Traffic Engineering in a Fiber Deep Gigabit World

A Technical Paper prepared for SCTE/ISBE by

John Ulm
Engineering Fellow, Broadband Systems
CTO - Network Solutions team
ARRIS
john.ulm@arris.com

Tom Cloonan
Chief Technical Officer, Network Solutions
ARRIS
tom.cloonan@arris.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Introduction

Many new innovations are finding their way into cable operators' plants and many others are on the way. With DOCSIS 3.1 deployments well underway, a new era of Gigabit services in the downstream are being introduced. With DOCSIS Full Duplex (FDX) on the horizon, operators have the promise of symmetric Gigabit services. At the same time, some operators are starting their migration to fiber deep networks and others are looking at Distributed Access Architectures (DAA) such as Remote PHY.

All of this will have significant impact in the way operators manage their traffic engineering and network capacity planning. As an industry, we need to re-evaluate and update our models. This starts with an intimate understanding of subscriber bandwidth behavior. This paper takes a detailed look at a year's worth of live consumer data collected from a single cable site. When sampling during peak busy hours, every packet was tracked to allow for traffic analysis down to the second.

The paper will investigate many different bandwidth trends uncovered. Some of the key variables of interest include traffic consumption based on:

- Differing Service group sizes
- SG to SG variation
- Subscriber service tiers
- Time of day
- Day of week
- Month to month

With these trends in hand, the impacts on existing network capacity models are discussed and how they might morph to provide traffic engineering in a Fiber Deep Gigabit world.

Broadband Bandwidth Trends

The Internet has been growing at a breakneck speed since its inception. And with it, we have seen a corresponding growth in dedicated network capacity. [ULM_2016] provided an overview of these trends which are highlighted and updated below.

1. Nielsen's Law and Cloonan's Curves

While Moore's Law is infamous in silicon realms, Nielsen's Law of Internet Bandwidth has become renowned in the broadband world. It basically states that network connection speeds for high-end home users would increase 50% per year. This law has driven much of the traffic engineering and network capacity planning in the service provider world. It has also led to much research on those topics.

In [CLOONAN_2014, EMM_2014], this research was expanded to also include traffic utilization in addition to the network connection speed. Nielsen's Law is shown in the figure below. Since the Y-axis is a log scale, the 50% Compounded Annual Growth Rate (CAGR) appears as a straight line. An interesting fact is that the graph starts in 1982 with a 300-baud phone modem. The industry is now in the fourth decade of closely following this trend.

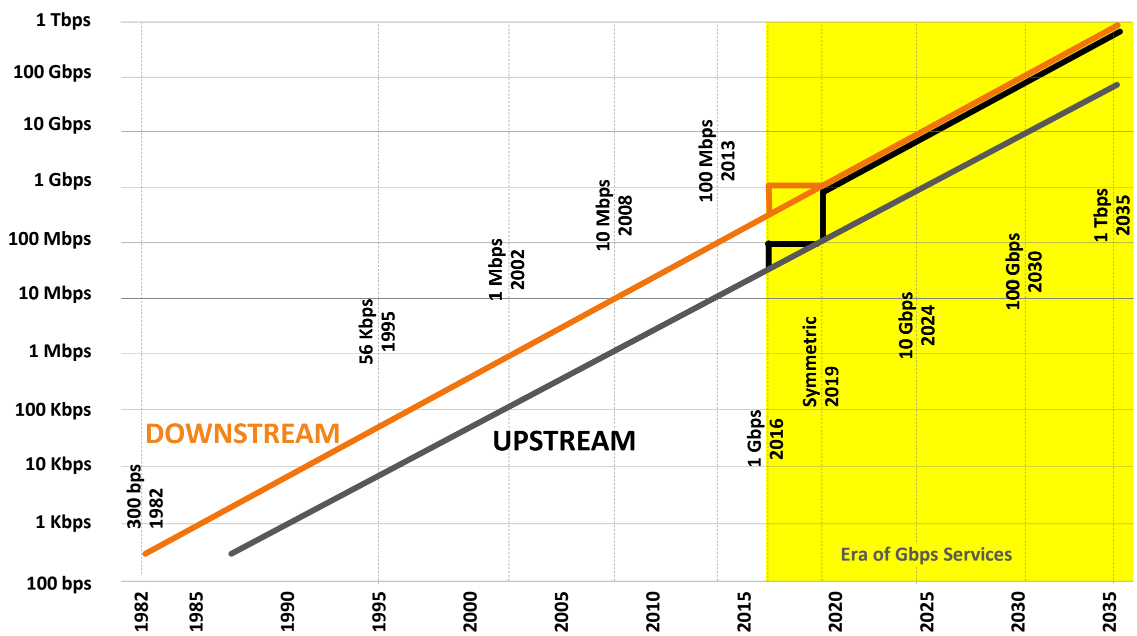

Figure 1 - Nielsen's Law – 50% CAGR

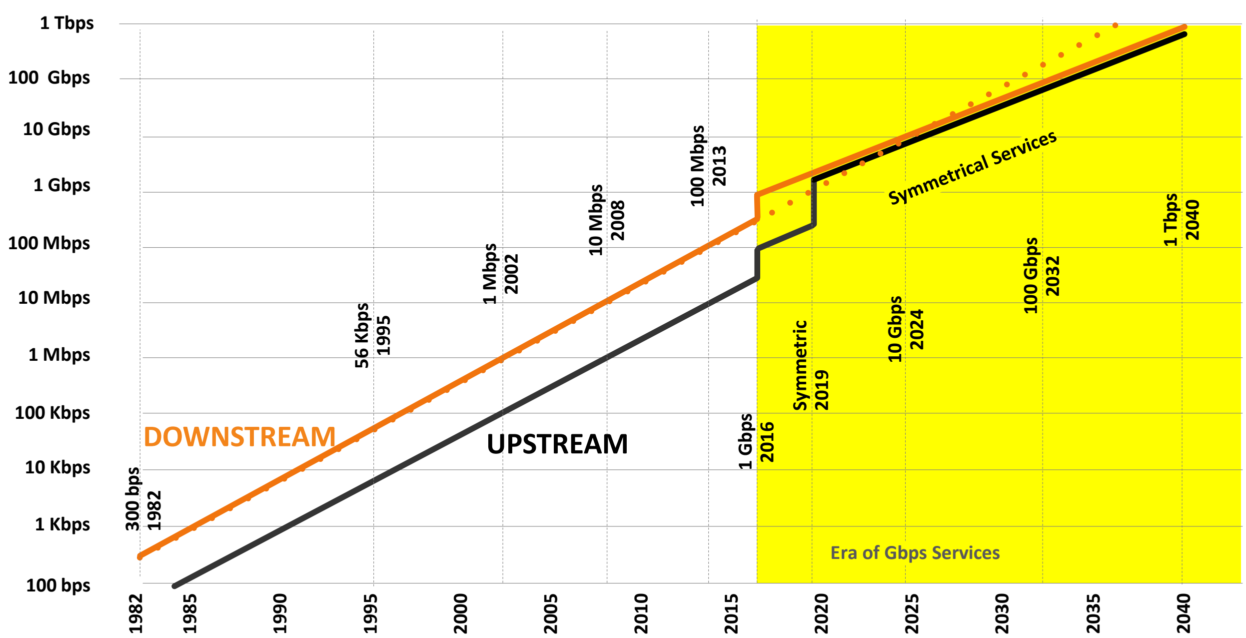

Figure 2 - Modified Nielsen's Law – 33% CAGR after 2018

While this trend shows a straight line increase, in reality, internet speeds will make a jump and stay there for a bit. There was recently a jump to 1 Gbps services that happened a couple years ahead of the Nielsen's Law projection. This is shown in the Figure 1. Service Tiers may stay here a couple years before the speed tier climb continues.

While Nielsen's Law focuses primarily on downstream speeds, it has been noted that upstream speeds have generally followed the same growth rate, but at about one-tenth the speed. However, with more Fiber to the Premise (FTTP) deployments and the upcoming introduction of DOCSIS Full Duplex (FDX), it is expected that the highest offered upstream speeds will take a step up as symmetric services become available.

Will Nielsen's Law continue its 50% growth unabated for the next couple decades? In recent years, there has been some suggestions on whether Moore's Law may be slowing down. Will this have a corresponding impact on Nielsen's Law? One could argue that Moore's Law is the fuel behind ever advancing Consumer Premise Equipment (CPE), and these CPE drive the need for Internet bandwidth.

Figure 2 takes a look at a modified Nielsen's Law where the CAGR is reduced to 33% going forward. This stretches the time for 10X growth from the original 5½ years up to 8 years. You can see that from a network capacity planning perspective, the overall impacts are similar. The changes over the next decade are minimal. Longer term, the time it will take to reach the 1 Tbps milestone gets pushed out about five years, from 2035 to 2040. So, even with a slowing in Nielsen's Law, there will be minimal impact in operators' long term network capacity planning.

Earlier work by Cloonan noted that the primetime average subscriber consumption (a.k.a. Tavg) has also been following this same basic trend as shown in the Figure 1. For service providers, an important metric is the traffic utilization in a Service Group (SG). The SG traffic utilization is a function of the number of subscribers (Nsub) times the average bandwidth per sub (Tavg). In [CLOONAN_2014, EMM_2014], this research was expanded to also include traffic utilization in addition to the network connection speed. This was shown in a chart known as Cloonan's Curve, where SG consumption is shown in addition to Nielsen's Law.

In the early DOCSIS days, many nodes were combined together and a SG might consist of thousands of subscribers. At that time, the SG traffic was an order of magnitude higher than the maximum network connection speed (a.k.a. Tmax after the DOCSIS parameter that dictates max network rates). Over time, the SG size has been shrinking and, with it, the ratio between Nsub*Tavg to Tmax. The SG traffic will eventually approach that of Tmax. As SG sizes dip below 100 subs, then Tmax starts to dominate the traffic engineering.

2. Broadband Subscriber Traffic Consumption

ARRIS has been monitoring subscriber usage for many years now. The chart below shows Tavg, the average subscriber downstream consumption during peak busy hours, for a number of MSOs over an eight year period. At the start of 2017, Tavg finally broke the 1 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 have an average downstream traffic growth that has been just under 40%. On a yearly basis, traffic growth can be very sporadic. It is not uncommon to see high growth in one year followed by little growth the next. So, the 40% trend should be used as a longer term guideline on downstream traffic consumption. This equates to roughly doubling every other year.

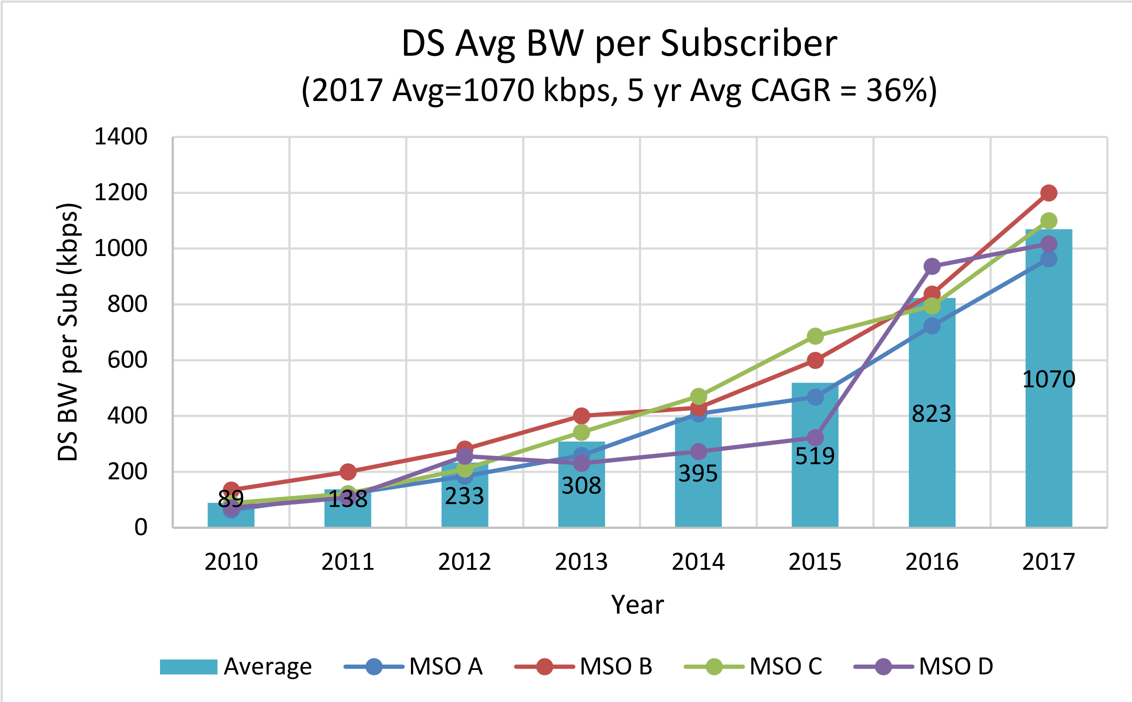

Figure 3 - Tavg, Average Subscriber Downstream Consumption

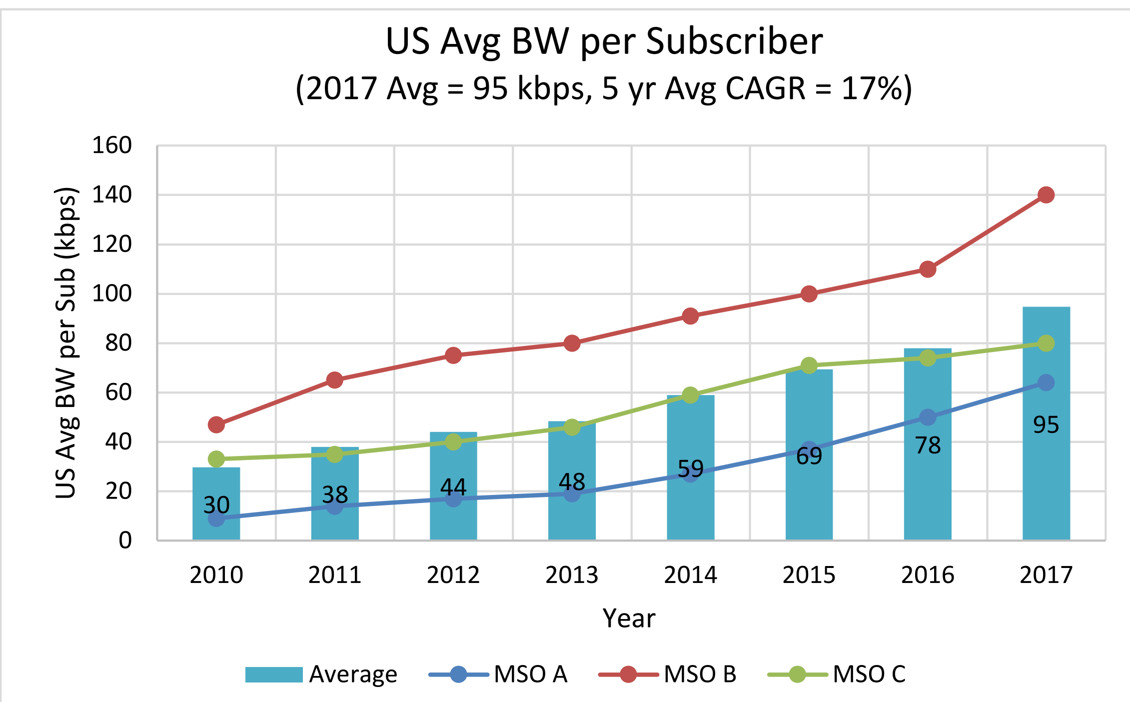

Figure 4 - Tavg, Average Subscriber Uptream Consumption

Interestingly, the upstream traffic is growing at a significantly slower rate. During the same eight year period, the upstream Tavg only grew at ~20% CAGR. Traffic is also becoming more asymmetric with video applications driving downstream consumption [EMMEN_2014]. While the average DS:US ratio is ~10:1, the MSO with the largest DS:US ratio seemed to have stabilized around a 15:1 ratio. It will be interesting to see what happens with other operators as they reach this point.

3. Selective Subscriber Migration Strategy

At first glance, Nielsen's Law is a scary proposition in that HFC networks might be obsolete in 5-7 years while it may take decades to build out an FTTP infrastructure. However, this is not the full story. As was shown in [ULM_2016, ULM_2014], Nielsen's Law applies to the top speed tiers which is only a very small percentage of the entire subscriber base, perhaps less than 1%. So, the key question then becomes, "What happens to the vast majority of subscribers on HFC who are not in the top speed tiers (a.k.a. billboard tiers) and when?"

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

The figure below maps out an example of the various service tier growth over the next two decades. While the 1% of subs in the top billboard tier hit 10 Gbps in ~2024, the 14% of subs in the performance tier don't hit that mark until ~2032. Notice that 85% of subscribers in the flagship basic tier and economy tier stay below this mark for several decades. It is important to note that 99% of the subscribers will still be comfortably using today's DOCSIS technology on HFC a decade from now. With a Selective Subscriber Migration strategy, it becomes very important from a traffic engineering perspective to understand the behavior of the individual service tiers. But with this understanding in hand, Selective Subscriber Migration can be used to extend the life of HFC for decades to come.

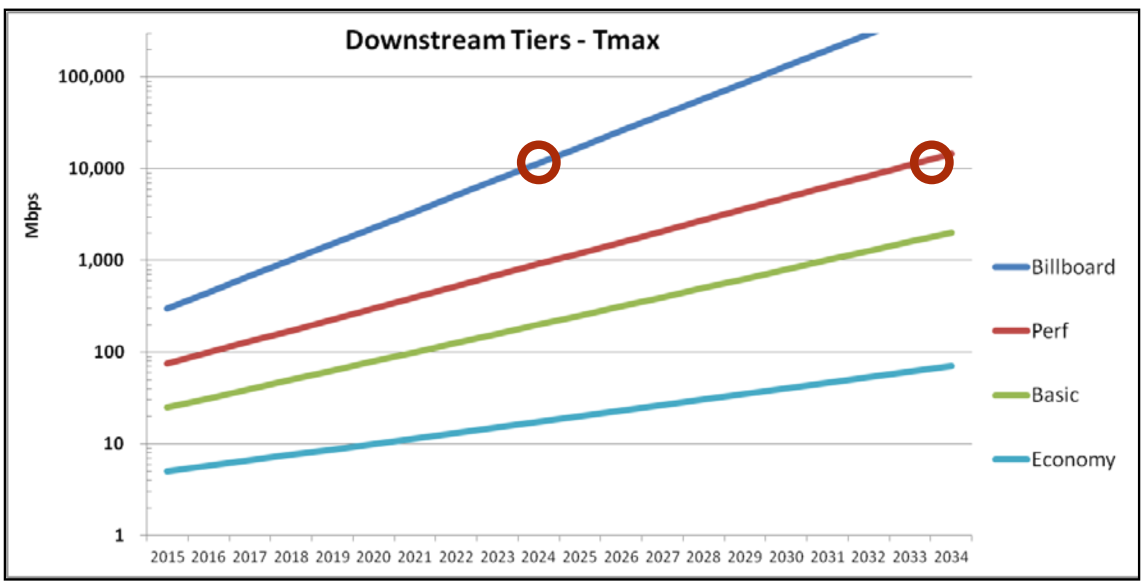

Figure 5 - Downstream Growth with Multiple Service Tiers

Review of Broadband Traffic Engineering

Previously, [CLOONAN_2014] provided an introduction on 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.

1. The "Simple" Traffic Engineering Formula

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

THE "2014" TRAFFIC ENGINEERING FORMULA (BASED ON Tmax_max):

C >= (Nsub*Tavg) + (K*Tmax_max), (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 Tmax offered by the MSO

There are obviously fluctuations that will occur (i.e. the "AC component" of traffic) which can force the instantaneous traffic levels to both fall below and rise above the DC traffic level. The second term (K*Tmax_max) is added to increase the probability that all subscribers, including those with the highest 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], a similar formula assumed that a K value of ~1.0 would yield acceptable and adequate QoE results. [CLOONAN_2014] goes on to provide simulation results that showed a value between K=1.0 and 1.2 would provide good QoE results for a service group of 250 subscribers. Larger SGs would need larger values of K while very small SGs might use a K value less than 1.0.

Using the simple Traffic Eng formula (1), it becomes possible to develop sophisticated network capacity models. Some results from the ARRIS Network Capacity model are shown in Figure 6. It provides an insight into both Tmax and SG Tavg behavior. During the next 5-7 years, the Tmax component dominates traffic engineering as it is driven by Nielsen's Law. The bandwidth needed by the top billboard tier dominates compared to the SG Tavg. But as top tiers are moved off the HFC, then eventually the Tavg component catches up again.

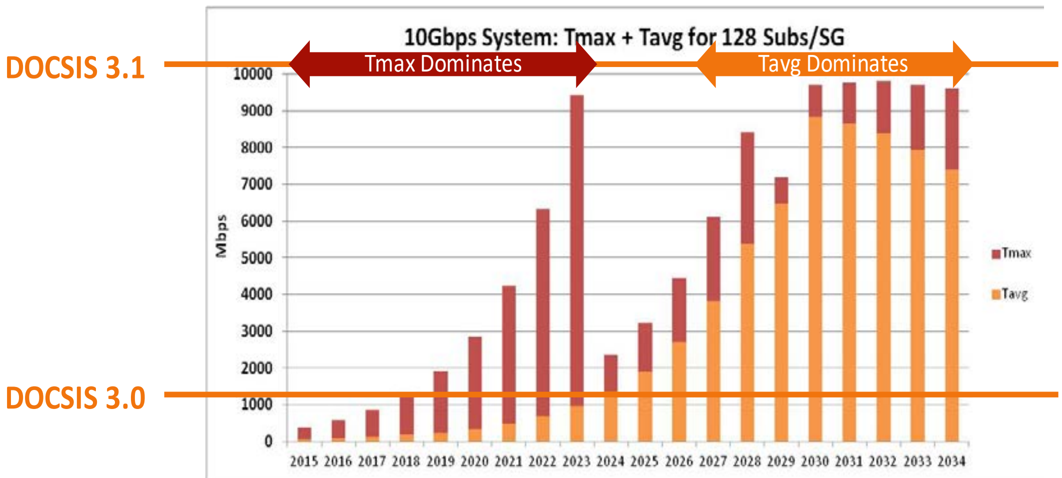

Figure 6 - Network Capacity Model Results

2. Limitations of the "Simple" Traffic Eng Formula

The "Simple" formula has been extremely effective, but it is very important to understand its limitations. The formula was developed for service groups on the order of a couple hundred subscribers, and found that a K value between 1.0 and 1.2 provided good QoE.

However, as the cable world migrates to Fiber Deep HFC designs jointly with Distributed Access Architectures, operators will need to perform traffic engineering on multiple different sized groups. On one extreme, the DOCSIS SG might shrink to less than 100 subscribers, perhaps only a couple dozen. On the other extreme, operators need to engineer the network links in and out of the Routers and CCAP Core with tens of thousands of subscribers. In between may be a multi-tiered Ethernet switching infrastructure where 100G Ethernet links cascade down to 40G links down to 10G links. Every link needs to be managed to make sure it is not a bottleneck to providing acceptable QoE.

The simulations in [CLOONAN_2014] show that the optimum value for K does vary with several different parameters. In reality, finding the optimum value of K becomes very complex and dependent on many variables.

The "DC" component of the formula (i.e. Nsub * Tavg) also has limitations. It appears to be fine for very large sizes but becomes less accurate for smaller SGs where there is much wider SG to SG variation. Going forward, the "Simple" formula will need to evolve to work across these wider ranges of variables.

Subscriber Bandwidth Behaviors

To enhance our Traffic Engineering formula, an intimate understanding of subscriber bandwidth behavior is needed. This paper takes a detailed look at a year's worth of live consumer data collected from a single cable site. A massive amount of data was collected during many peak busy hours. Every packet was tracked to allow for traffic analysis down to the second.

The paper will investigate many different bandwidth trends uncovered. Some of the key variables of interest include traffic consumption based on:

- Month to month
- Day of week
- Time of day
- Differing Service group sizes
- SG to SG variation
- Subscriber service tiers

1. Data Collection Methodology

Data was collected from a live DOCSIS system over the course of a year. Packet monitoring equipment allowed every packet in the system to be captured during each 30-minute sample interval. Typically, multiple measurements were taken during peak busy hours between 6 pm and midnight on a given night. This created massive amounts of raw data that filled disk drives. To make the analysis manageable, the data needed to be parsed into a more usable metric.

Previously in the simple formula, Tavg would typically be calculated across a timespan of many minutes or hours. However, the QoE must be measured on a much finer granularity for the applications people use such as web browsing, OTT video consumption or even running a speed test. All of these events are sensitive to latency on the order of a couple of seconds. We chose to analyze the data at 1 second intervals. This appeared to be the best compromise between observing QoE behavior yet minimizing file sizes to a manageable size.

2. Macro Trends

For a subset of the data, the major trends were reviewed as they varied by month, day and time. This particular data set was collected during the month of June 2016, and then from mid-September through early February 2017. Data was collected across every day of the week and from 6 pm to midnight.

2.1. Month to Month

As discussed earlier, Tavg per subscriber has been rising steadily over the years. The month to month variation for this data, shown in Figure 7 below, confirms this.

As can be seen, the overall trend is higher over time. However, it is not necessarily a smooth linear increase. September saw a big jump and then Tavg dropped a bit for October and November. This was followed by another big jump in December before Tavg slid a bit in January and February.

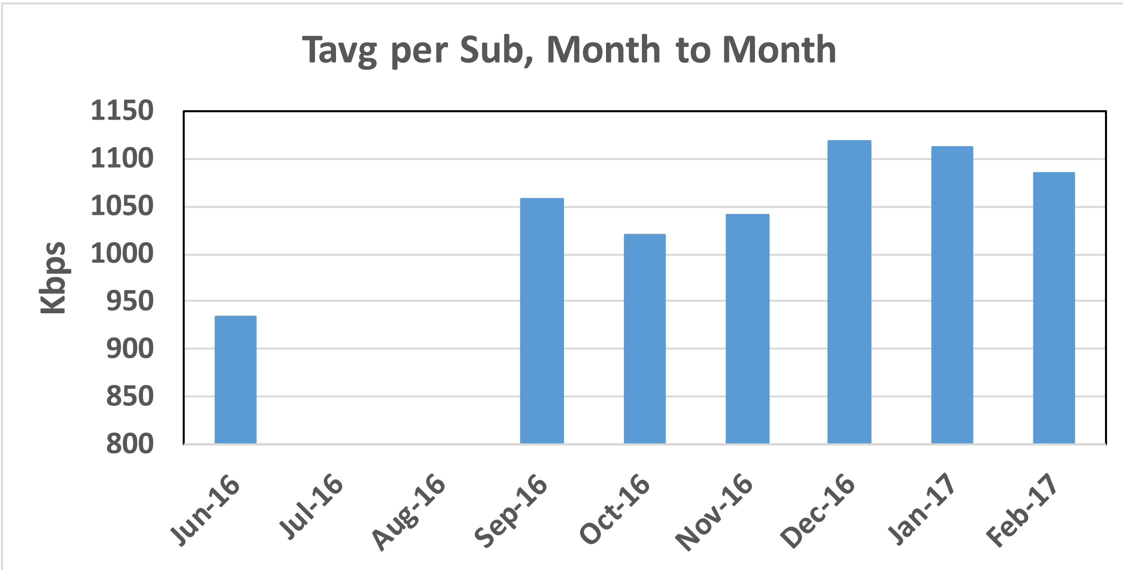

Figure 7 - Tavg per Sub, Month to Month variations

Tavg was slightly higher than 1100 Kbps in Jan 2017, which puts it a hair higher but in line with other data shown in Figure 3. The Tavg growth from June 2016 to Jan 2017 is just under 20% in 7 months. This equates to ~35% annual CAGR which is also right in line with the Figure 3 data.

2.2. Day of Week

What is the busiest evening for broadband usage? What is the least busy day? It turns out that on average for this data set, Sunday ties Thursday for the honor of busiest evening. This is shown in Figure 8. For least busy day, Saturday barely nudged out Tuesday.

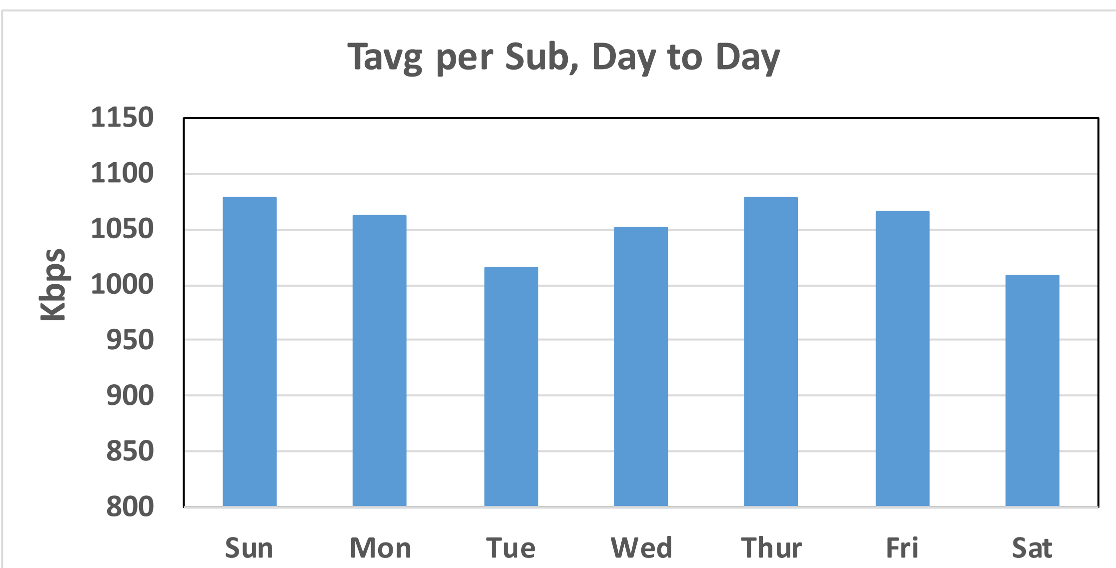

Figure 8 - Tavg per Sub, Day to Day variations

It should be noted that there is only a couple percent swing above and below the overall average, so the relative swings from day to day are not major.

2.3. Time of Day

For traffic engineering, there is often a reference to "peak busy hour". So exactly when is that? Figure 9 shows Tavg based on Time of Day across this very large data set. Each iteration lasted 30 minutes. The 6 pm bar in the figure represents the average of all iterations that were started between 6 pm and 6:30 pm.

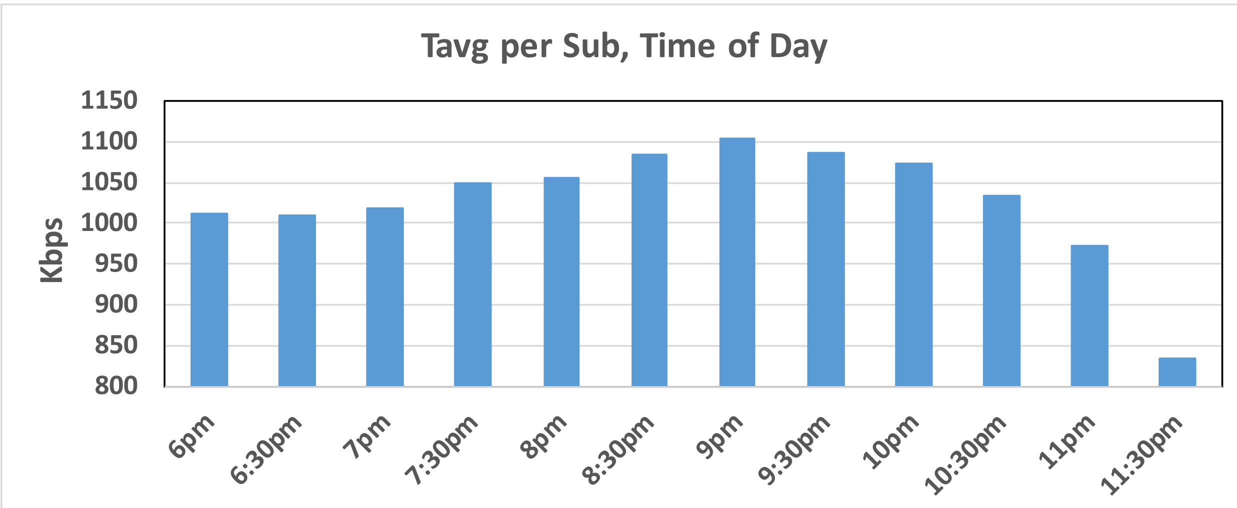

Figure 9 - Tavg per Sub, Hour to Hour variations

On average, the busiest time of day is between 9 pm and 10 pm. Usage is fairly constant during the dinner hours, and gradually increases up to the peak busy hour. The overall increase is ~10%. Usage then drops off as it gets past 11 pm.

3. Micro Trends

When looking across a large data set, the averages shown above provide an interesting data point, but do not have the resolution needed to understand the impacts of data bursts on subscriber QoE.

3.1. Micro view for Time of Day

To get a better understanding of fluctuations on a single day, an example from June is shown in Figure 10. As can be seen, there is ~33% swing from 8 pm to 10 pm which is much higher than the ~5% swing seen when averaged across several months of data. Its peak is about 15% above the monthly average, so this gives a sense for how a daily peak can be higher than the peak when averaged over a month.

The data in Figure 10 shows the variation in Tavg on 30 minute intervals. As was mentioned earlier, applications are concerned with latency impacts on the order of seconds, so even finer granularity is needed.

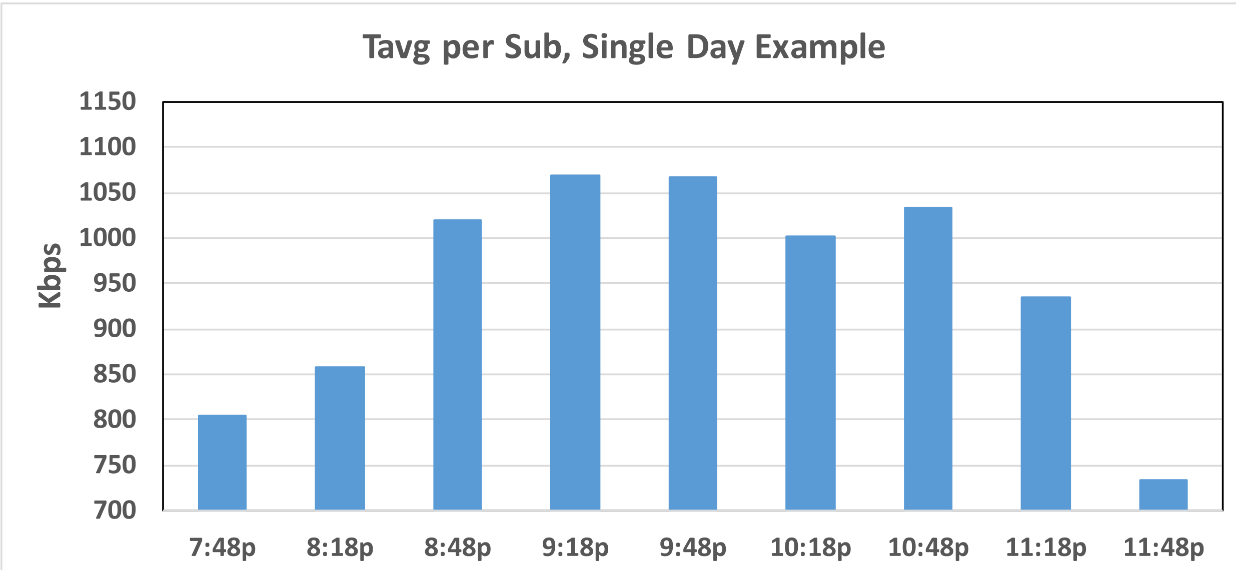

Figure 10 - Tavg per Sub, Single Day Example

As can be seen in Figure 10, there is a 2½ hour peak busy window from the start of the 8:48 pm interval to the end of the 10:48 pm interval where the peak usage is reasonably consistent, but falls off on either side. A deeper analysis in this window was done at 1 second intervals for more than 1000 subscribers. A histogram of the 1 second Tavg intervals using 10 Mbps bins is shown in Figure 11:

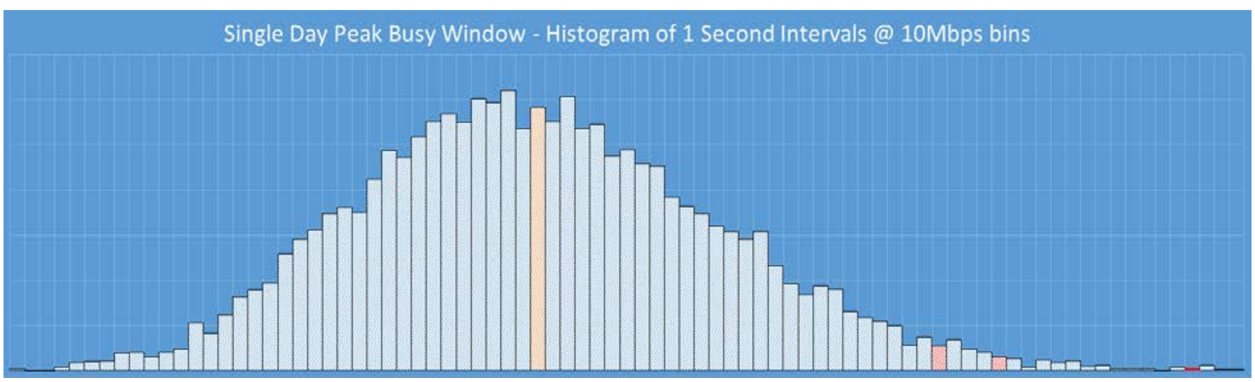

Figure 11 - Single Day Peak Busy Window – BW Histogram, 1 Second Intervals

The average per subscriber bandwidth during the 2½ hour peak busy window is just over 1000 Kbps. This is the highlighted bar near the middle of the chart. The maximum 1-second interval had a Tavg per sub that was just over 1500 Kbps. That's roughly 50% higher than the peak busy window Tavg. The minimum bandwidth seen in a 1 second interval was just under 700 Kbps. Therefore, the ratio of max to min is just over 2:1.

While knowing the max 1 second interval is useful information for understanding burst requirements, that 1 second interval still only represents 0.01% of the peak busy window. A more accurate characterization of the bandwidth distribution is needed for our QoE analysis. Some additional results from this data set are shown in Figure 12.

The standard deviation was calculated to be ~120 Kbps. Tavg + 2 standard deviations came in just under 1300 Kbps, while Tavg + 3 standard deviations were about 1400 Kbps.

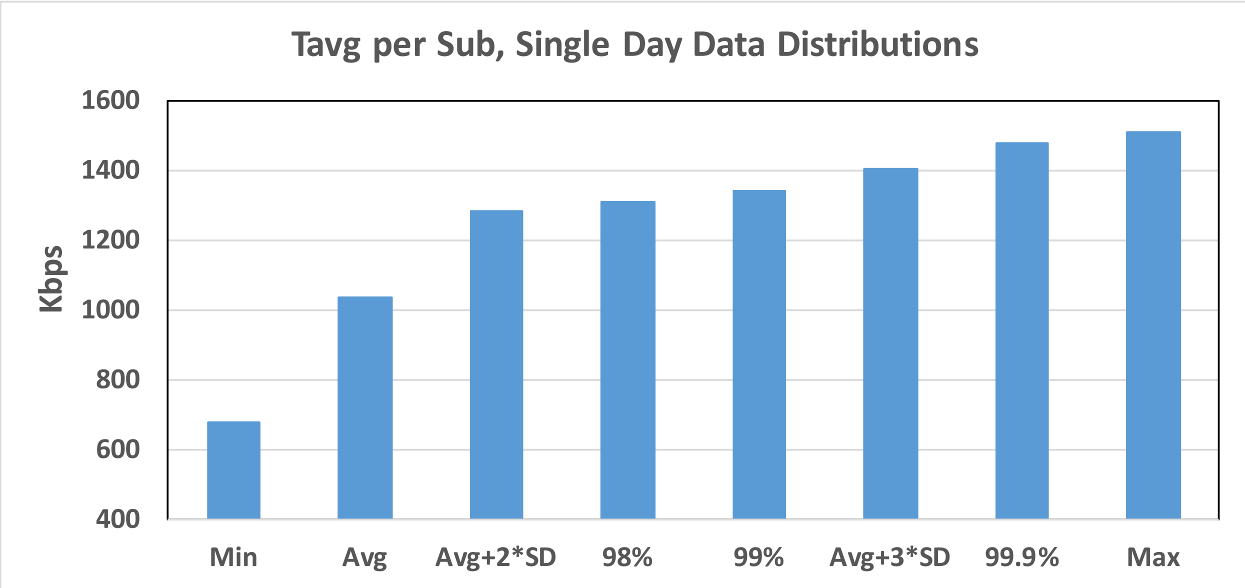

Figure 12 - Tavg per Sub, Single Day Data Distributions

In addition to standard deviation, an analysis was done using percentiles. The chart shows the percentile levels for 98%, 99%, and 99.9%. The percentile level shows the % of 1 second intervals that are at or below this level. The 98% percentile level was ~1300 Kbps. This means that only 2% of the number of 1 second bandwidth intervals exceeded this level. The 99% percentile was just marginally higher than the 98% level. Finally, the 99.9% level was very close to the ~1500 Kbps maximum. The data points associated with each percentile level are also shown in Figure 11 with the small highlighted bars on the right side of the chart.

4. Service Group Considerations

As discussed previously, the "Simple" traffic engineering equation has considerations as the Service Group sizes become either very small (e.g. Fiber Deep network) or very large (e.g. CCAP Core for Remote PHY). Our research looked at some of the impacts on traffic usage due to varying SG sizes as well as SG to SG variations given the same size and service tier distribution.

4.1. Service Group Sizes

Intuitively, as SG sizes become larger, there is expected to be less relative variation thanks to the benefit of large numbers of samples. Our research tried to understand the extent of this and quantify it. To illustrate this, Figure 13 and Figure 14show the same data set taken from a single day in Feb 2016 but organized as different sized SG. For Figure 13, the data was organized as one SG with ~1100 subs. For Figure 14, the data was organized as 11 SGs with ~100 subs each. The data provided bandwidth resolution in 1 second intervals. The X-axis in both figures varies from zero to the Max 1-second interval.

Looking first at Figure 13, one can see that it is a much tighter distribution. The maximum 1-second interval is only about 40% higher than the mean value. The minimum 1 second interval is about half of the maximum interval. The coefficient of variation (i.e. standard deviation divided by the mean) is less than 10%. In looking at many other data sets with ~1000 subs, the coefficient of variation ranged from 4% to 10%.

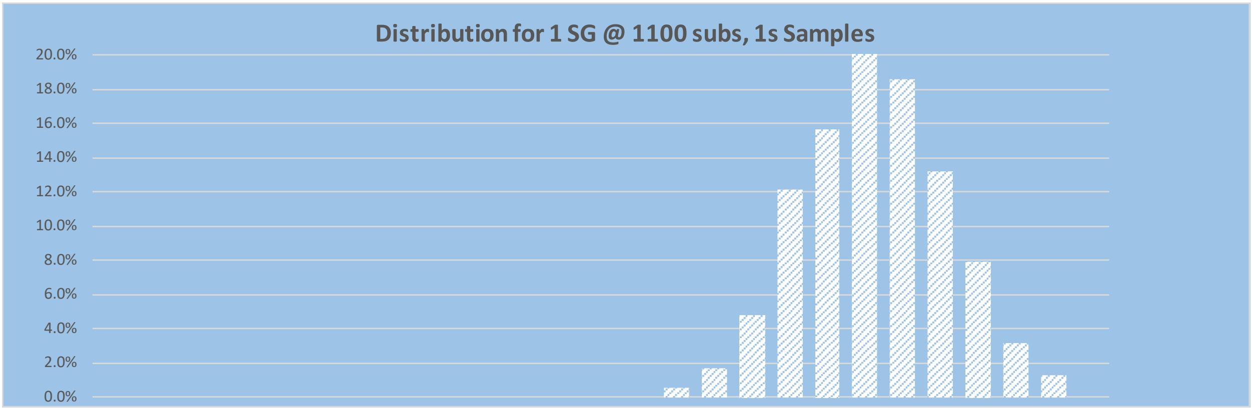

Figure 13 - Bandwidth Distribution for SG with 1100 subs, 1 sec intervals

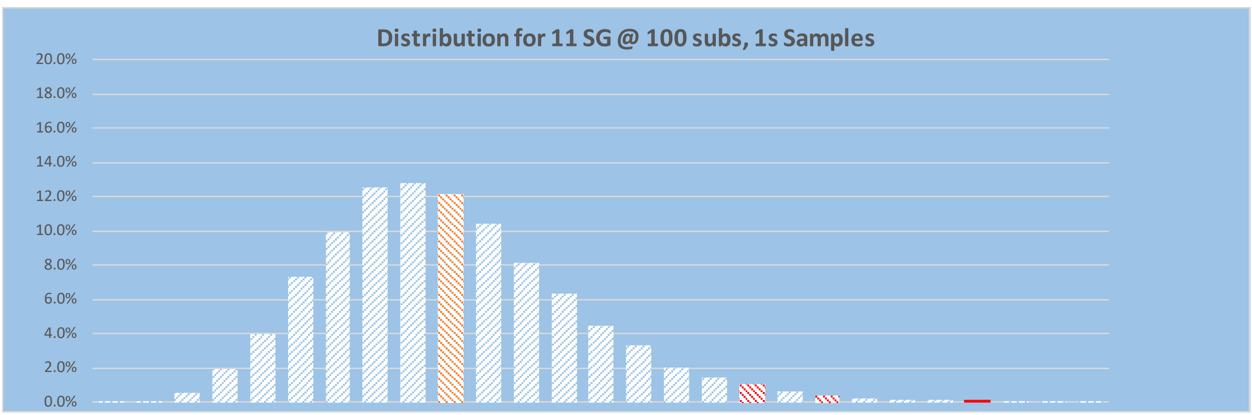

Figure 14 - Bandwidth Distribution for 11 SGs @ 100 subs, 1 sec intervals

Now looking at the 100 subs per SG data in Figure 14, it is apparent that it is a wider distribution. The maximum 1-second interval is almost triple the value of the mean. The coefficient of variation is much higher, around 35%. In some other data sets, it went over 50%. Figure 14 also indicates the average and 98%, 99% and 99.9% percentile values with shaded bars.

4.2. SG to SG Variations

Even for a given SG size, the traffic engineering must account for variations from SG to SG. Figure 14 shows the aggregated data for 11 unique SGs. But what is happening in each of these SG? Figure 15 helps give us an insight. This is using the same Feb 2016 data set as above.

The heavy blue line in Figure 15 shows the aggregated data from all 11 SGs. The thin lines show the individual SG.

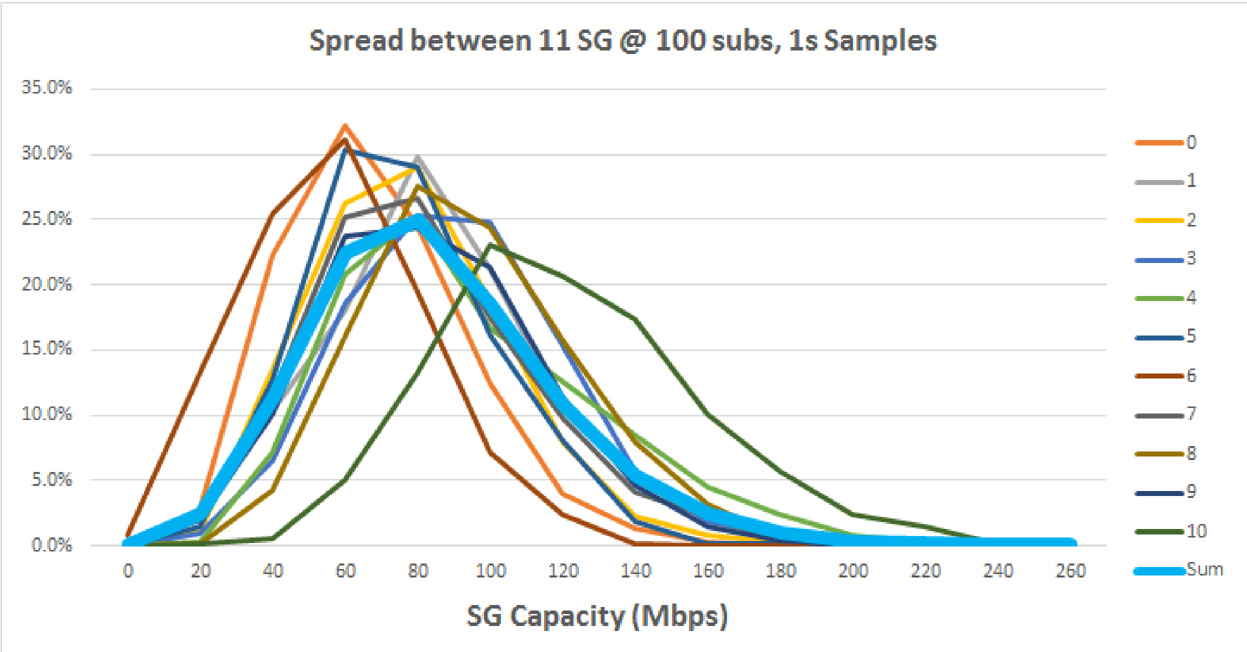

Figure 15 - Tavg per Sub, Single Day Data Distributions

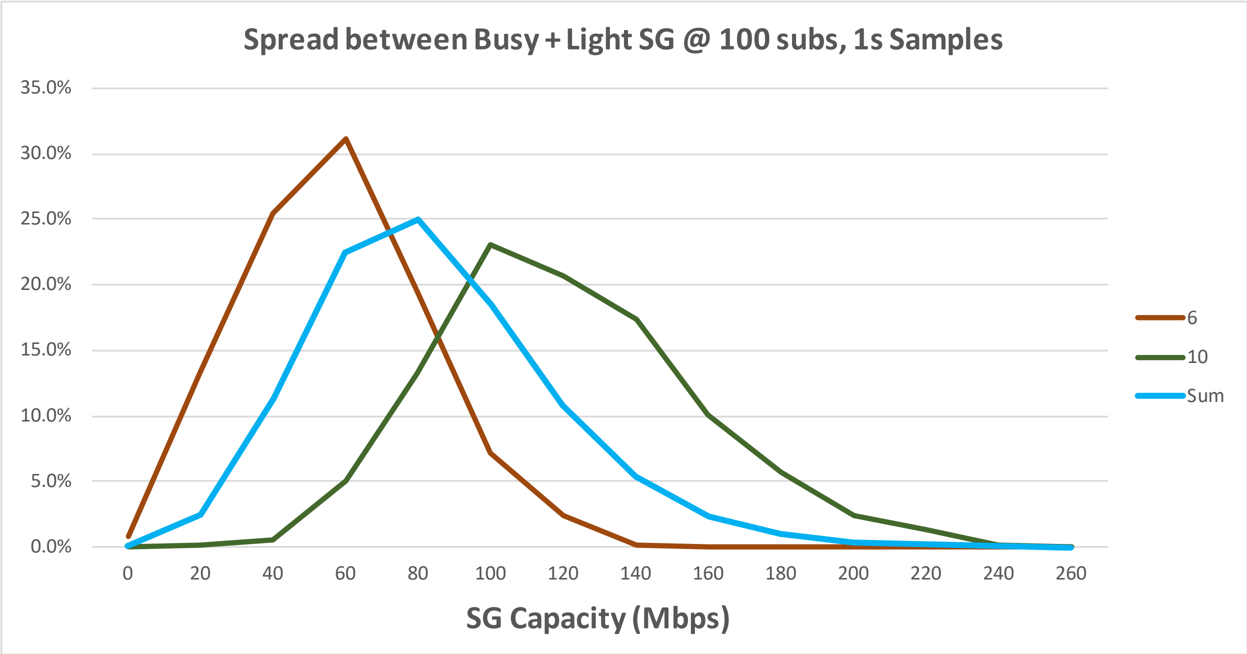

Figure 16 - Tavg per Sub, Single Day Data Distributions

Many of the SGs have very similar behavior. In Figure 16, the SGs with the two extremes are isolated: SG 6 with the lightest data traffic and SG 10 with the heaviest data traffic. Tavg for SG 10 is about twice that of SG 6; despite having virtually the same service tier distribution.

If the aggregated sum had been used for traffic engineering, the 98% percentile level was ~176 Mbps. While this would have been overkill for SG 6, it definitely does not address the needs of SG 10. SG 10 exceeds this value more than 10% of the time. Looking at the 98% level for the individual SG, then SG 6 would be ~124 Mbps while SG 10 would be ~220 Mbps.

5. Subscriber Variations – Heavy Users & Service Tiers

Knowing the average distribution is useful but not enough for understanding SG bandwidth behaviors. This is especially true for smaller SGs. As seen above, two 'similar' SG with 100 subs had bandwidth utilization that was different by a factor of two. The following sections explore some of the reasons for the SG to SG variations.

5.1. Active + Heavy Users

Within any given SG, there is a mix of active and idle subscribers, light and heavy users. As SG sizes approach 100 and shrink below that, then the types of user on any single SG can have a significant impact.

Even within a given service tier, the bandwidth usage can vary dramatically. Figure 17 looks at results for the most common service tier, 25M, for the Feb 2016 data set that was analyzed above. This is the mainstream tier and it contains about half of the total subscribers. Almost 62% of the subscribers were predominantly quiet during this 30-minute interval and consumed less than 250 Kbps. Figure 17 shows the Tavg bandwidth distribution for the remaining 200 active subs.

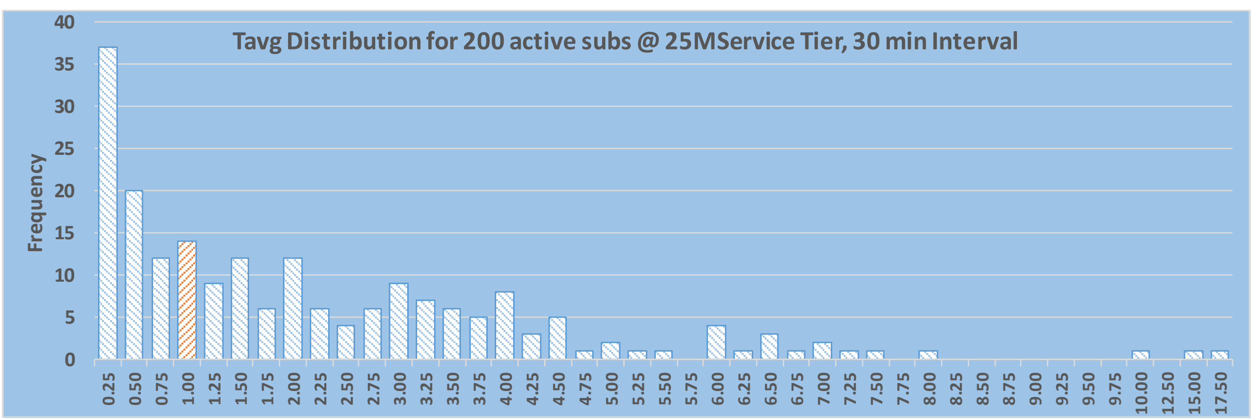

Figure 17 - Tavg BW distribution for 25M Service Tier

The bandwidth average across this entire group was about 1 Mbps. 75% of the total subs were below this level. The highest bandwidth consumer used almost 20 Mbps for an entire 30-minute interval. There were three subs that were over 10Mbps and 21 subs that averaged more than 5 Mbps over that interval. Obviously, if one SG gets a disproportionate number of heavy users (either too many or none at all), this can greatly influence the bandwidth utilization for a given SG.

In addition to analyzing 25M Service Tier, a look at the 100M tier provided a useful insight. SG 6 with the lightest utilization had three 100M subs but they were all relatively quiet. SG 10 had only two 100M subs but one was quiet and the other was extremely active. In fact, the active 100M subscriber averaged

~43 Mbps over the entire 30-minute interval! This was obviously a big factor in SG 10's high bandwidth utilization.

As can be seen for small SG sizes, the number of active and heavy users compared to the number of idle and light users can dramatically affect the SG bandwidth utilization. As SGs become very large, then the laws of statistics tend to even things out.

5.1. Service Tier Impacts

In addition to the active/idle ratio, another important factor in SG variation is the mix of Service Tiers among the various subscribers. Table 1 provides an example mix of Service Tiers with their respective bandwidth utilization.

Table 1 - Example Bandwidth Distribution by Service Tier

Service Tier	% of Subs	Tavg per Sub (Mbps)	Avg Burst Magnitude (Mbps)
6M	8%	0.49	6.9
12M	24%	0.67	7.9
25M	44%	1.01	11.8
50M	11%	1.68	17.6
100M	3%	2.66	26.4
Avg	100%	0.91	10.4

For this data set, the overall Tavg for subscribers was 910 Kbps. As can be seen, the Tavg when measured for each service tier can vary quite a bit. The Tavg for the lowest tier, 6M, came in at just under 0.5 Mbps, while the top 100M tier had Tavg = 2.66 Mbps.

To help manage QoE, our research also investigated the differences in traffic bursts between the service tiers. The bandwidth data was analyzed at 1 second intervals. Figure 18 shows a probability distribution function (pdf) for the bandwidth burst rates at any given second for each of the service tiers.

During relatively idle periods, all the service tiers behave reasonably similarly. This can be seen on the left-hand side of Figure 18. However, once the subscriber becomes very active, then their maximum burst capability is limited by their service tier Tmax value. As can be seen on the right-hand side of Figure 18, the higher service tiers have higher burst rates.

The magnitude of the average burst is also provided in the rightmost column in Table 1. The lowest 6M tier has a burst magnitude of 7 Mbps as its Tmax value (~8 Mbps) is slightly higher than the actual tier level. The 100M tier has a burst magnitude of 26 Mbps, even though its average utilization is only 2.7 Mbps. This gives an insight into the active to idle ratio.

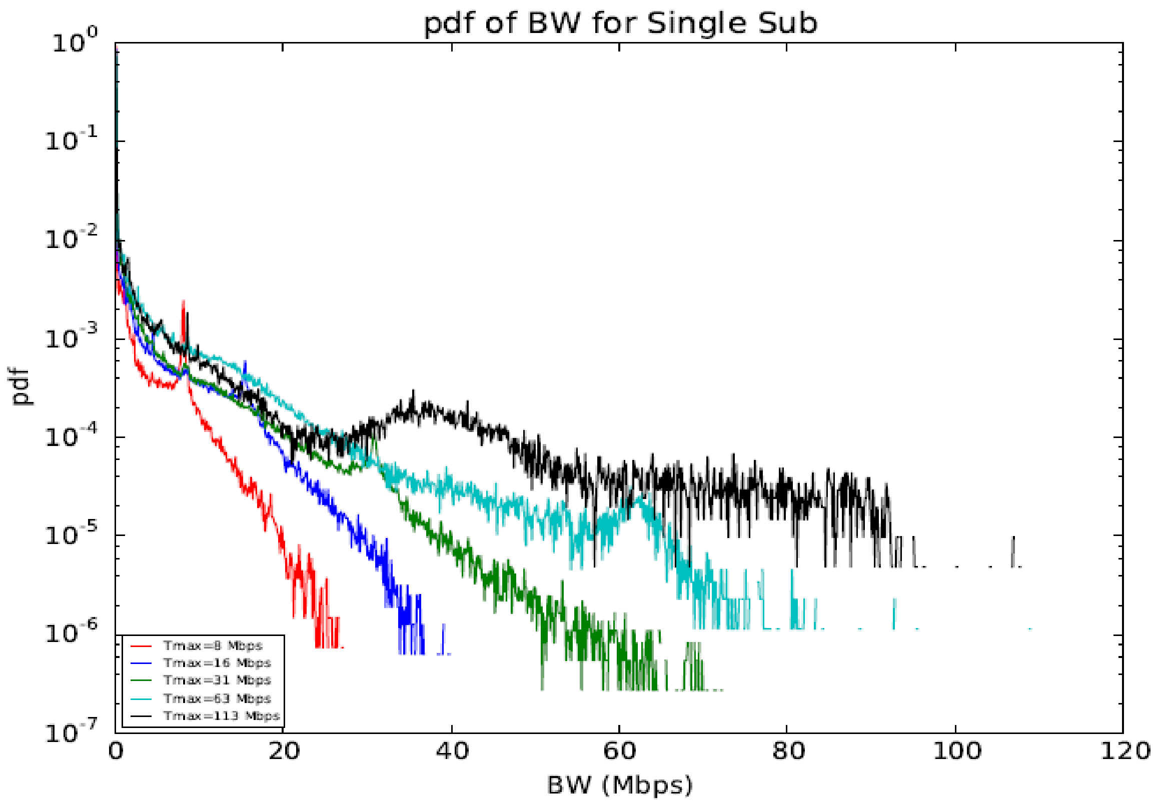

Figure 18 - Bandwidth Transmit Rate Probabilities for a single Subscriber

These differences in service tier bandwidth utilization may be further impacted if an operator has implemented data usage caps for some or all of their tiers.

Updating the Traffic Engineering Formula

The above trends can now be used to evaluate the impacts on existing network capacity models and see how they might morph to provide traffic engineering in a Fiber Deep Gigabit world. The "Simple" formula divided into two parts: a DC component (Nsub * Tavg) that is the average traffic utilization and an AC component (K * Tmax_max) to compensate for traffic fluctuations. However, the AC component has to account for many things including traffic fluctuations.

Stepping back and looking at this slightly differently, our objective is to determine a capacity threshold where applications get good QoE for the network traffic utilization. But how do we measure the QoE component? It turns out that most operators and most subscribers rely on some sort of speed test to determine whether the service is meeting its Service Level Agreement (SLA). Interestingly, it also turns out the speed test is one of the most sensitive applications to increased network latencies and utilizations so it is actually an ideal choice to monitor SLA.

At its most pure form, the traffic engineering requirements are:

$$C >= Tburst + Tdata \tag{2}$$

where:

C is the required bandwidth capacity for the service group
Tburst is the bandwidth target used to meet the SLA test
Tdata is the overall network bandwidth at the time of Tburst

For an operator who wants to have a pure best effort service with no SLA guarantees, then Tburst can be set to zero. However, for the many operators under a regulatory microscope, then Tburst will be equal to Tmax_max, the maximum offered service tier. There is a middle ground here as well. Some operators may choose to support a fraction of the advertised service rate. So, for example, an operator might want to guarantee 75% of the service rate, so Tburst would equal 0.75 * Tmax_max. For the remainder of this paper, it is assumed that Tburst equals Tmax_max as the typical scenario.

The Tdata component of the above equation is more complex. It must be an estimate of the data utilization during the SLA test. These tests might typically run from 15 seconds to a minute or two. The Tdata component will obviously vary from time interval to time interval. One can estimate Tdata by measuring the average bandwidth in the service group and then adding an additional margin to achieve our expected QoE.

1. New Basic Formula

Refining our traffic engineering formula now comes up with this basic one:

$$C >= Tmax\_max + Tavg\_sg + QoE\_margin \tag{3}$$

where:

C is the required bandwidth capacity for the service group
Tmax_max is the highest Tmax offered by the MSO.
Tavg_sg is the average bandwidth consumed by a service group during the busy-hour
QoE_margin is additional margin required due to data utilization fluctuations

This becomes our base formula going forward. The first two components, Tmax_max and Tavg_sg are readily available and/or measurable. Our traffic engineering research can now focus on defining the QoE margin component.

But how does this base formula relate to our earlier 2014 "Simple" formula? The answer is quite well. It turns out that the "Simple" formula estimates Tavg_sg using the number of subscribers and the average bandwidth per sub. And it is using a QoE margin of 0.2 * Tmax_max when K=1.2. See the example below:

$$C \geq Tmax\_max + (Nsub*Tavg) + (0.2*Tmax\_max) = (Nsub*Tavg) + (1.2*Tmax\_max) \quad (4)$$

where:

C is the required bandwidth capacity for the service group
Tmax_max is the highest Tmax offered by the MSO.
Nsub is the total number of subscribers within the service group
Tavg is the average bandwidth consumed by a subscriber during the busy-hour

2. Tavg_sg – Operational Considerations

There are multiple considerations for the Tavg_sg component of the new formula. For people in operations, Tavg_sg is more easily attained than either the Tavg or Nsub component in the old formula. The base formula (3) is something that they can measure on a SG by SG basis. Once the QoE margin is established, then this formula can be used as a basis to determine when a SG approaches its maximum capacity before it must be split. So Tavg_sg has operational advantages.

For network capacity planning, it may be desirable to predict required capacity for different SG with different service tier mixes. The generic (Nsub*Tavg) falls short in this respect. By expanding Tavg_sg and breaking out data utilization by service tiers or other groupings, required capacity can be estimated by:

$$C \geq Tmax\_max + \sum Nsub(i)*Tavg(i) + QoE\_margin \quad (5)$$

where:

C is the required bandwidth capacity for the service group
Tmax_max is the highest Tmax offered by the MSO
Nsub(i) is the number of subscribers on the i^{th} service tier
Tavg(i) is the average BW consumed per sub during the busy-hour on the i^{th} service tier
QoE_margin is additional margin required due to data utilization fluctuations

Referencing back to Table 1 shows an example of different Tavg for different service tiers. When combined with the service tier distribution, then a weighted average can be calculated to find Tavg_sg.

As can be seen, the Traffic Engineering formula can adapt as needed. Either using Tavg_sg when appropriate, or decomposing it down to individual service tier components.

3. QoE Margin – the Magic Delta

Tmax_max and Tavg_sg are relatively straightforward components, so that leaves most of the complexity for future traffic engineering research on the QoE margin. This is the AC component, or the Delta bandwidth on top of the static bandwidth utilization. Our job is to take the magic out of the Delta bandwidth.

3.1. “Simple” Formula – still valid after all this time

The “Simple” formula is still valid. As discussed above, it maps quite well to the new base formula where QoE_margin = 0.2 * Tmax_max. However, it is just as important to understand its limitations. It works quite well for SGs with a couple hundred subscribers. The “Simple” formula may be a bit of overkill as the SG size shrinks to 100 subs or less.

Its simplicity is its strength. It is well suited for planning and quickly getting a ballpark estimate of capacity needed. In a golf analogy, think of it as the drive that gets you much closer to the hole. Operators should feel comfortable to continue to use this formula.

3.2. Operational Thresholds

As mentioned above, Tavg_sg is often easily measured on a SG by SG basis and is an important metric from an operational perspective. But what should be used for the QoE margin component with it?

One thing to consider from our data analysis is standard deviation, or more specifically the coefficient of variation. Depending on how much margin an operator might want to build into their system, they might consider adding two to three standard deviations as their QoE margin. The standard deviation may come from CMTS monitoring tools or research as described earlier in this paper.

Our early research has shown that a 100 sub SG might have a coefficient of variation in the 35% to 50% range. 2-3 standard deviation would then mean that anywhere from 70% to 150% of the measured Tavg_sg would be added as the QoE_margin. For example, if this SG measured Tavg_sg = 100 Mbps, then QoE_margin would be between 70 and 150 Mbps depending on the coefficient of variation used and the number of desired standard deviations.

As another example, consider a Remote PHY CCAP Core port that supports 1000 subscribers. Suppose that it has a measured Tavg_sg = 1.2 Gbps. From our early research, its coefficient of variation might be 10%, so three standard deviations would require an additional 0.36 Gbps for the QoE_margin.

Instead of using standard deviations, an operator might use a percentile level (e.g. 98%, 99%, 99.9%) to determine the Tavg_sg + QoE_margin. This might be accomplished with CMTS monitoring tools that provide a histogram for a SG similar to the results in Figure 11. The QoE_margin would equal the difference between the selected percentile level (e.g. 98%) and the measured Tavg_sg.

3.3. Big Data Analytics

What our research has found is that there is a massive amount of data and many complicated variables at play here. It turns out that providing sufficient QoE for traffic engineering is a problem that is suited to Big Data Analytics. This work is still in its infancy. Our goal is that Big Data Analytics can be leveraged to not only select optimum QoE margins in existing networks, but become a tool to predict how our networks will morph and the QoE margins of the future.

Conclusion

Many new innovations are finding their way into cable operators' plants and many others are on the way, including DOCSIS 3.1, Remote PHY Distributed Access Architectures, and Fiber Deep networks with DOCSIS FDX. All of these will significantly impact how operators manage their traffic engineering and network capacity planning.

To enhance our Traffic Engineering formula, an intimate understanding of subscriber bandwidth behavior was needed. This paper took a detailed look at a year's worth of live consumer data collected from a single cable site. The massive amount of data took samples during peak busy hour and tracked every packet to allow for traffic analysis down to the second.

Statistics were gathered and many different bandwidth trends uncovered. Some of the key variables of interest include traffic consumption based on:

- Differing Service group sizes – size matters; significantly increased variation for small SG
- SG to SG variation – substantial for small SG
- Subscriber service tiers – you get what you pay for: higher Tavg for higher tiers
- Time of day – peak busy window stretches to 2-3 hours
- Day of week – not much difference but Sunday is the busiest while Saturday is the quietest
- Month to month – erratic growth, but in line with industry's 35% CAGR

With these trends in hand, the impacts on existing network capacity models showed how they might morph traffic engineering in a Fiber Deep Gigabit world. A new basic formula evolved into:

$$C >= Tmax\_max + Tavg\_sg + QoE\_margin \quad (3)$$

where:

C is the required bandwidth capacity for the service group
Tmax_max is the highest Tmax offered by the MSO.
Tavg_sg is the average bandwidth consumed by a service group during the busy-hour
QoE_margin is additional margin required due to data utilization fluctuations

Tmax_max and Tavg_sg are relatively straightforward. Tavg_sg is often easily measured on a SG by SG basis and is an important metric from an operational perspective. Tavg_sg replaces the Nsub * Tavg component from the older "Simple" formula. That leaves most of the complexity for future traffic engineering research on the QoE margin. This is the AC component, or the Delta bandwidth on top of the static bandwidth utilization. Our job is to take the magic out of the Delta bandwidth.

The "Simple" traffic engineering formula is still as valid as ever. It provides an easy method to quickly get bandwidth capacity estimates. However, the newer traffic engineering formulae have been developed to provide more accuracy and handle a wider range of conditions from small Fiber Deep SG to very large CCAP cores.

There are different ways to estimate the QoE margin. This paper discussed several of these. Some use statistical measurements for a SG such as standard deviation or the coefficient of variation as well as percentiles. These statistics might be derived from CMTS monitoring of that SG or from analysis from a very large collection of data over time. Some examples were shown. It is important to note that the operator can choose how much margin they would like to build in. This may change from region to region based on a particular country's regulatory environment.

What our research has found is that there is a massive amount of data and many complicated variables at play here. It turns out that providing sufficient QoE for traffic engineering is a problem that is suited to Big Data Analytics. This work is still in its infancy. Our goal is that Big Data Analytics can be leveraged to not only select optimum QoE margins in existing networks, but become a tool to predict how our networks will morph and the QoE margins of the future.

Acknowledgements

The authors would like to gratefully acknowledge the assistance of Ben Widrevitz for his countless hours of collecting the live data and processing the raw data. Without that, this paper would not have been possible.

Bibliography & References

[CLOONAN_2013] Tom Cloonan, Jim Allen, Tony Cotter, and Ben Widrevitz, "Advanced Quality of Experience Monitoring Techniques for a New Generation of Traffic Types Carried by DOCSIS," Proceedings, The NCTA Cable Show Spring Technical Forum (June, 2013).

[CLOONAN_2014NCTA] Tom Cloonan, Mike Emmendorfer, John Ulm, Ayham Al-Banna, and Santhana Chari, "Predictions on the Evolution of Access Networks to the Year 2030 and Beyond," Proceedings, The NCTA Cable Show Spring Technical Forum (April, 2014).

[CLOONAN_2014EXPO] "Predictions on the Evolution of Access Networks to the Year 2030 & Beyond"; T. Cloonan, M. Emmendorfer, J. Ulm, A. Al-Banna, S. Chari, The Cable Show NCTA/SCTE Technical Sessions Spring 2014

[CLOONAN_2015] "Lessons from Telco and Wireless Providers: Extending the Life of the HFC Plant with New Technologies," Tom Cloonan et. al., The NCTA Cable Show Spring Technical Forum, May, INTX 2015

[CLOONAN_2016] "Using DOCSIS to Meet the Larger Bandwidth Demand of the 2020 decade and Beyond" Tom Cloonan, Ayham Al-Banna, Frank O'Keeffe; The NCTA Cable Show Spring Technical Forum, INTX 2016

[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

[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

[ULM2_2014] "Scaling Traditional CCAP To Meet The Capacity Needs Of The Next Decade" John Ulm, Tom Cloonan; 2014 SCTE Cable-Tec Expo

[ULM_2016] “Giving HFC a Green Thumb: A Case Study on Access Network and Headend Energy & Space Considerations for Today & Future Architectures” John Ulm, Zoran Maricevic; 2016 SCTE Cable-Tec Expo

[ULM2_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

Abbreviations

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

MSO	Multiple System Operator
N+0	Node+0 actives
Ncast	Narrowcast
NFV	Network Function Virtualization
NSI	Network Side Interface
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiplexing Access (Upstream)
OLT	Optical Line Termination
ONU	Optical Network Unit
OOB	Out of Band
OPEX	Operating Expense
OTT	Over the Top
PHY	Physical interface
PNM	Proactive Network Maintenance
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
QoS	Quality of Service
RF	Radio frequency
R-OLT	Remote OLT
RPD	Remote PHY Device
R-MACPHY	Remote MAC-PHY
R-PHY	Remote PHY
RX	Receive
SDN	Software Defined Network
SG	Service Group
SCTE	Society of Cable Telecommunications Engineers
SNR	Signal to Noise Ratio
TaFDM	Time and Frequency Division Multiplexing
Tavg	Average bandwidth per subscriber
Tmax	Maximum Sustained Traffic Rate – DOCSIS Service Flow parameter
TX	Transmit
US	Upstream
VOD	Video on demand
WDM	Wavelength Division Multiplexing