BITS, BIG SCREENS, AND BIOLOGY

Dr. Robert L Howald and Dr. Sean McCarthy
Motorola Mobility

Abstract

High definition television (HDTV) has dramatically improved the consumer viewing experience. As such, despite its hunger for precious bandwidth, increased HD programming continues to be a key industry objective. However, as evidenced by the exhibits and technology on display at the Consumer Electronics Show (CES) this past January, today's HD, is just a step in the progression of consumer video. In addition, today's video processing and delivery is also undergoing significant change. In this paper, we will explore new generations of HD video and the innovative enabling technologies that will support them. We then roll-up the components and project their impact to network architecture.

Specifically, we will consider advanced formats, beyond just emerging 1080p60 (blu-ray) HD. Recognizing the expectation of a very long HFC lifespan, we will quantify how QFHD (aka 4k) and even proposed "Super Hi-Vision," or UHDTV, stack up for consumer services. We will assess practical and human factors, including those associated with HD-capable second screens, such as tablets. We will quantify physiological variables to the optimization of the video experience, such as personal through immersive screen sizes, viewing environment, and high frame-rate television.

On the encoding side, we discuss H.265 High Efficiency Video Coding (HEVC) against its own "50%" objective. And, just as we considered human variables associated with the user experience, we can take advantage of human biology to deliver the highest perceived quality using the smallest number of bytes. Using new signal processing models of the human visual system (HVS), the ultimate arbiter of video quality, a unique combination of bandwidth efficiency and high perceived video quality can be achieved. This technique, called Perceptual Video Processing (PVP) will be detailed, and its impact on video quality and bandwidth quantified.

In summary, we will evaluate long-term network prospects, capturing the potential trajectory of video services, innovative encoding techniques, emerging use cases and delivery, and shifting traffic aggregates. In so doing, an enduring network migration plan supporting multiple generations of video and service evolution can be projected.

INTRODUCTION

Decades of broadband growth and an ever-increasing range of video services has given operators a sound historical basis upon which to base future growth trends, which is critical for business planning. Service growth and subscriber satisfaction with the portfolio of media delivered to them provides new revenue opportunities. To meet these demands, key decisions must be made for upgrading hubs, homes and the access networks. The prevailing MSO approach has been a very successful pay-as-you grow approach, capitalizing on technologies as they mature and as consumer demands require. This has worked extremely well because of the latent HFC capacity, which incrementally was mined as necessary by extending fiber, adding RF spectrum, incorporating WDM optical technologies, and delivering digital and switched services.

As IP traffic has grown aggressively, video quality has also moved ahead, albeit at a more gradual pace. The appetite for HD is being fed at this stage of the evolution, but the HD lifecycle itself has only just begun. As cable systems deliver 720p and 1080i formats, the ability to support and deliver 1080p quality already exists in the CE and gaming worlds. Flat panel televisions continue to become larger, more capable, and lower cost. Their size already is breaching the boundary of where a "normal" viewing distance would benefit from a yet higher quality video signal. 2k and 4k (Quad Full HD or QFHD) formats have entered the conversation and the demonstration rooms. These formats are being explored and seemingly will inevitably lead to a new service offering. Beyond QFHD is the Ultra-High Def (UHDTV, 4x QFHD)) format, or Super Hi-Vision, invented by NHK in Japan in the mid-1990's. At that time, it was foreseen by NHK to be a consumer format in the 2030 time frame.

EVOLUTION OF VIDEO SERVICES

Spatial Resolution

With the advent in particular of HDTV, development of QUAD HD (2x in each dimension) and UHDTV, the video and CE industries have a strong understanding of the relationship among resolution required, screen size, and viewing distance.

Just as visual acuity is measured and referenced to object sizes at defined distances, the display size and placement relative to the viewer is a key piece of the resolution requirement equation. Figure 1 is a straightforward way to see how these factors interact [40] based on recommendations provided by multiple professional organizations, home theatres experts, and retail manufacturers. Generally, for a fixed resolution (linear trajectories on the plot), a larger screen size is best viewed further away. 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. Finally, for a fixed distance from the display and the higher the format resolution, the larger the screen size should be.

As a simple example, a 50 inch screen, if viewed more than 20 ft away or greater, will begin to lose the benefit of HD at 720p, and provide 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 distinguishing of pixels. This chart thus also explains the increased pixel count of UHDTV based on a 100" display recommendation and wider viewing angle (closer).

The guidelines come from different organizations and retailers, and while they tend to cluster around similar recommendations, they are not in complete agreement. This is generally due to the varied perspectives of the organizations, such as, for example, what sells more TVs. The range of recommendations varies from about 1.5x-2.5x of display size for viewing HD content, with the lower end corresponding to 1080 resolution.

The recommendations are also correlated to an assumption about visual acuity as it relates to the ability to resolve the image detail. They are also associated with viewing angle considerations. For example, the recommended optimum fields of view are given as about 30° (SMPTE) or 40° (THX) in the horizontal plane. In the vertical plane, simpler guidelines are designed around avoiding neck strain, and so describe maintaining at least a 15° vertical field of view. The maximum recommended, beyond which neck strain is a risk, is a 35° viewing angle.

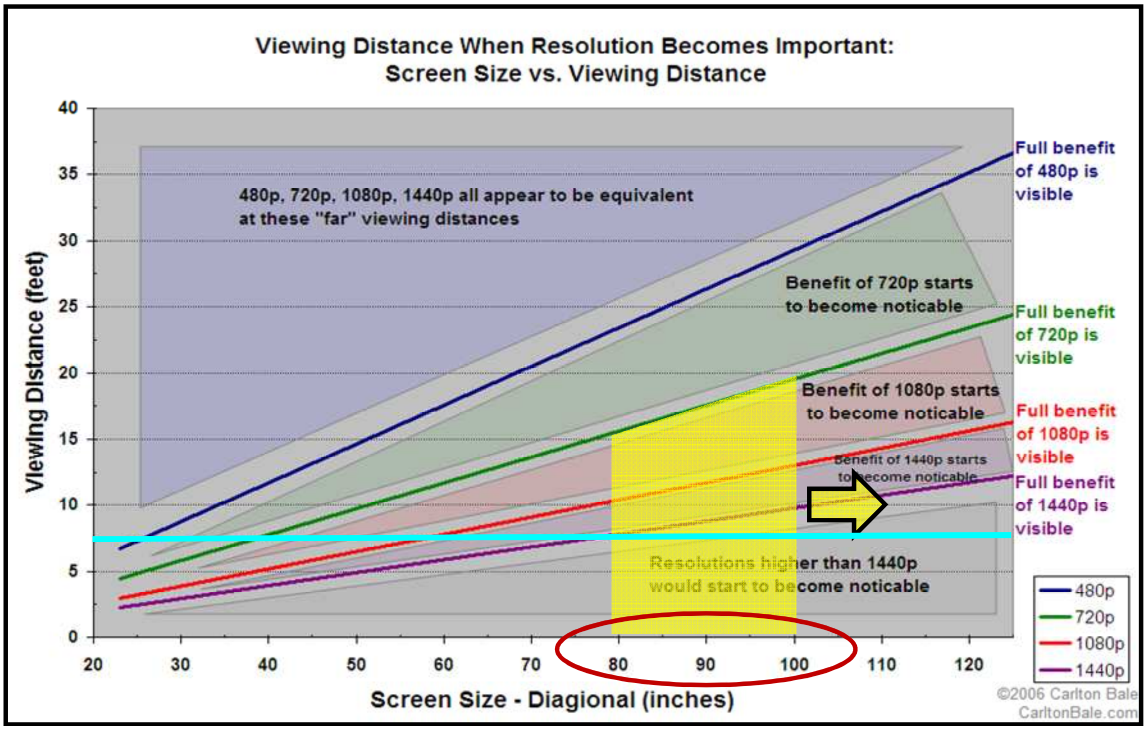

Figure 1 – Screen Size, Viewing Distance, and Spatial Resolution

Let's ponder modern display capabilities. Consider the bottom right corner of Figure 1, shaded yellow. A typical viewing distance in the home today is about 7.5-9 feet, which certainly has been driven in part by historical screen sizes. It is not surprising for anyone who has visited a big box retailer recently that flat panel screens are available now at ever-increasing sizes, such as those shown in the shaded yellow range of Figure 1. 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). At 80", flat panels have fully breached the 2560x1440 resolution threshold, sometimes referred to as Extreme HD (4x 720p HD resolution) in the gaming world. The next stop beyond this is QFHD at 3840x2160p. Based on this figure, there is potential viewing value for this format screen size and larger.

Note that UHDTV, was viewed as a 100 inch screen, but also viewed at only about 1 meter (3.3 feet). The intent was to generate the feeling of immersion. Studies by NHK concluded that feelings of discomfort often associated with immersive viewing such as IMAX level off with screen size at a certain point. In the case of UHDTV, the angle at which this occurs is for 80 inch screens. Therefore, a screen fully 100 inches is not expected to present an increased probability of discomfort, but yet yields the level of immersion and video quality desired in the experience.

Now consider Figure 2. Not only do larger primary screens translate into the need for better spatial resolution, our secondary screens also have gotten larger, simultaneously more portable, *and* capable of high quality video such as HD. The explosion of tablets has put an entire new generation of high-quality video capable screens literally at our fingertips for deployment virtually anywhere relative to our viewing perspective.

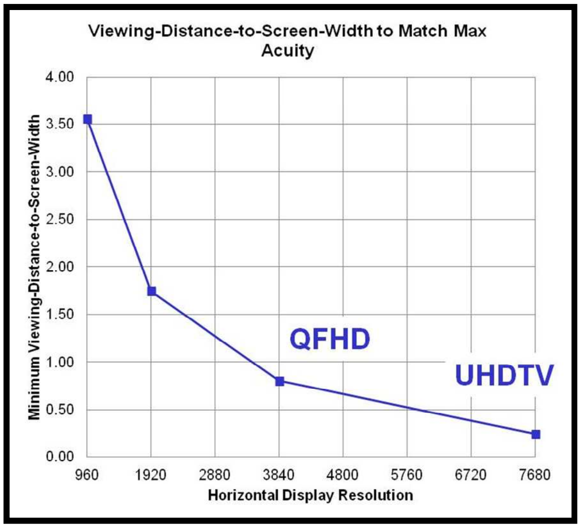

Figure 2 – Screen Size, Distance, and Resolution – Mobile Viewing

In the figure above, it is easy to see how the 1920x1080 resolution can be improved upon for reasonable viewing conditions. For a 10" Motorola Xoom tablet, for example, if the screen is about 17" away, its spatial resolution can be perceptibly improved with a higher resolution format. It is not hard to envision this scenario, for example, on an airplane or with a child in the backseat of a car.

The case for full QFHD or UHDTV on the 10" tablet would be difficult to make based on this figure without some other accompany variables. Nonetheless, clearly screen sizes and portability in this case are combining to change the paradigm of mobile viewing environments far, far away from the legacy of QVGA resolution at 15 frames per second.

Dynamic Resolution

The term "dynamic resolution" refers to the ability to resolve spatial detail of objects in motion. The 30 Hz (interlaced), 50 Hz, and 60 Hz frame rates have origins in AC line rates, and thus are not scientifically tied to video observation and testing. They simply exceeded what was known at the time about 40 Hz rates causing undesirable flicker. Most early analysis on frame rate was to ensure that motion appeared realistic (seamless), as opposed to a sequence of still shots. There was less focus on eliminating other artifacts of motion, such as smearing effects. Yet, as spatial resolution has improved, temporal resolution has not. Interlaced video itself is a nod to the imbalance of 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.

The above frame rates have since become embedded in tools and equipment of the production, post-processing, and display industries, and so, with respect to frame rate, we are hostages to the embedded infrastructure and scale of change that would be required to do anything else. As such, HDTV standards today are based on the 60 Hz interlaced or progressive frame rate. It has been suggested [1] that with larger and brighter displays of higher resolution, the frame rates in place based on practical implementation limitations of the 1930s era ought to be reconsidered, of course while recognizing a need to maintain some level of backward compatibility.

As displays become larger and of higher resolution and contrast, the challenges to effectively displaying motion increases, because the edges to which movement is ascribed are now sharper. What is optimal? There is not a firm answer to this question. The human visual system streams video continuously in a physiological sense, so the question is around the processing engine in the brain. Various sources describe tests where frame rates of 100-300 fps show perceived improvements compared to 60 fps [1, 41]. The difficulty of performing this type of testing – content and equipment – limits how much has been learned. There are potentially positive encoding implications to these higher frame rates. Intuitively, more rapidly arriving frames ought to be consistent with better coding efficiency, as it is likely that there is less variation frame-to-frame.

We will not consider any changes to frame rate beyond interlaced/30 to progressive/60. But this is a variable to keep an eye on as larger screens and live sports viewing collide.

Formats and Bandwidth Implications

High Definition has had a major impact on the industry in multiple ways. On the positive side, the Quality of Experience (QoE) delivered to the consumer is tremendously improved. HD has enabled cable operators to strengthen the service offering considerably. And, like the DVR, HD has very much the "once you have it, you never go back" stickiness to it.

Conversely, while HD services certainly act to increase revenue, they also create a significant new bandwidth burden for the operator. Whereas 10-12 standard definition (SD) programs can fit within a single 6 MHz QAM bandwidth, this number drops to 2-3 HD programs in a 6 MHz QAM. This loss in efficiency is compounded by the fact that HD today represents a simulcast situation – programs delivered in HD are usually also transmitted in the SD line-up. For all of the subsequent analysis, we will base MPEG-2 SD and HD program counts per QAM on averages of 10 SD/QAM and 2.5 HD/QAM. Obviously there cannot be a fractional number of programs n a QAM slot. The 2.5 assumes that for MPEG-2 encoded HD, an operator may chose 2 or 3 in a 6 MHz slot based on content type, and the QAMs are equally split with both.

Perhaps most worrisome with respect to bandwidth is that current services are basically HD 1.0. Only the first generation of formats are deployed – 1280x720p and 1920x1080i. The improvement over SD is so vast that it is easy to wonder what could possibly be the benefit of even higher resolution. However, as we showed in Figure 1, it is relatively straightforward to show how the continued advancement of display technology at lower and lower costs, in particular consumer flat panels, leads to reasonable viewing environments where resolution beyond 1080-based systems would be perceptible. In addition to the flat screen scenario, similar analysis in Figure 2 showed similar conclusions for"2nd screen" tablet viewing. All modern tablets support HD quality viewing. Coupled with realistic use cases that are likely to include close viewing distances, higher resolution scenarios may add value here as well.

We will consider the effects of two next generation video formats on the HFC architecture's ability to support them – QFHD and UHDTV. QFHD has had prototype displays being shown since approximately 2006 and has entered the conversation as the big box retailers now routinely display 80" screen sizes. Analyst projections have placed QFHD in the 2020 time frame for deployment timeframes. A comparison of these two formats against standard HD, and other formats, is shown in Figure 3 [19].

Note that QFHD works out to 4x the pixel count as 1080 HD, and UDHTV works out to 16x the pixel count. In each case, there is the possibility of higher bit depth (10-bit vs. 8-bit) as well, which translates into more bits and bandwidth. We will assume this is taken advantage of in the latter case only. As a result, we arrive at the following set of potential scaling factors, without any assumptions about possible latent compression efficiencies on top of conventional gains projections for new display formats.

SD to:

1080i – 4x
1080p – 8x
QFHD – 32x
UHDTV – 160x

It is of course premature to know precisely what compression gain may be available for advanced formats, since these enhanced formats are in their infancy. For now, we will rely on the resolutions to correlate with bandwidth, with the 8-bit to 10-bit pixel depth for UHDTV and the frame rate for p60 (doubling the information rate) as the only other variations quantified.

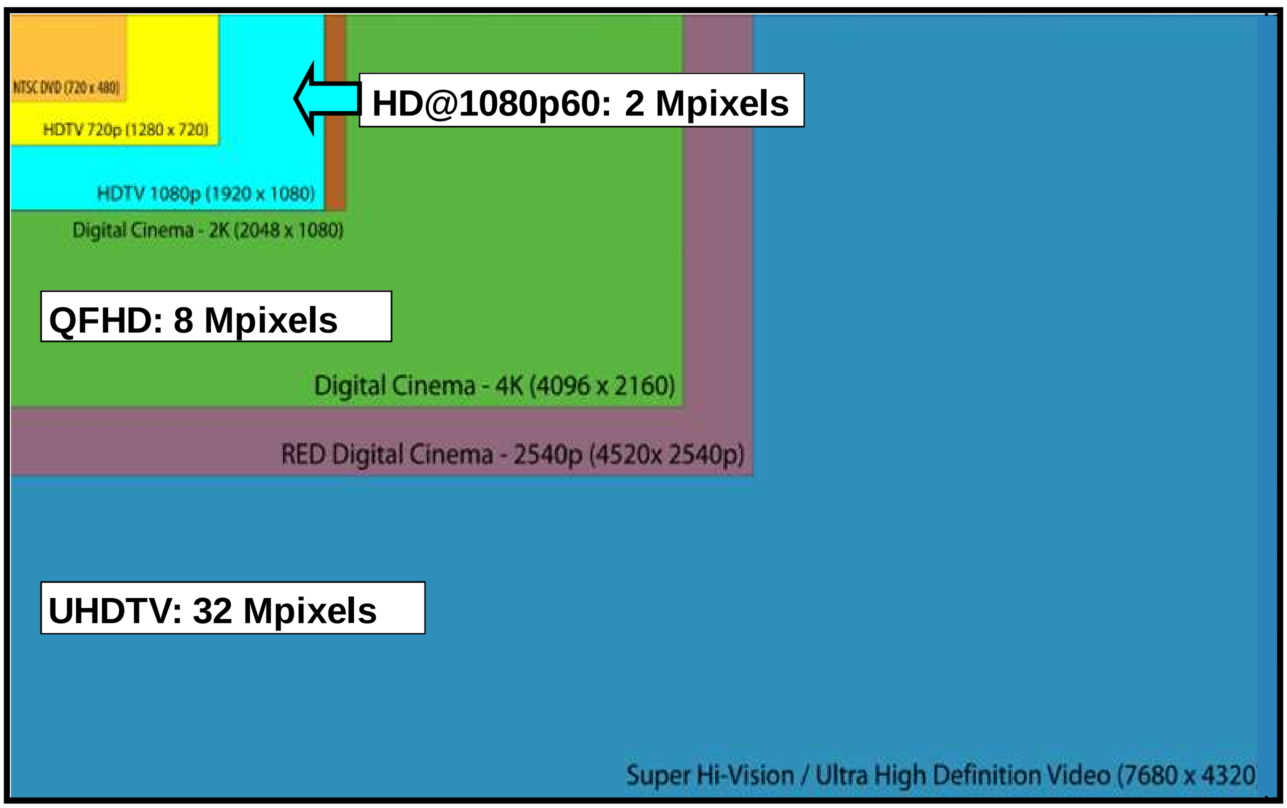

Figure 3 – Beyond High Definition Formats Comparison

VIDEO COMPRESSION – STILL ON THE MOVE

It may not seem like long ago, but it is nearly 10 years since the Advanced Video Coding (AVC) [27, 30] international standard was completed in 2003. AVC – also known as H.264 and as MPEG-4 part 10 – has been a remarkable success. It has enabled IPTV and HDTV to take hold and grow commercially. It has enabled Blu-Ray video quality at home. And it has been powering new models of delivering digital video over the internet. AVC 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'll be soon joined by a new entrant to the international standard portfolio -- High-Efficiency Video Coding (HEVC) [6, 10, 18, 22, 31, 32].

In many ways, HEVC is a close cousin to AVC. Both are of the same genus of hybrid block-based compression algorithms that incorporate spatial and temporal prediction, frequency-domain transforms, data reduction through quantization, and context-sensitive entropy encoding. Where HEVC stands out is in the wealth and sophistication of its coding tools, and in its superior compression efficiency.

Figure 4 captures the state of the set of core MPEG compression standards in the context of their lifecycle.

Efficiency

First and foremost for any compression standard is the simple question of how much more efficient it will be at compressing video streams. HEVC aims to double the compression efficiency of its AVC predecessor. AVC itself doubled compression efficiency compared to MPEG-2. That means that a consumer quality HDTV program delivered using 16 Mbps today with MPEG-2 (like a cable TV QAM channel supporting 2-3 HD channels) would need only about 4 Mbps using HEVC. As we will see in subsequent analysis, it also means that we might reasonably expect to be able to deliver Super HD (4kx2k) over the bandwidth we use today for regular HDTV, enabling yet another generation of enhanced video services.

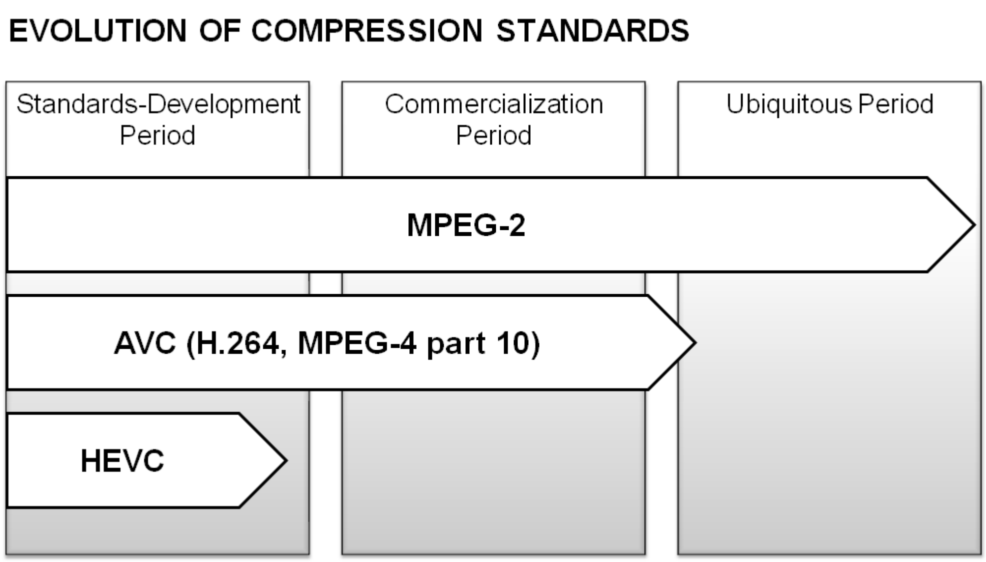

Figure 4 – State of Current Video Compression Standards

At the onset of the HEVC development process, the ITU-T and MPEG issued a joint call for proposals [33]. Twenty-seven proposals were received and tested in the most extensive subjective testing of its kind to date. Scrutiny of the proposals entailed 134 test sessions involving 850 human test subjects who filled out 6000 scoring sheets resulting in 300,000 quality scores. The conclusion [34] was that the best proposals yielded 50% bit rate savings compared to AVC at the same visual quality. The potential for another 50% gain launched the Joint Collaborative Team for Video Coding (JCT-VC), and HEVC development formally got underway.

In late 2011, JCT-VC reported another series of compression-efficiency tests [35] using objective rather than subjective methods. Those test showed that HEVC had not yet hit the 50% mark with scientific certitude, but was very close and had excellent prospects for additional gain. Table 1 shows the results from objective tests comparing HEVC to AVC High Profile. The tests were conducted using various constraints to examine the efficiency of HEVC for several important potential use cases: broadcast such as over cable, satellite, and IPTV that need random-access features to support fast channel change and trick modes, low-delay applications such as video conferencing, and intra-only compression that uses only spatial prediction within each frame of video to support applications such as contribution-quality video.

The bit rate savings listed in Table 1 represent the point at which HEVC and AVC High Profile produce the same peak-signal-to-noise ratio (PSNR). Though PSNR can be a sometimes inaccurate metric of subjective video quality [25, 39], the data in Table 1 are consistent with the earlier extensive subjective testing [35] and are thus expected to be valid predictors. The data of Table 1 represent overall average performance of the various HEVC use cases for a wide range of resolutions from 416x260 to 2560x1600 [13]. It is clear from Table 1 that HEVC substantially outperforms AVC High Profile.

Other results from the JCT-VC report on objective tests are displayed in Table 2. These results provide insight into how HDTV might differ from mobile devices with regard to HEVC efficiency.

Table 1 - Compression Efficiency of HEVC compared to H.264/MPEG4 part 10 AVC

NOTE: Relative Compression Efficiency is calculated as 1/(1 - Bit Rate Savings)

Example Use Case	Encoding Constraint	Bit Rate Savings	Relative Compression Efficiency
Broadcast	Random Access	39%	164%
Video Conferencing	Low-Delay	44%	179%
Contribution	All-Intra	25%	133%

.

Table 2 – Current Compression Efficiency of HEVC for HDTV and Smartphone

Display	Width	Height	Bit Rate Savings Compared to AVC High Profile	Relative Compression Efficiency
HDTV	1920	1080	44%	179%
Smartphone	832	480	34%	152%

Table 2 points out that HEVC's gains for HDTV resolutions are greater than for smartphone resolutions. They are also greater than the average over all random-access results shown in Table 1. These results hint that HEVC may become relatively *more* efficient for emerging resolutions beyond HDTV, such as 4K (4096 x 2048) and Ultra HD (7680x4320). If such proves to be the case, market forces might help accelerate deployment of HEVC as a way for operators and display manufacturer to offer new beyond-HD options to consumers.

It is important to note that both MPEG-2 and AVC improved significantly as they moved from committee to market. Even today, MPEG-2 and AVC continue to become more efficient as competition pushes suppliers to find new ways of improving quality and squeezing bits. The same dynamic is expected with HEVC. It should experience additional improvements, rapidly, when it emerges from the standardization process, followed by long-term, continuous honing through commercial competition. It is common in industry circles to project that HEVC will achieve its targeted doubling in compression efficiency – it is simply a matter of time.

For purposes of our subsequent analysis of HFC capacity and services, we will assume that HEVC will indeed ably achieve its 50% goal when commercially available.

Under the Hood

Some of the AVC efficiency gains were the result of new coding techniques such as context-adaptive binary arithmetic entropy coding (CABAC). Yet a large part of the gains came from making existing tools more flexible. Compared to MPEG-2, for example, AVC provided more block sizes for motion compensation, finer-grained motion prediction, more reference pictures, and other such refinements.

HEVC also gains its performance edge by using newer versions of existing tools. One of the most significant enhancements is that the concept of a macroblock has morphed into the more powerful concepts of Coding Units (CU), Prediction Units (PU), and Transform Units (TU).

Coding Units are square regions that can be nested within other Coding Units in a hierarchical quad-tree like manner to form an irregular checkerboard. The advantage is that smaller Coding Units can capture small localized detail while larger Coding Units cover broader more uniform regions like sky. The result is that each region in a picture needs to be neither over-divided nor under-divided. Avoidance of excessive segmentation saves bits by reducing the overhead of signaling partitioning details. Judicious subdivision saves bits because the details within each terminal Coding Unit can be predicted more accurately.

Prediction Units extend the "just-the-right-size" coding philosophy. Prediction Units are rectangular subdivisions of Coding units that are used to increase homogeneity – and thus predictability – within Coding Units. If a particular Coding Unit encompasses a region of grass and a region of tree bark, for example, an encoder might attempt to arrange the boundary between Prediction Units so it matches the grass-bark boundary as closely as possible. Together, Coding Units and Prediction Units create a quilt of more homogeneous patches that are easier to compress than regions of heterogeneous textures.

Transform Units are also subdivisions of Coding Units. The objective is to position and size Transform Units such that a picture is subdivided into mosaic of self-similar patches when viewed from the frequency domain. One of the dominant visual artifacts in MPEG-2 and AVC is the distortion that sometimes occurs near sharp edges and around text. This artifact is a result of performing a transform and quantization across radically different textures on either side of the edge. In HEVC, Coding, Prediction, and Transform Units work together to more precisely decouple the textures flanking the edge thereby reducing spillover and avoiding the visible defect.

Other coding tools also get a makeover in HEVC. Intra prediction supports many more directional modes to discriminate the angular orientation of lines, edges, and textures more exactly. Inter prediction has improved interpolation filters to yield higher quality motion vectors. And there are less costly ways of sending motion vector information to the decoder. HEVC also gains at least one new kind of loop filter targeted at improving both objective and subjective visual quality.

Not all the enhancements in HEVC are incremental. HEVC will be capable of delivering high-quality video to every conceivable device from the size of a thumbnail to a wall-filling 8k x 4k display in wide-gamut color palette that rivals the natural world. That is an opportunity for unparalleled consumer experiences.

Commercialization -- Profiles & Levels

Compression standards of the caliber of HEVC are complex amalgams of sophisticated algorithms and protocols. In the past, specific subsets of capabilities and features of MPEG-2 and AVC were organized into Profiles with Levels to aid commercial adoption and facilitate interoperability between vendors. It would be unsurprising if HEVC also adopted a family of Profiles, but at the moment the HEVC Committee Draft [37] specifies only Main Profile, which roughly corresponds to AVC High Profile.

Within the HEVC Main Profile, the Committee Draft does specify a number of Levels. Each Level corresponds to a maximum picture size (in terms of number of samples) and maximum pixel rate for the luma component. From these constraints, it is possible to indicate the minimum Level that would correspond to various consumer devices, as we do in Table 3 for smartphones; HDTV on tablets and flat panels at home; and next-generation beyond-HDTV displays.

Table 3 - How HEVC Main Profile Levels Might Correspond to Various Displays

	Example Format	Width	Height	Frame Rate	Minimum Level
Smartphones	QCIF	176	144	15	1
	CIF	352	288	30	2
	480p	854	480	30	3
	QHD	960	544	60	3.1
HDTV	720p	1280	720	60	4.1
	1080p	1920	1088	30	4.1
				60	4.2
Beyond HDTV	4K	4096	2160	30	4.2
				60	5.1
	Ultra HD	7680	4320	30	6
				60	6.1

Note that most smartphones and sub-HD resolutions would be supported starting at Levels 1 through 3, depending on the picture size and frame rate. Note that any Level above the minimum Level could be used. HD resolutions would be supported starting at Level 4. Beyond-HD resolutions would require at least Level 5 & 6 with one interesting exception. Super HD 4k x 2k resolution at 30 frames per second shares Level 4.2 with 1080p 60 frames per second. It may turn out that operators will be able to leverage Level 4.2 in the future to provide consumers with both 1080p60 sports content and Super HD 4k film content (24 frames per second).

Next Steps

The process of earnest creation of HEVC began in 2010 with a Call for Proposals (CfP). There have now been nine JCT-VC meetings in which approximately 200 attendees per meeting created and debated over 2000 input documents. In **February 2012**, JCT-VC issued a complete draft of the HEVC standard called the Committee Draft [37] which will be refined over the coming months. The Committee Draft also serves as a starting point from which to explore development of commercial HEVC products. The Final Draft International Standard is scheduled to be made available in January 2013 for formal ratification.

HEVC is well on its way. And, as we shall see in the next section, it will be an essential component of future advanced video services for cable operators, based on what we are able to project today for service mix, spectral constraints, and likely migration strategies.

TRAFFIC AND SPECTRUM

Dynamics of the Shift to IP Video

While video resolution affecting bandwidth requirements presents an enormous capacity challenge, it is not the only variable driving spectrum use. In addition to bandwidth growth of the video itself, the nature of the traffic aggregate being delivered is changing as well. There are many variables in play, virtually all of which are driving towards increasing unicast delivery of video content:

- More content choice
- Time-shifting
- Trick play expectations

- Network DVR (nDVR),
- Video capable IP device proliferation (tablets and smartphones)
- Shrinking service groups

And, of course, over-the-top (OTT) viewing from web-based content providers is already unicast delivery.

As a result of these shifts, the gains typically afforded by multicast capability, or bandwidth reclamation gains associated commonly with SDV architectures, begin to evaporate. Consider Figure 5 [23]. On the right edge of the curve, we can see by comparing the DOCSIS channel count required for delivery of unicast compared to multicast that for a large group of active users and predominantly linear content, there is significant, exploitable gain. This converts to important bandwidth savings. This has been the lesson of SDV widely deployed in HFC networks today. However, these deployment advantages are based upon the content choice and the size of service groups of the time. Today, as node splits occur, the growing use of a variety of IP clients consuming video, increased choice etc., the operating point on the curve shifts to the left.

The crosshairs in the figure (60% penetration x 60% peak busy hour viewing on a 500 hhp node) represents a reasonable operating point in a system outfitted with 200 HD and 200 SD programs available as switched services. There is clearly much less gain at this point, suggesting only a modest savings in exchange for the complexity of multicast. Some optimization steps may be taken to most efficiently allocate spectrum, but with an eye toward simplicity of architecture as well. This approach is shown in Figure 6 [23]. This diagram illustrates the concept of broadcasting the very popular content to take advantage of programming where simultaneous viewing is likely to occur regularly, optimizing use of bandwidth while maintaining simplicity in the architecture. Analysis in [23] suggests that the vast majority of gain, around 80-90%, occurs in approximately the first 20 programs.

Thus, a combination of broadcast and unicast may be the end result of an IP Video system weighing the tradeoffs of efficiency and complexity. The modest loss of efficiency of “all unicast” in the figure is recovered through the use of a small tier of broadcast services. And, as service groups continue to shrink, there will be virtually no bandwidth efficiencies lost at all. This is illustrated in Figure 5, for example, for the 80 active IPTV viewers.

<u>Next Generation Video Formats: Parallel Characteristics to IP Video</u>

The dynamics commonly associated with 2nd screen viewing may also come to pass in the next generation primary screen video world. There is a large permutation of video formats for mobile viewing, being usurped today by high quality formats. The likely similar dynamic to emerge for primary screens is simply that new formats will get introduced well before other formats are retired. Historically, this would suggest a need to simulcast formats to ensure all customers have their video needs served based on what formats they can support on the TV sets in their home. With more formats arriving, and an overall accelerated pace of change, this could create a bandwidth Armageddon given the nature of the advanced formats relative to bandwidth consumption. However, as we shift into the IP Video world today built around 2nd screen compatibilities, we are developing and deploying tools for discovery and delivery of a large permutation matrix of formats and protocols based on the different capabilities and interfaces of IP client devices.

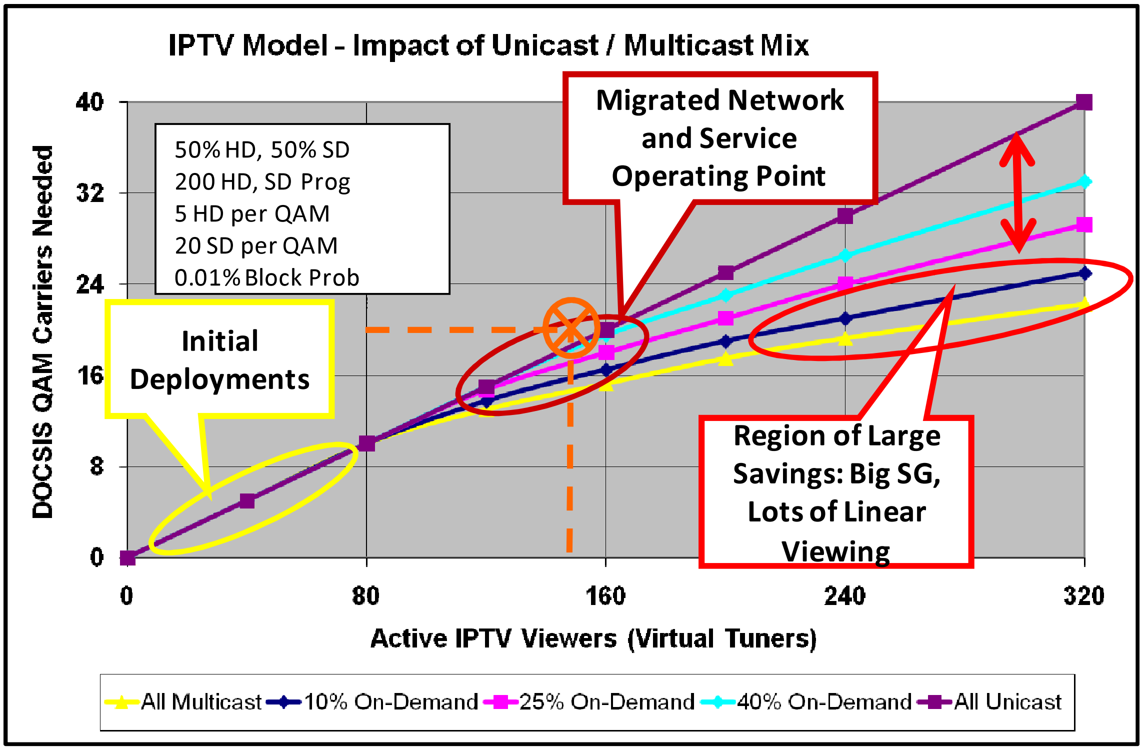

Figure 5 – IP Video Shifts the Spectrum Allocation Methodology

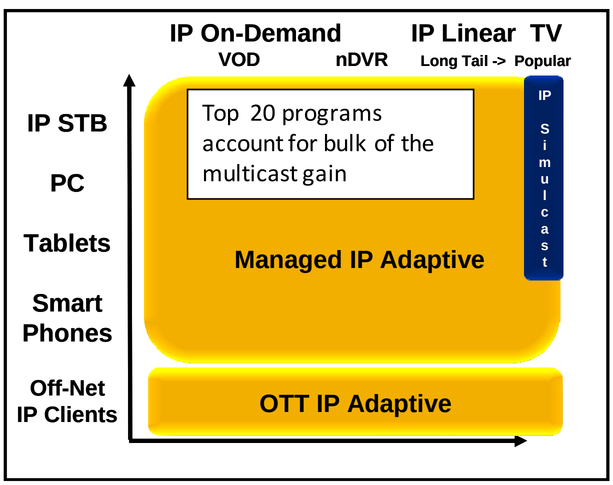

Figure 6 – Optimizing IP Video Delivery

This same dynamic could occur in the future with new high-resolution formats and smart TVs, with the only difference being that the process will take place with respect to discovering and adaptation to *primary* screen capabilities. The intelligence required is being built today to serve those 2nd and 3rd IP screens. By the time, for example, QFHD is a video format scaling in volume, the migration characteristics driving traffic to nearly all unicast will have taken place. As such, primary screen format discovery will be timely for keeping simulcast requirements at bay.

The model that we will assume as we assess the network implications is one of a small set of broadcast (conservatively quantifying with 40 total broadcast programs), with all other traffic as unicast. We will assume that the remaining traffic for video – the video unicast – is inherently captured in the traffic projections as part of 50% CAGR on the downstream. It may, in fact, be precisely what the CAGR engine of growth *is* for IP traffic over the next decade. A contrasting view would be to project HSD growth at 50% CAGR, but add to this video traffic aggregates representative of video service rates of an aggregate [13].

NETWORK IMPLICATIONS

It is quite simple to illustrate a network capacity problem in the face of increasing video quality and resolution, which directly translates into more bandwidth required. In Figure 7 we find the intersection of traffic growth, video services, and time in order to help guide MSO decision timelines. The trajectories moving upward from left to right show a commonly assumed Compound Annual Growth Rate (CAGR) of 50% over a period of ten years offset with two breakpoints over the course of the decade where a (perfect) node split takes place.

While the HFC available capacity in the downstream is over 5 Gbps when considering the highest order modulation profile currently utilized (256-QAM – the yellow horizontal threshold) it is of course not all available to support IP traffic today. The vast majority of today's spectrum is set aside for video services. Figure 7 charts the growth of IP services, but also quantifies the setting aside of spectrum used for video services. These video services that are the moving target that we are looking to quantify here. The bandwidth set aside for video services is subtracted from the 870 MHz capacity to identify the threshold for when the IP traffic would exceed the available spectrum to support it. These thresholds are the horizontal lines on Figure 7.

Four thresholds are shown bounding the available capacity over the course of 10 years. The first, baseline case (red) identifies the available spectrum for data services growth if the video service offering is made up of 60 Analog carriers, 300 SD programs (30 QAM slots), 50 HD programs (20 QAM slots), and 8 VOD slots. The math for this distribution of broadcast and VOD is quite simple: 60+30+20+8 = 118 slots consumed for video services, leaving 18 slots for DOCSIS.

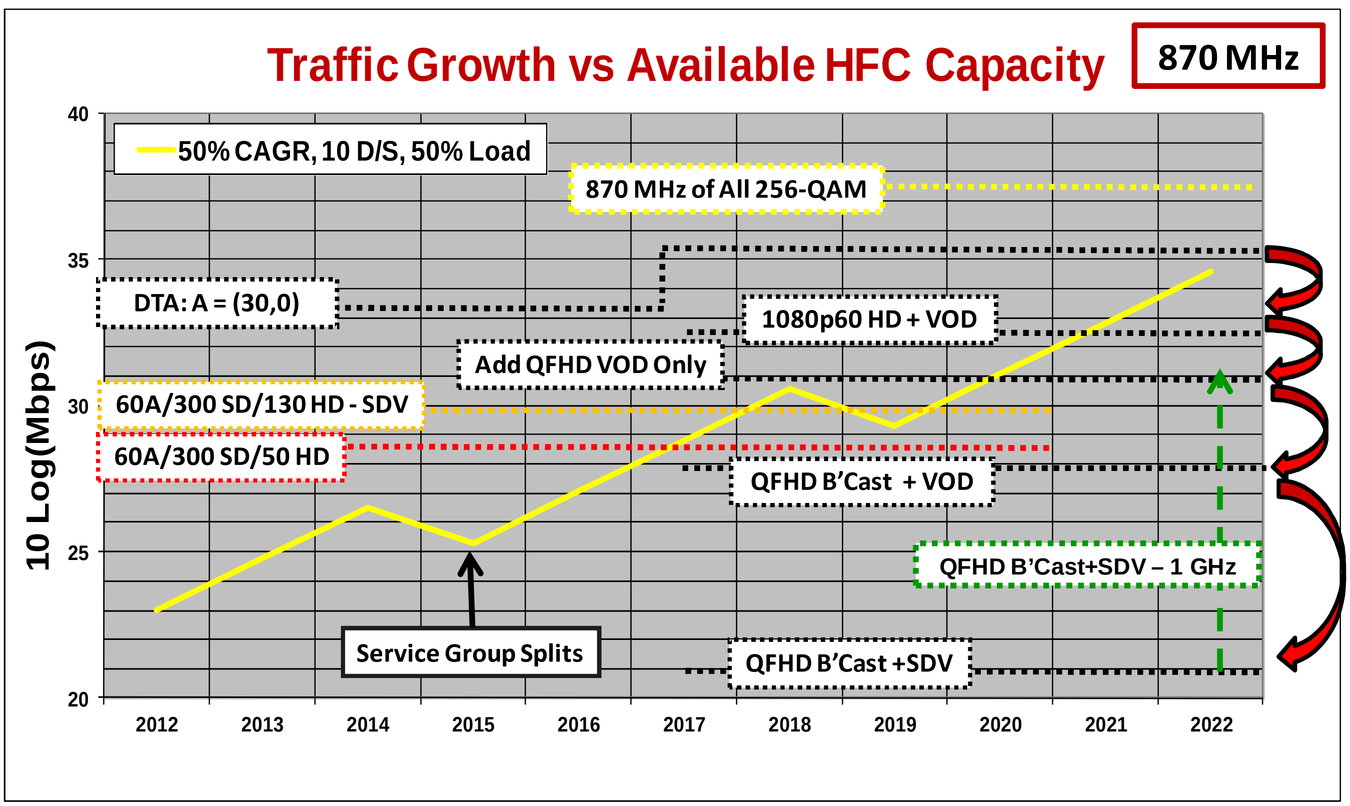

Figure 7 – New Resolutions Project to Massive Spectrum Management Concerns

Under an assumption that today's downstream DOCSIS carriers consume 200 Mbps of capacity (50% peak busy hour usage of 10 deployed downstream slots), then this video service architecture supports IP traffic growth through the year 2016, assuming there is one service group split along the way.

The orange threshold identifies the available headroom for IP growth if Switched Digital Video (SDV) is deployed, and the SDV achieves 3:1 gains for both SD and HD. Also, the HD program count is increased to 130 (modeled after a specific operator example objective). The broadcast tier in this case is limited to 60 Analog carriers, and the top 40 most popular channels offered, which are broadcast in both HD and SD. All other programs are on the SDV tier (about 20 SDV slots). The benefits of SDV are clear in Figure 7. Despite more than doubling of the bandwidth-intensive HD programming, we nonetheless gain new capacity for IP growth.

As powerful as SDV is for reclaiming spectrum, it is only reclaiming QAM spectrum, which is already inherently efficient in delivering digital video. There are further, large spectrum gains available by instead reclaiming spectrum from the Analog carriers through the use of digital terminal adaptors (DTAs). In Figure 7, the implementation is a phased approach. Phase 1 is a reduction of Analog slots from 60 to 30 – the black threshold that extends through 2017. In 2017, it is suggested that Phase 2 kicks in, whereby all Analog carriers are removed. This is the second black threshold, where now well over 3 Gbps has been freed up as capacity for IP growth. This chart and analysis process also identifies the flexibility available in downstream spectrum management. There are many knobs and levers associated with decisions on service mix and use of tools available for bandwidth growth.

Of course, the core issue as new video services evolve is that a 10-year plan demands consideration of these bandwidth-hungry next generation video possibilities. Ten years of tools and projections are encouraging, but the projection is based on video services and technology as we know them today. The plan can quickly implode by considering the capacity when including the integration of new generations of HD.

Four phases of next generation video evolution are identified by the red arrows on the right side of Figure 7. First, consider simply that all of the HD becomes 1080p60 HD – broadcast, SDV, and VOD. It is assumed this format does not require a simulcast phase (existing STBs and HDTVs support the format if it is available to them). The drop in available capacity (the first red arrow on the right hand side of Figure 7) reflects about a lost year of lifespan, all other assumptions the same.

Next consider that a Quad Full HD (QFHD) format is made available on VOD as an introduction to this format in its early days, as capable televisions become available to early adopters. The current VOD allocation remains (1080p60) in this case, so this advanced VOD service is completely additive in terms of spectrum. It is assumed that this format is deployed only using MPEG-4 compression. Nonetheless, as revealed with the second red arrow, we see a larger step downward in available capacity, which now is just over 1 Gbps. Roughly another year of lifespan is lost, all other assumptions the same. Furthermore, this would drive the timing of the second node split for downstream in 2018 if a QFHD VOD tier were to become viable in that time frame.

Now consider the third step, whereby QFHD was used for Broadcast HD and VOD, but not SDV. Note we have not included a simulcast of standard HD, even though it is TBD at this point whether a QFHD format can be "down-resolutioned" to standard HD. Certainly this is not the case in today's televisions or STBs feeding televisions, but it is likely to be a consideration in future iterations. The 4x scaling of standard HD is of course, in part, to make it more likely that systems can take advantage of current processing in the video chain through the simple integer scaling factor of pixels. Not delving into the details of how this might play out, we quantify the impact of a change in the broadcast and VOD to QFHD. The effect identified by the third red arrow is to drop network capacity down to about 600 Mbps, and clearly this is eating into any hope for supporting long-term IP traffic growth.

Lastly, now consider that the SDV tier is converted, but the VOD is not. An example of why this might be practical is that as the IP migration takes place, it might be determined that the legacy VOD infrastructure is not permitted to grow with new MPEG-2 TS based investment. These investments would be made instead in the IP domain, with VOD being one of the first phases of the video services migration to IP. In this case of Broadcast and SDV supporting QFHD as opposed to standard HD, we clearly see, in the form of the lowest black threshold on the chart at about 21 dB (just over 100 Mbps of capacity available for IP traffic, or three DOCSIS channels), the hopeless situation for next generation video without some new ideas and evolutionary approaches to be supported over the HFC network.

To point out a measure of hope that hints at some of the consideration we will account for later in the paper, the upward pointed green arrow shows where this situation would instead fall if there was 1 GHz worth of spectrum to work with. The spectrum freed up by 1 GHz of HFC compared to 870 MHz is about 22 slots, which works out to almost 900 Mbps using 256-QAM. New spectrum is but one tool we will evaluate as a means to enabling the migration of next generation video services

Note also that we have as yet not even attempted to factor in any capacity effects associated with Ultra-High Definition Television (UHDTV) as a potential format.

LONG-TERM VARIABLES: GOOD NEWS – BAD NEWS

We observed in Figure 7 that there was an obvious problem brewing under the assumptions made based on considering HFC architectures, services, and technology, as we know each today.

In Table 4, we begin to make the case for why the situation may not be as dire as these projections. On the left hand side of Table 4, "Losses," we quantify in the decibel language of the projection analysis the potential bandwidth penalty of the new formats, quantified in the row identified based entirely on the resolution difference. Again, it may be determined in practice that the encoding process more favorably compresses the formats than is portrayed in Table 4, but for now we will simply rely on encoding efficiency gains consistent with the average savings attributed to H.264 and H.265 using today's HD format. In each case, this amounts to 50% savings, based on early evaluations of H.265 and our prior discussion on HEVC.

Table 4 – Video & Network Variables: Losses and Gains

Losses	dB	Gains	dB
1080p60	3.00	H.264	3.00
QFHD	6.00	H.265	3.00
UHDTV	6.00	Split	3.00
10-bit	0.97	N+0	9.21
Frame Rate	0.00	Mod Profile	0.97
Total	15.97	VBR/D3	1.55
		Total	20.73
Difference	**4.76**	(All)	
	7.76	(HD to QFHD)	

The conservative assumption for 1080p60 is that it is 2x the bandwidth required of 1080i30. For the purposes of this study, as is generally done in practice as well, we will not distinguish between bit rates of 1080i and 720p although the former is roughly 12% more bits of transport rate.

We consider a 10-bit depth of field for UHDTV, but no additional overhead associated with changes to subsampling. We also do not consider any additional frame rate impacts on transport bandwidth. While scan rates of television rates have increased, and, as discussed, studies [1] reveal that frame rates higher than 60 Hz are perceptible by humans, there appears to be no move afoot to standardize in the market place on anything higher. Additionally, UHDTV is standardized around a 60 Hz frame rate. While research noted above has shown perceptibility by human of up to 300 Hz, it would not be fair to impart new bandwidth associated with new frame rates at this stage, even if they are to take shape. It is intuitively likely that higher frame rates lend themselves to more similarities between adjacent frames. We leave the variable in the chart because we believe that in time, formats will begin to experiment with higher frame rate delivery. We should keep this variable in the back of our minds as a possible wildcard.

The "Losses" when added together in the worst case of UHDTV as the final phase is 15.97 dB, which we will round to 16 dB for discussion purposes.

Now, let's take a look at the "Gains" in Table 4. Some of these we have already observed as "HFC as we know it" gains in our projection chart. We identified service group splits by the traffic growth breakpoints in the chart, which recognized the virtual doubling of bandwidth (ideally) in a typical node split. The average bandwidth allocated per subscriber in the split service group is now twice as much.

We also capture the service group split function here identified as N+0. In this case, we are recognizing that rather than perform further business-as-usual node splits after another round of this expensive activity, an "ultimate" split is executed instead, where the fiber is driven deepest – to the last active. The impact to the average bandwidth made available per subscriber is much greater in this case, with the homes passed per N+0 node assumed to be 40. Note that we identify one split prior to N+0 in Table 4. In the actual timeline-based model we will capture the move to N+0 as an extra split (two total) prior to the migration to N+0. We captured the decibel effect (3 dB) within the N+0 adjustment in the table to match what we will show on the subsequent projection analysis.

Lastly, we applied the benefits of MPEG-4 encoding in introducing QFHD – obviously better than MPEG-2, but also clearly not enough itself to compensate the bandwidth growth. This is intuitively obvious enough, seeing as the MPEG-4 gains do not offset the resolution increase in pixel count. However, it should be pointed out that the 1080p60 case shown in the trajectory of Figure 7 may indeed be offset by the introduction of MPEG-4 to deliver that service. In fact, it is reasonable to consider that 1080p60 as a service does not become a video service offering *until* MPEG-4 is available.

We showed in the Figure 7 a hopeful sign in the form of a different total available spectrum – 1 GHz vs. 870 MHz. However, because our starting assumption of 870 MHz may be optimistic or pessimistic, and because the spectrum expansion discussion is a wide-ranging one, we will address the physical bandwidth component in a subsequent discussion dedicated to spectrum.

We identify three other “Gain” variables – the subsequent generation of encoding, H.265, the use of IP Video delivery using bonded DOCSIS channels, and the opportunity to be more bandwidth efficient in an evolved (i.e. N+0) HFC architecture

<u>High Efficiency Video Coding (HEVC, H.265)</u>

As described, HEVC is in the heavy lifting phase of development and standardization, has as an objective a 50% better bandwidth efficiency of video transport, all while also yielding a higher quality. It appears to be on the track to achieve these targets.

The time-to-market for encoding standards and time-to-scale of advanced video formats follow roughly similar temporal cycles in terms of years. They are not necessarily in phase, but in both cases long evolution cycles have been the norm. As shown in Figure 8, it has been to the case in the past that the encoding gains served to continually drive down the rate of video (all SD for a time), even as slow as the pace of encoding development was. This singular fact explains the rise of over-the-top video. Data speeds raced ahead while video rates continuously dropped, crossing paths about seven years ago.

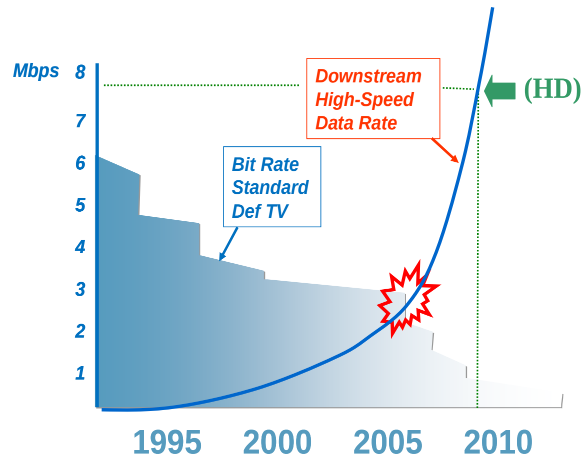

Figure 8 – Compression Meant Video Rates Only Decreased for Many Years

Now, however, demand for more HD has exploded, and display technology advanced significantly as well. It appears that the continuously accelerating pace of technology development will mean that higher quality, better resolution video will proceed faster than the process of standardizing encoding techniques. There is no accelerant to such a process, and arguably the increasingly competitive technology environment could lead to a slower standardization process, with service providers caught in between.

As indicated, early evaluation of H.265 and the conclusions drawn around this work described previously suggests that it will indeed achieve its target objective of 50% savings in average video bandwidth.

<u>IP Video</u>

Legacy architectures are based on simple traffic management techniques that allot an average of 3.75 Mbps per standard definition video stream to fit 10 such streams within a 40 Mbps single-carrier downstream QAM pipe. The heavy lifting of bit rate allocation is done at the MPEG level, whereby video complexities are estimated, and a fixed number of bits in the pipe are allocated to the ten streams under the constraint not to exceed 37.5 Mbps total. The same process plays out over High Definition slots, but in this case only two or three HD streams are part of the multiplex.

The introduction of DOCSIS 3.0 adds channel bonding to the toolkit, which, with the addition of MPEG-4 encoding, increases the stream count by over and order of magnitude relative to the transport pipe size. The net effect is the ability to use law of large number statistics for both SD and HD to the favorable advantage of less average bandwidth. So many independent streams competing for so much more pipe capacity results in a self-averaging effect that yields more efficient use of an N-bonded channel set when compared to MPEG-2 based video over N single channel QAM slots.

Self-averaging suggests that variable bit rate (VBR) streams can be used, recognizing the peaks and valleys will be handled inherently by the statistics (actually a capped VBR). Several prior analyses [16] of DOCSIS-based delivery, taking advantage of favorable statistics of wide channels to better handle the peaks and valleys of video traffic, shows that capped variable bit rate transmission yields a bandwidth savings that can be exploited. We use a 70% scaling as the bandwidth required for VBR-based channel bonded DOCSIS video in comparison to CBR-based single carrier QAM transport.

Fiber Deep Migration

"Business as Usual" HFC migration has been shown to be well-suited to about a decade of video and data traffic growth, without any new or special tools or techniques to accomplish this lifespan [13]. As discussed, 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. Regardless of the name, the architectural implications have two core components: small serving groups - on the order of 20-40 – and the opportunity to exploit new coaxial bandwidth becomes much more straightforward (30 assumed). The lifespan provided by BAU splits will not only make N+0 more cost effective due to RF efficiencies, but it will also leave operators within a stone's throw of FTTP should the need arise as an end state.

An important "side" benefit of an N+0 architecture is that the quality of the RF channel improves dramatically without the noise and distortion contributions of the RF cascade. The result is a higher SNR HFC link in the forward path. Because of this, we then consider more bandwidth efficient modulation formats. In Table 4, we have assumed that 1024-QAM will be readily accessible in such architectures, and in particular if new FEC is also implemented.

Finally, the removal of all RF amplifiers in the plant leaves only taps, passives, and cabling between node and subscriber, a much simpler scenario for flexible and expanded use of new coaxial bandwidth. Prior analysis [12] has shown how 10 Gbps (GEPON) and higher downstream capacities become conceivable in this architecture.

As fiber penetrates deeper into the HFC architecture, ultimately perhaps landing at N+0, the possibility of exploiting more bandwidth efficient modulation profiles exists, especially if the forward error correction (FEC) is updated from J.83 to modern techniques with substantially more coding gain. Here, we assume 1024-QAM supplants 256-QAM, for 25% added efficiency [15].

The dB Balance Sheet

Adding up the "Gain" side of Table 4, we find a total of about 20.7 dB, vs. 16 dB of "Loss." The encouraging information here is that this implies that, in principle, we can convert our current HD lineup fully to UHDTV and this would still be supported over the HFC network. All else equal, HFC lifespan would not be compromised in the face of IP traffic growth – the trajectory thresholds would not drop. This is so because the net of the gains and losses is a positive 4.7 dB. Thus, the thresholds would actually rise. Better yet, if the only format we bother concerning ourselves for business planning purposes is QFHD, then we have another 3 dB or headroom in our net gain.

The flaw in this good news story, of course, is that by the time we are considering QFHD, the IP CAGR is already threatening video service thresholds. We are at or near the end of the ten year window of migration. We are looking to extend HFC lifespan *beyond* this decade to the next while introducing these advanced video services. The excess gain can be viewed as available overhead for a simulcast transition. Based on Table 4, there is 4.8-7.8 dB to work with as part of enabling the possibility. While the services our transitioning, the IPV evolution is taking place, and the network is undergoing BAU migration, there are some "Not Business As Usual (NBAU)" evolutions expected to be taking place as well related to spectrum and architecture.

We will use Table 4 and these NBAU evolution factors to extend the projection through another decade, and draw conclusions on the intersection of video evolution, traffic growth, capacity, and the role of CAGR.

NEW SPECTRUM CONSIDERATIONS

Now let's consider the spectral aspects that were discussed in the last section, but not quantified in Table 4.

Figure 9 illustrates the anticipated spectrum migration of the HFC architecture long-term. A key driver discussed in great depth in [13, 14] is the necessity of operators to do something to address the limited upstream for the future. There are no easy answers to new upstream spectrum, and this figure describes the most effective approach and best performing from a modulation efficiency and flexibility standpoint, and which also yields the most efficient use of spectrum long term. The later is perhaps *the* key long-term primary objective for HFC spectrum evolution.

Because of the reasons outlined in [13] and [14], 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 9, where some level of asymmetry consistent with what supports the downstream/upstream traffic ratio, will remain. No matter where the Frequency Domain Diplex (FDD) architecture lands in terms of diplex split, it is most assuredly going to yield a downstream capable of over 10 Gbps, and an upstream capable of over 1 Gbps.

While Figure 9 represents the most likely evolution scenario, other versions may come to pass. However, for any implementation, it is virtually guaranteed that the 10 Gbps/1 Gbps targets, at least, will be achieved. We will use this certainty in our projections in determining the ability of the evolved HFC architecture to deliver next generation video service in the face of continued growth in high-speed data services.

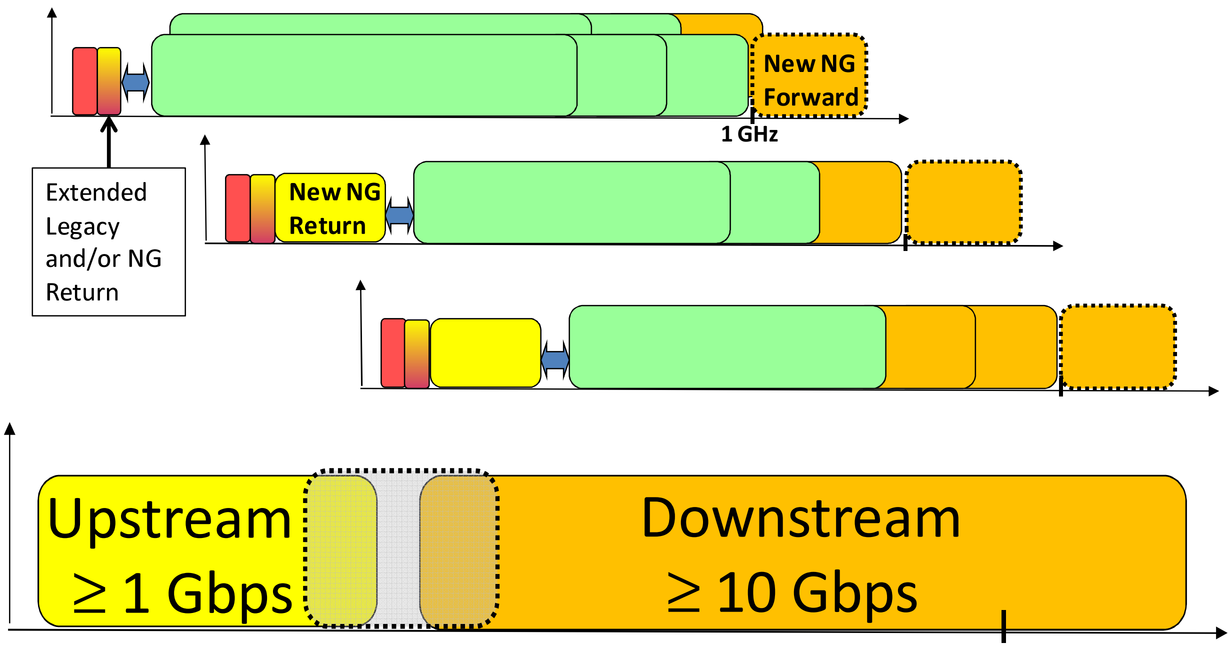

Figure 9 – Probable Evolution of the Cable Spectrum

PUTTING IT ALL TOGETHER

We now revert back to our original problem of capacity growth, and extended timeline of Figure 7 to account for the introduction of new generation of video formats. Beginning with Figure 10, we take into account all of these factors of video bandwidth growth and capacity preservation, placed in the context of HFC lifespan.

Video Service Delivery Assumptions

As we discuss video formats such as QFHD and UHDTV, it is reasonable to assume that HEVC has a key role, that fiber deep migration has continued to take place and is quite far down the path, and that the IP Video transition is in full swing, and possibly even complete. It is also reasonable to suggest that *unless* these evolutions take place, it is not practical to consider new tiers of advanced video services. Under this assumption, QFHD and UHDTV only become service in the cable network over IP, and only when HEVC is available in products for deployment.

The transition model is, of course, critical, as every new format introduces a period of simulcast if a service represents a broadcast. Conversely, in a full IP transition and a fully unicast architecture, the resolution and format become part of control plane and discovery. There is no wasted simulcast bandwidth, just any bandwidth penalty paid if the migration of video service delivery from the "legacy" efficiencies of broadcast to a dominantly unicast architecture is not properly managed (see Figure 5 and 6).

The Intersection of Video Services and Traffic Growth

Now let's consider Figure 10. Figure 10 is a modified Figure 7, extended through the end of the next decade, managed with an N+0 migration, and accounting for various capacity enhancing techniques discussed above.

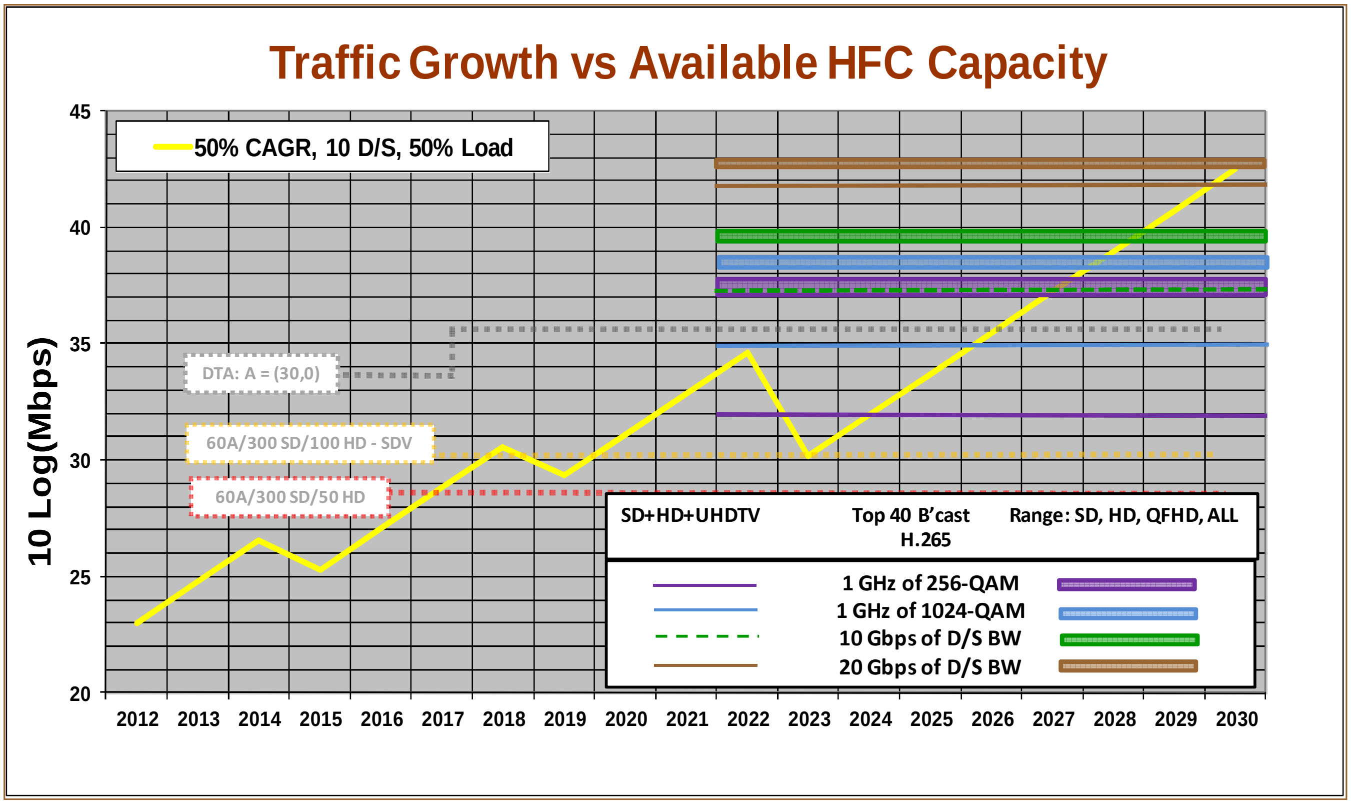

Figure 10 – Next Gen Video, Traffic Growth and HFC Capacity Limitations

The CAGR description is no different than Figure 7, it only goes on for longer, and sees a steeper breakpoint in 2022, representing the final phase migration to N+0. The Figure 7 thresholds are shown, in faded form, for reference against the cases to follow.

The legend at the bottom right is described as follows.

In all cases, we are talking about thresholds set by having a *static IP broadcast of the Top 40* channels. We are therefore taking advantage of IP video efficiencies, only as we know them today and previously identified in Table 4. Recall, we indicated that for a 200/200 lineup of SD/HD, then the top 20 programs would amount to 80-90% of the multicast gain in a switched IP system capable of multicast. The conclusion from that analysis was that a simplified, near-optimal architecture might instead be a mix of full broadcast and unicast, recognizing that all dimensions of network and service evolution are towards more unicast. From that, we have conservatively used a Top 40 program broadcast, which essentially would account for all of the multicast gain. At 40, it will likely come at the expense of some inefficiency of spectrum use versus multicast, but we prefer to err on the side of setting aside more spectrum for the purpose of a conservative analysis.

Also, because we are ultimately after the second-decade phase of HD evolution, we implement the next phase of compression evolution, HEVC, in calculating the long-term thresholds.

Four cases of available capacity are identified:

1) 1 GHz of spectrum carrying all 256-QAM, or 6.32 Gbps (purple)
2) 1 GHz of spectrum carrying all 1024-QAM, or 7.9 Gbps (blue)

3) A 10 Gbps downstream, in light of our prior conversation about the evolution of cable spectrum and key objectives (green)
4) A 20 Gbps downstream – enabled only through an N+0 architecture with a further extended use of coaxial bandwidth, requiring additional plant evolution of the passive architecture, including tap changes (brown)

Four cases of video formats are also analyzed. However, three of them fall close to one another in net capacity impact, and are lumped together in a "range" identified by a *rectangle* of the associated color on Figure 10. The fourth, most burdensome case is, not surprisingly, that which includes the introduction of UHDTV under the bandwidth assumptions we have identified previously – 160x the bandwidth requirement of the SD resolution and format. These UHDTV cases are identified by *lines* of the associated color – note that the green, 10 Gbps line is dashed, only because it overlaps the rectangular threshold range of the 256-QAM case.

The three cases in proximity whereby a rectangle is used to identify the threshold range are (in each case a simulcast of the Top 40):

1) SD + 1080p60
2) SD + 1080p60 + QFHD
3) SD + QFHD only

The latter, for example, makes sense if we consider that the integer relationship of formats (4x scaling of pixels) makes for the potential that next generation QFHD screens are also capable of displaying a "down-res" to 1080p60, or the STB/CPE function is capable of performing this function for the television. The range of remaining capacity in these three cases seems intuitively very close, and in fact is always within about 1 dB. This is a product of three things:

- Large capacity made available by all-QAM to 1 GHz, at least
- HEVC whittling down SD and 1080p60 rates by a factor of one-quarter
- The nature of the chart, based on nonlinear CAGR, is decibel units which tend to compress large numbers, which is illustrated by recognizing we are quantifying the impact of 18 years of aggressive compounding of traffic.

Let's examine what Figure 10 reveals.

First, consider UHDTV as a format that is mid-to-late next decade in scale at the earliest. It is not realistically able to be supported by HFC, at least under the assumptions we have used here. Even the most favorable of evolution deployments shown here – 20 Gbps of downstream capacity – suggests that persistent CAGR coupled with this broadcast video service runs out of room before the end of the decade. The vast majority of the bandwidth is the UHDTV itself, so eliminating the simulcast component is negligible to this conclusion.

By contrast, if we look at the QFHD scenarios, and view this as a format eligible at the end of this decade, then even the least capable case of 1 GHz of 256-QAM bandwidth offers nearly a decade of support for this scenario, with a range reaching exactly to the end of the next decade (2030) before a threshold breach of HFC capacity. This bodes well for the ability of tools available – just as we understand them today – to manage through an aggressive combination of video service evolution and persistent CAGR of IP traffic. It remains to be seen if this form of the evolved HFC network is the most cost-effective approach to enabling this service mix. But, it is surely comforting to know that the possibility exists to support such services with a 2012 understanding of technology, recognizing in addition the long time window of observation in which to adapt strategy and technology accordingly.

Note, of course, that since we have used 10 Gbps and 20 Gbps and not QAM calculations, these threshold apply equally to any access network that would set aside IP bandwidth for 40 channels as described herein. However, since other architectures may be full multicast, a broadcast adjustment (removing this lost capacity) might be in order for an accurate comparison. This is quite easy to accommodate by noting that 10 Gbps is simply 40 dB on Figure 10, while 20 Gbps is 3 dB higher at 43 dB. It is clear that there is very little difference in lifespan implied between these thresholds and those with broadcast allocations in this stratosphere of bit rates and continuance of CAGRs.

<u>Settling of CAGR</u>

In Figure 11, we show a modified case, whereby the assumption is made beyond this first decade that CAGR *decreases* to 32%. We chose this settling of CAGR at 32%, such that the net CAGR for the period through the end of the next decade is an 18-year average CAGR of <u>40%</u>.

The logic behind this assumption is that we have seen this aggressive march forward of CAGR driven primarily by over-the-top (OTT) video services. In the model developed here, we are already allocating spectrum for most popular video services, and thus using video also as a driver for CAGR could be considered double counting, at least in part (the most-watched part) of this phenomenon. In addition, the vast history of CAGR growth has been around *catching up* with our ability to download and/or consume media – audio, then video. Once these media consumption appetites are satisfied, then it is possible that a CAGR settling will take place, with limits set by behaviors and eyeball counts [11]. Of course, it may simply be replaced by as-yet-to-be-determined non-media consumption applications, or altogether different kinds of media consumption that is bandwidth-busting, such as volume displays. That, however, seems beyond even the extended time frames we are evaluating here.

The above reasoning was completely qualitative, and it may in fact turn out that aggressive 50% CAGR persists indefinitely, or possibly increases. Nonetheless, because of the ramifications of long-term CAGR variation, we thought it useful to show this perspective, and that an 18-year average of 40% CAGR was a reasonable amount of settling to consider. Note that only at the year 2030 exactly would the 40% average and the 50%-32% model meet. The trajectories along the way getting to those points will, of course vary.

Now let's evaluate what Figure 11 below says about video services evolution, capacity, and time.

First, observe now that *every* QFHD case indicates a lifespan of the network *through* the end of the next decade, even the 1 GHz, 256-QAM only case. This is a very powerful statement about the impact a settled CAGR may have on the support of advanced video services. It is also a reminder about the dramatic mathematical and planning implications of 18 years of compounding.

For UHDTV, there still does not appear to me much hope for a lasting solution to broadcast support, under what seems like the reasonable assumption that it does not make a large-scale service appearance until 2025 or beyond. The best case scenario in Figure 11 only suggests that 20 Gbps of network capacity covers the UHDTV scenario plus traffic growth into 2032-2033, which is then very shortly after it would have been introduced.

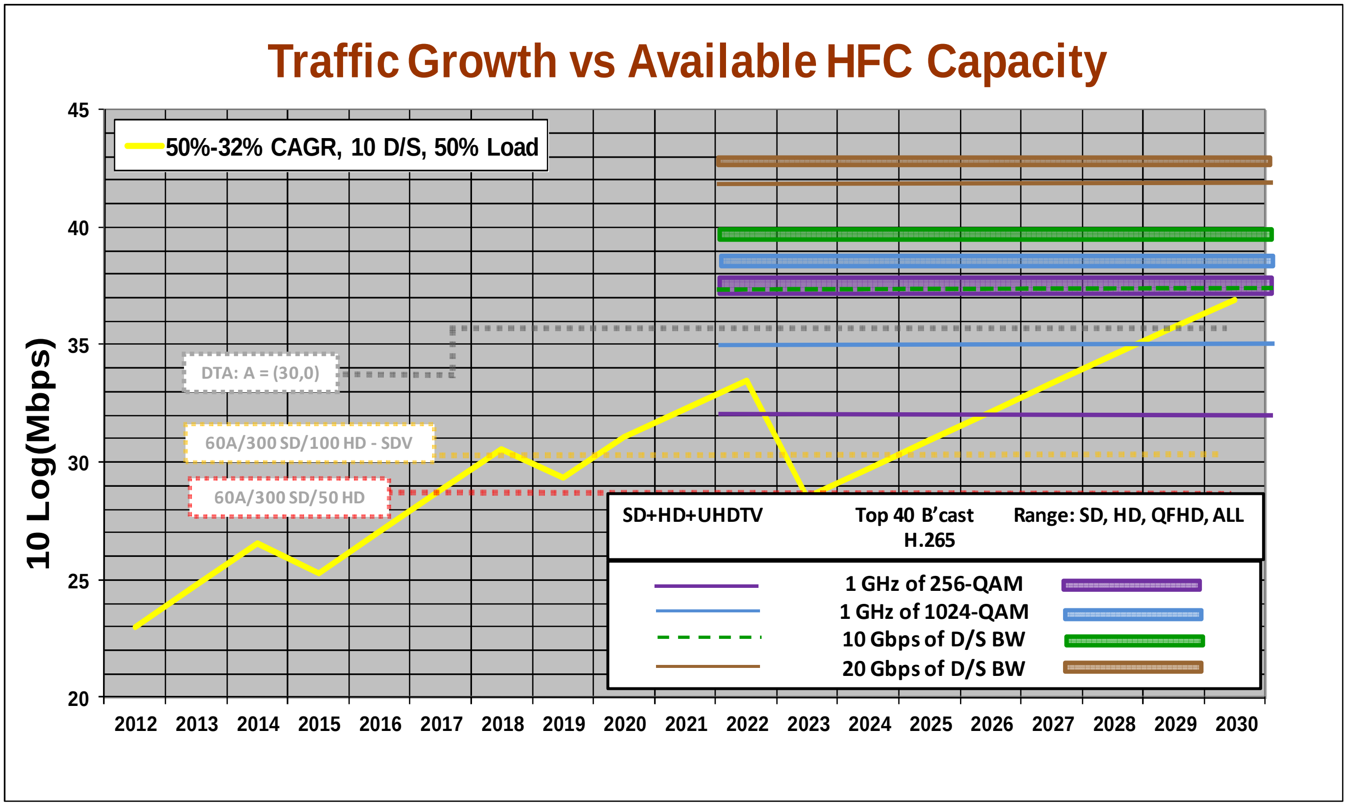

Figure 11 – Next Gen Video, Traffic Growth + CAGR Settling, and HFC Capacity

Conversely, this conclusion might more optimistically be stated by noting that HFC that manages a capacity of 10-20 Gbps *can* clearly support an early phase of UHDTV experimentation and deployment, and provide some cushion of years over which its significance as a scalable service can be evaluated. Does it become a niche scenario, where a very select number of channels become part of a programming lineup, much like 3D is today? For a mid-2020 time frame of experimentation, there are enough years of support in an early, modest phase of deployment where these kinds of questions can be asked and answers provided. These answers can then be used to guide a phase of network evolution, such as Fiber-to-the-Home, if scalability of the service is required. Or, it may lead to the conclusion that UDHTV is not an every-household type of consumer service, but associated with, for example, the penetration of home theatre-type owners. If so, it likely remains largely on the IP unicast service tier, and never become a broadcast scenario to worry about. Though, if this latter situation comes to pass, then this could exactly be the kind of thing that keeps CAGR chugging at 50%, while this model reflects the 18-year, 40% average case.

There are clearly many interrelated variables to consider and scenarios to quantify. Our assessment of the results leave inevitably to the conclusion that these projections are best viewed as living documents, and must be periodically re-assessed for the validity of the assumptions as trends and service mixes evolve over time. Advantageously, though, the projections indicate there are valuable windows of time near term, and again in the long term as efficiency improves. These windows offer the opportunity to observe and make methodical decisions to manage the evolution, without the pressure of an urgent congestion problem on the horizon.

THE EYES HAVE IT

While developing HEVC, compression science was not standing still elsewhere. Recently, a new technology – Perceptual Video Processing (PVP) -- was incorporated into broadcast encoders to improve the efficiency of both MPEG2 and AVC significantly. PVP technology [20] leverages the biology of *human vision* itself to enhance the encoding process. It can be thought of as a compression co-processor. Performance improvements typically range from 20% for moderately-easy-to-encode content to up to 50% for hard-to-encode content. Given the close familial resemblance of HEVC to its predecessors, it's quite possible that PVP could grant similar bonus improvements on top of HEVC's innate high compression efficiency as it has for MPEG-2 and AVC [6, 34, 35].

Signal Processing and Human Vision

Perceptual Video Processing (PVP) technology is an encapsulation of design principles that are thought to be at work in the visual system based on decades of research into the biology of human vision [2, 3, 4, 7, 9, 20, 24, 26]. Though biological in origin, these design principles are rooted in concepts that are familiar to signal processing engineers, namely, the ideas of noise reduction, signal estimation, and error signals. What is unique is that PVP is based on a model [21] of early visual signal processing, which has the following key components:

- Vision is tuned to the scale-invariant statistics of natural images [8]
- First stages of visual processing act as optimal filters designed to minimize the impact of noise
- A second stage of processing makes an estimate of the error associated with the first stage and uses that error signal to self-adapt to changing lighting conditions
- The output stage of processing is a coded form of the error signal, which can be thought of as a visual map of statistical uncertainty associated with the estimation process.

A key insight is that statistical uncertainty equals perceptual significance. The output error signal – the "uncertainty" signal -- highlights two kinds of information:

1. Image features that are uncertain because local correlations in the image are as likely to be attributable to noise as to actual variations in the signal. These are the features that are likely to be ambiguous from a signal estimation point of view and thus may require more attention.

2. Image features that contain local correlations that deviate from statistical expectations associated with natural scenes. In some sense, these are the "unexpected" correlations that might be worthy of closer inspection.

The notion that the output of early vision correlates with local statistical uncertainly provides a potential clue about higher-level perception and visual behavior. Eye-tracking and saccades, for example, might be considered behaviors intended to spend more time inspecting areas of high uncertainty to minimize overall uncertainty. Similarly, areas of high activity in retinal output might correlate with areas of high perceptual significance because they are the most suspicious in terms of statistical expectations – this is a clue that it may be worthy of special attention.

This model of the early visual system might also provide a context for understanding why edges are perceptually significant. According to the model's key components, edges are not perceptual important because they are edges, rather because they are localized correlations that deviate from the global expectation of scale invariance and thus require longer inspection to reduce uncertainty. It is not in fact the edge that has maximum uncertainty, rather it is the area around the edge, which itself might provide insights into the fundamental nature of perceptual masking and Mach bands – the illusion of heightened contrast near edges.

<u>The Engineering View of Retinal Processing</u>

The signal processing described occurs through the biological processing of the retina. The retina is made up of specialized cell layers, and each has a specific task. These can be classified as follows [20]:

Photoreceptors – The rods and cones we learned about in primary school health class. Photoreceptors are the first line of processing, are very densely aligned, and convert light (photons) into neuroelectrical signals.

Horizontal Cells – This second stage of processing cells collect the output of the photoreceptors and share them with adjacent horizontal cells as kind of a spatial low-pass filter operation on the discrete photoreceptor inputs.

Bipolar Cells – In the third stage of processing, bipolar cells collect both photoreceptor and horizontal cell inputs, and essentially acts to subtract the photoreceptor cell inputs, performing a differentiator type of mathematical operation.

Amacrine Cells – Bipolar cell inputs are received by amacrine cells, which come in different types. One important type acts as an electrical rectifier and gives a measure of the mean activity in the bipolar layer. A second type provides feedback to the first two layers to adjust their response properties according to this mean activity observed.

Ganglion Cells – The final stage of retinal processing, these cells take input from both bipolar cells and amacrine cells, and process and package them for delivery over the optical nerve to the brain.

Figure 12 illustrates the visual processing stages as a signal processing operation, described using tools analogous to functions common in signal estimation applications [20].

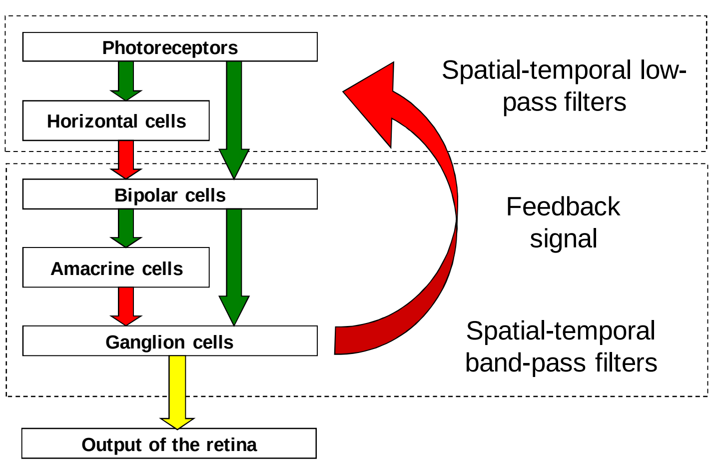

Figure 12 – Visual Cells as Signal Processing Functions

PVP Technology

Considerations for the biology of vision has proven to be very effective in improving compression efficiency in professional broadcast encoders. The key design principles have been extended to encompass space, time, and color and collected into a set of tools and software and hardware implementations collectively referred to as the Integrated Perceptual Engineering Guide (IPeG™). PVP is a particular commercial implementation of IPeG designed to operate in real time to reduce compression entropy and improve predictability in coding.

Internally, PVP identifies features in video that are likely to have high perceptual significance and modifies those features to reduce the number of bits required while preserving video quality. In its first commercial incarnation [20, 38], PVP performs two noteworthy complimentary operations: 3-Dimensional Noise Reduction (3DNR) and Adaptive Detail Preservation (ADP). The 3DNR operation is a combination spatial/temporal nonlinear adaptive filter that is very effective at reducing random noise in areas the eye may not notice. The ADP element preserves visually important detail and attenuates quantization noise, impulse noise, stochastic high-contrast features, and other hard-to-compress detail difficult for the eye to track.

An example of PVP used to improve compression efficiency for statistical multiplexing is illustrated in Figures 13 and Figure 14. The central concept in statistical multiplexing (aka "statmux") is that more and better channels can be delivered over a limited bandwidth by allocating bits intelligently between the various channels that comprise a statistical multiplexing pool. Channels that are easy to encode at a given point in time are given fewer bits than channels that are hard to encode. This traditional "statmux" operation is illustrated in Figure 13.

Using PVP, this operation is modified with this additional intelligent processing as shown in Figure 14. The statistical multiplexer still does its bit rate allocations, as always, but it now does so based on an enhanced set of inputs from the IPeG processor. PVP improves statistical multiplexing by selectively reducing the greediness of hard-to-encode channels in real time. High compression entropy means more bits would be needed to achieve a target video quality. Low compression entropy would require fewer bits to achieve the same video quality. PVP preferentially reduces the entropy of hard-to-encode features thereby making tough content kinder and more generous neighbors in the pool.

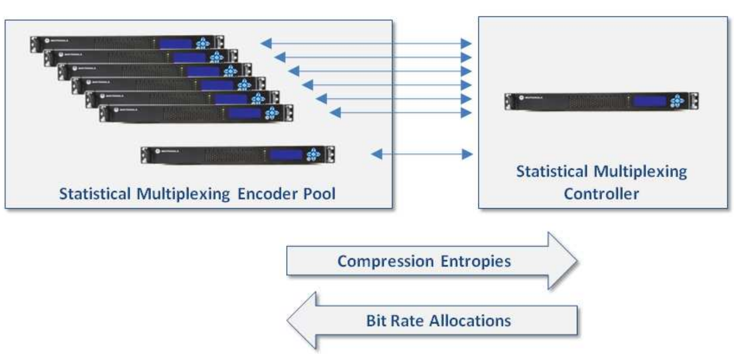

Figure 13 – Traditional Statistical Multiplexing

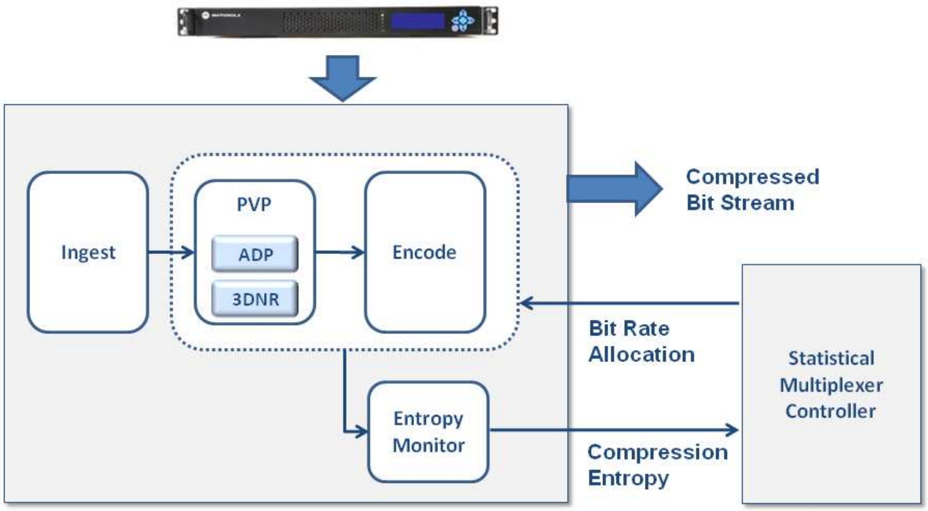

Figure 14 – PVP: Perception-Guided Adaptive Modification of Compression Entropy

An example of the graded impact of PVP on compression entropy is shown in Figure 15. Note that the relative impact of PVP is largely independent of the operational bit rate, which could prove to be a useful feature in statistical multiplexing pools that contain premium channels with higher targeted operating bit rates than other channels in the same pool. The data shown in Figure 15 are typical of moderate-to encode and difficult-to-encode broadcast content. The actual reduction in compression entropy may be optimized for particular use cases by adjusting the strength of PVP from weakest to strongest.

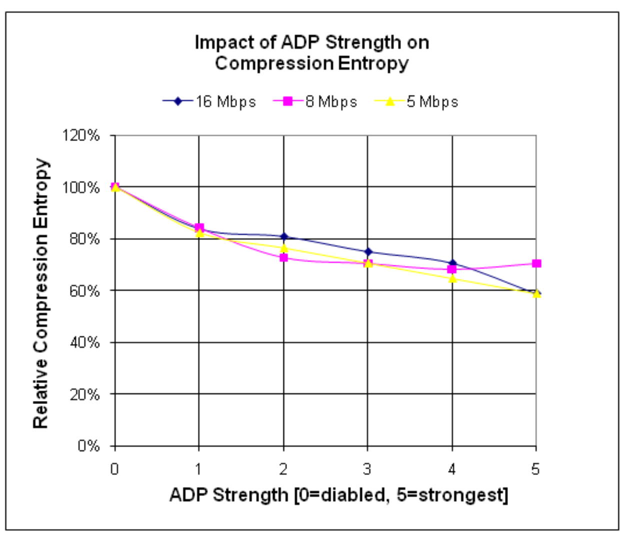

Figure 15 – PVP Reduces Compression and Can be Tuned to Requirements

Complementing HEVC

One of the key advances of HEVC is "just-the-right-size "processing in which each Coding, Prediction, and Transform Unit is sized precisely to capture the self-similarity within the picture detail they encode. It is without question a highly efficient way to squeeze bits -- but it's *not* the way the eye sees.

There are two key questions to examine to predict the impact of PVP on HEVC efficiency:

1) Would PVP enhance predictability and thus promote regions of self-similarity that can be captured efficiently by HEVC Units?

2) Are HEVC's "just-the-right-size" Coding, Prediction, and Transform Units also "just-the-right-size" for the natural scale of vision? If they are, then we would expect PVP to have less of an impact for HEVC than it does for AVC and MPEG-2.

The first question is straightforward, and the answer is *yes*. PVP nudges video towards statistics that would be expected of clean natural scenes when those nudges would not be very noticeable. In other words, the PVP promotes predictability and regional self-similarity. It does this by reducing unpredictable random noise and slightly modifying stochastic high-contrast features that are "unexpected" as described previously. On this basis, we would expect PVP to improve HEVC's innate compression efficiency to approximately the same extent that PVP improves AVC and MPEG-2 efficiency.

The second question is a bit more involved. Getting a handle on the natural scale of vision entails comparing the size of retinal images to the resolving power of the eye.

The visual angles subtended by various kinds of displays are listed in Table 5. The size of the visual field depends on the physical size of the display and its distance from the viewer. For QFHD (4k) and Ultra HDTV, we use the dimensions of recently announced displays [29] and predict that comfortable viewing distances will be only moderately larger than they are for traditional HDTV.

Table 5 -- Expectable Visual Angles for Various Display Types

Display Type	Format	Resolution		Dimensions (inches)			Viewing Distance (inches)	Visual Angle (degrees)
		Horizontal	Vertical	Diagonal	Width	Height		
Smartphone	QHD	960	544	5	4	2	12	19
Tablet	1080p	1920	1080	11	10	6	16	35
HDTV	1080p	1920	1080	55	48	27	76	35
Super HDTV	4K	4096	2160	70	62	33	88	39
Ultra HDTV	8k	7680	4320	85	74	42	90	45

The fovea of the retina sees the central 2 degrees of the visual field with high acuity [17]. It is the part of the retina with the greatest resolving power. We watch television by continually moving our eyes around to bring our fovea in line with particular features one after the other. Our brains integrate this sequence of focal observations into a unified seamless experience.

In Figure 16, we examine the size of the foveal image relative to the size of the visual field subtended by various display types. Our fovea spans only about 1/10th the width of a smartphone display, which means we must still move our eyes about even for the smallest display type. For 1080p and finer resolutions, our fovea sees at any moment in time only a disc of pixels having a diameter about 5% of the width of the whole display. It is worth noting that area of the disc comprises less than 1% of the total pixels in the display. We only see that small 1% of the display in detail at any instant. Research into bit rate reduction of video in other circles has been around trying to figure out how to take advantage of the fact that so little of a screen is actually processed at any given instant [5].

In Table 6, we quantify these relationships across a range of display types. An important insight comes about when we analyze the number of physical pixels seen by the fovea as a function of display size and resolution. A disc about 100 pixels in diameter contributes to foveal vision for smartphones, 1080p tablets, and HDTV. If brightness and contrast were put aside, the equal density of pixels would suggest that we would notice about the same level of visual detail – and same level of compression artifacts – on smartphones and tablets as we would see on HDTV when viewed from normal distances. Visual details and artifacts would likely be less noticeable for 4K and UHDTV because they would be 2-3x less magnified in the foveal image according to the pixel diameters shown in Table 6.

Figure 16 – Size of Projected Foveal Image (yellow) vs. Display Type

Table 6 -- Size of Foveal Field of View Relative to Size of Coding Units

Display Type	Format	Size of 2-degree Foveal Field of View						
		Percent of Screen Width	Pixels (dia.)	Macroblocks or Coding, Prediction, and Transform Units				
				4x4	8x8	16x16	32x32	64x64
Smartphone	QHD	11%	101	25	13	6	3	2
Tablet	1080p	6%	111	28	14	7	3	2
HDTV	1080p	6%	110	27	14	7	3	2
Super HDTV	4K	5%	211	53	26	13	7	3
Ultra HDTV	8k	4%	343	86	43	21	11	5

The scale of MPEG-2 and AVC macroblocks and sub-partitions relative to the size of the foveal image for smartphones, tablets, and HDTV is illustrated in Figure 17a. The homologous HEVC Coding, Prediction, and Transforms Units are depicted in Figure 17b. We noted previously that the fovea covers only a tiny fraction of a display screen at any moment. Smaller yet are macroblocks, sub-partitions, and HEVC Units. Even the Largest Coding Unit (LCU) presently allowed in HEVC (64x64) is significantly smaller than the fovea's field of view.

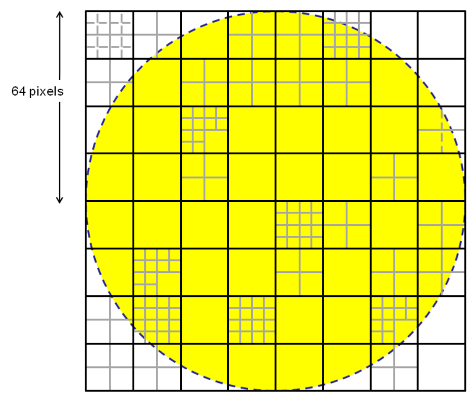

Figure 17a – MPEG-2 and AVC Macroblocks (dark) and Sub-partitions (light) Relative to Foveal Image (yellow) for smartphones, tablets, and HDTV

The homologous HEVC Units for 4k and UHDTV are illustrated in Figure 17c. The difference between HDTV and beyond-HD is a matter of visual scale. The foveal image of a LCU becomes 2-3 times smaller in 4k and UHDTV, respectively, compared to HDTV. Other smaller HEVC Units become visually diminutive, and the smallest 4x4 HEVC Units become tiny.

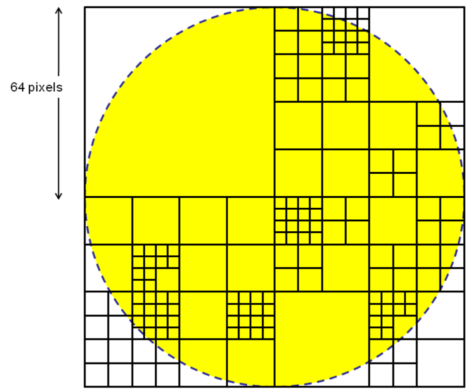

Figure 17b – HEVC Coding, Prediction, and Transform Units Relative to the Foveal Image for smartphones, tablets, and HDTV.

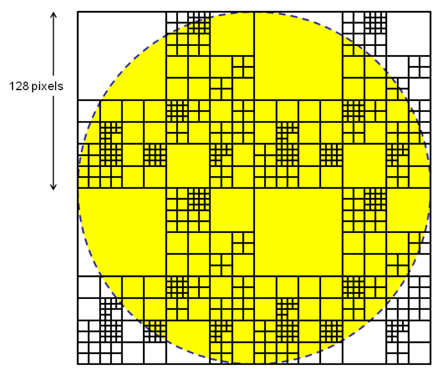

Figure 17c – HEVC Coding, Prediction, and Transform Units Relative to the Foveal Image for 4K and Ultra HD (note the relative size of the Units are smaller than in Figure 17b)

HEVC and AVC use rectilinear segmentation. The specific architecture is different, but the motivating philosophy is the same. More important, the visual scale of the rectangular segments is not dramatically different. HEVC provides a few larger block-size options that are better able to isolate regions of self-similarity without over segmentation, but those block sizes are still smaller that the fovea's field of view.

We can conclude from the above analysis that HEVC Units are, in fact, not always visually "just-the-right-size." Like AVC macroblocks and sub-partitions, HEVC Coding Units will have discrete boundaries within the foveal field of view even when encoding video that is visually smooth across the fovea. Compression artifacts tend to gather around discrete boundaries because those are the places that prediction is weakest. When those boundaries lay within the retina's high-acuity foveal field of view, they will be noticed. HEVC would need larger Largest Coding Units (LCU) to prevent over segmentation of the foveal image and meet the "just-the-right-size" visual ideal. For of smartphones, 1080p tablets, and HDTV the LCU would need to be at least 128x128. For 4K and Ultra HD, LCU would need to be at least 256 x256.

<u>PVP and HEVC Together</u>

Given the overall similarity of HEVC and AVC in terms of coding philosophy and visual scale, we project that PVP will improve HEVC coding efficiency to much the same extent that is improves AVC and MPEG-2 coding efficiency. HEVC's intrinsic compression efficiency is reported in [6, 34, 35]. Relative bit rates expected are listed in Table 7 and plotted in Figure 18. The impact of PVP is very content specific. Nonetheless we have found that PVP provides an overall average bit rate savings of ~20% in national-scale commercial deployments. We use that value in Table 3 to calculate the benefit of PVP to HEVC.

Table 7 -- Expected Bit Rate for Various Coding Modes and Display Types

Coding Method	Expected Bit Rate (Relative to AVC alone)		
	Smartphones	1080p Tablets & HDTV	4K & UHDTV
AVC	100%	100%	100%
AVC + PVP	80%	80%	
HEVC	66%	56%	50%
HEVC + PVP	53%	45%	40%

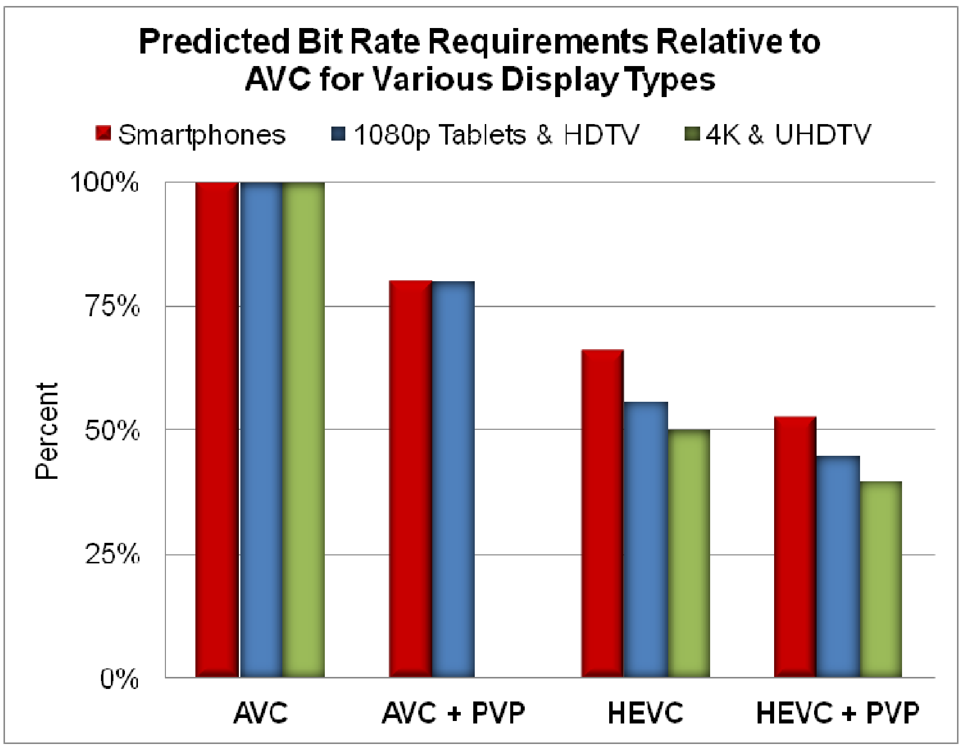

Figure 18 – Projected PVP Efficiencies Bit Rate for AVC and HEVC vs. Display Type

We can take this new knowledge and apply it to the prior figures that quantify traffic growth impacts. Figure 19 does so for the last case evaluated previously (Figure 11, 18-yr average CAGR of 40%). Of course, we do not anticipate tremendous new lifespan effects of PVP with a projected 20% of added efficiency. The expected value, at least early in PVPs evolution, is improved QoE of AVC and eventually HEVC video.

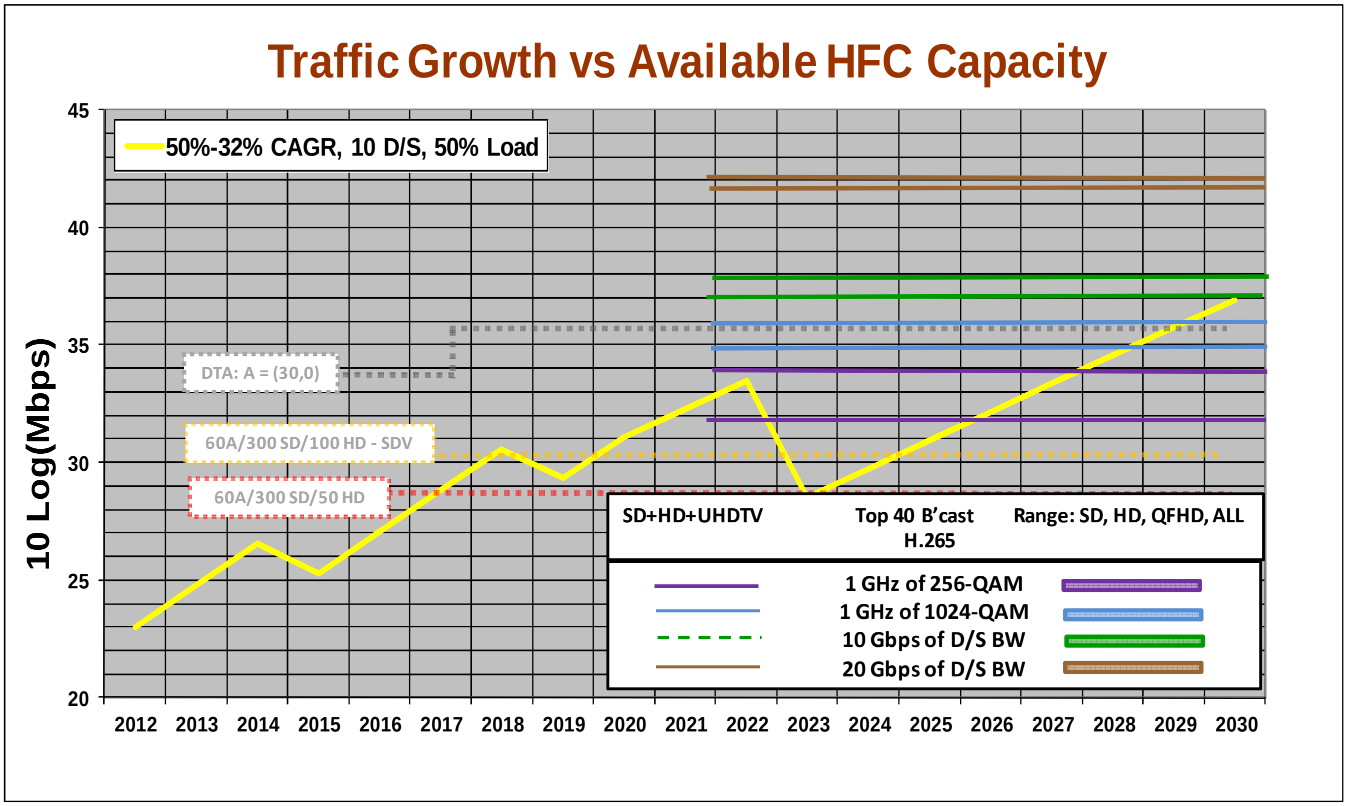

Figure 19 – HEVC + PVP, Traffic Growth and HFC Capacity (Settled CAGR Case)

Figure 19 indicates that the 20% of added efficiency at least has made the least-capable architecture evaluated (1 GHz of 256-QAM, purple) theoretically capable of weathering UHDTV services, or any substitute, similarly bandwidth-hogging applications that might beat it to market, into the middle of the next decade without the threat of breaching the threshold of capacity within a ten-year time frame under the assumptions used here. For that architecture, it also amounts to two extra years of lifespan, with the added burden on the non-PVP case that the final N+0 segmentation must also occur at least two years earlier.

For the higher capacity cases (1024-QAM, 10 Gbps, 20 Gbps), the impacts are less dramatic. Given that the existing network is, in fact, based on 256-QAM and outdoor plant equipment is 1 GHz capable only today, that impact carries more weight regarding preparation for a next generation of video bandwidth utilization.

Now consider Figure 20. Figure 20 is a redo of Figure 7, with the anticipated 20% benefits of PVP rolled up on the case of MPEG-4 AVC used in the Figure 7 analysis. In this case, we can observe a pretty significant impact of the extra 20%, largely because modest increases translate into large dividends when there is so little latent network capacity to begin with. These are shown in the upward pointing black arrows, which show the before/after of PVP being added for each scenario previous calculated. For example, in the worst case scenario in Figure 7 (and shown also in Figure 20) – QFHD in both the broadcast and the SDV tier as next generation HD, the network capacity was essentially completely consumed. Three available slots remained for IP traffic.

Because 20% of that tremendous amount of bandwidth is also a good chunk of bandwidth itself, adding it back to the pool for IP growth is pay substantial dividends as shown in Figure 20. With the savings, QFHD could actually be supported with some data growth runway. And, with a 1 GHz network, the network supports this level of enhanced HD with IP growth through 2020 under the migration assumptions used here of two segmentations. It is very unlikely that enhanced HD resolutions will be this pervasive in the market is such a short period of time. The introduction as VOD may be more practical in the timeframe of Figure 20. However, it is comforting to apply a bandwidth hungry, yet practical, "killer" application example to analyze in the projection analysis, and come out with a conclusion that the system does not only not break, but in fact enabling of such an application to a degree before any new steps or technologies are applied that could increase network capacity.

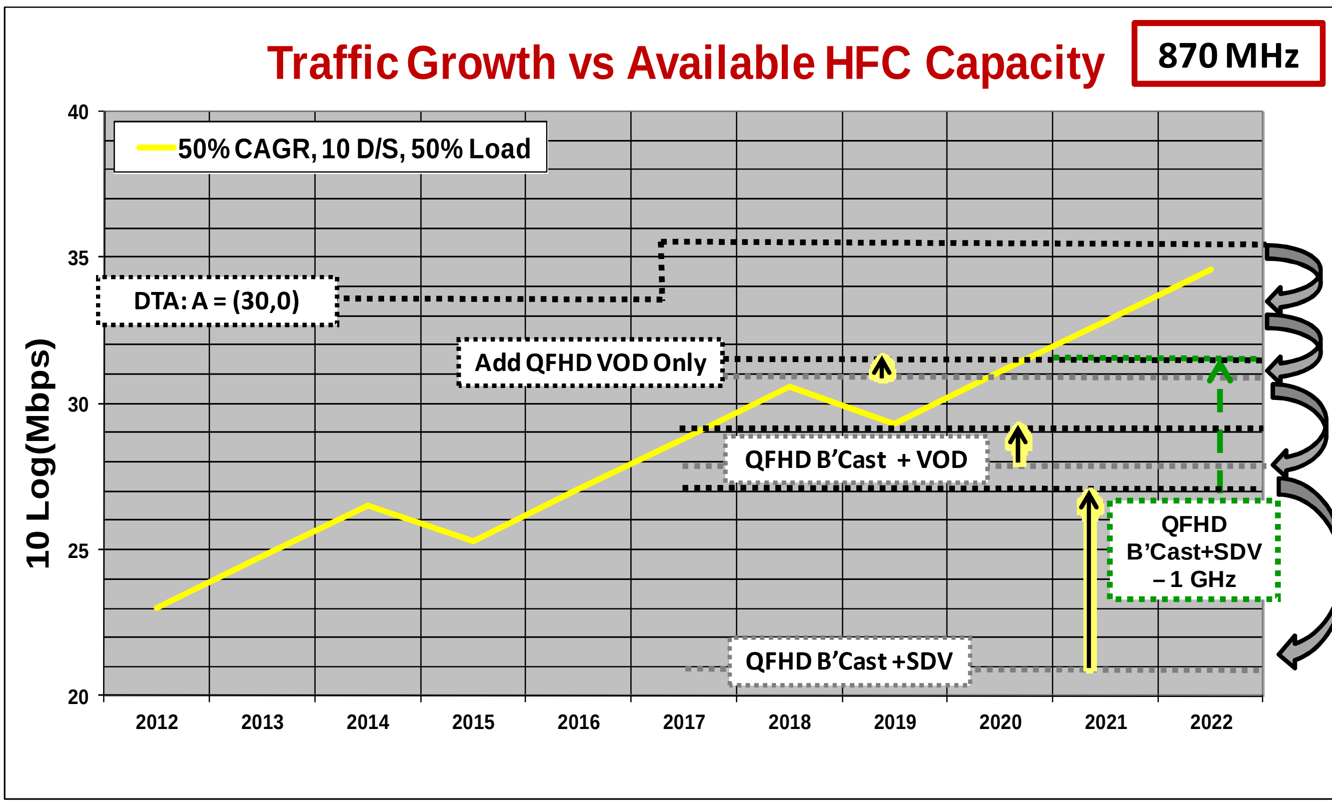

Figure 20 – Added Capacity with 20% PVP Efficiency, QFHD Format Cases (aggressive CAGR Case)

SUMMARY

In this paper, we evaluated network projections for the long-term, including many permutations of scenarios that included current and future services. We included technology and architecture options that are likely to come into play during the time windows observed, and applied these to quantify their effect. These include the shift to IP delivery, "beyond HD" video services, standards-based and innovative new encoding techniques, emerging use cases and delivery, and architecture, spectrum, and RF delivery enhancements. The result is a blueprint for an approach to preparing network service and migration plans – a blueprint that is, however, a "living document" given the accelerating pace of change in technology and services.

It is clear that there are many interrelated variables. However, any solution approach must include a comprehensive understanding that quantifiably describes the effects of network, technology, and service changes, such as shown in this paper. This is critical to properly engage in effective scenario planning, bound the problem, and prepare solution paths suited to an operator's circumstances and expectations.

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] Attneave, F., *Information Aspects of Visual Perception*, Psychol. Rev. 61 183-93, 1954.

[3] H.B. Barlow, The Coding of Sensory Messages: Current Problems in Animal Behaviour, Ed. W. H. Thorpe and O.L. Zangwill, Cambridge: Cambridge University Press, 331-360, 1961.

[4] Barlow, H.B., *Redundancy Reduction Revisited*, Network: Comput. Neural Syst. 12:241-253, 001.

[5] Deering, Michael F, "*The Limits of Human Vision*," Sun Microsystems, 2nd International Immersive Projection Technology Workshop, 1998.

[6] De Simone, F et al., *Towards high efficiency video coding: Subjective evaluation of potential coding Technologies*, J. Vis. Commun. (2011), doi:10.1016/j.jvcir.2011.01.008

[7] Dowling, J.E., The Retina: An Approachable Part of the Brain, Harvard Univ. Press. 1987.

[8] Field, D.J., *Relationship Between the Statistics of Natural Images and the Response Properties of Cortical Cells,* JOSA A, 4 (12): 2379-2394, 987.

[9] Hare, W. A. and W.G. Owen, *Spatial Organization of the Biplolar Cell's Receptive Field in the Retina of the Tiger Salamander,* J. Physiol. *421*:223-245, 990.

[10] 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.

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

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

[13] 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.

[14] 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.

[15] Howald, Dr. Robert L**,** Michael Aviles, and Amarildo Vieira, *New Megabits, Same Megahertz: Plant Evolution Dividends*, 2009 Cable Show, Washington, DC, March 30-April 1.

[16] 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.

[17] Helga Kolb, et al, *Webvision: The Organization of the Retina and Visual System. Part XIII: Facts and Figures Concerning the Human Retina*, WorldPress, http://webvision.med.utah.edu

[18] 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 Techn., Vol. 20, Nr. 12 (2010) , p. 1676-1687.

[19] McCann, Ken, and Jeff Gledhill, Adriana Mattei, Stuart Savage, *Beyond HDTV: Implications for Digital Delivery*, An Independent Report by ZetaCast Ltd, July 2009.

[20], *A Biological Framework for Perceptual Video Processing and Compression*, SMPTE Motion Imaging Journal, Nov/Dec 2010.

[21] 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.

[22] 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.

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

[24] Vu, T.Q., S.T., McCarthy, and W.G Owen, *Linear Transduction of Natural Stimuli by Light-Adapted and Dark-adapted Rods of the Salamander, J. Physiol. 505(1):* 193-204, 1997.

[25] Wang, Zhou and Alan C. Bovik, *Mean squared error: Love it or leave it? A new look at Signal Fidelity Measures*, IEEE Signal Processing Magazine, January 2009.

[26] Watanabe, S., *Information-Theoretic Aspects of Inductive and Deductive Inference*, IBM J. Res. Dev. 4. 208-231, 1960.

[27] 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.

[28] Yoshika Hara, *NHK Bets on Super Hi-Vision as Future TV*, EE Times, Sept 17, 2007.

[29] *CES 2012: 4K TV Sets Make Their Debut, Minus the Hoopla*, Los Angeles Times, January 11, 2012.

[30] 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).

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

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

[33] ISO/IEC JTC1/SC29/WG11 and ITU-T Q6/16, "Joint Call For Proposals on Video Compression Technology", WG11 document N11113 and Q6/16 document VCEG-AM91, Kyoto, January 2010.

[34] JCTVC-A204, "Report of Subjective Test Results of Responses to the Joint Call for Proposals (CfP) on Video Coding Technology for High Efficiency Video Coding (HEVC)," Dresden, DE, April, 2010.

[35] JCTVC-G339, "Comparison of Compression Performance of HEVC Working Draft 4 with AVC High Profile," Geneva, Nov. 2011.

[36] JCTVC-F900, "Common test conditions and software reference configurations," Torino, IT, 14-22 July, 2011.

[37] JCTVC-H1003, "High efficiency video coding (HEVC) text specification draft 6," Geneva, CH, November, 2011.

[38] Motorola Mobility SE6601 Encoder, http://www.motorola.com/Video-Solutions/US-EN/Products-and-Services/Video-Infrastructure/Encoders/SE-6300-6500-Series-US-EN.

[39] "Video Quality Experts Group Report on the Validation of Video Quality Models for High Definition Video Content" VQEG HDTV Final Report, vers. 2, June 2011.

[40] www.carbonbale.com

[41] www.100fps.com