Managing the Coronavirus Bandwidth Surge

How to Cope with the Spikes and Long-term Growth

A Technical Paper prepared for SCTE•ISBE by

John Ulm
Engineering Fellow, Broadband Systems
CommScope – CTO Network Solutions team
john.ulm@commscope.com

Dr. Thomas Cloonan
CTO – Network Solutions
CommScope
tom.cloonan@commscope.com

Table of Contents

Title **Page Number**

List of Figures

1. Introduction

"With all due respect, sir, I believe this will be our finest hour." That quote from the movie *Apollo 13* is strangely appropriate for what happened with the broadband industry after a global lockdown at the start of the COVID-19 pandemic. The Coronavirus pandemic turned the whole world upside down, with entire countries forcing their population to live, work, learn and play from home. There was a year's worth of bandwidth growth (or more) in a 2-week interval. The broadband industry held its breadth to see how our broadband infrastructure would cope.

Broadband kept society and the economy running (as best we can) in these hard times. This is arguably its most significant contribution to society in its short life-span. Hats off to everyone who has created and helped make broadband and the Internet available to society.

In this paper, the bandwidth (BW) impact to our broadband networks from around the globe is reviewed and showed how it was handled. The upstream got crunched much more than the downstream. Network capacity planning was key to having sufficient headroom to withstand an unexpected jolt. Some network capacity guidelines are reviewed and show how the Quality of Experience (QoE) level varies with different margins.

For those severely congested networks, some helpful tips when in crisis are provided. These overnight quick fixes can give a little breathing room until more permanent capacity can be added. The relative merits of some near-term solutions to deploy over coming months are discussed. These include: more DOCSIS channels, especially OFDM/OFDMA, more CCAP ports, segment congested nodes, rapidly increase DOCSIS 3.1 (D3.1) modem penetration, and deploying Wi-Fi 6 services.

Some bandwidth projections and mid- to long-term migration strategies are reviewed that operators should consider moving forward before the next bandwidth surge hits. These include: Migrate to 1218/85 or 1218/204 MHz plant today, with a transition to DOCSIS 4.0 over time; reduce legacy Video QAMs using IPTV/SDV/compression; Fiber Deeper; DAA; and Wi-Fi 6E. It is more important than ever to plan and start to implement our long-term strategies like 10G, especially for the upstream – cable's Achilles heel.

Finally, the longer-term impact of the pandemic is considered. What is the new normal? We don't expect bandwidth levels to ever go back to pre-COVID days. We'll touch on how user's behavior patterns changed and what new applications had the door opened.

2. Pre-Pandemic – the Calm before the Storm

Before we dive into the bandwidth impacts of the pandemic, let's take a look at the world pre-COVID. CommScope/ARRIS has been monitoring subscriber usage for over a decade now from the same group of MSOs. The data from this set has been compared and maps closely to many other MSOs globally.

The chart below, Figure 1, shows the average subscriber downstream consumption, Downstream (DS) Average BW per subscriber (Tavg), during peak busy hours for a number of Multiple System Operators (MSOs) over a ten year period. At the start of 2020, DS Tavg had surpassed the 2 Mbps barrier. It turns out that the Tavg growth rate was higher at the start of this decade and has tailed off a bit in recent years. Over the last 3-4 years, this group of MSOs had an average downstream traffic growth that had been around 30%.

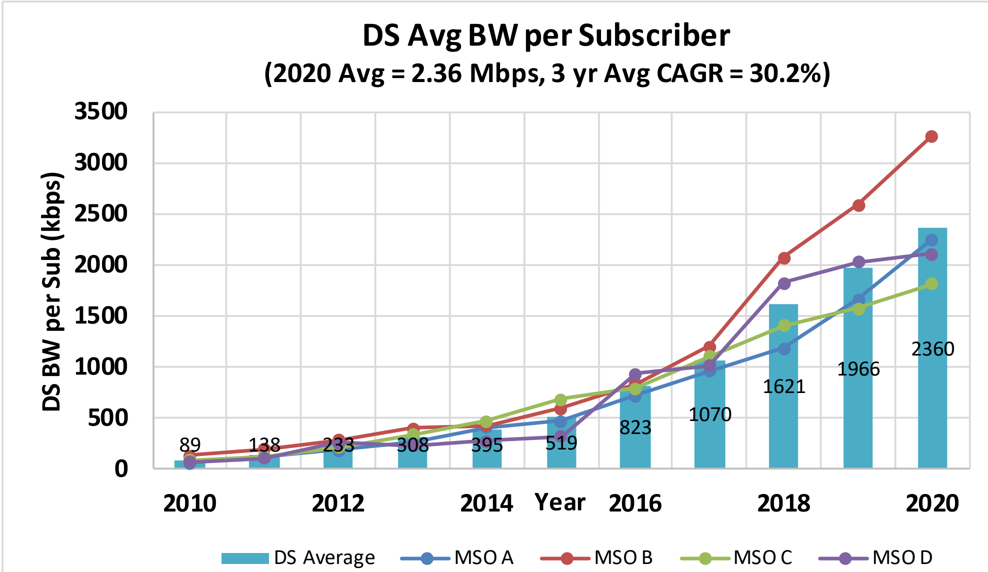

Figure 1 – Downstream Average Bandwidth per Subscriber through Jan '20

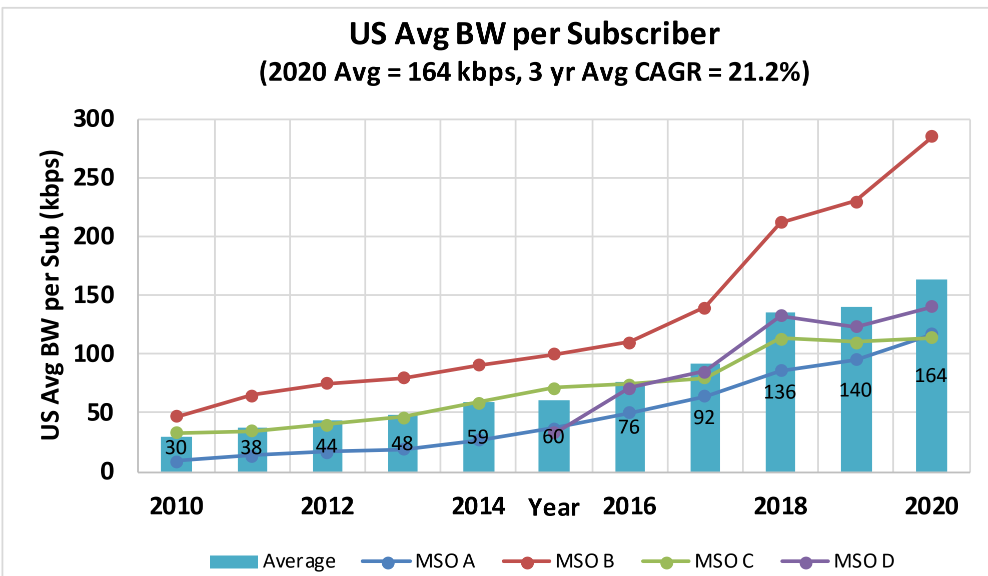

Figure 2 – Upstream Average Bandwidth per Subscriber through Jan '20

Interestingly, the upstream (US) traffic is growing at a significantly slower rate than the downstream as shown in Figure 2. During the same ten year period, the upstream Tavg generally grew at less than 20% compound annual growth rate (CAGR).

Over the years, traffic has also become more asymmetric with video applications driving downstream consumption [EMMEN_2014]. The average DS:US ratio is about 14:1, the MSO with the largest DS:US ratio seems to have stabilized around an 18:1 ratio. The MSO with the lowest DS:US ratio was about 11:1 ratio.

3. COVID-19 Pandemic hits – the Storm Surge

During mid-March, our lives got turned upside down and the entire world suddenly had to work, learn and play from home. To monitor the impact of the Coronavirus BW surge, we had access to live data through the CommScope ServAssure program and were able to monitor 1000's of nodes to study the effects.

Below are some very representative data samples taken during late March to demonstrate some of the changes seen. Broadband usage started picking up earlier in the morning and slowing down later at night such that the lull in the wee hours of the morning shrunk. Daytime usage was up significantly. Some networks saw daytime downstream usage increase as much as 86% while daytime upstream usage increased as much as 150%. The shape of the upstream was also dramatically altered. The upstream would reach its peak around 10am and then stay plateaued for the rest of the day until about midnight.

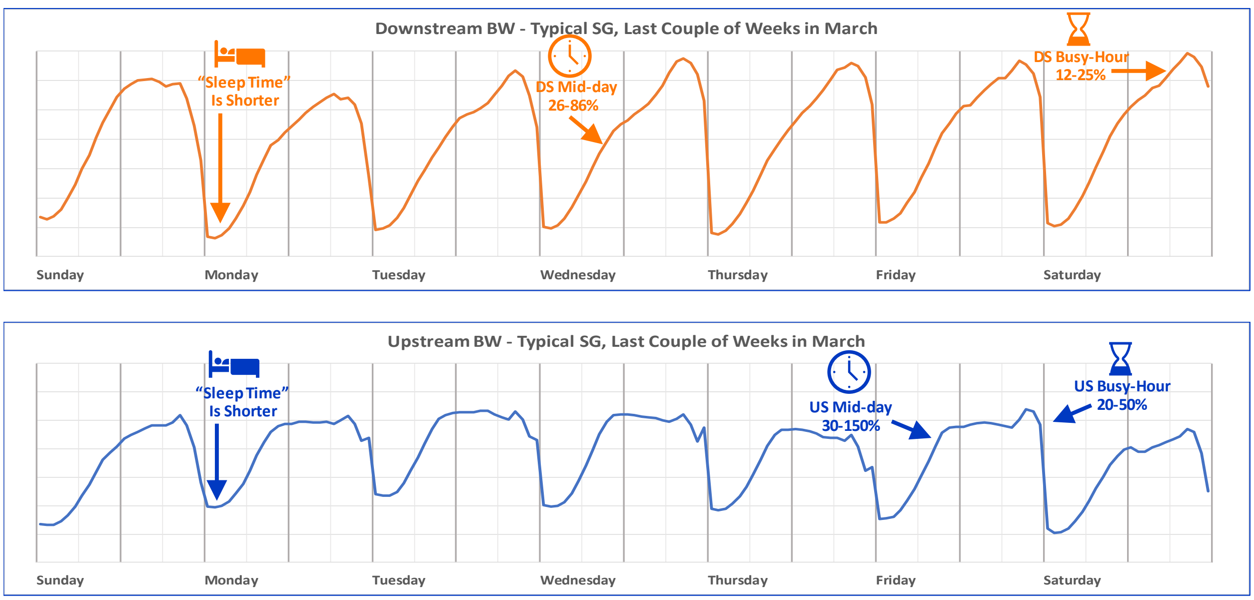

Figure 3 -Coronavirus BW Surge impacts across Time of Day

While the staggering daytime usage increases made for good press, it did not impact the overall network capacity requirements. Networks must be designed for peak utilization. The daytime usage was starting from a much lower point, so its percentage increase is exaggerated. Our focus going forward will be on the increase to the peak busy period for the day. For the above nodes that we monitored, we saw peak downstream increases in the 12-25% range while peak upstream increases were even higher in the 20-50% range.

3.1. NCTA COVID-19 Dashboard website

During the pandemic, the NCTA (The Internet & Television Association) was particularly proactive in publishing its members' network status. It collected nationwide information from the following MSOs:

- AlticeUSA, CableOne, Charter, Comcast, Cox, GCI, Mediacom, Midco, Sjoberg's

Altogether, this group represents 94% of the entire U.S.A. cable subscriber base. The NCTA published its data on the website: https://www.ncta.com/COVIDdashboard. The National Upstream Peak Growth and National Downstream Peak Growth charts are shown in figure 4 below.

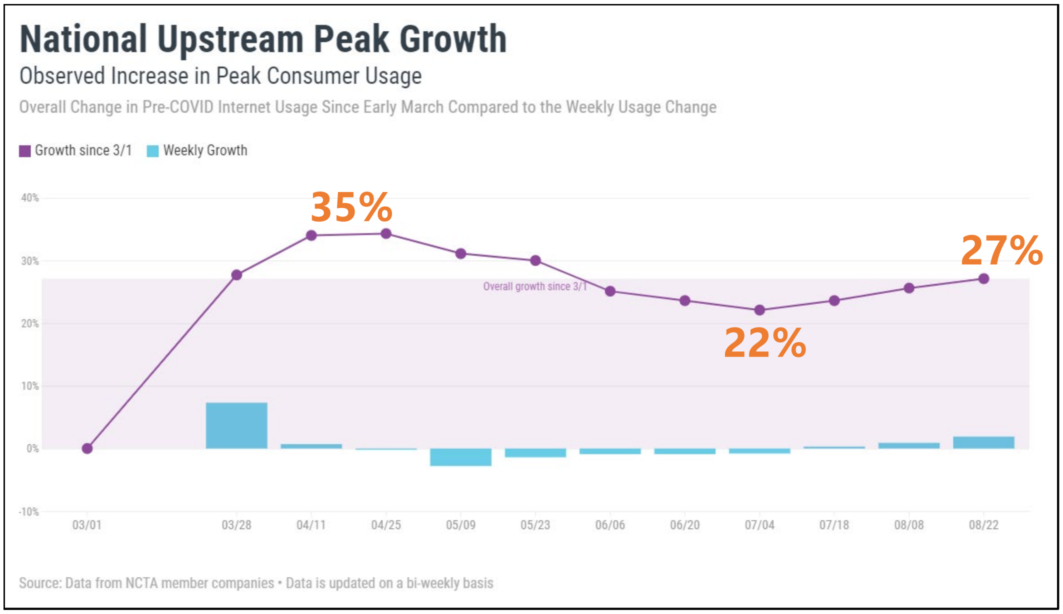

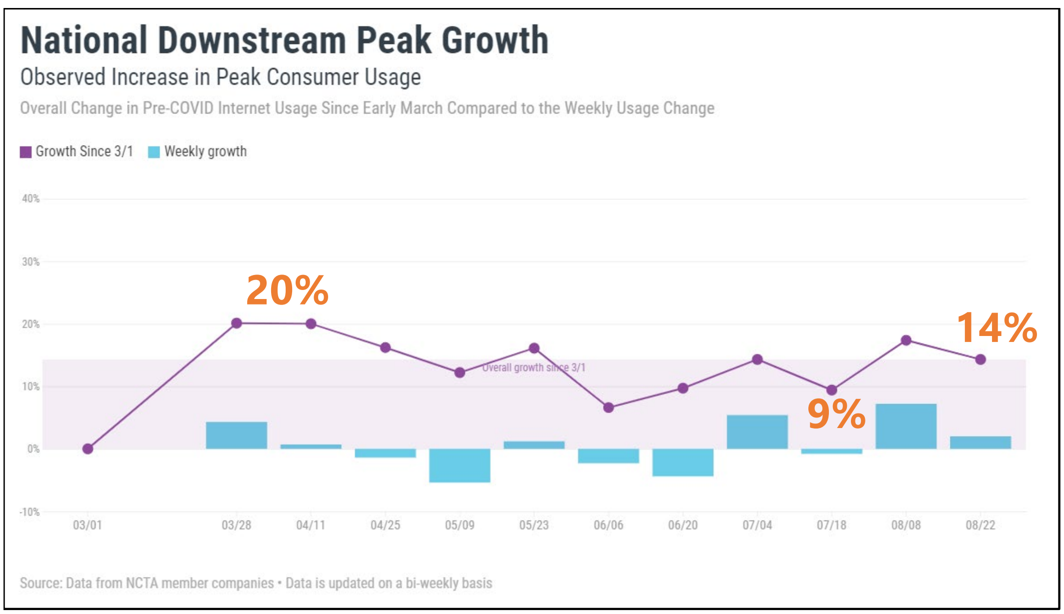

Figure 4 – NCTA COVID-19 Dashboard: National Peak Internet Growth

These charts show the change in peak consumer usage relative to March 1st, before the pandemic lockdowns began. By early April, the upstream peak was up 35% while the downstream peak was up 20%. After flattening out for a bit, these peak increases started to gradually drop. As of mid-July, the upstream peak was still 24% higher than March 1st, while the downstream peak was down to just a 9% increase from the start. There has been a slight uptick with the end of summer.

At this point, it is not clear if the dip in the peak increases is a result of the gradual re-openings happening around the country, or whether this is just a typical summer slow-down as people start to do more outside activities instead of being in front of their computers. It will also be interesting to observe the behavior of the above curves in the Autumn of 2020 with many schools planning to re-open with both face-to-face and on-line classes; this may result in an interesting bandwidth surge with more ubiquitous push towards on-line learning than was carried out in the Spring of 2020. Another interesting observation in the above curves is that the downstream peak increase of 9% after four months is actually in line with a 30% annual growth rate. So the downstream may be close to being back to normal levels. The upstream peak is definitely still elevated considering its historical 20% annual increases.

The NCTA website provides a wealth of additional information. It shows a dashboard with peak usage network status for the nation or for a given state. A snapshot of the National dashboard for mid-May is shown in figure 5. As can be seen, 97% of networks nationally were within their normal peak usage range with 2.8% at an elevated peak usage (i.e. yellow zone). Only 0.1% of networks were in the orange zone of substantially elevated peak usage and hardly any were in the red zone with significant elevated peak usage. This shows how well our broadband networks fared.

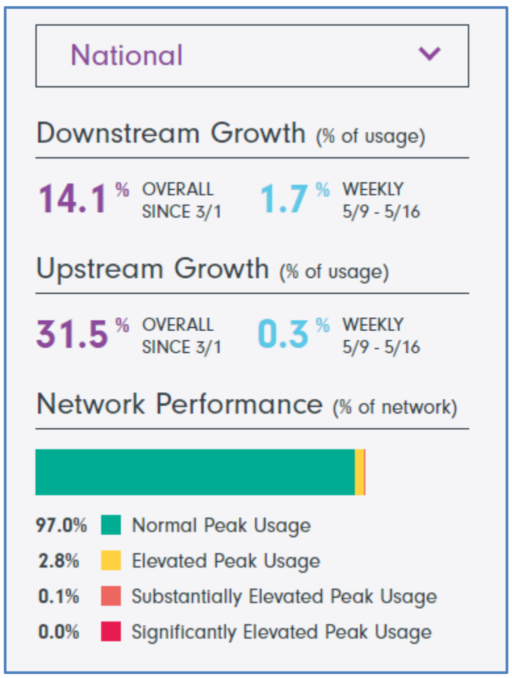

Figure 5 – NCTA COVID-19 Dashboard: Peak Usage status as of mid-May

The NCTA website also provided some additional key takeaways. Here were some given in mid-May:

- National US peak ***growth remains mostly flat*** with slight dip from 35% Peak
- National DS peak ***growth receding*** over last 2 months from 20% Peak
- Provider backbone networks ***have significant capacity***
    - Show no signs of congestion
- US peak hours in many regions have shifted from late evening to afternoon
- Wi-Fi data traffic & Wi-Fi calling are increasing as compared to mobile
- Networks are supporting more Wi-Fi-connected devices

3.2. Comcast COVID-19 Dashboard website

During the first few months of the pandemic, Comcast also published information on their own website, https://corporate.comcast.com/covid-19/network. Figure 6 below shows the Comcast Bandwidth Chart from mid-May. The upstream peak growth was holding fairly steady around 30-32%, while the downstream peak growth shot up to ~20% in late March and then slid back down to ~10% in mid-May.

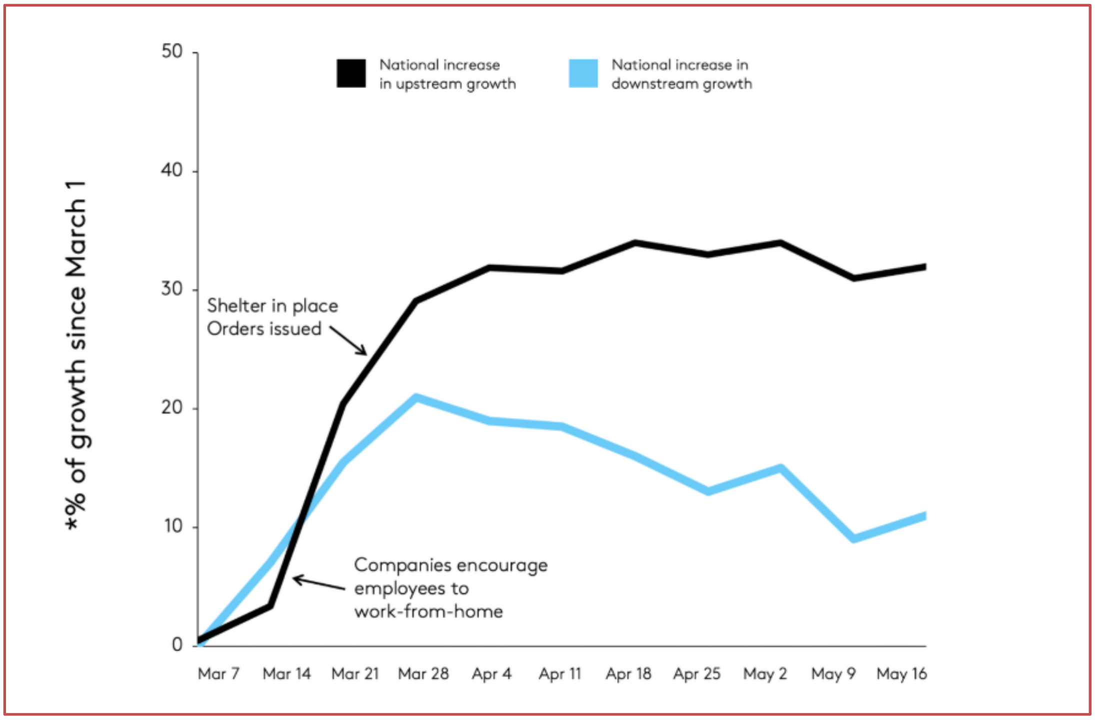

Figure 6 – Comcast COVID-19 Bandwidth Chart

Some observed effects on bandwidth as seen by Comcast include:

- Weekday usage is up:
  - VoIP & Video Conferencing is up 210-285%
  - VPN traffic is holding steady, up 30-40%
- Evening & weekend usage is up:
  - Gaming downloads are up 20-35% generally, up to 80% during new releases
  - 20-40% increase in streaming and web video consumption
  - Linear video consumption increased +2 hours per day per household
  - Video OnDemand (VoD) hitting record highs, up 50% YOY
- Xfinity Mobile sees a 36% increase in mobile data usage over Wi-Fi
  - But a 17% decline in LTE Data usage

Comcast also provided a number of interesting quotes, some of which are shown in figure 7. Of particular note is the quote from Comcast President Tony Werner: "We engineer the network to handle spikes and shifts in usage, and what we have seen so far with COVID-19 is within our capacity". The authors have seen this numerous times –

> *Networks that have been designed with sufficient capacity have been resilient and has handled the Coronavirus BW surge.*

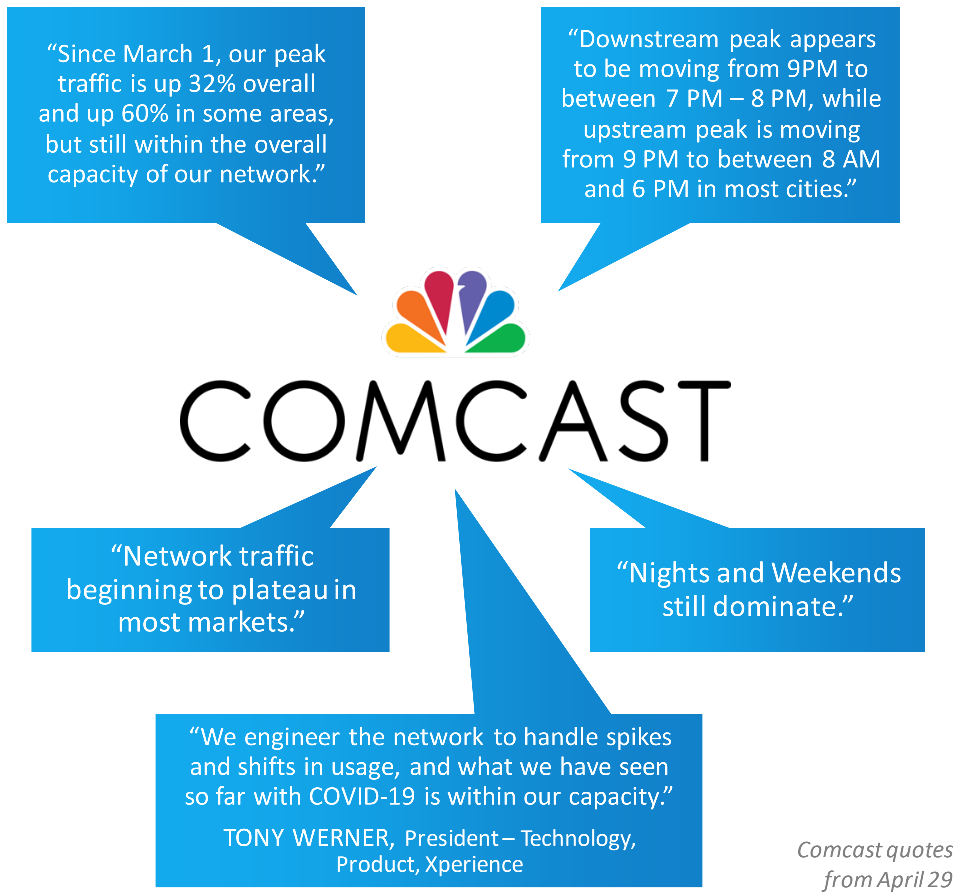

Figure 7 – Comcast COVID-19 Quotes

3.3. Other data gathered from various MSOs around the Global

The CommScope ServAssure program allowed us to collect data from multiple MSOs from around the globe. We were able to look at data from before and after the start of the pandemic.

One set of data came from a North American metro-area and covered 1000's of nodes. From that we were able to get the network utilization levels for every service group during the four weeks in March. The relative downstream (DS) service group (SG) utilization levels are shown in figure 8. The left side of the chart shows the entire node population while the right side zooms in on the "yellow" to "red" regions.

During the first week before the lockdown, almost 90% of the SGs were in the "green", i.e. <65% DS utilization. These networks were in great shape and in no immediate need for additional capacity. Around 10% of SG were in the "yellow" or "orange" zones (i.e. between 65% to 80%) indicating that they might need extra DS capacity in the next 3-9 months. Only ~1% of SG were in the "red" that needed immediate attention (i.e. "red" means >80% utilization).

As utilization increased dramatically with the pandemic lockdown, there was a noticeable change in the DS SG utilization numbers by the last week in March. The number of DS SG in the "green" dropped to about 80%, while DS SG in the "yellow" almost doubled to 18%. The DS SG in the "red" tripled to around 3%.

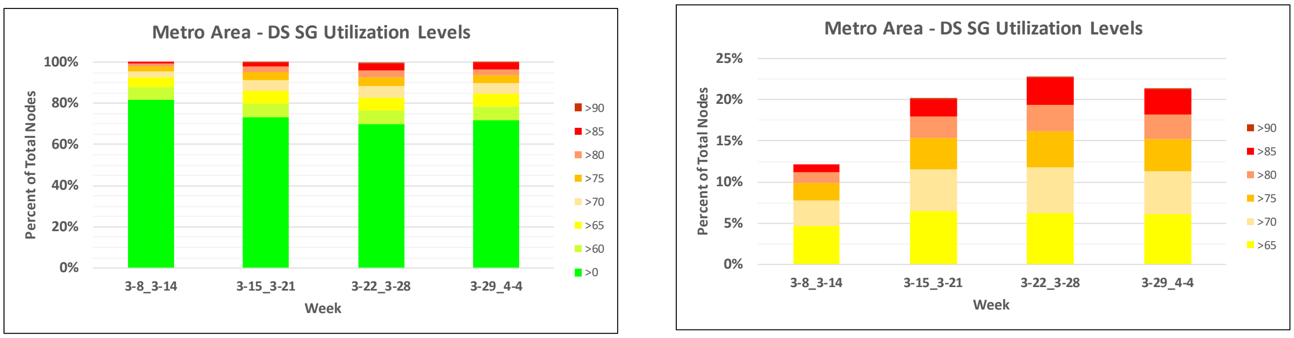

Figure 8 – Downstream Service Group Utilizations from No. American Metro-area

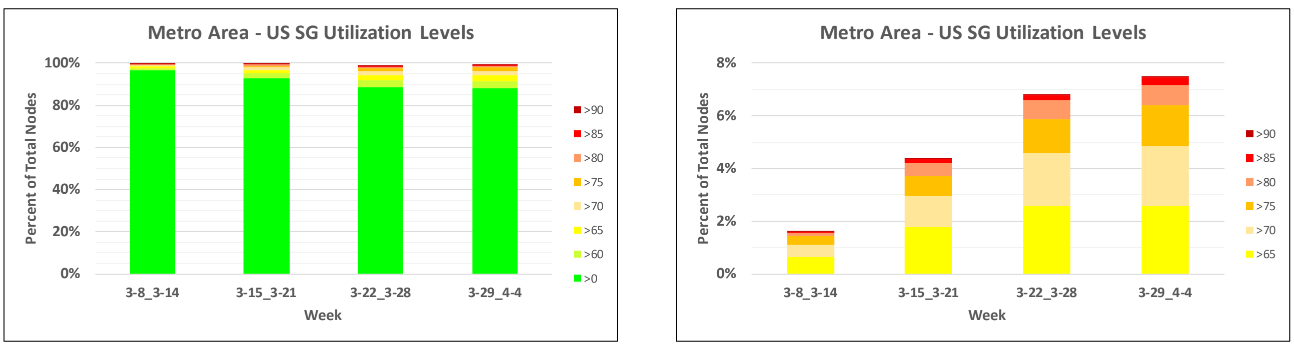

Figure 9 – Upstream Service Group Utilizations from No. American Metro-area

Figure 9 provides the upstream (US) service group (SG) utilization for the same North American metro-area. The entire node population is shown on the left while the right chart zooms in on the SG in the "yellow" to "red" zones.

This particular MSO has been meticulous with plant maintenance and reducing cascade lengths. Most of their upstream networks have at least four DOCSIS ATDMA channels. So in general, they have made sure there is plenty of upstream capacity. The early March data bares this out as approximately 98% of the US SG were in the "green"; ~1.5% were in the "yellow" and less than 0.25% of US SG were in the "red".

As the pandemic hit, it took its toll. The "green" US SG dropped to almost 90%. The number of US SG in "yellow" quadrupled to ~6% while "red" US SG also quadrupled to 1%. While these percentages look low, this is across 1000's of nodes that were impacted. The number of nodes that needed immediate attention increased dramatically, as well as the number of nodes needing extra capacity in the next 3-9 months.

Figure 10 takes a closer look at the downstream impacts using a typical DS SG. The chart on the left shows the utilization levels for different hours during the day and how it changed from week to week during March. The chart on the right shows similar utilization levels for the evening hours.

Notice how the DS utilization builds during the course of the day, from 10am to 2pm to 4pm to 6pm. By the 6-7pm hour, the DS SG is approaching its peak busy period. In general, 9-10pm was the busiest hour and used for our Tavg calculations. Although, 8-9pm and 10-11pm were very close. It shows that the DS peak busy period is close to a 3 hour window. There was a noticeable drop-off in utilization after 11pm.

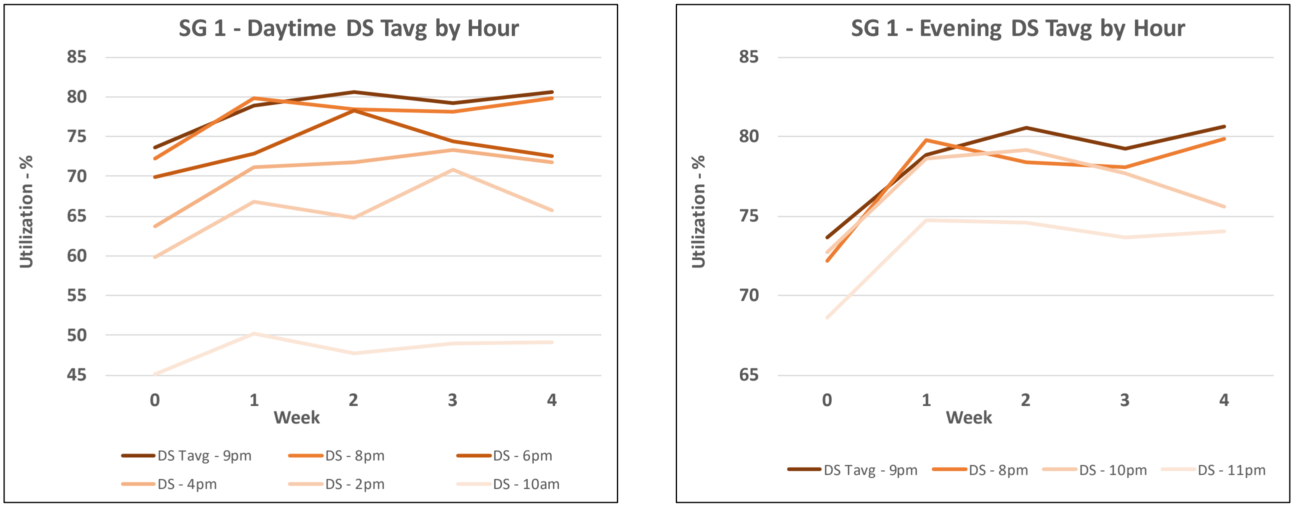

Figure 10 – Downstream Service Group Utilizations per Hour from N.A. Metro-area

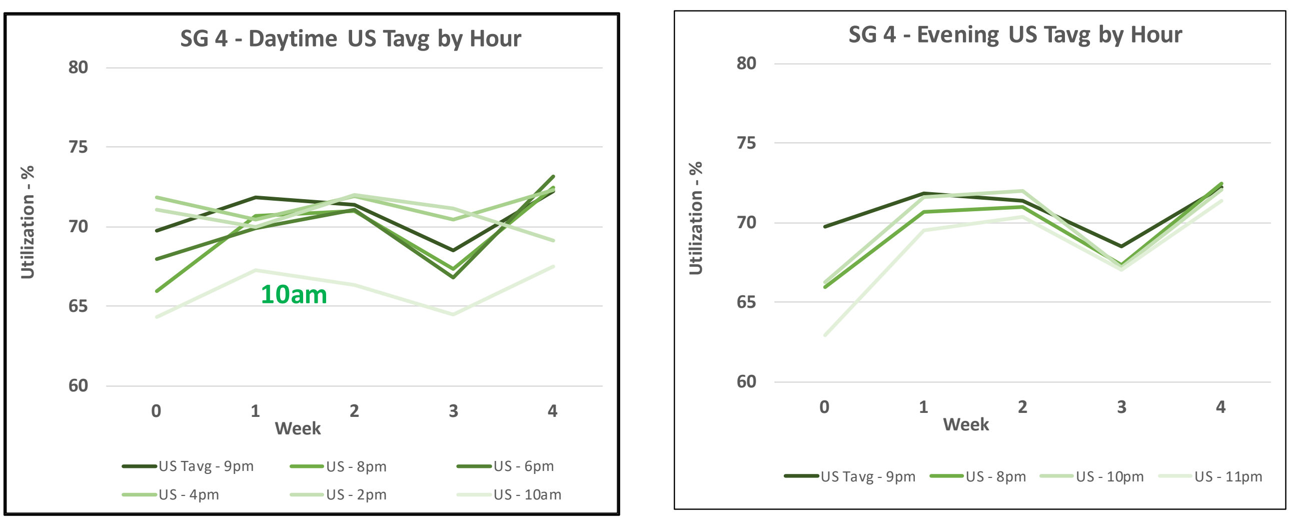

Figure 11 – Upstream Service Group Utilizations per Hour from N.A. Metro-area

Figure 11 takes a closer look at the upstream impacts using another typical US SG. The chart on the left shows the US utilization levels for different hours during the day and how it changed from week to week during March. The chart on the right shows similar US utilization levels for the evening hours.

Right away, the differences from the DS utilizations in figure 10 are obvious. By 10am hour, the US utilization is up to 65% and then the rest of the day it stays near peak levels. A closer look at the evening hours on the right shows how steady the US utilization remains for the entire evening. The peak busy period is now 11am to 11pm for this SG.

4. The Cause of the Surge – the usual Suspects

Taking a closer look at the various applications that caused the Coronavirus BW surge, it turns out it is basically just a lot more of things that are already been seen.

4.1. Work @ Home

The first very obvious impact of the lockdown was the significant increase in work@home employees. Some estimates showed that pre-COVID only 5%-10% of the workforce would work@home on a full or part-time basis. Overnight with the lockdown, this number skyrocketed to the 50%-90% range depending on locale. Working from home increases the use of Virtual Private Networks (VPN) and file transfers. Video conferencing would also jump up using programs such as WebEx, Microsoft Teams, Skype, Zoom, etc.

In general, many of these work applications are symmetric in nature, so they would tend to have a much bigger impact on the upstream capacity. Many current work environments are leveraging a cloud infrastructure and common applications such as Microsoft Office will routinely auto-save files to the cloud. This has big ramifications as employees work from home. If a worker is updating a 20MB PowerPoint or Word document, the cloud-sync function may cause it to get uploaded a dozen times over an hour of working. Again, adding additional burden to our upstream networks.

4.2. Remote Learning

Schools were one of the first institutions to close down when the pandemic hit, and school districts scrambled to set up their remote learning programs. Zoom and Skype became two of the favorite video conferencing methods to reach students during the day. Again, a symmetric application putting additional burdens on the upstream.

YouTube usage also saw a big uptick as a way for teachers to distribute their lesson plans. And with everyone at home, YouTube became a big source of educational material as people researched the pandemic and looked for ways to pass the time.

4.3. Video Streaming

As was noted by Comcast above, viewers had more time on their hands and hence the viewing hours per week went up. Many video content providers offered free trial subscriptions to entice folks to sample their content. It remains to be seen how many of these consumers will stick with the service after the trial is over.

4.4. Social Networking

With chatting around the water cooler removed, there was a definite increase in social networking apps as folks looked at ways to stay in touch. Facetime, Skype and Zoom sessions became very popular, not only to keep in touch with current contacts, but to re-establish with some long lost folks too.

4.5. Gaming

With everyone at home and more spare time available, gaming applications such as Twitch also a sparked increase. For many of today's on-line games, a download might be on the order of 100GB. There were several instances of a new release of a popular game that directly impacted network performance.

While, in general, gaming applications tend to impact the downstream more, they may often be accompanied by some social networking interaction and possibly some video conferencing.

The gamers are the ones that may have been impacted the most as these networks became congested and latencies increased.

5. QoE Impacts of the Surge – the Good, the Bad & the Ugly

With network congestion comes impacts on subscribers Quality of Experience (QoE). The Coronavirus BW surge introduced 12-18 months' worth of upstream BW growth in a 1 week period. The downstream was only slightly more forgiving with 8-12 months of BW growth when the lockdowns started.

Some of the negative QoE effects on subscribers include:

- Packet delays for gamers from higher latencies and buffer-bloat
- Video tiling in streaming & video conferencing from packet delays and drops
- File Download & Upload times increased (a little bit)

In addition to the above impacts, some subscribers who are working/learning/playing from home have insufficient service tiers to support their traffic load and/or are running up against any data caps that their service provider may have.

5.1. Upstream Capacity – Achille's Heel of Cable

And this heavy usage is lasting longer throughout the day. The upstream peak period now lasts 12-14 hours from mid-morning to midnight. Because the upstream pipe is much smaller while having the more significant BW surge, it has come under the most scrutiny. At this point, the upstream is considered the Achille's Heel of the cable industry. Current North American 42 MHz systems might support 60-100 Mbps from 3-4 DOCSIS 3.0 ATDMA channels. European 65 MHz HFC systems have a little more breathing room with up to 175 Mbps from up to 7 DOCSIS 3.0 ATDMA channels. But this capacity pales compared to the 5+ Gbps of downstream capacity found on an 870 MHz plant. Because of this, a lot of the focus on following sections will be on enhancing the upstream capacity.

5.2. BW Impacts as measured by Speed Tests

For many consumers, when network performance starts to lag, they run a speed test. One of the most common available is Ookla's www.speedtest.net.Ookla. Ookla tracks and publishes its speed test results from around the globe. They did a very good job of showing how every country fared during the pandemic:

https://www.speedtest.net/insights/blog/tracking-covid-19-impact-global-internet-performance/

Figure 12 shows a snapshot of the USA Internet performance based on the Ookla speed test. It covers a timeframe from December 2019 to May 2020. When the lockdowns first occurred in mid-March, the Fixed broadband networks recorded a 9% drop. Note that this was actually less than the drop seen during the holiday week of December 23-29.

The downstream internet performance then gradually recovered over the following weeks until it was only down about 1% from March 2nd until mid-May.

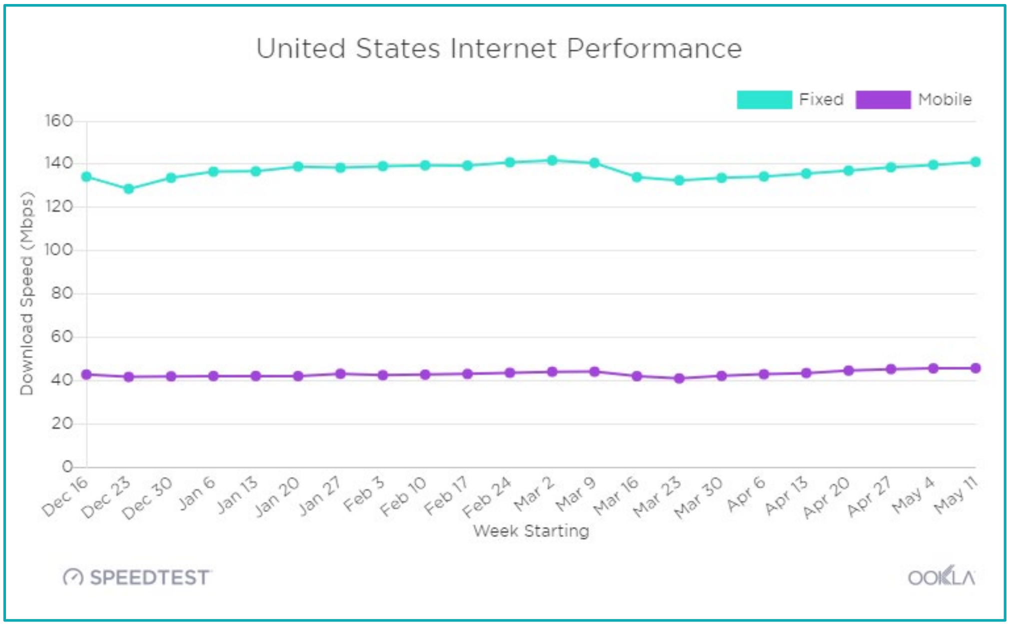

Figure 12 – Ookla Speedtest: U.S. Internet Performance during Pandemic

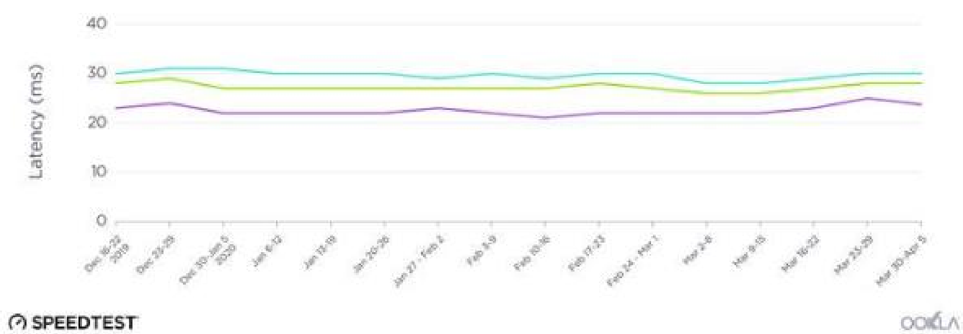

Figure 13 – Ookla Speedtest: U.S. Internet Latency during Pandemic

Figure 13 shows the latency measurements from Ookla speed test during this same period. Notice that there was no perceived impact on latencies due to the Coronavirus BW surge. More evidence that broadband networks are holding their own!

5.3. Handling Congestion – DOCSIS to the Rescue

Overall, most of these bad QoE effect have been relatively minimal to date for most service groups. Only a handful of SGs that were pushed to the limit before the BW surge had detrimental QoE. In general, the DOCSIS networks are holding up to this sudden, stressful packet load. Why, despite heavily congested upstreams causing delays and drops? Several reasons:

- Good network capacity planning added plenty of headroom
  - e.g. enough headroom to absorb the largest Service Level Agreement (SLA) bursts
- Excellent CMTS Scheduling algorithms
  - These are sophisticated AI Engines that have evolved over the last two decades
  - They are excellent at adapting to congestion, with fair BW distribution

- Most internet applications are elastic and forgiving; TCP and ABR recover from thruput reductions & packet loss

And it also doesn't hurt that subscribers are more tolerant about small lapses when they are more worried about the virus et. al.

To bring home this point of having good bandwidth capacity planning with an excellent CMTS scheduler, figures 14 and 15 show a before and after picture of a hypothetical upstream channel. These upstreams will have traffic from many different service flows (SF) from many different modems. These will present a fairly static load on the system. But at any instant in time, any one of these SF might burst up to its maximum SLA capacity. This is shown in the yellow spikes in the figures.

Before the BW surge with good network planning, then the available US capacity is sufficient to handle the largest SLA burst. Everyone is happy!

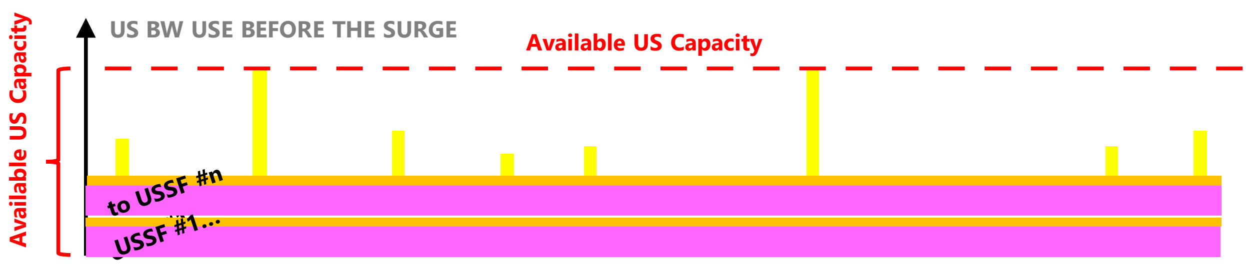

Figure 14 – Upstream Bandwidth depiction before the Coronavirus BW Surge

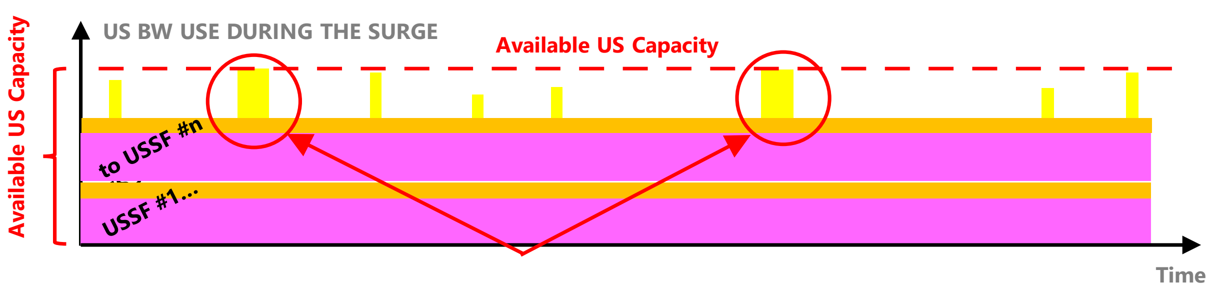

Figure 15 – Upstream Bandwidth depiction after the Coronavirus BW Surge

After the Coronavirus BW surge, many of the service flows increase their traffic load and the tide rises. There is no longer sufficient room to handle the largest BW bursts. This is shown in the red circles in the figure. This is also when the CMTS schedulers kick into high gear.

Great CMTS schedulers will throttle down these high bursts, delaying packets and perhaps dropping packets when needed. They are also looking across all SFs and prioritizing as needed. The bottom line is that these bursts get spread out slightly in time with barely a notice from the subscriber. Many subscribers with smaller SLA's may not see any impacts at all. This is the beauty of statistical multiplexing traffic and intelligent US scheduling during the Coronavirus BW surge.

6. Repairing Leaks during the Storm – Short-term Fixes

For those handful of SGs that became super-congested during the BW surge, operators were looking for some quick fixes that might gain an additional 5% or 10% or 15% of capacity gain. They also wanted something that they could implement in the matter of hours or days to relieve the immediate heartburn.

Below is a smattering of small things that might help. By themselves they may not add up too much, but collectively a handful of these might be enough to get an operator through the rough patch until more permanent capacity can be added to the system.

6.1. Simple I-CCAP related changes

Today's Integrated CCAP (I-CCAP) are very powerful and flexible networking devices with a lot of capabilities. The most straightforward solution to the BW crunch is to add more DOCSIS channels which can be done under software (SW) control using licensing. This would be the first choice solution provided there is spectrum available.

If the operator has a substantial mix of DOCSIS 3.0 and 3.1 modems, then they will need to consider whether to add 3.0 SC-QAM channels or 3.1 OFDM/A channels. In general, the 3.1 OFDMA US channels provide 65% to 100% better spectral capacity (Mbps/MHz) than 3.0 ATDMA channels. The 3.1 OFDMA channels are also more robust and can operate over a much wider frequency range, even well below 20 MHz in the "muck". In general, the 3.1 OFDMA channels have been under-utilized by many operators to date.

But even some 3.0 ATDMA channels could be operated below 20 MHz as well. The operator might need to use a lower modulation (e.g. 16-QAM or QPSK) and/or lower channel widths with increased FEC and interleaving. Some I-CCAP might have an integrated US agility capability to auto-adjust QAM levels.

If there is not sufficient available spectrum in the downstream, the 3.1 OFDM channel is robust enough to run in the roll-off region. E.g. put an OFDM channel from 750-942 MHz in a 750 MHz system. There might be 1 Gbps or more of capacity available in the roll-off.

In a very congested network, the CMTS scheduler is using the SF Tmax value (i.e. its SLA burst value) to fairly assign bandwidth. So, for example, maybe all SF will receive 60% of Tmax during congestion. But there might be a big discrepancy between the highest and lowest Tmax SLA. For example, the top tier in the upstream might be 50 Mbps while the lowest tier is 5 Mbps. The lowest tier will only get 10% of the capacity of the high tier. One strategy might be to temporarily boost the lowest tier (e.g. up to 20 Mbps) to ensure that these customers get adequate US performance to work successfully from home.

The DOCSIS CMTS has dozens of other configuration parameters that could be tweaked to help things. One could enable "Power-Boost" modes; configure Tmin parameters to provide minimum BW guarantees; adjust SF priorities; or tweak any scheduler parameters that might be available.

6.2. CM related changes

There are a handful of cable modem (CM) based configuration items that can be done to help. First, the operator should turn on the TCP Ack suppression on all modems if they haven't already. This will help reduce the number of TCP Acks in the upstream and every little bit helps when severely congested.

Buffer-bloat is another problem that has been studied for many years. This is the large increase in queue depths during congested intervals. This then results in very high latencies seen by the applications. The operator can enable some Buffer Control TLVs and turn on DOCSIS 3.1 Active Queue Management (AQM) to help suppress buffer-bloat and reduce latencies when congested.

6.3. In-Home changes

From the subscriber's home perspective, there are a number of items that fall under the quick fix umbrella. First, an operator might encourage subscriber to upgrade to a higher service tier. A new promotion that is geared towards families that are working at home and remote learning might gain a lot of traction while bringing in extra revenue.

With so many people in the home trying to actively use the internet simultaneously, many Wi-Fi routers may be caving under the load. An operator might provide some guidelines to their customers on improving the Wi-Fi router location within the house.

Finally, some I-CCAP might have capabilities like Integrated Service Class Agility (ISCA) that can dynamically reduce the Tmax value temporarily for any "bad actors".

6.4. Plant Optimizations

Sometimes operators will use a common configuration across all of their downstream and upstream CMTS ports. This may result in using the lowest common denominator and many SG may be leaving unused capacity on the table. For heavily congested SG, the operator may need to optimize the configuration parameters for those channels.

Depending on the Signal-to-Noise Ratio (SNR), an operator may be able to increase the QAM modulation and/or optimize the Forward Error Correction (FEC) for a given channel. For example on a cleaner plant, the US ATDMA channel might get bumped from 32- to 64-QAM modulation. An US OFDMA channel might go from 256- to 512- or even 1024-QAM modulation. In the DS OFDM channel, an operator might jump from 1024- to 4096-QAM.

The I-CCAP may ship with a default FEC setting that is very robust, but not efficient. Optimizing the US FEC could increase US capacity by as much as 13%. That may be enough to get that SG out of the dog house.

Warning – only change the QAM-modulation and FEC to values that can be supported by the SNR. Being too aggressive may introduced undetected packet errors causing dropped packets and making performance worse than before. An operator should use tools like the Intelligent Channel Optimizer (ICO) to optimize settings. For OFDMA channels, optimize bit-loading settings in Modulation Profiles using PMA tools (while still maintaining adequate Packet Error Rate, PER).

7. Cleaning up after the Storm – Near-term 6-12 month Strategies

The previous section talked about some quick fixes that could be done to relieve pressure on SGs in the "red" zone. This section takes a look at what can be implemented over the next 6-12 months to address SGs in the "yellow" zone. These can prepare an operator in case another BW surge arises.

7.1. CMTS related changes

Rinse and repeat – add more DOCSIS channels is the recurring theme here. Any additional capacity means adding more DOCSIS channels. The questions become: where does the spectrum come from and what is the mix of 3.0 and 3.1 channels.

Once the available spectrum is used including the roll-off region, then the operator will need to look at reclaiming spectrum from Legacy video QAMs; upgrading plants to 1218/85 or 1218/204 MHz; and/or segmenting service groups in either hubs or in nodes. Segmenting SG implies that the operator will have to add more CCAP ports to their system.

These options will be detailed below.

7.2. Legacy Video

For many operators, Legacy video still consumes a significant portion of their downstream spectrum. This is predominantly broadcast video QAM channels but also may include a handful of Video-on-Demand (VOD) QAMs.

Operators should consider converting any remaining older MPEG-2 encoded broadcast programs over to MPEG-4 encoding. This cuts their bandwidth requirements in half. This also means they will need to retire any MPEG-2 only settop boxes (STBs) that are still out there.

With all the improvements in encoding and stat-muxing technology, the operator may be able to increase the number of programs per QAM channel. For example, they may be able to jump from 2-3 HD programs per QAM up to 4-5 HD programs per QAM.

Perhaps the most powerful tool available to the operator for reducing Legacy video spectrum is Switched Digital Video (SDV). This technology has been around for 15 years but has undergone a bit of a renaissance by migrating to cloud-based infrastructures. This mature technology is now more cost effective than ever. SDV plays well with I-CCAP that integrate the SDV video QAMs directly into it. You can turn on an SDV system in a matter of weeks. And by aligning the SDV group with the shrinking DOCSIS SG, the operator could free up enough spectrum for a pair of 192 MHz OFDM channels.

7.3. CM & In-Home changes

The operator’s biggest weapon in its arsenal is DOCSIS 3.1. But it doesn’t help if there are not enough 3.1 modems in the field. The first priority for the operator should be to upgrade ALL modems for high tier customers (e.g. 200Mbps+ DS). By getting all the high tier customers onto 3.1 OFDM/OFDMA channels, they get better service. And this off-loads the potentially congested 3.0 SC-QAM channels for the lower tiers.

As mentioned before, the upstream is the cable operator’s Achilles heel. 3.1 OFDMA channels can more than double the upstream capacity in a 42MHz system. Operators should be rolling out more and more OFDMA channels. As the 3.1 modem penetration starts to overtake 3.0 penetration, then 3.0 SC-QAM channels can be reclaimed and converted to OFDMA spectrum.

Another big in-home improvement that an operator can make in the near-term is to start upgrading their customers to the new Wi-Fi 6 routers (i.e. 802.11ax). These can significantly increase the in-home Wi-Fi capacities and reach, which are both critical with so many people working from home.

The Wi-Fi 6 routers could be bundled with 3.1 modems for new service tiers that target work and learn at home families.

7.4. RF Plant Upgrades

There are a number of things that an operator can do with respect to the outside plant to enhance capacity in the near term. Perhaps the lowest hanging fruit with minimal investment is to segment existing nodes. Typically nodes are deployed with a single downstream module and a single upstream module. This configuration is referred to as 1x1 and maps to a single CMTS US port and a single DS port. The node may have two or four RF ports, and this DS and US are shared across all RF ports.

Some nodes with a digital return can easily be re-configured to support a second US. This effectively splits the node into 1 DS and 2 US or a 1x2 node. This will help relieve US congestion on that SG.

Later, another DS module can be added to the node to make it a 2x2 configuration. Eventually more modules can be added so the node evolves into a 4x4 configuration. All of this can be done at the node location without touching the rest of the outside plant. Additional CCAP ports will be needed in the hub site as well.

For severely congested SG, an operator might consider upgrading those first to 1218/85 or 1218/204 MHz plants. This is part of the longer term strategy and provides enough capacity for the rest of the decade. The 85 MHz return almost triples the amount of usable US spectrum. Combined with 3.1 OFDMA, it can support 400-500 Mbps US tiers. The 204 MHz return enables 1 Gbps US services. Upgrading all of your plants to 1218 MHz will take a while, so start today with those SG that are most congested.

DOCSIS 3.1 is very powerful and optimizes capacity according to the quality of the plant. It turns out that long fiber runs (e.g. >40km) with AM optics can have a detrimental effect on 3.1 OFDM capacity. This can be remedied by replacing those AM fiber runs with digital optics when Remote PHY devices (RPD) or Remote MAC-PHY devices (RMD) are placed in the node. And again, this might be part of a longer term MSO strategy to migrate to a complete Distributed Access Architecture (DAA).

8. Preparing for the next Storm – Network Capacity Planning

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

8.1. The "Basic" Traffic Engineering Formula

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

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

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

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

where:

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

The MSO data collected and shown in section 2 provides a good indication of Tavg, at least before the Coronavirus BW surge hit. Note that figures 1 and 2 are very generalized results that are averaged across millions of subscribers. A given operator will need to take a look at their own networks to ascertain what is an appropriate value for them. Regarding SG congestion, they need only look at the total SG consumption (i.e. Nsub*Tavg product term) and there is no need to break it down into the individual Nsub and Tavg components.

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

The second term in the formula (K*Tmax_max) has an adjustable parameter defined by the K value. This parameter allows the MSO to increase the K value and add bandwidth capacity headroom that helps provide better QoE to their subscribers within a service group. In addition, the entire second term is scaled to be proportional to the Tmax_max value, which is the maximum Tmax value that is being offered to subscribers. A change in the K value results in a corresponding change within the QoE levels experienced by the subscribers who are sharing the service group bandwidth capacity (C). Lower K values yield lower QoE levels, and higher K values yield higher QoE levels).

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

8.2. Providing Sufficient Headroom for a BW Surge

It will be useful to provide some examples on how the CommScope Basic Traffic Engineering formula can be applied to provide adequate headroom to sustain a BW surge. The first DS example is shown in figure 16. It assumes that the SG size is 100 subs. Before the BW surge (on the left in the figure), it uses a Tavg value = 2.36 Mbps taken from figure 1. The highest tier is 1G (i.e. Tmax_max = 1 Gbps) with a QoE constant K = 1.2. The basic formula says that the operator needs at least 1436Mbps to maintain that Kvalue. This might be a combination of 3.0 SC-QAM and 3.1 OFDM channels bonded together. The Nsub*Tavg component equals 236 Mbps (shown in blue) with the K*Tmax_max component = 1200 Mbps (shown in yellow).

After the surge is shown on the right hand side of figure 16. It assumes DS usage grows by 25% to Tavg = 2.95 Mbps. The Nsub*Tavg has now increased to 295 Mbps while overall capacity has remained fixed at 1436 Mbps. This reduces the K*Tmax_max component to 1141 Mbps which is an effective K value equal to 1.14. This is still in the "good" QoE range, despite seeing an instantaneous 25% increase in usage.

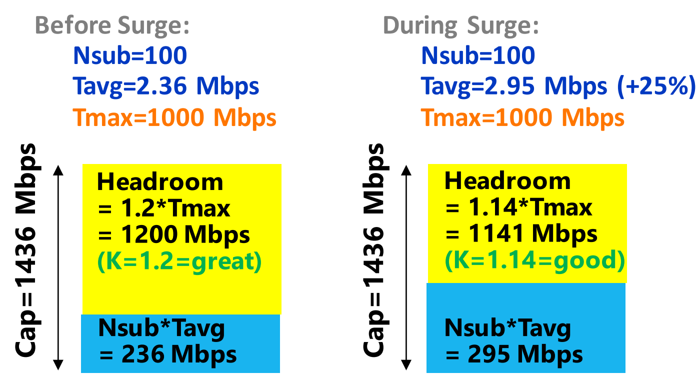

Figure 16 – QoE-based Traffic Eng Formula at Work during DS BW Surge, small SG

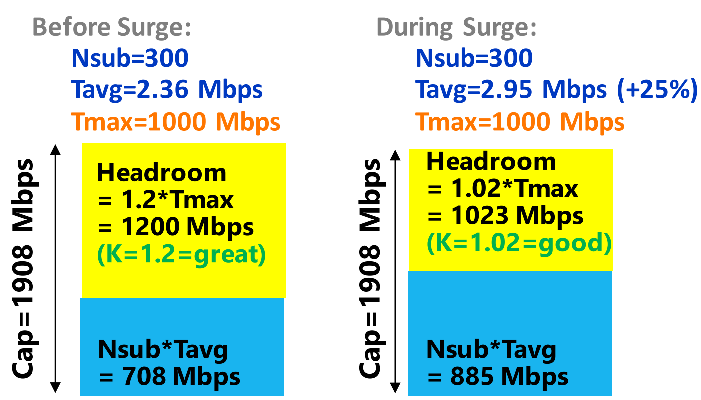

Figure 17 – QoE-based Traffic Eng Formula at Work during DS BW Surge, large SG

The second example in figure 17 increases the number of subscribers in the SG to 300 subs. Before the surge, this increases the Nsub*Tavg component to 708 Mbps. The overall required capacity is now increased up to 1908 Mbps to still maintain a QoE constant K = 1.2.

During the surge, the increased Tavg now has a bigger impact due to the increased SG size. The Nsub*Tavg component has now swelled to 885 Mbps. This now leaves 1023 Mbps for the K*Tmax_max component which provides an effective QoE constant K = 1.02. This is still in the "good" region, although near the bottom. So for larger SG, it is more important to have the cushion associated with K=1.2.

The third example in figure 18 takes a look at what happens if sufficient headroom is not applied. This example assumes the operator has rolled out the same 1436 Mbps overall capacity across their entire footprint, regardless of the SG size. This example looks at a SG with 250 subs. The Nsub*Tavg component before the surge is 590 Mbps. This makes for an effective K = 0.85 which is in the "fair" region. This means that the 1G customers will have good QoE much of the time, but may see some congestion periodically during peak busy periods.

Once the Coronavirus BW surge hits this SG, its Nsub*Tavg jumps to 738 Mbps. This reduces the K*Tmax_max component down to 551 Mbps which results in K=0.7. This is in the "poor" region. The 1G customers QoE impact becomes more noticeable.

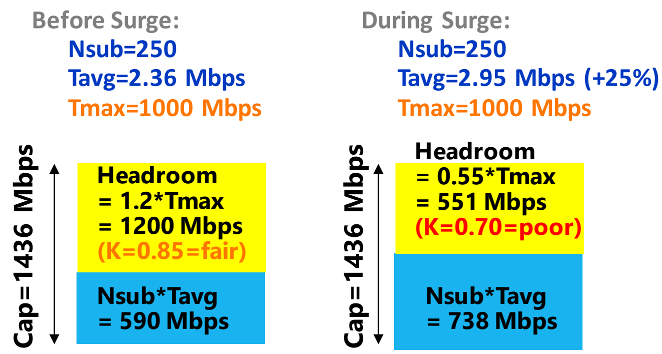

Figure 18 – QoE-based Traffic Eng Formula at Work for Marginal SG

It is important that operates not only monitor their network utilizations, but they must also understand the highest tier being offered and adjust their thresholds accordingly. For example, just having a utilization threshold of 80% to take action is not good enough by itself. That might have been adequate when the top tier was 100 Mbps, but then starts to fail when the top tier is 1 Gbps.

> *As the top tier is increased, the utilization threshold to take action must be correspondingly decreased to maintain consistent QoE.*

9. Building the Storm-proof Network – Mid- to Long-term Strategies

The Coronavirus BW surge is a glimpse at our future BW trends. We will not return to our pre-COVID days and will need to deal with a "new normal". This surge should serve as a wake-up call to make sure our industry continues to push its networks forward. We'll take a look at some mid-term solutions that can be implemented over the next couple years and then look at what will be coming later this decade.

9.1. Mid-term Solutions – for 1-3+ Years of Capacity Needs

Over the next couple years, operators need to focus on getting switched over to DOCSIS 3.1. This is the industry's workhorse right now. Operators should have plans in place to try and become 100% 3.1, phasing out their 2.0 and 3.0 modems. This can be done over a 5-8 year window, but needs to start in earnest now. Remember, DOCSIS 3.1 is your biggest weapon with your biggest bang for the buck.

9.1.1. CCAP related changes

As time goes on, operators should continue to maximize their 3.1 OFDM/OFDMA capacity. Any older CMTS should be upgraded to support the full 3.1 capabilities, including OFDMA in the upstream with both 85 and 204 MHz splits; and multiple 192 MHz OFDM channels in the downstream.

Over time, SG sizes will continue to shrink with node segmentations, so new CCAP ports will continually be added. Segmenting SG size is especially helpful in containing Tavg growth as was shown in the examples of the previous section.

During this timeframe as 3.1 penetration increases and high tier customers are moved to 3.1 modems, the operator should start considering switching some 3.0 SC-QAM channels to OFDM/OFDMA spectrum.

9.1.2. IP Video Migration

As part of the strategy to completely remove the Legacy Broadcast video spectrum, operators should have their IP video migration plans in full swing. It may take 5 to 8 years or more to get all of the Legacy STB swapped out for IP STB or other IPTV devices.

Ramping up the IPTV consumption while it coexists with Legacy Broadcast QAMs may create a sizable bandwidth bubble. The operator may also need to consider SDV or encoding enhancements as described in section 7.2.

Some other technologies to consider when deploying Adaptive Bit-Rate (ABR) IP video is multicast-assisted ABR (M-ABR) and Smart ABR technologies. Using multicast in an IP video world may still be very important when considering special events with large viewer populations tuned in. M-ABR also significantly reduces the BW requirements between the CDN and the CCAP. For a CCAP with 25,000 subscribers, it could potentially reduce IP video capacity originating at the CDN from several hundred Gbps down to 10's of Gbps.

9.1.3. CM & In-Home changes

During the mid-term phase, operators should get the remaining high tier subs completely on 3.1 modems. Not only does this make it easier to provide them with good QoE during times of congestion, it will free up precious 3.0 bandwidth that will be needed by the lower service tiers.

And with higher penetrations of 3.1 modems, the operator can make better use of OFDMA spectrum. OFDMA has the potential of increasing the 42 MHz upstream capacity from 100 Mbps with 3.0 to the 200-250 Mbps range depending on plant quality.

In addition to the migration to 3.1 modems, operators should also be deploying Wi-Fi 6 throughout the rest of its customer base.

9.1.4. RF Plant Upgrades

During the mid-term years, operators should be upgrading their older 550/750/870 MHz plants to 1218/85 or 1218/204 MHz plants. These technologies are available today and will carry the operator through the rest of this decade. [ULM_2019] takes a look at how 1218 MHz networks can achieve our Cable 10G DS goals. DOCSIS 3.1 specification was written in 2014 and is just now becoming widespread. DOCSIS 4.0 specification just became available and requires an even bigger plant investment with new taps. So don't expect 4.0 to become mainstream until the end of this decade. And if growth rates continue to slow, maybe the need for 4.0 will be pushed out further in time.

The long term strategy for every cable operator should be to push fiber deeper. Eventually we may see fiber passing every home, but that level of investment takes multiple decades. However, MSOs should continue the fiber deep march in the mid-term and continue to reduce serving area size and cascade lengths.

There are some significant benefits with a fiber deep approach (e.g. N+0, N+1). In addition to the reduced SG size and increased 1.2GHz spectrum, the total number of active components can be reduced increasing reliability AND cutting OPEX costs in half. Fiber deep networks are also much closer to homes and become a jumping off point to selectively offer FTTH services.

Finally, some operators may be on the road to DAA. The mid-term period will be a good time to start that transition and start ramping up their deployments of RPD or RMD in the field.

9.2. Long-term Solutions – for 5-10 Year Horizon

What does the distant future hold? Nobody knows for sure and the crystal ball is cloudier than ever. Bandwidth usage may not snap back to its original pre-Coronavirus levels when Coronavirus ends. Why? Because this novel social experiment we are all involved in may foster a new social paradigm. Workers and companies may decide to explore more work-at-home activities and students and schools may decide to explore more on-line education activities.

Bandwidth will obviously continue to grow into the 2020's, although the rate of growth has a lot of debate. We are starting to see new applications like eSports getting a boost right now. The problem experienced during Coronavirus is only a sampling of what will happen in the future when that bandwidth growth crosses certain thresholds

> ***MSOs & Vendors need to begin working now to upgrade the HFC Plant to support the future Bandwidth Growth.***

While DOCSIS 3.1 can carry us a very long way, the industry needs to work on new technologies today to be ready for the future. This is what the Cable 10G initiative and DOCSIS 4.0 specification are all about. DOCSIS 4.0 will give us a choice of two options. Full-Duplex (FDX) technologies for N+0 plants that support dynamically shared spectrum from 108-684 MHz. This will enable 4-5 Gbps US tiers.

Alternatively, 4.0 also specifies some ultra-high split options paired with extended spectrum (ESD). These can offer service tiers similar to FDX on a 1794/492 MHz plant using simple FDD techniques. Both of these technologies are enabling multi-gigabit upstream tiers to be competitive with 10G PON symmetric offerings.

Inside the home, there will also be a migration from today's Wi-Fi 6 routers to the next generation of Wi-Fi 6E. The new generation will have a new spectral band from 5.9-7.1 GHz that enables new use cases and opens consideration of 10 Gbps around the home wirelessly that pairs nicely with 4.0 modems.

10. Conclusion

The Coronavirus BW surge gave us 12-18 months of upstream usage growth overnight, and about half that in the downstream. And despite the surge, we have seen that the DOCSIS network continues to work very well. But this event has shown weaknesses in the system that need filling to make it through the near-, mid- and long-terms.

Some operators needed some immediate quick fixes for severely congested service groups to relieve the pressure. These needed to be rolled out in a matter of hours or days. Adding DOCSIS channels with the I-CCAP was first choice if spectrum was available, with the roll-off being a good option for the 3.1 OFDM channels. Turning on OFDMA in the upstream should be a priority as well. There are a number of various configuration changes that could be done to improve operation including:

- Enabling TCP Ack suppression
- Reduce buffer-bloat latencies with Buffer-control TLVs, 3.1 AQM

- Optimizing QAM-modulation or FEC
- Re-locating in-home Wi-Fi

These might only provide 5% or 10% or 15% improvements, but everything helps in a congested network.

In the near-term over the coming months, operators can add more substantive improvements to their network capacity. To find spectrum for additional DOCSIS channels, operators can use SDV or improved encoding to reduce Legacy Video spectrum. Node segmentation might split SG to help manage BW. Nodes might go from 1x1 to 1x2 to 2x2/2x4 to 4x4 segmentation. This can be done simply by upgrading the node and not touching anything else in the outside plant. Meanwhile, more DOCSIS 3.1 modems and Wi-Fi 6 routers should be deployed to customers' homes.

Network capacity planning is key to being able to withstand another bandwidth surge. The CommScope/ARRIS Basic Traffic Engineering formula provides guidelines on how much capacity is required to provide good QoE. We saw some examples on what happens before and after the BW surge and why the QoE headroom is important.

> *Networks that have been designed with sufficient capacity have been resilient and has handled the Coronavirus BW surge.*

In the mid-term over the next several years, many plants will be upgraded to 1218/85 or 1218/204 MHz. Fiber Deep networks will continue to be our long term goal as we look to push fiber closer to the home while increasing plant reliability and reducing OPEX. Distributed Access Architectures including RPD and RMD will take hold, especially on longer fiber runs where they provide the most performance gains. And the IP video migration should be in full swing with technologies like M-ABR supplementing it.

Looking into the distant future, BW will not snap back after the pandemic. It will continue to grow through the decade. MSOs need to begin upgrading their networks for the future. While DOCSIS 3.1 continues to be the workhorse today, development of 4.0 products will help to enable symmetric multi-gbps services later in this decade.

Broadband kept society and the economy running, as best we can, in these hard times. This is arguably its most significant contribution to society in its short life-span. Hats off to everyone who has created and helped make the Internet and Broadband available to society.

While collectively we should all be proud of the infrastructure we have created, we still have much work to do. The Coronavirus BW surge has exposed cracks in the system. This should serve as a wake-up call. We still have improvements to make to the infrastructure. It's time to start upgrading the network for the demands of the 2020s.

Bibliography & References

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

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

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

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

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

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

[RAM_2020] Ram Ranganathan et. Al., "Decoding the Bandwidth Surge during COVID-19 Pandemic – an Indepth Study on DOCSIS Upstream Bandwidth Surge and their Impact on Video Conferencing", SCTE Cable-Tec 2020, SCTE

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

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

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

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

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

Abbreviations

ABR	Adaptive Bit Rate
AI	Artificial Intelligence
AQM	Active Queue Management
ATDMA	Advanced Time Division Multiplex Access
BAU	Business as Usual
Bps	Bits Per Second
BW	Bandwidth
CAA	Centralized Access Architecture
CAGR	Compounded Annual Growth Rate
CAPEX	Capital Expense
CCAP	Converged Cable Access Platform
CDN	Content Distribution Network
CM	Cable Modem
CMTS	Cable Modem Termination System
CPE	Consumer Premise Equipment
D3.1	Data Over Cable Service Interface Specification 3.1
DAA	Distributed Access Architecture
DEPI	Downstream External PHY Interface
DOCSIS	Data Over Cable Service Interface Specification
D3.1	DOCSIS revision 3.1
DS	Downstream
EPON	Ethernet Passive Optical Network (aka GE-PON)
EQAM	Edge Quadrature Amplitude Modulator
ESD	Extended Spectrum DOCSIS
FD	Fiber Deep
FDD	Frequency Division Duplex
FDX	Full Duplex (i.e. DOCSIS)
FEC	Forward Error Correction
FTTH	Fiber to the Home
FTTLA	Fiber to the Last Active
FTTP	Fiber to the Premise
FTTT	Fiber to the Tap
FTTx	Fiber to the 'x' where 'x' can be any of the above
Gbps	Gigabits Per Second
GHz	Gigahertz
HFC	Hybrid Fiber-Coax
HD	High Definition
HP	Homes Passed
HSD	High Speed Data
I-CCAP	Integrated Converged Cable Access Platform
ICO	Intelligent Channel Optimizer
IEEE	Institute of Electrical and Electronics Engineers
IPTV	Internet Protocol Television
ISCA	Integrated Service Class Agility

M-ABR	Multicast-assisted Adaptive Bit Rate
MAC	Media Access Control interface
MACPHY	DCA instantiation that places both MAC & PHY in the Node
Mbps	Mega Bits Per Second
MDU	Multiple Dwelling Unit
MHz	Megahertz
MSO	Multiple System Operator
N+0	Node+0 actives
NCTA	The Internet & Television Association
NFV	Network Function Virtualization
NSI	Network Side Interface
Nsub	Number of subscribers
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiplexing Access (Upstream)
OPEX	Operating Expense
OTT	Over the Top
PER	Packet Error Rate
PHY	Physical interface
PMA	Profile Management Application
PNM	Proactive Network Maintenance
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
QoS	Quality of Service
RF	Radio frequency
R-OLT	Remote OLT
RPD	Remote PHY Device
R-PHY	Remote PHY
R-MACPHY	Remote MAC-PHY
RMD	Remote MAC-PHY Device
RX	Receive
SCTE	Society of Cable Telecommunications Engineers
SDV	Switched Digital Video
SF	Service Flow
SG	Service Group
SLA	Service Level Agreement
SNR	Signal to Noise Ratio
STB	Set-top box
SW	Software
Tavg	Average bandwidth per subscriber
TCP	Transmission Control Protocol
Tmax	Maximum Sustained Traffic Rate – DOCSIS Service Flow parameter
TX	Transmit
US	Upstream
VOD	Video on demand
VoIP	Voice over Internet Protocol
VPN	Virtual Private Networks