The Power of DOCSIS 3.1 Downstream Profiles

John T. Chapman, CTO Cable Access BU & Cisco Fellow
jchapman@cisco.com
Cisco

Abstract

DOCSIS 3.1 will offer a more robust PHY that supports technologies such as OFDM and LDPC. With these technologies comes the opportunity to customize modulation profiles for groups of CMs to take advantage of the variation in SNR on an HFC plant.

This paper will explore the use of the modulation profiles in the downstream path.

TABLE OF CONTENTS

INTRODUCTION

Disclaimer

The ideas described in this paper are being considered to become part of DOCSIS 3.1 specifications.

DOCSIS 3.1 specifications are still under development. The following represents the author's current thoughts on what downstream profiles in DOCSIS 3.1 might look like, and do not represent actual decisions made regarding the final form of the specifications or technology.

Anything could change. Seriously.

OFDM

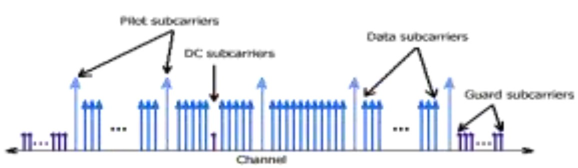

Figure 1: OFDM

As of this writing, DOCSIS 3.1 – the next version of DOCSIS – is being defined. While the ink is still drying and the mystery of what DOCSIS 3.1 will hold has yet to be revealed, one of the more certain outcomes is that DOCSIS 3.1 will move from relying on the tried and true single-carrier QAM (SC-QAM) technology for the PHY, and will incorporate OFDM – Orthogonal Frequency Division Multiplexing.

OFDM is really just a collection of very narrow QAM subcarriers as shown in Figure 1. Because they are orthogonal, these narrow QAM subcarriers can be placed very close to each other or even overlap each other, making for much more efficient use of the available spectrum.

For example, an OFDM channel might be as large as 192 MHz and contain 3840 QAM subcarriers, spaced every 50 kHz, or even 7680 carriers, spaced every 25 kHz.

OFDM allows each subcarrier to have its own modulation value and amplitude (or none at all in the case of a muted subcarrier). It is this interesting property that allows OFDM to fit a downstream transmission path like a glove.

The description of the modulation type for each subcarrier and its amplitude value is stored in a data structure simply referred to as a profile.

Once you can tailor such a glove, just how many gloves, or profiles, are needed for a downstream HFC plant? How are these profiles managed?

LDPC

A second part of the DOCSIS 3.1 story that is interesting is the adoption of LDPC (Low Density Parity Check). LDPC is a FEC (Forward Error Correction) technique that can be be used to correct bit errors due to noise. LDPC is more powerful than its DOCSIS predecessor Reed-Solomon. As such, LDPC is a major factor in allowing higher order modulation to be used.

So while LDPC allow a great density of bits per hertz, OFDM allows for a better customization of the frequency spectrum and noise mitigation.

DEFINING THE NEED

Is One Profile Good Enough?

The answer to this question is not obvious. In fact, there are proponents on both sides of this debate.

The argument for one profile would be to pick a modulation value that is good enough. For example, today's plant is almost entirely 256-QAM. There is rarely a need to do anything more or less. Some operators will put 64-QAM in the roll-off regions of the HFC plant, but that is more of an exception rather than a rule.

Thus, in a HFC plant with OFDM and an improved FEC, the new norm could be 1024-QAM. To enforce that, any part of the plant that could not support 1024-QAM with the new CMs would not be upgraded until it did support it. In such a scenario, no more – and no less – would be needed.

The other side of the argument is that either the plant cannot or is not always upgraded everywhere, or that even with a well-maintained plant, there is enough natural variation is SNR (signal-to-noise ratio) that extra performance can be squeezed out.

Axiom 1

It is easier to downgrade than upgrade.

If some of the CMs in the HFC plant are not working well, either due to interference by noise-like carriers such as digital TV ingress, or due to an insufficient SNR margin, the profile can be modified to accommodate and thus alleviate the problem. This action would result in a profile that has lower throughput but higher reliability.

But what if the system is already working? Hypothetically, the profile can be upgraded as well. In practice, if upgrading a profile is not done right, then CMs may drop off. This improvement often flies in the face of "if it ain't broke, don't fix it."

The likely outcome is channel degradation over time that is hard to recover from.

The take-away is that if profiles are to be updated, then a robust system for measuring, testing, and deploying profiles is needed. Such a system would benefit from the ability to create additional profiles that can be tested before deployment.

Axiom 2

The decision to sort is easier than the decision to deny or downgrade.

For a single profile system, when a CM is in a bad part of the HFC plant, the choice is to downgrade everyone or deny service to that particular CM. When modifying an OFDM profile in a system with only one profile, <u>all</u> CMs are impacted by definition. So, if one CM is having problems, scenarios could exist where fixing the service for one CM may break the service for other CMs.

With a multi-profile system, unlucky CMs can be moved to a lower throughput profile without disturbing the rest of the CMs. In a system with multiple profiles, CMs that are in trouble can be moved to isolated profiles, or to profiles that have a lower throughput and higher robustness.

The OFDM Paradox

To truly take advantage of OFDM, there needs to be a channel profile that can be changed over time to allow the system to learn and adapt to channel conditions.

Changing that profile requires good decision-making. The challenge is that the feature that potentially makes OFDM work better – the ability to optimize a profile – can also make it work worse due to the decision-making process coupled with operational realities.

For a single-profile downstream, when one CM has a channel problem, the impact of the solution to that problem (changing the profile) is directly felt by all other CMs. For a multi-profile system, when one CM has a problem, the solution (moving the CM to a different profile) does not directly impact the remaining CMs.

A single-profile system works by providing the worst service to all CMs.

A multi-profile system works by providing the best service to all CMs.

This paper will pursue this latter line of thinking and examine how to manage multiple profiles.

New Spectrum Opportunities

With advances in optical and RF equipment, it is now possible to extend the spectrum of the plant further in both the upstream and downstream. This is of particular interest when more spectrum is needed and where higher throughput is required.

In these new areas of spectrum, the HFC plant may have more variation or more micro-reflections, and other impairments where OFDM and LDPC may prove quite useful.

Below cutoff:

Below the downstream cutoff frequency (750 MHz, 1 GHz), the plant is generally well engineered. Even so, the amount of plant SNR differences within a downstream can be as much as 8 to 12 dB. This may allow for higher order modulation for CMs that are in areas of high SNR, such as homes after the first amplifier, compared to homes after the last amplifier where there is higher culmative noise.

Above cutoff:

To operate above the cutoff frequency, the plant will require upgraded amplifiers and optical nodes. This effectively moves the cutoff frequency higher. The only element that does not get upgraded is the taps. Most new taps have a viable frequency response up at least to 1 GHz and often past 1.2 GHz. Thus, OFDM would be very useful for a service from 1 GHz to 1.2 GHz where tap performance will cause channel variations.

QUANTIFYING THE VALUE

Channel Study

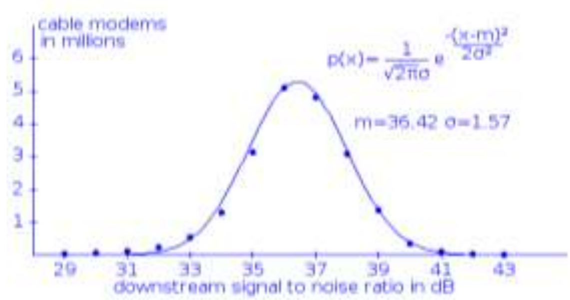

Figure 2: SNR Variation in an HFC Downstream

A study conducted by Comcast [1] and shown in Figure 2 measured the downstream

SNR for a large population of CMs on an HFC plant. As a general conclusion of this study, the consensus was that in an average HFC plant, there would be an 8 dB variation in SNR between various CMs.

The primary contribution to this variation has to do with the location of the CM along the HFC plant (topology constraints). CMs that connect to the coax segment nearest to the optical node and before the first amplifier see higher SNR values. A CM located at the end of the longest coax link after the last amplifier (typically 5 amps but could be more) typically has lower SNR.

With different profiles, different groups of CMs could be provided with a higher order modulation (where SNR is higher) while other CMs are provided with a lower order modulation (where SNR is lower).

A basic example use of profiles to study would be the following:

- Profile A: 256-QAM (2%)
- Profile B: 1024-QAM (25%)
- Profile C: 2048-QAM (64%)
- Profile D: 4096-QAM (9%)

In this model, the predominant modulation is shown. The model is also hierarchical where each level is progressively better than the previous level.

In parenthesis is the percentage of CMs from the study [1] whose maximum receive capability matched one of the profiles.

So, if a single profile were used, then all CMs would have to use the 256-QAM profile or the 1024-QAM profile (if the plant were upgraded to eliminate 256-QAM). With multiple profiles, over two-thirds of the CMs can have profiles that exceed the baseline profile.

Efficiency Analysis

OFDM Profile (percent of tones at each modulation level)

	QAM Levels					CM
Profile	4096	2048	1024	256	unused	Distribution
D	100%				0%	9.0%
C		100%			0%	64.0%
B			100%		0%	24.5%
A				100%	0%	0.0%
					non-D3.1:	2.5%

Throughput with a Spectrum of: 192 MHz

Profile	4096	2048	1024	256	Total Mbps
D	2,304				2,304
C		2,112			2,112
B			1,920		1,920
A				1,536	1,536

Aggregate Raw Throughput (Mbps):	2,029
Spread Highest to Lowest:	50%
Increase compared to worst-case channel:	32%

Figure 3: Throughput with Pure OFDM Profiles

OFDM Profile (percent of tones at each modulation level)

	QAM Levels					CM
Profile	4096	2048	1024	256	unused	Distribution
D	60%	25%	10%	5%	0%	9.0%
C		60%	30%	10%	0%	64.0%
B			80%	20%	0%	24.5%
A				100%	0%	0.0%
					non-D3.1:	2.5%

Throughput with a Spectrum of: 192 MHz

Profile	4096	2048	1024	256	Total Mbps
D	1,382	528	192	77	2,179
C		1,267	576	154	1,997
B			1,536	307	1,843
A				1,536	1,536

Aggregate Raw Throughput (Mbps):	1,926
Spread Highest to Lowest:	42%
Increase compared to worst-case channel:	25%

Figure 4: Throughput with a Blended OFDM Profiles

The data in Figure 3 and Figure 4 assume four profiles where each profile has a different percentage of bit loading. Figure 3 uses the rule where all subcarriers use the same modulation whereas Figure 4 modifies the modulation of each subcarrier to match the SNR of the plant at that frequency. Then a percentage of the total CM population is assigned to each profile. (Input data values in the study are not measured values)

Note that a higher average throughput is maintained even though a notable percentage of CMs are in lower performing plant segments.

As an example, a single profile implementation may have had to go down to 1,536 Mbps throughput whereas a multi-profile implementation was able to achieve 1,926 Mbps aggregate throughput. That is over a 25% improvement in raw throughput.

BUILDING THE MECHANISM

Codeword Builder

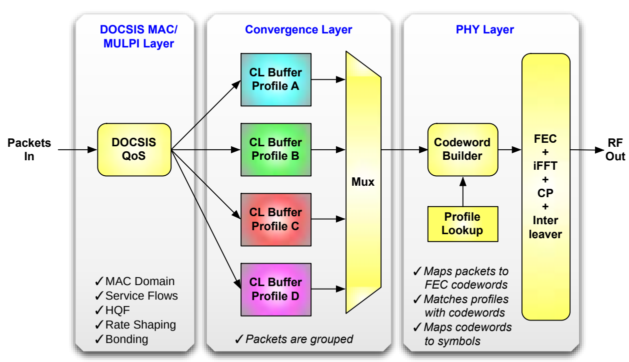

Figure 5: Codeword Builder

The OFDM channel generates a bit stream to the convergence layer. The convergence layer groups these bits into FEC codewords. DOCSIS encapsulated packets are then placed into the payload of these codewords.

Each path, as designated by a profile, is multiplexed at the codeword level. That is, each consecutive codeword can belong to a different profile.

An example implementation is shown in Figure 5. Packets exit the DOCSIS QoS process. As part of this process, they receive a tag that indicates what profile they are associated with.

The codeword builder collects these packets and sorts them into shallow collection buffers depending upon their tag. These buffers allow the PHY to accumulate enough bytes to fill a codeword. Packets are mapped directly to codewords and may be split across codeword boundaries.

A full-length codeword might have a payload of 1777 bytes and a total size of 2025 bytes. Thus, each codeword can contain on the order of one full length DOCSIS frame.

Latency

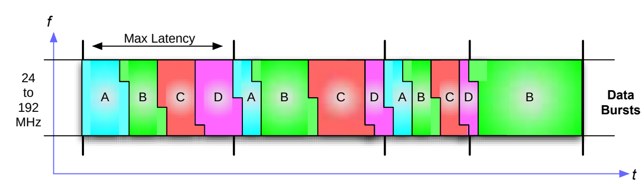

Figure 6: Profiles in the Downstream Path

The codeword builder can multiplex codewords in any manner it wants. This gives it the freedom to dynamically allocate data capacity to each profile depending upon the packet arrival rate.

The codeword builder does not actually care about data capacity which is pre-calculated by the rate-shaping mechanism in the DOCSIS MAC layer. Rather, the codeword builder is focused on minimizing latency.

Latency is dependant upon how often a path can be serviced. So, if a basic round-robin scheduler were used, as in A-B-C-D, then each channel would have to wait at least 4 codewords. This is shown in Figure 6.

If one path, say path B, has a higher bandwidth demand, the codeword builder may schedule A-B-C-B-D-B where the B path is serviced every second codeword.

Let's calculate how long it takes to send a codeword. Then we can figure out typical latency in terms of codewords. Assuming:

- 24 MHz min OFDM channel size
- 192 MHz max OFDM channel size
- 1024-QAM (10 bits/s/hertz)
- 16200 bit (2025 byte) codeword

The latency is calculated by:

latency = (codeword bits) / (bits/s/Hz) / BW

Then the time for one codeword varies from 8.4 µs minimum to 67.5 µs maximum. For four profiles at one codeword each, the latency per profile would be 34 µs and 270 µs. This latency would be in addition to any latency created by the downstream interleaver.

This leads to the following observations:

1. *There is a trade-off between the max number of profiles and latency.*
2. *This trade-off depends upon the scheduling algorithm of the codeword builder.*
3. *Larger channels provide lower latency and could accommodate more profiles for a given latency budget.*

SYSTEM OPERATION

Details

Profiles and OFDM Channels

Let's review the definition of a profile. First, an actual profile contains the dynamic configuration values for an OFDM channel. This configuration information includes modulation level for each subcarrier. A zero modulation level would imply a muted sub-carrier.

For convenience, each profile is assigned a letter. So, there will be Profile A, Profile B, etc.

Each downstream OFDM channel will have a range size, such as a minimum RF bandwidth of 24 MHz to a maximum of 192 MHz. Thus, the whole downstream 1 GHz spectrum could be covered with about five OFDM channels.

Each OFDM channel has its own unique set of profiles.

Thus, OFDM channel 1, Profile A is different than OFDM channel 2, Profile A. The reason for this is simple. Profiles describe subcarriers at a particular frequency, and each OFDM channel occupies different frequencies.

Profiles and Paths

From a packet forwarding perspective, each profile creates a unique path through an OFDM channel from the CMTS to the CM.

Packets travel paths. That path is described by a profile.

Thus, the two words “profile” and “path” often get used interchangeably when talking about how packets are sent from a CMTS to a CM across an HFC plant.

The term “profile” is a PHY level description; the term “path” is a MAC level description.

Multiple Paths

A CM may receive packets on more than one path. Thus, there are multiple forwarding paths, each with its own profile, from the CMTS to the CM.

Not all CMs have to receive on all profiles. One CM may receive profiles A and B while another CM may receive profiles A and C.

The CMTS will keep track of all the active paths to all the CMs. The CMTS will have to manage these multiple paths and determine through policy and forwarding rules what packets are to be placed on what path.

Path Assignment

The rules for profile usage are simple. The CMTS may use any path/profile for any task. The CM must accept anything addressed to it by any active path it may have.

For example, Profile A could be the common profile that all CMs can receive. It would be used for booting since the performance characteristics of new CMs are not known. Profile A could also be used for most or all MAC Management Messages (MMM), since it is by design the most robust and common profile.

However, the CM must be able to accept a MMM on any profile. This actually keeps things simple for the CM. The CM uses multiple paths and profiles to ensure packets get delivered through the PHY. It then uses MAC filtering to determine what to do with the packets.

This rule set also allows policies to change over time at the CMTS without impacting the implementation of the CMs in the field.

Service Flows

A service flow (SF) is a collection of packets that match a classifier. A typical classifier consists of the source and destination IP address, source and destination port, and the protocol type (this is the classic “tuple”).

The most common SF is the default SF that contains anything not in a specific SF. Voice over IP usually has its own SF per voice call. Managed Video over IP will likely also have its own SF per video flow.

The convention followed will be that all SFs are fully contained within a profile for a given OFDM channel. That means that if a video flow is assigned to profile C in a given OFDM channel, then all packets in the flow remain on profile C.

Bonding

Since SFs are assigned to a single profile, there is no need to bond across profiles within an OFDM channel.

Bonding can be used to allow a SF to traverse multiple OFDM channels. For example, if a 5 Gbps SF was desired, that SF might be mapped across three OFDM channels.

Within each OFDM channel, a profile would be chosen. Since profiles are unique per OFDM channel, it could be any profile that the CMTS chooses. So, in this example, the three profiles across each of the three OFDM channels could be B-B-B or B-C-D.

Bandwidth Management

MAC or PHY?

So where does the decision get made as to what profile a packet should belong? Where does the bandwidth get calculated?

The MAC does all the bandwidth management. The MAC figures out if the packets will fit in the OFDM channel. The PHY packages up the packets and sends them out.

Thus, in theory, the packet buffers in Figure 5 should never overflow because the MAC will only send the packets to the PHY if there is room for them to hit the wire. That allows these buffers to be shallow, low cost, and low latency.

Channel Capacity

The aggregate capacity of the OFDM channel depends upon the configuration of each of the profiles contained in that OFDM channel and the relative usage of each profile.

There is no limit set on how much data capacity a profile may have. Any profile could be assigned 100% of the OFDM channel capacity at any time. The only rule is that all the profiles share the aggregate channel capacity. They can share that data capacity anyway that the CMTS sees fit.

If a profile is updated, then the channel capacity will change. This is a slow change since profiles are not updated frequently. However, when a profile is updated, the instant in time of the update to the rate-shaping mechanism must be synchronized to the instant in time of the update of the profile.

Time-Varying Channel Capacity

To calculate the throughput of an OFDM channel, you need to know the RF bandwidth and its modulation order or bit density.

This presents a challenge. Each profile has a different modulation order definition. Further, the amount of packets per profile depends entirely upon how many CMs are on each profile, the CMTS forwarding policy for packets, and the instantaneous traffic profile per CM.

That means that the bandwidth per profile is time-varying. It follows that the total channel bandwidth added across all profiles is also time-varying. In fact, if the modulation profiles ranged from:

- Profile A: 256-QAM (8 bits/s/Hz),
- Profile D: 4096-QAM (12 bits/s/Hz),

then the data capacity could vary as much as 50% at any point in time.

The MAC downstream rate-shaping mechanism must take this into account. To do this, the MAC downstream scheduler keeps track of which profile each packet has been assigned to and what the average bandwidth per profile is. It then calculates how much bandwidth is available on an instantaneous basis.

This is different than in DOCSIS 3.0. In all the early versions of DOCSIS, the max data capacity of the physical channel was constant. Now it is not.

Frequency-Varying Channel Capacity

The previous time varying calculation used the average throughput of a profile. But what if the modulation values of the subcarriers are different over frequency? For example, in the lower half of an OFDM channel, a profile may use 4096-QAM but in the upper half it may have dropped back to 1024-QAM.

From this, an average channel capacity can be calculated. But that average is only valid if the distribution of payload bits within a path is uniform across the entire channel. The channel is up to 5 codewords per symbol wide. If the profile definition is changed per codeword then it depends upon where that codeword lands within the channel and thus as to what the actual resulting data capacity of the profile is.

The difference between the predicted bandwidth by the rate-shaper based upon the average profile data capacity and the actual data capacity that results can be considered an error factor. This error can cause a slow drift in buffer levels. If there is less actual data capacity than calculated, then the PHY buffers may build up.

In practice, there should be enough randomness and buffering in the system that this averages out. A quick fix is to allow for this error in the rate-shaping calculations by subtracting out a few percentage points from the total data capacity. A good system design should also have a way of checking and correcting for this error.

Profile Management

How are profiles managed? When do profiles get updated? How do they get updated? Which CM goes to which profile?

For a profile to be useful, it should have the following measurable characteristics:

1. *Each profile should have a measurable and significant difference from another profile.*

2. *Each profile should serve a measurable and significant number of CMs.*

As the number of profiles increase, the system complexity tends to increase. Profiles are meant to be used sparingly and they should produce measurable and tangible results.

Static Profiles

The simplest solution would be to set all modulations levels within a profile, and perhaps even all amplitude levels, to the same value. For the small percentage of subcarriers that get into trouble, the LDPC feature of DOCSIS 3.1 will correct for a certain amount of errors.

For example, Profile A could be 256-QAM for all subcarriers. Profile B could be 1024-QAM for all subcarriers, and so on.

This approach is simple. It also creates and guarantees a strict hierarchy among profiles. Each profile will have a defined performance. If a CM does not work on one profile, it can be downgraded to a lower profile until it does work.

This is a good default operation mode and one that is predictable enough that it can be deployed with good success.

Dynamic Profiles

Despite the fact that static profiles will work, the flexibility of OFDM is almost completely lost. It seems like with a good set of test and measurement tools, these profiles could be updated, and performance could be optimized. Further, actual field problems that cause trouble tickets and truck rolls could be found and adjusted for automatically.

This is a tricky business. As mentioned before, it is easier to downgrade a profile than improve it. Thus, the algorithm should make sure that profiles do not degrade over time.

The algorithm also has to decide what to do when a CM is in trouble. Should it move the CM to a different profile, or should it update the profile?

One answer is that if the trouble is a single CM, then the CM gets moved. If the trouble is seen across a large number of CMs within a group, then the profile should get updated.

This leads to questions like:

- How much time should be used to detect common errors across CMs? The longer the measurement time, the more CMs with common errors that may be discovered. However, it also takes longer to fix the problem.

- How many CMs constitute a significant group to where the profile would get updated instead of moving CMs? Five? 10? 50?

- What is the nature of the problem? Is it frequency specific or is it channel specific? A clever algorithm must have the right information in order to make a judgment call.

If there is specific interference, such as interference from LTE, the best approach may be the nulling or lowering of modulation on specific subcarriers. If the problem is across the entire channel, then the SNR that a particular CM sees may have changed and the profile is no longer valid for that CM.

Measure and Sort

An ideal system with dynamic profiles would have the following capabilities:

1. The ability to measure the SNR for each subcarrier at each CM. This could be as many as 7680 measurements, one for each subcarrier for a 192 MHz OFDM channel with 25 kHz spacing. For CMs with multiple channels, there could be 2x to 5x the number of measurements (up to 38,400 values per downstream).

   This measurement would occur at boot time and then again at scheduled intervals or when there is a problem.

   Note that many of the measured SNR values will be the same. Thus, the resulting measurement data can be easily compressed prior to transmission.

2. The ability to predict from the SNR measurements what profile to assign a CM to.

3. The ability to test how well the CM performs on that profile.

4. The ability to analyze error scenarios for root cause.

5. The ability to fix those problems either by adjusting the profile or moving the CM.

6. The ability to do all this and yield a system that is efficient yet is also stable and reliable.

CHOOSING THE NUMBER OF PROFILES

Earlier, it was discussed that fewer profiles could provide a simpler system with lower latency. In this section, we will discuss use cases that impact the number of profiles required. We will then reconcile these two needs with a proposal.

The question is really separate for the CMTS and CM. A CM only needs to know about profiles that impact the CM, whereas the CMTS needs to know about all the profiles that impact all CMs. Thus, the CMTS probably has to support more profiles than a CM does.

HFC Plant Variations

MMM and Data

MAC management messages can be either multicast or unicast. Either way, it is important that they be delivered in a robust manner. The MMM traffic is also low bandwidth, allowing the MMM to be sent on a common profile like profile A.

Data is largely unicast and is thus sent directly to a single CM. It makes sense therefore to send data on the highest profile available to the CMTS and CM.

This requirement suggests that the CM must support at least two profiles.

Penalty Box

Sometimes there is a change in the HFC plant or even a home coax network that causes the SNR to be enough out of whack that the CM will either not work or partially work with lots of errors.

The right answer is a truck roll to fix the problem. However, prior to that truck roll, would it be possible, and of value, to offer an error-free lower speed of service to that subscriber until the problem can be fixed?

That is the concept of the penalty box. It would be a lower-order profile into which a CM would be sorted to when it cannot work on the higher order profiles. Once the CM is in the penalty box, the cable operator would be alerted so that the problem could be pro-actively fixed before the customer complained.

In the earlier example, the penalty box profile would be satisfied by the 256-QAM profile since the expectation is that the plant should work with 1024-QAM. However, in some cases, it might be necessary to add, say, an additional 64-QAM profile.

The penalty box may require support of an extra profile on the CMTS.

Modulation

In the generic example in this paper, four modulation levels were suggested. These were 256-QAM, (512-QAM skipped), 1024-QAM, 2048-QAM, and 4096-QAM. Each increase in modulation order (each extra bit/s/Hz) requires an extra 3 dB of SNR. Thus, this range of four modulation orders covers a 12 dB variation in plant. That is 50% more than the target 8 dB variation mentioned earlier and thus should be enough. Note that four distinct profiles could be defined with only a 9 dB variation in SNR if no modulation orders were skipped.

Inevitably, questions arise about also supporting 512-QAM. Additionally, there may also be future modulation orders such as 8192-QAM and 16384-QAM. These latter modulation orders may only work for short passive coax plants with home gateways.

Despite these additional modulation levels, it does not seem realistic that every modulation level would be needed within a particular fiber node's service area.

This requirement suggests that the CMTS must support at least four profiles and potentially a few more.

Geographical Differences

Since an HFC plant spans a geographical area, it is feasible that one part of a plant may be subject to interference that another part of the plant is not. This would have to do with the physical location of an interference source and how it couples into the HFC plant.

Let's say that a plant spans the east and west side of an area. Let's assume that there is an interference source on the east side of the plant that does not exist on the west side. Thus, the west side of the plant may have an average SNR that is, say, 6 dB higher than the east side.

Should there be separate profiles for the west side and east side? Well, it depends upon the fix.

If the fix is to not null out subcarriers and instead just run a lower modulation order, separate profiles are not needed. What may be needed is a large range of profiles at the CMTS.

If the fix is to null out a subset of subcarriers on frequencies where the interference is, then there may be benefit of using separate sets of profiles for the west side and east side.

The system would have to have an accurate way of measuring the interferers and of determining a set of suppressed carriers that works for all the CMs.

A simpler solution may be to use one set of profiles, but null out the subcarriers so that both the west and east sides are impacted equally.

This requirement suggests that if the CMTS could support more profiles, it could help with plant management. This would not impact the CM since the CM is on a specific plant segment.

IP Multicast

For IP multicast to work, all the CMs that are subscribed to a common multicast must be on the same path. That means a common profile.

The easiest solution would be to put multicast on the default profile. In this case, profile A would be used. By plan, all CMs can receive profile A, so multicast will always work.

This has the negative effect of multicast running slower than the unicast flows that exist on higher order modulation profiles. If IP multicast is not going to be deployed in any serious amount, then perhaps this is okay. At the same time, it becomes a self-fulfilling prophecy. If IP multicast is designed to always run slower than IP unicast, it is less likely to get used.

Since CMs are capable of receiving multiple profiles, a more clever algorithm could assign the IP multicast stream to the highest order profile that is currently shared among CMs. So if CM 1 supported profiles A-B-C-D and CM 2 supported profiles A-B-C, the multicast could be put on profile C. If a third CM that only supported profiles A-B joined, then the CMTS has the option of moving the multicast to profile B, or duplicating the multicast on profile B.

This requirement suggests that if the CM can support multiple profiles, the CMTS has complete flexibility in implementing multicast, and there is a lower or even no throughput penalty for using IP multicast.

Profile Update

How do you change a profile without risking a large disruption on the network if something goes wrong?

Whole Profile Update

The easiest way to update a profile is to issue a new profile and then hit the update bit on the profile descriptor to tell the CMs to start using the updated profile.

A more cautious way would be to change a small set of subcarriers at a time. The idea would be that the change would be minor enough that the CM FEC would still correct it, but the CMTS could discover that the FEC had to correct errors.

Migrate to a New Profile

Another approach would be to establish the new profile and then move CMs over to it over a period of time while measuring performance.

This requirement suggests that the CMTS and CM would need to support a spare profile.

Profile Changing

When a CM changes to a new profile, there is always a chance that it will miss some packets that got stuck in a CMTS queue. There are two basic solutions to this problem.

Pause / Play

Similar to the RF channel change techniques in DOCSIS 2.0, one approach for a profile change is for the CMTS to shut down the flow of packets to the CM, move the CM to a new profile, and then restart the flow of packets once the CM has properly synced to the new profile.

This design approach was taken because the DOCSIS 2.0 CM could only receive one downstream RF channel at a time.

This is an undesirable solution because CMTSs do not like to halt the flow of packets. Halting the packets of an unknown flow requires buffering. Buffering can overflow.

Add / Drop

A more elegant solution would be for the CM to first add the CM to the new profile so that it is receiving the old and new profiles. The CMTS would then move the traffic from the old profile to the new profile. Once this has occurred, and the buffers have emptied, then the CMTS can remove the CM from the old profile.

This technique may require sequencing of packets to ensure that packets are not delivered out-of-order.

This requirement would indicate that the CM supports one extra profile.

Probing

How can it be determined that a CM can be moved to a higher profile?

Predict

If the SNR of the CM is known, then it can be predicted what modulation level the CM can support.

The problem with this approach is that there are really 7680 SNR levels per OFDM channel. That is, the SNR varies with frequency. In addition, the LDPC FEC is very good at fixing problems where there are a small number of subcarriers that have SNR that is too low for the assigned modulation.

Thus, a good algorithm is required that can look at the average SNR and take into the account the strength of the FEC when making decisions.

Test & Measure

A more reliable technique would be to get the CM operational on a known profile, say, profile A. Then, in the background, tell the CM to receive packets on another profile, say, profile B, and to measure and report FEC and CRC error statistics.

The other profile could even carry a specific bit pattern that the CMTS generates so that the measured results are statistically accurate.

This requirement suggests that the CM needs to support at least one extra profile.

The 4+1 Proposal

Note that supporting more profiles does not mean that there are increased throughput requirements, since even one profile could be set to receive the full channel bandwidth. Rather, multiple profile support is about feature support.

It can be demonstrated that there are many reasons why it is useful for the CM to support more than two profiles. Many features require the support of at least three profiles at the CM. Simultaneous support of several features, and the desire to allow features to operate independently suggest an even higher number.

One proposal would be to support four active profiles and one spare profile. The spare profile could be used for temporary things like profile changes, probes, etc., while the other four would allow full, simple, and reasonable support of multicast.

The spare profile is like a spare tire. It is used only when changing something.

SUMMARY

OFDM allows for customization of a frequency profile to allow for optimum throughput. This profile could use the optimum modulation order for each subcarrier and would mute subcarriers where there was interference. LDPC is an error correction technology that effectively allows even higher orders of modulation to be used and allows for an “average” SNR to be used for setting modulation levels.

Despite the higher degree of flexibility that OFDM has, the HFC plant is fairly well characterized and maintained, and the amount of variation among CMs is minimized. Thus, a small number of profiles, such as four, may be sufficient to allow a compromise between system optimization and retaining simplicity.

REFERENCES

[1] Dave Urban, "The Case for Multiple MCS," **P8023.bn EPoC PHY Task Force**, 2012, http://www.ieee802.org/3/bn/public/adhoc_mmp/adhoc_mmp_urban_01.pdf