Collision-Free Hyper-Speeds on the Bi-Directional FDX Highway

A Technical Paper prepared for SCTE by

Dr. Robert Howald
Fellow
Comcast
+1 (267) 398-8104
robert_howald@comcast.com

John Ulm
Engineering Fellow
CommScope
+1 (978) 609 6028
John.ulm@commscope.com

Saif Rahman, Comcast
Dr. Zoran Maricevic, CommScope

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

The promise of 10G is emerging, as Full Duplex DOCSIS (FDX) and Extended Spectrum DOCSIS (ESD) solutions make their way through laboratory testing to field trials towards first market launches. DOCSIS 4.0 consists of these two complementary technologies aimed at dramatically increasing upstream capacity and, correspondingly, upload speeds that customers will enjoy.

The FDX option for DOCSIS 4.0 is based on use of spectrum in the 108 MHz-684 MHz band (or subset of the band) for both downstream and upstream. Intuitively, signals overlapping in frequency and time interfere with one another. However, FDX has two key innovations that prevent this. Echo Cancellation (EC) technology removes potential interference where possible. Interference that cannot be removed by EC is *avoided* by controlling transmission timing within the DOCSIS scheduler. FDX sizes up the network and creates groups of users – "Interference Groups"(IG) – that should not access the same chunk of spectrum at the same time. This knowledge is incorporated into algorithms for allocating downstream and upstream resources.

Because the scheduler's job is to fairly and efficiently allocate time and frequency resources, the effect on traffic engineering arises. In this paper, a deep dive into the mathematical modeling and analysis, based on empirical data from real DOCSIS HSD systems, will be described. The analysis will show:

- How large can an IG be before it effects performance
- What are the implications to IG size and FDX service group size
- What are the relationships among common traffic engineering variables – bandwidth, speed, penetration of services and tiers, service groups size – with the IG element introduced
- How do the results impact network migration strategy

This pioneering analysis breaks new ground on traffic engineering of FDX systems, speaks to key aspects of N+x FDX systems, and promises to be a foundation for field implementation guidelines of DOCSIS 4.0 FDX.

2. DOCSIS 4.0 Full Duplex Overview

Figure 1 illustrates the essential spectrum goal of FDX– enabling significantly more upstream. More interestingly, these new FDX upstream bands (Red "FDX" in **Figure 1**) are also available for downstream! This is quite different than typical Frequency Domain Duplex (FDD) operation, the approach taken in DOCSIS 4.0 FDD. How can both downstream and upstream data exist in the same spectrum? There are two essential innovations in DOCSIS 4.0 that enable this.

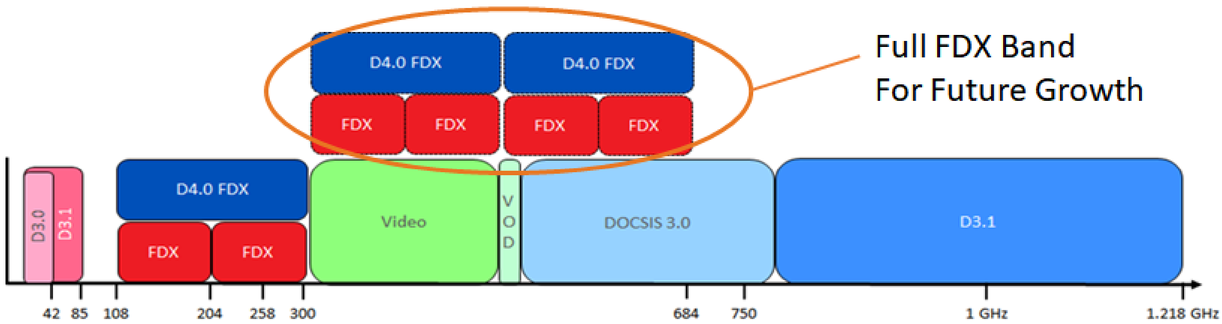

Figure 1 – Massive New Upstream Bandwdith by Sharing Downstream and Upstream in the Same Spectrum using DOCSIS 4.0 Full Duplex (FDX)

2.1. Key FDX Innovations

Although more upstream bandwidth is defined, it is the same 96 MHz OFDMA physical layer blocks that are defined in DOCSIS 3.1, except that the band over which they can operate is expanded. Thus, DOCSIS 4.0 leverages the power of the DOCSIS 3.1 PHY completely. Six additional 96 MHz blocks are added across the 108-684 MHz band, complementing an 85 MHz “mid-split” system.

Downstream and upstream can occupy the same band using a technology known as Echo Cancellation (EC). Echo Cancellation, in general, is a mature technology used in other telecom networks, such as xDSL and wireless. It has not yet been implemented in cable networks. The EC concept is very similar to those other applications, although the cable does introduce some new implementation challenges. EC is the first of the two critical innovations that power FDX.

The second key innovation is based on a fundamental architectural difference in cable systems compared to telco xDSL systems. Twisted pair telco networks are point-to-point connections from the DSL Access Multiplexer, or DSLAM, whereas HFC is a point-to-multipoint system. This logical architecture difference creates the need for another layer of innovation for FDX. This is the creation of Interference Groups (IGs) and Transmission Groups (TGs) for the scheduler to manage.

Figure 2 illustrates these innovations using a passive coaxial network (i.e. N+0) for simplicity. N+0 is NOT a requirement for FDX. In fact, one of the drivers for this paper is the evolution of FDX for N+x networks, whereas the DOCSIS 4.0 specification uses N+0 model types only as reference architectures. FDX-capable amplifiers will support FDX signals over N+x networks, allowing FDX to be implemented over a much wider swath of the footprint more quickly, and more of which is N+x than N+0. As we know from standard HFC networks, RF amplifier cascades impact end-of-line fidelity, and this is no different for FDX amplifiers. However, with amplifiers for FDX, an additional effect comes from the formation of IGs and TGs. As a result, an additional trade space created by N+x FDX networks are maximum speed tiers that can be supported for what amount of user penetration versus amplifier cascade depth. We will briefly discuss FDX amplifiers themselves in a subsequent section.

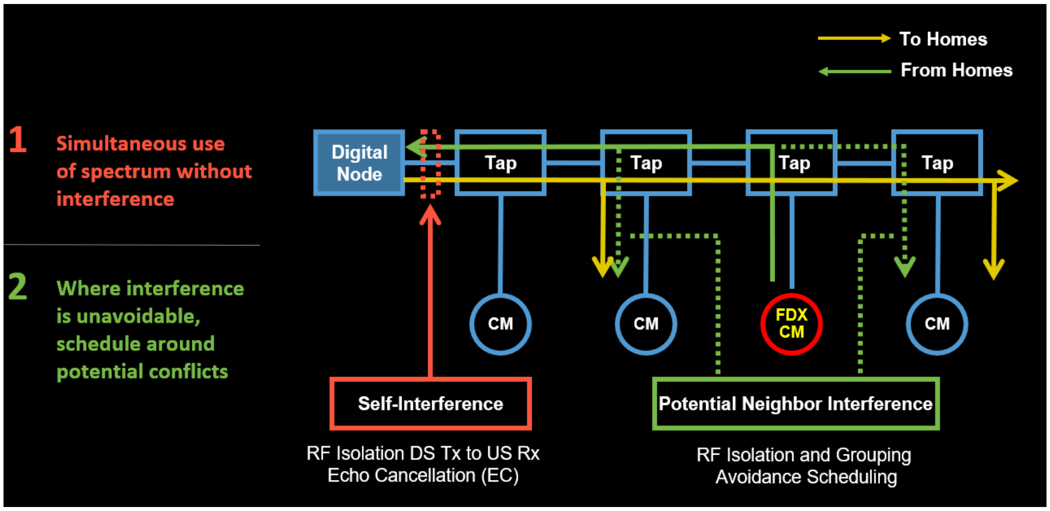

Figure 2 – Two Key New Innovations Power DOCSIS 4.0 FDX: Echo Cancellation and Interference Group / Transmission Group Formation

2.1.1. A Closer Look: Echo Cancellation

Referring to Figure 2, adding upstream signals at frequencies where downstream signals exist requires that the downstream signal be "subtracted" before the upstream (US) OFDMA receiver. This requires high RF isolation and strong EC of the much higher downstream signal that is reflected back into the US receiver. While the implementation details may be complex, the EC concept is a quite simple, and the digital signal processing (DSP) principles to build it are very mature. A simplified diagram illustrating the EC concept is shown in Figure 3.

The node downstream transmit signal will have some of its energy impose onto the US Rx simply through imperfect isolation characteristics of real hardware of an RPHY Node. It will also have some energy reflected back by imperfect RF interfaces as it travels down the coaxial plant, such as from the return loss of a tap, for example. These are the so-called "Echoes" that give the EC function its name. What is distinctive to EC for cable is the high cancellation required across a broad, multi-octave, bandwidth.

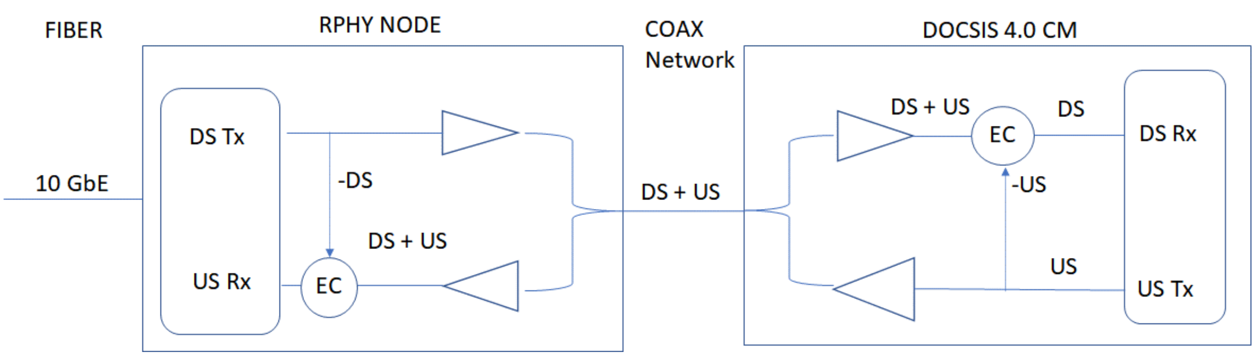

Figure 3 – Basic View of Echo Cancellation Concept

For more details describing Echo Cancellation and performance observed, please refer to **[6][7][8][9][10].**

2.1.2. A Closer Look: Interference Groups and Transmission Groups

As noted, the HFC network is a point-to-multipoint architecture. While an FDX modem knows its own upstream transmission, it cannot know that of his neighbor, and thereby has no way to "cancel" interference from a neighbor. It requires sufficient RF isolation among the homes sharing a coaxial RF segment to prevent FDX-band upstream users from interfering with a neighbor using that band for downstream. Unfortunately, RF isolation among homes cannot always be guaranteed to be high enough. Isolation relationships among homes are determined as part of FDX "sounding" process for DOCSIS 4.0 and DOCSIS 3.1 devices that support FDX-Light (aka FDX-L) functionality. A DOCSIS 3.1 CM that supports FDX-L becomes aware it is on an FDX-enabled network and can participate in the sounding process as a DS "measurer" device.

In situations without sufficient RF isolation, we instead call on the virtual CMTS (vCMTS) scheduler to avoid an interference scenario. Note that the vCMTS refers to the Comcast implementation of the CMTS function using commercial off-the-shelf servers hosting CMTS code and integrated with a DAA system based on a switched Ethernet architecture [7][8]. In an FDD system, the scheduler does not need to pay close attention to the relationship of downstream and upstream access to the coax. This changes in FDX. During FDX "sounding," the FDX system determines these isolation relationships. Potentially interfering users are lumped into "Interference Groups," or IGs. A logical set of IGs is called a Transmission Group (TG). Transmission Groups are created because not every IG needs to be treated independently – its overkill to do so, as we shall see later in this paper – and the vCMTS workload can be simplified by aggregating IGs into TGs. The scheduler assures that potentially interfering pairs are not accessing the same spectrum in the same time slot.

Because there are six OFDMA blocks, the vCMTS can service multiple IGs with uniform capacity and speeds by assigning different Resource Block Assignments (RBAs) to each IG. Figure 4 shows an example of how the 108-684 MHz FDX band might be allocated to simultaneously support a case with three TGs. These RBAs can adapt with time based on traffic and peak speed demand.

Note that the FDX band is not all of the DOCSIS spectrum available. Non-FDX DOCSIS 3.1 spectrum and DOCSIS 3.0 spectrum will also exist. Furthermore, not all of the FDX band needs to be assigned at all – more US bandwidth is allocated to FDX as peak upstream speed requirements increase.

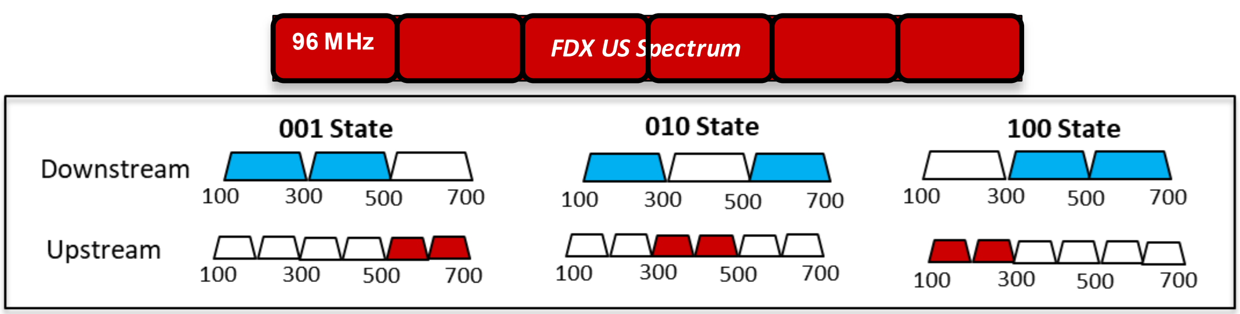

Figure 4 – FDX Resource Block Assignment and Sample Band Allocation "States"

2.2. FDX-Capable Amplifiers

The FDX system specifications were written using the assumption of an N+0 network. However, as mentioned previously, FDX is not technically limited to an amplifier-free plant and the specification does not prevent it. In fact, shortly after the FDX specifications were completed, a CableLabs Study Group was formed to evaluate methods for implementing amplifiers that support FDX. The foundational EC technology developed for the FDX RPD can be applied at any point in the network to manage overlapping spectrum, and this can include amplifiers. Of course, these are not traditional amplifiers, but a new class of device that includes digital signal processing (DSP).

An EC-based amplifier concept is shown in Figure 5. The nature of overlapping spectrum and gain in both directions creates a full-circle loop gain path in the FDX band. The EC must be capable of suppressing the FDX loop gain, such that the net gain around the path is < 0 dB across all frequencies to maintain a stable device. The EC must further be designed to act on the echo it is suppressing sufficiently that the aggregate residual echo noise, which becomes part of the amplifier's own noise floor, supports the US MER requirements effectively for DOCSIS 4.0, without introducing unacceptable MER degradation and subsequent loss of bandwidth efficiency.

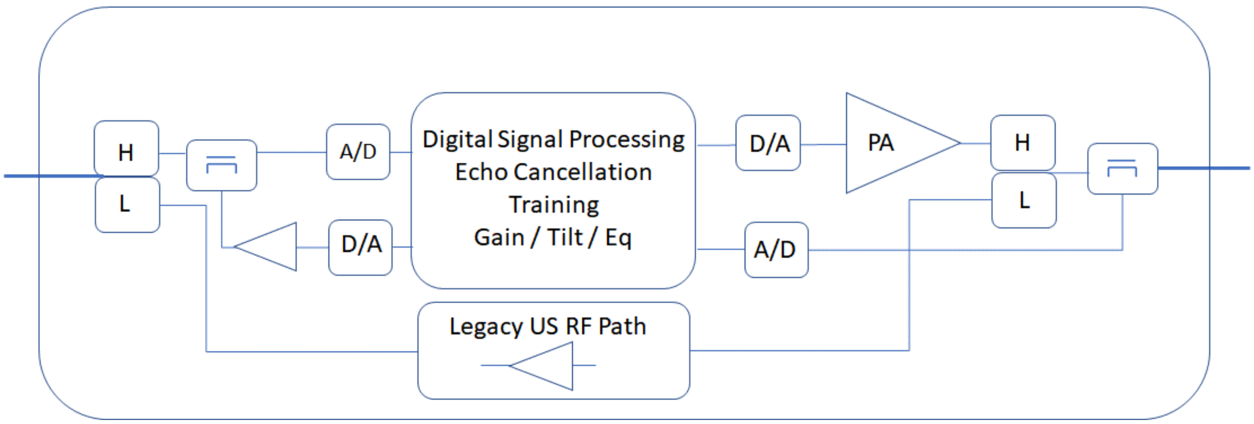

Figure 5 – Topology of a DSP-Based FDX Capable Amplifier

As noted previously, as an HFC amplifier cascade increases, noise contributions aggregate and the MER decreases. For the FDX amplifier, in the FDX band, there is an additional noise contributor in the form of residual echo. The amount of acceptable degradation due to the amplifier is a system engineering parameter that flows from performance specifications ultimately to the performance of the EC itself. A significant advantage for amplifier EC in comparison to an N+0 node output is that the levels on the DS Tx port are lower, and on the US Rx port higher. Thus, the DS to US level ratio is smaller, which is favorable for the EC.

While noise analysis is relatively straightforward for N+x with FDX, traffic engineering requires additional scrutiny due to the shared DS and US and the effect of amplifiers in expanding the size of an IG. Consider the N+1 system shown in Figure 6. When an amplifier is included, there is an expansion of Interference Group 4 (IG4) to the "south," or home-facing, side of the amplifier. These users become part of the last IG of the tap string before the amplifier. This is because of the limited drop-to-output isolation characteristics of today's taps. This parameter can be optimized for high drop-to-output isolation, but until there was FDX to consider, there was no reason to drive more aggressive specifications for this parameter.

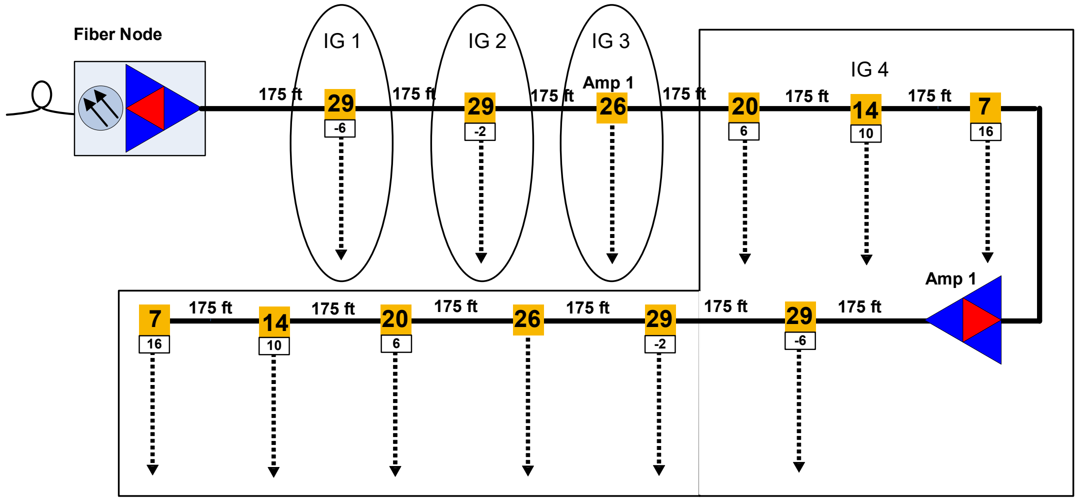

Figure 6 – Potential Interference Group (IG) Elongation due to an FDX Amplifer

3. FDX Traffic Engineering: An Operational Perspective

If we think about network segmentation triggers in current HFC networks, rules have been developed that link service group size, speeds, percent capacity utilization, and total capacity required. Then, with an awareness of device penetrations, this can be translated to DOCSIS 3.0 and DOCSIS 3.1 spectrum requirements and subsequently used to compare to thresholds that trigger a network augmentation.

Similar capacity-based analysis and empirical rule making will now apply to FDX, with one additional nuance. The FDX band, of course, can be allocated for both downstream and upstream. Infrequent bursts of peak speeds can therefore be called upon to service the downstream or the upstream, but they have a new shared resource dependency. The introduction of new terminology for FDX, in particular the terms "Interference Group" and "Transmission Group," created an unfortunate aura of mystery around how the FDX band operates. This new terminology placed outsized significance on what turns out to be, from a traffic engineering point of view, a relatively benign, and certainly manageable phenomenon. What we have with FDX is a relatively modest "joint access" twist on HSD traffic engineering that we otherwise have been doing successfully for an independent downstream and independent upstream for decades.

Through these decades, operators have learned how to aggregate subscribers in a way that efficiently uses available capacity of the downstream and upstream while still delivering a high Quality of Experience (QoE) to the customer. Generally, operators have created empirically-based guidelines based on the key traffic engineering variables in play – capacity, service group size, HSD penetration, maximum speed, distribution of tiers, D3.1 and D3.0 mix, and also typically including some allowance baked in to account for a projected Compounded Annual Growth Rate (CAGR). Network augment rules have been developed over time to govern node splits, spectrum re-allocation, and device mix policies.

Instead of independent statistics of DS and of US, we now have a portion of the spectrum where resources are shared. However, we already have mature, reliable traffic models for downstream and upstream users! In either case – FDX or FDD – a set of users are sharing finite bandwidth and time resources, and the network design is deliberately NOT a non-blocking architecture, or it would be massively overdesigned. Users "compete" for the finite resources in the DS and in the US. The vCMTS makes decisions on packets sent DS and grants allocated US based on a set of QoE criteria determined empirically, as noted previously. Any individual user's packets can be deferred for another's *as part of* *normal CMTS operations* – we just refer to it as *load balancing* and/or *fair scheduling*, and we developed rules and configuration settings for how to manage this effectively.

Again, for FDX, the only difference is there are users of DS and US both looking to access bandwidth and time resources, and the CMTS scheduler must accommodate them both. The mathematics and modeling of this joint access situation are in the sections to follow. They will underscore that an irrational "fear" of the Interference Group – seeking to minimize its size to a single Tap or two – was developed without any analysis or evidence to tell us just how big of a TG is too big. It just "felt" bad that US could interfere with DS unless we did something about it, that the two would both be looking to access resources, and one could be denied access in some mini-slot in deference to another. Of course, this happens ALL OF THE TIME in an independent DS or US!

A helpful *qualitative* way to think of the FDX band is to consider the DS users of the FDX band (D3.1+D4.0 CMs) as a service group (SG). Then, the US FDX band users can be considered a smaller set (smaller because D4.0 users < (D3.1 + D4.0) – and largely so for many years) of "new" users looking to, less frequently (because US utilization << DS utilization), access FDX band frequency and time resources. These FDX US "users" just happen to have the same MAC address and billing address of some of the DS users – which of course traffic engineering does not know nor care about. Because the upstream average and peak utilization is so much lower than downstream, and because the aggregate US traffic of a SG fits comfortably into the 85 MHz legacy band for many years of CAGR, except for brief, low-likelihood moments of peak speed bursts >300+Mbps, the FDX Band will be dominated by its use as a downstream channel. If the band was entirely dedicated to the upstream as it would be in an ultra-high split FDD system, it will be *idle* spectrum the vast majority of the time.

In summary, the DOCSIS 4.0 FDX Working Group (note – author's included!) studied, in depth, the RF side of IGs and TGs, how to determine them, and incorporated it all into the specification. BUT – the working group did NOT do the math on practical implications of this phenomenon to the traffic engineering of the FDX band. We are now doing that Math! THE traffic engineering questions become:

1) How large can an IG be before there is an impact to the customer experience?

2) What service speed / IG size / spectrum rules exist when downstream and upstream traffic engineering become co-mingled in FDX?

These key questions are addressed in the traffic engineering analysis and modeling to follow. Before diving right into to answering these questions, however, we first take a step back to build up the foundation of burst traffic statistical behavior.

4. Network Capacity Planning

Determining the required amount of capacity needed for a service group (SG) is critical for providing customers with the appropriate quality of experience (QoE). Figure 7 shows an example of a SG upstream (US) capacity usage over a 16-minute window sampled at 1-second intervals. If an operator samples the bandwidth (BW) usage once every 16-minutes, then the purple line represents that average bandwidth for that measurement interval. This is a very useful datapoint but insufficient to determine the required capacity for the SG.

Sampling at 1-minute intervals would capture some of the variations, or ripples in the system but would still miss the many spikes which are individual modems bursting on top of the average BW. The gray line in the figure shows the high-water mark from 1-minute samples.

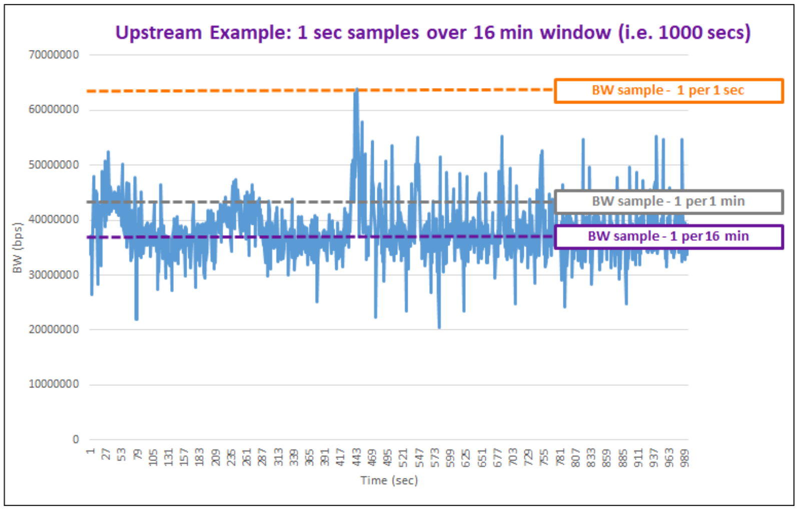

Figure 7 – Example of Upstream Capacity Usage

If an operator could sample at 1-second intervals or faster, then they could more closely determine the required SG capacity. The 1-second high-water mark is shown by the orange line. But sampling at these rates is not feasible in real systems, so another method is needed.

4.1. The "Basic" Traffic Engineering Formula

The CommScope (formerly ARRIS) team has been providing industry leading research in traffic engineering for many years which was most recently highlighted in [12]. The network capacity analysis in [12] provides an insight into how to calculate the SG capacity requirements. Previously, [4] introduced traffic engineering and QoE for broadband networks. From there, the paper develops a relatively simple traffic engineering formula for service groups which is easy to understand and useful for demonstrating basic network capacity components. Some additional references of note include [3],[5],[11],[13],[14], and [15].

The "Basic" formula shown below is a simple two-term equation. The first term (Nsub*Tavg) allocates bandwidth capacity to ensure that the aggregate average bandwidth generated by the Nsub subscribers can be adequately carried by the service group's bandwidth capacity. The first term is viewed as the "DC component" of traffic that tends to exist as a continuous flow of traffic during the peak busy period. The growth rate of Tavg has seen much research. Some new revelations are discussed in [11].

"CLOONAN'S CLASSIC CAPACITY EQUATION" Traffic Engineering Formula:

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

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 $0 \leq K \leq$ infinity, but typically $1.0 \leq K \leq 1.2$
Tmax_max is the highest Service Tier (i.e. Tmax) offered by the MSO

There are obviously fluctuations that will occur (i.e. the "AC component" of traffic) which can force the instantaneous traffic levels to both fall below and rise above the DC traffic level. The second term (K*Tmax_max) is added to increase the probability that all subscribers 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 service tier being offered to subscribers.

In previous papers [1], found that a K value of ~1.0 yields acceptable and adequate QoE results. [4] provides simulation results that showed a value between K=1.0 and 1.2 provides good QoE results for a service group of 250 subscribers. Larger service groups (SGs) would need even larger values of K while very small SGs might use a K value near or less than 1.0.

4.2. The "Modified" Traffic Engineering Formula

Over time, it was discovered that the optimum value for K would vary based on all the inputs: Nsub, Tavg and Tmax_max. Some of these limitations were noted in [13] along with some refinements to the basic formula above. This resulted in the following which is still algebraically equivalent to the basic formula:

Modified "CLOONAN'S CLASSIC CAPACITY EQUATION" Traffic Engineering Formula:

$$C \geq (Nsub*Tavg) + (K-1)*Tmax\_max + Tmax\_max \qquad (2)$$

The subtle change is that there are now three main components to the traffic engineering formula:

1. Peak Busy Period Average Consumption (i.e. Nsub*Tavg)
2. Peak Busy Period Ripple (i.e. (K-1)*Tmax_max)
3. Headroom for maximum Service Tier Burst (i.e. 1 * Tmax_max)

Figure 8 shows how the modified formula maps to the US capacity usage example given above. The basic formula might have used a value of K=1.2 in this example. This is now broken into a burst component equal to Tmax_max plus a ripple component that is estimated by 20% * Tmax_max.

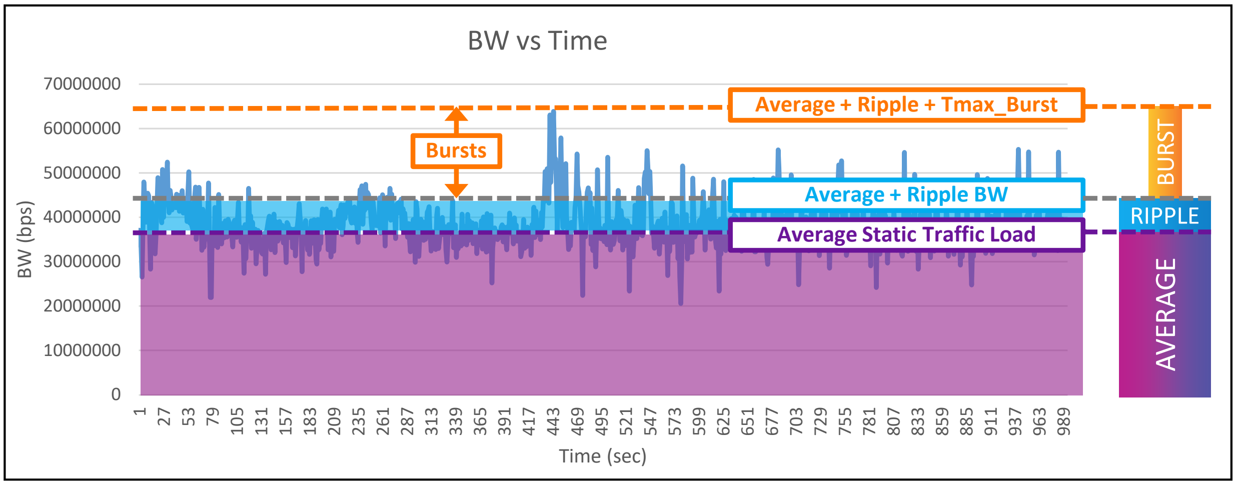

Figure 8 – Mapping Traffic Eng Formula to US Capacity Usage Example

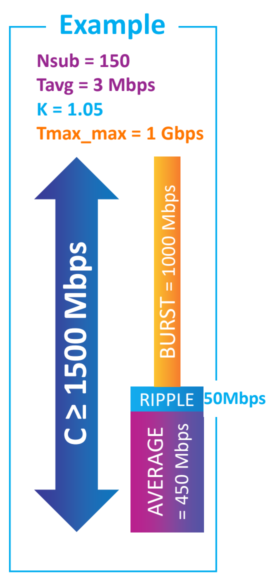

Figure 9 – Network Capacity Example for 1G Service Tier

Figure 9 shows a network capacity example for an operator who wants to support a 1 Gbps downstream (DS) service tier on a SG with 150 subs and an average peak consumption of 3 Mbps per subscriber. This requires at least 1500 Mbps of SG capacity. This includes a ripple component of 5% * Tmax_max which is 50 Mbps. Note that in this example, the ripple component as a function of Tmax_max is much smaller than the initial K=1.2 values because Tmax_max is significantly higher.

The Peak Busy Period Average Consumption and maximum Service Tier Burst components are well known and easily obtained. Traffic engineering research has since focused on quantifying the Peak Busy Period Ripple component as it is impacted by all inputs, not just Tmax_max.

4.3. Determining Service Group QoE based on Probabilities

It became apparent that quantifying the SG subscribers QoE needed to focus on the probabilities for SG capacity. And to predict behavior across different SGs with different parameters, a network capacity transmit model for individual subscribers was necessary.

4.3.1. Individual Subscriber Bandwidth Probabilities

Data was collected across an entire CMTS for multiple years during prime-time evening hours. Every packet was captured; associated with a particular modem and service tier; and timestamped down to a millisecond. Big data analytics was then used to find the patterns across the different types of subscribers.

It was discovered that all subs had very similar patterns. Figure 10 shows a very typical sequence that occurs frequently. Around 70% to 80% of internet traffic these days are video which pre-dominantly uses adaptive bit rate (ABR) encoding. The ABR streaming (#1 in figure) is a “grassy” region where traffic is limited by the video resolution (e.g. 5-10 Mbps). A second ABR starts later (#5 in figure) and begins with a burst to pre-fill its video buffers. This burst size is relatively small, limited by the video rate and buffer depth.

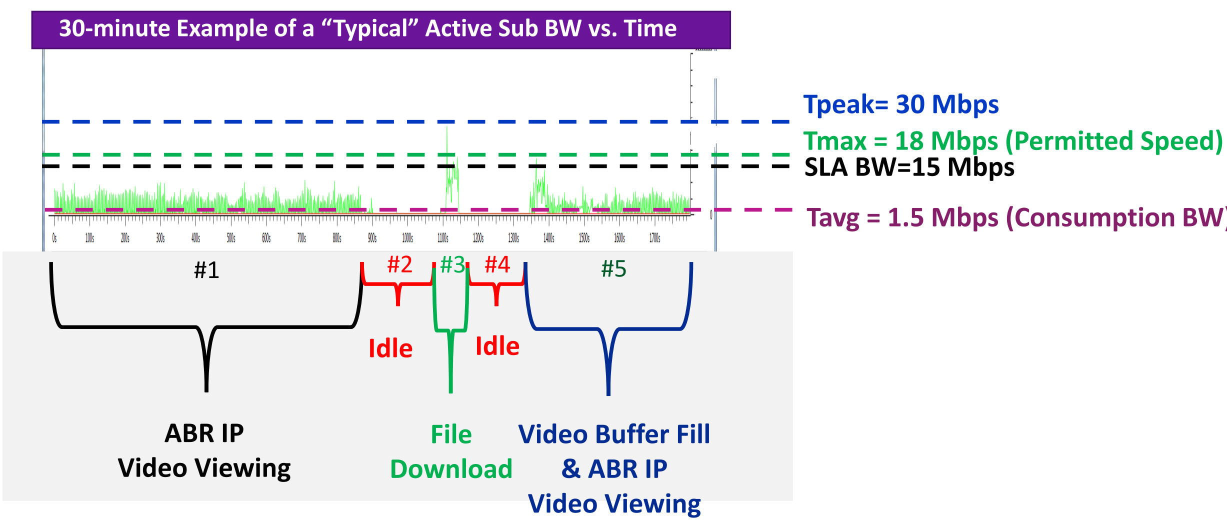

Figure 10 – Typical Subscriber Traffic Scenario

Sandwiched between these ABR streams is a file download. For a short instance, it may burst above the modem’s Tmax rate before the DOCSIS token bucket algorithms kick in to limit the transfer to its Tmax value. In today’s world, a 12 MB file download might represent a couple high resolution pictures or a PDF/Word/Powerpoint document. A modem with a 100 Mbps tier might download this in ~1 second while a 1 Gbps tier modem only needs 0.1 seconds.

The Big Data Analytics then determined the BW transmit probabilities for each identified type of subscriber using 1-second windows. It was found that using a window too fast (e.g. milliseconds) made it difficult or impossible to see the patterns while too slow (e.g. minutes) caused the bursts as shown above to be missed. Figure 11 shows the probability mass function (pmf) for a subscriber with a 1 Gbps tier and an average peak consumption of 3 Mbps. [note – the pmf is just a discrete version of the probability density function (pdf).]

Using the 12 MB file download example above, the 1G sub transfers 96 megabits in 0.1 seconds. Assuming the remainder of the 1-sec window is idle, then that burst would correspond to 96 megabits in a 1-sec window below even though the instantaneous rate is 1 Gbps.

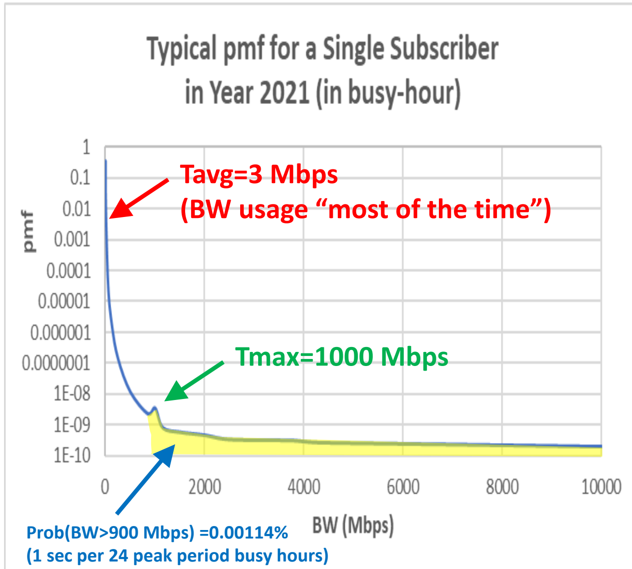

Figure 11 – Single Sub BW DS Transmit Probabilities, 1G @ 3 Mbps, 1-sec windows

Being a log scale shows how quickly the probabilities drop for large bursts. Looking at the area under the curve >900 Mbps gives us a probability of 0.00114% which equates to this occurrence of a sustained burst >900 Mbps happening for 1 sec for every 24 peak period busy hours. With ~3 peak period busy hours per day, this means 1 sec for every 8 days.

The Big Data Analytics also found patterns as the Tavg and Tmax are varied from subscriber group to subscriber group. This allows subscriber behavior to be predicted for future Tavg and Tmax values. Figure 12 shows the pmf for a hypothetical user in 2030 with a 4G tier consuming 15 Mbps during peak. Not only does the "Tmax bump" get pushed to the right, but the probability tail gets pulled lower too. The probability of a burst >3600 Mbps is now an order of magnitude smaller at 0.000117%. This maps to 1 sec every 240 peak period hours which might only be 1 sec every 80 days.

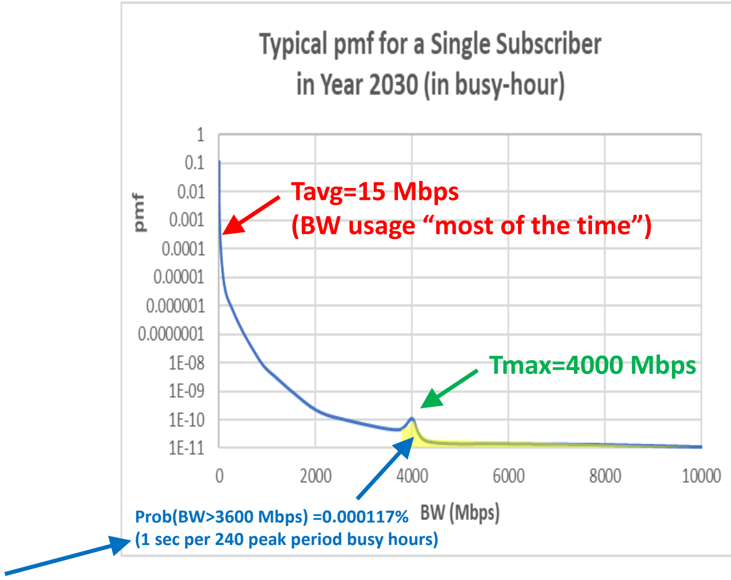

Figure 12 – Single Sub BW DS Transmit Probabilities, 4G @ 15 Mbps, 1-sec windows

In the FDX Amplifier Working Group previously described, an example was given for probabilities for subs with a 1 Gbps speed tier and a 2 Mbps Tavg:

- Probability (a single random sub creates a DS burst > (0.9*Tmax) in a 1-sec sample window) = 3.4e-6 = 1 second out of every 3.4 days
- Probability (50 subs sharing BW create a single DS burst > (0.9*Tmax) in a sample window) = (3.4e-6)*50 = 1 second out of 1.6 hours
- Probability (2 subs simultaneously create a DS burst > 2*(0.9*Tmax) in a sample window) = (3.4e-6)^2 = 1 second out of 2743 years

As can be appreciated, these bursts to >90% of Tmax are relatively rare.

4.3.2. Service Group Subscriber Bandwidth Probabilities

Once single subscriber BW pdf's were in place, it is then possible to create a SG BW pdf to analyze the behavior at the SG level. An example of the SG BW DS probabilities is shown in Figure 13.

This is the output of a Monte-Carlo simulation with 100K trials. It is a SG that consists of 128 subs, all of which have a 1G DS service tier. The Tavg = 15 Mbps which represents a time later this decade.

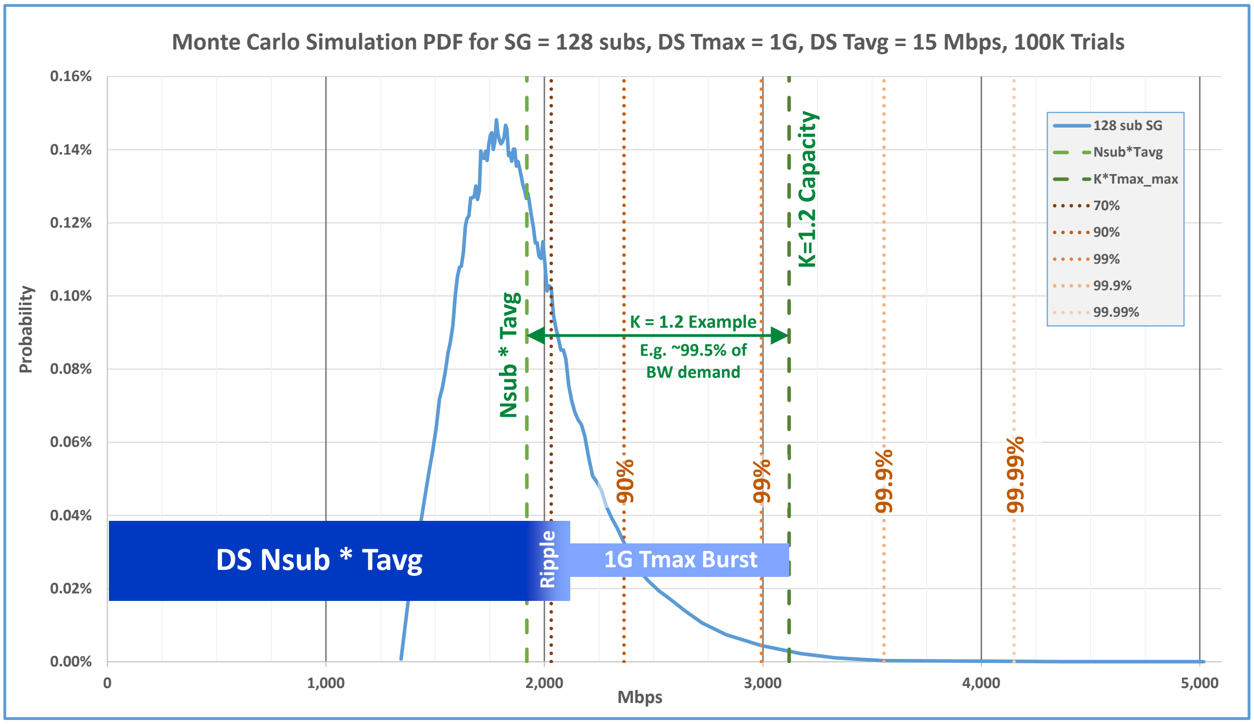

Figure 13 – Network Capacity Example for 1G Service Tier

Note the asymmetry in the curve – there is a much longer tail to the right. Various cumulative distribution function (CDF) probability thresholds are shown (i.e. 90%, 99%, 99.9%, 99.99%) to provide an insight as to the probability the SG BW hits different capacities.

The blue candlesticks in the figure show how the traffic engineering formula overlays the SG PDF. In this instance, the formula said 3,120 Mbps is required. Note that this is sufficient for ~99.5% of the peak busy period time. For the remaining 0.05% of the time, the buffers may be temporarily filled, and some latencies introduced. Note that 0.05% represents ~50 seconds out of every evening. For most of those 50 seconds, the delays should be insignificant and not noticed by users. This is normal network behavior and is why many users can share a single network pipe.

The remainder of the document uses these candlesticks often. Please keep in mind how low the probabilities are for their burst regions as shown above.

4.4. FDX Traffic Eng for Large SG with Overlapping Upstream + Downstream

The FDX Amp use case is leveraging existing plant and may contain many 100’s of subscribers. When FDX was conceived, this was not envisioned as a viable use case. In the early days, it was thought that there would be significant overlap between the US and DS in the FDX band. This is shown in Figure 14.

The more that the DS and US peak period average consumption overlaps, then the more important it becomes to have multiple FDX transmission groups (TG).

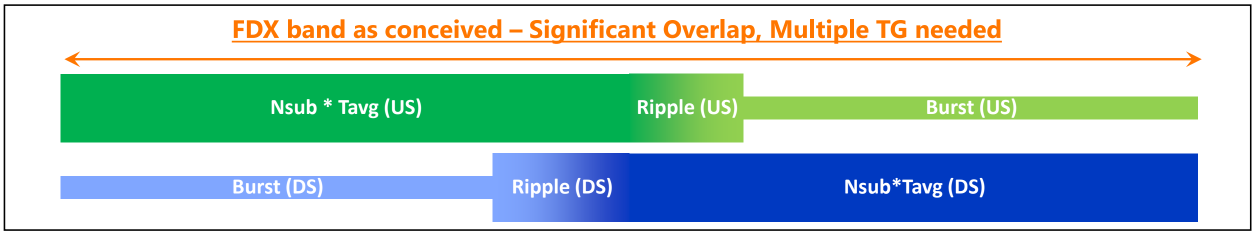

Figure 14 – FDX Band as conceived, significant overlap, Multiple TG needed

What has subsequently come to light is that the multi-gigabit burst region is dominating our traffic engineering formula. This means that the burst region (i.e. Tmax_max) is much larger than the consumption component (i.e. Nsub*Tavg). Figure 15, drawn to scale from one of our upcoming scenarios, shows how only the DS + US burst regions overlap in the FDX band. Most of the US traffic fits below 85 MHz while most of the DS traffic stays in the dedicated DS spectrum above 684 MHz.

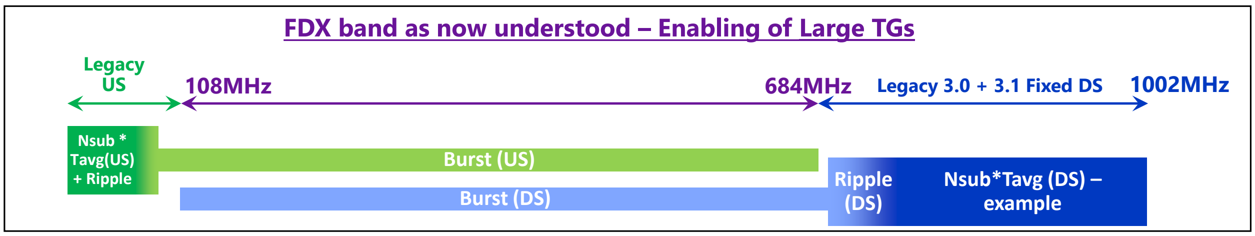

Figure 15 – FDX Band as understood now, only burst overlap, Single TG needed

In one of our examples, the US stay below 85 MHz for 99.8% of the time and only needs to access the FDX band for 0.2% of the time. Meanwhile, the DS stays above 684 MHz for 99.9% of the time and only needs the FDX band for 0.1% of the time. So, the FDX band is basically providing the burst bandwidth (i.e. 5G DS, 4G US in this case) and the probabilities of US + DS bursts overlapping is microscopic.

There are certain SG capacity requirements that must be met in any system, whether it is FDX or not. These include:

1. DS SG Capacity ≥ Nsub(SG) * DS Tavg + DS Tmax_max + DS Ripple
2. US SG Capacity ≥ Nsub(SG) * US Tavg + US Tmax_max + US Ripple

These two traffic engineering (TE) conditions are shown pictorially in Figure 16**.**

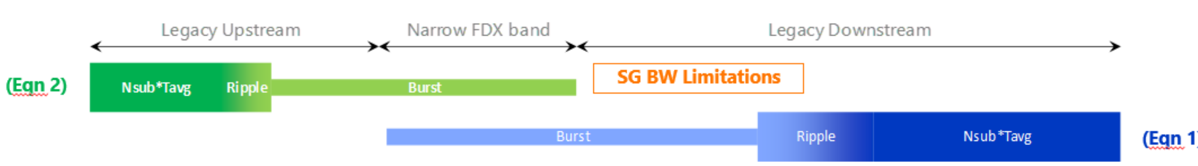

Figure 16 – FDX Band Size – SG Limits

In an FDX world, that means that the DS capacity above 684 MHz plus the DS capacity in the FDX band must be sufficient to meet TE condition #1. This must be met, totally independent of US requirements. Also, the US capacity below 85 MHz plus the US capacity in the FDX band must be sufficient to meet TE condition #2. This must be met, totally independent of DS requirements. In this respect, an FDX network's capacity requirements is no different than any other network.

However, FDX introduces a new twist with a TG sharing US + DS spectrum inside the FDX band. *Our traffic engineering analysis makes a basic assumption that it is acceptable for US + DS burst regions to overlap within a TG, but it is NOT acceptable for one burst region to overlap with the other's consumption region* (i.e. Nsub*Tavg + Ripple). It assumes that usage in the consumption region is too frequent to effectively share FDX BW.

For FDX TG, there are now two additional TE conditions that must be met:

3. DS TG Capacity ≥ Nsub(TG)*DS Tavg + DS Tmax_max + DS Ripple – excess US consumption
4. US TG Capacity ≥ Nsub(TG)*US Tavg + US Tmax_max + US Ripple – excess DS consumption

These are shown in Figure 17:

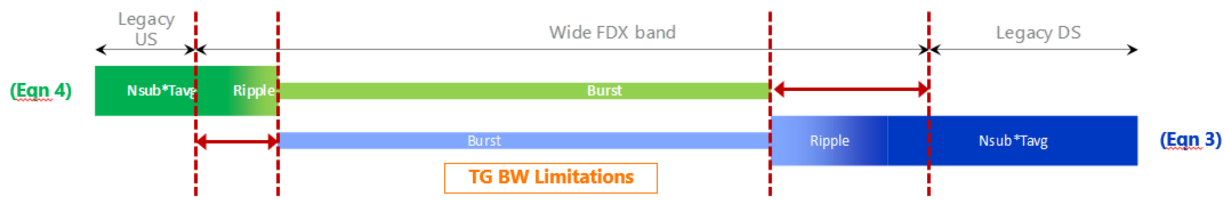

Figure 17 – FDX Band Size – TG Limits

Normally, the consumption BW is kept in the fixed legacy spectrum. This is below 85 MHz for the US and above the top of the FDX band (e.g. 684 MHz) for the DS. As long as these conditions are met, then neither of the TG conditions kick in.

The DS TG Capacity, TE condition #3, only comes into play when the US consumption (i.e. US Nsub*Tavg+Ripple) exceeds the 85 MHz capacity. As this excess US consumption fills into the FDX band, it effectively takes away from FDX BW that the DS can use within this TG.

Similarly, the US TG Capacity, TE condition #4, only comes into play when the DS consumption (i.e. DS Nsub*Tavg+Ripple) exceeds the dedicated DS spectrum capacity above the FDX band (e.g. 684 MHz). As this excess DS consumption fills into the FDX band, it effectively takes away from FDX BW that the DS can use within this TG.

5. FDX Network Capacity Modeling for Large SG with FDX Amps

Business decisions about the timing and pace of introducing multi-gig symmetric speed tiers will govern the pace of the migration of the network. There is uncertainty in the timing of launching these new tiers, the pace at which they roll-out nationally, and in the traffic growth to anticipate what will have occurred by the time they are launched. Therefore, a range of spectrum scenarios and target speeds were teed up for evaluation to cover what today appears to be realistic representative situations for FDX launch and years into the future of its use.

5.1. FDX Use Case Overview

5.1.1. Traffic Growth vs Time Considerations

Table 1 identifies two traffic utilizations that can apply for four different growth scenarios based on CAGR assumptions. If the CAGR is 30% in 2023 and persists,, then, as shown in the table, an average downstream user at peak busy hour will consume approximately 7-8 Mbps downstream and 500 kbps upstream. That same DS/US ratio will also be approximately correct in 2026 if the CAGR is instead 15%.

Similarly, if the 30% or 15% persists through 2026 or 2029, the DS and US will be approximately 15 Mbps and 1 Mbps, respectively. Of course, the further out the model is taken, the more uncertain it is, as CAGR typically wobbles year-to-year. However, in recent years, with the exception of the COVID-induced spike which is still settling, residential CAGRs have dropped from the 50% range of the 2010 time frame through the low 30% range pre-COVID, and it continues to decline. The upstream had been flattening out, decreasing to the <15% range pre-COVID, and appears to is returning to that range.

With these utilization boundaries, we have run models to draw out the acceptable TG size for speed tiers ranging from 1 Gbps to 5 Gbps for US and DS. We have also added some potential higher speed DS projections for 6 Gbps and 7 Gbps. Note that the speeds possible have a dependency on how efficient the bandwidth can be used. The models therefore also show the TG threshold range as a function of achievable QAM format for DS and US. Minor variations of QAM efficiency are expected, just as they are in HFC, as RF amplifiers incrementally reduce DS and US signal fidelity (MER) at the respective receivers. These fidelity losses translate to relatively minor variations in capacity. For example, if 4096-QAM is unachievable, but 2048-QAM is, then the capacity loss is about 8.3%.

Table 1 – CAGR Effects on per-user Peak Busy Hour (pbh) Average Utilization

Near term (2023) w aggressive CAGR (30%) or longer term (2025) with modest CAGR (15%): **7 Mbps DS / 500 kbps US**

Longer term (2026) w aggressive CAGR (30%) or even longer term (2029) with modest CAGR (15%): **15 Mbps / 1 Mbps**

			Network Migration Plan Year							
			1	2	3	4	5	6	7	8
			2022	**2023**	**2024**	**2025**	**2026**	**2027**	**2028**	**2029**
	CAGR %	**Tavg**								
DS	**15**	**4.5**	5.18	5.95	6.84	7.87	9.05	10.41	11.97	13.77
	20		5.40	6.48	7.78	9.33	11.20	13.44	16.12	19.35
	25		5.63	7.03	8.79	10.99	13.73	17.17	21.46	26.82
	30		5.85	7.61	9.89	12.85	16.71	21.72	28.24	36.71
US	**15**	**0.3**	0.35	0.40	0.46	0.52	0.60	0.69	0.80	0.92
	20		0.36	0.43	0.52	0.62	0.75	0.90	1.07	1.29
	25		0.38	0.47	0.59	0.73	0.92	1.14	1.43	1.79
	30		0.39	0.51	0.66	0.86	1.11	1.45	1.88	2.45

5.1.2. Spectrum Allocation Considerations

There is a time element to spectrum allocation as CAGR and new speed tiers – forcing functions for new blocks of spectrum that must be large enough to deliver the highest speed tier, with margin – require continuous maintenance of the channel maps. This has been an HFC story for decades: managing spectrum as HDTV took hold, alongside VOD and niche/specialized video channels grew even as HSD CAGR was consistently high. This ultimately led to the elimination of inefficient analog carriage altogether, as well as technologies such as Switched Digital Video (SDV) to perform the spectrum balancing act.

The contemporary version of this balancing act is DOCSIS 3.1 vs less than DOCSIS 3.1 devices. Operators are looking to add more spectrum that supports DOCSIS 3.1, to take advantage of the increased bandwidth efficiency (reminder note: DOCSIS 4.0 uses the same DS and US OFDM and OFDMA technology), and to reduce video single carrier QAM as channels move to IP delivery. The IP video Nirvana is removing single carrier QAM video altogether, but more realistically a long term scenario is reducing it to a small block that is sufficient to support the segment of the customer base that relies on this service.

Referring to Figure 18, five spectrum allocation cases are shown. The five are broken into two categories, near term (scenarios 1-3) and longer term (scenarios 4-5). Scenarios 1-3 will use as their peak busy hour average usage for DS and US as 7 Mbps and 500 kbps, respectively per the above. Scenarios 4 and 5 will use the peak busy hour average utilization for DS and US as 15 Mbps and 1 Mbps, respectively.

Scenario 3 is based on near term CAGR values, but presumes that an All-IP conversion has taken place – i.e. no video single carier QAMs. This scenario is to serve more as a boundary, "what's possible" scenario under these particular average DS and US utilizations. Scenario 5 is also an All-IP conversion, in a longer term and potentially practical time frame, for comparison, using the increased utilization that would occur with the CAGR values described in the model descriptions above

The near term scenarios 1 and 2 show a spectrum still heavily weighted by video single carrier QAM spectrum – 48 slots to be exact. There is also significant bandwidth set aside for DOCSIS 3.0 (168 MHz) to support the still large volume of D3.0-only modems in the network. While the balance is rapidly shifting, as of this writing, there are still more D3.0 modems in the network than D3.1 modems at Comcast.

Scenario 1 can be viewed as an initial launch scenario, supporting 2 Gbps symmetric speed. Scenario 2 provides on additional OFDMA block that enables 3 Gbps symmetric services. And again, scenario 3 is a "what if" under these lighter DS and US traffic loads.

Scenario 4 and 5 presume that significant harvesting of video single carrier QAM and D3.0 modems has taken place, and all or most of the spectrum is DOCSIS 3.1 enabled DS and the FDX US band is maximized to 684 MHz. In these longer term scenarios, access to higher symmetric speeds – 4 Gbps and 5 Gbps – is evaluated, as well as DS speeds above 5 Gbps. In the DS, there is more total spectrum to allocate – out to 1218 MHz – than in the DOCSIS 4.0 upstream.

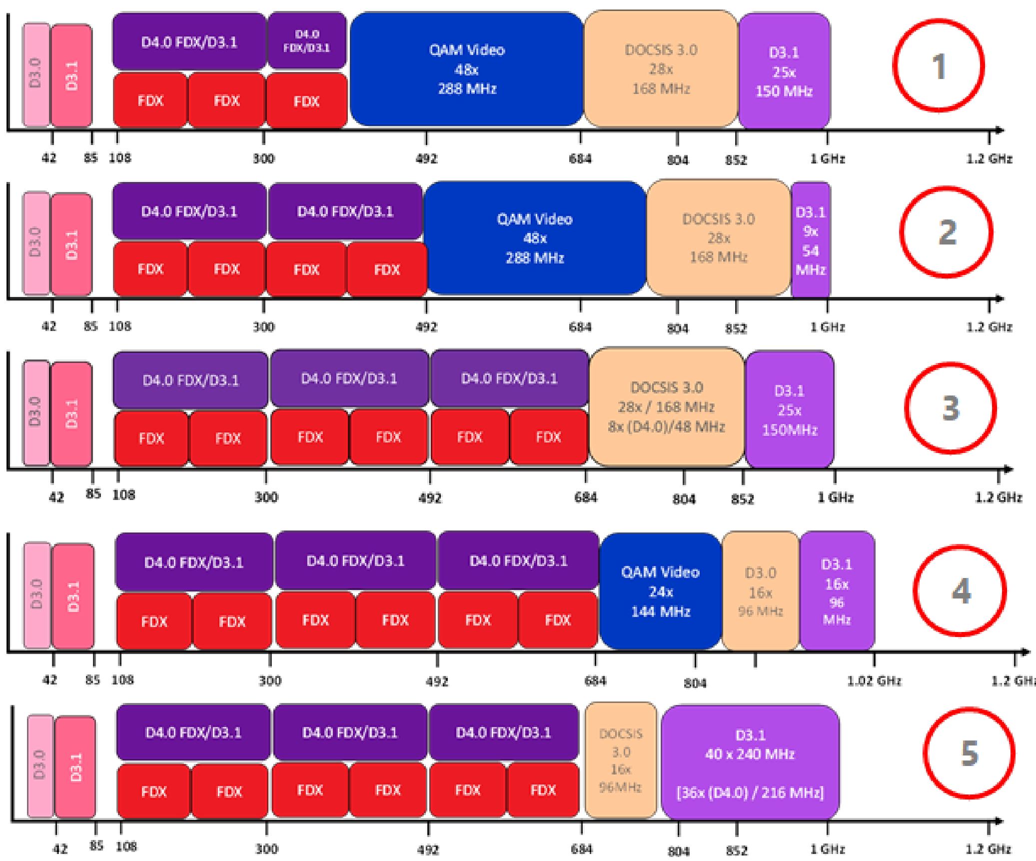

Figure 18 – Example Phases of Spectrum Migration vs Time

5.1.3. Network Capacity Modeling Assumptions

For our network capacity modeling, the following inputs were used in all scenarios:

- DS modulation = 4096-QAM for dedicated DS spectrum above FDX band
- Maximum DS frequency = 1002 MHz
    - Roll-off region not considered for this analysis, but could help
- DS modulation = 1024- to 4096-QAM inside the FDX band
- DS OFDM parameters: 1.25 usec Cyclic Prefix, 8K FFT
    - 4096-QAM => 9.70 bps/Hz

  - 1024-QAM => 8.08 bps/Hz
- US fixed capacity below 85 MHz = 450 Mbps
  - Equates to 4 x 6.4 MHz SC-QAM + ~46 MHz OFDMA @ 1024-QAM
- US modulation = 64- to 1024-QAM inside the FDX band
- US OFDMA parameters: 1.25 usec Cyclic Prefix, 4K FFT
  - 1024-QAM => 7.70 bps/Hz
  - 256-QAM => 6.16 bps/Hz
  - 64-QAM => 4.62 bps/Hz
- SG sizes are varied from 32 to 384 subs
- Maximum service tier (Tmax_max) varied to fit up to 5G DS, 4G US
  - Every subscriber (i.e. 100%) takes the maximum service tier

5.2. FDX Use Case 1 – Nearer Term Years 2023-25

Case 1 assumes a DS Tavg = 7 Mbps with US Tavg = 0.5 Mbps. Three scenarios were considered: two with legacy video and one with 100% video over IP (IPTV) – meaning no QAM video carriers, all have been reclaimed for DOCSIS 3.1.

For our network capacity modeling, a SG CDF was created to allow us to determine the probability of various capacity thresholds. The 70th percentile point was used to effectively estimate the consumption (i.e. Nsub*Tavg) plus the ripple components.

Many folks are still familiar with the original traffic engineering formula using the K-value. To help relate our network capacity modeling results, Table 2 shows the effective K-value that would generate very similar results. As mentioned earlier, the optimum K-value varies with Nsub, Tavg and Tmax_max. The Tavg is fixed for Case 1, so the table shows the range of the other inputs.

For all US combinations, the service tier (Tmax_max) is >1,000 times larger than the Tavg. Thus, the effective K-value is extremely small, 1.025. For the DS, the Tmax_max to Tavg ratio isn't nearly as wide. For 4G & 5G tiers, the effective K-value is 1.05. For 3G tiers, the effective K-value is 1.05 until the SG size reaches 384 subs where it increases to 1.1. For 2G tiers, the effective K-value is 1.05 until the SG size reaches 256 subs where it increases to 1.1.

Table 2 – Effective K values used for Case 1

K-value Used	Original K		QoE modeling – Effective K			
Nsub/SG	**DS**	**US**	**2G DS**	**3G DS**	**4G,5G DS**	**All US**
32	1.2	1.2	1.05	1.05	1.05	1.025
64	1.2	1.2	1.05	1.05	1.05	1.025
128	1.2	1.2	1.05	1.05	1.05	1.025
256	-	-	1.1	1.05	1.05	1.025
384	-	-	1.1	1.1	1.05	1.025

5.2.1. Case 1 with 288 MHz Legacy Video, 396 MHz FDX Band

This scenario has 288 MHz of Legacy Video spectrum and the FDX band is limited to 108-396 MHz. This leaves DOCSIS with enough spectrum to fit 28 SC-QAM channels for D3.0/D3.1 modems plus 150 MHz OFDM channel for D3.1/D4.0 modems. These channels total ~2.5 Gbps of fixed DS capacity.

Figure 19 shows the DS BW Capacity required for several different Tmax_max tiers. The Y-axis is the capacity in Mbps while the X-axis is a log scale of the number of subs in a SG (i.e. Nsub). The yellow curve shows the DS peak period average consumption (i.e. Nsub*Tavg). The various green curves show the maximum capacity required for various Tmax_max, from 2G to 5G. Note that the burst plus ripple components is added to the average consumption.

The dashed lines show various key capacity marks. The lowest one is the SC-QAM capacity of 1.05 Gbps. This is the upper capacity limit for 3.0 modems. A separate calculation must be done to ensure the 3.0 pool does not exceed their capacity limit. The next blue dashed line shows the Fixed DS capacity limit of ~2.5 Gbps. As long as the DS peak period average consumption is less than this, then the US TG limit (i.e. TE condition #4) never kicks in. The figure highlights where Nsub*Tavg crosses this point at 358 subs. Above this point, the DS average consumption starts taking away usable FDX BW from the US.

The top 3 blue dashed lines show the total DS capacity available for the various FDX DS QAM modulations (i.e. 1024-, 2048-, 4096-QAM). The point where a green curve crosses the blue dashed line is the maximum number of subs that can be supported for that combination, i.e. TE condition #1.

This scenario supports:

- 2G DS tier to 375+ subs
- 3G DS tier to 240 subs @ 1024-QAM; 306 subs @ 4096-QAM
- 4G DS tier to 90 subs @ 1024-QAM; 157 subs @ 4096-QAM
- 5G DS tier does NOT fit

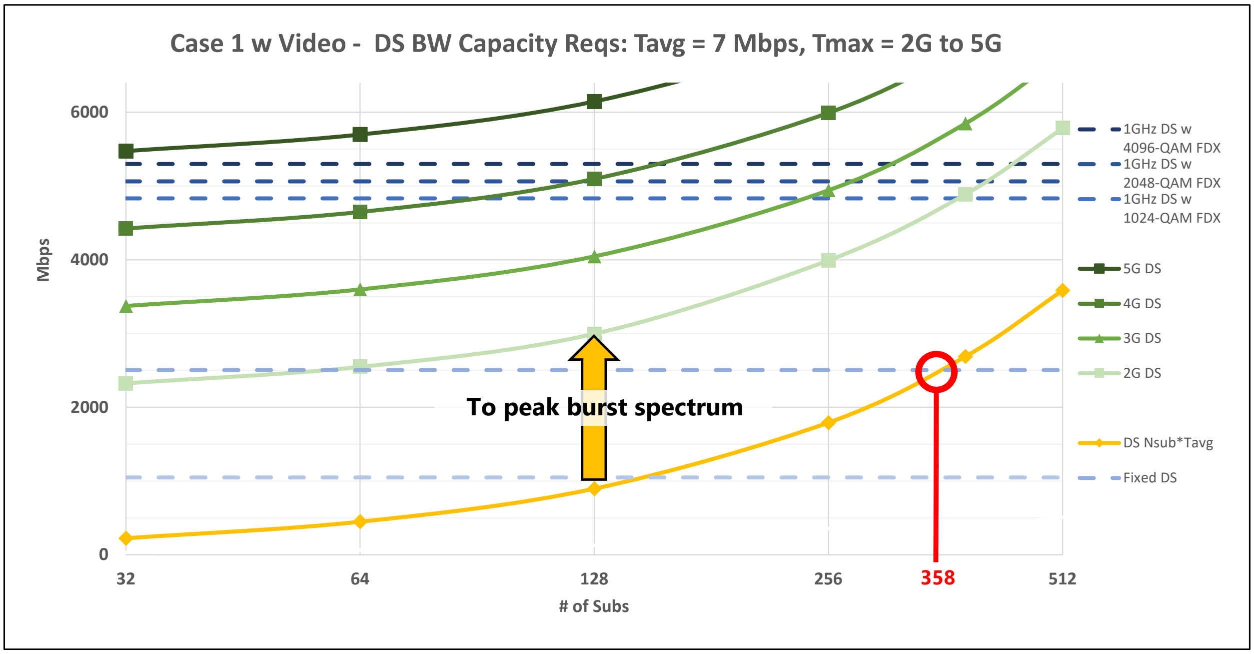

Figure 19 – Case 1 w Video, 396 MHz FDX – DS BW Capacity Limits

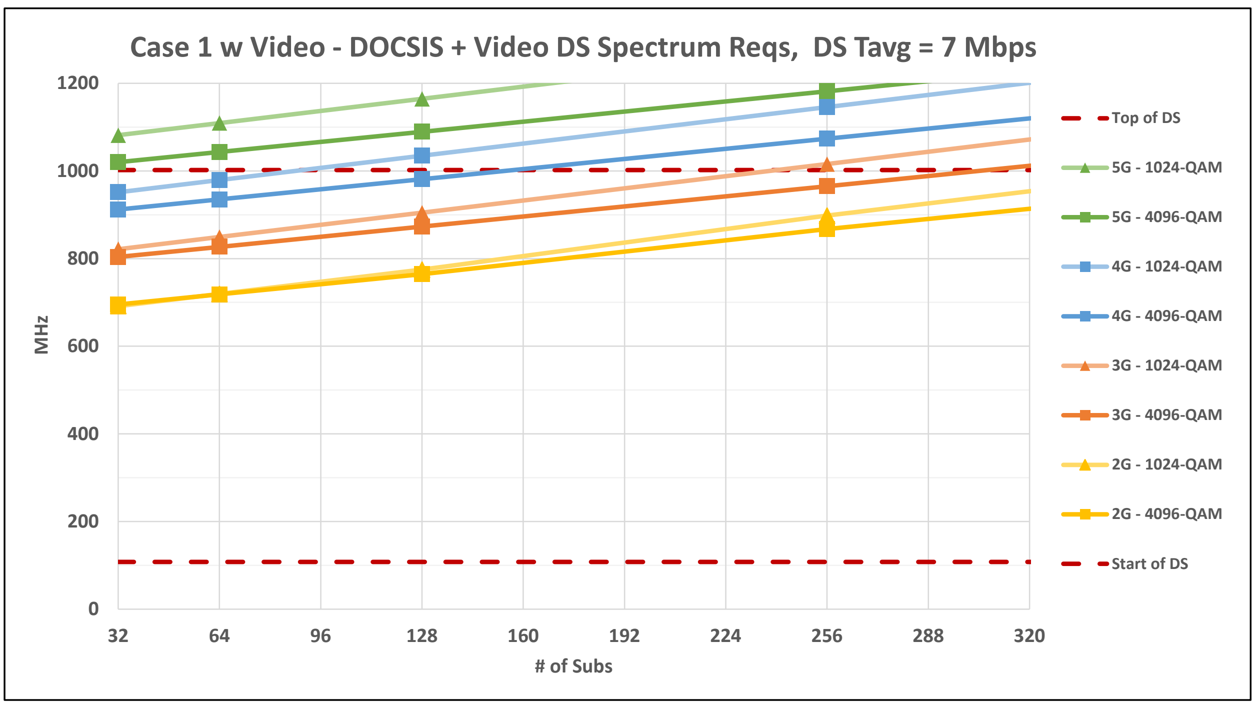

Figure 20 – Case 1 w Video, 396 MHz FDX – DS Spectrum Requirements

Figure 20 shows the DS spectrum requirements for each of the above combinations. Note – this includes both the DOCSIS DS spectrum and 288 MHz Legacy Video. This analysis assumes 1002 MHz is the top spectrum. If the roll-off for the 1 GHz is usable or the taps are upgraded to 1.2 GHz, then the figure also shows what can be achieved up to 1200 MHz. Even at 384 subs where Nsub*Tavg has overflowed into the FDX band, it only needs >96MHz of FDX BW <0.25% of the time.

To put the DS in context, here are examples of FDX band probability usage. A SG with 256 subs and a 4G DS tier will stay within the fixed DS BW outside FDX band for 98% of the time and only need some additional FDX BW for just 2% of the time.

The US SG limits (i.e. TE condition #2) are examined in Figure 21. Like the DS chart, the yellow curve depicts the US Nsub*Tavg while the various orange curves show the total US capacity required for various US tiers (1.5G – 2.5G). The red dashed line on the bottom shows the fixed US capacity of 450 Mbps below 85 MHz. Notice that the Nsub*Tavg component never crosses this line. This means that the DS TG limit (TE condition #3) will never kick in.

The three upper dashed lines show the total US capacity available for the various FDX US QAM modulations (i.e. 64-, 256-, 1024-QAM). The point where an orange curve crosses the blue dashed line is the maximum number of subs that can be supported for that combination, i.e. US TE condition #2 from above. As can be seen, each jump in modulation effectively allows an additional 0.5 Gbps to be added to Tmax_max.

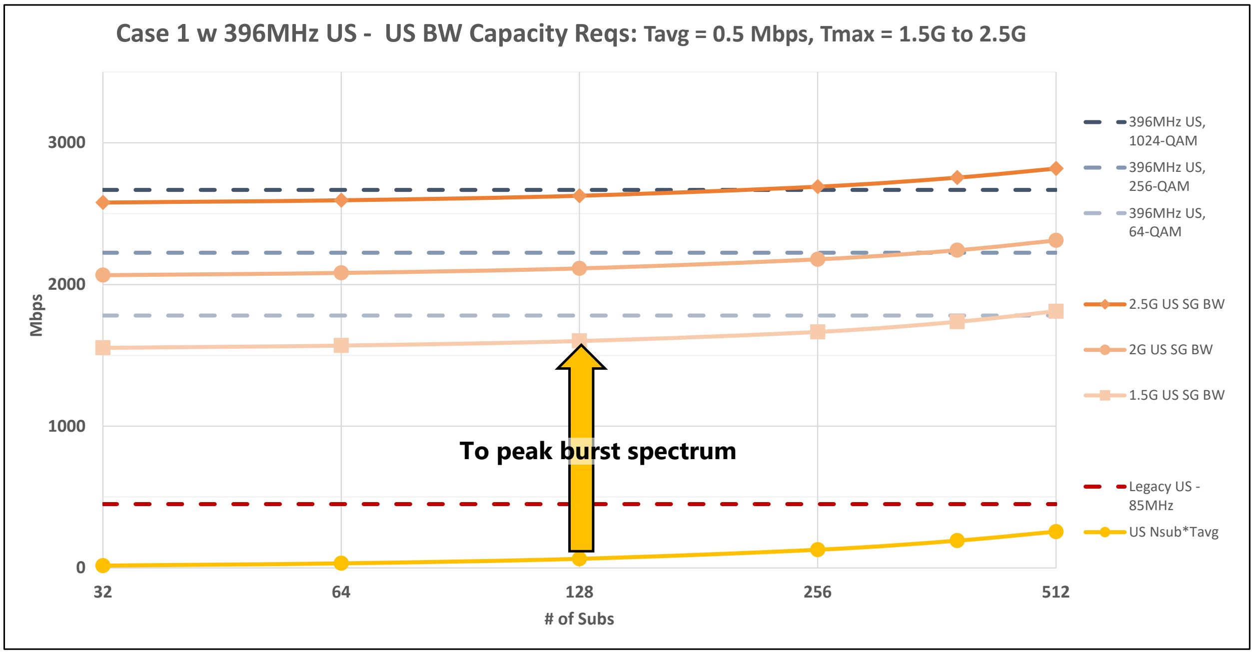

Figure 21 – Case 1 w Video, 396 MHz FDX – US BW Capacity Limits

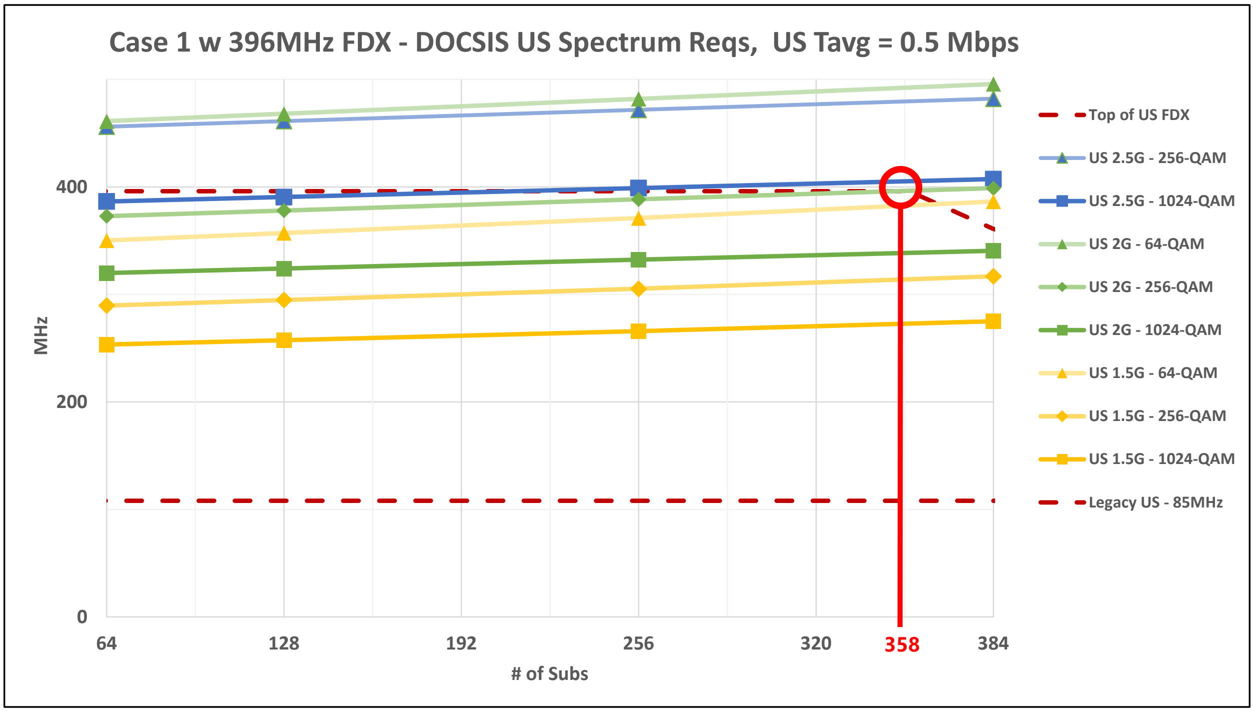

Figure 22 – Case 1 w Video, 396 MHz FDX – US Spectrum Requirements

Figure 22 shows the US spectrum requirements. The upper red dashed line is showing the top of the usable FDX spectrum. This is nominally 396 MHz when the DS peak period consumption does not overflow into the FDX band. However, Figure 22 shows that this point occurs at 358 subs. That is why this red dashed line starts to drop as Nsub increases past this point.

This scenario supports:

- 1.5G US tier to ~370 subs @ 64-QAM; 400+ subs @ 256- or 1024-QAM
- 2G US tier to ~350 subs @ 256-QAM; 400+ subs @ 1024-QAM
- 2.5G US tier to ~210 subs @ 1024-QAM
- 3G & 4G US tier does NOT fit

To put the US in context, a SG with 256 subs and a 2G US tier will stay within 85 MHz BW outside FDX band for 98% of the time and only need some additional FDX BW for just 2% of the time.

5.2.2. Case 1 with 288 MHz Legacy Video, 492 MHz FDX Band

This scenario increases the FDX band by 96 MHz to 108-492 MHz at the expense of the top OFDM channel. This leaves DOCSIS with enough fixed spectrum above the FDX band to fit 28 SC-QAM channels for D3.0/D3.1 modems plus 54 MHz OFDM channel for D3.1/D4.0 modems. These channels total ~1.57 Gbps of fixed DS capacity. It also maintains 288 MHz of Legacy Video spectrum.

Figure 23 shows the DS BW Capacity required for several different Tmax_max tiers. The second blue dashed line shows the Fixed DS capacity limit has now dropped ~1.57 Gbps. This causes the US TG limit (i.e. TE condition #4) to kicks in at 225 subs where Nsub*Tavg crosses this point. The DS average consumption starts taking away usable FDX BW from the US at a much lower point in this scenario.

The top 3 blue dashed lines show the total DS capacity available for the various FDX DS QAM modulations (i.e. 1024-, 2048-, 4096-QAM). This scenario supports:

- 2G DS tier to ~354 subs @ 1024-QAM; 400+ subs @ 2048- & 4096-QAM
- 3G DS tier to ~218 subs @ 1024-QAM; 306 subs @ 4096-QAM
- 4G DS tier to ~68 subs @ 1024-QAM; 157 subs @ 4096-QAM
- 5G DS tier does NOT fit

The DS 4096-QAM results are the same as before, while the DS 1024-QAM results are reduced from the previous scenario.

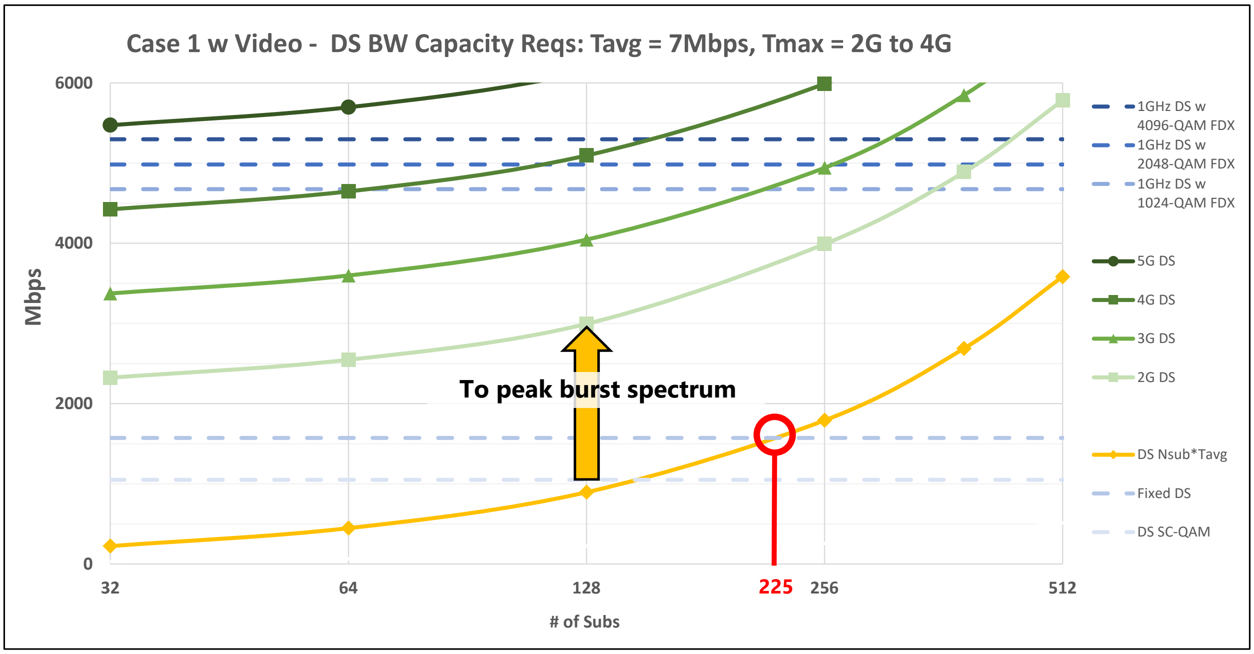

Figure 23 – Case 1 w Video, 492 MHz FDX – DS BW Capacity Limits

Figure 24 shows the DS spectrum requirements for each of the above combinations. Note – this includes both the DOCSIS DS spectrum and 288 MHz Legacy Video.

To put the DS in context, a SG with 128 subs and a 4G DS tier will stay within the fixed DS BW outside FDX band for 99% of the time and only need some additional FDX BW for just 1% of the time. Even at 256 subs where Nsub*Tavg has overflowed into the FDX band, it only needs >96MHz of FDX BW <5% of the time.

The US SG limits (i.e. TE condition #2) are examined in Figure 25. The various orange curves show the total US capacity required for various US tiers (1.5G – 3G). Like the previous scenario, each jump in modulation effectively allows an additional 0.5 Gbps to be added to Tmax_max.

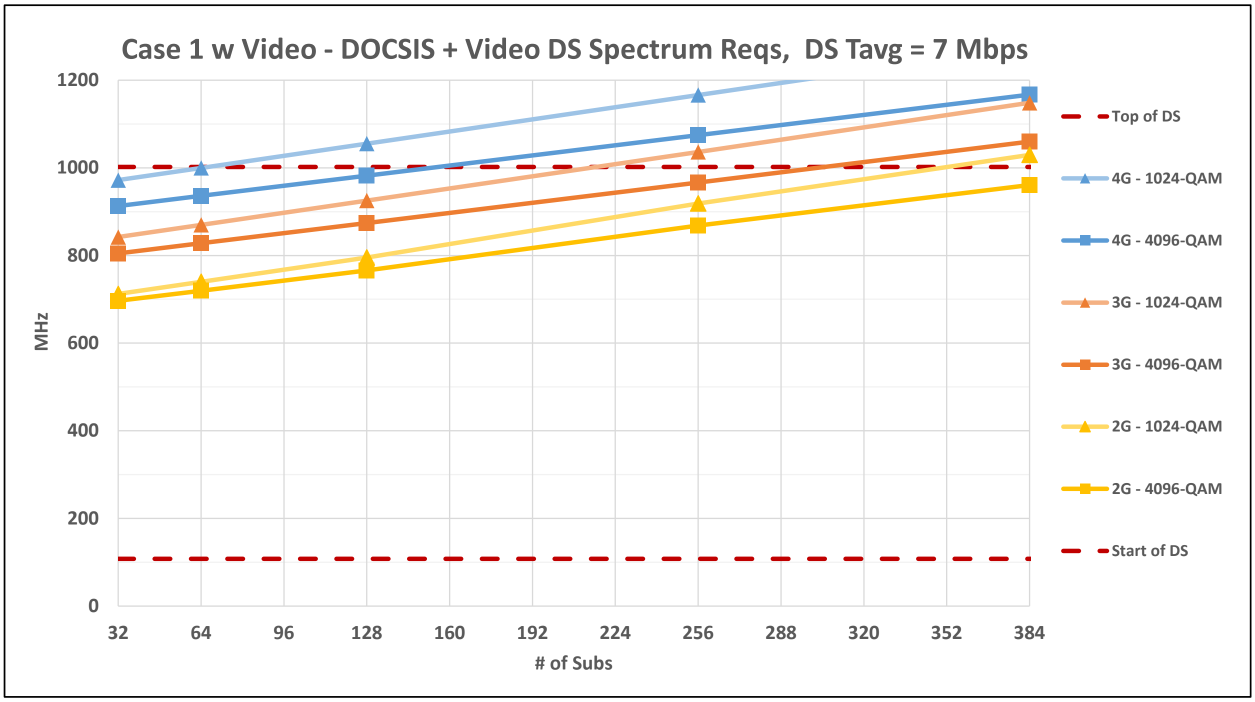

Figure 24 – Case 1 w Video, 492 MHz FDX – DS Spectrum Requirements

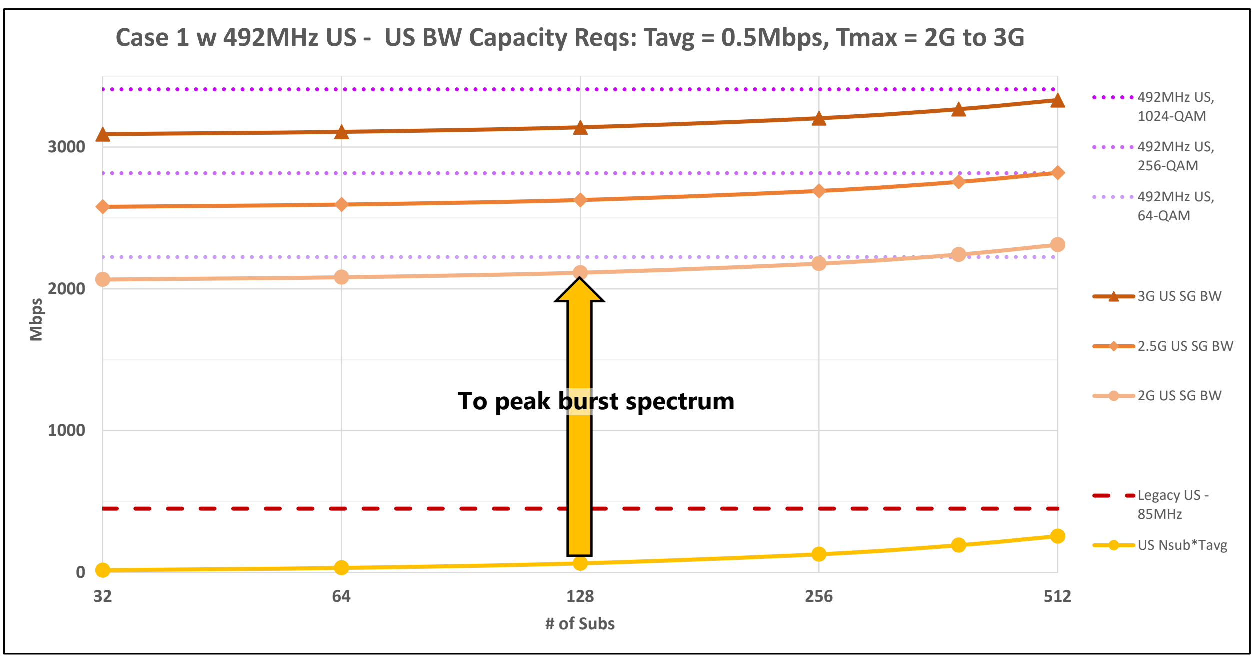

Figure 25 – Case 1 w Video, 492 MHz FDX – US BW Capacity Limits

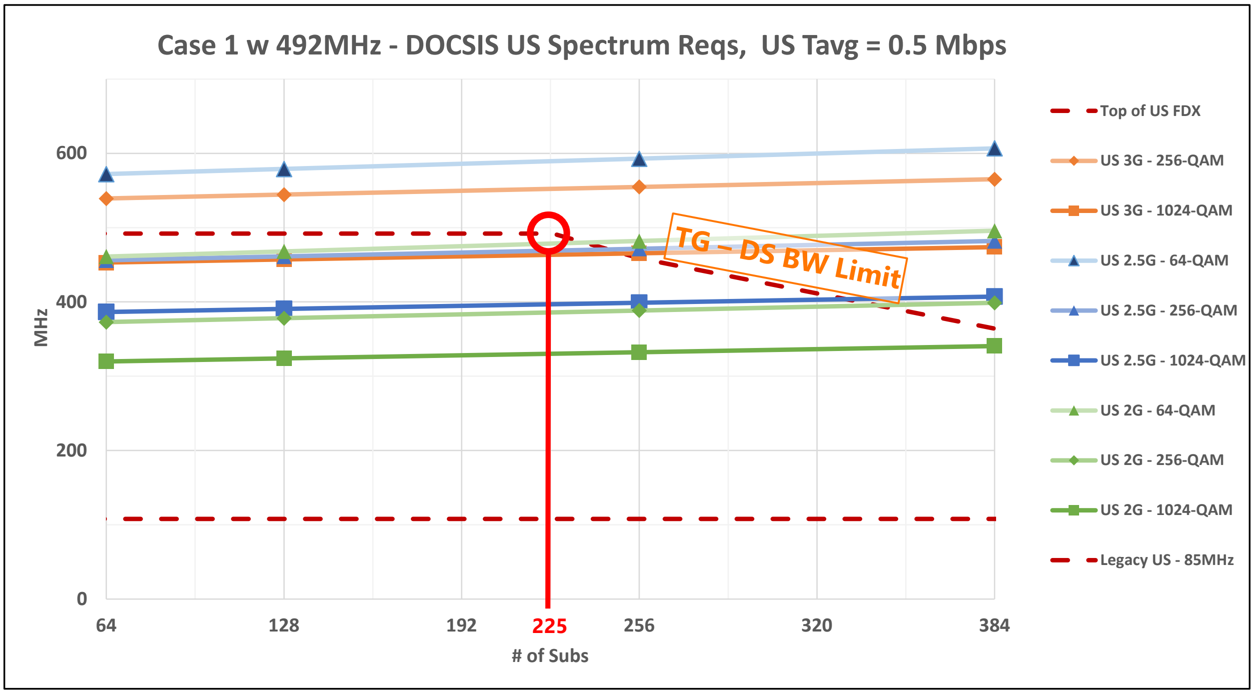

Figure 26 – Case 1 w Video, 492 MHz FDX – US Spectrum Requirements

Figure 26 shows the US spectrum requirements. The upper red dashed line is nominally 492 MHz until the DS peak period consumption overflows into the FDX band at 225 subs. The drop in usable US BW above 225 subs is the result of TE condition #4. As Nsub increases, the impact becomes greater.

This scenario supports:

- 1.5G US tier to ~360 subs @ 64-QAM; 400+ subs @ 256- or 1024-QAM
- 2G US tier to ~235 subs @ 64-QAM; ~340 subs @ 256-QAM
- 2.5G US tier to ~245 subs @ 256-QAM; 330 subs @ 1024-QAM
- 3G US tier to ~250 subs @ 1024-QAM
- 4G US tier does NOT fit

To put the US in context, a SG with 256 subs and a 2G US tier will stay within 85 MHz BW outside FDX band for 98% of the time and only need some additional FDX BW for just 2% of the time.

5.2.3. Case 1 with IPTV, no Legacy Video, 684 MHz FDX Band

In this scenario, 100% IPTV migration is reached which frees up the 288 MHz Legacy Video to be added to the DOCSIS BW pool. The FDX band is now a full 108-684 MHz while the fixed DS capacity above 684 MHz returns to 28 SC-QAM channels for 3.0/3.1 modems plus 150 MHz OFDM channel for D3.1/D4.0 modems. These channels total ~2.5 Gbps of fixed DS capacity.

Figure 27 shows the DS BW Capacity required for several different Tmax_max tiers. The second blue dashed line shows the Fixed DS capacity limit has now back up to ~2.5 Gbps. The US TG limit (i.e. TE condition #4) kicks in at 358 subs like the first scenario.

Notice that the top 3 blue dashed lines showing the total DS capacity available has increased significantly to 7-8 Gbps. The Case 1 IPTV scenario supports:

- 2G, 3G & 4G DS tier to 400+ subs @ all modulations
- 5G DS tier to ~272 subs @ 1024-QAM; 400+ subs @ 4096-QAM

Figure 28 shows the DS spectrum requirements for each of the above combinations. The probability of using the FDX band is the same as scenario 1 above. The US SG limits (i.e. TE condition #2) are examined in Figure 29. The various orange curves show the total US capacity required for various US tiers (2G – 4G). Figure 30 shows the US spectrum requirements. The Case 1 IPTV scenario supports:

- 2G US tier to 400+ subs @ all modulations
- 3G US tier to 400+ subs @ 256- & 1024-QAM
- 4G US tier to 400+ subs @ 1024-QAM

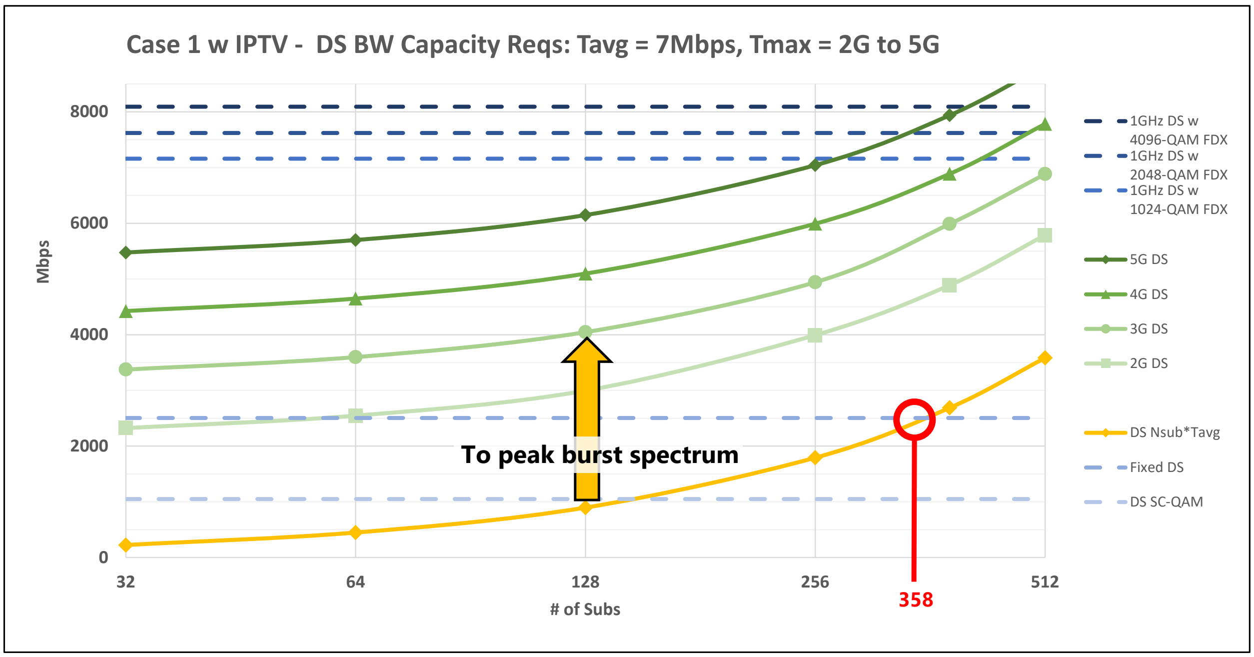

Figure 27 – Case 1 w IPTV, 684 MHz FDX – DS BW Capacity Limits

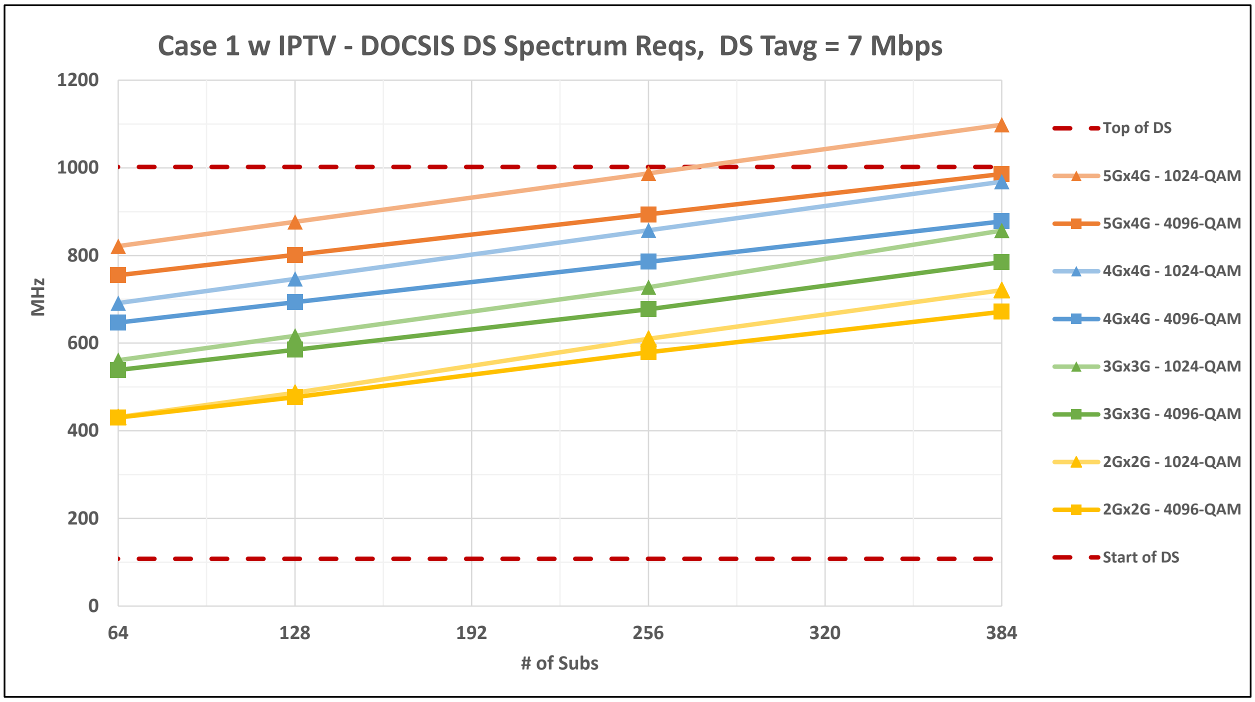

Figure 28 – Case 1 w IPTV, 684 MHz FDX – DS Spectrum Requirements

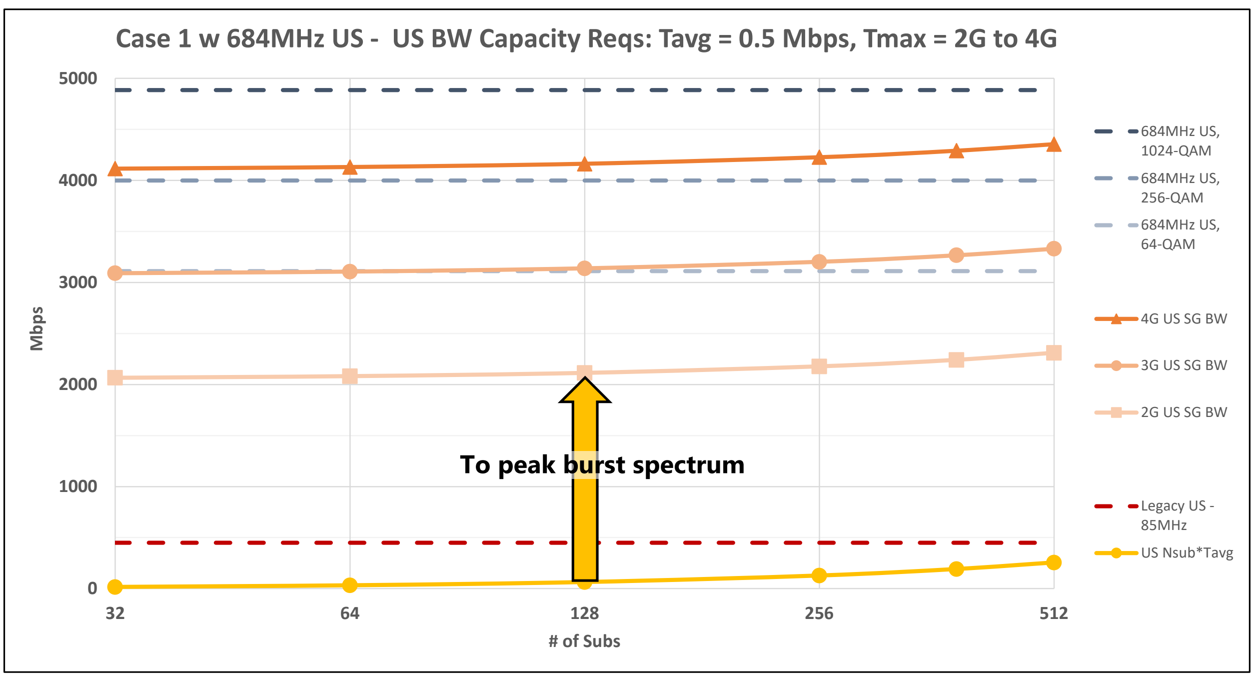

Figure 29 – Case 1 w IPTV, 684 MHz FDX – US BW Capacity Limits

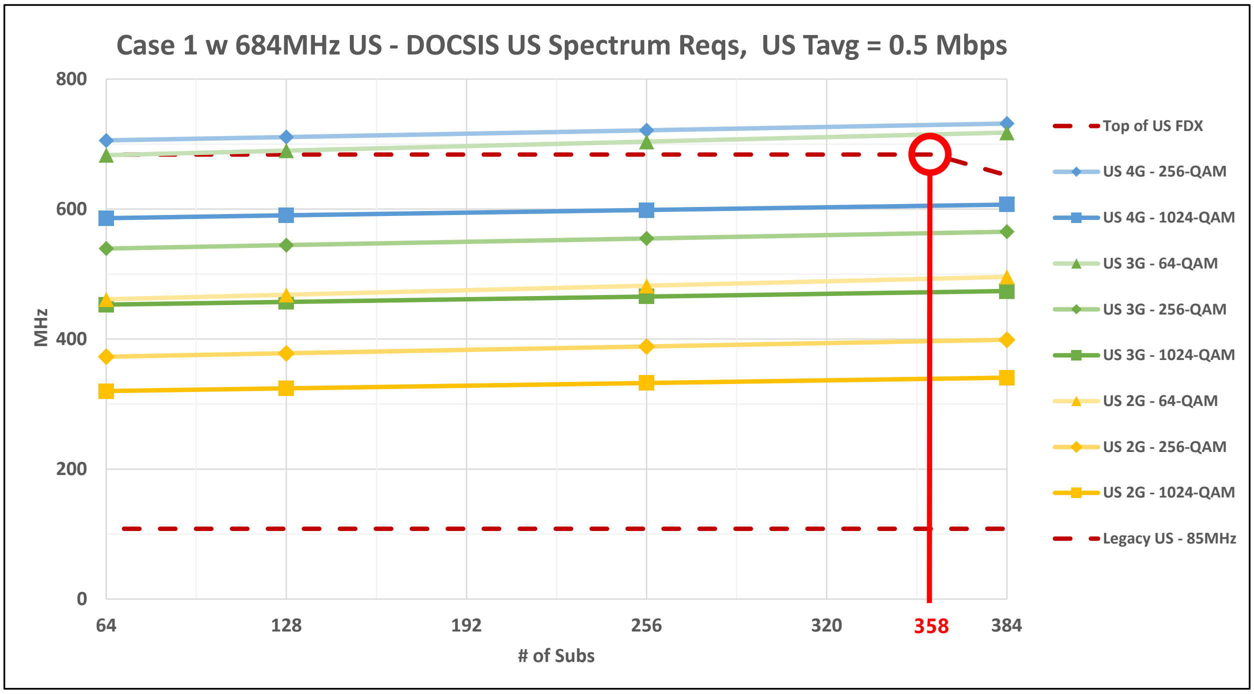

Figure 30 – Case 1 w IPTV, 684 MHz FDX – US Spectrum Requirements

5.2.4. Probability of DS + US Overlapping Tails – Case 1

The analysis until now has focused on capacity requirements and how many subscribers can receive adequate service for a given configuration. Now let's take a closer look at the probabilities of a DS tail overlapping with an US tail in the shared FDX band. Again, this assumes this is one single large TG, so the US and DS need to timeshare any FDX channels they need to use.

Figure 31 is a complex chart showing the probabilities of spectrum usage for both US + DS CDF with 200 subs. Our model assumes that the US capacity fills from the lowest channel in the spectrum and increases to the right until it reaches the top of the FDX band. The DS capacity is the opposite where it fills from the top channel in the spectrum down to the left. The fixed DS capacity is completely used before the DS uses the FDX band.

The yellow curve on the left is the US 'CDF'. Technically, it is the (1 – CDF) function. It tells us the probability that the US will need MORE than that amount of spectrum. The curve crosses the red line (85MHz) at 0.2%. This means that 99.8% of the time, the US capacity remains completely within 85 MHz, and only 0.2% of the time does it even need to request FDX BW. That equates to only 20 seconds every evening!

The blue curve on the right is the DS 'CDF'. It tells us the probability that the DS will need spectrum below that point. The curve crosses the top of the FDX band (684MHz) at 0.1%. This means that 99.9% of the time, the DS capacity remains completely above 684 MHz, and only 0.1% of the time does it even need to request FDX BW. That equates to only 10 seconds every evening!

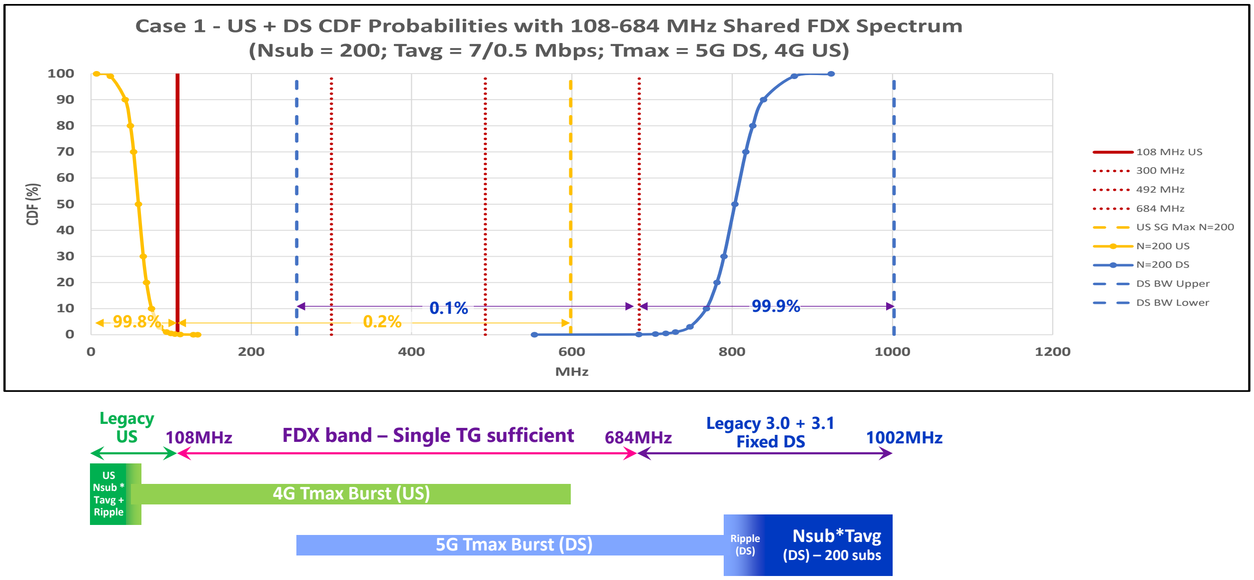

Figure 31 – Probability of DS + US Overlapping Tails – Case 1 IPTV, 200 subs

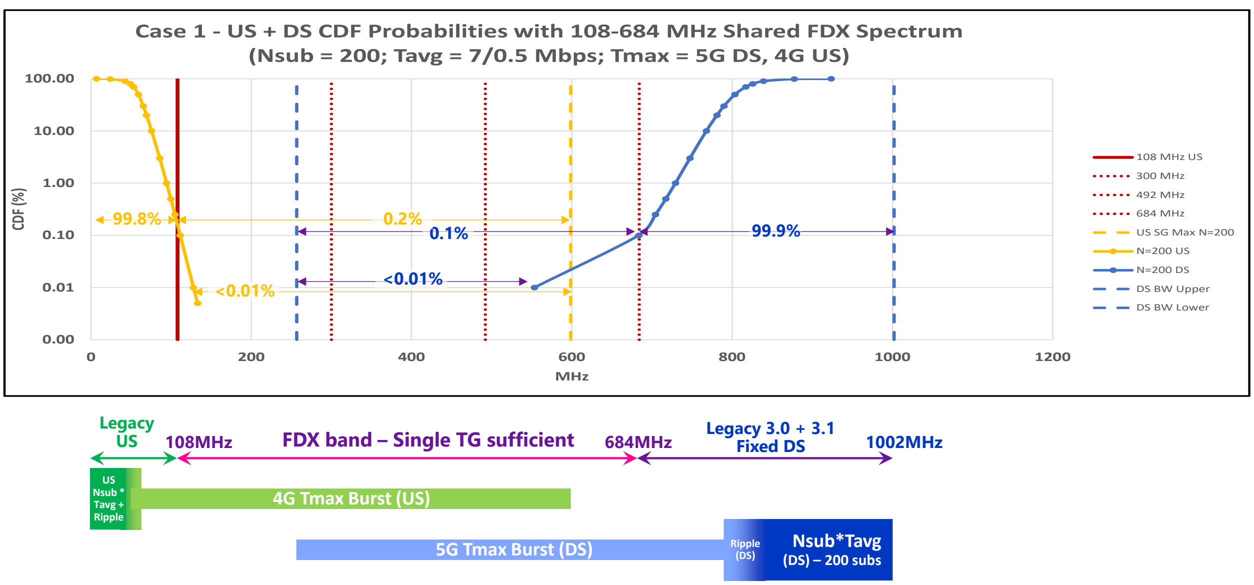

Figure 32 – Probability of DS + US Overlapping Tails – Case 1 IPTV, 200 subs (Log Scale)

To take a closer look at the overlapping tails, Figure 32 shows the same data but on a log scale. The US datapoint for 0.01% is at 128 MHz. What this means is that the US needs more than 20 MHz worth of capacity from the FDX band 0.01% of the time. For 20 MHz worth of capacity, the scheduler could assign this as 25% of a single 96 MHz FDX channel or as 4% of the entire FDX band. So, the US can still burst to the top of the FDX band, it is just that it is very infrequent and short bursts that shouldn't congest the system. In the rare event that the US does have a sustained burst (e.g. speed test once a month), the yellow dashed line at 599 MHz shows the US upper capacity limit based on the traffic engineering formula. The US will not exceed this BW point over a sustained period.

Looking at the DS CDF, the 0.01% datapoint is at 533 MHz. This means that 99.99% of the time, the DS BW needs are above this point and only 0.01% of the time does the DS need to go below this point. 533 MHz represents ~80% use of a single 192 MHz FDX DS channel or ~26% of the entire FDX band. The blue dashed line at 257 MHz represents the lowest spectrum BW point that the DS needs based on the traffic engineering formula (e.g. a speed test on top of normal traffic).

During the normal course of events, the probability of the US + DS tails overlapping is on the order of 1 second every 100+ years! Case 1 is the easier one, let's see how things change with Case 2.

5.3. FDX Use Case 2 – Longer Term Years 2026-29

Case 2 assumes a DS Tavg = 15 Mbps with US Tavg =1 Mbps. Two scenarios were considered: one with legacy video and one with 100% IPTV.

Table 3 shows the effective K-value for Case 2. As mentioned earlier, the optimum K-value varies with Nsub, Tavg and Tmax_max. This table is identical to Table 1except for the 3G DS tier. For 3G tiers, the effective K-value is 1.05 until the SG size reaches 256 subs where it increases to 1.1.

Table 3 – Effective K values used for Case 2

K-value Used	Original K		QoE modeling – Effective K			
Nsub/SG	**DS**	**US**	**2G DS**	**3G DS**	**4G,5G DS**	**All US**
32	1.2	1.2	1.05	1.05	1.05	1.025
64	1.2	1.2	1.05	1.05	1.05	1.025
128	1.2	1.2	1.05	1.05	1.05	1.025
256	-	-	1.1	1.1	1.05	1.025
384	-	-	1.1	1.1	1.05	1.025

5.3.1. Case 2 with 144 MHz Legacy Video, 684 MHz FDX

This scenario increases the FDX band to 684 MHz while Legacy Video shrinks to 144 MHz. This leaves DOCSIS with enough fixed spectrum above the FDX band to fit 16 SC-QAM channels for D3.0/D3.1 modems plus 96 MHz OFDM channel for D3.1/D4.0 modems. These channels total ~1.53 Gbps of fixed DS capacity.

Figure 33 shows the DS BW Capacity required for several different Tmax_max tiers. The second blue dashed line shows the Fixed DS capacity limit has now dropped ~1.53 Gbps. This causes the US TG limit (i.e. TE condition #4) to kicks in at 102 subs where Nsub*Tavg crosses this point. The DS average consumption starts taking away usable FDX BW from the US at a much lower point in this scenario.

If the system is upgraded to 1218 MHz, then more than 200 MHz of OFDM capacity is added to the fixed DS capacity. See the lower purple dashed line in Figure 33. This moves the crossover point for TE condition #4 up to 242 subs.

The top 3 blue dashed lines show the total DS capacity available for the various FDX DS QAM modulations (i.e. 1024-, 2048-, 4096-QAM) up to 1002 MHz. This scenario supports:

- 2G DS tier to ~265 subs @ 1024-QAM; ~327 subs @ 4096-QAM
- 3G DS tier to ~200 subs @ 1024-QAM; ~260 subs @ 4096-QAM
- 4G DS tier to ~132 subs @ 1024-QAM; ~195 subs @ 4096-QAM
- 5G DS tier to ~62 subs @ 1024-QAM; ~125 subs @ 4096-QAM

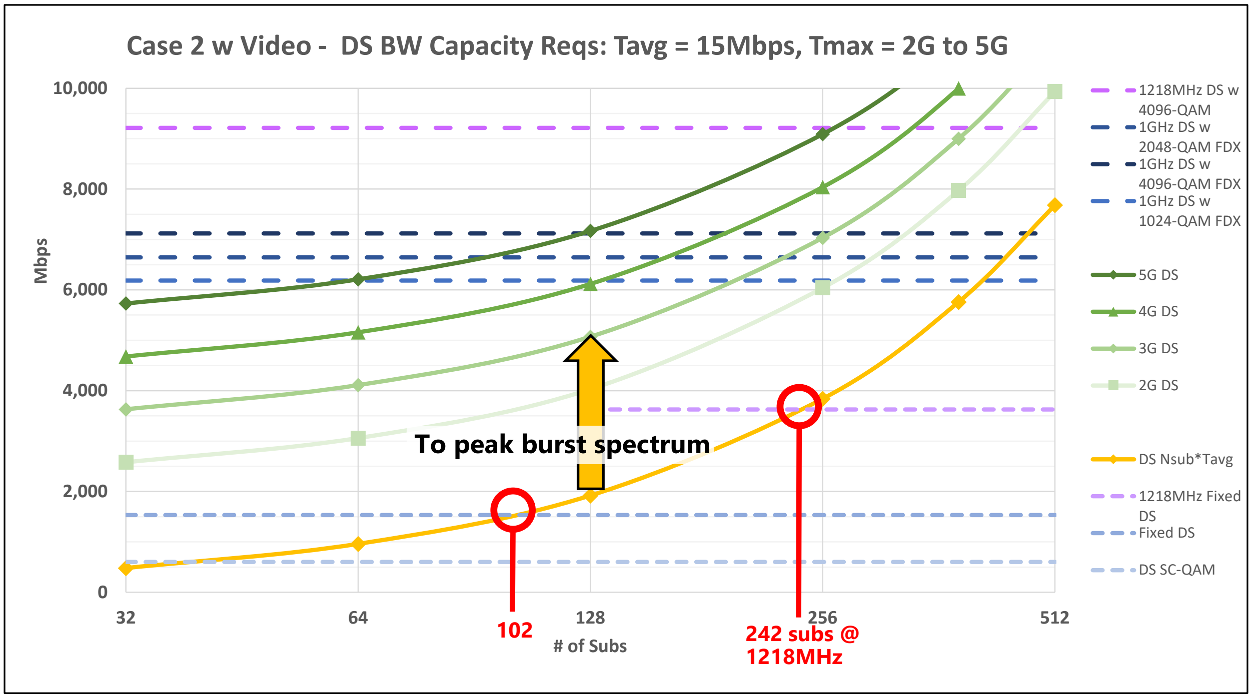

Figure 33 – Case 2 w Video, 684 MHz FDX – DS BW Capacity Limits

Figure 34 shows the DS spectrum requirements for each of the above combinations. Note – this includes both the DOCSIS DS spectrum and 144 MHz Legacy Video.

To put the DS in context, a SG with 128 subs and a 4G DS tier will stay within the fixed DS BW plus one 192 MHz FDX band for 99.5% of the time and only need some additional FDX BW for just 0.5% of the time.

The US SG limits (i.e. TE condition #2) are examined in Figure 35. The various orange curves show the total US capacity required for various US tiers (2G – 4G). Each jump in modulation effectively allows an additional 1 Gbps to be added to Tmax_max. Note that the US Nsub*Tavg finally exceeds the 85 MHz BW beyond 450 subs which is much higher than the DS.

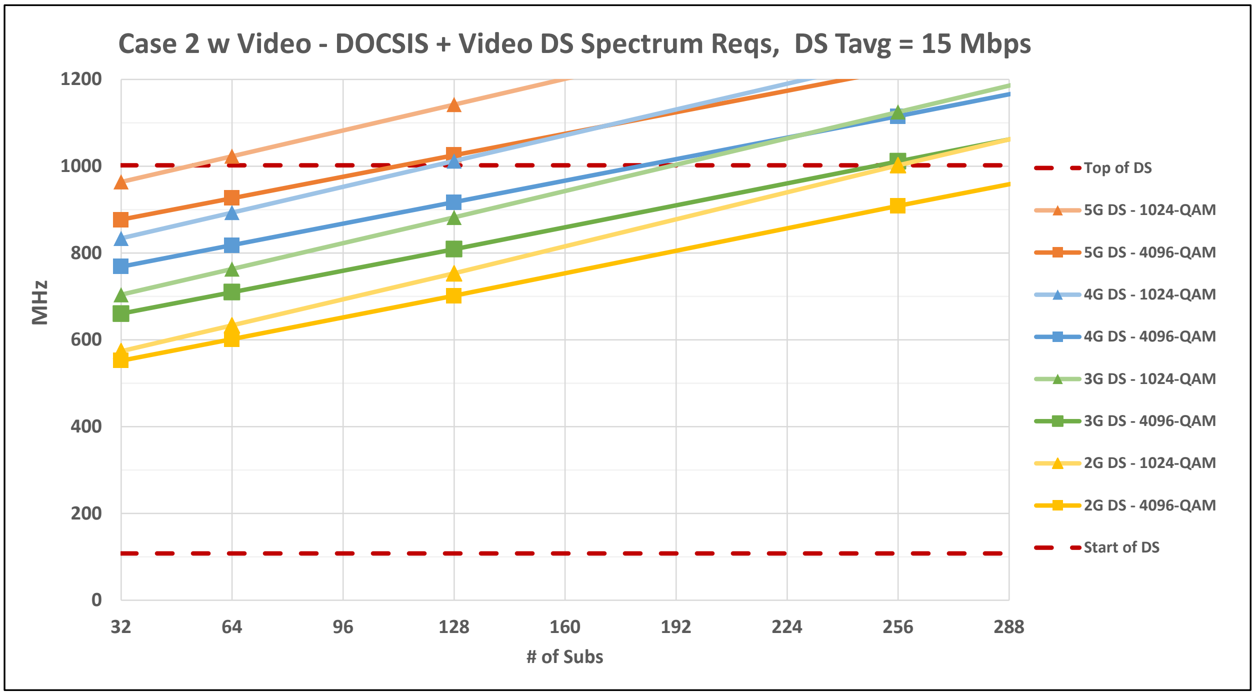

Figure 34 – Case 2 w Video, 684 MHz FDX – DS Spectrum Requirements

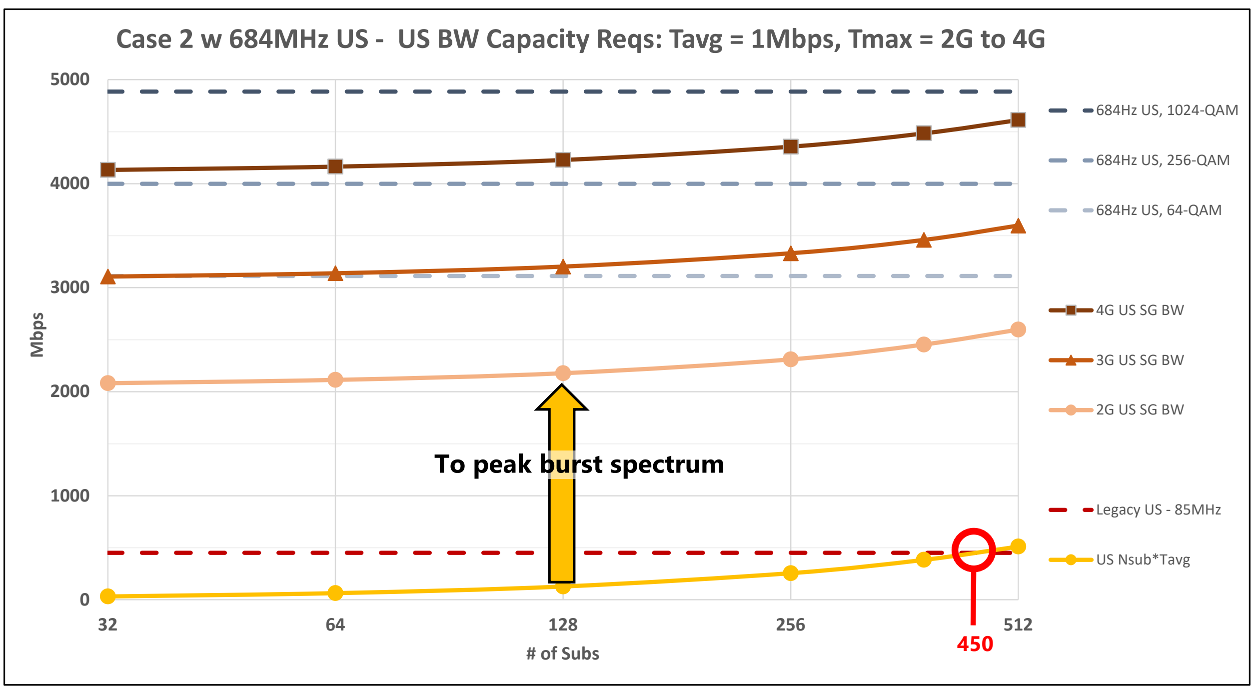

Figure 35 – Case 2 w Video, 684 MHz FDX – US BW Capacity Limits

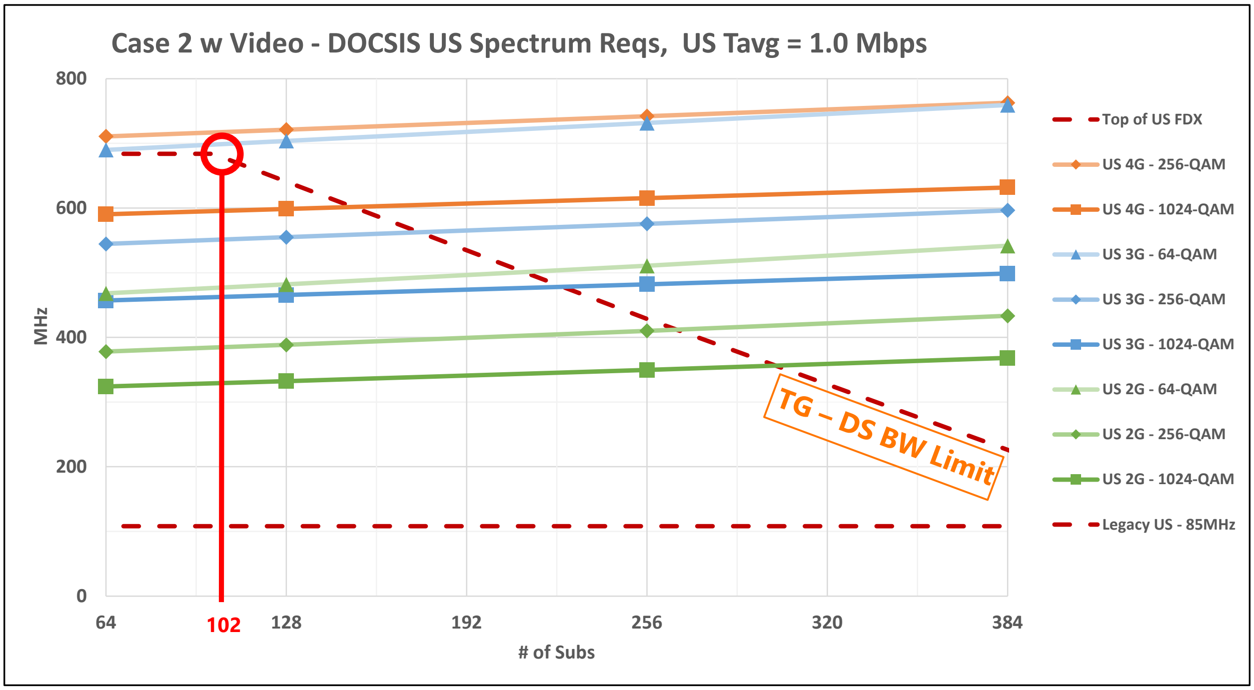

Figure 36 – Case 2 w Video, 684 MHz FDX – US Spectrum Requirements

Figure 36 shows the US spectrum requirements. The upper red dashed line starts at 684 MHz until the DS peak period consumption overflows into the FDX band at 102 subs. The drop in usable US BW above 102 subs is the result of TE condition #4. As Nsub increases, the impact becomes greater. The slope is also steeper in Case 2 with the higher Tavg.

This scenario supports:

- 2G US tier to ~210 subs @ 64-QAM; ~265 subs @ 256-QAM; 310 subs @ 1024-QAM
- 3G US tier to ~175 subs @ 256-QAM; 225 subs @ 1024-QAM
- 4G US tier to ~150 subs @ 1024-QAM

TE condition #4 has now become the limiting factor in how many subs can be supported in a SG.

To put the US in context, a SG with 256 subs and a 2G US tier will stay within 85 MHz BW outside FDX band for 92% of the time and only need some additional FDX BW for just 8% of the time.

5.3.2. Case 2 with IPTV, no Legacy Video, 684 MHz FDX Band

Migrating to 100% IPTV frees up the 144 MHz Legacy Video which helps with the previous TE condition #4 limitations. The fixed DS capacity above 684 MHz is now 16 SC-QAM channels for 3.0/3.1 modems plus 240 MHz OFDM channel for 3.1/4.0 modems. These channels total ~2.93 Gbps of fixed DS capacity, almost double the previous scenario.

Figure 37 shows the DS BW Capacity required for several different Tmax_max tiers. The second blue dashed line shows the Fixed DS capacity limit has now back up to ~3 Gbps. The US TG limit (i.e. TE condition #4) now kicks in at 195 subs.

Notice that the top 3 blue dashed lines showing the total DS capacity available has increased significantly to 7.5-8.5 Gbps. The Case 2 IPTV scenario supports:

- 2G DS tier to ~353 subs @ 1024-QAM; 400+ subs @ 4096-QAM
- 3G DS tier to ~285 subs @ 1024-QAM; ~348 subs @ 4096-QAM
- 4G DS tier to ~225 subs @ 1024-QAM; ~287 subs @ 4096-QAM
- 5G DS tier to ~155 subs @ 1024-QAM; ~218 subs @ 4096-QAM

Figure 38 shows the DS spectrum requirements for each of the above combinations. The US SG limits (i.e. TE condition #2) are examined in Figure 39**.** The various orange curves show the total US capacity required for various US tiers (2G – 4G). Figure 40 shows the US spectrum requirements. The Case 2 IPTV scenario supports:

- 2G US tier to ~275 subs @ 64-QAM; ~325 subs @ 256-QAM; 350 subs @ 1024-QAM
- 3G US tier to ~250 subs @ 256-QAM; ~290 subs @ 1024-QAM
- 4G US tier to ~230 subs @ 1024-QAM

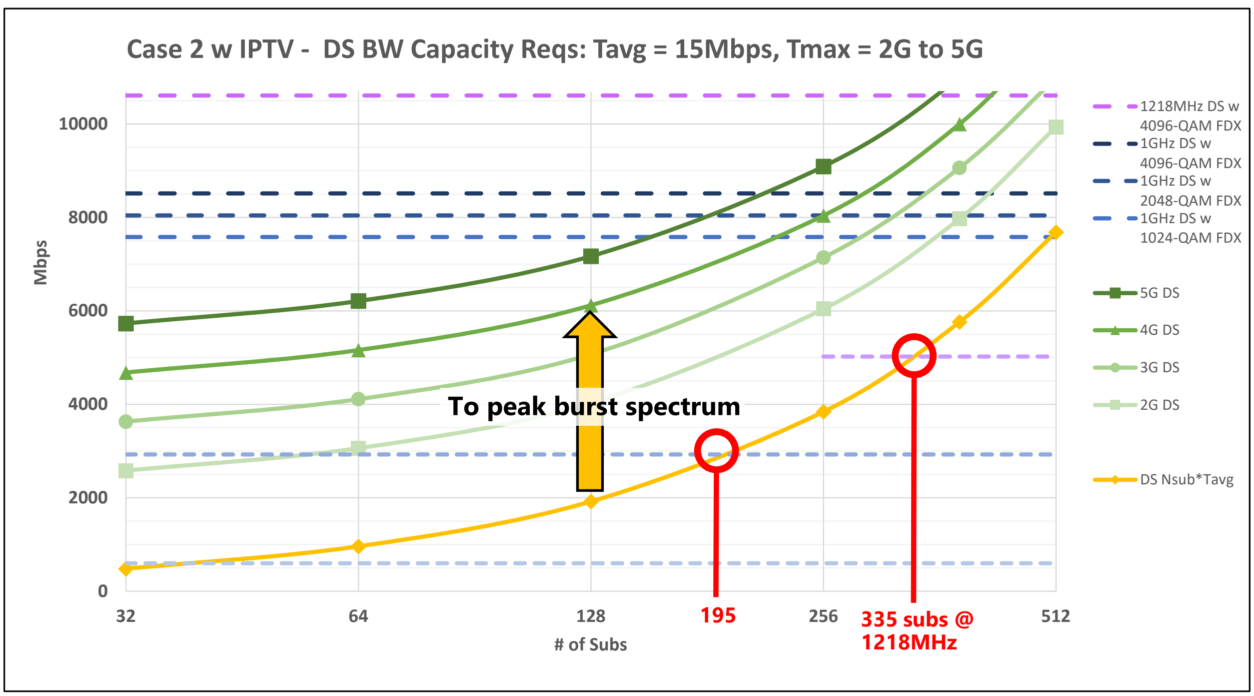

Figure 37 – Case 2 w IPTV, 684 MHz FDX – DS BW Capacity Limits

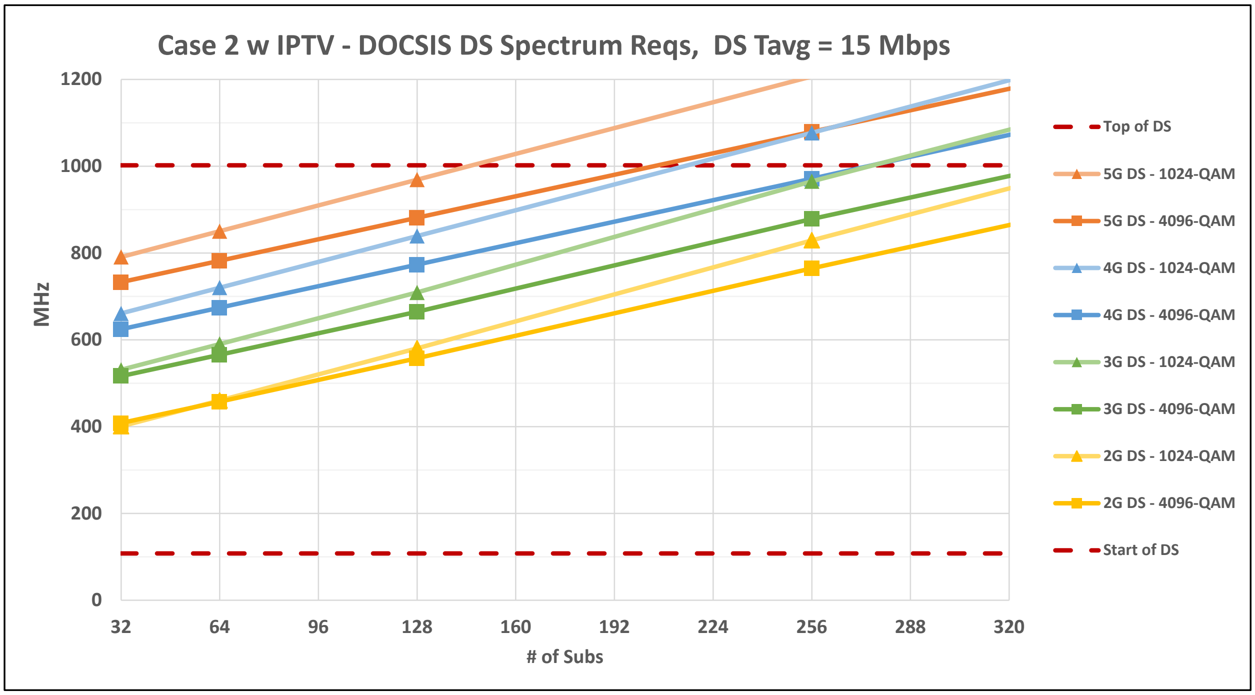

Figure 38 – Case 2 w IPTV, 684 MHz FDX – DS Spectrum Requirements

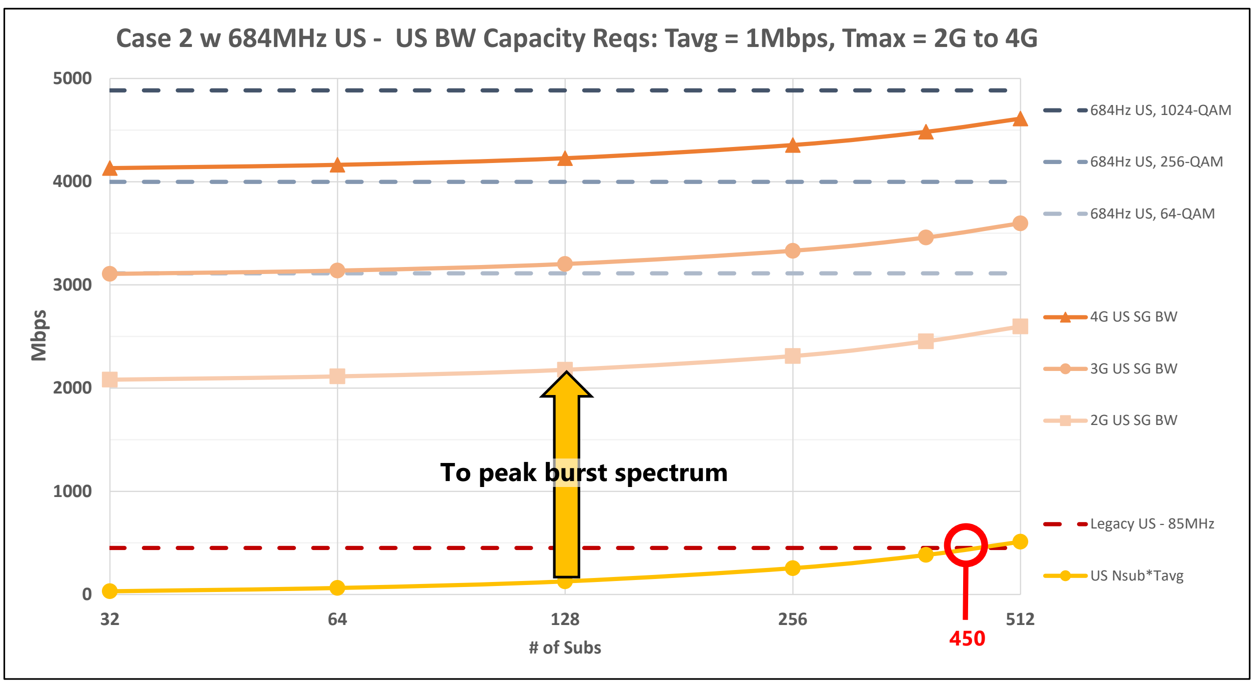

Figure 39 – Case 2 w IPTV, 684 MHz FDX – US BW Capacity Limits

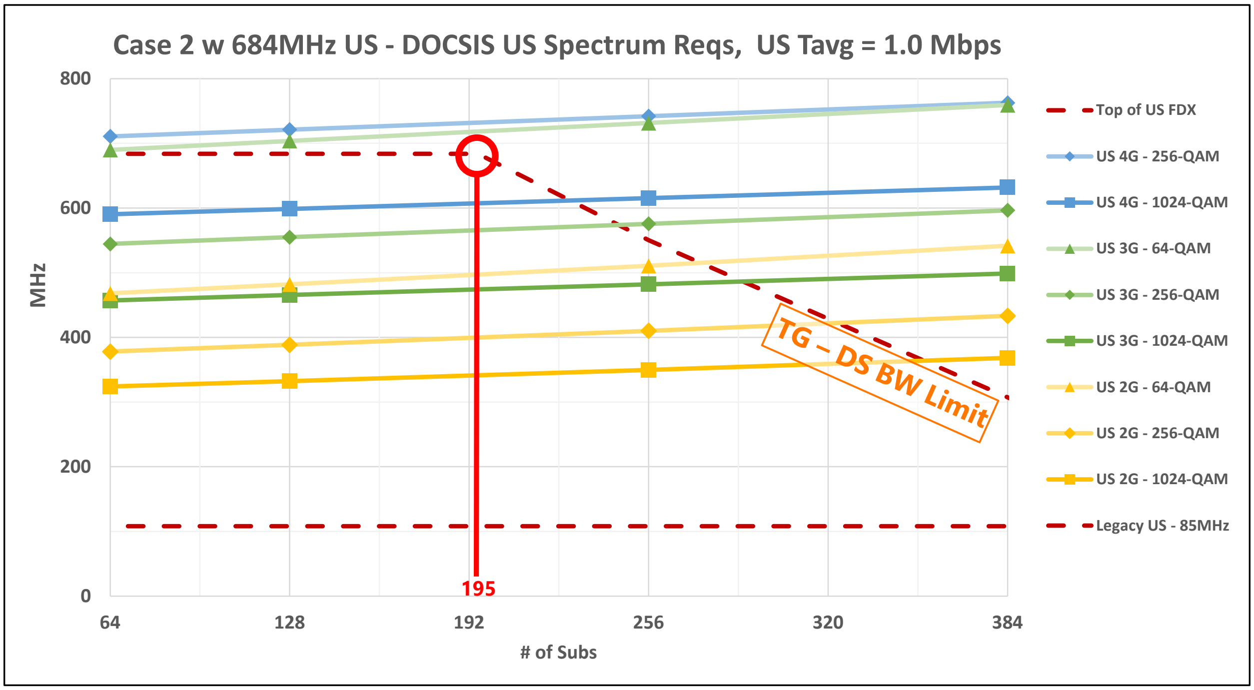

Figure 40 – Case 2 w IPTV, 684 MHz FDX – US Spectrum Requirements

Figure 41 shows how three different service tier combinations map to spectrum using the candlesticks from the traffic engineering formula. The top example is a SG with 256 subs and 3G DS x 3G US tiers. Notice that the 3G US burst requires more spectrum than the DS 3G burst due to its lower modulation. This is a non-typical example of when US limits the DS from growing any further.

The middle example shows 256 subs with 4Gx2G tiers. The fully symmetric would not fit. As can be seen, there is still some headroom in both US + DS for additional Tavg growth. The bottom example shows 128 subs with 5Gx4G tiers with headroom.

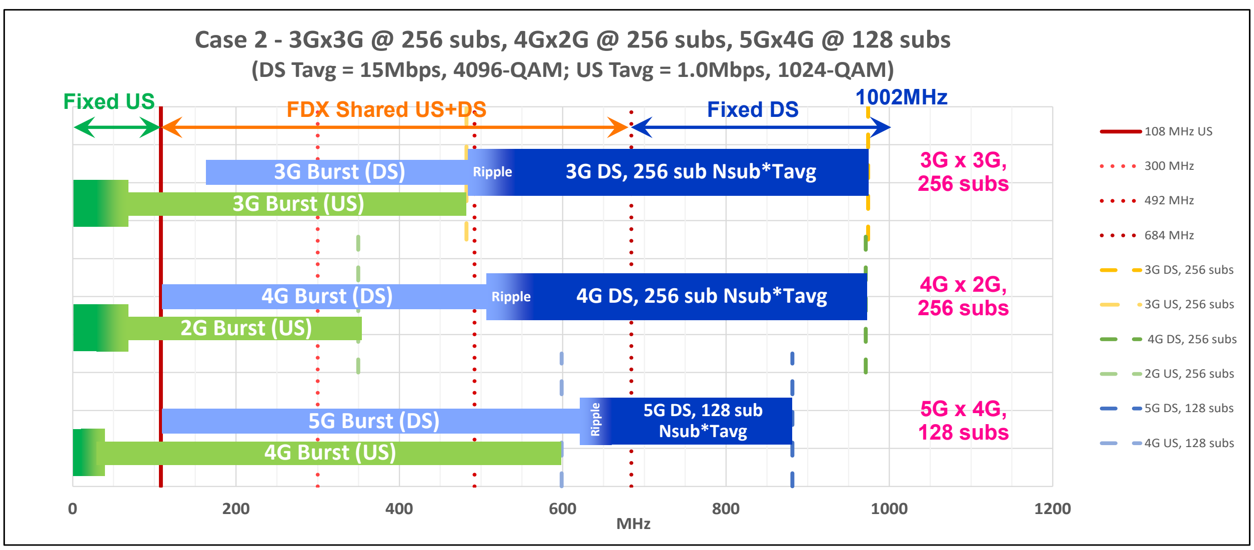

Figure 41 – Case 2 w IPTV, 684 MHz FDX – Various Service Tier Examples

5.3.3. Probability of DS + US Overlapping Tails – Case 2

As seen in the Case 2 scenarios with increased Tavg, the DS capacity needs more and more of the FDX BW. Figure 42 looks at the Case 2 IPTV probabilities of spectrum usage for both US + DS CDF for 200 subs.

As before, the yellow curve on the left is the US 'CDF'. For Case 2, the curve crosses the red line (85MHz) at ~3%. This means that 97% of the time, the US capacity remains completely within 85 MHz, and ~3% of the time it needs to request FDX BW.

The blue curve on the right is the DS 'CDF'. The DS curve crosses the top of the FDX band (684MHz) at ~50%. This means that 50% of the time, the DS capacity remains completely above 684 MHz, while 50% of the time it will need to request FDX BW. A much different situation than Case 1!

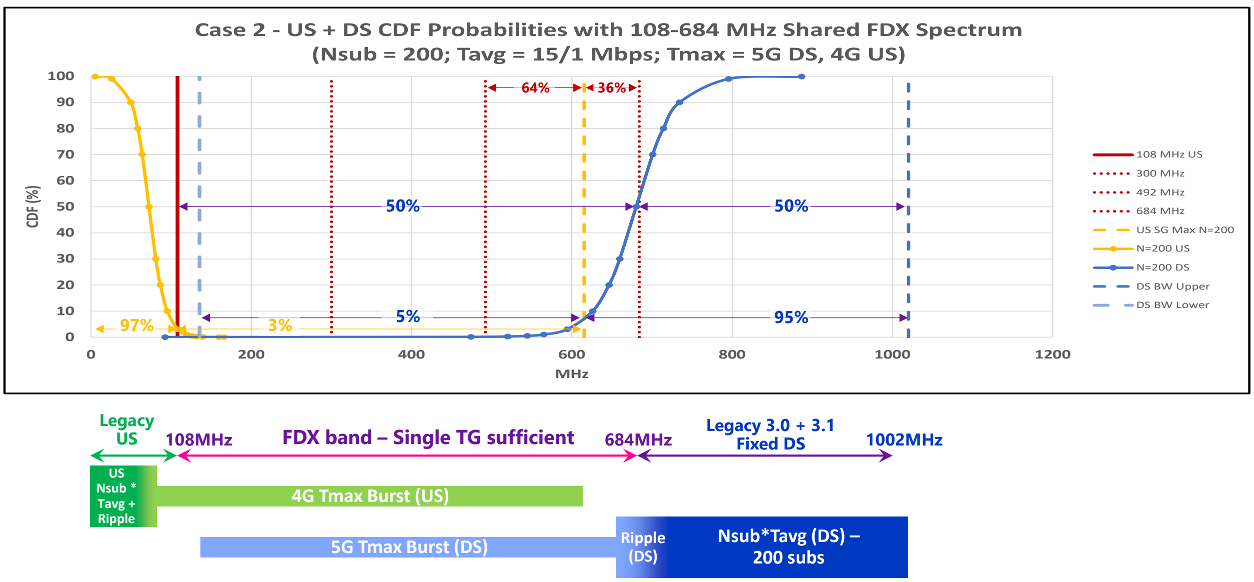

Figure 42 – Probability of DS + US Overlapping Tails – Case 2 IPTV, 200 subs

To take a closer look at the overlapping tails, Figure 43 shows the same data but on a log scale. The US datapoint for 0.01% is at 160 MHz. What this means is that the US needs more than 52 MHz worth of capacity from the FDX band just 0.01% of the time. For 52 MHz worth of capacity, the scheduler could assign this as 55% of a single 96 MHz FDX channel or 9% of the entire FDX band. Even in Case 2, the US can still burst to the top of the FDX band, it is just that it is very infrequent and short bursts that shouldn't congest the system. In the rare event that the US does have a sustained burst (e.g. speed test once a month), the yellow dashed line at 615 MHz shows the upper capacity limit based on the traffic engineering formula. The US will not exceed this BW point over a sustained period.

Looking at the DS CDF, it crosses 492 MHz just above 0.1%. This means that 99.9% of the time, the DS BW needs are above this point (i.e. DS Fixed spectrum + a single 192 MHz FDX DS channel) and only 0.1% of the time does the DS need to go below this point into a 2nd &/or 3rd FDX DS channel. The blue dashed line at 135 MHz represents the lowest spectrum BW point that the DS needs based on the traffic engineering formula (e.g. a speed test on top of normal traffic).

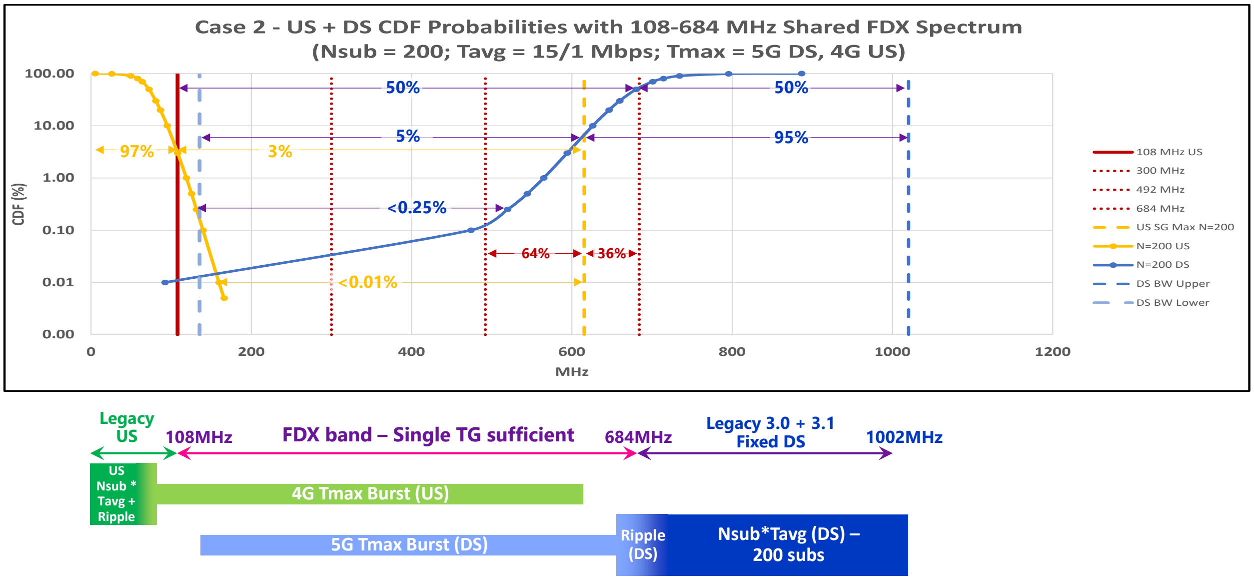

Figure 43 – Probability of DS + US Overlapping Tails – Case 2 IPTV, 200 subs (Log Scale)

While the probabilities of the US + DS tails overlapping becomes infinitesimal, the operator still needs to engineer the system for that instance when one modem does burst to its maximum (e.g. speed test). The worst case happens when the US bursts to its maximum, which is the yellow dashed line at 615 MHz. At this point, the US is completely consuming 108-492 MHz FDX band and needs 64% of the 492-684 MHz FDX channel (and 36% is available for DS BW). The critical question then becomes how much is the DS overlapping this?

To answer that, look to see where the blue DS curve intersects with US max at 615 MHz. The DS CDF intersects at the 5% point on its curve. This means that 95% of the time, the DS stays above this BW point, and everyone has sufficient capacity. However, 5% of the time there is a BW conflict with the DS overlapping the US Tmax_max burst. But how big is it? Breaking the 5% window into segments shows:

- US burst gets 4000+ Mbps for 95% of time
- US burst gets 3840-4000 Mbps for 2% of time
- US burst gets 3615-3840 Mbps for 2% of time
- US burst gets 3450-3615 Mbps for 0.5% of time
- US burst gets 3270-3450Mbps for 0.25% of time
- US burst gets <3270 Mbps for 0.25% of time

The impact of this contention is minimal. On average, it will reduce the US 4G burst by <0.5%. Just adding a small amount of over-provisioning into Tmax_max will easily compensate for this slight shortfall.

6. Results Summary

As we noted at the outset and have shown herein, traffic analysis of the FDX band is tractable using known statistical methods, tools, and empirical data. We have gone onto significant depth explaining the statistical foundation of burst traffic, the empirical basis of these statistics, describing modeling scenarios used to create operational guidelines, and creating parametric curves that quantify the relationships among bandwidth allocated, speeds, bandwidth efficiency, and transmission group sizes.

The foundational principle enabling large transmission group sizes is that, because of the statistical characteristics of real traffic, networks are able to be oversubscribed – meaning that many customers can effectively share capacity that is a number less than the sum of their individual peak usage needs. This has been the case for decades. As the peak speeds to average utilization skew to larger ratios in the Gigabit speed era, we see that the higher the peak burst, the increasingly infrequently it occurs. Years of evidence from broadband service delivery have been used to generate probability distributions that describe the traffic behavior, and which are leveraged and extrapolated to predict the expected performance in these models, and in particular how that applies to the dual-use FDX band.

The consolidated results are shown in Table 4. Results show that large TGs can be supported without loss of QoE. Furthermore, the TG sizes align well with what are current node sizes and service groups sizes that will result from network augments that are already part of the Comcast HFC upgrade plan.

As a point of context, the average size of a fiber node today is 350 hp and getting smaller. As nodes get split during DAA network upgrades, they are roughly halved – or become approximately 175 hhp. Adjusting then for residential penetration, these numbers would represent a service group of subscribers of about 200 and 100, respectively – very well aligned to the values in Table 4. *This suggests that, with proper spectrum management, and even under aggressive CAGR, N+x FDX – including with elongated TG – will comfortably multi-gig symmetrical speeds well within today's network architecture and expectations of future augments.*

Table 4 – TG Size Summary: Spectrum Scenarios and Speeds

TG Size Summary*

	Scenario	2G / 2G	3G / 3G	4G / 4G
Near Term	1	>350	306 / X	157 / X
Near Term	2	>400	250	157 / X
Near Term	3	> 400	>400	>400
Longer Term	4	310	225	150
Longer Term	5	350	290	230

*4k-QAM DS / 1k-QAM US

5 Gbps DS @ TG Size = 218
6 Gbps DS @ TG Size ~ 125
7 Gbps DS @ TG Size ~ 60

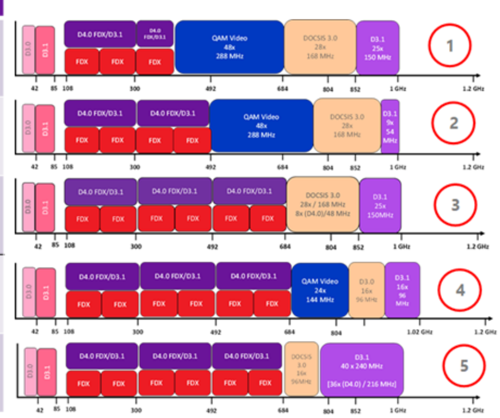

60

7. Conclusion

Traffic engineering of the HFC network has been at the root of our success deploying HSD services cost-effectively, and without over-engineering the network. The industry has decades of successful deployment experience. Over time, operators have developed mature empirically-based statistical models of DS and US. This has translated into robust processes for operating and augmenting the network in the face of DS and US CAGR and continually increasing speeds tiers. The essential relationships of capacity supply vs peak speed and user demand have led to predictable thresholds for configuring service groups. These same learnings can be applied to traffic engineering the FDX band.

However, we have even more tools, and more powerful ones, at our disposal today to manage capacity. Utilization data can be captured and delivered to the cloud on a service-group-by-service group basis. With algorithms to crunch the localized data and traffic trends, extremely granular, and thereby more efficient, cost-effective, and targeted migration planning can take place. On top of this, because there are now decades of aggregate traffic data history, the empirical data can be used confidently to create probability distributions using classical goodness-of-fit tools of statistical analysis and pattern recognition theory. Once these are crystallized into mathematical expressions, we have the foundation to quantifiably predict burst characteristics sliced and diced into durations and size, and use these to guide business decisions. The modeling exercise described in this paper is an outcome of this process.

From a traffic engineering perspective, an FDX US population represents a small number of less aggressive "users" looking to infrequently access frequency and time resources in a specific band that is managed jointly by the vCMTS. Similar to independent DS and US criteria, thresholds can be quantified for TG size versus peak speeds (for a fixed set of assumptions on CAGR, total capacity, penetration, BW efficiency). Infrequent (statistically) peak bursts and DS Avg BW >> US Avg BW means the FDX Band's RBAs will be deployed as DS blocks the vast majority of the time.

By contrast, if a large chunk of the coaxial spectrum such as 5-204 MHz, 5-300 MHz, 5-396 MHz, etc., is dedicated to upstream traffic, it will be *idle* the vast majority of the time. Furthermore, for the increasingly high "ultra-high split" options, while this high-quality spectrum mostly idles as an outlet for a very occasional US burst, the approach will force more downstream into the least predictable and never before activated part of the coaxial spectrum above 1 GHz, and only *after* all of the taps and passives have been upgraded.

Traffic analysis of the FDX bandwidth has shown that

1) The initial intuitive instinct to minimize the size of an Interference Group turns out, in practice, to be evidentially unfounded.
2) Large TGs can be supported while maintaining customer QoE, similar to how oversubscription models have worked for operators for decades of broadband services.
3) The TG sizes determined for the multi-gigabit symmetric speed tiers of interest align well with current node sizing and the expected network augments in the Comcast upgrade plan in the years ahead.
4) Capacity, speed, penetration, and TG relationships can be used, as they are similarly used today, to provide guidance to network operators' network augmentation and business (speed) planning

With these findings, and with the innovative development of FDX amplifiers already in the works and showing promise, the industry can now feel confident that multi-gigabit symmetrical speeds can quickly be enabled in their existing N+x HFC deployments, founded on a deep understanding of the practical realities of burst traffic engineering and how it applies in the FDX band.

Abbreviations

ABR	Adaptive Bit Rate
BW	Bandwidth
CAGR	Compounded Annual Growth Rate
CAPEX	Capital Expense
CCAP	Converged Cable Access Platform
CDF	Cumulative Distribution Function
CM	Cable Modem
CMTS	Cable Modem Termination System
CPE	Consumer Premises Equipment
D3.1	Data Over Cable Service Interface Specification 3.1
D4.0	Data Over Cable Service Interface Specification 4.0
DAA	Distributed Access Architecture
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
DSLAM	Digital Subscriber Line Access Multiplexer
DSP	Digital Signal Processing
EC	Echo Cancellation
EOL	End of Line
EPON	Ethernet Passive Optical Network (aka GE-PON)
ESD	Extended spectrum DOCSIS
FDD	Frequency Domain Duplex
FDX	Full Duplex DOCSIS
FDX-L	Full Duplex DOCSIS Light
FTTH	Fiber to the Home
FTTx	Fiber to the 'x' where 'x' can be any of the above
Gbps	Gigabit per second
GHz	Gigahertz
GOA	Grey Optics Aggregator
GOT	Grey Optics Terminator
HEO	Headend Optics
HFC	Hybrid Fiber-Coax
HP	Homes passed
HSD	High Speed Data
HW	Hardware
I-CCAP	Integrated Converged Cable Access Platform
IG	Interference Group
IP	Internet Protocol
IPTV	Internet Protocol Television
LDPC	Low Density Parity Check (FEC code)
LE	Line Extender
MAC	Media Access Control
MB	Multi-port Bridger
Mbps	Megabit per second
MHz	Megahertz
MSO	Multiple System Operator

N+0	Node+0 actives
N+ x	Node + x actives (amplifiers)
NCTA	The Internet & Television Association
OFDM	Orthogonal Frequency-Division Multiplexing
OFDMA	Orthogonal Frequency-Division Multiple Access
OPEX	Operating Expense
PDF	Probability Density Function
PHY	Physical interface
PMF	Probability Mass Function
PON	Passive Optical Network
PSD	Power Spectral Density
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
RBA	Resource Block Assignment
RF	Radio Frequency
RPHY	Remote PHY
SC-QAM	Single Carrier QAM
SDV	Switched Digital Video
SG	Service Group
SCTE	Society of Cable Telecommunications Engineers
SNR	Signal to Noise Ratio
STB	Set-top Box
Tavg	Average bandwidth per subscriber
TCP	Total Composite Power
TG	Transmission Group
Tmax	Maximum sustained traffic rate – DOCSIS Service Flow parameter
TX	Transmit
US	Upstream
xDSL	Digital Subscriber Line, unspecific type

Bibliography & References

Technical Papers

[1] Cloonan, Dr. Tom et. al., Advanced Quality of Experience Monitoring Techniques for a New Generation of Traffic Types Carried by DOCSIS, NCTA Spring Technical Forum 2013, NCTA

[2] Cloonan, Dr. Tom, John Ulm, Ayham Al-Banna, Frank O'Keefe, and Ruth Cloonan, Capacity Planning, Traffic Engineering, And HFC Plant Evolution For The Next 25 Years, SCTE Cable-Tec Expo, Sept 30-Oct 3, 2019, New Orleans, LA.

[3] Cloonan, Dr. Tom et. al., The Big Network Changes Coming with 1+ Gbps Service Environments of the Future, SCTE Cable-Tec Expo 2017.

[4] Cloonan, Dr. Tom et. al., Simulating the Impact of QoE on Per-Service Group HSD Bandwidth Capacity Requirements, SCTE Cable-Tec Expo 2014.

[5] Cloonan, Dr. Tom et. al., Using DOCSIS to Meet the Larger BW Demand of the 2020 Decade and Beyond, NCTA Spring Technical Forum 2016, NCTA

[6] Howald, Dr. Robert and Jon Cave (Comcast), Olakunle Ekundare (Comcast), John Williams (Charter), Matt Petersen (Charter), Developing the DOCSIS 4.0 Playbook for the Season of 10G, 2021 SCTE Expo Oct 11-14 (Virtual Event).

[7] Howald, Dr. Robert, Roaring into the 20's with 10G, 2020 SCTE Expo, Oct 13-16 (Virtual Event).

[8] Howald, Dr. Robert L, and Dr. Sebnem Ozer, Robert Thompson, Saif Rahman, Dr. Richard Prodan, Jorge Salinger, What is 10G – The Technology Foundation, 2019 SCTE Expo, Sept 30-Oct 3, New Orleans, LA.

[9] Prodan, Dr. Richard, 10G Full Duplex DOCSIS Implementation Exceeds Expectations, 2021 SCTE Expo Oct 11-14 (Virtual Event).

[10] Prodan, Dr. Richard, Optimizing the 10G Transition to Full Duplex DOCSIS 4.0, SCTE Expo, Oct 13-16, 2020 (Virtual Event).

[11] Ulm, John and Zoran Maricevic, Ram Ranganathan; Broadband Capacity Growth models – will the End of Exponential Growth eliminate the need for DOCSIS 4.0?, SCTE Cable-Tec 2022, Sept 19-22, Philadelphia, PA.

[12] Ulm, John and Dr. Tom Cloonan, The Broadband Network Evolution continues – How do we get to Cable 10G?, 2019 SCTE Expo, Sept 30-Oct 3, New Orleans, LA.

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

[14] Ulm, John and Zoran Maricevic, 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

[15] Ulm, John and Dr. Tom Cloonan J. Ulm, T. J. Cloonan, Traffic Engineering in a Fiber Deep Gigabit World, 2017 SCTE Expo, Oct 17-20, Denver, CO.

[16] Wall, Dr. Bill, Practical Considerations For Full Duplex Deployments In N+x Environments, 2019 SCTE Expo, Sept 30-Oct 3, New Orleans, LA.

Press Releases

[17] *Comcast Demonstrates Fastest-Yet Speeds Over a Complete 10G Connection on a Live Network*, Comcast Corporate Press Release April 28, 2022.

[18] *Comcast Announces World-First Test of 10G Modem Technology Capable of Delivering Multigigabit Speeds to Homes*, Comcast Corporate Press Release Jan 13, 2022.

[19] *Announcing Another 10G Milestone Amidst a Flurry of Innovation*, Comcast Corporate Press Release Oct 14, 2021.

[20] *Comcast Conducts Groundbreaking Test of 'Full Duplex' Chip that Will Support Future Multigigabit Upload and Download Speeds*, Comcast Corporate Press Release Apr 22, 2021.

[21] *Comcast Reaches 10G Technical Milestone Delivering 1.25 Gig Symmetrical Speeds Over a Live, All-Digital HFC Network*, Comcast Corporate Press Release Oct 8, 2020.

Other

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

