

VARIABLE BIT RATE VIDEO SERVICES IN DOCSIS 3.0 NETWORKS

Xiaomei Liu, Cisco Systems
Alon Bernstein, Cisco Systems

Abstract

DOCSIS 3.0 and Modular CMTS promise to provide ten times the bandwidth at one tenth of the cost, compared to existing CMTS technology. With these forward-looking trends, it is becoming increasingly viable to consider a channel-bonded DOCSIS network as a fully-converged network to transport video, voice and data. A bonded, converged and asynchronous data pipe, married with variable bit rate (VBR) video coding, can deliver the full potential of IPTV over cable.

This paper examines the technical and economic implications of VBR over DOCSIS. It proposes an IP-level VBR network statmux to deliver VBR video over channel-bonded DOCSIS and quantifies the efficiency of the network statmux with the results of lab tests. It provides insights into various architecture issues related to VBR delivery. Finally, it explores mechanisms to improve robustness and enhance the subscriber viewing experience.

INTRODUCTION

Cable networks are experiencing an explosion in demand for increased bandwidth. A significant amount of bandwidth pressure comes from High Definition Television (HDTV) service expansion, which MSOs have used as a strategic move to compete with satellite and telco video service providers. Today, the 100+ HD channel service is on the horizon as more HD content is offered. Meanwhile, content personalization and targeted advertisement are gradually transforming the video delivery vehicle from broadcast to unicast. Yet over-the-top video services and user-generated video content simultaneously

drive bandwidth demand with millions of video assets streamed or downloaded to PCs.

Cable operators have many tools to address the overwhelming bandwidth crunch problem. Some of these approaches in the MSO toolkit include: analog channel reclamation, switched digital video, node splitting, plant upgrades to 1GHz, and MPEG-4 part 10 video coding. Although cable operators can drill for additional bandwidth in HFC networks with major capital expenditure, there is work that can be done to eliminate any bandwidth inefficiency in HFC networks first.

Starting with video sources, it is common knowledge that variable bit rate (VBR) encoding of video is significantly more efficient than constant bit rate (CBR) encoding. In MPEG video encoding, while the CBR video keeps the bitrate constant, the VBR video attempts to keep the video quality constant. The nature of MPEG video encoding allows encoders to use fewer bits for simple scenes and more bits for complicated and motion rich scenes. With comparable video quality, VBR can yield 40 percent or more bandwidth savings over CBR [1]. As a result of its coding efficiency, VBR video is widely used in DVD and in broadcast video applications such as digital satellite and cable.

The introduction of broadcast-oriented MPEG statmuxes paved the way for delivery of VBR streams in broadcast video. Since most transmission channels have fixed bandwidth, MPEG statmuxes combine a number of VBR streams into a single aggregated constant bitrate channel. The statistical distribution of bitrate peaks and valleys allows the combined streams to use less bandwidth than what is needed if each

VBR stream is sent individually. At any given point in time, if the bandwidth of a VBR bundle exceeds the capacity of an MPEG transmission channel, the MPEG statmux applies requantization at the MPEG level to reduce the instantaneous bitrate of video streams to fit the transmission pipe. This action does come at the expense of a non-zero impact to video quality.

Ironically, in advanced video services such as Switched Digital Video (SDV) and Video on Demand (VOD) where the last mile bandwidth efficiency is needed the most, CBR instead of VBR video is deployed universally today. This is because SDV and VOD present a challenging case for traditional broadcast oriented MPEG statmuxes:

1) MPEG statmuxes are computationally intensive and costly. MPEG statmuxes achieve rate reduction by transrating selected MPEG frames. Transrating is an expensive operation as macroblocks in pictures are re-quantized and re-encoded at the MPEG level. Although the cost of MPEG statmuxes is not a concern in a broadcast network as the per stream statmux cost is shared among all the subscribers in the network, quite the contrary is true in the SDV and VOD world. In an increasingly unicast-based video delivery network, the per stream statmux cost now becomes a per-subscriber cost, which makes current MPEG statmuxes economically impractical to deploy at the network edge.

It is already a complicated operation to statistically multiplex MPEG2 encoded VBR streams; it is an even more daunting task to perform statmux on MPEG4 encoded VBR streams because of the incrementally intensive video computations involved.

2) MPEG statmuxes apply extensive stream analysis in order to mitigate the video quality degradation caused by the transrating operation. The stream analysis as well as the

transcoding operation introduces delays, typically on the order of 1 second. This long latency is more noticeable and undesirable for on-demand and interactive video services.

3) It is also no surprise that traditional MPEG statmuxes have difficulties dealing with encrypted content considering that the rate reduction techniques involved need to analyze and re-encode the stream content. For example, pre-encrypted VOD content makes the elementary MPEG stream inaccessible for transrating.

The business and technical issues pointed out above have forced cable operators to give up VBR efficiencies and opt instead for CBR video delivery in switched and on demand video services.

These assumptions change as video over DOCSIS becomes a reality. Not only does wideband DOCSIS provide an IP transport to MPEG video, it also brings along a promising new way of VBR statistical multiplexing.

DOCSIS 3.0 AND NETWORK STATMUXES

DOCSIS 3.0 is perhaps the most anticipated technology of the year in cable industry. DOCSIS 3.0 takes the DOCSIS beyond just an IP transport for data and voice services. IP video over DOCSIS is rapidly gaining traction with DOCSIS 3.0. In fact, some MSOs already have started market trials and deployments of IPTV over DOCSIS are in the planning stages.

Partly because of DOCSIS 3.0 channel bonding and DOCSIS 3.0 enhanced multicast, IP video over DOCSIS becomes a feasible technical possibility. Channel bonding, the most important feature of DOCSIS 3.0, makes the channel bandwidth a magnitude higher than before by allowing CMTSS to bond multiple downstream and/or

multiple upstream RF carriers in order to deliver higher bandwidth to the home. Today, eight channel bonded cable modems are already available in the market, enabling downstream bandwidth rates of around 300Mbps. This increased bandwidth capacity is essential to bandwidth-hungry applications such as standard definition (SD) and high definition (HD) video.

DOCSIS 3.0 enhanced multicast adds source-specific multicast (SSM) and Internet group management protocol (IGMPv3) support. In addition, multicast sessions can also be managed with quality of service (QoS) guarantees. Switched multicast in an IP/DOCSIS transport increases bandwidth utilization efficiency in the same way as switched digital video does in an MPEG transport.

Equally important in terms of their potential technological impacts on the industry are modular CMTS (M-CMTS) and universal quadrature amplitude modulation devices (QAMs). The separation of the DOCSIS media access control (MAC) and physical layer protocol (PHY) allows independent scaling of upstream and downstream bandwidth. Economics of scale will drive down the costs of universal QAMs, lower the overall solution cost of M-CMTS and make DOCSIS economically viable for IP video delivery.

IP Transport vs. MPEG Transport

IP transport distinguishes itself from MPEG transport in a number of ways. IP transport is asynchronous packet-oriented transport. IP networks also introduce jitter. MPEG transport, on the other hand, is synchronous transport. When MPEG transport streams are delivered over IP networks, receivers must remove network jitter in order to recover the original video source clock. IP set-top boxes (STBs)

normally have dejittering buffers that can tolerate 100ms of network jitter.

IP/DOCSIS 3.0 transport supports wideband transmission with bandwidth upper capacity limited only by cable modem technologies. Downstream bandwidth speeds of 300Mbps are enabled by today's eight-channel cable modems. It is only a matter of time before much higher bandwidths are available as Moore's Law keeps bringing down the cable modem cost. MPEG transport, on the other hand, which does not support channel bonding, has a bandwidth limitation of a single QAM channel. In North America, the bandwidth of a single QAM channel is capped at 38.8Mbps with QAM256 modulation. The Law of Large Numbers implies that statistical multiplexing efficiency improves as the number of VBR streams in the transmission channel increases. The proliferation of HDTV in households and the increasing number of HD streams in the network make wideband transport much more attractive for the purpose of statistical multiplexing. While the narrowband MPEG transport struggles to provide efficient statistical multiplexing of HD streams without compromising the video quality, the bonded DOCSIS 3.0 transport can easily support statistical multiplexing of HD streams with good statistical multiplexing gains and video quality.

IP networks also have built-in quality of service (QoS) capabilities. Cable modem termination systems (CMTSs) implement advanced DOCSIS QoS features such as admission and policy control, priority queuing, traffic policing, traffic shaping etc. These IP network features are readily applicable for VBR video delivery.

IP networks also are converged networks. They allow data, voice and video to be simultaneously delivered through the network. Converged networks provide great bandwidth savings just by enabling all

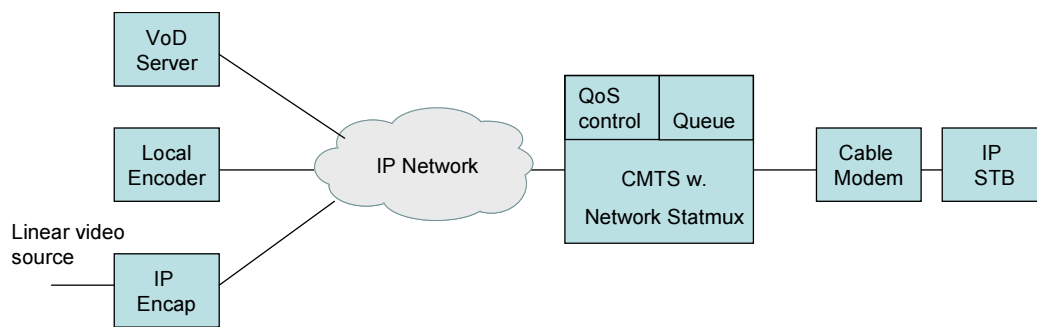


Figure 1 The VBR Network Statmux Architecture

services to share a single bandwidth resource pool. If the slightly different peak hours of these different services are also considered, the bandwidth savings are even greater. Field data indicates that bandwidth savings of 20 to 30 percent are achievable with a truly converged network. Even better is the fact that these savings are realized on top of the VBR over CBR bandwidth savings. In a converged IP network, if at a certain instant there is leftover bandwidth after all the VBR video traffic, then lower priority data traffic can consume the unused bandwidth. In other words, not a single bit of bandwidth is wasted! MPEG transport networks, however, are special-purpose networks used for video delivery exclusively. In an MPEG transport network, either due to MPEG virtual buffer constraints or low instantaneous video bitrate, MPEG statmuxes must insert NULL packets to fill the MPEG transmission channel. The NULL packet filled bandwidth is simply wasted.

VBR Network Statmuxes

The characteristics of IP transport make it a perfect match for VBR statistical multiplexing. The essence of an IP-level VBR network statmux, or simply network statmux, is to avoid transrating at the MPEG-level in an attempt to solve the bandwidth oversubscription problem. Instead, queuing and buffering are used at the network edge.

Figure 1 depicts a high level system diagram of network statmuxes.

In this architecture, all IP video sources transmit VBR streams. VOD servers store and stream VOD content in single program transport stream (SPTS) VBR format directly. The coding efficiency improvement of VBR over CBR bodes well for VOD as it brings 40 percent or more storage and streaming capacity savings to VOD servers. Local encoders encode real time video and send out SPTSs with desired VBR mean rates and peak rates. Linear video sources from satellite are converted from multi-program transport stream (MPTS) to SPTS and IP encapsulation is added at the same time.

The IP video streams then travel from the video sources through the network and arrive at the cable network edges where the last mile is the DOCSIS path. The CMTS, be it modular or integrated CMTS, is the starting point of the DOCSIS path. Powered with QoS control and advanced queuing features, the CMTS is an ideal candidate to implement VBR network statmuxing. At the CMTS, if there is no congestion, i.e. when the combined VBR stream bandwidth is less than the bandwidth limitation of the IP/DOCSIS pipe, all video streams pass through packet buffers inside the CMTS with minimum delays. When congestion occurs, the bursty VBR streams will be queued up at the CMTS

temporarily. In the extreme case that the CMTS buffer is full, packets will be dropped by the CMTS, though the probability of such drop to occur is controlled by an admission control function that limits the number of video flows admitted to a single pipe. The CMTS queuing will be further discussed in more details.

The IP video streams eventually terminate at IP STBs. The IP STB can either be a standalone IP STB behind a cable modem, or a hybrid STB with an embedded cable modem. As a result of CMTS queuing, additional network jitter is introduced to video streams. This jitter is absorbed by the IP STB as it dejitters and buffers packets before video is sent to video decoders. In order for VBR network statmuxing to work, there must be a limitation of the jitter introduced by the CMTS queuing so that the end-to-end network jitter is tolerable to the IP STB. In today's well managed service provider network, end-to-end network jitter is a magnitude lower than the dejittering capability of IP STBs. This leaves a big jitter budget for the CMTS to implement VBR statmuxing. For instance, the CMTS introduced jitter can be limited to 60ms. The jitter limitation is enforced on the CMTS by restricting the CMTS queue buffer size. The buffer size is chosen so that packets will stay in the buffer for less than the jitter limit time. The CMTS queue can be drained at the maximum bandwidth of the bonded DOCSIS channel.

The VBR video traffic should be marked with higher priority than data traffic either by differential services control point (DSCP) marking, or video flows should be explicitly identified through flow specs. In a converged IP/DOCSIS pipe, any underutilized bandwidth resources after the VBR video traffic can be consumed by lower priority data traffic.

Besides the number of streams participating in statistical multiplexing, another crucial factor that affects the statistical multiplexing efficiency is the VBR peak (bandwidth) to mean (bandwidth) ratio. The higher the peak to mean ratio, the less efficient the statistical multiplexing is. Early research [2] has shown that the MPEG-2 encoded broadcast quality VBR video has a typical peak to mean ratio from 1.3 to 2.4. IP network statmux design assumes video sources have peak to mean ratios within the limit of 2.4. This peak to mean ratio should be enforced at the video encoder for best video quality. Although service providers can use rate clampers along the video transmission path, the method of rate clamping at the network is inferior to the encoder peak rate enforcement solution as a consequence of additional video processing cost and degraded video quality.

The buffering and queuing scheme brings significant advantages to IP network statmuxes over traditional MPEG statmuxes. First, network statmuxes preserve original video quality keeping the video unchanged at MPEG level. Video is produced by video sources and consumed at STB receivers. No network components within the transmission path re-encode the video content between video sources and STBs. For the same reason, pre-encrypted content can now be easily multiplexed by network statmuxes. Second, network statmuxes introduce delays bounded by the network jitter limit, which make them ideal for low latency video services such as VOD and interactive TV. Lastly, by avoiding the expensive transrating operation and by leveraging the built-in QoS capabilities of IP networks, network statmuxing turns out to be a cost effective approach to VBR video delivery.

As a summary, Table 1 below highlights the key differences between network statmuxes and traditional MPEG statmuxes.

	Network Statmux	Traditional MPEG Statmux
Statmux Efficiency	More efficient w. wideband DOCSIS	Limited by the narrowband MPEG QAM
Bandwidth Overflow	Buffer and delay	Transrate
Video Quality	As good as original stream	Quality degradation from transrating
Latency	Less than 100ms, e.g. 60ms	0.5-1 second
Pre-encryption	Transparent	Have difficulty
Cost	Buffering and QoS already built into IP transport No deep packet processing	Additional system components Deep packet MPEG level transrating
Bandwidth Utilization	100% with converged network	Null packet filling, suboptimal

Table 1. Key Differences - Network and Traditional MPEG Statmuxes




Figure 2 VBR over DOCSIS System Test Diagram

THE EFFICIENCY OF NETWORK STATMUXES

Exactly how efficient are network statmuxes? Perhaps nothing answers the question better than lab results from a proof of concept project. The basic design is to deliver VBR video streams into a controlled IP/DOCSIS channel with embedded QoS features. Streams are added to the channel one by one until the video quality is affected by packet drops. Based on the maximum number of VBR streams supported and the transmission channel bandwidth capacity, VBR statmux efficiency and bandwidth utilization improvement over CBR are calculated. To further study the effect of QoS control on VBR statmux efficiency, network buffers with different sizes are used to smooth out the VBR traffic.

Figure 2 presents the system diagram of VBR over DOCSIS testing. In this experiment, video sources are obtained from satellite feeds. Video streams are converted from MPTS to VBR SPTS and IP encapsulation is added by a Digital Content Manager (DCM). The VBR video streams then pass through a network delay emulator before they reach the DOCSIS CMTS. In the DOCSIS path, a pre-DOCSIS 3.0 modular CMTS based wideband solution is used. In this solution, the CMTS, the EQAM and the cable modem together form the DOCSIS last mile. VBR video streams are terminated at the IP STB. A video quality analyzer is added to the path to monitor video impairments. A PC is used as a receiver of best effort data traffic.

Video Source

Live SD-only satellite feeds are used as the VBR sources. In the video industry, SD CBR MPEG2 video streams are encoded at a nominal rate of 3.75Mbps. The video quality associated with 3.75 Mbps CBR coding is well accepted as the standard for broadcast video. Assuming 40 percent VBR coding efficiency improvement over CBR, the VBR stream with equivalent video quality should have an average bitrate around 2.25Mbps. Two HITS satellite feeds selected for this experiment have just the right video characteristics. HITS1 and HITS9 each comes as an MPTS bundle from the satellite at 27Mbps. Each bundle has 12 video programs. The average bitrate is 2.25Mbps and the peak to mean ratio is 2 to 2.4.

Since these VBR streams originate from MPTS bundles, they are not good sources due to the correlated bitrate peaks and valleys. To create independent VBR streams with uncorrelated bitrate peaks and valleys, the network emulator is used to insert different delays to individual streams. For example, the first stream of the MPTS bundle is delayed 300ms, the second stream is delayed 600ms, the third stream is delayed 900ms and so on until the last stream is reached. The emulated delays combined with the live feeds generate the desired VBR sources.

Video Quality Measure

The IP video delivery quality can be measured with both packet drops and network jitter. As discussed earlier in the paper, the buffer sizes in the CMTS are chosen so that the maximum jitter introduced is bounded. The CMTS-introduced jitter will be removed by IP STBs without affecting video quality. Unfortunately, IP packet drops will cause considerable video quality impairments. To ensure that video quality degradation due to IP transport network is negligible from a subscribers point of view, most carriers allow

the transport network to introduce at most one visible degradation in video quality every two hours. This criteria translates to 1E-6 maximum packet drop probability. Packet drops are detected from QoS counters either in the CMTS or in the video quality analyzer. In the test, VBR streams are inserted to the DOCSIS path one by one while packet drops are monitored. A VBR stream can only be added if this addition will not cause any packet drops for a two hour period. Besides, the probability of packet drop is measured over long term tests and is close to 5E-7, which is better than the well accepted 1E-6 criteria.

DOCSIS 3.0 Channel and QoS Options

Different DOCSIS 3.0 bonding group sizes are used to investigate the statmuxing efficiency with regard to channel bandwidth. To make it simple, the bandwidth increment is 38.8Mbps – the bandwidth equivalent of a single QAM.

The CMTS wideband default class queue is assigned to the VBR streams. Since all VBR video streams are delivered through multicast in the test, the classifier at the CMTS classifies all multicast video traffic to the default class queue. The default class queue size is adjusted to reflect the maximum jitter introduced by the CMTS. Best-effort IP data traffic is mixed with video traffic to drive the IP bandwidth utilization to 100%. While the default class is used to transmit video traffic, the best effort class queue is used to transmit unicast traffic to PC.

Results and Discussions

To facilitate discussions, two quantitative measures are defined. VBR statmux efficiency is defined as

$$\text{Efficiency (\%)} = (\# \text{VBR streams} * \text{VBR average rate}) / \text{channel capacity}$$

where *#VBR streams* is defined as the maximum number of VBR streams that can fit in a transmission channel without causing packet drops for a two hour window. The theoretical limitation of VBR statmux efficiency is 100% if one considers CBR as a special case of VBR with a peak-to-mean ratio of one.

The bandwidth utilization improvement is also derived by comparing VBR statmuxing with a CBR based solution. The bandwidth utilization improvement is defined as

$$\text{Improvement (\%)} = (\#VBR \text{ streams} - \#CBR \text{ streams}) / \#CBR \text{ streams}$$

where the *#CBR streams* is defined as the maximum number of CBR streams with equivalent video quality that can fit in the same channel.

The experiment results are displayed in Figure 3 and Figure 4. Figure 3 highlights the VBR network statmux efficiency vs. DOCSIS channel bandwidth. Figure 4 shows the VBR network statmux bandwidth utilization improvement over CBR vs. DOCSIS channel bandwidth.




Figure 3. VBR Network Statmux Efficiency in DOCSIS




Figure 4. VBR Network Statmux Bandwidth Utilization Improvement

Several conclusions related to network statmux efficiency can be readily drawn from these results. First, the network statmux efficiency improves when the transmission channel bandwidth increases. With 60ms QoS buffers, the statmux efficiency is 81.2 percent for a DOCSIS channel of 38.8Mbps and the efficiency rapidly reaches 91.3 percent when the DOCSIS channel bandwidth is 232.8Mbps with four bonded QAM channels. Quite contrary to our original assumptions, VBR network statmuxes are efficient for SD-only content even when used with unbonded DOCSIS channels.

Next, network statmuxes dramatically improve the bandwidth utilization over the CBR solution. While the unbonded DOCSIS channel delivers 40 percent more streams if VBR and network statmuxing are utilized, the four channel bonded DOCSIS can boast a 57.5 percent enhancement. This improvement is superior to what traditional MPEG statmuxes can achieve in the MPEG transport path. Because network statmuxes only require queuing and buffering instead of the heavy MPEG level processing required by traditional MPEG statmuxes, the same or more bandwidth savings are achieved with only a fraction of the cost.

In addition, as the QoS buffer size increases, the network statmux efficiency further improves. This aspect of the testing was implemented on a simulated Gigabit Ethernet (GE) path with QoS buffer control. To demonstrate the QoS buffer effect, the buffer size is represented in terms of the maximum jitter introduced by the buffer. In Figure 5, when the bandwidth is two QAM equivalent, the statmux efficiency is 75.4 percent with a 4ms buffer and the efficiency is 86.9 percent with a 60ms buffer. However, as the channel bandwidth increases, the buffer size introduced improvement is reduced. For instance, if the channel bandwidth is 38.8 Mbps and the buffer size is increased from 4ms to 60ms, the statmux efficiency improves 11.5 percent. In contrast, when the channel bandwidth is 232.8Mbps, increasing the buffer size from 4ms to 60ms only yields about 4 percent improvement in statmux efficiency.




Figure 5. The Buffer Size Effect

This is not unexpected since the statmux efficiency is already very high when the channel bandwidth reaches a high threshold, thus, room for additional improvement is limited. Note that the statmux efficiency of an IP/DOCSIS path is slightly worse than that of the simulated IP/GE path as a result of additional DOCSIS overhead at layer one and layer two.

Finally, bandwidth utilization should not be confused with VBR statistical multiplexing efficiency. Bandwidth utilization of a transmission channel can reach 100 percent as long as the lower priority traffic can be mixed with video data and no single bit of bandwidth is wasted. There is no doubt that a converged IP pipe holds the promise to fully utilize the transmission channel. To prove the point, the PC behind the cable modem pulls big files from the CMTS through the bonded DOCSIS channel when the VBR video bandwidth utilization is as high as 90 percent. No packet drops are detected during the process due to the QoS features of the CMTS and the priority treatment of VBR video traffic.

RELIABLE VIDEO DELIVERY

Video streams are particularly sensitive to packet drops. Because of the high level of compression used in video delivery, even a single packet drop could result in significant video artifacts. There are three main sources of packet drops in IP networks:

- Because the core of the network is usually rich with bandwidth, packet drops at the core are usually not related to congestion. Instead, load balancing actions, route changes and/or temporary equipment failures could cause packet drops.

- The edge is relatively bandwidth-poor. It is the pipe to the subscriber which is at risk of being congested and as a result, it is where packets are most likely to be dropped.

- Media errors: though technically packets might be dropped on the Ethernet part of the network, most drops occur on the HFC network itself due to RF issues.

To minimize and possibly eliminate the video artifacts caused by packet drops, we

propose a three-tier approach which addresses all major sources of packet drops:

- Admission control: the role of admission control is to make sure that the network, or in the context of this paper, the CMTS specifically, can deliver content reliably.

- Scheduling: while admission control makes sure that we deliver content reliability, it is the scheduler that does the actual work of delivering the content in real time in a reliable way.

- Error repair: Error repair was designed to help in cases where packets are dropped because of media errors and/or network flaps, however, it could be used to help in cases where both admission control and scheduling could still not guarantee packet delivery.

ADMISSION CONTROL AND SCHEDULING

With CBR services, admission control is trivial. If a CBR flow requires X mbps, and the bandwidth of the channel is $10X$ Mbps, then 10 flows can be admitted. The CMTS would track the number of flows that the cable segment has to carry, and reject any request to activate an 11th flow.

When it comes to VBR the picture is more complex. As explained previously in this paper, VBRs can be oversubscribed because it is “reasonable” to assume that not all flows send at peak traffic rate all at the same time. But what does “reasonable” mean? There is a probability that enough traffic peaks occur at the same time and in such an event the channel will not have enough capacity to carry all the traffic. In such an event the CMTS scheduling queuing and admission control disciplines will help to minimize (or eliminate) packet drops in the event of traffic congestion.

To illustrate how a scheduler works, we can start with a simple example: assume a CMTS with a backhaul link of 1Gbps, and two video flows, each one with a traffic peak of 6mbps. Furthermore, we can assume that admission control limits the CMTS to accept only two flows for this case (though more could possibly have been admitted). Both flows drain onto a single output that is capable of 10mbps. We assume that the targets of these flows are two cable modems as depicted in Figure 6.




Figure 6. Queuing example

Since the video flows are 6Mbps each there is no congestion risk on the backhaul. However, when they get to the cable interface the worst case aggregate rate they can reach is 12Mbps while the cable interface in this example can support only 10Mbps. The tools the CMTS can use are queuing and scheduling:

- queuing will buffer up the packets in a “queue” until the 10 Mbps channel is available again to send them.
- scheduling will decide which queue to service and in what order

The CMTS can use the DOCSIS tools to define the queuing/scheduling structure needed to deliver these flows reliably:

- A classifier that will uniquely identify the video flow. For example, the combination of a destination IP address of

the client device, and a destination user datagram protocol (UDP) port are a good way to identify a packet stream that belongs to a single video flow. This approach can also be used for multicast flows. The detailed discussion of how this classifier is created is outside the scope of this paper.

- A service level agreement that defines how to queue the flow. For example, in our case it's a flow that has a 6 Mbps peak rate.

The CMTS manages queuing by dedicating a queue to each one of the video flows and by controlling queue scheduling.

Naturally all the queuing/scheduling can do is to mitigate the cases where the aggregate traffic bursts are above 10Mbps. If the bursts are too long then packets will experience an unacceptable delay in the CMTS queues (and eventually will get dropped as the CMTS queues have limited size). By putting a limit on the number of flow admitted by admission control, the MSO can control the tradeoff of how many flows can be admitted to a channel vs. what the packet drop probability would be.

An additional tool that can reduce the risk of packet drops (at the expense of having less video flows committed) is the use of "guaranteed minimum rate". This parameter in DOCSIS defines the rate that a scheduler MUST deliver even in a case of congestion. In a way, one can view CBR as a case where the the peak rate equals the committed rate. Based on this, the smaller the difference between the committed rate and the peak rate, the smaller the risk of packet drops.

Another tool in the DOCSIS toolkit is "priority". This parameter is critical in an environment where we have mixed video with other services such as data. The priority parameter, along with the "guaranteed minimum rate" parameter, gives an assurance that even if data services in a given channel are congested to the point where packets are dropped, there is still enough bandwidth dedicated to high priority video flows.

However, even with proper admission control and scheduling, packet drops could still occur. Facilitated with extremely low packet drop rate, error retransmission and IP-based packet-level forward error correction (FEC) are promising cost-effective solutions for the packet drop problem. Both technologies are well understood and have been applied to IP video applications to protect video streams from common impairments of IP networks, such as packet loss.

MPEG over UDP/IP is widely used to transmit real time video traffic through IP networks. With UDP transmission, packet drops are not reliably detected due to the lack of a sequence number at the UDP layer and due to the limited capabilities of the MPEG transport stream level continuity counter. By adding a real time protocol RTP above UDP, packet drops are easily identified through a 16-bit sequence number at the RTP layer. Error retransmission and FEC leverage the RTP encapsulation of the video stream and repair dropped packets at the network edge. An example architecture for error retransmission and FEC and is shown in Figure 7.




Figure 7. Error Repair Architecture

In the above error repair architecture (Figure 7), an error repair client is located at an IP STB. When FEC is applied, the video source sends out an FEC stream along with the video stream. In this scheme, periodically selected media packets are used to generate FEC packets. The error repair client is responsible for detecting packet loss and recovering the lost packets utilizing the additional FEC stream. When error retransmission is utilized, a repair server at the network edge caches video content. The error repair client utilizes standard based RTP/RTCP toolkit defined by IETF to request retransmission of lost packets. The same toolkit can also be used to accelerate channel changes in IPTV.

CONCLUSION

The ever increasing demand for bandwidth requires efficient HFC bandwidth utilization. DOCSIS 3.0 and IP video are shifting the VBR video delivery to a new paradigm. IP level VBR network statmuxes overcome the shortcomings of traditional MPEG statmuxes and provide the least expensive, low latency and best video quality approach to reap the benefits of VBR video.

The proof of concept work introduced in this paper not only proves the feasibility of this VBR network statmuxing in today's DOCSIS 3.0 networks, but also demonstrates

the tremendous value and potentials of wideband DOCSIS in video delivery. DOCSIS CMTSs, with their built-in advanced QoS capabilities play an important role in VBR network statmuxing.

The trend in tomorrow's video delivery is more HD content and more advance coded video content. VBR network statmuxes respond to this trend by leveraging the channel bonding capability of DOCSIS 3.0 and generate unprecedented multiplexing efficiency as the wide channel is promising to get wider. By avoiding deep packet processing, VBR network statmuxes scale easily to future video coding technologies.

REFERENCES

1. Si Jun Huang, "Principle, Applications of Variable Bit Rate Coding for Digital Video Broadcasting, w. Statistical Multiplexing Extension", NAB 1999
2. Daniel P. Heyman and T.V. Lakshman, "Source Models for VBR Broadcast Video Traffic", IEEE Transactions on Networking, Feb. 1996