THE COMPLETE
TECHNICAL PAPER PROCEEDINGS
FROM:

A QUANTITATIVE APPROACH TO MULTI-TIERED STORAGE AND STREAMING FOR VOD SERVERS

Robert Duzett
C-COR Inc.

Abstract

The uneven popularity of titles within a VOD content library, as expressed by its content demand profile, opens the door to improved efficiencies from multi-tiered storage architectures. Quantitative methods are shown for addressing the design and provisioning of such architectures. A single constant 'r' can describe the content demand profile. Analysis using 'r', along with storage media pricing ratios 'p' and system streaming capacity 'S' results in a model for optimizing and balancing the mix of storage across multiple storage tiers. This analysis is applied to both 2-tier and 3-tier architectures, and is extended to consider the effects of future trends in content profiles, streaming demand, and storage media.

VOD Demand Profile

Data gathered from VOD deployments show that there is a steep curve representing the distribution of content to streaming demand ("hot" content feeds a lot more of the streams). One very useful way to represent this Content Demand Profile is to plot a curve mapping cumulative library content hours (sorted in order of popularity) on the X-axis to cumulative fraction of total streams on the Y-axis. Thus, for a selected number of the most popular content hours you can look up what fraction of the streams are driven by that content (see Figure 1).

We have observed repeatedly from actual VOD profile data that the curve representing this relationship can be mapped closely to an exponential formula: 1-e^(-rc), where 'c' represents the number of content hours along the sorted content list and r is a ratio characterizing the steepness of the curve. In this way, a single constant, 'r', can represent very closely the content demand profile for a given time window of VOD usage for a given content library ('r' characterizes the library and the demand it generates). A higher value of 'r' means a steeper curve. As will be shown hereafter, we can use 'r' to help calculate other useful relationships based on the demand profile. The graph of Figure 1 shows a demand profile from an actual deployment as well as a curve fitted closely to it using the formula 1-e^(-rc). For this data, r = 0.007.

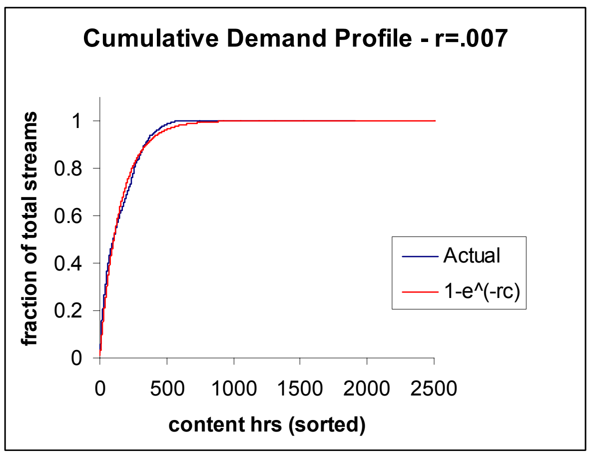

Figure 1

Tiered Caching Principles

The nonlinear nature of the VOD stream-from-content demand profile, much like other resource demand profiles in general computing, suggests the opportunity for improved efficiencies via caching. Placing a small amount of "hot" content into a small cache of expensive but fast memory could potentially make the whole system more efficient.

Today's leading edge VOD servers, for example the C-COR n5 server, make use of this caching opportunity to create two or more "tiers" of storage based on different technologies or different performance/density points – e.g. RAM, Flash, fast disk, slow disk, etc. The basic assumption is that for each higher storage tier, storage gets cheaper while bandwidth gets more expensive. Thus the lowest tier is likely the fastest but the most-expensive per byte (& the least dense).

At each storage tier, we are trading off additional STORAGE (cached, copied) at this tier to replace STREAMING from the next higher tier.

The ideal model for making this tradeoff allows all tiers unrestricted scaling for streaming or storage. In practice, however, there are physical and architectural limitations. Also, the requirements of a VOD deployment, as mapped to a given storage tier or technology, will be unlikely to yield a perfect balance of Bandwidth and Storage. The storage will be either "content-limited" or "streaming-limited".

"Content-limited" means that the bandwidth available from the required content exceeds the bandwidth required for streaming – in other words, storage is being added for content, not streaming. "Streaming-limited" means that the storage capacity from the required streaming storage exceeds the storage required for the current content library – in other words, storage is being added for streaming, not content.

This tension between streaming and content requirements can lead to inefficiencies. For example, consider a content-limited situation in which a centrally-located storage system could provide all necessary streaming bandwidth but for limited transport bandwidth to the edge. In this case, the content must be pushed out to the edge and duplicated at various headends. On the other hand, consider a streaming-limited case in which every headend has more than sufficient content storage because of streaming bandwidth requirements placed on the storage. In this case, an excessive amount of storage is paid for but a portion goes unused.

Tiered storage can ameliorate these kinds of imbalances and make overall operations more efficient.

A 2-tier Caching Model: Disk vs. DRAM

The Hard disk drive is a commodity high-DENSITY storage. DRAM is a commodity high-BANDWIDTH storage. The ideal storage would have disk density and DRAM bandwidth. Based on today's pricing:

Density-per-$ ratio of disk:DRAM = 60:1
Bandwidth-per-$ ratio of disk:DRAM = 1:30

We can somewhat balance these two ratios by storing the content library on hard disks, while caching hot content in DRAM.

Until recently, RAM caching was not economical ($cost/density was too high). However, because the RAM density growth trend (~40% per year) is much steeper than the disk performance growth trend (12-15% per year), RAM caching will prove more and more cost-effective as time goes on.

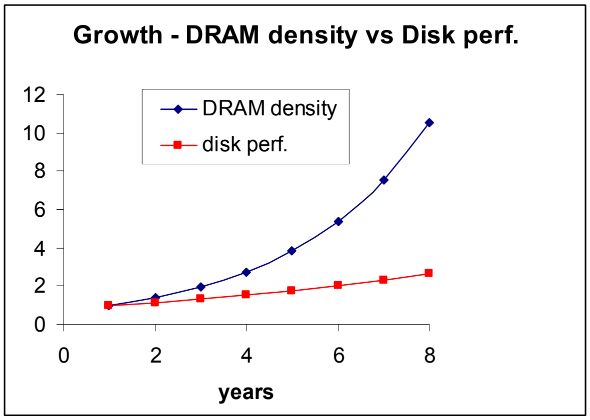

Figure 2

If we're stream-limited, RAM cache provides an opportunity to remove disks or increase server performance (to the limit of the platform). If we're content-limited or transport-limited, RAM cache provides an opportunity to radically reduce content duplication at the edge servers while centralizing the complete content library on high-density disks.

How Much Cache is Cost-Effective?

To provision the storage subsystems of a large VOD server for optimum cost-performance, one can expect the hottest titles will be stored in, and streamed from, cache while the rest of the content will be streamed from higher storage tiers. But, what is the optimum balance of cache content and higher-tier streams? That is, what is the optimal cache size?

If one were to attempt to get ALL streaming from RAM, then ALL the content would have to be stored in RAM, which would obviously be too large and expensive for even moderate content libraries. So, one must attempt to achieve a reasonable portion of streaming, as cost-effectively as possible, from RAM.

The demand profile curve indicates how much CONTENT must be cached to achieve a given HIT RATIO. As content is added incrementally to the cache, the hit ratio rises and incrementally more streams can be fed from the cache. So, the hit ratio, and therefore the CONTENT SIZE, of the cache, NOT the BANDWIDTH CAPACITY of the cache, determines how many streams it can feed (assuming of course sufficient bandwidth capacity from the cache).

We can create a cumulative hours vs cumulative streams graph by multiplying the demand profile curve (cumulative hours vs cumulative *fraction* of streams) by the total number of streams for the system (for example, 8000 streams). See Figure 3. This graph shows how many streams will be sourced by the cache for any given size of cache. In effect, for a given cache size, the streams underneath the curve come from the cache while the streams above the curve come from disk.

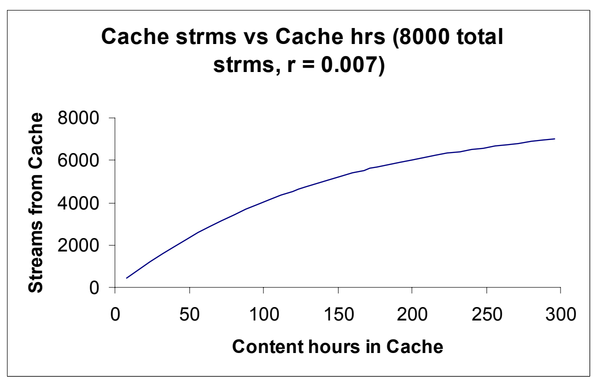

Figure 3

Considering this graph, adding cache is cost-effective as long as the incremental cost for cache content is LESS than the corresponding decremental cost from streams displaced from the next higher storage tier. Therefore, adding cache is cost-effective as long as the slope of the curve is GREATER than the pricing ratio (p):

p = \$-per-hour for cache / \$-per-stream for higher-tier

So, cache is cost-effective while the slope of the streams vs hours curve is greater than p. The slope is the derivative of the curve, so we have:

Sre^(-rc) > p (S=total streams, c=cache content hrs)

Solving for c, we can determine the maximum cache size that is cost-effective for any system of ‘S’ streams and a content library with demand profile ‘r’, given a pricing ratio ‘p’ between two storage tiers.

Cache_max = ln(rS/p)/r (in hours)
Hit-ratio_max = 1-p/(rS)

For example, graphing hit-ratio_max vs total streams for p=47 and r=0.007 shows that caching is cost-effective for VOD servers larger than about 7000 streams; and a cost-effective hit-ratio of about 50% is reached with a server size of about 13000 streams. That 50% hit ratio corresponds to about 100 hours of content in cache (160 GB). See figures 4 and 5.

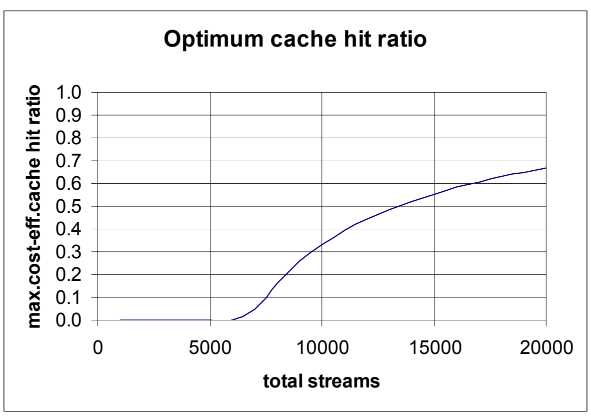

Figure 4

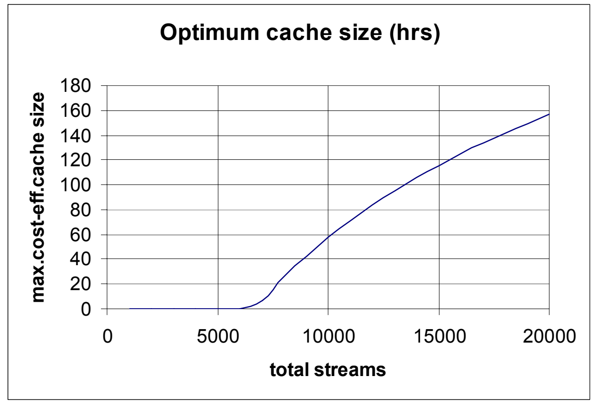

Figure 5

So, we have determined four major factors that determine the appropriateness and size of a given caching storage tier:

1) the demand PROFILE of the content library, represented by ‘r’.

2) the PRICING ratio ‘p’.

3) the total #STREAMS of the system

4) the total CONTENT library size (this dictates the size of the highest storage tier and thus the minimum bandwidth that may be streamed from it, ie whether we are content-limited)

We have also determined that RAM is not a cost-effective way of achieving streaming bandwidth for systems smaller than 7000 streams, because disk bandwidth is cheaper for those systems; and that even above 7000 streams cost-effectiveness places limits on the amount of RAM that is desirable to displace disks for streaming bandwidth.

For this reason, VOD servers should be designed such that the disk vs RAM tradeoff can be made in a balanced and flexible way based on the size of system to be deployed. The server architecture should not place unreasonable restrictions on storage tier provisioning. For example, this is why the C-COR n5 VOD server was specifically designed to make both disk and RAM independently scalable, and thus balance a wide range of potential needs from both disk and RAM.

A 3-tier Caching Model

This 2-tier storage model (tier-1 is the cache, tier-2 is the disk array) can be extended to a 3-tier storage architecture. For analysis purposes we will consider here a 3-tiered global architecture in which all the storage on all tiers is globally accessible by all streams of the system. Tier 1 is the fastest storage; tier 3 is the densest storage. Tiers 1 and 2 cache content from the global library at tier3.

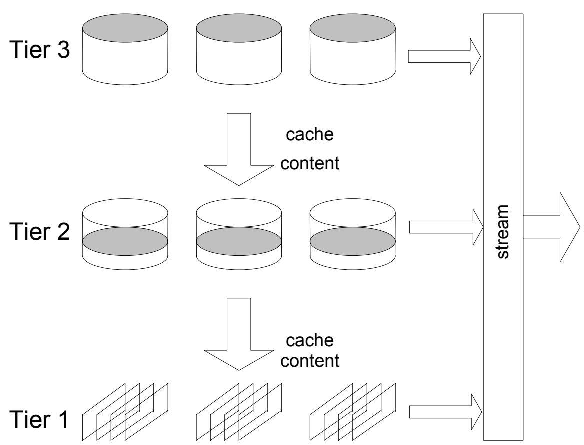

Figure 6

We consider all practical/reasonable storage technologies – e.g. RAM, SCSI, SATA – and various devices from each. We consider costs for each device ($-per-GB, $-per-Mbps) and pricing ratios (p) between tier candidates, and then choose 3 reasonably-priced devices that reflect increasing streaming costs and decreasing storage costs. See Table 1.

Tier#:	**1**	**2**	**3**
	DRAM	**SCSI**	**SATA**
	DDR266	15K73	7.2K320
p=$/hr_this / $/strm_nxt	48.5	1.8	
$/hr_this / $/hr_nxt	22.3	8.8	
$/strm_nxt / $/strm_this	18.2	1.2	
Raw unit capacity (GB)	2.1	73.0	320.0
net hrs/unit	1.3	29.2	128.1
net strms/unit	1120	64	27

Table 1

For a given content library size and profile (‘r’), we analyze the cost-effective tier boundaries for various stream counts, keeping in consideration both content-limited and streaming-limited effects.

This analysis allows us to determine the optimum storage balance and hit ratios for the 3 storage tiers, for systems of any total stream count. We can also determine the optimum storage cost, and compare this with various 2-tier and 1-tier storage technologies.

Figure 7 shows the optimum storage mix for a 3-tiered 5000-hour system, across a wide range of system sizes (characterized by maximum stream counts, on the x axis), using a content demand profile of r=.007. Note that tier-3 is used to archive the entire library and is designed to be content-limited, while tiers 2 and 1 are used to provide the necessary streaming bandwidth. Tier-2 storage is the most cost-effective streaming storage up to 7000 streams and ramps up to that point, beyond which tier-1 takes over. Figure 8 shows the corresponding hit ratios for the three tiers.

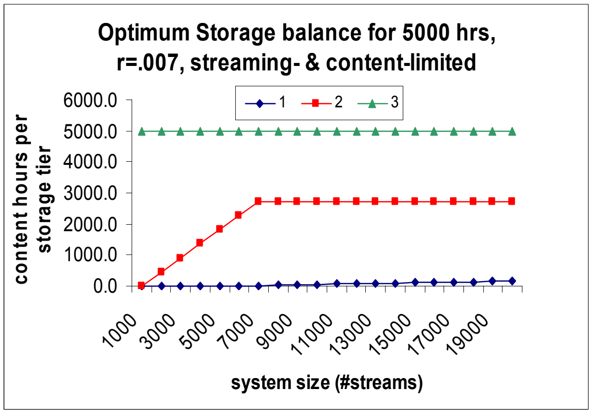

Figure 7

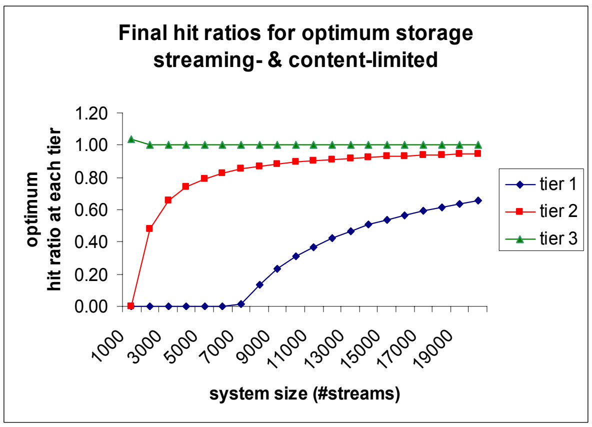

Figure 8

Figure 9 shows the total per-stream storage cost for this 3-tier, 5000 hour architecture. For comparison, it also shows the 1-tier equivalent, which applies the most cost-effective disk drive to meet content and streaming requirements for each stream-count point on the x-axis. Note the cost-savings achieved by architecting a 3-tier storage hierarchy. This difference widens further when you consider the real-world physical constraints generally faced by a single-tier storage system. While this analysis assumed unlimited scalability, the number of drives that can realistically be supported often forces the use of faster, more-expensive drives to meet bandwidth requirements within the allotted density.

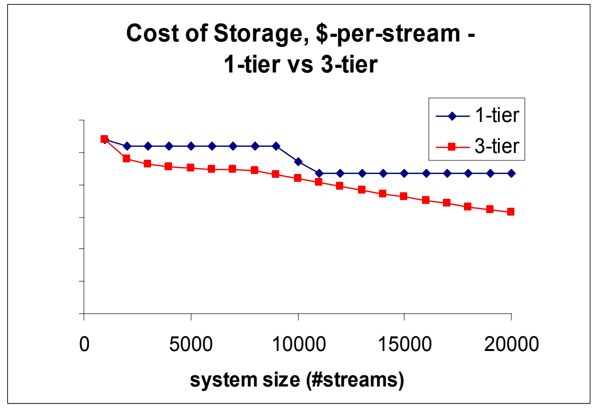

Figure 9

Other Tiered Architectures

Tiered architectures other than the global hierarchical architecture can also be analayzed. Among these are:

- central global library (t3) with distributed isolated local servers (t2 & t1 caches) (=edge servers);
- central global library (t3) with distributed switched local servers (t2 & t1 cache content distributed among servers);
- central global library (t3) + site global cache (t2) + distributed server local cache (t1); and
- content distributed optimally among 3 tiers – no caching, just shuttling among tiers (monolithic or switched servers).

as well as others. All of these are derivatives of the central tier-3 model.

Practical Considerations & Further Quantitative Analysis

All storage tiers, devices, and technologies have scaling limitations and overheads. These include mechanical and packaging limits, controller design tolerances, interconnect bandwidth &

latency limits, transport capacities, etc. In addition, the costs, bandwidths, and capacities of server platforms, storage systems, and other infrastructure can have a significant effect on the final cost-effectiveness of any tiered architecture, beyond the storage devices themselves. Many of these limitations, overheads, costs, and capacities can be built into a multi-tiered model such that their effects can be felt and accounted for in the architectural analysis.

For example, there exist server architectures today that narrowly limit the interconnect bandwidth coming from the disk array while maintaining highly-scalable bandwidth from RAM. This unfortunate bottleneck restricting disk-sourced streaming creates a severe imbalance in the architecture and a consequent cost premium, as shown in figure 10. Note that the optimum balance of disk and cache is broken by an architectural limit of 1000 streams from disk, which causes the storage costs of small systems with this disk bottleneck to be more than double those of systems that are well-balanced. A significant premium is paid even for large configurations.

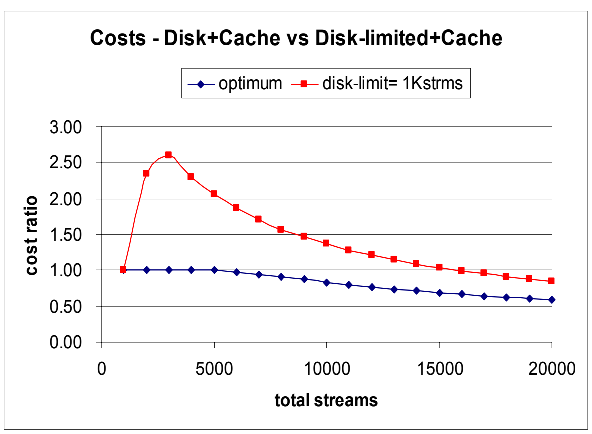

Figure 10

A successful tiered-storage server architecture will maximize the scalability and flexibility of each storage tier, and the storage system as a whole, within reasonably anticipated ranges, so that a wide variety of VOD deployments can be configured as close as possible to the optimum balance of tiered resources using the most appropriate storage media and technologies.

Other concerns, not directly affecting capital economics and difficult or impossible to include in a mathematical model, could alter architectural decisions. These include reliability, operational costs & considerations, interoperability, legacy, etc. and must be duly considered in all architectural design and development.

Tiered Caching Trends over the Next 3-5 yrs – Technology & Economics

It has been shown above that caching effectiveness hinges on content library characteristics (size and profile ‘r’); the price ratios (‘p’) between various storage types; and system stream counts (‘S’). It is very interesting to consider trends and expectations for these parameters over the next few years to see where caching and tiered storage may take us in the future. By extrapolating historical and predictive numbers for such things as device costs, bandwidths, and capacities; library sizes and content mixes; bit rates, take rates, penetration ratios, and HD ratios; effects of Moore’s law on platform capacities; etc. and then applying a practical best-case and worst-case range to each of these, one can build a model that looks at caching effectiveness and/or storage costs over several years as well as its sensitivity to particular parameters or groups of parameters. And this can be done for various tiered or non-tiered architectures. Below are graphs for a 2-tiered, disk+RAM, caching model (figures 11 and 12):

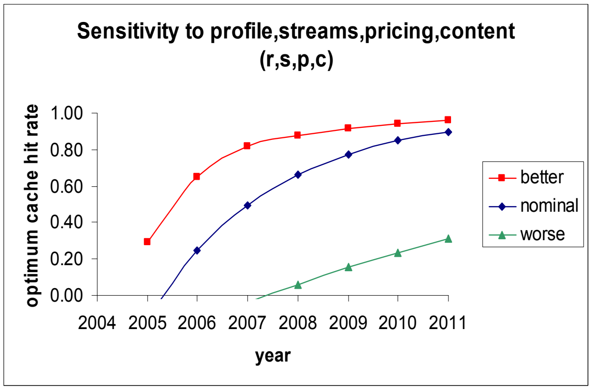

Figure 11

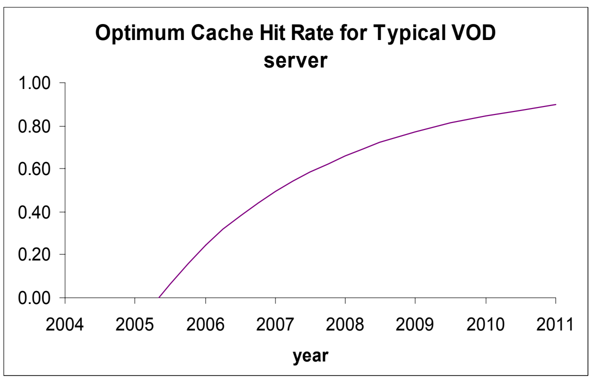

Figure 12

It is also important to consider the effect of new technologies on future tiered storage architectures. For example, NAND Flash has become denser and cheaper than RAM, as well as, in the right format, faster than disk. It therefore has the potential to become a cost-effective storage tier for video systems.

<u>Measurement and Predictability in a Hierarchical Age</u>

The addition of new storage tiers to the VOD architecture creates new complexities for the designer as well as the intergrator and system manager. The storage and streaming requirements of a VOD system or deployment now invoke multi-dimensional parameters. Storage is no longer determined by simple questions of "how much?" and "how fast?", but also by "which tier?" and "what hit-ratio?", etc.

This paper has offered some basic tools for identifying and talking about the key parameters that characterize a multi-tiered architecture. An integrator can anticipate and design for a required range of 'r', 'p', 'S', and 'C' values for a specific deployment or for a general architecture over many deployments. Architects and managers can characterize content demand profiles with a simple 'r' number so they can be discussed quantitatively.

As libraries grow and become more diverse, 'r' values will undoubtedly fall, though not linearly. At any given time, however, one may discuss how 'r' values change over the years, weeks, or months; for time of day; and across content management policies and marketing approaches. Variations in 'r' can also characterize the effect of the response times of cache algorithms being implemented or studied.

In other words, use 'r' as a measure not just of static content libraries but of server and cache efficiencies, and content management efficiencies and marketing efficiencies. Which titles are marketed and how they are marketed can make a big impact on the content demand profile and thus on storage and transport and streaming efficiencies. Measuring hit rates at various tier boundaries of a specific system tells you things about that system only, but then translating those hit rates into an 'r' value now describes the overall profile of access demand for the content library from the attached subscribers.

An integrator or system architect will specify and test a VOD deployment against an expected range of content profiles, as characterized by ‘r’ values; and against a range of expected pricing trends for various storage media, as measured by ‘p’ values; and against a range of system sizes, as specified by ‘S’ stream counts.

CONCLUSIONS

- A Content Demand Profile can be fairly characterized and quantified with a single number, which can then be used to drive a tiered caching model

- A tiered caching model can accurately model effects of both content- and streaming-limited cases.
- A caching model can be used to find an optimum cost-effective balance of 2 or 3 storage tiers.

- Cache effectiveness is determined by 1) the content profile, 2) total streams, 3) relative price-performance of storage devices, and 4) content library size.

- DRAM will be an increasingly cost-effective caching technology for VOD; current economics support it for medium-to-large systems.

- A successful tiered storage architecture will strive for balance, flexibility, and scalability across all tiers, so that the resulting VOD system can be cost-effectively and efficiently applied to a wide range of deployment opportunities.

Contact information:
Robert Duzett
rduzett@c-cor.com
503-690-6305

ADAPTING ADVANCED VIDEO CODING TO CABLE NETWORKS

Mukta Kar, Ph.D. and Yasser Syed, Ph.D.
Cable Television Laboratories, Inc.

Abstract

This paper describes the current cable environment and what demands it has on development of a cable-friendly AVC video coding stream. It considers what the further restrictions on AVC are for recreating the same services already provided by MPEG-2 plus coexistence in the same network with MPEG-2 video streams and other digital services. The paper also looks at what are the requirements for multiplexing, random access to support existing broadcast services, and how this affects the coding design. Furthermore it looks at additional emerging needs to address services like Personal Video Recorders (PVRs), Video-on-Demand (VOD), and Digital Program Insertion (DPI) and how this can constrain the design on the video stream and the video coding layer. This paper concludes by saying it is not enough to just have a highly efficient video coding standard developed, the cable environment (or any application environment) imposes its own set of demands that trades off the advanced video coding design and some compression efficiency for expected viewer experience.

INTRODUCTION

In the last century television contents (video and audio) have been delivered to consumers mostly in analog format. It was realized then that the fidelity of analog reception is good in noiseless or less noisy environments, and also the analog system is either inefficient or inflexible in the important areas such as service security, bandwidth usage, interactivity, etc. Around the 80's and 90's of the last century, researchers understood the power and flexibility of migrating to digital technology, and it started being practical to do so.

However, the delivery of uncompressed digital audio/video to consumer homes is not an acceptable solution as it takes a huge amount of bandwidth to deliver. To tackle this problem, researchers had been working on how to compress digital audio/video effectively for consumer-oriented applications. Towards the end of the last century, audio-video compression technology matured to the level that it can be practically used for consumer applications, and MPEG-2 standards were created in 1994 to broadcast digital audio/video to consumer homes and offices. MPEG-2 technology takes much less bandwidth than analog delivery (today approximately 10:1 bandwidth efficiency) while providing same or better fidelity (quality of audio/video) even in a noisy environment.

AVC: EMERGING DEMANDS FOR VIDEO CODING

Over the last few years MPEG-2 technology has become very popular and received global acceptance which has led to many new innovations such as HDTV, PVR (DVR), VOD, streaming media, interactive TV, etc. These applications were unthinkable in the analog format and did not exist as MPEG-2 standards were created. While these new applications can help generate more revenues for the broadcast industries (cable, satellite and terrestrial), they also require more digital bandwidth to deliver. For example, an HDTV channel requires a bandwidth 15-20 Mbits/sec whereas a SDTV channel takes only 2-3 Mbits/sec using the MPEG-2 video standard [1,2].

Besides broadcast industries, other video industries such as DVD, streaming media, etc. also needed an advanced video codec that has significantly better compression than MPEG-2. For example, an HD movie cannot be stored on a current DVD disc. As MPEG-2 video needs more bandwidth, today's streaming media technologies use proprietary codecs to deliver their content but encounter storage and software challenges due to a lack of convergence on a compression standard. In another industry, conversational video applications (video conferencing and video telephony) use another video codec standard, but are also facing similar challenges to streaming media. This shows that each of the major video industries today use diverse video standards and different hardware platforms for delivery and storage which impede consumer acceptance of multiple video applications.

One of the answers to the problems above is to compress digital audio/video even more efficiently than that in MPEG-2 in order to be useable across different industries. A better answer will be to create a new video standard that will not only provide significantly better compression than MPEG-2 but also a convergence in hardware platforms across major video applications.

To have better compression efficiency as well as to eliminate divergent requirements in hardware, the video experts in ITU and ISO/MPEG formed a joint video team (JVT) to investigate a solution. Their collaborative work over two years resulted in the creation of an advanced video coding (AVC) in ISO/MPEG (MPEG-4 part 10) and H.264 in ITU in May, 2003 [4]. AVC may be considered as an aggregation of tools, some of which are enhanced MPEG-2 tools and a few additional new tools that have been added to increase compression efficiency while providing the same or better picture quality. In brief, AVC achieved two important objectives - better compression efficiency than MPEG-2 and convergence in hardware across major consumer video applications - but sacrificed backward compatibility with the MPEG-2 video compression standard. However, backward compatibility has been addressed by the silicon designers at the AVC codec chip level such that it becomes transparent to users.

AVC tools consist of enhanced MPEG-2 tools, some new tools, and error resiliency tools. MPEG-AVC or AVC is a large tool box which can meet the needs of most major video applications such as broadcast video, conversational applications, and streaming video. The AVC system has created an opportunity to change the paradigm of coding in AVC. The new tool box allows multiple reference frames to be used for predictions as opposed to the two reference frames allowed in MPEG-2. It also introduces spatial prediction in intra-coded frame and adaptive motion compensation. Motion compensation resolutions' block size may be 16x16, 8x16, 16x8, 8x8, 8x4, 4x8 or 4x4 pixels. It supports block transforms of 8x8 and 4x4 pixels for luma which may be used adaptively. It also introduces in-loop filtering to smooth the sharp edges which are an artifact of using block transforms. AVC also added a few error-resilient tools (not present in MPEG-2) to support video delivery over non-QOS networks such as public IP networks and wireless networks.

CURRENT STATE OF CABLE: THE NEED FOR AVC IN CABLE

About ten years ago, the cable industry was delivering premium TV programs primarily in analog only format. Over the last ten years evolutionary changes have been taking place in cable. Today cable delivers not only analog video but also a few digital enabled services such as digital TV, high speed data, and telephony. Cable is also in the process of using non-linear video services in addition to scheduled types of services. At present cable delivers linear applications like broadcast television but also is adding non-linear services such as VOD, PVR, VOIP. Also we are shifting from the broadcast television model to a personal television one. As a result cable faces a real bandwidth crunch as it has to support the legacy analog and digital broadcast television services due to regulatory and logistical reasons, while adding an array of the new digital services to actively engage with new competition. VOD, MOD, and PVR services require a significant amount of storage at the headend and in client devices. So we need better compression of the digital content than MPEG-2 can provide. AVC will certainly help in delivering more channels over the same digital bandwidth as well as in storing the digital content at the headend or in PVRs.

It was mentioned earlier that MPEG-2 technology has been accepted widely across the global broadcast industry. It provides tremendous business continuity in terms of content, equipment (transmission, storage, and test & measurement), etc. It is expected that the general video industry will migrate to AVC in a next few years, especially in the area of HDTV delivery and storage. AVC may also provide opportunity for adding new services, such as video conferencing and video telephony, using the same AVC hardware platform. In a way, to future-proof our industry with respect to delivery of multiple applications, the cable industry needs to consider all these factors mentioned above.

ADAPTING THE AVC STANDARD TO CABLE

Although AVC created profiles and levels geared towards major video applications, still it is not well tailored towards the cable environment. The video stream may have a much better compression efficiency using all tools in the standard, but the expected viewer experience for cable TV has not been designed in the AVC standard. The adaptation of MPEG-2 to cable was an easier process than what is now being considered because the standard was optimized more towards a broadcast application, while the AVC standard tries to encompass more than that. Additionally cable has also been changing at the same time to address other applications than just broadcast.

Constraints need to be written on the AVC standard (or any future advanced coding standard) to maintain current expected services and viewer experience while improving the plant bandwidth usage efficiencies and expanding on these services. This constraint document will be similar to the standard document SCTE 43 2005 [6] that is written for MPEG-2 video. The network infrastructure and legacy equipment need to continue operating without disruption. Existing services such as the Electronic Program Guide, Emergency Alert Messages, and audio-video synchronization need to be maintained and be agnostic to any new video coding standard. Viewer experiences like channel switching/surfing, PVR recording and interactive video services still need to be maintained transparent to different video coding technologies. At the same time, the splicing of video streams to perform local commercial insertion should be visibly seamless as done in MPEG-2. The cable viewer should not be aware of any changes in video coding technologies, but should benefit from the increased services. Designing constraints may have a minor negative effect on compression efficiency, but will allow the viewer experience to remain intact.

DELIVERY OF AVC CODED VIDEO

It is important to note that currently all digital services are delivered over MPEG-2 transport over the HFC cable network. To avoid expensive changes to the existing MPEG-2 transport-based infrastructure, any new digital services that will be added to cable should also be delivered using existing MPEG-2 multiplexing capability with other services without causing any side effects. This means that the existing services based on analog, MPEG-2 video, and high speed data services will remain unaffected. So it is probably the most important requirement that any new service based on AVC coded video (or any advanced coded video that may be added to cable) should be deliverable over the MPEG-2 transport infrastructure that exists in the cable network today.

It may be noted that AVC is just an advanced video coding standard. Unlike MPEG-2, no specific transport standard has been created for delivery of AVC coded video. The MPEG committee realized that a matured MPEG-2 infrastructure should not be reinvented. Keeping that in mind, MPEG created an amendment to MPEG-2 Systems [3] to deliver AVC-coded video over existing MPEG-2 transport. Figure 1 shows how AVC coded video can be packetized in MPEG-2 transport packets which will enable multiplexing with other elementary streams and other cable digital services. Transport over MPEG-2 will ensure audio-video synchronization (no lip-synch error). Additionally, it will support existing cable applications such as channel change, trick-mode, and DPI. Some constraints will be necessary at the video level (SCTE 43 [6] for MPEG-2 video). The digital transport / multiplex standard SCTE 54 2004 [7] has to be amended. These are discussed below.

It is mentioned earlier that the AVC standard will not only benefit broadcast applications, it will also be used in other major video services such as conversational and streaming media applications. ITU H.320 and Internet Protocol (IP) are used today to deliver conversational and streaming media applications, and these transports can be utilized for AVC as well. Referring back to Figure 1, AVC coded video allows for supporting these 3 major transport protocols (MPEG-2, H.320 & IP), using a network abstraction layer (NAL). First, the coded video and related information (PS and SEI) are packetized in NAL packets which can be variable in size. Each NAL packet will have a type value that tells the type of payload it is carrying (VCL, PS or SEI). Those NAL packets are then re-packetized in MPEG-2, IP and H.320.

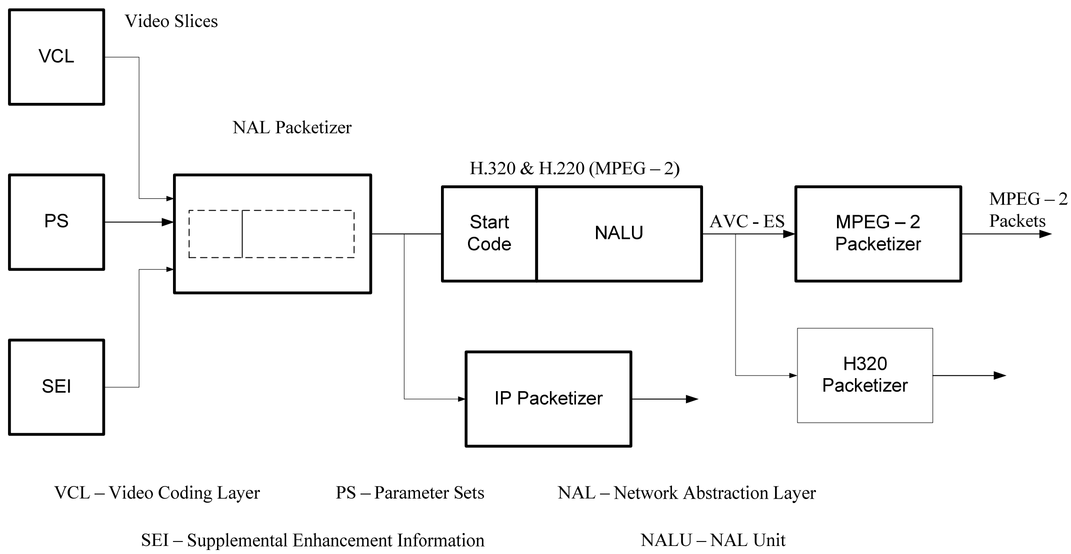

Figure 1. Carriage of AVC Compressed Video

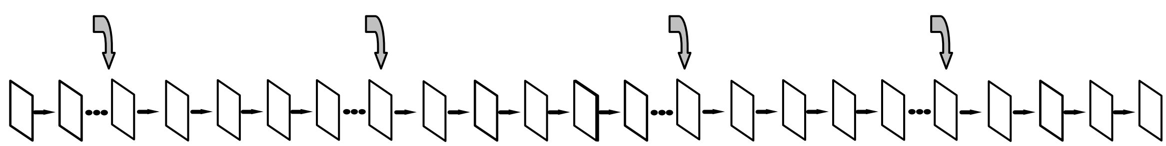

Figure 2. Random Access Points into Video Stream

RANDOM ACCESS REQUIREMENTS FOR CABLE

A critical feature needed for cable is the ability to randomly access the video stream at regular intervals [see Figure 2]. This is an essential need for fast consistent channel changing/ channel acquisition in the broadcast environment, but it also can be quite useful for DVR and commercial (ad) insertion applications. As will be pointed out in the following sections, constraints to design the video stream and transport headers will affect how a video compression standard is used and will tradeoff some potentially very high coding efficiency gains in exchange for supporting a viewer expected experience.

Channel Change requires random access points to allow switching into an existing broadcast channel without an undue amount of delay in receiving the video picture. The viewer after he switches the channel will normally tolerate 1-2 seconds at most of not seeing video (because that is what a viewer already expects from existing analog change times) before assuming something is wrong with the channel and switch away. The channel change time is actually due to a combination of factors such as physical RF tuning, conditional access, parsing streams, and decoding the video stream. Decoding of the video is dependent upon the design of the video stream to allow for regular periods of random access while the rest of the channel change contributing factors are outside factors. A good strategy will allow channel changing to be fast, frequent, and output undisrupted clean video once the decoding process starts or in other words a graceful transition from one channel to the other. Additionally to work in an existing cable environment, a viewer should be able to switch to channels that have alternate video coding technologies (MPEG-2, analog) for the channel with again a graceful transition.

Channel Change requires a combination of signaling at the transport layer to identify the pertinent video transport packets and designing of the video coding pattern in the video layer to create a clean continuously displayable output video only using data after the access point. Some compression efficiency in the video stream will be sacrificed in order to allow for a consistent and fast channel change between video streams.

HOW MPEG-2 DOES RANDOM ACCESS

An MPEG-2 stream uses the Random Access Indicator (RAI) bit in the MPEG-2 transport (4byte) header adaptation field shown in Figure 3 to identify that this is the transport packet (188 bytes) that identifies an access point in the video stream. Once that transport packet is decoded, sequence header and picture header information is given to identify the dimensions and frame rate of the video. This also starts the process of identifying the next transport packets in the sequence.

At the video coding layer, the random access point identifies the start of a GOP (group of pictures) that consist of a pattern of I, P, and B coded pictures [see Figure 4]. The next random access point will be at the beginning of the next GOP which is normally encoded at intervals of 15 to 30 frames (1/2 second or 1 second) with some leeway for making some occurrences of the interval smaller to adapt to stream changes or stream conditioning situations. The types of pictures indicate the different type of prediction that is used in MPEG-2 video:

I-Pictures - These pictures use intra-frame compression and are not predicted from other pictures (no temporal compression). This is also a

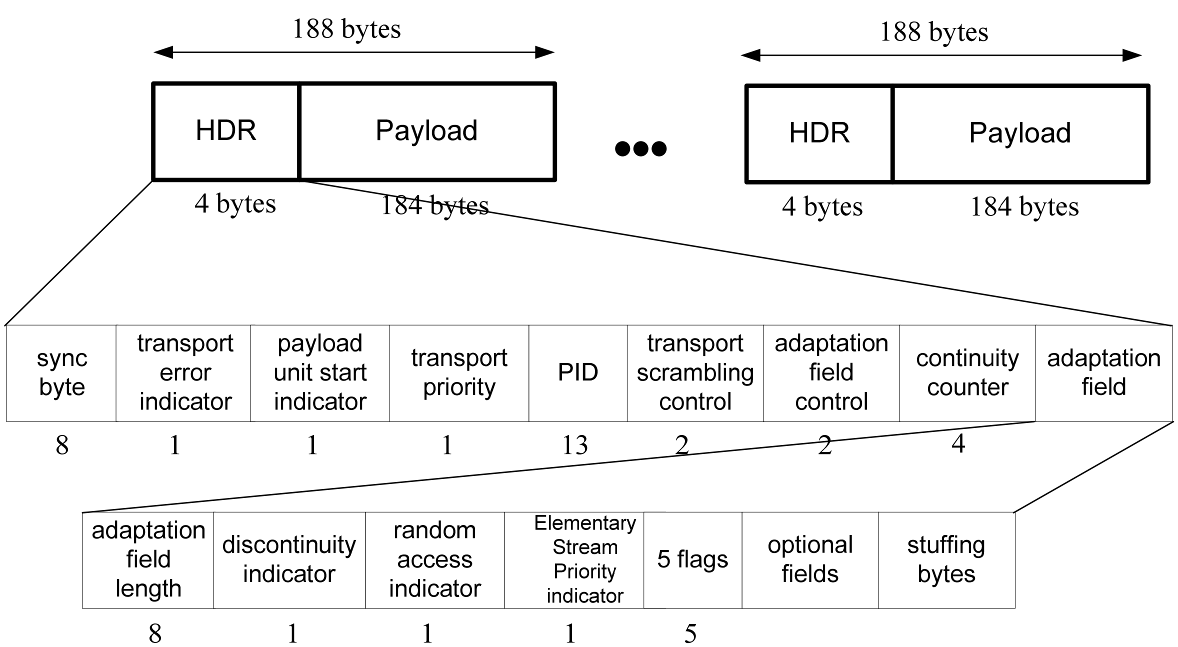

Figure 3. MPEG-2 Transport Stream Packet with Header Information

type of anchor picture (in AVC this could also be called a reference picture) which are pictures that can be used for prediction. The I-picture can be the start of a random access point in the video stream.

P-Pictures - These pictures can either use intra-frame compression or prediction from the nearest past anchor frame (forward predicted). The P-picture is also another type of anchor frame.

B-Pictures - In addition to intra-frame compression and forward prediction, these pictures can also predict from the nearest anchor frame in the future (backward predicted) as well as predict both forward and backward at the same time (bi-directionally predicted). The B-picture cannot be used in turn for prediction and is not an anchor frame.

There are a couple of things to note. Prediction is only dependent on the nearest anchor frames. This means that to decode a picture, at a maximum, one needs to store is a segment of the video stream that contains that particular picture and all pictures up to and including the nearest anchor pictures in both directions (so long as the anchor pictures are already decoded). In the reference picture buffer this allows for a natural "bumping" process to occur such that when an anchor picture gets stored and decoded then the oldest anchor picture can be discarded. Also of note, is that any picture acting as an anchor picture must come earlier in the transmit/decode order than the picture that requires that prediction (even though in cases of backward prediction that anchor picture is actually displayed later). Lastly, exiting (switching channels) before any anchor picture in transmit/decode order allows for the video displayed before the point to have no jumps in time (temporal continuity).

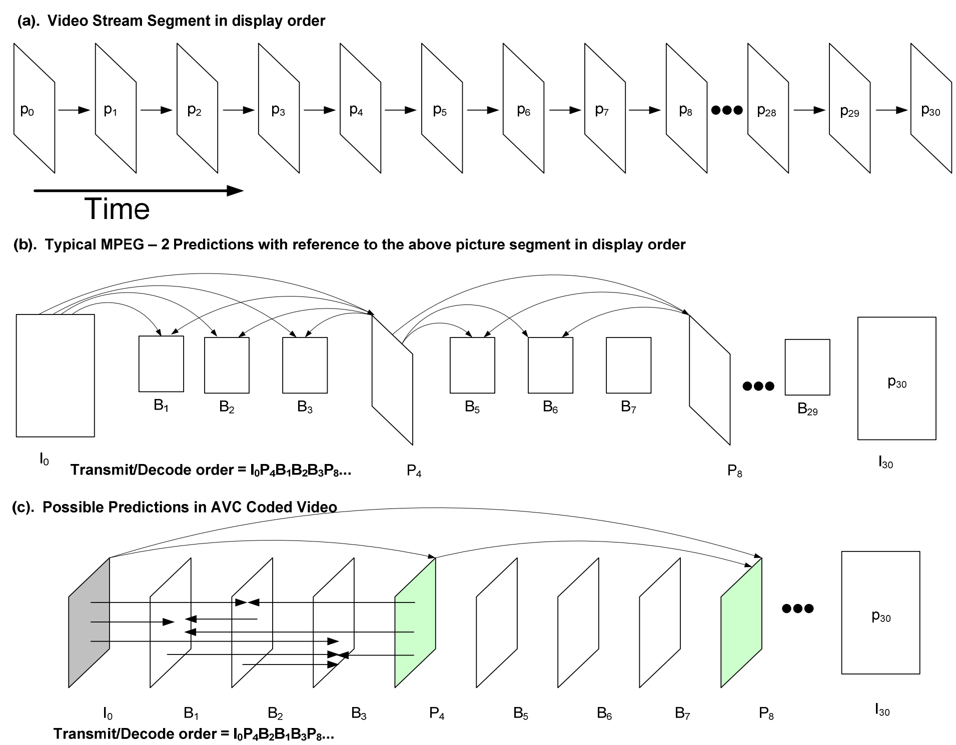

Figure 4. Possible MPEG-2 and AVC Prediction and Coded Pictures of Video Stream

The MPEG-2 GOP structure length and pattern indicates the video stream segment guarantees that all anchor pictures contained within are decodable and which allows for any pictures in smaller segments (mini-GOPS) predicted from the anchors pictures to be decodable. Having a random access point start at a beginning of a GOP allows for continuously displayable video from that point on. This in turn directly affects channel change time and the video quality in the channel change.

ADAPTING RANDOM ACCESS TO AVC

The MPEG-2 GOP structure is a very powerful mechanism for random access and the AVC encoding standard can accommodate this GOP structure, but this substantially limits temporal compression improvements that AVC can offer. Even though the concept of I, P, B are called the same as in MPEG-2, AVC has changed what they meant and also changed prediction rules between I,P, and Bs [see Figure 4]. These changes are:

Pictures – These are made up of combination of slices which can be of the same type or a combination of different types. Prediction from a slice requires storage of the entire picture. Pictures are differentiated between referenced (anchor) pictures and non- referenced pictures which are not used as predicted pictures. Using MPEG-2 Terminology (not really correct for AVC but often used), an "I" picture can have only I slices, "P" Picture can have I slices and P slices, a "B" picture can have I, P, and B slices.

I, P, B – These refer to actual slice types in the pictures rather than the pictures themselves. A single picture can contain different combination of these slices depending on the number of slices allocated to pictures. An MPEG-2 like picture can be emulated by limiting the combination of slices in the picture.

Predicted Pictures – These pictures whether forward or backward predicted do not need to predict from the nearest reference (anchor) frame. They can refer to any reference frame stored in the Decoded Picture Buffer (DPB) and are limited to maximum number of reference frames storable at any one time. Though one may only need to predict only from one slice of the picture, the entire picture must be stored. To emulate MPEG-2 behavior simply limit prediction to nearest anchor neighbors.

IDR – This is an intra-frame picture using intra compression techniques that also use spatial prediction between macroblocks. An IDR (Instantaneous Decoder Refresh) is a special intra-frame picture that indicates that any pictures subsequent to an IDR cannot reference (predict from) pictures before the IDR in transmission order (see Figure 5). Without an IDR tag, an intra-frame picture allows for pictures after it to reference or predicted from one that came before the intra-frame picture in transmission order. An MPEG-2 I picture can be emulated by an IDR. Some differences between using an IDR tag or not is that the bit rate spikes up around an IDR than if the tag was not used.

P - These slices can now be both forward or backward predicted. For MPEG-2 behavior, this would need to be limited to forward prediction to only I or P slices.

B - These slices can now be reference (anchor) frames and predicted in either direction from any two slices and can be weighted. To emulate MPEG-2 behavior, B's should not be able to predict from a B and forward and backward dual prediction should be used without weighting.

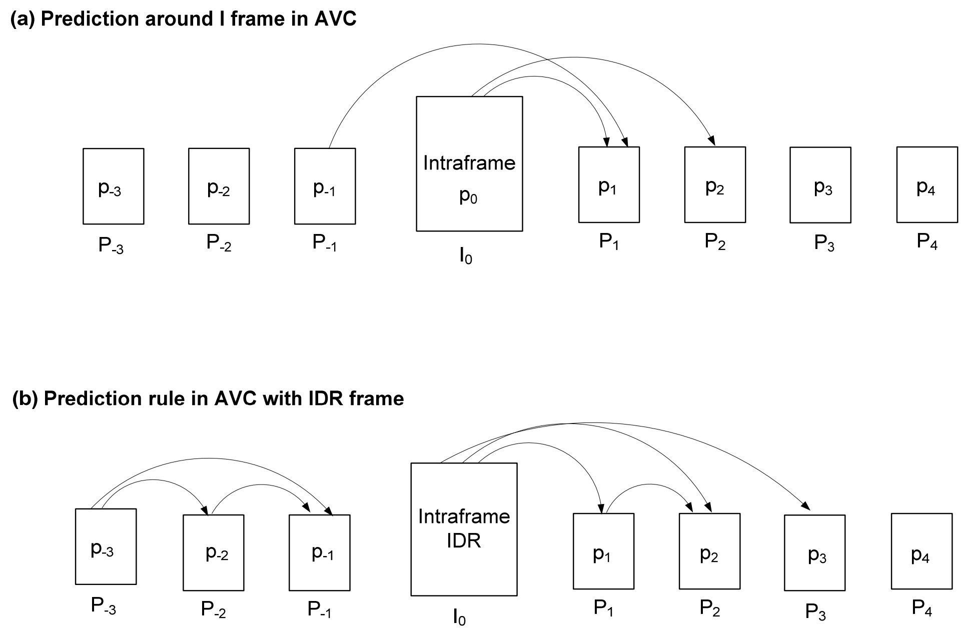

Figure 5. Difference of Intra-frame with and without IDR Signaling

The MPEG-2 GOP structure can be contained in the design of a cable-friendly AVC video stream, but many of the temporal compression improvement techniques that the AVC standard provides would be lost.

Providing for random access in cable should not make the AVC stream conform to an MPEG-2 structure but rather accommodate it and other approaches. In order to do this, rules need to be developed (not implementations) that allow for fast channel change times and good video quality during the change while not preventing too much loss in compression improvements. In order to do this, a concept of a segment still needs to exist and it is important that the transmit/decode order and the display order do not get too far out of alignment.

Some possible things to consider are creating something similar to a GOP length in AVC which would bound the channel change time. Another thing would be to limit the time segment/ stream segment that contain reference pictures that could be used for decoding a particular picture. The use of backward prediction in AVC can be extended over longer windows which can give more compression efficiency, but may cause problems in some cable applications. Backward prediction needs to be constrained such that it does not affect existing cable applications or such that stream conditioning can handle infrequent trouble spots that might occur. To keep the decode and display order from getting too far out of alignment, bounding the presentation time from the decode time of the random accessed picture could be used. Additionally setting a bound between the random accessed picture and the first clean picture that allows for continuously displayable video from that point on would also be useful in this process. Lastly providing some rules to allow for points of time continuity in the stream would be helpful.

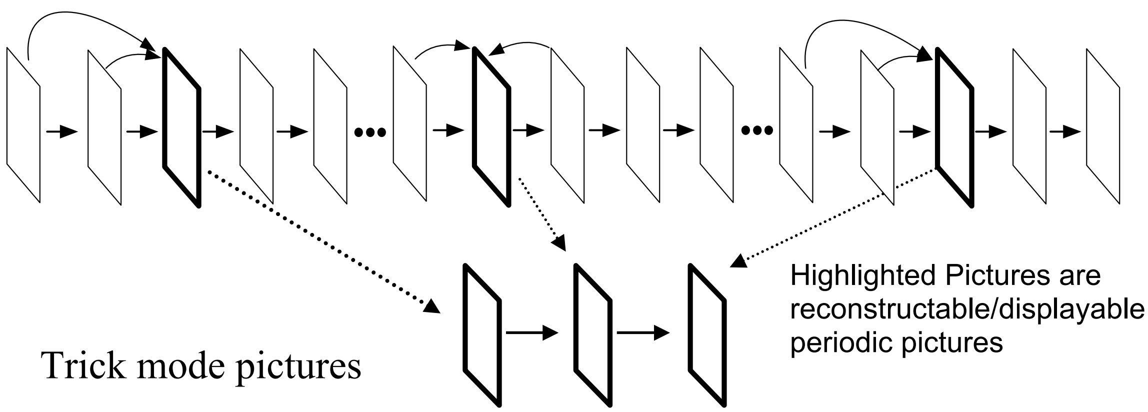

Figure 6. Generation of Trick Modes using Small Subset of Pictures

EMERGING APPLICATIONS: DIGITAL VIDEO RECORDERS

In the present cable environment, cable friendly video streams may also have to address other applications often in the same stream meant for broadcast. This would mean more stringent requirements on the RAP interval for streams that address these additional applications.

A popular application that is used with broadcast streams is the DVR storage and navigation of that stream. This is a service that viewers expect to have because it exists in the analog television and MPEG-2 domains today. In the AVC domain this is not as easy and unless constraints are used to ensure a compatible stream is provided, the DVR design may require a full decode/local encode of the stream prior to storage.

For the stream itself, regular self contained pictures or field pictures (for interlaced) need to be available at less than or equal to 1 second intervals. To avoid running decoder chips at faster than normal decode rates, the displayable pictures for DVR FFWD/RWD need to depend on a smaller subset of pictures in the RAP interval [see Figure 6 above].

Additionally slower FFWD/RWD rates that approach real-time playback need to consider smooth motion techniques between these displayable frames which would warrant a periodicity between sampled frames. On the transport level, several things need to be signaled about the video stream. One of these areas is identifying regular video starting points such as a start of the clean video display in the RAP interval. Another areas is the identification of reference versus non- reference pictures which identifies some pictures that may not be needed for a DVR FFWD/RWD decode. Lastly doing transport signaling in the clear can allow for manipulation of the video stream even though it may be encrypted on the hard drive.

AVC or any other advanced video codec needs to have the ability to do this. This may require consideration of the DVR application while designing the RAP interval. It would also require developing signaling in the transport layer which is a factor usually not considered in a video coding standard.

EMERGING APPLICATIONS: DIGITAL COMMERCIAL INSERTION

The digital program insertion (DPI) or insertion of local commercials or short programs is achieved by splicing from a network program stream to a local advertisement (ad) stream at a specific point in time. At the end of the ad a second splice occurs to get back to the network program. The timing of splice points is signaled by SCTE 35 [5] splice_info_section() messages carried in the transport stream (TS) packets of network program stream. In splicing, the exit point from the network program stream is called the OUT point or out of the network point. The start of the ad or the second stream is called the IN point. Similarly, at the end of the local ad-insertion window, the splicer exits the ad-stream and returns back to the network stream. So one DPI opportunity involves two splicings or two pairs of IN and OUT points [see Figure 7 below].

The requirement of splicing from one compressed stream to another compressed one is to ensure temporal completeness or temporal continuity of the exited stream. This is mainly achieved in MPEG-2 splicing by exiting on an anchor frame such as I or P picture. Sometimes this is called completing the mini-GOP. If one exits a stream in the middle of a group of B-frames, there will be a temporal jump in presentation space between the surviving B-frame(s) and the earlier anchor frame in decode order. When splicing between two AVC coded streams, this temporal continuity at an OUT point must exist as well. The use of IDR pictures can be helpful (but may not be necessary) to create this temporal continuity.

In the case of the IN point the stream is a new one to a decoder, so the stream should have all the parameters such as profiles and levels, picture resolution, etc., at the start that is needed for a decoder before starting to decode any picture. This is insured in MPEG-2 coded video stream starting with the sequence header, picture header, etc., followed by an I-frame. A similar solution in AVC would possibly be to start with a sequence parameter set (SPS), picture

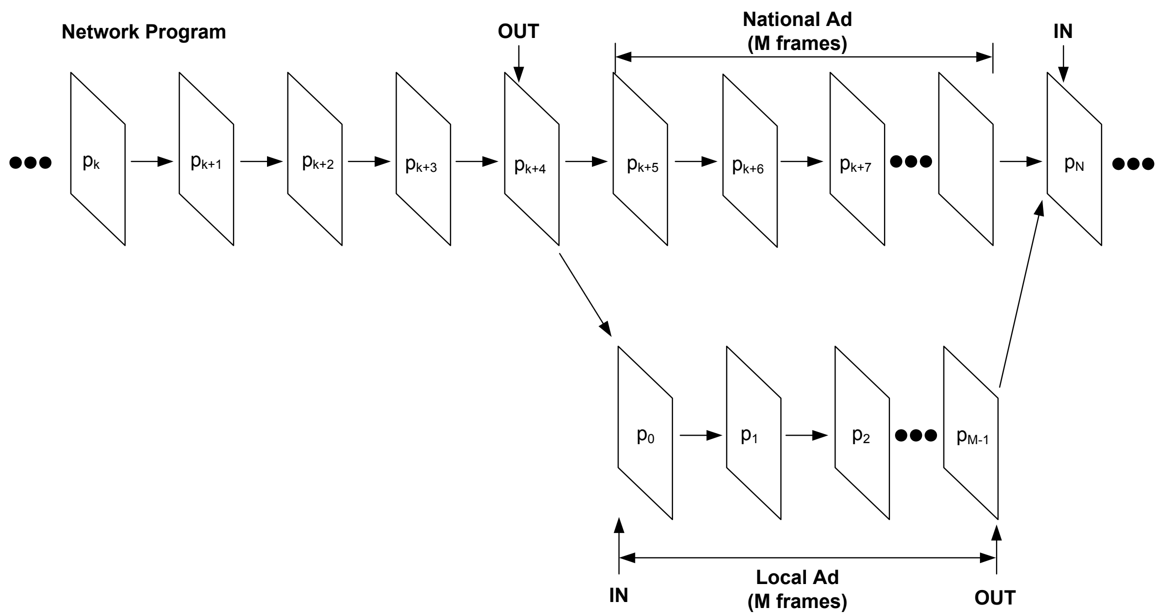

Figure 7. Simplified Diagram of Local Commercial Insertion

parameter set (PPS), etc., and an IDR picture. Also it may be recommended that one or more mini-GOPs may be encoded around any splice points just to avoid any splicing complexity due to delayed signaling for the splice points (ad-insertion). At the end of a commercial when the splicer returns to the network, the exited ad should also have- temporal continuity when returning to the network feed.

CONCLUSIONS

The AVC standardization process is in progress. Issues/requirements related to major broadcast applications such as channel change time, VOD, PVR, DPI have been investigated and defined. It has been apparent that some amount of bindings will be required to make it cable application-friendly. However, constraints will be minimized to achieve maximum compression efficiency while maintaining higher quality video than that produced by other similar standards such as DVB, Blue-ray DVD, HD-DVD, etc. The least amount of constraints will also leave room for encoders to produce a bitstream with potentially higher compression efficiency without compromising quality in the future.

ACKNOWLEDGEMENTS

The authors wish to thank CableLabs management for their support and encouragement in performing this work.

They would also like to thank the members of the advance codec drafting group in SCTE for many helpful discussions.

REFERENCES:

1. ISO/IEC IS 13818-2, International Standard (2000), MPEG-2 Video.

2. ISO/IEC 13818-2: 1996/Cor. 1:2000 (E) MPEG 2 Video Technical Corrigendum 1.

3. ITU-T Rec. H222.0 | ISO/IEC 13818-1 (2nd edition, 2000): “Information Technology – Generic Coding of moving pictures and associated audio: Systems”- and ITU-T Rec. H.222.0 | ISO/IEC 13818-1:2000/AMD.3: (2003): “Amendment 3: Transport of AVC video data over ITU-T Rec. H222.0 | ISO/IEC 13818-1 streams”.

4. ITU-T Rec. 264 | ISO/IEC 14496-10 AVC: “Advanced Video Coding for Generic Audiovisual Services” (May 2003) and Amendment-1 to 14496-10 “Fidelity Range Extensions (2004) ”.

5. ANSI/SCTE 35-2004, Digital Program Insertion Cueing Message for Cable.

6. SCTE 43-2005, Digital Video Systems Characteristics Standard for Cable Television.

7. ANSI/SCTE 54-2004, Digital Video Service Multiplex and Transport System Standard for Cable Television.

ADAPTIVE SWITCHED DIGITAL SERVICES IN CABLE NETWORKS

Michael Adams
Terayon Communication Systems

Abstract

Switched Digital Broadcast (SDB) is a new technology that has generated considerable interest because it promises better utilization of HFC bandwidth combined with the potential to offer the customer a virtually unlimited number of broadcast (linear) channels. However, progress on implementation has been slow, and the focus has been on vertical solutions that do not inter-work with other services such as digital ad insertion, VOD, and network-based PVR. Moreover, there is no clear revenue model for SDB; it is a technology rather than a service.

SDB promises the greatest increase in bandwidth efficiency when it is applied to the 'long-tail' content which is rarely viewed. However, it is difficult to predict which channels will represent this group and how real-world events dynamically change the popularity of a given channel. What is needed is an intelligent system that can adapt to instantaneous changes in channel viewing behavior.

This paper describes an approach to providing Switched Digital Services that combines the video digital processing required for ad insertion (both into broadcast and on-demand programming) with the ability to switch video in unicast or multicast modes. Therefore future digital switched services must seamlessly include support for advanced features such as targeted Ad Insertion, Fast Channel Change, VOD, and Network PVR.

INTRODUCTION

Switched digital services (often called switched broadcast or switched digital broadcast) are conceptually very simple; *"Only transmit those channels that are being watched"*.

In a cable system of any significant size every broadcast channel will have at least one viewer. However, deployment of Video on Demand (VOD) and High Speed Data (HSD) services into cable systems has naturally led to the creation of "service groups". Essentially, the cable system is segmented into a large number of HFC segments, each with its own dedicated bandwidth. By doing this, narrowcast services can be supported at high penetrations by re-using the same spectrum over and over. This is often termed "spatial re-use".

Within these service groups, which are typically between 500 and 2,000 homes-passed, the broadcast channels have a much lower probability of being watched. Trials conducted by Time Warner Cable in Austin, Texas [Brooks] indicate that, even at peak times, only 45 out of 170 channels will be simultaneously viewed in a service group that serves 1000 homes. The other 125 channels are just occupying spectrum and are completely unused. If 12 standard definition programs are carried in each 256-QAM channel, the total spectrum required for 170 programs is 84 MHz (14 QAM channels). And 74% of this spectrum is not being actually used – a tremendous waste of resources in a 750 MHz plant – in fact, about 9% of the total downstream capacity!

Switched Digital Services Algorithm

So how does *"Only transmit those channels that are being watched"* actually work in practice?

1. When a set-top 'tunes' to a switched program:
    - If that program is currently being transmitted in the service group, the set-top selects it as normal.
    - If that program is not currently being transmitted to the service group, the set-top signals to request it. The system then starts transmission of that program, and acknowledges the request. Then the set-top selects it.
2. When the set-top tunes away, the set-top signals to release the program.
3. When all set-tops within a service group tune away from a program, transmission can be stopped (if necessary) to recover channel capacity.

Benefits of Switched Digital Services

The biggest advantage is that new linear programming services can be added without limit to a switched digital services line-up. This "virtually infinite" linear programming line-up allows many niche programs to be added. This will be an important addition to a cable operator's ability to respond to competition from satellite and telco providers.

Another significant advantage is to reclaim bandwidth for other services, such as VOD, High-Speed Data, VoIP or Digital Simulcast.

In addition, there are several other advantages that can be gained by deploying switched digital services including targeted advertising, fast channel change, and integration with VOD services.

Service Group Size

In practice, a block of QAM channels is assigned to each service group, and the following "VOD math" is used to calculate the optimal service group size as follows. First we need to calculate how many digital set-tops can be supported simultaneously:

$$\text{\#Digital Set-tops} = \frac{\text{\#QAM Channels x 10}}{\text{Peak Usage}}$$

Let's plug some numbers into this equation; if a block of 4 QAM channels per service group is chosen, then each service group receives a maximum of 10 constant bit rate VOD streams, each encoded at 3.75 Mbps. If the peak usage is 7.5%, we have a service group size of 533 digital set-tops. Taking into account the digital set-top penetration gets us to the number of homes passed in the service group, as follows:

$$\text{Service Group Size} = \frac{\text{\#Digital Set-tops}}{\text{Digital Set-top Penetration}}$$

Let's take an example, where the digital set-top penetration is 50% of homes passed. In this case, we have a service group size of 1,066 homes passed.

SDS REFERENCE ARCHITECTURE

Figure 1 shows reference architecture for switched digital services.

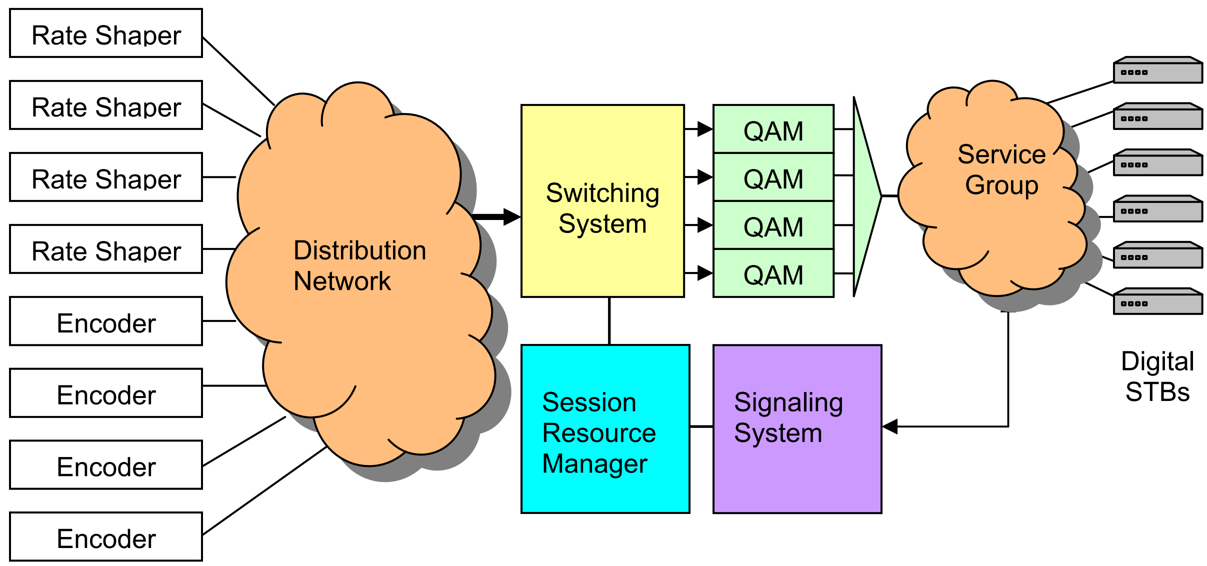

Figure 1. SDS Reference Architecture

The key features of the architecture are:

<u>Constant Bit Rate (CBR) encoders</u>

Before being switched, all programs are converted into MPEG-2 single program transport streams (SPTS) at a constant bit rate of 3.75 Mbps. This is the same format as is used for standard definition VOD content.

There are two ways of doing this:

1. If the channel is already available as a variable bit rate (VBR) program, it may be possible to rate shape it into a CBR program. However, this may not always produce good results. For example, a demanding sports feed such as ESPN, might peak at a significantly higher bit rate, and clamping the bit rate to 3.75 Mbps may cause a significant loss of picture detail. Visually, this will appear as tiling – where all the picture coefficients have been approximated to the point where macroblock boundaries become evident to the human eye.

2. A second option is to re-encode the program from a baseband (uncompressed) signal. This will typically produce a better result because the encoder has more flexibility in coding the stream than the rate shaper. Nevertheless, it can be very challenging to encode certain feeds into a 3.75 Mbps SPTS. In practice this equates to cost – requiring a more sophisticated and expensive encoder.

Despite these drawbacks, 3.75 Mbps was selected in early implementations because it makes the session resource management simpler. At 3.75 Mbps, 10 programs will comfortably fit into a 256-QAM channel.

In future, it is likely that multiple CBR rates might be supported. For example, 3.75 Mbps for movie channels, 4.5 Mbps for sports channels, and, of course, significantly higher rates for HD channels.

Distribution System

After the SPTS are built at the master headend, they must be distributed to the switching systems. An efficient way to accomplish this is to use IP multicast service. Each switching system will join the multicast group for those channels that need to be delivered to that service group.

Switching System

The output to each service group comes from a switch that essentially has n inputs, where n is the total number of available channels, and m outputs, where m is the maximum number of channels that can be viewed simultaneously.

In some implementations, the switch has been collapsed into the modulation system [Bombelli]. In others, it is a separate piece of equipment.

Modulation System

The modulation system is a bank of QAM modulators. Fortunately, high-density Edge-QAMs, developed for VOD, are ideal for this purpose. The number of QAMs assigned to each service group must be carefully chosen to reduce the probability of blocking to an acceptable level.

Set-top Client

For switched digital services to be deployed, existing set-top boxes must be able to support them. This is achieved by developing a special piece of client software that monitors channel selection, and signals to the system whenever a channel is selected or released.

In early implementations, a dedicated switched digital services client was developed. In future, an integrated client that supports VOD and switched digital services will be required to simplify software integration.

Signaling System

The signaling system allows each set-top to request a program when it is selected by the user. The signaling system that already exists for VOD signaling can be re-used for this purpose. In fact, even the protocol can be re-used with some minor extensions.

Session Resource Manager

The session resource manager (SRM) is responsible for making sure the correct channels are made available in each service group according to the set of signaling requests received over time. In implementations to date, a dedicated SRM has been developed for switched digital services (as opposed to re-using the existing VOD SRM). There are several reasons for this:

1. The switched digital services SRM must understand the concept of shared programs within a service group. In other words, if a second set-top requests a particular channel that is already being transmitted, no action is necessary.

2. The switched digital services SRM must provide an extremely rapid response to the user's program request. A typical response time of 100 milliseconds is desirable, which is an order of magnitude faster than most VOD SRMs.

3. Implementation simplicity – simply put, it is usually easier to start from scratch to develop an SRM from the ground up than to modify existing code.

IMPLEMENTATION CHALLENGES

Switched digital services are not without significant implementation challenges. Some are similar to VOD and others are unique to switched digital services. Fortunately, there are solutions to each of these challenges and, for clarity that is how we've organized the next sections.

One-way Consumer Electronics agreement

Digital cable-ready receivers built to the one-way "plug and play" agreement [NCTA], cannot signal to request a switched broadcast channel. The only solution is to keep a set of programs in the broadcast tier. Fortunately, not all linear programming needs to be constrained to this tier. In the case of digital simulcast, the same programs are provided in analog, and are therefore available in that format to a cable ready receiver, so these channels can be switched without impact. In addition, any new linear programming could be made available in a switched digital services tier. Obviously, it would be important for customers to be informed that a set-top box would be required to receive them.

The long-term solution is a two-way capable digital cable-ready receiver. Although progress on the definition of a single standard has been slow, several bi-lateral agreements with specific manufacturers have been developed and first units have been deployed. A switched broadcast services application would have to be developed and downloaded to these devices in order to allow them to participate in a switched digital services tier.

Program Blocking

Program blocking will occur at any time when the number of simultaneous program requests exceeds the capacity that has been allocated in a service group. This is the same effect that happens in any switched system, whether it is a telephone system or a VOD system. The difference here is that the customer's expectations have been determined by the behavior of broadcast cable, which doesn't normally block. In light of this what probability of blocking is acceptable? A lot of operators would say "zero!" Unfortunately, engineering the system to ensure that blocking will never occur means building in significant additional headroom.

However, it is possible to use some other tools that are already in the bag to ameliorate this situation. Statistical multiplexing is commonly used in cable systems to combine multiple variable bit rate (VBR) programs into a constant bit rate (CBR) QAM channel. Each program is processed in real-time to ensure that the instantaneous sum of the bit rates never exceeds the total capacity of the QAM channel. (If it does, even for a few milliseconds, MPEG packets will be dropped and the video will literally fall to pieces, as compressed digital video is extremely vulnerable to packet loss.)

Unfortunately, statistical multiplexing is too expensive to be applied to every service group because there are so many of them. However, if we observe that blocking is unlikely to occur in all service groups simultaneously, we can see that only opportunistic statistical multiplexing is needed. This could be implemented at a realistic price point by sharing the statistical multiplexing resource across a number of service groups. This leads us towards a model where the switching system incorporates a statistical multiplexing resource.

Impact of Channel Surfing

Imagine a user is channel surfing through the switched digital services tier. For each channel change, the system must establish a session for the new channel and tear-down the session for the old channel. Now imagine that a large number of viewers all start to channel surf at the same time (for example, when a popular channel goes to commercial).

The impact of channel surfing can be minimized by transmitting popularly watched channels even when they are not being viewed. (Essentially this is exactly the way the CPU cache works in a computer system.) Caching could be implemented by building a learning algorithm into the SRM. A simple first-in, first-out algorithm will work quite effectively and minimize the number of sessions that have to be established and released. A more sophisticated algorithm could weight channels according to the number of viewers watching them simultaneously.

In addition, a fast channel map indicating which the channels are in the cache can be sent via a broadcast carousel to all set-tops. On a channel change, the set-top client checks to see whether the program is already being transmitted, in which case the set-top can tune to it directly. (Note that it must still signal to request the program because otherwise the SRM would have an inaccurate picture of the system state, and the session carrying the program could be unexpectedly released at some future time.)

However, the signaling system and SRM will still have to keep up with the signaling load, even if the session establishment/release rate is reduced. In addition, we have to be sure that the cache (switched digital services tier size) is reasonably large, otherwise channels will be added and released constantly and we will be back where we started. (In computer science terminology, this behavior is aptly termed "thrashing".)

The best way to build a high-performance signaling system is to distribute it across multiple CPUs, each serving a set of service groups. This also has the nice property of reducing the failure group size. There is also the possibility for building a fault-tolerant system if each set-top client maintains a list of multiple signaling sub-systems. In this arrangement, if a request from the first signaling sub-system doesn't yield a timely response, the client re-tries the request with an alternate signaling sub-system.

Operational Impacts

There is no doubt that the introduction of switched digital services changes the model of a broadcast channel. SDS requires zero non-responders and 24/7 upstream network availability.

However, some additional robustness and graceful degradation can be built into the service:

1. SDS session manager periodically broadcasts tuning information for current state to ensure that the system stays in a consistent state. The set-top client is responsible for monitoring these messages and taking appropriate action. For example, if it is connected to a session but detects that its identifier is not in the reference list for that session, it must re-request the

session to prevent unexpected termination of the session.

2. In the event of a failure to reach the SDS session manager due to a communications path failure or failure of the SDS session manager itself, the STB client may choose to use the last tuning information received for the desired program.

3. In the event of an SDS session manager failure, the sessions can be left untouched. After the SDS session manager recovers, it should be possible for it to recover the system state by polling the set-top clients to retrieve their current state. Thus the impact of an SDS session manager failure does not cause each viewer to have to re-tune their set-top after the outage.

Program Hold Time

How do you know if anyone's watching a program (or if the STB was just left on)? This is analogous to making a phone call to someone, mid-way through the conversation you have to answer the door, and you completely forget about your phone call. Because there is an intelligent party on the other end, they will hang up and the call will be released. But in the case of switched digital services, the system doesn't know whether you are still watching the program or not.

One solution is to monitor the viewer's activity by keeping track of remote control commands. If no remote control commands are received for a preset time-out period, then an on-screen message can be generated to test is the viewer is still there (and paying attention). After a further time-out, the session would be released. Obviously this approach could cause confusion and irritation for some viewers; however it would only become necessary at some threshold where the number of free sessions is becoming almost exhausted, so during periods of lower contention for resources it is less likely this algorithm would have to be executed.

A related problem is that some viewers will manually tune to a channel and set their VCR to record a program on that channel at some later time, leaving the system unattended. In this case it is quite possible that the session would be released either before, or, more irritating still, during the program recording period. The solution is to educate the customer in the use of VCR control functions or to deploy more DVR capable set-tops.

Cost of Switching and Modulation Equipment

From the reference architecture, it is apparent that a lot of additional switching and modulation equipment will be required to deploy switched digital services. Let's take the example of a medium-sized cable system that passes 300,000 homes. Assume that we decide to use a service group size of 1,000 homes-passed, and that 50 sessions per service group are required. For every 10 sessions, we will need to modulate a 256-QAM channel, thus requiring 5 QAM channels per service group. With 300 service groups in the system, a total of 1500 channels of QAM modulation will be required.

Fortunately the solution for this already exists; high-density edge-QAM modulators were developed to provide a low cost modulation for VOD. Given the scale of modulation resources, any viable switched digital services deployment will leverage a highly commoditized edge-QAM market.

Switching cost is less of an issue because a single switch can easily serve a large number of service groups. In our example there are a total of 15,000 streams at 3.75 Mbps, or 56 Gbps of total switching capacity.

In addition, there are a great many advantages to be gained from deploying a narrowcast architecture such as switched digital services. Each session is consumed by a small group of viewers (in some cases only one viewer), and so there is a tremendous potential for customization of the content to that audience. This could be done using a combination of Digital Program Insertion (DPI) and graphical overlay technologies.

ADAPTIVE SWITCHED DIGITAL SERVICES

One of the most difficult challenges in a switched digital services deployment may be choosing which programs can be placed into the switched digital tier. Some programs must stay in the broadcast tier because of programming agreements or because they must be made available to one-way cable ready receivers. Others will be clear candidates for switched digital service tier, such as niche programming.

In the middle is a gray area of channels that could fit into either group according to system dynamics. In an adaptive SDS implementation, the caching algorithm will effectively promote switched channels to broadcast channels if they become heavily used. Moreover, the operator can manually promote them to un-switched mode on a per service group basis if necessary.

Benefits of Adaptive SDS

SDS is an interesting technology, but with significant deployment and operational impacts, operators are looking for additional features and benefits to justify its deployment.

SDS will be deployed first in bandwidth-constrained systems that need to add digital simulcast for competitive reasons but do not have the available spectrum for the additional QAM channels required to carry them. But as a pure bandwidth management tool, SDS is best viewed as a stop-gap measure.
Removing a channel from the analog tier creates tremendous bandwidth saving, but it is a slow process. Nevertheless, each channel removed from the analog tier allows 12 new SD digital channels or 2 HD digital channels.

We can start to see more potential benefits from switched digital services if we observe that, for the first time, the system has a direct way of measuring who is watching broadcast programs and when they are watching them.

Targeted Advertising

This information is of great interest to advertisers because it could support targeted advertising. In turn, targeted advertising promises increased revenues for the cable operator. Of course, special care must be taken to preserve the individual viewer's privacy but this is not incompatible with matching the advertisements to the viewer's interests based on established demographic classification.

How might this work in practice with SDS? When the first viewer tunes to a program, the system would start to insert commercials in that channel according to the demographic classification into which that viewer falls. When a second viewer joins the session, the system might deliver commercials that intersect the two viewer classifications. Alternatively, if the system is lightly loaded, the SDS resource manager could choose to deliver the program over a new session to the second viewer, enhancing the capabilities for targeting commercials.

Fast Channel Change

Fast channel change (FCC) is being offered by some competitive IPTV products as a benefit over conventional digital cable implementations. In FCC, the transition from one channel to another is accomplished more quickly due to two changes to the architecture:

1. The channel change is accomplished by switching in the network, rather than in the set-top, so there is no re-tune latency.

2. The new channel is conditioned such that it starts with an intracoded picture – so that all of the information to display the picture arrives immediately on channel change.

The combination of these techniques can produce a maximum channel change time that is quite fast (i.e. better than analog cable and maybe one fifth that of digital cable; for example, 0.2 seconds versus 1 second is achievable).

In the case where a viewer switches to a SDS program that is not yet being transmitted in the service group, a similar technique could be employed to commence from an intracoded picture, avoiding the delay of waiting for the next I-frame.

In the case where a viewer 'tunes' from one SDS program to another, and is the sole viewer of the program in the service group, the existing session could be re-used and a program switch achieved in the network.

Transition to Network PVR

Some operators are experimenting with PVR services, where the network allows time-shifting of programming using VOD technology. An example of this is Time Warner Cable's "Startover" service [Yahoo!]. The Startover service allows the viewer to watch a particular program from the beginning after selecting it part-way through. To achieve this, the viewer is switched to a VOD session for a delayed version of the program.

A Startover session is actually a type of SDS session, where the program is unicast and the program is streamed from a VOD server rather than a real-time programming source. In fact the SDS session can be re-used when the viewer transitions from a broadcast program, where he or she is the only viewer within the service group, to a Startover session. However, to do this the modulation resource must be shared between Startover (the VOD system) and the rest of the SDS system.

In addition, Startover sessions are another great opportunity to increase the value of advertising by targeting the commercials so they are more likely to be related to the viewer's interests. That same SDS targeting infrastructure can be seamlessly re-used for inserting commercials into a program regardless of whether it is a broadcast or VOD session. In fact, the same SCTE 35 digital cue messages that indicate the ad "avails" can be used to support commercial insertion.

Ad-supported VOD sessions can also re-use the same infrastructure required for SDS. Once a SDS system is deployed it only makes sense to share modulation resources across all switched services, whether they are broadcast programs, Startover sessions, or VOD sessions.

Requirements for SDS systems

Now that we have discussed some of the benefits that can accrue for the deployment of SDS, we can start to develop a set of requirements for SDS systems. In order to use SDS as more than a stop-gap measure for bandwidth management, the system should support the following requirements:

1. Support targeted advertising by incorporating DPI capabilities into the switching system.

2. Support fast channel change by incorporating the ability to synchronously deliver an intracoded picture to the newly selected channel.

3. Seamlessly integrate with Startover by sharing modulation resources with the VOD system.

4. Support opportunistic statistical multiplexing to provide additional, "elastic", headroom when necessary under periods of peak demand.

The common thread of all these requirements is that the switching system is the ideal point in the system to provide the additional functions and benefits we have discussed. This is an interesting conclusion because it flies in the face of the conventional wisdom of some proposed switched broadcast implementations, where the switching system is incorporated into the modulation system in order to minimize cost.

CONCLUSION

In this paper we have summarized the various challenges, solutions, and benefits to deploying switched digital services.

We have argued that enhancements to the switching system can generate significant incremental revenue by supporting targeted advertising in broadcast and VOD programming. In addition, competitive features such as fast channel change could be introduced by adding flexibility into the switching system. Finally, by incorporating opportunistic statistical multiplexing, the bandwidth management properties of the system can be enhanced.

Therefore, in a SDS solution, we argue that the operator should invest additional resources into the switching system in order to realize these benefits.

REFERENCES

1. Brooks, Paul: Time Warner Cable Digital Broadcast Field Trial Results Austin, Texas. SCTE Cable-Tec Expo 2005 Proceedings and Collected Technical Papers.

2. Bombelli, Lorenzo and Brooks, Paul: An Open Architecture for Switched Digital Services in HFC Networks. SCTE Conference on Emerging Technologies 2006 Presentations and Collected Technical Papers.

3. NCTA, Cable MSO-Consumer Electronics Industry Agreement on "Plug & Play" Cable Compatibility and Related Issues. December 19, 2002.

4. Yahoo! Entertainment: Cable Comes to the Rescue of Channel Surfers, Dec 23, 2004.

APPLYING DIGITAL WATERMARKING TO HOME ENTERTAINMENT CONTENT DELIVERY NETWORKS

Reza Rassool, Chief Engineer
Widevine Technologies Inc.

Abstract

Watermarking technology, answering TV industry requirements, (1) has emerged in trials over the last few years. While the initial deployment focus has been largely on future digital cinema applications (2), there are significant challenges of scalability, performance, and economy in adapting the same technology to today's content delivery networks (CDN) to the home. This paper, drawing from RFI (3) responses for candidate technology for the Widevine 'Mensor™' (4) digital forensic system, discusses the compromises necessary to engineer an economical watermarking solution for home entertainment.

INTRODUCTION

Anti-piracy, digital rights management and copy protection technology have long been the subject of scientific endeavor in the entertainment industry. Specifically the quest for an imperceptible means to deter and track unentitled usage of content has resulted in technology called *watermarking*.

Since 2004 the content owners have increasingly included the term 'watermarking' in their questionnaires to service owners who seek content. While this language has not yet translated into specific requirements in carriage contracts, they do give service operators and manufacturers notice of impending conditions for obtaining content in the future.

In 2005, Widevine Technologies Inc., a supplier of security solutions to the entertainment industry issued a request for information (RFI) to core technology suppliers for watermarking components to be included in its Mensor digital forensic product line. This paper charts the technical inquiry to derive a set of requirements for watermarking for home entertainments networks. It draws from the RFI responses while respecting the confidentiality of the respondents.

REQUIREMENTS

Digital Cinema Initiatives, LLC (DCI) was created in March 2002, as a joint venture of Disney, Fox, MGM, Paramount, Sony Pictures Entertainment, Universal and Warner Bros. Studios. DCI's primary purpose is to establish and document voluntary specifications for an open architecture for digital cinema that ensures a uniform and high level of technical performance, reliability, and quality control. The Digital Cinema System Specification V1.0 includes requirements for the watermarking of content. In September 2002 the European Broadcasting Union (EBU) reported (1) on the trials of watermarking technology to meet the needs of digital television. The EBU differentiates the watermarking requirements of the W1 (contribution), W2 (distribution) and W3 (consumption) segments of the content pipeline.

Of the two sets of requirements, the digital cinema business appeared to capture the immediate focus of technology suppliers. Even the emphasis of the EBU was on the W1 and W2 segments of the content pipeline. Now, as IPTV emerges as a significant market for entertainment content, technology suppliers are having to modify their strategies concerning forensics as, they either retool or design anew solutions for W3 – watermarking content right to the home and beyond.

Why Extend Watermarking to the Home?

The commercial rationale for deploying digital forensics to the home entertainment network can be simply stated – It may become a contractual pre-requisite for obtaining content.

1) Theft deterrence

Watermarking is proposed as a theft deterrent – to keep honest people honest. In this respect users should be made aware of the watermark's existence. Ironically the most carefully engineered quality of a watermark's invisibility detracts from its deterrent effect. To augment the deterrent effect it is important to use other means to alert the user that the content they are accessing is watermarked with a payload that uniquely identifies their client device and the time of access. This can include a visible mark or a warning introduction to the content.

2) Carriage contract questionnaires

Without divulging confidentiality it can be said that the studios have increasingly employed language in questionnaires ranging from a general query,

> "Has your company deployed any forensic watermarking (invisible) technology? If so, please describe in detail."

to very specific questions regarding the watermarking capabilities of service operator's network.

> "Is the STB capable of session based watermarking for high value content?
> 1. Please identify the watermarking technology used and the payload.
> 2. Please describe the forensic marking process.
> 3. Please describe the forensic marking detection and recovery process.
> 4. Please describe the robustness of the watermark in terms of survivability to obscuration, down rez, or overmarking."

3) Tracking

Watermarking is also useful in providing evidence of theft in criminal proceedings. The Mensor technology has already proven itself in court. Twenty-eight of the world's largest entertainment companies brought the lawsuit against the makers of the Morpheus, Grokster, and KaZaA software products. Evidence for the plaintiff was provided by early Mensor technology. The case was decided on June 27, 2005 in favor of the plaintiff. [7]

4) Leak detection

Watermarking is further used to detect breaches of security in a content delivery network. Leak location in a multi-node CDN requires that the watermarking scheme supports:

- embedding at multiple nodes
- multiple overlapping marks or mark replacement

5) Security renewal targeting

Through the extraction of the payload from pirated content the service operator, aggregator, or content owner is able to identify the node of the CDN from which the content was obtained. This allows for the targeted renewal of security in the conditional access systems of the CDN. The cost of security renewal can be prohibitive without such targeting.

6) Copy protection

Watermarking has also been promoted as part of a copy protection system. This scheme, which has had limited success, requires that a client device is capable of reading watermarks and respecting the copy control information contained therein.

7) Content tracking

Watermarking additionally provides a means to embed metadata into content that will survive numerous transport mechanisms. This usage of watermarking is used to monitor and audit the delivery of paid content such as advertising. (5)

8) QoS optimization

A hidden bonus for service operators is the promise of using the feedback of extracted metadata from the edge of a CDN to optimize the serving of content. VOD content is transmitted by an aggregator, such as TVN, to hundreds of nationally distributed VOD servers. Local usage data could be used to tune the forward transmission of the most popular titles in a particular area.

9) Enabling new business models

As peer to peer files sharing network companies scramble to legalize – watermarking technology offers a solution. One can imagine that a P2P player could extract and respect an embedded payload, and then allow the user to purchase an entitlement for legitimate access to the content.

ENGINEERING CHALLENGE

The challenge for engineers is to glean rational requirements from the esoteric messages from the industry and sparse data points in related fields. Specifically, the challenge facing the author's colleagues was to design a digital forensic system that would be economically and computationally appropriate for deployment in content owner/aggregator facilities, operator head-ends, and customer devices.

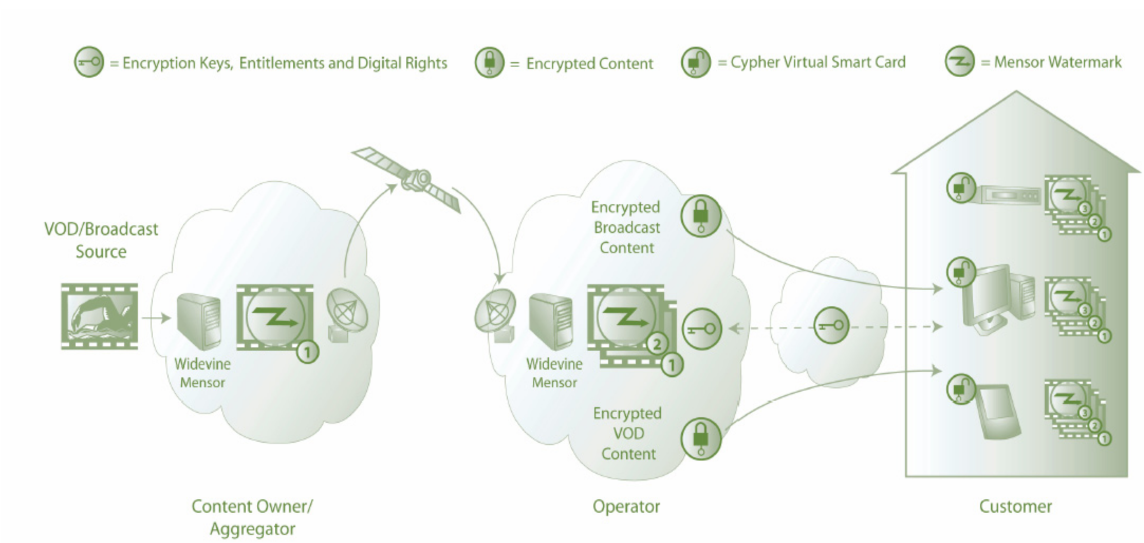

Figure 1. Watermarking applied to a multi-stage content delivery network

Figure 1 shows a multistage content delivery network. The content originates with a content owner or content aggregator. The content is transmitted to a service operator. The diagram shows a satellite transmission, however, other means are also used. The service operator serves both broadcast and VOD content to numerous customers.

Build or Buy

There are numerous attractions to originating core technology specifically tuned for an application; however engineers rarely have that luxury. In the case of watermarking there exists a wealth of prior art and core technology that has already gained industry acceptance. In the security field there is a distinct disadvantage to 'home grown' algorithms. Specifically, cryptographic products with industry standard algorithms such as AES have greater credibility than those with proprietary algorithms. Another factor particular to watermarking is the reality that Digimarc Corporation has a strong patent portfolio with far reaching claims. These factors encouraged us to adopt the course of licensing and integrating 3rd party watermarking components into a design for a digital forensics system.

Adapting Existing Technology

It is clearly unfeasible to place a $20,000 watermarking engine, designed for post production applications; into a sub $100 set top box for session based watermarking. The technical challenges include:

- scaling the head-end embedder to handle hundreds of streams
- implementing watermarking algorithms within the CPU budget of a set top box processor
- integrating watermarking into a security framework.

In order to understand how to adapt the existing algorithms, it is necessary to formulate a model to quantify the variables by which we can judge a watermarking scheme.

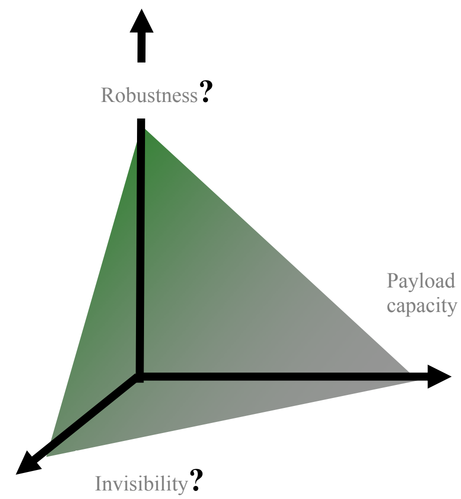

Figure 2. Watermarking dimensions

The concept in the figure above is often expressed but very rarely substantiated. Of these three dimensions only payload capacity is quantified.

Invisibility is a Pre-requisite

Even though the EBU requirements state:

"Note that the invisibility of the watermark, which is a typical subjective consideration, is different for W1 (at contribution level) and W2 (at contribution end-user level)." [1]

The result of analysis of watermarking algorithms shows that it is not practical to vary the invisibility of the watermarking algorithm. The method of spread spectrum coding of a signal through selective DCT modifications, exhibits a cliff effect with regards to invisibility (6). The mark is either invisible or obviously visible.

Through interviews with Hollywood content owners we found little acceptance for watermarks that were anything but invisible.

By accepting invisibility as a prerequisite, we now have the variables of payload capacity and robustness to tune to arrive at an optimally engineered solution.

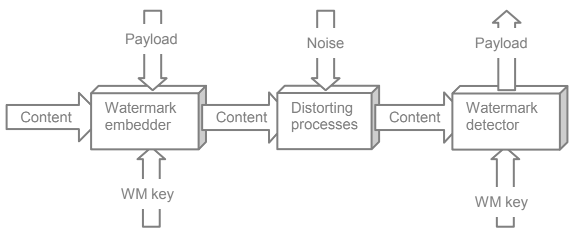

Figure 3. Watermarking dataflow

The figure above shows the typical end-to-end dataflow for watermarking. The watermarking embedder uses a watermarking key, WM key, to securely embed a payload into the content. The payload is embedded into a segment of content define in time as the watermarking minimum segment (WMS) (1).

We define here the notion of watermarking strength, S_w in terms of signal to noise ratio.

If the signal rate is R_s and the payload rate is R_p. Then with a payload size of P bits:

$$R_p = \frac{P}{WMS} \quad \textbf{(b/s)} \qquad \text{Equation 1}$$

$$S_w = \log\left(\frac{R_s}{R_p}\right) \qquad \text{Equation 2}$$

Additionally in the figure we see that a watermark's robustness can be defined in terms of the watermark to survive noise added into a distorting process. Robustness is not easily quantifiable. Among enumerations of various transformations or signal distortions to which the watermarked content should be subjected, the industry requirements commonly define robustness in terms of the lowest compressed rate, R_c, from which the payload must be recovered.

We focus on this quantification of robustness and then as in equation 2 express the maximum distortion, D_{max}, in terms of a signal, R_s to noise, R_n, ratio.

$$D_{\max} = \log\left(\frac{R_s}{R_n}\right) \qquad \text{Equation 3}$$

Now we can go further and express robustness, R, as the ratio of watermarking strength and maximum distortion.

$$R = \frac{D_{\max}}{S_w} \qquad \text{Equation 4}$$

Restated, this becomes

$$R = \frac{\log\left(\frac{R_s}{R_n}\right)}{\log\left(\frac{R_s}{R_p}\right)} = \log_{Rp} R_n$$

Or

$$R_p{}^R = R_n \qquad \text{Equation 5}$$

This eliminates the content signal rate from the relationship. R is a characteristic of a specific watermarking algorithm. Equation 5 indicates that for a given engineered robustness, R, as the payload size or WMS is increased there is a logarithmic in Rn, the lowest compression rate from which the watermark can still be extracted.

Robustness is increased by redundantly inserting the mark into the content which entails increased computational expense.

Requirements	DCI System Spec V1.0	EBU W1	EBU W2
WMS	5 mins	1 sec	5 sec
Payload size	35bits	64bits	64 bits
Signal data rate	JPEG 2000 @250Mb/s	SDI @270Mb/s	MPEG-2 @ 8Mb/s
Robust to compression rate	1.1Mb/s	2Mb/s	2Mb/s

Table 1. Quantifiable EBU and DCI requirements

The DCI and EBU requirements are shown in Table 1. From these we can calculate values of robustness.

The head-end watermarking embedder, Mensor Server embeds a payload of 21 bits while the Mensor Client embeds a payload of 64 bits in session based watermarks. The following table shows the resulting robustness. Note: The results have been scaled using content bit rate in Mb/s and payload bit rate in bit / second.

	Robustness
EBU W1	0.322
EBU W2	0.055
DCI	0.240
Mensor Server	0.041
Mensor Client	0.039

Table 2. Robustness

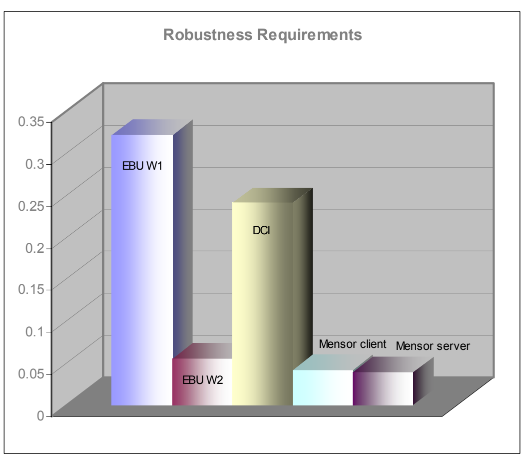

Figure 4. Robustness requirements of various watermarking systems

CONCLUSION

Robustness can be quantified and related to payload bit rate by

$$R_p{}^R = R_n$$

In engineering a cost effective watermarking solution for home entertainment, we utilize savings in reducing robustness. We see that EBU W1 and DCI require high robustness while EBU W2 and Mensor call for an order of magnitude lower robustness. Reduction is robustness translated into a reduction in cost and has enabled us to arrive at a cost effective home entertainment watermarking solution.

BIBLIOGRAPHY

[1] Choosing A Watermarking System For Digital Television - The Technology And The Compromises
L. Cheveau, European Broadcasting Union, Switzerland, IBC 2002 Conference Proceedings

[2] DCI Digital Cinema System Specification v1.0a

[3] Content Watermarking System Request For Information (RFI) V1.1, Widevine Technologies Inc.

[4] Antipiracy - Trends And Technology (A Report From The Front)
R.P. Rassool, Widevine Technologies Inc., IBC 2002 Conference Proceedings

[5] Teletrax™ Technical Specification, Philips.

[6] Robust Hash Functions for Digital Watermarking, Jiri Fridrich and Miroslav Goljan, Center for Intelligent Systems, SUNY Binghamton

[7] MGM Studios Inc., et al v. Grokster, Ltd., No. 04-480 on the Supreme Court's docket, June 27, 2005.

ARCHITECTURES FOR ADVANCED ADVERTISING – COMMERCIAL DRIVERS AND ENGINEERING SOLUTIONS

Guy Cherry, C-COR Incorporated
Rick Bohrer, Atlas On Demand

Abstract

A key to the future success of the cable industry will be the industry's response to emerging opportunities in advanced advertising. This paper will explore four key technological leverage points expected to exhibit significant change over the next five years. Today's work in each area will be outlined and an attempt will be made to predict the advertising infrastructure required over the course of the next five years.

INTRODUCTION

A key to the future success of the cable industry will be its response to emerging opportunities in advanced advertising. Traditional advertising infrastructure has already seen changes in the first half of this decade brought on by computer technologies, digital delivery, and changing viewer habits. It is our belief that these changes will accelerate in the next five years bringing about a very different relationship between viewers, advertisers and product suppliers.

The cable advertising relationships that will develop over the coming years are likely to mirror, and perhaps improve upon, the relationships that are already in place and supporting the dynamic growth seen today on the Internet. This paper will explore and compare the evolving landscape of the cable advertising infrastructure with that being used in Internet advertising today.

There are at least four significant areas where change has begun to occur. The first and most fundamental is the replacement of pure broadcast delivery with digital pointcasting and narrowcasting (e.g. VOD (Video On Demand) and SDB (Switched Digital Broadcast) services). These technologies allow customized viewer experiences and can allow advertising campaigns to target specific individuals and populations with greater effectiveness. This population targeting is similar to a technique called behavioral targeting being explored today in the Internet advertising community.

A second area of change is the establishment of an "advertising return path" allowing subscriber behaviors, such as purchase history and program viewing history, to be anonymously tracked and used to deliver a more relevant advertising message at the next opportunity. Again note the similarity to Internet advertising where purchase and site-visit history can be used to select advertising that is more relevant and useful to individual subscribers. Note that ad placement decisions are enabled by a return path that delivers very granular data on each ad and each subscriber.

By 2010 the cable advertising industry will require flexible and standardized architectures for leveraging population targeting and return path capabilities. Future back-office architectures must allow any subset of the subscriber population to be targeted and to allow them to respond with purchase requests and viewing preferences in real time. The SCTE DVS (Society of Cable Television Engineers Digital Video

Subcommittee) Working Group on Digital Advertising has generated a number of successful standards supporting digital advertising and is currently working on standards in these key areas.

Note again the similarity to the Internet today where advertisers, portals, web sites, 3rd party ad-servers and analytics providers cooperate to deliver and optimize online advertising campaigns.

These two changes will be complemented by changes on two additional fronts. The third is the shortened timeframe for making ad placement decisions, and a near-immediate ability to update the logic that drives those decisions. The cable architectures being proposed mirror those operating on Internet, and include real-time "Ad Selectors" which are part of the control flow. These ad-decisions will be driven by changes in programming (e.g. "live" events) and by measuring the effectiveness of ads via the return path mentioned above. This shortened timeframe will be yet another key element in future architecture and standardization efforts.

The fourth area considered in this paper is the changing "composition model" for advertising campaigns. The traditional 30 or 60 second "spot" will continue to be augmented by other delivery models such as barkers, scrolling marquees, and other graphical overlays. The capability to follow through with email, phone contacts and pointers to longer-format content is also beginning to play a key role. A coordinated delivery of mixed advertising content, with follow-through, will continue to give advertising campaigns greater leverage.

This paper will explore these four broad architectural areas, will examine the key interoperability points that will benefit from industry standardization, and will compare these to advertising on the Internet today. Today's work in each area will be outlined and an attempt will be made to predict the advertising infrastructure and capabilities needed over the course of the next 5 years.

ADDRESSABILITY AND INTERACTIVITY

Addressability is the ability to direct specific streams to specific sets of viewers. This capability provides a powerful leverage point for advertising systems.

Some cable systems utilize addressability but not an interactive return path. For example, zone-based linear advertising is implemented to address specific sets of ads to specific geographic zones. The intention is to target particular populations with specific advertising campaigns. However, this is a "blind" operation – there is no mechanism to measure the effectiveness or interest level in the advertisements on a zone-by-zone basis.

A return path can leverage addressability in many ways. Consider the following examples:

1) Traditional VOD is based on a digital network's ability to pointcast streams in response to individual requests, and this is directly analogous to the delivery of Internet advertising. This allows the targeting of ads based on individual subscriber preferences.

2) Switched Digital Broadcast (SDB) also leverages the presence of a return path allowing bandwidth to be tailored to user channel selection preferences in real time.

In both cases tuning choices are processed by back-office servers that make bandwidth allocation decisions and setup the resulting sessions. At the point where these actions are taken (sometimes referred to as the Session Resource Manager or SRM) advertising decisions can also be made and carried out.

Advertisement decisions can be based on the content selected and/or knowledge about the specific individual. On the Internet, these two types of addressability are referred to as contextual-targeting and behavioral-targeting respectively. There are variations of each. For example, contextual-targeting might be dependent only on the channel or VOD category while behavioral-targeting might be based on population groupings and aggregate statistics. Targeting of this type can be highly effective and is enabled by a return path providing detailed data on each and every interaction between each individual subscriber and each advertisement.

SDB presents interesting opportunities for both contextual targeting and behavioral targeting. Contextual targeting is enhanced due to the greater diversity of channels available. In fact, when bandwidth allows, it may be desirable to manage multiple versions of the same network channel but with different ads targeting different demographics. Behavioral targeting is possible whenever the subscriber count of a SDB channel is one or, possibly, a low number.

Recent studies of Internet ad awareness have shown that behavioral targeting is considerably more effective at making an impression than contextual targeting. One recent study, which tracked eye movement, measured a 50 percent greater impact with behavioral ads than contextual ads after two displays of the ad.

It is likely that many experiments will be attempted in the area of addressable advertising over the next few years. This will necessitate flexible architectures and active standards activities.

Further variations on addressable advertising include:

1) Allowing the user to launch nested sessions to obtain further information. This is called "telescoping"; it provides a direct feedback loop and is reminiscent of the browse dynamics of Internet surfing.

2) Simply tracking subscriber trick play operations can also give an indication of user interest. An example of this type of feedback is described in greater detail in the next section. Again, this is analogous to the Internet where various levels of interactivity are tracked, e.g. an impression, a click, or more engaged interaction with "rich" advertisements.

Addressability and interactivity are the key leverage points for change over the next five years. We will therefore interject an example before discussing the third and fourth axis of change: market dynamics and composition.

AN EXAMPLE

The VOD advertising architecture shown below in Figure 1 supports dynamic addressability.

The key components of the system are:

1) An Ad Execution System (AES) capable of managing sessions composed of

advertisements and on-demand content. The AES can easily modify which ads are shown in which avails and can change these associations dynamically as required by the ADS. The AES is also capable of collecting detailed session data and reporting this data to the ADS.

2) A (possibly remote) Ad Selector (ADS) which communicates with the AES through an early variation of a SCTE/DVS/WG5 draft protocol. The ADS is responsible for all ad selections and placement decisions.

The dynamics of this system can be grouped into three stages as shown in the diagram. The steps labeled A1, A2, and A3 are initialization steps which occur before sessions are launched. Content (both advertisements and entertainment content) must be loaded onto video servers' inventories and must be communicated between ad sales organizations and the delivery system. In addition the ADS must share initialization data with the AES.

Session setup steps are labeled S1 through S5. These steps very closely resemble the classic VOD session-setup steps managed by a Session Resource Manager (SRM). In this extended message flow the SRM either references pre-loaded ad bindings or communicates directly with an ADS to access a set of ads to be pre-pended to the viewer's request.

The exchange labeled P1 denotes the periodic reporting of viewer data from the AES to the ADS(s). This data includes trick-play statistics for all advertisements and can be used by the ADS(s) to intelligently select future ads.

The three principal components shown in Figure 1 closely parallel the current components seen in the field of Internet advertising. In the Internet architecture the Ad Management infrastructure communicates with a web site (as opposed to a Headend) which in turn supports a browser. The interactions labeled A1, A2, A3, and P1 directly analogous. The session setup steps, S1-S5 serve the same function in cable as content delivery on the Internet.

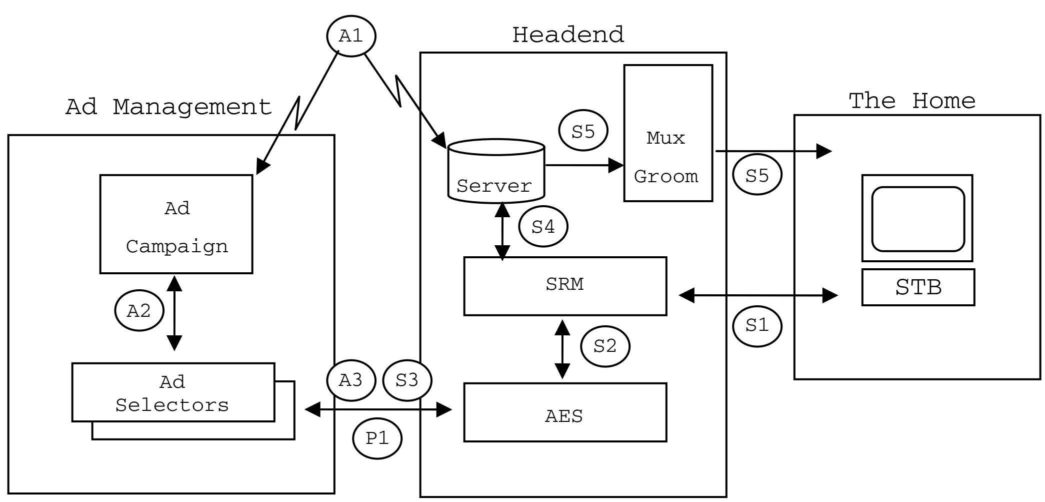

Figure 1. Adaptive VOD Advertising Example

Note that the ADS must be aware of the set of avails and advertisements available to it at all times (communication labeled A2 in the diagram). In the early implementations these sets can be static and downloaded manually into the ADS and VOD server. However as these systems become more sophisticated market dynamics will require standard APIs for coordinating the changing sets of avails and advertisements.

In this system subscribers have full trick-play control of the content including the advertisements. Thus subscribers can fast-forward through ads or rewind to watch an ad again. Each of these user actions, however, is noted and the interaction can be measured for each advertisement. As on the Internet, these interaction metrics can then be used to adjust future ad selections for the individual.

This architecture can support advertising placement ranging from basic to advanced:

1) Initial systems user bumper ads placed before or after user selected on demand content. Bumper ad selection can be entirely content-based or, with the right system capabilities, can change rapidly based on knowledge of the actual subscriber at session setup time.

2) Interstitials ads are a possible future addition to this architecture. It is not clear, however, whether interstitials will be a welcome or useful advertising tool for movie content. Interstitials will be needed, however, for other types of content such as sports events.

3) There are a variety of ideas and open questions for handling pause, resume and trick-play operations. When resuming a session subscribers can be shown the same ads or updated ads. These may be displayed at the resume point or only at the original avail points. Other policy decisions include behavior after a rewind and fast-forward: should the system show the same ads or replace the originals?

MARKET DYNAMICS

As discussed in the previous section there are many ways to organize addressability and two-way control traffic. As leverage points are established that show signs of a solid return on investment, a dynamic supply and demand scenario is likely to emerge.

Again, we can look to the Internet as an example of the kind of market dynamics that are possible with an effective feedback loop. Internet ads are measured for impact and swapped out quickly if deemed ineffective. In cable it can take weeks to swap out an ad; on the Internet it takes minutes.

Effective and near-immediate feedback on the Internet allows campaigns to be scrutinized, altered and optimized quickly, resulting in campaigns whose effectiveness increases throughout execution of the campaign.

Changes to a campaign on the Internet are often accompanied by a renegotiation of rates between the buyers and sellers of inventory. The near-immediate feedback provides concrete information on which mutually beneficial pricing can be negotiated.

Standards bodies may play a significant role in this area. Automated measuring, modifying, and the re-pricing of ads may all benefit from uniform data formats and APIs.

COMPOSITION

The advertising architecture described in the above example is designed to be flexible in terms of the selection and placement of ads. However it does not address the changing visual composition of television ads.

Digital delivery of compressed video and audio also allows the delivery of supporting data streams. As envisioned when MPEG 2 was standardized, these streams can be used to send enhancements to the standard video display.

The DVB MHP and OpenCable OCAP efforts have focused on new client mechanisms for enhanced ads. The OpenCable initiative has also defined the Enhanced TV or ETV specifications for use on existing, or "legacy", boxes.

There is still a fundamental question of whether viewers will actively engage when presented with interactive options. By comparison the Internet is based on user interaction and regarded as a "lean forward" medium. Entertainment television may require more subtle invitations to "click back".

At an architectural or system level there is little doubt that web-based display technologies based on HTML will eventually prevail. Web-based tools offer all of the interactivity and connectivity required for future content guides and service interfaces.

This is not to say the television will grow into a lean-forward or browse-based medium; instead it is likely that the Internet toolset will be used to support a new usage model – call it the "lean-back Internet".

STANDARDIZATION EFFORTS

Given the range of possibilities for building advanced advertising systems there is an industry need for a flexible architectures, data formats, and component interfaces.

The SCTE Digital Video Subcommittee (DVS) has supported standardization in the area of Digital Program Insertion (DPI) since 1998. Figure 2 is a high-level block

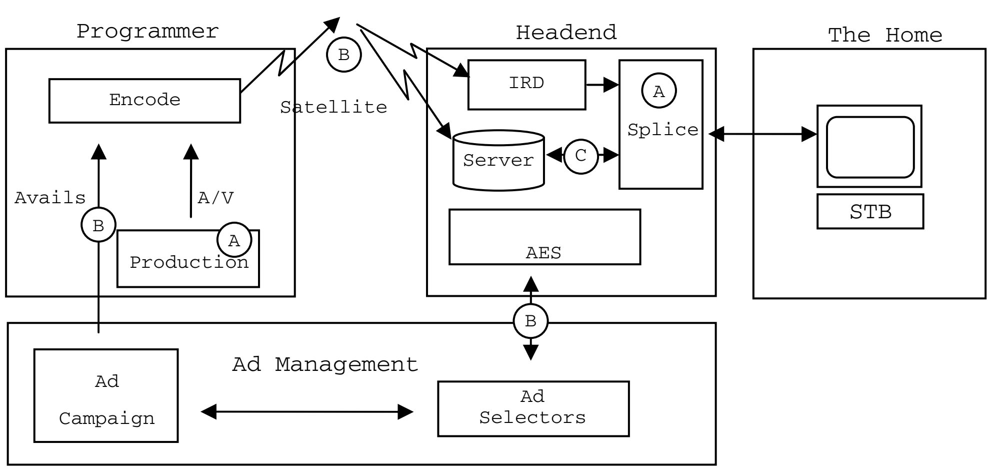

Figure 2. System Block Diagram for Standardization

diagram showing the data and control flows for advanced advertising systems.

The committee's initial work centered on the precise structure needed at the MPEG layer to insure clean splices (labeled A). This was followed closely by a standard (SCTE 35) to identify "splice points" (i.e. avail boundaries) in a network feed (B). A realization followed, however, that digital splicer vendors were emerging separately from ad server vendors; thus, a standard API between a generic video server and splicer would ensure interoperability while allowing competition among splicer vendors (C). With these standards in place, a first wave of digital ad insertion products emerged for inserting ads in digital broadcast feeds.

To accommodate the advertising directions described in this paper new work has begun in a number of key areas. These include:

1) Communication between the AES and ADS components to support dynamic ad decisions.

2) The standardization of metadata in the network feed that identifies the content (called program identification fields). This metadata can be monitored to so that ad placement can rapidly react to schedule changes. For example a sports event might run beyond its scheduled time; with program id metadata in place, additional high-priced sports ads can still be placed.

3) Standards and technologies to support client based advertising (CBA). CBA refers to architectures that rely on the STB to select and present ads that are either already resident on the STB disk or simultaneously available on other frequencies.

Future work may address data formats and APIs for managing the dynamic and automated pricing decisions as well as the ad composition changes discussed above.

CONCLUSION

Cable advertising is changing in fundamental ways. This change will be fueled by opportunities for increased revenue. It will require considerable experimentation at the system level and may benefit from continued standards efforts.

Throughout this period many lessons can be learned from the converging efforts in the field of Internet advertising. In fact it is likely that that end of this era will be marked by a technology convergence with IP and web-based technologies.

BEYOND BANDWIDTH MANAGEMENT: BUSINESS BENEFITS OF SWITCHED DIGITAL VIDEO IN CABLE

Lorenzo Bombelli
Director, Product Strategy and Management
Subscriber Network Systems
Scientific Atlanta, A Cisco Company

Abstract

Fueled by the advent of an open IP-based architecture that has catalyzed the development of reliable, cost-effective, and scalable solutions, Switched Digital Video (SDV) technology promises to fundamentally change how digital video is delivered over cable networks, enabling MSOs to offer consumers a wider variety of programming while effectively managing HFC network bandwidth.

Although bandwidth management is the primary driver for SDV in cable, from a business perspective SDV offers MSOs a number of additional benefits that reach beyond its pivotal role as a bandwidth management tool.

Unlike its telco cousin, the cable version of SDV is designed to operate over existing HFC infrastructure and to enable delivery of switched video services on the existing installed base of some 40 million MPEG set-tops that cannot decode a DOCSIS or IP stream, giving MSOs significant scale advantage in the introduction of new video services. Key switching features, like the ability to share QAM bandwidth between multiple services on a per-stream basis on a single QAM carrier give MSOs the same granularity and flexibility of service delivery to existing MPEG set-top boxes as any competitive service.

In contrast to traditional storage-based video on demand (VOD) technology, SDV is transparent to consumers. By preserving the consumer experience, SDV immediately helps MSOs give consumers more of what they want when they want it without changing the way they get it.

Because SDV fundamentally changes the model for bandwidth consumption from a linear model based on the program offering to one based on program viewership, SDV helps MSOs offer an extensive lineup of niche content ranging from local and other premium sports packages to ethnic programming, thus leveraging the "long tail" phenomenon to improve customer satisfaction, reduce churn, and generate new revenue streams from premium tiers.

A powerful byproduct of SDV systems is that they generate detailed viewership data that give MSOs unprecedented and direct visibility to consumer viewing behaviors, while maintaining the privacy of individual customers. This data can be used to improve program offerings, maximize the return on investment for each program offered, and drive growth in ad revenue by targeting ads more granularly.

Ultimately, switched digital video networks can be configured for switched unicast, enabling MSOs to deliver an individual copy of broadcast content to each consumer complete with highly targeted advertising and to offer features that enhance the viewing experience such as faster channel changes.

Finally, the new IP-based SDV architecture offers MSOs the opportunity to extend the reach of switched digital services to address the explosion of IP-based devices that are capturing an increasing share of video consumption.

OPEN SWITCHED DIGITAL VIDEO ARCHITECTURE

Over the past year, industry leaders have collaborated to develop an open architecture for delivery of SDV to the existing installed base of MPEG-based set-top boxes as highlighted in Figure 1.

Some of the key components and features of the SDV architecture include:

- An **SDV Client** protocol that can be natively integrated into set-top program guides and digital navigators from multiple vendors.

- A **Master Session and Resource Manager (SRM)** that is able to control and arbitrate bandwidth for multiple applications. Initially the master SRM manages the dynamic allocation of bandwidth between switched digital video and existing video on demand services. However, the architecture also allows for additional application servers to share system bandwidth.

-

- A **Switched Digital Video (SDV) Server**, a distributed SRM that is client-aware and constitutes the brains of the switched digital video network. The SDV server is designed to handle a large volume of real-time channel change and other service requests, and for high system availability and redundancy. It keeps track of all user requests and can use the information to improve system performance or to report viewership data and other key system parameters. The switched digital video server can be distributed or centralized, and is built on standards-based server hardware (e.g. Dell) and operating systems (e.g. Linux).

- **Session-Based QAM Modulators** that support IP Multicast (IGMPv3) and interoperability with standard IP switch/switch-routers. The session-based

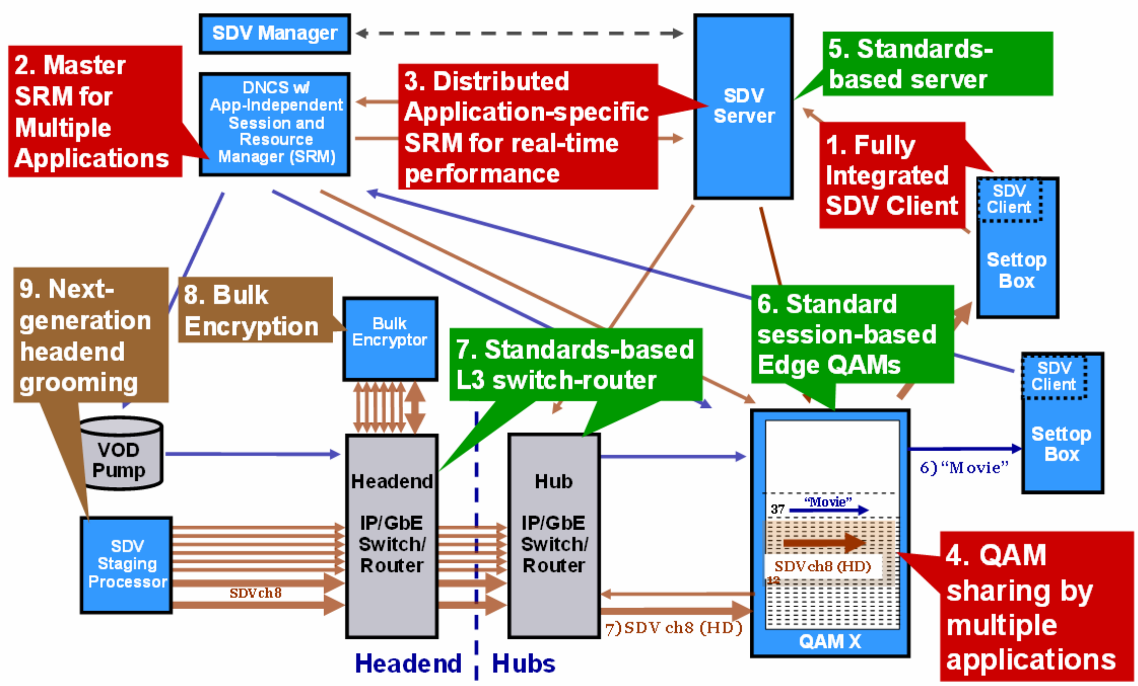

Figure 1. Key Components and Features of the Open Switched Digital Video Architecture

QAMs interface with both the Master SRM and SDV server using open, published interfaces to dynamically deliver switched video services over the HFC network, and enable QAM sharing between SDV and existing video on demand services.

- **Standards-Based Switch-Routers** that provide content distribution over local and regional fiber backbones. The key requirement that the SDV architecture places on the IP-based content distribution network is that it supports IP Multicast (IGMPv3) with Source-Specific Multicast (SSM) in order to provide source-level redundancy.

- Network-attached **Bulk Encryptors** that allow for independent scaling of encryption and edge QAM resources. Like the SDV server, the network-attached bulk encryptor is connected to the IP network over Ethernet (the bulk encryptor has multiple bi-directional Gigabit Ethernet ports) and can be centralized or distributed to achieve optimized cost and performance.

- A network-attached **Staging Processor** for grooming and processing of switched video streams. Initially, the role of the staging processor is to clamp the video streams to a known maximum bit rate to enable deterministic bandwidth allocation in the switched network. However, in the SDV architecture the staging processor is a network-attached appliance that may be able to perform additional functions such as ad insertion, rate shaping, transcoding between different digital video formats and alignment of key MPEG video frames (I-Frames). Although today the cost of MPEG processing is much higher than the cost of switching, and would burden the overall system if applied at the network edge, as higher density, lower cost video processing platforms are developed, the open SDV architecture allows for video processors to be applied seamlessly,- connected by IP across the network and scaled according to business requirements.

- **Standards-Based Reliability and Redundancy Mechanisms** that address both content and signaling networks to provide redundancy and service resiliency. These mechanisms protect against outages in sources (including encoders, staging processors and bulk encryptors), content distribution network, QAM modulators, and SDV servers. The SDV architecture also support service resiliency in the event of an HFC reverse path outage.

- An **SDV Manager** that features standard SNMP MIBs for automated configuration and management of system components from a single user interface.

- Architectural fit with future service convergence architectures such as Packet Cable Multimedia and IP Multimedia Subsystem Architecture.

The referenced white paper, *“An Open Architecture for Switched Digital Services in HFC Networks”* [1], provides a detailed overview of the key elements and operation of a switched digital video system using this open architecture.

The advent of an open SDV architecture has accelerated the development of technologies and solutions for delivery of SDV services over HFC networks. It is more reliable, more cost-effective and more scalable than existing MPEG-based architectures. As such, to MSOs, a switched digital video system is more than just a powerful bandwidth management tool. It is a key stepping stone in the evolution to next generation networks that not only pays for itself by delivering an immediate return on investment based on bandwidth management, but also delivers a number of key benefits that add business value to MSOs both immediately and going forward.

ENABLING EXISTING MPEG SET-TOP BOXES

One of the most important features of the open SDV architecture is that it enables switched delivery of video services over IP network infrastructure to the existing installed base of more than 40 million MPEG set-top boxes. The alternative would be to overlay switched services to IP-based set-tops. An IP overlay would not only fail to deliver immediate bandwidth savings (additional bandwidth would be required to support the overlay), but also to address the delivery of new services to existing digital subscribers without replacing their set-top box and incurring significant capital, operational and opportunity cost.

Interestingly, because the open, IP-based SDV architecture reduces the bandwidth required to deliver both existing and new services to existing MPEG set-tops, it also catalyzes the evolution to all digital and end-to-end IP-based networks by freeing up bandwidth for services that require a video overlay such as digital simulcast and video over DOCSIS.

QAM BANDWIDTH SHARING

As shown in Figure 2 below, the open SDV architecture is designed to enable the coexistence of SDV and existing VOD services, and enables dynamic bandwidth sharing between multiple services sharing the same edge resources.

An application independent SRM enables sharing of edge resources (QAM modulators for an HFC system) among the various interactive applications requiring HFC bandwidth. Such sharing of QAM resources is normally considered to occur in one of two ways: inter-carrier sharing, and intra-carrier sharing.

Inter-carrier QAM sharing, the simplest and most straightforward method of sharing QAM resources among applications, involves the assignment of individual RF carriers to specific applications. Some of the RF carriers within a QAM chassis, for example, may be assigned to VOD while others may be assigned to SDV. Inter-carrier QAM sharing is less dynamic, but is simpler to implement and does not require the intelligence to arbitrate between more than one application demanding share of a given carrier.

Intra-carrier QAM sharing is more complex but more powerful than inter-carrier QAM sharing. It relies on some level of intelligence to dynamically make decisions on the priority of two or more applications requesting access to stream capacity within a QAM carrier. To accomplish this, carriers are designated on the SRM as sharable among more than one application. It is then up to the SRM to assign bandwidth on each carrier upon request based on some set of rules or business intelligence. It may be as simple as prioritizing the applications as a whole, or as complicated as prioritizing individual streams based on the value of content, time of day or any of a number of factors. The requirement for decision-making intelligence, intra-carrier QAM sharing calls for QAM modulators to be entirely session-based.

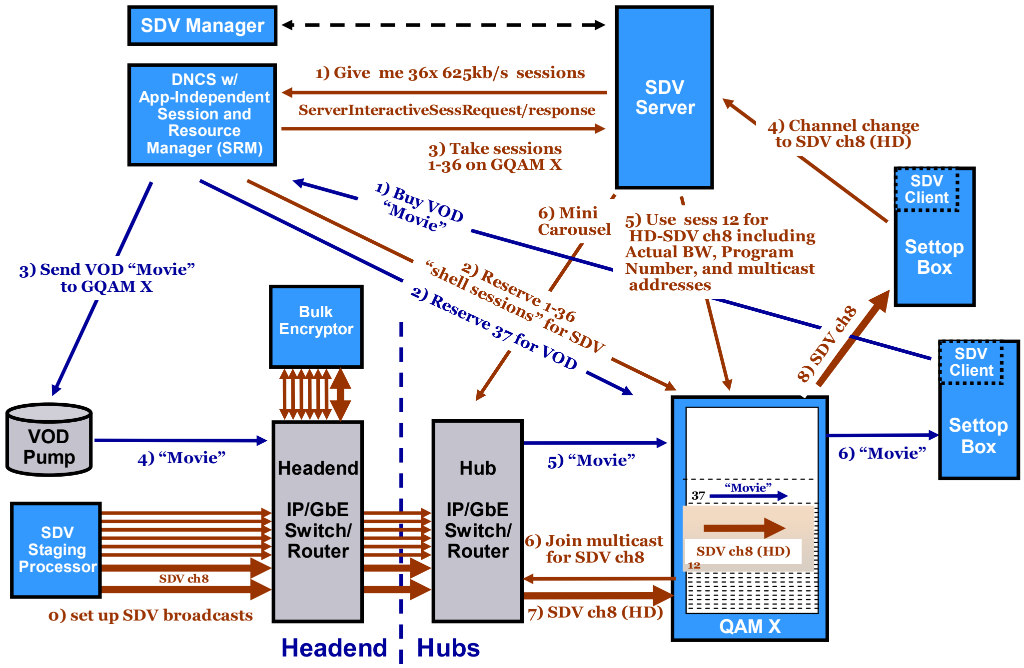

Figure 2. SDV and VOD Coexistence

TRANSPARENCY, CUSTOMIZATION AND "LONG TAIL" CONTENT

Switched digital video extends the capability of MSO networks to enable the delivery of content on demand without changing the average consumer's "lean back" viewing experience.

In order to view video on demand on a traditional storage based VOD system, a consumer searches for a given piece of content, selects it, and then waits approximately 5-10 seconds for the selected content stream to appear on the screen. The VOD experience is very different than typical TV viewing, wherein a consumer tunes to a favorite channel or surfs through channels to see "what's on". He then might stay tuned to that channel as long as it continues to entertain him.

Because SDV is designed to behave, from the consumer's perspective, exactly like broadcast TV, it must deliver the same experience as regular TV viewing. However, because it takes advantage of viewing statistics to deliver video streams only if they are being watched, it enables the delivery of many more channels, allowing MSOs to offer content that provides high value to small groups of users, sometimes referred to as "long tail" content [11].

Opportunities to deliver such programming, ranging from local and other premium sports tiers to international and ethnic programs (as illustrated in Figure 3), is now offered by satellite and can drive growth in new subscribers and improvement in customer retention if offered as part of an MSO's bundle of services.

One of the challenges in delivering these channels over cable networks is that there are pockets of demand for many of these channels in most cities, and certainly in any metropolitan area. As a result, in order to

Figure 3. Examples of International Programming Options

meet the needs of the entire subscriber base, an MSO would have to offer many or even all of the packages shown above. In the traditional linear broadcast delivery model this plan would be impossible. With switched digital video—which capitalizes on geographic diversity, viewing statistics and the efficiency of multicast—however, most MSO networks can deliver most (if not all) of the programs shown above.

VIEWERSHIP DATA AND CLIENT AWARENESS – GETTING CLOSER TO CONSUMERS

A valuable byproduct of switched digital video systems is very detailed viewership data. The SDV system has visibility to every channel change request in real-time, whether it is for a switched, broadcast digital, or analog channel. It also receives and stores non-real-time reports from set-tops on all other remote control activity. Although the viewership and user activity data is primarily used to optimize system performance, the data can be reported from the SDV server in a secure manner that protects consumer privacy.

The enhanced visibility to consumer behavior can be used to continually fine tune and target service offerings, improve marketing campaigns, and increasingly target advertising to drive ad revenue growth.

The SDV server is also aware of the capabilities of each SDV client (set-top or other device) that it serves. This client-awareness combined with rich viewership and activity data can be used for a number of key operational benefits including facilitating the rollout of new set-tops or devices that support advanced video formats, or delivering different streams to different devices depending on the bandwidth or format capabilities of each device.

Combined with powerful network-attached video processing and digital ad insertion, these capabilities may prove to be key drivers of revenue growth for MSOs.

SWITCHED UNICAST

Taken to the limit, targeted advertising leads ultimately to a switched unicast model where an individual copy of each program is delivered to each customer even if the program is viewed by more than one customer.

As with any capability, switched unicast comes at a cost. First, switched unicast by its very nature forfeits much of the efficiency offered by switched multicast, and, as a result, fundamentally costs more than switched multicast. Second, the cost of the video processing technology required for digital program (ad) insertion (DPI) today is significantly higher than the cost of switched video, and makes it prohibitively expensive to deliver individualized ads.

The business model for highly targeted TV ads has yet to be proven and may not become viable for some time. However, as switched video systems give MSOs greater visibility to consumer demand, as video processing technology continues to improve at Moore's Law rates, and as competition for consumer attention intensifies, it is not a stretch to envision a day when highly targeted, if not unicast, streams become a business reality.

Given the current cost of video processing technology and the lack of a proven business model, it may not make sense to broadly deploy switched unicast capability initially. Nevertheless, an open, IP-based SDV architecture allows MSOs to introduce higher performance, more cost-effective, higher density network-attached, next-generation video processors to switched video networks to deliver increasingly customized content.

FAST CHANNEL CHANGE

The SDV architecture offers the opportunity to significantly improve channel change times. There are a number of methods available for speeding up channel change times. One method that is gaining momentum involves delivering a switched unicast stream to each subscriber and starting each stream with an I-Frame, an MPEG frame that can immediately resolved by the receiver. This approach, combined with other techniques, can significantly reduce channel change times and make them faster than they are in digital broadcast networks today.

Like individualized advertising, fast channel change is a powerful tool that is enabled by SDV and is available to MSOs to be used as required by business and competitive needs. Similarly, fast channel change will benefit from performance, cost and density improvements in video processing, and may become more viable with the advent of next-generation network-attached video processors.

A key additional consideration in evaluating the fast channel change feature is that the fast channel change capability will only be available on channels in the switched video tier. As a result, consumers may become less satisfied with their experience in viewing programs that are left in the traditional linear delivery tiers. Preservation of a consistent consumer experience is a key consideration in rolling out any new service or feature. In the case of fast channel change, it is likely to be driven, as are many new features in cable, by competition, and may accelerate a wholesale transition to switching all digital video services in order to preserve a consistent consumer experience.

EVOLUTION TO NEXT GENERATION ARCHITECTURES

As we described above and shown in Figure 4, the SDV architecture described here is an important element of a comprehensive and evolutionary next generation architecture strategy.

First and foremost the SDV architecture described here provides a critical first step from today's linear broadcast delivery model to a non-linear switched delivery model based on IP core infrastructure, and acts as a catalyst for the evolution to all digital and end-to-end IP-based services by freeing up bandwidth in the network.

In addition, because the open SDV architecture is based on end-to-end IP infrastructure using open protocols and interfaces, it enables the delivery of multicast video streams from the content source all the all the way to the IP network edge, where session-based edge QAM modulators output MPEG video over RF to existing digital set-tops.

As such, it offers MSOs the option to distribute existing and new services including video on demand, digital simulcast, and DOCSIS with switched digital video on a converged IP network infrastructure (whether it be converged on a single or multiple wavelengths is up to the MSO).

It also provides an immediate business driver (bandwidth management) to extend the end-point of IP convergence beyond the headend. Traditionally video has been broadcast over 1310 or 1550 nm fiber as an overlay to the IP data backbone. The SDV architecture extends the IP network to the hub where IP services are launched today

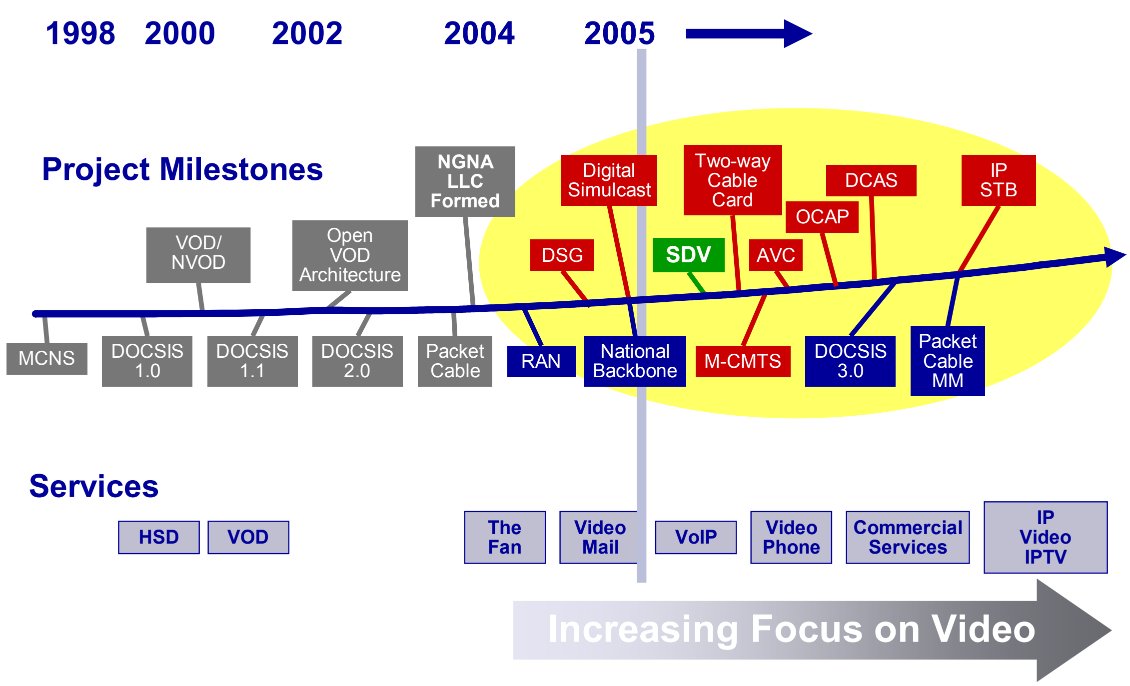

Figure 4. Evolution to Next Generation Architectures

over CMTS, an important step toward end-to-end IP convergence and next generation architecture.

Because one of the main goals of the SDV architecture is to address the existing installed base of MPEG set-tops, it features an out-of-band signaling mechanism that is structurally consistent with the Packet Cable Multimedia (PCMM) and IP Multimedia Subsystem (IMS) architecture models for non-IP clients. However, the SDV architecture can also be evolved to support IP-based clients and, in conjunction with the DOCSIS, PCMM and IMS architecture models, to provide a true multi-service, multi-client architecture that enables any-source to any-user IP-based video delivery.

SUMMARY

In summary, although bandwidth management is the primary driver for switched digital video in cable, the open, IP-based SDV architecture that is being embraced by MSOs offers a number of additional benefits and capabilities, including:

- The delivery of switched services to existing MPEG digital set-tops
- QAM bandwidth sharing between SDV, VOD and future on demand services
- A transparent consumer experience compared with traditional digital broadcast cable TV
- The ability to launch new premium and niche programming tiers
- Rich viewership data to help get closer to customers and better understand their needs
- A business-driven path toward increasing customization of content and advertising
- Powerful competitive features such as fast channel change
- A sure-footed evolutionary path to IP-based next-generation architectures

Thus, switched digital video not only delivers immediate benefits in the form of spectrum savings and increased efficiency, but it also presents great opportunities for sources of competitive advantage going forward.

REFERENCES

1. An Open Architecture for Switched Digital Services in HEF Networks, Luis Rovira, Lorenzo Bombelli, Paul Brooks, January 2006
2. Switched Multicast Services for HFC, Luis Rovira, October, 2005
3. Time Warner Cable Switched Broadcast Channel Change Message Interface Specification TWC-SDB-CCMIS-2_12, October 12, 2005
4. Time Warner Cable Switched Digital Broadcast Mini Carousel Message Interface Specification TWC-SDB-MINI CAROUSELMIS-2_1, October 12, 2005
5. SDB Server/RM – Edge Device Interface 1.2, September 9, 2005
6. Master SRM – Edge Device Interface 0.5, October 19, 2005
7. Session Setup Protocol Extension for Server Initiated Setup, Version 1.8 September 14, 2005

8. ISO 13818 part 6, Digital Storage Media Command and Control (DSM-CC)

9. TWC Pegasus Session Setup Protocol (SSP) Version 2.3

10. ISO 13818 part 1, MPEG Systems Specification

11. The Long Tail, Chris Anderson, Wired Magazine, October, 2004

BROADCAST QUALITY VIDEO OVER IP NETWORKS: CHALLENGES AND SOLUTIONS

Santhana Krishnamachari, John Hartung, Junfeng Bai
EGT, Inc

Abstract

Deploying video services over IP networks is a significant hurdle because of the stringent requirements on bandwidth, loss rate, and delay jitter imposed by video traffic. Today, most of the video traffic over the Internet and private networks is in the form of data downloads in non real-time, or real-time transmission of low resolution video. Scaling this to a full-resolution high quality video and to a large number of flows continues to be a significant challenge. In this technical case study we present the results of experiments that we have performed to study the feasibility of transmitting broadcast quality video over the Internet. Several point to point communication links were set up with nodes at various geographical locations. Statistics pertaining to video transport, namely available bandwidth, delay jitter, and loss rate were collected over long periods of time. These statistics are presented in this paper and they show that although the network paths are reasonably well behaved for most of the time, intermittent variation in characteristics will require additional techniques such as passive measurement, path diversity, rate adaptation, and error recovery to achieve robust reception.

1. INTRODUCTION

The main motivators for distributing video content over IP networks for a multi-channel operator are : (i) Substantial cost reduction in dollars per bits transmitted compared to other modes of transmission, e.g., over a dedicated satellite link. (ii) Flexibility to converge multiple services such as high speed data, voice and video over a single network. This convergence offers better statistical utilization of the channel resource and allows the amortization of the cost of multiplexing, transport gear over different services. (iii) Deployment of switched digital video broadcast over the last mile is expected to increase consumer interest in niche, on-demand content. Bringing such niche content to the headend can be achieved in a cost effective manner over an IP network or the Internet.

The Internet in its current form and most private networks cannot guarantee the quality of service required for broadcast quality video. Over the last two years, we have been witnessing an increasing trend of using IP networks for transport of audio and voice for commercial purposes. Rhapsody and a multitude of radio stations on the Internet stream audio services. VoIP providers like Vonage and cable MSOs use the public or private network for voice traffic. Audio streaming can tolerate longer network delays that results in a start-up latency. VoIP has stringent limits on one way delay and loss rate so that the interactivity of conversation is not impaired. The expected evolution of using IP networks for video transport is significantly more difficult because of the higher bandwidth requirements, in addition to the restrictions on delay jitter and loss rate.

Prior experiments performed to measure internet characteristics show a wide variation in loss rate. It was reported in [1] that the average loss rate was a very low 0.42 %, but during the worst one-hour period the average loss rate was over 13%. Published reports in [2] indicate that the convergence of inter-domain routes can take upto tens of minutes in case of a link failure. Although the performance of private IP networks can be drastically improved by traffic engineering, as the loading on these network increases the performance characteristics will be similar to the Internet.

In addition to the loss rate characteristics discussed above, variations in delay jitter and available bandwidth in the network significantly affect the robustness of video transmission. In this paper we present the results of the experiment we have performed to measure the loss rate, delay jitter and available bandwidth of network links. Similar to the reported results on loss rate, other network characteristics also exhibit large variations over short time periods.

The paper is organized as follows: Section 2 presents the results of our measurements. Section 3 presents details on existing technologies that can be used to improve the quality of service and their limitations. Section 4 presents the additional supporting technologies that are required to achieve robust video transmission and concludes the paper.

2. NETWORK MEASUREMENTS

In this case study, we deployed 6 nodes globally to measure the statistics of the Internet links between different locations. The geographical locations of these nodes are illustrated in Figure-1. We collected the available bandwidth, delay jitter and loss rate statistics for all the links over a period of about one month. In this paper, we will show some of measurement results for the link from Gsnet.ch to EGT at Atlanta. Other links exhibit very similar characteristics.

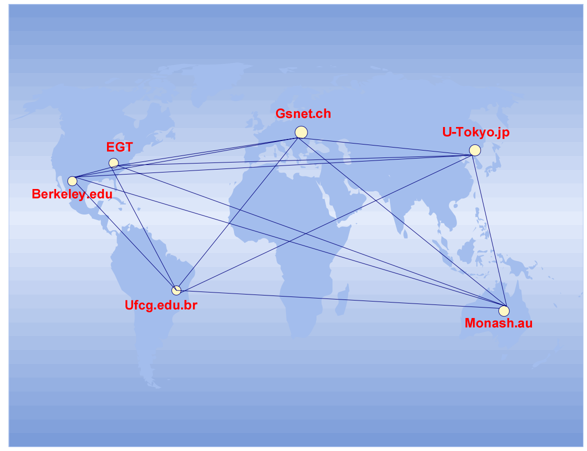

Figure 1. Geographical locations of nodes used for measurement

Packet Loss Rate Measurement

Figure-2 shows the packet loss rate experienced by a 1Mbps video stream from Gsnet.ch to EGT. The time duration of this experiment is from a Friday afternoon to the following Monday afternoon. The packet loss rate is collected every 1 second. Figure-3 shows the corresponding histogram for the loss rate.

The histogram shows that about 90% of the time the packet loss rate is less than or equal to 1%. Such small packet losses can be tolerated by video decoders by using advanced error concealment techniques. Forward Error Correction (such as Pro-MPEG [3]) techniques with about 5~10% overhead can also be used to counter small packet losses.

Figure-2 shows that the packet loss is relatively small during the weekend. For the period corresponding to Monday morning (between 65-75 hours in the plot) there are large bursty packet losses that continue for several hours. In such bursty loss cases, FEC or concealment techniques cannot recover all the lost packets.

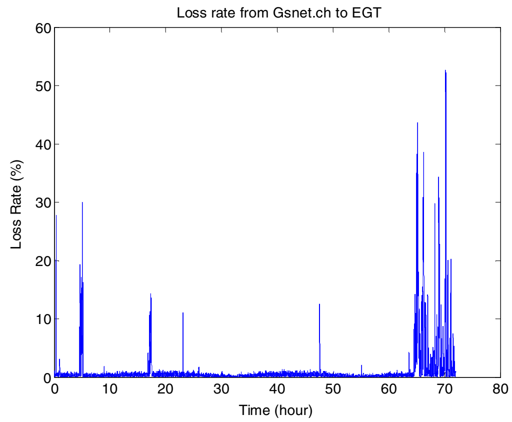

Figure 2. Packet loss rate statistics

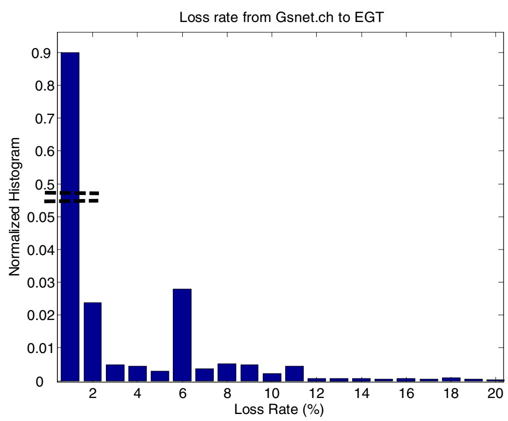

Figure 3. Histogram of packet loss rate

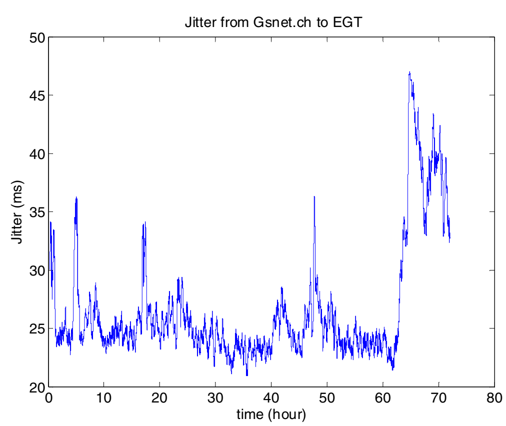

Figure 4. Delay jitter statistics

Delay Jitter Measurement

Figure-4 shows the jitter over same time duration and Figure-5 shows the corresponding histogram for the jitter. We measure the jitter at the receiving side with the following equation

Jitter = [Tr(i)-Tr(i-1)] – [Ts(i) – Ts(i-1)]

Here, Ts(i) is the sending time of the *i*th packet at the sending side and it is inserted into the packet header as a timestamp. Tr(i) is the time when receiving side receives the *i*th packet.

Each packet may experience different transit delay because of different queueing delay introduced by routers along the path from source to destination. We collect the maximum of absolute value of jitter for every 100 received packets and show the result in Figure-4. This maximum jitter is critical for determining the size of dejittering buffer at the receiving side to correctly recover time interval between packets while preventing the buffer from overflowing or underflowing.

It can be observed from Figure-2 and Figure-4 that as the network loss rate increases, the delay jitter also gets worse as expected.

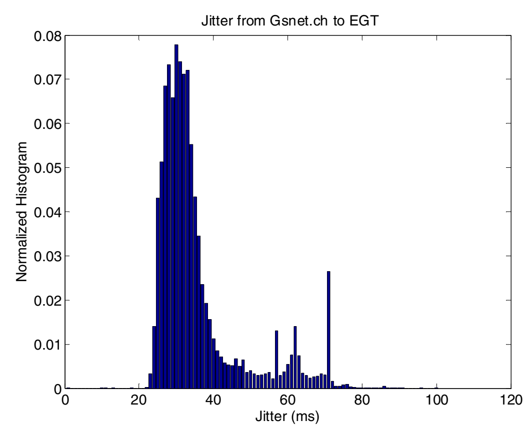

Figure 5. Histogram of delay jitter

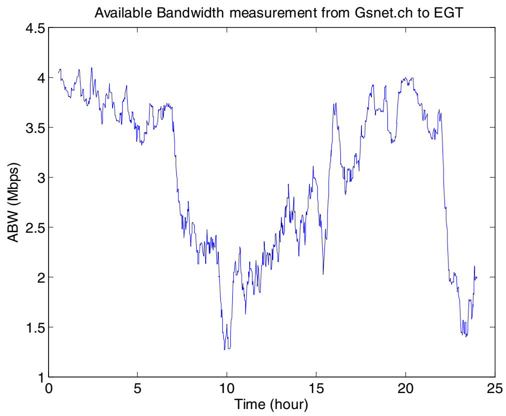

Figure 6. Available bandwidth statistics

Available Bandwidth Measurement

The available bandwidth estimation tool Pathload [4] was used to collect the bandwidth statistics. The result of available bandwidth measurement for 24 hours between Gsnet.ch to EGT is shown in Figure-6. The available bandwidth changes dramatically in the range of 1.5 – 4 Mbps. If the video rate is higher than the available bandwidth, this will result in increased congestion in the network links resulting in large packet losses and delay jitter. This plot clearly shows the need for adapting the video bitrate to match the available bandwidth.

3. EXISTING TECHNOLOGIES

Packet switched networks were originally developed to provide best effort delivery of data and achieve high efficiency through statistical multiplexing. However, the stochastic nature of traffic through these networks leads to unavoidable congestion at switches that is inconsistent with the requirements of high quality video delivery [5]. Two approaches have been taken to reduce or mitigate the effects of this congestion. The first accepts the performance of the network and attempts to provide acceptable QoS for video by compensating for network characteristics at the ingress and egress of the network. In the second approach new protocols have been standardized to allow prioritization of video flows through the switching elements of the network, thereby reducing the probability of congestion for that data. The following two sections give an overview of these two approaches.

Endpoint QoS

The available bandwidth in a default internet route varies widely, and the effect of exceeding this bandwidth is increased packet loss and a rapid degradation of video quality. One solution to this problem is source rate control where the sender adapts its rate to match the available bandwidth. This assumes that there is an accurate measurement of the bandwidth. This technique works for point to point transmission, however, for multicast streams the available bandwidth will usually vary for each endpoint. In this case receiver rate control is used to adapt the rate as a function of the available bandwidth to each receiver. This is implemented by encoding and packetizing the media in multiple layers, with a base layer providing the minimum acceptable quality. The receiver adjusts its rate by connecting to one or more layers whose sum satisfies the bandwidth constraint [6].

The transmission latency of packets between two endpoints varies when congestion occurs due to changing queueing delays in routers along the path. This delay jitter can lead to jerkiness in the playback and packet losses when a packet is delayed beyond its presentation deadline. The introduction of a playout buffer is used to relax the timing constraint, however, this leads to a delay in playout that can be unacceptable when a new stream is started.

Commercial streaming players like the Microsoft Windows player and the Real Network Real player typically introduce 5-15 seconds of delay. An alternative approach that minimizes latency and startup is adaptive media playout (AMP) [7]. In this scheme the rate at which the decoder buffer is emptied, and the media is presented, is varied in order to avoid losses due to missed presentation deadlines. This can combined with retransmission to avoid packet losses as described in the next section.

Packet losses can be dealt with using channel coding to recover from losses and error concealment and resilience to minimize its effect. There are two basic channel coding techniques, retransmission and forward error correction (FEC). Retransmission consists of detecting the lost packet at the receiver and signaling the sender of the loss. A minimum delay of one round trip time is incurred in addition to the time needed at the receiver to detect the loss. This technique has the advantage of using additional bandwidth only when losses occur, however, it requires a back channel that may not be available in applications such as multicast. Alternatively lost packets can be recovered without a back channel using forward error correction. This is accomplished by interleaving a group of packets and adding an FEC code to each. The FEC from the group can be used to recover a number of lost packets within the group. FEC has a disadvantage in that it incurs a rate increase due to the addition of the code words and a delay due to the interleaving of multiple packets even when there are no packet losses.

Error concealment makes use of spatial and temporal correlation to recover lost video information caused by packet loss. These techniques are not standardized and many techniques have been developed making use of motion and spatial information to improve the estimate. Error resilience attempts to encode and packetize the encoded video bitstream in order to minimize the effect of synchronization loss and error propagation. In general, a single bit error can prevent decoding of a video stream, due to variable length codewords, until the next synchronization word. Techniques such as application level framing (ALF) [8] allow the bitstream to be packetized so that each packet is independently decodable. This prevents individual packet losses from propagating errors to the following packet.

The above techniques can be combined to take advantage of the fact that not all bits in the video stream are of equal importance. For example, ALF can be used to form two packet types; one containing high importance data such as I-frames, and another containing lower importance B and P-frames. Unequal loss probabilities, and transmission cost, can be obtained for the two types by applying unequal error protection. This combined source and channel coding achieves lower distortion at an equivalent transmission rate as compared to a system using equal protection for all video bits.

Network QoS

Several types of network protocols have been standardized to enable end to end QoS capabilities in large scale networks (internet). An early protocol to support the requirements of individual flows was called intserv. Intserv works with a resource reservation protocol (RSVP) to set up a path through the network meeting the flow requirements. However, it scales poorly and two types of service aggregation protocols have been subsequently developed. The first type, called diffserv, provides a mechanism to label packets according to their required class of service (COS). The diffserv protocol uses an IP header field to label packets according to their transport requirements so that routers can apply different forwarding algorithms to meet those requirements. A second protocol that serves a similar purpose is multiprotocol label switching (MPLS). In addition to specifying the COS for the packet, this protocol also specifies the forwarding path.The forwarding path is set up in advance using a label distribution protocol (LDP). The primary difference between diffserv and MPLS is that diffserv uses the default routing (e.g. open shortest path first (OSPF)), while MPLS enables engineered routes to be specified.

These protocols have several shortcomings for both private networks and the public internet. The first is that traffic engineering is needed to allocate sufficient bandwidth for the aggregate traffic in each link. Because of the stochastic nature of the flows, however, large over-provisioning is needed. The second problem is that QoS guarantees are only possible if all routers in the network implement the protocols. It is possible to build private networks with these capabilities, however, due to the heterogeneity of the internet, this is unlikely to be supported for many years. In addition bridging protocols are needed to maintain the COS labeling across domains such as different internet service providers (ISPs) and types of facilities (e.g. DWDM and ethernet).

4. NEEDED SUPPORTING TECHNOLOGIES

As presented in the previous section, end-point and network QoS techniques offer improved performance, but still cannot deliver guaranteed resiliency against varying network characteristics. The following techniques or combination of technologies are required to achieve robust video transmission over IP networks.

Estimation of network characteristics: Robust measurement of network parameters is essential to proactively use the techniques discussed in Section 3. It should be clear from the plots of Section 2 that detecting and responding to variation in network characteristics can lead to frequent and in some cases long interruption in service. It is necessary to ensure that the measurement techniques have relatively short time constants to respond to short term variations and in addition be able predict the variations in network statistics as well.

Passive measurements: The measurement of network statistics needs to be achieved in a passive mode where no additional probing data is required for measurement. As the loading on the network links increases, the additional load introduced by the active measurement techniques should be minimized.

Path diversity: To avoid interruptions caused by catastrophic link failure it is necessary to introduce route diversity for video transmission that can be implemented in a fashion scalable to large number of streams. One way of achieving route diversity is to combine source coding technique such as multiple description coding (MDC) and use MPLS COS labeling and forwarding to ensure the transport of the two different descriptions is over paths without common links. MDC is an encoding technique where the video is coded into two or more streams each of which is independently decodable. Jointly decoding multiple descriptions increases the quality of the received video.

In summary, we have presented the results of network measurements between several nodes distributed over diverse geographical locations. These measurements clearly indicate that the characteristics essential for video transport vary widely over time. Two different existing approaches, namely the end-point QoS and network QoS, to improve the quality of video transport over IP networks were presented. Finally, a set of new techniques pertaining to measurement and path diversity that are required to provide broadcast quality distribution were presented.

REFERENCES

[1] D.G. Andersen, A.C. Snoren and H. Balakrishnan "Best Path vs multi-path Overlay Routing", in *Proc of IMC* 2003, Oct 2003.

[2] C. Labovitz, A. Ahuja, and F. Jahanian, "Delayed Internet Convergence", in *Proc. Of SIGCOMM* 2000, Aug 2000.

[3] Pro-MPEG Forum code of Practice #3 - release 2: Transmission of Professional MPEG-2 transport streams over IP networks. http://www.pro-mpeg.org

[4] M. Jain and C. Dovrolis, "Pathload: A Measurement Tool for End-to-End Available Bandwidth", *Proc of PAM* 2002, 2002

[5] Ion Stoica, "Stateless Core: A Scalable Approach for Quality of Service in the Internet", *Ph.D. thesis, Carnegie Mellon Univ., Pittsburgh*, PA, Dec. 15 2000.

[6] S. McCanne, V. Jacobsen, and M. Vetterli, "Receiver-driven layered multicast", *ACM SIGCOMM,* Aug. 1996.

[7] M. Kalman, E. Steinbach, and B. Girod, "Adaptive Media Playout for Low Delay Video Streaming over Error-Prone Channels", *IEEE Transactions on Circuits and Systems for Video Technology,* June 2004.

[8] D. Clark and D. Tennenhouse, "Architectural Considerations for a New Generation of Protocols", *ACM SIGCOMM,* September 1990.

Contact Information:
Santhana Krishnamachari
EGT
75 5th Street, NW
Atlanta GA 30308
santhana@egtinc.com

DIVERSITY VIA CODE TRANSFORMATIONS: A SOLUTION FOR NGNA RENEWABLE SECURITY

Yongxin Zhou, Alec Main
Cloakware Inc

Abstract

The Next Generation Network Architecture (NGNA) has been designed to dynamically download and renew secure software components. Such renewability has reduced effectiveness without code diversity, which is the creation of semantically equivalent yet structurally different code components.

A promising solution for diversity is code transformation technology, which provides a large number of highly obfuscated, different representations of the same digital system. Such transformations provide protection of secret information and operations and effectively hide small updates, while preventing automated attacks. This forces attackers into time-consuming manual reverse engineering to break each diverse instance. When components are renewed faster than the time it takes to break an instance, then a truly secure and flexible system has been created.

We outline a mathematical foundation of generating code transformations based upon rich algebra systems, including finite modular ring, finite Galois ring, 2-adic algebra, and Boolean algebra. Then we give some functional properties of Boolean Arithmetic (BA) algebras and show that Mixed Boolean and Arithmetic (MBA) transforms provide efficient and secure methods for data and operation hiding. Lastly, we provide concrete examples of data transforms, operation transforms, and their compositions to generate diversity.

INTRODUCTION

The Need for Renewable and Diverse Software

Software – although often considered the weakest link when creating a secure content protection system – is unavoidable in content protection devices and systems of the future. Ever-more powerful processors and the supply of free, powerful operating systems means that designers will continue to leverage the flexibility of software to future-proof designs and devices while standards change and markets evolve.

The notion that software is insecure is also under debate – especially for network connected devices. IPTV networks have largely chosen lower cost, software-renewable-based security over traditional smart cards. DirecTV, the largest satellite digital broadcaster, is also fighting piracy via software updates. Of course, their system would have been more effective if it had been designed for such countermeasures from the beginning.

Most security experts now accept that a defense in depth is required. You cannot build one, super strong defense and expect it to withstand attack forever. Technology changes and defenses will fail. Similarly, the reality is that neither hardware nor software is immune to being hacked. The more important questions are:

- How do you recover from an attack?
- How do you mitigate the impact of a successful attack?

The simple answer is to renew the security. Since hardware is by nature more costly to replace, this ultimately means renewing the software. Microsoft's Windows Media Digital Rights Management (WMDRM) has remained hack-free for a number of years, largely due to Microsoft's software renewability strategy. While smart cards are also renewable, the cost has proven unacceptable and arguably their form factor has made it easier for pirates to distribute their product.

Downloadable and renewable software is a key component of the Next Generation Network Architecture (NGNA) [1] project and its new Downloadable Conditional Access System (DCAS) [2]. A secure yet generic microprocessor combined with renewable software allows for future flexibility and portability across Multiple Service Operator (MSO) networks, while improving security.

However, the software must also change when it is renewed. This is the basic premise behind software aging to prevent piracy of consumer software [Anck] and DirecTV's approach toward piracy. But the time between updates and the nature of the changes are also critical factors: i) if it takes months between software updates, then the hacker (and his millions of closest friends on the Internet) will have free content until the next update; and ii) if the change to the security is minimal, the hacker will perform a differential analysis on the new software and quickly isolate the changes and re-hack the system.

Software diversity is the creation of semantically equivalent yet structurally diverse software – the key to solving piracy. Software diversity can be deployed across an installed base (spatial diversity) to mitigate the effect of a successful attack to a subset of the users. In addition, software diversity can be deployed over time (temporal diversity). Software obfuscation and diversity is vital to hiding small changes between software releases and preventing differential analysis attacks against the software.

The resistance of a system against attack is a function of both initial resistance and diversity (Figure 1). While well-designed hardware provides higher initial resistance, it typically lacks diversity. Software has lower initial resistance but can easily be diversified and renewed at low cost [Main].

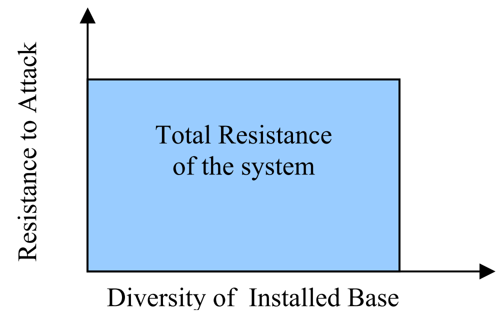

Figure 1. Resistance of a System vs. Diversity

Another effect to consider is that the incentive of a hacker to attack a system (e.g. fame, monetary gain, impact) increases as diversity decreases (Figure 2). This is the corollary of the Microsoft "monoculture" described by some security experts [Greer]. It is also the effect of mass produced hardware security (e.g. smartcards, Xbox and PSP).

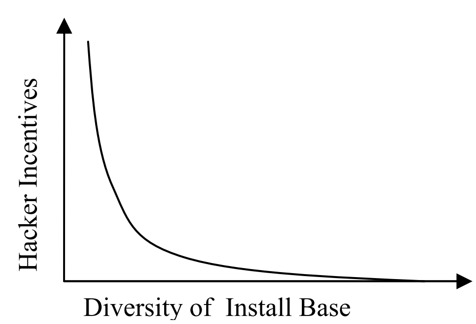

Figure 2. Hacker Incentive

Automated Diversity

Code transformation is an automated approach to generating a large number of highly diverse and obfuscated software instances that is fundamental to any renewable software security.

Based on computational complexity theory, such transformations provide protection of secret information and operations and hide small code changes between updates, while preventing automated attacks via diversity. This forces the attacker into time-consuming manual reverse engineering to break each diverse instance. When components are renewed faster than the time it takes to break the instance, then a truly secure and flexible system has been created [Schn]. Clearly, automated techniques are required to allow for fast updates when necessary.

In this paper, first we build our mathematical foundation for generation of code transformations over existing electrical architectures, which inherit rich algebra systems, including finite modular ring, finite Galois ring, 2-adic algebra, and Boolean algebra.

Then, by selecting the closed algebra systems, we define Boolean Arithmetic (BA) algebras, which applies to real machine computations. We then show that Mixed Boolean Arithmetic (MBA) functions defined over these BA algebra provide an efficient and secure method for data and operation hiding.

Thirdly, we provide concrete examples of code transformations using MBA functions applied to data transforms, operation transforms, and their compositions to obfuscate some common operations, such as addition.

Finally, we briefly discuss general methods of attacking transformed code and counterattacks.

Note that since crypto software components in Conditional Access and Digital Rights Management (DRM) systems are related to integer computations only, we will not discuss floating-point computations. Floating point does not preserve the algebraic structure discussed and more hybrid and practical approaches are required.

THE MATHEMATICAL FOUNDATION

Basic Binary Data Operations in Processors

Since binary code contains all functionality of software and is what is exposed to the hostile environment, it is natural to address software security from this level of code. Almost all processors, both real time and general purpose, have a set of common binary integer instructions. This set includes basic arithmetic operations, bitwise operations, comparison operations, branch operations, and memory access operations.

Arithmetic operations are addition +, subtraction -, multiplication *, and division /. Bitwise operations are and &, or |, exclusive or ^, not ~, bit shift right >> and left <<. Comparison operations include greater than >, greater than or equal >=, less than <, less or equal than <=, not equal !=, and equal ==. In most systems, comparisons are built together with branch operations. Since memory access operations involve no computation, they can simply be regarded as value assignments, which can impact the resulting code, but does not impact the mathematical theory.

Introduction of Boolean Aritmetic Algebra

Software is composed of those basic operations computing on binary data. The most common data format used in all operation systems is two's complement format [Dornhoff].

In this paper, we assume all data has a fixed bit size n. We use B^n to represent all n-bit binary data, where B = {0, 1}. For example, B^{32} is the set of all 32-bit binary values. To simplify the discussion, we assume the address space of the digital system is also n-bit, which is the case of real machines with a 32-bit address space.

We have defined a new term called Boolean Arithmetic (BA) algebra to encapsulate all concepts above. BA algebra is defined on set B^n with arithmetic, bitwise, and comparison operations, thus, the algebra is defined as (B^n, +, -, *, &, |, ^, ~, >>, <<, >, >=, <, <=, ==, !=). For simplicity, hereafter, we use B^n to represent this algebra system.

Digital machines work on BA algebras. That is, a binary program can be regarded as a mathematical function from domain B^n X B^n X … X B^n to co-domain B^n. A large number of diversified code formats can be generated for a given BA algebra function.

Algebra Systems over a BA Algebra

For a BA algebra, B^n, there are several known algebraic systems as its subsystems. Here we mention some of them, which are useful for our code transformations:

(B^n, &, |, ~) is a Boolean algebra structure. It is isomorphic to the one-bit Boolean algebra and has all identities of logical relationships [Dornhoff].

(B^n, ^, &) is a Boolean ring [Dornhoff]. (B^n, +, *) is isomorphic to the modular integer ring $Z/(2^n)$, where Z is the integer ring. Since this is two's complement format, then we define the Abelian group (B^n, +). The modular ring (B^n, +, *) can also be regarded as a Galois ring $GR(2^{s,},2^{st})$, where s and t are integers such that s*t=n [Wan].

(B^n, +, *, /) can be interpreted as a truncation of the 2-adic ring Z_2.

Over the binary data values B^n, we can define a finite field $GF(2^n)$ = (B^n, ^, #) algebra system, where # is the multiplication of the field [Lidl].

Over rings $Z/(2^n)$, $GF(2^n)$ and $GR(2^{s,},2^{st})$, we have polynomial rings $Z/(2^n)[x]$, $GF(2^n)[x]$, and $GR(2^{s,},2^{st})[x]$, as well as multivariate polynomial rings [Jacobson]. More algebra structures can be defined over those rings, such as group rings $Z/(2^n)G$, $GF(2^n)G$ and $GR(2^{s,},2^{st})G$, and loop rings $Z/(2^n)L$, $GF(2^n)L$ and $GR(2^{s,},2^{st})L$, where G is a group and L is a loop [Passman] [Goodaire]. Matrix rings can also be defined over all these rings as can many more algebra systems [Jacobson].

In summary, we demonstrated that over the set B^n, a large number of algebra systems is defined and can be represented by operations in BA algebra. This fact is the theoretical foundation for us to create diversified code. This is a huge area to explore further and apply to software, but in the remainder of this paper, we restrict ourselves to the algebra properties from Boolean (B^n, &, |, ^, ~) and modular ring (B^n, +, *) and use them to create examples of code transformations.

CODE TRANSFORMATIONS

Transforms for Operands and Operators

Software is defined by its operations and orders of the execution of those operations. Code transformations are transforms of operands, called data transforms; and transforms of operators, called operation transforms; and rearrangements of the execution order, most likely related to different basic blocks, called control flow transforms.

This paper focuses on the first two transforms and does not discuss control flow transforms. Specifically, we explore protection and diversity by applying bitwise and arithmetic operations of a BA algebra.

Data Transforms

To generate functionally equivalent code, the requirement of invertibility of a data transform is necessary. Any one-to-one mapping f(x) from B^n to B^n can be used as data transforms. For simplicity, we will not consider multivariate cases, such as matrix transforms, in this paper.

To illustrate data transforms, let's start with linear transforms. Suppose a is an odd number and b is any integer of $Z/(2^n)$. Define a linear transform f(x) = a*x + b. Then f(x) is an invertible data transform with inverse $g(x) = a^{-1}*x + (-a^{-1}*b)$. For instance, for 32-bit words:

f(x) = 674529845*x + 944280100

is a data transform with inverse function

g(x) = 998260765 *x + 490631660

It is easy to verify that f(g(x))=x and gf(x)=x, for any x in $Z/(2^n)$.

We also have invertible quadratic polynomial transforms over $Z/(2^n)$. Let a, b, c be three integers and define function $f(x) = a*x^2 + b*x + c$. f(x) is invertible if a is even and b is odd [Rivest]. To have an invertible quadratic function $g(x) = u*x^2 + v*x + w$ with integer u, v, w, solve equation over $Z/(2^n)$ based on the condition g(f(x))=x. It gives us a formula for u, v and w:

$u = -ab^{-3}$,
$v = 2a*b^{-3}*c+b^{-1}$,
$w = -b^{-1}*c - a*b^{-3}*c^2$,

if we assume the coefficient a of f(x) satisfies the condition $2*a^2 = 0$. For 32-bit words, a concrete example is:

$f(x) = 1502216192*x^2 + 3387143129*x + 1221118466$

with inverse function:

$g(x) = 113639424*x^2 + 841194601*x + 3173903662$.

Again, it is easy to verify f(g(x)) = g(f(x)) = x , for any x in $Z/(2^n)$.

In general, for any invertible polynomial function f(x) of degree n, we can construct its inverse in polynomial format with the given degree n under certain conditions of coefficients of f(x).

Other than polynomial functions, there exist enormous invertible functions with MBA expressions over BA algebra, as shown in [Klimov]. All those transforms can be used as data transforms.

Operation Transforms

Operation transforms replace one operation or multiple operations by some other operations. For example, over BA algebra B^n, addition x + y can be replaced by expression x^y + 2*(x&y), because they are equal over B^n. This is an example of MBA expression, and there are other interesting MBA identities that can also be used, some of which have been used in hardware implementation, such as x^y = (x|y)–(x&y). [Warren]

Such identities are quite useful for operation transforms, and we think they deserve a new definition. Over B^n, identities of the format $\sum a_i*e_i = 0$, where a_i are non-zero constant integers and e_i are bitwise expressions of certain variables, are linear MBA (Mixed Boolean and Arithmetic) identities. Linear MBA identities can be generated as follows.

For a fixed number of variables, if the columns of the truth tables of some Boolean expressions e_i of these variables form a set of linearly dependent vectors over modular ring $Z/(2^n)$ with coefficients a_i, we can construct a linear MBA identity $\sum a_i*e_i = 0$.

For any given singular (0,1)-matrix M, we have a linear MBA identity $\sum a_i*e_i = 0$, where e_i are Boolean expressions whose truth tables are columns of M, and vector a_i is the non-zero solution from the linear system M*X = 0 over $Z/(2^n)$.

Since there exists a large number of singular (0,1)-matrices, the number of linear MBA identities is enormous. Following methods mentioned above, we can also show that any bitwise expressions can be replaced by a nontrivial linear MBA expression from a linear MBA identity.

For two variables x, y from B^n, here we list more linear MBA identities:

x + y = (x^(~y)) + 2*(x|y) + 1;
x | y = x + y + 1 + ((-x-1)|(-y-1));
x & y = ((~x)|y) + x + 1;
x ^ y = x – y - 2*(x |~ y) - 2.

We encourage readers to write code to verify an interesting non-linear multiplication MBA identity:

x*y = (x&y)*(x|y) + (x&~y)*(~x&y).

EXAMPLES OF CODE TRANSFORMATIONS

Compositions for Arithmetic and Bitwise Operations

By using functional compositions of data and operation transforms, a given function can be reformatted into a new one with the same functionality. Because of the large number of diversified data and operation transforms, diversified code of different formats can be created.

In this section, we give examples of transformed code for some basic operations of a BA algebra. These examples are over BA algebra system B^{32} of 32-bit data and operations. They are generated based on the data transform f(x) = a*x + b and some linear MBA identities of two variables. Some random coefficients are assigned for integers a and b which are either split or folded with other constants in the expression. All code can be easily verified using computer programs, while noting that signed and unsigned data types make no difference in binary code.

In the following examples, we assume x and y are two input operands and z is the output operand of the given operation. Other variables are intermediate ones with the same variable size.

ADDITION, z is x + y:

t1 = (4211719010 ^ 2937410391 * x) + 2 * (2937410391 * x | 83248285) + 4064867995;

t2 = (2937410391 * x | 3393925841) + 638264265 * y - ((2937410391 * x) & 901041454);

z = 519915623 * t1 - ((3383387769 * t2 + 129219187) ^ 2756971371) – 2 * ((911579527 * t2 + 4165748108) | 2756971371) + 4137204492;

SUBTRACTION, z is x - y:

t1 = (268586306 ^ (904621911 * x)) + 2 * ((904621911*x) |4026380989) + 3383600763;

t2 = (904621911 * x | 898293889) + 961858761 * y - ((904621911 * x) & 3396673406);

z = 2385439847 * t1 -((673378951 * t2 + 2646655957) ^ 4036393323) -2 * ((3621588345 * t2 + 1648311338) | 4036393323) + 2704000620;

OR, z is x | y:

t1 = (223550072 ^ 1783698419 * x) + 2 * (1783698419 * x |4071417223) + 865773809;

t2 = (1783698419 * x | 3200160250) + 3498694157*y - ((1783698419 * x) & 1094807045);

z = 1905442107 * t1 -((94202053 * t2 + 1827854967) ^ 1206196916) + ((2389525189 * t1 + 1837740140) | (4200765243 * t2 + 2621826401)) -2 * ((4200765243 * t2 + 2467112328) | 1206196916) + 3508709997;

AND, z is x & y:

t1 = (543752565 ^ 3040826005 * x) + 2 * (3040826005 *x|3751214730) + 3865950549;

t2 = (3040826005 * x | 3870539833) + 1950617889*y - ((3040826005 * x) & 424427462);

z = ((830084931 * t1 + 2725441253) & (2151370015 * t2 + 843551768)) + 3464882365 * t1 - ((2143597281 * t2 + 1431133401) ^ 2824232692) -2 * ((2151370015 * t2 + 2863833894) | 2824232692) + 2119073563;

XOR, z is x ^ y:

t1 = (913079019 ^ 864001891 * x) + 2 * (864001891 * x |3381888276) + 1912011113;

t2 = (864001891 * x | 1111987895) + 3067869469 * y - ((864001891 * x) & 3182979400);

z = 4281784907 * t1 - ((722243893 * t2 + 758131618) ^ 2456696338) – 2 * ((3572723403 * t2 + 3536835677) | 2456696338) + 3205073209 + 2 * ((13182389 * t1 + 2473633363) | (3572723403 * t2 + 314825442));

Compositions for Comparison Operations

Comparison operations are used as decision makers in the execution path of the code. For example, in a Conditional Access application, it plays an important role for policy checks.

Here we use equal == comparison as an example to show how to use MBA functions to transform the code. We also apply the comparison results into a computation with linear MBA functions such that inequality of the two values will cause malfunction of the original computations.

Suppose we have two variables p1 and p2. In the following example, if p1 equals p2 the output z is x + y, otherwise it has a very high probability that the value of z is not x + y:

t1 = (224336580 ^ 1271166401 * x) - ((~p2) | (3572723403*x)) + 2 * (1271166401 * x | 4070630715) + (p1 & (3572723403 * x)) + 3753888605 - p1;

t2 = (1570335793 * p1 | 3302094725) + 1182396209 * y - ((1570335793 * p2) & 992872570);

z = 4131953217 * t1 -((687540689 * t2 + 3822681666) ^ 3209267133) – 2 * ((3607426607 * t2 + 472285629) | 3209267133)+ 56754764;

By saying high probability, we recognize that in some special cases (e.g. for x = 0 or -1), then the equation may still compute. These cases can be avoided in practice, or added to create further ambiguity and frustration for the attacker.

Again, interested readers can verify this with a computer program.

<u>Attacks and Counterattacks</u>

In a hostile environment, where the cryptographic keys and code of Conditional Access (CA) systems operate, there exist many possible attacks. Common ones are static and dynamic code analysis, powered with sophisticated debug tools, as well as tampering and emulation type attacks [Oorschot].

Although many attacks exist and more will be developed, we have focused on analysis and tampering attacks designed to either recover secrets or algorithms or change the behavior of the software. These types of attacks (as opposed to emulation type attacks) require reverse engineering (or analysis) of the code.

Reverse engineering attacks fall into two basic types: automated and non-automated. The first one is using general algorithms to attack code, while the second one is the combination of using some algorithms and manual tools to determine the functionality of a given program.

It can been proven that the code recognition problem, that is, to classify different code formats based on their functionalities, is an NP-complete problem. Therefore, it would be difficult, if not impossible, for attackers to find an efficient general algorithm to determine the functionalities of all possible code.

Since creating a general automated attack is difficult, attackers would focus on special properties of the code and tools in an attempt to simplify portions of the code. Taking examples in the previous section that use Boolean arithmetic operations, an attacker could try to use symbolic simplification packages in computer algebra systems, such as commercial products Maple [3], or Mathematica [4]. It is easy to find that those examples cannot be simplified because of the mixture of arithmetic and bitwise operations. These tools would need to be adapted to treat this special situation, but countermeasures need to either apply more algebra structures into the code expression, or to inject more MBA identities into the code. Determining such identities is a hard problem.

Only highly skilled individuals can attempt a manual attack against diversified code generated from code transformations based on different algebraic systems. While they may learn certain techniques when reverse engineering the first instance, the ability to generate specialized helper tools or utilities will be difficult. The variety of algebra systems available is extensive and the composition of such systems result in an extraordinarily high number of combinations. If the only viable attack is to manually reverse engineer the resulting code, then clearly we have achieved our objective.

Furthermore, code obfuscation techniques have been developed to seamlessly inject instances of NP complete problems, such as 3SAT, into transformed code. This makes it impossible to perform manual attacks on individual instances.

CONCLUSION

Rich algebraic structures compatible with digital processors guarantee the existence of a large number of code transforms.

As demonstrated in our examples, code transformations introduce obfuscation and diversity, which are vital for hiding secrets, hiding small code changes and preventing automated analysis which in turn prevents automated attacks against the installed base of devices.

By using Mixed Boolean Arithmetic transformations to data and operators, the resulting diverse code can be renewed and be expected to withstand attack for a reasonable period of time. These transformations are not susceptible to analysis using commercial tools such as Mathematica or Maple.

While the new DCAS system allows for renewable software, this must be diverse to prevent piracy. Hardware does not provide suitable diversity, leaving this role to software. Total system resistance to attack is a function of resistance and diversity. Code transformations provide a practical automated solution for such low-cost, renewable, software security.

ACKNOWLEDGEMENTS

The authors thank Phil Eisen, Song Chen and Heather MacIntosh for their helpful comments and remarks.

END NOTES

[1] Next Generation Network Architecture (NGNA):www.cabledatacomnews.com/ngna /ngnaprimer.html

[2] CableLabs: www.opencable.com/dcas/

[3] Maple: www.maplesoft.com

[4] Mathematica: www.wolfram.com

REFERENCES

[Anck] Anckaert, Bertrand; De Sutter, Bjorn; and De Bosschere, Koen. Software Piracy Prevention through Diversity. In *Proceedings of the 4th ACM Workshop on Digital Rights Management*, 2004.

[Dornhoff] Dornhoff, L.L., Hohn, F.E., Applied Modern Algebra, *Macmillan Publishing Co., Inc.,* 1978.

[Goodaire] Goodaire, E. G., Jespers, E., Polcino Milies, C., Alternative Loop Rings, *North Holland,* 1996.

[Greer] Greer, Dan, et al. Cyber Insecurity: The Cost of Monopoly. September 23, 2003.

[Jacobson] Jacobson, N., Basic Algebra I, *W H Freeman & Co 2nd edition,* 1985.

[Klimov] Klimov, A., Shamir, A., A New Class of Invertible Mappings, in *Cryptographic Hardware and Embedded Systems - CHES 2002, Lecture Notes in Computer Science, Vol 2523, pp 470-483,* 2003.

[Lidl] Lidl, R., Niederreiter, H., Finite Fields, *Second edition, Cambridge University Press*, 1997.

[Main] Main, Alec. Security and savings: Going digital and getting both, NCTA Technical Papers. May 2004.

[Oorschot] Main, A. and van Oorschot, P.C. Software Protection and Application Security: Understanding the Battleground. Springer LNCS. June 2003.

[Passman] Passman, D. S., The Algebraic Structure of Group Rings, *Wiley-Interscience,* 1977.

[Rivest] Rivest, R. L., Permutation Polynomials modulo 2^w, in *Finite Fields and their Applications, Vol7 (2001), pp287-292.*

[Schn] Schneider, F. B., Zhou, L. Implementing Trustworthy Services Using Replicated State Machines, *IEEE Security and Privacy, Vol3 (5), pp 34—43*,2005.

[Wan] Wan, Zhe-Xian, Lectures on Finite Fields and Galois Rings, *World Scientific Pub Co Inc*, 2003.

[Warren] Warren, H. S., Hacker's Delight, *Addison- Wesley*, 2003.

About the Authors

Yongxin Zhou is a senior mathematician at Cloakware. His name is on the author list of two published US patents on software security and 12 research papers on algebra theory. He holds a PhD in Mathematics from Memorial University of Newfoundland, Canada.
(yongxin.zhou@cloakware.com)

Alec Main is Cloakware's Chief Technology Officer. He has published numerous papers and articles on software protection and has spoken at key conferences including The National Show, RSA Conference, Intel Developer Forum, Information Highways, DRM Strategies, Mobile DRM and Certicom-PKCS.
(alec.main@cloakware.com)

ETV – CABLE'S INTERACTIVE PLATFORM FOR LEGACY AND ADVANCED RECEIVERS

Frank Sandoval – Director, OCAP Specifications, CableLabs
David Hooley – CableLabs

Abstract

Enhanced Television (ETV) is a cable defined platform for delivering interactive video applications. Interactive games, game shows, reality programming, awards shows, sports, and advanced advertising are some of the anticipated applications. The platform has been defined by CableLabs within the OpenCable project, with the hearty support of the MSO community.

This paper will describe the technical aspects of the platform, including the application format and on-the-wire signaling. CableLabs has also developed User Interface and Operational Guidelines to address deployment issues.

MSOs and programmer partners, along with tools vendors, have collaborated on several interoperability events, and technical field trials are being contemplated.

Finally, a set of anticipated business models will be introduced.

ETV – ENHANCED TELEVISION

Enhanced and Interactive TV

As seems common in the emerging digital TV landscape, new terms almost always mean different things to different people. For our purposes, the term 'Interactive TV' will be used to denote interactive features that augment the viewing experience but are a part of the content of a program; Electronic Program Guides, or stock/news/weather crawls are examples.

'Enhanced TV' will refer to video programming that has an interactive application bound to it; an application that executes concurrently with the program, and terminates when the program is no longer being presented. Examples include an awards show that allows you to vote on upcoming awards in real-time, and an advertisement that lets you request more information about a product.

An enhanced television program is one that has had an interactive component produced to magnify a viewer's experience. The interactive component is called, naturally, an enhancement.

OCAP Bound Applications

OCAP, the OpenCable Applications Platform, is another MSO initiative managed by CableLabs. OCAP is a Java environment, meant to abstract system details away from applications, so that interactive applications may be deployed on a wide range of televisions and set-top boxes. OCAP contains definitions for 'service bound' applications. These applications are embedded into the same MPEG-2 Transport Stream as the associated audio and video. When a receiver tunes to a program, it reads the data PIDs that contain the application, and where an application is present, loads and launches it.

ETV applications are transported and executed in a similar way. They are accessed when a viewer tunes into an enhanced program, and are terminated when the show ends or the viewer changes channels.

Project Overview

The CableLabs Enhanced Television project was initiated in early 2004 to provide a technical solution for ETV on legacy thin-client cable receivers that would also be appropriate for advanced receivers. MSO and vendor participation has been consistently strong and progress has been rapid.

Mission

Provide a complete technical definition of a national cable ETV platform, to enable deployments near year-end 2006.

Deliverables

ETV applications are compact sets of files that contain the programming logic and resources, like images, that comprise an interactive experience. These files describe application User Interface (UI) and logic and support a set of UI 'widgets', return channel capabilities, and real-time event synchronization mechanisms.

A 'user agent' software module on a receiver reads and renders ETV applications based exclusively on the content of these resource files. The resources are multiplexed with their associated A/V programs in the A/V stream. On-the-wire signaling instructs user agents to acquire and process these resources.

Both legacy (written in native platform code) and advanced (written using OCAP APIs) user agents are supported

The project deliverables are a set of guidelines and specifications that provide an end-to-end platform. The deliverables and timelines are:

- Binary Interchange Format (BIF) Specification. An application format that can be executed on the widest range of receivers, including limited capability devices such as the Motorola DCT 2000, while taking advantage of the advanced capabilities of more advanced devices. This format allows applications to be network and device independent, providing a nationwide cable footprint for ETV. The initial version, Version I01, was issued Q2 '05. Version I02 was issued Q3 '05 and included a number of clarifications and improved features.

  Version I03 will be issued Q1 '06 and will include: support for on-demand sessions, greater graphic display capabilities such as translucency, and an XML format of BIF that enables software tools to easily exchange applications and application elements.

- Application Messaging Specification. Defines network signaling to enable application acquisition by receivers and application synchronization with programming. Based upon OCAP signaling mechanisms, it adds features to support ETV applications and to target limited capability devices.

  Version I01 issued Q2 '05. Version I02 issued Q3 '05.

- Operational Guidelines. This document describes the overall end-to-end architecture and provides information on deployment mechanics, such as downstream and upstream bandwidth management, access to receiver resources, and other relevant topics. A key component of this architecture describes possible national broadcast distributions. Since the initial work was begun, the scope has expanded to include return channel, VOD, testing and other deployment-related issues.

  Version I01 is expected to be issued Q1 '06

- User Interface Guidelines. This document describes how ETV applications are presented to viewers and describes the basic ground rules for providing simple and consistent user experiences. This document defines the cable equivalent of BSkyB's 'red button' commonly used in Britain.

  Version I01 issued Q1 '05

- Metrics Gathering. This document describes an OCAP-oriented approach to gathering metrics data from ETV receivers. Still in its formative stages, this effort will focus on such measurements as click-stream, application lifecycle, and others.

  Version I01 targeted Q4 '06

Features and Functionality

The ETV platform is specifically designed to support relatively simple applications that are delivered with a TV program. Network bandwidth and receiver memory is assumed to be quite limited, constraining the typical size of applications, and they are not generally allowed to perform sensitive operations on a receiver. The platform is not designed to provide a full programming environment for workhorse applications like EPGs.

ETV application make very efficient use of bandwidth and memory resources, through special encoding and data sharing, can be targeted to many receiver classes, executing in a limited fashion on lower end devices while accessing advanced features of others. They can be extended to support innovative receiver and network specific functions.

ETV applications present a user interface through the use of common widgets. Widgets, such as images, text boxes, selectable buttons and so on, can be customized to suite the application's needs. ETV applications can communicate with head-end modules, can save local files if resources permit, and can be tightly synchronized with their associated programs.

The current, baseline ETV platform includes support for some network-specific functions like VOD.

Some Implementation Details

The user agent is responsible for moving an ETV application between these lifecycle states:

- Unloaded
- Loaded
- Suspended
- Running
- Terminated.

A running application is in one of the following navigation states:

- Unrealized
- Disabled
- Enabled
- Focused
- Armed
- Selected.

Multiple applications can run concurrently, and a suspended or not-in-focus application can execute while in the background.

Two-way processing is supported via form submission and form response. Both scheduled and do-it-now stream events (triggers) are delivered to the user agent via in-band data.

An application is characterized by the resources at its disposal, consisting of page, data, platform and other resources such as images and MPEG stills. Each resource structure consists of a header, one or more tables, and a data heap.

All application elements are referenced through tables, and table entries reference data structure offsets in the heap. Some example tables and their contents are:

- *Action Table:* application script references (eg, LoadApp, SelectService, SubmitForm, AddWidget, StoreVariable)

- *Generic Data Table:* self-defining data structures

- *Metadata Table:* application metadata

- *Platform Directory Table:* pointers to platform-specific code and data sections in the resource

- *Reference Table:* acts as a variable store

- *Resource Locator Table:* pointers to external resources and to response data

- *Trigger Table:* trigger schema definitions

- *Widget Table:* provides the hierarchy of pages in the application (eg, Container, Form, Button, Video)

Messages are conveyed to the application using the in-band Enhanced TV Integrated Signaling Stream (EISS). There are three descriptors conveyed in the EISS table:

- *Application Information Descriptor:* used to autostart, present or destroy the application

- *Stream Event Descriptor:* used to send 'do-it-then' stream event and generic payload data to the application

- *Media Time Descriptor:* provides the clock for 'do-it-then' stream events

ETV and OCAP

ETV complements OCAP by enabling what OCAP calls 'bound' applications on non-OCAP receivers. Although the application format is new, the network signaling is common with OCAP. Recall that a receiver 'user agent' renders ETV applications, so ETV support can be added to OCAP by defining how a user agent may operate as an OCAP application, and/or be defining how a user agent can be built into an OCAP implementation.

Interops

The first ETV Interoperability event was held in Aug 2005. There were three User Agents hosted in two major set-top environments and several stream generators and application toolkits represented. This first event provided an excellent demonstration of initial developments, ranging from simple "Hello World" type applications to a sophisticated application designed to accompany a nationally broadcast network program.

The second event was held in February 2006. The target here was for a complete implementation of the platform, and that goal was largely achieved. More extensive use was made of EBIF capabilities, including bitmap images and transparency. Full application signaling was supported, including the sending of real-time stream events. Two-way communication was performed over the return channel, and included retrieval of a VOD asset from a live VOD server. Some basic metrics gathering from receivers was also demonstrated.

Deployment Goals and Field Trials

The CableLabs Executive Committee (CEOs) has directed CableLabs to facilitate and provide project oversight for an ETV field trial in early 2006, with the strategic goal of providing a baseline for national deployment by several MSOs in late 2006.

A key portion of the Operational Guidelines specification is a set of potential national broadcast scenarios, with the issue of affiliate pass-through of ETV programming a particularly thorny one. A field trial is essential in helping to select the viable choices from this list, and ensuring that the affiliate pass-through problem has a solution prior to a deployment decision.

Another gating factor for deployment is a solid framework for gathering and reporting metrics from participating ETV receivers. Again, this is an essential capability that a field trial must demonstrate, and includes considerations of timeliness, chain of flow and security and privacy.

The first phase of the field trial will involve native user agents and alpha versions of applications. Later phases will demand Java user agents and "baked" applications.

Programming Partners

For both the field trials and subsequent deployment, it is vital to have participation from a range of representative programmers and advertising partners. The field trials planned include a broad selection of programmers, broadcast networks, national cable networks, and advertisers.

The revenue-generating business model for ETV deployment must clearly satisfy all of these participants before deployment can become a reality.

Business Models

The notion of enhanced programming has been both widely praised and ridiculed for many years. Time will tell which camp is more closely right, but without a mass medium roll-out one's opinions remain safe from an empirical test. We do have the benefit of some very successful examples on which to pin our hopes. PC/Internet based 'two-screen' applications have been deployed by many programmers that have shown extremely high take rates. Mobile phone applications have also been deployed to accompany popular awards and games shows. These examples demonstrate a hearty appetite for interactive services among TV audiences.

Another important example is the BSkyB platform in England. Very wide usage of interactive features incorporated into a wide range of programming has been observed. The revenue prospects for BSkyB are improved by the twin facts that return channel communications involve a toll call, which provides a revenue split between the phone and video carrier, and that betting is less regulated that in the US, allowing cash cow applications like real-time horse racing to prosper.

The great hope for enhanced programming on cable is perhaps the opportunity to add interactive features to advertising, thereby increasing the value of ads to both sponsors and networks.

Do not count out the creativity, and pecuniary interest, of the television industry. We may very well see some ‘killer-app’ emerge from an unexpected quarter.

EXPLOITING HFC BANDWIDTH CAPACITY TO COMPETE WITH FTTH

Tony Werner
Senior Vice President and Chief Technology Officer, Liberty Global, Inc.
Oleh J. Sniezko
Chief Technical Officer, Aurora Networks, Inc.

Abstract

Several different Fiber-to-the-Home (FTTH) architectures are starting to emerge and be deployed. Proponents of FTTH recognize the high cost of installation, but stress the capacity advantages over Hybrid Fiber Coax (HFC) networks. While it is true that optical fiber has almost unlimited capacity, the practical capacity of these networks are not superior to HFC networks in most cases. This paper presents capacity comparisons of popular FTTH architectures with that of a modern HFC network. In addition to this comparison, the paper also explores several methods for exploiting the significant unused capacity of HFC networks.

This paper presents CWDM downstream and upstream technologies that allow for low-cost segmentation of the optical serving area to the levels below 100 homes. As demonstrated this can be accomplished without additional fibers between the node and the headend or hub. This architectural modification of the HFC networks is also non- service interrupting and can provide capacities that meet or exceed today's PON architectures at a fraction of the cost.

Additionally the paper lies out how the HFC architecture can efficiently provide this additional bandwidth on a geographically granular basis, up to and including Fiber-to-the-Building where it makes sense for business applications or large multi-tenant buildings.

COMPETITION

FTTH Deployments

Fiber to the Home (FTTH) has been deployed in varying degrees throughout the world. Each region has its own set of variables for choosing FTTH and what works in one country does not necessarily make sense in another.

Japan is leading the world in FTTH deployment. As of September 2005, Japan had 3.98 million FTTH subscribers and is adding 100K new subs a month. NTT is the largest carrier of FTTH and has reported that it will be investing $47 billion through 2010 to upgrade 30 million homes and businesses.[1] Even though FTTH is gaining subscribers, it is not having a material impact in areas where modern HFC networks exist. The HFC networks are typically offering 15 to 30 Mbps tiers combined with rich video offerings which satisfy most of the Japanese consumers.

Korea has the highest broadband penetration rate in the world, aims to have 100 Mbps service available to 5 million subscribers by 2007 and to 10 million subscribers by 2010 mainly over a FTTH network.[2]

Limited FTTH trials exist in Europe and are usually being introduced by local municipalities. As of Q2 2005, there were 166 FTTx trials/projects, 72% of which were initiated either by the municipal or the local power utility.[3]

Five countries are responsible for 97% of the activity, they are: Sweden, Italy, Denmark, and Norway.

Other than Verizon, FTTH deployment in the US has also mainly been rolled out by municipalities. Today FTTH has been trialed or rolled out in 652 communities across 46 states, but only accounts for 322,700 connected homes.[4] Nearly 50% of those projects do offer some form of triple-play services.

Verizon announced aggressive plans with the launch of FiOS. They reported to have 3 million homes passed by the end of 2005 and plans for another 3 million homes passed in 2006.[5] Several of these homes are apartment complexes which have not been wired for service. Subscriber success has not been clear in these markets. Most of the incumbent cable operators have preempted the FiOS offer of 15 Mbps by 2 Mbps with an equivalent or higher speed offer. Verizon has been quoted as achieving over 20% penetration in certain markets. It is unclear, however, whether this is an actual penetration number or a sales to contact number as the competing cable operators claim that Verizon has only actually achieved low single digits.

HFC: Capacity Overview

FTTH perhaps has more marketing appeal than it does technical appeal. It is true that optical fiber's theoretical bandwidth capacity is nearly unlimited. It is also true that 10 Gbps per wavelength with 50 GHz can be commercial deployed for long haul networks today, providing 800 Gbps in the 1550 nm window. These technologies are not practical for access architectures for several reasons and as such the practical capacity of FTTH networks is very similar to that of modern HFC networks.

Traditional 870 MHz HFC networks are capable of over nearly 5 Gbps of downstream capacity and more than 150 Mbps of upstream capacity. European and Japanese cable systems are capable of 270 Mbps in their upstream and both US and international HFC networks can be configured in a Next Generation Network Architecture which can provide over 6 Gbps in the downstream and 360 Mbps in the upstream.

It must also be recognized these speeds can be accomplished at a fraction of the cost of FTTH networks.

CWDM, DWDM Complement HFC

The optical links between the headend and the node match the capacity of the traditional telecommunications networks. Time Division, Wave Division and Frequency Division multiplexing techniques can and are all employed on this portion of the network. This allows for extremely efficient use of the optical fiber. Thanks to short distances limited noise contributions the CNR/SNR requirements of Shannon's bandwidth theorem can be easily met.

HFC Network Capacity

The HFC network has significant capacity and is an excellent position to compete with FTTH networks Figure 1 and 2 compare the capacity of A,B,G, and E, PONS to HFC networks of varying node reductions.

As figure 1 indicates even a traditional 870 MHz HFC network with 500 home passed nodes has more downstream bandwidth per customer than most FTTH architectures today. After segmentation, downstream bandwidth can significantly exceed that of most PONs and upstream bandwidth can achieve parity.

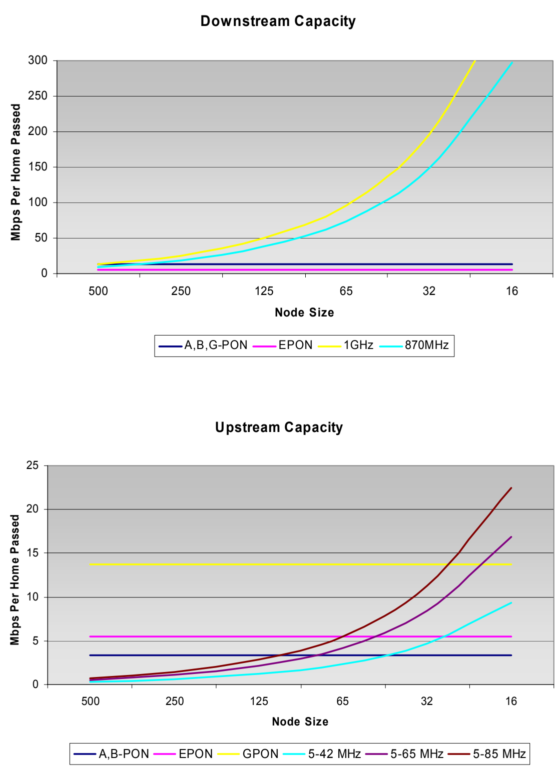

Figure 1. Bandwidth Capacity Comparison PON vs. HFC

In addition to HFC networks being able to eloquently evolve to increased bandwidth, it is also well suited to evolve to FTTx at any point based on demand. This demand can be very granular and thus the economics quite attractive. Key things to consider are the bandwidth limiting devices in the coaxial network and the distance limit for low-cost optics in the optical network. Typically the bandwidth limitions in the Coaxial networks are a result of RF actives. As fiber is deployed deeper, several options are available to overcome these limitations, including removing the RF actives all together.

LINKS TO THE NODES

Analog Links: CWDM vs. DWDM and Digital Baseband Technologies

DWDM capability of the optical links between headends and hubs has been documented in Figures 1, 2, 3, and 4. This technology is most suited for high-level aggregation (40 nodes can be fed from 3 fibers) over long distances thanks to cost-effective optical amplification. Designs reaching the ranging limits of the DOCSIS systems have been implemented and operated for several years. However, for shorter distances and segmentation applications more cost-effective techniques exist. The following reviews some of these techniques.

Forward CWDM Links

The Coarse Wave Division Multiplexing (CWDM) technology, especially when used for FDM analog and QAM signals, encounter at least two major challenges. One of them is SRS-caused crosstalk between CWDM wavelengths on the same fiber. The other is high level of dispersion in SMF-28 or equivalent fiber above water peak. This fiber type is dominant in access networks today.

The theoretical description of the Stimulated Raman Scattering and its relation to the phenomena of Raman gain is well understood and is used in optical amplification. However, the same phenomenon leads to undesired amplification (shorter wavelength pass their energy to longer wavelengths) in multi-wavelength systems. If a wavelength is modulated, this amplification results in bi-directional (theoretically asymmetrical) crosstalk.

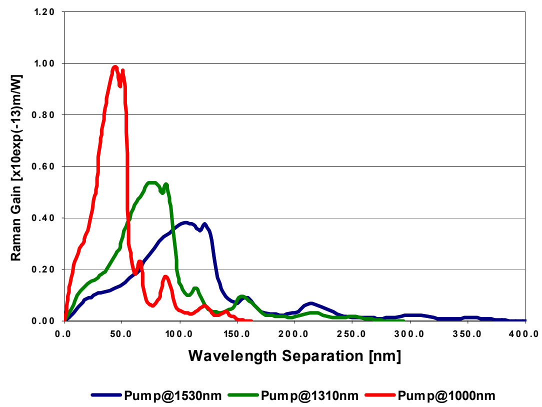

a) Raman Gain vs. Wavelength Separation

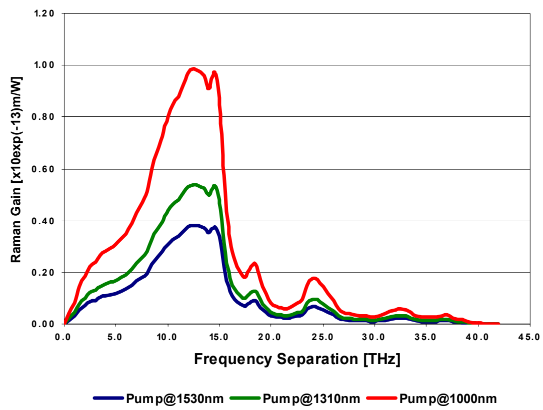

b) Raman Gain vs. Frequency Separation

Figure 2a & 2b. Raman Gain for Three Different Pump Wavelengths

Figure 2 shows the theoretical plots of the Raman gain. This theoretical relationship was closely matched during measurements of crosstalk at several pump wavelengths. (An example of the test results is shown in Figure 3.) Other contributions to crosstalk (e.g., XPM) at lower separations between wavelengths and higher RF frequencies cause the crosstalk values to deviate from the theoretical Raman gain relationship. (Detail descriptions of test methodologies and test results are beyond the scope of this paper.)

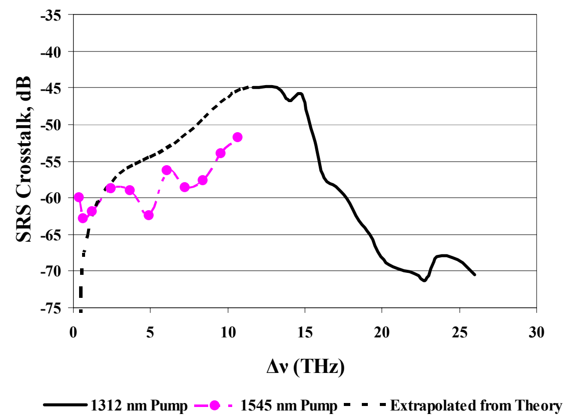

a) 499.25 MHz

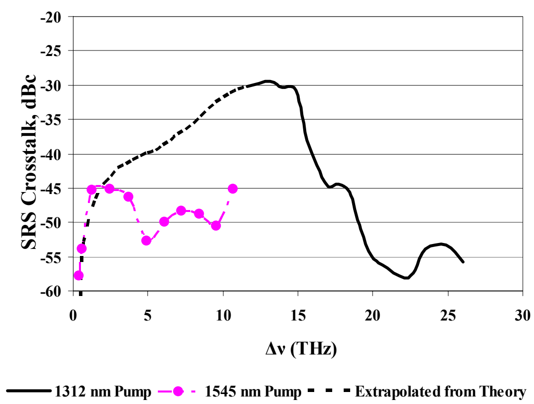

b) 55.25 MHz

Figure 3. Crosstalk vs. Δν (Test Results for +10 dBm Pumps and 25.3 km of Fiber)

The crosstalk test results indicate that low frequency NTSC analog carriers of different content cannot be carried without the possibility of interference on any combination of two wavelengths unless they are separated by more than 30 THz. However, since QAM channels can tolerate higher level of interference, the crosstalk between CWDM wavelengths at higher RF frequency (where QAM channels are usually placed) can be low enough to allow carrying QAM channels of different content on different wavelengths. The test results were used to calculate the limits (under most conservative assumptions) of CWDM systems from QAM signal crosstalk point of view. Figure 4 shows cumulative crosstalk for a CWDM system when the fiber loading starts with 1270 nm and consecutive wavelengths are added (except water-peak wavelengths: 1310, 1390 and 1410 nm).

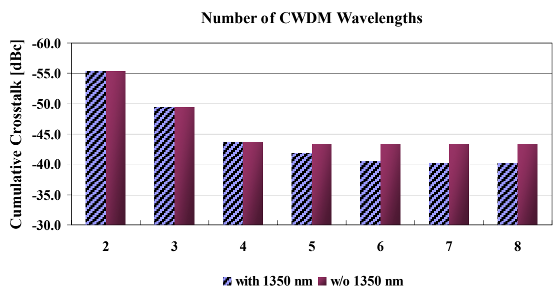

Figure 4. Calculated Cumulative Crosstalk at 499.25 MHz

Depending on the acceptable level of interference, a combination of several wavelengths with QAM loading of different contents above 500 MHz can be supported. Moreover, the higher the RF frequency, the lower the crosstalk.

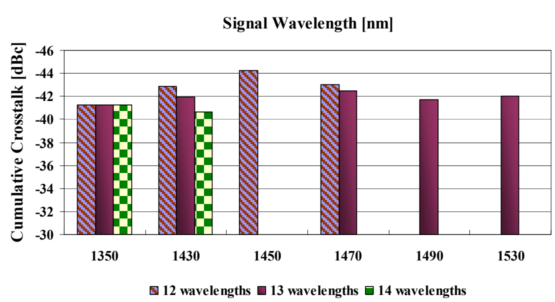

Figure 5. Cumulative Crosstalk for Multi-Wavelength Loads

Figure 5 shows that up to 12 CWDM wavelengths (up to 14 wavelengths could be acceptable) can be combined onto a single fiber as long as the NTSC analog video channels carry the same content and QAM channels are placed above 500 MHz. Unfortunately, Raman gain crosstalk is not the only impairment that can cause problems in analog optical links.

<u>Dispersion</u>

Dispersion in SMF-28 or equivalent type fibers can actually introduce stringent limits on the number of CWDM wavelengths carrying NTSC analog video channels in a single fiber. Figure 6 shows typical dispersion relationship for SMF-28 fiber.

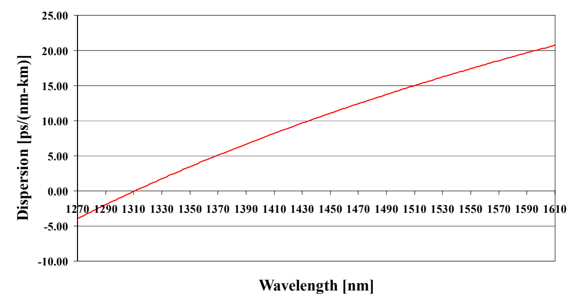

Figure 6. Dispersion of SMF-28 Fiber

The combination of high dispersion and laser chirp that is inherent in typical direct modulated lasers with FM efficiency of 100 MHz/mA will cause second order distortions (including CSO).

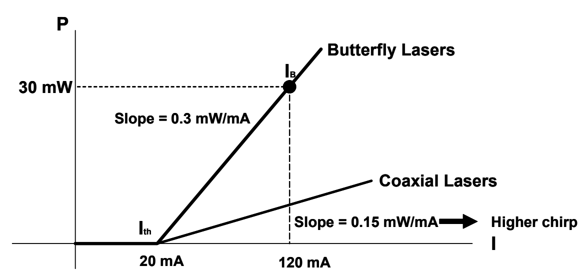

Figure 7. Typical Laser Characteristic

The total chirp will depend on the laser FM efficiency and slope (see Figure 7 for typical laser characteristic). Under the assumptions presented above, 30 mW lasers will have 10 GHz chirp (at 100% modulation) and 3 mW lasers will have 1

GHz chirp. Figure 8 shows what levels of distortions can be caused solely by laser chirp and dispersion at 1450 nm wavelength over 20 km of SMF-28 fiber.

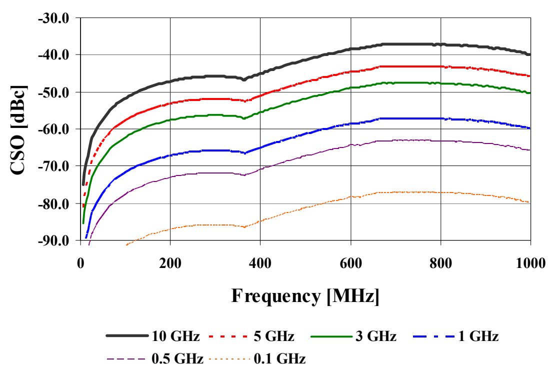

Figure 8. Second Order Distortion Caused by Dispersion at 1450 nm Wavelength in 20 km of SMF-28 Fiber

Figure 9 presents expected second order distortion levels at different wavelengths for different lasers in 20 km of SMF-28 fiber.

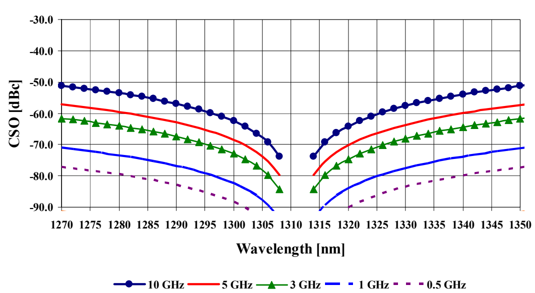

Figure 9. Second Order Distortions vs. Wavelength (20 km of SMF-28 Fiber)

If we assume that the contribution to the laser CSO from chirp/dispersion generated CSO cannot exceed –70 dBc (to avoid significant degradation of laser nonlinearity-generated CSO), then we can calculate the maximum number of CWDM wavelengths per fiber in the forward direction.

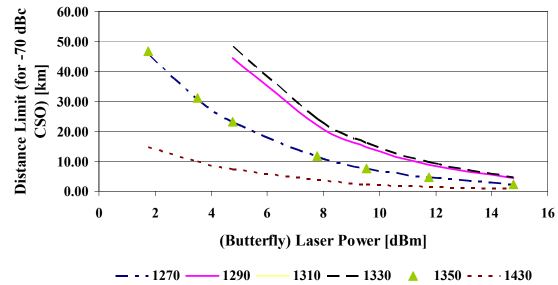

Figure 10. Dispersion CSO Distance Limits vs. Laser Power (Chirp)

Figure 10 presents the results of these calculations at 548.5 MHz, assuming FM efficiency of 100 MHz/mA for laser with characteristics presented in Figure 7. Under these assumptions, 6 wavelengths can be placed on a single fiber of 10 km length (loss budget permitting), 5 wavelengths on 15 km long fiber and 3 wavelengths on 20 km long fiber.

The SRS and dispersion considerations above were based on these assumptions:

1. The transmitter loading is hybrid analog/digital QAM
2. Analog channels on all wavelengths are the same.
3. Analog load consist of 77 NTSC video channels between 54 and 552 MHz.

If the number of channels or the loading type changes, the values in Figure 10 will change as well. In the extreme, with digital-only loading, the dispersion may be less limiting than Raman crosstalk (crosstalk is highest at the lowest frequencies, dispersion generated CSO is highest at 725 MHz) unless the low frequency QAM channels carry the same information on all wavelengths.

Dispersion Remedies

Under the assumptions used in the CWDM analysis, the limiting factor is dispersion combined with the laser chirp. To ease these limitations, the following can be implemented:

1. Lower chirp lasers used,

2. Dispersion compensation circuitry added in the transmitter for high-dispersion wavelengths,

3. Dual receiver configuration used in the node.

A detailed analysis of pros and cons for each of these solutions should decide about the selection of the optimal solution for a particular application. The remedies #1 and #3 can be easily implemented. Figure 11 shows a simplified diagram of the dual receiver links.

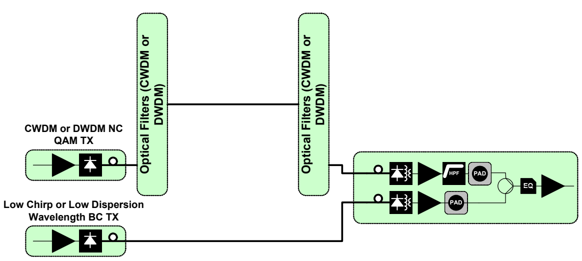

Figure 11. Dual Receiver System

One must note that even at 20 km, two fibers can feed 6 independent forward areas with 3 forward wavelengths per fiber and with the reverse signals counter-propagating on the same fibers that carry forward signals.

Analog Reverse Links

Analog reverse transmitters are lower power (typically 3 dBm or lower). Moreover, only digital (64 QAM max) channels are transmitted in the reverse links. For these reasons, SRS crosstalk can be disregarded. Similarly, CSO problems in DFB analog links can be disregarded (low chirp for low power laser), even in coaxial lasers. However, CSO in-links with FP lasers of much higher chirp must be analyzed at wavelengths above the water peak (1430 nm and longer). Even for –30 dBc C/I requirements, these links may be limited in distance.

Digital Baseband Links

Baseband digital links do not show problems attributed to the analog links. The power into the fiber for digital links is usually low (no Raman crosstalk problems) and dispersion does not lead to second order distortion. Instead, it results in pulse spreading but digital lasers are designed for a specific dispersion limits and digital technology developed many different remedies against dispersion (chromatic and PMD)

Several manufactures have been deploying digital baseband transport technology in reverse links. This transport allows for using both CWDM and DWDM technology to support reverse bandwidth up to 85 MHz (NGNA specified reverse upper frequency limit).

Moore’s Law provides for increased access to low-cost, high-data rate components. An example of this are Quad Fiber Channel and 10 Gbps transceivers, which are now readily available. This can add to the capacity of fiber between the headends/hubs and the nodes supporting digital reverse and providing additional bandwidth in the first mile plant.

The baseband digital link can share the forward and reverse fiber either in a counter-propagating manner or co-propagating manner if the frequency spectrum of the signals carried does not cause Raman crosstalk interference to other signals.

Applications for CWDM

Figures 13 and 14 present two examples of applying CWDM technology in the forward (analog transmitters) and reverse (with digital multiplexing of two reverse segments per wavelength) paths to segment fiber deep node clusters. The first implementation with the distance to the first node limited to 10 km uses full-load CWDM transmitters. The second implementation lends itself architecturally to a dual receiver configuration due to operator choice of Broadcast and Narrowcast equipment locations and the distance to the first node. The bandwidth capacity gains per household are presented in Figure 12 (refer also to Figure 1). Capacity per user will depend on the service penetration levels. The following assumptions were used:

1. Both clusters serve 500 households.
2. Initial sub split is 42/54 MHz; final sub split is 85/105 MHz.
3. Reverse used 16 QAM modulation initially and 64 QAM finally
4. Forward NC bandwidth is 192 MHz initially and 288 MHz finally.

The segmentation and other modifications listed above multiplied the forward NC bandwidth per household by a factor of 7.5 and the reverse bandwidth by a factor larger than 40. This still leaves 105 to 582 MHz bandwidth for broadcast signals. If this bandwidth is filled with digital 256 QAM signals, it will provide more than 3 Gbps broadcast capacity.

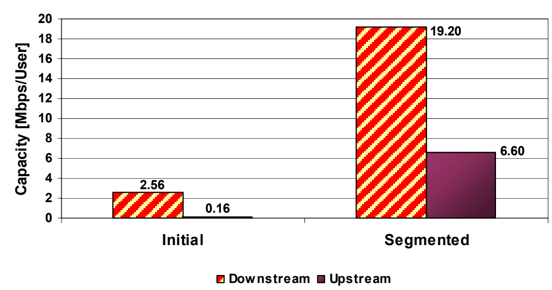

Figure 12. Bandwidth Capacity Gains per Household due to CWDM Segmentation Technology

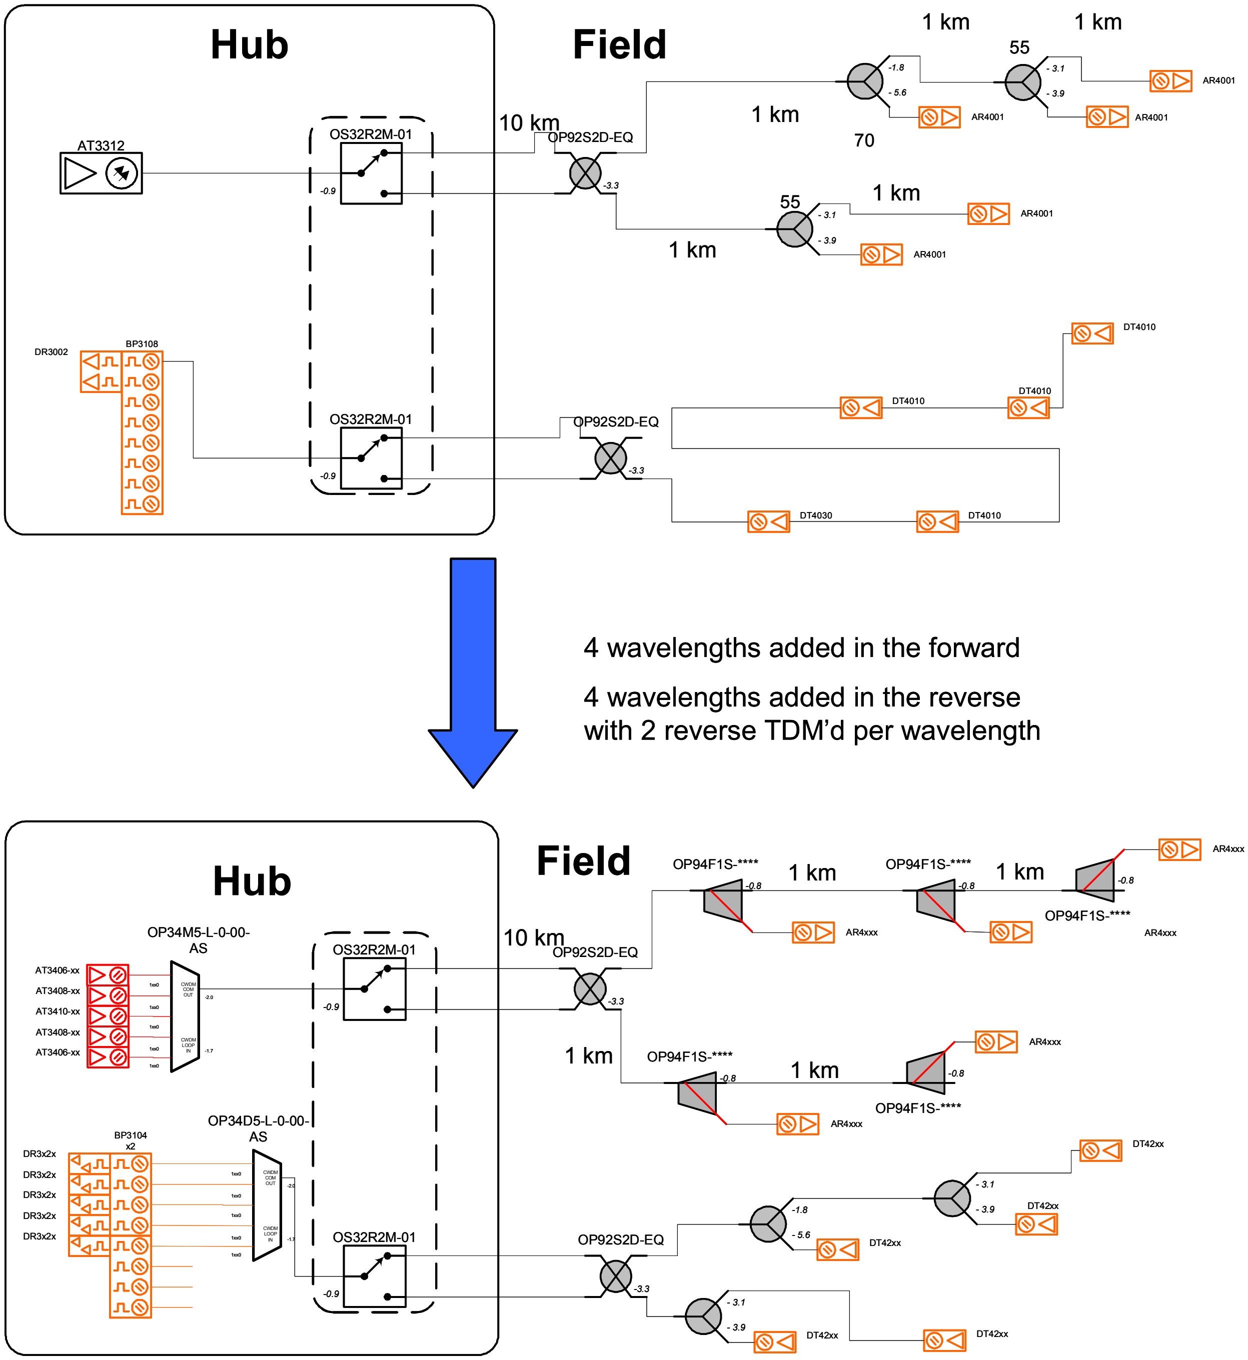

Figure 13. Segmentation with CWDM Technology: Forward with CWDM Full-Load Transmitters; Reverse with CWDM Digital Reverse with TDM'd Paths

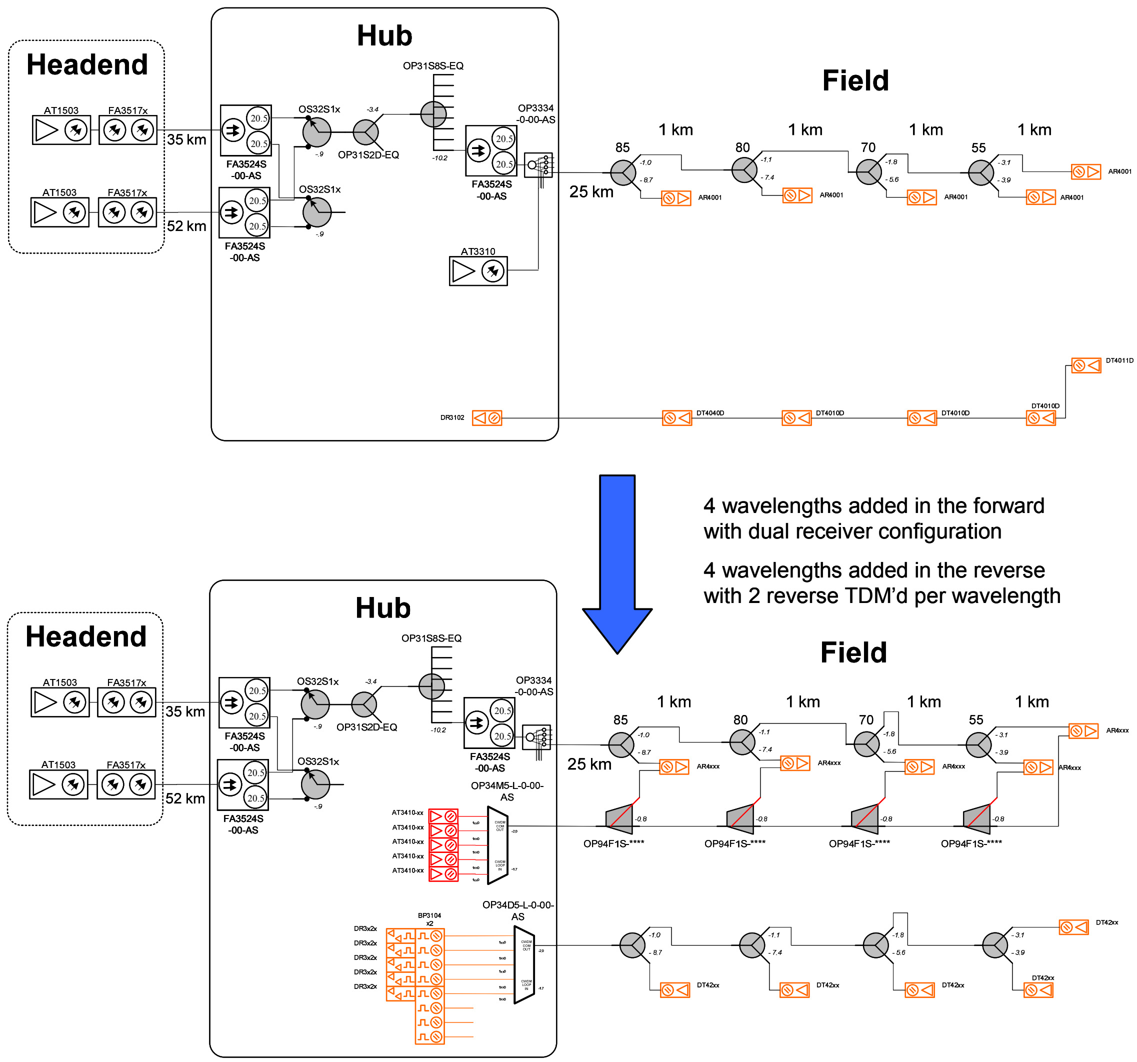

Figure 14. Segmentation with CWDM Technology: Forward with CWDM NC Transmitters and Dual Receivers; Reverse with CWDM Digital Reverse with TDM'd Paths

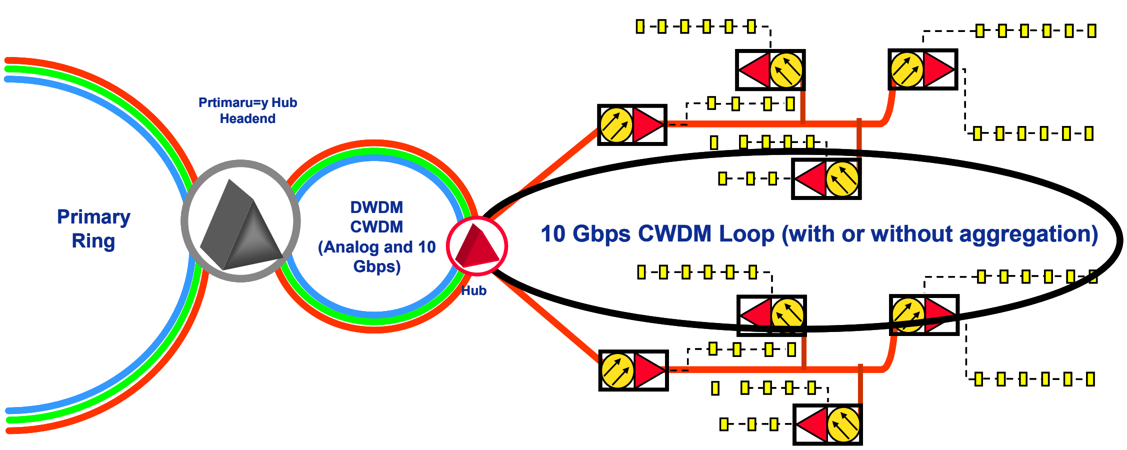

Figure 15. Digital 10 Gbps Transport Technology between Headend/Hub and Nodes

Applications for Digital

Figure 15 presents a method of implementing digital baseband transport technology between headend/hub and the nodes. The multi-wavelength fiber capacity allows for filling up the unused CWDM and DWDM wavelengths in a counter-propagating or co-propagating manner to support bandwidth capacity enhancement in the first mile plant beyond those offered by traditional forward and reverse HFC technologies.

FIRST MILE TECHNOLOGIES

HFC architectures have significant fiber capacity based upon these optical technologies. The challenge is delivering this capacity over the coaxial cable.

Absolute Bandwidth and Bandwidth per User

The efforts of exploiting coaxial plant capacity progresses in two dimensions: bandwidth expansion and expansion of bandwidth per customer. In the first category are such efforts as:

1. Increasing system capacity towards 1 GHz with a combination of traditional analog video and digital QAM channels

2. Using spectrum above the existing nominal design limits of the broadband subsystem.

In the second category are efforts to:

1. Segment nodes into smaller serving areas by using fiber capacity and by extending fiber deeper into the coaxial plant to the point of eliminating RF actives after the optical node,

2. Replace analog channels with digital channels,

3. Improve digital signal efficiency by:

   a. Increasing QAM modulation levels for digital signals,

   b. Increasing coding capacity for digital video signals,

   c. Reclaiming broadcast digital bandwidth with switched digital architecture, and

   d. Increasing stat-muxing efficiency of digital video signals.

All these efforts can lead to reclaiming up to 288 MHz of forward bandwidth for narrowcast data signals. The hope is that DOCSIS 3.0 will allow using this bandwidth in a manner similar to FTTH where very high- capacity forward and reverse channels are shared among multiple users to improve statistical muxing gains.

CMTS Cost

Today, DOCSIS CMTS channels typically serve between 500 and 1500 users with a single forward channel and multiple reverse channels. Average cost per user ranges from $5 to $20. Many sub-systems of the CMTS card are under-utilized, for example, typical reverse channel capacity exceeds forward channel capacity.

To match the capacity of APON and provide 640 Mbps downstream capacity and 150 Mbps of upstream capacity, 12 forward 6 MHz channels and 4 reverse 6.4 MHz channels are required. This accounts for the true network capacity of the APON architecture.

To match APON electronic costs of $300 per link, the DOCSIS CMTS configuration described above must drop to approximately $6,000. Note that GPON can provide higher capacity so the CMTS pricing will need to be even lower than this. The DOCSIS 3.0 CMTS configuration must include all components downstream of network interface (including QAM modulators and burst receivers) within the cost target indicated.

The HFC plant has significant advantage in its scalability. FTTH plant must be built from day one for the final penetration level due to loss budget requirements. The HFC plant can add equipment in the headend as the penetration levels increase. This will allow for taking advantage of declining prices and allows the operator to deploy bandwidth when and where it is needed.

Gbps over Coax

If and where required, fiber can be extended to the last active (fiber deep deployments). This can be done efficiently and in a non-interruptive manner (no changes to the coaxial plant). With the actives eliminated from the coaxial network several options exist for increasing its capacity.

The coaxial cable spectrum is not limited to 870 MHz or 1 GHz. Most of the current deployments of HFC networks use 1 GHz passives and the passive section of coaxial plant can be easily used to 1.5 GHz and even to 3 GHz as long as it is not restricted by RF actives. This can allow for a use of the bandwidth above 870 MHz for point-to-multipoint (P2MP) technology deployment over passive coaxial plant.

Figure 16 presents one of many possible ways of spectrum utilization above the traditional HFC bandwidth. More optimal arrangements are being designed to simplify the implementation of this system and its integration with passive coaxial plant.

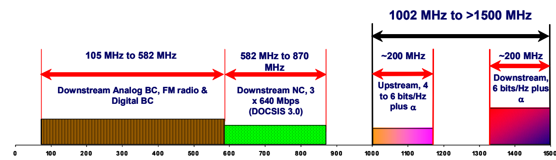

Figure 16. One of Possible Spectrum Utilization of Coaxial Passive Network

Beside physical layer, data and MAC layers are being developed to allow for adding Gbps capacity to the HFC capacity. The optimal solution would allow for several different data rates in forward and reverse to allow for adaptation to passive coaxial network conditions above the HFC operational frequency. It is mostly designed to take advantage of very short distances between the optical node and the farthest customer.

This technology can be deployed in a selective or opportunistic manner in areas where extreme capacity is required.

Fiber on Demand

The digital capacity of fiber to the node can support much more bandwidth and many more applications. Moreover, the fiber in HFC network and especially in fiber deep HFC network is deployed to the proximity of the residential and business neighborhood. It is closer to the premises than the fiber in FTTN architecture where the node is designed to serve an area of 2,000 households. Therefore, at very low additional construction cost, a P2P (point-to-point) or P2MP fiber links can be deployed from the node to the premises.

Initial deployments of FTTP can serve businesses, schools and other public building with service expansion to SOHO premises and MDUs. These deployments can initially start with point-to-point (star) topology from the node to the premises. Modules providing an interface between the node digital uplink and standard FE or GigE are deployed in the optical node. They allow for installation of IEEE 802.3 standard compliant CPE devices and support 802.3ah capability. The CPE devices can be purchased off-the-shelf and self-installed by the customer or installed by an operator.

For residential deployments, an xPON compliant OLT is installed in the node and standard based ONTs are installed on customer premises. Note that the distance is significantly shorter than the distances in traditional FTTH deployments. This has a direct impact on loss budget and hence can result in selection of lower cost components and subassemblies.

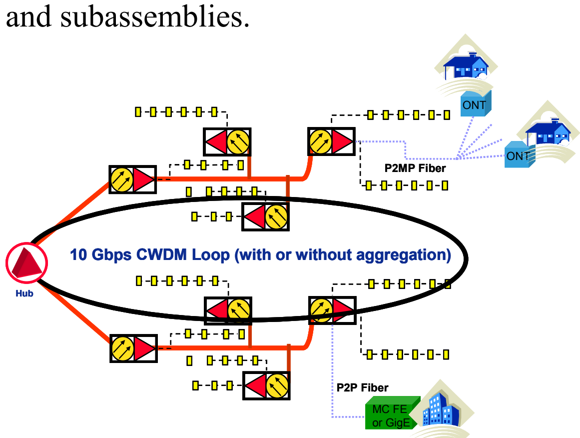

Figure 17. P2P and P2MP Opportunistic Deployments of FTTP in Fiber Deep HFC Networks

Both types of deployments can be implemented in an opportunistic manner, very similar to manner of deployment provided by many business service operators today. To further lower the future cost of deploying fibers, provisioning for fiber in the access plant during green-field construction can be implemented. A significant advantage of the evolutionary approach is the fact that all services supported by HFC are still being provided and the investment in all facilities and service equipment is fully utilized. Only when and where required or beneficial, is FTTP/FTTH from optical nodes deployed.

SUMMARY AND CONCLUSIONS

Node Uplink Capacity

The uplink in HFC nodes is based on the fiber optical technology and all developments that happen in this technology can be applied in those links. CWDM and DWDM for analog and QAM signals and for digital signals allow for exploiting the fiber capacity to its full potential. The optical nodes can be connected to the network via links that will not become a bottleneck for the traffic generated by the user connected to these nodes. The requirement is to deploy fiber deep enough into the access network so the first mile plant can put the uplink capacity to full use.

Passive Coaxial Network Capacity

Coaxial network capacity is utilized only partially due to a simple fact that the loss of coax increases with frequency. In the past, the loss was compensated with RF amplifiers but at the same time RF amplifiers were limiting the bandwidth potential of the passive coax. With fiber deep into the HFC network, it is possible to support the traditional HFC bandwidth delivery to the customers without additional RF amplification between the node and the customer outlets. This traditional HFC bandwidth can provide increased capacity per user with the help of uplink technologies and DOCSIS and digital video technologies.

In passive coaxial network, the capacity above the traditional HFC bandwidth is now open to easy mining with advanced digital coding and modulation techniques. This bandwidth capacity can be expanded to 1.2 to 1.6 GHz in the existing plant and to 2 to 3 GHz with deployment of expanded bandwidth passives. Implementing Gbps bi-directional capacity over coax becomes possible in this scenario.

Evolution to FTTP

The fiber push to the last active shrinks the distance between the fiber and the farthest customer. This distance stays below 1,000 m and most customers are within 500 m and drastically closer in densely populated areas. An opportunistic deployment of fiber to the premises becomes affordable, especially when provisioning for the future deployment took place during the plant construction. P2P dedicated links to businesses and high-bandwidth users and P2MP PON deployments from the node will allow for a new dimension added to the HFC network to further improve its competitiveness.

Other Factors to Consider

In the environment of increased competition, other factors besides bandwidth and capacity are important. In a perfectly competitive environment, the variable cost of providing a unit of outcome becomes critical. While HFC deployments have an advantage in capital outlay per household, operational costs are also important parameters. As operators extend fiber deeper it both increases network reliability and lowers maintenance and other operational cost.

ACKNOWLEDGEMENT

The authors wish to express their gratitude to Sudhesh Mysore for his contribution into analysis and testing of the phenomena deciding about the use of CWDM analog links for HFC plant segmentation.

[1] Ministry of Public Management – Information and Communications in Japan
[2] Ministry of Information and Communication, Republic of Korea, e-Korea Vision 2006, April 2002
[3] IDATE, January 2006
[4] Broadband Properties, November 2005
[5] Verizon Vice Chairman Lawrence Babbio, Merrill Lynch 2006 Communications Forum, February 28, 2006

EXPLORING THE IMPACT OF MIXING DELAY SENSITIVE AND DELAY INSENSITIVE TRAFFIC ON DOCSIS NETWORKS

Stuart Lipoff
IP Action Partners

Abstract

Results are presented, and discussed, of discrete event simulation and mathematical modeling exploring the quality of service (QoS) cross impacts of mixing various types of delay insensitive (e.g. web surfing) and delay sensitive (e.g. VoIP telephony) on a DOCSIS network. Of particular interest is the impact on over-the-top streaming traffic when the operator asserts QoS control over streaming services offered and managed by the cable operator.

INTRODUCTION

Historically most DOCSIS cableplant IP traffic was largely downstream intensive and undifferentiated– mainly consisting of delay insensitive web surfing and file downloads. However in today's situation we already have new applications (e.g. peer-to-peer) vastly increasing traffic intensity. As we look toward a future traffic mix of delay sensitive carrier managed services combined with delay sensitive 3rd party over-the-top services; there is the possibility of negative impacts on the user experience resulting from the new traffic mix.

The DOCSIS protocol was designed from the start to be future proof. By employing node splitting and/or exercising protocol options, the operators can scale the service to support increased traffic. DOCSIS also has hooks that allow highly flexible management of the shared up and down stream bandwidth to support service level guaranties, quality of service (QoS), and the overall user experience. Under the assumption that cable operators (MSO) act rationality, they will take advantage of the future proof capabilities of DOCSIS to add capacity and activate features (e.g. QoS) as the operator rolls out new services (e.g. voice over IP telephony) that require QoS and/or increased bandwidth. The flip side of a rational business decision however, is that an MSO would not be motivated to spend capital to add extra capacity to their cablemodem network simply to support a 3rd party service offering; especially when such a third party offering may require special treatment or this third party offering consumes an excessive share of the limited up and down stream system bandwidth without additional compensation to the MSO.

This study was designed to understand which types of services can gracefully coexist on the same plant versus those services that consume plant resources out of proportion to their economic value. In order to develop these insights, mathematical and discrete event modeling was employed to simulate the impact of adding various example services to a system that was rationally engineered to support a baseline web surfing traffic + MSO provided telephony load. The analysis explores the impacts on the user experience as a function of increasing traffic by adding new delay sensitive traffic to the legacy delay insensitive system designed for web surfing. In a similar fashion, the cross impact is studied between delay sensitive traffic provided by the MSO and delay sensitive traffic created by the subscriber adopting over-the-top delay sensitive services without the knowledge and support of the MSO.

Examples of the delay sensitive traffic include: streaming music, streaming video, voice telephony, and multiplayer games. Of particular interest are the current over-the-top third party service offerings of public switched telephone network (PSTN) voice telephony. The 3rd party providers of such services represent them to be comparable to legacy circuit switched PSTN services offered by local exchange carriers (LECs). The reason for my particular focus upon over-the-top PSTN voice telephony is the following:

- MSOs are now in the process of rolling out Packet Cable managed PSTN voice telephony services that are expected to consume significant up and downstream DOCSIS resource beyond today's legacy situation.

- Since the traffic and user expectations for PSTN voice telephony is well known, it is possible to understand the impacts of traffic mix changes on PSTN voice telephony using objective criteria that can be modeled and analyzed

- An MSO provided PSTN voice telephony service can take advantage of DOCSIS features to manage QoS and reduce the management message overhead of up and down stream resource allocation by means of Packet Cable features that closely integrate the PSTN packet switch with the CMTS. For example, RFC2748-COPS messages from the call management server to the CMTS to allocate Unsolicited Grant Service.

- An over-the-top PSTN voice telephony service can not take advantage of the MSO's efficient resource allocation mechanisms and instead must appear to the CMTS in the same way that other best effort traffic (e.g. web surfing) appears to the CMTS. Given the special needs of streaming delay insensitive voice telephony in both the up and down stream directions, it would suggest careful study to determine if a cable system designed for best effort delay insensitive web surfing would provide a satisfactory user experience for an over-the-top delay sensitive service offering.

An important question of interest to the cable industry MSO and technology providers is: what absolute capacity limits can a particular system provide for a particular mix of traffic. Because DOCSIS offers such a wide variety of configuration options taken together with HFC network partitioning variations, the infinite possibilities of service mix, subscriber adoption rates, and traffic intensity; the absolute capacity limits of DOCSIS is not a topic I address in this study. Instead the focus of this study is upon the shape of the operating characteristic that relates offered traffic to the user experience performance. To explain the meaning of shape in the context of this study, I refer to figure 1 below. Ideally you would like a communications system to have a type 1 linear operating characteristic where there is a linear relationship between the offered traffic and the system performance as experienced by the user. The system performance parameter(s) will vary somewhat depending on the type of traffic. For web surfing, peak down load speeds would be the primary system performance parameter.

For PSTN voice, call attempt blocking and packet loss during the call would be among the important performance parameters. For most multiplayer games, latency time would be the primary performance parameter of interest. Although a linear operating characteristic is ideal, most real multiple access systems, such as DOCSIS, have curves that look non-linear like type 2 where there are limits and thresholds. The objective of this study is to understand where the limits kick in and how steep is the curve's slope as the offered traffic mix is increased beyond the threshold.

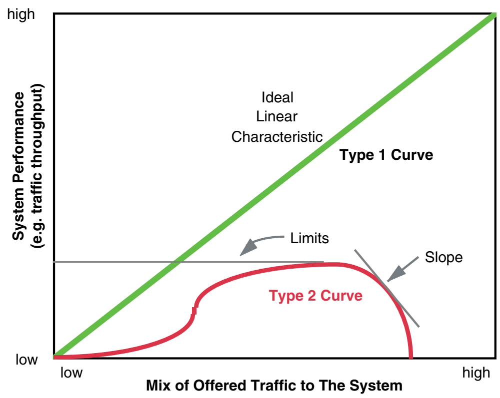

Figure 1. Illustrative Operating Characteristics

BASELINE DOCSIS SYSTEM ASSUMPTIONS

In parallel with rapidly growing consumer adoption of broadband cablemodem services we also have rapid growth in bandwidth hungry multimedia traffic. With that in mind, picking a single point system configuration as a baseline that is representative of all present and future cable systems is not possible. Non-the-less with so many possible variables, modeling requires some assumptions be made and a subset be held constant while other subsets are varied to develop insights. Since the stated objective of this study is not to develop a perspective on absolute traffic capacity of DOCSIS in general, but rather to understand the figure 1 nature of the shape of the operating characteristics, the choice of a particular baseline is somewhat less important as long at the assumptions are reasonable. For all the models analyzed in this study, the network layout shown in figure 2 is assumed. In the figure 2 system each optical transition node (OTN) has 500 homes passed (HP) with an 80% cable take rate of which 20% of the cable subscribers take DOCSIS cablemodem (CM) service resulting in 80 CMs/OTN. At the CMTS, each downstream port of the CMTS splits to serve 6 OTNs while there is a 1:1 mapping of one OTNs into one CMTS upstream port. The resulting load per CMTS port is 480 CMs/Port downstream and 80 CMs/Port upstream. As the source note in figure 2 indicates, these parameters are based on assumptions in a 2004 Cisco publication. However, in analyzing the Cisco assumptions in the light of today's typical traffic loads, for the purpose of this study I split the traffic on each CMTS by two times over Cisco's assumptions.

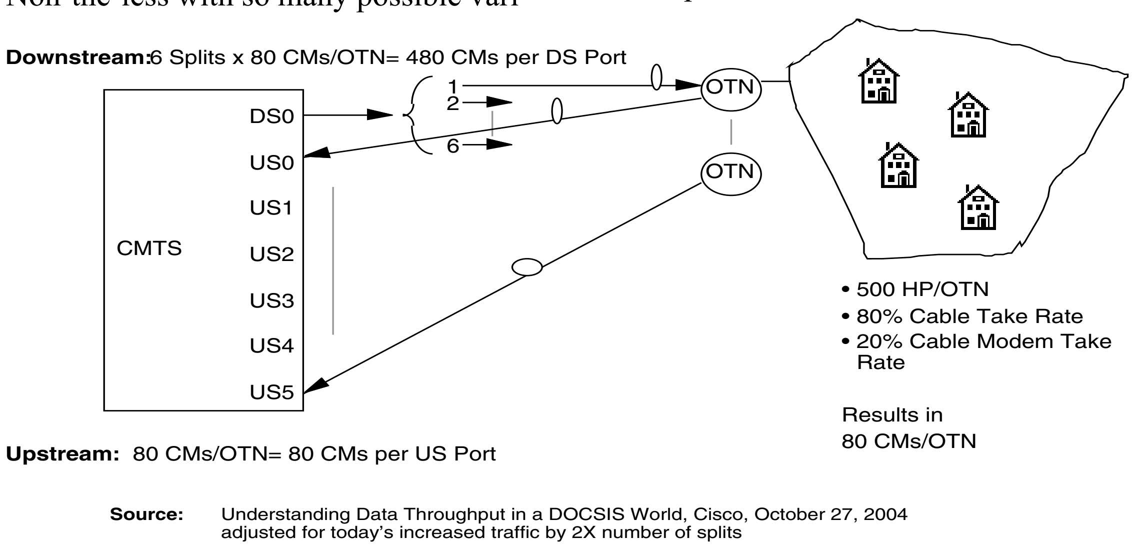

Figure 2. DOCSIS Over HFC Network Layout Assumed for Modeling in this Study

This ability to split ports at the CMTS or HP/OTN demonstrate the ability of an HFC based DOCSIS system to scale to meet traffic needs. However a key assumption made in this study methodology is that MSOs acting rationally will invest in adding capacity to their system when there is an economic justification to do so. However when third party service providers promote over-the-top services that might consume significant transmission resources, no rational MSO should make an uneconomic investment in adding capacity to support such over-the-top services.

The figure 2 network layout is the first of three categories of basic system assumptions I employed in the modeling and analysis. Figure 3 shows the assumptions made regarding raw and useful payload data rates on each CMTS port.

In any actual DOCSIS system down and up stream raw data rates can be set and the overhead can vary somewhat, but for the purpose of modeling, figure 3 payload parameters will be assumed as available for the purpose of carrying the various mixes of traffic studied.

	Modulation	Raw Data Rate	Overhead (note 1)	Available for Payload	Number of CMs/Stream
DownStream	64-QAM	30.3 Mbps	18.60%	25 Mbps	480
UpStream	QPSK	2.56 Mbps	18%	2.1 Mbps	80

Source: Understanding Data Throughput in a DOCSIS World, Cisco, October 27, 2004 adjusted for today's increased traffic by 2X number of splits

Notes **1)** Aprox dowstream overhead: Reed Solomon FEC=4.7%, Trellis Coding=6.7%, MPEG2=2.4%, Ethernet, DOSIS, IP = 2.8%, DOCSIS MAP=2%
Aprox upstream overhead: FEC=8%, Maintenance, acks=10%

Figure 3. Data Rate Settings for Downstream and Upstream CMTS Ports

Figure 4 contains CMTS parameters that are the third category of assumptions that are needed for the modeling I performed. None of the figure 4 parameters are fixed by DOCSIS but instead can be set by the

Assumption	Value
Bytes/MiniSlot	16
Contention MiniSlot Grants/ MAP under heavy load as a % of # of Upstream Modems/ CMTS Port	10%
Time Interval between MAP Msgs	2mSec
Interleaver Depth	32

Figure 4. CMTS Parameters That Impact Modeling and Limit Transmission Rates in this Study

operator and dynamically adjusted, if desired, by scheduling and resource management software running on the CMTS. In the simulation and modeling in this study, I picked the parameters in figure 4 to represent numbers that one might find during a period of high traffic intensity. Using these parameters as fixed in the simulations is justified since I am only interested in the behavior of the system when it is heavily loaded, so the operating characteristic the models predict should be valid during the heavily loaded system time of interest. Each CMTS supplier is free under DOCSIS to develop proprietary scheduling algorithms that attempt to maximize the user experience as a function of the traffic. During period of light traffic load, such a scheduling algorithm will adjust the ratio of contention to reservation minislots such that there are more contention slots available during light traffic loads to minimize latency in servicing best efforts contention upstream bandwidth requests, but during periods of heavy traffic loads the scheduler will assign more of the upstream transmission minislots to carry reserved payloads rather than be used for upstream bandwidth requests. The number of contention minislots for each MAP interval in this high traffic load modeling is set to 10% of the number of upstream modems per CMTS so as to balance efficient use of the upstream bandwidth but still leave a reasonable number of contention slots so that idle CMs can have opportunities to request bandwidth. A typical 2mSec is used in the modeling for the length of each MAP interval. This 2mSec inter MAP internal taken together with the 32 interleaver depth results in any one cablemodem needing to wait for every other MAP interval (i.e. 4mSec) to request bandwidth if the first contention request is not received. A interleaver depth of greater than 32 would increase latency and could require a cablemodem having to wait three or more MAP intervals before making a second request for bandwidth.

OFFERED TRAFFIC ASSUMPTIONS, QUALITY OF SERVICE, AND USER EXPERIENCE EXPECTATIONS

This section identifies the possible mixes of traffic on the DOCSIS system, the likely traffic intensity, and the different quality of service criteria that would impact the user experience. It is important to note that each traffic type will generally have very different QoS criteria and very different thresholds for acceptable versus unacceptable user experiences. For example some traffic types (e.g. web surfing) will maximize the user experience by providing very fast peak downstream data rates while the user will be generally insensitive to brief interruptions in the data stream. Multiplayer, *shoot first* games will care most about latency while at the other extreme public switched telephone network voice will care about the probability of blocking, latency, and packet loss.

Websurfing and Other Best Effort Data Traffic

The bulk of today's operational DOCSIS systems have been engineered for Websurfing traffic and that is the first traffic stream considered in this study. Although many general characteristics of this traffic are the same today as when systems were first launched, it is also the case that the magnitude of this traffic is rapidly growing. The general characteristics of websurfing traffic include:

- It is downstream intensive resulting in an asymmetrical traffic load between the down and up stream paths
- It is delay insensitive in that brief interruptions of the packet stream or short delays in serving upstream best-efforts contention requests for bandwidth will not be noticeable to the end user

Parameter	Downstream (note 1)	Upstream (note 2)
Average Busy Hour Data Rate/CM Subs (Active+Inactive)	29 kbps	7.25 kbps
Ratio Active/Total Subscribers	50%	50%
Peak Rate Limited	4.0 Mbps	384 kpbs

Source: MSO Website and conversation with an MSO

Notes 1) Downstream traffic growing about 33%/year
2) Upstream based on US/DS ratio of about 1:4
Upstream traffic growing about 25%/year

Figure 5. Websurfing Average Traffic Parameters

Figure 5 shows the assumptions employed in the modeling for this baseline web surfing load. Of particular note in figure 5 is that traffic is growing between 25% to 33% per year and that 50% of subscribers are assuming to be actively surfing during any measurement interval. The MSO interviewed indicated there is no well defined busy hour but instead this is an afternoon to early evening busy period of 4-6 hours in duration. With this in mind, the modeling assumes that the voice telephony busy hour will fall somewhere within the web surfing busy period assuring a 100% overlap in traffic. The rapid growth rate of web surfing traffic also implies a regular need to scale the system to accommodate growth. The average data rate shown in figure 5 was reported to me by an MSO based on actual system measurements. The inter-arrival time between web pages is given by the viewing time of a web page before the next page is requested and that parameter is shown in figure 6. Of note in figure 6 is a page viewing time of about 40 seconds but notice the very large standard deviation suggesting there is considerable variation about the 40 second average. The impact of this 40 second time for websurfing is that on the average of every 40 seconds, a subscriber in best efforts mode will need access to a contention minislot for the purpose of requesting bandwidth. As will be shown, this 40 second time is a comfortable many orders of magnitude larger than a much shorter interval for best effort telephony traffic.

Parameters	Mean	Standard Deviation	Best Fit Probability Distribution
Viewing Time (seconds)	39.5	92.6	Weibull

Source: A Behavioral Modem of Web Traffic by Choi and Limb, Georgia Institute of Technology, 1999 Proceedings of the Seventh Annual International Conference on Network Protocols

Figure 6. Websurfing Traffic Interarrival Times

As noted in the introduction to this section, the main performance measure impacting the websurfing user experience will be the peak limited downstream data rate. Of secondary interest to websurfing would be latency associated with sending upstream page requests that exceed on the order of a second. For email and file up and down loads, the peak data rates in both the up and down stream direction would most impact the user experience. Unlike websurfing however, latency would generally be unimportant. For user consumption of real time audio, video, or multimedia streaming downstream average data rates of less than 100 kbps are typical so the peak data rate is a less important factor; however, for this streaming content the key user experience impacting parameter is an interruption of the downstream longer than the buffer size (settable for most players generally in the range of 10-30 seconds).

Multiplayer games are an emerging important category of best efforts traffic. It was reported that one game, CounterStrike, was the third largest generator of UDP traffic on the internet behind only DNS and RealAudio traffic[1]. There are many variations of multiplayer online games including: peer-to-peer and client server variants. CounterStrike is a client server game and as the reference paper notes, the players consume an average of 40 kbps during the sessions that can last up to about 2 hours in duration. CounterStrike is in the popular category of a *shoot first* game which means that the key user experience impacting parameter will be response time. Packets are typically small on the order of 50mSec indicating that response times should be on the same order. In the case of a DOCSIS system, the mechanisms that could likely control response time would a lower bound on the interleaver depth and propagation delays to an upper bound set by congestion in upstream contention minislots.

Although best efforts data traffic on DOCSIS networks is more than just web traffic, since there were no measuring tools available to separate this best efforts traffic into categories such as web page downloads, email, file uploads, over-the-top multimedia, etc; for modeling purposes the traffic is assumed to be 100% web surfing in terms of average data rate. Also for accessing the user experience acceptable threshold, only latency delays associated with web page requests at inter-arrival time is considered further in this study.

PSTN Voice

As will be shown, the findings of interest in this study relate to interactions between best efforts websurfing discussed above with PSTN voice traffic. The PSTN voice traffic of interest is of two types: 1) QoS managed voice traffic that is provided by MSOs and 2) Best efforts voice traffic provided by third party service providers on an over-the-top basis. Assuming that the target user experience is the same for type 1) and 2) PSTN voice services, the payload data rates and user experience acceptability criteria will be the same. Figure 7 shows a variety of waveform and low bit rate source codecs and their raw data rates along with the resulting up and down stream payload data rates on the network.

Codec Type	Codec Rate	OPSK UpstreamRate	Downstream Rate	Packet Size
G.711	64 kbps	115.2 kbps	109.6 kbps	10 ms
G.728	16 kbps	57.6 kbps	61.6 kbps	10 ms
G.729E	12 kbps	57.6 kbps	57.6 kbps	10 ms

Source: White Paper: Engineering CMTS and HFC for VoIP with Capital and Operating Expense in Mind, Strater & Nikola, Motorola Broadband Communications Sector, December 2004

Figure 7. PSTN Voice Data Rates

Also shown is a typical number used in the modeling for the packet size. Longer packet sizes can provide more efficient use of the bandwidth resource but add latency and therefore numbers on the order of 10mSec are assumed as an acceptable tradeoff. In recent years the voice quality of low bit rate codecs has improved to the point that they are more than acceptable for voice only services. However, since the PSTN voice services of interest are those comparable and fully competitive with those offered by wireline local exchange carriers (LECs), for this study I assume the G.711 waveform codec which would assure quality comparable to an end office last mile wireline circuit and be compatible with fax and analog data modem traffic. As figure 7 shows, this implies payload rates on the DOCSIS system of 115.2 kbps upstream and 109.6 kbps downstream with 10mSec packet sizes.

While type 1) and 2) PSTN voice traffic have different QoS management and resource allocation mechanisms, for a comparable user experience they share the same underlying mechanisms that impact the user experience. Under the constraint of the same maximum acceptable round trip path delay of 300mSec, the user experience criteria fall into two categories[2]:

- Blocking of call attempts

  —This is typically managed by admission control in which a maximum number of voice circuits are reserved in order to keep the probability of blocking by a subscriber to <1.0%

  —For mathematical modeling purposes the Erlang B criteria was employed to compute blocking probability

  —Voice activity detection can be employed to improve system efficiency but was not assumed in this modeling

- Packet loss following call establishment

  —Once a call is established, excessive loss of packets will impair the channel. In the case of voice the conversation is interrupted and for fax and analog modem data, bits are lost

  —Packet loss can be compensated for by buffering and retransmission but this adds latency impacting path delay

  —Toll quality services typically aim to keep round trip path delay under 300mSec (150mSec end-to-end) and this was assumed as limit in the modeling

Figure 8 reveals the impact of packet loss rate on the user experience. Consistent with the <1% blocking criteria for a service offering comparable to a wireline LEC, I have chosen a maximum tolerable packet loss rate of <0.01% in order to provide toll quality voice while maintaining fax and analog data modem compatibility.

	Total Packet Loss Rate		
Feature	0.1% to 1%	0.01% to 0.1%	<0.01%
Voice Quality	Sub Toll	Toll	Toll
Call Completion Rate	98%	99.5%	99.99%
Fax	Dropped connections & error per page	Dropped connections & error per page	Analog Wireline Comparable
Dialup Modem	Dropped connections	Dropped connections	Analog Wireline Comparable

Source: White Paper: Engineering CMTS and HFC for VoIP with Capital and Operating Expense in Mind, Strater & Nikola, Motorola Broadband Communications Sector, December 2004

Notes Jitter is another important factor that could impact QoS in addition to packet loss rate

Figure 8. Impact of Packet Loss Rate on PSTN Voice End User Experience

In building the simulation model for the impact of packet loss on best effort over-the-top services, I also needed to understand the maximum contribution to latency from the uncertainty of timely grants of payload packet transmission requests in a contention minislot. Note that this is not an issue with respect to QoS managed MSO provided services since, the assumption is that the MSO provided services operate on an Unsolicited Grant Service basis where once the call is setup, the bandwidth is granted to the voice call for the entire duration without having to resort to using contention minislots for a bandwidth request. Figure 9 below shows a typical allocation of latency in a VoIP network by the contribution of each network segment. If the end-to-end delay limit is set to 105mSec, then the maximum allowable delay due to grant uncertainty would be 105-142.5=7.5mSec. As noted earlier, for the choice of interleaver depth and 2mSec MAP interval, a particular CM would be able to request bandwidth on a contention minislot only every other MAP interval resulting in 4mSec between requests. This in turn implies that if the delay contribution from grant uncertainty is held to 7.5mSec, then no more than two contention requests can fail before a packet will be lost.

Segment	Network	Component	Function	Nominal Delay (mSec)
Local	Upstream HFC Access	MTA	Voice Packetization	10.0
			DSP Operations	5.0
			Packet Encryption	0.5
		HFC	Upstream Trasmission	0.5
		CMTS	CMTS Forwarding	1.0
	Local IP Network	Routers		5.0
		Trunking GW	IP Processing	2.0
			RTP Decryption	0.5
			DSP Processing	3.0
			Jitter Buffer	15.0
			Subtotal Local	**42.5**
Long Distance	PSTN		Propagation Delay	100.0
			Total end-to-end	**142.5**

Source: White Paper: Engineering CMTS and HFC for VoIP with Capital and Operating Expense in Mind, Strater & Nikola, Motorola Broadband Communications Sector, December 2004

Figure 9. Typical Contribution to VoIP Path Delay by Function and Network Segment

MODELING AND ANALYSIS FINDINGS

Included within the scope of this study is two separate analysis which share in common the impacts of other traffic on best-efforts over-the-top PSTN voice services. As indicated above, if the goal is to provide an over-the-top PSTN voice service that is comparable to wireline LEC and MSO provided QoS managed VoIP, then the over-the-top service must not exceed the following two limits:

- The probability of blocking during a busy hour call attempt must be <1.0%
- Under a constraint of a maximum 150mSec end-to-end delay (e.g. 7.5mSec grant uncertainty contribution) the packet loss during a call must be <0.01%

Blocking Analysis

The blocking analysis was performed using a mathematical model based on Erlang B traffic theory. This theory is based on the statistical probability of finding available bandwidth (for DOCSIS reservation minislots) available given a finite number of voice circuits (i.e. those minislots assigned to a user for the call), a finite number of subscribers, and information on traffic intensity (e.g. Erlangs) during the busy hour of interest. The Erlang B theory is an approximation since it assumes infinite sources, Poisson arrivals, exponential holding times, and blocked calls cleared. The mathematics of this theory outputs the number of voice circuits required to maintain the level of blocking probability under 1% given the random arrivals of subscriber call attempts.

The Erlang B theory is employed in study as follows:

- Unlike websurfing that has an asymmetrical traffic load, voice telephony has a nearly symmetrical up/down stream data rate; therefore the upstream load of 80 CMs/CMTS port is used as the tight constraint

- The baseline load of figure 5 best efforts websurfing and other data traffic is assumed as being consumed and therefore unavailable to PSTN voice telephony users

- A busy hour traffic intensity of 0.1 Erlangs/subscriber was assumed[3] for computation of The Erlang B blocking at the 1% QoS

- In examining the impact of mixing type 1) MSO provided QoS managed telephony with type 2) third party provided over-the-top best-efforts telephony; the modeling assumes the QoS management system will give priority to MSO telephony. The resulting effect is that the combination of best efforts websurfing plus MSO QoS managed telephony traffic will consume upstream resources that will not be available for the over-the-top best efforts telephony

Employing the points above, what was studied was the impact of type 1) MSO telephony upon type 2) 3rd party telephony to see if they can both co-exist on a system that a rational MSO has engineered to be of a size that supports MSO provided services. Figure 10 provides the results of this analysis.

MSO Provided PSTN Voice	%of Downstream Resource Utilized	%of Upstream Resource Utilized	Downstream Limit % Adoption of BE PSTN Voice	Upstream Limit % Adoption of BE PSTN Voice	Constraint Limit % Adoption of BE PSTN Voice
0%	56%	33%	42%	13%	13%
1%	59%	33%	39%	13%	13%
2%	60%	33%	38%	13%	13%
3%	62%	54%	36%	8%	8%
4%	63%	54%	35%	8%	8%
5%	64%	64%	34%	3%	3%
6%	66%	64%	33%	3%	3%
7%	67%	64%	32%	3%	3%
8%	68%	69%	30%	3%	3%
9%	69%	69%	29%	3%	3%
10%	71%	80%	28%	0%	0%

Figure 10. DOCSIS Capacity for Other-the-Top versus MSO Provided PSTN Voice at 1% Blocking

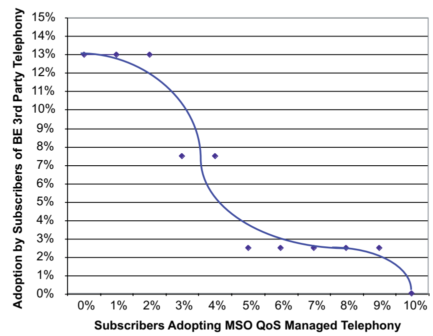

Figure 11. Graph of DOCSIS Capacity for Other-the-Top versus MSO Provided PSTN Voice at 1% Blocking

The left most column in figure 10 shows the percent adoption by the 80 subscribers on the CMTS port of MSO provided PSTN voice services. Notice that at zero percent adoption, 33% of the upstream resource is already utilized by just best efforts websurfing and other data traffic. At this starting level of no MSO provided PSTN voice service, the system will accommodate up to 13% adoption of over-the-top PSTN voice (13% computes to 10 subscribers on this 80 CM/CMTS port system). As the number of MSO subscribers grows from zero the assumption is that the CTMS and Packet Cable QoS mechanisms will give priority to the MSO traffic and the number of voice circuits available for the 3rd party over-the-top services will need to be limited. Figure 10 shows that as the MSO provided PSTN voice adoption grows to 10% (10% computes to 8 subscribers), there are not enough voice circuits available to support a 1% blocking rate for the over-the-top PSTN voice service provider. Figure 11 shows the shape of the operating characteristic by fitting a curve to the figure 10 data. Note the steep slope of the curve at the threshold as the MSO provided service adoption approaches 9%.

Lost Packet Analysis Theory

The blocking analysis above is but one of two impairments considered in this study. The second analysis examined the impact of lost packets in over-the-top PSTN voice services from 3rd party providers. Unlike the QoS managed PSTN voice services from MSO providers, the mechanism for requesting and receiving bandwidth grants for over-the-top providers is assumed to be best-efforts, as available, bandwidth. Unlike the previous blocking analysis which examined the limited up and down stream reservation minislots carrying payload traffic, this lost packet analysis looks instead at the bandwidth allocation process as the primary mechanism leading to lost payload packets.

For this analysis, the assumption is that the DOCSIS network is heavily loaded and the minimum number of contention minislots possible is allocated to maximize payload carrying reservation minislots. A well designed CMTS scheduling algorithm would reduce the actual number of contention minislots to the lowest level suitable to support the normal best-efforts data (mostly websurfing) traffic.

I further assume that MSO provided voice telephony for subscriber originated calls occurs at the very low rate of 0.75 call attempts during the busy hour.[4] As figure 6 shows, the mean time between best-efforts requests for bandwidth on contention minislots for websurfing is on the order of 40 seconds. On the other hand, figure 7 shows the packet size for PSTN voice is 10mSec. Assuming that this over-the-top 3rd party PSTN voice is treated just like other best-efforts data traffic such as websurfing it will require timely access to a contention minislot on the order of every 10mSec to avoid packet loss. While lost bandwidth requests made on a contention minislot can be repeated there is a limit on acceptable latency that would imply a finite buffer size so that only a limited number of bandwidth request repeats are possible prior to losing a packet. It is further assumed that no attempt will be made to give over-the-top best efforts PSTN voice any priority mechanisms (e.g. will disable piggyback requests in request/transmission policy) over less frequent websurfing traffic.

To arrive at an estimate for lost over-the-top PSTN voice lost packets my model reduces to looking at how many of the upstream contention minislot requests for bandwidth are received without corruption from contention for the minislot. Since there is a maximum end-to-end latency limit of 150mSec, there will also be some limit on how many sequential contention minislot bandwidth requests can be corrupted prior to losing the packet. To understand how many requests in sequence can be lost, I examined the total contribution to end-to-end latency as shown in figure 9.

Of the maximum limit of 150mSec end-to-end delay there is a maximum of 7.5mSec available for grant uncertainty delay in the contention minislot mechanism. In this system, using the figure 4 parameters, a missed contention request must wait for every other 2mSec

MAP interval; therefore only two sequential contention minislot grant requests can be corrupted before the packet is lost.

Lost Packet Analysis Model

The Extend™[5] discrete event modeling software package was employed to build a simulation model to determine the traffic intensity that would cause a specific over-the-top best-efforts PSTN voice subscriber under test (SUT) to experience a packet loss in excess of the toll quality limits of 0.01%.

This modeling was performed at a macro level in order to reduce the computational complexity so that alternative parameters could easily be explored. In this context a "macro level" means that there was not an attempt to model the performance of the system at the detailed level of operation of each logical elements and decision process in the real system. Instead a number of lower level logical elements are aggregated into large macro blocks that are designed to provide a good estimate of the actual system.

The analysis was based on the following:

- Only the upstream was modeled assuming 80 cablemodems per OTN per CMTS port
- Minislot size of 16 bytes
- Time interval between MAP messages of 2mSec
- With interleave depth of 32, a subscriber under test (SUT) would have best effort access to every other MAP at 4mSec intervals
- The over-the-top VoIP service would make a best effort (BE) request for each packet

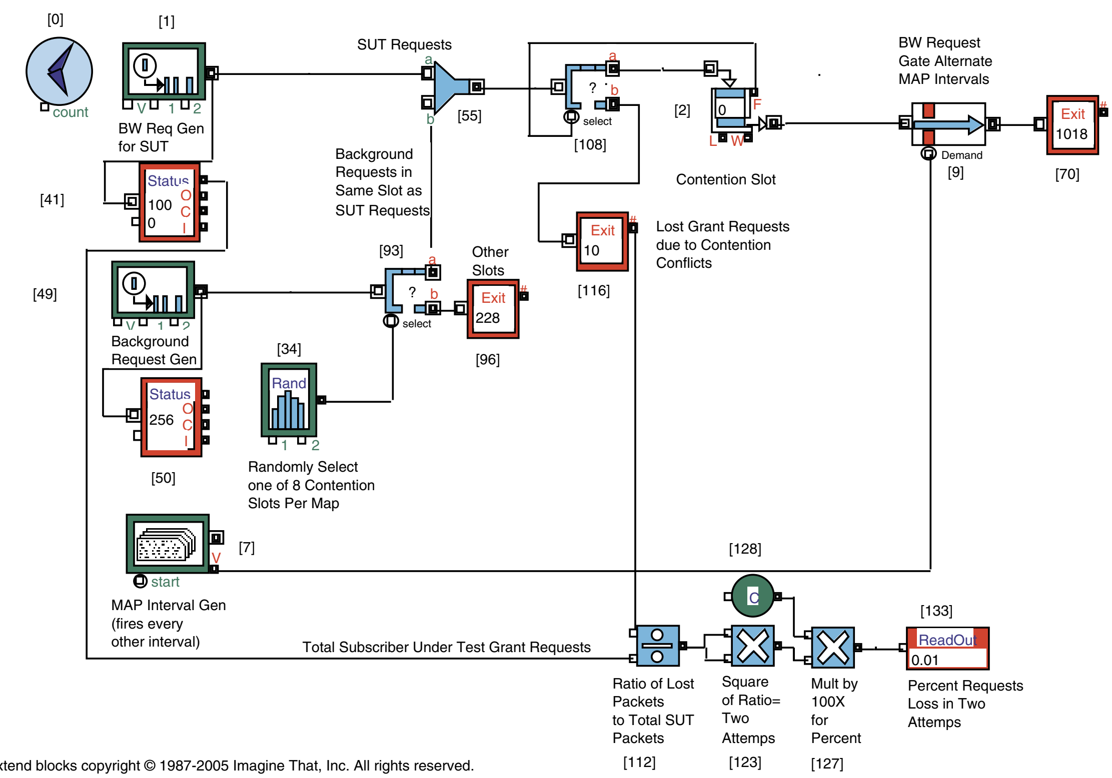

Figure 12. Discrete Event Model to Estimate Lost Packet Rate in Over-the-Top PSTN Voice Service

- The MSO provided VoIP service would be accommodated using unsolicited grant service with 8 minutes between call attempts and would therefore would have minimal impact on traffic in the bandwidth request contention slots

- The busy hour traffic is 0.1 Erlangs with 8 minute average call duration

The model architecture consisted of:

- Examination of a single SUT to determine the impact of other BE traffic upon the packet loss rate to the SUT

- There is contention between one SUT traffic generator and a background traffic generator

- A maximum of two bandwidth request attempts is permitted before the packet is lost in order to maintain low latency

Figure 12 shows a block diagram of the simulation model. This model operates as follows:

- Block [1] represents grant requests every 10mSec by a single subscriber under test (SUT) for bandwidth on a single contention minislot represented by a queue of length one in block [2]. The 10mSec interval is given by the 10mSec packet length and for best-efforts service the bandwidth must be requested after each packet.

- Block [49] represents background contention minislot requests by other over-the-top subscribers who have active calls that overlap with the SUT. The arrival rate varies as will be explained based on the number of other active calls.

- Block [7] generates the timing for 4mSec map intervals and controls the gate block [9]. If there is only one non-corrupted request in the contention minislot, the bandwidth request will get through the gate to the exit block [70]

- Any background requests from block [49] that make their way through block [55] to the contention slot in block [2] will block the SUT request resulting in the SUT request switched by block [108] to the lost requests to block [116] simulation exit.

- Note that not all background bandwidth requests from block [49] make their way to contend with the SUT. Since there are 8 contention minislots per map interval, on an equal random basis, one in eight background grant requests will fall into other slots. The random generator block [34] working with the switch block [93] configures the switch to route only one in eight background grant requests to the same contention slot as the SUT.

- The lost SUT grant requests in block [116] are divided by the total SUT grant requests from block [41] to compute the ratio of lost grant requests to total grant requests. To compute lost packets the loss grant ratio is squared to account for a maximum of two requests by block [123] and converted from a decimal to a percentage by block [127] for readout of lost percent packets in block [133].

All of the parameters necessary to exercise the model have been stated with one exception. The remaining parameter is the block [49] background traffic contention grant request arrival rate. This parameter is in turn controlled by the expected number of simultaneous over-the-top calls that overlap with the SUT. Figure 13 shows the results of estimating these overlaps as a function of the number of over-the-top subscribers on the upstream CMTS port. Recall that this system supports a maximum of 80 CM subscribers, so figure 13 explore the overlaps from 2% to 100% (i.e. 1 to 80 subs). By employing the previously stated assumptions of 0.1Erlangs/sub during the busy hour, and a mean call length of 8 minutes; a single background subscriber active during the 60 minute busy hour would present a mean 0.11 number of average simultaneous background calls to the SUT. The background grant request rate parameter entered into the model block [49] background traffic generator is shown in the next to last column in of figure 13. This number is computed as twice the 10mSec packet rate divided by the average simultaneous calls. The basis for the factor of two, is that the SUT MAP interval is 4mSec (i.e. every other MAP) but the background traffic will be evenly distributed between every 2mSec MAP interval. Therefore the SUT will never be corrupted during half of the MAP intervals.

Number of BE VoIP Subs	Total Number of Background Calls Placed During the Busy Hour	Average Simultaneous Calls	Calculated Call Arrival Rate per 4mSec Interval	Subscriber Adoption Percentage
1	1	0.11	182	2%
6	4	0.51	39	7%
27	20	1.69	12	33%
40	30	2.92	7	50%
53	40	4.15	5	67%
67	50	5.18	4	83%
80	60	5.83	3	100%

Figure 13. Estimation of Average Number of Simultaneous Over-the-Top PSTN Voice Calls During the Busy Hour

Recall that the background traffic contention grant requests are not only spread randomly between each of the two 2mSec SUT MAP intervals, but they can also be expected to be uniformly distributed between the eight contention minislots during each MAP interval. This means that for the call arrival rate from background traffic to become comparable to the SUT traffic in the contention minislot, the overlapping background subscribers need to number about 16 relative to the SUT.

Lost Packet Analysis Findings

The model was exercised using the figure 13 parameters at four levels of over-the-top subscriber adoption. Since the percent packet loss threshold is a small number, on the order of 0.01%, the model was set to run for ten thousand steps. It appeared to converge at about to three thousand steps to the stable findings as shown in figure 14.

% of 80 Subscribers who adopt over-the-top PSTN Voice	Average Number of Active Calls Overlaping with SUT	Background Traffic Contention Request Rate per MAP in msec	% Probability of Packet Loss for SUT After TwoContention Attempts to Request Bandwidth
7	0.51	39	0.01
33	1.69	12	0.07
50	2.92	7	0.53
100	5.83	3	3.31

Figure 14. Results of Discrete Event Modeling to Estimate Lost Packet Rate in Over-the-Top PSTN Voice Services

The model results suggest that the packet loss rate approaches the figure 8 toll quality limit (i.e. OK for voice but unsuitable for fax and analog modem data) at about 7% adoption of over-the-top PSTN voice. This 7% adoption level corresponds to a mean number of about 0.51 background calls that overlap with the SUT. Just above 50% adoption by over-the-top subscribers, the packet loss rate shown in figure 14 approaches the 1% level of noticeable voice quality impairments.

OPPORTUNITIES FOR FURTHER WORK

This study focused upon the interactions between best efforts data traffic, MSO QoS managed PSTN voice, and third party over-the-top best-efforts PSTN voice. As mentioned the best efforts data traffic will be a mix of delay insensitive websurfing and other delay sensitive and insensitive other data traffic (e.g. streaming music, videoconferencing, file uploads, multiplayer games). An obvious extension of this study would be to attempt to understand the nature and impact of other than websurfing best efforts data traffic. It can be expected that there may well be significant impairments to best-efforts PSTN voice due to large peak downstream data bursts in addition to the upstream resource limitation impacts exposed by this study.

Other areas for additional study would involve refinement of the discrete event model to increase the level of complexity from the high macro level to one where the random nature of some of the input parameters are explored. In particular, the parameter employed in the model for simultaneous background calls was based on the offline (i.e. outside of the model) computation of the mean number of simultaneous calls. A model refinement would use not just the mean but would explore the probability distribution to yield a packet loss rate estimate at a specific probability (e.g. to address what level of adoption would keep packet loss rates under 0.01% for 90% of the time).

CONCLUSIONS

There is significant promotion by over-the-top providers of PSTN voice services who suggest these services can be used in a best efforts mode on DOCSIS cablemodem networks. At the same time these over-the-top providers are promoting their services, MSOs are rapidly rolling out QoS managed PSTN services. Both the MSO and over-the-top providers are representing their services as toll quality comparable to the wireline services now offered by incumbent local exchange carriers (LEC).

However, if a rational MSO economically dimensions their DOCSIS service to meet their service obligations for websurfing and MSO provided PSTN voice, it is unlikely that third party over-the-top PSTN providers will be able to maintain LEC comparable quality of service.

There appears to be two mechanisms that will impair the quality of over-the-top PSTN voice. One mechanism derives from limited upstream reservation (i.e. traffic payload) minislots. For this mechanism, as MSO provided PSTN voice grows, there will not be enough voice circuits for over-the-top PSTN voice to meet the <1% probability of blocking QoS criteria. The analysis findings for the system examined in this study suggests that there will be a steep decline in over-the-top PSTN voice service when the MSO traffic levels reach adoption levels as low as 9% of cablemodem subscribers.

The second mechanism impairing over-the-top PSTN voice is independent of the growth of MSO provided PSTN voice. Unlike the limited reservation minislot mechanism above, this second mechanism is controlled by limited contention minislots available for a best-efforts bandwidth grant requests. For the example system modeled, it appears that an over-the-top PSTN voice service in best efforts mode will experience packet loss that degrades voice quality when there are about 3 simultaneous other active calls competing with one subscriber under test. Given typical residential subscriber call arrival rates and holding time, the average adoption rate at which the impairment becomes noticeable occurs at about 50% adoption of over-the-top PSTN voice by cablemodem subscribers.

ENDNOTES

[1] "A Traffic Characterization of Popular On-Line Games" by Feng et al, IEEE/ACM Transactions on Networking, Vol 13, No. 3, June 2005

[2] The source for these parameters is the same as the source note shown in figure 8.

[3] "Multimedia Traffic Engineering for HFC Networks", Cisco Systems, 1999 states typical traffic intensity between 3-6 CCS, so a value of 0.1Erlang was selected at the conservative low end between 3-4 CCS

[4] ibid, "Multimedia Traffic Engineering for HFC Networks", Based on of 0.1 Erlangs of traffic during the busy hour and mean residential call duration of 8 minutes. 0.1 Erlangs is average 6 minutes of busy hour traffic/subscriber of 8 minutes duration is 0.75=(6/8)

[5] Extend™ is a trademark for discrete event modeling software claimed by the firm of ImagineThat.

HIERARCHICAL INTER-CMS ARCHITECTURE USING STANDALONE SIP ROUTE PROXY

Rupal Desai and Rob Kissell
Cisco Systems, Inc.

Abstract

Cable multiple service operators (MSOs) are rapidly ramping up their VoIP subscriber base. As they deploy regional Call Management Servers (CMSs), they also deploy costly interconnects with the PSTN for routing calls either to another CMS or to the PSTN. A hierarchical inter-CMS architecture using a standalone SIP route proxy enables MSOs to avoid the unnecessary PSTN interconnection costs and gives them flexibility to deploy CMSs regionally in a way that best utilizes their network resources. This paper will discuss the role of a SIP route proxy in interconnecting CMSs, media gateway controllers (MGCs) and peer VoIP providers and illustrate the business case behind deploying one such architecture.

INTRODUCTION

Voice has become an integral part of triple play offered by the Cable multiple service operators (MSOs). According to the latest data, North American MSOs have announced over 3M VoIP subscribers and have plans to add many more. With this blistering growth comes the complexity of scaling the network. Also, Cable MSOs are looking for ways to keep voice calls on their IP network which not only improves voice quality but also significantly reduces call-termination costs. Additionally, offering multimedia services to enhance subscriber experience is also a priority. Hierarchical Inter-CMS architecture can not only be used to effectively scale the network but also to provide policy-driven PSTN interconnect strategy which can significantly reduce operating costs and help evolve MSOs network to PacketCable 2.0 / IMS architecture.

TODAY'S ARCHITECTURE

In order to fully understand the potential of the Inter-CMS architecture, let's have a look at the current architecture deployed today. Cable MSOs have deployed VoIP using the integrated softswitch architecture. This architecture is composed of CMS and MGC functionality, as defined by CableLabs™, into one single element. The key benefits of the current architecture are: scalability (compare to TDM switches), significant reduction in capital expenditure and the ability to introduce features at a reasonable cost. Most of the VoIP deployments by MSOs today are based on the PacketCable compliant multimedia terminal adapters (MTA) which can communicate with softswitches via network call signaling (NCS) protocol. The generic high-level diagram below shows the current VoIP architecture.

As shown in Figure1, the CMS and MGC are combined in a single element known as softswitch. When a call is received by a softswitch, it determines, from its internal database, the type of call. If the call type is long-distance, then it routes the call out to PSTN interconnection trunks. If the call is local, then it terminates the call to itself. In this scenario all off-net calls are routed out using PSTN interconnection. Even the calls between some subscribers served by the same MSOs have to traverse through PSTN.

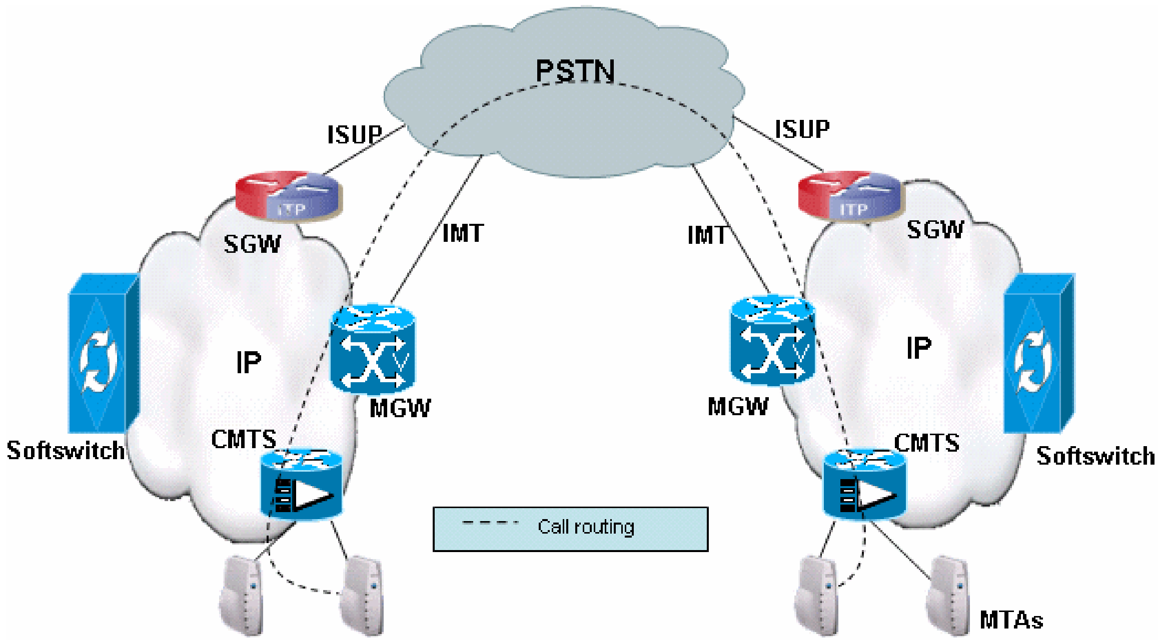

Figure 1. PacketCable™ architecture without SIP routing

This architecture has allowed MSOs to quickly and profitably deploy VoIP service to tens of thousands of subscribers in different regions. Multiple softswitches have been deployed in multiple points of presence (PoPs). To connect the VoIP subscribers in different PoPs and to the PSTN, most MSOs use inter-machine trunks (IMTs) for bearer connectivity and SS7 network for call signaling. This approach is reliable and widely deployed today but it has few shortcomings.

1) High cost associated with terminating calls to PSTN. The inter and intra lata calls must traverse the PSTN before reaching their destination. Also, MSOs have to pay a huge sum of money to buy the port capacity on media gateways and connecting signaling links and linksets to the SS7 network.

2) During interconnect, packet-to-TDM and TDM-to-packet conversion adds processing delays and hence affects the voice quality.

3) TDM and SS7 facilities are dedicated and can't be used for anything else but Voice traffic.

4) No sharing of ILEC interconnection between softswitches is allowed.

One way to potentially overcome the above mentioned shortcomings is to deploy SIP trunk signaling between softswitches. Most of the softswitches today support standardized version of SIP protocol – RFC 3261. By connecting softswitches with SIP trunks, it eliminates the need to interconnect to PSTN for "on-net" calls which not only improves the voice quality but also drives down the call terminating cost since fewer PSTN interconnections are required. The diagram below shows the SIP trunk signaling between softswitches.

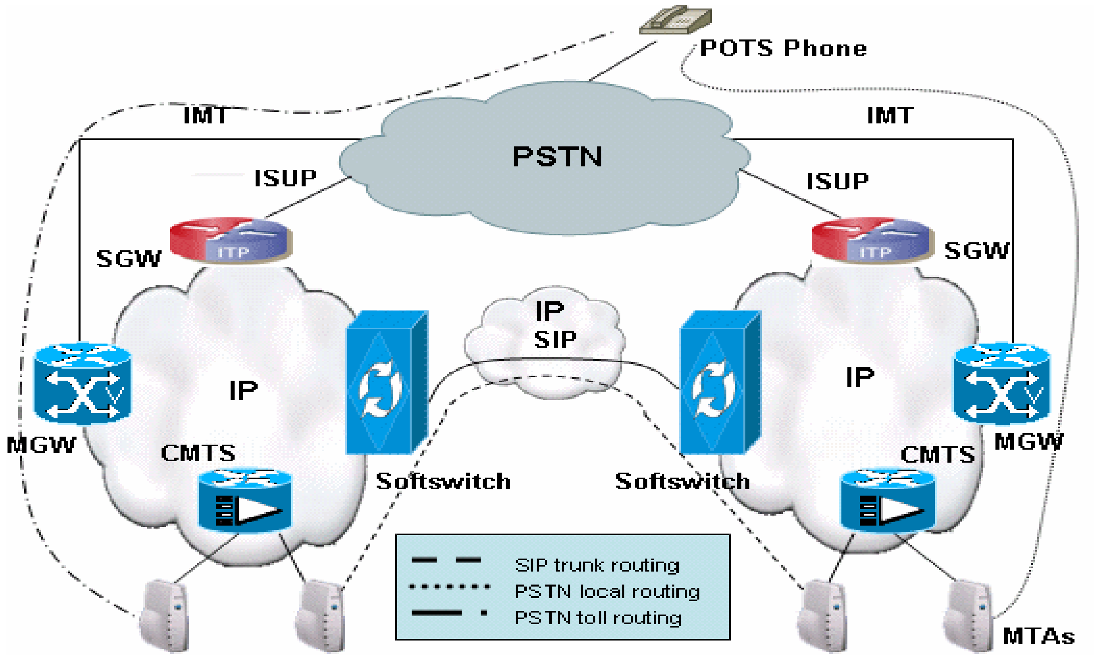

Figure 2. PacketCable™ architecture with SIP trunk signaling

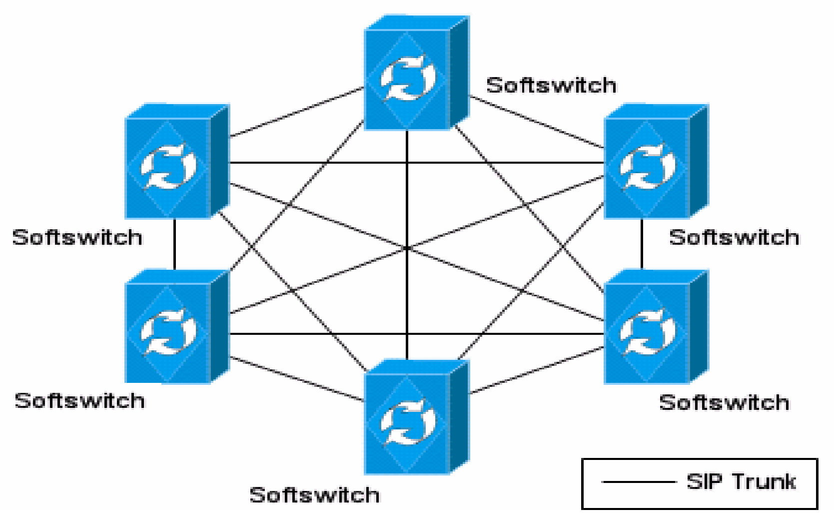

Figure 3. Full Mesh routing using SIP trunk between softswitches

As shown in the Figure 2, when a call is received by a softswitch, it determines, from its internal database the type of the call. If the call type is long-distance, then it routes the call out to PSTN interconnection trunks only if the call is terminating to a foreign subscriber (PSTN toll routing in Figure 2). But, if the call is terminating to MSO-owned subscriber then call is routed out to SIP trunk connected to terminating softswitch (SIP trunk routing in Figure 2). In this scenario not all "on-net" calls have to use PSTN interconnection for routing which results in operational cost savings.

Many MSOs have deployed SIP trunk signaling between softswitches. As this solution reduces the need for PSTN interconnections and improves voice quality, it raises other operational and technical dilemmas. MSOs have deployed multiple softswitches in a PoP and multiple PoPs across different geographic locations to support the growing demand of VoIP. This approach requires a fully meshed network – connecting every softswitch not only to every other softswitch in a PoP and to softswitches in other PoPs as well. Figure 3 below shows a fully meshed network.

As it is apparent from Figure 3, fully meshed networks are not scalable due to the complexity involved in configuring and maintaining large networks. For every n CMSs in an MSO's network, they'll need n times (n-1) – i.e. n-squared – number of trunks. Provisioning these trunks as the network grows to 10 CMSs and beyond can result in an unduly burdensome number of man-months or even man-years spent provisioning trunks on each new CMS added. This is to say nothing of continually managing and modify translation tables etc. as numbers or number ranges served by each CMS are continually modified.

Additionally, and more importantly, some of the formerly co-operative inter-exchange carriers have been acquired by the cable operators' competitors. Cable operators are therefore searching for ways to define new network architectures that provide operationally viable and cost-effective alternatives for both CMS deployment and PSTN interconnect to deliver high quality VoIP services to their subscribers. Introducing tandem-like hierarchical means of routing calls to and from NCS endpoints using stand-alone SIP route proxy could very well solve all the operational, technical and economical issues that MSOs face today.

INTER-CMS ARCHITECTURE USING STANDALONE SIP ROUTE PROXY

Hierarchical inter-CMS call routing using standalone SIP route proxy allows MSOs to more intelligently and economically route calls on their network. As we will see in later sections, inter-CMS architecture, despite its incremental capital cost, offers steep OpEx savings. But, let's first take a look at how a SIP route proxy can be utilized to offer highly-available, scalable, centralized and intelligent solution to enhance MSOs PSTN strategy.

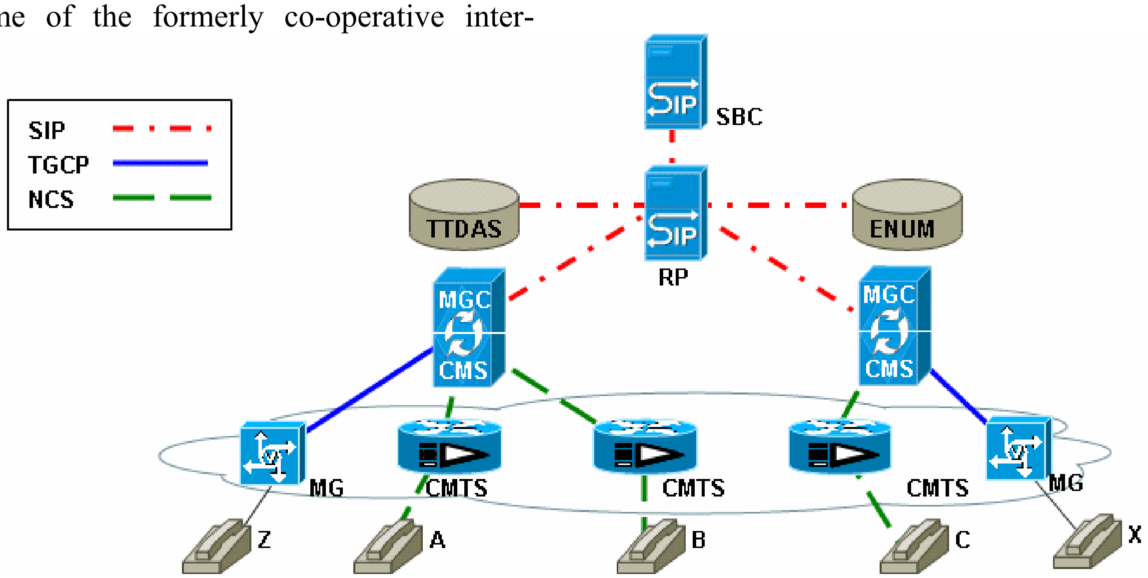

Figure 4. Inter-CMS architecture using SIP route proxy

The figure below shows a generic Inter-CMS architecture. It consists of a standalone SIP route proxy, external database, for e.g. TTDAS, and an ENUM server along with all the required components of PacketCable architecture.

With the help of SIP route proxy, call routing can be accomplished in variety of ways as described later in this section. This gives traffic engineers greater flexibility as to how different types of calls can be routed in the network. By gaining control over call-routing, MSOs can achieve two exceedingly important goals of keeping calls on-net and avoid costly PSTN interconnections. For example, in the diagram above, call routing can be accomplished in following ways:

<u>On-net routing:</u>

<u>Local:</u> User A calls B. Since there is IP connectivity between A and B and both are on the same CMS, the CMS route from A to B is NCS to NCS.

<u>Long-distance:</u> User A calls C. There is IP connectivity between A and C but they are on different CMSs. The CMS route table for A routes the call to SIP Route Proxy (RP). The SIP RP routes the call to the CMS that C is attached to.

<u>Off-net Couting:</u>

<u>Local Origination:</u> User A calls POTS Z. CMS determines that call is local and hands the call off to MGC serving the local market. If all the PSTN trunks are currently in use then a route advance is performed and the call is forwarded to closest MGC using the SIP RP.

<u>Long-distance Origination:</u> User C calls POTS Z. CMS determines that call is long-distance and routes the call to SIP RP. SIP RP determines the nearest MGC and hands the call off to that MGC for call termination.

<u>Local Termination:</u> PSTN subscriber Z calls user A. MGC determines that call is terminating to directly-connected CMS and thus hands the call off to that CMS.

<u>Long-distance Termination:</u> PSTN subscriber Z calls user C. MGC determines that call is not local so, routes the call to SIP RP. SIP RP determines which CMS to route the call and hands the call off to that CMS.

As it is evident from the examples above, the SIP route proxy facilitates the routing between CMSs and supports various routing options for calls destined to/from PSTN. All different call types such as local, toll, and long-distance can be routed by SIP route proxy such that calls can be retained on-net for better voice quality and reduced operating expense.

<u>IMPLEMENTING INTER-CMS ARCHITECTURE</u>

The first phase in implementing Inter-CMS architecture is logically and physically separating softswitch into CMS and MGC as shown in Figure 5. CMS is used for line side features such as controlling MTAs via NCS protocol, providing subscriber features and accessing subscriber database. MGC is used for trunk-side features such as outbound route selection, PSTN interconnection and providing PSTN features such as LNP.

Once the CMS and MGC are separated, the next step is to introduce a SIP route proxy in the middle. Figure 6 shows a standalone SIP route proxy that can route calls directly between CMSs or between CMSs and multiple MGCs in the network.

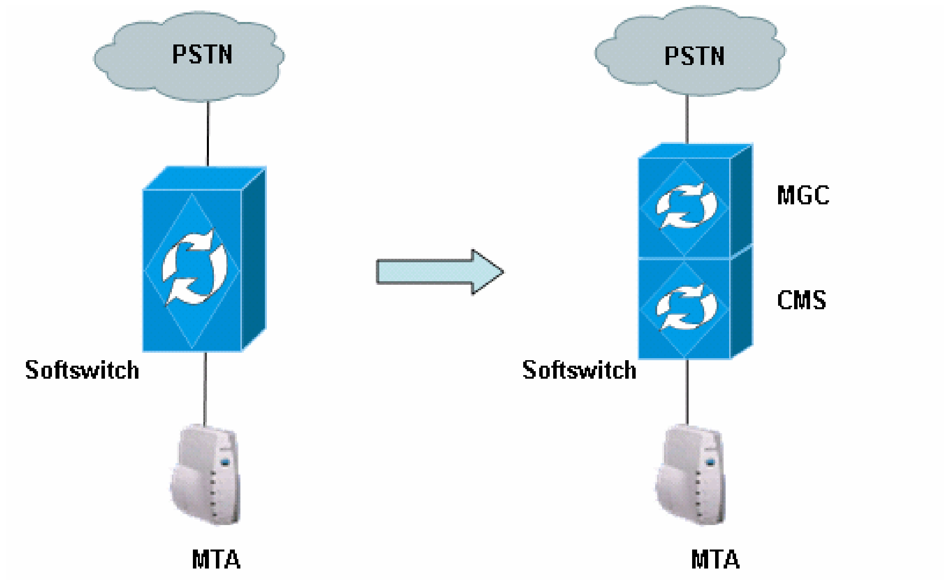

Figure 5. Step 1: Separate Softswitch in CMS and MGC

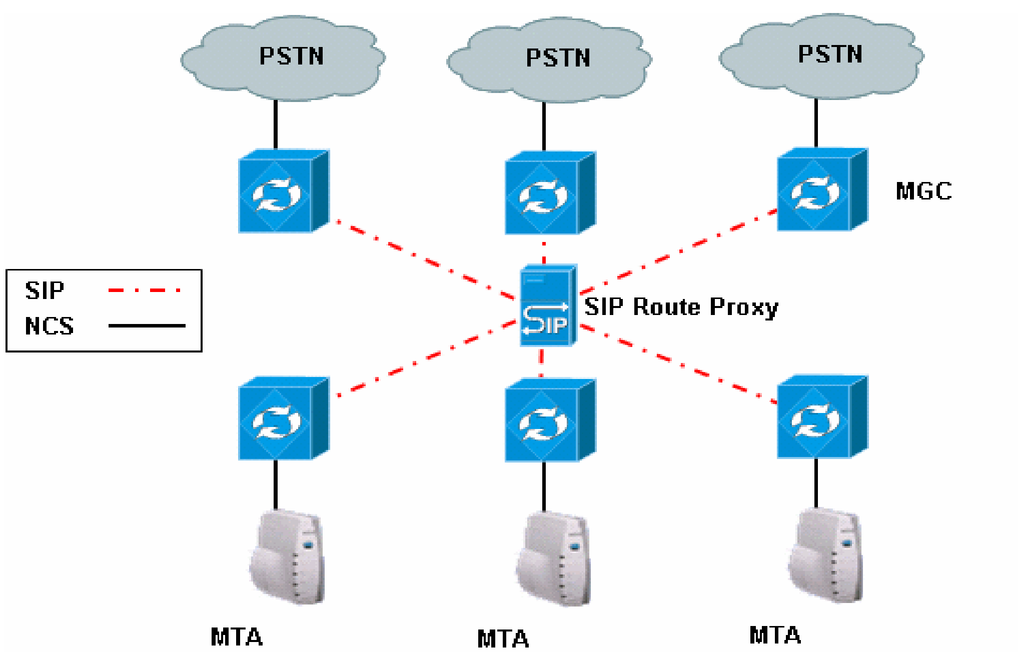

Figure 6. Step 2: Deploy SIP Route Proxy

The SIP route proxy performs multiple roles in the network –

1) It allows PSTN interconnections to be shared by multiple CMSs such that CMS on the left, in the figure above, can route calls to PSTN by interconnecting with MGC on the right or on the left depending upon the least cost routing logic set in the SIP route proxy.

2) Since CMS and MGC are separated in two different physical

components, they can be scaled independent of each other. You can have difference number of MGCs, depending on number of PSTN interconnections required, then CMSs which depends on growth of VoIP subscribers.

3) The SIP route proxy can be used to provide policy-driven IP interconnection strategy to Peer MSOs' or partner IXCs' networks. The IP interconnections can be managed by connecting SIP route proxy with the session border controllers (SBC).

Routing Logic and Capabilities

The SIP route proxy is typically a centralized function in the network. It can be shared by multiple PoPs and is usually deployed as a cluster of servers for high availability. Its primary role is to handle the routing of calls such that calls between MSO-owned subscribers can be kept on-net and for all other calls needing PSTN interconnection, can be routed based on least cost option available. Routing requests can be received from any element in the network which either do not provide this resolution capability, or in cases where operating policy dictates this logic be centrally located and administered.

SIP route proxy can be configured to do ENUM lookup on every request it receives. When a request is received by SIP route proxy, it queries the ENUM database to determine if the call is to a MSO-owned subscriber. If that is the case then it routes the call to the CMS responsible for the subscriber. Otherwise it can route the call either to a session border controller (SBC) or to a local media gateway controller (MGC) depending on the type of desired PSTN interconnection. For terminating calls in the network, SIP route proxy can query the external database such as TTDAS. When a request is received, SIP route proxy can query the database on NPA-NXX of the called number to determine and to route the call to CMS attached to that subscriber.

SIP route proxy can also be used to route advance in case of congestion or downstream failure. When a failure response is received, SIP route proxy can attempt to route the call over to the alternate PSTN interconnect point such that the call can be successfully completed.

Enhanced Routing

As the subscriber base of VoIP service grows, MSOs are not only witnessing increase in "on-net" traffic but the inter-MSO traffic is growing as well. Hence, MSOs are seeking for ways to directly route calls between their networks to reduce OpEx. SIP route proxy combined with ENUM server can be effectively used to route calls between providers' networks. All "off-net" calls can be routed to SIP route proxy which can determine, after an ENUM query, where the call needs to be routed – either to a TDM interconnect point or to a directly connected peer-MSO network.

Besides reducing the OpEx, SIP route proxy can be used to effectively manage and in some cases deploy new revenue generating services. For e.g. Virtual Number service which provides its subscriber the freedom to choose the area code of their choice in return for a fixed subscription fee. SIP route proxy, upon receiving a request to terminate a call to the Virtual Number, performs ENUM query and determines the actual number and routes the call to the CMS the subscriber is attached to.

Evolution to PacketCable 2.0 Architecture

Cable operators today are not only thinking about reducing OpEx but also are looking to evolve towards next generation architecture to offer their subscribers revenue-generating SIP based multimedia services.

For today's MSOs, it is not the question of "if" but "when" and "how" they will move from PacketCable 1.X architecture to tomorrow's PacketCable 2.0 / IMS architecture. It is important for MSOs to achieve this evolution through addition of individual capabilities that solve a business need, such as inter-CMS call routing, and also add components of the target architecture. Gracefully introducing individual pieces of next-generation architecture is a key to successful network evolution. Inter-CMS architecture, if you will, is the first step in that evolution process. By adding SIP route proxy, which can function as interrogating call session control function (I-CSCF) or border gateway control function (BGCF), and separating CMS and MGC, where MGC can perform media gateway control function (MGCF), completes the first step in evolving towards the next generation architecture.

Adding an I-CSCF function to a cable company's VoIP network becomes increasingly important as MSOs look to evolve and deploy parallel SIP-based voice (e.g. business line services to SMB customers) and SIP-based multimedia (e.g. video phone services to consumer customers) service networks leveraging their current deployed PacketCable infrastructure for PSTN interconnect, 911 calling, etc. These services need a common point of attachment – or the full meshing problem described above becomes infinitely worse – and all services need a way of routing calls or "sessions" that are addressed to a SIP URI (e.g. bob@anycableco.net) as well as to e.164 telephone numbers. Cisco believes these should be capabilities of any Route Proxy added to MSO's networks from Day One.

BUSINESS CASE

To demonstrate the benefits of pursuing inter-CMS routing via a standalone SIP route proxy, a business case has been developed. This business case considers alternative deployment scenarios for routing traffic between CMSs, thereby keeping more calls on-net. This business case captures the investment (both CapEx and OpEx) required to implement each scenario and details the resulting economic savings for each.

In total, three scenarios were evaluated

- Case A – No Inter-CMS Routing
- Case B – Inter-CMS Routing via Meshed SIP Trunk Network
- Case C – Inter-CMS Routing via Standalone SIP Route Proxy

The economic analysis for the scenarios is comprised of the following factors:

- CapEx: Investment in TDM gateways, session border controllers, SIP route proxy servers, router ports required for VoIP backhaul between PoPs
- OpEx: IP circuit expenses for backhauled traffic, PSTN interconnection charges, and deltas in operational staffing expenses. (Note: capturing full operational costs for running a VoIP network is not part of this exercise, only the

	Emb Base	Year 1	Year 2	Year 3	Year 4	Year 5
Subscribers						
Number of Subscribers	500,000	600,000	720,000	864,000	1,036,800	1,244,160
Network Layout						
Number of PoPs	5	5	5	5	5	5
Total Number of CMS	9	10	12	15	18	21
Call Volume						
Total Minutes of Use (M)	5,475	6,570	7,884	9,461	11,353	13,624
Calls per Second (BH)	278	333	400	480	576	691
Simultaneous Calls (BH)	50,000	60,000	72,000	86,400	103,680	124,416

Table 1. Network Topology

relative differences between scenarios)

- Revenue: Incremental revenue associated with enhanced services.

Network Topology

The sample network modeled has an existing base of 500,000 voice subscribers. Over the next 5 years, this network is expected to grow with a CAGR of 20%. An existing network was selected versus a greenfield network in order to reflect the current state of the industry and to recognize that inter-CMS routing often becomes a problem after initial launch. Further statistics about the network topology and traffic are provided in Table 1.

Traffic Distribution

A critical factor in determining PSTN interconnection expense is characterizing traffic as “on-net”, where both originating and terminating points of the call are on the cable network vs. “off-net”, where the call either originates or terminates on the PSTN. Also, whether the call is local or long distance has a significant effect on interconnection fees.

Table 2 illustrates the percentage of calls by type over the life of the model. It is apparent that as more subscribers use the cable voice network, the probability of ‘on-net’ calls increases. This model assumes symmetric patterns of originating and terminating calls.

PSTN Termination Fees

A cable MSO typically has agreements in place with a competitive or inter-exchange carrier that handles PSTN interconnection. The MSO pays a per-minute termination to this interconnect partner for off-net calls. PSTN termination fees assumed in this model are:

Long Distance (including Toll) Termination Fee: $0.02 per call minute

Local Termination Fee: $0.003 per call minute.

A ten percent reduction in termination fees was assumed in the case where the cable operator interconnects to the carrier network directly via IP trunks.

Scenario Descriptions and Call Handling

Case A (No Inter-CMS Routing) is illustrated in Figure 7. In this case each softswitch is isolated from the others. The only calls that are kept “on-net” are intra-CMS calls, calls that originate and terminate on the same CMS. All other calls are routed to the interconnect partner via MGC and MG.

	Year 1	Year 2	Year 3	Year 4	Year 5
On-Net Calls					
Intra-CMS Calls	2%	3%	3%	4%	4%
Intra-PoP Calls	2%	4%	5%	7%	8%
Inter-PoP Calls	3%	4%	4%	5%	6%
MSO On-Net Calls					
MSO Peer Calls	0%	0%	2%	4%	6%
Off-Net Calls					
Local	46%	44%	42%	40%	38%
Long-Distance + Toll	47%	47%	44%	41%	38%

Table 2. Traffic Distribution

Case B (Inter-CMS Routing via Meshed SIP Trunk Network) is illustrated in Figure 8. In this case, all of the CMSs are interconnected to one another via a meshed network of SIP trunks. This offers the benefit of keeping significantly more calls on-net: intra-CMS calls, inter-CMS calls within the same PoP, inter-CMS between PoPs. All remaining calls are handed off to the interconnect partner either via TDM trunk (MG) or IP trunk (SBC).

Case C (Inter-CMS Routing via Standalone SIP Route Proxy) is illustrated in Figure 9. In this case, each of the CMS is interconnected to a set of SIP route proxies that provide the routing logic for all calls. This not only enables the carrier to keep the following calls on-net: intra-CMS, inter-CMS calls within the same PoP, and inter-CMS calls between PoPs, but it also retains MSO peer traffic on-net. This latter category is enabled by providing a single network point for all ENUM queries. Interconnection is handled via a combination of SBCs (MSO peers and PSTN) and MGs (PSTN).

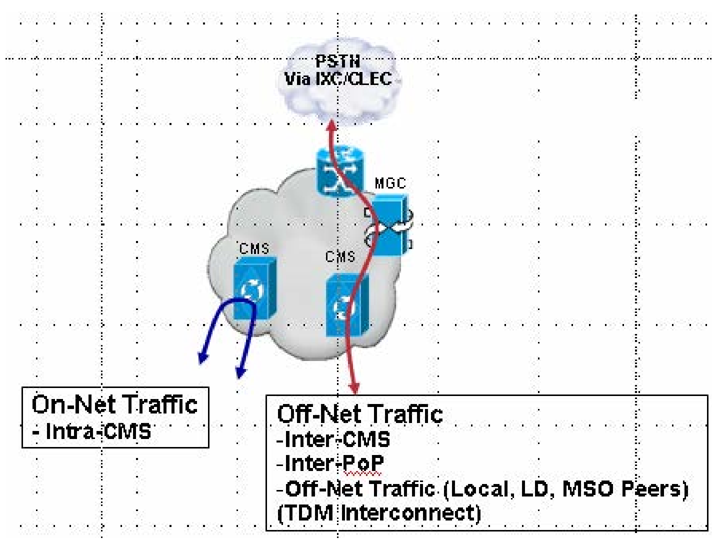

Figure 7. Case A - No Inter-CMS Routing

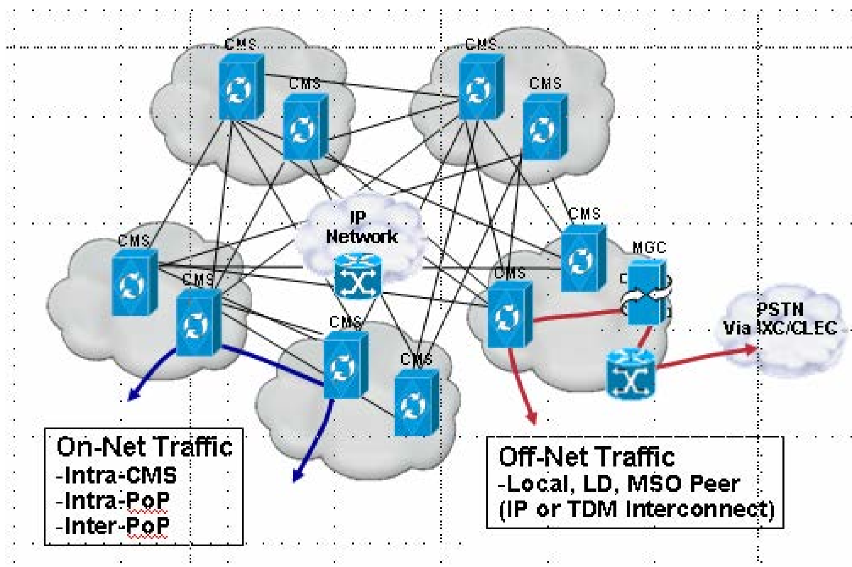

Figure 8. Case B – Inter-CMS Routing via a Meshed SIP Trunk Network

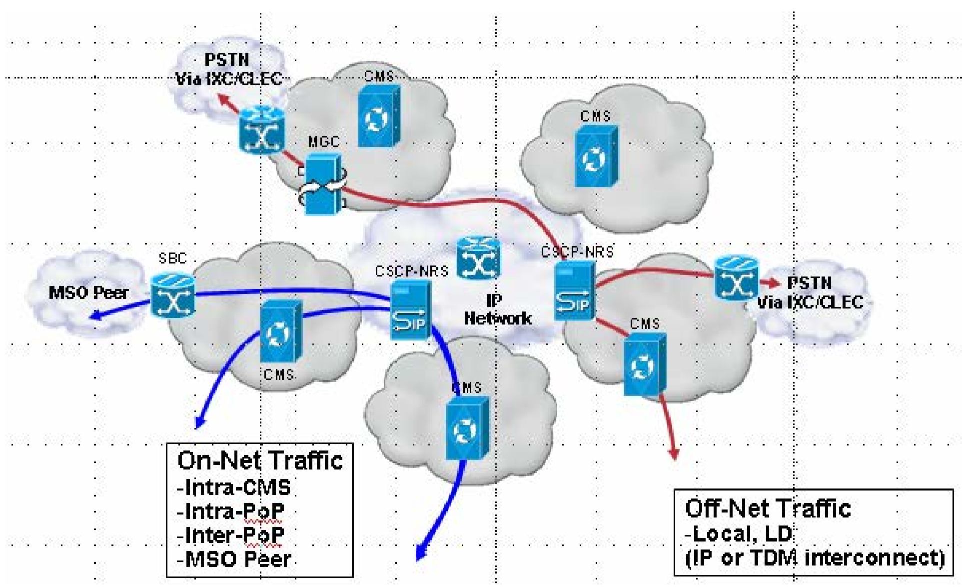

Figure 9. Case C – Inter-CMS Routing via SIP Route Proxy

Financial Investment

The amount and nature of investment required for each scenario differs.

Case A requires relatively little investment. Since most of the traffic is handed off to the PSTN carrier at the originating PoP; the only CapEx investment required is the purchase of incremental TDM gateway ports as traffic increases. A modest staffing expense is incurred to establish and maintain the routing tables in each softswitch.

Case B requires a modest investment. While more calls are staying on-net, driving a reduction in gateway capital expenses, there is incremental expense associated with backhauling IP traffic (routers and IP circuits). Also, there is a significant expense associated with setting up and maintaining a meshed network between CMSs (N-squared problem), and maintaining unique routing tables in each CMS.

Case C also requires a modest investment. Operationally, this scenario is easier to implement than Case B because fewer SIP trunks are required and there are fewer routing databases to provision and maintain. Consequently, there is a noticeable reduction in operational expense. But there is an incremental investment in SIP route proxies in order to support this routing function.

A summary comparing the investment required (excluding PSTN interconnection fees) for each of the three scenarios is illustrated in Figure 10.

Financial Return

The return on investment for each scenario can be articulated by calculating how much cost avoidance was achieved thru the reduction in PSTN interconnection fees. Additional value is captured by accounting for incremental revenue achieved by selling virtual number services.

In Case A, the cable operator pays a total of $495M in PSTN interconnection expenses over 5 years. This figure understandably increases in each of the years, as the total number of subscribers and traffic increases. This fee, $495M, represents the largest expense item for the cable operators rolling out cable voice service.

Case B reduces that expense by a considerable amount. By keeping more calls on-net and beginning to employ SBCs for IP interconnect, the total PSTN interconnection expense for the 5 years is $425M, a reduction of $69M or 14%.

Case C, however, offers the optimal return. Keeping even more of the traffic on-net allows the cable operator to reduce the

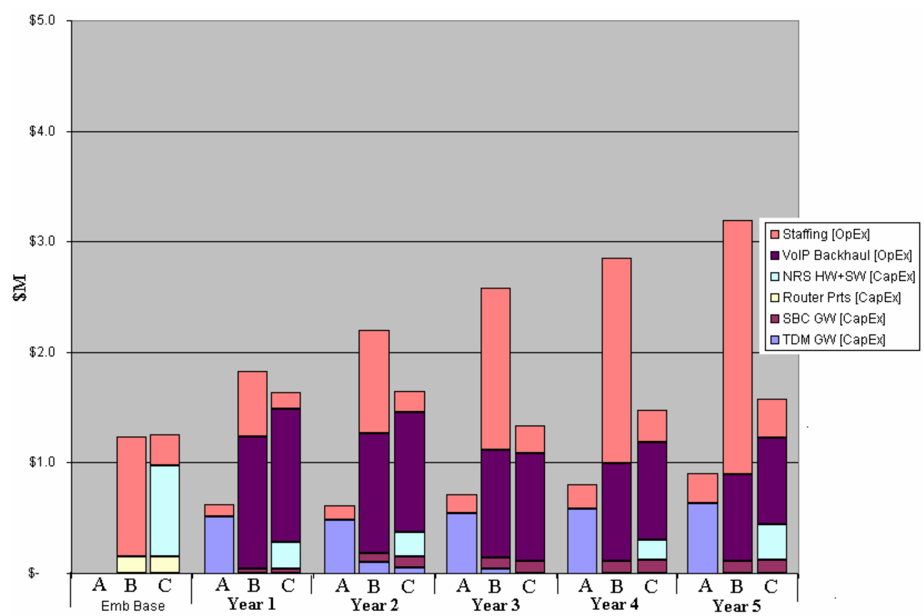

Figure 10. Comparative Financial Investment
(Excluding PSTN Interconnection Expense)

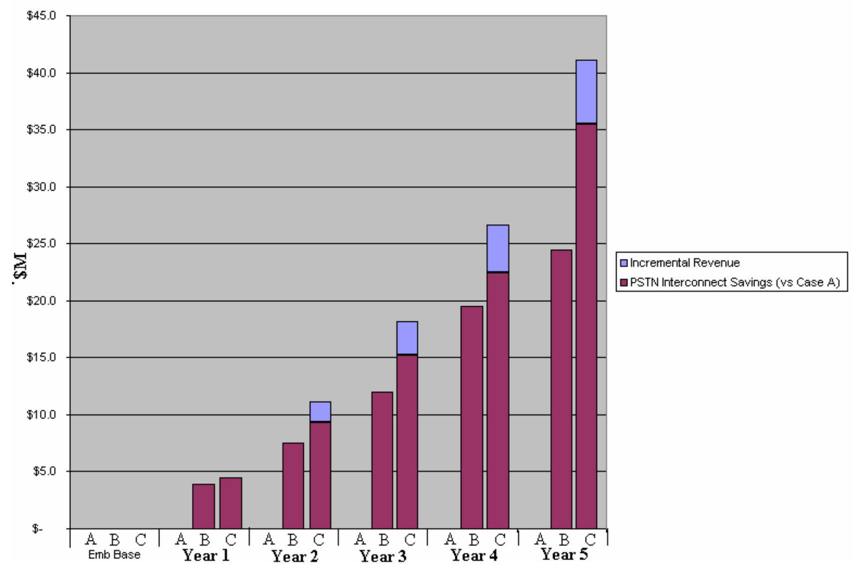

Figure 11. Comparative Financial Return

	5 Year Investment	5 Year Return	NPV	IRR
Case A	$ (3.65)	$ -	$ (1.87)	NA
Case B	$ (13.89)	$ 69.12	$ 35.08	392%
Case C	$ (8.93)	$ 121.34	$ 70.62	475%

Table 3. Financial Metric Summary

PSTN interconnection expenses for the 5 years to $388M, a $37M improvement over Case B (9%), and a $107M improvement over case A (22%). The case is further enhanced by incremental revenue associated with virtual number service adding an incremental $14M to the top line over 5 years.

The summary comparing the financial return for the three scenarios is shown in Figure 11.

<u>Financial Metrics</u>

Table 3 summarizes the investment and return for each of the scenarios. Case C, which uses the standalone SIP route proxy for inter-CMS routing offers the best return on investment.

<u>SUMMARY</u>

The hierarchical inter-CMS architecture using a standalone SIP route proxy helps MSOs scale their cable voice networks, while effectively managing their operating costs.

The benefits achieved are:

Operational Benefits:

- Maintaining more traffic ‘on-net’ improving overall service quality

- Avoiding operational complexity associated with growing cable voice networks

- Scaling subscriber counts independently of PSTN interconnection

Financial Benefits:

- Dramatically reducing PSTN interconnection expenses

- Optimizing CapEx investment

- Reducing ongoing staffing expense

- Facilitating new revenue streams

Strategic Benefits:

- Reducing dependency on PSTN interconnect partner who is often a competitor

- Gracefully evolve towards NGN based on PacketCable 2.0 / IMS architecture.

The inter-CMS architecture is based on highly-available and intelligent technology which can immensely help cable operators increase their profitability, efficiency and level of control over call routing both from, and within, their network.

HOME CO-AX METAMORPHOSIS:
FROM TV DISTRIBUTION TO CONTENT NETWORK

Steve Jenkins and Stephen Beales
Pace Micro Technology plc

Abstract

Signals being distributed through consumer premises often embark on an unpredictable journey through old cable with badly made extensions, low quality signal splitters and aging amplifiers. However, issues over cost, access to premises and liability hold operators back from upgrading home distributions to the same standard as their delivery networks.

This paper briefly recaps what a home network must provide from a home owner's perspective. It is assumed that the home owner has an existing co-ax distribution connected to a CATV network, and now wants to upgrade it to a full bi-directional home network to carry TV, data and voice between the ever growing mass of home infotainment equipment.

A range of typical distribution components are characterised for home network use to determine their performance in the reverse direction and frequencies outside of their intended band of operation. The impact of some non-ideal installation techniques is also examined.

Downstream signals to CPE (consumer premises equipment) are shown not to require any special attention since the requirements are the same for legacy purposes. Upstream signals require special attention and predictions are made to show how data rates of 60Mb/s can be achieved without modification to a badly installed home distribution system, although in-line filters are required to protect legacy receivers from upstream signals.

SETTING THE TARGET

With the adoption of multiple personal computers, distributed audio and video, gaming and a host of other networkable devices within the home it is no surprise that people want to link up all their devices over one network, and not have to run one co-ax to connect to the external network then another for their own equipment. It's important that the PC can connect to the TV and that the security system is connected to both. It's also important that their own local communications are kept private from the delivery network and neighbors. All too often systems are designed from an operator's perspective with scant regard for the end user needs.

The whole network installation needs to be painless for the home owner, so it's safe to assume that the equipment must be inexpensive, quick to install, need no specialist instrumentation, and leave the decorating intact.

SO YOU WANT TO USE YOUR ANALOG DISTRIBUTION AS A DIGITAL NETWORK?

In the eyes of the average DIYer, converting an existing co-ax network from analog to triple-play should be a piece of cake. If it worked for the old analog stuff then surely it must work for the digital stuff, after all it's the latest technology isn't it?

Ideal home networking technologies must be able to live up to this premise, which means they must be able to operate over any home network architecture where analog TV was successfully distributed with no modification to the infrastructure, even if the analog picture was noisy or distorted. Equally, it must also be able to reliably deliver return path signaling even though the network was never originally intended to be used for this purpose.

Of course, once the bi-directional digital communications are running over the co-ax, the home owner needs them to operate without causing interference to any legacy equipment that is not part of the digital network.

THE EVOLVING DISTRIBUTION

At some point in the co-ax distributions life it may have been able to handle bi-directional digital signaling, but as the systems grow organically over time, sometimes professionally and sometimes not, it is likely to contain problems that can scupper a seamless upgrade.

In the oldest installations the home co-ax system may have started out as a CATV feed into the living room and the master bedroom. Over time, the co-ax was split to feed other bedrooms, the study, the kitchen and maybe even the back yard. With a professional installation this would have been designed as a star system or 'branch and tap' configuration, but in many cases a splitter is simply inserted in the nearest available co-ax and, so long as the TV can get some kind of picture, then the job's done.

Figure 1 shows a substandard home-grown installation which largely matches the previous description. One of the splitters was accidentally installed the wrong way round which resulted in no picture in bedroom 3, and lower signals in bedrooms 1 and 2. To resolve this, a 20dB amplifier was installed at the star point. Finally, when the content network was upgraded for data, a cable modem was installed in the office.

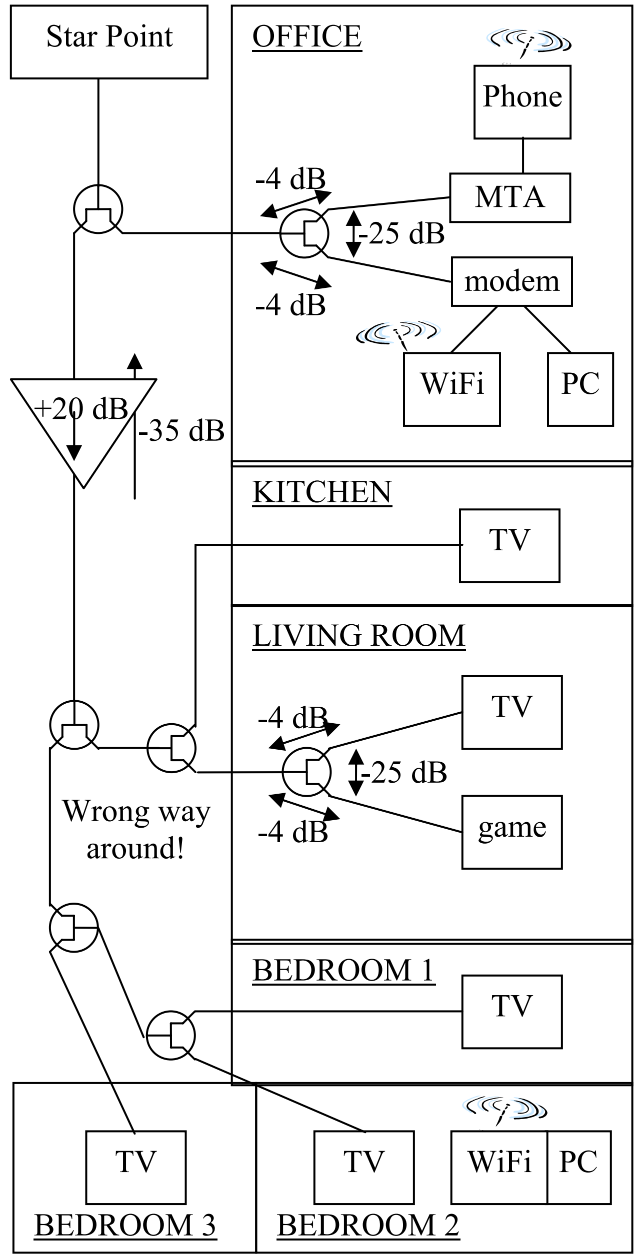

Figure 1. Reference legacy distribution

IT WORKS JUST FINE FOR ANALOG

At this time the network is carrying a full suite of analog, digital TV and data signals which consumes the entire spectrum from 5 to 870MHz.

From star to:	Signal gain
Office MTA	-7
Office CM	-7
Kitchen TV	6
Living room STB	9
Living room games	9
Bedroom 1 TV	5
Bedroom 2 TV	5
Bedroom 3 TV	-13

Figure 2. Signal gain between star point and all outlets

Figure 2 shows the averaged signal gain at each outlet for every path in the distribution in the traditional downstream 54 to 870 MHz band. The distribution is adequate for distributing signals from the star point with all outlets levels within 14dB of each other, apart from bedroom 3 with the splitter installation fault which results in a grainy but acceptable picture.

DISTRIBUTION BECOMES NETWORK

The biggest step change in converting the distribution is the ability to send signals in other directions. This means that new high-loss paths are created across some splitter outputs and in the reverse direction across amplifiers.

UNDERSTANDING THE NETWORK COMPONENTS

Co-ax

The most common cable choices over recent years are CT100 and CF100. The impedance is matched for 75 ohms and the attenuation is 10dB per 100 meters at 870MHz and 0.2dB at 5MHz. In general any other co-ax found in domestic installations will be of lower quality in terms of attenuation, frequency roll-off and impedance matches.

Following the notion that this co-ax has been good enough to pass legacy signals, then the cabling will be suitable for digital communications that use the same frequencies or lower.

Splitters

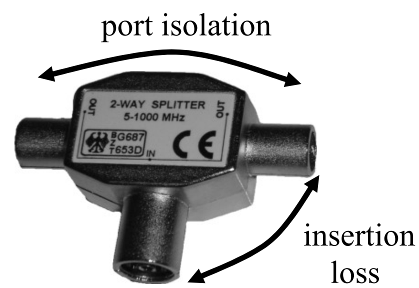

Figure 3. splitter sample #1

Almost all modern power splitters are the transformer type as these give minimum signal loss. When a signal is split, all the power is divided equally at each of the outputs and therefore the output signals are in theory 3dB lower than at the input. In practice there is some additional loss, typically between 0.5 and 1.5dB. The extent of this additional loss being indicative of the quality of the splitter.

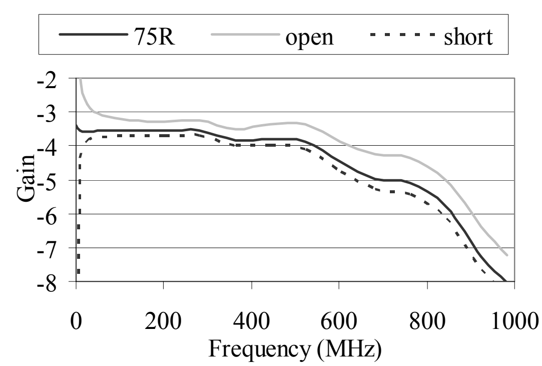

Figure 4. Transformer splitter insertion loss sample #1

The graph in figure 4 show that when the unused port is correctly terminated in 75 ohms the insertion loss is between 3.5 and 5.5dB up to 800MHz, but then rapidly increases to more than 8dB at 1GHz. Again, high frequency loss impacts on the potential of this unused spectrum. When the unused port has nothing connected to it, then the insertion loss is reduced by 0.5 to 1dB. Similarly a short increases the loss slightly. If an unterminated co-ax is connected to the unused port then these effects are reduced in proportion to the length of the co-ax.

High port isolation reduces the interference that is caused when two TV sets are tuned to the same channel. Figure 5 shows that when the input (common) port is correctly loaded in 75 ohms, then the port isolation is generally greater than 20dB. In practice it is unlikely that the common port will ever have a significant impedance mismatch at the common port as either an open or short here would result in no signals further down the distribution. However, the effects are shown here as they highlight a transformer-splitter's extreme sensitivity to common-port loading.

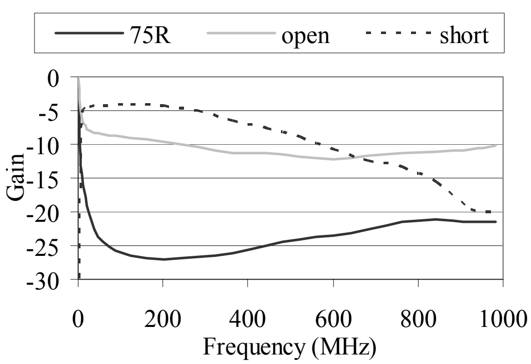

Figure 5. Transformer splitter port isolation sample #1

When using a distribution with splitters, consideration of the return path signaling must be taken. In this case a device transmitting across the output ports has the potential to saturate other receivers further up the distribution as a result of the low insertion loss in comparison to the port isolation.

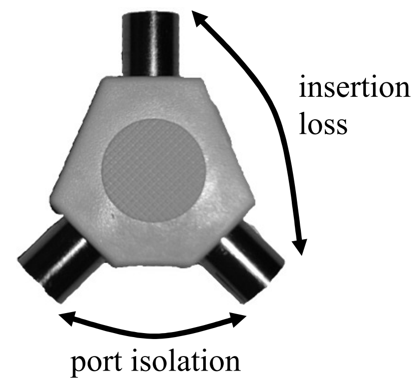

Figure 6. Splitter sample #2

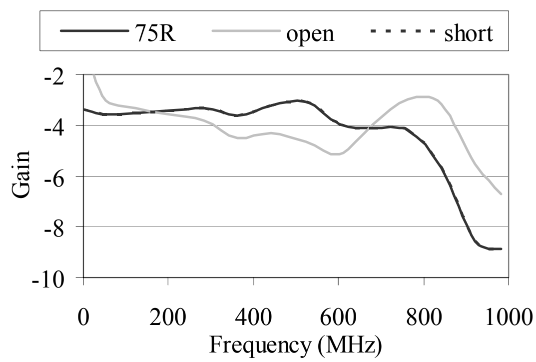

Figure 7. Transformer splitter insertion loss sample #2

There are a number of techniques used to construct a transformer-splitter. Figure 7 show that the insertion loss of a second type is much flatter than sample #1 to 800MHz with similar high frequency roll-off to 1GHz.

Trunk and Tap

This technique uses unequal splitters where the loss in each split along the trunk is very small, and each tap off the trunk has around 20dB of insertion loss. This system is rarely used in a single domestic installation and is more likely to be found in MDUs (multi dwelling units) or in the co-ax part of a HFC network.

Daisy-Chain

A simple, though incorrect way for an uninformed DIYer to install a new outlet is to run a spur from an existing outlet in a similar way a mains outlet might be wired.

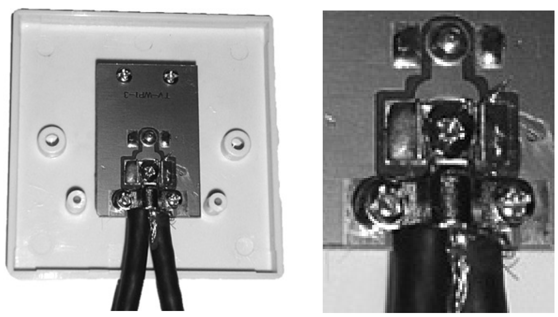

Figure 8. straight spur

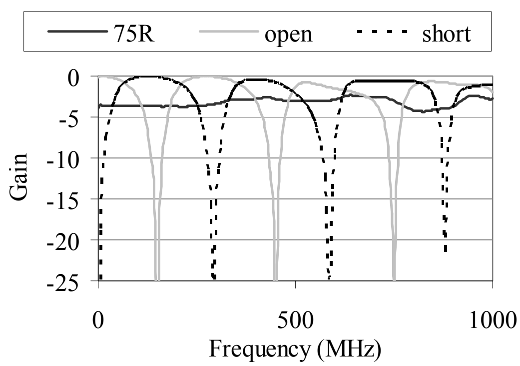

Figure 9. Daisy chain 17 cm spur

Figure 9 shows the frequency response between the incoming co-ax and the wall plate socket when there is a 17cm spur. This length is fairly typical when wall plates are mounted back-to-back either side of a wall. When the spur is loaded with 75 ohms, the insertion loss is approximately 3dB as theory predicts. Problems occur when the spur is not correctly loaded, the most typical instance being an open circuit. In this case nulls are presented in the frequency response. Figures 10 and 11 show that as the stub length increases then the effect is reduced.

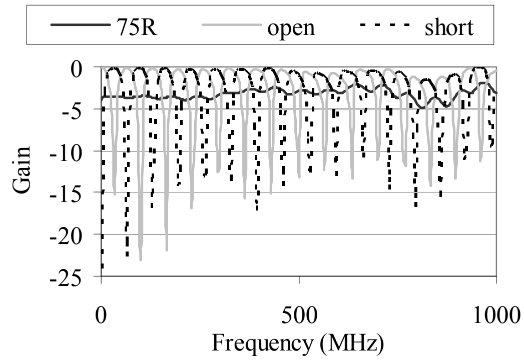

Figure 10. Daisy chain 300 cm

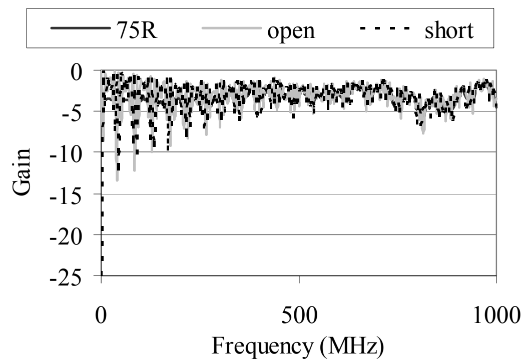

Figure 11. Daisy chain 900 cm

The major problems caused by unterminated spur outlets are loss of signal level, and channel tilt. Figure 12 is a snap-shot of the frequency response for two different spur lengths.

Signal loss can be more than 40dB within the nulls at certain frequencies which means that problems may have been masked in the past if legacy signals happened to avoid the nulls leading to the potential for misdiagnosis of issues when the distribution is switched to network.

Of equal importance is the gradient of the nulls which can result in channel tilt. The choice of communication system must take this into account either in its ability to avoid these frequencies, or simply to be robust enough to handle them.

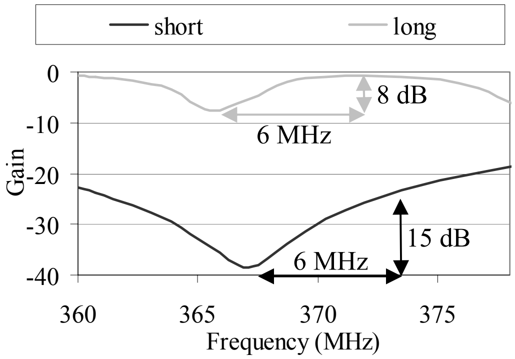

Figure 12. Unterminated cable lengths

One-way Amplifiers

Figure 13. variable gain amplifier

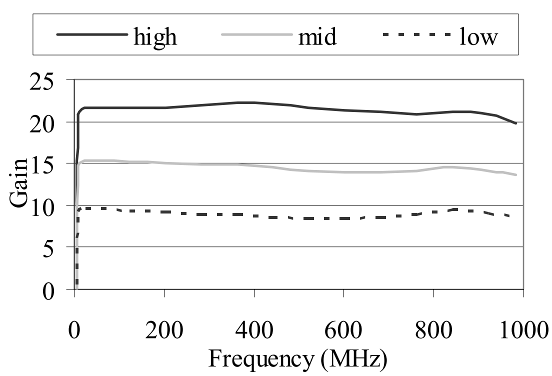

Figure 14. Amplifier Gain response

The sample characterized is a variable gain device with a range from 10 to 20dB. The gain response is flat within ±2dB from 5MHz to 1GHz as shown in figure 14.

The most important question is how a one-way amplifier handles upstream signals, that is, signals launched at high enough power to overcome the reverse isolation shown in figure 14.

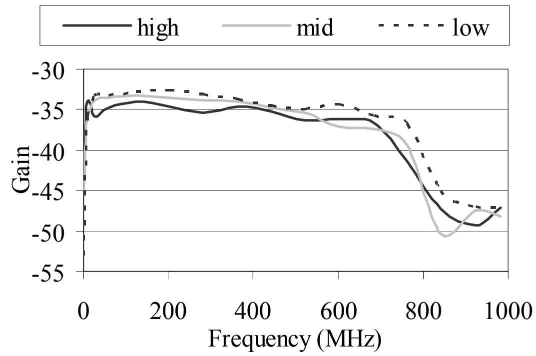

Figure 15. Amp reverse response

The main concern is signal distortion in the reverse direction. For this measurement the amplifier is intentionally characterized in reverse, so the input and output are interchanged. Figure 16 shows the third order intercept response measured at the output, i.e., the distortion on the signal once it has passed backwards through the amplifier. This particular amplifier uses current starvation to vary the gain which is why the best results are achieved at the highest gain setting. The intermodulation performance is also poor at the higher frequencies where signal levels have to be increased to overcome the amplifiers reverse isolation.

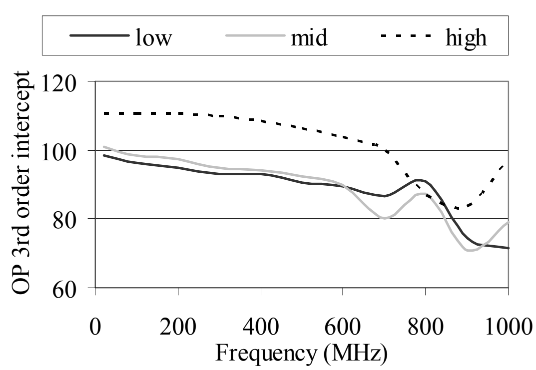

Figure 16. Amplifier IP3 response

UPSTREAM POSITIONING

There are three general frequency ranges to consider for upstream signaling, below 50MHz, 50 to 870MHz and above 870MHz.

Upstream in the 50-870MHz Band

Assuming that the downstream frequency range is fixed so that analog signals can pass to legacy receivers, table 2 (see next page) shows that the 50 to 870MHz band is not suitable for upstream because the high launch levels required will saturate downstream tuners with signal levels potentially being over 70dB greater than the wanted downstream. For example, to transmit from bedroom 2 to the office there is 80dB of loss to overcome. However, a downstream signal from the star point loses 8dB so for both signals to appear at the office at the same level, bedroom 2 must transmit 72dB higher than the downstream signals at the star point.

The gain between the star point and bedroom 2 is 2dB so the power of the upstream signal it is transmitting will be 70dB higher than the downstream signal which will cause its tuner to saturate, this is self interference. Similarly interference is caused to other client devices such as in bedroom 1 where the transmit signal intended for the office will be 72 – 1 – 26 = 45dB more than its downstream signal in bedroom 1.

Segregating the Signal Bands

Seperating the upstream and downstream into two bands allows the downstream tuners to be protected from the high power upstream signals using band filters or diplexers, therefore tuner saturation is avoided.

Tables 1 and 3 use the average characteristics of the distribution network operating in the 0 to 50MHz and 870MHz to 1GHz bands respectively. The overall difference in the network's performance in these two bands is that, at higher frequencies, there always greater signal loss to overcome. Signal loss impacts the carrier-to-noise ratio and thereby limits the efficiency of the modulation schemes available.

In addition to increased signal loss it has also been shown that when a one-way amplifier is present, the reverse intermodulation performance is also impaired at high frequencies. This limits the maximum transmit signal level which again affects the carrier-to-noise ratio thereby reducing the ability to use efficient modulation schemes.

HOW MUCH BANDWIDTH?

There are many different usage profiles to consider which can include internet usage, bulk file transfers between PCs, multi-room digital video recorders (DVRs) to name but a few. On top of this there are various usage profiles for all these applications. For the purpose of this paper, the model used to estimate the bandwidth budget internal to the home is to consider two 12Mb/s digital video files being played across the network at the same time, one of which can be at 4 x speed for trick modes giving a combined data rate requirement of 60Mb/s. This aggressive case is likely to cover most other cases.

Within the reference model, the restrictions imposed by incorrectly installed splitter 4, and distortion in the reverse direction across a one-way amplifier limits the choice of modulation efficiency to two bits per Hz. A 20 per cent provision for overhead reduces this to 1.7 bits per Hz data rate which requires 36MHz of spectrum. With the amplifier removed and the splitter installed correctly then a modulation efficiency of 3.4 bits per Hz can be used thereby reducing the sectrum required to 18MHz.

Low band		To								
		star	office MTA	office CM	Kitchen	Living TV	Living Games	Bed1	Bed2	Bed3
From	star	-	-7	-7	6	9	9	5	5	-13
	office MTA	-7	-	-25	-16	-20	-20	-20	-20	-38
	office CM	-7	-25	-	-16	-20	-20	-20	-20	-38
	Kitchen	-46	-71	-71	-	-29	-29	-36	-36	-54
	Living TV	-49	-75	-75	-29	-	-25	-40	-40	-58
	Living games	-49	-75	-75	-29	-25	-	-40	-40	-58
	Bed1	-50	-75	-75	-36	-40	-40	-	-25	-7
	Bed2	-50	-75	-75	-36	-40	-40	-25	-	-7
	Bed3	-68	-93	-93	-54	-58	-58	-29	-29	-

Table 1. Average signal gain for 0 to 50 MHz for all outlet routes

Mid band		To								
		star	office MTA	office CM	Kitchen	Living TV	Living Games	Bed1	Bed2	Bed3
From	star	-	-8	-8	2	6	6	2	1	-15
	office MTA	-8	-	-25	-19	-24	-24	-25	-25	-42
	office CM	-8	-25	-	-19	-24	-24	-25	-25	-42
	Kitchen	-49	-74	-74	-	-30	-30	-39	-39	-56
	Living TV	-53	-79	-79	-30	-	-25	-43	-43	-60
	Living games	-53	-79	-79	-30	-25	-	-43	-43	-60
	Bed1	-53	-80	-80	-39	-43	-43	-	-26	-10
	Bed2	-54	-80	-80	-39	-43	-43	-26	-	-9
	Bed3	-70	-97	-97	-56	-60	-60	-31	-30	-

Table 2. Average signal gain for 50 to 870 MHz for all outlet routes

High band		To								
		star	office MTA	office CM	Kitchen	Living TV	Living Games	Bed1	Bed2	Bed3
From	star	-	-12	-12	-5	1	1	-6	-7	-14
	office MTA	-12	-	-20	-20	-27	-27	-28	-28	-36
	office CM	-12	-20	-	-20	-27	-27	-28	-28	-36
	Kitchen	-65	-85	-85	-	-27	-27	-40	-40	-49
	Living TV	-69	-92	-92	-27	-	-20	-46	-46	-54
	Living games	-69	-92	-92	-27	-20	-	-46	-46	-54
	Bed1	-71	-93	-93	-40	-46	-46	-	-21	-13
	Bed2	-72	-93	-93	-40	-46	-46	-21	-	-12
	Bed3	-79	-101	-101	-49	-54	-54	-28	-27	-

Table 3. Average signal gain for 870 MHz to 1 GHz for all outlet routes

CONCLUSIONS

There are three major factors to be understood in order to provide an ideal solution that can use the home broadcast distribution networks as a bi-directional digital network. These factors are all regarding upstream as any downstream signals are assumed to work in line with legacy services.

Attenuation paths within the distribution mean that upstream and downstream channels must be separated into discrete bands of frequency so that receiver tuners can be protected by filters from becoming saturated by return path signals.

Attenuation within the distribution is highest at higher frequencies, therefore the 0 to 50MHz region is the best choice for the return path since this allows the use of more efficient modulation schemes.

Upstream signaling can be blasted the wrong way across one-way amplifiers but the levels are restricted by reverse intermodulation performance. This is made worse when a variable-gain amplifier is set to low gain but only becomes a significant when this event is in combination with another extraordinary case such as a splitter installed incorrectly, or an unterminated spur.

IMS AND THE IMPORTANCE OF ACCESS NETWORK MANAGEMENT

Rich Higgins
C-COR, Inc.

Abstract

IP Multimedia Subsystem (IMS) is positioned to become the primary next generation IP network services architecture for broadband service providers worldwide. IMS has the potential to increase revenue opportunities by allowing greater flexibility in both services and the types of devices that consumers will use to access these services. IMS also has the potential to lower service development costs and reduce time to market.

The need for access network management both to prevent theft/abuse and ensure a quality user experience grows as more services come to depend on this shared infrastructure.

This paper will examine some of the IMS architecture/service elements and the implications of IMS for the cable access network architecture with particular emphasis on the Policy Decision Function (PDF) and access network security.

INTRODUCTION

Background

IP Multimedia Subsystem is an architecture which defines a platform for the delivery of multimedia services via the IP protocol. It is largely derived from specifications developed by the 3rd Generation Partnership Project (3GPP) but incorporates many standards from the Internet Engineering Task Force (IETF). IMS was originally developed to provide mobile operators the ability to offer IP based services to subscribers via cellular networks.

A number of other standards bodies have adopted IMS as a basis for their ongoing work. The European Telecommunications Standards Institute (ESTI) Tispan project is extending IMS to include support for legacy telecom systems. ESTI has agreed that Tispan developed specifications will be submitted to the 3GPP body for approval and inclusion in future IMS releases. CableLabs, via PacketCable 2.0, is focused on extending release 6 to include support for cable HFC access networks. Other standards groups such as the International Telecommunications Union (ITU), the Alliance for Telecommunications Industry Solutions (ATIS), the Open Mobile Alliance (OMA), the GSM Association, and 3GPP2 are also following the progress of IMS closely and are expected to be heavily influenced by it. It is expected that IMS release 7 will include some of these extensions as well as important enhancements such as better end-to-end quality of service and policy control definitions.

The following section provides a very brief look at the IMS architecture for those who may be unfamiliar with the basic principles of operation.

IMS Overview

Traditional telecom operators have long sought a Service Delivery Platform (SDP) which would be based on a common infrastructure and would allow for the rapid and inexpensive development and deployment of new services. IMS incorporates many of the features desired of an SDP. In the IMS model both the Core (or control) layer and the Service (or application) layer are access agnostic and are accessible using standards based protocols.

Any access mechanism that is IMS compliant can make use of the complete suite of control functionality offered by the core as well as applications in the service layer. New services need be developed only once as they are independent of the access network. It no longer matters whether a subscriber is requesting service via a cable network or over a cell phone. All that matters is the availability of the necessary resources to provide the requested service.

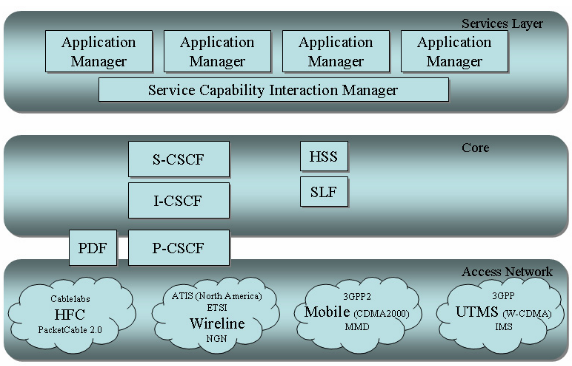

Figure 1. IMS Architecture

As shown in

Figure 1 immediately above the access network layer is the core or control layer. The core contains the Call Session Control Function (CSCF) as well as specific systems which support this general processing function. It is the Proxy CSCF (P-CSCF) which acts as the interface between the access network and the IMS core. All requests for services from any end user device (referred to in the IMS specifications as User Equipment or UE) regardless of access mechanism are initially handled by the P-CSCF.

The Interrogating CSCF (I-CSCF) provides an access point into a specific operator network. If the P-CSCF is located in the subscriber's home network then requests for services may (or may not) route directly to the Serving CSCF (S-CSCF) depending on the number and configuration of S-CSCF's in the network. If the P-CSCF, which is requesting services, is located in a network that is not the UE's home network the I-CSCF of the home network is contacted and will determine the appropriate S-CSCF to route the service request.

It is the S-CSCF that is responsible for actually servicing the request(s) of the subscriber that is associated with that server. The S-CSCF interacts directly with the Services Layer and the various Applications Managers. The S-CSCF also interacts with the Home Subscriber Server (HSS) which is a database of information about the subscriber. Specifically, the HSS maintains the mapping between the subscriber and the S-CSCF, subscriber service profile, and security (identity) information.

The Policy Decision Function (PDF) is responsible for authorizing quality of service for the media in accordance with the parameters of the service request and the operators established business rules. In a PacketCable 2.0 implementation it is likely that the PDF functions will be performed by the PCMM Policy Server.

The Subscriber Location Function (SLF) is also defined in the IMS specifications and serves to identify the HSS that contains information about a specific subscriber. Also identified by the IMS specs is a function referred to as the Service Capability Interaction Manager (SCIM). The SCIM was envisioned as a service broker between application mangers. As of release 6 this component was ill defined and there is considerable debate as to the necessity for a separate SCIM.

Increased Revenue Opportunities

The disassociation of access, control, and services creates new service opportunities which would be difficult, at best, to host in a traditional "vertical" service environment. Because the control and services layers operate in the same way regardless of access network, development of services which span access devices will be simplified. New services could include multi-mode IP phones that operate wirelessly via Bluetooth or WiFi connected to an HFC network while in range of a home networking device but which switch seamlessly to a cellular network when the handset leaves the home.

IMS also incorporates the concept of "presence" in which the network itself contains intelligence about the subscriber's location and access device capabilities. It is possible to envision future service in which a voice only call may be initiated via a relatively low bit rate cellular connection while the user is driving to work but switches automatically to a video conference as the user enters their office environment.

The IMS infrastructure will allow subscribers to access content regardless of their broadband access network (HFC, DSL, WiFi, WiMAX, 3G, etc.). This unfettered service access will lead to service roaming beyond simple voice. Subscribers no longer need be tied to their desk, or their television, or their telephone in order to access the services or content that they desire.

Reduced Development Costs

In proprietary service environments providers have had to resort to developing their own applications from scratch. At best, they could rely on a few select vendors to develop services which frequently required long, expensive, and sometimes arduous integration periods.

In an IMS environment interfaces and the underlying protocols are well known and very actively used by the development community. Development environments and simulation tools exist that allow applications to be developed independent of the provider or network. It is reasonable to expect that application development will move more toward an "Internet like" model where new services are rapidly created and deployed.

Reduced Time to Market

Because application development is no longer necessarily confined to a specific platform the number of potential developers increases dramatically. Competition for new applications will not be confined to large vendors only but will be open to startups and even individuals. This model begins to more closely resemble the Internet model of application development and deployment.

If, indeed, application development becomes more "Internet like" it is reasonable to assume that the speed with which the applications are developed and deployed in an IMS environment will be similar. Groups are already working on standardizing interfaces to the IMS services layer to ease development of IMS applications. The Java Community Process describes their JSR 281 standard as follows, *"This JSR is intended to enable application programmers to easily write applications that can integrate with the IP Multimedia Subsystem (IMS). The specification will expose IMS functionality through high-level*

APIs in an integrated and consistent way. The API hides IMS implementation details to the maximum extent. The API abstracts the underlying technology and at the same time provides the developers with maximum flexibility. This approach secures conformance to IMS related standards and at the same time gives developers possibility to focus on the functionality of the services and not on the IMS technology implementation details. In this way IMS domain will be revealed to the broad J2ME developer community and will encourage faster adoption of the IMS services provided by the wireless networks." [1]

ACCESS NETWORK SECURITY

Opening the network to a multiplicity of access devices and (potentially) external service provider networks raises concerns particularly regarding user authentication.

Security threats to the network generally fall into one of three categories:

- Theft of service
- Denial or disruption of service
- Information theft (subscriber or provider)

IMS security standards are extensive and designed to deal effectively with each of these threats. A complete discussion of IMS security is well beyond the scope of this paper.[2] Instead, this discussion will focus only on the access portion of the IMS network and the security enhancements required for cable.

Authentication

In keeping with its wireless roots, release 6 of the IMS standards authentication of the user equipment (UE) is handled by an application running on a Universal Integrated Circuit Card (UICC). Authentication between the UE and the P-CSCF in handled via a challenge response mechanism known as Authentication and Key Agreement (AKA) specifically UMTS AKA. Similarly, authentication between the UE and the core network is handled via ISM AKA.[3]

During initial registration the UE sends a SIP Register message to the P-CSCF. The P-CSCF forwards the Register message to the appropriate S-CSCF (with the assistance of an I-CSCF as required) which contacts the HSS to obtain the user information necessary to complete the authentication. Using the information obtained from the HSS the S-CSCF responds to the UE (via the P-CSCF) with a challenge. If the UE responds correctly the UE will have successfully authenticated and a security association will have been established between the UE and P-CSCF. Additional message traffic including requests for service will now be allowed.

As mentioned above, the HSS contains profile information which includes both the subscriber's private and public identities. The private identity along with a long term shared secret is used to authenticate the user to the network. It is the responsibility of the HSS to maintain the shared secret key information necessary to secure communications between the UE, the P-CSCF, and the S-CSCF.

In a cable environment some devices may use a UICC based application for authentication (e.g. a GMS phone requesting access via a cellular network) while others may not (e.g. streaming audio or video to a PC). Extensions to define authentication mechanisms that provide support for non-

UICC devices will likely be addressed by PacketCable 2.0.

IMS makes use of the concept of private and public identities. The private identity is typically assigned by the home network at the beginning of the subscription process. It is permanent and persists for the entire duration of the user's subscription with the service provider. The public identity is used by the subscriber to request services and/or communication with other users. There is typically only one private identity associated with a UE although there may be many public identities.

The one-to-one relationship of private identity and UE can be problematic. It is certainly possible to envision a single user owning multiple UE's with different capabilities (e.g. cell phone with multimedia capabilities and a PC). Difficulties also arise in the instance where a single UE requires multiple security credentials for accessing services via different access mechanisms. As in the earlier example it is certainly within the realm of possibility that an operator will wish to offer a service that allows customers to move seamlessly between a cellular network and a wireless (e.g. WiFi) network in a home environment. In this case the UE would require simultaneous security associations with both the P-CSCF that resides in the cellular operator's home network and the P-CSCF that resides in the MSO network.

POLICY DECISION FUNCTION (PDF)

Quality of Service (QoS)

In the IMS specification QoS within the core and services layers are handled with traditional IP bandwidth reservation or packet tagging methods.[4] Bandwidth reservation is accomplished via the Resource reSerVation Protocol (RSVP)[5] Integrated Services (IntServ)[6]. Packet tagging is accomplished via Differentiated Services (Diffserv)[7] or Multi-Protocol Label Switching (MPLS)[8].

In the HFC portion of the access network QoS will likely be provided as specified by PacketCable Multi-media (PCMM) with some modifications. In PCMM the application requesting service communicates directly with the Application Manager (AM). It is the AM that communicates with the Policy Server (PS). The PS, in turn, communicates with the policy enforcement point (PEP) which, in the HFC, network is the CMTS.

In an IMS implementation the application requiring QoS would first communicate that request to the P-CSCF. The P-CSCF would then forward the request to the AM which would determine the resources required for that service request. The AM would then generate a PacketCable Multimedia request which is forwarded to the PS. The PS will verify that the resource requests are within acceptable limits as set by the operator.[9] The PS will then act as a proxy with respect to the application manager and the PEP by forwarding policy requests and returning responses.

The ability of the access network to support a dynamic level of QoS is especially important in an IMS network. One of the most interesting features of the Session Initiation Protocol (SIP), which is the backbone of the IMS signaling network, is the ability to dynamically add and/or modify service flows. From a services perspective this means that the operator can offer a greater range of services than are possible today. For example, most early implementations of VoIP support some sort of multi-party calling (typically 3-way calling). Currently, 3-way calling functionality is an application that must be written for specific vendor hardware. In a SIP environment, additional call legs can be added to a media stream simply by making additional requests for resources. Also, there is no particular reason why a subscriber should be limited to only 3-way calling (except as dictated by business rules). With SIP a caller could conference 4, 5, or more other callers just as easily.[10] Similarly, a user may initiate a session as voice only and switch at some point to voice with video or vice versa.

SERVICE MONITORING

Another area of growing importance in the delivery of IP services via the IMS network is the ability of the operator to translate measured network performance into an understanding of service performance. Many content owners are now including some reference to service monitoring in their contractual agreements. Content providers want to be able to ensure not only that content was played out to a subscriber as specified but also that the quality of the user experience was acceptable. In an IP environment effective service monitoring will require changes to the network and changes to the monitoring process.

Release 6 of the IMS specifications includes support for Real-Time Transport Control Protocol (RTCP)[11] as the mechanism for monitoring audio service quality in the access network. RTCP provides feedback about specific session performance including measurements of packet loss, inter-arrival jitter, and delay or latency. These parameters can be translated directly to audio quality for example when judging the quality of a VoIP call. Additional audio monitoring information is available via the RTCP Extended Reports (RTCP-XR)[12] standard.

There are as yet no specified standards for monitoring video quality.

CONCLUSION

IMS is potentially very important to the cable industry. Depending on implementation details it may provide a powerful new means for rapid deployment of new services. It is also likely to be the mechanism by which the primary competitor (telco's) will deploy new services including IPTV.

Although there seems to be general agreement on the high level standards there is some disagreement on implementation details. One vendor may emphasize a specific component while another describes additional components or lacks some entirely. It is vital that the industry remain focused on creating cable specific standards to allow truly uniform multi-vendor compatibility if IMS is to reach it's full potential.

REFERENCES

1. Java Specification Requests (2005) JSR 281: "IMS Services API"

2. 3GPP TS 33.210: "3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G Security; Network domain security; IP network layer security"

3. 3GPP TS 33.203: "3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G Security; Access security for IP-Based services"

4. 3GPP TS 23.207: "3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; End-to-end Quality of Service (QoS) concept and architecture (Release 6)"

5. IETF RFC 2205 (1997), "Resource ReSerVation Protocol (RSVP)"

6. IETF RFC 2210 (1997), "The Use of RSVP with IETF Integrated Services"

7. IETF RFC 2474 (1998), "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers"

8. IETF RFC 2702 (1999), "Requirements for Traffic Engineering Over MPLS"

9. PKT-TR-MM-ARCH-V02-051221, "PacketCable™ Technical Report Multimedia Architecture Framework"

10. IETF RFC 4353 (2006), "A Framework for Conferencing with the Session Initiation Protocol (SIP)"

11. IETF RFC 3550 (2003), "RTP: A Transport Protocol for Real-Time Applications"

12. IETF RFC 3611(2003), "RTP Control Protocol Extended Reports (RTCP XR)"

INDIVUALLY TARGETED ADVERTISING IN A SWITCHED SERVICES ENVIRONMENT

Steve Riedl, Time Warner Cable
Paul Delzio, BigBand Networks

Abstract

Enhanced network switching technologies are becoming available throughout cable plants and providing operators with an opportunity to dramatically change the way advertisements are delivered and viewed by subscribers. Instead of subscribers being exposed to advertisements chosen on assumed demographics, ads can be forwarded based upon more personally relevant data such as viewing preferences, hobbies, interests, personal demographics and other parameters. Targeted advertising offers marketers greater confidence that the promotional messages they are paying for will result in improved response rates among their intended audiences.

Service providers benefit by having the ability to sell more valuable advertising content to more refined subsets of subscribers. This could also satisfy subscriber demands to receive only relevant content and potentially less cumulative advertising.

Additionally, by moving away from a "push" model that delivers ads of little interest to viewers, to a scenario in which subscribers are more likely to be engaged, opportunities exist for greater interactivity. Within their own streaming experience, subscribers can request additional information or other follow-up from the vendor regarding what's being advertised, possibly going so far as to complete a transaction. For example, after viewing a promotional message about a new sports car, a subscriber can be invited to request a longer form of the advertisement and schedule a test drive.

This paper compares two broadcast architectures, switched and stored. Each can analyze data about what programs are being viewed and deliver ads that have the highest probability of achieving the desired responses, based on subscriber profiles and household classification data. Furthermore, the paper considers multiple models for delivery of targeted advertising in switched environments. Direction of advertising towards its correct destination can be achieved from multiple network locations, ranging from switched digital broadcast techniques based at headend and hub locations to dynamic digital tuning within set-top boxes.

INTRODUCTION

Just as the Internet is driving new consumption patterns of video, it is radically transforming advertising models. Google has strongly established that delivering contextually relevant ads increases the interest of the user and value to the advertiser. Recent studies of TV viewers have confirmed that subscribers are more likely to respond favorably to ads if the topic is of interest to them. So, contrary to concerns that increasing use of personal video recorders and on-demand services will decrease the effectiveness of ads, an opportunity is rapidly emerging for individually targeted advertising to increase the response rates of consumers.

Unlike legacy advertising methodologies that broadcast the same ads to all subscribers, targeted addressable advertising achieves a much closer pairing between the interests of consumers and the promotional messages they receive. Ads can be selected and forwarded to subscribers based on viewing preferences, hobbies, interests, demographic information, and other relevant data.

Targeted advertising works in a switched digital broadcast environment and aptly demonstrates that adoption of a switched services model not only allows cable operators to reclaim bandwidth, but offers a compelling array of additional benefits.

OVERVIEW OF SWITCHED DIGITAL BROADCAST

Switched digital broadcast is a method of delivering selected broadcast programming only within nodes, or service groups, where subscribers actively request that programming. Unlike traditional broadcast service where all programming is delivered to all subscribers all the time, switched digital broadcast delivers programming only when and where requested by subscribers.

When a subscriber in a service group wants to watch a program that is currently being delivered to other subscribers within the same group, the new subscriber simply joins the existing multicast program delivery. As a result, there is no need for additional bandwidth to deliver the same program. Using such a methodology enables the creation of virtual programming capacity without the correlated physical expense of creating and dedicated precious bandwidth resources. Figure 1 shows the bandwidth savings that can be achieved using switched digital broadcast.

Switched unicast is a form of switched digital broadcast in which each subscriber receives a unique program stream. Switched unicast consumes dramatically less bandwidth than traditional broadcast programming, and though it doesn't have the same bandwidth reclamation

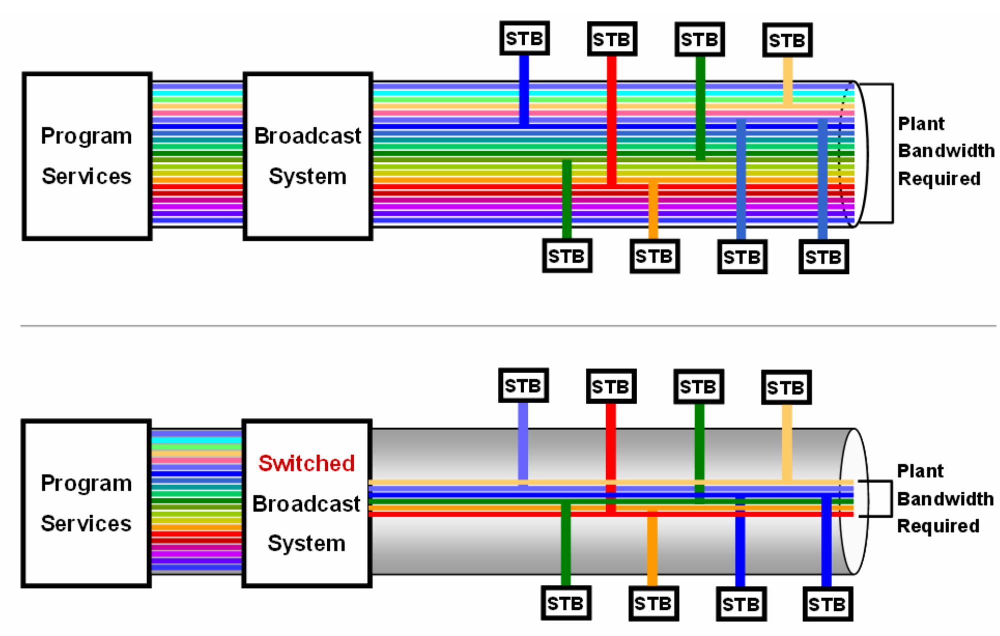

Figure 1. Switched digital broadcast significantly reduces bandwidth consumption

capability as switched multicast, it still enables cable operators to significantly increase the amount of programming offered. The greatest benefit of switched unicast, however, is in enabling operators to individualize programming and establish one-to-one subscriber relationships.

INTRODUCTION TO TARGETED ADVERTISING

Targeted addressable advertising offers cable operators the opportunity to increase advertising revenues by delivering promotional messages that more closely match subscribers' interests: marketers are willing to pay a premium rate if their ads achieve improved response rates among their intended audiences. Several recent research studies indicate that advertisers have a strong willingness to pay twice as much for an advertisement if there is a guarantee that it reaches a targeted audience (source: On Demand report, CTAM, June 2005). It benefits cable operators by allowing them to enhance advertising revenues, especially since marketers may increase the variety and quantity of ads they run. For example, an automobile manufacturer that traditionally provided one or two ads to an operator for broadcast to a large audience may choose to run five or six ads tailored to match the interests of more refined audiences.

An example of targeted addressable advertising is shown in Figure 2. In this example three subscribers, all watching the same program on the switched tier, receive different ads during the commercial breaks. Subscriber #1, an avid snowboarder in his early twenties, receives an ad about low-cost winter cabins in Lake Tahoe. During the same commercial break, subscriber #2, a thirties-something bachelor, views an ad about an upcoming motor show. Subscriber #3, an older couple planning a second honeymoon, receive information about cruises in the Caribbean.

Targeted addressable advertising systems allow subscribers to undertake ad telescoping activity. Telescoping enables a viewer to press a button on the remote control for more information or request a longer form of the ad or a brochure. Interested viewers can schedule an appointment with a retailer, receive a demo

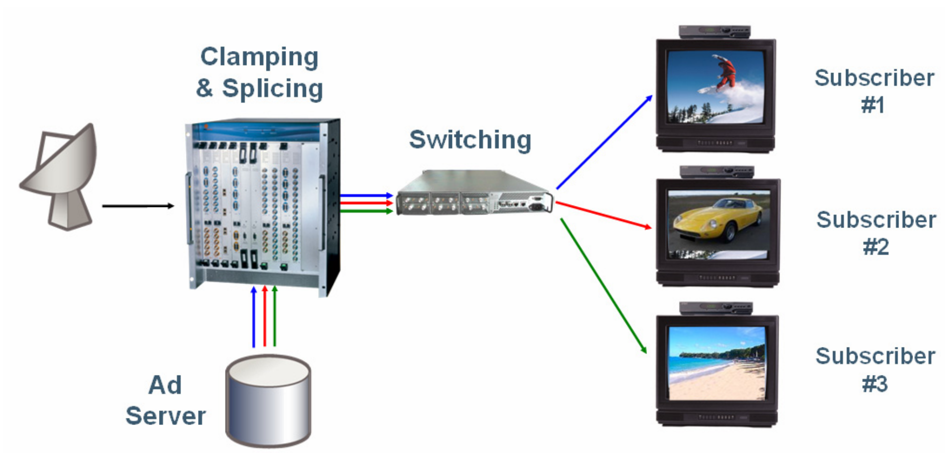

Figure 2. Delivery of targeted advertising, here within a switched digital broadcast architecture, provides ad content of personal relevance even to multiple subscribers watching the same program

in the mail, or take part in a promotional contest.

Advertising can be further enriched using PVRs (personal video recorders), networked or otherwise. These enable subscribers to pause near-live broadcasts while telescoping. Once the ad has been viewed, the viewer can return to programming, at the place where it was paused.

There is also the opportunity to increase the pool of marketers willing to advertise on television if they now are able to reach a very targeted audience independent of the program or geography.

The technical key to targeted advertising is switching. In order for subscribers watching widespread broadcast programming to see specific advertising content, that content must be dynamically directed to, or accessed by, the subscriber through switching at some level. This can occur within the cable network since unicast methods of switched broadcast are able to dynamically deliver separate copies of live digital program streams to each active subscriber, allowing each stream to be spliced with the right advertisement.

Alternatively, all subscribers watching the same content can receive the multiple ads in sync, with precisely timed dynamic digital tuning by each set-top box to reveal the correct ad. There are variations on both these network-based and subscriber-based models of ad targeting.

MATCHING ADVERTISEMENTS AND SUBSCRIBER INTERESTS

Switched digital broadcast systems track which programs subscribers are watching in real-time. This provides cable operators with information that can be used to determine which promotional messages a subscriber is more likely to be interested in. For example, a subscriber watching a sports network has a higher probability of being interested in hearing about an upcoming local sporting event.

Naturally, privacy restrictions surround the gathering of individual subscriber data, limiting cable operators to matching ads to the programs that subscribers are watching in real-time. However, the ability to build profiles about subscribers' interests adds tremendous value to targeted advertisements. Accordingly, the more enterprising cable operators could explicitly ask their subscribers about their ad preferences. In return for providing a cable operator with a list of the subjects and categories they'd be interested in viewing ads on, subscribers could receive a complimentary gift or upgrade to an expanded service package, or some other incentive.

Switched digital broadcast systems can be configured to store information about subscribers' viewing patterns anonymously. This can provide cable operators with precise viewership statistics without relying on third parties. The value of this information is high because it provides insights into the viewing patterns of all subscribers on the switched tier not just the subset of viewers that have been enlisted by audience research firms.

IMPLEMENTING TARGETED ADRESSABLE ADVERTISING

The core components of a switched digital broadcast system are:

- Clamping sub-system;

- Edge sub-system;
- Management server;
- Client software.

The clamping sub-system converts incoming broadcast programs to a predetermined constant bit rate value. Though not absolutely required, the normalization of all programs to a constant bit rate allows a faster and simpler switching mechanism, and paves the way for QAM resource sharing with similar services such as VOD. The bit rate most often chosen is one that resembles the program parameters for VOD streams. The clamped programs are multicast to all hubs using a Gigabit Ethernet transport network.

The edge sub-system replicates program streams and directs them to the appropriate edge QAM, in response to directions from the switched broadcast manager. Edge switching capabilities can be integrated into the QAMs or can be a stand-alone entity upstream of the QAMs.

Client software resides on a subscriber's set-top box. When a program is selected, the client conveys the channel request upstream along with information that uniquely identifies the node-group location of the set-top box. The client functionality can be easily integrated into the tuning firmware of future set-top boxes.

The switched digital broadcast manager uses the received channel number to identify the requested program and, consequently, the input port on the switch where the program is being received. Similarly, the switched broadcast manager uses the set-top box ID and associates the node group information to determine the downstream path that connects to the subscriber.

When an available downstream QAM and program resource are identified, the frequency and program information is returned via the downstream out-of-band channel to the set-top box, which decodes and displays the program using the normal tuning mechanisms. While the specific

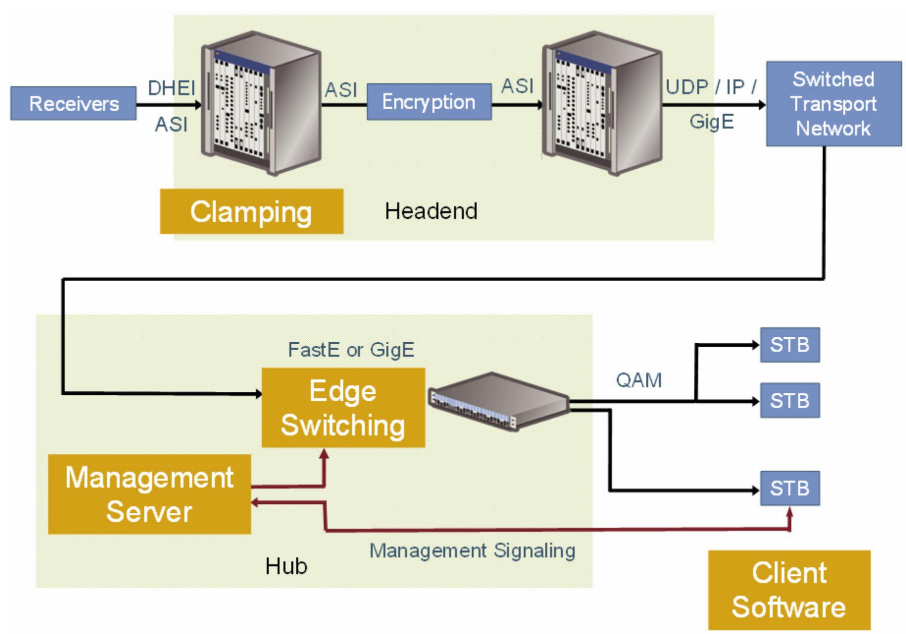

Figure 3. Core components in a switched digital broadcast network

frequency and program number may vary with time, the channel number seen by the subscriber's set-top box is unchanged.

The building blocks of a switched digital broadcast network are shown in figure 3.

When a subscriber changes programs the following sequence occurs:

1. Subscriber's set-top box sends an upstream request to the switched digital broadcast server;
2. Switched digital broadcast server allocates the bandwidth needed for the program and instructs the set-top box where to put the program;
3. Switched digital broadcast server sends a message to the appropriate QAM to join the IP multicast that has the required program;
4. Program appears on subscriber's TV set.

AD INSERTION IN SWITCHED NETWORKS

The process used to insert ads in a switched environment is similar to the process used in legacy broadcast networks except for a few minor differences. The key steps are the following:

1. Ad splicer in the program path detects an embedded SCTE 35 cue packet and sends an SCTE 30 cue request to the ADM (ad decision manager) that an ad avail will appear shortly. The splicer is typically the clamping device in the acquisition sub-system, and may also replicate the input program multiple times for geographic zoning of the advertisements.
2. ADM queries the switched digital broadcast manager to find out how many viewers in each node are watching the network. The ADM also requests the MAC addresses of all viewers in a node, enabling unique subscriber demographics to be identified for each MAC address. The added demographic information, coupled with the active programming information, provides the ADM with important data from which to select the most appropriate advertisement.
3. ADM chooses an appropriate ad selector from a list of pre-registered ADS (ad decision selector) servers. The ADM requests ads to fill the avail from the selected ADS using the DVS/629 API. It's worth noting that ad selection is traditionally based on the owner of the ad avail slot or the particular network.
4. The ADM then triggers splicing of the chosen ads into the program stream. Splicing can be implemented by a currently available splicer, or other devices.
5. After receiving verification that the ad has played and been properly spliced in and out, the ADM notifies the ADS of the success, or failure, of the insertion. The ADM can be configured to record more information about the event in its log files. This information is then recorded and sent to the traffic and billing system for invoicing, or other analysis reporting.

CLIENT-BASED AD INSERTION

There is an emerging method of targeted addressable advertising that includes client-based ad selection. In this form of targeted advertising the subscriber's set-top box contains an agent that selects the most relevant advertisement from a pool of available ads. This requires that a set-top client be deployed and maintained with relevant subscriber profile data and be capable of dynamically tuning when time for the insertion occurs. This sort of functionality must be programmed for every digital client environment in use, but such functionality is generally achievable on most digital devices. The hypothesis here is that the STB (set-top box) contains household information that is additive to the demographic and geographic information, thereby increasing the potential relevancy of a targeted advertisement. These extra targeted ads can either utilize additional plant bandwidth or can be stored on DVR set-tops.

One approach to client-based targeting, shown in Figure 4, is based on dynamic tuning within a multiplex. In this scenario a digital multiplex does not carry its maximum load of programs, but instead maintains some spare capacity so that whenever time for an advertisement arises for any program, multiple ads are inserted within the multiplex. Each set-top box dynamically accesses the digital ad corresponding to its profile, and after the ad returns to the program that was being watched. This method is relatively easy to implement, as it is not contingent on switched digital broadcast deployment, but it does present bandwidth challenges in reducing the programming load of the multiplex. It also requires detailed synchronization of insert times within a multiplex so that not too many individual ads are required to be carried simultaneously. For similar reasons, the variety of profiles that can be simultaneously supported is limited.

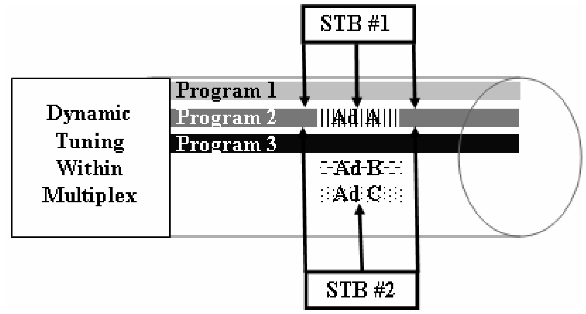

Figure 4. Client-based targeting in which each set-top box dynamically access a digital ad within the same multiplex as the programming

An alternative method of client-based ad insertion, shown in Figure 5, dedicates a complete multiplex to carriage of the advertisements. In this case a carousel methodology can be utilized for a constant rotation of ad content. When time arises for a targeted ad, the subscriber's set-top box re-tunes to the multiplex with advertising and selects the ad best matching the subscriber's profile. This allows for full programming multiplexes, only requiring that a single multiplex, or a few multiplexes, be reserved for ad carriage. It also enables a large number of overall profiles to be maintained, and a large number of programs to simultaneously receive targeted splicing. There is some efficient sharing among these, as is illustrated in the image, by STBs #1 and 2 accessing the same ad stream despite watching different programs. It does increase the burden on the set-top box to seamlessly tune to the multiplex with advertising and re-tune to the original programming. And ultimately there is some limitation on the advertising that can be delivered.

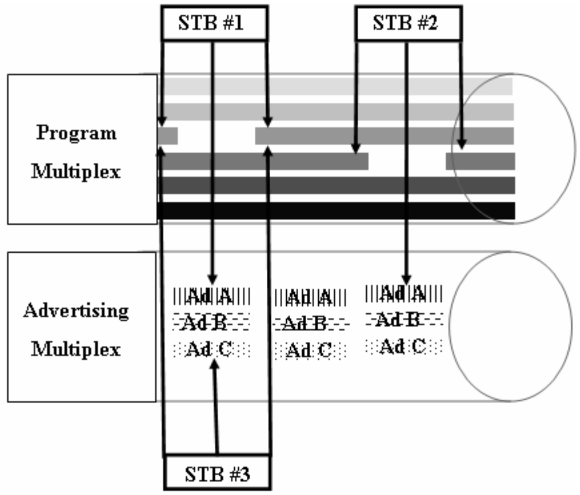

Figure 5: Client-based targeting in which each set-top box dynamically tunes to a different, ad-dedicated multiplex at times for insertion

Targeted addressable advertising practices are just emerging and there are several other models being considered besides those described here. For example, some are considering loading ad content onto set-top boxes' PVR (personal video recorder) hard drives so that targeted ads can be dynamically accessed without any need for altered tuning or additional content carriage. A drawback of this approach is that it's only applicable to the subset of viewers who have PVR-enabled set-top boxes.

MAXIMIZING NETWORK RESOURCE UTILIZATON

While targeted advertising has clear business benefits, cable operators must first consider how to cope with the additional load burdened on their system in terms of bandwidth consumption, headend resources and/or set-top box processing. For example, consider a reasonably sized cable system deploying switched digital broadcast. There could be more than 500 service groups, each switching digital programming within 8-12 QAMs. This can require about 60,000 unicast streams each with different ads. This will overburden legacy equipment. Part of the gain can be achieved through the fact that there will be replication of subscribers watching the same programming, or receiving the same ads, or both. These can be leveraged to consolidate certain steps of media processing and switching.

As resources scale up, operators can consider dynamic approaches to maximize precision of targeting when total load is lighter, and maximize efficiency when it's heavier. In this case switched unicast can be used when feasible, but when viewership is high, the plant can dynamically revert to switched multicast in order to conserve on stream capacity and resources like splicing.

An potential hybrid approach could be for a multicast to be broken into sub-groups by like profile, each of which receives a unique stream, so that while ads are not personally targeted, they are not bluntly broadcast either. Since many programming networks appeal only to a few demographics, this methodology can fulfill both the targeting and efficiency objectives of operators.

There are other efficiency gains that can be achieved through switching. If no subscriber is watching a specific program, there is no reason to run ads or perform splicing associated with that program, thereby saving not only network capacity but splicing resources too. The adoption of targeted advertising from a PVR within the set-top box can also enhance efficiencies, by loading ads to the DVR during periods of limited live plant utilization.

In a switched environment, a cable operator has insights into how subscribers and the program they are watching correlate.

This introduces some interesting opportunities, including the following:

- When there are multiple viewers, one option is to select a dominant demographic and use that as the basis for selecting an ad. Since people within a specific demographic are more likely to watch the same programs, this approach may be very successful. After all, this is essentially the basis of traditional cable ad sales;

- Additionally, when there are multiple viewers, an SDB system can determine if spare capacity is available and, if so, forward advertisements targeted at the unaddressed demographic groups within the node;

- If only one viewer of the network in SDB or if operating in unicast or nPVR, a one-to-one targeting opportunity appears, enabling the ultra-personalization of ads, even including the insertion of the subscriber's name in the promotional message;

- Finally, if no one is watching a specific program, there is no reason to run the ad associated with that program, thereby saving not only network capacity but splicing resources too.

In general though, targeted addressable advertising will increase the quantity of splicing resources required. For example, a marketer that was previously content to provide a single ad for a large audience with poor response rates decides to run multiple ads to more closely address the interests of her targeted audiences.

A smarter stream allocation methodology can be utilized by comparing the availability of ad content and the limited demographics that need to be targeted at any given time. Consider, for example, that many cable networks already target certain demographics, increasing the likelihood that a new viewer will be similar to current viewers. With that in mind, the following algorithm can be applied when a subscriber requests to a stream:

1) If a program is requested that is already being broadcast to one or more viewers, the new subscriber can be added to a multicast session - the SDB system compares the new user's demographics to all current multicasts of a particular network, and if a close enough match can be made, the new user joins that stream.

2) If a program is requested but either the demographics of the new viewer don't match those of any existing multicast session or the program is not already being broadcast, a new stream is generated.

Depending on the available ad inventory and current usage of a given network, this will utilize much less hardware while still making the most of unicast targeting ability.

TARGETING WITHIN ADDITIONAL SERVICES

On-demand networks, also referred to as non-linear networks, can attain similar advertising benefits as switched digital broadcast networks. An nPVR (networked-PVR) model is similar to switched unicast one, but uses a VOD server to create stream as needed. In this model, networks are acquired in real-time, with an on-demand session being created between the VOD server and the subscriber's set-top box.

Advertisers have been particularly concerned that ads are not included in many types of on-demand programming, or can be skipped over. However, subscribers have revealed strong interest in free on-demand content, and some operators are beginning to provide them with ad support. Most likely there will not be much tolerance for the same quantity of advertising that generally occurs in live programming, but targeting ads could make up for this. Subscribers would be more willing to sit through fewer, more on-target ads. Advertisers would be prepared to pay more in order to be an exclusive, or one of few, sponsors of personally selected content to an audience that can be captive if fast-forward capabilities are suspended during advertisements.

The intelligence of VOD systems and unicast switching of them to subscribers allows the switched unicast method of ad targeting to be easily replicated for on-demand. Certain of the client-side models could also be utilized. Also, on-demand order patterns could be combined with other sources of subscriber data to enhance profiling.

Ads can be effectively inserted into on-demand programming in a variety of ways. These include:

- Bookend ads;
- Interstitial ads;
- Pause ads;
- Telescoping ads.

Bookend ads, also known as bumper ads, appear at the beginning and end of a VOD program. Although akin to the commercials that accompany the start of a movie on a DVD, bookend ads do not need to be random, but can be targeted based on information gathered from viewer surveys, or other relevant data.

Interstitial ads can be spliced into on-demand content such as a past season of a popular TV program, during what had been the commercial breaks in the original broadcasts. Additionally, the use of nPVRs enables interstitial ads to be paused during near-live broadcasts, allowing a subscriber to undertake some type of telescoping activity.

Interstitial ad replacement may actually be an imperative in some PVR content as the original ad content may have expired or no longer be relevant.

Pause ads can be used any time a subscriber pauses an on demand session. Some studies have shown that pause ads are the least intrusive to the subscriber, as they have already self-interrupted their viewing session. These studies have also shown multiple pauses per program giving a rich new source of ad inventory.

Telescoping enables a viewer to press a button on the remote control for more information or request a longer form of the ad, or peruse a topic of interest. PVR technology, networked or otherwise, enables subscribers to pause near-live broadcasts while telescoping. Once the ad has been viewed, the viewer can return to programming, at the place where it was paused.

In addition to live and on-demand video programming to televisions, cable operators can also consider additional directions to expand targeted advertising throughout expanding triple-play and quad-play activities. Data for ad targeting can extend to services like Web surfing. And targeted insertion can occur to additional services including on-line and mobile video, or data services in which cable operators can aim to capture some of the success achieved in search advertising.

ENHANCING FUNCTIONALITY OF TARGETED ADS

Non-linear programming and switched digital advertising are two of many innovations increasingly available on digital cable networks. Multiple interactive and multimedia functions can be leveraged so that when an ad is targeted to carefully selected viewers, those viewers are well positioned to easily respond.

Targeting can be combined with logo insertion or other data overlays to further customize content. This could include messaging to appeal to different profile types viewing the same ad. Regional data could be overlaid on video to indicate a local retailer or provider.

Interactivity can make television advertising completely actionable, extending from just making an impression to effecting a transaction. As interactive technologies mature, these could be used to request more information, get a call from a provider, or even effect a purchase. There is a multi-billion dollar market of direct response advertising and home television shopping with even bigger markets achieved on the Web. A subscriber with a remote control in hand could prove to be fully interested in making a purchase without ever departing the television experience. There could also be means for a subscriber to port the interaction of an interesting video ad to a PC or portable device.

Actionable advertisements could enhance operators' business models in a number of ways. They could claim commission on sales or even sell directly themselves. Certainly their own offerings including customer support and premium service sales could occur by interactivity keying off of well targeted ad placements. Additionally, the Internet model of pay-per-click could be implemented such that advertisers are charged a premium whenever interactive functionality on an ad is accessed.

AD ENCODING FOR SWITCHED AND STORED NETWORKS

Switched and VOD advertising require ads to be encoded differently than on traditional ad insertion systems. Linear systems typically encode ads at full resolution with typical bit rates of 6 Mb/s CBR. In the case of stored content, ad encoding methods need to more closely follow VOD encoding specifications.

The industry is moving towards encoding all advertisements based on the CableLabs VOD encoding specifications. As the current switched digital broadcast servers and session resource managers use CBR 3.75 Mb/s standard definition and 15 Mb/s high definition streams, encoding all ads at this specification allows development of simpler splicers that are able to handle larger numbers of splices.

New session resource managers will be capable of handling an assortment of CBR bit rates. One method is to divide the current 37.5 Mb/s available in to sixty segments and allocate the bandwidth on a segment-by-segment basis. However, this may require some additional work from the SBM or session resource manager, if the network has been allocated less than the ads' 3.75 Mb/s or the 3.5 Mb/s that is typically in the ad. In this case a request will need to be made to the SBM for more bandwidth. Another option might be to encode ads at a slightly lower video bit rate so that the video and

AC3 audio is closer to 3.125 Mb/s, if twelve programs need to be accommodated in each QAM.

INTELLIGENT AD SCHEDULING SYSTEMS

With a large number of avails from the current linear broadcast, switched digital broadcast, nPVR and various VOD advertising methods such as pre- and post-bookend ads, the current systems were not intended to effectively sell advertising in this environment. One place to look for insight is Internet advertising, which various methodologies, including banner ad and video that can be dynamically selected. There are some unique issues currently with cable advertising that are not faced in the Internet world that need to be considered. The most important issue that Internet advertising has not really had to deal with is provisioning high quality content on to the appropriate server, it is very easy to pull an ad from a web site thousands of miles away, but the bandwidth does not yet exist to do this for content that is megabytes in size.

Another issue is that in this new model the operator may decide to start serving targeted impressions for content owners in addition to advertising on content and avails it has rights to. It is likely that these content owners will have certain vendors that they want to work with in this space. It is important that the ADM allow and control multiple ADS servers.

These new advertising methods will also likely require new sales and campaign management tools to effectively manage these campaigns as initially the number of ad avails will not be known. As the system plays ads though, it will need to acquire statistical data as to the expected inventory that will be available in the future.

The campaign manager should also be able to update current campaigns if they are not executing properly. It may need to alter the target demographics or other parameters to reach the desired advertising goal.

EXPANDING AD SALES MODELS BEYOND LOCATION

Other achievements in Internet advertising that could be considered for television include national campaigns and online sales of placements. Cable operators are beginning to sell ad insertion on macro-zone, in addition to micro-zone, geographic bases. This means that a national advertiser could purchase time across all of an operator's systems. Combining this with targeting could enhance this business model by allowing that advertiser to select its advertisement best catering to each profile type across an operators' subscribership, distinguishing household from household, nationwide.

In addition to national advertisers being interested in accessing all of an operator's subscribers, tiny businesses could have the same interest. Their appeal could be so precise that it's not worth their purchasing content within a single system, but there is economic justification across all of an operator's footprint. Consider a destination event for a small but highly dedicated affinity group: if the provider could access a highly fragmented and esoteric collection of people with aligned interests for video advertising, that could be highly valuable.

A new, technology-based sales and marketing method could enhance such nationalized ad sales. Real-time transactions could be implemented, resulting in an ad being instantly sourced and distributed by transport nationwide.

CONCLUSION

Adoption of switched digital broadcast networks continues to gain momentum in the industry. As switched digital broadcasting networks become more pervasive in the marketplace it enables a whole new level of enhanced services of which targeted advertising is a powerful force.

As the technology becomes more widely deployed, it will increase the ability of cable operators to successfully compete for local advertising dollars, especially compared to operators using legacy broadcast advertising models.

New technologies such as VOD and PVR have led to speculation of the decline of the TV advertising industry. However, those active in the advertising business, including cable operators and the marketers who run ads, have emerging technologies at their disposal. These can be leveraged to assure that advertisements are relevant and valuable to their audiences.

The Internet industry has flourished with advertising by directing relevant and actionable content to consumers. Switching technologies applied to video can be leveraged by cable operators to similar effect. The result is targeted advertising wherein subscriber information is incorporated into the determination of what ads are delivered, or selected, to the subscriber. This approach can be combined with other industry advances such as multi-service integration, interactivity, and ongoing data analytics to continue advancing the prominence of television as the seat of the most compelling advertising platform.

MANAGEMENT OF STIMULATED RAMAN SCATTERING IN CATV WDM REVERSE PATH SYSTEMS

Bryant A. Best
Scientific-Atlanta, Inc.

Abstract

Fiber nonlinearities can significantly limit the performance of WDM optical systems. Crosstalk due to Stimulated Raman Scattering (SRS) can potentially reduce carrier-to-interference ratios to unacceptable levels. The effects of SRS include crosstalk that may produce interference at levels up to 30dB below desired carriers. The reverse path is particularly susceptible to this impairment.

This paper presents theoretical models for SRS. Empirical and simulation data is presented showing impairments from SRS in fiber links. Management of the factors that contribute to SRS is a critical part of system design. Specific techniques are described that will ensure that acceptable levels of system performance are obtained.

INTRODUCTION

New services requiring two-way traffic on HFC networks (such as cable telephony, VoIP, cable modem and two-way set-top proliferation, IPPV, etc.) have created additional demands on reverse path transmission capability. System architecture considerations, focused on the management of this additional traffic, must take into account frequency allocation, homes passed per node, and segmentation in the reverse path. Until recently, the generally accepted standard fiber optic reverse plant design was built upon the FP (Fabry-Perot) or DFB (distributed feedback) laser at 1310nm or a DFB at 1550 nm with a dedicated return fiber. The emergence of CWDM laser technology allows for a low cost strategy to expand the capacity of the reverse path plant by allowing multiple (typically up to eight) return path transmitters to share a single return fiber. This approach serves to "future proof" the network by allowing additional transmitters to re-use the same fiber as well as facilitate a fiber-deeper architecture in which the node supports a smaller number of subscribers, thereby boosting customer performance as transmitter bandwidth is shared with fewer end-users. CWDM offers these benefits at a much lower cost, size, and power consumption than an equivalent DWDM solution.

A WDM solution is not without its drawbacks, however. Optical losses from the combining and splitting of wavelengths at the HFC network endpoints, as well as degrading signal distortions due to channel-to-channel interactions caused by nonlinearities in the fiber itself must be understood and accounted for to achieve acceptable network performance.

It is convenient to assume that each wavelength on a multi-channel WDM optical system will behave and perform as if it was on its own dedicated fiber. A system designer that uses this viewpoint would calculate loss budgets, compute SNR, etc. for each service, and would expect service performances accordingly. The well-known performance-limiting degradations suffered by a telecom system's high-powered, tightly spaced, amplified DWDM systems seemingly would not apply to CWDM in an HFC reverse path application. Given a typical reverse path CWDM system's modest launch power into a shorter fiber span, with wide channel spacing and narrower modulation bandwidths, one would expect each of the wavelengths to behave as an independent link. Unfortunately, the reality is that even these systems can have wavelength interaction affecting system capability.

Interactions between optical signal carriers propagating along a fiber can cause a multitude of potentially degrading effects and much work has been done to understand their behavior and impact [1]-[4]. Long haul digital DWDM transmission system designers, whose challenge is to extract maximum link performance under extreme optical budgets, are well aware of the pitfalls and consequences of improper management of fiber effects.

Any fiber optic system must deal with the possibility of one or more of the following nonlinear fiber impairments.

- Four Wave Mixing (FWM, is similar to CTB with optical carriers)
- Self-Phase Modulation (SPM)
- Cross-Phase Modulation (XPM)
- Stimulated Brillouin Scattering (SBS)
- Stimulated Raman Scattering (SRS)

Each of these nonlinearities is dependent upon system parameters that govern how much of an effect it will have on system performance. Individual transmitter launch power, number of wavelengths, total power in the fiber, transmission distance, modulation type and bandwidth, laser wavelengths and spacing all have an influence on which nonlinear mechanism will have significance on the system. These parameters, as they apply to reverse plant systems using CWDM, play into making one type of fiber non-linearity dominate the others and have a significant influence on system performance. Our studies show that reverse path CWDM systems are most susceptible to stimulated Raman scattering, or SRS.

STIMULATED RAMAN SCATTERING

Stimulated Raman scattering is a fiber nonlinear effect which results in the transfer of optical power from one wavelength to another. This is characterized as a broadband scattering effect [6] and it influences all other wavelengths on a fiber within approximately 125 nm due to an intrinsic property of fused silica glass. SRS occurs when the laser signal is scattered by natural molecular vibrations in the fiber (phonons). Interaction between the laser signal and the vibrating glass molecules scatters light from the signal in all directions. Nearby molecules absorb the scattered light and re-emit a photon with energy roughly equal to the original photon. SRS-induced crosstalk (Raman crosstalk) occurs when another signal at a different wavelength is co-propagating on the fiber and causes the molecule stimulated by the first wavelength to emit a photon at the second. As photons of shorter wavelength light (higher frequency) contain higher energy, the transfer of energy due to Raman scattering causes energy to transfer from shorter to longer wavelengths. In terms of measured crosstalk, the effect is equal in both directions, meaning the shorter wavelengths experience crosstalk by power reduction, the longer wavelengths by power addition.

For Raman crosstalk to occur, optical power must be simultaneously present between interfering signals. This implies that on digital systems, crosstalk will only

$$Crosstalk(SRS) \approx 10\log\left\{\left(\frac{\rho_{srs} g_{12} P_{\text{int}}}{A_{eff}} \cdot \frac{m_{\text{int}}}{m_{CATV}}\right)^2 \frac{1+e^{-2\alpha L} - 2e^{-\alpha L}\cos(\omega d_{12} L)}{\alpha^2 + (\omega d_{12})^2}\right\} \quad (1)$$

$$\Theta_{SRS} = \tan^{-1}\left(\frac{-\omega d_{12}}{-\alpha}\right) + \tan^{-1}\left(\frac{e^{-\alpha L}\sin(\omega d_{12} L)}{e^{-\alpha L}\cos(\omega d_{12} L) - 1}\right) (2)$$

where,

$Aeff$	effective area of fiber
ρ_{srs}	effective polarization overlap factor
g_{12}	Raman gain coefficient
P_{int}	optical power of the interfering signal
m_{CATV}	optical modulation index (OMI) of signal
m_{int}	OMI of the interfering signal
α	attenuation coefficient of the fiber
d_{12}	group velocity ($\approx D (\lambda_1 - \lambda_2)$)
D	dispersion coefficient
L	fiber length
ω	RF angular frequency

occur when both wavelengths are transmitting a '1'. No crosstalk occurs when either path is transmitting a '0'. In analog modulated systems, where the optical carrier is always present, Raman crosstalk can be present continuously.

SRS Equations

Phillips and Ott studied the governing mathematics of channel-to-channel interference caused by Raman crosstalk on WDM CATV systems, and derived equations for magnitude (in dB) and phase (in radians) of crosstalk between two wavelengths [5].

The Raman gain coefficient, g_{12}, denotes the Raman gain between interacting wavelengths and is inherent to the fiber. It is typically approximated by a line with slope of 6.0E-14 m/W from 0 to 100 nm [7]. The gain peaks at approximately 100nm of wavelength separation and falls off sharply thereafter.

From analysis of Equation (1) it can be noted that Raman crosstalk is predominantly dependent upon optical power of the interfering signal, wavelength separation between the signals (see P_{int} and g_{12} terms in numerator), and the modulating frequency of the interfering signal (see ω term in the denominator).

SRS CROSSTALK MEASUREMENTS

CATV systems are particularly sensitive to SRS effects due to demanding CNR requirements. The focus of SRS-induced crosstalk has primarily been on video overlay PON (Passive Optical Network) architectures, where Raman crosstalk exists between the 1490 nm data carrier wavelength and the 1550 nm wavelength carrying amplitude modulated video signals, and DWDM forward path video distribution systems [8]-[12]. SRS effects on a CWDM return path system have been somewhat overlooked as most services on these systems have greater tolerance of noise and unwanted interference than an AM-VSB video signal. Despite the greater tolerance to impairment, SRS effects on services in a CWDM return path system must be quantified and limited to acceptable levels.

Experiments were constructed to show crosstalk as a function of the critical parameters of optical launch power, wavelength spacing, and modulating frequency on a typical CWDM CATV reverse path system.

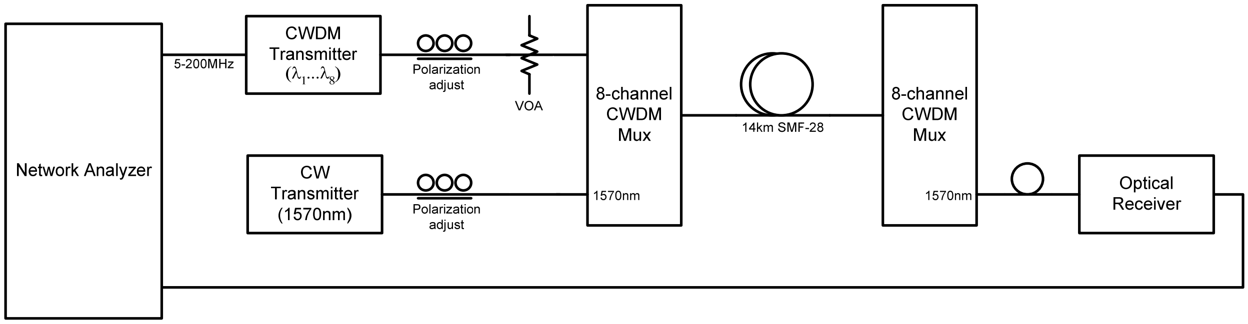

Figure 1. Experimental setup with two transmitters over 14km SMF.

Test Setup

The measurement setup shown in Figure 1 consists of two CWDM DFB transmitters combined through a CWDM mux to 14km of single mode fiber. One wavelength is directly modulated with a swept RF signal while the other is unmodulated (CW). An optical attenuator, inline with the modulated wavelength, varies the optical power of the interfering channel. A CWDM demux at the end of the link selects the test channel wavelength which is fed to the receiver. Polarization was adjusted on each wavelength for worst case behavior. Crosstalk due to the optical mux and demux were determined not to contribute to the SRS crosstalk measurements. Crosstalk measurements are made by measuring the ratio of RF signal power on the unmodulated CW wavelength to the RF signal power on the modulated wavelength, with each referenced to the same average power.

Test and Simulation Results

Figure 2 shows crosstalk from 5 to 200 MHz between two CWDM transmitters spaced 100 nm apart (1470 nm and 1570 nm) with 3 dBm launch powers over 14 km of single mode fiber. Combining mux loss was approximately 2 dB per port. Worst case crosstalk was measured at -46 dB at 5 MHz. Local nulls seen in the crosstalk over frequency are due to the cosine term in the numerator of Equation (1). Changes in fiber length or wavelength spacing of the carriers will result in a shift of the null frequencies.

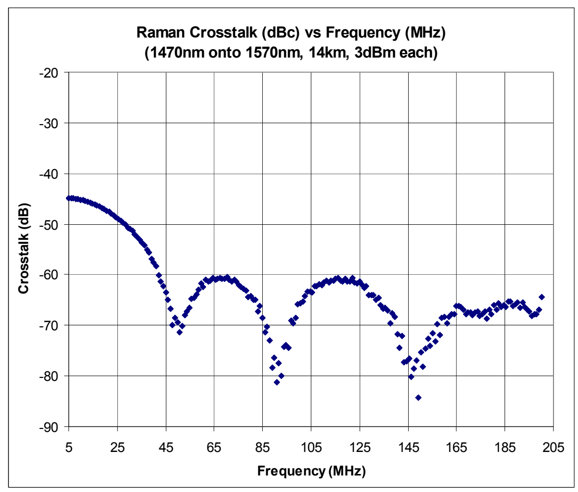

Figure 2. Measured 1470nm crosstalk onto 1570nm (3dBm each, ~2dB mux loss, 14km SMF-28)

Figure 3 shows the dependence of crosstalk on optical power and wavelength separation. The plot shows measured crosstalk at 5 MHz onto a CW channel at 1570 nm from each of the other wavelengths in an eight channel CWDM system (1470 nm to 1610nm) with the output power of the interfering signals varied from -6 dBm to +3 dBm. It is seen that SRS crosstalk increases 2 dB for each 1 dB increase in interfering wavelength power and increases with frequency separation. The plot also illustrates that crosstalk from a neighboring wavelength is a function of the difference in wavelength and is largely independent of whether the interfering signal is higher or lower in wavelength.

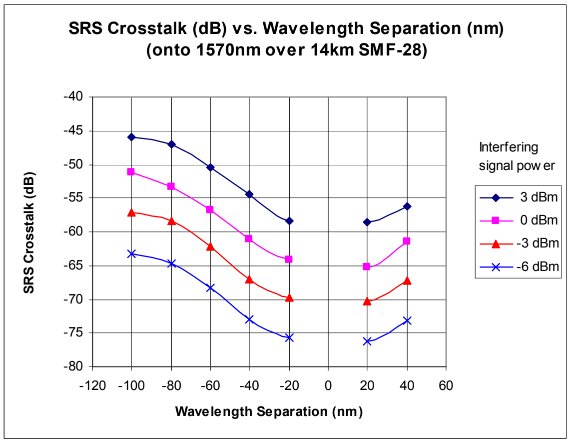

Figure 3. Measured crosstalk at 5MHz onto 1570nm vs. wavelength separation for various interfering signal output powers (1570nm at +3dBm out, ~2dB mux loss, 14km SMF-28).

At a given modulation frequency, total SRS-induced crosstalk on a system will be due to the vector addition (amplitude and phase) of all interfering signals and governed by equations (1) and (2). Simulations using these equations predict worst case crosstalk approaching -30 dB in an eight-wavelength reverse path CWDM system over 14 km when the optical power per wavelength is +7 dBm. Additional increases in transmit optical power or fiber length can increase crosstalk even further.

As shown, Raman crosstalk is strongly dependent upon optical power in the fiber. From Equation (1), this crosstalk is also shown to be a function of fiber length (see term L). Using Equation (1), Figure 4 shows simulated plots of crosstalk between wavelengths of 1470 nm and 1570 nm at various modulating frequencies as a function of fiber length. For each frequency plotted, the crosstalk is shown to increase from a minimum (with no fiber) to an initial maximum value and then vary about a final crosstalk value as fiber length is increased. Others have shown under certain conditions, using a more complex model [12], that crosstalk may reach another peak (in excess of the initial peak) at a longer fiber distance.

Raman crosstalk is also very sensitive to signal polarization states [12]. Crosstalk is maximized between two wavelengths when the signals are in perfect polarization alignment during propagation. As the polarization overlap factor between the interfering wavelengths can vary between 0 and 1, we have assumed the overlap to be 0.5 for simulation purposes. In actual systems, time-varying polarization states between the signals while traversing the fiber span will cause the measured crosstalk to vary.

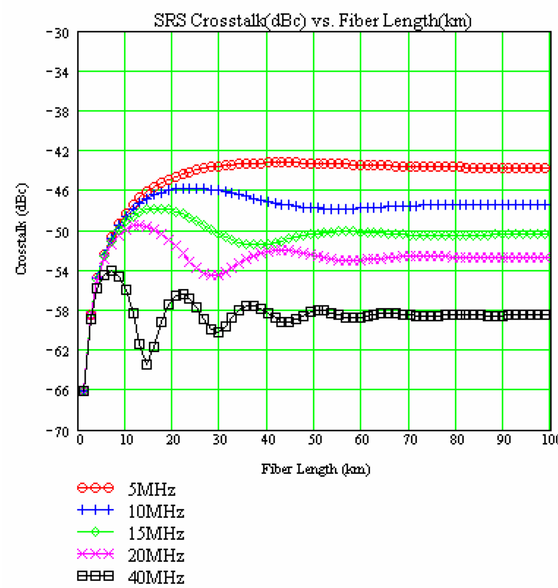

Figure 4. Crosstalk vs. fiber length for various frequencies (1470nm to 1570nm at 3dBm output per wavelength, ~2dB mux loss, SMF-28)

CWDM SYSTEM DESIGN GUIDEINES

Desired crosstalk performance for a WDM reverse path system is generally better than -40 dB. Careful system planning is necessary to achieve this performance for CWDM systems where crosstalk due to SRS tends to dominate. SRS crosstalk can be minimized by paying heed to the critical parameters that govern it. Below are guidelines to help reduce crosstalk effects based on the system parameters of optical power, wavelength spacing, and modulating frequency.

- Optical launch power per transmitter should be kept to the minimum necessary to achieve required CNR goals. Keep in mind that CNR margin may need to be added due to crosstalk effects.
- Services requiring high carrier-to-interference ratios should be located in that part of the RF spectrum not occupied by modulation on another wavelength.
- If all wavelengths are not being utilized, use those wavelengths that are closest together or greater than 100 nm apart.
- Place the most sensitive services at the high end of the RF modulation spectrum where crosstalk effects are reduced.

REFERENCES

[1] Rogers H. Stolen, “Nonlinearity in Fiber Transmission,” Proceedings of the IEEE, vol. 68, issue 10, pp. 1232-1236, October 1980.

[2] W. J. Tomlinson, Rogers H. Stolen, “Nonlinear Phenomena in Optical Fibers,” IEEE Communications Magazine, vol. 26, no. 4, April 1988.

[3] Andrew R. Chraplyvy, “Limitations on Lightwave Communications Imposed by Optical-Fiber Nonlinearities,” Journal of Lightwave Technology, vol. 8, no. 10, pp. 1548-1557, October 1990.

[4] Dogan A. Atlas, John J. Kenny, “Multiwavelength Analog Video Transport Network,” International Topical Meeting on Microwave Photonics 1999, MWP ’99 Technical Digest., pp. 189-192 vol.1.

[5] Mary R. Phillips, D. M. Ott, “Crosstalk due to optical fiber nonlinearities in WDM CATV Lightwave systems,” Journal of Lightwave Technology, vol. 17, pp. 1782-1792, Oct. 1999.

[6] R. Ramaswami, K. N. Sivarajan, “Optical Networks – A Practical Perspective,” Second Edition, San Francisco: Morgan-Kauffman, 2002. pp. 80-81.

[7] R. Ramaswami, K. N. Sivarajan, “Optical Networks – A Practical Perspective,” Second Edition, San Francisco: Morgan-Kauffman, 2002. pp. 326-329.

[8] H. Kim, K. H. Han, Y. C. Chung, “Performance Limitation of Hybrid WDM Systems Due to Stimulated Raman Scattering,” IEEE Photonics Tech. Letters, vol. 13, no. 10, October 2001.

[9] Dogan. A. Atlas, "Nonlinear optical crosstalk in WDM CATV systems", Nanostructures and Quantum Dots/WDM Components/VCSEL's Microcavaties/RF Photonics for CATV HFC Systems, 1999 Dig. LEOS

Summer Topical Meetings, 1999, pp. IV23-IV24.

[10] F. Coppinger, L. Chen, D. Piehler, “Nonlinear Raman Cross-Talk in a Video Overlay Passive Optical Network,” Optical Fiber Communications Conference, 2003. OFC 2003, 23-28 March 2003, Page(s): 285 - 286 vol.1.

[11] Z. Wang, A. Li, C. J. Mahon, G. Jacobsen, E. Bodtker, “Performance Limitations Imposed by Simulated Raman Scattering in Optical WDM SCM Video Distribution Systems,” IEEE Photonics Tech. Letters, vol. 7, no. 12, December 1995.

[12] F. Tian, R. Hui, B. Colella, D. Bowler, “Raman Crosstalk in Fiber-Optic Hybrid CATV Systems With Wide Channel Separations,” IEEE Photonics Tech. Letters, vol. 7, no. 12, December 1995.

Bryant A. Best received the Bachelor and Master of Science degrees in electrical engineering from the Georgia Institute of Technology in 1989 and 1992, respectively. From 1992 to 1999, he was a design engineer at Scientific Atlanta, Inc. working in the Optoelectronics Group where he developed CATV optical transmitters and receivers. In 1999, he joined Ciena Corporation to develop OC-192 transceivers for long-haul telecom applications. In 2002, he rejoined Scientific Atlanta and is currently working on research and product development in the Access Networks Group.

MIGRATING TO IMS AND PACKETCABLE 2.0: HOW TO TRANSITION TO NEW MULTIMEDIA AND FIXED MOBILE APPLICATIONS

Jonathan Rosenberg
Cisco Systems, Inc.

Abstract

With ever-increasing competition from telcos, "over-the-top" competitors, and direct broadcast satellite, Multiple System Operators (MSOs) must differentiate with consumer-driven, rich IP multimedia services. How can operators deliver advanced offerings like fixed mobile convergence while simultaneously improving the operational efficiency of their networks? How can MSOs deliver these services today, while ensuring a future roadmap to a next-generation infrastructure?

This paper reviews the forces that are influencing the development of relevant standards to meet these challenges. The standards include IP Multimedia Subsystem (IMS), Session Initiation Protocol (SIP), PacketCable™ 1.x, PacketCable 2.0, and PacketCable Multimedia™ (PCMM). Also included is an example of how an MSO can gracefully migrate from PacketCable 1.x voice services to IMS/SIP multimedia services and the key benefits that such a gradual transition offers.

TECHNOLOGY DRIVERS

As evidenced by the joint ventures between major cable providers and mobile carriers, MSOs are serious about adding mobile services to their cable bundles. Maintaining parity with telcos (with strong mobile offerings already) and changes in consumer media consumption patterns have brought attention to the technology alternatives for multimedia services over cable infrastructures. More and more subscribers want their mobile phones to serve as their primary phones and want the ability to access Internet, video, and other multimedia services from their handsets, PCs, or TVs.

Talk of "triple play" services – data, voice, and video – has evolved to "quadruple play" with the addition of wireless services. But it would be more accurate to describe the vision in terms of delivering "triple play on the move." With that vision comes the desire to introduce blended services such as video phone, multimedia chat, and gaming to a myriad of devices. Subscribers also want to simplify the hodge-podge of multiple phone numbers, subscriptions, passwords, buddy lists, e-mail addresses, etc. In short, they want to be able to get whatever services they want, whenever they want them, regardless of where they are.

To remain competitive, MSOs must provide a variety of new multimedia services in the near term, and in a faster, more efficient way. Examples of upcoming services include "one-number" phone service, unified messages (combining e-mail and voicemail), dual-mode cellular Wi-Fi handsets, video on demand (VoD) streaming to cell phones, and cross-device functions such as controlling electronic program guides with a cell phone, or viewing caller ID using a set-top box (STB).

Looking further out, MSOs will be asked to deliver next-generation capabilities. Examples include buddy lists that span multiple devices (mobile phones, soft phones, and dual-mode phones), content on demand to any device, content sharing (between a personal video recorder, a cell phone, or other device), and remotely programmable personal video recorders for cell phones.

To deliver services that compete in this rapidly changing market, MSOs need an infrastructure that supports a wide range of devices including soft clients, PDAs, STBs, and game consoles. Services must support presence (the ability to know if a person is available for communications) and identity management and personalized services. Furthermore, high-quality voice and video applications require built-in quality of service (QoS) capabilities in the network. Even more challenging, MSOs must ensure they also continue to provide their existing primary-line voice services without interruption.

STANDARDS

As with any technology that spans multiple vendors, multiple network architectures, and multiple media, standards must play a vital role in meeting the challenges faced by today's MSOs. In particular, there is a functionality gap between today's PacketCable 1.x specification and what is needed. In response, CableLabs® and a consortium of participating industry players have evolved the Network-based Call Signaling (NCS) standard and developed PacketCable 2.0.

Comparing the old and the new, PacketCable 1.x defines a centralized architecture based on a variant of the Media Gateway Control Protocol (MGCP), called Network Call Signaling (NCS). Call management servers (CMSs, or softswitches), media gateways, and Media Terminal Adapter (MTA) endpoints are used to mimic the functionality of the PSTN. This standard supports only MTA endpoints (for connecting legacy phones and facsimile machines), and does not go beyond POTS over IP at the application and service levels.

PacketCable 2.0, by contrast, will include SIP, a protocol standardized by the IETF, to achieve a distributed multimedia signaling architecture with presence and identity capabilities. A broad range of SIP endpoints are supported within a PacketCable 2.0 environment, including SIP MTAs, SIP phones and soft phones, STBs, video phones, PCs and game consoles, and dual-mode cellular Wi-Fi handsets. SIP also enables enhanced services including videoconferencing, presence-based messaging, and IP Centrex functions.

SIP, however, is a protocol, and not a network architecture. A consortium of players in the mobile market, the 3rd Generation Partnership Project (3GPP), developed the IP Multimedia Subsystem (IMS) specification. IMS is a standardized network architecture for SIP-based services. It has been adopted by other standards bodies, major carriers, and equipment vendors alike.

While IMS originated in the mobile market, it uses SIP to support services across virtually any access technology.

IMS has evolved so that multiple access technologies can be blended: Wi-Fi, Worldwide Interoperability for Microwave Access (WiMAX), DSL, broadband cable access, and even enterprise-level T1. The IMS architecture includes an agnostic control plane that can work over cable or over mobile networks, and serves to bridge packet, circuit, wired, and wireless worlds. A layered approach decouples the network infrastructure from services with a standardized, horizontal approach (compared to traditional vertical silo service approaches), and enables a flexible platform for integrated IP multimedia service delivery. The profitability appeal of IMS for service providers lies in its ability to provide a standard platform to respond rapidly to marketplace dynamics (i.e., new service requirements) and the need to better address service personalization (e.g., self-subscription, buddy lists, etc.) and control (e.g., QoS, class of service, charging, security, content filtering, etc.). Essentially, IMS is an application-centric concept appealing to all types of service providers.

Because of the momentum of SIP (already accepted by wireline and mobile carriers, Internet service providers, and many PC and CPE product vendors) and IMS, it is logical that MSOs and CableLabs would decide to adopt SIP and IMS as a foundation for next-generation networking services as specified in the evolving PacketCable 2.0 standard. Note that PacketCable Multimedia (PCMM) will be the mechanism for implementing QoS for SIP services as part of PacketCable 2.0. Outside of PacketCable 2.0, PCMM can also provide QoS for non-SIP or non-IMS multimedia applications such as gaming or IPTV. PacketCable 2.0 brings IMS and SIP together to help ensure a high-quality experience for all subscribers.

MIGRATING TO IMS/SIP MULTIMEDIA SERVICES

Protecting current investments and leveraging lessons already learned with PacketCable 1.x requires a gradual, evolutionary path to get to a PacketCable 2.0 infrastructure. Such a path can be realized by using current business needs to define and drive the addition of incremental capabilities to existing networks. Each incremental capability can bring a carrier one step closer to an IMS architecture.

This type of graceful transition toward IMS offers several advantages for cable operators:

- A transparent path for subscribers. Incremental capabilities can be introduced while protecting the subscriber experience and shielding them from underlying architectural changes.

- Reduced deployment times. For many operators, the introduction of PacketCable 1.x required the integration of the CMS, media gateway, and MTAs with their provisioning, billing, auditing, operations, and other systems. If these existing interfaces can be taken advantage of during the introduction of PacketCable 2.0 capabilities, deployment times can be reduced.

- Reduced operating costs. Packetcable 1.x also required operators to train operations staff in an entirely new set of skills. Leveraging that skillset can help reduce the costs of deploying PacketCable 2.0.

- Increased revenue. Incremental services can generate revenue during the migration to IMS.

The evolution to an IMS architecture will involve the incorporation of the three core IMS control functions (see Figure 1). The Serving Call Session Control Function (S-CSCF), also referred to as the home proxy or subscriber proxy function, manages access to the subscriber database and uses the information stored in that database to invoke features and applications in response to subscriber requests.

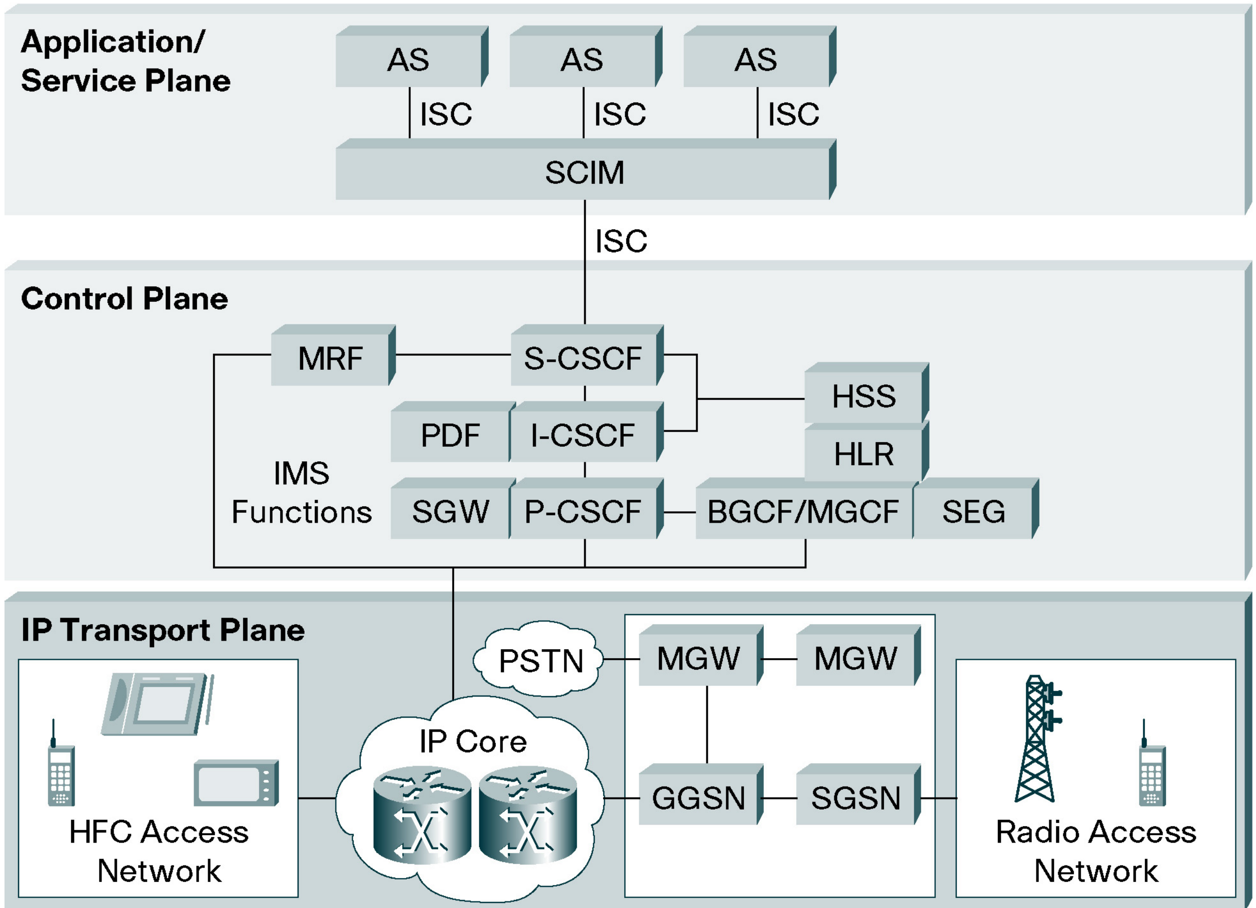

Figure 1. IMS for PocketCable 2.0

The second control function, the Interrogating Call Session Control Function (I-CSCF), controls the boundary to the network and is responsible for routing requests to the right S-CSCF. The third control function, the Proxy Call Session Control Function (P-CSCF), acts as an interface to clients , secures the link to the client, and facilitates roaming.

These IMS control functions sit on top of the IP transport plane. For GSM mobile communications, the IP layer includes the Gateway GPRS Support Node (GGSN) and Serving GPRS Support Node (SGSN) mobile access routers. On the cable side, this layer includes the HFC access network components.

In the application/service layer of the network, an IMS architecture introduces the IMS Service Control (ISC) interface for connecting the S-CSCF to a Service Capability Interaction Manager (SCIM). The SCIM performs feature interaction management, and connects to application servers using the same ISC interface.

The IMS is an important part of a service provider's network, providing a control plane for SIP-based services. However, IMS is only part of the story. As shown in Figure 2, a complete IP-based next-generation network (NGN) is composed of three distinct layers. These are the network layer, which includes the

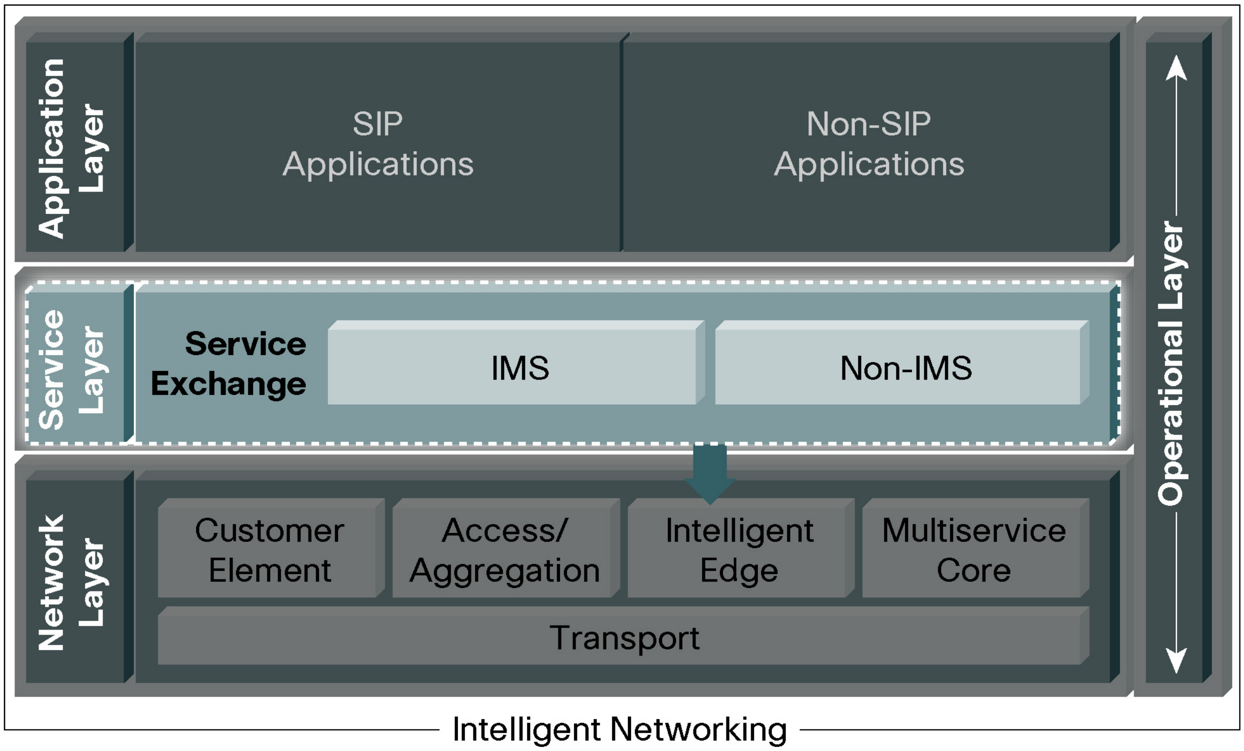

Figure 2. Next-generation network

entirety of the IP network, the control layer (of which IMS is a part), and the applications that reside on top. The control layer needs to provide a link between the applications and the underlying IP network for both SIP applications and for non-SIP applications. Indeed, the vast majority of IP-based applications today are not SIP-based, and a layer of service control is needed for these, too.

The importance of IMS for future services, and the importance of growing and protecting existing PacketCable 1.x services, calls for a phased introduction of IMS. Consider the following example as one path an operator might take.

Phase 1 – PSTN Bypass

Focusing on a reduced dependency on the PSTN addresses a relevant business need for today's MSOs. By making better use of Media Gateway Controllers (MGCs), MSOs can move more of their long-distance voice service off the PSTN and onto the IP network. This reduces the dependencies on competing carriers and reduces costs simultaneously.

To accomplish this first step, the soft switch is decomposed into two logical components – a subscriber-facing CMS and a PSTN-facing MGC. (See Figure 3.) By separating the soft switch into these component parts, the network can be more easily scaled for better overall network efficiencies.

Once PSTN and subscriber control functions are separated, MSOs can then introduce the first IMS element: a combined I-CSCF and Breakout Gateway Control Function (BGCF) function. (BGCF is the

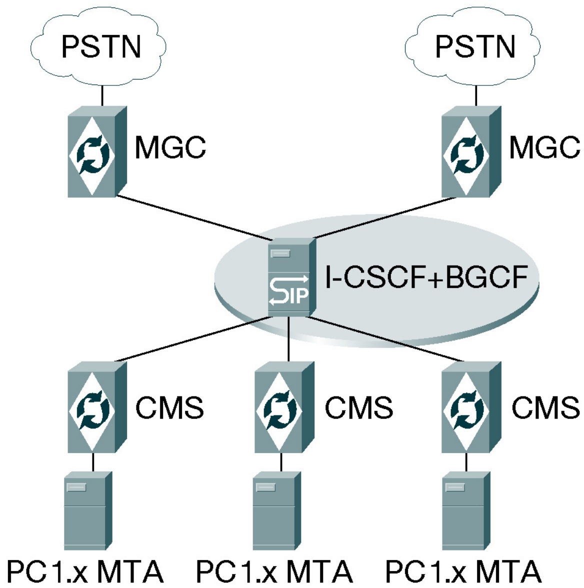

Figure 3. Inserting I-CSCF + BGCF (Bypassing the PSTN)

interface for interconnecting IMS with legacy networks.) This element serves multiple roles. First, it allows PSTN interconnects to be shared by multiple CMSs. CMSs can be added as needed, allowing the network to scale with increases in subscribers. Similarly, PSTN interconnects can be added as traffic requires. Calls between subscribers can stay on-net, routed to the correct terminating CMS by the I-CSCF function.

This configuration offloads calls from the PSTN. By keeping voice traffic on the IP core network, instead of requiring high-cost hand-offs to the PSTN, MSOs can not only reduce costs, but also increase call quality, reduce latencies, and allow the use of high-rate, wide-band codecs for improving speech quality. These features add up to improved differentiation and service quality.

Phase 2 – Add SIP-Based Services

The addition of the first IMS component in Phase 1 offers many benefits, but does not result in increased revenue. Phase 2 introduces new revenue streams for IMS-enabled services.

The introduction of SIP-based services maintains the previously mentioned focus on business needs. Specifically, MSOs need to introduce new services to provide feature differentiation from telcos and incumbent telephone providers that are competing in the same markets. Retaining customers can be more successfully accomplished by introducing SIP-based services that are not available from these competitors. Revenues are also increased from the monthly charges associated with the new services, which can include capabilities such as voice dialing, caller ID on a set-top box (STB), and click to dial.

With the changes already made in Phase 1, new SIP-based services can now be rapidly introduced and delivered by connecting new application servers to the CMSs. IMS introduces the 3GPP-specified ISC interface, a SIP-based interface for interfacing to application servers. Using these constructs, multiple application servers from multiple vendors can all interconnect to CMSs over the IMS ISC interface. Here, a true blending has occurred. The SIP-based ISC interface from IMS and PacketCable 2.0 is used, but to provide features to existing PacketCable 1.x endpoints. This allows a minimally disruptive transition with no forklift upgrades and no replacement of endpoints. Application servers can be added without changes in the CMS, allowing for faster rollout of services.

Phase 3 – Business/Commercial Voice

For the next phase of the transition to IMS, an MSO might focus on the business needs related to the expansion of the commercial

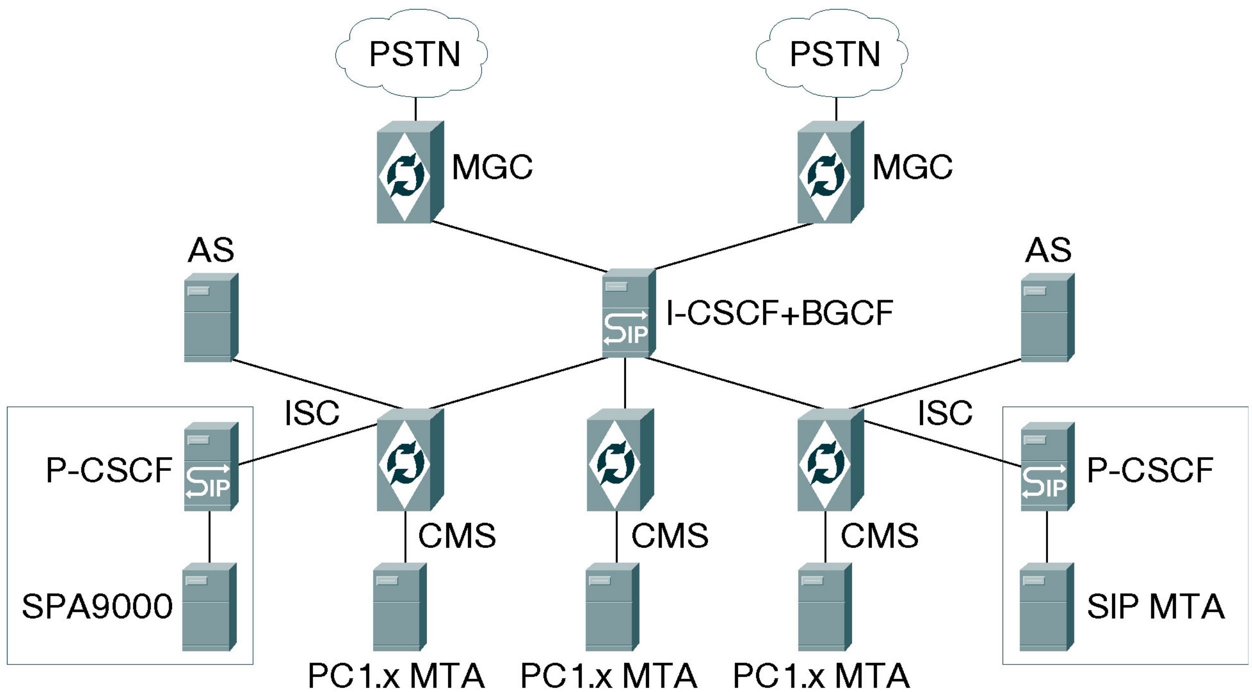

Figure 4. Insert P-CSCF (enable business voice)

subscriber base. These customers require a high-quality, business-grade experience and an expanded feature set with capabilities such as conference calling, call forwarding, and integrated voicemail and messaging.

To address this business goal, MSOs can build on the IMS environment by adding SIP endpoints (e.g., the Linksys SPA9000, a residential standalone box that provides a small, integrated IT PBX solution with support for up to 16 lines, or SIP MTAs).

To interface to these SIP endpoints, the operator introduces another IMS component during this phase: the P-CSCF (see Figure 4). This proxy control component provides PCMM QoS features that ensure business-grade voice services. The P-CSCF component will be connected to the underlying CMTSs and policy servers to provide QoS and security functions such as Network Address Translation (NAT) traversal and secure access – features that help ensure the protection of the new smarter endpoints.

As in the previous phase, this phase uses the in-place CMSs, allowing them to control the SIP endpoints. Through the ISC interface, the SIP endpoints also gain access to the new applications deployed from the previous step..

With the completion of Phase 3, an MSO has actually introduced a significant portion of an overall IMS architecture. The CMS has taken on much of the role of the S-CSCF, the P-CSCF was introduced to talk to SIP endpoints, and the I-CSCF was introduced for enhanced routing functionality. A big part of making this migration graceful was the blending of S-CSCF functions with the CMS.This phase also increases the revenue flow from new services.

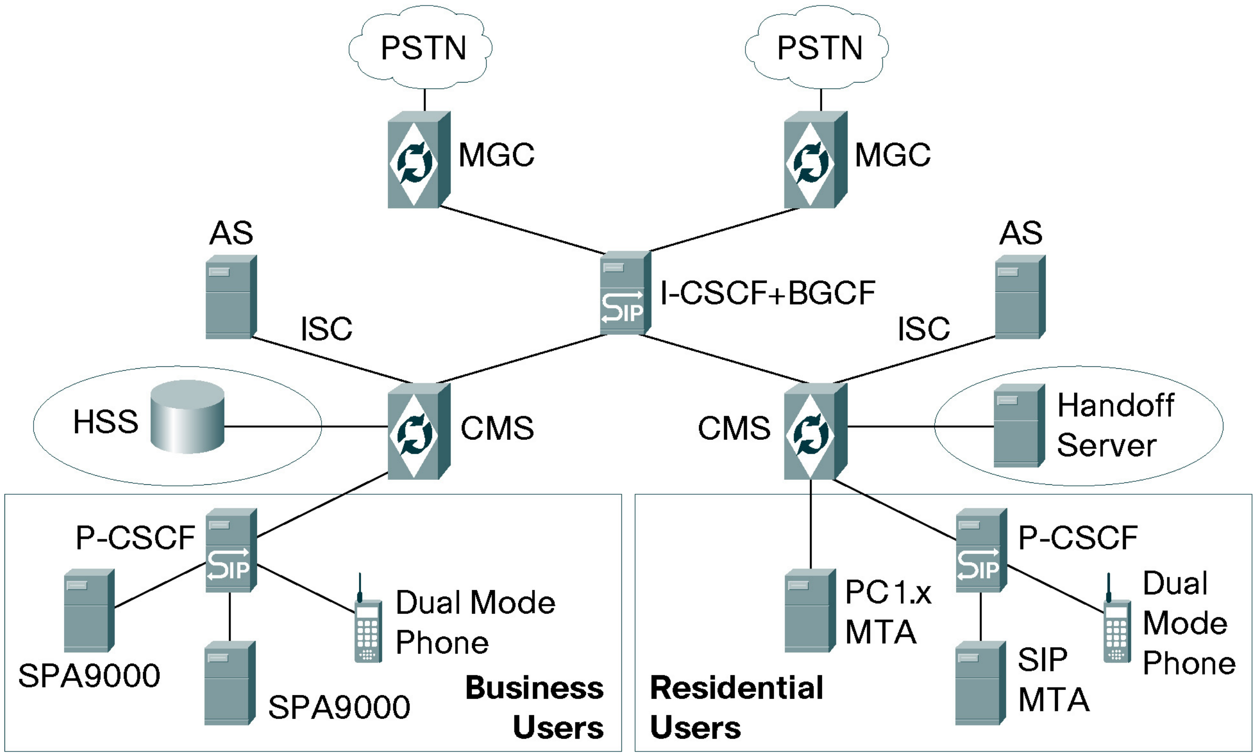

Figure 5. Add handoff server (FTC)

<u>Phase 4 – Fixed/Mobile Convergence</u>

Moving toward fixed/mobile convergence (FMC), an MSO can address several business needs relating to the introduction of "triple play on the move." MSOs can provide a superior user experience with better at-home coverage, begin to converge the fixed and wireless experience, and unify broadband VoIP and wireless with single-number access. Users gain access to new applications that can only be delivered over high-speed networks, and MSOs make progress toward the goal of all three services (data, voice, and video) on all three devices (TV, PCs, and handsets).

Accomplishing this phase involves the support for dual-mode handsets, and the introduction of two servers (see Figure 5). The dual-mode devices can communicate over the cellular network, or act as a new endpoint on the IP network. The Home Subscriber Server (HSS), the last missing piece of the IMS architecture, is introduced. It is needed to manage subscriber data uniformly between the cellular and IP worlds. The Handoff Server is also introduced in this phase. It runs on top of the ISC interface, and provides a seamless experience when subscribers move from the cellular network to a Wi-Fi network.. The CMS remains the functional center of the network, but with the introduction of the HSS, has added the Cx and Sh interfaces defined by the IMS, taking it a step further to becoming a complete S-CSCF. By continuing to take advantage of the CMS in each phase, MSOs accomplish a truly evolutionary move to IMS.

SUMMARY

By introducing a common suscriber data model and standardized interfaces for application servers, the IMS promise is to create a flexible platform for quickly launching new services. IMS also allows MSOs to take full advantage of their existing IP core network, and avoid overloading the cable infrastructure because very few applications need to extend beyond the IP network.

The IMS component of PacketCable 2.0 provides many opportunities for cost savings including bypassing the PSTN, more flexibly scaling the network, and more quickly integrating new application servers. The resulting new services can greatly enhance revenues while differentiating an MSO from competing telcos.

MSOs require a Service Exchange Framework (SEF) that supports a transparent migration to IMS, with full support for both IMS and non-IMS endpoints. This framework must allow service providers to deliver today's voice, video, and data services efficiently while also creating a foundation for new rich multimedia services. With a robust service exchange solution for IMS, cable MSOs will be able to support multiple applications on a common infrastructure that supports both SIP and PacketCable 1.x, while offering subscribers a customized service experience based on real-time state information and profile preferences.

The Service Exchange Framework eliminates dependence on a single application vendor, and allows cable providers to reuse solution components across multiple applications and even other access networks for full subscriber mobility. This framework will also enable cable operators to generate revenue by offering FMC services to their subscribers, thereby providing access over any network to a complete array of real-time, multimedia business and consumer services such as VoIP, video content sharing, presence-based services, and video telephony.

ABOUT THE AUTHOR

Dr. Jonathan Rosenberg is a Cisco Fellow in the Routing and Service Provider Technology Group at Cisco Systems. There, he is responsible for guiding the technology directions for their service provider VoIP products. Jonathan is active in the IETF; he is the lead author of the Session Initiation Protocol (SIP), the inventor of SIP for presence and IM, known as SIMPLE, in addition to other Internet technologies, such as STUN, XCAP and TRIP. Jonathan is also a member of the Internet Architecture Board. Jonathan received his PhD from Columbia University, and his Bachelors and Masters from MIT. Jonathan has been awarded the Voice on the Net Pioneer Award by pulver.com, and was named one of the one hundred most innovative young technologists in the world by Technology Review Magazine.

Other papers by Jonathan may be found at http://www.jdrosen.net. For more information about IMS and SIP, visit: http://www.cisco.com/en/US/netsol/ns537/networking_solutions_announcement0900aecd80381291.html. For more information about broadband cable, visit: http://www.cisco.com/en/US/products/hw/cable/index.html

QUALITY OF EXPERIENCE IN CABLE NETWORKS: CHALLENGES, TRENDS, AND SOLUTIONS

Alon Bernstein
Cisco Systems

Abstract

Managing quality of experience (QoE) in a multiservice network is one of the most challenging aspects of network design; even more so, in point to multipoint (P2MP) cable networks.

This paper explains how QoE relates to quality of service (QoS), outlines what challenges cable operators see in managing QoE today, identifies the challenges waiting in the near term, and relays how to address them.

QUALITY OF EXPERIENCE VERSUS QUALITY OF SERVICE

A service provider typically sells a service level agreement (SLA) in terms of QoS. Typical parameters are:

- Peak rate – the maximal rate a service can reach in bits per second.
- Committed rate – the minimal rate guaranteed to the service. Even if the network gets congested the user will get at least "committed rate" bits per second.

Parameters that are internally set, but not typically communicated to the user as part of an SLA are:

- Burst size – how many bytes can be sent at line rate. Note that for the duration of the burst, the peak rate is violated, but over a longer period of time, it's maintained.
- Priority – what level of preferential treatment the service gets over other services.

The SLAs above could provide an easy way to compare Internet Service Provider (ISP) offerings: The more megabits per second, the more desirable the service would be. But QoS is only part of the whole user experience.

The data rate of a service is a not a good enough metric of the "experience" for the user. For example, loading a 500 K web page at 100 Mbps (which takes about 40ms) and loading the same page at 1 Gbps (about 4ms) provides the same experience for a user (Note that a full rate video at 30 frames per second switches a image every 33ms – no point in loading pages faster than that).

Quality of experience is, in many cases, a subjective measure which makes it hard to compare one service offering to another. However, several guidelines can be defined per application. For example, voice has Perceptual Speech Quality Measurement (PSQM) and Mean Opinion Score (MOS) as tools that are used to measure quality of experience in an automated way instead of polling a group of listeners each time a system test is run.

Similar attempts have been made to quantify the experience with video as well; for example, Video Quality Metric (VQM) and the "voice quality group" effort.

Quality of experience may also impact quality of service. As an example, a higher codec compression rate might require lower QoS, but may also result in lowered user experience because of the degraded image quality.

Another aspect in which QoE is different than QoS is that it includes all aspects of providing a service. While QoS focuses on the data plane (how many bits per second a service can provide), QoE captures the user experience in its totality, including functions that are associated with the control plane such as the time it takes to change a channel when channel surfing, or the time it takes to establish a voice call.

For control plane QoS, there are "magic numbers" that define the service as acceptable. For example, a channel change time of less than half a second is considered sufficient.

CURRENT BANDWIDTH AND RESOURCE MANAGEMENT CHALLENGES

The following is a list that tracks current issues with bandwidth and resource management in cable networks:

(1) Oversubscription and Modeling: Cable operators count on oversubscription in order to reduce expenses. In other words, the number of users per downstream or upstream is such that if everyone became active at once, their QoE would be unacceptable. A typical number of subscribers per downstream today is 800. With a downstream rate of 38 Mbps, that would be 47.5 Kbps per user – less than an analog modem! However, there is a statistical assumption that not everyone is going to be active at once. Unfortunately, one size does not fit all: traffic patterns are unpredictable. They depend on such things as demographics, penetration levels and other factors that are not within the cable operator's control. All the above makes modeling and capacity planning a challenge.

(2) Admission control of video flows: Video flows can be long lived – 90 minutes or more. If service is blocked because all the bandwidth dedicated for video is used, then this blockage might persist for a while. This means that oversubscription estimates for video need to be more conservative than for voice since the admission patterns are less dynamic.

(3) End-to-End QoS: QoS is typically guaranteed only within the cable operator's network. No uniform end-to-end QoS infrastructure has been standardized. An end-to-end QoS is needed if MSO A routes a phone call (or other service) directly to MSO B without going through the PSTN first.

(4) Symmetry: Voice, file sharing and video conferencing, drive bandwidth allocation to symmetry. The way the cable plant is wired, however, is inherently asymmetrical with much more downstream than upstream bandwidth.

(5) Shortage of bandwidth: There is not enough bandwidth in general and, in particular, for the next killer application—video. Certain applications, such as peer-to-peer (P2P) file sharing constantly drive bandwidth utilization up.

(6) Statistical multiplexing: The relative size of an upstream (US) or downstream (DS) channel is small compared to the bandwidth consumed by the services offered; for example, 6 Mbps service on a 38 Mbps downstream, reducing statistical multiplexing gain.

(7) Simple billing: Service is typically a flat rate best effort. Best effort might not be the optimal scheduling mode for all applications.

(8) Viruses and DoS (denial of service) attacks are a constant risk since they can reduce usable network bandwidth. QoE can be impacted by viruses and DoS attacks that specifically target the control plane. A DoS attack that floods the network with bogus IGMP JOINs can slow channel change times.

(9) Common language for creating QoS enabled flows: Many applications that may benefit from a QoS service flow do not have the proper hooks in them to trigger a service flow creation automatically. For example, most gaming consoles do not have a standardized way of requesting "delay sensitive" service for their flows.

(10) Benchmarking: Users have tools such as DSL reports, to grade an MSO service. These are primarily bandwidth driven. These are QoS tools and do not reflect the QoE that an MSO has to offer.

(11) Stressing the control plane: QoS is a data plane issue. However, highly dynamic services, such as video channel change or large number of voice calls, strain the control planes as well and impact QoE.

(12) Stressing the data plane QoS: With the DOCSIS® per-flow queuing, a very large number of queues have to be managed for every serving group. This stretches the limits of proprietary and off-the-shelf network processors to their limit.

(13) Commercial services: Cable networks have traditionally targeted consumers. There is a strong drive for providing commercial services over the cable network, including T1/E1 emulation, which is not a natural fit for a packet network that operates in an RF environment.

TRENDS IN QOS AND QOE

The following list is an attempt to predict where the QoE/QoS trends are:

(1) Increased number of services: With PCMM, the range of applications supported by the cable plant will increase, and so will the unpredictability. Tiered services are likely to become more common.

(2) The networked home: Each household is going to have more and more end users and PCs. Each one may be using a different set of applications (for example: data, gaming, video) which means tiered services in the home.

(3) The DOCSIS upstream: Some services on the DOCSIS upstream (UGS/UGS-AD/RTPS) have specific jitter requirements. When the number of different services increases—each with its own

periodicity and jitter requirements—the algorithmic complexity of scheduling these services, and performing admission control on all these varied services, becomes unmanageable. It is mathematically known to be an "NP-complete" problem, meaning that its impossible to find an optimal solution in a finite time).

(4) Increased bandwidth per-user: Because of competition with DSL, and the drive for bandwidth hungry Video On Demand (VoD) applications, the bandwidth demand per-user is increasing. A single high-definition (HD) stream is at about 8 Mbps. Since a cable downstream is about 38 Mbps at the most, there is not much room to grow. Some service providers are already singing up subscribers for 20 Mbps and above. Clearly at these rates, there can not be much statistical multiplexing gains.

(5) Delay sensitive applications: As more delay-sensitive applications, such as voice, video conferencing and gaming, are widely deployed, the need for end-to-end QoS will increase. The need to properly prioritize and schedule all these different services that are all delay-sensitive will also increase.

(6) Billing: As more services are added, a flat rate best effort service would become only a baseline service.

(7) Positive trends in DoS: There seems to be a decline in successful DoS attacks and viruses—most likely as a result of users becoming smarter about protecting their machines, and Microsoft actively working to reduce the number of vulnerabilities that Windows OS has. Router companies, as well, have come up with automated systems to detect and defuse DoS attacked. DoS attacks, however, are still QoS/QoE risks.

(8) Higher compression: While higher compression reduces the need for bandwidth, it also increases the sensitivity to packet drops because each bit of information becomes critical as compression ratio goes up. This is not an issue for "over the top" services which are likely to use Transmission Control Protocol (TCP), but will be an issue for a more optimized real time protocol (RTP) which currently does not have re-transmission capabilities (an RFC draft only).

(9) Multicast: Multicast greatly improves bandwidth utilization when the same content is viewed by a large number of users at the same time. Multicast services are going to be more common all the way to the Customer Premise Equipment (CPE).

(10) "Over the top" services: Content and services can be provided by outside companies. The same way that voice is delivered by Sktype and Vonage, video distribution would follow the same path. Providers external to the MSO network will provide video services. For these services, end-to-end QoS may also become critical.

(11) P2P file sharing: P2P file sharing networks can be viewed as systems for delivering video over the Internet, since video is (in terms of traffic volume) what is driving the

P2P usage. In other words, even if a service provider manages to clamp down on P2P traffic by some means, it does not mean that the bandwidth demands would decrease– at least on the downstream direction. They will be replaced by legitimate bandwidth hogs in the form of VoD streams.

(12) More subscribers, more bandwidth: For each household passed, each subscriber is going to have more bandwidth. If one assumes that penetration rates are becoming saturated that would mean that the number of users will remain fixed. But this is not the case. DOCSIS Set-top boxes (STBs), as well as other DOCSIS-enabled devices (power meters for example), will increase the number of devices that the network has to support, and with it the constraints on QoS.

SOLUTIONS

Before outlining a set of solutions, a basic question has to be answered: Why does QoE need to be managed ? Can't all QoE problems be solved by providing enough resources to the network? After all, QoS is not an issue if there is enough bandwidth in the network and QoE should not be as well.

There are several answers to this question. First of all, a network can be built more economically if it is not built on worst case assumptions regarding bandwidth utilization. For example, we can assume that each household needs 50 Mbps to allow for 3 HD streams and Internet services. For 500 households passed, that would be 25 Gbps – exceeding the capacity of a single fiber node – and more than what is needed in practice.

Even for a network that is over-provisioned, having QoE/QoS enforcement is needed for mission critical services (such as 911 calls) so that even if the statistical model fails on extreme circumstances, the critical services are not impacted.

Because of the shared media nature of cable networks, a service might degrade as more users are added to the same serving group. In the past, customers have preferred a QoE where bandwidth availably is restricted, but constant, over one where bandwidth availability fluctuates between high and low. Actively managing bandwidth achieves the target of having predicable and stable bandwidth.

The concept of "net neutrality" is hotly debated these days. It's a question on what types of preferences an MSO can give its own traffic versus externally sourced traffic. There seems to be a general agreement that it is not okay to intentionally degrade external traffic, and that it is acceptable to provide a higher level of service to internally sourced traffic.

However, even the latter is debated since an MSO can end up leaving very little bandwidth for external services – enough for web browsing, but no more than that. It's still not clear what type of regulation, if any, would be enforced in these cases. It's likely that QoE tools will be needed to manage it, if indeed it is enforced.

The following list outlines solutions to the challenges and trends presented in earlier sections. Note that not all items can be addressed. This list covers the ones that can:

(1) DOCSIS 3.0 downstream channel bonding: CableLabs® is close to releasing a first version of the

DOCSIS 3.0 specifications. DOCSIS 3.0 increases the bandwidth of a DOCSIS channel by means of channel bonding. This means it does not increase the physical capacity of the channel. Instead, it uses a Multi Link PPP (MLPPP) like technique to spread packet across independent L2 links. The overall effect from an L3 perspective is a faster link.

DOCSIS 3.0 comes with its own set of QoE issues. Because packets are sequenced and sent on independent links, they might be delayed until all of them are received in sequence. This issue is exacerbated when a packet is lost. In certain cases, the re-sequencing engine will have to wait a full re-sequencing window (up to 18 ms) before resuming operations.

Another complexity in DOCSIS 3.0 comes from the fact that not all devices and not all services use the same number of channels. This creates a system of multiple, partially overlapping, groups on which a packet can be stripped. Scheduling and load balancing across these groups is a complex task. However, if executed correctly, it will provide both higher bandwidth per modem, as well as statistical multiplexing gains.

Bonding provides a new set of knobs that can be fine tuned to provide different QoE levels: the number of channel a flow is sent across, the timeout for re-sequencing, including/omitting sequence numbers, etc.

DOCSIS 3.0 upstream channel bonding will also address some of the issues with symmetry. Although the physical bandwidth of the plant is still heavily skewed towards downstream bandwidth, upstream bonding gives the ability of sending up to a 100 Mbps on the upstream bandwidth, so selected cable modems can have symmetry.

(2) DOCSIS 3.0 multicast: Another improvement in DOCSIS 3.0 is enhanced multicast support. This will further enhance the bandwidth savings that multicast enables. It's important to note that even though the trend is to have a stream per subscriber (VoD), multicast still has a role to play for real-time viewing events, and for off-line downloading of popular content to PVRs.

(3) DOCSIS 3.0 commercial services: DOCSIS 3.0 also addresses the issue of transporting T1/E1 services over DOCSIS, which require specific clocking to maintain the appropriate QoE.

(4) Applying PCMM to non-PCMM applications: The same deep-packet-inspection tools that are typically used to filter unwanted traffic can be used to detect certain types of flows: for example, gaming, and request a specific QoE level on their behalf so that even non-PCMM enabled applications can have dedicated QoE services.

(5) Dynamic bandwidth management: Since QoE is not about bits per second, it's possible to tailor an SLA that fits a user profile instead of selling a limited set of SLAs. A user

that is doing only web browsing does not need a high SLA. One option an MSO has is to sell packages such as "web browsing" and "video" instead of selling a number of mega bytes per second. Another option is to sense what applications a user is running and dynamically assign a profile that fits the users traffic patterns. This will help in modeling and optimizing bandwidth utilization.

(6) End-to-end bandwidth reservations: The issue of end-to-end bandwidth reservation is beginning to be addressed in various forums, but no standard is emerging. There are many issues with customer ownership, responsibly in case of failures and communication of the service levels (not all ISPs mark "high priority" the same way) that are not fully resolved yet. Of special note is the IP sphere effort, which tries to dynamically set business agreements, on a call-by-call basis, between service providers.

(7) Improved modeling: Several tools for advanced modeling are becoming available from CableLabs® and other sources. In addition to that, a good approach to planning network capacity is to start an iterative process: make some assumptions about network utilization and the number of subscribers that can use the network. Raise the right flags as the system begins to cross certain thresholds (for example, aggregate data rate, or CPU utilization or control plane load) so that the MSO has enough time to properly update the network before a QoE disruption occurs.

(8) Flexible jitter bounds for the upstream: A way to address the upstream scheduling problem is to provide a statistical guarantee on jitter/delay, rather than an absolute guarantee. This approach would make the upstream look more like a packet transport (same as the downstream) where the jitter/delay is a function of utilization levels and not guaranteed in any way. Using admission control tools, it is possible to keep the jitter/delay at a certain statistical guarantee.

CONCLUSION

The vision of a converged video/voice/data network has been around for years. We have reached a point where this vision is becoming a reality. The competition between DSL and cable has driven cable to provide voice services, and DSL to providing video services. While cable operators still manage video as a separate network, they will soon provide video services over their data pipes—either from externally sourced servers, or as an upgrade to their current video offerings (IP video over DOCSIS can provide better bandwidth usage, enhanced interactivity, better network connectivity, etc). Other services will be added: gaming, backup services, video conferencing and others. Managing the QoE for these services will be critical for future successful deployments.

SWITCHED UNICAST: IT'S NOT JUST ABOUT CAPACITY

Ran Oz,
Founder and Chief Technology Officer
BigBand Networks

Abstract

Switched unicast is a form of switched broadcast in which each subscriber receives a unique program stream. The benefits of switched unicast extend beyond reclaiming network capacity. This paper discusses the benefits of a switched unicast implementation, exploring issues such as addressable advertising, content personalization and transformation, along with other opportunities for overall enhancement of the subscriber's viewing experiences. Additionally, cable operators can get precise viewership statistics without relying on TV audience rating firms, since intelligent switched unicast systems monitor and respond to all subscriber activity, and can store data about which programs subscribers are watching.

The paper also examines the role of switched unicast in enabling more open conditional access systems, expanding choices of set-top boxes and other subscriber devices, and driving innovations such as edge-based rate shaping and adoption of universal edge QAMs that further enhance bandwidth efficiencies.

REVIEW OF SWITCHED BROADCAST

Switched broadcast came of age in 2005 with major cable operators beginning deploy the technology in several locations. Switched broadcast leverages the field-proven concept that the actual viewership of programming at a specific time within a local area is a fraction of the total number of offered programs.

Switched multicast delivers programs only when and where requested by viewers, unlike legacy broadcast systems that deliver all programming to all subscribers. If a subscriber wants to watch a program that is currently being delivered to other subscribers within the same node, or service group, the subscriber simply joins the existing switched multicast session. As a result, there is no capacity consumed in delivering the program to the additional subscriber. This methodology optimizes bandwidth and resource efficiency in enabling the creation of virtual programming capacity and bandwidth without the correlated physical expense of creating and dedicating precious bandwidth and associated QAM resources.

Bandwidth reclamation is an imperative for cable operators wanting to expand services and enhance their competitiveness. Switched broadcast enables cable companies to offer services spanning HDTV, VOD IPTV and higher speed broadband, without needing to upgrade HFC capacity.

Additionally, by moving some, or all, of the broadcast digital tier to the narrowcast spectrum, cable operators are able to significantly increase the amount of programming offered. The so-called "long tail effect" demonstrates that there is a large aggregate demand when many specialized offerings are made available.

Adoption of a switched services model, therefore, enables operators to both increase revenue opportunities by successfully motivating subscribers to purchase premium packages, while mitigating competitive threats from satellite broadcasters and telcos, and media companies using the Internet to sell video directly to consumers.

The key building blocks of a switched multicast network are shown in figure 1. These are:

- Clamping sub-system;
- Edge sub-system;
- Management server;
- Client software.

The clamping sub-system converts incoming broadcast programs to a predetermined constant bit rate value. Though not absolutely required, the normalization of all programs to a constant bit rate allows a faster and simpler switching mechanism, and paves the way for QAM resource sharing with similar services such as VOD. The bit rate most often chosen is one that resembles the program parameters for VOD streams. The clamped programs are multicast to all hubs using a Gigabit Ethernet transport network.

The edge sub-system replicates program streams and directs them to the appropriate edge QAM, in response to directions from the switched broadcast manager. Edge switching capabilities can be integrated into the QAMs or can be a stand-alone entity upstream of the QAMs.

Client software resides on a subscriber's set-top box. When a program is selected, the client conveys the channel request upstream along with information that uniquely identifies the node-group location of the set-top box. The client functionality can be easily integrated into the tuning firmware of future set-top boxes.

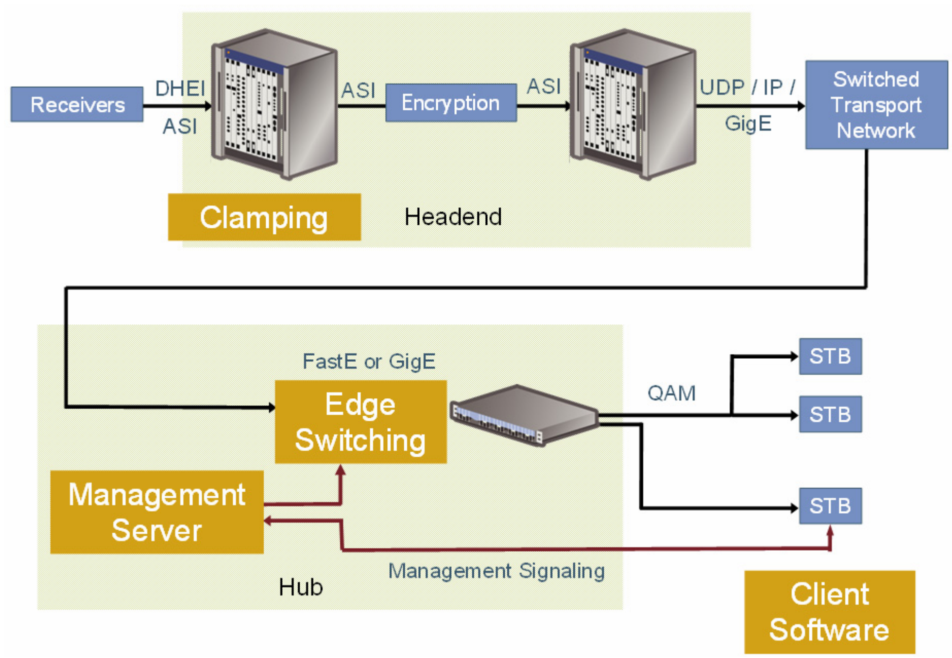

Figure 1. Components of switched broadcast systems

The management server uses the received channel number to identify the requested program and, consequently, the input port on the switch where the program is being received. Similarly, the switched broadcast manager uses the set-top box ID and associates the node group information to determine the downstream path that connects to the subscriber.

When an available downstream QAM and program resource are identified, the frequency and program information is returned via the downstream out-of-band channel to the set-top box, which decodes and displays the program using the normal tuning mechanisms. While the specific frequency and program number may vary with time, the channel number seen by the subscriber's set-top box is unchanged.

BANDWIDTH SAVINGS WITH SWITCHED UNICAST

Switched unicast is a form of switched broadcast in which each subscriber receives a unique program stream. The technology employs a similar tuning process to switched multicast except that a subscriber is allocated new dedicated bandwidth whenever switched broadcast programming is requested. Switched unicast consumes significantly less bandwidth than traditional broadcast methods while allowing operators to offer highly customized programming. This is demonstrated by the two examples that follow.

Figure 2 shows the results of a recent field trial.

In this example 172 programs were available to viewers, but the maximum number ever watched a one time was 70. This demonstrates the potential for bandwidth savings in a switched multicast model. The bandwidth benefits of switched unicast stem from the fact that the maximum number of tuners active at any one time was 110, so that even when each tuner received a dedicated stream, the total amount of bandwidth consumed was less than with in a legacy broadcast environment.

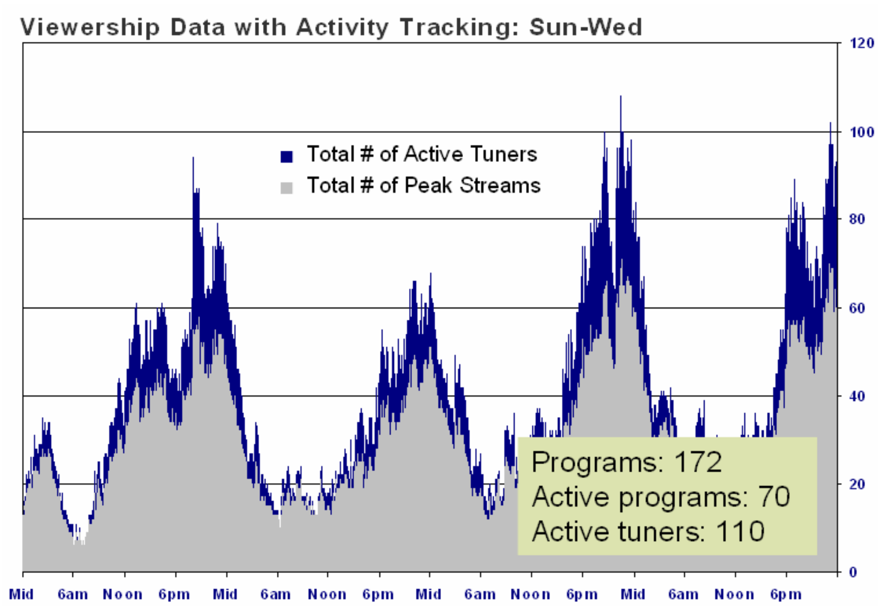

Figure 2. Bandwidth savings with switched unicast (Deployment A)

Another example of the bandwidth savings that switched unicast offers is provided in Figure 3. Just as in the previous example, the data presented here come from an actual field trial. In this example there were 79 programs available to viewers but a maximum of only 23 watched at any time, and the number of active tuners peaked at 49.

In both trials switched unicast significantly reduced the amount of bandwidth required to deliver the programming that subscribers wanted. The ratio of dedicated program streams to program choices peaked at 0.64 in deployment A and 0.62 in deployment B. Additional field trials currently underway will yield more insights into the bandwidth savings that switched unicast enables. Switched unicast, however, offers cable operators much more than merely bandwidth savings. Several compelling benefits of the technology are described in the following sections.

In both trials switched unicast significantly reduced the amount of bandwidth required to deliver the programming that subscribers wanted. The ratio of dedicated program streams to program choices peaked at 0.64 in deployment A and 0.62 in deployment B.

It's worth noting that in neither field trial was the switched tier thoroughly analyzed to determine which programs should be switched and the optimum size of service groups. Consequently, the bandwidth savings potential of switched unicast is even greater than that shown in both figures. Ongoing field trials will yield more insights into the bandwidth savings that switched unicast enables.

Switched unicast, however, offers cable operators much more than merely bandwidth savings. Several compelling benefits of the technology are described in the following sections.

SWITCHED UNICAST: IT'S NOT JUST ABOUT CAPACITY

In addition to bandwidth reclamation, switched unicast offers a variety of benefits. These include:

- Content personalization;

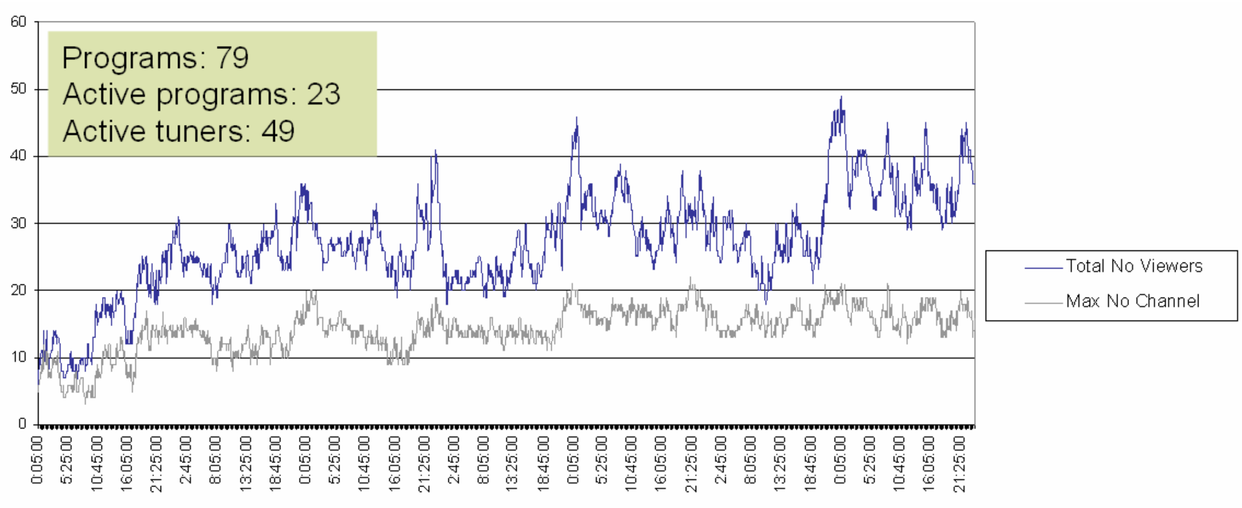

Figure 3. Bandwidth savings with switched unicast (Deployment B)

- Targeted addressable advertising;
- Precise viewership data;
- Expanded operator choices of CAS;
- Wider consumer choices for CPE.

Each of these is discussed in the sections that follow, along with an explanation of how innovations at the edge of the cable network are broadening the capabilities of switched unicast systems.

CONTENT PERSONALIZATION

The cable industry continues to look beyond traditional broadcast TV models, and is examining initiatives that will allow subscribers a higher degree of content customization. Switched unicast offers unique opportunities to enhance subscribers' viewing experiences by allowing content to become increasingly personalized. By enhancing their customers' viewing experiences, cable companies will be better equipped to fend off competitive threats from emerging Web-based video services.

Personalized news is one example of how programming can be chosen to reflect the interests of individual subscribers. Switched unicast systems enable content and subscriber interests to be correlated, allowing news summaries to be created that are more likely to retain the attention of viewers than traditional broadcast TV news programs. Enterprising newsrooms could record a series of short news stories on a wide range of topics, allowing cable operators to combine into personalized bulletins that address a subscriber's specific interests, whether it's baseball teams or Internet start-ups. A personalized version of a music network is another example of how switched unicast provides cable operators the opportunity to offer increasingly customized content.

This type of personalized service could be offered as a premium service, potentially earning the operator additional revenues. Alternatively, personalized content could be offered at no additional charge, with the expectation that increased customer loyalty results.

TARGETED ADDRESSABLE ADVERTISING

Switched unicast provides cable operators with the ability to increase advertising revenues by delivering promotional messages that more closely match subscribers' interests: marketers are willing to pay a premium rate if their ads achieve improved response rates among their intended audiences. In fact, several recent research studies indicate that advertisers have a strong willingness to pay twice as much for an advertisement if there is a guarantee that it reaches a targeted audience.

Figure 4 illustrates how three subscribers, all watching the same program on the switched tier receive different ads during the commercial breaks. For example, subscriber #1, an avid teen snowboarder, receives an ad about snowboard sales. During the same commercial break, subscriber #2, a thirties-something bachelor, views an ad about an upcoming motor show. Subscriber #3, an avid traveler in her fifties, receives information about cruises in the Caribbean.

Switched unicast systems remember subscribers' viewing patterns and can use this information when determining which ads to forward. For example, a subscriber hopping to a soap opera after a significant amount of time watching a sports network can receive an ad about an upcoming local sporting event.

Switched unicast also enable subscribers to undertake forms of ad telescoping activity. Telescoping enables a viewer to press a button on the remote control for more information or request a longer form of the ad or a brochure. Interested viewers can schedule an appointment with a retailer, receive a demo in the mail, or take part in a promotional contest.

Targeted advertising can be further enriched using nPVRs (networked PVRs). These enable subscribers to pause near-live broadcasts while telescoping. Once the ad has been viewed, the viewer can return to programming, at the place where it was paused.

By storing information about which programs individual subscribers are watching, switched unicast technology enables cable operators to collect precise viewership data. This not only yields insights into which programs are most popular, it also allows cable operators build profiles about their subscribers' interests. This helps cable operators determine which promotional messages a subscriber is more likely to be interested in. For example, a subscriber that routinely tunes to a home improvement network may be interested in an ad about a sale at a local hardware store even though they may be currently watching something else. the home improvement network.

The ability to build profiles about subscribers' interests adds tremendous value to targeted advertisements. Accordingly, the more enterprising cable operators could explicitly ask their subscribers about their ad preferences. In return for providing a cable operator with a list of the subjects and categories they'd be interested in viewing ads on, subscribers could receive a complimentary gift or upgrade to an expanded service package, or some other incentive.

VIEWERSHIP STATISTICS

Switched digital broadcast systems can be configured to store information about subscribers' viewing patterns. This allows cable operators to get precise viewership statistics without relying on third parties such as TV audience research firms. The

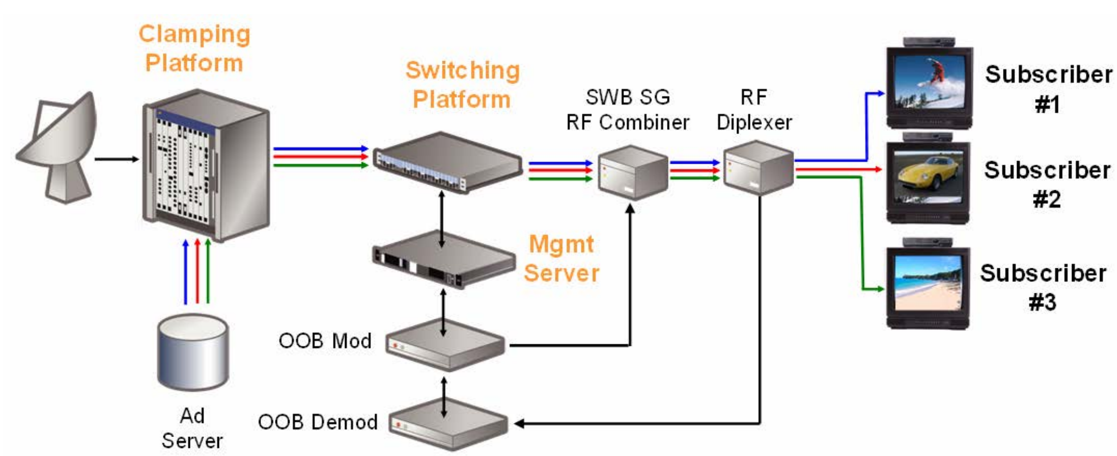

Figure 4: Subscribers receiving different ads though all view the same programming

value of this information is high because it provides insights into the viewing patterns of all subscribers on the switched tier, not just the subset of viewers that have been enlisted by audience rating firms, whose viewing habits may not necessarily represent those of the majority. Moreover, unlike viewers who track their viewing habits using diaries, switched unicast systems provide precise records without human biases.

Naturally, privacy restrictions surround the gathering of individual subscriber data, and the gathering of individual subscriber data could be implemented on an "opt-in" basis, with subscribers receiving a complimentary gift or some other incentive for participating.

EXPANDING CHOICES FOR CONDITIONAL ACCESS AND CPE

At present cable operators are limited to using proprietary system-based encryption methodologies. Switched unicast introduces the opportunity for a wider range of algorithms to be used when encrypting content. It achieves this by creating more than one version of each program stream and routing each one through parallel encryption servers. Client software on each subscriber's set-top box recognizes which encryption protocol it is compatible with, and requests that version from the switched broadcast manager. This concept is represented by Figure 5.

This liberates cable operators to select conditional access systems from a wider range of vendors, not just a few. It can also accelerate the transition from cable operators leasing and depreciating CPE (customer premise equipment) to subscribers buying their own devices. In turn, subscribers benefit from being able to choose from a wider variety, and price range, of set-top boxes, digital-ready TVs, PCs and other IP devices.

Switched unicast systems can serve as gateways to wider CPE choices by leveraging the OpenCable application platform, a middleware software layer specification intended to enable the developers of interactive television services and applications to design products that run successfully, independent of set-top or television receiver hardware or operating system software choices. The advanced encryption standard assists in the adoption

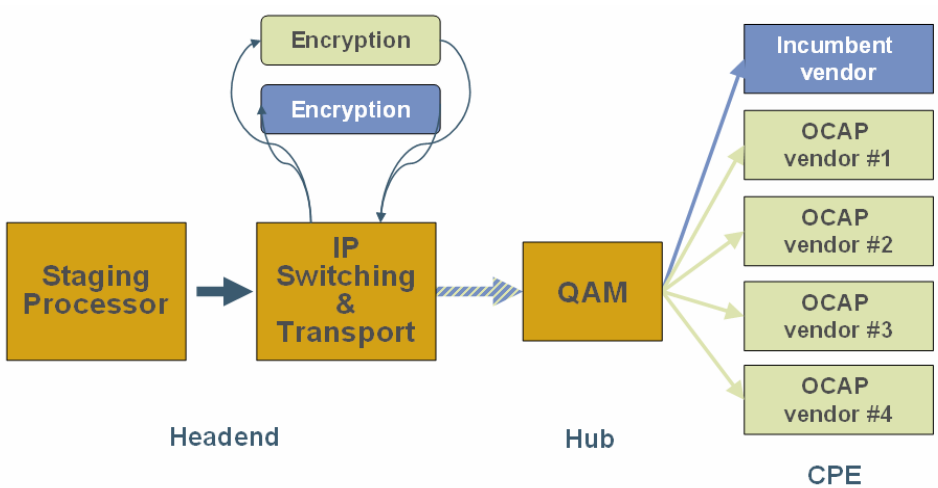

Figure 5: OCAP STB deployment with AES encryption

of open systems by defining a block cipher that allows encryption and decryption of programs irrespective of equipment vendor.

INSTANT CHANNEL CHANGE

The ability for subscribers to change programs without any perceived delay is a requirement of any type of network. It becomes more challenging, however, when each subscriber receives a dedicated program stream containing a single channel. In a switched unicast network each subscriber's set-top box is "parked" on a particular frequency and PID (packet identifier). When a subscriber requests a program, the unicast systems delivers the program to the subscriber's STB on the frequency that his set-top box is already tuned to. This helps eliminate latency from retuning.

An additional speed benefit is derived from the fact that the switched unicast system immediately sends the subscriber the most recent i-frame for the new program, avoiding the tuning delay that would occur if the subscriber's set-top box had to join a program stream already being broadcast and wait for next i-frame to be transmitted.

EDGE INNOVATION FOR BANDWIDTH BENEFITS

While switched broadcast provides cable operators the ability to reclaim large amounts of bandwidth, a combination of expanding subscribers demand for greater personalization, the growth of high definition TV and the need for high data speeds, will drive the need for continous evolution of network infrastructures.

Aggregation of services at the edge of the network demands offers the opportunity for innovations that results in versatility and improved bandwidth utilization. Edge rate shaping is the first of these innovations.

Rate shaping at the network edge begins with information collected at headend. Intelligent video platforms analyses the complexity of each program and creates metadata that characterizes each program. In addition, each progam is clamped at a variety of different bit rates. Metadata and clamped program streams are forwarded to hub locations.

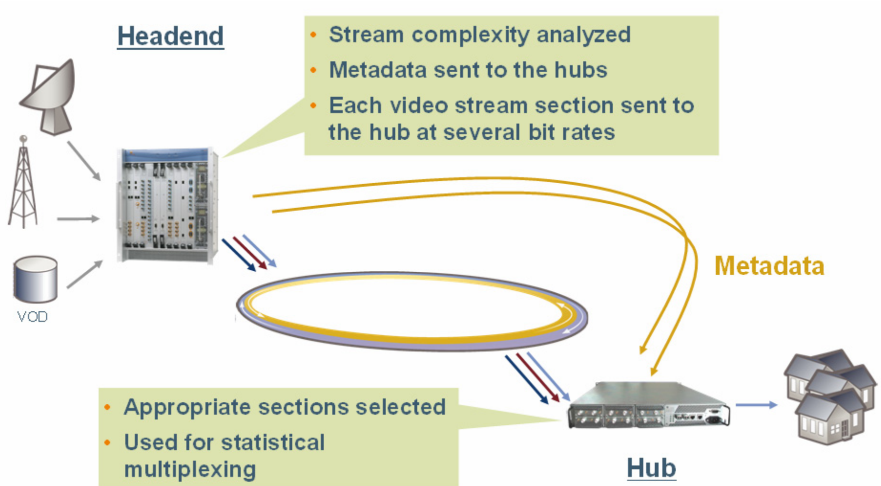

Figure 6: Rate shaping of program streams at edge locations

At each hub, edge video platforms dynamically examine the metadata and bandwidth availability on the network, and select the version of the program stream that both maximizes video quality and fits within available capacity. Edge rate shaping is shown in figure 6.

Universal edge QAMs evolve network infrastructures and lead to significant cable operator benefits. UEQs (universal edge QAMs) are enabled by the adoption of modular CMTS (cable modem termination system) architectures. An M-CMTS has disaggregated functionality for best-of-breed performance, resulting in improved economics, and serving as a catalyst for both capacity increases and more flexible allocations of upstream and downstream bandwidth.

The edge QAMs used in M-CMTS platforms have the potential to be the same ones used in switched multicast and unicast systems. This introduces the possibility of eliminating "nailed-up" allocations of QAMs across data, VOD and switched broadcast services, replacing them with QAM capacity that is shared across services. Versatile edge QAMs enable dynamic sharing of edge resources, in response to time-varying subscriber demands.

Figure 7 shows the UEQs can be dynamically allocated to support data, VOD and switched video services.

Universal edge QAMs are a vital milestone in the migration of cable infrastuctures from fixed silos that are relatively bandwidth inefficient, to networks that are more versatile and more fully utilized.

CONCLUSIONS

Switched unicast is the evolution of switched broadcast systems that begun being deployed in 2005. Switched broadcast provides cable operators the ability to expand offerings to include HDTV, wider selections of on-demand content, increasing amounts of peer-to-peer file sharing, and other services, by dramatically reclaiming network spectrum.

Switched unicast, however, is about much more than bandwidth savings. The benefits of switched unicast include greater content personalization and targeted addressable

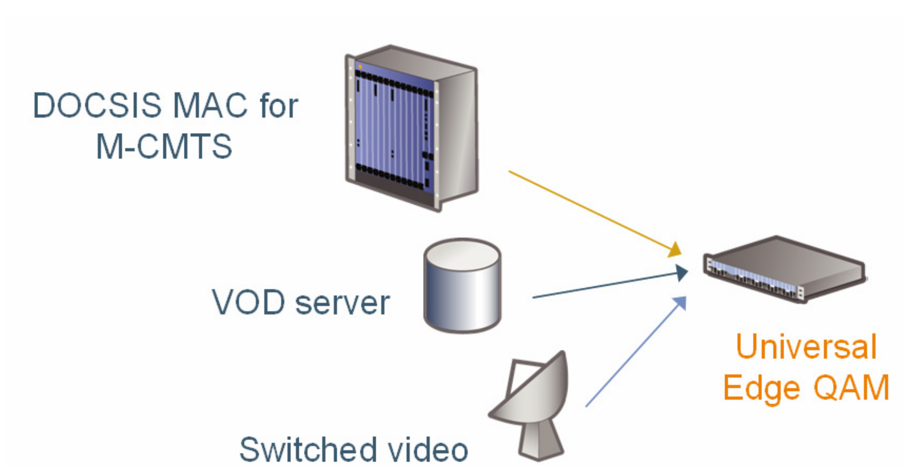

Figure 7: Versatile edge QAMs dynamically share data, on-demand and switched video

advertising opportunities, along with the ability for cable operators to obtain precise viewership statistics directly. These capabilities ultimately improve the viewing experiences of each and every viewer.

Additional operator benefits that can be derived from switched unicast include the freedom to select conditional access systems from an expanded choice of vendors, while subscribers benefit from having wider selections of set-top boxes and other CPE.

The cable engineer deploying switched broadcast primarily to alleviate capacity limitations will find that once a switched network has been implemented entire new areas of service benefits appear, as described above, with accruing competitive benefits.

SWITCHING TO OVERDRIVE ON THE INFORMATION SUPER-HIGHWAY

Oleh J. Sniezko
Chief Technical Officer, Aurora Networks, Inc.
Ray S. Thomas
Principal Engineer, Network Infrastructure, Time Warner Cable

Abstract

The authors presented analysis and processing results of the vast data collected during characterization of the broadband optical links used for expanding broadband bandwidth to 1 GHz and for supporting full digital load.

Three different types of load for optical links were analyzed:

1. *A combination of the CW load with test QAM channels,*
2. *A combination of the NTSC modulated analog channel load with test QAM channels, and*
3. *Full digital load of QAM channels.*

The results indicate that the dominant types of distortion to digital 256-QAM signals are nonlinear products generated in optical links. These distortions affect the digital signals before laser clipping occurs. Alignment of the optical laser transmitters is critical to realize full capacity of the optical links up to 1 GHz. Similarly, the parameters of the QAM signal FEC, and especially the interleaver settings, are important to achieve this goal. The paper presents the conclusions and lists some pointers to apply in laser transmitter alignment.

INTRODUCTION

Drive for Bandwidth

The Cable TV industry is continuously striving to increase the maximum rate of data and video throughput capability of their networks. These efforts focus on absolute bandwidth expansion and/or on expansion of bandwidth per customer. In the first category are such efforts as:

1. Moving system capacity towards 1 GHz FDM limit with a combination of traditional analog video and digital QAM channels, and
2. Using spectrum above the existing nominal design limits of the broadband subsystem.

In the second category are efforts to:

1. Segment nodes into smaller serving areas,
2. Replace analog channels with digital channels (also known as analog bandwidth reclamation),
3. Increase QAM modulation levels for digital signals,
4. Increase coding capacity for digital video signals (introduction more efficient encoding and digital compression algorithms),
5. Reclaim digital bandwidth with switched digital architecture,
6. Increase stat-muxing efficiency of digital video signals.

Critical Link

The optical links between headends/hubs and nodes are a critical part of these efforts. Their quality determines the absolute usable bandwidth and the levels of modulation for the QAM signals (bit/symbol capacity). Therefore, these links were tested and the test results were analyzed to determine their capability to deliver increased bandwidth to broadband customers.

Clipping or Nonlinear Distortions

There have been several mis-understandings about what is causing impairments to digital signals of higher modulation orders (256 QAM and higher). Several publications [1, 2, 3] presented in-depth analysis on this issue in the past. Notwithstanding those findings, the popular believe persists that it is clipping that causes all the problems with QAM signals.

We tested this theory against the theory that clipping is not the main contributor to the problem of digital signal impairments and that different distortions affect the digital signals before clipping really contributes to BER.

Figure 1 represents the most likely explanation of the problem. The analog lasers are selected for their transfer function linearity and they are further linearized to achieve high level of performance required for analog links. However, for very practical reason, those linearization efforts are limited to an operational range of analog links. For 1310 DFB lasers, this practical limit is around 35% of peak OMI values. Beyond this range, linearization is not practical. Henceforth, the linearization efforts stop there. Even within these limits, the linearization techniques allow for higher digital modulation levels if they are applied with understanding of the problem. The test results allow for clear problem definitions and for focusing the linearization efforts on optimization of the optical links for all the signals (analog video and digital QAM) transmitted over analog optical links as opposed to the efforts of achieving the best results within the bandwidth occupied by analog signals. This assertion has been tested during the data gathering process and confirmed in several ways.

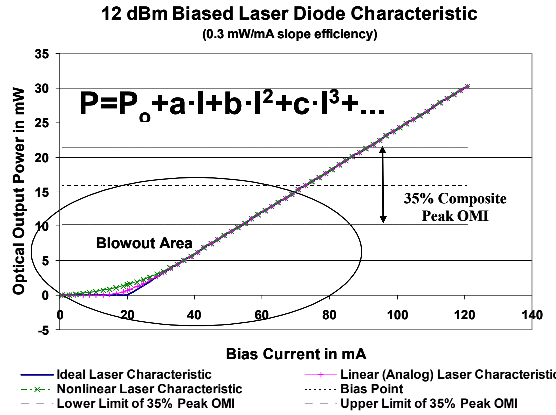

a) Laser Diode Characteristic

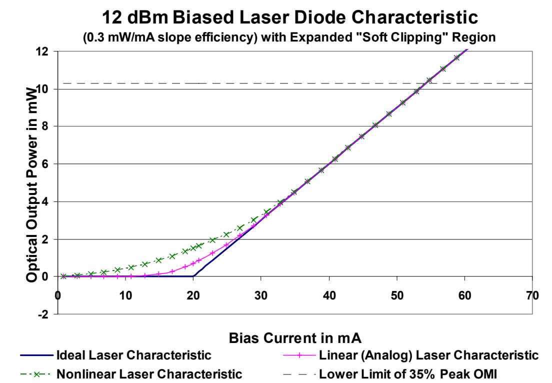

b) Blowout of the Threshold Area

Figure 1. Laser Diode Characteristics – Linear Area for Analog Lasers

TESTING METHODOLOGY AND TEST PROGRAM

Test Setup

Three different types of load for optical links were analyzed:

1. 72 CW carriers between 120 and 552 MHz,
2. 72 NTSC modulated analog channels between 120 and 552 MHz, and
3. 72 digital QAM channels between 120 and 552.

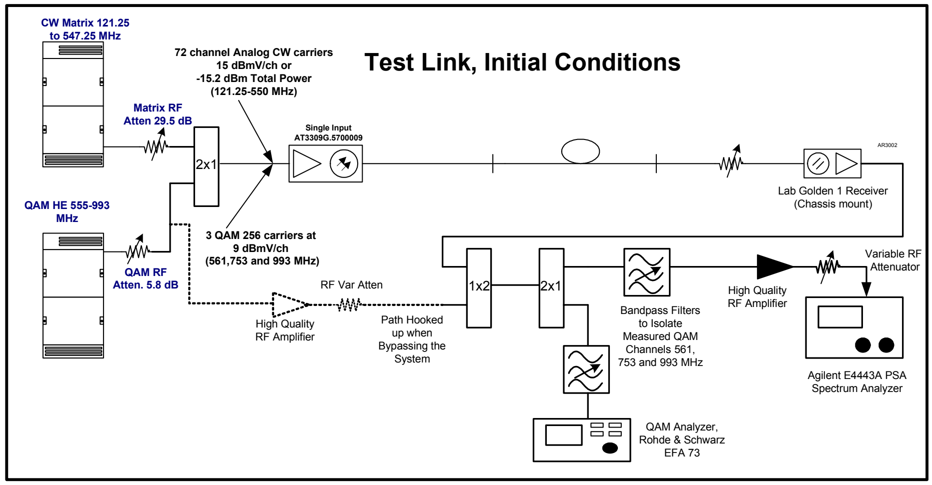

a) CW Load with QAM Test Channels

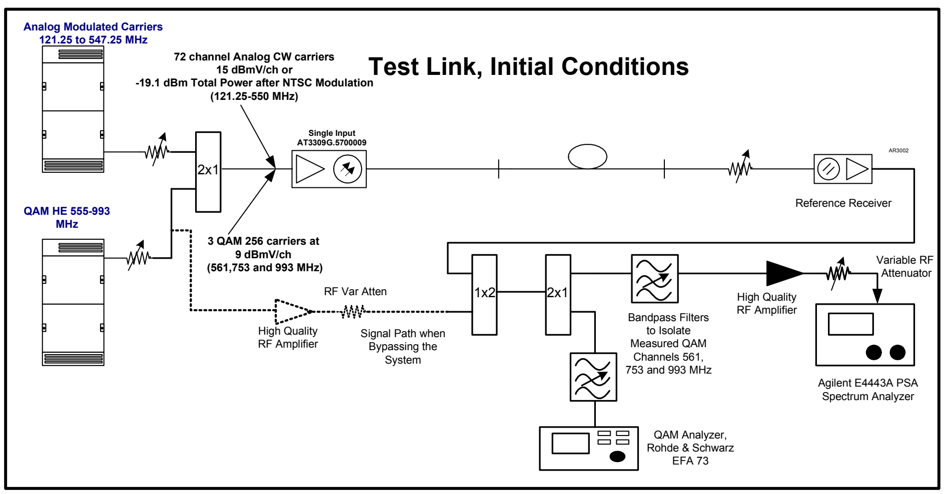

b) Analog Modulated Load with QAM Test Channels

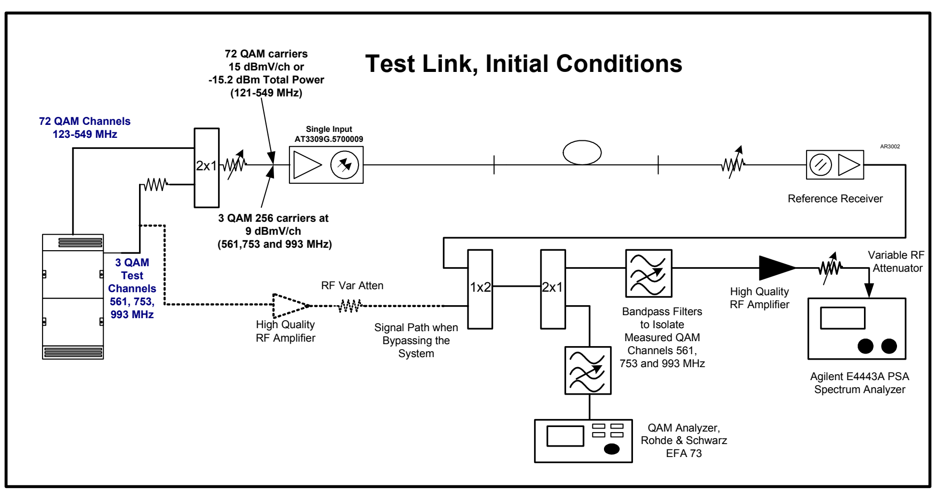

c) QAM Load with QAM Test Channels

Figure 2. Test Set Up

The channels between 54 and 88 MHz were not included in the testing for the following reasons:

- To eliminate nonlinear distortions at non-standard frequencies;
- To account for a trend in NGNA frequency subsplit evolution towards 85/105 MHz;
- For ease of alignment of NTSC and QAM upconverters.

Three 256 QAM test signals at 561 MHz, 753 MHz and 993 MHz were used throughout the testing to measure digital transmission performance. The test QAM signals were compliant with ITU Recommendation J.83 Annex C or Annex B. The BER for these signals were tested in pre-FEC mode and in post-FEC mode. Annex C (implemented for example in Japan) signals were used to test pre-FEC BER. These signals allow for real pre-FEC BER testing while signals compliant with Annex B (implemented for example in the USA and Canada) have embedded Trellis coding in their constellation mapping and hence there is no access to errors before Trellis coding. However, signals compliant with Annex B allow for setting interleavers with different parameters (see Table 1). Therefore, these signals were used to test interleaver effectiveness in correcting bit errors. The test signals were:

- 6 dB lower in level relative to CW carrier level for CW load,
- 6 dB lower in level relative to peak levels of the analog channels with NTSC analog video load, and
- 6 dB lower in level relative to the QAM channel levels for the QAM load.

<u>Litmus Test for Clipping</u>

To test the hypothesis that the laser clipping is not a dominant source of impairments to 256 QAM signals, a simple test was conducted: performance of the three test QAM channels was tested in two configurations, once while passing through the optical link with all other signals and second time after combining with the signals that were passed through the link after the optical receiver. If the clipping were the source of impairments, the BER in these two different configurations should be drastically different as the test QAM signals combined with the signals after the receiver (the second test) would not be subject to clipping.

To further test the clipping hypothesis, only the test QAM signal levels at the laser transmitter input were increased in 3 dB incremental steps and the BER performance was noted for all three levels. If clipping were the major contributor, the BER should not improve.

<u>Cause of QAM Signal Impairments</u>

An alternative hypothesis was tested based on the collected measurement results: that the QAM signal impairments were caused by the nonlinear distortions (second, third and higher orders).

<u>Character of Impairment and Their Effect on QAM Signals</u>

Three different types of load also allowed testing how the character of the nonlinear distortions affects QAM signals. The knowledge about the differences in the effect of distortions caused by different loads could allow for optimization of the links (laser transmitters) and the settings of the QAM signals.

Interleaver Efficiency against Different Impairments

Finally, the capability of adjusting interleaver settings for Annex B compliant QAM signals allowed testing the interleaver setting effectiveness in correcting errors caused by distortions of various characters.

CLIPPING OR OTHER PROBLEMS

QAM Signal Quality: Passed through Optical Link and Bypassing Optical Link

To determine whether clipping of the laser transmitter is a major source of QAM signal impairments, the performance of the test QAM signals was tested:

- when the test signals were transmitted on the same link as the load signals, and
- when the test signals were combined with the load signals after the optical receiver (only load signals were transmitted through the optical link).

The results of this simple test for all three types of load are presented in Figure 3.

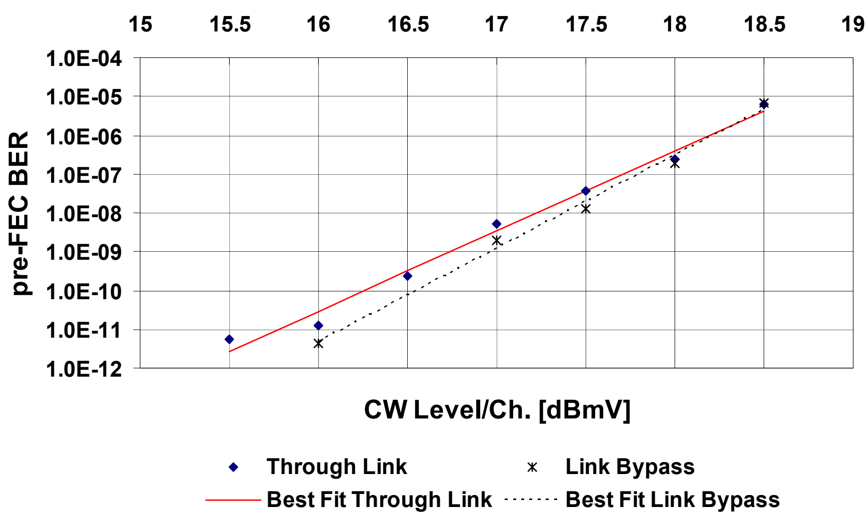

a) Test Signal BER vs. CW Signal Levels into the Transmitter

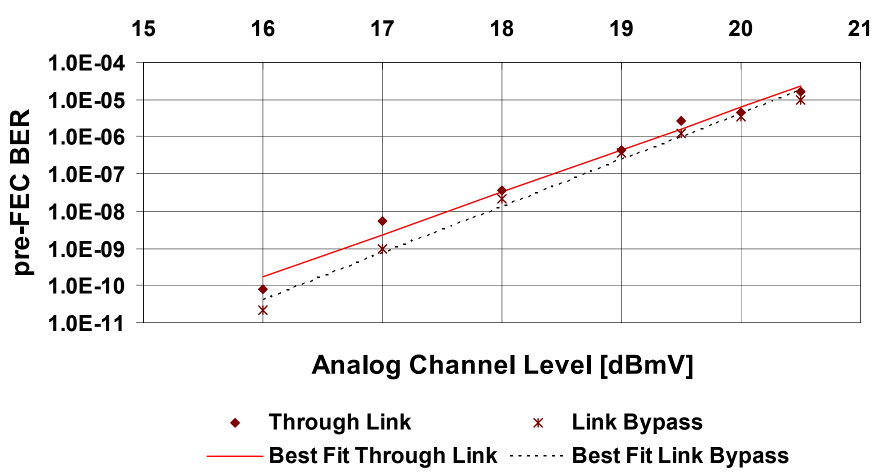

b) Test Signal BER vs. Analog Video Carrier Levels into the Transmitter

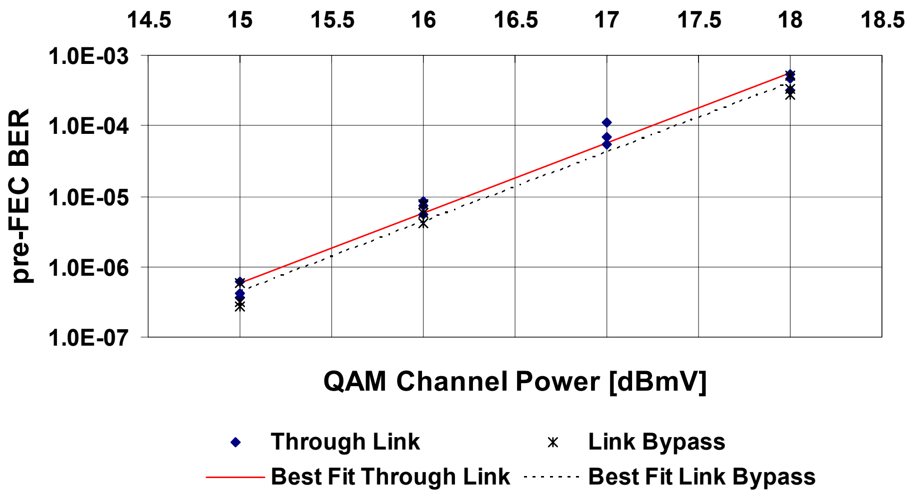

c) Test Signal BER vs. QAM Carrier Levels into the Transmitter

Figure 3. Performance of Test QAM Signals

The test results show that the performance of the test QAM signals are practically the same when they are passed through the transmitter and when they are combined after the optical receiver with the signals transmitted through the optical link and with all distortions generated by the load in the optical link. The logical conclusion from this test is that the distortions are the major contributor to the QAM signal impairments. At low levels of distortions (low input levels for CW and NTSC analog video carriers into the transmitter), some other impairments in the optical link may contribute to the slightly worse performance of the test QAM signals transmitted through the link but when distortions become dominant, the test QAM signal performance is the same in both test arrangements. This is opposite to the situation that would be caused by clipping (if clipping were dominant, the difference in performance of the test QAM signals in two different arrangements would increase with the increase of the transmitter load levels).

To re-confirm this hypothesis, an additional test was performed: at the high levels of load with CW carriers, the levels of the test QAM signals (compliant with Annex B) were increased in 3 dB steps (from 6 dB lower than CW carriers to the same power levels as the CW carriers) and the performance of the test QAM signals was recorded.

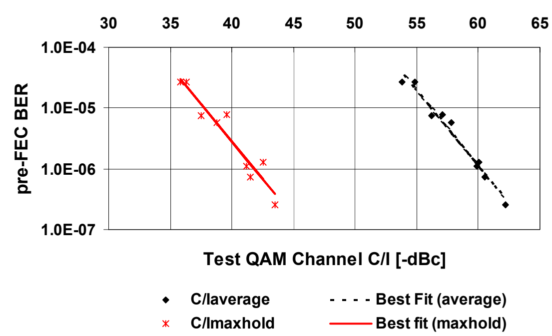

Figure 4. BER Improvements with Increase in Test QAM Signal Levels

The results were plotted versus C/I for QAM signals. Figure 4 shows BER performance as a function of C/I where “I” is a sum of power levels of all distortions within the channel. Two different sums were used to express I:

1. average RMS power of all distortions (measured as specified by NCTA Recommended Practices on Cable Television Systems, Third Edition),
2. 1-minute max-hold power with 1 kHz VBW setting of all distortions.

The second setting was used as a proxy of the peak value of the distortions throughout the entire testing process.

The performance of the test QAM channels improved linearly as their levels were increased (and the C/I was increased).

Carrier-To-Interference: Main Contributor to BER Degradation

The test results presented above clearly indicated that clipping was not a dominant source of QAM signal impairments at the transmitter driving levels tested. However, the BER performance of the QAM signals was strongly related to the levels of distortions. The next three figures present this strong log-log dependence with very high probability that this dependence is linear (on the log-log scale).

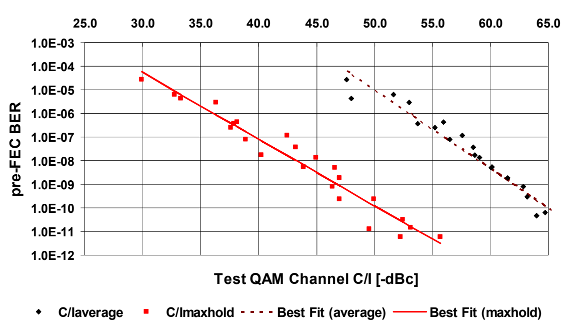

Regression Statistics for CW Load		
Parameter	Best Fit Average	Best Fit Maxhold
Multiple R	0.97164498	0.97175989
R Square	0.94409396	0.94431729
Adjusted R Square	0.94098807	0.94166573
Standard Error	0.43462985	0.49676825
Observations	20	23

Figure 5. BER versus C/I for CW Load (with Regression Analysis Statistics)

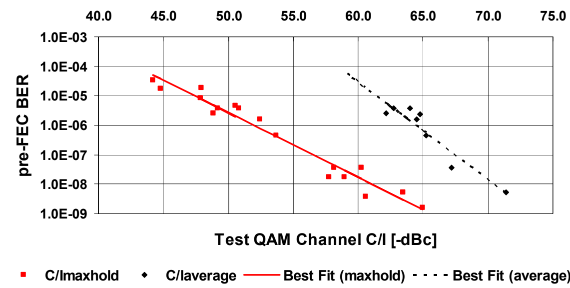

Regression Statistics for NTSC Modulated Load		
Parameter	Best Fit Average	Best Fit Maxhold
Multiple R	0.97822689	0.96476750
R Square	0.95692785	0.93077633
Adjusted R Square	0.95405637	0.92448327
Standard Error	0.31362594	0.29806038
Observations	17	13

Figure 6. BER versus C/I for NTSC Analog Video Load (with Regression Analysis Statistics)

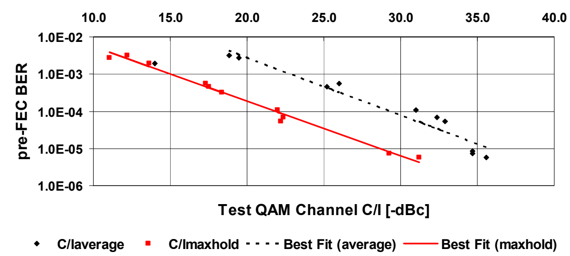

Regression Statistics for QAM Load		
Parameter	Best Fit Average	Best Fit Maxhold
Multiple R	0.97243812	0.99324842
R Square	0.94563589	0.98654243
Adjusted R Square	0.93959544	0.98504714
Standard Error	0.25499442	0.11693500
Observations	11	11

Figure 7. BER versus C/I for QAM Load (with Regression Analysis Statistics)

As previously described, “I” represents sum of average power levels or max-hold levels of all distortions within the test channels for CW and NTSC analog video loads. For QAM load, “I” represents integrated power of noise and noise-like intermodulation products within 6 MHz test channel.

The plots for CW and NTSC load show that BER vs. $C/I_{average}$ line is steeper than the line for BER vs. C/$I_{maxhold}$. This can be explained by the fact that the max-hold values increase faster at higher load levels (see Figure 8). For QAM loading, the difference between average and max-hold level of distortions is approximately constant.

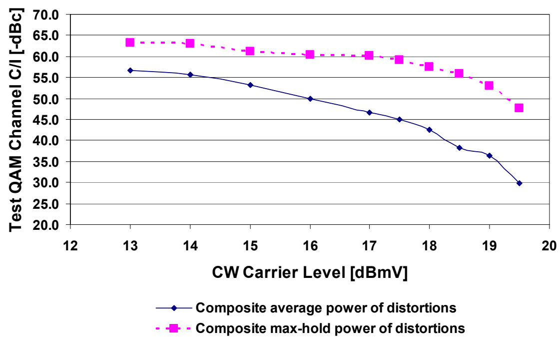

Figure 8. C/I versus Carrier Level for CW Load

Character of Impairments and Their Effect on QAM Signals

A comparison of plots in Figures 5 through 7 shows that BER values are different for the same C/I when "I" is caused by different loads. Obviously, the distortion characteristics are drastically different for distortion caused by CW load from those caused NTSC analog video carriers and those caused by QAM load. Figure 9 depicts screen shot examples of distortions for three different loads. Distortions generated by CW load are narrowband in nature (usually the power is concentrated within 30 kHz or narrower bandwidth), whereas distortions caused by NTSC analog video carriers are spread due to modulation sidebands and aeronautical frequency offsets and occupy up to several hundred kHz. Distortions caused by QAM load are wideband in nature and occupy entire bandwidth (noise-like character). The distortions caused by QAM signals are the result of the same nonlinearities of the optical link that causes nonlinear distortions if the loads are CW or NTSC analog carriers. The nonlinearity is clearly noticeable as their level rises with the input level rise into the transmitter and this rise accelerates at higher levels (higher order nonlinearities contribute to the distortion levels).

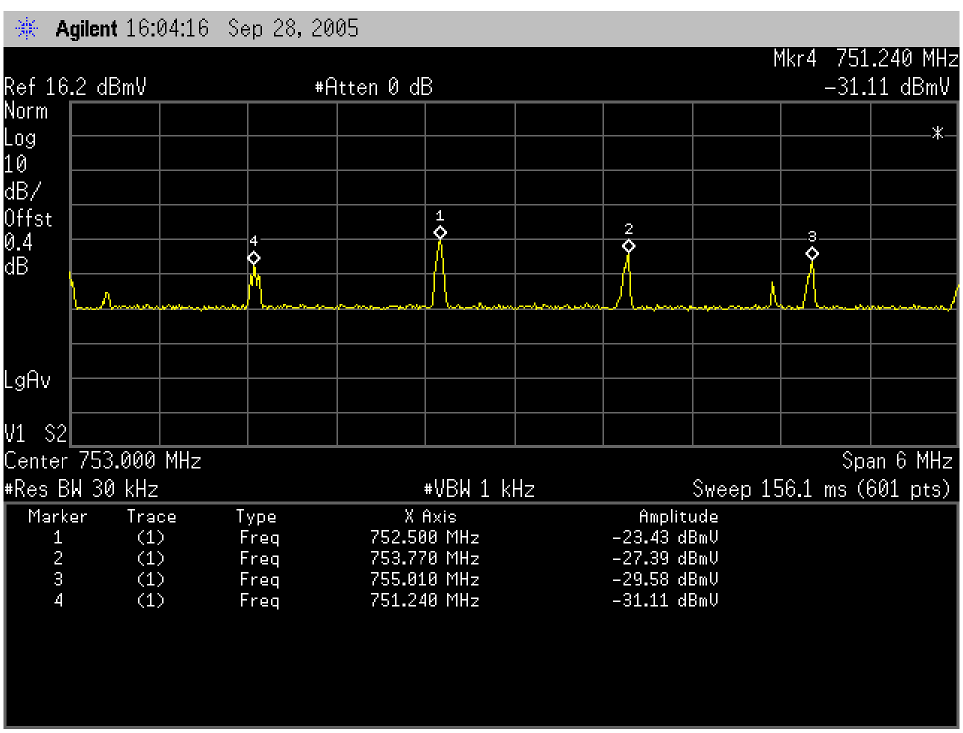

a) Distortions Caused by CW Load

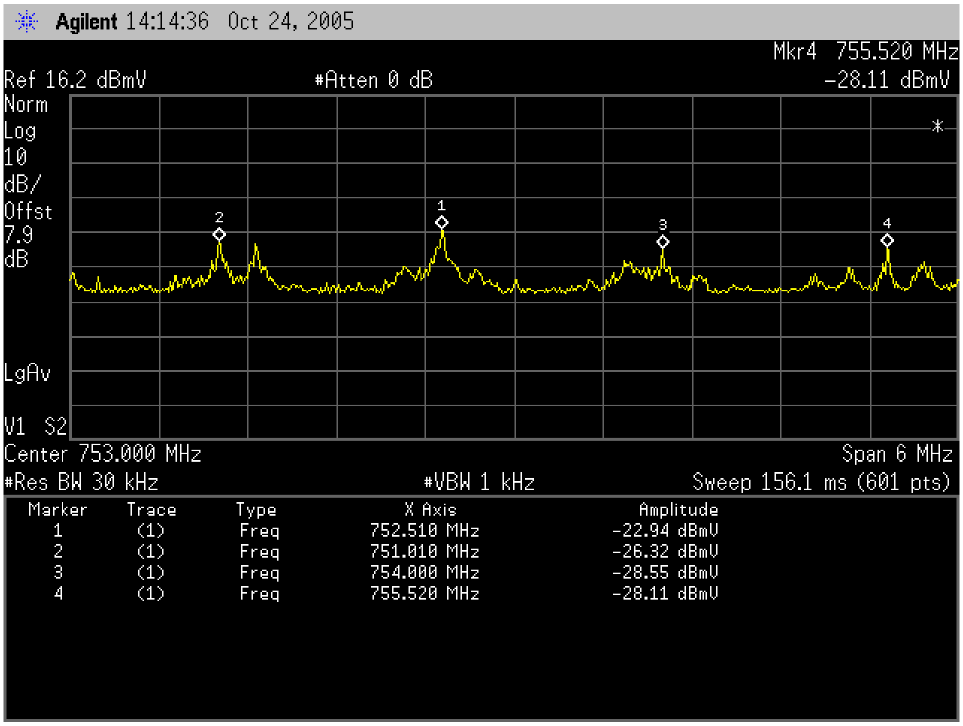

b) Distortions Caused by NTSC Analog Video Load

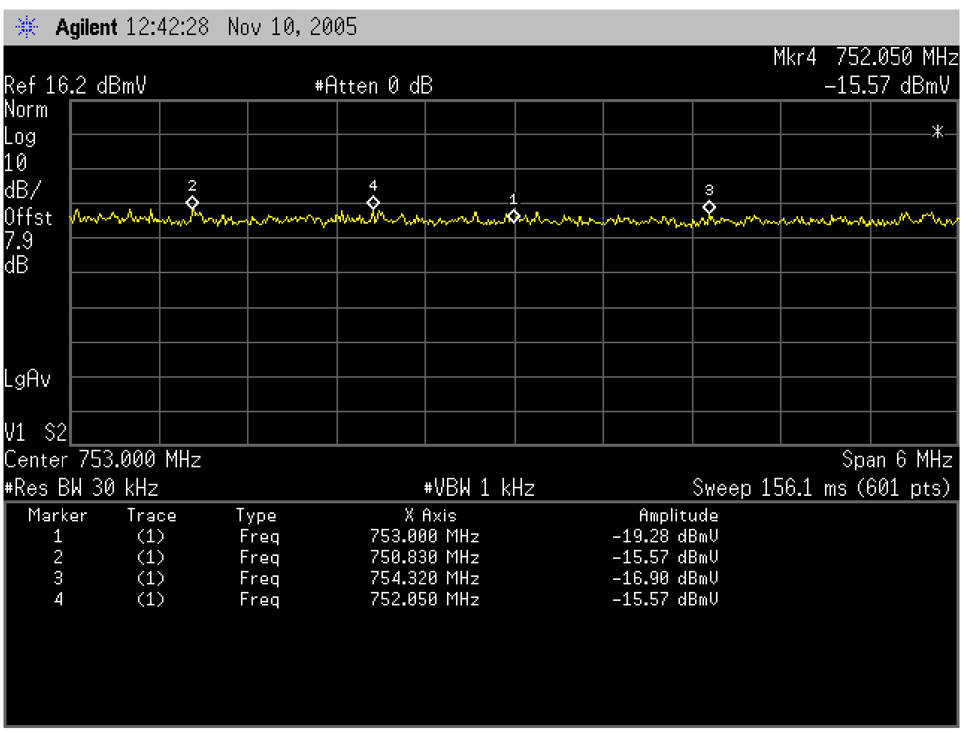

c) Distortions Caused by QAM Load

Figure 9. Different Distortion Characters for Different Load

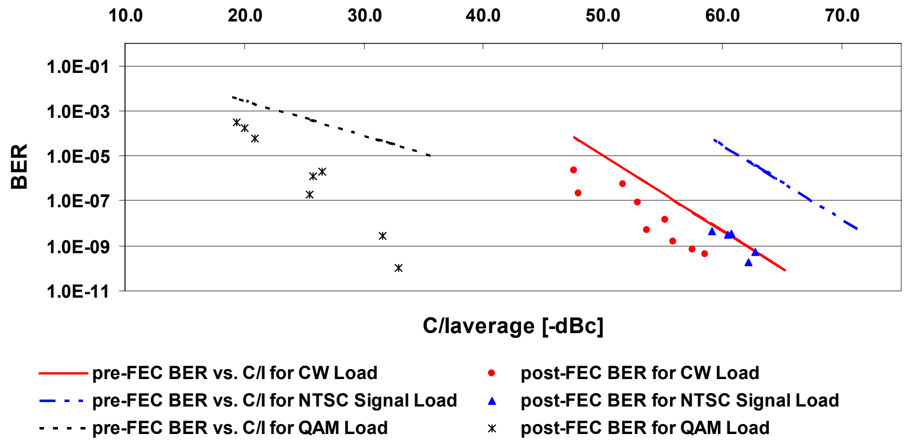

Figure 10. BER vs. C/I for Different Distortions

Figure 10 compares pre-FEC BER of QAM channels compliant with Annex C vs. C/I plots for all three types of load. It is apparent that the QAM signals are more sensitive to distortions generated by NTSC modulated video carriers than to distortions generated by CW carriers. This finding has been confirmed in the past [2]. The test results also show that QAM signals are least sensitive to the noise-like distortions caused by QAM load. This is also understood in light of commonly known and documented BER vs. SNR (or CNR) characteristics.

Figure 10 also presents post-FEC BER measured vs. C/I for Annex C signals. The test results clearly indicate different FEC effectiveness in correcting errors caused by distortions generated in optical links by different loads. This finding is investigated below.

Interleaver Effectiveness

Figure 11 presents FEC effectiveness in correcting errors for Annex C QAM signals (see Table 1). For distortions caused by QAM load, the interleaver setting in Annex C (I=12, J=7) allows FEC to correct pre-FEC errors occurring at a rate of 1.0E-5 and improve BER to levels below 1.0E-12. It is significantly less effective in correcting errors caused by distortions generated by NTSC load. Pre-FEC errors occurring at a rate of 1.0E-5 are corrected to BER levels below 1.0E-9. The effectiveness in correcting errors caused by CW load is very low for this interleaver setting (from 1.0E-5 pre-FEC BER to just below 1.0E-6 post-FEC BER)

Why such a big difference. Let us first analyze the distortion characteristics.

1. The distortions caused by CW carriers are the narrowest in bandwidth. The correlation times for a process are inversely proportionate to the process bandwidth. For 30 kHz CTB, the correlation time is 33 µs. This means that once the interference reached the peak values, they do not change appreciably over the time equal to the correlation time.
2. The distortions caused by modulated carriers (some of them with 12.5 kHz offset) are of much wider bandwidth. Hence, they have lower correlation times and the peak burst, if happens, is shorter in duration.
3. The intermodulation noise generated by QAM load is practically uncorrelated and hence it is not bursty in nature.

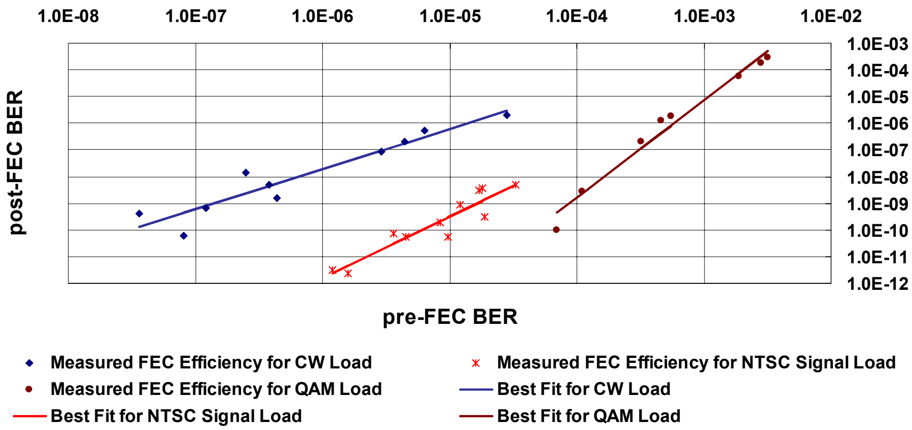

Figure 11. FEC Effectiveness

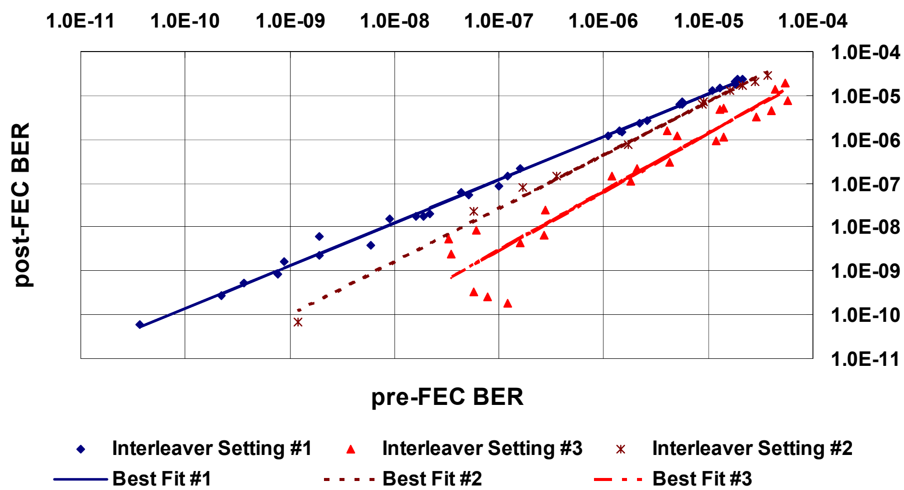

a) FEC Effectiveness for CW Load Distortions

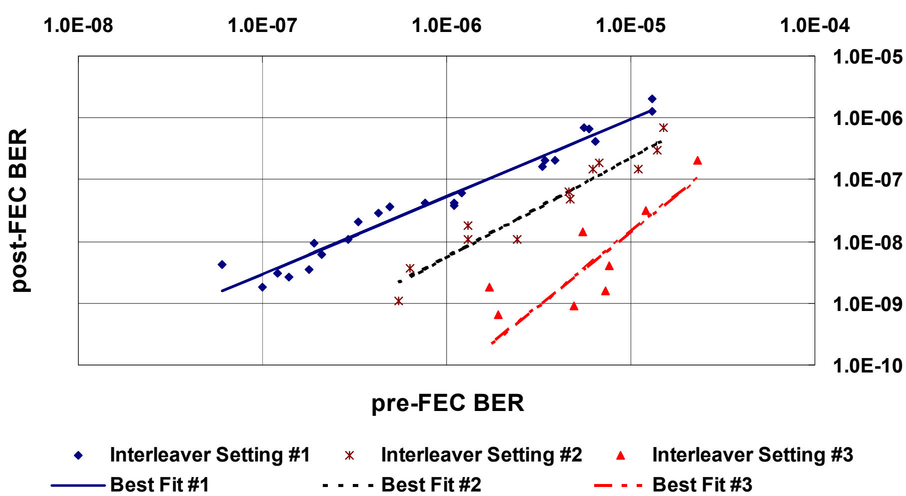

b) FEC Effectiveness for NTSC Load Distortions

Figure 12. Effectiveness of FEC for Different Interleaver Settings

For more detail explanation of this analysis, please refer to [3].

The conclusion of this analysis is that the depth of the interleaver is quite important in correcting errors caused by nonlinear distortions. To investigate effectiveness of the interleaver in correcting errors, Annex B QAM signals were used as they allow for flexibility in interleaver parameter setting. The following interleaver settings were tested:

1. I-16 and J=8,
2. I-32 and J=4, and
3. I-128 and J=1.

The test results indicated strong dependency of the FEC effectiveness in correcting errors caused by bursty interferences (distortions generated by CW and NTSC loads) on the depth of the interleaver. They also indicate that the FEC with the same interleaver setting is more effective in correcting errors caused by distortions generated by NTSC load (shorter burst duration). This finding bears significantly on the settings of the interleaver parameters for different services (see the latency contributions for different interleaver settings in Table 1). This is a well known finding (see also [4]) but applied here to the errors caused by nonlinear distortion of bursty character.

I (Number of Taps)	J (Increment)	Burst Protection (µs) 64 QAM/ 256 QAM	Latency (ms) 64 QAM/ 256 QAM	Comments
128	4	379/264	16/11	Annex B, for digital video only
128	3	285/198	12/8.4	
128	2	190/132	8/5.6	
128	1	95/66	4/2.8	Annex B
64	2	47/33	2/1.4	
32	4	24/16	0.98/0.68	
16	8	12/8.2	0.48/0.33	
8	16	5.9/4.1	0.22/0.15	
12	7	23/18	0.196/0.147	Annex C
12	7	18/14	0.156/0.116	Annex A

Table 1. Interleaver Parameters

Test results showed that an increase in the interleaver depth improved effectiveness of the FEC in correcting errors caused by nonlinear distortions. The importance of this effectiveness is obvious if we analyzed data from Table 2. To achieve QEF (Quasi Error Free) reception as defined for Annex B in North America, we need to achieve better than 2.6E-11 post-FEC BER.

Symbol Rate/Bit Rate	Average Error-Free Period	Required BER	Comments
5.3605370 Msps/42.885 Mbps (256 QAM Annex B)	2.5 sec	9.33E-8	Commonly Accepted
	1 min	3.89E-10	
	5 minutes	7.77E-11	
	15 minutes	2.59E-11	QEF for Annex B (Quasi Error Free)
	1 hour	6.48E-12	QEF time for Annex A (requires lower BER due to higher symbol rate)

Table 2. BER Requirements for Error-Free Reception

SUMMARY AND CONCLUSIONS

Load

The findings of the extensive testing effort are obvious: QAM signals, and especially 256 QAM signals, are very sensitive to distortions caused by analog load (see also [2]). Fortunately, only at levels exceeding the nominal input levels by 5 dB for NTSC analog video load, the pre-FEC BER approaches 1.0E-5 level for well aligned laser transmitters (see Figure 3). Even for CW load, the nominal input levels would need to be exceeded by 2 dB to approach this level of pre-FEC BER.

The encouraging test results for purely QAM load show that the QAM signals can be 6-7 dB higher than the levels they are set today on hybrid analog/digital optical links to approach 1.0E-5 pre-FEC BER levels and that the FEC is quite effective in correcting errors caused by distortions generated by QAM loads (see Figure 11). It seems that the optical links can be safely aligned 3-4 dB above their normal levels (6 dB below analog carriers) over the entire operational bandwidth for pure-digital load links at very good performance levels. This translates to lower transmitter cost (2 dB lower power for the same RX output) for pure-digital systems.

Alignment

One critically important conclusion for laser transmitter alignment is to concentrate on optimizing levels of nonlinear distortion throughout entire operational bandwidth as opposed on achieving the best performance within the analog load bandwidth.

Interleaver Setting

The crucial finding of the testing was that the specification requesting 1.0E-5 or better performance at all frequencies for 256 QAM Annex C signals is sufficient to achieve QEF transmission. An analysis of Figure 3 and Figure 12 allows for this conclusion. Figure 3 shows that for the level of CW load for which pre-FEC BER reaches 1.0E-5, the pre-FEC BER for NTSC load only slightly exceeds 1.0E-7. Figure 12 shows that for 256 QAM Annex B signals with interleaver setting of I-128 and J=1 and pre-FEC BER caused by NTSC load equal to 1.0E-7, the post-FEC BER is lower than 1.0E-11. This outcome is a result of several factors:

1. Significantly lower distortion levels for NTSC analog video load (10 to 12 dB lower);
2. Drastically improved effectiveness of the FEC in correcting BER caused by distortions generated by NTSC analog video carriers.

Although these improvements are partially offset by higher sensitivity of QAM signals to distortions generated by NTSC load (see Figure 10), the net result is still positive.

However, it is also important to note that the interleaver setting resulting in lower interleaver depth (for example, at I=32 and J=4), this conclusion does not hold true and the post-FEC BER would be higher than 1.0E-10, which may be not acceptable for video reception (see Table 2). Therefore, it is critical to select the interleaver settings that are optimal for the service provided with higher interleaver depth for error-sensitive services that can tolerate high latency and lower interleaver depth for services sensitive to latency. There are some services that may not allow for this optimization. Among these are IP video streaming and video games. For these services, signal placement at frequencies with lower level of interferences may be the only choice (if the operator does not wish to select better quality laser transmitters or higher power laser transmitters modulated at lower OMI).

ACKNOWLEDGEMENTS

The authors wish to express a deep gratitude to Ricardo Villa, who spent long hours to built all the signal sources and to collect a complete set of test equipment, familiarize himself with the test processes and to collect extensive set of data that allow for this analysis.

REFERENCES

[1] Gerard Terreault: *QAM Signal Impairments and Their Effects on MER and BER*, Version 1.04, Sunrise Telecom Broadband

[2] D. Stoneback, et al: *Distortion Beat Characterization and the Impact on QAM BER Performance*, NCTA 1999 Technical Papers

[3] Ron Katznelson: *Statistical Properties of Composite Distortions in HFC Systems and Their Effects on Digital Channels*, NCTA 2002 Technical Papers

[4] Shlomo Ovadia: *The Effect of Interleaver Depth and QAM Channel Frequency Offset on the Performance of Multichannel AM-VSB/256-QAM Video Lightwave Transmission Systems*, NCTA 1998 Technical Papers

[5] Darryl Schick: *Characterization of Peak Factor Effects in HFC Systems Using Phase Controlled Carriers*, NCTA 1996 Technical Papers

[6] ITU Recommendation J.83 Annex B

[7] ITU Recommendation J.83 Annex C

[8] O Sniezko: *Composite Triple Beat: CW versus Modulated Carriers – Perceptibility Levels and CTB Improvement*; Rogers, January 1991.

THE NEW BROADBAND EDGE: SETTING THE STAGE FOR THE "*CONVERGED SERVICES ARCHITECTURE*"

John Treece
Juniper Networks, Inc.

Abstract

Since the introduction of the first cable modem termination systems for MCNS support in the mid to late '90s, the access network has continued to challenge Moore's Law: Rapidly advancing technologies, developed to meet the growing needs of cable operators worldwide. The industry now finds itself at the inception of a new and even more complex phase/era. It includes, arguably, the most dramatic and potentially rewarding transformation the cable access network has ever seen.

INTRODUCTION

The genesis of the modular CMTS (M-CMTS) was originally conceived during the cable industry's Next Generation Network Architecture (NGNA) effort. This new architecture works to de-couple the historically integrated platforms, most notably the PHY and MAC layers.

By de-coupling the distinct functions of these devices, suppliers can focus on core competencies, thus creating greater innovation and expedited roadmaps.

At the same time, cable operators benefit from this increased innovation, and gain a more cost effective and flexible approach to scaling the last mile infrastructure and introducing new services. As an example, a near-term benefit with the introduction of these specifications will allow operators to leverage the same QAM modulators used today for MPEG video transport, for data -- essentially creating a "universal edge device." Longer term benefits include implementing a platform that will allow MSO's to migrate to an all-IP environment with the ability to support quad-play services.

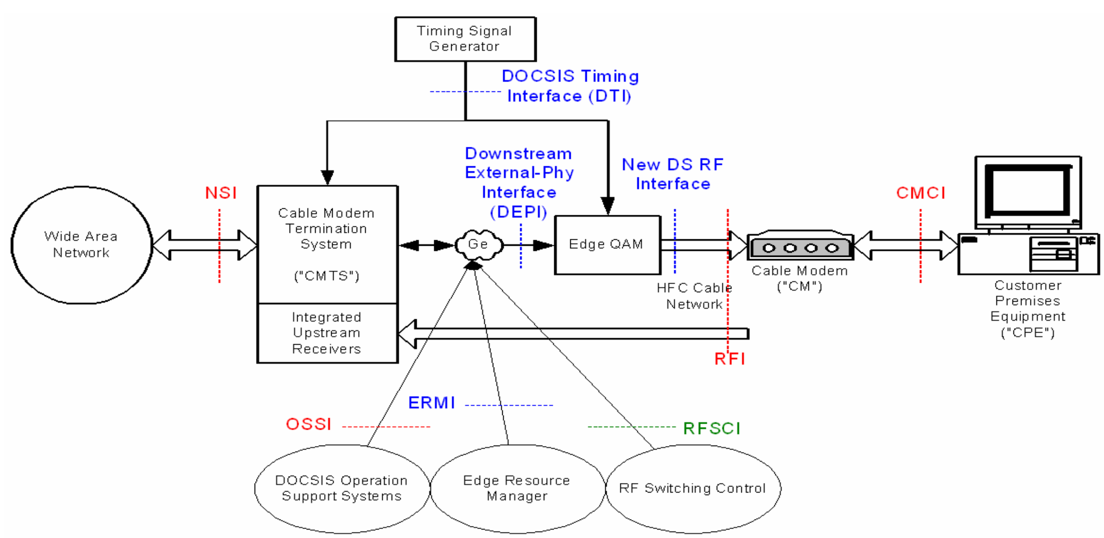

Figure 1

ARCHITECTURE

MSOs have invested upwards of $85 billion dollars over the last 30-40 years on the last mile infrastructure, and they are now starting to focus their capital, both intellectual and economic, on making better and more intelligent use of this existing infrastructure. One could argue that the value of the HFC networks is in its transparency. In other words, it's simply an analog signal pipe. It doesn't care about DOCSIS or QAM for that matter, and becomes a huge advantage for operators as you simply have to change the end points to run new and profitable applications. What this means however is, although we have this great and flexible HFC architecture, really the capabilities and the efficiency and the value of the HFC are dictated by the intelligence of the systems on either end.

This is evident by the emergence of many new initiatives such as, Digital Simulcast, Switched Digital Broadcast, the deployment of Next Generation On-Demand architectures, the move towards downloadable security, new ad insertion technologies, and finally new access network initiatives such as the modular CMTS.

The M-CMTS architecture is comprised of several different components as shown in Figure 1. They include the Timing server which supplies a synchronous DOCSIS clock and timestamp for each component; the Universal Edge QAM, capable of supporting both the DEPI and the DTI Timing interface; the MAC Core which contains all legacy CMTS downstream components, with the exception of the PHY; and the Edge Resource Manager which allocates bandwidth for both native MPEG video and DOCSIS which is primarily used for VoD today.

This new architecture will not only allow for reduced cost in hardware, but also reduced operations cost, because of associated reductions in complexity. An immediate savings in hardware can be seen by comparing today's model to a "pay as you go" model. When introduced, the current generation CMTS added improved density, high availability capabilities, as well as other feature upgrades from the first generation CMTS. These served the needs of the primary services at the time, essentially data services. Today, if MSOs need an additional downstream or upstream port for capacity relief, tied to it are either an additional physical downstream port or multiple upstream ports -- that operators may *not* need.

In addition, the move towards true convergence which will pave the way to "all IP" transport is necessitating the requirement to couple together robust Layer 3 functionality and traditional DOCSIS processing. A common delivery mechanism such as IP will allow for greater innovation and the evolution of new applications that work to support the "anything, anytime, anywhere" mentality that is rapidly emerging in today's consumer.

Figure 1 represents a logical overview of this architecture taken from Cable Labs who produced the initial four interface specifications. They are; Downstream external PHY (DEPI), DOCSIS Timing interface (DTI), M-CMTS Operations Support System Interface (MOSSI), Edge Resource Management Interface (ERMI), and the DOCSIS Radio Frequency Interface (DRFI).

An additional advantage to de-coupling the MAC and PHY portions of the architecture is the ability for the universal edge device to evolve without the burden of being physically linked to the MAC processing components allowing suppliers to innovate and create solutions that were once unfeasible. Because the PHY portion of the network is relatively simple, operators may now look at different physical media that may provide increased efficiency, such as the use of optics. For example, this would eliminate the need for optical-to-RF and RF-to-optical transitions respectively that take place in the remote hub sites today.

FEATURE SUPPORT

With required support for new features such as IPv6, Channel Bonding, new advanced security algorithms, and advanced Layer 3 support, these new platforms will require industry-leading expertise in both DOCSIS and IP. Whereas the features of the CMTS were primarily service-driven in the past, they will now be required to adapt to the advancements in DOCSIS, and ensure interoperability with the complex routing schemes that operators are putting in place. These new networks are being designed to efficiently transport video, data, and voice traffic locally, regionally, and nationally.

This means that next-generation access platforms will require seamless interoperation with these new robust metro networks to provide better scale and flexibility. They'll also need to offer multiple operational configurations to meet the needs of the MSO. And, they'll need to do all of this in a more cost efficient manner.

Of all the new features introduced in DOCSIS 3.0, IPv6 and Channel Bonding in particular have the potential to act as a catalyst for growth and future opportunity. All access network components that support voice, video, and data, will be required to support the legacy IPv4 address space, as well as IPv6. With the number of cable modem users in North America approaching 30 million and the rapid deployment of new services such as VoIP, operators are finding it more and more difficult to procure public address space for the growing number of IP devices. At the same time, the private IP space used for management is also becoming more difficult to manage. With some OSS systems evolving to a more centralized approach, the use of overlapping federated address space becomes less of an option. For these reasons, it is conceivable to think that we will see infrastructure support for IPv6 introduced for some operators in the 2006-2007 timeframe. This is another example of the critical requirement for new access network technologies to not only support historical DOCSIS processing, but to have the ability to support such new features as IPv6.

The need for greater downstream capacity has been analyzed intensely over the past few years as applications have emerged that require more and more bandwidth. This development, combined with increasing competitive pressures, has made Channel Bonding a priority, as evidenced by its inclusion in the now pending DOCSIS 3.0 specification.

Currently, a single 6MHz channel is used with the capability of producing up to 38Mbps (using 256QAM modulation) minus overhead. Operators today have service offers ranging from 1.5-15Mbps downstream. Channel Bonding has the ability to take that 38Mbps and increase it to speeds over 100Mbps with room to grow. Bonding also affords some additional technical advantages such as statistical multiplexing gains within a bonding group. However, data is not the only driver that can take advantage of this increased capacity to the home. As MSOs begin the deployment of IP capable set-tops boxes for video services, the new devices are being developed to support DOCSIS and will have the ability to tune to bonded streams. This becomes increasingly important as high definition content, which requires nearly triple the bandwidth of a standard definition stream, continues to grow as consumer demand also increases.

IMPLEMENTATION

Although MSO's are heavily engaged in the development during this innovative period, they are equally sensitive to protecting their investments and reducing stranded assets. Figure 2 shows a potential migration path that accomplishes the goals of introducing a flexible and scalable architecture, while re-purposing existing deployed assets.

The first diagram to the far left (Phase I) represents today's infrastructure with an integrated CMTS and separate QAMs for uni-cast video traffic. The second diagram (Phase II) shows the initial phase of implementation where the downstream PHY is separated from the legacy CMTS with the MAC core being introduced for Layer 3 and DOCSIS processing. In this phase, the downstream edge QAM device is now used for data, voice, and uni-cast video traffic. Since the upstream interface has yet to be defined, the upstream traffic is handled by the existing CMTS which is re-purposed to act as an upstream termination point.

During this phase, both upstream demodulators or burst receivers and MAC processing is handled within these devices. The final diagram (Phase III) represents a complete separation of MAC and PHY for both the downstream and the upstream.

Although formal work has yet to begin on the upstream interface, there still seems to be a division of thought relative to remote or centralized MAC layer processing on the upstream. Independent analysis continues to weigh the pros and cons of a MAC-less upstream PHY termination device.

Another interesting opportunity is the idea of consolidation of DOCSIS and IP routing functionality in the remote distribution hubs. Today, the CMTS is a consolidated MAC/PHY device that performs limited Layer 3 routing functions. These devices then route to a local aggregation router that usually sits within the same facility. However, the ability to combine robust Layer 3 routing capabilities with the DOCSIS processing can be an

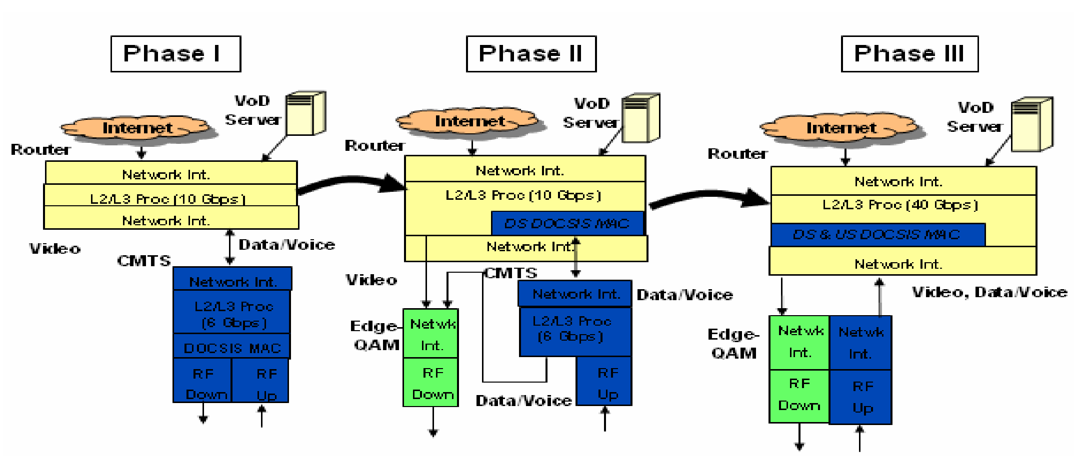

Figure 2

attractive alternative to today's architecture, particularly as operators move towards and all-IP environment. Because DOCSIS has reached a state of maturity within the industry, combined with the separation of the physical layer, it may now be possible to build a DOCSIS processing component or "blade" that lives within a traditional aggregation router. This option has the potential to reclaim valuable real estate within the headends, reduce power and environmental needs, and sets the stage for convergence with one device providing routing and processing for video, voice and data services.

This option provides a true solutions architecture with the ability to scale and evolve with the growing needs of today's consumer.

SUMMARY

Convergence has been a common topic of discussion throughout the last few years, however it seems now more then ever it is becoming a reality. Convergence has historically been categorized and discussed independently as; network, service, or in-home devices. However, it is now quite evident that convergence must and, in fact is, emerging within each of these domains, which is necessary to achieve the "anything, anywhere, anytime", content delivery objective. With this imminent evolution, it is imperative that the edge platforms provide the robust capabilities that will be needed to provide reliable performance for a plethora of services without being limited in their ability to grow with the network.

TRANSITIONING TO M-CMTS

Andrew Page, John T. Chapman, and Carmelo Iaria of Cisco Systems, Inc.
Glenn McGilvray of Scientific-Atlanta, Inc.

Abstract

The CableLabs® Modular Cable Modem Termination System (M-CMTS™) series of specifications introduces powerful new tools to enhance the architecture of the CMTS. Cable operators have several options for migrating from the current DOCSIS architecture to the M-CMTS architecture, and must evaluate the options relative to their individual success criteria to determine how best to migrate to M-CMTS.

To assist cable operators in formulating an M-CMTS migration strategy, this paper provides an overview of the options and highlights the advantages and disadvantages of each approach. The paper identifies and analyzes key criteria for selecting the M-CMTS Core and Edge QAM (EQAM) options and concludes with a recommended approach for migrating to an M-CMTS architecture.

1.0 INTRODUCTION

The Next Generation Network Architecture (NGNA) framework defined in 2004 by Comcast, Time Warner Cable and Cox Communications asserted that the DOCSIS® CMTS could and should be modularized. These cable operators had the foresight that next generation broadband competition was mounting in the form of xDSL, FTTH and other technologies. To remain competitive, they needed to radically increase the capacity of their DOCSIS networks, while reducing their costs and improving their operational efficiencies. They also needed to maximize the return on investments from their DOCSIS networks by deploying a flexible infra-structure, capable of supporting DOCSIS 1.x, 2.0 and 3.0 features simultaneously, as well as transitioning all services to a converged, end-to-end IP network. The operators recognized that a modular CMTS architecture could fulfill these requirements and be the key ingredient of the next-generation cable network. In addition, with a modular CMTS architecture, best-of-breed solutions could be deployed, allowing operators to choose components from multiple vendors based on individual product strengths and the criteria deemed most important by each operator.

The M-CMTS series of specifications published in 2005 by Cable Television Laboratories (CableLabs) codified the CMTS architecture envisioned by the NGNA architects. The specifications clearly define the functional components of the architecture and the external interfaces supported by each component. As with most technical speci-fications, vendors were given the freedom to take different approaches to implementing the M-CMTS architecture. Cable operators must, therefore, evaluate each approach relative to their success criteria to determine which approach to follow. To assist cable operators in formulating an M-CMTS deployment strategy, this paper provides an overview of the available options and highlights the advantages and disadvantages of each option.

2.0 M-CMTS ARCHITECTURE

Prior to the release of the M-CMTS specifications, the DOCSIS CMTS was defined as a single chassis that integrates the DOCSIS MAC and PHY functions and all upper layer protocols. Herein, this type of CMTS, is referred to as an Integrated CMTS (I-CMTS).

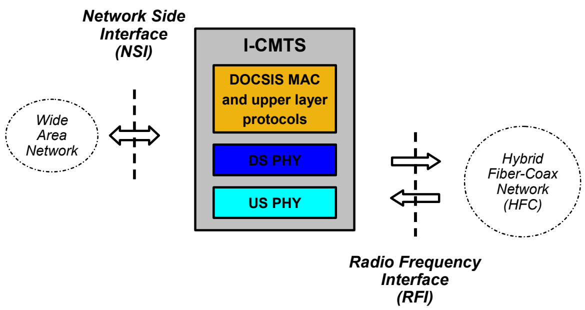

Figure 1. I-CMTS

According to Synergy Research Group, cable operators, worldwide, purchased more than $2.0 billion of I-CMTS equipment from 2002 to 2005, including nearly 35,000 I-CMTS chassis. With such a large and valuable embedded base of I-CMTSs, cable operators must leverage these assets to their full extent as they transition to an M-CMTS architecture.

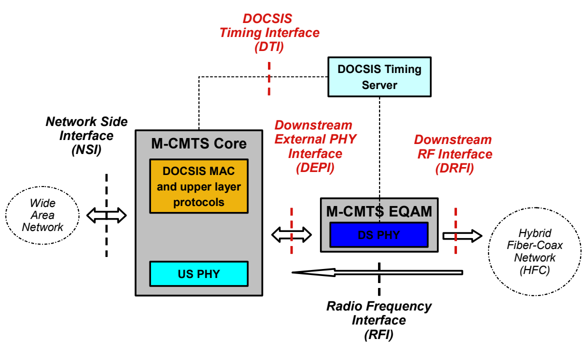

Figure 2. M-CMTS Phase 1

The M-CMTS architecture partitions the I-CMTS into a set of at least two physically separate components with standard interfaces between them. The M-CMTS Core was so named because it retains the core functionality of the I-CMTS, namely the DOCSIS MAC and upper-layer protocols (referred to collectively as Layer 2+). The M-CMTS Edge-QAM (EQAM) supports the downstream PHY function and the Downstream External PHY Interface (DEPI), which is the interface and associated protocol by which the M-CMTS Core and EQAM communicate with each other. The M-CMTS EQAM must comply with the Downstream Radio Frequency Interface (DRFI) specification, which is an update to the original DOCSIS RFI specification and includes requirements for transmitting multiple QAM channels per physical RF port, a common practice with state-of-the-art EQAM technology. Both the M-CMTS Core and EQAM must comply with the DOCSIS Timing Interface (DTI) specification which defines the interface and associated protocol for synchronizing the M-CMTS components with a common timing reference.

The upstream PHY function can either be supported by the M-CMTS Core or in a physically separate component. However, the current release of the M-CMTS specifications do not define a separate upstream PHY component or the protocol by which the M-CMTS Core and upstream PHY component would communicate with each other. The upstream PHY component and interface [referred to herein as upstream receiver and Upstream External PHY Interface (UEPI)] will be addressed in a future release of the M-CMTS specifications. Of course, vendors will have the freedom to implement the upstream receiver as a standalone device or integrate it with the M-CMTS EQAM.

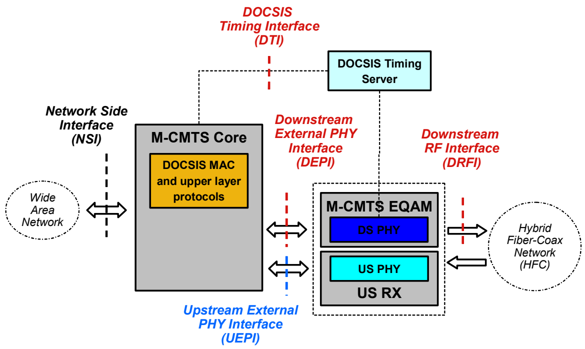

Figure 3. M-CMTS Phase 2

The fact that DEPI has been standardized, while the upstream receiver and UEPI are not yet standardized, implies a phasing of M-CMTS solutions. Such phasing should be considered when formulating an M-CMTS deployment strategy. For the purposes of this paper, we define two M-CMTS phases, as shown in Table 1.

	M-CMTS Core	**M-CMTS EQAM**	**Upstream Receiver**
Phase 1	Layer 2+ Upstream PHY	Downstream PHY	(none)
Phase 2	Layer 2+	Downstream PHY	Upstream PHY

Table 1. M-CMTS Phases

In Phase 1, the M-CMTS Core performs all functions of an I-CMTS, except for the downstream PHY function. In this phase, the upstream ports are integrated in the M-CMTS Core, while the downstream ports are supported in the EQAM. In Phase 2, the M-CMTS Core is purely a Layer 2+ component. An upstream receiver is introduced to support the upstream PHY function.

3.0 MIGRATING TO AN M-CMTS ARCHITECTURE

The M-CMTS architecture splits the I-CMTS into an M-CMTS Core, EQAM and (optionally) upstream receiver. There are two approaches to transition from an I-CMTS to an M-CMTS solution:

- The first approach (which we will refer to as Option A throughout the text) is to add the required interfaces to the I-CMTS so it can operate as an M-CMTS Core and utilize an external EQAM and (optionally) upstream receiver.

- The second choice (which we will refer to as Option B through the text) is to add the required interfaces to the I-CMTS so it can operate as an M-CMTS EQAM and (optionally) upstream receiver, and install a new M-CMTS Core.

This section provides an overview of these options and discusses the key criteria for assessing the M-CMTS Core product choices (section 4.0) and EQAM product choices (section 5.0).

In addition to the M-CMTS Core and EQAM implementation options, cable operators may also have options for implementing DTI when building an M-CMTS solution. The DTI system is based on a client/server architecture. The DTI client must be implemented on both the M-CMTS Core and EQAM components, while the DTI server theoretically can be hosted on the M-CMTS Core, on the EQAM, or on a separate computing platform. However, in this paper, we assume a new DTI Sever is required in all cases.

3.1 Transitioning the I-CMTS to an M-CMTS Core (Option A)

Since the M-CMTS Core retains most of the functionalities of an I-CMTS, the most logical approach to building an M-CMTS solution is to evolve the existing I-CMTS into an M-CMTS Core. With this option, the migration can be gradually accomplished in two phases: Utilize a low-cost, external M-CMTS EQAM to perform the downstream PHY function (in Phase 1) and then add an upstream receiver to extract the upstream PHY function (in Phase 2).

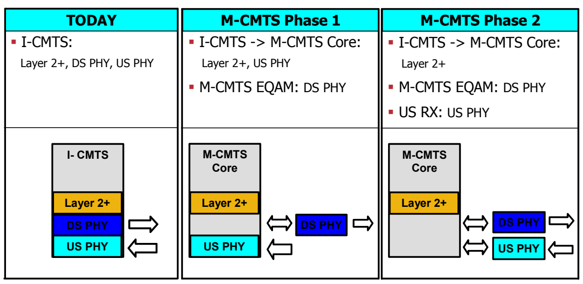

Figure 4. M-CMTS Migration Option A

Executing this option involves the following high-level tasks.

Repurposing the I-CMTS into an M-CMTS Core (Phase 1):

1. Upgrade the I-CMTS to M-CMTS Core

a) Add new line card(s) that support the downstream MAC function and Gigabit Ethernet (GE) ports to the existing I-CMTS.

b) Add/upgrade the I-CMTS hardware to comply with DTI.

c) Upgrade the I-CMTS software to support DEPI, DTI and other M-CMTS-related features.

2. Install a new M-CMTS EQAM(s) that is compliant with DEPI, DTI, and DRFI (or upgrade an existing video EQAM's hardware and software to comply with the M-CMTS specifications).

3. Install a new DTI Server.

Repurposing the I-CMTS into an M-CMTS Core (Phase 2):

1. Enhance the M-CMTS Core to support external upstream PHY

a) Add new line card(s) that support the upstream MAC function and GE ports to the existing I-CMTS (if not supported on the downstream MAC line card).

b) Upgrade the I-CMTS software to support UEPI.

2. Install a new external M-CMTS upstream receiver or add upstream-only PHY line card(s) to the M-CMTS EQAM that supports DTI and UEPI.

The I-CMTS is the centerpiece of the DOCSIS network today. It possesses the intelligence for managing all DOCSIS features and services. The success of a cable operator's DOCSIS service offerings hinge upon the performance and reliability of the I-CMTS. Evolving the I-CMTS into an M-CMTS Core retains the value of the I-CMTS and avoids the risks associated with introducing a new M-CMTS Core.

In Phase 1 of this approach, cable operators are able to use the existing upstream ports on the I-CMTS as part of the M-CMTS solution. In fact, there is no change to the upstream functionality of the I-CMTS in Phase 1. Cable operators can gradually transition to M-CMTS with lower risk of service interruption.

This M-CMTS migration approach is a direct application of a fundamental premise of the M-CMTS architecture: the cost of growing the downstream channel capacity of a DOCSIS network can be dramatically reduced by extracting the downstream PHY function from the I-CMTS and utilizing external EQAM devices that have been cost-optimized for price-sensitive video-on-demand (VoD) applications. Although the M-CMTS EQAM must support new features and functionality specific to the M-CMTS architecture, it is very similar to the video EQAM. Therefore, equivalent pricing is expected. The use of external EQAM devices are not only cost-effective, but also present cable operators with an opportunity to deploy purpose-built products in their DOCSIS networks from the leaders in the EQAM market, and even use the same EQAM device to support both video and DOCSIS applications.

<u>3.1 Transitioning the I-CMTS to an M-CMTS EQAM (Option B)</u>

The second option for building an M-CMTS solution is to install a new M-CMTS Core chassis and evolve the existing I-CMTS into an M-CMTS EQAM and upstream receiver.

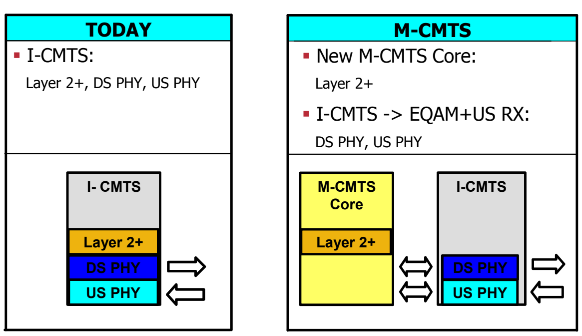

Figure 5. M-CMTS Migration Option B

Executing this option involves the following high-level tasks.

<u>Repurposing the I-CMTS into an M-CMTS EQAM</u>

1. Install a new M-CMTS Core chassis with DOCSIS Layer 2+, DEPI, UEPI and DTI support.

2. Repurpose the I-CMTS into an M-CMTS PHY Layer chassis

a) Add downstream-only PHY line card(s) with DRFI-compliant RF ports to the existing I-CMTS (or utilize the existing line cards in the I-CMTS to support the downstream PHY function)

b) Add upstream-only PHY line card(s) to the I-CMTS (or utilize the existing line cards in the I-CMTS to support the upstream PHY function)

c) Add/upgrade the I-CMTS hardware to comply with DTI

d) Upgrade the I-CMTS software to support the DEPI, UEPI and DTI protocols and other M-CMTS-related features.

3. Install a new DTI Server.

Considering that the EQAM and upstream receiver are the low-value components of the M-CMTS architecture, and are the targets for cost optimization, repurposing the existing I-CMTS into an M-CMTS EQAM and upstream receiver devalues the cable operator's investment in the I-CMTS. Introducing a new M-CMTS Core forces the cable operator to re-purchase the high-value Layer 2+ function of the I-CMTS and could introduce risk if the new M-CMTS Core is unproven in a DOCSIS network.

To independently scale the downstream and upstream channel capacity in this approach requires new PHY line cards in the I-CMTS. Considering the advantages of a standalone EQAM product (as described in section 5.0), it is probably more prudent to invest in the standalone EQAMs than in additional PHY line cards for the I-CMTS.

4.0 M-CMTS CORE CHOICES

Making the right choice for the M-CMTS Core is a fundamental step to successfully migrate from today's I-CMTS to an M-CMTS architecture and meet the strategic imperatives that drive such a migration.

In review, the two options for implementing an M-CMTS Core are:

- Option A: Evolve an existing I-CMTS to become the M-CMTS Core
- Option B: Choose a completely new M-CMTS Core

This section evaluates the pros and cons of these two options. Underlying assumptions are: (1) The cable operator has successfully deployed High-Speed Data (HSD) and Voice over IP (VoIP) DOCSIS services, and (2) Throughout the transition to an M-CMTS architecture, the operator must sustain or grow these services and offer new services.

Each option will be tested against the following key evaluation criteria:

- Support for existing DOCSIS features and services.
- Scalability for existing DOCSIS services.
- Support for new DOCSIS services (next 3-5 years).
- Capital budget impact (next 5 years).

The first objective in any network architecture migration is to implement the migration with minimal or no impact to existing services. For this reason, when migrating from an I-CMTS to an M-CMTS architecture, it is critical to ensure that the M-CMTS Core can support all DOCSIS features and services required by the cable operator.

Primarily, this means that all Layer 2 and Layer 3 features currently supported by the
I-CMTS must be supported by the M-CMTS Core so that the same level of service, security and performance can be offered to subscribers. At the same time, and perhaps more critical to the success of the migration, it also means that the cable operator must be able to deploy and manage the M-CMTS Core without significant changes to back-office systems and operating procedures.

Another consideration in assessing the ability of an M-CMTS Core to support existing DOCSIS features and services is the stability of the design. Considering that the primary (Layer 2+) functions of the existing I-CMTS do not change in the M-CMTS architecture, operators can be confident that the I-CMTS will be able to perform equally well in supporting DOCSIS features and services when operating as an M-CMTS Core. If the new M-CMTS Core platform that would be deployed under Option B is unproven in DOCSIS networks, this option adds risk to the operator's ability to maintain existing service levels and operational effectiveness.

Based on the above considerations, choosing Option A (Phase 1) as the first step toward the migration to an M-CMTS architecture reduces the risk of impacting existing DOCSIS services and is a safer approach.

Scaling Existing DOCSIS Services

If seamless support of existing DOCSIS service is a key decision criteria for the M-CMTS Core, migrating from the existing I-CMTS to an M-CMTS architecture would not make much sense unless the M-CMTS architecture could meet the increasing scaling requirements posed by the projected growth of the HSD and VoIP services.

Obviously such growth will depend on a number of factors which are very specific to each region of the world; there is no single rule that can be applied everywhere. Generally speaking though, the trend for the broadband cable market is toward much higher download speeds. HSD tiers of 10 to 30 Mbps in the downstream and from 1 to 5 Mbps in the upstream are expected to be quite common in the next 12 to 24 months, with a number of geographies in the 50 to 100 Mbps range for the downstream. At the same time, penetration rates of HSD service are expected to reach the 50% to 60% range as more affordable broadband tiers are offered to entice dial-up users to upgrade to broadband. What the 'H' in HSD will stand for in just a year or two will be dramatically different than what we're experiencing today.

On the voice front, more and more cable operators are now successfully deploying VoIP services. Bundling VoIP with HSD and video services has proven to be a driving force behind the continued growth of the VoIP subscriber base, as well as a powerful tool for reducing subscriber churn.

The success of these DOCSIS services translates into three fundamental requirements relevant to scaling existing services that cable operators must consider when it comes to the M-CMTS Core decision:

- Enhanced Forwarding Capability
- Strong Control Plane Scalability
- Support for key features such as channel bonding and IPv6

Each of these requirements will now be examined in detail.

Enhanced Forwarding Capability

Enhanced Forwarding Capability means that the overall forwarding capacity needs to be much higher than today's I-CMTS. If an I-CMTS is today providing 1 Gbps downstream capacity to a given serving area, a good benchmark for an M-CMTS Core is the ability to offer at least 5 Gbps to the same serving area, with the ability to migrate to 10 Gbps. Although there will be deployments that require a significant increase on the upstream bandwidth, it is anticipated that, at least in the short term, the HSD services will remain predominantly an asymmetrical service. This is especially true in light of the fact that the content being accessed by the residential broadband subscribers is becoming richer and richer and that music and video downloads are becoming a standard on any Internet portal.

Control Plane Scalability

While the evolution of the HSD service offerings will dictate the forwarding capacity requirement for an M-CMTS Core, growth in broadband penetration rates, and increased adoption of VoIP among broadband residential subscribers, will require the M-CMTS Core to have much greater control plane scalability than today's I-CMTS. Projections show that an M-CMTS Core must be able to support the control plane traffic for at least 50,000 to 60,000 DOCSIS CPE devices, and ideally up to 80,000 devices.

Support for Channel Bonding and IPv6 Features

Last, but not least, a number of upcoming features such as channel bonding and IPv6 are critical to scale existing DOCSIS services. The channel bonding functionality is fundamental to enable cable operators to offer HSD services with maximum download speeds above 30 Mbps—the practical limit of a traditional DOCSIS 1.x/2.0 cable modem using a single downstream channel. IPv6 support will also become critical due to IPv4 IP address exhaustion as the number of subscribers and their connected devices grows, as new services are offered, and as users demand seamless connectivity in their home network environments.

The ability of an M-CMTS Core product to scale for existing DOCSIS services depends more on the product implementation than whether it is evolved from an existing I-CMTS or is an entirely new product. Cable operators should carefully consider both options, but be aware of the fact that some I-CMTS products in the market today will be able to scale to meet the requirements listed above, while others will not. For example, an I-CMTS with a passive backplane can grow with the service requirements without introducing fixed bottlenecks in the system. An I-CMTS with a passive backplane architecture can therefore scale without a forklift upgrade. Scalability is achieved by simply adding the newest-generation DOCSIS interfaces, backhaul interfaces and/or routing engine as needed.

Support New DOCSIS Services Planned in the Next 3-5 years

If the ability to scale deployment of existing services is an important driver for the migration towards an M-CMTS architecture, it is equally critical that the new architecture be able to support new types of services that the operator plans to deploy over the next 3-5 years. The range of future services currently being investigated by cable operators is quite wide and depends on the particular broadband competitive landscape in each particular region and the business focus of each operator. In this section, we'll focus on a subset of such services where the choice of an M-CMTS Core is particularly important.

One of the major threats to the cable operator's core business is the move of large Telcos into the residential entertainment video market with both on-demand and scheduled broadcast video content offered over their broadband infrastructure. While the IPTV technology is still at the initial deployment phase, few have any doubts that it will open significant business opportunities and change the way residential users experience video services. We believe it will be critical for the long term success of cable operators to make their migration to the M-CMTS architecture, keeping open the possibility to deliver Video over DOCSIS (VDOC).

A detailed description of possible VDOC architectures is outside the scope of this document, but it is relevant to highlight key aspects that impact the ability of an M-CMTS Core to support a VDOC service offering. In addition to the characteristics described earlier (such as Enhanced Forwarding Capacity, Control Plane Scalability and availability of features such as IPv6 and channel bonding), a number of fundamental requirements to enable VDOC services exist, including:

- Enhanced IP Multicast support (including Multicast QoS)
- Intelligent Multicast Routing to optimize video bandwidth allocation
- Ability to enable fast channel change
- Integration with Entitlement Servers to provide secure authentication
- Advanced Admission Control capabilities

When choosing an M-CMTS Core that can enable VDOC services, it is critical to carefully evaluate the product's DOCSIS, IP and Video features. If VDOC will expand the cable operator's service offering, it will be critical to continue to introduce innovation in the communications segment with services such as video telephony, video conferencing or integration with mobile services.

We will group these new services under the name of Next Generation Communications Services (NG-COMMs). Among the DOCSIS technologies that will enable NG-COMMs services, PacketCable Multimedia (PCMM) and PacketCable 2.0 are becoming increasingly important.

- Some of the functions being defined by PacketCable 2.0 include the capability for enhanced residential VoIP services and video telephony, feature integration across service platforms, and mobility services.
- PCMM opens up the possibility for cable operators to launch very innovative services such as multi-party game playing and video-conferencing by extending the QoS mechanism beyond the VoIP application. Figure 6 presents the CableLabs PCMM architecture diagram.

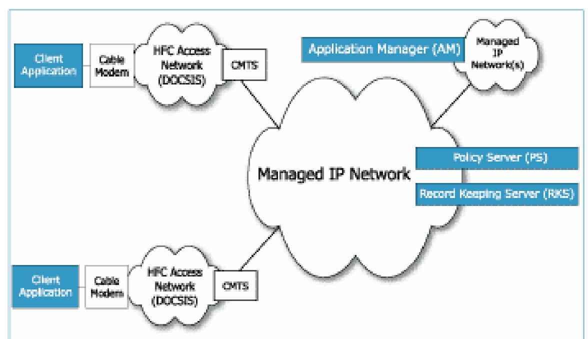

Figure 6. PacketCable Multimedia Diagram

A key aspect for the support of PCMM and PacketCable 2.0-based services is the ability of the M-CMTS Core to integrate well with the Application Manager, the Policy Server, and the Record Keeping Server. In the case where the M-CMTS Core is an evolution of an I-CMTS (Option A) that has already been through PCMM integration steps, this will translate into much greater service velocity than the case where such application servers need to be integrated with a new CMTS Core.

To summarize, some may argue that a newly developed M-CMTS Core platform (Option B) can support new services better than an I-CMTS since the I-CMTS was originally designed for today's services. However, before drawing conclusions, it is important to recognize the flexibility of an I-CMTS platform and its ability to evolve to meet future service requirements. In addition, the optimization of the underlying DOCSIS infrastructure and the integration of the I-CMTS with advanced service architectures that has likely occurred during field deployment might prove a key asset to move faster towards the launch of such services.

Capital Budget Impact (Next 5 Years)

Perhaps the most pressing business driver for a cable operator to migrate to an M-CMTS architecture is the ability to add DOCSIS capacity and meet increasing bandwidth requirements at a fraction of today's cost. In this section, we consider how well the two migration options considered in this paper meet this.

It is important to realize that, with the separation of the downstream PHY into an external EQAM, the M-CMTS architecture has simultaneously achieved two key objectives. These objectives are:

- Separate the downstream PHY from the MAC functionality
- Separate the downstream PHY from the upstream PHY

The first objective is critical to leverage existing video EQAM technology for DOCSIS traffic, and therefore, reduce the cost of DOCSIS downstream ports. The second objective is as critical since it removes inefficiencies introduced by the fixed ratio between downstream and upstream ports in today's I-CMTS.

The impact on CapEx achieved through the first objective is easy to grasp by comparing the cost between a downstream port in a video EQAM and on a DOCSIS CMTS, as well as the expected decline.

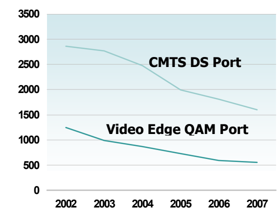

Figure 7. CMTS DS Port vs. EQAM Port

To better understand the impact of the separation of downstream and upstream PHY functions, let's consider the following capacity sizing exercise: a cable operator wants to deploy 10 Mbps downstream and 2 Mbps upstream service in a network with 500,000 households passed. It projects to reach 10% take rate (50,000 subs). Assuming 200 kbps downstream and 40 kbps upstream are allocated to each subscriber, 10 Gbps in the downstream direction (50,000 * 200 kbps) and 2 Gbps in the upstream direction (50,000 * 40kbps) will be required. This translates into 250 downstream and 200 upstream ports (assuming 40 Mbps per downstream and 10 Mbps per upstream port).

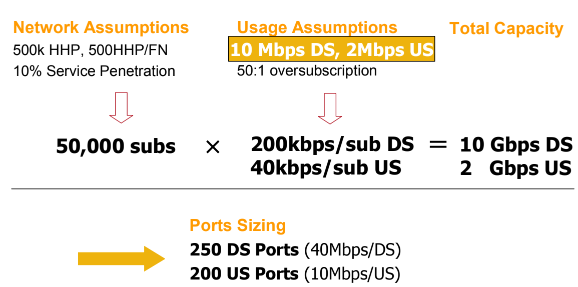

Figure 8. Example 1

Today downstream and upstream ports come packaged in one line card with 4 to 6 upstream ports per downstream. While the earlier traffic assumptions show the M-CMTS requirement to be approximately 1.25 downstream ports for every upstream port, this means that upstream capacity that is not required will be purchased, leaving the upstream ports under-utilized. For example, if the line cards come with 5 downstream and 20 upstream ports, 50 of such line cards will be required to meet the downstream capacity requirement, leaving the operator with 1000 upstream ports while the traffic requirement indicates that 200 are sufficient (20% upstream port utilization). If the line cards have 2 downstream and 12 upstream ports, this will mean that in order to purchase the 250 downstream ports required, 125 line card with 1200 upstream ports will be purchases (16.7% utilization). In reality, the HFC topology puts additional requirements beyond pure capacity in order to connect each fiber node to the CMTS. Although we've chosen to ignore this factor for simplicity, this doesn't impact the validity of the argument presented.

By breaking the fixed ratio between upstream and downstream ports, an M-CMTS architecture allows upstream and downstream capacity to be purchased independently so that CapEx can be optimized. In the earlier example, only ten 5x20 line cards are required to meet the upstream bandwidth requirements providing 50 downstream ports. The additional 200 downstream ports can be purchased independently by purchasing a MAC engine for the M-CMTS Core and a number of EQAMs for the downstream PHY interfaces.

As the example shows, significant savings can be achieved due to the separation of upstream and downstream ports in the M-CMTS architecture. There is

10 Mbps DS - 2Mbps US

Ports Sizing
250 DS Ports (40Mbps/DS)
200 US Ports (10Mbps/US)

Today's I-CMTS	M-CMTS Core + Edge QAM	
5x20 LC	5x20 LC + 24DS M-CMTS Edge QAM	
50 Cards	10 Cards	9 MAC Blades + 9 Edge QAM
250 DS	50 DS	200 DS
1000 US (20% utilization)	500 US (100% utilization)	-
7 Chassis	2 Chassis	

Figure 9. Example 2

a key aspect of the above analysis that is important to stress. Due to the asymmetry of today's high-speed Internet service offering when compared with the available capacity on a DOCSIS line card, cable operators most likely have excess upstream capacity they've already purchased (upstream port under-utilization described above) for an I-CMTS deployment. With regard to this aspect, migrating an existing I-CMTS into a M-CMTS Core (Option A) clearly shows strong advantages from a CapEx standpoint. No additional upstream ports are required in addition to what is already available until those ports are fully utilized and the system is balanced with the addition of the required downstream capacity. For this reason, we expect Option A to provide the minimum impact to CapEx since it fully leverages the assets already purchased.

M-CMTS Core Conclusions

This section analyzes the key selection criteria a cable operator should consider when making a decision on the M-CMTS Core during a transition to an M-CMTS architecture. Table 2 summarizes important points.

Although a newly developed M-CMTS Core might look attractive when it comes to support for future DOCSIS services, this potential advantage must be balanced against the key downsides of the approach:

- Risk of impacting existing services if M-CMTS Core is unproven in DOCSIS networks
- Risk of delaying migration to M-CMTS if back-end integration effort is extensive
- Larger upfront CapEx negatively impacting cash flow

Choosing to migrate an existing I-CMTS into an M-CMTS Core might be the safest approach. The key is to choose an I-CMTS that can scale to support existing and new DOCSIS services, while keeping the changes to turn it into an M-CMTS Core at a minimum.

5.0 M-CMTS EQAM CHOICES

The EQAM is a key component of the M-CMTS architecture and is an important factor in the success of an M-CMTS deployment strategy. This section contains an analysis of the key evaluation criteria to be considered in choosing which EQAM option best meets the operator's objectives for deploying an M-CMTS solution.

As explained earlier, cable operators have two fundamental options for deploying an EQAM as part of an M-CMTS solution. These two options are:

- Option A: Install a new standalone EQAM
- Option B: Repurpose the I-CMTS into an EQAM

Criteria	**Critical factors**
Support for existing DOCSIS features and services	• Support all Layer 2 and Layer 3 features required for today's HSD and VoIP services • Integrate seamlessly in today's operations structure • Have the design stability of a field- proven implementation
Scalability for existing DOCSIS services	• Provide 5-10 Gbps forwarding capability • Support 50k-80k DOCSIS CPEs • Support Channel Bonding and IPv6
Support for new DOCSIS features and services (next 3-5 years)	• VDOC • NG Communications such as video telephony or video conferencing
Capital budget impact (next 5 years)	• Fully leverage existing capital expenditure in I-CMTS deployment

Table 2. Key Evaluation Criteria for M-CMTS Core

The key criteria for assessing these EQAM options are:

- Cost optimization
- Implementation risk
- Vendor and product choice
- Density
- Scalability

- Flexibility
- Cost optimization

Cost is a primary consideration in assessing the options. The EQAM product must be designed for cost optimization; not just at initial deployment, but throughout the product life cycle. Product upgrades such as adding QAM channel capacity, supporting new features, implementing redundancy, and replacing components must be cost-effective.

EQAM costs are directly related to the power and signal quality of the outputs. Compared to the power and signal quality of today's video EQAMs, the DRFI specification raises the bar significantly. It will be a challenge for vendors to meet the DRFI requirements at the same price per QAM as video EQAMs. However, the leaders in the video EQAM market will be able to leverage a high-volume cost structure to optimize their M-CMTS EQAM costs, and thus, will foster a competitive pricing environment.

By standardizing the interfaces and functionality of the EQAM, the M-CMTS specifications drive a common, well-defined feature set, which enables additional cost optimization in EQAM products. With a standalone EQAM, features like redundancy, scalability, density, manageability and reliability can all be targeted for the EQAM application.

<u>Implementation Risk</u>

This section highlights some of the potential risks of the EQAM options that might impact the decision on which M-CMTS migration strategy to choose. For example, an I-CMTS runs on a large, complex code base that supports features which are superfluous to an EQAM. Repurposing an I-CMTS into an EQAM is therefore a more difficult task, and thus riskier than implementing an application-specific standalone EQAM. Considering that an existing I-CMTS that is repurposed to perform as an M-CMTS EQAM is also likely to perform as an M-CMTS upstream receiver, the product development effort will be more complex (and thus riskier) than developing a standalone EQAM.

Implementing both the EQAM and upstream receiver functions on the I-CMTS also adds risk to achieving interoperability with the new M-CMTS Core. This risk is even greater if a pre-standard version of UEPI is implemented by the M-CMTS Core and EQAM.

<u>Vendor and Product Choice</u>

The M-CMTS specification opens the door to the DOCSIS downstream channel market for EQAM vendors and significantly increases the total addressable market for their EQAM products (possibly by as much as twofold or higher). For this reason, it is expected that several EQAM vendors will take this opportunity to introduce standalone M-CMTS EQAM products.

Cable operators who choose Option A (standalone EQAM) will have the freedom to choose EQAM products from multiple vendors, including the EQAM market leaders which to date have not been able to offer products for the DOCSIS network. The freedom to choose from multiple vendors ultimately fosters competition and aligns with the M-CMTS objective of using cost-effective EQAMs.

Enabling multiple vendors to offer M-CMTS EQAM products also results in greater product choice for cable operators. A standalone EQAM can be implemented with either a stackable or chassis form factor, giving operators the freedom to choose the best form factor for their particular environment. A stackable EQAM is likely more suitable in small hub sites where the growth can be accommodated with a few EQAMs, whereas a chassis-based EQAM is likely more suitable in large headend and hub sites where high growth is anticipated. The form factor of the existing I-CMTS may not be suitable where the operator intends to deploy the M-CMTS EQAMs. Of course, operator preference dictates which form factor is most suitable. Having EQAM products available in both form factors is necessary to fulfill market requirements worldwide.

Density

The dwindling available rack space in a typical cable headend or hub site and the ongoing migration from analog to digital services is escalating the priority of EQAM density (in terms of QAM channels per rack unit) as a key evaluation criterion for EQAM products. Recalling that a primary objective of the M-CMTS architecture is to enable significant growth in DOCSIS downstream channel capacity, it is imperative that M-CMTS EQAM products provide the density required to support the projected growth within the cable operator's limited rack space.

As an example of the potential growth of DOCSIS downstream channel capacity, consider the case of a cable operator with a DOCSIS 3.0 network who wants to offer a HSD service with 100 Mbps maximum download speed. Such a service would have to be delivered on a four-channel bonding group (using 64 QAM Annex B modulation). This essentially matches the number of QAM channels typically installed for VoD service today. As new high-bandwidth services such VDOC are offered, the number of DOCSIS QAM channels required can quickly surpass the number of narrowcast video QAM channels installed.

The earlier example illustrates the point that M-CMTS EQAMs must at least match the density of video EQAMs. Video EQAMs typically support up to 24 QAM channels per rack unit today, and are expected to at least double in density within the next 2-3 years. Ideally, M-CMTS EQAMs will achieve much greater density to allow for the dramatic growth in downstream channel capacity envisioned by the M-CMTS and DOCSIS 3.0 initiatives.

Standalone M-CMTS EQAM products will likely be based on or modeled after video EQAMs, and thus can be expected to offer similar QAM channel density. Stackable EQAMs are likely to achieve greater QAM channel density than chassis-based EQAMs because chassis-based products typically incorporate high availability features that are not normally found in products with a stackable form factor.

When repurposing the I-CMTS into an M-CMTS EQAM, the number of QAM channels supported by the downstream PHY card is the critical factor in determining the overall QAM channel density of the I-CMTS chassis. The overall QAM density of the I-CMTS must be comparable to the density of standalone EQAM products. Otherwise, the cable operator will essentially be wasting rack space by using the I-CMTS as an M-CMTS EQAM. The overall QAM density of the I-CMTS is particularly critical if the new M-CMTS Core will be co-located with the I-CMTS, and thus require additional rack space.

Scalability

There are numerous factors to consider in assessing an M-CMTS EQAM's scalability, including:

- Output power level and signal quality requirements (dictated by DRFI)
- Number of RF ports (may be limited by connector type)
- Number of QAM channels per RF port (block upconversion ratio)
- Number of line card slots (in modular designs)
- Backplane capacity per line card slot
- Total backplane bandwidth
- Power supply capacity
- Air flow/cooling capacity
- Number and speed of wide area network (WAN) interface ports
- Processor performance (for data and control plane traffic)
- System memory

It is beyond the scope of this paper to address each of the scalability factors listed above. However, some general observations about the scalability of the different types of M-CMTS EQAM products being considered are provided.

Due to their small form factor, stackable EQAMs are not scalable. Today's stackable video EQAMs are typically able to support up to 24 QAM channels, with the expectation of doubling in scale in the next 2-3 years. When additional QAM channels are required, additional stackable units are deployed in the network.

Scalability is important for chassis-based EQAM products, since these products occupy more rack space and have longer life spans than stackable products. Since the useful life of a chassis-based product is typically at least 5-10 years, the product must be designed not only for today's services and technology, but also in anticipation of tomorrow's services and technology. For example, a chassis-based EQAM should be able to provide the total system power required to support full loading with current-generation QAM modules, as well as full loading with next-generation QAM modules which will support twice the number of QAM channels. All of the scalability factors listed above should be considered in selecting a chassis-based EQAM.

The best way to meet the scalability requirements of an M-CMTS solution in a chassis-based EQAM is to deploy a new purpose-built chassis that addresses the scalability factors for current and future DOCSIS services over the duration of the product life cycle. A chassis originally designed as an I-CMTS will likely have significant shortcomings in several of the scalability factors listed above.

Flexibility

The M-CMTS architects took great care in defining the M-CMTS EQAM so that it could be seamlessly integrated with a traditional video EQAM, and thus, enable both DOCSIS and digital video applications to share a common EQAM device. As a result, a new product concept emerged from the M-CMTS architecture, the Universal EQAM. A Universal EQAM is able to support traditional digital video services, as well as DOCSIS services, and enables cable operators to make more efficient use of network resources.

Both of the M-CMTS EQAM options being considered in this paper can be implemented as a Universal EQAM by integrating the video EQAM functionality with the M-CMTS EQAM functionality. However, the choice of M-CMTS EQAM options does have implications on the flexibility achieved with a Universal EQAM.

There are three critical factors that should be considered in assessing the overall flexibility of a Universal EQAM:

- Range of Services Supported
- Resource Allocation Model
- Granularity of Resource Sharing

Range of Services Supported

Since the interfaces and functionality of the M-CMTS EQAM is standardized, all Universal EQAM products that are compliant with M-CMTS specifications will be able to support the full range of standard DOCSIS services. However, since the video EQAM is not standardized, each video service requires unique features on the EQAM. Therefore, the range of video services a Universal EQAM can support will depend on the product's feature set. For maximum flexibility, a Universal EQAM must support video on demand (VoD), digital simulcast, and switched digital video (SDV), as these services are critical to the cable operator's success.

Resource Allocation Model

The ability to dynamically allocate shared EQAM resources to services as needed greatly enhances the flexibility of a Universal EQAM. Dynamic resource allocation enables the most efficient use of EQAM resources, especially when the EQAM resources are shared by a wide range of services. To support dynamic resource allocation, a Universal EQAM must interoperate with a resource management system. Without dynamic resource allocation, EQAM resources must be statically allocated to each service, potentially stranding resources and affecting service availability.

Granularity of Resource Sharing

There are multiple levels of granularity for sharing Universal EQAM resources among services. Services may be able to share a Universal EQAM per chassis, per RF port, per QAM channel, or even share the same QAM channel. The flexibility of a Universal EQAM is maximized when each QAM channel can be shared by multiple services.

Considering the above factors for maximizing the flexibility of a Universal EQAM, and the fact that existing video EQAM products can not be directly integrated with an I-CMTS EQAM, the best approach for achieving flexibility is to deploy a standalone Universal EQAM which integrates the M-CMTS EQAM functions into a video EQAM product capable of supporting VoD, digital simulcast and SDV.

M-CMTS EQAM Conclusions

This section analyzes the key evaluation criteria to be considered when assessing the options for deploying an EQAM as part of an M-CMTS solution. Table 3 summarizes the key criteria and the critical factors discussed in the analysis.

Criteria	Critical factors
Cost optimization	• Designed to optimize cost • High-volume cost structure • Standard EQAM feature set
Implementation risk	• I-CMTS code base • Design complexity • UEPI standardization • Interoperability
Vendor and product choice)	• Vendor choice with standalone EQAM • Standalone EQAM can be stackable or chassis-based
Density	• High priority feature • Match video EQAM density
Scalability	• Critical for chassis-based EQAMs • Must scale with services and technology • Requires purpose-built EQAM
Flexibility	• Universal EQAM • All video and DOCSIS services • Dynamic resource allocation • Granularity of resource sharing

Table 3. Key Evaluation Criteria for M-CMTS EQAM

In summary, for each of the key criteria identified, deploying a standalone M-CMTS EQAM has significant benefits compared to repurposing the I-CMTS into an EQAM. The cost and scalability advantages of a standalone EQAM is of critical importance, considering that these are primary objectives of the M-CMTS architecture. In addition, offering greater choice and flexibility with lower risk are important operational benefits of the standalone EQAM option.

CONCLUSION

The Modular CMTS (M-CMTS) series of specifications published by CableLabs introduces powerful new tools to enhance the architecture of the CMTS, the fundamental building block of all IP-based services over the cable network. The high-level objectives of the M-CMTS architecture are to improve the flexibility of the current DOCSIS network, enable independent scaling of downstream and upstream channel capacity, and improve the overall efficiency of the cable network with best-of-breed solutions. The solutions allow cable operators to prepare their DOCSIS infrastructure to enter the new broadband era.

This paper presents various approaches cable operators can follow to migrate from the current DOCSIS architecture to an M-CMTS architecture. The high-level tasks for executing the migration are described for each approach, followed by a discussion of the options for implementing an M-CMTS Core and EQAM as part of an M-CMTS solution.

A thoughtful analysis of the key criteria for assessing each M-CMTS migration strategy leads to the recommendation that cable operators should follow a two-phased approach that builds upon the current CMTS installed base to implement the M-CMTS Core function, and utilizes new standalone Edge QAMs in the first phase and external upstream receivers in the second phase. This strategy meets the objectives of the M-CMTS architecture, while providing cable operators greater choice and flexibility with lower risk.

UNDERSTANDING THE IMPACT OF BITTORRENT ON CABLE NETWORKS

Jim Martin
Clemson University

Terry Shaw
CableLabs

Abstract

BitTorrent is a peer-to-peer (P2P) protocol for sharing media (audio, video and software) files that accounts for 18% of traffic on cable high-speed data networks. In order to better understand the network dynamics of traffic associated with this protocol, we analyzed roughly 1 Gbytes of network trace data from BitTorrent clients running in three different environments. Based on this data, we performed a simulation-based analysis of a DOCSIS®-based cable network involving a varying number of simulated BitTorrent users and web browsing users. Our results suggest that as few as 10 BitTorrent clients can double the access delay experienced by other non-BitTorrent users.

INTRODUCTION

BitTorrent is a peer-to-peer (P2P) protocol that provides scalable file sharing capabilities. As in other P2P systems, BitTorrent provides an overlay network that runs over the Internet. BitTorrent is unique from other P2P protocols in several aspects. The most notable of these are: 1) the use of swarming; and 2) the use of incentives to prevent free-riding. BitTorrent, frequently the dominant application in a network, poses a dilemma for cable network service providers. In general, P2P applications are broadband applications that contribute to the demand for high speed broadband access. This usage is driven by younger Internet users. In a recent study, it was estimated that about 58% of American teenagers with broadband access have downloaded audio and/or video files from the Internet using P2P networks[1]. However, BitTorrent can consume tremendous amounts of bandwidth in both the upstream and downstream directions. The primary incentive used by BitTorrent to prevent free-riding is that the protocol mandates that a file-downloader must make some content available for upload. One study from several cable operators found that 18% of all traffic is BitTorrent [WSJ05]. The government and the cable industry are addressing the evolving economic and policy issues. In this study we explore the impact that BitTorrent users can have on a cable network. We provide insight to this issue by presenting an analysis of BitTorrent involving live network measurement and simulation[2].

In addition to file sharing, P2P has been used for grid computation [Entropia], storage [DKK+01], web caching [IRD02] and directory services [RFH+01, SMK+01]. BitTorrent, however, is designed specifically for file distribution. The algorithms contained in BitTorrent were heavily influenced by previous P2P protocols and applications. We briefly overview two such protocols, Gnutella and Kazaa, before presenting BitTorrent.

---

[1] In a recent study, it was estimated that about 66% of American teenagers have downloaded audio and/or video files from the Internet. It was also found that 30% of teenagers admit to currently using P2P networks and 28% admit to using P2P networks in the past [PIP05].

[2] The observations that we report represent preliminary work. We expect our findings and conclusions to evolve as the study progresses.

Gnutella

Gnutella is an overlay network superimposed on the Internet [MAR+04,IVK]. Unlike the centralized Napster approach, Gnutella is decentralized because meta-data (i.e., the indexes of where files are located) is distributed throughout the network. At initialization time, a Gnutella client accesses a host server to obtain a list of peers currently in the network. After this step, peers learn about the network and available content using Ping and Pong messages. Once a client has identified a file to download, it sends the request to the network using a query message. Query replies will contain information needed for the client to download the file.

Kazaa

Few details of the Kazaa protocol are available. Several 'black box' measurement studies do however provide some insight [LKR04,LRW03]. There are two types of nodes: Super Nodes (SN) and Ordinary Nodes (ON). Each ON must be told or must learn the address of at least one SN. SNs maintain lists of other SNs that can be thousands large. Each SN maintains records for files located at the ONs in its domain. ONs contact their SN on initialize, providing metadata that describes the content located at the ON. When an ON requests a file, the SN looks locally and also propagates the request to other SNs in the overlay network. To improve performance, files can be stored in fragments on multiple nodes and they can be transferred in parallel.

BitTorrent

While other P2P systems provide distributed object location in addition to file downloading, BitTorrent assumes that the user has a URL of the file to be downloaded. BitTorrent is therefore a protocol for distributing files. Files are broken up into chunks and downloaded in pieces. This concept is borrowed from [RB02] whose authors experimented with several parallel access schemes to download a file from multiple servers concurrently. BitTorrent files are represented by a torrent which is a meta-information file that identifies the chunks associated with a file and a tracker address that knows about the file. Downloaders, also referred to as clients or peers, download the torrent that is available on a web site. The download engages the BitTorrent client software that was installed on its machine which contacts the tracker that was identified in the torrent to obtain a number of IP address/port pairs corresponding to peers that have one or more chunks of the file. The client periodically reports its state to the tracker (about every 30 minutes). The client attempts to maintain connections with at least 20 peers (referred to as the peer set). If it can not, it asks the tracker for additional peers. When a client first connects to a tracker, the tracker attempts to reverse connect to it. If it succeeds, the client is added to the list of peers for the torrent maintained by the tracker. This ensures that the peers in the client's peer set will accept incoming connections. At least one of the peers must contain all of the file. Once a peer has downloaded the entire file, it becomes a seed and provides chunks of the file to lechers for free. The system is self-scalable in that as the number of downloaders increase, the number of nodes that can function as seeds also increase.

BitTorrent breaks a file into chunks, or pieces. To enhance performance, BitTorrent further breaks pieces into small (typically 16 Kbytes) sub-pieces (sometimes called blocks) and maintains a strategy of having at least 5 concurrent block requests pipelined to a given peer. The client contains a piece selection algorithm that decides the next piece to obtain. The algorithm has three modes of operation: starting, downloading, or finishing. When the client starts it has nothing to upload. The algorithm randomly selects the initial pieces. Once one or more pieces are received, the algorithm switches to a rarest first algorithm. Based on information learned from its peers, the client chooses to download the chunk that is least available. This ensures that peers will have pieces that they will need. Once the peer has downloaded all but the last few pieces, it might experience significant delay if the remaining pieces are located at nodes that are available only over low-bandwidth paths. To avoid this, it sends requests for these sub-pieces to all peers.

A peer actively controls its downstream bandwidth consumption by granting upload requests to peers from which it is currently downloading a piece. In essence, it provides a 'tit-for-tat' bartering approach to distributed resource management. By choosing to upload to peers that provide the best download rates, freeloaders will perform poorly. A peer effectively 'chokes' another peer by denying download requests from a peer. Snubbing is when a peer is choked by all peers with which it was formerly downloading.

A BitTorrent peer uploads only to a limited number of peers (usually four). Peers monitor download rates and decide who to choke over time intervals of every ten seconds (note that some implementations might use a different rechoke period). When the next piece is to be downloaded, the peer that has been performing the worst will likely be replaced by a higher performing peer. To probe the network in search of better performing peers, BitTorrent supports an 'optimistic unchoke.' Every third rechoke period, it unchokes a peer by uploading to it regardless of its download rate. Therefore, a BitTorrent client is allowed to upload to a maximum of five peers, four that implement tit-for-tat bartering to maximize downstream throughput and a fifth that probes the network searching for better performing peers.

BitTorrent is unique from other P2P protocols in several aspects: 1) the use of swarming; 2) the use of incentives to prevent free-riding. BitTorrent, frequently the dominant application in a network, poses a dilemma for cable network service providers. It is a broadband application which contributes to the demand for high-speed broadband access. However, BitTorrent can consume tremendous amounts of bandwidth. In this study, we explore the impact that BitTorrent users can have on a cable network. We provide insight to this issue by presenting an analysis of BitTorrent involving live network measurement and simulation. This paper is organized as follows. First, we overview related performance studies of peer-to-peer systems. We then describe our modeling efforts: a measurement study designed to characterize bandwidth consumption of BitTorrent clients and a simulation analysis designed to assess the impact of BitTorrent on an HFC cable network. We end the paper with conclusions and our future directions.

RELATED WORK

There has been a great deal of prior research on P2P protocols and systems. The majority of the studies have focused either on P2P deployments in the Internet [SGG02, NCR+03, RIP01] or on evaluating protocol issues [RFK+01, SMK+01,GKT03]. There have been several analytic models proposed. Several notable efforts have been based on queuing theory [GFJ+03] and on fluid flow models [CN03,CNR03].

The existing studies of BitTorrent indicate that the protocol is indeed scalable and robust [PGE+05, SGP04, IUB+04, BLS04, BHP05]. In [IUB+04] the authors examined the log file from the tracker associated with the popular Redhat Linux 9 distribution torrent. They observed that peers continue to participate in the BitTorrent network as a seed for an average of 6.5 hours after the entire file has been downloaded. The authors also showed that the average download rate is over 500 Kbps and that nodes do consume symmetric amounts of bandwidth. The authors observed that 81% of all file downloads were incomplete. Of these aborted downloads, 90% had retrieved less than 10% of the file. The average download rate of the 19% of the sessions that completed was 1.3 Mbps which is larger than the average download rate of all sessions at 500 Kbps. The authors studied an individual peer by running an instrumented client. The client downloads a 1.7 Gbyte file in 4500 seconds. During the download period, the client interacted with roughly 40 peers. Upon completion, the client remained on line as a seed for 13 additional hours. During this period, roughly 90 leechers were served (but only 4 at a time while the others are choked). They found that the volume of traffic in the upstream and downstream directions at the client was correlated, but throughputs were not correlated. 85% of the file was sent by only 20 peers including 8 seeds that provided 40% of the file. This set of 20 peers were not a part of the initial peer list provided by the tracker which suggests that to enhance performance NAT must not prevent nodes from connecting with a downloader.

In [PGE05], the authors examined the access of 60000 files and find an average download bandwidth of 240 Kbps and that only 17% of the nodes stay on line for one or more hours after they complete the download. The authors suggested that there might be a shortage of seeds in typical usage scenarios. The authors assessed performance based on availability, integrity, responsiveness to flash crowds and on download performance. To obtain an availability data point, they tracked the activity associated with a popular file from its initial offering until when it died (i.e., when it is no longer available). During the file's three month lifetime, 90155 peers downloaded at least one piece. Of the 53883 peers that were not behind firewalls, only 17% have an uptime longer than 1 hour after they finish downloading. By trying to upload corrupt files but failing due to the moderator's inspection, the authors deduce the BitTorrent network is relatively pollution free. By correlating system activity with the introduction of a new file (The Lord of The Rings III) that caused a flash crowd effect, the authors observe that the system remained stable.

In [QS04], the authors applied fluid flow modeling techniques to BitTorrent to obtain a model that predicts average number of seeds, downloaders, and download time.

To assess the effectiveness of BitTorrent, the authors in [BP05] defined link utilization as the aggregate capacity of all nodes participating in the overlay network. Fairness is defined in terms of the relative number of pieces served by a node to the number downloaded. Another aspect of fairness is if all nodes, in particular seeds, are equally utilized. Finally, a diversity measure assesses how effectively pieces are distributed throughout the overlay network.

In [BLS04] the authors analyzed BitTorrent activity over a four month period involving thousands of torrents. They observed a mean and median file size of 760 and 600 Mbytes, respectively. The mean and median session duration was 13.25 and 8.4 hours, respectively. The authors plan on developing a BitTorrent application for distributing large network traces to the research community.

In [SGP04] the authors instrumented the standard BitTorrent client and monitored the activity associated with a popular tracker. They observed an average download rate of 200 Kbps among all clients that were not behind a firewall. The authors proposed a streaming content delivery network based on a BitTorrent-like protocol. They pointed out that two areas that need improvement for this application are better protection from freeloaders and better support for nodes behind firewalls.

MODELING BITTORRENT

To the best of our knowledge no one has characterized in detail the workload generated by a BitTorrent client. As this was required for our simulations, we performed an analysis involving network traces of real BitTorrent clients.

BitTorrent Traces

We obtained three sets of traces. They are summarized in Table 1. All involved downloading the same torrent (a 4.3 GByte file actively traded on the BitTorrent network) using version 4.04 of the BitTorrent client at [BT]. The client machine was a WindowsXP machine. According to the statistics at the tracker, the torrent consistently had hundreds of downloaders and tens of peers. In our experiments, firewalls prevented remote peers from initiating TCP connections with the client. The BitTorrent client allows the user to specify the maximum bandwidth to be consumed by the host. We set this to the maximum value of 45 Mbps.

We refer to the three data sets as Set1, Set2 and Set3 respectively. Two of the sets (Set1 and Set2) are from a machine on Clemson University's campus and the other set (Set3) is from a machine located at the author's home. To help manage P2P traffic generated by students, Clemson deployed a Packeteer bandwidth management device that limits individual TCP sessions (except for web traffic) to a maximum rate of 64 Kbps[PACK]. Set1 traces were subject to this control. For Set2, we configured the Packeteer device so that packets associated with the IP address of the BitTorrent client had high priority and were not subject to rate limits. The speed of the access link connecting Clemson's gigabit Ethernet campus network to the Internet is 100 Mbps. Set 3 traces were from a machine connected to the Internet by DSL with service rates of 1.5 Mbps downstream and 256 Kbps upstream.

To calibrate our trace and analysis methodology we performed a set of TCP transfers between the DSL connected machine and the host located on campus. We transferred 500 Kbytes ten times from the DSL connected machine to the campus machine (i.e., upstream from the perspective of the BitTorrent client). Then, we transferred 500 Kbytes ten times from campus to the DSL connected machine. For these traces, and for all traces described in this report, all TCP packets sent and received by the DSL connected host were captured using tcpdump[TCPD]. Table 2 summarizes the calibration results. The top row shows the average of a set of performance measures collected from each of the 10 upstream transfers and the bottom row shows the downstream results. The columns labeled 'AvgCx BW', 'Avg TCP RTT', and 'Avg TCP LR' indicate the average TCP throughput, round trip time and loss rate, respectively, experienced by the 10 TCP connections. The column labeled 'Bottleneck Link Speed Estimate' examines the interarrival times between ACK packets that arrive at the client and, based on a packet-pair algorithm, estimates the bottleneck link speed in the upstream direction [KLD+04]. Since this flow will not be classified as P2P traffic by the Packeteer device, we expect that the bottleneck link speed in both the upstream and downstream directions to be the DSL access link speeds. The results from Table 2 confirm this although the estimated bottleneck bandwidth is roughly 20% lower than the actual. This is due to framing overhead that is not included in the bandwidth estimate. The TCP throughput was slightly less than the bandwidth estimate. This is because of additional overhead due to TCP. The uncongested RTT over the path based on a 56 byte Ping probe is roughly .11 seconds. The mean TCP RTT of 1.16 seconds observed in the upstream connections along with no packet losses implies that sustained queuing exists over the path, presumably at the DSL router's upstream queue. Further analysis shows that the TCP congestion window stays at a maximum value of 32 Kbytes that was set by the receiver's advertised window. The transmission time of 32 Kbytes over a 256 Kbps link is roughly 1 second making the observed RTT reasonable. The algorithm that we use to obtain the average TCP RTT is based on previous work [MNR03]. It requires a send side trace and consequently we were not able to obtain a RTT result for downstream connections.

The first two traces of each set were taken while the client was downloading the torrent. The third trace was taken once the client became a seed. While downloading, data flowed in both directions over most of the TCP connections. Our analysis tools operate only on the data sent in the upstream direction. While seeding, data flowed primarily in the upstream direction. Each trace in Set1 and Set2 contains about 400000 packets comprising roughly 60 unique connections of which 20 to 30 were active throughout the lifetime of the trace. The Set3 traces each contained about 60000 packets comprising slightly more (40-60) active connections while downloading and significantly more once the client became a peer (180 connections). More than half of the connections in all sets transferred a very small amount of data (presumably BitTorrent messages).

Table 3 summarizes the observed behavior of upstream TCP transfers contained in the Set1 traces. The columns labeled 'Avg Aggregate US BW' and 'Avg Aggregate DS BW' show the aggregate bandwidth consumed in the upstream and downstream direction, respectively, by the client during the trace. As with all traces described in this report, more bandwidth is consumed in the upstream direction than in the downstream. The column labeled 'US/DS BW ratio' captures this. The mean ratio for the 4,5,7,8 traces was 2.5 and increased to 24 once the client became a peer (i.e., traces 6 and 9).

The Set1 results clearly show the impact of the Packeteer device. Without the rate limits imposed by Packeteer we would expect to see a bottleneck bandwidth estimate in the 1 to 5 Mbps range (i.e., typical broadband access downstream service rates). As expected, with the Packeteer device engaged, the estimated bottleneck bandwidth was close to the configured service rate of 64 Kbps. The transmission time of a 1500 byte packet over a 64 Kbps link is .188 seconds making the observed TCP RTT reasonable. The average upstream and downstream bandwidth consumption for each trace is significant, especially once the client becomes a seed. The average upstream TCP throughput varied widely although never exceeded 50,000 bps.

The last four columns in Table 3 indicate the number of concurrent transfers. We define an active connection to be a connection that sends data in a one-way direction at a rate that is greater than a threshold level. We used thresholds of 10 Kbps, 40 Kbps and 100 Kbps. With Set1, since TCP connections are limited to 64 Kbps, we did not see any connections with a transfer rate greater than 100 Kbps. We see many low rate (10 Kbps) flows and smaller number of flows that consume at least 40 Kbps. Once the client became a peer, we observed 28 concurrent 10 Kbps upstream flows and 22 concurrent 40 Kbps flows.

In the Set2 traces we see an order of magnitude increase in the bottleneck link speed which confirms that the Packeteer service rates are no longer in place. The bottleneck link speed ranged from a minimum of 256 Kbps to 10 Mbps. The RTT is reasonable except for Trace 22 which had an RTT of .802 seconds. It's unclear why this value is so high. We see roughly 5 to 7 active high speed flows in the upstream direction. In the downstream direction, the number of active high speed flows was very low. Presumably, peers were not capable (or willing) of sending at very high rates.

The Set3 traces also seem reasonable although the upstream bottleneck link estimates were off by almost 50%. The error associated with packet-pair estimate increases if the senders are not constantly sending. The average upstream TCP throughput is low because the connection was bidirectional and but there were periods of time when data was flowing in only one direction. This "idle" time is reflected in the TCP throughput estimate. The TCP RTT matches the value we saw in the calibration tests (refer to Table 2). Once the client becomes a seed the RTT becomes significantly higher. However, the TCP loss rate gets lower. Further investigation is required to explain this.

Simulation Analysis

In prior work we implemented a simulation model of a DOCSIS network using the NS-2 simulation tool [MW05, ns2]. We have extended this model to support a simple BitTorrent traffic model. Figure 1 illustrates the simulated network. There are 200 cable modems (labeled CM-1 through CM-n) that share one downstream channel and one upstream channel. Fifty CMs are configured with the client side of a web traffic model. A simulated web user at the CM generates requests to web servers (nodes S-1 through S-x) following the model described in [BC98]. Figure 2 identifies the DOCSIS network configuration parameters and the web model settings that were used in the experiments. Refer to [MW05] for further details of the model.

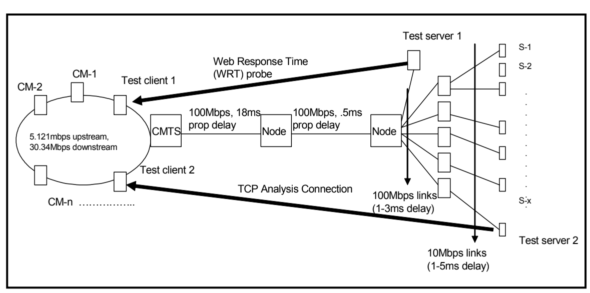

Figure 1. Simulation Network Model

Model Parameters
Upstream bandwidth 5.12Mbps
Preamble 80 bits
Downstream bandwidth 30.34Mbps
4 ticks per minislot
Default map time: 2 milliseconds (80 minislots per map)
Fragmentation Off, MAP_LOOKAHEAD = 255 slots
Concatenation ON
Backoff Start: 8 slots, Backoff stop: 128 slots
12 contention slots (minimum), 3 management slots

Web Traffic Model Parameters
Inter-page: pareto model, mean 10 and shape 2
Objects/page: pareto model, mean 3 and shape 1.5
Inter-object: pareto model, mean .5 and shape 1.5
Object size: pareto model, mean 12 (segments) shape 1.2

Figure 2. DOCSIS Network and Web Traffic Model Simulation Parameters

A varying number of the CMs are configured with the client side of a BitTorrent traffic model. In the model one or more TCP connections transfer data in the upstream direction and one or more additional TCP connections transfer data in the downstream direction. An on/off traffic source is attached to each TCP sender associated with the client. For US traffic, the senders are located at the designated CM and interact with the TCP sinks located at the servers (S-1 through S-x). For DS traffic, TCP senders are located at the servers and the sinks are attached to the CMs. A BitTorrent source periodically transfers a large amount of data to the server. The message size is exponentially distributed with a mean of 1 Gbyte. This value is based on a measurement study that found an average torrent size of about 800 Mbytes [BLS04]. The off time of the traffic source is exponentially distributed with a mean of 2 seconds. To model *n* concurrent TCP connections at a BitTorrent client, we create n TCP flows between the server and the client but reduce the amount of data sent by each traffic source by a factor of *1/n.*

We ran four sets of simulation experiments. In each set, the number of simulated BitTorrent clients in the HFC network was increased from 0 to 50 in increments of 10 over 6 different runs. The four sets are summarized as follows:

Symmetric: For each BitTorrent client, there was one upstream and one downstream flow.

Downstream asymmetric: The total number of upstream BitTorrent file transfers is limited to five streams. We increase the ratio of downstream BitTorrent traffic to upstream traffic by increasing the number of downstream connections from 5 to 50.

Upstream asymmetric: The total number of downstream BitTorrent file transfers is limited to five streams while the number of upstream connections increases from 5 to 50.

Symmetric and parallel connections: For each BitTorrent client, instead of one upstream flow and one downstream flow, we adjust the traffic source accordingly and use four concurrent connections upstream and four concurrent connections downstream.

To calibrate our simulation model with observed behavior, we compare the symmetric simulation results (summarized in Table 6) with the results from measured traces 7 and 8 shown in Table 5. The average TCP connection throughput, the loss rates and the aggregate US bandwidth are similar. The most significant difference is that the aggregate downstream bandwidth is much higher in the simulation. This is due in part because our BitTorrent model does not capture the application 'tit-for-tat' dynamics, nor does the model support user specified rate limits that are available on most BitTorrent clients. However, the more significant reason for the asymmetry is because our simulation network model assumes peers are connected to the network with very high speed links (10Mbps).

Summary of simulation results

For each of the four sets of simulation experiments, we obtain a number of performance measures that fall into two categories.

- Those that assess performance of CMs that are not running BitTorrent. We ran several network performance monitor applications on the Test client 1 and Test server 1 nodes. These nodes were not running the BitTorrent or the web traffic generators. For brevity we report only the results of a CBR flow between these two nodes. The objective of the monitor is to simulate a best effort VoIP flow and to monitor the UDP packet jitter, latency and loss. The CBR source is attached to the Test client 1 node and is configured to send a 350 byte message every .05 seconds (i.e., 56Kbps). Figures 3, 4 and 5 visualize these results for each of the four sets of experiments.

- Those that assess DOCSIS network dynamics. Figure 9 visualizes the mean collision rates reported by the CM's. Figures 6, 7 and 8 provide insight to how the CMs obtained upstream bandwidth. Figure 6 plots the percentage that contention requests were used for all packets (IP packets, fragments and management messages) that were sent upstream. Figure 7 plots the percentage of upstream packets that were sent in a concatenated frame. Figure 8 plots the average number of packets per concatenated frame.

Figures 3 through 9 plot the performance measures. Each figure contains four plots representing the results of the experiments. The following observations apply to all four sets:

- As expected, the load on the upstream channel increased linearly as the number of BitTorrent clients grew. Over a 5.12 Mbps upstream channel, as few as 10 BitTorrent clients caused the access delay experienced by other users to double. When 10 BitTorrent clients were active, the US channel was 55% utilized. This result was observed with no service rate limits and also when 384 Kbps upstream service rates were imposed.

- As the BitTorrent load increased, the majority of bandwidth requests was accomplished using concatenation with an average of 3.5 IP packets (primarily TCP acknowledgement packets) inserted into each concatenated frame. If concatenation was disabled, most bandwidth requests would rely on piggybacking which is not as efficient and would lead to significantly higher mean access delays.

We note the following differences between the experiments:

- The UDP packet jitter and latency is double for the symmetric, parallel connections compared to the symmetric, single connection runs (i.e., Figures 3 and 4). At the same time, the collision rate is lower by roughly 20% once the network becomes congested (Figure 9). More study is required to explain this behavior. However, the result suggests that BitTorrent's use of concurrent connections might have a greater detrimental impact on real-time traffic such as voice and video.

- All statistics confirm that performance is significantly worse when BitTorrent traffic is bi-directional rather than asymmetric.

- The downstream asymmetric configuration delivers upstream data by concatenation less often (by 50%) than the other sets. This is because the US utilization was lower than in the symmetric experiments.

CONCLUSIONS

The objective of this study was to gain insight in how varying levels of BitTorrent traffic on an HFC network impacts other CMs and on the dynamics of a DOCSIS network. Different parameters and model assumptions will change the results. In particular, the availability and locality of the torrent, peer behaviors, and network loads all determine the impact on the network and subsequently, on subscriber's perceived performance. Our goal was to obtain sufficient data permitting us to build a simulation model of a BitTorrent traffic generator. Our simulation model is most similar to the behavior observed in a residential network. We configured four concurrent upstream and downstream flows that consumed symmetric levels of bandwidth.

Based on simulation we have shown that a small number of BitTorrent users can impact other users, even when CMs are provisioned with low upstream service rates. Cable companies can no longer rely on static rules of thumb when provisioning an HFC network. This issue becomes more urgent if the provider plans to reliably support best effort VoIP and video applications. The provider must monitor network performance and adapt the network as needed.

There are many directions for future work including further development of our BitTorrent traffic model and subsequent analysis in cable environments. The impact that high-bandwidth applications such as BitTorrent have on other subscribers becomes more problematic as service rates increase. Therefore, we are developing adaptive bandwidth management techniques that implement fairness policies through a combination of performance and monetary incentives.

REFERENCES

[BC98] P. Barford, M. Crovella, "Generating Representative Web Workloads for Network and Server Performance Evaluation", ACM SIGMETRICS '98, July, 1998.

[BHP05] A Bharambe,C. Herley, V. Padmanabhan, "Analyzing and Improving BitTorrent Performance", Microsoft Research Technical Report MSR-TR-2005-03, Feb 2005.

[BITPR] BitTorrent Documentation: Protocol available at http://www.bittorrent.com/protocol.html

[BLS04] A Bellissimo, B. Levine, P. Shenoy, "Exploring the Use of BitTorrent as the Basis for a Large Trace Repository", Technical report 04-41, Department of Computer Science, University of Amherst, June 2004.

[BT] BitTorrent client and tracker software, http://www.bittorrent.com

[BTSPE] The BitTorrent Specification: http://wiki.theory.org/BitTorrentSpecification

[CN03] F. Clevenot, P. Nain, "A Simple Fluid Model for the Analysis of SQUIRREL", Technical Report, Inria RR-4911,2003.

[CNR03] F. Clevenot, P. Nain, K. Ross, "Stochastic Fluid Models for Cache Clusters", Technical Report, Inria RR-4815, 2003.

[COH03] B.Cohen, "Incentives Build Robustness in BitTorrent", Workshop on Economics of Peer-to-Peer Systems, Berkeley CA, May 2004. Available at http://www.bittorrent.com/bittorrentecon.pdf

[DKK+01] F. Dabek, M. Kaasheok, D. Karger, R. Morris, I Stoica, "Wide-area Cooperative Storage with CFS", SOSP01, Oct 2001.

[ENTRO]Entropia, http://www.entropia.com

[GFJ+03] Z. Ge, D. Figueiredo, S. Jaiswal, J. Kurose, D. Towsley, "Modeling Peer-to-Peer File Sharing Systems", IEEE Infocom 2003.

[GKT02] L. Gao, J. Kurose, D. Towsley, "Efficient Schemes for Broadcasting Popular Videos", Multimedia Systems, Vol 8, No4, July 2002.

[IRD02] S. Iyer, A. Rowstron, P. Druschel, "Squirrel: A Decentralized Peer-to-Peer Web Cache", PODC02, 2002.

[IUB+04] M Izal, et. Al., "Dissecting BitTorrent: Five Months in a Torrent's Lifetime",Passive and Active Measurements, April 2004.

[IVK] I. Ivkovic, "Improving Gnutella Protocol: Protocol Analysis and Research Proposals", unpublished report available at http://www9.limewire.com/download/ ivkovic_paper.pdf.

[KBB+04] T. Karagiannis, A. Broidi, N. Brownlee, K. Claffy, M. Faloutsos, "Is P2P Dying or just Hiding?", IEEE Globecom, November 2004.

[KLD+04] S. Kang, X. Liu, M. Dai, D. Loguinov," Packet-pair Bandwidth Estimation: Stochastic Analysis of a Single Congested Node", ICNP-04, 2004.

[LKR04] J. Liang, R. Kumar, K. Ross, "Understanding KaZaA", unpublished, available at http://cis.poly.edu/ ~ross/papers/.

[LRW03] N. Leibowitz, M. Ripeanu, A. Wierzbicki, "Deconstructing the KaZaA Network", Proceedings of the Third IEEE Workshop on Internet Applications (WIAPP'03), June 2003.

[Mark02] E. Markatos, "Tracing a large-scale Peer to Peer System: an hour in the life of Gnutella", CCGrid 2002.

[MNR03] J. Martin, A. Nilsson, I. Rhee, "Delay-based Congestion Avoidance for TCP", IEEE/ACM Transactions on Networking, 11(3), 356-369 (2003).

[MW05] J. Martin, J. Westall, "Validating an 'ns' Simulation Model of the DOCSIS Protocol", under review, available at: http://people.clemson.edu/~jmarty/papers/docsis-model.pdf.

[NCR+03] T. Ng, Y. Chu, S. Rao, K. Sripanidkulchai, H. Zhang, "Measurement-Based Optimization Techniques for Bandwidth-DemandingPeer-to-Peer Systems, Infocom, 2003.

[NS2] The Network Simulator. Available at : http://www-mash.cs.Berkeley.EDU/ns/.

[PACK] Packeteer bandwidth management devices, http://www.packeteer.com

[PGE+05] J. Pouwelse, P. Garbacki, D. Epema, H. Sips, "The BitTorrent P2P File-Sharing System: Measurements and Analysis", International Workshop on Peer-To-Peer Systems, Feb 2005.

[PIP05] A. Lenhart, M. Madden, "Teen Content Creators and Consumers", Pew Internet & American Life Project, November 2, 2005.

[QS04] D. Qiu, R. Srikant, "Modeling and Performance Analysis of BitTorrent-like Peer-to-Peer Networks", Proceedings of the

2004 conference on Applications, technologies, architectures, and protocols for computer communications, 2004.

[RB02] P. Rodriguez, E. Biersack, "Dynamic Parallel Access to Replicated Content in the Internet", IEEE/ACM Transactions on Networking, Vol 10, no 4, Aug 2002.

[RFH+01] S. Ratnasamy, P. Francis, M. Handley, R. Karp, S. Shenker, "A Scalable Content Addressable Network", ACM SIGCOMM 2001

[RFI02] M. Ripeanu, I. Foster, A. Iamnitchi, "Mapping the Gnutella Network: Properties of Large-scale Peer-to-Peer Systems and Implications for System Design", IEEE Internet Computing Journal, 6(1), 2002.

[RIP01] M. Ripeanu, Peer-to-Peer Architecture Case Study: Gnutella Network, 2001.

[SGG02] S. Saroiu, P. Gummadi, S. Gibble, "A Measurement Study of Peer-to-Peer File Sharing Systems", Proceedings of Multimedia Computing and Networking 2002.

[SMK+01] I Stoica, R. Morris, D. Karger, M. Kaashoek, H. Balakrishman, "Chord: A Scalable Peer-to-Peer Lookup Prototocol for Internet Applications", ACM SIGCOMM 2001.

[SGP04] K. Skevik,V. Goebel, T. Plagemann, "Analysis of BitTorrent and its use for the Design of a P2P based Streaming Protocol for a Hybrid CDN", Technical report, Department of Informatics, University of OSLO, 2004.

[TCPD] tcpdump program available at http://www.tcpdump.org

[YV04] X. Yang, G. Veciana, "Service Capacity of Peer-to-Peer Networks", IEEE Infocom

[VY03] G. Veciana, X. Yang, "Fairness, Incentives and Performance in Peer-to-Peer Networks", Proceedings of the Forty-first Annual Allerton Conference on Communication, Control and Computing, Oct 2003.

[WSJ05] P. Grant, J. Drucker, "Phone, Cable Firms Rein in Consumers's Internet Use", The Wall Street Journal, October 21, 2005.

Set	Comment	US bandwidth	DS bandwidth	# Traces in Set
1	Campus, but rate limited	64 Kbps	64 Kbps	3
2	Campus, high speed	100 Mbps	100 Mbps	6
3	DSL home network	256 Kbps	1.5 Mbps	3

Table 1. Summary of the three data sets

Trace Number	US Bottleneck Link Speed Estimate	AvgCx BW (bps)	Avg TCP RTT	Avg Loss Rate (%)
Calibration US	214484	210161	1.164	0
Calibration DS	1221614	1123897	*	0

Table 2. Residential Network Calibration Results

Trace Number	US Bottle-neck Link Speed Estimate	AvgCx BW (bps)	Avg TCP RTT	Avg Loss Rate (%)	Avg aggregate US BW (bps)	Avg Aggregate DS BW (bps)	US/DS BW ratio	US > 100Kbps, 40Kbps, 10Kbps flows	DS > 100Kbps, 40Kbps, 10Kbps flows
1	54873	18655	.554	2.43	326079	129626	2.5	0, 2.8, 7.5	0, .2, 4.9
2	55523	36705	.407	.431	462271	312500	1.48	0, 6, 10.1	0, 2.6, 8.3
3 (seed)	63655	47509	.447	.96	1379270	30947	45	0, 22, 28.5	0, 0, 0

Table 3. Set1 BitTorrent Client Measurement Results

Trace Number	US Bottle-neck Link Speed Estimate	AvgCx BW (bps)	Avg TCP RTT	Avg Loss Rate (%)	Avg Aggregate US BW (bps)	Avg Aggregate DS BW (bps)	US/DS BW ratio	US > 100Kbps, 10Kbps flows	DS > 100Kbps, 10K bps flows
4	3948836	191714	.366	2.6	2180090	1292200	1.69	5.8,10.0	4.9,8.4
5	3809208	211886	.361	2.9	2525240	885066	2.85	5.8, 12.3	3,10.8
6 (seed)	2425737	203143	.802	1.1	2671200	111122	24	6.4, 16.1	0,3.7

Table 4. Set2 BitTorrent Client Measurement Results

Trace Number	US Bottle-neck Link Speed Estimate	AvgCx BW (bps)	Avg TCP RTT	Avg Loss Rate (%)	Avg Aggregate US BW (bps)	Avg Aggregate DS BW (bps)	US/DS BW ratio	US > 100Kbps, 40Kbps, 10Kbps flows	DS > 100Kbps, 40Kbps, 10Kbps flows
7	123797	23347	.989	4.35	189623	161176	1.18	.1, 1.2, 3.8	.1, 1.2, 3.8
8	111115	17840	1.11	4.8	201635	156835	1.28	0, 1.2,4.7	0, .8, 4.4
9 (seed)	140998	20869	3.93	.013	214107	5851	37	0, .4, 6.2	0, 0, 0

Table 5. Set3 BitTorrent Client Measurement Results

Set Identifier	**Avg TCP Cx BW**	**Avg TCP Cx RTT**	**Avg TCP Cx LR**	**Aggregate US BW**	**Aggregate DS BW**	**US/DS BW ratio**	**Number of concurrent US and DS flows**
Symmetric	60000	.05	4%	160 Kbps	1.4 Mbps	.11	1
Symmetric and parallel cxs	17500	.2	6.5%	185 Kbps	1.3 Mbps	.14	4

Table 6. Summary of the Symmetric Simulations

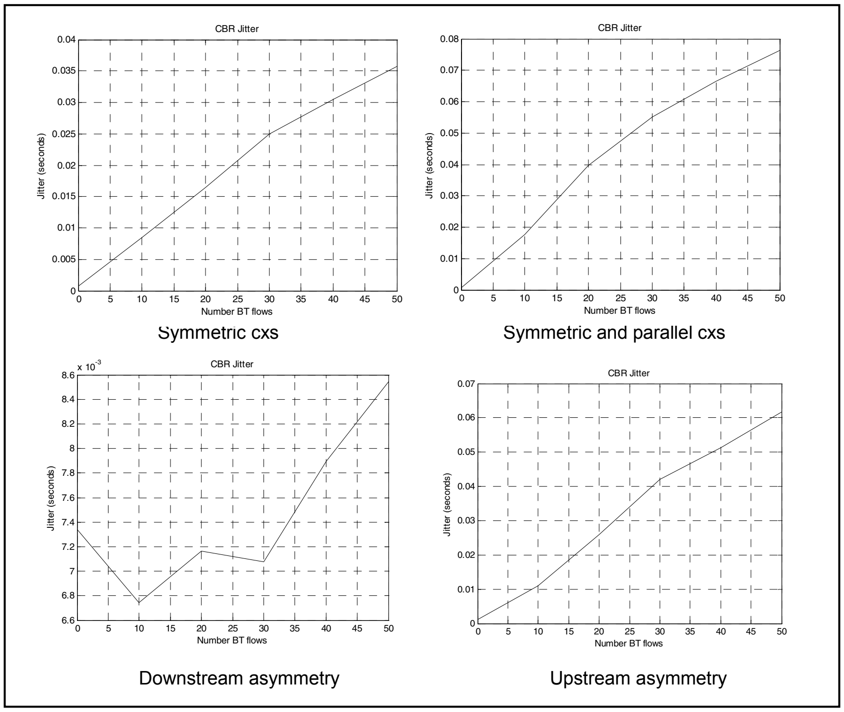

Figure 3. CBR Jitter as BitTorrent Load Increases

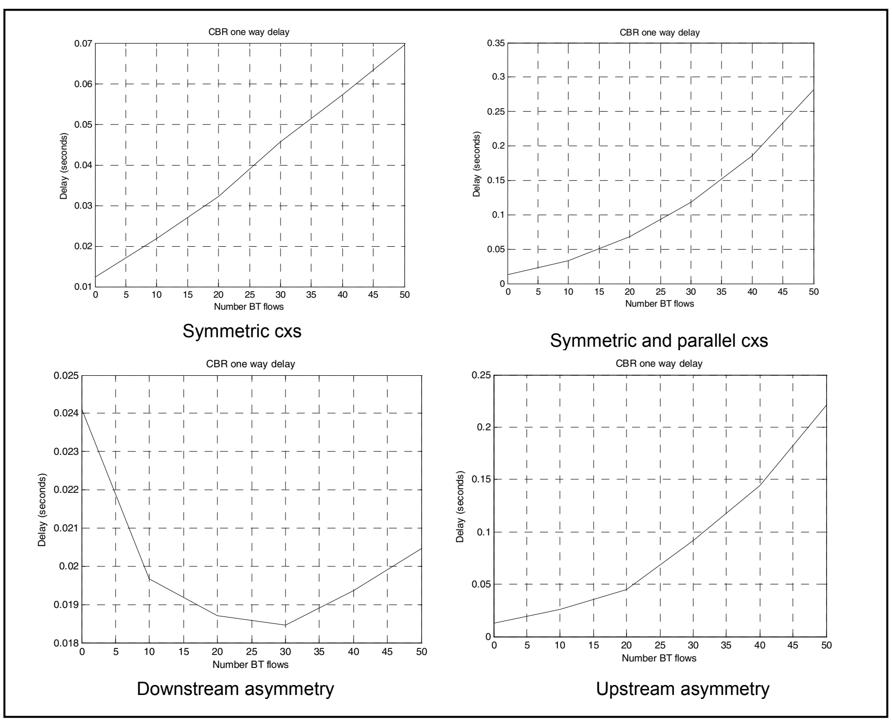

Figure 4. CBR One-Way Latency as BitTorrent Load Increases

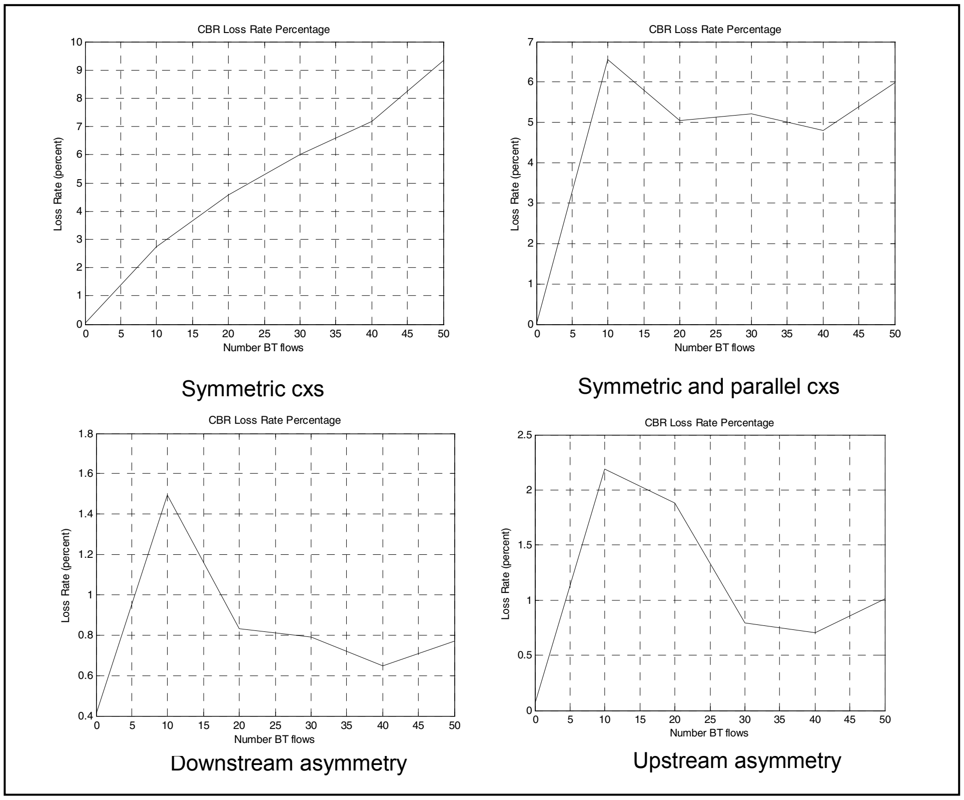

Figure 5. CBR Loss as BitTorrent Load Increases

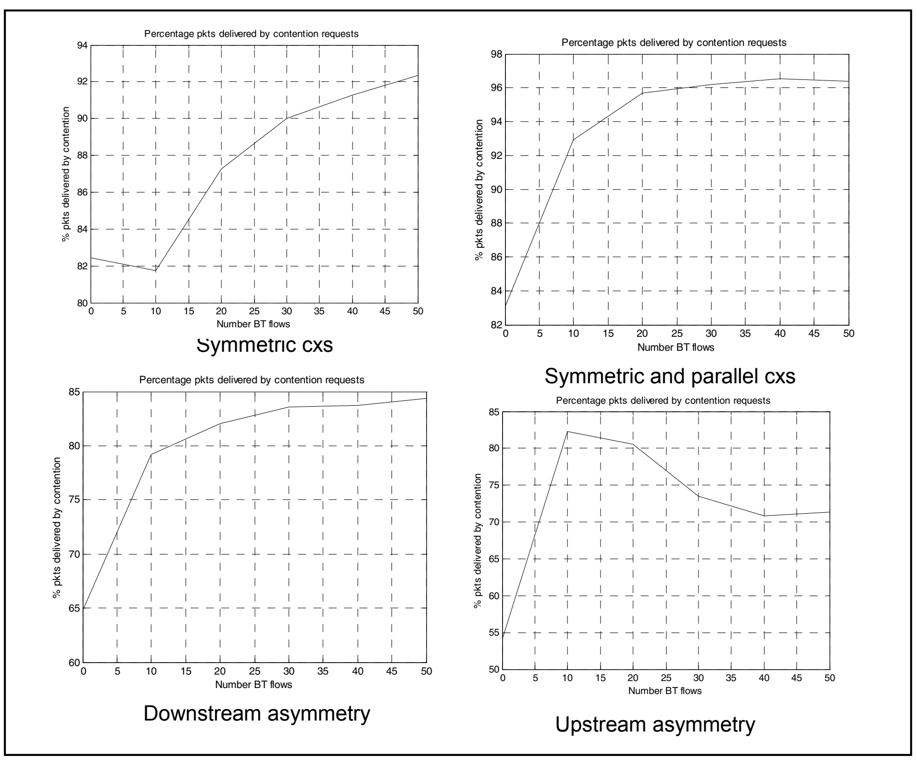

Figure 6. Percent Contention Requests as BitTorrent Load Increases

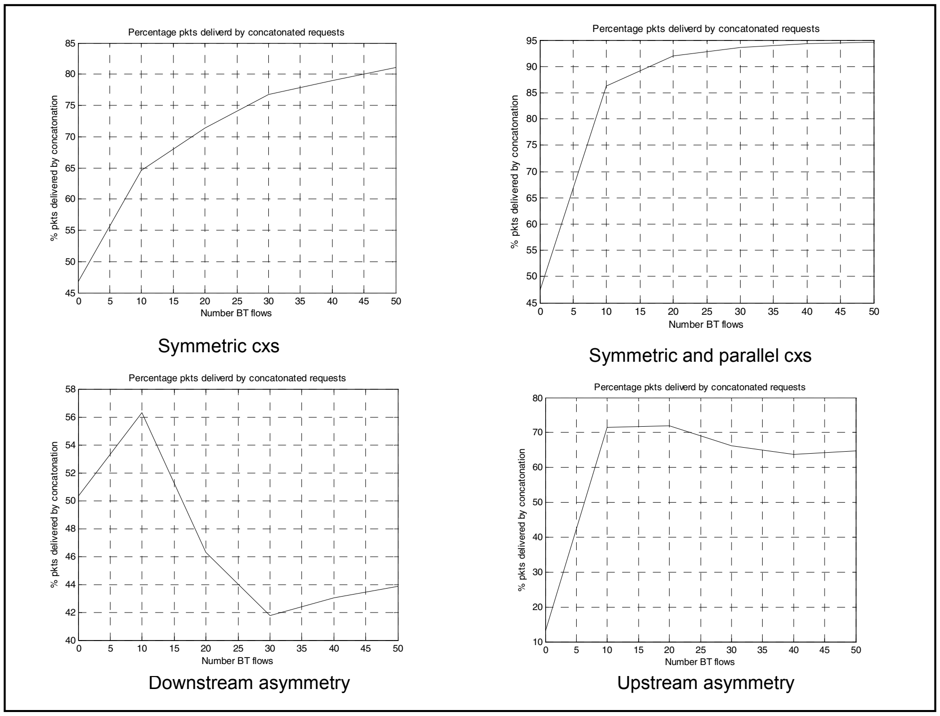

Figure 7. Percent Concatenation Requests as BT Load Increases

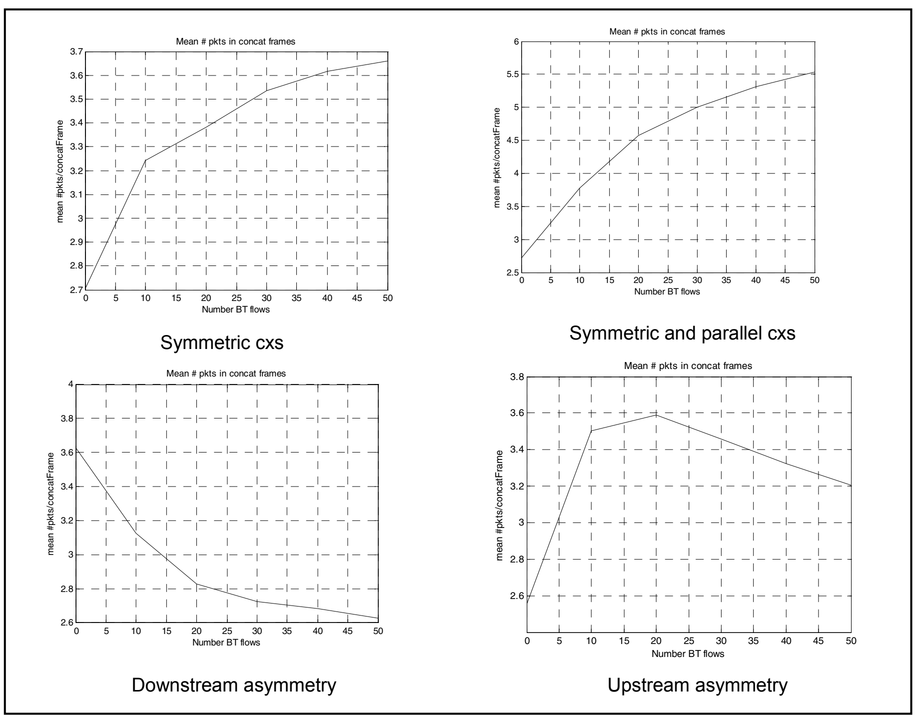

Figure 8. Mean Number of Pkts in Concatenated Frames as BitTorrent Load Increases

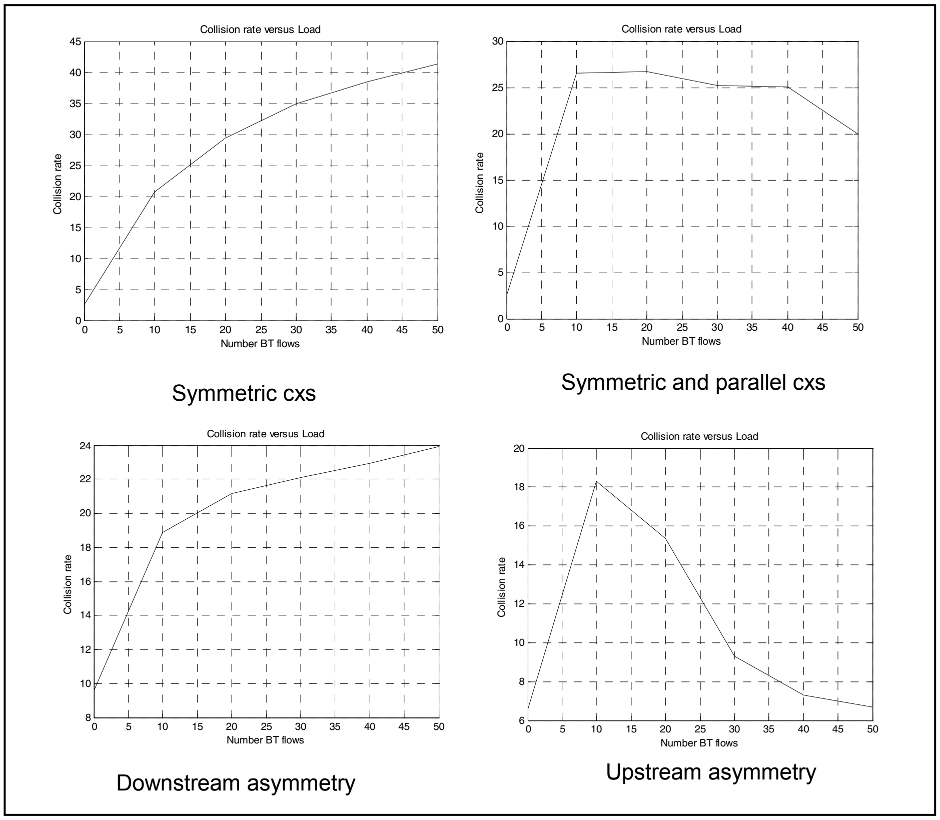

Figure 9. Collision rate as BitTorrent load increases

ISBN 0-940272-01-6; 0-940272-08-3; 0-940272-10-5; 0-940272-11-3; 0-940272-12-1; 0-940272-14-8; 0-940272-15-6; 0-940272-16-4; 0-940272-18-0; 0-940272-19-9; 0-940272-20-2; 0-940272-21-0; 0-940272-22-9; 0-940272-23-7; 0-940272-24-5; 0-940272-25-3; 0-940272-26-1; 0-940272-27-X; 0-940272-28-8; 0-940272-29-6; 0-940272-32-6; 0-940272-33-4; 0-940272-34-2; 0-940272-35-0; 0-940272-36-9; 0-940272-37-7; 0-940272-38-5; 0-940272-39-3; 0-940272-40-7; 0-940272-41-5; 0-940272-42-3; 0-940272-43-1; 0-940272-44-X; 0-940272-45-8; 0-940272-46-6; 0-940272-47-4; 0-940272-48-2; 0-940272-49-0; 0-940272-50-4; 0-940272-51-2; 0-940272-52-0; 0-940272-53-9; 0-940272-54-7