DOCSIS 3.1 Downstream Early Lessons Learned

A Technical Paper prepared for SCTE/ISBE by

John J. Downey
CMTS Technical Leader
Cisco Systems Inc.
RTP, NC
919-392-9150
jdowney@cisco.com

Table of Contents

List of Figures

List of Tables

Introduction

DOCSIS 3.1 (D3.1) systems have been deployed in operating cable networks and some lessons have been learned. Allocating dedicated spectrum for new D3.1 cable modems (CMs) is a scary proposition and cable operators know what's at stake. To compete with the marketing hype of fiber-to-the-home (FTTH), the cable industry has many tricks up its sleeves such as D3.1.

The D3.1 downstream (DS) has been the focus for the last year with the upstream (US) just around the corner. D3.1 offers speeds required by our customers and the tools and feature-set to make it work properly for the highest availability possible. A few features that help achieve "self-healing" are: DS resiliency, graceful profile management, mixed-modulation profiles, and more. For on-going operational monitoring, the DOCSIS standard has built-in proactive network maintenance (PNM) features to identify and locate hybrid fiber-coax (HFC) impairments quickly before customers even know they exist.

This paper and presentation explore the intricate details of real-world scenarios of DOCSIS 3.1 deployments using best practices to optimize the cable modem termination system (CMTS) for maximum throughput even in an impaired environment: the type of environments every cable operator eventually experiences.

Content

1. General Basics

The D3.1 DS uses orthogonal frequency-division multiplexing (OFDM) with spectrum options of 24 to 192 MHz in block/channel widths. The actual spectrum is 204.8 MHz, but that is not what is seen from an RF standpoint when viewing on a spectrum analyzer. A byproduct of this is that one may notice CM time offsets 20 times larger because DOCSIS 1.x, 2.0 and 3.0 are based on 10.24 MHz clocking. This D3.1 204.8 MHz clocking can create a misconception of much higher time offsets.

D3.1 DS modulation schemes range from 16-QAM to 4096-QAM with even higher options possible in the future. Five data profiles can be supported for modems to choose from with automatic profile selection. Some of the profiles could consist of mixed modulation with five to seven different modulation schemes in the entire DS block/channel. This may be advantageous for known plant roll-off issues above 750 MHz or 860 MHz, etc.

Exclusion bands are also possible to avoid ingress or create a null space to inject a carrier for amplifier automatic gain control (AGC), leakage testing, or alignment tones for rough balancing.

D3.1 CMs can cross-bond traffic between legacy DOCSIS carriers known as single carrier QAM (SC-QAM) signals with one to two OFDM channels. This will be the typical deployment scenario for many years to come to support legacy CMs and provide very high peak speeds for D3.1 CMs with 1 Gbps as a goal. The aim for these D3.1 CMs is to prefer the OFDM ch before utilizing the SC-QAM signals so as not to "starve" out legacy CMs. "Primary capability" can be either a SC-QAM or the OFDM block with its coinciding physical layer link channel (PLC).

At this time, few systems have any purely D3.1 service groups (SGs). SC-QAM signals will probably be present for many years for D2.0 and D3.0 CMs. With that said, it doesn't mean a SC-QAM signal is necessary for D3.1 CMs to operate, but it is advantageous to cross-bond for the following reasons:

1. Bigger channels are always best. Higher peak rates can be achieved and less need for load balance to occur from legacy CMs and between D3.0 bonding groups.
    a. D3.1 prefer OFDM for traffic before SC-QAM signals are used anyway, so no concern for "starving" legacy CMs.
2. Having a SC-QAM primary allows DS resiliency to work a bit better. More on that later.
3. SC-QAM signals for voice over IP (VoIP) traffic and other low-latency queuing (LLQ) traffic may be needed.
4. Not all SC-QAM signals should be primary-capable when we start having 24 and more channels. Almost 2-3 Mbps of overhead is created when a SC-QAM signal is primary.
5. In case of an emergency situation like overheating, one could shut down the OFDM first as it could be a major contributor.

Current D3.1 CMs on the market support 32 SC-QAM plus 2 OFDM blocks. They also support up to 4096-QAM (also called 4K-QAM) even though 8K-QAM and 16K-QAM are options in the specification. CM spectrum support is 1 GHz with most at 1.218 GHz. 1.794 GHz is an option in the spec, but may or may not be pursued depending on full duplex DOCSIS (FDX), which is under development at Cablelabs.

In regards to US support, D3.1 CMs support eight SC-QAM signals along with two 96 MHz (max) orthogonal frequency-division multiple access (OFDMA) blocks. Even if the CM US chipset supports a return path upper frequency limit of 204 MHz, the internal diplex filer may be hardware or software limited to 42 MHz or 85 MHz.

Another potential issue to be aware of is DS power levels supported from a CMTS connector. This is specified by the DOCSIS Downstream Radio Frequency Interface (DRFI) spec. The level range vs channel loading is listed in Table 1. From this table, one can see that more spectrum (chs) located on a DS connector, whether video, legacy SC-QAM or D3.1 OFDM, etc., will lead to lower max output available and precautions need to be taken. Headend wiring may need to shortened or changed from mini coax back to Series 6, passives swapped/replaced, etc.

Table 1 - Max Carrier/OFDM Loading to Channel Level Range

Max Carrier	No OFDM	24 MHz OFDM	48 MHz OFDM	96 MHz OFDM	192 MHz OFDM	384 MHz OFDM
8	**41 – 50**	39 – 48	37 – 46	35 – 44	32 – 41	29 – 38
16	**37 – 46**	36 – 45	35 – 44	34 – 43	31 – 40	29 – 38
24	**35 – 44**	34 – 43	34 – 43	32 – 41	31 – 40	28 – 37
32	**34 – 43**	33 – 42	32 – 41	31 – 40	30 – 39	28 – 37
48	**31 – 40**	31 – 40	31 – 40	30 – 39	29 – 38	27 – 36
64	**30 - 39**	30 – 39	29 – 38	29 – 38	28 – 37	26 – 35
96	**28 – 37**	28 – 37	27 – 36	27 – 36	26 – 35	25 – 34
128	**26 – 35**	26 – 35	26 – 35	26 – 35	25 – 34	24 – 33
158	**25 – 34**	25 – 34	25 – 34	25 – 34	24 – 33	Not Possible

The ranges listed in Table 1 are in dBmV and the maximum value is 1 dB above the DRFI spec. A CMTS with a tilt function could help engineers achieve flat inputs to the DS optical lasers with options of linear or non-linear tilt along with power level offsets as well.

2. Spectrum Allocation and Thoughts

Current deployments are looking to extend the DS out to 1 GHz or 1.218 GHz for actives/passives. D3.1 allows a higher option of 1.794 GHz, but current CMs don’t support it.

The D3.1 192 MHz block(s) starts at 108 MHz optional with a 258 MHz starting frequency mandatory. The D3.1 CM’s filter and/or scanning table may negate any D3.0 operation < 261 MHz center frequency for SC-QAM signals. Potential ingress sources reside throughout the spectrum from long term evolution (LTE) 4G mobile to Multimedia over Coax Alliance (MoCA), garage door openers, off-air broadcasters, etc. Proper planning for channel allocation/location in the spectrum should be done for the intended market.

US decisions also affect the low-end of the DS. One also needs to consider DS amplifier AGC/ALC and potentially legacy settop box out-of-band control channels. Figure 1 displays potential spectrum allocation and ingress sources along with typical diplex filter splits.

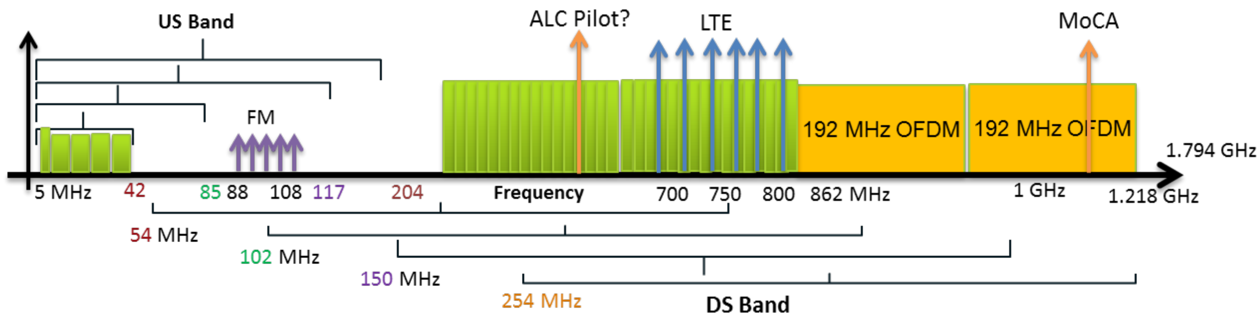

Figure 1 - DS Spectrum Allocation

3. Interference Resiliency Testing

To prove the resiliency of D3.1 with its added robustness from low density parity check (LDPC) forward error correction along with time and frequency interleaving, some testing was done with SC-QAM signals simulating wideband ingress. Some baseline configs and steps are listed below:

- An OFDM channel with 192 MHz width using 25 kHz subcarrier spacing.
- A DS bonding group with one SC-QAM plus OFDM block with both as primary.
- Data traffic set to ~ 1.4 Gbps through a few CMs.
- A test CM forced to use 4096 QAM.
- An interference source comprising eight SC-QAM signals from a DS port and combined with the working port.
- Variable padding added to the interfering signal starting with 30 dB.
- A speed test was observed to make sure no packets were lost.
- Receive MER (RxMER) readings were documented along with the “break point”.

Table 2 is taken from the DOCSIS 3.1 PHY spec’s Table 7-41. This table lists the suggested operational RxMER for the corresponding modulation schemes and bit loading (bits/symbol) for that modulation. These values are used by the CMTS to make the decision for graceful profile management.

Table 2 - RxMER to Bit Loading Mapping

RxMER (in ¼ dB)	RxMER (in dB)	QAM	Bit Loading
60	15	16	4
84	21	64	6
96	24	128	7
108	27	256	8
122	30.5	512	9
136	34	1024	10
148	37	2048	11
164	41	4096	12
184	46	8192	13
208	51	16384	14

Note: It's in this author's opinion that these thresholds are ~6 dB too conservative and it may be worth having a correction factor of 2-4 dB configured in the CMTS.

Figures 2 through 6 display the test results from various ingress levels and show RxMER vs subcarrier frequency. Figure 2 is the baseline 192 MHz wide OFDM block with no ingress. The average RxMER for all the active subcarriers is 47.4 dB.

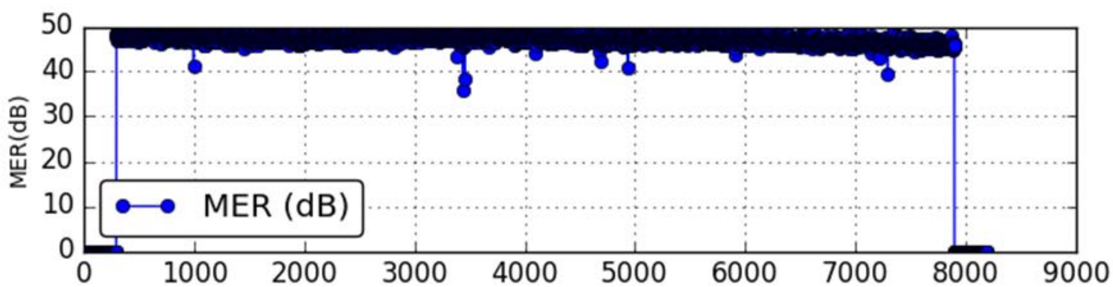

Figure 2 – 47.40 dB Avg RxMER with No Interference

Note: It's this author's opinion that this modem is reporting the pilot RxMER without the proper 6 dB correction factor and hence reading 6 dB higher than it should.

Figure 3 shows an average RxMER of 41.88 dB with the ingress activated with 20 dB padding vs the 30 dB start value. As can be seen in the figure, the eight SC-QAM signals of ingress are causing lower RxMER values of ~ 28 dB in that region. With such low RxMER and affecting 48/192 = 25% of the block, the internal threshold table would suggest this modem drop from 4K-QAM to 256-QAM!

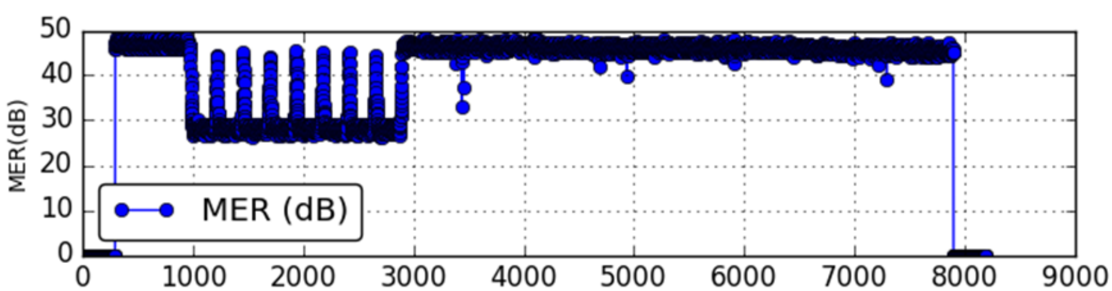

Figure 3 - 41.88 dB Avg RxMER with 20 dB Padding

Interesting enough, this modem did not drop any packets and continued at 4K-QAM at full throughput. All active subcarriers were being used at 4K-QAM even in the region where RxMER values were ~ 28 dB! This is believed to be possible because of the fast Fourier transform (FFT) functionality along with LDPC and time and frequency interleaving.

Note: The secondary observation was a very slight RxMER decrease for all subcarriers even where there was no ingress. The belief here is that the ingress is being spread across all subcarriers even though it is actually localized in specific spectrum.

Figure 4 shows the average RxMER dropping to 39.13 dB and all the subcarriers getting affected in addition to just the ingress localized subcarriers.

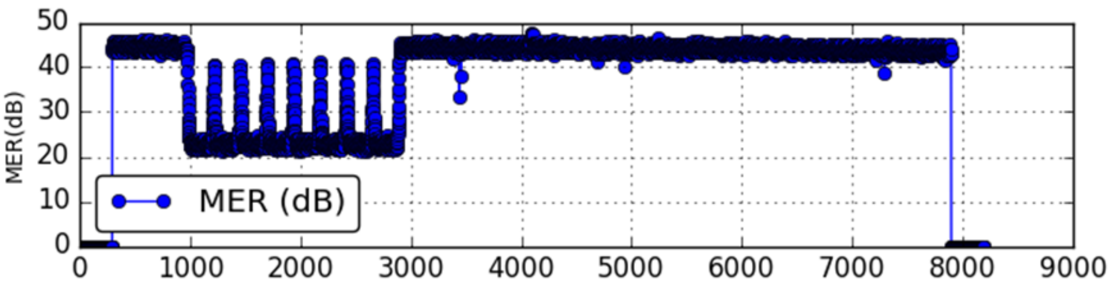

Figure 4 - 39.13 dB Avg RxMER with 5 dB More Ingress

Note: All RxMER values dropped after the RxMER in the 48 MHz interference spectrum dips below~ 25 dB. This was observed even when the ingress was just one SC-QAM signal.

Figure 5 continues this trend and we still did not observe any packet drops!

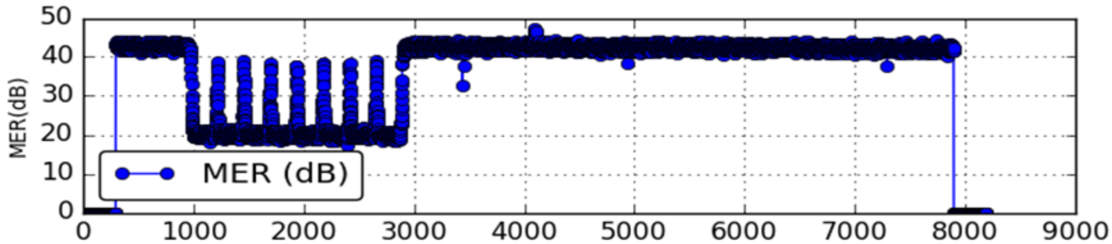

Figure 5 – 37.14 dB Avg RxMER with 8 dB More Ingress

Figure 6 increase the ingress so that the localized subcarriers drop below 20 dB RxMER and still operating without dropped packets. Now all the other subcarriers are definitely getting affected and almost exceeding where 4K-QAM threshold is listed. The avg RxMER is definitely past the threshold and very near the actual "break-point".

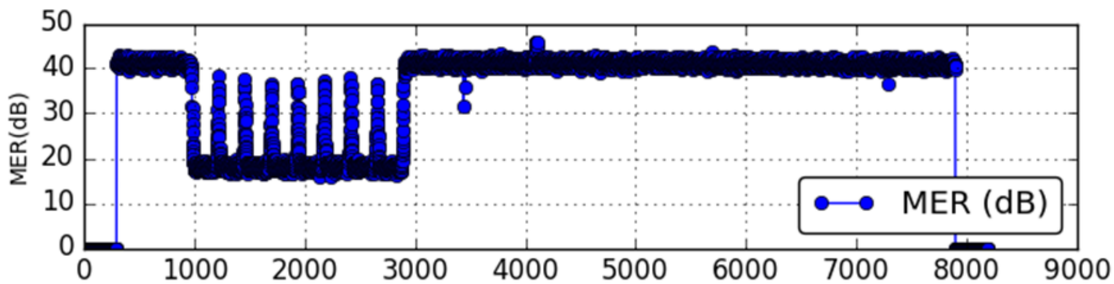

Figure 6 - 35.59 dB Avg RxMER with 10 dB More Ingress

The following observations were made:

- Steady-state interference did not affect throughput as anticipated even at very low RxMER readings. Possibly an added effect of LDPC, freq interleaving and FFT functionality.
- Overall subcarrier RxMER values dropped when the impaired subcarrier RxMER got < ~25 dB.
    - Actual impaired spectrum may be a deciding factor, but not observed.
    - Not an issue when exclusion bands used.
    - Possibly an added effect of FFT functionality and/or overdrive.
- Some CMs appear to incorrectly report RxMER ~ 6 dB higher than other CMs and may be based incorrectly on pilot readings.
- Mixed-modulation profiles probably not necessary for ingress, but for roll-off, maybe.
- The PLC uses 16-QAM and is very robust, but avoiding known ingress spectrum is still a best practice worth following.

4. Graceful Profile Management

The CMTS incorporates a feature defined in D3.1 as profile management. This can be done with a software defined network (SDN) application or within the CMTS itself. The CMTS transmits to the CM on a control profile (profile 0) until the CM RxMER information for all active subcarriers is received by the CMTS to determine the optimal configured profile. The CMTS periodically polls each CM to gather RxMER of every active subcarrier and recommends the best profile configured in the CMTS.

Even if thresholds are aggressive, too lenient, or no updated RxMER readings from a CM, a “catch-all” is still used and that is the processing of a cm-status message of “unfit profile”.

From testing results and “real-world” deployments, the following deviations from Cisco CMTS defaults are suggested: (refer to your CMTS vendor for specific command syntax)

1. Changing modulation for an entire block for a small amount of subcarriers is not in our best interest. The default is to ignore 2%, but increasing that to 10% is suggested.
    - o
      ```
      cBR8(config)#cable downstream ofdm-prof-mgmt exempt-sc-pct 10
      ```
2. The internal threshold table is a bit conservative. The default RxMER offset is 0 dB with a value in quarter dB steps. Increasing that to 12 is suggested. 12/4 = 3 dB correction factor.
    - o
      ```
      cBR8(config)#cable downstream ofdm-prof-mgmt mer-margin-qdb 12
      ```

One can also statically map a CM to particular data profile, if so desired.

- ```
  cBR8(config)#cable down ofdm-flow-to-profile profile-data <1-5> mac-addr <>
  ```

D3.1 CMs can lock to four profiles plus the control profile and store this in NVRAM. The CMTS and CM can support more, but the CM will need a dynamic bonding change (DBC) to support the fifth and could cause dropped packets. With this knowledge, it may be best to:

3. Use 256-QAM for the control profile and 4K-, 2K-, and 1K-QAM profiles along with one mixed profile (1K-QAM with 256-QAM) for potential roll-off issues.
    - o **Note**: CMs will only use these profiles if needed and will automatically go in and out of upgrade/downgrade.

Figure 7 is a sample output from a D3.1 CM showing current subcarrier RxMER readings. It shows the percent of subcarriers that would be able to run a given modulation. **Note**: This is not based on the D3.1 threshold table.
  ```

```
CM> /cm_hal/ofdm_analyzer 32 0
--------|---------|------------------------------------------------------
  RxMER |   Max   |          RxMER Histogram for 2749 Subcarriers
   dB   | Bitload |
--------|---------|--------10--------20--------30--------40-----50-----60-
  30 dB |  1K-QAM |
  31 dB |         |*
  32 dB |  2K-QAM |*
  33 dB |         |*
  34 dB |         |*
  35 dB |  4K-QAM |*
  36 dB |         |*
  37 dB |         |***
  38 dB |  8K-QAM |******
  39 dB |         |************
  40 dB |         |****************************
  41 dB | 16K-QAM |************************************
  42 dB |         |********
  43 dB |         |*
```

Figure 7 - Bit Loading Information from D3.1 CM

5. DS Resiliency / Partial Mode

The D3.1 OFDM PLC uses 16-QAM and is very robust, but it’s still a best practice to find “clean” spectrum to avoid potential issues. If the OFDM channel is used as the primary channel, losing the PLC would be catastrophic. Cross-bonding with SC-QAM signals may be in our best interest to provide higher peak speeds and have a means for the CM to enter partial mode.

When a D3.1 CM reports via a cm-status message a non-primary (secondary) RF channel impairment for SC-QAM or OFDM, the following things happen on a Cisco CMTS:

- The CM is marked `p-online` for easy identification.
- If the RF channel impairment is below the configured DS resiliency thresholds, the D3.1 CM’s service flows are moved to a resilient bonding group (RBG) or a narrowband (NB) interface (primary DS).
- If the RF channel impairment exceeds the configured DS resiliency thresholds, the impaired RF channel is temporarily de-activated from all the bonding groups (BGs).

The CM can move in and out of partial mode automatically. This feature is independent of the graceful profile management feature.

6. Capacity

Another tool in the toolbox is the use of Powerboost[tm] and the peak-rate TLV. Comcast trademarked the name, but it is a simple manipulation of the DOCSIS standard for CM rate limiting. This can be exploited to satisfy speed tests without over-provisioning by the typical 10%.

Over-provisioning is typically done to satisfy speed tests and the inherent difference in reporting at Layer 2 of the Open System Interconnection (OSI) model vs Layer 3. The CM and CMTS are reporting at Layer 2, but speed tests are reporting at Layer 3. The Layer 2 Ethernet overhead of 18 bytes per frame is not being counted, leading to a misconception on the end-users part of lower speed.

Figure 8 displays the amount of peak speed achieved while only over-provisioning this CM by 2%. The customer was sold a 500 Mbps service. The CM file was configured with 510 Mbps max rate, 600 Mbps peak rate, and 70 MB DS max burst.

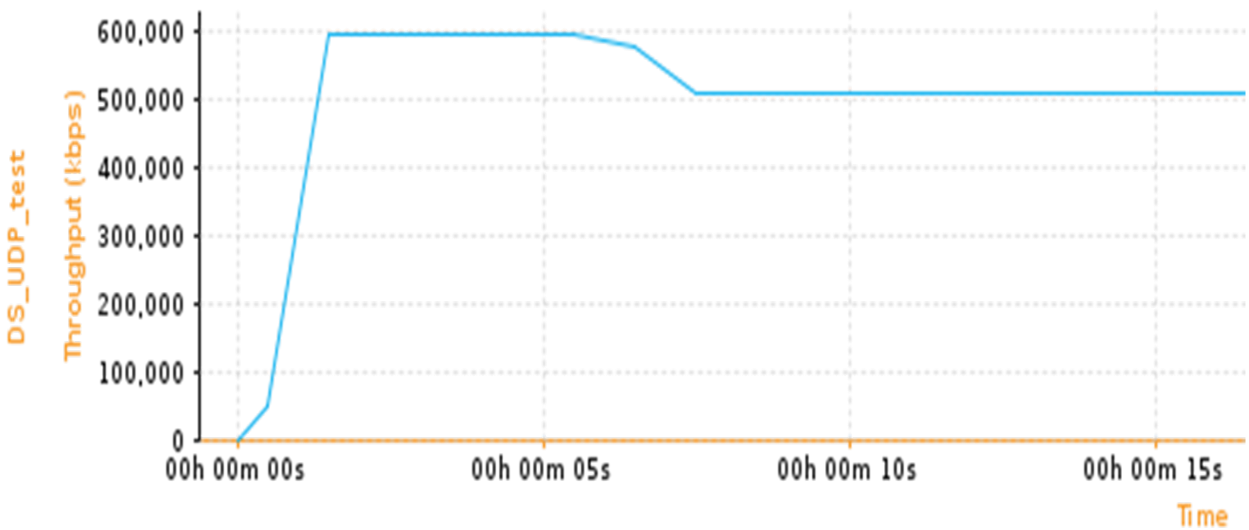

Figure 8 – Powerboost Example

This achieved approximately 6 seconds of 600 Mbps peak speed. The DS max burst and peak rate in the CM file could be manipulated even further to achieve longer times and/or peak rates.

7. OFDM Settings to Maximize Speeds

To maximize D3.1 speeds, the following settings are suggested (example values are for a Cisco CMTS):

- `cyclic-prefix 192` - Make this value as low as the HFC plant will support. 192 is lowest while 1024 is the default.
- `pilot-scaling 48` - Keep at the lowest setting, which is the default of 48.
- `roll-off 128` - Make as large as possible but must be less than the cyclic prefix value.
- `subcarrier-spacing 25KHZ` - Less overhead for 25 kHz vs 50 kHz.
- `profile-control modulation-default 256-QAM` - Configure 4k, 2k, & 1k-QAM data profiles with maybe one mixed profile.
- `profile-ncp modulation-default 64-QAM` - Make the next codeword pointer (NCP) as high as the plant supports.
- `guardband-override 1000000` - OFDM guardband override in 50 kHz increments.

The cyclic prefix and roll-off settings are likely the parameters that will need to be adjusted based on individual HFC plant characteristic as well as the OFDM channel width. As mentioned previously, the roll-off is required to be less than the cyclic prefix per the D3.1 PHY spec. A large roll-off value is desirable as this decreases the guardbands (taper regions) providing more spectrum for active subcarriers. However, to have a larger roll-off, a larger cyclic prefix setting is needed which leads to additional overhead on every subcarrier. It is not possible to optimize both of these settings.

In general, for smaller OFDM channel widths, more speed will be achieved by maximizing the number of subcarriers with a higher roll-off setting even though it creates more overhead per subcarrier with a higher cyclic prefix.

The Cisco CMTS also has a D3.1 DS guardband override feature. Remember, the actual DS is 204.8 MHz and the CMTS will null out a bunch (according to spec) to make 192 MHz, then there is a taper region for maybe 188-190 MHz of actual usable bandwidth. It was decided that the channel width (actual spectrum) would not exceed 192 MHz and allow the guardband to be overridden from the default settings which are based on cyclic prefix and guardband settings. Some settings can make the guardband 1.5 to 2 MHz on each end. If you have no adjacent carriers or even if another OFDM block adjacent, you may get away with less guardband without interfering with each other. The guardband must be symmetrical on the left and right.

You can override the default behavior where the guardbands are based on the roll-off setting. The larger the roll-off setting, the smaller the guardbands needed and therefore more of the OFDM channel can be used for transmitting data. The roll-off must be less than the cyclic prefix per the PHY specifications. A lower cyclic prefix is desirable as it reduces overhead and allows the OFDM channel to run faster. Your HFC plant needs to have minimal micro-reflections to support this lower cyclic prefix. Using a low cyclic prefix means you have to use a lower roll-off which will create larger guardbands. You end up with two variables that can’t both be optimized. The override feature allows you to use lower cyclic prefix and roll-off but then set the guardbands manually.

The whole reason Cisco defaulted the guardband size based on roll-off is because that is recommended in the PHY specifications as to prevent adjacent channel interference. The values were picked based on testing. If you start overriding these values, you can cause adjacent channel interference. It’s not recommended to change the guardband setting from the default unless the customer understands the consequences.

The maximum an active OFDM channel can be is 190 MHz so you will be limited to 1 MHz guardbands if configuring a full 192 MHz channel. If you configure a smaller channel size, you can use no guardband – but again look out for adjacent channel interference.

These settings will achieve the highest speeds, but some settings may need to be adjusted for real plant settings. Table 3 shows the speeds that can be achieved with various settings.

Table 3 - OFDM DS Speed Estimates (25 kHz)

Channels	Spectrum	DOCSIS 3.0	DOCSIS 3.1 (25 kHz subcarrier)		
		256 QAM	1024 QAM	2048 QAM	4096 QAM
4 channel	24 MHz	151 Mbps	172 Mbps*	189 Mbps*	206 Mbps*
8 channel	48 MHz	302 Mbps	373 Mbps*	410 Mbps*	448 Mbps*
16 channel	96 MHz	603 Mbps	776 Mbps*	853 Mbps*	931 Mbps*
24 channel	144 MHz	905 Mbps	1178 Mbps*	1296 Mbps*	1414 Mbps*
32 channel	192 MHz	1206 Mbps	1584 Mbps**	1742 Mbps**	1910 Mbps**
	2x192 MHz		3168 Mbps**	3484 Mbps**	3802 Mbps**

*** 25 kHz subcarriers**, running same modulation, 1.175 MHz guardbands, roll-off 192, cyclic prefix 256, 2 x NCP (64-QAM)

**** 25 kHz subcarriers**, running same modulation, 1.725 MHz guardbands, roll-off 128, cyclic prefix 192, 2 x NCP (64-QAM). The red line in table 3 indicates the channel width inflexion point where more speed is achieved via manipulation of roll-off and cyclic prefix. The guardband could be reduced to 1 MHz as well for even more speed.

8. General Thoughts for D3.1 Upstream (US)

I would be remiss to not mention D3.1 US at least in brief. D3.1 US uses OFDMA. Even though this feature is not deployed in production networks to date, lab testing is on-going and some observations have been made.

Most cable plants are limited to 5 MHz to 42 MHz with some upgrades happening to expand to 85 MHz. Because of the limited spectrum and the need to still provide services to legacy CMs, D3.1 US is not being heavily pursued. Until higher US speeds > 50 Mbps are required, US expansion to 85 MHz or even 204 MHz may be stalled. Part of the D3.1 spec allows time and spectrum sharing between SC-QAM US and OFDMA, which could allow higher peak speeds for D3.1 modems at the expense of more overhead.

For the interim, it's this author's belief that this time sharing also known as time and frequency-division multiple access (TaFDMA) may not be utilized at first and cross-bonding of four advanced time-division multiple access (ATDMA) with an OFDMA channel in an 85 MHz plant would suffice. Even with only one SC-QAM signal bonded with the OFDMA signal provides benefits as listed here:

- The CM would still have T4 multiplier of 2 allowing RF tech maintenance of up to 30*2 = 1 minute without losing DS lock from a T4 timeout.
- The CM would be more resilient in that it would have at least two channels doing station maintenance (SM).
- D3.1 US power is based on a formula for spectrum used and not just how many chs are in the US bonding group. So, max Tx doesn't change much when adding a small SC-QAM.
- Having a SC-QAM signal with its own scheduling may be simpler for VoIP and other scheduled flows that are latency and jitter sensitive.

One major observation made when deploying D3.1 CMs was how the US Tx level is being reported. The US max Tx level is 65 dBmV and translates to 53 dBmV for same bandwidth as four 6.4 MHz ATDMA chs. This value is 51 dBmV/ch for D3.0 without the Cablelabs' engineering change notice (ECN) for extended power. The D3.1 CM reports its US Tx level based on 1.6 MHz of bandwidth, which leads to a 6 dB correction factor compared to a 6.4 ATDMA channel.

Note: When replacing a D3.0 CM with a D3.1 CM (even in D3.0 mode) in the same location, you may notice the US Tx level reports 6 dB lower! This is normal and expected.

More data and testing needs to be gathered for US resiliency and how modems are assigned different interval usage code (IUCs). This means modems could use different modulation depending on the CMTS upstream RxMER for the respective subcarriers. Depending on the number of IUCs supported, this may negate the usage of lower spectrum with lower modulation so as to save those IUCs for modems with poor performance and nothing to do with lower spectrum.

Conclusion / Summary

DOCSIS 3.1 is being deployed today with great success; even better than expected. With features such as graceful profile management and resiliency, deploying higher modulation schemes has become more realistic. With the introduction of remote PHY architectures, expected performance increase in RxMER, and potentially less amplifier cascades, even higher modulation schemes will be possible for the DS and

US. When even higher speeds are required for the upstream, operators have a choice to change the diplex filter split to 204MHz/254 MHz and stay with D3.1 CMs or upgrade to FDX with compatible CMs. The FDX spec is under development at Cablelabs as of this writing.

In regards to on-going monitoring of plant performance, one can utilize the CM full bandwidth capture (FBC) for DS "sweepless sweep" and ingress testing and verification. One could activate CMTS RF and use that for test signals. Other aspects of PNM can also be utilized such as US spectrum viewing at the CM and CMTS along with US pre-EQ information. In the case of remote PHY, that information will be gathered at the remote PHY device (RPD). One could also use test equipment with a built-in CM to balance amplifiers by looking at CM Tx levels as they work their way downstream from the node/RPD. This may be necessary until US sweep is supported with a distributed access architecture (DAA) system.

Everyone should be aware that correctable forward error correction (FEC) errors with D3.1 will show very high, if not 100%! This is expected when using OFDM with LDPC and such large block sizes. Even though there are potentially thousands of subcarriers, they are all processed with the FFT.

Note: There is no special CM file needed for D3.1, but the CM must be in multiple transit channel (MTC) mode also referred to as US bonding. Sometimes US level issues could force the CM to non US bonding and the CM will not register properly and may not even downgrade to D3.0 or 2.0 mode. Always keep in mind that performance and features supported can vary dramatically with CM firmware.

Fiber deep architectures and remote PHY will allow much higher speeds via the higher order modulation schemes supported with D3.1. Node plus 0 could also facilitate an easier transition to FDX. Coax attenuation is our biggest hurdle at higher frequencies, so limiting the coax to eventually drop cable only (think fiber-to-the-tap) is a proposition that still works in our favor.

By utilizing CMTS features for robustness and "self-healing", we can successfully operationalize these more complex architectures, multiplexing technologies and modulation schemes. Some of these features include: utilization load balance (2.0 & 3.0), US & DS resiliency, dynamic modulation, graceful profile management, etc.

Future SDN of OFDM profile management may not be as critical as first thought, but it will give us even more granularity when needed or justified. Be sure to utilize PNM to be more proactive in the monitoring of your customers' quality of experience (QoE).

Abbreviations

4G	4^{th} generation mobile technology (LTE)
AGC	automatic gain control
ALC	automatic level control
ATDMA	advanced time-division multiple access
avg	average
B	bytes
b	bits
BG	bonding group
bps	bits per second
ch	channel

CM	cable modem
CMTS	cable modem termination system
coax	coaxial cable
corr	correctable
D1.x	DOCSIS 1.0 & 1.1
D2.0	DOCSIS 2.0
D3.0	DOCSIS 3.0
D3.1	DOCSIS 3.1
DAA	distributed access architecture
dB	decibel
DBC	dynamic bonding change
dBmV	decibel millvolt
DOCSIS	Data-Over-Cable Service Interface Specifications
DRFI	Downstream Radio Frequency Interface
DS	downstream
ECN	engineering change notice
FBC	full bandwidth capture
FDX	full duplex DOCSIS
FEC	forward error correction
FFT	fast Fourier transform
freq	frequency
FTTH	fiber-to-the-home
Gbps	gigabits per second
HFC	hybrid fiber-coax
Hz	hertz
IP	internet protocol
IUC	interval usage code
LDPC	low density parity check
LLQ	low latency queue
LTE	long-term evolution (4G)
Mbps	megabits per second
MER	modulation error ratio
MHz	megahertz
MoCA	Multimedia over Coax Alliance
MTC	multiple transit channel = US bonding
NB	narrowband
NCP	next codeword pointer
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
OSI	Open Systems Interconnection
PLC	physical layer link channel
PNM	proactive network maintenance
QAM	quadrature amplitude modulation
QoE	quality of experience
RBG	resilient bonding group
RF	radio frequency

RPD	remote PHY device
RTP	Research Triangle Park, real-time transport protocol
Rx	receive, receiver
SC-QAM	single-carrier QAM
SCTE	Society of Cable Telecommunications Engineers
SDN	software defined network
SG	service group
T4	timer 4 = 30 second CM timer for DS lock
TaFDMA	time and frequency division multiple access
TLV	type length value
Tx	transmit, transmitter
US	upstream
VoIP	voice over IP

Bibliography & References

Google Hangout - #1 DOCSIS Podcast (The only one)

http://volpefirm.com/

https://plus.google.com/u/0/+Volpefirm/videos