IS NIELSEN READY TO RETIRE?

LATEST DEVELOPMENTS IN BANDWIDTH CAPACITY PLANNING

JOHN ULM, ENGINEERING FELLOW, BROADBAND SYSTEMS

TOM CLOONAN, CTO, NETWORK SOLUTIONS

MIKE EMMENDORFER, SENIOR DIRECTOR OF SOLUTION ARCHITECTURE AND STRATEGY, OFFICE OF THE CTO

JEFF FINKELSTEIN, EXECUTIVE DIRECTOR OF STRATEGIC ARCHITECTURE, COX COMMUNICATIONS

JOHNPAUL FIORONI, SENIOR DIRECTOR, ENGINEERING OPERATIONS, COMCAST CORPORATIONS

ARRIS™

ARRIS

TABLE OF CONTENTS

ABSTRACT

For several decades now, Nielsen's Law has been a trusted guide to the broadband industry in predicting our highest tier of service rates (a.k.a. Billboard rate), and traffic consumption has merrily followed suit. However, some operators have recently experienced an unexpected dip in growth rates. Is this a sign of things to come or just a 1-year hiccup?

The paper takes a deeper look at underlying factors that are driving Broadband growth rates. For recent history, it has been video driving the growth engine. But the video world is changing as the industry starts migrating to 4K Ultra-HD (UHD) using High Efficiency Video Coding (HEVC) encoding. We'll shed some light on how far that can carry growth into the future.

A closer look is taken at the impacts of Nielsen's Law, as it predicts the growth of the highest service tier level being offered. But this might only represent a tiny fraction of the overall subscriber base. What about other service tiers, will they always follow suit? We analyze service tier distribution and consumption to see its impact on HFC longevity. This analysis is combined with recent novel developments in traffic engineering calculations to look at the potential impact on the longevity of HFC. Could it be that the vast bulk of subscribers might stay on HFC forever?

The primary author, John Ulm, would like to heartily thank his co-authors for providing the research material, data and insights necessary to pull this paper together. It must be noted that the opinions, views and conclusions reached in this paper are those of the primary author and not necessarily those of the co-authors.

BACKGROUND

Nielsen's Law of Internet Bandwidth

Our industry is fully aware of Nielsen's Law, also known these days as the "Billboard Internet Speed". Nielsen's Law of Internet Bandwidth simply states that a high-end user's connection speed grows by 50% per year. This means it doubles roughly every 21 months. It turns out that this "Law" has held fairly consistently for three decades now, having started with 300 baud phone modems back in 1982. Nielsen's Law is important because the highest data speed offered is a determining factor for sizing the network. Is it possible that a service provider will offer a residential 10+ Gbps Internet service by 2025? When will Nielsen's Law Break? It is unlikely that a 50% compounded annual growth rate (CAGR) for the Max offered service tier will last forever. Moore's Law broke and Nielsen's Law will too, but when and at what rate?

Ultimately, the Internet service tier being offered is a decision of the service providers and they have the power to simply pull back the lever of growth provided competitive pressures permit them to do so. As the Cable Service Providers reach the limits of HFC, they must decide if the continued investment for 50% CAGR is warranted based on their revenue return for the services requiring the growth. What user applications will drive this investment? Do they need to switch to Fiber to the Premise (FTTP) for their entire customer base or for only a small percentage of their Internet customer base that demand the highest Service tiers?

Cloonan's Curves

Some research of Dr. Thomas J. Cloonan, CTO of ARRIS and a co-author, combined with Nielsen's Law is captured in the "Max Internet Data Services Tier Offering Downstream" graph in this section. Dr. Cloonan begins with the data rate offered since 1982, charts growth through to the present day and finally out to the year 2030. This data, referred to as Cloonan's Curve, also reflects the historical 50% CAGR as does Nielsen's Law. The data service portion of the model is predictable but at some point, as with Moore's Law, Nielsen's Law may alter and not continue on this 50% CAGR trajectory forever.

While the 50% CAGR downstream assumption may or may not always be true, it is used as a baseline to compare with other scenarios. These predicted, future values from [Cloonan_2014] are shown in Figure 1 for a traffic model that assumes the 50% CAGR of the past continues for both Tavg and Tmax values moving forward into the future.

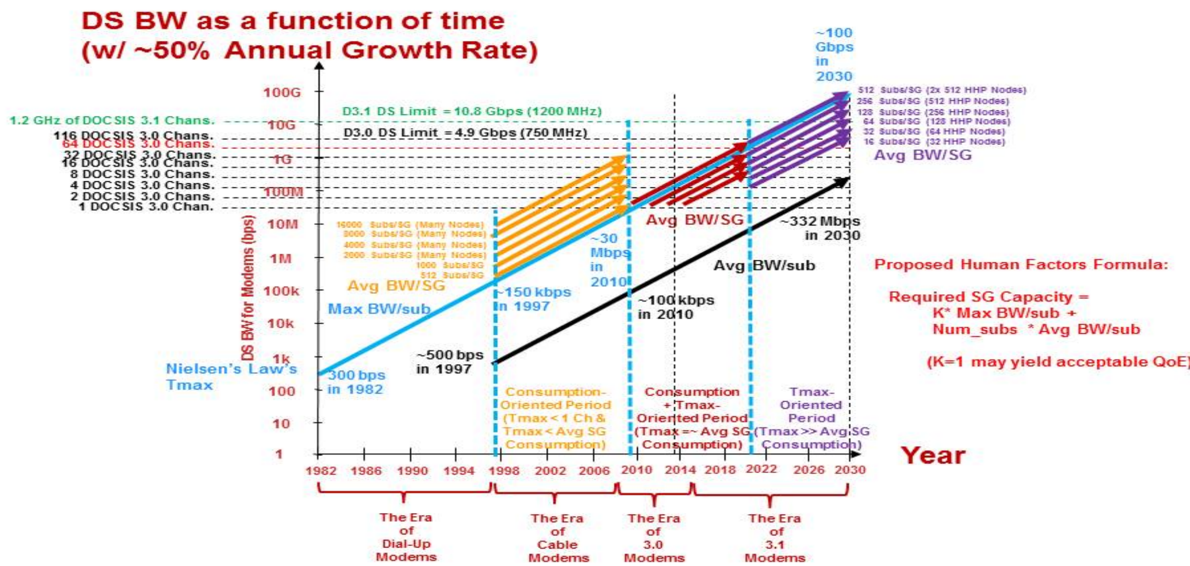

Figure 1 – DOCSIS HSD Downstream Traffic Engineering Predictions with 50% CAGR in Future

The years range from 1982 to 2030 and cover several different eras of modem service, including the Dial-Up modem era, the cable modem era, the DOCSIS 3.0 (D3.0) modem era, and the upcoming DOCSIS 3.1 (D3.1) modem era. The first plot to consider is the light blue plot, which shows Nielsen's Law Curve that identifies the expected Billboard Bandwidth (Tmax) values on a year-by-year basis. If these trends continue, then Nielsen's Law curve predicts that the Tmax value for a high-end modem may be on the order of 100 Gbps by 2030. The max capacity of a HFC plant with D3.1 is reached in the middle of next decade, about 10 years from now.

The second set of plots illustrates average bandwidth consumption rates as a function of time. The black plot illustrates the approximate average per-subscriber bandwidth consumed by a single subscriber during the busy-hour period of time (typically 6pm to 12am). This chart assumes that the Tavg values also grow with a 50% CAGR. The multi-color plots show the average bandwidth consumed per service groups (SG), for a variety of SG sizes. For each Era, the collective SG sizes continue to shrink.

The family of curves illustrating the growing service group average bandwidth levels is divided into three different operating regimes, which in turn define three different periods of time in the life of the HFC plant. The three periods are:

- The Consumption-oriented Period from 1997-2009 (illustrated in orange)

- The (Consumption+Tmax)-oriented Period from 2009-2021 (illustrated in maroon)
- The Tmax-oriented Period from 2021-2030 (illustrated in purple)

The Consumption-oriented Period occurred in the early days of DOCSIS deployment. During this period, service groups were extremely large and the SG aggregate average bandwidth dominated bandwidth capacity planning. Splitting SG effectively doubles the capacity and was the primary tool for traffic engineers.

Many MSOs transitioned to the (Consumption+Tmax)-oriented Period roughly around the 2009 time-frame in conjunction with the MSO transition to DOCSIS 3.0 channel-bonding. This is where we are today. In this period, The SG Aggregate Average Bandwidth roughly equals the rate of the highest service tier identified by the DOCSIS value: Tmax.

This has allowed traffic engineers to use a very simple rule to estimate the size their High Speed Data (HSD) networks:

$$\text{Required Service Group Bandwidth Capacity} = 2 \times \text{Tmax} \qquad (1)$$

The Rule of Thumb can often be used as a quick guideline for estimated capacity needs. However, MSOs have developed tools and optimizations based on volumes of empirical data and traffic analysis that allow them to engineer networks for high quality of experience as key variables change. This Rule of Thumb has been morphing over time. It was originally 3X and has been shrinking with time (and smaller Service Groups, SG). Currently, some operators have even started using 1.6X as their rule of thumb for a potential very high speed service. Since the Rule of Thumb was not always consistent, [Cloonan_2014] provided a simplified formula for estimated Service Group bandwidth capacity:

$$\text{Required Service Group Bandwidth Capacity} = S * \text{Tavg} + \text{Tmax} \qquad (2)$$

where S is the number of subscribers within the Service Group, Tavg is Per-Subscriber Average Busy-Hour Bandwidth, and Tmax is the Maximum Sustained Traffic Rate (Billboard Bandwidth) offering to the subscribers. To make use of this formula, one needs to make some estimated predictions about the future values of S, Tavg, and Tmax. This is challenging, but historical data can help make these predictions.

Coincident with this Consumption+Tmax-oriented period, operator’s may reach a 1:1 mapping between HSD SG and fiber nodes. Once this happens, further SG splits require node splits in the HFC plant. During these years, MSOs will undoubtedly add more bandwidth capacity (i.e. DOCSIS channels) to the HSD service tier to stay ahead of the

Tmax growth, and they may also perform node splits to periodically reduce the service group’s aggregate average bandwidth consumption (i.e. Tagg = S * Tavg) level.

However, with shrinking SG size, we are now approaching a new time called the Tmax-oriented Period (see purple plots). During this period, the SG aggregate average bandwidth is significantly smaller than Tmax. This can also be derived from equation (2) as S becomes very small making the Tavg term insignificant.

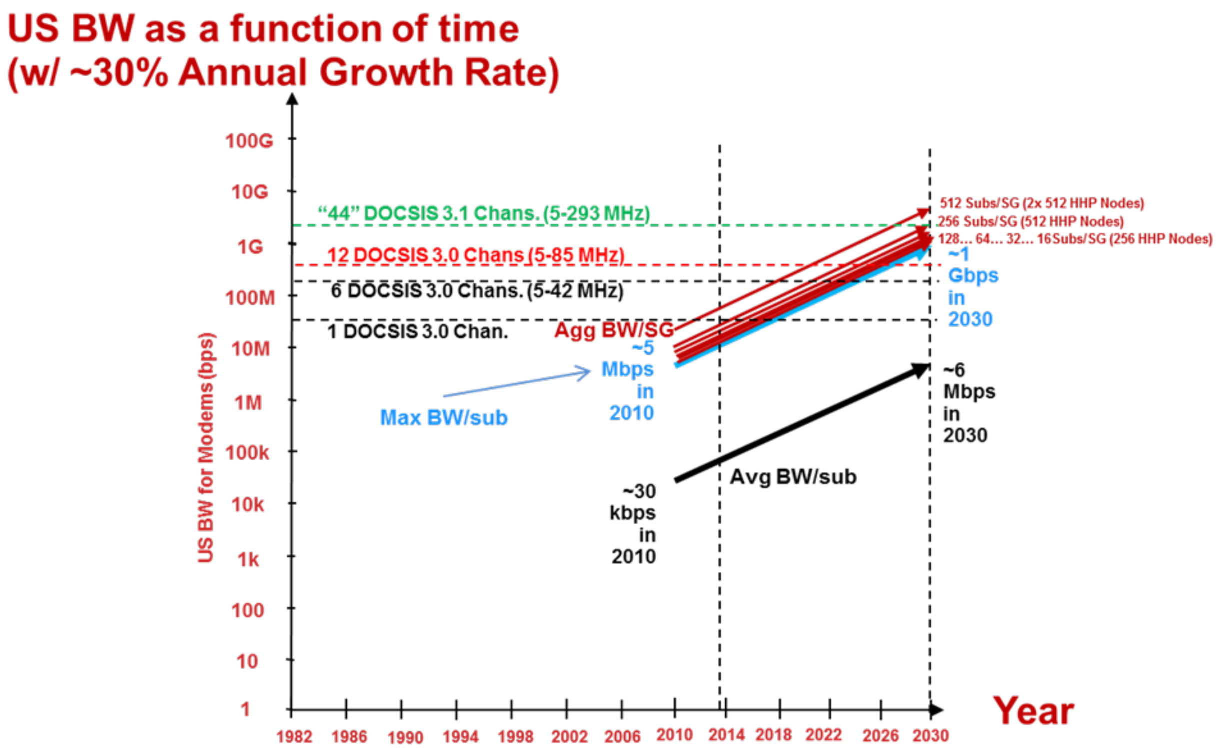

Figure 2 – DOCSIS Upstream Predictions with 30% CAGR in Future

The previous discussion focused on Cloonan’s Curves for the downstream. Figure 2 shows the Cloonan Curves for the Upstream. The industry has only started tracking the upstream in detail relatively recently, so there is not as much history. What is known is that Nielsen’s Law of 50% CAGR does not apply to upstream speed growth.

While upstream growth rates vary significantly from operator to operator or site to site, collectively the upstream seems to have been following a 30% CAGR. Figure 2 maps out potential upstream capacity growth to the year 2030 using this CAGR. At this point, it appears that the HFC upstream capacity is reached in roughly the same timeframe as the downstream dependent on whether an operator migrates to an 85 MHz or 204 MHz upstream split.

From a traffic engineering perspective, Tmax will eventually dominate and splitting SG size further to reduce aggregate average bandwidth buys very little relief while getting constantly more expensive. In the near future, MSOs may find it interesting to explore some new approaches:

1) Having different service group sizes for Video on Demand (VoD) and Switched Digital Video (SDV) and DOCSIS HSD
2) Using FTTP (PON or Extended-Spectrum RFoG) for select subs with high Tmax values to reduce the effective Tmax on the HFC plant
3) Using FTTP (PON or Extended-Spectrum RFoG) for select subs with high average bandwidth utilizations to reduce the Tavg on the HFC plant
4) Simply limit the expansion of Tmax (if permitted by the competition)

A paradigm shift in traffic engineering is coming that moves away from node splits and towards enhancing Tmax burst capabilities with a transition to DOCSIS 3.1 and increased spectrum, e.g. 1218 MHz.

Emmendorfer's Network QoE Capacity Formula

While some MSOs currently size their network on a method of multiplying the billboard speed by either doubling (2X) or tripling (3X) the billboard speed to determine the amount of DOCSIS capacity per service group, there are potential issues with using this Rule of Thumb method for DOCSIS network sizing. While this method worked for the most part, it was not particularly accurate, needed tweaking over time and will eventually break as shown in [Emmendorfer_2014, Cloonan_2014].

As a result of all this, [Emmendorfer_2014] developed a formula to definitively describe the required SG bandwidth capacity over time. This paper unveiled a new traffic engineering and capacity planning formula to help MSOs properly size their networks to accommodate Nielsen's Law, Traffic, and Competition. This is called the Network Quality of Experience (NQoE) Capacity Formula and is a unit of measure that may be used to size any service provider network and network technology.

The NQoE Capacity Formula goals are:

- Achieve Max Service Tier even during busy periods
- Allocate appropriate amount of network resources
- Configurable to accommodate any data network
- Accommodate estimates of service tier and traffic growth rates
- Achieve Max Service Tier through next network capacity adjustment

From this work, the amount of DOCSIS HSD capacity which might be required to provide an "adequate" Quality of Experience level to a Service Group can be roughly described by the following NQoE Capacity formula shown in Figure 3. The details of each term are explained in [Emmendorfer_2014]. It is noted here that this formula, like Cloonan's simplified formula, has two major components: Max Service Tier being offered to consumers and the service group aggregate average bandwidth during Peak hours.

The NQoE Formula more accurately forecasts network capacity requirements. It uses the combination of service tier plus traffic per service group as well as other factors to estimate network capacity needs. Since the Rule of Thumb method does not take into account traffic or service group size among other factors it is not as accurate.

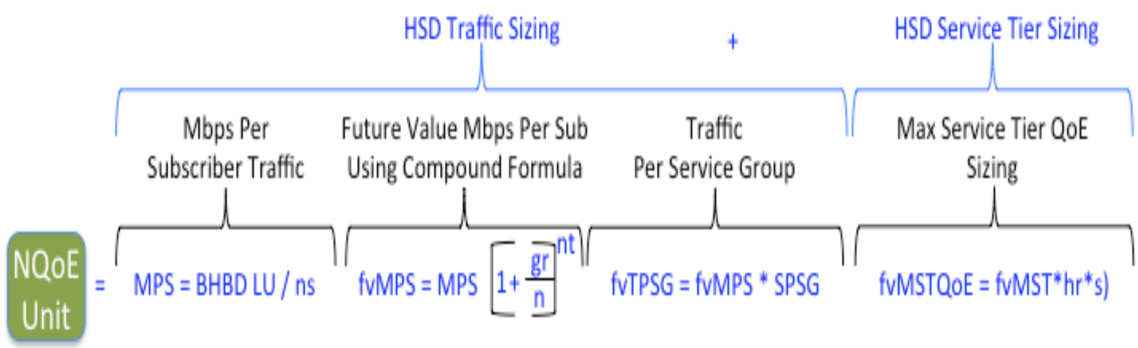

Figure 3 – Network Quality of Experience (NQoE) Capacity Formula

The Max Service Tier component is straightforward and well known. This is the Service Providers "Billboard" rate that has been following Nielsen's Law for three decades. The effects of this were clearly shown in the Cloonan Curves above. The service group capacity is less obvious and requires more considerations.

Operators will need to track these key drivers and levers that force network change, like Nielsen's Law of Max Services Tier Growth Rate and also traffic growth Rates for proper network planning.

Measuring and Estimating Customer Traffic

In addition to the service tier offered to consumers, the actual usage of the network by the consumers is a critical factor for network planners. This measurement of the total amount of bandwidth or traffic in a serving area is often measured during busy hour(s) to drive operator traffic engineering limits.

The Internet Traffic CAGR is determined by the combination of consumer usage patterns, service provider speed tiers offered, and various technologies that will influence network usage. The upstream traffic has its own CAGR that is different than downstream traffic. The growth rate of traffic will vary widely even within a service provider, because usage patterns will be different between demographics. Traffic growth rates are very hard to forecast because there are many possible influences to drive traffic growth.

According to the 2013 Cisco Visual Networking Index (VNI), Global IP Traffic is forecast to expand three-fold from 2013 to 2018. Note that this is roughly a 21% CAGR which is much less than Nielsen's Law. However, this may not account for a significant amount of local content that is being cached by operators in their private CDN.

Some Key Observations from [Emmendorfer_2014] considered for forecasting capacity:

- Traffic growth has not always grown at 50% CAGR
  - Industry observed 20-30% through most of 2000s
- Symmetry Downstream-to-Upstream was not always 7:1 DS:US
  - 2:1 through most of 2000s was observed

[Emmendorfer_2014] posed several Key Questions that must be considered:

- When will Nielsen Law of 50% CAGR for max service tier offered break?
- Which "Downstream Traffic" CAGR do you believe and when will it break?
- Which "Upstream Traffic" CAGR do you believe and when will it break?

Making the Most of HFC

[Cloonan_2014] looks into the crystal ball and uses basic trend analyses and traffic engineering analyses to create predictions that attempt to answer challenging questions like when will HFC to FTTP transition take place. By its very nature, traffic engineering is an imperfect science requiring many guesses and assumptions and approximations to be utilized. Sometimes, these guesses can be wrong. However, when implemented properly, traffic engineering predictions can be very useful in developing rough plans for future HFC networks and in the sizing of future HFC networking equipment.

Capacity Impact of DOCSIS Migration Plans

The existing HFC cable network infrastructure offers great flexibility and provides high levels of digital capacity, on the order of multiple Gbps to the home but significantly less from the home. The industry recently released the latest generation of specification: DOCSIS 3.1 (D3.1). This technology provides many new capabilities that promise to make the most of HFC. From a traffic engineering perspective, it is important to understand the capacity of each generation of consumer premise equipment (CPE), as this ultimately limits the maximum service tier that can be offered to subscribers. The North American CPE downstream capacity migration is shown in the Table 1 below:

CPE HSD DS Capacity	D3.0 in 2012	Current D3.0 Phase	D3.1 DS Gen 1	D3.1 DS Gen 2 – Full Spectrum		D3.1 Futures (TBD)
DS Range (MHz)	54 - 1002	108 - 1002	108 - 1218		258 - 1218	500 - 1794
DS QAM Level	QAM: 6 MHz, 256-QAM		OFDM: 24-192 MHz, To 4096-QAM			To 16384
# DS CPE Channels	8 x 6 MHz	24 – 32 x 6 MHz	24 – 32 x 6 MHz + 2 x 192 MHz	6x192 MHz	5x192 MHz	6-7x192 MHz
CPE DS Capacity	300 Mbps	~1 Gbps	~5 Gbps	10+ Gbps	9+ Gbps	>15 Gbps

Table 1 – N.A. CPE Downstream Capacity by DOCSIS Generations

The first column shows net data capacity with 8-bonded D3.0 downstream channels of about 300 Mbps for CPE that have been deployed heavily over last several years. This capacity is sufficient to offer 100 or 150 Mbps Service Tiers in the downstream using the Rule of Thumb method of 3X or 2X respectively. The cable industry is now just starting to deploy the next generation of 24 to 32 bonded D3.0 channels CPE. As shown in column 2, these CPE allow net data capacity to increase to about 1 Gbps, a factor of 3-4 over the 8-bonded channel CPE.

On the horizon is the first generation of D3.1 CPE. These devices will be backwards compatible and support 24-32 bonded downstream (DS) channels that are D3.0 compatible. In addition to this, the D3.1 CPE will support two ultra wide 24 to 192 MHz wide OFDM (Orthogonal Frequency Division Multiplex) channels and be able to bond between the OFDM and the D3.0 Single Carrier QAM (SC-QAM) channels. The OFDM channels will also support up to 4096-QAM modulation for a 50% performance improvement in spectral density over D3.0 SC-QAM channels dependent on HFC plant conditions. This means the first gen D3.1 CPE can eventually provide net data capacities in the 4-5 Gbps range (column 3). The D3.1 devices will also support extended HFC spectrum to 1218 MHz.

Looking into our crystal ball, Moore's Law should help 2nd generation D3.1 CPE devices to achieve full spectrum operation by supporting up to six 192 MHz wide OFDM channels. By this time, the net data capacity will also be dependent on where the upstream split has been moved. For an 85 MHz system (column 4), the CPE can utilize 108 MHz to 1218 MHz in the downstream over six OFDM channels for net data capacity that exceeds 10 Gbps. If the upstream split moves to 204 MHz (column 5), the downstream capacity is reduced to five OFDM channels in 258 MHz to 1218 MHz for a net data capacity of around 9 Gbps.

The D3.1 specification put hooks in for potential future systems that might support spectrum to 1794 MHz and modulation rates of 16,384-QAM. Seven of these enhanced OFDM channels might provide more than 15 Gbps capacity in the downstream. Spectral widths larger than 1794 MHz might be feasible, e.g. in an RFoG system; and this could provide even higher capacities.

The net upstream capacity for DOCSIS CPE devices are shown in the Table 2 below:

CPE HSD US Capacity	D3.0 in 2012	Current D3.0 Phase	D3.1 US Gen 1 – 8x6.4 MHz, 2x96 MHz			D3.1 Futures (TBD)
US Range (MHz)	5-42	5-85	5-42	5-85	12-204	12-400
US QAM Level	ATDMA: 6.4 MHz, 64-QAM		OFDMA: 96 MHz, 1024-QAM			To 4096
# US CPE Channels	4 x 6.4 MHz	8 x 6.4 MHz	8 x 6.4 MHz, 1x37 MHz	8 x 6.4 MHz, 1x80 MHz	8 x 6.4 MHz, 2x96 MHz	4x96 MHz
CPE US Capacity	100 Mbps	200 Mbps	~250 Mbps	~600 Mbps	~1.5 Gbps	>3 Gbps

Table 2 – CPE Upstream By DOCSIS Generations

Today, the 8x4 cable modems have been deployed in volume in the field. They contain 4-bonded D3.0 upstream channels to go with the 8-bonded downstream channels. Each SC-QAM upstream channel has a net data capacity of about 25 Mbps at 6.4 MHz for total upstream (US) capacity of about 100 Mbps (column 1). The latest D3.0 CPE are 24x8 or 32x8 devices with 8-bonded D3.0 US. This doubles upstream capacity to about 200 Mbps (column 2), but may require a 65 MHz or 85 MHz upstream to fully support eight US channels of 6.4 MHz each.

The 1st gen D3.1 devices will also support 8-bonded SC-QAM US channels for backwards compatibility plus two 96MHz OFDMA (Orthogonal Frequency Division Multiplex Access) channels. These OFDMA channels will support up to 4096-QAM modulation rates in the CPE, but the spec only requires 1024-QAM operation in the CMTS. Because the upstream is burst oriented, the D3.1 OFDMA channels can be time division multiplexed with the 3.0 SC-QAM channels. This means D3.1 CPE can burst across the entire US spectrum with the higher capacity OFDMA modulation rates.

So, the total upstream capacity now becomes very dependent on the available upstream spectrum. For a 42 MHz system, the D3.1 CPE could potentially burst at rates of 200-250 Mbps (column 3). That's more than twice the capacity of a 3.0 SC-QAM system. For an 85 MHz upstream split, the D3.1 maximum burst capacity is in the 500-600 Mbps range (column 4). Finally, a 204 MHz upstream split (column 5) allows both OFDMA channels to run at full tilt and achieve a net data capacity of 1.5G bps. Using the

NQoE formula with smaller SG will allow operators to offer a 1Gbps upstream service tier.

Finally, looking out a decade at potential D3.1 futures, it may be possible to see CPE with a 400 MHz upstream split that support more than 3 Gbps utilizing 4096-QAM modulation.

How will cable operators be able to handle this much capacity growth? Reducing average Service Group size from 500 Households Passed (HHP) down to 50 HHP means that they must now support 10 times as many Service Groups. These issues around scaling Converged Cable Access Platforms (CCAP) are discussed in detail in [ULM_2014].

VIDEO – DRIVING THE NIELSEN'S LAW GROWTH ENGINE

The Over the Top (OTT) video providers such as Apple TV, Amazon, Hulu, Netflix, YouTube and others entering the On-Demand video market are driving traffic and peak period consumption upward. Data from the Sandvine Global Internet Phenomena Snapshot: 1H 2014 North America, Fixed Access report illustrates that 64% of the peak hour downstream traffic is real-time entertainment from these OTT providers.

The percentage of real-time video entertainment has gone up significantly over the last several years whereas a decade ago this was a rounding error. Traffic in that time frame was dominated by web browsing and file sharing. It is very important to understand those types of traffic and the volume of traffic as a percentage.

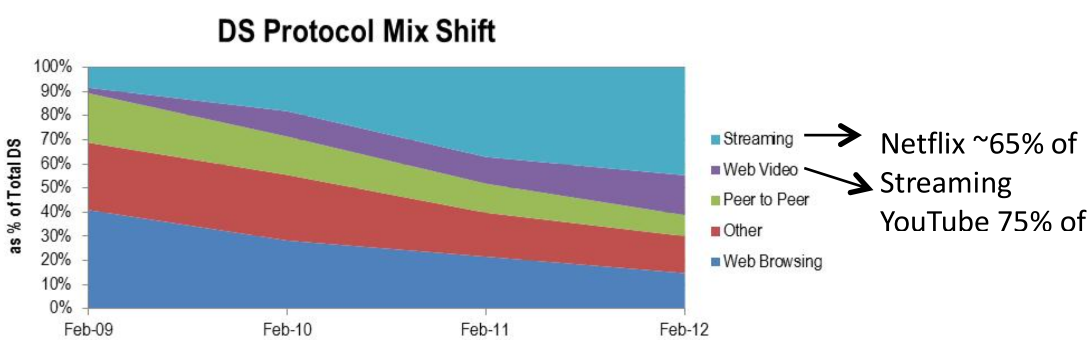

Figure 4 – Downstream Application Mix

Data from some operator sites corroborates these trends. Over a three year period from 2009 to 2012, streaming and Web videos increased their downstream bandwidth consumption share from a paltry 10% to almost 60% share in three years. Netflix accounted for 65% of the streaming traffic while YouTube representing 75% of the Web video traffic. These gains came mostly at the expense of Web browsing and Peer-to-peer (P2P) that dropped their collective share of downstream bandwidth from 61% to 24% share of the traffic. Meanwhile, the traffic mix in the upstream has remained relatively unchanged over the last three years.

One key for network planning is to separate out video traffic from the traffic data and estimate the amount of users watching video over the top (OTT) as they are most likely not simultaneously watching an MSO delivered video offering. OTT video services driving the explosive growth in Internet traffic have seen annual growth rates range from 40% to over 100%.

The increase in real-time video entertainment has also moved the symmetry between downstream and upstream traffic from 2:1 or even 4:1 from a decade earlier to now over 10:1. We will examine the impact of video service as a key driver for traffic consumption and growth rates.

Some key video service observations from [Emmendorfer_2014]:

- OTT Video is currently the driver for HSD traffic growth
- OTT is becoming an overlapping video service with each MSO's managed video service
  - A person watching OTT is not watching CATV service (and vice versa)
- Video Capacity Planning Factors
  - 2.61 = number of people per household (2008-2012 U.S. Census)
  - 2.5 – 2.93 = # of TVs per household (2010-2011 according to Nielsen)
  - Today – 66% are SDTV and 34% are HDTV according to Nielsen
    - When might it be 66% are HDTV and 34% are 4K UHDTV?

The main conclusion is that the migration to a single network for video delivery will likely extend the life of the HFC network and perhaps the service group size levels. A typical node and customer take rate could be served if capacity and spectrum are allocated to a single video delivery network. If there were two video delivery networks, MPEG Transport Stream (TS) and IP Video over DOCSIS, the allocation of capacity would likely be higher, since it may be hard to predict which network a customer would use in a given evening.

Improvements from New Blends of Video Compression Techniques

MSOs will definitely be looking to various technology paths to squeeze the bandwidth of MSO-managed video into a smaller portion of the HFC spectrum. Using these technologies will permit the MSOs to accommodate the rapid growth of HSD bandwidth on their HFC plant for a longer period of time. In addition to changes in the blend of video transmission techniques, MSOs will also be looking very seriously at changes in the blend of video compression techniques.

Many digital video compression technologies are now available for MSOs to consider and utilize. These include the MPEG-2, H.264/AVC (a.k.a. MPEG-4), and H.265/HEVC compression techniques. Each new generation of compression technology provides approximately a 2X gain in capacity. In conjunction with compression, statistical-multiplexing technologies that simultaneously pack multiple streams in a single channel are also a key part of potential capacity gains.

It is very informative to take a look at the relative bandwidth requirements for various screen resolutions for these different compression technologies:

Resolution	MPEG-2	MPEG-4/H.264	HEVC/H.265
SD: 480i30	**2.5 – 4 Mbps**	1.5 – 2.5 Mbps	-
HD: 720p30	5 – 10 Mbps	2.5 – 5 Mbps	1.5 – 3 Mbps
HD: 720p60 (or 1080i60)	**9 – 18 Mbps**	4.5 – 9 Mbps	**2 – 5 Mbps**
Bluray:1080p60	12 – 25 Mbps	6 – 12 Mbps	3 – 6 Mbps
UHD: 4Kx2Kp60	-	25 – 40 Mbps	**12 – 20 Mbps**
UHD: 8Kx4Kp60	-	-	25 – 60 Mbps

Table 3: Video BW Requirements with Different Resolutions and Compression Techniques

The table above is a rough estimate based on equivalent video quality for a given screen resolution. The compression efficiencies of these various technologies continue to evolve and improve over time, but the comparative quality of each compression technique varies from technology to technology, from vendor to vendor, from month to month, and from video quality expert to video quality expert.

From the table, it is apparent that a move towards HEVC encoding would provide the best compression ratios. As a result, MSOs will likely want to move in that direction. However, they will need to migrate in that direction from their current MPEG-2 equipment, and there are limits placed on them by the fact that the existing equipment deployed in the field may not support the improved compression technologies. As a result, this migration path may be a two-step path, moving from MPEG-2 to H.264 first for installed CPE, and then moving from H.264 to HEVC with newer CPE capable of that format.

Leading video expert Dr. Sean McCarthy describes in [McCarthy_2014] the science in how Ultra HD approaches the limits of the human visual system.

To understand the long term impact of Video on traffic growth, one must consider both the evolution of screen sizes in conjunction with the evolution in compression technologies. Over the next decade or two, we'll continue to see HDTV replace SDTV for the 2nd & 3rd TV screen in the home. This will also be supplemented by portable devices like tablets and smart phones with HD resolution screens. The next decade will see the Ultra HD (UHD) screens replacing HD for the large Primary screen in the home. Even if this transition completes in a decade, it will occur faster than the HD transition. So by the end of next decade, we can expect SDTVs to be essentially retired with the majority of viewing screens to be HD resolution with anywhere from one quarter to one half of the screens being UHD.

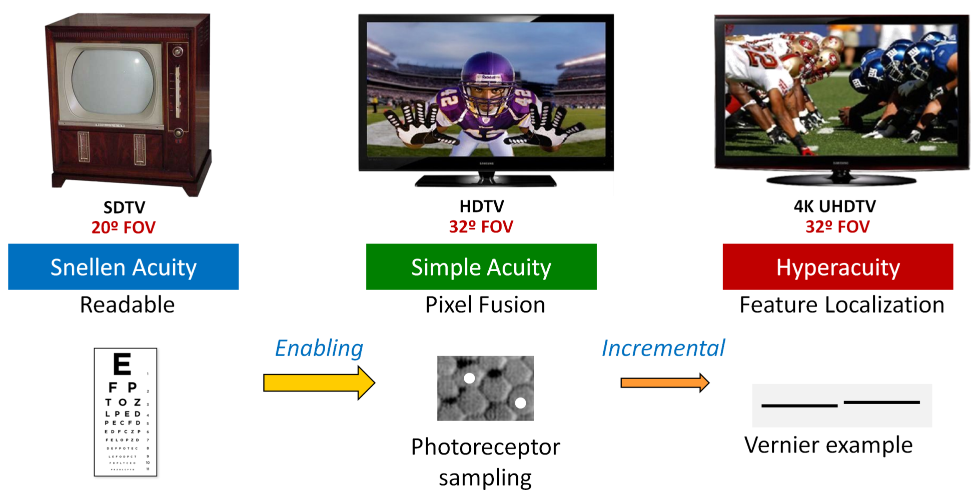

Figure 5 – Video Resolution Trends

Earlier this year, Parks Associates announced new research "4K Today: Bringing Ultra HD to Market" that predicts Ultra HD 4K TVs will follow a similar, but faster, growth pattern as HDTVs. The HDTV market took 15 years to reach 82% penetration of U.S. broadband households. This report predicts that 4K TVs will reach mass-market pricing in the next 2-3 years and top 80% of households in approximately 10-12 years. Note that while HDTV has reached 82% household penetration today, it still only represents 34% of total TVs in homes as SD TVs are typically found for the 2nd, 3rd & 4th TVs in the house. So even if Ultra HD 4K TVs reach 80% penetration in less time than HDTV, one can expect to see that Ultra HD 4K TVs might only account for one third of total TVs in 10-12 years from now. And this percentage is dropped even further once tablets are also considered as a viewing device inside the home.

The other interesting data point is to look at the video bit rates highlighted in red in the above table. It turns out that the bit rates for HD and 4K UHD using HEVC a decade from now will be very comparable to the SD and HD video bit rates being used today in MPEG-2 TS video delivery systems. What this means to Nielsen's Law is that video will continue to drive growth for a few more years as the industry maintains dual video delivery infrastructures and finishes the transition to UHD for the primary screen and HD screens for secondary viewing. However, after this video delivery transition, its growth rate will flatten very quickly. Some other application will then need to step up if Nielsen's Law is to continue!!

Challenges from New Blends of Video Resolution Techniques

Many MSOs will begin to migrate towards IP Video over DOCSIS transport for their video in the next few years. There are numerous drivers behind this migration. For example, the use of DOCSIS for video transport will provide much larger channels than the 6 MHz channels of MPEG-TS, so better packing of HD and UHD programs is permitted. Also, the use of D3.1 will offer higher spectral efficiencies that will permit more programs to be carried within a given chunk of spectrum. In addition, this migration may drive costs lower as MSOs will be able to utilize a single head-end and CPE infrastructure for all of their services (voice, data, and video). Plus, having all services on a single infrastructure may permit better sharing of bandwidth between the bursty services than is permitted when using separate, silo-based infrastructures for each service (as is done today).

But as those MSOs begin to move video streams to IP Video over DOCSIS delivery systems, they will want to get through the "simulcast window" more quickly, permitting them to eventually reclaim most of the bandwidth associated with MPEG-TS video. However using Switched IP Multicast for Linear IP Video transmissions can help operators to capitalize on bandwidth efficiencies during this bubble.

In terms of targeting a single video service delivery network, this would likely reach the same network capacity, yet other costs need to be considered. The main conclusion is that the migration to a single network for video delivery will likely extend the life of the HFC network and perhaps the service group size levels.

A typical node and customer take rate could be served if capacity and spectrum is allocated to a single video delivery network. If there were two video delivery networks, MPEG TS and IP Video over DOCSIS, the allocation of capacity would likely be higher since it may be hard to predict which network a customer would use in a given evening. The use of the Network Quality of Experience (NQoE) Formula will help service providers of all types and technologies size their network more accurately.

BEYOND NIELSEN'S LAW: THE OTHER SERVICE TIERS

Nielsen's Law applies to a service provider's Top Tier, or Billboard rate. Yet this tier may only represent a fraction of one percent of the provider's subscribers. As the Billboard rate approaches the capacity limits of HFC, service providers will need to initiate a migration plan to move the Top Tier to another technology such as Fiber to the Premise (FTTP). At this point in time, it becomes critical that the service provider understands what is happening to all of their other service tiers in order to map out an effective migration strategy for the entire subscriber base. This paper is the first to detail the trends and make projections on all of the other Service Tiers that Nielsen's Law is ignoring.

Service Tiers – Functional Decomposition

In the early DOCSIS days, service providers had a single speed service tier on the order of a few megabits per second. Over the years as Nielsen's Law pushed the Top Tier higher, room was made below that to start introducing new levels of service. Today, there may be a half dozen or more service tiers available from any given operator. This number will only go up over time. In general, we will group the service tiers into three broad categories:

Performance Tiers; Common or Popular Tiers; and Economy Tiers

The Performance Tiers offer the highest broadband speeds; with the Billboard Tier at the pinnacle of this group. The Common or Popular Tiers contain the majority of the provider's subscribers. This group tends to be in the middle with regards to internet speeds. The Popular category also tends to be the most focused on video entertainment as their main use case. Finally, the lowest performance group is the Economy Tiers. These go after cost sensitive markets and often compete against older DSL technologies.

Service Tiers – Broadband Speeds

The most obvious distinguishing characteristic of any service tier is its broadband speeds, most notably in the downstream. The table below collects various service tier broadband rates, upstream & downstream, from three different operators (labeled MSO 'A', 'B' and 'C') grouped according to the above categories. This data was sampled roughly over the last year or so and does not necessarily represent the same instant in time for each MSO.

Observe that the Top Billboard Tier is anywhere from 6X to 20X higher than the popular Common Tiers. Also, note that MSO 'C' recently announced a 305 Mbps downstream tier in a select market due to FTTP competition, but its Top Tier that is generally available everywhere else is 150 Mbps downstream.

DS, US rates (Mbps)	MSO 'A'	MSO 'B'	MSO 'C'
Top Tier, Billboard speed	250, 25	150, 20	150-305, 20
Performance Tiers	50-105, 10	50-100, 10-20	50-100, 5-10
Popular, Common Tiers	25, 5	25, 5	15, 1
Economy Tiers	3-6, 1	1-5, 0.5-1	3, 0.5

Table 4: Broadband Speeds for Various MSOs

We took a closer look at the service tier offerings of one of these MSO over a four year period. There were approximately a half dozen different service tiers being offered that got categorized in Table 5.

The Top Tier, downstream Billboard rate increased five-fold over this four year period from 50 Mbps to 250 Mbps. This works out to a 50% CAGR which aligns almost perfectly with Nielsen's Law. Let's take a much closer look at what else is going on with the various Service Tiers during this timeframe.

The first observation is that the upstream rate for the Top Tier is growing much slower than the downstream rate. The upstream rate increases half as much (2.5X) from 10 Mbps to 25 Mbps. This is a hair over a 25% CAGR for the upstream speed. The net effect of the slower upstream growth rate is that the DS-to-US asymmetry grows from 5X to 10X over this period.

DS, US rates (Mbps)	4 yrs ago	Couple yrs back	Current
Top Tier, "Billboard speed"	50, 10	105, 15	250, 25
Performance Tiers	25, 2	25-50, 4-10	50-105, 10
Popular, Common Tiers	8-12, 2	8-15, 2	25, 5
Economy Tiers	3-6, 1	3-6, 1	3-6, 1

Table 5: Broadband Speed Tiers as a Function of Time

At the other end of the spectrum, the Economy Tiers remain perfectly flat with no growth whatsoever. However, the Common Tiers see a comparable growth of 2.5X in both the upstream (2 Mbps to 5 Mbps) and the downstream (8-12 Mbps to 25 Mbps) directions.

Common Tiers are seeing a slower CAGR, e.g. 25%, in both US & DS

Since the majority of the subscribers reside in the popular Common Tiers, this has potentially profound impacts on the network traffic engineering. It should be noted that Service Tier Broadband speeds are updated sporadically, often one increase every 12-24 months; so it is difficult to pin a single CAGR percentage on the growth. It is more important to notice the overall trend and the observation that it tends to be half the growth rate of the Top Billboard speed.

CPE Deployments and Service Tier Restrictions

So what is preventing the popular Common Tiers from growing at the same Nielsen's Law Rate as the Top Billboard Tiers? One key economic factor is the service tier restrictions that are imposed by a given cable modem technology. When a service provider introduces a new Top Tier, it is also most likely the most expensive tier offering the highest broadband speeds. This impacts rollouts in two ways. First, the economics are favorable because there is substantial additional increase in revenue associated with this subscriber. Second, as will be shown in the next section, it only applies to a very small fraction of the entire subscriber base. So even if the operator needs to update the CPE device with the newest DOCSIS technology, the total number of truck rolls needed is very limited.

When an operator increases the broadband speeds for the popular Common Tiers, it is often a very big deal and is usually accompanied with a big marketing campaign and lots of fanfare. This is because it applies to the vast majority of the operator's subscribers AND it is being done with no or very little associated increase in revenue. The bulk of the Common Tier often has a CPE with a certain older DOCSIS technology (e.g. D2.0, D3.0). A given DOCSIS generation will have a hard limit on the data capacity it can support. Once an operator wants to exceed this, then the CPE device must be replaced. Now the operator has substantial capital outlay with no associated increase in revenue. This creates much resistance to the Common Tier growing at Nielsen's 50% CAGR and causing rapid CPE churn.

So what limits on the service tier are imposed by various DOCSIS technologies that are widely deployed today? DOCSIS 1.1 and 2.0 modems supported a single downstream channel of roughly 38 Mbps (N.A.). Using the Rule of Thumb method, operators might provision broadband service tier speeds up to 12 or 15 Mbps. The D2.0 modem upstream capacity could be as high as 25 Mbps but is often restricted to be lower due to network conditions. Operators are more conservative as well with lower speed upstream service tiers, so an upstream service tier of 1-5 Mbps is more common for D2.0 modems.

DOCSIS 3.0 modems introduced channel bonding with a significant leap in capacity. Four bonded DS channels produce a 150 Mbps capacity in N.A. enabling 50 or 75 Mbps service tiers by the Rule of Thumb method. Eight bonded channels bumps the data capacity to 300 Mbps allowing 100 or 150 Mbps service tiers. In the upstream, four bonded channels can produce up to 100 Mbps capacity depending on network conditions. This has allowed operators to offer upstream service tiers in the 20 to 25 Mbps range. By using the Emmendorfer NQoE Formula and reducing SG sizes, operators may be able to squeeze another 30-50% improvement out of these service tier levels.

As can be seen, these hard modem limits show up as typical service tier offerings. The process of transitioning from D2.0 modems to D3.0 modems for the popular Common Tiers is an expensive, long drawn out multi-year process. Because of this, the Common Tier may actually get spread across two speed bands during the transition.
Since operators want their CPE investment to stay in the field for as many years as possible, the replacement CPE may be several DOCSIS generations more current. This is why there is such keen industry efforts to get D3.1 devices deployed, even though existing service tier broadband speeds are well below the D3.1 capacities.

Service Tiers – Subscriber Distribution

The Table below summarizes data collected on subscriber distribution across the various service tiers from dozens of different markets. The data from MSO ‘A’ was averaged across all of its markets while the MSO ‘B’ data provides a range across the markets. This provides a glimpse into some market to market variations that can be experienced.

The subscriber distribution across the various service tiers varied quite a bit between markets but two key factors were noted: market demographics (low vs. high income) and market competition (e.g. DSL or FTTP). The Tier distribution is also very dependent on how each operator prices and markets the various service tiers.

As expected, lower income markets have a much larger percentage of subs in the Economy Tiers, while high income SG have a larger percentage in the Performance Tiers including the Top Billboard Tier.

Subscriber Distribution	MSO #A	MSO #B
Top Tier, Billboard speed	0.2%	0.1-1%
Performance Tiers	24%	5-18%
Common, Popular Tiers	64%	50-72%
Economy Tiers	12%	15-40%

Table 6: Subscriber Mixes across the Service Tiers

In looking at some historical data across several years, there appears to be a slight increase in the percentage of subscribers on the Performance Tiers, including the Billboard rate; at the expense of the Common Tier. This change is only on the order of a couple percentage points. Meanwhile, the Economy Tier subscriber % appears to be holding steady. This mix is something that the industry should continue to monitor.

Consumption by Subscriber, by Service Tiers

Besides service tier rates, the other major factor in network capacity planning from the NQoE Formula is the actual bandwidth consumption per service group. This is typically calculated from the average bandwidth per subscriber times the total # of subs per SG. However, if an operator starts migrating its highest Performance Tiers and/or the heavy usage subscribers off of HFC onto FTTP, what is the net traffic engineering impact for the remaining subs on the HFC?

The traffic consumption per subscriber is not uniform across the entire subscriber base. Looking at total consumption (upstream + downstream) in one of the markets, the Top 1% of users accounts for ~15% of total consumption; the Top 10% account for slightly more than half; while the Top 25% represents 70% of total. The peak hours showed similar patterns.

Not surprisingly, average traffic consumption per sub goes up with increasing broadband speed tier. The major question then is what is the relationship between consumption and service tier rates? If the service tier speed doubles, will consumption also double? Or not move at all? Or is it something in between?

Data was collected from a couple sites regarding consumption per sub for each of their service tiers, and then the data was normalized relative to the bandwidth per sub of the most popular Common Tier (i.e. 25 Mbps downstream). The results are shown in the chart below.

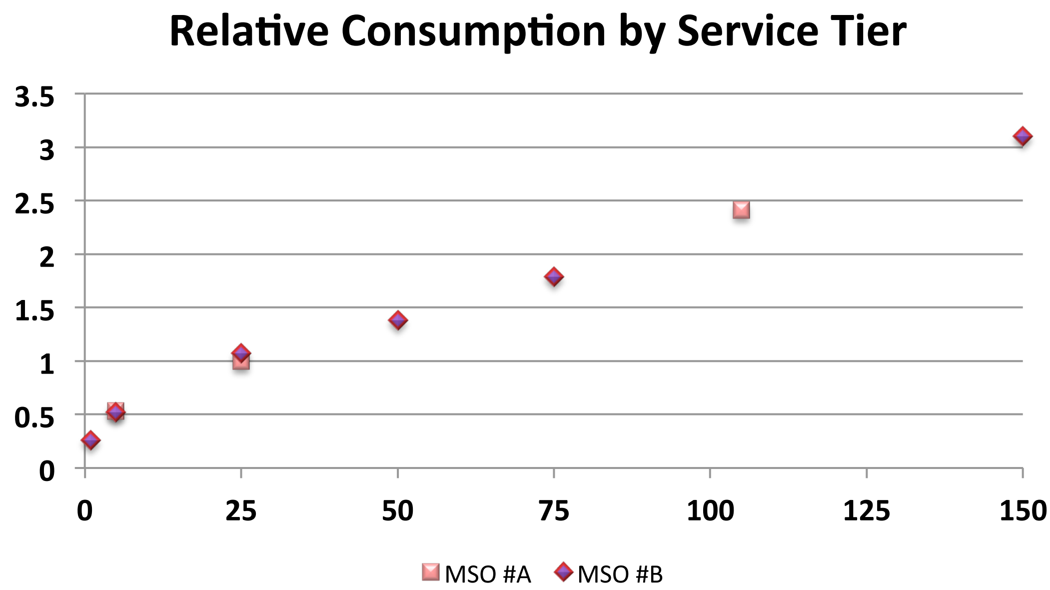

Figure 6: Relative Consumption by Service Tier Bandwidth

There appears to be a very linear relationship between service tier rate and consumption. Increasing the service tier rate from 25 Mbps to 50 Mbps (i.e. 100%) resulted in a ~40% increase in consumption per sub; increasing from 25 to 75 Mbps (200%) results in ~80% increase in consumption; and so forth. While extremely interesting, this sample data set is much too small to draw any definitive conclusions. It is highlighted here to raise questions and to help motivate future industry research in this area.

One trend that became apparent is that the traffic consumption at higher Performance Tiers is growing at faster rates than the other tiers. This should be monitored closely in the future.

The other wild card when considering long term consumption growth is the imposition of data usage caps. Several operators have been exploring the impacts of caps on very heavy data users. Widespread acceptance of data usage caps might make a noticeable change on long term consumption growth rates.

Upstream/Downstream Asymmetry

Nielsen's Law with its 50% CAGR only applies to the downstream Internet service rates. Since upstream spectrum is one of the most precious resources in the cable world, it is important to understand the relationship between the upstream and downstream.

In the figure below from the Sandvine Report, the Downstream-to-Upstream Ratio (DS:US) is shown during an average day. This data is useful to understand the ever increasing spread between the traffic directions. In the early 2000s, this number had far more symmetry. During the peak of peer-to-peer traffic, some traffic assessments showed a near 1:1 ratio but more often this ratio was 2:1. The increased amount of real time video entertainment is causing the ratio to spread dramatically. This spread in the ratio of DS to US traffic is expected to continue until there is some upstream application that will consume bandwidth at a faster rate and duration during peak period. Maybe Business or other Cloud Services is the symmetric application that will help drive upstream usage.

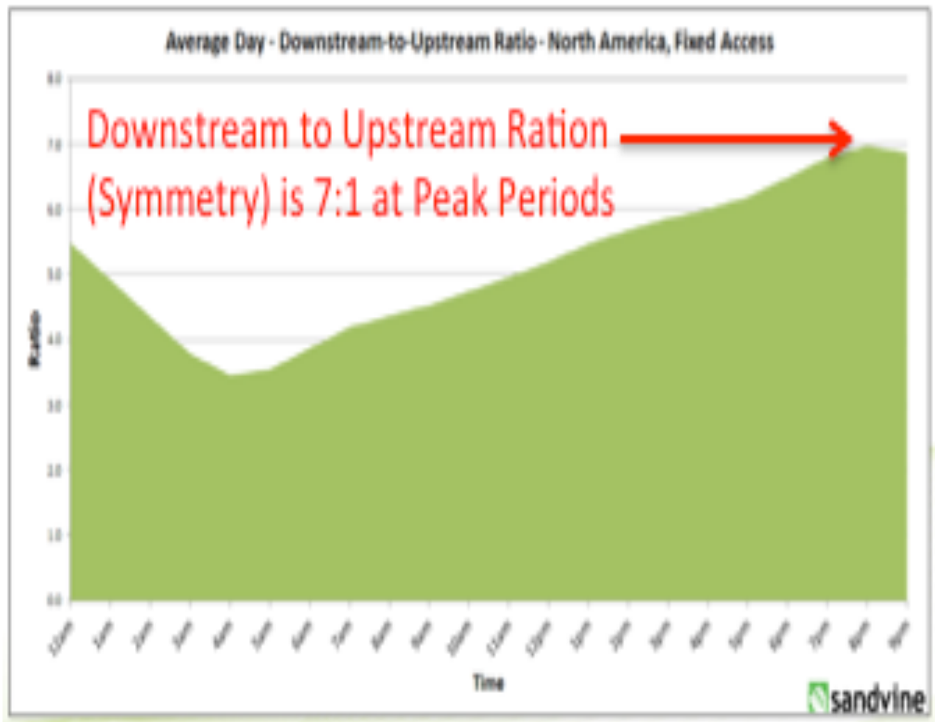

Figure 7 – Average Downstream-to-Upstream Ratio Complied by Sandvine

Sandvine also reports that BitTorrent continues to lose share and now accounts for just 6.0% of total traffic during peak period. In their last report, they revealed that for the first time File sharing as a whole accounted for less than 10% of total daily traffic, and that trend continues with File sharing now responsible for just 8.3% of daily network traffic. This demonstrates a sharp decline in share from the 31% of total traffic revealed in Sandvine's 2008 report.

With many cable and DSL providers considering implementing usage based billing, an examination of usage distribution is of interest to many. According to Sandvine, in North America, the top 1% of subscribers who make the heaviest use of the network's upstream resources account for 47% of total upstream traffic. The comparable downstream users account for 12% of downstream bytes. Meanwhile the Top 10% heaviest users consumed almost 80% of the upstream and about 50% of the downstream. At the opposite end of the usage spectrum, the network's lightest 50% of users account for only 7% of total monthly traffic.

Looking at some MSO historical data on upstream Year-over-Year growth rates per sub, the upstream growth rate peaked in 2008 with ~80% growth from the year before and then dropped significantly before bottoming out in 2011 at 15%.

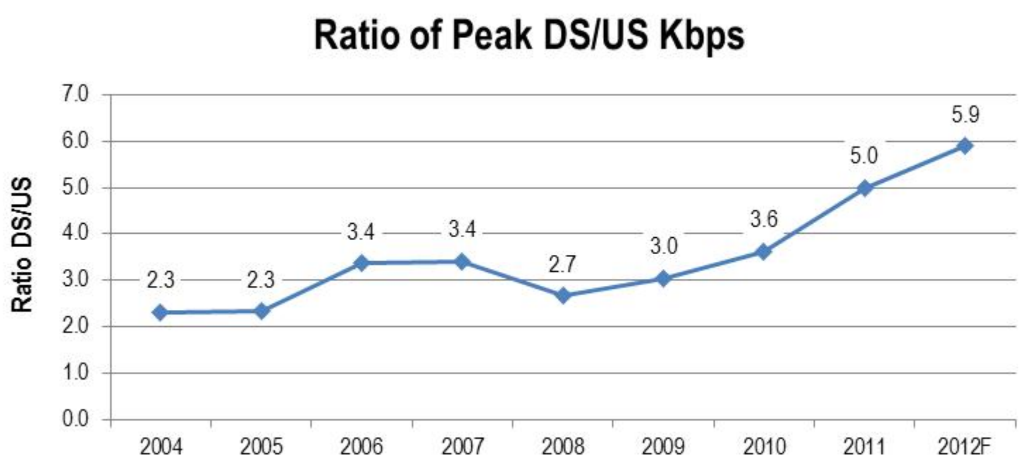

Figure 8 – Ratio of Peak DS/US

The DS:US ratio historical data is shown in the Figure 8. Back in 2004, the DS:US ratio was 2.3 and has been rising since then except for the 80% upstream growth spike in 2008. By 2012, the DS:US volume ratio reached almost 6. And early indications for 2013 indicate a continuation of this trend.

HFC – MAKING A GOOD THING LAST

If an operator only looks at Nielsen's Law, they might conclude that the HFC becomes obsolete in less than a decade from now. One may wonder how these dates can be extended. One obvious way is for the MSO to alter one or more of the terms within the NQoE Formula. As an example, some MSO marketing teams may have the option of slowing down the growth rates on Tmax values within that formula. This may or may not be possible depending on the Tmax challenges from the competition. But if it is possible, then this can have profound effects on the life-span of the HFC network.

As [Cloonan_2014] indicates that the change in the growth rate of Tmax has greatly extended the sunset of the 1.2 GHz HFC plant deep into the 2030 timeframe. The change in growth rates has also made node splits to smaller node sizes more valuable again, and it also extends the HFC sunset deeper into the future. If MSOs are able to initiate this change in Tmax growth (and a similar change occurs in Tavg values), then a lot of value can be obtained.

The authors will not describe a single path into the future, because it seems certain that the technologies used within HFC plants will likely be quite varied as MSOs move forward into the future. Market bifurcations will undoubtedly occur as different MSOs select different paths based on their different constraints, starting points, and biases.

- Most of these MSOs will take advantage of DOCSIS 3.1 as a means of increasing the available bandwidth capacity of their existing HFC plants.
- Some of these MSOs will change the split on their Upstream spectrum to be 85 or 204 MHz in an attempt to provide more Upstream bandwidth capacity. The 204 MHz split may permit Tmax = 1 Gbps Upstream service offerings.

Focus Shifts – From Node Splits to Spectrum Enhancements

Node splits have long been a trusted tool used by many MSOs throughout the years to reduce the bandwidth demands within a Service Group. Node splits offer no change in the Service Group bandwidth requirements for Broadcast services. They only offer benefits to Narrowcast services- including Switched Digital Video (SDV), Video on Demand (VoD), DOCSIS HSD, DOCSIS Voice over IP (VoIP), and DOCSIS IPTV.

In actuality, node splits only partially help with DOCSIS HSD services. Node splits only reduce the traffic consumption piece of HSD services, but does not affect the bandwidth required for the "Billlboard" Service Tier. [Cloonan_2014] shows that with shrinking node sizes to 256 HHP and below, operators are reaching the point of diminishing returns on node splits. Instead, service provider focus now must shift towards enhancing networks to increase the max burst rate for the Billboard Service Tier. This can be done with a migration to DOCSIS 3.1 and HFC plant spectrum enhancements to 1218 MHz downstream and 85 MHz or 204 MHz upstream.

How Much Upstream Spectrum will be needed?

Thanks to Nielsen's Law, most network capacity planning focus over the years has been on the downstream. This has been accentuated with the video growth in the downstream in recent years. Upstream Billboard bandwidths and average upstream bandwidth consumption rates display much more variation than downstream, with CAGRs at different MSOs showing rates ranging from 10% to 30%.

CPE Technology	DS Tier	42 MHz US Tier	85 MHz US Tier	204 MHz US Tier
D3.0 (32x8)	800 Mbps	65 Mbps (12:1)	125 Mbps (6:1)	-
1st gen D3.1	3 Gbps	150 Mbps (20:1)	400 Mbps (7:1)	1 Gbps (3:1)
2nd gen D3.1	6 Gbps	150 Mbps (40:1)	400 Mbps (15:1)	1 Gbps (6:1)

Table 7: Bandwidths (& DS:US Ratios) Provided By Different DOCSIS Generations

Whether the upstream becomes the limiting factor is very dependent on the DS:US ratio that might be seen in the coming years. The table below looks at the upcoming DOCSIS technologies from 32x8 3.0 modems to the first generation of D3.1 then a full spectrum (6 x 192 MHz) 2nd gen D3.1 device. Potential service tiers were calculated with the NQoE formula assuming fairly small SG of 125 subs or less and shown in Table 7. The upstream service tiers were calculated for 42 MHz, 85 MHz and 204 MHz US splits.

If the DS:US ratio continues to grow to around 12:1, then a 42 MHz system might be able to support an 800 Mbps DS with D3.0 32x8 modem or almost 2 Gbps DS with D3.1. If the ratio stays around the current 6:1 or 7:1 range, then an 85 MHz system is needed to support 800 Mbps DS with D3.0 32x8 modem or 3 Gbps downstream with D3.1. An 85 MHz system with a 15:1 ratio could support a full 1.2 GHz DS spectrum with a 6 Gbps service tier. Note that a 204 MHz split system is only needed if the DS:US ratio shrinks back to 4:1 or 2:1 over time. This requires a new upstream application (e.g. Business Services or maybe some Cloud based applications) to reverse the current trend in the ratio.

With Today's DS:US ratio, 85 MHz Upstream pairs well with D3.1 Systems

BW Capacity Planning for the Popular Common Tiers

Delaying Common Tier Migration away from HFC

Following Nielsen's Law, the Top Billboard Tier will hit the HFC ceiling by early next decade. But does that mean the entire subscriber base needs to be moved to FTTP shortly after this time? The Common Tier broadband speed was shown to be growing at roughly half the rate of the Top Tier. So when will the Common Tier hit the HFC ceiling?

Figure 9 below shows what will happen if the Top Tier continues to follow Nielsen's Law while the popular Common Tier tracks to a 25% CAGR over that same time. The chart shows that the popular Common Tier might survive another 15+ years on the HFC after the Top Billboard Tier has been forced to move off the HFC. This means the operators may have a significantly more time to migrate the bulk of their subscribers to new FTTP technologies.

What does this mean to the operator for a possible HFC to FTTP migration strategy? Our analysis of Service Tier distribution shows that the Top Tier may only account for 1% of the subscribers while the popular Common Tiers account for roughly two thirds to three quarters of the subscribers. This means that within six to eight years from now, an operator may migrate its Top Tier, about 1% of total, to get past 10 Gbps data delivery. The remaining Performance Tiers can then be migrated over the next several to half dozen years. The Performance Tiers may only account for 10% to 20% of the total subscriber base. The operator then has another decade to move the remainder of its subscribers. This gives the operator a 15-20 year window to make the switch.

Delaying Common Tier allows Operators to migrate over 15-20 years

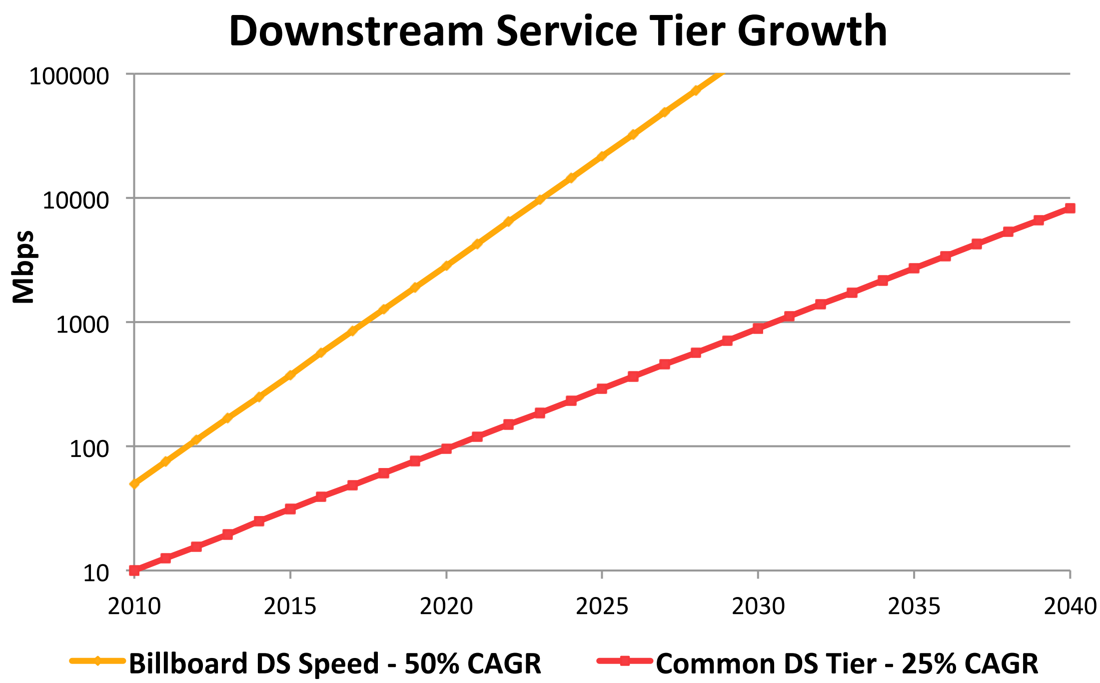

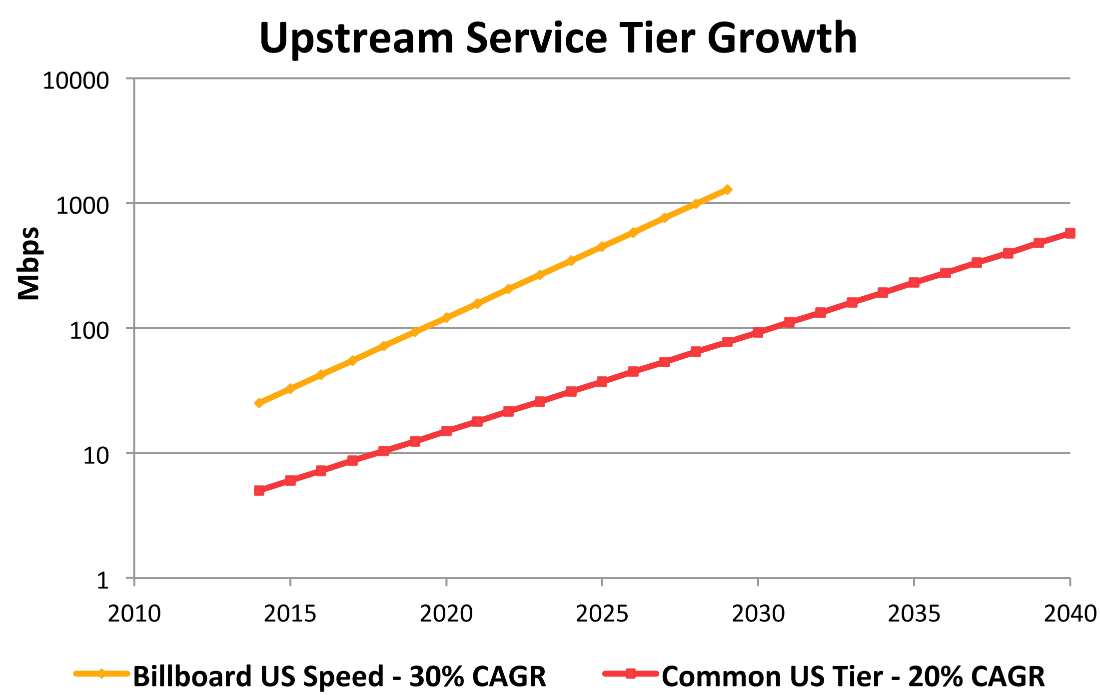

Figure 9 – Bandwidth Speeds per Year for Top Tier, Common Tier

During this lengthy transition, there are some other considerations for deciding which users to migrate next. Some users that need more symmetric capacity such as business services and cloud storage services may hit an upstream limit sooner. These would be good candidates to move. Another possible strategy is to move heavy users off of the HFC first. Some heavy users will be in the Performance Tiers but many will remain in the Common Tiers. As our results showed, the Top 1% of heavy users account for 50% of upstream capacity while the Top 10% account for 80% of upstream capacity and 50% of downstream capacity. Moving these heavy users can significantly reduce aggregate SG capacity and breathe years of additional life into the HFC for the subs that remain.

Video Services and HFC Longevity

The legacy of the cable industry is to provide entertainment to its customers. It started with cable video services and then continued with broadband internet services that has spawned the OTT industry such as Netflix and YouTube. Video entertainment, either operator managed or OTT, is the major application that consumers are willing to pay substantial dollars to receive.

From our earlier analysis, video bandwidth will start to flatten out in the next decade once a majority of the homes have an Ultra HD 4K TV. But how much bandwidth is this per service group? A quick analysis shows that an average 25 Mbps per home could provide a dedicated Ultra HD (UHD) and a couple HD streams to each. A total of 200 active subscribers in a service group would need 5 Gbps which is well within D3.1 downstream capabilities. In reality, there will be additional statistical multiplexing gains and multicast gains that can reduce this bandwidth further if needed.

While the D3.1 downstream is sufficient to meet long term video capacity needs, what about the upstream capacity? It turns out that very Asymmetric DS:US ratio is the norm for video centric applications. Assuming that the HFC contains primarily video centric subscribers long term, then the Upstream IP Acknowledgements (ACKs) for the 5 Gbps of Downstream HTTP-based video would only require ~100 Mbps of Upstream bandwidth, which can be reduced further with IP ACK suppression protocols. Thus, even with the addition of other Upstream traffic, a DS:US ratio of 10:1 to 15:1 may be quite acceptable. This means that an 85 MHz upstream split should be sufficient for this environment. This has the profound possibility that:

> *Video Centric Customers might live on HFC Forever*

Competitive Pressures – HFC vs. PON

One of today's key drivers for continued increases in broadband speeds is the very competitive environment between cable and FTTP providers such as Verizon FiOS and Google Fiber services in North America. Today's EPON/GPON deployments are basically 1 Gbps and 2.5 Gbps downstream capacity. Gazing into the crystal ball, where will PON be a decade from now? 10G PON is available today but not yet cost effective. This is expected to become the mainstream technology over the next decade. However, the big question becomes the accompanying upstream rates. Is it 10G/10G or 10G/1G PON? For residential applications, the 10 G upstream has some significant cost and distance limitations making 1 G upstream the likely candidate.

Getting beyond 10 G FTTP over metro distances (e.g. 40 km to 80 km) cost effectively is extremely challenging. ITU has worked on a 40 G PON solution that is actually 4x10 G WDM based where an ONU device connects to a single 10 G wavelength, so not a true 40 G to each user. A study group in IEEE 802.3 is looking into Next Generation EPON that goes beyond 10 G EPON but finding technology hurdles very challenging. Given the long cycle time to get IEEE 802 standards to market and then volumes to ramp up, this makes us believe that typical cost effective residential PON deployments a decade from now will be 10 G/1 G PON systems.

It turns out that a D3.1 cable system with 1.2 GHz plant roughly provides the equivalent capacity of a 10 G/1 G PON system with 10 G optics in the downstream path and 1 G optics upstream. The actual data capacity of the 10 G/1 G PON system after FEC and overheads are removed is closer to 8.5 Gbps DS and 700 Mbps US. D3.1 can actually exceed the 10 G PON downstream capacity. In an 85 MHz plant, the D3.1 upstream capacity can reach 600 Mbps which is within spitting distance of the 10 G/1 G upstream capacity. A 204 MHz plant can effectively double the 10 G/1 G PON upstream capacity and allow cable operators to offer a true 1 Gbps upstream service tier over cable. In theory, D3.1 spectrum could even be extended beyond 1794 MHz to offer much higher bandwidths (e.g. 20-40 Gbps) within RFoG systems. The use of narrow OFDM sub-carriers within D3.1 could actually help solve optical dispersion issues that may complicate high-bandwidth OOK-based PON systems.

Residential Deployments: D3.1 + 1.2 GHz can match 10 G/1 G PON capacities

With fairly comparable capacity between D3.1 1.2 GHz HFC plant and 10 G/1 G PON systems, the broadband industry may eventually reach a détente that helps to finally break Nielsen's Law.

CONCLUSIONS

Nielsen's Law has dictated network capacity planning for many years. Existing Rule of Thumb models are starting to break as shown by the Cloonan Curves. This has led to the development of Emmendorfer's Network Quality of Experience (NQoE) Formula to more accurately calculate the required capacity over time.

The cable industry is now entering a new era where the DOCSIS Tmax burst parameter starts to dominate the capacity planning. This means that the industry will see a paradigm shift away from node splits to manage capacity towards enhancing burst rate capacity by transitioning to DOCSIS 3.1 and extending spectrum to 1.2 GHz or beyond. These enhancements will give HFC the capacity to rival 10 G/1 G PON in the residential arena, even with an 85 MHz upstream split.

As we investigate how much life is left in HFC, we take a closer look at the impact on Nielsen's Law. It turns out that it is applicable to the Top Billboard Tier, but that might only be 1% of the total subscriber base. Our research shows that the most popular Common Tiers are showing growth rates half that of Nielsen's 50% CAGR. This means that the bulk of subscribers (e.g. 80% to 90%+) might have another 15-20 years left on HFC. This leads to a strategy of slowly migrating subscribers to FTTP as needed. This starts with the Top Billboard Tier and is then followed by the other Performance Tiers. Additionally, operators may move Business Services and other users with upstream heavy applications like cloud storage services. Moving heavy users also extends the longevity of the HFC for the remaining cable subs.

Our research also takes a closer look at video growth that is currently the primary source of the 50% Nielsen's Law growth. It turns out that this will be flattening out over the next decade as Ultra HD 4K TVs get widely deployed. Since Ultra HD approaches the limits of the human visual system, we might reach the point where we can satisfy a small SG of video centric subscribers (e.g. 200 subs) on HFC forever.

The paper's title poses the question on whether Nielsen's Law is ready to retire. While Nielsen should keep very busy for several more years, if these shadows are not altered, then Nielsen will be tending his garden full time sometime in the next 5-10 years.

MEET ONE OF OUR EXPERTS: John Ulm

John Ulm is a Fellow of the Technical Staff within the ARRIS Network Solutions CTO group. In this role he has been investigating strategic technical directions for multi-screen services and broadband bandwidth expansion. Recent activities include research into Multiscreen Adaptive Streaming IP Video solutions; next generation CCAP architecture; next generation HFC architectures; and new HFC protocols including DOCSIS 3.1 and IEEE 802.3bn EPoC.

John's two+ decades in the Broadband industry began as designer, architect and MAC protocol developer at LANcity, pioneering the industry's first cable modem systems. He was one of the primary authors for the Cable Industry's DOCSIS 1.0/1.1 specifications that drove cable modem success during its early days. He also spent time as a Network Processor architect for Nortel and as a senior technical consultant to the Broadband industry with YAS Corp.

REFERENCES

1) [EMMENDORFER_2014] – "Nielsen's Law vs. Nielsen TV Viewership For Network Capacity Planning"; M. Emmendorfer, T. Cloonan, The Cable Show NCTA/SCTE Technical Sessions Spring 2014

2) [CLOONAN_2014] – "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

3) [MCCARTHY_2014] – "Ultra HD – Connecting the Present to the Future"; Dr. S. McCarthy, SCTE Cable-Tec Expo, Fall 2014

4) [MUTALIK_2014] – "Partnership for Extended Capacity: DOCSIS 3.1 with RFoG"; V. Mutalik, M. Schemmann, A. Al-Banna, Z. Maricevic, The Cable Show NCTA/SCTE Technical Sessions Spring 2014

5) [ULM_2014] – "Scaling Traditional CCAP to Meet the Capacity Needs of the Next Decade"; J. Ulm, T. Cloonan, The Cable Show NCTA/SCTE Technical Sessions Spring 2014

6) [ULM_EXPO_2013] – "Unmanaged ABR – How to Control those Unruly Teenagers"; J. Ulm, A. Eshet, N. Radian, SCTE Cable-Tec Expo, Fall 2013

ADDITIONAL MATERIAL:

(1) Nielsen's Law Max Internet Service Tier Offered – Downstream. Source: http://www.nngroup.com/articles/law-of-bandwidth/

(2) Sandvine, Global Internet Phenomena Snapshot: 1H 2014 North America, Fixed Access, Available for Download at www.sandvine.com

ABBREVIATIONS AND ACRONYMS

CAGR	Compounded Annual Growth Rate
CCAP	Converged Cable Access Platform
CMTS	Cable modem termination system
CPE	Consumer Premise Equipment
DOCSIS	Data over cable service interface specification
D2.0, D3.0, D3.1	DOCSIS 2.0, DOCSIS 3.0, DOCSIS 3.1
DS	Downstream
DSL	Digital Subscriber Line
FTTP	Fiber to the Premise
HD, HDTV	High Definition, High Definition Television
HEVC	High Efficiency Video Coding
HFC	Hybrid Fiber-Coaxial
HSD	High-speed Data (broadband service)
HHP	Households Passed
IP, IPTV	Internet Protocol, IP based Television
MPEG	Moving Picture Experts Group
MSO	Multiple system operator
NQoE	Network Quality of Experience (Capacity Formula)
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiplexing Access
OTT	Over the top (video)
PON	Passive Optical Network
QAM	Quadrature amplitude modulation
QoE	Quality of Experience
QoS	Quality of Service
RF	Radio frequency
RFoG	RF over Glass
SCTE	Society of Cable and Telecommunications Engineers
SDV	Switched digital video
SG	Service Group
UHD	Ultra High Definition
US	Upstream
VoD	Video on-Demand
WDM	Wavelength Division Multiplexing