Field Experiences with US OFDMA and using US Profile Management

A Technical Paper prepared for SCTE•ISBE by

Karthik Sundaresan
Distinguished Technologist
CableLabs
858 Coal Creek Circle, Louisville, CO 80303
3036613895
k.sundaresan@cablelabs.com

João Pedro Fernandes
HFC Engineer
NOS
Portugal
+ 351 931005187
joao.fernandes@nos.pt

Jay Zhu
Senior Engineer
CableLabs
858 Coal Creek Circle, Louisville, CO 80303
3036613312
j.zhu@cablelabs.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

DOCSIS® 3.1 is now largely deployed in the field, but so far it has been an Orthogonal Frequency Division Multiplexing (OFDM) Downstream only endeavor. Operators are beginning to test Upstream Orthogonal Frequency Division Multiple Access (OFDMA) and are discovering various intricacies in getting the US OFDMA to work robustly. We CableLabs and NOS, a cable operator in Portugal have been working together and have been focused on DOCSIS 3.1 OFDMA in general and collaborating on the Upstream Profile Management Application (PMA). Defining appropriate Upstream profiles for an OFDMA channel affects the stability of the modems on the channel and also the capacity realized. We implemented data collection agents through the CMTS CLI to collect Upstream RxMER for each CM, or on other CMTSs collect the data manually for each CM. US RxMER looks very different than the Downstream RxMER, due to the noise funneling characteristics on the HFC plant. NOS engineering and CableLabs collaborated on various US PMA algorithms and this paper will describe some of the new methodologies we have developed. The NOS operations team is doing a field trial with the US data collection and then configuring the profiles/IUC (interval usage code) generated by PMA on a live plant to understand the impact and behavior. This paper focuses on the lessons we have learnt from the DOCSIS 3.1 upstream field trial and production systems.

2. US OFDMA background

DOCSIS 3.1 introduces OFDM downstream signals and OFDMA upstream signals to achieve robust operation and provide more efficient use of the spectrum than previous DOCSIS versions. OFDMA for the upstream path is a multi-user version of OFDM, and assigns subsets of subcarriers to individual CMs.

2.1. US splits

The DOCSIS 3.1 system will have options of several split configurations that can be exercised based on traffic demand, services offered and the capability of the cable plant. In the upstream direction, the cable system may have a 5-42 MHz, 5-65 MHz, 5-85 MHz, or 5-204 MHz pass bands. A D3.1 CM supports one or more of the following upstream upper band edges, (as long as one is 85 MHz or greater): 42 MHz; 65 MHz, 85 MHz, and/or 204 MHz. The DOCSIS 3.1 Network supports a minimum of two independently configurable OFDMA upstream channels with each occupying a spectrum of up to 95 MHz. A DOCSIS 3.1 CM is capable of transmitting on OFDMA channels and legacy single carrier-QAM channels (SC-QAM) at the same time (as controlled by the CMTS). There are no legacy SC-QAM channels above a frequency of 85 MHz.

2.2. OFDMA channel basics

The OFDMA upstream multicarrier system is composed of either 25 kHz or 50 kHz wide subcarriers. In the upstream, the subcarriers are grouped into independently configurable OFDMA channels each of up to 95 MHz encompassed spectrum, totaling 3800 25 kHz spaced subcarriers or 1900 50 kHz spaced subcarriers. When configured for 2K FFT (Fast Fourier Transform), the CMTS uses the subcarriers in the range 74 <= k <= 1973, where k is the index of the subcarrier defining the OFDMA signal. When configured for 4K FFT, the CMTS number uses subcarriers numbered in the range 148 <= k <= 3947.

The parameters of the two OFDMA channels can be independently configured thereby optimizing configuration based on channel conditions. The table lists the Upstream Channel parameters, from [PHYv3.1]

Table 1 – D3.1 Upstream OFDMA Parameters

Parameter	Value	
Upstream Sampling Rate	102.4 MHz	
Upstream Elementary Period Rate	1/102.4 MHz	
Channel bandwidth	10 MHz, …, 95 MHz	6.4 MHz, …, 95 MHz
IDFT size	2048	4096
Subcarrier spacing	50 kHz	25 kHz
Symbol duration	20 µs	40 µs
Maximum active subcarriers (95 MHz)	1900	3800
OFDMA Cyclic Prefix size	0.9375 µs, 1.25 µs , 1.5625 µs , 1.875 µs 2.1875 µs , 2.5 µs, 2.8125 µs 3.125 µs, 3.75 µs, 5.0 µs, 6.25 µs	
OFDMA Roll-off Period Size	0 µs, 0.3125 µs , 0.625 µs , 0.9375 µs 1.25 µs , 1.5625 µs, 1.875 µs, 2.1875 µs	
OFDMA Modulation orders	(BPSK), QPSK, 8-QAM, 16-QAM, 32-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, 1024-QAM, 2048-QAM, and 4096-QAM	

2.3. US OFDMA Frames & minislots

DOCSIS 3.1 Upstream transmission uses OFDMA frames. Each OFDMA frame is comprised of a configurable number of symbols (K = 6 to 36). Several transmitters may share the same OFDMA frame by transmitting on allocated subcarriers of the OFDMA frame. The structure of an OFDMA frame is depicted in Figure 1. The upstream spectrum is divided into groups of subcarriers called minislots. Minislots have dedicated subcarriers, all with the same modulation order ('bit loading'). [PHYv3.1] specifies two minislot sizes by specifying the number of subcarriers per minislot. There are 8 or 16 subcarrier minislots, a minislot is always 400KHz wide (25KHz subcarrier *16, or 50KHz subcarrier *8). Minislots have dedicated subcarriers, all with the same modulation order ('bit loading'). Though the span of the minislot is always 400KHz, the length of the minislot in the time is the same as the number of symbols(K) of the frame. An operator can configure the number of symbols in an OFDMA frame to pick an appropriate size for the minislots on a channel. A CM is allocated to transmit one or more minislots in a transmission Burst. The modulation order of a minislot, as well as the pilot pattern used may change between different transmission bursts and are determined by the profile definition. Several transmitters may share the same OFDMA frame by transmitting on their allocated minislots on the OFDMA frame.

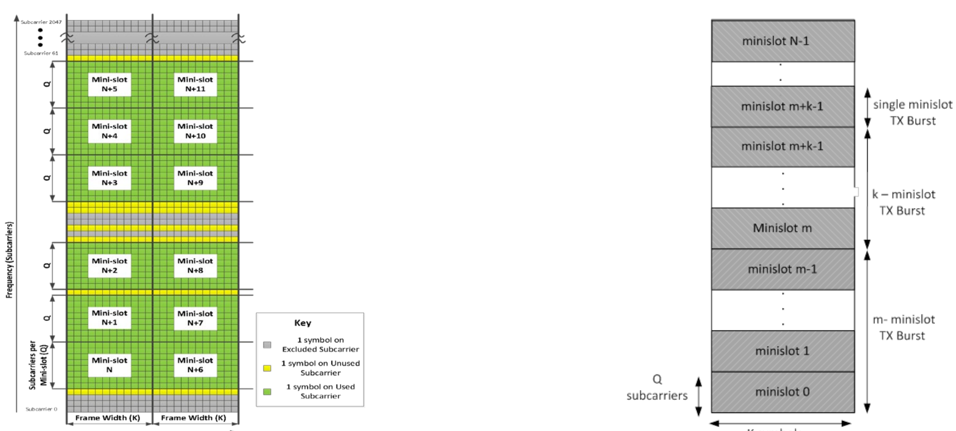

Figure from [MULPIv3.1]

Figure 1 – Upstream OFDMA Minislot Layout and Grants across Minislots

There are two types of minislots edge minislots and body minislots. An edge minislot is the first minislot in a transmission burst, and body minislots are used for all other minislots in a transmission burst. See [PHYv3.1] for minislot usage with exclusion bands etc.

Each minislot is comprised of pilots (P), complementary pilots (CP), and data subcarriers. Pilots are used by the CMTS receiver to adapt to channel conditions and frequency offset. Pilots are subcarriers that do not carry data and encode a pre-defined BPSK symbol known to the receiver. [PHYv3.1] also specifies complementary pilots which are subcarriers that carry data, but with a lower modulation order than other data subcarriers in the minislot. If the modulation order used for data subcarriers in the minislot is M, the complementary pilots are used with modulation order equal to the maximum between M-4 and 1 (BPSK). For example, if the bit loading in a minislot is 10, Complementary Pilots use 6 bits.

For each minislot size, seven pilot patterns are defined, see figure 2. Pilot patterns differ by the number of pilots in a minislot, and by their arrangement within the minislot. The different pilot patterns enable the operator to optimize its performance (physical layer rate and pilot overhead) according to different conditions and variations of SNR with frequency. Each pilot pattern defines edge and body minislots.

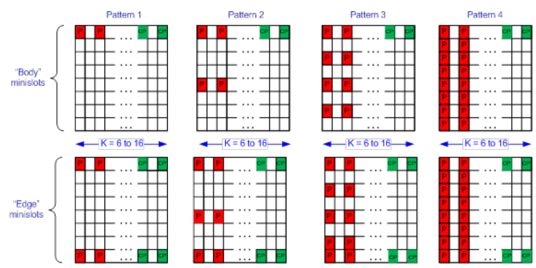

Figure from [PHYv3.1]

Figure 2 – Sample Pilot Patterns

2.4. US FEC

An upstream grant from the CMTS indicates which minislots are assigned to a given burst and which upstream profile is to be used by the CM. 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. Per [PHYv3.1], OFDMA use three Quasi-Cyclic Low-Density Parity-Check codes (QC-LDPC) for the upstream transmission, as depicted in Table below

Table 2 - FEC coding parameters

Code	LDPC Code Rate	Codeword size in bits	Information bits	Parity bits
Long	8/9 = 89%	16200	14400	1800
Medium	28/33 = 85%	5940	5040	900
Short	¾ = 75%	1120	840	280

2.5. Upstream Noise funneling

The DOCSIS upstream behaves differently than the DOCSIS downstream. As the DOCSIS downstream signal is transmitted (from the CMTS to the CM) through the branched cable topology, the signal becomes attenuated and weaker at each branching and with cable distance. The strong signal at the headend now requires amplification to adequately reach the subscriber. Noise or interference in downstream frequencies that enter the cable system say at a node or an amplifier only affects customers past that point.

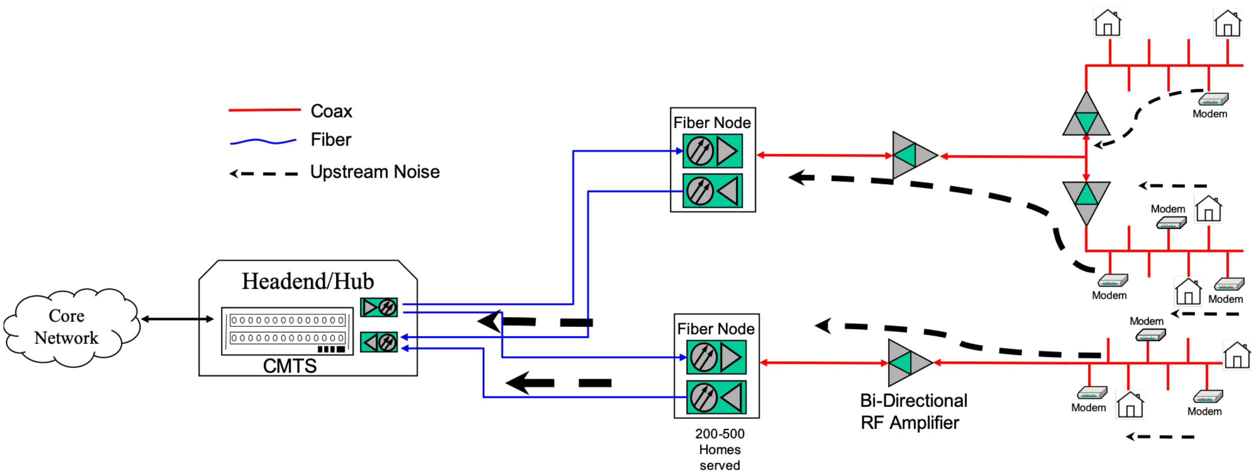

Figure 3 – US Noise funneling

In the downstream direction, there is one location from where the signals enter the HFC plant, specifically the cable modem termination system (CMTS) in the headend or at a node location with distributed access architectures (e.g. Remote PHY). The operator has control of the signal at that point and along the network, to ensure it reaches every CM. From the headend to the CM, the RF signal fans out in a star topology network in a point to multipoint fashion. It is the opposite on the upstream/return path: the RF signals enter the plant from every home that is attached to the plant, and all of those signals combine together as they travel to the headend. Typical of all point to multipoint networks, the noise from every device on the network gets combined as it travels upstream and is finally received on the upstream port at the CMTS. This is known as the noise funneling problem, as shown in the figure 3. The additive nature of noise has a large impact at the CMTS upstream receiver. A small amount of noise from every unterminated cable or loose connector enters the upstream path and is combined with other ingress noise and then amplified as it travels upstream and gets funneled up to the CMTS causing it to be unable to decode the communication from the CM. HFC plant segmentation somewhat mitigates the noise funneling effect by reducing the number of combined signals traveling on a given upstream path.

2.6. Common Upstream Issues

There are many upstream impairments on the cable plant, some of the causes and ways operators tackle them are described here. Cable operators are constantly working to mitigate upstream ingress noise. Cable operators have a hard time pinpointing where ingress noise is entering the plant. A spectrum analyzer can allow an operator to visualize the noise problem, but it doesn't tell the operator the location of the problem. A common method, to track ingress, is disconnecting each segment of the plant at a node to see the impact on noise levels on the spectrum analyzer. This is laborious and customer impacting and as the noise sources are bursty and intermittent, there is a bit of luck involved in the technician finding the source. A technician needs to visit the worst ingress locations, and make physical changes to the network. Often the ingress locations are also a place where signals leak out of the cable plant. There are many tools available which are based on this concept of looking for test signals leaking out of the plant.

Thermal noise at the amplifiers and fiber optic link noise can be sources of upstream impairments. Other ingress noise sources on the upstream path include impulse noise from loose connectors, reflections from unterminated splitters or taps, cracked cables, common path distortion due to corroded connectors or cabling issues. Many of these are resolved by physical repair to the damaged cable network infrastructure (replacing drop cable, servicing defective network components such as power supplies). Reflections can be mitigated through equalization coefficients used within the DOCSIS technology. Identifying problem frequencies and avoiding them for upstream transmission is another method, to side step some problems at least temporarily. Using external data analytics an operator can potentially identify ingress locations and use frequency changes to dynamically avoid this type of noise. Impulse noise can be partially mitigated by the use of DOCSIS forward error correction (FEC) and interleaving algorithms. Again, a technician might need to visit the worst locations to make physical changes to the network. Correct amplifier alignment and setting up correct power levels in the upstream path also prevents issues like laser clipping.

2.7. CMTS Profile / IUC mangement

The CMTS assigns OFDMA Upstream Data Profile (OUDP) IUCs to the CMs based on the measured plant conditions. It is intended that the Data Profile IUC 13 is configured as a robust OFDMA profile usable by any DOCSIS 3.1 CM served by that upstream channel. Data Profile IUC 13 is used for all OFDMA data grants to modems which have not completed registration.

During or after modem registration, the CMTS has the option of assigning the CM to use any other configured data profile. Typically, the data profiles other than IUC 13 will be configured for higher performance than IUC 13, although not all of these profiles will be usable by all modems. The CMTS assigns the CM either one or two data profiles (IUCs) for each OFDMA channel in the modem's Transmit Channel Set. This can be assigned during Registration, and can be changed after Registration using dynamic messages (DBC transaction). After registration, the CMTS grants bandwidth on the OFDMA channel for data transmissions to a CM using one of the CM's assigned OUDP IUCs.

2.7.1. Upstream Profile Testing

Because it is expected that not all upstream data profiles will be usable by all CMs, a CMTS can evaluate a CM’s performance using a particular profile before assigning that profile to be used for live user traffic. The DOCSISv3.1 technology [MULPIv3.1] provides various tools to aid the CMTS in gathering information about upstream profile performance. A CMTS performs such an evaluation in vendor-specific ways, usually revolving around modulation error ratios or codeword error ratios. These tools are based on two types of transmissions: upstream probes, and upstream Data Profile Testing bursts.

2.7.2. Upstream Probes and RxMER Measurements

A CMTS uses upstream probes for ranging-related functions such as determining transmit pre-equalizer coefficients. A CMTS also has the option of using an upstream probe to take an RxMER (received modulation error ratio) measurement. The CMTS grants probe opportunities to a CM in a P-MAP message with the "MER" bit set. When the CMTS receives the probe transmissions from the CM corresponding to such a grant, it performs the RxMER measurement and uses the results in its decision making. It also populates the corresponding MIB object or can upload a RxMER per subcarrier file via TFTP, for the operator’s information.

2.7.3. Upstream Data Profile Testing Bursts

Some types of upstream profile performance cannot be measured using probe bursts. For example, a CMTS might wish to gather information on FEC performance or count CRC errors for a particular profile. Probe bursts cannot be used for these purposes as don’t carry any information, and instead the CM&CMTS can send/receive upstream Data Profile Testing bursts. Per [MULPIv3.1], to command a CM to send an upstream Data Profile Testing burst, the CMTS first assigns an OUDP Testing SID to the CM on one or more upstream channels. The CMTS then sends a grant to an OUDP Testing SID, the IUC of this grant is an Assigned OUDP IUC currently assigned to the modem. The modem responds to a valid grant to any of its OUDP Testing SIDs by sending a Data Profile Testing burst in the grant. The Data Profile Testing burst from the CM is a 64-byte Ethernet packet, with counting pattern in payload bytes beginning with 0x01, continuing with 0x02, 0x03, etc., and ending with 0x2E (count is re-started at 0x01 in each successive packet). The CM fills the grant with DOCSIS frames. The modem treats all grants to its OUDP Testing SID(s) as grants to a single flow existing across all OFDMA channels to the SID has been assigned.

2.7.4. IUC/ Profile change

The CMTS assigns one or two OUDP IUCs to a CM, once the assignment is successful, that CM is ready for transmitting data using the assigned IUCs. After registration, the CMTS grants OFDMA bandwidth for data transmissions to a CM using one of the CM's assigned OUDP IUCs.

A CM supports 2 US Profiles/IUCs at a time. A CM starts on the OFDMA channel with IUC 13 (e.g. say set to 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) dynamically via DBC messages. The CMTS continues to use the default IUC-Profile 13 to forward traffic to avoid any packet loss during IUC change, when the DBC is in process. In practice, this means that the upstream capacity for the CM is changing intermittently as it switches between profiles, which could lead to a degraded performance and user experience.

2.8. Upstream OFDMA : the need for Profile Management

The DOCSIS 3.1 specification fundamentally changes the nature of information delivery on the cable plant, and the way HFC networks will be maintained and managed. In a significant change from previous DOCSIS versions, the OFDM/OFDMA channel does not use a one-size-fits-all modulation scheme; rather, the modulation can be optimized based on actual plant conditions at different frequencies and individual devices. CMs and CMTS that communicate with a cleaner signal can utilize an efficient high-order-modulation, while devices that have a degraded signal will use more robust modulation, all on the same channel.

The DOCSIS 3.1 toolbox provides a wide range of modulation choices that can be used to fine-tune the transmissions to get the best performance from the current network conditions. To manage the optimization of these settings across the population of devices, the CMTS uses the concept of Upstream profiles. An upstream profile defines the modulation order (i.e., bit loading) and pilot pattern on each of the minislots on the channel (up to 237), spanning (up to 3800 or 1900 subcarriers) across the OFDMA channel.

DOCSIS 3.1 specifications [MULPIv3.1] provides for defining multiple upstream profiles, where each profile can be tuned to account for the specific plant conditions that are experienced by a set of CMs. A well-designed, optimized set of modulation profiles allows an upstream channel to operate with robustness and a lower SNR margin, potentially allowing a channel to deliver an overall higher throughput. In addition, it can allow for communication to devices by providing service even in situations where significant plant impairments exist.

The application that implements this optimization logic is external to a CMTS, enabling the most efficient use of profiles across channels and CMs. For an operator, it also allows uniform operation of such algorithms across different CMTS platforms. This profile optimization and profile creation functionality is implemented as an ‘application’ running outside the CMTS and is known as the Profile Management Application (PMA). Managing profiles manually for an operator for thousands of CMTS Upstream channels is a labor-intensive process requiring a deep understanding of the channel conditions. The calculation and recalculation of profiles would overload human operators and PMA simplifies this process significantly.

2.9. Designing Profiles for US OFDMA channel

Now in the upstream, the noise from every house and every network element gets accumulated and is seen at the upstream receiver on the CMTS. Now a CMTS receiver can measure the received modulation error ratio (RxMER) for each CM, see figure 4 for some example measurements from a live network. In the upstream, this signal to noise signature for each of the CMs (that are sharing the upstream channel) starts looking very similar, as they all share the same noise across the channel with slight differences due to the signal levels itself or some in house network problems. This means common profiles can be designed for many CMs experiencing similar noise conditions and most CMs will be able to use a common profile. (This is very different from the behavior seen in the downstream, where different sets of CMs have very different noise signatures.) The variation in the Upstream RxMER from sample to sample for a single modem itself is much greater than the mostly tight RxMER variations that we are accustomed to seeing in the downstream. For CMs which suffer more noise, they can be put into a different profile optimized for their particular noise environment. The modulation orders within a profile can vary appropriately across the spectrum as per the noise levels in that part of the spectrum.

Figure 4 – US RxMER Measured on multiple CMs on a upstream Channel

The upstream Profile Management Application (PMA) can automate this design of the profiles on upstream channels across various segments in the cable plant. Reading the upstream RxMER from the CMTSs on the network, processing the RxMER information with intelligent algorithms to create profiles, and then configuring the newly optimized profiles on the CMTS are the primary functions an upstream PMA solution accomplishes. Configuring optimized profiles brings solid reliability to the upstream network connection and also increases the capacity in parts of the spectrum which can accommodate higher modulation orders.

Given this understanding of the upstream plant behavior and the RxMER signatures, the question now is, what are the best algorithms to design upstream profiles which give operators robust upstream operation as well as increase the throughput? We address this problem with a few different solutions in chapter 4.

3. US OFDMA Deployment at NOS

In this chapter we share some of the Upstream OFDMA field rollout experiences by the NOS access engineering team in Portugal earlier this year. After some lab trials the effort quickly moved to limited field trials with internal employees and then to customers, in a phased approach. The support for OFDMA on CMTSs and CMs is maturing, but some of the initial trials led to a lot of lessons learned. Many of these bug fixes and feature improvements may take time to make it into operator networks, and we hope that sharing those notes here helps other operators with their Upstream trials and roll-outs. There may be gaps in system implementations which could use some thought and new solutions. Data collection for Upstream is based on the CMTS and these capabilities are still limited across the different implementations. This data is necessary to ensure that an operator can design the correct IUCs for each of the US OFDMA channels. Like many European HFC plants , the NOS plant has a 65 MHz EuroDOCSIS® split, which leaves open a reasonable amount of spectrum of OFDMA channels. We discuss the reasons we decided to keep the trials above 23.5 MHz and also share some trial experiences below that. The HFC upstream plant clean-up is always good practice for all operators and this trial needed more of that to make the OFDMA operations more robust. Monitoring the network and developing a few Key Performance Indicators (KPIs) for the DOCSIS Upstream are important to evaluate the changes made to the network and understand the performance of the network.

3.1. Channel location

The NOS plant in Portugal has upstream spectrum up to 65 MHz and currently there are 2 to 3 SC-QAMs per serving group and up to 4 upstream serving groups (US SGs) per downstream serving groups (DS SGs). With a 65MHz split and SC-QAM channels being already located at the upper edge, the natural choice to place OFDMA was the lower part of the spectrum, with a hope to take advantage of the enhanced flexibility of OFDMA.

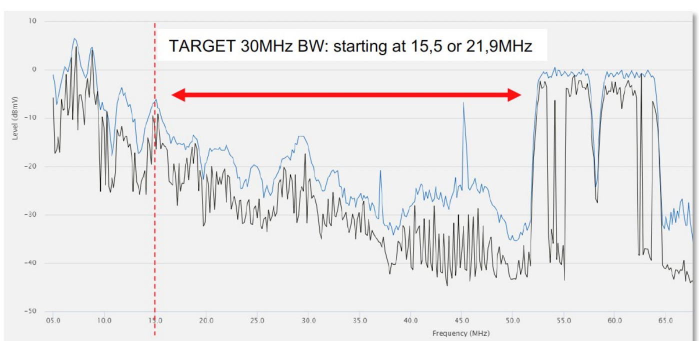

Figure 5 – Target spectrum for OFDMA channel placement

The initial trials started with OFDMA channels located at 15MHz up to 45.5 or 51.9MHz, depending on existing number of total SC-QAMs in the US SG, which were either 2 or 3. The spectrum below 40MHz was never used before for DOCSIS upstream and the initial analysis indicated that a high level of noise was sometimes present. Even after a fair bit of work to clean the outside plant, it turned out that using spectrum below 23.5MHz made the channel too susceptible to impulse noise. The result of this was that the OFDMA channel was affected, leading to channel impairment. So, the decision was made to locate the lower edge of the OFDMA channel at 23.5MHz for better stability.

At the time of this writing another approach was being tested: This approach chose to locate the SC-QAMs from 30MHz and use the upper part of the upstream spectrum for the OFDMA channel. As this is a much cleaner spectral area, more bits/hz can be extracted from the same bandwidth for the OFDMA channel. Now this will only work if the SC-QAMs can work error-free in the lower zone of the spectrum. To ensure the stability of the US SC-QAMs, upstream agility features will definitely help.

Another approach that needs to be explored is to use two smaller OFDMA channels instead of a larger, single one. The idea is to have each of these OFDMA channels straddle the 2-3 SC-QAM channels. This way a smalled but highly reliable OFDMA channel can be obtained, providing a baseline capacity that is always present. The second OFDMA channel, due to its location in the lower end of the spectrum, will be much more susceptible to impairment or US IUC downgrades/flapping. This channel can then act as reserve capacity available most of the time to most of the CMs.

The field deployment was on two different CMTS platforms, both with different upstream capabilities, the table below describes some of the parameters chosen for the deployment.

Table 3 – Upstream Channel Parameters Chosen on CMTSs

Parameter	CMTS vendor 1	CMTS vendor 2
OFDMA Bandwidth	20.1 – 45.5/51.9 MHz	15.5 - 45.5 MHz 21.9 - 51.9 MHz
Number of IUCs	2 IUCs – 13; 12	2 IUCs – 13; 12
Variable bit loading	Flat Profiles	Variable
K	18	16
Pilot Pattern	4	2
Subcarrier Spacing	50	50
Roll Off	224	96
Cyclic Prefix	320	192

85 MHz Trial Learnings: In limited areas we also tried out OFDMA channels on an 85 MHz plant, as that was available in some areas. Here the OFDMA channel was located from 22MHz to 85MHz – with an exclusion band for 2 SC-QAM channels. Initially for the OFDMA channel we needed to create a 1.6MHz band in each side of the QAMs with 64QAM modulation, as per vendor recommendations, to protect the rest of the OFDMA channels. This configuration was revealed to be unstable with loss of IUC12 and OFDMA impairment. What solved the problems was adding a 500kHz guard band at the exclusion zone. Once this was added and the Upstream performed well. The calculated OFDMA capacity was 344 Mbps, and along with the couple of SC-QAM channels, actual speed tests gave us numbers of ~400 Mbps. We plan to use this selectively in very noisy and tough upstream environments, reusing existing 85MHz-ready equipment when available.

3.2. Phases of turning on OFDMA

OFDMA is a new technology for the Operators and for the CMTS and CM vendors, so a very cautious approach was used to deploy this technology in the network and enabling traffic on those OFDMA channels. After some initial lab trials, a three-phase approach was used. This was in parallel with continuous testing and upgrading of the CMTS and CM software with bug fixes.

Phase 1 – Activating the OFDMA channel, but not using it for customer data

In this phase we were able to check CMTS and CM behavior, regarding registration, CM management stability, etc. No service flows were assigned to the OFDMA channel. This was done to minimize the customer impact as the upstream traffic was not allowed to use OFDMA channel and instead the traffic continued to use the existing SC-QAM channels. To accomplish this, different techniques had to be used on the two CMTS models deployed in the network. We were able to build the data collectors and processors, and start acquiring the KPIs selected and RxMER data files, which gives us visibility over the new upstream spectrum and the new OFDMA technology

Phase 2 – Using the OFDMA channel for customer data, but not upgrading their speed tier.

Once the data gathering and modem registration issues were crossed we moved into the phase of using the OFDMA channel for traffic. This phase allowed us to see the real impact of having live data traffic on the channel. Profile/IUC downgrades would happen in one CMTS just based on codeword (FEC) errors. So, with data traffic we could see the quality of the IUCs that we had created and the degree of readiness of the plant. During this phase IUCs were refined and there was a massive amount of work to go out and correct the plant when possible. At this point the customer was not given any speed upgrades, so any channel impairments on the OFDMA meant that the customers would fall back to DOCSIS3.0 capacity and still get the same capacity as they did before.

Phase 3 – Upgrading the upstream speed to customers

Once the IUC definitions were mature and robust, in this phase, the Maximum Sustained Traffic Rate was upgraded in the CM Config files for the eligible customer base. In this phase the impact could be significant as any channel impairments would directly mean less capacity and the customer might not be able to the reach the top speed offered in the service.

3.3. Technology maturity

At the time testing and field trialing started, both CMs and CMTSs Upstream implementations were quite immature. There were lots of feature limitations, and some features were not implemented at all, which constrained the roll-out across the footprint. Some examples in the CMTS were a very limited number of IUCs supported, CMTS management of IUCs not reacting to codeword errors, no Upstream PNM implementation, a very limited set of configuration parameters like pilot patterns, subcarrier spacing, cyclic prefix, etc. CMs were more comparatively mature from the beginning as regard support for the OFDMA set of features. Now both CMTS and CMs had numerous bugs that needed to be isolated and solved. This was by far the longest part of the work needed to get OFDMA out in to the field and make it generally available for customers.

Some challenges included issues like handling DBC messages to change IUC definitions as appropriate. Other CMTS devices had issues collecting US RxMER data reliably, and this was a key requirement to understand the plant and create appropriate profiles. Other initial limitations forced Service flows to be assigned to OFDMA channels, even when the operator was not planning on sending data traffic on those channels. Many CMTSs only support 2 IUCs one of which included IUC 13. The profile/IUC definition also only allowed 4 exception zones for modulation order changes, and allowed only one exclusion zone. What this meant was the that the allowed configuration options were somewhat limited. Load balancing of US traffic across SC-QAMs and OFDMA channels were also not fully mature.

3.4. Initial IUC definition for OFDMA chanel

The CMTS OFDMA channels need some initial IUC definitions to get the upstream channel up and running. For this initial phase of trials and the initial IUC definitions there was no US RxMER information available, as no OFDMA channel was provisioned in the network. Hence a PMA could not be used to create optimized profiles. So, another method was devised to define the initial IUCs for the 6000+ US SGs. The basic idea was to measure the noise on the plant when no SC-QAM or OFDMA channel was present and then use that to estimate the US MER.

3.4.1. Data collection for initial IUC

The Upstream data was captured with our data collection system (Viavi XPERTrak), which was continuously collecting signal power measurements in the spectrum where the OFDMA channel was to be located. The measurements were done before the OFDMA channel was turned on. This consisted of an average dBmV level value for each 250kHz bin of spectrum from 5 to 65MHz, collected every 15 minutes. Many days of historical data was used as input to a custom-made algorithm that returned IUC definitions for each US SG.

datetime	frequency	avg
2020-02-09T03:15:03Z	5001	-48,3
2020-02-09T03:15:03Z	5251	-48,4
2020-02-09T03:15:03Z	5501	-48,3
2020-02-09T03:15:03Z	5751	-48,5
2020-02-09T03:15:03Z	6001	-48,5
2020-02-09T03:15:03Z	6251	-48,4
2020-02-09T03:15:03Z	6501	-48,3
2020-02-09T03:15:03Z	6751	-48,5
2020-02-09T03:15:03Z	7001	-48,5
2020-02-09T03:15:03Z	7251	-48,7
2020-02-09T03:15:03Z	7501	-48,5
2020-02-09T03:15:03Z	7751	-48,5
2020-02-09T03:15:03Z	8001	-48,4
2020-02-09T03:15:03Z	8251	-48,2
2020-02-09T03:15:03Z	8501	-48,6
2020-02-09T03:15:03Z	8751	-48,5
2020-02-09T03:15:03Z	9001	-48,6
2020-02-09T03:15:03Z	9251	-48,5
2020-02-09T03:15:03Z	9501	-48,6
2020-02-09T03:15:03Z	9751	-48,7
2020-02-09T03:15:03Z	10001	-48,3

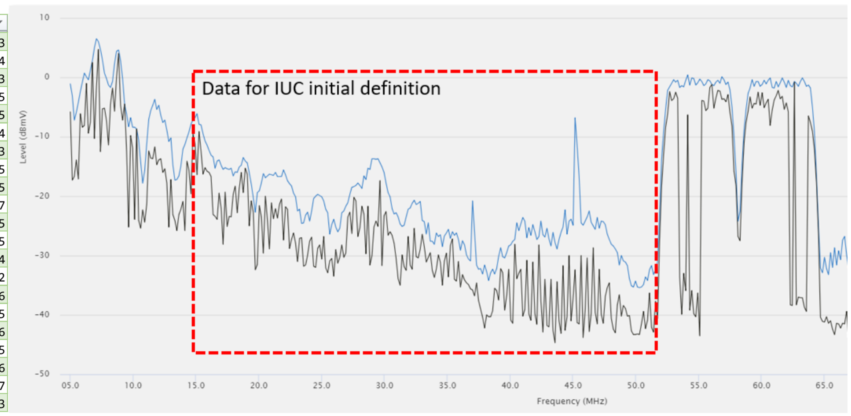

Figure 6 - Example power level data used to calculate intial IUC definitions

3.4.2. Initial IUC Definition

The MER was calculated from the noise level as measured in the signal power measurements in the previous section and the assumption that the CM signal arrives at the CMTS at 0dBmV. For IUC 13 we used the maximum power level found in the data set for each bin as the assumed noise level. For IUC12 we used the average level, choosing a bit more relaxed value for use in designing a higher profile. The power level data had to be tailored and adapted to minislot boundaries, and CMTS specific implementation rules needed to be accommodated on the design of profiles. This algorithm was implemented in a software program that could automatically read the data and output the needed CMTS commands for every serving group.

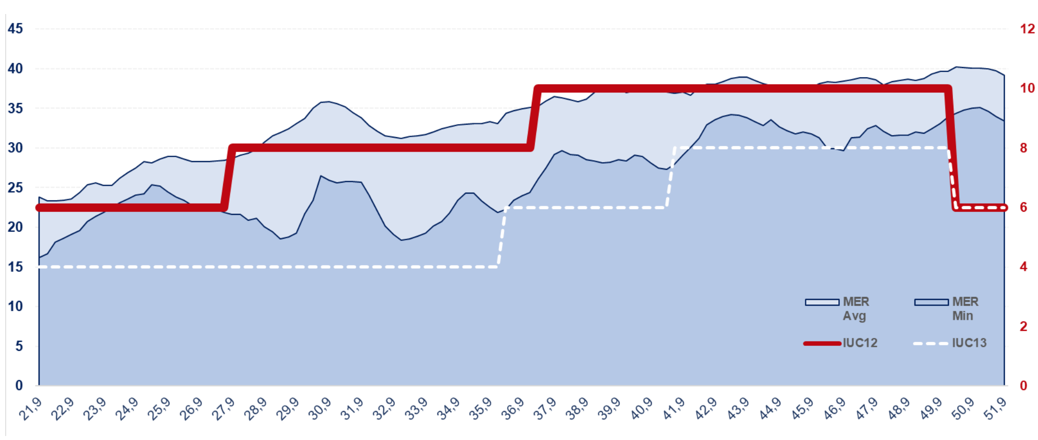

Figure 7 - Example IUC definition from the deduced MER

The OFDMA features available at the time along with the limitations of each CMTS vendor had to be taken into account for this algorithm to convert this power level measurements into MER and ultimately into modulation orders. Also, worth noting that this method only works before the OFDMA channel is turned on in the plant. As soon as CMs start transmitting in the OFDMA channel, this method cannot be used any longer as the external signal power measurements cannot separate the signal from noise.

3.5. Upstream Plant readiness

Level alarms were setup, in our data monitoring system, for the OFDMA band before OFDMA signals were activated. Whenever the threshold conditions were crossed, technicians were dispatched to check the plant, and correct noise sources or pinpoint customer installation problems. This happened for about 8% of the US SGs. Each and every US SG was again verified and acted upon just before provisioning a customer with a higher top speed that actually needs the OFDMA channel capacity.

Based on the KPIs described below, a technician truck roll was ordered for the SGs with high number of impairment-hours. Technicians would go through the plant, amplifier by amplifier, tap by tap, looking for the source of the noise impairment and trying to isolate it. If the source of the problem happened to be in the outside plant, the technicians would take steps to solve it. If the source was identified to be inside a home, a filter blocking part of or the complete upstream would be used on the drop, and a then a home network maintenance work item would be scheduled.

This is very similar to the normal maintenance of the upstream in DOCSIS3.0. The main difference here was the large amount of work it took to do this for the whole network in a short amount of time. The other difference from DOCSIS 3.0 is also the use of this new spectrum which is in the lower part of the upstream spectrum and very noisy. As with the legacy 3.0 spectrum, this is be a process which will continue on a daily basis, it is now even more demanding as this new lower portion of the spectrum is more susceptible to noise and interference. This will continue to be ongoing operations work, just like it was for the upper part of the spectrum before OFDMA. High pass filters that used to be deployed to block noise can no longer be used in the OFDMA region, making it harder and more expensive to maintain the outside plant.

3.6. KPIs for monitoring

We defined a new KPI to monitor OFDMA upstream channels: IUC-usage-hours. IUC-usage-hours is defined by number of hours in a day a CM uses each IUC on a particular OFDMA channel. This quickly became the main KPI that we tracked during these initial deployments and Upstream OFDMA roll-out. We tracked this metric of IUC-usage-hours over days, weeks and now months. As an example, a CM could be on IUC 12 for 22 hours a day and IUC 13 for 2 hours a day, while another CM perhaps could be in an OFDM-impaired state (partial service) for an hour a day.

Now in order to understand the reason for the variance in IUC-usage-hours, a homemade tool was built with the following additional KPIs:

- Codewords per IUC (total, errored): Per CM statistics on codeword errors, to understand the customer impact.
- Receive and Transmit Power: Per CM Rx/Tx power
- Data volume / Heavy user: The data volume metric was considered useful to diagnose different CM behaviors in the same US SG, as CMTS depends on codeword error ratios (CER) for IUC

management and its algorithm depends on traffic. Heavier users are more likely to hit thresholds and experience IUC changes.

These set of metrics proved essential, as this was the data used to fine tune IUCs, to make decisions along the trials and subsequent rollouts, to order plant cleaning or to dispatch a technician to the customer premises.

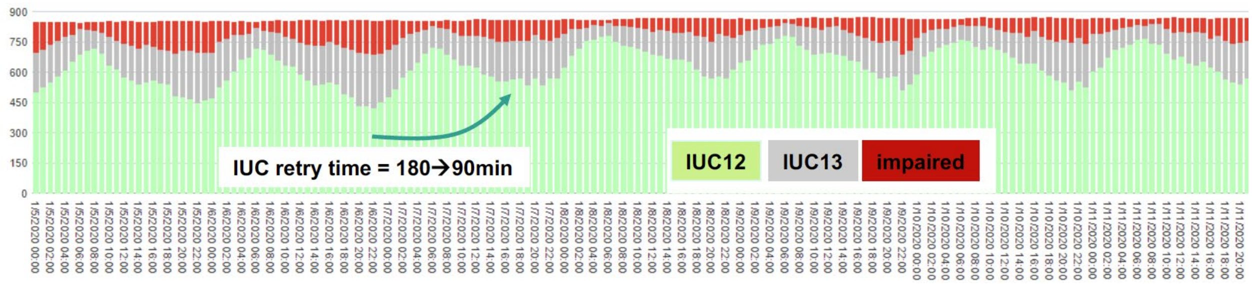

Figure 8 – Tracking IUC-usage-hours per day (across CMTS)

CMTSs automatically change the IUC used for each CM based on CER thresholds. The effect of one of the adjustments made during the process of deploying OFMDA can be seen in Figure 8 above. When changing the IUC retry interval from 180 minutes to 90 minutes, the percentage of time CMs spent in IUC 12 and 13 increased, and the time spent without OFDMA capacity decreased. At this point in time (i.e. prior to the PMA field trial) the CMTS IUC change mechanism algorithm was flawed and customers could sometimes have packet loss for a few seconds during IUC downgrade. So, the goal became not wanting to have many IUC switches. This testing allowed the operations team to decrease this time without any noticeable service degradation.

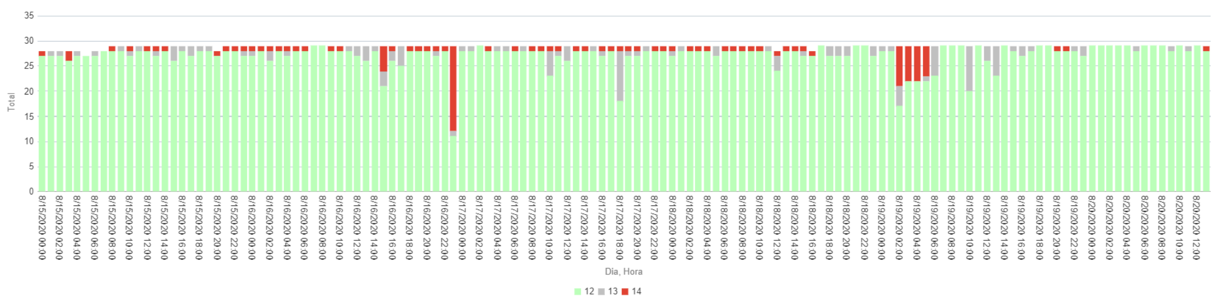

14 – denotes Impaired OFDMA operation

Cordas com maior % em IUC13 + IUC14 - UPSTREAM

Ano	Semana	Corda	CMTS	Total CMs	# Amostras	%IUC12	%IUC13	%IUC14	%IUC13 ou 14
2020	33	VCH13 C27D2	PNO1CMTS013	12	1884	34	17	49	66
2020	33	PNF03 3JD1	PDS1CMTS004	9	1502	22	29	48	78
2020	33	PTM25 2JD1	PTM1CMTS005	24	3666	35	18	47	65
2020	33	AGU03 C41D2	QUE1CMTS012	15	2518	38	22	40	62
2020	33	SES01 C35D1	FGT1CMTS005	4	672	43	19	38	57
2020	33	PVZ05 3JD1	VCD1CMTS003	11	1815	40	25	35	60
2020	33	TMA10 2JD1	FGT1CMTS004	20	3144	23	42	35	77
2020	33	BBA05 C35D1	PNO1CMTS009	6	1004	31	35	33	69
2020	33	SJM02 C53D1	RME1CMTS003	2	504	67	0	33	33
2020	33	FIG01 3JD2	FIG1CMTS005	11	1847	42	25	33	58
2020	33	PAR04 3JD1	PAR1CMTS006	23	3751	67	0	33	33

Figure 9 – Operational dashboards displaying CM IUC distribution in one US SG and top US SGs with CM impairments

3.7. US RxMER Data Collection from the CMTS

The DOCSIS 3.1 CCAP OSSI specification defines a standard procedure to use SNMP to start upstream RxMER measurement on the CMTS and have the encoded binary files that comply to the specification defined format uploaded to a TFTP server. However, this feature is not currently supported by all of the available CMTSs. In our field data collection practice, 2 other methods were used to collect upstream RxMER data based on what the CMTSs support.

3.7.1. Standard SNMP + TFTP method

The DOCSIS 3.1 CCAP OSSI specification defines MIBs and data formats to allow upstream RxMER data collection. The high-level steps are described below, on a per CM basis

- The data collector specifies the CM MAC address and CM upstream channel interface number, sets parameters for upstream RxMER data measurement, and specifies the TFTP server address and directory for file upload through the MIBs on the CMTS
- The CMTS performs upstream RxMER measurement for the specified CM and OFDMA channel
- The CMTS uploads the encoded file that contains the upstream OFDMA RxMER measurement data to the specified TFTP address and directory
- The data collector reads the uploaded binary file, decodes it and stores it into the data service for PMA profile calculation

3.7.2. SNMP + SFTP method

Some CMTSs may comply to part of the DOCSIS 3.1 CCAP OSSI specification and may support the SNMP MIBs for upstream RxMER measurement. However, they may require different methods to gather the collected data, such as SFTP. The steps are described below, on a per CM basis

- The data collector specifies the CM MAC address and CM upstream channel interface number, sets parameters for upstream RxMER data measurement through the MIBs on the CMTS
- The CMTS performs upstream RxMER measurement for the specified CM and OFDMA channel
- The CMTS stores the encoded file that contains the upstream OFDMA RxMER measurement data in the CMTS's local storage (the CMTS may only store a limited number of PNM files)
- The data collector retrieves the stored binary file through SFTP from the CMTS, decodes it and stores it into the data service for PMA profile calculation

3.7.3. CLI method

Some CMTSs may not support the MIBs for upstream data collection that are defined in the DOCSIS 3.1 CCAP OSSI specification. However, they may support measuring the upstream OFDMA RxMER data and presenting the data through the command line interface (CLI). In this case, we can parse the CLI output to collect the upstream OFDMA RxMER data. The steps are described below, on a per CM basis

- The data collector uses a CLI client to specify the CM MAC address and CM upstream controller/interface/channel number in the CMTS CLI and collects the CLI output (the CMTS may not support data collection triggering, instead, it may have a time interval to be configured to periodically perform upstream RxMER measurement)
- The CLI client automatically parses the CLI output and generates JSON formatted upstream OFDMA RxMER data
- The data collector stores the JSON formatted data into the data service for PMA profile calculation

3.8. US RxMER Data Collection in field deployment for PMA trials

CableLabs DOCSIS Common Collection framework (DCCF) software can collect PNM data from the plant. However, the CMTSs were not fully compliant to PNM specification for PNM testing and reporting. One CMTS could generate the RxMER files and store them in its filesystem, but it didn't implement the TFTP upload. So, for the sake of this field trial, some shell scripts were developed to trigger RxMER measurement through SNMP and fetch the encoded binary files from the CMTS using SFTP.

Data was collected for all the CMs and all the SGs in field trial every 2 hours, for a couple of weeks. A software limitation and the total number of CMs imposed this 2-hour limit. In total, more than 13,000 RxMER files were collected. For practical reasons, all SGs we collected data belonged to the same CMTS.

4. US PMA Algorithms

The upstream RxMER data has different features compared to what we see in the downstream. Although one could start with the downstream PMA's clustering algorithm to calculate upstream OFDMA modulation profiles, given the difference of the Upstream RxMER data (as explained in section 2.9) it become clear that a simple "clustering of CMs" approach would not be a good solution. Due to noise funneling since the US RxMER were similar many of the profiles are expected to have similar characteristics. So, we developed several new candidate algorithms for comparison and improving the robustness of the upstream. The immediate goal for this initial rollout was focused around developing IUCs which modems could stay on without profile flapping. Each profile is expected to be used by a group of CMs that have similar channel characteristics.

Each of the Minislots in the Upstream OFDMA channel can be configured to use a different modulation order. This allows the operator to optimize the upstream transmissions across the wide frequency band (10-96 MHz) of the channel. The specific choice of modulation order selected for each minislot is communicated to the CMs in the form of an IUC (modulation profile) which allows the CM to modulate the signal accordingly. An IUC/modulation profile consists of a vector of bit-loading values, an integer value for each active Minislot in the upstream channel. Since the modulation orders range from QPSK to 4K-QAM, the range of bit-loading values is from 2 to 12.

The PMA generates an IUC 13 that is the lowest common denominator profile, which can be successfully used by all CMs in the Service Group. A CMTS can support up to 7 modulation profiles, including IUC 13. Each CM can be assigned up to 2 modulation profiles at a time, including IUC 13 and an optimized profile for the CM's unicast traffic.

This capability, the ability to optimize the upstream transmission for the channel characteristics of the CM population, is a powerful feature that allows for a significant improvement in robustness and channel capacity. The CMTS and CM perform measurements and report network conditions as a part of supporting PNM functionality in the DOCSIS network. The DOCSIS 3.1 Upstream PNM Measurements includes: US active and quite probes, Triggered spectrum capture, US equalizer coefficients, Impulse noise statistics FEC statistics, Histogram, Channel Power and the RxMER per subcarrier. So far, we are basing the PMA profile creation algorithms on the US RxMER data, in the future one can include other upstream data sets to fine tune the profiles that we create.

4.1. Differences from Downstream Algorithms

The downstream PMA algorithm looks into the variances among the CMs, clusters the CMs into groups, and assign different modulation profiles to each CM group. The algorithm can use a data snapshot that's captured only once from the CMs, or it can use pre-processed data (average, minimum, percentile etc.) from each CM's multiple historical data captures.

On the upstream, each CM's RxMER data tend to have similar patterns when they are captured during the same time slot because of noise funneling. The variance is much more within the CM's RxMER captures over a relatively long period of time. Hence, the upstream PMA algorithm focuses more on optimizing the channels robust operation and CMs' upstream capacity over a certain period (when the operator does not plan to change upstream profiles frequently such as changing the profiles every few hours). The candidate upstream PMA algorithms also consider the fact that the CMTSs automatically upgrade/downgrade the upstream modulation profiles for each CM based on their monitored FEC performance, which makes the time clustering based profiles more effective.

Problem: Given a set of US RxMER data samples from CMs, return optimized profile definitions.

We have two classes of algorithms we developed and are field trialing, one is developing using the percentile method and the other is using time clustering methods. These methods and their variants are described below.

4.2. Percentile method

The CMTS is actively managing the CM's IUC/profile assignment as plant conditions change. The upgrades/downgrades to the CMs' upstream OFDMA modulation profile is based on the CMs' FEC performance or RxMER data. The percentile method for creating profiles is a simple statistical set of methods which can create robust profiles based on the past performance of the plant. The idea is to choose a conservative profile which can fit most of the CMs, most of the time. So, the algorithms arrange the CMs in descending order of their RxMER values and choose to design a few different profiles at a few different percentile values (e.g. 0.5 percentile and 2 percentile)

This method focusses on optimizing the overall channel robustness so that an CMTS can maintain good upstream service for most of the CMs.

4.2.1. *Algorithm 1A : Per CM Percentile*

Inputs: A list of CM RxMER per subcarrier, choice of a percentile numbers that an operator wants to choose at the per CM level and for the profile level.

Outputs: List of robust profile definitions for use on the upstream channel

Algorithm:

- Calculate a representative CM RxMER sample for each CM from the data captured over time
  - For each CM on the US channel
    - Create an artificial 'x^{th}' percentile sample of RxMER for this CM
      - Start with the first sub carrier in the channel
      - Find the 'x^{th}' percentile RxMER value across all the samples for that sub carrier for that CM

        - Repeat for all subcarriers
      - Repeat for each ‘x^{th}’ percentile chosen by the operator.
        - E.g. 0.5th percentile and 2nd percentile could be the two percentile values for ‘x^{th}’. The number of percentile values depends on the number of needed IUCs.
- Create profiles from CMs’ representative RxMER samples (percentile samples)
  - Group each of the ‘x^{th}’ percentile data samples from the CMs (from the previous step)
  - Create new ‘y^{th}’ percentile (normally 0.5% or 0%) samples across each of the CMs’ ‘x^{th}’ percentile groups.
    - Start with the first subcarrier in the channel
    - Find the ‘y^{th}’ percentile RxMER value across all the CM samples from an ‘x^{th}’ percentile group
    - Repeat for all subcarriers
    - Repeat for all ‘x^{th}’ percentile groups
  - Translate each of the ‘y^{th}’ percentile RxMER to Modulation orders per [PHYv3.1] spec.

4.2.2. Algorithm 1B: All CMs Percentile

Inputs: A list of CM RxMER per subcarrier, choice of a percentile numbers that an operator wants to use to create profiles.

Outputs: List of robust profile definitions for use on the upstream channel

Algorithm:

- Create a profile from all of the RxMER samples from all of the data
  - Create a new ‘y^{th}’ percentile sample across all of the CM’s RxMER samples.
    - Start with the first sub carrier in the channel
    - Find the ‘y^{th}’ percentile RxMER value across all the samples for all the CMs
    - Repeat for all subcarriers
  - Repeat for as many ‘y^{th}’ percentile values as many as IUCs are needed.
    - E.g. IUC 13 could be the 1st percentile, and IUC 12 could be the 5th percentile.
  - Translate each of the ‘y^{th}’ percentile RxMER to Modulation orders per the PHY spec.

4.2.3. Algorithm 1C: Remove Outliers + Percentile

Inputs: A list of CM RxMER per subcarrier, choice of a percentile numbers that an operator wants to use to create profiles, and a choice of how to detect outlier CMs.

Outputs: A list of robust profile definitions for use on the upstream channel

Algorithm:

- Find and remove Outlier CMs.
  - Method C1
    - For each CM, across all of its samples: Calculate an average RxMER number for the CM
    - Calculate Standard deviation for each CM, with respect to all of its samples.
    - Remove CMs outside the Confidence interval (e.g.: 98%) lower bound.
      - Remove all samples of all outlier CMs from the data set

- Method C2
  - Calculate average RxMER: Each CM, all samples
  - Remove CMs below 70% of average value (of all CMs)
    - Remove all samples of these CMs (in some cases, this may be an empty operation), from the data set.
- Create a profile set from the remaining RxMER samples
  - Use Algorithm 1A or 1B

4.3. Time Clustering Methods

The CMTS automatically upgrades/downgrades the CMs' upstream OFDMA modulation profile based on the CMs' FEC performance. The hypothesis for the time clustering method is that we can find clusters of CMs with similar RxMER over time and design IUCs based on those clusters to keep the CMs on IUCs for long periods of time. The CM clustering algorithm that's been previously developed and used in downstream PMA is repurposed to calculate clusters from the data captured over time on all CMs. The time clustering methods focus on optimizing the overall channel capacity through days or weeks of time without recalculating or modifying the upstream modulation profiles very frequently. The J value that's used in the downstream for capacity gain measurement is repurposed to measure the theoretical capacity gain from the time clustering based profiles over a time period.

4.3.1. Algorithm 2A: Time Clustering using Actual CM Samples

Inputs: A list of CM RxMER per subcarrier, choice of a number of profiles an operator.

Outputs: A list of robust profile definitions for use on the upstream channel

Algorithm :

- Find Time clusters from all CM's and all of their samples
  - Use actual sample of RxMER for every CM, every measurement
  - Use a clustering algorithm to find groupings across full-RxMER values.
    - e.g. reuse the PMA Algorithm (for Downstream)
    - Find 2-7 clusters
- For each cluster/group of CM-samples
  - Choose an RxMER value per subcarrier, from this groups samples as follows
    - The average value at each subcarrier
    - The minimum value at each subcarrier
    - A certain percentile value (picked by the operator)
    - Some other centroid definition
  - Translate each of the resulting RxMER vector to Modulation orders per the PHY spec.

4.3.2. Algorithm 2B: Time clustering using artificial Percentile Samples

Inputs: A list of CM RxMER per subcarrier, choice of a number of profiles an operator.

Outputs: A list of robust profile definitions for use on the upstream channel

Algorithm:

- For All CM's and all samples

    - Create 199 artificial samples of RxMER for each CM as described below
        - An RxMER value per subcarrier is chosen using a percentile value across all samples of that CM
        - Samples will use percentiles from 0.5% to 99.5% in increments of 0.5%
- Perform clustering with these artificial samples
    - Use Algorithm 2A (This step will find the most common percentiles)

4.3.3. Algorithm 2C: Time Clustering after removing outliers

- For All CM's and all samples
    - Use all samples of RxMER for each CM
    - Create an artificial threshold sample using a percentile value (e.g. 1%)
        - RxMER value per Subcarrier is chosen using a %tile value across all samples
    - For every sample
        - If any individual RxMER within the sample is below artificial threshold sample, then bump up those RxMER values to the threshold MER level using a ceiling-like calculation. This keeps the weight of the sample for other areas in the channel which are unchanged
        - Instead of bumping up samples, we could also ignore the samples
- Now with the outliers removed, use Algorithm 2A.

4.4. Other PMA Considerations

Modulation order of profile from minislot RxMER: There are a few different methods to choose the Modulation order of the profile from the RxMER values of all the subcarriers within that minislot. One could choose one of the following RxMER values to use when translating to the modulation orders

- Average RxMER of subcarriers within minislot.
- Majority RxMER of subcarriers within minislot
- X percentile RxMER of subcarriers within minislot
- Minimum RxMER of subcarriers within minislot

Margins when translating from RxMER to modulation order: Different margins can be applied at different frequencies within the channel. Lower frequencies could have a higher margin, while higher frequencies wouldn't need as much of a margin. Another way to apply different margins is to preprocess the US RxMER data and if there is a higher variation in the data, one can apply a higher margin. This way the operator can lower the US RxMER before sending to PMA algorithm

Capacity gain Calculation/Optimization function: The downstream has a J-Value calculation which can calculate the capacity achieved by a set of profiles. In the downstream the profiles are also weighted by the number of CMs associated with the profile to calculate capacity. Also, the gain is calculated with reference to 256-QAM in the downstream. For the upstream given the dynamic nature of the noise on the plant and the fact that many of the CMs have the same RxMER signature, the idea of an optimization function can be built around how much time is spent on each of the IUC/profiles, i.e. of the many RxMER data samples of all CMs fit which a particular profile, we can weight each profile by that percentage of samples, to calculate an overall weighted capacity. To figure out if an US RxMER sample can fit a profile, the idea would be to establish a threshold across all subcarriers: e.g. only 1 or 2 % of subcarriers can be below the required RxMER level for that modulation order. Also, this comparison should be done on a on minislot basis. A good reference in the upstream to capture the capacity gain would be to compare to 16-QAM or 64-QAM, depending on the plant.

PMA Exception handling: Due to adjacent channel interferences (from adjacent SC-QAMs), there is a need to handle exceptions for the Higher and lower edges of the OFDMA channel. The ideas is that a few number of minislots (2-4) need to be limited to one or two lower modulation orders than normal to account for interference from an adjacent SC-QAM

5. US PMA Field Trial Results

5.1. US RxMER Field Data

The below figure 10 shows a sample of the US RxMER collected at the CMTS, for 4 cable modems within one Serving group. The data was collected for every cable modem, every 2 hours for over 3 weeks.

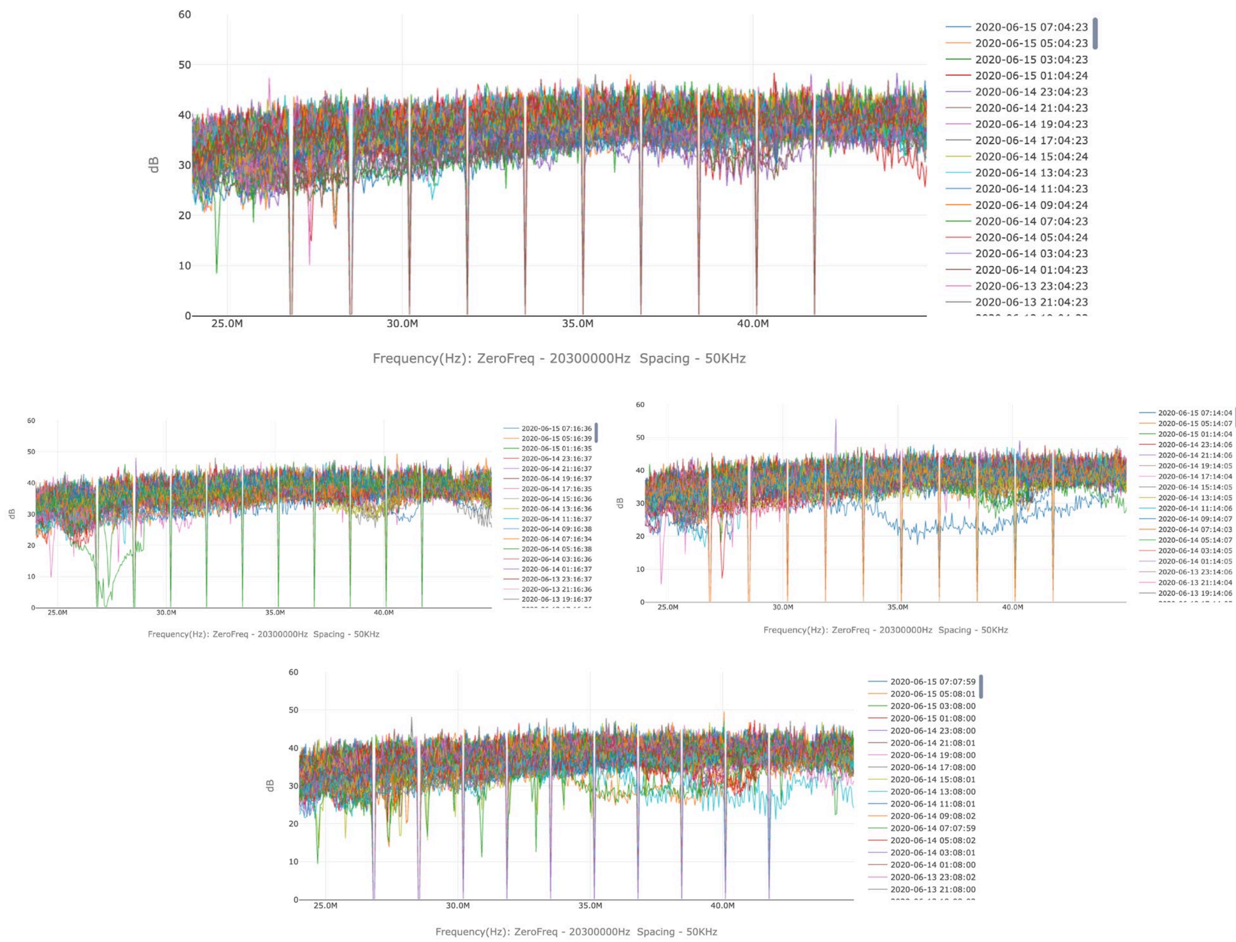

Figure 10 – Us RxMER Samples from 4 CMs

5.2. US IUCs/Profiles designed

The following figures show the profiles designed by the Algorithm 1A,1B,1C, 2A, 2B,2C from the section above. This iteration of profile creation was limited to 2 IUCs due to CMTS limitations. The profile 0 in blue is used as IUC 13 and the profile 1 in orange is used as IUC 12.

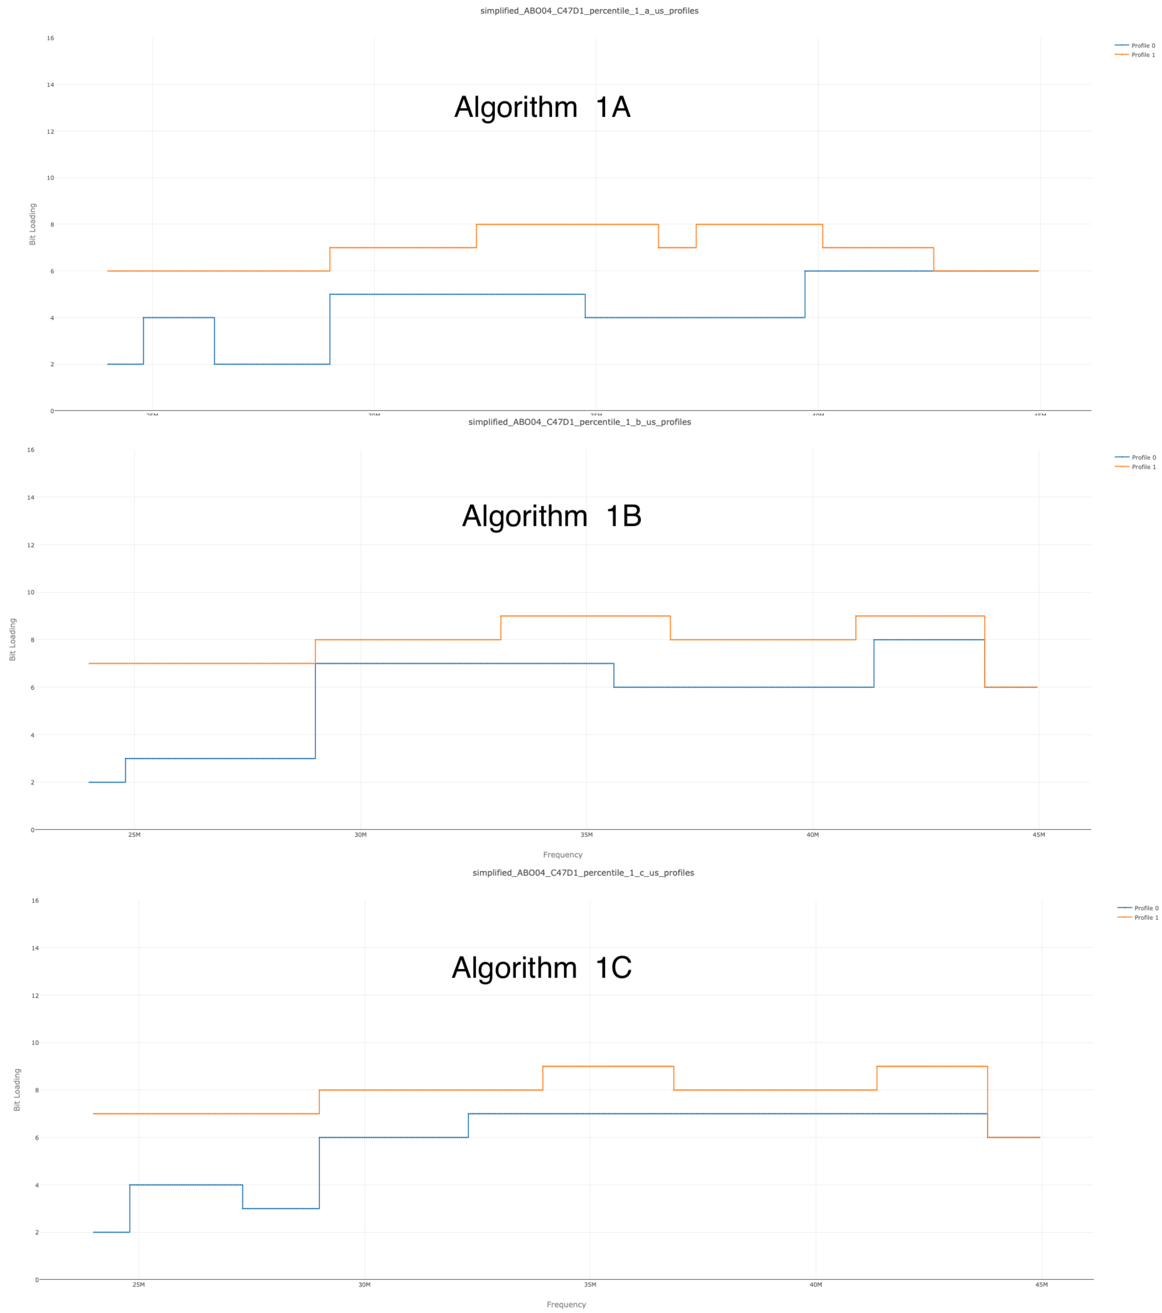

Figure 11 – IUCs created using Percentile Algorithms 1A,1B,1C

The profiles created by algorithm 1A are more conservative than the 1B and 1C as those algorithms remove the outliers or ignore the worst samples across all the CMs equally.

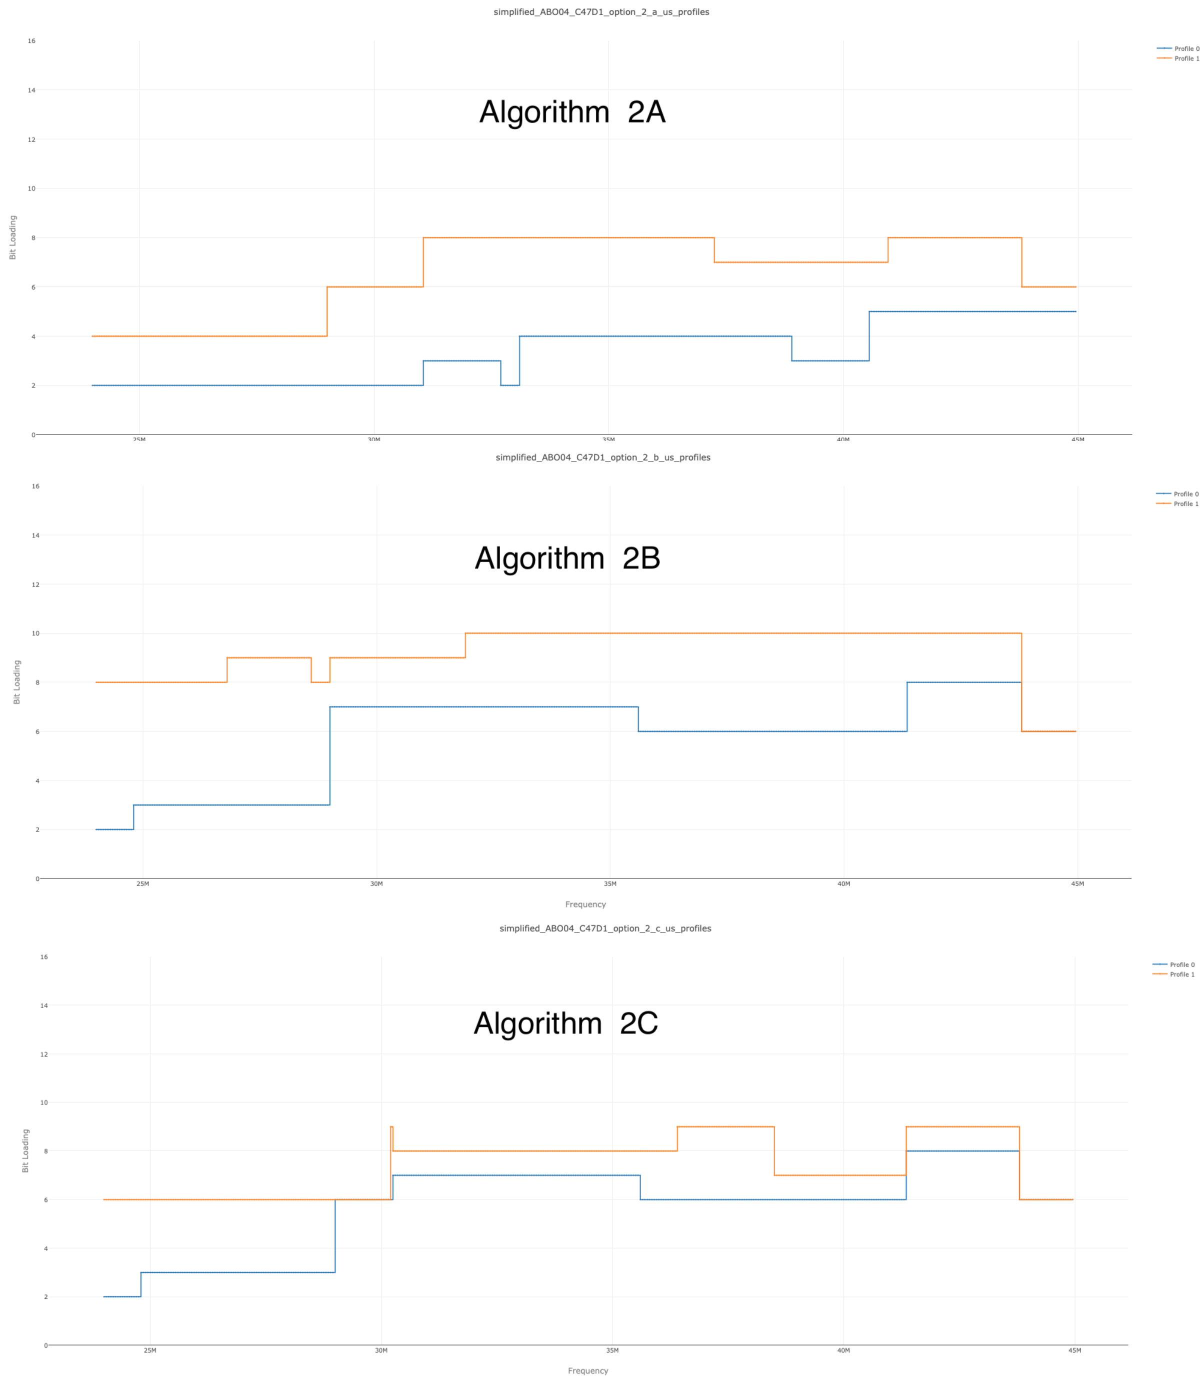

Figure 12 – IUCs created using Time Clustering Algorithms 2A, 2B, 2C

The profiles created by algorithm 2A is more conservative than the 2B as 2A operates on all the original data samples, whereas algorithm 2B operates on a subsampled percentile data set. Algorithm 2C is somewhat in between as it removes the outliers but keeps the weighting of all the samples.

The goal for the field trial is to try out each of these 6 methods of profile creation and uses the profiles they create in the field for a week and observe the performance of these profiles.

5.3. Performance and Stability

The table below depicts the field results for the profiles designed using algorithms 1A and 1B.

Table 4 – Field Results from using PMA Algorithm 1A, 1B

	capacity (Mbps/MHz)				Performance					
Upstream Serving group	Algorithm 1A		Algorithm 1B		Algorithm 1A			Algorithm 1B		
	IUC12	IUC13	IUC12	IUC13	IUC12	IUC13	impaired	IUC12	IUC13	impaired
ABO04 C47D1	5.9	3.7	6.8	4.8	97.8%	1.6%	0.6%	96.4%	3.4%	0.2%
PAR06 2JD1	4.6	2.4	4.1	3.1	99.0%	0.8%	0.2%	95.9%	3.8%	0.3%
PSL03 C53D2	4.4	2.3	4.8	3.6	92.0%	3.6%	4.4%	93.8%	4.5%	1.7%

The Table above shows the IUC-usage-hours percentage for the period and the channel capacity in Mbps per MHz for each algorithm.

IUC-usage data was collected each hour. This CMTS upgraded IUCs without any other consideration X minutes after it was downgraded and X was configured in this trial to 15 minutes. This means the results in the table are a sampling of the actual behavior, but still a good picture of the reality as each set of profiles (from each of the algorithms) was run for more than one week.

CMTS manages IUCs based on FEC errors. When the number of errored codewords crosses a threshold on a configurable codeword window, the IUC is downgraded (from 12 to 13, and from 13 to no-IUC, i.e., channel impaired). So, at this point we are not tracking FEC error rates because the CMTS will essentially downgrade the IUC up to the point that there are no errors.

Due to the working restrictions due to Covid, the field trials were delayed and slowed down a bit, and at the time of writing, we have not had a chance to evaluate the other 4 methods of profile creation. So far Algorithm 1A shows greater stability with the designed IUC 12. Our hypothesis is that profiles based on methods 2A,2B,2C will be likely give us more stable upstream operations. We hope to share those results in the near future.

5.4. US IUCs/Profiles designed with no IUC limitations

We ran the same algorithms on the data from the serving groups to create more than 2 IUCs, we chose 7 IUCs as that will be maximum number supported by the CMTS per channel, in the future. The figure 13 below gives us an idea of the profiles created for a single serving group. In these Serving groups we currently have only ~30 CMs each, so some of the profile definitions are overlapping (we chose percentile values of 0.5,3,6,9,12,15,18.)   For a larger set of CMs we expect the profile definitions to be more spread out.

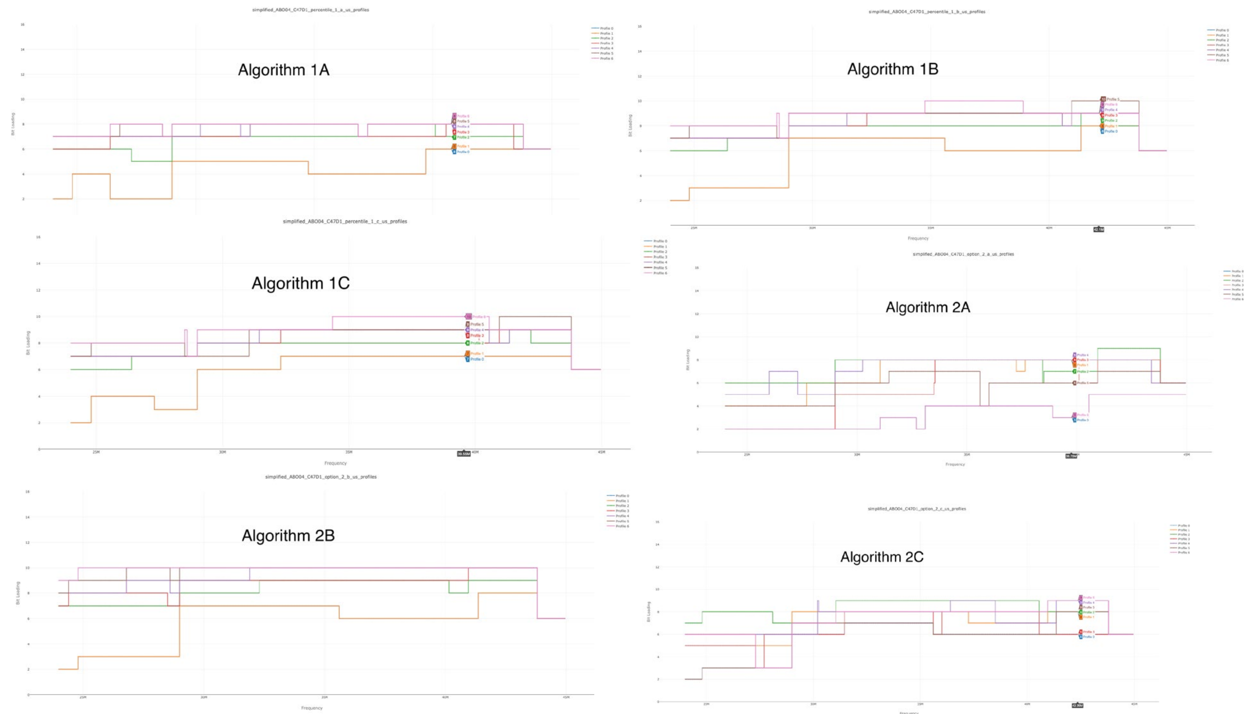

Figure 13 – 7 IUCs created using all Algorithms

6. Conclusion / Future Work

OFDMA is a newer technology with a lot of promise for operators looking to increase their upstream bandwidth capacity. At this point, OFDMA is a bit of a long walk to the finish line, because of novelty, the immature implementations of the technology (which is quickly being fixed) and the use of new-noisy spectrum which hasn't been used before. There will be a cyclic process of cleaning up the upstream plant, followed by degradation and new issues seen in the plant followed by more plant work and upgrades.

As it stands today, more flexibility in the CMTS OFDMA system functionality is needed, as this will enhance the upstream stability and ultimately the customer experience. There are many CMTS limitations, the most important of them being the limitation around just 2 IUCs per channel. The number of IUCs supported per channel needs to increase to accommodate the upstream variations. Improvements in the ability to define more than the limited (currently 4) modulation orders per profile would help design more granular profiles and improve the stability of the modems. There is an immediate need to measure and use the US RxMER to design and then select IUCs, and the use of FEC errors to downgrade IUCs. A fuller implementation of Upstream PNM functionality as defined in the DOCSIS 3.1 specifications will go a long way to ensuring a smooth transition from SC-QAM upstreams to OFDMA upstreams.

Lower frequencies in the upstream can be a very noisy and can make the channels using that portion of the spectrum unstable. DOCSIS 3.1 brought a high degree of adaptability with the new OFDMA technology. CMTSs can switch between upstream profiles (IUCs) to cope with the variability of the plant but the operator needs to create and design the IUCs optimally. So, a Profile management Application (PMA) is even more necessary on the upstream. The work done here shows some encouraging results for different upstream PMA algorithms and further testing is ongoing.

In terms of future work, we plan to complete the analysis on all the profile creation techniques and share the top methods with industry. Also, when more IUCs per channel are available, it would be interesting to see if which methods can be used to create better profiles which can increase capacity and keep the upstream operation robust. We also plan to experiment with pilot patterns and other OFDMA parameters like the number of symbols per frame, FFT size, Cyclic prefix, choosing exclusion bands etc. and see how best to optimize the OFDMA channel operation.

Abbreviations

CW	Codeword
CMTS	Cable modem termination system
CM	Cable modem
CER	Codeword error rate
DS	Downstream
DCCF	DOCSIS Common Collection Framework
HFC	hybrid fiber-coax
Hz	hertz
FEC	forward error correction
FFT	Fast Fourier transform
IDFT	Inverse discrete Fourier transform
IUC	Interval Usage Code
KPI	Key Performance Indicator
SG	Service group
SID	Service Identifier in the DOCSIS upstream
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
PMA	Profile Management Application
PNM	Proactive Network Maintenance
RxMER	Receive Modulation Error ratio
US	Upstream
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

[MULPIv3.1] DOCSIS MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I20-200407, April 7, 2020, Cable Television Laboratories, Inc.

[PHYv3.1] DOCSIS Physical Layer Specification, CM-SP-PHYv3.1-I17-190917, September 17, 2019, Cable Television Laboratories, Inc.

[CCAPOSSIv3.1] DOCSIS Converged Cable Access Platform Operations Support System Interface Specification, CM-SPCCAP-OSSIv3.1-I18-200610, June 10, 2020, Cable Television Laboratories, Inc.

DOCSIS 3.1 Profile Management Application and Algorithms, 2016 Spring Technical Forum Proceedings, Greg White and Karthik Sundaresan (CableLabs)

Practical Lessons from D3.1 Deployments and a Profile Management Application (PMA), 2019 SCTE Expo.

