Is "Unity Gain" Still the #1 Objective?

Maybe YES!

A Technical Paper prepared for SCTE•ISBE by

John Ulm
Engineering Fellow, Broadband Systems
CommScope – CTO Network Solutions team
Moultonborough, NH 03254
+1 (978) 609-6028
john.ulm@commscope.com

Dr. Zoran Maricevic
Engineering Fellow
CommScope – CTO Network Solutions team
15 Sterling Drive Wallingford, CT 06117
+1 (203) 303-6547
zoran.maricevic@commscope.com

Dr. Frank O'Keeffe
Distinguished Systems Engineer
CommScope – CTO Network Solutions team
2 Johnstown Close, Glounthaune, Cork, Ireland T45V094
(+353) 87 668 8789
frank.okeeffe@commscope.com

Table of Contents

List of Figures

List of Tables

Title **Page Number**

1. Introduction

As of April 2020, all the 1,000+ pages of the Data Over Cable Service Interface Specification (DOCSIS) 4.0 PHY and MULPI specs, combined, are out and about. That's step one in making the 1,794 MHz top frequency in the downstream, combined with up to 684 MHz for the top frequency in the upstream, a new reality for the hybrid fiber-coax (HFC) plants. As part of the greater 10G initiative, this is what will be needed to support the many Gbps in both the upstream and the downstream for nodes with amplifier cascades (a.k.a. N+X). The multiple system operators' (MSOs') successful business model and existing HFC broadband networks are our starting point. All that's needed is a simple upgrade to get from here to there, easy peasy!

Not so fast?! Well, OK, there are some "minor details" to sort in the process of getting there: Where is the Goldilocks zone – not too little but not too much either – for the RF amps gain and power? There are many items to consider.

With a comprehensive network model of a node + 6 cascade, should the "unity gain" concept be extended all the way up to 1.8 GHz? Once modeled, what would an optimal power spectral density (PSD) distribution across the forward spectrum look like? What happens as the distance between the amps increases? Does the tap position on a link make a difference? What is the impact of the drop cables? Should we continue to use end-of-line (EoL) throughput as our capacity benchmark?

There are many, many questions and this paper will start to look at some of the considerations with rolling out 1794 MHz HFC plants. All of this is distilled into conclusions and best practices guidance – to help make 10G and the DOCSIS 4.0 networks a reality for the time to come!

2. Unity Gain Overview

Figure 1 is a big picture view of a Hybrid Fiber Coaxial (HFC) network. It shows a "40,000-foot view" of an HFC plant. While the plant and its architecture kept evolving – for example, a head end to node fiber link took the place of a much longer cascade of trunk amps – it is also amazing how many things remain the same: fiber nodes feeding RF amplifiers, amplifiers feeding taps, and taps, via "drops" delivering signals to ever-important subscribers.

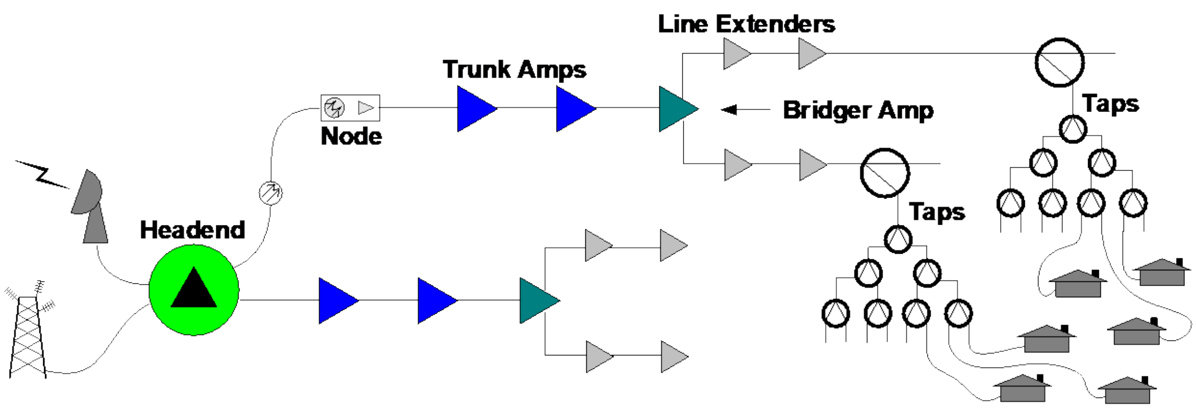

Figure 1 – A big picture view of a Hybrid Fiber Coaxial (HFC) network

Another HFC aspect that has remained the same is that of "unity gain" – a concept of setting every amplifier output signal to the same shape and magnitude, no matter if the amp is first or last in the cascade nor how "lossy" the span in front of the said amplifier was. This consistency of output is shown via blue "amplifier output" lines in Figure 2 for the downstream direction.

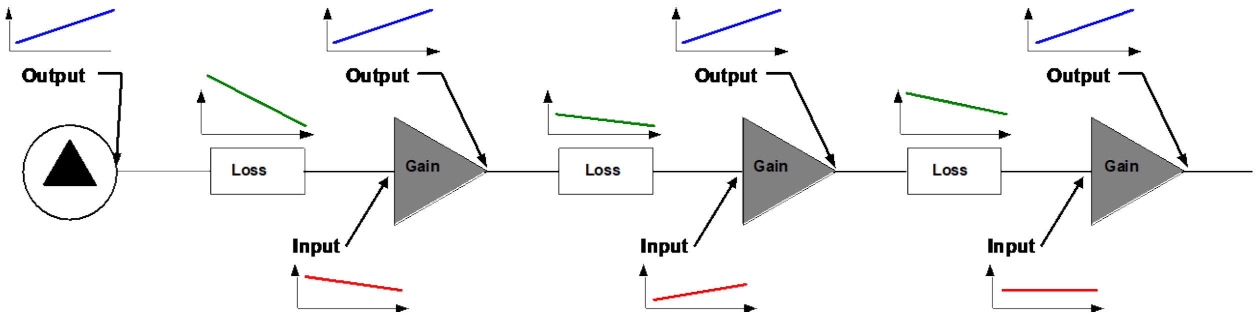

Figure 2 – Illustration of the "Unity Gain" concept, applied to an RF amplifier cascade

In this example of a three-amplifier stage cascade, each output stage (in blue) is the same. However, the coaxial cable and tap loss preceding each of the amplifiers is of a different value, as illustrated via green "loss over frequency" lines. These various length cable spans, driven by the same previous-amplifier output, will produce various inputs, shown via red "amplifier input" lines. It is the gain and slope of every amplifier that is tuned, traditionally via selection of proper attenuator, cable simulator, and/or cable equalizer, to make every amplifier output the same. One goal of the tuning is to set each amplifier's input stage signal as flat as possible to minimize noise figure contribution, the other is to up-tilt the final output, to minimize non-linear distortion effects. Nevertheless, the noise effects and the non-linear distortions still add up as the cascade length increases, despite keeping outputs of each amplifier the same.

Another rationale for unity gain is to provide a common input power and modulation order (e.g. 256-QAM) for every cable consumer premises equipment (CPE) including legacy video set top boxes (STB) as well as DOCSIS 2.0 &/or 3.0 cable modems (CM). This holds no matter what channel frequency is being used or whether the CPE device is next to the node or at the end of line (EoL).

The span length between the amps, and how much of the corresponding amplifier gain is needed to achieve "unity gain" has been the cable engineers' focus for many years. Many HFC plants were originally designed with amp spacing to accommodate lower frequencies, e.g. 450/550/750 MHz. Over time, these were then "pushed" to higher frequencies, e.g. 870/1002/1218 MHz. Due to much higher cable losses at higher frequencies, this resulted in amplifiers needing significantly higher output gains to maintain the unity gain in the system. If the cable losses were more than the increased amplifier output gains, then the HFC plant needed to be "re-spaced" where amplifier locations were moved, something that is highly undesirable. This issue now gets exaggerated with pushing the HFC to 1.8 GHz.

As a point of reference, consider 750 MHz HFC plant that many "first world" countries built quite a few years ago. At 750 MHz, these spans amount to ~20 dB attenuation following trunk amps, ~40 dB attenuation following multi-port bridgers and ~30 dB attenuation following single output line extenders.

3. Network Capacity Planning for Cable 10G

The network capacity analysis in [ULM_2019] provides an insight into the capacity requirements needed to support the cable 10G initiative. It first looks at the traffic engineering needed for a common 10G network using both PON and cable systems. Then a closer look is taken at the spectrum planning for an HFC system.

The CommScope (formerly ARRIS) team has been providing industry leading research in traffic engineering for many years which was recently highlighted in [ULM_2019]. Some additional references of note include [CLO_2014], [EMM_2014], [ULM_2014], [CLO_2016], [ULM_2016], [ULM_2017] and [CLO_2017].

3.1. The "Basic" Traffic Engineering Formula

Previously, [CLO_2014] introduced traffic engineering and quality of experience (QoE) for broadband networks. From there, the paper went on to develop a relatively simple traffic engineering formula for service groups that is easy to understand and useful for demonstrating basic network capacity components.

The "Basic" formula shown below is a simple two-term equation. The first term (Nsub*Tavg) allocates bandwidth capacity to ensure that the aggregate average bandwidth generated by the Nsub subscribers can be adequately carried by the service group's bandwidth capacity. The first term is viewed as the "DC component" of traffic that tends to exist as a continuous flow of traffic during the peak busy period.

The "2014" Traffic Engineering Formula (Based on Tmax_max):

$$C \geq (Nsub*Tavg) + (K*Tmax\_max) \qquad (1)$$

where:

C is the required bandwidth capacity for the service group
Nsub is the total number of subscribers within the service group
Tavg is the average bandwidth consumed by a subscriber during the busy hour
K is the QoE constant (larger values of K yield higher QoE levels)…
where $0 \leq K \leq$ infinity, but typically $1.0 \leq K \leq 1.2$
Tmax_max is the highest Service Tier (i.e. Tmax) offered by the MSO

There are obviously fluctuations that will occur (i.e. the "AC component" of traffic) which can force the instantaneous traffic levels to both fall below and rise above the DC traffic level. The second term (K*Tmax_max) is added to increase the probability that all subscribers, including those with the highest service tiers (i.e. Tmax values), will experience good QoE levels for most of the fluctuations that go above the DC traffic level.

The second term in the formula (K*Tmax_max) has an adjustable parameter defined by the K value. This parameter allows the MSO to increase the K value and add bandwidth capacity headroom that helps provide better QoE to their subscribers within a service group. In addition, the entire second term is scaled to be proportional to the Tmax_max value, which is the maximum Tmax value that is being offered to subscribers.

In previous papers [CLOONAN_2013, EMM_2014], found that a K value of ~1.0 would yield acceptable and adequate QoE results. [CLOONAN_2014] goes on to provide simulation results that showed a value between K=1.0 and 1.2 would provide good QoE results for a service group of 250 subscribers. Larger service groups (SGs) would need even larger values of K while very small SGs might use a K value near or less than 1.0.

3.2. Broadband Subscriber Traffic Consumption

CommScope/ARRIS has been monitoring subscriber usage for over a decade now from the same group of MSOs. The data from this set has been compared and aligns closely to many other MSOs globally.

Figure 3 shows the average subscriber downstream consumption, DS Tavg, during peak busy hours for several MSOs over a ten-year period. At the start of 2020, DS Tavg had surpassed the 2 Mbps barrier.

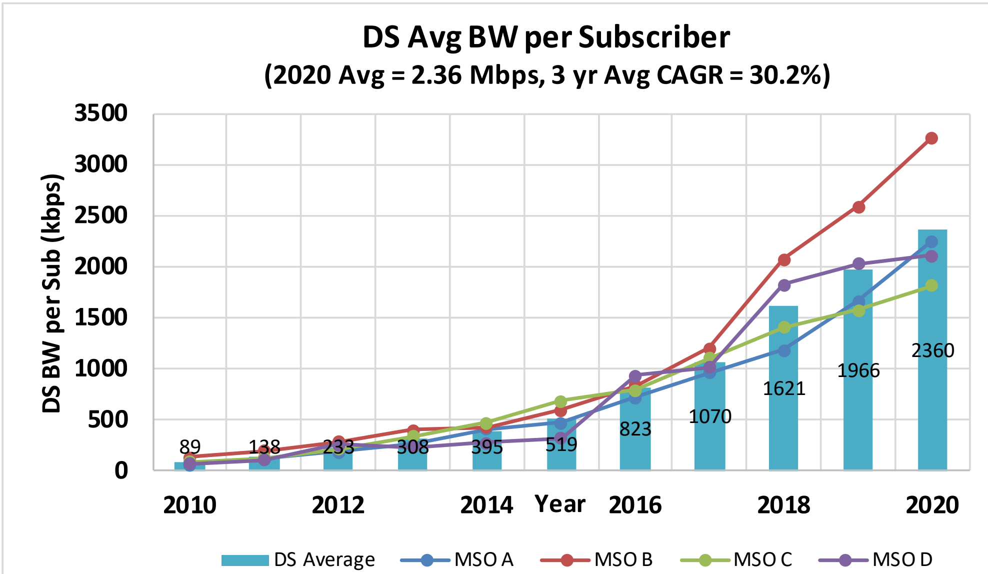

Figure 3 – Downstream Average Bandwidth per Subscriber through Jan '20

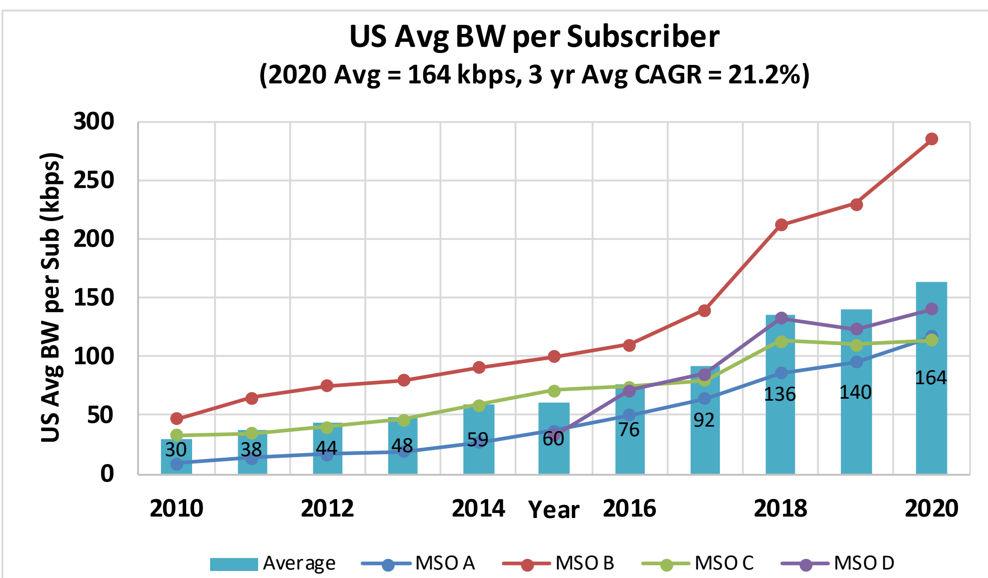

Figure 4 – Upstream Average Bandwidth per Subscriber through Jan '20

It turns out that the Tavg growth rate was higher at the start of this decade and has tailed off a bit in recent years. Over the last 3-4 years, this group of MSOs had an average downstream traffic growth that had been around 30%. Interestingly, the upstream traffic is growing at a significantly slower rate than the downstream as shown in Figure 4. During the same ten-year period, the upstream Tavg generally grew at less than 20% compound annual growth rate (CAGR).

This MSO data provides a good indication of Tavg, at least before the Coronavirus bandwidth (BW) surge hit. Note that figures 3 and 4 are very generalized results that are averaged across millions of subscribers.

Over recent years, there has been a slowing in the downstream usage growth rate (i.e. Tavg) compared to the service tier growth rate (i.e. Tmax). This has several consequences including that the networks become more "bursty". It also means that the overall utilization of the network is lower too. In this respect, it is important to try and maximize subscribers per service group (SG) in order to take advantage of statistical multiplexing and get better economics.

3.3. Network Capacity Modeling for Cable 10G Downstream

[ULM_2019] used the CommScope network capacity modeling tools to see how a 1218/204 MHz HFC plant could support the cable 10G downstream requirements. This analysis will first be reviewed and then look at what changes for the extended spectrum DOCSIS 4.0 (ESD) scenarios.

The network capacity results from [ULM_2019] show the potential capabilities for a 1218/204 MHz HFC plant. It begins with a 512-home passed (HP) service group with 256 subs (i.e. 50% penetration). The max downstream (DS) service tier starts at 1 Gbps and grows by 1 Gbps per year from 2022 until it finally reaches 8 Gbps DS SLA in the year 2028. For this case study, it is assumed the Tavg growth rate continues its gradual decline over the next decade. This will leave Tavg at ~15 Mbps by the end of the decade. If the Tavg growth rate does not decline, then these dates might get pulled in by 2-3 years.

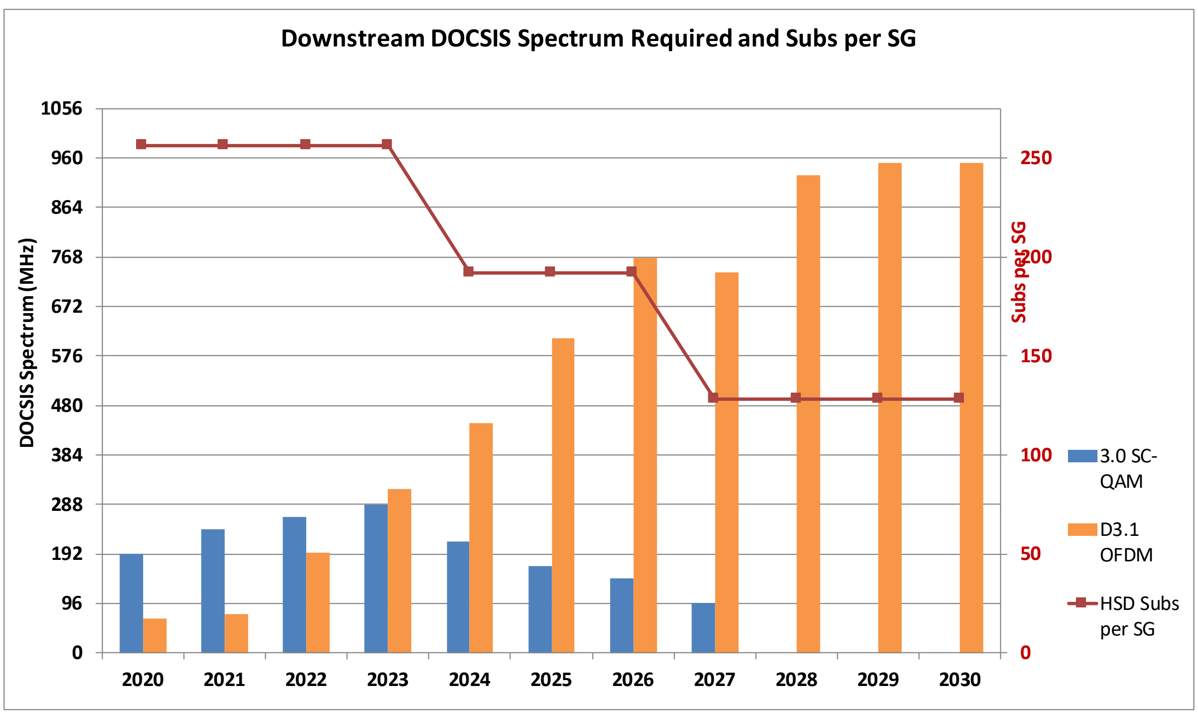

Figure 5 – 1218/204 MHz System – Subs per SG, DOCSIS Spectrum Needs

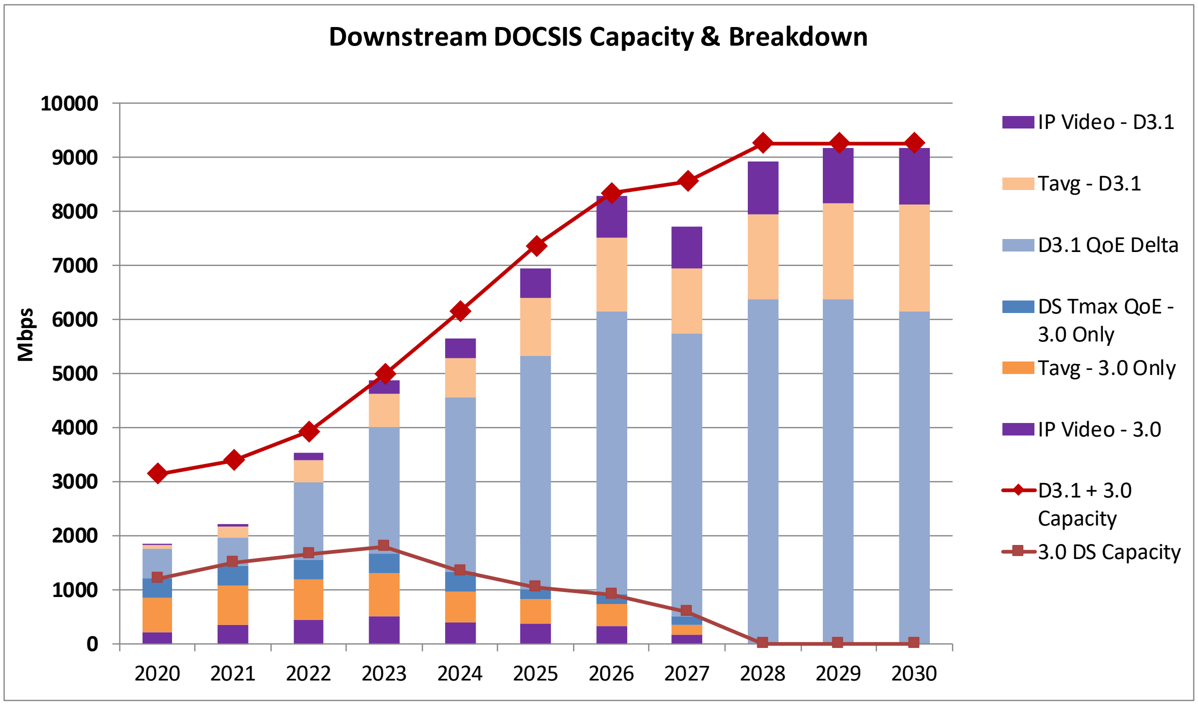

Figure 6 – 1218/204 MHz System – DOCSIS Usage: Tmax, Tavg, IP Video

This 1218/204 MHz scenario shows that by 2027, all SGs need to be at 128 subs or fewer. The max subs per SG is shown in Figure 5. The figure also breaks out the downstream spectrum amount needed for both DOCSIS 3.0 SC-QAM and DOCSIS 3.1 OFDM. Note that by 2028 in this case study, 100% of DOCSIS cable modems have been converted to DOCSIS 3.1 enabling OFDM channels across the entire spectrum to maximize capacity.

The DOCSIS capacity usage is broken out in Figure 6. It shows the amount of capacity needed for both DOCSIS 3.0 and DOCSIS 3.1. As can be seen, the Tmax component dominates over time. The upper red line shows the combined 3.0 + 3.1 total capacity for the system.

Perhaps a key point of this case study was that a reasonably clean 512 HP node can be upgraded to 1218/204 MHz and support a service tier of 8 Gbps x 1.5 Gbps for the next decade. The only change needed will be a SG segmentation (i.e. upgrade node from 1x1 to 2x2) somewhere in the middle of the decade. There is no pressing near term need to push the HFC to very small (but inefficient!) SG sizes found in N+0 systems.

3.4. Network Capacity Planning for DOCSIS 4.0 ESD

The cable 10G initiative strives for a more symmetric network with multi-gigabit upstream service tiers. DOCSIS 4.0 ESD has evolved as a way to offer even higher upstream tiers while remaining in an N+X environment. Thus, ESD also is trying to avoid the costs of an N+0 upgrade such as that needed with DOCSIS 4.0 FDX.

3.4.1. Cable 10G Service Tiers

As shown in [ULM_2019], the subscriber is actually getting an 8 Gbps SLA in a 10G world after accounting for all the different overheads (e.g. PHY, MAC, IP layers). Table 1 summarizes the various options from that study and their respective downstream (DS) and upstream (US) SLAs that service providers can consider offering. Because capacity in an HFC system can vary quite a bit based on many variables, the offered SLAs are actually a range of values.

Table 1 – Summary of SLA Options for 10G PON & 1218 MHz Plants

10G PON Options	DS SLA (Gbps)	US SLA (Gbps)
10G/1G EPON	8	0.8
10G/10G EPON	8	8
XG-PON	8	2
XGS-PON, NG-PON2 (single wavelength)	8	8
10G HFC Options	**DS SLA (Gbps)**	**US SLA (Gbps)**
1218/85 MHz	8 – 10	0.4 – 0.5
1218/204 MHz	6 – 8	1.0 – 1.5
1218/300 MHz	5 – 7	2.0 – 2.5
1218/396 MHz	4 – 6	2.5 – 3.0
1218/85 MHz + 108-684 MHz FDX/Soft-FDX	8 – 10	5 – 6

Perhaps the key motivating driver for DOCSIS 4.0 FDX and ESD is the ability to offer more symmetric multi-gigabit upstream tiers. The minimum goal for ESD would be to at least match the DS tiers for a 1218/204 MHz plant while offering substantially higher upstream tiers. Ideally, ESD would be able to match the DS tiers capabilities of 10G PON and 4.0 FDX systems.

The table above shows upstream splits up to 396 MHz. The DOCSIS 4.0 working group has also added a 492 MHz split and 684 MHz split as additional options.

3.4.2. DOCSIS 4.0 ESD – Tavg and Nsub targets

The [ULM_2019] study assumed that the Tavg growth rates continue to decline and estimated that Tavg would reach 15 Mbps per sub by 2030. Given the additional investments needed for ESD, a Tavg of 20-40 Mbps per sub will be considered to make sure there is additional lifetime for this investment.

Since ESD is targeted at N+X environments, our analysis will assume at least 125 subscribers per SG, which was the ending point in the 2019 study.

3.4.3. DOCSIS 4.0 ESD – Network Capacity Requirements

The 2019 study showed that ~9 Gbps of DS capacity is required for 125 subs @ Tavg=15 Mbps. If Tavg is bumped up to 25 Mbps, then total system capacity pushes above 10 Gbps. Pushing Tavg up to 40 Mbps increases the required system capacity to around 12 Gbps. These are key targets that will be used in our following analysis.

4. Simulation Model Overview

Our network performance estimates were done using a MATLAB simulation model. The model simulated the effects of attenuation in the hardline coaxial cable, taps, drop cable and home entry port, and amplification in the intermediate amplifiers, to calculate the received power levels at modems at different locations. The effects of noise and distortion accumulation in the cascade were also simulated to calculate the effective SNR at the modem. The results were compared with DOCSIS 4.0 received power and SNR thresholds to estimate the available modulation and bit-loading at each frequency, yielding the total gross throughput. Physical layer overheads such as cyclic prefix, FEC (LDPC), guard bands, pilot and PLC tones were subtracted from the integrated bit-loading curve (gross throughput) to calculate the net system throughput. The model used 78% efficiency for the PHY layer.

4.1. Cable Model – PSD and TCP

The DS power spectral density (PSD) at the node output was represented as a tilt/drop/tilt (a.k.a. lightning bolt) in order to achieve a sensible total composite power (TCP) (nominally 70 dBmV based on 4.0 working group discussions, but this parameter could be varied if needed) while providing adequate levels for legacy SC-QAM services at lower frequencies (up to 860 MHz). The PSD is shown in Figure 7.

Cable attenuation, for both hardline and drop cables, was calculated as a function of frequency and cable segment length from tabulated data from cable manufacturers. A large database of tap response curves from several tap families, taken from a combination of manufacturers' data and our own measurements, was available in the model. Note that no tap equalizers (or "cable simulators") were present in the model; all available RF power was used to optimize the bit-loading from each tap.

4.2. Amplifier Model

The amplifiers were modelled as non-ideal near-unity gain devices. The gain of each amplifier was tuned to approximately match the attenuation of the preceding cable span, subject to a configurable maximum gain. This approximate match used a linear-tilted gain over most of the spectrum, producing a small gain deviation; over longer cascades, this produced a slightly U-shaped end-of-line PSD, in contrast to a flat received PSD which would occur if perfect unity-gain amplification was possible.

The gain limit became significant for longer cable runs that is referred to as stretched plant. At higher frequencies, the inter-amplifier span attenuation sometimes exceeded the assumed maximum gain, resulting in progressively lower PSD at these frequencies after each stage in the cascade. Hence the stretched plant did not quite achieve unity gain at 1800 MHz. Figure 8 shows an example of the inter-amplifier attenuation and amplifier gain for a stretched plant.

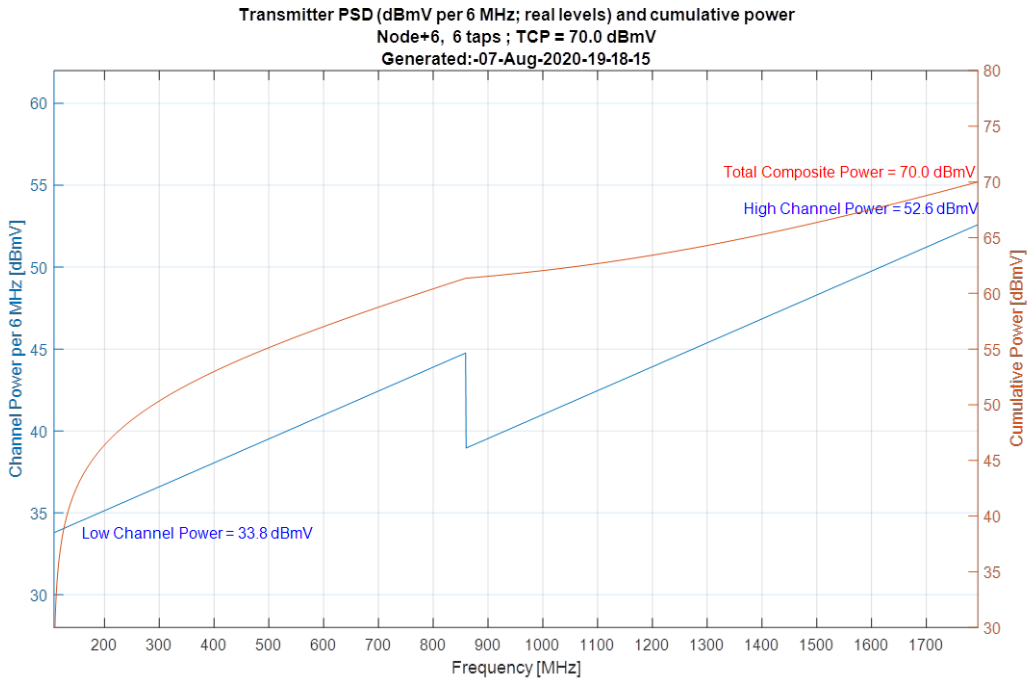

Figure 7 – PSD Profile Showing Tilt/Drop/Tilt Characteristic

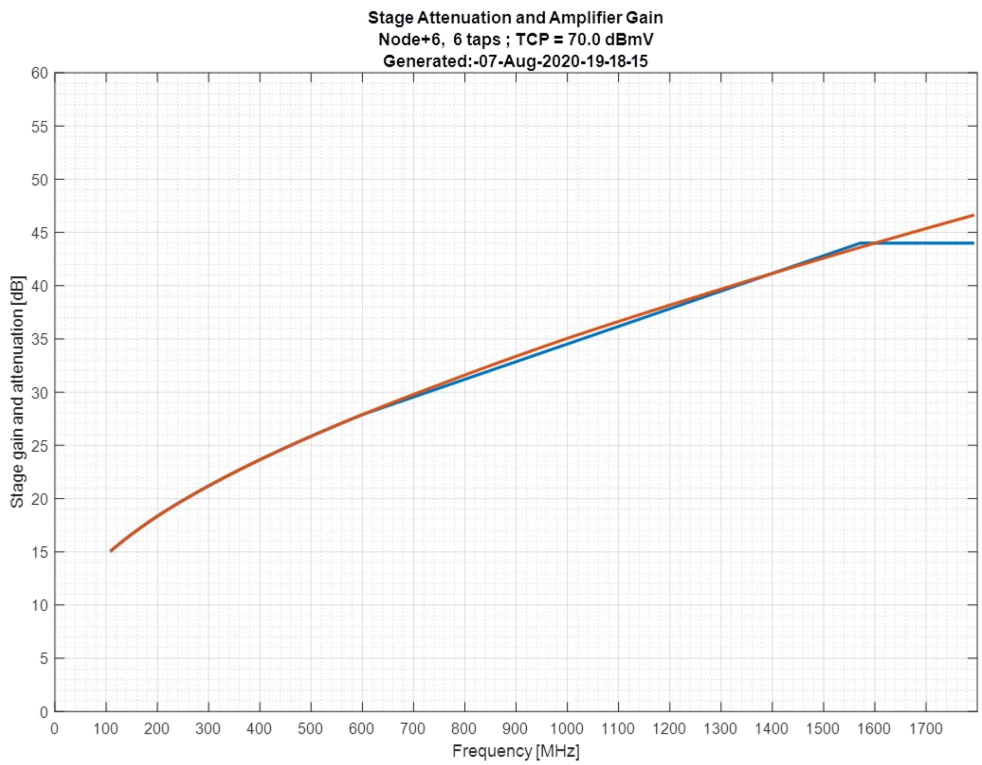

Figure 8 – Inter-amplifier Span Attenuation & Amp Gain (With 44 dB Gain limit) in a Stretched Plant

4.3. System Noise Performance

The noise performance of the system was modelled using an estimate of the node's output SNR, and noise figures for the amplifiers in the cascade, and for the cable modem. The various noise components were modelled through the successive attenuation and amplification stages of the cascade to produce a received noise estimate. Distortion was treated in a similar way; distortion at the node output and each amplifier output was modelled as a third-order function of TCP, and coherence effects were also calculated.

4.4. Amplifier Spacing – Typical and Stretch Networks

For our network simulations, it was important to model amplifier spacings that were representative of real-world data based on customer's amplifier spacing data. A cumulative distribution function (CDF) of some sample amplifier spacing for single output line extenders (LE) is shown in Figure 9, while the CDF for multi-port bridgers (MB) is shown in Figure 10. Line extenders often account for more than 60% the total number of amplifiers. This data is courtesy of Cox Communications.

Most of the simulations used either "typical plant" or "stretch plant" parameters for plant length. "Typical plant" parameters are chosen to include a majority of inter-amplifier span attenuations (for line-extender (LE) spacings). A total attenuation of 30 dB at 1 GHz and 40 dB at 1794 MHz is used, which is an attenuation value greater than that of ~65% of line-extender to line-extender spacings according to Figure 9. This attenuation is simulated with 6 segments of 175' of P3 625 hard-line and 5 taps, a total of 1050' between amplifiers. Note that 850' of P3 500 would give similar results.

Unless otherwise specified, each simulation used the same parameters for the node to first amplifier span and for the last amplifier to end tap as for the inter-amplifier spans.

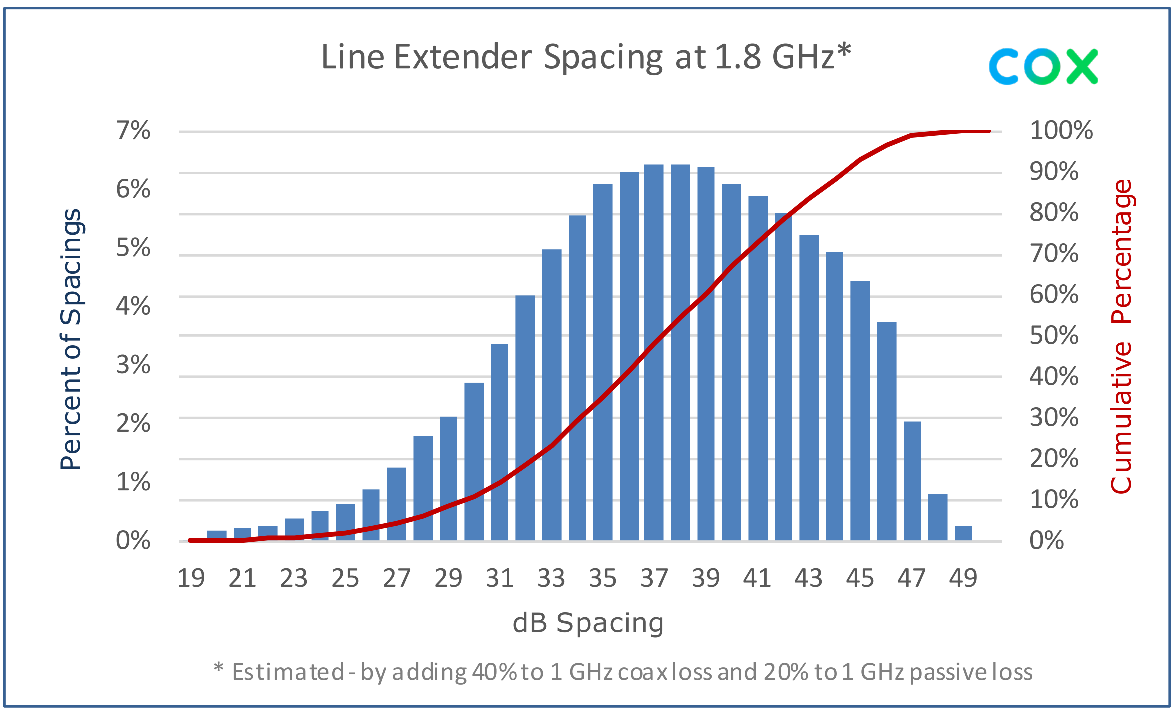

Figure 9 – Amplifier Spacing – Single Output Line Extenders (LE)

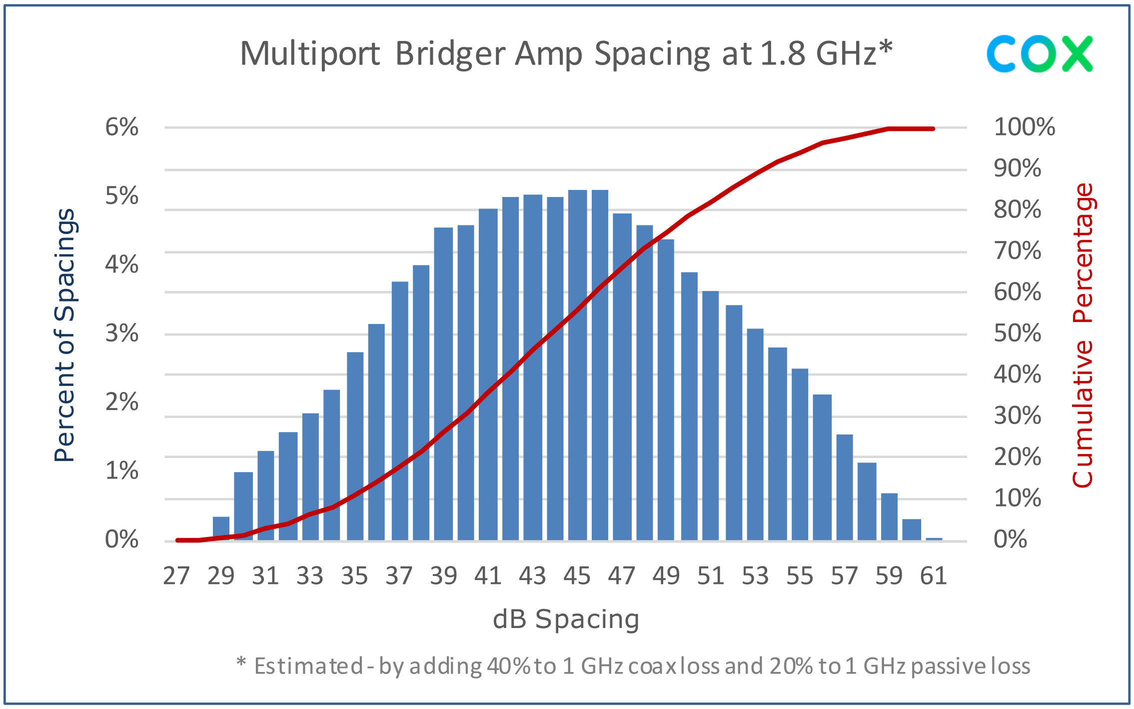

Figure 10 – Amplifier Spacing – Multi-Output Bridgers (MB)

"Stretch plant" parameters are chosen to exceed most inter-amplifier span attenuations; the corresponding attenuations are 35 dB at 1 GHz and 47 dB at 1794 MHz. This is greater than ~97% of LE-to-LE spacings, and two-thirds of bridger-to-amplifier spacings according to Figure 10. Overall, this would cover more than 85% of the amplifier links. This attenuation is simulated with 6 segments of 215' of P3 625 hard-line and 5 taps, a total of 1290' between amplifiers. Note that 1050' of P3 500 would give similar results.

For both typical and stretch plants, the baseline case used 150' of RG-6 drop cable along with a 3.5 dB loss inside the home (e.g. splitter at point of entry or 20' of RG-6 inside the home). Note that cable losses from 150' of RG-6 drop cable would be roughly equivalent to 210' of RG-11 drop cable.

5. Conventional Wisdom Using Unity Gain Approach

As discussed previously, unity gain helped maintain a constant QAM-modulation for legacy STB and DOCSIS 2.0/3.0 cable modems. System performance was defined by the lowest common denominator. The end of line (EoL) provided a key performance monitoring spot to determine whether every CPE device could successfully receive the 256-QAM downstream channel. But future systems will be migrating to DOCSIS 3.1 and 4.0 cable modems. Our network simulations calculate the maximum downstream network capacity achievable using 3.1/4.0 OFDM channels with variable-bit loading per sub-carrier with up to 4096-QAM modulation.

For our network simulations, the EoL performance is first checked for both 1218 and 1794 MHz systems to see how well unity gain performed. Figure 11 looks at the network capacity for a 108-1218 MHz downstream for various amplifier cascade lengths, from Node+0 up to Node+6. This was done for both a typical plant and a stretch plant conditions.

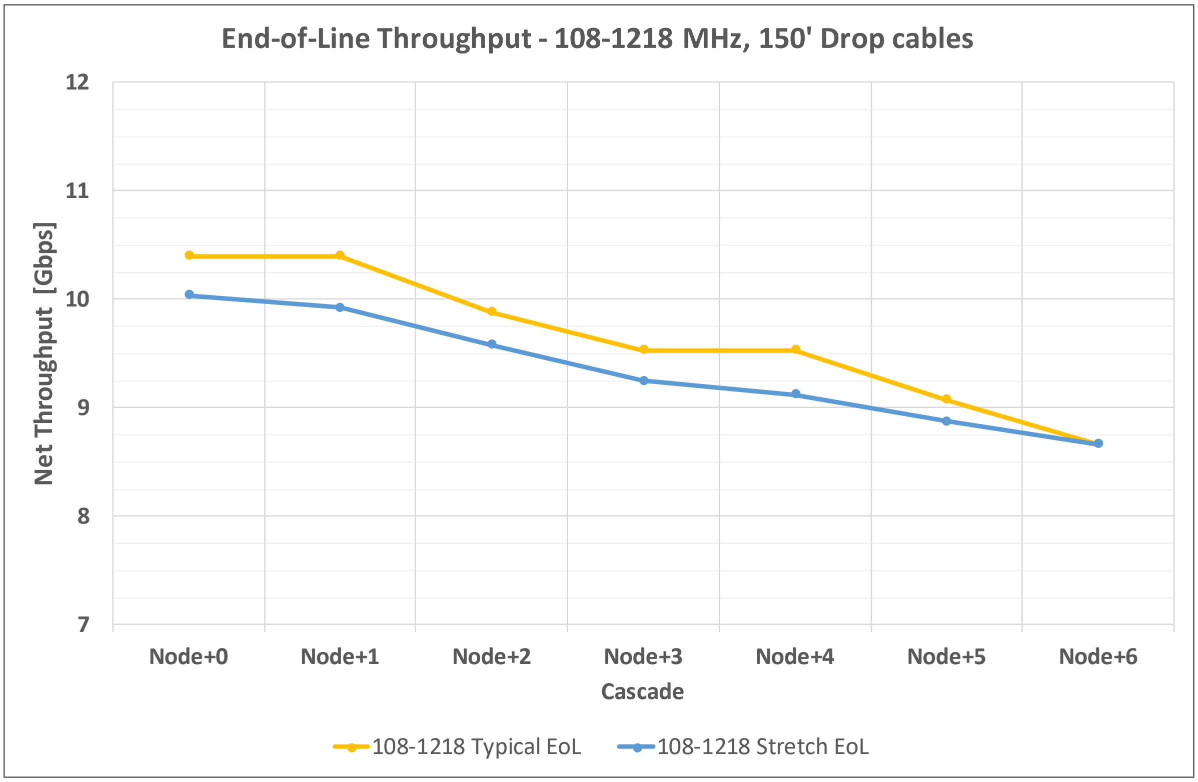

Figure 11 – End-of-Line Throughput - 108-1218 MHz, 150' Drop cables

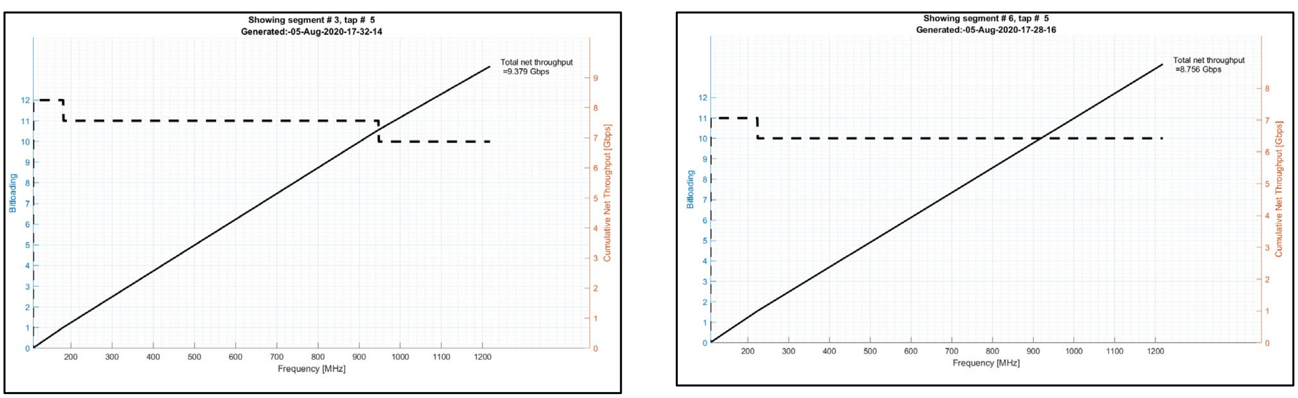

Figure 12 – Stretch 108-1218 MHz Plant Bit-loading – Tap 5 for N+2 and N+5

The 1218 MHz system maintains unity gain for both the typical and stretch plants. The network capacity degrades slightly as the amplifier cascades increases due to the accumulated noise and distortion from each additional amp stage. The additional cable loss in the stretch plant degrades network capacity slightly from the typical plant, but each amplifier stage is still able to maintain its unity gain.

Figure 12 shows the bit-loading for the 5th tap in a N+2 and N+5 stretch 108-1218 MHz system. As can be seen, the bit-loading remains fairly constant over the entire spectrum gain. This is what one might expect with unity gain.

Figure 13 now shows the EoL throughput for a 492-1794 MHz downstream system with cascade lengths varying from Node+0 to Node+6. This has 192 MHz more total downstream spectrum than the 108-1218 MHz system, but it is spectrum at higher frequencies with higher losses. Note that the downstream network capacity of the 492-1794 MHz typical plant in Figure 13 is very close to the 108-1218 MHz typical plant in Figure 12. However, the EoL throughput for the 1.8 GHz stretch plant in Figure 13 is substantial lower than the typical plant by around 20% or almost 2 Gbps. The stretch plant is also below the minimum 9 Gbps capacity target that was discussed in section 3.

Figure 14 shows the bit-loading for the 5th tap in a N+2 and N+5 typical 492-1794 MHz system. Since unity gain is maintained, there is no drop-off in bit-loading at higher frequencies. In fact, there is a slight increase in the bit-loading at the higher frequencies due to some non-linear effects in the amplifier.

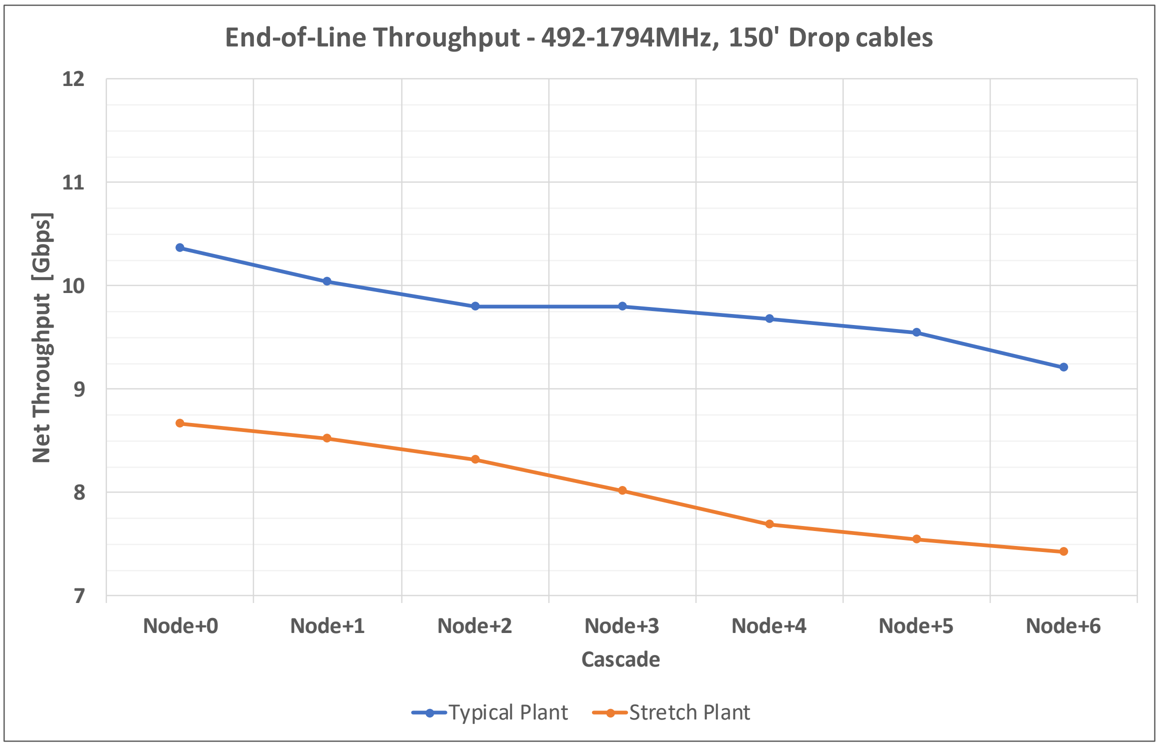

Figure 13 – End-of-Line Throughput - 492-1794 MHz, 150' Drop cables

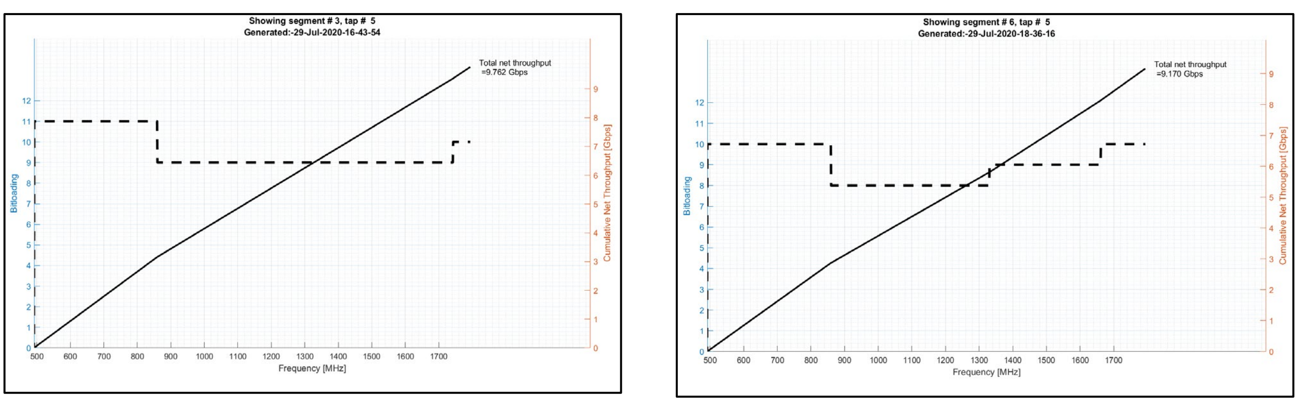

Figure 14 – Typical 492-1794 MHz Plant Bit-loading – Tap 5 for N+2 and N+5

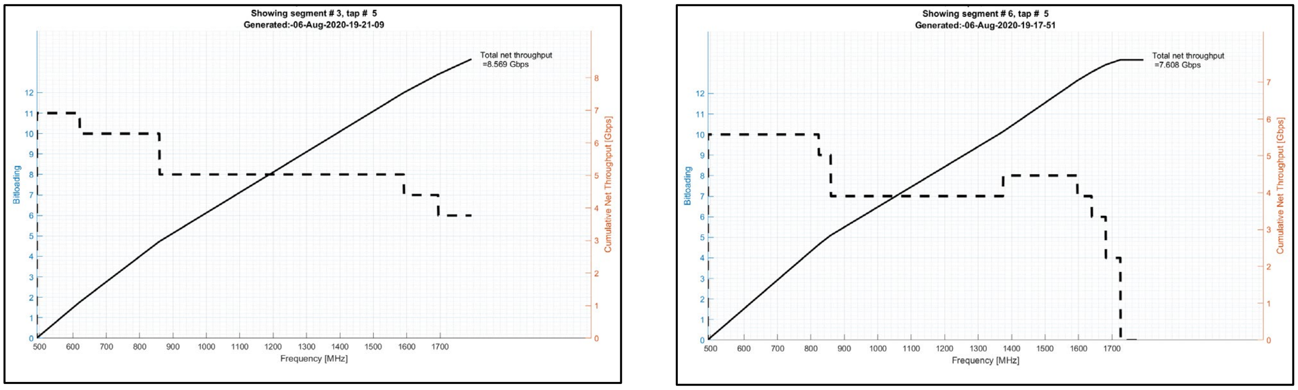

Figure 15 – Stretch 492-1794 MHz Plant Bit-loading – Tap 5 for N+2 and N+5

The stretch 492-1794 MHz system does not maintain unity gain. Figure 8 showed that the amplifier gain reaches its maximum output around 1600 MHz. Figure 15 shows the bit-loading for the 5th tap in a N+2 and N+5 stretch 492-1794 MHz system. It clearly shows how the bit-loading starts to drop-off above 1600 MHz. Note that the loss is relatively minimal after the 2nd amplifier output, only 1-2 orders of modulation. However, by the output of the 5th amplifier, the 5th tap is seeing a dramatic drop in bit-loading, losing roughly half the capacity in the 1602-1794 MHz OFDM channel.

Based on our earlier network capacity planning targets of needing 9-12 Gbps, the stretch 492-1794 MHz plant would not have acceptable system capacity using EoL throughput measure. Our conventional wisdom would dictate that these plants would need to have the amplifiers re-spaced or a mid-span amplifier added in order to maintain unity gain.

6. Game Changer – DOCSIS Compensates When Unity Gain Lags

When the DOCSIS 3.1 (D3.1) specification was written, the authors had a vision of a new DOCSIS PHY layer that squeezed every last once of capacity out of the cable plant. DOCSIS 4.0 inherits all these capabilities. Some of the most important D3.1 capabilities were wide OFDM channels with variable bit-loading along with the support of multiple profiles.

Previously, all modems would receive the identical downstream data stream. Therefore, the entire network had to operate for the lowest common denominator. Hence, the EoL capacity was a good metric for determining the plants capabilities. Each OFDM channel has a group of D3.1 profiles. Within that channel, modems are put into profile groups with similar modems and receive data at an optimal data rate. For example, modems in one profile could be receiving at 4096-QAM modulation while other modems in a different profile are receiving at 256-QAM modulation. This allows the DOCSIS system to optimize system capacity. This means that total system capacity depends on the capacity seen by every home on every tap in the system.

Figure 16 shows a 108-1218 MHz downstream example in a N+6 plant for both typical and stretch amplifier spacing. Network capacity is shown at every tap in the system. Modems further out on the cascade are seeing slightly lower performance and would be operating with D3.1 profiles that are using QAM modulations one or two steps below the best profile. In this example, the best profile would have 20% more capacity than the lowest profile.

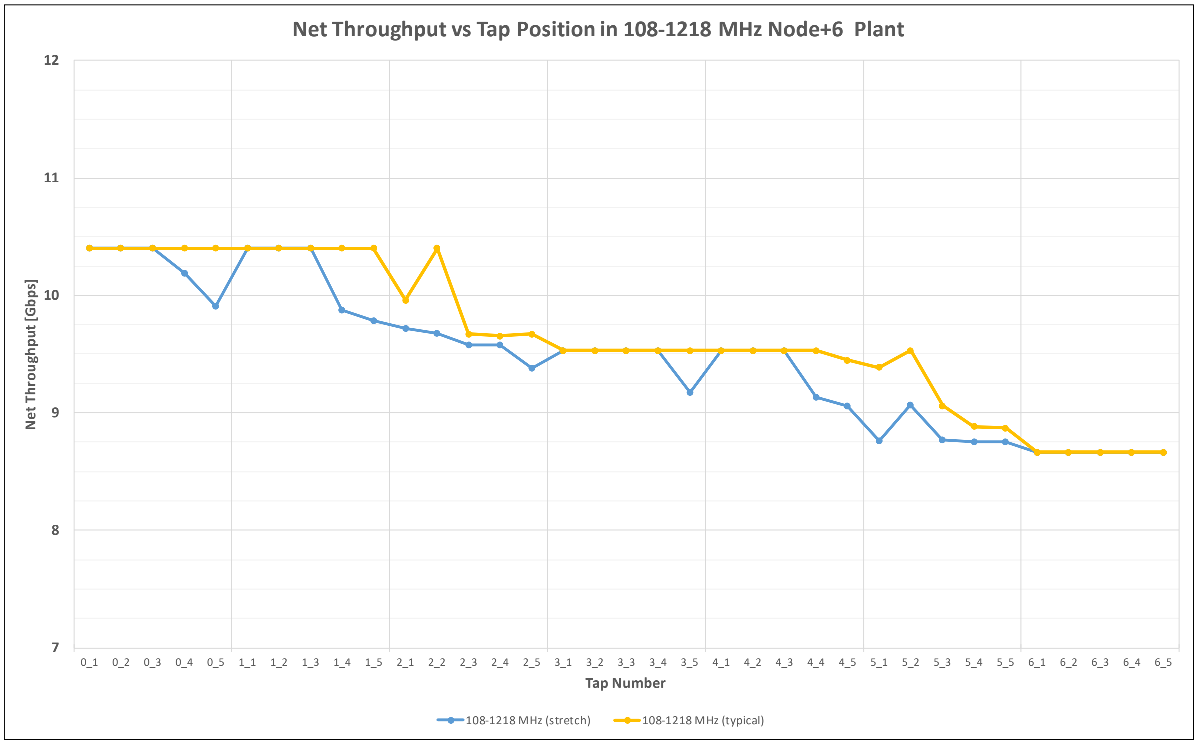

Figure 16 – 108-1218 MHz Net Throughput vs. Tap Position in N+6 Plant, 150' Drop cables

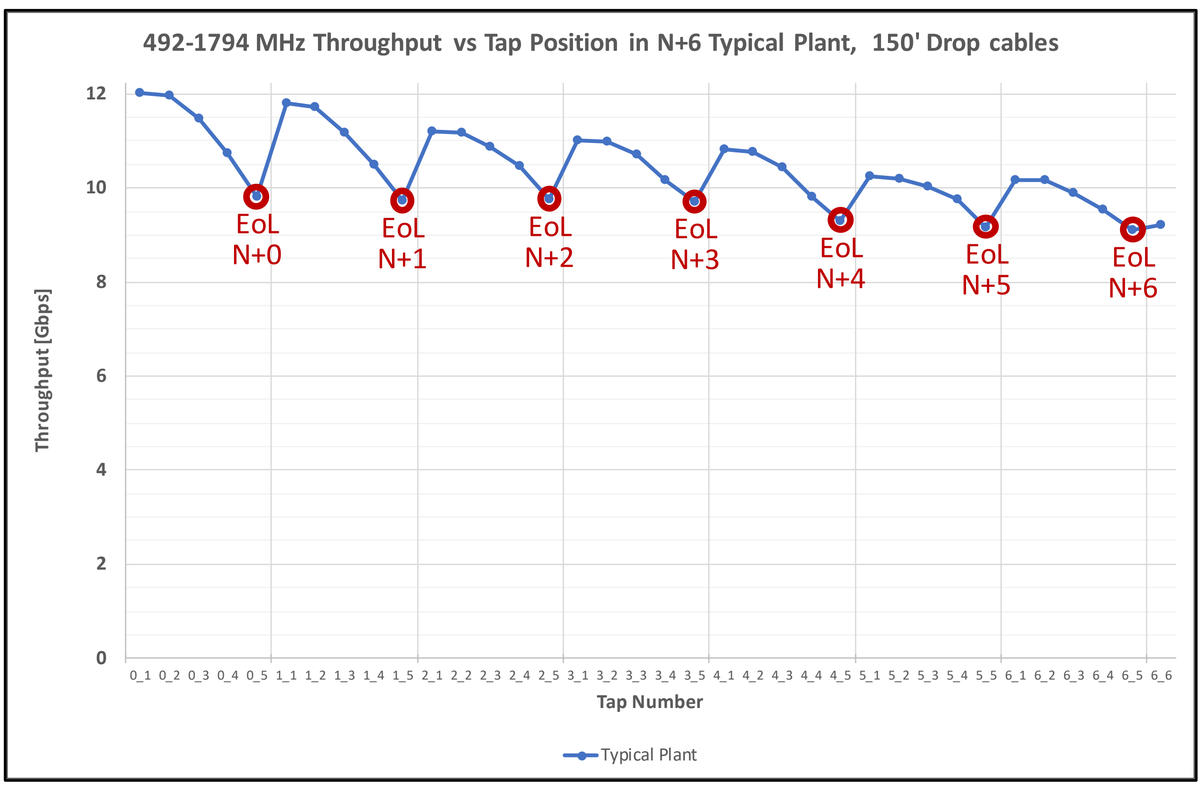

Figure 17 – 492-1794 MHz Throughput vs. Tap Position, N+6 Typical Plant, 150' Drops

For 1.8 GHz plants, the cable losses at the higher frequencies will magnify the differences between the best profiles and the lowest profiles. Figure 17 helps give a glimpse into the wider discrepancy seen by modems across an entire typical plant. Capacity changes based on two factors: # of amplifiers passed; and # of taps from the amplifier output. Figure 18 shows the stretch plant that has even more of a swing. In this example, the typical plant has a 32% swing from the lowest to highest capacities while the stretch plant sees a 60% swing from low to high! It goes from ~7.5 Gbps at tap 5 on the 6th amplifier up to almost 12 Gbps at the 1st tap from the node.

Remember, the EoL capacities discussed earlier basically aligns with the 5th tap value after the last amplifier, so it is basically following the troughs on this chart (see EoL for each N+X in Figures 17 & 18). These wide variations are just waiting to be optimized by DOCSIS. Many of the modems will have capacities significantly above the EoL capacities.

The profiles are set up separately for each DOCSIS OFDM channel. The OFDM channels are up to 192 MHz wide. So, the 492-1794 MHz downstream will have up to 7 OFDM channels with each channel getting its own optimized profiles.

In addition to optimizing each channel, the CMTS scheduler does load balancing between the channels. And an intelligent scheduler can factor the profiles for each modem to determine which channels it should use to optimize total network capacity.

So, if EoL is not a good measure of the network capacity, what should be used? It turns out that this is an extremely complex answer based on many different variables. The following sections will explore a number of these variables and make some estimates of network capacity for a ‘reasonable’ system.

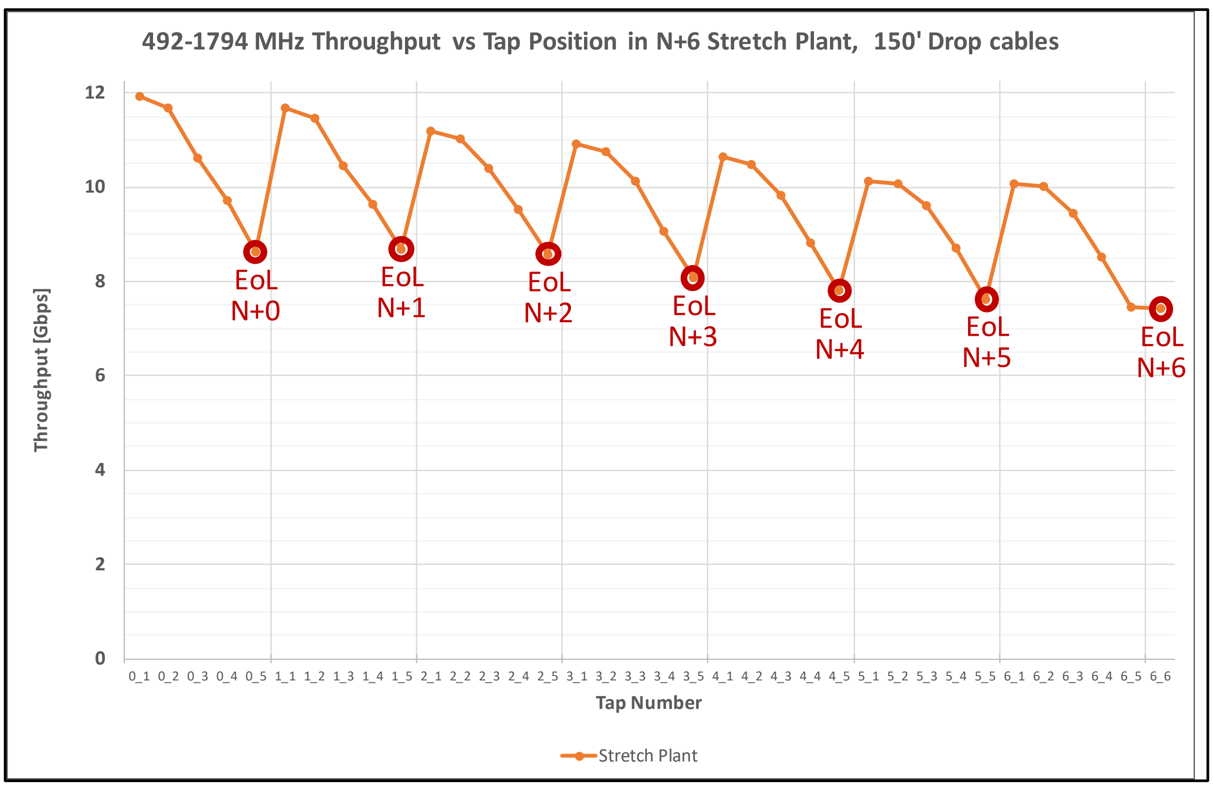

Figure 18 – 492-1794 MHz Throughput vs. Tap Position, N+6 Stretch Plant, 150' Drops

6.1. Estimating HP/Subscriber Distribution across N+X HFC Plant

As can be seen in Figures 17 & 18 above, the available capacity to any given modem is a function of where it is located on the plant relative to the fiber node. The HFC plant is a tree and branch topology that fans out for each additional level in the cascade. If one blindly assumed that each amplifier has 2 outputs feeding the next level of amplifiers, then a N+6 plant would have 64 amplifiers at the last stage in the cascade and a total of 126 amplifiers per node. If the outputs per amplifier is increased to 3 or 4, then it grows exponentially to where there could be 1000's of amplifiers from a given node.

But this is not reality. To understand what is out in the real world, the authors spoke with our customers and tapped into the knowledge of the in-house CommScope HFC design team. This in-house design team is perhaps the most experienced HFC design team anywhere, with a legacy going back 30+ years at some of our previous incarnations: ARRIS, C-Cor, Motorola, General Instrument, Philips.

The first thing we learned is that there is really no "typical" system. Real world HFC plants vary all over the map. So, we set out to define a "reasonable" scenario to get some baseline capacity estimates. From there, certain variables can then be changed to understand their impact on network capacity.

In general, amplifier spacing on average is relatively constant at 4-5 amplifiers per mile – especially if the networks considered include a mix of low density rural, medium density suburban and high-density urban areas. The homes passed (HP) density impacted how many homes might be off each tap (e.g. 1-2 HP per rural tap, 4-8 HP per urban tap). So, for our analysis, it is assumed that HP would be distributed evenly across all taps in the system.

As a starting point, a suburban build with 80-110 homes passed per mile (HP/mile) was chosen. From there, plants of various cascade lengths (i.e. N+0 to N+6) were analyzed. Table 2 shows an estimate of RF amps and homes passed based on input from our expert in-house HFC design team. This table shows that the number of RF amplifiers per node grows almost linearly with the cascade depth. It is NOT exponential! For each level in the cascade, another six to eight amplifiers are typically added.

Why is this??? It turns out there are several factors. First, the majority of amplifiers are single output line extenders (LE). These are amplifiers just cascading down a street without any branches. Next, multiple output bridgers (MB) would often have some outputs feeding side streets without connecting to another amplifier (i.e. a dead end). Finally, not all branches go out the entire cascade depth. In a N+6 plant for instance, many branches might terminate after 3, 4 or 5 amplifiers and not reach the sixth cascade level.

The homes passed per node also tended to follow the amplifier count with typically 16-20 HP per amp.

This knowledge then let us model how the homes passed (and hence subscribers) are spread across the HFC plant. After the first amplifier, a linear increase in HP per cascade level is used.

The next task was to understand how many taps are typically following each amplifier. As noted above, the amplifier spacing is fairly constant, and hence the number of taps were too. But there is some variation. The HFC design team might need to shorten or extend a leg based on the real-world geography. The vast majority of amplifiers have either 4 or 5 taps at its outputs. The number of fewer taps (e.g. 3) or more taps (e.g. 6) was trivial for this analysis. It turns out that percentage of amplifiers with 4 taps is roughly twice that of the amplifiers with 5 taps.

Table 2 – Example RF Amps and Homes Passed per Node for N+X

HFC Cascade Lengths	# of RF Amps per Node	# of Homes Passed per Node
Node + 0	0	40 - 60
Node + 1	4 – 6	80 - 150
Node + 2	10 – 12	150 - 240
Node + 3	15 – 20	240 - 360
Node + 4	22 – 28	360 - 480
Node + 5	28 – 35	480 - 600
Node + 6	35 - 45	600 - 720

6.2. DOCSIS Optimized Network Capacity Analysis

The first order of business was to look at the 108-1218 MHz plant shown in Figure 16 and see how much additional DOCSIS capacity might be available compared to the EoL measurements. Figure 19 shows the DOCSIS optimized weighted average capacity compared to the EoL capacity for both the typical and the stretch 1218 MHz plants.

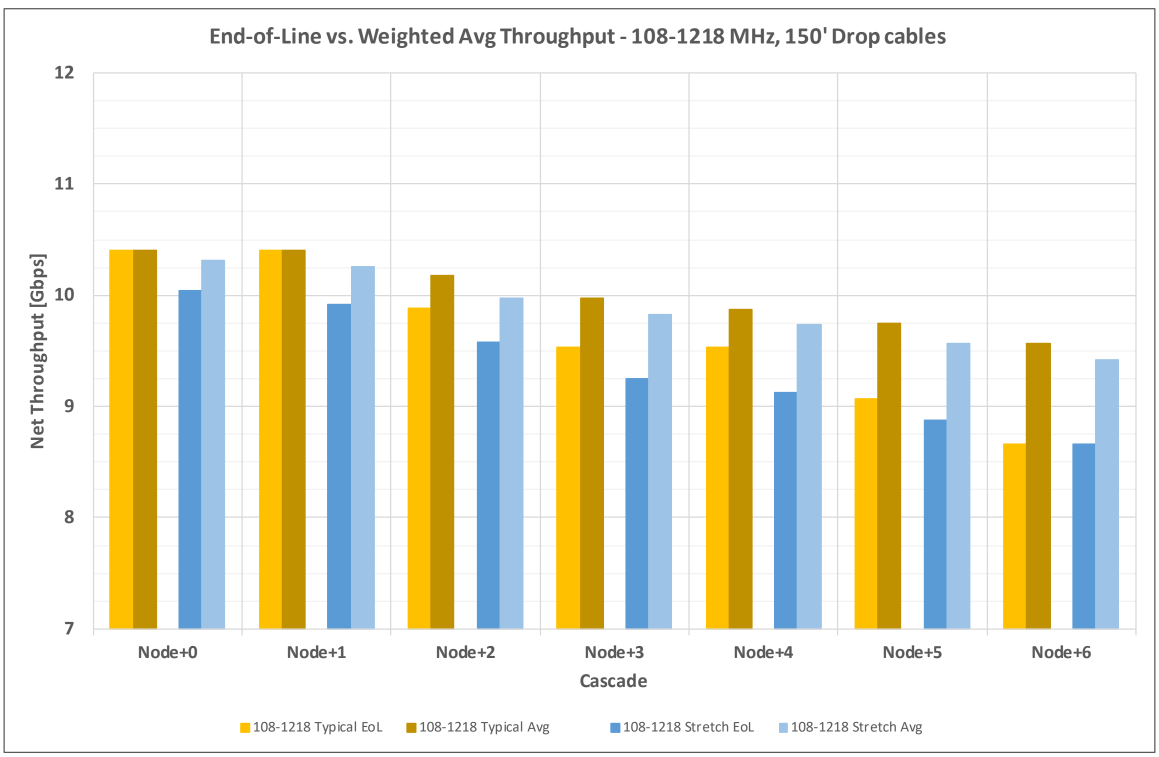

Figure 19 – End-of-Line vs. Weighted Avg Throughput - 108-1218 MHz, 150' Drop cables

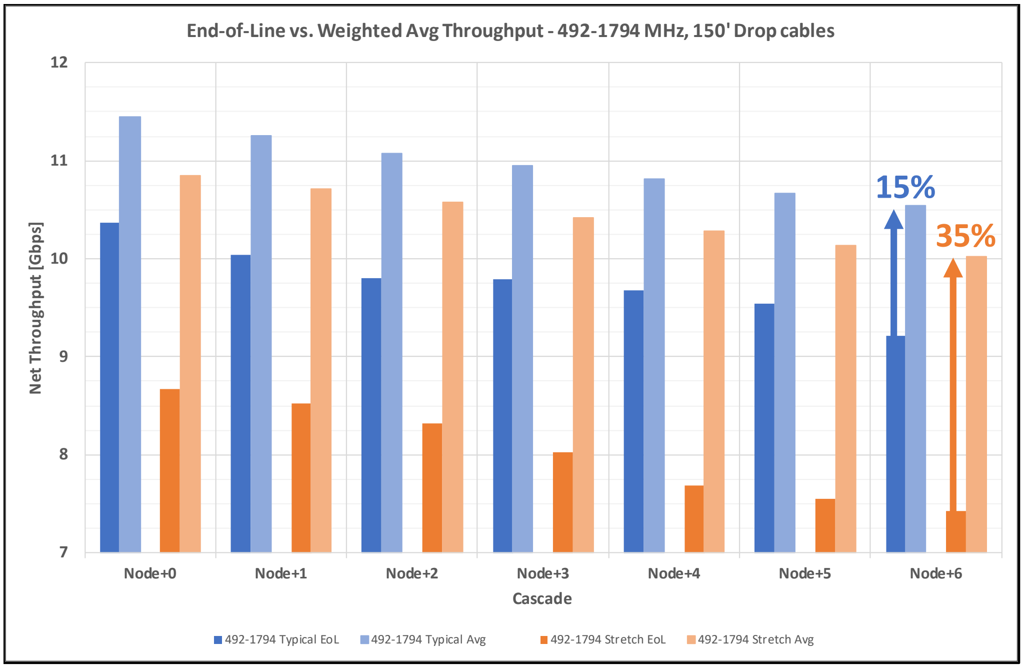

Figure 20 – End-of-Line vs. Weighted Avg Throughput - 492-1794 MHz, 150' Drop cables

Both typical and stretch plants saw very little gain for shorter cascades while N+6 plant saw up to 10% gains over the EoL calculations. So EoL appears to be reasonably accurate for the 1218 MHz plant.

The next scenario looked at the 492-1794 MHz plant previously shown in Figure 17 & 18 to see how much additional DOCSIS capacity might be available compared to the EoL measurements. Figure 20 shows the DOCSIS optimized weighted average capacity compared to the EoL capacity for both the typical and the stretch plants.

The difference between the DOCSIS optimized weighted average capacity and EoL capacity is significant. For the typical plant, the gains start around 10% for N+0 and increase to 15% for N+6. The stretch plant sees even larger gains, from 25% gains for N+0 plant to 35% for N+6 plant. The EoL measure is grossly underestimating the capacity of the stretch plant. The EoL measure showed that typical plants had 20% to 25% more capacity than the stretch plant. The DOCSIS optimized capacity results above shows that the gap between typical and stretch plant is significantly less, in the 5% range!

Perhaps one of the most significant conclusions from this study is that the stretch plant that was considered inadequate, is actually capable of supporting the network capacity requirements and does NOT need any extra help (e.g. re-spacing amps or adding mid-span amplifiers).

6.3. Impact of shorter drop cables (66’/100’/150’)

Drop cables are less costly than trunk lines, but also do not have as good RF characteristics. This becomes a big deal as the HFC plant pushes to 1.8 GHz. Per our in-house HFC design expert, most drop cables out there today are RG-6 cables. Higher performance RG-11 drop cables are typically used for extremely long runs, e.g. 200’ or longer.

Our expert felt that our 150’ drop cable assumption in the previous results was on the high side. Figure 21 and 22 show the impact on total network capacity by varying the drop cable length from 66’ to 100’ to

150'. This was done for the typical plant in Figure 21 and the stretch plant in Figure 22, using RG-6 cable for all cases.

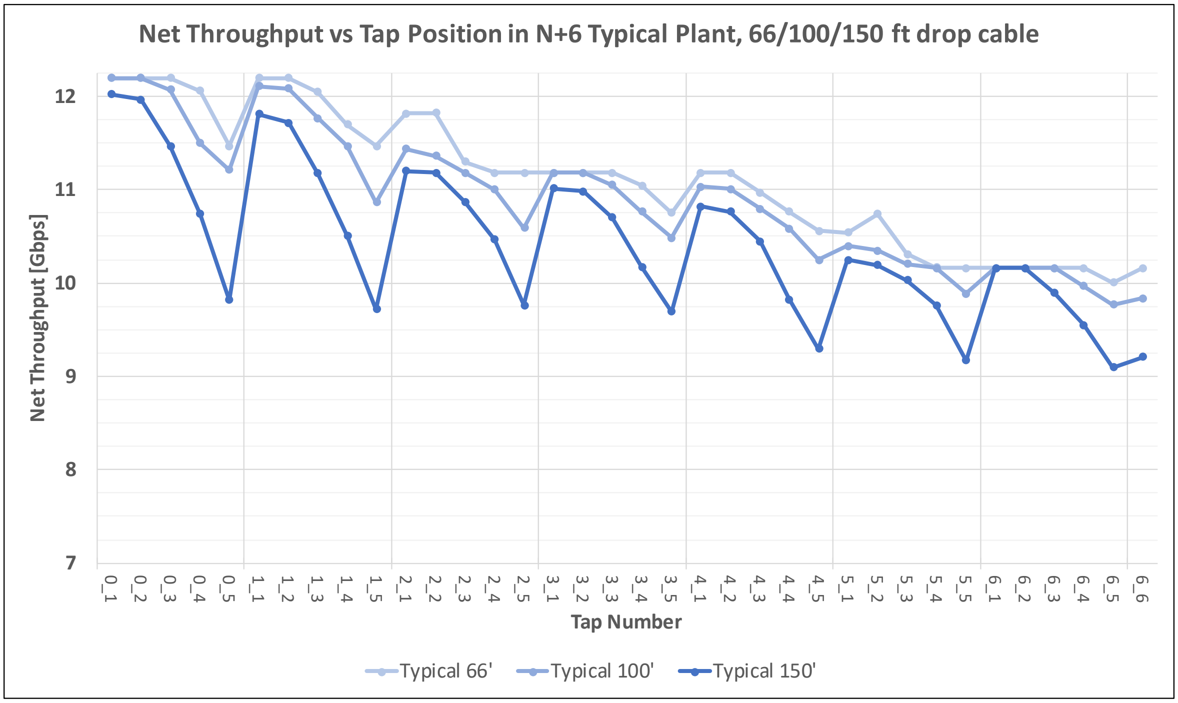

Figure 21 – Net Throughput vs Tap Position in N+6 Typical Plant, 66/100/150 ft drop cable

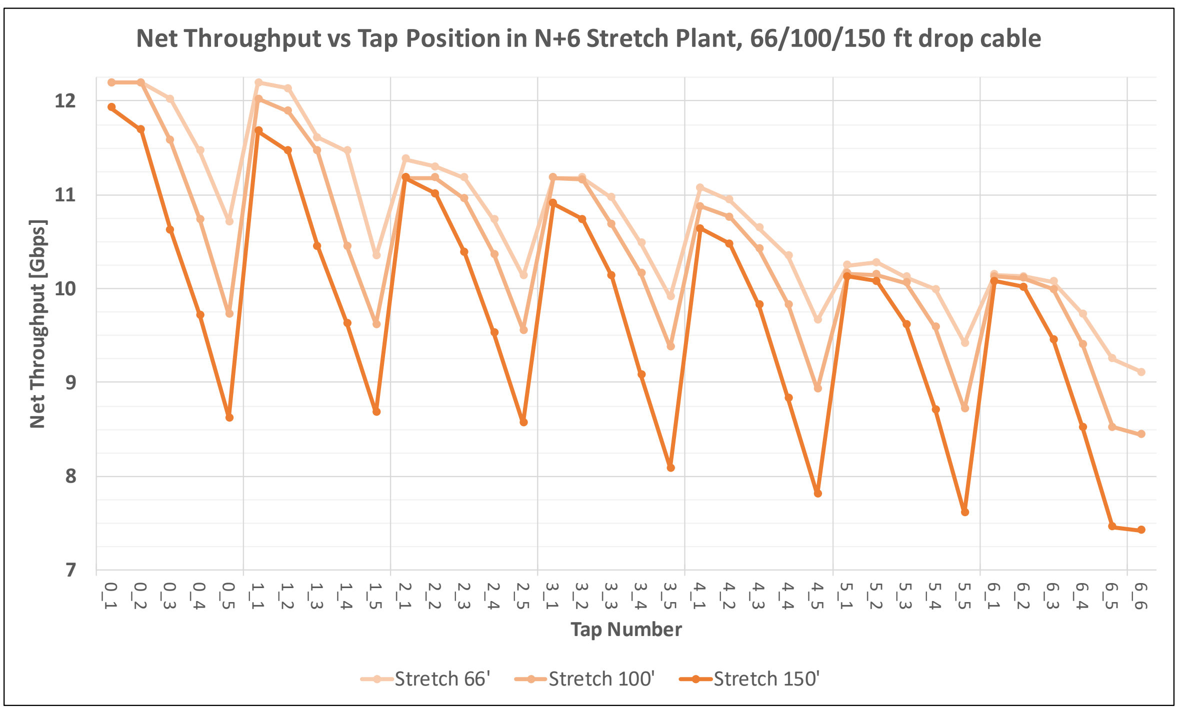

Figure 22 – Net Throughput vs Tap Position in N+6 Stretch Plant, 66/100/150 ft drop cable

In general, the short drop cable makes a significant difference, especially for the fourth and fifth taps. For the typical plant, there is less than a 3% difference on the first two taps after the amplifier for 100' compared to 150'. The 4th tap has gains in the 7-10% range while the 5th tap sees 7-15% improvements.

The stretch plant shows even wider spreads between 100' and 150'. The gains on the first two taps can be up to 4%. The 4th tap sees improvements up to the 8-12% range and the 5th tap sees 10-16% gains with the shorter 100' drops.

> *Note that replacing a 150' RG-6 drop cable with RG-11 will result in improved performance that is very close to the 100' RG-6 drop cable.*

Migrating from 100' to even shorter 66'drop cables did not have as big effect, especially on the typical plant. The 66' drop cable did show some gains over the 100' drop on the 4th and 5th taps of the stretch plant. The largest gains being 7-10% after the first two amplifier and then shrinking gains with more amplifiers in the cascade.

6.4. Impact of Super-stretched Links (56dB)

As shown on Figure 10, a small percentage of multi-port bridger amplifiers have extremely large dB spacings that are >55 dB. Figure 23 shows the impact of inserting one or two of these super-stretched links with 56 dB spacing into a typical plant. For reference, 56 dB spacing between amplifiers could represent any one of the following scenarios:

1. 1560 ft of 0.500 cable
2. 1920 ft of 0.625 cable
3. 2640 ft of 0.875 cable
4. Or 1560 ft of 0.625 plus 5 taps

In general, it would be expected that these super stretched links are express feeder runs over a long distance without any taps. For this example, Figure 23 shows scenario 4 with 1560' of 0.625 cable plus 5 1.8 GHz taps to demonstrate the impacts of the distance on these taps.

For the first case, a super-stretch 56 dB link is inserted following the node. As can be seen, the total capacity seen at the output of the 1st amplifier drops by ~0.5 Gbps at the 1st tap and over 1 Gbps by the 5th tap relative to the typical plant. After each additional amplifier in the cascade, the delta with the typical plant shrinks as the noise introduced by each amplifier accumulates and starts to dominate.

For the second case, the super-stretch 56 dB link is inserted after the 3rd amplifier. The outputs after the node, 1st and 2nd amplifiers are identical to the typical plant. The outputs after the 4th, 5th and 6th amplifiers are very close to the first case in which the super-stretch link is at the node.

The final case shows a scenario where two super-stretch links are inserted, one after the node and one after the 3rd amplifier output. Adding the second super-stretch link shows a noticeable degradation to the total network capacity after the 4th, 5th and 6th amplifier outputs.

Figure 23 – 492-1794 Typical Plant with 1 or 2 Super-stretch 56 dB Links

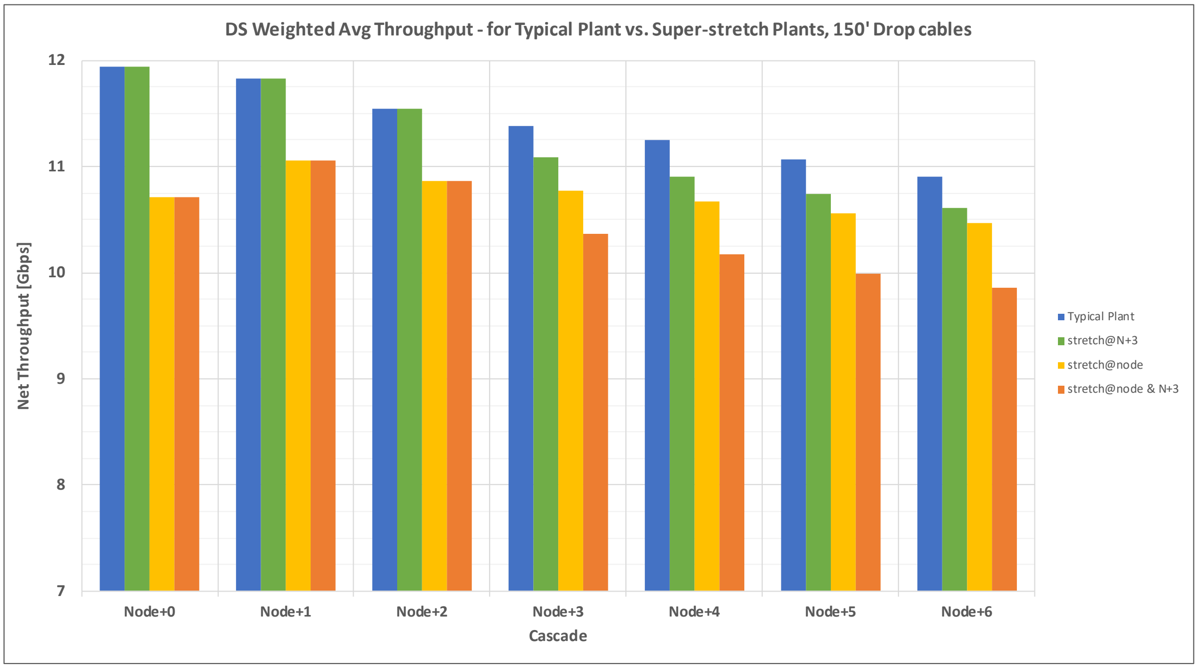

Figure 24 –Weighted Average DS Capacity with 1 or 2 Super-stretch 56 dB Links

Figure 24 compares the DOCSIS optimized weighted average DS capacity for the three super-stretch cases compared to the typical plant. With the CMTS scheduler effectively averaging capacity across all the taps, this helps to minimize the impact of these super-stretched links. For Node+6 plant, adding one super-stretched link drops average capacity by <5% while adding two super-stretched links only has ~10% impact.

6.5. Downstream ESD Capacity for various US/DS splits

The DOCSIS 4.0 specification allows for many different upstream (US) options. For a static ESD system, the possible diplexer splits are:

- 85/108 MHz
- 204/258 MHz
- 300/372 MHz
- 396/492 MHz
- 492/606 MHz
- 684/834 MHz

The first number is the upper edge of the upstream spectrum. The second number is the starting edge of the downstream. The guard band is the region between these two and is not usable. Notice that the guard band grows larger as the US/DS split frequencies increase.

Our network capacity modeling research has shown time and again that downstream bandwidth capacity is the limiting factor in almost every analysis. So, it is extremely important to try and maximize the downstream bandwidth capacity.

The DOCSIS 4.0 specification also supports a dynamic FDX option. For FDX, the shared band starts at 108 MHz and can go up to 204, 300, 396, 492 or 684 MHz. The FDX downstream goes up to 1218 MHz. So downstream bursts can potentially leverage the full 108-1218 MHz.

Figure 25 shows the downstream capacity at each tap location for various static ESD US/DS splits for a N+6 typical plant and 150’ drop cables. It also compares these to a 108-1218 MHz DS similar to that used by FDX (although full-duplex operation may add other impairments which are not factored into this example). Figure 26 then shows the DOCSIS optimized DS weighted average capacity for the various static ESD splits and the dynamic FDX for various cascade lengths from N+0 to N+6.

Because the changes between the US/DS splits are in the lower frequencies (i.e. below 684 MHz), the capacity differences remain constant. For each static ESD curve, there is a drop in capacity with increasing taps. This is the same for all static ESD scenarios, so they all have the same shape. Note that the 684/834 split has a noticeably larger gap than the other splits. This is a combination of two factors. First, it has the largest guard band. Second, its downstream is missing a lot of the lower frequencies which tend to have higher bit-loading as shown previously in Figures 14 and 15.

The 108-1218 MHz FDX DS spectrum is also shown for reference and appears to align closest with the 492/606 MHz static ESD DS capacity. The next section takes a closer look at that.

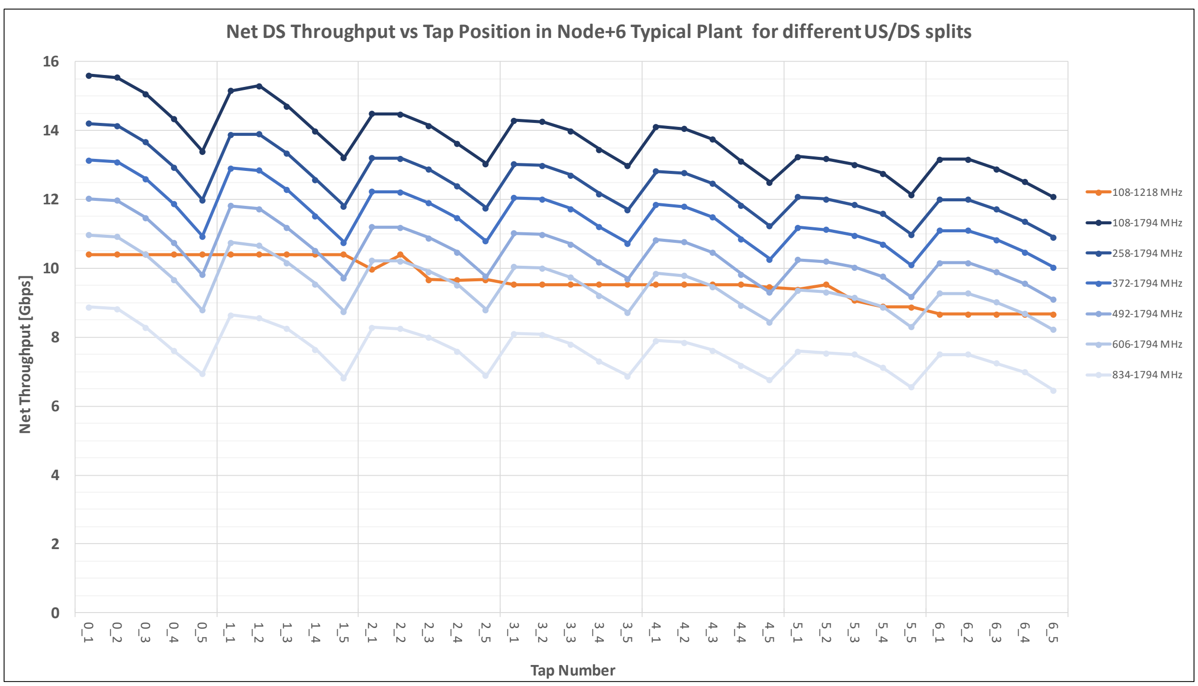

Figure 25 – Net Throughput vs. Tap Position in N+6 Typical Plant, different US/DS splits

Table 3 looks at some potential service tier combinations that could be considered for the various ESD split options. These SLAs are based on the basic traffic engineering formula discussed earlier in section 3. Remember that it is a function of the number of subscribers, Nsub, and the average peak period consumption per sub, Tavg. Both may change over time. For this table, the base assumption is Nsub=200, DS Tavg=20 Mbps, US Tavg=1.2 Mbps and K=1.0-1.2. Note that 100 subs @ DS Tavg=40 Mbps would be roughly equivalent.

By comparison, the SLA for 10G PON with 64 subs and same Tavg would be 6-7.5 Gbps for both the US + DS. The last row in Table 3 shows the theoretical best SLA's on a 1794 MHz plant. This uses dynamic soft-FDD where the CMTS scheduler switches between two diplexer settings (i.e. 85/108 MHz and 684/834 MHz) to time division multiplex the 108-684 MHz spectrum between US + DS.

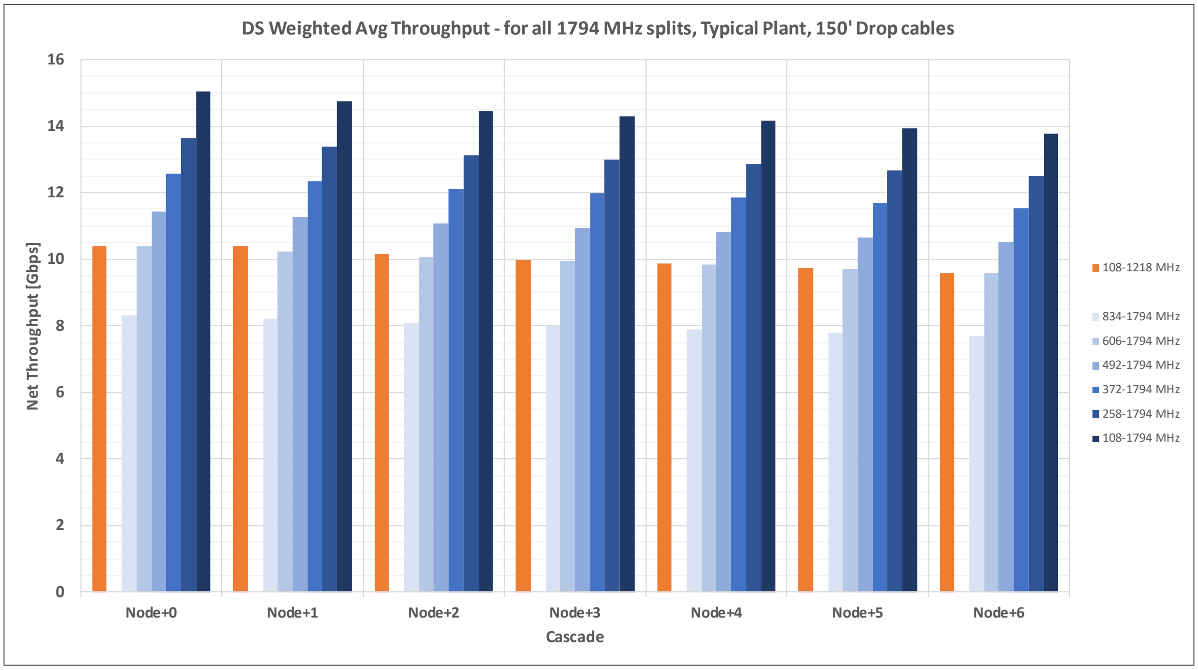

Figure 26 – Weighted Average DS Capacity for different US/DS splits

Table 3 – Summary of SLA Options for 1794 MHz Plants

10G HFC Options	Top of US (MHz)	Start of DS (MHz)	DS SLA (Gbps)	US SLA (Gbps)
1794/204 MHz	204	258	7 – 9	1 – 1.5
1794/300 MHz	300	372	6 – 8	1.5 – 2
1794/396 MHz	396	492	5 – 7	2.5 – 3
1794/492 MHz	492	606	4 – 6	3 – 3.75
1794/684 MHz	684	834	3 – 4	4 – 5
1794/85 MHz with 108-684 MHz dynamic Soft-FDD	684	108	8 – 10	4 – 5

6.6. Comparing 108-1218MHz FDX vs. 606-1794MHz ESD DS Capacity

The DOCSIS optimized capacity for both the static 606-1794 MHz ESD and the 108-1218 MHz dynamic FDX DS capacities are compared in Figure 27. This gives a better handle on how these two stacks up against each other. Note – FDX in N+X plant would require FDX-capable amplifiers which are still under investigation.

For typical plant conditions, the DOCSIS optimized weighted average capacity for the 606-1794 MHz ESD is almost identical to the 108-1218 MHz FDX system.

On a stretch plant, the 606-1794 MHz ESD takes a hit in the higher frequencies, especially above 1.5 GHz. This reduces its DS capacity by about 5% compared to the 108-1218 MHz DS in a stretch plant.

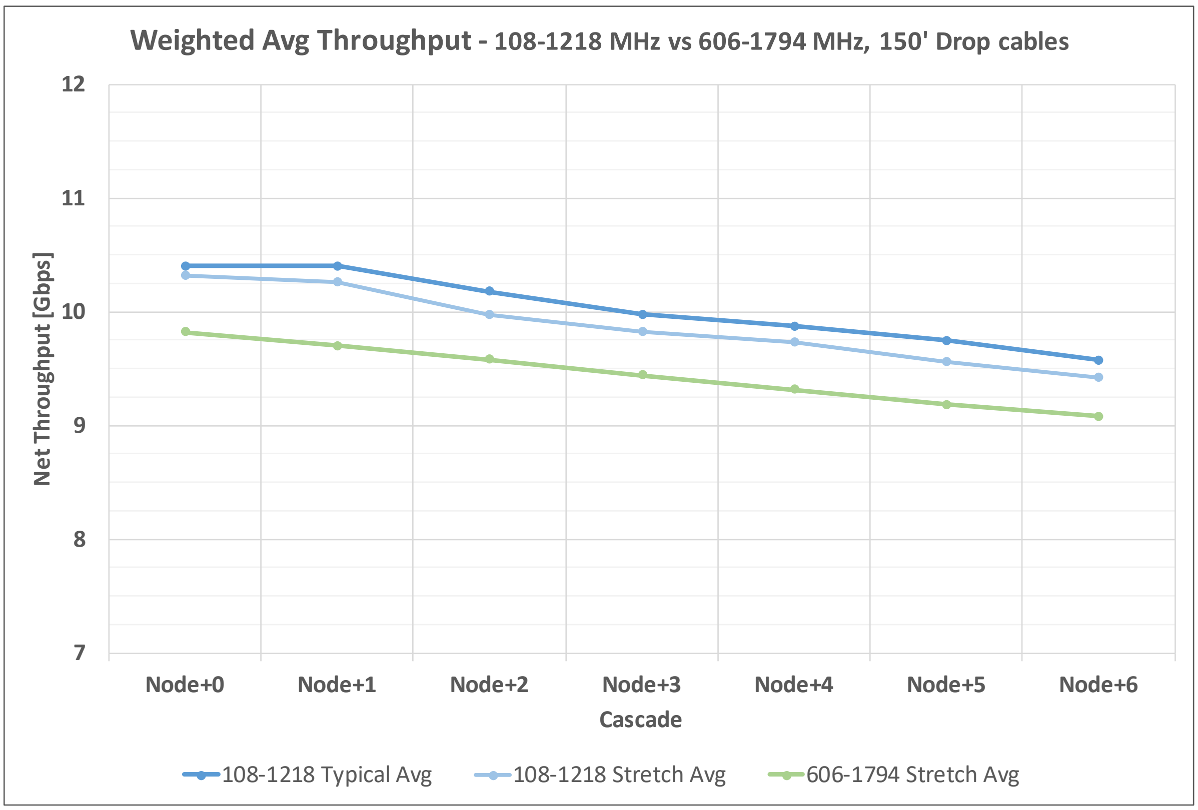

Figure 27 – Weighted Avg Throughput - 108-1218 MHz vs. 606-1794 MHz, 150' Drop cables

6.7. Improving on the DOCSIS Weighted Average Capacity

The DOCSIS weighted average capacity in the earlier analysis's assumes that the subscriber's data usage is relatively evenly distributed across the HFC plant AND the subscriber's data usage is evenly distributed across all the available OFDM channels. However, each channel has its own unique set of profiles that can vary significantly from channel to channel.

Figure 28 below looks at the bit loading for two extreme cases in the 492-1794 MHz plant. The first case is a home on the 1st tap after the 1st amplifier for a typical plant amp spacing. The second case is a home on the 5th tap after the 6th amplifier on a stretch plant.

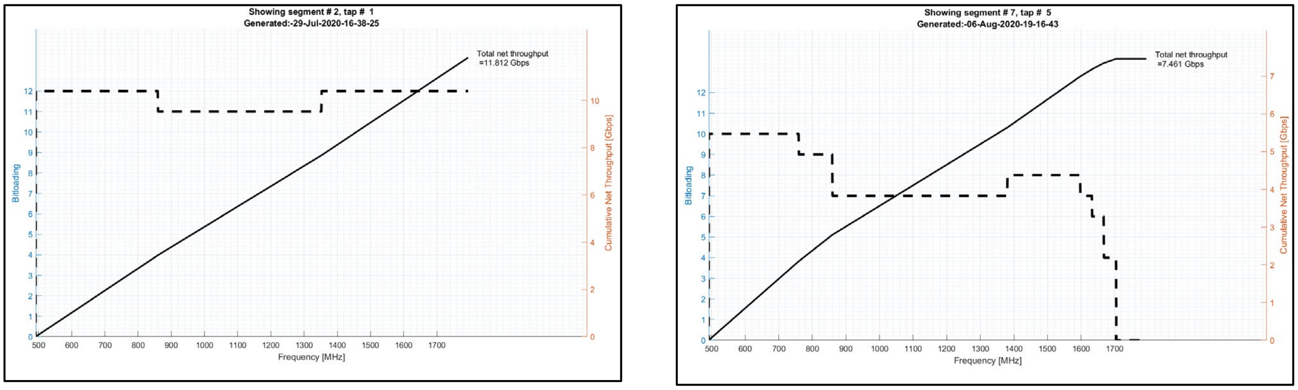

Figure 28 – Example Bit-loading for Near and Far Homes

As can be seen, the modem at the far home has vastly lower bit-loading capacity in the upper frequencies, and in particular, the top OFDM channels, i.e. 1602-1794 MHz. If the far modem has a 5 Gbps burst, the CMTS could schedule it in the 492-1218 MHz OFDM channels. Note that the near modem still has ~5 Gbps available to it in the top 3 OFDM channels from 1218-1794 MHz. So, both modems could be bursting to 5 Gbps simultaneously even though the EoL throughput for this plant is less than 7.5 Gbps. And the total burst rate is higher than the average of the two bit-load maps put together.

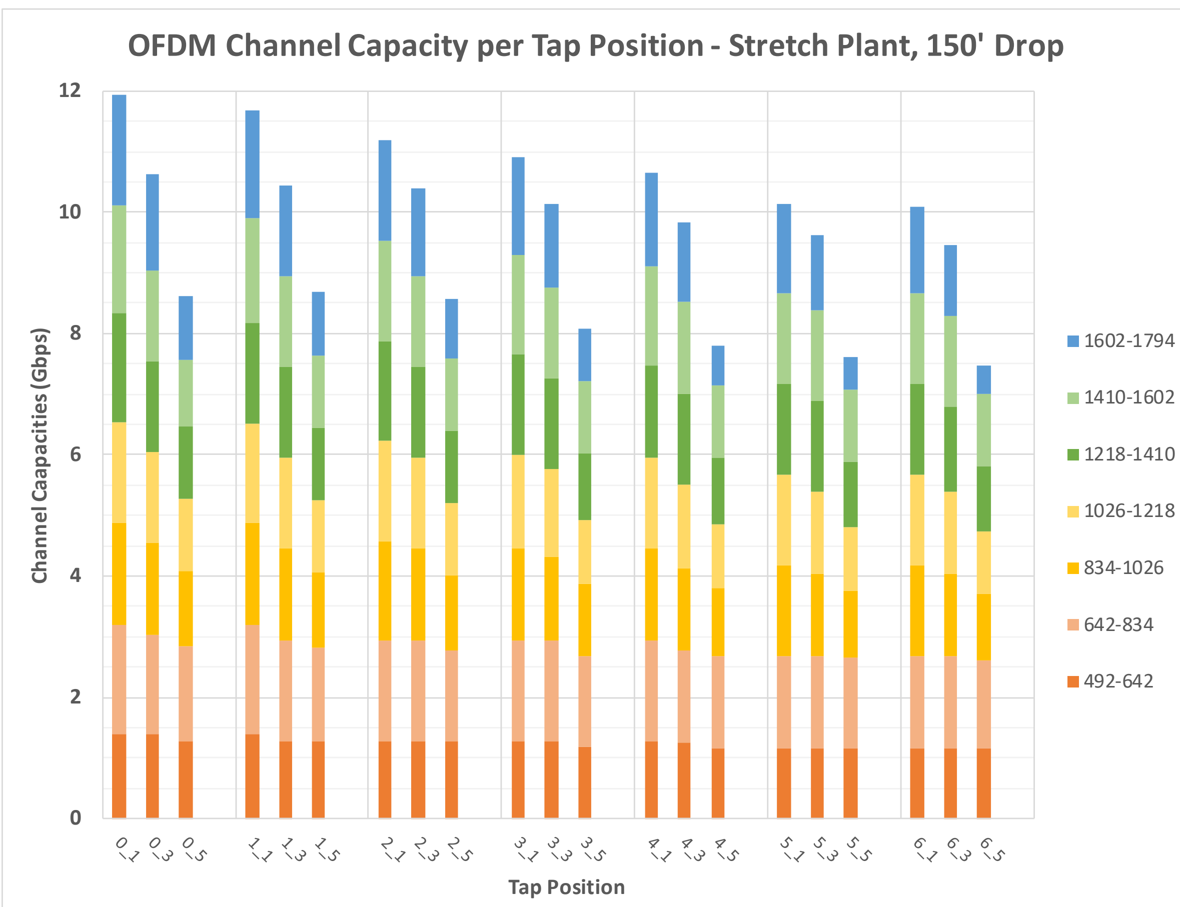

Figure 29 – Example Channel Capacities for various Tap Positions

Figure 29 shows how these bit-loadings map to the individual channel capacity for the 1st, 3rd and 5th tap on every amplifier output stage. As can be seen in the above channel capacities, the top OFDM channel for the 5th tap is practically unusable. For the other high frequency OFDM channels (i.e. 1026-1602 MHz) the 1st and 3rd taps can be 30% to50% higher capacities than the 5th tap location. Overall, the 1st tap position is getting reasonably good capacity across all the channels, even at the higher frequencies.

For over a decade, the CMTS schedulers have been load balancing across multiple DOCSIS 3.0 bonded channels. And they have been doing a great job at it. In the above scenario, the CMTS scheduler can tend to put data from the further taps in the lower frequency OFDM channels while the nearer taps utilize the relatively empty upper frequency OFDM channels. This means the CMTS can achieve even higher capacities than those shown with the weighted average capacity analysis above.

6.8. Handling the Corner Cases – Meeting Tmax Burst QoE

The DOCSIS weighted average capacity and the intelligent CMTS scheduler improvements addresses the total system capacity available to the scheduler. However, as seen previously, the capacity seen by any individual home can vary dramatically. The CommScope basic traffic engineering formula discussed in Section 3 can help determine the capacity needed for each individual home so it can obtain its appropriate Quality of Experience (QoE).

Consider a worst-case scenario where an operator wants to support a 7.5 Gbps DS service tier across their entire footprint, including 10G PON FTTH and DOCSIS 4.0 ESD 492/1794 MHz HFC plant. As seen in Figure 18, some home locations might only see total available capacity around 7.5 Gbps. That means that whenever that home wants to burst to its full 7.5G Tmax, it will need to use every ounce of spectrum available to it. That would leave practically nothing left for the remaining subs (i.e. Nsub*Tavg component from the formula).

In reality, the CMTS would not let these other subs starve, but will tend to allocate higher capacity lower spectrum to the bursting modem and others near the EoL. Meanwhile modems closer to the node and amplifier outputs will be using their higher capacity profiles in the higher frequency OFDM channels. So, the bursting modem might only get 5-6 Gbps out of its 7.5G tier during peak busy times.

So, how probable is this worst-case scenario? There are several factors that must all align for this to happen. These factors and the individual probability of each include:

- Subscriber takes the premium 7.5G DS top billboard tier [~1% to ~5%]
- Subscriber home is on the 5th tap [20%, i.e. 1 out of 5]
- The probability that this link even has 5+ taps [~30%]
    - Majority of amplifier outputs have 4 taps or fewer
- The tap is after the 5th or 6th amplifier [~25% of amplifier outputs on N+5/N+6 plant]
- % of MSO plants that are even N+5 or higher [1% to 20%???]
- The amplifier output is being stretched beyond unity gain [~5% to 15%]
- The RG-6 drop cable is >125' [~25% to 50%]

On N+5/N+6 plants, only 1 out of every ~200 to ~1,000 amplifier outputs might fall into this category. While it varies greatly by MSO, N+5 and N+6 plants are also becoming a smaller and smaller percentage of their total HFC plants. Of these potential problem amplifier links, it is only a problem if there is a subscriber on the 5th tap that takes the 7.5G DS top billboard tier. For all these stars to align, it looks like a probability on the order of four or five 9's that this won't happen.

It may be the end of this decade before the 7.5G DS tier is being offered and this worst-case scenario can potentially kick in. So, what options are available to the operator over this time window to handle this potentially rare event? Here are some possibilities:

1. MSO accepts slight QoE degradation (i.e. 7.5G drops to 5-6G during peak busy hour for these very infrequent customers)
2. MSO only offers 4G-5G DS service tier to this customer
3. Pull fiber to that subscriber's home and switch them to 10G PON
4. Reduce the N+5/6 cascade depth
5. Add mid-span amplifiers to boost higher frequency signals
6. Replace the RG-6 drop cable with RG-11

Option 1 might be very acceptable, but potentially dependent on your regulatory environment. The customer perceived QoE between 5 Gbps bursts and 7.5 Gbps burst capacity may not be perceptible. And this is only happening for rare times when the consumers bursts to the max and for a very small percent of the customer population. If the operator is still uncomfortable with this, they can choose option 2 and only provide the consumer with a slightly reduced Tmax (e.g. 5G instead of 7.5G). Note, this is something that the DSL world has had to always deal with, but this is not nearly as dramatic. In older DSL technology, customer capacity might have varied from 2 Mbps to 25 Mbps.

If the operator wants to correct the situation and provide additional capacity where needed, then options 3-6 would need to be considered over the next 5-8 years. An economic analysis of some of these choices is in the next section. Since this scenario does not need to get resolved until the end of the decade, the operator has plenty of time to correct the issue before it happens.

Everyone agrees that the long-term strategy is to eventually get to FTTH. However, many consider this a multiple decade transition to get to the point where fiber is pulled down every street. Option 3 jumps directly to FTTH as the solution. However, since this is a N+5/N+6 plant, chances are that the fiber is not nearly as close to the home as it might be in N+0/N+1 plant. It might be over a mile away from the customer. So, option 3 is expected to be the most expensive.

With the long-term strategy to continually push fiber deeper, then option 4 above is making a step in that direction. If these problem amplifier spans can be identified today (e.g. N+5/N+6 with stretched outputs), then the operator can spend the next decade gradually attending to these and pushing fiber deeper. Grey optics aggregator (GOA) / grey optics terminator (GOT) is a cost-effective fiber deep architecture that could be used to reduce the cascade length, without necessarily going all the way to N+0. The GOA functionality is placed at the parent node location. The downstream optical signal is split and passed on to the GOT node using short distance optics. The GOA also aggregates all of the GOT return signals. The GOT nodes are transparent to the head end and does not require any additional head end optics.

A low cost GOT fiber node can replace one of the multi-bridger amplifiers on the path to these problem spans to reduce the cascade length. If an operator has a 20 year plan to convert to a fiber deep network (e.g. 5% of plant per year), they can target these problem spans now to make sure issues with these are corrected before the end of the decade. Later in time, the GOT node can be upgraded to a full-fledged fiber node as needed.

Options 5 and 6 both fix the potential capacity issue but are not in alignment with the strategic direction of pushing fiber deeper towards eventual FTTH. Option 6, pulling a new drop cable, could be aligned with this strategy if it is a "siamese" cable that contains both fiber and RG-11 drops. Option 5, adding a mid-span amplifier, might make sense if the logistics do not allow option 4 to pull fiber deeper to reduce cascade lengths.

7. Cost and Logistic Analysis

Our starting point for economic analysis discussion is a 750 or 860 MHz HFC network case study depicted in Figure 30. The fiber node area, serving 400 HP, is also a single Integrated Converged Cable Access Platform (I-CCAP) service group. It is a 42/54 MHz "sub-split" system with a total of 20 RF amplifiers, 10 each of bridger and line extender type. There is a total of 100 taps, on 25,000 feet of hardline coax plant, for an average density of 84 HP/mile and ~4.2 amps/mile. The longest RF amp cascade length is N+4.

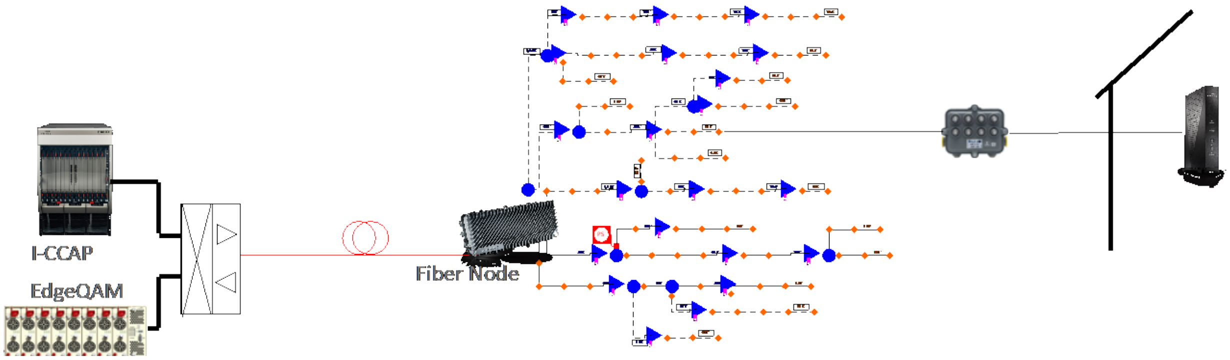

Figure 30 – I-CCAP N+4 HFC network with 42/54 MHz sub-split

The cost of upgrading this network to 1218 MHz with a 204/258 MHz high-split is analyzed by (a) replacing "ePack" node and amps modules and (b) adding appropriate I-CCAP license to augment both DS and US data capacity. The left-most column of Figure 31 depicts this network upgrade, along with the percent breakout of various elements. Furthermore, this 1218/204 MHz upgrade is normalized to 100%, in order to be the baseline to compare to the other cases, namely 1.2 GHz FDX N+0 and 1.8 GHz ESD N+4. The 2nd column shows the 1218/204 MHz upgrade where the tap faceplates are also upgraded to 1.2 GHz. This adds ~25% premium on top of the base case.

The network shown in Figure 30 is applicable to the 1.2 GHz high-split case. For the other two cases, however, networks of Figure 32 and Figure 33 are the respective representations. FDX requires a N+0 "fiber-deep" upgrade. The original fiber node plus 20 amplifiers end up being replaced by 4 fiber nodes with ~50% of the original hardline coax over lashed by new fiber. For the 1.8 GHz ESD N+4 case, whole-station amplifiers, as well as all 100 tap housings are replaced.

As shown in Figure 31, the FDX N+0 upgrade comes to 454% of the "base case", that is, the 1.2 GHz high-split upgrade. The fiber material plus labor costs really dominate for the N+0 upgrade. Our analysis assumed 80% aerial plant. This component could be much higher if the percentage of underground plant increases over 20%. The 1.8 GHz ESD upgrade comes to 214% of the "base case", or a 114% premium.

If only 1 to1.5 Gbps US service tiers are needed, the 1218/204 MHz upgrade is the most cost effective. However, once multi-Gbps US tiers are needed, then the FDX N+0 or the 1.8 GHz N+4 upgrade is required. Our study shows that the FDX N+0 upgrade is roughly twice the cost of the 1.8 GHz ESD N+4 upgrade for this N+4 plant case study.

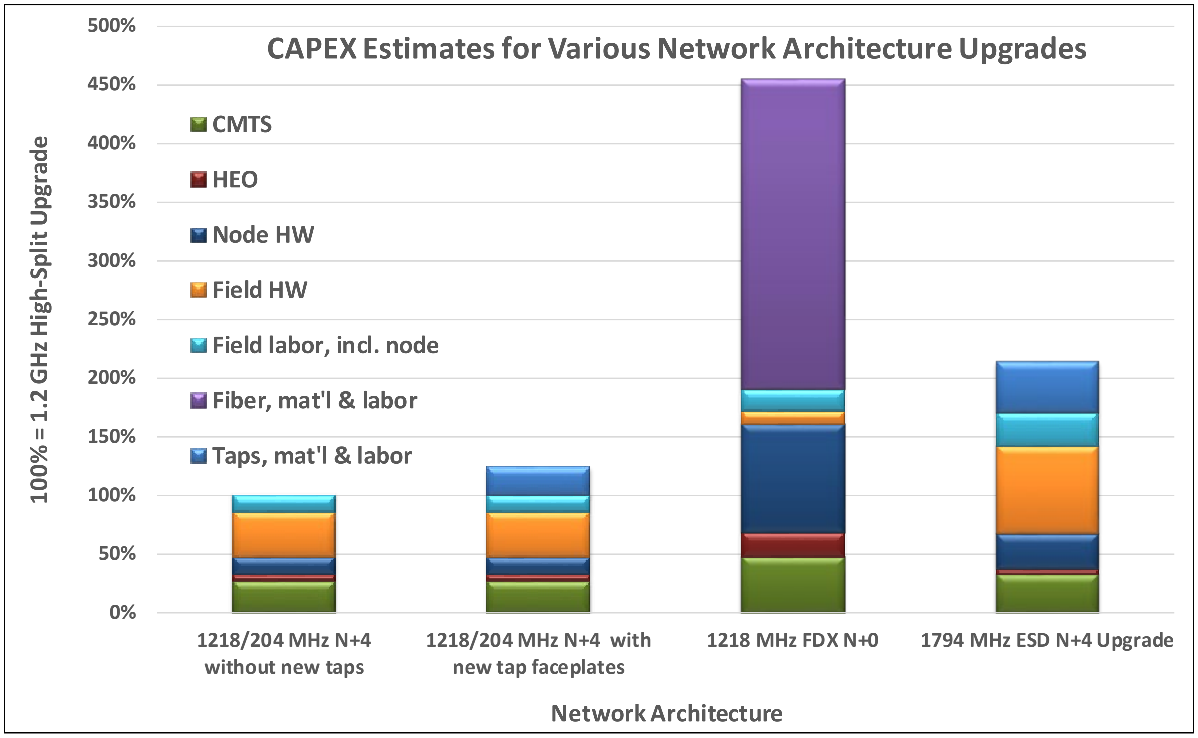

Figure 31 – CAPEX estimates for 1218/204 MHz N+4 (both without and with new tap faceplates), 1218 MHz FDX N+0, 1794 MHz ESD N+4 upgrades

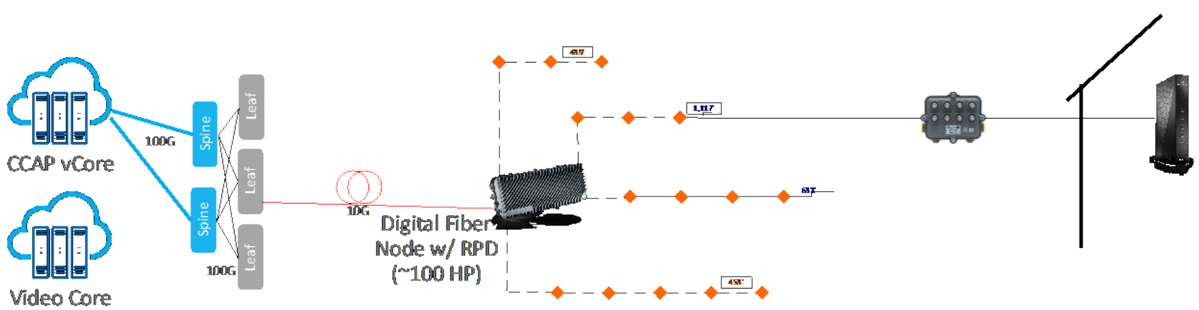

Figure 32 – DAA N+0 HFC Network Model for 1.2 GHz FDX Upgrade

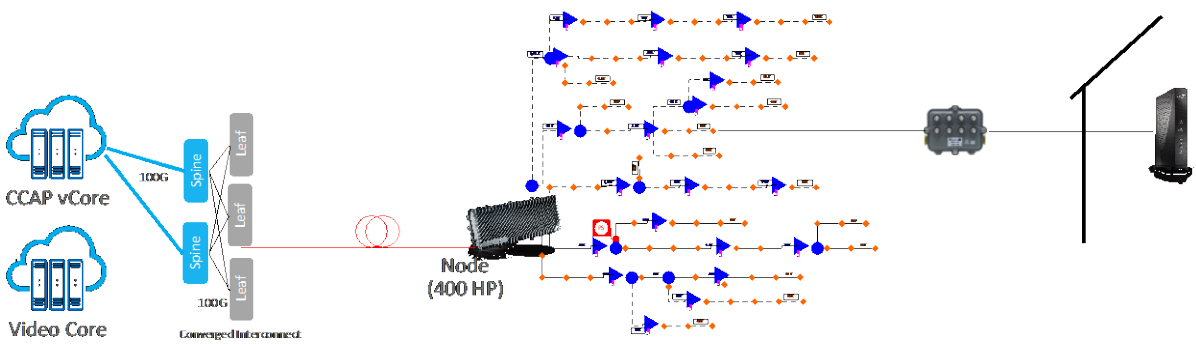

Figure 33 – DAA N+4 HFC Network Model for 1.8 GHz ESD upgrade

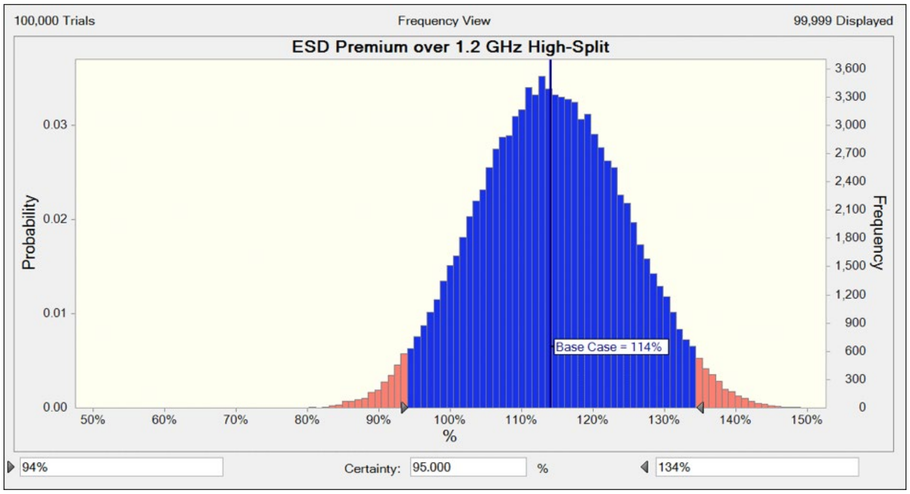

Figure 34 – Monte-Carlo analysis of "premium" for 1.8 GHz ESD case

There are other benefits to the ESD approach. For one, it is providing dedicated upstream capacity while FDX is relying on statistical time sharing of US+DS spectrum. The second is that ESD is a much simpler system than FDX which should result in OPEX savings through reduced maintenance and easier to diagnose plant issues.

Because many elements go into an evaluation of this type, a Monte-Carlo analysis was run with various assumptions varying over a reasonable range for the ESD upgrade case. A total of 100,000 Monte-Carlo samples were run to see how much the ESD "premium" might vary. Figure 34 shows that the ESD "expected" premium of 114% expands into a 94% to 134% "confidence interval".

What is driving the 114% ESD premium, while the FDX N+0 premium is 354%? Figure 31 offers some answers: fiber over lash required for N+0 accounts for 264% of the 354% premium for FDX! The 2nd most dominant FDX component is node hardware at ~90%. For ESD, RF amplifiers, labor to replace those, and a material and labor to upgrade tap housings and faceplates to 1.8 GHz account for most of the cost.

For the ESD case, a reasonable "workload" for the whole amplifier and taps upgrade were considered, especially allowing for the whole housing to get replaced. This assumed there is 1 tap per 4 HP; the RF amplifiers are closer to 1 to 20 HP ratio, and nodes are even further – in the assumptions made above it's 1 to 400 HP. If an amplifier module takes approximately 1 hour to replace and a tap housing takes 45 minutes to replace, how much of an additional labor force would an operator need in order to do these upgrades?

For an MSO serving a region with population of one million people (e.g. ~400,000 HP), Table 4 holds the answer. For the taps portion, it would require an additional full-time crew of ~10, working full-time, over a 5-year period; and doing nothing else. One may describe it as a large task, but 10 additional employees to do this task seems reasonable.

The previous section discussed a corner case where a subscriber with a long RG-6 drop cable at the EoL on a N+5/N+6 stretch plant does not have sufficient network capacity to obtain the highest service tier. Of the half dozen potential options mentioned, the most cost effective would be to replace the drop cable with RG-11. The operator might also consider putting it in a conduit for reliability and ease of later upgrading to a fiber drop as well as consider siamese cable/fiber pair. Adding a mid-span boost amplifier for high frequencies on the last link might triple or quadruple the cost of the drop cable option. And, more than one mid-span boost amplifiers might be needed on some of the other links in the cascade. Neither of these options pushes fiber deeper so they may not be aligned with the longer strategic direction.

Table 4 – Estimate of Labor force required to replace amps and taps in 5 years

	Nodes	RF Amps	Taps	Homes-Passed
Quantity	1,000	20,000	100,000	400,000
Task duration	1 Hour	1 Hour	45 minutes	
24/7 Person-Years	0.1	2.3	8.6	
40hr/week Person-Years	0.5	9.6	36.1	
30hr/week Person-Years	0.6	12.8	48.1	
Required size crew to complete the task in 5 years	**<<1**	**<3**	**<10**	

Trying to jump directly to FTTH can be extremely costly. Since this home is EoL on a N+5/N+6 plant, it might be up to a mile and a half from the fiber node. The costs of pulling fiber can skyrocket, especially if the plant is mostly underground. Some estimates show this to be 100 times the cost of just replacing the drop cable. This is hard to justify for a single customer.

One strategy is to push fiber deeper over a 20-year window to get most customers within 1000' or 1500' of fiber access. At this point, the operator can offer FTTH on demand when and even if needed (e.g. 90%+ of subs may stay on cable 'forever'). To achieve this, the operator can focus on the N+5/N+6 stretch plants in the near-term, with the goal of getting all of their HFC down to N+2/N+3 by the end of the 2020 decade. Then in the decade that follows, the operator can push fiber deeper to achieve N+0/N+1 that can enable FTTH on demand by the end of the 2030 decade. Analysis such as [ULM_2016] have shown that fiber deep N+0/N+small upgrades are more cost effective than FTTH for HFC brownfield upgrades.

8. Conclusion

The title of the paper posed the question "Is "Unity Gain" Still the #1 Objective?" as the cable world migrates to 1.8 GHz DOCSIS 4.0 plant. The paper has shown how much more difficult the 1.8 GHz unity gain task will be compared to 1218 MHz. Looking at stretch plant, the unity gain starts to come up short and the cracks become obvious at the longer cascades like N+5 and N+6.

Previously, conventional wisdom would have said that the amplifiers would need to be re-spaced; a potentially expensive proposition. Recently, there has been discussions of adding a mid-span amplifier to boost the gain for these links, which should not be taken lightly, as this adds more active components into the system.

Instead, inherent DOCSIS capabilities can leveraged to compliment unity gain when it comes up short. Previous thoughts of using EoL throughput as an indicator of plant capacity are no longer accurate at 1.8 GHz. DOCSIS 3.1/4.0 OFDM channels with variable bit-loading and multiple profiles enable the CMTS scheduler to maximize the system capacity. This was shown nicely in Figures 18 & 20, which are repeated in Figure 35 below. The N+6 stretch plant saw DOCSIS weighted average capacity gains that were 35% higher than EoL throughput.

The paper explored some of the variables that impact system capacity. In addition to amplifier spacings (e.g. typical vs. stretch plants), it looked at various cable drop lengths and the addition of super-stretched links into a typical plant. These are all scenarios that allow DOCSIS to optimize the system capacity some more.

The 1.8 GHz ESD plant offers several options for different upstream splits, varying all the way up to 684/834 MHz. Choosing different splits lets the operator balance between upstream and downstream bandwidth. The paper looked at the DS capacities associated with the different splits and then showed the range of DS + US service tier SLA combinations that might be supported near the end of this decade.

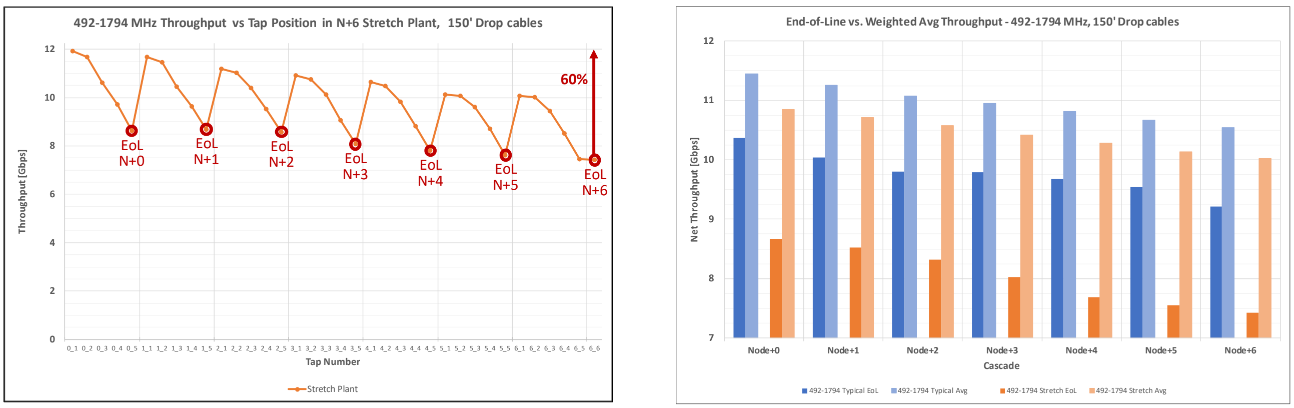

Figure 35 – EoL vs. Weighted Avg Throughput – Typical and Stretch Plants

The cost and logistic analysis section hopefully gave the reader some sense of the economic tradeoffs that will be encountered when looking at these various options. When the operator finally needs multi-Gbps US tiers, the 1.8 GHz ESD is more cost effective, potentially half of the cost of a FDX N+0 upgrade.

So, to answer the question from our title, the answer is: Maybe Yes! The authors still believe that trying to maintain unity gain is a key objective in any HFC design. It has shown its worth time and time again. However, as it approaches the breaking point, DOCSIS scheduling can fill the gaps on many of these plants on the bubble, reducing the number of HFC links that need drastic action to a significantly smaller amount. The two work well together.

Acknowledgements

The authors would like to thank Cox Communications and in particular David Job, Principal Eng of OSP Engineering at Cox, for providing the data and insights into the HFC amplifier spacing distribution.

The authors are also extremely grateful to everyone that gave us insights and advice on HFC design and doing a sanity check on what “reasonable” might be. In particular, we want to thank Stuart Eastman of CommScope for sharing some of network design wisdom; and Jay Lazorcik of CommScope for HFC network and unity gain diagrams.

Bibliography & References

[ALB_2019] A. Al-Banna et. al., "Operational Considerations and Configurations for FDX & Soft-FDD," SCTE Cable-Tec 2019, SCTE

[CLO_2019] T. J. Cloonan et. al., "Capacity Planning, Traffic Engineering, And HFC Plant Evolution For The Next 25 Years," SCTE Cable-Tec 2019, SCTE

[CLO_2017] T. J. Cloonan et. al., "The Big Network Changes Coming with 1+ Gbps Service Environments of the Future," SCTE Cable-Tec 2017, SCTE

[CLO_2016] T. J. Cloonan et. al., "Using DOCSIS to Meet the Larger BW Demand of the 2020 Decade and Beyond," NCTA Spring Technical Forum 2016, NCTA

[CLO_2014] T. J. Cloonan et. al., "Simulating the Impact of QoE on Per-Service Group HSD Bandwidth Capacity Requirements," SCTE Cable-Tec 2014, SCTE

[CLO_2013] T. J. Cloonan et. al., "Advanced Quality of Experience Monitoring Techniques for a New Generation of Traffic Types Carried by DOCSIS," NCTA Spring Technical Forum 2013, NCTA

[EMM_2014] "Nielson's Law vs. Nielson TV Viewership for Network Capacity Planning," Mike Emmendorfer, Tom Cloonan; The NCTA Cable Show Spring Technical Forum, April, 2014

[FDX_PHY] "DOCSIS 4.0 Physical Layer Specification", CM-SP-PHYv4.0-D01-190628, CableLabs 2019

[FDX_XSD_IBC] "Full duplex DOCSIS & Extended Spectrum DOCSIS Hold Hands to Form the 10G Cable Network of the Future", by F. O'Keeffe et. al., IBC 2019

[ULM_2019] J. Ulm, T. J. Cloonan, "The Broadband Network Evolution continues – How do we get to Cable 10G?", SCTE Cable-Tec 2019, SCTE

[ULM_2018] J. Ulm, "Making room for D3.1 & FDX – Leveraging Something Old that is New Again!", SCTE Journal of Network Operations : Find Fresh Approaches to Plant-Related Topics, Vol 4. No. 1. Dec 2018, SCTE

[ULM_2017] J. Ulm, T. J. Cloonan, "Traffic Engineering in a Fiber Deep Gigabit World", SCTE Cable-Tec 2019, SCTE

[ULM_2016] "Adding the Right Amount of Fiber to Your HFC Diet: A Case Study on HFC to FTTx Migration Strategies", John Ulm, Zoran Maricevic; 2016 SCTE Cable-Tec Expo

[ULM_2014] "Is Nielsen Ready to Retire? Latest Developments in Bandwidth Capacity Planning", John Ulm, T. Cloonan, M. Emmendorfer, J. Finkelstein, JP Fioroni; 2014 SCTE Cable-Tec Expo

Abbreviations

BW	bandwidth
CAGR	compounded annual growth rate
CAPEX	capital expense
CCAP	Converged Cable Access Platform
CDF	cumulative distribution function
CM	cable modem
CMTS	Cable Modem Termination System
CPE	consumer premises equipment
D3.1	Data Over Cable Service Interface Specification 3.1
D4.0	Data Over Cable Service Interface Specification 4.0
DAA	distributed access architecture
DOCSIS	Data Over Cable Service Interface Specification
DS	downstream
EOL	end of line
EPON	Ethernet Passive Optical Network (aka GE-PON)
ESD	extended spectrum DOCSIS
FDX	full duplex (i.e. DOCSIS)
FTTH	fiber to the home
FTTx	fiber to the ‘x’ where ‘x’ can be any of the above
Gbps	gigabit per second
GHz	gigahertz
GOA	grey optics aggregator
GOT	grey optics terminator
HEO	head end optics
HFC	hybrid fiber-coax
HP	homes passed
HW	hardware
I-CCAP	Integrated Converged Cable Access Platform
LDPC	low density parity check (FEC code)
LE	line extender
MAC	media access control
MB	multi-port bridger
Mbps	megabit per second
MHz	megahertz
MSO	multiple system operator
N+0	node+0 actives
NCTA	The Internet & Television Association
OFDM	orthogonal frequency-division multiplexing
OPEX	operating expense
PHY	physical interface
PON	passive optical network
PSD	power spectral density
QAM	quadrature amplitude modulation
QoE	quality of experience

RF	radio frequency
SG	service group
SCTE	Society of Cable Telecommunications Engineers
SNR	signal to noise ratio
STB	set top box
Tavg	average bandwidth per subscriber
TCP	total composite power
Tmax	maximum sustained traffic rate – DOCSIS Service Flow parameter
TX	transmit
US	upstream