Deploying PMA-Enabled OFDMA in Mid-Split and High-Split

A Technical Paper prepared for SCTE by

Maher Harb
Director, Data Science
Comcast
maher_harb@comcast.com

Dan Rice
VP, Access Architecture and Technology
Comcast
Daniel_rice4@comcast.com

Kevin Dugan
Data Scientist
Comcast
kevin_dugan2@comcast.com

Jude Ferreira
Principal Data Scientist
Comcast
jude_ferreira@comcast.com

Robert Lund
Principal Network Engineer
Comcast
robert_lund@comcast.com

Ramya Narayanaswamy
Director, Data Science
Comcast
ramya_narayanaswamy@comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

1. Introduction

Comcast has deployed DOCSIS® 3.1 (D3.1) orthogonal frequency division multiple access (OFDMA) in the mid-split (MS) and high-split (HS) bands of the upstream (US) spectrum on our virtual cable modem termination system (vCMTS) platform. D3.1 OFDMA allows higher modulation levels with up to 2x efficiency increases, with 2048-QAM today and anticipating 4096-QAM in the near future, compared to single carrier-quadrature amplitude modulation (SC-QAM) of 64-QAM. More importantly, D3.1 US technology allows configuring the modulation per 400 KHz mini-slot, enabling adaptation to the ingress and distortions discovered in the network. Several ingress and impairment sources that have the potential to degrade capacity and customer experience have been identified; examples include off-air very high frequency (VHF) broadcast, linear distortions, in-home analog TV modulators, frequency modulation (FM) radio, NOAA weather radios, and pre-equalization stability. As we deployed OFDMA initially on the network we quickly realized that in order to take full advantage of the capabilities OFDMA has to offer, most notably delivery of hundreds of Mbps US product speeds, a profile management application (PMA) system is required similar in concept to the one deployed downstream for managing orthogonal frequency division multiplexing (OFDM) profiles, which we reported on in previous SCTE contributions [1-4]. This paper presents the initial results from adapting the OFDM PMA system for OFDMA and testing in both lab systems and production field deployments. We share some examples of spectrum impairments, characterization methods, our thinking around how profiles are constructed based on device-level MER measurements, and the validation of the profiles against the vCMTS internal PMA function. We also comment on some of the intricacies of expanding upstream spectrum that warrant revisiting the core algorithm in the future.

2. The Upstream Split Change Evolution

DOCSIS-based broadband access has historically been operated with lower upstream capacity capabilities due to the amount of spectrum available in traditional splits of the 42 MHz low-split (also known as sub-split), 85 MHz MS, 204 MHz HS, as well as the upcoming full-duplex (FDX) with adjustable upstream and downstream spectrum usage, including simultaneous spectrum overlap. Today many hybrid-fiber-coax (HFC) networks still operate with a low-split configuration as shown in Figure 1 and are limited to upstream speeds under 100 Mbps.

As operators migrate to an 85 MHz upstream MS spectrum band as shown in Figure 2, they can offer increased product speeds, for example 300 Mbps, and with a 204 MHz HS spectrum band, 1 Gbps upstream products are possible. While updating the network to one of these upstream split scenarios, downstream spectrum may also be expanded to include up to 1 or 1.2 GHz of spectrum. This spectrum update allows additional downstream data capacity. Managing this spectrum is the topic of another Comcast 2022 Technical Forum Paper focused on delivery of multi-Gbps DS services [5]. Ensuring the quality and speeds of this new US spectrum requires using D3.1 OFDMA technology, which is designed to adapt to network conditions to maximize capacity and robustness based on the same platform described in prior SCTE papers on the Comcast PMA solution [1-4].

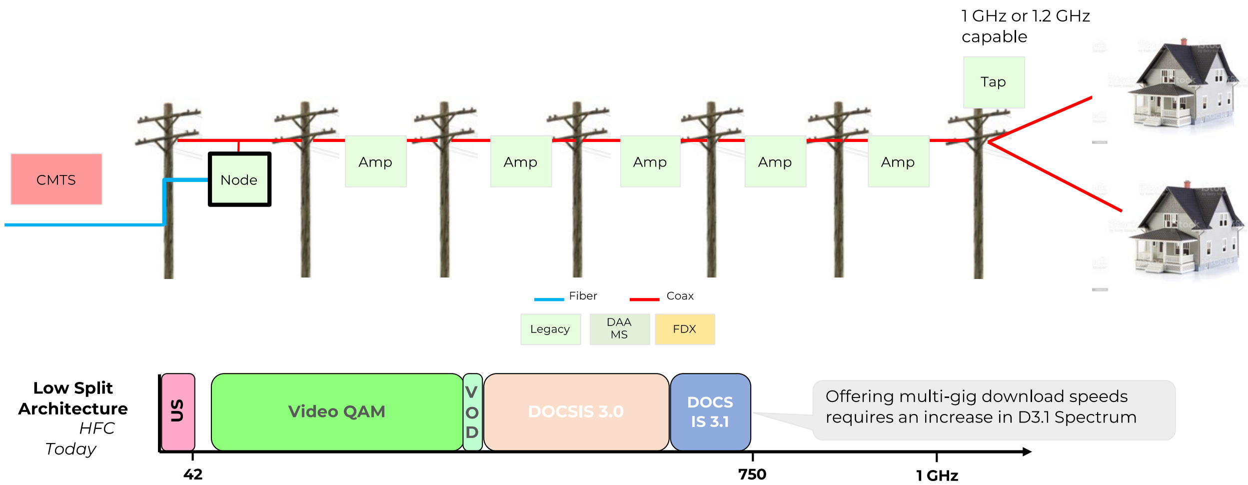

Figure 1. Yesterday's low-split (sub-split) network utilizes mostly 750 MHz analog fiber coax network, with typically Node +6 Amps or less.

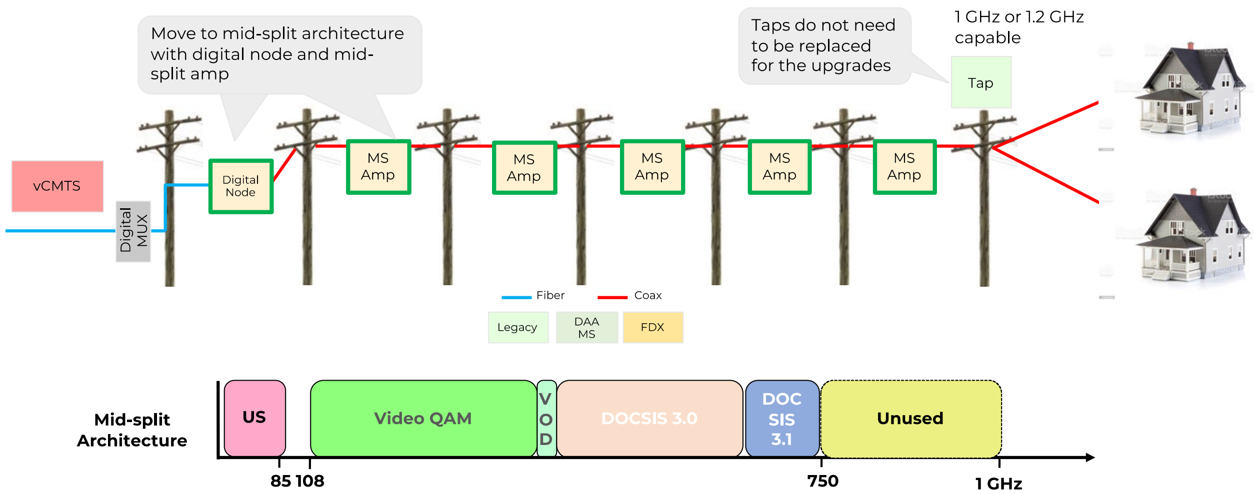

Figure 2. Mid-Split expanded upstream using D3.1, increasing capacity by >4x and enabling speeds up to 300 Mbps; overall network spectrum is extended to 1 GHz enabling multi-Gbps downstream speeds.

3. OFDMA Capacity Analytics

Deploying OFDMA in the mid-split spectrum increases capacity >4x compared to low-split pre-OFDMA US channel plans. During the COVID outbreak and the increase in bandwidth growth, particularly in the upstream direction, Comcast deployed six upstream channels and a PMA solution for D3.0 upstream channels to manage performance and rapidly add new capacity [2-3]. When deploying OFDMA spectrum on the nodes with high utilization and reducing the number of D3.0 channels, the resulting capacity needed to be equal to or greater bonded capacity than the six D3.0 channels in addition to enabling higher upstream product speeds. As a result, Comcast developed capacity analytics to understand in detail how the evolution to OFDMA would impact network performance from a capacity perspective.

It is instructive first to understand how this new capacity is provisioned to enable hundreds of Mbps product speeds. As shown in Figure 3 the channel plan was migrated from 6 US D3.0 channels to 4 D3.0 channels plus a single OFDMA channel. The blue line in the example node shows that the capacity of the 4 D3.0 channels is about 85 Mbps. This is less D3.0 capacity than the black line in the chart, which represents 100 to 120 Mbps of capacity available with the 6 US D3.0 channel lineup. As a result, Comcast needed to ensure that when we added the OFDMA capacity, enough traffic could be shifted into the OFDMA channel such that the net result would be equal or greater capacity available to the cable modem (CM) population on average. When the vCMTS bonds the D3.0 and D3.1 capacity for the D3.1 CMs in the service group (SG), it was possible – as depicted by the green line – to exceed 500 Mbps with the right PMA solution and network quality.

Depending on how the traffic is allocated across the channels, we are introducing a new metric known as "effective capacity" or "aggregate_speed" (yellow line in Fig. 3), which is the traffic-weighted capacity of the service group. This is a function of the amount of traffic enabled on the D3.1 OFDMA spectrum vs. the D3.0 spectrum along with the total capacity in the D3.0 and D3.1 channels based on the profile management solution optimization of the channels. Both the traffic allocation and the capacity of the channels change over time, based respectively on bandwidth demand and network quality. If the yellow line drops below the black line the goal of exceeding the D3.0 capacity when adding OFDMA is not met for that node and unit of time. This chart is one example node with 24% of the D3.1 modems able to access the OFDMA channel.

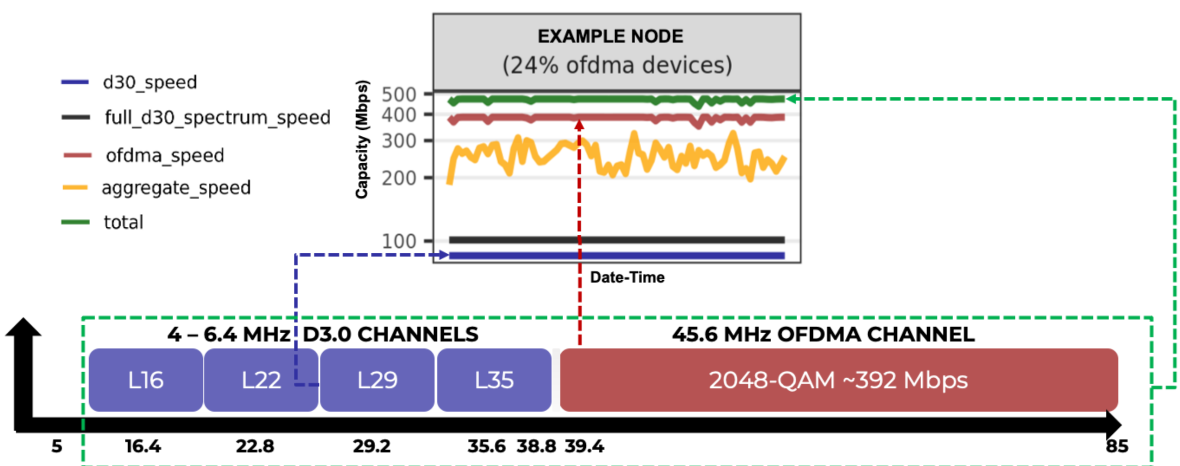

Figure 3. DOCSIS upstream capacities for the mid-split scenario. Total effective capacity is a function of PMA for OFDMA based on signal and noise quality, D3.1 penetration, and traffic, and consumption statistics.

The data of Figure 3 can be aggregated to the network level across a set of remote PHY device (RPD) Distributed Access Architecture (DAA) nodes. This aggregation for hundreds of nodes is illustrated in Figure 4 with the following key points:

- >1000 RPDs are mid-split activated (panel 1).
- ~20% of devices are OFDMA active with the potential to increase by > 10% as our in-home test solution [6-7] migrates each home to use of the OFDMA channel (panel 2).
- ~17% of traffic goes through the OFDMA channel on average; this is sufficient to increase the net effective capacity above the capacity of the full spectrum D3.0 (panel 3).

- Current traffic-weighted effective capacity (yellow) is higher than projected throughput if D3.0 spectrum was configured with 5th and 6th channels (i.e., the legacy six channel configuration shown as blak line in panel 6).
- The consumption of individual customers on smaller nodes can influence the results such as a heavy consumer with a D3.1 vs. D3.0 CM.
- The abrupt increase in OFDM capacity (red line) represents the switch from 32-QAM to 256-QAM. As PMA is added and higher order modulations are included, as detailed in this paper, the net effective capacity will further increase (panel 6).
- 200 Mbps speed test passes > 98% of time with only 256 QAM.

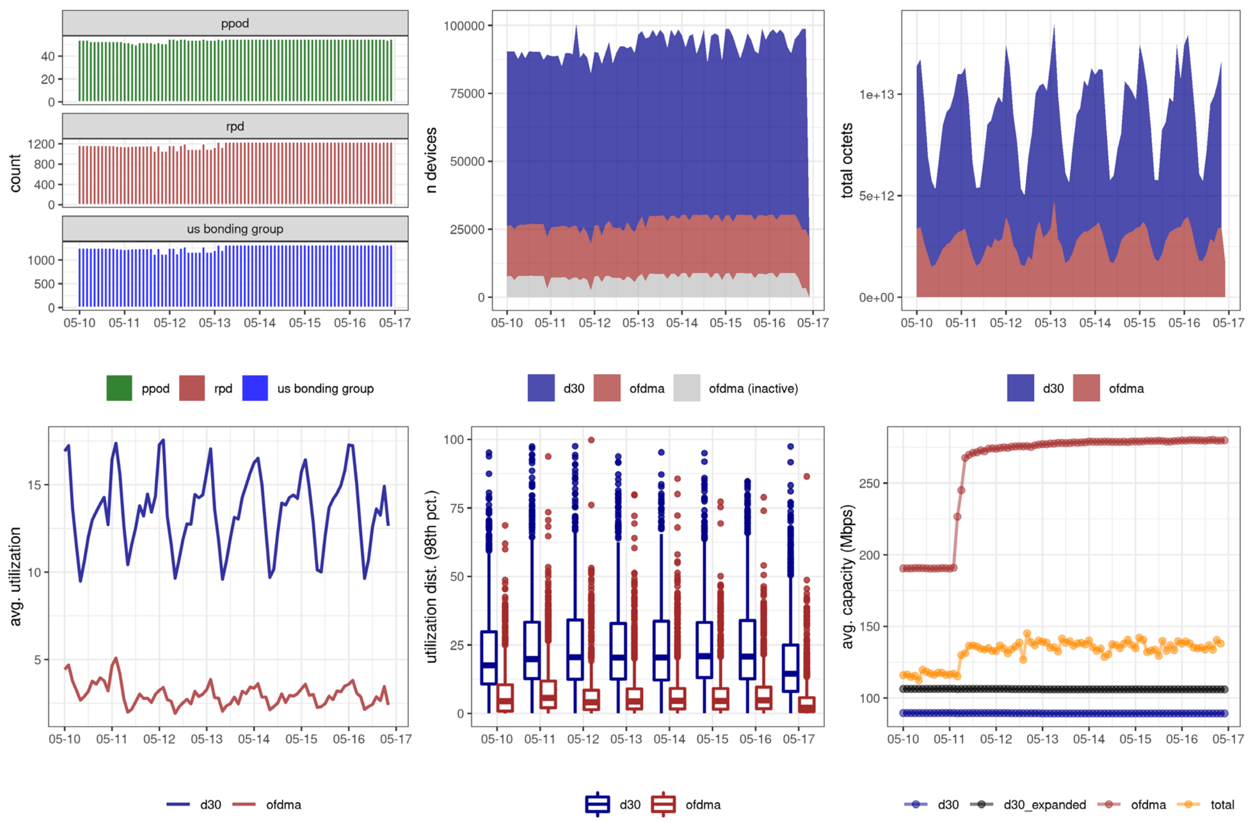

Figure 4. Critical mid-split set of performance and capacity statistics.

While aggregating across all RPD nodes, the net effective capacity is positive at only 256-QAM flat modulation configured in the OFDMA channel. There may be specific nodes that fall below the break-even line due to challenges in CM's connection to the OFDMA channel, or due to the low penetration of

D3.1 modems along with the utilization of the D3.0 vs. the D3.1 spectrum.

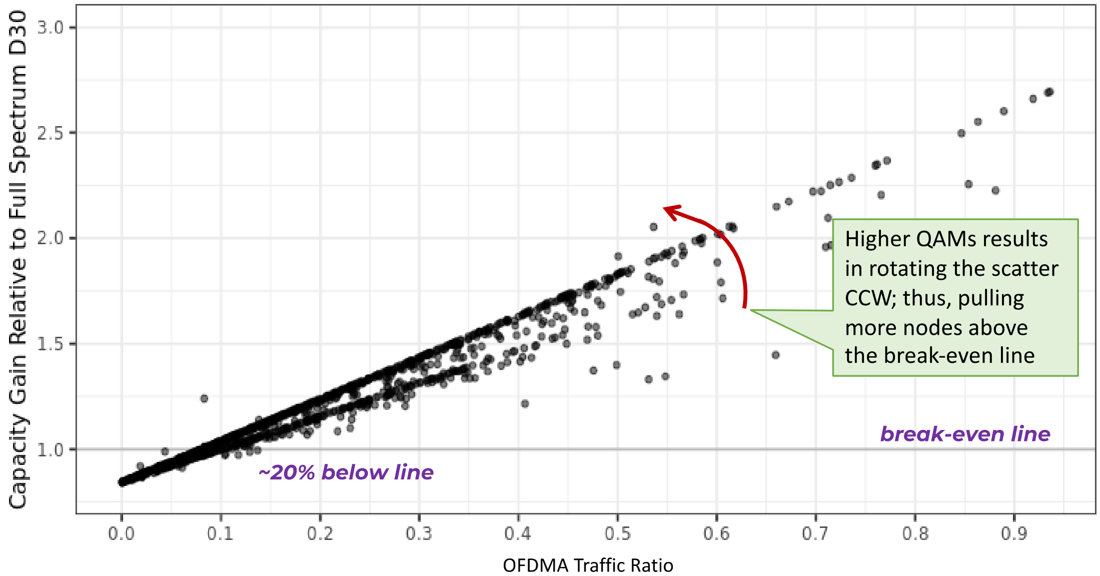

Figure 5 illustrates how at the above OFDMA configuration, and without migrating the traffic reliably to OFDMA, the effective capacity could fall below the break-even line for D3.0 spectrum. Fortunately, for ~20k nodes in this analysis, only 1 is projected to have D3.0 utilization go above alert level without gaining benefit from OFDMA due to low D3.1 penetration as shown in Figure 6. Increasing the OFDMA capacity with the use of PMA, as illustrated in Figure 5 moves the modems all above the break-even line, as will the organic increase of the D3.1 relative CM population.

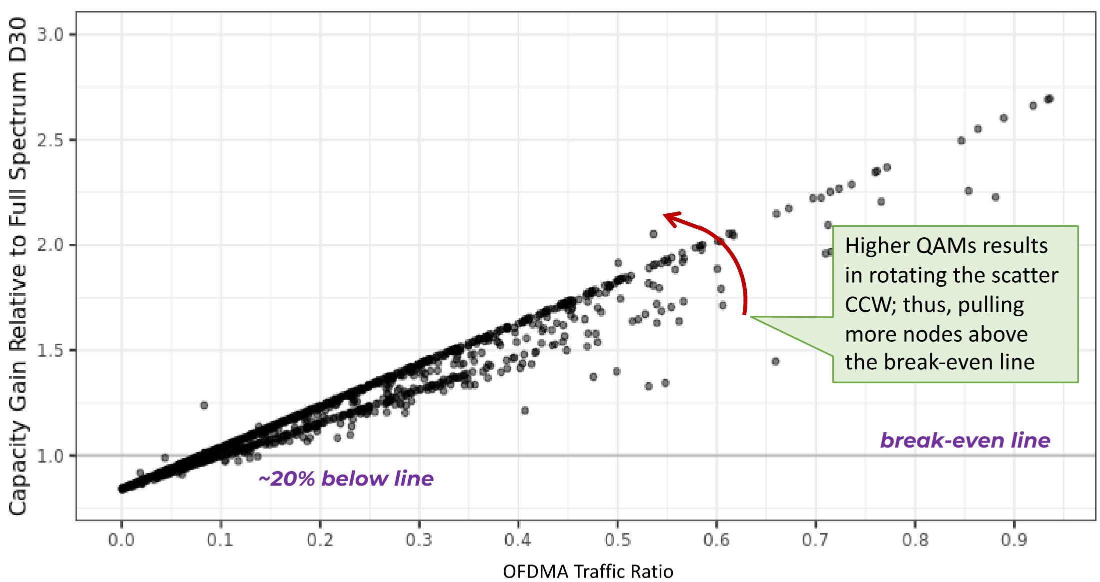

Figure 5. Effective capacity correlated to OFDMA-to-D3.0 SC-QAM traffic ratio. Enabling PMA is expected to lift most nodes above the break-even line.

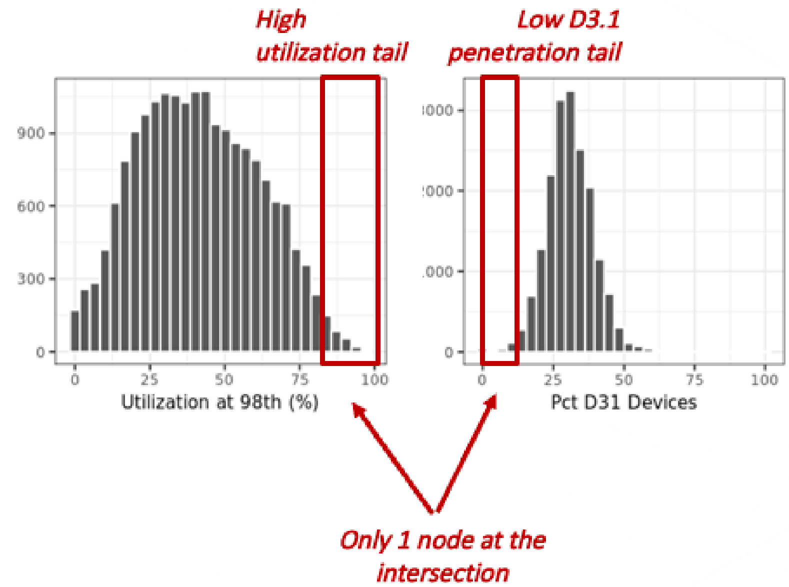

Figure 6. Utilization and D3.1 cable modem penetration for ~20,000 to-be-converted analog nodes. The intersection of the 2 tails of the distributions reveals that only 1 node will be driven into alert level due to loss of the 5th and 6th SC-QAM channels to accommodate OFDMA in the mid-split.

4. Mid-Split Deployment Data

OFDMA deployment in mid-split spectrum results in increased capacity and higher upstream speeds for D3.1 customers. It is critical to track and manage consumption and traffic from a capacity analytics standpoint to ensure that there are no capacity constraints driven by change in customer behavior and that spectrum utilization is within the normal operational range. There are additional factors like operator-initiated speed tests to check OFDMA connectivity and available speeds that could be adding additional consumption that need to be accounted for in this analysis.

In our tests, speed increase analysis was based on capturing the 98th percentile traffic and utilization for OFDMA and SC-QAM interfaces at a digital node and PPOD level for the test group (OFDMA-enabled and speed tier increases) and the control group (OFDMA-enabled and no speed tier increases) over a period of 3 months. For calculating traffic (consumption), the speed test-driven traffic was subtracted from the total per subscriber traffic. Pre and post-spead tier upgrade traffic was compared to assess the impact of the speed tier increase between test and control group.

As noted in

Figure 7. SC-QAM and OFDMA 98th utilization and traffic (Mbps).The vertical line in all panels marks the timestamp of the speed increases. The top row tracks the traffic on OFDMA (left) and SC-QAM channels (right) for test group (blue lines) and control groups (green and red lines). The bottom row tracks the utilization on OFDMA (left) and SC-QAM channels (right) for test group (blue lines) and control groups (green and red lines).

analysis of OFDMA utilization and traffic post-speed increase in one deployment area, there were no notable increases in both utilization and traffic for the test group as noted by the blue trend line.

Additionally, SC-QAM traffic and utilization dropped for the test group (as noted by blue line) due to customers utilizing available OFDMA capacity.

Analysis of test and control groups showed that consumption per device post-speed increase saw a marginal increase of up to ~20%, as noted in

Figure 8. Per device (CM) total consumption for the 3 groups pre and post speed increase.

. Additional speed increases within the mid-split spectrum can absorb customers' increased consumption; overall traffic (SC-QAM and OFDMA) utilizations are within the current utilization trend.

With additional capacity gains due to OFDMA enablement, speed increases are not resulting in increased utilization.

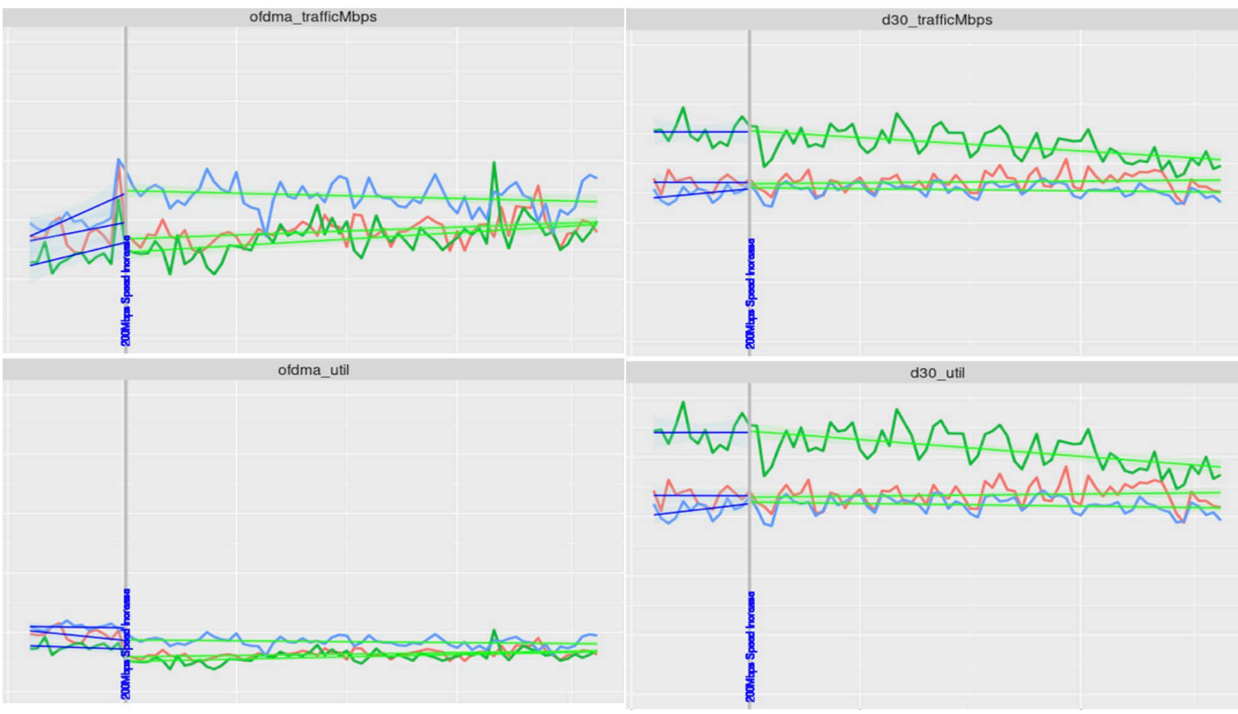

Figure 7. SC-QAM and OFDMA 98th utilization and traffic (Mbps).The vertical line in all panels marks the timestamp of the speed increases. The top row tracks the traffic on OFDMA (left) and SC-QAM channels (right) for test group (blue lines) and control groups (green and red lines). The bottom row tracks the utilization on OFDMA (left) and SC-QAM channels (right) for test group (blue lines) and control groups (green and red lines).

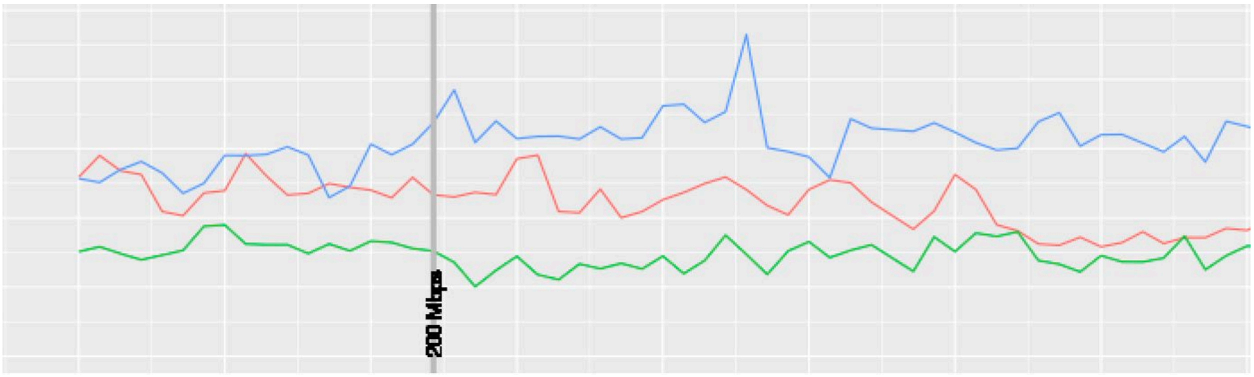

Figure 8. Per device (CM) total consumption for the 3 groups pre and post speed increase.

5. PMA for OFDMA

5.1. Why PMA for OFDMA

In addition to needing the capacity for bandwidth growth, and for higher product speeds, OFDMA managed through a PMA solution can significantly improve customer experiences, as we have discovered with OFDM in the downstream and D3.0 US profile management. During initial OFDMA deployments several challenging network impairments and distortions limited the capacity of the OFDMA channel without a PMA solution. In some cases, the ingress sources coupling into the network were so severe that modems were challenged even to connect to OFDMA and a more robust modulation was required for interval usage code (IUC) 13, which CMs use to try and perform initial ranging on the OFDMA channel. Several examples are shown here along with some analytics of different impairments discussed later in section 6 (mid-split ingress).

Figure 9 illustrates ingress from a very high-powered VHF transmitter that was impacting a variety of nodes. As can be seen, the lower VHF channels 2 and 6 are both coupled into the upstream at very high levels through loose connectors or damaged cables. During this initial deployment, a 256-QAM flat modulation profile was tested on the OFDMA channel. At the receiver, the spectrum analyzer showed noise at a level that would have required a flat 32-QAM, or even more robust configuration, to allow the modems to range on the OFDMA channel. As a result of this impressive ingress, Comcast did some of the network analytics detailed in section 6.1 (VHF TV Ingress) to identify the probabilities of this type of ingress impacting the wider network. Delivering the higher speed services cannot be done in the presence of this type of ingress without an OFDMA PMA solution.

Figure 9: Lower VHF ingress into mid-split channel would have required 32-QAM or lower modulation without PMA solution.

Figure 10 illustrates a second example of why PMA is required to achieve Gbps symmetrical speeds. This node had loose connections or cracks in the cable near three ingress sources:

1) An upper VHF spectrum transmitter with channels 9 & 10;
2) A NOAA weather radio transmitter, that although narrow, needed to be properly managed by PMA; and
3) FM radio ingress between 88 &108 MHz.

The lower channel in the high-split node is stopped at 85 MHz primarily to enable mid-split capable CMs on this node to achieve the hundreds of Mbps speeds as these CMs cannot bond to a channel above their capabilities while the HS capable modems can bond to all of the spectrum. While not intended or preferred for capacity reasons, this does avoid the FM radio band which can impact HS node performance. Technology development is happening currently to allow a mid-split modem to use an OFDMA channel that extends above 85 MHz while only being scheduled in minislots below 85 MHz with appropriate upstream channel descriptors.

If a channel is deployed across the FM spectrum, it is clear that a PMA solution would be required while remediating the network. The right chart in Figure 11 shows the PMA response to this ingress. The modulation levels of the OFDMA profile are tailored around the spectrum and shown as yellow segments. The modulation in the minislot containing the NOAA radio needed to be modulated at QPSK to avoid using error correction capabilities to resolve it. Note that most of the spectrum is modulated at 2048-QAM with the future option of 4096-QAM (when the feature is available on the vCMTS platform). Without a PMA solution to manage the capacity, this node would not have been capable of delivering 1 Gbps upstream speeds.

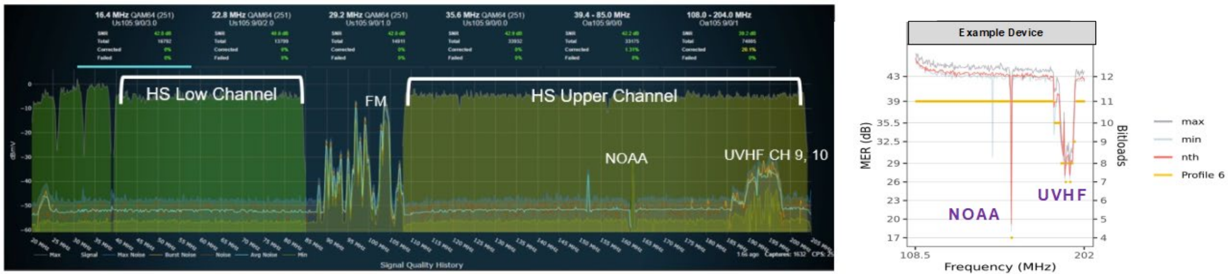

Figure 10: (right panel) Example high-split spectrum capture showing upper VHF and NOAA radio ingress. Mid-split channel has excellent signal quality. (left panel) PMA constructed profile for the high-split channel.

5.2. PMA Virtual Network Functions Development Starts and Ends in the Ingress Lab

In addition to the analytics of field data, a critical component of developing a PMA solution for OFDMA is the ability to recreate ingress from the field within the lab as a test vector for the PMA algorithm. In the example shown in Figure 12 the VHF ingress in Denver is measured at a Comcast lab, injected into a software defined radio (SDR) platform, and played back into the lab node. The OFDMA IUC or modulation profile design of the algorithm is then evaluated to verify that it resolved the packet loss and automatically configured on the lab vCMTS. This SDR platform is also capable of converting MER per subcarrier or per minislot into a signal that simulates the ingress in the field that would have caused that MER profile. The OFDMA is run periodically in the continuous integration and continuous deployment (CICD) environment to ensure that all the test vectors are well managed as vCMTS, CM, or PMA software is evolved.

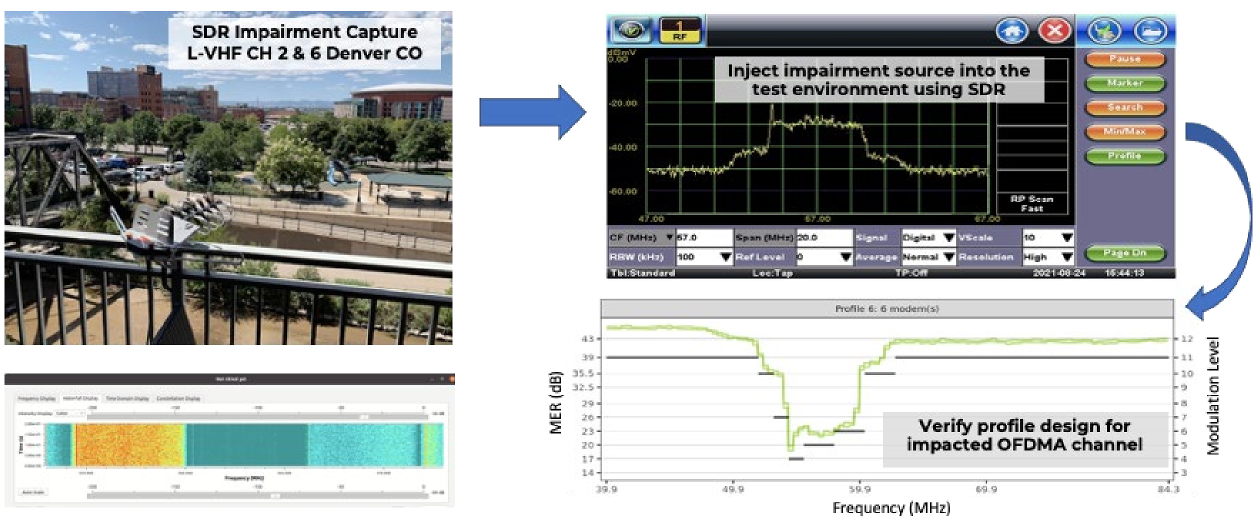

Figure 11. Lab testing flow. VHF Ingress is captured (left panel) and injected into the node using SDR platform (top right panel). PMA adjusts to the ingress by constructing the proper modulation profile (bottom right panel).

5.3. OFDMA PMA Comes Together

The PMA base platform has been described in several previous papers [2-5]. Figure 13 describes the base platform including 3 primary components: the configuration manager which automates the instantiation of new IUC/profile configurations; the analytics engine which recommends changes to configuration; and the data engine that collects the real time streaming OFMDA telemetry and makes it available for the PMA solution along with other proactive network maintenance (PNM) tools used to direct fix agents that remediate the network. Building on this same existing platform for OFDM has dramatically accelerated the time-to-market for the OFDMA solution.

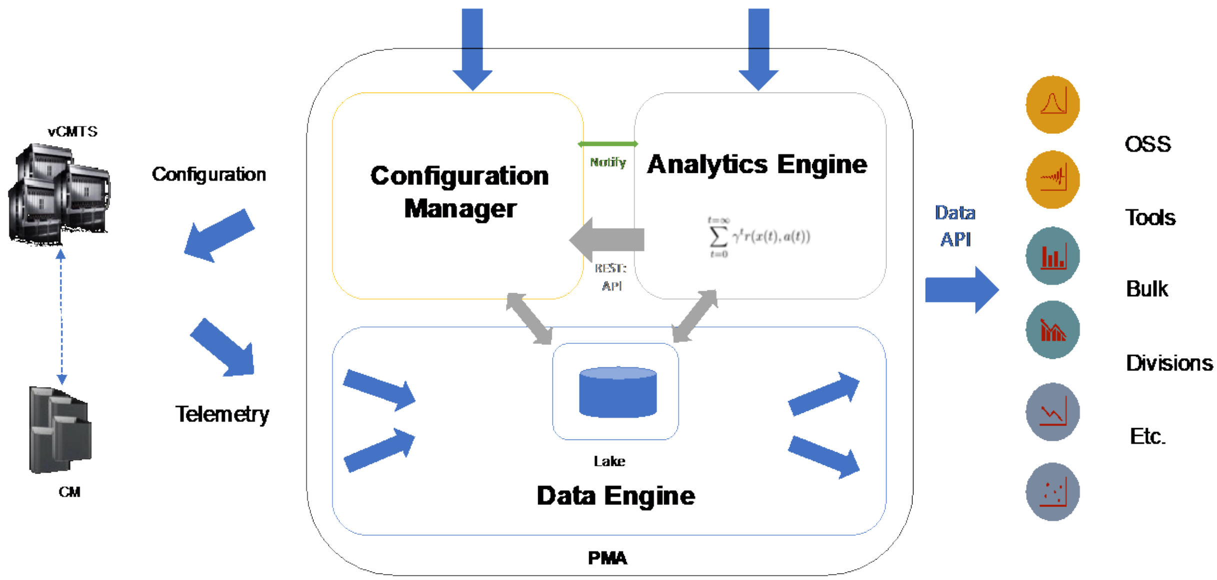

Figure 12. PMA architecture supports DS and US D3.1 and US D3.0 profile management.

Several examples of OFDMA support the platform descriptions in the previous papers are referenced in this paper. Figure 14 shows an example of a fairly consistent MER signature for a node in which most of the spectrum could support 2048-QAM except for a small region. Each modem MER per minislot is aggregated over several days of 5-min samples with statistics such as the 10th percentile calculated as representative of the minislot-aggregated MER value. In this example, a cable modem, labeled as “Device 9” was having some attenuation issues and could not operate with this IUC configuration, so a second configuration (row 2, column 1) was configured by PMA, along with a third to cover “Device 4” that was experiencing another drop in MER (row 1, column 4).

The PMA analytics engine compiles this data for each cable modem, clusters them into similar performance groups, and then sends a set of recommended configurations to the configuration manager. The configuration manager then schedules and automatically implements the new configuration along with pre- and post-checks and change management tickets (auto-open & close) with transactional integrity.

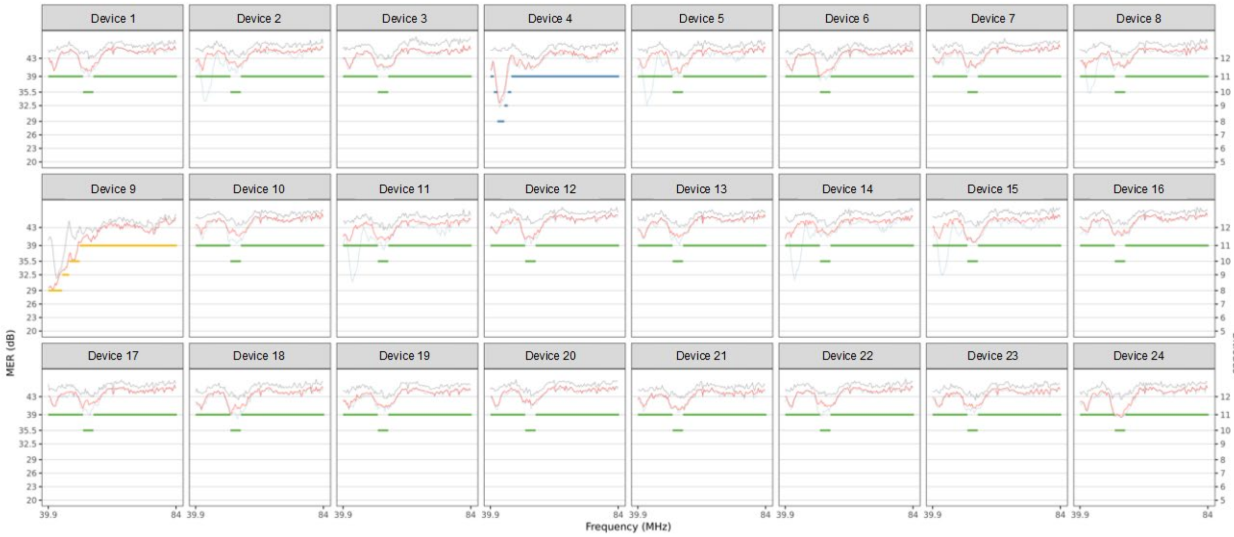

Figure 13. Example PMA IUC/profile design across mid-split spectrum that improves the channel capacity compared to adopting a "flat modulation" IUC design.

Figure 15 shows a HS channel with a single CM from one of the first HS trials. 7 IUC profiles were constructed by PMA, but only 1 was required for the HS modem. This is the same case shown at a different time for the spectrum analysis shown in Figure 10. This is a recommendation that was updated as the level of the ingress varied over time while the channel was being dynamically managed by PMA.

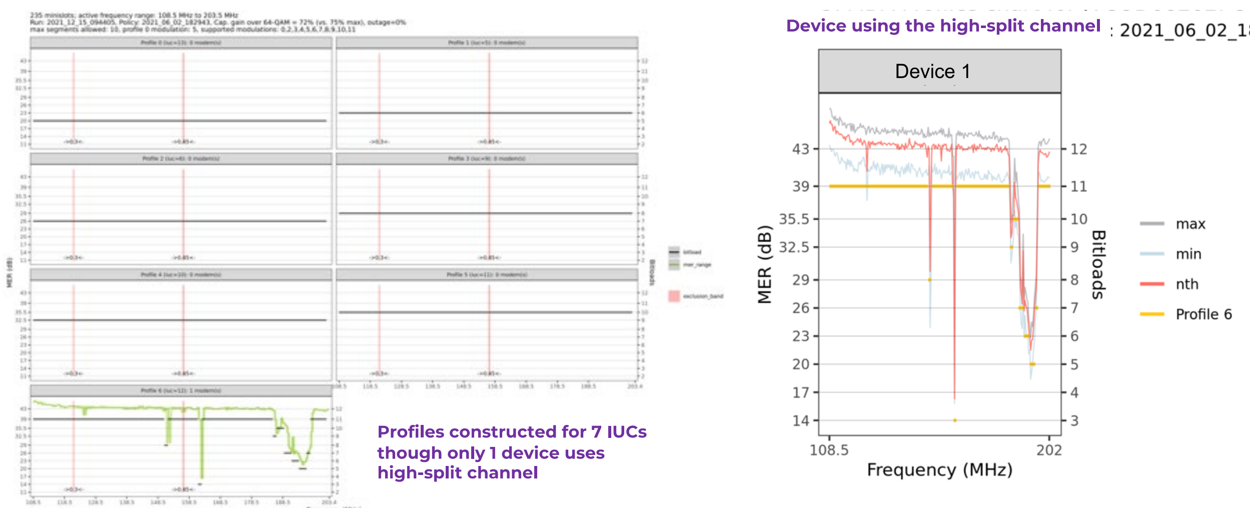

Figure 14. Example PMA design for the high-split channel in Figure 10. Left panel shows the 7 constructed profiles. Since only 1 cable modem is bonded to the channel, 6 profiles were configured with flat modulation. The right panel shows the profile tailored to the noise detected on the channel.

Figures 16 and 17 are dashboard type charts that track various metrics of interest relating to PMA performance. These include device counts, capacity, and traffic metrics both at the channel level (first row) and the IUC level (second row), as well as normalized traffic by IUC and MER statistics (last row). The first example shown in Figure 16 is for a field node with relatively clean spectrum. For this node, traffic is flowing mostly through IUC 12—the highest capacity profile at ~ 400 Mbps. MER statistics show levels that are consistently above 41 dB and the uncorrectable codeword rate is extremely low. As a result, profile switches by PMA are minimal (indicated by vertical solid lines).

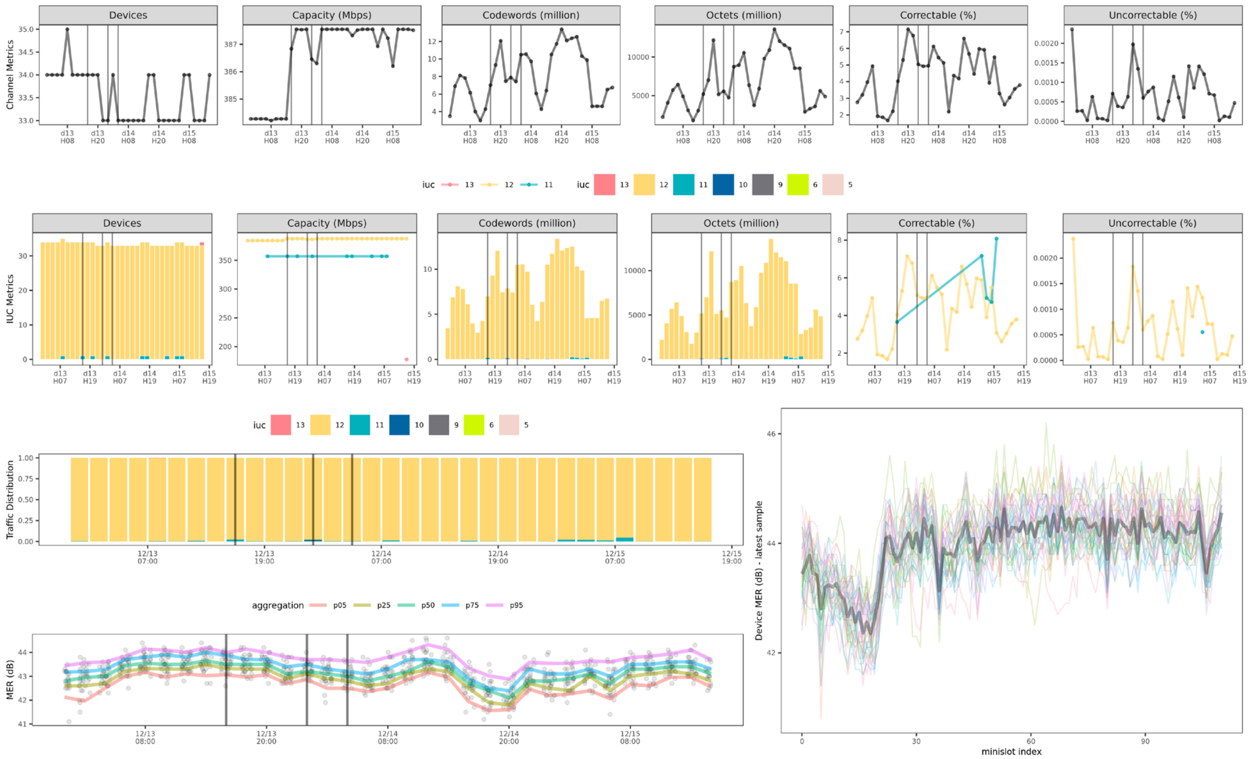

Figure 15. Example MS PMA performance charts for a clean spectrum.

In contrast, Figure 17 is for a node that exhibits variation in MER level within the 34-43 dB range. For this node, the internal PMA function responds by moving devices to suitable profiles. In this example, the dip in MER correlates with the shift of traffic from IUC 12 to lower capacity profiles. The PMA system is also responding by reconstructing the profiles in response to the dynamic nature of the noise. While the uncorrectable codeword rate is higher than the clean spectrum node, it is kept under 0.2%.

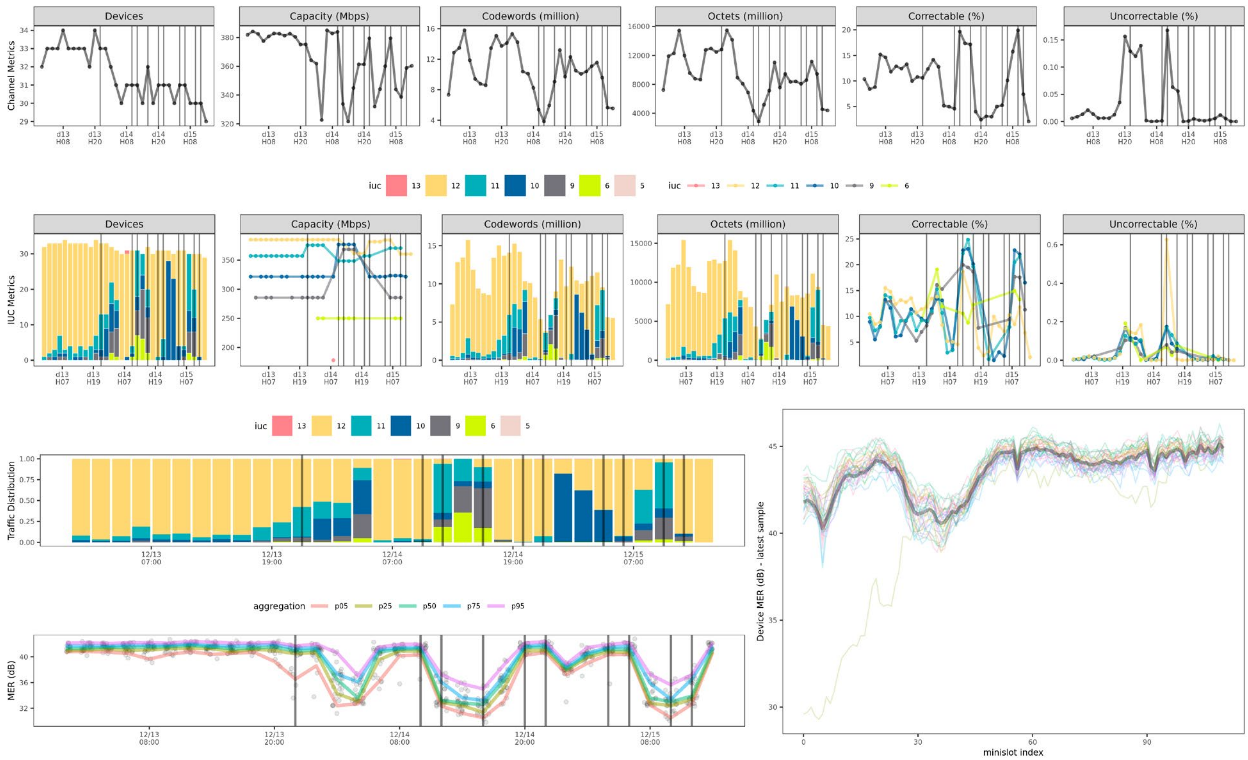

Figure 16. Example MS PMA performance charts for a noisy spectrum.

5.4. State of the Upstream Spectrum

While many of these ingress sources are problematic, the network is generally in great shape and easily supports the capacity and product speed expectations for the MS and OFDMA. Figure 18 describes the network from one of our very first OFDMA trials across 100 nodes in challenging RF environments. Even in these nodes it can be seen that:

- ~50% of network supports 4096-QAM US OFDMA based on DOCSIS PHY spec MER requirement.
- The IUC selection in the vCMTS is actually 3 dB better, implying that ~85% of OFDMA traffic can flow on 2048-QAM. This is consistent with lab testing that shows the OFDMA receiver performs at least 3 dB better than the CableLabs DOCSIS 3.1 PhHY specification requirements in the MER range of interest.
- The forecast is for ~70% of OFDMA traffic to use 4096-QAM based on this early data when that feature is available from the vCMTS.
- This performance exceeds our initial capacity models for both mid-split and high-split that were created to set goals for new product speeds.

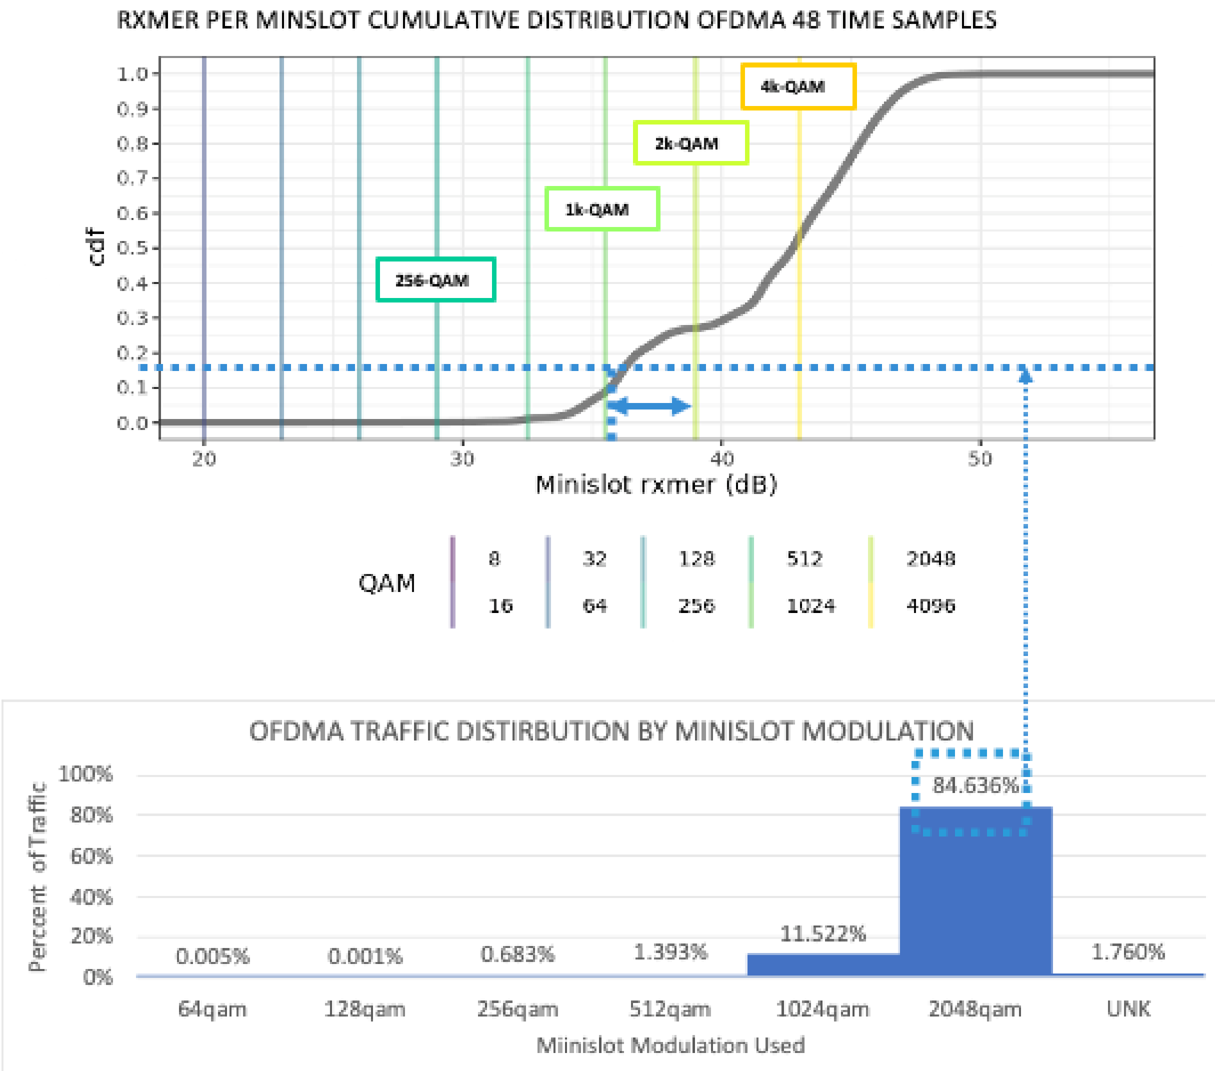

Figure 17. (Top panel) MER cumulative distribution from early trials on ~100 nodes shows that ~85% of mini-slots can support 2k-QAM. (Bottom panel) Actual traffic stats confirm this picture.

While Figure 17 was from an initial small trial, the network performance trend is improving now that we are delivering OFDMA to hundreds of thousands of CMs with OFDMA activated. Figure 18 shows the distribution of MER per minislot over 2 days of time across tens of thousands of modems early in the writing of this paper. Several key points to note include:

- MER Performance statistics are very promising; against specification thresholds, ~90% of mini-slots show 4096-QAM speeds.
- Lab and field testing of OFDMA PMA indicates we have 3 dB more additional margin relative to specifican thresholds.
- Even at specification thresholds, this distribution represents an average of ~520 Mbps US capacity per MS SG.
- Work is ongoing to tune the OFDMA PMA, once deployment challenges are stabilized.

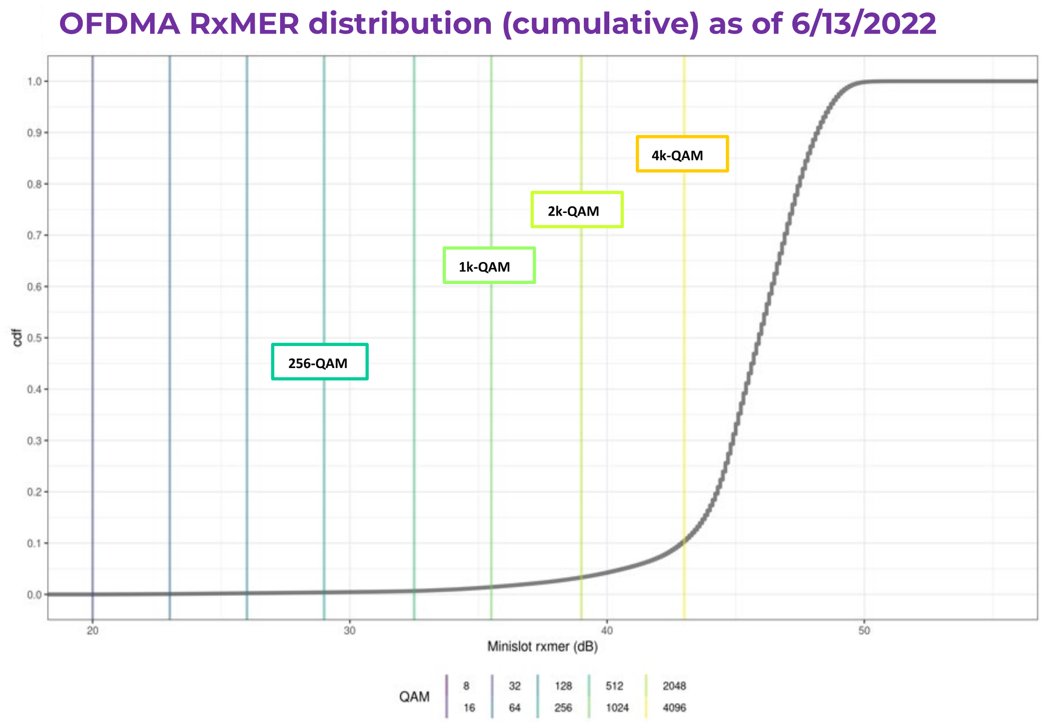

Figure 18. MER per minislot aggregated across tens of thousands of modems at 1-hour samples over 48 hours across varied network locations.

6. Mid-Split Ingress

6.1. VHF TV Ingress

VHF over the air (OTA) ingress is one of the more common ingress sources we discovered in our initial OFDMA deployments. As shown above this ingress can be problematic without a PMA solution. To understand how it might affect capacity, Comcast did an analysis of power levels expected at homes and within nodes. The analysis was based on a Federal Communications Commission (FCC) spectrum database for VHF transmitters and a propagation model informed by measuring actual ingress using our full band capture function. The bands evaluated are shown in Figure 19.

Channel Number	Frequency in MHz
Channel-2	54-60
Channel-3	60-66
Channel-4	66-72
Channel-5	76-82
Channel-6	82-88
FM Broadcast	88-108

Figure 19. One primary source of ingress which impacts mid-split is the lower VHF OTA television broadcast. The channels of interest are 2-6 and FM broadcast

A sample of VHF transmitters was selected and the overlapping channel power for devices at various distances from these transmitters was measured. Figure 21 shows the variation in channel power as a function of distance from the VHF transmitters. Key points to note include:

- Impact reduces as the distance from the VHF transmitter increases and it appears to follow a power trend line.
- The impact appears to vary across device types. A potential explanation for this variation may be due to shielding improvements in newer devices.
- Device Type 3 filters channels below 108 MHz and hence is minimally impacted.

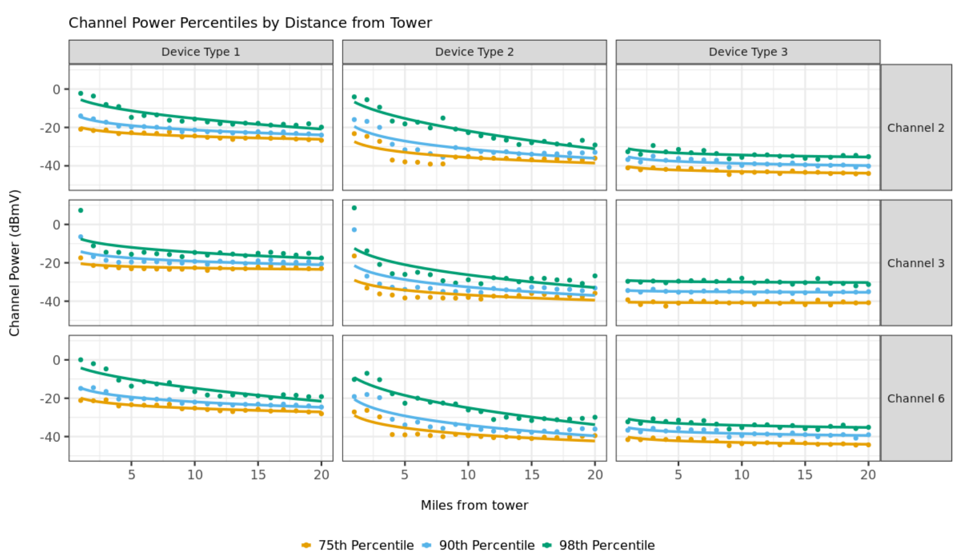

Figure 20. Variation in channel power vs distance for VHF transmitters for different device categories

6.2. Interesting Impairments to Keep in Mind for PMA

The impairment in Figure 22 was seen in several nodes and diagnosed to see if it was distortions from DS signals. After investigating we believe these narrow-band ingressors are an old VCR, video game (i.e., Atari) or the wrong connector on an older set-top-box connected to an outlet in the home. This is the signature and power levels that would be generated from a channel 3 amplitude modulation (AM) modulator at 61.25 MHz.

Figure 21. Example video modulator in home transmitting into OFDMA spectrum

While stabilizing the OFDMA technology deployments we have observed several cases of high tilt compensation across the OFDMA channel. Some of these responses affected MER of the OFDMA channel resulting in some compensation by the PMA solution. Examples of these are shown in Figure 23. Some of these were due to network devices that were not replaced in the mid-split upgrade process, such as inline equalizers in the lower half of Figure 23 that were tuned for a low-split network. Other cases were from equalization optimization issues requiring a periodic unequalized probe to reset the equalizer and interaction with upstream transmit powers. Other cases were a result of incorrect use of upstream conditioning in new mid-split amplifiers.

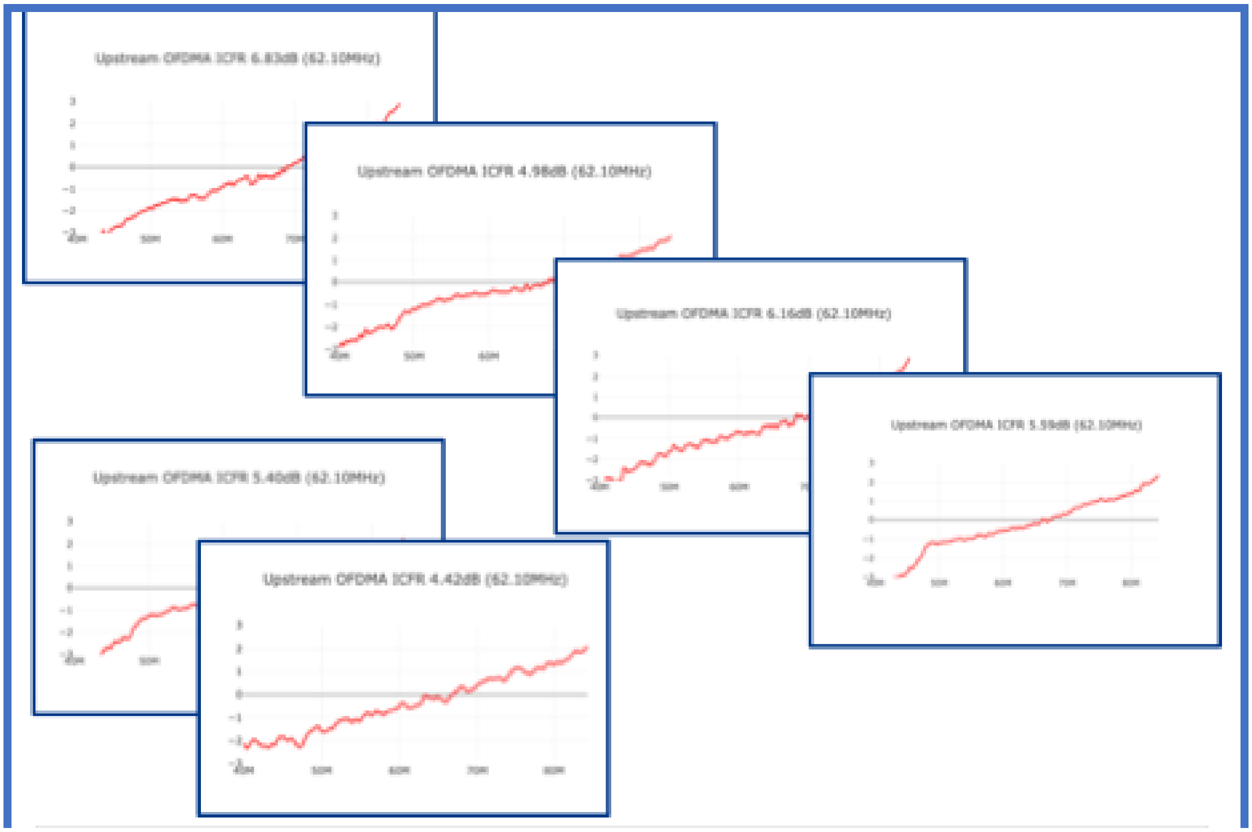

Figure 22. Various Pre-EQ responses for OFDMA channel showing titlt.

7. Conclusion

The paper details some insights into the challenges of deploying OFDMA technology. As part of the deployments and trials over the last 18 months we have shown the benefits and success of the OFDMA technology for delivering hundreds of Mbps of US speeds in MS networks and 1 Gbps US speeds in high-split networks with a positive customer experience. Comcast has determined that an OFDMA PMA solution is essential in delivering these new higher speed products. The PMA solution is essential because of the ingress and distortion challenges we have identified in a very small percentage of nodes that changes over time. Primary ingress sources include over-the-air broadcast transmission. These ingress sources are easily handled by a dynamic modulation profile management application that is an incremental evolution to the platform.

Abbreviations

AM	amplitude modulation
CM	cable modem
CMTS	cable modem termination system
D3.0	Data Over Cable Service Interface Specification 3.0
D3.1	Data Over Cable Service Interface Specification 3.1

DOCSIS	Data Over Cable Service Interface Specification
DS	downstream
FCC	federal communications commission
FDX	full duplex
FM	frequency modulation
HFC	hybrid fiber coaxial
HS	high-split
IUC	interval usage code
MER	modulation error ratio
MS	mid-split
NOAA	National Oceanic and Atmospheric Administration
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
OTA	over the air
PMA	profile management application
PNM	proactive network maintenance
PPOD	physical pod
RF	radio frequency
SCTE	The Society of Cable Telecommunications Engineers
vCMTS	virtual cable modem termination system
US	upstream
UVHF	Upper VHF band, Low VHF can be used to refer to channels 2-6, Upper VHF for channels above channel 6
VHF	very high frequency

Bibliography & References

1. "A Machine Learning Pipeline for D31 Profile Management", Maher Harb, Jude Ferreira, Dan Rice, Bryan Santangelo, and Rich Spanbauer, NCTA Technical Paper, 2019.
2. "Full Scale Deployment of PMA", Maher Harb, Jude Ferreira, Dan Rice, Bryan Santangelo, NCTA Technical Paper, 2020.
3. "Optimizing DOCSIS 3.0 Configuration in the Upstream through Applied Reinforcement Learning", Kevin Dugan, Maher Harb, Dan Rice, Rob Lund, NCTA Technical Paper, 2021.
4. "PMA Improvements - Strategies employed for faster mitigation, increased capacity, and cost savings", Leech et al, NCTA Technical Paper, 2022.
5. "Time is Ripe for D30 Farming: Achieving Optimal Spectral Efficiency by Allocating D30 Spectrum to OFDM", Maher Harb, Chad Humble, Sebnem Ozer, Dan Rice, NCTA Technical Paper, 2022.
6. "Rapid and Automated Production Scale Activation of Expanded Upstream Bandwidth", Rob Thompson, Rob Howald, John Chrostowski, Dan Rice, Amarildo Vieira, Rohini Vugumudi, Zhen Lu, NCTA Technical Paper, 2021.
7. "Bringing the Mid-Split Factory Online to Rapidly Produce Terabytes", Howald et al, NCTA Technical Paper, 2022.

