Towards Fully Automated HFC Spectrum Management

A Technical Paper prepared for SCTE by

Jay Zhu
Principal Software Engineer, Access Network Engineering
Comcast
4100 E Dry Creek Rd, Centennial, CO 80122
+1 (720) 955-6764
Jay_Zhu@comcast.com

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

Jeff Howe
Director, Access Network Engineering
Comcast
+1 (630) 988-9060
Jeff_Howe@comcast.com

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

Dan Rice
VP2, 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

1. Introduction

In 2022, Comcast introduced a fresh perspective on optimizing the spectrum to achieve desired outcomes for product speeds, capacity, node segmentation, and cost-effectiveness by maximizing spectral efficiency via reallocation of data over cable service interface specification (DOCSIS®) 3.0 (D3.0) single carrier-quadrature amplitude modulation (SC-QAM) channels to DOCSIS 3.1 (D3.1) orthogonal frequency division multiplexing (OFDM) channels while accounting for the utilization distribution [1]. Comcast also presented a virtual network function (VNF) concept to manage the balance automatically and dynamically between the SC-QAM and OFDM spectrum.

As the penetration growth in D3.1 devices continues and the deployment of D4.0 devices comes to the horizon, such a VNF solution can ensure our network resources are utilized at their optimal efficiency and robustness during the transformation. At a high level, a requirement of this VNF solution is to consider each individual service group's characteristics and constraints at scale and produce tailor-made recommendations as inputs into the automation layer where configuration change procedures, error handling and failover logic are implemented. A unified source of truth for the current state of the spectrum configuration is also maintained to support the VNF operations in the larger picture of the architecture.

Based on these requirements, late last year, we continued to make progress on this initiative and designed and developed algorithms to effectively produce spectrum configuration recommendations based on the given real-world constraints and objectives. Using the analytical engine of the VNF service with baseline constraints, we estimated the average capacity gain that can be achieved across our distributed access architecture (DAA) footprint is 231 Mbps per service group at the time of measurement, accounting for strict constraints and without adding any new spectrum. This capacity increase is enabled by the ~44% capacity gain from our profile management application (PMA) VNF [2][3] while the SC-QAM spectrum is converted into OFDM spectrum when the constraints are satisfied.

To demonstrate and test our work, we built an automation layer with a configuration manager and a closed-loop system in our lab that supports dynamic radio frequency (RF) spectrum changes using our virtual cable modem termination system (vCMTS) – remote physical layer (PHY) device (RPD) configuration application programming interfaces (APIs) without disconnecting the cable modems (CMs) on the service group. This demonstrates the ability of the system to make hit-less RF spectrum changes in an integration context. We also measured the peak throughput changes of the DOCSIS 3.1 CMs after the hit-less spectrum changes.

2. Problem Formulation and Algorithm Design

In the last year's paper, converting a static number of SC-QAM channels into OFDM spectrum across the entire footprint was discussed and analyzed with 4 scenarios, where converting 4 SC-QAM channels could add ~66 Mbps of capacity to each service group in the network, and converting 12 SC-QAM channels could add ~200 Mbps capacity to each service group, with potential increases in SC-QAM utilization on a subset of service groups.

This one-size-fits-all experiment is an effective starting point and could achieve significant gains with system simplicity. As we progress with the VNF implementation, the solution can now be further improved and scaled to fit each individual service group's characteristics such as their spectrum constraints and the most recent 98th percentile utilization values on SC-QAM and OFDM spectrum, without having to search for an optimal configuration that could fit the entire footprint.

2.1. Real-World Constraints

When designing the VNF, there are a few general real-world policy-imposed constraints the VNF's algorithm needs to account for as part of its design. For example,

Table 1 – Real-World Policy-Based Constraints

Attribute	Common Value	Description
SC-QAM/OFDM utilization upper limit	60%	The 98th percentile utilization on both SC-QAM and OFDM have their upper limits, such as 60% to ensure high service quality. This means that the algorithm must ensure the estimated SC-QAM utilization and OFDM utilization will not exceed the given limit after the optimization/rebalance. This threshold can account for some years of compound annual growth rate (CAGR) which could influence this threshold to avoid increases in consumption growth that would require adding back SCQAM spectrum. While these removals and re-additions have been proven out, more conservative thresholds could be chosen until a fully automated system exists.
RF tones		Existing tones such as automatic gain control (AGC) tones, local RF inserts, and video on demand (VoD) tones must be correctly handled for each service group. For example, analog AGC tones should not overlap with DOCSIS spectrum but may be "worked around" by adding exclusion zones in OFDM channels; digital AGC tones may require the presence of neighboring power within a certain frequency range to the tone for the best performance.
Exclusion zones		Known exclusion zones must be considered in the optimization. For example, if a frequency range is known to suffer from ingress impairments, the algorithm should ensure that the SC-QAM channels do not overlap with the ingress.
Minimum number of SC-QAM channels	24	A minimum number of SC-QAM channels, for example, 24 or 32 channels can be required to achieve maximum product offers on D3.0 modems even if the overall capacity would accommodate a channel reduction. This constraint may also vary on a per-service group basis.

Attribute	Common Value	Description
SC-QAM channel size	6 MHz	The allocation of SC-QAM channels must use a minimum unit size.
OFDM channel size limits	24 – 192 MHz	The OFDM channels have size limits. The DOCSIS 3.1 CM physical layer specification [4] allows an OFDM channel to be 24 megahertz (MHz) to 192 MHz wide, but this range could vary based on deployment requirements.
OFDM channel guard band size	1 MHz	The OFDM channels' guard band overhead should be accounted for by the algorithm. Such a requirement allows for optimization accuracy and more meaningful reasoning for the recommendations when necessary.
SC-QAM applicable frequency ranges		The SC-QAM channels can have preferred frequency ranges in deployment.
Contiguous spectrum allocation requirements		The SC-QAM channels can require minimum allocation sizes, for example, 4 contiguous channels per allocation block.
Spectrum ordering		The SC-QAM channels may be preferred to be allocated to the lower frequency end, relative to the OFDM spectrum.
Hardware limitations		Hardware limitations such as the RPD chipset's capabilities can affect the maximum supported number of certain types of channels.

The listed policy-based constraints support the definition of the problem space for the design of the algorithm.

2.2. Algorithm

To build an effective VNF for this optimization task, it is desired to design an algorithm that:

- Adjusts its operation based on simple policy attributes that allow for A/B testing to optimize without writing new code.
- Produces recommendations that are guaranteed to satisfy all given policy-based constraints if feasible solutions exist.
- Is simple to maintain and add/remove policies for different use case requirements.
- Produces optimal spectrum configurations based on the given constraints if feasible solutions exist.
- Is performant, can recompute optimal spectrum allocation plans for our entire footprint efficiently for performance and elastic cloud costs.

With the given complexity of the listed constraints, there are multiple options on how the algorithm can be implemented. In our design, this optimization problem is formulated as a mixed-integer linear programming (MILP) problem, and the constraints are modeled as linear inequality constraints. Using MILP results in an explainable model where the descriptive objective function and constraints can be used as the immediate representation of the algorithm, reducing the gap between a discrete, programmatic implementation and a theoretical design. In addition, such a model can have maintainability benefits when the policy, constraints and rules are modified, added, or removed.

2.2.1. Objective Function and Formulation of Linear Constraints

First, a definition of contiguous spectrum segment is created to represent a frequency range that can be freely used by DOCSIS downstream channels. With this definition, constraints can be contained within each contiguous spectrum segment and be simplified. Constraints 2 and 3 in Table 1 become preprocessing rules for identifying contiguous spectrum segments.

Because the objective of this optimization can be defined as maximizing the total spectral efficiency (linearly related to the total capacity). The objective function is defined as:

$$\max \sum_{i=1}^{N} x_i + \alpha \times y_i'$$

Figure 1 – The Objective Function

where N is the number of contiguous spectrum segments that can be reallocated and each of the segments can vary in their frequency ranges, x_i is the target allocation of SC-QAM spectrum, y_i is the target allocation of OFDM spectrum (with guard bands subtracted), α is the spectral efficiency ratio of OFDM / SC-QAM which suggested by the overall downstream capacity gain (~44%) observed on OFDM channels.

For each individual contiguous spectrum segment, the total allocated spectrum must be smaller than or equal to the size (W) of the segment:

$$x_i + y_i + K \times n_i \leq W_i$$

Figure 2 – Total Spectrum Segment Constraint

where x_i is the target allocation of SC-QAM, y_i is the target allocation of OFDM (with guard bands subtracted), K is the total OFDM guard band size per OFDM channel, n_i is the number of OFDM channels (to take OFDM's guard band overhead into consideration), and W_i is the size of the i_{th} contiguous spectrum segment.

The allocated SC-QAM within each continuous spectrum segment should be divisible by the size of an SC-QAM channel, hence:

$$x_i = p_i \times R$$

Figure 3 – SC-QAM Allocation Minimum Unit Constraint

where p_i is an integer variable that represents the number of SC-QAM channels, and R is the size of an SC-QAM channel.

Considering constraint 8 in Table 1, the allocated SC-QAM within a continuous spectrum segment may also subject to an upper limit due to the maximum frequency an SC-QAM channel can be allocated:

$$x_i \leq \eta_i$$

Figure 4 – SC-QAM Allocation Frequency Limit Constraint

where η_i is calculated by subtracting the continuous spectrum segment's starting frequency from the maximum frequency that an SC-QAM channel can be allocated.

We use a binary constraint to represent the presence of OFDM and take OFDM's guard band overheads into account:

$$c_i \times M \geq y_i$$

$$y_i + (1 - c_i) \times M \geq P - K$$

Figure 5 – OFDM Allocation Constraint

where c_i is a binary variable that indicates the OFDM presence in a contiguous spectrum segment, y_i is the target OFDM allocation (with guard bands subtracted), P is the minimum size of an OFDM channel (with guard bands included, e.g.: 24 MHz), K is the OFDM guard band overhead (both guard bands), M is a large enough number to ensure:

- When OFDM is present, y_i is greater than 0, c_i must be 1 to satisfy the first inequality, and in the second constraint the OFDM size must meet its minimum size requirement.
- When OFDM is not present, y_i equals to 0, the first constraint is always satisfied, c_i becomes 0 because of minimization of overheads, and because M is large, the second constraint is always satisfied and the OFDM size variable of the contiguous spectrum segment does not need to meet the minimum size requirement.

Similarly, we use a binary constraint to represent the presence of SC-QAM and ensure that enough number of SC-QAM channels is allocated sequentially within a contiguous spectrum segment:

$$b_i \times M \geq x_i$$

$$x_i + (1 - b_i) \times M \geq B$$

Figure 6 – SC-QAM Allocation Constraint

where b_i is a binary variable that indicates the SC-QAM presence in the contiguous spectrum segment, x_i is the target SC-QAM allocation, B is the minimum required allocation for SC-QAM channels within a contiguous segment, M is a large enough number to ensure:

- When SC-QAM is present, x_i is greater than 0, b_i must be 1 to satisfy the first inequality, and in the second constraint the SC-QAM allocation must satisfy the minimum SC-QAM allocation constraint within a contiguous spectrum segment.
- When SC-QAM is not present, x_i equals to 0, the first constraint is always satisfied, and because M is large, the second constraint is always satisfied.

In addition to the previous constraints, the number (minimum) of OFDM channels must be associated with the target OFDM allocation:

$$n_i \times (Q - K) \geq y_i$$

Figure 7 – Minimum Number of OFDM Channels Constraint

where y_i is the target allocation of OFDM (with guard-bands subtracted), K is the total OFDM guard band size used by each OFDM channel, n_i is the number of OFDM channels, and Q is the maximum size of an OFDM channel (with guard bands included, e.g.: 192 MHz). This constraint allows the minimum number of OFDM channels to be known as n_i and used in the constraint described by Figure 2 – Total Spectrum Segment Constraint to account for OFDM guard band overhead in the optimization.

For the minimum total SC-QAM allocation requirement, such as 24 or 32 SC-QAM channels, the following constraint is defined:

$$\sum_{i=1}^{N} x_i \geq \Gamma$$

Figure 8 – Minimum Total SC-QAM Allocation Constraint

where $\mathbf{\Gamma}$ is the minimum required SC-QAM allocation.

Finally, utilization constraints are specified to provide the upper limits of estimated utilization values after optimization. Respectively, there are SC-QAM utilization constraint and OFDM utilization constraint (i.e.: the utilization of SC-QAM/OFDM after optimization must not exceed 60% allowing 2 years of growth at a conservative CAGR before service group augmentation. This policy can be made more aggressive as full automation is completed):

$$\sum_{i=1}^{N} x_i * \Theta \geq L * \beta$$

$$\sum_{i=1}^{N} y_i * \Lambda \geq H * \gamma$$

Figure 9 – SC-QAM and OFDM Maximum Utilization Constraint

where L is the original SC-QAM allocation, β is the original SC-QAM utilization, Θ is the upper limit of SC-QAM utilization after optimization, H is the original OFDM allocation, γ is the original OFDM utilization, and ƛ is the upper limit of the OFDM utilization after optimization.

2.2.2. Performance

It is known that complex MILP problems can consume a significant amount of time to solve. A common implementation of solving a MILP problem is to programmatically generate the objective function and constraint representations of the problem and pass them to a dedicated MILP solver for high performance computing. Many of the solvers implement various heuristic algorithms to reduce the problem-solving time. The performance of a solver depends on several factors, including but not limited to:

- The number of variables in the constraints. Usually, the solving time increases super linearly as the number of variables grows.

- The objective function and constraints. Different solver implementations can perform differently on the same problem. Often, a solver is not consistently slow or fast on all problems.
- The implementation of the solver. There are open-source solvers and commercial solvers, either could perform well depending on the optimization problem and the heuristics and tuning in place. Commercial solvers often have advantages in scalability and solving time, but it is not a guarantee.

Because one of the requirements of the spectrum management VNF is to process the service group data and produce spectrum allocation recommendations for the entire footprint efficiently, there is high expectation on its algorithm's performance. As a result, given the current design of the constraints, the algorithm spends ~5ms on average to compute an optimal DOCSIS downstream spectrum plan for a service group on a single CPU core, which infers that it is possible to compute spectrum plans for the entire footprint within minutes using limited resource. Note that this performance result was evaluated using an open-source MILP solver library. Given this performance result, we have not investigated the performance of commercial solvers.

2.2.3. Testing

In addition to unit tests, manual validations, and lab tests, we developed fuzzing tests to reinforce the robustness of the algorithm and ensure that expected outputs are produced with a wide range of random given inputs and constraints.

The fuzzing tests cover a 1 GHz spectrum reserved for random initial SC-QAM and OFDM allocations and utilizations. And the tests check the optimization output under the following combinations of conditions:

Table 2 – Fuzzing Conditions for the Algorithm

Condition	Value
Randomized minimum SC-QAM allocation requirement	true + false
With random insertion tones	true + false
Allow OFDM AGC overlapping	true + false
Randomized initial OFDM size	24 – 192 MHz (0.5 MHz steps)

For each fuzzing condition, 2000 random test cases are generated. This helps automatically cover potential edge cases of the constraints as well as the validity of the optimization output under all conditions.

3. System Requirements and Architecture

The spectrum management VNF requires interfacing capabilities with other systems and services to retrieve the existing spectrum configurations and information of constraints. The VNF should also generate recommendation tasks for the automation layer to perform configuration management.

The VNF's core algorithm is developed as a production-ready microservice that has optimized performance and ideal CPU and memory efficiency. It is not directly interacted with by the users/engineers. Instead, we built a proof-of-concept (PoC) application with a graphical user interface (GUI) as a playground for convenient experiments, testing and demonstrations during development. This PoC application implements all required APIs to communicate with the VNF and data APIs for retrieving the service group data and applying optimized spectrum configurations. It also provides a summary dashboard and data tables and visualizations for quick optimization result analysis.

3.1. PoC Software

The PoC software has a simple architecture to allow the users to interact with its GUI, parallelize optimization requests to the spectrum management VNF, and store constraint settings and spectrum allocation data and statistics for before and after optimization.

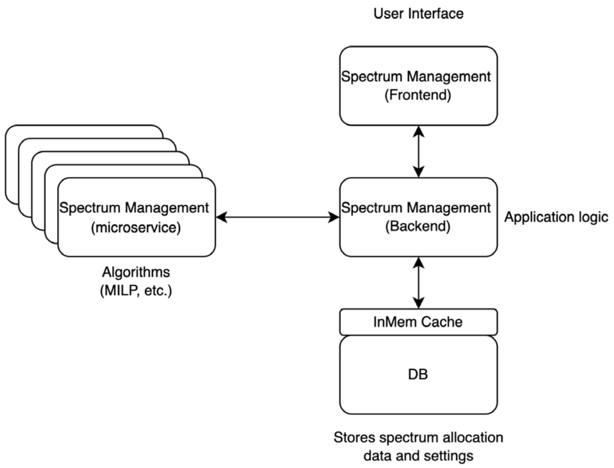

Figure 10 – PoC Software Architecture

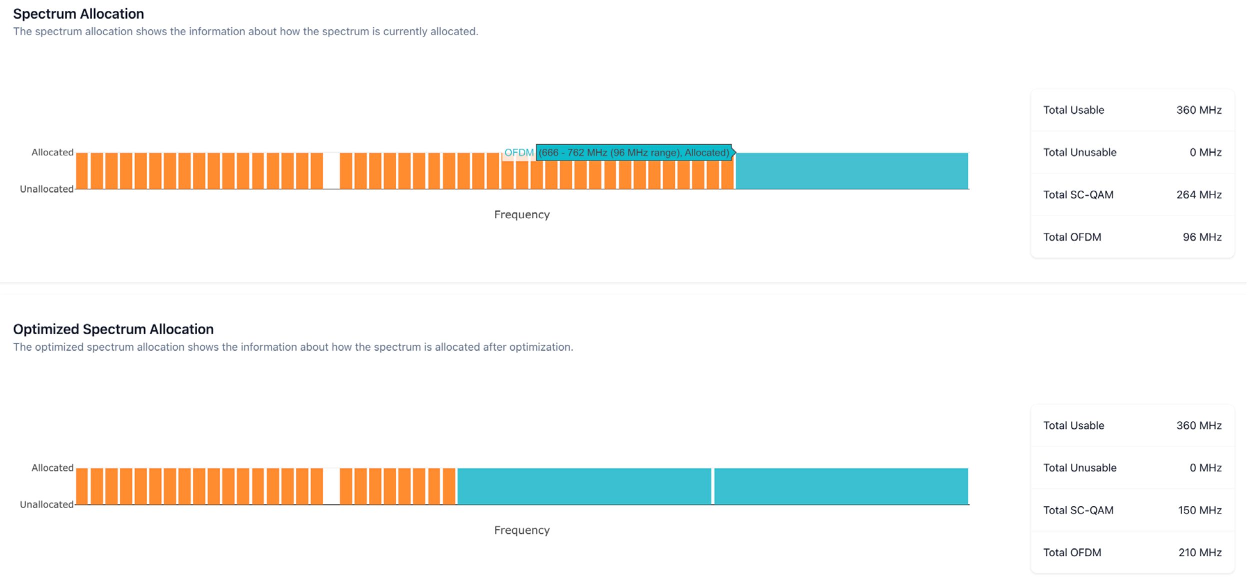

Spectrum Allocation

The spectrum allocation shows the information about how the spectrum is currently allocated.

Total Usable	360 MHz
Total Unusable	0 MHz
Total SC-QAM	264 MHz
Total OFDM	96 MHz

Optimized Spectrum Allocation

The optimized spectrum allocation shows the information about how the spectrum is allocated after optimization.

Total Usable	360 MHz
Total Unusable	0 MHz
Total SC-QAM	150 MHz
Total OFDM	210 MHz

Figure 11 – PoC Software: Per Service Group Spectrum View

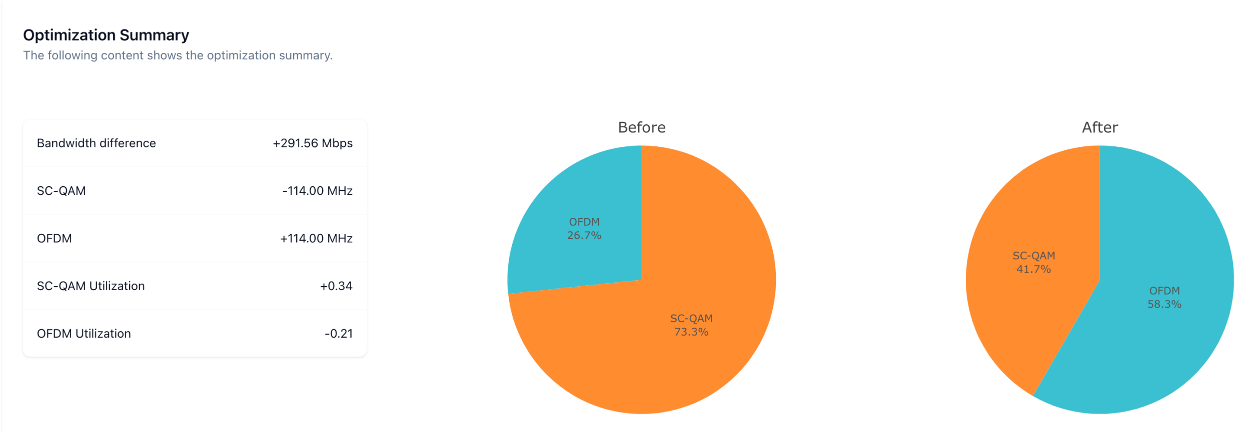

Optimization Summary

The following content shows the optimization summary.

Bandwidth difference	+291.56 Mbps
SC-QAM	-114.00 MHz
OFDM	+114.00 MHz
SC-QAM Utilization	+0.34
OFDM Utilization	-0.21

Figure 12 – PoC Software: Per Service Group Optimization Summary

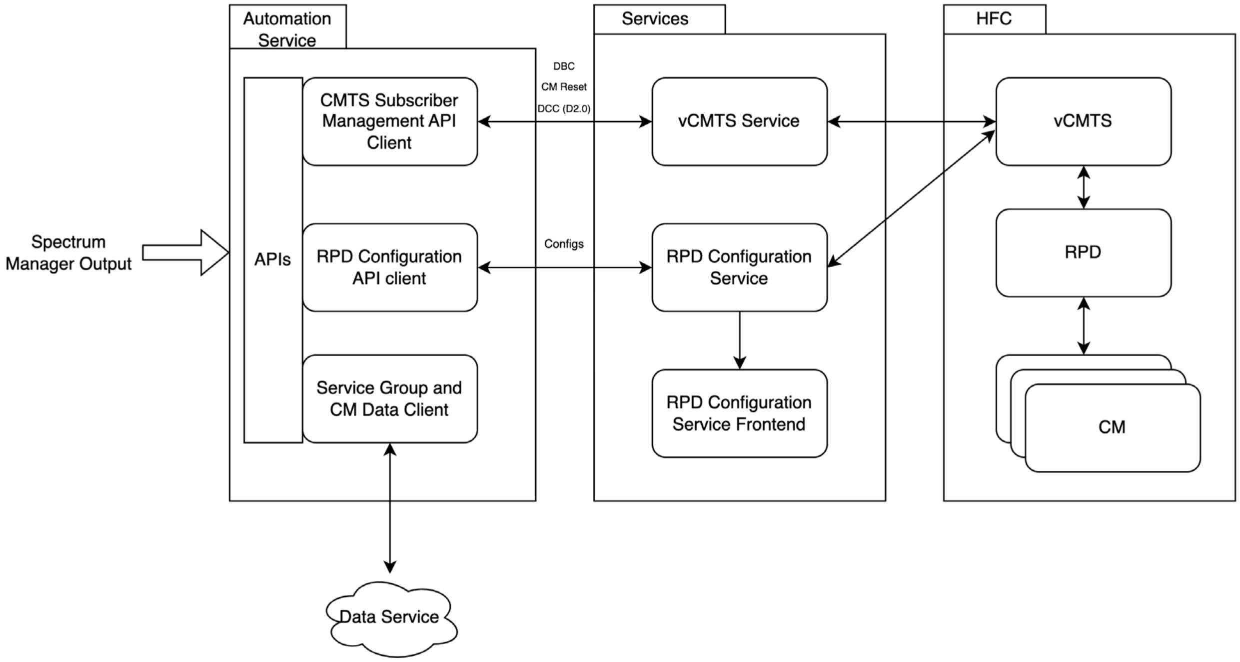

Figure 13 – PoC Software: Automation and Integration Service

A separate service illustrated in Figure 13 is developed as the automation layer for handling hit-less spectrum configuration changes, simple load balancing, and service group data retrieval for the PoC application. The PoC application's backend communicates with such service to extend its functionalities. More details of the automation service are discussed in section 5.

3.2. Integration Context

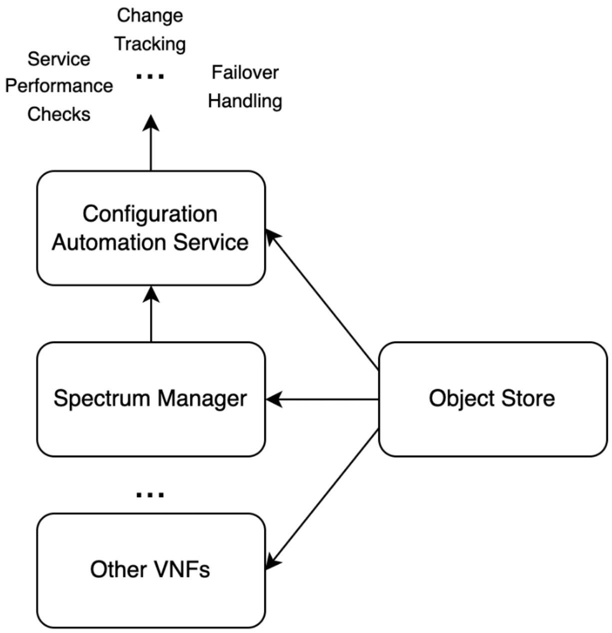

Figure 14 – Spectrum Manager: High-Level Deployment Integration

In the deployment integration context, the spectrum management VNF, along with other VNFs, are part of our event driven architecture centered around a unified object data store. The concept of this architecture is illustrated at a high-level in Figure 14. This architecture provides all required data for the computation performed by the spectrum management VNF and automation pipelines for scheduling and applying the configuration changes. The automation service illustrated in Figure 13 can be used as the initial implementation of part of the automation pipeline for spectrum management in this integration context.

4. Experiments and Results

To evaluate the performance and effectiveness of the VNF, we took a snapshot of the service group data of our entire DAA footprint and experimented with different constraint parameters. In all our experiments, the OFDM's spectral efficiency factor was configured to 1.44 which means 44% capacity gain based on RF compared to SC-QAM and OFDM channels that use 256 quadrature amplitude modulation (QAM) order. In addition, the experiments focused on rebalancing the existing configured SC-QAM and OFDM spectrum, which means no additional DOCSIS spectrum was added to the configurations after the optimization.

4.1. Using the Baseline Configuration

The baseline configuration is a set of conservative parameters. The upper limit of the SC-QAM and OFDM spectrum are set to 60% respectively. The minimum required SC-QAM channels is 24, which allows for up to 20 SC-QAM channels worth of spectrum to be converted into OFDM spectrum if the starting point is 44 SC-QAM channels. The minimum allocation block size of SC-QAM channels is 24 MHz (4 x 6 MHz channels). And finally, the minimum required OFDM channel size is 48 MHz.

Table 3 – Spectrum Allocation Optimization: Baseline Parameters

Parameter Name	Value
Maximum allowed SC-QAM utilization	60%
Maximum allowed OFDM utilization	60%
SC-QAM channel size	6 MHz
Base bandwidth per 6 MHz	37 Mbps
Minimum number of SC-QAM channels	24
Minimum SC-QAM allocation block	4 x 6 MHz = 24 MHz
OFDM guard band size (1 guard band)	1 MHz
OFDM size limit	48 – 192 MHz

Using the baseline parameters as constraints in the computation, it was estimated to improve each service group's capacity by 231 Mbps on average. Among all the service groups that were computed by the optimization algorithm, 94% of which received feasible solutions with improvements. The remaining 6% of the service groups were labeled "infeasible" because the solver could not find a solution that satisfies all constraints (such as the utilization limit constraints). The majority of the infeasible cases resulted from the high utilization on both SC-QAM and OFDM and the minimum SC-QAM allocation requirement (24 channels).

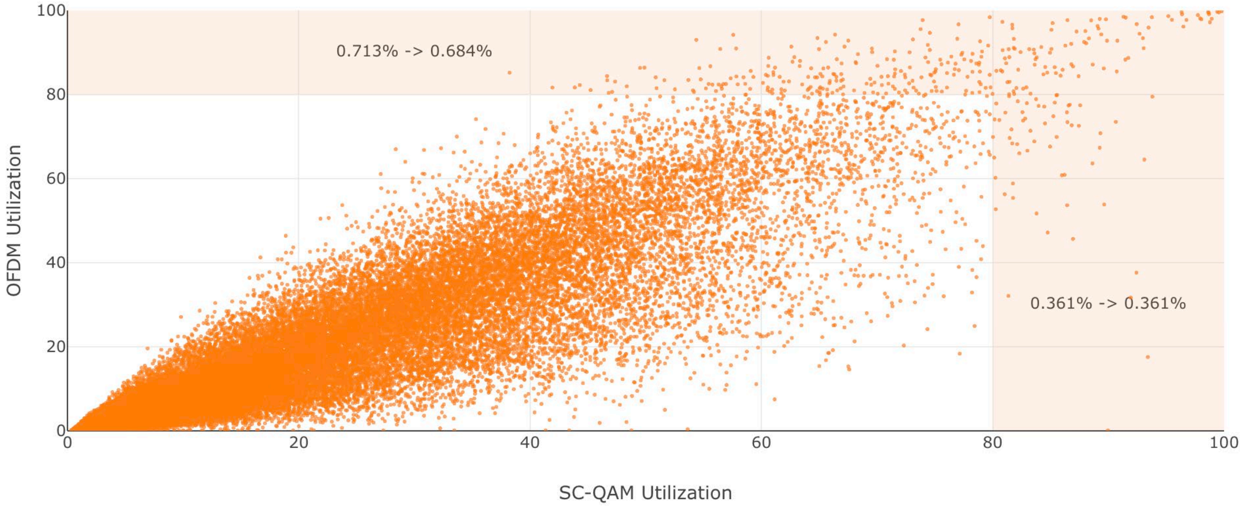

Figure 15 – Joint SC-QAM and OFDM Utilization Distributions (current)

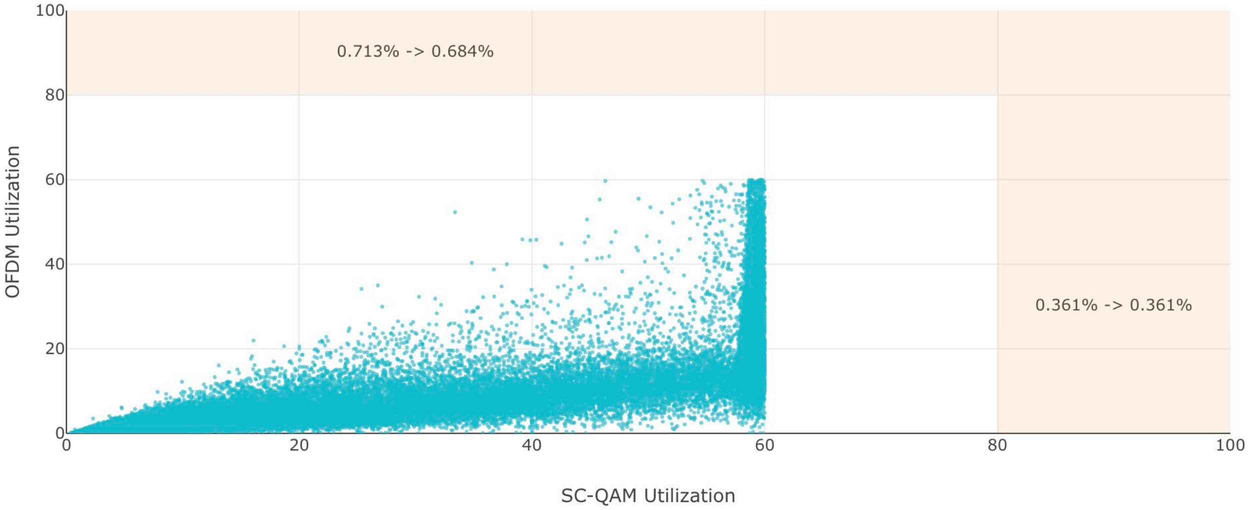

Figure 16 – Joint SC-QAM and OFDM Utilization Distributions (optimized)

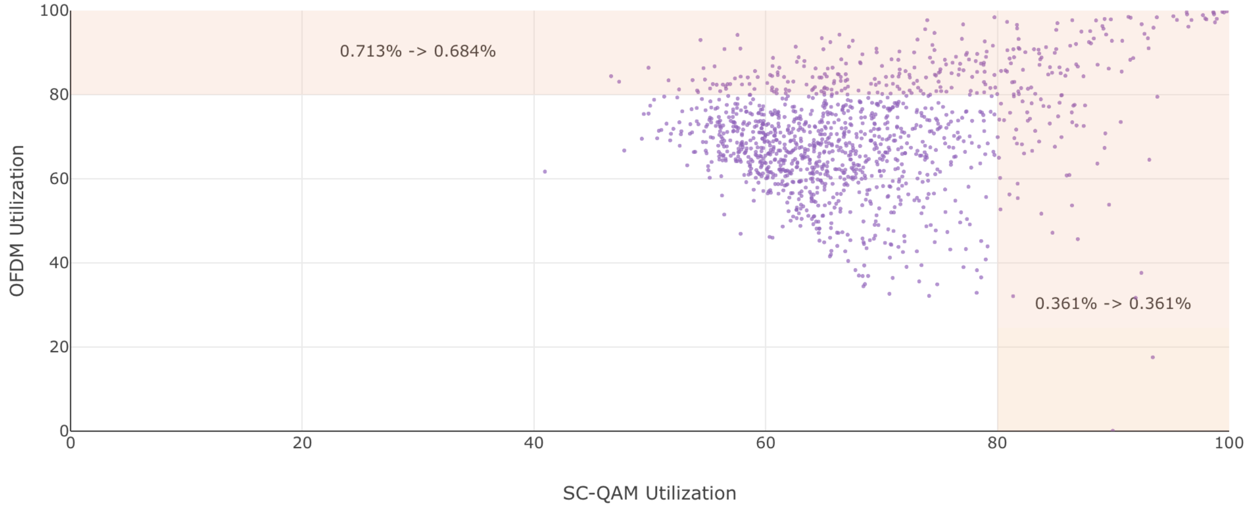

Figure 17 – Joint SC-QAM and OFDM Utilization Distributions (infeasible)

After the optimization, 0.029% of the service groups were relieved from high OFDM utilization (>= 80%) and had both SC-QAM and OFDM utilization contained within 60% by estimation. In Figure 16, the distribution shifted clockwise due to the lower OFDM utilization values, and the estimated utilization values on both SC-QAM and OFDM spectrum were capped at 60%. For most of the “infeasible” service groups, Figure 17 indicates that the existing high utilization is the cause.

4.2. Reducing the Minimum SC-QAM Allocation Requirement

The optimization using baseline constraint parameters produced promising results. However, the minimum number of SC-QAM channels (24) could be a limiting factor for the overall capacity gain. As the network transformation continues and newer technologies grow in adoption, we could expect this requirement to reduce over time.

In this experiment, all parameters were the same as the previous experiment, except the minimum number of SC-QAM channels was reduced from 24 to 16. This policy could be reduced as a result of a reduction in the maximum product speed offered to a D3.0 modem.

Table 4 – Spectrum Allocation Optimization: Reduced Minimum SC-QAM Allocation

Parameter Name	Value
Maximum allowed SC-QAM utilization	60%
Maximum allowed OFDM utilization	60%
SC-QAM channel size	6 MHz
Base bandwidth per 6 MHz	37 Mbps
Minimum number of SC-QAM channels	**16**
Minimum SC-QAM allocation block	4 x 6 MHz = 24 MHz
OFDM guard band size (1 guard band)	1 MHz
OFDM size limit	48 – 192 MHz

This constraint change increased the estimated average capacity gain from 231 Mbps to 321 Mbps.

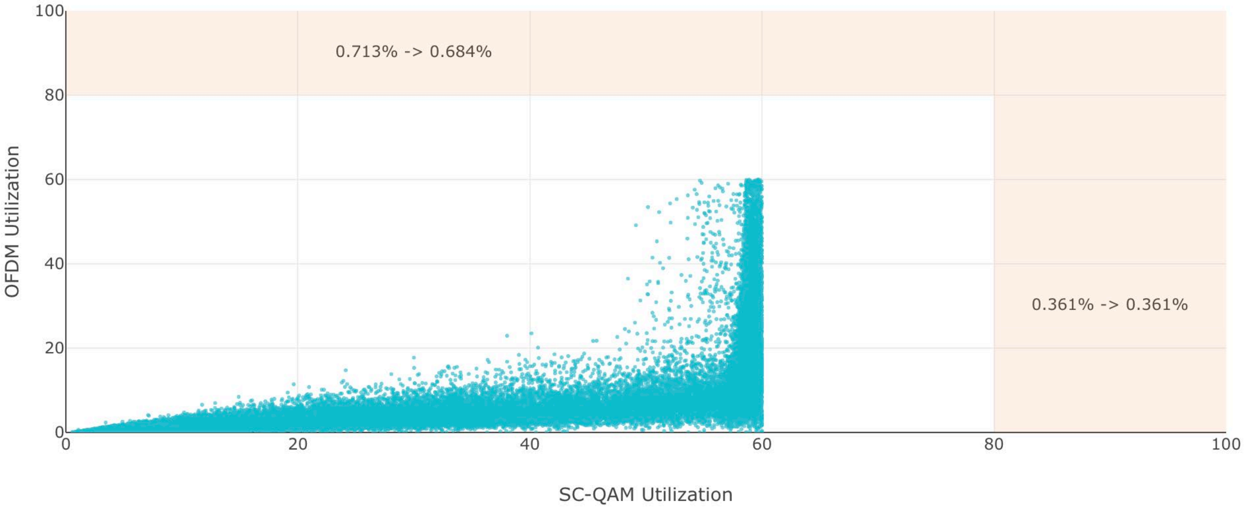

Figure 18 – Joint SC-QAM and OFDM Utilization Distributions (optimized with a minimum of 16 SC-QAM channels required)

From the optimization summary, there are still 6% of service groups that did not have feasible optimized spectrum configurations, but a significant amount of service groups could further reduce the SC-QAM spectrum within the allowed utilization limit to bring higher spectral efficiency and capacity gain.

4.3. Increasing the SC-QAM Utilization Limit

From a capacity management perspective, the full-scale deployment of the spectrum management VNF can potentially combine the SC-QAM utilization metric with the OFDM utilization metric and allow us to consider the DOCSIS spectrum as an integral spectrum resource. With this assumption, it is possible to reconsider the utilization constraints. The 60% SC-QAM utilization limit in the baseline configuration can be increased without risking overly high utilization on the SC-QAM spectrum. Meanwhile, this provides more room for SC-QAM to OFDM conversion for improved capacity and robustness.

In the next experiment, the SC-QAM utilization limit (after optimization) was increased from 60% to 80%.

Table 5 – Spectrum Allocation Optimization: Increased SC-QAM Utilization Limit

Parameter Name	Value
Maximum allowed SC-QAM utilization	**80%**
Maximum allowed OFDM utilization	60%
SC-QAM channel size	6 MHz
Base bandwidth per 6 MHz	37 Mbps
Minimum number of SC-QAM channels	**16**
Minimum SC-QAM allocation block	4 x 6 MHz = 24 MHz
OFDM guard band size (1 guard band)	1 MHz
OFDM size limit	48 – 192 MHz

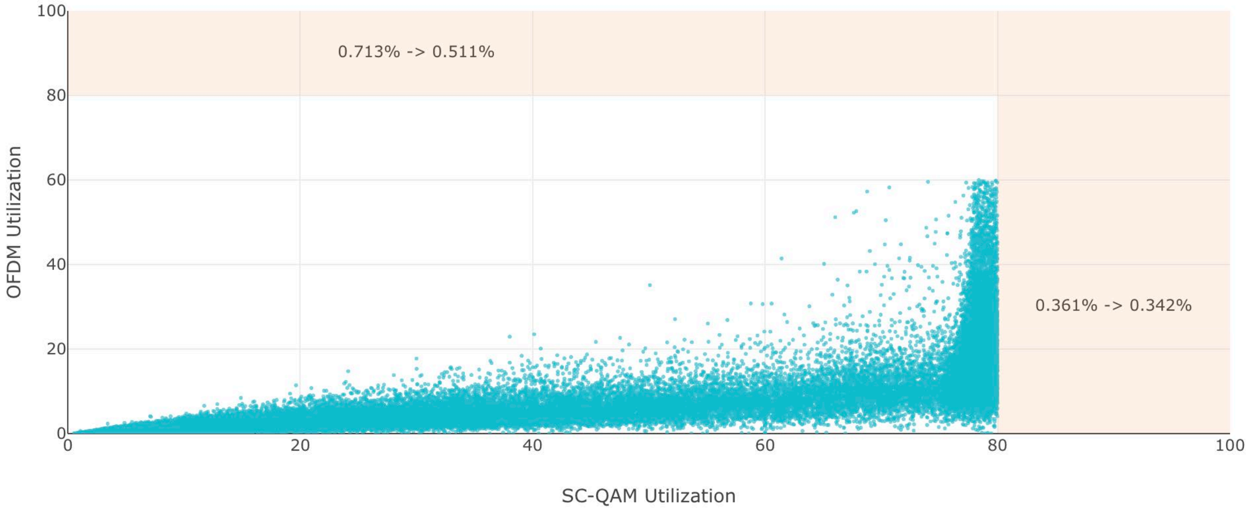

Figure 19 – Joint SC-QAM and OFDM Utilization Distributions (optimized with a minimum of 16 SC-QAM channels required and 80% SC-QAM utilization limit)

This constraint change increased the estimated average capacity gain from 321 Mbps to 347 Mbps. Meanwhile, the percentage of service groups that can be optimized increased from 94% to 98.3%, 0.102% of service groups were relieved from high OFDM utilization (>= 80%), and a small number of service groups (0.019%) were relieved from high SC-QAM utilization as a side effect from converting SC-QAM spectrum into OFDM spectrum.

5. Configuration Automation and Hit-less Changes

To apply the optimized spectrum configurations, we developed an automation service that is responsible for the translation and application of the updated spectrum configurations. This automation service coordinates with our other internal services to support seamless RF channel configuration changes without service disruption.

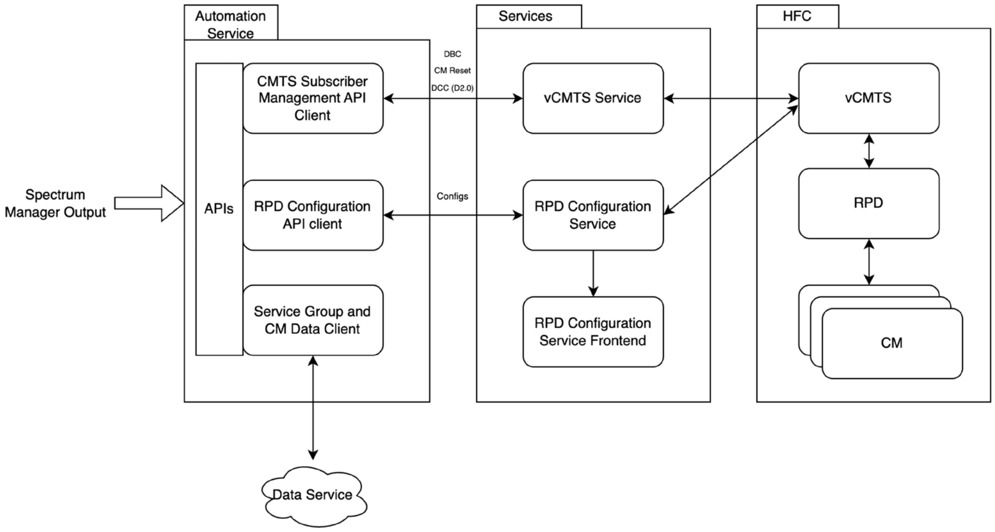

Figure 20 – PoC Software: Automation and Integration Service

5.1. Hit-less Spectrum Configuration Change Procedure

DOCSIS media access control (MAC) and upper layer interface specification [5] requires the CMs to support dynamic bonding change (DBC) and dynamic channel change (DCC) using DOCSIS MAC management messages. The DBC messages are supported by DOCSIS 3.0 and newer versions, and the DCC messages can be used for older versions of DOCSIS devices that exist in the network. In the downstream, DBC and DCC allow the CMs to dynamically change their receive channel set (RCS) without introducing significant service disruption (depending on the initialization technique used). This means that the CMs can dynamically change the frequencies they receive and can change their primary channel’s frequency. For DOCSIS 3.0 and newer CMs, channel bonding is supported, and the cable modem termination system (CMTS) maintains the information of logical bonding groups that refer to different sets of channels.

In our lab testing environment, the automation service is programmed to execute the following procedure to ensure hit-less configuration changes for each individual service group at scale:

1. Compute the differences in the spectrum configuration change, create or reuse a set of intermediate bonding groups that make references to SC-QAM channels that are not being changed. These bonding groups are used by the CMs temporarily during the configuration change.
2. Use DBC messages to move all the service group’s CMs with different capabilities, such as 4, 8, 16, 24, 32 SC-QAM CMs, to the intermediate bonding groups in parallel, with CM count-based load balancing by bonding capabilities. This step ensures that during the configuration change, the SC-QAM channels used by the CMs, including the primary channels, are not changed. In this step, the “direct” initialization technique is used. The parallel DBC messages can significantly reduce the wait time in a service group with many CMs. And the load balancing implemented by the automation service further reduces potential service impact due to the temporarily reduced total usable spectrum during the configuration change.
3. Apply RF spectrum configuration changes, respectively. These changes include SC-QAM and OFDM channel configurations, MAC domain configurations (primary channels, bonding groups etc.), and synchronization of any other associated configurations. This is also where vendor-specific configurations are handled.

4. Once the new RF spectrum configurations are accepted and stable, use DBC messages in parallel to move the CMs to their best-fit bonding groups, with load balancing.
5. As a cleanup step, once step 4 is complete, remove all redundant bonding groups created during the configuration change.

This procedure allows for RF channel changes without service disruption and with minimal and temporary capacity reduction during the change. Each individual configuration change procedure is executed sequentially and changes for different service groups are handled in parallel for scalability.

5.2. Closed-loop Integration and Results

Using this automation service, we conducted experiments to test the spectrum management VNF and the automation service with various CMs in our lab. These CMs include a mix of DOCSIS 3.0 devices with 4, 8, 16, 24, 32 bonding capabilities, and DOCSIS 3.1 devices. The total number of CMs is 32.

In the following test, the existing spectrum configuration has 44 SC-QAM channels and a 96 MHz OFDM channel, as shown in Figure 21.

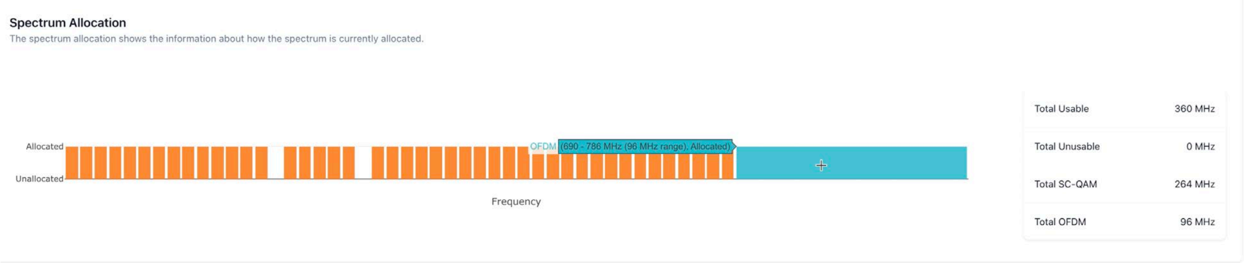

Spectrum Allocation

The spectrum allocation shows the information about how the spectrum is currently allocated.

Total Usable	360 MHz
Total Unusable	0 MHz
Total SC-QAM	264 MHz
Total OFDM	96 MHz

Figure 21 – Closed-loop Integration: Baseline Spectrum Configuration

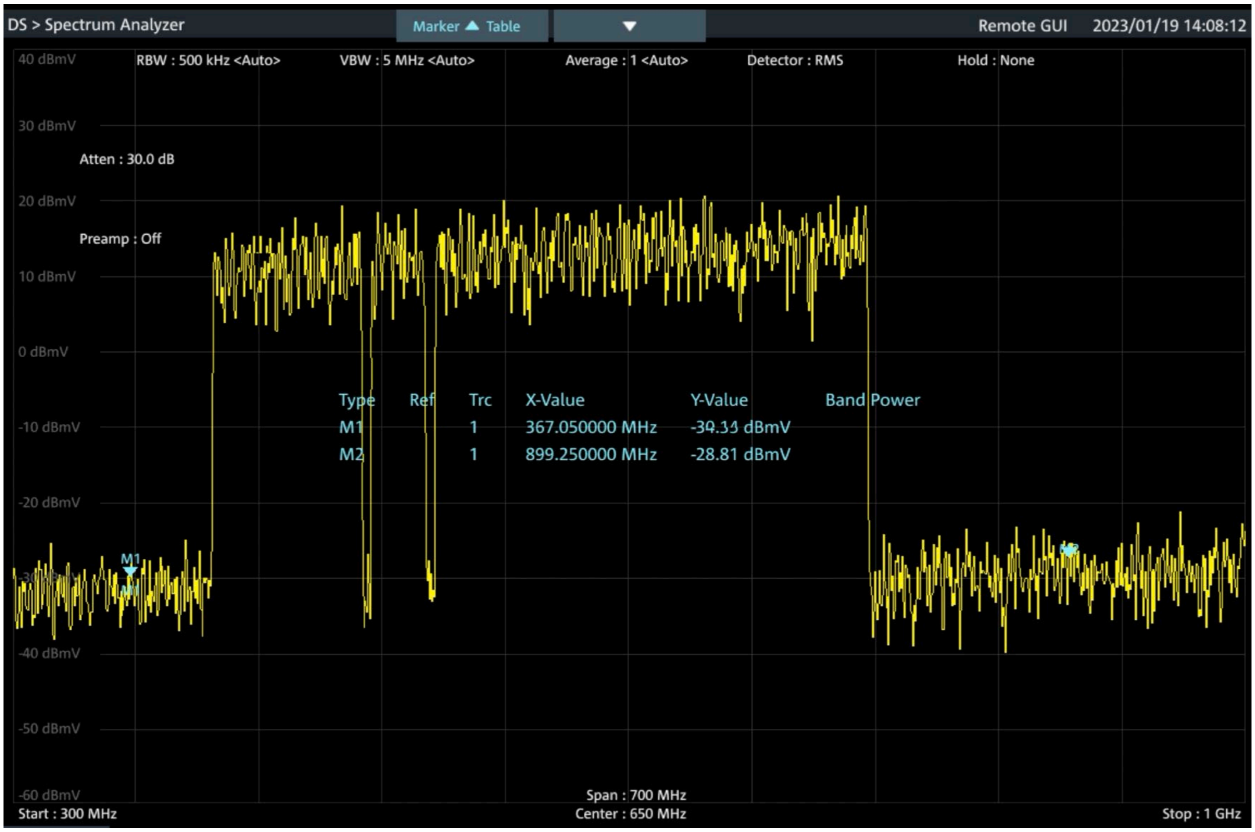

RBW : 500 kHz <Auto>
VBW : 5 MHz <Auto>
Average : 1 <Auto>
Detector : RMS
Hold : None
Atten : 30.0 dB
Preamp : Off

Type	Ref	Trc	X-Value	Y-Value	Band Power
M1		1	367.050000 MHz	-30.33 dBmV	
M2		1	899.250000 MHz	-28.81 dBmV	

Start : 300 MHz
Span : 700 MHz
Center : 650 MHz
Stop : 1 GHz

Figure 22 – Closed-loop Integration: Baseline Spectrum Configuration (Spectrum Analyzer View)

The peak throughput tested from a DOCSIS 3.1 device with a 2.5 Gbps Ethernet port is ~1.8 Gbps. The total capacity of this spectrum configuration is greater than 1.8 Gbps. However, because the DOCSIS 3.1 CM can only bond to up to 32 SC-QAM channels, 12 of the SC-QAM channels were not used by the CM during the speed test, reducing the peak throughput by 440 – 450 Mbps.

The optimized spectrum configuration from the VNF is:

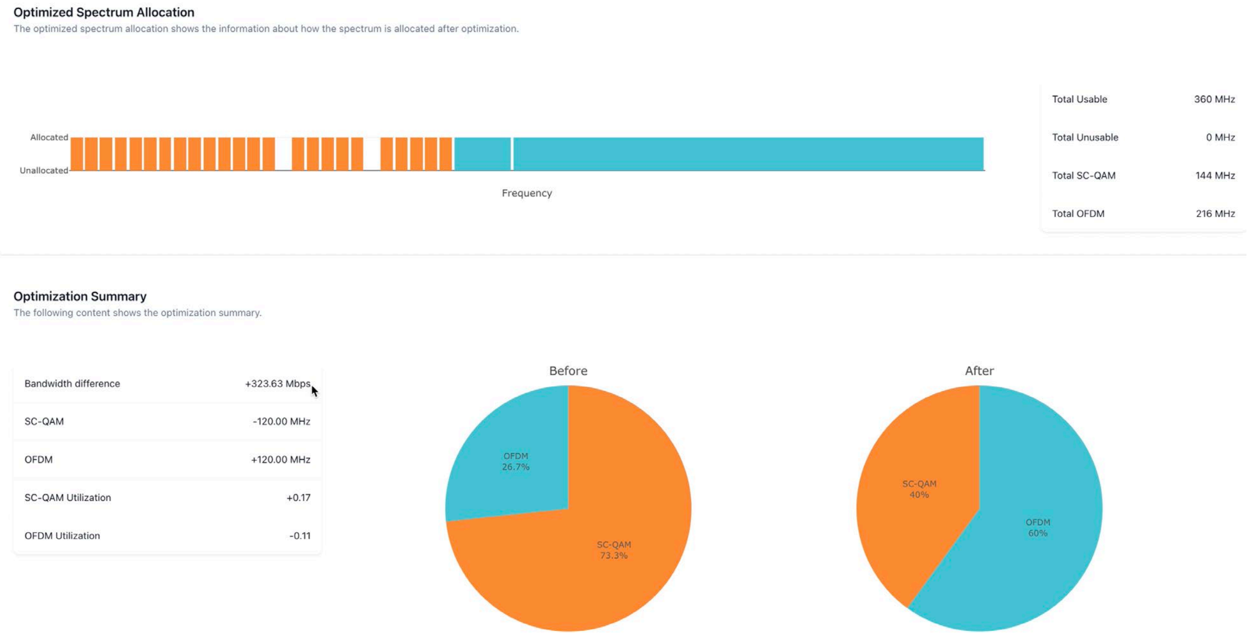

Optimized Spectrum Allocation

The optimized spectrum allocation shows the information about how the spectrum is allocated after optimization.

Total Usable	360 MHz
Total Unusable	0 MHz
Total SC-QAM	144 MHz
Total OFDM	216 MHz

Optimization Summary

The following content shows the optimization summary.

Bandwidth difference	+323.63 Mbps
SC-QAM	-120.00 MHz
OFDM	+120.00 MHz
SC-QAM Utilization	+0.17
OFDM Utilization	-0.11

Figure 23 – Closed-loop Integration: Optimized Spectrum Configuration

To test the configuration impact on connectivity and total capacity, background traffic in both downstream and upstream to/from all CMs was scheduled in combination with periodic speed tests. The real-time traffic is displayed on a dashboard shown in Figure 24.

Figure 24 – Closed-loop Integration: Background Traffic

Once the optimized spectrum configuration was sent to the automation service, steps 1 and 2 in the procedure were executed and the CMs were moved to intermediate bonding groups. The DBCs caused a brief 1-2 second traffic rate drop as the CMs had to temporarily use the limited downstream spectrum during the bonding group change. This can be found in the traffic rates history shown in Figure 25 between 7:13 AM and 7:14 AM. After the bonding group changes, the traffic rates stabilized as shown in Figure 25.

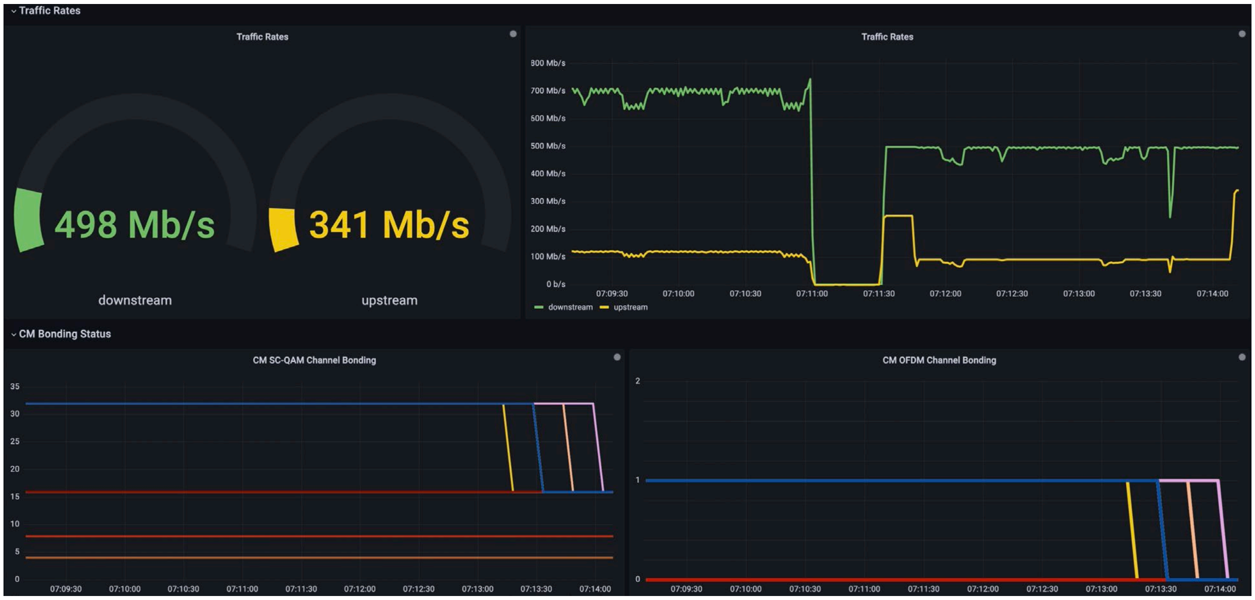

Figure 25 – Closed-loop Integration: Traffic Rates (After Moving to Intermediate Bonding Groups)

The automation service then sent the new RF channel configuration to the vCMTS API. As the first step, the RPD turned off the OFDM channel and removed 20 SC-QAM channels, as shown in Figure 26.

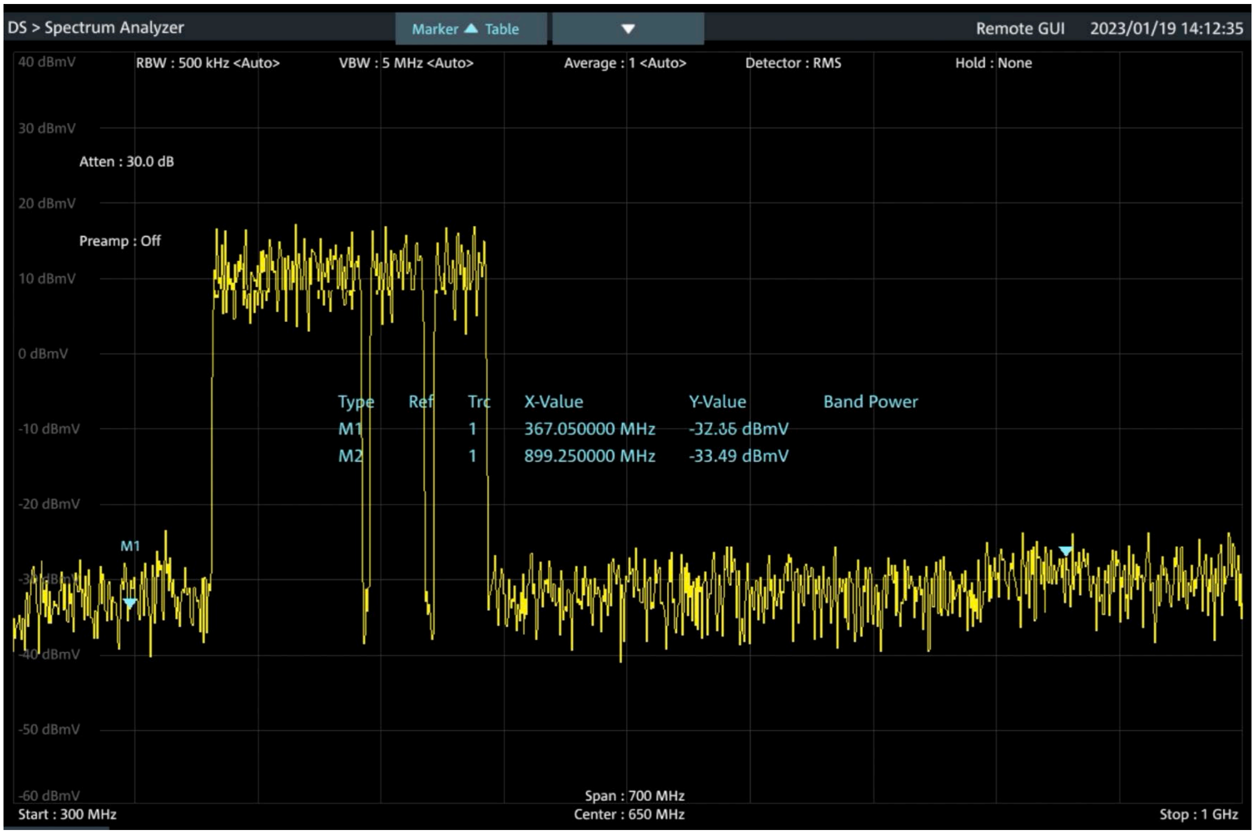

RBW : 500 kHz <Auto>
VBW : 5 MHz <Auto>
Average : 1 <Auto>
Detector : RMS
Hold : None
Atten : 30.0 dB
Preamp : Off

Type	Ref	Trc	X-Value	Y-Value	Band Power
M1		1	367.050000 MHz	-32.36 dBmV	
M2		1	899.250000 MHz	-33.49 dBmV	

Figure 26 – Closed-loop Integration: Spectrum Analyzer View During the RF Change

The RPD then turned on the 2 new OFDM channels configured as the VNF specified: a 24 MHz OFDM channel and a 192 MHz OFDM channel. As a result, this configuration change converted 120 MHz of SC-QAM spectrum into OFDM, as shown in Figure 27.

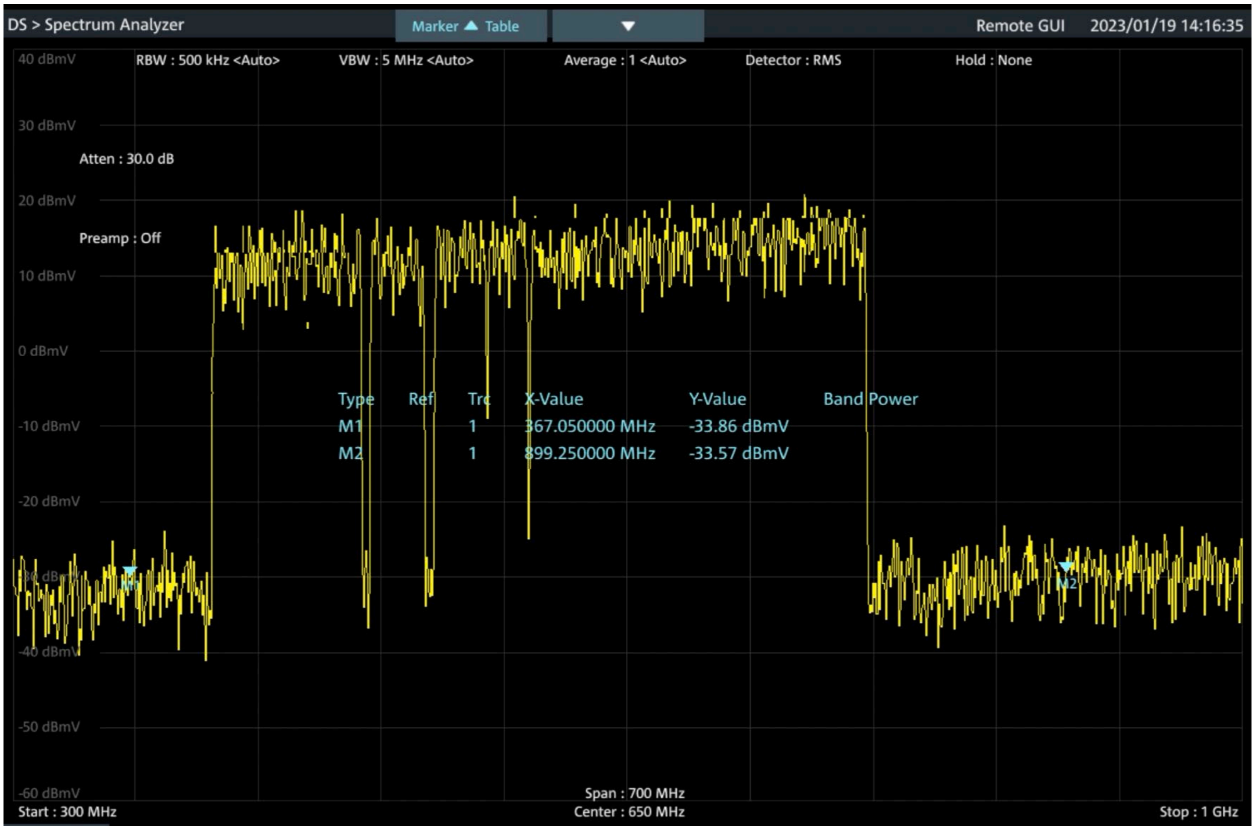

RBW : 500 kHz <Auto>
VBW : 5 MHz <Auto>
Average : 1 <Auto>
Detector : RMS
Hold : None
Atten : 30.0 dB
Preamp : Off
Start : 300 MHz
Span : 700 MHz
Center : 650 MHz
Stop : 1 GHz

Type	Ref	Trc	X-Value	Y-Value	Band Power
M1		1	367.050000 MHz	-33.86 dBmV	
M2		1	899.250000 MHz	-33.57 dBmV	

Figure 27 – Closed-loop Integration: Spectrum Analyzer View After the RF Change

Once the RF change was complete and stabilized, the automation service started to move the CMs back to their best-fit bonding groups based on their capabilities. No CM indicated partial service after the bonding change and the background traffic rates recovered to the starting point, as shown in Figure 28.

Figure 28 – Closed-loop Integration: Background Traffic Rates After the Configuration Change

After the configuration change was complete, the DOCSIS 3.1 CMs were fully bonded to the 24 SC-QAM channels and 2 OFDM channels (216 MHz total spectrum) on the downstream. This means it is expected to see a peak downstream throughput increase beyond the estimated 323.63 Mbps capacity gain coming from spectral efficiency improvements. The speed test measured ~2.38 Gbps application traffic throughput on a single CM with 2.5 Gbps Ethernet ports when the channels were under 100% utilization, as shown in Figure 29. As a result of re-allocating the DOCSIS technology across spectrum, the device was able to achieve almost 600 Mbps higher speeds above the baseline configuration. These higher speed capabilities come along with additional benefits to the total capacity and lower utilization of the service group. This was all accomplished with non-service affecting customer impact.

Figure 29 – Closed-loop Integration: Downstream Speed Test After the Configuration Change

6. Conclusion

In this paper, we presented our design of an efficient and effective algorithm for the spectrum management VNF, and various experiments and their results from applying this VNF to production service group data. We also introduced our PoC software and the automation service that we have developed to support closed-loop testing and experiments and a lab demonstration with real-time information captured during the automated, non-service-affecting spectrum change, and shared the peak downstream throughput improvement after the automated configuration change. We conclude that:

- The spectrum management algorithm can effectively account for various real-world constraints that have been identified so far, and it is performant.
- The experiments around applying the VNF to production service group data showed promising average capacity gain and peak speed improvements from automatically converting SC-QAM to OFDM based on each service group's characteristics, without adding new spectrum.
- The PoC software and the automation service helped envision and demonstrate the closed-loop integration for the spectrum management VNFs. And it is possible to perform DOCSIS spectrum changes without disrupting service.

As part of this effort, we also discovered improvements to the load-balancing algorithms that are now being implemented as a VNF outside of the core vCMTS microservice. As the load balancing VNF becomes mature, production trials of this spectrum management technique will be a key element associated with our FDX deployment scale for DOCSIS 4.0.

Abbreviations

AGC	automatic gain control
API	application programming interface
bps	bits per second
CAGR	compound annual growth rate
CM	cable modem
CMTS	cable modem termination system
DAA	distributed access architecture
DBC	dynamic bonding change
DCC	dynamic channel change
DOCSIS	Data Over Cable Service Interface Specifications
Gbps	gigabits per second
GHz	giga hertz
GUI	graphical user interface
HFC	hybrid fiber coax
MAC	media access control
Mbps	megabits per second
MHz	mega hertz
MILP	mixed-integer linear programming
OFDM	orthogonal frequency division multiplexing
PHY	physical layer
PMA	profile management application
PoC	proof of concept
QAM	quadrature amplitude modulation
RCS	receive channel set
RF	radio frequency
RPD	remote PHY device
SC-QAM	single carrier-quadrature amplitude modulation
vCMTS	virtual cable modem termination system
VNF	virtual network function
VoD	video on demand

Bibliography & References

1. *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.
2. *A Machine Learning Pipeline for D3.1 Profile Management*, M. Harb, J. Ferreira, D. Rice, B. Santangelo, and R. Spanbauer, NCTA technical paper, 2019.
3. *Full Scale Deployment of PMA*, M. Harb, J. Ferreira, D. Rice, B. Santangelo, NCTA technical paper, 2020.
4. *DOCSIS 3.1 Physical Layer Specification*, CM-SP-PHYv3.1-I19-211110
5. *DOCSIS 4.0 MAC and Upper Layer Protocols Interface Specification*, CM-SP-MULPIv4.0-I05-220328