ACCOUNTING FOR TECHIES: TAKING IT TO THE ULTRA

Dr. Robert L. Howald
Dr. Sean T. McCarthy
ARRIS

Abstract

High-Definition video has been a terrific success story for cable operators. It is irresistible to consumers once they're introduced to the vividness of the HD experience. As a high-demand premium service, it delivers precious bottom-line value to the MSO. However, as more subscribers flock to HD, pressure increases on operators to provide more HD programming. A key dilemma becomes allocating sufficient bandwidth and balancing HD demands with competing demand for spectrum from IP services. Fortunately, all operators are engaged in multiple activities that will substantially improve network efficiency. Over time, these various tools allow operators to see a way out of this resource limitation predicament……or do they?

High-Definition service today is HD 1.0. Already in labs, standards bodies, and in demonstrations from trade show floors to "big box" stores are emerging HD 2.0 and even HD 3.0 technologies. The operator's spectrum efficiency toolkit is not improving in isolation. Technology to further enhance the media experience is occurring in parallel, and it comes at the price of more bits-per-second. Not far around the next corner, Quad Format HD (aka 4k HD) aims to provide the next phase of display resolution. High Frame Rate (HFR) HD is poised to make subscribers think differently about dynamic resolution. Finally, Ultra High-Definition, an HD 3.0 candidate, represents a 32x resolution experience compared to HD 1.0.

In this paper, we raid the engineering labs and turn over their technical reports to the accountants to develop the long term Balance Sheet. We will account for the resource Liabilities of increased pixels per frame and increased frames per second. On the Asset side of the ledger, we analyze how new formats interact with available knobs and levers in MPEG-4 (H.264), anticipate the emergence of HEVC (H.265), consider architecture and spectrum evolution, and discuss the numerical implications of each. We also integrate IP Video and assess its role in the network transport transition, in addition to the service transition. The service transition is encumbered by quantifiable Liabilities brought about by the nature and burden of legacy support and multi-format simulcast. We assimilate compounded IP growth rates as part of a full IP transition. Using our itemized analysis, we reconcile the Balance Sheet, and draw conclusions in the context of HFC capacity constraints and optimization. Finally, we present strategies for the HD 2.0 era and consider an era of HD 3.0.

INTRODUCTION

Decades of service evolution – video, voice, and data alike – have given operators a sound historical basis for business planning of new growth scenarios. The prevailing MSO approach has been a very successful pay-as-you grow approach, capitalizing on technologies as they mature while meeting consumer service demands. This approach has benefited from the availability of new HFC capacity, incrementally exploited through fiber deep extensions, RF bandwidth upgrades in the distribution plant, use of WDM to fuel continued segmentation and service options, and migrating to either all-digital services and/or switched digital video (SDV) architectures.

However, while the appetite for all things HD continues to be strong, the lifecycle of already-bandwidth-burdensome HD itself has only just begun. Cable systems deliver 720p and 1080i formats today, while "Full HD" – a term that never had a uniformly understood meaning itself – 1080p already exists in the consumer electronics world of Blu-Ray and gaming consoles. Flat panel televisions continue to become larger, more capable, and lower cost. The CES show earlier this year had dozens of "4K HD" (3840 x 2160p), televisions on display, and sizes as large as 84" (7 feet!) are on the market. 4K HD television sales have surprised analysts, and they have upwardly adjusted their forecasts. By 2015, it is now projected that the majority of TVs sold, if not the vast majority, will be 4K HD capable. Since these TVs are likely to come with format upconversion, users can take advantage of them before there is mass content delivered in the 4K format.

Another accelerating factor is the cycle of television replacement. It has dropped to about 6 years more recently, down from a historical 10 year lifecycle. The move from HD to Ultra High Definition (UHD – to be defined) is anticipated to not be as large of a hurdle for consumers as the move from SD to HD was. This is somewhat ironic, since the relative enhancement to viewing performance from SD to HD is larger than HD to 4K.

The 4K HD format represents 4x the number of pixels of 1080-column HD. Formats definition and early technology also exist beyond 4K UHD in the form of 8K UHD (aka NHK's Japan's Super Hi-Vision). The pixel multiplier for 8K HD is another 4x over 4K HD, with the total pixel grid standing at 7680 x 4320p. As always, "p" stands for progressive scan, but the implicit "60" cannot necessarily be assumed as new video science takes place, as we shall discuss.

With IP CAGR's racing ahead, and quantifiable limits of the HFC architecture's capacity, it is important to account for these emerging video formats destined for the marketplace when we analyze long-term capacity management. In the case of 4K UHD at least, with televisions moving off the shelves more quickly than expected, it suggests a high likelihood of a mass-market technology in a way that the cumbersome size implications of 8K UHD may not.

In the analysis to follow, we will do the accounting – Bandwidth Assets and Bandwidth Liabilities – that give us insight into the future possibilities. Then, because of the limitations of this simplistic Balance Sheet approach, we will build out a long term service and architecture Capacity Management Timeline. A key component of the timeline is that the service evolution will occur in parallel with a major architecture evolution – the IP Transformation. As we shall see, it is valuable – critical – to understand this end objective and the implications. The path to Network Nirvana is a complex balance of new services arriving, old services phasing out, and architectural techniques and tools leveraged adroitly to survive the capacity management challenge of this multi-dimensional transition.

HD 2.0 & HD 3.0: WE'VE ONLY JUST BEGUN

Pixel Perfect

Last fall (October 2012) the Consumer Electronics Association (CEA) came out with definitions of Ultra High Definition to help to market higher resolution televisions. And, it is not just the big box stores and CE vendors that are in on the act – the ITU is developing recommendations for both 4K HD and 8K HD. And, the European

Broadcasting Union now defines 4K HD and 8K HD as UHD-1 and UHD-2, respectively. We will use UHD-1 to refer to 4K and UHD-2 to refer to 8K when it is not explicitly stated throughout the paper. Depending on discussion context, we will also refer to UHD-1 as HD 2.0 and UHD-2 as HD 3.0.

Since the advent of HD, the video and CE industries have gained a strong understanding of the relationship among resolution, screen size, and viewing distance. Video technology has continued to advance as the science of video perception and factors effecting Quality of Experience (QoE) are better understood. The science now spans the knowledge base created through the development of HD, as well as human vision biology and the neural processing that interprets the visual information sent to the brain. Basically, the science now incorporates the entirety of the Human Visual System (HVS).

Beginning with the fundamentals of the interplay among screen size, resolution, and viewing distance, Figure 1 captures the interaction in a straightforward way [23]. For a fixed screen size, higher resolutions are best viewed by sitting closer to allow for the full benefit of the increased detail on the display. For a fixed distance from the display, the benefit of higher format resolutions ensues with a larger screen size. These are well-understood principles.

As a simple example from Figure 1, if viewing a 50” screen from more than 20 ft away or greater, you will lose the benefit of HD at 720p and instead have an experience more akin to Standard Definition 480p. Sitting too close, such as 5 ft away on a 100” 1080p screen, threatens quality due to the distinguishing of pixels. This actually explains UHD-2 (8K HD) resolution. This format was originally envisioned as an immersive experience using a wider field of view (FOV), and to achieve that FOV by sitting closer demanded very high resolution to ensure high image quality.

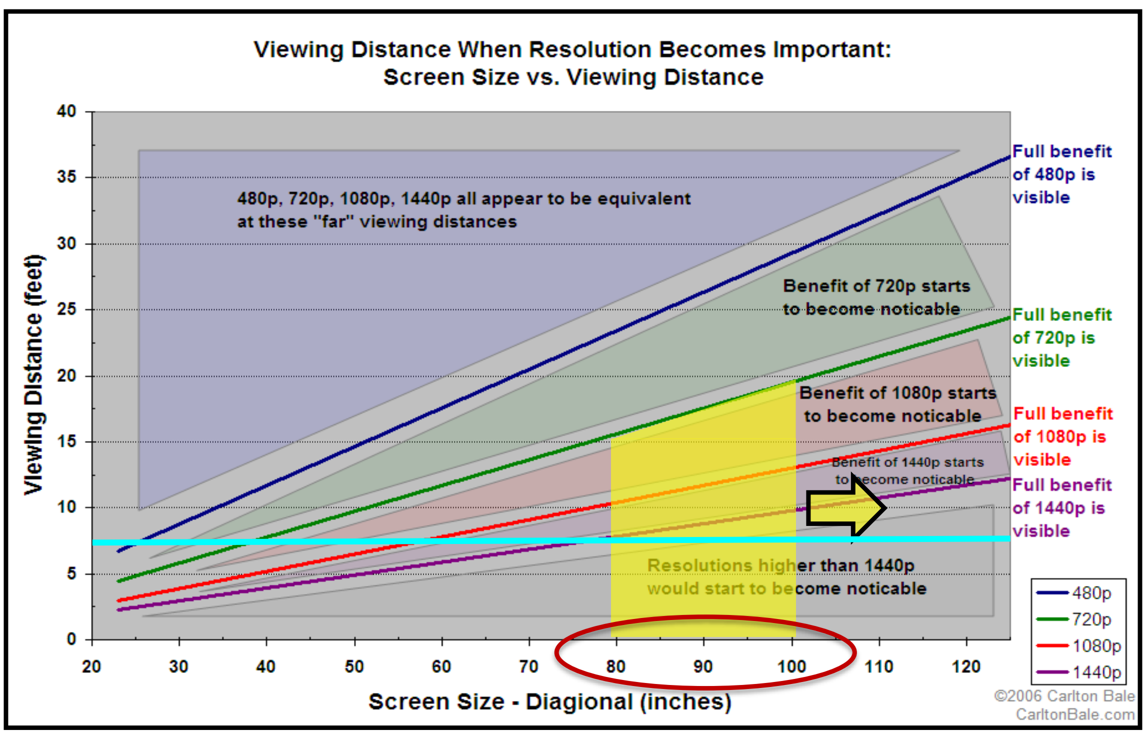

Figure 1 – Screen Size, Viewing Distance, and Spatial Resolution [23]

Consider the bottom right corner of Figure 1, shaded yellow. A typical viewing distance in the home today is about 7.5-9 feet. Flat panel screens are available now at ever-increasing sizes, such as those shown in this shaded yellow range. At 7.5 feet distance (light blue line), "only" a 55" screen could show perceptible benefits for resolutions better than 1080p (light blue line crosses red line). A 60" screen is sometimes considered the 4K TV threshold of benefit. Now 4K HD capable 84" flat panels are available (for a cool $25,000).

The Eyes Don't Have It

UHD recommendations are also correlated to visual acuity principles as they relate to the ability to resolve the image detail beyond simple optical acuity principles used to characterize the quality of our eyesight. This is an area where significantly better understanding has occurred in recent years. The common Snellen optics standard of 20/20 vision is a figure of merit based on visual acuity (VA) associated with accurately resolving high contrast, sharp edged objects (letters). However, it is well understood, for example, that the human visual system can perceive vernier or edge alignment at up to *ten times* the precision of standard VA, and such aspects of the brain's processing of video images translate to QoE.

Snellen acuity for "normal" vision works out to 30 cycles per degree. By comparison, as shown in Figure 2, HDTV satisfies a resolution consistent with 60 cycles per degree resolution, achieving "simple acuity," while UHD-1 achieves 120 cycles per degree, or the noticeable subjective threshold of "hyperacuity." This threshold accounts for the entire visual system's role in the viewing experience, beyond simple spatial resolution and vision receptor biology. It accounts for the role of neural processing. This experience achieves what's often referred to as "retinal" image quality [11].

Figure 2 – Tiers of Acuity Relate to the Quality of the Viewing Experience [11]

Simply put, treating the biology of the eye as a camera of a particular resolution and comparing it to image pixel density does not capture the essence of the brain's role in processing the information contained in the image, and which subjectively effecting our perception of its "realness." Indeed, studies to determine this relationship have concluded precisely this. An example of such a study is shown in Figure 3 [11].

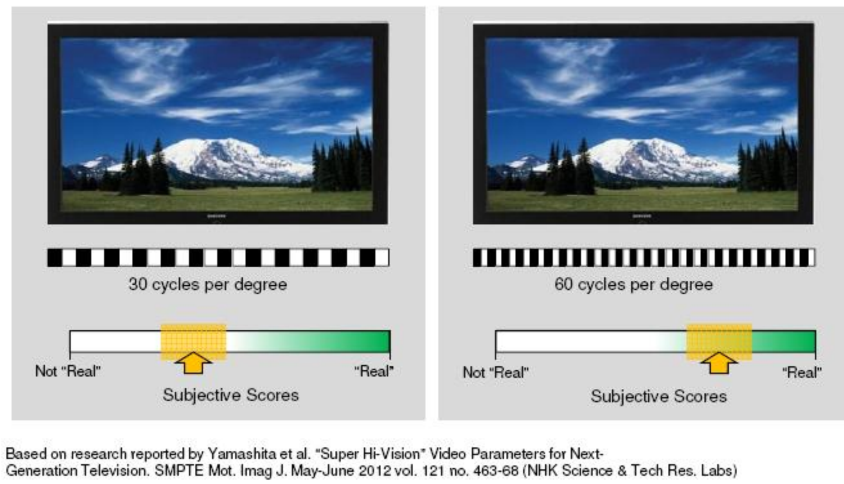

Figure 3 – "Normal" (20/20) Vision is Insufficient in Characterizing the Viewer's Subjective Video Experience [11, 18]

Studies also support the objective of a more *immersive* experience of a wide FOV (UDH-2). As shown in Figure 4, a wider FOV delivers the "being there" experience. To obtain this effect is a combination of a larger screen and close-up viewing, as in an IMAX theatre. And, of course, close-up viewing makes pixel density proportionally more important. According to ITU-R Report BT.2246: "UHDTV is … intended to provide viewers with … a wide field of view that virtually covers all of the human visual field."

The original UHD-2 recommendations by NHK are a careful balance of immersion and experience quality against the propensity for viewer discomfort as the brain tries to reconcile the imperfections of the artificial immersion environment.

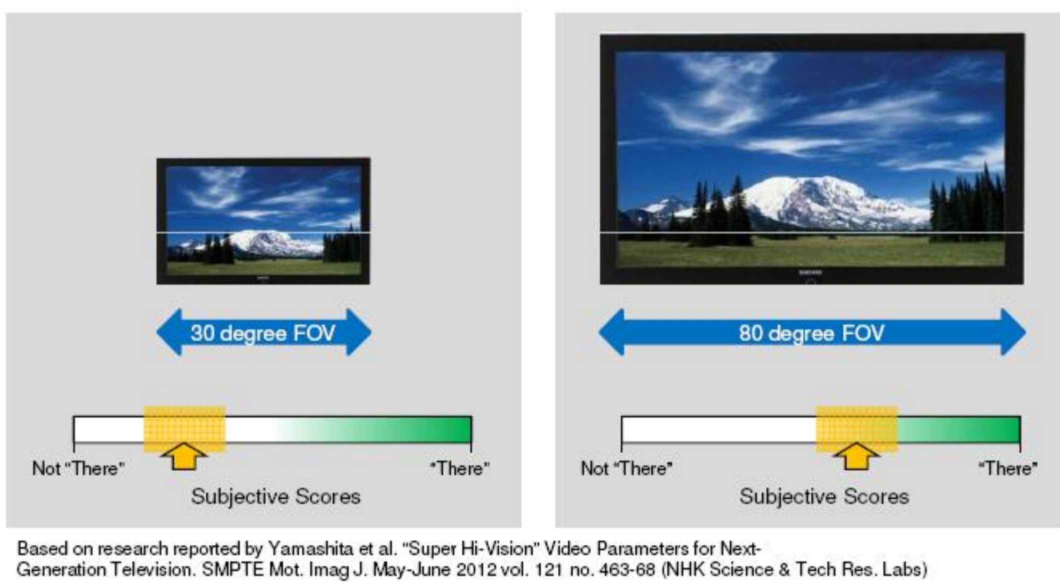

Figure 4 – Field-of-View and "Being There" Immersion [11, 18]

In addition to bit rates that will increase as pixel counts rise, a key component needed for quantifying screen sizes impacts for capacity management analysis is usage metrics. For UHD-2, the usage is constrained by its relationship to extremely large displays when compared to the inherent limitations of normal walls in a normal home. A 60” display is a five foot diagonal screen, and therefore a horizontal length of over 4 feet, while an 84” display (7 feet) has a horizontal length of over 6 feet! Though these display sizes (and projector systems) are suited to home theatres, they are not consistent with the living area viewing environment typical homes. They are very imposing companions in a normal living room. It is therefore deemed unlikely in our analysis that typical residential entertainment evolves in this way to take over the home.

In summary then, the long-term assumption that will be represented in subsequent capacity management analysis is that UHD-1 is a mass-market, scalable service for MSOs to introduce when the time is right. Data suggests this might be sooner than expected, with Credit Suisse projecting adoption by 135 broadcasters by 2017 [22]. We assume that UHD-2 is a service MSOs want to offer to the class of customers capable of enjoying it, but as an available IP program stream (or VOD title) only so as to avoid large spectrum penalties of broadcast.

Small Screens Pack a Big Punch

The HD 2.0/3.0 physics is not constrained to the primary screen. Consider Figure 5. Not only do larger primary screens suggest better spatial resolution, our secondary screens are now capable of high quality video such as HD. Tablets have changed the game for 2nd screen viewing, where a few hundred kilobits-per-second of low-quality video is no longer the norm for over-the-top (OTT) services. This changes how we quantify the impact of 2nd screen video devices receiving IP video around the home.

It is easy to see from Figure 5 that standard HD can be improved upon for reasonable viewing environments. For a 10” tablet, if the screen is about 17” away (airplane, back seat of a car), its spatial resolution can be perceptibly improved with a higher resolution format.

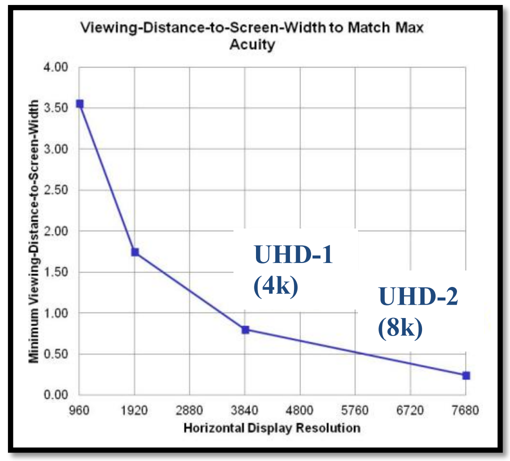

Figure 5 – Screen Size, Distance, and Resolution – Mobile Viewing

The case for UHD-1 based on the 10” tablet is difficult to make, but clearly screen sizes and portability have combined to change the paradigm of mediocre 2nd screen video.

The Need for Speed

The 30 Hz (interlaced), 50 Hz, and 60 Hz frame rates have origins in AC line rates, and thus are only loosely scientifically tied to video observation and testing. They simply were high enough to avoid the known issue of flicker. However, as spatial resolution has continued to improve, temporal resolution has not. Interlaced video itself is a nod to overcoming poor motion representation – exchanging spatial resolution for a higher rate of image repetition to better represent motion than a progressive scanning system of the same bandwidth.

Due to this legacy, HDTV today for progressive format is p60, or 60 fps. As displays become larger and of higher resolution and contrast, the challenges to effectively displaying motion increase because the edges to which movement is ascribed are sharper. Studies have shown that better scores for video QoE are given for higher frame rates for the type of video that intuitively would benefit most (high action, sports). Figure 6 shows such an example taken from ITU studies aimed at UHDTV recommendations.

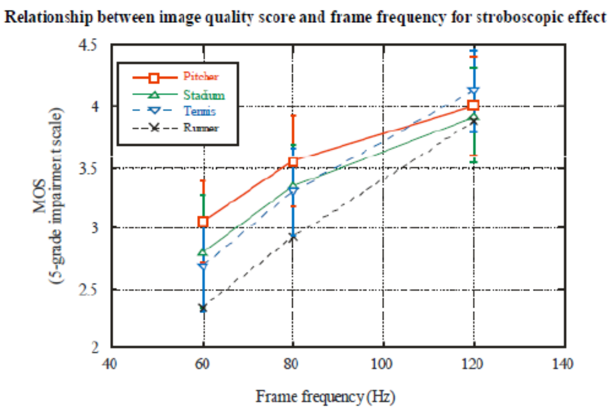

Figure 6 – Mean Opinion Score vs. Frame Rate and Video Type (Source: Report ITU-R BT.2246-1, 2012)

ITU recommendations for UHD include frame rates up to 120 fps. In addition to simply better representing high motion video, further ITU studies reveal that viewers are more susceptible to flicker for a wide Field-of-View (FOV). This suggests the need for higher frame rates for very large displays. The ITU judged 120 fps as necessary to minimize motion blur, stroboscopic effects, and perception of flicker for a wide FOV. Figure 7 from the same ITU report shows how the "flicker" frequency observable increases with a wider FOV.

Other studies have made similar conclusions about the perceptual benefits of increasing frame rates, including results as high as 300 fps [1].

We will consider 120 Hz as a component of the video evolution and quantify the effect. We will base use of this frame rate on the statistics of content type drawn from the most recent Motorola Media Barometer [24]. In this survey, 24% of viewing is sports, and another 24% is non-drama entertainment. We will consider these two categories as eligible for 120 fps. We thus assume that for UHD-1 viewing, eventually, 50% of it will be 120 fps. We will assume all UHD-2 viewing is 120 fps.

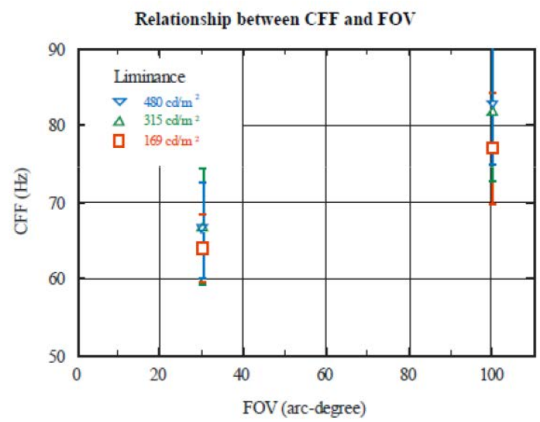

Figure 7 – Field of View vs. Critical Flicker Frequency (Source: Report ITU-R BT.2246-1, 2012)

ASSEMBLING LIABILITIES

Whereas 10-12 standard definition (SD) programs fit in a single 6 MHz QAM bandwidth, this number drops to 2-4 for today's HD, HD 1.0 – 720p and 1080i. Even at that, four is understood to be a nod to the trading off of video quality in favor of efficiency. Even worse, HD today represents a simulcast – programs delivered in HD are also transmitted in SD.

Moving beyond SD and HD1.0, UHD-1 (4K) works out to 4x the pixel count as 1080 HD, and UHD-2 (8K) works out to 16x the pixel count. These are major new Liabilities.

Relative to SD, we can summarize the pixel related bandwidth multipliers as:

1080i – 4x
1080p – 8x
UHD-1 (4K) – 32x
UHD-2 (8K) – 128x

Figure 8 captures the relationship amongst the formats. Note that the “Digital Cinema” 4K shown in green in Figure 8 is slightly wider than the format being discussed here (4096 vs. 3840), and thus the quotes around “UHD-1.”

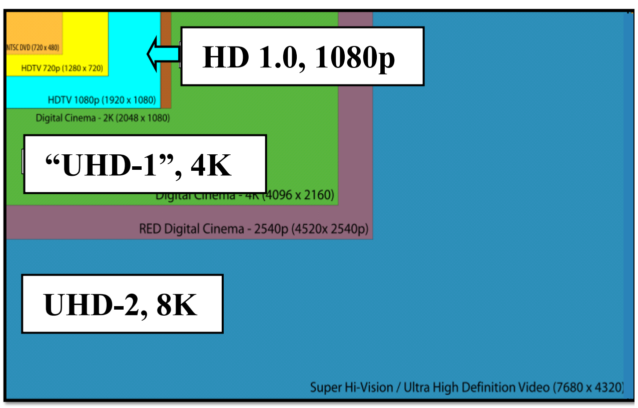

Figure 8 – Beyond HD 1.0: UHD-1 (4K) and UHD-2 (8K) [5]

The frame rate increase discussed is also an important Liability, scaling UHD-2 and, as described, about half of the UHD-1 programming.

There is one more Liability that we entertain. In UHD cases, there is the possibility of using higher bit depth (from 8-bit to 10-bit or 12-bit) color quantization, which also effects the bandwidth. In fact, the ITU Recommendation includes consideration of both 10-bit and 12-bit quantization depth. In the analysis, we apply 10-bit depth to UHD-1 only when high frame rate formats are used, and always for UHD-2.

All scale factors do not necessarily directly translate to bandwidth multipliers, but it is a useful first-order upper limit assumption. Our assembled *Liabilities* are therefore:

-1080p scaling from HD 1.0
-UHD-1 pixel scaling
-UHD-2 pixel scaling
-Frame Rate increase (some content)
-Quantization Depth increase (some content)

ASSEMBLING ASSETS

Encoding Efficiencies: H.264 & H.265

It has been 10 years since the Advanced Video Coding (AVC) [16, 18] international standard was completed in 2003. AVC – also known as H.264 and as MPEG-4 part 10 – and its equally successful predecessor, MPEG-2, are expected to continue to play an important role in the digital video economy for many more years, but they have been joined by the latest encoding standard -- High-Efficiency Video Coding (HEVC) [2, 3, 10, 14, 19, 20], or H.265.

The Final Draft International Standard (FDIS) for HEVC was ratified in January 2013. A deeper description of HEVC itself is given in [4]. Figure 9 captures the state of the set of core MPEG compression standards in the context of their lifecycle.

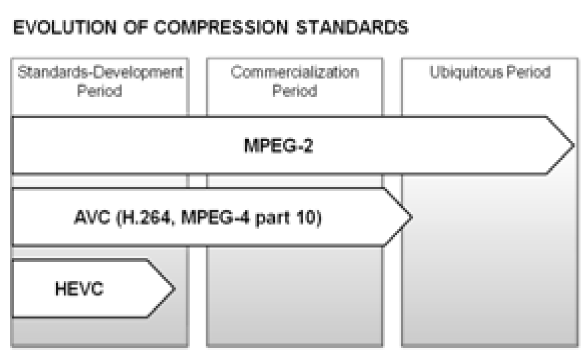

Figure 9 – State of Video Compression Standards [4]

H.265 roughly doubles the compression efficiency over its H.264 predecessor. H.264

itself doubled compression efficiency compared to MPEG-2. Thus, an 18 Mbps HD program using MPEG-2 would need only 4.5 Mbps using HEVC. Alternatively, UHD-1 can be delivered with HEVC at the same rate that HD over MPEG-2 is today. The savings represent key capacity management evolution Assets.

We will use 17 Mbps as our average data rate for a UHD-1 program at 60 fps using nominal bit depth (8 bits), based on recent internal studies of HEVC codec performance.

New Spectrum Considerations

Figure 10 illustrates the anticipated spectrum migration of the HFC architecture long-term. Because of many reasons outlined in [8] and [9], we foresee a phased approach to spectrum migration, consistent with the way operators incrementally deal with infrastructure changes in the context of dealing with legacy services and subscribers.

The end state of the spectrum migration is shown in the bottom illustration of Figure 10. This long-term end state (as an HFC-style architecture) maintains a level of asymmetry consistent with historically observed downstream/upstream traffic ratios.

Note that the objectives set out for DOCSIS 3.1 is to ensure at least 10 Gbps downstream and 1 Gbps upstream, and this phased evolution plan is meant to be aligned with that objective.

We evaluate spectrum evolution in our analysis of new capacity Assets, considering both the “Excess Bandwidth” case of 1.2 GHz and the extended bandwidth case of 1.6 GHz. We refer to a 1.2 GHz upgrade in subsequent tables as “Spectrum A” and 1.6 GHz as “Spectrum B.”

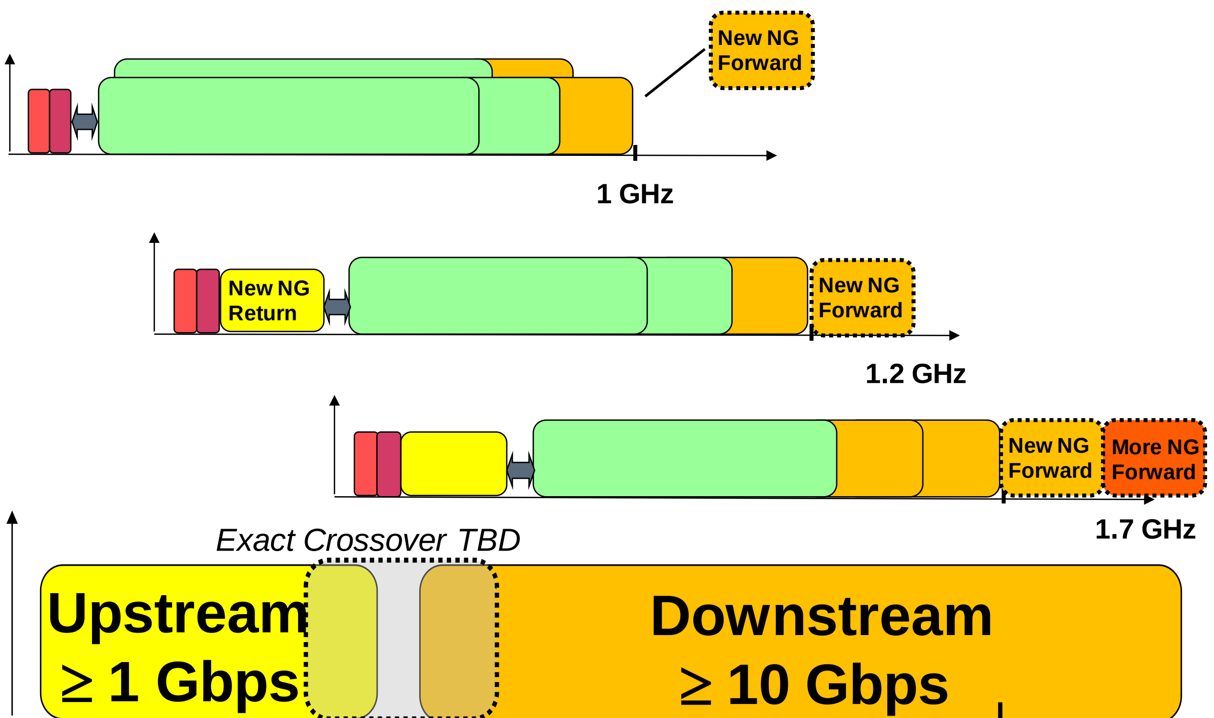

Figure 10 – Probable Evolution of the Cable Spectrum

THE dBALANCE SHEET – SIMPLE

In assembling our Liabilities, we quantified pixel count increases and accompanying video QoE parameters (frame rate and quantization depth) that translate into new bandwidth requirements. Going no further with media consumption bandwidth, these become our fundamental HFC capacity "Liabilities." What do we mean by going no further? We mean that we are excluding futuristic cases such as holograms or multi-dimensional displays. It seems reasonable to view these as beyond 15 years if they come to pass at all.

Also note that the analysis is limited to media consumption only, and not intended to address other potential services that may evolve to consume bandwidth. In our subsequent analysis, we *do* account for a separate Internet data service, but we do not attempt to quantify, for example, remote healthcare services, machine-to-machine activities, or other potential uses of the network that could affect bandwidth. We do not search for new "killer apps."

Assembling all of the Assets and Liabilities discussed above in Table 1 shows the Capacity Management Balance Sheet. At first glance, it appears threateningly out of balance from an operator perspective.

Table 1 – Capacity Management Balance Sheet, Simple Form

Assets	dB	Liabilities	dB
H.264	3.00	1080p60	3.00
H.265	3.00	UHD-1 (4k)	6.00
Spectrum A	2.43	UHD-2 (8k)	6.00
Total A	8.43	10-bit	0.97
Spectrum B	3.19	Frame Rate	3.00
Total B	9.19	**Total**	18.97

What does the table tell us exactly? By simply aggregating the columns, Table 1 suggests that if we take a snapshot of today's digital services that exist in the spectrum available (this example assumes a 750 MHz network – a value needed in order to determine new Spectrum Assets) and, fast forwarding to a future state where all are converted to UHD-2, High Frame Rate, enhanced bit depth, we'd have a (very) negative balance and thus not enough capacity to do so. It would be a great sign if the conclusion of such tabulation left us with a positive balance. But, since it does not, we ask: is this an evolution example that we really care about? In practice, the picture is not really this ugly – we just need to delve beyond this oversimplified arithmetic and analyze the problem a little deeper.

First, let's recognize that the more likely expectation, the long-term mass-market evolution assumption previously stated: UHD-1 (4K) takes hold as the significant wide penetration video service, with UHD-2 (8K) being an available format in a selective on-demand or VOD unicast fashion. We certainly do not envision a video evolution that takes all of the current programming line-up on the wire and converts it to UHD-2 to be put on the wire, which is what direct application of Balance Sheet arithmetic simulates.

Now consider Table 2. It quantifies the Balance Sheet outcome for various permutations of spectrum and service rather than this unrealistic one. Under the UHD-1 only scenario, as shown in Table 2, we identify a sliver of hope in the upper right-most cell of the table. In this case, assets just eclipse liabilities, the preference for capacity planning (clearly our version of a Balance Sheet would trouble actual Accountants!). However, we have yet to try and accommodate a portion of UHD-1 programs at higher video quality (VQ) such as frame rate or quantization depth. We can project that to do so would not be possible in any significant degree from a capacity perspective.

Table 2 – Achieving UHD-1 Balance

	Fixed Spectrum	Spectrum A	Spectrum B
UHD-1 Only	-3.00	-0.32	0.49
UHD-2 Only	-9.00	-6.32	-5.51
UHD-1 HQ	-6.97	-4.29	-3.48
UHD-2 HQ	-12.97	-10.29	-9.48

Diving deeper still, note that the current spectrum usage is actually a simulcast mix of SD and HD programs. The calculation in Table 3 converts these all to a single format. In other words, two of the same format of the same program is what the calculations in Table 2 show. However, the lower format simulcast represents a smaller percentage of the spectrum to begin with so the impact is limited. Significantly, however, as shown in Table 3, it does improve the outcome, as the assets now exceed the liabilities in the case of a 1.2 GHz spectrum (second yellow cell) expansion instead of only the 1.7 GHz case. Progress!

Table 3 – Correcting for Simulcast

	Fixed Spectrum	Spectrum A	Spectrum B
UHD-1 Only	-1.90	0.78	1.59
UHD-2 Only	-7.90	-5.22	-4.41
UHD-1 HQ	-5.86	-3.19	-2.38
UHD-2 HQ	-11.86	-9.19	-8.38

We can further estimate that any system that has at least 60 analogs today (750 MHz network) that is within 3.25 dB of balance will also be sufficient if the analog simulcast is removed and donated to the digital video pool. This additional set of successfully balanced options is circled in red on Table 3. Though there is little margin to be overly comfortable between Assets and Liabilities in virtually all cases, that we can emerge in positive territory is encouraging. It is not surprising that the extremely powerful use of analog reclamation shows significant Balance Sheet benefits. More Progress!

Encouraging as things are looking, and as satisfyingly simple as Balance Sheet tabulation can be, Table 3 still does not tell a complete evolutionary story. We do not anticipate that the approach for even HD 2.0 would be to implement every legacy digital program in the lineup today as a UHD-1 program, either broadcast or on SDV. It would be part of the broader IP transition – from a timing standpoint as well as from an efficiency of spectrum standpoint.

It is also insufficient to focus on final outcomes of service evolution only. The transition plan, which accommodates simulcast and interim architectural phases are critical for capacity management. A more comprehensive analysis approach will be required.

The Intersection of Video Services and IP Traffic Growth

The downstream rate of traffic growth has been chugging along at about 50% per year. The Balance Sheets approach of video service conversion does not account for this dynamic – it basically assumes that while video evolves the remaining spectrum allocation remains static.

Of course, historically distinct video and IP services are not staying that way. Indeed, it is one of the evolutions of video services as we know it – in the form of over-the-top (OTT) services – that is currently the engine of this 50% year-over-year Compound Annual Growth Rate (CAGR). In addition, MSOs are currently migrating their own video services to IP. Some have already begun the process. The IP Transformation is expected to take an extended period of time, as revenue-producing legacy services will not quickly be removed. As such, while tabulating Assets and Liabilities is very insightful, it is important for investment planning purposes to view an evolution timeline that quantifies all of the moving parts of video *and* data. The analysis approach is more complicated, but is still entirely tractable.

An example of the conceptual timeline approach was described in great detail in [4] and [8], and a sample transition analysis timeline from [4] is shown in Figure 11. We will summarize the process here as analysis later in this paper will use a similar methodology. Please refer to [4] for a complete description of this case and [11] for a walk-through description of how the analysis approach works.

As a brief context, Figure 11 charts the 50% CAGR number commonly used for downstream traffic growth (yellow trajectory) against various thresholds of video service evolution and architecture evolution through the year 2030. An inherent assumption of the 50% CAGR *in this version* of analysis is that the 50% CAGR includes the MSO IP Video transition – it is not separately accounted for. Another way to look at this is that the engine of growth driving 50% CAGR has been OTT video services (i.e. Netflix) and the shift from OTT penetration and growth to MSO introduced services is an invisible shift of eyeballs relative to CAGR.

The trajectory is broken up into four segments with three discontinuities that represent three service group splits – by half, by half, and then down to N+0 (approximately one-third at that point).

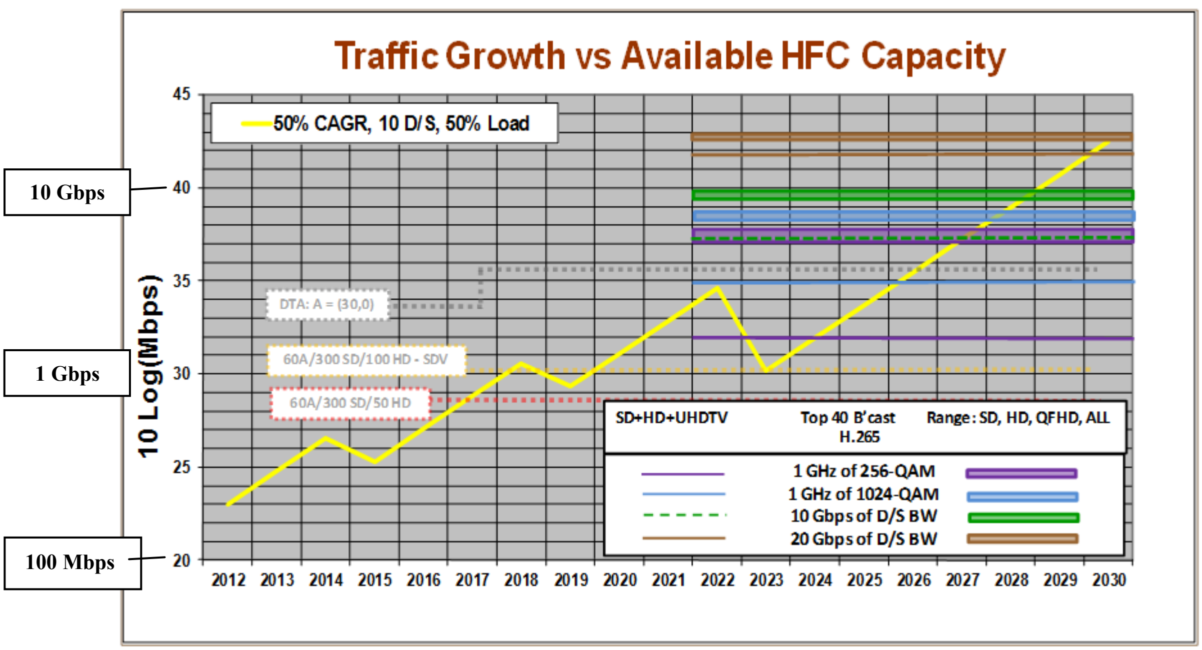

Figure 11 –HD 2.0 & 3.0, IP Traffic Growth and HFC Capacity Limitations [4]

Three labeled horizontal thresholds are shown on the left side with specific video service and implementation assumptions. For example, the red threshold bar is for a system with 60 analog video carriers, 300 SD, and 50 HD services in an 870 MHz system. If these services held steady, then the growth of IP traffic would breach the threshold in 2017, assuming one node split occurs in the intervening time period. The two horizontal thresholds above this represent the addition of HD programming (to 100) while also deploying SDV to manage the bandwidth growth, and the use of analog reclamation in two phases: 60 to 30 and then in 2017 the removal of all analog.

The remaining thresholds in the upper right are described in the legend at the bottom right. In all cases, there is an assumption of a static IP broadcast of the Top 40 channels, based on analysis and reasoning we shall describe in a subsequent section. Four cases of video format combinations are analyzed with different network architecture assumptions. See [4].

A key conclusion from [4] around Figure 11 was that the UHD-2 (8K) as a format with a broadcast component had major implications without significant new capacity exploitation. Of course, as previously stated, we do not envision UHD-2 as such a service. By contrast, for UHD-1 (4K) scenarios, even the least capable case (1 GHz of 256-QAM) had an extended lifespan. This bodes well for the ability of tools already available to manage through an aggressive combination of video service evolution and persistent CAGR of IP traffic. Figure 11 makes it abundantly clear that for a persistently high CAGR over a long period of time, CAGR eventually wins, a fact we revisit and reevaluate in the next section.

We will use a similar tool and analysis approach in the calculations ahead.

ASYMPTOTIC ALL-IP TRANSFORMATION

In the Figure 11 analysis, we approached the problem of CAGR and video services as largely orthogonal services sharing common spectrum to be managed. We accounted for the video services in the analysis approach in [4], after conversion to IP, as a spectrum block that was therefore unavailable to be used for CAGR associated with the HSD service growth.

In this paper, we instead consider the IP video and data services as a composite to reconcile capacity constraints and CAGR over a 15-year transition period. A core reason for this is based on the common assumption that the engine of 50% CAGR in the last several years has been video services – albeit over-the-top (OTT) video services. MSO managed IP Video services are on deck, and it is a central premise of the analysis approach taken here that streaming video has been and will continue to drive CAGR. This leads to a foundational premise of the analysis going forward:

With our knowledge of HD2.0 and HD3.0 visual acuity relationships and technology availability, and the assumption of video-driven CAGR, we can set some realistic boundaries on asymptotic capacity requirements for a given serving group aggregate. Armed with an asymptote, we can project long term CAGR from the media consumption standpoint that tapers to this projected asymptote over the transition duration.

Our approach, therefore, is an analysis aligned with the principles of [5], where theoretical limits of video representation were piled one on top of the other to develop a worst case media consumption asymptote per household. Here, we remove the theoretical aspects and consider more practical and technological aspects of HD 2.0/3.0 over a workable business planning horizon. The objective we execute on will now be to find these boundaries and develop a migration approach aligned with what we discover, using our Asset and Liability line items to guide a transition plan.

For the IP Video world, there are additional implications and valuable Assets to assess that are associated with managing streams instead of "channels," which we now describe.

Trends Effecting IP Video

The nature of video traffic being delivered is changing. There are many variables in play, virtually all of which are driving towards increasing unicast, and perhaps acting as coal for the CAGR engine:

- More content choice
- Time-shifting
- Trick play expectations
- Network DVR (nDVR),
- Video capable IP device proliferation (tablets and smartphones)
- Shrinking service groups

One important result of these shifts is that gains typically afforded by multicast capability or bandwidth reclamation gains associated commonly with SDV architectures begin to erode.

The benefits (or not) of multicast bandwidth savings can be determined quantifiably. A unique tool has been developed that takes into account the traffic, device, and format variables, as well as the known viewing behaviors of service group aggregates based on the history of SDV and IPTV networks. It quantifies these behaviors in the calculation engine to predict bandwidth and channel usage requirements. The tool is freely available for use on the web at http://www.motorola.com/Multicast-Unicast-Calculator/.

We will use this tool to determine to IP Video bandwidth requirements for services as we know them today – DOCSIS 3.0 with SD and HD 1.0 services. And, because the calculation engine is agnostic to formats specifically in favor of defined bit rates, we can use it with HD 2.0 and HD 3.0 inputs along with projected assumptions about the relative viewing behaviors of future simulcast services to determine IP Video bandwidth requirements in the UHDTV era.

Lastly, we can adapt the outputs to determine DOCSIS 3.1 requirements, adjusting for the increased spectral efficiency D3.1 entails.

IP Video Traffic Multiplexing Efficiencies

Legacy architectures are based on simple traffic management techniques that allot and enforce an average of 3.75 Mbps per SD video stream in to fit 10 streams (at least) per QAM carrier.

The introduction of DOCSIS 3.0 adds channel bonding to the toolkit. The net effect of bonding coupled with more streams/QAM of H.264 or H.265 is the ability to use the law of large numbers, reducing average bandwidth. Many independent streams competing for much more pipe capacity result in a self-averaging effect [6]. Smoothing out peaks and valleys are handled inherently by statistics operating with a capped VBR scheme, and most likely with Adaptive Streaming to ensure QoE. An example multiplex of independent peak-to-average video waveforms used in simulations to quantify this effect is shown in Figure 12.

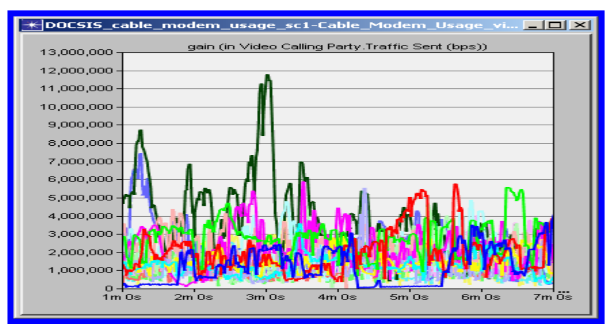

Figure 12 – Channel-Bonded VBR is a BW Efficiency Asset [6]

Based on simulations and observations, we use an 80% scaling as the bandwidth required for VBR-based channel bonded DOCSIS video in comparison to single carrier QAM transport.

Fiber Deep Migration

"Business as Usual" HFC migration has been shown to be well-suited to supporting a lifespan of at least a decade of legacy video evolution and aggressive IP data traffic growth [8]. The use of node splitting in the HFC architecture reaches its ultimate phase when the last active becomes a fiber optic node. This architecture goes by various names – Passive Coax, Fiber-to-the-Last-Active (FTLA), or N+0. Figure 13 illustrates this classic multi-phased operator migration strategy for segmenting a serving area.

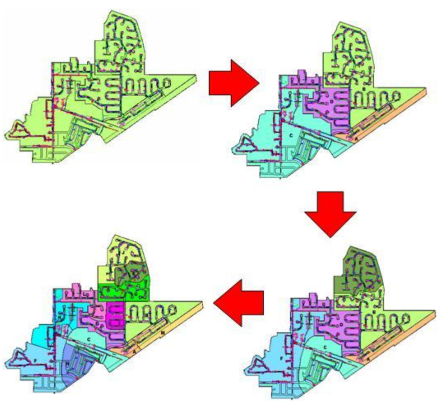

Figure 13 - Fiber Deep "Business as Usual" Migration is an Average BW/hp Asset

Regardless of the name, the architectural implications for N+0 have three core components:

1) Very small serving groups (40 hhp assumed)
2) The opportunity to exploit new coaxial bandwidth with no actives after the fiber optic node
3) A higher performing (higher SNR) HFC channel.

The latter two are both targeted by DOCSIS 3.1 – the first to enable the 10 Gbps worth of downstream capacity, the second to make use of the most bandwidth efficient modulation profiles (4096-QAM, possibly higher). This will offer the best opportunity to achieve 10 Gbps or more. Figure 14 compares today's 256-QAM format to the 4096-QAM format that clearly offers more bits per symbol, as envisioned in DOCSIS 3.1 for future bandwidth efficiency. They are shown at equivalent BERs of 1e-8 uncoded, which corresponds roughly to SNRs of 34 dB and 46 dB, respectively.

The N+0 architecture is essential for a growing forward spectrum to 1.6 GHz. Plant RF actives are unlikely to be stretched to 1.6 GHz, as they may be for the 1.2 GHz extension. N+0 will also leave operators within a stone's throw of FTTP.

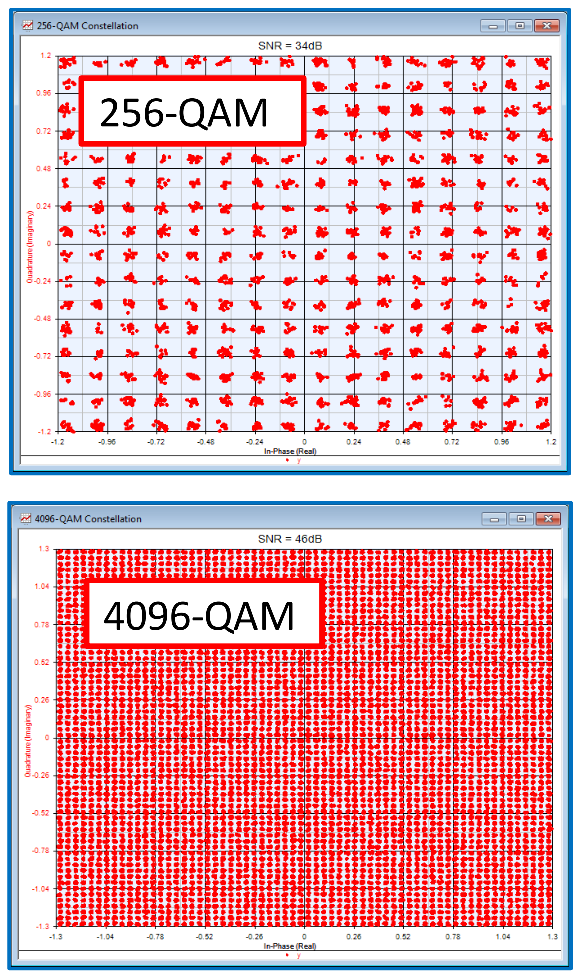

Figure 14 – DOCSIS 3.1 Modulation Formats as a BW Efficiency Asset

BEGIN WITH THE END IN MIND

An idea of where you want to be, to the best it can be known, is how to approach optimizing the path from A to B. In this case, "B" means an all-IP end state, and some projected mix of video services. Furthermore, we consider the end state as something accomplished over a 15-year transition period.

Importantly, in the all-IP case, we move away from channel-thinking and towards thinking in terms of streams. This is even more so the case considering DOCSIS 3.1. DOCSIS 3.1 will be a wideband OFDM system which literally removes the idea of channels in the conventional 6 MHz and 8 MHz sense. While DOCSIS 3.0 allows us to mathematically quantify the impact of wideband channels by allowing channel bonding, DOCSIS 3.1 remakes the physical layer QAM slots themselves in favor of OFDM subcarriers (also carrying QAM) able to be implemented as a single block over a very wide bandwidth.

What must be determined is "simply" how many streams and of what type to project as a 15-year assumption, from which the aggregate bits-per-second can be determined for the service group size envisioned. This is exactly what the modeling tool described above has an engine to calculate. However, the tool is for today's practical scenarios. This means a result based on DOCSIS 3.0 and input stimuli built around SD, 2nd screen viewing of varying bit rates, and HD 1.0. A snapshot of some of the parameters entered and used by the tool is shown in Table 4.

As noted previously, the stimuli used to drive the engine are merely names – any numbers can be put in that represent a scenario of interest, so long as the Mbps used in each category are aligned with the programming definition and usage defined for each.

Table 4 – IP Video Modeling Tool

STEP 1: DEFINE SERVICE GROUP SIZES, DOCSIS OVERHEADS, DOCSIS AND IP SERVICE PENETRATION	
Service Group Size HHP:	125
Plus the following SG sizes:	62.5 and 31.25
DOCSIS Packet and Header Overhead:	5
DOCSIS Channel Bonding Overhead:	1
Overall DOCSIS Penetration:	70
Penetration Range of IPTV:	0.1 to 100 of DOCSIS Households
This Translates to an IPTV Penetration Range of:	0.07 to 70 of All Household Passed
Percentage of IPTV Customers Actively Viewing at Peak:	70% (50% is typical for cable systems as shown by historical SDV data)
Average Number of IP Devices per Home Active at Peak Viewing:	3.7

STEP 2: DEFINE THE TYPES OF VIDEO PROGRAMMING TO BE DELIVERED, THE REQUIRED DATA RATES AND THE VIEWING MIX BY CHANNEL TYPE	
Required Data Rate by Program Type	**Mb/sec**
SD Bandwidth per Program MPEG4:	5
720p HD Bandwidth per Program MPEG4:	30
HD1080p/3DHD Bandwidth per Program MPEG4:	170
Medium Speed IPTV Bandwidth per Program MPEG4:	0
Low Speed IPTV Bandwidth per Program MPEG4:	0

We drive these inputs with the following video bit rate numbers and stream distribution assumptions, and also under an assumption of all HEVC at the end of the 15-year evolution:

For the 4K format, we will project that half are enhanced with 10-bit color depth and high frame rate (120 fps) with no additional encoding benefit assumed. More frames generally suggests less difference between frames and therefore potentially more coding gain, but the programming targeted is precisely the action-type video that is less likely to have that characteristic, or at least not to the same degree as a drama or news program. As we discussed previously, market analysis suggests this represents roughly 50% of viewing, so we account for this by having the enhanced UHD-1 as 50% of the total.

We project all of UHD-2 as 10-bit and higher frame rate, and assume this is a small percentage (10%) of pbh viewing associated with home theatre-type users. It is likely that the need for this service at all will be very serving-area sensitive (i.e. higher income suburban single-family home neighborhoods). Table 5 summarizes the values used as the viewing end-state parameters modeled.

Table 5 – End State (15-yr) Viewing Behaviors and Bit Rates

Format	Average Bit Rate	% Viewership
HD1.0	5 Mbps	20%
UHD-1	17 Mbps (50%)	70%
UHD-1 Enhanced	42.5 Mbps (50%)	
UHD-2	170 Mbps	10%

Two cases were examined for service group end states – operators who migrate to N+0 and others currently planning to be in the 100-150 hp "sweet spot" [25] which correlates to two more segmentations of the service group, whether virtual or actual fiber deeper. This leaves an "N+Small" HFC architecture in place, for example, such as an N+(1-3) cascade.

Other key assumptions include a 70% penetration (modest growth over the course of 15 years), a 70% peak-busy-hour (pbh) usage (aggressive), 1.5 streams per user per household, and the users per household governed by demographics associated with data extracted from recent 2010 census data, shown in Figure 15.

While we neglect mathematically household greater than 5 (all above 5 are treated as 5), we also neglect that about half of the households with children are those with children of an age unlikely to be independent viewers of multiple screens. There are other deeper weeds of demographic detail we could include such as this, but this model seems sufficient for better-than -ballpark estimates.

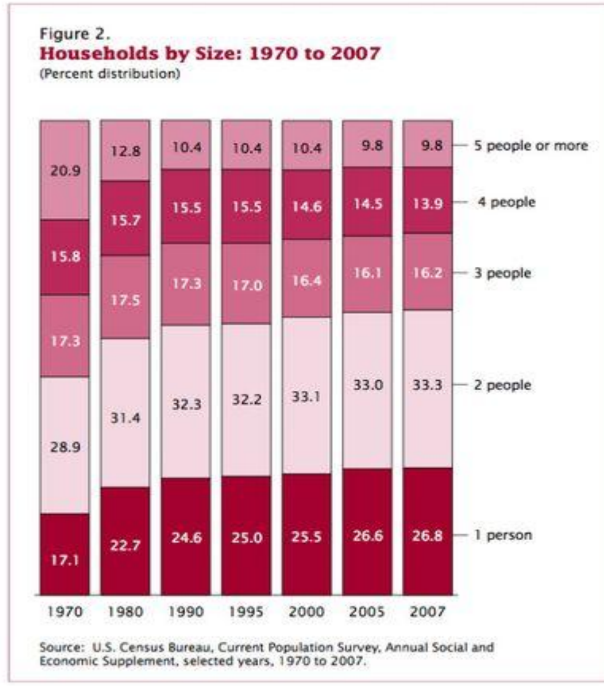

Figure 15 – Household Sizes to Govern Stream Counts [5]

Based on the above inputs for bit rates, viewership distribution, and service group size, we use a modified form of the model shown in Table 4 with HD 2.0 and HD 3.0 bit rates and usage metrics from Table 5. And, we recognize that the output is calculated in DOCSIS 3.0 channels, but which we can easily convert to Mbps or Gbps based on inputs for bit rate/channel of 256QAM and overhead losses. The two modeled cases and their results are shown in Tables 6 and 7.

Table 6 – End State (15-yr) Service Group Size of 125 HHP or "N + small"

Total DOCSIS Capacity Required to Deliver Managed IP Video Services

Total Bandwidth Requirements - Case 1

Service Group Size HHP:	125	Peak IP Video Devices / Home:	3.7 Active
Max Video Penetration:	70%	At Peak Penetration:	87.5 Homes
Percentage of Video Devices Active at Peak:	70%	At Peak Viewership:	227 Streams

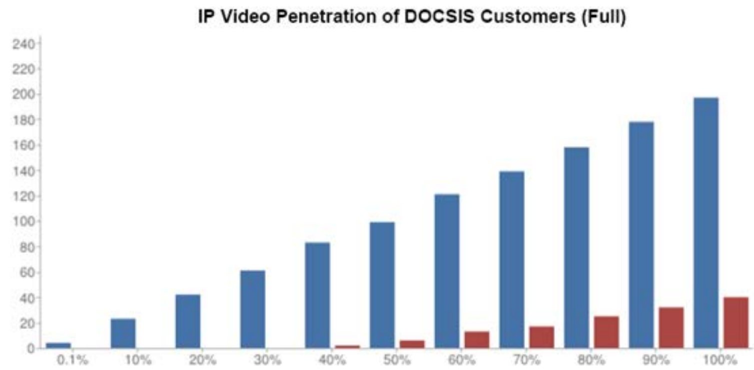

IP Video Penetration of DOCSIS Customers Charts

Total Downstream Channels @ 100% Unicast

Potential Channel Savings with Multicast

Table 7 – End State (15-yr) Service Groups Size of 40 HHP @ N+0

Total DOCSIS Capacity Required to Deliver Managed IP Video Services

Total Bandwidth Requirements - Case 1

Service Group Size HHP:	40	Peak IP Video Devices / Home:	3.7 Active
Max Video Penetration:	70%	At Peak Penetration:	28 Homes
Percentage of Video Devices Active at Peak:	70%	At Peak Viewership:	73 Streams

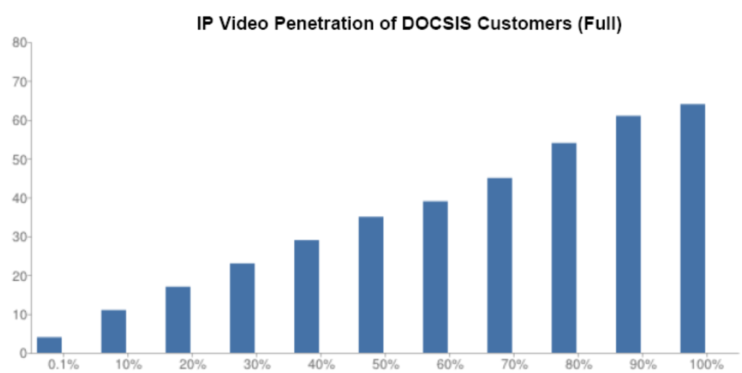

IP Video Penetration of DOCSIS Customers Charts

Total Downstream Channels @ 100% Unicast

Potential Channel Savings with Multicast

The analysis of the above two cases results in the following:

125 HHP (N+Small)

- 197 DOCSIS 3.0 Channels (~8.5 Gbps, overhead included)
- Potential savings of 40 channels available with multicast:
  - 157 D3.0 Channels (~6.7 Gbps)

40 HHP (N+0)

- 64 DOCSIS 3.0 Channels (~2.7 Gbps)
- No savings from multicast capability

Note from Table 7 that the serving group size combined with the unicast expectations and programming breadth has eliminated any multicast savings for the latter case.

Broadcast, Multicast, or Unicast?

One high level architecture result of this in-depth modeling of IP Video, and alluded to in the N+0 case above, is that because of the unicast trends, multicast gains may be limited and gradually erode. Because of this limited bandwidth benefit over time, and the complications brought about by multicast in the architecture (Wi-Fi, IP devices in the home, Adaptive Bit Rates), it may instead be simpler and nearly as effective to deploy a combined broadcast plus unicast architecture, since the vast majority of multicast gain is limited to the most popular programming. And, a broadcast component satisfies the "Superbowl" problem and even multi-channel major event scenarios (breaking major news story).

Analysis indicates that 80-90% of the multicast gain is obtained in the most popular 20 programs and fewer of course during "major event" scenario. This architectural concept is shown in Figure 16. We take a conservative approach in the analysis and will examine the case of 30 total broadcast programs, combined with the remaining all-unicast traffic as calculated by the model. We do not subtract any "broadcast" in the 125 HHP case where there is some available multicast gain, but will identify and compare the multicast example on the Capacity Management Timeline analysis. A detailed analysis of IP Video multicast and unicast architectures and implications is contained in [15].

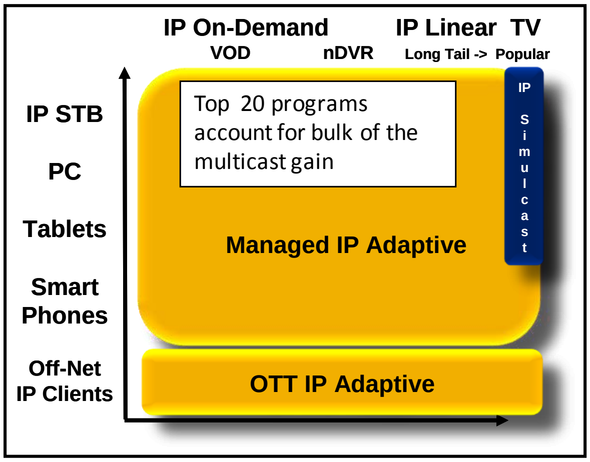

Figure 16 – Optimizing IP Video Delivery

Accounting for the IP Broadcast means allocating bandwidth for a simulcast of 30 IP channels in p60 HD and UHD-1 (4K), by assumption of desired broadcast service format mix in 15 years. For 4K content, we continue to assume half of the content benefits from enhanced quality (bit depth and high frame rate). Table 5 identifies the expected average bit rates this entails, and when it is aggregated over wideband IP channel it sums as shown in Table 8.

Table 8 – IPV Broadcast (30 Programs)

Format	Avg Bit Rate	Total
HD1.0	5	150.0
UHD-1	17	892.5
UHD-1 Enhanced	42.5	
Total	IPV Eff @ 80%	834.0

Under an assumption of *10 bps/Hz net* (payload) spectral efficiency we therefore allocate *85 MHz* for this broadcast spectrum. This efficiency is based on recent analysis done as part of the Channel Model ad-hoc in IEEE 802.3bn [26]. Analysis there has shown that over 90% of today's DOCSIS CPE report SNRs capable of 2048-QAM today, using current LDPC FEC technology to set QAM thresholds, as shown in Figure 17. 2048-QAM achieves a raw efficiency of 11 bits per symbol, and after allowing for overhead losses, a net efficiency of approximately 10 bps/Hz is expected.

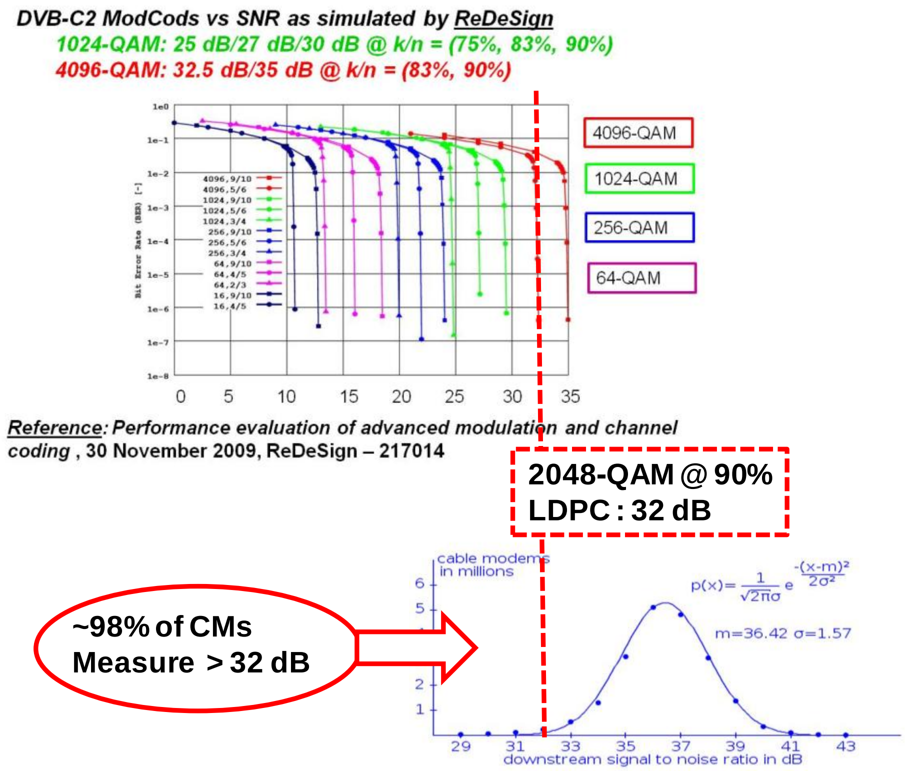

Figure 17 – Fielded CPE Report Much Higher Capacity Potential than Currently Implemented, with 90% Being 2048-QAM Capable [26]

Clearly, we will only march closer to capacity with FEC, and next generation CPE will be no worse in terms of sensitivity, fidelity, and implementation loss as today's – and should be significantly better.

Note that this implementation – broadcast + unicast – yields a total bits-per-second aggregate that represents a virtually non-blocking (Pb = .01%) unicast asymptote assuming our statistical basis of viewing.

<u>Let There be Data?</u>

Streaming video is the engine of CAGR growth in today's downstream, taking over for increased penetration of web browsing of sophisticated multimedia websites. This will remain the case with more bits required for HD 2.0 and the unicast driving trends previously described. However, in the analysis, we also set aside a complementary browsing experience. For this, we assume a

1 Gbps service with 1% concurrency (100:1 oversubscription) as a sufficient complement to the IP video service.

Given penetrations and serving group sizes, this works out to:

125 HHP: 875 Mbps (970 Mbps w/OH)
40 HHP**:** (N+0): 280 Mbps (310 Mbps)

We now have the components of our "End in Mind" requirements. We will analyze both by Balance Sheet and by the Capacity Management Timeline per the Figure 11 approach.

THE dBALANCE SHEET - REVISITED

We have now identified several new Assets and Liabilities in the prior sections. Let's add these to the Balance Sheet and assess their meaning. This is shown in Table 9 below.

Table 9 – Balance Sheet, New Entries

Assets	dB	Liabilities	dB
H.264	3.00	1080p60	3.00
H.265	3.00	UHD-1 (4k)	6.00
DOCSIS 3.1	1.76	UHD-2 (8k)	6.00
IPV Efficiency	0.97	10-bit	0.97
Split	3.00	Frame Rate	3.00
N+0	7.96	B'Cast	0.56
Spectrum A	2.43	**Total**	**18.97**
Total A	**22.12**	B'Cast A	0.33
Spectrum B	3.83	B'Cast B	0.24
Total B	**23.52**		

Here, "Broadcast X" represents the loss due to the choice to allocate 85 MHz of spectrum for 30 programs of IP Broadcast. Since this is an absolute set-aside, in this case an absolute network bandwidth was required. We again chose the 750 MHz downstream which incurs the most relative loss in the "Fixed Spectrum" case.

Repeating our matrix of possibilities, we can see in Table 10 that there are now *no deficit-only conditions*.

Table 10 – Capacity Balance, All Assets and Liabilities

	Fixed Spectrum	Spectrum A	Spectrum B
UHD-1 Only	10.13	12.79	14.28
UHD-2 Only	4.13	6.79	8.28
UHD-1 HQ	6.16	8.82	10.31
UHD-2 HQ	0.16	2.82	4.31

No deficit scenarios is an encouraging outcome. The massive benefit of service group splitting (nearly 11 dB total in the Asset column) is making a big difference. Of course, the use of all of the Assets listed in a capacity management calculation assumes a full transition to IP Video, where video streams and the associated traffic engineering of them allows the consideration of Asset parameters associated with service group sizing. Looked at another way, there is no bandwidth benefit to node splitting – to any serving group size – for broadcast spectrum. Segmentation of the network only has value in a switched architecture – whether that is classic legacy cable SDV or the IP Video case we are focusing on here.

Since IP Video *is* the plan for most MSOs long-term, *and* it is the end-game assumption we are working with here, the conclusions to draw from Table 10 should indeed be viewed as positive indicators for the long-term outlook. Table 10 suggests that there should be no scenarios where N+0 is an insufficient end-state solution. However, removing the N+0 (4.96 dB – replace it with a normal split) and perhaps some of the spectrum expansions (2-4 dB) would result in some negative balances in Table 10, suggesting capacity limitation for some of the 125 HHP scenarios.

Fully examining Tables 9 and 10, we find that this is so only for the cases where conversion to UHD-2 is the service assumption objective. As has been discussed, UHD-2 is viewed for this analysis as an available unicast streaming format to a small anticipated percentage of subscribers, and for which we have traffic engineered unicast IPV to support it based on Tables 6 and 7.

Encouraged by the long-term, the complete task then involves surviving the transition period of this capacity management challenge. Again, old services do not immediately die when new services are introduced. Likewise, new services rarely are introduced through massive service change all at once. However, allocating resources for them generally comes at service introduction. An IP approach helps to manage this resource allocation proportional to the penetration.

With the transition as the key challenge, we will walk through a timeline of potential service evolution-architectural migration scenario aligned with reasonable timings of each. First, however, we put our end state IP Video calculations to use and discuss their context and role in analyzing the transition process.

CAPACITY MANAGEMENT TIMELINE

With the above accounting for IP Video downstream bits and services, and our Asset and Liability line items to play with, we have all of the information we need to evaluate and develop a comprehensive Capacity Management Timeline for evolution planning. An important modification to the Figure 11 approach is used to bound problem by considering CAGR from a perspective other than blind allegiance to persistently aggressive growth.

Asymptotes of Behavior and Biology

There are *three key principles* to this perspective and implementation on a Capacity Management Timeline:

1) Recognition that detailed residential demographics are available and are more useful as metrics for an IP streaming world and multiple IP devices per home video connectivity than homes passed.

2) Recognition that humans have a limited ability to multi-task, in particular with video. While secondary screens as simultaneous and background content playing during a primary viewing experience may be common, humans have a limited ability to focus on multiple things at once with comprehension. Because of this, we cap simultaneous streams per home at $\leq$ 2 per individual (1.5 was used).

A counter argument to (2) is the "pub-style" home environment, with TVs just on 24/7 in rooms throughout a home, relatively independent of occupancy. The analysis bets against this, with "green" objectives perhaps a factor in this style of viewership evolution.

3) From the standpoint of *purely media consumption driven bandwidth*, or, alternatively aggregate bandwidth strongly dominated by media consumption, the IP Video tool output in the prior section coupled with the HSD service assumption calculation represents a *projected bandwidth growth asymptote*.

It is of course risky to suggest that bandwidth may stop growing, lest future work can look back and snicker about the naiveté of such an estimate – similar to early predications about the necessary memory requirements for a PC. Prior analysis [4, 5] considers assumptions otherwise, so consider this one in a series of analysis to weigh the possibilities and make a judgment. Of course, observation of trends over the course of time allow the industry to update the projections and react (if necessary) accordingly.

A logical reasoning besides visual acuity arguments for consumption bandwidth asymptotes can be built around the historical basis of much of the 50% CAGR number itself. From a media consumption standpoint, we can recognize that the speeds delivered from the Internet's mass-scaling outset were in large part associated with chasing the next required bit rate for increasingly higher level human media experiences:

1) Alphanumeric characters
2) Voice
3) Images (pictures)
4) Music
5) Low speed video
6) SD Video
7) HD 1.0

And now, per this analysis: HD 2.0 or HD 3.0 video. Based on our prior discussion on hyperacuity and studies such as [11], where it is foreseen that UHD delivers "retinal" image quality, and recognizing that normal households have wall sizes that limit the reasonable size of displays, there are several reasons to anticipate that *there is little practical benefit to services beyond the UHD format objective*s quantified in Table 8. In other words, our media consumption rate chase comes to an end (short of the holograms, etc.).

This then logically leads to our "asymptotic" limits for the service groups as shown in Table 11. Note that the IPV broadcast of 30 programs will be accounted for in a different manner on our Capacity Management Timeline. As we shall see, the broadcast spectrum offset plays a part in setting the threshold of available capacity for IP growth.

Table 11 – Projected Growth Asymptotes

	125 HHP	40 HHP
	(Mbps)	(Mbps)
Unicast Only	9470	3010
Multicast	7670	3010

Under these assumptions, it is insightful to recognize that the highest aggregate in Table 11 (125 HHP, Unicast only) is less than 10 Gbps, meaning it is lower than the objective called for as capacity for the DOCSIS 3.1 downstream – 10 Gbps or greater.

From these calculations, we can determine a 15-yr growth rate that ends up at 9470 Mbps/125 HHP from a starting point of 12 DOCSIS carriers (assumption) to a 500 hp serving group. The *resulting average CAGR for 15 years works out to about 33%.*

Over the 15 years, we have broken the growth rates into 3-year segments that could represent a possible play out of the entire 15-year transition. The incremental CAGRs used (average to 32.8%) are as follows:

Years: 1-3: 40%
Years: 4-6: 30%
Years: 7-9: 40%
Years: 10-12: 35%
Years: 13-15 (Complete Transition): 20%

The pattern recognizes that current CAGRs may settle after successive years of very rapid OTT growth. A new engine may be cable managed IP Video services penetrating alongside current OTT services, which will begin slowly and be many years in becoming highly penetrated, but will keep growth compounding steadily.

In the 7-9 year period, there will be some scale of 4K television sets (2015 being the threshold year) and a reason to move services in the direction of HD 2.0, increasing CAGR in the process before once again settling at the end of the cycle. This is when the final phase of the UHD service mix has been deployed, consumer usage patterns settled, and no clear growth engine for continued bandwidth expansion emerges for media consumption, at least as we know it and understand it today for business planning purposes.

Using the above CAGR segments and the boundary conditions calculated previously, the growth trajectories described and the asymptotes calculated can be visualized on the Capacity Management Timeline as the projection shown in Figure 18 for the two serving group sizes discussed. We have added two years at the end to allow us to envision a full tapering of CAGR. We will discuss the other Figure 18 markers and labels in the next section.

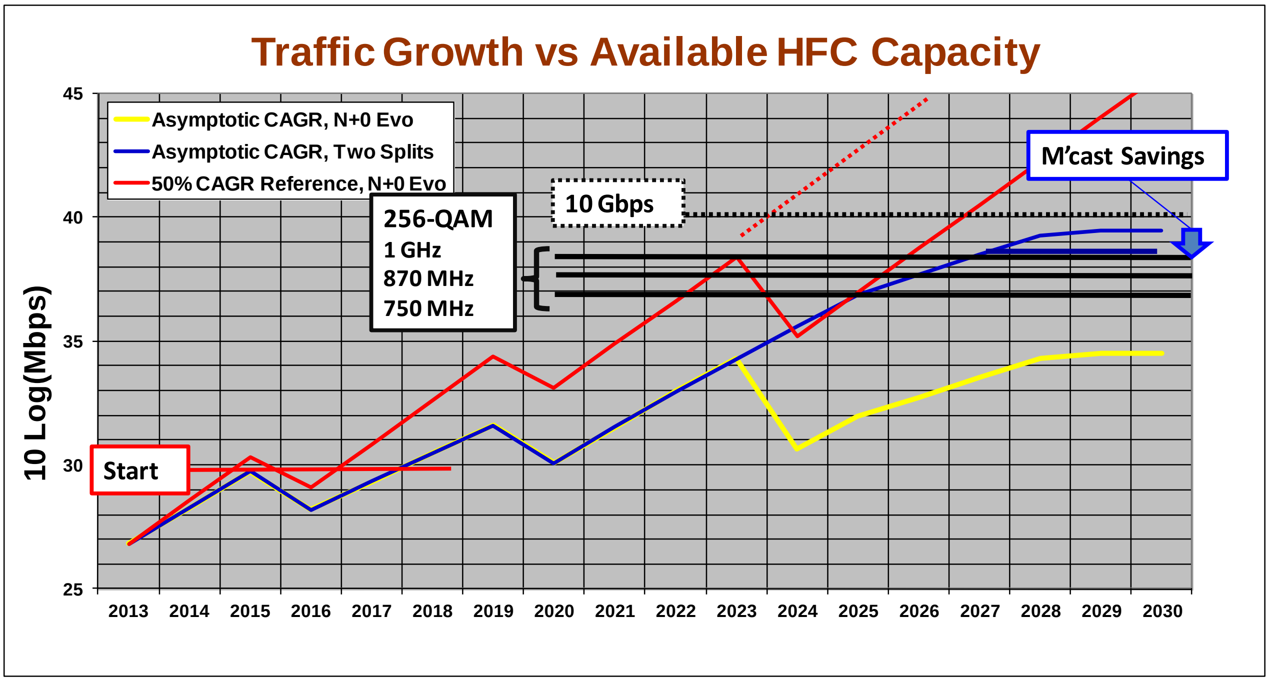

Figure 18 – Staggered CAGR Tapering Towards a Deterministic Video Services Evolution

As previously highlighted, it is very significant to point out that the assumptions of the CAGR rollercoaster ride and ultimate settling aligned with "retinal" video quality yields a sub-10 Gbps capacity aggregate. As Figure 18 shows, it does not leave very much margin beneath 10 Gbps for variations in assumptions, but then again DOCSIS 3.1's objective is a *minimum* of 10 Gbps.

Additionally, we have plotted the case that includes potential multicast savings – a small relative offset on Figure 18 as a decibel scale, but perhaps one that becomes important savings given the small margin between unicast and the 10 Gbps threshold.

Note that in prior work [5], as discussed, we estimated broader *theoretical* asymptotic boundaries of media consumption. The two key factors considered in [5] that are not addressed in this more practical perspective are consideration of the entire field of view of human visual system – because clearly sending more than this would not be sensible – and frame rates up to 300 fps based on some advanced research into how high the frame rate can go and yield an observable

VQ difference. Please see [5] for more details.

In this paper, we are taking the approach that such theoretical boundaries are unlikely to come into play in a time frame meaningful to be planning for them. A key conclusion even in that case was that, for an N+0 architecture, the aggregate *theoretical* capacity requirements could still be met, albeit with more aggressive RF distribution evolution [7]. This again is encouraging regarding our N+0 capacity fulfillment expectations.

Lastly, we re-emphasize that we are focusing on *media consumption-based service only as we understand them today*, and not trying to account for yet-to-be-determined bandwidth hungry applications, or even of the volume display or hologram sort of media consumption – considered outside the window of interest to compare.

AN EXAMPLE, PHASED MANAGED EVOLUTION OF SERVICES

Refer again to Figure 18. Several thresholds are drawn horizontally across the traffic trajectories shown.

On the left hand side ("Start"), the assumption used to guide the timeline is that we have fully utilized spectrum today. There is very little room to add new DOCSIS channels to accommodate growth. This is generally where MSOs are today in North America. We orchestrate the analysis in a way that is mostly agnostic to whether the network is 750 MHz or 870 MHz of downstream bandwidth. We only assume that in both cases, the spectrum is full. They will be different, obviously, in the types and amounts of services they offer. As a simple example, the 870 MHz network may carry 50 more broadcast HD channels.

The assumption threshold at "Start" is that up to 24 downstream DOCSIS channels can be squeezed out through some combination of tools – be that more SDV or removing some analog programming. In so doing, combined with a node segmentation plan over the next three years, the most aggressive CAGR situation stays below threshold for the duration, leaving a relatively short time window to execute on additional bandwidth recovery mechanisms. With the defined CAGR slow down (blue), this is extended by only about a year.

The upper right thresholds of Figure 18 are insightful. There are three traditional HFC downstream spectrum definitions, and the capacity associated with each if they were completely full of 256-QAM. These work out to, for the bit rate on the wire, as:

750 MHz – 116 slots; ~5.0 Gbps
870 MHz – 136 slots; ~5.8 Gbps
1 GHz – 158 slots; ~6.8 Gbps

Of course, these are not available capacities (yet) as they are consumed with legacy services. But, they offer immediate insight into the ability to architect for sufficient long-term capacity by simply comparing them to the trajectories in blue (125 HHP) and yellow (40 HHP).

The final threshold is the DOCSIS 3.1 objective of 10 Gbps. With DOCSIS 3.1 defining advanced modulation profiles and extended spectrum, there are multiple combinations of spectrum allocation and M-QAM order to achieve 10 Gbps.

+5 Years

As pointed out previously, while the end state of IP transformation may be encouraging, getting over the "simulcast bubble" is a complex capacity management challenge. Our groups of 256-QAM thresholds in Figure 18 look encouraging, but nearly all of that capacity is already spoken for. We have stitched together a possible scenario that looks at two 5-year snapshots of service and architecture evolution that encompasses current trends in service, HD 2.0 and HD 3.0 projections, and line items from the Assets toolkit.

Consider Figure 19. A new purple threshold range has been added for analog reclamation savings, with an assumption that 100% of analogs (60) will be removed at the end of five years. Most MSOs expect to remove all analog services, though fully doing so in 5 years may be aggressive for some. The analysis is easily adjustable for some small subset of analogs it may be desired to keep for traditional cable TV customers. However, it is readily apparent what the capacity bang-for-the-buck is in this most efficient of spectrum reclamation actions that can be taken.

The range associated with the analog reclamation rectangle is associated with the level of DOCSIS 3.1 migration – from none to an average 25% additional efficiency. This could be full migration enabled up to 1024-QAM, or partial with an average better than 1024-QAM on the D3.1 services.

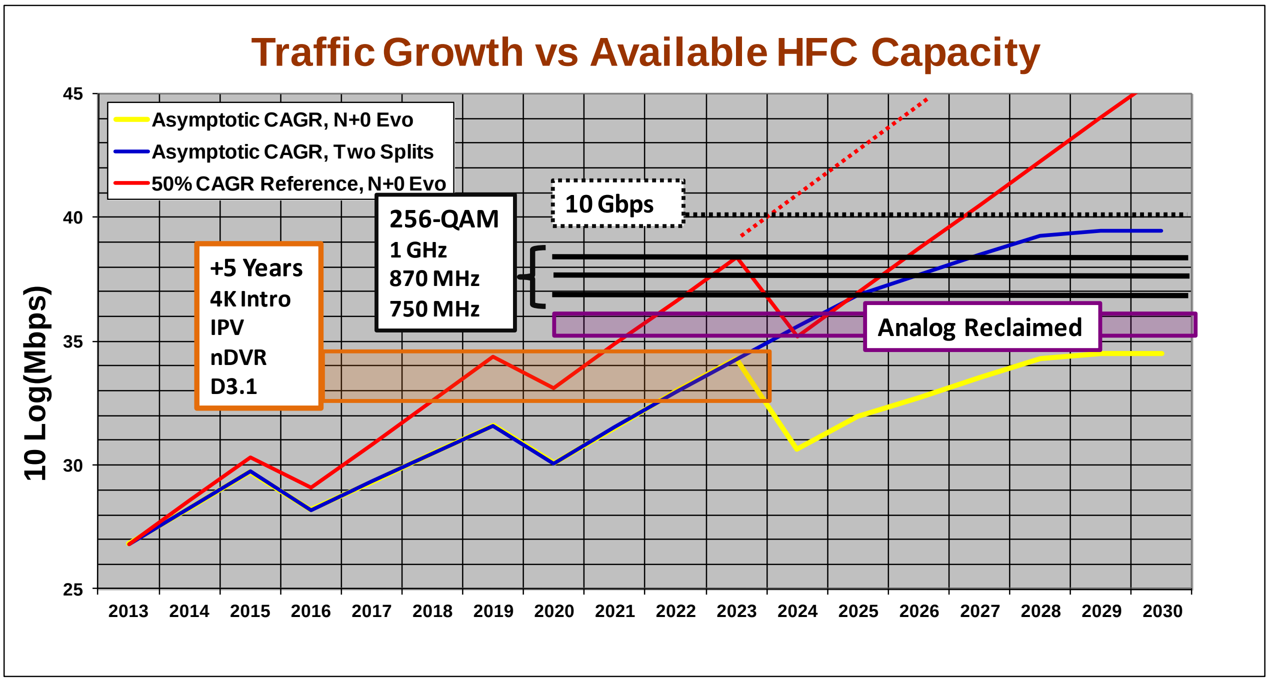

Figure 19 – Capacity Management Timeline – 5-year Snapshot

The orange rectangle represents what capacity remains for IP growth when the analog reclamation assumption is made, but with set-aside spectrum for legacy digital services and new video service evolution. The set of service and architecture evolutions assumed to have taken place by year 5 are:

Five-Year Snapshot

-Introduce 4K content into the VOD service offering (HEVC)
-Mix of VOD usage shift: 70/30 HD/SD to 30/60/10 UHD-1/HD/SD (no change in usage concurrency at pbh of 10%)
-Broadcast 10 programs in 4K (HEVC)
-IP Video Simulcast (25% penetrated, D3.0)

The same IP Video modeling tool previously extended for HD 2.0/3.0 services was used in its more common role of determining the number of DOCSIS 3.0 slots needed to support the legacy video mix, with the analysis resulting in 19 DOCSIS 3.0 slots required (and again not very much to be gained by multicast). At this phase of in the timeline, 250 HHP is the state of the serving group size. The results are shown in Table 12.

Table 12 – IPV BW Rolled Out over D3.0

Total DOCSIS Capacity Required to Deliver Managed IP Video Services

Total Bandwidth Requirements - Case 1

Service Group Size HHP:	250	Peak IP Video Devices / Home:	3.7 Active
Max Video Penetration:	17.5%	At Peak Penetration:	43.75 Homes
Percentage of Video Devices Active at Peak:	70%	At Peak Viewership:	113 Streams

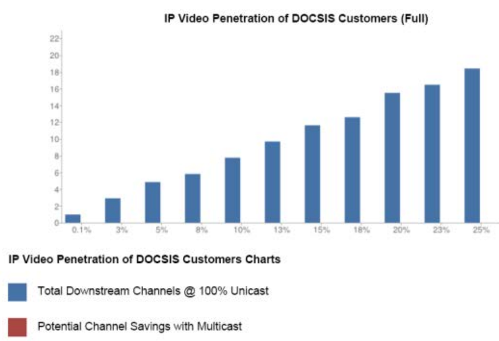

Creating the range of variation (the rectangle) are two considerations not deterministically assumed:

1) DOCSIS 3.1 penetrated in a meaningful way, or not at all. When so, the calculation assumption is enablement of 1024-QAM for bandwidth efficiency.

2) Rolling out of network DVR services, or not. Recorded content today in the US represents about 1/3 of the total content viewed [24]. Guaranteed unicast concurrency of video services today represented by VOD will increase with nDVR to this value plus VOD, worst case. We assume at +5 years that it has risen to 20% of viewing in addition to VOD. The same viewing format mix is assumed. This represents an aggressive assumption given the limited UHD-1 content available to record at this stage of the service evolution.

Observing Figure 19, what can we conclude? First, to make even reasonable room, a complete analog reclamation was assumed, and the 5-year plan may be aggressive for that assumption. However, it is likely that partial analog reclamation at least can be assumed.

Despite the resources freed up, the simulcast bubble that keeps legacy QAM in place while adding new bandwidth-draining services may leave the network vulnerable under persistently aggressive 50% CAGR at the end of the decade (2018). Implementing a second split by decade’s end, along with gains of DOCSIS 3.1 and removal of nDVR from the mix extends this by about three years (2021).

Alternatively, a settling of the aggressive 50% CAGR, as shown in the blue trajectory, does not threaten the entire range of “thresholds” in the orange box for the 125 HHP serving group sized over the next nine years. That is a comforting window of time, and emphasizes why it is also important to keep a continual eye on CAGR trends. Again, the introduction of DOCSIS 3.0 IP video does not create new eyeballs in the service group; it (mostly) shifts them. Accounting for the spectrum allotted to IP Video channels (19) *and* assuming full-speed ahead with 50% CAGR may be an unrealistic exercise in double counting. A settled CAGR coupled with the pre-allocated IP Video traffic-engineered spectrum may better describe the dynamics.

Overall, the subtle lag in the CAGR per the prior assumptions (6-yr CAGR becomes 35%) pushed the “Year 5” service mix out to at least 8 years before the modified trajectory breaches the capacity barrier. Without the introduction of nDVR, this extends to 10 years. And, of course, with a less aggressive nDVR assumption – either the viewing mix skewed more towards traditional HD or more modest penetration – something in-between would result.

The 5-year snapshot also points out how subscribers might first get a taste of the new UHD-1 service in VOD and limited broadcast. By doing so, MSOs can observe the reaction and interest, and determine possible ARPU avenues in scaling the hot new format accordingly. Capacity management at this stage involves primarily limiting early programming, expecting segmentation, and removing analog carriers. In this case, we have left two levers – the Asset of DOCSIS 3.1 and the Liability of nDVR – as examples of how strategic decisions reflecting the pace of service and technology infrastructure investment can impact network lifespan.

<u>+ 10 Years</u>

Now consider Figure 20.

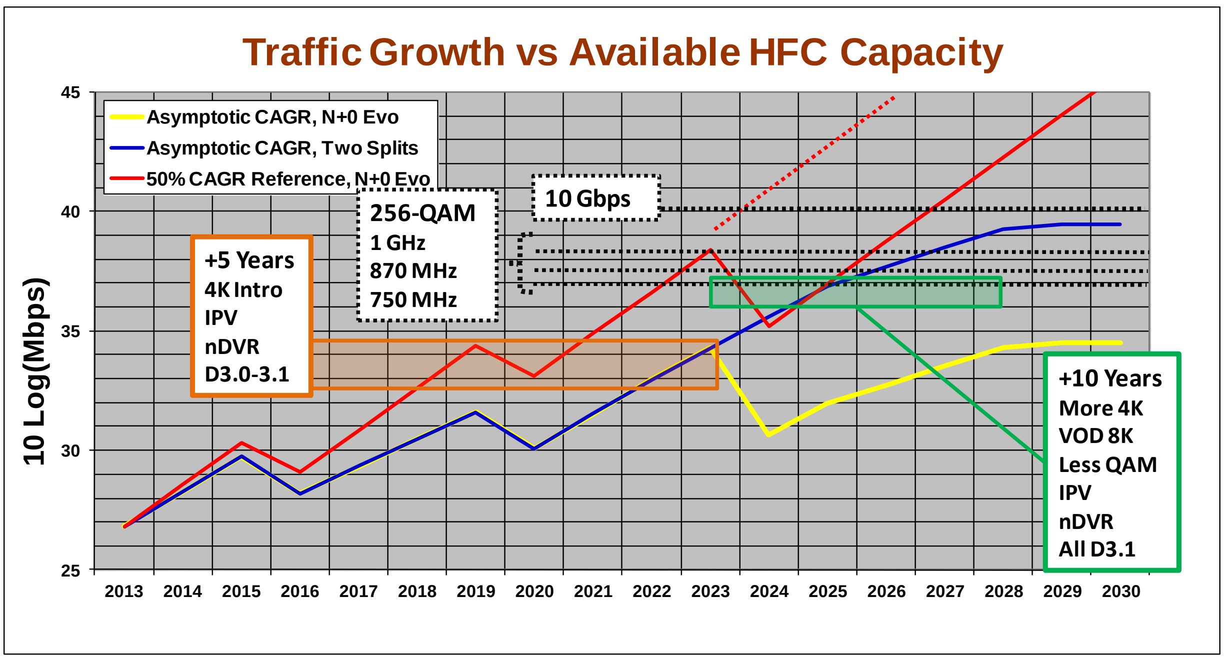

Figure 20 – Capacity Management Timeline – 5 and 10 year Snapshots

Here, a green rectangle represents a snapshot at +10 years, where the assumptions of the evolution state are as follows:

Ten-Year Snapshot

-VOD Mix 70% UHD-1, 25% HD, and 5% UHD-2 (8K home theatre, always HQ)
-VOD all IP/DOCSIS 3.1 (IP Unicast)
-Broadcast 30 programs in 4K (HEVC); 50% HQ (high frame rate and 10-bit), moved to IP
-IP Video is 50% penetrated
-All of DOCSIS 3.0 to 3.1
-DOCSIS 3.1 assumption 2048-QAM avg
-Total recorded viewing (nDVR traffic engineering) @ 30%
-50 legacy SD and HD channels broadcast

Most of the increase in available capacity comes from the massive reduction of legacy QAM content, making room for what now is a sizable broadcast of UHD-1 4K programming. The 50 programs of broadcast are along the lines of maintaining the basic service subscriber's offering today, just translated to the era of "you-can-no-longer-avoid-digital." There are also savings in D3.0 to D3.1 bandwidth efficiencies that are meaningful, though not as large.

Note that the entire programming offering of course still exists in SD and HD, it has just been migrated to IP, as has the 4K HD, which presumably has passed its consumer interest test at Year 5 as a service to scale to mass consumption.

The capacity gap between the orange and green rectangles identifies new lifespan to be engineered should the potential breach in capacity threshold that projects in Figure 19 come to fruition. Of course, the evolutionary steps are not discrete as pictured, and steps towards the ten year rectangle of capacity can be made that basically bridge the gap in the figure with a continuous threshold moving "Northeast" on Figure 20.

These bandwidth reclamation measures for the 125 HHP case, however, run out of steam by 2026, and there is no path under these assumptions that suggest two service splits suffice under either the continued aggressive CAGR of 50%, or the settled CAGR case which asymptotes just below 10 Gbps.

However, it is also notable that a prior-to-2024 evolution to N+0 meets the requirements of settled CAGR capacity growth, though not a persistently aggressive (red trajectory) 50% CAGR. Of course, this is really simply testament to the axiom that without infinite bandwidth, over time, CAGR always wins – it is just a matter of the time scale chosen.

In summary, for this 10-yr snapshot and assumptions, we have shown that an N+0 evolution meets the capacity needs of long term HD 2.0/3.0 evolution coupled with an HSD service, and meets the transitional needs of the simulcast period prior to getting to the asymptotic point. We have more deeply penetrated new services and technology by Year 10, while reducing legacy service offerings in a managed way that stays within available capacity with some level of certainty in the N+0 case. There is a 10+ year time window to evaluate whether a 125 HHP "sweet spot" is sweet enough to emerge as sufficient for the years that follow.

<u>Network Nirvana</u>

We calculated our 15-year end state IP transformation previously. There, legacy-free IP video bandwidth requirements under HD 2.0/3.0 evolution were determined using the modeling tool extrapolated to the UHD generation of video services.

As we did in the transitional 5-yr and 10-yr snapshots, we now account for the set aside spectrum at Year 15, in this case for 30 broadcast IP programs in UHD-1 and 1080p60 HD. Again, our UHD-1 assumptions are that approximately half of the programming would benefit from the enhanced video quality. UHD-2 is available as a unicast IP Video stream to the home theatre crowd. Recall, we calculated that the spectrum required for the IP broadcast would be 85 MHz.

Lastly, to truly account for all long-term network expectations, we must account for the expanded upstream, which may top out at 200 MHz. It is anticipated that an 85 MHz mid-split will be the first phase of upstream evolution, which has an impact on the downstream as well, but which is relatively modest. A life expectancy of at least ten years is foreseen at upstream growth rates [8], [9] under two splits and 85 MHz.

However, we will assume the more aggressive case for the "high split" comes to pass by Year 15. It is consistent with the DOCSIS 3.1 objective of enabling 1 Gbps of capacity upstream. Obviously, it also represents the most threatening case for managing downstream growth, since significant downstream spectrum is now sacrificed. In fact, the 1.2 GHz forward band was first conceived as a way to offset the loss of downstream spectrum as the upstream expanded [7].

For this Year 15 case, refer to Figure 21.

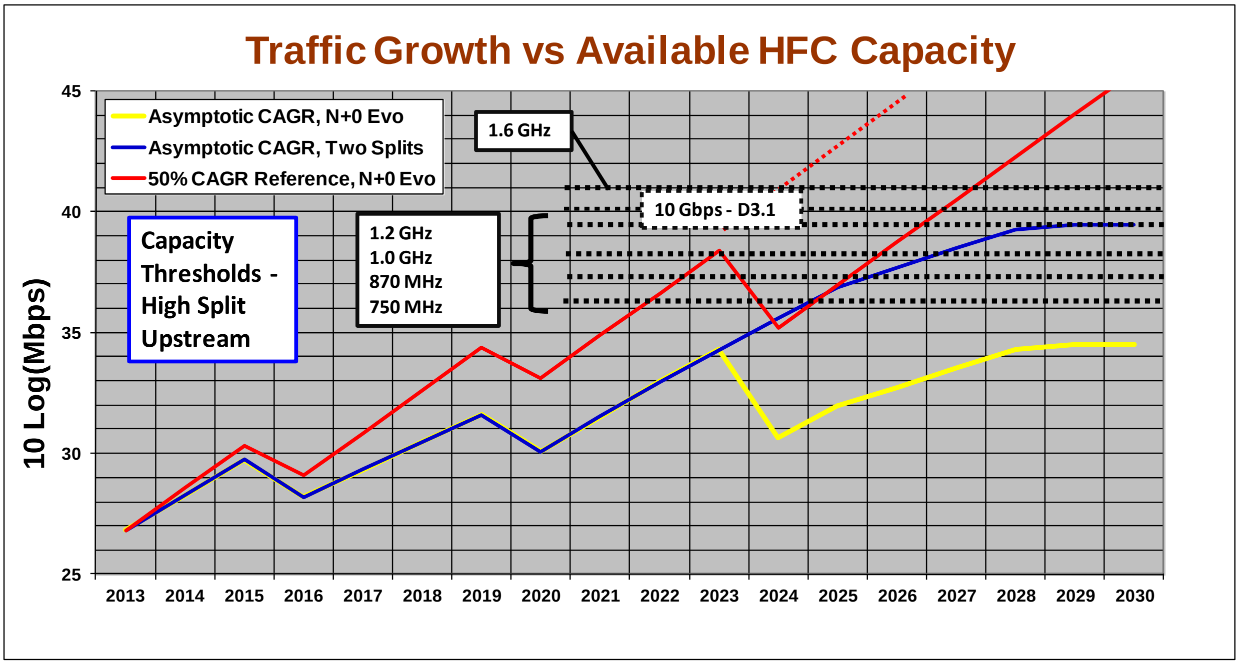

Figure 21 – Capacity Management Timeline
15 Yrs, IP Broadcast/Unicast Architecture, Upstream High Split

Figure 21 identifies the new capacity thresholds under the assumptions of the unavailable 85 MHz broadcast spectrum and spectrum donated to the upstream for several forward spectrum scenarios. The figure assumed that the forward path would begin at 250 MHz and achieves 2048-QAM.

As might be expected from all of the clues we have accumulated in our Balance Sheet tables and calculations along the way, an evolution to N+0 is sufficient from a capacity perspective in the end state of all-IP transformation, with asymptotic consumption behavior, even for the most constraining of forward bandwidth scenarios, although the gap closes noticeably with decreasing available downstream.

As important, it is *not* the case that capacity is always sufficient when the segmentation is limited to two splits. *Without bandwidth extension, capacity is insufficient*. Under the asymptotic growth scenario and an extension to 1.2 GHz, the capacity threshold and consumption asymptote are basically identical. The many variables in play could swing that scenario either way. For example, we might assume DOCSIS 3.1 makes it to 4096-QAM for all or enough to move the bar, or will we may squeeze a little more efficiency out of HEVC than foreseen.

However, the large impact variables involve assumptions of the average CAGR over long periods of time, and percentage viewership metrics for the most bandwidth-consumptive HD 2.0 and HD 3.0 streams. The accuracy of these assumptions is more likely to determine the sufficiency of defined thresholds, and especially so for borderline cases.

Perhaps most importantly, like any capacity management analysis, the work is a living document, with periodic updates associated with trends observed and technology shifts necessary to adapt the path forward.

SUMMARY

In this paper, we executed a long-term capacity management analysis, with permutations of scenarios of current and future services and architectures. While such long time windows are sensitive to assumptions, it is important to understand that all of the possibilities are quantifiable in straightforward fashion. The analysis undertaken here quantified each individual evolution variable in terms of its role as Liability or Asset, considering the extended time period of an HD 2.0/3.0 evolution and IP Transformation, and described the intricacies of deploying these Balance Sheet line items in a Capacity Management Timeline analysis.

From our perch today, the end state appears to be an attractive one in terms of available capacity for projected services over evolved HFC. We also examined a phased example transition to visualize the complex capacity balance and timing involved in crafting effective migration strategies. We reemphasize that, as insightful as the results herein may be, long-term capacity analysis is very much a living exercise and our team expects to continue to update our perspective periodically. Nonetheless, being able to comprehensively understand and methodically quantify the problem is essential to properly engage in effective scenario planning at every stage of the exercise, and to enable optimization of solution paths suited to an operator's particular circumstances. We hope this paper helps the industry in exactly this manner.

REFERENCES

[1] Armstrong, M and D Flynn, M Hammond, S Jolly R Salmon, *High Frame Rate Television*, BBC Research Whitepaper WHP 169, September 2008.

[2] De Simone, F et al., *Towards high efficiency video coding: Subjective evaluation of potential coding Technologies*, Journal of Visual Communications (2011), doi:10.1016/j.jvcir.2011.01.008

[3] Ho, Yo-Sung and Jung-Ah Choi, *Advanced Video Coding Techniques for Smart Phones*, 2012 International Conference on Embedded Systems and Intelligent Technology (ICESIT 2012), Jan. 27–29, 2012.

[4] Howald, Dr. Robert L, Dr. Sean McCarthy, *Bits, Big Screens and Biology*, The Cable Show Spring Technical Forum, May 20-22, Boston, MA.

[5] Howald, Dr. Robert L, *Boundaries of Consumption for the Infinite Content World*, SCTE Cable-Tec Expo, New Orleans, LA, October 20-22, 2010.

[6] Howald, Dr. Robert L, Dr. Sebnem Zorlu-Ozer, Dr. Nagesh Nandiraju, *Delivering Pixel Perfect*, The Cable Show Spring Technical Forum, May 11-13, Los Angeles, CA.

[7] Howald, Dr. Robert L, *Fueling the Coaxial Last Mile*, SCTE Conference on Emerging Technologies, Washington DC, April 2, 2009.

[8] Howald, Dr. Robert L, *Looking to the Future: Service Growth, HFC Capacity, and Network Migration*, 2011 Cable-Tec Expo Capacity Management Seminar, sponsored by the Society for Cable Telecommunications Engineers (SCTE), Atlanta, Ga, November 14, 2011.

[9] Howald, Dr. Robert L, and Phil Miguelez, *Upstream 3.0: Cable's Response to Web 2.0*, The Cable Show Spring Technical Forum, June 14-16, 2011, Chicago, Il.

[10] Marpe Detlev , et al., *Video Compression Using Nested Quadtree Structures, Leaf Merging, and Improved Techniques for Motion Representation and Entropy Coding*, IEEE Trans. Circuits Syst. Video Technology., Vol. 20, Nr. 12 (2010) , p. 1676-1687.

[11] McCarthy, Dr. Sean T., *Quantitative Evaluation of Human Visual Perception for Multiple Screens and Multiple CODECs*, 2012 SMPTE Annual Technical Conference & Exhibition

[12] McCarthy, Dr. Sean T, *A Biological Framework for Perceptual Video Processing and Compression*, SMPTE Motion Imaging Journal, Nov/Dec 2010.

[13] McCarthy, Dr. Sean T., and W.G. Owen, "Apparatus and Methods for Image and Signal Processing," US Pat. 6014468 (2000). US Pat. 6360021 (2002), US Pat. 7046852 (2006), 1998.

[14] Sullivan, Gary J. and Jens-Rainer Ohm, *Recent Developments in Standardization of High Efficiency Video Coding (HEVC),* SPIE Applications of Digital Image Processing XXXIII, Andrew G. Tescher (editor), Proceedings of SPIE Volume 7798, Paper number 7798-30, August, 2010.

[15] Ulm, John and Gerry White, *Architecture & Migration Strategies for Multi-screen IP Video Delivery*, 2012 SCTE Canadian Summit, March 27-28, Toronto, CA.

[16] Wiegand, T, G. Sullivan, G. Bjontegaard, and A. Luthra, *Overview of the H.264/AVC Video Coding Standard*, IEEE Trans. Circuits Syst. Video Technol., vol. 13, no. 7, pp. 560-576, July 2003.

[17] Yamashita, Takayuki et al. "Super Hi-Vision" Video Parameters for Next-Generation Television, SMPTE Mot. Imag J. May-June 2012 vol. 121 no. 463-68 (NHK Science & Tech Res. Labs)

[18] ITU-T and ISO/IEC, ITU-T Rec. H.264 | ISO/IEC 14496-10 Advanced Video Coding (AVC), May 2003 (with subsequent editions and extensions).

[19] ISO/IEC JCT1/SC29/WG11 (MPEG), "Description of High Efficiency Video Coding (HEVC)," doc. no. N11822, Daegu, KR, January 2011.

[20] ISO/IEC JCT1/SC29/WG11 (MPEG), "Vision, Applications and Requirements for High Efficiency Video Coding (HEVC)", doc. no. N11872, Daegu, KR, January 2011.

[21] ITU-R BT.2020 (2012) Recommendation -"Parameter values for ultra-high definition television systems for production and international programme exchange"

[22] advanced-television.com

[23] CarltonBale.com

[24] http://mediacenter.motorola.com/Content/Detail.aspx?ReleaseID=15389&NewsAreaID=2&ClientID=1

[25] http://www.cable360.net/ct/webcasts/2010_08_10/

[26] http://www.ieee802.org/3/bn/public/mar13/howald_3bn_01_0313.pdf