Full Scale Deployment of PMA

Lessons Learned from Deploying the Profile Management Application System at Scale and Considerations for Expanding the System Beyond OFDM

A Technical Paper prepared for SCTE•ISBE by

Maher Harb
Director, Data Science
Comcast
1800 Arch Street, Philadelphia, PA 19103
267.260.1846
maher_harb@comcast.com

Bryan Santangelo
Executive Director, Data Eng and Science
Comcast
1800 Arch Street, Philadelphia, PA 19103
918.640.8936
bryan_santangelo@comcast.com

Dan Rice
Vice President, HFC Architecture
Comcast
1401 Wynkoop St Ste 300, Denver, CO 80202
720.512.3730
daniel_rice4@comcast.com

Jude Ferreira
Principal Engineer
Comcast
1800 Arch Street, Philadelphia, PA 19103
215.286.4070
jude_ferreira@comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

Introduction

In 2019, Comcast developed a Profile Management Application (PMA) system for generating and transacting D3.1 downstream (DS) profiles tailored to the conditions of each Orthogonal Frequency Division Multiplexed (OFDM) channel in its network. The approach, machine learning algorithms and system architecture were described in a previous SCTE technical paper [1]. The initial plan for this follow-up paper was to focus on Comcast's PMA deployment journey, the success of which is evidenced by thousands of Cable Modem Termination Systems (CMTSs) managed by the PMA, yielding greater than 20 Tbps of added downstream (DS) capacity to the network.

With the onset of the COVID-19 crisis, some of that focus shifted, in lockstep with the shift of the U.S. and worldwide workforce from office to home. Figure 1 shows the 32% increase in upstream (US) traffic, post-COVID, and the shift in peak times for DS traffic from 9:00 PM to 7:30 PM, and from 9:00 PM to 8:00 AM & 6:00 PM for US traffic. Figure 2 shows the bandwidth demand growth, around time of the COVID crisis (Spring of 2020), for US traffic (black curve) and DS traffic (sky blue curve). With work-from-home traffic increasing on the network, and because of the earlier implementation of the PMA, the DS capacity was available, and the network was able to easily scale to the significantly increased demand.

The US is a different story. As a fraction of the total available spectrum, and even as it is being industrially widened from sub-split to high-split configurations, the fact remains that US capacity is a more difficult challenge. Commencing with shelter-at-home requirements, US traffic grew sharply, seemingly overnight. Comcast has publicly shared data on the increases in traffic scale since COVID started [2-4], along with transparency about the level of investment and technological attention that prepared us for "Black Swan" scenarios like a pandemic. This enabled more effective management of the additional traffic growth delivered over the Data Over Cable Service Interface Specification (DOCSIS) broadband network [5]. As this paper will ultimately show, by adding an upstream PMA focus to the existing PMA suite, we were able to boost upstream capacity by 36%, from 86 Mbps to 117 Mbps.

Given these extraordinary circumstances, with the COVID crisis in full swing, and with the shift in internet usage, we refocused this paper to share our accelerated efforts in developing and deploying PMA for the US using DOCSIS 3.0 technology. Fortunately, the technology that was brought to bear on the challenges of US capacity was already under development. The effort was shaped by the early concepts found in CableLabs member publications from the late '90s [6] into the early 2000s. Combined with state-of-the-art methods, including scaled cloud-based compute, and machine learning techniques such as Reinforcement Learning (RL), we were able to ensure system stability and optimize the network bandwidth as spectral conditions and demand changed.

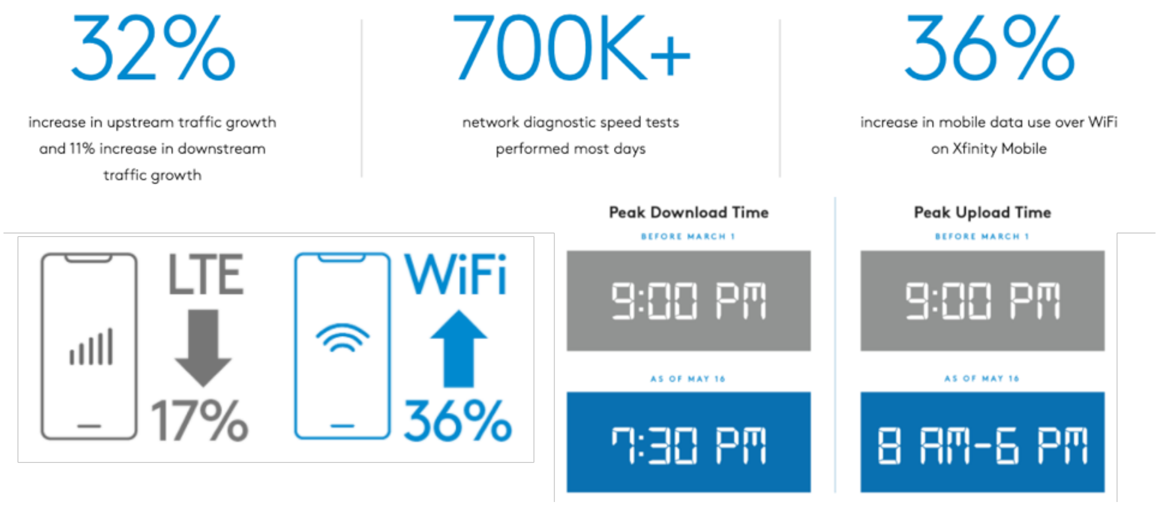

Figure 1 - COVID network impact on traffic peak times.

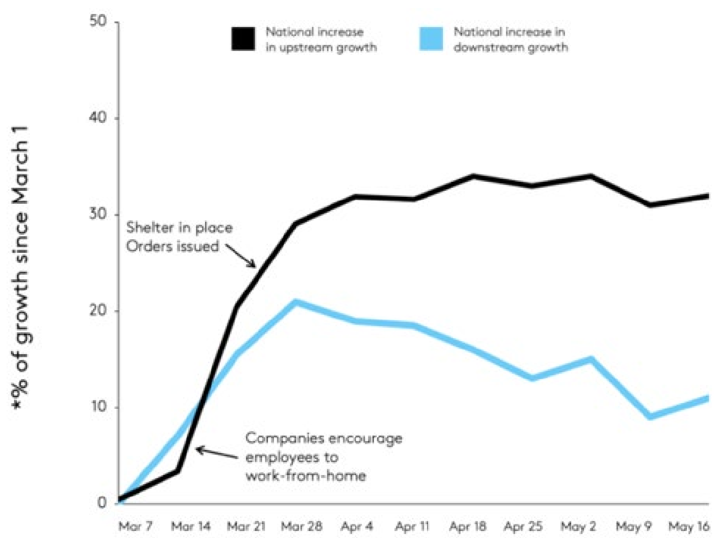

Figure 2 - COVID network impact on bandwidth demand.

Note that at the start of 2020, the US D3.0 profile management was under development as part of a holistic capacity management and long-range access network architecture plan, along with other efforts, such as passive hybrid fiber-coax (HFC) with deeper fiber, virtualized CMTS with remote PHY nodes, and spectrum augmentations, such as high-split and FDX. The D3.0 US PMA effort was initially targeted at optimizing the upstream spectrum, in conjunction with the deployment of additional US channels. The additional US channels were located in parts of the spectrum that would be difficult to manage without an autonomous modulation profile optimization system, given known and persistent levels of ingress. The D3.0 US profile management was also intended for those network segments without fiber deep and spectrum upgrades, to help offset the timing risks related to developing and deploying new technology by deferring or eliminating investment in legacy technologies.

The paper is organized as follows: Section 1 recaps our 2019 Expo paper by providing an overview the PMA system and its algorithms and describing the current state of the DS algorithm. Section 2 presents analyses and dashboards created to track system performance & health. Section 3 presents our vision for how the PMA system is expected to evolve in the future. Section 4 is dedicated to describing the US PMA system. Section 5 concludes with the lessons learned from the PMA deployment journey.

1. Overview of System and Algorithms

The PMA system architecture and DS algorithms have been described at length in the previous SCTE technical paper [1] A brief recap and description of the prior work is included here for professional context. The PMA system is composed of four separate components, shown in Figure 3: Data Collector, Data Storage, Analytics Engine, and Configuration Manager.

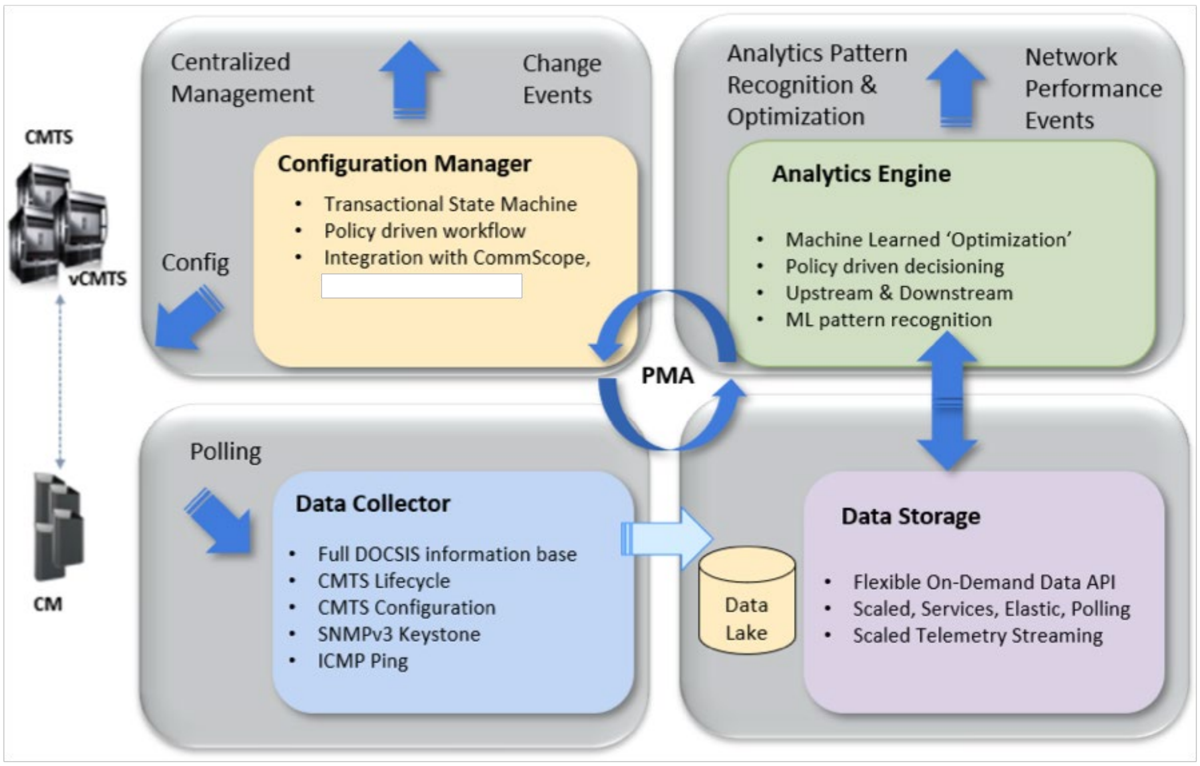

Figure 3. The PMA System Architecture.

The Data Collector is responsible for collecting telemetry data from CMTSs and gateway devices. The data is polled at different frequencies that range from every 5 min to hourly, and was designed to constitute a "comprehensive poller," enabling applications beyond the scope of PMA. From a PMA perspective, the data needed to support the construction of OFDM profiles falls into the following categories:

- **Network topology**: Establishes linkage between device, OFDM channel, and CMTS.
- **Configuration model**: Provides characteristics of the OFDM channel, e.g. number of subcarriers, subcarrier width, frequency range, position of exclusion bands, etc.
- **CMTS type**: Provides make, model, hardware & software versions of a given CMTS.

- **Telemetry**: Retrieves Modulation Error Ratio (MER), Forward Error Correction (FEC), signal, and traffic measurements from devices, and channel utilization measurements from CMTSs. This category constitutes the largest bulk of the data, given that MER spectra are measured at a per-device OFDM subcarrier resolution, with 4096 data points for each MER sample for each device.

The Data Storage is primarily comprised of a public cloud-based data lake, where the polled data listed above land in raw (unprocessed) form.

The Analytics Engine (AE) is a machine learning pipeline that uses the data to construct OFDM profiles suitable for use by the devices in the network—given spectral conditions measured over certain time windows. At its core, constructing profiles is a type of optimization problem in which the stated objective is to maximize channel capacity and minimize codeword error rates, subject to certain constraints. Thus, the problem contains an inherent trade-off between improving robustness and increasing network capacity, since reducing error rates is achieved by opting for lower modulation levels, at the expense of reduced channel capacity.

The constraints are dictated by the CMTS hardware and software versions, as different CMTSs support different numbers of profiles per OFDM channel. Within the construct of a profile, they may also support different numbers of modulation exception zones (segments), as well as imposing additional constraints on the attributes of a segment (e.g. segment width).

Algorithmically, the AE uses hierarchical clustering—a type of unsupervised machine learning algorithm—to group together devices that share common noise characteristics and assign them a common modulation profile. Additional smoothing algorithms are applied post-clustering, to reshape the segments according to given constraints. In the current version of the algorithm, the clustering objective function is designed to maximize capacity around a statistical decision boundary.

FEC rates are considered, indirectly, by imposing additional constraints on the mapping from MER values to modulation levels (e.g. a MER value > 27 dB supports 256-QAM at maximum). As an example, the plot in Figure 4 shows MER measurements alongside the constructed profiles on a dual *y*-axis plot. Since spectral conditions vary over time, multiple MER samples are captured over a time window dictated by AE policy. For each panel (device) we show 3 curves characteristic of the variation in MER: the max level (dark gray curve), the min level (light gray curve), and the 10th percentile (red curve). Also, per policy, it is the 10th percentile that is fed to the algorithm as conservatively representative of the device's MER state. The constructed profiles are overlaid on the plots and follow the scale of the right *y*-axis. In this specific example, the CMTS allows 4 profiles per OFDM channel, 4 segments per profile, and a segment width that is a multiple of 1 MHz. Profiles 1-3 are overlaid in yellow, blue, and green colors, respectively on the devices that are assigned to each of the 3 profiles. Profile 0 (not shown) is the control profile and is set to a flat 256-QAM by AE policy. Note that the impairments shown are generated in the lab and applied to select devices. Because of the CMTS-imposed limitation of 4 exception zones (segments), the algorithm overcompensates for the V-shaped impairment exhibited in the MER spectra of device #5.

Lastly, the Configuration Manager (CM) is responsible for transacting profiles generated by the AE. The output from the AE defines profiles according to a standardized intermediate JSON format that is agnostic to the CMTS make and model. The CM converts the output to commands that are specific to the CMTS. The CM is also responsible for validating the profiles, deciding on whether to reject or accept the AE recommendations, scheduling the transacting of the profiles according to a policy that defines allowed maintenance dates/times, and performing pre- to post-transaction checks to confirm that the configuration was successfully applied.

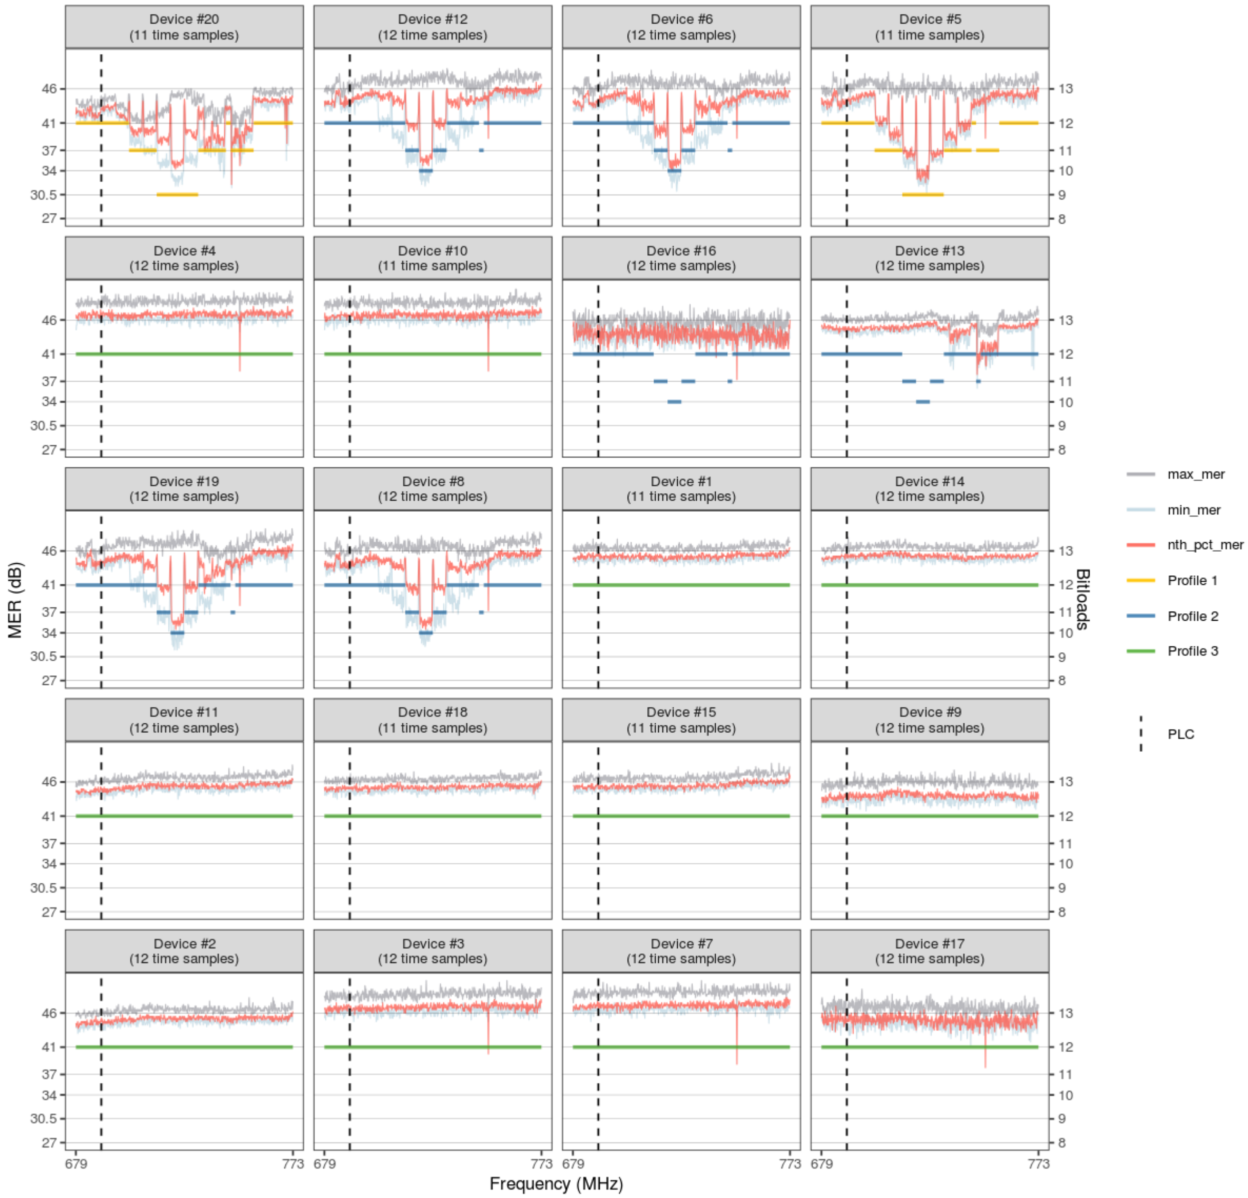

Figure 4. Example of PMA profiles constructed for a set of 20 devices.

2. Network Visibility

A prerequisite to deploying PMA was achieving a high level of visibility into the network, in order to respond to issues that may arise as deployment ramps up. Thus, this requirement opened a path to an independent effort on dashboarding and automated notifications. While the PMA's Data Collector is exhaustive in the breadth of data it collects, the focus of this effort was to provide insight and visibility into metrics that are PMA-related or perceived to be PMA-related. We start by introducing the basic design criteria for the PMA dashboards:

- The overall health of the PMA system is to be monitored through a common "home view" that constitutes an entry point to all other dashboards.

- One must be able to drill down from the "home view" to an "OFDM channel" detailed view.

- All dashboards must display near-real time information (limited by the Data Collector's polling frequencies).

- Dashboards must be informative in the operational sense, i.e. highlighting issues and allowing operations to take appropriate actions to remediate.

2.1. Home level view

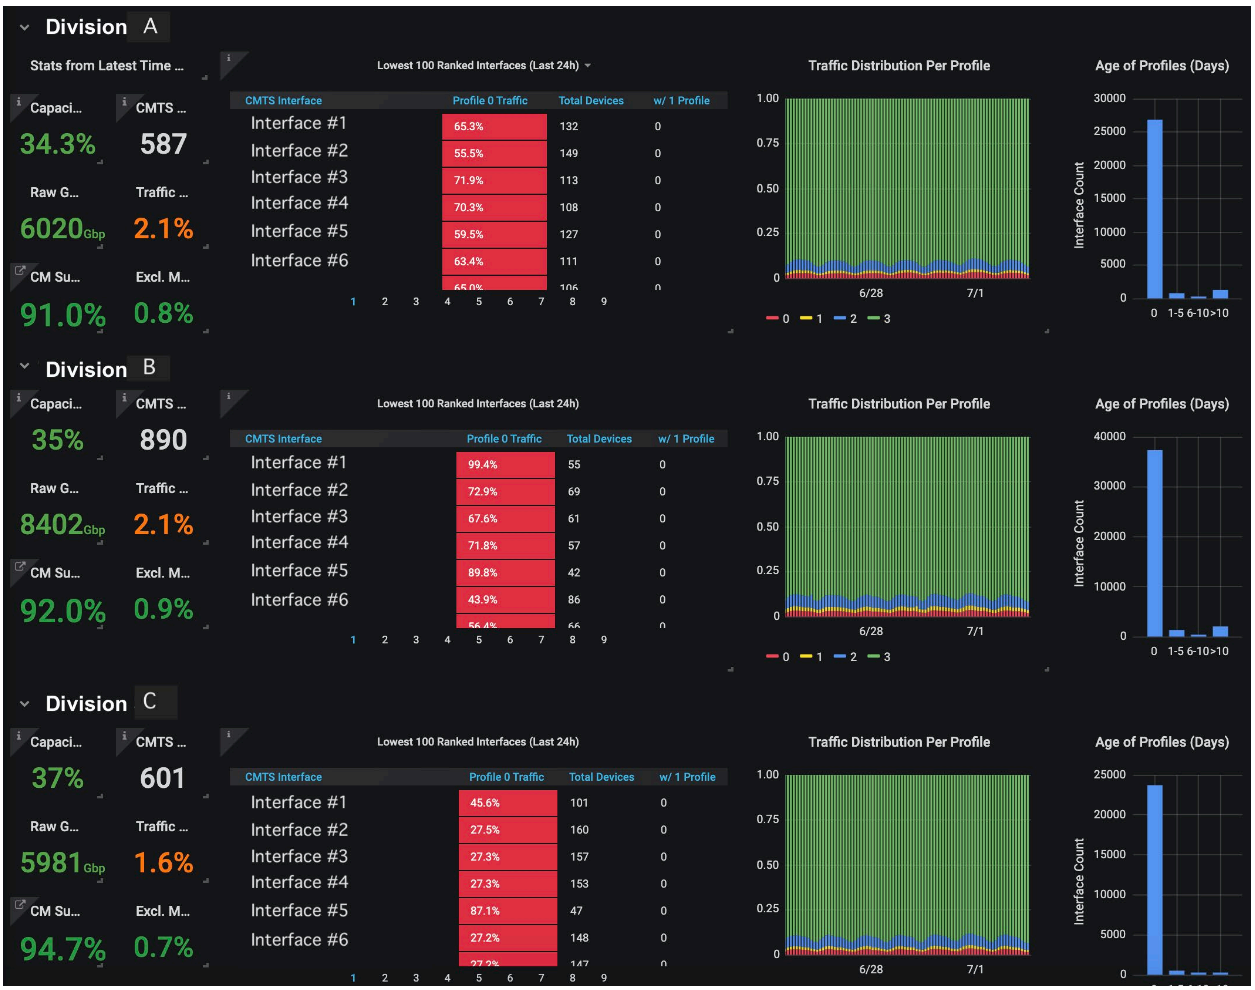

Figure 5. Screenshot of the DS PMA home view dashboard.

Given the design criteria, we next introduce some of the key metrics displayed on the home view dashboard shown in Figure 5. The home view aggregates all metrics by division—the highest-level organizational structure within Comcast's network. This view is meant as a first entry point to monitor the health of the PMA system. The "Lowest 100 Ranked Interfaces" table brings to attention interfaces (OFDM channels) that are low performing, from a PMA perspective and, thus require further investigation. Clicking on an interface name opens the interface-level dashboard with additional details.

Note that names of divisions and interfaces shown in the dashboard screen capture were anonymized. The metrics shown on the home view are:

- **Point-in-time Metrics**: Latest metrics on health and performance of the system that include:
    - **Capacity gain** (34.3% for Division A in Figure 5): The capacity gain attributable to PMA as measured relative to a 256-QAM baseline. Note that the capacity metric represents an instantaneous value calculated from the actual traffic distribution across different OFDM profiles as captured from telemetry in 100s of thousands of service groups across the network.
    - **Raw gain** (6020 Gbps for Division A in Figure 5): Same metric as above but calculated in units of absolute Gbps.
    - **Number of CMTSs** (587 for Division A in Figure 5): Number of CMTSs managed by the PMA system at the time of this snapshot.
    - **Traffic on Profile 0** (2.1% for Division A in Figure 5): Since Profile 0 is configured as 256-QAM or 64-QAM by policy, ideally no devices would use Profile 0. Thus, this metric is indicative of a combination of the health of the system and the freshness of the configured profiles. If spectral conditions degrade and data profiles are not updated in a timely manner, more traffic will flow on Profile 0.
    - **CM success rate** (91.0% for Division A in Figure 5): Percent of CMTSs that were successfully configured with updated profiles during the last configuration window, which is set by policy.
    - **Excluded modems** (0.8% for Division A in Figure 5): Percent of devices with severe impairments that were excluded from the clustering algorithm. The rationale for device exclusion is to avoid wasting a customized profile on severely impaired devices that require field work to effectively mitigate. The current criteria for excluding a device requires that the mean modulation that a device supports is less than the profile 0 capacity (256-QAM or 64-QAM).

- **Lowest 100 ranked interfaces**: The aggregate measures described above may conceal severe issues affecting only a small number of OFDM channels. Hence, this table ranks the 100 channels that require most attention. The ranking is based on a combination of metrics that include traffic on Profile 0, number of devices experiencing OFDM partial service issues, and total number of devices on the channel. In addition, this table provides an entry to detailed interface-level dashboard by clicking on the channel name.

- **Traffic distribution**: This is a time series showing the distribution of traffic across profiles aggregated for all OFDM channels in the topology context. The capture shows that >90% of traffic flows through Profile 3—the highest capacity profile. It also reveals some cyclical behavior with increased traffic on Profile 2 during peak usage times. Note that PMA configures profiles based on device clustering according to the MER characteristics. Even though the PMA assumes that each device should be assigned a specific profile, the ultimate assignment of profiles is left to the CMTS, according to its internal profile selection function as a policy. The fact that most of the traffic flows on Profile 3 indicates that the mapping between MER values and modulation levels is conservative. Section 3 discusses capturing additional capacity gain by adjusting the AE policy thresholds.

- **Age of profiles**: A distribution of profile age shows that most interfaces were updated within the last day (column labeled 0 days). This view allows monitoring profiles that may become stale. The current update frequency is daily. However, the AE will recommend profile updates only if there was a change in the spectral conditions that warrants updating the existing profile. Hence, it is expected that the age distribution will show profiles that are older than 0 days.

2.2. Interface level view

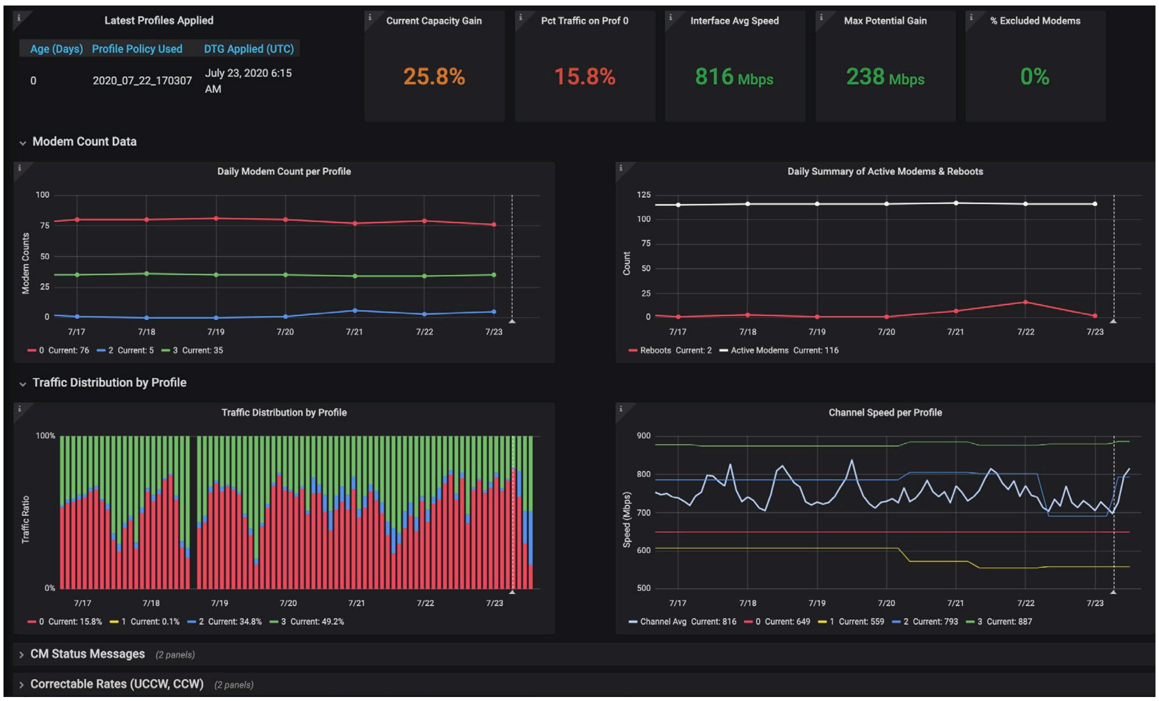

Figure 6. Interface level view dashboard for DS PMA.

The interface (OFDM channel) level dashboard offers a highly detailed view of the network dynamics occurring on that interface. A screenshot of the interface dashboard is shown in Figure 6. It shows point-in-time metrics as well as time series metrics covering various dimensions. This view shows device counts and reboots, traffic distribution, and channel speed as function of time. Other views including the CM status message statistics and error rates are "collapsed in" (not shown in screen capture) for conciseness. The view contains several sections covering the following areas:

- **Point-in-time metrics**: These are similar to the home-level view of point-in-time metrics but aggregated at the interface-level.

- **Modem counts data**: Counts of devices by profile, total device count and count of device reboots.

- **Traffic distribution**: Traffic distribution across profiles and the speed of each profile.

- **Modem CM-Status Messages**: Breakdown of D3.1 CM-status messages reported by devices related to the health of the OFDM channel.

- **FEC error rates**: Time history of the correctable and uncorrected codeword error rates.

- **Traffic volume data**: Time history of traffic volume (unicast octets) as reported both by CMTS and devices.

Note that similar dashboards were developed for the US PMA. Select views of these are shown in Section 4.

2.3. Notifications

Because networks tend to degrade over time from an impairment perspective, if not maintained, we developed methods to reduce the impact of the network on the customer experience. In addition to the operational views described above, impairment detection algorithms are provided as a core part of the architecture. Notifications from these are delivered across a messaging bus to other Comcast OSS tools, to ensure that technicians are dispatched to the right hubs, network segments and homes to remediate issues. The following data are provided as part of the notifications, to assist with event prioritization and triangulation:

- Interface details such as Physical Link Channel (PLC) location, start/end frequencies, total number of impacted cable modem counts.
- List of impacted cable modems with severity of impairments.
- Interface impairment ranking at the national, divisional, and regional levels.
- Reference to the API for historical data stored in the data lake related to the event.
- Reference for the API to enable fix agents to collect real-time on demand data, to confirm that the issue is still present, and to assist in isolation and to confirm mitigation.
- Other data sources to enrich the event, such as the mobile wireless carriers that overlap with the OFDM channel.

Notifications for severely impacted interfaces are also sent via email to divisional engineering and operations teams. Additional details, such as device/interface RxMER images, are posted to internal collaboration platforms. Note that another Comcast-authored SCTE paper describes the pattern detection associated with these notifications [7].

3. Next Generation Algorithms

The current PMA algorithm predicts a ~30% increase in network capacity, relative to a flat 256-QAM baseline, on average. Figure 5 shows that the capacity gain as measured from real-time traffic is around ~35%. This gap between predicted gain and realized gain is due to the internal profile selection function of the CMTS, based on CM-STATUS messages. The internal profile selection effect is clearly visible in the traffic distribution by profile shown in Figure 5, in which more than 90% of the traffic is observed to flow through Profile 3—the highest capacity profile. The logic for internal profile selection function varies across vendors, but the general idea is the same: the CMTS regularly sends test codewords to devices on all configured profiles and determines, based on the encountered errors, the highest capacity profile a device is capable of using. The CM responds with CM-STATUS messages descriptive of that specific device's perspective on the performance. These messages are further processed through proprietary selection algorithms with controls for the operator to optimize. The fact that most devices are able to use a higher capacity profile than recommended by the PMA algorithm indicates that the thresholds adopted for converting MER to modulation (shown in Table 1) are somewhat conservative. This added resiliency against errors is a known feature of the low density parity check (LDPC) error correction algorithm. While quantifying the LDPC benefit is not straightforward, prior studies suggest that it provides an additional 3-6 dB improvement over Reed Solomon—its D3.0 counterpart [8] Next, we discuss how to take advantage of the superior performance of the LDPC to capture additional capacity gain.

Table 1 - Minimum MER values that support the corresponding modulation from DOCSIS 3.1 specification.

MER Threshold (dB)	Modulation efficiency
0	0
9	2
15	4
21	6
24	7
27	8
30.5	9
34	10
37	11
41	12

The PMA pipeline offers several algorithmic tuning knobs that are configured by policy. One such knob is a global offset in dB applied to the thresholds listed in Table 1. We experimented with the knob by relaxing the threshold for 2 production CMTSs, while tracking key performance metrics against a control CMTS group within the same site. The results of the experiment shown in Figure 7 reveal that an additional 10% increase in capacity is garnered by relaxing the thresholds by 3 dB, without introducing a negative impact on the codeword error rates. In all panels, the control group is shown in red and the experiment group in light blue. The first 3 vertical dashed lines correspond to the points in time at which the thresholds were adjusted by 1, 2, and 3 dB respectively. For the 4th vertical dashed line, the adjustment was maintained at 3 dB (i.e. no change over previous state). The top left panel in Figure 7 shows the capacity of the experiment group increased per adjustment, with a total increase of ~10% at 3 dB. The top right panel shows the traffic on Profile 0 (256-QAM) shows similar levels across the 2 groups, indicating that shifting the thresholds is not causing data profiles to become unusable by devices. The bottom left panel shows the correctable error rates increased for the experiment group following the first adjustment by 1 dB. The bottom right panel shows the uncorrectable error rates reveal no trend related to the adjustments (notice that the uptick towards the end of timeline was experienced by both experiment and control groups).

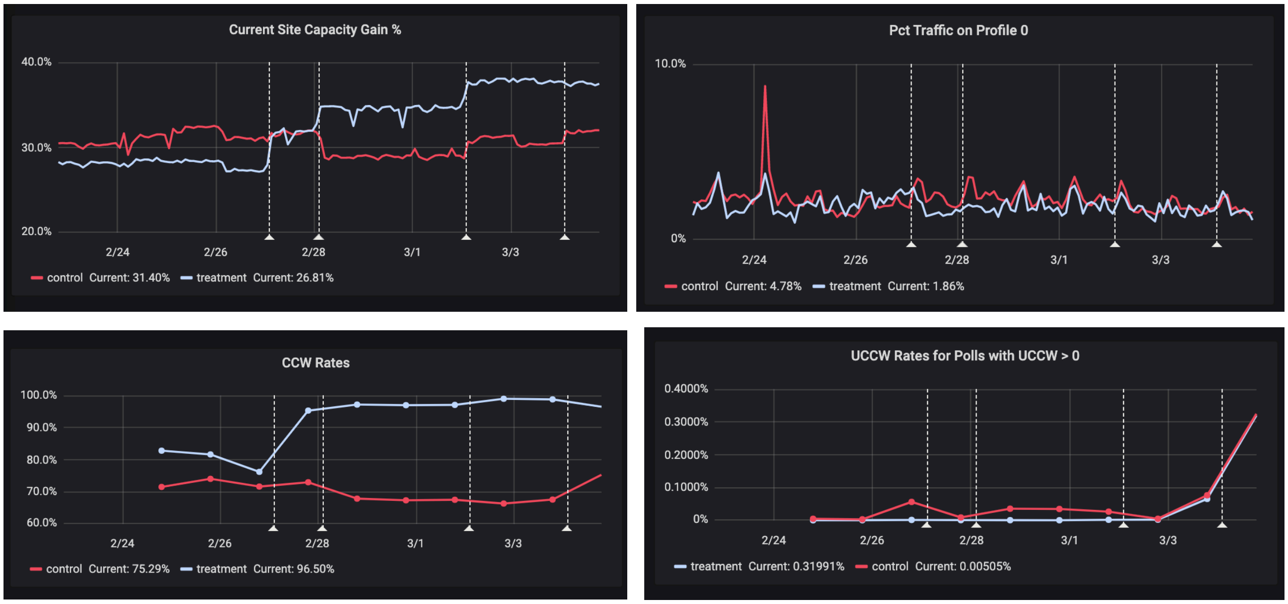

Figure 7 - Key results from the experiment in which modulation mapping thresholds were relaxed.

While the results shown in Figure 7 are encouraging and point to the possibility of capturing additional gains by tweaking the global policy, we opted not to adopt such strategy for optimizing capacity gains. The main motivation is the fact that plant conditions vary across the network. A one-threshold-fits-all policy may mask problems in areas were impairments exist and where certain population of devices may benefit from some added robustness (even moving the threshold in the opposite direction, to scarify capacity for robustness). Hence, our current effort focuses on developing a Machine Learning (ML) control system that offers much greater flexibility in configuring the profiles, including adjusting the modulation thresholds down to within a specific profile exception zone (segment). Other policy attributes may similarly be modified, such as the statistical aggregations used across time and frequency, before and after clustering of modems.

One methodology that shows promise in this realm is reinforcement learning (RL). In RL, the ML agent “learns” an optimal policy by interacting with the environment. The outcome is akin to allowing the agent to dynamically modify the MER mapping thresholds, or other policy attributes, per OFDM channel-profile-exception zone (segment) and based on feedback in the form of the FEC error rates encountered by devices. We are currently in the midst building a RL solution for US PMA. As will be shown in Section 4, US PMA has a limited action space, compared to DS PMA, and therefore it offers an opportunity to experiment with and refine the solution with the expectation that these methods will be subsequently adapted to be used for DS PMA and D3.1 US OFDMA PMA.

4. Upstream DOCSIS 3.0 Profile Management

4.1. Overview

The goal for the D3.0 US PMA solution was to add enough US capacity to allow at least 1 year of bandwidth demand growth, based on forecasted Compound Annual Growth Rates (CAGR), without having to segment nodes using legacy technology. Our engineering models calculate capacity based on MAC layer data rates, net of all physical layer and time-based overhead. Based on these models and progress to-date, using both D3.0 upstream PMA and adding the 5th & 6th channels, we have increased the upstream MAC layer data rate capacity from approximately 86 Mbps to approximately 117 Mbps, an increase of 36%, exceeding a year's worth of upstream growth even if COVID traffic levels remain as described in Figures 1 & 2.

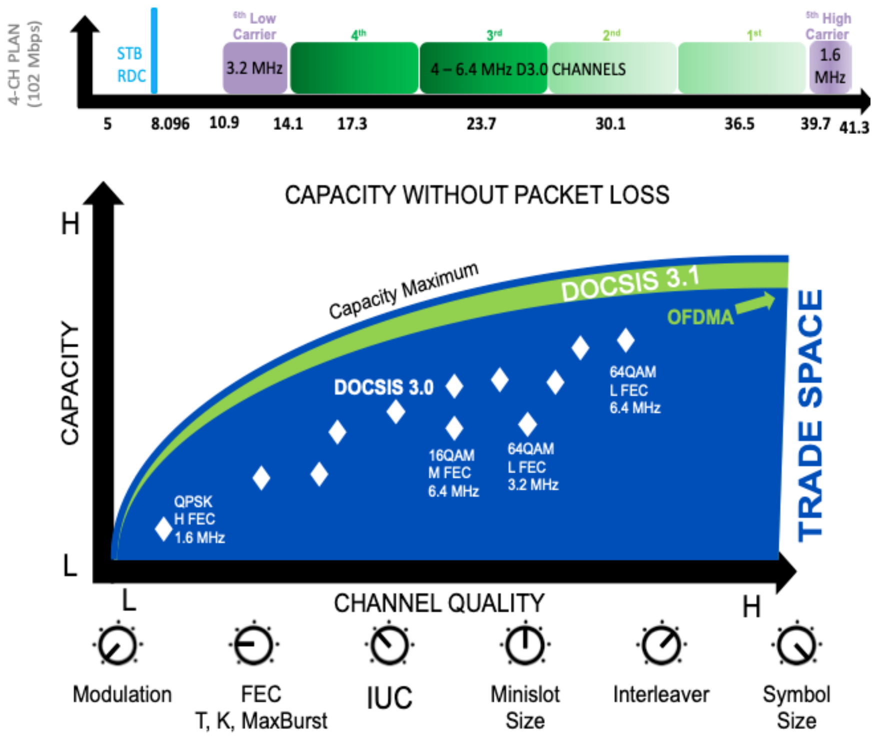

Figure 8 - Schematic demonstrating how upstream Profile Management would enable additional capacity in legacy low split HFC.

The profile management optimization process uses all the different knobs available in the upstream channel and modulation profile configuration to increase capacity, where channel quality allows it (see Figure 8). The top panel shows the addition of the 5th and 6th upstream channels; together, these are responsible for increasing capacity by approximately 15%. We are migrating to a standard channel plan across our sub-split nodes that enables a channel width of 3.2 MHz for the 5th US channel and 1.6 MHz for the 6th US, where both represent SC-QAM channels. Figure 8's bottom panel is a schematic highlighting the tradeoff between increasing capacity and increasing robustness. US PMA configures the knobs shown in the schematic to achieve the right balance. These tuning knobs are responsible for additional 20% capacity across all channels, above the robust defaults in use across the network in Q1 2020. Also, when needed, the solution can increase Forward Error Correction (FEC) and use more robust modulation to mitigate potential customer experience issues attributable to noise and ingress. These configuration updates are all done autonomously on the production network today. However, when the COVID-induced increase in upstream traffic began, the US automated execution flows had not been completed. Specifically, the configuration manager (CM) features were still under development.

As the capacity impact became apparent, the development team came together to work closely with the field operations and engineering teams, to calculate and statically configure more efficient modulation profiles with AE. The intent was to provide capacity relief and prove out the benefits of optimization algorithms on the production network, while simultaneously accelerating the development of the closed-loop autonomous system. The initial static deployments based on the profiles calculated by the AE, as the shelter-at-home orders peaked – more so in markets saturated with high-tech corporations -- had a significant positive effect on capacity. One example of an extreme upstream congestion case is shown in Figure 9. The top plot shows octet utilization for 4 channels, and the bottom plot the minislot and octet utilization of one of the optimized channels. Profile upgrades were done on the 1st and 2nd channels but not the 3rd and 4th channels. Consequently, approximately 10% more peak data is seen going through network with the initial efficient profiles. In this example, the minislot utilization and the octet utilization were at the maximum level for extended periods of time. When the static, more efficient profiles were applied to the 2 higher spectrum channels, the capacity was increased by 17% relative to the capacity of the two lower spectrum channels, which were not modified until the fully autonomous CM was available. Additionally, for the same level of minislot utilization, 10% more Mbps (octet utilization) were able to be sent by customers through the network than were previously transmittable. Correspondingly, the minutes of time at maximum utilization per day were subsequently reduced. Adding additional channels to further augment capacity on top of this example, along with other techniques, enabled a very quick response to and resolution of capacity hotspots.

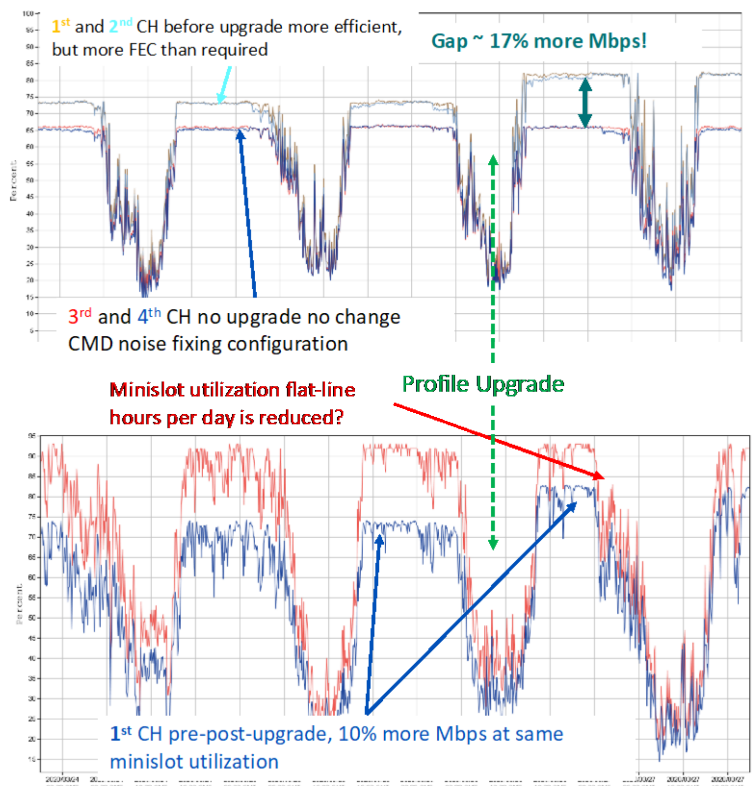

Figure 9. Peak COVID High Utilization Example.

Figure 10 shows an example of one of the operational dashboard's tracking system performance for the upstream PMA solution. Notice the activity in the profile distribution plots (identifiable by the colored bars around 7/18). This activity corresponds to the early stages of introducing a new CMTS, because the control system requires a few update cycles before converging to steady-state, at which time which most channels operate at the highest capacity profile (profile index 251). For reference, as additional CMTSs are added to the system, they are initialized to start on static, very robust profiles. These robust profiles are intended to combat a common transient noise source on the network in the lower 2 spectrum channels.

Additionally, and as noted earlier, a very robust but efficient modulation profile had also been statically deployed on the upper two spectrum channels, as described in the 2018 SCTE paper [9]. After accelerating the CM development and moving through trial and deployment, the automated system was turned on. It iterated through a set of modulation profiles, in small, low risk increments, and based on the channel's quality metrics. The metrics include correctable and uncorrectable codeword errors, signal-to-noise-ratio (SNR), partial service statistics, and minislot utilization. Over a period of approximately 24 hours, the modulation profiles were modified until they reached a steady-state based on the set control system strategy. The result is approximately 98% of the upstream channels are running effectively, without codeword errors, and around 2% are running on a profile addressing US noise sources such as transient switching power. The 21.2% capacity gain shown on the dashboard for these CMTSs is consistent across our network, approximately doubling the capacity gains of the example shown in Figure 9. The upshot is that the system autonomously optimizes all upstream channels, not just those in the higher quality spectrum.

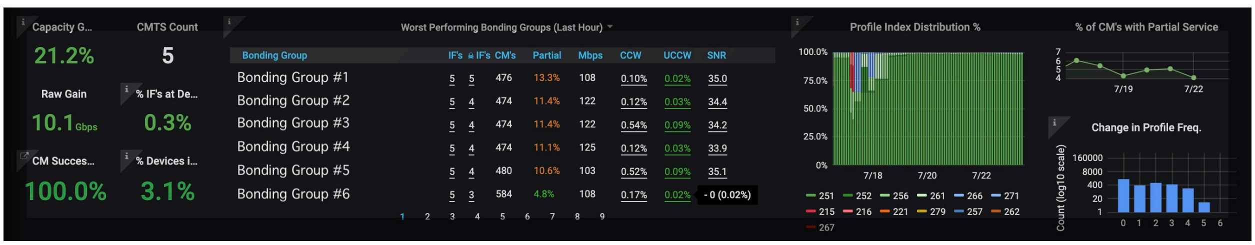

Figure 10 - US Dashboard example for a HUB of 5 new CMTSs managed by the US D3.0 PMA solution.

In addition to other key performance indicators, such as system settling time, capacity gain and error rates, there is the profile stability rate. Also shown on the dashboard is a distribution of how often an upstream channel has a profile update, highlighting remarkable network stability. In total, 98% of US channels require a profile update less than 1 time per week; 99% of US channels require a profile update less than 3 times a week; and < 0.5% of channels require profile updates more than once per day.

This stability also highlights the promise of D3.1 OFDMA profile management as the modem population grows, or the upstream spectrum available allows us to achieve the benefits of OFDMA. The authors plan to provide a future update on the OFDMA algorithm under development in 2020. What's exciting so far is that network quality is superior to that required for the best D3.0 capacities, and the profile management system will be able to take advantage of the higher order modulations available with D3.1.

4.2. How D3.0 US PMA works

The PMA system, as described previously, was extended to implement US D3.0 PMA functionality. The capacity is increased (or decreased) in small steps, while errors are fixed proportionally or predictively, either by increasing robustness or detecting transient noise indicators. When network issues are detected, operations staff is notified of network issues to drive remediation, ensuring maximum capacity is sustained and yielding operational savings. This is a "business as usual" step that happens regardless of whether the customer experience is being degraded due to packet loss.

The modulation profile capacities shown in the Figure 11 are based on compatible upstream channel configurations and channel widths. Similar templates exist for the narrower 3.2 MHz and 1.6 MHz channels. Figure 11 summarizes a subset of modulation profile attributes that must be set compatibly with the US channel attributes and other aspects, such as codeword size, preamble length, guard time, and interleaver settings. For example: profile 251 uses a 97 bytes payload and a 2 bytes parity for the short data grant, a 247 bytes payload and a 4 bytes parity for the long data grant, and designates the cutoff between short and long data grants to be 5 minislots in length (the burst size). The station maintenance and unsolicited grant service interval usage code (UGS IUCs) are similarly optimized to achieve the efficient use of minislots and required robustness. Each template consists of 25 profiles, constructed to comprehensively sample the parameter space, along the modulation and FEC regime dimensions. For example, profile 251 exhibits the highest modulation (256-QAM) and least robust profile (meaning the profile with the lowest FEC overhead.)

(lower) <-- Modulation --> (higher)

251 25.6 Mbps QAM64 short: 97/2, burst=5 long: 247/4 SNR for 1% error rate = 22.8 dB	256 24.5 Mbps QAM64 short: 91/5, burst=5 long: 239/8 SNR for 1% error rate = 22.1 dB	261 23.3 Mbps QAM64 short: 105/10, burst=6 long: 229/13 SNR for 1% error rate = 21.2 dB	266 22.5 Mbps QAM64 short: 99/13, burst=6 long: 223/16 SNR for 1% error rate = 20.5 dB	271 20.7 Mbps QAM64 short: 99/13, burst=6 long: 121/16 SNR for 1% error rate = 20.3 dB
252 21.4 Mbps QAM32 short: 98/2, burst=6 long: 247/4 SNR for 1% error rate = 20 dB	257 20.5 Mbps QAM32 short: 92/5, burst=6 long: 239/8 SNR for 1% error rate = 18.8 dB	262 19.4 Mbps QAM32 short: 102/10, burst=7 long: 229/13 SNR for 1% error rate = 18.2 dB	267 18.8 Mbps QAM32 short: 96/13, burst=7 long: 223/16 SNR for 1% error rate = 17.5 dB	272 17.8 Mbps QAM32 short: 96/13, burst=7 long: 138/16 SNR for 1% error rate = 17.3 dB
253 17.2 Mbps QAM16 short: 91/2, burst=7 long: 247/4 SNR for 1% error rate = 16.6 dB	258 16.5 Mbps QAM16 short: 101/5, burst=8 long: 239/8 SNR for 1% error rate = 15.9 dB	263 15.6 Mbps QAM16 short: 91/10, burst=8 long: 229/13 SNR for 1% error rate = 15.2 dB	268 15.1 Mbps QAM16 short: 101/13, burst=9 long: 223/16 SNR for 1% error rate = 14.6 dB	273 14.3 Mbps QAM16 short: 101/13, burst=9 long: 138/16 SNR for 1% error rate = 14.4 dB
254 13 Mbps QAM8 short: 100/2, burst=10 long: 247/4 SNR for 1% error rate = 14.7 dB	259 12.4 Mbps QAM8 short: 94/5, burst=10 long: 239/8 SNR for 1% error rate = 14 dB	264 11.7 Mbps QAM8 short: 96/10, burst=11 long: 229/13 SNR for 1% error rate = 13.2 dB	269 11.3 Mbps QAM8 short: 90/13, burst=11 long: 223/16 SNR for 1% error rate = 12.5 dB	274 10.8 Mbps QAM8 short: 90/13, burst=11 long: 146/16 SNR for 1% error rate = 12.4 dB
255 8.7 Mbps QPSK short: 93/2, burst=14 long: 247/4 SNR for 1% error rate = 10.3 dB	260 8.3 Mbps QPSK short: 95/5, burst=15 long: 239/8 SNR for 1% error rate = 9.3 dB	265 7.8 Mbps QPSK short: 85/10, burst=15 long: 229/13 SNR for 1% error rate = 8.6 dB	270 7.6 Mbps QPSK short: 87/13, burst=16 long: 223/16 SNR for 1% error rate = 7.9 dB	275 6.7 Mbps QPSK short: 89/15, burst=17 long: 104/16 SNR for 1% error rate = 7.6 dB

(efficient) <-- FEC Regime --> (robust)

Figure 11 - The profile configuration template for 6.4 MHz-wide channel.

Figure 12 describes how the whole PMA system came together. In the top schematic, theoretical models for noise are shown, as well as measured MER and traffic data that informs the construction of the US profiles. The bottom left schematic shows the result, which is a profile matrix configuration that spans a wide range of operating conditions. On the bottom right, system response is shown, in terms of error rates along with a defined policy that informs the selection of a suitable profile. This iterative process continues indefinitely as network conditions change. Note that the current profile update cycle is once every 6 hours. This value is configurable through policy and we are currently conducting A/B testing on a small subset of the CMTSs to discern if a higher update frequency (~hourly) further improves performance. In more detail:

- Models were developed for different transient and white noise channel models for each of the profiles.
- The range of network metrics were analyzed across the different spectrum locations, coming from over 50M+ DOCSIS devices, to understand the range of performance expected on the production network.
- Distribution of traffic packet and concatenated burst data was obtained to fine tune the codeword efficiency calculations.
- This data was fed into models of profile design, resulting in a set of static modulation profiles that cover a wide range of operating conditions and capacity yields.
- These profiles were tested in an automated lab system that automatically configures the profile on the upstream channel and injects the additive white Gaussian noise AWGN or transient noise, at increasing intensity, in increments of a fraction of a dB. Simultaneously, the system measures the CMTS upstream metrics for SNR and metrics related to codewords and traffic generators. The lab results were then compared to the theoretical models and matched to within +/- 1 dB. These lab and modeling results were then used to set the profile management application policy.
- The analytics engine, based on these models, analyzes the network data in real time. Based on statistics across time samples, and driven by policy, it selects the correct modulation profile for the current conditions. A policy, for example, could limit the operation to a subset of the matrix of profiles, such as limiting the low spectrum 6th channel to only operate in lower modulation transient noise profiles, or constraining the higher spectrum channels to only operate in the top 2 rows of the matrix, depending on channel minislot utilization levels.
- The analytics engine then closes the loop by continuing to measure network performance. It adjusts the modulation profiles as required, and based on network performance statistics intended to keep the errors at a healthy level, also based on policy.
- This closed loop control system is modifiable using Reinforcement Learning (see Section 4.4).

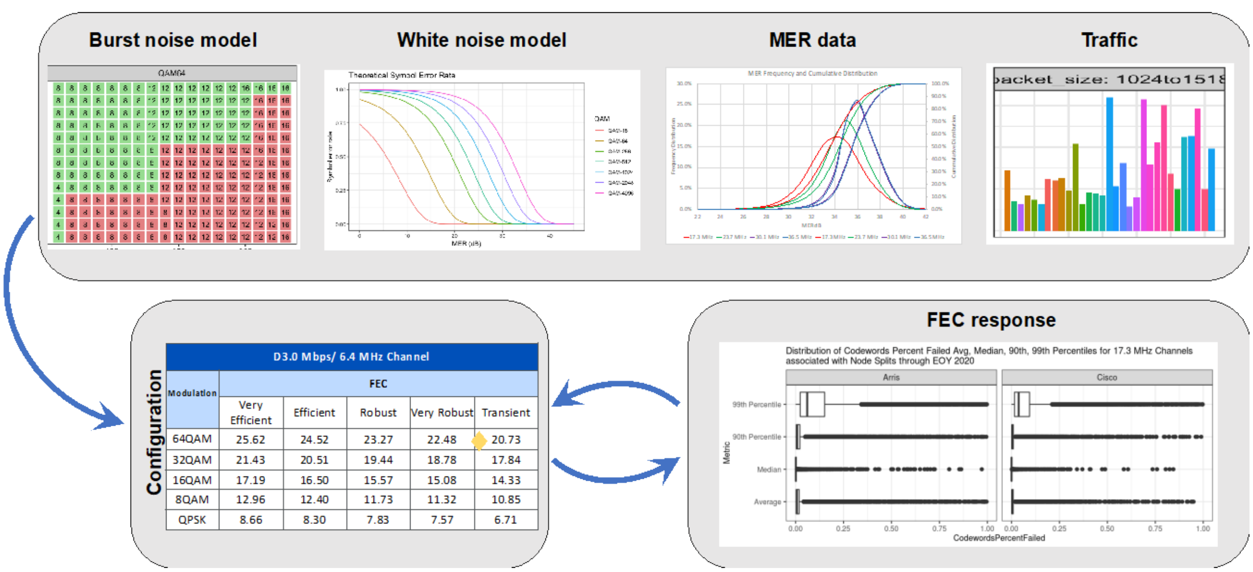

Figure 12 - D3.0 PMA System Comes Together.

4.3. Why D3.0 US Optimization increase network capacity?

The core idea behind the US profiles is based on changing the FEC configuration to play the trade-off between increasing codeword packing efficiency and increasing robustness against noise (be it white noise or high frequency transient noise). The concept is illustrated in Figure 13, which shows an illustration of the anatomy of a FEC codeword for a range of modulation profiles. In this example, moving from the transient to the very efficient profile would allow transmitting the same packet while leaving sufficient time for two additional DNS requests to a website (such as “www.ieee.org”) that otherwise would have been occupied by un-necessary physical layer and time overhead. The codeword includes a preamble at the beginning, a guard interval at the end, and portions for data (payload) and parity bytes in between. The lengths of all of these components are defined as part of an US profile configuration. In addition, D3.0 allows the definition of different FEC configurations for different traffic packet types (short data grants, long data grants, and voice.) This is where the problem becomes an optimization problem: Knowledge of traffic and noise patterns can inform proper profile configuration, to achieve high packing efficiency (defined as data payload length/codeword length), while maintaining error rates at an acceptable level. Rather than constructing the “right” profile for each US interface, we created a global template of 32 profiles covering a range of operating regimes and QAM modulations (a subset of which is shown in Figure 11 for 6.4 MHz wide channels). Notice how varying the packing efficiency is achieved by changing the FEC configuration: the long data grant for profile 251 (the most efficient 64-QAM profile) has a payload length of 247 bytes and a parity length of 4 bytes, while profile 271 (designed to deal with transient noise) has a payload length of 121 bytes and a parity length of 16 bytes. Consequently, profile 251 has a speed that is approximately 20% larger than profile 271. Note that the values stated above are not arbitrary. They were designed in consideration of the actual traffic packet size distribution on the network, and the type of transient noise typically experienced due to DC power supplies (the so-called common mode disturbance (CMD) noise described in Ref. 7).

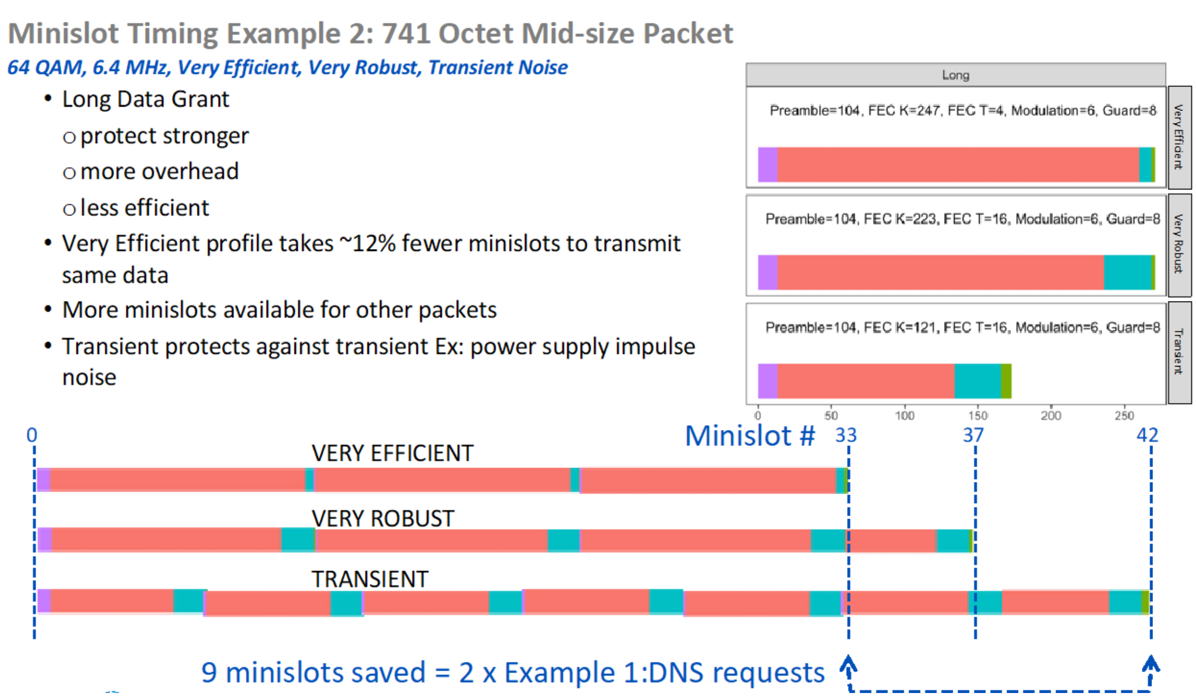

Figure 13 - Anatomy of a FEC codeword showing example of different codeword configurations for the D3.0 US long data grant.

4.4. The Reinforcement Learning approach to US PMA

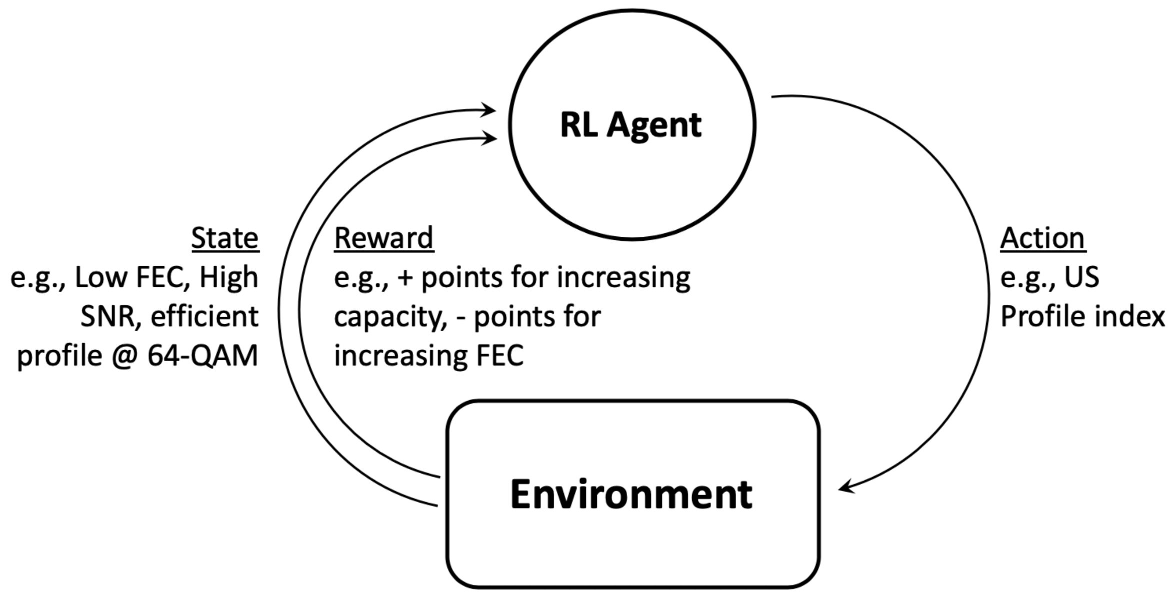

Figure 14 - Schematic of the Reinforcement Learning system for US PMA.

The general scheme for the Reinforcement Learning (RL) problem is shown in Figure 14. The goal of the agent, as it interacts with the environment, is to learn an optimal policy that maps states into actions. In the context of upstream PMA, states represent channel telemetry and configuration, while actions represent the upstream profile choice. The ideas behind RL are based on decades-old concepts in control theory and dynamic programming. What is unique to RL is the concept of learning from "trial and error." That is, the agent will at times purposely select a random, non-optimal action in order to continue exploration of the environment. This is useful for two reasons: (1) The number of states, in combination with possible actions, is too vast to sample in a systematic way. (2) Even with sufficient knowledge of the dynamics of the environment at a given point in time, the environment is typically non-stationary. Changes in these dynamics warrant that the agent maintains a level of exploration over time. Our first attempt at representing the US PMA control problem as an RL problem included defining the following components:

- **States**: FEC error rates, traffic volume, and device SNR aggregated to the channel level are discretized by binning the continuous variables according to some given thresholds (e.g. uncorrectable error rate > 1% is High, <1% but >0.1% is Medium, and <0.1% is Low). The profile modulation (QPSK to 64-QAM) and profile type (most efficient to most robust) are discrete to begin with. The resulting state space includes ~6,000 states.

- **Actions**: We limit the agent's actions to changing the profile level by +/-3 steps (in either direction, i.e. upgrading or downgrading by *n* steps), maintaining the same state, or transitioning in or out of the special profile designed for transient noise. Thus, the total number of possible actions is 9.

- **Rewards**: These are scalar values allocated to the agent following each action taken. They are designed to incentivize the agent to maximize returns in the long run. Rewards proportional to the change in profile level are awarded (e.g. upgrading the profile by 3 steps incurs a reward of +3). However, experiencing an uncorrectable error rate above 1% incurs a reward of -10, as this is deemed to be negatively impactful to the customer experience.

- **Policy**: The policy is a mapping from states to actions—i.e. it represents the decision process. The agent starts with an initial policy that resembles the fixed rule-based approach in the live production PMA pipeline. Example: if FEC rates are below a certain required threshold and the SNR level is above another required threshold, then the channel is reconfigured with a profile that is one step higher than the current one.

With the above in place, the goal of the agent is to discover and maintain an optimal policy. This is done by maintaining a degree of exploration. Typically, the agent would select a random action 1% of the time and the optimal action 99% of the time, under what is known as the epsilon-greedy policy (with epsilon=0.01). Initially, there will be a subset of channels running on a randomized profile between update cycles (currently 6 hours). The risk of causing an adverse effect is mitigated by the fact that the action space limits the magnitude if the upgrade to 3 steps. Furthermore, as more states are explored and convergence to an optimal policy is achieved, the amount of exploration could be dialed down to below 1%. Within the core RL algorithm, as new state-action pairs get explored, the agent collects its rewards based on changes in the environment (perceived through the state) and updates a table representing the long run return for each combination of state-action. This table can be trivially converted into a "most optimal" policy at the time. The update process itself is founded in theory relating to Markov Decision Processes (MDP). These are a special type of processes in which the current state is sufficient to describe all the preceding history of the process. Modeling the RL problem as an MDP allows the use of a host of sampling-based approaches to efficiently update the state-action values. In addition, MDP theory predicts convergence to an optimum policy, given a stationary environment and sufficient number of update iterations. Our current efforts are focused on implementing the RL algorithm and testing it on a small subset of production CMTSs, in order to assess how well it performs against the existing static policy.

5. Conclusion

As of this writing, Comcast is currently optimizing capacity across thousands of CMTSs and hundreds of thousands of OFDM and D3.0 upstream channels for more than 50 million DOCSIS gateways and cable modems. The system processes tens of terabytes of data and performs approximately 500,000 recommendations and transactions per day to optimize the upstream and downstream channels. To date, this has yielded capacity improvements of more than 30% in the downstream (towards customers) direction, and of approximately 20% in the upstream direction (from customers.)

To achieve scale and very high levels of reliability, PMA was developed as a serverless, elastic, cloud-native solution—taking advantage of distributed compute, storage, and network options. One of the primary design patterns was to allow for component independence: loose coupling with strong interoperability to promote feature velocity and a high cadence of delivery. Adopting these basic principles enabled each of the functional components and sub-components to have independent, automated, and continuous integration and deployment trains. This allowed for system updates to scale to business needs, often being updated seamlessly and in parallel several times a day.

There are real opportunities to capture additional efficiencies from invested capital in the network. Along this journey we learned a couple things:

- *Keep it simple: It matters to scale.* Network capacity is a complex environment. Modern access networks are complicated; DOCSIS technology is complicated; CPE devices are complicated and required many firmware updates. CMTS solutions are heterogenous and required many new features to be developed. The number of devices and variety of devices are complicated; their different firmware versions, also complicated. Complex systems often exhibit emergent and unexpected behavior and require special consideration when building to scale. Best to keep it as simple as possible (advice perhaps easier said than done!)

- *Leverage the cloud.* The benefits of the cloud are real, and public cloud is an accelerator. Leveraging public cloud resources allows you to focus resources on key problems. It allows you to reliably reach scale, at a velocity that meets business demands.
- *Keep learning through iterative builds.* For best results, build on platforms that can be updated seamlessly and constantly tested. We could not have anticipated all the conditions that we encountered and benefited greatly from a team and culture that was built on iterative learning.
- *Be data-driven by measuring what matters*. Along the entire journey, from early analysis to production dashboards and quality metrics, trusting the data and developing a data-driven strategy allowed us to match practice to theory and build confidence. Being data driven will drive you toward success.

Abbreviations

AE	Analytics Engine
AI	artificial intelligence
API	application programming interface
AWGN	additive white gaussian noise
CAGR	compounded annual growth rates
CLI	command-line interface
CMTS	cable modem termination system
CMD	common mode disturbance
DOCSIS	Data Over Cable Service Interface Specification
D3.0	DOCSIS 3.0
D3.1	DOCSIS 3.1
DE	Data Engine
DS	downstream
FDX	Full duplex
FEC	forward error correction
HFC	hybrid fiber-coaxial
JSON	JavaScript Object Notation
LDPC	low density parity check
MER	modulation error ratio
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
PHY	physical layer
PLC	physical link channel
PMA	Profile Management Application
QAM	quadrature amplitude modulation
RL	Reinforcement learning
SCTE	Society of Cable Telecommunications Engineers
SNR	signal to noise ratio
UGS IUC	unsolicited grant service interval usage code
US	upstream

Bibliography & References

1. *"A Machine Learning Pipeline for D3.1 Profile Management", M. Harb, J. Ferreira, D. Rice, B. Santangelo, and R. Spanbauer, NCTA technical paper, 2019.*
2. *"Our Network", Comcast corporate website, https://corporate.comcast.com/our-network*
3. *"Comcast Network Investment", Comcast press update, https://update.comcast.com/download/15781/*
4. *"As remote work exploded, Comcast turned to AI to keep the internet running", FastCompany article, https://www.fastcompany.com/90519167/as-remote-work-exploded-comcast-turned-to-ai-to-keep-the-internet-running, 2020.*
5. *COVID-19: How Cable's Internet Networks Are Performing: METRICS, TRENDS & OBSERVATIONS, https://www.ncta.com/COVIDdashboard.*
6. *https://www.nctatechnicalpapers.com/Paper/2002/2002-optimizing-the-last-mile*
7. *Convolutional Neural Networks for Proactive Network Management: Developing Machine Learning Models to Detect and Classify Impairments in D3.1 OFDM Channels, SCTE Cable-Tec Expo 2020, Ferreira, Harb, Subramanya*
8. *"DOCSIS 3.1 Downstream Early Lessons Learned", J. J. Downey, NCTA technical paper, 2017.*
9. *What Gets Measured Gets Done / What Gets Analyzed Gets Transformed, Mutalik, Rice, Subramanya, Wang, SCTE IBSE Cable-Tec Expo 2018, https://www.nctatechnicalpapers.com/Paper/2018/2018-analytics-for-a-wider-deeper-network-view*

Additional References

ANSI C63.5-2006: American National Standard Electromagnetic Compatibility–Radiated Emission Measurements in Electromagnetic Interference (EMI) Control–Calibration of Antennas (9 kHz to 40 GHz); Institute of Electrical and Electronics Engineers

The ARRL Antenna Book, 20th Ed.; American Radio Relay League

Code of Federal Regulations, Title 47, Part 76

Reflections: Transmission Lines and Antennas, M. Walter Maxwell; American Radio Relay League