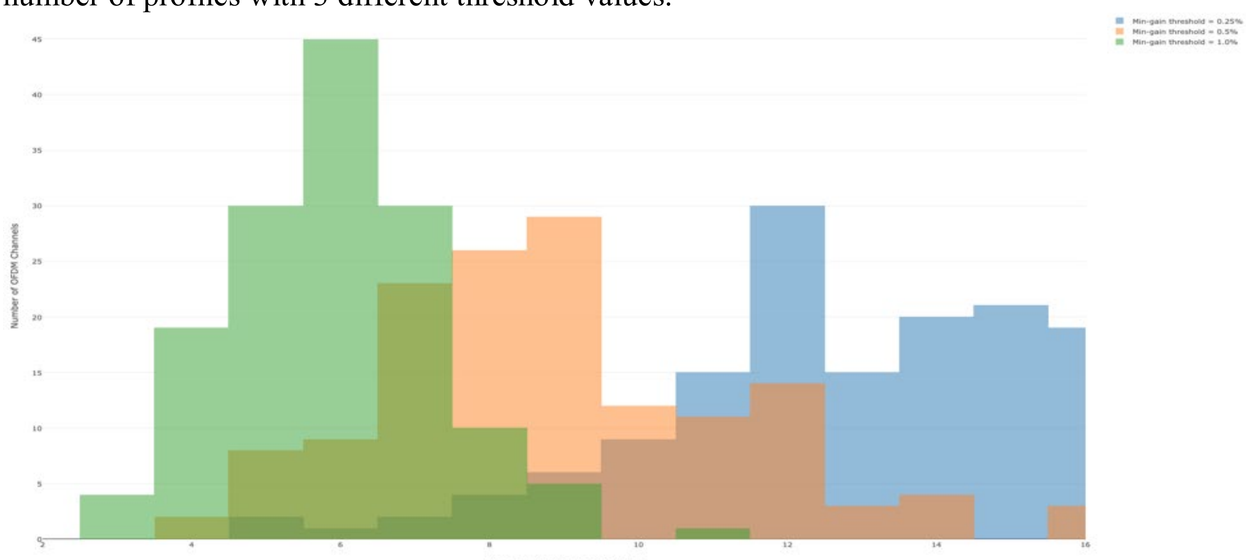

Figure 31 -Recommended number of profiles

As the increment threshold decreases from 1% to 0.5% to 0.25 %, then recommended number of profiles increases. For a minimum gain threshold of 1% majority of the recommended number of profiles is under 10. This is a good number of profiles for a D3.1 system to support, and many CMTS vendors are planning to increase the support for the number of DS profiles from ~3-4 to 8-10. The operators deploying D3.1 are coming to a consensus around this number of profiles.

DOCSIS 3.1 Upstream

7. D3.1 OFDMA Upstream FEC behavior

The D3.1 equipment supports a minimum of two independently configurable OFDMA channels each occupying a spectrum of up to 95 MHz in the upstream. The systems support upstream transmissions from 5 to at least 204 MHz and agile placement of the OFDMA channels within that range.

7.1. Noise Characteristics on an OFDMA Channel

US spectrum is in general noisier and more susceptible to interference than the DS spectrum. We are only beginning to get US RxMER data from field CMTSs. A few examples of these are shown in the figures below.
Depending on the plant, many CMs have relatively clean RxMER levels across a 20-55MHz OFDMA channel. (In the figure below, we are seeing some MER reporting issues from the CMTS every 5 MHz or so).

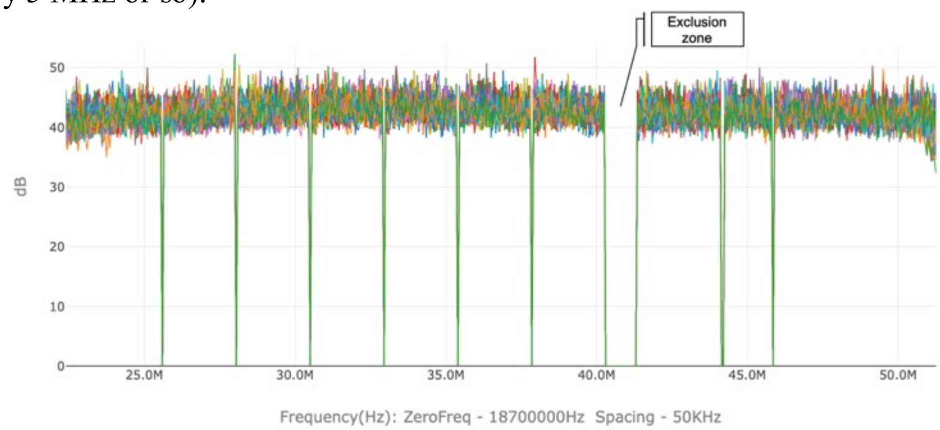

Figure 32 – Multiple measurements of D3.1 US RxMER from a CM

In the figure below, we are seeing some changes in the measured RxMER when one CM is present on the channel vs when multiple other CMs are also using the same OFDM channel.

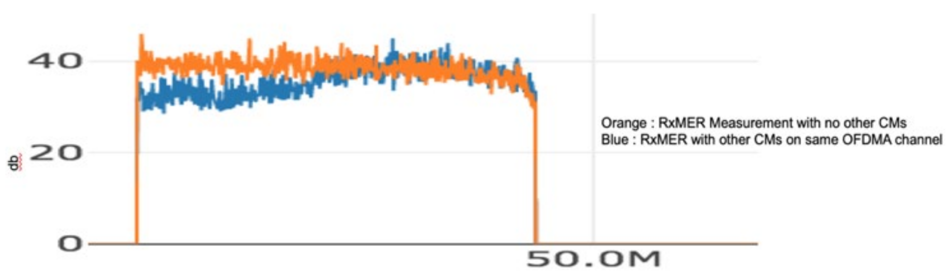

Figure 33 –D3.1 US RxMER from a lab CM

7.2. US FEC behavior

The choice of codeword sizes to be used in any given burst is based on the grant in the MAP message. The grant indicates which minislots are assigned to a given burst and which upstream profile is to be used. The CM and CMTS use this information to determine the total number of bits in the grant which are available to be used for FEC information or parity. Codewords are filled and transmitted in the following order, Full, Medium, Short, with codeword shortening when needed.

The ability of the system to support a given QAM level depends on the RxMER values and the mappings to an appropriate QAM level, when creating a profile. These mappings are defined in [PHYv3.1] and are summarized in the Table below.

Table 5 - US RxMER to QAM Level mapping

Upstream Constellation/ Bit Loading	US MER(dB)
QPSK	11.0
8 QAM	14.0
16 QAM	17.0
32 QAM	20.0
64 QAM	23.0
128 QAM	26.0
256 QAM	29.0
512 QAM	32.5
1024 QAM	35.5
2048 QAM	39.0
4096 QAM	43.0

7.3. Lab testing of US FEC behavior on D3.1 equipment

We wanted to understand the performance of the US FEC on D3.1 equipment, in an effort to understand at what points will an operator start seeing failures in the system. We are testing these with 1 CMTSs and 2 different CMs. We hope to increase the testing to include other CMTSs and CMs. The test was run on a 5-45 MHz OFDMA channel

7.4. Baseline test (no noise)

This test checks the RxMER at the CM and ran downstream traffic from the CM to the CMTS:

1. Send traffic from the modem starting at 150 Mbps for 30 seconds.
2. Repeat for the US IUC/profile configured at each QAM level between 16QAM – 64QAM-256QAM. We will extend the testing to other QAM level as time permits.

In this baseline test there were a few corrected FEC codewords for all of the 3 profiles.

7.5. Test to discover failure points, noise across entire channel

The next test determines at what points would the CMTS start seeing codeword errors. The idea here is to increase the (AWGN) noise floor on the channel and see how the system performs. For each modulation order: QAM 16, 64, 256 and each with 2 different pilot patterns, the task was to identify the average RxMER of the channel at which

- the first corrected/ uncorrected codeword is seen
- 100% corrected/uncorrected codewords are seen

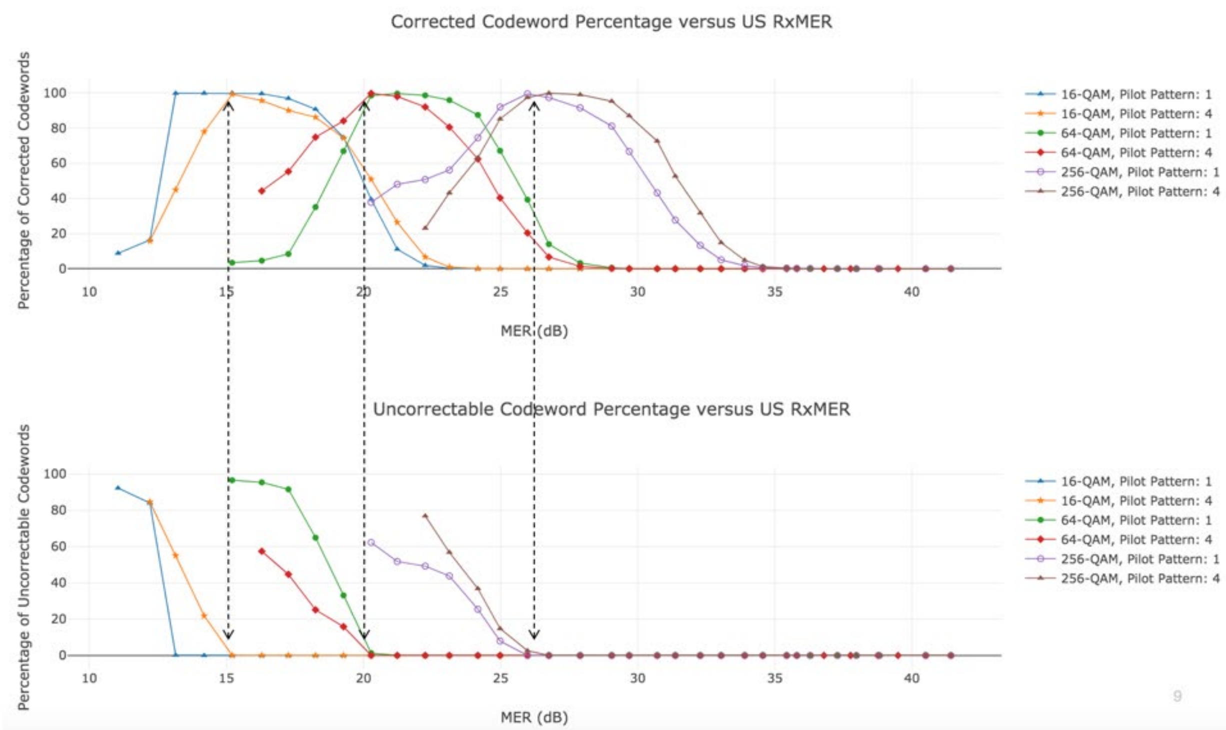

Figure 34 – Average D3.1 US RxMER levels for Correctable and Uncorrectables

We identified the different corrected codeword failure levels at the CMTS, and the corresponding RxMER values tied to those failure point. As seen in the figure there is an inverted S-curve growth in corrected codewords as noise increases. As the number of corrected codewords reaches 100%, the number of uncorrectable codewords start increasing. As expected the lower the modulation order, the more noise is needed to get to the first corrected, 100% corrected and first uncorrectable codewords. In the upstream, we can see the number of uncorrectable codewords go up to high levels, as this CMTS has been configured to operate only on this single OFDMA US. In addition, the CMTS does not change the affected IUC at this time.

8. OFDMA Channel Configuration

8.1. US Channel Location

Though the supported upstream frequency range starts at 5 MHz on a D3.1 plant, the few European operators we have been collaborating with, have seen noise issues in the lower part of the spectrum. Rather than debug one too many things at once, the operators have made a choice to start their OFDMA channels at 20 MHz or 23 MHz and higher. The European operators have the luxury of a higher split at 65 MHz, so can still fit in a 40 MHz OFDMA channel.
The channel also has excluded sub carriers for the return sweep generator. Also, typically the lower frequencies are excluded subcarriers or forced to be unused carriers.

8.2. TaFDM

DOCSIS 3.1 also supports simultaneous Time and Frequency Division Multiplexing (TaFDM) between SC-QAM and OFDMA channels. This means both OFDMA and SC-QAM can simultaneously operate on the same frequencies, divided in time. This allows for the use of OFDMA across the entire spectrum, while maintaining backward compatibility with legacy DOCSIS SC-QAM channels. The figure below from [MULPIv3.1] provides an example of how TaFDM can operate with an OFDMA channel sharing the same spectrum as four SC-QAM channels.

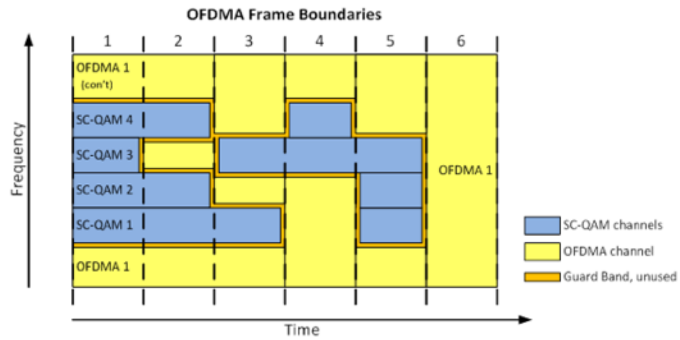

Figure 35 – Time and Frequency division multiplexing

In some initial testing, when using the TaFDM feature between SC-QAM & OFDMA channels, it has been observed that the throughput varies depending on the set of channels in use. In one example configuration, there was one 40 MHz OFDMA channel and 4 SCQAM channels in between as per above diagram. The throughput for the time with OFDMA-only portion in use is the highest, the SCQAM-only is the lowest and when they are time sharing the throughput is somewhere in between. To achieve a good OFDMA throughput, consecutive spectrum is needed without the use of SC-QAM channels. The CMTSs tends to schedule OFDMA traffic first in the OFDMA area before it schedules OFDMA bursts in the TaFDM area. If that area is in the lower part of the US spectrum, then the OFDMA also has to deal with the ingress noise which is more typical at the lower frequencies up to 20 MHz. Some operators have turned off TaFDM as that feature is not quite mature on the CMTS implementations and has not gone through enough CMTS-CM system debug.

8.3. Ranging location

On each CMTS, there needs to be a configuration for the ranging zone within the OFDMA channel. Initial field testing has shown that that moving this ranging zone above the 18-20 MHz range yields better stability for the D3.1 CMs. Moving the ranging zone to a cleaner and stable part of the spectrum ensures that the CMs come online and stay online on the OFDMA channel.

Figure 36 – OFDMA Channel configuration example

8.4. Minislots

Minislots are 8 sub carriers or 16 subcarriers wide. Number of symbols in time for an OFDMA frame(K) is in the range 6-36 symbols wide. When a single subcarrier is excluded, the CMTS needs to readjust the Minislot locations, as no excluded subcarriers are allowed within a minislot.

Exclusion-Bands or Zero-Modulated sub-carrier could be easier to configure in steps of whole minislots. At the time of testing earlier this year, some CMs were unable to handle complex UCDs. Along with the modulation order, the minislot also needs to be configured with the appropriate choice of pilot patterns. There is a choice of 7 pilot patterns (for each minislot size), among the pilot patterns tested, pattern 4 looked to be more robust and had relatively fewer lost traffic compared to pattern 1.

8.5. IUC / Profile management

It is intended that the burst descriptor associated with the data profile IUC 13 be configured as a robust OFDMA profile usable by any DOCSIS 3.1 CM served by that upstream channel. The CMTS uses data profile IUC 13 for all OFDMA data grants to modems which have not completed registration. The CM transmits data using the OFDMA Burst Descriptor for IUC 13 prior to registration.

During or after modem registration, the CMTS has the option of assigning the modem to use any data profile specified in the UCD. Typically, the Burst Descriptors for data profiles other than IUC 13 will be configured for higher performance than IUC 13, although not all of these Burst Descriptors will be usable by all modems. IUC 13 is the lowest common denominator profile, it used by all CMs before registration and after registration for sending mac management messages. Data Profile IUCs (IUC # 5, 6, 9, 10, 11, 12, (and 13)) can use the following modulation orders BPSK, QPSK, 8-QAM, 16-QAM, 32-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, 1024-QAM.

Initial CMTS implementations have a restriction on the number of profiles available on a channel. This is somewhat restrictive. Some operators have started with 2 IUC profiles, as an example: IUC-13 (Default/Fallback): using 16-QAM, and an IUC-12 (Data): using a mix of 64-QAM in low frequency, 256-QAM in higher frequency area

8.6. OFDMA Profile Flapping

After registration, the CMTS grants OFDMA bandwidth for data transmissions to a CM using one of the CM's assigned OFDMA Upstream Data Profile (OUDP) IUCs. The CMTS cannot grant data bandwidth to a CM using an IUC not specified as one of that CM's assigned OUDP IUCs. Upon successful completion of a transaction assigning one or two assigned OUDP IUCs to a CM, that CM needs to be ready for transmitting data using the assigned IUCs.

A CM supports 2 US Profiles/IUCs at a time. A CM starts on the OFDMA channel with IUC 13 (say for example 16 QAM). At a later point the CM is assigned an additional IUC (e.g. IUC 12, say 256 QAM). When CMTS sees US FEC errors on the secondary profile (IUC 12 in this example), it chooses to rectify the situation. A CMTS can reassign the CM a new IUC (say IUC 11, with 64 QAM in areas of high noise and 256 QAM elsewhere) via DBC messages. The CMTS continues to use the default IUC-Profile 13 to forward traffic to avoid packet loss during IUC change, when the DBC is in process. In practice, this means is that the US capacity for the CM is changing intermittently quite significantly which leads to a degraded performance and user experience. We observed this in multiple MSO lab trials and also in our testing.

8.7. Upstream channel evalution tools

Because it is expected that not all upstream data profiles will be usable by all modems, a CMTS might wish to evaluate a modem's performance using a particular profile before assigning that profile to be used. There are two tools to aid the CMTS in gathering information about upstream profile performance: upstream probes, and upstream Data Profile Testing bursts.

A CMTS uses upstream probes for ranging-related functions such as determining transmit pre-equalizer coefficients and additionally using an upstream probe to take an RxMER measurement. To do this, the CMTS grants P-IEs in a P-MAP message with the "MER" bit set. When the CMTS receives the probe transmission corresponding to such a grant, it performs the RxMER measurement.

Some types of upstream profile performance measurements cannot be performed using probe bursts, like FEC performance or count CRC errors for a particular profile. Probe bursts cannot be used for these purposes since they carry no information. D3.1 systems support sending/receiving upstream Data Profile Testing bursts. The CMTS first assigns a Data Profile Testing SID to the modem on one or more upstream channels. (Transmit channel set encodings can be sent as part of a DBC transaction.) The CMTS then sends a grant to a Data Profile Testing SID. The CM responds to this grant by sending a Data Profile Testing burst in the grant.

8.8. US profile change CER based vs RxMER based handling

There are a few different ways in which the CMTS can handle US profile issues for a CM.

The first method is using the codeword error rate, or the number of uncorrected codewords seen from a CM. In this case the CMTS knows that the IUC (modulation and pilot patterns) is not good enough for the CM to successfully transmit on the US. The CMTS can then chose to move the CM to a lower modulation/higher pilot pattern profile.

The second method is using the US RxMER, or the measured upstream RxMER seen from transmissions from a CM. In this case the CMTS based on the RxMER levels, makes a judgement based on thresholds that the signal is not good enough for the CM to successfully transmit on the US using that IUC. The CMTS then choses to assign an appropriate IUC/profile.

A third method is to combine the two metrics and decide on when the profiles need to change. As seen in this discussion the need for US RxMER is paramount, to design IUC/profiles and creating profiles across US channels across the plant. Currently US RxMER support on CMTSs is limited by accuracy issues and also not having any TFTP support to move the data to an external data lake.

9. US PMA

Similar to the downstream, we want to understand how effective a custom designed profile/IUC will behave on the upstream and the kind of benefits customized profiles will give us.

9.1. Baseline test on 256-QAM

This test runs upstream traffic from the CM to the CMTS and checks the RxMER at the CMTS

1. Send traffic to modem starting at two different packet sizes and rates.
2. Repeat for different QAM orders.

We notice a few US packets being dropped, we are still working through the reasons why we have these packet drops even in the baseline case. The figure below shows the US RxMER measured at the CMTS and the traffic rates and number of packets sent and received.

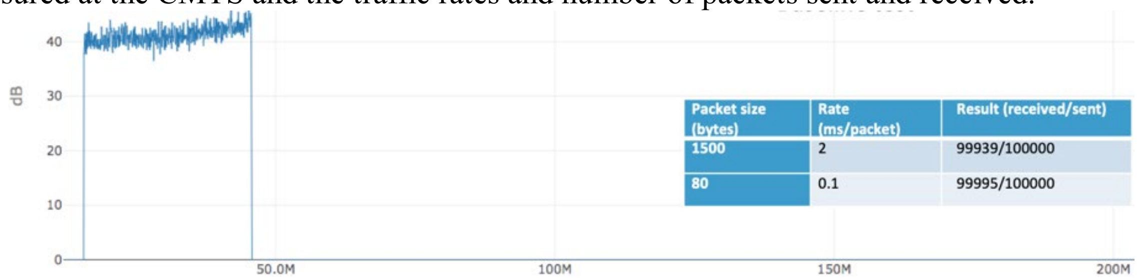

Figure 37 –D3.1 US Baseline test

9.2. US Noise injection test

The test here was to determine behavior of single IUC/profile after noise is injected at specific location within the channel. A narrow band noise was injected (at widths of 0.2Mhz, 1 MHz, 5MHz into the US channel. As expected the packet drops increase for additional noise width. The figures below show the US RxMER measured at the CMTS, each of which show the noise ingress, and the traffic rates and number of packets sent and received.

For all of these tests, we are still continuing to analyze the relationship between packet sizes, packet rates, minislot sizes/ number of symbols in a frame etc. (Also having an SC-QAM US as a primary channel makes it a bit hard to tease out the distribution of load across the channels.)

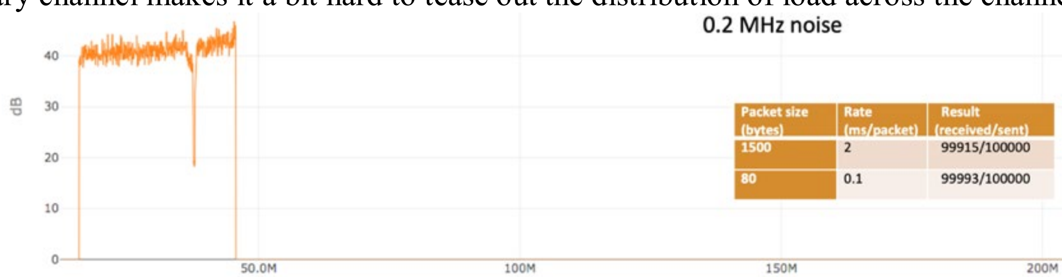

Figure 38 –D3.1 US Noise Injection test 0.2 MHz

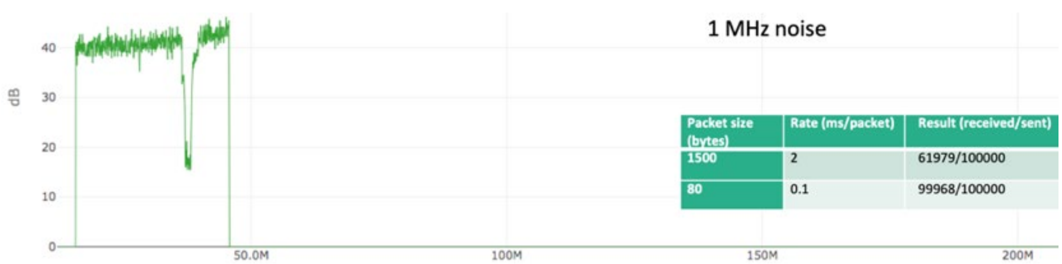

Figure 39 –D3.1 US Noise Injection test 1 MHz

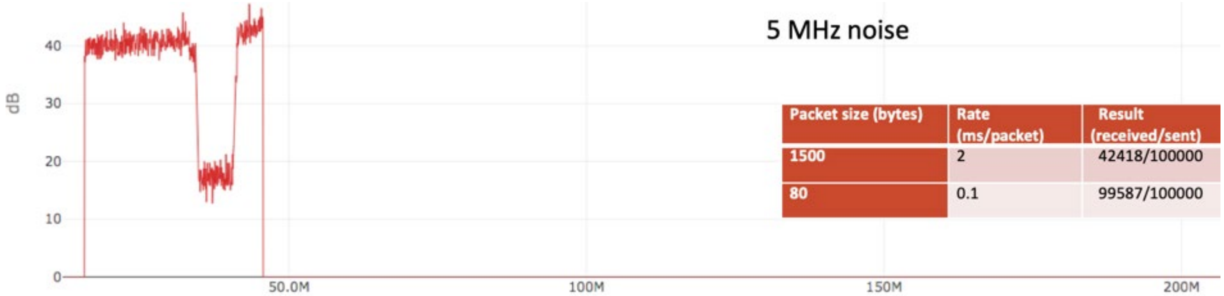

Figure 40 –D3.1 US Noise Injection test 5 MHz

9.3. US PMA test

This test repeats the 5 MHz noise injection test, but this time with a customized profile, which works around the noise ingress. This profile allows the channel to operate at the same level in the presence of the ingress, as the system did without any ingress. Upstream profile design gives the same basic benefits as it does in the downstream, it has the ability to allow the channel to operate with uncorrectable codewords, which means a more stable usage of the upstream.
The potential gains and benefits of the US PMA will be likely more than the Downstream, given that the upstream suffers from a lot more noise ingress and funneling effects from the plant. As we get more US field data from D3.1 CMs, we will gain a better understanding of the kind of profiles we need to create and the number of profiles needed etc.

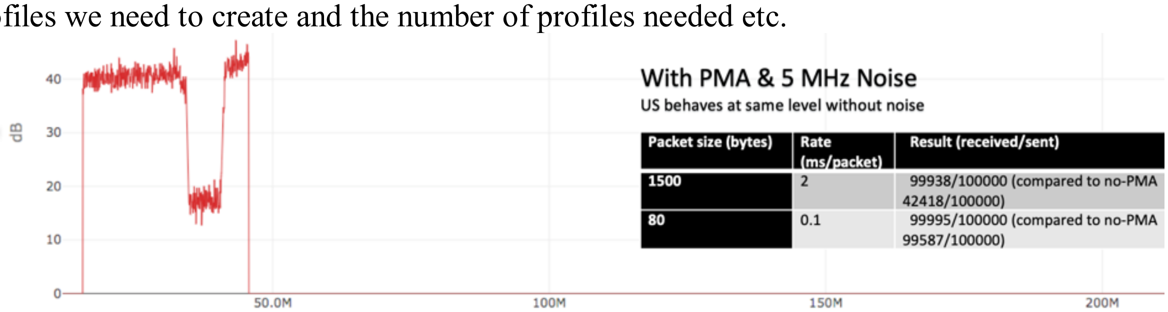

Figure 41 –D3.1 US Noise Injection test 5 MHz

Conclusion

A D3.1 network is a highly capable access network. The system features including the LDPC FEC, time and frequency interleaving and other signal processing enhancements have made a D3.1 network very robust to noise ingress. When the noise is severe custom modulation profiles using a Profile Management Application is absolutely need. The customized profile creation and configuration increases the reliability of the network (upstream and downstream) and in addition maximize the capacity. Support for 8-10 profiles on the downstream will be important in the years to come, as OFDM becomes the cornerstone of DOCSIS technology. MAC layer settings around CM-STATUS needs to be optimized to get the best traffic connectivity and experience for the customer. With increased size of the OFDM channels, the impact of optimizing them is also huge. Some good engineering is need in the configuration of the DS and US channels and their settings. Support for upstream signal quality data from a CMTS will become more important as operators start deploying more OFDMA.

Abbreviations

bps	bits per second
CM	cable modem
CMTS	cable modem termination system
DOCSIS	Data over Cable System Interface specification
FEC	forward error correction
HFC	hybrid fiber-coax
Hz	hertz
ISBE	International Society of Broadband Experts
PMA	Profile Management Application
RxMER	receive modulation error ratio
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

[PHYv3.1] DOCSIS 3.1, Physical Layer Specification, CM-SP-PHYv3.1-I16-190121, January 21, 2019, Cable Television Laboratories, Inc

[MULPIv3.1] DOCSIS 3.1, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I17-190121, January 21, 2019, Cable Television Laboratories, Inc.

[PMA-TR-CL] DOCSIS 3.1 Profile Management Application Technical Report CM-TR-PMA-V01-180530

[D31PMA-INTX16] DOCSIS 3.1 Profile Management Application and Algorithms (2016), Greg White and Karthik Sundaresan, CableLabs. https://www.nctatechnicalpapers.com/Paper/2016/2016-docsis-3-1-profile-management-application-and-algorithms

[LTEFreq] https://en.wikipedia.org/wiki/LTE_frequency_bands

Acknowledgements

Thanks to different operators across the world for sharing their DOCSIS 3.1 data with CableLabs. Thanks to Comcast, Shaw, Videotron, Vodafone Germany, NOS Portugal, for graciously sharing their data, and involving CableLabs in their D3.1 trials and deployment efforts. Thanks to Dan Rice, for inviting CableLabs to be part of the Comcast efforts for D3.1 testing and analysis, and thanks to Ray Hammer on running countless tests to understand the CM-STATUS behavior and thanks to Paul Schauer for collecting and sharing data. Thanks to Nader Foroughi, Shaw, for involving CableLabs with their PMA trials. Thanks to Peter Wittman at Vodafone and thanks to João Fernandes at NOS for involving CableLabs with upstream OFDMA testing.

