CALIBRATING THE CRYSTAL BALL FOR THE NEXT DECADE OF GROWTH

Dr. Robert L Howald
Motorola Home & Networks

Abstract

Today's HFC plants continue to be a powerful infrastructure for delivery of video, voice, and data services to residential and business customers. It has successfully evolved over time to support a broader suite of services, and these services continue to be enhanced. Where will it go next? While predicting what comes next is always risky, the alternative – moving ahead without a vision – is riskier still. A decade ago, Motorola embarked on a bandwidth projection analysis in order that the industry could prepare for a future full of new possibilities. That projection was published as part of the NCTA proceedings in 2002.

Now, here we are in 2010, with nearly the ten years of "the future" behind us. What predictions were accurate? What misfired? What factors contributed to divergence in estimated growth? This paper will assess those projections. The illuminating conclusions provide important lessons learned about subscriber behaviors, the pace of technology maturity, and how new services come to market. We can use such lessons to better project the next wave of services and technology. Such knowledge is critical to making the next ten years a success.

INTRODUCTION

Predicting technology over a 10-year horizon can be risky, at best. Nonetheless, it is valuable to do so, not only to identify potential game-breakers to fundamental assumptions, but also as part of continual assessments required to keep in touch with the changing dynamics of the industry. Rapid change, relatively speaking, was the name of the game in 2001 when this initial assessment began. A rapid, accelerating pace of change is part of the dynamic as well today as we assess the cable business. The changes in play 10 years ago were primarily driven by early action around new services (primarily DTV and data) and the maturation of key enabling technologies. Today's changes have very similar elements, but with the powerful new variable of wireline marketplace competition brought on by telco triple-play providers. This represents an overriding force of change unaccounted for in any significant way 10 years ago, and that undoubtedly has had an effect on the services evolution MSOs have undertaken. Another important element of the MSO picture today is a renewed focus on commercial services outside of the small business (best-effort DOCSIS) realm. In the assessment done 10 years ago, commercial services were deliberately *not* considered, in order that the analysis could focus on residential services. However, because of increasing service rates, deeper fiber runs, PON, and WDM technologies, synergies have been created between the residential and fiber architectures. The effect could impact residential network evolution through the "pull" of new fiber technologies.

The pages that follow will be organized quite simply:

1) Walk through the predictions that were made in each service area 10 years ago, including referencing the words from the original paper [1] directly.

"All reference [1] statements will be in italics and quotes."

2) Compare that prediction to today's reality, considering explanations and lessons learned from the disparities
3) Quantify and visualize the outcomes of predicted versus "actual" on a before-and-after spectral map
4) Identify areas that completely missed the mark
5) Discuss potential drivers for the next decade of growth

Note that the reference paper for the 10-year analysis [1] was published in the NCTA proceedings in May, 2002. Note also, however, that much of the MSO data gathered to support a 2002 publication was obtained in 2001. As such, it represents EOY 2000 to 2001 data. We describe these steps in order to clarify how a paper in the middle of 2010 captures a "decade" of growth. Arguably, one could consider it 9+ year comparison.

Now, let's get on with the post-mortem!

SERVICE-BY-SERVICE ASSESSMENT

Analog Broadcast

"Analog broadcast service is projected to remain largely unchanged over the next 10 years"

It is difficult to assess this singular statement very critically, as it is certainly true in many cases and most MSOs at this instant. As such, this prediction has objectively turned out to be quite accurate. However, looked at in a fuller context, there is not much in [1] that recognizes the trajectory of some MSOs moving away from analog carriage, and some, such as Comcast, in a very aggressive way. Most have some amount of reclamation under consideration as a minimum, with the question mostly about when. The following statement also appears:

"In a typical network, 14 analog channels are expected to migrate to digital, reducing the analog spectrum required from about 500 MHz to 400 MHz by year 2006"

Thus, while there is a recognition that the principle of reclamation would come into play, [1] does not foresee the congestion-based momentum to exchange analog bandwidth for digital. As we will see, this is mostly due to a significant underestimation of the growth in the area of High-Definition Television (HD).

Digital Video Broadcast

First, some context of "where we were" when we think about the starting point time frame for the 10-year projection:

"Digital video cable is currently in the mass adoption phase. By the end of year 2001 approximately 18 million digital cable set-top boxes were in use by US subscribers. Typical systems offer 10 QAM carriers."

It is impressive to note how far cable has come in DTV since this study! The study noted 10-12 QAM carriers was typical in launching into this prediction for where it was headed in 10 years:

"A net gain of 36 additional programs is expected over the next 10 years. Four additional QAM carriers will be added to cable plants, bringing the total number carrying Standard Definition (SD) content (including migration from analog broadcast) to 16 by year 2011."

This prediction for broadcast SD is off by roughly a factor of two (low), which seems unusual given that at this point in the DTV transition, we were still on the "hype" curve. Note also, however, that being off by a factor of two over a 10-year period means that the growth was underestimated by less than 10% per year on a compounding basis. An inkling into prevailing thoughts observed at the time and driving the underestimation were two subsequent statements:

"With VOD services emerging ..... and the cable modems competing for consumers' free time, it is hard to see a case for the addition of many new broadcast channels"

VOD was at an even earlier point on the "hype" curve, enough so that it was already bleeding attention away from the "mature" digital TV technology. We will evaluate the VOD piece in a later section. We now know that the addition of HSD to the service portfolio has done little to divert attention away from TV viewing hours. In fact, in the last few years, it appears that the capability of the PC medium to support video has in fact delivered *more* hours to the big screen through the association and loyalties built with broadcast programming through the PC.

Another statement made that caused second thoughts on continued broadcast SD growth:

"Beyond our 10-year period, 2-way interactive broadcast content could be the salvation of broadcast services in a world that is otherwise evolving to total content-on-demand"

The latter part of this prediction is prescient, and will be discussed later. However, on the initial postulate, there was a sense that once interactivity arrived (envisioned as the OCAP effort taking shape), the nature of the viewing experience would change in a way that negatively impacted the pure broadcast experience. This would be due to subscribers finding the interactive channels more compelling, stifling the growth of "POBS" – plain old broadcast services. In reality, while interactivity exists, it is still struggling to find its place in the way envisioned – a way pretty much everyone envisioned at that time. Perhaps also contributing was the general misunderstanding of just how much some of us might enjoy simply being couch potatoes! More seriously, part of the interactive aspect may have been connected to demographic changes, and the anticipation of the emerging behavior patterns of "connected youth." Consider the following statement:

"Consumer interest in interactive TV exists as evidenced by a growing number of consumers interacting with TV programs using PCs"

As we now know, demographic patterns have manifested themselves in multi-media experiences not necessarily onto TVs. Instead, multi-media has proliferated onto the myriad of other devices that advancing technology made possible, and where convenience has trumped performance quality, as similarly seen in the cell phone voice example. The PC, rather than merely an outlet for TV-viewing interactivity, is instead (or in addition to) a popular screen of over-the-top content – a trend noted 10 years ago, but underestimated in speed and magnitude. Its role in MSO-owned and managed content is being defined by many operators at this point.

Finally, the calculation of interactivity did not foresee the desire for the "lean-back" viewing experience likely increasing with HD penetration, which has grown as large screen TVs became affordable.
Consider once again the statement:

"Beyond our 10-year period, 2-way interactive broadcast content could be the

salvation of broadcast services in a world that is otherwise evolving to total content-on-demand"

The perceptive recognition of a "total content on demand" world in this early digital era was indicative of the anticipation of how VOD would transform the industry. While this means of unicast video distribution may not be the core technology around which "everything-on-demand" takes place, there was a general sense of engineering future systems for an increasing range of multi-cast and unicast delivery. Total content on demand is now more broadly captured by the industry mantra of serving the "four any's" – any content, anywhere, anytime, any device.

HDTV Broadcast

"Is the picture quality worth the price of an HDTV? How many consumers viewing a 42 inch screen at normal distances can discern the improvement in HDTV quality relative to DVD or MPEG 2 SD quality?"

Probably the most significant underestimate in terms of bandwidth repercussions was in the area of HD viewing. However, objectively, the error is not actually so large from the perspective of what is in the field in many places today, but more so in the context of where trends are headed this year and next. The above statement alludes to some of the perceived barriers to scaling HD – the high price of HDTV's at the time, and the associated value proposition for normal viewing. As expected, and predicted in the paper, the price for mass adoption did get to a tipping point in the 10-year time frame – relatively recently, in fact. The original analysis did not foresee that with HD would come a new class of display technology and an overall increase in "normal" size screens to enhance the value proposition of HD – the concern alluded to in the above statement which referred to what was essentially the largest "tube" sizes of the day of 42 inches. Associated with this is the increase in sports viewing (NFL Network, Golf Channel, MLB Network, ESPN19 anyone?) for which HD is fuel for the fire.

Finally, a key missing factor was the forces of external competition – specifically satellite broadcasters, who, without a VOD play, found a powerful, profitable refuge in racing to the front of the HD competition.

"With two HDTV programs per carrier, systems will begin carrying 4 to 6 HDTV video programs this year. By year 2011, 16 HDTV channels are forecast. The bulk of content will continue to be delivered in SD resolution."

The underestimation of HD content, because it represents the largest Mbps volume of all current services, adds up to the largest error in the bandwidth maps we will show later. However, note that we are in a period where HD is a relatively rapidly moving target – most MSOs are looking to add HD content in a big way, with physical bandwidth in the way in the near term. Bandwidth constraints revolve around both the Mbps associated with HD, but also with the need to simulcast alongside the SD version.

However, again, while most MSOs expect rapid addition of HD programming in the very near term, the prediction is actually not very far off at this exact moment for many systems. The above statement predicts 8 QAMs of HD, using a two-HD-programs-per-QAM relationship, where 2-3 is normal and content dependent today. As such, the prediction was for 32 HD programs by 2011. In fact, by my count, my home cable system, in a middle tier (top 30) metro area overbuilt by VDSL triple-play services, has between

35-40 HD programs depending on the channel lineup version I review.

Nonetheless, large MSOs today are looking to be HD-competitive beyond the VOD-based HD library often used to boost advertising campaigns. This generally means campaigns to achieve 100-program line-ups of HD, or greater, or approximately 40 QAMs worth (240 MHz) – about one-third of the bandwidth on a 750 MHz plant.

While the HD prediction may have been in fact quite accurate in the purely numerical context of systems today, the underestimating of the trend of accelerated HD deployment leads to an area that was completely missed – implementation of switched digital video (SDV) technologies in cable. Long the "only" realistic solution for telco video delivery via the overmatched xDSL wires, cable has seen this matured technology as another bandwidth tool in the toolbox – such as to increase an HD line-up without having to physically support the complete spectrum required to do so. Exploiting both the shrinking serving group sizes and natural statistical gains of popular, multicast content, bandwidth gains (and thus program count gains) of 2-3x of "virtual" bandwidth can typically be added.

MSOs are at different stages of SDV deployment. Allocations of 4, 8, or 16 SDV QAMs is roughly par for where competitive systems that have deployed the technology will likely be this year, moving towards 20-24 where aggressively underway already. It is well-documented that TWC has been the most aggressive of the large MSOs in North America deploying SDV.

Note that MPEG-4 gains were not factored in as elements of the 2002 analysis – nor has there been significant MSO activity moving towards MPEG-4 based HD for traditional MPEG-2 TS QAM delivery.

Video-on-Demand (VOD)

Again, as a point of reference context from [1]:

"At the start of Year 2002, operators had launched or planned to launch VOD (commercially or in trials) in almost 90 markets"

Thus, VOD service was very much a newly emerging service, and with that emergence, there was much hype in what it could become, and the impacts it might have on the fundamental broadcast-oriented nature of cable video delivery. Nonetheless, and in spite of the hype of the period, the analysis relied on key numbers in ultimately predicting conservative growth. An important factor was noting that VOD was necessarily tied to digital penetration, and there was enough available trend data and market research at that time to have some "expert" opinions rendered on that trajectory:

"Kagan's 2001 annual growth forecast shows digital cable penetration growing to 63% by year 2011"

The above analysis was actually relative to homes passed, so error can be attributed to properly capturing the actual MSO service penetration multiplier in today's more competitive world. It turns out to be a good representation as a function of MSO cable subscribers, however, and so reasonably captures the digital growth trajectory in the context of the cable customer base.

Also, the analysis relied on realistic models and research in viewership behaviors of like-content to that for which VOD would naturally support. To a first order, this would be the "Blockbuster/West Coast Video" replacement model, followed secondarily perhaps by likely very popular TV shows – those that inspire lifetime loyalty and repeat watching. These would be shows with the popularity of, say, M*A*S*H, Seinfeld, Friends, etc. Of course, the "West Coast Video" model did indeed have a noticeable impact on … West Coast Video!

Using peak-time viewership behaviors while concluding that availability of VOD does not lead to major changes in basic viewing behaviors of that type, the following guidelines were used to model its growth and impact:

"The estimate of simultaneous use during peak hours is 5% today and forecast to increase to 9 % by year 2011"

The 9% value is in line with ranges used today in system engineering of VOD. Peak hour concurrency factors from 5% to 15% are unofficial but observed values used by architects today. The net result of this penetration and peak-time concurrency is 3-6 VOD QAMs predicted by the analysis from 2002. The analysis assumes that node sizes (quite accurately) will be in the 500 hp range, justifying three QAMs. It does not further consider service group splitting of video and data groups, which today may result in sharing of VODs across nodes. Doing so would double the QAM count for the same traffic engineering parameters, resulting in 6 VOD QAMs. And, in fact, 4-8 VOD QAMs is a reasonable count on today's systems, with some MSOs expecting to increase this to perhaps 12-16 in the coming years.

VOD trajectories may grow slowly moving forward. A significant factor in where VOD heads, and recognized in the 2002 analysis, is the impact of IPTV on VOD. VOD, as a unicast video delivery mechanism, represents a natural service type to permit smooth migration to IP delivery from a technology standpoint – opposed of course by legacy infrastructure and HSD architectures built for data. Nonetheless, a prophetic statement from 10 years ago was:

"VOD can be streamed over the Internet using IP and DOCSIS………at these rates, audio and video quality is competitive to that offered over MPEG 2 multi-program transport streams to set-top boxes"

Of course, we now know that the video and HSD service rate relationships reached a watershed moment, shown in Figure 1. There was a separate section in the 2002 paper entitled "VOD over IP," which fits best as a discussion topic in this section of our analysis today. It is hard to have a discussion about video service trajectories and service expectations without devoting some time to video over IP. The impact brought about by the crossing trajectories in Figure 1 brought some inevitability to cable's video evolution path. So, let's dig into this idea of video over the HFC IP pipe, which today is implemented by DOCSIS.

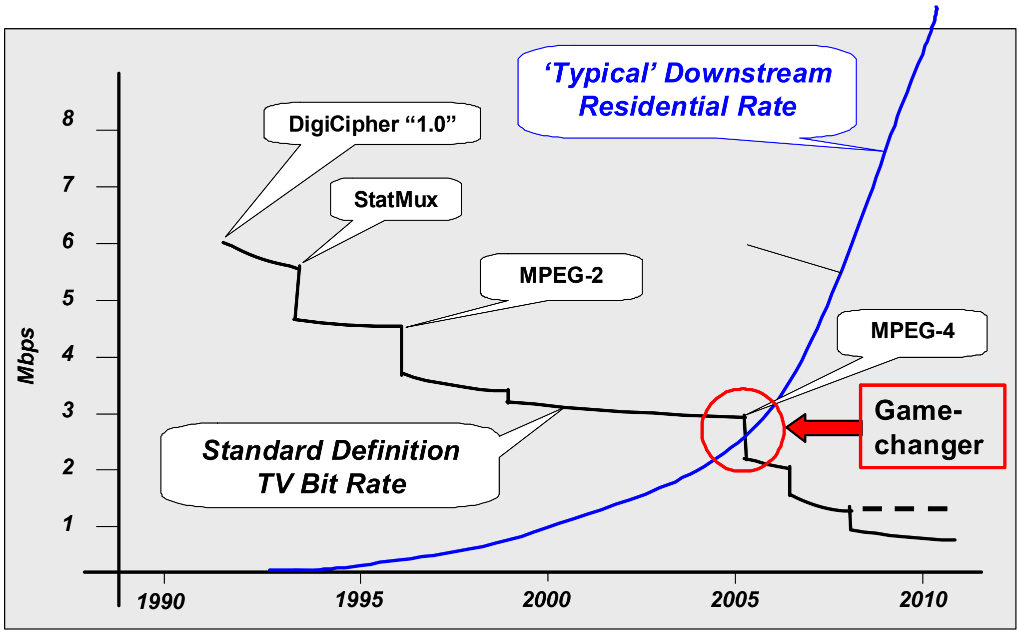

Figure 1 - HSD Access Speeds vs. Video Rates

Note from Figure 1 that the introduction of MPEG-4 encoding (H.264 AVC) brought with it another 50% of average video rate bandwidth efficiency – though this 50% value can be quite content dependent [2]. As a new technology, it would of course involve major infrastructure changes. As such, it naturally conjures up thoughts on how best to introduce it, and suggests a new opportunity to remake and improve future delivery systems, such as with IP delivery.

In addition to MPEG-4, another potential bandwidth efficiency exists, which was not considered when initially pondering IP delivery in [1]. This is the benefits of variable bit rate (VBR) delivery enabled by IP packet scheduling mechanisms. HSD schedulers are designed to efficiently make use of capacity when input data is of varying packet sizes and arrival rates. Coupled with more streams per carrier (MPEG-4 helps with the law of large numbers) and the introduction of DOCSIS 3.0 channel bonding, VBR promises additional bandwidth efficiencies over traditional CBR delivery – likely in the range of 20-40%.

This combination of bandwidth efficiencies is doing two things:

1) Providing a reason to hold off on what might otherwise be continued growth of MPEG-TS based narrowcast delivery for on-demand
2) Making reasonable the idea of "simulcast" of MSO channel line-ups (or portions thereof) for delivery of essentially the same video line-up over IP

Consider that a logical target of IPTV for MSOs initially may be PC screens with modest (VGA-like) resolutions. In this case, the bandwidth numbers come together quite nicely:

- Today's Broadcast Line-Up (MHz) = [FULL] MHz

- IP Simulcast Line-up (MHz) = [FULL][50%(MPEG-4)][75%(VBR)][50%(VGA)]= 18.75% x [FULL] MHz

Thus, an initial IP video offering for the PC could require just one-fifth of the full-service bandwidth being made available. Put another way, 300 Broadcast SD channels over 30 QAMs could be handled instead by 6 QAMs, or 36 MHz of spectrum…pretty powerful stuff. Of course, this must work its way up to 12 QAMs (72 MHz) to bring SD to the PC, but this could perhaps occur in time as the IP service is gradually rolled out by replacing less-efficient MPEG-2 VOD carriers.

So, as we can see, there are new and compelling reasons to consider the unicast "everything on demand" future as one derived from IP delivery even from an access network standpoint – with the access network being the last piece of the HE-to-end device architecture not already IP-centric.

Note also the issue that struck a chord 10 years ago is still under scrutiny today in envisioning this IP evolution in full:

"……but end-to-end QoS mechanisms are required to support continuous data rates in the range of 2 Mbps to 4 Mbps"

CMTS platforms with prioritization capabilities, and a myriad of IP QoS techniques have been developing in these last 10 years, but the concept of "guaranteed" services for IP services per the Intserv model has given way to simpler and more scalable means to offer statistical guarantees. However, with video services being intolerant to packet drops, and with constraints around jitter, these assurances are still to be proven out in a way that does not require severe underutilization of pipe capacity [2][3].

Another perceptive statement:

"HDTV VOD might be an interesting proposition. Early HDTV adopters are good candidates for higher priced VOD content."

While we have moved beyond "early HD adopters," it is certainly the case that service providers (and over-the-top-providers) have recognized that HD content can be charged at a premium – not because "early adopters" represent a wealthier segment of the population, but because the value proposition of HD video is that strong.

Finally, a still relevant postulate from 2002, mired in some regulatory obstacles:

"VOD has much more potential than just replacing the video store. MSOs could offer server based Personal Video Recording (PVR) capability."

This one still remains to be seen, given the potentially significant implications to CPE, storage, content delivery networks, and access bandwidth.

<u>Internet Access</u>

<u>Downstream</u>

This basic user expectation, which was at the origin of placing cable data services at the forefront, has not changed:

"Users' key expectations are low latency in delivery of web pages and downloads, rapid updates in games and seamless delivery of streaming content. They also expect "always on" service."

Of course, what qualifies as "low latency" and "seamless" has changed, as the bar has been raised in both cases. Additionally, "always on" today really, really means ALWAYS, as opposed to almost always. The most powerful example of this reality is the routine use of remote offices and workforces despite the increase in the quantity of information that is exchanged daily.

Once again, to put into perspective "where we were" at the start, consider this statement from [1]:

"It is estimated that YE 2001 a provisioned gross average bit rate of 21 kbps per subscriber was needed to achieve subscriber satisfaction"

With traffic engineering principles applied in the original analysis, this gross average worked out to a peak service rate of about 1 Mbps. The state of spectrum at this stage was one downstream DOCSIS carrier deployed, serving multiple nodes, which served a larger number of homes passed than today, creating serving groups in the multiples of thousands.

Let's compare where downstream data is today and where it is headed in the short term. Like HD, HSD is, relatively speaking, a moving target. Most MSOs are executing on plans now to add downstream DOCSIS carries, in some cases to add them and to bond DOCSIS 3.0 channels. They are simultaneously working on shrinking service groups on a market-by-market basis based on competitive need, engineering for more bps/sub to keep ahead of the growth trajectories. As such, downstream DOCSIS spectrum, which is moving towards no longer being shared across nodes if not the case already, can be considered somewhere between 2-4 channels today, on average, with plans to increase to 6-8 shortly thereafter where spectrum is available or can be cleared. Some MSOs are looking more aggressively still to execute on the transition to an all-IP architecture, in which case DOCSIS carriers would take on a larger role and consume yet more spectrum than 8 slots more quickly. The trend towards all-IP is undeniable, but the speed at which that can occur for video is hindered by several key legacy factors. Thus, for the purposes of DOCSIS QAM count, to assess the prediction in [1], ranging up to 8 in the 2010-11 near term is a reasonable high end. However, recognize that the accelerated pace of 2 to 4 to 8 would be expected to continue to take place in 2011-12 and 2012-13 in more aggressive transitions, which is an important immediate consideration for planners.

Looking at the numbers, then, it was recognized even early in the HSD business that trying to guess the next big application was less likely to capture the growth requirements than a compounded growth rate. In other words, over 10 years, it is smarter to base projections on the smoothed curve of growth as opposed to the more realistic series of step functions underlying the average growth trajectory, which led to the following long range projection:

"The 10-year forecast, therefore, is for consumption to grow at 50% per year"

The assumption of historical growth rates continuing has generally come true, although the actual growth as calculated in terms of 4 DOCSIS carriers served over today's node sizes turns out to be closer to 40% compounded on average. While this reflects a pretty good prediction, over 10 years this means being off by a factor of two. Using this growth premise, the analysis in [1] concluded that 6-8 DOCSIS carriers would be required to satisfy demand. That is, in [1], it was anticipated to be a moving target as well, and in so doing anticipated 6 DOCSIS carriers for 2010 and 8 for 2011. This is quite an accurately painted picture, given that we are looking at 2-4 at the moment, and moving to 4-8 in the near term. The difference between 6 versus 4 can be traced to the difference noted above in the impact of being off by 10% in the compounding rate for 10 years. Nonetheless, this is a pretty perceptive projection.

With respect to peak rates – an issue coming to the forefront as we introduce DOCSIS 3.0 channel bonding – and applying traffic engineering parameters predicted, it was anticipated that the peak service rate to the consumer would be approximately 10 Mbps. This is indeed in the ballpark of where downstream speeds offered by large MSOs in competitive environments, where tiers in the 5-20 Mbps represent high end downstream services.

Overall, DOCSIS downstream growth was quite accurately predicted, in particular given the limited amount of cable legacy to draw upon for HSD services.

Upstream

"Average upstream consumption increases from 7 kbps to 700 kbps"

This represents about 59% if calculated as compounded growth. However, the figure calculated in [1] to represent growing traffic is actually an estimate based on *downstream* compounding (actually further broken into compounded volume@25% and compounded concurrency @20%) multiplied by a factor representing the traffic mix trending towards more symmetry. The compounded growth tied to symmetry is described as follows:

"Upstream bandwidth increases more than downstream due to the expectation that rate asymmetry (the ratio of downstream to upstream rates) will decrease from 6:1 to 3.5:1"

We would likely not take this approach today, having noticed that this symmetry trend has actually reversed itself with the introduction of video clips as core drivers of HSD bit volume. On average, however, and recognizing that the above statement is on a per-sub basis (penetration accounted for as in [1]), this growth rate to 700 kbps overstates average rates today. However, as with downstream, and perhaps more so than downstream, adding upstream immediately is a high priority for MSOs today. It is a recognized potential bottleneck.

For the 700 kbps to closely represent the situation would mean getting the 4 upstream carriers that some MSOs are looking to turn on going, all at 64-QAM. Alternatively, it comes close also by considering a next stage reduction of average service group size shrinking, creating new virtual bandwidth – i.e. more bps/sub by reducing subs. With neither of these two elements quite in place, and not likely to happen until next year at best, this estimate appears off by about a year or so, which is not bad. This is not surprising when recognizing that aggressive compounded growth rate used to generate it, and noticing that the symmetry trend in fact did not continue as anticipated. All in all, the estimate represents a reasonable prediction of where things headed for upstream. They are certainly not there yet, but a noted objective in many camps is to do exactly what it would take to get to this range in the near term. And, it is preferred in any case to slightly miss on the high side than the low side.

"As for upstream.....the peak rate is expected to increase from 200 kbps to about 3.2 Mbps in year 2011"

This is a very high-end upstream service tier range, but nonetheless in the field of play of today's offerings.

All in all, then, the upstream predictions have turned out to be quite solid.

Streaming High Quality IP Audio and Video

We covered much about the implications of streaming video in the VOD section when discussing IP video. The context of the 2002 paper recognized over-the-top video services as a bandwidth driver in calculating the HSD growth, and identified the means by which this occurred and would continue to occur:

"PC based multimedia decoders (Windows Media Player, RealPlayer, QuickTime, and ultimately MPEG4) are widely used to deliver low resolution, low rate VOD over "Best Effort" Internet access service"

Note that, because of the obvious value to the IP world for cable, MPEG-4 encoding was anticipated.

This over-the-top video was distinguished in [1] from the "high-quality IP video" which we would, today, probably think of as delivery of managed MSO content on the IP pipe. This was not necessarily broken down into these segments in 2002, but given the overall novelty of the concept at that time, it is difficult to assess that added detail critically. Nonetheless, this farsighted statement was made in 2002:

"Mass-market penetration of streaming will likely wait until solutions are in place to move the content into the entertainment center and other places within the home. Lacking solid QoS guarantees, entertainment quality video and audio cannot be delivered reliably enough to satisfy paying consumers."

Indeed, precisely at this time we are seeing intense activity with MSOs and at CableLabs around MoCA, DLNA, and UPnP to ensure PHY throughput and QoS delivery around IP-enabled homes for multi-media content distribution and delivery. MSOs have suggested, modified, re-visited, and updated various approaches to developing home "gateways" to ensure high quality delivery, with multiple IP avenues of distribution throughout the home to IP client devices.

What MSOs want to avoid is a newly rolled out video services architecture that can be tainted by improperly subscriber "engineered" home networks. An IP home architecture, potentially managed, provides some assurances that are not available today when homeowners are combined with STBs, cable modems, routers, splitters, and coaxial cable.

Hi QoS Audio Streaming

"A successful service might grow to a saturation penetration of 20% HP by 2005"

This prediction is simply a swing and a miss. This can be attributed to the ease of high quality audio delivery over the top (Internet Radio examples), in part because of the modest bit rates when compared against the increasing downstream tiers. Figure 1 is a good reference point – when video becomes supportable, audio becomes nearly insignificant.

In addition, the expectation bar has been lowered somewhat for audio through the years. Audio services on the web are often likely background or in concert with other multi-tasking functions – something not behaviorally similar in a video environment. And, analogous to voice services, music on the go in the form of IPODs and PDAs with mp3 players have lowered the bar on consumer accepted audio quality. Perhaps audiophiles are fewer and further between, but consumers, en masse at least, have chosen convenience over hi-fidelity.

Fortunately, because audio is so bandwidth non-intensive, this misfire does not significantly impact overall bandwidth results.

IP Telephony

"IP telephony is estimated to grow from a 2% penetration in year 2002 to about a 30% HP penetration by year 2011"

VoIP growth has been robust since its introduction, although perhaps not quite as robust as predicted 10 years ago across the board for Cable VoIP services. Intervening factors were the introduction of over-the-top voice (i.e. Vonage), and the shift, in particularly demographically, towards consumers choosing a single voice service, and choosing the most convenient one – their cell service. Nonetheless, voice traffic is a rounding error in traffic analysis. It requires proper treatment (highest priority) in the DOCSIS world, but in terms of bandwidth consumption, it not significant.

Node Segmentation

Figure 2 shows a figure directly from [1], suggesting the time frame and justification for node splitting throughout the past decade.

This figure turns out to be quite prophetic. Its essential conclusion is that node sizes would be cut from the range of 2000 hp to 500 hp, which would be sufficient for downstream into 2011:

"By choosing to leave the downstream node size at 500 HP, more carriers are required but equipment cost is saved. This configuration supports expected traffic requirements through year 2011."

Most operators are pondering, planning, or executing that next post-500 hp split. On average, numbers are beginning to drop below 500 hp. This prediction turned out to be quite accurate as far as macro bandwidth trends driving downstream node sizing. Note also that the DOCSIS downstream count shows 6 carriers moving to 8. Again, while a larger set of DOCSIS carriers than in use in general today, most MSOs see this number of DOCSIS downstreams in their relatively near future. The introduction of DOCSIS 3.0, in which bonding technology enables higher service tiers @N x 40 Mbps, has likely accelerated the addition of channels. Basically, it makes it more probable that chunks of 4 channels at a time will be added than single new DOCSIS downstream.

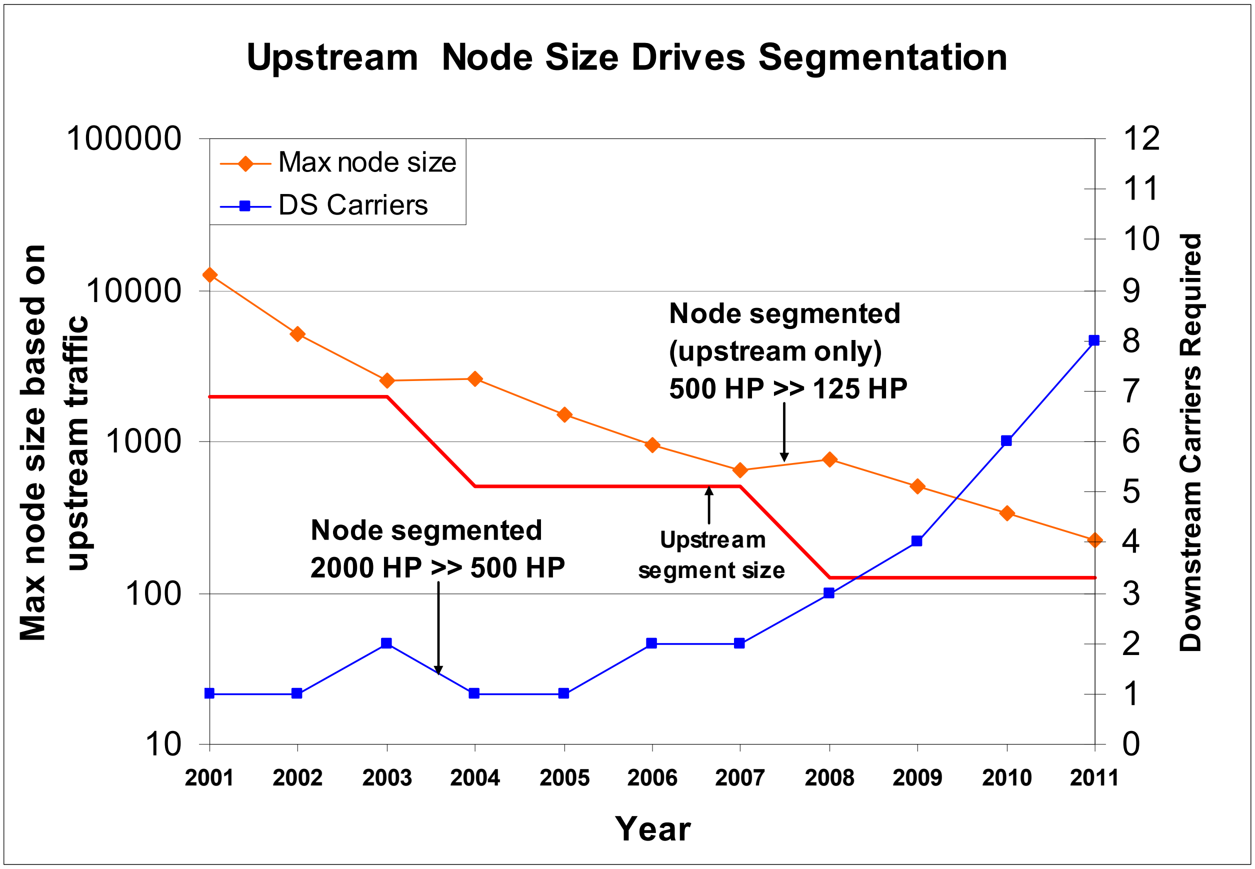

Figure 2 – Node Splitting Projections from [1]

COMPOSITE BANDWIDTH

Figure 3 shows a comparison of the cable spectrum, based on these components:

1) The historical basis from [1]
2) The predicted spectrum usage for 2011 from [1]
3) Three cases of "actual" meant to cover the range of "typicals"

Since the range of "typical" matters in a way that impacts available HFC bandwidth as defined by the upper band edge, it was felt that highlighting how service mix choices matter to this key parameter over a range would be valuable. This was also suggested in [4], where the ability to support new growth as a function of bandwidth available or created, and the broadcast/unicast mix, was quantified in years.

"It can be see there is spare capacity in HFC plants built out to at least 750 MHz. Downstream carriers are added to satisfy downstream demand through year 2011."

Notice that it was observed 10 years ago that the downstream spectrum supported the bandwidth growth needs of the anticipated service mix, given that serving groups sizes would be shrinking. As we know today, that has proven to be the case, generally, even with the underestimation of HD in the projection. Although, depending on the mix of other services, that the mix always fits comfortably in 750 MHz, which was the prediction, is not foolproof. However, actual growth has not exceeded the trends that already existed at the time that were resulting in the introduction of 870 MHz and 1 GHz plant equipment.
Consider now the upstream projection once again, and note that there was an expectation

that upstream growth would have forced a node split by this point. This is associated with the relatively accurate prediction that there would be more pressure on upstream bandwidth versus downstream bandwidth, and thus the upstream would drive the need for further segmentation. This has turned out to be the general scenario across MSOs offering HSD services, especially in competitive environments.

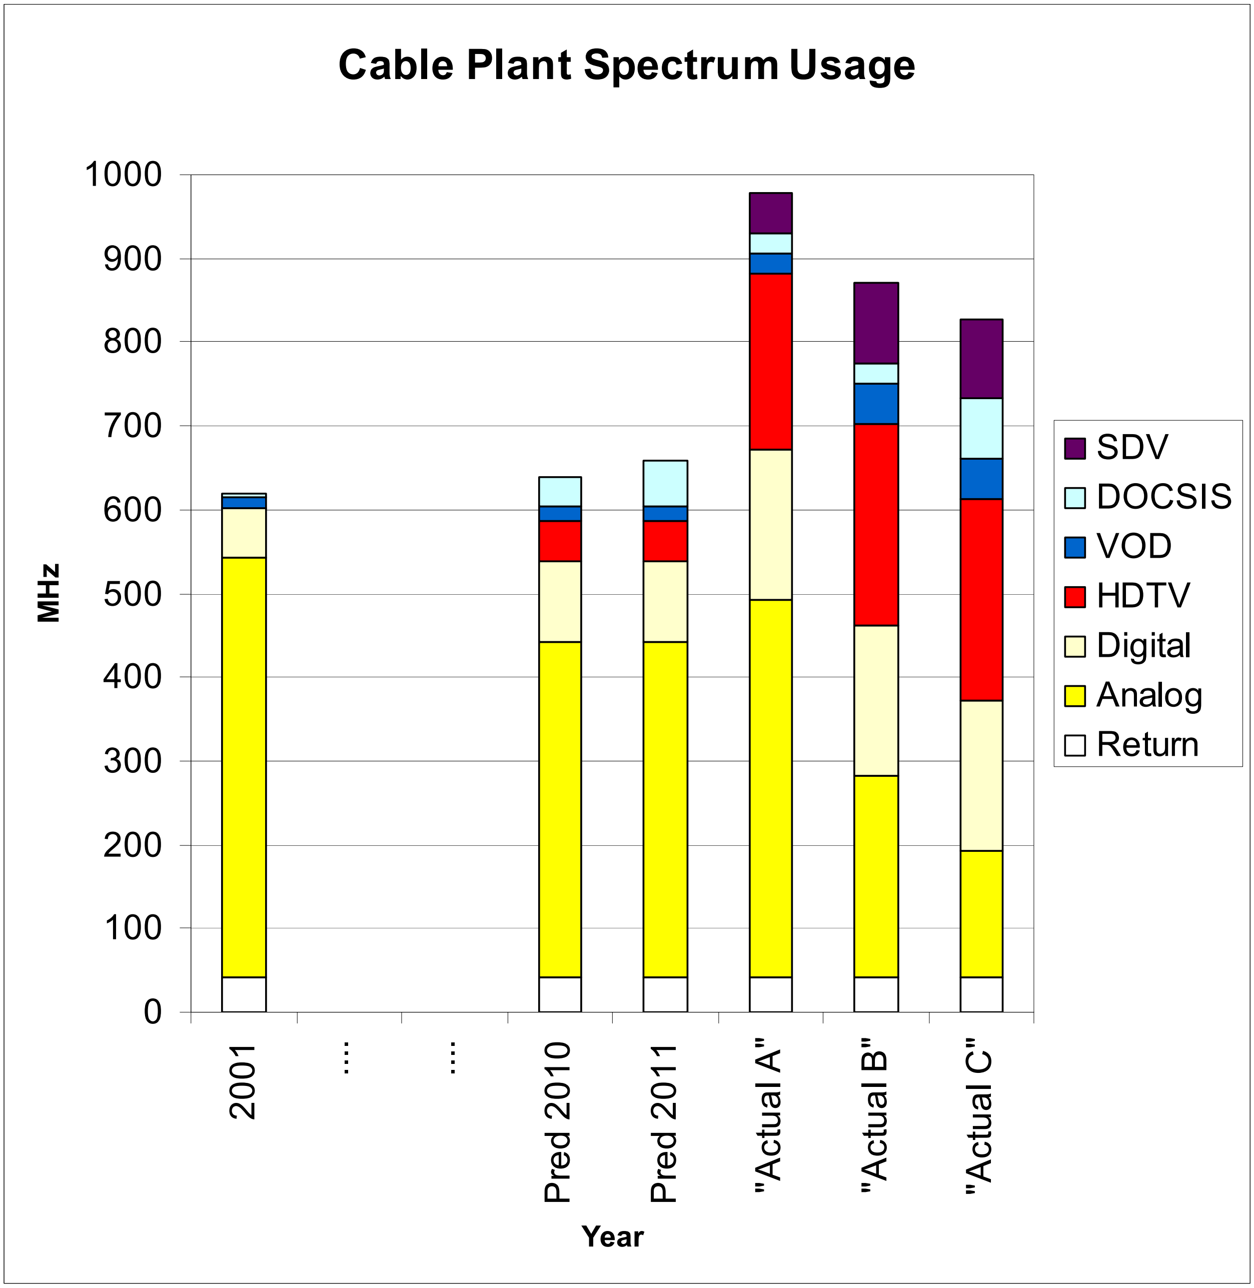

Figure 3 – Historical, Predicted, and Current Spectrum Usage

However, the prediction that an additional split to 125 hp would be necessary for the upstream was driven in part by an underestimation of what DOCSIS would evolve into. The original analysis assumed that the upstream would be capable of about 80 Mbps maximum, based on using 16-QAM

@ 3.2 MHz bandwidth, or about 10 Mbps per channel. It did not foresee the implementation of 64-QAM, nor S-CDMA to turn on the low end of the band. As such, the prediction is roughly one-half of what can be squeezed through a fully optimized and exploited upstream today. Thus, being off by almost one-half, there is nearly one additional "traffic doubling period" missing from [1].

Now, at 25% compounded growth, traffic doubles in roughly three years, while it doubles in roughly two years at 40% compounded growth. The original analysis from [1], based on approximately 59% compounded growth, is thus approximately 1.5 years off in time with respect to recommending a node split for upstream bandwidth. That is, traffic doubling at a 59% clip takes about 1.5 yrs. Applying this correction, the analysis would have concluded that a node split would be required at year "2008.5." And, indeed, dropping node sizes below 500 hp average lines up in time with this trend in competitive markets on an as-needed basis. The granularity of node-splitting was in fractions of one-fourth in [1], merely because the ability to segment most core node families in fourths. One segmenting visit per node was assumed.

Looking back, rather than compounding the concurrency year-on-year, we can recognize that increased speeds lead to lower concurrency for the same service rate on web-browsing type services, which has some logic to it – things get there faster to consume, but consumption time (human oriented) is about the same. This phenomenon is shown in Figure 4.

Of course, as web browsing continues to become multi-media oriented, this could reverse course in favor of supporting committed streaming requirements. Nonetheless, the point of this lesson learned was that, should compounded concurrency growth be removed, the upstream compounded growth rate predicted drops to approximately 30%, and thus roughly 2.5 yrs of growth time. In this case, the node split prediction would have recommended that 2009.5 would be go-time for the next split – again well within the ballpark on where upstream growth has taken us.

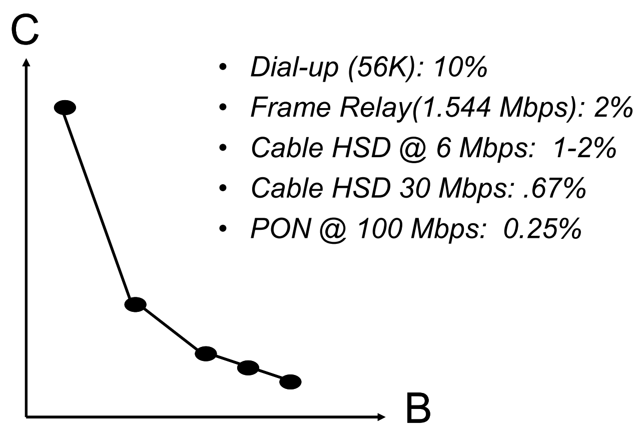

Figure 4 – Historical Concurrencies Observed

All things considered, although the mix had some miscalculations, both the upstream and downstream macro bandwidth requirements were reasonably well-aligned with where things have gone. The largest "miss" associated with HD penetration that underestimates its contribution to overall bandwidth is in part offset by the introduction of SDV, which allows for additional HD programs to be added to the "broadcast" line-up without requiring an allocation of spectrum, at least not 1:1. Also, while the paper notes that analog reclamation will not be observed in a large way by 2011 – true in cases, not true in others – the consideration of this highly efficient means of exploiting HFC allows for the net spectrum predicted versus "Actual" (case C) to be similar.

While a key miss was the introduction of SDV into cable these last couple of years, there was certainly commentary suggesting the need for more and broader content, and the potential for the eventual role of HD content. This 10-yr old thought is the kind of logic that led to the march towards adopting SDV:

"….there are systems that would like to provide more broadcast channels. Needs include serving multi-lingual and ethnic populations with international programming and lots of HDTV, eventually"

FTTP ARCHITECTURES

"Beyond year 2011 MSOs will have to decide between pushing HFC capacity further or re-trenching to bring fiber to the home"

"Capital costs for FTTH are expected to become competitive for green fields deployments well with the 10 year forecast period"

Both of the above statements have turned out to be sound predictions. Many MSOs are looking at whether or when HFC naturally migrates to FTTP, and are developing transition technologies (i.e. RFoG and DOCSIS back-office for PON) to enable such a possibility. There is little argument that fiber and coax builds (parts & labor) cross in optimal choice as densities decrease, but the build out to the FTTP home still comes at a CPE premium under more typical HFC densities, although the gap is shrinking. Other factors, however, have driven "greenfield" environments to become based on FTTP, such as the real-estate development market, and the general perception that if FTTP is an expected endpoint, it does not make sense to be installing new coaxial cable.

"FTTH is considered the "end game" since it offers enormous bandwidth to the home"

It is hard to argue the point "more bandwidth is better." Of course, costs and legacy obstacles make this not a cut-and-dry question when considering time frames. It can be easily shown [4] that HFC architectures can be incrementally improved and exploited to deliver capacities that enable tremendous new growth that converts into years and years of life span. The question is whether that span exceeds the time frame for which practical business planning needs to take place, or if it is on the horizon in a way that plans need to be put into place now to enable this FTTP transition. Or, even if not an obvious "endgame," is "just in case" investment and planning warranted. Most operators fall into this latter category – its too big to ignore, and no one wants to be left without adequate bandwidth given the historical inability to recognize the onset of the new killer applications.

A DECADE AGO IT WAS CONCLUDED…..

"In 10 years the number of bits pouring into the home will be over 50 times the amount delivered today"

This was stating that straight line growth in bits *consumed* would likely continue at a compounded growth rate close to 50% (i.e. 1.48^10, or 48% growth for 10 years, equals 50x). It does *not* mean that there is 50x more QAM carriers, of course. This QAM count number is on the order of 10x. Of these 10x more carriers, the consumption patterns of consumers has shown a steady compounded growth along the lines of Moore's law for computing power. This simplistic model is generally associated with data services, but as video and data services begin to blur in the digital realm, the proper traffic growth model to use for consumption becomes less clear.

"Rich interactive multimedia video will be commonplace"

While interactivity has increased, a decade ago it was felt that by now it would have emerged from the shadows and be representative of the "typical" viewing experience. This has not occurred for a myriad of reasons, but initiatives such as EBIF, OCAP, and architectures in CE circles still exist with expectations to do so. Of course, the fact that multiple initiatives are still in the mix is part of the reason interactivity has not scaled as expected.

"HDTV will succeed as one of the many services"

Though not much of a reach to predict at the time, HD taking hold had become a major topic of discussion at the time of [1] because of how slow this seemed to be taking place. Nonetheless, this is a clearly true prediction, and, as indicated, in fact was underestimated in [1], albeit not by very much in years at this instant of time. While not foreseeing the trend trajectory as aggressively as it has played out, HD mass adoption has been closer to the back end of the 10-yr period analyzed.

"Telephony will become a rounding error in the traffic analysis"

Indeed, once HSD scaled with year-on-year compounded growth, voice bps became insignificant from a bandwidth perspective.

"This growth has been shown to be easily supported by continuous upgrades to the HFC infrastructure"

The relative ease of scalability of the HFC plant (yes, field folks, I know what is involved here!) has been proven out yet again as new services are added with incremental changes to access networks. Also, there is much room left in HFC in terms of capacity:

"Much more can be squeezed out of HFC, if and when needed…."

– so much so that the discussion on where to "end" hinges more on capex spending/opex maintenance questions:

"MSOs will continually be faced with capital investment trade-offs between infrastructure upgrade costs vs. how much excess capacity to install…."

– as well as hinging on developing strategies for retiring legacy infrastructure:

"Legacy equipment will tend to make upgrade trade-offs more complex and optimum timing will vary"

TEN MORE YEARS OF POSSIBILITIES

Moving forward, there are a bundle of new service opportunities that can keep the bandwidth growth line trajectory moving northbound, as well as some practical and less sexy reasons bandwidth will be on the rise.

Simulcast Redundancy

Somewhat ironically, the fact that there are multiple opportunities aggravates the situation because of the need to support consumers on the network for existing services. This plays out primarily on the video side, where, because of the notion of "TV Everywhere" and the continuing enhancement of the video experience, it will be important to have a "simulcast" strategy. That is, today's network uses simulcast to support digital and analog carriage of a subset of available channels to support the different tiers of cable services. In addition, digital channels broadcast in SD are simulcast HD for some (increasing) subset of the channel offering. We can expect this same situation to play out for 3D services, since essentially no consumers today have 3D TVs, so the initial availability of content for it will require a 2D version. Finally, supporting multiple devices typically means smaller screen versions being available (VGA, CIF, etc). While these are expected to be delivered over the IP infrastructure, they do represent redundant streams of bandwidth. Fortunately, for both access networks and storage architectures, the "small screen" bandwidths are much smaller than their HD counter parts, so they carry less impact. Furthermore, introduced via the IP network, the opportunity exists to introduce the service as MPEG-4 encoded only, another roughly 50% average savings.

This bandwidth logic above is one reason that multi-screen edge transcoding has given way to multiple-stream storage models. At this stage, real time transcoding at the consumer edge is to costly to envision as a way to resolve the simulcast bandwidth issue, and, initially at least, the bandwidth premium for IP video streams may be minor. In addition, storage is relatively cheap, and the incremental additional storage, even for multiple formats, adds up to something palatable given that there is a growing library of HD content that must also be managed.

In time, of course, with the transition to IPTV becoming the video delivery endgame, and DOCSIS being the HFC vehicle for so doing at least in the foreseeable future, the bandwidth management bubble will involve the pace of the retiring of MPEG-TS delivery while increasing IPTV delivery. Clearly, introducing the latter (IPTV) must come before the acting on the former (MPEG-2 TS) if the transition is not to be an abrupt one. Equal or better video QoE in the IP domain will be necessary, and that means SD and HD delivery, creating a non-trivial problem. With the deployment of SDV, both technologies can take advantage of switched multicast statistically, so each has these built-in efficiencies. However, it is worth pointing out that the SDV QAM count for virtually "infinite" content for SD must be multiplied by roughly four for HD content.

Thus, a key element of an MSO strategy we expect to encounter moving ahead is an effective strategy for managing the "simulcast bump" that new services create. The bump is likely to encountered as a series of smaller bumps to hurdle at different times.

Video Services – Still on the Move

Let's move onto the content itself. The bandwidth hog is, and will continue to be, what is required to support consumer video expectations. Today, that expectation is satisfied by 1080i HD. However, 1080p HD is not far behind, adding to the per-stream average rate. Resolutions higher than HD are also on the drawing board (e.g. Ultra-HD – 4k x 2k) that can be 4x or more the total pixels of HD, or 4x the raw bandwidth. It is likely to be accompanied by compression advances, but also likely not 1:1 with resolution increases, in particular given the time it takes to develop encoding technology.

The maturation of MPEG-4 is the bit rate quiver available to battle the bandwidth bulge, in addition to the continuing use of current HFC tools of SDV, reclamation of analog, and plant expansion to 1 GHz. To the extent that a system objective may be to engineer for everything-on-demand, full unicast delivery, fiber deep supporting service group splitting supports the necessary per-subscriber bandwidth to enable this. That is, in order to support the demands of a unicast architecture from a traffic engineering of spectrum perspective, node splitting and segmenting to smaller and smaller serving groups makes it quite reasonable to deliver full downstream unicast under some pretty aggressive consumption assumptions.

Consider the following scenario first described in [4]:

Five simultaneous (viewing + recording) Ultra-HD streams in 3D (first generation) over MPEG-4. This scenario contains a mix of bandwidth killers and helpers:

- Ultra-HD is hungry as described above, but not "Super-Hi Vision," which is hungrier still
- Use of 3D adds bandwidth to capture the left-right eye perspectives. We do NOT account for the reduction of bandwidth over time.
- H.264/MPEG-4, though not widely used today in cable, will be in the time frames of this projection
- Five streams is clearly aggressive, but of course U-Verse today advertises four simultaneous streams (not all HD)

This adds up to slightly below 135 Mbps of CBR video services. Consider a very aggressive penetration @75%, a very aggressive concurrency of use @ 50%, and a 1 Gbps data service @1% oversubscription, typical for data access. For services provided by only today's 256-QAM, 6 MHz QAM channels, we can show that HFC can ultimately support this, as serving group sizes are reduced, as shown in Table 1 (Red-to-Yellow-to-Green having the obvious implication). In [4], techniques to exploit more coaxial bandwidth are also described, such that we can go beyond conventional wisdom of 5-6 Gbps of RF access bandwidth.

The use of CBR-only delivery is a significant conservative factor – adding the increased efficiency for VBR delivery previously described moves the solid bar of comfort level (Green-to-Yellow) northbound. This also implies that, similar to the prior discussion of using the introduction of MPEG-4 as an opportunity to inject transformative changes to the infrastructure, introduction of advanced services such as 3D, at lease in scale, might best be implemented within the context of the IPTV transition.

HHP/Node	Req's GBPS	QAMs/Node	Spectrum (MHz)
250	14.53	384	2304
180	10.46	276	1656
125	7.27	192	1152
75	4.36	115	690

Table 1 – Supporting an Extreme Services Mix Over HFC

Data …… or Video?

As described above, we can scale downstream data services to 1 Gbps under reasonable traffic metrics and support that growth over HFC. Whether bonding of DOCSIS QAMs is the most effective way to do that longer term is questionable, or even whether such a service rate should be an RF solution, as opposed to a fiber solution. For commercial service to large enterprises, the Gbps rates make sense. However, for this customer set, FTTP solutions overlaid onto the HFC are an effective alternative to burning RF residential bandwidth. There is an expectation with residential services that rates will continue to climb on average at some 20-40% compounded rate. Perhaps more importantly, however, is that concurrency factors will shift as the HSD content shifts from web pages to video clips to OTT video. This would require allocating more bandwidth, linearly with the increase in concurrency. Fortunately, if allocated for a service of 1 Gbps at 1%, the same math holds true for 100 Mbps at 10%. Similar to how data growth was estimated in [1], a compounded growth rate assumption looks still to be a useful way to capture growth. If it can be reliably broken into component pieces of rate growth and concurrency, it perhaps becomes a better tool for understanding total bandwidth needs going forward.

On the upstream, while audio file sharing was all the rage driving bandwidth in [1], it is less so today because of, well, the law, and itunes. Nonetheless, rapid upstream traffic and service tier expansion has continued as a strong element of HFC planning. As in downstream, video services upstream could be bandwidth busters, and many possible applications that might accelerate this have been kicked around for many years. They may or may not take hold in a big enough way to matter, but over the next several years at least, there is a more important concern for upstream than postulating about new applications. That issue is simply staying ahead of normal growth in the available 5-42 MHz of very imperfect spectrum, and preferably working the service rate up to 100 Mbps – a logical market target to support broadly, and recently set as an FCC national objective. It is not a simple thing to accomplish 100 Mbps within today's spectrum. And, it is not difficult to show that under compounded growth assumptions, the upstream lifespan runs out of steam before the downstream, and within a period of time to be planning for next steps for upstream bandwidth strategy [5].

CONCLUSION

Ten years ago, an analysis was undertaken to project bandwidth requirements on HFC, with the intent to derive bandwidth needs for a decade's worth of growth. It was hoped that MSOs could use the information as a planning tool. Several of the individual services were predicted quite well in how they would scale (DTV, VOD, HSD), some were significantly underestimated (HD), and some were missed completely (SDV). The analysis offered many unquantified trend projections that turned out to be quite prophetic. Finally, macro bandwidth projections due to the services growth described turned out to be reasonable, the prediction that RF plant had runway to support the growth was on target, and the projection of timing and logic for node splitting for service group reduction also turned out to be pretty accurate, all things considered.

The resulting assessment leads to a couple of important, confidence-building conclusions:

1) A sound understanding exists on the important, larger picture of the behaviors of consumer broadband consumption and growth, and the business implications

2) The flexibility and scalability of the HFC architecture has held to be as powerful as anticipated. This can be traced to HFC being built with just the right number of component parts, each of individual scalability and interoperability, enabling incremental investment that can be simply implemented, and ultimately rapidly paid for and profited from.

3) We have our work cut out for us for the next ten years to be so prophetic!

The job now revolves around item 3): quantifying the next set of projections based on the incoming mix of new possibilities described – video services, data speed trajectories and potential for "killer apps," and the introduction of commercial and wireless services support. At Motorola, we continually update our thoughts on where the bandwidth comes from and where the access network goes over time to support it. As we put behind us this last decade of growth, we similarly project the next decade by combining the lessons learned herein, the anticipated mix of new services, along with a dose of reality check of where on the hype curve these new services exist, and what these factors are likely to project to relative to mass adoption.

REFERENCES

[1] Randy Nash, "A 10-year Residential Bandwidth Demand Forecast and Implications for Delivery Networks," 2002 National Cable Television Association (NCTA) Show Proceedings Technical Papers, New Orleans, La., May 5-8, 2002.

[2] Dr. Robert Howald, "Web Surfing to Channel Surfing: Engineering the HSD Edge for Video," 2009 Cable-Tec Expo, sponsored by the Society for Cable Telecommunications Engineers (SCTE), Denver, CO, Oct 28-30, 2009.

[3] John Ulm and Patrick Maurer, "IP Video Guide – Avoiding Potholes on the Cable IPTV Highway," 2009 Cable-Tec Expo, sponsored by the Society for Cable Telecommunications Engineers (SCTE), Denver, CO, Oct 28-30, 2009.

[4] Dr. Robert Howald, "Fueling the Coaxial Last Mile,"2009 Society for Cable Telecommunications Engineers (SCTE) Emerging Technologies Conference, Washington, DC, April 3, 2009.

[5] Dr. Robert Howald, "The Broadband Horizon," Cable Next-Gen Broadband Strategies: Docsis 3.0, Wireless, Fiber & Beyond, Denver, Co., Feb 25, 2010.

ACKNOWLEDGMENTS

The author would like to thank the following individuals for their insights and support in gathering data critical to developing this paper: Curtiss Smith, Mohsen Manoochehri, Tony Stormer, Michael Brannan, Jack Moran, John Ulm, Fred Slowik, Chris Bastian (Comcast).