Accounting for Every MHz of Bandwidth: Data & Algorithms for Artifact Discovery and Close-Packing of QAMs in Support of Spectrum Activation

A Technical Paper prepared for SCTE by

Maher Harb
Distinguished Engineer, Data Science
Comcast
1800 Arch St, Philadelphia, PA 19103
+1 (215) 990-8376
maher_harb@comcast.com

Wenlong Shen
Senior Engineer, Machine Learning
Comcast
1800 Arch St, Philadelphia, PA 19103
+1 (267) 260-3798
wenlong_shen@comcast.com

Matt Stehman
Senior Engineer, Machine Learning
Comcast
1800 Arch St, Philadelphia, PA 19103
+1 (267) 260-3956
matthew_stehman@comcast.com

Sanket Walavalkar
Executive Director, Next Generation Access Network
Comcast
1800 Arch St, Philadelphia, PA 19103
+1 (917) 592-6789
sanket_walavalkar@comcast.com

Dan Rice
VP, Next Generation Access Network
Comcast
4100 E Dry Creek Rd, Centennial, CO 80122
+1 (720) 512-3730
daniel_rice4@comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

Comcast's network is undergoing an evolution towards the virtualized Cable Modem Termination System (vCMTS) and remote physical layer (PHY) architecture as part of the journey to 10G [1]. As of early 2023, millions of homes passed can enjoy the multi-Gbps downstream and a few 100 Mbps upstream speeds made possible by deploying a full-width orthogonal frequency division multiplexing (OFDM) channel at 810-1002 megahertz (MHz) and an orthogonal frequency division multiple access (OFDMA) channel in the mid-split (40-85 MHz) region of the spectrum across the vCMTS platform. The path to 10G involves incrementally deploying additional OFDM/OFDMA channels under the Data Over Cable Service Interface Specification (DOCSIS) 4.0 Full Duplex (FDX) technology [2]. However, reconfiguring the spectrum across markets and localities in support of the 10G roadmap is a daunting task because of the complex spectrum management involving the linear video and single carrier-quadrature amplitude modulation (SC-QAM) channels that need to be accommodated, moved, or phased out by converting linear QAM-based video to IP-based video. Adding to the challenge is the prevalence of local video insertions used by multi-dwelling units (MDUs), managed by the local regions, that may not be properly documented in a central location. In this technical paper, we introduce a methodology for discovering spectrum occupancy based on an object-oriented model of the cable modem's full band capture (FBC). Alongside the configuration data collected from the vCMTS, the FBC model allows the discovery of a host of artifacts that, if not addressed, would be problematic for turning up energy in new regions of the spectrum. The introduced methodology served as the basis for a machine learning pipeline established for automating the required spectrum enablement pre-checks and for generating detailed reporting on all discovered issues.

Furthermore, the deployed pipeline for artifact discovery was repurposed to support the effort for deploying a second OFDM channel on the vCMTS platform within the 618-810 MHz range. For this second phase of spectrum activation, the pipeline was also tasked with generating remote PHY device (RPD) configurations tailored for each node, with the goal of eliminating the vacant spectrum (gaps) between video and downstream SC-QAM and creating OFDM exclusion bands informed both by regional policies and by the discovery of insertions on the node.

Lastly, we also present our early explorations of the third phase of spectrum activation, which involves repacking video channels to free up the spectrum for FDX. Traditionally, video programming templates have been constructed locally at a market level and were not directly subject to centralized planning and optimization. These early explorations reveal opportunities to free up the spectrum by ensuring that each video QAM is fully packed with multiplexed programs (TV channels). Developing an artificial intelligence (AI) driven pipeline to automatically calculate the capacity requirements and optimal spectrum allocations given a variety of variables such a customer premise equipment (CPE) and network technology penetrations, traffic load, infrastructure and technology constraints, dynamic efficiencies from PHY and media access control (MAC) layer optimization micro-services is our vision. Fully automating the spectrum management per service group not only optimizes capacity and speed, but removes the inevitable human errors, time to configure and inefficiencies required when operating networks at the scale of 10s of millions of broadband customers.

2. Phase 1 of Activations: Mid-split and OFDM at 810-1002 MHz

In a previous NCTA technical paper, we shared our experience around enabling the mid-split spectrum by activating OFDMA on the vCMTS platform [3]. The focus of that contribution was on the profile management application (PMA) solution for constructing interval usage code (IUC) profiles, the capacity performance of OFDMA, and the upstream speeds that can be made accessible to our customers with OFDMA. Since then, thousands of nodes are now mid-split enabled. The deployment of mid-split occurred in tandem with standardizing the OFDM channel placement to the 810-1002 MHz frequency range. The rationale behind syncing the 2 efforts (mid-split enablement and OFDM expansion) is to (1) offer higher upstream & downstream speeds simultaneously and (2) minimize the number of incremental configuration changes toward the end-state, in which FDX is fully deployed and symmetrical multi gigabit per second (Gbps) speeds are possible.

The 2 main challenges in the deployment journey were "cleaning up" the spectrum channel map to make it compatible with the newly opened regions of the spectrum and implementing automations around pre-enablement checks, and around the generation & transaction of the spectrum configurations that are tailored to each node on the network. To this end, one obvious exercise involved swapping legacy amplifiers for ones that are mid-split and 1 gigahertz (GHz) capable. However, issues surrounding deployment go beyond missing amplifier upgrades. To highlight those challenges, consider the high variability in spectrum configuration before mid-split deployment depicted in Figure 1. The figure shows that the top 25 configurations cover ~46% of the pre-mid-split RPD population. Note that an active carrier was deliberately placed at 999 MHz to aid in discovering potential roll-off. In all the permutations shown in Figure 1, the 810-1002 MHz spectrum is almost free for placement of the OFDM channel. Though there are edge cases not included in the top 25 in which SC-QAM channels need to be relocated to free up the 810-1002 MHz region. Furthermore, some permutations explicitly show the existence of a local video insertion while others may have an active insertion that is undocumented.

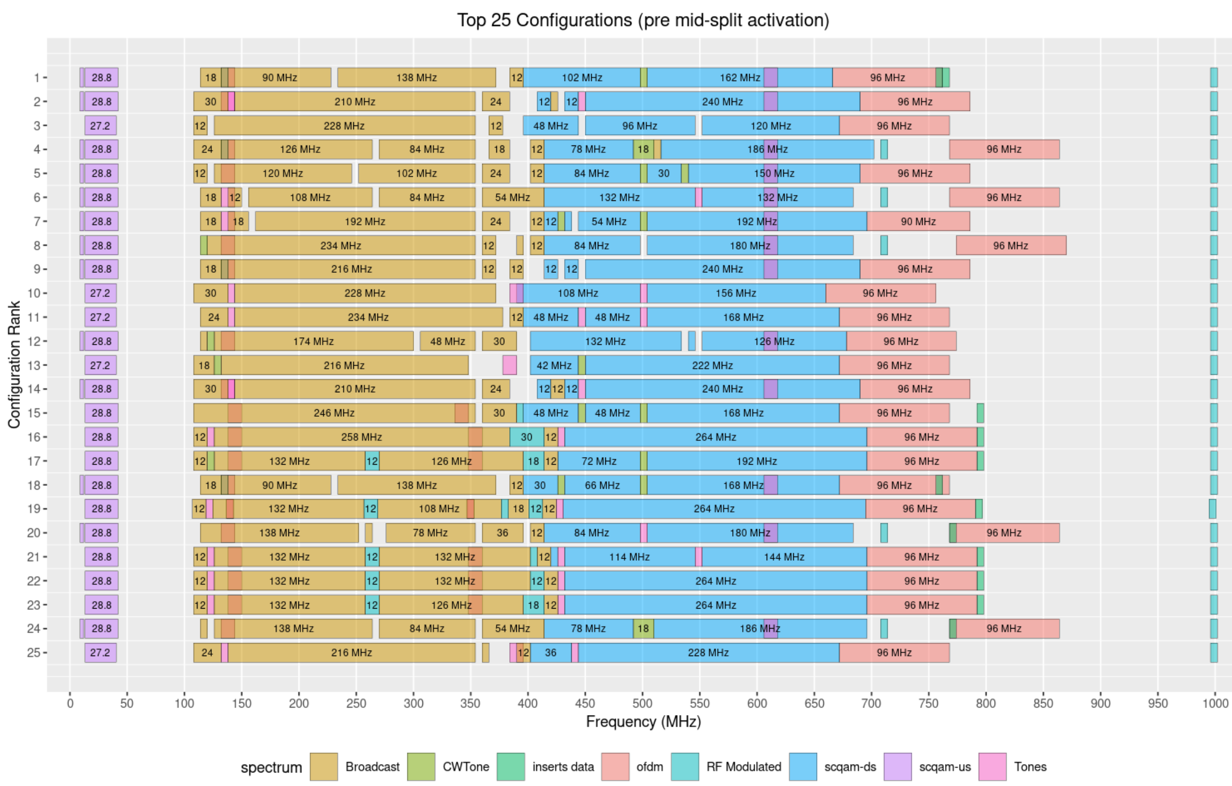

Figure 1 - Top 25 configurations for RPDs that haven't yet undergone mid-split activation. ~46% of this RPD population is covered by the top 25. Aside from the upstream SC-QAM block, there exists high variability across configurations.

In contrast, Figure 2 shows the top 25 configurations post mid-split activation, covering ~57% of this RPD population. The upstream and OFDM regions are now standardized across this population, while the high variability remains in the video and downstream SC-QAM regions (to be addressed over the next iterations of spectrum evolution).

Given the high variability in configurations, transforming the spectrum from the picture of Figure 1 to that of Figure 2 involved dealing with the following list of potential problems:

- Existence of undocumented local video insertions (usually serving an MDU) within the mid-split or the 810-1002 MHz region. In general, this type of problem extends to any new region of the spectrum to be used for activating DOCSIS channels.
- Existence of severe roll-off in the 810-1002 MHz region. This could be indicative of a missed legacy amplifier or localized to a tap serving a few dwellings.
- Existence of severe ingress in a currently vacant region of the spectrum, which will be used for relocating downstream SC-QAMs channels to free up the 810-1002 MHz for OFDM.

It was imperative to detect each of the scenarios above pre-deployment to avoid negatively impacting customers. Henceforth, we will refer to the existence of any of the above as an "artifact". In another previous technical paper, we shared a few concepts around artifact detection, based on an analysis of the cable modem's FBC [4]. These early explorations were ad-hoc in nature and involved developing a separate algorithm for the detection of each artifact. Next, we present a unified approach for artifact detection based on an object-oriented model of the FBC.

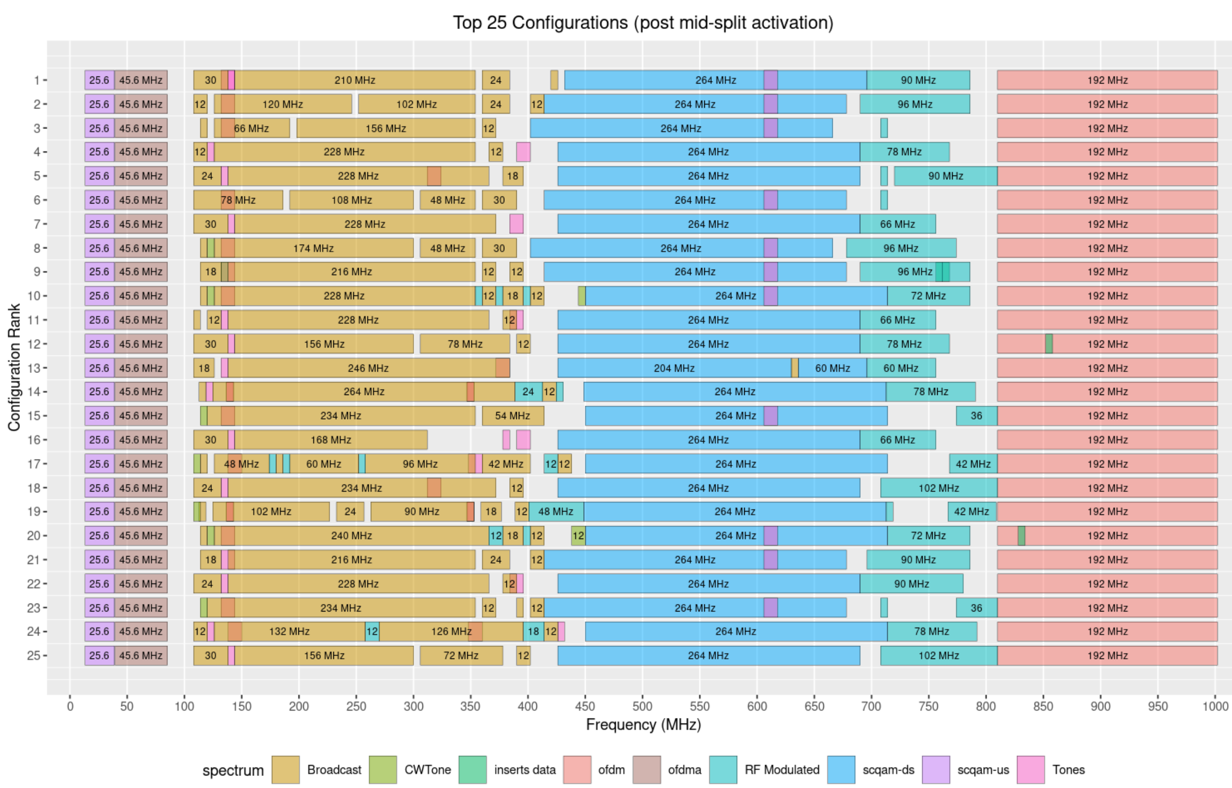

Figure 2 - Top 25 configurations for RPDs that are mid-split active. ~57% of this RPD population is covered by the top 25. Though upstream and OFDM blocks are aligned across permutations, there remains high variability across video and downstream SC-QAM.

3. An Object-Oriented Model for the Full Band Capture

The idea of the new approach to pre-enablement checks is to conduct these systematically rather than what was previously a set of ad-hoc rules, each relating to a specific region of the spectrum and a specific type of artifact. The new approach recognizes that any of the checks performed falls broadly under a single task: matching the spectrum configuration against the cable modem's FBC and deciding on whether the two views agree or conflict. The latter is an indication of a host of problems which we are interested in detecting. The high-level concept is summarized in the 2x2 decision matrix shown in Table 1.

It follows that we benefit from a generic approach to processing the FBC, in which signal is aggregated across each 6 MHz channel and a Low/High label is assigned to each based on some mechanism for designating that the channel exhibits an energy level commensurate with signal vs. noise.

Table 1 - 2×2 decision matrix showing how to interpret a channel's combination of energy & occupancy states.

		Channel (QAM) Energy (based on FBC)	
		Low	**High**
Channel (QAM) Occupancy (based on RPD configuration)	**Vacant**	• No evidence of noise or local insertions • Clear for placing QAMs/OFDM	• Possible local insertion • Possible ingress, etc. (some other RF (Radio Frequency) pattern)
	Occupied	• For a 999 MHz carrier, indicates existence of roll-off • Possible ingress, etc. (some other RF pattern)	• Healthy spectrum in general • For a 999 MHz carrier, indicates no evidence for roll-off

An object-oriented model was built for this purpose. The FBC object is instantiated with 3 inputs:

- **FBC data**: an array (float) representing the raw measurement typically at ~117 kilohertz (KHz) resolution.
- **Configuration data**: a dictionary in which keys are the type of spectrum and values are arrays of the (start,end) frequencies of the contiguous spectrum blocks. Example:

```
{'broadcast': [(114.0, 120.0), (126.0, 222.0), (228.0, 240.0), (246.0, 414.0),
               (426.0, 438.0)],
 'rf modulated': [(222.0, 228.0), (240.0, 246.0), (414.0, 426.0), (438.0, 450.0),
                  (714.0, 786.0)],
 'tones': [(120.0, 126.0), (138.0, 150.0), (348.0, 360.0)],
 'local inserts': [(792.0, 798.0)],
 'ofdm': [(810.0, 1002.0)],
 'ofdma': [(39.4, 85.0)],
 'scqam-ds': [(450.0, 714.0)],
 'scqam-us': [(13.2, 38.8)]}
```

- **HRC flag**: a Boolean indicating whether the plant follows the standard (False) or the harmonic related carrier (HRC) plan in which QAM channels are offset by 1.25 MHz (True).

Once instantiated, the following host of pre-processing steps are invoked by the class:

- The spectrum is "cleaned" by removing guard regions. These have fixed locations depending on whether the plan is standard vs. HRC.
- Energy is aggregated per channel/QAM according to the standard EIA mapping.
- A clustering algorithm determines the threshold that best separates the channels across spectrum into 2 groups: noise and signal. Consequently, each channel is labeled as Low/High depending on whether it belongs to the noise/signal group. It is critical to create such a designation without knowledge of the spectrum configuration according to the RPD.
- An additional label is added to the channel based on knowledge of the RPD configuration. It contains all spectrum types assigned to the channels.

Here's an example of instantiating the FBC object and querying a single channel:

```
example_spectrum = FullBandCapture(raw_spectrum=x, config=config_dict, hrc=is_hrc)
example_spectrum()[89]
```

Executing the code above returns the attributes of electronic industries association (EIA) channel 89 (corresponds to a center frequency of 615 MHz according to the EIA plan):

```
{'frequency': 615, 'range': (5174, 5218), 'energy': -7.98, 'configuration': ['scqam-ds', 'Tones'], 'label': 'High'}
```

We are now ready to "interrogate" the FBC object to check for the presence of artifacts of interest. Below are some examples.

3.1. MDU Insert within 810-1002 MHz

The example shown in Figure 3 represents detection of an MDU insert in the 810-1002 MHz region intended for OFDM expansion. With the instantiated FBC object, detecting inserts involves the following single line of (P*ython*) code:

```
{eia:attributes for eia,attributes in example_spectrum.channel_data.items()
if not any(attributes['configuration']) and
   attributes['label']=='High' and
   attributes['frequency']>810}
```

Executing the code above returns the suspect insert at 855 MHz:

```
{134: {'frequency': 855, 'range': (7222, 7266), 'energy': -19.98, 'configuration': [], 'label': 'High'}}
```


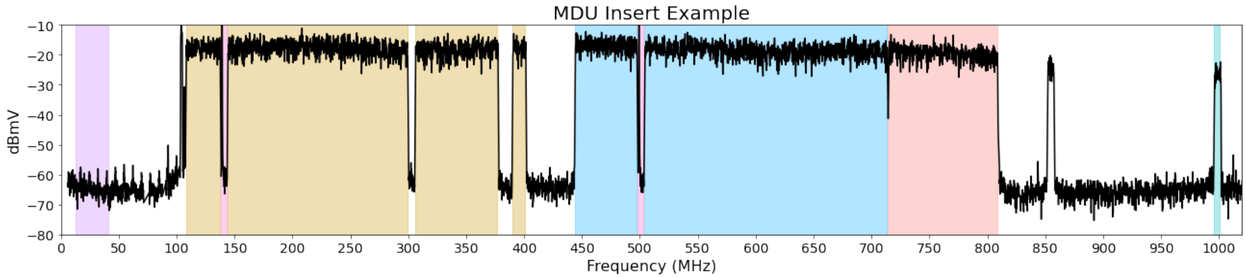

Figure 3 - Example FBC that shows an undocumented local insert at 855 MHz. The overlaid colored shadings correspond to the RPD configuration according to the color map in legend of Figure 1.

3.2. Roll-off within 810-1002 MHz

The example shown in Figure 4 represents detection of potential roll-off in the 810-1002 MHz region intended for OFDM expansion. The roll-off is inferred from the relatively low signal of the 999 MHz carrier (15 dBmV lower than the average signal level). Once again, with the instantiated FBC object, detecting roll-off involves executing the following single line of (P*ython*) code:

```
{eia:attributes for eia,attributes in example_spectrum.channel_data.items()
if attributes['configuration']==['RF Modulated'] and
   attributes['energy']<(example_spectrum.signal_level-15) and
   attributes['frequency']==999}
```

Executing the code above returns the low signal 999 MHz carrier:

```
[{158: {'frequency': 999, 'range': (8451, 8495), 'energy': -37.77, 'configuration':
['RF Modulated'], 'label': 'Low'}}]
```

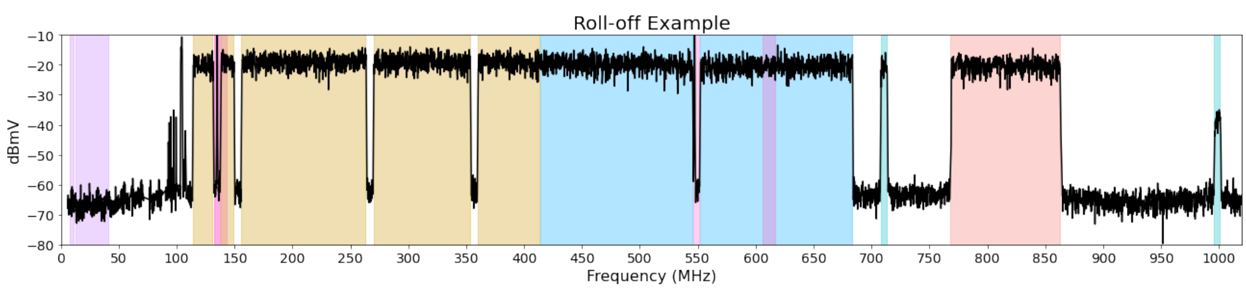

Figure 4 - Example FBC that shows roll-off at the 999 MHz carrier. The overlaid colored shadings correspond to the RPD configuration according to the color map in legend of Figure 1.

4. The Pipeline for Artifact Detection

The pipeline for artifact detection scans the network once a day, searching for the type of artifacts described in the previous section. Note that the “heavy lifting” is done by the FBC model, which, in effect, processes the raw data in the same way regardless of the intended artifact. Since the detection mechanism based on FBC discovers artifacts at a cable modem’s level, this data is aggregated to the RPD level for decision-making on whether to move forward with spectrum activation or block the RPD until the discovered issues are resolved. This aggregation is rule-based, requiring a certain number of cable modems to exhibit the artifact for the RPD to be blocked from spectrum activation. A typical threshold used is ~10 cable modems. In setting the threshold, we attempt to strike a delicate balance between False Positives and False Negatives. One approach for threshold selection is based on constructing a sensitivity curve such as the one shown in Figure 5. A rule-of-thumb for threshold selection follows the "elbow method"; i.e., setting the threshold close is the elbow within the arm-resembling sensitivity curve. This method does not consider the unequal impact of generating False Positive and False. For example, if the consequence of False Negatives is catastrophic then it makes sense to shift the threshold further to the left on the curve (i.e., threshold < 10).

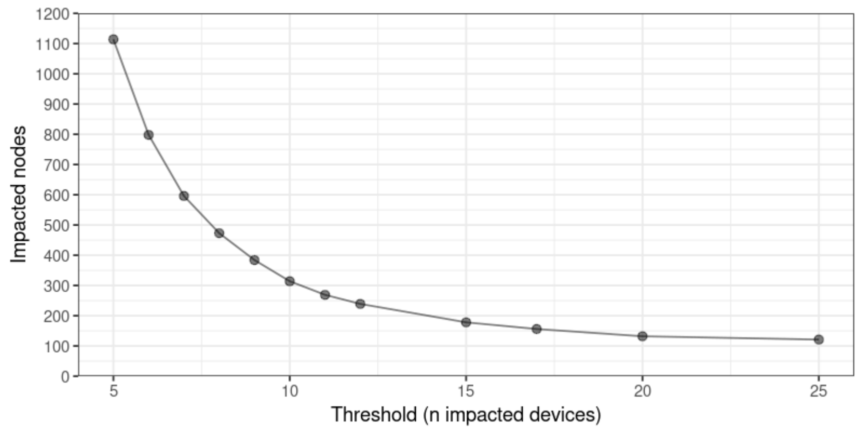

Figure 5 - Sensitivity curve for insertion detection. The *x*-axis is the selection threshold in the number of cable modems that show roll-off. The *y*-axis is the corresponding number of impacted nodes. The "elbow method" would set the threshold close to ~10.

Since the pipeline uses fresh telemetry, once the issue is fixed the blocked RPD will automatically be removed from the blocked list for spectrum activation on the following pipeline run. Therefore, manual management of a blocked/cleared RPD list is not required.

The pipeline also employs algorithms for narrowing down the root cause of the issue where possible. For example, when severe roll-off is observed impacting many devices on the node, we invoke graph algorithms powered by the graph representation of the network [5, 6] to determine if the issue is related to a missed amplifier vs. multiple separate and unrelated problems at each dwelling. This additional context helps the field techs to quickly narrow down the search for the root cause of the problem or work with the customers to resolve their home issues.

Figure 6 shows an example root cause analysis in which 23 cable modems exhibited severe roll-off at the 999 MHz carrier (depicted as pink-shaded rounded boxes on the graph). A least common ancestor algorithm determined that the source of the problem is the amplifier marked with an asterisk. This sort of finding is invaluable to the field techs to swap the missed amplifier in a timely manner.

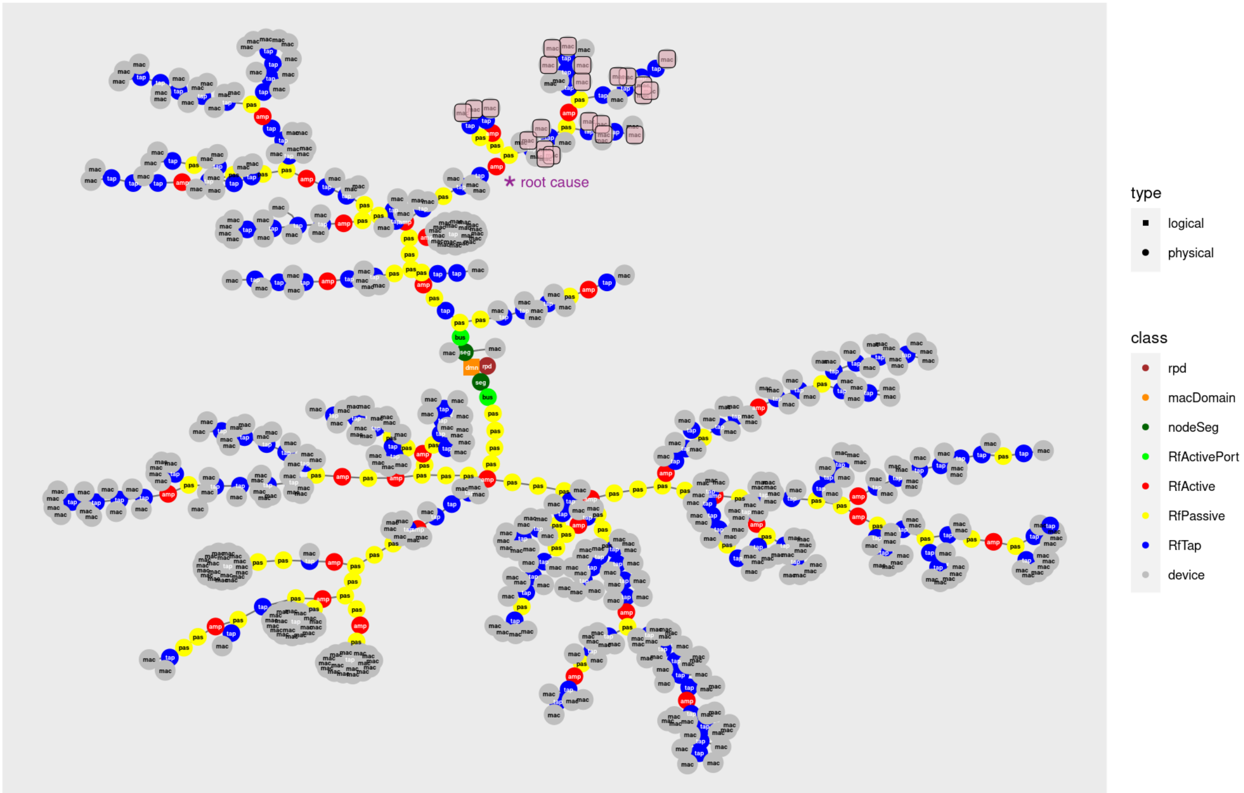

Figure 6 - A graph of the network showing all elements under a node. Cable modems with roll-off are highlighted in pink (top center area). The amplifier (red circle) responsible for the roll-off is marked with an asterisk. This amplifier was missed in the upgrade construction activity.

5. Phase 2 of Activations: 2nd OFDM

Phase 2 of spectrum activation is centered around adding a 2nd OFDM channel. This time however, the OFDM channel will not have a standardized location but will be constrained by the availability of spectrum from the edge of the video block (varies by node) to the edge of the first OFDM channel (810 MHz). The basic set of rules that govern the placement of the 2nd OFDM channel is the following:

- Move the downstream SC-QAM spectrum to the top edge of the video and remove any gaps within the downstream SC-QAM blocks, if available and permissible.
- Place a second OFDM channel with an upper edge at 810 MHz (lower edge of the first OFDM channel) and a lower edge at either the top of the downstream SC-QAM or 618 MHz (whichever is higher); this ensures that the OFDM channel width is no larger than 192 MHz.
- Create exclusion bands within the OFDM channel to accommodate detected MDU inserts as well as static exclusion frequencies dictated by policy. These static exclusions are intended to support several standard locations for video channels that will facilitate the video shift for FDX by simulcasting some video controller data and a standard location for QAM video in MDUs compatible with DOCSIS 4.0 spectrum.
- Check that created channel satisfies DOCSIS specifications; for example, in relation to the ratio of exclusion regions to usable OFDM spectrum, the proximity of exclusion bands to edges, etc.

- Create a JavaScript Object Notation (JSON) formatted file to include the new configuration (to be transacted by an automated flow).
- The physical link channel (PLC) and OFDM location are further configured to support cellular band leakage measurement markers based on learning of legacy analog nodes configuration clusters so that leakage is detectable across all nodes in the region without more complicated leakage configuration geo-zones.

An example of the above logic in action is shown in Figure 7. The top panel shows the current configuration of an RPD. Notice the wide vacancy between the broadcast video and the downstream SC-QAM block. For this RPD, regional policy dictates establishing exclusion bands reserved for future local insertion usage at 747, 753, and 795 MHz. The middle panel, however, reveals an undocumented insert active today at 777 MHz. The bottom panel shows the recommended configuration, in which SC-QAM is moved to the edge of the video, a 2nd OFDM is placed between 648 and 810 MHz, and 3 DOCSIS standard compliant exclusion bands are created to cover the 4 insert locations (note, 747 and 753 are adjacent, requiring a single exclusion band).

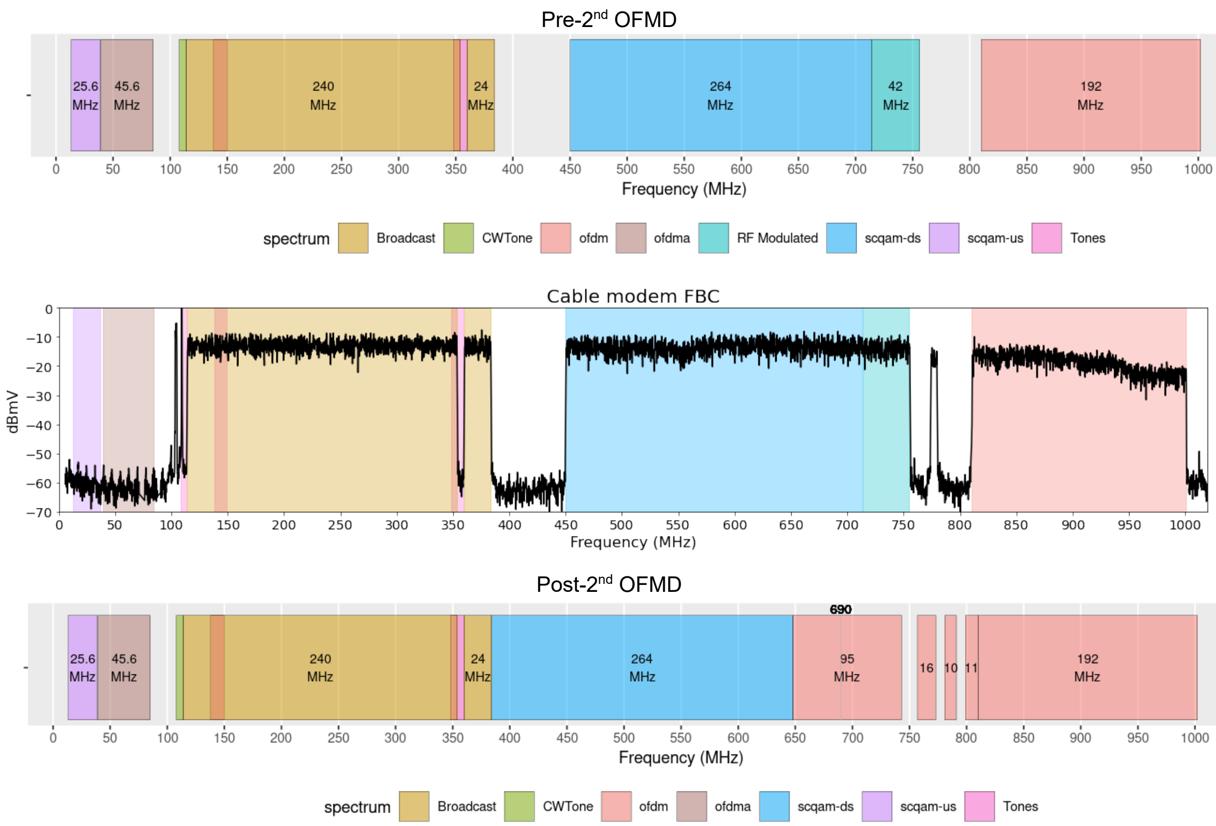

Figure 7 - Top panel: Spectrum configuration pre-2nd ofdm shows a “wasted” gap between video and downstream SC-QAM. Middle panel: an FBC from a cable modem reveals an undocumented MDU insert at 777 MHz. Bottom panel: The recommended configuration moves the SC-QAM block to the edge of the video, places a 2nd OFDM channel and creates exclusion bands per policy and for the detected insertion.

With the established logic, Figure 8 offers a view of the top 25 configurations post-deployment of the 2nd OFDM channel. The spectrum is becoming more packed. Next, we'll shift our focus to the broadcast video regions.

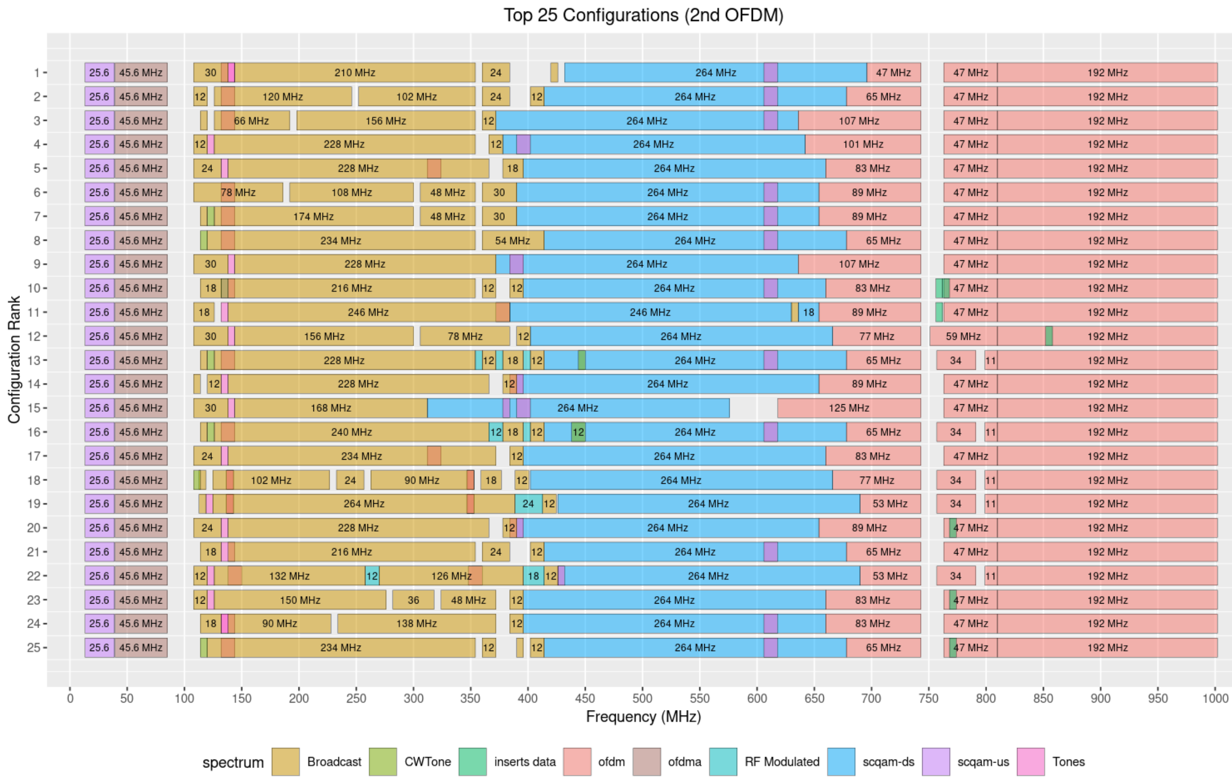

Figure 8 - Top 25 configurations for RPDs post 2nd OFDM placement.

6. Phase 3 of Activations: Close-Packing of Video

Phases 1& 2 of spectrum activation do not impact the video configuration. As seen in Figures 1-2 & 8 the video blocks are somewhat fragmented and can be made to occupy a narrower range of spectrum. But, making video use of spectrum more efficient goes beyond tighter channel packing. This is because within each QAM, several programming single program transport streams (SPTS) (TV channels) are multiplexed together in a multi-program transport stream (MPTS). Depending on the video encoding level, e.g., music, standard definition (SD), high definition (HD), of each program there may be room for moving content around such that each QAM is fully utilized.

For the ensuing discussion, to avoid confusion, we use the term "QAM" to refer to the 6 MHz frequency channel and "program" to refer to the TV channel. We present some findings based on exploratory data analysis performed on an example programming template. The first task was to examine how each video QAM fill factor varies by program encoding definition. The assumption is that a single QAM carries up to 15 SD programs, 9 HD programs, or 50 music choice programs. Figure 9 (top panel) shows the distribution of content by QAM for a given template highlighting the encoding level (fill color). The view confirms the assumed limits. However, we also see examples of QAMs having programs with a mix of different encoding levels. We also see a QAM that extends beyond the limit as it contains 50 music choice programs and 1 SD program (at ~340 MHz).

It is also assumed that programs of similar categories (e.g., Sports, Family, Movies, Latin, etc.) are typically multiplexed together. To validate this hypothesis, we assigned each program a category obtained by performing text mining on the full program name. These categories shown in the bottom panel of Figure 9 represent our classification of the type of programming. We do indeed see some clustering of similar type programs on the same QAM.

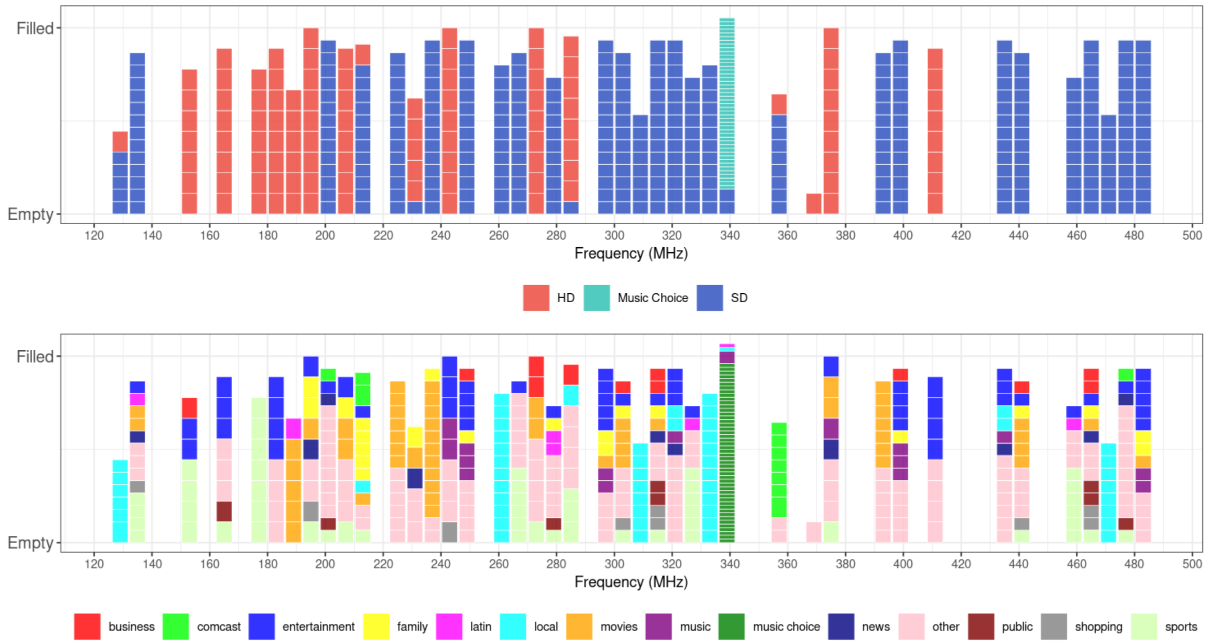

Figure 9 - Top panel: packing of video QAMs with programs in which color mapping indicates encoding level. Bottom panel: same data color-coded to indicate the program category.

We propose an algorithm for repacking video that is a “greedy” approach to relocate programs from scarcely populated QAMs to densely populated QAMs, with the preference for filling each QAM with programs of the same type of encoding, if possible. In more detail:

- Partially populated QAMs are ranked according to fill factor (least to most).
- Programs from the least populated QAMs get migrated to the most populated QAMs (preferably the same encoding level).
- The process repeats iteratively while tracking a performance metric on the overall fill factor of the template (entire video spectrum).
- Once no further gains are made, the process terminates.
- Lastly, QAM frequencies are reassigned to remove any gaps between video QAMs. This step (bottom panel of Figure 10) can be informed by the policy.

Figure 10 is a representative example of the above logic in action. The top panel represents the current state, in which the video occupies 42 QAMs. The middle panel shows the result of applying the greedy algorithm to improve the fill factor. For this template, 7 QAMs were freed, and the final lineup is almost filled, with few exceptions. The bottom panel shows the state after reassigning QAMs to frequencies such that the entire video is one contiguous block.

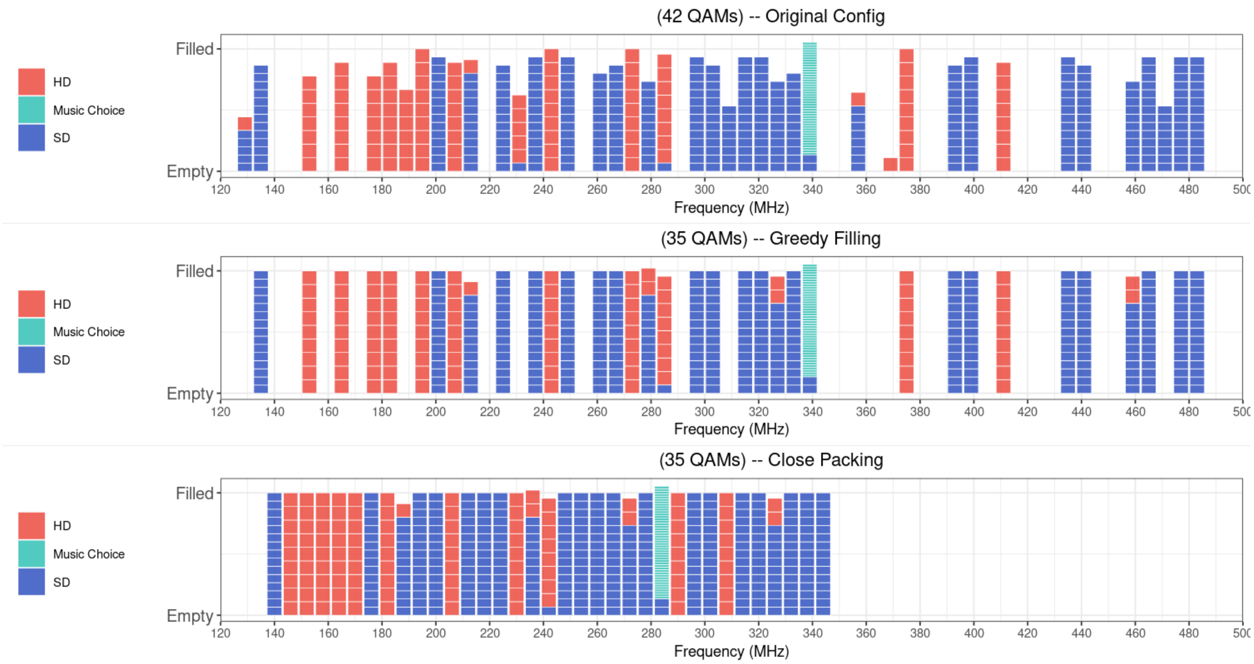

Figure 10 - Top panel: starting state in which video QAMs are neither completely filled nor closely packed. Middle panel: the view after invoking the greedy algorithm for improving the fill factor. Bottom panel: The view after closely packing the video QAMs.

This allocation and efficiency of video management is helpful in the following ways.

- Video can be consolidated to reduce the number of QAMs so that the spectrum can be used more efficiently, enabling increased DOCSIS IP (Internet Protocol) spectrum, for higher speed 10G broadband products.
- It allows us to manage the spectrum appropriately as we evolve to DOCSIS 4.0 based on the variety of broadband and set top box (STB) device capabilities.
- It allows us to easily evolve QAM Video to IP video efficiently over time as we expand the FDX spectrum.

This Machine learning pipeline is currently under development aligned with our launch of multi-Gbps symmetrical service around the time of the SCTE conference. We will report on the FDX version of this pipeline that is being prepared to calculate optimal allocations of spectrum for D3.0 modems, D3.1 modems, D4.0 modems , IP Video and QAM video spectrum collectively required at each stage of broadband speed increases that will be deployed in the next two years as we evolve to full FDX Band where the video channels will be reduced and shifted to the spectrum above the FDX band. This will result in a fully automated spectrum activation and management pipeline that we will describe in a future SCTE forum.

7. Conclusion

We presented a host of algorithms in support of activating the spectrum. In phase 1 of activations (mid-split and 810-1002 MHz), we focused on developing a pipeline for artifact detection based on an object-oriented model of the cable modem full-band capture. In phase 2, the pipeline was extended to generate the spectrum configuration for the second OFDM channel including exclusion bands informed by the discovery of local insertions. Finally, phase 3 of activation, which is currently in the development stage, is concerned with freeing up the spectrum for FDX by reducing the footprint of broadcast video. This is achieved through re-multiplexing of programs such that each video QAM is fully utilized.

This base spectrum management pipeline for activation of new spectrum will be converged with the real time spectrum management introduced in last year's SCTE program [4] and in other Comcast papers in this year's SCTE program [7] to achieve the fully automated AI-supported pipeline for deploying, activating and ongoing management of our valuable spectrum assets.

Abbreviations

AI	artificial intelligence
CPE	customer premise equipment
DOCSIS	data over cable service interface specification
EIA	Electronic Industries Association
FBC	full band capture
FDX	full duplex
Gbps	gigabits per second
GHz	gigahertz
HD	high definition
HRC	harmonic related carrier
IP	internet protocol
IUC	interval usage code
JSON	JavaScript object notation
KHz	kilohertz
MAC	media access control
MDU	multi dwelling units
MHz	megahertz
MPTS	multi-program transport stream
NCTA	National Cable & Telecommunications Association
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
PHY	physical
PLC	physical link channel
PMA	profile management application
QAM	quadrature amplitude modulation
RF	radio frequency
RPD	remote PHY device
SC-QAM	single carrier quadrature amplitude modulation
SCTE	Society of Cable Telecommunications Engineers
SD	standard definition
SPTS	single program transport streams
STB	set top box
TV	television
vCMTS	virtual cable modem termination system

Bibliography & References

1. *Roaring into the '20s with 10G*, R. Howald, R. Thompson, S. Ozer, D. Rice & L. Wolcott, T. Cloonan, R. Cloonan, J. Ulm, and J. Ariesen; NCTA technical paper, 2020.

2. *Developing the DOCSIS 4.0 Playbook for the Season of 10G*, R. Howald, J. Williams, J. Cave, O. Ekundare, and M. Petersen; NCTA technical paper, 2021.

3. *Deploying PMA-Enabled OFDMA in Mid-Split and High-Split*, M. Harb, D. Rice, K. Dugan, J. Fereirra, and R. Narayanaswamy; NCTA technical paper, 2022.

4. *Time is Ripe for D3.0 Farming: Achieving Optimal Spectral Efficiency by Allocating D3.0 Spectrum to OFDM*, M. Harb, C. Humble, S. Ozer, and D. Rice; NCTA technical paper, 2022.

5. *How Network Topology Impacts Rf Performance: A Study Powered By Graph Representation Of The Access Network*, M. Harb, K. Subramanya, R. Narayanaswamy, S. Walavalkar, and D. Rice; NCTA technical paper, 2021.

6. *Graph Algorithms and Real Time Telemetry for Intelligent Plant Operations*, B. Lutz and M. Stehman; NCTA technical paper, 2023.

7. *Towards Fully Automated HFC Spectrum Management*, J. Zhu, M. Harb, D. Rice, J. Howe, and C. Humble Insert; NCTA technical paper, 2023.

