Improving Upstream Efficiency

A Technical Paper prepared for SCTE by

Karthik Sundaresan
Distinguished Technologist
CableLabs
k.sundaresan@cablelabs.com

Tom Williams
Distinguished Technologist
CableLabs
t.williams@cablelabs.com

Sheldon Webster
Lead Engineer
CableLabs
s.webster@cablelabs.com

Alberto Campos
Fellow
CableLabs
a.campos@cablelabs.com

Doug Jones
Principal Architect
CableLabs
d.jones@cablelabs.com

CableLabs
858 Coal Creek Circle, Louisville, CO,80027
3036619100

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

Cable operators are facing an unprecedented increase in upstream traffic usage because of the shift to working/schooling from home, and the reliability, capacity, and efficiency of the upstream is top of mind for all operators. This paper is focused on improving the efficiency of the DOCSIS upstream (DOCSIS 3.0 SC-QAM and DOCSIS 3.1 OFDMA) channels. The paper will document some of the upstream engineering problems seen by operators and make robust recommendations on how improve those situations. There are thousands of upstream parameter settings and control knobs available to an operator and much of the time operators leave those settings at their default values. This paper will investigate areas (for SC-QAM & OFDMA) such as minislot size, FEC, Cyclic Prefix, FFT size, frame size, pilot patterns, channel size, guard bands, profile definition, CMTS settings for changing IUCs, etc., and make recommendations on each of those parameters and settings. This paper will also explore configuration file service parameters for a cable modem (CM) and make recommendations for better performance on the upstream. Ultimately, this paper will help operators understand optimized upstream configurations and settings for production deployments.

This paper assumes some background understanding of both DOCSIS 3.0 and DOCSIS 3.1 upstream technologies. Here we try to provide an understanding of different upstream channel parameters (which an operator can tweak) and then the pros and cons of choosing different values for that parameter. Based on these tradeoffs we make recommendations on optimal values for each of these channel parameters or network configurations.

2. DOCSIS 3.0 SC-QAM Upstream Recommendations

This section of the paper is a “how-to” on improving SC-QAM upstream efficiency. Access network, outside plant and system engineers would like to understand the benefits and how to realize those benefits as they continue to maintain the DOCSIS SC-QAM upstream technology.

2.1. Minislot Sizes

The DOCSIS bandwidth allocation MAP uses time units of "minislots." On SC-QAM channels, the size (duration) of the minislot is a multiple of the DOCSIS SYNC time ticks (6.25 us). A minislot is the unit of granularity for upstream transmission opportunities and represents the time allowed for CM transmission of a fixed number of symbols. Figure 1 shows the range of values of minislot sizes for an upstream SC-QAM channel, assuming 64 QAM as the modulation order.

To determine the best minislot sizes, one needs to understand how the network upstream traffic behaves and model it accurately. Based on the traffic patterns seen on the upstream one can optimize the DOCSIS minislot size for an upstream channel.

There are overheads involved with different settings of the minislot size. Setting minislot sizes to one of smaller settings can increase the scheduling granularity for the CMTS. Also, for larger grants, the overhead of bytes wasted in the grant will reduce. On the other hand, larger minislot sizes increases the largest grant that can be made to DOCSIS 2.0 modems (which support only one upstream SC-QAM channel). Prior to DOCSIS 3.0 technology modems requested bandwidth in minislots, whereas starting with DOCSIS 3.0 technology, modems (with support for channel bonding) request bandwidth in bytes using the queue-depth based requests. Also, a larger minislot size, may make the computations a bit simpler on the CMTS upstream scheduler and receiver though the benefits may vary by implementation.

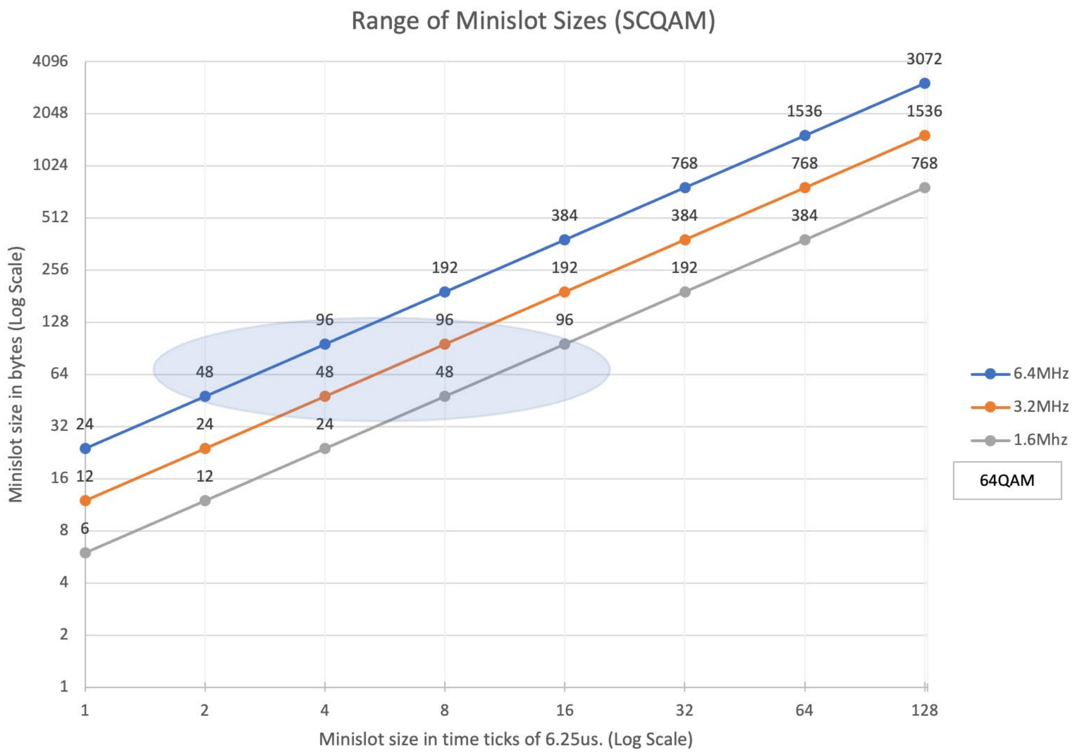

Figure 1 – SC-QAM Minislot Sizes Range

As seen in the graph above, for a 6.4 MHz channel (5.12 MSyms/sec) using 64 QAM modulation, a minislot size of 1 has a raw capacity of 24 bytes. The maximum grant size (a 255-minislot grant) will be 6120 bytes. A minislot size of 4 will have a raw capacity of 96 bytes and allows a maximum 255-minislot grant of 24480 bytes. Based on current patterns of upstream packet sizes, and the size of an Ethernet frame carrying a TCP ACK (64 bytes), the consensus is that minislot sizes in the range of 48 bytes to 96 bytes will enable an appropriate variance in grant capacity and efficiency for upstream scheduling. This translated to a minislot size configuration of 2 or 4 for a 6.4 MHz channel or a size of 4 or 8 for a 3.2 MHz channel.

2.2. FEC Settings

DOCSIS specifications allow for a range of FEC settings for an upstream SC-QAM channel. This is enabled by two settings in the upstream channel which have the following range restrictions:

- FEC codeword information bytes (k),   16 <= k <= 253
- FEC codeword correction setting (T),   0 <= T <= 16.
    - The number of parity bytes is 2T and ranges from 0 bytes to 32 bytes
    - Thus, the total codeword length (k+2T) would be in the range of 16 to 285 bytes

A short data grant may use FEC parameters that are appropriate to shorter packets while a long data grant may be able to take advantage of greater FEC coding efficiency. FEC codeword lengths for Request, Request/Data, and Ranging IUCs can be shorter while the codeword lengths should be longer for Short Data grant (IUC 5) and Long Data grant (IUC 6) and UGS grants (IUC 11).

The FEC parameter selection is a trade-off between channel utilization and robustness to noise. In practice a codeword efficiency of 75%-to-90% (i.e., 10%-25% FEC overhead) looks to be the optimum for a DOCSIS upstream SC-QAM channel. An operator can use the lower end of that range for noisier channels (e.g., lower in the cable upstream spectrum) and the higher end of the FEC efficiency range for cleaner channels. Of all the possible values of k from 16 to 253, many k & 2T values lead to a codeword efficiency which fall outside this optimal range. The below graph (Figure 2) shows the zone (gray) of these optimal values.

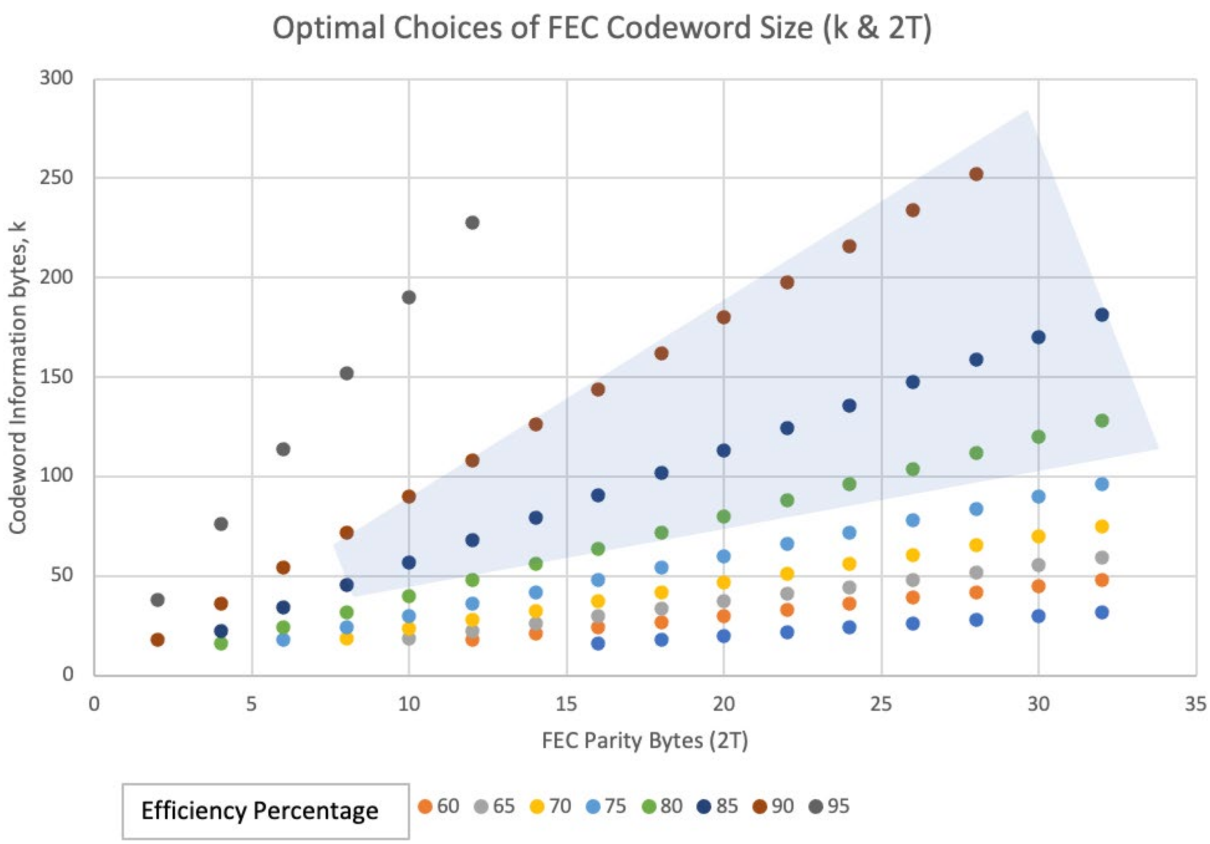

Figure 2 – SC-QAM FEC Sizes Range

Table 1 shows the range of k values (information bytes) for various choices of 2T (FEC Parity bytes). Combinations of k & 2T which give below 50% codeword efficiency are ignored, as they are not useful in the normal cable plant. Shown in green are the range of k values which fall in the 75-90% efficiency range and are the meaningful codeword sizes to choose for Data IUCs.

Ideally an operator would read the RxMER numbers for each of the CMs using the upstream channel and based on those would choose an appropriate FEC setting (k & 2T) and vary the modulation order for that IUC.

Table 1 – Choosing values of k and 2T for SCQAM FEC

Information bytes (k)	FEC bytes, 2T															
Codeword Efficiency	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32
50								16	18	20	22	24	26	28	30	32
60						18	21	24	27	30	33	36	39	42	45	48
65					19	22	26	30	33	37	41	45	48	52	56	59
70				19	23	28	33	37	42	47	51	56	61	65	70	75
75			18	24	30	36	42	48	54	60	66	72	78	84	90	96
80		16	24	32	40	48	56	64	72	80	88	96	104	112	120	128
85		23	34	45	57	68	79	91	102	113	125	136	147	159	170	181
90	18	36	54	72	90	108	126	144	162	180	198	216	234	252		
95	38	76	114	152	190	228										

2.3. Interleaver settings

The DOCSIS SC-QAM upstream interleaver supports two modes, one operating mode in which the block size is fixed, and a second dynamic mode in which the interleaver depth is determined based on the burst size. The FEC encoded data bytes of the packet are first divided into interleaver blocks. In the fixed mode, the interleaving depth of the last interleaving block of a packet can be small, resulting in low burst noise robustness for this block. In dynamic mode, the depths of the interleaver blocks are chosen such that all blocks have approximately the same depth to achieve nearly optimal burst noise robustness (for the given block size). The current point of view is to choose "dynamic mode" and let the system decide the appropriate interleaver settings. An MSO who uses data analytics on the channel performance metrics could potentially outperform the CMTS but until operators have such software systems built, the dynamic options look to be the best.

2.4. Upstream performance

This section provides guidance on the operational conditions and recommends which modulation orders are recommended for various levels of SNR to have an FEC error-free operation. The ability of the system to support a given QAM level depends on the RxMER values and the mappings to an appropriate QAM level when creating a profile. These mappings are used for DOCSIS 3.0 SC-QAM channels and are summarized in Table 2 below.

Table 2 – Upstream RxMER to QAM Level mapping

Upstream Constellation/ Bit Loading	Upstream MER (dB)
QPSK	10
8 QAM	14
16 QAM	16
32 QAM	19
64 QAM	22+

3. DOCSIS 3.1 OFDMA Upstream Recommendations

This section of the paper is a "how-to" on improving OFDMA efficiency. Access networks, outside plant and system engineers would like to understand the benefits and how to realize those benefits as they plan to deploy DOCSIS 3.1 OFDMA technology.

3.1. OFDMA Channel Location

Many operators have shared data indicating the spectrum below 20 MHz is burdened by the presence of ingress noise as shown in Figure 3. In some plants which are well maintained and cleaner, operators have successfully reported running SC-QAM channels down to 9 MHz. In some cases, if the plant is clean, the spectrum may be amenable to the presence of an OFDMA channel in those lower frequencies. In many cases if the spectrum contains bursty or impulsive noise ingress, then we do not recommend placing OFDMA channels in that region. A few operators have reported OFDMA codeword errors leading to system stability issues with the OFDMA channel in spectrum below 20 MHz, whereas these issues become non-existent when the channel is moved to higher portion of spectrum.

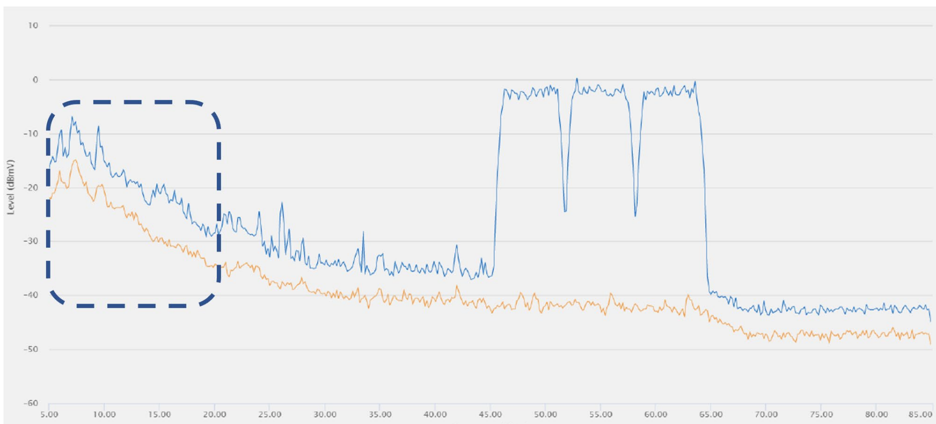

Figure 3 – Noisy Spectrum Below 20 MHz

When initially deploying OFDMA channels it would be prudent to initiate the OFDMA trial in a cleaner (higher) part of the spectrum. The idea is to work out the inevitable system issues and bugs and get the OFDMA channel working reliably first in a cleaner part of the spectrum and then evaluate OFDMA below 20 MHz.

3.2. Number of symbols in a frame (K)

The number of symbols in an OFDMA frame, K, is configurable between 6 (minimal value) and one of the following values:

- With 20 µs FFT duration (2K FFT)
  - Kmax = 18 for BW≥ 72 MHz
  - Kmax = 24 for 48 MHz ≤BW < 72 MHz
  - Kmax = 36 for BW < 48 MHz
- With 40 µs FFT duration (4K FFT)
  - Kmax = 9 for BW ≥ 72 MHz
  - Kmax = 12 for 48 MHz≤ BW < 72 MHz
  - Kmax = 18 for BW < 48 MHz

Figure 4 shows the layout of minislots within an OFDMA channel.

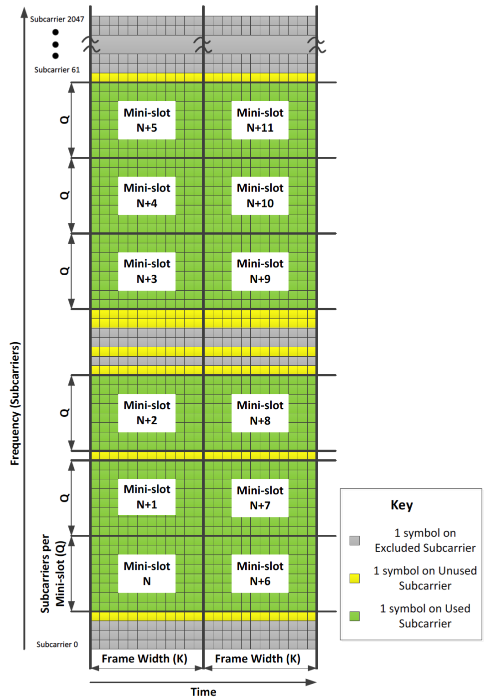

Figure 4 – Minislots in OFDMA

We performed various throughput tests in the lab with 32 MHz and 75 MHz channels, at 2K FFT and 4K FFT with different values of K. We consistently found a 3% to 5 % increase in data throughput across all CMTS platforms when the value of K is increased from 6 to 9/18 (for the larger 75MHz channels) or 6 to 18/36 (for smaller 32 MHz channels).

There are two minislot types, edge minislots and body minislots. Edge minislots have a bit more overhead and carry less data than body minislots. An edge minislot is the first minislot in a transmission burst and also used after an exclusion band or a set of skipped subcarriers. Body minislots are used for all other minislots in a burst. Hence it is advisable for an operator to avoid gaps in the OFDMA channel if possible.

3.3. FEC – Unreliable/Uncorrected Codewords

We tested various DOCSIS systems to verify the FEC performance under different noise level/ type (flat AWGN) to build expectations under different plant conditions. Figure 5, 6, and 7 show the FEC error rate (both correctable and uncorrectable) over different RxMER levels. The lab testing was run with traffic of 200- and 1500-byte packets and a combination of the two. The noise (AWGN) was introduced to reduce the RxMER levels across the OFDMA channel, and the levels of FEC statistics were captured at every level of increasing noise. We also tested with two different pilot patterns, the most dense pattern 4 and the least dense pattern 1.

Note that in most CMTS implementations when unreliable codewords are detected on an OFDMA channel, the upstream receiver discards the unreliable codewords considering them as uncorrectable. In other implementations they are sent to the MAC layer for additional processing using the MAC HCS and Ethernet CRC.

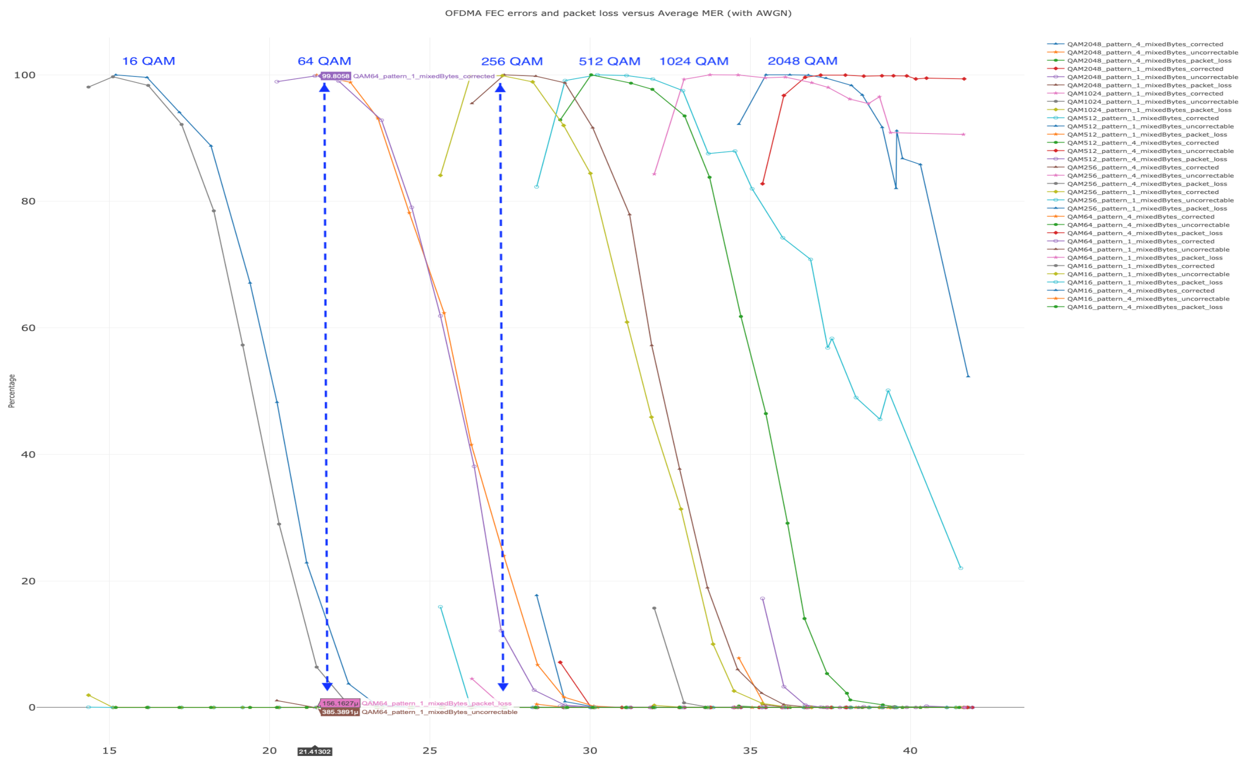

Figure 5 – OFDMA FEC Performance (mixed size packets, Pilot Pattern 1,4)

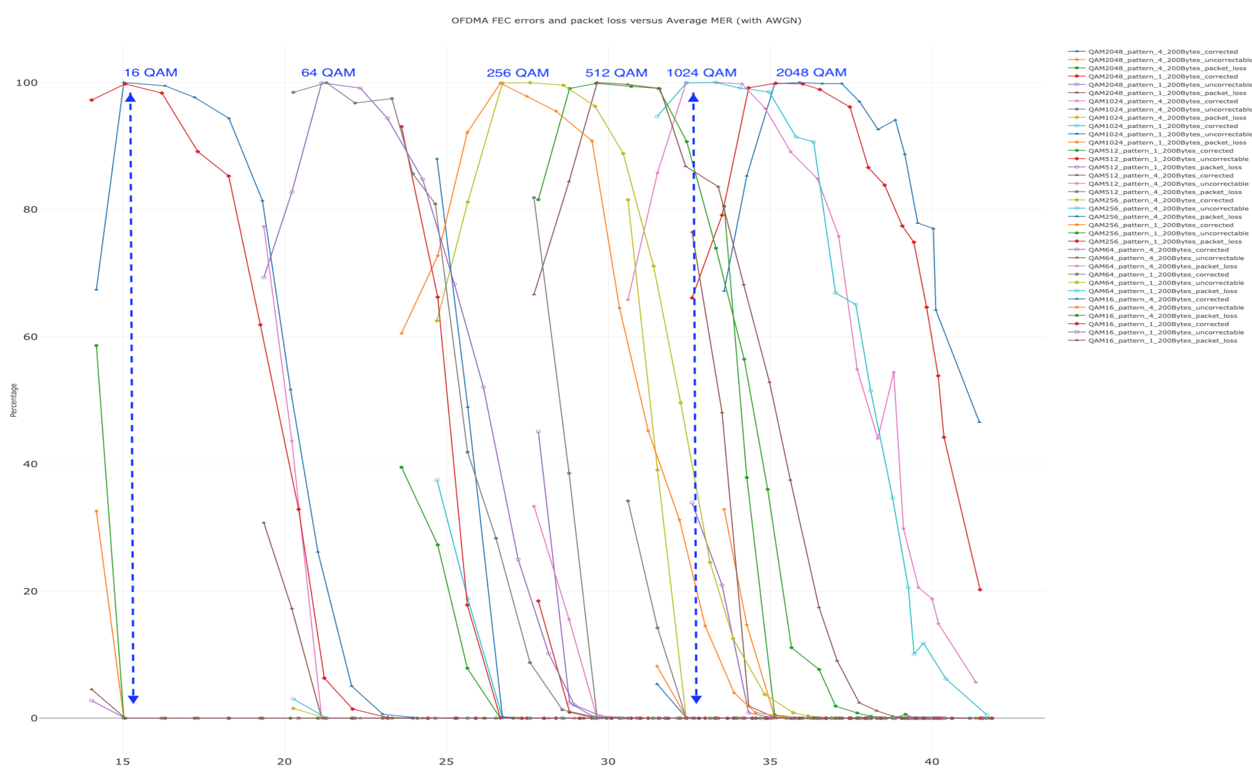

Figure 6 – OFDMA FEC Performance (200 byte packets)

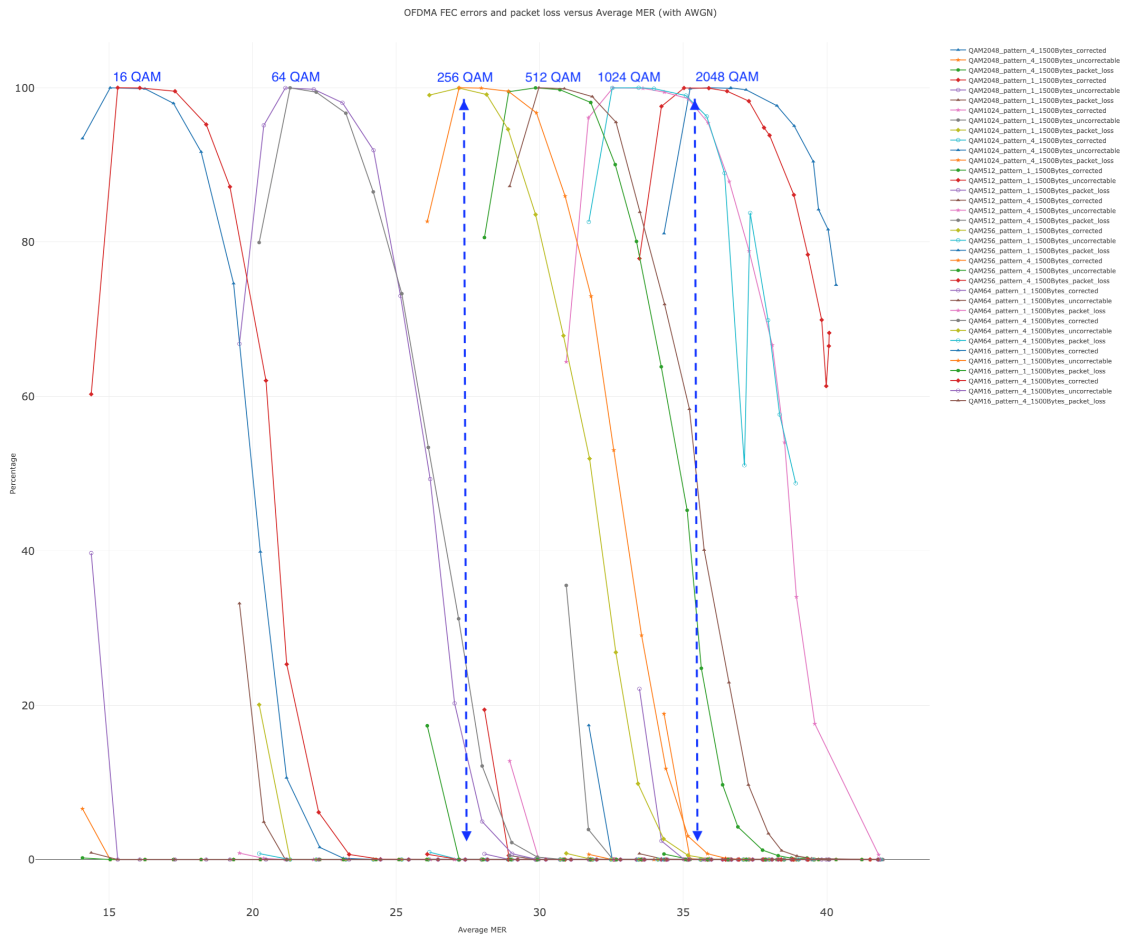

Figure 7 – OFDMA FEC Performance (1500 byte packets)

3.4. Cyclic Prefix (CP) and Roll-off Period (RP)

Two options exist for upstream OFDMA FFT size, either 2048 (2k FFT) or 4096 (4k FFT). A cyclic prefix (CP), which precedes an OFDMA symbol, should be slightly longer than a longest significant echo to be encountered in the channel, which normally is in the range of a few microseconds for coaxial cable. (See Figure below from [DOCSIS PHYv3.1]. The addition of a cyclic prefix enables the receiver to overcome the effects of inter-symbol-interference caused by micro-reflections in the channel.

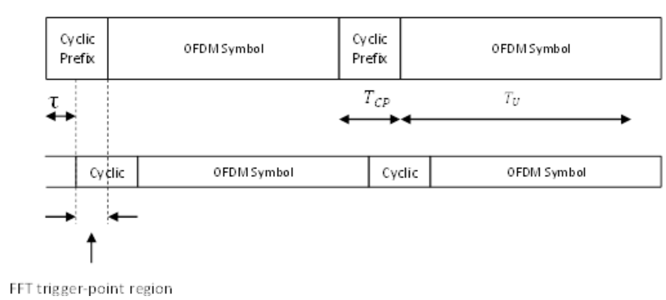

Figure 8 – CP and microreflections

The duration of the FFT useful symbol duration is 20 µsec for the 2k mode, or 40 µsec for the 4k mode. Therefore the percentage of overhead for a 2.5µsec CP can be either 12.5% (2k FFT) or 6.25% (4k FFT). CP sizes vary from 0.9375 µsec to 6.25 µsec, for which the overhead will vary from 4.6% to 31.25% for the 2k FFT, or 2.3% to 15.6%for the 4k FFT. Figure 1 shows the change in capacity for a 40 MHz channel, 2k FFT size, for the range of possible CP values.

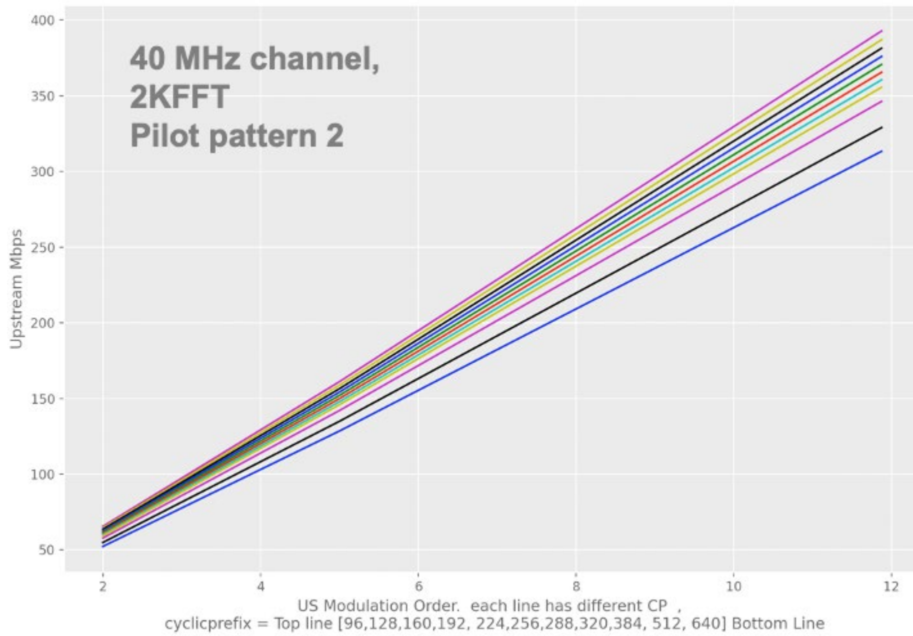

Figure 9 – OFDMA capacity change for different CP

Windowing maximizes channel capacity by sharpening the edges of the spectrum of the OFDMA signal (this applies to an OFDM downstream signal as well). Spectral edges occur at the two ends of the spectrum of the OFDMA symbol, as well as at the ends of internal exclusion band. The roll-off period (RP) setting is another OFDMA overhead. An RP is defined by a number of samples and essentially is a gradual rise or fall in energy and precedes or trails an OFDMA transmission. The purpose of the RP is to limit the spread of interference of an OFDMA transmission to adjacent single frequency carriers. See Figure 14. A longer duration RP produces less interference, and like the CP, decreases efficiency. Note that the effect of a RP is to decrease effective CP time, so if a long RP is chosen, the CP duration should be increased. The RP does nothing to protect an OFDMA channel from energy from adjacent SC-QAM channels.

3.5. 2K vs 4K FFT size or (subcarrier spacing)

The OFDMA upstream multicarrier system is composed of either 25 kHz or 50 kHz wide subcarriers. In the upstream, the subcarriers are grouped into independently configurable OFDMA channels each encompassing up to 95 MHz of spectrum, totaling 3800 25 kHz spaced subcarriers or 1900 50 kHz spaced active subcarriers. Figure 10 shows the possible size of the minislots with 4k FFT (25 kHz subcarriers) and with 2k FFT (50 kHz subcarriers)

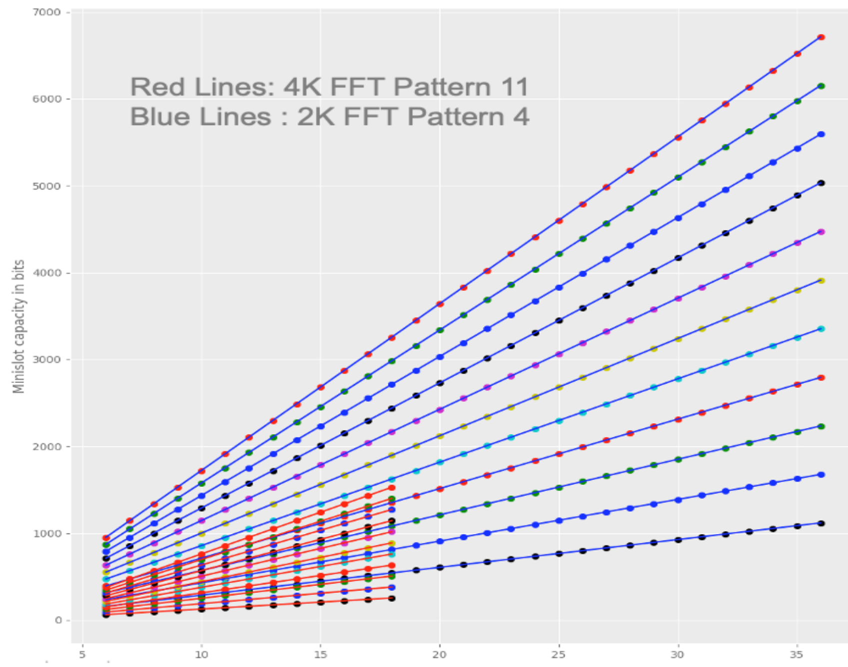

Figure 10 – OFDMA Minislot capacity change 2K vs 4K FFT

Figure 11 shows the capacity difference in OFDMA channels using either the 2K FFT or the 4K FFT, with different pilot patterns.

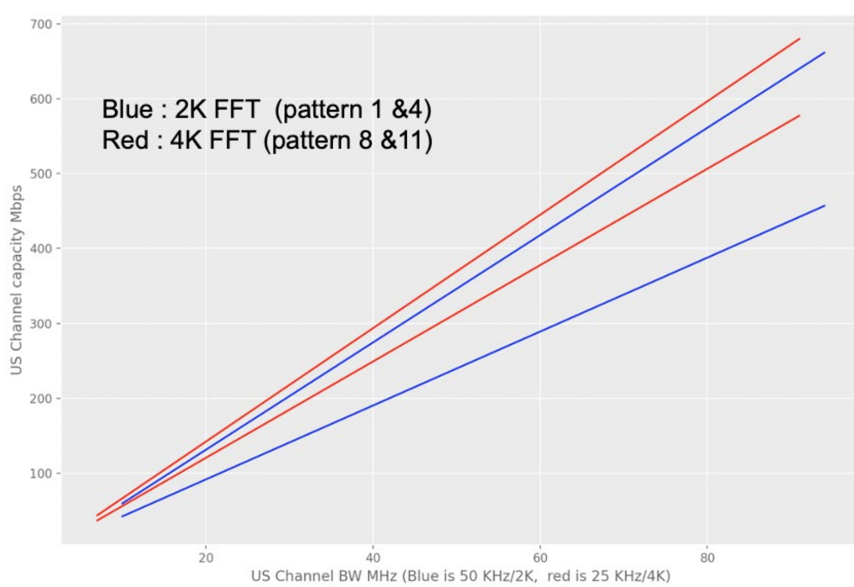

Figure 11 – OFDMA capacity change 2K vs 4K FFT

When processing an OFDM or OFDMA frame with CP, an exact integer number of symbols are chosen for processing by a FFT, either 2k or 4k. The abrupt truncation of symbols in the time domain creates a window in the frequency domain which has frequency domain sidelobes. See Figure 12. If spurious energy, such as a single carrier or a continuous wave (CW), lands on a sidelobe it will cause inter-symbol interference, generally affecting the subcarriers on the edges of the OFDM(A) band the most. One solution is to provide vacant bandwidth adjacent to the OFDM(A) transmission to prevent sidelobe vulnerability. Another solution that has been used is to use a lower modulation order for subcarriers that are on the lower and upper portions of an OFDM block. Using lower order modulation, such as 256-QAM instead of 1024-QAM reduces the chance of symbol errors.

On sidelobe vulnerability, 2k is twice as susceptible because each sidelobe is twice as wide, relative to 4k. Note that this vulnerability affects upstream OFDMA as well as downstream OFDM.

An adjacent OFDMA channel will not cause harm to an OFDMA channel if the interfering frame timing is the same. This implies that both channels are using the same timing including FFT size, CP, and RP.

4k does have a vulnerability relative to 2k. It is more susceptible to phase noise. That is because any wander or movement of phase while an OFDM(A) frame is being captured is detrimental. Phase noise can be caused by a transmitter symbol clock, a receiver symbol clock, or any frequency conversion used in the OFDM(A) signal processing, such as block up and down conversion. Lab tests indicate this effect is a couple of dBs of MER worsening for 4k around the 35-40 dB level.

3.6. Upstream Pilot Patterns

Various pilot patterns are available for OFDMA transmissions. Not all CMTS vendors have implemented all pilot patterns, and different implementations support different default pilot patterns. Generally boosted pilots give a better channel characterization in the presence of noise, and denser pilot patterns improve estimates for signal time, frequency, phase, and amplitude offsets. But more pilots come at the expense of efficiency. The pilot pattern chosen should be sufficiently dense to characterize the ripples in the frequency response caused by echoes. If transmit pre-equalization is not used, less efficient pilot patterns should be used.

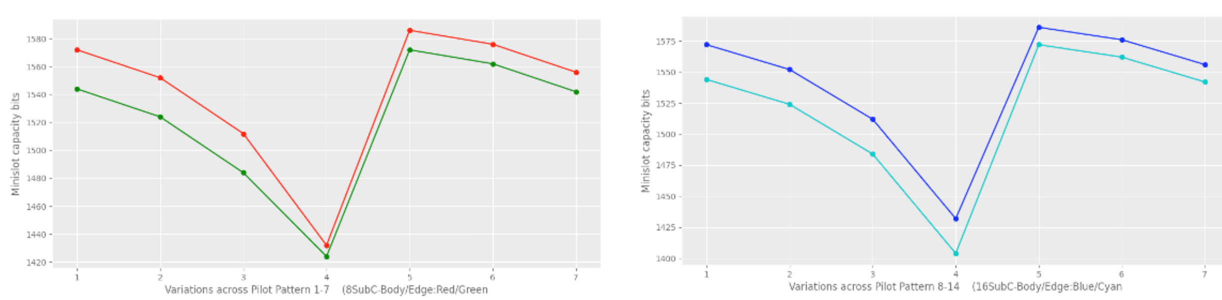

K= 20 symbols. 2K FFT
Patterns 1-7

K= 10 symbols. 4K FFT
Patterns 8-14

Figure 12 – OFDMA Minislot capacity change w Pilot Patterns

It is recommended that pre-equalization be turned on for both OFDMA and SC-QAM channels because not using pre-equalization requires a use of more dense pilot patterns, which are less efficient. Operators need to watch for this as in many cases pre-equalization, for OFDMA and SC-QAM channels, is off by default.

3.7. Efficiency of OFDMA channel

The throughput and efficiency (e.g., bits/Hz) of the OFDMA channel changes with respect to channel width and other configuration parameters, e.g., CP, RP, frame size, etc. The expected range has been detailed in a previous paper [D31 Capacity].

4. Upstream Interference Analysis

Generally, these technical discussions apply to both upstream and downstream single carrier (SC, such as SC-QAM) and multicarrier signals (MC signals, such as OFDMA and OFDM), although upstream efficiency is the general topic of this paper.

A problem arises when a single carrier signal and a multicarrier signal are located next to each other in adjacent frequencies. For this reason, when doing spectrum planning, for best efficiency it is generally best not to "mix up" single carrier and multi carrier signals but keep them separated as much as possible

On the upstream, as previously discussed, due to upstream noise funneling, the lower part of the spectrum, especially below 20 MHz, is hostile to all signals, and particularly hostile to OFDMA signals. This is due to impulsive energy spreading by an OFDMA receiver's FFT. This is particularly detrimental when there is switching regulator impulsive noise ingress in the upstream. This example of impulsive energy is created by poor switching power supply filtering combined with a loss of coaxial shield integrity (e.g., shield break). This allows both radiated and conducted energy from the power mains in homes to get onto the coax center conductors and travel upstream. Many switching regulated power supplies operate at frequencies close to the OFDM/A frame rates, such as 50 kHz, so there may be an impulse in every OFMDA FFT, with impulsive energy being transformed (spread) to each OFDMA symbol.

In the downstream both SC-QAM and OFDM signals are continuous, hence if there is adjacent channel interference, it is also continuous. In contrast, the upstream uses bursts of energy, so if a SC-QAM burst occurs when there is no adjacent OFDMA burst, no interference occurs. And if a OFDMA burst occurs when there are no adjacent SC-QAM bursts, again there is no interference. An interference problem only arises when the SC-QAM and OFDMA bursts occur at the same time(s) in adjacent frequency bands.

Interference needs to be evaluated in both directions, OFDMA into SC-QAM, and SC-QAM into OFDMA. Generally, an adjacent SC-QAM signal is more likely to damage an OFDMA signal than vice-versa.

4.1. OFDMA interference into SC-QAM

When an OFDM or OFDMA signal is created by a D-A converter and transmitted, out of band sidelobes are created. This is because, when transmitting data, the OFDMA signal abruptly transitions to different amplitude (voltage) levels. The sidelobes can be seen and measured on a conventional scalar spectrum analyzer and can cause interference with adjacent single carrier signals. Usually, time domain symbol damage is limited to just the one or two symbols that occur at the start and stop of the OFDM frame. The sidelobes interference can be reduced by using a roll-off period specified in the DOCSIS specification, along with providing a guard band between an OFDM(A) block and a single carrier signal. This causes inefficiency both in time for the roll-off period, and lost bandwidth.

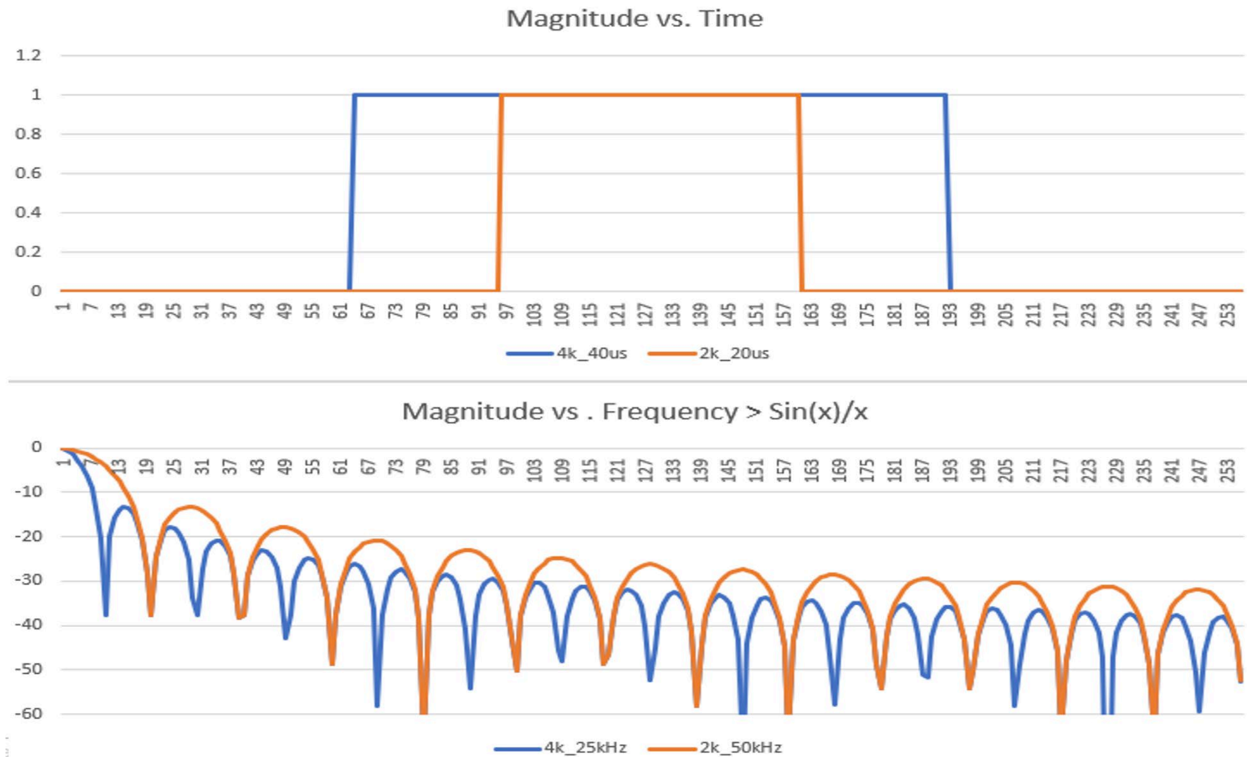

Figure 13 – OFDM/A Sidelobes created by rectangular time domain energy

In Figure 13, the upper plot is a rectangular pulse of magnitude voltage vs. time. The vertical axis is linear voltage, and the horizontal axis is time index. The rectangular pulses have a duration of either 20 µsec (2k mode) or 40 µsec (4k mode). If this waveform is transformed into the frequency domain with an FFT, the lower plots are obtained. The lower plot's vertical axis is dB, and the horizontal axis is frequency. The spectral shape consists of a DC term at 0 dB on the left, and then decreasing sidelobes over the rest of the plot. The sidelobes decrease as an absolute value of a sine(x)/x function. The separation between sidelobes is 50 kHz for 2k mode or 25 kHz for 4k mode. As you place more vacant spectrum between a SC-QAM signal and an OFDMA signal, the OFDMA interference to the SC-QAM signal is less.

The DOCSIS specifications provides another way to reduce interference besides vacant spectrum. The method is to shape OFDMA interference sidelobes and reduce their spectral reach over into the SC-QAM signal. That is done by employing a roll-off period (RP), which is tapering of the time domain envelope so that it gradually rises and falls, not with an abrupt transition, as illustrated in Figure 14. This is also referred to as a Tukey filter. Tapering is illustrated in Figure below (from the [DOCSIS PHY v3.1] specification). The roll-off period is subtracted from the cyclic prefix (CP) length, so CP length may need to be increased if a roll-off period is used. Unfortunately, employing a roll-off period and a CP is a time overhead penalty which decreases efficiency.

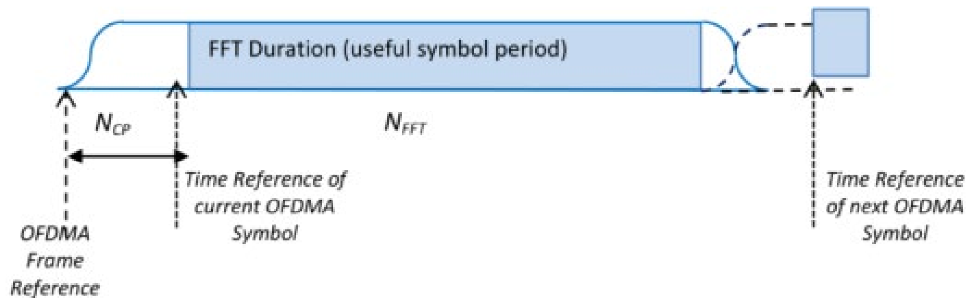

Figure 14 – OFDMA Signal, Useful symbol period and roll off

The above figure shows a time diagram of an OFDMA signal showing roll-off period and useful symbol period.

As mentioned above, for OFDM modulation, in general, there are spectral sidelobes beyond the edges of the modulated spectrum. This is a physical property of OFDM modulation and comes in the form of non-zero modulation energy beyond the edges of the configured modulated spectrum (i.e., below the lowest active subcarrier and above the highest active subcarrier). Figure 15 below, illustrates this modulated energy (sidelobes) and how it varies over different settings of the roll-off period (Nrp).

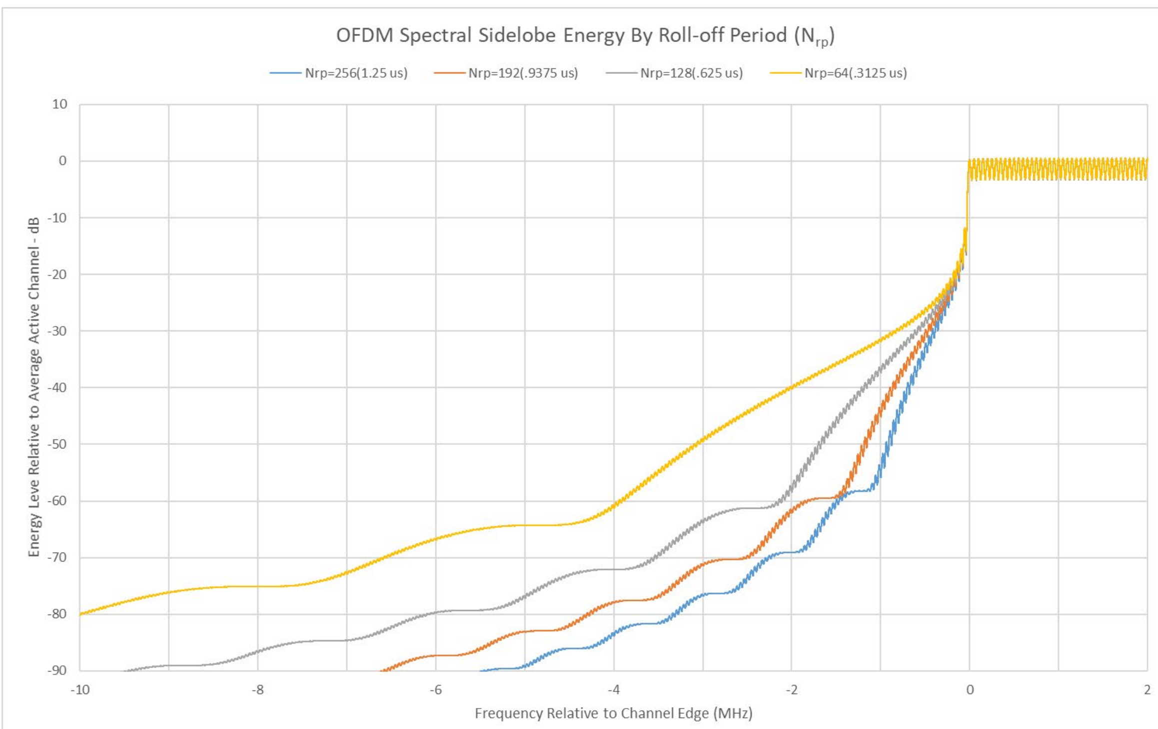

Figure 15 – OFDM/A Spectral Sidelobes (impact on Adj. SC-QAM)

** Figure 15: MATLAB simulation data provided courtesy of Roger Fish & Thomas Kolze, Broadcom*

Note that there is a trade-off between different Nrp settings. Smaller settings may offer more efficiency in the time-domain but have the cost of additional spectrum used by the OFDM channel modulation (larger spectral sidelobes/guard bands).

If you consider the top time plot of Figure 13 carefully, the interference to adjacent SC-QAM channels only occurs in time at the instant when the voltage is abruptly changing to or from zero, occurring every 20 or 40 us. So, if the SC-QAM bursts could be timed to "miss" those abrupt steps, no interference would occur. That could be an opportunity for a CMTS scheduler. Using 40 us (4k mode) generates half as many abrupt steps as 2k mode, so is half as harmful to an adjacent SC-QAM.

There is also a possibility that the damage to adjacent SC-QAM symbols can be eliminated by using good FEC settings on the SC-QAM bursts. The SC-QAM symbol rate is 5.12 MSymbols per second, and symbols occur every 195 ns. The abrupt transitions will damage one or two SC-QAM symbols out of every 256 symbols. Upstream FEC uses a Reed-Solomon block code, so setting the value of T to a large number (e.g., T= 16) can undo symbol damage. Furthermore, the increased FEC code overhead increases the robustness of the SC-QAM signal in general.

If two OFDM(A) signals are using the same exact timing, start time, stop time, cyclic prefix, and roll-off periods, they do not interfere with each other. So, no guard band or roll-off period needs to be used when signals are originating from the same slot/port on a CMTS and are using the same synchronous clock and the same OFDM & OFDMA frame timing. (The same CP etc. on the OFDMA frame)

4.2. SC-QAM interference into OFDMA

When receiving an OFDM(A) channel a comparable timing process occurs, where a first sample is taken continuously to the final sample. Unfortunately, time sampling is abrupt, not tapered. As a result, the same leakage phenomena that occurs with transmission occurs with reception, but this shows up as a susceptibility to interference from foreign energy at frequencies outside of the OFDM(A) channel. This energy can be a nearby SC-QAM and shows up as poor MER on subcarriers near the band edge as seen in the experimental results in Chapter 5.

Back a few years ago, it was observed in the field that the RxMER (receive modulation error rate) per subcarrier for lowest and highest subcarriers appeared worse than MER for subcarriers in the center of the OFDM or OFDMA band. It was determined that the presence of a SC-QAM above or below the OFDM or OFDMA spectrum was the culprit. Figure 17 and Figure 18 are plots of MER per subcarrier with an adjacent SC-QAM signal. Observe that subcarriers nearest the OFMDA band edges were automatically eliminated (not used) by the CMTS to avoid damage. Figure 19 is a MER per subcarrier plot with no adjacent SC-QAM signals. MER is uniformly good.

Why was this happening? An OFMDA receiver uses a rectangular sampling window in the time domain. The rectangular sampling window is illustrated in Figure 14 as the time interval with a blue shading. A resulting receive spectral window is not rectangular, but also extends into adjacent frequencies with sidelobes.

Figure 16 illustrates two receive spectral windows, one in red for 2k FFT and one in blue for 4k FFT. Because the sidelobes associated with the 4k FFT sidelobes decay more rapidly than the sidelobes associated with the 2k FFT, the 4k FFT has a narrower receive spectral window, and thus is more efficient than the 2k FFT. That is, the 4k FFT pulls in less energy from an adjacent SC-QAM signal relative to a 2k FFT.

Can using a roll off period on a transmitted signal improve efficiency by reducing interference? The answer is no, because the tapering is not inside the blue sampled region and is discarded (not used in the transform). Using a long cyclic prefix or a long roll-off period do not improve the problem of poor receive MER on OFMDA band edges. Choosing a 25 kHz (4k) subcarrier spacing over a 50kHz (2k) subcarrier spacing does make an improvement.

This advantage is illustrated in Figure 16. This narrower receive spectral window advantage for a larger transform size applies for both OFDMA and OFDM.

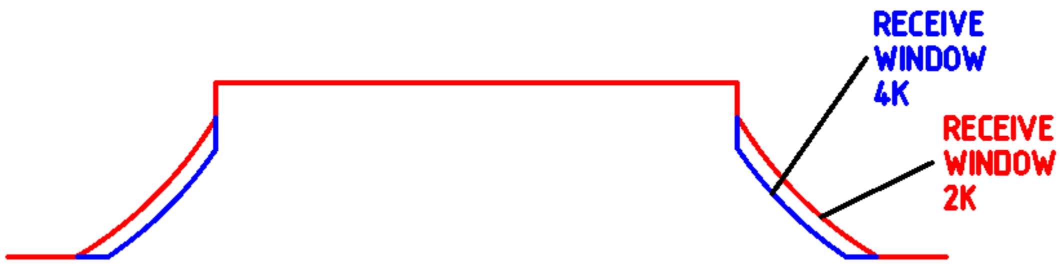

Figure 16 – Receive Windows 2k vs 4k

Receive spectral windows extend into the sidelobes and are wider for 2k (red) relative to 4k (blue) FFT.

4.3. Conclusions on Upstream Interference

The following are some conclusions reached based on the above discussion.

1. Generally OFDMA is harmed more by an adjacent SC-QAM than the SC-QAM is harmed by an adjacent OFDMA.

2. 4k is more efficient than 2k for reasons of narrower receive spectral window and required CP being half the symbol period. 4k also interferes with adjacent SC-QAM less, because of half as many abrupt voltage transitions.

3. Putting on more CP or RP on an OFDMA does not protect it from an adjacent SC-QAM signal. That is because its CP or RP do not modify the receive spectral window.

4. Adding a RP to an OFDMA channel does protect an adjacent SC-QAM channel from interference.

5. This discussion applies to upstream or downstream channel adjacencies.

5. OFDMA and SC-QAM Upstream Recommendations

5.1. Guard Bands

To examine the treatment of spectral guard bands in the DOCSIS specifications for OFDM/OFDMA channels and within different implementations, it is instructive to look at fundamental properties of OFDM modulation, in general, as well as a comparison between DS/OFDM and US/OFDMA specification requirements. As described in chapter 4, it is important to point out that the additional spectral sidelobe energy beyond the edges of the modulated spectrum is part of the channel and carries useful information. Therefore, another signal encroaching on this spectrum, such as a non-synchronous adjacent channel, has the potential to interfere with the OFDM channel and, vice-versa, the spectral sidelobe energy of the OFDM channel can also cause some interference on the adjacent channel.

5.1.1. DS/OFDM Requirements

We think of a full-width DS/OFDM channel as 192 MHz, although the maximum modulated spectrum width is only 190 MHz. This is convenient since, for downstream OFDM channels, the specifications require a minimum 1 MHz spectral guard band at the lower and upper channel edges. This guard band is required to accommodate the spectral sidelobe energy at the edges of the OFDM channel.

The nominal 192 MHz OFDM channel width includes 1 MHz (minimum) unmodulated spectrum at each edge as guard bands. It should be noted that more guard band could be needed, depending upon the setting for the roll-off period (N_{rp}). [DOCSIS PHYv3.1] spec (Appendix V) provides suggested additional guard band sizes (referred to as Taper Regions) for different N_{rp} settings. The MATLAB model used for the plot in Figure 15 was used to obtain the recommended guard bands (Taper Regions). Note that there is a trade-off between different N_{rp} settings. Smaller settings may offer more efficiency in the time-domain but have the cost of additional spectrum used by the OFDM channel modulation (larger spectral sidelobes/guard bands).

5.1.2. US/OFDMA Requirements

The physical properties of the OFDM modulation process still apply for bursted OFDMA upstream signals. Spectral sidelobe energy still occurs during bursts. Since the US/OFDMA channels are scheduled bursts of minislots, adjacent channel interference due to the spectral sidelobes of the OFDMA channel would only occur when these adjacent channels are scheduled and granted to burst simultaneously. Interference can be avoided through the scheduling and granting of these adjacent channels.

Based on our lab experiments, we found that CMTS implementations enforce some amount of guard band at the edges of OFDMA channels. Different CMTS implementations enforce OFDMA guard bands in different ways. In some cases, a minimum 0.5 MHz guard band is enforced by the CMTS at each edge of the OFDMA modulated spectrum. In other cases, 1 MHz or more guard bands are enforced by the CMTS implementation, either fixed or variable, depending upon the configuration parameters of the adjacent channels.

The important take-away here is that different CMTS implementations enforce different guard bands in different ways, even though the specs do not require enforcement of guard bands for US/OFDMA channels. It’s important for the operator to examine and understand whether and/or how their specific CMTS implementation(s) enforce(s) guard bands at the edges of US/OFDMA channels.

This automatic enforcement of guard bands by the CMTS, which is not required by the spec, could result in part of the OFDMA channel not being used. For example, the CMTS might not grant a number of minislots at the edges of the channel so that its own minimum guard band requirements are met.

For the OFDMA channel, as an example, the interference comes in the form of reduced RxMER on the OFDMA subcarriers at the edge close to the adjacent channel whenever both channels are bursting simultaneously. Figure 17 shows the degraded RxMER per subcarrier at the edge of a 32 MHz OFDMA channel when an adjacent SC-QAM channel is bursting simultaneously (SC-QAM channel is to the left of the OFDMA channel). Note that, although the SC-QAM is configured immediately adjacent (guard band = 0 MHz), the CMTS in this case enforces a minimum 0.5 MHz guard band by leaving 2 minislots ungranted/unused at the edge of the channel.

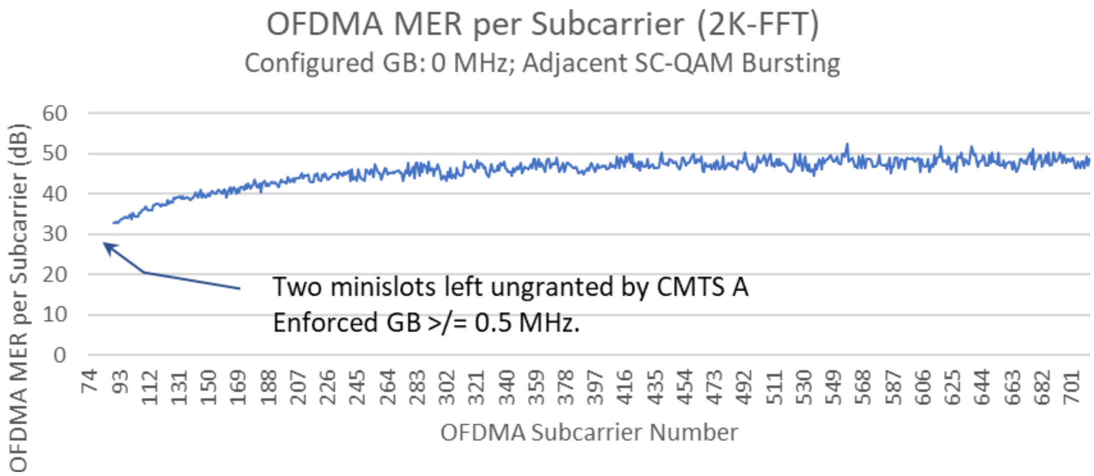

Figure 17 – RxMER : Non-synchronous Adj Channel Bursting Simultaneously (CMTS A)

In Figure 18, after further lab experiments using a different CMTS but with identical channel configuration, we observed that the first 10 minislots are left unused by this CMTS to automatically enforce its own guard band rules.

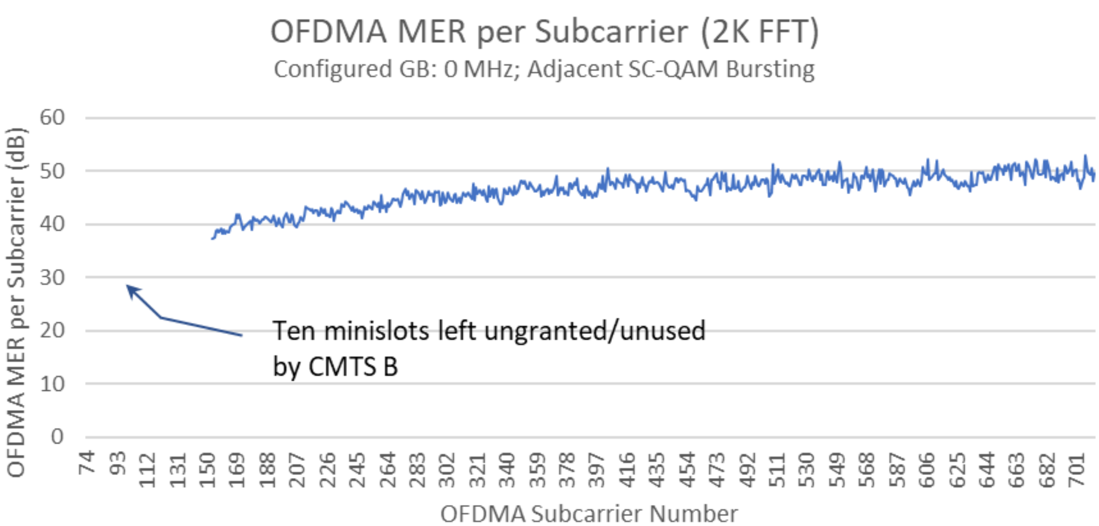

Figure 18 – RxMER : Non-synchronous Adj Channel Bursting Simultaneously (CMTS B)

It is important to remember that the interference affected RxMER and automatically enforced guard bands only occur when both adjacent channels are bursting simultaneously. When bursting alone as shown in Figure 19, there is no degraded RxMER at the edge and there are not automatically enforced guard bands.

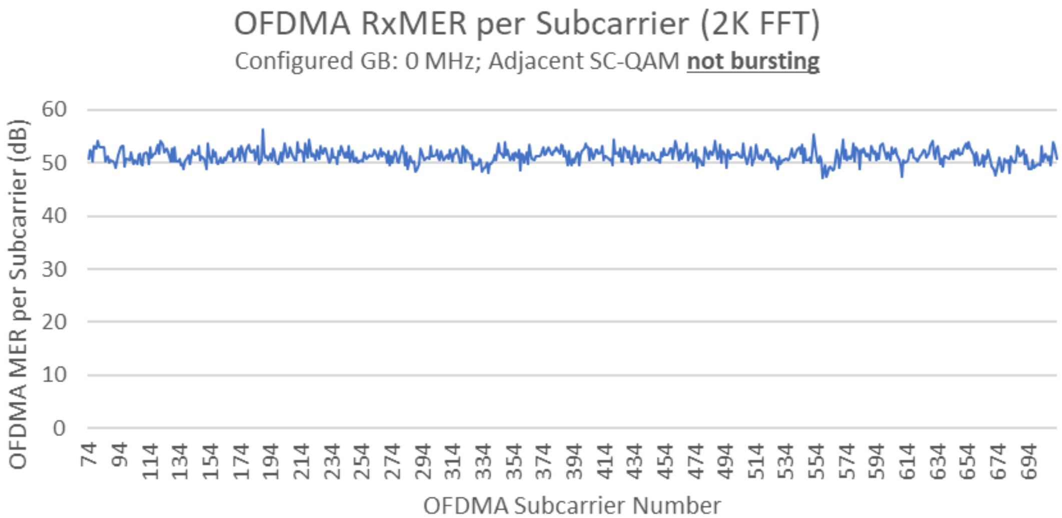

Figure 19 – OFDMA RxMER (OFDMA bursting alone)

5.1.3. Recommendations

It's important for the operator to know and understand how their specific CMTS systems treat adjacent channels with respect to their own enforcement of guard bands. If their own minimum guard band limits are not met by the spectral location of adjacent channels, parts of the OFDMA channel edges might be unused (minislots ungranted) by the CMTS scheduler.

An adjacent SC-QAM channels' interference on the OFDMA channel largely affects the subcarriers on the edge of the OFDMA channel, the impact is that the subcarriers on the edge of the OFDMA channel have worse RxMER than the rest of the channel. The OFDMA interference on SC-QAM can be remedied with a large guard band or moderate guard band with taper (Roll off period). Tapering/Roll-off period does not protect the OFDMA channel from an adjacent SC-QAM channel.

5.2. TaFDM (Time and Frequency Division Multiplexing)

It's useful to examine the use of TaFDM in the context of the discussion of guard bands (above). As previously noted, DOCSIS specifications do not mandate minimum guard bands for US/OFDMA channels. In fact, it follows that different non-synchronous channels may even occupy the same frequency spectrum. This is allowed by the specification and can be handled by the CMTS scheduling the use of overlapping spectrum by the included channels at separate times.

5.2.1. TaFDM and CMTS Guard Band "Rules"

In practice, we find that CMTS implementations tend to adhere to their own guard band enforcement processes, either partially or in whole, to implement TaFDM. In fact, we have seen that some CMTS implementations might automatically trigger their own TaFDM mode when adjacent channels are configured closer than their own expected minimum guard band. Even though the modulated spectrum may not overlap, it could be considered overlapping by the CMTS if a portion of their expected minimum guard band is overlapped.

Other CMTS implementations may extend their own guard band rules to zero out larger portions of the OFDMA channel that is also occupied by, for example, an overlapping SC-QAM channel. This could result in half or even more of the OFDMA channel being zeroed out (not granted for OFDMA bursts).

Like the examples above, Figure 20 shows the MER per subcarrier for the same 32 MHz OFDMA channel, but with a 6.4 MHz SC-QAM channel configured in TaFDM mode just above the center of the OFDMA channel.

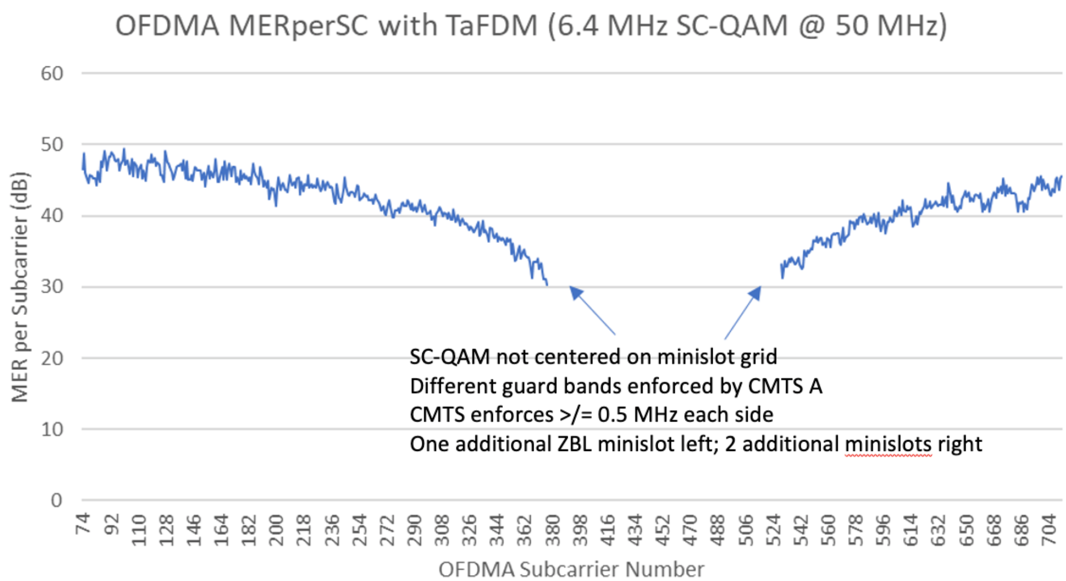

Figure 20 – OFDMA RxMER w TaFDM (SC-QAM bursting in center of OFDMA) (CMTS A)

As shown in Figure 21 when using a different CMTS, nearly half of the OFDMA channel is unused in order to enforce a combination of its own guard band and TaFDM rules.

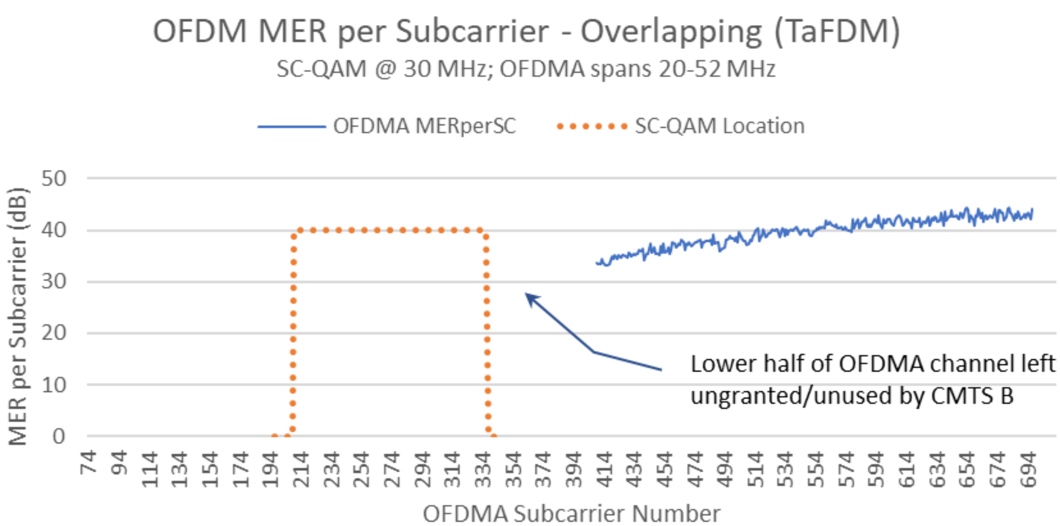

Figure 21 – OFDMA RxMER w TaFDM (SC-QAM bursting in center of OFDMA) (CMTS B)

It is worth mentioning again that when the OFDMA is bursting alone, the full width of the channel is used without MER degradation, as shown in Figure 19.

5.2.2. Recommendations

It is important to understand how a CMTS implementation handles TaFDM to judge whether the associated trade-offs are worthwhile. If a particular CMTS only enforces a small guard band at each end of a fully overlapping spectrum and, therefore, allows simultaneous bursting of overlapping channels using TaFDM, then there may be situations where fully overlapping channels are useful. Of course, it's possible that the use of TaFDM could zero out half or more of the useful modulated spectrum of an OFDMA channel, depending upon the location of the overlapping channel and the specific CMTS "rules" for handling guard bands and TaFDM, and in these cases may not be an efficient use of the spectrum.

6. DOCSIS Protocol and Configuration Efficiency Considerations

DOCSIS protocol relies on several processes for proper operation. These processes include initialization, request for data, the multiple aspects of ranging and the processes to assign CMs to the right channels and their configuration to the right profile. Most of these processes will have some impact on efficiency that are worth further analysis.

6.1. Initialization

The initialization of a CM consists of downstream scanning, upstream channel acquisition, ranging, IP connectivity, time of day acquisition, config file download in addition to registration verification and the establishment of connection privacy. These processes by themselves are part of the expected behavior, however efficiency problems arise when these initializing CMs take a longer time or struggle to complete any of these processes. There may be conditions in the plant, sub-optimal configuration or problems with the implementation that forces a CM to repeat one or more of these processes to recover normal operational state.

6.2. Initial Maintenance

DOCSIS 3.1 CMs use an initial maintenance region that is of significant size in frequency and time. This region is used to allow CMs to perform initial ranging which takes place in a broadcast mode and is used with a very coarse time adjustment as initially the CMs do not know what timing offset to apply to their transmissions. The design of this time and frequency resource area (ranging zone) must be done very carefully to minimize collisions that occur when that initial maintenance region is not big enough or to avoid wasting resources if it is too big. The frequency of occurrence of this initial maintenance region, which is called the insertion interval, must be carefully defined. In conjunction with the insertion interval, the back-off process is configured by adjusting the random range back-off parameter that is used after a collision occurred.

Above all the location of the ranging zone needs to be chosen in a higher part of the spectrum to minimize interference from the noise sources on the upstream. This has been a single point of failure in some early deployments of OFDMA, where the default ranging zone created by the CMTS was not in a robust region to allow all the CMs to range successfully.

Figure 22 depicts the initial ranging zone in a DOCSIS 3.1 OFDM channel, time and frequency map.

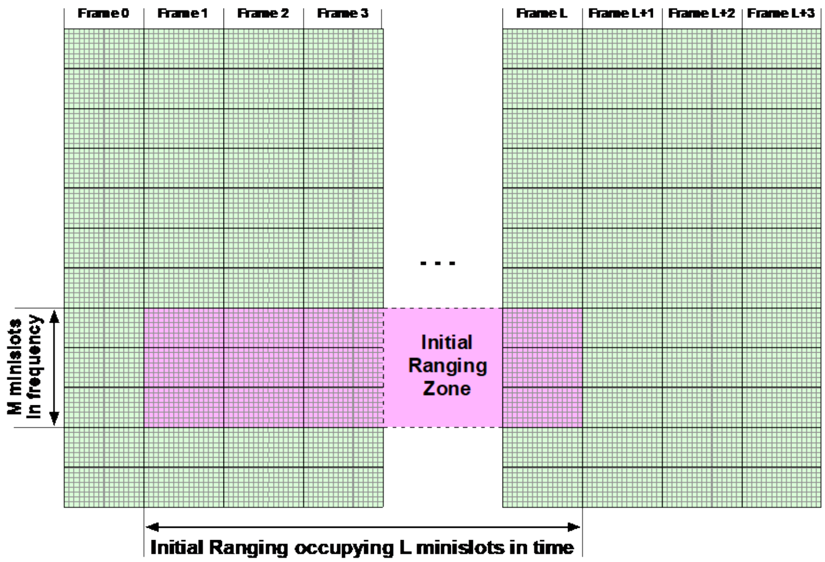

Figure 22 – Initial ranging zone in DOCSIS 3.1 in symbol versus frequency view

When ranging power levels, since initially the CM has no knowledge of what the right power level should be, it must go through many power levels before converging on a final value. The time this process lasts can be significant and any improvement mechanisms are worth implementing. One popular approach is to try the last good known state first. So, if a CM reboots it uses the last good known power and frequency configuration parameters that worked the last time the CM successfully initialized. Once the initial maintenance is successfully achieved, the CMTS and CM can communicate through unicast messages. Nevertheless, impaired channel conditions can force CMs into repeating this process.

6.3. Station Maintenance

The ranging response message is used to adjust transmit timing information, transmit power level information, transmit equalization information among other things. A CM that requires such an adjustment must wait until the CMTS acknowledges that the CM is ready to transmit by changing the ranging response messages from “RNG-RSP continue” to “RNG-RSP success”. The CMs cannot transmit data while still in “RNG-RSP continue”, a state which results in upstream inefficiency. These processes can extend for a while. Every time that a RNG-RSP message is not acknowledged, the T3 timeout counter is increased. After 16 successive T3 timeouts a T4 timeout occurs. The US channel is then deemed unusable, and the US channel enters in partial service mode or if it is the only channel, it forces re-initialization which represents an even greater impact on efficiency. In a good stable environment periodic maintenance is only used to conduct minor adjustment. Monitoring and quantifying the extend a CM spends beyond an uneventful ranging, helps you assess the impact on CM performance and efficiency.

6.4. Upstream Partial Service

The CMTS puts a CM in partial service mode when a CM is not able to range on a particular channel. This CM in partial service stops using the degraded channel. While station maintenance is the default mechanism to put CMs in partial service mode, alternative metrics can also be used to place a CM in partial service which include thresholds in SNR, as well as thresholds in uncorrectable as well as correctable codewords. This provides the operator with different tools to manage which CMs should have channels in partial service mode. These alternatives mechanisms also provide a hysteresis function so that a CM does not waste time going in and out of partial mode. The CMTS periodically will have the CM attempt to range on the degraded channel and check if it is still impacted. Going into partial service not only has the impact of not having US channel available to a CM, but also these transitions to and from partial service impact efficiency. CMTS algorithms on resolving partial service is a big factor in how quickly a CM returns to full service.

If a system struggles to achieve operational stability due to constant adjustment of power levels, constant adjustment of timing offset and a lack of convergence of equalization coefficients, efficiency is significantly impacted. Proper maintenance of a healthy US and DS channel prevents such occurrences.

6.5. Contention and Piggybacking Bandwidth Requests

The process of requesting upstream bandwidth in DOCSIS systems could also introduce inefficiencies. To request bandwidth, (minislots is the upstream resource currency), the CM takes advantage of a contention period designed for that purpose. This contention period allows CMs to send a request message containing how much data is desired. The size and periodicity of this contention period must be configured carefully by the CMTS to limit the number of collisions. Like contention in initial maintenance, when request messages collide, a random back-off is used to reduce the probability of further collisions

In addition to the contention request messages, DOCSIS systems also leverages a piggybacking process to reduce the number of contention requests needed thereby having a side benefit of reduce the likelihood of collisions. In piggybacking a bandwidth request is appended behind the packet just transmitted thereby not requiring a contention request.

Monitoring the number of regular request messages sent and number of piggyback requests sent in addition to monitoring the number of bytes requested and bytes granted are ways of determining how efficiently bandwidth is being allocated.

6.6. Traffic Characteristics and Signatures

The amount of traffic as well as its characteristics such as packet rate, packet rate variability and packet size distribution, stream duration as well as resources available are useful in determining how many mini-slots and how often to allocate request message opportunities. Since the traffic characteristics are dynamic, the CMTS will likely react and adapt according to these changing traffic conditions. The intelligence and efficiency of the CMTS scheduling algorithms will result in lower or higher transport efficiencies.

In the upstream, in DOCSIS 3.1 technology there are short, medium, and long codewords, each with different efficiencies. Very short packets will have lower efficiency and longer packets will have higher efficiencies. Also, in SC-QAM upstream operation each packet is preceded by a preamble and a guard time. This represents a time overhead that is fixed and for short payloads the payload size could be comparable to the preamble and guard-time, while with longer payloads such as is the case of 1500 Bytes, the preamble and the guard-time is only a small fraction of the packet duration. This inefficiency is aggravated with higher order modulations as the payload is further reduced in time. Efficient concatenation reduces the percentage of shorter packets thereby resulting in higher efficiency.

In DOCSIS systems the CMTS measures consumption and allocates resources in minislots, not bytes. One can have an accurate estimate of the resources being used when we leverage the MIBs that provide minislot utilization. An operator needs to make sure that the system is not starving for minislots, simply tracking the bytes can be deceiving.

Figure 23 shows an example of packet size distribution characteristics for some applications and aggregate packet size distribution of CMs on a service tier to highlight the specific impact that packet size distribution could have on performance. Some other packet size studies have found that 80% of upstream packets have a size below 80 bytes, and 90% of upstream packets are below 160 bytes.

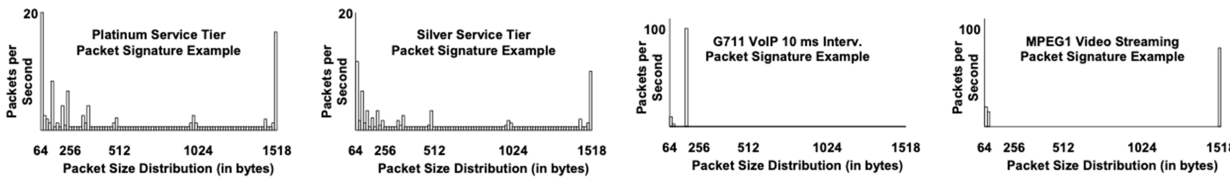

Figure 23 – Traffic Signatures for specific applications and for sample service tiers

In today's very high service tiers, bursty traffic is more prevalent. Tweaking the token bucket parameters is a tool that has been used in the past to optimize performance and can play well under bursty conditions, this could get large bursts of traffic out of the way to free up the channel for other transmissions.

6.7. Service Flow Configuration

The service flow configuration for the upstream directly affects the user experience. There are many Service flow parameters which an operator configures for an upstream service flow. These include CM Service Flow Parameters such as the Max sustained rate, Peak traffic rate, Max traffic burst, SF Priority, R/T Policy, Buffer control, AQM etc. Using the Buffer control setting (50 ms recommended) and enabling AQM will reduce the latencies observed by the user. The Max Sustained traffic rate for the service flow is set by the Service tiers offered by the operator.

One recommendation that will improve the user experience across all the service tiers is setting the Maximum Traffic Burst to something beyond the default of 3044 bytes, which is too small. The recommendation is to bump up the Max Traffic Burst (B) setting for a service flow to something in the range of 50 kBytes to 65 kBytes. Higher values of B can give users a boost in the initial data rate seen by the user, until the CM uses up the tokens allowed by the B value.

6.8. Profile Management / Upstream IUC management

DOCSIS 3.1 specifications introduced the concept of profiles so that performance of CMs could be optimized based on the SNR that they have available on the plant. DOCSIS systems have thereby evolved from the lowest common denominator conditions operation to a more dynamic opportunistic resource usage approach.

The CMTS assigns OFDMA IUCs (profiles) to the CMs based on the measured plant conditions. It is intended that the Data Profile IUC 13 is configured to be a robust OFDMA profile usable by any DOCSIS 3.1 CM served by that upstream channel and is used for all OFDMA data grants to modems which have not completed registration and for transport of MAC management messages and data grants after registration.

During or after modem registration, the CMTS has the option of assigning the CM to use any other configured data profile. Typically, the data profiles other than IUC 13 will be configured with higher modulation orders than IUC 13, although not all these profiles will be usable by all modems. The CMTS typically assigns IUC 13 and an additional data profile (one of IUC 5/6/9/10/11/12) to the CM for an OFDMA channel. The CMTS grants bandwidth on the OFDMA channel for data transmissions to a CM using this data profile (e.g., IUC 5). Now at some point the plant may have a noise ingress and interfere with the upstream transmissions. If the CMTS detects codeword errors on the IUC/profile used for a CM, it will temporarily use IUC 13 for data transmission, and initiate a dynamic change process to swap out the errored data profile with lower modulation data profile (e.g., swap out IUC 5 with IUC 6). This process continues (and the CMTS may further downgrade profiles) until the CMTS can receive codewords from the CM without errors. At a later point in time when the noise ingress goes away, the CMTS based on US RxMER measurements (using probes) or OFDMA Upstream Data Profile (OUDP) testing bursts, can evaluate the CM performance on a higher modulation profile and then choose to assign those profiles to the CM if appropriate. (e.g., upgrade from IUC 6 to IUC 5)

On noisy upstream spectrum, ingress noise issues may interfere with the signal enough to cause FEC codeword errors, to initiate the whole process described above, where a CM's profile is continuously being upgraded or downgraded, or in the worst case where the CM goes into partial service and is unable to use that channel for periods of time. Intermittent noise will cause this profile flapping behavior in the upstream OFDMA channel, such instability may be triggered by plant conditions and ultimately ends up impacting CM performance. (Similar profile flapping has been observed in noisy downstream channels as well).

DOCSIS CMTS settings need to be configured correctly to minimize the profile flapping. This includes the thresholds for the number of FEC errors, the RxMER dB thresholds for a particular modulation order, the time for which a CMTS downgrades a CM before re-evaluating if a profile is appropriate for a CM, the choice of looking at RxMER vs FEC errors in the upgrade/downgrade decision.

A well-designed Profile Management application (PMA) or system can create a good set of profiles, adjusted to the noise characteristics of that node, and minimize the profile flapping for each CM. A lot of the profile management concepts, algorithms, field experiences and deployment lessons have been described in a previous paper [US PMA].

The ability of the system to support a given QAM level depends on the RxMER values and the mappings to an appropriate QAM level, when creating a profile. These mappings are defined in [PHYv3.1] and are summarized in the Table below.

Table 3 – Upstream RxMER to QAM Level mapping

Upstream Constellation / Bit Loading	Upstream MER (dB)
QPSK	11.0
8 QAM	14.0
16 QAM	17.0
32 QAM	20.0
64 QAM	23.0
128 QAM	26.0
256 QAM	29.0
512 QAM	32.5
1024 QAM	35.5
2048 QAM (optional on CMTS)	39.0
4096 QAM (optional on CMTS)	43.0

6.9. Channel Conditions and PNM

In the previous sections, we have seen how unstable operations of DOCSIS processes directly impacts transport efficiency. This change from a stable to an unstable DOCSIS operation may have been triggered by changes in channel conditions. The DOCSIS system configuration may not be robust enough to overcome the channel impairments. These impairments may be sporadic or time of day sensitive. Even though these conditions may be difficult to detect, there are different parameters that are worth monitoring, including:

RxMER on a per channel and per subcarrier basis, FEC statistics, transmit and receive power levels and the different PNM tools such as equalization coefficients/channel estimate, full band capture, upstream triggered spectrum analysis, quite probes and CM specific probes analysis.

The goal is to have a good understanding of the conditions of the plant, troubleshoot and correct problems that are performance impacting and/or verify performance and profile assignment is as expected.

Of the upstream PNM functions defined in the specifications, we have obtained and looked at the following data sets: upstream RxMER per subcarrier, upstream Equalizer Coefficients (CMTS), upstream Triggered Spectrum Capture Analysis, and upstream FEC Statistics (aggregated over the channel / all CMs). All of these have useful PNM applications in understanding the upstream performance. Upstream and downstream equalization coefficients contain valuable channel information, please refer to [PreEq Analysis] for further details.

So far, we have not been able to get access to data from CMTSs for the following PNM features: upstream Capture for Active and Quiet Probe, upstream Impulse Noise Statistics, upstream Histogram, upstream Channel Power.   These upstream PNM functionality needs to be implemented on CMTS platforms and then further investigated and applications developed.

7. Upstream Split Migration Scenarios

There are a few different approaches for increasing the upstream bandwidth within the cable plant. The DOCSIS 3.1 system will have options of several split configurations that can be exercised based on traffic demand, services offered and the capability of the cable plant. These include the classic Low-split, the mid-split, and high-split frequency plans for DOCSIS 3.1 equipment and going to the ultra-high split options with DOCSIS 4.0 technology. To reach the target service goal in the upstream direction, plant changes on the upstream/downstream spectrum split are expected.

7.1. Adding Upstream Spectrum

Information theory says there are two ways to increase channel capacity:

- Add spectrum
- Increase the signal-to-noise ratio of the spectrum in use

In brief, doubling the upstream spectrum will double the upstream capacity; that is, a 100% increase in capacity. Staying in the same spectrum and increasing the Gaussian SNR by 6 dB to go from 256 QAM to 1024 QAM provides a 25% increase in capacity. Below 20MHz the noise is usually bursty, so a Gaussian noise assumption is not quite appropriate. Furthermore, the impulsive noise is spread to all frequency domain symbols by the FFT at the CMTS receiver. At some point, increasing the SNR has diminishing returns and the big gains will come from increasing the amount of spectrum allocated to the upstream.

Hence DOCSIS 3.1 technology includes both mid-split and a high-split options will be discussed in more detail. DOCSIS 4.0 technology includes ultra-high splits which provide additional upstream spectrum. Figure 24 shows the relative amounts of upstream spectrum with sub-split, mid-split, and high-split plant in DOCSIS 3.1 systems.

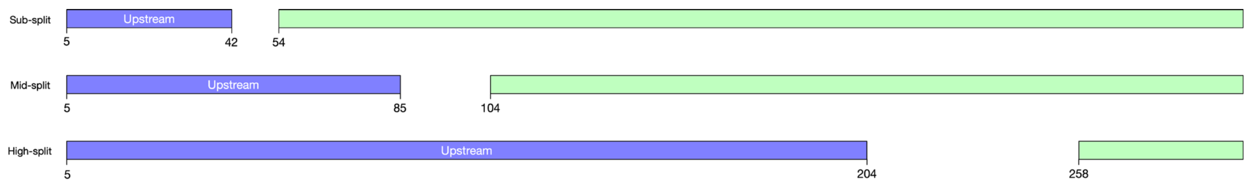

Figure 24 – Split Scenarios

7.1.1. Impact on Legacy Customer Premise Equipment

Existing cable modems deployed in the field may not support different splits, however, they will operate with the DOCSIS channels that reside in their existing upstream spectrum option. Set-top boxes (STBs) are limited to forward data channel (FDC) frequencies between 70 MHz and 130 MHz in the downstream direction as described in the SCTE 55-1 and 55-2 standards. As a result, these set-top boxes can continue to operate on a mid-split network but not a high-split network.

7.1.2. More Upstream spectrum and Downstream Spectrum

As more upstream spectrum is added, the downstream starts higher in the spectrum. For example, with a sub-split the downstream spectrum starts at 54 MHz. With a high-split network, the downstream starts at 258 MHz, and to maintain enough downstream spectrum the top end of the useful spectrum is usually increased, and the typical number is 1000 MHz or 1218 MHz.

As a summary, Table 4 shows the aggregate upstream throughput for the examples given in the following sections.

Table 4 – Upstream Throughput Examples

Case	Channels	Approximate Throughput
Sub-split Use Cases		
1	Four 3.2 MHz SC-QAM at 32 QAM	44 Mbps
2	Two 3.2 MHz SC-QAM at 32 QAM and Two 6.4 MHz SC-QAM at 64 QAM	72 Mbps
3	Four 6.4 MHz SC-QAM at 64 QAM	100 Mbps
4	Four 6.4 MHz SC-QAM at 64 QAM and One 3.2 MHz SC-QAM at 64 QAM and One 1.6 MHz SC-QAM at 64 QAM	115 Mbps
5	Two 6.4 MHz SC-QAM at 64 QAM and One 3.2 MHz SC-QAM at 64 QAM and 13 MHz OFDMA at 1024 QAM	160 Mbps
Mid-split Use Cases		
6	Ten 6.4 MHz SC-QAM at 64 QAM	250 Mbps
7	Four 6.4 MHz SC-QAM at 64 QAM and 40 MHz OFDMA at 1024 QAM	400 Mbps

8	70 MHz OFDMA at 1024 QAM	525 Mbps
High-split Use Cases		
9	Four 6.4 MHz SC-QAM at 64 QAM and One 66 MHz OFDMA at 1024 QAM and One 96 MHz OFDMA at 1024 QAM	1350 Mbps (1.3 Gbps)
10	Two 96 MHz OFDMA at 1024 QAM	1700 Mbps (1.7 Gbps)

7.2. Sub-split HFC network

A sub-split HFC network has a return path up to 42 MHz and has room to carry multiple upstream DOCSIS channels. In North America, sub-split is a hold-over from the days of analog television channels and cable-ready TV sets (both of which are rapidly going extinct, if not already there). Sub-split allowed the old Channel 2 to be carried on the coaxial cable starting at 54 MHz. The old analog Channel 2 has been replaced by digital TV carriers and IPTV, so the reasons for sticking to sub-split are diminishing. The following sections will show examples of how to get the most out of various configurations of an upstream channel.

The configurations in Cases 1-5 shown below will fit within the spectrum of a sub-split HFC network and as can be seen the upstream capacity can more than double by increasing both the amount of spectrum used for broadband and increasing the modulation order of the SC-QAM channels.

7.2.1. Making the most of Upstream SC-QAM Channels

For sub-split operation operators should experiment with various recommendations to get more upstream capacity. Strategies include:

7.2.1.1. Case 1: Remove Guard Bands

Removing large guard-band between upstream SC-QAM carriers; these guard bands are not needed. The DOCSIS specifications are written such that upstream SC-QAM channels can be directly adjacent to each other. Deployments exist where the upstream SC-QAM carriers have large guard-bands between them as shown in Figure 25 below.

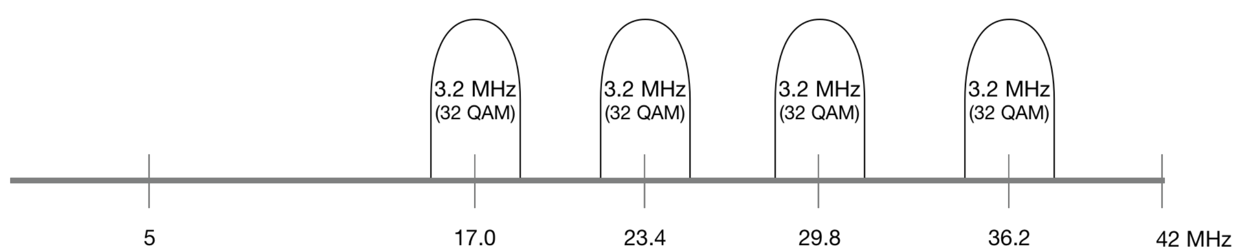

Figure 25 – Large Guard Bands Between Upstream Channels

With four 3.2 MHz wide channels, operating at 32 QAM, the aggregate throughput of this upstream channel is about 44 Mbps which leaves a lot of unused capacity on the table. The guard bands between upstream SC-QAM carriers are not needed; the carriers may be directly adjacent to each other on each side with no guard-band. However, as you place a carrier near the roll-off starting at 42 MHz it is best to leave around a 500 kHz guard-band from 41.5 MHz to 42 MHz to avoid non-linearities related to amplifier cascades.

7.2.1.2. Case 2: Increase Modulation Orders

The next step for an operator would be to increase the modulation orders to 64 QAM for the upstream SC-QAM carriers. The DOCSIS SC-QAM technology is quite resilient, and the recommendation is to increase carriers to 64 QAM modulation, widen the carriers to use available spectrum, and remove the guard-band between the carriers. As shown in Figure 26 below, the moves should be done stepwise to confirm proper operation.

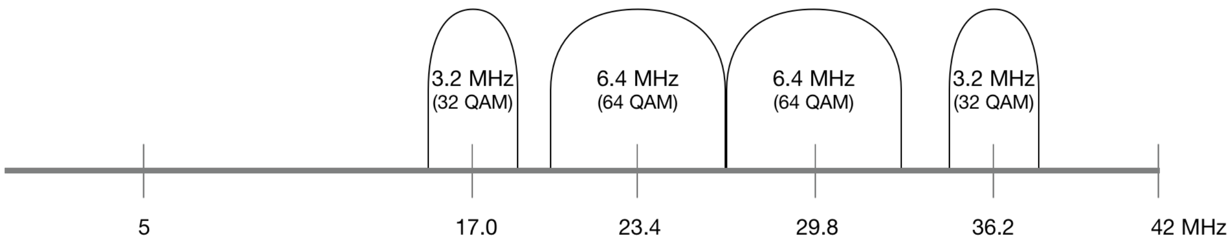

Figure 26 – Increasing Channel Width To Increase Upstream Capacity

With both two 3.2 MHz wide channels operating at 32 QAM and two 6.4 MHz wide channels operating at 64 QAM, the aggregate throughput of this upstream channel is about 72 Mbps.

7.2.1.3. Case 3: Use available spectrum

Another option for an operator would be to use all the available spectrum and using the widest channels. Figure 27 shows using all available spectrum, which could include widening all upstream carriers and abutting those carriers.

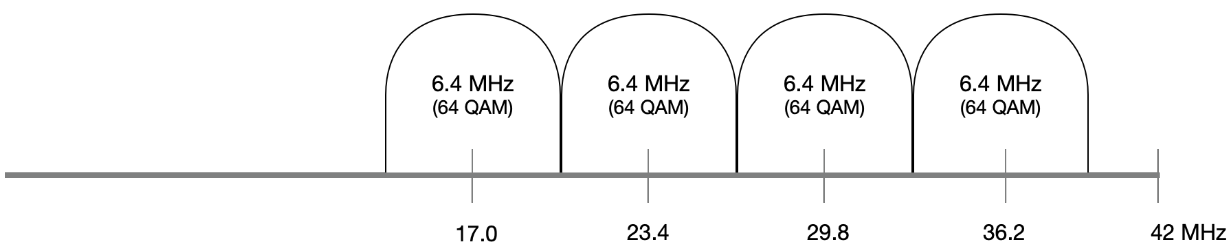

Figure 27 – Maximizing Upstream Capacity With Wide Upstream Channels and Small Guard Bands

With four 6.4 MHz wide channels operating at 64 QAM, the aggregate throughput of this upstream channel is about 100 Mbps.

7.2.1.4. Case 4: Squeeze in an additional carrier

As a final step for a sub-split network, operators could try adding one or more additional upstream SC-QAM carriers, up high or down low as shown in Figure 28 . The paper [Bandwidth Growth] draws the conclusion that even adding 10% additional upstream capacity can help alleviate upstream congestion.

Operators have been successful at adding new carriers, which is a testament to maintaining the plant more diligently over the last decade.

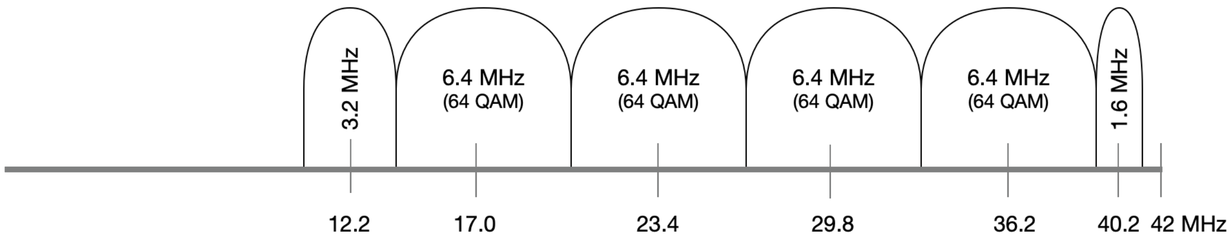

Figure 28 – Adding Additional Upstream Channels

This example adds both one 3.2 MHz wide channels operating at 64 QAM and one 1.6 MHz wide channels operating at 64 QAM, the aggregate throughput of this upstream channel is about 115 Mbps.

The DOCSIS technology has many options for optimizing channel layout and is different from even a decade ago. The digital transceivers have increased in sensitivity and capabilities to enable optionality not thought of even 5 years ago. Homes passed has decreased which lowers the effect of noise funneling at low frequencies. Cascades have shortened which lessens the impact of group delay close to the diplex filter cut-off frequency. Operators have been successfully running narrow carriers (typically with lower order modulation) both down to 10 MHz and closer to the diplex filter.

7.2.2. Sub-Split and OFDMA Channel

If the deployment must stay sub-split and there are enough DOCSIS 3.1 modems on the network, consider removing upstream SC-QAM channels and replacing them with an OFDMA channel. OFDMA technology makes better use of spectrum because it can operate nominally at 1024 QAM whereas an upstream SC-QAM channel is limited to 64 QAM.

An OFDMA channel low in the spectrum will be susceptible to impulse noise. Very strong noise impulses in the time domain might result in the loss of blocks of symbols at the receiver because that impulse noise is spread across multiple symbols.

7.2.2.1. Case 5: Add a small OFDMA Channel

The recommendation here is to put the OFDMA channel high in the spectrum where there is less impulse noise, (at least initially until the operators get comfortable with OFDMA) as shown in Figure 29.

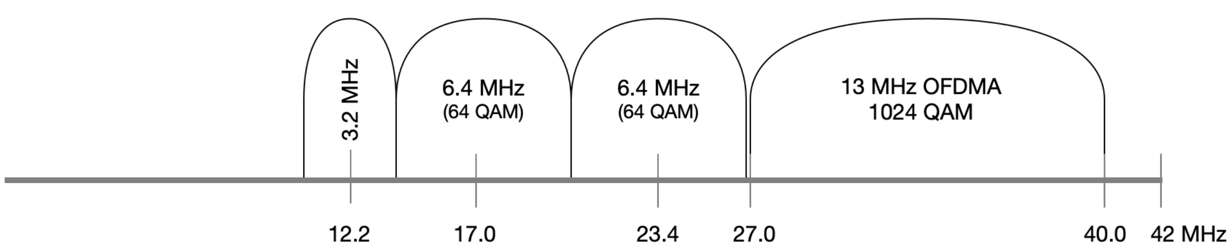

Figure 29 – Sub-split HFC Network with OFDMA Channel

The arrangement of channels shown above can provide an aggregate upstream capacity of 160 Mbps, showing the efficiency of the OFDMA spectrum as compared to SC-QAM spectrum. Note though that only DOCSIS 3.1 modems can take advantage of the OFDMA channel, therefore, there should be a large percentage of DOCSIS 3.1 modems on that plant segment.

Time and Frequency Division Multiplexing (TaFDM) is discussed in Section 5.2.

7.3. Mid-split HFC network

A mid-split network has a return path up to 85 MHz, or two times the spectrum of a sub-split network. With a mid-split, the forward path begins around 108 MHz so set-top box-based video services should be able to be maintained because the forward data channel can be up to as 130 MHz.

The additional upstream spectrum provided by a mid-split network can provide about 500 Mbps of capacity.

7.3.1.1. Case 6: Start with SCQAMs

In terms of additional upstream spectrum, a mid-split can be configured in several ways. As shown in Figure 30, a mid-split can fit ten traditional upstream SC-QAM carriers of 6.4 MHz width (for a total of 64 MHz of upstream spectrum allocated to broadband).

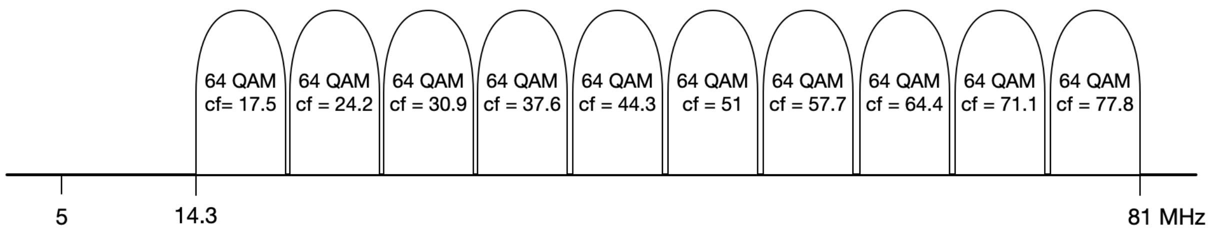

Figure 30 – Mid-split with 10 SC-QAM channels

This configuration can yield up to 250 Mbps of aggregate upstream capacity. Note that DOCSIS technology allows SC-QAM to be modulated up to 256 QAM and no higher. That is, with SC-QAM the options of 512 QAM and 1024 QAM are not available. However, the newer DOCSIS 3.1 technology does allow higher order QAM modulation.

7.3.1.2. Case 7: Add an OFDMA channel

As shown in Figure 31, another example is the traditional 4 upstream SC-QAMs up to 40 MHz, and then 40 MHz of OFDMA running at 1024 QAM from 42 MHz to 82 MHz.

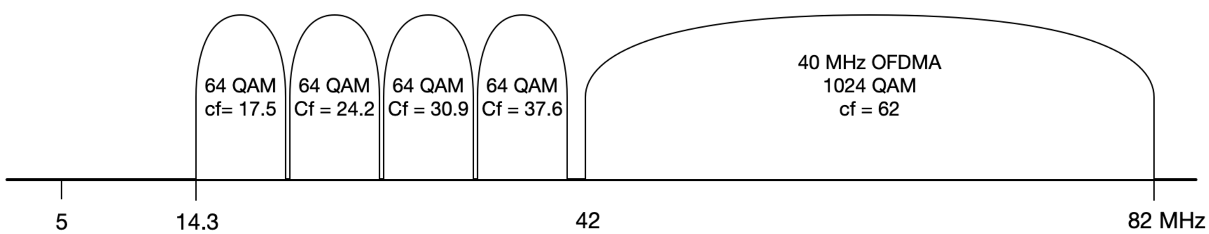

Figure 31 – Mid-split with Both SC-QAM and OFDMA channels

This configuration can yield up to 400 Mbps of aggregate upstream capacity and increase of 150 Mbps using the same spectrum because the OFDMA carrier can operate at a higher order of QAM modulation than the single-carrier QAMs.

7.3.1.3. Case 8: Replace with OFDMA

As shown in Figure 32, another example is shown using 70 MHz of OFDMA at 1024 QAM which can offer 525 Mbps of capacity.

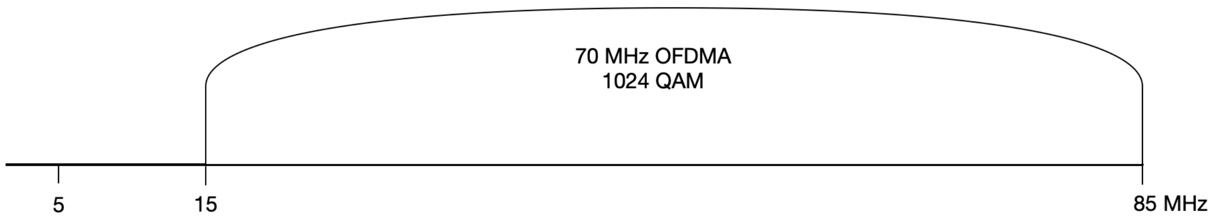

Figure 32 – Mid-split with All OFDMA channel

Though it may be necessary to retain a single SC-QAM upstream channel for DOCSIS 2.0 modems.

7.4. High-Split HFC Network

A high-split network has a return path up to 204 MHz, or four times the spectrum of a sub-split network. With a high-split, the forward path begins around 258 MHz hence set-top box-based video services cannot be maintained because the forward data channel (FDC) can only be moved as high as 130 MHz as described in the SCTE 55 standards.

However, the additional upstream spectrum provided by a high-split network can provide more than 1.5 Gbps of capacity.

7.4.1.1. Case 9: Add a second OFDMA

Figure 33 shows a configuration that retains the four SC-QAMs for DOCSIS 3.0 (and earlier) modems and uses the rest of the spectrum for OFDMA using up to 1024 QAM modulation.

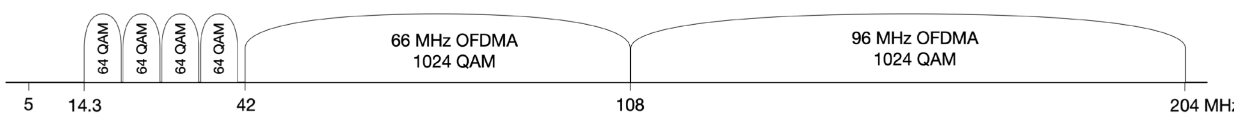

Figure 33 – High-split with Both SC-QAM and OFDMA Channels

This configuration can provide up to 1.3 Gbps of capacity while retaining backward compatibility for DOCSIS 3.0 and earlier modems.

7.4.1.2. Case 10: Full OFDMA

Figure 34 shows all the high-split spectrum using OFDMA and this configuration is capable of over 1.6 Gbps of capacity.

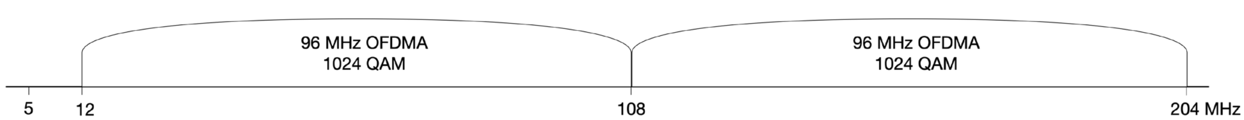

Figure 34 – High-split with All OFDMA Channels

7.5. Ultra-high split Networks

The DOCSIS 4.0 specification builds upon DOCSIS 3.1 OFDM and OFDMA technology with an extended Frequency Division Duplex (FDD) DOCSIS alternative. DOCSIS 4.0 FDD supports both mid-split and high-split and provides extended upstream splits up to 684 MHz in an operational band plan that is referred to as Ultra-high Split (UHS). DOCSIS 4.0 FDD also introduces expansion of usable downstream spectrum up to 1794 MHz to support the higher upstream splits.

8. Conclusion

For the past 20 years in most of the world, the DOCSIS upstream has included up to four SC-QAM channels. With DOCSIS 3.1 technology, OFDMA technology became available for the upstream. As compared to SC-QAM, OFDMA has more configuration options which can be used to optimize upstream efficiency.

As DOCSIS technologies evolve, so do the tools and understanding the operators need to ensure that their networks are running as efficiently as possible, while still maintaining robust service offerings. There is a fundamental trade-off between robustness and throughput. This paper analyzes those trade-offs and discusses methods to increase the efficiency of the DOCSIS upstream. Included are discussions on:

- SC-QAM technology,
- OFDMA technology,
- Interference considerations when combining SC-QAM and OFDMA on the same upstream,
- DOCSIS protocol efficiency,
- Adding more spectrum to the upstream.

A first step to increase upstream efficiency is to use as much of the available spectrum as possible with existing SC-QAM technology. SC-QAM technology has a well-established history and experience, though potentially less flexible when compared to OFDMA technology. There are still multiple parameters like minislot size or FEC size choices which when optimized will improve the upstream efficiency.

A second step includes consolidating SC-QAM technology and allocating spectrum to OFDMA technology. Mixing both SC-QAM and OFDMA technology on the same upstream allows support for older cable modems (DOCSIS 3.0 technology and earlier) as the penetration of DOCSIS 3.1 modems increases. Adding OFDMA allows the DOCSIS 3.1 modems to take advantage of the newer technology to increase the overall throughput of the upstream.

OFDMA technology can provide more flexibility for spectral efficiency but carries more configuration complexity to be considered across all networks as well as individual network conditions. A prerequisite for deploying OFDMA technology is working through the optimal configuration of the upstream parameters which include trade-offs for robustness and capacity. For example, an operator needs to understand how a CMTS enforces guard bands in the upstream and think through the implementations of TaFDM on the CMTS. The overhead of the cyclic prefix can gain robustness but at the cost of capacity. The 2K FFT is more robust and the 4K FFT provides more throughput. 4k is more efficient than 2k for reasons of narrower receive spectral window and required CP being half the symbol period, 4k also interferes with adjacent SC-QAM less. Designing profiles to match the channel conditions on a particular part of the plant, using a profile management application (for OFDMA and SCQAM channels) allows an operator to gain robust operation and extra capacity. OFDMA is worth the effort because it is the new foundation for upstream technology just as OFDM has become the new downstream technology.

Lastly, to get more out of the upstream consider adding new spectrum by changing the upstream split. Just continuing to segment the HFC network does not provide more speed, rather, segmentation replicates the existing spectrum and provides the same speeds for fewer users which does have benefit but does not enable the true capability of the HFC network. To get to both more upstream capacity and really fast speeds, more spectrum must be allocated to the upstream. This new spectrum should be allocated to OFDMA technology to make the best use of it.

Abbreviations

AP	access point
AWGN	average white gaussian noise
bps	bits per second
CM	cable modem
CMTS	cable modem termination system
CP	cyclic prefix
CW	continuous wave
dB	decibel
DOCSIS	data over cable service interface specifications
FDD	frequency division duplex
FDX	full duplex
FEC	forward error correction
FFT	fast Fourier transform
GB	guard band
Gbps	gigabits per second
HD	high definition
Hz	hertz
IUC	interval usage code
kHz	kilohertz
MAC	media access control
Mbps	megabits per second
MER	modulation error ratio
MERperSC	modulation error ratio per subcarrier
MHz	Mega hertz
MSyms/sec	Mega symbols per second
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
PMA	profile management application
PNM	proactive network maintenance
QAM	quadrature amplitude modulation
QPSK	quadrature phase shift keying
RNG-RSP	ranging response message
RP	roll-off period
RxMER	received modulation error ratio
SC-QAM	single carrier QAM
SCTE	Society of Cable Telecommunications Engineers
TaFDM	time and frequency division multiplexing
UHS	Ultra high split

Bibliography & References

[DOCSIS PHYv3.1] *DOCSIS 3.1 Physical Layer Specification,* CM-SP-PHYv3.1-I18-210125, 2021, Cable Television Laboratories, Inc.

[DOCSIS MULPIv3.1] *DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification*, CM-SP-MULPIv3.1-I21-201020, 2020, Cable Television Laboratories, Inc.

[US PMA] *Field Experiences with US OFDMA and Using US Profile Management,* K. Sundaresan, J. Zhu, and J. P. Fernandes, *SCTE Expo 2020*

[D31 Capacity] *Accurately Estimating D3.1 Channel Capacity,* Karthik Sundaresan, SCTE Expo 2017

[PreEq Analysis] OFDMA predistortion coefficient and OFDM estimation decoding and analysis, Tom Williams, Jason Rupe, Alberto Campos, SCTE Expo 2021

[Bandwidth Growth] *Managing the Coronavirus Bandwidth Surge: How to Cope with the Spikes and Long-term Growth,* John Ulm & Dr. Thomas Cloonan, CommScope, SCTE Expo 2020

