Time is Ripe for D3.0 Farming: Achieving Optimal Spectral Efficiency by Allocating D3.0 Spectrum to OFDM

A Technical Paper prepared for SCTE by

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

Chad Humble
Senior Engineer, Next Generation Access Network
Comcast
4100 E Dry Creek Rd, Centennial, CO 80122
+1 (720) 951-8092
Chad_Humble@comcast.com

Sebnem Ozer
Senior Principal Engineer, Next Generation Access Network
Comcast
1800 Arch St, Philadelphia, PA 19103
+1 (215) 286-8890
SEBNEM_OZER@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

List of Figures

1. Introduction

Spectrum in cable networks is often statically configured to support a mix of DOCSIS® 3.0 (D3.0), single carrier-quadrature amplitude modulation (SC-QAM), and DOCSIS 3.1 (D3.1) Orthogonal Frequency Division Multiplexing (OFDM) channels. Generally, more spectrum is allocated to D3.0 based on infrequent adjustments to add additional channels, when possible, informed by utilization and modem technology trends as part of long-range capacity and capital planning processes. These capacity management processes consider many business and policy elements including supply chain, procurement, warehousing and inventory, construction permitting, and work order management for node segmentations and cable modem termination system (CMTS) configuration. These important elements can be made even more effective when enabled through software-defined virtual network functions (VNF) that remove many of the manual procedures via automation and deferral of other procedures and capital investment based on real-time channel and spectrum optimization. Automating these solutions also leads to fewer instances of human error, reduces service interruption, and increases availability.

In one illustrative example, the downstream (DS) configuration consists of 44 D3.0 channels (264 MHz) and 1 OFDM (96 MHz) channel in a node that supports 750 MHz of total hybrid fiber coax (HFC) infrastructure spectrum. A configuration that favors D3.0 from a spectrum perspective is counter-intuitive since OFDM supports higher modulation efficiencies (up to 4096-QAM) and a superior low-density parity check (LDPC) error correction algorithm. Using our Profile Management Application (PMA) VNF [1-2], which adjusts physical layer capacity in real time, Comcast is able to improve DS Mbps/MHz by > 44% on average using D3.1 technology. Yet, conventional wisdom states that expansion of the OFDM spectrum is limited by the D3.1 device penetration. Multi-Gbps products, which are enabled only by OFDM/OFDMA, can help drive this device penetration based on consumer demand.

This paper introduces a fresh perspective on optimizing the spectrum to achieve desired outcomes for product speeds, capacity, node segmentation, and cost effectiveness. The goal of the optimization exercise is to maximize spectral efficiency via reallocation of DS D3.0 channels to OFDM while accounting for any increase in utilization, which may lead to costly node augments. We show that the current configuration and CMTS load balancing place a higher burden on the OFDM spectrum, thus making the conditions ripe for spectral reallocation. This reallocation creates a material positive benefit in support of multi-Gbps products. Lastly, we present a VNF concept to farm the D3.0 spectrum and manage load balancing in an automated fashion.

2. How Downstream Spectrum is Utilized

A common spectrum configuration assigns 44 SC-QAM channels for D3.0 and a single 96 MHz D3.1 OFDM channel (equivalent spectrum to 16 SC-QAM channels). Historically, before D3.1 device penetration accelerated, the OFDM channel was thought of as a “bonus” that offers some relief to the D3.0 spectrum and allows the D3.1 cable modem (CM) devices access to exclusive speed tiers (e.g., 1 Gbps and higher). However, D3.1 device penetration has been steadily increasing, mostly by organic means (i.e., new customers opting for the “latest and greatest” cable modems with the best Wi-Fi capability). The distribution of D3.1 CM penetration across all service groups (equivalent to a downstream port on a CMTS) and hardware platforms at Comcast is shown in Figure 1. The median D3.1 penetration has recently crossed a critical level with the consequence that the traffic loads on D3.0 and D3.1 technologies in the spectrum are now comparable on average (but vary per service group). This balance means that a more thoughtful spectrum allocation strategy is now required. Such a strategy may be tailored to the specific characteristics of a service group and considers aspects that include D3.0 and OFDM channel utilization, traffic forecast, projected growth in D3.1 device penetration, and plant configuration, as well as all the nuances of the current spectrum configuration (location of video, video-on-demand, tones, local inserts, etc.).

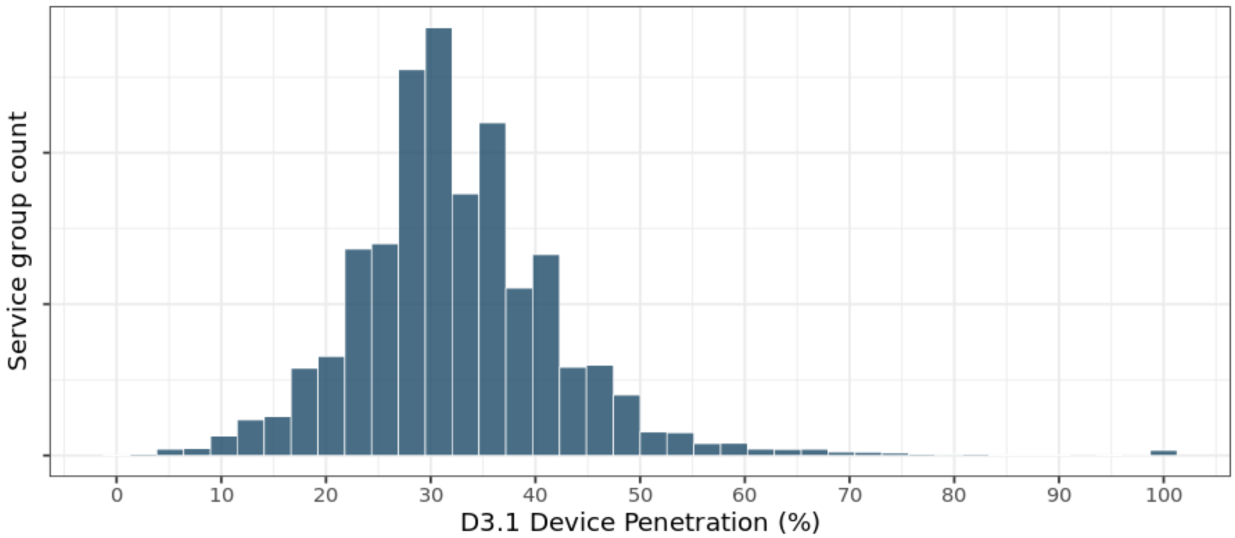

Figure 1 - Histogram of D3.1 device penetration across the network. Service group counts omitted from the y-axis on purpose as they do not inform the discussion.

Additional D3.0 distributions of interest are shown in Figure 2. Panel 1(left) confirms that a 44 SC-QAM configuration is the most common, though other variations exist (namely, 40, 36, 32, 24). Panel 2 shows the distribution of the number of DOCSIS devices per service group. Panel 3 is the distribution of service group utilization for the D3.0 part of the spectrum. Service group utilization was calculated by first averaging the point-in-time utilization (sampled at 5-min intervals) across all D3.0 channels in the service group then taking the 98^{th} percentile sample within a 30-day period as representative of the utilization. Typically, service groups that fall above the 80% utilization level are considered to be highly utilized and may require some form of remediation (it could involve, for example, “splitting the node” to create two service groups out of the original one). Finally, Panel 4 (right) shows the D3.0 channel utilization variation range (max – min) within a service group. This metric is a gauge of the CMTS internal load balancing function within the D3.0 spectrum given that each device is allocated a subset of the 44 channels to use (each unique subset is referred to as a downstream bonding group). Though not shown here, we found a large discrepancy between our primary hardware platforms in terms of D3.0 load balancing ability, with some achieving much tighter utilization spread among the SC-QAM channels than others.

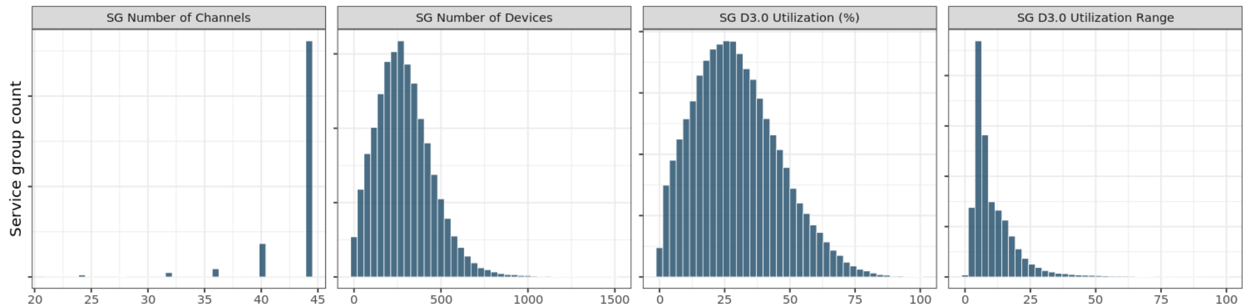

Figure 2 - Distributions of (1) number of D3.0 channels per service group, (2) number of devices per service group, (3) D3.0 service group utilization, and (4) D3.0 utilization range across channels within the service group.

What is more important for the concept of spectrum allocation is the load balancing between the D3.0 and the OFDM parts of the spectrum. The CMTS will tend to shift traffic for D3.1-capable devices heavily in favor of OFDM on some of our hardware platforms. For example, on our virtual CMTS (vCMTS) we can influence the load balancing through other VNF and micro-service opportunities. This is beneficial as it allows us to be maximally spectrally efficient, but the higher utilization of D3.1 should be anticipated and accounted for as part of the capacity planning process and policy described in the introduction. To further investigate this feature of CMTS functionality, consider the joint distribution for D3.0 and OFDM service group utilizations shown as scatter plot in Figure 3. In the ideal scenario, with evenly balanced D3.0 and OFDM, the distribution would be centered around the 45º bisector line (shown as black line). Instead, the joint distribution reveals higher OFDM utilization compared to D3.0 (orange line-of-best-fit rotated counterclockwise (CCW) relative to the 45º line). A few points regarding how this data is interpreted:

- At the current median D3.1 penetration level, balancing the D3.0 and OFDM utilizations is challenging for the CMTS owing to the limited spectrum allocated to OFDM. The imbalance is further affected by the fact that the highest speed tiers can only be accessed using D3.1-capable devices.
- More heavily used service groups are driven in large part by the OFDM portion of the spectrum, implying that the D3.0 part has excess bandwidth to yield.
- The joint distribution will continue to rotate counterclockwise (CCW) as D3.1 device penetration continues its organic growth and as CMTS scheduler management is improved, causing more service groups to cross the high utilization threshold due to high OFDM utilization.

Given all of the above, shifting some spectrum from D3.0 to OFDM is a win-win opportunity:

1. It relieves service groups that are heavily utilized in the OFDM spectrum.
2. It adds capacity to the network enabling even higher speed offerings.

Yet, this is not a one-size-fits-all situation and must be targeted within a well-considered spectrum reallocation program.

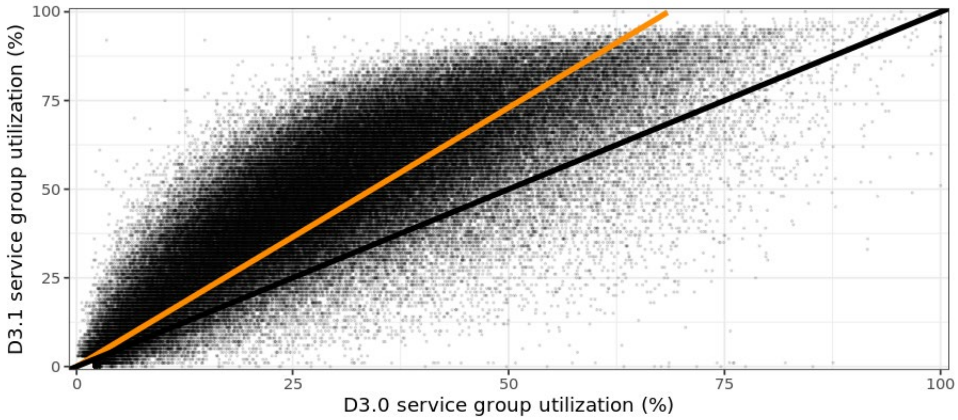

Figure 3 - Joint D3.0 and OFDM distribution shown as a scatter plot. The black line is the bisector (45 deg. angle) and the orange line is line-of-best-fit for the scatter.

3. D3.0 Spectrum Farming Approaches

This section explores two D3.0 spectrum farming approaches and estimates their impact on service group utilization and capacity. First, presented below are some guiding principles that apply across all analyses:

- It is understood that farming D3.0 spectrum presents a trade-off: overall capacity is increased because of the higher OFDM efficiency compared to SC-QAM, but the D3.0 service group utilization may also increase because of the lost D3.0 spectrum and the fact that not every device in the service groups is D3.1-capable.
- Utilization is assumed to scale linearly with the amount of added/removed spectrum, which is a reasonable assumption that applies to both SC-QAM and OFDM. Thus, this assumption is used as a recipe to project the impact of adding/removing spectrum on utilization.
- For the time being, we ignore second order effects. For example, the fact that increased capacity will enable higher speed tiers, which will in turn lead to higher traffic and increased service group utilization. We also ignore that the load balancing function may adjust given the re-allocation of available capacity between technologies. Note that such changes in speed offerings and customer behavior are tackled within the traffic forecasting and capacity planning functions.

The first approach considers farming the same number of SC-QAM channels across the entire footprint. While such an approach is not favorable from the perspective of the ultimate optimization, it is appealing from the perspective of its simplicity: maintaining a single standard configuration and avoiding development of what could be a complex spectrum management system. Estimating the impact of network wide reallocation of spectrum is straightforward. At any given number of converted SC-QAM channels, the utilizations of the D3.0 and OFDM components of a given service group are scaled proportionally to the lost or added spectrum. The efficiency assumptions of the D3.1 channel are based on years of PMA [1,2] experience across tens of millions of modems.

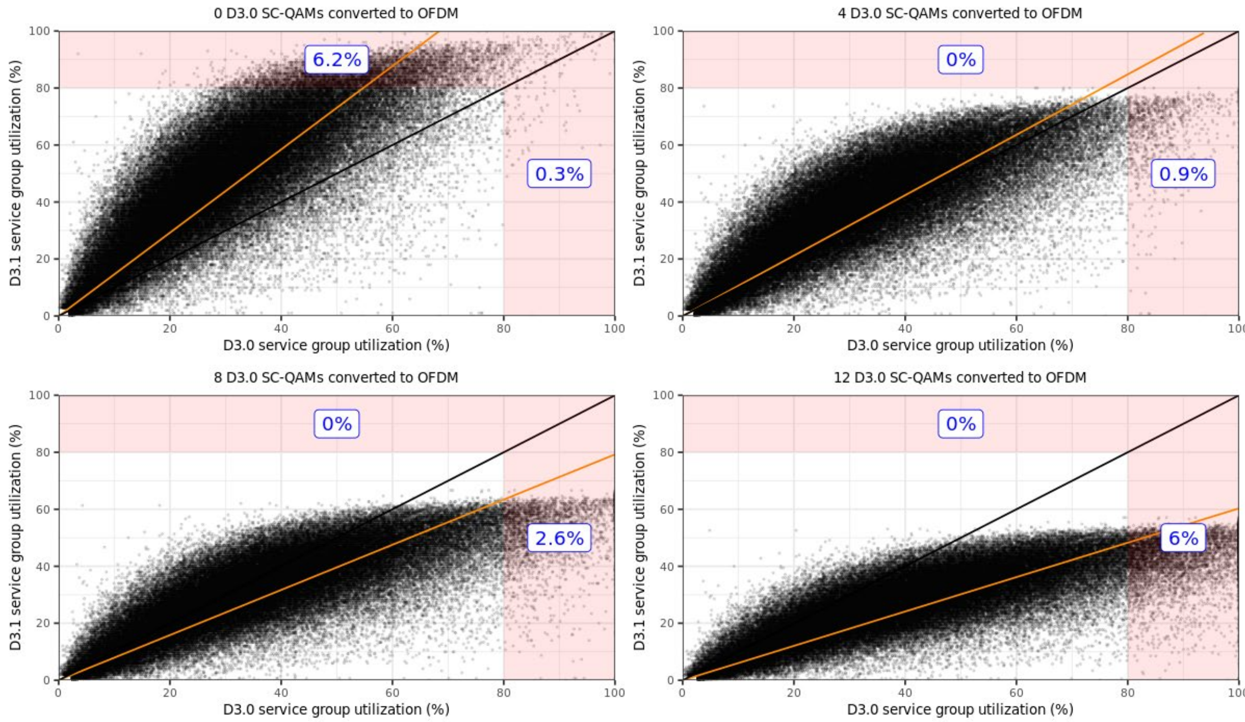

Figure 4 - Joint D3.0 and OFDM distributions at different number of SC-QAM channels converted to OFDM. (Top left) 0 channels converted—representing the baseline, (top right) 4 channels converted, (bottom left) 8 channels converted, and (bottom right) 12 channels converted.

Results of the impact analysis are shown in Figure 4. Each of the panels explores the impact of converting *n* SC-QAM channels to OFDM, with *n* being 0 (baseline), 4, 8, and 12. Highlighted on the plots are service groups that are heavily utilized (>80% utilization) falling in the shaded pink area with the percent figure annotated. It is evident that as more SC-QAM channels are converted, the joint distribution rotates clockwise (CW) as expected. At the highest level of converting 12 channels, we end up at roughly the same ratio of service groups that are heavily utilized (~6%) yet these are now driven by the D3.0 utilization rather than OFDM. Under this scenario, no relief is gained in terms of the number of service groups that require remediation as a proxy for equipment and labor capital for adding network capacity. However, ~200 Mbps of capacity was added to each service group in the network (OFDM is estimated to be 16.6 Mbps more efficient than D3.0 per converted 6 MHz channel—this estimate is based on current PMA performance where OFDM is deployed). Also, based on the results in Figure 4, one can argue that converting 4 SC-QAM channels is an easy decision: it alleviates almost all utilization issues while adding a modest ~66 Mbps of capacity to each service group in the network.

To support multi-Gbps downstream product speeds, additional Mbps are needed than can be accommodated with the baseline configuration scenario. This speed requirement is particularly challenging with HFC design spectrum limits such as 750 MHz or 860 MHz node deployments. The only way to accommodate the capacity needed by a multi-Gbps product speed is to re-allocate more OFDM spectrum by reducing D3.0 spectrum in these nodes or via a labor-intensive spectrum upgrade.

The second approach manages spectrum on an individualized service group level. The idea is to convert any number of channels such that:

- The SC-QAM block is not reduced beyond 24 channels as this number corresponds to the maximum number of D3.0 channels to which a D3.0 device is able to bond. Thus, a 44 SC-QAM block will have a maximum of 16 channels to convert.
- The projected D3.0 service group utilization does not exceed 60%, which is a level that allows a safety buffer of several years compounded annual growth rate (CAGR) of 20% before exceeding 85%. Similar models have been developed based on peak period available bandwidth targets.

The analysis shown in Figure 5 explores this approach when applied to service groups that are currently configured with 44 SC-QAM channels. It can be seen from Panel 1 (left) that most service groups are able to release 16 channels for conversion without driving D3.0 utilization above the 60% level. For those service groups that are already utilized above 60%, no channels are converted. Other service groups fall between the 2 extreme cases and are able to convert some number of channels between 0 and 16. Comparison between the current D3.0 utilization (Panel 2) and the projected D3.0 utilization (Panel 3) shows a stark difference. The projected distribution does not exhibit a normal shape as it is artificially constructed from different sub-populations; most notable, the sub-population that was driven to the 60% limit. Lastly, Panel 4 reveals the extent of added capacity: most service groups increased their capacity from just under 2.5 Gbps to above 2.8 Gbps.

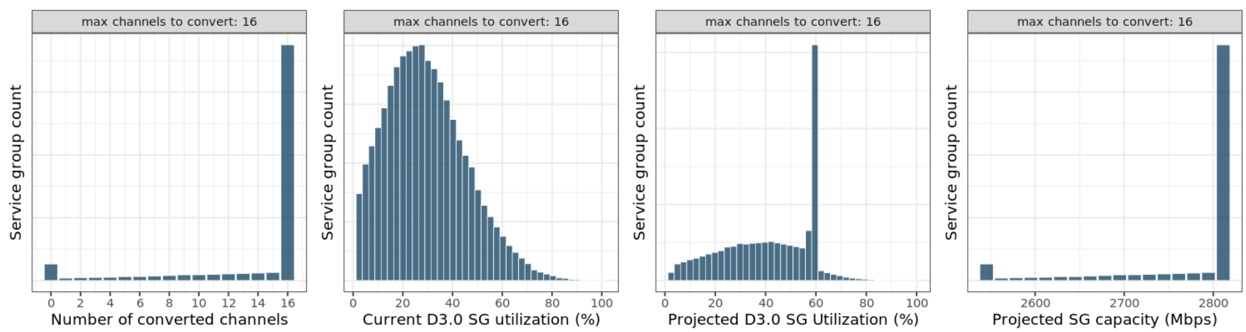

Figure 5 - Analysis exploring converting up to 16 SC-QAM channels to OFDM. The number of converted channels is specific to each service group and selected such that projected D3.0 service group utilization does not exceed 60%.

4. View of the State of the Spectrum

The concepts discussed so far present the problem of spectrum farming in an idealized fashion whether the approach is based on a completely dynamic policy optimizing capital investment lifetime or is based on policy constraints. It sufficiently covers the core dimensions pertaining to D3.1 CM penetration, service group utilization, and service group capacity. However, it ignores some important practical aspects to spectrum farming. One of these foundational pragmatic requirements is the challenge around having high confidence visibility into current spectrum allocations on every node. For instance, the Comcast network includes a variety of different hardware platforms, different plant configurations, and different spectrum allocations across localities – e.g., placement and number of video quadrature amplitude modulation (QAM) channels, DOCSIS channels, local video insertions, amplifier automatic gain control (AGC) signals, leakage markers, etc. In a 1 GHz plant, it might not be required to convert SC-QAM channels as the configuration is expected to contain enough vacant spectrum to deploy multiple increased bandwidth OFDM channels. Yet, even in this scenario, it is necessary to understand how the spectrum is currently utilized and the kind of housekeeping that is needed to free up contiguous spectrum blocks. More importantly, striving for a standardized spectrum configuration is a goal in its own because standardization simplifies both day-to-day operations before software management can take over and deployment of future technologies (e.g., full duplex (FDX), and high split). Given these dynamics, it becomes obvious that automation is a prerequisite to any effort directed towards:

- Optimizing capacity in HFC network segments with limited spectrum by converting D3.0 channels to OFDM.
- Freeing up contiguous blocks of spectrum to deploy OFDM where plant configuration allows it.
- Relocating specific channels in order to move towards a steady state with a few standardized configurations across the network.

Presented in this section are early explorations into concepts that are foundational to the spectrum management tool. The first exercise involved gaining a complete picture of the state of the spectrum configuration across the network. The initial phase of data collection was limited to the vCMTS platform. Data from various sources were merged to provide a unified view on the downstream configuration for every remote PHY device (RPD), which is used in this paper as synonymous to a node, managed by the vCMTS platform. This data analysis task confirmed the lack of standardized configuration as it revealed that ~700 distinct permutations exist for the tens of thousands deployed RPDs. Figure 6 shows three example configurations that represent a small part of the diversity of what is deployed in the field. These examples fall within the top 25 configurations, which covers most of the nodes, creating a "long tail" in the distribution of custom DS channel configurations. Notice the sharp contrast between the three configurations. The first (Rank 1 in RPD count) has a block of vacant spectrum above 750 MHz ready to be used for OFDM expansion. The second (Rank 2) has a lot of vacant spectra but in a more fragmented way, with 2 SC-QAM blocks surrounding the OFDM channel. The third (Rank 21) has vacant spectrum in the expansion region as well but except for a 6 MHz "local insert" channel at ~850 MHz. The local insert could be thought of as a video channel combined locally with the cable signal to serve a specific use (e.g., a security camera system within a hospital or a housing development). This type of spectrum use presents a problem for OFDM expansion. Yet it is often the case that these types of exceptions are not easy to trace since the configuration found in the data is registered for an entire site (large set of RPDs) even if the insert is physically available only on a single RPD or on a node leg. This situation highlights the importance of adding a validation layer on top of the picture that was constructed solely from various configuration data sources.

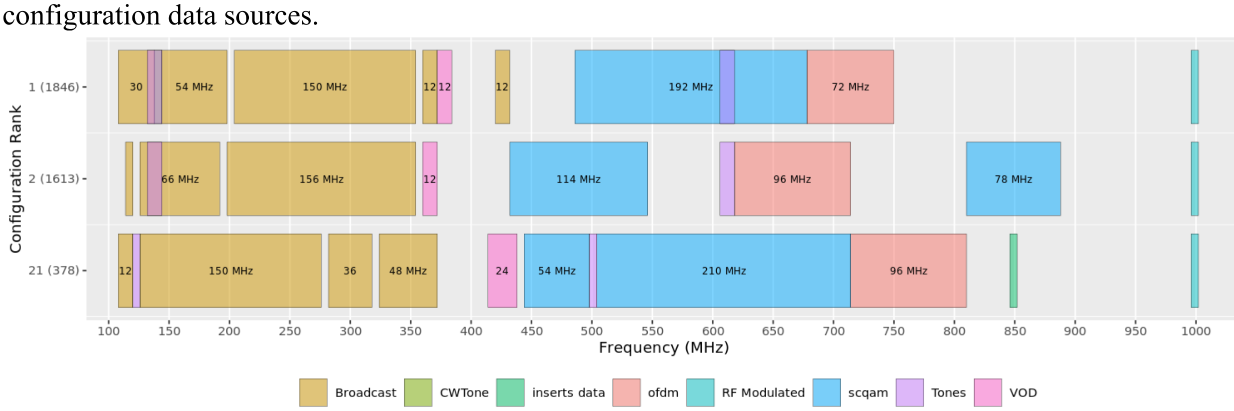

Figure 6 - Example of three common downstream spectrum configurations.

A validation methodology was adopted that is based on using the full band spectral capture from CM gateways to infer whether a range of spectrum is occupied or vacant. Briefly:

- Raw spectra (117 KHz resolution bandwidth) are captured from every DOCSIS device on the to-be-validated nodes.
- Channel power is calculated from raw spectra for the standardized 6 MHz EIA channel plan.

- A threshold is adopted for determining whether the channel is in use or vacant. For the results presented here, the threshold was set to -40 dBmV.
- Device-channel level data is aggregated to the RPD-channel level and to the configuration-channel level.

The outcome of the vacancy detection process is shown in Figure 7 in the form of a dark gray band overlaid on the configuration map. It is seen that the vacancies detection algorithm performs well as the spectrum gaps neatly correspond with documented vacancies. There are few exceptions. Tones for amplifier AGC, local modulator, or leakage usually overlap with inferred vacancies since these are much narrower than the 6 MHz channel. Furthermore, leakage tones are allowed to be placed within a DOCSIS block, so they should be ignored. What is more interesting is the overlap between the local insert at 850 MHz (Configuration 21) and the inferred vacancy. Since this configuration is shared by 378 RPDs, the local insert may be present in only a small subset of these or even just a single smaller multi-dwelling unit (MDU) on a node that also serves other customers. A deep dive into the device and RPD level data reveals 2 interesting findings shown in Figure 8:

- The local insert was confirmed on one RPD out of the 378 that share the same configuration. 100% of devices on the RPD exhibited energy above the noise level on that RPD.
- Energy was picked up by 50% of devices on a different RPD on a channel that does not correspond to any known use in the configuration (at 885 MHz). Further investigation revealed that the signal represented ingress from a recently installed police repeater.

The first case requires field work to move the inserted channel to a different location and certify that the RPD can proceed with OFDM expansion. The second case does not require any further action since ingress is more appropriately handled by customized OFDM profiles within the PMA system along with plant maintenance dispatch for downstream ingress remediation.

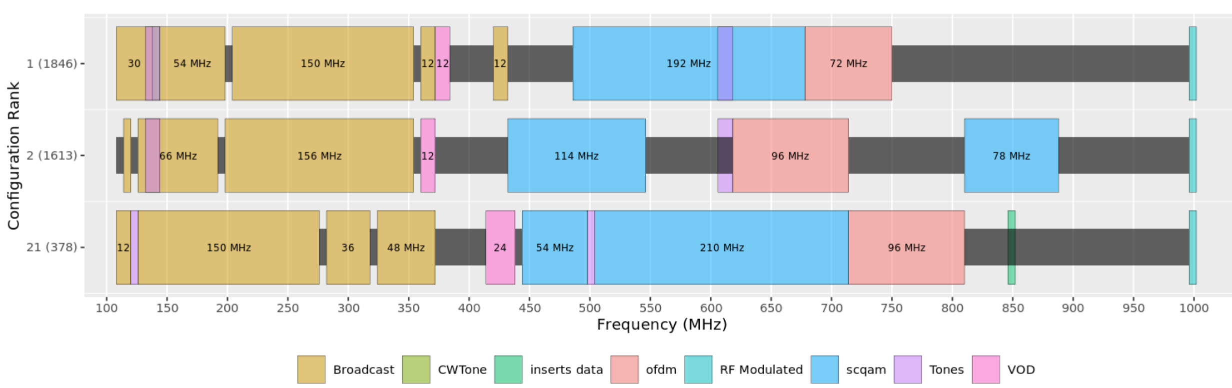

Figure 7 - The same configuration shows in the previous figure with the outcome of vacancy detection overlaid as dark gray band.

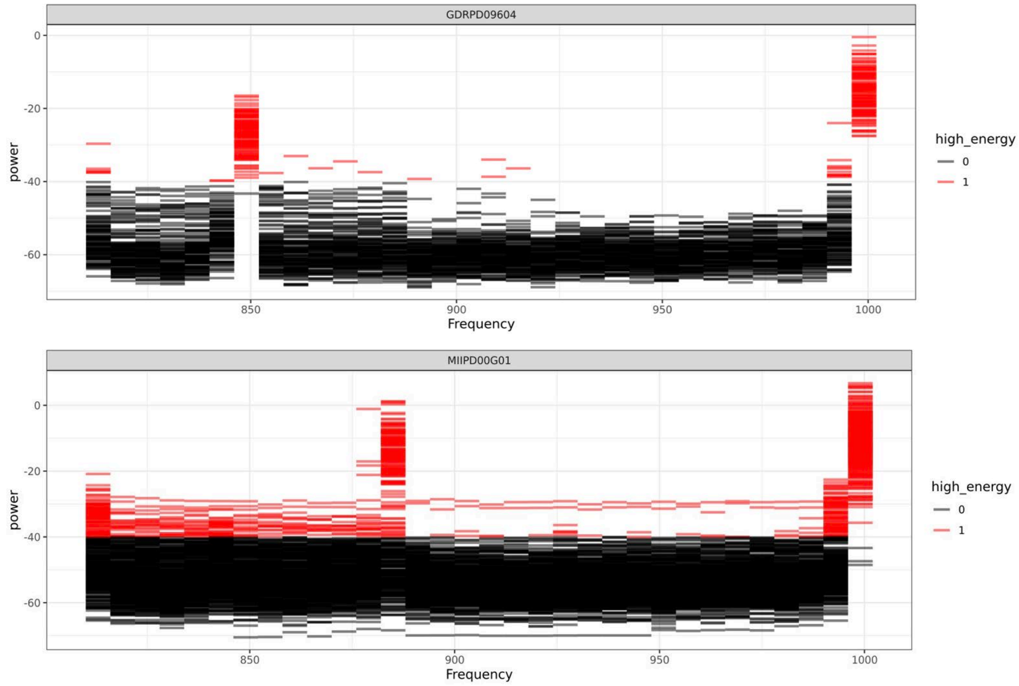

Figure 8 - Example of two "conflicts" picked up by the vacancy detection algorithm: (1) a local insert was confirmed at 850 MHz, and (2) ingress due to police repeater was picked up at 885 MHz on a different RPD.

These early explorations show that we now have a methodology for capturing and validating the state of the spectrum configuration. The subsequent steps involve:

1. Automating the pipeline for building and validating the spectrum configuration for every node in the network (both integrated CMTS (iCMTS) and vCMTS platforms).
2. Implementing an algorithm for moving spectrum around to free up contiguous spectrum blocks for OFDM expansion.
3. Implementing the algorithm for farming D3.0 spectrum where plant configuration restricts use above 750 MHz.
4. Developing a layer that translates the output from both algorithms into the CMTS-specific configuration.
5. Ensuring visibility into real-time spectrum allocations and the technology used in that spectrum is available to all stakeholders. This includes the capacity planning tools and process to ensure optimized investment in network expansion only where and when it is needed.

5. 2 Gpbs Testing

In this section, we introduce some of the preliminary speed test results for a 2 Gbps downstream service to demonstrate how additional OFDM spectrum could be used for higher broadband product speeds. The potential of higher speeds is one of the primary motivations behind the D3.0 farming development. The downstream spectrum configuration for these experiments was 28 SC-QAM channels & a single 192 MHz OFDM on the iCMTS and 44 SC-QAM channels and a single 192 MHz OFDM on the vCMTS.

Figure 9 shows that a single D3.1 cable modem can reach 2.3-2.4 Gbps data throughput over iCMTS and vCMTS systems when there are no congestion or radio frequency (RF) noise issues.

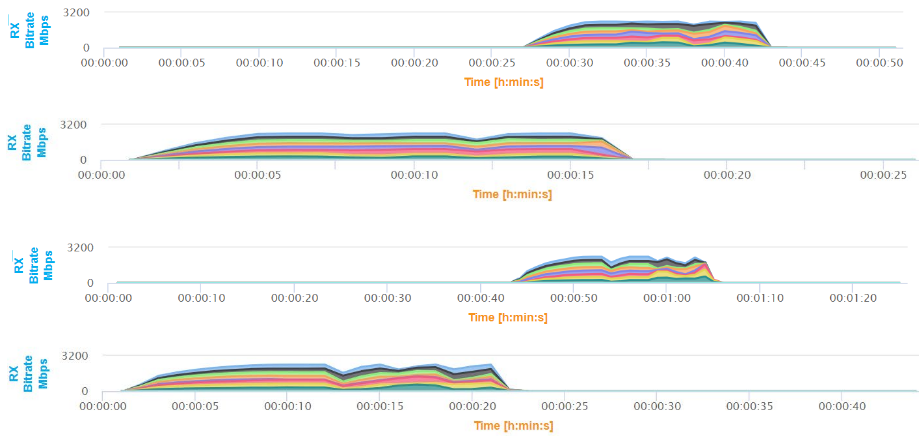

Figure 9 - The data throughput over iCMTS (top 2 panels) and vCMTS (bottom 2 panels) for a single D3.1 cable modem demonstrating speeds in the 2.3-2.4 Gbps range. The different colors indicate different TCP flows utilized by the speed test.

The serving group utilization and RF noise issues may affect the maximum achievable rate at any given time. Figure 10 displays an example for a D3.1 cable modem in a service group with different utilization levels measured within a 15 second window. The results are used to analyze the average and peak utilization levels of the service group to confirm the service group availability for the 2 Gbps speed tier and to drive service group upgrade strategy accordingly. The graphs on the left correspond to relatively high utilization levels and correspondingly impact 2 Gbps speed test results compared to the graphs on the right in which the service group utilization is low. This is analyzed further in this section with more examples to show the impact of utilization burstiness and transmission control protocol (TCP) behavior on supporting higher downstream rates.

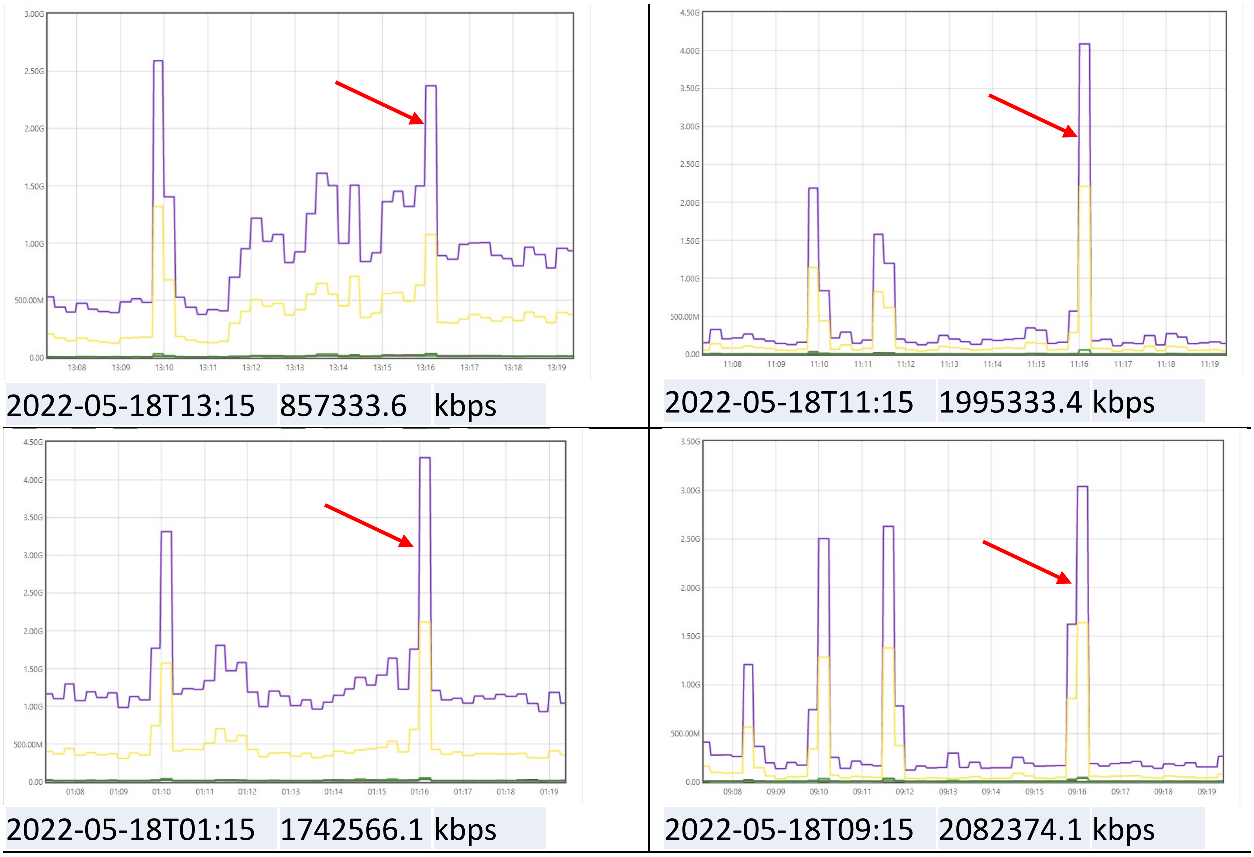

Figure 10 - Speed test results for a 2 Gbps speed tier. There are no RF issues observed for this cable modem and corresponding node. However, utilization levels during the speed test have an impact on the results. Purple lines are total octets at service group level while yellow lines are total octets for the OFDM channel. The load telemetry data is displayed with a delay compared to speed test timing displayed in the bottom of each graph and indicated by a red arrow.

With the increased speed tiers and emerging applications, new characteristics have been observed for the ratio of peak-to-average traffic load and peak duration per subscriber and per serving group. As discussed in [3], the higher service tiers have higher burst rates. Although traffic load averaged over longer time periods such as 5 minutes may be low, bursts observed within smaller time windows in the order of seconds can be 10-20 times higher for the current pre-FDX rates. For Gbps rates, microbursts can be observed within windows on the order of microseconds. Microbursts may happen due to service characteristics, network segments with different rates and classic TCP protocol response during congestion [4,5].

Traditionally, utilization in each channel at a CMTS is reported per 5-minute interval. New telemetry models can provide utilization values at 15 second intervals. Measuring traffic utilization and burstiness at smaller window sizes may not be feasible for large scale systems. However, this information is crucial in network analysis. As discussed in [6], bursts and their overlap by multiple users can be infrequent and these short-term bursts may have little impact on aggregate values, but the spikes may have a significant impact on services like gaming, videoconferencing, and augmented reality/ virtual reality (AR/VR) applications. Internet service providers (ISPs) must develop effective measurement and prediction systems to project granular measurements at small scales to aggregated scales.

An example is shown for an iCMTS with the same settings described for Figure 9. A speed test is performed for a subscriber with 2 Gbps DS speed tier rate, between 7 and 22 seconds as shown in Figure 11. Initially, the serving group load is close to 0.9 Gbps before the speed test starts. After the speed test starts, a set of bursts occurs between 12 and 18 seconds and affects the speed test outcome as the service group is highly utilized during this timeframe. Figure 12 displays the total average load measured at iCMTS. The measurements are collected via a command line interface (CLI) command to get the bit rates for each DS channel. One set of the measurements provides the load averaged over a 30 second window while the other set provides the load averaged over a 1 second window, which is a more accurate representation of the burst that impacts the speed test as shown in Figure 13.

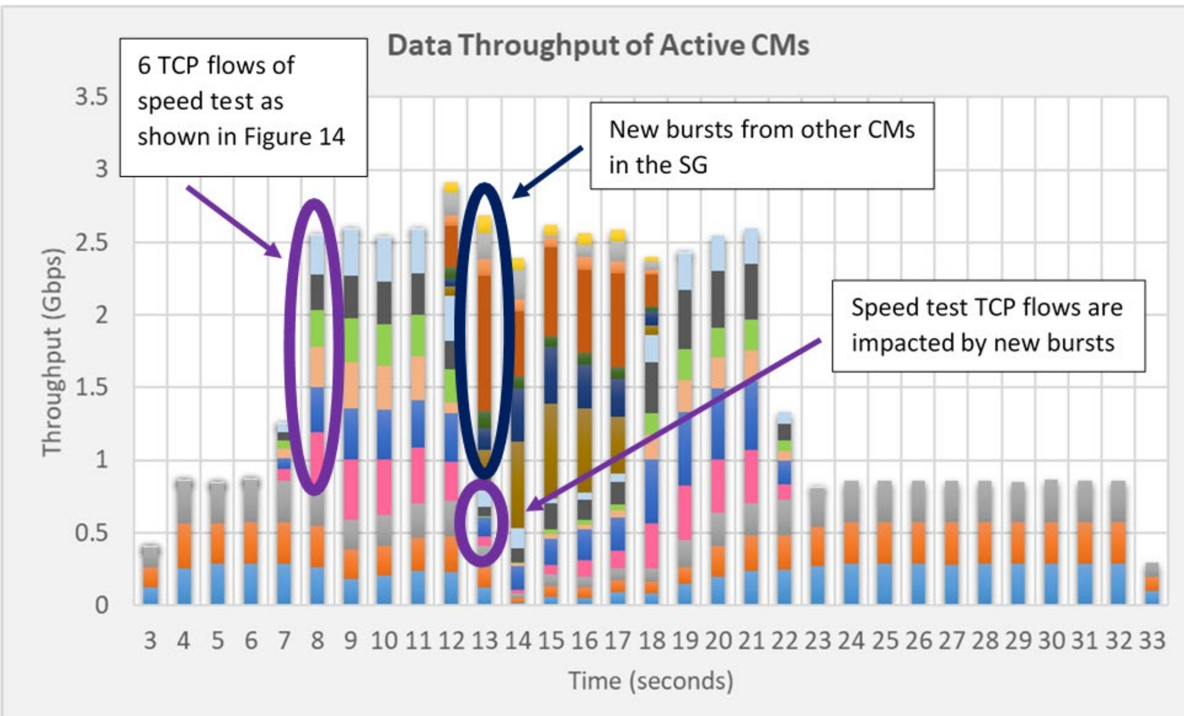

Figure 11 - Total data load computed at the traffic generator's customer premises equipment (CPE) ports connected to the CMs in the serving group. The speed test runs between 7-22 seconds.

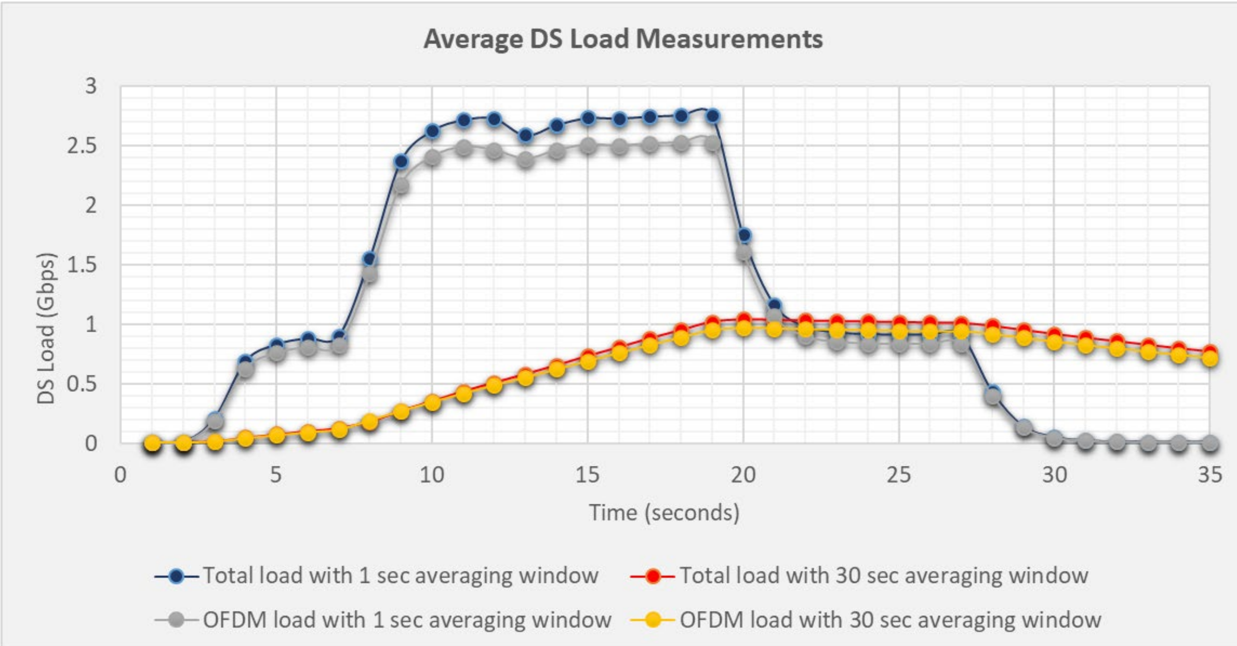

Figure 12 - Total average load measured at the iCMTS by summing 30 sec and 1 sec channel loads in the serving group.

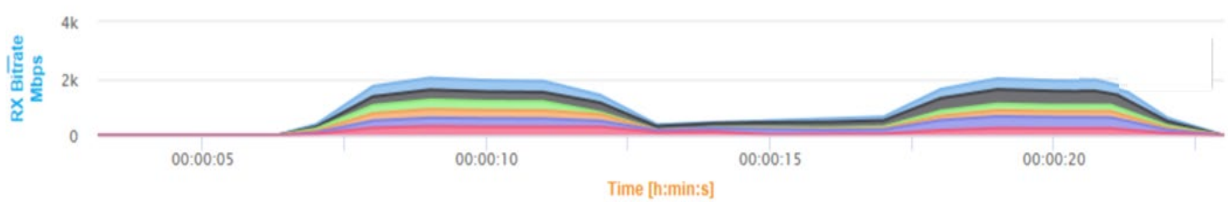

Figure 13 - Speed test result from the traffic generator GUI.

This example shows that utilization and other metrics must be analyzed at more granular levels with increased speed tier rates. Burst level analysis is crucial in network planning, cost analysis, quality of experience (QoE), and speed assurance as well as optimization of network functionalities for efficiency and fairness.

6. Conclusion

We presented our perspective on reallocating D3.0 spectrum to OFDM, motivated by the need to increase service group capacity to support higher speed tier offerings. We conclude that:

- The growth in D3.1 device penetration has now reached a level where reconsidering how downstream spectrum is allocated between D3.0 and OFDM is necessary.
- The objective of reallocation is twofold: increase service group capacity to enable higher speed offering and provide relief to a currently highly utilized OFDM spectrum.

- To this end, two approaches were presented: a one-size fits all farming of D3.0 channels, vs. a customized recommendation that considers service group characteristics.
- The latter, due to its dynamic nature, requires building a spectrum management tool; a cornerstone of such a tool is building a unified source of truth for the current state of the spectrum configuration.
- Such a tool also feeds into additional spectrum housekeeping efforts that are targeted towards freeing up space for expanded OFDM region in support of enabling 2 Gbps and higher downstream speeds.
- These spectrum management and farming VNFs will be critical to managing spectrum effectively as the future evolves towards DOCSIS 4.0 and multi-Gbps symmetrical services.
- Initial speed test results reveal the importance of increasing service group capacity via an approach that considers the speed test result sensitivity to utilization and the impact on TCP flow control behavior when testing for these very high speed tiers.

Abbreviations

AGC	Automatic gain control
AR	Augmented reality
CCW	counterclockwise
CLI	Command line interface
CM	cable modem
CW	clockwise
CMTS	cable modem termination system
CPE	customer premise equipment
D3.0	DOCSIS 3.0
D3.1	DOCSIS 3.1
dBmV	Decibels relative to 1 milliVolt
DOCSIS	Data Over Cable Service Interface Specifications
DS	downstream
EIA	Energy Information Administration
FDX	full duplex
Gbps	Gigabits per second
GHz	GigaHertz
HFC	hybrid fiber coax
iCMTS	integrated cable modem termination system
ISP	Internet Service Provider
KHz	KiloHertz
LDPC	Low density parity check
Mbps	Megabits per second
MDU	multi-dwelling unit
MHz	MegaHertz
OFDM	orthogonal frequency division multiplexing
OFDMA	Orthogonal frequency division multiple access
PMA	profile management application
QAM	quadrature amplitude modulation
QoE	quality of experience
QoS	quality of service
RF	radio frequency
RPD	remote PHY device

SC-QAM	single carrier-quadrature amplitude modulation
TCP	Transmission Control Protocol
vCMTS	virtual cable modem termination system
VNF	Virtual network function
VR	Virtual reality

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. *Full Scale Deployment of PMA*, M. Harb, J. Ferreira, D. Rice, B. Santangelo, NCTA technical paper, 2020.
3. *Traffic Engineering in a Fiber Deep Gigabit World*, John Ulm & Tom Cloonan, SCTE-NCTA Workshop, 2021.
4. https://support.huawei.com/enterprise/en/doc/EDOC1100086962
5. IETF L4S Architecture Draft: https://datatracker.ietf.org/doc/draft-ietf-tsvwg-l4s-arch/
6. *Roaring into the '20s with 10G*, Robert Howald, Robert Thompson, Sebnem Ozer, Daniel Rice, Larry Wolcott, Tom Cloonan, Ruth Cloonan, John Ulm, Jan Ariesen, SCTE-NCTA Workshop, 2020.
7. https://github.com/BroadbandForum/obudpst

