The Broadband Network Evolution Continues – How Do We Get To Cable 10g?

A Technical Paper prepared for SCTE•ISBE by

John Ulm
Engineering Fellow
CommScope
22 Topside Rd, Moultonboro NH 03254
john.ulm@commscope.com

Tom Cloonan
CTO- Network Solutions
CommScope
2400 Ogden Ave.- Suite 180, Lisle, IL 60532
tom.cloonan@commscope.com

Table of Contents

Title	Page Number

List of Figures

Title	Page Number

Introduction

On January 7, 2019 at the Consumer Electronics Show (CES), NCTA – The Internet & Television Association®, CableLabs®, and Cable Europe® introduced the cable industry's vision for delivering 10 gigabit networks, or 10G™ – a powerful, capital-efficient technology platform that will ramp up from the 1 gigabits per second (Gbps) offerings of today to speeds of 10 Gbps and beyond – to consumers across the globe in the coming years.

10G is a goal, a lighthouse in the distance towards which all MSOs and vendors can use as a beacon in the night to steer their boats safely away from danger. It will take some time to get there. It is actually only a single point in a continuum of improvements that will occur in the future. It is not the end-point, it is an interim point. We will likely push on beyond that 10G point in the future. The focus of this paper is the migration needed over the next 7-10 years using existing technologies to achieve the 10G goals. Another paper [CLO_2019] takes a deeper look into the future with some potential new technologies to see where the industry might be in 15-25 years beyond 10G.

To reach the 10G goals might require an aggregate of last-hop technologies, including: HFC, DOCSIS, PON, and Wireless which are all likely candidate technologies. It turns out that 5th generation (5G) mobileand 10G HFC are quite complementary and will help one another. This is discussed further in [ULM_2019].

But in getting to 10G, what does it take for the DOCSIS/HFC system to actually deliver on the 10G Service Level Agreement (SLA) promise? Many MSOs do not really know what it takes to do 10 Gbps downstream, let alone 10 Gbps symmetrical.

The cable industry is going through unprecedented technology changes starting with the introduction of DOCSIS 3.1 (D3.1); then Distributed Access Architectures (DAA) such as Remote PHY (R-PHY); DOCSIS Full-Duplex (FDX); Low Latency DOCSIS (LLD) initiative; and now Extended Spectrum DOCSIs (ESD) efforts. These are all building blocks to help us reach Cable 10G. But what are the logistics to get on the right path?

This paper will discuss the importance of these technologies and address the following migration topics:

- 10G network capacity planning
- Outside plant considerations and logistics
- CPE considerations

Cable 10G is an ambitious initiative, but it is in keeping with the broadbandheritage to which we've grown accustomed. We can get there with the right roadmap.

Content

1. Cable 10G™ – A Journey, Not a Destination

A quick introduction to 10G™ can be found on the CableLabs website www.cablelabs.com/10g and given below:

Figure 1 – 10G™ Logo

What is 10G™?

The 10G platform is a combination of technologies that will deliver internet speeds 10 times faster than today's networks and 100 times faster than what most consumers currently experience. Not only does 10G provide faster symmetrical speeds up to 10 Gbps, but also lower latencies, enhanced reliability, and better security in a scalable manner.

Why Do We Need the 10G Platform?

Our digital future will stall without a platform that can meet our needs. While we don't know what the next trend will be, we do know the internet will be central to it. By advancing device and network performance to stay ahead of consumer demand, 10G will provide a myriad of new immersive digital experiences and other emerging technologies that will revolutionize the way we live, work, learn, and play. Like the saffron in paella, or the milk in a latte, our industry's networks and innovations are the crucial ingredients in creating a better future for humanity.

In the downstream (DS) direction, 10G will be helpful in providing delivery of immersive video services (virtual reality & augmented reality for "Holodeck Experiences"). It will also be useful for providing more "snappy" services. For example, downloading:

- A two-hour HD movie in 3-4 seconds (vs. 5-7 min @ 100 Mbps)
- A two-hour 4K UHD movie in 12-15 seconds (vs. 20-25 min @ 100 Mbps)
- A large gaming program such as Call of Duty's Black Ops (~100 GBytes) in 90 seconds
    - instead of an 2½ hours @ 100 Mbps.

In the upstream (US) direction, it could be used for providing more "snappier" services again, but it could also prove to be very useful in enabling low-latency transport. The extra bandwidth (BW) helps enable Proactive Grant Services (PGS) to accelearate US delivery. That lower latency can permit 5G mobile backhaul and midhaul, and if the latencies drop low enough, it could even permit 5G fronthaul.

10G may also enable many different commercial applications such as remote medical procedures.

There are four key attributes to the 10G platform:

1. Speed
2. Latency
3. Security
4. Reliability

10G's promise of faster speeds, more capacity, lower latency and greater security will enable and help fully realize a wide variety of new services and applications that will change the way millions of consumers, educators, businesses and innovators interact with the world. Much of the underlying technology has already been specified or is a work in progress.

The speed attribute will leverage technologies such as :

- DOCSIS 3.1
- DOCSIS 4.0 - FDX, Extended Spectrum DOCSIS
- PON
- Coherent Optics
- Advanced WiFi including WiFi 6 (a.k.a. 802.11.ax)

Some applications driving the need for lower latency DOCSIS include: gaming; VR/AR (avoiding nausea); CoMP; and autonomous navigation. Latency is a function of packet processing times, queeng times, transmission times, and propagation times. We can improve all areas. CableLab's Low Latency DOCSIS (LLD) project includes ideas in all these areas. The existence of 10G BW capacity also helps, because higher bandwidth capacity leads to less congestion and also permit new techniques like Proactive Grant Service (PGS) to expedite BW grants in the upstream (US). Work is also being done in low latency mobile X-haul and low latency Wi-Fi.

Security is an integral part of 10G. Work continues at CableLabs in Micronets, secure downloads, & MACsec. This will become part of the new DOCSIS 4.0 specification. Vendors are also working to make more secure systems with separate, isolated processors & memory in chips.

With respect to reliability, new DAA node designs of the future will likely be adding in redundancy in processing and redundancy in NSI-Side links to Leaf Switches in CIN as Moore's Law improvmenets in gate density help. Reliability is being addressed by proactive network maintenance (PNM) and dual channel Wi-Fi. This will improve observability and redundancy and A/I to monitor (PNM, more data analytics in DAA nodes, redundancy in ring of DAA)

And for all four 10G attributes (i.e. speed, latency, security, reliability), it is equally important that they scale… on all markets.

The remainder of this paper will now focus on the speed and capacity required to achieve the 10G goals.

2. Network Capacity Planning for 10G

The "10G" in the announcement is for 10 Gigabits per sec (Gbps). But what exactly does that mean as there are different ways of measuring capacity? For example, Liberty Global's Virgin Media division in the U.K. ran tests earlier this year over a 10G EPON network and demonstrated users getting 8.5 Gbps throughput. It turns out that this is extremely close to the theoretical maximum throughput for 10G PON technology. "10G" PON has a raw physical rate of 10 Gbps but there is ~15% overhead from the PHY and MAC layers. So the customer actually nets 8.5 Gbps from the 10G PON.

Our analysis first takes a look 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. Finally, some subscriber migration strategies are discussed that will help the transition to 10G platforms.

2.1. Traffic Engineering for 10G Networks

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

2.1.1. The "Basic" Traffic Engineering Formula

Previously, [CLO_2014] provided an introduction to 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 cable 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 busy-hour period.

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

$$C >= (Nsub*Tavg) + (K*Tmax\_max) \quad (1)$$

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 <= K <= infinity, but typically 1.0 <= K <=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. A change in the K value results in a corresponding change within the QoE levels experienced by the subscribers who are sharing the service group bandwidth capacity (C). Lower K values yield lower QoE levels, and higher K values yield higher QoE levels).

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 SGs would need larger values of K while very small SGs might use a K value less than 1.0.

2.1.2. Broadband Subscriber Traffic Consumption

ARRIS/CommScope 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 maps closely to many other MSOs globally.

The chart below, Figure 2, shows the average subscriber downstream consumption, DS Tavg, during peak busy hours for a number of MSOs over a ten year period. At the start of 2019, DS Tavg was approaching the 2 Mbps barrier.

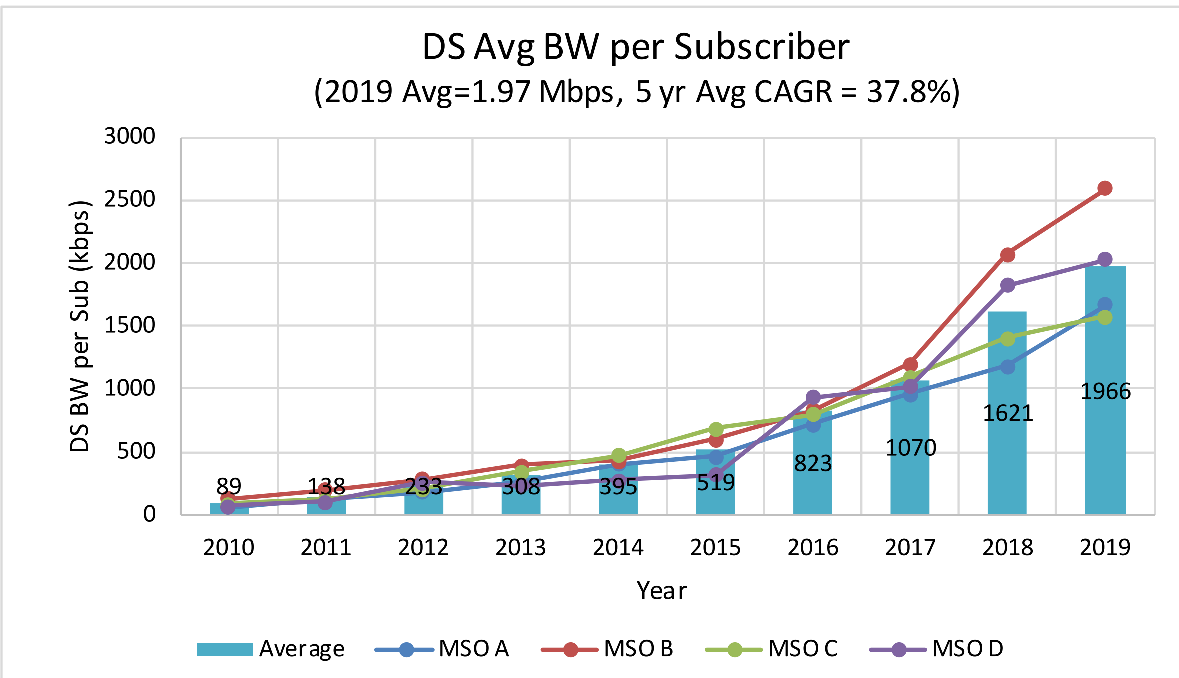

Figure 2 – Tavg, Average Subscriber Downstream Consumption

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 4-5 years, this group of MSOs had an average downstream traffic growth that had been just under 40%. On a yearly basis, traffic growth can be very sporadic. It is not uncommon to see high growth in one year followed by little growth the next. This is exactly what happened in 2018 and 2019. The 35%-40% trend could be used as a longer term guideline on downstream traffic consumption, but others feel the growth rate may continue to decline over time. DS Tavg could reach ~20 Mbps in roughly 7 to 8 years assuming growth rate remains constant, but it might only be ~15 Mpbs in 10 years if growth rate trends continue to decline.

Interestingly, the upstream traffic is growing at a significantly slower rate as shown in Figure 3. During the same ten year period, the upstream Tavg generlly grew at less than 20% compound annual growth rate (CAGR). Notice that there was a significant spike in 2018 followed by a dip in 2019, but the resulting long term CAGR still stayed under 20%.

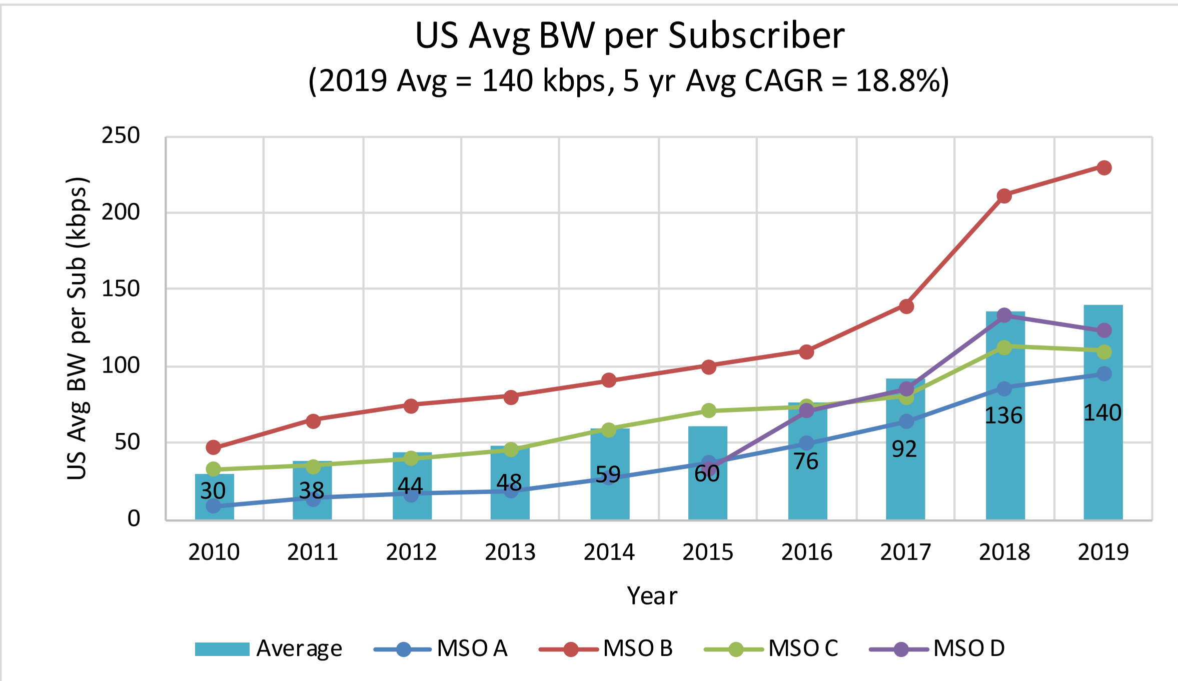

Figure 3 – Tavg, Average Subscriber Uptream Consumption

Traffic has also become more asymmetric with video applications driving downstream consumption [EMMEN_2014]. The DS:US ratios are shown in Figure 4. As of 2019, the average DS:US ratio is about 14:1, the MSO with the largest DS:US ratio seems to have stabilized around a 18:1 ratio.

The 10G platform will be opening up a world of new applications. Many of these will start driving upstream bandwidth consumption. This may come from a plethora of IOT devices in the home, or maybe inexpensive HD resolution video cameras pushing content to the cloud. For the purposes of this study, we will assume the upstream growth will roughly match the downstream growth over the next 7-8 years. This will push US Tavg up to ~1.5 Mbps in that timeframe.

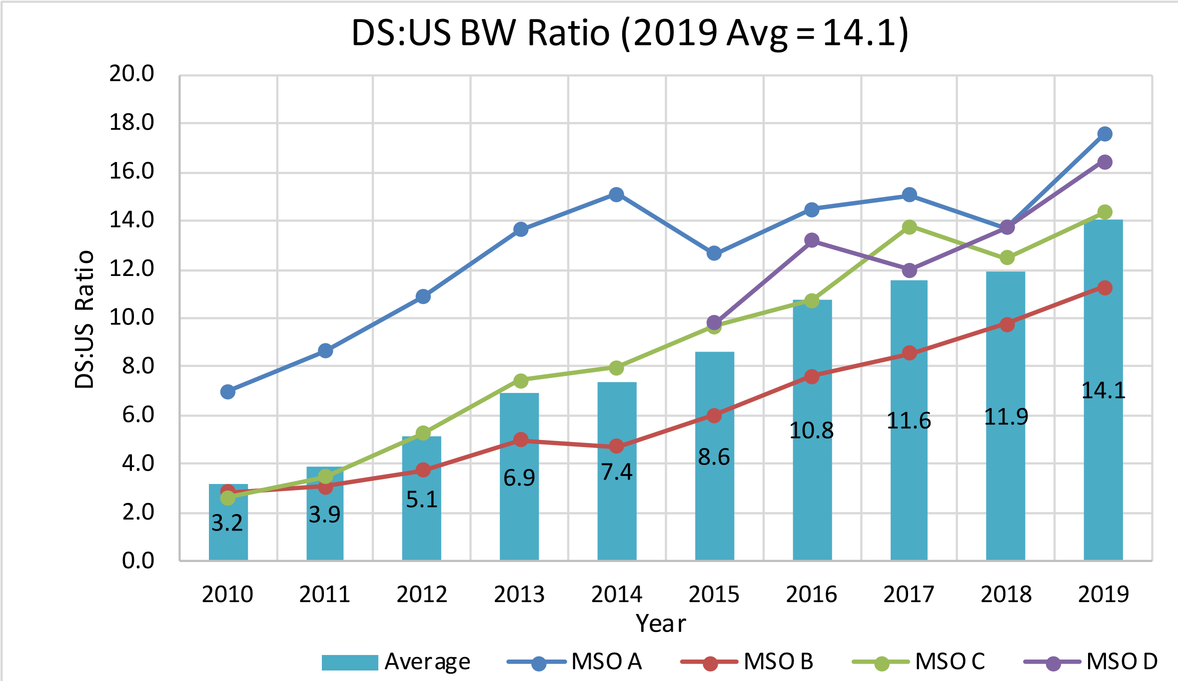

Figure 4 – Tavg, Downstream:Upstream Ratio

2.1.3. *Max Service Tiers in 10G Era*

So, what kind of service tiers will subscribers enjoy in this new high 10G bandwidth era? The Tmax value from the basic formula helps define the Service Level Agreement (SLA) that the MSO can offer to their customers.

As previously discussed, the 10G PON provides a net downstream capacity of ~8.5 Gbps to the consumer. This capacity might support a downstream SLA of 8 Gbps. The service group (SG) utilization (i.e. Nsub * Tavg) for a 64 subscriber PON might grow to a bit over 1 Gbps in the 7-8 year window. That means a consumer with a 8 Gbps SLA will have a QoE coefficient of K=0.9 to 1.0 which is reasonable for this relatively small SG size. The next section will discuss what is needed on an HFC system to support this same DS SLA.

Getting to a true 10 Gbps downstream SLA that is equivalent to 10G Ethernet will mean providing slightly greater than 10 Gbps network capacity. This will push the PON networks into next generation PON technology (e.g. 20+ Gbps). Because HFC can incrementally add capacity with additional spectrum, there are certain downstream scenarios that will be discussed where existing 1218 MHz HFC might be able to hit this target. In general, future technologies such as 1.8 and 3.0 GHz HFC plants, are out of scope for this paper and are discussed further in [CLO_2019].

Choosing the upstream SLA is a more complicated matter. As can be seen with the DS:US consumption ratio, there might be a 20:1 difference between the two. However, in the new 10G era, there may be a need for gigabit US SLA tiers with high burst rates, even if the US consumption might be much lower than downstream.

Looking at PON systems, they offer both symmetric and asymmetric data rates. GPON provides 2.5 Gbps downstream data rates with 1.2 Gbps upstream data rates for a 2:1 ratio. The IEEE 10G EPON downstream might be paired with either a 1G or 10G upstream for 10:1 or 1:1 ratios. In the ITU world, XG-PON pairs 10 Gbps downstream with 2.5 Gbps upstream (i.e. 4:1 ratio) while XGS-PON provides a symmetric 10 Gbps in both directions for 1:1 ratio.

HFC systems have traditionally been extremely asymmetric, but these trends are changing. In the following sections, a range of upstream SLAs are considered to pair with the 8 Gbps DS SLA with a discussion on the technology trade-offs needed for each.

2.2. Spectrum Planning – What to Do with the Legacy Video Spectrum and Other Questions?

In order to achieve 10G goals on HFC systems, it will require MSOs to get on a path to converting most or all of its legacy video spectrum over to DOCSIS high speed data (HSD) spectrum. Most MSOs have already taken the step of removing analog video channels from their plants. That is the first big step and biggest bang for the buck. Now is the time to start transitioning the rest of the digital video spectrum and there are a number of possible options.

The long term goal is to get to a true IP video infrastructure where every home has IP set-top boxes (STB) or IP capable devices for playing video. This provides a common infrastructure across all access technologies including PON, DOCSIS and wireless. However, the economic realities will mean that legacy STBs will be around for many years to come and need consideration.

For some MSOs with pre-dominantly MPEG-4 capable set-top boxes (STB), they can convert their MPEG-2 digital video content to MPEG-4 and roughly cut their legacy video spectrum usage in half. Most of that gain is in the HD content, so having all HD customers with MPEG-4 capable STBs is important.

If an MSO still has a significant amount of MPEG-2 only STBs, then they could relegate these to basic only video tiers and encourage customers who want advanced video tiers to migrate to newer technology (e.g. IP STB).

It turns out that an older technology, switched digital video (SDV), is having a bit of a renaissance and provides some very powerful capabilities to reclaim legacy video spectrum. This is detailed in [ULM_2018]. Cloud-based SDV can be used to save as much as 400 MHz of spectrum. These savings will be critical as cable operators look to add multiple additional 192 MHz OFDM channels downstream and/or expand the upstream splits to higher thresholds. Pairing Cloud-based SDV with IP video migration can be very powerful.

2.3. Network Capacity Modeling for 10G Downstream

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 a number of consequences including the network 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.

Below are some network capacity results from the CommScope network capacity modeling tools. This shows the potential capabilities for a 1218/204 MHz HFC plant. It begins with a 512 homes passed (HP) service group with 256 subs (i.e. 50% penetration). The max DS service tier starts at 1 Gbps and grows by 1 Gbps per year starting in 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 two to three years.

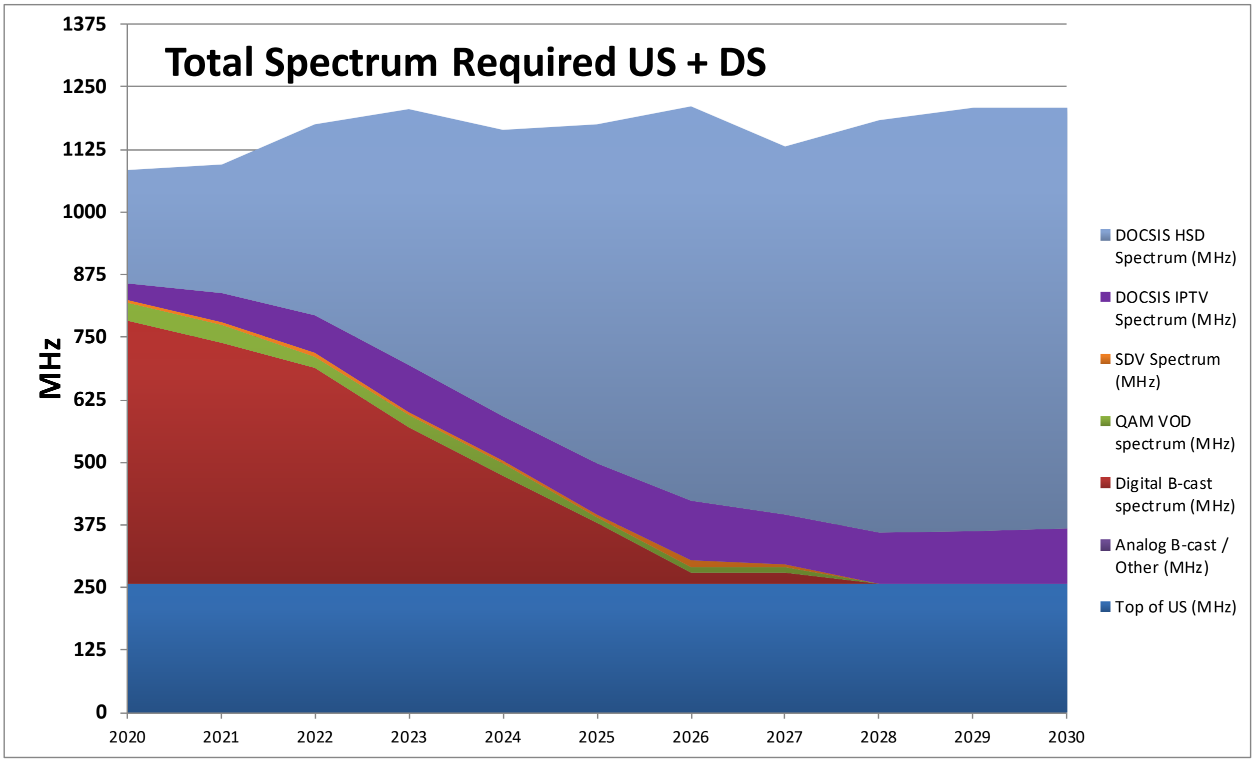

Figure 5 – 1218/204 MHz System – Spectrum Utilization

Figure 5 shows the spectrum utilization for the 1218/204 MHz plant. With the high-split, the downstream spectrum starts at 258 MHz. For this case study, a combination of IP video and SDV are used to reduce the legacy video spectrum requirements over a 5-6 year window. The figure shows how digital broadcast spectrum drops much faster then the corresponding growth in IP video. This spectrum savings allows the DOCSIS spectrum to grow with increasing Tavg and Tmax for a couple years without the need for a SG split.

By the year 2024, the 256 sub SG has maximized its downstream capacity. A larger SG will need a SG split by this time. However, a SG with 192 subs or less can survive for another couple years. Finally by 2027, all SGs need to be at 128 subs or less. The max subs per SG is shown in Figure 6. 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 7. It shows the amount of capacity needed for both DOCSIS 3.0 and DOCSIS 3.1. Each of these is then further broken out into separate components:

- QoE Delta (i.e. K*Tmax), Nsub*Tavg and IP Video

As can be seen, the Tmax component dominates over time.

The lower red line in Figure 7 shows the amount of DOCSIS 3.0 capacity in the system while the upper red line shows the combined 3.0 + 3.1 total capacity for the system.

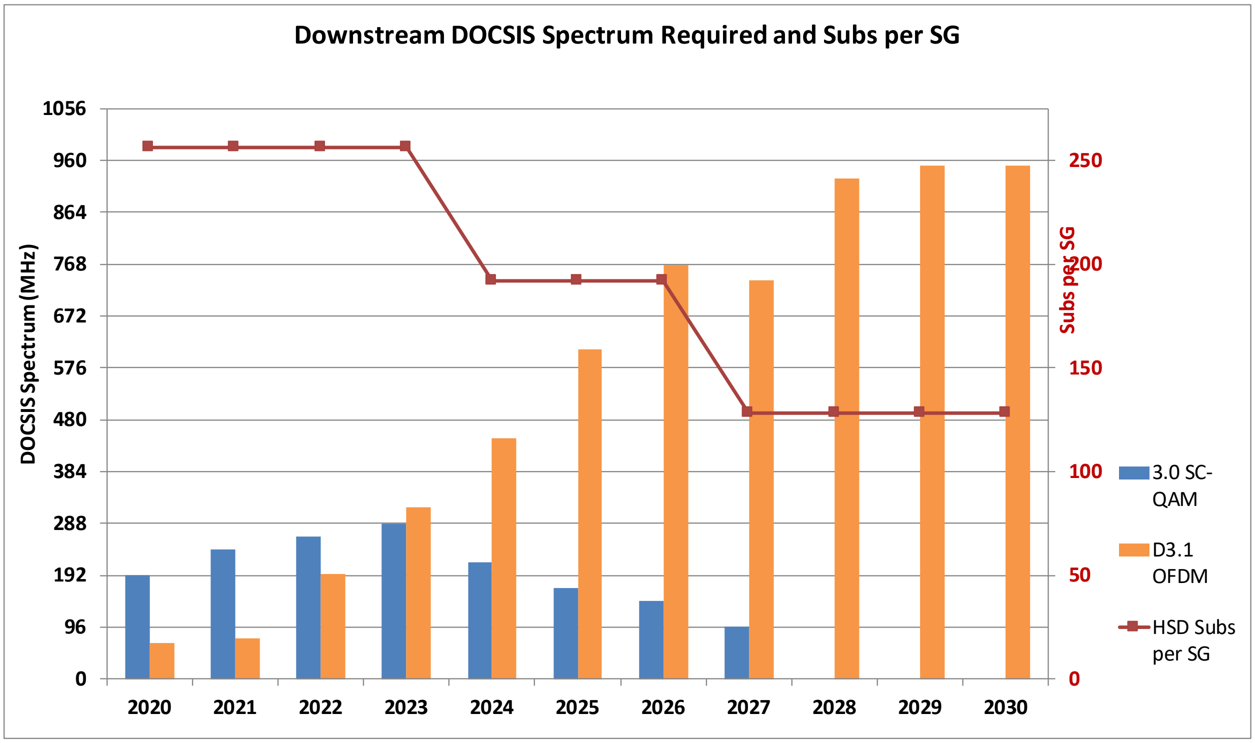

Figure 6 – 1218/204 MHz System – Subs per SG, DOCSIS Spectrum Needs

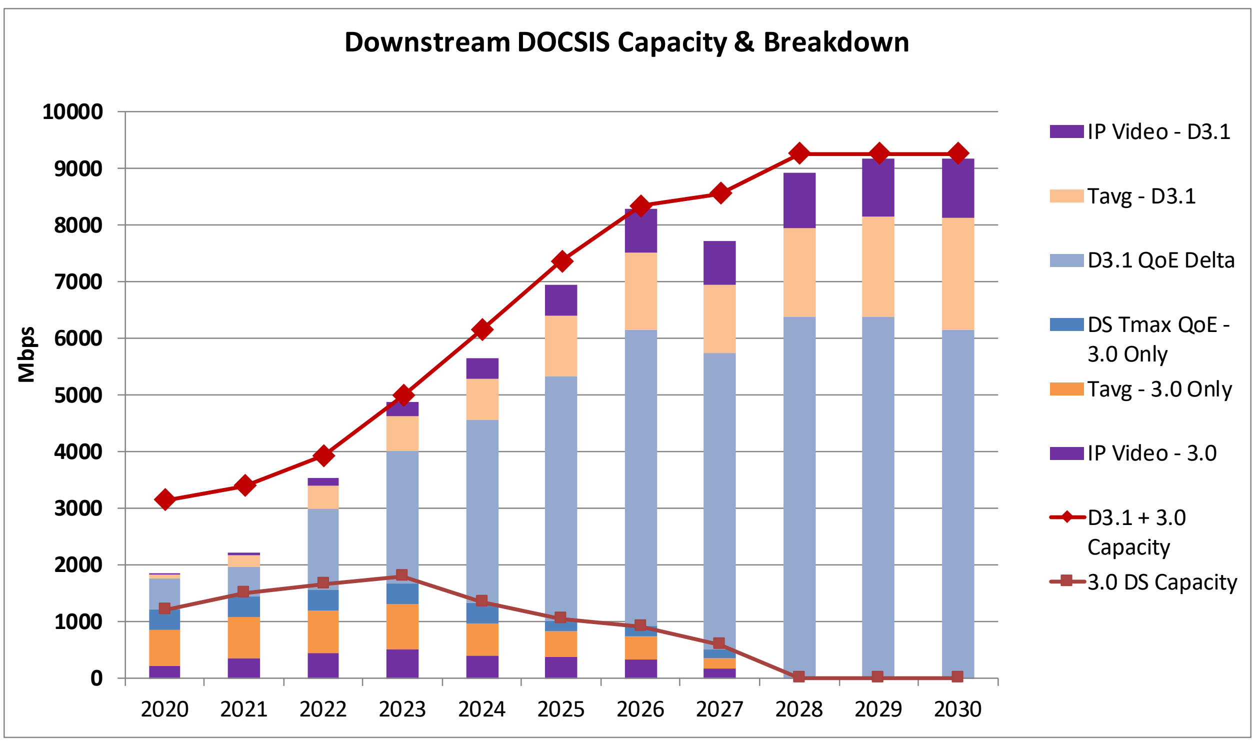

Figure 7 – 1218/204 MHz System – DOCSIS Usage: Tmax, Tavg, IP Video

Perhaps the key point of this case study is that a 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.

2.4. Selective Subscriber Migration Strategies

At first glance, Nielsen's Law and 10G are a scary proposition such that HFC networks might be obsolete in five to seven years while it may take decades to build out an FTTP infrastructure. However, this is not the full story. As was shown in [ULM_2016, ULM_2014], Nielsen's Law only applies to the top speed tiers which is a very small percentage of the entire subscriber base, perhaps less than 1%. So, the key question then becomes, "What happens to the vast majority of subscribers on HFC who are not in the top speed tiers (a.k.a. billboard tiers) and when?"

The [ULM_2014] case study looked at service tier evolution at a few cable operators. It turns out that for many service providers, their mix of service tiers follows a distribution similar to the technology adoption life cycle that was pioneered by Everett Rogers' *Diffusion of Innovations* theory in the early 2000's. This is where the term 'Early Adopters' was coined. See Figure 8.

Looking at a typical cable operator today, they may have a top billboard tier of 1 Gbps that is being taken by a small percentage of innovators. A slightly larger group of early adopters may be taking a 400 Mbps performance tier while the majority of subscribers are taking a 50-100 Mbps basic or flagship tier. Finally, there may be the technology laggards who have older modems (e.g. 2.0, early 3.0) that still have 10-20 Mbps economy tiers.

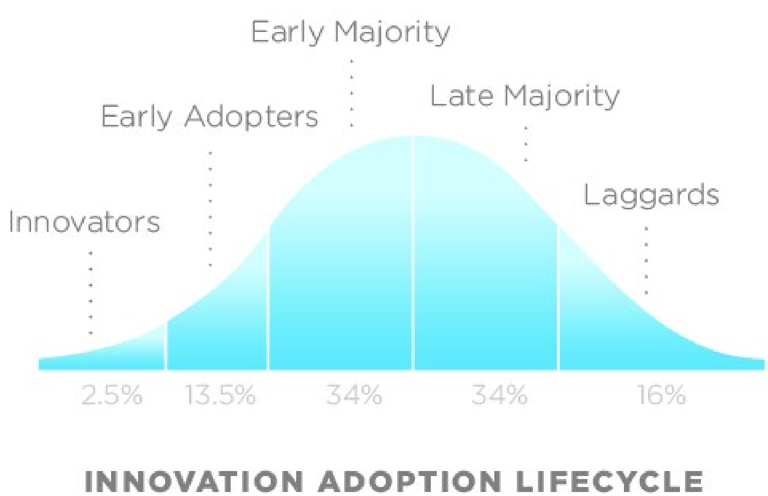

Figure 8 – Innovation Adoption Lifecycle

Perhaps the key finding from the [ULM_2014] study is that the different service tiers are growing at different rates. While the top billboard tier continues to follow Nielsen's Law 50%, each subsequent lower speed tier is growing at a slower rate. Hence, the lower the service tier rate, the lower its CAGR.

Figure 9 maps out a possible scenario from the case study of the various service tier growth over the next two decades. While less than 1% of subs in the top billboard tier hit 10 Gbps in ~2024, the 15% of subs in the performance tier doesn't hit that mark until ~2032. Notice that 85% of subscribers in the flagship basic tier and economy tier stay below this mark for several decades.

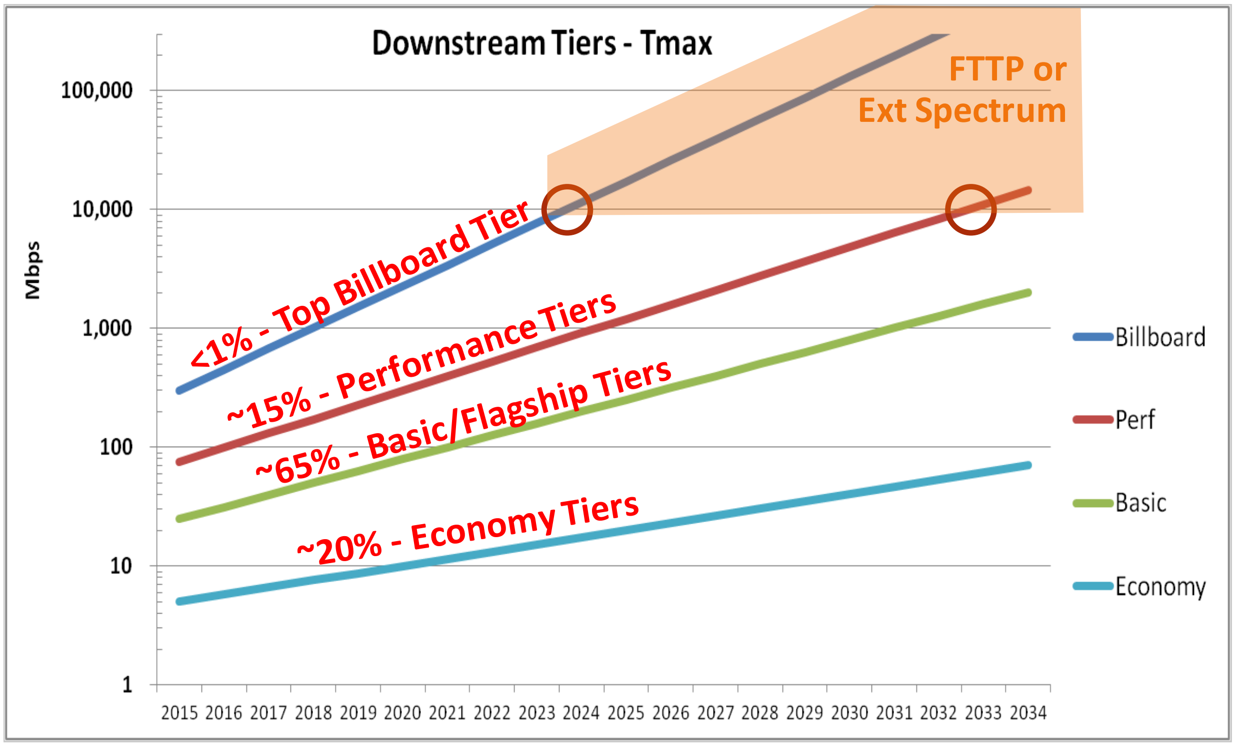

Figure 9 – Downstream Growth with Multiple Service Tiers

With a selective subscriber migration strategy, only the heavy users and subscribers on the top tiers need to get migrated to the new technology. For example, a leading customer (i.e. Innovator) wanting 10 Gbps DS SLA might be given a next generation PON or DOCSIS extended spectrum network connection. Over several years time, early adopters will start to join in and more given the new technology. The vast majority of subs remain on the existing technology.

It is important to note that 99% of the subscribers can still comfortably use today's DOCSIS technology on HFC a decade from now. With a selective subscriber migration strategy, it is very important from a traffic engineering perspective to understand the behavior of the individual service tiers. But with this understanding in hand, selective subscriber migration can be used to extend the life of HFC for decades to come. When the time comes to migrate the premium tiers from traditional HFC to new technologies, it might be done with either FTTP technology or maybe extended spectrum HFC if it is viable by that time.

This is a crucial concept for our 10G migration strategy. Initially, 10G will only be for a fraction of a percent representing the innovators. This stage will then be followed by the early adopters that might represent 5% to 15% of the subscribers. Then finally 10G will be broadly consumed by the entire customer base. The 10G migration strategy can take advantage of this shifting over time which might take a decade.

3. Outside HFC Plant Considerations and Logistics

The cable industry has invested much into newer technology such as DOCSIS 3.1 and 1218 MHz capable components. How well will this support the 10G goals? It turns out that there are many variables involved that will be discussed here.

3.1. Technology Options to Enhance HFC Bandwidth Capacity

The cable industry has many tools in their toolbox, each with its own strengths, advantages, and costs. A summary of possible options is shown in Table 1.

Table 1 – Technology Options to Enhance Bandwidth Capacity

Typical Options for last few years:	**Ease & Cost**	**Tmax**	**Nsub*Tavg**	**Notes**
Analog Reclamation	☺	☺☺☺ (DS only)	☺	Need DTAs
Node Segmentation, node splits	☺ to ☹		☺☺☺	HFC plant upgrades
DOCSIS 3.1	☺☺	☺☺		Need D3.1 CPE
More Options for coming years:				
Fiber deep – N+0, N+1; 1.2GHz DS,85-204MHz US	☹	☺☺☺☺☺	☺☺☺☺☺	Long term strategic direction, often combined with DAA
MPEG-4 transition	☺☺	☺☺ (DS only)		If mostly MPEG-4 capable STB
IPTV transition	☺	☺ -> ☺☺☺ (DS only)	☺ (DS only)	Need IP capable CPE; gains depend on IPTV strategies
Cloud-based SDV	☺☺☺	☺☺☺ (DS only)	☺ (DS only)	Leverage existing legacy STB
Future Options:				
DOCSIS FDX	☹☹	☺☺☺ (US only)		Assumes DAA & N+0, too
FTTH, Ext Spectrum, FTTTap	☹☹☹	☺☺☺☺☺	☺☺☺☺☺	Next steps after fiber deep

Over the past decade or so, most operators have reclaimed analog spectrum and continued to do node segmentation and node splits as business as usual. It turns out that each addresses a different component in our traffic engineering formula. The analog reclamation frees up spectrum which is critical to offering higher service tier SLA (i.e., Tmax) while node segmentation and splits reduce SG size which directly addresses the Nsub*Tavg component in our traffic engineering formula.

More recently, operators have started using DOCSIS 3.1 as they migrate to 1 Gbps downstream service tiers. D3.1 operates in today's existing HFC plants without any changes. It improves spectral efficiency (i.e., bps/Hz) enabling more capacity from a given amount of spectrum. D3.1 is also robust enough to operate in the roll-off region to provide additional bonus capacity. It is straightforward to implement without plant updates by deploying D3.1 cable modems and some SW/HW upgrades to the existing CMTS/CCAP platforms.

Over time, the D3.1 benefits will continue to grow as the plant is improved. The spectral efficiency improves significantly with fiber deep networks and it enables operation over wider frequencies (e.g., 1218 MHz downstream, 85/204 MHz upstream).

From a strategic perspective, all operators plan to push fiber deeper until it eventually becomes fiber to the home (FTTH). However, this is a multi-decade process. Some operators are considering a fiber deep HFC (with N+0, N+1 cascades) as the next major step along the way. But this can be a monumental task with pulling fiber deeper in the plant and increasing nodes counts by a factor of 10X to 20X. So, while this option gets top marks for increasing spectrum AND significantly reducing SG size for both upstream and downstream, its major costs and complexities will force this option to be done slowly over time.

If an operator is focused on more easily increasing DOCSIS downstream bandwidth capacity, there are several other options available in the near-term including:

- Migrating broadcast video to MPEG-4/H.264 on existing legacy STBs
- Migrating broadcast video to IP video delivery over IP STBs
- Migrating broadcast video to Cloud-based SDV over existing legacy STBs

The MPEG-4/H.264 option makes sense if all or the clear majority of the legacy STBs support MPEG-4 decoding. Overall, it provides a 2:1 gain in broadcast spectrum converted. Problems arise if there is still a substantial MPEG-2 only STB user base. That may limit the broadcast programs that can be converted (e.g., limited program tiers) or require those older STBs to be replaced.

The IP video migration includes a wide array of potential solutions. Some operators may only move video on demand (VOD) and select programming content to IP video while other operators may move aggressively to IP video everywhere. The potential bandwidth capacity gains from reducing legacy broadcast video can vary dramatically as well. In addition to an all new video infrastructure, the operator also needs to replace its CPE with IP video capable boxes. During the transition window, the operator must continue to support the legacy STBs, creating a bubble in their bandwidth needs.

The third option for reducing legacy broadcast spectrum is SDV. In the early days with large SGs and limited number of SDV EQAM modulators, the spectrum gains were somewhat limited (e.g., 84 MHz). However, with today's current technologies, Cloud-based SDV can be used more aggressively and achieve gains in excess of 400 MHz. A SDV case study is described in[ULM_2018]. SDV also has the advantage over the other approaches in that it works on existing legacy STBs (MPEG-2 and/or MPEG-4) and requires minimal video infrastructure.

The final options for increased bandwidth capacity are looking much further into the future. Operators will be considering options such as FDX DOCSIS to enhance the available upstream spectrum and network architectures such as FTTH, fiber to the tap (FTTT) and DOCSIS extended spectrum (e.g., 1.8 and 3.0 GHz).

3.2. Cleaning Up the Plant to Maximize bps/Hz for D3.1

HFC networks have always been an evolving and changing infrastructure, repeatedly delivering bandwidth capacity increases to accommodate the needs of their various services (video, high speed data, and voice) in a "just-in-time" fashion. MSOs have long recognized that the HFC plant contains vast quantities of un-tapped bandwidth capacity that can usually be enabled in a gradual fashion using minor evolutionary transitions to various sub-systems within the plant. This low-cost evolutionary approach to network transition has long been preferred over more expensive revolutionary changes that attempt to change a large amount of the HFC plant equipment all at once.

The HFC plant still has an incredible amount of un-tapped bandwidth capacity that can be enabled in a very smooth and cost-effective fashion in the coming years without causing any unnecessary and premature plant augmentations. Let's take a look at how existing technology can be exploited to meet some or all of the 10G goals.

DOCSIS 3.1 brings many benefits to the table, with vastly improved spectral efficiency being a key advantage over DOCSIS 3.0 technologies. A target configuration to support the max 10G downstream SLA of 8 Gbps would be 4x192 MHz OFDM channels bonded with 32 SC-QAM channels. This consumes 960 MHz of downstream spectrum with total capcity of ~8.6 Gbps at 4096-QAM modulation. The HFC system now provides basically the same total downstream capacity as 10G PON while providing backwards compatibility for DOCSIS 3.0 modems that will be around for a while.

The total amount of downstream capacity from the OFDM channels might vary by as much as 20%-30% depending on the communication channel quality as originally shown in [ALB_2014]. There are several techniques that the operator can use to improve OFDM capacity. First, they can replace long analog fiber links (e.g. 40+ km) with digital optics of a Distributed Access Architecture (DAA) system. The 2nd item is to push fiber deeper and to shorten the cascade length (e.g. N+6/N+3 down to N+1/N+0).

If an operator has a 1218/85MHz HFC plant, they have enough spectrum to support the 8 Gbps DS SLA, provided they can squeeze the legacy video spectrum down to 150 MHz over time. This is very feasible using the techniques described earlier. With an 85 MHz upstream split, the operator may still be able to pair the 8 Gbps DS tier with a 400-500 Mbps US tier, which provides a reasonable 15:1 to 20:1 DS:US ratio.

However, some operators may want to support a 1 to 1.5 Gbps US tier to go with the 8 Gbps DS tier. At this point, the operator will need to consider a 204 MHz upstream split. Note that there is still 960 MHz of downstream spectrum available, but this requires that the legacy video spectrum has been migrated to all-IP video delivery over the DOCSIS channels.

3.3. Can We Achieve Cable 10G with Existing 1 GHz Taps?

As mentioned above, cable has been very successful at growing incrementally as capacity is needed. As the HFC system electronics are pushed to 1218 MHz, a question arises to what are the capabilities of the passives in the networks? It turns out that most HFC plants today have had 1 GHz taps installed over the last couple decades. This represents an extremely large installed base. Can we achieve the Cable 10G goals with these 1 GHz taps? That would make this transition much more cost effective.

The ARRIS/CommScope team analyzed a number of the most common taps in our research labs to get a better understanding of how they behave above their specified 1 GHz limit. Figure 10 shows the behavior for one of the 'middle of the road' taps.

As can be seen, it can support 12 bps/Hz (i.e. 4096-QAM modulation) up to and slightly beyond 1000 MHz. It dips slightly to 11 bps/Hz (i.e. 2048-QAM modulation) for another ~100 MHz. The roll-off starts to accelerate above 1200 MHz, but it can still support 9 bps/Hz (i.e. 512-QAM modulation). The bottom line for all the taps is that the 192 MHz OFDM channel above 1000 MHz had anywhere from 75% to 90% of the capacity of OFDM channels below 1000 MHz.

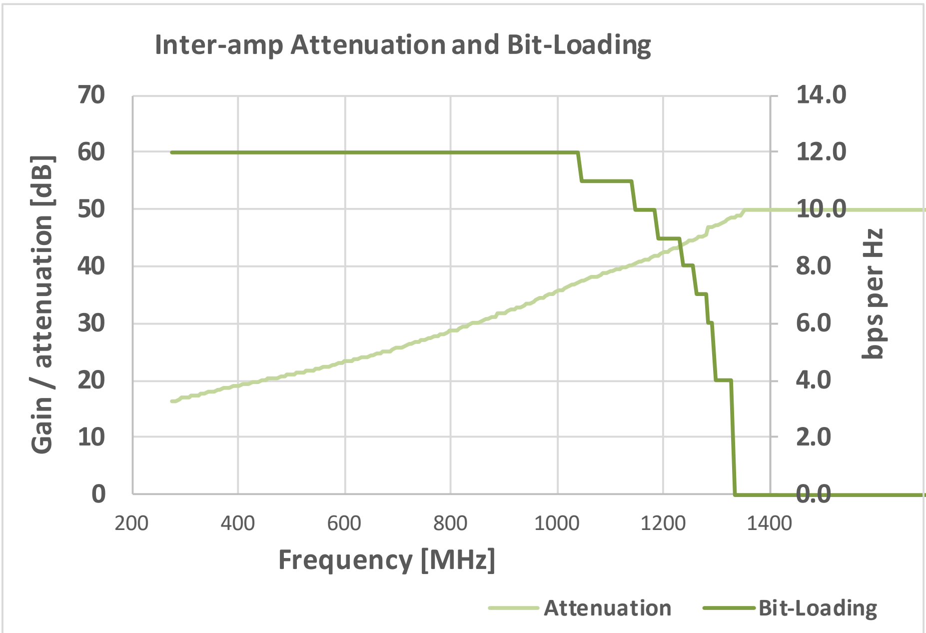

Figure 10 – Sample Tap: Inter-amp Attenuation and Bit-Loading

In looking at the impact of various 1 GHz taps, the team also analyzed how tap capacity varied as the cascade length increases. As can be seen in Figure 11, all five tap types were able to achieve 8.5 Gbps capacity for a N+0 plant. Introducing a single amplifier in a N+1 system causes a slight degradation but overall capacity is still above 8 Gbps.

As the cascade length is increased to N+3 and N+6, the capacity degradation continues. The N+3 capacity is roughly around 8 Gbps while N+6 drops even further close to 7 to 8 Gbps. In both cases, there is a ~500 Mbps difference between the best and worst tap.

All of these capacities were calculated for an amp spacing of 1050 feet (e.g. 6 taps at 175' each). The team then analyzed the impact of closer amps spacings of 700' and 875' (e.g. 4 or 5 taps at 175' each). This is shown in Figure 12.

Note that the capacity for 700' and 875' spacing is very similar. This is due to cable attenuation on shorter cables not being sufficient to pull the signal close to the thermal noise floor. So, the SNR of the node (plus distortion added by the amplifiers) dominates and determines the received SNR. Therefore, a 100' cable would give similar performance to a 500' cable.

For longer cable spacing, the attenuation is sufficient that the EOL (End of Line) SNR is influenced by the thermal noise floor. One you reach this point, any further increase in length reduces throughput. So there is a soft threshold effect in operation. In this example, the longer spacing of 1050' (e.g. six taps @ 175') results in about a 10% drop in capacity.

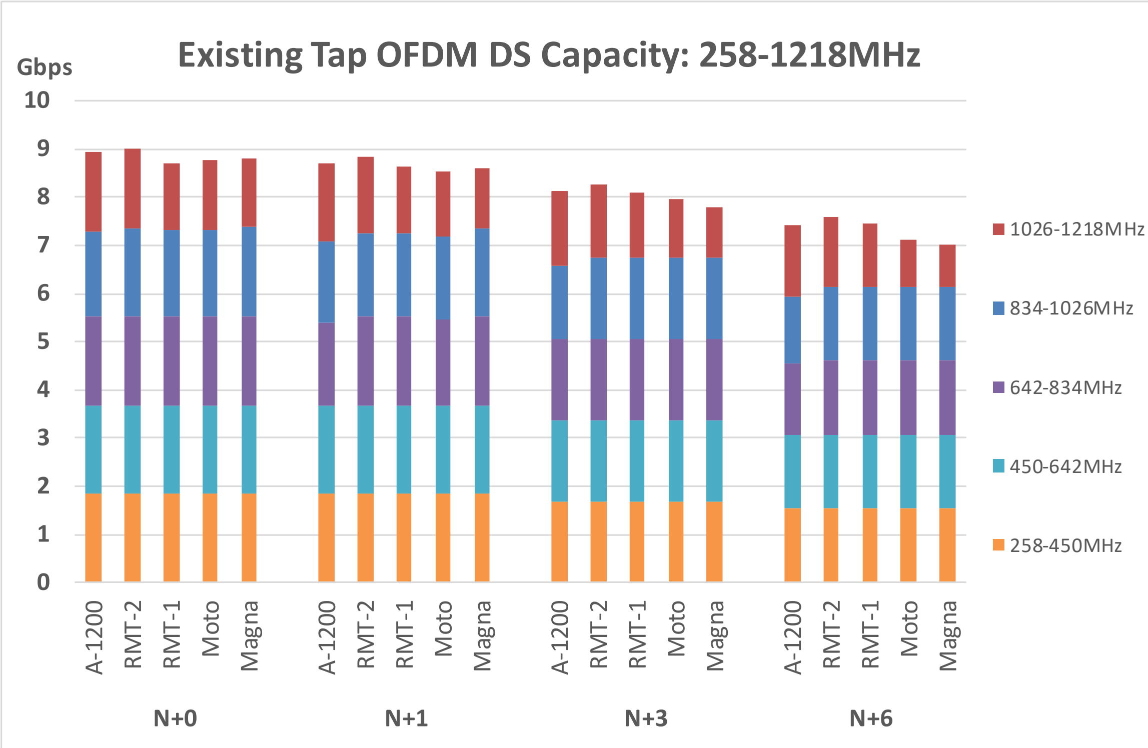

Figure 11 – Tap Capacity for Various HFC Cascade Lengths

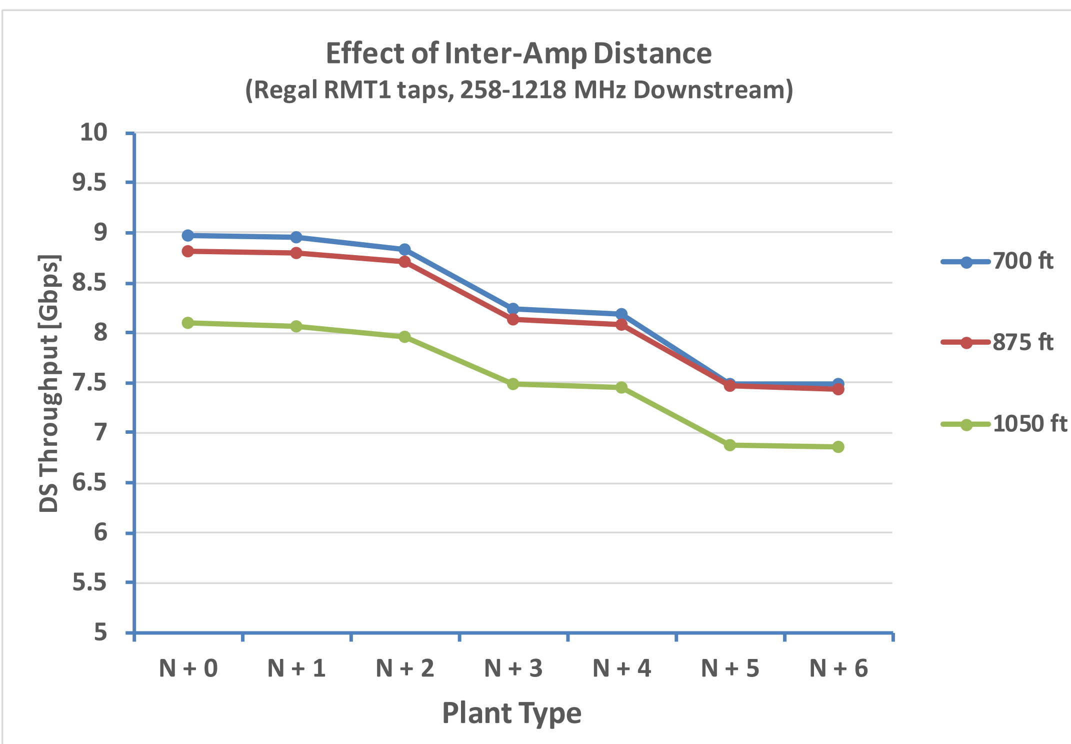

Figure 12 – Impact of Amp Spacings

The conclusion is that operators can still get decent performance up to 1218 MHz using the existing 1 GHz taps. However, the capacity of a particular plant is going to vary based a number of key variables including:

- Cascade length (e.g. N+0 to N+6)
- Amplifier spacing
- Optical link (i.e. DAA or Analog Fiber links)
- Tap type

To reach the 10G DS goals, the operator will need to look across all these variables to see how they might reach 8 Gbps DS SLA. Depending on plant variables, it may be necessary to slightly lower customer tiers to 5 or 6 Gbps DS SLA until the plant can be upgraded.

4. Migrating from DS Only 10G to More Symmetrical 10G

To support cable 10G DS SLAs, the operator will need a minimum of an 85 MHz upstream split. This supports a 400-500 Mbps US SLA and a DS:US ratios of ~15:1 to ~20:1. Becoming more symmetrical and supporting Gbps US SLAs will require much more upstream capacity and newer technologies.

4.1. 204 MHz Frequency Division Duplex (FDD) High-Splits

The original DOCSIS 3.1 specification supports a 204 MHz Frequency Division Duplex (FDD) upstream split, resulting in a usable spectral range that is up to five times larger then current 42 to 65 MHz widths. The 204 MHz split supports two 96 MHz OFDMA channels that can enable 1 to 1.5 Gbps US tiers. This gets the 8 Gbps DS SLA to a DS:US ratios of about 5:1.

This approach continues to use Frequency Division Duplex technologies that separate upstream spectrum and downstream spectrum (with a guard-band in between). The change to 204 MHz FDD high-split operation typically requires changes to both the existing nodes and the existing amplifiers on the HFC plant. Sometimes, plug-in modules for new diplex filters can be utilized to offer this upgrade path. The diplex filters typically permit the downstream spectrum to pass signals beginning at 258 MHz, so the resulting guard-band creates a diplex filter “spectral penalty” of 54 MHz (i.e.54 MHz of spectrum is unusable from 204-258 MHz). The downstream spectrum has also been reduced by 150 MHz, which is close to ~1.5 Gbps of downstream OFDM capacity.

One of the potential side issues that must be dealt with when 204 MHz FDD high-split is utilized is the passing of downstream out-of-band (OOB) signals to existing set-top boxes. To support this OOB capability wherever it is required, vendors are exploring options that would permit these downstream OOB signals to be passed through the high-split nodes and amplifier in the HFC network (even though it is in the upstream portion of the spectrum).

4.2. Full-Duplex DOCSIS (FDX)

Getting beyond a 1 to 1.5 Gbps US tier will require additional technologies beyond the 204 MHz FDD upstream. Some recent work at CableLabs has focused on a new technology called Full Duplex DOCSIS (FDX) and recently moved into the DOCSIS 4.0 domain. FDX leverages echo canceller technology to allow simultaneous upstream and downstream operation in the FDX band. FDX is targeted at a fiber deep Node+0 DAA environment.

The FDX capability offers a fundamental benefit that permits upstream spectrum expansions to occur without causing reductions in downstream spectrum. One of the key FDX technology enablers, Echo Cancellation, is required in both the optical node and potentially in the cable modem (CM).

Echo cancelling is a well-known technology. FDX was originally an addendum to the DOCSIS 3.1 specification but has since become part of the new DOCSIS 4.0 specification [FDX_PHY]. The technology proposes to have downstream and upstream D3.1 transmissions occurring in the same frequency band at the same time at the CMTS. In the FDX specification, the overlapping frequency bands can be in any of the following ranges starting at 108 MHz and ending at: 204 MHz, 300 MHz, 396 MHz, 492 MHz, 588 MHz, or 684 MHz as shown in Figure 13. These FDX bands shown in Figure 13 are in addition to the standard 5-85 MHz upstream that can be utilized as well.

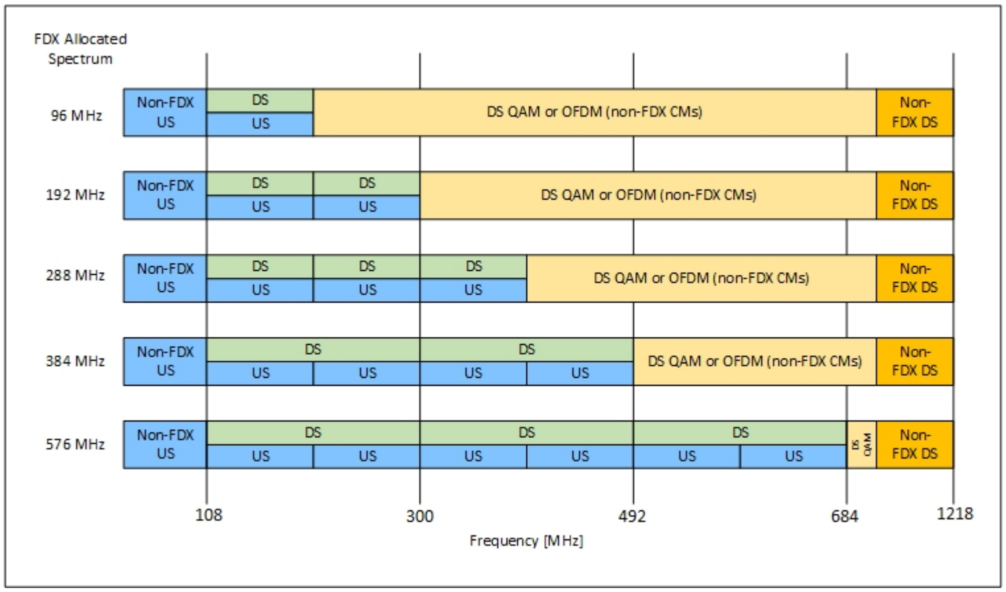

Figure 13 – Full-Duplex DOCSIS (FDX) Spectrum Band Options

On a fiber deep Node+0 plant, the upstream OFDMA channel might net capacity as much as 10 Mbps per MHz. This means that a 108-300 MHz FDX system might support a 2 Gbps US SLA while the full spectrum 108-684 MHz FDX system might support a 5-6 Gbps US SLA. Using the full FDX band would enable the operator to offer such DS/US service tiers as 8 Gbps x 2.5 Gbps or a fully symmetric 5 Gbps x 5 Gbps SLA.

With the Node+0 architecture, the 108-1218 MHz of downstream spectrum might actually net over 10 Gbps of downstream capacity which means the 1218 MHz FDX system could be pushed to a true 10 Gbps DS SLA with 5 Gbps US SLA. The downstream is now on par with 10G Ethernet while supporting a 2:1 DS:US ratio.

Current FDX work is moving along well. Initial field trials were in 2018 and continued in 2019 with very promising results. Real-world deployments are targeted to take place in 2020.

FDX should work fine in Node+0 environments, but its ability to perform in Node+X environments is still under study. A number of operators are reluctant to jump to N+0 but are still interested in achieving more symmetrical upstream service tiers. The issues with FDX in this environment and possible alternatives are explored in [ALB_2019]. Because of the FDX challenges in N+X HFC plant, other technologies are under consideration for those scenarios.

4.3. Frequency Division Duplex (FDD) – 300 or 396 MHz Upstream Splits

Perhaps the simplest way to increase upstream capacity is to just raise the upstream split even higher then the 204 MHz defined in the current DOCSIS 3.1 specification. Pushing it up to 300 or 396 MHz will add one or two more 96 MHz OFDMA channels respectively. Each OFDMA channel will add almost 1 Gbps to the upstream capacity. This would let an operator support 2, 2.5 or 3 Gbps US SLA.

However all of this comes at the expense of downstream spectrum. As the upstream split is pushed higher, the guard band region also grows larger. So every increase of 96 MHz to upstream spectrum results in the loss of ~120 MHz of downstream spectrum, which is why this option is often paired with an extended spectrum downstream (e.g. 1.8+ GHz). But what can be managed from a 1218 MHz system? Figure 14 shows the impact to the 1218 MHz downstream capacity for various cascade lengths as the upstream split is increased.

If the cascade length is short (e.g. N+0, N+1), then a 300 MHz split might still be able to squeeze out a 7.5 Gbps DS x 2.5 Gbps US SLA. Or maybe an operator would rather have a 396 MHz split to offer a more symmetric 6 Gbps DS x 3 Gbps US SLA.

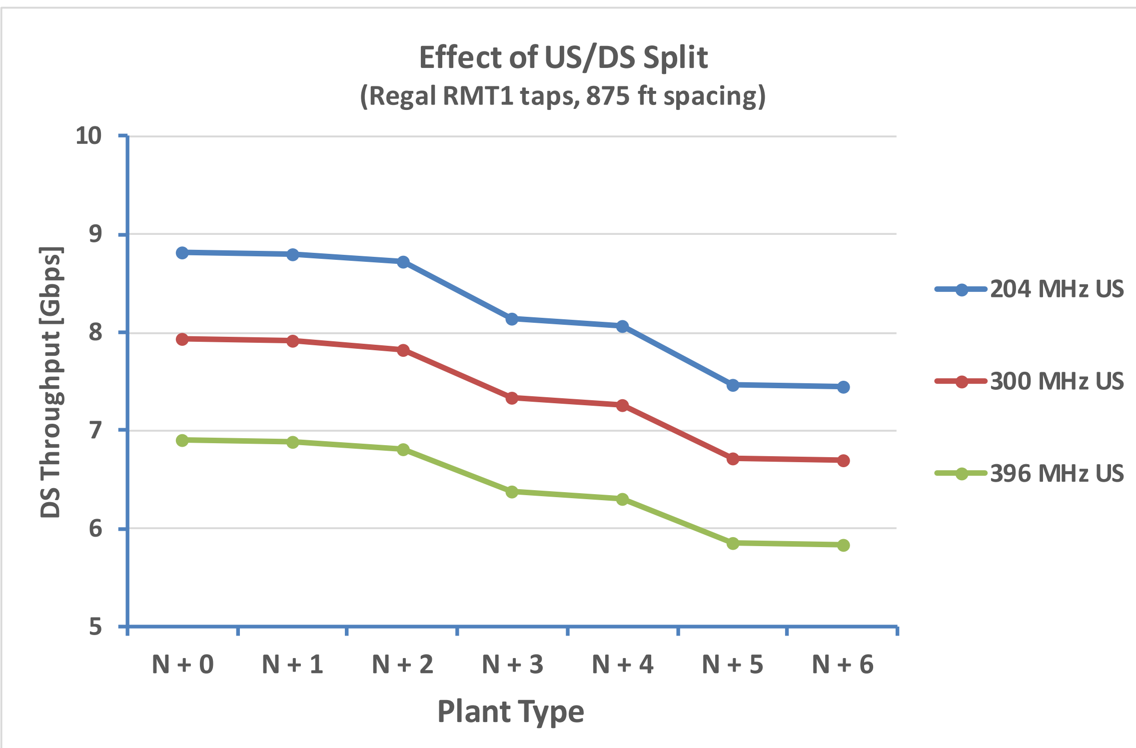

Figure 14 – 1218 MHz DS Capacity Impact from Ultra-High Upstream Splits

With the reduction in downstream capacity due to the higher upstream splits, is there any additional way to squeeze out extra DS capacity? In Figure 10, some 1 GHz taps actually continue to provide bandwidth above 1218 MHz. In the example above, the tap actually can support 4 bps/Hz (i.e. 16-QAM modulation) up around 1300 MHz. Our analysis for all of the taps was extended to include an additional OFDM channel from 1218 to 1410 MHz, which would also require next generation modems too.

A couple of the better taps might provide an additional 1 Gbps of capacity above 1218 MHz in an N+0 plant, but this plant is FDX capable and the extra spectrum may not be needed. Those capacity gains might be cut in half for N+3 and be fairly negligible at N+6. The lesser taps have almost no additional capacity above 1218 MHz, even at N+0. So extra capacity is very dependent on tap type.

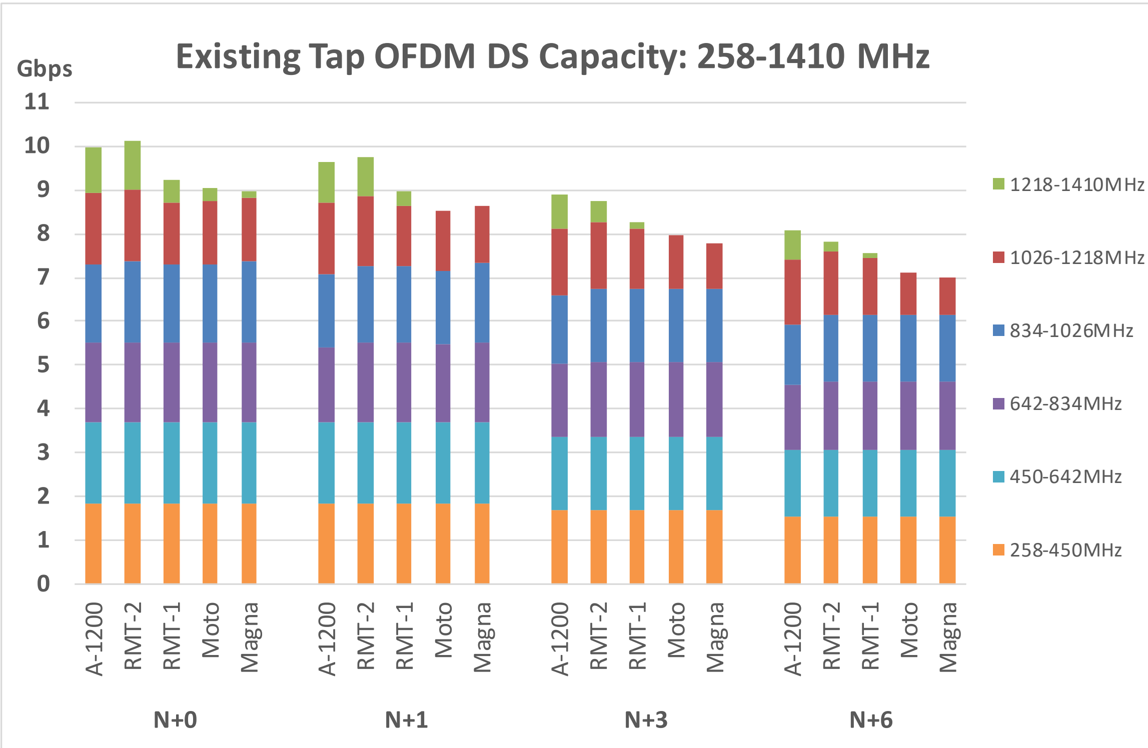

Figure 15 – Tap Capacity to 1410 MHz for Various HFC Cascade Lengths

4.4. Soft-FDX – Static or Dynamic

Although promising research has been conducted at CommScope and other vendors on FDX capable amplifiers in the past year, a problem called Interference Group Elongation has been identified, see [ALB_2019]. This problem causes serious issues with this traditional FDX in a Node+X proposal. In the end, it causes large Interference Groups to be created that span most of the length of each RF Leg on a node. If we still want to share upstream and downstream spectrum, this drives us to time division duplex (TDD) operation. We will discuss two variants of this called static soft-FDX and dynamic soft-FDX operation.

Soft-FDX can be considered a special mode of FDX where the interference group includes an entire RF leg that might span multiple amplifiers. The 'soft' adjective refers to the ability to change the location of the US/DS split easily (potentially via software). Soft-FDX helps in supporting high US speeds, which are occasionally demanded by users, without permanently locking the spectrum to the US which can severely affect the valuable DS spectrum that is used to offer many services including video and high DS speeds which are demanded more frequently than the US.

Soft-FDX can be either static or dynamic. Static soft-FDX refers to the case where the US/DS split location does not change frequently (e.g., on the order of months or years). On the other hand, dynamic soft-FDX refers to the case where the US/DS split location changes in real-time based on traffic demand (on the order of milliseconds or seconds). For instance, in the dynamic soft-FDX case, when there is a need for more US spectrum to run an US speed test or upload as an example, the split changes to accommodate that and when the need for the added US spectrum goes away, the split changes back to reclaim the valuable DS spectrum. Both static and dynamic soft-FDX can be implemented using special assignment of the FDX RBA messages.

From a burst speed perspective, soft-FDX has the same burst capabilities as traditional FDX. So they can support the same SLAs as FDX. However, this is fine as long as both upstream and downstream are not bursting for sustained periods of time. Based on current upstream utilizations, this should not be a problem. Traffic engineering studies are currently underway to determine if soft-FDX provides adequate bandwidth capacity. It may require that nodes, amplifiers and CMs be able to quickly switch the directionality of the frequency band between the upstream and downstream directions (via rapid RBA switching).

In reality, traditional FDX uses TDD concepts within the HFC plant. The creation of Transmission Groups within FDX defines groups of neighboring modems that are "noisy neighbors," and they are required to transmit using TDD approaches. Soft-FDX simply extends the size of a Transmission Group to be an entire RF leg on the node instead of a sub-section of the RF leg.

This soft-FDX approach may permit the use of soft-FDX amplifiers in Node+X systems (X>0), and it may permit sharing of the upstream and downstream frequency bands inside of Node+X systems (X>0).

In conclusion, the traditional FDX Node+X solution does not have a clear path to success in the 2020-2028 time-frame for Node+X (X>0) architectures. On the other hand, if its traffic engineering performance is deemed to be acceptable, soft-FDX is probably a technology that could be implementable by the mid-2020s if bandwidth capacities require it.

4.5. Addressing Cable 10G with Blended Fiber Deep and PON Systems

With all of the previous HFC options (FDX, Soft-FDX), the cable system has become much more symmetric, but still falls short of a true 8 Gbps x 8 Gbps SLA. In addition, questions start to be raised on how to finally get to a true 10G Ethernet equivalent of 10 Gbps x 10 Gbps SLA? Even 10G PONs can not meet that goal.

One option that an operator can consider is to use a Selective Subscriber Migration strategy and start moving to a blended system with HFC fiber deep plus PON system. As discussed earlier, only a small percentage of early adopters may go for these fully symmetric 10G systems. Rather then switching their entire subscriber based to fiber to the home (FTTH), an operator might choose to only provide FTTH to these innovators and early adopters.

To accomplish this on demand (i.e. install FTTH within a few days or weeks of customer requesting that service), the operator will need to have the fiber extremely close to every customer's home. It turns out that a fiber Deep HFC upgrade is an excellent stepping stone for this strategy. Not only does it provide cable 10G for its vast majority of subscribers, but the operator can now deliver beyond 10G via FTTH in a timely and cost effective manner.

5. CPE Considerations

First generation DOCSIS 3.1 modems supported downstream capabilities of 2x192 MHz OFDM bonded with 32 SC-QAM (i.e. 3.0 channels) along with 2x96 OFDMA upstream channels. These D3.1 modems have five times the capacity of 32x8 D3.0 modems. These might provide subscribers with up to a 4 Gbps DS x 1.5 Gbps US SLA.

However, our network capacity modeling results show that it is still very important to upgrade the majority of D3.0 modems to D3.1 over the next decade. This will allow the operator to convert SC-QAM spectrum to more efficient, higher capacity OFDM channels. If an operator chooses to keep a large percentage of their subscriber base on D3.0 modems, then they may lose 1-2 Gbps from their total downstream capacity.

Only customers needing more than 4 Gbps DS SLA will require a new 2nd generation D3.1 modem that has at least four 192 MHz OFDM channels in the downstream. At this point in time, it is likely that the next generation of modems may also be FDX capable with up to 684 MHz upstream support.

One key tenet in all of the proposed upstream extensions above is that every approach can leverage an FDX capable modem. So, ultra-high splits (300/396 MHz) and soft-FDX will all be relying upon and using the same modem technology as traditional FDX. This should help the industry drive modem volumes and economies of scale.

There was a brief discussion regarding use of the roll-off range from 1218 to 1410 MHz. If that path is chosen, it would require a future generation modem with Extended Spectrum capabilities (e.g. 1.8/3.0 GHz) for the innovators and early adopters who need that extra bandwidth burst.

As operators ramp up cable 10G services, the home networking piece will also become extremely important to be able to move up to 10 Gbps throughout the home. Recent advances with Wi-Fi 6 (a.k.a. 802.11ax) should be up to the task. These wireless data rates should match up well with cable 10G targets.

There are also advances with home wireless routers. The newest ones are tri-band where they use one band for dedicated mesh backhaul inside the home. This will become even more powerful as additional spectrum becomes available in the 6 GHz band.

Conclusion

This paper takes a look at how cable operators can reach the 10G goals with 10G PON and 1218 MHz HFC technologies within the next decade. [CLO_2019] takes a longer term look over the next 25 years past 10G. After accounting for all the different overheads (e.g. PHY, MAC, IP layers), the subscriber is actually getting a 8 Gbps SLA in a 10G world. Table 2 summarizes the various options 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 2 – Summary of 10G Access Network Options

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

As can be seen by these options, an operator can choose how symmetric they want their system to be. This will be driven by competitive market forces as well as new yet unknown upstream applications that may appear in the future.

These SLAs will be quite adequate for the vast majority of subscribers over the next decade. There may be a small number of innovators and early adopters that want to go beyond these service tiers latter in the decade, but that can be handled with a Selective Subscriber Migration strategy that moves this small percentage to a next generation PON (e.g. 20+ Gbps) or to an extended spectrum HFC (e.g. 1.8 or 3.0 GHz).

Our traffic engineering and network capacity analysis shows that 1218/204 MHz technology meets the needs through the end of the next decade. While getting to fiber deep N+0 is a good long term strategic goal, a 500 HP node size, N+X system is still reasonable as long as it can be segmented.

If more symmetric upstream services are needed or desired (i.e. greater than 1.5 Gbps), then a migration to traditional FDX for N+0 or Soft-FDX for N+X is a reasonable path. These also restore some downstream spectrum (i.e. 108-258 MHz) that may actually enable a true 10 Gbps DS SLA.

Our investigations into 1 GHz tap technology show that operators can achieve the 10G goals with the existing installed base of taps. This will buy the operator more time before they need to pull the trigger and replace them. Hopefully the 1.8/3.0 GHz future taps will be cost effective by that time.

Finally, fiber deep and DAA become more important technologies at helping operators to achieve the 10G goals.

Acknowledgements

The authors would like to gratefully acknowledge the assistance of Frank O'Keefe for his PHY Layer modeling and analysis, in particular for the 1 GHz tap data. The data was priceless.

Abbreviations

BAU	Business as Usual
Bcast	Broadcast
Bps	Bits Per Second
CAA	Centralized Access Architecture
CAGR	Compounded Annual Growth Rate
CAPEX	Capital Expense
CCAP	Converged Cable Access Platform
CM	Cable Modem
CMTS	Cable Modem Termination System
CPE	Consumer Premise Equipment
D3.1	Data Over Cable Service Interface Specification 3.1
DAA	Distributed Access Architecture
DCA	Distributed CCAP Architecture
DEPI	Downstream External PHY Interface
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
DWDM	Dense Wave Division Multiplexing
E2E	End to end
EOL	
EPON	Ethernet Passive Optical Network (aka GE-PON)
EQAM	Edge Quadrature Amplitude Modulator
FD	Fiber Deep
FDX	Full Duplex (i.e. DOCSIS)
FTTH	Fiber to the Home
FTTLA	Fiber to the Last Active
FTTP	Fiber to the Premise
FTTT	Fiber to the Tap
FTTx	Fiber to the 'x' where 'x' can be any of the above
Gbps	Gigabits Per Second
GHz	Gigahertz
HFC	Hybrid Fiber-Coax
HP	Homes Passed
HSD	High Speed Data
I-CCAP	Integrated Converged Cable Access Platform
IEEE	Institute of Electrical and Electronics Engineers
IEQ	Integrated Edge QAM
LDPC	Low Density Parity Check FEC Code
MAC	Media Access Control interface
MACPHY	DCA instantiation that places both MAC & PHY in the Node
Mbps	Mega Bits Per Second
MDU	Multiple Dwelling Unit
MHz	Megahertz
MSO	Multiple System Operator
N+0	Node+0 actives
Ncast	Narrowcast

NFV	Network Function Virtualization
NSI	Network Side Interface
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiplexing Access (Upstream)
OLT	Optical Line Termination
ONU	Optical Network Unit
OOB	Out of Band
OPEX	Operating Expense
OTT	Over the Top
PHY	Physical interface
PNM	Proactive Network Maintenance
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
QoS	Quality of Service
RF	Radio frequency
R-OLT	Remote OLT
RPD	Remote PHY Device
R-MACPHY	Remote MAC-PHY
R-PHY	Remote PHY
RX	Receive
SDN	Software Defined Network
SG	Service Group
SCTE	Society of Cable Telecommunications Engineers
SNR	Signal to Noise Ratio
TaFDM	Time and Frequency Division Multiplexing
Tavg	Average bandwidth per subscriber
Tmax	Maximum Sustained Traffic Rate – DOCSIS Service Flow parameter
TX	Transmit
US	Upstream
VOD	Video on demand
WDM	Wavelength Division Multiplexing

Bibliography & References

[ALB_2014] A. Al-Banna et. al., "The Spectral Efficiency of DOCSIS® 3.1 Systems," SCTE Cable-Tec 2014, SCTE

[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

[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, Z. Maricevic, "Cable 10G vs. Wireless 5G/CBRS – Foe or Friend? A Survey of Next Gen Network Directions", 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

