DOCSIS 3.1 OVERDRIVE: DYNAMIC OPTIMIZATION USING A PROGRAMMABLE PHYSICAL LAYER

Saifur Rahman
Senior Engineer
Comcast

Jason Schnitzer
Founder & CEO
Applied Broadband

Joe Solomon
Principal Engineer
Comcast

Dr. David Early
Data Scientist
Applied Broadband

Abstract

DOCSIS 3.1 promises to dramatically increase data rates of the HFC infrastructure through the introduction of downstream Orthogonal Frequency Division Multiplexing (OFDM) and upstream Orthogonal Frequency Division Multiple Access (OFDMA).

However, in order to realize the full potential of these powerful new modulation methods, DOCSIS 3.1 channel parameters must be dynamically adjusted to optimize for the greatest possible throughput given the time varying error performance of the downstream signal path.

This paper describes an experimental software-defined profile optimizer for DOCSIS 3.1 networks. In it, we describe the theory of operation driving automated data collection, optimization analysis, and programmatic control of the DOCSIS 3.1 physical layer.

To illustrate our approach and learnings to date, we'll examine the current state of the art of D3.1 systems to support profile optimization capabilities; which form the basis of an early implementation of SDN architectural principles in cable broadband networks.

BACKGROUND

This section explores the capabilities of DOCSIS 3.1 modulation profiles and builds upon earlier works describing these protocol features [1][2][3].

Modulation Profiles

Traditionally, modulation parameters in the DOCSIS physical layer were statically configured in the CMTS and rarely, if ever, modified. A DOCSIS 3.0 (D3.0) [4] channel provides 6 MHz (North America) of data transport at a fixed modulation order for all cable modems (CMs) that can receive the channel.

In DOCSIS 3.1 (D3.1) [5], significant flexibility in the size and modulation of a channel has been introduced. Channels range in size from 24 MHz to 192 MHz in the downstream and 6.4 MHz to 96 MHz in the upstream. The D3.1 channel is separated into discrete subcarriers; each subcarrier is only 25 or 50 kHz wide, with thousands of subcarriers per channel. Each subcarrier can have a separate modulation order applied, allowing multiple modulation orders to be applied across the spectrum band. As the constellation density increases (from 256-QAM to 4096-QAM or better), the number of bits transmitted per hertz (bits/Hz) also increases. However, higher modulation orders are more susceptible to channel interference than lower modulation orders. Therefore, higher modulation orders (e.g., 4096-QAM) should be used on subcarriers in areas of spectrum where channel error performance is favorable, and less efficient modulation orders (e.g., 256-QAM) should be used on subcarriers in areas of spectrum where channel error performance is less favorable.

The modulation order applied to a subcarrier is managed through the creation of modulation profiles. Each modulation profile defines, on a per-subcarrier basis, the modulation order for all the subcarriers in that channel. A downstream channel can support up to 16 modulation profiles; an upstream channel can support 8. Why so many? Because the CMTS can assign each CM that uses that channel a different modulation profile from the set of profiles configured (to a maximum of 4 per CM). More robust modulation profiles (e.g. lower modulation order) can be used for important, low bitrate data to ensure it gets through; profiles with more aggressive modulation orders can be used to send data at the highest possible bit rate for that CM. This allows modulation profiles to be designed in a way that accounts for the downstream channel characteristics and performance at each CM. Profiles can be defined to use higher order modulations in parts of the spectrum where the signal-to-noise ratio is high, and use more robust modulation orders in parts of the spectrum where channel conditions are not as optimal. This approach optimizes the amount of data that can be transmitted and received, with a goal of providing the highest data rate possible to each individual CM.

The CMTS decides which modulation profile to apply to a channel for a given CM based on two factors: The quality of the downstream channel signal measured at the CM receiver, and the performance of the forward error correction (FEC). If the quality of the signal being received is high, the CMTS can choose a modulation profile for the CM that uses higher order modulation. If the quality of the signal is less than optimal, a modulation profile that uses more robust modulation orders can be chosen.

The CMTS also monitors the ability of the receiver to correct codeword errors using FEC. A high incidence of FEC errors may indicate an overly ambitious modulation profile; a complete lack of FEC errors suggests a higher performing modulation profile can be used.

One interesting property of D3.1 is the fact that there will be FEC correctable errors even during normal (acceptable) operation. The use of FEC coding gain to push higher modulation orders is made possible by the system’s ability to adapt to a lower bit loading profile if the coding gain can no longer provide unerrored frames. This is a key mechanism in D3.1 to provide the maximum capacity of the channel while also running at a lower operating margin.

Modulation Profile Design

In current D3.1 implementations, the CMTS does not participate in the design or creation of modulation profiles – this task is left to the operator. This is where the challenge of D3.1 optimization begins. Determining the modulation profiles that can be used by the best performing CMs and the worst performing CMs is relatively straightforward: The profile for "top tier" CMs, where plant performance is high across spectrum, uses high modulation orders (increased bit rate) across the spectrum. The profile for "bottom tier" CMs, where the plant performance is consistently poor, uses more robust modulation orders (lower bit rate). The CMTS can then assign these profiles to the top and bottom CM performers accordingly.

What about the middle performance tier, where the majority of CMs reside? Given that most service groups show a normal (Gaussian) distribution of downstream MER, setting one modulation order across all subcarriers is not operationally possible, or is suboptimal in that the lowest common MER value dictates the modulation order in that service group. Therefore, some means of dynamically adjusting modulation profiles is necessary for optimal network throughput. The number of profiles and the bit loading of the profile are based on several factors which are discussed in this paper.

To create effective and efficient modulation profiles across this population, the quality of the signal must be analyzed across the entire downstream spectrum of the channel for all CMs. Patterns in the quality of the signal across the spectrum should be identified and evaluated using criteria including:

- CM grouping: Which CMs have similar enough performance (commonality) across the spectral width of a channel to share the same modulation profile?
- Topological grouping: Are there areas of the plant (in certain geographic regions) where downstream spectral performance is not as strong as in other locations of the plant?
- Diurnality: Are there daily patterns where plant performance degrades or improves?
- Seasonality: Are the plant characteristics consistent between summer and winter, considering the different environmental factors?

Identifying these patterns in signal performance helps to design profiles that work for a large number of CMs on a given channel. Because a given CMTS serves tens of thousands CMs, there is a large amount of data to gather and analyze.

In addition to designing modulation profiles based on CM performance, we must consider the processing load imposed on the CMTS from having a large number of active profiles per channel. CMTS scheduler efficiency can be impacted by the number of profiles that are in use on a channel, especially on wide channels. What is a 'reasonable' number of profiles to have active at any given time? What is the number of profiles that a CMTS can support without impacting overall system performance?

Modulation Profile Optimization

The questions described in the previous section are difficult to answer and, given that conditions in the plant are not static, present challenges to keep the answers fresh and accurate. This necessitates a system that dynamically manages the modulation profiles that are associated with a D3.1 channel and to which CMs each are assigned. The D3.1 profile optimizer will be responsible for the following tasks:

1. Develop a set of modulation profiles for a channel that allows the CMs on that channel to use the most efficient bit loading across the channel.
2. Make recommendations for profile assignment, communicating to the CMTS the best fit profile for each CM.
3. Monitor CM performance using a given profile over time and change profile assignment recommendations, as appropriate.
4. Monitor channel performance across the entire CM population, as measured at each CM, and adjust existing profiles or create new profiles to adapt to those conditions.

This profile optimizer will continuously monitor plant and CM performance. A mature system will consider performance over time, identifying intermittent anomalies and recurring interferers. In addition to designing modulation profiles and managing their assignment in such a way that the best possible bit loading can be achieved, the profile optimizer will also be able to identify parts of the plant where operation is suboptimal, flagging issues for operational teams to investigate and fix. The system augments existing Proactive Network Maintenance (PNM) [6] toolsets, optimizing performance where there are issues while also shedding light on those issues.

The remainder of this paper focuses on work completed to date to define an experimental profile optimizer and create the mechanisms and programmatic interfaces necessary to control the management of modulation profiles dynamically. Preliminary work focuses on the downstream, as many operators are focusing on deploying D3.1 downstream channels first.

MODULATION PROFILE OPTIMIZATION USING PROGRAMMABLE NETWORK PRINCIPLES

The Role of SDN

Central to the concept of SDN is the idea of the separation of the *control plane* (how to direct traffic) from the *data plane* (forwarding of traffic according to control plane directives). The advantages of this new architectural approach have been broadly researched and documented [7][8][9], but commonly describe accelerating network innovation while enabling greater resource efficiencies within network deployments (e.g., network and system virtualization).

In a SDN configuration, the control plane asserts all influence over the configuration and state of the all data plane network elements, both intermediate (routers, switches, CMTSs) and endpoints (CM devices). The control plane realizes this via a well-defined Application Programming Interface (API) [Feamster].

Traditionally, DOCSIS CMTS implementations have contained *both* the control plane and data plane tightly coupled and embedded in a single integrated network element. However, specific management tasks such as complex D3.1 modulation profile management, taken the fullest potential of Shannon's limits, are both computationally rigorous and data intensive. [1] Removing these complex profile optimization tasks and the long term data management function from the CMTS reduces the burden on the CMTS focusing resources on its principal role as a termination system within the data forwarding plane.

Figure 1 illustrates this conceptual separation of control and data planes in the context of a D3.1 network supporting an OFDM profile optimization use case.

A second concept central to SDN is the idea of abstracting network layer complexity, like different types of devices and protocols, through a single interface. The control plane interacts with the data plane through a well-defined API. This contains the complexity of vendor-specific implementations and a plurality of diverse protocols within a software layer de-coupled from higher-level application logic.

In developing this experimental D3.1 programmable modulation profile management system, the following SDN principles were successfully implemented:

- Separation of the control plane and the data plane.
- Abstraction of network complexity in terms of both vendor-proprietary CMTS features and low level network management instrumentation through the use of a CMTS Vendor Abstraction Layer (CVAL).
- Interaction between control and data planes through the application of a programmatic RESTCONF [10] API based on a standard YANG [11] data model.

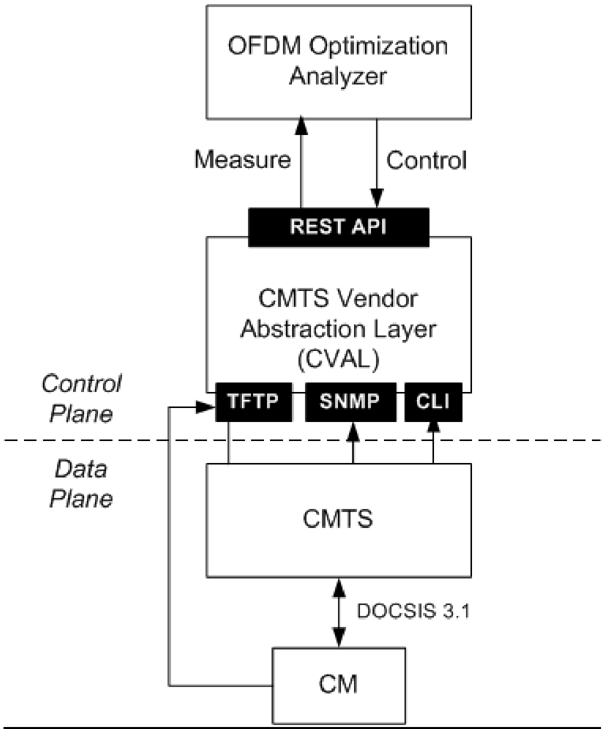

Figure 1: An experimental D3.1 profile optimizer architecture

Theory of Operation

The profile optimizer implements the key functions described in this paper in three iterative steps:

1. **Data Collection:** Key channel performance metrics are derived from network data collected for each channel and CM device.
2. **Optimization Analysis:** Per-CM downstream channel error performance is evaluated using key metrics and network optimization algorithms combined with a profile configuration policy.
3. **Profile Assignment:** Once optimization analysis indicates the need for a CM channel profile to be changed, then assignment occurs through the control plane.

Details underpinning these three conceptual steps are provided below.

Data Collection

Periodically, data is collected from the D3.1 network (both CMTS and CM) to gather key metrics. This is performed using two forms of low-level network layer instrumentation:

SNMP: Simple Network Management Protocol, using the data models described in the D3.1 Management Information Base (MIBs) defined in [5]. Though it is largely held that more contemporary network management protocols exist, SNMP remains widely adopted and universally supported by CMTS vendors. Collection of key metrics described in this paper via SNMP polling is performed on both a periodic and real time basis. SNMP is also used to configure the CM's measurement and uploading of data delivered via TFTP.

TFTP: An important tool to gather channel performance metrics is the Trivial File Transfer Protocol (TFTP), which is used to gather file-based measurement data on downstream channel performance on a per-OFDM-subcarrier basis. Though TFTP was originally defined by the IETF in 1981 [TFTP-RFC], it remains the preferred choice within the DOCSIS specification for the retrieval of CM-sourced data describing PNM file contents as defined in [5].

The data collected to derive key metrics is discussed later in this paper.

Optimization Analysis

Optimization analysis automates the evaluation of CM downstream performance based on key metrics derived from the network performance data collected. Operational thresholds defined for the range of acceptable values are compared to those calculated from measurements. If it is determined that CM's current profile is sub-optimal, the parameters of the desired profile are forwarded to the control plane for CM assignment.

Profile Assignment

CLI: The Command Line Interface (CLI) exposes vendor-specific instrumentation using a remote shell connection. In this system, the CLI is used to instruct the CMTS which profile to assign to a given CM; this interface will be used until such time that the standards-based interface is available. At the time of this paper's writing, the CableLabs PMA working group is evaluating different options for data models and protocols that together define the behavior of a standard CMTS API.

Key Metrics

To effectively design modulation profiles the health of the plant has to be understood. D3.1 defines new MIB objects to provide visibility into plant health. The primary metrics of use for the evaluation of downstream channel error performance are:

1. Receive modulation error ratio (MER)
2. Receive power
3. Forward Error Correction (FEC) statistics

MER: Due to the size of downstream D3.1 channels, an average MER reading for the channel is not an accurate measurement of a channel's health; areas of degraded performance can be hidden when averaged with measurements of good performance, making the overall performance of the channel appear to be better than it actually is. For this reason, a method to measure and record the MER for each downstream channel subcarrier has been defined [5].

To make this measurement, the CM is commanded via SNMP to start measuring the MER of each subcarrier and write that value to a file saved in NVRAM on the CM. Once complete, TFTP can be used to transfer the file to a location accessible by the profile management system. The `docsPnmCmDsOfdmRxMerTable` of the `DOCS-PNM-MIB` provides the mechanism necessary to execute this measurement. The `docsPnmBulkFileControl` object is used to upload the resulting file. The D3.1 Operations Support System Interface Specification details how these objects are implemented and used [5].

Receive Power: Downstream receive power measurements taken by the CM are available on a per-channel basis. Like the MER measurements, an average receive power measurement across the entire width of a downstream channel is limiting because it does not provide a detailed enough view into areas of the channel spectrum where receive power is below operational thresholds, possibly requiring more a robust modulation order. For this reason, D3.1 provides receive power measurements per 6 MHz block within the channel's spectral width. These measurements are available via SNMP from the `docsIf31CmDsOfdmChannelPowerTable` in the `DOCS-IF31-MIB`. The D3.1 channel is divided into 6 MHz bands and a row entry with the measured power is provided for each.

The CM can also perform downstream spectrum analysis, providing the the energy level of the signal at each frequency within a specified frequency range. The CM is signaled to begin the measurement via the `docsIf3CmtsSpectrumAnalysisMeasTable`, which was extended for D3.1. The resulting measurements are written to a file on the CM that can be uploaded using TFTP.

FEC Statistics: Downstream FEC statistics in D3.1 are maintained on a per-CM, per-modulation-profile, basis providing insight into the performance of all active profiles. For each profile, the CM keeps a running count of:

- The total number of codewords received by the CM.

- The number of codewords that were corrected by the FEC.
- The number of codewords that were not correctable.

The FEC counters are sampled on a periodic basis; the amount that each counter increases indicates how the profile is performing. D3.1 uses a new scheme based on Low-Density Parity Check (LDPC) FEC, which is a more robust error correction algorithm compared to that used in previous versions of DOCSIS (Reed-Solomon). As a result, having a high number of corrected codewords indicates that the FEC is working well and that the modulation profile is operating near its performance edge. If the number of uncorrectable code words is high, then the modulation profile is likely not performing well and should be reevaluated.

The FEC counters are provided in the `docsIf31CmtsCmUsOfdmaProfileStatusTable` of the `DOCS-IF31-MIB`. Codeword statistics can also be collected over a period of time (either a 10 minute period or a 24 hour period) on a per-profile basis and written to a file within the CM that can be uploaded via TFTP. When collecting for a 10 minute period, the CM records codeword data every second for a total of 600 measurements. When collecting for a 24 hour period, the CM records codeword data every 60 seconds for a total of 1440 measurements. Each recorded entry is time stamped to indicate when the measurement was performed. Each record contains the total number of codewords, the number of correctable codewords, and the number of uncorrectable codewords received during that interval. The `docsPnmCmDsOfdmFecTable` of the `DOCS-PNM-MIB` is used to conduct these measurements.

The profile management system could collect all of the data discussed here directly from the network layer as required. However, this data is also collected by operators in the context of other network health and maintenance systems, so the data may already be available in an existing data store, queryable by the profile optimizer. The schedule for refreshing this data needs to be considered before this data can be used, since the age of the data could impact its usefulness.

STATE OF THE ART & CURRENT LIMITATIONS

Today’s Implementations

At the present time, D3.1 systems are just beginning to become available from equipment vendors with early field deployments underway. Given the richness of the total D3.1 feature set, many operators have chosen a phased approach to deploying the new protocol’s full capabilities. The current focus of initial deployments is to increase downstream channel capacity and service speeds by introducing a single OFDM channel to the forward spectrum. OFDM channels deliver data in a more spectrally efficient way (bits/Hz), allowing MSOs to get more bits through the plant by using higher modulation orders.

With compressed time-to-market for CMTS vendors to implement the core functions of the D3.1, the industry has yet to fully focus on developing value-add features such as a full PMA support within the CMTS. Implementations currently support 3 to 4 profiles with statically set bit loading. The modulation order assigned to a subcarrier cannot be changed while the profile is active. This number of profiles should be sufficient to adequately assign groups of CMs to profiles that provide capacity gains in these early deployments. Given the other limitations noted below, an increased number of profiles would not be currently useful.

A significant limiting factor in how many profiles are useful in practice is the bit loading supported by the CMTS. Most implementations today only support a limited number of square constellations. Non-square constellations or mixed modulation orders are not supported. This limitation yields a step size of ~6bB MER between modulation orders, limiting the number of usable profiles as shown in Figure 2.

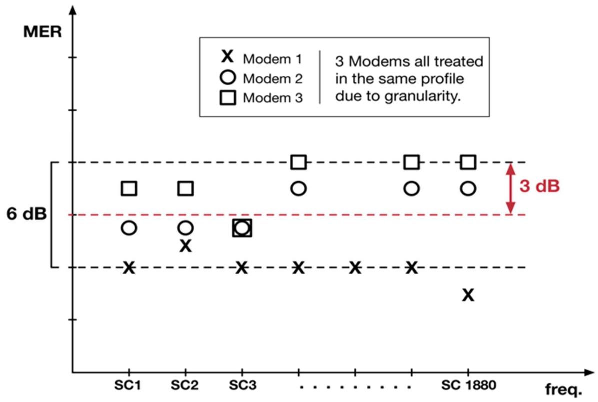

Figure 2: Limited Step Size Challenges

A second limiting factor is that the bit loading for each modulation profile cannot be assigned at the subcarrier level – implementations support 3 or 4 zones (spectral widths) of subcarriers that have the same modulation level. Techniques are under evaluation that would use a "pseudo-mixed modulation" that could produce a step of ~3dB: by using zones with different modulation orders, the average MER needed can be controlled. With this method, 4 profiles could support a total variance of 12dB in MER between CMs within a service group.

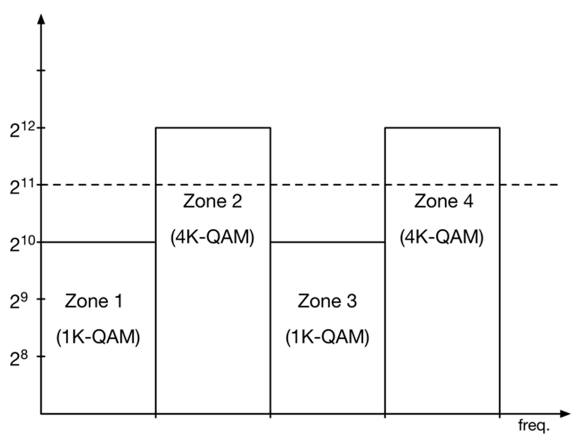

Figure 3: Using Zones to Provide Smaller Step Sizes

Possibly the largest obstacle to realizing full OFDM optimization in practice today is the relatively simple logic a CMTS uses to determine when to dynamically move a CM from one profile to another: the CMTS moves a CM if there are deteriorating RF conditions (move from a higher to a lower bit loaded profile) or if there are improving RF conditions (move from a lower to a higher bit loaded profile). In addition, current CMTSs may only support moving in one direction. This logic uses limited data, such as receiving a CM_STATUS message from the CM with a reason code of high uncorrectables. There is no trend analysis or use of prior events in determining if the status message indicates just a one-time anomaly or is indicative of a more systemic issue that requires a change; the CMTS pessimistically moves the traffic to the lower profile on receipt of the message. Current implementations do support a method to limit cycling of profiles (moving to a worse profile, then moving quickly back to the original) via simple protection mechanisms.

Support for generating a new profile at runtime based on is currently not supported. Today's D3.1 implementations support statically defined profiles that the operator preconfigures on the CMTS. There is no evaluation conducted by the CMTS to determine the suitability of the profile for existing plant conditions.

The final limitation we will discuss is the use of the test profile to optimize modulation profiles. The D3.1 specification [5] recommends that the CMTS supports at least 16 profiles for every channel and that CMs must support simultaneous reception of traffic on 4 profiles, not including the test profile. The specification allows any of the 16 profiles' bit loading configuration to change dynamically in order to better perform in varying channel conditions. The test profile must be supported by the CMTS and CMs; it allows the CMTS to use emulated test data to determine if the CM receiver can function optimally given a candidate profile.

Ideally, these protocol capabilities should be leveraged by CMTS implementations to optimize the generation of profiles, the bit loading within each profile, as well as to determine the optimal profile for each CM. To determine if a profile is suitable for use, the CMTS should continually use the test profile and/or periodically probe the receiver's MER performance. Current implementations do not support either of these methods to determine the bit loading of a profile.

As mentioned above, the bit loading of the channel in current implementations is not applied at the subcarrier level. Non-uniform bit loading in zones of the frequency band is supported instead. Whether this is approach is sufficient depends on the frequency response of the channel. Currently, operators are deploying D3.1 downstream channels in areas of spectrum that have a relatively flat (ideal) response over hundreds of subcarriers where a broader resolution of subcarrier control is sufficient. The need for finer granularity on a per-subcarrier bitloading basis becomes critical as the operating band is pushed into plant roll-off spectral regions, where there is likely to be more variability in the response due to fewer subcarriers in these regions.

Ongoing HFC plant characterizations have shown that hundreds of megabits of capacity could be gained via proper bit loading of the channel in the roll-off region of the usable downstream spectrum. See Figure 4, which shows a sharp roll-off on the right side of the plot. This will become an important feature as more spectrum is needed to offer higher tiers of service.

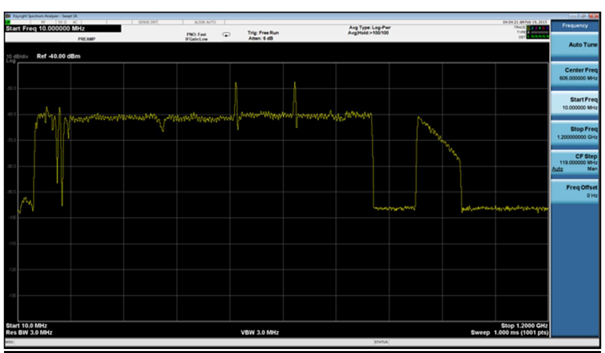

Figure 4: Bit Loading in the Roll-Off Region

Proposed Enhancements

Given these limitations of the current state-of-the-art, what things would be useful to have in the future? There are several functions needed within the CMTS, along with further development of the PMA.

The goal of an external profile optimizer is to remove the burden of optimizing plant performance from the CMTS. Instead, the CMTS could implement just the base mechanisms to react quickly to a problem with a CM and rely on an external entity to make longer-term decisions. For example, the CMTS could move the CM to a very robust profile that is known to work under the worst case MER when the modem reports FEC issues via the CM_STATUS message. Once error free performance is restored, the profile optimizer can determine a more optimal profile and provide this guidance to the CMTS. This requires a new interface between the CMTS and the profile management system. CableLabs is currently working to define new interfaces and APIs to allow such an system to interface with the CMTS.

Regardless of the level of embedded optimization capability, the CMTS will continue to play a part in the process of optimizing and managing profiles. The CMTS needs to provide information via a well-defined interface for things such as the service group to CM mapping, the profile to CM mapping, and the ability to act on suggestions from the system to change a CM's modulation profile.

Test profiles will continue to serve as a powerful method to understand how a candidate profile will perform on a CM. Again, the profile optimizer needs an interface to the CMTS in order to direct it to configure and use the test profiles. The FEC results from the test profile then need to be conveyed back to the system, where the data can be analyzed and the profile optimized. One of the main reasons the test profile is important is that it provides a direct measure of a particular bit loading configuration to support error-free performance. Since the CMTS only supports one test profile, the profile optimizer will need to communicate different bit loading settings for the test profile to allow direct measurement of the effectiveness of different candidate profiles.

Future CMTS support of mixed modulation formats, as defined in the D3.1 PHY specification [12], would improve the granularity; mixed modulation will support a step size of 1.5 dB in MER, making it possible to differentiate between groups of CMs within a service group. Mixed modulation, coupled with the support of more profiles in the CMTS, can further optimize channel capacity.

FUTURE WORK

What additional steps are needed to develop a full-featured autonomous D3.1 profile optimization system?

Optimizing CM throughput performance requires analysis of several statistics over time in addition to evolving algorithms that evaluate key metrics. Given the early stages of D3.1 deployments, limited operational knowledge exists to fully understand all the behaviors of this new modulation scheme in a modern HFC plant. Having this optimization logic executing in an external application (rather than embedded in the CMTS) allows the algorithms to be adapted more quickly as new trends are understood without having to update the CMTS software.

For example, the following behaviors need to be better understood before the logic of a complete-the-profile optimizer can be fully developed:

- How often should the re-assignment of CM profiles occur?
- What MER value ranges and margins should we use for each profile to limit the cycling of profile changes?
- How frequently should test profiles be used to probe FEC performance?
- How does MER vary from service group to service group?
- What is the frequency response of the plant (MER variation across frequency)?

Although initial estimates for each of these questions are available, more data needs to be collected to characterize the long-term RF dynamics of the HFC plant in which OFDM channels are operating. There are several key D3.1 features that can aid in this data collection and characterization. Key performance from the receiver has to be collected with enough frequency over long durations. As discussed earlier, these are the key metrics that are needed:

- FEC uncorrectable performance for both data and control channels
- Receive power levels of all channels via full band capture
- Receive MER (RxMER) per D3.1 subcarrier

In addition to SNMP-based data collection from the CMTS, data can be captured directly from the CM via periodic TFTP uploads to a data store. This allows every CM in a service group to upload downstream channel error performance data with enough frequency and over a long duration of time without adversely impacting the CMTS performance. Trends in the data can then be analyzed and the characteristic of the plant better understood. In addition to the measurements discussed earlier, the spectrum capture feature of the CM provides the RF power level at the subcarrier resolution for the entire RF band. Also available via file upload is the receiver equalizer coefficient data; this can be used to estimate tilt, ripple, and group delay.

All of this data could be housed in a data store, either external or incorporated within the profile optimizer. This rich set of data will then need to be analyzed to develop the logic of the system. Since these stats are just beginning to be analyzed and given the breadth of information, an iterative approach will be needed. The exact heuristics and algorithms in the profile optimizer will be in an ongoing state of development as more insight is gained.

CONCLUSION

In this paper we have discussed progress to date on developing a experimental system for the the optimization of modulation profiles for D3.1 networks. We have described the challenges in managing modulation profiles over a large set of CMs with a diversity of downstream channel quality performance levels. We described the use of an SDN-based approach to creating an experimental profile optimizer, and detailed the primary functions of this programmable system for use in D3.1 networks. Finally, we cataloged the obstacles that currently exist to develop and deploy such a system, given the relative maturity of D3.1 CMTS implementations. Given these limitations, we have defined areas of focus and proposed next steps for the evolution of a programmable D3.1 profile optimizer.

We see a great need for a programmable profile optimization software system as a critical component of the full D3.1 service delivery infrastructure. Without a robust system, the full potential of D3.1 will be difficult to achieve.

Work will continue to evolve the experimental implementation described in this paper in order to deploy a robust production system based on both the progression of D3.1 and on practical SDN principles.

ACRONYMNS AND ABBREVIATIONS

API – Application Programming Interface
CLI – Command Line Interface
CM – Cable Modem
CMTS – Cable Modem Termination System
CVAL – CMTS Vendor Abstraction Layer
DOCSIS – Data Over Cable Service Interface Specification
FEC – Forward Error Correction
HFC – Hybrid Fiber Coaxial
LDPC – Low Density Parity Check
MER – Modulation Error Ratio
OFDM – Orthogonal Frequency Division Multiplexing
OFDMA – Orthogonal Frequency Division Multiple Access
PMA – Profile Management Application
SDN – Software Defined Networking
SNMP – Simple Network Management Protocol
TFTP- Trivial File Transfer Protocol
YANG – Yet Another Next Generation

REFERENCES

[1] P. Sowinski. Leveraging external applications for DOCSIS 3.1 HFC plant optimization. NCTA 2014

[2] J. Chapman, The Power of DOCSIS 3.1 Downstream Profiles. NCTA, 2013.

[3] N. Pantelias .Monitoring the Variable Network: Impact of DOCSIS 3.1 on Plant Utilization Metrics. NCTA, 2015.

[4] DOCSIS® 3.0 specifications: Cable Television Laboratories, Inc. "CableLabs Specifications" [Online] Available: http:/cablelabs.com

[5] DOCSIS® 3.1 specifications: Cable Television Laboratories, Inc. "CableLabs Specifications" [Online] Available: http:/cablelabs.com

[6] DOCSIS Best Practices and Guidelines: Proactive Network Maintenance Using Pre-Equalization, CM-GL-PNMP-V02-110623. CableLabs, 2011.

[7] N. Feamster, J. Rexford, E. Zegura. The Road to SDN: An Intellectual History of Programmable Networks. IEEE ACM Queue, December 2013.

[8] Virtualization and Network Evolution Open Networking - SDN Architecture for Cable Access Networks Technical Report VNE-TR-SDN-ARCH-V01-150625, CableLabs.

[9] Bruno Nunes Astuto, Marc Mendon_ca, Xuan Nam Nguyen, Katia Obraczka, Thierry Turletti. A Survey of Software-Defined Networking: Past, Present, and Future of Programmable Networks. Communications Surveys and Tutorials, IEEE Communications Society, IEEE, 2014, 16 (3), pp.1617 - 1634.

[10] A. Bierman, M. Bjorklund, K.Watsen. RESTCONF Protocol, draft-ietf-netconf-restconf-09. Internet Engineering Task Force, December 15, 2015.

[11] M. Bjorklund, Ed.. YANG – A Data Modeling Language for the Network Configuration Protocol (NETCONF). Internet Engineering Task Force, October 2010, RFC-6020.

[12] DOCSIS 3.1 PHY Physical Layer Specification, CM-SP-PHYv3.1-I02-140320, CableLabs