Fighting FTTH with DOCSIS

A Technical Paper prepared for SCTE by

Mike Darling
Principal Engineer
Rogers Communications
mike.darling@sjrb.ca

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

1. Introduction

Cable operators are preparing to, or have already begun, upgrading their hybrid fibre-coax (HFC) networks in order to offer higher-speed service tiers. There are many upgrade paths and corresponding strategies that cable operators could take, each with different considerations. In this paper each upgrade path will be analyzed from a strategic point of view contemplating topics such as capacity, tier capability, and traffic growth. These circumstances could be compared to those faced by telcos a decade ago when they had to decide whether to upgrade their digital subscriber line (DSL) networks to maintain competitiveness with HFC, although unlike DSL, HFC networks have a long life ahead of them. This paper will consider if lessons can be learned by cable operators from those circumstances. This round of network upgrades represents a significant investment, underlying the importance of a well-thought-out strategy to optimize return on investment.

2. The DOCSIS Ecosystem

The share of global broadband subscribers served with HFC technology has plateaued in recent years. This is primarily due to three reasons:

1. Most cable operators have moved to fibre-to-the-home (FTTH) in their greenfield builds, limiting the number of new HFC homes passed built each year.
2. Government subsidy programs are geared toward FTTH in underserved areas, increasing the overall share of FTTH homes passed.
3. A subset of HFC operators are upgrading their HFC plant to FTTH.

This trend can be seen in the OECD broadband data [1] in Figure 1. While both FTTH and HFC had been increasing market share until 2020 at the expense of DSL, in the last two years HFC's share has been reduced as well.

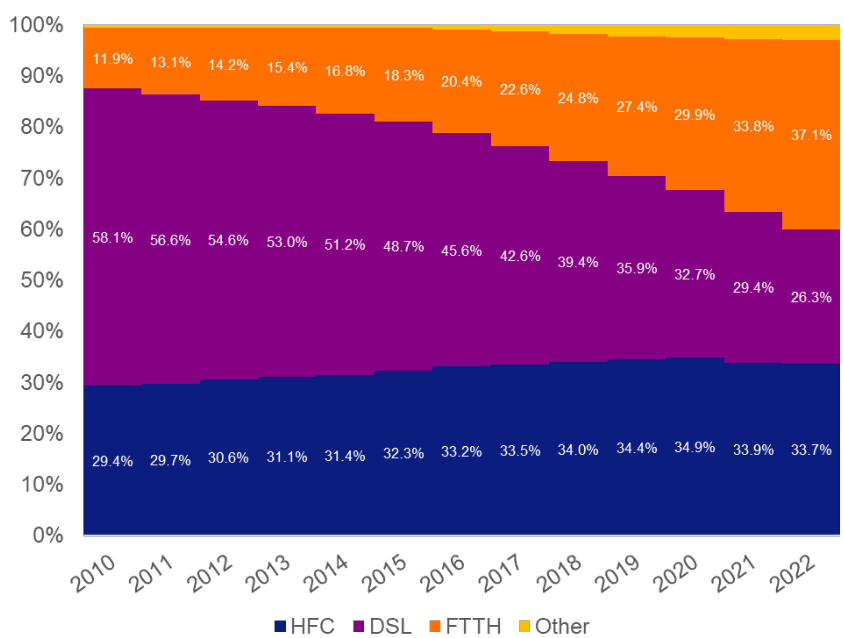

Figure 1 – Broadband Subscriptions by Technology - OECD [1]

Looking at this trend, some are questioning the future of HFC technology and whether there will be continued innovation in the ecosystem. The myriad potential upgrade options and publicly announced plans by cable operators have added to the complexity faced by cable operators considering upgrade paths. This is similar to the situation telcos found themselves in years ago as the DSL market began to contract, with a few very important differences. An analysis of this situation has the potential to inform the current discussion.

2.1. Comparison to DSL

Telcos in Canada began to transition away from their legacy DSL networks approximately 10 years ago. It is worthwhile examining the circumstances that led them to that decision. Just as data over cable service interface specifications (DOCSIS) technology improved through DOCSIS 1.0, 1.1, 2.0, 3.0, 3.1 and most recently DOCSIS 4.0, broadband DSL had its own path. Starting with Asymmetric DSL (ADSL), the technology moved to very-high-speed DSL (VDSL), and finally G.fast [2], adding variants that improved speeds along the way, as shown in Figure 2.

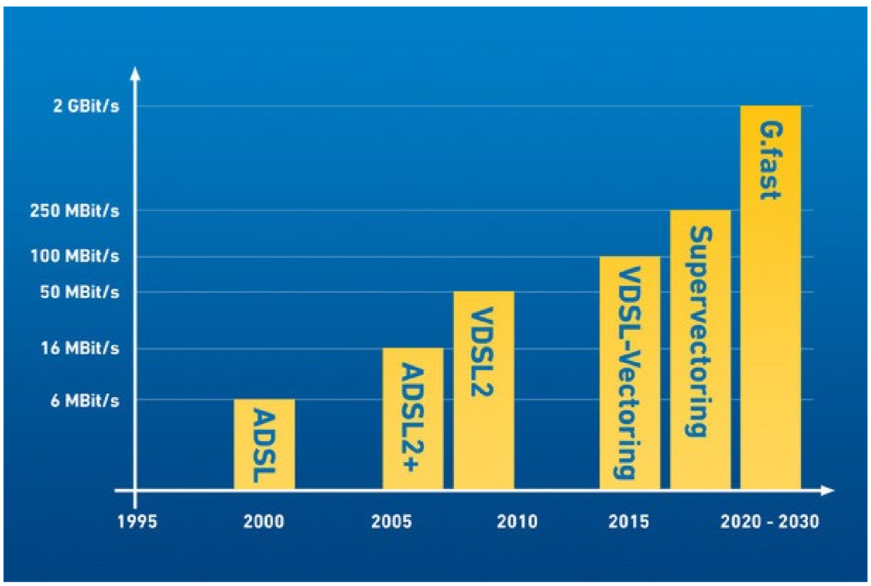

Figure 2 – DSL Speed Evolution [2]

As with HFC technology, DSL standards incorporated additional spectrum over time. Unlike HFC technology, one of the main challenges for DSL is crosstalk, or interference between subscriber lines, which are carried in bundles as shown in Figure 3.

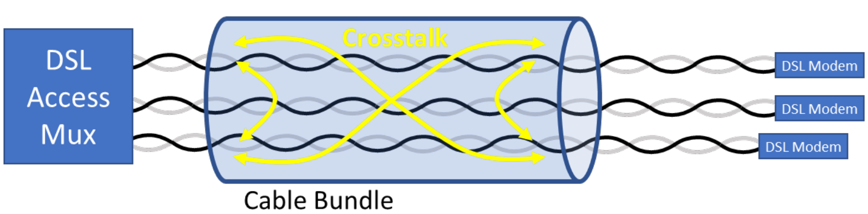

Figure 3 – Crosstalk in DSL Networks

To increase speeds over the DSL network, crosstalk had to be minimized. To this end, vectoring, which pre-distorts the signal before being sent through the carrier bundle, began being used to balance out the impact of crosstalk. There is only so much crosstalk reduction to be accomplished through vectoring however, and to support higher speeds the distance between the access node and CPE had to be minimized, as shown in Figure 4 – Fibre-to-the-curb (FTTC) Speed Graph.

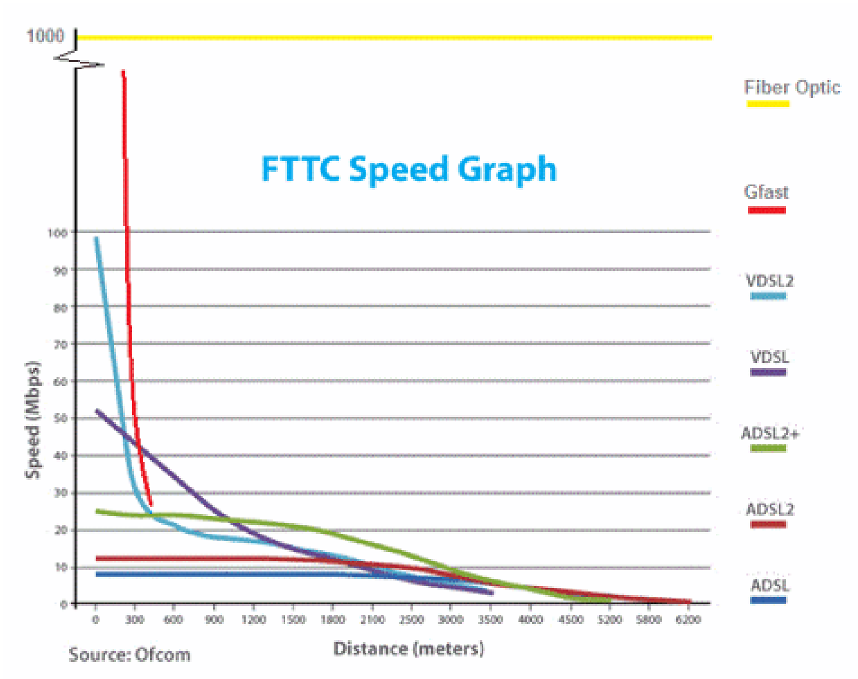

Figure 4 – Speed vs Distance in DSL Networks [3]

To reduce distances, telco operators were required to change their outside plant architectures, which began with loop bundles originating from central offices (COs) and travelling several kilometers to subscriber premises. The new architecture made use of DSL access multiplexers (DSLAMs) installed in the outside plant, which brought fibre connectivity closer to the subscriber, limiting the distance of the bundle and thus reducing crosstalk.

Over time, subscriber speed requirements increased due to internet protocol (IP) video adoption and competition from HFC operators who could offer speeds higher than VDSL2 technology could support. This forced the telcos to make a decision whether to upgrade their DSL networks to G.fast, which could increase speeds offered, but required DSLAMs to be moved even closer to subscribers; or to move directly to FTTH.

In Canada the telcos decided to forego a large-scale upgrade to G.fast, deciding instead to begin upgrading their networks directly to FTTH. The dynamic for both cable and telco operators is shown in Figure 5.

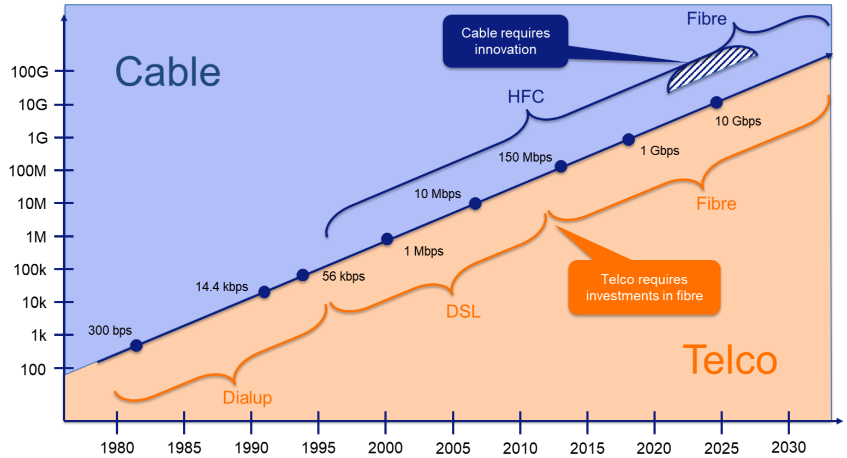

Figure 5 – Subscriber Speeds Over Time

There are some who would argue that cable operators are in the same place now as the telcos were approximately 10 years ago, having to decide whether to perform another upgrade to their current technology, or make the move to an all-fibre network. There are important differences in this situation, however.

A significant difference is that a G.fast upgrade to telco DSL networks would require a large number of net-new DSLAMs to be installed in the network, along with the temperature controlled cabinets that house them in order to reduce the loop lengths between DSLAMs and subscribers. On the HFC side, a DOCSIS 4.0 upgrade only requires the exchange of equipment to newer versions.

Another difference is the additional speed achieved by the upgrade. The speed increase realized by G.fast would strongly depend on the loop distances, but it is likely to have topped out at several hundred Mbps. This would not have been sufficient to match DOCSIS 3.0 speeds beginning to be available at that time, which had a downstream capacity of greater than 1 Gbps. DOCSIS 4.0 has the potential to upgrade downstream speeds to over 10 Gbps which is on par, or slightly better than, 10 Gbps Passive Optical Network (XGS-PON), the current top of the line FTTH technology being deployed.

The last difference is the traffic growth environment. In the mid-2010s year-over-year (YoY) traffic growth was routinely over 50%, as shown for the Western Canadian cable footprint in Figure 6, causing network demand to double in less than two years. The solid line is peak traffic, shown for reference but without an axis, while the dashed line is YoY growth, measured monthly. This is drastically different from today, where growth has slowed to below 20% and is forecasted to drop even further.

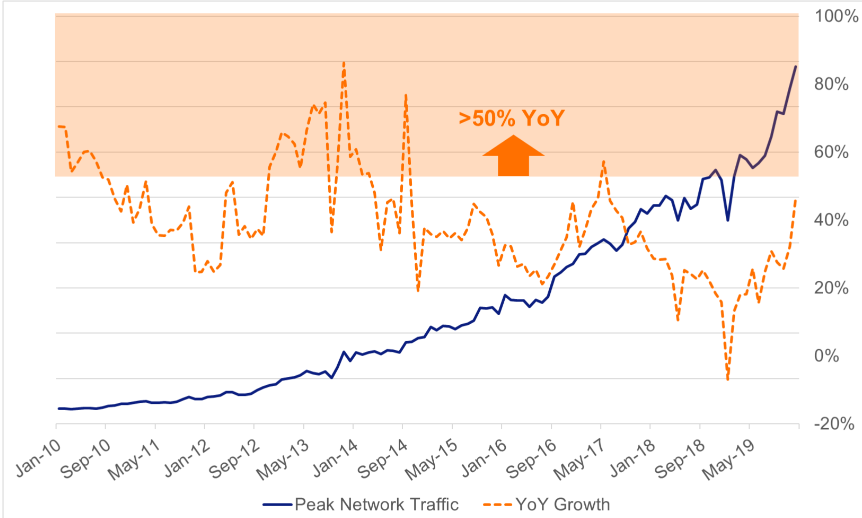

Figure 6 – Downstream Peak Network Traffic and YoY Growth in 2010s

The important question that cable operators must answer is which upgrade path is optimal. The answer will depend on what cost per home passed the operator forecasts for each type of upgrade, and what kind of traffic growth each operator has forecasted for the coming years. The former topic has been the subject of much debate, but will be specific to each operator's context, while the latter topic will be explored in the next section.

3. Traffic Growth

The COVID-19 pandemic generated a large amount of network traffic in a short time due to remote-learning, work from home, and people spending more time in their residences. As the pandemic started to wane, traffic growth slowed, resuming a trend that had started in the years before COVID-19. Figure 7 shows the total network traffic trend in the Western Canadian cable footprint in the last six years. The large growth due to COVID-19 can clearly be seen in the higher YoY growth in the months from roughly March 2020 to March 2021. In the two years since that time, the YoY growth rate has hovered around 10%.

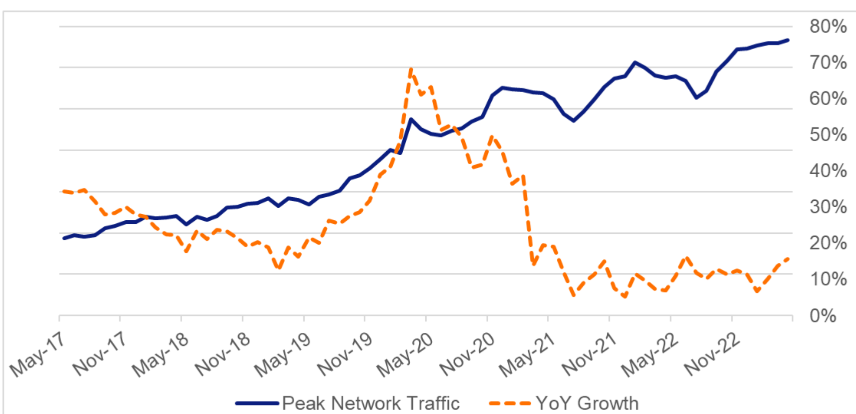

Figure 7 – Downstream Peak Network Traffic and YoY Growth

In the upstream, the story is much the same. Pre-pandemic YoY growth was slowing, a trend that returned at the end of the pandemic, with YoY growth below zero for some specific months as shown in Figure 8.

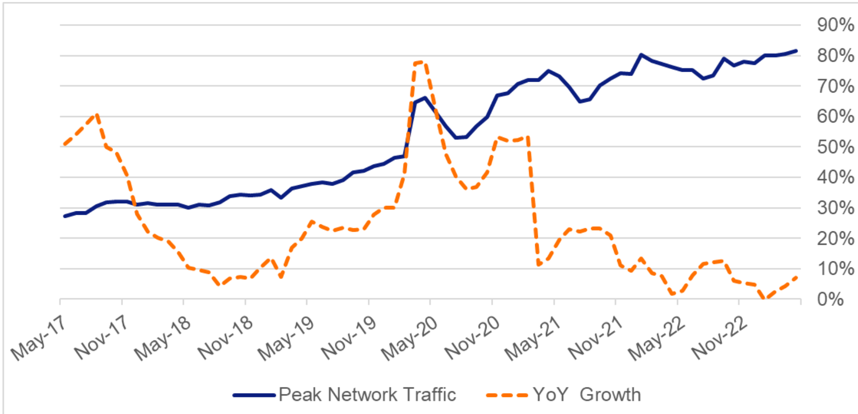

Figure 8 – Upstream Peak Network Traffic and YoY Growth

Forecasting network traffic growth is difficult, as it tends to come from new and unpredicted network use. However, most network growth has been triggered by some variety of IP video. In the mid-2010s this was Netflix, YouTube, and the greater adoption of IP Television (IPTV), while during the COVID-19 pandemic it was video conferencing that especially increased upstream growth rates. Given this we can analyze the potential impacts of video in the coming years.

3.1. IPTV transition

While IPTV adoption is high, cable operators continue to offer broadcast video based on Quadrature Amplitude Modulated (QAM) signals. Most operators have chosen to move to a unicast or multicast IPTV platform as there is no limitation on the number of channels carried. As subscribers are transitioned from the QAM infrastructure to IPTV, additional traffic is added to the DOCSIS network. The impact of the transition depends to what extent the subscribers were watching broadcast video and to what extent they were watching over-the-top (OTT) IPTV services that were already being carried over the DOCSIS network. In the worst case, subscribers can be considered to have increased their network use from a non-IPTV user to match the average IPTV subscriber. Most broadband video subscribers watch OTT IPTV services prior to switching to linear IPTV, thus reducing the impact of the transition. In addition, operators can influence the speed at which the IPTV transition happens through marketing strategies. Lastly, the upside of the IPTV transition is when all subscribers are transitioned, the spectrum currently used to carry broadcast video can be freed up for DOCSIS signals, increasing the capacity available. At current DOCSIS serving group sizes, this transition ends up using spectrum more efficiently, although creating an accordion effect where both IPTV and broadcast video must be supported until the latter is no longer required.

3.2. Higher Resolutions

In addition to increased IPTV subscriber counts, the resolution of IPTV signals, and thus the quantity of traffic, is increasing over time. 4K video signals create roughly four times the traffic as high-definition (HD) signals. Although 4K television adoption is high and increasing, content continues to be limited. This is partially due to 4K having limited perceived benefit over HD in most viewing environments. As shown in Figure 9, 4K or Ultra HD only has a perceived benefit either for large televisions or for viewers sitting close to their TV.

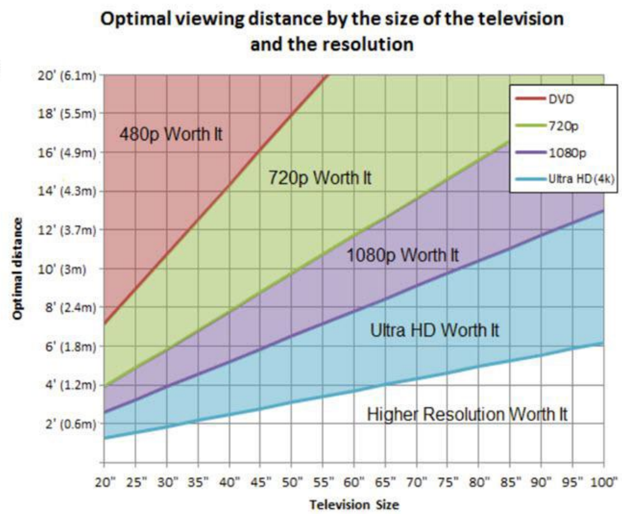

Figure 9 – Optimal Viewing Distance [4]

The first content category that will drive more 4K adoption is sports, which benefits from higher resolutions and higher bitrates in general due to quick motion. This could be observed in IPTV network traffic statistics during the National Hockey League (NHL) playoffs this year, where games were carried in 4K. Especially when local teams were involved, 4K viewership drove increased peak network traffic.

3.3. Increased Coding Efficiency

Video bitrates for the same resolution have decreased over time, owing to the increased coding efficiency of new compression standards. Many operators are still carrying Motion Picture Experts Group Version 2 (MPEG-2) signals in their broadcast video platforms but have moved to MPEG-4 for IPTV. As shown in Figure 10, the high efficiency video coding (HEVC) standard will allow for a further approximately 50% reduction in bitrates when compared to MPEG-4.

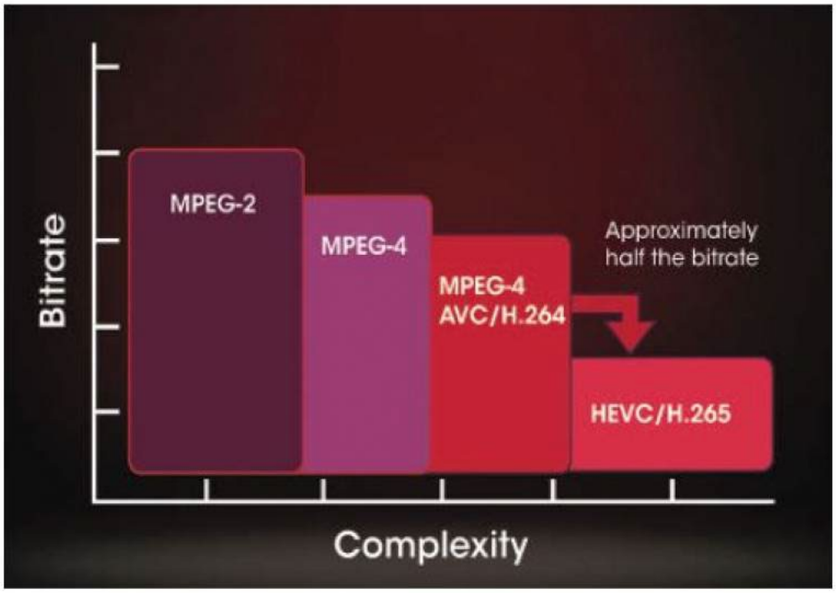

Figure 10 – Coding Efficiency of Video Standards [5]

The two trends of increased video resolution and increased video coding efficiency are likely to mostly cancel each other out over time. Future network growth is likely to be driven by some form of video. Luckily for network operators, they have influence over video bitrates and speed of IPTV transition and thus some measure of control over network traffic growth.

4. DOCSIS Capacity Calculations

Network capacity calculations can be done using several different methods. There is the physical data rate (PHY rate), which includes all overhead, there is also the theoretical maximum which takes the PHY rate and subtracts out the theoretical overhead, and finally there is the practical data rate observed in real networks. While marketing material may make use of PHY rates or theoretical maximums, network planners need to use real-world capacity calculations.

In calculating DOCSIS capacity, achievable modulation rates must be considered. While in most circumstances a modem may receive a signal with high enough signal-to-noise ratio (SNR) to be able to use 4096-QAM signals, there may be a material number which cannot. In Figure 11, the SNR distribution of several hundred thousand modems in Western Canada is displayed.

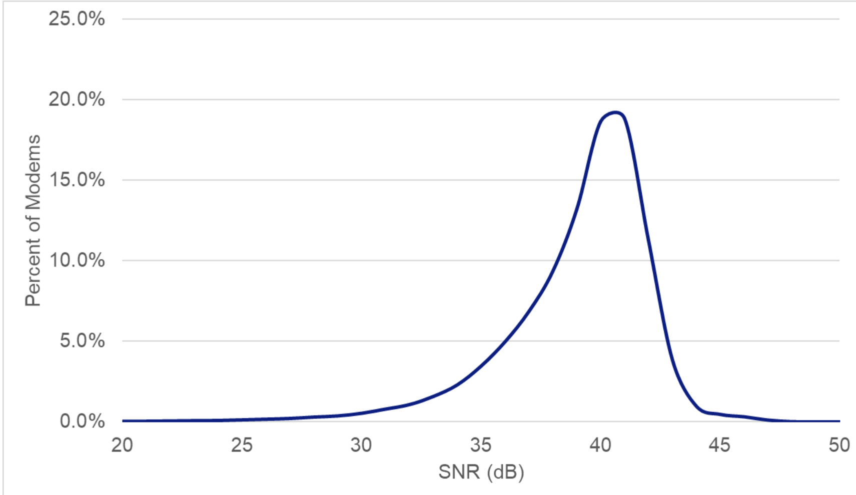

Figure 11 – Downstream SNR

Assuming 4096-QAM is achieved when SNR is greater than 38 dB, the majority of the network would support that modulation rate. However, a significant portion of the network has an SNR below 38 dB. Care must be taken when declaring the capacity of the network, as using 4096-QAM will overcount the capacity to some of the network and using 1024-QAM would undercount the majority. Overcounting may lead to circumstances where operational teams are pushed to increase SNR to achieve 4096-QAM in areas of the network where it is difficult and expensive. One method of estimating capacity is to decide what percentage of the network is less than 4096-QAM acceptable. In this case capacity calculations can be done using the 90th percentile, or what performance 90 percent of the network is able to achieve. In this paper 9 bps per Hertz (bps/Hz) is used in the downstream and 7 bps/Hz in the upstream.

5. Traffic Patterns

ADSL was the first version of DSL technology aimed at multimedia access, including video. Early standards had downstream to upstream (DS:US) capacity ratios of 10:1 or greater. It was only with the introduction of VDSL2 that high-speed symmetric services were offered.

FTTH technologies have symmetric and asymmetric variants, with the most common in use in Canada currently being Gigabit Passive Optical Network (GPON) with 2.5 Gbps downstream and 1.25 Gbps upstream PHY rates. While this is asymmetric, the telcos offer mostly symmetric services, only offering asymmetric tiers when the upstream cannot match the downstream, such as 1.5 Gbps downstream by 940 Mbps upstream. The next generation of FTTH technology is XGS-PON, with a PHY rate of 10 Gbps symmetric, which will allow operators to offer symmetric tiers greater than 1 Gbps. Offers up to 8 Gbps symmetric over XGS-PON are in market today.

Starting with VDSL2, telcos have tried to highlight their symmetric services as an advantage over HFC technology, which has always been asymmetric. The success of this strategy has been limited as network traffic is inherently asymmetric.

In the Western Canadian cable footprint the network has exhibited a DS:US traffic ratio of approximately 15:1 since the mid-2010s. This has remained stable despite some major changes in the DS:US ratio of offered tiers, which aggregated for all subscribers makes up the network provisioned ratio. Figure 12 shows both the DS:US traffic ratio and the DS:US provisioned ratio. There were two major changes in DS:US provisioned ratio, first when downstream provisioned rates for several popular tiers were doubled overnight, and next when upstream provisioned rates were increased as a mid-split upgrade was close to completion. As can be seen, the DS:US traffic ratio stayed the same, despite the provisioned ratio changing quite drastically. This points to the traffic ratio being an inherent product of subscriber behaviour rather than a product of network characteristics.

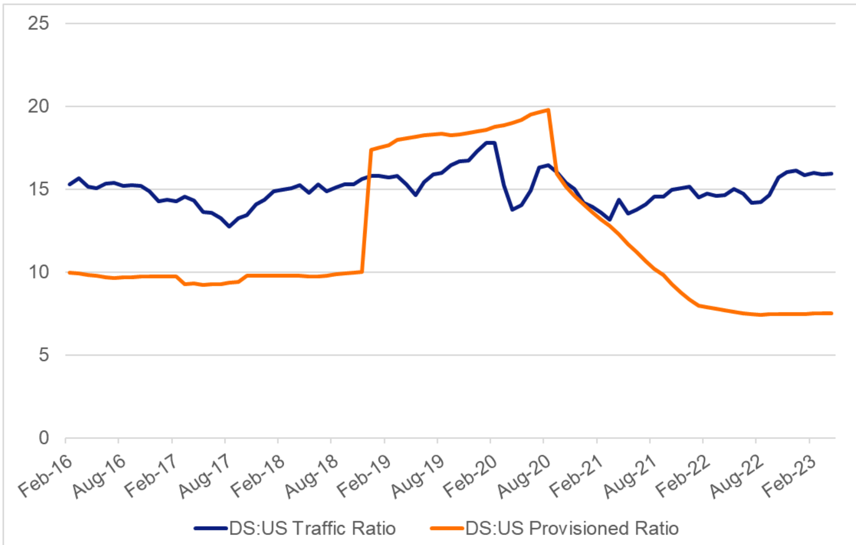

Figure 12 – DS:US Traffic and Provisioned Ratios

This leads to the question as to whether symmetric tiers are necessary, and if not, whether they will become so in the future. As with traffic growth it is not easy to predict what applications might cause a change in DS:US traffic ratios, but it is difficult to envision a wide-spread scenario where subscribers share more information with the world than they consume.

6. FTTH Competition

Canadian telcos continue to aggressively upgrade their DSL networks to FTTH. As of the end of 2022, Bell's network was 70% FTTH [7], and Telus' network was 91% FTTH [8]. They currently offer speeds of 1.5 Gbps downstream and 940 Mbps upstream over GPON, while quickly upgrading their platforms to XGS-PON. They are able to upgrade their networks to XGS-PON by adding new network equipment that can coexist with installed GPON equipment as shown in Figure 13 [6].

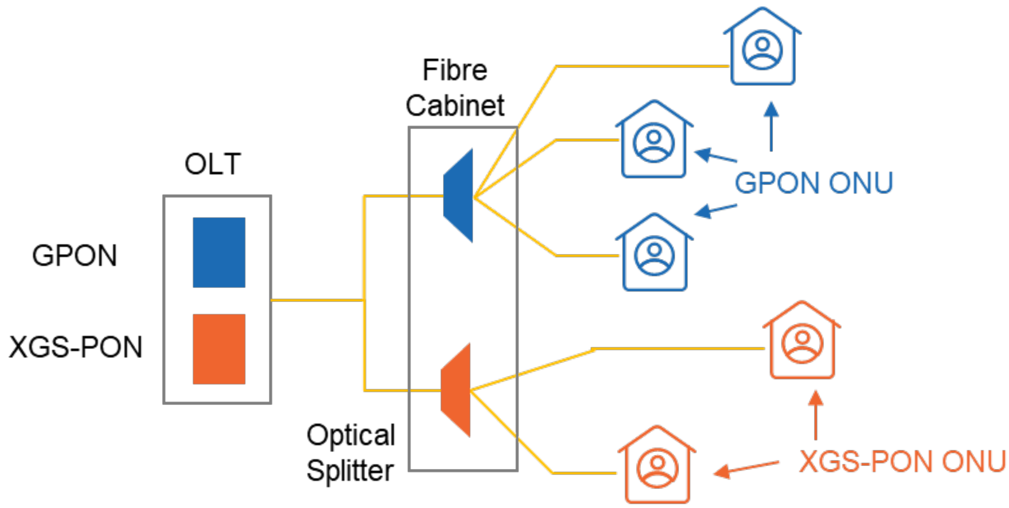

Figure 13 – Passive Optical Network Coexistence

In the future, 50 Gbps PON and subsequent versions will coexist with GPON and XGS-PON, allowing for higher speed offerings.

7. HFC Upgrade Options

HFC networks have long operated in a frequency division duplex (FDD) mode using a sub-split, where the upstream is carried from 5 to 42 MHz with the downstream starting at 54 MHz. Previous HFC network upgrades were focused on adding downstream spectrum, 550 MHz then 750 MHz, 860 MHz and 1 GHz. Using DOCSIS 3.1 cable modems (CMs), gigabit speeds are possible with this network setup, allowing cable operators to offer competitive speed tiers. From here operators have several options to upgrade the HFC network, including remaining with a sub-split plant.

7.1. Sub-split

As mentioned, a sub-split plant uses upstream spectrum to 42 MHz and downstream spectrum from 54 MHz to either 750 MHz, 860 MHz or 1 GHz. Figure 14 shows an example with a top frequency of 750 MHz along with a possible breakdown in spectrum use.

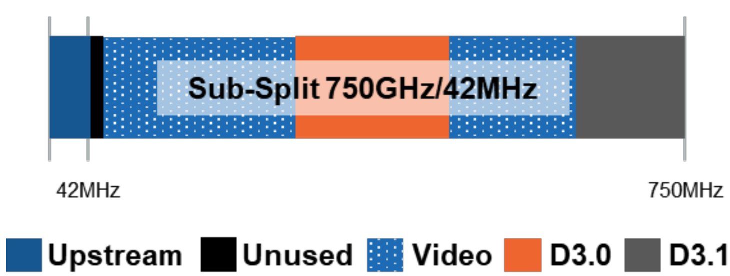

Figure 14 – Sub-split HFC Network

The capacity of a sub-split network depends on how much of the spectrum is used for DOCSIS vs broadcast video, and whether and how much spectrum is being used for DOCSIS 3.1 signals, which use spectrum more efficiently than DOCSIS 3.0 signals. Assuming all spectrum is allocated to DOCSIS, the spectrum available is shown in Table 1.

Table 1 – Available Spectrum in Sub-split Plant

Upper Frequency (MHz)	Upstream Spectrum (MHz)	Downstream Spectrum (MHz)
750	37	696
860	37	806
1,002	37	948

Using a capacity estimate of 9 bps/Hz in the downstream and 7 bps/Hz in the upstream as discussed in section 4, we arrive at the capacities shown in Table 2, with related DS:US ratios.

Table 2 – Capacity in Sub-split Plant

Upper Frequency (MHz)	Upstream Capacity (Mbps)	Downstream Spectrum (Mbps)	DS:US Ratio
750	259	6,264	24:1
860	259	7,254	28:1
1,002	259	8,532	33:1

A sub-split plant in combination with DOCSIS 3.1 allows operators to offer services fulfilling the requirements of the vast majority of subscribers. The DS:US capacity ratios in a fully utilized sub-split plant, however, are higher than the peak traffic ratios which, as discussed earlier, have stabilized at approximately 15:1.

7.2. Mid-split

A mid-split plant uses upstream spectrum to 85 MHz and downstream spectrum from 108 MHz to either 1 GHz or 1.2 GHz. Figure 15 shows an example with a top frequency of 1 GHz.

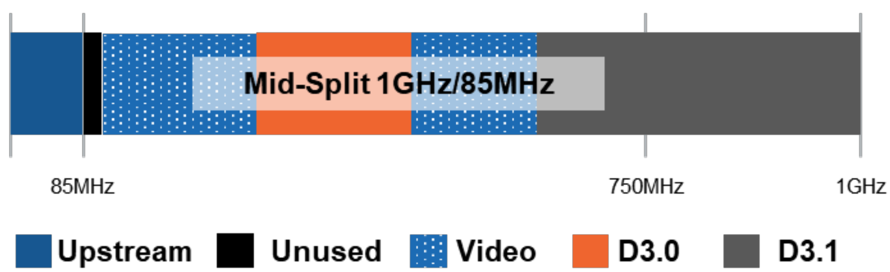

Figure 15 – Mid-split HFC Network

Assuming all spectrum is allocated to DOCSIS, the spectrum available is shown in Table 3.

Table 3 – Available Spectrum in Mid-split Plant

Upper Frequency (MHz)	Upstream Spectrum (MHz)	Downstream Spectrum (MHz)
1,002	80	894
1,218	80	1,110

Again, using a capacity estimate of 9 bps/Hz in the downstream and 7 bps/Hz in the upstream, we arrive at the capacities shown in Table 4, with related DS:US ratios.

Table 4 – Capacity in Mid-split Plant

Upper Frequency (MHz)	Upstream Capacity (Mbps)	Downstream Capacity (Mbps)	DS:US Ratio
1,002	560	8,046	14:1
1,218	560	9,990	18:1

A mid-split upgrade to 85 MHz in the upstream and to 1 GHz or 1.2 GHz in the downstream allows operators to offer services that meet the requirements of virtually all subscribers and has a DS:US capacity ratio in the same range as DS:US traffic ratios.

7.3. High-split

A high-split plant uses upstream spectrum to 204 MHz and downstream spectrum from 258 MHz to either 1 GHz, 1.2 GHz, or 1.8 GHz. Figure 16 shows an example with a top frequency of 1.2 GHz.

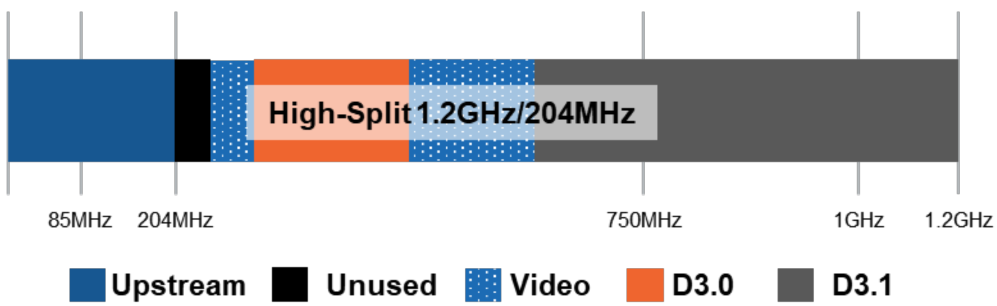

Figure 16 – High-split HFC Network

Assuming all spectrum is allocated to DOCSIS, the spectrum available is shown in Table 5.

Table 5 – Available Spectrum in High-split Plant

Upper Frequency (MHz)	Upstream Spectrum (MHz)	Downstream Spectrum (MHz)
1,002	199	744
1,218	199	960
1,794	199	1536

Again, using a capacity estimate of 9 bps/Hz in the downstream and 7 bps/Hz in the upstream, we arrive at the capacities shown in Table 6, with related DS:US ratios.

Table 6 – Capacity in High-split Plant

Upper Frequency (MHz)	Upstream Capacity (Mbps)	Downstream Capacity (Mbps)	DS:US Ratio
1,002	1,393	6,696	5:1
1,218	1,393	8,640	6:1
1,794	1,393	13,824	10:1

A high-split upgrade allows operators to offer services that meet the requirements of most subscribers and has a lower DS:US ratio than is seen in DS:US traffic ratios. A milestone that a high-split upgrade potentially enables is a 1 Gbps upstream service. This represents greater than 70% of the capacity of the upstream which is higher than operators are generally comfortable with, and any interference from over the air (OTA) radio or TV signals may reduce this capacity. This is similar to FTTH operators offering a 1 Gbps upstream service on asymmetric GPON, which has a PHY rate of 1.25 Gbps.

An additional challenge with high-split upgrades is regarding broadcast video platforms using SCTE 55-1 or 55-2 standards, where the downstream telemetry does not operate in the frequencies required with a high-split plant. This requires that all SCTE 55-1 and 55-2 broadcast video platforms be swapped out for DOCSIS set-top gateway (DSG) platforms which use the DOCSIS network for command and control, or IPTV platforms which use the DOCSIS network for command and control as well as to transfer video content.

7.4. Ultra-high-split

An ultra-high-split (UHS) plant uses upstream spectrum to 300 MHz, 396 MHz, 492 MHz, or 684 MHz and downstream spectrum from 372 MHz, 492 MHz, 606 MHz, or 834 MHz to 1.8 GHz. Figure 16 shows an example with upstream to 396 MHz and a top frequency of 1.8 GHz.

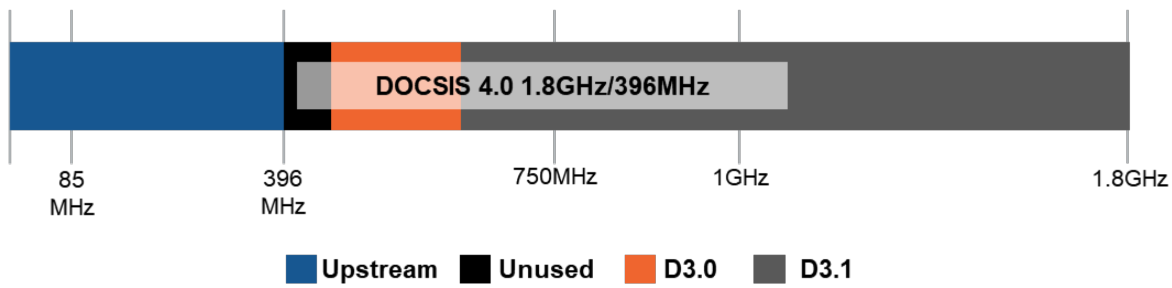

Figure 17 – Ultra-high-split HFC Network

Assuming all spectrum is allocated to DOCSIS, the spectrum available is shown in Table 7.

Table 7 – Available Spectrum in Ultra-high-split Plant

UHS Variant	Upper Frequency (MHz)	Upstream Spectrum (MHz)	Downstream Spectrum (MHz)
UHS-300	1,794	295	1,422
UHS-396	1,794	391	1,302
UHS-492	1,794	487	1,188
UHS-684	1,794	679	960

Again, using a capacity estimate of 9 bps/Hz in the downstream and 7 bps/Hz in the upstream, we arrive at the capacities shown in Table 8, with related DS:US ratios.

Table 8 – Capacity in Ultra-high-split Plant

UHS Variant	Upstream Capacity (Mbps)	Downstream Capacity (Mbps)	DS:US Ratio
UHS-300	2,058	12,798	6:1
UHS-396	2,737	11,718	4:1
UHS-492	3,409	10,692	3:1
UHS-684	4,753	8,640	2:1

As can be observed, the services as well as the DS:US ratio vary significantly depending on the choice of diplex frequency. Crucially, all UHS variants have a downstream capacity in excess of the estimated 8,600 Mbps that XGS-PON is practically capable of [10].

8. Strategies

8.1. Maintaining a Sub-split Plant

Operators may choose to maintain their sub-split network and perform node splits to control congestion. This strategy may be used as an end-goal if the operator believes that network traffic has saturated, or will soon, and that upstream tiers supported over sub-split are sufficient. It also makes sense if investment needs to be minimized for various reasons, or to buy time to allow the network to be upgraded to FTTH, a capital intensive and slow process.

8.1.1. Capacity Perspective

There are two major considerations when analyzing capacity - the service tiers that can be offered and the ability to manage congestion. If an operator desires to offer symmetric tiers to compete with FTTH it is better to have symmetric capacity, while if an operator desires to manage network congestion most efficiently it is better to have asymmetric capacity, in-line with observed DS:US ratios, discussed earlier. The capacity calculations performed in section 7.1 assume that all spectrum is allocated to DOCSIS and reveal a DS:US ratio from 24 to 33:1 depending on the high frequency. This DS:US ratio is not optimal for either offering symmetric services or managing congestion. In practice, however, all downstream spectrum is not allocated to DOCSIS, with a portion continuing to support broadcast video, meaning that the actual DS:US ratio is lower. The number varies from operator to operator depending on how much spectrum is dedicated to video. In addition, the upstream and downstream capacity will depend on whether DOCSIS signals are DOCSIS 3.0 or DOCSIS 3.1, as the latter are capable of higher orders of modulation and thus are more spectrally efficient. This is in turn dictated by the mix of DOCSIS customer premises equipment (CPE) in the network.

8.1.2. Node Splits

A major benefit of node splits is that resources can be targeted where they are needed. A node split is when one or more logical serving groups (SGs) are added to a pocket of homes connected to an existing serving group, decreasing the number of homes sharing DOCSIS capacity. This can involve adding DOCSIS ports in a hub site, adding additional optics feeding the optical node, or building fibre to a subset of homes, usually down a trunk run, and installing a new optical node. The latter scenario is shown in Figure 18.

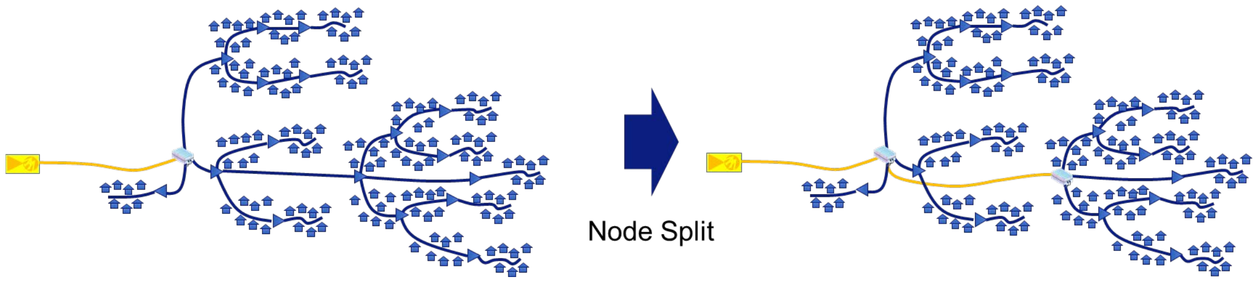

Figure 18 – Node Split Example

The aim of a node split is to precisely divide up the homes from a single serving group into two or more serving groups, theoretically equally dividing up the traffic. In practice, node splits are never perfectly balanced due to physical characteristics of the outside plant, and to the extent that they are balanced in terms of numbers of homes this does not necessarily equate to traffic balance. Figure 19 shows traffic distribution by modem at the network daily peak, averaged over a month, which illustrates that most users' traffic demands are small, even at network peak. The larger the serving group, the more likely its distribution looks like the network aggregate distribution and the more likely a node split will result in two balanced nodes.

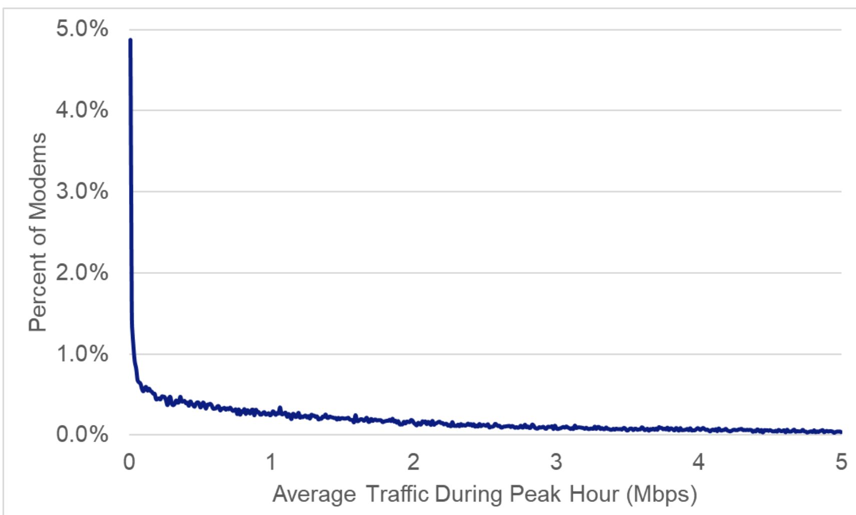

Figure 19 – Peak Hour Traffic Distribution

Due to the statistical nature of traffic use, it is not uncommon to split a congested serving group only to find that one of the new serving groups remains congested. This is more likely to occur as serving group sizes shrink, as a small serving group that is congested is more likely to be dominated by a small number of heavy users. From a capital perspective, node splits also increase in terms of cost per home passed as serving groups shrink, as the costs are primarily labor costs incurred during fibre builds. The result is that node splits have increasing costs and diminishing returns over time as serving groups shrink. At this point, the strategy is no longer cost efficient, and another approach should be explored.

Especially with larger serving group sizes, node splitting is a simple but effective strategy which allows for capacity to be added where it is needed, reducing wasted effort. The eventual outcome of node-splitting is small node sizes in areas where demand is great, and large node sizes where demand is small. In recent years the Cloonan Formula [9] has been used to estimate the effort required to upgrade the network over time. The simple version is shown below:

Serving Group Capacity >= Nsub*Tavg + K*Tmax_max

Where Tavg is the average traffic per subscriber in bits per second during the network busy hour, Nsub is the number of subscribers in the serving group, Tmax_max is the maximum provisioned traffic rate in bps and K is a quality of experience (QoE) factor.

Using the Cloonan Formula with a known serving group capacity, Tavg and Tmax_max, a target number of subscribers per node can be calculated. This however is only useful to the extent that Tavg and Nsub are independent variables, which after years of adding capacity where it is required, has largely ceased to be the case. This may lead operators to the conclusion that they need to split larger nodes to hit a target, even in the absence of congestion, reducing the strategic benefit of targeting capacity where it is needed. Figure 20, based on data from Western Canada, shows the relationship between Tavg and Nsub, which shows that Tavg is smaller for larger Nsub.

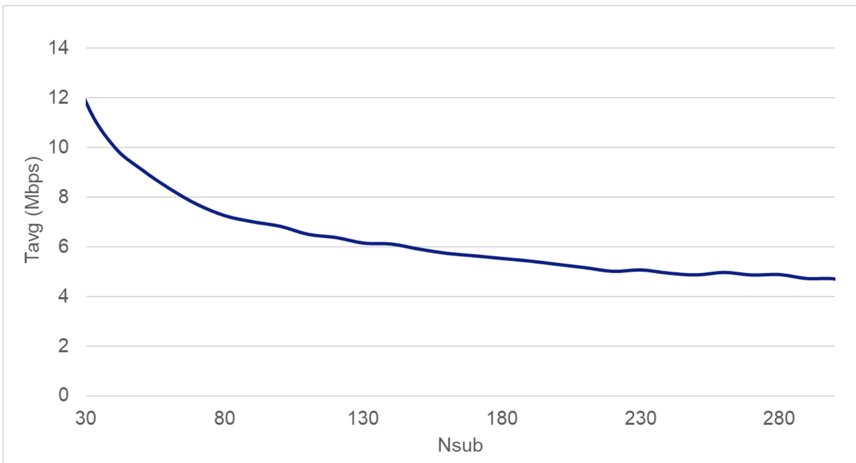

Figure 20 – Tavg vs Nsub

8.1.3. IPTV Transition

Transitioning broadcast video subscribers over to an IPTV platform can free up spectrum that can be allocated to DOCSIS. If the subscriber count is not too high, there will be a net benefit in efficiency by making the transition. The transition requires swapping out subscriber CPE and can be done on a system-by-system or node-by-node basis. By transitioning on a node-by-node basis capacity upgrades can be targeted where they are needed, but the complexity of the upgrade is increased.

8.1.4. DOCSIS CPE Upgrades

Transitioning DOCSIS 3.0 signals to DOCSIS 3.1 will increase capacity and can be done on a node-by-node basis, as with the IPTV transition. Depending on service tiers supported over DOCSIS 3.0 CPE, a small number of DOCSIS 3.0 signals may be left to support low tier services, as well as low-bandwidth services such as phone and broadcast video, allowing legacy phone and video CPE to be kept in the network for longer.

8.2. Mid-split Upgrade

Operators may choose a mid-split upgrade if node-splits cease being an efficient means to control congestion, or if there is desire to increase upstream services. Assuming all spectrum is allocated to

DOCSIS, the DS:US ratio of a mid-split network is either 14:1 or 18:1, which matches the observed DS:US traffic ratio well. Although Gbps symmetric services are not supported with a mid-split network, upstream services of 200 Mbps are in market currently, leveraging mid-split networks.

A mid-split upgrade was completed in the Western Canadian cable footprint, allowing for new higher tiers, increased speeds in current tiers, and congestion to be brought to near-zero. Mid-split upgrades have the advantage that the majority of broadcast video hardware continues to function, limiting upgrade costs to the video platform. The outside plant upgrade can be performed as a drop-in upgrade, meaning that re-spacing amplifiers, a costly endeavor, is avoided. Depending on the status of the network, much of the sub-split network hardware can be retained. Costs can be limited if amplifier and node housings can be reused, and analog node optics maintained. Similarly, tap and passive housings can be reused, and only the faceplates changed out.

If an operator has a large installed base of broadcast video CPE that makes use of SCTE 55-1 or SCTE 55-2 signaling, then a mid-split upgrade may be an optimal strategy that can eliminate congestion and allow for increased tiers without a large increase in capital intensity and in a relatively short time frame.

8.3. High-split Upgrade

If operators do not have a large installed base of SCTE 55-1 or SCTE 55-2 based broadcast video, or want to offer Gbps upstream services, then a high-split upgrade makes sense.

There are a few decisions to be made when upgrading to high-split. The first is whether to move to a DAA platform at the same time. While it is possible to continue using analog optics in a high-split network up to 1.2 GHz, the upstream transmitter in the optical node may be a performance bottleneck. Upstream HFC systems are designed to maximize SNR without saturating the upstream transmitter. If signal levels into the upstream transmitter are too high, the laser will saturate or clip, reducing the SNR or causing bit errors. If the signal levels into the transmitter are too low, the SNR will not support higher order modulations. For this reason, it may make sense to move to a DAA architecture which replaces the analog optical link with a digital link. DOCSIS 4.0, and by extension 1.8 GHz plant, assumes the use of DAA optics. Transitioning to DAA can cause challenges, as current analog nodes may be using a CWDM or DWDM system which does not interoperate with digital optics. In this case all nodes using the same multiplexing system might have to be upgraded in a similar timeframe.

The next decision is whether to upgrade the taps and passives. This will depend on whether the current devices are 1 GHz or lower, and whether spectrum over 1 GHz is free to be utilized. If the devices are lower than 1 GHz, then they must be replaced, either to 1 GHz, 1.2 GHz, 1.8 GHz or higher if available and cost effective. If the current devices are already 1 GHz, and spectrum above 1 GHz is used by Multimedia over Coax Alliance (MoCA) in-home video signals, then it is likely not worth upgrading the passives. If current devices are 1 GHz and the spectrum above 1 GHz is free, then it may be worth an upgrade to unlock the additional spectrum.

Another decision is whether to upgrade to 1.8 GHz in the downstream, which is likely to be dictated more by equipment availability than by strategy. Assuming the equipment cost is only incrementally more than 1 or 1.2 GHz, it makes sense to upgrade to the highest bandwidth amplifiers, taps and passives, even if spectrum above 1 GHz will not be used until a later date.

A high-split 1.8 GHz network has the advantage of being able to offer downstream service tiers competitive with, or better than, XGS-PON and a 1 Gbps upstream. As mentioned earlier, however, upgrading to an ultra-high-split variant provides more upstream headroom and is likely to be a better strategy in this case.

8.4. Ultra-high-split Upgrade

A high-split 1.8 GHz plant offers downstream capacity greater than 10 Gbps and upstream capacity in excess of 1 Gbps. An upgrade to UHS from high-split 1.8 GHz trades off downstream capacity for upstream capacity, and so the decision depends upon what maximum tier an operator wishes to offer and what percent of the total capacity the operator is willing to offer as a service. The capacity tradeoff is shown in Figure 21, using 9 bps/Hz in the downstream and 7 bps/Hz in the upstream.

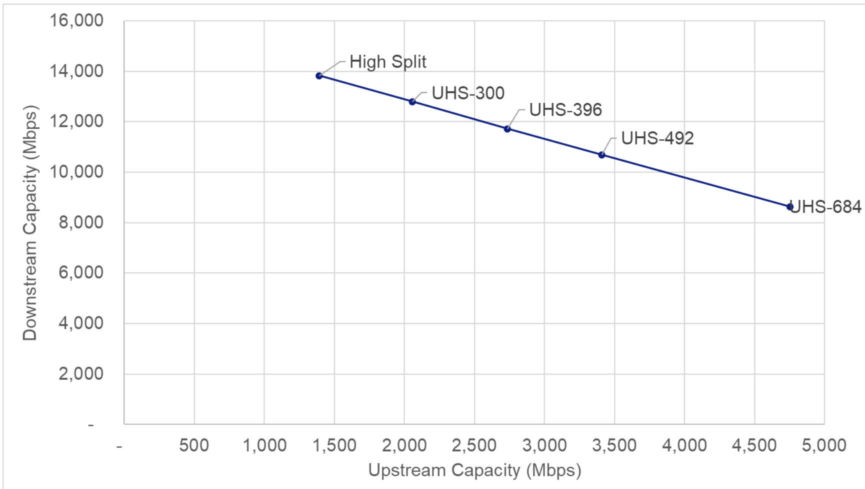

Figure 21 – High-split and UHS Capacities

The two strategies that make sense in terms of service offerings are to either maximize the downstream tier, as this is what subscribers tend to value, or maximize the symmetry of the tier. If an operator wants to offer 10 Gbps downstream and 1 Gbps upstream then high split, UHS-300, UHS-396 and UHS-492 have the capacity to support the offering. Which variant is most supportable and future-proof depends upon the operator's network starting point as well as their forecast for future growth. Here the Cloonan formula, referenced earlier can help as a guide.

Serving Group Capacity >= Nsub*Tavg + K*Tmax_max

If a K factor of 1 is used for simplicity, a downstream Tmax_max of 10 Gbps can be used, leaving Nsub and Tavg. If the network is expected to provide sufficient capacity for 10 years, then an Nsub and Tavg must be forecast for that time. Using an example of 100 for Nsub and 20 Mbps for Tavg the serving group capacity must be greater than or equal to 12 Gbps as calculated below.

Serving Group Capacity >= 100*20 + 1*10,000 = 12,000 Mbps

Of the initial group of four variants, only high-split and UHS-300 meet this requirement. Doing the same exercise with the upstream, using a K factor of 1, an upstream Tmax_max of 1 Gbps and example values of 100 for Nsub and 2 Mbps for Tavg, the serving group capacity must be greater than or equal to 1,200 Mbps.

Serving Group Capacity >= 100*2 + 1*1,000 = 1,200 Mbps

Both high-split and UHS-300 meet this requirement, assuming that no capacity is lost due to OTA interferers or similar.

XGS-PON operators currently have 8 Gbps symmetric tiers in market, so if an operator wants to match the tier in the downstream while keeping the tier as symmetric as possible then UHS-684 or UHS-492 are likely to be most desirable. Using the same downstream numbers as above but substituting in 8 Gbps for the tier we arrive at a serving group capacity greater than or equal to 10 Gbps.

$$\text{Serving Group Capacity} >= 100*20 + 1*8{,}000 = 10{,}000 \text{ Mbps}$$

In this case UHS-684 does not have sufficient capacity to meet this requirement, but UHS-492 does. To determine the maximum upstream, we can reorder the Cloonan formula to calculate the upstream Tmax_max as shown below.

$$\text{Tmax\_max} <= (\text{Serving Group Capacity} - \text{Nsub}*\text{Tavg})/K$$

Using the UHS-492 serving group capacity of 3,409 Mbps we arrive at a Tmax_max of 3,209 Mbps, making the most likely highest upstream service tier 3 Gbps.

$$\text{Tmax\_max} <= (3{,}409 - 100*2)/1 = 3{,}209 \text{ Mbps}$$

UHS-396 is also a potential middle ground, where an operator can offer multiple Gbps upstream and a higher downstream than what XGS-PON can offer, such as 10/2 Gbps.

8.5. FTTH Upgrade

An operator may also decide to upgrade their HFC network straight to FTTH, and indeed some operators have announced that intention publicly. This strategy makes sense if the operator believes an FTTH upgrade is not significantly more capital intensive than an HFC upgrade, or that network growth is great enough and sustained over time such that an HFC network upgrade will not delay an upgrade to FTTH long enough to justify the investment.

Making some assumptions, we can calculate under what traffic growth rates and FTTH upgrade cost conditions the total cost of ownership (TCO) of an HFC upgrade is more desirable than an FTTH upgrade and vice-versa. The analysis assumes a serving group with 200 households with 50% penetration, an UHS-396 HFC upgrade with a cost of $100-$300 (based on publicly released estimates), Tavg of 5 Mbps and Tmax_max of 10 Gbps.

FTTH Upgrade Cost per Household Passed

Traffic Growth Rate	$ 300	$ 400	$ 500	$ 600	$ 700	$ 800	$ 900	$ 1,000
5%	HFC	HFC	HFC	HFC	HFC	HFC	HFC	HFC
10%	HFC	HFC	HFC	HFC	HFC	HFC	HFC	HFC
15%	FTTH	HFC	HFC	HFC	HFC	HFC	HFC	HFC
23%	FTTH	FTTH	HFC	HFC	HFC	HFC	HFC	HFC
32%	FTTH	FTTH	FTTH	HFC	HFC	HFC	HFC	HFC
50%	FTTH	FTTH	FTTH	FTTH	FTTH	FTTH	HFC	HFC

Figure 22 shows the results, where purple cells represent growth and cost conditions where an HFC upgrade provides a lower and thus more desirable TCO, and orange cells represent conditions where an FTTH upgrade has a lower TCO. In this example we can see that the combination of low FTTH upgrade cost and high growth favours an FTTH upgrade, while high FTTH upgrade cost and low growth favours an HFC upgrade.

FTTH Upgrade Cost per Household Passed

Traffic Growth Rate	$ 300	$ 400	$ 500	$ 600	$ 700	$ 800	$ 900	$ 1,000
5%	HFC	HFC	HFC	HFC	HFC	HFC	HFC	HFC
10%	HFC	HFC	HFC	HFC	HFC	HFC	HFC	HFC
15%	FTTH	HFC	HFC	HFC	HFC	HFC	HFC	HFC
23%	FTTH	FTTH	HFC	HFC	HFC	HFC	HFC	HFC
32%	FTTH	FTTH	FTTH	HFC	HFC	HFC	HFC	HFC
50%	FTTH	FTTH	FTTH	FTTH	FTTH	FTTH	HFC	HFC

Figure 22 – TCO Comparison Between FTTH and HFC Upgrades

There are many factors that determine the cost of an FTTH upgrade, but the most influential are whether the plant type is aerial or underground, ownership and regulatory rules around access to infrastructure, and labour costs having to do with installing fibre. While operator forecasts for traffic growth will vary, the trend is for slowing growth, as discussed earlier.

With generally high estimated costs for FTTH upgrades and low forecasted traffic growth, most operators would find themselves on the top right of the chart, making an HFC upgrade the optimal decision from a financial perspective. Within operator networks there are likely to be areas which can be upgraded at a lower cost, or traffic growth is higher than average due to demographic shifts or similar, and operators may decide to upgrade those areas directly to FTTH.

While from a financial perspective an HFC upgrade may represent the option with the best TCO, there are other reasons to choose an upgrade path. Operators may perceive FTTH as the only upgrade path that achieves technology leadership instead of asking whether subscriber needs can be met with an HFC network, and whether they will continue to be met over the expected lifespan of the network. Subscribers do not inherently ask for a technology solution; they ask for services. To the extent that they do ask for a specific technology, it tends to be driven by marketing activities.

There also exists confusion over what technology is most appropriate in new builds and what is capable of meeting the needs of subscribers. Thanks to the transition to all-IP services and improvements in CPE and CPE install practices, FTTH is the most future-proof technology, and at build time it is the most appropriate technology, assuming reasonable density of homes. This is why most HFC operators are building FTTH in new areas. This, in addition to the fact that many HFC operators are upgrading to FTTH add to the misleading messaging that FTTH is required to offer services that subscribers need, including traffic symmetry. Members of the US Senate recently urged the Federal Communications Commission (FCC) to change their broadband definition from 25 Mbps downstream by 3 Mbps upstream to 100 Mbps symmetric [11], a definition that would needlessly exclude HFC sub-split plants. Although an argument for 100 Mbps downstream can be made, it is difficult to envision what residential use cases require high-bandwidth symmetric services.

9. Conclusion

In order to offer higher-speed service tiers, cable operators will be obliged to upgrade their HFC networks. There are many upgrade paths and strategies to choose from, and each operator must decide what elements to prioritize. Strategies range from maintaining a sub-split network to minimize capital intensity to upgrading straight to FTTH. In between are strategies to upgrade to mid-split, high-split or ultra-high-split using DOCSIS 4.0 technology. The choice of diplex frequencies has implications on what tiers can be offered and how closely network capacity matches the needs of subscribers.

Abbreviations

ADSL	Asymmetric DSL
bps	bits per second
CM	Cable Modem
CO	Central Office
COVID-19	Corona Virus Disease of 2019
CPE	Customer Premises Equipment
CWDM	Coarse WDM
DAA	Digital Access Architecture
DOCSIS	Data Over Cable Service Interface Specifications
DS	downstream
DSG	DOCSIS Set-top Gateway
DSL	Digital Subscriber Line
DSLAM	DSL Access Multiplexer
DWDM	Dense WDM
FCC	Federal Communications Commission
FDD	Frequency Division Duplex
FTTC	Fibre-to-the-curb
FTTH	Fibre-to-the-home
Gbps	Gigabit-per-second
GHz	Gigahertz
GPON	Gigabit PON
HD	High Definition
HFC	Hybrid Fibre-Coax
Hz	Hertz
IP	Internet Protocol
IPTV	IP Television
K	QoE Factor
Mbps	Megabit-per-second
MHz	Megahertz
MoCA	Multimedia over Coax Alliance
MPEG	Motion Pictures Expert Group
NHL	National Hockey League
Nsub	Number of Subscribers
OLT	Optical Line Terminal
ONU	Optical Network Unit
OTA	Over the Air
OTT	Over-the-top
PHY	Physical Layer
PON	Passive Optical Network
QAM	Quadrature Analog Modulation
QoE	Quality of Experience
SCTE	Society of Cable Telecommunications Engineers
SG	Serving Group
SNR	Signal-to-noise Ratio
Tavg	Average Traffic at Busy Hour
TCO	Total Cost of Ownership

Tmax_max	Maximum Provisioned Traffic Rate
UHS	Ultra-high-split
US	upstream
VDSL	Very-high-speed DSL
WDM	Wavelength Division Multiplexing
XGS-PON	10 Gbps Symmetric PON
YoY	Year-over-year

Bibliography & References

[1] OECD Data - https://data.oecd.org/broadband/fixed-broadband-subscriptions.htm

[2] https://en.avm.de/guide/avm-explains-connection-types-dsl/

[3] https://seekingalpha.com/article/1922581-calix-the-local-telcos-favorite

[4] https://www.forbes.com/sites/kevinmurnane/2017/10/26/upgrading-to-4k-hdr-tv-how-far-you-sit-from-the-screen-is-critical/?sh=693155541801

[5] https://www.thebroadcastbridge.com/content/entry/2746/hevc-decoded

[6] *Maximizing Returns on the Path to DOCSIS 4.0*, Mike Darling; SCTE Cable-Tec 2021, SCTE

[7] https://www.bce.ca/investors/AR-2022/2022-bce-integrated-annual-report.pdf

[8] https://www.telus.com/en/about/investor-relations/reports/annual-reports/2022?intcmp=tcom_about_annual-reports_link_2022-annual-report

[9] *Predictions on the Evolution of Access Networks to the Year 2030 & Beyond,* Cloonan, Tom, Emmendorfer, Mike, Ulm, John, Al-Banna, Ayham and Chari, Santhana; SCTE Cable Tec 2014, SCTE

[10] *Network Capacity Options on the Path to 10G*, Karthik Sundaresan; SCTE Cable-Tec 2022, SCTE

[11] https://www.pcmag.com/news/senators-broadband-speed-minimum-should-be-100mbps-for-downloads-and-uploads

