THE COST OF FAIR BANDWIDTH: BUSINESS CASE ON THE RISE OF ONLINE VIDEO, P2P, AND OVER-THE-TOP; PROACTIVE STEPS OPERATORS CAN TAKE TODAY TO MEET BANDWIDTH DEMANDS IN THE NEAR FUTURE

Michael Eagles, UPC Broadband
Robert F. Cruickshank III, C-COR

Abstract

The growth of popular video sharing sites such as YouTube, P2P file sharing applications such as Bit Torrent, and over-the-top broadband services has placed unprecedented strain upon the HFC network. Demand for the broadband pipe is growing daily and in order to ensure that all services and applications are delivered reliably, operators must plan for the traffic increase and take proactive steps to mitigate the rising demands.

This paper will take a multi-pronged view of the demands placed on today's network, including a recent history traffic analysis of sample systems and projections of where networks are headed tomorrow. Additionally, the paper will investigate both hardware and software options for proactively working to meet bandwidth demands today and over the next three to five years.

INTRODUCTION

Increasing Bandwidth Needs

With operators launching many bandwidth intensive services including High Definition Television (HDTV), Video on Demand (VOD) and ever-increasing High Speed Data rates, there is an ongoing need for greater capacity. Here we take a close look at High Speed Data and the rise of online multimedia devices and applications to better understand current and future demand trends.

In this paper we use the term “bandwidth” to mean “RF bandwidth” and “information rates” to mean advertised product speeds. We discuss “consumption” in the context of usage trends and subscriber “bit-rates” or average speed in Kilobits per second (Kbps).

Why Look at High Speed Data?

High Speed data is one of the few services where forces, largely external to the cable operator, shape the demand for capacity. In the recent rise of online video, P2P and over the top multimedia services, we notice a trend from the use of High Speed Data for “access” toward its use for “entertainment”. In this paper we seek to understand the impact of this shift on bandwidth needs.

Methodology

We considered the level of congestion on today’s network by looking at a snapshot of global consumption patterns in addition to a specific per system example.

In assessing the recent and projected growth in demand and bandwidth needs we considered both historical trends in per subscriber average capacity, per subscriber profiles and device and application trends.

Finally we considered a tool-kit of alternatives for the cable operator, grouped into “manufactured bandwidth” or capacity expansion and “technology bandwidth” or capacity management.

AN ANALYSIS OF TODAY'S NETWORK DEMANDS

In this first section of the paper we explore a snapshot analysis of the demands placed upon the HFC network by identifying global trends and considering data services in a major metropolitan city in the United States.

Average Usage per Subscriber

We started with a review of global usage trends using a base of 2.75 million DOCSIS® Cable Modem (CM) devices. As shown in Table 1, the average high speed subscriber sends/receives about 300 Megabytes (MB) of IP traffic per day. As one might expect, average daily usage varies among different regions of the world, but not drastically.

	Average Daily Directional Usage [MB]		
Region	Avg. Up	Avg. Down	Avg. Up+Down
USA	108	210	318
W Europe	123	192	315
E Europe	95	183	278
S America	138	176	314
Asia	98	190	288

Table 1. Typical Average Daily CM Usage by Region during November 2006[i]

In order to identify sources of congestion on today's network we take a closer look at how this average usage is distributed across the user base and what attributes contribute to this usage.

Usage by "Power Users"

A common belief is that some subscribers send/receive more or less traffic than other subscribers. A logical follow-on question is "What percentage of subscribers send/receive more than others?" Figure 1 shows the cumulative percentage of traffic sent/received by ~500,000 CMs during June 1995 in a major metropolitan city in the United States.

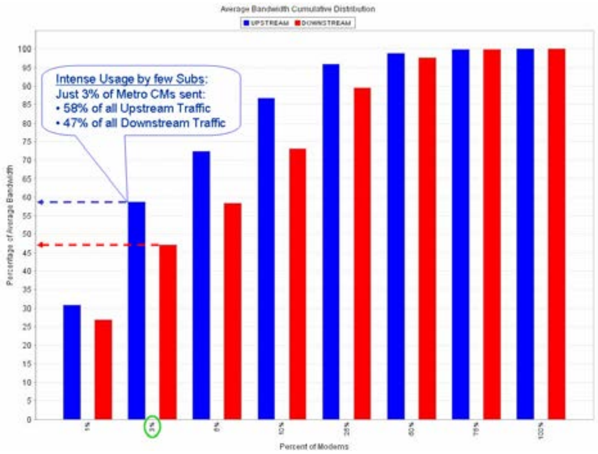

Figure 1. Average Bandwidth Cumulative Distribution in June 2005

Referencing Figure 1, one can see that just 3% of the CMs (on horizontal axis) sent 58% of upstream traffic (on vertical axis in blue) and received 47% of the downstream traffic (on vertical axis in red).!

Fast-forwarding almost two years to today, we see similar results in the upstream—just 3% of the CMs sent 56% of upstream traffic. But look, over the same time period the downstream number has dropped from 47% to 37% as shown in Figure 2.

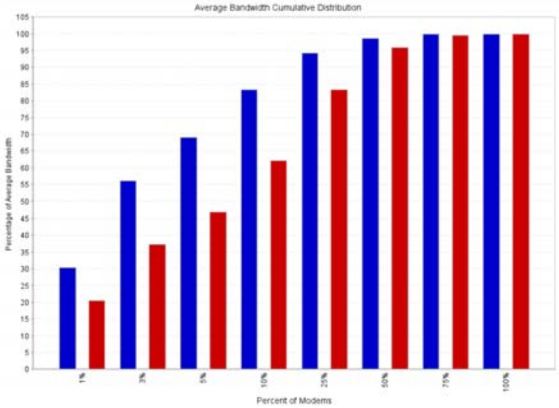

Figure 2. Average Bandwidth Cumulative Distribution March 2007

The takeaway messages from Figures 1 & 2 are: a) half to three-quarters of all IP traffic is sent/received by just 5% or less subscribers, and the remaining 95% of subscribers send/receive the balance of IP traffic on MSO DOCSIS® networks; and, b) the distribution of volume has changed due to an increase in downstream consumption by the "average" user.

We believe a possible reason for the distribution change is that bandwidth intensive, multi-media applications are becoming widely used by mainstream subscribers rather than only early adopters.

Active vs. Total Cable Modems

Another question is "What percentage of CMs are active throughout the day and night?" Figure 3 shows the number of active and number of total DOCSIS devices (CMs) on a typical Metro upstream hour by hour over 1 week. The total number of devices (across the top) is relatively consistent and ranges from 216 to 220 (the number varying as CMs become unreachable, for example, when subscribers shut down/power up their CMs). The active number of devices (lower periodic trace) varies throughout day from 63 to 148. Dividing the number of active devices by the number of total devices at any hour yields the % of active devices (upper periodic trace) which is read on the right-hand vertical axis.

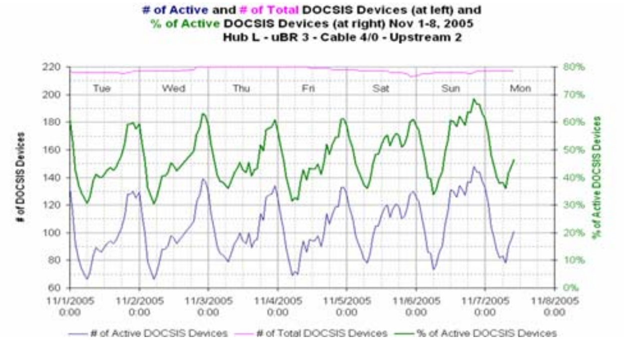

Figure 3. Active & Total Devices on a Typical Metro Upstream in 1 week

Notice in Figure 3 that the percentage of active devices never falls below 31%; put another way, about one-third of all CMs actively send/receive traffic at all times.

Zooming out to a whole year in Figure 4, one can see the percentage of active devices grew more than 10% in 12 months. The takeaway message is that over time a greater percentage of the overall CM population actively send/receive traffic (are active) around-the-clock. We discuss possible reasons for this in a later section.

Fast-forwarding almost two years to today, we see results quite similar to the trends shown in figures 3 and 4.

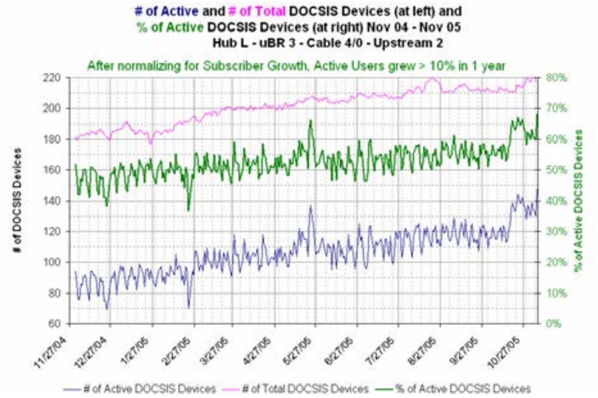

Figure 4. Active & Total Devices on a Typical Metro Upstream in 1 year

Rise in Measured Congestion

A greater percentage of CMs send/receive more IP traffic more of the time, resulting in a rise in HFC congestion. The top half of Figure 5 shows the hour-by-hour rise and fall of congestion across all of Metro's 500,000 CMs over one week. Notice that peak congestion occurs daily at 8-10 PM. The highest congestion level was late Sunday night when about 40,000 of the 500,000 Metro CMs experienced slowdowns due to congestion.

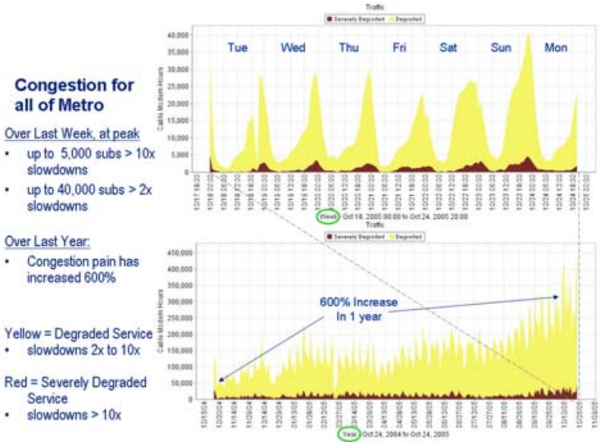

Figure 5. Weekly & Yearly effects of Congestion across all of Metro

The bottom half of Figure 5 shows the day-by-day rise and fall of congestion across all of Metro's 500,000 CMs over an entire year. The vertical axis now represents the number of modem congestion hours (the sum of the number of CMs congested per hour for all 24 hours in the day). Over the course of the year congestion levels rose nearly 600% from an average of ~50,000 modem hours per day to ~300,000 modem hours per day.

What is the Relationship between Congestion and Bandwidth?

Even when it is assumed that the traffic characteristics do not change, an increase in the available bandwidth to the subscriber - the "information rate" - results in peak rate requirement increases. As a result, the network can serve fewer subscribers at the same level of congestion or the same subscribers with increased congestion.

In considering the recent and projected growth in demand, we consider both historical and projected information rate increases that contribute to congestion as the speed of the sources (CM's) increases. In addition, we consider the impact on traffic characteristics caused by new applications and devices that create/consume IP traffic.

RECENT AND PROJECTED GROWTH OF THESE DEMANDS

In this section of the paper we explore a historical view of the increase in information rates and the impact of new multimedia online services such as YouTube. By exploring recent growth, we will extrapolate the magnitude of traffic that can be expected within the next three to five years for residential broadband data.

The Digital Household

We believe there are two primary factors at work contributing to the increase in measured congestion on HFC networks.

(1) Growth in Information Rates: Using historical trends as a base-line we note continued increase in information rates that lead to an increase in measured congestion. This is fueled by the competitive nature of the broadband access market.

(2) Growth in Multimedia Applications and Connected Devices: New multimedia applications, including embedded multimedia, P2P television, and TV place-shifting, result in more bits being transported. In addition, wireless home gateways are entering the mainstream, and we see growth in the number of IP-connected, traffic-generating devices behind each CM.

We see these two factors working together to create significant additional demand for network capacity.

Predicted Information Rates

To predict the recent and projected growth of demands on the HFC network we also considered an analysis by Bob Scott [ii] that extrapolates historical growth in data product advertised speeds since 1982. This analysis notes a close fit to Moore's Law and provides the following high speed data product information rate scenarios shown in Figure 6.

Using this model we start in 1982 with dial-up telephone Modem speeds of 300 bits per second (bps) and climb to nearly ~20 megabits per second (Mbps) today.

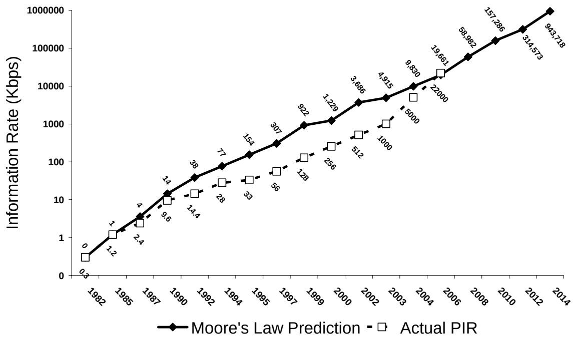

Figure 6. Moore's Law Growth in Historical and Predicted Information Rates

A sensitivity analysis from this work indicates a range of information rate or "speed" possibilities as shown in Table 2. We assume a subscriber take-rate of the different speeds in Table 2 - High Case of 5%, Probable Case of 10%, and Low Case of 85% - reflecting a mix of subscribers using different broadband product speeds to create a weighted information rate each year.

	2006 (Kbps)	**2008** (Kbps)	**2010** (Kbps)
High Case	20,000	60,000	160,000
Probable Case	10,000	20,000	60,000
Low Case	7,000	10,000	20,000
Weighted Average	8,000	14,000	30,000

Table 2: Predicted Information Rates in Kbps

We expect this increase in CM information rates to increase congestion and required bandwidth even if the bits transported and the number of CMs stays the same between 2006 and 2010; and we expect the impact on required bandwidth and congestion to be higher in the case that the bits transported increases through new applications and additional devices per CM.

Demand Model and Usage Profile

An analysis of traffic profiles described by using John T. Chapman's Multimedia Traffic Engineering Model published in 2002 indicates 1 to 2 traffic generating devices per CM [iii]. Recent research suggests that this has increased to over 3 devices [iv].

We assume an increase from 1.5 devices per CM in 2006 to 3 devices per CM by 2010 explained by increased market usage for new broadband connected applications.

In addition to considering historical per subscriber Kbps data we considered traffic contributions from existing applications on 1.5 connected devices per CM plus contributions from an additional 1.5 connected devices supporting several new applications including Embedded Multimedia (i.e. YouTube); P2P Television (i.e. Joost); VOIP Applications (i.e. Skype); TV

Placeshifting (i.e. Slingbox), and CE Devices (i.e. Connected Handycams).

Using this traffic engineering model in conjunction with recent new applications and devices we are able to estimate the future application traffic in average Kbps per data subscriber from approximately 80 Kbps in 2006 to 350Kbps in 2010 as shown in Figure 7.

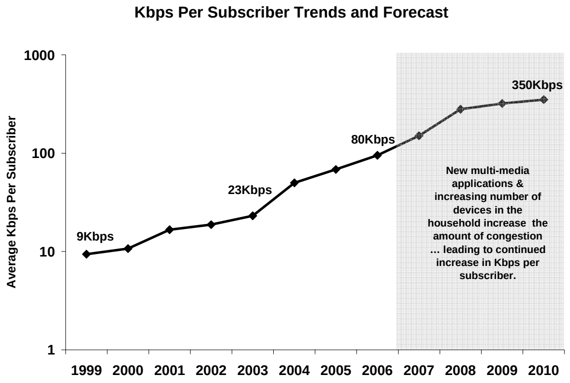

Figure 7. Kbps Per Subscriber Trends [v]

We look in more detail below at changes in the application mix over time to support this.

Changes in Application Mix over Time

We asked the question "What application changes are likely considering historical trends?" Based on information illustrated in Figure 8 we believe that P2P in its current 'download' form may have reached its peak and that a 'TV centric' and new multimedia web applications may provide the next wave of traffic.

Considering the rapid rise of P2P traffic, we note that the primary traffic generating application may not even exist at the time the network is being planned! The overall mix and contribution of various traffic-generating applications rises and falls over time as shown in Figure 8 and Table 3.

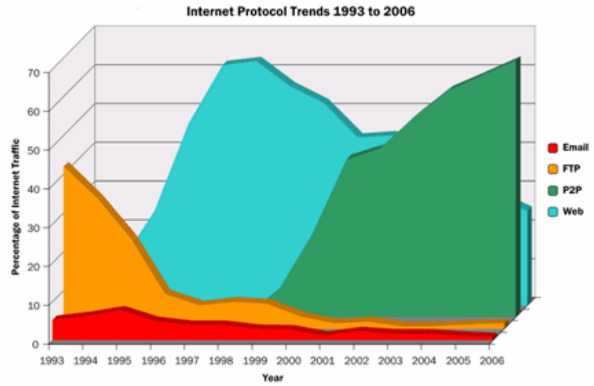

Figure 8. Internet Protocol Trends [vi]

1993	**1999**	**2006**	**2010 ?**
FTP	Web	P2P	P2P-TV
Email	FTP	Web	Rich-Web
Web	Email	News [vii]	P2P-Other
Other	Other	Other	Other

Table 3. Internet Protocol Trends [viii]

To understand two possible dominant applications we explore the growth of Embedded Multimedia Traffic and PC-Based Television and P2P Multimedia which in our model contribute a significant amount of the new traffic.

Embedded Multimedia Traffic:

Recent developments in web applications include the addition of video-based embedded multimedia. Examples include Brightcove, Coull, Google Video, MediaZone, MetaCafe, PermissionTV, Revver, VideoJug, YouTube, Yahoo! Video, and Ziddio. We explore YouTube in more detail for this category.

Google's YouTube founded in Feb 2005 [ix] provides flash-based video clips embedded in a web browser. Measurement of Google's YouTube traffic in Europe indicates that YouTube represented 2.5% of all downstream traffic year end 2006 [x], and growth of YouTube traffic was 400% for the 6 months ending 2006.

Normalized for subscriber growth this

translates into 1 Kbps per subscriber today growing to over 100 Kbps per subscriber in just 12 months. Is growth as shown in Figure 9 realistic?

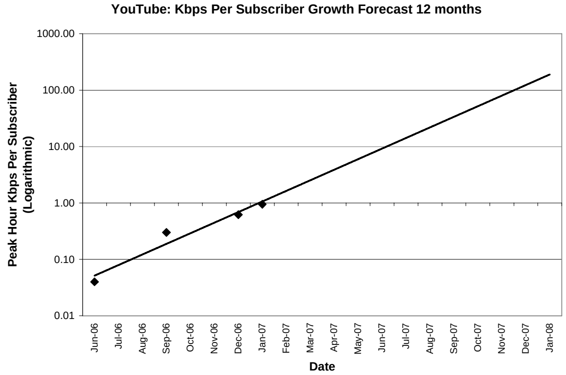

Figure 9. Recent You Tube Traffic Growth

We believe that, more than likely, embedded multimedia traffic growth per subscriber will slow in the future based on reaching maximum market usage for a given broadband subscriber population.

Using the Chapman Multimedia Traffic Engineering single user profile, we assume that by 2010 an increase in embedded multimedia web content including adoption of Google's Click-to-play[xi] advertising model, results in an embedded multimedia user base reaching market usage of 20% of broadband subscribers on average; and that the average bit rate is 250Kbps. This equates to 50Kbps of downstream traffic at maximum market usage using our assumptions.

PC-Based Television and P2P Multimedia

Another category of growth in multimedia applications is PC-based Television and P2P Multimedia. Recent developments include the introduction of several PC based multimedia online platforms including, Arvato, AOL's In2TV, Babelgum, BitTorrent, Grid Networks, Kontiki, Itiva, Joost, JumpTV, MediaZone, Network Foundation Technologies (NFT), Peer Impact, Red Swoosh, and Rawflow [xii].

We note that many P2P applications generate traffic from a series of machine queued requests. Machine-to-machine traffic, as the name suggest, indicates that the application usage continues even without human activity (i.e. mouse-click) present as shown in Figure 10. Perhaps this helps explain the rise in active modems shown in Figures 3 & 4.

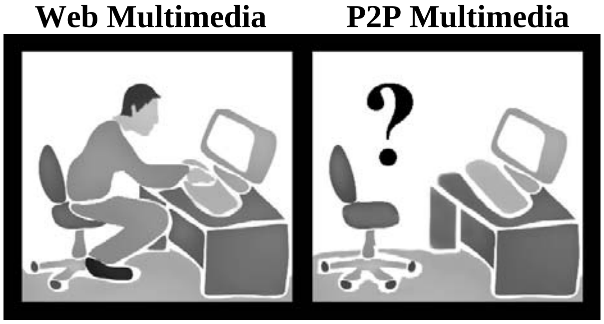

Figure 10. Web/P2P Multimedia Usage

Using the Chapman Multimedia Traffic Engineering single user profile we considered that users watching television have a peak time (on-peak) market usage of 60% and an off-peak value of 33% of broadband subscribers by 2010 based on the Nielsen "Households Using Television Number" [xiii]. We assume the average for P2P video is 40% and a blended average bit rate of ~550Kbps in the downstream and ~125 Kbps in the upstream. This equates to 220Kbps per subscriber in the downstream and 50Kbps in the upstream.

Today P2P video applications are PC-based. However, in the future such an application may find a low cost hardware "host" vehicle such as a set-top, thus significantly increasing market usage levels.

We believe a future statistical projection of Joost[xiv] traffic would assist in identifying growth patterns and a likely market usage peak.

Comparing Each Approach

We considered both Moore's Law historical information rates, keeping congestion constant, and Chapman's Multi-Media Engineering model to create a single CM profile.

	Historical Information Rates and Moore's Law	Single CM Profile using Chapman's Model
% Increase in Bit-rate 2006 to 2010	300% (320 Kbps)	341% (353 Kbps)

Table 4: Comparing Each Approach

For the purpose of determining capacity requirements we noted that Chapman's model generated a greater increase in the per user bit-rate and was therefore used as the basis for our 2010 assumptions.

Predicted Capacity Requirements

As a result of the application analysis above we are able to estimate the amount of additional future downstream and upstream capacity that will be required.

Assuming the information rates offered by the Moore's Law projection and the Kbps per subscriber generated by new applications using the Multimedia Engineering model, this suggests the following over-booking factors as shown in Table 5.

	2006	**2010**
Avg. Speed	8,000	30,000
Kbps	80	350
Over-booking	100	85
Subs/8MHz	517 (~500)	116 (~100)

Table 5: Speed and Kbps per Subscriber

We assume that 20% of the 50 Mbps (8 MHz) downstream port capacity in 2006 and 2010 is reserved for: a) other services (i.e. eMTA based VoIP and STB return path), and b) additional growth. This results in 40Mbps available for use per 8MHz channel. We also assume statistical multiplexing gains of 1.1 from bonding channels.

In 2006 we support 500 subscribers per 8MHz downstream at 80Kbps per sub compared to 2010 where we are able to support 100 subscribers per 8MHz downstream due to the additional of new applications and new devices. To support the equivalent density per service group we require 32MHz of downstream capacity and 25.6MHz of upstream capacity as shown in Table 6.

	2006	**2010**
Number of DS	1 x 8MHz	4 x 8MHz
Number of US	2 x 3.2MHz	4 x 6.4MHz

Table 6. Additional Capacity Required

That's 4 channels x 50 Mbps/ channel = 200 Mbps or 32MHz of total capacity in the downstream as outlined in Figure 11 below.

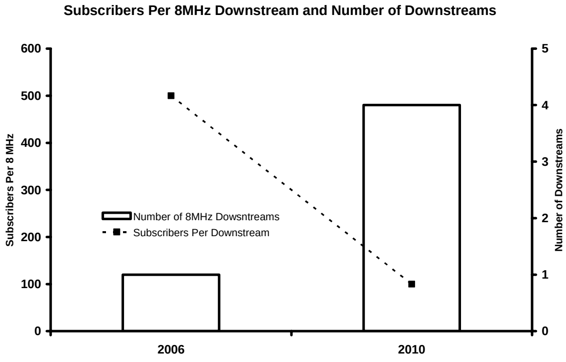

Figure 11. Additional Downstream Capacity

Additionally we would require 60 Mbps or 25MHz of upstream capacity as outlined in Figure 14 below.

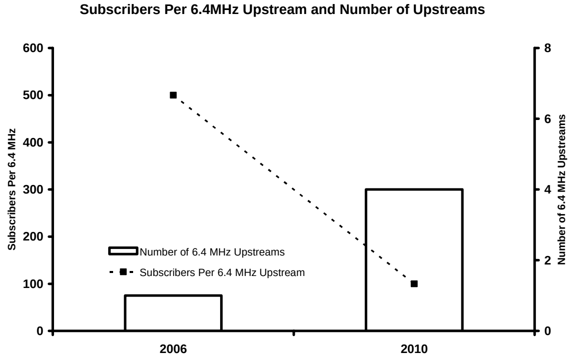

Figure 12. Additional Upstream Capacity

In summary, the growth in usage per subscriber and changes in applications together reduces the net number of subscribers per 8MHz downstream from 500 to something more like 100 as shown in Figure 12.

MANUFACTURED BANDWIDTH - EXPAND CAPACITY TO MEET DEMAND

Given the usage expected today and within the next few years, how can operators prepare from an HFC perspective? This section outlines the cost of various Capacity Management tactics including:

(a) Reducing Serving Groups

(b) Balancing Users across channel line-up

(c) Increasing Data rates

(d) Up-Selling Subscribers

(e) Managing Subscriber Traffic

Each of the above tactics may have one or more options as shown in Figure 13.

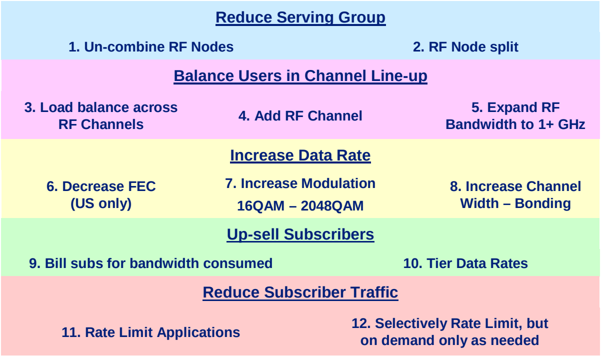

Reduce Serving Group
1. Un-combine RF Nodes
2. RF Node split

Balance Users in Channel Line-up
3. Load balance across RF Channels
4. Add RF Channel
5. Expand RF Bandwidth to 1+ GHz

Increase Data Rate
6. Decrease FEC (US only)
7. Increase Modulation 16QAM – 2048QAM
8. Increase Channel Width – Bonding

Up-sell Subscribers
9. Bill subs for bandwidth consumed
10. Tier Data Rates

Reduce Subscriber Traffic
11. Rate Limit Applications
12. Selectively Rate Limit, but on demand only as needed

Figure 13. Capacity Management Tactics & Options

Reduce Serving Group

We observe there are two ways to reduce the size of a serving group: Option 1 is to Un-combine RF nodes so that fewer nodes (and subscribers) are connected to an upstream or downstream port. Option 2 is to split a RF node into two or more serving groups with fewer subscribers. In the limit, Option 2 would be considered a “Fiber Deep” approach to reducing serving group size.

Balance Users across Channel Line-up

As penetration grows it is safe to assume the existence of multiple channels. Option 3 involves spreading users across multiple existing channels to evenly distribute traffic loads. Option 4 involves adding additional channels (assuming spectrum is available), and Option 5 requires expanding the HFC delivery infrastructure to accommodate the transmission of higher frequencies (e.g., 1GHz).

Increase Data Rates

In DOCSIS systems, upstream traffic ‘packets’ may be ‘protected’ from packet errors by different levels of Forward Error Correction (FEC). In fact, it is possible to not use any FEC at all, though field experience dictates that some level of FEC should always be used. To successfully transport un-errored packets, clean portions of the HFC require less FEC and noisy portions require more FEC. While ensuring critical packets (such as voice) are delivered unerrored, FEC comes at a price; greater FEC (overhead) results in less User Data (payload).

Clean HFC plants can also support higher order modulation which means more bits (a larger alphabet) can be sent at a given symbol rate. The take-home message in Options 6 & 7 is that cleaner HFC carries more User Data than noisy HFC.

Option 8 is the much talked about solution known as channel bonding wherein multiple channels are 'bonded' into a wider (super) channel with higher data rates and higher statistical multiplexing efficiencies

Up-Sell Subscribers

Options 9 & 10 imply that 'heavy' users who send/receive large amounts of traffic could/should pay more for service than 'light' users who don't send/receive much traffic. The thought is that pricing can/will discourage heavy use. Option 9 could be a unit measure that can be thought of as 'Bill by the Byte Transferred' which is akin to how many utilities/consumables are billed today (e.g., milk, water, gasoline, etc., each cost so much per litre.) Option 10 is similar to Option 9, but instead of billing per byte, bills on chunks of bytes (e.g., 0 ≤ 10 GB per month is $40, 10 ≤ 20 GB per month is $50, etc.).

Reduce Subscriber Traffic

Options 11 & 12 refer to methods used to limit certain or all users' traffic and/or traffic types. Before getting alarmed at either Option, take a moment to consider a few practical realities. All highways slow down when they get full; Internet pipes slow down when they get full, too. Option 11 involves rate limiting specific applications (such as P2P) around-the-clock, whereas an example of Option 11 is an on-demand as-needed rate limiting algorithm that enforces application or individual subscriber rate limits ONLY as needed (i.e., at the onset of congestion or when congestion persists).

The Cost of Capacity Expansion

We list the relative capital cost of some of the different capacity expansion tactics & options in Table 7. The important factor is that costs can vary dramatically between expansion options, and that a combination of different options can be considered (i.e. Switched Digital Video & Bonded Channels).

	$ Per Node Per MHz of Capacity	$ Per HP at 2,000 HP Per Node	$ Per Data Sub (25% pen.)
Upgrading to 1GHz [xv]	$714	$50	$200
Fiber Deep [xvi]	$107	$46	$184
Node Split [xvii]	$30	$13	$52
Switched Digital [xviii]	$312	$5	$20
Bond Channels [xix]	$632	$18	$70

Table 7. The Cost of Manufactured Bandwidth

TECHNOLOGY BANDWIDTH - MANAGE SCARCE RESOURCES WITH NEW TECHNOLOGY

Given the projections, expansion approaches in isolation will not be enough to maintain near-term demand. In addition, a software management approach to directing and managing traffic on the network will stretch the increases gained with HFC equipment upgrades, node splits, etc. This section outlines sample traffic management applications that ensure bandwidth is fairly provisioned so that high priority traffic such as VoIP calls are given top priority, while minimizing the impact of P2P and over-the-top applications.

We see four possible instruments to allocate scarce bandwidth:

(1) Rationing

(2) Delay

(3) Pricing

(4) Access

(5) Supply

Rationing

Operators are able to limit volume and access to certain applications. This can be hard rationing or soft (e.g. first warn heavy users and if behavior doesn't change, then deliver a penalty).

For example, dynamically re-provisioning a subscriber at a lower speed when consumption reaches a pre-set limit is a possible way of making additional bandwidth available to remaining subscribers.

Rationing has the disadvantage of penalizing all applications even those provided by the operators such as email, and also penalizes the subscriber for heavy off-peak consumption that may not impact operator capacity at peak time.

Delay

It is possible to delay the delivery of non-revenue generating consumption though the use of traffic shaping which may be time-based, user-based, protocol or application-based—and therefore 'fair' in the sense that it addresses specific areas of bandwidth consumption.

It is critical that the delayed application can be accurately identified and that a mechanism exists to detect a move by users to alternative applications for the same use. (i.e. a move from a P2P protocol to Network News (NNTP) protocol when wanting to download multimedia content).

Pricing

Pricing may be used to impact consumption and may be based on volume, nature of the service and across the day (as with the electricity industry).

Application specific charges can also impact traffic consumption. For example, it may be possible to offer certain P2P services within the monthly broadband service fee and to charge for access to others.

Alternatively it may be possible to charge a carriage fee for premium P2P content providers in return for a fair allocation of bandwidth.

We believe that pricing is one of the most effective mechanisms for managing heavy bandwidth consumption.

Access

With the growth of wireless home access points, theft-of-service can lead to increased traffic levels associated with a subscriber, even without the knowledge of the subscriber. Offering a managed home network enables the operator to reduce the risk of theft and improves the security offered to the end user.

Supply

Packet Cable Multimedia (PCMM) enables bandwidth supply control at the CMTS when applied to specific applications or service tiers. This enables the operators to monetize Quality of Service. In fact, it would be possible to package Quality of Service as a product.

Additionally as noted in Figure 14 it is possible to manage the supply of bandwidth through the use of a virtual or real two-tier backbone. Using this structure content providers and subscribers that pay for Quality of Service are able to access additional bandwidth.

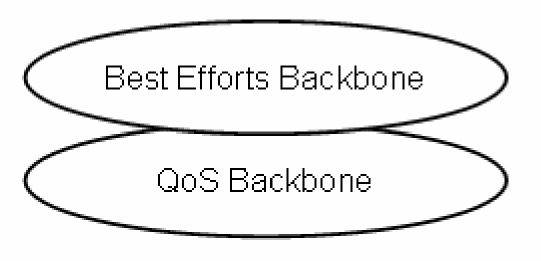

Figure 14. Two Tier Backbone Helps to Manage Supply

The Cost of Capacity Management

Reviewing the capital cost of capacity management in more detail we see that the cost of effectively managing capacity is far lower than capacity expansion. It follows logically then that options for managing this scare resource should be explored before/as expansion is considered. Table 8 provides examples for consideration.

	$ Per Node Per MHz of Capacity	$ Per HP at 2,000 HP Per Node	$ Per Data Sub (25% pen.)
PCMM [xx]	$166	$1.25	$5
Traffic Shaping [xxi]	$20	$0.15	$0.60

Table 8. The Cost of Technology Bandwidth – Managing a Scarce Resource.

© Mike Baldwin / Cornered

"You wouldn't believe how many people try to drive away without paying."

Figure 15. Bandwidth Expansion and Management is Not Cost-Free to the Cable Operator.

CONCLUSION

We outlined a trend in the rise of measured congestion and provided an overview of the historical and future application demands causing this congestion, and offer suggestions and relationships (a toolkit) that the cable operator may want to consider. We note that while HFC networks offer great flexibility in expanding bandwidth, not all alternatives are created equal; some are easier and more cost effective than others. A systematic and thoughtful approach to addressing congestion should include evaluating all possible options—considering both the subscriber and the content provider in view of both manufactured bandwidth and technology bandwidth alternatives.

REFERENCES

---

[i] UPC Broadband & C-COR internal reports based on CableEdge® measurements.

[ii] Bob McIntyre, "The QAM Before The Storm", SCTE Emerging Technologies Conference, January 2007

[iii] John T. Chapman, "Multimedia Traffic Engineering: The Bursty Data Model" , Cisco, January 8, 2002

[iv] Ian Fogg, "Home Networks in Europe: Leverage Strong Growth Potential in Congested Digital Device Landscape", Jupiter Research, Jan 2007

[v] Based on historical Kbps per subscriber growth in Europe. Forecast Kbps numbers based on High Speed data applications only and ignores VOIP and Digital Set-top return path traffic which are assumed to be included in the remaining 20% port capacity.

[vi] Andrew Parker, "The True Picture of Peer to Peer File Sharing", CacheLogic, 2004. http://www.cachelogic.com/home/pages/studies/2004_01.php

[vii] A revival in NNTP 'news' traffic has been observed recently in Europe primarily as a mechanism to share multi-part binaries.

[viii] CacheLogic study: The True Picture of Peer to Peer File Sharing. Andrew Parker, 2004. http://www.cachelogic.com/home/pages/studies/2004_01.php

[ix] YouTube website http://www.youtube.com/t/about

[x] Based on UPC Broadband Netflow measurements in Europe using YouTube's AS number.

[xi] Google press release: "Lights, Camera, Action!! - Google Introduces Click-to-Play Video Ads on the Google Content Network" May 23, 2006. http://www.google.com/intl/en/press/annc/clicktoplay.html

[xii] Wired Magazine Feb 2007 and industry sources.

[xiii] James Hibberd, "Cable Late-Night Growth Outpacing Prime - Competition From Other Platforms Creates Opportunity for Daypart", TelevisionWeek, August 7, 2006. Chart notes Nielsen Media Research 2005 Prime-time households using television rates at 60.2% in 2005 and Late-night households using television levels at 33.9%. http://www.tvweek.com/article.cms?articleId=30364

[xiv] Joost is a newly announced television sharing application that utilizes P2P communication to deliver high-quality television across the broadband network.

[xv] We assume the average cost to upgrade to 1GHz is around $45-$55 per HP. We assume the 1 GHz amplifier gain is higher than those at 870 MHz and permit a direct drop-in at existing amp locations.

[xvi] Fiber Deep calculation based on an assumed a Labour of 41,000 and materials of 50,000 per 2,000 Home Passed or a total of 91,000 per 2,000 Home Passed or $46 per Home Passed. Assumes Greenfield … cost would be lower if fiber is already in place.

[xvii] Assumes the worst case costs for a new node, including HE electronics, and new fiber if needed, would be $30,000. Numbers may vary depending on specific labor costs, electronics, and market.

[xviii] Switched digital assumes 80% digital penetration, 2,000 homes passed per service group, 4 QAM channels per service group, 100 switched channels, 10 channels re-encoded. Components include QAM modulators, SDV servers, Re-encoding, Multiplexing, Routing and transport. Assumes 4 x 8MHz of capacity can be released for other services.

[xix] Bonding Channel to Increase Bandwidth based on an assumed $10,000 per 8MHz DOCSIS channel today and assuming $10,000 per 4 x 8MHz DOCSIS channels in the future. We assumed 500 subscribers per downstream or $20 per subscriber. We also assumed $50 per CPE device.

[xx] PacketCable Multimedia: Assuming $5 per subscriber for Turbo Button and Gaming Application. Includes cost of Record Keeping Server.

[xxi] Adding Per-Protocol Traffic Shaping: Assuming 25 CMTS port per shaping device based on 900 Mbps per device and 50Mbps per downstream at 70% utilization or 35Mbps. Each shaping device costs $30,000 / 25 ports = $1,200 per CMTS / 4 nodes per downstream serving segment = $300 per node / 2,000 Homes Passed per node = $0.15 per home passed. We assumed that 25% of the total capacity could be reclaimed though shaping.