Qualifying Network Performance and Impairment Priority in Modern DOCSIS Networks

Healing the Self-Healing Network

A technical paper prepared for presentation at SCTE TechExpo24

Mike O'Dell
Distinguished Engineer
Comcast
michael_odell@cable.comcast.com

Maher Harb
Distinguished Engineer, Data Science
Comcast
maher_harb@comcast.com

Mike Spaulding
VP, Plant Maintenance
Comcast
michael_spaulding@cable.comcast.com

Scott Johnston
Director, Network Engineering
Comcast
bryan_johnston@cable.comcast.com

Westley Darby
Engineer, Quality Assurance
Comcast
westley_darby@cable.comcast.com

Alex Zayvelev
Senior Access Engineer
Comcast
alexander_zayvelev@comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

1. Efficiency as a Concept

The concept of efficiency in a system is not new. Almost all systems and machines intend to be as efficient as possible, whether it's a windmill, an automobile, a refrigerator, or a hybrid fiber/coax (HFC) network. The designers and manufacturers of almost everything create with efficiency as a factor. There are savings implications for both the manufacturer and the consumer when efficient things are made and used. Efficient things generally result in lower operating costs and potentially longer useful life spans. But what does efficiency mean in the context of modern Data-Over-Cable Service Interface Specifications (DOCSIS®) networks? To understand efficiency in this context, we looked at the definition of efficiency in several dictionaries, and found two definitions that add value to our concept:

Efficiency – noun: The quality of achieving the largest amount of useful work using as little energy, fuel, effort, etc. as possible (Cambridge Dictionary)

Efficiency – noun: (b): the ratio of the useful energy delivered by a dynamic system to the energy supplied to it (Merriam-Webster Dictionary)

In both definitions, there are elements that are useful for describing the optimal operation of a DOCSIS network. In the Cambridge Dictionary definition, the idea of the *largest amount of useful work*, *using as little energy as possible,* is intriguing. It leans into the conducted valuable work to make our networks more energy-efficient and sustainable. In the Merriam-Webster dictionary definition, the concept is that there is a ratio-driven relationship between the energy applied to the system and the energy delivered by the system. Also, the fact that the system is dynamic is particularly relevant in this application.

In this paper, we explore the various ways to use these definitions of efficiency, how they can be used to define impairments impacting a DOCSIS network, and how they can be a useful method of prioritizing remediation efforts.

2. Efficiency in DOCSIS Networks

The idea of efficiency in DOCSIS networks is not new either. In 2021, CableLabs published a technical report - *Improving DOCSIS Efficiency*. This report outlines the variables available to configure upstream and downstream channels to maximize throughput and efficiency, i.e., carry as many bits per second per hertz (bps/Hz) in single carrier quadrature amplitude modulation (SC-QAM) channels, multi-carrier signals such as orthogonal frequency-division multiplexing (OFDM) channels, or orthogonal frequency-division multiple access (OFDMA) channels. The impetus for this work, initially released in 2021, was to address the unprecedented increase in upstream traffic usage due to the shift in consumer internet usage driven by the COVID-19 pandemic.

The CableLabs report focused on improving the efficiency of the DOCSIS upstream SC-QAM and OFDMA channels. It outlined some of the challenges that operators faced and suggested how to make their networks more robust in the face of impairments. Some of the suggestions offered in the technical report were recommendations for configurable values, such as mini-slot size, forward error correction (FEC) codeword size (Figure 1), cyclic prefix, fast Fourier transform (FFT) size, frame size, pilot patterns, channel size, guard bands, profile definition, and cable modem termination system (CMTS) settings for changing interval usage codes (IUCs). The report illustrates the relationship between signal quality, modulation order, and packet loss in various configured states (Figure 2). It further recommended downstream OFDM channel configuration and spectrum placement (Figure 3) to achieve the highest data rates, with significant resilience built into the channels.

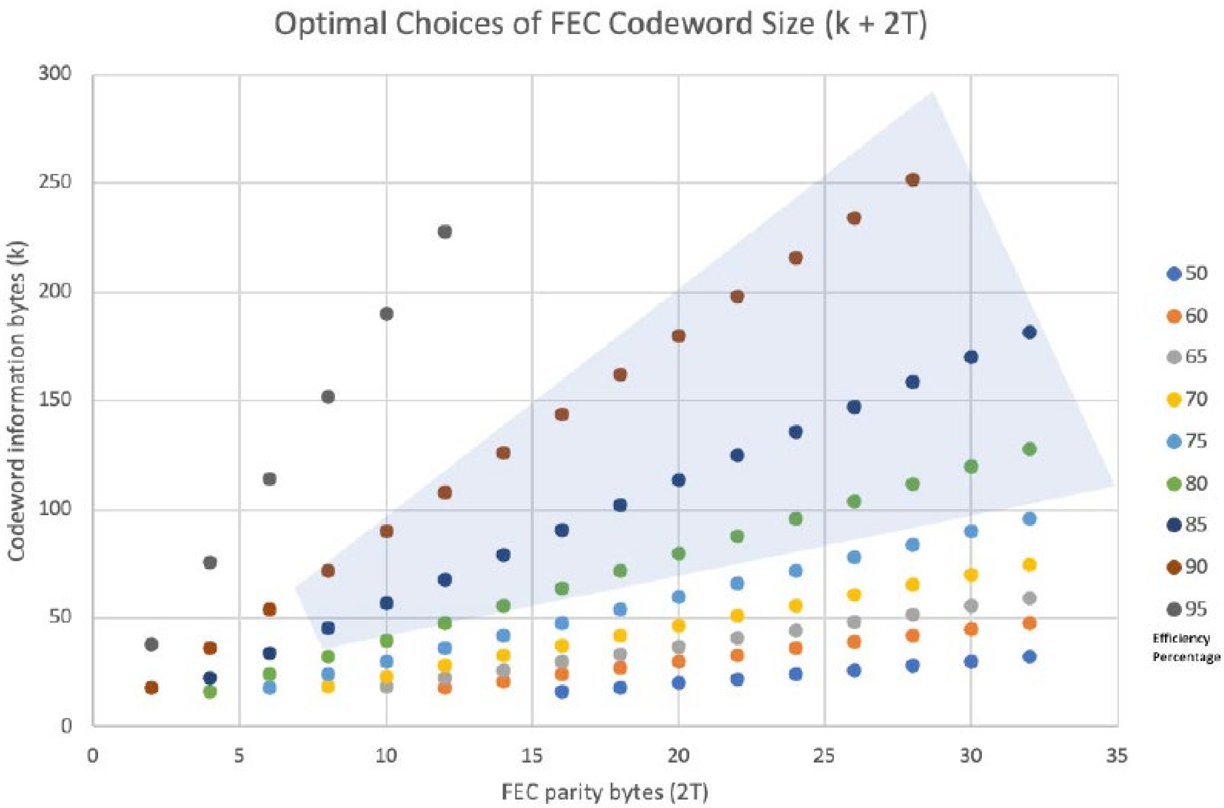

Figure 1 - Range of SC-QAM FEC codeword sizes

There is also a meaningful amount of content dedicated to the ability to construct profile management application (PMA) tools. This feature was introduced with the DOCSIS 3.1 specification and allows for configuring specific modulation profiles or bit-loading characteristics for each subcarrier within the OFDM and OFDMA channels. This extraordinary technology allows a network to adapt to the dynamic nature of the transmission medium within it and meaningfully ensure the customer experience (CX).

There are implications to this technology, however. The report clearly states in the conclusion: "the operators need to ensure that their networks are running as efficiently as possible while still maintaining robust service offerings, but there is a fundamental tradeoff between robustness and throughput" [1].

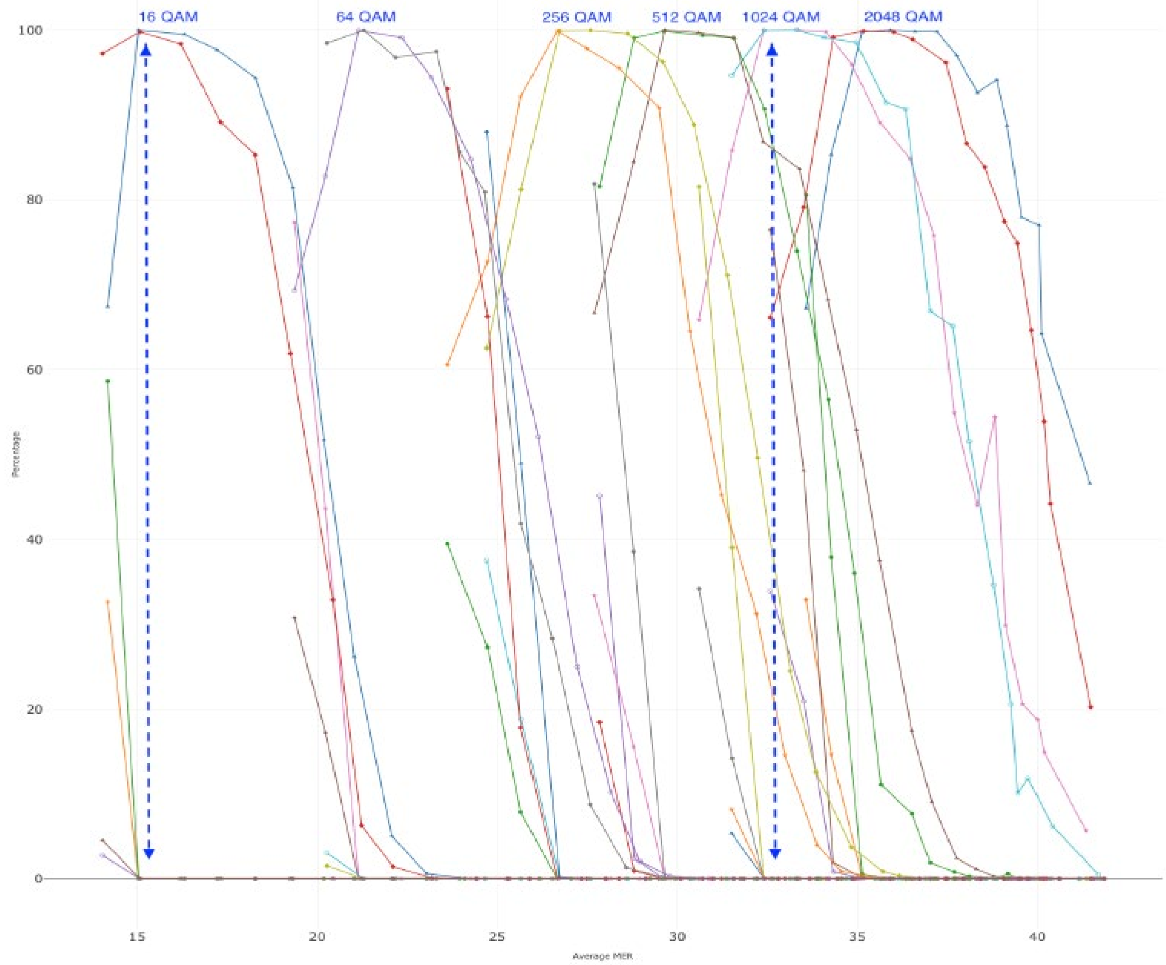

Figure 2 - OFDMA FEC Errors and Packet Loss Against Average RxMER (with AWGN) 200-Byte Packets

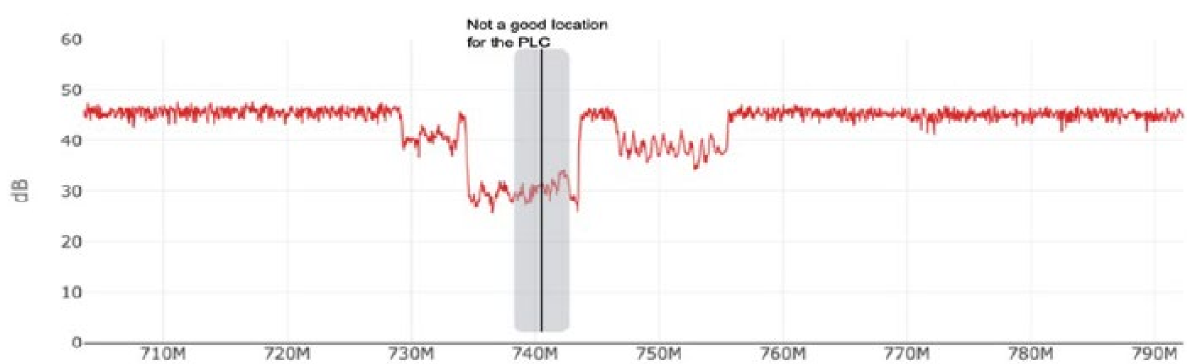

Figure 3 - CM RxMER Impact Due to LTE Ingress (729 MHz to 756 MHz)

However, the report does not cover how to measure and prioritize the impact on the network when PMA is actively remediating the symptom of network impairment.

3. Energy Efficiency in Modern Networks

Operators like Comcast and others are implementing comprehensive energy management systems (EMS) to address the sustainability and energy consumption required to operate large-scale networks. The development of metrics or key performance indicators (KPIs) measures the effectiveness of these programs, from the energy efficiency of data centers, headends, and other critical infrastructure facilities to the outside plant DOCSIS network power consumption and fleet vehicles. Energy data is collected, aggregated, and broken into significant energy users (SEUs). These complex measures can provide insight into the cost of operating large-scale networks and drive focus on becoming more efficient, sustainable, and profitable in multiple areas of the operator's business strategies.

> "Cable Operators consume energy in unique ways relative to other industries, making a structured, data-driven approach to energy management a requirement to meet Carbon Neutrality and Network Energy Efficiency goals in an optimized manner. Managing the unique energy footprint of a network can often feel like a massive challenge (like the Desmond Tutu metaphor of eating an elephant) but implementing an energy management system (EMS) that can evolve over time makes controlling network energy possible. The foundation for an energy management program is data. Energy consumption must be measured at actionable levels so that performance can be quantified." [2].

One reference focuses on one energy metric in particular, energy per consumed byte (EPCB):

> "The Green Grid (TGG) has developed a framework approach to productivity called data center energy productivity (DCeP). DCeP is an equation that quantifies useful work that a data center produces based on the amount of energy it consumes. 'Useful work' is a sum of tasks that are completed in a specified period of time. DCeP allows each user to define useful work and the weighting for various forms of useful work (if measuring more than one type) that apply to that user's business. The definitions can get as granular as necessary for the entity using the equation: web pages served, database transactions executed, emails served, etc.
>
> "In keeping with the TGG approach outlined previously of leaving it to the users to define what 'useful work' is, either useful work shall be consumed bitrate or consumed bytes. The resulting metric, to be detailed in the sections following, is 'Energy per consumed byte,' or EPCB [*sic*] " [3].

In the outside plant (OSP) footprint of the cable operator's DOCSIS network, there is typically one or more power supplies connected to the commercial electrical grid that provide power to the operator's cable access network to enable the transmission of signals through active devices, allowing multi-path connectivity to the customer base. The power consumption should remain largely constant within that node footprint, provided that the outside plant network is free from significant defects. Much work has gone into designing and optimizing the power supply networks to minimize inefficiency and continually right-size the power supplies to match the demand of the network itself. Further, the node service area is largely static with respect to the physical geography served by a particular fiber optic node service group. The customer usage behaviors in that node service area are not likely to vary significantly over short periods, barring another unforeseen circumstance like the COVID-19 pandemic. From this perspective, the EPCB is relatively fixed when correlated to the radio frequency (RF) performance within the footprint served by an optical node. So, in this application, energy efficiency and plant or spectral efficiency are not directly proportional in a meaningful way.

4. Characteristics of Modern DOCSIS Networks

DOCSIS 4.0 implementations are being deployed using different methodologies. The specification allows full duplex (FDX) DOCSIS and frequency division duplexing (FDD).[1] Both methods seek to expand the usable bandwidths of the most efficient DOCSIS technologies available, which are currently OFDM and OFDMA. In contrast to current deployments of sub-split networks laden with linear video, video on demand, and switched digital video bandwidth (Figure 4), DOCSIS 4.0 implementations seek to allocate the significant majority, if not all, of the available spectrum to data carriage channels, as a means of achieving multi-gigabit speeds. As we remove DOCSIS 3.0 and older premise equipment from the network, it is also feasible that DOCSIS SC-QAM channels could be deprecated in favor of additional OFDM channels. These networks would increasingly rely on IP video as the necessary video delivery methodology, using the more efficient OFDM channels.

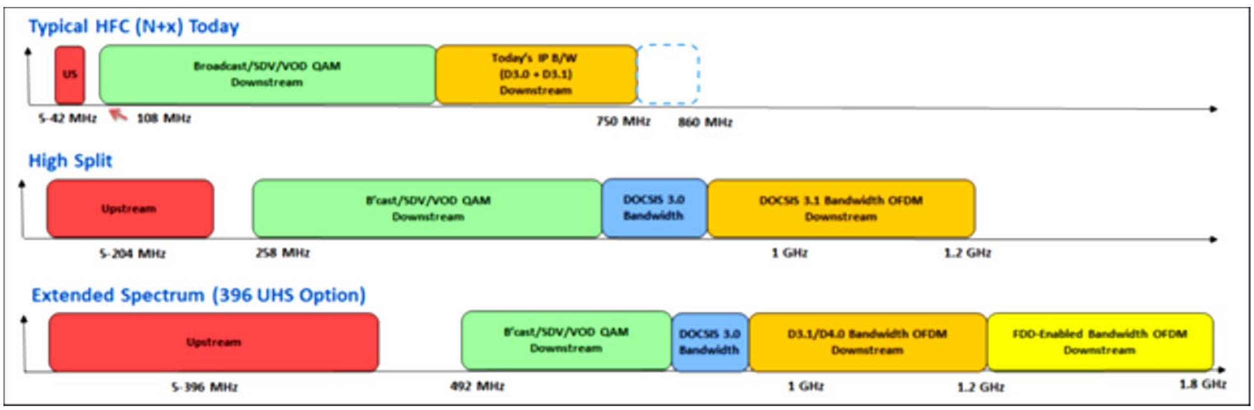

Figure 4 - Typical HFC vs. High-Split vs. DOCSIS 4.0 Extended Spectrum Operation

While the different methodologies within DOCSIS 4.0 expand these signal types in different ways, they share some common DOCSIS resources in both the upstream and downstream (Figure 5**).** Is this a glimpse of the future?

> "Observing the two DOCSIS 4.0 options, we see that both add a common set of upstream OFDMA blocks over a common frequency range. One, FDD, pushes a significant physical bandwidth extension into the network relying on common RF practices executed for several generations of HFC migration. The other, FDX, relies on the introduction of new digital signal processing (DSP) technology to complement the existing plant technology in order that the spectrum asset be operated more efficiently over a bandwidth that the plant is built to support.

> "Importantly, these are not mutually exclusive approaches! In fact, they are complementary technologies that, in principle, could be merged in some future DOCSIS extension. With the DOCSIS 4.0 upstream capacity as defined for FDX or FDD in the 5-

[1] Originally called ESD or extended spectrum DOCSIS, the DOCSIS 4.0 specifications use the terminology "DOCSIS 4.0 frequency division duplexing," often called just FDD.

6 Gbps range, merging the two technologies may be the next step toward pushing the upstream also towards the 10 gig goalpost" [4].

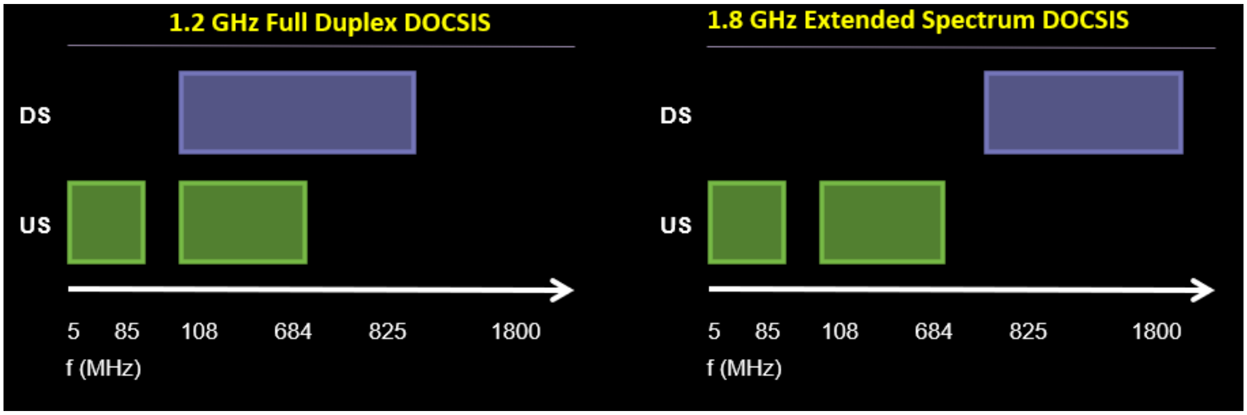

Figure 5 - Common DOCSIS Resources in FDX and FDD

As the "pipe" becomes more predominantly filled with the dynamic OFDM and OFDMA signal types, the more significant part that PMA plays in configuring and managing the signals in the network is to ensure the customer experience. More on the implications of that is discussed later in Section 6.

5. Operating in a Hostile Environment

Whether you are extending the bandwidth of your system to new operating frequencies or just reorganizing your existing spectrum to deploy additional DOCSIS bandwidth, there are perils associated with data channel bandwidth expansion. Cable networks are designed and intended to be closed systems, which, by rough definition, have sufficient shielding to keep their transmitted signals in and over-the-air signals out. This enables cable systems and over-the-air providers to use the same frequencies to operate their networks, also known as frequency reuse.

As the connectivity landscape evolves and the needs of a connected world increase, more and more over-the-air signals are activated to meet those needs. This creates a very crowded and hostile over-the-air environment. When the components of the cable system are properly installed and maintained, the shielding effectiveness of the system is quite good. It can withstand high levels of over-the-air energy and maintain a high integrity path for the system's intended signals. The reality, however, is that the network lives in a harsh environment. Constant exposure to wind, rain, UV rays, heat, cold, snow, and other environmental factors can cause the plant to move, expand, contract, and ultimately fail in some way. When the shielding integrity of the coaxial portion of the network fails, the potential for impact of some portion of the system's intended signals is relatively high. Ingress can occur and impact signals across nearly the entire spectrum typically used by cable operators (Figure 6).

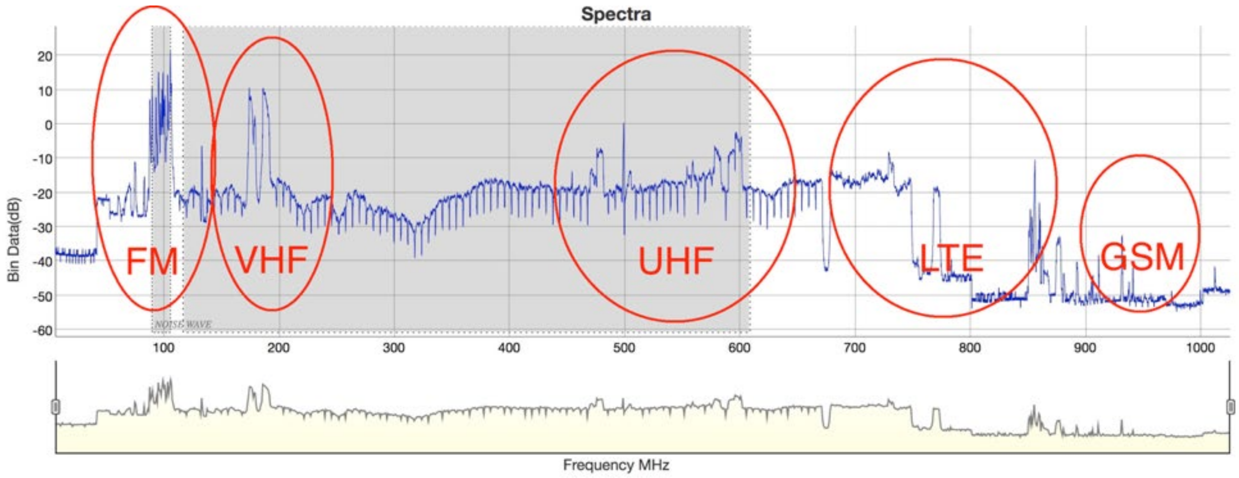

Figure 6 - Symptoms of Ingress as Displayed by Full Band Spectrum Capture

It is practically inevitable that as operators expand their DOCSIS channel footprints to allow for greater throughput in both the upstream and downstream directions, introduced are new collisions with impairments previously not common for DOCSIS channels. Logically, one would assume that these impairments already existed and impacted other services, such as video performance. That is undoubtedly true; however, the difference would be the transactional relationship operators have built to remediate the issue in the customer premise with a trouble call or an outside plant issue with a network truck roll. In instances where the customer experience is noticeably impaired in a video SC-QAM channel without the technical benefit of self-healing, the only available option would be a truck roll. The desired outcome of expanded DOCSIS bandwidth is to deploy OFDM or OFDMA channels with high-order modulation and self-healing properties. Their potential spectral efficiency is much greater than SC-QAM channels; however, significantly higher signal quality is required to operate at the highest modulation orders, as illustrated in Figure 7, from [5]. These same impairments that would have driven truck rolls to repair impact to legacy 64- or 256-QAM channels may now reside under an OFDM or OFDMA channel. A PMA could identify the symptoms of the impairment, recommend a profile change to one or more mini-slots, and mitigate any impact on the customer. Great, right? Well, maybe. From the customer experience lens, this is a desired outcome. From a network management perspective, however, the self-healing properties of the network also have the potential to hide or mask impairments in the network unless identified and alarmed by a network monitoring system. As discussed in the next section, adding PMA tools to expanded DOCSIS channels is an operational prerogative but comes with risks.

Modulation Order	RxMER
256-QAM	≥ 29~30 dB
512-QAM	≥ 31~33 dB
1024-QAM	≥ 34~36 dB
2048-QAM	≥ 37~39 dB
4096-QAM	≥ 40~42 dB

OFDM

Constellation	CNR[1,2] (dB)	Set Point (dBmV/6.4 MHz)	Offset
QPSK	11.0	-4 dBmV	0 dB
8-QAM	14.0	-4 dBmV	0 dB
16-QAM	17.0	-4 dBmV	0 dB
32-QAM	20.0	-4 dBmV	0 dB
64-QAM	23.0	-4 dBmV	0 dB
128-QAM	26.0	0 dBmV	0 dB
256-QAM	29.0	0 dBmV	0 dB
512-QAM	32.5	0 dBmV	0 dB
1024-QAM	35.5	0 dBmV	0 dB
2048-QAM	39.0	7 dBmV	0 dB
4096-QAM	43.0	10 dBmV	0 dB

OFDMA

Figure 7 - Signal Quality Requirements for OFDM and OFDMA

6. Self-Healing Networks

Profile management applications have been successfully deployed for years in some systems. The overarching goal of any PMA can be described in three stages: detect, mitigate, and fix. The fundamental elements of a profile management system can consist of data collectors, data storage, an analytics engine, and a configuration manager (Figure 8). The purpose of this paper is not to offer best practices for building or optimizing a PMA; several very good papers have been published to that end. Instead, it offers a new way of thinking about the performance of a network with a PMA that is actively implemented. Comcast has its implementation of PMA [6,7], known internally as Octave, which has been operational for several years. Comcast's Jay Zhu [8] describes their approach this way:

> “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... 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, […] 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 (PMA) VNF… 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 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.”

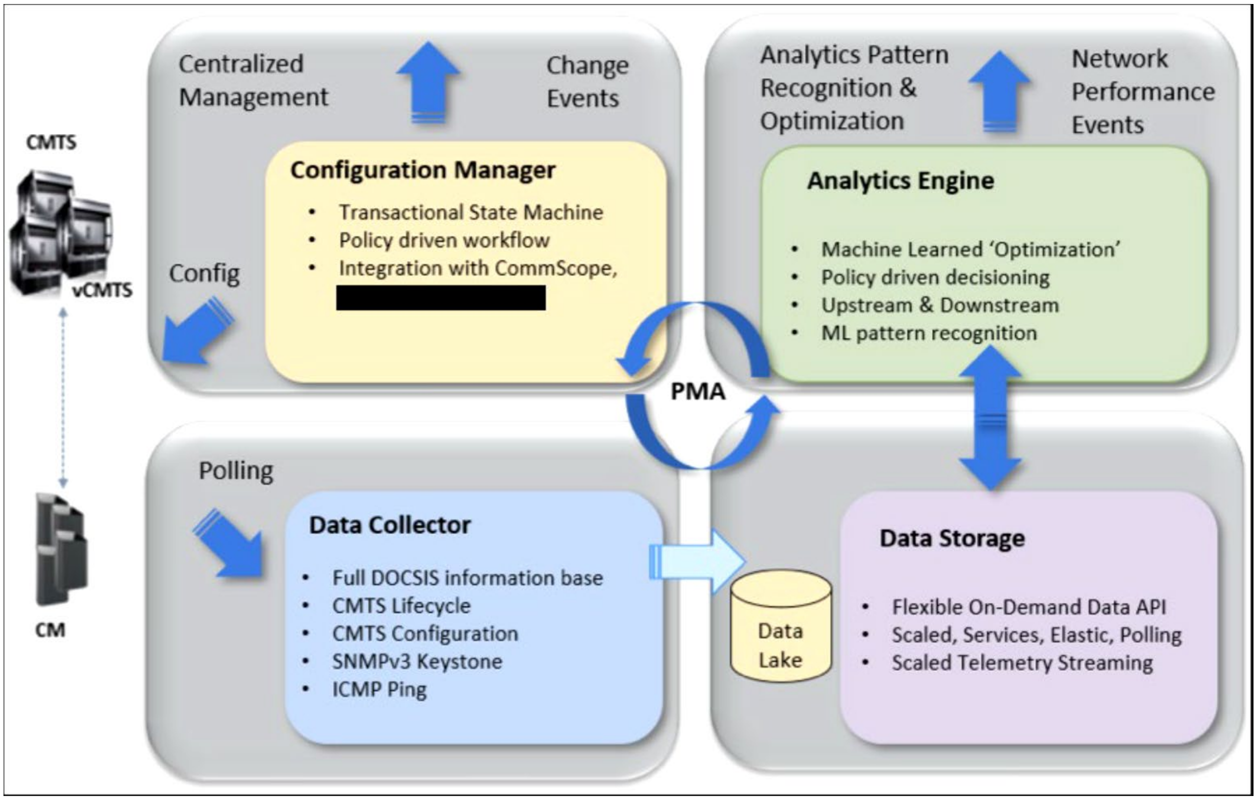

Figure 8 - Example of a PMA System Architecture

Several key points in this description are important to recognize. The reallocation of spectrum from SC-QAM channels to OFDM is essential, as OFDM channels are more spectrally efficient. However, they also benefit from being dynamic on the modulation profile, whereas SC-QAM channels in the downstream do not. The more spectrum that is allocated to OFDM, the more impaired bandwidth that a PMA can detect and mitigate, meaning the implications of impairment are not felt immediately by the customer base. Secondly, the virtualized network function can dynamically adjust at the service group level as the distribution of DOCSIS 3.1 and DOCSIS 4.0 devices incrementally increases, resulting in a more dynamic spectrum that the technical workforce has to interact with and understand. This technical complexity should not be overlooked; it has implications for technician-facing tools and training. The more applications like Octave make changes to mitigate the impact of impairments, the less efficient the network is. The greater the level of impairment, either in amplitude, bandwidth, or both, the lower the spectral efficiency of the network as a whole!

Challenges and considerations aside, the successful implementation of this application has resulted in significant improvements in upstream and downstream throughput, described as capacity gain, even without adding spectrum to the network. As new bandwidth is operationalized, that figure surely increases, but so will the operational complexity of managing the network. From a business perspective, the successful implementation of Octave could be measured as a significant reduction in network and customer-facing truck rolls. Rough estimates on network truck rolls alone, in the tens of millions, reduced since implementation to combat spurious or transient noise and persistent noise sources. The cost savings associated with that truck roll reduction figure are likely hundreds of millions of dollars.

The results of the Octave deployment overwhelmingly support the success of the detect and mitigate stages. Operationally, it is more challenging to determine how to evaluate when to generate a truck roll to enact the fix portion of the equation. The fact is that many of the impairments that result in a modulation profile adjustment are temporary and eventually self-clear, allowing Octave to upgrade or normalize the modulation profile back to its ideal or highest operating state. More importantly though, some require technician intervention to repair. The impairment's mitigation masks the traditional RF metrics that detect and identify an impairment within the channel's overall performance. The automated process easily overlooks impairments in the network, but the customer base experience is seamless. What prevents a serious challenge, is the tool suite's ability to recognize when a portion of the network is operating in a less-than-ideal efficiency state and create a process of scheduling remediation appointments with appropriate prioritization.

While this is good for customers and network technicians on the surface, are we putting our heads in the sand and ignoring real problems? It is feasible that these problems manifest themselves in DOCSIS 4.0 deployments as speed issues because the network cannot support muti-gigabit speeds due to the downgraded modulation orders in the impacted channels. DOCSIS 4.0 deployments promise much higher and, in some cases, symmetrical speed offerings. This is a fantastic opportunity to create capacity. However, the additional bandwidth may give the perception of capability that is actually far below the theoretical maximum of the network's capability. We know from experience that capability precedes consumption. At every point in the evolution of DOCSIS network deployments, as we've deployed greater capability, the consumer applications to consume that capability soon followed. Deploying these new DOCSIS 4.0 networks with PMA and self-healing properties allows for amazing possibilities for our customers. However, to meet the consumption and quality of service demands, operators may find themselves in a position to think about the efficiency of their deployed spectrum differently.

7. Characterizing Historic Network Performance

DOCSIS specifications and PNM tooling enable network performance monitoring and categorized measurements in many categories. From the most basic functions like registration state and transmit and receive level to more complex coefficients like in-channel frequency response (ICFR) and many other data points, we have an increasingly large volume of data points to choose from. Additionally, many telemetry and data collection platforms have transformed to acquire data more frequently and granularly. Operators have mountains of data they must aggregate, store, categorize, threshold, and prioritize to meaningfully determine which impairments, or more importantly, the symptoms of impairments, take the highest priority in the network remediation process.

Symptoms of access network impairments can be categorized in many ways but are generally grouped into three categories: signal quality performance measurements, data integrity performance measurements, and transmission medium performance measurements. In the first category, signal quality, we have historically measured the downstream channel's quality according to the channel type. Analog video channels, for instance, were qualified using baseband signal-to-noise (SNR) or RF carrier-to-noise (CNR), as well as measurements characterizing distortions like composite second order (CSO) or composite triple beat (CTB). On the other hand, digitally modulated channels more commonly describe their signal quality using receive modulation error ratio (RxMER). In the second category, data integrity measurements, the key performance indicators are the packet or bit error ratios, which describe an impairment's impact on the successful packet reception at the destination device. The final category, transmission medium performance, is most commonly frequency response. Sometimes called the peak-to-valley response, it indicates an impairment in an active device, passive device, or coaxial cable, which can create easily recognizable signatures in the frequency response, similar to the "water wave" (caused by water in the cable) shown in Figure 9**.**

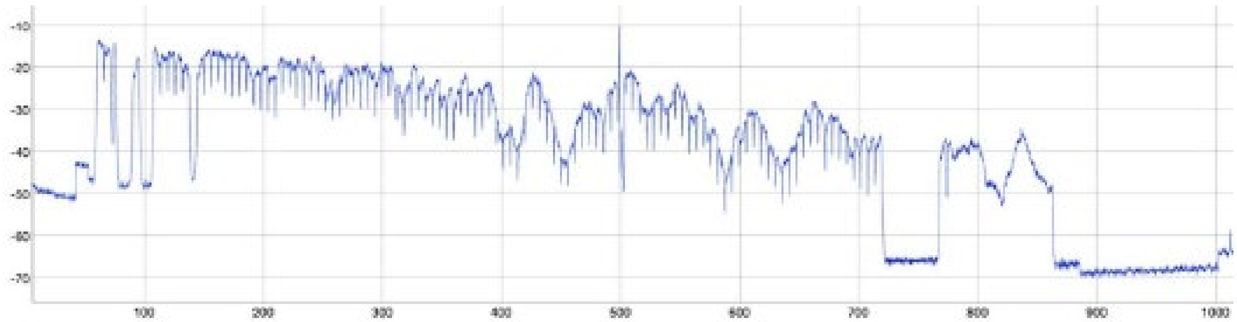

Figure 9 - Wideband Spectrum Capture Illustrating a Water Wave Impairment

The data points in these categories grow as operators and equipment manufacturers improve their hardware and software. The ever-increasing challenge is determining which of these symptoms, or combination of symptoms, is the most critical to repair. As access network technology evolves, shouldn't the way we measure the impairments affecting them also evolve?

8. Changing the Impairment Narrative

From the inception of broadband DOCSIS networks, the technicians maintaining those networks have understood the linear relationship between signal quality and data integrity, particularly in the sub-split upstream portion of the network. As noise or ingress is introduced into the signal path, the signal quality, typically measured in SNR or RxMER, declines relative to the level of the impairment. As the signal quality degraded, channels using Reed-Solomon error correction counters would begin incrementing correctable FEC errors. If the impairment continued to decline to the extent that Reed-Solomon error correction could not repair it, the interface would begin to increment packet loss or uncorrectable FEC errors.

Field technicians have long understood that this relationship defined the prioritization and urgency with which the impairment is addressed and resolved. Similarly, in the downstream SC-QAM channels, that linear relationship existed between RxMER and correctable and uncorrectable FEC errors in channels using Reed-Solomon error correction. As described in previous sections, there is an inarguable benefit to migrating spectrum from SC-QAM channels to OFDM and OFDMA channels. This is not only because of OFDM/A's spectral efficiency, which is measurably higher, but also because of their ability to dynamically manage the impact of impairments using profile management applications. But here's the rub: In effectively mitigating the impact of impairments, applications such as Octave impact the capacity and potentially the network's ability to deliver the operational maximum payload it is capable of.

As we deploy these high-speed, high-bandwidth DOCSIS 4.0 networks, we need to begin to think of their performance not against the customer experience solely or the legacy RF impairment markers that we have always used, but also on the dynamic management systems' impact on the performance of the network as a ratio of its operational maximum capability. Figure 10 shows the full band capture response of an OFDM channel superimposed over the per-subcarrier RxMER of that same device. This OFDM channel does have some exclusion bands to accommodate local insertion points in the network. The channel looks healthy using the traditional frequency response and receive power values. Underlying those measures, however, you can see that the RxMER is low and has a region of the OFDM channel that exhibits roll-off behavior in the RxMER.

If we dive deeper into this device, the OFDM channel uses the lowest modulation profile available due to the channel's poor RxMER performance. While the power level, error rate, and frequency response appear normal, the actual health of the channel is feeble. A significant percentage of this node's devices exhibit the same symptoms, but the customer experience is unaffected. Does this mean there isn't a problem in this node? Of course not. It does mean that the automated PMA is working very hard to maintain the quality of the customer experience but at the expense of data throughput or capacity.

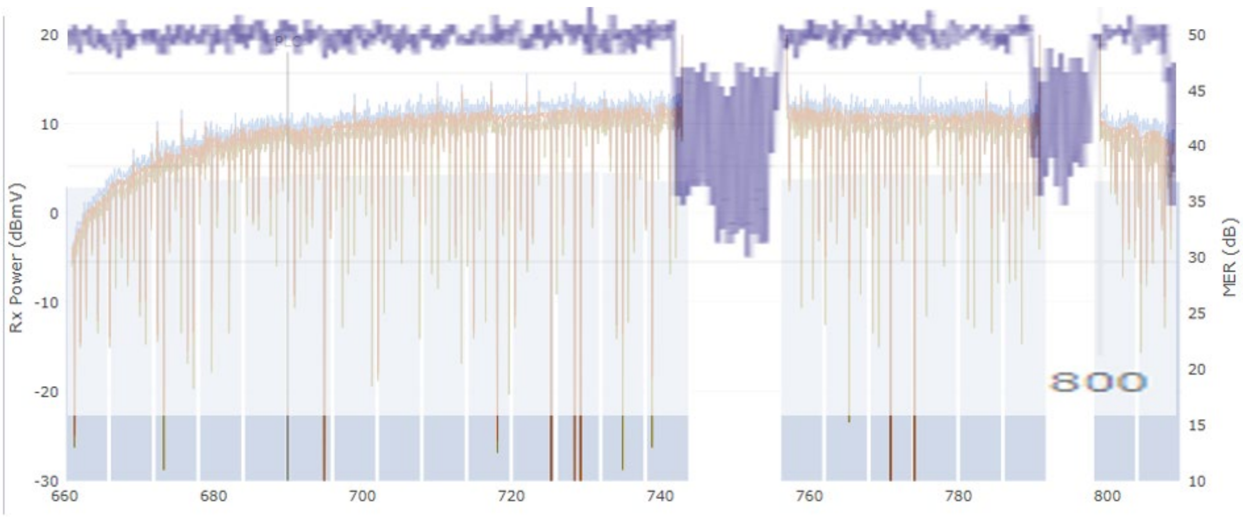

Figure 10 - Per Subcarrier RxMER of an Impaired OFDM Channel with Full Band Capture Frequency Response Superimposed

What if we measure the performance of a network based on how hard the PMA tools *are* working to mitigate the impact on the customers? What would that measurement look like? How would it impact the prioritization of impairment remediation?

9. Plant Efficiency as a Measure of Network Impairment

We have introduced efficiency and defined it as a general concept, as a technical report by CableLabs to configure DOCSIS channels for throughput and resilience, and as an initiative to understand energy consumption. We further propose that efficiency can be used as a figure of merit to define how well a network performs relative to its designed capability. Plant efficiency in this context can be defined as *the percentage at which a node or service area performs relative to its configured operational maximum*. A node with no impairments that allows all devices to use the highest orders of modulation possible on all available channels is considered to have a plant efficiency of 100%. The benefits of this figure as a measurement are numerous. While it is future-looking and leans into DOCSIS 4.0 deployments, it is also relevant to every iteration of DOCSIS deployment because it is based on the configured operational maximum at the service group or node level.

Consider a sub-split node with a 750 MHz downstream passband. Even this node has a measurable operational maximum value for the configured upstream and downstream DOCSIS channels. In this model, we can calculate the impact of impairment for every node configuration and channel plan variation, up to and including FDX/FDD, against the maximum capability of the service group or node with no impairments. Calculating plant efficiency for any node, with any channel plan variation, is based on the following parameters:

Downstream

- OFDM Channel Efficiency

•	SC-QAM Channel Efficiency

Upstream

•	SC-QAM Channel Efficiency
•	OFDMA Channel Efficiency

Using these individual components develops a spectral efficiency value for each channel type, which could further aggregate into an upstream and downstream efficiency and possibly into a single figure of merit for the cumulative node service area’s efficiency.

Considered data elements for an OFDM channel could include the channel width minus exclusion bands (MHz), channel utilization (%), channel capacity (Mbps), profile speed (Mbps), and the number of devices using each profile (count). For an OFDMA channel, the channel width - exclusion bands (MHz), channel utilization (%), channel capacity (Mbps), IUC speed (Mbps), and the number of devices using each IUC (count). For an upstream SC-QAM channel, the channel width, channel utilization (%), channel profile (Mbps); downstream SC-QAM channel (per device): the cable modem’s bonding capability against the bonded number of SC-QAM channels per modem summed across all devices in the service area (%). As a future consideration for FDX, the efficiency formula is yet to be determined for an FDX band. However, initial indications suggest that the same data for the FDX band exists for the OFDM and OFDMA channels in currently deployed states.

At the time of this paper’s writing, this method of validating network performance is manual, using ad-hoc reports and spreadsheets to validate the efficacy of the methodology. The following figures (Figure 11 and Figure 14) illustrate the allocated spectrum and the logical considerations in determining each channel type efficiency value for a currently deployed node. The example node was identified by correlating nodes with low upstream and downstream efficiency values in the OFDM and OFDMA.

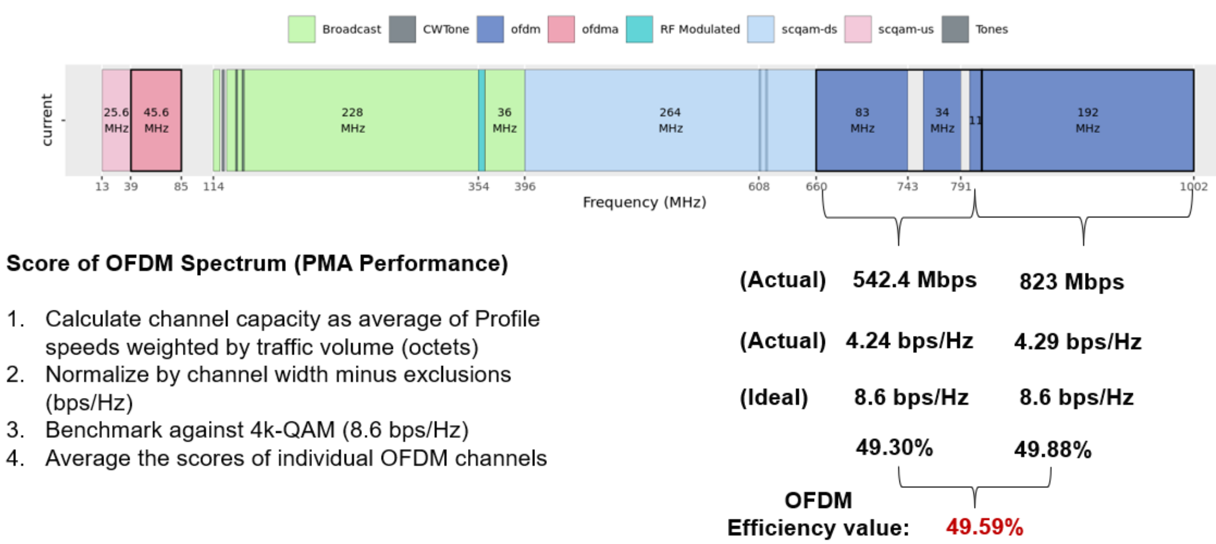

Figure 11 - Example of OFDM Channel Efficiency Value Considerations

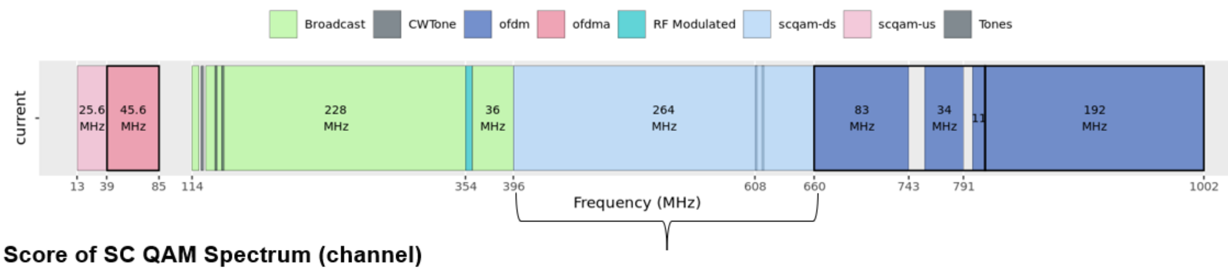

Score of SC QAM Spectrum (channel) Performance)

1. Calculate channel capacity as sum of active channels x 38.8 Mbps/6MHz
2. Benchmark against 44 SC QAM channels (6.46 bps/Hz)
3. Deduct efficiency as a percentage for any channel in a dormant (partial bonded) state

38.8 x "n" = 1707 Mbps (Ideal)
(Where "n" = 44)

6.46 bps/Hz (Actual)

98.03%

SC-QAM
Efficiency value: 98.03%

Scenario:
Node has 123 devices using a sum-total of 2290 SC QAMs.
20 devices contain a sum-total of 45 SC QAM channels in a "dormant" or partial state.
X=(((2290-45)/2290)*100)

Figure 12 - Example of Downstream SC-QAM Channel Efficiency Value Considerations

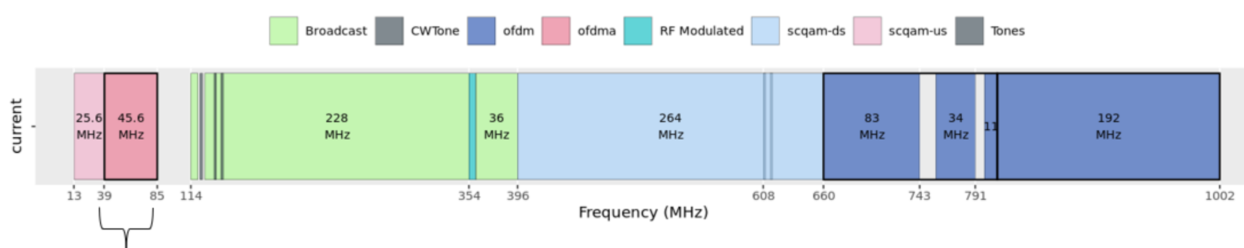

423 Mbps (Actual)

9.3 bps/Hz (Ideal)

6.5 bps/Hz (Actual)

69.8%

OFDMA
Efficiency value: 69.8%

Score of OFDMA Spectrum (PMA Performance)

1. Calculate channel capacity as average of Profile speeds weighted by traffic volume (octets)
2. Normalize by channel width minus exclusions (bps/Hz)
3. Benchmark against 4k-QAM (9.3 bps/Hz)
4. Average the scores of individual OFDMA channels

Figure 13 - Example of OFDMA Channel Efficiency Value Considerations

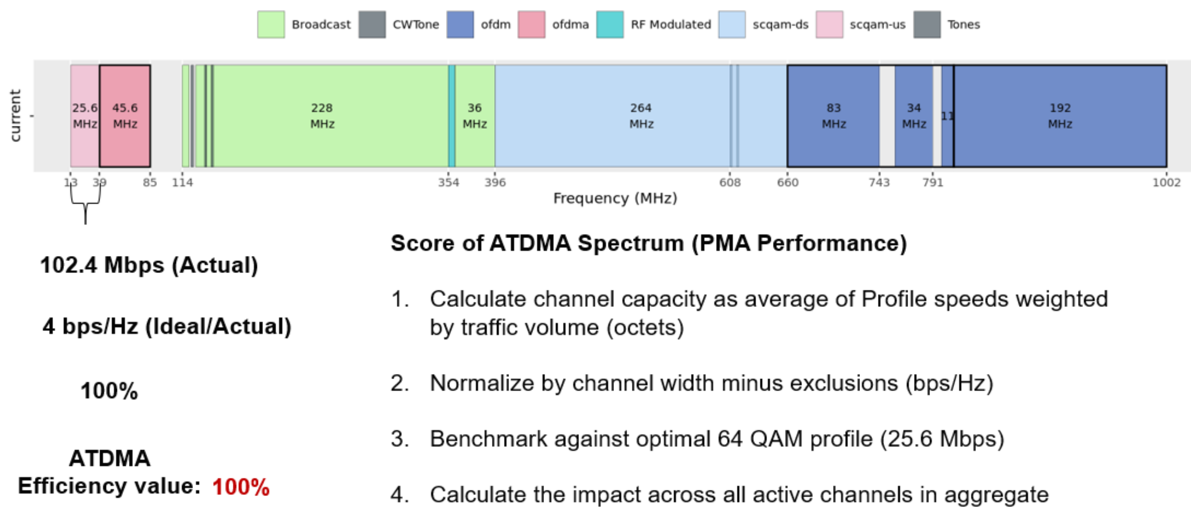

Figure 14 - Example of Upstream SC-QAM Channel Efficiency Value Considerations

All the values calculated in the previous examples came from a single node service area. The efficiency values displayed in the figures were derived using the telemetry points described and calculated in an un-weighted fashion using one week's worth of data. This example node displayed very few of the historical RF symptoms of impairment mentioned earlier in the paper. Upon inspection through the lens of efficiency, however, we recognize that there are significant signs that our PMA is working very hard to mitigate the impacts of the network impairments on the customer experience. In so doing, the spectral efficiency of our node is significantly degraded, both in the upstream direction and the downstream. We are still considering the best way to aggregate the individual efficiency values into singular upstream and downstream values and whether there is additional value to developing a method to aggregate them all into a singular value for the entire node. From an operational priority perspective, it makes sense to weigh each upstream and downstream component into a single directional value, considering the spectral efficiency of each signal type in the weighting. Signal types with more data transport potential could be weighted more highly, proportional to their optimal configured throughput.

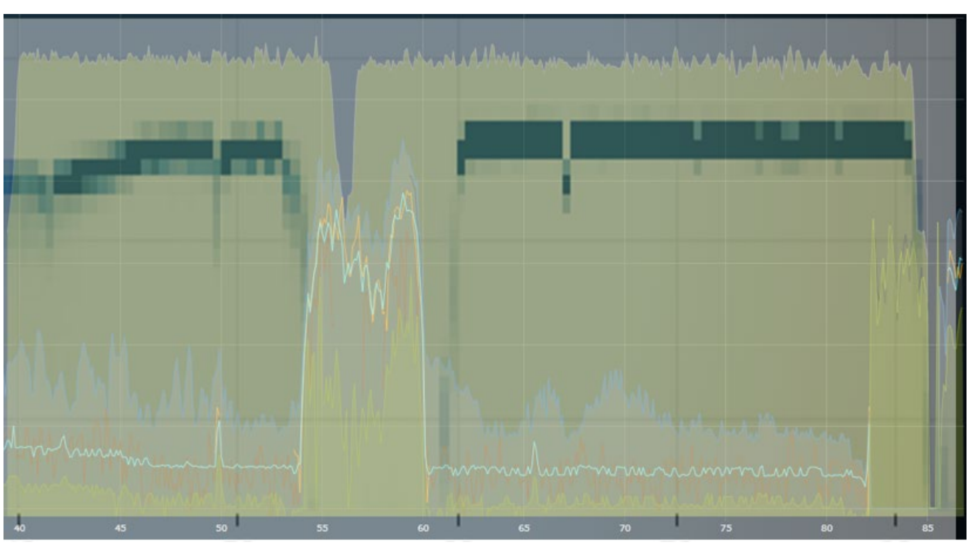

Figure 15 - An Upstream OFDMA Channel with VHF OTA Ingress and Modem RxMER Value Superimposed

While there is still value in measuring the traditional RF metrics to identify and illustrate defects in the network, leveraging these as the primary factors of repair prioritization might have diminishing returns when deploying more OFDM and OFDMA channels. As more of the spectrum is transitioned to these signal types and self-healing program automation becomes more prevalent, using plant efficiency as a primary figure of merit for prioritizing network repairs could add value to an operator's efficient workforce routing system.

9.1. Implementation Considerations

Deciding to use plant efficiency to measure network performance significantly changes how we think about network impairments. Some telemetries to consider when developing and deploying an automated prioritization system include:

- Channel Width (minus exclusion bands)
- Channel Utilization
- Channel Capacity
- Profile Speeds
- Device Profile Distribution
- IUC Speed
- Device IUC Distribution
- US SC-QAM Channel Width
- Channel Profile
- Channel Bonding Data
- DS SC-QAM Channel Bonding Data

Also critical is the ability to visualize the underlying impairment, resulting in low channel or plant efficiency. In OFDM and OFDMA channels, the ability to acquire and visualize RxMER data at the device level, signature match, and correlate geographically to device and system design maps enable the workforce to localize and isolate the primary causes of impairment that are resulting in poor performance, and low plant efficiency.

10. Business Value

For decades, we have operated in an environment that addresses network maintenance in one of two ways. First, demand or reactive maintenance identifies some measurable impact on the customer experience and requires immediate response. The second and more elusive is the preventive maintenance work. By name, it is meant to prevent something from happening. To prevent problems, we have developed standards and practices such as sweep and balance to keep the plant operating at peak performance. These standards were often time-based and had no real outcome-based measurements. These standards were born from a lack of visibility into actual network performance. Also, what does it mean for the network to perform at its peak, and how do we measure it? Service calls and customer contacts can tell us how well we address these demand issues and if we stay on top of fundamental preventive maintenance. Neither tells you definitively how the network is performing today.

With the evolution of spectrum management and PNM tools, we need to move beyond our traditional time-based preventive maintenance practices. These time-based methods served us well, but with software tools like PMA, they can cover the customer-impacting impairments much faster than we can roll a truck to resolve them. But there is a cost, and that cost is the efficiency of bits to traverse the network. Plant efficiency allows us to quantify that efficiency loss in a way that directs and prioritizes the parts of our network that need attention to restore the designed capability of the network. So now, as we look at the network, we see beyond the customer impairment and can genuinely prevent some impairment from impacting the customer. More importantly, we can quantify the repairs and the value they created for the network.

For the longest time, we struggled to identify preventive maintenance activities that we could point to and show how they create real value. Questions from finance like “What happens if we don't do this work?” were always elusive. So, we nervously stood our ground on tried-and-true practices, trying to make correlations to the larger service call or customer contact benefits. Today, in our PMA world, we can measure the impact on our network’s efficiency by mitigating customer impact. As a result, we see the value of the lost bits per second per hertz. We are giving up this real capacity on our network, which is valuable to our customers and companies. By virtue of this relationship to capacity, we can make accurate assertions of value created by preventive work to fix these impairments.

As previously shown, we can calculate the efficiency by comparing the current profile capability to the maximum profile capability represented as a percentage. So, if you are planning for capacity and have a node running at 49% efficiency but overall capacity is at 85%, we can improve overall capacity by ~50%. This has real dollars tied to it, and the total value depends on how many nodes you have at moderate to high capacity and moderate to low efficiency. For example, if you have 5% of nodes that fit this condition in a 100k node system, we have 5k nodes. If a virtual node split (segmentation or de-combining) in the same housing, the optics and labor could cost $10k to $20k. Addressing the impairments causing the low efficiency could save you ~$50M to $100M in capital across the 5k nodes. This savings is much more as we consider the costs of physical splits requiring new fiber placement and new nodes.

We can show straightforward correlations between work completed and efficiency gained. The value created through gained efficiency shows the real value of this preventive maintenance work. We can see other opportunities by using this same value on an individual subscriber to manage customers’ expected vs delivered speeds. Having this single figure of merit to evaluate efficiency opens the door to value-based, data-driven discussions on the merits of preventive maintenance.

11. Conclusion

As DOCSIS networks evolve, bandwidths expand in both directions and self-healing automation continues to mitigate the impact of impairments, measuring the spectral efficiency of every node could become increasingly important, both as a means of identifying the severity of impairments and prioritization of network remediation activities. The legacy indicators of network impairment no longer have a linear relationship to the quality of the customer experience and should be re-rationalized in the context of their impact to plant or spectral efficiency. In summary, plant efficiency could add value to a network operator’s suite of measurements not only to valuate network performance, but also to add prioritization to technician involved remediation activities.

Abbreviations

API	application programming interface
AWGN	additive white Gaussian noise
bps/Hz	bits per second per hertz
CM	cable modem
CMTS	cable modem termination system
CNR	carrier-to-noise ratio
CSO	composite second order [distortion]
CTB	composite triple beat [distortion]
CX	customer experience
DSP	digital signal processing
DOCSIS	Data-Over-Cable Service Interface Specifications
EMS	energy management systems
EPCB	energy per consumed byte
ESD	extended spectrum DOCSIS
FDD	frequency division duplexing
FDX	full duplex [DOCSIS]
FEC	forward error correction
FFT	fast Fourier transform
HFC	hybrid fiber/coax
ICFR	in-channel frequency response
IUC	interval usage code
KPIs	key performance indicators
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
OSP	outside plant
OTA	over-the-air
PMA	profile management application
RF	radio frequency
RxMER	receive modulation error ratio
SC-QAM	single carrier quadrature amplitude modulation
SEU	significant energy users
SNR	signal-to-noise ratio
VHF	very high frequency
VNF	virtual network function

Bibliography & References

[1] Sundaresan, Karthik et al., "Data-Over-Cable Service Interface Specifications DOCSIS 3.1: Improving DOCSIS Efficiency," CableLabs, CM-TR-Efficiency-V01-220627, Jun. 27, 2022.

[2] Flynn, Katie and Batchelder, Budd, "Eating the Energy Elephant One Byte at a Time: Managing Energy Consumption Across the Cable Network," SCTE Cable-Tec Expo 2023, Denver, CO, USA, Sept. 16-19, 2023

[3] SCTE, "Edge and Core Facilities Energy Metrics," SCTE, Exton, PA, USA, Tech. Rep. ANSI/SCTE 213 2020, 2020.

[4] Howald, Dr. Robert, Williams, John et al., "Developing The DOCSIS 4.0 Playbook For The Season Of 10g," SCTE Cable-Tec Expo 2021, Virtual, Oct. 11-14, 2021.

[5] "Data-Over-Cable Service Interface Specifications DOCSIS® 3.1.: Physical Layer Specification," CableLabs, CM-SP-PHYv3.1-I09_160602, Jun. 2, 2016.

[6] Harb, Maher, Ferreira, Jude et al., "A Machine learning Pipeline for D3.1 Profile Management," SCTE Cable-Tec Expo 2019, New Orleans, LA, USA, Sept. 30th - Oct. 3rd, 2019.

[7] Harb, Maher, Santangelo, Bryan, et al., "Full Scale Deployment of PMA" SCTE Cable-Tec Expo 2020, Virtual, Oct 12-15, 2020.

[8] Zhu, Jay et al., "Towards Fully Automated HFC Spectrum Management," SCTE Cable-Tec Expo 2023, Denver, CO, USA, Sept. 16-19, 2023

