THE COMPLETE
TECHNICAL PAPER PROCEEDINGS

FROM:

NCTA
Technical
Papers

A 10-YEAR RESIDENTIAL BANDWIDTH DEMAND FORECAST AND IMPLICATIONS FOR DELIVERY NETWORKS

Randy Nash
Motorola

Abstract

Today's Hybrid Fiber Coax (HFC) cable plants provide a cost effective infrastructure for delivery of video, voice and Internet data services to residential customers. Consumers expect Multi-System Operators (MSOs) to continue providing more content and innovative services at competitive prices. What types of service changes are likely to occur over the next 10 years? What are the implications of increased bandwidth consumption to the delivery network? Does Hybrid Fiber Coax (HFC) meet the capacity and future service needs and when do we need Fiber-to-the-Home (FTTH)? What are the value propositions for average consumers? Will there be a "killer app?" How many broadcast channels do we really need, can afford to deliver and can pay to produce? Who wants High-Definition Television (HDTV)?

The answers to these related questions depend not only on technological advances that change economics, but also on consumer expectations and adaptation to new technology. This paper takes a macroscopic approach to estimating demand changes in the following categories:

- *Analog broadcast*
- *Digital video broadcast*
- *HDTV broadcast*
- *Video-on-Demand (VOD)*
- *IP data and services*
- *IP Telephony*

A team of Motorola product engineers, applied researchers and marketing staff developed a forecast to better understand requirements and timing for next generation products. As with all attempts to predict the future, there are dimensions of uncertainty. However, the alternative is to march ahead without any vision of future needs. Industry analyst predictions were useful, but none put the pieces together from a bandwidth perspective. MSOs in North America, Europe and Latin America were consulted for plans and expectations. The forecast was updated and conclusions are presented here.

The bandwidth forecast categories are aggregated to determine RF bandwidth required on HFC nodes. This leads to a possible scenario for HFC node segmentations over the next 10 years.

INTRODUCTION

This paper examines the current state of CATV and broadband services to the home and offers an opinion on changes considered likely to occur over the next 10 years. The primary focus is on the center of the mass market in North America, with some commentary on other regions of the world.

Consumer Adoption of Technology

Before jumping into the forecast, it is helpful to consider, some historical perspectives. We have had over 50 years to develop the TV viewing habits which are engrained in our society. Probe Research[1] analyzed the household adoption of various consumer devices and services. The VCR took about 6 years from introduction to early adopter takeoff and then 10 years before adoption by the late majority. Premium cable service took 12 years from early adopter takeoff to late majority. Some products and services take a generation to become mainstream. The bellwether group is younger consumers, because they are always first to accept new technology. For example, the shortest adoption cycles in the Probe Research[1] study were for two generations of game machines which were targeted at children and teenagers. These each took 4 years from early adopter takeoff to late majority adoption.

FORECASTS BY BANDWIDTH CATEGORY

The bandwidth needs for the major services delivered over cable were considered and a most likely forecast scenario for the next 10 years was developed

Note that, this paper is deliberately colloquial in it's use of the term "bandwidth." In some cases it literally means usage of RF spectrum in the CATV plant. In others, "digital transport capacity" or "bit rate" might be more precise.

Analog Broadcast

At the close of year 2001 76% of US cable plants had a bandwidth of 750 MHz or greater and typically provided 80 channels of analog TV[2]. Analog broadcast service is projected to remain largely unchanged over the next 10 years. This is largely because there are approximately 267 million TV sets in the US, most of which are cable ready and still will be working in 10 years. Without a driving need to reclaim bandwidth, MSOs are likely to add new services in spectrum above 550 MHz.

Although not necessarily to reclaim spectrum, some migration of analog programming to digital is expected by year 2006. Analog scrambled programs are good candidates, since digital cable provides better security and obviates the need for two access control systems. In a typical network, 14 analog channels are expected to migrate to digital, reducing the analog spectrum required from about 500 MHz to 400 MHz by year 2006.

Digital video broadcast

Digital video cable is currently in the mass adoption phase. By the end of year 2001 approximately 18 million digital cable set-top boxes were in use by US subscribers. Systems offering digital video had between 3 and 12 digital Quadrature Amplitude Modulation (QAM) carriers; typical systems had 10. Each QAM carrier provides about 8 to 10 video programs, resulting in 80 to 100 digital channels on a typical system. With VOD services emerging (discussed later) and the cable modems competing for consumers' free time, it is hard to see a case for the addition of many new broadcast channels. A net gain of 36 additional programs is expected over the next 10 years. Some may be unique cable content and some will be digital cable versions of local DTV broadcast that operators opt to carry. Four additional QAM carriers will be added to cable plants, bringing the total number carrying Standard Definition (SD) content (including migration from analog broadcast) to 16 by year 2011.

Interactive TV still coming?

With an installed base of 18 million digital subscribers the business case for interactive TV applications is becoming interesting, despite past industry disappointments. Complicated business relationships and conflicting priorities may have been part of the problem, but the size of the addressable market was a likely factor. Consumer interest in interactive TV exists as evidenced by a growing number of consumers interacting with TV programs using PCs. Beyond our 10-year period, 2-way interactive broadcast content could be the salvation of broadcast services in a world that is otherwise evolving to total content-on-demand.

From a bandwidth perspective, set-top boxes share the 5 MHz – 40 MHz return spectrum with Data Over Cable System Interface Specification (DOCSIS) cable modems. DOCSIS upstream channels typically avoid the band below 20 MHz to avoid ingress noise and to avoid conflicts with other uses (e.g. set-tops and plant monitoring equipment). Most first-generation set-tops are capable of providing real time interactivity using a 256 kbps return modem in the 8 to 15 MHz band. Currently, most set-top return modems are used only for collecting monthly PPV purchases, but they have the potential to do much more. For example using a simple contention protocol with a conservative 10% channel loading, each set-top return channel supports 54 interactive packets per second (payload = 53 bytes). With a channel spacing of 192 kHz, 36 channels are available in the 8 MHz to 15 MHz band, yielding a system capacity of 2,000 interactive packets per second. Both subscriber equipment and sufficient bandwidth are available for real time TV-based interactivity. Successful interactive based businesses are expected to develop involving MSOs and third party application providers.

HDTV broadcast

Broadcasters debuted HDTV in the US in 1998. Consumer adoption of HDTV sets has been slow. Broadcasters offer relatively little of their programming in HD format. HDTV enthusiasts in the US have bought approximately two million HDTVs. The take off point for mass adoption of consumer devices has historically occurred between 10% and 25% adoption[1]. Given there are about 106 million US TV households, HDTV penetration would have to reach at least 11 million units before adoption take off is anticipated. The first barrier is the value proposition. Is the picture quality worth the price of an HDTV? How many consumers viewing a 42 inch screen at normal distances can discern the improvement in HDTV quality relative to DVD or MPEG 2 SD quality? The second issue is scarcity of HD content. DBS and cable MSOs are beginning to address this issue, but only a small fraction of programs are in HD format. D-VHS digital videotape players have recently been introduced at a retail cost of $2000.

In March of 2001, the FCC clarified “Must Carry” rules for digital TV and limited the MSO liability to one “primary video” for each local station. Liabilities aside, some operators have begun to offer HDTV to the enthusiasts on their systems. Rather than carry one HD program using VSB in a 6 MHz channel, most operators are choosing to provide two programs in a single 256 QAM carrier. Since most HDTVs have separate tuners and displays, the cable and broadcast modulation formats can be accommodated by having a decoder for each.

Although mass adoption is not expected in the next few years, increasing amounts of HD content are expected to appear on cable. With two HDTV programs per carrier, systems will begin carrying 4 to 6 HDTV video programs this year. By year 2011, 16 HDTV channels are forecast. The bulk of content will continue to be delivered in SD resolution and channels may even alternate between HD and SD for different programs.

Video-on-Demand (VOD)

Every year for the last few years, VOD seemed poised to put video rental shops out of business. The economics have been proven and deployments are growing, but infrastructure upgrades have taken time and are ongoing. At the start of Year 2002, operators had launched or planned to launch VOD (commercially or in trials) in almost 90 markets[3]. The Time Warner trial of "Subscription HBO" demonstrated strong consumer interest and willingness to pay $9.95 a month for the feature[3]. Subscription VOD provides the capability to pause, resume and rewind broadcast programs. It offers a selection of past programs that can be viewed at the consumer's convenience.

VOD has much more potential than just replacing the video store. Extrapolating subscription models further, MSOs could offer server based Personal Video Recording (PVR) capability. As VOD services gain popularity, MSOs will experiment and identify consumer interest and associated usage patterns to determine which can cost effectively be offered. HDTV VOD might be an interesting proposition. Early HDTV adopters are good candidates for higher priced VOD content.

The estimate of cable plant bandwidth for VOD begins with the take rate of digital subscribers. At the end of year 2002 digital penetrated approximately 17% of HP. Kagan's 2001 annual growth forecast[4] shows digital cable penetration growing to 63% by year 2011. The estimate of simultaneous use during peak hours is 5 % today and forecast to increase to 9 % by year 2011. Combining take rate, simultaneous use and program bandwidth (4 Mbps per program) the bandwidth required per HP is calculated. The assumption is that 100% of digital subscribers are provided VOD service. This may not be the case, but it yields sufficient bandwidth for some likely scenarios. If, for instance, only 50% of digital subscribers take VOD, then the simultaneous use could double to 18 % and we end up with the same aggregate bandwidth demand.

Internet Access

Internet Access is defined to cover "Best Effort" transport of IP packets between a subscriber and Internet Service Provider (ISP). The ISP service could be provided by the cable MSO or by a third party ISP with the MSO providing only access transport services. This is the broadest category and most difficult to forecast because of endless possibilities for new applications. Before talking about growth in user demand, it is necessary to define a starting point. A model bandwidth profile is established that typifies data link performance necessary to user to "satisfy" current users.

Users' key expectations of the broadband Internet are low latency in delivery of web pages and downloads, rapid updates in games and seamless delivery of streaming content. They also expect "always on" service. There is no clear consensus as to the data rate that defines broadband. US MSOs typically limit downstream data rates at 2 Mbsp and upstream at 384 kbps. Customers of those MSOs who do not rate limit can experience 5 Mbps or more. In contrast a large MSO in the United Kingdom offers two levels of service, 128 kbps and 512 kbps. Some believe always-on is the key attraction for English consumers and a data rate of 128 kbps is satisfactory.

Consumer cable Internet access is offered as a "Best Effort" service, usually with rate limits. Some MSOs' service agreements prohibit certain uses such as web server hosting and Virtual Private Networks (VPNs). The vast majority of DOCSIS packets are presently attributable to web surfing. Audio sharing is a significant component of the traffic mix on some systems, with peer-to-peer applications like Morpheus replacing Napster. Email and chat are popular, but messages are too small to significantly affect the overall bandwidth consumption.

Capacity planning for Internet Access services is more complex than that required for broadcast services. Usage is driven by user demand on an instantaneous basis, rather than by a more-or-less constant rate stream of content. Usage demand is bursty, and accurate traffic models of traffic are extremely complex. Interactions between network dynamics and the Transmission Control Protocol's (TCP's) congestion control algorithms affect both network utilization and subscriber-perceived performance. One cannot even talk about "bits per second" without asking "measured over how long a period?" Fortunately, crude approximations are good enough for our purposes.

Actual cable modem usage data was difficult to come by. What information was available varied considerably over time and by MSO. The most common provisioning practice seems to be a gross average. The total usable system capacity is simply divided by a target average capacity per provisioned modem to determine the maximum number of modems on a segment. Operators seem to set target capacity per modem by some combination of rules-of-thumb, customer satisfaction indicators and measurements.

It is estimated that YE 2001 a provisioned gross average bit rate of 21 kbps per subscriber was needed to achieve subscriber satisfaction. For example, DOCSIS with 256 QAM modulation in the downstream provides a data rate of approximately 38 Mbps (ignoring overhead). In typical traffic engineering practice, links are provisioned so as to be loaded to 50% of their raw capacity (as averaged over the busiest 15 minute interval of the day); this yields 19 Mbps usable capacity on a DOCSIS downstream carrier. This can be divided by 21 kbps per modem to support approximately 900 modems. If take rate for cable modems is 10% of HP, a single downstream carrier could support nodes totaling 9000 HP. Many systems have been provisioned at approximately this capacity.

For bursty traffic, gross averages are somewhat misleading in provisioning and measurement. From the consumer's perspective, the broadband experience is defined by the approximate peak data rate seen by a receiver during a burst (taking into account TCP congestion control, buffering in the network and packet loss). This makes peak data rate at busy hour a good target service objective. Since service economics depend upon large statistical gains, a simple on-off traffic model is employed as an approximation, and duty cycle (on/off ratio) is used for a provisioning metric. Capacity is provisioned as the usable system capacity divided by product of number of modems, estimated percentage of active users at busy hour, target peak rate and duty cycle. For purposes of this model, it is estimated that at YE 2001, 50% of subscribers are active at busy hour, with a 3.5% duty cycle, and that a satisfactory broadband experience requires a peak rate of about 1.2 Mbps. Fortunately, the long term average model is convertible to the peak burst model by dividing the average provisioned capacity by the product of busy hour percent active and duty cycle. Thus, in terms of our YE 2001 estimates, 21 kbps/(0.5 activity ratio * 0.035 duty cycle) = 1.2 Mbps.

In trying to forecast per-subscriber usage growth, trends in various applications were considered. For example, increases have been noted in streaming media objects in web pages, "post-Napster" peer-to-peer file exchange applications like Morpheus, broadband multiplayer games on the X-Box and (starting in the summer) Playstation 2, consumer-to-consumer exchange of digital video clips and snapshots. Distributed computing programs such as Search for the Extraterrestrial (SETI) encourage users to donate CPU power and bandwidth when they are not using their computers. This blurs the lines between active and in-active users. Chat programs are moving to offer audio and video clips that would boost the size of a message by at least a couple of orders of magnitude. Other trends suggest more web based applications and web based personal data storage.

That said, it was rapidly realized that by looking for “killer apps”, one could not see the forest for the trees. The real power of the Internet is in providing a communications substrate that enables innovation and rapid deployment of new and previously unimagined applications that in themselves become “killer apps”. None of us truly has the prescience either to pick winners and losers among the applications that are now emerging, or to predict the emergence and traffic characteristics of applications that have yet to emerge.

Instead, our forecast extrapolates from historical experience in the Internet[5] and from Moore’s law in allied technologies such as microprocessor computing power and memory density. This leads to an exponential growth model (or, to be more accurate, a logistical growth model which is indistinguishable from exponential growth in the near term). The 10-year forecast, therefore is for consumption to grow at 50% per year. Specifically, peak rate grows at 25% per year, and duty cycle at 20% per year. In practice, actual growth is not expected to be closely fitted to these growth curves, but it is expected to be a good fit over time with appropriate smoothing. Thus, by YE 2011, average active users will demand a peak rate of 11 Mbps, and a duty cycle of 22%.

As for upstream from the home to the network, the peak rate is expected to increase from 200 kbps to about 3.2 Mbps in year 2011. Average upstream consumption increases from 7 kbps to 700 kbps. Upstream bandwidth increases more than downstream due to the expectation that rate asymmetry (the ratio of downstream to upstream rates) will decrease from 6:1 to 3.5:1.

Technology savvy users will seek to push the envelope of what can be done with their broadband connections. Such subscribers can create a “tragedy of the commons” by taking unfair shares of the system capacity and ruining the broadband experience for others. MSOs can treat this as a threat, and put service agreement restrictions and mechanisms in place to keep problematic applications out of the network. Alternatively, they can treat it as an opportunity and put tiered service agreements and mechanisms in place to maximize total subscriber satisfaction and revenue. DOCSIS 1.1 and PacketCable offer architectures and protocols needed to coordinate QoS mechanisms for dynamic and provisioned service flows. However, the specific mechanisms are, to the extent possible, left to implementers. Even for “Best Effort” service, stronger or weaker CMTS implementations can greatly affect fairness and performance, both as seen by individual users and as seen by the MSO. Advanced CMTS features such as per flow queuing, longest queue push-out, hierarchical scheduling with rate guarantees, and fine-grained flow classification can ensure fairness and isolation of service flows within a service class, performance to service level agreements and performance of delay sensitive applications.

Considering all the forces at work, consumers will demand and broadband providers will have to deliver more bandwidth per user each year. With broadband service prices relatively flat, technology and economics need to continue to drive cost per bit down, thus keeping the business healthy.

Streaming high quality IP audio and video

Leveraging the QoS capabilities of DOCSIS 1.1, MSOs are uniquely positioned to offer high quality subscription audio and video streaming services. There is a segment of the population that is happy to listen to audio and watch video on their PC, but mass-market penetration of streaming will likely wait until solutions are in place to move the content into the entertainment center and other places within the home. Lacking solid QoS guarantees, entertainment quality video and audio cannot be delivered reliably enough to satisfy paying consumers.

Streaming audio bandwidth demands are modest (see service example in the next section) for DOCSIS downstream data rates and devices exist to move the audio from a PC to where consumers want to listen. For example, Motorola offers a wireless product called SimpleFi that links the PC to home stereo systems. A transmitter unit plugs into the USB connector of a computer and transmits streaming audio over a 2.4 GHz link. A user-friendly device with a IR remote control connects to RCA jacks on a stereo amplifier. A menu allows user to select between streamed Internet audio and stored personal audio files on their PC.

Hi QoS Audio Streaming

As an example bandwidth demand scenario, assume that 128 kbps is sufficient for high quality audio. Most people find that MP3 encoding at 128 kbps has acceptable sound quality. One could argue for a higher or lower data rate, but audio codecs are still improving, and it is unlikely that 128 kbps will be too low a rate in the future[6]. For provisioning estimates, assume the service is rolled out in year 2002 with a take rate of 2% HP. A successful service might grow to a saturation penetration of 20% HP by 2005. Assume the extreme case where each subscriber stream is on 100% of the time, hence maximum plant bandwidth requirement is 128 kbps x 20% take = 25.6 kbps/HP.

VOD over IP

VOD can be streamed over the Internet using IP and DOCSIS, but end-to-end QoS mechanisms are required to support continuous data rates in the range of 2 Mbps to 4 Mbps. At these rates, audio and video quality is competitive to that offered over MPEG 2 multi-program transport streams to set-top boxes.

The previous section on VOD forecasted VOD bandwidth and accounts for both delivery methods. The number of viewers in the home is the same regardless of the video delivery method. Some VOD service will likely migrate to IP delivery, but there are too many uncertain factors to predict how much and when.

PC based multimedia decoders (Windows Media Player, RealPlayer, QuickTime, and ultimately MPEG4) are widely used to deliver low resolution, low rate VOD over “Best Effort” Internet access service. The Internet Access forecast accounts for this type of VOD.

IP Telephony

Plans to rollout IP telephony in North America stalled as the economy faltered mid year 2000. PacketCable has had time to refine standards and equipment manufactures have had time to test and mature designs. Rollouts of IP telephony are expected to begin gradually late this year. Where offered, IP telephony is estimated to grow from a 2% penetration in year 2002 to about a 30% HP penetration by year 2011. Subscribers are expected to purchase an average of 1.5 IP telephony lines initially tending toward 2.0 in the later years. Penetration estimates take into account the likelihood that Voice-over-DSL will emerge as a second line telephony competitor to cable, and also take into account substitution competition from cell phones. The forecast for an eventual 30% penetration, assumes that MSOs will offer price competitive service.

For calculating bandwidth, a data rate of 100 kbps full duplex per call is assumed. The short packet size required to minimize voice latency adds considerable DOCSIS overhead. This assumption may be slightly low if only G.711 (64 kbps PCM CODEC) is deployed, but high if any other CODECs allowed by PacketCable are used. Typical traffic engineering practice allocates 0.1375 Erlangs per residential subscriber line and allocates enough trunks to ensure a call blocking probability a less than 1%. In populations of 60 to 150 lines, the number of required trunks can be approximated as 20% the number of lines. The node segmentation model (described later) tends to maintain the number of lines within this range. These assumptions allow comparison of telephony bandwidth requirements to other data service components on a HP basis.

SUMMING THE BANDWIDTH DEMANDS

In the next sections demands for IP delivered data and VOD are totaled. The aggregate amount determines node size. The average user consumption is multiplied by projected take rates to normalize data requirements to bits per home passed (HP). A user demand driven node segmentation example is presented.

Downstream to the home

The four downstream components that are destined for individual subscribers are:

- Internet Access
- Hi-QoS Audio Streaming
- IP Telephony
- VOD.

The graph in Figure 1 provides a sense of the growth in bandwidth demand. This result is a combination of increased per-user consumption and a growth in user population. Note that VOD is a significant factor in the near term, but levels out and eventually Internet Access dominates. Hi-QoS Audio and IP telephony are never relatively significant components because by the time subscriber take rates become significantly large, VOD and Internet Access have grown much larger. Figure 2 zooms in on the forecast out to Year 2006 and highlights the VOD growth in the near term. Figure 3 takes out VOD to better show the relative growth in data components delivered over IP. Note that in figure 3 the starting point for Internet Access is 4.2 kbps per HP.

Upstream from the Home

Upstream demand from the home has two components, IP Telephony and Internet Access. The graph in figure 4 shows the growth in demand. Note that upstream data grows even faster than downstream due to the trend toward more symmetrical applications

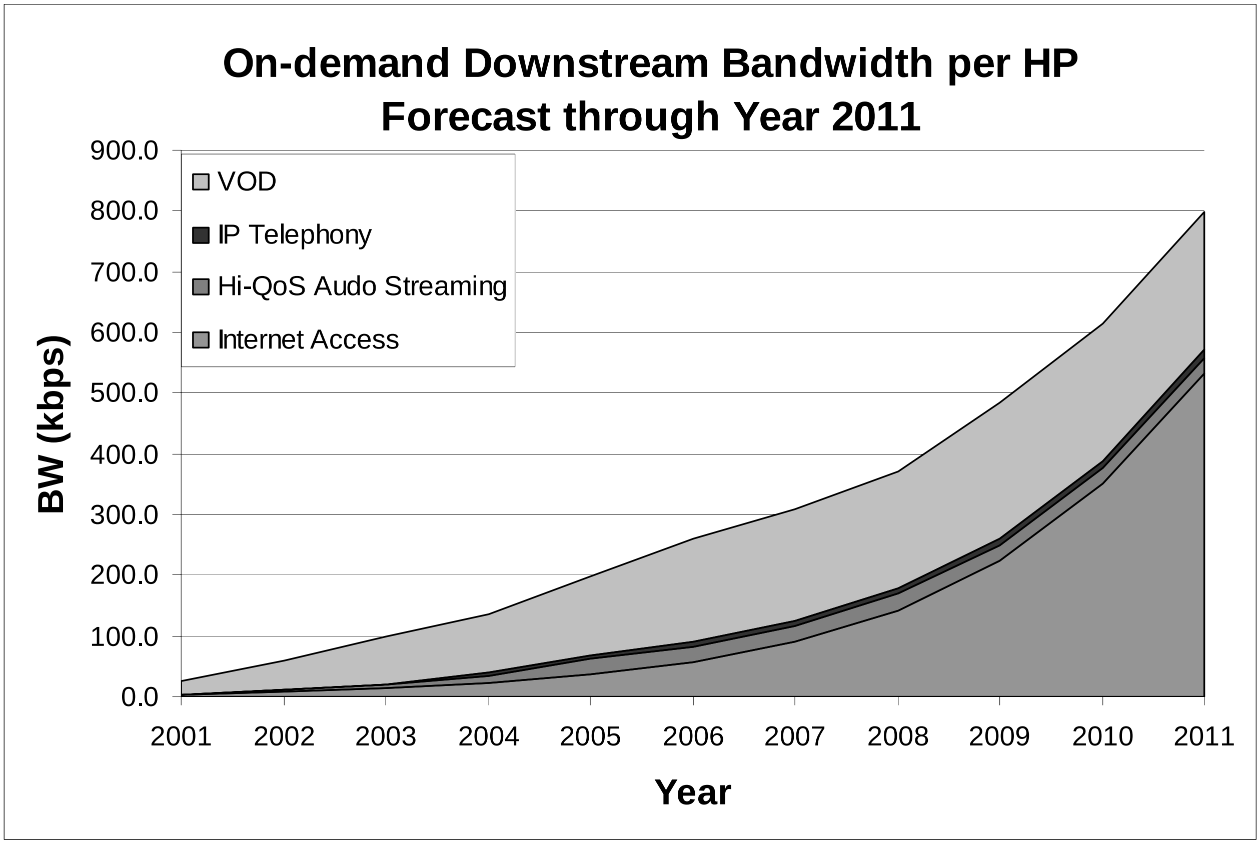

Figure 1

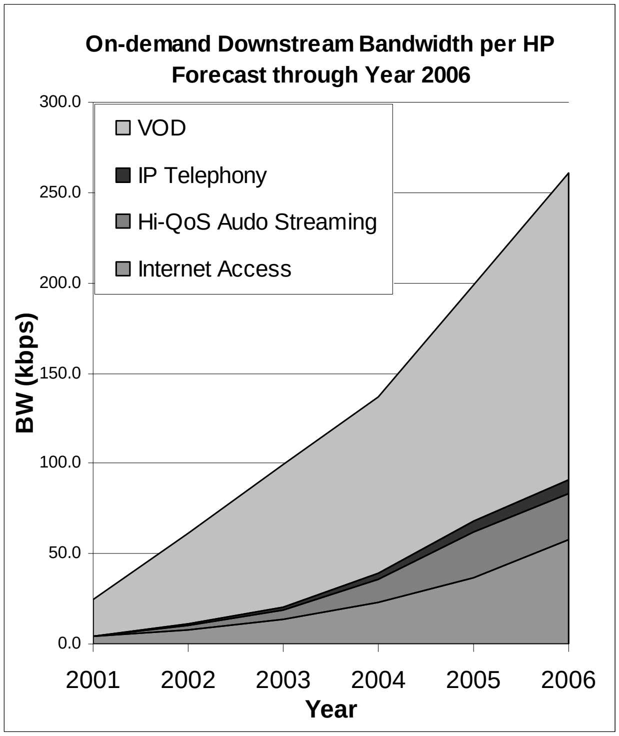

Figure 2

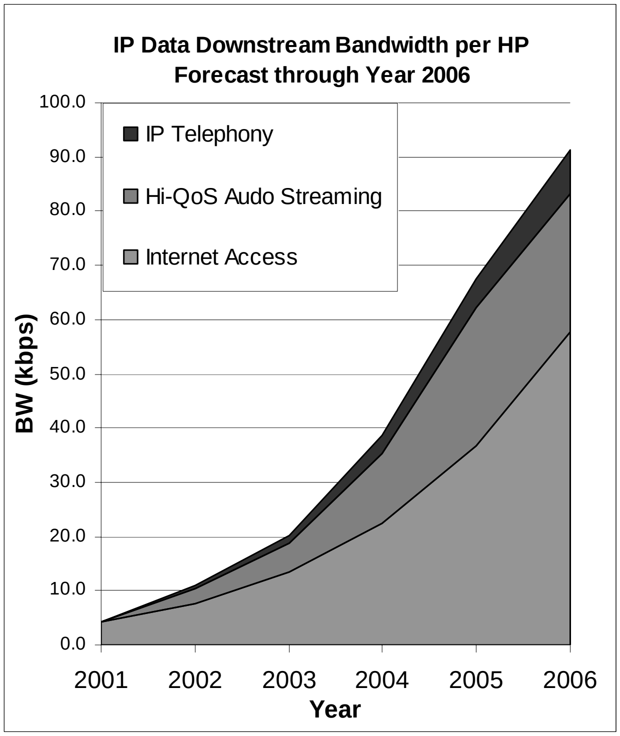

Figure 3

HFC node segmentation example

HFC cable plants deliver optical signals to "nodes", where signals are converted to radio frequencies (RF) and delivered to residences over coax. A tree and branch structure of two-way amplifies and splitters strives to provide each customer with a consistent signal quality. This tree structure also functions to collect upstream data transmitted back from each residence to the headend. The challenge for upstream operation is to overcome the ingress noise that gets funneled up the tree structure from every extremity. Noise is reduced as node size is reduced. As node size is reduced, DOCSIS cable modems can operate at higher symbol rates and higher order modulations. The chart below is a guideline that can be used to determine the upstream aggregate data rate capacity based on node size.

Node Size Homes Passed	Carrier BW (MHz)	Modulation	Number Carriers	Throughput (Mbps)
2000	1.6	QPSK	8	20
500	3.2	QPSK	8	40
125	3.2	16 QAM	8	80

Using this guideline with the forecast a scenario was developed for node segmentation driven by upstream traffic demand. The graph in figure 5 shows a system beginning with a 2000 Home Passed (HP) node and the actual number of HPs that could be supported based on upstream traffic demand. Before the "max node size" curve drops below the number of HPs on the segment, the node is segmented into 4 smaller nodes. The "DS carriers" graph in figure 5 shows the number of downstream DOCSIS carriers that are needed to complement the upstream traffic. In year 2003, the 2000 HP node is segmented to 500 HP in both forward and reverse directions. In year 2007 only the reverse direction is segmented from 500 HP to 125 HP. By choosing to leave the downstream node size at 500 HP, more carriers are required but equipment cost is saved. This configuration supports expected traffic requirements through year 2011.

THE TOTAL PLANT BANDWIDTH PICTURE

Combining broadcast and non-broadcast components an overall bandwidth loading picture was developed. Based on node size the aggregate requirements for non-broadcast traffic is calculated and fit into an integer number of 6 MHz QAM carriers. The cable plant spectrum usage chart in Figure 6 shows how the requirements stack up. It can be see there is spare capacity in HFC plants built out to at least 750 MHz. Although the primary factor driving node size is upstream data, the downstream is segmented simultaneously from 2000 HP to 500 HP. Note figure 6 shows a reduction for DOCSIS and VOD, but subscriber demand is actually increasing. Downstream carriers are added to satisfy downstream demand through year 2011.

A likely scenario has been presented for the typical advanced cable plant. It should be noted that upscale neighborhoods could easily demand much higher amounts of non-broadcast components. Some North American nodes are reported to have had cable modem take rates in excess of 50% HP. Also, there are systems that would like to provide more broadcast channels. Needs include serving multi-lingual and ethnic populations with international programming and lots of HDTV, eventually. 870 MHz plants offer significantly more bandwidth insurance for unplanned demand than 750 MHz. The top of the Figure 6 stops at 870 MHz to provide a relative picture of spare capacity. Current prices for 870 MHz node equipment are a small premium over 750 MHz, hence, a small price for insurance.

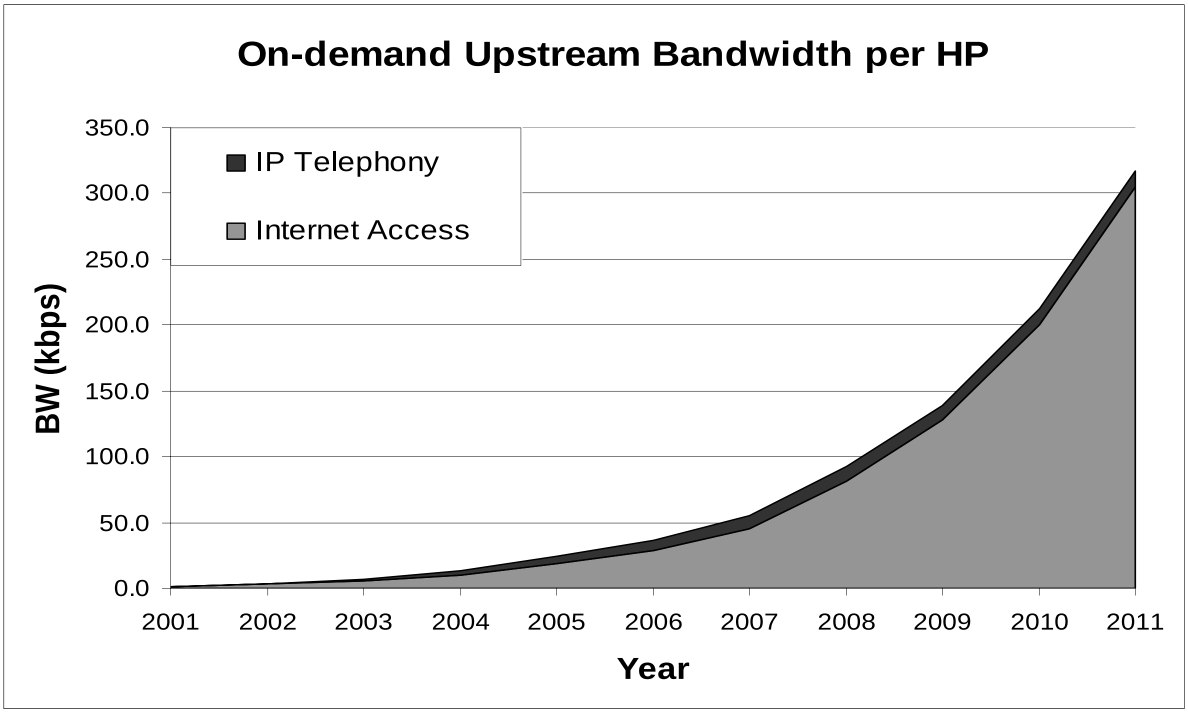

Figure 4

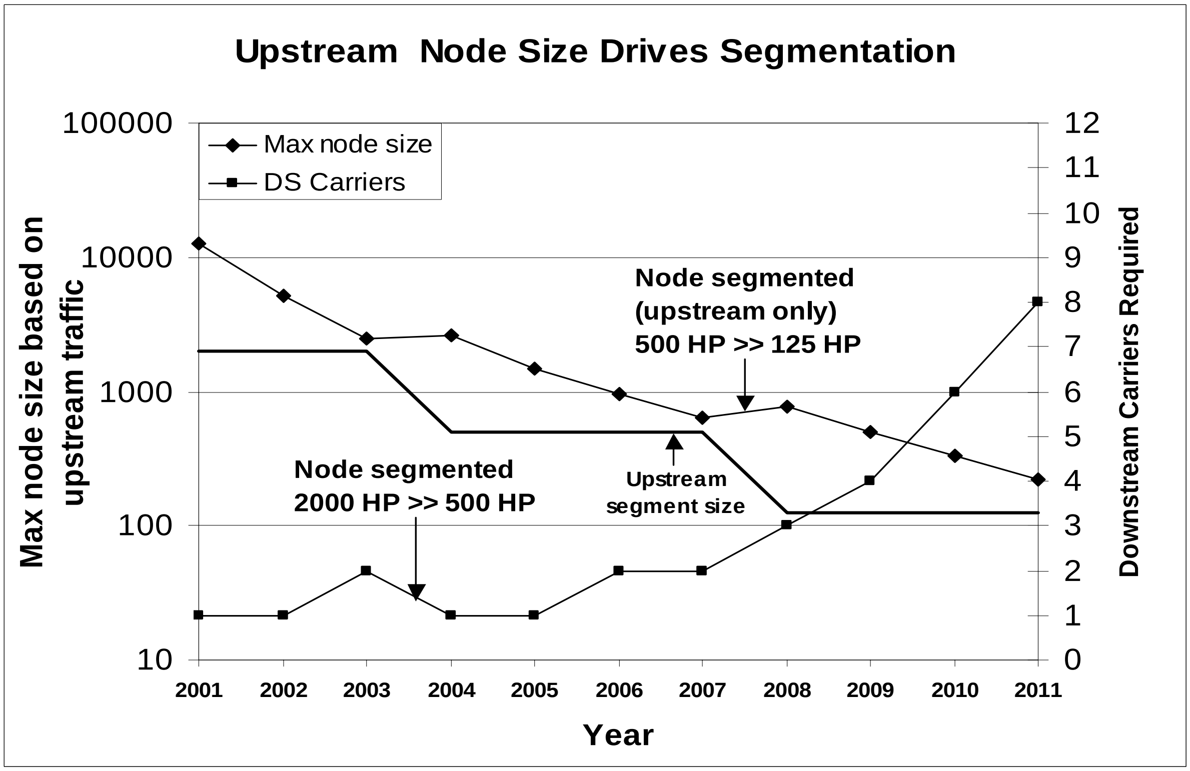

Figure 5

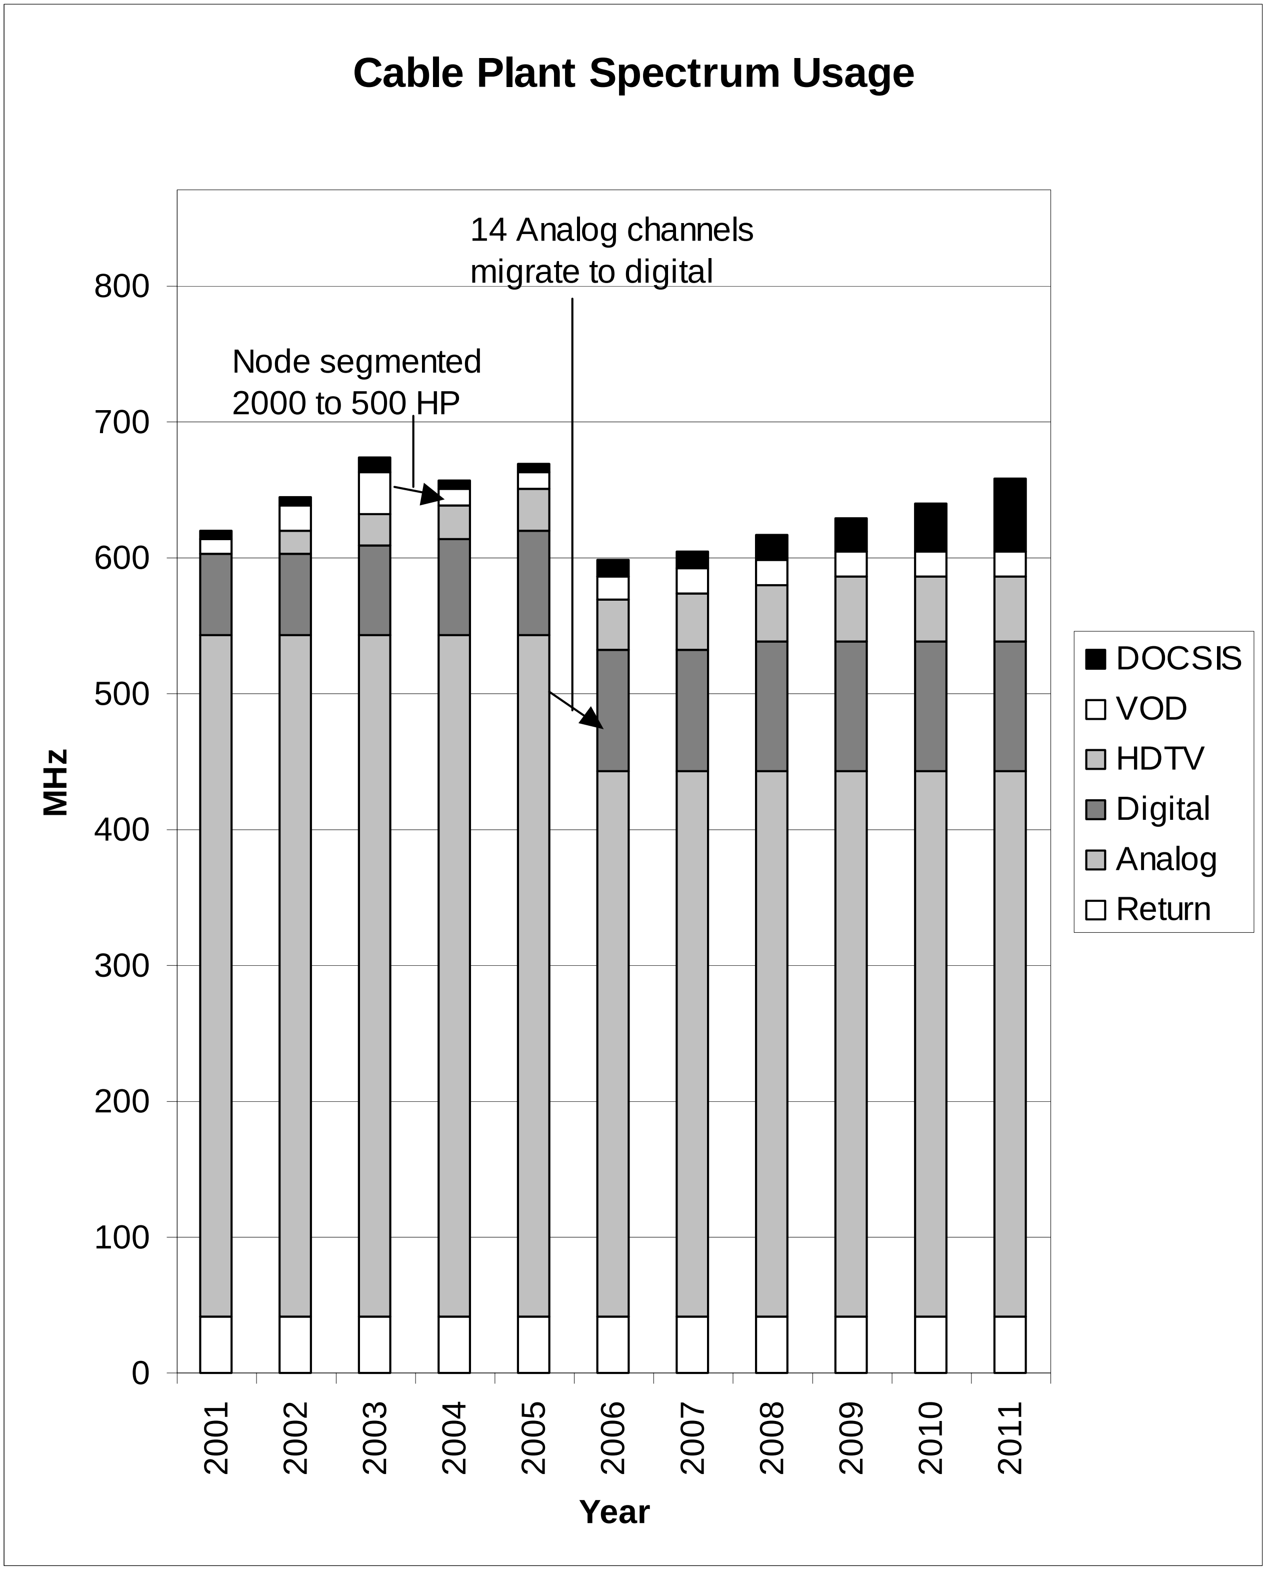
Cable Plant Spectrum Usage
14 Analog channels
migrate to digital
Node segmented
2000 to 500 HP
MHz
0
100
200
300
400
500
600
700
800
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
Year
■ DOCSIS
□ VOD
▨ HDTV
■ Digital
▨ Analog
□ Return

Figure 6

Everything-on-Demand – An Extreme Bandwidth Demand Example

Suppose an MSO wanted to offer all content "on-demand" plus maintain basic analog TV services. A quick calculation will show how this is still possible with a 750 MHz plant! Assume a nominal 500 HP node for the example and that 60% take everything-on-demand service. This is 300 subscribers. With an estimate 2.5 TV sets per home, at most 300 x 2.5= 750 individual streams are required for non-blocking service. The actual number required would be lower due to statistical gains. Dividing by 9 programs per digital carrier we determine that 84 channels are required. The band from 50 to 750 MHz supports (112) 6 MHz channels. Allocating 2 channels for DOCSIS downstream carriers still leaves 26 channels for basic analog service.

In this scenario all digital video content is delivered in 4 Mbps MPEG program streams to each TV. The headend provides broadcast programs either through digital switching or through the VOD server. All digital users have the option of watching live programs or pausing and replaying with an arbitrary time shift. They can choose from a library of stored content that may be included in a subscription or purchase on demand.

It has been shown HFC bandwidth is sufficient for this service. The real questions are how much the remaining infrastructure to support such services will cost and whether consumers are ready and willing to pay enough to justify those costs. The headend equipment would be very different and more complex than today's broadcast equipment. Emerging VOD offerings and PVRs will should begin to whet consumers appetite for this type of service.

Fiber to the Home (FTTH)

Beyond year 2011 MSOs will have to decide between pushing HFC capacity further or re-trenching to bring fiber to the home. Capital costs for FTTH are expected to become competitive for green fields deployments well with the 10 year forecast period. Fiber cable costs are competitive with coax today. Tools and techniques are improving that dramatically reduce the time to make splices. The biggest cost hurdle is the residential and field opto-electronics and the headend infrastructure. If analog TV is to be delivered in addition to digital services, the optics cost is higher for both plant and residential equipment. If only digital television is delivered, then set-tops are required for all TVs and an infrastructure like that described in the everything-on-demand scenario is required. It is not clear which of these options will be optimum, but it seems that voice, video and data services are needed to recover costs of building FTTH. FTTH is considered the "end game" since it offers enormous bandwidth to the home. The fiber itself is expected to last over 40 years, and more capacity requires only headend and subscriber equipment upgrades. Fiber-to-the-curb or fiber-to-the-building (FTTB) for multi-dwelling housing units offers an intermediate more cost effective step towards FTTH. FTTB solutions are expected to be cost competitive in the next few years.

International comparisons to North America

Penetration of Direct-to-Home (DTH) satellite video services in Europe is much higher than in North America but digital cable lags. Europe is impossible to completely generalize as situations vary greatly from country to country. From a recent, but limited sampling of European operators it was noted most are just beginning to upgrade plants beyond 550 MHz and deploying RF return capability.

Analog offerings are often less than 40 channels (as compared with 80 in typical US systems). There is interest in moving toward digital cable TV, but little interest in HDTV. Some believe that because PAL & SECAM offer superior resolution to NTSC, HDTV is less necessary. Services using Euro-DOCSIS are expected to expand, but deployments will be paced by plant upgrades

South America tends to have limited analog content and one-way plants of less than 550 MHz. Little change is expected in the near term due to limited capital for upgrades.

Asia-Pacific with the exception of a few countries tends to lag in all dimensions of service offerings. South Korea is one exception where cable modems and DSL are highly penetrated. Japan is planning aggressive deployments of fiber, but the economic model seems unclear.

Conclusion

The first 50 years of cable were about providing more signals and better picture quality. The next 50 will be about data services. In 10 years the number of bits pouring into the home will be over 50 times the amount delivered today. Rich interactive multimedia video will be commonplace. HDTV will succeed as one of the many services. Telephony will become a rounding error in the traffic analysis. This growth has been shown to be easily supported by continuous upgrades to the HFC infrastructure. Capacity estimates are conservative based on current technology. Much more can be squeezed out of HFC, if and when needed. The critical issue for MSOs will be monitoring usage and keeping ahead with provisioning. A good understanding of what users are doing and how the applications consume bandwidth is essential to anticipating changes in demand. Evolving to more sophisticated service metering will ensure customer satisfaction and QoS based services present a wealth of revenue opportunities.

MSOs will continually be faced with capital investment trade-offs between infrastructure upgrade costs vs. how much excess capacity to install. The node segmentation example divides nodes by four in years 2003 and 2007. Smaller divisions would result in more frequent upgrades, higher labor cost and more disruptions. Larger divisions would incur larger spending and excess capacity would represent a waste of capital employed. Legacy equipment will tend to make upgrade trade-offs more complex and optimum timing will vary. It is hoped that this forecast is helpful to long term planners.

Fiber optic technology is expected to advance at a rapid rate. Breakthroughs in performance and manufacturing cost are potentially disruptive to the scenario presented. If end-to-end costs for delivering bits over the Internet were to drop dramatically and Internet backbone latency and packet loss were considerably reduced an entirely new set of bandwidth intensive applications and businesses could emerge. With delivery costs tending toward distance insensitivity, independent businesses could cost effectively serve a geography disperse customer base.

Acknowledgements

First I would like to thank the numerous staff members at Motorola who participated in our internal discussions. Special thanks go to my colleagues Curtiss Smith, Javier Borlando, Dan Grossman and Jeff Ollis for the help in reviewing, gathering references and customer inputs. Thanks go to the MSO staff who were forthcoming with information on cable modem usage trends and provided great feedback on our predictions.

Contact Information

Randy Nash
Motorola
101 Tournament Blvd.
Horsham, PA 19044
Phone 215-323-1475
Email – randy.nash@motorola.com

[1] And Broadband for All: Future Residential Networks & Services, Module 2, Probe Research Inc., April 2001

[2] Annual Assessment of the Status of Competition in the Market for the Delivery of Video Programming FCC CS Docket No. 01-129, December 27, 2001.

[3] On-Demand, CED In Depth, February 2002

[4] Kagan World Media, 10-Year Cable TV Industry Projections, 2001

[5] Internet Growth: Is there a “Moore’s Law” for data traffic” – Coffman and Odlyzko ATT July 2000

[6] Flying hi-fi, Commverge, January 2002

A CABLE OPERATOR'S GUIDE TO CABLEHOME™ 1.0 FEATURES

Kevin Luehrs, Steve Saunders
CableLabs®
Nancy Davoust
YAS Broadband Ventures

ABSTRACT

CableLabs released the CableHome 1.0 specifications to home networking device vendors and to the general public in April 2002. The specification standardizes a suite of residential gateway functions enabling cable operators to deliver managed broadband services to their high-speed data service subscribers over the subscribers' home networks. This paper introduces the CableHome initiative at CableLabs and the Portal Services (PS) Element as a foundation, and then discusses component functions of the PS Element in terms of setup and configuration, and alternatives for operation. Each cable operator will have the opportunity to configure CableHome-compliant devices in a manner consistent with its business objectives. Although many of the options provided by the CableHome 1.0 specifications are described, specific configuration details are beyond the scope of this paper.

OVERVIEW

Introduction to CableHome

CableHome is an initiative undertaken by CableLabs at the direction of its member cable television companies to develop an infrastructure enabling cable operators to extend high-quality, managed, value-added broadband services to subscribers in their homes in a fashion that is as convenient as possible for the subscribers. The CableHome 1.0 specifications are a set of functional and messaging interface requirements describing cable-industry standard methods for implementing address acquisition, device configuration, device management, network address translation, event reporting, remote diagnostic procedures, secure software download, firewall monitoring and policy file download, as well as other functions in a residential gateway device or element connecting networked devices. These devices are located in a subscriber's home and are connected to the Internet through a DOCSIS cable modem and a cable operator's hybrid-fiber coaxial (HFC) network. Figure 1 illustrates a number of key CableHome network elements and concepts, which are described below. CableHome 1.0 specifications introduce and use concepts of Wide Area Network (WAN) (cable network) and home Local Area Network (LAN) address realms, translated (LAN-Trans) and non-translated (LAN-Pass) address realms within the home LAN, IP addresses intended to be used for management traffic (WAN-Man IP) or for user/application data traffic (WAN-Data IP), and Embedded (with a cable modem) versus stand-alone residential gateway functions. The specifications refer to LAN IP Devices, which are the elements connected to a subscriber's home network communicating using the TCP/IP protocol suite. The specifications also define a Portal Services (PS) Element, which is a collection of functions providing the capabilities listed in the previous paragraph between the WAN and LAN realms, serving networked devices in the translated address realm in the home, and providing management capabilities for monitoring and configuring the various functions of the PS.

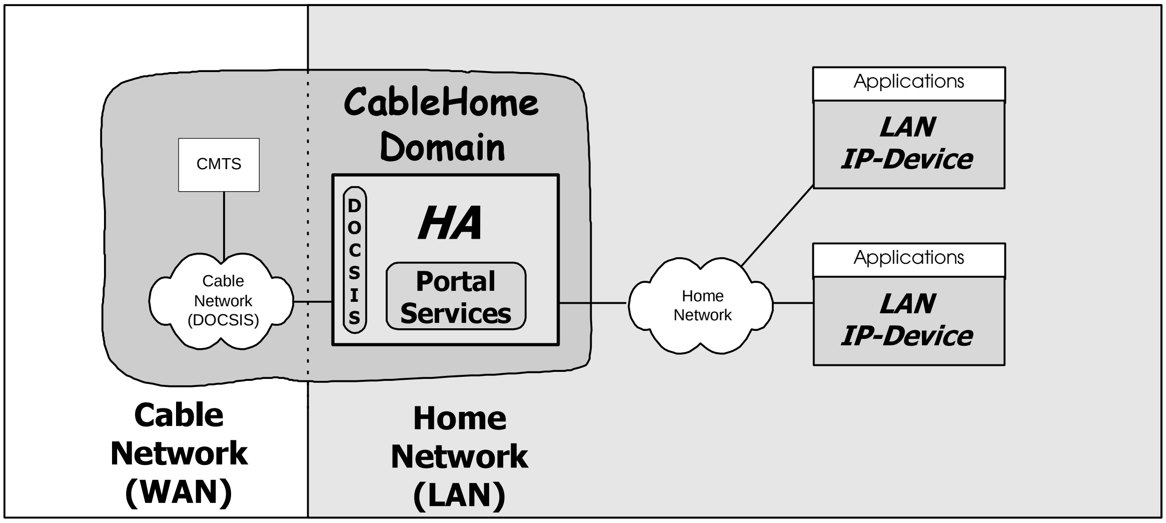

Figure 1: CableHome Network Elements

CableHome functions, in the form of the PS, always operate in conjunction with DOCSIS functionality. The PS functions may reside within the same physical device as the DOCSIS functionality (Embedded PS) or in a separate device (Stand-alone PS). In all cases, CableHome identifies the Home Access (HA) Device as the collection of DOCSIS and CableHome functionality that connects the Cable Network to the Home Network.

The CableHome Domain is the extent of the WAN and LAN networks which carry CableHome messaging and includes devices that implement CableHome functionality. CableHome extends DOCSIS functionality further to the edge of the cable network and provides a set of tools the cable operator can use to better support his or her high-speed data subscribers. In the following sections, the PS Element is described to the next level of detail.

Portal Services Element

The CableHome 1.0 PS Element is a collection of eight functions intended for implementation in a residential gateway device with a broadband connection through a DOCSIS cable modem. The PS supports multiple IP clients in the home, and gives them controlled access to the Internet via the cable modem's CPE interface. Most of the PS functions are similar to functions implemented in residential gateway products in the market today, with additional features to enable cable operators to provide quality, managed, value-added service to their high-speed data service subscribers.

Seven of the eight CableHome PS functions are referred to as portals since they link the cable operator's WAN to the subscriber's home LAN. The CableHome Portal functions are listed below:

1) CableHome Dynamic Host Configuration Protocol (DHCP) Portal (CDP):
Provides network address information functions including a server for the networked elements in a subscriber's home.

2) CableHome Address Portal (CAP):
Interconnects the WAN and LAN address realms for application data traffic via network address translation and bridging.

3) CableHome Management Portal (CMP):
Provides an interface between the cable

operator and manageable parameters in the PS through the CableHome-specified Management Information Base (MIB).

4) CableHome Naming Portal (CNP):
Provides a simple DNS service for networked home devices requiring naming services.

5) CableHome Test Portal (CTP):
Provides a means for the cable operator to initiate remote ping and loopback tests.

6) CableHome Security Portal (CSP):
Participates in authentication, and exchanges keying material with the Key Distribution Center (KDC) server for CableHome security functions in the PS.

7) CableHome Quality of Service (QoS) Portal (CQP):
Provides transparent bridging for QoS messaging between PacketCable applications and the PacketCable QoS infrastructure on the cable network.

The eighth PS function is a firewall function that provides protection of the home network from malicious attack.

Relative to residential gateway products currently available through retail, a CableHome 1.0-compliant residential gateway with PS functions will allow the cable operator to better manage subscribers' broadband experiences and minimize subscribers' home network down time. CableHome 1.0 also allows cable operators to provide security and privacy to their high-speed data subscribers.

Compatibility With Existing CableLabs Specifications

An important design goal of the CableHome 1.0 specifications was to be compatible with existing CableLabs and industry specifications to the greatest extent possible. The CableHome specifications were developed for networked elements that will connect to cable operators' HFC network through a CableLabs Certified DOCSIS cable modem. The CableHome PS Element is an extension of the DOCSIS infrastructure, and employs procedures very similar to those required in DOCSIS specifications for management, event reporting, configuration file download, and secure software download.

CableHome 1.0 specifies features compatible with a DOCSIS 1.0 infrastructure, as well as additional features that support advanced capabilities of DOCSIS 1.1 and PacketCable™ infrastructures. CableHome-specified features therefore provide cable operators with a migration path as their facilities evolve over time.

PROVISIONING

CableHome 1.0 specifications define provisioning as the device initialization and initial configuration required to enable the PS Element and networked devices in the home to exchange meaningful information with one another and with elements connected to the cable network and to the Internet. CableHome specifications define a set of provisioning tools to accomplish this so that the cable operator can add value to the process. A goal of the CableHome specification is to define provisioning processes that enable all CableHome functionality without the need for subscriber interaction.

CableHome provisioning tools consist of a DHCP client, a DHCP server, a bulk configuration tool, and a time of day client. These tools have been designed to work on cable networks implementing DOCSIS 1.0 or

DOCSIS 1.1 cable modems (CM) and cable modem termination systems (CMTS), as well as on cable networks where advanced features such as those defined in the PacketCable specifications, are deployed.

Address Assignment

The first step in the device initialization process is the acquisition of a network address. The PS plays a dual role with respect to network address acquisition. The CableHome DHCP Portal (CDP) function is comprised of a CableHome DHCP Client (CDC) to acquire one or more network address(es) from the cable operator, and a CableHome DHCP Server (CDS) to assign private IP address leases to networked elements in the home.

DHCP Options 43, 61, and 177

To fully support addressing capabilities specified by CableHome 1.0, a cable operator's DHCP server must interpret DHCP Option 61 (client identifier), DHCP Option 43 with sub-options, and DHCP Option 177 with sub-options, plus 18 additional standard DHCP options.

Option 61 allows the PS to uniquely identify itself to the headend DHCP server when requesting multiple WAN-Data IP addresses using a single unique hardware (MAC) address. Support of Option 61 is required when operating the PS in NAT Primary Packet Handling Mode.

Option 43 with sub-options allows the PS to provide more detailed information about its capabilities. For example, CableHome 1.0 defines Option 43 sub-option 2 for the PS to indicate whether it is embedded with a cable modem or is a stand-alone device. Sub-option 11 indicates whether the address the PS is attempting to acquire is for the WAN-Man Interface or for the WAN-Data Interface.

The PS uses Option 177 sub-option 3 to request the location of the cable operator's SNMP manager, and uses sub-option 11 to request the KDC server's IP address.

CDC Operation

CableHome 1.0 specifications require the PS to implement two unique hardware (Media Access Controller - MAC) addresses. The WAN-Management (WAN-Man) MAC address allows the PS to uniquely identify itself to the address server (DHCP server) in the headend for acquisition of an IP address to be used for the exchange of management messages between the cable operator's network management system (NMS) and the management entity in the PS. The CDC will always attempt to acquire this WAN-Man IP address. Depending upon which primary packet handling mode the PS is configured to operate (described later in this document), the WAN-Man IP address may be the only IP address the PS acquires. The second hardware address specified by CableHome 1.0 is the WAN-Data MAC address, intended to be used in conjunction with DHCP Option 61 to uniquely identify one or more PS WAN Data Interface(s) for the acquisition of one or more global IP address(es) to map to private IP addresses in the home. The factory default value for the number of WAN-Data IP addresses the PS is required to request is zero. The cable operator must modify a WAN-Data IP Address Count parameter in order to configure the PS to request one or more WAN-Data IP addresses. The cable operator can configure the PS to use the WAN-Man IP address for application data traffic as well as for management traffic. Alternately, the PS can be configured to use one WAN-Data IP address to share among one or more LAN IP Devices when operating in port translation mode, and one or more WAN-Data IP addresses for 1:1 mapping to private LAN IP addresses when operating in address translation mode.

CDS Operation

Unless the cable operator chooses to serve all LAN IP Devices in the home with network addresses directly from the headend DHCP server, he must configure the PS to assign private IP addresses to the subscriber's home LAN elements. This is the function of the CDS. The CDS grants leases for private IP addresses in response to DHCP DISCOVER messages issued by LAN IP Devices, within constraints defined by three cable operator-configurable management parameters: a LAN Address Threshold limit and LAN Address Pool Start and End parameters defining the range of private IP addresses available for assignment.

The CDS supports 18 standard DHCP options. The cable operator can provision values for these options, or allow the PS to assign factory default values defined in the CableHome specifications. The CDS does not "pass through" DHCP options received from the headend DHCP server to LAN IP Devices.

CableHome Provisioning Modes

Two Provisioning Modes are defined in the CableHome 1.0 specifications: DHCP Provisioning Mode and Simple Network Management Protocol (SNMP) Provisioning Mode. The cable operator configures the PS to operate in these modes as described below.

DHCP Provisioning Mode follows closely the provisioning method defined for a cable modem in the DOCSIS 1.0 and DOCSIS 1.1 specifications, and is intended for compatibility with DOCSIS 1.0 and DOCSIS 1.1 systems. Characteristics of DHCP Provisioning Mode are listed below:

- PS configuration file name and location are provided to the PS in the DHCP OFFER
- The PS is required to download a PS configuration file
- The PS will default to using SNMP version 1 and version 2 for management messaging, but can be configured by the cable operator to operate in SNMP version 3 coexistence mode

The cable operator configures the PS to operate in DHCP Provisioning Mode by including the location of the Trivial File Transfer Protocol (TFTP) server containing the appropriate PS configuration file in the *siaddr field* and the name of the PS configuration file the file field of the DHCP OFFER message, AND by not including DHCP Option 177 sub-option 51 (Key Distribution Center (KDC) server location) in the DHCP OFFER. If Option 177 sub-option 51 and either or both of the PS configuration file parameters are not present, or if all three parameters are present in the DHCP OFFER message received from the headend DHCP server, the PS will generate an event indicating an error condition, and re-issue DHCP DISCOVER to try again for a valid combination.

SNMP Provisioning Mode allows the PS to take advantage of advanced features similar to those defined in the PacketCable Multimedia Terminal Adapter (MTA) specifications. Characteristics of PS operation in SNMP Provisioning Mode are as follows:

- The PS will authenticate itself to a Key Distribution Center using the Kerberos protocol, and will exchange security keys with the KDC to use when exchanging management messages with the NMS via SNMP version 3
- PS configuration file download is optional. If no PS configuration file is provided, the PS will operate using factory default settings.

- The PS will default to SNMPv3 Coexistence Mode operation and will use SNMP version 3 for management messaging
- The cable operator may optionally trigger the PS to download a PS configuration file by writing the PS configuration file location and file name in URL format via SNMP version 3

The cable operator configures the PS for SNMP Provisioning Mode by not including PS configuration file information (file name and TFTP server location) and by including DHCP Option 177 sub-option 51 (KDC server location) in the DHCP OFFER message sent to the PS. When configured to operate in SNMP Provisioning Mode, the PS will exchange messages with the KDC server to acquire keying material to authenticate itself with the KDC server. Once authentication has been completed, the PS is capable of exchanging secure SNMP version 3 management messages with the NMS in the headend. When secure management message exchange is enabled, the cable operator has the option of modifying a parameter in the PS via an SNMP message to trigger the download of a PS configuration file. However, since the CableHome specifications define factory default values for all necessary parameters, the PS does not necessarily require a PS configuration file to operate, and could potentially operate indefinitely on the cable network without receiving a configuration file.

PS Configuration File

The PS configuration file provides a means for the cable operator to issue configuration instructions in bulk to a PS Element. It also provides the means for providing the PS with code verification certificates (CVC) used for secure software image download procedures.

CableHome 1.0 specifies TFTP for the transfer of configuration files (PS configuration file and firewall configuration file) from the cable operator's headend to the PS.

Any configuration file downloaded by the PS should be authenticated to ensure that the file is not corrupt. The PS configuration file is authenticated with a hash value, which is a code compared to the result of a calculation performed on the file itself. Correct correlation between the hash value and the calculated value indicates that the file is valid. When the PS is operating in DHCP Provisioning Mode, the hash value is passed to the PS with the configuration file name in the *file* field of the DHCP OFFER message. Download of the PS configuration file to the PS operating in DHCP Provisioning Mode is triggered by the presence of the configuration file name and location in the DHCP OFFER.

When the PS operates in SNMP Provisioning Mode, the cable operator must provision the hash value in the PS by writing it to the PS via SNMP, before the configuration file download is triggered by a second SNMP message writing the configuration file name and address to the PS.

Once triggered to download the PS configuration file, the PS will continue trying to download the file until it successfully downloads and processes the file. If the PS encounters an error when processing the PS configuration file, it will report the failure as an event and try again to download the file.

The cable operator is responsible for correctly sequencing configuration parameters in the PS configuration file, for not creating conflicts between configuration file-set parameters and SNMP-set parameters, for providing the correct PS configuration file name and location to the PS, for providing the correct hash value to the PS, and for correctly triggering the file download.

MANAGEMENT

CableHome 1.0 specifications describe several features allowing the cable operator to monitor and configure PS parameters, format event reporting, and initiate remote testing for diagnosing problems on the home LAN. The CableHome Management Portal (CMP) function of the PS is the entity that responds to SNMP management messages from the cable operator's NMS. Access to the CMP is through the PS WAN-Man Interface, which is bound to the WAN-Man IP address.

The PS is capable of operating in two Management Modes. The cable operator can configure the PS to operate in NmAccess mode for DOCSIS 1.0 compatibility, or to operate in SNMP v3 Coexistence Mode, which is supported by DOCSIS 1.1 and PacketCable specifications.

Management Mode

A PS operating in DHCP Provisioning Mode defaults to operating in NmAccess Mode, and to using SNMP v1/v2. In this mode the cable operator can control access to management parameters by writing to the NmAccess Table of the DOCSIS Device MIB [RFC 2669]. The cable operator can also put the PS into SNMPv3 Coexistence Mode by writing the appropriate parameters to the snmpCommunityTable [RFC 2576] through the PS configuration file or via direct SNMP messaging. Once in Coexistence Mode, the PS will respond to SNMP v1, v2, or v3 messages.

When the PS is operating in SNMP Provisioning Mode, it defaults to operation in SNMP v3 Coexistence Mode for management messaging. All three versions of SNMP are supported but by default version 1 and version 2 are disabled. The cable operator can enable them by writing appropriate parameters to the snmpCommunityMIB [RFC 2576].

When operating in SNMPv3 Coexistence Mode, access to manageable parameters is controlled by View-based Access Control Model (VACM) Views [RFC 2575] and User-based Security Model (USM) [RFC 2574] Users. CableHome 1.0 defines one User, CHAdministrator, and one View (read and write access to all parameters), which are assigned to the cable operator. With the rights afforded to the CHAdministrator User, the cable operator can create additional Users and Views. In this way, the cable operator can allow access on an object-by-object basis to the subscriber or other parties.

Event Reporting

CableHome 1.0 specifies over 50 defined events for asynchronous reporting of errors and pass and fail conditions for several processes. Most of these are events also defined in the DOCSIS specifications, and some are CableHome-specific events. The cable operator can control how the events are reported, and can throttle individual events by writing to appropriate objects of the DOCSIS Device MIB, support for which is required in the CableHome 1.0 specifications. Events can be reported as local (to the PS) log entries, system log entries, or traps. The cable operator can retrieve local log entries by accessing the appropriate MIB objects using SNMP.

CableHome Test Portal

The CableHome Test Portal (CTP) consists of two remote diagnostic testing functions: CTP Connection Speed Tool and CTP Ping Tool. The cable operator initiates these tests by issuing appropriate SNMP commands to the CMP.

The Connection Speed Tool is a form of loopback test in which the CTP sends packets, the length, number, and frequency of which are specified by the cable operator, to a privately-addressed element connected to the home LAN. If the loopback function is supported in the home network device, it will echo the packet(s) back to the CTP, which will log statistics such as round trip time, packets sent, and packets received for the cable operator to retrieve from a set of MIB objects. The Connection Speed Tool enables the cable operator to gain some performance statistics about the link between the PS and any connected device with an IP address that supports the loopback function.

The Ping Tool allows the operator to ping a privately-addressed device in the home to verify connectivity between the PS and the device. The cable operator configures the destination IP address, number of packets, packet size, frequency, and timeout parameters, and retrieves test results by accessing MIB objects using SNMP.

CableHome-Defined Parameters

Access to PS parameters is one of the values CableHome provides by enabling the cable operator to provide managed service to high-speed data subscribers.

CableHome defines five MIBs for this purpose: PS Device (PSDEV) MIB, CDP MIB, CAP MIB, CTP MIB, and Security (SEC) MIB. MIB objects are accessible through the CMP, via the PS WAN-Man Interface.

The PS DEV MIB provides access to device information such as serial number, hardware version, and MAC addresses; device reset control; provisioning mode control; configuration file parameters; provisioning state information; notification (trap) objects; and software image download parameters.

CDP MIB objects include information about home LAN elements (IP addresses, client identifiers, lease times, host names, and DHCP options); server addresses; LAN address control parameters (address limits and address pool range); and a table of client identifiers associated with the WAN-Data Interface. The cable operator has visibility on privately-addressed home LAN elements through this MIB.

WAN-to-LAN IP address mappings stored in the PS are accessible through the CAP MIB. This MIB also provides access to timeout parameters for the mappings, a table of hardware addresses of LAN elements assigned address leases directly from the headend DHCP server, and the primary packet handling mode control parameter. The cable operator can provision WAN-to-LAN IP address mappings by writing to the appropriate parameters in the CAP MIB.

The CTP MIB contains parameters that control and configure the Connection Speed and Ping Tools, and provides access to the results of those tests.

Firewall parameters including the firewall policy file name and location, hash value, and enable/disable control are accessible through the Security MIB. This MIB also provides control over firewall-related events that allows the cable operator to be notified about various types of attacks.

PACKET HANDLING

A collection of functions within the CableHome Address Portal (CAP) provide packet handling capabilities that enable IP packet flow from the WAN to the LAN, and vice versa (the NAT, NAPT, and Passthrough functions described below). In addition, the CAP provides a function that protects the HFC network from intra-home traffic (the USFS function described below).

The packet handling functions that are applied will be dependent upon whether public, private, or mixed addressing is desired for LAN IP Devices. If the cable operator has chosen to address LAN IP Devices privately, then network address translation (NAT) functions must be employed in order to enable packet flow between the LAN and WAN. If public addressing has been chosen for LAN IP Devices, then the CAP must ensure that all traffic (including DHCP messaging) is transparently bridged between the LAN and WAN. In addition, packet handling for mixed public and private addressing of LAN IP Devices is supported.

In order to control which packet handling functions are active, the cable operator can configure the PS to operate in one of four primary packet-handling modes, which are listed below:

- Passthrough Bridging Mode
- CableHome Network Address Port Translation (C-NAPT) Transparent Routing Mode
- CableHome Network Address Translation (C-NAT) Transparent Routing Mode
- Mixed Bridging/Routing Mode

These modes determine the mapping functions applied between WAN-Data IP addresses and the addresses of devices connected to the subscriber's home LAN.

The cable operator configures the PS to operate in one of these modes by modifying the CAP Primary Mode parameter passed as a PS configuration file parameter or via an SNMP set message. The factory default value of the CAP Primary Mode is C-NAPT Transparent Routing Mode.

The cable operator will configure the PS to operate in Passthrough Bridging Mode when one or more of the following considerations are relevant:

- Public addressing is desired for all IP devices in a home (in anticipation that applications will be running in the home that are C-NAT/C-NAPT intolerant)
- It is important to preserve existing home device address assignment models (i.e. public addresses served directly to CPE by cable network DHCP servers)

It should be noted that, when in Passthrough mode, addresses supplied to a home might reside on different logical IP subnets.

In Passthrough mode, the CAP acts as a transparent bridge for packets flowing between the LAN and WAN. Forwarding decisions are made primarily at OSI Layer 2 (data link layer) and C-NAT/C-NAPT Transparent Routing functions are not applied. Like all other traffic between WAN and LAN elements in this mode, DHCP messaging is bridged, resulting in direct address acquisition communications between home devices and cable network DHCP servers. As a result, all LAN IP Devices will receive public IP addresses, and no address translation will be required.

The cable operator will configure the PS to operate in C-NAPT Transparent Routing Mode if one or more of the following considerations are relevant:

- Conservation of public IP addresses is important
- Same-subnet addressing for devices on the home LAN is important

C-NAPT address translation is a one-to-many mapping function. A single public WAN IP address is mapped to multiple private LAN addresses via port multiplexing. In C-NAPT Mode the CDC acquires one WAN-Data IP address and uses this address for each of one or more WAN address - private LAN address tuple(s). The single WAN-Data IP address can be shared between two or more networked elements connected to the home LAN. C-NAPT address mappings can be created dynamically or the cable operator can provision them.

Dynamic C-NAPT address mappings are created when a privately addressed element in the home network sources a packet destined for an IP address outside the private address space in the home. When the packet reaches the PS, the CAP will determine whether a mapping exists for the source address and, finding none, will create the mapping, replace the packet's source address with the WAN-Data IP address, and forward the packet to the PS Element's default gateway.

If the cable operator creates a C-NAPT mapping, the CAP will find the mapping when an "outbound" packet arrives from the home LAN, replace the packet's private source IP address with the corresponding WAN-Data IP address, and forward the packet to the upstream router in the cable operator's network. C-NAPT mapping creates a 1-to-many association between the WAN-Data IP address and the private IP addresses bound to elements connected to the home LAN.

The cable operator will configure the PS to operate in C-NAT Transparent Routing Mode if one or more of the following considerations are relevant

- Same subnet addressing for devices on the home LAN is important
- Applications that cannot tolerate C-NAPT Routing will be running in the home.
- Source based routing within the cable network will be employed in conjunction with network address translation

It is possible that a set of public IP addresses provided to a home may not reside on the same subnet. C-NAT address translation is a one-to-one mapping function. The PS will own all of the public address supplied to the home, and they will be uniquely mapped to the same-subnet private addresses that have been assigned by the CDS to LAN IP Devices. This one-to-one mapping function enables a privately addressed home device to be uniquely associated with a public WAN IP address, and thus source based routing techniques used in the cable network will not be compromised.

Dynamic C-NAT address mappings are created when a privately addressed element in the home network sources a packet destined for an IP address outside the private address space in the home. When the packet reaches the PS, the CAP will determine whether a mapping exists for the source address and, finding none, will create the mapping, replace the packet's source address with the WAN-Data IP address, and forward the packet to the PS Element's default gateway. If there is not a WAN-Data IP address available against which to create the C-NAT mapping, an event will be generated.

If the cable operator creates a C-NAT mapping, the CAP will find the mapping when an "outbound" packet arrives from the home LAN, will replace the packet's private source IP address with the corresponding WAN-Data

IP address, and will forward the packet to the upstream router in the cable operator's network. C-NAT mapping creates a one-to-one association between the WAN-Data IP address and the private IP addresses bound to elements connected to the home LAN.

The cable operator will configure the PS to operate in Mixed Bridging/Routing Mode if the cable operator wishes to use private addressing for some of the devices in the home (thus requiring C-NAT/C-NAPT Routing functionality) while concurrently using public addressing for other devices (thus requiring the Passthrough Bridging function).

To operate in this mode, the cable operator sets the primary mode to C-NAT or C-NAPT Transparent Routing. In addition, one or more MAC addresses, belonging only to those LAN IP Devices whose traffic is to be bridged, are written into what is known as the Passthrough Table.

When in this mode, the CAP examines MAC addresses of received frames to determine whether to transparently bridge the frame or to perform any C-NAT or C-NAPT Transparent Routing functions at the IP layer. In the case of LAN-to-WAN traffic, the PS examines the source MAC address, and if that MAC address exists in the Passthrough Table, the frame is transparently bridged to the WAN-Data interface. In the case of WAN-to-LAN traffic, the PS examines the destination MAC address, and if that MAC address exists in the Passthrough Table, the frame is transparently bridged to the appropriate LAN interface. If the MAC address does not exist in the Passthrough Table, the packet is processed by higher layer functions, including the C-NAT/C-NAPT Transparent Routing functions.

The Upstream Selective Forwarding Switch (USFS) function prevents intra-home communications from affecting the HFC network and is in place primarily for the case in which devices in the home are publicly addressed and reside on different logical IP subnets.

The USFS routes traffic that is sourced from within the home network and is destined to the home network directly to its destination. LAN IP Device sourced traffic, whose destination IP address is outside the LAN address realm, is passed unaltered to the CAP bridging/routing functionality.

The USFS functionality makes use of the ipNetToMediaTable [RFC-2011], which contains a list of MAC Addresses, their corresponding IP Addresses, and PS Interface Index numbers of the physical interfaces with which these addresses are associated. The USFS will refer to this table in order to make decisions about directing the flow of LAN-to-WAN traffic. In order to populate the ipNetToMediaTable, the PS learns IP and MAC addresses and their associations. For every associated physical interface, the PS learns all of the LAN-Trans and LAN-Pass IP addresses along with their associated MAC bindings, and this learning can occur via a variety of methods. Vendor specific IP/MAC address learning methods may include: ARP snooping, traffic monitoring, and consulting DHCP table entries.

The USFS inspects all IP traffic received on PS LAN interfaces. If the destination IP address is found (via the ipNetToMediaTable) to reside on a PS LAN interface, the original frame's data-link destination address is changed from that of the default gateway address to that of the destination LAN IP Device, and the traffic is forwarded out the proper PS LAN interface. If a match to the destination IP address is not found in the ipNetToMediaTable, the packet is passed, in its original form, to the

C-NAT/C-NAPT transparent routing function or the Passthrough bridging function (depending on the active packet handling mode).

NAME RESOLUTION

CableHome 1.0 specifications define a basic name resolution service for devices connected to the subscriber's home LAN. This function, embodied in the CableHome Naming Portal (CNP), establishes a table of host names, client identifiers, and private IP addresses for home LAN elements. This feature allows the home user to refer to networked devices in the home by a human-readable name rather than by an IP address.

The CNP obtains host name and associated private IP address information for LAN IP Devices from the CDP LAN Address Table in the CDP MIB. The CDP LAN Address Table is populated when devices in the home are served by the CDS with private address leases.

When a DNS query is issued to the PS Element's DNS server IP address from a privately-addressed element in the home, the CNP refers to its table and if the requested host name is found, the CNP replies with the appropriate IP address. If the CNP does not locate the requested host name in its table, it replies to the querying device on the home LAN with the globally-routable IP address of the cable operator's DNS server. It is then the responsibility of the home LAN device to direct its query directly to the cable operator's DNS server for host name resolution.

Queries from devices in the passthrough realm, i.e., those devices served directly from the cable operator's headend DHCP and DNS servers, will be addressed directly to the cable operator's headend DNS server and will not be served by the CNP.

CableHome 1.0 requires compliance with standard DNS message formats described in [RFC 1034] and [RFC 1035].

SECURITY

The security in CableHome can vary widely depending upon which system the cable operator deploys and how the operator configures options for each system. The security considerations discussed here are intended to give the operator some insight into what security configuration settings exists for each system. Security settings will be discussed for

- Items that need to be configured in the back office prior to network operation
- PS Element security configuration during the provisioning process
- PS Element security configuration via SNMP
- Firewall configuration
- Secure Software Download

The first step to understanding the security considerations is for the operator to decide which network environment CableHome will be deployed upon. The CableHome specification was created to operate in three environments. These three environments are referred to as DOCSIS 1.0 system, DOCSIS 1.1 system, and CableHome Enhanced environment. A DOCSIS 1.0 system is a PS configured to operate in DHCP Provisioning Mode and NMAccess management mode. A DOCSIS 1.1 system is a PS configured to operate in DHCP Provisioning Mode and SNMPv3 Coexistence Mode for management messaging.

This discussion will point out security for each of these options and explain the relevant security setting for each.

Back Office Security Configurations

In the back office the cable operator will need to set up and configure the CableHome network, prior to running the CableHome system. There are several security considerations and configurations needed.

Configuration for secure software download is required in the DOCSIS 1.0 system for the stand-alone PS Element and for either the embedded or stand-alone PS Element in the DOCSIS 1.1 system. The operator must insert one or more CVCs into the PS configuration file to enable the secure software download. The operator may choose to apply for a service provider CVC with CableLabs, and must keep secret the private key associated with the CVC. The cable operator may then optionally insert their CVC in the configuration file to give the operator control over all software download images that will be accepted by the PS Element. The options for control of software download are discussed in more detail in the software download section. Secure software download provides the same security and uses the same method for all three systems.

The CableHome Enhanced environment requires security configuration for secure software, mutual authentication and secure management messages. For CableHome, secure software download requirements are the same as for DOCSIS 1.0 or 1.1 systems. To meet the requirement for mutual authentication, CableHome uses X.509 certificates and the Kerberos protocol with the PKINIT extension.

The cable operator must set up the KDC for Kerberos functionality. The KDC must have a KDC certificate issued by the Service Provider CA certificate, which means the cable operator must apply for a CableLabs Service Provider CA certificate. The KDC must also be provisioned with the CableLabs Manufacturer Root CA certificate and the Local System CA certificate if one exists. If the cable operator is planning on using the same KDC for PacketCable then the KDC will need the MTA Root CA Certificate as well. With the CA certificate comes the responsibility of running a public key infrastructure (PKI), and the cable operator will need to implement the appropriate security policies and procedures.

The Kerberos set up involves more than just provisioning the box with certificates. It involves planning for the Kerberos infrastructure, configuration of the infrastructure (multiple KDCs) and configuration of CableHome management message security for SNMPv3.

To create a Kerberos infrastructure several key issues need to be resolved. The Kerberos protocol has a master and slave relationship within for its server hierarchy. The cable operator will need to decide how many Kerberos realms are needed, the name of each realm, how many slave KDCs are needed and where they should be located, hostnames for each KDC and how to map the hostnames into the Kerberos realms. The operator will also need to set up and manage the Kerberos database.

The KDCs are then configured for the infrastructure settings along with the Kerberos message settings and the CableHome specific settings. For the KDC messages, encryption with DES3 and authentication with MD5 is required and may need to be configured. The CableHome Kerberos configurations include various message parameter options. For example the operator will need configuration for the desired Kerberos ticket duration.

The KDC supplies the PS Element with the initial information used to set up the SNMPv3 keys, one key for authentication and one key for privacy (encryption). For this reason, the KDC needs to be provisioned with the correct encryption and authentication algorithms for the KDC to negotiate with the PS Element within the Kerberos message exchanges on behalf of the SNMP manager in the NMS. SNMPv3 authentication is required, and must use a default value for the MD5 hash algorithm. The operator may support other hash algorithms and can add those to the list of acceptable or preferred hash algorithms. Encryption on management messages is currently optional and the operator will either need to list the null algorithm, if no encryption is desired, or the DES algorithm if encryption is needed. DES is currently the only SNMPv3 supported encryption algorithm. Both the KDC and the NMS will need to be configured to choose the appropriate hash and encryption algorithm for the NMS and PS Element to communicate securely.

PS Element Security Configuration During the CableHome Provisioning Process

Security within the PS Element provisioning process includes security on the information provided in the message exchanges, establishment of security configuration settings as a result of the information extracted from the messages, and completion of the mutual authentication process. Security for the message exchanges was set up in the back office configuration and the KDC and NMS were also appropriately provisioned with the security configuration settings needed to communicate with the PS Elements. The cable operator should not need to configure anything during the provisioning process itself.

On DOCSIS 1.0 and 1.1 systems within CableHome mutual authentication is not available. Security on management messages is possible in SNMPv3 Coexistence Mode if the operator uses SNMPv3. The cable operator will set up security for the PS Element the same way it is described in DOCSIS 1.1 specifications for the cable modem. The PS Element will receive its initial keying material for SNMPv3 from the Diffie-Hellman kick start. To provide the Diffie-Hellman kick start with all the appropriate parameters, the CableHome Administrator calculates the values for security name and public number and populates the usmDHKickstartTable with these values.

In SNMP provisioning mode, the PS Element completes mutual authentication with the KDC, key management for SNMPv3, configuration file settings, and configuration file security. There are three parts to authentication: the identity credential, the checking of the identity credential for validity and the common means to communicate the identity information. CableHome specifies an industry standard identification credential, X.509 certificates, in conjunction with RFC 2459 for certificate use, and Kerberos, which is a common communications protocol for mutual authentication. X.509 certificates are exchanged between the PS Element and the KDC during the Kerberos PKINIT exchange, which is wrapped in the AS Request and AS Reply messages. Each side validates the information in the certificate and verifies the certificate chain back to the CableLabs root for each chain. Once the trust has been established, the information for the SNMPv3 keys is sent from the KDC to the PS Element.

If an operator wishes to enable secure software download, the trigger is a CVC and must be sent in the PS configuration file. The operator has five choices for which CVCs to place in the configuration file. These are discussed in the secure software download section. The CVCs placed in the PS configuration file must match the CVCs sent in the code image download. Prior to placing any CVCs in the PS configuration file, the CVCs must be verified and validated as part of the back office security procedures.

The configuration file requires a SHA-1 hash to protect it from being changed in transit. The NMS will add the hash to the configuration file prior to placing it on the TFTP server for download. The PS element will check the hash prior to processing the configuration file.

SNMP PS Element Security Configuration

Some security information may be updated via SNMP. This allows the operator flexibility in managing the network and does not require the PS Element to get a new configuration file each time these items need to be updated. The operator may initiate and monitor secure software downloads, update CVC certificate information, and monitor firewall events via SNMP MIB variables.

Firewall Configuration

Firewall configuration follows the same method as specified for PS Element configuration. In DHCP provisioning mode, the cable operator provides information to trigger a firewall configuration in the PS configuration file. If the IP address of the firewall configuration file TFTP server, the firewall configuration file filename, the hash of the firewall configuration file, and the encryption key (if the firewall configuration file is encrypted) are included in the PS configuration file then the PS Element will request the firewall configuration file. After the firewall configuration file is received the hash of the configuration file is calculated and compared to the value received in the PS Configuration File. If encrypted, the file is decrypted. The file is then processed. In SNMP provisioning mode, the cable operator may trigger a firewall configuration by passing firewall configuration file information in either the PS configuration file or via SNMP, but the rest of the process is the same.

Firewall attack events are monitored via SNMP MIBs. The cable operator must configure the firewall attack time limit and maximum number of events allowed within that time limit. An event is logged if more than the maximum number of attacks occurs. The operator security policy and procedures manual should instruct the administrator with an appropriate response to attacks.

Firewall rule sets are not defined in CableHome. A variable is defined to track the rule set version and an additional variable allows the operator to enable or disable the security policy for the firewall.

Secure Software Download Configuration

The DOCSIS CM controls embedded downloads. If the PS Element is embedded with a DOCSIS cable modem, then the software image must be a single image and the download will be controlled by the cable modem according to the DOCSIS specifications. If the PS Element is not embedded with the cable modem, then the PS Element has its own code image and is responsible for the secure software download process as specified in CableHome.

Inclusion of CVC(s) in the configuration file to match the CVC(s) in the code image enable secure operator controlled downloads. The cable operator will enable secure software download to the PS Element by placing the location of the file, the filename and the CVC(s) in the PS configuration file. Once secure software download is enabled, the operator can trigger a software download by sending the location of the file and the filename from the NMS. The download will only be triggered if the filename does not match the name of the current code image in the PS Element. In the configuration file the operator will choose between the five possible CVC combinations to provide authorization for specific images according to their security policy.

- Option 1: Send the manufacturer's CVC
- Option 2: Send the manufacturer's CVC and CableLabs CVC.
- Option 3: Send CableLabs CVC
- Option 4: Send the manufacturer's CVC and the service provider's CVC
- Option 5: Send the service provider's CVC

The manufacturer's CVC and corresponding signature ensure the code has not been altered since it left the manufacturer's facility. The signature can legally bind the manufacturer to any claims made about this particular code image. The CableLabs signature means that the code has been through the certification testing program and was passed by the certification board. The CableLabs certified sticker on the outside of the box is now accompanied by the CableLabs digital signature over the certified code image for certified product. The operator's signature will bind the signed code image according to the operator's security policy. Code images will only be downloaded into the PS Elements according to the operator's security policy, and it is up to the operator to define if the manufacturer, CableLabs, or operator approved code is appropriate for the PS Elements on their network.

Back office CVC validation is necessary. Prior to providing images for secure software download the operator must 1) validate and verify the CVC(s) and signature(s)on the code image, 2) validate and verify the CVC(s) prior to placement in the configuration file, 3) make sure the CVC(s) to be placed in the configuration file match those in the code image, 4) insert the CVC(s) into the configuration file, and if using the service provider CVC, then 5) attach the service provider CVC to the code image and sign the code image. The configuration file and code file is then ready to be placed on the TFTP server for the provisioning process to start. It is imperative for the PS to check the software image CVC(s) and signatures prior to use to make sure the code file has not been tampered with and to ensure the code has come from a trusted source.

The operator can either upgrade the PS Element to the next authorized version of the code or revert to a previous version of the authorized code at any time using the same process. To revert to a previous version of code, the manufacturer signing time on the code must either be equivalent to or later than the current version of the code installed in the PS Element. Logistically this means the cable operator will need to track signing times on all the valid code images for compliance to the security policy.

If secure software download fails for any reason, an event will be sent to the NMS and the cable operator will need to decide on the appropriate course of action.

There are many other aspects of security for CableHome that are not discussed in this paper. The goal was to address from a high level some of the items the cable operator will need to configure in order to deploy and maintain a secure CableHome network. It is critical for the cable operator to create and use a security policy and procedures manual for all aspects of its network access, information storage, technical configuration, security breaches, security auditing and reviews as well as the daily maintenance of security technology.

REFERENCES

CableHome 1.0 Specification, CH-SP-D01-020131, Cable Television Laboratories, Inc., January 31, 2002

PacketCable Security Specification, PKT-SP-SEC-I05-020116, Cable Television Laboratories, Inc., January 16, 2002, http://www.PacketCable.com./

Kerberos V5 Installation Guide, Release 1.2, Document Edition 1.1, Massachusetts Institute of Technology, January 9, 2002, http://web.mit.edu/kerberos/www/krb5-1.2/krb5-1.2.3/doc/install.html#SEC7

CableHome PSDEV MIB Specification, CH-SP-MIB-PSDEV-D02-020131, Cable Television Laboratories, Inc., January 31, 2002

CableHome CAP MIB Specification, CH-SP-MIB-CAP-D02-020131, Cable Television Laboratories, Inc., January 31, 2002

CableHome CDP MIB Specification, CH-SP-MIB-CDP-D02-020131, Cable Television Laboratories, Inc., January 31, 2002

CableHome CTP MIB Specification, CH-SP-MIB-CTP-D02-020131, Cable Television Laboratories, Inc., January 31, 2002

CableHome Security MIB Specification, CH-SP-MIB-SEC-D02-020131, Cable Television Laboratories, Inc., January 31, 2002

RFC 1034, IETF, Domain Names - Concepts and Facilities, November 1987.

RFC 1035, IETF, Domain Names - Implementation and Specification. November 1987.

RFC 2011, IETF, SNMPv2 Management Information Base for the Internet Protocol Using SMIv2, November 1996.

RFC 2459, IETF, Internet X.509 Public Key Infrastructure Certificate and CRL Profile. January 1999.

RFC 2574, IETF, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMP), April 1999.

RFC 2575, IETF, View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP), April 1999.

RFC 2576, IETF, Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework, March 2000.

RFC 2669, IETF, DOCSIS Cable Device MIB - Cable Device Management Information Base for DOCSIS Compliant Cable Modems and Cable Modem Termination Systems, August 1999.

AN ANALYSIS OF THE TDMA AND S-CDMA TECHNOLOGIES OF DOCSIS 2.0

Fabien Buda, Emmanuel Lemois, and Hikmet Sari
Juniper Networks
4-14 rue Ferrus, 75014 Paris, France

Abstract

This paper analyzes the performance of TDMA and S-CDMA, which are the two modula-tion and multiple access techniques included in the newly developed DOCSIS 2.0 specifications. The two techniques are analyzed in the presence of linear signal distortion, ingress noise, and burst noise, which characterize cable plants. We also discuss how the transmission parameters can be adopted in both multiple access techniques to plant conditions in order to increase capacity or robustness. The results give a good summary of the relative merits of TDMA and S-CDMA.

1. INTRODUCTION

Data over Cable System Interface Specifications (DOCSIS) [1] elaborated under the leadership of Cable Television Laboratories, Inc. (CableLabs) have established themselves as the major industry standard for two-way communications over hybrid fiber coax (HFC) cable plants. The first specifications, referred to as DOCSIS 1.0, have been largely deployed in cable networks, and the second specifications (DOCSIS 1.1) are now in the early phase of deployment. The physical (PHY) layer of DOCSIS 1.1 specifications is the same as that of DOCSIS 1.0, the difference between the two sets of specifications lies on the medium access control (MAC) layer, which includes a quality of service (QoS) in DOCSIS 1.1 specifications.

The demand for increased capacity on the upstream channel has stimulated CableLabs to draft a new set of RF interface specifications referred to as DOCSIS 2.0. These specifications, which were developed with the vendor community, also aim at increasing the robustness to various impairments on the upstream channel in cable plants. The basic features of the new specifications include an increased channel bandwidth (6.4 MHz), several additional modulation schemes including 64-QAM which gives a 50% increase of the throughput with respect to the 16-QAM modulation that is in the current DOCSIS 1.0 and 1.1 specifications, and an improved forward error correction scheme.

One of the most controversial issues during the elaboration of the DOCSIS 2.0 specifications was the multiple access scheme, which describes how different cable modems (CMs) access a particular physical channel. The two techniques in compe-tition were the conventional time-division multiple access (TDMA) used in DOCSIS 1.0 and 1.1 as well as in many other international standards, and the synchronous code-division multiple access (S-CDMA) scheme used in some proprietary systems. The decision was eventually made to include both techniques in the specifications, and make both of them mandatory for both the CM side and the cable modem termination system (CMTS) side.

The two multiple access technologies included in DOCSIS 2.0 have very basic differences in terms of their operating principles and robustness to channel impairments and equipment imperfections. This paper gives the results of an exhaustive perfor-mance evaluation of DOCSIS 2.0 TDMA and S-CDMA in hybrid fiber/coax (HFC) cable networks. The evaluation is performed in the presence of most common impairments in cable plants which include linear signal distortion, additive white Gaussian noise (AWGN), ingress noise (characterized as narrowband interference) and burst noise. We also discuss how the transmission parameters can be adopted to plant conditions in order to increase capacity or robustness.

The paper is organized as follows: In Section 2, we give a brief review of DOCSIS 2.0. Section 3 discusses the noise performance, and more particu-larly the trade-off between data rate and perfor-mance, and the optional trellis code in the S-CDMA mode. In Section 4, we study the influence of channel impairments including linear distortion, ingress noise, and burst noise. Finally, we give our conclusions in Section 5.

2. A BRIEF DESCRIPTION OF DOCSIS 2.0

Working with the vendor community, CableLabs has recently released the interim DOCSIS 2.0 RF interface specifications [2], which aim at increasing the capacity and robustness to various impairments of the upstream channel in cable plants. These specifications do not affect the downstream channel, neither do they affect the medium access control (MAC) functions except for the changes required to accommodate the new physical (PHY) layer. The basic features of the new specifications include an increased channel bandwidth (6.4 MHz), several additional modulation schemes including 64-QAM which gives a 50% increase of the throughput with respect to the 16-QAM modulation that is in the current DOCSIS 1.0 and 1.1 specifications, and an improved forward error correction scheme.

The previous DOCSIS specifications (DOCSIS 1.0 and 1.1) were based on time-division multiple access (TDMA). More specifically, time-division multiplexing (TDM) is used on the downstream channel (from CMTS to CMs) and TDMA is used on the upstream channel (from CMs to CMTS). For DOCSIS 2.0, two multiple access proposals were made. One of these is to keep the TDMA used in previous DOCSIS specifications, and the other is synchronous code-division multiple access (S-CDMA), previously used in some proprietary modems. The decision was made in August 2001 to include both schemes in DOCSIS 2.0.

2.1. TDMA

TDMA is a simple and popular multiple access technique used today in many international stan-dards and proprietary systems. It consists of assig-ning different time slots to different users. During the time slot assigned to one user, all other CMs remain silent, and therefore, there is no interference between users. Since DOCSIS 1.x is based on TDMA, the use of this technique in DOCSIS 2.0 is natural and its implementation requires little effort. Furthermore, since DOCSIS 2.0 requires backwards compatibility with DOCSIS 1.x, implementation of TDMA in DOCSIS 2.0 equipment is unavoidable.

The upstream spectrum in cable networks extends from 5 MHz to 42 MHz (to 65 MHz in Europe). The upstream channel width in DOCSIS 1.x is 2^{n-1} times 200 kHz, with n = 1, 2, 3, 4, 5. DOCSIS 2.0 adds a 6.4 MHz channel bandwidth corresponding to n = 6. In all cases, the raised-cosine Nyquist roll-off factor is 0.25, and the symbol rate is 0.8 times the channel bandwidth. A number of upstream channels may be available in the 5 – 42 MHz upstream spectrum depending on what part of this spectrum is used for legacy (analog or non-DOCSIS digital) channels and what part of it is not usable due to excessive noise, interference, or distortion. TDMA on the cable upstream channel is actually a combination of frequency-division multiple access (FDMA) and TDMA. When a CM makes a resource request, the CMTS grants time slots on one of the upstream channels, each of which accommodates a number of CMs in the TDMA mode. That is, the multiple access scheme is actually

FDMA/TDMA, but it is referred to as TDMA for simplicity.

The upstream modulation schemes in DOCSIS 1.x are the quaternary phase-shift keying (QPSK) and the quadrature amplitude modulation (QAM) with 16 constellation points (16-QAM), which doubles the data rate. To these, DOCSIS 2.0 adds 8-QAM, 32-QAM, and 64-QAM, but only the latter modu-lation is mandatory at the CMTS. With respect to 16-QAM, this modulation increases the data rate by 50% but loses 6 dB in signal-to-noise ratio (SNR), defined as the transmitted energy per symbol to the noise spectral density ratio (E_s/N_0). The inclusion of 8-QAM and 32-QAM in the specifications allows reducing the granularity in the trade-off between data rate and performance.

Since DOCSIS 2.0 also aims at increasing robust-ness, it extends the error correction capacity of the Reed-Solomon (RS) code from 10 to 16 bytes per block. Furthermore, a byte interleaver is included in the TDMA specifications so as to break the error events caused by burst noise and uniformly distri-bute the resulting symbol errors. The interleaver is a block interleaver whose length is equal to the RS block length and depth is a configurable parameter, which distributes the error bursts over a selected number of RS blocks.

2.2. S-CDMA

The S-CDMA of DOCSIS 2.0 is actually not a true CDMA, but rather a mix of code-division multip-lexing (CDM), CDMA, and TDMA. The incoming data is organized in mini-slots, which have two dimensions (spreading codes and time). The time duration of mini-slots is one S-CDMA frame that spans a programmable number of S-CDMA symbol intervals. (The maximum frame length is 32 S-CDMA symbol intervals.) Symbol spreading is performed through multiplication by a spreading code (spreading sequence) of 128 chips taken from a set of 128 orthogonal codes that are generated by a quasi-cyclic shift. Spreading is in the time domain, which means that an S-CDMA symbol interval is equal to 128 TDMA symbol intervals.

A mini-slot contains a programmable number of spreading codes, which can be as low as 2 and as high as 128. A mini-slot contains symbols from a single CM. Suppose that the number of codes per mini-slot is 4 and that the frame length is 16. Then, the mini-slot contains 64 symbols, and a given code is assigned to the same user for a time duration of 16 consecutive S-CDMA symbols. The 4 symbols transmitted in parallel within the same mini-slot in this example are code-division multiplexed. If all mini-slots of an S-CDMA frame are assigned to the same cable modem, S-CDMA is reduced to pure CDM during that interval.

To the other extreme, if the mini-slots contain two codes only, and each mini-slot is assigned to a different CM, then there is code-division multiple access between 64 CM signals during that frame. The spreading code orthogonality ensures that there is no interference between symbols transmitted in parallel by the same CM, since these symbols are perfectly time synchronized. But interference arises between signals generated by different CMs, due to non-ideal timing synchronization. To limit the resulting degradation, DOCSIS 2.0 specifies that the maximum timing error between a CM and the CMTS must not exceed 1% of the chip interval.

In addition to the RS code, S-CDMA specifications also include trellis-coded modulation (TCM) as an option. Trellis coding reduces the number of information bits per symbol by one, and so trellis-coded 64-QAM (referred to as 64-TCM) is equivalent to uncoded 32-QAM in terms of spectral efficiency. Therefore, S-CDMA specifications also include 128-TCM, which is strictly equivalent to uncoded 64-QAM as far as information bit rate is concerned. But the S-CDMA specifications neglect to include an interleaver between the external RS code and the internal trellis code, which reduces the benefit of trellis coding. In contrast, S-CDMA specifications include some interleaving after the trellis encoder to reduce the effect of burst noise. This interleaver operates on subframes, and interleaving is different for uncoded bits and coded bits. Subframes are independent of mini-slots, and a subframe is always contained within a single frame.

Symbol spreading in S-CDMA is used for the traffic mode only. That is, the spreader is turned off (S-CDMA is deactivated) during initial ranging and periodic station maintenance. The implication of this is that whatever this multiple access technique may offer with respect to TDMA, the benefit is restricted to the traffic mode.

3. NOISE PERFORMANCE

3.1. Flexibility of the Standard

DOCSIS 2.0 is a toolbox that gives full flexibility in trading off data rate against performance and robustness to channel impairments. Indeed, it includes all QAM signal constellations from 4-QAM (QPSK) up to 64-QAM, and also the RS code correction capability can take all values from *RSt = 0* (no coding) up to *RSt = 16*, for different block lengths. First, it is well known that with respect to 2^m-QAM, the 2^{m+1}-QAM signal constellation increases the number of bits per symbol by one, but requires 3 dB higher signal-to-noise ratio (SNR) to achieve the same bit error rate (BER) performance.

This means that the constellation alone offers a trade-off between data rate and performance with a granularity of 3 dB in SNR or transmitted signal power. The only exception to this general rule is the step between 4-QAM (QPSK) and 8-QAM, which is approximately 4 dB. For example, if the cable network is operating with 64-QAM and the noise level increases, then switching to 32-QAM increases noise margin by 3 dB, and switching to 16-QAM increases it by 6 dB. In practice, it is desirable to have a finer granularity in this trade off. The second set of parameters that make this possible are the block length and error correction capacity of the RS code. Since there is a large degree of freedom in selecting the code parameters, the granularity of the trade off between data rate and performance can be made as fine as desired.

At this point, it is instructive to discuss the additional flexibility in the S-CDMA mode. As presented in Subsection 2.2, the total number of spreading codes is 128, but this number can be arbitrarily reduced in order to improve performance. This property of S-CDMA is actually often put forward as an advantage with respect to TDMA. While it is correct that S-CDMA allows this type of trade off, it is of little interest as it sacrifices too much data rate compared to what is offered by the modulation and code parameters. Indeed, reducing the number of codes from 128 to 64 in S-CDMA gives an SNR gain of 3 dB, but this is achieved by sacrificing 50% of the data rate. If we assume that the network is using 64-QAM, the same gain is achieved by switching down to 32-QAM, and this gain only involves a 16.6% of the data rate. In fact, the noise margin can be increased by 9 dB through modulation (by switching from 64-QAM to 8-QAM) if one is prepared to sacrifice 50% of the data rate.

Fig. 1 shows the spectral efficiency (number of information bits per symbol) vs. the SNR required for BER = 10^{-8} for different signal constellations, RS code parameters, and also the number of spreading codes in S-CDMA. Rather than giving some discrete points corresponding to a few values of the *RSt* parameter, this figure gives for each constellation a continuous curve corresponding to all *RSt* values from *RSt = 0* to *RSt = 16* with an RS block length of 255, and then to lower RS block sizes with *RSt = 16*. For uncoded 64-QAM and RS-coded QPSK, the figure also gives a curve that illustrates performance vs. spectral efficiency in S-CDMA with a reduced number of codes.

This figure clearly shows that for each signal constellation, there is a range of the number of information bits per symbol over which that constellation is the best to use. For example, the 64-QAM constellation is optimum down to 4.8 bits per symbol, and 32-QAM is optimum between 4.8 and 3.8 bits per symbol. It is also clear from this figure that reducing the number of spreading codes in S-CDMA to improve performance is not attractive as long as such a compromise can be made by the modulation and coding functions.

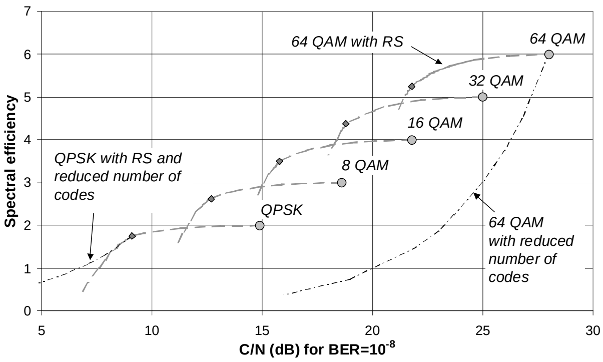

Fig. 1: SNR required with different signal constellations, RS code rates, and number of spreading codes in S-CDMA.

The only case where reducing the number of spreading codes in S-CDMA may be of interest is when the lowest level modulation (QPSK) together with *RSt = 16* in the RS code are not sufficient to get the desired performance. But even if this were a situation of strong practical interest, the current S-CDMA specifications would not help, because symbol spreading is disabled during initial ranging and periodic station maintenance, which are most critical to overall system performance.

3.2. Optional Trellis Coding

As mentioned earlier, the DOCSIS 2.0 specifica-tions include an optional TCM in the S-CDMA mode. The trellis code is an 8-state code, which gives an asymptotic coding gain of 4 dB. But what is more significant is the gain provided by the trellis code when cascaded with the mandatory RS code. Depending on the *RSt* parameter of the RS code, a BER of 10^{-2} - 10^{-4} at the TCM decoder output is reduced to less than 10^{-8} by the RS decoder. The asymptotic gain of the concatenated coding scheme is therefore the TCM coding gain somewhere in the range of 10^{-2} - 10^{-4}. With the trellis code used, this gain is 1.1 dB for *RSt = 16* and 2.5 dB for *RSt = 2*.

But the coding gain indicated above assumes an infinite interleaver between the RS encoder and the trellis encoder, so that the deinterleaver on the receive side can break the error events at the TCM decoder output and distribute them over a large number of RS blocks. The number of RS symbol errors per block is then small, and these errors can be located and corrected by the RS decoder. Unfortunately, the DOCSIS 2.0 specifications do not include such an interleaver but instead a sub-frame interleaver that keeps a part of the information bits uninterleaved. This significantly reduces the TCM gain, as illustrated in Fig.2 for 128-TCM. The RS code used in this figure is of length 255 and its correction capacity is *RSt = 16*.

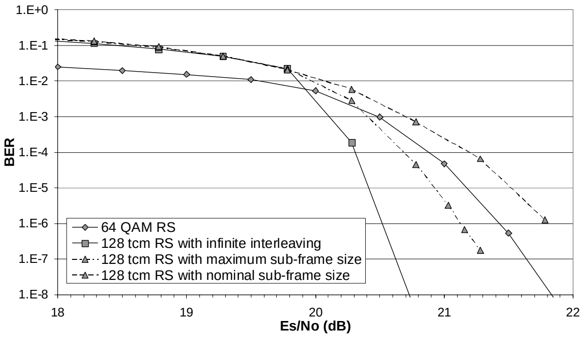

Fig. 2: BER performance of 128-TCM.

We can see that at the BER of 10^{-8}, 128-TCM gains 1.1 dB with respect to 64-QAM when it employs an infinite interleaver prior to the trellis encoder. With the subframe interleaver of DOCSIS 2.0, the gain is a function of the subframe size. With the maximum subframe size (equal to one frame), the gain is 0.4 dB, and with the nominal subframe size (equal to one RS block) 128-TCM actually loses 0.5 dB with respect to 64-QAM.

4. CHANNEL IMPAIRMENTS

The three major impairments on the upstream channel in HFC networks are microreflections, narrowband ingress noise, and burst noise [3], [4]. In this section, we analyze the performance of the DOCSIS 2.0 TDMA and S-CDMA in the presence of these impairments.

4.1. Channel Microreflections

A common model for the microreflections on the cable upstream channel (see [3], [4]) is to use an amplitude of –10 dBc up to 0.5 μs, –20 dBc up to 1.0 μs, and –30 dBc beyond 1.0 μs. To simplify the simulations, we used a discrete channel model with 3 echoes, where the first echo is 10 dB below the main signal path and has a delay of one symbol period T, the second echo is 20 dB down and has a delay of 2T, and finally the third echo is 30 dB down and has a delay of 3T. Using this model, we have investigated the impact of channel microreflections on TDMA and S-CDMA, and the results are shown in Fig. 3 for uncoded QPSK. These results show that the difference between the two multiple access techniques is rather small, but the SNR degradation due to intersymbol interference (ISI) is smaller in TDMA even when the number of codes is reduced by 50% in S-CDMA.

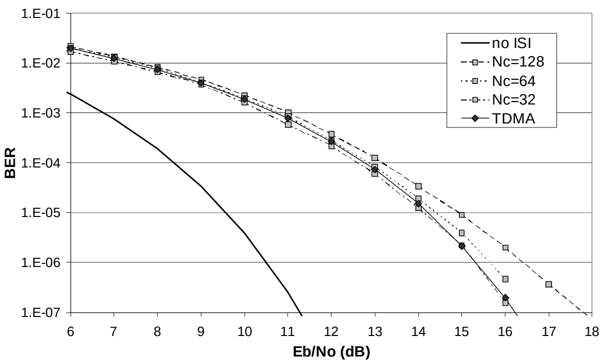

Fig. 3: Impact of channel microreflections on TDMA and S-CDMA.

In the ranging mode, the CMTS identifies the upstream channel impulse response, computes the optimum equalizer coefficients, and sends them to the corresponding CM. But the pre-equalizer setting is never perfect, and there is always some residual ISI at the threshold detector input of the upstream demodulator. Note that any residual ISI is easily handled by the equalizer in a TDMA burst receiver, by identifying the channel impulse response during the preamble. But the individual signals transmitted in parallel by different CMs in S-CDMA having different distortions, there is no way to cancel the residual ISI before the despreader. Furthermore, any processing after the despreader would involve an extremely complex multiuser detector.

4.2. Ingress Noise

Ingress noise originates from man-made sources like shortwave AM and HF radio emissions which leak into the cable and affect the upstream channel. It is modeled as narrowband interference, generally of 20 kHz bandwidth. In a 3.2 or 6.4 MHz wide upstream channel, typically two or three interferers can be encountered, and the carrier-to-interference ratio (CIR) can be as low as 0 dB.

The influence of narrowband interference on code-division multiple access (CDMA) signals is an issue on which there are erroneous ideas in the technical community, because CDMA is often assimilated to direct-sequence spread spectrum (DS-SS) signaling. But as indicated in [5] and [6], the TDMA vs. CDMA issue (whether CDMA uses orthogonal spreading codes or pseudo-noise codes) can not be assimilated to DS-SS vs. non-DS-SS signaling. It is shown in these papers that CDMA and TDMA have the same performance in terms of the CIR at the threshold detector input, and that TDMA actually has better BER performance.

The reason for this surprising result is that while the interference at the threshold detector input has a quasi-uniform amplitude distribution in TDMA, the despreading operation makes it Gaussian in CDMA, and the Gaussian distribution is more sensitive than the uniform distribution in the range of interest. In fact, the most robust multiple access technique to narrowband interference is orthogonal frequency-division multiple access (OFDMA) [7] which can inhibit the carriers that are affected by interference.

In conventional TDMA receivers, compensation of narrowband interference is performed by means of notch filters, and the resulting ISI is compensated using an equalizer. One problem with this approach is it mixes the channel distortion and ingress noise problems, ignoring the fact that the former depends on the transmitting CM while the other does not. A better approach, implemented in the advanced CMTS receiver presented in [8], is described in [9]. It consists of implementing a noise prediction filter to estimate the ingress noise and subtract it from the received signal, and an adaptive equalizer optimized under the zero-forcing criterion to estimate the channel impulse response, and compute and send the pre-equalizer coefficients to the CM of interest. In this receiver structure, the noise prediction filter coefficients are saved from burst to burst, while the pre-equalizer coefficients are recomputed at each burst using the preamble.

Performance of the ingress noise canceller of [9] is illustrated in Fig. 4. This figure shows the influence of three 20-kHz wide interferers of individual power of –15 dBc on 16-QAM, and their compensation using the ingress noise canceller. We can see that the system operates at the BER of 10^{-5} with an SNR degradation limited to 2.1 dB, which is an outstanding performance for this environment.

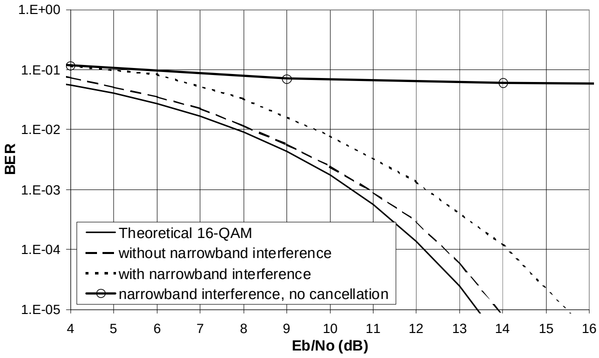

Fig. 4: Effect of narrowband interference on 16-QAM TDMA and its compensation.

In S-CDMA, the ingress noise becomes wideband (and its samples uncorrelated) at the despreader output, and therefore it must be cancelled before the despreading operation. Although this is possible in principle, noise prediction using the receiver decisions in this case involves a significant delay and cannot be made as reliable as in TDMA.

4.3. Burst Noise

Burst noise in cable plants occur mostly at frequen-cies below 10 MHz. We will use a simple model where bursts occur periodically at a predetermined repetition rate. The three parameters in this model are the burst duration, the burst amplitude, and the repetition frequency.

The basic countermeasure against burst noise is error correction coding with interleaving. Such a mechanism is included in both modes of DOCSIS 2.0. In TDMA, noise immunity is determined by the RS code and the byte interleaver. The RS code parameters are the number of information bytes *RSk* and the number of correctable byte errors *RSt*. The block length *N* is given by *N* = *RSk* + *2RSt*. The block interleaver parameters are its width *IntW* and its depth *IntD*. The interleaver width is equal to the RS block length. At the interleaver, the input data is written row by row, and read column by column. The deinterleaver performs the inverse operation.

The bytes affected by a noise burst of *TBurst* seconds will be located in *NC* columns in the deinterleaver. *NC* is related to the symbol rate *R* and the number of bits per symbol *m* by the relation

$$NC = \left\lceil \frac{(TBurst.R + 2)\times m}{8\times IntD} \right\rceil$$

where $\lceil x \rceil$ denotes the smallest integer larger than or equal to *x*. Since burst noise is generally assumed to have a large power with respect to the useful signal, we can assume that the codeword error rate is 0 if $NC \le RSt$ and 1 if $NC > RSt$.

In S-CDMA, in addition to the RS code and to interleaving, mitigation of burst noise also benefits from signal despreading. But as mentioned earlier, S-CDMA does not use the byte interleaver of TDMA, but instead it uses symbol interleaving over subframes of height (number of spreading codes) *F*. Assuming that each subframe contains one codeword, the bytes affected by burst noise are uni-formly distributed over the codewords of the frame. In the absence of trellis coding, it can be shown that the number of affected bytes per codeword is

$$Nbytes = \left\lceil \left\lceil \frac{TBurst.R + 2}{128} \right\rceil \times \frac{F.m}{8} \right\rceil$$

But note that *Nbytes* may exceed *RSt* in S-CDMA, while still ensuring an acceptable codeword error rate. The latter is given by

$$CwER = \sum_{i=0}^{\max(Nbytes,RSt)} C_{Nbytes}^{i}\, ByteER^{i}\left(ByteER\right)^{(Nbytes-i)}$$

where *ByteER* is the error rate of the bytes affected by burst noise, taking into account the despreading gain on noise power. With optional TCM, calculation of the error rate is very involved, and therefore we resorted to computer simulation. Fig. 5 shows the BER results obtained using 16-QAM with no trellis coding and 32-TCM. The burst noise power in these simulations was 0 dBc and the subframe height was 15.

The results show that the BER is lower with 16-QAM, which exhibits a very sharp drop when the burst duration gets shorter than 256 chips (2×15 S-CDMA symbols per RS codeword). The higher BER of 32-TCM is due to the behavior of the TCM decoder, which actually leads to error events of increased length with respect to that of the noise burst. This suggests that the TCM option should be disabled to operate in the presence of burst noise.

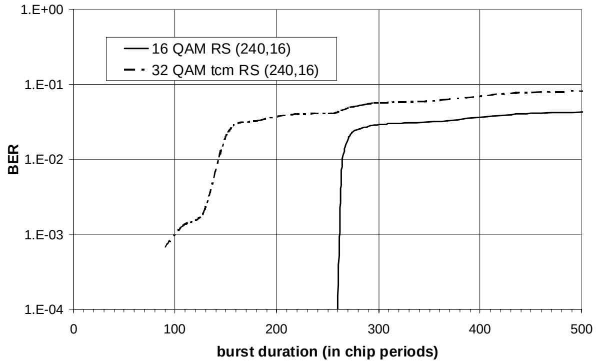

Fig. 5: Performance of TCM with no interleaving in the presence of burst noise.

Next, performance of TDMA and S-CDMA was investigated using the impulse repetition rate vs. impulse duration leading to a BER of 10^{-8} as criterion. The results are depicted in Fig. 6 for 16-QAM and 1518-byte and 64-byte packets. The RS code in these calculations was RS(250, T=16) for the long data packets, and RS(74, T=5) for the short packets. The depth of the byte interleaver used in TDMA was *IntD = 7* for long packets and *IntD = 1* for short packets. Finally, the frame length of S-CDMA was *K = 32*, and the subframe height was *F = M/K*, where *M* is the number of QAM symbols per RS block. Fig. 6 also indicates the minimum and maximum values of the burst repetition rate vs. burst duration that are encountered in cable plants.

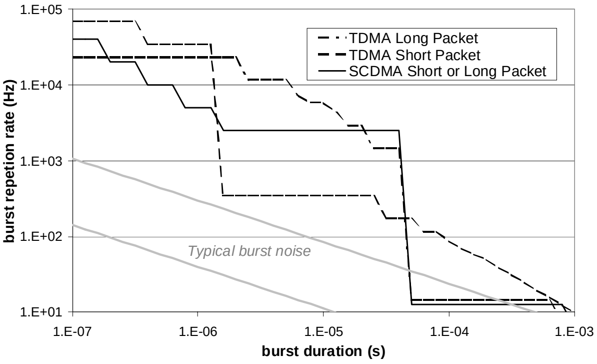

Fig. 6: Performance of TDMA and S-CDMA in the presence of burst noise.

We observe that with exception of a small region near the burst duration of 10^{-4} seconds, both TDMA and S-CDMA can cope with the burst noise of this model. The other observation is that the comparison of TDMA and S-CDMA in this environment is a function of the burst noise characteristics and also of the packets length.

5. CONCLUSIONS

We have analyzed the performance of DOCSIS 2.0 TDMA and S-CDMA on AWGN channels and in the presence of channel impairments, which include microreflections, narrowband ingress noise, and burst noise. Our analysis showed that reducing the number of spreading codes in S-CDMA to increase performance loses too much in data rate to be of practical interest. The set of constellations and RS code parameters included in the standard give the desired flexibility in trading off performance against useful data rate. Also, trellis code option without interleaving gives little improvement on one hand and degrades performance in the presence of burst noise on the other hand. Finally, TDMA and S-CDMA have similar sensitivities to channel impairments, but it is easier to compensate for microreflections and narrowband ingress noise in burst TDMA receivers than in S-CDMA receivers.

REFERENCES

[1] Data-over-Cable Service Interface Specifica-tions – RF Interface Specification, SP-RFI-I05-991105, Cable Television Laboratories, November 1999, Louisville, Colorado.

[2] Data-over-Cable Service Interface Specifica-tions – RF Interface Specification, SP-RFIv2.0-I01-011231, Cable Television Laboratories, Nov. 1999, Louisville, Colorado.

[3] B. Currivan, "Cable Modem Physical Layer Specifications and Design," *in Cable Modems: Current Technologies and Applications*, *John Fijolek & al.* (*Editors*), pages 135-142, IEEE Press, 1999.

[4] T. J. Kolze, "An Approach to Upstream HFC Channel Modeling and Physical-Layer Design," *in Cable Modems: Current Technolo-gies and Applications, John Fijolek & al.* (*Editors*), pages 135-142, IEEE Press, 1999.

[5] H. Sari, F. Vanhaverbeke, and M. Moeneclaey, "Extending the Capacity of Multiple Access Channels," IEEE Communications Magazine, vol. 38, no.1, pp. 74-82, January 2000.

[6] M. Moeneclaey, M. Van Bladel, and H. Sari, "Sensitivity of Multiple Access Techniques to Narrowband Interference," IEEE Transactions on Communications, vol. 49, no. 3, pp. 497-505, March 2001.

[7] H. Sari and G. Karam, "Orthogonal Frequency-Division Multiple Access and its Application to CATV Networks," European Transactions on Telecommunications & Related Technologies (ETT), vol. 9, no. 6, pp. 507-516, November/ December 1998.

[8] F. Buda et al., "Design and Performance of a Fully-Digital DOCSIS CMTS Receiver," 2001 NCTA Technical Papers, pp. 212-220, June 2001, Chicago, Illinois.

[9] A Popper, F. Buda, and H. Sari, "An Advanced Receiver with Interference Cancellation for Broadband Cable Networks," Proc. 2002 International Zurich Seminar on Broadband Communications, pp. 23.1-23.6, February 2002, Zurich, Switzerland.

ANALYSIS OF BANDWIDTH-CONSERVATIVE SERVICE OPPORTUNITY

Gregory E. Feldkamp, Ph.D.
Vice President – Engineering and Development
@Security Broadband Corporation

Abstract

A significant challenge for cable operators is to leverage the high-speed data infrastructure to implement services for which customers are willing to pay an additional monthly charge without the burden of a significant increase in capital costs. @Security Broadband's SafeVillage^{TM} system (patents pending) creates such an opportunity through a residential monitored security service, which provides audio and video alarm verification by monitoring personnel in a centralized center, and live viewing and two-way audio by homeowners over the public Internet.

A critical design consideration has been the capacity implications of carrying streaming video and audio over the high-speed data infrastructure and thus the impact on the cable operator's capital expenditures. To characterize this impact, @Security conducted an extended study as part of an 85-home technology trial carried out in cooperation with a major cable operator. This paper describes the data collection and analysis methodology, and provides a quantitative description of the findings.

INTRODUCTION

The introduction and rapid growth of high-speed data (HSD) services have been enabled by the large investment by cable operators in two-way, hybrid fiber-coax (HFC) systems. Currently, these services center largely on public Internet access.

In addition to the growing penetration of HSD service, significant increases in per-subscriber usage are occurring. While growth in penetration generally is accompanied by corresponding increases in revenue, growth in per-subscriber traffic does not produce commensurate revenue growth for the cable operator. Bandwidth-intensive applications deployed by end users such as peer-to-peer music and video file sharing and network-enabled gaming produce no incremental service revenues, but increase cost by accelerating the need for capacity expansion through fiber node splits, upgrades to modulation hardware, or additional DOCSIS channel assignments.

Opportunities to create additional revenue streams through services provided over the high-speed infrastructure are attractive to cable operators, especially services with a relatively modest impact on upstream data-carrying capacity. Potential synergies from bundling, and possible reductions in HSD churn, are also of interest.

The SafeVillage^{TM} system was designed with those considerations in mind. It leverages the HSD infrastructure to provide a service comprising:

- Video and two-way audio enabled, centrally monitored residential security.

- Web-based communication using video and two-way audio between family members in the home and those at remote locations such as their places of work.

The service is designed to do this in a manner that minimizes the impact on HSD bandwidth utilization.

Because implementation reality often diverges from design intent, it was important to verify the bandwidth performance of the service through field measurements. This was done as part of a technology trial of the service hosted by Cox Communications using systems installed in 85 homes in the Las Vegas area. The methodology and results of this bandwidth utilization study are described below.

SERVICE OVERVIEW

The centrally monitored security aspect of the service allows abnormal situations in the home, such as intrusions, to be detected. The triggering event, which might be a trip of a motion detector, is reported to a central monitoring station staffed on a 24x7 basis by professional security operators. Upon reception of an alarm report, the operator must make a timely decision as to whether to dispatch police, fire, or emergency services personnel in the appropriate jurisdiction. To aid in this decision process, the service provides a live video feed from the home to the operator. The operator also can review potentially relevant recorded clips such as video from an outdoor camera overlooking the front door. In addition, the operator can use the system to listen to the home and, if appropriate, talk to people on the premises to verify the nature of the situation.

Authenticated members of the homeowner's family also can use the system on a demand basis to view the home from a Web-connected PC. In the technology trial, audio communication for such customer access was not supported, though two-way audio is supported in the commercial service. The customer also can select cameras and perform other control actions remotely.

Privacy and information security are extremely important elements of the service but are not discussed in this paper because they have little bearing on the question of HSD bandwidth utilization. Similarly, account and service management functions are not described.

DESIGN GOALS

System design goals with significance for bandwidth usage included:

- Provide video and audio suitable for use in rapid assessment of alarm situations at the central monitoring station (i.e., acceptable resolution, quality, and frame rate).
- Ensure that upstream traffic (alarm, control, audio, video, heartbeat) fits within a channel rate - limited to 128 Kbps.
- Limit the length of on-demand video viewing sessions to prevent abuse (e.g., to prevent continuous streaming to the customer's PC at work).
- Avoid excessive management traffic that contributes to background load on the HSD system.

SYSTEM ARCHITECTURE

An overview of the technology trial system architecture appears in Figure 1.

The service-enabling equipment installed at the customer premises included:

- A cable modem, access to which was shared with the customer's PC.
- Wireless motion detectors and door/window contact sensors.
- Video cameras.
- A user control unit for arming/disarming.
- An audio intercom and a wireless event receiver, both of which were integrated with the user control unit.
- A video transmission unit (VTU) for video capture and compression, audio analog/digital conversion, and video, audio and alarm communications over the IP infrastructure.

At the cable headend, in addition to the HSD Cable Modem Termination System, a digital video recorder (DVR) was installed. The DVR received and forwarded alarm event information reported from the customer premises equipment; received, stored, and forwarded video and audio transmitted to and from the customer premises; and maintained heartbeat communication with the equipment at the premises.

At the central monitoring station in Orlando, Florida, an automation system provided monitoring and business application support to security operators. The automation system received alarm event reports from the home and displayed them to the operators at PC-based workstations along with site and emergency contact information on a 24x7 basis. The operator workstations also allowed operators to view video streamed from the home during an alarm condition and to use two-way audio to listen to and interact with the home.

At the data center, which was located in a commercial co-location facility, a Web site mediated access by the user to remote viewing and control functions.

From an Internet-connected PC at an arbitrary location such as the homeowner's office, the user could interact with the Web site and with the DVR to carry out remote viewing and remote control such as camera selection. Both live viewing and viewing of video recorded at the DVR, e.g., a video clip showing a package being delivered to the front door, could be performed.

IP-based communications between the VTU in the home and the headend-based DVR took place over the HSD infrastructure, including event reporting, video, audio, control, and heartbeat traffic.

Communications between the DVR and the systems in the central monitoring station used IP over a frame relay permanent virtual circuit.

Communications between remote PCs and the Web site, and between remote PCs and the DVR, occurred over the public Internet.

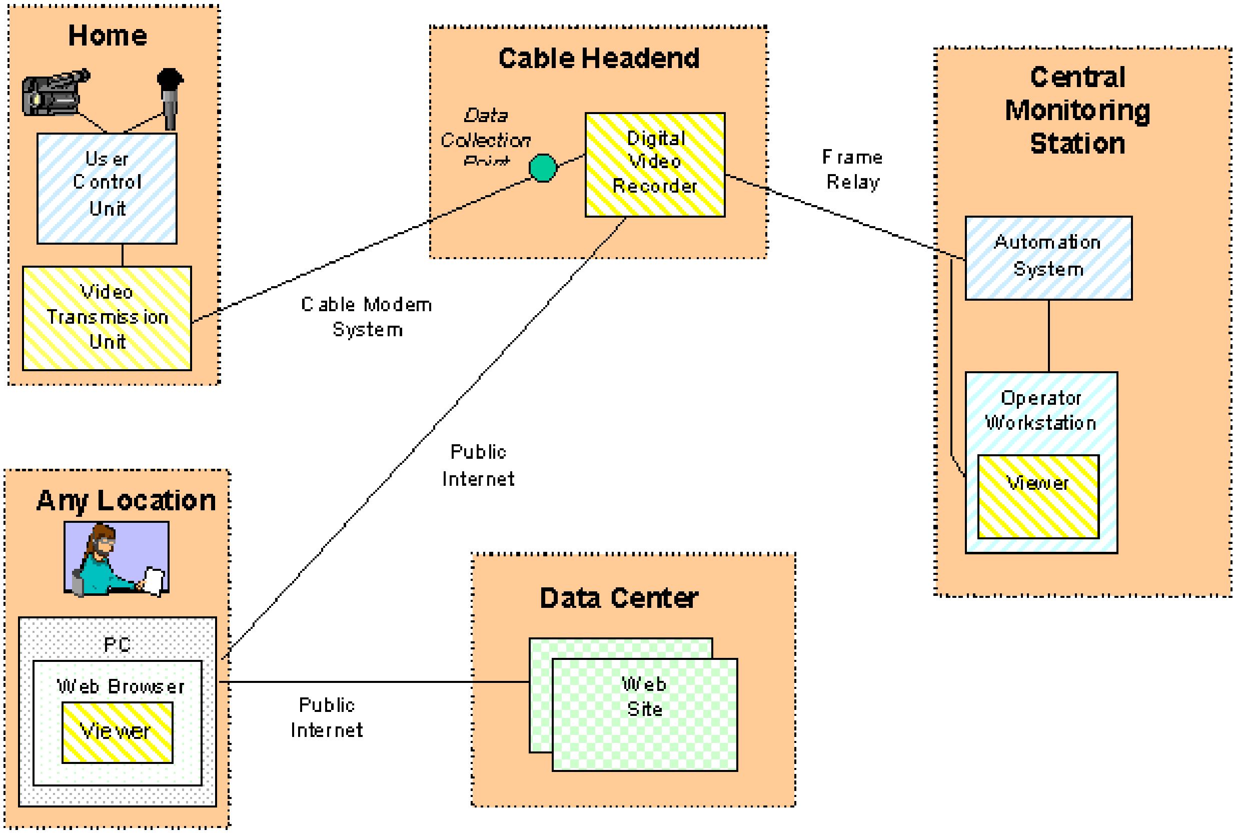

Figure 1. Technology Trial System Architecture

BANDWIDTH STUDY OBJECTIVES

Beyond the need to gain empirical evidence of the capacity impact of the service, a field study was warranted for additional reasons:

- Laboratory measurements cannot quantify individual user behavior such as the frequency of on-demand viewing and the length of a viewing period (although sessions had automatic timeouts, a user could request session renewals to extend the viewing period).
- Aggregate or collective behavior across a user population cannot be gauged in the lab to understand, for example, high levels of on-demand traffic and the relationship between service busy periods and overall HSD busy periods.

With this in mind, key objectives of the trial service bandwidth study were to measure:

- the effect of the service on the upstream HFC path,
- the effect of the service on the downstream HFC path, and
- the effect on the headend Internet connection (both outbound and inbound).

An additional objective was to estimate the potential capacity impact on the HFC infrastructure that would result from high service penetration in an area served by a single fiber node.

DATA COLLECTION METHODOLOGY

In the trial system configuration, all service traffic was routed through the DVR in the cable headend, which provided a convenient data collection point, as noted in Figure 1. This was exploited for these purposes by enabling port spanning in an adjacent Layer 2 switch, allowing all Ethernet frames into and out of the DVR to be sent to a port sniffer on a PC attached to the switch.

Data reduction software allowed time-stamped Ethernet frame headers and IP packet headers to be captured without recording of payload data. Daily uploads and post-processing (filtering, correlation, etc.) were performed off-line using various tools.

MEASUREMENT RESULTS

Typical measurements of upstream traffic from the entire trial population of 85 installations appear in Figure 2 using 30-minute averaging intervals. The aggregate traffic floor, largely consisting of heartbeat messages, typically totaled around 5-6 Kbps (about 70 bps per installation). Traffic peaks correspond to video streaming activity between one or more VTUs in homes and the DVR in the headend. Nearly all of this activity was due to on-demand viewing by homeowners or automatic, non-alarm video recording (triggered by exterior, front-door activity).

Peak on-demand viewing periods tended to occur on weekdays during the afternoon and early evening.

The 30-minute averaging in this chart does not provide a good indication of average per-video session data rates during streaming, but measurements using one-minute granularity showed a consistent average of around 95 Kbps per video or video/audio stream.

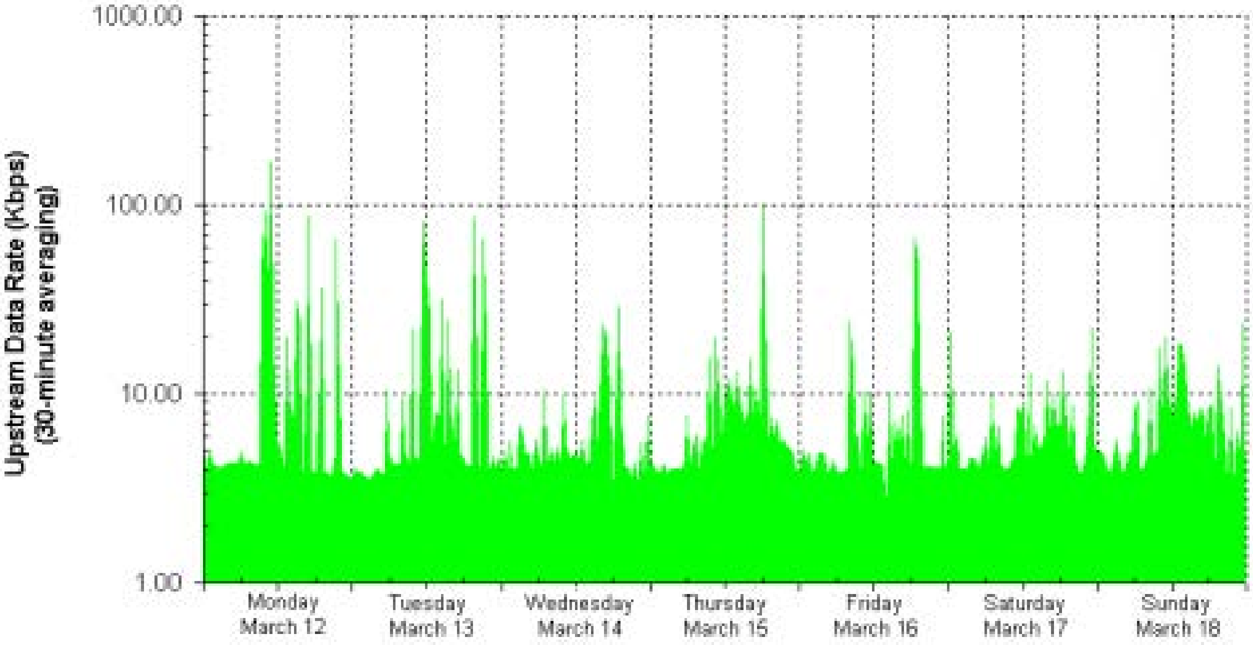

Figure 2. Aggregate Upstream Service Traffic

As shown in Figure 3, the number of concurrent video sessions across the trial population was small, ranging from zero to three in the week shown (and never rising above four actual sessions during the study). This is consistent with the fact that session durations for on-demand viewing and non-alarm event video recording were typically no more than a few minutes, and with the fact that there were no obvious influences that would cause a strong correlation in the timing of viewing by users from different households.

It should be noted that scenarios can be anticipated that could induce correlated behavior, e.g., a tornado warning on a weekday that stimulates many users to attempt to view their homes around the same time. Service-based admission control and congestion management techniques may be appropriate for dealing with such situations.

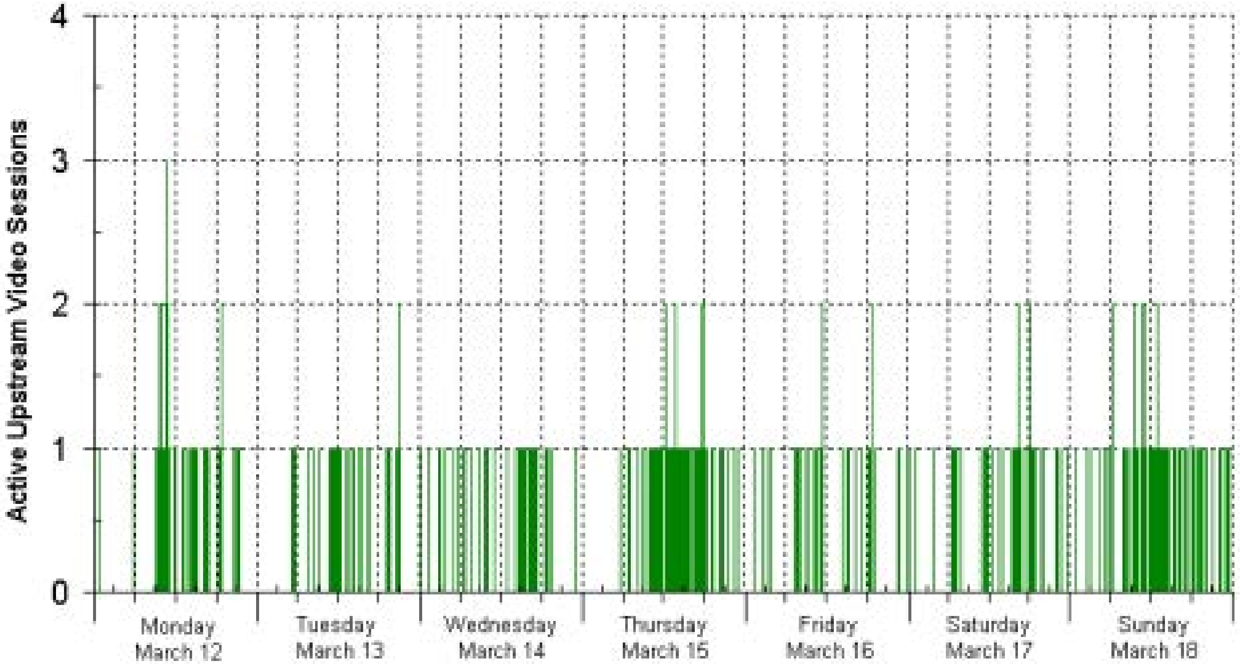

Figure 3. Active Video Sessions

From Figure 4 it may be seen that ten out of 85 installations (11 percent of the trial population) accounted for about 42 percent of the total traffic in the week shown. Over the entire course of the study (roughly ten weeks), the top ten accounted for about 36 percent of the total traffic.

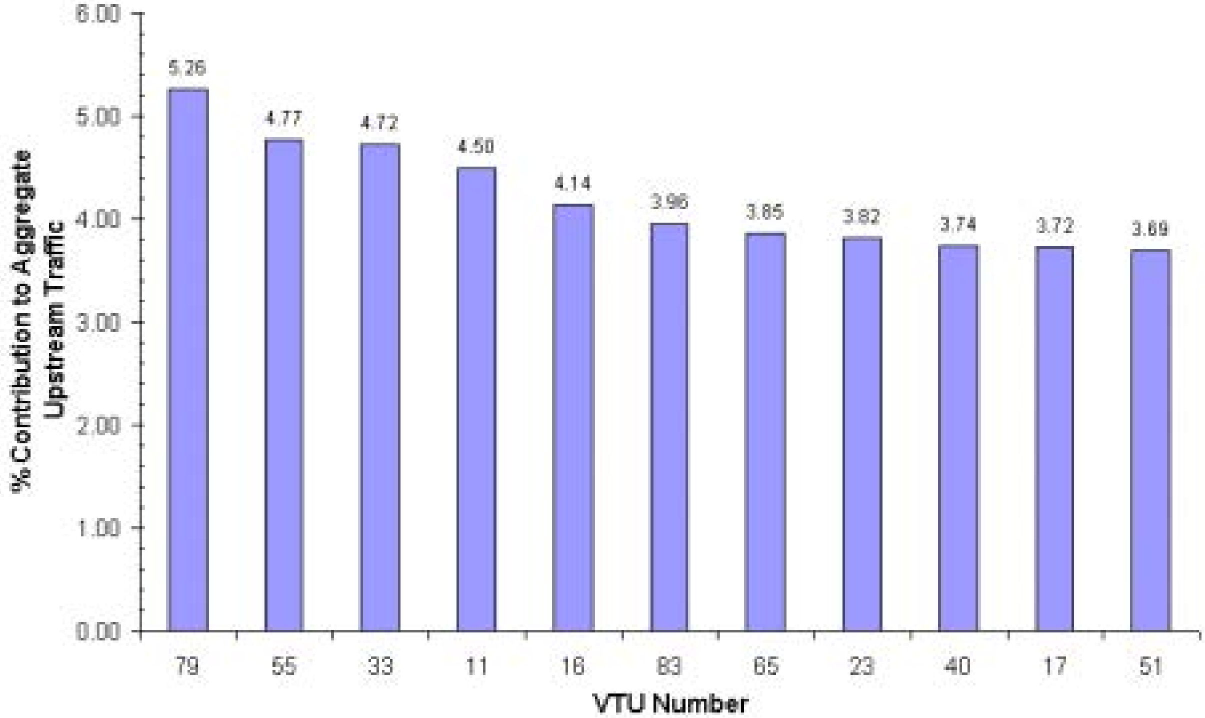

Figure 4. Contributions by Individual Users to Aggregate Upstream Service Traffic

As indicated in Figure 5, most video streaming sessions were only a few minutes in length. It should be noted, however, that the distribution is skewed downward by video recording performed automatically as a result of front-door activity (as opposed to user-initiated, on-demand viewing); these two forms of streaming could not be distinguished with the data collection instrumentation used.

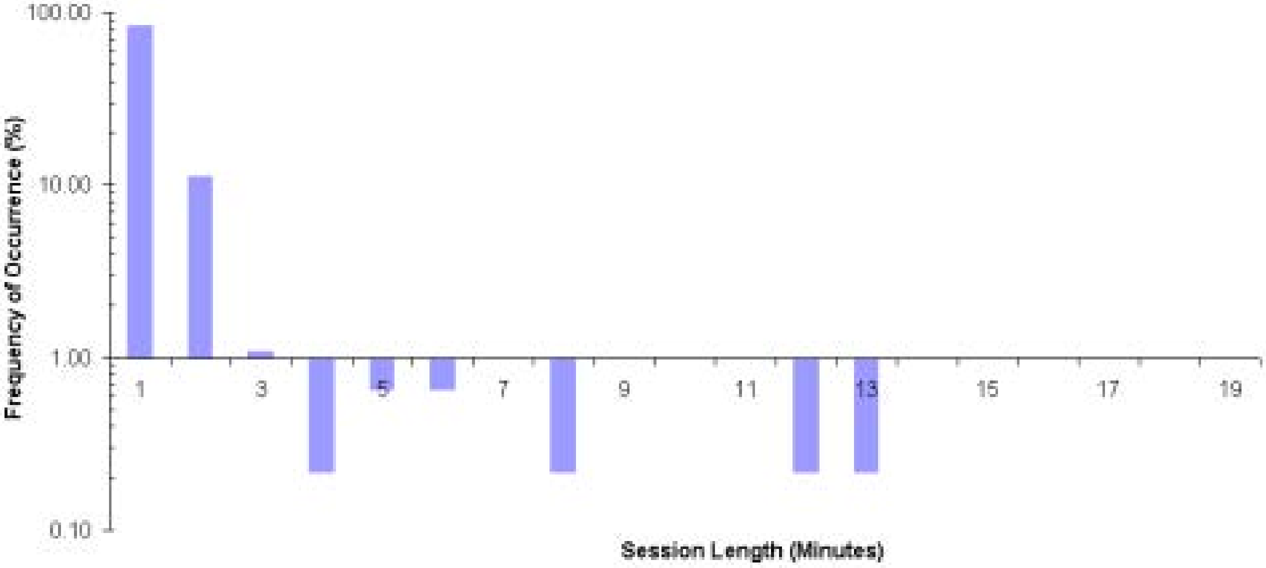

Figure 5. Video Session Durations

Aggregate downstream traffic is shown in Figure 6. The traffic floor averaged 11-12 Kbps across the trial population. The larger peaks reflect system management activity (e.g., software downloads to VTUs installed in homes).

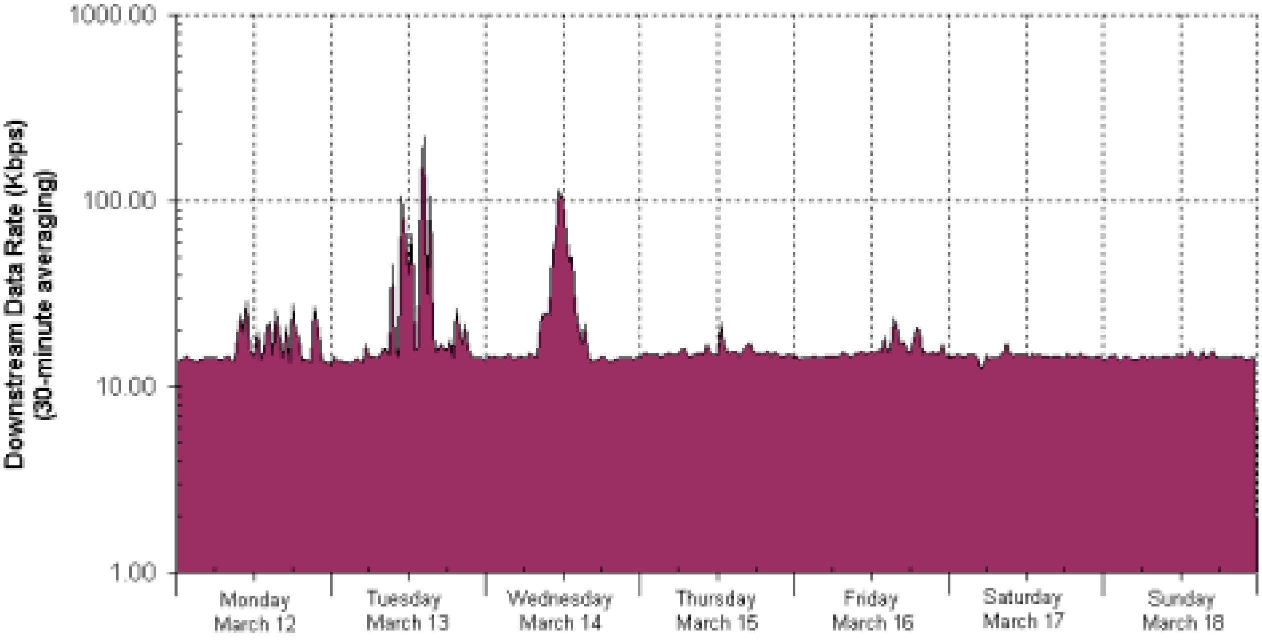

Figure 6. Aggregate Downstream Service Traffic

Aggregate outbound Internet traffic from the DVR in the headend is shown in Figure 7. A comparison with Figure 2 highlights the absence of the 5-6 Kbps floor from heartbeat traffic, which was not relayed by the DVR onto the public Internet connection. Also, peaks in Figures 2 and 7 differ to some extent due to the fact that non-alarm video recorded on the DVR was not relayed over the Internet connection.

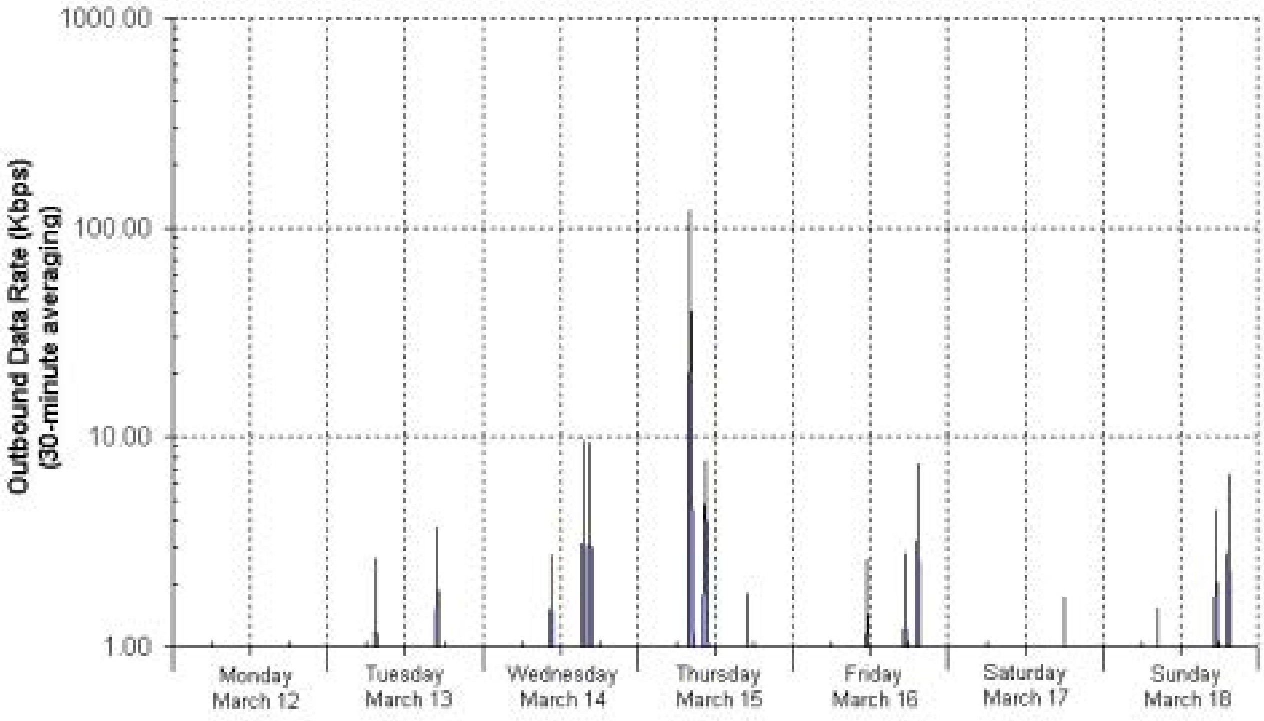

Figure 7. Aggregate Service Traffic Outbound to Internet

Total inbound traffic from the public Internet across the trial population as depicted in Figure 8 was minimal because it comprised only simple control messages such as viewing and camera change requests.

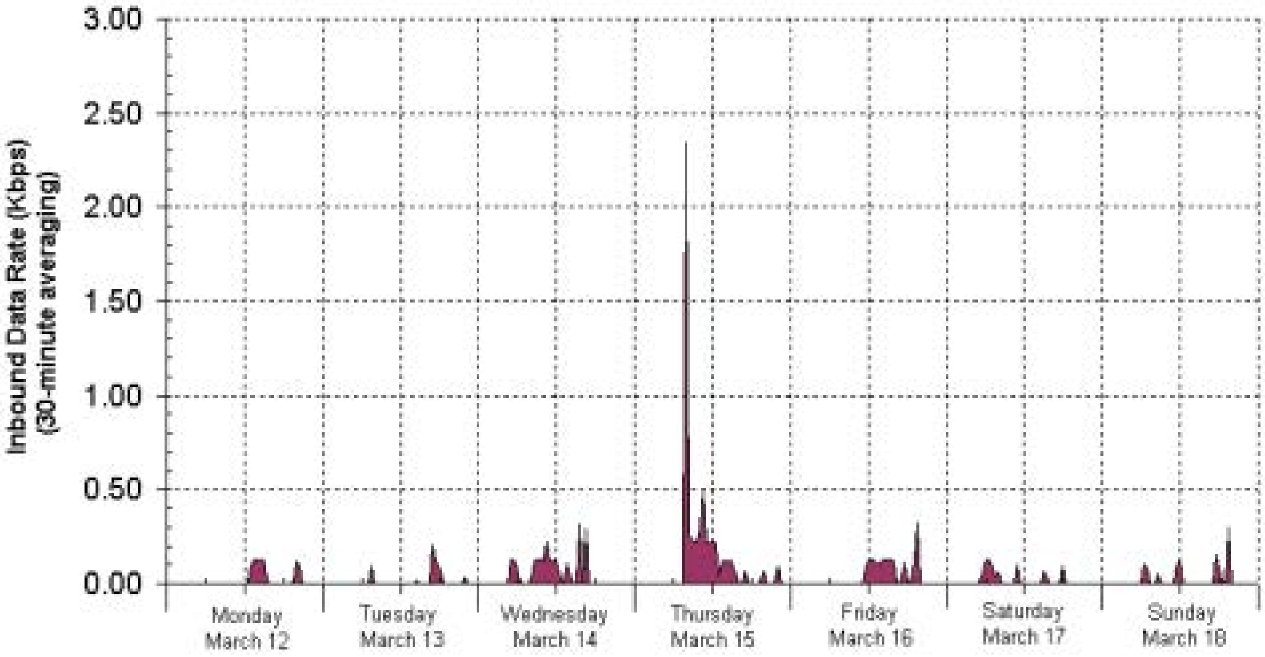

Figure 8. Aggregate Service Traffic Inbound from Internet

CONCLUSIONS

Essential findings – based on the trial service measurements and analysis – are summarized below.

Effect of the service on the upstream HFC path:

- The upstream bandwidth impact of the service was minimal with a traffic floor of 5-6 Kbps (largely heartbeat data) totaled across all 85 trial participants, and peaks rarely exceeding 120 Kbps.
- The top ten traffic generators (11 percent of the trial population) produced roughly 36 percent of the upstream traffic.
- The typical peak number of concurrent video sessions was two to three.
- The average length of a video session was 2 minutes 48 seconds.
- The typical data rate for a single video stream was 95 Kbps.
- Most on-demand viewing occurred during weekdays in the afternoon and early evening. The majority of the traffic seen in other parts of the day and during the weekend was from non-alarm and system testing activity.

Effect of the service on the downstream HFC path:

- The downstream impact of the service was insignificant with a traffic floor of 11-12 Kbps (primarily heartbeat data) totaled across all 85 trial participants.
- Downstream traffic peaks rarely exceeded 120 Kbps.

Effect of the service on the headend Internet connection:

- Both outbound and inbound Internet traffic were minimal. Outbound traffic peaks rarely exceeded 120 Kbps. Inbound traffic peaked at 4 Kbps and consisted of control messages only.

Projected impact of high penetration of the service on the HFC infrastructure under a given node:

To estimate this, assume an average fiber node size of 1,200 homes passed, an available upstream data rate per node of 1.9 Mbps, and a service penetration rate of 15 percent. Based on the technology trial bandwidth utilization measurements, estimates of the loads on the upstream HFC using one-minute averaging are as follows:

- Aggregate traffic floor on the HFC upstream: 11.6 Kbps.
- Peak concurrent live video user sessions: six to seven.
- Peak HFC upstream bandwidth utilization: 603 Kbps or 31 percent of the available upstream bandwidth per node.

Furthermore, use of the trial service measurements overestimates what the commercial service would exhibit assuming similar user behavior. This is because the architecture for the commercial service differs from that of the trial service. In the trial, all traffic was routed through the DVR located in the headend, including recorded non-alarm traffic from outdoor cameras overlooking the front entrance of each home, which used HFC bandwidth. In the commercial service, recording of all normal, non-alarm video is local to the equipment at the customer premises.

Thus, even with this high level of penetration, the service would not impose a significant burden on the upstream bandwidth. It therefore is an example of the type of service that can be very attractive by virtue of its potential as an additional source of recurring revenue with relatively modest impact on the HFC cost structure.

FUTURE WORK

Similar bandwidth utilization studies of the commercial service are planned to quantify the effects of changes to the system architecture such as the one described above. In addition, such studies are needed to understand the effect of possible differences in user behavior due to the addition of on-demand audio, and from the fact that users will be paying customers rather than trial participants.

ACKNOWLEDGEMENTS

The results reported in this paper reflect the efforts of a sizable team. Special recognition, however, is made of the following individuals:

- *Venkatesh Iyer*, for carrying out the majority of the data collection and analysis.
- *Kyle Cooper*, for trial system deployment and data collection instrumentation.
- *Richard Jennings and Harris Bass,* for driving the technology trial definition, implementation, and management.
- *Michael Hale and John Fountain (Cox Communications)*, for guidance and review of the bandwidth study design and execution.
- *Jayson Juraska and Anthony Lee (Cox Communications),* for sponsorship and project management on behalf of the technology trial host company.

CONTACT:

Gregory Feldkamp
Vice President – Engineering and Development
@Security Broadband Corporation
(512) 391-4444
gfeldkamp@atsecurity.net

ARCHITECTING THE HOME GATEWAY TO DELIVER BUNDLED SERVICES AT THE LOWEST COST AND HIGHEST PERFORMANCE

Andrew M. Parolin
SiGe Semiconductor, Inc.

Abstract

This paper provides an overview of the various home-networking technologies, along with their associated advantages and disadvantages. An in-depth analysis of the rollout of 802.11 (a,b,g) is performed

A cable and 802.11 gateway architecture is examined, analyzing the trends in component technology, integration and power consumption, relative to system cost.

INTRODUCTION

Broadband to the home and office has become a reality. With any technology introduction there are two factors for its mass success, and thus economic benefit to both the technology developers/providers and end users: (1) Ease and cost of installation and operation; and (2) worthwhile applications that compel the subscriber/purchaser to use the technology.

With the introduction of home installation kits and the retail market for hardware and software, the technical barriers to setting up broadband services have been greatly reduced. Receiving broadband to the home is becoming as easy as ordering cable television or local telephone service. The "killer application" for this technology is high-speed Internet service, which alone could most likely drive next-generation technology. Applications such as telephony, digital video, and Video On Demand (VOD) are all reaching points of maturity where they will also become prime drivers. Once the enabling technology is installed and subscribers become familiar with one or two basic applications, the technology becomes almost "infectious", with subscribers using more and more of the technology.

In recent years, technology has focused on the "last mile". As broadband to the home has reached the masses, and upgraded cable plants can provide the required bandwidth for new and desired subscriber applications, the early adopters of this technology are now ready for the next generation of broadband technology to the home – the Home Gateway. For the purposes of this paper, a Home Gateway is a unit that can distribute data/services to multiple sources for multi-service applications. Users want the flexibility to take this large "bandwidth pipe" and its applications and seamlessly distribute it to computers, entertainment units and appliances in the home. The primary drivers for this technology are to minimize costs in hardware/systems, ease installation, and allow portability within the home/office. This technology offers a multi-service operator (MSO) the opportunity to increase revenue per subscriber, either by (1) supplying more applications; (2) enhancing the subscriber experience; and (3) offering networking services to the subscriber.

This paper will discuss the trends in the architecture of the Home Gateway, with an emphasis on semiconductors required to provide these services. Specific emphasis will be given to the requirements of semiconductors in the Home Gateway.

Standard	Technology	Approximate Max Data Rate	Advantages	Disadvantages
Ethernet (802.3)	Wired (Cat 5 Cabling)	100 (Mbps)	• Fast • Inexpensive hardware	• Requires new wiring • Not portable
WLAN	Wireless	54 (Mbps)	• Portable • Reducing in cost quickly • Relatively quick	• "Coverage" problems. • Interference • Range
HPNA	Wired (Twisted Pair Telephone lines)	32 (Mbps)	• Some existing wiring in place	• Not portable • Dependent on quality of phone line • Limited number of phone lines in homes
HomePlug	AC/Mains Wires	14 (Mbps)	• Pseudo-portable – plugs are relatively convenient but still require connection	• Data rates are not high enough for multi-media applications. • Dependent on quality of phone line • As most equipment is AC powered, NIC built already built in.

Table 1: Last 100 Feet Technologies

"LAST 100 FEET" TECHNOLOGIES

There are many technologies and standards available for implementing the "Last 100 Feet", as shown in Table 1. These include wireless LAN (WLAN) technologies, Electrical distribution standards such as HomePlug and twisted pair standards, such as Home PNA (HPNA) and Ethernet.

Although all networking technologies have their advantages and disadvantages, it is believed that wireless, and specifically the 802.11 (a,b,g) wireless standards, will provide the major technology thrust for the last 100 feet. Features such as "no cables", ease of portability, available bandwidth and the historical quick adoption of wireless devices such as pagers, cell phones, and PDAs, show that wireless will be the major technology in this market. With respect to bandwidth alone, few of the technologies can support multimedia applications simultaneously delivering multiple sessions of Internet connections at 3Mbps and video/audio applications at 10Mbps.

Table 2 gives an overview of the wireless networking standards. IEEE 802.11 has become the dominant wireless networking technology. In fact, companies promoting other wireless technologies are including support for 802.11 into their product lines to track this market transition. The shaded areas represent the three relevant 802.11 standards.

The rollout of the 802.11 (a,b,g) standards began with 802.11 (b), providing basic data connectivity within the home and office. The 11Mbits/s data rate, approximately the same as 10Mbps Ethernet, was satisfactory to meet subscribers' early data networking requirements. The access rate was similar to the service they received by being attached to a wired Ethernet network solution.

	Infrared	Bluetooth	DECT	HomeRF	HiperLan2	802.11 (a)	802.11 (b)	802.11 (g)
Speed(Mbits/s)	4	1	0.032	1-10	54	54	11	22/54
Launch	-	1997	1988	1998	1999	1999	1999	July/2002
Range (ft)	15	30	1,000	150	400	400	300	300
Modulation	N/A	FHSS	FHSS	FHSS	OFDM	OFDM	DSSS	DSSS/OFDM
Frequency(GHz)	IR Band	2.4	2.4	2.4	5	5	2.4	2.4

Table 2: Wireless Networking Standards

802.11 (a) was intended to rollout next, providing a much higher data rate for advanced services (video, VOD, telephone) and larger networks. However, this rollout strategy has changed over the last year. 802.11 (g), expected to be ratified in the summer of 2002, will rollout simultaneously, as shown in Figure 1, bridging the 802.11 (a) and (b) standards. According to Cahner's In-Stat, 802.11 (g) (or a combination of 802.11 (a) and 802.11 (b)) will capture 63% of the market by 2005, compared to 802.11 (a) at 33%.

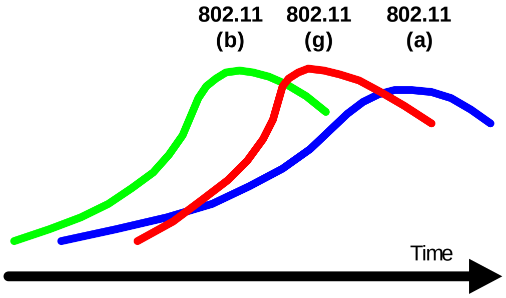

Figure 1: 802.11 Technology Implementation Roadmap

802.11 (g) provides many advantages over 802.11 (a) while transitioning from 802.11 (b). Because the operating frequency is at 2.4GHz, in the Instrument, Scientific and Medical (ISM) Band, 802.11 (g) leverages the technology that has been developed for 802.11 (b), Bluetooth and the cellular market. This includes backward compatibility of the standard 802.11 (g) to 802.11 (b) standards, as well as components on the reference design. This minimizes technical risk and component costs due to large volumes provided by these other markets, as shown in Figure 2.

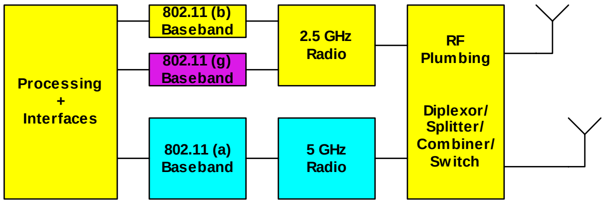

Figure 2: 802.11 (a,b,g) System Diagram

Another advantage of operating at 2.4GHz is to minimize propagation issues that have been a major cause of concern for using the 5GHz 802.11 (a) standard. 5GHz operation has had problems reaching the entire network footprint. Operation at 2.4GHz has had much empirical analysis, trials and customer feedback. If an MSO is going to promote the use of a wireless network technology, its operation will reflect on their total service offering, so issues such as this must be bulletproof before being released. The 2.4GHz ISM band, used by 802.11 (b) and (g) is also recognized around the world, so equipment does not have to be customized according to geography. Portions of the 5GHz spectrum used by 802.11 (a) has been allocated to other services, and although work is in process to make it usable worldwide by a 802.11 (a) variant called 802.11 (h), this will take time. At the access point level, total system costs can be reduced as components can be shared between the 802.11 (b) and (g) systems. Microsoft has previously stated that any new wireless LAN standard including the 5GHz 802.11 (a) and (h) should be backwards compatible with 802.11 (b) in order to have the software drivers incorporated within the Windows XP software package. So, access points and nodes can be sold with 802.11 (g) compatibility at a small price increase. This makes the transition to the higher data rate services much easier as subscribers do not have to replace their equipment.

Concerns with 802.11 (b) and (g) largely center around potential "interference" issues because the 2.4GHz ISM band is used by other wireless network technologies such as Bluetooth, HomeRF and DECT. Microwave operation has also been known to cause interference issues but these issues have been largely resolved.

In the end, the 2.4GHz technologies, 802.11 (b) and (g), will be pushed as far as possible until capacity begins to drop due to over usage. At this point, wireless access devices will have had the opportunity to begin offering 802.11 (a) compliance in their 802.11 (g) product lines, minimizing the switch over costs when the transition begins to occur.

LOCATION OF THE HOME GATEWAY

Figure 3 represents the opportunities for the Home Gateway to provide access to the last 100 feet. Each system has advantages and benefits to the subscriber and MSO, so there will be no one winner in the Home Gateway market. Technology selection will depend on the maturity of the subscriber base, and the goals and core competence of the individual MSOs.

(1) *Pole Side Network Access Point.* Although a novel Home Gateway solution, this solution has the advantage that it allows the MSO to maintain control of configuration, provisioning, operation and upgrades of the network access point. This relates directly to the solution cost with respect to maintenance and support. Other advantages related to revenue and costs include: (i) minimizes the technical competence required by a subscriber, increasing the total overall subscriber base, (ii) requires subscribers to contact the MSO to add features and additional nodes, so the MSO can more easily charge for these additions and (iii) hardware costs can be minimized by sharing access points among multiple subscribers depending on their consolidated bandwidth requirements. Security issues can be solved via encryption.

(2) *Outdoor Wall Mounted Access Point.* This solution is similar to (1) above with the added benefit that it is commonly used for cable telephony solutions today,

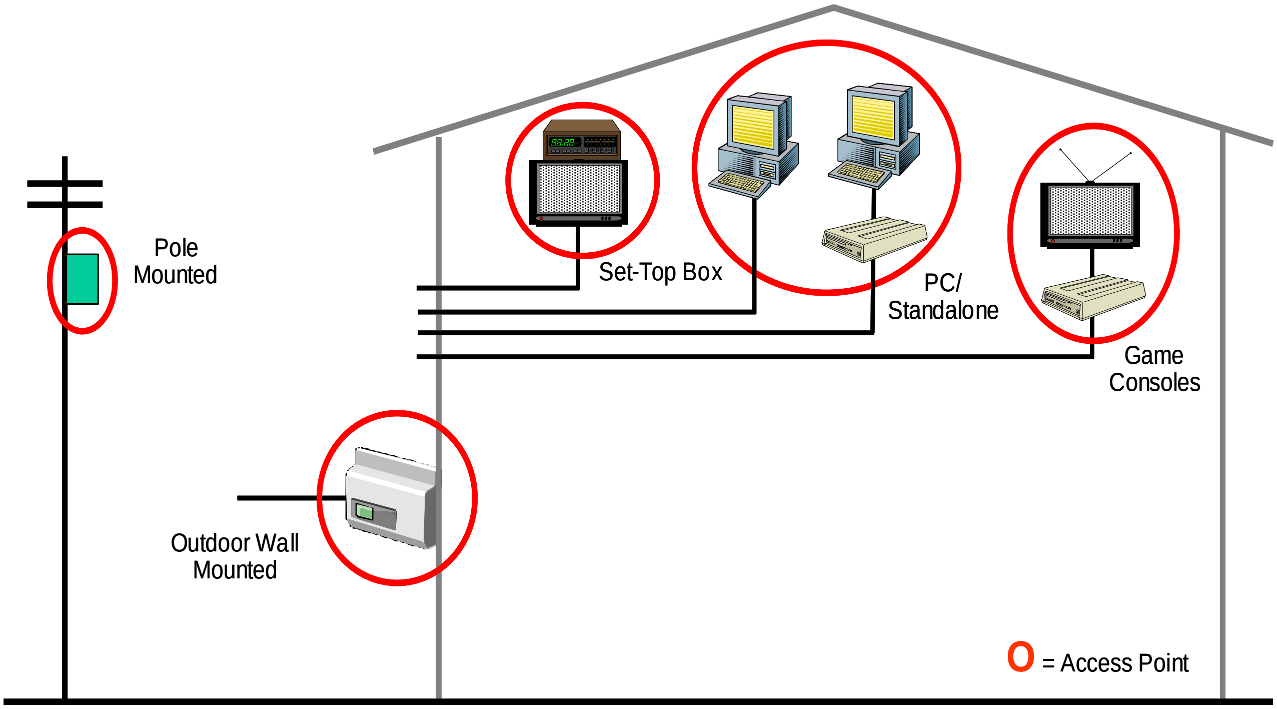

Figure 3: Home Gateway Opportunities

minimizing technical risk in the rollout. The advantages and disadvantages are similar to (1) above with an additional advantage being that the MSO has the ability to select from multiple access solutions such as Wireless, HPNA and HomePlug.

(3) *PC/Standalone Access Points.* These products can be commonly found in the retail market today and require a computer with connected access point. Benefits to the MSO of these access points are that the user must purchase the hardware and is responsible for maintaining their own network. Disadvantages with this model is that (a) it is difficult for the MSO to collect additional revenue for this service and additional nodes and (b) the customer will not have the skill to properly setup their network.

(4) *Set Top Box (STB) Access Points.* These products can already be found on the market, with STBs containing both traditional video and cordless telephone service. The advantages of this access point are that (a) subscribers are already familiar with the STB product in the home, minimizing the “new technology” scare that may limit penetration; and (b) subscribers purchase/rent their individual equipment, minimizing MSO costs. With entertainment typically being the primary driver for new applications within the home, the STB is well positioned to be a strong winner.

(5) *Game Console Access Point.* Although another novel Home Gateway product, who best to drive the next generation broadband technology than the “next generation” of subscribers. New networking technology can be driven into the home via gaming features.

SYSTEM ARCHITECTURE

Figure 4 shows a typical architecture, and major components, of a Home Gateway using broadband cable connection as the front-end and 802.11 for the back-end, in-home distribution. The cable front-end architecture requires multiple tuners to receive all the available applications. Semiconductor manufactures have been under pressure to lower the total solution cost of the system. This translates not only to lower IC prices, but also to the ability to offer application-specific semiconductors that contain higher integration, thereby absorbing more of the discrete solution, and minimizing board/product size. In the end, the solution cost must be minimized. Increasing integration is occurring with both digital and radio frequency (RF) technologies. With digital technology, the functionality of the MAC, PHY, processor and back-end chipsets are constantly being integrated on one IC. With RF technology, IC tuner and amplifiers are increasing integration. This minimizes the RF solution to a two IC solution with minimal discrete component support.

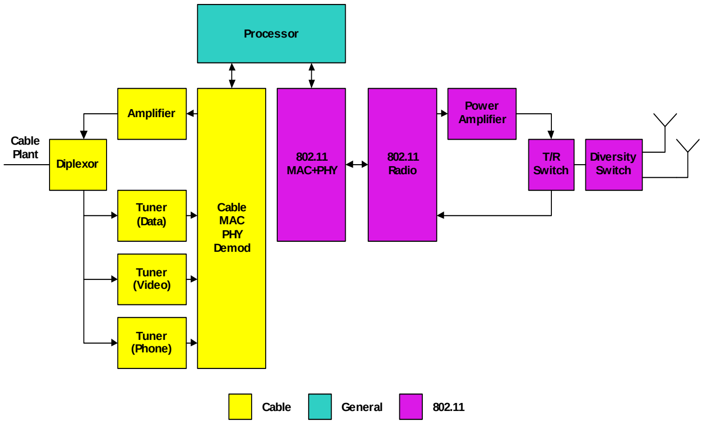

Figure 4: Cable-802.11 Home Gateway Architecture

SYSTEM DESIGN CONSIDERATIONS

Outdoor verse Indoor

Outdoor products have cost and revenue advantages over indoor, subscriber purchased units. However, there are two primary concerns with producing equipment for outdoor operation: (1) operation over a wide temperature range; and (2) powering the equipment. Product offerings marketed to outdoor applications typically have to perform over a temperature range of –40 to +85 Celsius. Developing a product to work at one particular temperature extreme is not difficult. However, designing one product to work at both extremes requires creative and robust designs. Historically, RF designers needed to choose components specifically designed for this wide temperature range or choose components with very low variation. This shortens the list of available components. These components were typically more expensive. In addition, due to the lack of options, they could cause supply problems during times of high demand, requiring excess inventory costs to be incurred. Designers are also required to add increased performance margin for these harsher requirements.

The second major concern, powering equipment, comes from the fact that MSOs must economically power remote equipment. Local AC/main powering raises safety issues that must be addressed, and increases the cost of the equipment as well as liability concerns. If the gateway is supporting primary telephony, the MSO is required to ensure that the gateway remains operational during a local power failure. Battery powering is not a preferred solution due to equipment costs and requirements to eventually replace batteries. Remote powering, via the cable, has been selected as the preferred solution, with battery operation in lower density subscriber situations. However, this makes lowering the power consumption of the equipment critical to minimizing the remote powering infrastructure costs. Lowering the power consumption of a Home Gateway allows the MSO to deploy more gateways for a given infrastructure cost.

For the cable front-end, these two issues have largely been solved due to the aggressive push by some MSOs to deploy cable telephony, where the preferred deployment is outdoor, remote powered customer premise equipment. Gateway manufacturers and component manufacturers have years of experience developing products meeting these requirements. The outdoor component design issues have been minimized by the development of integrated circuit process technologies, such as silicon germanium (SiGe), and increased integration by IC manufacturers. The downstream portion that historically was composed of 300 components has now been integrated onto one IC tuner, with some of these new IC tuners operating over –40 to 85 Celsius. An added benefit of these process technologies is that their power consumption is quite low, facilitating the use of efficient remote powering over the cable. With the tuner now integrated onto one IC, sourcing issues for wide temperature range components is minimized, and design issues relative to matching components is easier. Baseband vendors have also developed products for the cable telephony market, offering lower power consumption.

802.11 (a,b,g) products have been following a similar trend as cable products, with semiconductor manufacturers supporting extreme temperature ranges and low power.

Power Consumption

Both the cable and wireless markets have individually been driving for low-power solutions. Cable applications must have low-power operation for lifeline cable telephony, remote indoor power, USB powered cable modems, smaller form factors and increased applications without the requirement for an internal fan. 802.11 technologies have been driving for lower power because of PC peripheral power supply constraints, and longer battery life for portable applications. Both industries have lowered power consumption by leveraging new silicon processes, largely driven by the cell phone market, increased integration and better design.

Apart from the low power consumption required for outdoor, remote powering applications, as mentioned above, there is also the requirement for low power consumption for indoor applications. In applications where multiple access points may be connected to a central unit, Remote powering, over the indoor data connection, can be used to minimize deployment time and costs for access points in indoor applications. As no AC re-wiring is required, installation can be done by the corporate IT group, technician or subscriber. An electrician is not required saving time and cost. Secondly, portable in-home/business solutions, such as PDAs, laptop and portable tablets, require battery operation and usage time must be maximized to ensure subscriber satisfaction. This puts a constraint on the wireless standard and its available chip-sets to offer extremely low power consumption.

The current generation of cable modem applications requires approximately 4-5 Watts. The 802.11 (b) wireless backend currently requires approximately 2 Watts to operate with the power budget currently dominated by the power amplifier, which consumes about 1 Watt.

Now, the market is beginning to see components that will allow cable modem functionality to approach a power consumption of 2.5-3 Watts. This drop in power consumption is largely based on the drop in power consumption of the major cable components, with some components dropping their power consumption by 50% over the last year. Current power budgets for the major components, in sampling, of the cable front end are: Downstream Tuner = 600 mW, Upstream Amplifier = 500 mW, Processor/Mac/Phy = 500-600 mW.

New 802.11 (b) solutions have been able to achieve operation at under 1 Watt, largely driven by more power efficient designs and new power amplifier technology, which reduces the power amplifier current consumption, the most power intensive portion of the design, to 45% over current generation solutions.

So, combining the 3-Watt cable and 1-Watt Wireless sub-systems a basic access point could draw under 4 Watts. This can be powered, using today’s technology, by a remote source over the data cable.

Bills of Materials

Both cable and wireless industries have been moving towards higher integration and thus a reduction in components. IC manufacturers have been racing to increasingly integrate more of the total solution into one IC. This is most evident in the cable modem industry, wherein (1) silicon tuner manufacturers integrate the complete receiver on one IC; and (2) baseband manufacturers integrate support for the complete range of back-end distribution standards including, PCI, Ethernet, 802.11, and HPNA. Some baseband vendors have even attempted to integrate the cable upstream amplifier into their IC. For 802.11 (a,b,g), the radio, which was typically two separate ICs, has been integrated onto one radio IC, making it a three chip set solution. New power amplifiers have also been designed, which eliminate the output filter by integrating it within the matching circuitry.

This can be most dramatically seen in the cost of basic cable modems, which have dropped in price from $179 to $89 dollars within a year. 802.11 (a,b,g) devices are expected to follow similar aggressive pricing as they enter the consumer market.

CONCLUSION

An overview of Home Gateways technologies was given, with the most favorable architecture being a broadband cable front-end, with multiple tuners for multi-media applications, and an 802.11 (a,b,g) back-end, for the in-home distribution. Power consumption, standards, integration and operational environments were all examined with respect to cost and performance.

CONTACT INFORMATION

Andrew Parolin
Marketing Manager, Broadband
SiGe Semiconductor, Inc
2680 Queensview Drive
Ottawa, Ontario, Canada K2B 8J9
Phone: (613) 820-9244 x524
Fax: (613) 721-7726
E-Mail: amp@sige.com

BANDWIDTH UTILIZATION ON CABLE-BASED HIGH SPEED DATA NETWORKS

Terry D. Shaw, Ph. D.[1]
CableLabs

Abstract

The CableLabs Bandwidth Modeling and Management project addresses the use, management, performance simulation, and network economics of cable high-speed data systems. On this project, we have analyzed consumer usage patterns based upon data collected on live cable-based high-speed data systems as well as network simulations.

Usage data on live cable networks indicate that traffic flows are fairly predictable over DOCSIS™ networks. Some of the primary patterns emerging include:

- ***Skewed distribution of bandwidth consumption**. As a general rule in many systems, 30% of the subscribers consume about 60% of the data.*
- ***Students drive seasonal characteristics**. System usage is appreciably higher during local school holidays and vacation periods.*
- ***Usage rapidly evolving.** In the data we collected, the average per capita use of data shows steady growth over time in both the upstream and the downstream, with upstream use growing more rapidly. This results in the average downstream/upstream symmetry ratio trending downwards, with a significant number of subscribers using more upstream than downstream data.*

These observations indicate the importance of deploying DOCSIS 1.1 in order to meet this increasing demand for upstream capacity. Preliminary simulation results indicate DOCSIS 1.1 will enable a significant increase in upstream system carrying capacity over the already substantial capacity of DOCSIS 1.0. These simulations indicate that DOCSIS 1.1 will allow almost 20% more upstream capacity than DOCSIS 1.0.

Simulations have also been used to study the characteristics of specific types of traffic found on cable networks. For example, simulations of peer-to-peer applications indicate that one user without rate limits can consume up to 25% of upstream capacity with a usage pattern resembling a very high-speed constant bit-rate application. These simulation results highlight the potential benefits for cable operators to manage their bandwidth using tools such as rate limits, service tiers, and byte caps on usage.

INTRODUCTION

The usage of DOCSIS high-speed data networks by a mass audience is expected to grow dramatically over the next several years. In order to gain a better understanding of the nature of growth to be expected, CableLabs has initiated the Bandwidth Modeling and Management project. This project has several aspects including the collection and analysis of bandwidth use on live cable high-speed data systems, the development of tools for the modeling and management of data traffic, and assessment of the strategic technical and economic implications of these measurements.

To date, this project has focused on the aggregate and individual characterization of the use of data by subscribers on DOCSIS and other cable-based high-speed data networks. We have studied patterns of behavior based on data collected from live cable systems and the performance of DOCSIS 1.0 and 1.1 systems in simulation. The Bandwidth Management project enjoys operational support from a number of cable operators in North America. The analyses in this paper are comprised of a composite view of these data collection efforts.

The primary focus of our data collection and analysis work during 2001 has been to characterize the data usage patterns over cable systems. As will be shown, usage data on live cable networks indicate that traffic flows are fairly predictable over DOCSIS and other cable-based high-speed data networks. However, the high-speed data market is rapidly evolving and continuing study is necessary in order to understand the network implications as new network applications and services are developed.

DATA COLLECTION AND ANALYSIS

Data has been collected from a variety of types of sources. The three primary sources used to date have been the web page MRTG[2] graphics used for bandwidth management by cable operators, data logs used to archive information by MRTG, and specially instrumented data collection systems.

A number of parameters can be measured and collected from high-speed data systems. However, the specific parameters collected vary from system to system depending on the available tools. Also, the means of presentation of the data are varied. Data collected on different systems must be normalized (converted to the same basis of measurement) so that results can be compared on an "apples-to-apples basis." Three parameters that are of particular use for the analysis of system performance and the analysis of user behavior are defined as:

- **Volume of data used.** The amount of data consumed (uploaded, downloaded, and uploaded + downloaded) during the measurement interval. This parameter can be collected for an aggregate (for example, all of the users on a given upstream or downstream frequency).
- **Data Rate.** Practically, the average data rate is merely an expression of volume of data used within a specified short time period, usually a second. It varies from volume of data rate primarily in the time units of the measurement interval. It is usually collected for an aggregate (for example, all of the users on a given upstream or downstream frequency).
- **Symmetry.** This parameter is the ratio of (Downstream)/(Upstream) where Downstream and Upstream refer to either volume of data used or data rate. It is used to characterize user behavior and has many system architectural implications.

General Observations

The access to live network traffic data provided to CableLabs by its members provides insight to a number of facets of network performance and operational methodologies. These include:

- **Daily (diurnal) usage patterns**
- **System capacity loading**

- **Seasonality of use**
- **Overall trends in volume and symmetry of network use**

All of the findings in this report should be regarded as preliminary because of the limited amount of data analyzed, and the way individuals are using cable-based high-speed data systems is constantly evolving. Continuing study is necessary in order to understand the evolution of system use. The mathematical characteristics of several of these parameters are discussed in the appendix.

Diurnal Usage Patterns

Our traffic measurements indicate that regular or predictable diurnal, or daily, variations of the volume of traffic flow over a network are present. These variations are basically sinusoidal in nature with a period of 24 hours, with possible variations of behavior between weekdays, weekends, and holidays.

While the specific pattern diurnal variation observed will vary from system to system, the pattern shown in Figure 1 is fairly typical. This graph, taken from a MRTG site used by a system operator for on-going system monitoring, represents the aggregated traffic of cable modem subscribers during the period 11/26/01 through 12/4/01 in a major metropolitan market. The green area represents the downstream data use and the magenta line represents the upstream data use. Each point graphed represents the average aggregate data rate for a one-hour interval. In this market, traffic begins to build in mid-afternoon on school days and typically reaches its peak around midnight local time. It then rapidly falls to a minimum at about 5 a.m. local time. In this case, the daily usage peak "busy hour" is from 8 p.m. to midnight on Friday evening. Weekend use is similar to weekday use with more traffic seen earlier in the day (the spike occurring on Thursday morning is caused by a rollover in the counter used to collect this data).

Figure 1. Daily system usage pattern for a metropolitan system.

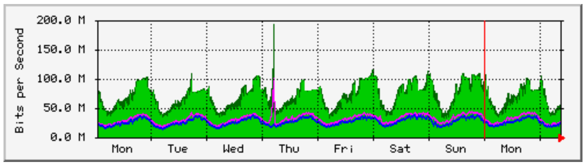

Indications of Seasonality

Seasonal variations are the macroscopic variations in the way a system is used throughout the course of a year. During the early stages of system deployment, the seasonal variations are usually obscured by the overall increase in traffic due to the addition of new subscribers on existing network segments. This behavior can also be masked as users are re-allocated to different network systems as the subscriber traffic increases. Hence, seasonality can only be observed in systems where the same user population has been fairly stable or slowly increasing over a longer (more than six months) period of time. In systems where this has been the case, however, the seasonal effects can be clearly visible. However, Internet usage is continuously evolving, and these variations must be tracked over a long period of time so that the effects can be clearly understood.

The seasonal variations are clearly visible in Figure 2. In general, high traffic times appear to correspond well with school vacations and major holidays[3]. This traffic data shows a marked increase about Thanksgiving (end of November), slows during mid-December, and upsurges again near the end of December (Christmas). After the December school break, traffic decreases with a brief upsurge in early April (Spring Break) and early May (Finals Week). When school lets out for the summer (early June), there is a tremendous sustained increase in the traffic until early August.

After schools reconvened for the Fall term, traffic reverted to a level similar to that seen in the Spring. There is one minor upsurge on Columbus Day 2001 (early October) that corresponds to a similar upsurge on Columbus Day 2000 (not shown on this graph).

Figure 2 is a MRTG graph of the upstream traffic on a well utilized upstream for the time late-October 2000 through November 2001. This upstream serves an area in which cable modem service has been available for several years, and has had fairly stable (but slowly growing) user population over the course of the period May through November 2001. Each data point represents the 24-hour average traffic on the upstream for this aggregate of users.

Figure 2. Seasonal high traffic appears to correspond to school vacations for late-October 2000 through November 2001.

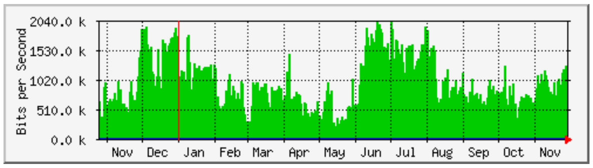

Overall Trends in Volume and Symmetry

The overall trends of the use of data (volume of data consumed per user, symmetry of data flow on the network) will vary depending on a number of factors, including:

- The demographics of a particular region,
- User sophistication,
- “Neighborhood effects” such as the popularity of an interactive game played among residents of a neighborhood,
- Geographic and weather factors, and
- System management policies (such as tiers of service based on data consumed).

These effects will vary from system to system and from node[4] to node within a system. Overall, across a large aggregation of users, the volume of data consumed per user appears to be increasing, and the downstream/upstream ratio of data consumed appears to be decreasing. Figure 3 shows the average symmetry, the ratio of downstream to upstream use, calculated for the months of May through September 2001 for a very large aggregate of subscribers. During this period of time, the symmetry ratio decreased from just over 2 (2 times as much downstream data as upstream) to 1.4. This value was calculated based on the data flow on the network. Other measurements taken on an individual basis on a node indicates that individuals vary widely in their symmetry of use. The decrease in the overall average symmetry of traffic has been observed on a number of different systems. Early indications are that, on an individual node basis, the symmetry is relatively high (about 3) when a system is first deployed and decreases as the users on the system begin to more fully understand the capabilities of the cable high-speed data communications system.

Figure 3. Overall symmetry of network use based on volume of traffic flow for a large aggregate of users.

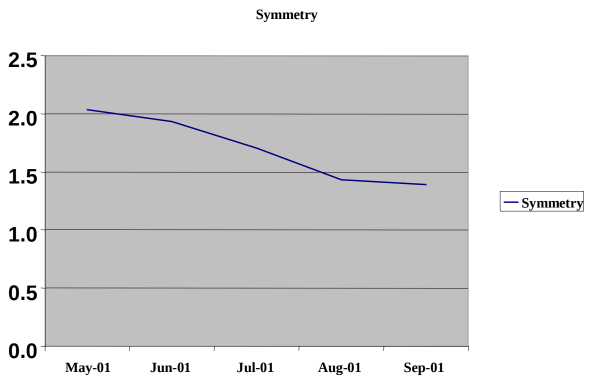

SIMULATION of DOCSIS NETWORKS

An important reason for performing data collection and analysis is to have a basis of empirical data to use for the construction of predictive tools. Analysis of measurement data, such as the probability distributions discussed in the Appendix, are essential for making simulations realistic and accurate. Predictive tools can simulate network performance given the behavior of specific network elements based on these data. These tools can also be used to study the behavior of new applications and services on the network. One such predictive tool is the Modeler software developed by OPNET.

CableLabs developed a model of the behavior of the Gnutella peer-to-peer file transport protocol on a DOCSIS 1.0 access network using the OPNET Modeler software program in order to determine the types of traffic characteristics that can be expected from these applications.

Figure 9. Network topology used to examine the effects of peer-to-peer file sharing on a DOCSIS access network.

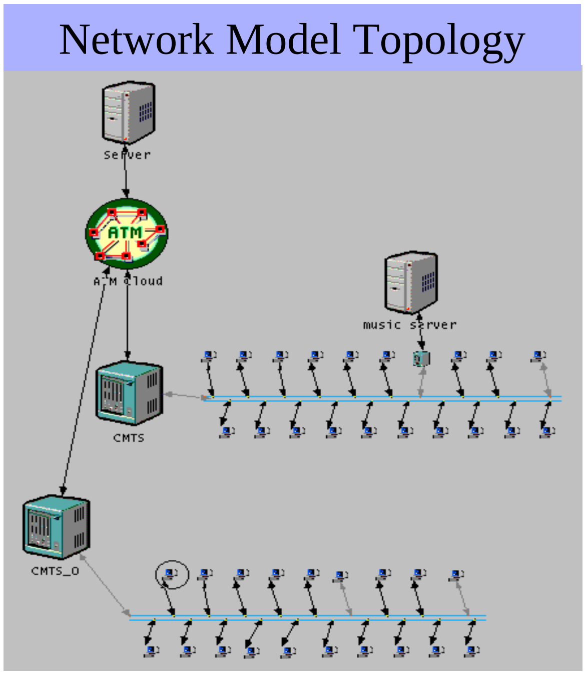

In this simulation, two DOCSIS access networks were configured. Figure 9 provides a diagram of the network topology studied in this simulation. One of these access networks (connected to ***CMTS*** in Figure 9) hosted one cable modem that supported a music server. In the simulation, this modem-server (***music server*** in Figure 9) combination would replicate the performance of a single subscriber that was acting as a source of files (such as MP3 audio files) to other users of the peer-to-peer application. The MP3 file transfer was simulated as a 5 Mbyte HTML file download from this server to users of the peer-to-peer application. The users of the peer-to-peer application were simulated by 12 of the cable modems connected to ***CMTS_0*** in Figure 9, each of which would download five of the simulated MP3 files per hour. These modems served as data sinks (or destinations) for the files transferred from ***music server***. The primary role that these modems played in the simulation was to send requests for file transfers to ***music server***. The module ***server*** acted as the source of typical background traffic for this simulation.[5]

In this simulation, the DOCSIS access networks were configured without any rate limits in order to study how much traffic could be generated by the single subscriber hosting the ***music server***. The DOCSIS upstream data rate simulated was 5.12 Mbps. The simulation was configured to simulate the activity generated in 20 minutes of network activity.

The effect on the upstream traffic generated by the ***music server*** in this limited simulation is shown in Figure 10.

Figure 10. Upstream traffic load generated by a single user serving MP3 files on a DOCSIS 1.0 access network. In this simulation, the CMTS is configured as a router.

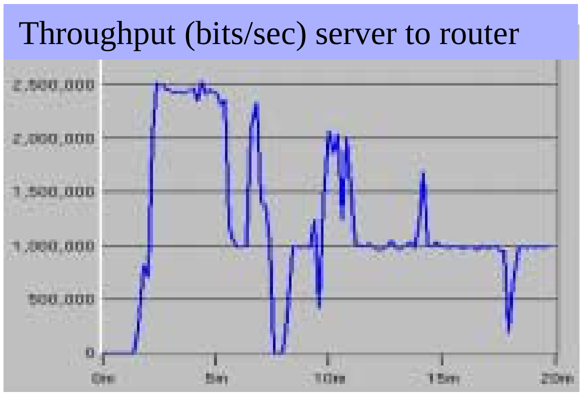

As can be seen in this graph, this single user consumed nearly 50% of the upstream bandwidth available for a substantial proportion of the simulated period, and nearly 25% for the rest of the period. Therefore, a single user of peer-to-peer applications will have a persistently high data rate resembling a very high-speed constant bit rate application. The results of this simulation correspond closely with network behaviors that have been observed in the real-world environment. This simulation highlights the importance of bandwidth management tools such as rate limits, volume limits, and traffic prioritization and differentiated services (from DOCSIS 1.1 and PacketCable™) for assuring a well-functioning network and good customer experience for all subscribers.

PACKETCABLE SIMULATION

CableLabs developed PacketCable protocol extensions to the commercially released DOCSIS 1.1 OPNET model library. A number of scenarios were developed in conjunction with the PacketCable project to test the capabilities of the model and explore the efficiency of the DOCSIS 1.1/PacketCable protocols. The basic scenario topology used for testing is shown in Figure 11. In the scenarios tested, the upstream channel data rate was set to 1.28 Mbps.[6]

Figure 11. PacketCable™ Scenario Topology

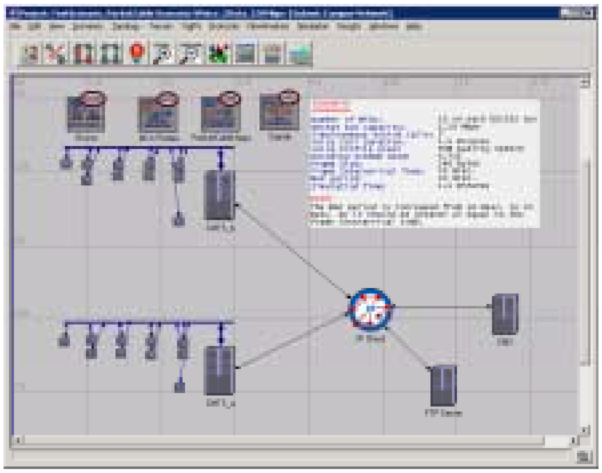

In these scenarios, two CMTS supporting 10 multimedia terminal adapters[7] (MTA) each were configured that were able to communicate through an IP cloud. A combination of voice (using the G.711 codec) and data traffic was modeled in order to study specific load conditions. In the simulation, calls originated on an MTA on a bus supported by one CMTS were terminated at an MTA supported by the other CMTS, and all of the voice calls were initiated and terminated simultaneously in order to produce “worst case” results. Due to the low upstream data rate (1.28 Mbps) and the high data rate required by the G.711 codec (nominal line rate is 64 kbps, but encoding inefficiencies create an actual line rate closer to 128 kbps), the upstream channel was saturated with relatively few simultaneous calls. In fact, in simulation, the upstream channel could only support 9 simultaneous calls; a tenth call specified in the simulation was not completed due to lack of available bandwidth at a quality level needed to support the call. Figure 12 provides a graph of the upstream channel utilizations expressed as a percentage of the total possible load throughput rate for the scenarios with 8, 9, and 10 specified voice calls. (Note that the number of calls was limited due to the low data rate of the underlying DOCSIS system. Much higher data rates, with resulting call carrying capacity would be used in actual system deployments).

Figure 12. Upstream channel utilization for PacketCable™ voice usage simulation.

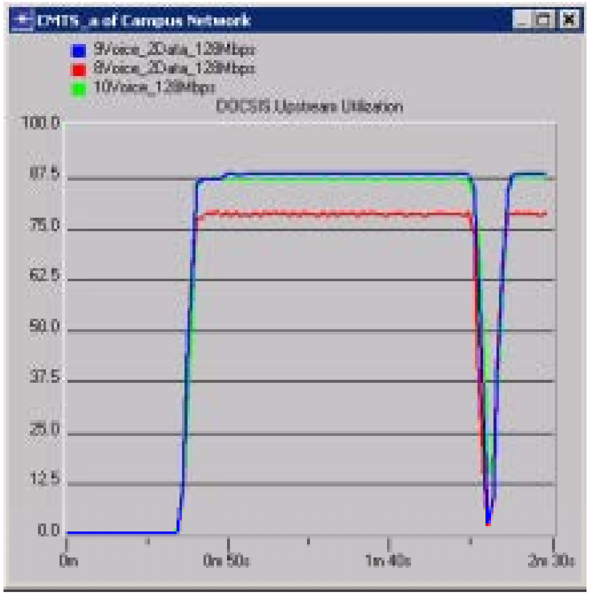

Figure 12 shows that in this simulation, the DOCSIS upstream has about 78% utilization (or about 1 Mbps in this case) with 8 voice calls and 87.5% utilization with 9 and 10 specified calls (Call 10 was not completed). The simulation was also tasked with simultaneously collecting information on the performance of the data applications (modeled as uploads of file transfer protocol, or FTP, traffic from two cable modems to the CMTS).[8] Figure 13 provides a graph of the amount of data uploaded to the CMTS for the same time interval. Note that in the case with 8 voice conversations and two simultaneous data uploads, the performance is fairly constant. With 9 voice conversations, the data rate peaks sharply in the brief interval between voice conversations. Another view of the performance of the data transfer is provided by Figure 14, the response time for data traffic uploaded to the CMTS. This graph shows that in the case of 8 speakers, the file transfer takes place with a fairly constant delay of about 1 second.[9] In the case of 9 speakers, the data is buffered at the cable modem until enough bandwidth is available to transfer it to the CMTS.

Figure 13. Rate of data uploaded to CMTS while voice conversations are taking place..

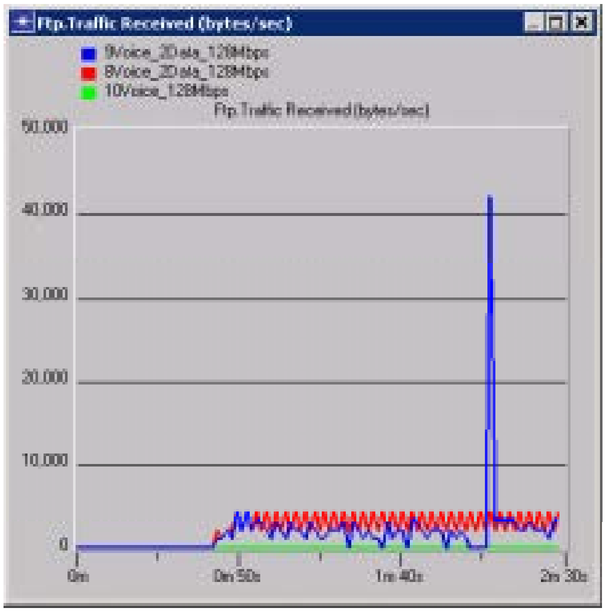

Figure 14. Data upload response time for various loading conditions.

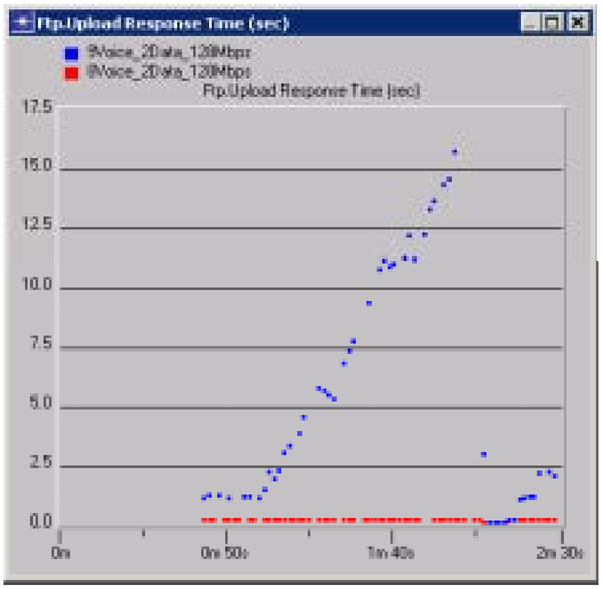

Taken all together, these preliminary results indicate that the DOCSIS 1.1 system with PacketCable protocol can operate at a utilization level of over 78%. This can be compared to earlier simulation results (circa 1999) that indicated that the DOCSIS 1.0 system performance would degrade rapidly at a utilization level of 65%. Thus, these preliminary simulation results indicate that DOCSIS 1.1 will have a 15% increase in carrying capacity over DOCSIS 1.0. Furthermore, Figures 13 and 14 indicate that DOCSIS 1.1 allows graceful degradation of the upstream under peak loading conditions in that bandwidth for support of lower priority applications is cannibalized to support higher priority applications.[10]

CONCLUSION

The use of cable-based high-speed networks has evolved substantially over the course of the last year, and will continue to evolve as new applications and services are deployed on the networks. This paper marks the first, preliminary steps in understanding this evolution of use. In order to manage this evolution, the cable industry needs to continue the collection of data, not only on the access network, but also at higher levels of aggregation so that we can determine the nature of local content management that will be required. The cable industry also needs tools ranging from specialized data collection and analysis techniques to innovative ways to manage service deployments in a tiered service environment. To this end, CableLabs will continue analysis of bandwidth use on live cable high-speed data systems, the development of tools for the modeling and management of data traffic, and assessment of the strategic technical and economic implications of these measurements in the coming year.

APPENDIX: Observed Parameters and Their Distributions

Appendix A describes the parameters used for the detailed analysis of data on the Bandwidth Management Project. This appendix also reports preliminary research findings on the probability distributions of these parameters based on live network traffic data.

A.1 Volume of data use

A.1.1 Definition

The amount of data consumed (uploaded, downloaded, and uploaded + downloaded) during the measurement interval. Specifically, this parameter corresponds to the number of bytes of information delivered to, or originating from, a single cable modem as represented by the MAC address[11]. This parameter can be collected for an aggregate (for example, all of the users on a given upstream or downstream frequency) and an individual (per MAC address) basis.

A.1.2 Use

This parameter is useful for examining clusters of users and pattern of use. It can also be used as a measurement mechanism for volume-based tiering (where a usage tier is defined in terms of, for example, number of Gigabytes per month).

A.1.3 Units of Measurement

This value is recorded in terms of bytes, and typically measured in Mbytes/(unit period) for cable modem networks. The unit period may be seconds, hours, days, or months.

A.1.4 Distribution

If **V** = Volume of data used by an individual user, the probability

$$P[\mathbf{V} < v]$$

has been observed to have an exponential distribution for upstream, downstream, and total volume of data used.

A.2 Data Rate

A.2.1 Definition

Theoretically, the data rate is the derivative of the volume of use with respect to time. Practically, the average data rate is merely an expression of volume of data used within a specified short time period, usually a second. It is usually collected for an aggregate (for example, all of the users on a given upstream or downstream frequency). The use by an individual is difficult to instrument at this time and is usually derived from a longer term volume measurement for the user.

A.2.2 Use

This parameter is typically collected by network management tools such as MRTG, and is the fundamental parameter for system capacity planning due to the small time interval of measurement and use in the definition of communication systems. The time trajectory of this parameter for successive measurement intervals for an aggregate of users, particularly to those users assigned to the same upstream and downstream frequency of operation, is a fundamental tool for use in capacity analysis. The distribution characteristics of this parameter are useful in predicting the performance and QoS characteristic of a network segment.

A.2.3 Units of Measurement

Typical measurement units for this parameter are bits/second (bps) and bytes/second (Bps) (8 bits = 1 byte).

A.2.4 Distribution

For aggregates of users over typical measurement intervals (5 minutes, 15 minutes), the time trajectory of this parameter has the characteristics of a normally distributed, or Gaussian, white noise. That is, individual observations of this parameter are normally distributed with mean equal to the average data rate. The standard deviation of this parameter can be used in conjunction with the mean as the parametric basis for system capacity estimation. The diurnal variation $\mathbf{X_t}$ in the data rate load on an upstream or downstream has the formal functional form (at time t) as a stochastic process:

$$\mathbf{X_t} = \mathbf{R(t)} + \mathbf{n_t}$$

Where

R(t) is a sinusoidal function expressing the mean data rate, and

$\mathbf{n_t}$ is a zero mean Gaussian white noise with standard deviation much less than **R** since the data rate will always be positive.

The expected accumulated volume of use is related to data rate by the expression:

$$\mathbf{V(T)} = \int_{[0,T]} \mathbf{R(t)}\, dt$$

Since the expected mean of $\mathbf{n_t}$ is 0.

A.2.5 Derivation

The MRTG tool is commonly used to collect and report data flow rates. In order to collect the data, MRTG will poll periodically a router to collect the value of a byte counter to get a value V(n) (the value of the counter at polling interval n). In order to determine the data flow in bytes per second during the period, MRTG computes

F(n) = [V(n)-V(n-1)]/(number of seconds in reporting period)

In order to get the flow value in bits per second for the reporting period, it is enough to compute 8*F(n). The normality of this distribution is to be expected from the Central Limit Theorem which states, in effect , that if one takes random samples of size ***n*** from a population of mean ***m*** and standard deviation ***s***, then, as ***n*** gets large, then the distribution of the average of these samples ***X*** will approach the normal distribution with:

- Mean = ***m***
- Standard Deviation = $\boldsymbol{s/\sqrt{n}}$

A.3 Symmetry of Data Use

A.3.1 Definition

Downstream Data used per measurement interval/Upstream data used per measurement interval.

A.3.2 Use

This parameter has many implications for system architecture and a fundamental indicator of system usage by subscribers.

A.3.3 Units of Measurement

This parameter takes a dimensionless positive value. Values in the range of 0.5 to 4 are typical. It can be expressed as a ratio R or in Decibels (DB, 10 $\log_{10}$(R)). The data used can be measured in volumetric or flow-rate terms, and can be used to measure populations and individuals. The key item is to maintain the same interval of measurement for comparison.

A.3.4 Distribution

If **R** = symmetry of data use by an individual then

$$P[\mathbf{R} < r]$$

Has been observed to have a log normal distribution when expressed as a ratio. When expressed in DB, the value $\mathbf{R_{DB}}$ = 10 $\log_{10}$(**R**) follows a normal distribution.

REFERENCE NOTES

[1] The views expressed in this paper are solely those of the author. The author appreciates detailed manuscript comments from Simon Krauss, Dorothy Raymond, and David Reed and manuscript preparation assistance from Kathy Mitts. The author, however, bears full responsibility for any errors of fact or interpretation in the paper.

[2] MRTG (Multi Router Traffic Grapher) consists of a Perl script which uses SNMP to read the traffic counters of routers and a C program which logs the traffic data and creates graphs representing the traffic on the monitored network connection. These graphs are embedded into web pages that can be viewed from any modern Web-browser. MRTG is available as a download from: http://mrtg.hdl.com/

[3] A working hypothesis for this is that, in many cases, the primary residential users of the system will be of school age. A holiday from school provides an opportunity to use the system. This effect is also seen in the diurnal variations. It has also been observed that systems serving communities with 12-month schools do not exhibit seasonal fluctuations to the same extent as communities with 9-month schools.

[4] In this paper, the term **node** refers to the group of subscribers that receive their cable modem service from the same blade in the CMTS.

[5] Typical background traffic consisted of a mixture of simulated Email, web page requests and downloads, and file transfers.

[6] The model is capable of using upstream channel data rates of 1.28, 2.56, 5.12, and 10.24 Mbps. 1.28 Mbps was used in order to reduce complexity during scenario development and testing and to explore the inefficiencies inherent in lower data rate systems. In general, DOCSIS systems work better when run at higher data rates due to the benefits of statistical multiplexing.

[7] The MTA is specified in the PacketCable specifications. In the simulation, it is modeled as a cable modem that can support both data and PacketCable protocol sessions.

[8] The file transfers used in this simulation were made at a simulated data rate of 1.28 Mbps (the label stating 128 Mbps in the figure legend is a misprint. The legend will not support the use of a decimal point).

[9] This delay is an artifact of the structure of the simulation.

[10] The issue of when an upstream DOCSIS 1.0 upstream channel is nearing its capacity limit is a difficult one to address. Most of the evidence in this area is anecdotal. The capacity limit depends on a number of factors including: The sophistication of the scheduling algorithm in the CMTS; the upstream channel line rate; the types of applications run by network users; and the number of MAC addresses (cable modems) that a CMTS blade can efficiently support

[11] The MAC address (Media Access Control Address) is the unique hardware number assigned to network connection devices such as cable modems.

BROADCAST TRIGGERS INSERTION FOR DIGITAL ENVIRONMENTS

Hervé Utheza
Vice President Product Management & Solutions Marketing
Liberate Technologies

Abstract

As deployments of interactive television gain momentum, triggers that enable enhancements to TV programs and commercials will increasingly be inserted into the broadcast stream so the enhancement is synchronized with the video programming. Enhanced broadcast and synchronization will become the norm in digital, as well as analog environments.

Delivery of enhancements poses great challenges for existing networks. Primary considerations include the traffic bursts to the network invited by triggers; upstream and downstream system bandwidth; Internet backbones access and servers; and the limitations of some set-top boxes to receive/extract triggers.

This paper provides an overview of a solution for the economical delivery of enhancements on existing networks. As this solution could be applied globally, this paper is focused on solutions for systems in North America.

ENHANCED TELEVISION

Enhanced television consists of providing additional information and interactivity (enhancements) along with the normal television programming. Viewers can access these enhancements as they watch a chosen program. Such enhancements are implemented using consumer device software, such as the Liberate TV Navigator^{TM} along with network servers and broadcast equipment that complements the traditional video/audio broadcast network. The consumer device is a set-top box that is connected to the video distribution network (generally cable or satellite). The enhancement itself can be created offline to the TV programming, or can be created on the fly for a live enhancement (such as a sports broadcast). Content authoring and insertion is a major part of the enhanced broadcast chain.

Enhanced television covers material that is, or not (depending on the case) *synchronized* to the video and audio component of a program; synchronization can be of various degrees of accuracy, and does not often need to match the concept of *frame accuracy*. It also tends to have some form of interactivity, implemented either using features of the receiver alone (one-way broadcast) or using servers via a two-way network connection. With a two-way network connection, this interactivity can be used to enable e-commerce transactions – a combination of television and e-commerce often called *t-commerce*.

CHALLENGES FOR DELIVERY OF ENHANCEMENTS ON EXISTING NETWORKS

The end-to-end delivery of enhancements synchronized with the video programming introduces a wide variety of new services for the viewers and new business opportunities for the broadcasters and network operators. However, the delivery of such services has inherent technical challenges, which must be solved for a successful launch.

Enhanced Television synchronized content creates burst in the network

Insertion of enhancements at a very specific time in a video broadcast causes peak subscribers demand for bandwidth on the return channel. At a given point in time, hundreds of thousands of subscribers will receive a trigger and then request the enhanced content, increasing drastically the upstream and downstream peak bandwidth requirements on the local network.

Synchronized enhancements involves multiple players in the end to end video distribution chain

Adding synchronized enhancements to video programming involves every entity from the broadcast studio all the way to the subscriber. The broadcast studio usually has responsibility for the creation and insertion of TV enhancements. Synchronization is achieved through insertion of triggers with the video signal (assuming no alterations from the video distribution network). From content authoring to synchronization to delivery, it is crucial to define the various interfaces allowing proper end-to-end integration.

The content insertion requires a proper scheduling mechanism and content caching scheme. National networks also have local affiliates and the insertion of local enhanced content brings a new level of complexity to the picture.

There are multiple network topologies and video transmission schemes

The number of players in today's video distribution chain complicates the enhanced data delivery mechanism. Broadcast studios can usually send the content to a satellite uplink in analog or digital format. Network operators receive the content from the satellite downlink, package it and distribute it to their subscribers through their specific networks. Each distribution network has its own characteristics in terms of possible downstream and upstream bandwidth configurations, terminal (set-top box) capabilities, and network management (for example, it may not be possible to extract triggers on a given digital set-top). Depending on the end-to-end network configuration, transmission can be in analog or digital, potentially altering the form of the enhancement as it is converted between transmission formats

END-TO-END SYSTEM CONSIDERATIONS

Enhanced television consists of providing additional information and interactivity (enhancements) along with the normal television programming. Viewers can access these enhancements as they watch a chosen program. Such enhancements are implemented using consumer device software (possibly running a client middleware emgine such as the Liberate TV Navigator^{TM}) along with network servers and broadcast equipment that complements the traditional video/audio broadcast network. The consumer device is a set-top box that is connected to the video distribution network (generally cable or satellite). The enhancement itself can be created offline to the TV programming, or can be created on the fly for a live enhancement (such as a sports broadcast). Content authoring and insertion is a major part of the enhanced broadcast chain.

Enhanced television covers material that is, or not (depending on the case) *synchronized* to the video and audio component of a program; synchronization can be of various degrees of accuracy, and does not often need to match the concept of *frame accuracy*. It also tends to have some form of interactivity, implemented either using features of the receiver alone (one-way broadcast) or using servers via a two-way network connection. With a two-way network connection, this interactivity can be used to enable e-commerce transactions – a combination of television and e-commerce often called *t-commerce*.

Enhanced television brings together divergent pieces of technology and content. An end-to-end broadcast solution has to address the issues related to developing, delivering (broadcasting) and managing the different pieces of content (or assets) that make up an enhancement, and the applications that support them (for instance, the payment systems).

From an end-to-end system standpoint, a number of key functions need to be addressed in order to assure proper delivery of enhancements. Those functions include:

- Authoring, Hosting and Archiving the content
- Management of user responses (client and server)
- Transaction fulfillment
- Security
- Branding
- Enhancement advertising
- Injecting Triggers (and enhancement)
- Receiving Triggers

This paper focuses on the last two items from this list: Injection of triggers / enhancements and receiving triggers.

Injecting Triggers (and enhancements)

Scheduling and injecting triggers associated with the enhancement into the broadcast stream; in some cases the enhancement itself may also be injected into the broadcast, to improve scalability (when bandwidth is available). Insertion in both analog and digital transport carriers is possible, but the two differ significantly.

Receiving Triggers

Functions in the receiver provided by the middleware (set-top box) to allow it to receive and appropriately handle triggers. A distinction needs to be made between triggers that *announce* the presence of enhancements that might be offered to the viewer and triggers that *cause* the enhancements to function, once selected by the viewer.

SOME BACKGROUND ON ATVEF

The Advanced Television Enhancement Forum (ATVEF)[i] is of a cross-industry group of companies[ii] representing major television programmers, technical platform providers, broadcasters, and transport providers. ATVEF, which is not a standards organization, was created to specify the design, delivery and display of enhanced and interactive TV programming that can be authored once using a variety of tools, and deployed to a variety of television, set-top, and PC-based receivers. ATVEF's creation was driven by a much-needed standardization effort at a time when broadcast TV and the Internet are converging to deliver rich, cross-platform, cross-network services and content. Older Enhanced Television systems did not have the benefit of Internet content and technology. ATVEF utilizes the Internet as the cornerstone of the specification.

The ATVEF specification defines two aspects of the enhanced television platform: the Content and the Transport. The ATVEF content defines different level of content profiles, which are designed to target different range of receiver capabilities. The ATVEF transport defines different types of transport mechanisms for carrying ATVEF triggers and content supporting different network architectures.

The following section provides a brief overview of the ATVEF transport mechanism.

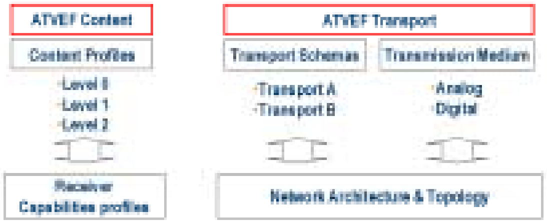

Figure 1 ATVEF Content and transport

ATVEF Transport

The display of Enhanced Television content consists of two aspects: 1) delivery of data resources and 2) display of named resources synchronized by triggers. All forms of ATVEF transport involve data delivery and triggers. The capability of networks for one-way and/or two-way communication drove the definition of two models of transport for ATVEF.

ATVEF defines two kinds of transport. Transport A is for delivery of triggers by the forward (or broadcast) path and the pulling of data by a (required) return path. Transport B is for delivery of triggers and data by the forward (or broadcast) path where the return path is optional.

ATVEF Transport A

With the Transport A delivery scheme, only the triggers are sent with the video / audio broadcast. In order to get the enhancement, the set-top must have a live network connection (two-way) via an ordinary telephony modem or cable return channel. Usually Internet protocols are used to load content in this manner.

ATVEF Transport A highlights:

- Triggers (URL) only are sent to the set-top box embedded with the video/audio stream.
- Corresponding content is accessed from the two-way network (e.g. phone, cable return)

Transport A is usually seen as a transport mechanism for analog video where triggers are carried over the VBI part of the video signal. Note that the VBI can also be carried on a digital MPEG video through the MPEG user data field.

ATVEF Transport B

In the Transport B delivery scheme, the triggers and the enhancements are both sent with the video/ audio broadcast. The enhancements are included in the video signal itself and do not require a real-time network connection for delivery of interactive content, however, this does not allow for targeting, personalization, or customization (a return channel is required for this). Transport B provides a less expensive alternative where delivery of a single enhancement to all set-top boxes is appropriate, and is particularly useful for satellite-delivered programming that lacks a real time connection.

The advantages of Transport B are that it places no additional load on the two-way network (the data is sent to all receivers), and it does not require a two-way connection (so it can run on video networks with no (or limited) return channel. As the triggers and enhancements are part of the broadcast feed, the level of interactivity provided with Transport B is limited to broadcast content (enhancements are interactive but there is no communication with the two-way network). Once a viewer selects the trigger, the set-top box retrieves the interactive content locally and displays the enhancement in real time, allowing the subscriber to see the enhanced content immediately without additional load on the two-way network.

For analog video transmission, the VBI is the only channel to carry triggers and enhancements. Since the VBI is not a high bandwidth channel, Transport B with rich enhancements is challenging on analog video. For that reason, Transport B scheme is usually seen as a transport mechanism for digital video where triggers and content are carried on MPEG transport.

ATVEF Transport B highlights:

- Triggers and Content are sent to the set-top box packaged together as insertions into the video/audio stream.

SOLUTION FOR SCALABLE DELIVERY OF ENHANCEMENTS ON EXISTING NETWORKS

In an environment where enough bandwidth is available from the broadcaster, it is possible to broadcast video, audio, triggers and enhancements from the origination point. In this scheme, the broadcaster would use a data broadcast server (such as Liberate Mediacast™) coupled with a scheduling system allowing them to add enhancements to their programming fairly easily. For end-to-end digital systems, this is possible assuming that enough bandwidth is available on the uplink.

For the case where the programming is originated in analog format, the VBI is the only channel to carry programming related information. This channel is limited and not very suitable to deliver rich content.

As the number of digital set-top boxes deployed on cable systems across North America is getting bigger and bigger, a solution for the delivery of enhancements on those existing networks is crucial in order to avoid hardware upgrade of large number of set-tops.

Model for end-to-end delivery

The remainder of this paper presents a hybrid model for end-to-end delivery of enhancements that:

a) Allows broadcasters to keep creating triggers and enhancements using the same infrastructure.

b) Allows support of triggers and enhancements of networks where the return channel and the processing capabilities of the digital set-top are limited.

This hybrid model carries timing related information (triggers) with the video programming and publishes the enhancement on a generally accessible channel (such as a proxy server). The model reduces the bandwidth requirements from the broadcaster origination point while maintaining the timing / synchronization information with the video transport. As the industry is migrating from analog to analog / digital hybrid to full digital a solution is needed to ensure the delivery of enhancements synchronized with the video.

Depending on the network capabilities either a Basic of Advanced model will allow a scalable deployment of synchronized enhancements.

The Basic model, which is the simplest, requires a stronger network infrastructure in terms of two-way connection. The Advanced model introduces data carousels (such as through the Liberate Mediacast™ server), reduces peak bandwidth requirements on the two-way network, and allows a more scalable delivery of synchronized enhancements without delivering the enhancement with the video signal from the origination point. This model also allows local enhanced content filtering and insertion.

In both models, the broadcaster inserts timing information (triggers) along with the video signal and then publishes the enhancement (usually via a web interface). As far as the broadcaster is concerned, this embodies the ATVEF Transport A model. Depending on its network topology, the network operator will either pass through the triggers to its set-tops or will extract the triggers allowing a pre-fetch and pre-caching of the enhancements.

Description of Basic and Advanced synchronization models

Basic model

For networks allowing it, the Basic model is the simplest. This is ATVEF Transport A end to end. In this case, Liberate Connect^{TM} servers are installed as part of the network infrastructure at the headend and Liberate TV Navigator^{TM} is loaded onto the set-top box at the subscriber site. *Figure 2* below shows the flow diagram for the Basic model.

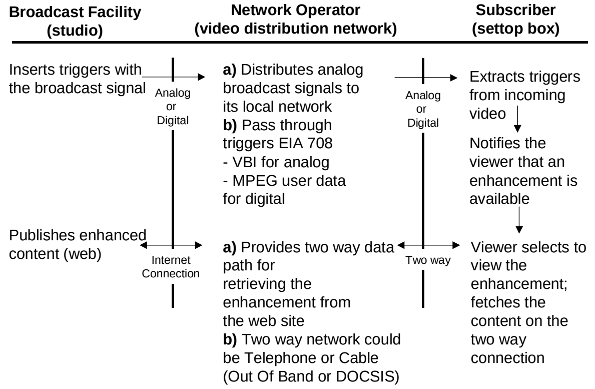
Broadcast Facility (studio)
Network Operator (video distribution network)
Subscriber (settop box)
Inserts triggers with the broadcast signal
Analog or Digital
a) Distributes analog broadcast signals to its local network
b) Pass through triggers EIA 708
- VBI for analog
- MPEG user data for digital
Analog or Digital
Extracts triggers from incoming video
Notifies the viewer that an enhancement is available
Publishes enhanced content (web)
Internet Connection
a) Provides two way data path for retrieving the enhancement from the web site
b) Two way network could be Telephone or Cable (Out Of Band or DOCSIS)
Two way
Viewer selects to view the enhancement; fetches the content on the two way connection

Figure 2 Basic model flow diagram

In this model, the network operator passes through the video/audio with embedded triggers to the set-top. Triggers are embedded in the VBI. For digital, the VBI is delivered through the user data field in MPEG. The software on the set-top requests enhancements over a two-way connection, typically telco or cable return. This model assumes there is enough downstream and upstream bandwidth to accommodate the bandwidth and latency demand for a large number of subscribers.

Basic model assumptions and limitations

This basic model is suitable only for networks that can sustain the peak bandwidth required by the retrieval of synchronized content.

Basic model assumptions

- The network operator infrastructure must have two-way capability with enough capacity and minimum return channel latency to support concurrent sessions for a large number of subscribers.
- The set-top must be able to extract triggers from the analog or digital video signal (Liberate TV Navigator, like other implementations, needs a minimum functionality from the set-top to achieve this).

Basic model limitations

- Limited scalability: synchronized content will introduce synchronized peaks demands from the subscribers. Requires strong 2-way network.

Advanced model

Today's network topologies with infrastructures that can support the Basic synchronization model are limited. In order to enable deployments of such synchronized services without engaging in the expense of a network upgrade, Liberate proposes an enhancement to the Basic model. The Advanced synchronization model achieves greater scalability and makes the end-to-end model feasible on today's networks.

Figure 3 below illustrates the information flow for the Advanced model.

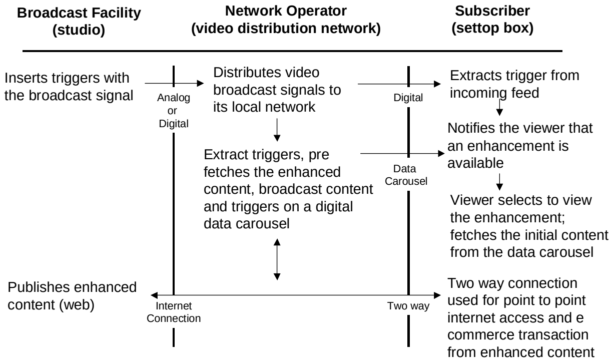
Broadcast Facility (studio)
Network Operator (video distribution network)
Subscriber (settop box)
Inserts triggers with the broadcast signal
Analog or Digital
Distributes video broadcast signals to its local network
Digital
Extracts trigger from incoming feed
Extract triggers, pre fetches the enhanced content, broadcast content and triggers on a digital data carousel
Data Carousel
Notifies the viewer that an enhancement is available
Viewer selects to view the enhancement; fetches the initial content from the data carousel
Publishes enhanced content (web)
Internet Connection
Two way
Two way connection used for point to point internet access and e commerce transaction from enhanced content

Figure 3 Advanced model flow diagram

In this model, the network operator passes through the video and audio to the subscriber in the analog or digital format. At the network operator site, the synchronization with triggers can be achieved by introducing infrastructure allowing greater scalability. The triggers are extracted at the head end, and then the enhancement can be pre-fetched and stored on an in-band carousel. This carousel is carried on the physical layers supported by the specific network. Examples are MPEG transport , Out Of Band channel or IP Multicast in a DOCSIS environment. At this point, the enhanced content is available to all subscribers. As soon as a trigger is sent through the broadcast server (such as generated by Liberate Mediacast™), the viewer is notified and, if the viewer selects the enhancement, it is retrieved from the carousel. The two-way connection is used only when the enhanced session goes interactive (e.g. user requests a link that is not on the carousel by going to an e-commerce or secure session).

The specific deployment for the Advanced model relies on carrying the data carousel over the local network. The physical layer to carry the data carousel is determined by the set-top capabilities in terms of data delivery. Carousels can be delivered over MPEG and over a DOCSIS channel (IP Multicast).

The following section provides an end-to-end case study for the Advanced synchronization model.

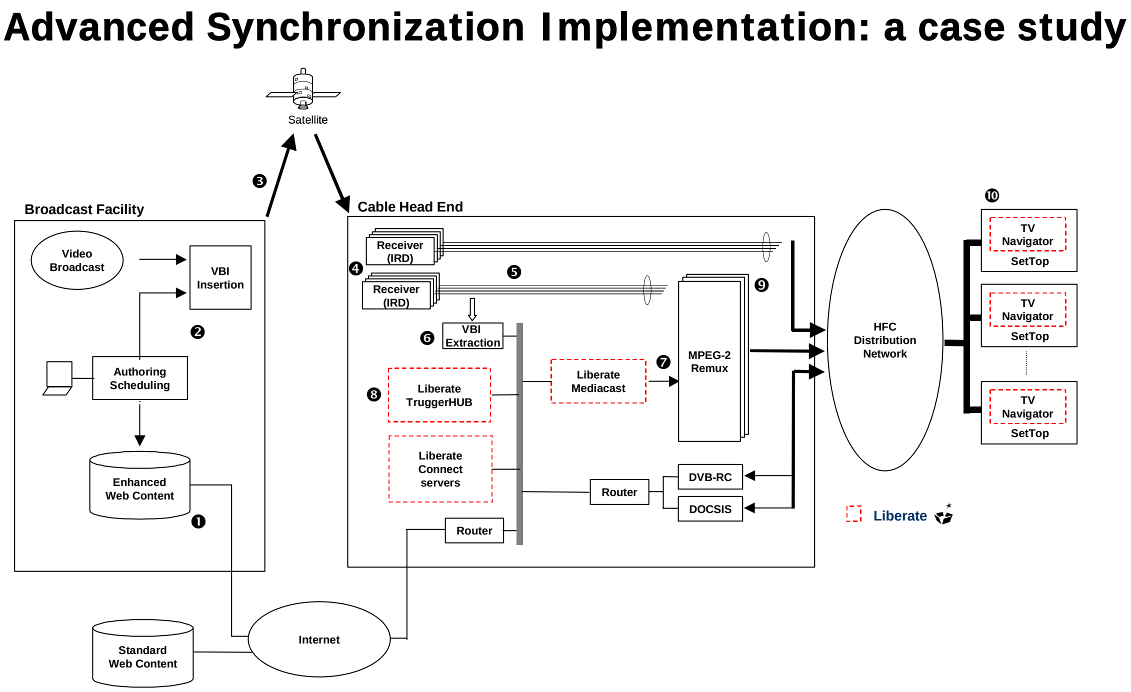

Figure 4 Advanced model end-to-end case study

This is a typical scenario of operation for the Advanced synchronization model for cable systems in North America:

❶ Content generation

The broadcaster (usually in conjunction with an ETV ASP (Application Service Provider) partner) creates the enhancement for a specific program. This enhancement is described through a play list (using XML standards) and provides location and timing characteristics (date, time, expiration) for the enhanced content. The enhancement is stored on a location accessible through the Internet backbone. It will be pulled at a later time from the Liberate infrastructure server that resides at the cable operator premises.

❷ Synchronization of information and trigger insertion

Using an authoring and scheduling system (multiple companies provide such systems), the triggers are created and inserted with the video signal, thus embedded in the broadcast feed, typically in the VBI (ATVEF Transport A). Scheduling information and enhancement (through play lists) can be made available (published) ahead of time (for the cable headend to fetch schedules and content). The broadcaster performs the insertion of triggers in the analog or digital domain. Signal transmission can be analog or digital but will preserve the embedded ATVEF triggers (ATVEF transport A, VBI, on analog or digital video).

❸ Broadcast of video content

The video is sent from the broadcaster through a network (analog or digital); typically satellite. The signal contains the ATVEF triggers embedded with the video signal. Again, the video signal could be analog or digital.

❹ Receive broadcast

The broadcast signal is received at the cable headend through a receiver (IRD). Depending on the network architecture, the network operator will either:

a) Broadcast the analog video to the network and let the set-top extract triggers from the analog video (Basic model).

b) Extract the triggers* at the headend, pre-fetch the enhanced content and send it on a data carousel through Liberate Mediacast™ (Advanced model).

* Trigger extraction can be done with external devices.

❺ Video distribution

Typically, the video programs (channels) are distributed to the subscribers in analog or digital format (some channels analog, some channels digital). The video delivery scheme (analog or digital) impacts the ATVEF support on the platform since the enhancement is transported to alternate paths. For example, for a single tuner set-top that is tuned to an analog channel, the enhanced content has to be delivered through the Out Of Band (OOB) in order to avoid tuning away and to be able to overlay the enhancement on the video programming.

❻ Triggers extraction and processing

For the specific Advanced synchronization model, the triggers are extracted at the headend. The Liberate TriggerHUB™ server is the link between the VBI extraction devices and the Mediacast carousels. Triggers are sent to the set-tops through Mediacast carousels using a notification mechanism. The Liberate TV Navigator™ allows the set-top to receive and appropriately handle triggers.

The user interface for offering enhancements to the viewer can be customized by the service operator through the TV Navigator™.

❼ Broadcast content

Liberate Mediacast™ is used to broadcast content through a data carousel scheme. The carousels are carried either in-band with the MPEG (single tuner set-tops) or through IP Multicast on the DOCSIS channel (dual tuners set-tops). The enhanced content is provisioned as part of a service that includes MPEG audio, video, and enhanced data. Mediacast™ enhances the system scalability since the enhanced content is made available to all subscribers served by the specific headend.

❽ Liberate TriggerHUB™ server

The TriggerHUB™ server resides at the cable head end and works in conjunction with Liberate Mediacast™ server. TriggerHUB™ pulls the schedule for the enhanced content associated with each video channel (from the broadcaster site). The TriggerHUB™ is the interface between the incoming triggers and the data carousel. It monitors triggers from multiple incoming sources, schedule them to the Liberate Mediacast™ server.

The first trigger of a program could include a <LINK> tag that refers to a play list. The TriggerHUB™ fetches the initial page, extracts the link, fetches the play list, creates a schedule from it, and hands it to the Mediacast™ server for carousel broadcast. This mechanism allows the system to prepare the content in advance, hence reducing the latency on subsequent triggers.

The broadcaster can push play lists to the cable headend's TriggerHUB™ server allowing real time support of changes in the program schedule.

❾ MPEG re-multiplexing

When the target set-top is a single tuner set-top, the carousel is carried on MPEG and combined with the video and audio programs. An MPEG re-multiplexing device is used to combine the Video and Enhanced content (data carousel) on a single transport. The audio, video, and data are combined into a specific service allowing simple management and authorization through an existing Conditional Access System (CAS). The MPEG re-multiplexing device is generally part of the network for grooming video programs on specific MPEG transports.

❿ Subscriber

The set-top runs Liberate TV Navigator™ allowing it to receive, synchronize, and display the enhanced content. The set-top, when tuned to an enhancement-enabled channel, will then be able to notify the subscriber of enhanced content and retrieve it from the Mediacast™ carousel without tuning away. The carousel either resides on the same digital transport stream (MPEG) for single tuner set-top, or on the DOCSIS channel (IP Multicast) for dual tuner set-tops.

CONCLUSION

The successful deployment of Enhanced TV services requires a close coordination of Network operators, Content developers and Head End equipment vendors to overcome the many inherent problems in distributing synchronized enhanced content. Both the content and the delivery methods must accommodate the realities of the network. A platform based on standards such as ATVEF goes a long way towards ensuring a homogeneous content, allowing consistent processing. To handle the network requires a system that is capable of efficiently delivering content over networks with varied types of two-way backchannels.. Liberate facilitate this integration by offering a platform based on existing standards (such as ATVEF). From a content standpoint, the Liberate Enhanced TV platform allows common content and tools across different network topologies and set-top boxes capabilities. From a networking standpoint, the platform, through Liberate Mediacast™, ensures a delivery of enhancements which does not rely on a back channel and offers the network operator a direct control over the bandwidth allocation for the delivery of such enhancements.

FOOTNOTES

[i] ATVEF Web site: http://www.atvef.com

[ii] ATVEF was founded by a group of 14 companies: CableLabs, CNN Interactive, DIRECTV, Discovery Communications, Inc., Intel Corporation, Liberate Technologies, Microsoft and WebTV Networks, NBC Multimedia, NDTC Technology, Inc., Public Broadcasting Service (PBS), Sony Corporation, Tribune, The Walt Disney Company, Warner Bros. Over 130 companies worldwide have signed licenses to implement the ATVEF content specification.

CABLE INDUSTRY CONSIDERATIONS IN CHOOSING WIRELESS HOME NETWORKING TECHNOLOGY

Yigal Bitran, CTO
Broadband Communications Israel, Texas Instruments

Abstract

Home networking evolved over the last few years into an affordable technology that can be applied to cable customers. Many home networking alternatives were proposed over the last few years. Wireless home networking emerged as the most promising solution in terms of consumer and operator benefits.

We will compare the wireless home networking variants to other alternative home networking technologies and discuss the wireless home networking options. We will focus especially on IEEE 802.11 and its various extensions, and discuss what features are important for addressing the Cable industry needs and their importance to enable new revenue opportunities.

HOME NETWORKING OVERVIEW

Home networking has rapidly emerged in the last three years and is considered today as an important element for home connectivity.

Many home networking alternatives exist to date with each alternative further divided into flavors, which makes it very confusing for the consumer and operator. The main home networking alternatives are:

- Wireless home networking: Utilizing RF technology to create a wireless local area network. This category includes IEEE 802.11[1], HomeRF[2], and HiperLAN2

- Home Phoneline networking: Using phone lines to create a home network. This category includes HPNA1.0 [3] and HPNA2.0

- Powerline home networking: Utilizing AC powerline network as the home network media. This category includes HomePlug [4] and X-10 technologies.

- Ethernet: This good and mature technology can be used for networking the home, however unlike the previous technologies, requires addition of CAT5 wires across the home

Table 1 details the technical attributes of the key alternatives

Table 1 – Home Networking Standards

Standard	*PHY rate*	*QoS*	*Wiring*
802.11	11-54Mbps	Full (11e)	No
HPNA2.0	16-32Mbps	Weak	Phoneline
HomePlug	14Mbps	Weak	Powerline
100BaseT	100Mbps	No	CAT5

Among the various technologies, wireless LAN is gaining momentum as the consumer's technology of choice. Its most important feature is avoiding the need to install new wires, while still having the freedom to roam throughout the house and use a computer where desired. In addition, many laptops are already equipped with WLAN cards used in the office environment. This drives higher availability of WLAN stations. Powerline networking may become attractive depending on the ability to overcome the Powerline noise issues and the availability of very low cost NICs that connect to the power line network. Phoneline networking was attractive in the past as the only 10Mbps technology available but is losing momentum as users prefer using wireless technologies.

Among the various wireless technologies, IEEE 802.11 is emerging as the technology of choice mainly for its wide industry support and higher rate compared to HomeRF. Compared to HiperLAN2, 802.11 is more suitable for North America and can provide equivalent rate and range depending on the chosen 802.11 extension.

IEEE 802.11 FLAVORS AND ADVANCED TECHNOLOGIES

Most people are well familiar with IEEE 802.11b, or Wi-Fi, as the standard for wireless LAN. There are, however, many extensions to IEEE 802.11 aiming to improve rate, range, QoS and security compared to the 802.11 baseline.

Table 2 summarizes the relevant IEEE 802.11 standards and draft standards (*italicized*).

Table 2 – IEEE 802.11 Standards

Standard	*Layer*	*Features*
802.11b	PHY	Baseline, 2.4GHz
802.11a	PHY	5GHz OFDM modulation up to 54Mbps
802.11e	MAC	QoS features
802.11g	PHY	Up to 54Mbps in 2.4GHz
802.11i	MAC	Improved security

802.11b provides up to 11Mbps physical layer rate with carrier frequency at 2.4GHz range. The modulation is based on Direct Sequence Spread Spectrum (DSSS) using Complementary Code Keying (CCK) and optional Packet Binary Convolutional Code (PBCC) single-carrier technologies. IEEE 802.11b standard defines PBCC at 5.5 and 11Mbps. There is an additional 22Mbps extension supported by current generation silicon solutions. A receiver that implements a PBCC convolutional decoder properly provides 3dB-coding gain (over CCK), which translates into either 70 percent extended coverage or into a higher rate at a given range (e.g. 5.5 to 11Mbps).

The effective TCP/IP [5] throughput of 802.11b at 11Mbps rate is reduced from the theoretical 11Mbps to 5-6Mbps when considering the MAC layer overhead (see figure 1). Each transmission of a 1460-byte IP packet (1060 µsec) is preceded by inter-frame spacing (60 µsec), 300 µsec back-off time on average, 72 or 144 preamble bits transmitted in 1Mbps (72/144 µsec), MAC header (24/48 µsec) and 10 MAC overhead bytes (55 µsec). An 802.11 acknowledge follows each data packet in addition to IP acknowledge packet adding more overhead. The overall overhead accounts for 50 percent of the total time and reduces the effective IP throughput of 802.11b to 5-6Mbps.

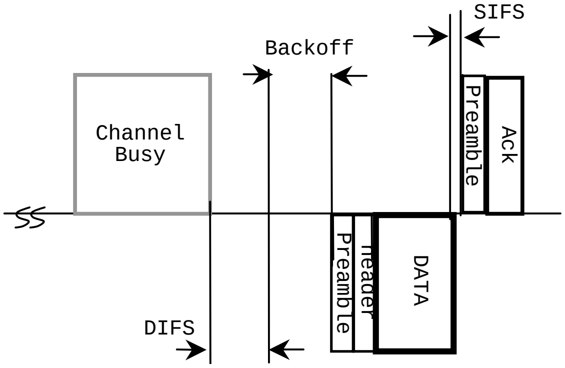

Figure 1 - 802.11 Distributed Access Control

IEEE 802.11a provides a PHY layer based on 5.2GHz RF frequency and high rate modulations from 6 to 54Mbps. IEEE 802.11a is based on multi-carrier OFDM technology where each individual tone can support up to 64QAM modulation, allowing very high rates.

A tradeoff exists between range and rate at 5.2GHz RF such that the rate falls back rapidly as the distance between the station and the access point increases. A disadvantage of IEEE 802.11a is the use of 5.2GHz frequency and OFDM modulation, which are not interoperable with the widespread 2.4GHz 802.11b technology.

The new IEEE 802.11g draft specification defines two technologies that can increase bit rates at 2.4GHz while keeping backward compatibility to the widespread 802.11b technology. One technology is CCK/PBCC-11, used in 802.11b, while the other is OFDM, similar to 802.11a OFDM. IEEE 802.11g allows a mixed mode called CCK-OFDM and another mode called PBCC-22/33. PBCC-22/33 allows coexistence with legacy 802.11b stations while providing higher rates of 22 and 33Mbps for new stations. This 22Mbps mode is already supported in Texas Instruments 802.11b-only silicon solutions (e.g. ACX100) and wireless LAN access points and NIC solutions.

IEEE 802.11e working group focuses on enhancing the MAC layer QoS capabilities in order to support multimedia applications such as video, audio and voice. The MAC scheme used in 802.11 is a contention-based carrier-sense multiple-access with collision avoidance (CSMA/CA) mechanism with binary exponential backoff (BEB) called Distributed Coordination Function (DCF). In DCF, there is no need for a central coordinator, and each station can attempt accessing the network based on some DCF rules. In low network loads, most access attempts are successful on first attempt and stations are able to send packets over the network with low latency. In some cases, the channel is occupied by another station for a few milliseconds and the station needs to defer until the channel is idle again. Since the access method is based on "First-come-First-serve" a very high latency in the order of tens of milliseconds can be created, especially in very high network load scenarios. A well-known example would be a file transfer between two PCs or between a PC and a printer that would block QoS-sensitive traffic. Another drawback of a contention based access mechanism is the low efficiency usage of available channel bandwidth due to collisions and backoff mechanisms. Results attained over Ethernet networks using a similar contention based access mechanism show that for multiple device networks the effective throughput can go as low as 10 percent of the actual payload data-rate.

IEEE 802.11e offers improved mechanisms to solve the above issues. The change from legacy 802.11 MAC to 802.11e can be compared to the change made in the DOCSIS [5] MAC when moving from DOCSIS 1.0 to DOCSIS 1.1. The most relevant mechanism is Hybrid Coordination Function (HCF), which supports a mix of contention based as well as centrally coordinated access. In HCF mode, a central coordinator (called hybrid coordinator, HC) is defined. HCF supports both prioritized QoS as well as parameterized (sometimes referred to as guaranteed) QoS. This is done through the support of prioritized traffic categories (TCs) as well as parameterized traffic streams (TSs), which are similar to service flows in DOCSIS 1.1. Stations can request bandwidth reservation from the HC, and a scheduler that resides in the HC controls the admission into the channel and may support QoS-critical applications such as voice, audio and video using the parameterized QoS mechanisms. Other types of traffic can be supported using the prioritized contention based mechanism. The centrally controlled channel access also makes the channel usage much more efficient, allowing for much higher effective data rates. This is achieved by eliminating contention intervals and ACK overhead or by using burst acknowledges mechanism. In an HCF centrally controlled scenario of streaming video for example, 80 percent efficiency can be easily achieved.

IEEE 802.11i introduces additional security features beyond the 802.11 baseline. IEEE 802.11i supports 40-bit Wireline Equivalent Privacy (WEP), 128-bit WEP and Advanced Encryption Standard (AES) algorithms as well as improved mechanisms for authentication, significantly reducing the risk for hacker attack on private data.

There are other IEEE 802.11 extensions offering additional improvements over the baseline 802.11, however, they are less relevant for the Cable industry.

IEEE 802.11 APPLICATION TO CABLE REQUIREMENTS

IEEE 802.11 is a generic technology that can be applied to enterprise, consumer and even access environments. When considering IEEE 802.11 for Cable, one needs to consider the specific requirements relevant to the Cable operators and end-users:

- Installation: This is one of the most important factors affecting both the end users and the Cable operators. Obviously, wireless LAN does not require installation of new wires, which makes it appealing as a technology. However, the issues of installing drivers and configuring the network need to be addressed. Robust installation software is key to smooth installation process. Provisioning of network addresses can be resolved by CableHome[7] and smooth installation can be resolved by utilizing plug and play technologies[8].

- Supporting multimedia applications: Unlike enterprise environments, home environments are expected to have richer multimedia traffic including audio, video and multi-player gaming. Supporting video is critical for Cable operators as this is traditionally their key business and value proposition over competing access providers. Initially it is expected that supporting low rate, streaming IP video traffic will be required. Rates will vary starting from 100kbps up to 750kbps per stream, with moderate requirements for QoS. At a later stage, operators and service providers would like to offer broadcast quality video with MPEG rates between 4-6Mbps. Video conferencing also could be offered either using PC Web cameras or special device. Ultimately, multiple video streams mixed with data, gaming and music could be envisioned, requiring very high rate and QoS performance of the wireless network.

- Voice support: As many cable operators would like to offer voice as an additional service, Cable Gateways will include PacketCable [9] functionality as a standard feature. PacketCable phones will be connected to RJ11 jacks at the Cable Gateway. Another possibility is using 802.11-based cordless phones that will provide both normal cordless phone functionality and smart terminal functionality, enabling additional service revenue.

In order to support the requirements above, the Cable industry will need to utilize wireless LAN technologies beyond the basic 802.11b. This could happen in phases as services offered evolve (see figure 2):

- Phase I (today): Support basic data connectivity with streaming IP video at low-moderate rates. IEEE 802.11b is sufficient to support this level of service

- Phase II: In addition to basic data connectivity service, phase II will support broadcast quality video distribution from residential gateways or set-top-boxes to remote TVs and PCs. This level of service requires higher physical rate going from 11Mbps to 22/24Mbps (with 12-16Mbps effective payload rate) and IEEE 802.11e for QoS support and higher payload efficiency. This combination of high effective throughput and QoS support will enable reliable distribution of broadcast quality video even when the network is loaded with other traffic. For backward compatibility with phase I solutions, 802.11 networks need to use 2.4GHz and thus 802.11g is preferred over 802.11a as it will support backwards compatibility with legacy equipment.

- Phase III: In addition to Phase II, Phase III will support multiple video streams between PCs, DVD player and TVs, few voice and audio streams and data traffic. This level of service will require the highest rate, as in IEEE 802.11g 54Mbps mode, and full usage of 802.11e HCF in order to provide guaranteed bit rate service to the end user.

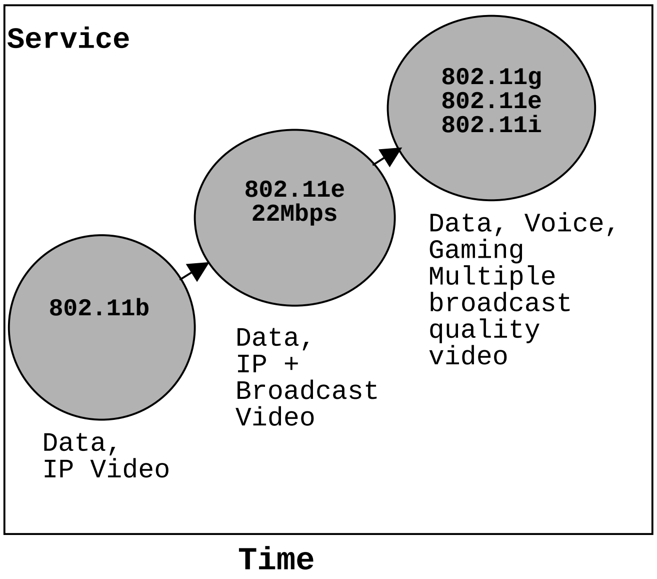

Figure 2 - Evolution of Wireless LAN for Cable Applications

CONCLUSION

From the multiple home networking technologies that can be applied to the Cable environment IEEE 802.11 seems to provide the best solution to the Cable industry needs. IEEE 802.11 different extensions were analyzed in the context of Cable industry requirements. An evolution of features is expected, starting with IEEE 802.11b for basic service, continuing through 22Mbps rate for providing broadcast quality video and finally onto IEEE 802.11g and IEEE 802.11e to provide ultimate rate and QoS features supporting multiple broadcast quality video streams, voice, audio, gaming and data services.

REFERENCES

[1] http://standards.ieee.org
[2] http://www.homerf.org
[3] http://homepna.org/
[4] http://www.homeplug.org
[5] http://www.cablemodem.com/
[6] http://www.ietf.org/
[7] http://www.cablelabs.com/cablehome/
[8] http://www.upnp.org/
[9] http://www.packetcable.com/
[10] http://grouper.ieee.org/groups/802/11/

[11] “Delivering High QoS Broadband Services into the Home via DOCSIS Cable Access and Wireless Home Networks”, Ofir Shalvi and Noam Geri, Texas Instruments, NCTA 2001

CLONED IDENTITY THREATS IN PACKETCABLE^{TM}

Alexander Medvinsky, Jay Strater
Motorola Broadband

Abstract

MTA clones might lead to service theft, breach of privacy, and denial of service. This paper proposes techniques that may be utilized by an IP Telephony service provider to detect and disable cloned MTAs and investigates what MTA configurations make sense for tamperproof hardware. It also considers techniques involving CMTS and DHCP server configuration and filtering options for limiting the geographic distribution of MTA clones.

INTRODUCTION

PacketCable [2] provides a set of specifications for VoIP services layered on top of DOCSIS-based HFC networks [7]. Denial of service threats that could disrupt an IP Telephony network, phone service theft and user privacy issues were all considered in the PacketCable security design. The PacketCable security specification [1] provides cryptographic protection that addresses these threats at a protocol level. But is protocol-level security enough to address these threats?

This paper considers a particular class of threats due to illegal duplication of the PacketCable client (MTA) identities. Since PacketCable provides cryptographic security, in order to duplicate an MTA identity one would need to make a copy of the MTA private keys and certificates in addition to copying the MTA host name, IP address, and MAC address. Let us say that an owner of a legitimately purchased (or leased) MTA proceeds to duplicate its identity into a number of illegal clones. What kind of threat does it pose to IP Telephony service providers? The paper discusses scenarios where the use of MTA clones might lead to service theft.

The PacketCable security team also took these cloning scenarios into consideration and the PacketCable security specification includes a discussion of these threats. Two main techniques that could be used to prevent clones are Fraud Detection/Prevention services and tamperproof hardware inside the MTA that would make duplication of cryptographic keys difficult. Because these techniques do not require inter-operability and because both of them affect either the cost of running an IP Telephony network or the cost of the MTAs, PacketCable does not provide specific requirements in this area.

This paper proposes techniques that may be utilized by an IP Telephony service provider to detect and disable cloned MTAs and investigates what MTA configurations make sense for tamperproof hardware. The proposed fraud detection and disabling techniques are based on particular properties of the Kerberos/PKINIT protocol that is utilized by PacketCable to distribute cryptographic keys to the MTAs. Fraud management puts an additional burden on the operator and if improperly administered could result in an uncomfortably large number of false alarms. Therefore, it is desirable to complement fraud management with tamper-resistant key storage in the MTAs.

Cost effectiveness of tamper-resistant storage in the MTA seems to largely depend on what other services are provided by that MTA. If it is a stand-alone MTA device that provides nothing but interactive VoIP services, secure storage can only be used to protect PacketCable cryptographic keys. If it is an MTA that is integrated with a settop box, it is resident on a platform that already has a very high motivation for secure key storage in order to prevent theft of broadcast video. In such an environment, an MTA benefits from secure key storage that is already present on that platform at little or no added hardware cost. Other integrated MTA platforms are also considered in this paper along with the corresponding motivation to utilize secure key storage.

This paper also addresses an MTA cloning threat that is better addressed with cloning prevention at the DOCSIS level rather than at the PacketCable application level. The paper discusses a denial of service scenario where a malicious adversary using a false identity is able to fool a CMTS into dropping valid downstream packets destined for some MTA. This threat is based on the fact that a CMTS will allocate a gate for each phone call that is authorized for a specific quality of service and has a specific bandwidth limit. If an MTA were to receive packets at too high of a rate, the CMTS would be forced to drop some of them. In order to orchestrate such an attack, this adversary need not know any of the cryptographic keys of another MTA. This paper proposes a solution to the problem that involves the CMTS with a cost of some administrative burden.

Finally, this paper considers techniques involving CMTS and DHCP server configuration and filtering options to severely limit the geographic distribution of MTA clones and, therefore, the viability of MTA cloning operations.

PACKETCABLE^{TM} ARCHITECTURE OVERVIEW

PacketCable is a project conducted by CableLabs. The project goal is to identify and define standards used to implement packet based voice, video, and other real time multimedia services over cable systems. PacketCable products are a family of products designed to deliver enhanced communication services using packetized data transmission technology over the HFC data network using the DOCSIS protocol. PacketCable products overlay the 2-way data ready, broadband cable access network. Initial offerings are packet voice. Packet video and other multimedia are longer term goals that are just now starting to be addressed by the PacketCable MultiMedia project.

The following diagram shows the PacketCable reference architecture (also see [2]).

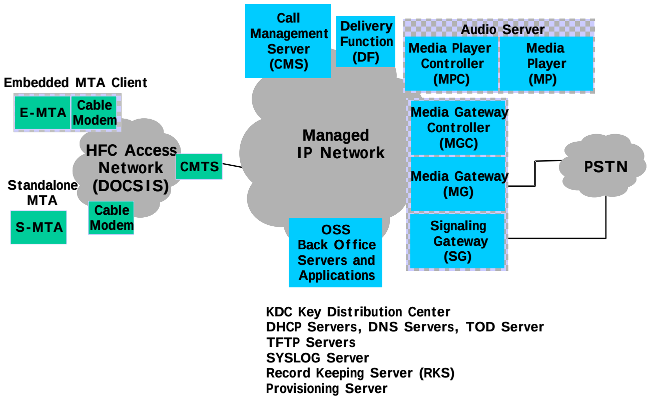
Call Management Server (CMS)
Delivery Function (DF)
Audio Server
Media Player Controller (MPC)
Media Player (MP)
Embedded MTA Client
E-MTA
Cable Modem
HFC Access Network (DOCSIS)
CMTS
Managed IP Network
Standalone MTA
S-MTA
Cable Modem
Media Gateway Controller (MGC)
Media Gateway (MG)
PSTN
Signaling Gateway (SG)
OSS Back Office Servers and Applications
KDC Key Distribution Center
DHCP Servers, DNS Servers, TOD Server
TFTP Servers
SYSLOG Server
Record Keeping Server (RKS)
Provisioning Server

Figure 1. PacketCable Reference Architecture

For the purpose of subsequent MTA security discussion, key elements in this architecture are the MTA, CMTS, CMS, MTA Provisioning Server, and KDC. The PacketCable signaling [3], bearer [5], and management protocols [4] between these components are shown in the following figure.

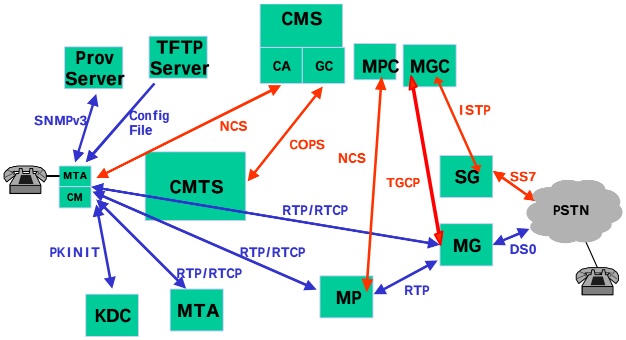

Figure 2. PacketCable Signaling, Bearer, and Management Interfaces

Following are a listing of the security protocol and key management techniques on the MTA interfaces (see [1] for details):

- Public key enabled Kerberos protocol between MTA and KDC
- Kerberized SNMPv3 on SNMP protocol between MTA and Provisioning Sever
- Hash and encryption of MTA configuration file retrieved from TFTP server
- Kerberized IPsec on NCS protocol between MTA and CA
- Cipher + MAC (Message Authentication Code) with NCS key distribution on RTP and RTCP protocol between MTA and MG, MP, and other MTA

IP TELEPHONY THREATS OVERVIEW

The IP Telephony system threats fall into three general categories:

1. **Service Theft**. An adversary manipulates the IP Telephony system in order to gain some financial benefit. For example, an adversary impersonates a valid VoIP subscriber and is able to make free long distance phone calls and charge them to the victim's account.

2. **Breach of Privacy**. An adversary is able to snoop on IP Telephony traffic (either signaling, management, or bearer channel) without a proper authorization.

3. **Denial of Service**. An adversary disrupts IP Telephony service, making the network completely non-functional, decreasing Quality of Service (QoS) below an acceptable level, or corrupting MTA configuration content.

PacketCable^{TM} security addresses all known theft of service threats on IP Telephony system interfaces that are within the scope of the PacketCable^{TM} project. Similarly, PacketCable^{TM} security addresses breach of privacy threats on PacketCable^{TM} interfaces that require privacy. Major denial of service threats resulting from unauthorized protocol manipulation are also addressed.

However, protocol and interface manipulation are not the only means by which an adversary may attack an IP Telephony system. Hacking into an IP Telephony server and disabling it would be an attack that should be prevented using various techniques such as firewalls, local access control, etc. However, since these local security measures do not require interoperability, they would fall out of the scope of protocol specifications such as PacketCable^{TM}. Similarly, identity cloning threats, where secret cryptographic keys are illegally extracted from a VoIP client and then distributed to other VoIP clients, should be addressed but are normally not covered by protocol specifications. The rest of the paper specifically addresses the cloning threats and how they may be prevented.

MTA CLONING THREATS

An MTA clone is a copy of an original, legally configured MTA that possesses the original MTA's identity (e.g., MAC address) as well as the original MTA's secret cryptographic keys. This enables the MTA clone to falsify its identity as if it were the original MTA.

In order to better understand the MTA cloning threats, it is helpful to identify the use of the various MTA identities within PacketCable™:

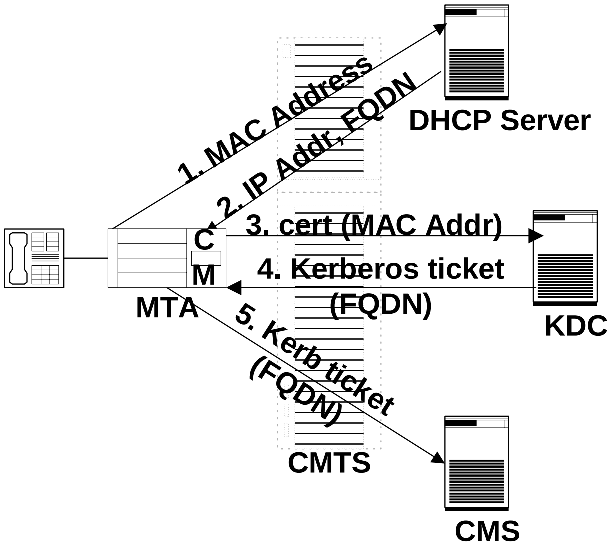

Figure 3: Usage of MTA Identities

Figure 3 shows that an MTA is authenticated by the CMS using a Kerberos ticket issued to this MTA, where the ticket contains the MTA Fully Qualified Domain Name (FQDN). The MTA FQDN along with the MTA port number identifies a specific VoIP subscriber.

The ticket certifying MTA FQDN is obtained from the Kerberos Key Distribution Center (KDC). The MTA provides its digital certificate with the MTA MAC address to the KDC and the KDC verifies the mapping of the MTA MAC address to its FQDN by performing a lookup into a subscriber database. (PacketCable™ defines a secure interface from the KDC to a subscriber database in order to perform this lookup.) Once the KDC verifies the mapping between the MTA MAC address and the FQDN, it returns a ticket to the MTA.

In addition to the MAC address and the FQDN, the other MTA identity used within PacketCable is its IP address. The DHCP server assigns an MTA an IP address based on its MAC address. Also, figure 3 shows that there is a CMTS located in the middle of all of the MTA's interfaces to an IP network. The CMTS has the visibility of the MTA's MAC address and its IP address inside the frame and packet headers respectively. The MTA is shown to be an embedded MTA, where the Cable Modem (CM) is integrated with the MTA as a single device. For the purpose of this paper, this is only an example – the same cloning threats and prevention measures apply for a standalone MTA that is not integrated with the CM.

An MTA clone in this architecture would:

- Somehow obtain a copy of the original MTA's device certificate and RSA private key.
- Use this certificate and private key to sign a ticket request for the KDC and the KDC would map the MAC address in the certificate to the original MTA's FQDN.
- Use the ticket with the original MTA's FQDN to establish security associations with the CMS.

As a result, the cloned MTA would make phone calls using the original MTA's subscriber account.

MTA clones in an IP Telephony system present the following threats:

- A subscriber authorized only for subscription services such as local calls and unlimited long distance minutes within a limited area can freely share the account with the clones. This subscriber could be a pirate that makes money from selling clones.

- A pirate might sign up using a false subscriber account with a stolen credit card number and then allow clones to make long distance calls. The pirate in this case has no intent to pay the phone bill. Eventually, the pirate account would be closed and the pirate might try to open another one with another false identity.
- In the case of a soft MTA implemented on a PC platform, the private key might be stolen by hackers on the Internet and then used to create clones. These clones could be used to charge phone calls to the victim's account or to disrupt the telephony network operation.

MTA CLONE DETECTION AND DISABLEMENT

Kerberos-Based Clone Detection and Disablement

The characteristics of the Kerberos key management protocol can be utilized to detect and disable MTA clones. Clone detection can be performed either by the KDC or by the CMS as explained in subsequent sections.

Clone Detection and Disablement by the KDC

A PacketCable™ KDC delivers a Kerberos session key to an MTA clone using a method called Diffie-Hellman key agreement. The Diffie-Hellman key agreement is part of a PKINIT extension to Kerberos that allows KDC clients to authenticate to the KDC using public key cryptography. The Diffie-Hellman key agreement is illustrated in the following figure:

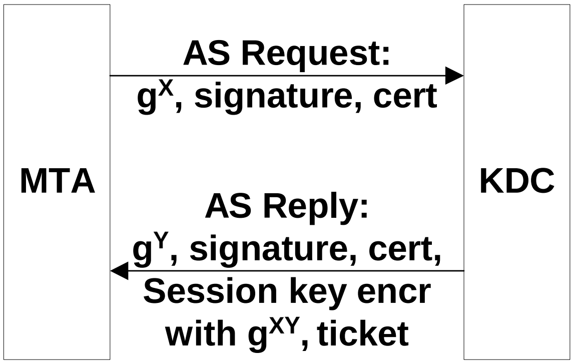

Figure 4: PKINIT with Diffie-Hellman Exchange

During a Diffie-Hellman key agreement, the MTA generates a secret value X, computes g^X and sends it to the KDC in the AS Request. It is not feasible to compute X from g^X within a reasonable amount of time.

The KDC in turn generates a secret value Y and computes both g^Y and $(g^X)^Y = g^{XY}$. The KDC then generates a unique session key and a ticket for this client and encrypts the session key with g^{XY}. The encrypted session key, ticket and g^Y are all sent back to the client in the AS Reply message.

After receiving the AS Reply, the client computes $(g^Y)^X = g^{XY}$ and decrypts the session key. A snooper that doesn't know the value of X or Y cannot figure out g^{XY} and thus cannot decrypt the session key. An MTA clone does not know X because X is generated on the fly for each ticket request. Therefore, MTA clones cannot snoop on the AS Reply message and determine the session key that was received by the original, legally authorized MTA.

In order for MTA clones to obtain their own tickets they each have to send their own AS Request and obtain their own unique session key. This makes the clones detectable at the KDC. When a KDC issues a ticket, it puts in a lifetime that specifies when this ticket is no longer valid. A MTA should keep reusing this same ticket until it expires. The following occurs when multiple MTA clones attempt to obtain tickets:

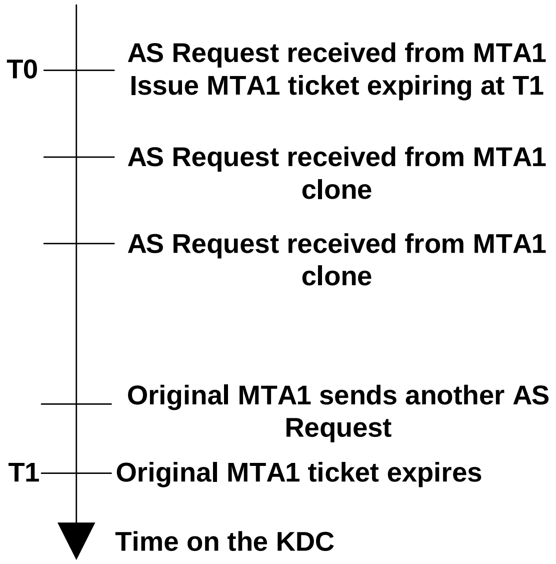
T0
AS Request received from MTA1
Issue MTA1 ticket expiring at T1
AS Request received from MTA1 clone
AS Request received from MTA1 clone
Original MTA1 sends another AS Request
T1
Original MTA1 ticket expires
Time on the KDC

Figure 5: Clone Detection by the KDC

After issuing a ticket to an MTA, the KDC could save the expiration time of that ticket. Normally, the KDC would not expect the MTA to request another ticket until the original ticket is almost expired. If the KDC receives AS Requests with the same MTA identity too early, it could be one of the following situations:

1. MTA clones are obtaining tickets
2. The MTA somehow lost its ticket and had to get another one.
3. Since PacketCable^{TM} uses Kerberos over UDP, the AS Reply packet can be lost and the MTA would time out and retry.

If the MTA has to retry because of UDP unreliability, the retries will be allowed only within a few seconds of the original request. Furthermore, the MTA can use an identical value of g^X while sending the retries. Therefore, the KDC should be able to distinguish the UDP retries from other unexpected AS Requests.

The MTA should not accidentally lose tickets very often if at all. Even if this were the case with a bad MTA implementation, such implementations should not be deployed until the software bugs are fixed. Therefore, if a KDC receives an early AS Request before the original MTA ticket expired and this request is not a retry, there is a very high probability that this is a request from an MTA clone. After one or two such early requests it is probably safe to flag this MTA as a clone threat and deny it any further tickets.

Clone Detection and Disablement by the CMS

Within the PacketCable^{TM} architecture, when the CMS receives an MTA ticket, it uses it to establish IPsec security associations with that MTA but does not need to save this ticket. If the CMS were to save at least the session key and the ticket expiration time, that would enable the CMS to perform clone detection.

The CMS would normally not expect the same MTA to send a different ticket with a different session key until the old ticket is either already expired or is close to its expiration time. When the CMS notices that the MTA changed its session key too early, it can be one of the following:

1. MTA clones are alternating at establishing security associations with the CMS in order for each to make a phone call.
2. The MTA somehow lost its ticket and had to get another one.
3. MTA had to retry due to the unreliability of the UDP transport.

The cloning detection at the CMS is analogous to that at the KDC. UDP retries would be limited to a short period after the original message and can be identified by a common attribute such as an IPsec SPI (Security Parameters Index). Also, a reasonable MTA implementation should not be accidentally losing tickets. So, after one or two such early session key changes, a CMS can assume that a particular MTA has been cloned and flag it in its database.

The limitation of this cloning detection method is that there could be many CMSs on the same IP Telephony network and when each clone is assigned a different CMS, cloning will not be detected. Since in general there are fewer KDCs than there are CMSs, it is easier to catch clones at the KDC. However, if the KDC does not support this functionality, it may still be useful to perform clone detection at the CMS.

Tamper-Resistant Key Storage in the MTA

Although it is possible to detect MTA clones at the KDC or CMS, cloning detection is not a PacketCable^{TM} requirement and may not be available in a particular CMS or KDC brand. Cloning detection would also add cost and complexity to a server. It must be properly implemented and tuned to avoid false alarms and does not guarantee 100% detection. Therefore, other anti-cloning measures should also be considered.

Another way to prevent MTA cloning is to build an MTA where the cryptographic keys are protected inside tamper-resistant hardware. A secure hardware module must not expose protected keys on its external interfaces, which means that the module must internally implement all cryptographic operations associated with the protected keys.

This protected key storage does come at some cost. This cost can be significantly reduced if the cryptographic module does not have to support internal generation of public keys. Public/private key pairs as well as the corresponding digital certificates can be generated and installed into the secure hardware module during the manufacturing process.

There are a number of integrated client platforms where an MTA may be integrated with other functions in the same device that may also have a need for secure hardware. In such cases, the cost of secure hardware may be a lesser factor or maybe not a factor at all if a particular integrated platform is already required to include secure hardware for functions other than the MTA. These integrated platforms are discussed in the following subsections.

Advanced Digital Television Set-Top

Most digital television set-tops available today already include some form of secure hardware. This secure hardware may take the form of a Smart Card, PCMCIA card or an embedded secure co-processor.

The reason for this is that, in the case of broadcast television, set-tops are not required to send any upstream messages and therefore clones are not detectable in the network. Furthermore, premium television programming has a significantly large revenue stream that attracts pirates.

Some of the more recent advanced set-top models support not only digital broadcast television, but also Internet connectivity and email services with an integrated DOCSIS cable modem. In addition, integrated set-tops may also support MTA functionality to provide VoIP services.

For this type of integrated platform, it makes sense to take advantage of the already available secure hardware modules to protect PacketCable^{TM} MTA keys, including the 1024-bit RSA key as well as the Kerberos session keys.

There are some additional short-lived keys that could also be protected inside the secure hardware module, although there is less risk in losing the shorter-lived ones.

Home Gateway

A home gateway platform would be located in a consumer's home and would sit between the HFC network and a subscriber home network. A home gateway may, for example, obtain entertainment content from the Internet and then distribute the content over a home network, subject to protections provided by Digital Rights Management (DRM).

A DRM system generally includes local enforcement of content usage rules. For example, content copying outside of the home network may be prohibited, or the content may be downloaded for only a limited time period after which it must be erased from the home network. Commonly, DRM is enforced by encrypting the stored content and allowing a client to access a decryption key only when content usage rules are satisfied.

Since the evaluation of content usage rules is performed locally inside a home gateway (and inside other home network devices), the home gateway may already contain secure hardware for enforcing Digital Rights Management. An integrated home gateway may also include an MTA and provide VoIP services. In this case, it again makes sense to share the secure hardware element for protection of both the DRM keys and the PacketCable^{TM} MTA keys.

Soft MTA

An MTA can also be implemented in software, running on a PC that is connected to the Internet via a cable modem. This PC may also be running other unrelated software, or may be downloading software from the Internet and would therefore be a potential target for hackers on the Internet.

It is therefore conceivable for the hackers on the Internet to extract MTA keys without the owner's knowledge and then install them into clones. It would therefore be prudent for a soft MTA to store its keys inside a secure hardware module such as a USB token or a Smart Card.

IP ADDRESS CLONING THREATS AND THEIR PREVENTION

Up to this point, the paper described MTA cloning threats in which the subscriber's identity is impersonated, where the subscriber is linked to an MTA FQDN. In order to impersonate a subscriber, an MTA FQDN, an associated set of cryptographic keys, an MTA MAC address and possibly an MTA IP address are copied from a legitimate MTA into a clone. In addition, cloning of only an MTA's IP address can lead to denial of service. Two such threats are explained in the following subsections.

Loss of QoS at an MTA

Media stream (RTP) packets for a voice conversation between two MTAs (or between an MTA and a PSTN Gateway) may only be authenticated end-to-end. PacketCable^{TM} provides an optional MMH MAC that can be added to each RTP packet to verify that it came from a legitimate source and was not modified in transit.

Because the MMH MAC is verified by a VoIP endpoint (MTA), the CMTS cannot distinguish between good and bad downstream RTP packets and will pass them all through to an MTA. At the same time, the CMTS enforces a rate limit for each MTA and will start dropping downstream packets if the allocated bandwidth for a particular MTA is exceeded. The same applies to the upstream packets, although the CMTS limits the upstream packets to a particular MAC address domain associated with a specific CMTS line card.

Also, the PacketCable^{TM} DQoS specification requires the CMTS to pass only those VoIP packets that are associated with a particular VoIP QoS gate, where a gate is associated with a specific source MTA IP address and a specific destination MTA IP address. In order for a CMTS to forward a downstream VoIP packet to an MTA, the source IP address must correspond to a previously allocated gate.

A possible attack would be where an adversary:

1. Determines the IP addresses of the two MTAs (or MTA and PSTN Gateway) that have a current voice conversation.
2. Impersonates the IP address of a PSTN Gateway or of one of the MTAs.
3. Starts sending garbage packets to the other MTA.

In this case, the CMTS would match the garbage packets against one of the gates and pass them through to the MTA. But once a rate limit for that MTA is reached, the CMTS will start dropping packets – both good and bad.

This attack can be addressed by making it difficult for VoIP clients to falsify an IP address. Assuming that the adversary is located on an HFC network, the following prevention steps can be taken:

1. The CMTS verifies that the HFC MAC address and IP address match for each upstream packet. This forces an adversary to have to impersonate both the IP address and MAC address at the same time.

2. The CMTS matches up an MTA MAC address against a Cable Modem MAC address. (Even an embedded MTA is required to have a separate MAC and IP addresses.) This check forces an adversary to impersonate the Cable Modem MAC address as well.

3. The Cable Modem provides physical security for the BPI+ keys. BPI+ provides Cable Modem authentication. By physically securing the Cable Modem keys, it makes the impersonation of the Cable Modem MAC address very difficult.

A simplification of steps 1 and 2 can also be applied in what is known as the “DHCP authority” function. This function has the DHCP server only assign long-term IP addresses to CM and MTA with provisioned CM and MTA MAC addresses respectively. Furthermore, it has the CMTS store IP address to CM MAC address associations based on DHCP requests/acknowledgements. In this case the CMTS acts as a DHCP relay agent for CM and MTA, allowing it to sniff and direct DHCP packets passing through. With the DHCP authority function, upstream packets must match IP and CM MAC address associations or be dropped. This applies to CM IP address to CM MAC address mapping as well as MTA IP address to CM MAC address mapping.

An adversary that is sending bad VoIP packets can also be located somewhere else on the Internet where the upstream packets do not go through a CMTS. In that case, an impersonation of a legitimate MTA’s IP address can be prevented as follows:

1. The operator of the managed IP backbone would have some Edge

Router that connects to the Internet at-large, knowing that MTAs don't connect through that interface.
2. The Edge Router receiving packets from the general Internet would mark the TOS (Type-Of-Service) byte in the IP header to distinguish them from other packets.
3. When a CMTS gets incoming packets with this TOS byte value, it knows they didn't come from an MTA and would therefore not allow any such packets to match any of the gates, regardless of the IP address values.

Loss of IPsec Security Associations

IPsec keys are normally associated with specific IP addresses. A CMS keeps a list of IPsec Security Associations, where each one has a different MTA IP address.

An adversary could:

1. Take a certified PacketCable^{TM} MTA (MTA-A) and spoof an IP address of another legitimate MTA (MTA-B).
2. MTA-A with MTA-B's IP address sends MTA-A's ticket to the CMS to establish new IPsec SAs.
3. The CMS replaces IPsec SAs of MTA-B's IP address with new ones, based on the session key in MTA A's CMS ticket. Since the real MTA-B did not initiate this key management transaction, it will no longer share IPsec keys with the CMS and will temporarily lose service.

This attack can be addressed by having the CMS conduct a DNS query of the IP address corresponding to the MTA FQDN in the AP Request CMS ticket. If the IP address from this interaction differs from the IP address of the AP Request the request is dropped. Unfortunately, this approach may be CMS processing intensive.

A better approach of mitigating this attack is to have the KDC place the MTA's IP address into a ticket. In this case the CMS would not accept a ticket if the IP address inside the ticket doesn't match the address in an AP Request IP header. If a KDC client falsifies its IP address during a ticket request, usually the KDC will not be able to route a ticket back to that client. So, it would be difficult for an adversary to falsify the IP address inside the ticket. This protection could be strengthened further by verifying the IP address to MAC address mapping at the CMTS. Alternatively it could be strengthened by having the KDC determine the MTA IP address via DNS lookup using the MTA's FQDN, based on MTA MAC address and returned by the provisioning server.

LIMITS TO MTA CLONING

MTA subscriber cloning consists of extracting the MTA FQDN, device certificate and private key and copying them into clones. A cloned MTA could have its own MAC address, since in general the KDC looks at the MAC address in the MTA device certificate and does not know if it is the same as the MAC that the CMTS encountered in the MAC frame header. Current PacketCable specifications do not require the KDC to check the MTA IP address, so each MTA clone could also have its own IP address.

Such threats can be mitigated through clone detection and/or with tamper-resistant key storage in an MTA. But what happens if these approaches are not feasible? Can a cloning threat still be mitigated? The answer is “possibly”, if the clones can be restricted to a small portion of a cable plant and forced to operate under operating conditions inconvenient to the pirate.

If MTA clones and their associated CM can be restricted to the same CMTS upstream, then they will be limited to a small cloning population and will be forced to have only one clone operate at a time. In the latter case the MTA's associated CM would also have to be cloned (MAC and BPI key included) so that more than one CM clone would experience conflicting upstream synchronization messages. CM cloning could be forced through use of the "DHCP authority" function as described previously.

Still, how can an MTA/CM clone be forced onto a single CMTS upstream? First, MTA clones cannot be allowed to use their own IP address. As already mentioned in this paper, the KDC can map the MAC address in the MTA device certificate to the IP address that was previously assigned by the DHCP server and then verify that it is the same as the source IP address in the Kerberos AS Request message. Alternatively, the KDC can first map the MAC address to an MTA FQDN and then to an IP address. Either way, all MTA clones would be forced to share the IP address of the original MTA.

The sharing of an IP address requires some out-of-band coordination between MTA clones since they will not be able to make phone calls at the same time without interfering with each other. This already creates inconvenient operating conditions for a pirate.

Once we know that all clones of the same MTA have to share the same IP address, the subnet component of the IP address could be utilized to restrict the geographical location of the MTA clones. This requires that a CMTS, as part of its DHCP relay operation, convey the IP subnet of an upstream interface associated with a CM or MTA in the "giaddr" field of their DHCP discover messages (see [6]). It also requires that the DHCP server that is configured with the MAC address of the CM and MTA sending an offer message (per "DHCP authority" function) have these MAC addresses assigned to the IP address pool corresponding to the subnet of the CMTS upstream channel in which the CM and MTA are located.

Such restrictions would come at the expense of added DHCP and CMTS configuration complexity. It would also come with the restriction that CM and MTA locations be known for the provisioning. However this restriction could be avoided if the CMTS were to record the upstream interface on the first CM or MTA registration, and make sure that this interface does not change without the customer calling a CSR.

SUMMARY

MTA Cloning attacks could potentially result in loss of revenue and disruption of IP Telephony service. In order to fully address cloning, one needs to fully understand the PacketCable^{TM} architecture and in particular the use of multiple MTA identities that include an MTA MAC address, IP address and its FQDN. Different cloning attacks may be based on the duplication of a different MTA identity.

While most MTA cloning attacks are detectable, cloning detection still has to be built into the IP Telephony network and would potentially affect server performance and complexity. Cloning detection has to be carefully implemented so as not to cause false alarms.

In addition to cloning detection and disablement, it is also possible to protect cryptographic keys with the secure key storage inside MTAs. These two anti-cloning measures can be complementary to each other. The use of secure key storage is particularly attractive on integrated client platforms where it is already utilized for other functions such as decryption of broadcast television and Digital Rights Management.

Cloning may also be effectively mitigated by forcing them to operate under a single CMTS upstream channel. This requires DHCP and CMTS configuration and filtering options as well as DHCP exchange measures. The approach comes at the expense of added configuration complexity and location knowledge, but may be attractive when cloning detection and/or secure key storage is not feasible.

REFERENCES

[1] PacketCable Security Specification, *PKT-SP-SEC-I05-020116, January 16, 2002, Cable Television Laboratories, Inc., http://www.PacketCable.com/*

[2] PacketCable 1.0 Architecture Framework Technical Report, *PKT-TR-ARCH-V01-991201, December 1, 1999, Cable Television Laboratories, Inc., http://www.PacketCable.com/*

[3] PacketCable Network-Based Call Signaling Protocol Specification, *PKT-SP-EC-MGCP-I04-011221, December 21, 2001, Cable Television Laboratories, Inc., http://www.PacketCable.com/*

[4] PacketCable MTA Device Provisioning Specification, *PKT-SP-PROV-I03-011221, December 21, 2001, Cable Television Laboratories, Inc., http://www.PacketCable.com/*

[5] PacketCable Audio/Video Codecs Specification, *PKT-SP-CODEC-I03-011221, December 21, 2001, Cable Television Laboratories, Inc., http://www.PacketCable.com/*

[6] Dynamic Host Configuration Protocol, *IETF (R. Droms), Internet Informational Standard, RFC 2131, March 1997.*

[7] Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification, *SP-RFIv1.1-I07-010829, August 29, 2001, Cable Television Laboratories, Inc. http://www.CableLabs.com/*

ACKNOWLEDGEMENTS

This paper is in part based on the work done by the PacketCable^{TM} security team. Several of the cloning threats discussed in this paper were uncovered during discussions within the PacketCable^{TM} security team.

CONTACT INFORMATION

Alexander Medvinsky
Motorola Broadband
6450 Sequence Dr.
San Diego, CA 92121
Tel: (858) 404-2367
smedvinsky@gi.com

Jay Strater
Motorola Broadband
101 Tournament Dr
Horsham, PA 19044
Tel: (215) 323-1362
jstrater@gi.com

CONSIDERATIONS FOR DIGITAL PROGRAM INSERTION OF MULTIPLE-VIDEO PROGRAMS

Aldo G. Cugnini
SpotOn^{SM}, an ACTV Company

Abstract

Applications are now being deployed that give multichannel video programming distributors the ability to deliver interactive and/or addressable (targeted) advertisements to homes equipped with digital set-top boxes. These applications, in concert with appropriately encoded bitstreams, provide the viewer of an enhanced program with an interactive experience. The initial use of this technology will be to bring the user enhanced advertisements, and to provide the advertiser with feedback on how viewers interacted with their advertisement. In order to enable this functionality, the enhanced advertisements must be inserted or "spliced" into the network programming in a seamless fashion at the headend. This can be accomplished by implementing the system described here.

INTRODUCTION

Digital Program Insertion (DPI) of Targeted Advertisements provides one means for accelerating the transition to digital cable by decreasing the complexity of local ad insertion equipment, and increasing revenue by providing additional spot sales opportunities. Targeted Ad Insertion can help subsidize the transition to fully Digital Cable, and can yield more bandwidth and MSO revenue by enabling a path to ultimately reclaim analog bandwidth. This technology will also further motivate the deployment of digital set top boxes as MSOs increasingly utilize non-customer based revenue.

An end-to-end system for inserting interactive ads into a network feed obtained from a satellite downlink will be described. Two different scenarios will be discussed, involving analog and digital source material. One scenario represents the situation where the downlink network is analog, or the feed must be decoded down to the analog level, e.g., to extract the analog cue-tone information needed by an ad server. The second scenario describes a system where the downlink signal format is digital, and DVS/253 cue information has been inserted into the stream. In this latter case, the video does not need to be decoded and re-encoded, a situation which greatly reduces complexity and cost.

In order to provide for seamless splicing, certain requirements must be satisfied in an MPEG bitstream. In addition to the proper handling of video frames and frame types, the set-top box video decoder buffer must be managed appropriately at the splice points. Several schemes for meeting this requirement will be discussed.

OVERVIEW OF DIGITAL PROGRAM INSERTION

Techniques for conventional single-program DPI are well known and will be summarized here. A typical DPI system is shown below in *Figure 1*. Because legacy equipment must be considered, a hybrid system is depicted that supports insertion into an analog network. A Traffic and Billing (T&B) System maintains the overall schedule of ads to be inserted. The ad server and ad splicer exchange timing and asset information concerning the scheduling of an ad insertion.

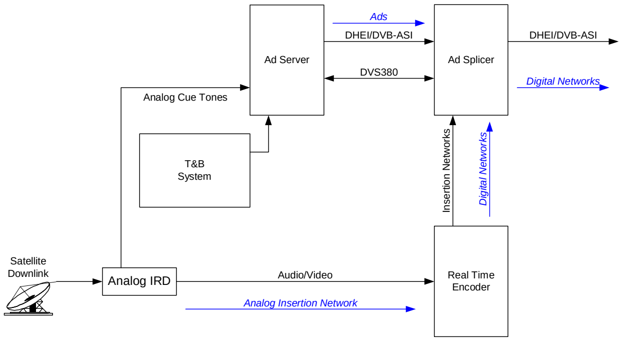

Figure 1: Digital Program Insertion.

In this scheme, analog cue tones are received from the insertion network and passed on to the ad server. Using DVS-380 protocol, the ad server and ad splicer establish a dialog to schedule and carry out a DPI event. Pre-compressed ads are delivered to the ad splicer at the appropriate time for insertion into the Insertion Network. A DVB/ASI or DHEI interface is then used to deliver the bitstream to the cable plant.

Several different technologies must be employed to realize this system, including bitstream multiplexing, bit-rate management, and cue message detection, as well as system intercommunication and event timing.

INSERTION OF MULTIPLE VIDEO PROGRAMS

In *Figure 2*, we see a new element, the SpotOn Controller, which is responsible for adding the necessary information to support Targeted Ad Insertion (TAI) in the set-top box (STB). Although the controller is shown as a separate element, its functionality could be integrated into the T&B System or the ad server; in some cases it could also be remotely located. Its purpose is to create a new data service, which will instruct the STB which of the various alternate programs (ads) should be viewed. These instructions are coded using a proprietary syntax called ACTV Command Language (ACL). This configuration can also be used to provide national distribution of targeted ads.

The STB program selection is enabled through the use of a small client resident in the STB, which decodes and acts upon the ACL commands. Using a user profile stored in the STB, the client switches the decoded video based on the ACL commands it receives.

There are two distinct “splices” which are necessary to implement a Targeted Ad Insertion: the headend ad insertion, where a multiplex of several ad programs is inserted into the program stream; and the STB program switch, where the STB switches between the different ads in the multiplex. Each of these splices must be done seamlessly, so that the viewer is unaware (if so desired) of the splice.

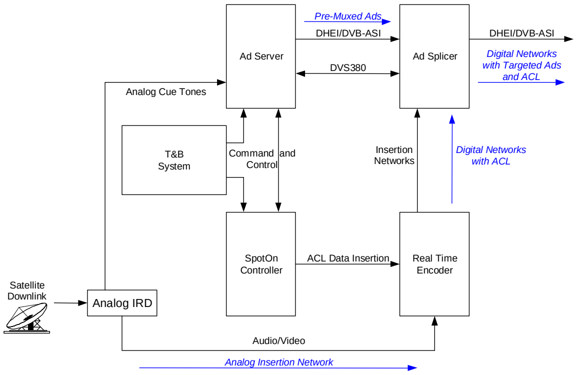

Figure 2: Targeted Ad DPI over Analog Insertion Network.

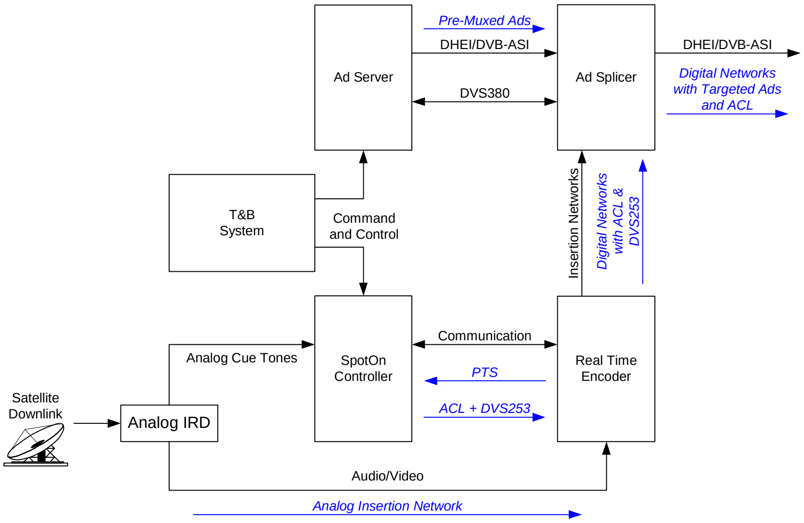

Figure 3. Insertion using DVS-253 Compliant Splicer.

The first-generation TAI system eases the real-time burden on the ad splicer by pre-multiplexing the “bundle” of ads before storing them on the ad server. Development of a real-time ad multiplex system will allow last-minute assembly of the ad multiplex.

INSERTION USING DVS-253

As ad splicers with DVS-253 capability become available, the system can evolve to support such devices, as shown in *Figure 3*. Since the insertion network is still analog here, there is no DVS-253 network messaging present. However, all the necessary information is available locally to synthesize these messages. The SpotOn controller can receive scheduling information and the analog cue tones upon which to base an insertion. After retrieving PTS information from the encoder, it can then assemble the appropriate DVS-253 message to hand off to the encoder, which then inserts the message into the stream.

DIGITAL INSERTION NETWORK

Ultimately, the evolution towards an almost fully-digital plant will enable the realization of the simple architecture shown in *Figure 4*. With the exception of locally produced analog programming, all video can ultimately be handled in compressed digital form. An IRT interfaces digital programming with embedded DVS-253 cue messaging directly to the ad splicer.

SPLICING CONSIDERATIONS

Bitstream Integrity

In order to provide for seamless splicing, certain constraints must be met in an MPEG bitstream, to assure that the STB video decoder is presented with an MPEG-compliant bitstream. The subject of proper switching of video frames and frame types is well known and will not be discussed at length here. For the purposes of our discussion, it is sufficient to assume that the streams must be “Closed GOP” and all splices must occur at GOP boundaries, so that the bitstream transition is from one intact “old” sequence to an intact “new” one.

In general, these conditions must be met at the input to the STB; hence, the input streams to the ad splicer can be

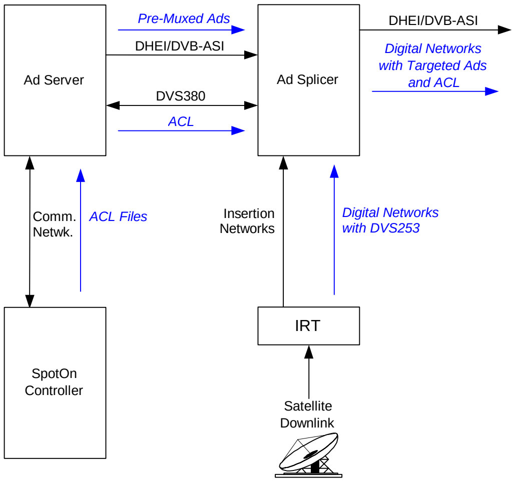

Figure 4: Fully-digital Program Insertion.

unconstrained, given a sophisticated splicer that can create these conditions at its output. However, certain new conditions must be present in the stored ad streams in order to facilitate splicing.

Multiplexing

Targeted Advertising is accomplished in existing digital STBs by commanding the Transport Demultiplexer to switch to a different video program (PID); see *Figure 6*. When the appropriate ACL command is received, the ACL Client instructs the Demultiplexer to switch to a different PID. The action of the client ensures that this switch is performed only at a GOP boundary.

In order to produce a seamless switch, the bitstream at the ad splicer must be assembled in such a way that any switch from one video (ad) program to any other results in an MPEG-compliant stream. (This includes a PID switch from the network video to one of the ad videos.) This requires all contemporaneous video sequences in the multiplex to start after and end before the transmission of any video from a previous or subsequent sequence, respectively. This "gap" is shown below in *Figure 5*.

This requirement must be met at any point deemed a "splice opportunity" at the STB. Note that, since a typical ad insertion will span several sequences, only the first and last sequences must meet this requirement. Thus, the required gap may be created at the time the video is encoded, and an entire multiplexed "package" can be stored on the ad server, greatly reducing the processing requirements of the ad splicer. Since this can be done in advance of airtime, and even in non-real-time, the process can be performed offline using a post-processing (software) algorithm after the video is encoded.

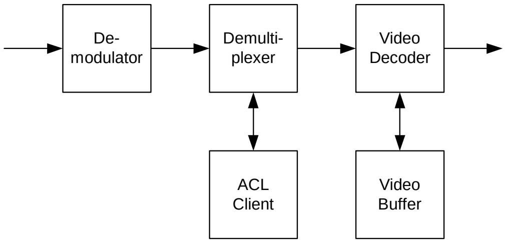

Figure 6: STB Decoder.

Alternatively, the individual ad videos can be encoded and stored on the ad server, so that a last-minute multiplexing can be performed. Of course, this requires more real-time processing in the ad splicer at airtime, but the added complexity may be desirable in order to support the added flexibility.

Bit-rate management

The video decoder buffer must be managed carefully at the splice points in order not to create an overflow or underflow condition. An MPEG video encoder ensures this by maintaining a predictive model called the Video Buffering Verifier (VBV). This model calculates how the decoder buffer will behave when presented with the MPEG bitstream. Because the action of the buffer can be completely predicted based on certain

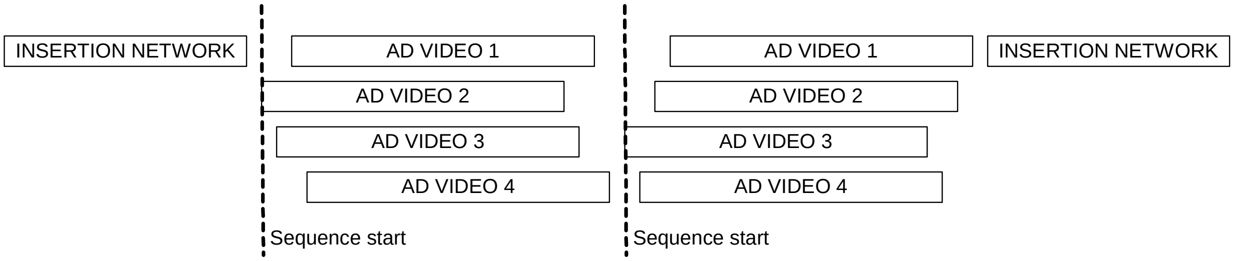

Figure 5: Bitstream Multiplex.

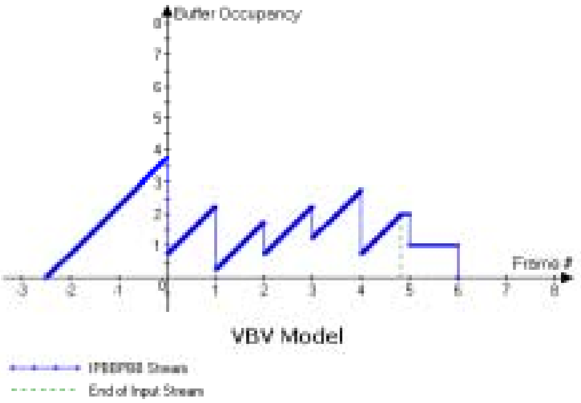

Figure 7. Video Buffer Trajectory.

bitstream parameters, and MPEG bitstream compliance requires that the VBV never overflow or underflow, proper operation of the decoder can be achieved by keeping an accurate VBV model in the encoder. See *Figure 7* above.

When presented with a Constant Bitrate (CBR) stream, the video buffer fills at a constant rate (linear slope of diagonal lines), and is emptied instantaneously by an amount equal to the size of each picture. The initial time that the buffer fills before the first picture is removed is the *vbv_delay* of that picture. It is important to note that if this sequence were followed immediately by another sequence, then the first bit of the new sequence would enter the buffer after the end of the first sequence. This would occur in the example at about time *t* = 4.8. In effect, this point becomes an available *splice point* in the old stream.

The result of blindly appending a new sequence is apparent in *Figure 8*. As an example, if we append the original sequence *to itself*, we see an undesired effect. Focusing on the region around the splice point, we can see that the buffer continues to fill at the video rate (as the size of each frame is unchanged from the previous example). However, the first frame of the new sequence must be removed at time *t* = 7. This requirement causes the buffer to fill for a time *less* than that originally specified for the first frame of the new sequence. In this example, we can see that the buffer fills for roughly 2.2 frames, whereas the original sequence called for a *vbv_delay* of 2.5 frames.

The consequence is that, upon the removal of the frame at time *t* = 8, the buffer underflows, i.e., not enough data has entered the buffer to ensure it is ready to be removed at the next access time. Since the new stream was encoded with the expectation of a specific VBV trajectory, this condition must not be violated at any downstream point. The *vbv_delay*, the bitrate, and the size of the new frames can only be modified if the resultant stream maintains VBV compliance.

It is important to note that the error will still exist at the end of the ad stream—it does not “flush” after one GOP, and we should not expect it to do so at any time in the future. The only way to correct the error is to re-establish the correct *vbv_delay* of a subsequent sequence.

One solution is to present the new sequence to the buffer at a time in advance of the decode time equal to the amount

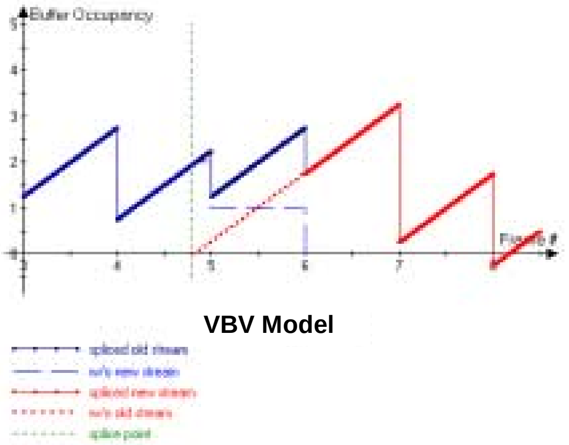

Figure 8. Buffer Occupancy at Splice Point (Bad Splice)

specified in the *vbv_delay* for the first frame of the new sequence. In other words, the last bit of the outgoing stream should remain in the buffer for a time equal to the *vbv_delay* parameter of the first frame of the new stream, minus the display time of the last frame of the outgoing stream.

One way to do this is to reduce the size of (i.e., re-code) the last picture(s) of the old stream. This is shown in *Figure 9* below. The size of the frame at time *t* = 4 was reduced from 2 to 1.5. This causes the old stream to end sooner, and allows the new stream to enter the buffer at the appropriate time, *t* = 4.5.

The specific solution depends upon the conditions at the splice point. If the outgoing stream terminates *before* the new stream should start, then null padding can be used to extend the life of the old stream in the buffer. If the outgoing stream would otherwise terminate *after* the new stream should start, then the last few frames of the old stream can be re-coded with fewer bits.

In this particular solution, all of the re-coding operations are accomplished in the ad splicer at airtime. Although the *vbv_delay* of the ad stream can be known in advance, the same parameter cannot be known of the network stream until it enters the ad splicer.

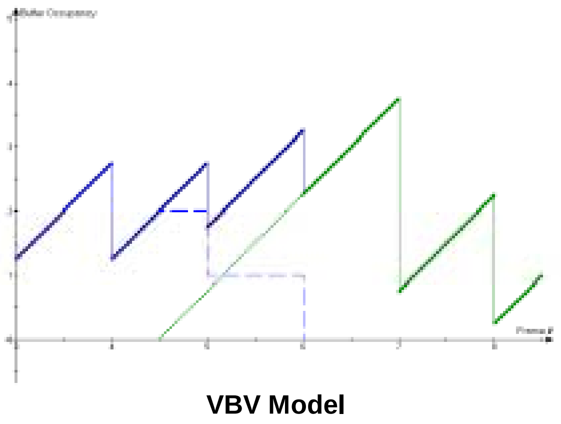

Figure 9. Buffer Occupancy at Splice Point (Good Splice).

However, there are other solutions which distribute the complexity elsewhere. One such method is to force the *vbv_delay* of the first frame of the incoming network stream to a known specific value, so that the splicer can set up the conditions for a proper splice in advance of the new stream being available. This is the rationale for the *splice_decoding_delay* MPEG parameter, which is the specific value of *vbv_delay* needed to perform seamless splicing without post-modifying the bitstreams. This parameter essentially defines the point at which the old stream has finished entering the video buffer.

For an MP@ML stream at video rates up to 7.2 Mb/s, SMPTE 312M specifies a *splice_decoding_delay* of 250ms. Although this solution removes the constraints on the splicer, it places a large one on the bitstreams—one that can compromise the quality of the video. For this reason, it is not a preferred solution.

Another method is to use a large value of *vbv_delay* for the ad stream, coupled with relatively small picture sizes. This is shown below in *Figure 10*. Because the buffer is constrained in its excursion, any splicing errors will be less likely to result in over- or underflow. One advantage here is that this places no burden on the splicer; however, there is a large price paid in picture quality due to the limited picture sizes.

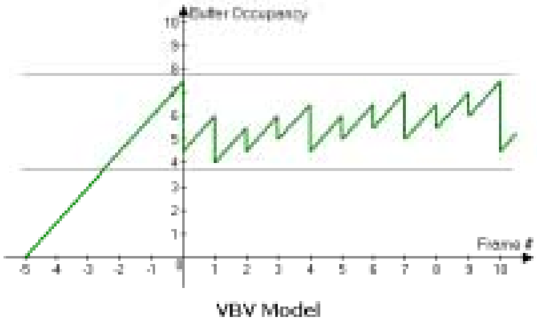

Figure 10. Effect of Large vbv_delay.

Due to the fact that the ad streams are pre-encoded in advance, other splicing solutions can be employed that maximize video quality, including dynamic combinations of the above and other proprietary techniques.

MULTIPLE VIDEO SPLICING

So far, we have focused our bitrate management discussion on that of single-video stream splicing into another single-video stream. We will now expand the discussion to that of a single-video program (or transport) stream splicing to and from a multiple-video stream. In order to properly perform these splices, we modify our recoding solution as follows. (Note that there are other solutions, as well.)

Single-video stream to multiple-video stream

1. The *vbv_delay* of each of the first pictures of the new stream is calculated. The longest *vbv_delay* sets the point at which the outgoing stream must terminate.
2. The end of the outgoing stream is shortened as needed by recoding the data into fewer bits.

Multiple-video stream to single-video stream

1. The *vbv_delay* of the first picture of the new stream is calculated. The *vbv_delay* sets the point at which the latest outgoing video component must terminate.
2. The end of each video component of the outgoing stream is shortened as needed by recoding the data into fewer bits.

Multiple-video stream to multiple-video stream

1. The *vbv_delay* of each of the first pictures of the new stream is calculated. The longest *vbv_delay* sets the point at which the outgoing stream must terminate.
2. The end of each video component of the outgoing stream is shortened as needed by recoding the data into fewer bits.

SUMMARY

An overview of digital program insertion of multiple video programs has been presented. Various solutions were described considering both legacy and emerging architectures, such as hybrid analog-digital systems and DVS-253 cue messaging. Bitstream considerations have been analyzed, including bitstream integrity, multiplexing, and bit-rate management. Splicing of multiple video programs can be accomplished by logical extension of single-program splicing techniques.

ACKNOWLEDGEMENTS

The author would like to thank his colleagues Mike Cristofalo, Frank Deo and Art Greenberg. for their useful discussions and contributions regarding this paper.

NOTICE

The reader’s attention is called to the possibility that commercialization of the technologies described herein may require the use of inventions protected by patent and intellectual property rights.

REFERENCES

1) ISO/IEC 13818-1:1996(E), "Information technology - Generic coding of moving pictures and associated audio information: Systems," *MPEG-2*, April 4, 1996.

2) SMPTE 312M-1999, "Television—Splice Points for MPEG-2 Transport Streams," *SMPTE*, 1999.

3) SCTE 30 2001 (Formerly DVS 380), "Digital Program Insertion Splicing API," *SCTE*, 2001.

4) SCTE 35 2001 (formerly DVS 253), "Digital Program Insertion Cueing Message for Cable," *SCTE*, 2001.

5) Mukta Kar, Sam Narasimhan and Richard S. Prodan, "Local Commercial Insertion in the Digital Headend," *NCTA Technical Paper*, 2000.

6) D.T. Hoang, E. Linzer, J.S. Vitter, "Lexicographically Optimal Rate Control for Video Coding with MPEG Buffer Constraints," *Dept. of Computer Science, Duke University*, February 1996.

CUSTOMIZED BROADBAND – ANALYSIS TECHNIQUES FOR BLENDED MULTIPLEXES

Dr. Robert L. Howald
Erik Metz
Rob Thompson
Motorola Broadband Communications Sector

Abstract

With HFC reaching saturation levels in North America, much of the focus of the vendor community for infrastructure development is in the international arena. As a result, there is a wide variation in frequency plans, bandwidths implemented, transport methods, and cascade requirements. This variation occurs simply because of global differences in demographics, topology, and standards used. As a basic example, it is now quite common to require system analysis and design for deployments featuring a mixed set of analog and digital signals, where digital channels of multiple bandwidth and modulation formats exist between groupings of analog channels, within different segments of the forward band. This situation means that a mechanism for analyzing, characterizing, and understanding the effects on performance of these systems is necessary to optimally develop system solutions in these cases.

Because of these needs, a user–friendly tool has been developed that predicts performance based on manual inputs on all of the key system variables – analog and digital frequency multiplexes, relative levels, slopes, output levels, and distortion variables. Perhaps the simulator's most important task is its ability to calculate the distortion performance of any set of channels – from all analog to all digital – at any relative level of each, for any slope and output level, and for any particular distortion baseline. The simulator output can be delivered in numerical tables of every analog and digital distortion component, or plotted. Output plots can be broken down into each individual contributing analog and digital distortion component, by the order of distortion component. They can be combined into a single composite plot that includes both the analog and digital components individually by order. Or, the output can be combined into one composite map. These outputs are critically important to understanding performance characteristics, and ultimately in designing and recommending hardware for such systems.

INTRODUCTION

The purpose of this paper is to present a summary of a distortion modeling tool, its capabilities, and applications. The primary function of this simulation tool is to graphically and numerically report the second and third-order distortion at any frequency, for any analog and digital multiplex. Throughout this paper, graphical examples are discussed, and some general relationships regarding the behavior of distortion are given.

TRADITIONAL LINEUP

A traditional, forward-path, NTSC channel lineup occupies the band from 55.25 MHz to an upper limit of 865.25 MHz. The lower 500 MHz of is usually reserved for 6 MHz wide, analog channels, while the upper 320 MHz is reserved for 6 MHz wide, digital channels. The digital channels are typically operated anywhere from 6-10 dB below their analog counterparts. These lower levels for digital channels are sometimes referred to as digital-derate levels. The entire forward path is tilted, such that the lowest analog channel is 12 dB lower than it's analog equivalent at 865.25 MHz, 6 dB if that high channel is actually digital, with a 6 dB digital-derate. This is a "typical", North American system whose performance can easily be predicted with measured specifications and some generalized rules. General relationships are what allow us to predict the distortion effects of changing RF levels and cascading RF amplifiers.

General Distortion Relationships

Typically, we can expect third-order distortion, CTB, to degrade by about 2 dB for every 1 dB increase in channel level. We can also expect second-order distortion, CSO, to degrade by about 1 dB for every 1 dB increase in channel level. Cascading RF amplifiers requires that we add CTB performance on a 20-LOG scale, though many times it measures less. Cascaded CSO performance adds on a 20-LOG scale, though a 15-LOG scale is more typical. Additionally, changes in distortion level due to changes in channel loading can be approximated. However, calculating changes in the distortion levels requires that the channel frequencies, which see the most number of beats, are known. Generalizations like these create a simple system for predicting CTB and CSO performance and they are all based upon specifications obtained from laboratory measurements and simple analog beat mapping tools.

Classical Analog Beat Mapping Tools

Without the aid of beat mapping algorithms, generalizations regarding distortion behavior would be difficult to make. Beat maps generate the number of beats (frequency tones) that fall at any frequency as well as the distortion magnitude at that frequency. Simons [1] made this easy to determine when he showed how to calculate distortion magnitude and location for second/third-order distortion. The proof is nothing more than an algebraic manipulation of the time-domain representation of the channel spectrum.

Consider channels A and B. Mixing these channels produces second-order harmonic distortion at frequencies 2A and 2B, as well as sum and difference distortion at frequencies A±B. The harmonic distortion is about 6 dB below the sum and difference distortion. CSO is a cumulative sum of all the second-order harmonic and sum and difference distortion that fall at any frequency.

Now consider channels A, B and C. Mixing produces third-order harmonic distortion at frequencies 3A, 3B and 3C, triple-beat distortion at frequencies A±B±C and intermodulation distortion at frequencies 2A±B, 2A±C, 2B±A, 2B±C, 2C±A and 2C±B. Intermodulation distortion is about 6 dB below triple-beat distortion. Third-order harmonic distortion is about 15.5 dB below triple-beat distortion. CTB is the sum of all third-order harmonic, intermodulation and triple-beat distortion that fall at any frequency.

An analog beat mapping tool keeps track of all these distortions by storing them in a table of increasing frequency. This can be easily implemented and plotted in a spreadsheet program such as Microsoft Excel. Figure 1 shows an example plot of all output distortion components throughout frequency. Different colors identify the distortion contributor. Figure 2 is the cumulative sum of the CTB and CSO components.

Both Figure 1 and Figure 2 are relative plots, which means the levels shown are not absolute levels. The levels shown use second-order sum and difference beats, A±B, and third-order triple-beats, A±B±C as a starting point, with all other types of distortions either about 6 dB or 15.5 dB below in order to get the total relative plots you see in Figure 1 and Figure 2

Figure 2 also shows worst case CTB occurring near the middle of the band at 397.25 MHz. Worst case, in band CSO is at the low frequency edge, at 54MHz.

The worst case distortion occurs at frequencies with the greatest accumulation of beats and power. Analog beat mapping tools provide insight on where to expect worst case distortion as well as distortion levels throughout the band.

Real Hardware Considerations

It is valuable to know theoretical analog distortion performance. However, it does not compare to the accuracy of testing in the lab and observing the output on a spectrum analyzer. However, predicting distortion with real hardware has drawbacks. Lab testing takes time, particularly given multiple channel scenarios with different levels and cascades. Therefore, despite the fact that using real hardware is the most accurate characterization of performance, it may not be the most efficient route, especially when considering many scenarios and trying to quickly respond to customer inquiries.

MODELING DIGITAL MIXING

Accurate modeling is an effective and efficient way to get answers regarding performance. Analog modeling is well understood today, however, that’s not enough. A way to incorporate digital channels into the mix as well as making relative distortion models predict absolute distortion levels versus frequency would enhance the modeling required for today’s applications.

Modifications to Traditional Tools

Predicting and measuring analog distortion is different from digital distortion primarily in the effect it has on other channels. Digital distortion is noise-like. However, the same rules used to calculate analog distortion still apply to digital distortions. The trick is how to describe the digital channel. A simulator with a frequency domain engine could easily model digital channels as a series of discrete tones. A simulator with a time domain engine could represent digital signals as a sinc or raised-cosine function.

After calculating all the distortions, including mixing between analog and digital signals, the three categories must be distinguished – analog/analog, digital/digital, and analog/digital. The reason will become clear once intermodulation noise is defined.

Composite Intermodulation Noise, CIN

Composite intermodulation noise is the noise-like digital distortion that is generated from mixing analog/digital and digital/digital signals together [2,3]. CIN is the combination of CIN2 and CIN3. CIN2 is the second-order digital distortion and CIN3 is the third-order digital distortion. Carrier-to-Noise ratio, CNR, is thermal noise associated with a specific bandwidth and the noise figure of the RF amplifier. All of these noise ratios combine together into composite Carrier-to-Noise ratio, CCN.

Low Power, Low Bandwidth Digital Loads

You can manage CIN by controlling the amount of power and bandwidth associated with your digital channels. Generally, the lower the power and bandwidth, the lower your CIN will be. This may seem intuitive based on what we know about analog distortion, but may be better understood by considering the total power load of the digitally loaded portion of the forward spectrum.

Total digital power load is the sum of the power of all the individual digital channels. Digital distortion varies as a function of the total power load. Therefore, digital bandwidth and derate can be exploited to reduce total power load, which reduces digital distortion and improves CIN. This will be shown in following two cases.

550 MHz Analog Plus 100 MHz Digital

The analog distortion for this case is shown in Figure 1 and Figure 2. The digital distortion, for about 100 MHz of digital loading, is illustrated in Figure 3 and Figure 4 Figure 3 has the digital channels 10 dB below the highest analog carrier and Figure 4 has the digital channels 6 dB below the highest analog carrier. The total digital power load is roughly 4 dB lower for 10 dB digital-derate than it is for the 6 dB digital-derate. The second/third-order distortion is about 4 dB worse in the 6 dB derate system. Therefore, increasing the digital signal level increases the distortion, which is the same as degrading CIN. The one-to-one correspondence indicates that the dominant CIN3 contribution is from 2A+1D

550 MHz Analog Plus 320 MHz Digital

Adding digital bandwidth shifts the maximum third-order distortion. For 320 MHz of digital bandwidth, the maximum CTB is at 706.5 MHz, compared to the maximum at 544.5 MHz for 100 MHz of digital bandwidth. CSO distortion is at its maximum at 315.25 MHz for 320 MHz of digital loading, while 100 MHz digital loading reaches maximum at 99.25 MHz. Therefore, expect shifts in the maximum and minimum locations of CTB and CSO with changes in the bandwidth.

Increasing the digital bandwidth to about 320 MHz will increase the total digital power load by about 7 dB. This will increase the third-order distortion by about 7 dB and increase the second-order distortion by about 6 dB, as illustrated in Figure 5

In Figure 6, with the digital-derate level changed from 6 dB to 10 dB, the performance is 4 dB better in both CTB and CSO.

These examples show that by increasing the total power load, either through increasing channel level or bandwidth, distortion increases, ultimately degrading CIN.

Models and Measurements

Up to this point, it has been shown that relative changes to total power load results in relative changes in the distortion. Measured performance can be tied to the distortion model through nonlinear gain coefficients [1]. Nonlinear gain coefficients effectively scale the distortion levels from a relativistic to an absolute value. For example, assume an RF amplifier has a measured performance of –60 dBc CTB and –66 dBc CSO. The input levels to the amplifier are 17 dBmV per channel. For 550 MHz worth of analog channels, assume that the gain of the device is 27 dB and there is no tilt, so the output power of the carriers is 44 dBmV. Now, from those values of CTB and CSO, calculation of the intercept points yields [5,6], IP_2=128.45 dBmV and IP_3=92.86 dBmV. The intercept points enable us to calculate k_2 and k_3, which are the second/third-order nonlinear gain coefficients. For this case, k_2=1.34E-4 and k_3=3.14E-6. Approximating the absolute distortion for CTB and CSO requires scaling each second/third-order distortion beat by k_2 and k_3, respectively. This results in the beat map shown in Figure 7. The maximum CTB is about –16 dBmV (–60 dBc), the maximum CSO is about –22 dBmV (–66 dBc), with the carrier outputs at 44 dBmV.

CUSTOM MULTIPLEXES

Summarizing, modeling allows prediction of the distortion spectrum for any amplifier of specified performance. However, the real strength of this modeling capability lies mainly in its ability to predict distortion for atypical situations. Custom multiplexes are common in the international arena, offering a wide variety frequency plans, implementation of bandwidths, transport methods, and cascade requirements. The tools described will allow prediction of distortion for any channel configuration. Prediction becomes much more efficient with the aid of a programming language like Visual C++ or Visual Basic, or using MATLAB.

There are some general rules that are valuable to keep in mind. CIN behaves more like a noise floor even though it's generated the same way analog distortion is. Therefore, expect to see a nearly flat spectrum, at least in third order distortion. Also, expect CIN to increase with any increase in total power loading. It may not be a one-for-one increase for all cases, but should increase either way. Predicting performance for custom multiplexes is no different than for the traditional case, just simply a matter of managing the input load and separating the analog and digital components as necessary.

Using programming as described above, simulation time is about 1 minute for every MHz of digital signal. MATLAB provides an output in as little as 4 minutes, regardless of the amount of digital bandwidth. This is because the FFT feature in MATLAB doesn't care what the input spectrum is. It does an FFT based on the number of data-points of the input vector. The only trick with MATLAB simulation is being sure to obey Nyquist.

Using numerical tools, many different scenarios of channel plans can be determined in a reasonable amount of time. This makes for quick turn-around answers with confidence.

Figures 8 and 9 are plots of the relative analog distortion generated for two such cases. Figures 10 and 11 are their respective cumulative distortions. These two cases have digital intermixed with analog throughout the forward band, with digital channel bandwidths of 8 MHz. As indicated, both relative and absolute results can be obtained.

Frequency Dependent Effects

Despite these conclusions, there is one more wild card at work. With today's RF electronics that are often optimized for performance at particular frequencies, the distortion effects across frequency are not constant. The result is that nonlinear gain coefficients will change over frequency. Therefore, the model must account for how k_2 and k_3 will change over frequency. Figure 12 shows how two tones, of equal power and separated by 2, 20 and 100 MHz, resulted in varying distortion performance depending upon where they were located in frequency.

CONCLUSION

Today's link analysis requires very flexible modeling tools. Both relative and absolute second/third-order distortion for any analog and digital multiplex can be modeled. With efficient numerical tools, many variations of an analog and digital multiplex can be predicted. The capability to analyze, characterize and understand the effects of these systems allows development of optimal system solutions.

REFERENCES

[1] Simons, K. (1968). “Technical Handbook for CATV Systems”. Philadelphia: Jerrold Electronics Corporation.

[2] Hamilton, J., & Stoneback, D. (1994). “How digital carriers affect analog plant”, Communications Engineering and Design, 80-93

[3] Waltrich, J. B. (1993). “Distortion produced by digital transmission in a mixed analog and digital system”, Communications Technology, 40, 64-76

[4] Matrix Test Equipment. (1998) “Notes on composite second and third order intermodulation distortions”. Available http://www.matrixtest.com//Literat/mtn108.htm

[5] Matrix Test Equipment. (1998) “The relationships of intercept points and distortion”. Available http://www.matrixtest.com/Literat/MTN109.htm

FIGURE 1. Analog Individual Components

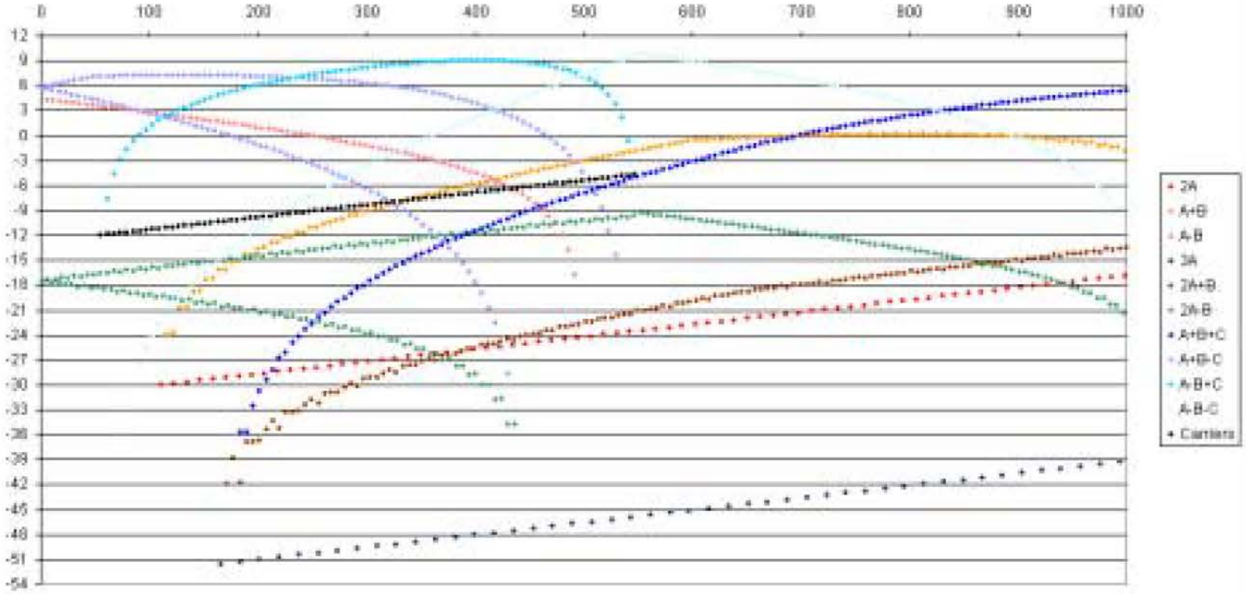

FIGURE 2. Analog Composite

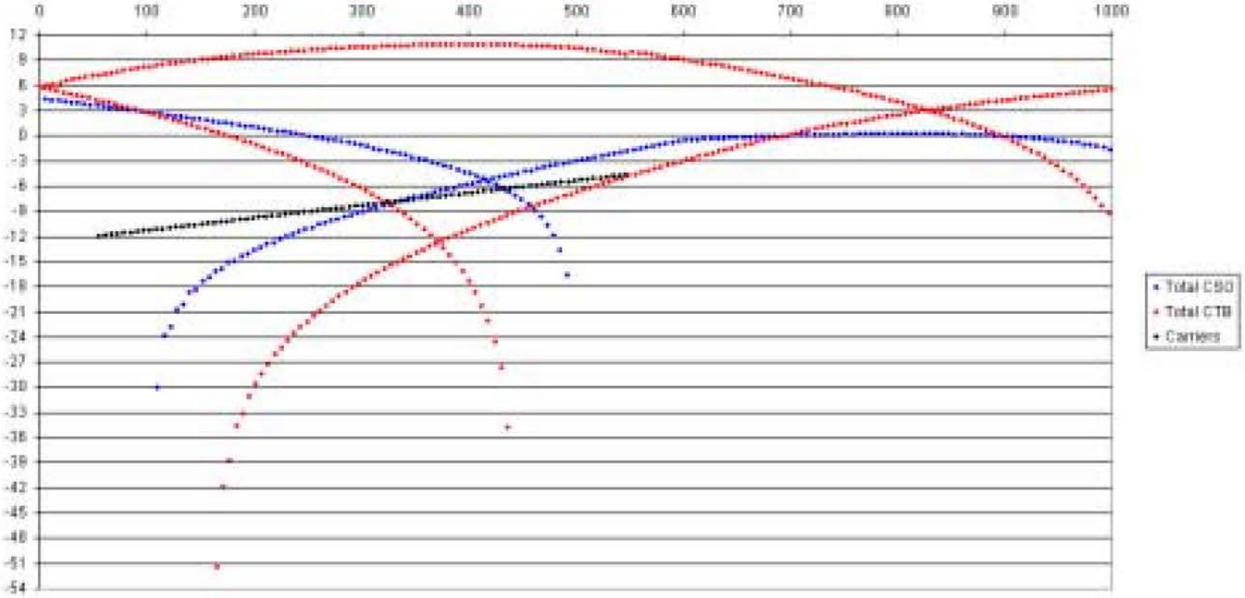

FIGURE 3. Digital Composite

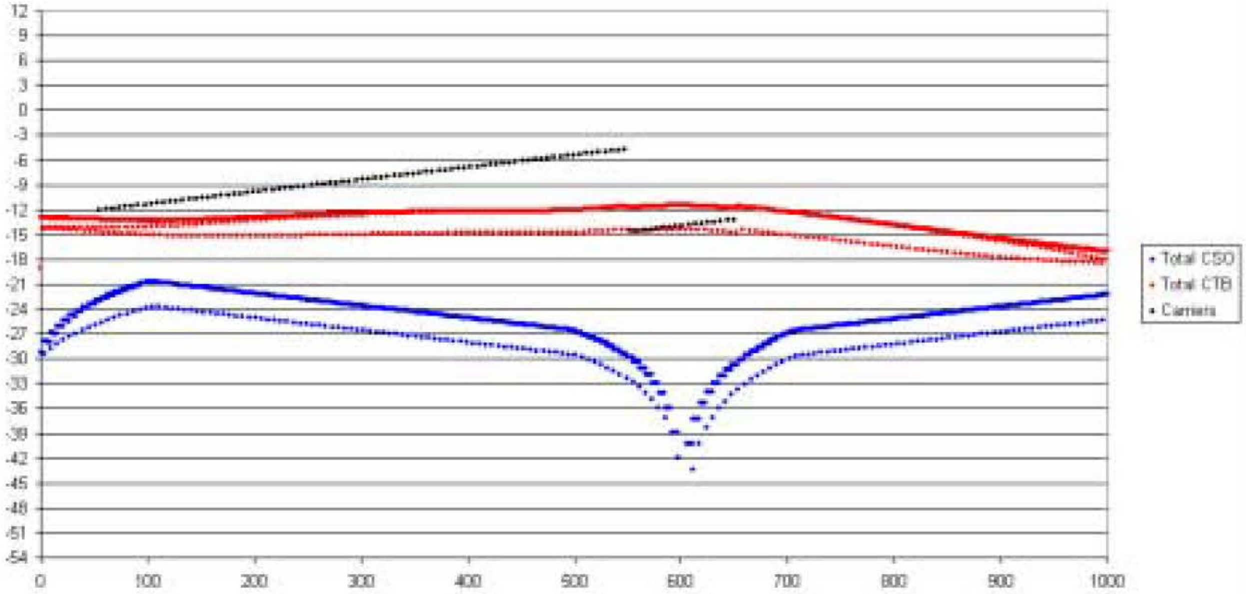

FIGURE 4. Digital Composite

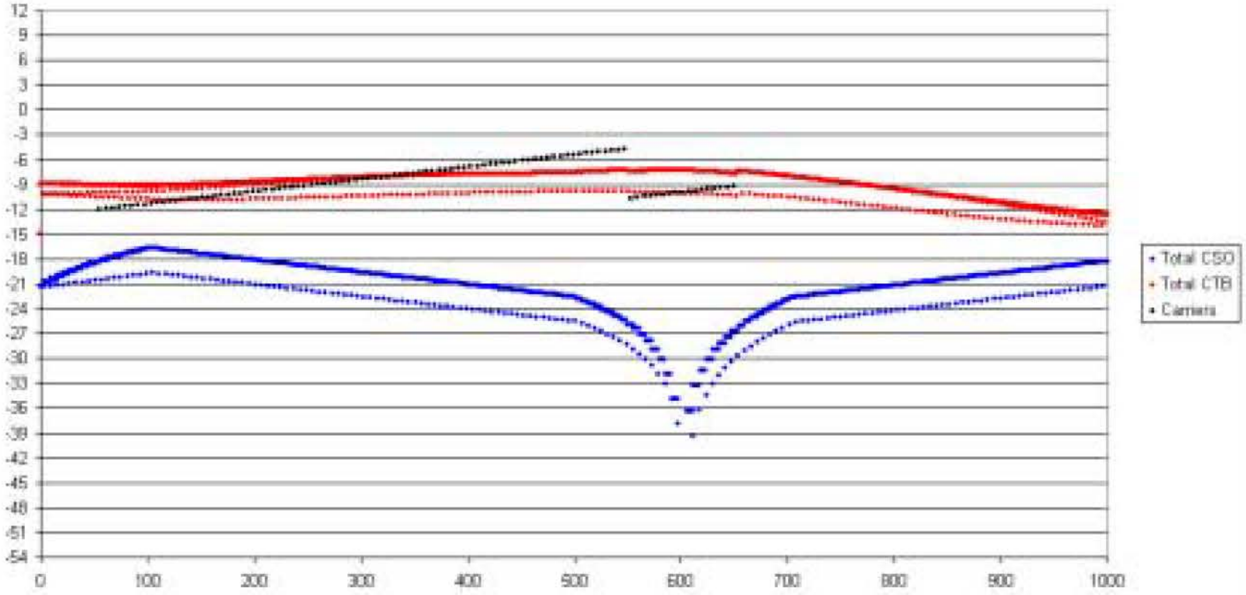

FIGURE 5. Digital Composite

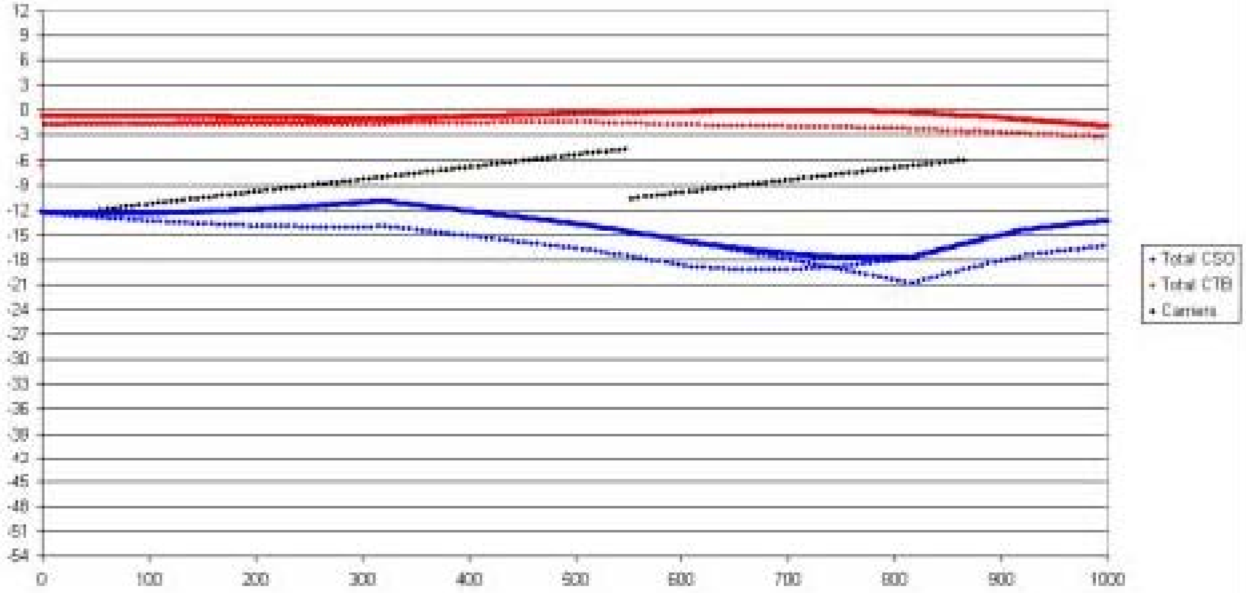

FIGURE 6. Digital Composite

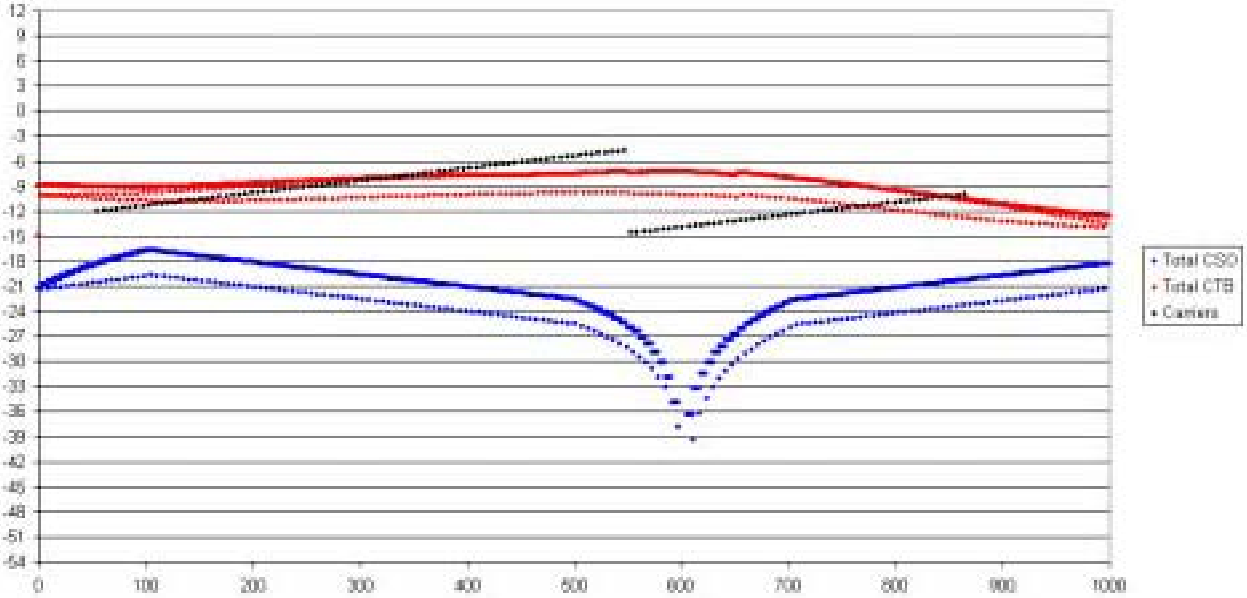

FIGURE 7. Analog Composite

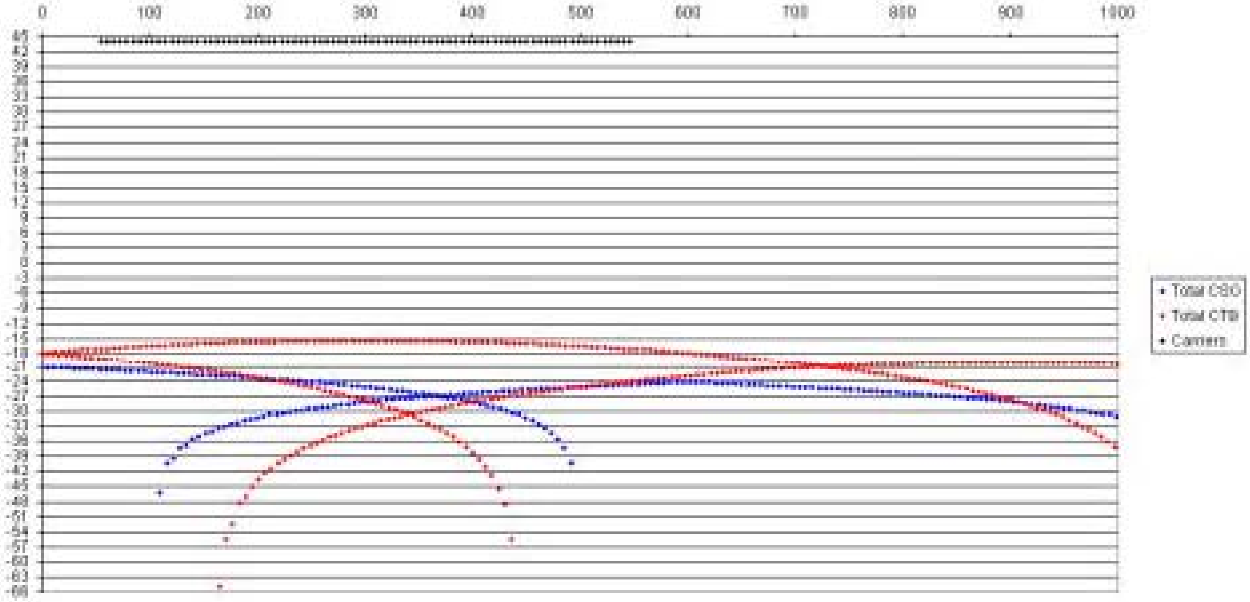

FIGURE 8. Analog Composite

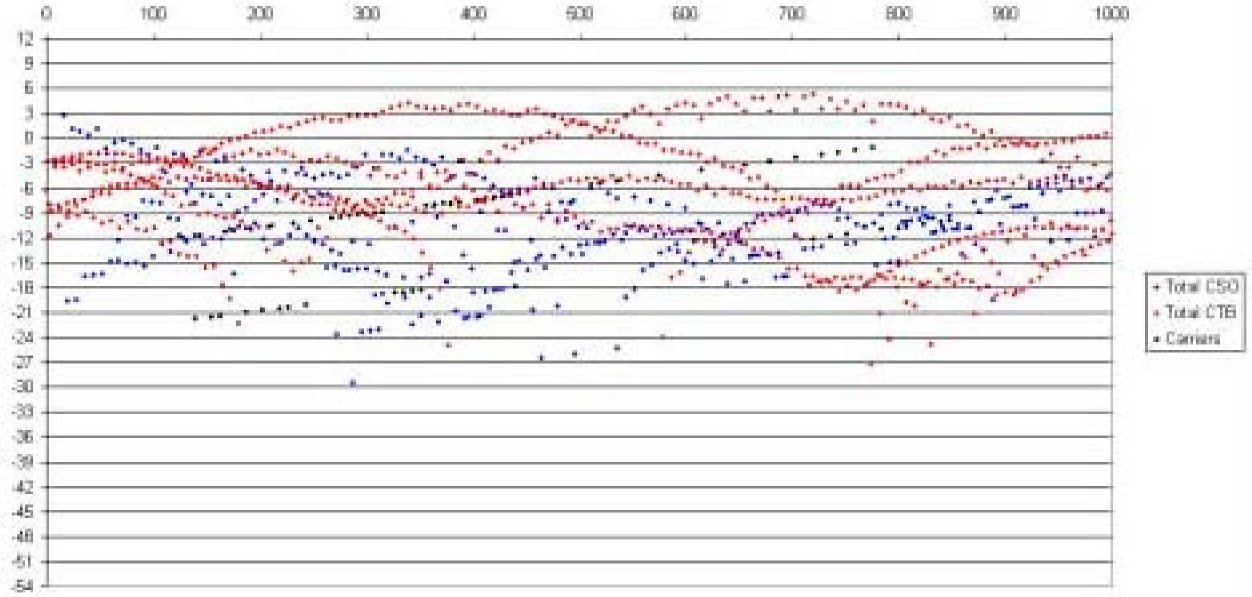

FIGURE 9. Analog Composite

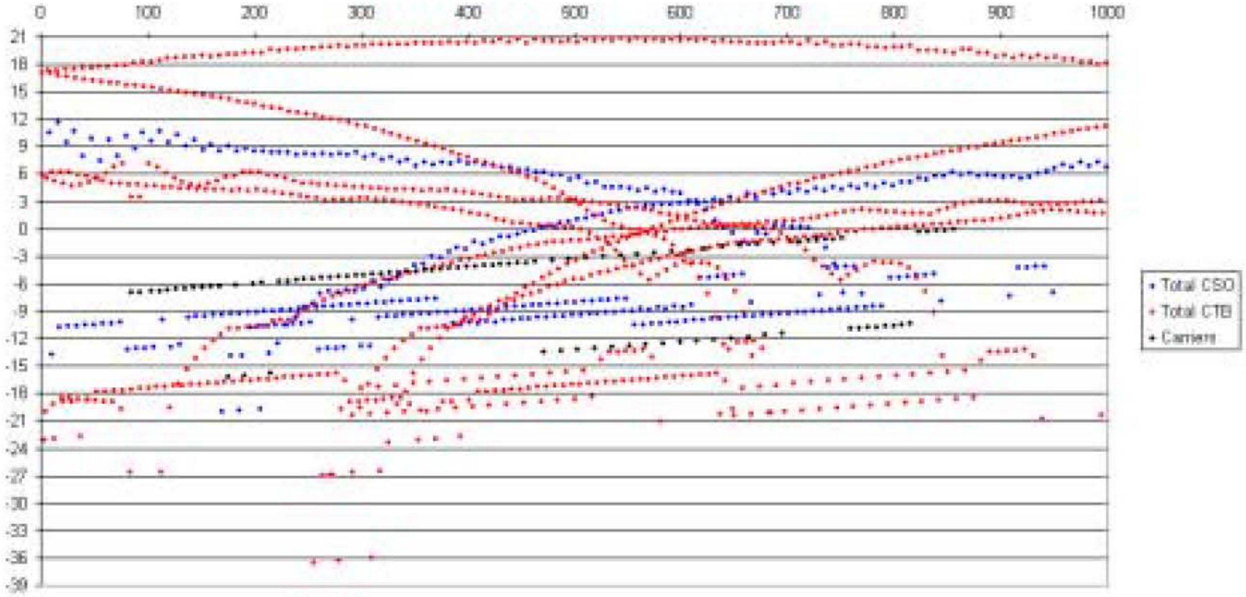

FIGURE 10. Digital Composite

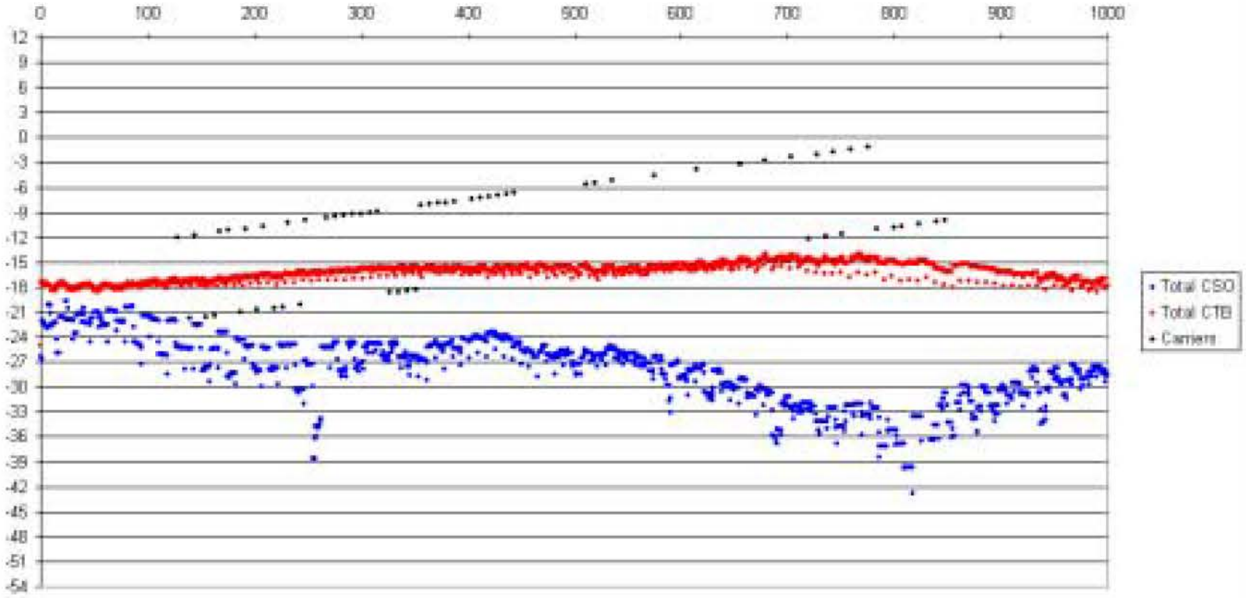

FIGURE 11. Digital Composite

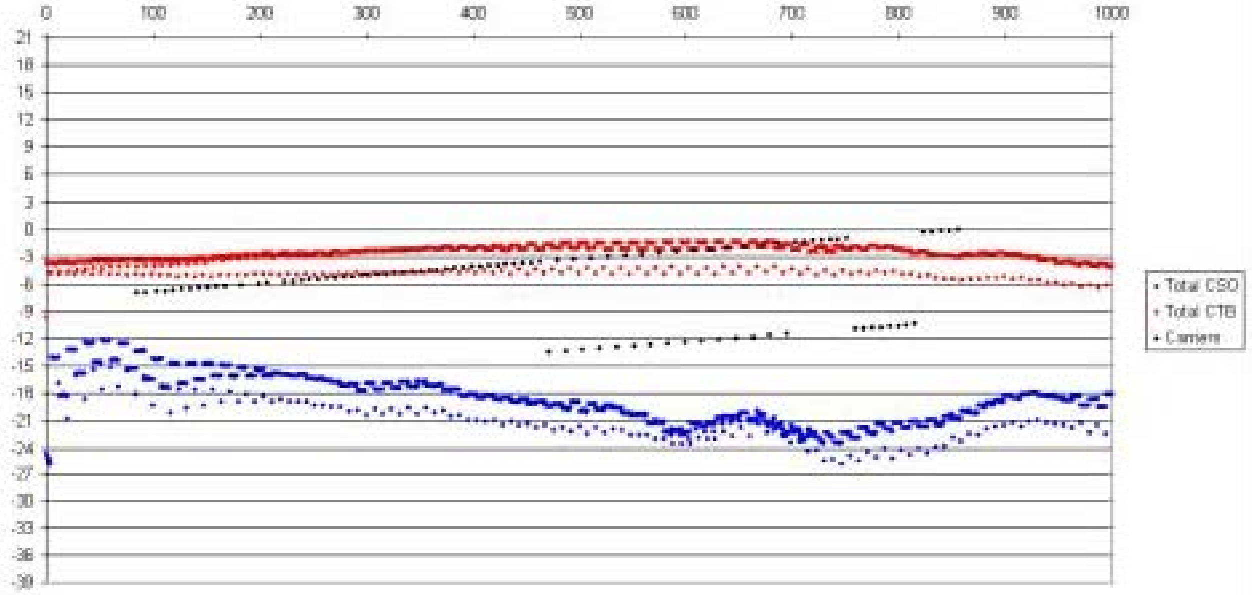

FIGURE 12. 2A Distortion vs. Frequency

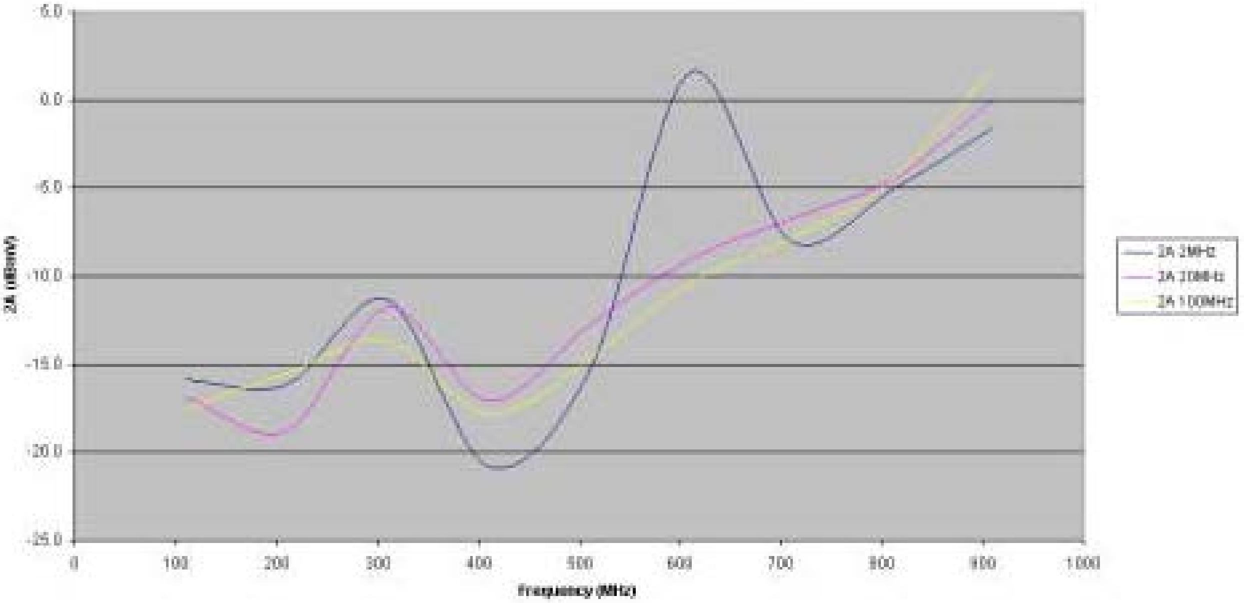

DETERMINING READINESS FOR 2-WAY INTERACTIVE SERVICES

Bob Harrison
Spyglass Integration

Abstract

A team of RF, system integration, quality assurance testing, and software development engineers at Spyglass Integration has recently created a comprehensive suite of testing and measurement tools and methodologies that characterize the downstream and return path bandwidth utilization for different classes of 2-way interactive services such as VOD, tCommerce, and unified messaging. These characterizations can be mapped to an operator's existing broadcast and return system network as a means to identify and mitigate bottlenecks and realize a balanced delivery of services for both steady-state and peak subscriber access.

In this paper, we will introduce these tools and methodologies that may enable operators to determine their current network's suitability for deploying 2-way interactive services, and identify where infrastructure investment or bandwidth allocation modifications may be considered to achieve required quality of service for subscribers.

Background

For interactive TV applications it all comes down to the subscriber's perceived quality of the service. Does the interactive guide fully populate with program information? Is a VOD purchase request properly provisioned, enabled, and billed? Is an acknowledgement for a commerce transaction quickly provided? Do messaging services such as multi-player gaming chat and e-mail provide a responsive user interface?

When the data transport bandwidth for downstream and return systems in cable plants is exploited for emerging interactive services, it is important to know how well the systems which support these services function, perform, and scale. Bandwidth bottlenecks in the network topology need to be anticipated and identified. Application server response, as it is integrated within the network, needs to be stressed and measured. The consumer set-top terminal's ability to receive and transmit messages influences perceived performance.

It is possible to functionally test an interactive service by configuring a test bed consisting of an application server which hosts an interactive service integrated with a digital head-end on an isolated RF network with a representative set top terminal. Anticipated command and control messaging and data flow between the application server and the set top may be observed and analyzed, and a service's functional goals may be validated with respect to an agreed upon service specification. But this functional validation is not sufficient for deployment consideration by a network operator. The operator is concerned with the stability, performance, and scalability of the service functionality as tens or hundreds of thousands of customers subscribe to and use the service. Will the newly introduced service fail? Even worse, will the resources absorbed by the new service break existing and stable revenue generating services such as core digital video broadcast and PPV?

How can a network operator or a vendor of a new interactive service predict that the new service will not impact current network operations, and secondly, provide a level of service quality that will meet the expectations of all the subscribers who are offered the service?

Current Service Evaluation Practices

Once a vendor of an interactive application has demonstrated that the service meets its specified functionality through a thorough validation or acceptance test plan, operators engage in a phased approach to understand issues of performance, stability, and scalability of the service, as it applies to their unique network environment, without negatively impacting currently deployed services.

Operators have created laboratories that attempt to replicate their operating network so that they may stage the service in a familiar environment. For the first time, the service is integrated in a head-end which maps the component versions, configurations, third party video distribution and data network infrastructure products that represents the operator's deployed operations. Service functionality may be revalidated at this point, but what about performance, stability, and scalability?

Internal "Friendlies"

Possibly ten to a hundred operator employees will be given access to the new service and asked to "give it a try". If the service fails or performs poorly, these non-subscribers (friendly users) will report their observations and impressions in a qualitative way. Rankings on a scale of "1 (poor) to 10 (excellent) " are solicited. These friendlies are not quality assurance specialists performing evaluations based on formal test procedures. They are considered representative of exercising the kind of service interaction that can be expected of subscribers. Are the friendlies all accessing the service at the same time ? Are they accessing all the features offered by the service? Are they examining boundary conditions or service inflection points as a means to examine extreme stress scenarios? Not necessarily.

The goal of this internal friendlies trial process is often to ascertain the stability of the application server and set top client application over a long period of time (weeks to months) and to understand major issues of service stability (does the service crash or become unavailable) to anticipate subscriber acceptance of the service. This level of testing, performed on an isolated network (an internal laboratory head-end) does not predict service performance, stability or existing network integrity as downstream and return path data communication bandwidth by the service approaches the nominal or peak utilization of a subscriber population in a specific property. Nor does it address the load of the application server itself (ability to service transaction requests). However, the level of confidence that the service may one day be considered deployable may be enhanced, because the service is consistent with the configuration and version of deployed network elements.

Bank of Set Tops

Within the laboratory evaluations, operators (with cooperation from their network infrastructure vendors and the interactive application service provider) often attempt stress testing by configuring many set-top boxes in a scripted or automated test harness. Using tools such as TestQuest, Inc.'s

TestQuest Pro that can replay streams of scripted IR commands, monitor the results produced on screen, and provide comparison with reference images, it is possible to repeatedly and deterministically emulate viewer behavior and create a methodology to invoke all service features across a finite number of set tops which have been allocated for the task. Even if hundreds of set tops are provisioned for this process, it still falls short of the subscriber population that will be expected to be supported by an operational head-end system.

Limited Operational Field Trials

Once the internal friendlies evaluation has been performed and (optionally) laboratory stress techniques have been analyzed, the service may become a candidate for a field trial. The operator selects a candidate property, and the service in integrated within an operational head-end. A small subscriber population is selected to evaluate the service. These subscribers are again friendly to the evaluation; it is not expected that they will discontinue service should they experience service disruption or other anomalies. The greatest value of the limited operational field trial is that the service functionality may be validated within an operational network. Again, confidence for total subscriber scalability has not been gained.

A New Approach

When evaluating a 2-way cable plant's suitability to support the introduction of an interactive service, several characteristics need to be studied:

1. The service introduction will not impact the actual or perceived delivery existing deployed services.
2. The server that supports the interactive application service must be shown to scale for the expected subscriber population request load (both nominally and during peak utilization)
3. Bandwidth limitations in the data network (downstream in-band and out-of-band) and return system must be identified so that bottleneck issues may be alleviated by network element upgrade or addition or topology reconfiguration.

A *Meaningful* Load Tester

Raskin and Stoneback suggest, "HFC network performance monitoring is likely to be done most effectively by collecting and coordinating communication performance information from the applications running over the network"[i]. This implies that network performance monitoring needs to be performed in the context of the applications that the network is expected to support, not simply loading a network with variable volumes and frequencies of data payloads. In response to this suggestion, Spyglass Integration created an application load tester and IP network interactivity tester which provides the flexibility to coordinate application oriented communication messaging and collect the relevant statistics with the goal of understanding HFC data network performance in a meaningful context.

Load testing addresses the objective to interject significant packet data in a cable data network in an attempt to load the system with the level of transaction traffic that can be expected by a realistic subscriber population. Load testing can be designed to be a vehicle to provide insight and analysis for throughput for either a single set top or many concurrent set tops. Throughput in this context is defined as the time it takes to receive a response at a set top for each message request or acknowledgment sent by the set top to an application server.

Throughput analysis comprises a technique and measurement capability to create a *meaningful* request from a set-top to an application server and measure the time for a *meaningful* response to be received by the set-top. The actual interactive service is invoked and satisfied by the application server.

Concurrent throughput analysis is the ability to measure processing speed from multiple set tops (the time it takes to receive a response at a set top for each command request or acknowledgement sent by "N" number of set tops to an application server.

Concurrent throughput analysis comprises a technique and measure capability to create a *meaningful* request from multiple set-tops to an application server and measure the time for a *meaningful* response to be received by the set-tops. The actual interactive service is invoked and satisfied by the application server. The behavior for each set top configured for concurrent throughput analysis must be separately identified and measurable.

Packet Characterization

Message packet characterization comprises the ability to record downstream packet characteristics from an application server communicating with a single set top for the following attributes:

- Size of packets
- Frequency of packets
- Information contained in the packets (content sensitive)

Concurrent packet utilization requires the ability to record downstream packet characteristics for "n" number of set tops for the following attributes:

- Average size of packets
- Average frequency of packets

To capture packet characterizations it is necessary to probe (or sniff) communication between a set top and an application server. A model can be built which represents the size, frequency, and content for a message set between the set top and application server. This model can be used to build script testing scenarios to generate meaningful requests and responses. Scripts that use these packet characterizations may be invoked to create repeatable network load that is representative of true interactive service messaging.

System Loading

In addition to characterizing a specific interactive service under test and evaluation, it is important to interject load that represents the delivery other data or video services. Therefore, a load testing environment must also provide the ability to simulate traffic unrelated to the application service under investigation, yet representative of other network functions (conditional access messaging, program guide data carouseling, PPV purchase polling, etc.).

A Load Tester for Motorola Networks

A load testing simulator had been designed and built by Spyglass Integration to drive application level requests in a Motorola DigiCable environment. The DCT 2000 set top is a proprietary platform that provides fundamental services on a network (UDP message packetization, DAC-6000 communications, and NC-1500 communications) as well as providing the fundamental RF network interfaces with out-of-band modulators (OM) and return path demodulators (RPD). It is possible to create a communication proxy application for the DCT

2000 that enables a PC based application simulator to use the DCT 2000 as a HFC RF gateway. The proxy essentially provides the out-of-band and return path network communication services requested by an application on the PC.

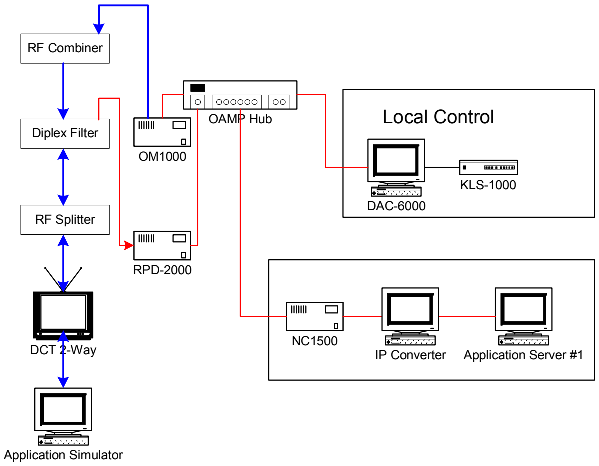

Figure 1

Figure 1 illustrates the integration of a PC based application server with the DCT 2000 through the DCT 2000's serial port. The DCT 2000 can be thought of as a tethered out-of-band and return path modem for the PC. The DCT 2000 is used only as a network modem. There is no need for compute intensive on screen display functions, IR remote interrupt service routine handling, or other data manipulation.

The application simulator may be scripted to create the various request messages a specific interactive application server expects. Commands to the DCT 2000 proxy place the messages on the HFC network for routing to the application server. The scripting is specific to each application server, and generally takes the form of an ordered series of messages as defined by the messaging protocol used by the set top client application developed for the service. For example, if the application server is a VOD server, the messages scripted for the PC application simulator would be the series of VOD commands that can be expected by a native VOD client on the DCT 2000 (session establishment, stream control, etc.).

A single application simulator PC and DCT 2000 proxy can be used to emulate multiple set top sessions if the IP address associated with the DCT 2000 proxy can be altered prior to delivering a message to the application server. This has been accomplished by creating a hardware / software gateway (indicated in figure 1 as the "IP Converter"). By changing the IP address of the packet prior to delivering it to the application server, the server can be spoofed to handle multiple logical sessions with multiple set tops. Through an IP mapping management technique, application server responses can be redirected back to the originating DCT 2000 proxy / application simulator, or another network device to collect the characteristics of the response.

The proxy service for the DCT 2000, the IP converter, and data (packet) collection and analysis utilities resident on the application simulation are created once and are independent of the class of interactive service which is being evaluated. The only variable is the messages and state transition protocol for the interactive service as defined by the vendor of the application server, emerging open standards, or by packet characterization probing.

Although we have implemented this environment with a DCT 2000 in a Motorola DigiCable network, the technique may be applied to virtually any infrastructure provider's network. The prerequisite requirement is for a set top application development environment to expose the network services required for return path and downstream (in-band and out-of-band) communications, as well as serial or Ethernet communications to a tethered PC.

Scalability of the Load Tester

The load tester can be used to simulate the realistic request / response application messaging of hundreds of set tops from a single set top. If one can characterize the anticipated consumer generated frequency of message requests (often represented in units of requests per minute or requests per hour), the scripting engine within the application simulator PC may be configured to generate the requests expected in nominal and peak utilization times by hundreds of subscribers. The IP converter spoofs the application server to believe that these requests originate from different set tops.

By positioning multiple load testing systems as front ends to independent collections of multiple out-of-band modulators, return path demodulators, and network controller elements in a network, it is possible to stress and identify bottleneck conditions with respect to independent network segments. The capacity, as indicated by set top population, of a out-of-band or return path segment can be measured, with respect to the application service being evaluated and the nominal loading of the delivery of existing data and video services. This is determined by creating scripts for each application emulator that provides realistic application server requests. Data (message packet size for each request and response and server response time) is collected for independent network segment to evaluate the equilibrium point between request and response message cycles. This equilibrium point represents the traffic scenario where the response time delay is deemed to be unacceptable to user's experience. This equilibrium point correlates to the number of set tops that can support a required level of service quality given the network configuration of an independent network segment. It is possible to use this data to develop a model of how set-tops may be distributed among multiple OM / RPD / NC network configurations to empirically achieve the performance goal required for an expected subscriber population.

We have found that the communication proxy application within the DCT 2000 can be extended to provide more than a protocol gateway. Since the DCT 2000's serial port can serve as a bottleneck for requests issued by the application simulator, request messages from the application emulator may be pre-cached in the DCT 2000, sequenced, and issued on a scheduled basis to an application server.

Discounting Application Latency

Sometimes the response latency of the application server is the bottleneck. Degradation of application response times is a function of application server performance, rather than the network. To remove application latency from a network performance characterization study a second tool has been developed, called the "Interactivity Tester".

The Interactivity Tester can be though of as a client message generator and server reflector of known UDP packets between a set top and the physical network location of an application server. Figure 2 illustrates the components of the Interactivity Tester, and its integration in an HFC network.

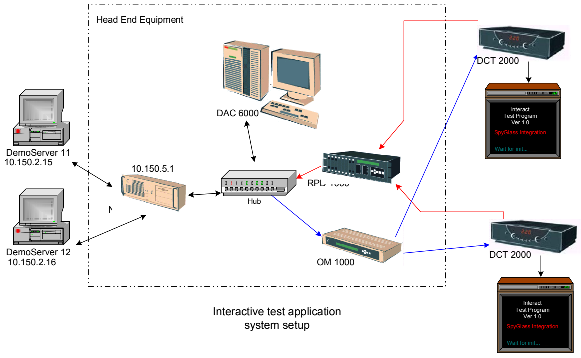

Figure 2

To remove application server latency, a PC-based network echo server (labeled “Demoserver” in figure 2) is positioned at the same network topology location as the application server. This echo server may be flexibly configured to provide UDP packets with message patterns identical to those that are expected to be provided by an application server in response to a request. A client component, installed on the set top, can be configured to initiate requests to one or more echo servers. The messages sent comprise payloads that emulate application message requests, and can be throttled at variable periodic frequencies (within a 1 second resolution). Figure 3 represents the typical on screen display that the set top client agent provides.

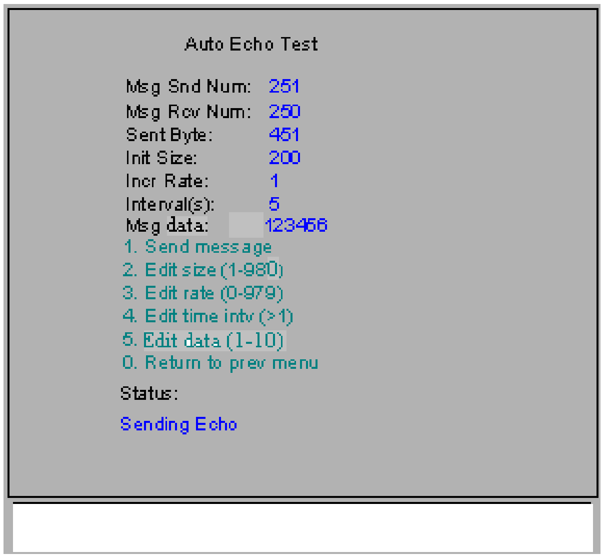

Auto Echo Test

Msg Snd Num: 251
Msg Rcv Num: 250
Sent Byte: 451
Init Size: 200
Incr Rate: 1
Interval(s): 5
Msg data: 123456

1. Send message
2. Edit size (1-980)
3. Edit rate (0-979)
4. Edit time intv (>1)
5. Edit data (1-10)
0. Return to prev menu

Status:

Sending Echo

Figure 3

The echo server is configured to provide a selected response to the client’s request. Figure 4 illustrates the method to build the response message.

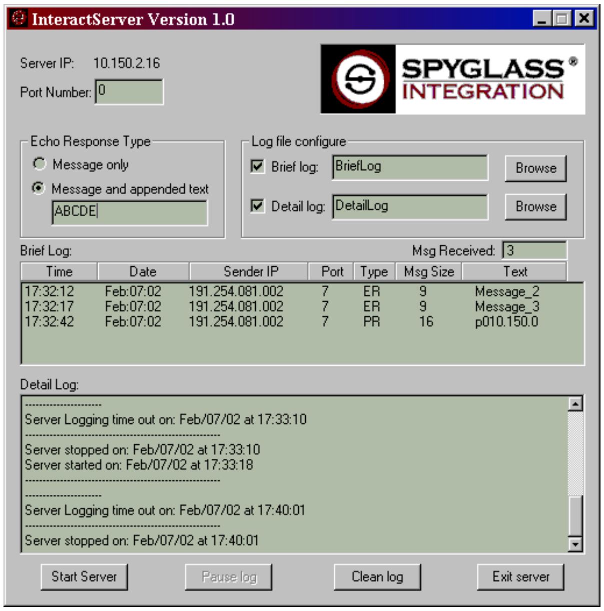

Server IP: 10.150.2.16

Port Number: 0

Echo Response Type
○ Message only
● Message and appended text
ABCDE

Log file configure
☑ Brief log: BriefLog
☑ Detail log: DetailLog

Brief Log:

Msg Received: 3

Time	Date	Sender IP	Port	Type	Msg Size	Text
17:32:12	Feb:07:02	191.254.081.002	7	ER	9	Message_2
17:32:17	Feb:07:02	191.254.081.002	7	ER	9	Message_3
17:32:42	Feb:07:02	191.254.081.002	7	PR	16	p010.150.0

Detail Log:

```
--------------------
Server Logging time out on: Feb/07/02 at 17:33:10
------------------------------------------------
Server stopped on: Feb/07/02 at 17:33:10
Server started on: Feb/07/02 at 17:33:18
------------------------------------------------
--------------------
Server Logging time out on: Feb/07/02 at 17:40:01
------------------------------------------------
Server stopped on: Feb/07/02 at 17:40:01
```

Figure 4

The echo server also provides a detailed logging facility with respect to transactions generated by the set top client, and responses issued by the echo server. Figure 5 illustrates an example of this detailed log file.

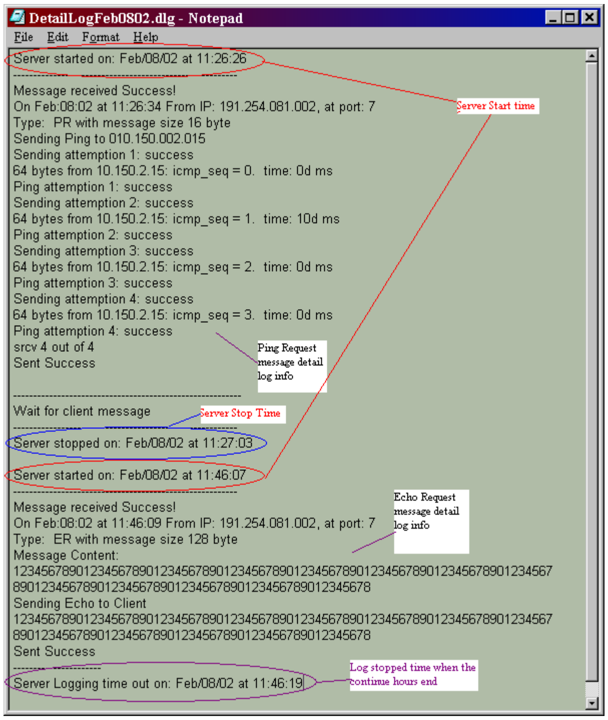
DetailLogFeb0802.dlg - Notepad
File Edit Format Help
Server Start time
Ping Request message detail log info
Server Stop Time
Echo Request message detail log info
Log stopped time when the continue hours end

```
Server started on: Feb/08/02 at 11:26:26
------------------------------------------------
Message received Success!
On Feb:08:02 at 11:26:34 From IP: 191.254.081.002, at port: 7
Type:  PR with message size 16 byte
Sending Ping to 010.150.002.015
Sending attemption 1: success
64 bytes from 10.150.2.15: icmp_seq = 0.  time: 0d ms
Ping attemption 1: success
Sending attemption 2: success
64 bytes from 10.150.2.15: icmp_seq = 1.  time: 10d ms
Ping attemption 2: success
Sending attemption 3: success
64 bytes from 10.150.2.15: icmp_seq = 2.  time: 0d ms
Ping attemption 3: success
Sending attemption 4: success
64 bytes from 10.150.2.15: icmp_seq = 3.  time: 0d ms
Ping attemption 4: success
srcv 4 out of 4
Sent Success

-------------------------------------------------
Wait for client message
------------------------------------------------
Server stopped on: Feb/08/02 at 11:27:03

Server started on: Feb/08/02 at 11:46:07
------------------------------------------------
Message received Success!
On Feb:08:02 at 11:46:09 From IP: 191.254.081.002, at port: 7
Type:  ER with message size 128 byte
Message Content:
12345678901234567890123456789012345678901234567890123456789012345678901234567
890123456789012345678901234567890123456789012345678
Sending Echo to Client
12345678901234567890123456789012345678901234567890123456789012345678901234567
890123456789012345678901234567890123456789012345678
Sent Success
--------------------
Server Logging time out on: Feb/08/02 at 11:46:19
```

Conclusion

The Load and Interactivity Testers are tools that may be applied to a laboratory HFC network to economically introduce a level of interactive service utilization to evaluate the network's and application service's abilities to scale for a subscriber population representative of a deployment environment. The tools provide the ability to characterize the data packet messaging between a set top and an application server, to script and invoke multiple concurrent sessions, and collect empirical data that represents messaging behavior. By analyzing this data, it is possible to identify either network bandwidth constraints or the capability of an application server to support a specific subscriber load over varying utilization assumptions.

Credits

The following Spyglass Integration team members contributed to the concept, design, and implementation of the Load and Interactivity Testers: Robert Cleroux, Tony Curran, Jim Desmond, Mike Foss, Bob Frankel, Ben Li, Ed MacDonald, George Sarosi, Martin Wahl, and John Zhang.

Bob Harrison is a Solutions Architect with Spyglass Integration located in Lexington, Massachusetts. He may be reached at *bharrison@spyglassintegration.com*

[i] Donald Raskin and Dean Stoneback, *Broadband Return Systems for Hybrid Fiber/Coax Cable TV Networks*, (Upper Saddle River, NJ : Prentice Hall PTR, 1998) p. 241

DEVELOPING MIDDLEWARE FOR THE OCAP MARKET

Joel Zdepski, Ph.D.
OpenTV, Inc.

Abstract

CableLabs, through the Open Cable process, has decided to adopt certain elements of the MHP standard. In this paper, we will provide perspectives on the nature of the differences between the MHP and its OCAP equivalent. Some of the implications of the changes as viewed from a middleware vendor and an application vendor will be discussed. OpenTV's road map for implementation of an OCAP/MHP solution that will also afford backwards compatibility with interactive solutions for existing thin set top boxes. The content migration path from thin middleware to OCAP and standards-compliant product architecture will be reviewed. We will also produce a forward-looking assessment of market trends in broadband/iTV content and resulting implications for set top box configurations and the software stack.

BACKGROUND

The DVB is a European industry group with over 250 member companies who have significant interests in the development and promulgation of standards for distribution of digital video services. The group has developed many standards since its inception in the early 1990's, including modulation formats for Cable, Satellite and Terrestrial delivery, System Information formats for use in EPG implementations, and electrical interfaces for equipment interconnection, just to name a few. More recently its members donated considerable of their resources in the specification of an API for set top boxes in order to foster interoperability of content.

The technical specifications developed to date by the DVB consist of two primary sub-components, an Execution Engine (EE) based on Java, and a Presentation Engine (PE) based on HTML/JavaScript. Cable Labs has also been developing similar specifications under the headings of OCAP 1.0 and OCAP 2.0. In a significant development, Cable Labs has aligned its standards with those developed by the DVB as much as possible. This decision is important because it offers the hope the vendors of Applications, Middleware, and Infrastructure will be able to leverage their technical developments across several markets.

The DVB released Version 1.0.2 of MHP in late February 2002, and the specification will be published as an ETSI standard. Version 1.1 of the MHP specification has also been approved. While version 1.0.2 only contains Java elements, version 1.1 includes HTML/JavaScript elements as well. It must be pointed out that these HTML/JS elements are optional, while all MHP receivers must be Java capable. It must also be noted that certain functions of the receiver, such as channel change, are only accessible through Java interfaces.

The next milestone to be reached to enable deployment of MHP products is the release of MHP Test Suites, which will be made available through ETSI to all implementers of MHP run-times and receivers. Implementers can then self certify their products by issuing an official statement that their product has successfully passed the Test Suites and claim MHP compliance. Only compliant products can be put on the market. Non-compliant products would be violating the legal rights of the DVB concerning the MHP mark as well as the IP rights of third parties in the MHP specification, because licensing of these rights under the DVB MoU only covers compliant products. OpenTV's belief, a view shared by the DVB, is that MHP will be successful only if all MHP receivers are fully interoperable. As a consequence, OpenTV is in favor of a strong compliance regime supported by a very comprehensive Test Suite. Such regime will guarantee that any MHP compliant application will run on any MHP compliant receiver, just like any VHS tape is expected to play on any VHS VCR.

There is currently no schedule for the development of a Test Suite for MHP 1.1. As a consequence, it is impossible to predict when 1.1 compliant receivers will be introduced in the market

A second milestone relates to the arrangement of IPR pooling organizations patent rights that will enable equipment implementers so secure from a single entity a significant fraction of the required IPR for providing a licensed implementation.

OCAP 1.0 DEVIATIONS FROM MHP

There are several different types of changes that adopter of the MHP specification may need to make in order to adapt to a local market. These changes can be divided into the following classes:

- Underlying-media/legal – changes required by the underlying transport infrastructure or by the region's laws. These changes generally require significant modifications or additions of code to support and as such are justifiable regional modifications to MHP.

- Language/cultural – changes required by the cultural and language differences of a region. Though justifiable regional modifications to MHP, these changes generally require minimal code changes and are usually embodied by changes in the data the code uses.

- Extensions/enhancements – changes desired by the region to offer services beyond those provided by the originating specification. These changes are acceptable as long as these extensions enhancements do not impact upon the specification as they form superset of MHP. In order to maintain the integrity of the receiver population, it is recommended, however, to put in place a process that will review these extensions, reject unnecessary derivatives and include approved extensions into future versions of the MHP specification.

- Cost – changes desired by the region to minimize the cost of implementing the specification. The number of issues under this category is very large and can range from minor to very significant differences in cost. E.g. Using DVB-SI in the US is technically possible but not practical in terms of cost. As such, some cost issues are justifiable regional issues while others are near term expedients and should be avoided.

- Technical improvements – changes desired by the specification writers based upon their belief that they have a better technical solution. These types of changes are usually not wise and should be resisted, since they lead to incompatibilities with MHP. It is better for all involved if truly superior technical solutions are submitted to the DVB for inclusion in future versions of the standard.

- Business Model – changes desired by the region to preserve or develop a given business model. Such issues often not advisable as they may lead to loss of application interoperability.

- Error – there is an error in the specification. As with improved technical design, corrections to error in the MHP specification are best fed back into the DVB for inclusion in future versions.

The changes made to OCAP 1.0 encompass most of these categories. Some of the key changes include (but are not limited to):

- Excising non-Cable protocols
- Substitution of AC3 Audio for MPEG-1 Layer 2 audio.
- Removal of DVB SI.
- Removal or Modification of subtitling and teletext.
- Replacement of AIT with XAIT.
- Modification to the application lifecycle.
- Prohibition of DVB HTML descriptors.
- Namespace changes for Xlet’s, etc.

While many of these deviations are logical regional variations, we believe that some of them should be revisited during the Corrigenda process for the OCAP 1.0 specification. In particular, changes, which result in undesirable side effects for applications being ported from a pure MHP market to an OCAP market should be reconsidered and in many cases, revised.

Perhaps one of the changes with the most far-reaching effect is the prohibition in OCAP 1.0 for signaling DVB HTML content. HTML is the most pervasive and important of the Internet protocols after TCP-IP. Restricting its use in the standard removes a key tool in achieving content interoperability across heterogeneous STB platforms. The US cable market continues to convert from analog to digital at a robust rate, and for the foreseeable future this is being achieved with STB’s purchased by the MSO and decidedly below the capabilities required to deploy the MHP specification. Content authoring and distribution in HTML/JavaScript affords many possibilities for achieving interoperability through transcoding to formats that can be rendered by the lightweight receivers that have already been deployed by US MSOs.

OPENTV’s MHP/OCAP IMPLEMENTATION

OpenTV’s Advanced STB implementation is shown in Figure 1. In addition to our MHP stack it includes a light-weight virtual machine and the Device Mosaic Browser. We distinguish four main layers in our MHP stack, where upper layers leverage resources shared by lower layers:

- The driver layer, which provides interfaces to the hardware through generic portability layers that guarantee interoperability between different hardware implementations. Because of OpenTV’s leading market share, our portability layer has become the de-facto standard in the industry, which is now emulated by our competitors. The design of our portability layer allows multiple clients to share the same resource through a single set of APIs. More specifically, C, Java and HTML components co-exist on the receiver and access the hardware through a common set of APIs that handle arbitration and serialization of requests where required.

- The kernal layer, which provides interfaces to the processing resources through generic portability layers that guarantee independence between different operating systems. C, Java and HTML components access the operating system through a common set of APIs.
- The Interactive TV libraries layer, which implements the core functionalities of our ITV run-time, such as communication, graphics, security, etc. C, Java and HTML components co-exist on the receiver and access these libraries through a common set of APIs. Arbitration and serialization of requests is implemented through policies such as application life cycle that are captured in our Control Task driver.

The Execution layer, which provides independence of the application binaries from the CPU of a particular receiver through an interpretive abstraction. We currently offer ANSI C, Java and HTML/JavaScript execution environments. It is likely that we will also introduce a Flash execution environment at a date to be determined. OpenTV is the only ITV middleware company offering an ANSI C execution environment, deployed on over 24 million receivers. Our Java environment is based on Sun Microsystems VM. Our HTML/JS environment is based on our Device Mosaic technology, which has already been licensed to PowerTV, Sony, Motorola, Tivo and WorldGate deployed on over 6 million cable receivers

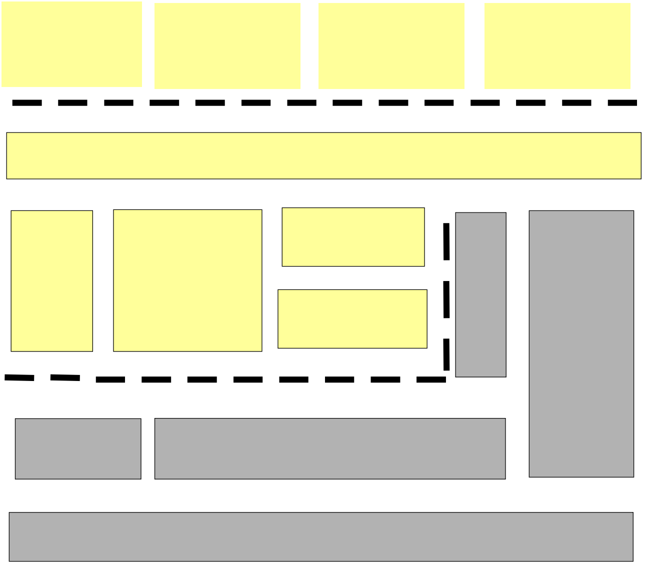

Figure 1: MHP/OCAP Software Stack

The main benefits of an integrated architecture include a smaller footprint, as well as flexibility in possible evolutions of the product, such as embedding Java scripts in an HTML page or carrying Flash content in DSM-CC carousels. It also allows us to quickly introduce new standards as they become available. One example is the implementation of an ARIB compliant BML module as an extension of our XHTML engine for the Japanese market. Finally, this architecture also allows us to make all features of our ITV libraries available to all execution environments. For example, our ITV libraries can support PVR functionalities (see for example the integrated PVR product introduced by Via Digital later this year), which can be exposed to our Java module. The benefit is a Java execution environment that is fully MHP compliant but can also offer features that are not currently covered by the MHP standard, such as PVR. In other terms, our architecture can continue to progress at the forefront of the state of the art while incorporating standard components, as they emerge.

CONTENT INTEROPERABILITY

OpenTV sees the emergence of networks with multiple tiers of receivers, which will offer different levels of capabilities. Basic receivers with limited processing power and memory will remain dominant. Some of them will be able to render HTML content, but most of them, for the short to medium term, will not have the hardware capabilities required to render Java based content such as MHP. On the other hand, we expect the emergence of high to very high-end receivers with mass storage. These receivers will have enough hardware capabilities to execute MHP applications. OpenTV offers a number of solutions to enable delivery of content on such hybrid networks.

One option is to develop content around OpenTV’s C based APIs. Since both low-end and high-end receivers include OpenTV’s C player, C based content can be executed on the entire population of receivers.
A second option is to develop content around OpenTV’s HTML/JavaScript based APIs. OpenTV’s HTML engine can run on both mid-range and high-end receivers. In addition, OpenTV is currently developing an extension to its Publisher product that compiles HTML/JavaScript content into OpenTV’s lightweight byte code (named o-code). Since low-end receivers include OpenTV’s C player (which includes the o-code interpreter), it is possible to execute this content on these receivers. As a consequence, it is possible to create ITV content once, and deliver it to the entire range of receivers, either through Publisher for low-end receivers, or directly for mid to high-end receivers. As market demand arises, OpenTV will consider extending Publisher to support the HTML profile of the MHP 1.1 specification.

Another option for content migration is to create multiple executables for different classes of receivers, while sharing the data for all classes of receivers. In this scenario, a C based executable would be created for low-end receivers and an MHP version would be created for MHP capable high-end receivers. These receivers can already share data provided through the return path, since both classes of receivers support the same communication protocols (HTTP, TCP). The benefit there is to use a single Web server infrastructure for all receivers. It would also be possible to for the receivers to share broadcast data. While our C player and our MHP extension currently support different carousel formats, as market demand arises, OpenTV is ready to implement a common broadcast stack for its C and MHP players in order to share carousels. The benefit would be to reduce broadcast bandwidth consumption.

CONCLUSIONS

The OCAP 1.0 specification is well aligned with the MHP 1.0.2 specification. Software Vendors will be able to significantly leverage their MHP development when developing for the OCAP market, however maintenance of two test regimes for MHP and OCAP remains a costly by product of the deviations which OCAP takes from the MHP specification. HTML/JavaScript is expected to form the largest body of interoperable content in interactive TV and fact that OCAP 1.0 does not provide for DVB HTML signaling.

DYNAMIC ADAPTATION TO IMPAIRED RF UPSTREAM CHANNELS USING ADVANCED PHY

Daniel Howard Broadcom Corporation
Hal Roberts ADC

Abstract

The use of advanced receiver processing and system adaptation in the cable modem termination system (CMTS) can easily quadruple the upstream capacity by opening up new RF spectrum and by more efficiently using existing RF spectrum. The advent of DOCSIS 2.0 provides a 'toolkit' of physical layer features that allow this potential. However, since the upstream spectrum in previously unused portions of the band is highly dynamic in the level and type of interference present, it is critical that the CMTS be able to dynamically sense and adapt to changing channel conditions. Such dynamic adaptation ensures that the channel remains active, even in the presence of strong interference, and (as importantly) ensures that as the channel conditions improve, the capacity is restored to higher levels. The DOCSIS 2.0 'toolkit' provides many more options for the CMTS to handle ingress while maintaining high bandwidth, but without well designed adaptive algorithms the toolkit will be unused or worse, ill-used.

In this paper, an intelligent CMTS with advanced receiver processing and advanced system algorithms for dynamic adaptation is shown to provide significant benefits to existing deployments of DOCSIS 1.0 and 1.1 cable modems, as well as set the stage for future improvements using DOCSIS 2.0 technology. The performance improvements will be demonstrated in the presence of all of the most common upstream plant impairments: additive white Gaussian noise (AWGN), ingress, common path distortion (CPD), and impulse/burst noise. Mitigation of these impairments will be shown to open up spectrum below 20 MHz that may previously have been considered unusable. Further, the reliability of interactive services is increased by such dynamic adaptation, improving the market appeal of applications such as voice over IP over cable. Since the CMTS cost per subscriber is far less than either the cost of the cable modem itself and/or further plant upgrades, the solution described in this paper provides the lowest cost and fastest time to market approach for quadrupling the upstream capacity of existing cable modem networks.

INTRODUCTION

The spectral efficiency of current cable modem RF upstream channels can be greatly increased now that advanced PHY technologies such as higher orders of modulation, increased error correction, interleaving, ingress cancellation processing and better equalization are readily available in modern CMTS systems. Twice the spectral efficiency can be obtained by operating at 16 QAM instead of QPSK, more if the Reed Solomon (RS) forward error correction (FEC) overhead can be reduced. Three times the spectral efficiency can be obtained if new advanced PHY modems are deployed which can operate at 64 QAM. Further, RF spectrum that was previously unusable due to ingress, impulse/burst noise, or lack of sufficient equalization can now be used by all modems on the network, thereby creating even more capacity on the upstream for cable operators.

To obtain the greatest increase in spectral efficiency on the RF upstream, cable operators will have to deploy both advanced PHY cable modems as well as an advanced PHY Cable Modem Termination System (CMTS). However, if current plants are operating at QPSK levels on the upstream, the greatest incremental increase will come from increasing the signal constellation from QPSK to 16 QAM, and opening up new RF spectrum that was previously unusable. This incremental increase can be obtained merely by upgrading the CMTS since current DOCSIS 1.x modems can operate at 16 QAM as well as QPSK, but may have previously been unable to use 16 QAM either due to ingress or due to lack of sufficient equalization. Further, the most likely deployment scenario is a mixture of 1.x and advanced PHY modems since significant numbers of 1.x modems have already been and are still being deployed. Hence, the question becomes how to use an advanced PHY CMTS with a mixture of current and future advanced PHY modems to obtain the most capacity out of the cable network without requiring redeployment of modems or plant upgrades. This is the question addressed by this paper, and the answer turns out to be via adaptation of the mixed cable modem network to the various RF impairments that exist on the network.

The paper begins with a description of advanced PHY features, with emphasis on those that apply to DOCSIS 1.x modems as well as advanced PHY modems. Next, a discussion of channel impairments and mitigation strategies will be used to show how an adaptive CMTS using advanced PHY features can keep channel capacities high most of the day, and only increase robustness (thereby dropping capacity) when the channel conditions require it. To accomplish this, an adaptation system is described which incorporates spectrum monitoring with control of CMTS burst profiles for the various traffic types being transported. The result is operation at high spectral efficiencies for the great majority of the day, meaning more throughput to users or alternately more users possible on current upstream channels.

ADVANCED PHY FEATURES AND APPLICATION TO DOCSIS 1.X MODEMS

As described in a companion paper in this session [1], DOCSIS 2.0 advanced PHY technology includes both advanced TDMA (ATDMA) and synchronous CDMA (SCDMA). These advanced PHY technologies increase the capacity in clean upstream channels by providing higher spectral efficiencies, with up to 64 QAM for ATDMA in the specification, (and up to 256 QAM in some vendor implementations). Advanced PHY also provides significant increases in the robustness of upstream signaling against the most common RF impairments: additive white Gaussian noise (AWGN), ingress of radio/navigation signals, common path distortion (CPD), and impulse/burst noise. For a detailed analysis and modeling of upstream impairments, the reader is referred to a previous NCTA paper by one of the authors, which also has references to other upstream measurements and modeling papers [2].

However, there are several features in advanced PHY CMTS systems that improve not only performance with advanced PHY CMs, but also the performance of existing 1.x CMs. A listing of several advanced PHY features, some of which apply to 1.x CMs is shown in Table 1.

Table 1. Advanced PHY Features

Feature	Improves 1.x ?
Improved AWGN Mitigation	
Lower implementation loss	YES
Better receive equalization	YES
Improved burst acquisition	YES
More FEC (T=16)	NO
Ingress/CPD Cancellation	
Cancellation of ingress	YES
Cancellation of CPD	YES
Improved Mitigation of Impulse and Burst Noise	
Cancellation of spectral peaks in periodic impulse noise	YES
More FEC (T=16)	NO
RS byte interleaving	NO
SCDMA mode	NO

As is seen from the table, existing DOCSIS 1.x modems will benefit significantly from deployment of a modern advanced PHY CMTS in mitigation of AWGN, ingress/CPD, and impulse/burst noise. The most significant improvement for 1.x modems is in the area of cancellation of ingress and/or CPD: over 30dB of narrowband ingress can be cancelled from the received spectrum. Since all of the processing for cancellation resides in the new CMTS, existing modems of any DOCSIS version will benefit from the advanced PHY technology. Figure 1 shows the improvement possible in 16 QAM mode against multiple ingressors, and Figure 2 shows the improvement possible against wideband ingressors. The latter impairment case is also applicable to a group of ingressors that must be cancelled as a zone as opposed to individual cancellation.

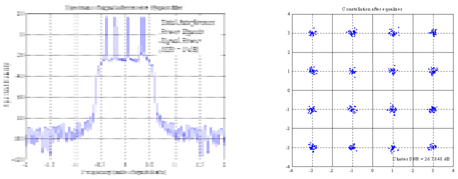

Figure 1. Cancellation of 5 Ingressors in 16 QAM mode, SIR=0 dB

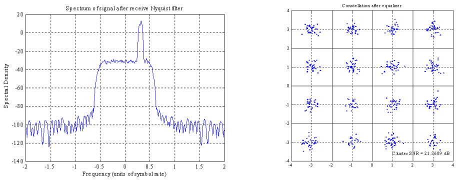

Figure 2. Cancellation of Wideband Ingress in 16 QAM mode, SIR=0 dB

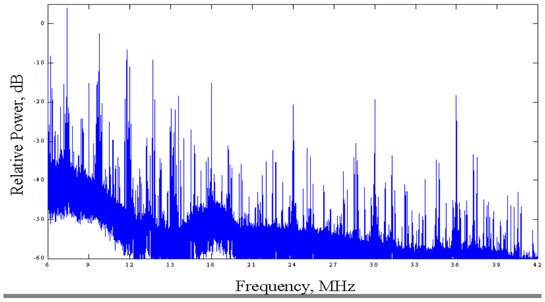

Figure 3. CPD may be seen as multiple narrow band ingressors, capable of being cancelled by an adaptive ingress filter.

The impact of the ability to cancel ingress and CPD from 1.x modems is significant: DOCSIS 1.x CMs can now be operated with less FEC in channels heavy with ingress and/or CPD, such as channels below 20 MHz which may have previously been considered unusable.

Further, the improved equalization and the lower implementation loss in advanced PHY CMTS hardware means that ingress-free channels, which previously could not support 16 QAM, can now easily support it.

The improvement in mitigation of impulse/ burst noise is less substantial than that of ingress and CPD, but nonetheless exists. For example, a previous paper by one of the authors showed that periodic impulse noise could be tracked and avoided by intelligent scheduling in an advanced CMTS. Note also that most impulse noise has a non-uniform spectral signature. This characteristic may be exploited and significant amounts of the noise energy reduced via the ingress cancellation processing. If either of the latter capabilities are not supported in the advanced PHY CMTS, the benefits of advanced PHY can still be reaped via deployment of data-only service in channels which have high impulse/burst noise. The additional packet loss due to impulse noise can easily be low enough for the most commonly occurring events that the result is a degradation that most users would seldom notice.

Thus, an advanced PHY CMTS can result in current DOCSIS modems operating at 16 QAM, and additional spectrum below 20 MHz being usable by these modems. If data-only modems are moved below 20 MHz, there would be more room at higher frequencies for services requiring higher quality of service or higher channel capacities.

EVOLUTION OF DEPLOYMENT STRATEGIES

In the previous section, the notion of how MSO's could deploy advanced technology was introduced, especially as it relates to the existence of legacy systems on the plant. In this section, a more detailed look at deployment strategies is presented; in particular, the transition from *hardware-limited* capacity to *spectrum-limited* capacity is described. Given the constraint of legacy modems on the plant and the desire to use spectrum below 20 MHz, it will be seen that the logical transition path is from upstream channels with relatively fixed center frequencies, symbol rates, and modulations to channels which can adapt the modulation, center frequency, and symbol rate to the instantaneous conditions of the channel.

Currently, many MSOs have more RF bandwidth available than upstream receivers, as evidenced by node combining and the lack of utilization of all RF upstream spectrum on the plant. Since often the modulation scheme currently in use is QPSK with maximum levels of FEC, the only adaptation strategy available is to hop the cable modem frequency. With the advent of guaranteed pre-equalization in DOCSIS 1.1, 16-QAM may be used for the majority of spectrum. Note that the fact that that the upstream can support high QAM levels is evidenced by successful deployments of ADCs OFDMA based cable telephony system in 32-QAM mode. OFDMA is inherently robust against multipath, however multipath may be mitigated by pre-equalization in TDMA/SCDMA systems.

INCREASING UPSTREAM BANDWIDTH

Already some MSOs have run out of upstream bandwidth due to a variety of factors: low frequency ingress, legacy FDM set top boxes, and bandwidth allocated to other entities, including government agencies and schools.

At the same time, due to the advent of more symmetric services, such as 'Voice Over IP', the demand for upstream bandwidth is beginning to increase. The MSOs are either currently faced with or soon will be facing the choice of improving the usage of the limited upstream resource, or to engage in expensive plant upgrades or node splitting.

As we shall see, the improved modulation techniques of advanced PHY systems opens the potential of optimally using this upstream resource.

Before utilizing advanced PHY, for many MSOs the first step to increasing upstream capacity is node 'decombining'. In systems with lower penetration rates, a common practice has been to combine the upstream signals from multiple node into the same CMTS receiver. As penetration rises, decombining the upstream nodes results in a low cost upstream bandwidth expansion. This allows separate optical nodes to be serviced by separate CMTS channels. To achieve this it is necessary to evolve to higher density CMTS receivers. Ultimately, however, if costs are to be contained it is necessary that the CMTS utilize the spectrum assigned to it in the most efficient manner.

MAXIMUM SPECTRAL EFFICIENCY IN DYNAMIC IMPAIRMENTS

As was stated above, the DOCSIS 1.x and 2.0 PHY 'toolkits' provide the potential for maximum utilization of the upstream channel. However, this cannot be realized without intelligent CMTS sensing, analysis and reaction mechanisms.

In addition, it will be necessary to have spare upstream receivers in the CMTS to take advantage of 'divide and conquer' strategies that will be seen below.

A well-known technique used in DSL (digital subscriber loop) to achieve "Shannon limited" capacity in available channels is to slice up the available bandwidth into microchannels and adjust the modulation parameters to the maximum bandwidth that the microchannel conditions will allow. With the use of spare upstream CMTS receivers, this same technique may be used in a coarser fashion to maximize the capacity of HFC upstream bandwidth (see Figure 5 below).

Dynamic Adaptation

Obviously dynamic adaptation only makes sense if the channel is changing dynamically. That this is true of HFC upstream channels is evidenced by many studies [3]. These studies show that impulse and burst noise is often higher at lunch and dinner times. Ingress is often higher at night, while CPD varies with temperature, humidity and wind due to the source of CPD, cable connectors. (CPD, or Common Path Distortion, is often caused by a diode effect resulting from the corrosion of cable connections.)

These same studies show that the percentage of time that impairments exist on the plant is usually well below 10%. While the new modulation technologies embodied in DOCSIS 2.0 provide increased robustness to tolerate these higher levels of impairments, the increased robustness may be had at the expense of spectral efficiency. On large packets for example, the spectral efficiency can be reduced from the maximum of about 4.7 bits/Hz (64 QAM, minimum FEC) to about 1.2 bits/Hz (QPSK, maximum FEC), a 4x reduction in capacity (see Figure 4 below). On a plant with 12 MHz of bandwidth currently used for upstream data service, the capacity can thus be varied from 14 Mbps to 57 Mbps with DOCSIS 2.0 modems on the plant. Clearly, dynamic adaptation is a key strategy in the optimization of upstream bandwidth.

The variation is slightly less when the modems on the plant are all DOCSIS 1.x, but is still dramatic. The maximum capacity would result from using 16 QAM, which has a spectral efficiency of about 3 bits/Hz. In this case, the plant capacity can be varied from 14 Mbps to about 36 Mbps. Further, dynamic adaptation can also open up new spectrum for service and increase capacity. If 6 MHz below 20 MHz were useable at 16 QAM with advanced CMTS processing most of the day, this would give an additional 18 Mbps of upstream capacity with existing DOCSIS modems. Thus, by upgrading the CMTS to one that uses dynamic adaptation to leverage ingress cancellation, improved equalization, burst acquisition and lower implementation loss, existing DOCSIS 1.x networks can be expanded from 14 Mbps to 54 Mbps, again a four-fold increase in capacity.

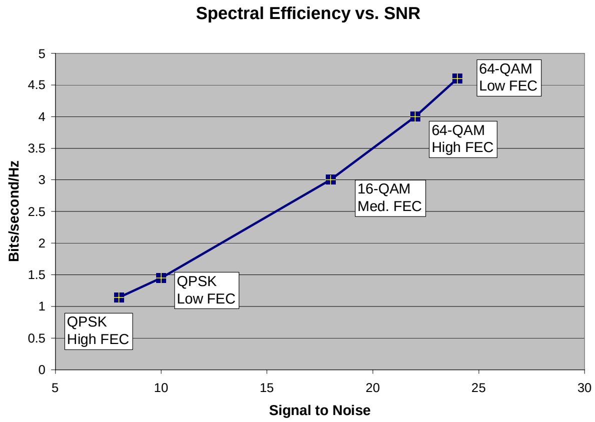

Figure 4. Spectral Efficiency vs. SNR

ADAPTATION TECHNIQUES

Adaptive Modulation

The examples of capacity improvements in the previous section point to simple adaptation techniques. First, the level of FEC used on packet transmissions can be increased as impairments increase. Over 6 dB improvement in robustness is possible with this technique, albeit with a 20-30% drop in spectral efficiency. The other simple technique is to decrease the order of modulation. As previously discussed, with 2.0 modems, 64QAM can be reduced to 16 QAM for 6 dB of improvement in robustness, or all the way down to QPSK for 12 dB of improvement. Taken together, changing FEC and modulation order provide up to 18 dB of improvement in robustness at the expense of 75% of network capacity. As can be seen by the plots of ingress from 24 nodes below, the ingress noise appears to vary from one end of the spectrum to the other by 15 dB to 20 dB, confirming that spectrum may be opened up by using the modulation choices available in advanced PHY.

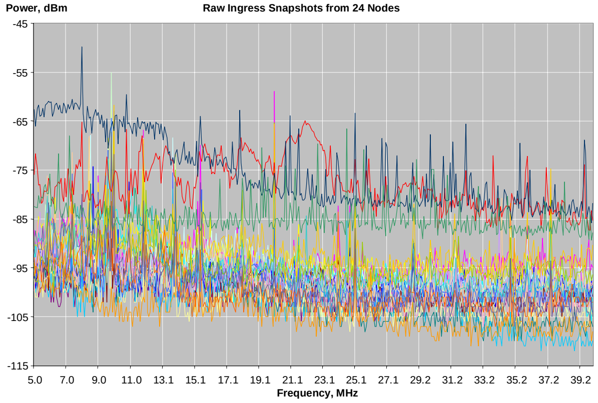

Figure 5. Upstream Plots from 40 Optical Nodes shows Ingress Decreasing by 15 dB to 20 dB at Higher Frequencies.

Channel Hopping

The next level of adaptation involves changing the channel center frequency to avoid significant levels of impairments such as ingress. While the availability of ingress cancellation technologies in the CMTS will reduce the necessity of changing channels much of the time, this adaptation technique remains viable for MSOs with spare RF upstream bandwidth. However, since now only the highest levels of ingress need be avoided, the frequency hop adaptation technique may now involve slight shifts in carrier frequency as opposed to hopping to an entire new block of spectrum on the upstream.

Decreasing Symbol Rate

Next, the symbol rate can be reduced for increased robustness against all types of impairments, again at the cost of reduced capacity. Assuming the transmit power is maintained at the original level, a reduction in symbol rate by a factor of 2 will add 3 dB more robustness against AWGN, ingress, and burst noise. Further, the length of an impulse event that can be corrected is doubled by the fact that in the time domain, the symbols are twice as long as before, therefore fewer symbols are corrupted by the same impulse event. Both aspects of symbol rate reduction are shown in Figure 6.

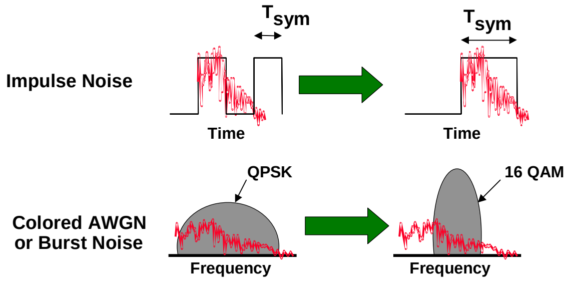

Figure 6. Adaptation via Symbol Rate Reduction.

Thus, for MSOs with spare RF spectrum but without spare upstream receivers, the following adaptation techniques apply:

Increase FEC
Reduce order of modulation
Reduce symbol rate
Change carrier frequency

If, on the other hand, the operator has run out of available bandwidth, and wishes to avoid costly plant upgrades, an effective next step is to employ backup CMTS receivers in the adaptation process. In particular, if the symbol rate is reduced to mitigate impulse/burst noise, the capacity on the plant can be maintained by dividing the channel into smaller subchannels with the same spectral power density. Thus the symbol rate and center frequency must be changed for this adaptation technique. Note that when altering the center frequency of current upstream channels, the following conditions apply:

1) Reranging is generally required.
2) If the order of modulation is already reduced to QPSK, reranging will likely not be required as pre-equalization can be avoided.

The benefits of channel dividing against narrowband and broadband impairments are shown in Figure 7 below, where it is seen that dividing an RF channel which previously could only support QPSK produces subchannels which support much higher orders of modulation. The technique merely requires the availability of backup upstream receivers to optimize the capacity of the network under impaired conditions.

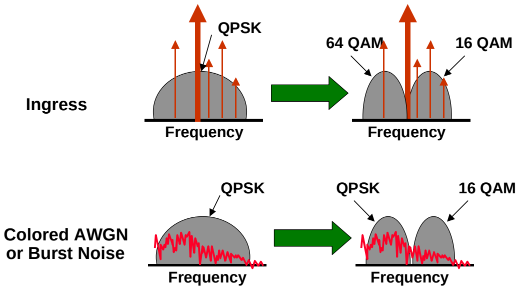

Figure 7. 'Divide and Conquer Strategy' - Channel Dividing to Combat Ingress and Broadband Noise.

Note that special considerations exist for mixed DOCSIS 1.x/2.0 channels. The CMTS must not select an adaptation technique that only works for 2.0 modems, although 2.0 specific techniques can be applied to the 2.0 modems as long as an alternative for the 1.x modems is applied as well. For example, if moderate impulse noise is detected, the CMTS could increase the interleaving on 2.0 modems while maintaining the order of modulation at 64 QAM, and reduce the order of modulation on 1.x modems. Alternatively, the 2.0 modems could switch to SCDMA mode, if impulse conditions warrant. If the impulse noise is too long for simple constellation changes, the symbol rate of all modems on the network may need to be reduced so that the 1.x modems stay active. There are also differences in the equalization capabilities of 1.x and 2.0 modems, and this may also lead to a different adaptation strategy when mixed networks are deployed.

Finally, additional adaptation techniques will likely exist depending on vendor implementations. For example, a smart scheduling alternative to periodic burst noise exists if the noise can be detected and tracked. In this case, the packets from 1.x modems (and the 2.0 modems if necessary) can be scheduled around the impairment without requiring a symbol rate reduction. And the modems may similarly have vendor specific performance and/or adaptation capabilities. Hence, the rules for adaptation should take into account any and all differences in CMTS and modem capabilities.

SYSTEM ADAPTATION

The heart of any scheme to dynamically adapt to changing channel impairments is the ability to detect and classify RF impairments on the upstream. Figure 8 depicts a basic adaptation system, with key components being the spectrum monitor and a lookup table of burst profiles. The spectrum monitor can be internal or external to CMTS, but it is important that RF impairment detection and classification process use rules based on plant measurements and impairment models, such as those presented in [2].

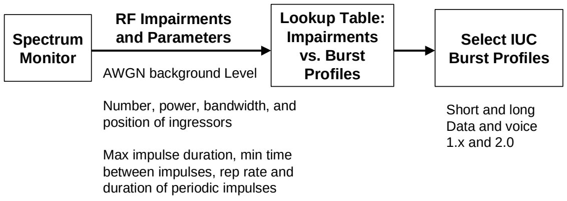
Spectrum Monitor
RF Impairments and Parameters
AWGN background Level
Number, power, bandwidth, and position of ingressors
Max impulse duration, min time between impulses, rep rate and duration of periodic impulses
Lookup Table: Impairments vs. Burst Profiles
Select IUC Burst Profiles
Short and long Data and voice 1.x and 2.0

Figure 8. Adaptation Process.

In particular, the spectrum monitor should classify each impairment separately, since different adaptation strategies exist for different impairments. For example, if the total interference power used to characterize channel, then ingress cancellation and FEC/interleaving will not be leveraged to their fullest extent. Consider the case with an AWGN background noise floor that is 22 dB down from the signal power level, but an ingress signal is present that is 10 dB above the signal power. A 2.0 modem could easily operate at 64 QAM and a 1.x modem could operate at 16 QAM in this level of noise. But if the total interference power were used to characterize the channel, the system would erroneously assume the channel was unusable due to SNR being too low for even QPSK operation.

Further, the spectrum monitor should examine both in-band and out-of-band impairments to be most effective. In the case of a single strong in-band ingress signal that is near the channel edge, a slight shift of center frequency only may be required to keep the channel active and at peak capacity. If the symbol rate is to be reduced without the creation of additional subchannels, the best position for the signal with the reduced symbol rate must be determined. Finally, for impulse/burst noise adaptation, the spectrum monitor should also have the capability to measure impairments in the time domain as well as in the frequency domain.

Once the RF impairments have been detected and classified, the results must be used to determine the burst profiles for the channel that optimize capacity while maintaining sufficient robustness against the impairments. A lookup table is one approach to this requirement, where the system performance is characterized in a lab against a variety of impairments and levels, and optimum burst profiles determined for each impairment and level of impairment. Table 2 shows an example lookup table for AWGN impairments with coarse parameter changes. In reality, a finer table would be desired so that the burst profiles can truly be optimized for the channel conditions that exist.

Table 2. Lookup Table for AWGN

SNR	Modulation	FEC
35 dB	256 QAM	Low[1]
30 dB	256 QAM	Med
25 dB	64 QAM	Low
20 dB	64 QAM	High
15 dB	16 QAM	High
10 dB	QPSK	Med

As the FEC overhead is increased and the modulation type reduced, the spectral efficiency will drop, but for the benefit of greater robustness. The actual FEC used in the burst profile will depend on the packet size, quality of service required, and so on. For example, one set of tables could apply to a packet error rate of less than 1%, while another set of tables could allow error rates of up to 5%. The former could then be applied to voice packets and the latter to best effort data packets. Hence, there could be several lookup tables for each type of service and packet size that optimizes the burst profile subject to the main constraint of tolerating the given level of AWGN with a selected packet error rate.

Similar lookup tables can be developed for each impairment and even combinations of impairments. In this manner, when any previously seen (or postulated) combination of impairments are detected on the cable upstream, the CMTS can use the optimum burst profiles for those particular impairments.

Customization of Algorithms

The algorithms described above, which are embedded in the CMTSs, should allow customization by the MSO. This is due to the fact that HFC plants may differ greatly in the typical ingress signature. Plants with large impulse noise compared to ingress may require different optimization than the reverse situation. In addition, algorithms will need modification over time as experience grows and/or the mix of 1.x/2.0 modems migrates towards all 2.0 modems.

Intelligent CMTS Initialization

How might such a system operate in a global sense? Upon boot-up, the CMTS would characterize the upstream channel using the spectrum monitor. Next, an initial burst profile based on the detected AWGN background level could be selected. Again, ingress and impulse/ burst noise power in the channel must not be used for this decision if it is to be optimal.

Next, if ingress is present and levels are too high for cancellation in the measured AWGN background, a new burst profile can be selected that the ingress canceller can handle. The same process can be used for other impairments such as impulse/burst noise, although some adaptation techniques such as interleaving are fairly independent of the background AWGN and ingress power levels.

CONCLUSIONS AND FUTURE DIRECTIONS

The need for, benefits of, and basic design aspects of an adaptation system for mixed DOCSIS cable modem networks have been described in this paper. It was shown that adaptation can quadruple the bandwidth on the network during the great majority of the day. Adaptation also reduces the capacity during detected impairments but in a manner that keeps spectral efficiency as high as possible for the detected impairments and returns to the highest spectral efficiency when the impairment diminishes.

As MSOs transition from hardware limited capacity to spectrum limited capacity, more complex adaptation schemes will be employed, for example channel dividing. This is

[1] Proprietary 256-QAM Mode

necessary to adapt 1.x modems while maintaining capacity, but as more 2.0 modems are deployed, the need for channel dividing may be reduced. Further improvements in the robustness and capacity of DOCSIS modems can also lead to modifications of the strategies described here.

REFERENCES

[1] Legault, B., “A Technical Analysis of DOCSIS 2.0,” NCTA Proceedings, 2002.

[2] Howard, D. “Detection and classification of RF impairments for higher capacity upstreams using advanced TDMA,” NCTA Proceedings, 2001.

[3] Prodan, R., et al., "Analysis of Two-Way Cable System Transient Impairments," CableLabs®, NCTA Proceedings, 1996 and "Two-Way Cable Television System Characterization", Cable Television Labs, 1995.

ECONOMIC IMPLICATIONS OF AN ADVANCED OPERATIONAL SUPPORT SYSTEM

Bruce Bahlmann
Alopa Networks

Abstract

An Advanced Operational Support System (AOSS) offers Broadband Service Providers (BSPs) the means to activate (provision) multiple broadband services and streamline numerous tasks. This paper will review the benefits of AOSS while quantifying the cost savings, productivity gains, overall economic benefit, and compounded revenue that can be realized from its use.

Introduction

There are least five different areas impacted by AOSS that BSPs should consider when exploring such an investment. They include:

- *Time to Market:* The reduction in the time required to create and deploy new products and services as a result of an AOSS deployment.
- *Market Penetration & Scalability:* The increased installation capacity that results from an AOSS deployment and the ability to scale its performance in step with subscriber growth.
- *Operational Cost Savings:* The reduction in operating costs (e.g. labor, phone support, service calls, etc.) as a result of an AOSS deployment.
- *Future Proofing:* The ability to leverage/reuse the interfaces and functionality to grow the number of services (e.g. data, voice, video, gaming, etc.) supported by the AOSS.
- *Service Assurance:* The increase in service availability and reliability that results from an AOSS deployment.

Let's take a closer look at each area as well as the challenges and benefits they represent to the BSP.

Time to Market:

Too often in business, advances in technology drive product offerings. Companies acquire raw technology along with suggestions or hints from developers, vendors, and the media to create new products that they have been led to believe will sell.

Providing broadband services is one area where product offerings are excessively driven by technological advances. Many of these advances do not come in any order and, more often than not, BSPs seeking to benefit financially from these advances are faced with expensive re-fitting of their infrastructure. BSPs need to focus on meeting the needs of their customers rather than reacting to technological innovation.

An AOSS that is structured around a clear and identifiable business model enables BSPs to quickly bring new products to market. In this way, the business needs drive the technology requirements lessening the reliance on innovation while maximizing the return on existing capital investments. In other words, the business people say, "This is what we need” to the technology people who then go about leveraging existing technology and readily accessible information about their infrastructure to build and maintain the highest quality of service.

Vendors supplying AOSS products address the needs of their customers by offering one of the following solutions:

Tightly Integrated Approach: A system of related components that have been designed (often from the ground up) to all work as a unit and provide superior functionality. These AOSS products offer the foundation for multiple service capability, reduced complexity, lower overall costs, and vast functionality in exchange for a single vendor dependency.

Component-Based Approach: A collection of individual components (or applications) that are assembled according to the needs of the customer to provide the required amount of functionality. These AOSS products permit the flexibility of somewhat interchangeable components (or reduced single vendor dependence) in exchange for higher overall costs (largely due to multi-vendor profitability requirements), more complexity, and less functionality.

While the benefits of the component-based approach are certainly important, the trend of all major AOSS vendors has been to bundle their provisioning applications together. They do this out of the need to offer full-functionality and a tightly integrated solution. Vendors still offering a component-based approach fail to match the functionality of the major vendors, as the best they are able to offer is mean functionality – only that functionality that works across all disparate components.

BSPs who do not seek a single vendor solution to their AOSS needs face a dying breed of vendors attempting to offer full-featured component-based systems. These component-based systems tend to be complex, expensive to maintain, and feature deprived. BSPs using them must take on outside consultants or develop unique integration experience among their employees to manage the complex blend of vendors required by such a system. These BSPs must also take steps to ensure these integration consultants or internal employees stay around for years to come if the resulting solution is to have any kind of shelf life. Troubleshooting between components quickly becomes burdensome and costly. In addition new features require coordination and help with integration among multiple companies resulting in long lead times and costly development. Even with acquired expertise and close relationships with all vendors, these solutions will always lack the functionality of tightly integrated solutions and find the road to new broadband service offerings slow and cumbersome. That is, unless BSPs deploying them become full-fledged development houses creating their own software. However, most BSPs want to keep subscriber and service focused so the distraction of internally designing, developing, and supporting software goes beyond their desired core competency.

If BSPs want to expand their number of service offerings but don’t want to keep purchasing entirely new applications and equipment for each service offering, they need to lay down a framework on which they can build. This framework should not be an individual component of the system that BSPs must add other components to and then continually re-glue them to create new service offerings. Rather, a fully functioning framework should support all the basic aspects of AOSS – perhaps like those described in Data Over Cable Service Interface Specification (DOCSIS). Equally important, but not addressed in DOCSIS, the framework should support interfaces to billing, troubleshooting and trouble ticketing, administration and management, etc. BSPs can use this framework to grow and expand their service offerings by merely reusing and/or expanding their existing AOSS system. Treating AOSS as a system also buys BSPs a lower incremental cost to enter new service offerings.

Deployments of entirely new systems can come with hidden costs such as the potential for impacting existing services that are fully deployed and operational. These deployments may well require lengthy and expensive risk assessments be completed and carefully reviewed before they can proceed. Fork lifting these entirely new systems into place also requires months of preparation, testing, integration, and field trials before they are ready for prime time.

Time to Market is a major consideration for BSPs exploring an AOSS investment. However, quantifying *Time to Market* can be difficult and subjective depending on the presence of any existing AOSS capability. Here are some things to consider when evaluating AOSS and *Time to Market*:

- *AOSS Installation and Deployment:* AOSS vendors vary on the time they require to install and deploy their AOSS. These times could range from 30 days to six months (or more). Since the start time can have a compound effect on your Return on Investment (ROI) moving forward, BSPs should keep this in mind.

- *New Service Rollout:* Non-AOSS supported new service rollout will generally take between six months to a year to complete - this figure could be less if minimal AOSS capability exists. An AOSS supported new service rollout may take between 60-90 days depending on amount of hardware that needs to be configured or added. As a result, AOSS empowers BSPs to offer new services in one-third to one-fourth the time it takes a non-AOSS BSP.

- *New Variation of Service Rollout:* This type of service rollout merely involves creating a new combination of service parameters to meet the needs of a different subscriber population. Non-AOSS BSPs would still take between 60 days and six months to complete. Note that completing this involves a number of organizations (e.g. sales, marketing, engineering, operations, field fulfillment, etc.). An AOSS supported rollout of this nature may take between two weeks and a month to complete with much of the effort focused on aspects other than technical (e.g. organization tasks such as training call centers, field fulfillment, marketing, etc.). As a result, AOSS empowers the BSP to offer a new service variation in one-fourth to one-sixth the time it takes a non-AOSS BSP.

- *Leverage Existing Deployed Hardware:* While this may be implied in the previous bullets on service rollouts, there are economic benefits to being able to simply reconfigure existing components to offer something entirely new to subscribers. The potential savings generated by this capability allows BSPs to prolong their investments in new hardware while potentially creating alternative ways to meet subscriber requirements within the functional means of their existing distribution network and AOSS.

The impact that AOSS can have on a BSP's *Time to Market* can be a very subjective. Perhaps the easiest way to quantify this benefit is that AOSS enabled BSPs will reach their *Time to Market* an average of 13 weeks before non-AOSS enabled BSPs. In addition, the whole process will be cheaper and more precise.

Market Penetration & Scalability:

An ever-present fact in the business of providing new broadband services is the need to increase the number of subscribers. This is typically done through some type of installation process that, depending on the type of service, may include wiring, configuration, and/or provisioning. The size and scope of these activities depends on how well the service(s) was/were designed and built. Many broadband services require as little as a simple activation (which can be easily handled remotely) whereas others require much more, including skilled labor. The initial installation of broadband data as well as other similar services requires all three activities as well as skilled labor. The dependency on manual intervention with each installation has slowed the accumulation of these subscribers.

Seeking efficiencies is paramount when rolling out new services. New services create growth in subscriber installation/activation and require coordinated efforts between marketing, sales, and field fulfillment groups to more rapidly place these services in the hands of subscribers. Without efficiencies, growth of subscribers and profitability in the new services are slow to develop. Out of this need grew something known as "auto-provisioning". The concept of autoprovisioning was simple -- create a way to automatically activate/enroll subscribers on broadband the way the dialup industry does it. However, it did not completely work in practice and thus "Auto-Provisioning" has now branched off into a number of different installation options for broadband operators – each installation option requiring a varying degree of BSP involvement and automation. As a result, a whole spectrum of installation options has evolved for prospective subscribers ranging from totally subscriber driven installs (known as self-install) to the traditional BSP employee driven installs.

BSPs are extremely interested in exploring ways to speed and automate installations. However, many products on the market do not represent a complete AOSS because they have elected to address only certain aspects of the installation process (e.g. automating the computer configuration portion) – few take the approach of trying to streamline the installation as a whole or address multiple installation methods. As a result, the benefit from these products is minimized by the additional need to glue all these systems together to benefit only a single installation option. BSPs stand to substantially increase (between 20-40%) their installation capability by increasing the number of installation options. Augmenting their employee driven efforts to install new subscribers offers BSPs new ways to increase subscriber growth without hiring more employees. Here are some things to consider when evaluating AOSS and *Market Penetration*:

- *Multiple Installation Option Support:* Increasing the number of installation options that BSPs make available to their subscribers distributes the responsibility of installing new subscribers. As a result, BSPs are not solely dependent on the their current number of installers to gain market share. Instead, this chore falls upon the technology (AOSS) as well as multiple BSP subscriber support organizations. By augmenting traditional BSP employee driven installs, BSPs can increase their install capacity by and average of 20 percent. This increase represents the additional install capacity realized through the addition of self-install (contributes a 5% increase in install capacity) and semi-manual installs (contributes a 15% increase in install capacity).

- *Enhanced Employee Driven Installs:* Introducing an AOSS provides a number of efficiencies to a BSP's existing employee driven installs. These efficiencies enable the BSP installer to streamline a number of ordinary tasks including streamlining verification of the subscriber's equipment meeting minimum specifications, provisioning, software installation, and operational checks. The result of these efficiencies results in an increase in install capacity by and average of 20 percent.

- *Packaging Support:* The ability to offer the same set of services over all markets or to have the flexibility to offer customized service offerings to targeted markets. Packaging support offers BSPs the ability to increase demand for services by tailoring services to subscribers, select markets, etc.

- *Compounded Revenue Effect:* The additional revenue generation that results from the increased installation capacity of AOSS that is realized over traditional BSP installation methods. Essentially AOSS creates opportunities to install more subscribers each year over traditional BSP installation methods. This surplus of installation capacity will generate revenue that can be above and beyond what would have other wise been possible using traditional BSP installation methods. This revenue can be compounded and saved or used to fund additional subscriber acquisition enhancement mechanisms.

Closely related to increasing the number of subscribers maintained by the BSP is the capacity of the BSP AOSS to grow/scale with the BSP needs. AOSS *Scalability* is not something that sticks out when evaluating an AOSS and unfortunately there is not any standardized means to measure it or compare one vendor's means of addressing this with another. Instead, it's a process of due diligence that must be completed for each vendor to ensure that their product/solution will indeed be able to support the numbers of subscribers that are being projected.

If the AOSS does not grow/scale with the BSP needs, it will ultimately limit future growth in the various services it supports. Here are some things to consider when evaluating AOSS and *Scalability*:

- *Consider its Composition:* Explore what makes up the AOSS, which parts are homegrown and which parts are third party, and consider the framework of the underlying code.

- *Review the Range of Solutions:* Determine how far each AOSS vendor can grow and determine which pieces get added/replicated/or reused (if any) as the AOSS is called upon to address increasingly larger numbers of subscribers.

The impact that AOSS can have on *Market Penetration* is very well understood. Depending on the capability of the AOSS the BSP should expect to see between a 20-40% increase in their installation capacity. What this means is that by installing an AOSS a BSP can increase their installation output by 20-40% (assuming demand is there) without bringing in additional install personnel.

Operational Cost Savings:

Rarely does a newly installed broadband service meet the highest efficiencies possible. This only happens over time, as various details about a service are better understood. These efficiencies may well be organizational or procedural or both. Other efficiencies will have technical or operational communications dependencies – it is these efficiencies that can be addressed with AOSS.

BSPs who invest in an AOSS find that they are not just buying technology that can help them interface with equipment and other applications – they are also buying expertise. Each AOSS represents an accumulation of best practices obtained through years of exposure to BSPs' products, services, and personnel. These best practices can range from the addition of new technology that streamlines a particular task to complex associations and communications that can connect a number of different tasks eliminating some employee's responsibility.

BSPs who are just beginning to explore an AOSS investment or deploy a new broadband service find that they can be effectively catapulted into greater operational efficiency through deploying a new AOSS. This ultra efficient state allows BSP employees to worry less about technology and more about the most important thing – their subscribers. By streamlining tasks the employees are able to spend more time one-on-one with the subscriber – reinforcing their decision to choose the BSP over some other provider.

Here are some things to consider when evaluating AOSS and *Operation Cost Savings*:

- *Account maintenance:* This is an activity where information about the subscriber's account, devices, service, etc. can be modified. The modification performed includes add, modify, or delete and traditionally this has been done over the phone. AOSS allows a subset of these operations (those frequently asked by customers) to also be completed via the subscriber's personal computer. By augmenting call centers, BSPs can decrease their phone support costs by an average of nine percent.

- *Employee Training:* Since AOSS has the ability to automate a number of tasks associated with BSP field technician's work, the size and scope of the technician's duties at the customer residence can be reduced. In fact, this reduction in duties can even result in a change to the field technician skill level required for new services (e.g. voice and data). It is conceivable that through continued advancements in AOSS that field technicians will require increasingly less specialized skills to install voice and data services. At that point, all installs will appear as standard video installs with the technician performing all the wiring and connections and then technology playing a key role in simplifying and completing the service specific tasks. In this way, BSPs can leverage the AOSS to lower their risk of training employees for their competitors while standardizing their install costs across multiple broadband services. AOSS is well on the way to achieving this goal while reducing the labor costs for installs.

- *Subscriber Education:* One of the keys to offloading calls into BSP call centers is a better-educated subscriber. Subscribers constantly seek information about changes to the service, service availability/performance, and how to do various tasks. AOSS can help augment BSP employee driven efforts to provide this information to subscribers by placing this information on line as well as on the subscriber's computer. With this information readily available a good percentage of subscribers can resolve problems independently. As a result, BSPs may see a reduction in call center information requests and potentially fewer service calls depending on how aggressively this area is approached by the BSP.

- *Service Calls:* BSPs track several service specific indicators that help them gauge how well each of their organizations are doing. One area of particular interest is service calls to new installs that happen within 30 days of the install. This particular service call is very costly to the BSP as it typically represents a poor quality initial install and can be detrimental to the subscriber's confidence in the overall service quality. While the reasons for these service calls may vary, they are due in large part to the lack of standardization of the install process. Each install is unique and handled by any one of several field technicians. Although all field technicians receive the same training, tools, and install the same service, their finished product (a completed install) varies from one installer to anther. These variances in the completeness and quality of the overall install occasionally show up when a service call is requested from a newly installed subscriber. When this happens the original installer may be advised on what (if anything) they did wrong. However, this feedback loop, that can include additional training for the field technician, is not a sure fix. The AOSS can help standardize installs by automating various steps the field technician takes regarding subscriber computer installation and configuration to ensure that all necessary steps are completed. If the field technician uses such a tool, their BSP employer can be assured that a portion of the install was completed the same as any other subscriber's computer. *Note: the wiring will remain a task that varies from one field technician and subscriber dwelling to another*. By standardizing subscriber installs, BSPs can decrease their service calls to newly installed subscribers by an average of 18 percent. *Also note that*

> *as service calls are reduced due to AOSS capability (which results in increased revenue and decreased costs) it creates a compound revenue impact.*

The impact that AOSS can have on *Operational Cost Savings* is mixed. Some aspects of the projected savings have been well researched and their numbers confirmed while other areas are still subjective. Depending on the capability of the AOSS as well as how aggressively the BSP pursues these efficiencies, the BSPs that forge ahead in this area should expect to see at least a 20% reduction in their overall support costs.

Future Proofing:

Investing in an AOSS should not be something done for a single service rollout. Rather, it is an act of planning for the future. Essentially, AOSS lays down a framework on which one can build a number of services. Similar to DOCSIS that builds upon former releases; so should the AOSS in providing an increasing number of services that effectively reuse as many existing components as possible. It represents the thinking of interfacing with all the necessary BSP service components (e.g. trouble ticketing, network operations, subscriber management systems, billing systems, network/hardware appliances, etc.) and then builds on top of these interfaces a strong, scalable, and reliable base that will support any number of additional services and functions.

To *Future Proof* means to create something that can endure and assimilate with changing BSP needs over time. This area is subjective only to the point that the BSP finds that offering a new service requires a completely new system. A completely new system involves purchasing and installing new hardware and software to run the system as well as all the other necessary components to operate it – network management software, troubleshooting software, training, billing interface software, etc. Since these would be entirely new applications, the impact on training as well as equipment/rack space would be significant. Instead, if this only required a new module of an existing, completely implemented system, the impact of change would be "manageable". All the training would be within an environment that is already familiar to the employees. Additionally, new services would generally not represent additional hardware but rather reuse that which is already deployed. One other aspect of trying to install a number of independent systems for each new service is that it places varied dependencies on existing applications maintained by the BSP (e.g. network management, trouble ticketing, billing, etc.). So if any of these systems became outdated, replacing the system might break one or more of the services dependent on the interface. Having fewer interfaces into these BSP applications reduces the risk that something will break during a transition.

The impact that AOSS can have on *Future Proofing* is subjective but only to the point where offering multiple services becomes a priority to the BSP. If the BSP does not feel a need to offer multiple services the need for an AOSS is minimal. However, since this is a priority for most BSPs, there is a need to justify the expense of investing in an AOSS in terms of the overall savings it could generate as well as the number of broadband services it can deliver.

Service Assurance:

Running a broadband service requires a high degree of technical expertise and most importantly -- consistency. As BSPs provide service to an increasing number of subscriber dwellings, the pressure to maintain initial service quality and reliability requires substantial attention of BSP technical staff. The case where the BSP delivers an "always on" broadband service provides the most challenging aspect of maintaining performance and scalability.

The introduction of AOSS within this environment provides opportunities to tie previously autonomous systems together. The resulting tightly integrated system provides real-time information about various network devices with a particular focus on those directly used by subscribers (modem, media terminal adapter, set top box, computer, residential gateway, etc.). Unlike network management systems that oversee the health of fixed assets (such as routers, servers, etc.) within a network, the AOSS focuses its attention on the health of subscriber specific devices as well as the communications between network elements used by the AOSS. Subscriber specific devices fall out of the space network management systems can comfortably observe as they can change over time and the network management system is not the authoritative source for this information.

The impact that AOSS can have on *Service Assurance* varies between AOSS vendors and can be very subjective depending on the functionality provided. While it makes sense that BSP subscribers want to know when and why their service is unavailable, it is very difficult to provide a figure that justifies the savings a BSP should expect as a result of deploying any given AOSS. Regardless, the BSP should expect that an AOSS with built-in *Service Assurance* should provide operational efficiencies that extend beyond the capabilities of their network management system.

Conclusions:

Exploring the need for an AOSS is not just an exercise in determining what services are next in line to deploy and what components are available to assist in this deployment. AOSS is something that, after it is deployed, continues to grow along side the demand/need for additional services and variations of service.

The economic implications of deploying an AOSS reach areas of operations, field fulfillment, call centers, and even engineering. While some of the benefits of AOSS remain subjective, an increasing number of them have been thoroughly researched and verified in the field. In fact, a number of these benefits have been incorporated into ROI models that AOSS vendors are now circulating to BSPs in order to help them more fully comprehend the impact of introducing an AOSS.

Contact Information:
Bruce Bahlmann
Dir, Technical Market Development
Alopa Networks
bruce@alopa.com

Special thanks to Steve Grzelak of Alopa Networks for contributing some of the material in this paper.

EFFICIENT SERVER DISTRIBUTION IN ALL DEMAND SYSTEMS

Doug Makofka, Bob Mack
Motorola BCS

Abstract

AllDemand systems are systems that only carry services on the plant when some client demands (tunes) them. VOD (sub)systems are the closest well-known analogy to the AllDemand system. Although AllDemand systems are similar to VOD systems, there are differences between the two. These differences require an AllDemand system designer to consider several potential architecture models in order to assure efficient media server distribution.

When considering server distribution in AllDemand service environments, one must take into account the types of AllDemand services that will be offered as well as the bandwidth available at different points in the network. This paper provides an overview to help in the selection of a trade-off of AllDemand system facets. A summary of AllDemand service types, and architectures is provided.

AllDemand SERVICE TYPES

In addition to the standard VOD, SVOD type services that exist in current systems, new service types are possible in AllDemand systems.

Broadcast AllDemand

Broadcast AllDemand services save plant bandwidth by only carrying a service on the plant if some set-top device has actually tuned the service. Once the service is tuned, it is placed on the plant. Every client then accesses (tunes) the same copy of the service. There is no motion control or time shifting of these services.

AllCached

AllCached services break the service schedule paradigm by allowing the subscriber to watch programs on an as-desired basis over some time period. All services delivered as AllCached are cached (surprise!). Tuning a service causes it to be played from the caching device to the viewer.

Network PVR

Network PVR can be approached two ways. One can think of providing a virtual disk in the network for use by the client. One can also think of it as adding motion control, and longer-term storage to the AllDemand service type.

Points of Note:

1. Broadcast AllDemand saves plant bandwidth. Other service types require more plant bandwidth;
2. The first ‘tuner’ of a Broadcast AllDemand service experiences two-way communications latency. Other service types carry two-way communications latency with every tune;
3. Broadcast AllDemand services are shared between subscribers. Other AllDemand services are served uniquely to each subscriber;
4. Local Insertion into cached material (e.g., locally-sourced advertising) must be accounted for in each of the service types, although it is not directly analyzed here.

ALLDEMAND ARCHITECTURE TYPES

Centralized Architecture

Figure 1 provides schemas for centralized architectures. The media servers are co-located in the head-end, and their content is distributed from the head-end to remote hubs. The re-multiplexing, encryption, QAM modulation and up-conversion equipment can be co-located with the media servers or located within a remote head-end. Centrally located equipment that requires skilled technicians to operate and maintain is a major benefit of a centralized architecture. The downside of the centralized architecture is the need for high bandwidth, high reliability, and low delay variation (i.e. jitter) connections to remote hubs. Usually, these kinds of networks are expensive; however, many of the larger MSOs have already deployed these high capacity networks. Customers with SONET rings with data capacities to OC-48 or OC-192 are not uncommon. Dense Wave Division Multiplexing (DWDM) over dark fiber is also gaining favor. This is because it is comparatively less expensive to deploy than SONET or ATM networks.

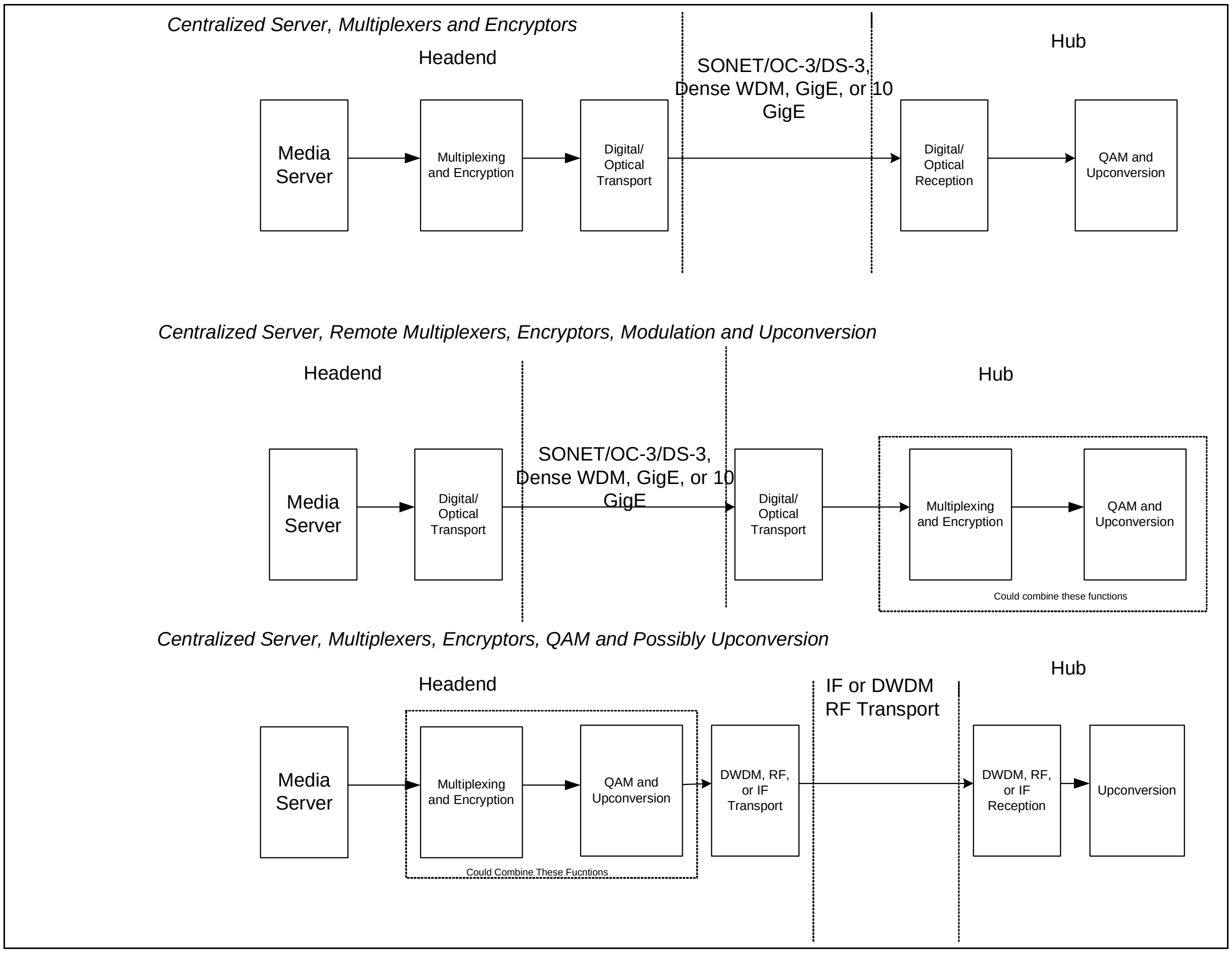
Centralized Server, Multiplexers and Encryptors
Headend
SONET/OC-3/DS-3, Dense WDM, GigE, or 10 GigE
Hub
Media Server
Multiplexing and Encryption
Digital/ Optical Transport
Digital/ Optical Reception
QAM and Upconversion
Centralized Server, Remote Multiplexers, Encryptors, Modulation and Upconversion
Headend
Hub
Media Server
Digital/ Optical Transport
SONET/OC-3/DS-3, Dense WDM, GigE, or 10 GigE
Digital/ Optical Transport
Multiplexing and Encryption
QAM and Upconversion
Could combine these functions
Centralized Server, Multiplexers, Encryptors, QAM and Possibly Upconversion
Headend
IF or DWDM RF Transport
Hub
Media Server
Multiplexing and Encryption
QAM and Upconversion
Could Combine These Fucntions
DWDM, RF, or IF Transport
DWDM, RF, or IF Reception
Upconversion

Figure 1 - Centralized Architecture Schemas

Physical Transport Type Comparison

Table 1 and Table 2 compare the number of physical transport channels required when distributing content using GigE, 10 GigE, ASI, and DS3. Table 1 is based upon a digital penetration rate of 60% with an on-demand usage rate of up to 80%. Table 2 is based upon a digital penetration rate of 30% with an on-demand usage rate of 30%. In each case, it was assumed that each user would require 3.75 Mbps of capacity for their on-demand service. Table 3 provides nominal assumptions about media server characteristics. With actual media server values, the media server resource requirements for a give system can be sized to the degree that it is the number of simultaneous streams that drive media server needs, rather than aggregate storage requirements.

Homes	Digital	Simultaneous	Number of	Number of	Capacity	Capacity	OC-192s	SEMs	SEMs	SEMs
Connected	Connects	Users	(Gig Es)	10 GigEs	(ASIs) [2]	DS3s [1]	Required	w/ GigE	w/ ASI	w/DS3
5000	3000	2400	15	2	58	232	2	15	15	58
10000	6000	4800	29	3	116	464	3	29	29	116
15000	9000	7200	44	5	174	696	4	44	44	174
20000	12000	9600	58	6	232	928	5	58	58	232
25000	15000	12000	73	8	290	1160	7	73	73	290
30000	18000	14400	87	9	348	1392	8	87	87	348
40000	24000	19200	116	12	464	1856	10	116	116	464
50000	30000	24000	145	15	580	2320	13	145	145	580
60000	36000	28800	174	18	696	2784	15	174	174	696
70000	42000	33600	203	21	812	3248	17	203	203	812
100000	60000	48000	290	29	1160	4640	25	290	290	1160
250000	150000	120000	725	73	2900	11598	61	725	725	2900
500000	300000	240000	1450	145	5799	23196	121	1450	1450	5799

Table 1 - Comparison of Transport Types (60% Digital Penetration, 80% Usage)

Homes	Digital	Simultaneous	Number of	Number of	Capacity	Capacity	OC-192s	SEMs	SEMs	SEMs
Connected	Connects	Users	(Gig Es)	10 GigEs	(ASIs) [2]	DS3s [1]	Required	w/ GigE	w/ ASI	w/DS3
5000	1500	450	3	1	11	44	1	3	3	11
10000	3000	900	6	1	22	87	1	6	6	22
15000	4500	1350	9	1	33	131	1	9	9	33
20000	6000	1800	11	2	44	174	1	11	11	44
25000	7500	2250	14	2	55	218	2	14	14	55
30000	9000	2700	17	2	66	261	2	17	17	66
40000	12000	3600	22	3	87	348	2	22	22	87
50000	15000	4500	28	3	109	435	3	28	28	109
60000	18000	5400	33	4	131	522	3	33	33	131
70000	21000	6300	39	4	153	609	4	39	39	153
100000	30000	9000	55	6	218	870	5	55	55	218
250000	75000	22500	136	14	544	2175	12	136	136	544
500000	150000	45000	272	28	1088	4350	23	272	272	1088

Table 2 Comparison of Transport Types (30% Digital Penetration, 30% Usage)

The Gigabit Ethernet solution requires the fewest number of physical interfaces. ASI requires four times as many interfaces when compared to GigE, and DS3 requires a 14 to 16-fold increase in the number of interfaces as compared to GigE. GigE is also considerably less expensive than either DS3 or ASI.

Parameter	Value	Notes
Interfaces	Up to 5 Gigabit Ethernets Up to 8 ASIs	
Transport Streams	Up to 16 Transport Streams/Gigabit Ethernet Output Up to 4 Transport Streams per ASI	Each TS is a 38.8 Mbps multiplex
Stream Rate	3.75 Mbps	

Table 3 – Media Server Assumptions

Building Blocks

AllDemand systems require different building blocks than 'broadcast-based + VOD' systems. The needs of AllDemand services require that the dense mux/mod chain packaging used in the VOD environment be extended to the entire system. Real-time encryption devices are needed to protect content served from the media caches with encryption that is associated with the access control scheme of the client set-top devices. Also, depending upon the scale of the system, switching between the media servers and the encrypt/mux/mod chain may be needed. Switching allows the media servers to scale naturally, rather than having to devote separate media servers to specific encrypt/mux/mod chain.

The centralized architecture schemas are fleshed out below. In current VOD environments the mux/mod chain is often integrated tightly with the VOD server. In the AllDemand system, the tight integration may lead to less than optimal media server scaling. In the AllDemand system, encrypt/mux/mod chain is split out. It can then be used for any service type in the AllDemand system (Broadcast AllDemand, VOD, cached, Network PVR etc), and it can scale with the transport needs of the network itself. A description of the building blocks used in building the following architecture diagrams follows:

SEM - Super Encryptor Modulator

The Super Encryptor Modulator (SEM) combines session-oriented encryption multiplex creation and QAM modulation in a single device. This gives the functional density needed to support the number of streams required in an AllDemand system. For this analysis, the SEM is assumed to support 16 QAM outputs (the number of outputs supported by an actual SEM may be different). Multiplex management is included in the SEM.

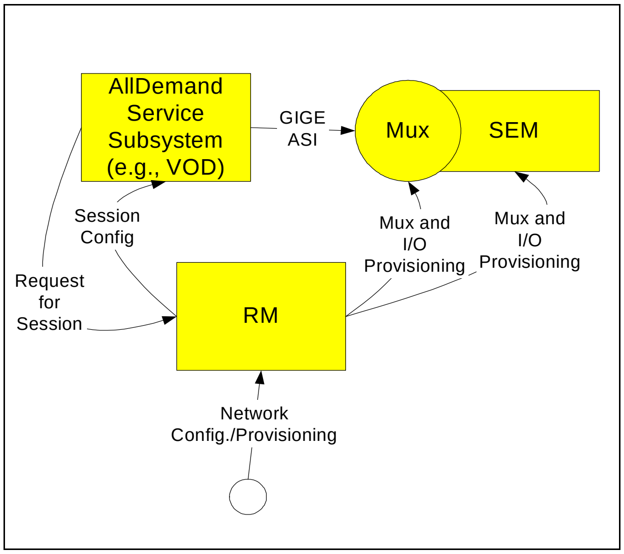

Figure 2

Resource Manager (RSM or RM)

The Resource Manager (RSM or RM) controls the SEM. It also performs bandwidth management. It provides an interface to the Service subsystems that allows sessions to be established between media servers and client devices.

Figure 2 shows a high-level view of the flows between a typical AllDemand service subsystem, the RM, and the SEM. The Mux circle in front of the SEM is normally part of the SEM. In some environments it may be desirable to have a separate Mux block.

DWDM Transport Carry Options

DWDM technology can be used to carry baseband digital, IF or RF modulated signals over the transport network. This provides flexibility in selecting locations for equipment in AllDemand systems. Of course other transport types can be combined to yield the same effect.

The following figures show architectures transporting different signals over DWDM. These diagrams also show different options tradeoffs in equipment locations. Fig. 3 shows GIGE over DWDM, with some level of switching available in the remote hub. Fig. 4 shows IF over DWDM, with just an optical to IF conversion and up conversion in the remote hub. Fig. 5 shows ASI over DWDM, with QAM modulation and up conversion in the remote hub.

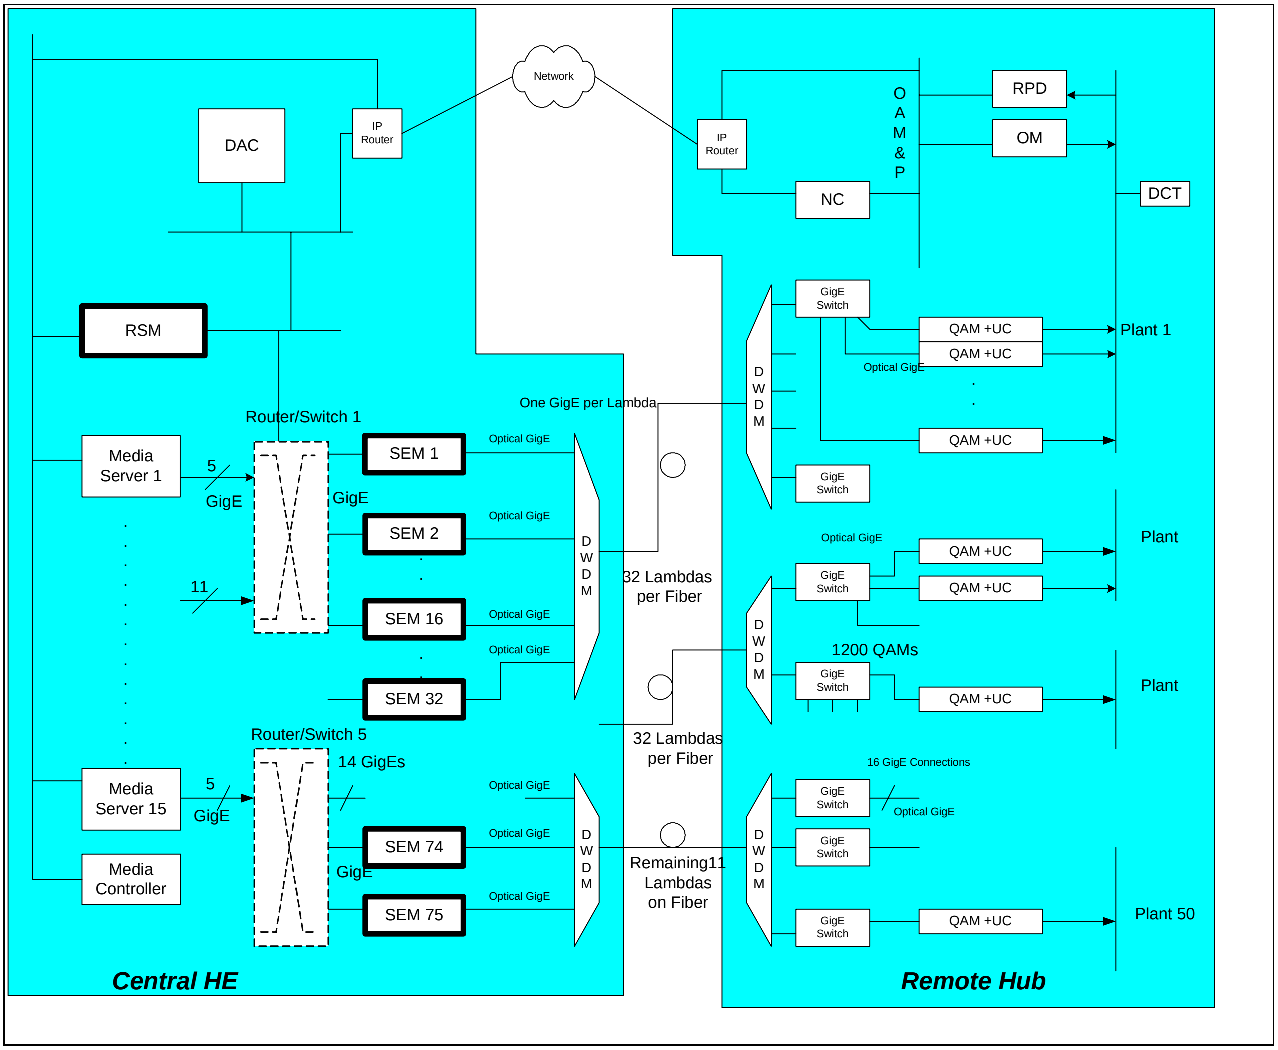

Figure 3 - Centralized Multiplexing, Encryption, and Gigabit Ethernet Transport

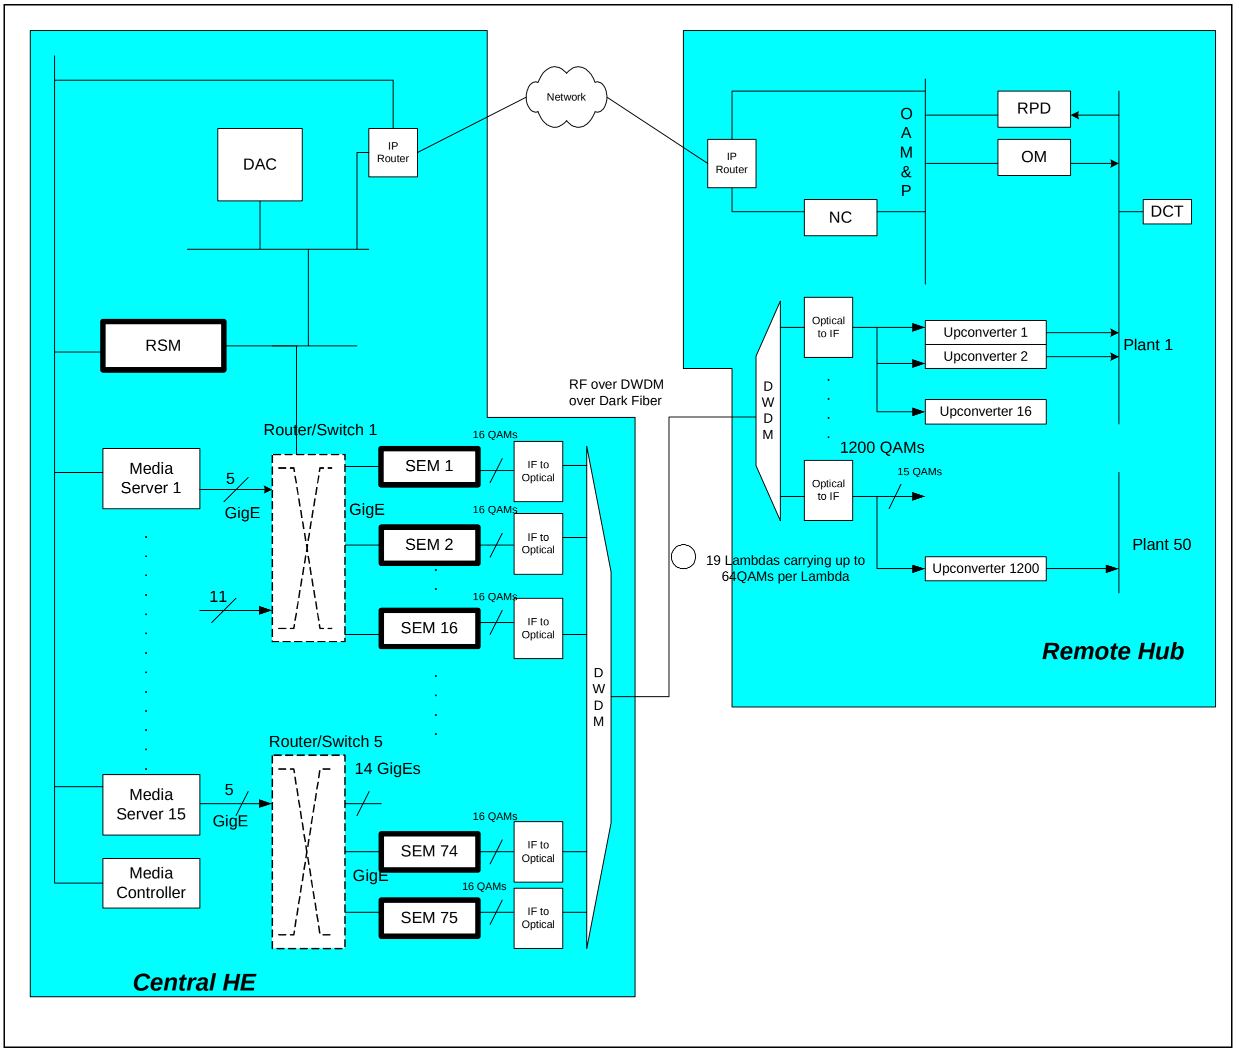

Figure 4 - Centralized Multiplexing, Encryption and QAM – IF over DWDM

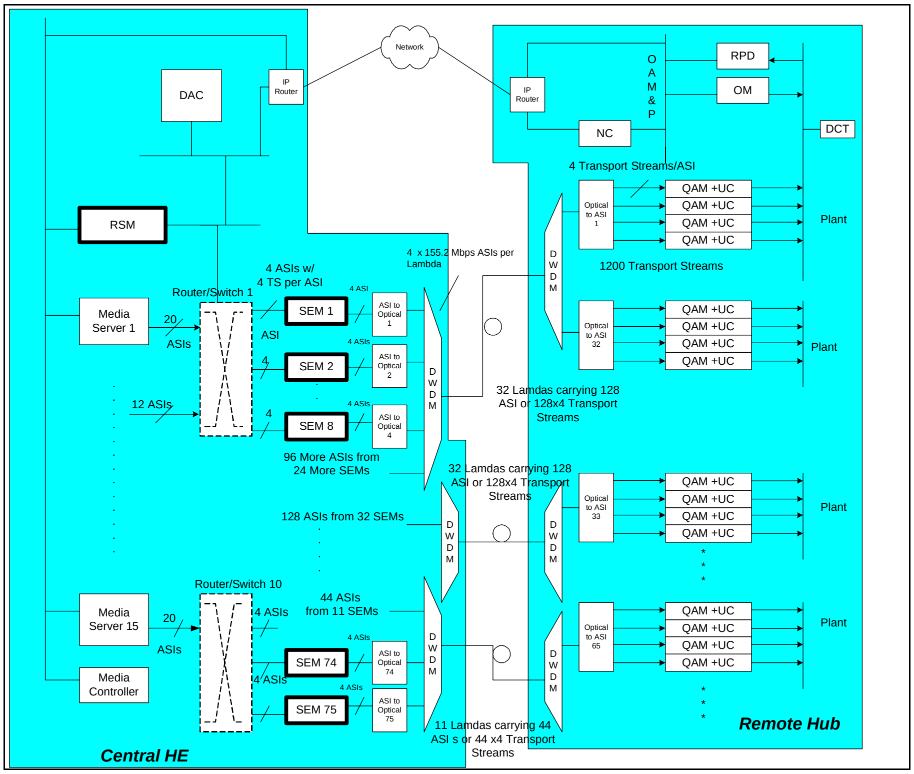

Figure 5 - Remote QAM Modulation – ASI over DWDM

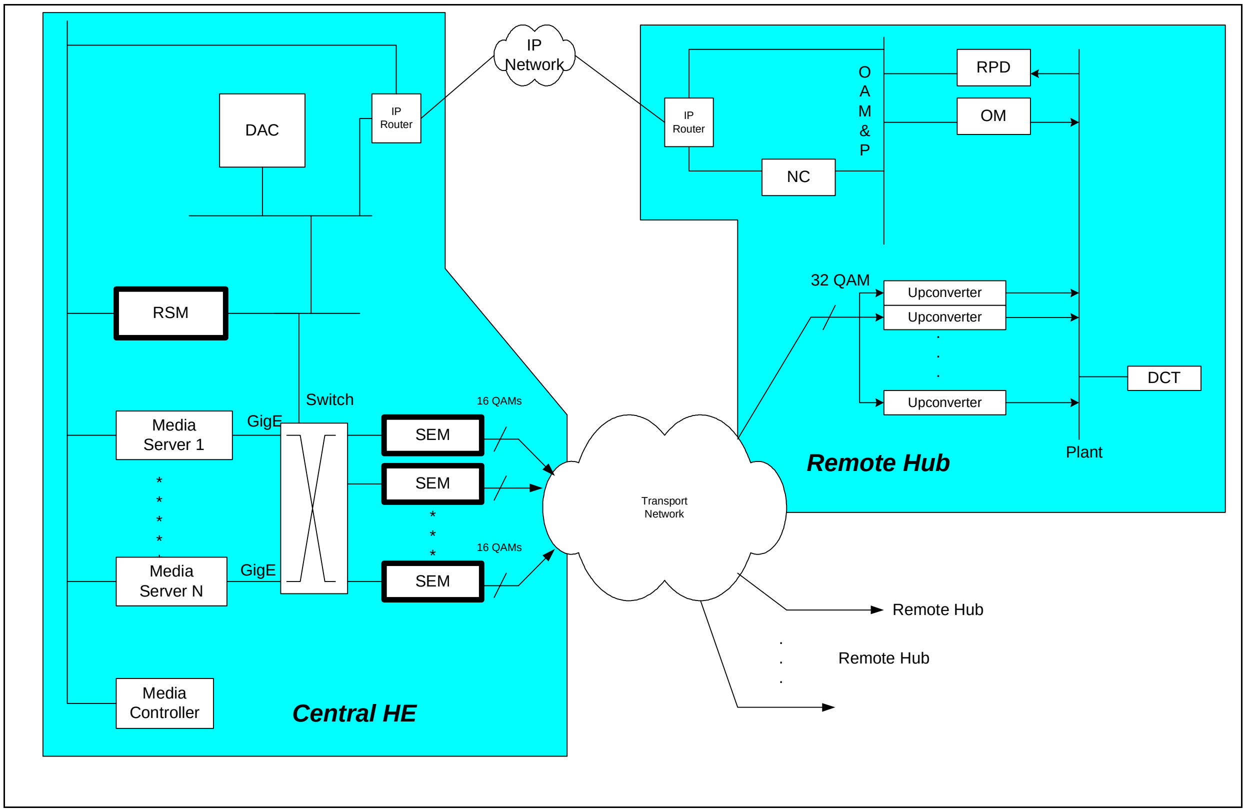

Figure 6 - Centralized AllDemand Nominal Architecture

Figure 6 shows the nominal Centralized AllDemand system architecture. This architecture supports the most flexible, efficient use of media server resources. It requires a high-capacity transport network, but it gains ease of support as well as low support overhead.

NON-CENTRALIZED ARCHITECTURE TYPES

Other architecture types are presented here for completeness. Although there are clear advantages to the centralized approach, there may be reasons to use a distributed or routed approach to the AllDemand system. These are summarized in Figures 7, 8, and 9.

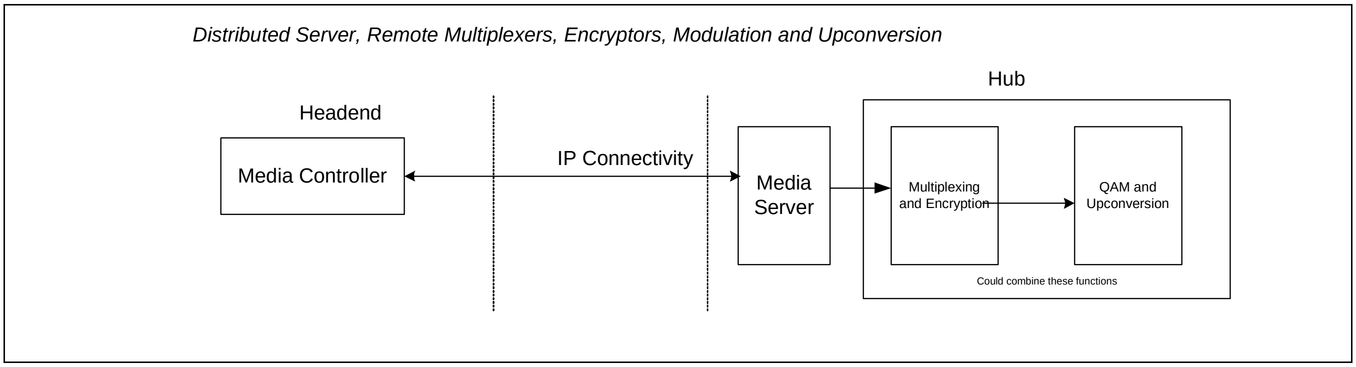
Distributed Server, Remote Multiplexers, Encryptors, Modulation and Upconversion
Headend
Media Controller
IP Connectivity
Media Server
Hub
Multiplexing and Encryption
QAM and Upconversion
Could combine these functions

Figure 7 - Distributed Architecture Schema

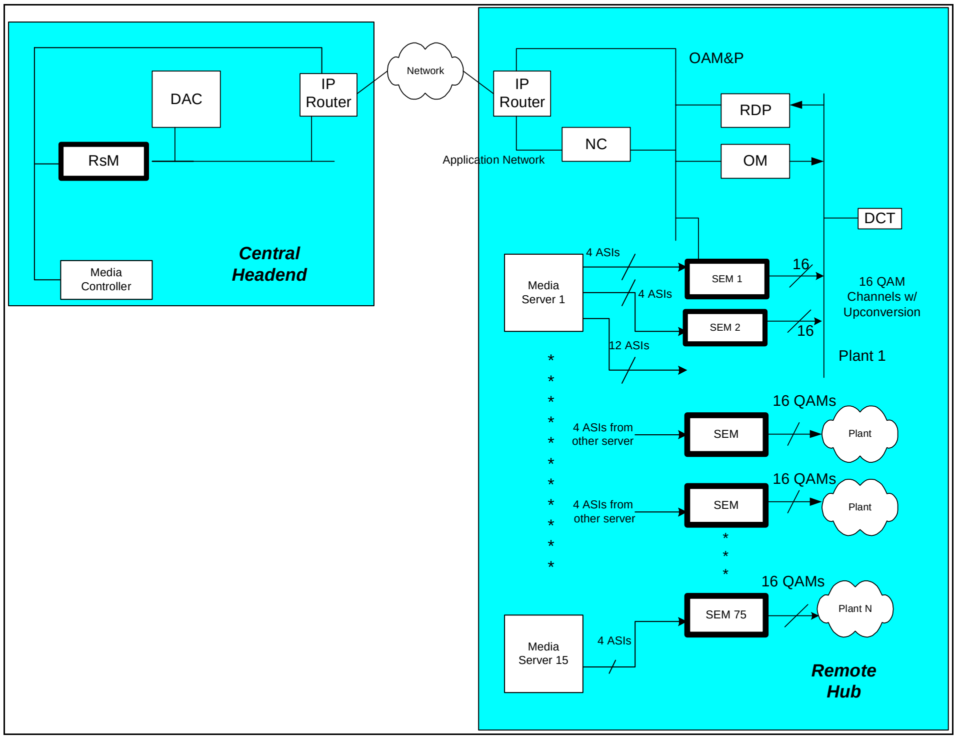

Figure 8 – Distributed System

Distributed Architecture

In the distributed architecture, the media servers are placed in the remote hubs. This may be necessary in systems that lack the transport resources to distribute streams from a central head-end. Some remote, distributed servers may also find themselves in centralized system to support regional insertion activities that are actually managed on a regional level. In general, unless the media server resources exactly match the needs of the local hub population, this scheme cannot be as efficient as the centralized approach:

- Media will need to be distributed (duplicated) in remote environments;
- On-going service, support, provisioning, and maintenance will need to be done at remote sites;
- 'Slack' resources to handle demand peeks cannot be broadly shared. They can only address a local population.

Routed Architecture

If the media servers support a routable output protocol (e.g., GIGE), then the output of the media servers can be routed to encrypt/mux/mod chains in remote hubs. This gives the advantages of a centralized system, but:

- Requires a transport network with QoS support at a higher protocol level (IP, or GIGE);
- Requires routing/switching hardware;
- Extends the provisioning/network management network for the encrypt/mux/mod chain to the remote hubs.

The routed architecture may be desirable in systems that are built primarily to serve IP data and media distribution over broadband environments. Systems built primarily to deliver Voice over IP (VoIP) and/or internet streaming media may already be designed in this fashion.

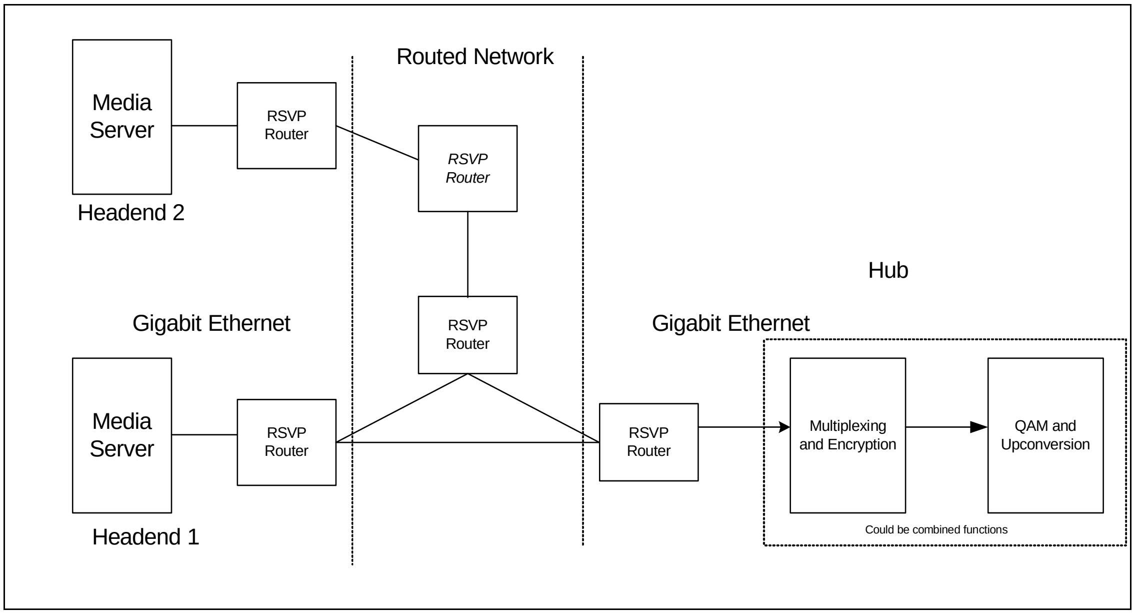

Figure 9 - Routed Architecture Schema

Summary

There are three categories of AllDemand services. Only BroadcastOnDemand saves plant resources. The rest demand more simultaneous services to the subscriber – culminating ultimately with a session per subscriber.

Network architecture drives the efficient distribution of media server resources in AllDemand systems. The GIGE protocol carried over DWDM enables a large reduction in the number of channels needed to transport fully formed streams to remote hubs. These transport improvements allow the AllDemand system to be centrally located in the head-end. Encrypt/mux/mod chains should be loosely coupled to the media servers. This allows the media servers to scale naturally – with total media storage and total simultaneous stream requirements, as opposed to being primarily partitioned by hub/transport architecture.

IMPLEMENTING THE NCTA-CEA PSIP AGREEMENT

Mark Corl, Glen Myers, Nandhu Nandhakumar, Jian Shen, and Gomer Thomas
Triveni Digital, Inc.
Princeton Junction, New Jersey

ABSTRACT

In ATSC-compliant DTV broadcasts, PSIP data provide DTV receivers with tuning information, electronic program guide data, as well as other information supporting a variety of functions. When the DTV broadcasting signal is received by a cable provider, multiplexed and carried on a cable network, the PSIP data need to be preserved and transformed to comply with cable standards. Otherwise, cable-ready DTV receivers may not be able to tune to the signal. In light of this need, the NCTA and CEA have reached an agreement regarding the carriage of PSIP on cable. This paper explores the issues surrounding the adaptation of PSIP data to the cable environment, including the consequences of merging multiple terrestrial transport streams into a single cable transport stream. The paper also presents a sample design and implementation of a system to support PSIP carriage on cable as highlighted by the NCTA-CEA agreement.

INTRODUCTION

In DTV broadcasts, a single MPEG-2-compliant transport stream (TS) can simultaneously carry multiple video, audio, and/or data programs [1]. In ATSC-compliant broadcasting, metadata describing the contents of the transport stream multiplex are carried by PSIP – the Program and System Information Protocol [2]. The PSIP data consist of collections of “table sections”, which are also encapsulated into MPEG-2 TS packets. These packets have unique PID values that can be used to distinguish the PSIP tables from each other, as well as from the audio, video, and data streams.

The PSIP tables are defined in ATSC standard A/65A [2]. These tables enable a number of important features for digital television (DTV) receivers:

- Tuning to programs by virtual channel numbers, rather than physical broadcast bands;
- Selecting language tracks;
- Creating interactive electronic program guides;
- Applying “V-chip” restrictions on viewing based on content advisory ratings.

The ability for a DTV receiver to tune to the terrestrial broadcasting signal based on virtual channel number rather than physical frequency band is very important for terrestrial broadcasters. Network broadcasters typically have one frequency band for broadcasting an analog signal and a different frequency band for broadcasting a digital signal. They have often invested substantial resources over the years in brand recognition for their analog channel number. Moreover, a DTV broadcast may consist of multiple programs, so it is necessary in many cases to identify multiple “virtual channels” within a single digital broadcast band.

The PSIP standards identify virtual channels (VCs) by the combination of “major channel number” (mandated to be equal to the analog channel number for stations with existing NTSC licenses) and “minor channel number.” . The virtual channel numbers allow DTV receivers to tune the DTV signal using the same analog channel number, even though the physical channel is at a different frequency. In addition, PSIP data provide details about alternative language tracks, content advisory ratings, and audio and video streams within the broadcast bands.

PSIP data also support interactive electronic program guides (EPGs) in standard, off-the-shelf TV sets and set-top boxes (STBs) by supplying information on upcoming programs. The standards and mechanisms described thus far allow the user to purchase a standard DTV receiver, hook it up to an antenna, turn it on, and have it work over-the-air with a minimum of manual configuration.

When a DTV signal is carried on cable to the consumers, the PSIP information needs to be preserved in order for consumers to retain the ability to tune to the signal using a cable-ready DTV receiver, analogous to using an antenna. For analog signals, consumers can purchase "cable-ready" TVs that can receive signals from an already-established cable connection. The same holds true for digital television; consumers will expect equivalent commercially available "cable-ready" DTVs or STBs that they can deploy without direct intervention from their cable provider. This is especially important for those consumers that only want the free or low-cost, entry-level "antenna extension" service to view the unscrambled DTV programs. These services do not require the use of the Point of Deployment security module (POD). These consumers will not be able to access the out-of-band channel navigation information and must rely on the in-band PSIP information for tuning and EPG services.

NCTA-CEA CABLE PSIP CARRIAGE AGREEMENT

With the intention of achieving this type of environment, the FCC issued a Report and Order on Cable Carriage of DTV (Docket 98-120), which in paragraph #83 requires carriage (if present) of PSIP data related to the primary video service.

To meet the needs of PSIP carriage on cable and to make the data meaningful for cable-ready receivers, an agreement has been established between the Consumer Electronics Association (CEA) and the National Cable Television Association (NCTA). The agreement describes the carriage of PSIP on cable in support of consumer digital receiving devices (digital receivers) connected directly to the cable TV system [4]. The intent is that consumers will be able to purchase a cable-ready receiver that can process unscrambled cable channels immediately, as well as play "host" to a cable-system specific decryption processor referred to as a POD module.

Key provisions of the NCTA-CEA agreement include:

- If a digital transport stream includes services carried in-the-clear, that transport stream must include VC data in-band in the form of ATSC A/65 (PSIP), if present in the stream.
- VC table data are also sent out-of-band to the POD module in the receiver.
- VCs are identified by a one- or two-part channel number, and a textual channel name.
- At least twelve hours of PSIP event data must be included, if received from the broadcaster.
- The cable provider has the option to limit the total bandwidth for PSIP data to 80 Kbps for a 27 Mbps multiplex and 115 Kbps for a 38.8 Mbps multiplex.
- Event data may be transported in-band and/or out-of band. The in-band data may be used to augment out-of-band data at the receiver.
- For access-controlled services, the out-of-band SI channel number may or may not match the channel number identified with in-band PSIP data.
- The channel number identified with out-of-band SI data should match the channel number identified with in-band PSIP data, for in-the-clear services.

The NCTA-CEA agreement does not preclude the possibility that cable providers may choose to implement alternative agreements with specific stations, station groups, or networks. These "private agreements" could have cable providers including different amounts of PSIP data from different terrestrial/broadcast sources, with different types of manipulations and bandwidth limitations allowed for each.

One implementation challenge from the PSIP agreement involves the handling of PSIP data when multiple transport streams containing PSIP data are multiplexed in cable plants. Terrestrial DTV channels have fixed bandwidth of approximately 19.39 Mbps, but digital cable systems use modulation methods that allow carriage of 27 Mbps or even 38.8 Mbps per transport stream. In addition, cable providers may want to select which programs (or VCs) from a given transport stream to carry. In order to optimize the usage of cable bandwidth, the cable providers naturally want to combine programs from several terrestrial transport streams together in one multiplex.

However, traditional MPEG-2 multiplexers are not designed to handle the PSIP data. Thus, a new PSIP-aware device needs to be developed that can process the PSIP data from the input transport streams and generate the in-band PSIP and out-of-band service information (OOB SI) using the input PSIP data according to SCTE standards.

DESIGN OF METADATA PROCESSING SYSTEM FOR CABLE HEAD END

Comparison of Terrestrial PSIP, Cable In-Band PSIP, and Cable Out-of-Band SI

The terrestrial PSIP metadata consists of a number of tables:

Table 1. Terrestrial PSIP Tables

Table	Description
Master Guide Table (MGT)	gives PIDs, sizes, and version numbers for all other PSIP tables.
System Time Table (STT)	gives current time, convertible to wall clock time at receiver.
Rating Region Table (RRT)	describes system(s) for rating broadcast content, referenced by "content advisory descriptors" in EITs (not sent when RRT is fixed, as in U.S.).
Virtual Channel Table (VCT)	provides details about the VCs in the stream, including channel name and number (different forms for terrestrial and cable PSIP).
Event Information Tables (EITs)	describe upcoming program "events," including title, time, captioning services, rating information.
Extended Text Tables (ETTs)	give extended descriptions of VCs and events.
Data Event Tables (DETs)	describe upcoming data "events,".
Directed Channel Change Table (DCCT)	provides definitions of virtual channel change requests.
DCC Selection Code Table (DCCSCT)	carries code values & selection criteria names for reference from DCCT.

These tables are encapsulated as private data into MPEG-2 TS packets and multiplexed along with the video, audio and data streams. The tables are identified by PID and table ID. The PID for MGT, STT, RRT, VCT is 0x1FFB, the so-called PSIP base PID. The PIDs for EITs and ETTs can be arbitrarily selected as long as they do not conflict with other PIDs in the same transport stream. The MGT provides the information necessary to discover what PIDs have been used for the EITs and ETTs. The DETs, DCCT and DCCSCT will not be discussed in this paper because they are beyond the scope of current PSIP carriage agreements.

In the cable environment, the service information is carried in two forms – in-band PSIP and out-of-band SI. The cable in-band PSIP differs in some subtle but significant ways from the terrestrial PSIP described above. Cable metadata relies much more on information in the PMT. For example, both the caption service descriptor and the content advisory descriptor (when present) must be carried in the EITs and may optionally be included in the PMT in the terrestrial world. For cable, these descriptors (when present) must be located in the PMT, and may be carried in the EITs. Additionally, in cable the AC-3 audio descriptor is found in the PMT, and so is optional in the cable EITs.

Another difference is in the VCT. The VCT comes in two forms, one for terrestrial broadcasts (the TVCT) and one for cable broadcasts (the CVCT). They are mostly similar with a few differences. Both list the virtual channels that appear in the broadcast stream and give information for each one including: Channel name, Channel number (two-part for TVCT, one- or two-part for CVCT), MPEG-2 program number (used by receivers to coordinate with entry in PAT), Service type (video, audio, or data-only), and Source_id (used to coordinate VCs with EIT entries). The TVCT supplies PID values of all the video / audio / data streams in the channel.

Terrestrial broadcasters may include the CVCT in their transport streams, in addition to the mandatory TVCT. Cable providers will have instances when they receive the CVCT within a terrestrial stream and pass it through, and other occasions when they will need to generate it locally based on the content of the TVCT and other tables. In terrestrial broadcasts, it is required to include a service location descriptor in the TVCT. This descriptor identifies the various elementary streams (video, audio, and data) included in the complete program. However, the CVCT does not require the presence of a service location descriptor—the information is expected to be present in the PMT.

The OOB SI is defined in SCTE standard DVS 234[3]. While similar in nature to PSIP, the DVS 234 tables are optimized for the cable environment. Tuning relies much more heavily on the data in the PMT. Aggregate EITs (AEITs) and ETTs (AETTs) are used to reduce the number of PID values that the POD host will need to process (MGT table types and corresponding tag values associate and distinguish the various table sections, rather than multiple PID values). Other notable differences include:

- The SI base PID value is 0x1FFC (in contrast to 0x1FFB for ATSC PSIP).
- The Network Information Table (PID 0x1FFC) delivers the Carrier Definition Subtable (CDS) and the Modulation Mode Subtable (MMS). CDSs define number of carriers used in the system and their frequency locations. MMSs define the modulation format (e.g. QPSK or 64QAM) for each carrier in the system.
- Two alternative types of Virtual Channel Table, Short-form (S-VCT) *and* Long-form (L-VCT), may be present in the transport stream, depending on selected profile (see Table 2).
- The Network Text Table (PID 0x1FFC) carries Source Name Subtables to associate names with each service listed in an S-VCT.
- The S-VCT and L-VCT deliver the Virtual Channel Map, Defined Channels Map, and Inverse Channel Map – the keys to channel navigation using the OOB metadata. VCTs also identify the physical cable carrying the transport stream. The L-VCT also includes carrier frequency and modulation mode information.
- Up to 30 days of event information may be carried in AEITs and AETTs. These use a maximum of four PIDs.
- Multiple MGTs corresponding to distinct channel maps may be included in the transport stream, distinguishable within the POD module (the POD identifies the "correct" MGT using the included map_id value and discards the others).

In-band PSIP Processing

According to the Agreement, when DTV services are carried over cable in-the-clear, PSIP data must be provided in-band if PSIP data is available in the original input signal. When multiple transport streams that contain PSIP data are multiplexed in the cable head-end, care must be taken to preserve the PSIP data from input streams and merge the data for the multiplexed output transport stream. As shown in the previous section, all PSIP data packets have the same base PID 0x1FFB, and use standard protocols to arrange EITs, ETTs and DETs. Thus, simply merging the streams is unlikely to work. Intermingling packets from different transport streams with a common PID value (for example PSIP's 0x1FFB) results in a stream that is not MPEG-2 compliant, and is certain to confuse receivers.

In order to meet the requirements of both the PSIP agreement and in-band cable PSIP carriage standards, the metadata processing system must provide the following two functions:

- Resolving the conflicts of PSIP data between multiple terrestrial streams and merging the data for a single output transport streams.
- Making sure the new PSIP data for the output transport stream complies with cable standards.

To do these, the metadata system must extract the PSIP data from the input transport streams, decode it, and parse the PSIP data to obtain the semantic contents. The decoded PSIP data are then modified to reflect the characteristics of the new output transport stream. Next, the PSIP data from different inputs are consolidated and merged into a single set of PSIP data. Finally, the new PSIP data are re-encoded into MPEG-2 packets and multiplexed back into the output transport stream with consolidated video and audio packets.

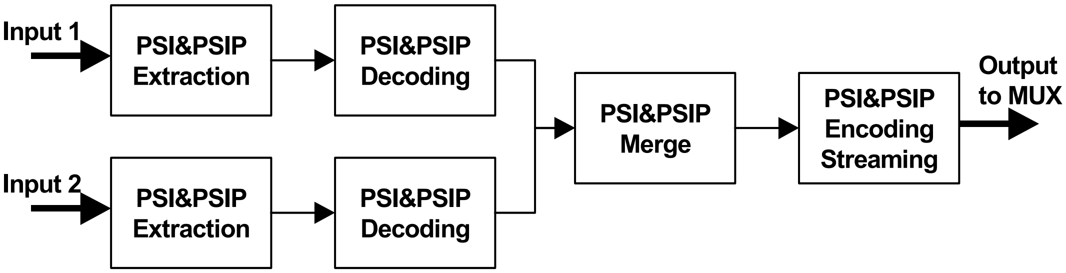

Figure 1: Flow chart of in-band PSIP processing

Figure 1 illustrates the flow chart of the in-band processing routine.

Specific changes have to be made for certain individual PSIP tables. In the VCT for each input Transport Stream, the list of Virtual Channels needs to be decoded. The description of each Virtual Channel will be translated with a new virtual channel number and the frequency of the output ransport Stream. The channel numbering system in cable can be different from that of terrestrial broadcasts. Terrestrial DTV channels are designated by a two-part, major-minor channel number to preserve branding and viewer familiarity, as described above. Broadcasters are likely to want to preserve their major channel numbers onto the cable system, much as they do on analog systems today. However, cable providers will sometimes have to re-map terrestrial broadcasters' channel numbers within the cable system. Cable providers typically do not use the two-part numbering scheme, and may want to convert the two-part number to a one-part number to coordinate it with the out-of-band guide information, though this is not recommended. Once the translation is made, the VCs from different input streams will be merged together to form a single CVCT for the output transport stream.

The EITs from the various input streams must be consolidated. Terrestrial broadcasters are required to carry at least four EITs, namely EIT-0, EIT-1, EIT-2, and EIT-3 that cover 12 hours of programming information. Additional EITs are optional; the number of additional EITs carried will vary between broadcasters. The ATSC recommended practices suggest carrying at least 3 days worth of EITs. In addition, for each EIT group, multiple instances of the tables may exist, each of which is related to a single VC. Because EIT PIDs are arbitrarily selected by each broadcaster, EIT packets from different input sources may have different or overlapping PIDs.

To process the EITs, the MGT from each input transport stream must be decoded in order to find all the EITs contained in the transport stream. Each EIT is decoded to find its association to a specific VC. When EITs from multiple transport streams are merged, the source_id in the EITs may need to be modified in order to resolve any conflicts. In addition, if the input transport streams contain more than four EITs, the additional EITs may be filtered out to reduce bandwidth. Finally, all EITs will be encoded into a new set of PIDs.

The cable operators have the option to carry ETTs or block ETTs. If the carriage of ETTs is selected, the ETTs have to be processed in much the same way as the EITs are.

Due to the changes in EITs and ETTs, the MGT for the output transport stream essentially has to be regenerated to reflect the presence of the EITs and ETTs, new PID selections, and table lengths of all PSIP tables. In addition, the RRT and STT may be updated if necessary.

Because of changes to the program line-up after transport stream re-multiplexing, it is obvious that the PAT and PMT will also require modification or regeneration. Although MPEG-2 multiplexers are designed to handle any PAT and PMT changes, the difference between terrestrial PSIP and cable in-band PSIP may require additional functionalities that do not exist in traditional multiplexers. For example, under certain situations, some descriptors may need to be copied from PSIP tables to the PMTs.

Finally, the system must provide a bandwidth estimation and reduction function in order to meet the bandwidth requirement specified by the PSIP agreement. There are several optional features that will affect the PSIP bandwidth, including the number of EITs, whether or not to include ETTs, as well as the time interval of EITs, that are not specified in the PSIP standards. Furthermore, multiple tables with the same PID may be packed into a single MPEG-2 packet. Finally, compression technology may be used to further optimize the bandwidth usage.

Out-of-band Service Information Generation

Cable systems have traditionally sent EPG data in an out-of-band (OOB) channel that describes all the programming available to the viewer in a single feed. Cable providers will continue to deliver guide information in the OOB channel in the digital domain, but most of the presently used methods are proprietary. Off-the-shelf cable-ready DTV receivers need a standards-based mechanism for delivering out-of-band metadata.

The metadata processing system will generate the OOB SI according to DVS 234. Although the cable OOB SI delivers information similar to in-band PSIP, the protocol used to format the data is significantly different from that of terrestrial or in-band cable PSIP, as described previously. The cable OOB SI contains several unique tables, including the NIT, NTT, S-VCT, and Aggregated EITs and ETTs.

The NIT and NTT are easy to generate and do not change once they are created. Both S-VCT and L-VCT are created to include all virtual channels in a cable network. In addition, links between DTV services in input streams and VCs in the output channel map will be maintained.

To generate the aggregated EITs and aggregated ETTs, a similar process to that of in-band PSIP processing will be used. First, the EITs and ETTs in the input streams will be decoded. The decoded tables will be updated for any changes in the program source_id. Finally, the aggregated EITs and ETTs are created by combining multiple tables from different sources into single aggregated tables.

However, not all OOB SI tables may be created and delivered depending on the profiles selected by the cable operator. DVS-234 defines six profiles for delivery of the service information via the out-of-band channel, described in Table 2.

Table 2. DVS 234 Metadata Delivery Profiles

Profile	Attributes
1 – Baseline	uses Short-form VCT, Modulation Mode, and Carrier Definition subtables for navigation
2 – Revision Detection	builds on Profile 1 by adding a revision detection mechanism
3 – Parental Advisory	builds on Profile 2 by adding RRT support for compliance with FCC-mandated content advisories
4 – Standard EPG Data	adds AEITs and AETTs to Profile 3 for non-proprietary EPG support
5 – Combination	allows navigation based on Long-form VCT, in addition to Profile 1 navigation
6 – PSIP Only	navigation is restricted to Long-form tables as in terrestrial PSIP

None of the mechanisms described thus far preclude the use of the proprietary service selection and navigation systems frequently used in cable today. These proprietary systems are likely to remain in the mix for the foreseeable future, supported by system-specific decryption functions in POD security modules. During this interim period, proprietary system suppliers will likely develop methods for ingesting the most-timely metadata available – delivered to the head-end "live" by PSIP in the case of terrestrial broadcasts.

IMPLEMENTATION EXAMPLE

Figure 2 illustrates the metadata processing system in the cable head-end environment. DTV services arrive at the cable head-end over satellite and terrestrial links, as well as via other means (over an ATM network, for example). Multiple transport streams originating from different sources are merged into a higher bandwidth transport stream and are then modulated by a QAM modulator before being sent out to customers via the cable plant. Depending on the bandwidth of the input streams, a variable number of transport streams can be multiplexed into a single output stream.

To protect the investments already made by the cable operators, the system we implemented works together with the MPEG-2 multiplexers, leaving those multiplexers responsible for the audio and video streams, while the metadata processing system is responsible for processing PSI and PSIP data.

The metadata processing system takes either full transport streams or “composite” SI streams – created by the multiplexers – as inputs. Depending on their source, some incoming transport streams may contain PSIP data while others may not. For example, transport streams received from off-air terrestrial broadcasts will typically contain PSIP, while encrypted streams may not include PSIP. The metadata processing system monitors each of the incoming streams in real-time and provides detailed information about the contents of the stream, the association between MPEG-2 programs and virtual channels, and program guides.

When multiple streams containing PSIP data are multiplexed into a single transport stream, the system processes the PSIP and creates new in-band PSIP data for the output stream. As described previously, the system decodes the original PSIP data obtaining the semantic contents, translates the data to a form consistent with the local cable environment and merges the metadata at the content level. The resultant PSIP tables are then encapsulated into MPEG-2 packets. A streaming device, which is a part of the metadata processing system, outputs the PSIP MPEG-2 packets based on the standard table requirements taking into consideration the bandwidth limitations prescribed by the Agreement. The output PSIP data stream from the metadata system is then multiplexed back into the output transport streams along with audio and video and other elementary streams. This forms the in-band PSIP data required by the NCTA-CEA Agreement.

In addition to handling in-band PSIP, the metadata processing system also generates an OOB SI stream. The aggregated SI data contains the information for all the “in-the-clear” virtual channels in the cable network, as well as any VCs the cable provider chooses to include for the purposes of discovery. For incoming streams that contain PSIP data, the system optionally extracts the EIT and ETT data and converts them to the aggregated SI format described in DVS 234. For incoming streams that do not contain PSIP, the system allows the input of the VC information so that it can be included in the OOB virtual channel map.

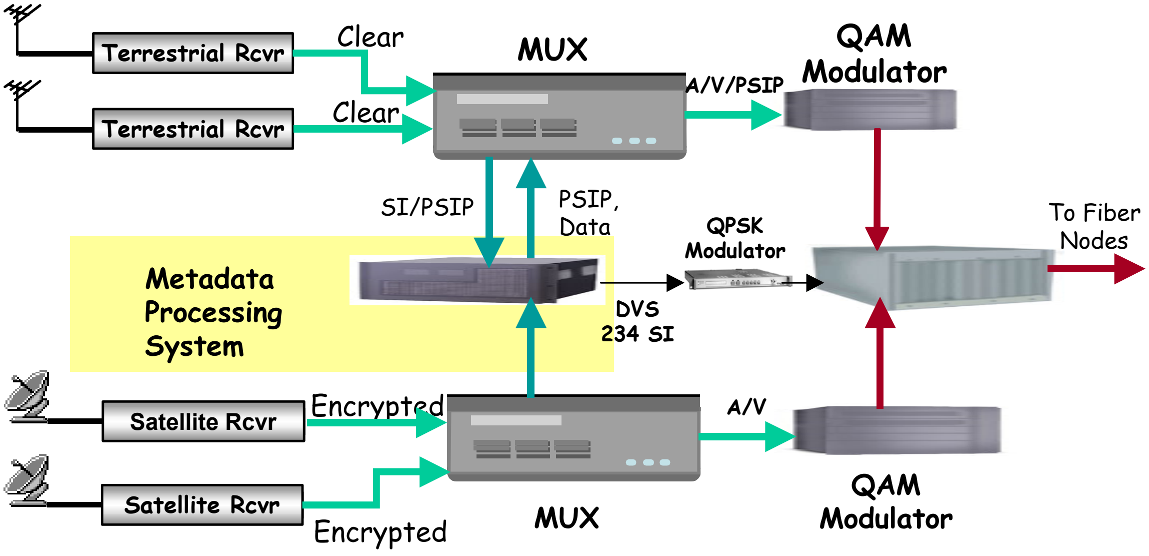

Figure 2: Metadata Processing System in Cable Head-End Environment

The metadata system also provides control of the associated multiplexers. Since typical multiplexers are not designed to handle PSIP structures, it is important that the metadata processing system can be tightly coupled with a multiplexer to provide direct control for handling PSIP packets. In addition, because the metadata system has the complete information about the elementary components of the input and output streams, the system can be used to automatically set up multiplexer functions, such as PID pass-through, PID blockage and PID mapping. The control feature is useful for the cable operator since it precludes the need to manually discover and manipulate the elementary stream PIDs.

In the future, the metadata processing system could also be linked to the proprietary program guide service to perform real-time updates of the service information. Typically, the database used by the cable operator for EPG service is days or weeks old. When a program, such as a sporting event, runs over time, the EPG information following the overrun program event is out-of-date. If the incoming stream contains updated PSIP information, this information could be used to update the cable guide.

SUMMARY

PSIP data in DTV channels provide broadcasters a tool for maintaining their analog brand and promoting their services through EPG information. Therefore, there is a strong incentive for broadcasters to include rich, timely PSIP data in their DTV transmissions.

When DTV services are delivered through a cable system, the PSIP data in the original broadcast signal must be updated and carried though the cable network so that consumers with off-the-shelf cable-ready DTV receivers can view DTV services in the clear. This is a key motivation for the FCC Report and Order mentioned above, and drove the NCTA and CEA to create a PSIP carriage agreement. In addition, service information that is based on SCTE standards and contains the authoritative, up-to-the-minute program guide information has clear advantages over proprietary EPG systems.

Cable providers need a device capable of handling PSIP and PSI data from multiple transport streams. An external metadata processing system, as presented in this paper, minimizes the impact on the cable head-ends, while still allowing them to comply with industry standards, carriage agreements, and FCC mandates. The metadata processing system in the cable head-end manipulates, aggregates, and harmonizes metadata for inclusion in the in-band and out-of-band channels. By directly ingesting terrestrial PSIP data, the cable system delivers relevant, accurate metadata to the viewer and enables off-the-shelf commercial receivers while maintaining compatibility with legacy receiver equipment.

REFERENCES

[1] ITU-T Rec. H. 222.0 | ISO/IEC 13818-1:1994, Information Technology — Coding of Moving Pictures and Associated Audio — Part 1: Systems.

[2] Program and System Information Protocol for Terrestrial Broadcast and Cable (Revision A), ATSC Document A/65A, 23 Dec 1997, Rev A, 31 May 2000.

[3] Service Information Delivered Out-of-Band for Digital Cable Television, SCTE Document DVS 234, Revision 2, 28 March 2000.

[4] NCTA-CEA Agreement reference.

CONTACT INFORMATION:

Jian Shen, Triveni Digital, Inc., 40 Washington Road, Princeton Junction, NJ 08550

IP STREAMING AND BROADBAND, TOGETHER ON THE ROAD TO THE TWO-WAY IP HIGHWAY?

Chris Dinallo
Pace Micro Technology Americas

Abstract

*Today's broadcast transport mechanisms limit the services that can be offered to subscribers***.** *The content itself must be manipulated (encrypted, multiplexed, modulated) in many stages before being sent down the cable. At the end of the cable, the set-top box needs special capabilities to make that content usable to the subscriber. All this is done using costly, time-consuming proprietary techniques. Enter IP (Internet Protocol): an open standard transport mechanism that allows operators to deliver existing program content with little or no manipulation. In addition, IP streaming allows operators to offer new and enhanced value-added services to subscribers, such as photo galleries, home movies, audio jukebox, email, advanced personal video recording (PVR), Web access, chat, gaming, Web-enabled programs, e-commerce/t-commerce, video on demand (VOD), subscription video on demand (SVOD), voice over IP (VoIP), and streaming media over home networking.*

Attendees will learn the following from this session:

- *Why is IP streaming a better alternative?*
- *What will it take from a technology and business standpoint to move toward IP?*
- *What can be done with IP streaming that can't be done today?*
- *Benefits to operators.*
- *Benefits to consumers/subscribers.*
- *Deployment challenges.*

> *"Much the way the Compact Disc and Internet has changed this planet forever, so will IP Streaming." C. Dinallo*

IP STREAMING CRITICS

Sure some will say that they have no use for Internet Protocol (IP) Streaming (also commonly referred to as Streaming Media) because broadcast technologies can do everything they need. Plus, who's to argue providers have been delivering content to millions of users for over 50 years. Much the way the LP has been surpassed by the CD, IP Streaming will too have its day. To understand why, one has to look deep into the characteristics that cause a new technology to succeed. The CD didn't put the LP out to pasture just because it eliminated the poor quality of pops and hisses. It offered much more. Value-adds like durability, portability, enhanced content (text & graphics), and yes good quality sound. Vinyl LP records just couldn't possibly compete with this new paradigm shift in listening technology. To understand if and how IP Streaming can flourish, let's explore beyond the surface of all the benefits of what it brings. And to do justice for those hard-liners still listening to their LPs, we'll discuss the costs and challenges associated with this technology that the network operators have to deal with regardless if they are cable, satellite, terrestrial, or DSL providers. The focus of this paper will be on cable since it has the most investment in legacy transport networks.

WHY SWITCH TO ANOTHER TRANSPORT MECHANISM?

Today's broadcast transport mechanisms constrain the services that can be offered to the subscribers by virtue of being mostly a one-way pipe of information. Audio and video content flows downstream from the head-ends to all set-top receivers. It's an all or none condition where each set-top receives the same data flow and is very restricted in how it can communicate back to the head-ends. In addition, the content itself must be manipulated (encrypted, multiplexed, modulated) in many stages before being sent downstream. Similarly, once downstream the cable the set-top receiver needs special capabilities to make that content usable to the subscriber. This is done using proprietary techniques. And although this works fine for broadcast services it does not address how to get custom services to *specific* set-top receivers. Nor does it leverage open standards such as Internet-Protocol (IP) in which the large development community could be leveraged to develop more enhanced services in a timely fashion. What we're getting to is how IP Streaming can be utilized to achieve the vital two-way interaction between the provider and a *specific* subscriber.

It is this two-way interaction that is driving providers to get excited about IP Streaming. For it enables the additional revenue they can create from their subscriber base. It allows for new services to be deployed that have never been possible before. Services like; sharing your digital photos and home movies over the broadband network. These are value-adds that are new, novel, and provide a value to the subscriber. Yet, cool new concepts don't always succeed in the market without the proper business model to carry them forward.

Such enhanced value services are best introduced to market utilizing a bundled approach. Where the "bundle" is comprised of familiar needed services combined with the new enhanced services. To do this, the provider creates a "bundled service" in which many of the provider's standalone features are combined into a package and offered to the subscriber at a cheaper rate than if the subscriber bought the same services individually. Bundled services are a proven mechanism that achieves greater value across many industries. Telephone Local Exchange Carriers (LECs) have been doing it with great results for years. For example, some LECs offer "premium" packages that include: Call Waiting, Caller ID, Caller ID w/ Name, Call Block, 3-Way Calling, Call Return, and Call Transfer. It not only provides a great value to the subscriber, but produces a substantial deterministic repeating revenue stream for the provider, sometimes as much as 50% of the basic service.

Providers want to leverage this same mechanism, albeit using services tailored to their industry: Photo Galleries, Email, Web Access, Chat, Web Enabled Programs, SVOD (Subscription Video on Demand), VoIP (Voice over IP), and yes, IP Streaming Media. IP can and has served all of these capabilities. Whether or not it makes fiscal sense to have streaming media over the cable plant is yet to be determined. In addition, because the Internet has become ubiquitous, it makes good business sense to leverage that infrastructure as a transport mechanism for content delivery.

Even as one decides IP Streaming is the right thing to do, there is still the question of how to do it and what transport to use. The obvious choice would be to encapsulate the IP data packets into the MPEG-2 payloads for broadcast (known as IP over MPEG-2), but doing so would not unleash the full potential that IP Streaming can offer. To merely broadcast IP data packets to all set-top receivers does not help enhanced services that need to target specific set-tops. What's needed is a mechanism for direct point-casting. Furthermore, MPEG-2 transport does not offer a return communication path. Here again, enhanced services do not gain any advantage over legacy mechanisms for fast two-way interaction. Thus, two criterias become apparent: (1) point-cast for direct addressing as opposed to broadcast transmission and (2) two-way communication for interaction. Given this, a better approach is to use the IP over DOCSIS transport layer. For DOCSIS addresses these two criterias and much more. Refer to sidebar on *IP Streaming over DOCSIS* to gain more insight into two-way communication.

To understand what IP Streaming can offer, let's briefly discuss how it came about and then we can explore how it enables the advantages of two-way interaction.

INTERNET PROTOCOL – IP

The Internet was created by U.S. defense and academic institutions to facilitate the sharing of data. It was conceived as a computer-based system in the data communication domain. Since then it has evolved tremendously to affect everyone.

The evolution was started by the Internet Engineering Task Force (IETF), an Internet standards group responsible for the design and upgrade of all Internet communication protocols, Internet Protocol (IP) standards were designed to solve the issues of communications between computers across heterogeneous networks. These standards deal primarily with issues such as networking, routing, and congestion control and do so by specifying a means fordata to be converted into smaller manageable sized data packets. IP is also a layered architecture. Residing on top of IP, many other powerful standards were created. These standards were used in developing applications like the World Wide Web, File Transfer Protocol (FTP), Usenet, and e-mail. These applications implement one standard IP transmission type named unicast. Unicast is the method used for two distinct endpoints to communicate directly (i.e., point-to-point). As IP routers became more powerful, multicast transmission came into existence, which offered more efficient use of the shared bandwidth across the network (i.e., point-to-multipoint). Multicast enables a single endpoint to communicate with many endpoints in one transmission session. This is much the same as the way broadcast TV or radio arrives at everyone's home or car within the network's geographic area, yet in a digital medium. Applying this technique of multicast to transporting multimedia data (audio, video, graphics, text, & data) now has the added benefit of reaching many subscribers while not overloading the digital bandwidth available in the network. Providers, especially cable operators, have a huge pipe to deliver multimedia data. Furthermore, since their network topology was designed as an 'edge' network, IP transmission gives them an advantage to efficiently transport their services to millions of subscribers in a dynamic and interactive fashion.

IP also has two transport delivery methods that reside on top of the layered IP architecture. The first is *Transmission Control Protocol* (TCP). TCP is connection-based where the data packets are transported using guaranteed delivery approaches that have automatic retries. Providers can use this for situations where data corruption or total loss is not tolerated. TCP does have associated overhead with it, so nothing is for free. The second transport method is *User Datagram Protocol* (UDP), which is connection-less. UDP is a highly efficient transfer mechanism of data packets that actually contain and can distinguish between multiple destination addresses. Unlike TCP, there is no receipt acknowledgment of delivery thus making UDP an “unreliable” protocol. Some networks choose UDP over TCP is because of the packet overhead savings by not opting for guaranteed deliveries. This is the automatic retryof retransmitting data packets if the sender does not receive an acknowledgement. Hence, the trade-off is delivery reliability vs. bandwidth. Many providers today choose UDP because it more closely emulates the efficiency of broadcasting.

Despite IP Streaming’s benefits, there are still critics in the broadcast industry. Some view IP Streaming as both an encroachment and redundancy of how content is delivered to the subscriber today. On the surface, these criticisms are valid because IP streaming does what their broadcast technologies have done for quite some time. However, drilling down one sees that IP streaming can enable new capabilities to many devices and is not just limited to televisions. Any device that can obtain connectivity to the network whether it be wired or wireless can become an IP Streaming client. All this is accomplished in an open more efficient environment which, at the end of the day, drives better viewing experiences, faster, and more accessible flow of information of all types that yield associated increased revenue opportunities for the providers.

One might question what an open environment offers to typically closed networks such as the ones that cable operators in North America have built. In relation to IP Streaming, open standards are the preferred approach for it allows more choices for the operators, faster development times, interoperability across heterogeneous networks, and cost benefits from having more open competitive supply vendors.

***IP Streaming over DOCSIS*:**
If the goal is to have an open standard that promotes two-way communication over broadband networks then the standard bodies unanimously agree that Data Over Cable Service Interface Specification (DOCSIS) is the preferred cable modem implementation over proprietary techniques. DOCSIS provides a high-speed, two-way communication path between the set-top receiver and the head-end plant. Although now the head-end equipment will consist of an additional piece of hardware called a Cable Modem Termination System (CMTS). The CMTS has the role of back hauling transmit and receive requests across the network. It performs this role in a IP-centric fashion including direct addressing to a specific set-top. It’s this direct addressing that unleashes the power for IP Streaming to capitalize on and providers to offer new services. Hence, with DOCSIS, the broadband network is transitioned into a two-way high capacity data independent carrier that goes beyond the reach of traditional broadcast television while still preserving the QAM infrastructure the provider has invested in over the years. Pace was the first and still the leader in DOCSIS deployments with over 2 million digital set-tops with integrated DOCSIS modems in the field.

As broadcasters solidify around a single approach, whether it is MPEG-4 and its multifaceted parameters or the more straightforward MPEG-2, IP Streaming will become more ubiquitous.

COMPRESSION ALGORITHMS

The whole concept behind streaming media is focused on compressing data (all types) into efficient packet sizes that minimize bandwidth, yet still preserve data integrity. The encoding side is only half of the challenge. The other half deals with ensuring that the receiving edge device can decode the encoded streams thus reconstructing back to the originals. It's this reconstruction process that drives the need for open standards and its interoperability goals.

Today, MPEG-2 dominates the industry with nearly all streaming media being comprised in some form of MPEG-2. There is a strong need among cable operators to achieve high quality video at significantly less than 1MB per second data rates. In fact, deployed schemes today are merely at data rates of 384KB/sec, 15 frames per second, 32bit color, and ¼ screen resolution. Clearly, not premium quality that subscribers demand. To address this quality deficiency while still preserving bandwidth capacity there are some techniques emerging that claim to have high quality MPEG-2 at 1MB/sec data rate.

In addition to emerging MPEG-2 techniques, of late, MPEG-4 is the successor of MPEG-2 that has been getting many headlines. MPEG-4 was created to address better quality at less bandwidth consumption of its predecessor along with adding enhanced features like non-rectangular objects known as sprites and animation. However, with these enhancements come complexities such as: file and transport formats, and control protocol. Unfortunately, it's these flexibilities that permit various streaming media implementations to have interoperability issues.

Once broadcasters and providers can adopt common formats within the MPEG parameters, the IP Streaming adoption rates will dramatically increase. These same issues also exist for the computing environment. To better predict where the entertainment market will go with respect to IP Streaming, one needs to see who the players are today.

Interoperability is key to driving this technology.

TECHNOLOGY DRIVERS

Today, the major technology movers in IP Streaming are Microsoft, Real Networks, and Apple Computer. There are others with products in market niches, like Video on Demand (VOD) that offer streaming solutions tailored to their environments. Yet, they all have one thing in common, each has adopted a proprietary version of the MPEG-2 standard. Given the benefits of open standards one may wonder why a proprietary standard is being used. That doesn't sound like it's in the best interest of moving a technology forward. This single reason alone could cause a slow adoption rate of this technology all because these companies have each adopted proprietary encoding schemes and associated streaming clients. Interoperability becomes non-existent.

There is also a trend to support MPEG-4. Microsoft's streaming solution actually uses an MPEG-4 compliant algorithm, however, its file format and multiplexing technique known as Advanced Streaming Format (ASF), is not compliant to MPEG-4. In contrast to Apple's Quicktime format, which follows the MPEG-4 format more exactly.

Some vendors are exploiting embedding techniques for MPEG-2 content that mimics capability designed into MPEG-4. It is called SMIL (Synchronized Multimedia Integration Language). SMIL is a text based markup language (really XML based) that allows a given stream to embed and/or link in other streams. The concept is not new and commonly referred to as metafiles (describes files that tie sets of other files together). The W3 Consortium has recently proposed a recommendation for SMIL 2.0 Animation. Techniques such as these that support metafiles are gaining momentum.

Of the top 3 vendors previously mentioned, the following tags identify their file and metafile formats:

- `Microsoft’s WindowsMedia player - .ASF and  .ASX for metafiles`
- `Real’s Realplayer - .RM and .RAM & .SMI for metafiles`
- `Apple’s Quicktime player - .MOV`

Yet, with MPEG-2, MPEG-4, and SMIL standards actual implementations today have inter-operatorability issues.

> *Web content drives enhanced services. Enhanced services drives revenue.*

ENHANCED SERVICES

Let’s be clear about one thing, enhanced service revenue generation is the overwhelming factor why providers make the investments they do.

For IP Streaming to become pervasive, providers must find value-add in doing what previously could not have been done or done easily and cost effectively. One approach is to evolve current services into rich enhanced services of all types. For example, VoD, VoIP, Audio Jukebox, Home Movies, and Photo Galleries are all types of enhanced services that can both benefit from IP Streaming while building on proven business models that subscribers will and (currently) do pay for.

Given this, let’s explore how we get to the money. What are the benefits and costs to the providers. Does the subscriber really gain value-add from this enhanced service enabling technology?

> *Reaching outside the TV box!*

BENEFITS TO PROVIDERS

First and foremost, providers benefit economies of scale of infrastructure by adopting an open transport model, i.e. Internet Protocol. Proprietary infrastructure delivery systems don't interoperate with other systems. And as such, providers are locked into specific technologies, costs, timeframes, and capabilities. Furthermore, some providers have many different "closed" systems in geographically isolated networks rendering it impossible to obtain economies of scale. In this situation, everything is different from unique equipment to back-office billing and support tools, right to and including the set-tops!

Another benefit of IP Streaming is the demand it drives for high-speed access. This is directly related to increased revenue opportunity. IP Streaming also allows for more dynamic and tighter Web content integration. This reduces costs by both equipment and manpower overseeing the content delivery operations. Less content data manipulation is another benefit. This comes into play with add insertions and utilizing off-the-shelf web servers which boils down to less specialty proprietary equipment needed.

Yet, one of the biggest benefits of all is expanding the provider’s customer base. This is realized because content delivery can now be to any IP connected device and no longer dependent on TVs only.

BENEFITS TO SUBSCRIBERS

From a subscriber’s perspective, they will need to see the benefits of IP Streaming in order to justify the possible rate increase associated with the new IP services. The proponents of IP Streaming believe the viewing experience will be more enriched because the providers can enhance the IP streams very easily. By using techniques like metafile support the content is much more dynamic. The provider can also choose to cache the streaming media content locally on hard drives within a set-top receiver, thus tremendous amounts of information are literally in the hands of subscribers. This cache can also be tailored to each subscriber's preference. Imagine having the data you want a mere push-button away without the time-consuming web surfing. This is definitely an enhanced service worth paying for.

CHALLENGES TO PROVIDERS

Where's the downside? IP Streaming still has it challenges. In North America, cable operators are still faced with deploying DOCSIS. There is still a large challenge in implementing Quality of Service (QoS), although this is getting resolved with IETF technologies focused on traffic management such as MultiProtocol Label Switching (MPLS) , Resource Reservation Protocol (RSVP) , and policy management like the Common Open Policy (COPS) protocol. COPS and RSVP work together in managing intelligent routing of network traffic based on priorities, traffic type, and user subscription level. This allows for tiered level services such as basic and premium. MPLS comes into play by allowing these route decisions to happen very quickly. In essence, MPLS acts as a lookahead on each packet to determine the fastest way to route it to its destination. It does this by integrating the data link layer and network layer. The integration point inserts a small "label" in the packet header that instructs the MPLS-enabled switches/routers how best to route to the destination address. As these technologies gain more deployment, the QoS issue becomes a non-issue.

Other areas of the technology have to do with security. Providers must make their infrastructure more secure by adding proper firewalls and other techniques that combat against denial of service attacks. There is also the much-debated topic of Digital Rights Management (DRM) techniques. All parties must be protected from pirating abuse. This topic is highly controversial and is beyond the scope of this discussion. A good starting reference is Linden deCarmo's article on DRM (see http://www.zdnet.com/products/stories/reviews/0,4161,2766381,00.html).

Lastly, providers need a common IP Streaming format and edge device client that can interoperate seamlessly. If this is made a priority, then the differentiating factors would merely be thin vs. thick clients and not the format of the content.

COSTS TO PROVIDERS

All of these challenges do come at a cost to the provider. This is why the business model must be right for IP Streaming to take-off. Here's a high level view of what kind of costs, both capital and personnel, providers can expect in this transition to IP.

- IP proficient technical personnel
- They should plan on consolidating network operations & platforms in order to reduce complexity of managing the NOC (Network Operations Center). This will entail forklifting legacy equipment out and putting in more data centric equipment.
- Revamping their IP network to employ a highly reliable network topology: including making it highly available to a minimum of 4-9s of availability (99.99% min. uptime), fault tolerant schemes, redundancy, failover, and load balancing.
- Investing in enhanced security techniques and equipment.

Yet, there is good news here for the cable operators. Besides having having done a great job of upgrading their Hybrid Fiber Coax (HFC) network to give them unsurpassed 24x7 high-speed bandwidth directly into the subscriber's home they have built up much of the needed infrastructure and expertise to deploy IP Streaming. And like their HFC upgrade, IP networks will take time and come on-line in increments. And rightly it should, because there is a ton of legacy equipment that still has return on investment life. IP and broadcast QAM technologies will co-exist for many years. Yet, now is the time to initiate the rollout of IP Streaming Media services.

IP HIGHWAY

As one can see, there are some hurdles along the IP highway to reach a solid business model in which both providers and subscribers benefit. However, the benefits listed above outweigh the hurdles. In fact, the outlook appears good from what the market researchers are forecasting. The prediction is for the streaming media market to skyrocket as compression algorithms improve and high-speed access gains momentum. Forrester Researchers predict that by 2003, 33 percent of all households will have broadband access. More specifically, according to DFC Intelligence, a research firm for interactive and digital entertainment, video streaming on the Internet grew 215% in 2000 to over 900 million total streams accessed. This includes broadband streams, which made up almost 29% of total accesses. As broadband moves toward ubiquity, operators are uniquely positioned to make streaming media what it's meant to be – a viable, revenue-generating business.

And rest assured that during the co-existence period of IP and Broadcast technologies Pace has products available today to meet the needs whether they are an incremental approach to IP utilizing our DOCSIS set-top technology as in our Di4000 and 700 series set-tops or a pure IP solution such as the DSL4000 and IP500 digital gateways.

Chris Dinallo, Chief Technologist

Mr. Dinallo's responsibilities for Pace focus on US Cable and include digital set-top box development, future technological directions, participation in standard bodies such as Cable Labs, SCTE, and TV Linux Alliance. Chris brings 17 years' experience developing innovative solutions in software and firmware. Since 1989, his expertise has been in the discipline of multimedia technologies. Prior to joining Pace, Dinallo has held engineering director positions in DVD and Voice over IP companies. In the (VoIP) telephony solution space, Dinallo's team architected next generation networks with a focus on enhanced services for telecom and cable service providers utilizing open standards and leading internet technologies.

References

Michael Adams, OpenCable Architecture, Cisco Press, 2000

W. Richard Stevens, TCP/IP Illustrated, Vol1 The Protocols, Addison-Wesley, 1994

DOCSIS, Cable Labs Specifications, http://www.cablemodem.com/specifications.html

PacketCable, Cable Labs Specifications, http://www.packetcable.com/specifications.html

RFC 2205, Resource ReSerVation Protocol (RSVP) Version 1 Functional Specification, Internet Engineering Task Force (IETF), September 1997, http://www.ietf.org/rfc.html

RFC 2326, Reatime Streaming Protocol , Internet Engineering Task Force (IETF), September 1998, http://www.ietf.org/rfc.html

RFC 2748, The COPS (Common Open Policy Service) Protocol, *Internet Engineering Task Force (IETF), January 2000, http://www.ietf.org/rfc.html*

Linden deCarmo, New digital rights technologies protect content creators' interests, but what about users' rights?, *http://www.zdnet.com/products/stories/reviews/0,4161,2766381,00.html, June 2001*

IP VIDEO TO THE HOME USING AN ETHERNET PASSIVE OPTICAL NETWORK

T. Neel, Alloptic

One of the key challenges to Internet protocol (IP) video in the past has been the large amount of bandwidth needed to deploy it. Optical fiber is recognized as the most effective means for transporting voice, video, and data traffic; however, it is expensive to deploy and manage point-to-point (PP) fiber connections between every subscriber location and the central office. Ethernet passive optical network (EPON) is an emerging broadband fiber infrastructure that addresses the high cost of PP fiber solutions. Gigabit EPONs were also developed to address the shortcomings of asynchronous transfer mode passive optical networks (APON), which are inappropriate for the local loop as they are complex and expensive and lack sufficient video capabilities and bandwidth capacity. EPONs provide greater service capabilities and higher bandwidth at reduced costs. Though EPONs are in the early stages of development, they figure to become the primary means for delivering converged video, voice, and data over a single optical access system. This paper discusses the benefits, features, and technological foundation of EPON, in comparison with APON, as the best alternative end-to-end architecture for IP video.

The Case for IP Video Services

Video Delivery Options

Over the past few years, digital-cable and digital broadcast satellites (DBS) have created a new business revenue model for the television (TV) industry. Similarly, IP multicast has the ability to generate new revenue models for the Internet. There are many ways to deliver video to the home. Video delivery options to the home include the following: cable TV (CATV) over hybrid fiber/coax (HFC), DBS, digital cable over HFC with a set-top box, multi-channel multipoint distribution service (MMDS), IP video over digital subscriber line (DSL), analog video overlay over fiber, and IP video over optical EPONs. Despite this plethora of choices, the trend is toward convergence on two different layers (see *Figure 1*).

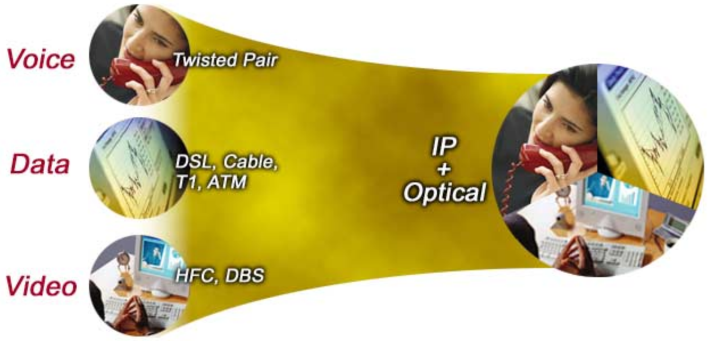

Figure 1: *Convergence of all Services over One Network*

In the first layer of convergence, three networks are converged into a single optical-fiber network. The second layer of convergence occurs in the IP layer. With the introduction of new services to IP, there is convergence to a single infrastructure on a single protocol. IP is thus the ubiquitous protocol for network convergence (see *Figure 2*).

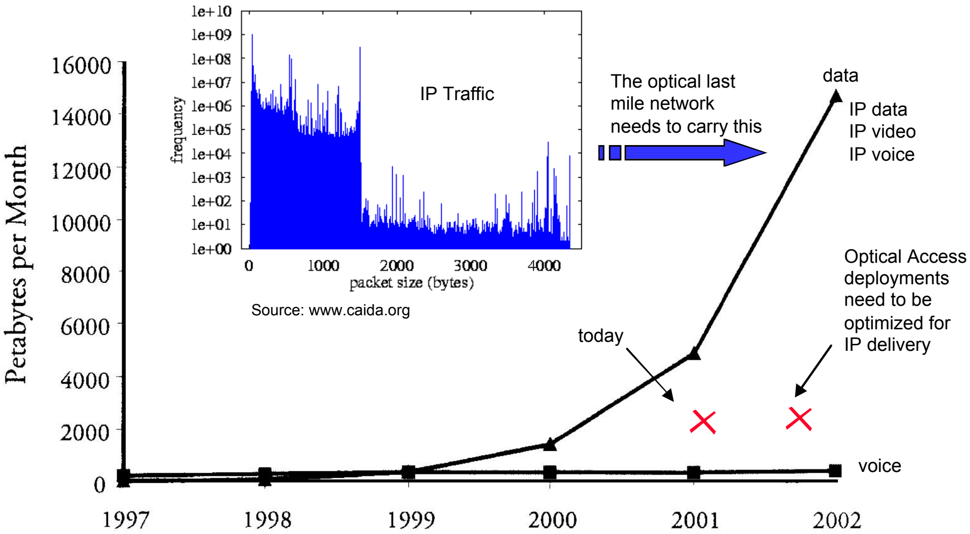

Figure 2: *IP Traffic Will Dominate the Last Mile*

Because of the increase in IP traffic on networks, service providers must be able to offer different types of services, such as video, data, or voice. Optical access deployments need to be optimized for IP delivery, and Ethernet is just that—a packet-based network, optimized to carry IP traffic.

Services Enabled by IP Video

IP video enables a growing number of services. Existing services for real-time video include broadcast and pay-per-view TV, and streaming video is an example of buffered video on the Internet. The simple act of adding bandwidth enables new video services. New services for real-time video include interactive TV, on-demand TV, videoconference, integrated Web content, digital TV (DTV), and high-definition TV (HDTV). New services in the buffered video space include personal videocassette recorders (VCR), time-shifted TV, and full-screen streaming video. Viewers can now enjoy interactive features of video, Web-based video, and services that allow them to watch programs at their convenience. The market for personal VCRs is growing at the moment because of TiVo boxes, which act locally or remotely in an IP video head-end. As this technology matures, more enablers will become available. Service providers benefit from IP video services, as it is easier to manage a converged single network for voice, data, and video. IP video services create multiple revenue streams, improved competitiveness, and a simplified network at reduced cost that is easy to manage, administer, and customize. IP video is in the process of becoming an incremental investment over data and voice, which are already in IP format, and it allows service providers to amortize the cost of broadband buildout over multiple services. IP video reduces provisioning costs for the service provider, as interactive media allows the subscriber to perform self-provisioning relatively easily. Converged networks minimize equipment and network management costs and eliminate the need to train a large number of technicians to maintain the network. The benefits of IP video services for the end user include more video services, better picture quality, bundled service packages, Web-enabled services, and ultra-fast Internet access. In addition, end users can choose from a large team of service providers to provide any service at any time.

EPON Primer

There are three primary ways to run fiber to a home. The first option is to establish PP links (see *Figure 3*).

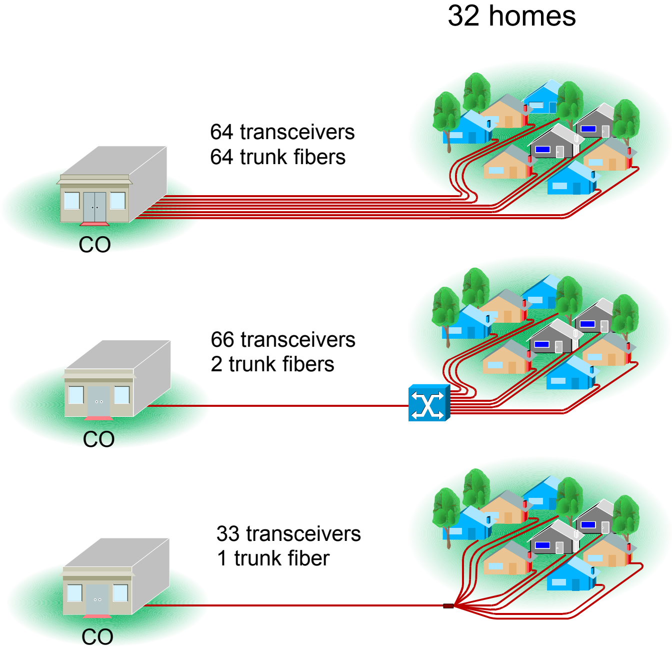

Figure 3: *PON—a Natural Step in Access Evolution*

In *Figure 3*, each of the 32 homes uses two fibers, which correspond to 64 transceivers and 64 trunk fibers. Convergence to a single fiber reduces the number of fibers and transceivers to 32 and 64 respectively. Because it is a fiber-rich infrastructure, companies have begun to install curb switches, such as Ethernet switches, in the network. A potential problem with these switches is that they require power, which creates a separate network with power needs. Using a curb switch reduces the number of trunk fibers; however, the number of optical transceivers is not diminished. The cost of curbside switch boxes is dominated by the cost of the optical transceivers. Therefore, the most logical solution would be to replace this box with an optical splitter to serve as a passive directional coupler. The optical splitter reduces the number of transceivers needed by nearly 50 percent and conserves fiber in the trunk, on the order of 1 to 64. PON is therefore a natural step in access evolution, as it simplifies the infrastructure and requires little maintenance.

PON Architectures

Bandwidth is increasing on long-haul networks through wavelength division multiplexing (WDM) and other technologies. However, there is a gap between metro-network capacity and the end user’s needs, separated by this last-mile bottleneck. PONs respond to this last mile of the communications infrastructure between the service-provider central office, head-end, and customer locations. PONs are point-to-multi-point (PMP) networks, which have the capacity to downstream broadcast media to residential homes. The following figure illustrates PON architecture (see *Figure 4*).

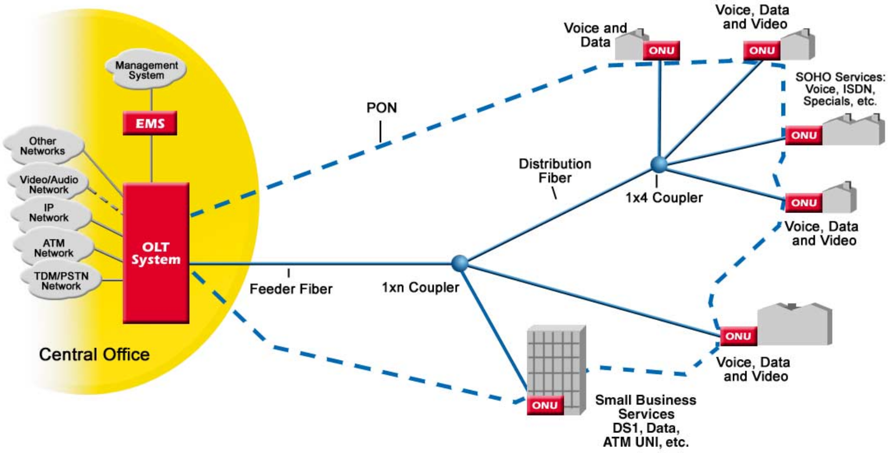

Figure 4: *PON Architecture*

There are some benefits to using PONs for video overlay. Although there is a cost to install an amplifier, which is shared by multiple locations, PONs are passive and eliminate all power in the field. The two primary types of PON technology are APON and EPON.

APON

APONs were developed in the mid-1990s through the work of the full-service access network (FSAN) initiative in the interest of extending high-speed services, such as IP data, video, and Ethernet over fiber, to homes and businesses. In APONs, protocol conversion is required for Ethernet. APON has been the traditional choice, especially the optical carrier (OC)–3, which has a bandwidth of 155 megabits per second (Mbps). In APONs, data is transmitted in fixed-length, 53-byte cells with a 48-byte payload. At the time that APON was developed, ATM was considered best suited for multiple protocols, and PON appeared to be the most economical broadband optical solution.

EPON

EPON was developed in 2000 and 2001 through the Ethernet in the First Mile

(EFM) initiative, the standardization effort of the Institute of Electrical and Electronics Engineers (IEEE). EPON was developed because the APON standard proved to be an inadequate solution for the local loop. EPON has a gigabit-per-second (Gbps) bandwidth and yields eight times more bandwidth than APON. While EPON offers greater bandwidth and broader service capabilities at reduced costs than APON, it has a similar fiber infrastructure. It is most effective to transport data, video, and voice traffic via fiber. However, EPON has a point-to-multipoint (PMP) architecture, which eliminates the need for regenerators, amplifiers, and lasers from an outside plant and minimizes the number of lasers required at the central office. Unlike PP fiber-optic technology optimized for metro and long-haul applications, EPONs respond to the needs of the access network in a simpler and more cost-effective manner. EPONs allow service providers to run fiber into the last mile at lower cost in order to provide an efficient, highly scalable, end-to-end fiber-optic network that is easy to manage. Ethernet has evolved significantly in the past 15 years, and everything about it has changed except the frame. *Figure 5* illustrates how Ethernet has evolved for metropolitan-area-network (MAN) and wide-area-network (WAN) applications.

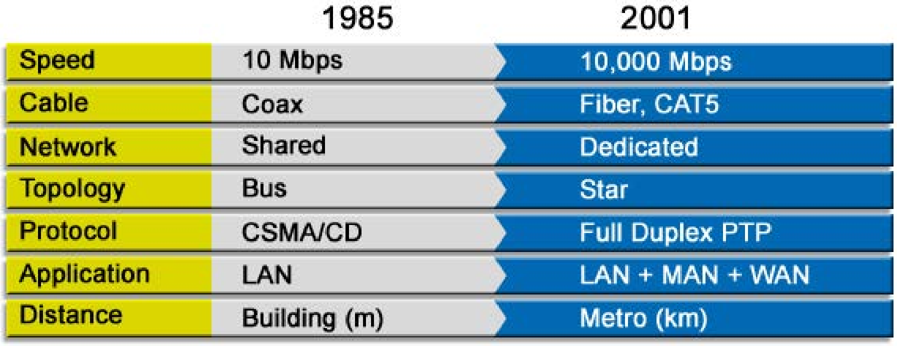

	1985	2001
Speed	10 Mbps	10,000 Mbps
Cable	Coax	Fiber, CAT5
Network	Shared	Dedicated
Topology	Bus	Star
Protocol	CSMA/CD	Full Duplex PTP
Application	LAN	LAN + MAN + WAN
Distance	Building (m)	Metro (km)

Figure 5: *Ethernet Has Evolved for MAN and WAN Applications*

In EPONs, data is transmitted in variable-length packets of up to 1,522 bytes per packet in a 1,500-byte payload (according to the IEEE 802.3 protocol for Ethernet), distinguishing it from APONs. Ten-gigabit Ethernet (10 GbE) products had begun to appear in the market before standards. Ethernet has evolved into a different medium as a dedicated PP, full-duplex-type lens and is beginning to challenge metro networks and local area networks (LAN) quite rapidly. Ethernet also eliminates protocol conversion. An important industry objective is full-service fiber to the home (FTTH) for delivering data, video, and voice over a single platform. The first instances of broadband in the home appeared with the advent of DSL and cable modems. Ethernet acts as an end user for these boxes. The paradigm shift in the industry has occurred with the use of Ethernet as the ubiquitous broadband port with a registered jack (RJ) 45 connector (see *Figure 6*).

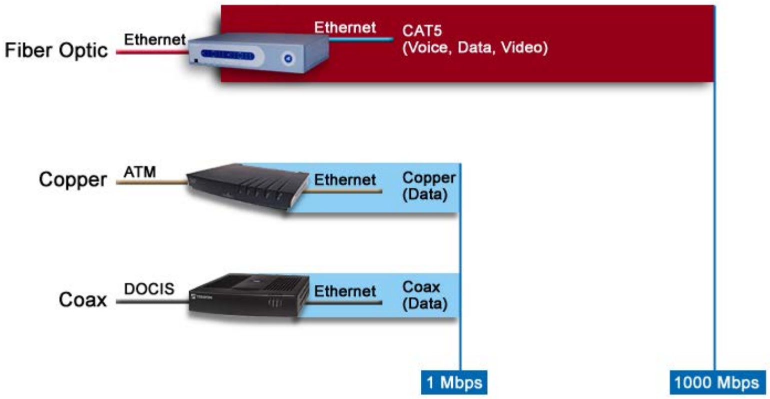

Figure 6: *Ethernet to the Home—the Only Choice*

With regard to these FTTH networks, there must be some consideration for the cost and size of these units. These boxes can be decoupled, leaving a simple box outside and another inside the home.

EPON Downstream

APON can only carry IP traffic by breaking packets into 48-byte segments, which is a time-consuming, complicated, and expensive task. EPON can carry IP traffic more effectively without overhead costs. In EPON the process of transmitting data downstream from the optical line terminal (OLT) to multiple optical network units (ONU) is different from transmitting data upstream from multiple ONUs to the OLT. EPON broadcasts data downstream to all ONUs, which use media access control (MAC) addresses to extract designated packets. *Figure* 7 illustrates the techniques used to manage downstream traffic in an EPON.

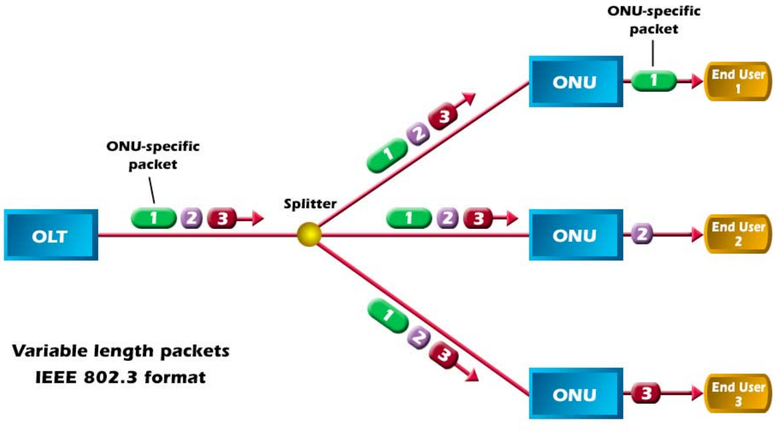

Figure 7: *Ethernet PON Downstream*

Data is broadcast downstream from the OLT to multiple ONUs, and each packet is able to designate the data to the appropriate ONU. The splitter divides the traffic into three separate signals, each one transporting ONU–specific packets. The ONU accepts only those packets that are intended for it and leaves the other packets for other ONUs. The process for EPON downstream is the same as it is for any shared-medium Ethernet LAN.

EPON Upstream

The transmission of data upstream over an EPON is largely the same as downstream transmission, with one key difference: ONUs transmit data upstream to the OLT in Ethernet frames with *ONU−specific time slots* to avoid transmission clashes. Statistical multiplexing is achieved by adjusting the size of a time slot to the amount of available data (see *Figure 8*).

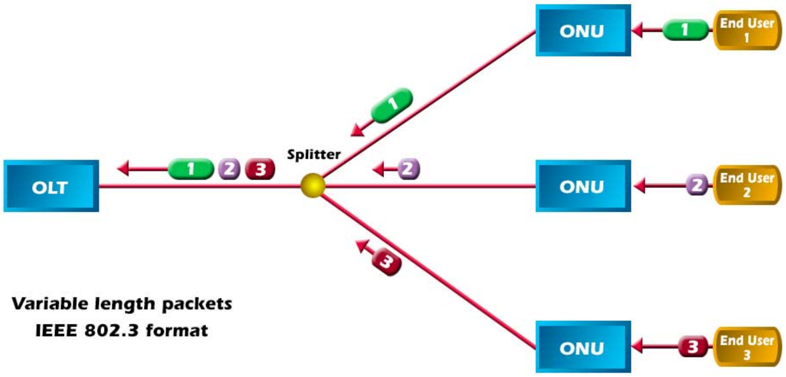

Figure 8: *Ethernet PON Upstream*

The time slots are synchronized to prevent upstream packets from interfering with one another once the data are joined to a common fiber. Each ONU has a separate time slot, the upstream traffic is divided into frames, and each frame is then separated into ONU−specific time slots.

Advantages of EPON for IP Video

EPON provides an Ethernet pipe that transports IP more effectively and at the lowest cost. It leverages off-the-shelf IP and Ethernet component solutions and a steep component cost curve. Ethernet is the most widely deployed network, with approximately 300 million ports worldwide; 80 percent of all networks in the world are Ethernet. LANs are approximately 95 percent Ethernet today, and have virtually eliminated all other networks including fiber distributed data interfaces (FDDI), asynchronous transfer mode (ATM), and token ring. Ethernet is a universal standard, with no variations. Another advantage of EPON for IP video is that it is scalable from LANs to MANs to WANs. Traffic is IP. Byte life begins and ends as IP and Ethernet, and cable and DSL modems have Ethernet interfaces. New competitive local-exchange carriers (CLEC), enterprise local exchange carriers (ELEC), and data local-exchange carriers (DLEC) can start with IP−centric networks. EPON is a plug-and-play environment with fewer arcane parameters. It is a reliable system with structured wiring and an optical plant that has management and troubleshooting tools. As equipment costs are decreasing, workforce costs are not coming down. EPON for IP video serves as a workforce solution, and many LAN technicians already understand and know how to work with Ethernet.

Quality of Service

There are cost and performance advantages to EPON as well, which allow service providers to deliver profitable service over an economical platform. There are a number of techniques that allow EPONs to offer the same reliability, security, and quality of service (QoS) as more expensive synchronous optical network (SONET) and ATM solutions (see *Figure 9*).

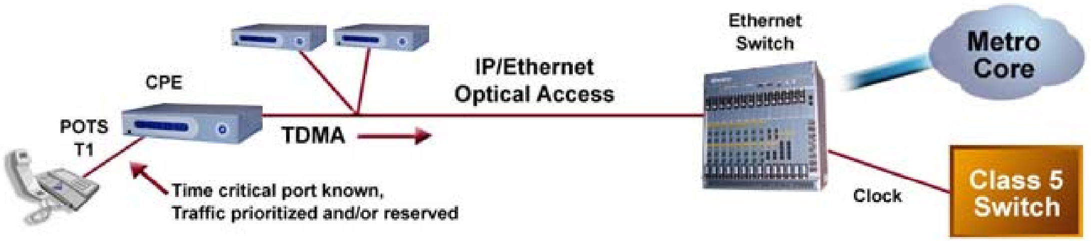

Figure 9: *Quality of Service*

The techniques include guaranteed QoS using type of service (ToS) field and differentiated services (DiffServ). Port-aware bandwidth reservation, queuing techniques, and traffic shaping are also necessary to guarantee QoS. Redundancy and security are other important aspects of QoS. Full system redundancy provides high availability and reliability. Diverse-ring architecture with full redundancy and path protection is another available element of EPON. Another important aspect of QoS is multi-layered security, such as virtual local-area network (VLAN) closed-user groups and support for virtual private network (VPN), Internet protocol security (IPSec), and tunneling.

How IP Video Works

A key element for successful IP video delivery is a high-bandwidth access network, such as GbE PON. High bandwidth is required to support high quality and also multiple channels being delivered to televisions in a household or a multi-dwelling unit. QoS, such as prioritization and bandwidth reservation and management, is another key element for successful IP video delivery. Another important factor is IP video processing, which includes video coding and compression, such as that defined by the moving pictures experts group MPEG–2 standard. A related video-processing standard, MPEG–4, provides features that are similar to those found in personal computer (PC) TV. A standard-definition quality TV (SDTV) will run at 2 to 12 Mbps in comparison, and a high-definition TV (HDTV) will run at 10–60 Mbps in compression. Another element of IP video processing is TV conversion, such as IP set-tops, home gateways and personal computers with traditional analog video output.

IP Multicast

IP multicast provides efficient delivery of selective broadcast IP TV by sending a single data stream (IP video channel) to multiple users simultaneously. IP multicast is more efficient than unicast as a protocol because it conserves bandwidth on the network. It sends only one copy of a channel regardless of the number of requests, and channels are forwarded only according to customer requests. IP multicast leverages broadcast features of PON (see *Figure 10*).

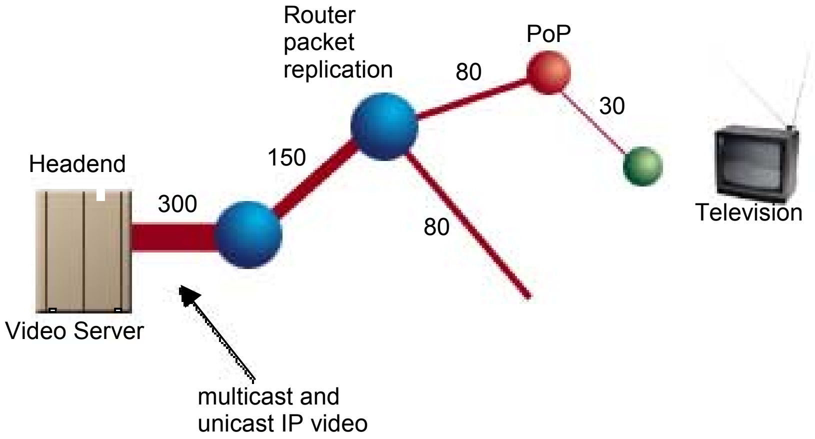

Figure 10: *IP Multicast*

For example, traffic, such as real-time video, comes in over broadcast satellite, and it might already be MPGEG-2 encoded or come in as off-air content. It passes through an encoder, which performs multiplexing and wraps it in IP packets. A stream flow of IP multicast is then released from the IP headend boxes.

End-to-End Architecture for IP Video

Customers are also interested in video on demand, which comes from stored video. Upon request, video servers send IP unicast traffic out to end user subscribers. A converged switch then processes multicast and unicast traffic. An example of end-to-end IP video-network architecture is shown in *Figure 11*.

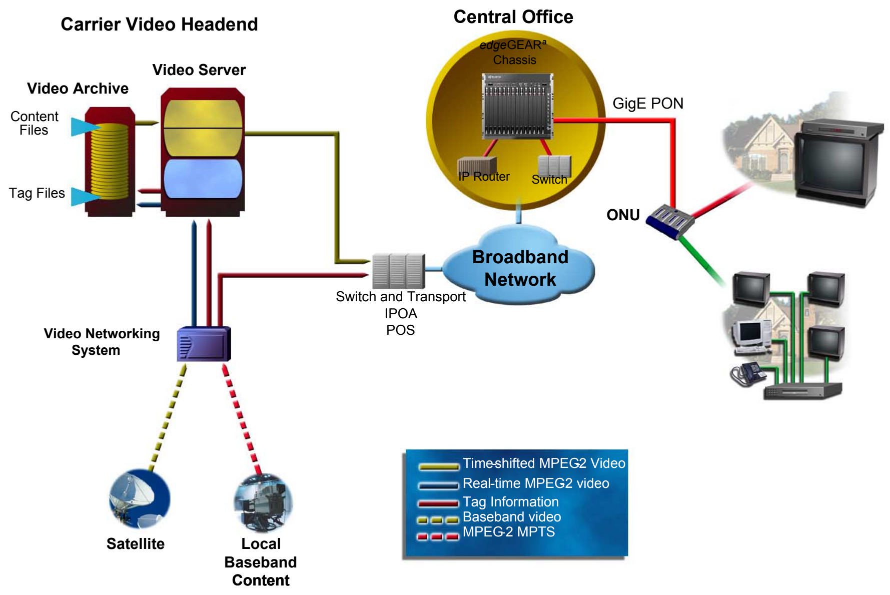

Figure 11: End-to-End IP Video Network Architecture

The Internet TV (ITV) manager is an important component, as it sends the menu and other elements to the set-top box. It interfaces with billing and orchestrates the workings of the IP head-end.

IP Video Head-End Equipment

IP video head-end equipment converts video content to packets and prepares it, adjusting bit rates for distribution across an IP video network. IP video head-end architecture is illustrated in *Figure 12.*

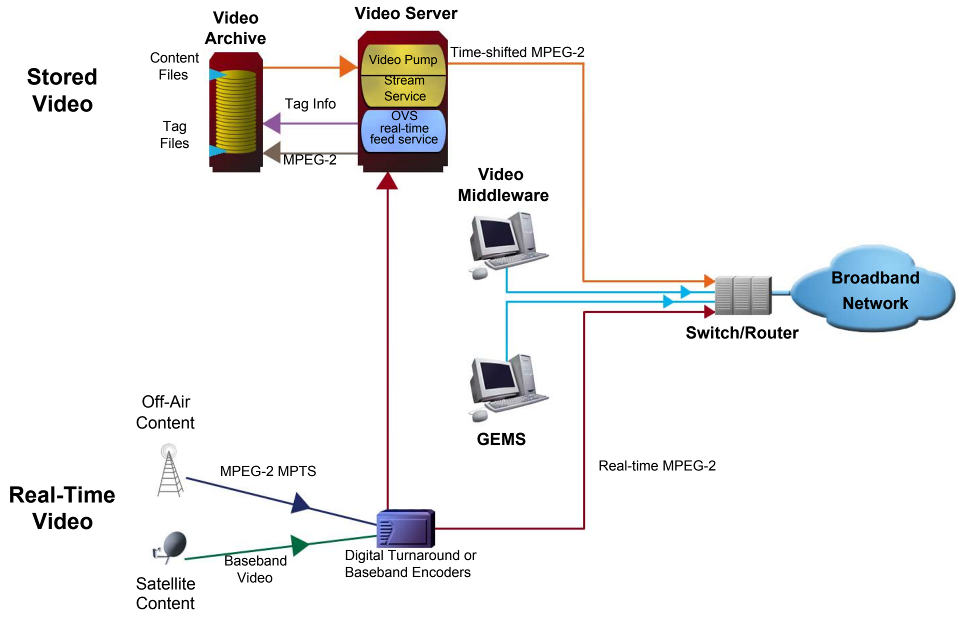

Figure 12: IP Video Head-End Architecture

IP video head-end equipment provides video content injection with real-time IP encoders and video servers, delivering satellite and off-air content to baseband encoders or digital turnaround. The video signals are then converted into individual MPEG packets, and video servers play out unicast video on request. Head-end equipment associates streams with unique IP multicast addresses. It also provides video content management through IP video middleware platforms and video customer premises equipment (CPE), which allows it to interface with video servers and billing systems. Head-end equipment authenticates and authorizes access to video content, and manages set-top boxes. Another important function of head-end equipment is its ability to provide video content distribution via core and metro IP multicast and unicast delivery. IP head-end culminates in the access network, which is a Gigabit Ethernet PON. IP multicast is a point-to-multipoint (PMP) network, which uses EPONs to deliver IP video. One key function of EPON is as a central office chassis that has a multi-service interface to the core WAN. It has a GbE interface to the PON with Layer-2 and Layer-3 switching and routing, and it handles QoS issues, service-level agreements (SLA), and traffic aggregation. Acting as a point to multipoint network, EPON splits optical signals traveling across a single fiber onto multiple fibers. It is a cost effective means of distributing high-bandwidth data, video, and voice traffic across the last mile of the network to customer-located ONUs (see *Figure 13*).

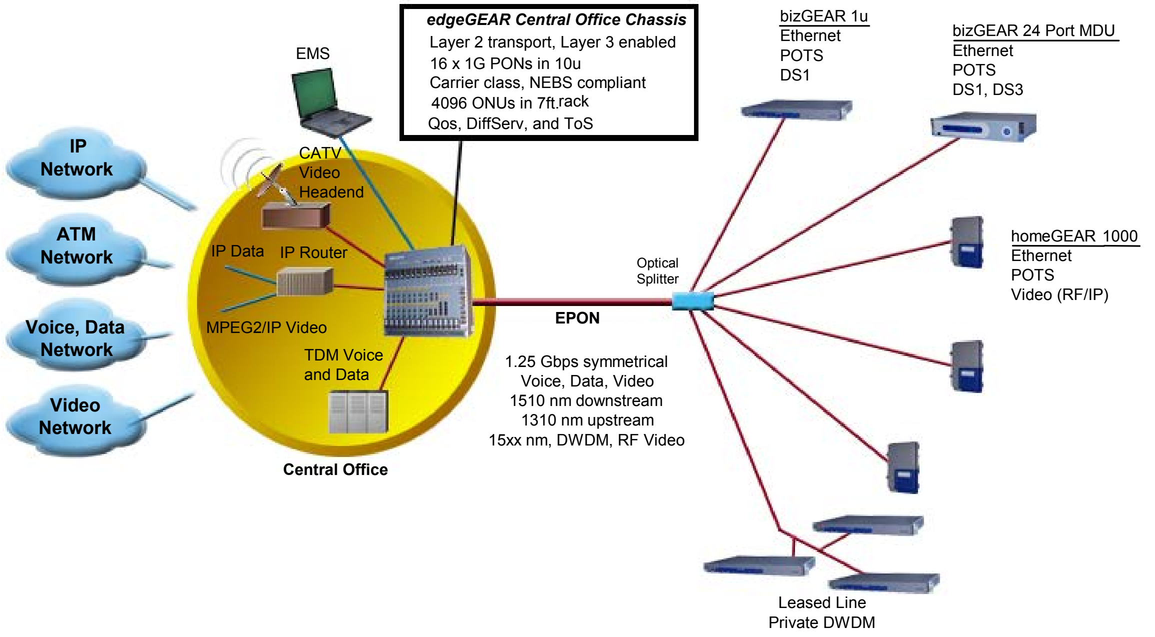

Figure 13: *Access Network and EPON Architecture*

Business and curb applications are possible with these types of networks, and it is possible to have up to 64 users on a single EPON fiber.

CPE

CPE includes ONUs, and a set-top box or a residential gateway. Traffic is received by the ONU, which acts as an interface between the EPON and a customer's data, video, and telephony equipment. It receives traffic in optical format and converts it to the customer's desired format. For example, an EPON ONU may have RJ11 POTs ports for voice, and RJ45 100 Mbps Ethernet ports for Internet access and IP video. The set-top box receives IP packets through a standard 10/100 Base T port and decodes the MPEG–2 stream for a single TV. In a sense, the IP set-top box acts like an Ethernet port, translating TV into a suitable format. A user can have a separate set-top box per TV or install a residential gateway that will connect to multiple TVs. The residential gateway decodes MPEG–2 streams for multiple TV sets and may offer value-added services (VAS). If a user receives 100 multicast streams, the ONU will only let one signal through to the TV, or up to four signals in the case of residential gateways (see *Figure 14*).

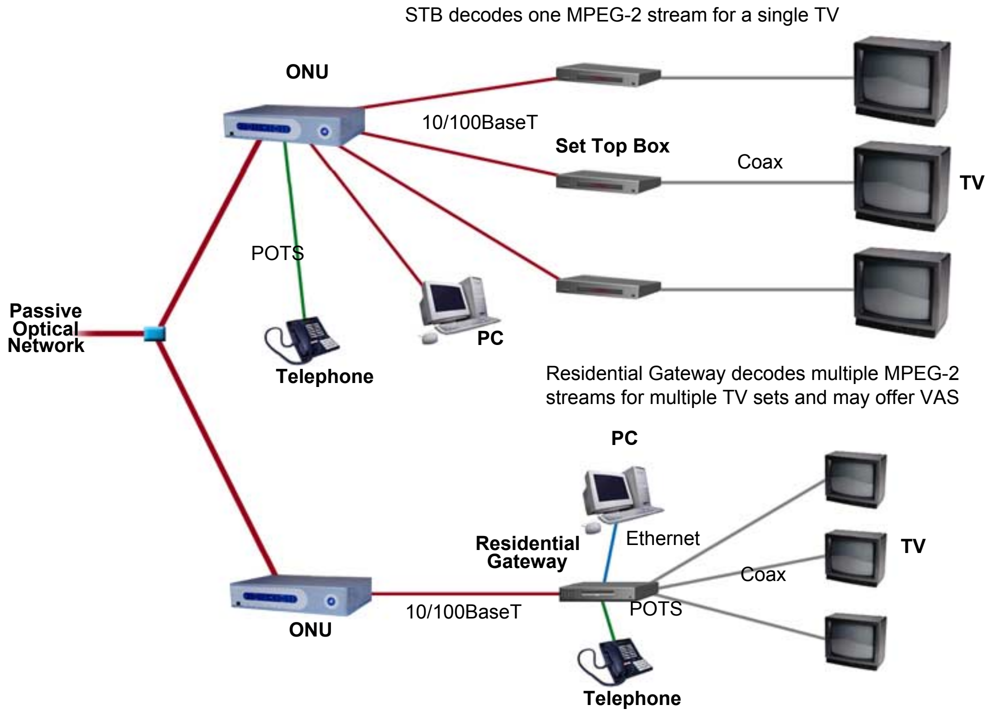
STB decodes one MPEG-2 stream for a single TV
ONU
10/100BaseT
Set Top Box
Coax
TV
POTS
PC
Passive Optical Network
Telephone
Residential Gateway decodes multiple MPEG-2 streams for multiple TV sets and may offer VAS
PC
Residential Gateway
Ethernet
TV
Coax
10/100BaseT
POTS
ONU
Telephone

Figure 14: *IP Video CPE*

The end user's experience is similar to watching CATV, as the IP streams, at around 4 Mbps, are indistinguishable from analog video.

Summary

IP format treats video as data and supports any type of TV service. It fits into a unified, single-network solution, though it requires significant bandwidth and QoS techniques. IP video enables integrated and interactive Web applications. EPON is an IP–based Ethernet broadband network, which delivers 1,000 times the speed of DSL or cable modems. EPON is available at a low infrastructure maintenance cost. It is scalable and designed to be a plug-and-play network. *Figure 15* provides an IP video demonstration to further illustrate the equipment being used.

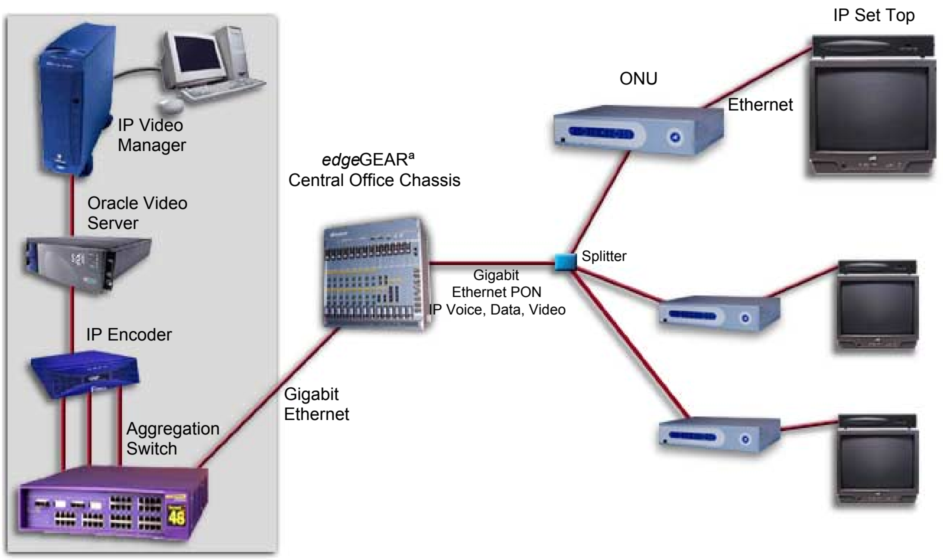

Figure 15: *IP Video over EPON Lab Demo*

The equipment includes IP set-top boxes, the ONU, and a box that sits in the local exchange. Servers store video, and encoders take real-time traffic and code it as IP. The aggregation switch gathers it together into a single GbE port and delivers it over the metro access. EPONs systems are in deployment/trials; a significant number of the trials underway will quickly migrate into full deployment. Many companies view this optical IP Ethernet architecture as the most effective means for transporting voice, data, and IP video services over a single network.

MANIPULATING MEMORY IN SET-TOP BOXES: GETTING THE MOST FOR YOUR MONEY

Haig Krakirian
Pioneer Digital Technologies, Inc.

Abstract

Everyone knows that memory size in a set-top box does matter. What most people don't know is that memory configuration, type, and architecture are key factors in deploying successful digital services. Memory capability within a set-top is a complicated and varied proposition. Many customers of set-top boxes may not fully comprehend the critical deployment of multiple memory types and advanced application requirements when evaluating purchasing decisions. Quite simply, memory configuration in the box affects the performance and the price. We believe the decision-makers need to understand when the technology is worth their dollar.

We propose that readers may want to learn a bit more about memory models offered in set-tops before they conclude that the only difference in set-top boxes is the price. We intend to discuss the ways in which memory architecture can be designed for optimal efficiency and economy in a set-top box, why it matters to today's customer, and how it can affect tomorrow's service offerings.

DON'T FORGET ABOUT MEMORY

In the computing industry, memory capacity has always been driven by software demands. As software becomes more complex, the need for greater and greater amounts of memory escalates. This trend is similar but even more pronounced in the set-top space, however, because in addition to processing digital video, set-tops must also process ITV applications which are rich in multimedia content.

While there are many parallels between the technological advancements of the set-top and computer industries, the pace of development has lagged on the set-top side. Cost is a primary reason for this. MSOs are always on the lookout for that "sweet spot" between price and functionality. In order to keep set-top prices as low as possible, memory requirements traditionally have been based on the types of applications available at the time of development. For many MSOs, therefore, set-tops have been geared toward MPEG decoding and navigation software (including program guides), rather than for more advanced interactive applications. The only business model available to MSOs has been that of video delivery, making it hard to justify the extra cost for memory when there is no clear indication that there would be a return on that investment.

Meanwhile, software vendors who were hoping to find a dynamic new market in set-tops have been hesitant to commit resources to that market because they are not seeing adequate memory on which to base a decent application. This dearth of applications has further diminished the desire of MSOs to invest in memory because there simply is not enough to offer customers. In short, we're stuck in a classic chicken and egg scenario.

Fortunately, today's temporary hurdles are forcing the bar to be raised for both the application developer and the MSO. Rich content requires more color and animation, but it is also memory intensive. As in the computer industry, it will take cooperation and coordination between software developers and hardware manufacturers to bring about the compelling, eye-catching content that digital cable networks are capable of.

MEMORY OPTIONS

In order to understand the evolution of memory in set-tops, it is important to understand what types of memory are available and their functions. Today’s memory types come in many forms and each has its own characteristics that can affect and improve the performance of set-top tasks. For set-top boxes, two primary types of memory are used: Random Access Memory (RAM), and Flash memory.

Random Access Memory or RAM, is memory that is available to the system processor when the set-top is booted up and running. RAM is essential to run and execute applications, but RAM memory loses its contents when the set-top is turned off, or if the content is not refreshed by an external charge. Thus, RAM is considered "readable/writable," meaning information, such as applications, can be written to, and read from its silicon, but not stored (without power), making it "volatile."

However, advances in RAM have blurred traditional definitions. The most common types of RAM today are Static RAM (SRAM) and Dynamic RAM (DRAM).

SRAM is faster than DRAM and accordingly more expensive. Because of the way DRAM is constructed, with individual cells composed of a transistor and capacitor, the contents must be continually refreshed by an external circuit to be retained in memory. SRAM requires no refreshing and retains data.

Because it's cheaper, DRAM is typically used in larger sizes. There are several different types of DRAM, including Synchronous DRAM (SDRAM), Double Data Rate RAM (DDR-RAM) and Extended Data Out RAM (EDO-RAM).

A large percentage of set-tops use SRAM as a digital channel receiving buffer, or the "interleaving buffer." SRAM serves this purpose well because the set-top processor needs fast access to RAM when processing video, but doesn't need to process a huge amount of data at once. On the other hand, DRAM is used for application execution, as a graphics rendering buffer and an MPEG decoding buffer.

RAM AND INTERACTIVITY

DRAM is quickly becoming the engine that caches and drives interactive applications in the set-top. Interactive TV applications are stored on servers and are broadcast, or streamed, over the network on specific channel frequencies. Set-tops "listen in" to those streams and begin caching, or saving, them in DRAM. Once the applications are fully loaded in DRAM, they are executed when a subscriber chooses to run that application.

One existing channel navigation paradigm is where each channel is an application. An email client is an example of an application that can be assigned to a channel allowing a user to simply go to a channel to check their email. In this case, when a user surfs over that channel, the application is loaded and executed out of DRAM.

Besides loading and executing applications, DRAM is used for loading data files that are used by those applications. A primary example is program guide data, which, because it changes daily, is best stored in DRAM. The set-top reads the program guide data files in much the same way it watches for interactive applications, and stores them in DRAM.

BACK IN A FLASH

The basic characteristic of flash memory is its "non-volatile" nature, meaning it can retain information without a backup battery or electric charge. When a box is turned off, flash memory will keep its data and content available for use when the box is turned back on. At the same time, flash memory can be re-programmed and erased, although a high voltage is needed to erase or make changes to the data stored in flash. Flash memory, then, is "persistent" in nature.

There are two types of Flash ROM, NAND type and NOR type. NAND is typically larger in memory size, but has slower access speeds.

Random access speeds of NOR can be from 70 ns to 100 ns per byte, and NAND can take 25 micro seconds to read 512 bytes. In contrast, page mode access, which allows quick read access from a specific range of memory, is generally faster in NAND. NOR can access at 20 ns to 25 ns per byte, but NAND can access at 50 ns per 512 byte. So, NOR can be faster than NAND for byte access, but NAND can be faster for bulk data access. Due to the nature of their individual capabilities, they are utilized for different tasks. NAND type is used for data storage, while NOR type is used for quick random access functions, such as program storage and execution.

Both NOR and NAND use FG (Floating Gate) cell technology. FG structure memory, especially in NOR type, may limit the potential growth of memory size because the memory cell is so complex. This is why Flash ROM is such an expensive element of a memory configuration. Fortunately a new type of Flash ROM called MONOS (metal oxide nitride silicon) type was recently introduced which allows for a more simplified cell structure, and therefore more potential for increases in memory size. Many manufacturers are betting on this technology to provide a high-memory, low-cost Flash ROM solution.

In the absence of hard drives in today's set-tops, flash memory has become the repository of the core applications of a cable video service, which include firmware, middleware, and a so-called "resident" application (comprised of a program guide, parental control, pay-per-view services, the core channel surfing application, etc.). In the event headend application servers go offline for any reason, subscribers will still be able to watch TV and change channels.

Over time, operators may determine that Web browsing is a critical application to their service offering, and placing the actual Web browsing application in flash memory will become more of an attractive option. Turning to this option, of course, means that more flash memory will be needed, beyond today's 4 megabyte (MB) standard.

MEMORY ARCHITECTURES

In addition to the types of actual memory available to set-top makers, there is a choice between memory architectures. A unified memory architecture uses the same memory to process video, graphics and other digital information. Unified has the advantage of providing a larger pool of available memory, depending on how many simultaneous functions the box is performing. If there is no MPEG video being processed, for example, there is more available memory for decoding graphics.

For applications developers, unified memory can be a drawback in that there is no easy way to know exactly how much memory will be available on a given set-top at a given time. This makes it tough to add the kinds of rich features desired by consumers because they push the envelope in memory usage. Unified memory can also be a drawback for the user who may start to see sluggish performance or scrolling hiccups if digital MPEG video is tuned, which takes priority in the decoding hierarchy.

A dual memory architecture solves the latter problem by providing graphics and video with their own dedicated memory. While this makes it impossible to take advantage of unused memory in, say, the video side of the box, it does allow developers a fixed memory limit when deciding what features to add to a particular application. The disadvantage to the set-top maker is cost. Duel memory cores require separate busses.

THE MARKET TODAY

The majority of advanced digital set-top boxes deployed today have a memory configuration of 4 MB of Flash and 8 MB of DRAM. For the MSOs who are deploying these boxes, the challenge is to utilize them to their fullest capacity.

With a 4/8 configuration, the operating system and core resident applications – such as the channel guide, pay-per-view, and video-on-demand – eat up about 5 MB of DRAM. That only leaves 3 MB of DRAM for the more advanced interactive applications, which is not a lot of memory when you're talking about advanced interactivity.

Compression techniques and intelligent caching techniques allow applications to take full advantage of a 3 MB footprint. However, decompressing the application data files can put a strain on CPU processing and bog the system down making a set-top unresponsive to user interaction. In addition, when you consider that any form of rich graphic content will require multiple 640x480 screens of at least 600 KB each, it becomes clear that 3 MB of memory can be eaten up in no time.

To be sure, there are a number of compelling applications that can run on 3 MB of memory. The question is which application should an MSO choose? MSOs are understandably hesitant to dedicate all of their remaining memory resources available in a set-top to just one application when their goal has always been to offer a host of interactive services as opposed to just one.

VOD IS KEY

Relying on the built-in MPEG video decoding capabilities of today's digital set-tops is most likely the solution to the limited memory problem. Interactive applications can leverage the capabilities of the current base of deployed digital set-tops to deliver rich and compelling content. After all, video is still king in the cable universe, and any new application will have to compete with it.

As long as cable operators are continuing their plant upgrades to provide streaming capability for every set-top, VOD will be the core technology for ITV applications. The cable infrastructure can be utilized to present rich, colorful video and animated content, rather than the more static web content. Best of all, VOD-streaming solutions do not require extra memory because you are using the video-decoding side of the box.

E-commerce, t-commerce, informational applications and other offerings would become more video-centric, rather than just text and pictures. A user could click through a few introductory pages and then get hit with a video stream. The PC has always had trouble with full-motion video, even under MPEG. The set-top has full-motion video built right in, so why not take advantage of that?

Most of the applications to-date have been data-centric, such as stock ticker information, sports scores and utility applications to manage programming options. There will have to be a re-thinking of the type of content presented and innovative re-use of graphics so as to offer a richer interactive experience from a minimal memory footprint. There is also a need to leverage the network more efficiently and manage content more wisely to cut down on the number of times the set-top has to access the information.

JUST AROUND THE CORNER

Clearly, then, the more advanced interactive content that requires larger chunks of RAM will have to wait for the next-generation boxes. But how long will it take? That is the million dollar question, but there are a number of factors that could propel the market a lot faster than most people think. High-definition television (HDTV) and personal video recorders (PVRs) will push the set-top to evolve much quicker than if current applications alone were leading the charge.

With PVRs, the caching of content requires a hard disk inside the box, which is memory that can be taken advantage of. The PVR may be the wedge that entices consumers into thinking of their set-tops as more than just dumb video terminals. After all, without a disk drive, web browsing on PC's would not be nearly as compelling as it is today.

HDTV could be another factor that propels additional memory in set-tops, mainly because it requires 32 MB of memory to decode. HD can be delivered via unified or dual memory architectures. Since most set-tops utilize unified memory, much of that capacity will not be used unless you are decoding HD.

Certainly, increasing set-top memory size is crucial to deploying advanced applications. While it is understood that memory is a scarce resource in set-tops today, there are a number of exciting developments right around the corner that may indeed accelerate the deployment of set-tops with expanded memory.

The key to success in the near term for MSOs is to understand current memory limitations and ways to leverage VOD technologies to offer enticing interactive services. It is also just as important to build a coherent roadmap for future interactive services based on the capabilities of next-generation set-tops just around the corner.

Haig Krakirian,
VP Software Engineering
Pioneer Digital Technologies, Inc.
2210 West Olive Avenue, 2nd Floor
Burbank, CA 91506
T 818.295.6628
F 818.295.6797
haig@pioneerdigital.com

MPEG STANDARDS EVOLUTION AND IMPACT ON CABLE

Dr. Paul Moroney, Dr. Ajay Luthra
Motorola Broadband Communications Sector

Abstract

The MPEG standards process has evolved from its beginnings in 1988 through today to cover a wide range of multimedia delivery technology. The entire phenomenon of digital television and the standards involved have affected our culture through increased content choices for the consumer, increased competition, and new types of products. Cable delivery has seen a huge impact, in particular from MPEG-2 based digital television. If we can extrapolate from the evolution of the MPEG process, more change is coming!

INTRODUCTION

No one would dispute that the cable industry has been dramatically impacted by the transition to digital television and related services, whether one considers the advent of effective DBS services, or the distribution of digital television to cable headends, or the distribution of digital services through the cable plant to consumers. Although many factors have played a role in the process, certainly the development of true standards was and continues to be a true driver. The ISO/IEC sponsored Motion Picture Experts Group [MPEG] has been the primary organization for the formation of the basic multimedia source coding standards. This paper examines where MPEG has been, and where it is going, to build a vision of the future of digital services over cable. Particular attention is paid to video compression.

The MPEG process began in October of 1988, under the very effective leadership of its convener, Leonardo Chiariglione, who continues in that role to this day. MPEG-1 [1] targeted stored content up to a 1.5 Mbps rate, matching the parameters necessary to store on CDs. MPEG-1 addressed audio compression, video compression, and a storage-oriented systems layer to combine them with other higher-level information, and the resulting standard earned an Emmy award. MPEG-2 [2] shifted the focus to entertainment television, added more video compression tools, multi-channel audio, and a new systems layer more tuned to the need for a transport definition for broadcast. This standard has been deployed very successfully worldwide, and earned an unprecedented second Emmy for MPEG. The designation "MPEG-3" had been set aside for HDTV extensions, but MPEG-2 worked so well for this application that MPEG-3 was never needed.

MPEG-4 [3] began slowly, while MPEG-2 based systems were deployed, and has developed into a true next generation multimedia system. MPEG-4 covers a much broader scope than its predecessors, and achieves dramatic improvements in the underlying source coding technology as well.

MPEG has also begun efforts in two related areas, not focused per se on multimedia compression. MPEG-7 [4] is targeted as multimedia search and retrieval, enabling the creation of search engines for such content. MPEG-21 [5] (the 21 suffix representing the 21st century) addresses the goal of true universal multimedia access. This group focuses on defining a multimedia framework so that broad classes of content can be created, delivered, and consumed in an interoperable manner.

MPEG-1 AND MPEG-2

MPEG-1 and MPEG-2 are defined in three main parts, addressing video source coding, audio source coding, and a systems layer.

Video

MPEG video source coding is designed to remove redundancy in video content both frame to frame, and within a frame. Algorithms that achieve these goals without degradation in the reconstructed video images are "lossless," that is, reversible. Algorithms that sacrifice quality in the interests of bit rate reduction are termed "lossy." Lossy algorithms attempt to hide errors, or artifacts, based upon models of human video perception. For example, it is more difficult to detect artifacts under high motion. It is also true that human vision chroma resolution is lower than human vision luminance resolution. Compression systems, such as MPEG, that seek to address a wide range of applications typically employ both lossless and lossy algorithms.

MPEG video compression at its base employs transformation, quantization, and variable length coding (VLC), in that order. Pixel domain representations of video frames are transformed through an 8 by 8 Discrete Cosine Transform (DCT) to a frequency domain representation. The resulting spatial frequency coefficients are scanned in a zigzag manner, and quantized. The resulting sequence of amplitudes and runs of zero values are then Huffman coded. The bits generated are grouped under a syntax organized roughly at the *sequence* level (a number of consecutive frames), the *picture* (typically frame) level, the *slice* level (rows of blocks), the *macroblock* level (4 luminance and 2 chrominance blocks) and the DCT block level, including header information appropriate to each.

The last key aspect is the inclusion of motion estimation and compensation, where the above approach applies to only the difference between a current frame (macroblock) region, and its best possible match in the prior frame. The offsets representing the best match, to the nearest x and y half-pixel, become motion vectors for the region, and are included in as header information.

Figure 1 shows the most basic block diagram of this compression process. Note the presence of the decode (decompress) loop, so that the compressor will be able to transform the pixel differences required for motion compensation.

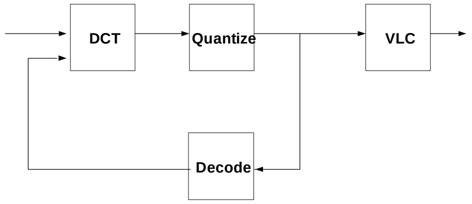

Figure 1: Basic MPEG Compressor

Although not shown in the Figure above, MPEG also allows motion estimation from a *future* frame, or from both the previous and future frames, in combination (See Figure 2). In MPEG parlance, a frame without motion compensation is an *I* frame, one with compensation from only the previous frame is a *P* frame, and one with bi-directional motion is a *B* frame. Note that in order to decode a B frame, coded frames must be sent out of (display) order. A decoder must have decoded the future frame before that frame can be used to "anchor" any frame decoded with bi-directional motion.

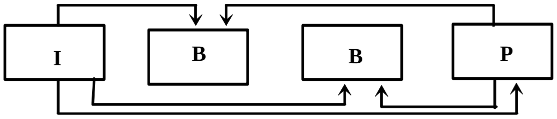

Figure 2: B Frame Prediction

The MPEG-2 standard introduced various enhancements, but two major changes were the addition of frame/field mode and direct support for film content. In order to compress interlaced video signals intended to be ultimately displayed on interlaced televisions, (see Figure 3) efficiency is enhanced by allowing the 8 by 8 pixels to be assembled from alternate (interlaced) fields. In fact, throughout the interlaced content, some regions compress better in this fashion [frame mode] and some compress better by taking the 8 by 8 pixel block from each field separately! Thus the standard allows either approach, signaled by header bits.

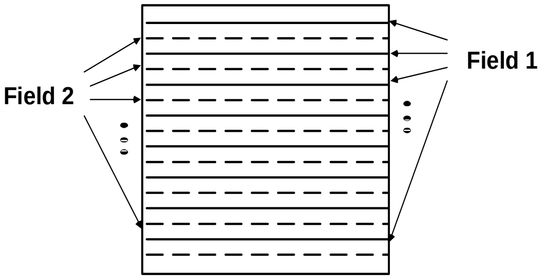

Figure 3: Interlaced Block Pair

For film content carried in NTSC television signals, common practice has been to convert the 24 frames per second progressive film to the NTSC 29.97 frames per second interlace (59.94 fields per second) through a *telecine* process known as 3:2 pull down. The first film frame is divided into two fields, and the first field is repeated in the NTSC signal as a "field 3". Frame 2 of the film becomes "field 4" and "field 5," without any repeat. Film frame 3 becomes fields 6 and 7, and field 6 is repeated as a field 8. To complete the cycle, film frame 4 becomes NTSC field 9 and 10. MPEG allows compressors to drop the duplicate fields, and code the film frames directly. An overhead bit called "repeat field" informs the decompressor, or decoder, how to restore NTSC for display.

Given the descriptions above, one can deduce the philosophy adopted by the MPEG planners. The maximal benefit to equipment suppliers, and consumers, taken together, was determined to be achieved through a standard that defines decoders to the extent required to be able to receive all signals that are "MPEG compliant" according to the standard syntax. Thus all aspects of the decoder are defined, with the exception of error recovery and the specifics of the display generation.

The compression process, however, provides substantial room for creativity and invention, within the context of the standard. Any compressor can be built, so long as an MPEG compliant decoder can decode the signals produced. Compression is, thus, far more of an art form, than decompression. Consider the various degrees of freedom. Which frame regions are to be predicted, and how? For predicted regions, what are the best motion vectors? Which frames do you select as I frames, or which slices are not predicted at all, so that decompression equipment can acquire a signal after a channel change, or after an error? How much should any given block or region coefficients be quantized to minimize perceptual loss yet achieve some desired net output bit rate (*rate control*)? How does the equipment estimate such loss? Which regions are best coded as fields, rather than frames? What is the best way to scan the coefficients (there is more than one option provided in the standard)?

MPEG-2 introduced the concept of *profiles* and *levels*, to allow the broad range of compression tools and bit rates to be targeted at classes of applications. This allowed products to be *compliant* at only a specific profile and level; thus low cost consumer devices were possible. No one device was burdened with implementing all the tools of MPEG, at all bit rates.

Audio

MPEG audio compression bears several similarities to the video compression described above. The overall compression approach involves lossless techniques, as well as lossy quantization. As with video, the approach is designed based upon a well-understood model of human auditory processing, so that errors are hidden as much as possible. Specifically, the ear tends to act as a multi-band analyzer. Within any *critical band*, distortion can be masked if there is enough signal content present. Thus quantization should be applied to these bands, according to this masking phenomenon.

MPEG-1 defined three layers of audio compression algorithms, with increasing complexity and performance. Layer 2 is a sub-band compression approach known as *MUSICAM*, and is broadly deployed in many existing MPEG-2 based systems. Layer 3 adds more compression tools, and performs better at a given bit rate. This layer carries the abbreviated name *MP3*, and has seen extensive recent use over the Internet for audio track distribution.

The MPEG-2 audio standard introduced a backward compatible multi-channel surround version of MPEG-1 defined audio. Thus advanced receivers could decode surround, while earlier equipment could process the stereo audio defined for MPEG-1. As this approach is not completely optimal for multi-channel surround, the MPEG audio sub-group also defined a new non-backward compatible algorithm called *Advanced Audio Coding (AAC)*. AAC provides the best quality overall audio coding approach, at added complexity, both for multi-channel surround and for regular stereo audio. With AAC, 64 kbps stereo audio is generally perceived as excellent.

Systems

MPEG defined a systems layer to combine the various audio and video streams into a higher-level *program*, and to provide additional information, such as time references. In both MPEG-1 and MPEG-2, the compressed audio and video are termed *elementary streams*. These streams reflect the syntax defined for each media type. The systems layer produces a *Packetized Elementary Stream* (*PES*) from each elementary stream, adding header information such as time markers (*timestamps*), and a type/label description. The resulting PES packets are then combined in one of two structures to build a higher-level construct.

In both MPEG-1 and MPEG-2, PES packets can be combined into a *program stream*, which is more targeted at storage applications. Program streams are characterized by large packet sizes, such as might be appropriate for a sector of a storage medium. The format has little or no error resilience, again, such as would be well matched to a file system. Program streams include a directory structure that defines various general information about the program, and describes all the component packetized elementary streams within the program stream.

Transport streams are MPEG-2 specific, intended for the transport of multiplexes of multimedia programs. MPEG-2 transport divides the component PES into 184 byte segments, prepended with a 4 byte header. Each resulting 188-byte transport packet thus includes data from a single PES, and a multiplex of such packets would mix together a number of programs. Content within a multiplex is grouped by a set of identifiers and tables within the multiplex.

Each packet header contains a 13-bit *Packet ID* (*PID*) for reference. The *Program Map Table* (*PMT*) lists the PIDs associated with that program, along with a type value (audio, video, etc.), and optional descriptive information. The PMT is carried in transport packets with its own PID. A single *Program Association Table (PAT)* lists all the programs contained in the multiplex, and includes a pointer to each PMT PID. The PAT is defined to be carried in packets of PID 0.

MPEG-2 transport also includes mechanisms to reference the conditional access data that applies to each program, and the 4 byte packet header includes encryption related bits, sequence numbers, and a provision to include an adaptation field within the body of the packet. Among other functions, the adaptation field must be present often enough to carry a *Program Clock Reference* (*PCR*) for the program. PCR reception allows an MPEG decompressor to rebuild system time, and manage buffers for proper decode and display. Features of this type allow MPEG-2 to support broadcast quality reconstruction of video, which was not an emphasis of MPEG-1.

MPEG-4

The MPEG-4 process initially focused on developing a new video coding standard for low bit rate coding. During the course of development of the standard, it was realized that a much broader scope was appropriate to support evolving new classes of multimedia applications, including those suited to the Internet. Thus MPEG-4 not only describes new, improved video coding tools, but a much broader multimedia support for these new applications.

First, MPEG-4 was targeted at a broad range of compression resolutions and bit rates. MPEG-4 was designed to be as efficient, or better, than MPEG-1 and the H.263 family of standards at very low bit rates, as efficient or better than MPEG-2 at mid to high bit rates, and so on. MPEG-4 syntax overhead had to be flexible enough to allow efficient operation for all these rates.

Second, to satisfy the needs of such broad applications it was considered important to have independent representation of each media type, e.g. video, graphics, images, speech, animation, synthetic audio, textures and text. This provides much better coding efficiency, since converting media types like text and graphics to video and compressing them as video causes a great loss in quality. (Just consider the rolling text credits at the end of an MPEG-2 heavily compressed movie!) Therefore, in addition to developing *natural* video coding tools, MPEG-4 also developed *Synthetic-Natural Hybrid Coding (SNHC)*, *Texture Coding* and *Sprite Coding* tools.

Third, these broader applications and rich media types drove MPEG-4 to the concept of objects. Rather than considering video as a progression of rectangular frames, scenes are now built from a composition of arbitrarily shaped objects. Certainly this shifts complexity to the decompressor/receiver device in a system, since this device must now compose scenes from these objects, but the potential for new service offerings that this enables is incredible. Not overstated, MPEG-4 allows the content creator to shift a (controlled) portion of the creative process, a portion of the studio, to the user's device.

Hyperlinking of these objects to other objects, such as parameters and text descriptions, offers the cable industry for the first time a true standards based approach to interactivity, with multimedia features only seen today in movie theaters.

MPEG-4 Natural Video Coding

At a very high level, MPEG-4 video coding, like MPEG-2, is motion compensated DCT based. However, MPEG-4 video coding in general includes more compression tools, and more complex tools, than its predecessors, as such tools are now cost effective to implement.

Coded frames can be of four types – I, P, B and *S*. As MPEG-4 went beyond the concept of coding rectangular video frames to arbitrarily shaped video objects, I, P and B frames are called I, P and B *Video Object Planes (VOPs)*. For video in a rectangular window, VOPs are the same as frames. Similar to MPEG-2, I-VOPs are compressed without any motion compensation, P-VOPs are compressed with motion compensation based on the past I or P-VOPs, and B-VOPs are compressed with motion compensation based on both past and future I or P-VOPs. In addition to I, P and B-VOPs, where the motion compensation is block based, MPEG-4 also allows S-VOPs with *Global Motion Compensation (GMC*), where a single global motion vector is estimated between two VOPs. It is helpful mainly for video sequences with large similar motion across the picture, such as when a camera is panning or zooming.

As mentioned above, MPEG-4 also developed tools to represent and compress arbitrary shaped video objects. Once pictures are broken into multiple objects, each object is compressed independently. Objects can thus have different quality, "frame" rate, and bit rate. In addition, two types of arbitrarily shaped video coding capabilities are defined: *binary* and *gray scale*. In binary shaped coding, objects can be composed to be either in the foreground or in the background. In gray scale shape coding, objects can be composed with 256 levels of transparency, or blending. With object representation, an encoder (now much more than a compressor) needs to have the capability to describe how a particular scene is composed, based on the multiple objects within it, and a receiver device needs to have the capability to recompose the scene in addition to simply decoding the objects and presenting them. Furthermore, the term "scene" is generalized to include all the media types in the content, not only video objects. To facilitate the capability of efficiently describing and sending the dynamically changing scene, a *BInary Format for Scene (BIFS)* description was also developed.

In addition to MPEG-2 type temporal and spatial scalability, a new type of scalability, *Fine Granularity Scalability (FGS)*, is also defined in MPEG-4. Scalability tools allow video to be compressed in multiple layers. To decode a picture, one does not need to receive all the layers; however, the picture quality of the decoded picture can be incrementally improved by decoding more layers. In FGS, as the name suggests, many multiple layers with small incremental numbers of bits can be sent. This technique is very helpful for adapting the compression rate and picture rate to the time varying available bandwidth of a network like the Internet.

MPEG-4 also defined new error resilience tools, which allow bit streams to withstand relatively large bit loss. As an example, the VLC tables can be decoded in both a forward and a reverse direction! All these tools are contained in Part 2 of MPEG-4. As this part describes the standard for coding more than natural video, it is called MPEG-4 *Visual* instead of MPEG-4 Video.

Profiles and Levels

MPEG-4 contains a plethora of compression tools that are useful for many different applications. As with MPEG-2, all the tools are not necessary or required for all the applications. Thus MPEG-4 has defined several profiles and levels to define interoperability points where a good compromise is made between implementation complexity and the needs of certain classes of applications. To achieve interoperability, decoders are required to implement all the tools in a given profile. Each profile then has multiple levels. Levels are defined by keeping in mind how much processing power is needed for real time decoding. They are mainly defined according to the sizes of the pictures, such as *QCIF* (176x144), *CIF* (352x288), *HHR* (360x576) and full resolution (720x576). Three main application sets, used for defining profiles related to the coding of natural video, were: (1) Wireless and Wireline Video Phone and Conferencing, (2) Streaming Video and (3) Interactive TV / Web enabled multimedia.

The most commonly known and used profile is the *Simple Profile (SP*). SP was defined for two-way video communication and very low complexity receivers, such as wireless videophones. The tools were selected by giving high priority to low delay, low complexity and error resilience. SP includes very basic coding tools including I-VOPs, P-VOPs, *AC/DC* prediction, *Unrestricted Motion Vectors (UMV*) and error resilience tools such as *Data Partitioning, Slice Resynchronization* and *Reversible VLC.* This profile is a subset of all other video profiles, that is, a decoder compliant to any other video profile is also capable of decoding SP. Due to its simplicity and the lack of any other profile for rectangular video, this also became the most commonly used profile for streaming video. However, the coding efficiency of this profile is low. In addition, levels are defined only up to CIF size pictures.

To make available a profile that is more suitable for streaming video applications over the Internet and has higher coding efficiency, *Advanced Simple Profile (ASP)* was defined [6, 7]. As a delay on the order of hundreds of milliseconds is not an issue for those applications, and targeted platforms have higher processing power, ASP coding tools include B-VOPs, GMC, *Quarter Pixel Interpolation*, and *Interlaced Video* tools (in addition to the SP tools). Streaming video applications generally use only the rectangular video window. Therefore, to control the complexity of implementation, shape-coding tools are not used in ASP. ASP thus provides the highest video coding efficiency (for rectangular video) among all the profiles in Part 2, significantly more than SP. Figure 4 shows a comparison of the performance of ASP with SP, for the case of the Stephan sequence and CIF resolution. The x-axis represents bit rate, and the y-axis represents the *Peak Signal to Noise Ratio (PSNR),* a measure of the distortion (error) in the decoded pictures (High PSNR means better quality.) Reference [8] provides additional

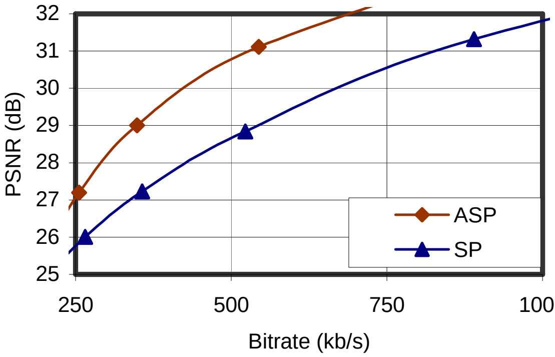

Figure 4: ASP vs. SP Performance, Stefan sequence (CIF)

performance data. In general, ASP provides a good tradeoff between video quality and coding complexity. In this profile, levels are defined that allow all the way up to full resolution size pictures. Therefore, this profile is equally applicable to both the PC and TV environments.

The FGS profile supports scalability for applications where bandwidth and processing power vary, such as is typical for Internet distribution to PCs. This profile allows the use of ASP as a base layer, with scalable FGS extension layers on top that allow improved quality for those receivers that can receive and/or process those extension layers. [7]

To promote the Interactive / Web enabled class of applications, the *Core* and *Main* Profiles are defined. Core Profile is a superset of SP and Main Profile is a superset of Core. The Core Profile adds B-VOP and Binary Shape coding tools on top of SP. The Main Profile adds Gray Scale shape coding, Interlaced Video Coding and Sprite Coding tools. It should be noted that Core and Main Profile do not have Quarter Pixel and GMC coding tools, as they were not developed at that time.

MPEG-4 also defines many other profiles. Reference [3] provides a full description.

Advanced Video Coding (AVC)

During 2001, the MPEG committee conducted some informal tests to see whether video coding technology had progressed since the development of the MPEG-4 Visual Part 2 standard. It was concluded that, although there were not fundamentally new coding tools discovered since the completion of Part 2, many coding tools that were rejected during the development of MPEG-4, due to their high implementation complexity, should be reconsidered for inclusion to provide higher coding gain. Basically, advancements in implementation technology, such as the availability of low cost high-performance microprocessors and high-density high speed integrated circuits and memory, justified inclusion of more complex tools.

In addition, the ITU-T had already started developing a next generation (H.26L) for low bit rate coding. Tests also showed that the ITU-T's H.26L effort had done a good initial job of coherently assembling a basic structure of coding algorithms with new coding tools. Thus, a joint video team (JVT) was formed with ITU-T, starting from the defined H.26L structure, to develop jointly a new video coding standard to provide significantly higher coding gain, a standard that was not necessarily backward compatible with MPEG-4 Part 2. The first phase of this standard will be completed by December 2002. Its targeted goal is to provide at least 2 times the coding gain over MPEG-2 for hard-to-compress video sequences. Once the JVT completes its task, ISO will adopt it as MPEG-4 Part 10 and will call it Advanced Video Coding (AVC), and ITU-T will adopt it most probably as H.264.

At a high level, AVC is also motion compensated block transform based coding. Some of the video coding tools that are likely to form a part of AVC are: I, P and B frames, *variable block size* motion compensation, a *4x4 integer* (DCT like) transform, *multi-frame* motion compensation, interlaced video coding tools, quarter-pixel interpolation (eighth-pixel in discussion), a *de-blocking (in-loop) filter*, an *adaptive block size transform* (also under discussion), global motion compensation and *global motion vector* coding (also under discussion), *switch-P frames* (SP) to allow switching from one bit stream to another at specific locations, *universal* variable length coding and *context based adaptive binary arithmetic coding*. AVC has not yet defined any specific profiles and levels, but these are expected in May 2002.

It is expected that this new standard will be capable of sending high-quality movies in the 500 kbps to 750 kbps range, acceptable quality movies in the 250 kbps to 500 kbps range, and high-quality video below 2 Mbps. Examples of the coding efficiency of the MPEG-4 AVC (JVT) standard are provided in Figures 5 and 6. As in Figure 4, x-axes in these figures represent bit rates and y-axes represent PSNRs. Further examples can be found in [8]

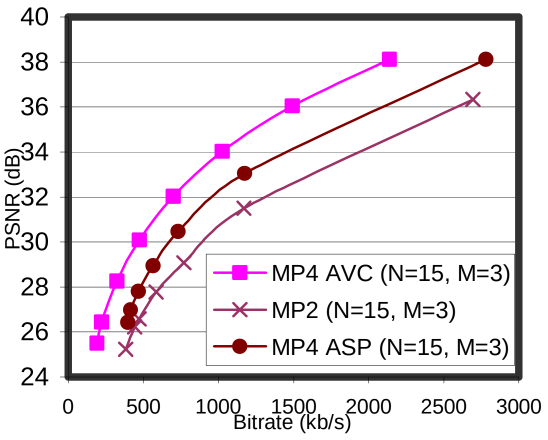

Figure 5: MPEG-2, MPEG-4 ASP, and MPEG-4 AVC, Bus sequence (CIF)

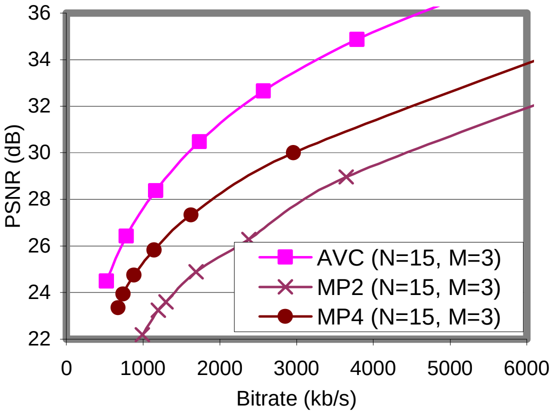

Figure 6: MPEG-2, MPEG-4 ASP, and AVC, Mobile & Calendar sequence (HHR)

The impact of a two-to-one efficiency improvement on cable and other forms of broadcast distribution is clearly one of capacity. More choices, more opportunities to narrowcast content, and better HDTV carriage will all be beneficial to the consumer. For IP carriage, where last mile bandwidth and quality of service guarantees are still major concerns, the efficiency gain may spell the difference between carriage of acceptable quality content at acceptable cost, and the more typical postage stamp sized Internet images. Furthermore, less video bandwidth consumption allows a set of new interactive services; such services can now supply additional streams of text, scene composition, graphics, etc., all under control of the overall application.

Other MPEG-4 Media Types

MPEG-4 extends AAC audio with several new tools, broadening its range of application. Speech coding is supported through a *twin vector quantization* algorithm, and a *CELP* (Code Excited Linear Prediction) algorithm, and text-to-speech conversion is defined. Synthetic audio can be supported through the *FM wavetable* or *model-based* synthesis algorithms, also including MIDI formats, essentially providing an enhanced “score.”

MPEG-4 supports face and body animation through a set of parameters that can be updated over time. As an example, avatars can be represented with very low bandwidth streams describing the animation changes and the accompanying coded speech.

Still images can be coded in MPEG-4 with the *zerotree wavelet* algorithm, and text can be coded in Unicode, with a font and color, for example. Rolling movie credits would be best handled in this fashion. MPEG-4 graphics objects are compressed (*geometry compression*) as a two-dimensional or three-dimensional mesh, with texture mapping. Figure 7 provides an example 3D mesh. Vertex positions, attributes, and connectivity (static or dynamic) would be coded as the geometry representation.

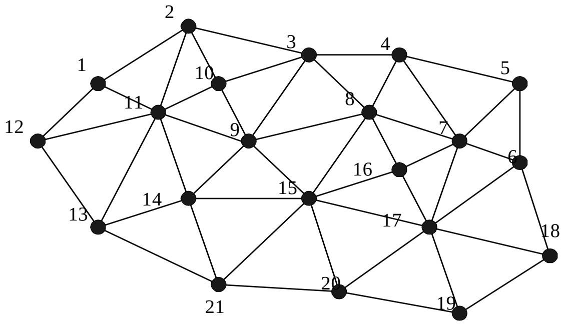

Figure 7: MPEG-4 Mesh

A full description of MPEG-4 media types can be found in [3].

MPEG-4 Systems

MPEG-4 systems provides various options for combining objects of different media types, or multiple objects of the same media type, into higher level structures, as well as techniques for combining and carrying multiple services. Figure 8 shows an example scene and Figure 9 its decomposition, which can be represented with BIFS. Timestamps can be supported through a *sync layer*, and services can be premultiplexed using a *flexmux*. MPEG-4 structures such as these can be carried in MPEG-2 transport, or directly in IP protocols [9]. See reference [3] for a complete description of MPEG-4 systems.

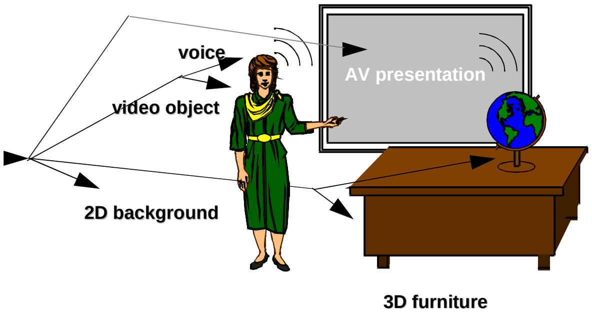

Figure 8: Multimedia Scene Example

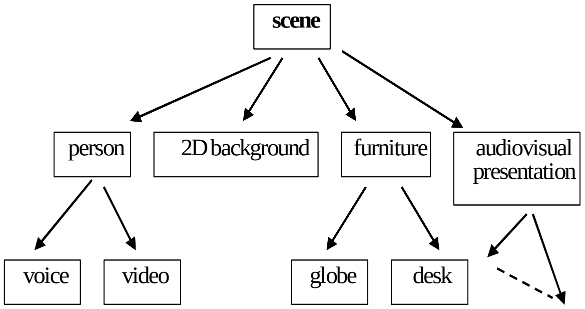

Figure 9: Scene Decomposition of Figure 8

MPEG-7 AND MPEG-21

MPEG-7 and MPEG-21 were not begun as efforts to find yet another new generation of coding techniques. Rather, when one surveys the broad landscape of multimedia applications, there are areas not covered by the earlier MPEG standards.

Search and retrieval of media objects is an important area for future growth of services. Images and sounds are proliferating every day on the Internet, and in the home, and the text-oriented search engines now available can only locate such content if an author has provided the right textual key words or titles. Longer term, users need to be able to find objects through their intrinsic multimedia attributes.

MPEG-7 standardizes such attributes to enable a new class of search engines, similar to the way earlier MPEG standards describe syntax decoding. Thus MPEG-7 does not define how one extracts such features from an object, nor does it define search engine technologies such as pattern recognition. This is the province of the creative process, such as in MPEG video *encoding*.

Examples of the over 100 features defined in MPEG-7 include dominant color or texture, color histograms, thumbnails, shape or edge description (well suited to trademark searches), and various motion oriented attributes (well suited to video surveillance). For audio, examples include waveform and spectral envelope, which can assist in finding similar voices; spoken word features, such as speech recognition employs; timbre description; and other features that might help locate content based upon a whistled or hummed portion. MPEG-7 also includes more basic aspects of content, such as content author, or owner, access rules, storage format, and so forth.

Structurally, features are defined by *descriptors* with defined syntax and semantics. *Description schemes*, expressed in XML schema, also provide relationships between these descriptors.

MPEG-21 began in 2000, and is still in its earlier stages. This work addresses the overall multimedia framework, filling in any other missing elements for providing a complete systems approach to creating, storing, sending, and using multimedia content. One key area being addressed by MPEG-21 currently is the definition of a digital rights description to support standardized access to content. This work will produce a *Rights Expression Language*, and a *Rights Data Dictionary*.

SUMMARY

The nearly 14-year MPEG process has addressed a wide range of multimedia signal processing and systems issues. From its inception as an audio/video compression standard for storage of content, through its evolution to its current overarching support of multimedia content generation, transmission, storage, and consumption, MPEG has succeeded in producing useful, successful, standards. As the earlier MPEG-2 standard has been widely deployed in cable systems and other broadcast distribution networks, the newer MPEG standards can be expected to support new classes of (revenue generating) services and applications for those industries.

REFERENCES

[1] MPEG-1, ISO/IEC 11172: Information Technology - Coding of Moving Pictures and Associated Audio for Digital Storage Media at up to about 1,5 Mbit/s.

[2] MPEG-2, ISO/IEC 13818: Information Technology - Generic Coding of Moving Pictures and Associated Audio Information.

[3] MPEG-4, ISO/IEC 14496: Information Technology - Coding of Audio-visual Objects.

[4] MPEG-7, ISO/IEC 15938: Information Technology - Multimedia Content Description Interface.

[5] MPEG-21, ISO/IEC TR 21000: Information Technology - Multimedia Framework.

[6] A. Luthra, “Need for Simple Streaming Video Profile,” ISO/IEC JTC1/SC29/WG11 M5800 Noordwijkerhout, March 2000.

[7] ISO/IEC 14496-2:1999 / Final Draft Amendment (FDAM4), WG-11 N3904, January 2001.

[8] A. Luthra, “MPEG-4 Video Coding Standard – An Overview,” SCTE 2002 Conference on Emerging Technologies, January 2002.

[9] "RTP Payload Formats for Carriage of MPEG-4 Content over IP Networks," ISO/IEC JTC 1/SC 29/WG 11 N4428, December 2001.

ACKNOWLEDGEMENTS

This paper is in part based on the work done by the video coding team in the Advanced Technology department of Motorola Broadband Communications, and of course on the work done by the entire MPEG committee. Our thanks go out to all these engineers, without whose commitment and expertise the MPEG process would never have produced such impressive successes.

CONTACT INFORMATION

Dr. Paul Moroney, Dr. Ajay Luthra
Motorola Broadband
6450 Sequence Dr
San Diego, CA 92121

MPEG-4 VIDEO-ON-DEMAND FOR CABLE SYSTEMS: AN OVERVIEW

By R. Jordan Greenhall
DivXNetworks, Inc.

Abstract

It is the purpose of this document to discuss the features and advantages of MPEG-4 video-on-demand solutions for cable systems. In so doing, the document will provide an introduction to the MPEG-4 ISO standard, review different video-on-demand products and different MPEG-4 video-on-demand solutions for those products.

INTRODUCTION TO MPEG-4

MPEG-4 is a new international ISO multimedia standard designed to be a complete and comprehensive standard for all multimedia. Unlike the MPEG-1 and MPEG-2 standards that were targeted at relatively narrow applications, MPEG-4 has been specifically designed to support a very broad array of applications across a large number of media and multimedia requirements. As a consequence, the standard is large and has a wide array of applicable tools that distinguish it from MPEG-2. However, the principal foci of MPEG-4 are:

- Superior coding efficiency. MPEG-4 is designed to provide video and audio quality indistinguishable from MPEG-2 at 1/3rd to 1/8th the bitrate.
- Interactivity. MPEG-4 natively supports object-based video (e.g., video "hotspots") and data back-channels, allowing it to provide the foundation for comprehensive interactivity.
- Multiple-platform. MPEG-4 is designed to deliver multimedia content across virtually all delivery media and to virtually any form of device. MPEG-4 can be delivered over wireless or wireline, or on physical media; it can ride on top of MPEG-2 Transport Streams, ATM, IP, and other transport protocols. Moreover, MPEG-4 is designed for playback on devices ranging from wireless handsets to digital cinema projectors.
- Massively scalable bitrate. MPEG-4 is designed to provide high levels of audiovisual quality at bitrates scaling from the micro-scale (sub 20Kbps) all the way to lossless Digital Cinema. This scalability is reflected both in the video and audio codecs, as well as in the file format and systems layer that has specific tools to allow for efficient dynamic scaling across bitrates.

VIDEO-ON-DEMAND PRODUCTS

To understand the impact of MPEG-4 on Video On Demand economics, it is important recognize that the current business environment is deploying three different kinds of "Video On Demand" – each of which has different properties and requirements.

Pay-per-view (Near VOD)

Pay-per-view, or Near Video On Demand is not, properly speaking, a Video On Demand product at all. Rather, pay-per-view combines the traditional broadcast video model where a single channel is addressed universally to all homes on a network with selective access technologies that limit the viewability of content from the client end. As a consequence, NVOD is able to restrict viewing to the granular "pay-per-movie" level and, by pushing out content on a staggered basis over multiple channels, generate an experience that is similar to "true" Video On Demand. However, while NVOD has proven attractive to subscribers, its relatively narrow content profile and, more importantly, its restrictions on subscriber control over the viewing experience (starting and stopping a movie, for example), have limited its subscriber appeal.

Pay Video-on-Demand (PVOD)

Thanks in large-part to MPEG-2 digital video technologies, service providers have recently been able to replace Pay-per-view systems with true Video-On-Demand solutions. With Pay Video-On-Demand, a subscriber is able to access a content library, select which content he wants to watch, pay for that content on a "per-view" basis and then watch it entirely under his own terms (i.e., begin watching when he wants, stop, rewind and fast forward, etc.). Because of these advantages, PVOD has proven substantially more popular than pay-per-view. Early trials of Pay VOD have shown buy-rates of 3 – 4X that of PPV, driving $16 - $22 monthly revenue per subscriber.

The principal weakness of Pay VOD is cost. Current PVOD video servers and storage facilities are expensive to deploy. Moreover, because each discrete PVOD session requires an entire dedicated MPEG-2 digital channel, the opportunity costs of PVOD are relatively high. (This is particularly exacerbated by the fact that most PVOD demand is currently focused in a relatively narrow time window during the week – requiring a large allocation of system bandwidth to meet peak demand which ends-up underutilized during the rest of the week.)

As a consequence, the utility of PVOD has been restricted to content that can command a relatively high price per view – premium events and blockbuster movies.

Subscription Video-on-Demand (SVOD)

Subscription Video-on-Demand (SVOD) is a relatively new VOD model that has arisen in response to the weaknesses of PVOD. With SVOD, a subscriber pays a fixed monthly price for access to a certain content library and is then able to view content from that library as much (or as little) as he wants during the month. The predicted utility of SVOD is fourfold. First, that bundling content libraries under a single monthly fee will provide the kind of value that will attract subscribers to different kinds of VOD content – archival movies, television content, educational content, etc. Second, that SVOD subscription fees will enable content providers to more effectively rationalize their production and revenue risks, thereby increasing ROI; Third, that SVOD usage will be less focused on nights and weekends than PVOD, thereby utilizing allocated VOD infrastructure more efficiently. Finally, that content usage and subscription fees can be appropriately calibrated to earn the system Operator a positive ROI on this alternative.

The principal weaknesses of SVOD are threefold. First, that SVOD content libraries can be very large – orders of magnitude larger than PVOD libraries, thus presenting substantial storage challenges. Second, that each SVOD session requires a dedicated VOD server channel. Thus while SVOD will likely fill VOD capacity troughs during non-peak hours, SVOD will also add to peak VOD consumption, thereby requiring expensive additional capacity. Similarly, the third weakness of SVOD is that under current MPEG-2 technology, each SVOD session requires an entire dedicated digital channel. As a consequence, a single subscriber watching an SVOD episode of Seinfeld would require as much bandwidth as 100,000 subscribers watching NBC.

DIFFERENT METHODS OF DELIVERING VIDEO-ON-DEMAND

MPEG-4 technologies can be used to deliver Video-on-Demand more efficiently and flexibly than current MPEG-2 based VOD systems. Cost savings using MPEG-4 can be found in storage, server infrastructure, bandwidth utilization and even in customer-premise equipment. Because of its flexibility, MPEG-4 can be delivered both via a cable system's digital video infrastructure as well its IP data infrastructure. In both of these modes, MPEG-4 is significantly more efficient than MPEG-2 for VOD. However, as discussed below, delivery of MPEG-4 VOD over cable IP data infrastructure is particularly exciting and presents the most compelling argument for MPEG-4.

The principal weakness of MPEG-4 is simply that it is a new technology and must compete with legacy systems and legacy investments in the older MPEG-2. Although some of this legacy investment is found in the headend (storage and servers), these technologies can be largely repurposed to MPEG-4 use and in any event represent the smallest portion of legacy investment. Rather, the largest hurdle faced by MPEG-4 is in the customer premise: digital set-top boxes. None of the currently deployed digital set-top boxes support MPEG-4. As a consequence, any decision to deploy MPEG-4 within a cable infrastructure must present a compelling argument to replace digital set-tops, or to supplement or upgrade them with collateral equipment (e.g., more functional cable modems, gateway devices, set-top upgrades, etc.).

MPEG-4 Digital Video

As discussed above, MPEG-4 is designed for delivery over many different protocols – including the MPEG-2 Transport stream. As a consequence, MPEG-4 technologies can be used within the MPEG-2-based digital video infrastructure to leverage as much of the current infrastructure as possible while taking advantage of some of the benefits of MPEG-4. The key advantages of such an approach are found in reduced storage requirements and reduced network bandwidth burdens allowed by MPEG-4's superior compression.

By compressing VOD content by a factor of 3 - 5X without any loss of quality, MPEG-4 gives cable operators significant flexibility in their VOD strategy. A cable operator can choose to simply use less storage for their current VOD content; provide a larger library of content over VOD; distribute their content more widely over their network, thereby reducing overhead on their costly ATM backbone; etc. Similarly, because MPEG-4's superior compression means 3 – 5X as many simultaneous streams on a given portion of allocated VOD bandwidth, a cable operators opportunity cost to scale VOD services is considerably reduced.

However, it is very important to note that in order to take advantage of these efficiencies, a cable operator must have in-place customer premise equipment that is capable of decoding MPEG-4 video. Although current MPEG-2 set-tops will be able to receive and decode the MPEG-2 Transport Stream carrier, the "supercompressed" MPEG-4 video signals within that Transport Stream will be inaccessible to the set-top without a dedicated MPEG-4 decoder.

The advantages of MPEG-4, including its native support for interactivity and ability to carry compression efficiencies to high definition and beyond, are compelling for any MSO that is serious about VOD (and interactivity) in its value-added strategy. However, in the event of such an upgrade to CPE, a service provider would be well-disposed to pursue more multi-functional "gateway" devices which, in addition to providing a platform for multiple bundled services in addition to cable and VOD, provide a pathway for the delivery of MPEG-4 video over an a cable system's IP data path.

MPEG-4 Video Over Data Networks

MPEG-4's ability to supercompress video, combined with its ability to delivery video over multiple transport protocols (specifically IP), allows an MPEG-4 enabled MSO to exploit their IP data infrastructure as truly effective video delivery channel.

The advantages of doing so are multiple:

1. Commoditized Backend Infrastructure – delivering VOD over IP Data means that the cable MSO can utilize its backend data infrastructure to distribute content to headends. This backend infrastructure can be built on-top of commoditized standard IT hardware such as Gigabit Ethernet, rather than dedicated and much more expensive ATM.
2. Commoditized Server and Storage – similarly, using MPEG-4 over IP enables VOD systems that utilize standard commodity IT hardware for both storage and video servers. This means storage and per stream costs that are a fraction of current VOD systems (over and above the savings driven by smaller files and reduced bandwidth).
3. Repurposed Hardware – because MPEG-4 over IP is treated by an MSO's infrastructure as "just more packets" a great deal of the infrastructure required for distributing and delivering MPEG-4 video content is the same as is used for all other IP data: email, files, web browsing, etc. As a consequence, capital investment in that infrastructure is paid for by multiple services, not just VOD.
4. Repurposed Bandwidth – similarly, and perhaps more importantly, MPEG-4 over IP uses the same bandwidth within the cable plant as all other IP data. Thus, the cable operator using MPEG-4 over IP as their VOD delivery medium is able to **better utilize** their existing IP data bandwidth, rather than attempt to figure out how to deal with underutilized dedicated MPEG-2 VOD bandwidth.
5. Flexible Delivery Approaches – as discussed above, there are many different kinds of VOD products, each of which presents its own tribulations for efficiency and ROI-conscious network operators. Delivery of MPEG-4 over IP can be done in three different ways, each of which has strengths that give network operators tremendous flexibility in efficiently and effectively delivering content to their subscribers.

MPEG-4 IP Data Streaming

Streaming is a technique used in data networks to provide an end-user experience that is identical to MPEG-2 based VOD. Using streaming, a video is delivered on-demand, in real-time and has complete "VCR"-style control (stop, rewind, fast-forward, etc.). Within a private network with adequate bandwidth, IP-streaming can provide quality of service indistinguishable from MPEG-2 VOD.

The principal drawbacks of streaming are that it is more server-intensive than other techniques (described below), thereby requiring more server infrastructure per simultaneous stream, and that streaming creates a virtual "channel" of dedicated bandwidth that lasts throughout the playing of the video content. As a consequence, streaming cannot take advantage of IP data's ability to "burst" delivery and provide more effective bandwidth shaping (described below).

The biggest advantage of streaming is found in multi-user events (particularly sporting events) where many of the features of VOD (such as when the content will be viewed, and fast-forward) are not applicable. In this case, network operators can take advantage of a variety of "IP Multicast" approaches to deliver an effective VOD experience to multiple simultaneous subscribers while dramatically reducing both server and network bandwidth overhead.

MPEG-4 IP Data Downloading

For many kinds of VOD content, the most efficient mode of delivery is "downloading" rather than "streaming." This is because downloading separates the act of viewing from the act of delivery and enables the network to deliver VOD content in the most efficient way for the network. Thus, where a piece of content is delivered at an average of 800 Kbps and the network is capable of delivering 3 Mbps to the subscriber, the network can "choose" to burst a 120 minute movie to the subscriber in just over 30 minutes. As a consequence, rather than tying down a server session and 800Kps of network bandwidth for the entire viewing, network resources are rapidly freed-up for other uses within the IP Data pipe (other VOD sessions, e-mail, web-browsing, etc.). Moreover, because the content resides at the client after the download is complete, reviews of the content take place entirely on the client and impose no additional burden on the network infrastructure.

The principal drawbacks of Data Downloading are:

- Storage Requirement – for content to be pushed to the client, the client must have some significant storage capacity (in the range of 350 MB per hour of content). In a gateway device equipped for PVR (and datacasting as described below), this is a non-issue, but this can be a serious hurdle for very inexpensive set-top devices.
- Limitations to Trick-Play – data downloading is fully capable of all of the features of trick-play with the single exception that fast-forward cannot go beyond where the content has been downloaded. Thus, if a subscriber begins viewing a movie and wants to fast-forward to the end, he will have to wait some amount of time (30 minutes in the above example) before he can do this. Because of the way that consumers typically use trick-play features, this limitation is usually unimportant, but should be considered when choosing which delivery approach to take.

MPEG-4 Datacasting

One of the more unique and compelling applications enabled by MPEG-4 over IP is "datacasting". With datacasting, the network operator "pre-loads" certain content on the subscriber's CPE storage by downloading that content during network down times before the content is available for viewing. A typical application, for example, would be to datacast a VOD version of a very popular blockbuster movie before the first day of the VOD window. When the content is available for viewing, all requests actually come from the pre-loaded content on the subscriber's client – meaning that the operator's VOD infrastructure takes no load whatsoever from the request.

This technique can be very efficient when used for content that is either likely to be extremely popular and the operator wants to avoid peak strain on his system, or for content that is uniquely targeted to the specific subscriber and the operator wants to deliver the content during network downtime (and avoid competing for resources during peak times).

The principle weaknesses of datacasting are lack of storage on the subscriber's CPE. With the CPE almost certainly serving double or triple duty as PVR, ad server and datacasting server, there is only so-much content that can be pre-loaded by datacasting.

However, this drawback points to the larger issue that each of these three methods of delivery over the IP Data infrastructure are complementary, not mutually exclusive. The best approach will require a mixture of methods appropriate to the content offering, the network and the subscriber. The key advantage of delivering MPEG-4 video over IP is that it has the flexibility to enable the operator to select the mixture of methods that is most appropriate and efficient for its profile.

(Note, it should be mentioned that the advantages of MPEG-4 compression are doubly applicable where the network operator will be providing a gateway device with PVR functionality – an MPEG-4 enabled PVR can store 3-5 times as much content as a typical MPEG-2 PVR. This means that a much smaller hard-drive can be used for the same amount of hours stored, at great savings to the MSO.)

ECONOMICS

Clearly, the actual economic footprint of a Video-on-Demand system is highly dependent upon specific network topologies and a mix of technologies used to deliver VOD. The below cost comparison between an MPEG-2 based VOD system and an MPEG-4 based VOD system provides a baseline that can be used to derive more detailed estimates based on real conditions in the cable network.

This model assumes a comparison between two cable networks with 100,000 subscribers (10,000 simultaneous streams at peak capacity). Both networks present a simplified topology of 100 nodes, each with a headend serving 1000 subscribers and a central supernode that serves the entire subscriber base through those nodes. One network is assumed to be using an MPEG-2 VOD system networked with an ATM backbone between nodes and the supernode, while the other network assumes an MPEG-4 based VOD system using Gigabit Ethernet to connect nodes with the supernode.

This basic topology assumes that each headend node stores roughly 20% of the total content library and absorbs 80% of the total VOD requests directly, while the supernode stores 100% of the content library, but handles only 20% of the total VOD requests (i.e., 2 out of 10 VOD requests aren't satisfied by the local headend and have to be delivered from the larger archive at the supernode). In the MPEG-2 system, this delivery from the supernode is via an ATM backbone, while in the MPEG-4 system, it is via a Gigabit Ethernet backbone.

MPEG-2 Based System

Total VOD Network Cost	$	5,917,139
Cost Per Simultaneous Stream	$	592
Cost Per Sub	$	59
Total Cost Less ATM Backend	$	4,717,139
Cost Per SS Less Network	$	472
Cost Per Sub	$	47

MPEG-4 Based System

Total VOD Network Cost	$	1,586,840
Cost Per Simultaneous Stream	$	159
Cost Per Sub	$	16
Total Cost Less GigE Backend	$	1,386,840
Cost Per SS Less Network	$	139
Cost Per Sub	$	14

Ratios

Total MPEG-4 vs. MPEG-2	27%
Total Less Networking Backend	29%

This model reflects cost elements associated with content storage at both the supernode and each headend node; VOD servers at both locations; and the ATM or Ethernet networking to support the VOD system. The model assumes a least-efficient MPEG-4 system that uses streaming technology and does not exploit download "bursting". The model also does not contemplate the cost savings and opportunity cost advantages associated with repurposing of IP data bandwidth contrasted with dedicated MPEG-2 VOD bandwidth. Finally, the model does not contemplate the tremendous efficiencies associated with datacasting to push the most popular or unique content to the end consumer device.

Thus, under the assumptions of this model, an MPEG-4 based IP video system is at least four times as cost-effective as an equivalent MEPG-2 based system – with tremendous flexibility to achieve even more substantial cost savings as more tools from the MPEG-4 toolbox are utilized.

Finally, it should be mentioned that this whitepaper addresses only the advantages of MPEG-4 for delivering VOD in a manner that is roughly identical to current MPEG-2 systems. This, of course, is only the tip of the iceberg for MPEG-4. The new MPEG-4 standard is a complete platform technology for a variety of next-generation applications including robust interactive and dynamic content and innovative distributed content delivery models. Consequently, upgrading to an MPEG-4 system makes sense both on a short-term ROI basis and on a longer-term strategic basis.

Contact Information:
Jordan Greenhall
DivXNetworks
10350 Science Center Drive.
Building 14, Suite 140
San Diego, CA 92121
Phone: 858.909.5300
Fax: 656.909.5301
Email: jgreenhall@divxnetworks.com
Website: http://www.divxnetworks.com

NETWORK PVR VIDEO SERVER ARCHITECTURE

Jay Schiller, Senior VP Broadband Strategy and Product Management
Michael Fallon, Senior Technical Writer
nCUBE Corporation

Abstract

Set-top Personal Video Recording (PVR) at home and network PVR services from cable operators allow consumers to record and watch television programming at their leisure. PVR also gives consumers full VCR-like control (pause, rewind, and fast forward) over live television.

This paper discusses the broadband video server architecture as a platform for network-based PVR services. Network PVR allows cable operators to offer PVR services to digital subscribers using existing digital set-top boxes without local hard disk drives. Instead, network PVR uses video on demand (VOD) systems with video servers and disk storage located in cable headends and hubs.

Network PVR Video Server Architecture describes the technical infrastructure - at the hardware and software level - involved with capturing cable networks into a video server while simultaneously streaming to subscribers and offering full VCR-like control. The paper looks at network PVR architecture in detail, describing the video server requirements and capabilities, as well as the additional equipment and software required for deploying network PVR services. The paper also looks at the Electronic Program Guide (EPG) application integration that is required for consumers to interact with the network PVR service.

WHAT IS NETWORK PVR?

PVR technology, also known as Digital Video Recorder (DVR) technology, is a relatively new product making rapid advances in the interactive television (ITV) marketplace. Consumers can purchase PVR products branded by Tivo, UltimateTV and ReplayTV today in consumer electronics stores and in conjunction with digital broadcast satellite (DBS) services. Today, PVR set tops are installed in the home just like – and alongside – VCRs and DVD players. A set-top PVR is yet another 'box.'

PVR significantly changes the television watching experience. Similar to a VCR, PVR gives consumers the ability to record television programs and watch them at their own convenience. PVRs can often be programmed to automatically record favorite programs. PVR also lets consumers have full VCR-like control (pause, rewind, and fast forward) over a live television broadcast. PVR lets you pause the Super Bowl, rewind and watch thrilling plays over and over again and then fast forward back to real time. You can also fast forward over commercials just like a VCR or jump past them using 'skip ahead' features that instantly advance the program by 30 seconds each time this feature is selected on the remote control.

PVR devices ingest and record programs, and store them on disk drives for playback with full VCR control. Playback can be within a second of the broadcast network feed or much later in time. In summary, PVR is essentially two applications:

- Record TV programs for viewing later with VCR-like controls - Users choose the programs and series they want to record. The PVR records the programs and displays a menu of recorded titles from which the consumer can select any recorded program, any time

- Watch currently broadcast networks with VCR-like controls - Users can pause and rewind live television programs

PVR dramatically changes the way people watch TV: People watch what they want when they want and they watch more TV. A NextResearch study revealed that the vast majority of respondents report changing their viewing habits because of their PVR. Of significant importance to cable operators is that 44 percent of PVR owners report having more premium channels than they had before PVR, and 43 percent have more total channels than before PVR. This is a clear indication that PVR owners seek more television choices than they did before owning a PVR set-top box.

Cable operators are currently exploring how to offer PVR functionality to subscribers, including whether to introduce subscribers to PVR using new, enhanced PVR set-top boxes with hard disk drives in the home or using headend-based video-on-demand systems enabled for network PVR. Everything that can be done with set-top based PVR can be done with a video server.

ADVANTAGES OF NETWORK PVR

Network PVR can be deployed on existing digital set-tops without a trip to digital subscribers homes. Network PVR offers economies of centralization that allow cable operators to extend the PVR offering to more subscribers with less capital and operational expense, in addition to a considerably shorter deployment phase. Because the entire digital subscriber base shares the storage footprint of network PVR, the aggregate storage requirement may be lower and the overall cost per subscriber reduced. Many subscribers can share one copy of each program or network broadcast. The centralization of hardware and operations minimizes staffing and technical support requirements while significantly reducing the investment in PVR-enabling equipment and its operational costs as compared to a PVR set-top box solution.

Consider some of the benefits for cable operators:

- Network PVR – more so than VOD – is the must-have ITV product offering that not only competes with DBS, but exceeds what DBS can offer
- Unlike set-top approaches offered by DBS, network PVR provides unlimited tuners and storage to consumers – all programs can be recorded all the time, so all programs can be available all the time even if subscribers did not remember to record it in advance
- Network PVR increases the number of digital subscribers and reduces churn
- Network PVR drives increased premium channel subscription rates and total channel subscription rates
- One server platform can be used for VOD, SVOD and PVR just by adding storage and streaming capacity
- Network PVR can be rolled out to existing digital subscribers without any truck rolls
- Network PVR operational costs are lower than set-top PVR costs – fewer total disk drives and no truck rolls or home repair calls
- High customer satisfaction – no lost content due to disk failures, no waiting for the cable guy, no recording conflicts

- ❑ Incremental revenue through PVR service fees, new digital subscribers and premium package subscriptions, and advanced addressable advertising sales targeted to subscribers

PVR SYSTEM OVERVIEW

With a PVR-enabled Electronic Program Guide (EPG) and a remote control, subscribers can navigate upcoming programs and past programs. Subscribers can then select programs or networks for on-demand viewing. Upcoming programs can be marked for recording, and once the program is recorded, it appears in the subscriber's personal library or "virtual locker,” which is accessible through the EPG. Personal libraries for each subscriber are maintained in a PVR preference server, which consists of a database and server-side business logic and is accessible through the programming guide. The preference server can also have applications that track a subscriber's favorite programs and alert them when programs matching their interests are available.

Subscribers can also select programs currently in progress and watch them on demand with VCR-like control. For example, a subscriber surfing through channels could begin watching a broadcast program already in progress, back up to the beginning of the program using the remote control, and then watch the program in its entirety. The subscriber could also channel off the initial PVR session, tune to another channel and enact PVR controls on another program without setting up a new session with the video server. Importantly, subscriber requests to time-shift the broadcast program result in a single video on demand session in which the streaming video content is simply switched on the video server. This single session per subscriber capability significantly decreases session set up and teardown traffic.

When subscribers select content from the server, whether it is for PVR, SVOD, or movies, the server streams the content to the QAM modulators and upconverters and into the hybrid fiber coaxial cable network cable plant. The content then streams into the subscriber homes where the signal is decoded by their set-tops and the programs displayed on their televisions. The QAM modulators and upconverters may be integrated into the video server, collocated with but external to the video server, or located remotely, many miles from the video server. If located remotely, the video streams are output from the server in either Gigabit Ethernet or DVB-ASI format, and then optically transported over fiber optic cable to remote QAM modulators and upconverters.

PVR SYSTEM ARCHITECTURE

There are a number of interdependent hardware and software systems necessary for network PVR, but the video server is the heart of the network PVR system. Understanding video server capabilities is important in deciding the appropriate server platform to choose. Given the massive amount of potential content storage possible with network PVR, storage scalability is a critical video server characteristic that must be fully understood to deploy a network PVR system cost-effectively.

In addition to the video server platform, there are architectural considerations for the deployment, which will affect the overall cost. Video servers can be deployed in either centralized or distributed architectures. In most cases, due to the massive amounts of content storage required for PVR, centralizing will be more economical.

The principal system components in an network PVR service include:

- Video encoders and demultiplexors to receive and format the television networks for ingest by the video server
- Video servers with capacity to ingest, store and stream content to subscribers
- Transport network from video servers to QAM modulators in headends and digital hubs
- Modulation and upconversion equipment to feed the combining networks
- Broadband two-way HFC network with bandwidth for network PVR streams
- Digital set-top boxes with PVR-enabled programming guides
- Subscriber preferences and recorded programming lists integrated with the video server and programming guides
- Billing system and conditional access systems
- Advanced addressable advertising applications

Figure 1 illustrates a typical network PVR implementation.

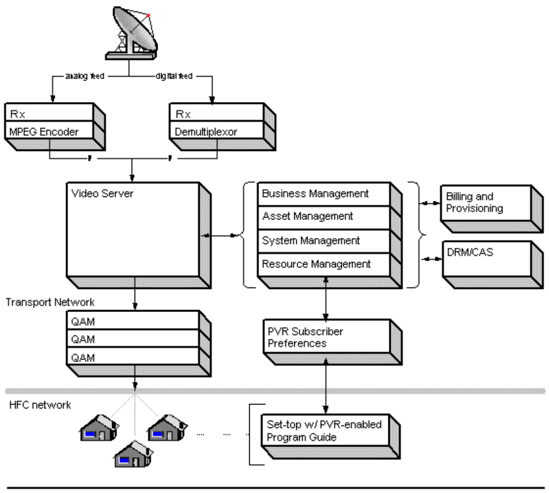

Figure 1: PVR System Architecture

As shown in Figure 1, analog and digital television network feeds are received at headends over satellite or over fiber at digital hubs. Analog signals are MPEG-2 encoded and ingested into the video server. Digital signals are demultiplexed and ingested into the video server. During ingest the fast forward and rewind view data is created, while the television network data is simultaneously written to disk storage. The television network data is then available - almost immediately - to be streamed out to subscribers just like video-on-demand (VOD) content.

Business management and other back-office applications manage the system bandwidth, coordinate video session setup and teardown, monitor system health, and report business and technical performance. During session management, coordination with conditional access and billing systems also occurs. The nCUBE back-office product is the nABLE Interactive Management Platform, and nABLE provides all the aforementioned functionality.

The PVR Video Server

The video server is the cornerstone of the network PVR system. Network PVR requires a video server to capture, store and stream video. The video server's ability to receive video in real time, prepare the visual fast-forward and rewind views, and write the video file to disk is essential. It also needs to identify start and stop record times within a network feed so that it can distinguish recorded programs for playback. In addition to recording and preparing network feeds, the video server will be simultaneously handling requests from the downstream network for streams, fast forward, rewind, pause, skip ahead and skip back on live broadcasts and recorded content.

PVR storage requirements are massive and a video server that cost-effectively scales to accommodate the video data is necessary to keep both capital and operating costs to a minimum. For example, to provide PVR storage for 200 networks for more than one month, the video server must be able to store more than 140,000 hours of unique content.

To scale storage-wise is not sufficient by itself, however. The video server must be able to stream out any of the captured networks and programs to any and all subscribers simultaneously. The goal is to minimize the number of copies of the same content needed to serve the maximum number of subscribers.

When considering PVR and the video server that will stand as the cornerstone of the PVR deployment, cable operators will want to look at several key characteristics of any one particular server platform:

- ❑ Storage Capacity
- ❑ Streaming Capacity
- ❑ Content Buffering and Recording

Capacity to Store Content

To store one cable network for one month requires 1.5 terabytes (TB) of MPEG-2 data ate 3.75 Mbps. A content library of 200 networks stored for one month each requires 300 TB of storage. Video servers that have streaming limitations must replicate content across video servers to satisfy stream requests. Video servers that have storage limitations must shuttle content from one server to the next to satisfy streaming demands. Shuttling content takes up enormous network bandwidth, and while content shuttling may work for movies, it is an impractical and very expensive solution for network PVR. The cost of duplicating content in multiple servers within a single headend or digital hub must be avoided or minimized.

Network PVR storage requirements also influence the deployment architecture in favor of a centralized approach, which avoids replicating content at multiple sites. Duplicating the same cable networks in each cable hub is more expensive than leveraging commodity fiber optic transport technologies and the fiber capacity commonly available in modern HFC cable plants. In particular, Gigabit Ethernet transport is rapidly becoming the most economical transport technology.

Capacity to Stream Content

More content equates to more on-demand use and more digital subscribers. In other words subscribers will watch more content from the video server when more choices are available. In addition, more content will attract more people to subscribe to digital services. PVR services also lower churn. The net effect is more digital subscribers using the system more often, resulting in higher stream concurrency than with VOD or SVOD.

Maximum simultaneous PVR stream capacity can easily vary from 30% to 90% of digital subscribers depending on how many cable networks are available on PVR and depending on whether they are available for record only or for full PVR even for the Super Bowl. For a region with 50,000 digital subscribers the video server may need to support 15,000 to 45,000 streams.

Content Buffering and Recording

The video server must be able to ingest and store a configurable time window for each television network. For instance, the most recent hour, week or months can be maintained. As new content is ingested, old content is deleted. This process is illustrated in Figure 2.

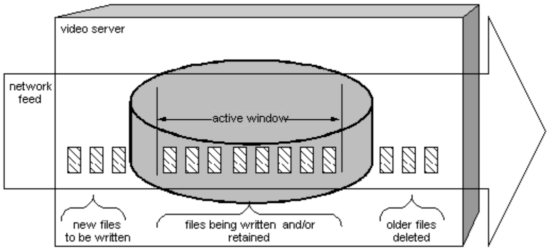

Figure 2: Content Buffering and Recording

Individual programs can be identified with the buffered MPEG-2 data in the window. Any portion of the buffer window that is identified as a recorded program can be stored on the server indefinitely. A program recorded for one subscriber can be made available to every subscriber a single copy of the recorded file.

The buffer window for each cable network and the record start and stop times can be accessed through the PVR system management application, as well as through Record and Play application programming interfaces (APIs).

The n4 Video Server

The current nCUBE video server is the n4. The n4 video server scales from small to large in five rack-unit (RU) building blocks called MediaHUBs. A single n4 video server can consist of anywhere from 1 to 256 MediaHUBs with stream capacity and storage scaling linearly. Table 1 shows the specifications for a fully configured n4 video server built with 256 MediaHUBs, assuming files encoded at 3.75 Mbps. Both the storage and streaming capacities are sufficient for a full PVR offering.

Table 1: n4 Video Server Specifications

Feature	**Specification**
Maximum Output Streams	> 33,000 streams
Maximum Storage in Hours	> 218,000
Maximum Storage in GB	> 368,000

Each MediaHUB can stream approximately 130 streams, assuming 3.75 Mbps. A single video server comprised of 256 MediaHUBs can serve 33,000 streams all from the same MPEG-2 file. The n4 video server provides 218,000 hours of unique storage, assuming 180 GB disk drives, which results in an aggregate of 368,000 GB or 368 TB.

All 218,000 hours of storage is available for unique content. There is no need to replicate content within an n4 server, whether it is configured for 100 streams or 33,000 streams. And 33,000 subscribers can access the same MPEG-2 file. This is very important for PVR because of the massive amount of content possible.

Deploying Network PVR

There are a number of unique characteristics to a PVR offering that operators will need to consider as they move into this advanced stage of interactive television.

Network Bandwidth Considerations

Operators must allocate enough bandwidth in the network to accommodate the PVR streams. Network PVR can easily require 30 to 90 MHz of bandwidth or more, depending on how many networks are available for PVR and whether the networks are available for record only or for record with VCR control.

Return path hardware is also required to accommodate upstream-request traffic; however, most of this hardware should already be in place for existing VOD and SVOD services.

Combining Set-top and Network PVR

If network PVR has an Achilles heel, it is that dramatic simultaneous spikes in bandwidth demand for certain live broadcasts could exceed the bandwidth allocated for on-demand services. This is the "Super Bowl" example where hundreds or even thousands of viewers of a live sporting event simultaneously stop the broadcast feed and rewind to watch a spectacular event. In this example, the system takes a substantial hit due to the large number of simultaneous stream requests.

One solution to the bandwidth spike issue is to combine network PVR with set-top PVR, offering subscribers a standard service of record-only network PVR for all networks and live PVR for a few networks. Disk-based set-tops could then be offered to premium customers paying a higher subscription rate for the full PVR experience for all events and all networks.

CONCLUSION

The subscriber experience defines PVR. Subscribers need to be able to select a program currently in progress and have full VCR-like control over it. They should be able to skip back to the beginning of a program already in progress. They need to be able to record programs in the future or choose to watch past programs. Current implementations and market research data show that PVR is a compelling technology for which consumers are willing to pay. Cable operators have the ability to deliver PVR functionality to their subscribers, driving incremental revenue from their subscriber base while they reduce churn and increase their digital subscriber base.

Because the two-way interactive network is typically already built for VOD and SVOD services, the incremental investment for PVR is primarily in the video server and video server storage. The video server storage and streaming requirements to deliver a compelling PVR offering are significant. Finding the most cost-effective video server solution that is capable of offering the acceptable level of services to subscribers is critical.

As PVR systems scale upward in storage and stream requirements, and as business models increasingly favor centralized architectures, the n4 video server emerges as an economically viable PVR platform. When video servers are integrated with content encoder and demultiplexors, nABLE management applications, PVR-enabled EPGs and subscriber preferences, cable operators can offer complete PVR services from their headends and digital hubs.

Contact: Jay Schiller, jschiller@ncube.com

ON PRESERVING THE QUALITY OF INTERNET STREAMING THROUGH A DOCSIS NETWORK

Matt Haines and Asha Vellaikal
Aerocast, Inc.

Abstract

Internet-based streaming media services can deliver high-quality audio and video to PC users at speeds from 500 kbps to over 1 Mbps. However, delivering a large number of these streams through a DOCSIS network can push the boundaries of the network capacity. The result is increased congestion over the DOCSIS channel that impacts all downstream users and reduces the quality of the real-time streams being delivered through the network.

In this paper we discuss solutions that can be used to preserve the quality of streaming media services through a DOCSIS network. These solutions can be divided into two categories depending on whether or not the underlying DOCSIS network supports Quality of Services (QoS) features. For DOCSIS 1.0, which does not support underlying QoS features, we discuss methods for adapting the streams in response to network conditions. For DOCSIS 1.1, which does support underlying QoS features, we discuss methods for streaming media applications to utilize the underlying QoS capabilities.

INTRODUCTION

Cable providers using the Data Over Cable Services Interface Specification (DOCSIS) now provide broadband Internet service to a growing number of homes and businesses. As of September 2001, 30 Multiple System Operators (MSOs) served 7.6 million cable modem subscribers [1].

In a typical configuration (Figure 1), a single Cable Modem Termination System (CMTS) provides a dedicated 27/38 Mbps downstream data channel that is shared by up to 1000 cable modem homes [2]. If 10% of the homes are equally sharing the bandwidth at any one moment, then each home would receive approximately 270/380 kbps of downstream bandwidth. However as emerging "high bandwidth" applications take hold, such as IP telephony and streaming media, this allocation falls below the threshold of adequate bandwidth.

The Internet Streaming Media Alliance (ISMA) [3] Profile 1 is targeted towards broadband users who want to view entertainment quality Internet media streams over personal computers or set-top boxes. Profile 1 is based on MPEG-4 [4] and supports video encoding rates from 500 kbps up to 1.5 Mbps. This means that a small number of users who are consuming high quality streaming media can dominate the total aggregate downstream bandwidth from the cable head end unless steps are taken to limit their bandwidth usage. However, aggressive bandwidth limiting techniques, such as throttling the cable modems to specific levels, often results in a poor quality streaming media experience. Unlike web traffic, the end-user experience for real-time streaming protocols degrades as bandwidth is limited below the encoding rate.

Thus the problem facing cable operators offering DOCSIS services is how to effectively control downstream bandwidth usage while preserving the quality of streaming media and other real-time, high-bandwidth services. In this paper, we address this problem by presenting a number of possible solutions for preserving the quality of streaming media. These solutions are grouped

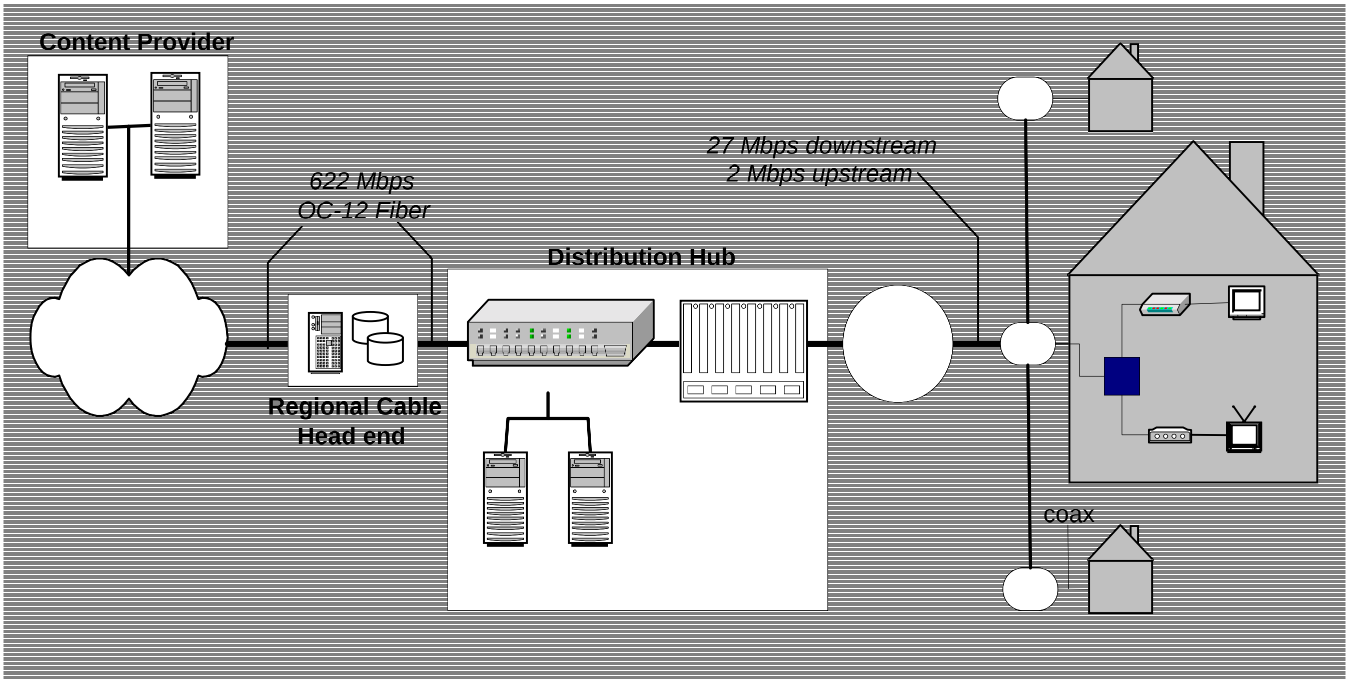

Figure 1: Sample DOCSIS Network Architecture

into two broad categories corresponding to DOCSIS 1.0 and DOCSIS 1.1, where the latter provides dedicated QoS features.

DOCSIS 1.0

In DOCSIS 1.0, there are no provisions for Quality of Service features such as bandwidth reservation. However, there are other means for preserving the quality of streaming media under demanding network conditions.

Co-Location with CMTS

Even with plenty of downstream bandwidth from the CMTS to the user, poor quality streaming will occur if the cable distribution hub (or regional head end) experiences congestion from the content provider origin servers through the Internet backbone. By co-locating streaming media servers within the distribution hub, the cable operator effectively shortens the end-to-end transmission and retains total control of the bandwidth needed to provide high quality streaming media services.

Co-location requires that all origin server content be replicated and sent to the distribution hubs, which may or may not be possible depending on the agreement between the MSO and the content provider. In the event that distributing origin content is not possible, an alternative solution is to co-locate streaming media caching servers [5]. The caching servers act to proxy the user's requests, serving the information from local cache if possible and otherwise requesting the information from the content provider origin server. While the first user to view a streaming media object may experience the Internet bottleneck, subsequent requests for the same content would be served from local cache, just as if the origin server were co-located. It is even possible to eliminate the bottleneck for the first user by pre-loading content that is expected to be popular into the caches.

The upside to co-location is total control over the end-to-end streaming bandwidth, eliminating any Internet backbone bottlenecks. The downside to co-location is the cost for additional servers (origin or proxy) in each distribution hub and the content coherence problem for distributed origin servers.

Adaptive Streaming

Another method to ensure reasonable quality under bandwidth variations is to use adaptive streaming, which modifies the encoding/streaming rate in response to network conditions. This approach requires a feedback mechanism between the client and the server to exchange up-to-date information regarding the bandwidth or network conditions experienced by the client. The server uses this feedback information to adapt the streaming rate appropriately by either reducing the rate when network congestion increases or by increasing the rate when congestion clears, thereby achieving reasonable quality under dynamically varying network conditions.

One simple method of adaptive streaming is to keep a single encoded file at the server and drop frames to reduce the overall rate. This frame dropping technique is often referred to as “stream thinning.” However, better quality can be achieved by keeping multiple bit rate encodings of the same stream with dynamic “up-shifting” or “down-shifting” between these encodings at key frames.

Examples of commercial products that support adaptive bit rate encodings are SureStream from Real Networks [6] and IntelliStream from Microsoft Windows Media [7]. Real Networks allows up to eight encoding rates in a single file with video window sizes and audio sample rates fixed for all bit rates.

The upside to adaptive streaming is that the streams can be adapted to changing bandwidth conditions to provide the optimal viewing experience. The downside of adaptive streaming is that the streams are not optimally encoded for a given bit rate since some of the encoding parameters are kept constant over all bit rates, and that encoded file sizes are increased for each encoding rate offered.

Access Limiting

The quality of streaming media applications suffers greatly when the available bandwidth for a session drops below the expected (encoding) rate. This would not be so bad if existing streaming media sessions were protected from new users consuming the last available bandwidth. Unfortunately, this is not the case. Rather, when a new user begins a session that consumes the last percentage of available downstream bandwidth, all existing sessions will suffer and begin “thrashing,” or missing packets and sending out more re-transmit requests that further exacerbate the problem.

The solution to thrashing is to not allow the available bandwidth to be completely consumed. This requires changing the access policy so that rather than always granting access to new sessions, it is possible to reject or limit access when the available bandwidth drops below a certain “safe” level.

Limiting access without the cooperation of all applications using the downstream pipe is only a partial solution. This is the principal behind formal QoS solutions that will be addressed in the next section. Still, limiting access for a single class of applications, such as streaming media, can still offer a significant benefit to preserving quality.

A prerequisite to implementing access limiting is to utilize a single funneling device through which all streaming media sessions flow. This provides a single point for gathering information on bandwidth conditions and deciding which requests are allowed or rejected. Devices which can play this role include a streaming server, proxy server, or intelligent router. In all cases, the device would be co-located close to the CMTS for gathering up-to-date information about the downstream bandwidth conditions through the CMTS's Network Management System (NMS) interface.

The upside to access limiting is the ability to limit the total number of streaming sessions or downstream bandwidth usage. This prevents thrashing and allows admitted users to preserve the quality of their sessions. The downside to access limiting without QoS support is the limitation to a single class of service, such as streaming, and the requirement of gathering real-time bandwidth information from the CMTS.

Excess Bandwidth Utilization

Many streaming applications are not constant with respect to their required bit rate. This provides an opportunity for creative ways for utilizing extra bandwidth that may be available at one moment to compensate for a reduction in available bandwidth at a later moment.

Skip Protection is a technology to improve the quality of playback at the client end. Skipping refers to pauses during playback caused by the need to re-buffer packets due to network congestion. Servers can utilize excess bandwidth available to buffer data faster than real-time (also referred to as "bursting") on the client machine. Thus if a larger buffer is available at the client, servers employing skip protection can utilize information about bandwidth conditions to fill up that buffer. Many servers including the QuickTime Streaming Server and the new Microsoft Windows Media Corona server now provide skip protection.

Forward Error Correction (FEC) techniques add redundancy to the original data stream so as to provide error resiliency to packet loss/corruption in the absence of a feedback channel between the client and the server. The DOCSIS physical layer uses FEC techniques to ensure reliable transmission over a noisy medium. However, it is also possible to use an application-level FEC mechanism to take advantage of excess bandwidth conditions. The Streaming Fountain product offered by Digital Fountain [8] utilizes an application-level FEC technique to encode excess information into the streaming packets. In the event that bandwidth is later constricted, the prior excess information can be used to re-construct lost packets without having to request retransmission.

The upside for excess bandwidth utilization is the ability to preserve the quality of streaming sessions in the presence of downstream bandwidth fluctuations. The downside of excess bandwidth utilization is the requirement of a client-side component to decode and utilize the extra bits being transmitted.

DOCSIS 1.1

Delivering applications with guaranteed quality of service (QoS) requires network-level components that provide end-to-end packet delivery with specified constraints, and QoS mapping to translate application-level quality of experience parameters to network-level QoS parameters.

QoS Network Components

In DOCSIS 1.0, all IP traffic from a single cable modem is grouped together under a single Service Identifier (SID). This means that all traffic types, including data, voice, and video, are treated equally by the Cable Modem (CM) and CMTS. DOCSIS 1.1 introduces the ability to separate different traffic types into different Service Flows and allow for different service parameters to be applied to each of the flows. In addition to adding service flows, DOCSIS 1.1 introduces new components for service flow management, downstream packet classification, and dynamic MAC messages, which together provide the basis for true QoS capabilities.

Differentiated Services (diffserv) is a QoS mechanism for supporting a limited number of QoS behaviors and aggregating all possible flows into this smaller set of behaviors. Diffserv defines a set of per-hop behaviors (PHB) that are applied to packets as they move through diffserv capable routers, such as a DOCSIS 1.1 CMTS. Though PHB only defines behavior for a single router, it is possible to combine multiple routers with the same PHBs and apply admission control to limit the number of PHB packets entering the system, thereby achieving end-to-end QoS.

An MSO can control the PHBs for all routers within its domain, but to provide true end-to-end QoS that spans multiple Internet domains, the MSO needs to negotiate bilateral agreements at domain boundaries called Service Level Agreements (SLAs). The SLA defines how a PHB from one domain will be carried through another domain.

Now that it is possible to give some traffic preferential treatment over other traffic, a policy system is needed to decide which packets receive the preferential treatment at the expense of other packets. The policy components include a policy database for keeping track of all relevant information; a set of policy decision points (PDPs) for inspecting resource requests and accepting or rejecting them; and a set of policy enforcement points (PEPs), which enforce the decisions made by the PDPs. For MSOs, the DOCSIS 1.1 CMTS will serve as the PEP since it has ultimate control over all packets into or out of the DOCSIS network.

The CMTS can receive policy information in two different ways. The first is "configured QoS," where policy is specified in the form of static classification information (packet IP/port) mapped to corresponding PHBs. In this case, the CMTS acts as both the decision point by performing the classification and the enforcement point by applying the correct PHB. The second method to receive policy information is "signaled QoS," where policy information arrives dynamically in the form of ReSerVation Protocol (RSVP) messages. In this case, the CMTS extracts the resource request from the RSVP message and presents it to the specified PDP for classification. Signaled QoS provides direct feedback to hosts by either rejecting the RSVP message or by accepting the RSVP message, in which case the CMTS is automatically configured to classify and handle the appropriate traffic.

QoS Mapping

End-to-end QoS using diffserv components (PHB, SLA, PDP, PEP, etc.) provides proper end-to-end packet delivery. However, true QoS requires support up and down the protocol stack on each side as well. That is, the streaming media clients and servers need to be able to communicate their quality needs to the underlying QoS network that will deliver the packets. Translating QoS specifications between different levels of the protocol stack is called *QoS mapping*.

Streaming audio/video application users express quality of experience in terms of parameters such as frames per second, resolution, and sampling rate. In addition, highly interactive applications also include delay as an important aspect of the user experience. These application parameters must then be mapped into network parameters such as bandwidth, packet loss, packet latency, and packet delay variation (jitter). For video applications, bandwidth and packet loss are typically more important than latency and jitter. While bandwidth demand is typically specified by the encoding rate, recent user studies have put absolute packet loss rates for VOD at 5% [9].

Once an application has determined how to map its quality of experience parameters onto network QoS parameters, it must employ a network API that allows these parameters to be specified. The latest version of the Microsoft Windows Socket (winsock2) API [10] allows for QoS parameters that utilize underlying QoS services from the operating system. This includes support for RSVP signaling, QoS policies, and invocation of traffic control over several protocol suites.

However, requiring an application to code directly to an underlying QoS network provides a dependency on the structure of that network. If the same application is to be executed atop various QoS networks, the dependencies become burdensome. To alleviate this mapping problem, the MPEG committee has defined an abstract QoS network within the Delivery Multimedia Integration Framework (DMIF) [11]. DMIF and its API (DAI) hide network-level details from the application programmer, including QoS signaling and transport mechanisms. DMIF-based QoS has already been studied for MPEG-4 streaming over IP networks with RSVP signaling and ATM networks with Q.2931 signaling [12].

Practical Improvements

Similar to the previous section that gave techniques to improve video quality in a non-QoS-enabled network, we now discuss how similar techniques can also be beneficial when DOCSIS 1.1 is available.

Co-locating video servers with the CMTS allows for end-to-end QoS during the transition period when Internet-wide QoS is not available but the user access network is QoS ready (DOCSIS 1.1 enabled). Even after Internet QoS is widely available, co-location allows for end-to-end QoS within the MSO domain, thereby removing the requirement of establishing and maintaining SLAs with Internet backbone providers.

Adaptive streaming techniques such as the availability of multiple bit rate encodings can be incorporated into QoS reservation decisions, thereby allowing for more choices. For example, if bandwidth reservation for the highest quality encoding fails, the server/client can re-negotiate with the QoS management for a lower bit rate version.

Access limiting is a fundamental aspect of any QoS enabled network that manages its resources for competing flows. However, unlike a DOCSIS 1.0 network where access limiting has to be explicitly introduced using specialized servers that regulate access, a QoS enabled network has native support for access limiting across all servers and application types. In addition, QoS-enabled networks allow advance reservations, which are not possible with simple access limiting.

Since QoS generally guarantees a certain constant bandwidth level for admitted applications, excess bandwidth utilization techniques would seem to offer little advantage. However, techniques such as skip protection with large client-side buffers can be used to reduce the dependence of proper QoS mapping for parameters like jitter. Excess bandwidth utilization techniques might also be used for variable bit rate encodings to reserve the “average” bandwidth requirement rather than the “peak” bandwidth requirement.

CONCLUSIONS

As DOCSIS networks continue to add subscribers and services, failure to preserve the quality of bandwidth critical applications will result in network congestion and poor user experience. However, there are several approaches for preserving the quality of streaming media in both a DOCSIS 1.0 and DOCSIS 1.1 network.

For DOCSIS 1.0, which does not support native QoS features, there are a number of techniques to preserve streaming media quality. Co-location places the entire end-to-end delivery route under the MSO domain, allowing for complete control over bandwidth policy decisions. Adaptive streaming and excess bandwidth utilization techniques try to preserve the optimal user experience under shifting bandwidth conditions. Access limiting provides a primitive level of QoS within a single application class, such as streaming media.

For DOCSIS 1.1, native QoS capabilities make it possible to offer differentiated service to streaming media applications, ensuring a certain level of bandwidth and latency tolerance. However, complete QoS requires support from all levels of the protocol stack as well as end-to-end network delivery. As these pieces start to unfold in a DOCSIS network and the wider Internet backbone, co-location, access limiting, adaptive streaming, and excess bandwidth utilization techniques can offer assistance in bridging the gaps and improving the overall user experience.

REFERENCES

1. *Cable Modem Market Stats and Projections*. Cable Datacom News. www.cabledatacomnews.com.
2. *Cable Data Network Architecture*. Cable Datacom News. www.cabledatacomnews.com.
3. *ISMA Specification*. Internet Streaming Media Alliance. www.isma.tv.
4. *Overview of the MPEG-4 Standard*. Moving Picture Experts Group. www.mpeg.telecomitalialab.com.
5. *Inktomi Traffic Server with Media IXT*. Inktomi Corporation. www.inktomi.com.
6. *SureStream*. Real Networks Corporation. www.real.com
7. *IntelliStream*. Microsoft Corporation. www.microsoft.com/windowsmedia.
8. *Streaming Fountain*. Digital Fountain Corporation. www.digitalfountain.com.
9. *QoS Requirements to Support Video and Audio Applications*. Dave Price. JANET QoS Workshop 2001.
10. *Microsoft Windows Quality of Service Platform Development*. Microsoft Corporation. www.microsoft.com/hwdev/tech/network/qos/default.asp
11. *An Overview of the Delivery Multimedia Integration Framework for Broadband Networks*. Jean-Francois Huard and George Tselikis. IEEE Communications Surveys 2(4), 1999.
12. *DMIF based QoS Management for MPEG-4 Multimedia Streaming: ATM and RSVP/IP Case Studies*. Victor Marques, Ricardo Cadime, Amaro de Sousa, and A. Oliveira Duarte. 3rd Conference on Telecommunications, April 2001.

OpenCable APPLICATION PLATFORM – STATUS AND ROADMAP

Allen R. Schmitt-Gordon, Ph.D.
Frank Sandoval
Cable Television Laboratories, Inc.

ABSTRACT

The OpenCable Application Platform (OCAP™) is a software middleware layer that resides functionally on top of the operating system of an OpenCable™ terminal. It provides an interface and enables application portability. A fundamental requirement is that applications written for OCAP be capable of running on any OpenCable hardware, without recompilation.

Two profiles of OCAP have been identified. OCAP 1.0 is a minimal platform that supports procedural applications with an Execution Engine (EE). OCAP 2.0 is a super set of OCAP 1.0, and includes support for declarative content with the inclusion of a Presentation Engine (PE), that supports HTML, XML, ECMAScript, and a bridge between the PE and the EE. The bridge enables PE applications to obtain privileges and directly perform EE operations.

OCAP 1.0 has been publically released, OCAP 2.0 is scheduled for around 1Q02. CableLabs plans to draft a family of OCAP specifications, each being backward compatible and defining different feature sets.

The OCAP specifications are based upon the DVB MHP specifications with modifications for the North American Cable environment that includes a full time return channel. OCAP 1.0 corresponds to MHP 1.0.2, and OCAP 2.0 will correspond to MHP 1.1.

INTRODUCTION

OCAP is part of a concerted effort, called OpenCable, by North American cable operators to provide the next generation digital device, encourage supplier competition, and create a retail hardware platform. A cable receiver provided at retail must provide portability of content and applications across networks and platforms, and be geared towards the full range of interactive services. Current devices are network specific and operate proprietary software that is not portable across platforms or networks. With OCAP, applications are written to a middleware API so that a single application can be deployed on any OpenCable host device. Such applications might include:

- Electronic Program Guide (EPG)
- Impulse Pay Per View (IPPV)
- Video On Demand (VOD)
- Interactive sports, game shows
- E-mail, Chat, Instant messaging
- Games
- Web Browser: Shopping, Home banking
- Personal Video Recorder (PVR)

OCAP provides applications an abstracted view of the receiver and network, hiding vendor and network specific characteristics that would tie an application to a given system. OCAP is operating system and hardware agnostic, so that applications can be run on a variety of CPUs and operating systems. The OCAP middleware also simplifies content development by encapsulating common operations within the API. Another essential requirement is that the middleware be secure and robust. Stability in the cable terminal or receiver is imperative as resets are not acceptible.

Background

The OCAP specifications include Application Programming Interfaces (API) as well as definitions regarding platform behaviors. The APIs define the syntax of the platform, while normative text define semantic guarantees.

The architecture of the OCAP 2.0 middleware comprises two parts: a Presentation Engine (PE) and an Execution Engine (EE). The PE is generally composed of an HTML engine and ECMAScript. The EE includes a Java virtual machine. This architecture is shown in Figure Figure 1. It shows that native applications are supported as well as applications written to the middleware via the OCAP interface.

In order to expedite development of the OCAP specifications, it was necessary to utilize existing standards and architectures as much as possible.

A key decision was to base the OCAP EE on standard Java APIs. In particular, the OCAP EE comprises pertinent portions of Sun's Java Virtual Machine (JVM) and JavaTV API specifications. The Sun technology is licensed by CableLabs and is available to all implementers of OCAP royalty-free. CableLabs will also incorporate the Sun Technology Compatibility Kit as part of the OCAP compliance test suite.

Another key decision was to adopt the DVB project's Multimedia Home Platform (MHP) specification. The DVB Project is a European consortium of manufacturers, content developers, broadcasters, governmental agencies, and system operators. MHP is designed to apply to a wide set of networks and devices, including cable networks. CableLab's recognized that DVB was trying to solve a similar set of problems. OCAP 1.0 is a delta from MHP 1.0.2, that is, the OCAP 1.0 document identifies conformance with MHP 1.0.2 section by section, and includes material to identify differences and extensions.

OCAP Roadmap

Rationale for Middleware

The current software model used by the cable industry is similar to the more general pattern in which applications are developed for a specific operating system (OS) and compilied for a specific device. This model does not lend itself to application portability where a variety of devices can be attached to several different network infrastructures. Currently, applications such as EPG, VOD, mail, etc. are compiled to the application programming interface determined by the operating system and associated hardware.

By providing a middleware layer that abstracts the functionality of the OS, hardware device and network interfaces, applications can be written that will run on any conformant platform. Java was chosen in order to avoid recompilation of source code as would be the case with languages such as C or C++.

OCAP 1.0

The primary architectural component of OCAP 1.0 is the Execution Engine (EE). It is composed of a Java Virtual Machine and a set of Java packages. Application portability is achieved through the "write once, run anywhere" nature of Java. Java source code is compiled into bytecode, and the JVM interprets the bytecode in real time on the target machine. Once the JVM is ported to a target machine, any EE application will run. The Java packages that comprise the EE include basic support functionality from Sun, such as net, io, util, and lang packages, and JavaTV and Java Media Framework. The EE also includes Java APIs identified in DAVIC, HAVi, and MHP 1.0.2. In addition, there are OCAP specific APIs.

This collection of APIs allows access to system resources, such as the cable tuner and input events, and access to network resources, such as System Information, in-band and out-of-band resources, and IP flows.

Because of the interpreted nature of Java, application security is easily maintained. Applications by default are given limited access to system resources. An application can gain extended access by requesting resources via a permissions file, and by being authenticated by the operator.

OCAP 1.0 specific APIs have been designed to address the business requirements of the cable industry. Unbound applications can exist outside of the context of a given service, or channel. Also, an optional monitor application can be created by the system operator. The monitor application allows to operator to control the execution lifecycle of other applications, and perform some simple resource management functions, in order to prevent harm to the cable network.

OCAP 2.0

OCAP 2.0 refers to OCAP 1.0 for core functionality and adds features found in MHP 1.1. A Presentation Engine and Bridge are included to support so-called declarative applications. The PE renders content such as graphics, text, animations and audio based on formatting rules in the PE itself and formatting instructions contained in markup language. The PE enables the use of tools that have been widely used for internet content.

OCAP 2.0 primary components consist of XHTML 4, XML 1.0, CSS 2, DOM 2, and ECMAScript.

In order to extend the functionality of the PE without adding the burden of requiring extensive plugins, a bridge is defined between the two environments. The bridge allows PE applications, through ECMAScript calls, to access functionality defined in the EE. Conversely, EE applications, through the DOM interface, can interact with concurrently running PE applications.

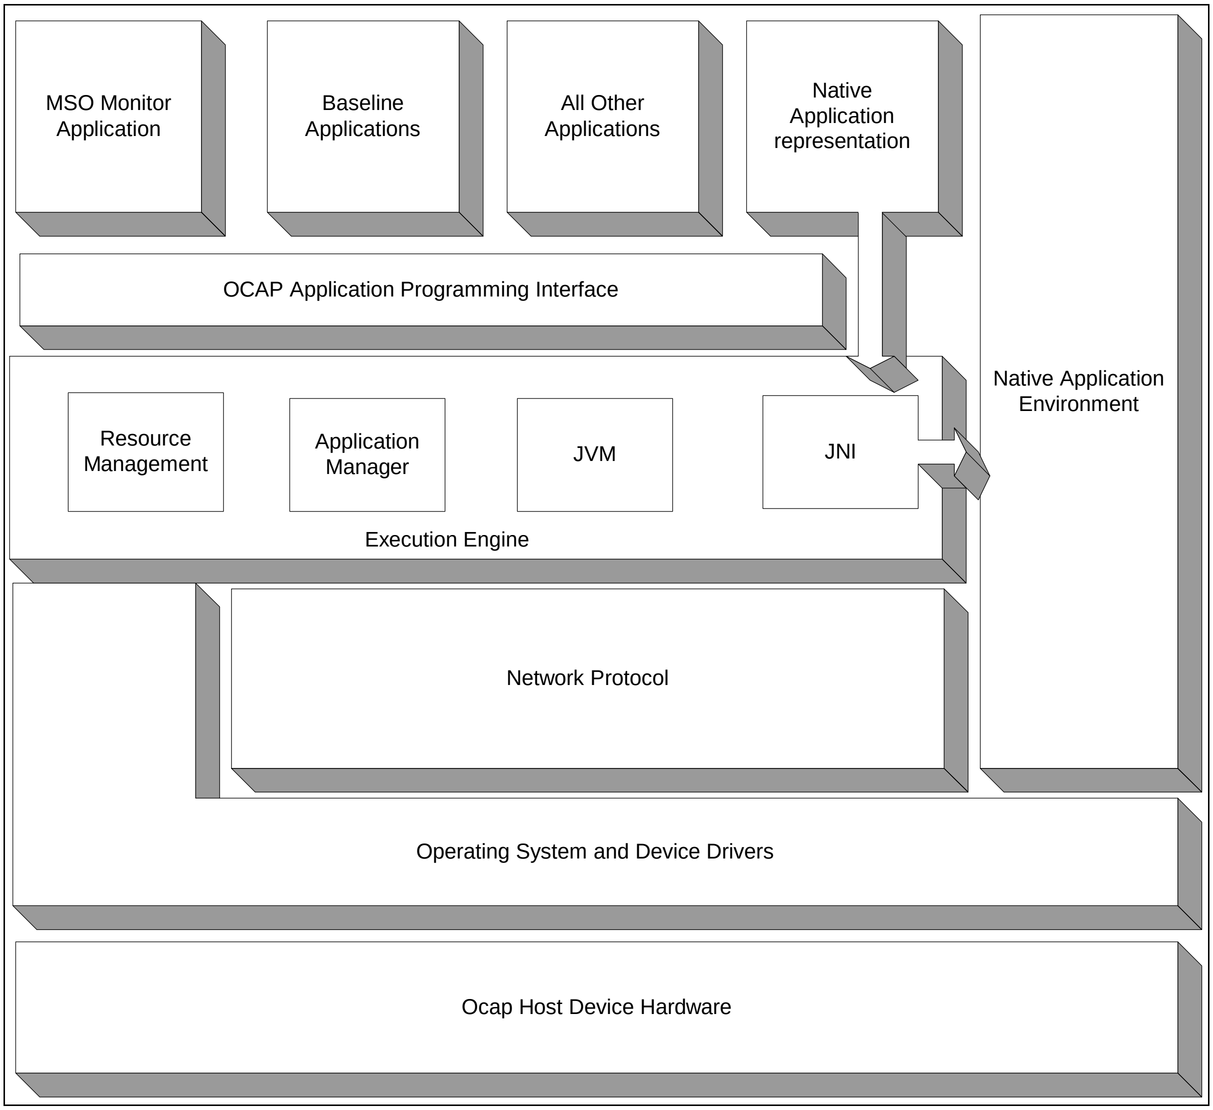
MSO Monitor Application
Baseline Applications
All Other Applications
Native Application representation
OCAP Application Programming Interface
Native Application Environment
Resource Management
Application Manager
JVM
JNI
Execution Engine
Network Protocol
Operating System and Device Drivers
Ocap Host Device Hardware

Figure 1 - OCAP 1.0 Software Architecture

For example, as part of the EE that is accessible via the bridge, JavaTV offers a common point of control and management of various system resources, includeing tuning. Thus, a PE application will access device resources through the bridge. This ensures that device resource contention is managed through a common control point for a PE and EE application that may be vying for the same resources.

Execution Engine

The EE provides a general application programming evironment for networking, file I/O, graphics, etc. The OCAP EE provides a full TV application environment. OCAP specific extensions have been added to MHP to cover elements such as a full time return channel, application management, resource contention mangement, and service information.

Major elements of the EE include control of application management through the pJava APIs, service information and selection through the JavaTV API's, media control through the JMF, and broadcast data through the MHP DSMCC APIs. Native applications are supported by creating an OCAP application that calls into native code via the Java Native Interface (JNI). In addition, the EE provides network management and IP data access and extensions from HAVI, DAVIC and DASE.

A fundamental feature of the EE utilizing Java is that security is built into the architecture from the ground up.

The components of the EE, some of which are shown in Figure 1, are described in detail in the OCAP 1.0 specification. In order to understand one of the key features of OCAP 1.0, which distinguishes itself from MHP 1.0.2, a discussion of the OCAP Application Model is warranted.

Application Model

OCAP 1.0 relies very heavily on the Application Listing and Launching APIs defined by DVB-MHP. This set of APIs enables lifecycle management of those applications that are bound to a service or program. OCAP 1.0 extends this model to include management of unbound applications or those applications that are not bound to a service or program. Signalling of such applications is done by a mechanism similar to the AIT specified by MHP 1.0.2, called the XAIT that is delivered via the traditional OOB or the DOCSIS channel (when available).

A special unbound application is called the Monitor Application The monitor application has a number of specific capabilities that can over-ride baseline functionality of the the OCAP 1.0 implementation (see Figure 1).

This functionality includes:

- Registration of unbound applications with the applications database.
- Validation of the starting of all applications through the setting of application filters.
- Registration of system errors that are propagated from the OS middleware, and/or OCAP 1.0 implementation, including OCAP applications.
- Request system reboot, and regitration of system reboot event.
- Control of copy protection bits and output resolution using the org.ocap.application.CopyControl interface.
- Filtering of User Input events and change their value before sending them to their final destination.
- Management of storage of any MSO unbound application, including itself using the persistent storage API, as defined by MHP 1.0.2 .
- Application lifecycle management, including that of the Monitor Application
- Resource Contention management of resource contention deadlocks.

If the Monitor Application implements any or all of the above functionality, it over-rides the implementation on a per MSO basis. Otherwise, the OCAP 1.0 implementation performs these functions in a generic manner.

OCAP-MHP Comparisons

The table in the appendix was generated by enumerating all of the java methods specified by the OCAP 1.0 and MHP 1.0.2 specifications. OCAP 1.0 shares a number of java packages with MHP 1.0.2. These packages are listed in section 12 of the OCAP specification. The packages come from SUN java JDK 1.1.8, PJAE 1.2, JavaTV, JMF, DAVIC, HAVi, and DVB-MHP. It is a useful but daunting task to enumerate the methods and interfaces that both OCAP 1.0 and MHP 1.0.2 share in common. Utilizing the javadoc tool and the javadoc APIs provided by SUN with the

JDK, the task is made somewhat less difficult. Doclets were written utilizing the javadoc APIs to produce custom javadoc output that list and count the methods in each package. These results are summarized in the table attached hereto as an appendix.

It should be noted that the number of MHP-specific methods is less than 10% of the total number of methods specified by MHP 1.0.2. Thus, even if OCAP 1.0 did not utilize any of the MHP-specific methods, it would still be highly compliant with the MHP 1.0.2 specification, with an 89% correspondence. The SUN java, javaTV and JMF components represent 71% of the total number of methods. However, the methods specified by OCAP 1.0 and MHP 1.0.2 add the useful and necessary functionality needed by applications.

Both specifications share the following packages:

org.davic.media
org.davic.mpeg
org.davic.mpeg.sections
org.davic.net
org.davic.net.dvb
org.davic.resources
org.havi.ui
org.havi.ui.event

OCAP 1.0 does not currently utilize the following packages. Ongoing efforts to harmonize OCAP with MHP may effect this list:

org.dvb.net.ca
org.dvb.net.tuning
org.dvb.si
org.davic.net.ca
org.davic.net.tuning

The org.dvb.event package is not currently specified by OCAP 1.0 but is under consideration for addition to the specification via the ECR process.

Conclusions

The OCAP family of specifications were developed in response to specific needs of the cable industry. In order to minimize time to market for new services, and to enable downward price pressure and facilitate innovation via a retail market for cable receivers, OCAP provides a full featured and robust software platform. It offers a very high degree of portability and uniformity for content display as well as offering a platform for the broadest possible range of application support. The OCAP architecture ensures security and robustness.

References

Schmitt-Gordon, A., OpenCabel Application Platform Architecture. Proceedings NCTA, June 2001.

OpenCable Application Platform Specification, OCAP 1.0 Profile, OC-SP-OCAP1.0-I01-011221, found at http://www.opencable.com

Kar, M.L., Vang, S., and Brown, R., Architecture of Retail Set-Top Box Application Platform for Digital Cable Netork, Proc. ICCE, June, 2001.

Zundel, J-P, Emergence of Middleware in Home Telecommunication Equipment, IEEE Communications, June, 2001.

Draft Multimedia Home Platform Draft 1.0.2 for review, Document TM2208r7, TAM 232r29, can be found at http://www.dvb.org

Digital Video Broadcasting (DVB) Multimedia Home Platform (MHP) Specification 1.1, tm 2485, tam 668r12, can be found at http://www.dvb.org

ATSC DASE AEE, Doc. T3-530 09, Feb 2001, Rev 1.

Havi Level 2 User Interface, section 2.5.2, http://www.havi.org

Documents relating to the PE including DOM, CSS, HTML can be found at http://www.w3c.org

Documents relating to the Davic specification can be found at http://www.davic.org

Appendix – Syntactical Comparison of OCAP 1.0 and DVB-MHP 1.0.2

PACKAGES	total methods	# methods used by OCAP	# methods used by DVB-MHP	syntactical compliance with DVB-MHP %
org.dvb.*	488	357	488	73
org.ocap.*	207			
javax.*	414	414	414	100
java.* (see note 1)	3440	3440	3440	100
org.havi.*	*	785	785	100
org.davic.*	*	139	268	52
TOTALS		5135	5395	95
note 1 - number of java.* methods actually used by DVB-MHP is somewhat less than required by SUN for the implementation				

OPTIMIZING THE LAST MILE

Rice, D. and Schnitzer. J.
Stargus, Inc.

Abstract

Several sources of inefficiency common in current DOCSIS network deployments are identified and provided as motivation for cable operators to optimize their network and operations. Through an investigation of DOCSIS technology, two complementary methods for the optimization of last mile cable networks are proposed.

INTRODUCTION

Multiple System Operators (MSOs) can optimize their access networks for new and existing services by efficiently managing "last mile" network quality and capacity. Obtaining visibility into, proactively acting on, and quantifying service-affecting HFC issues can raise the quality of the broadband experience and ensure the success of advanced, IP-based services. These successes can be measured in customer satisfaction, revenue growth, efficient operational expenditures, and reduced capital expenditures.

Today many last mile networks are operating with impairments and inefficiencies that are tolerated by email, web surfing, and other less data intensive or less time critical applications. Because of the error correction built into DOCSIS™ (Data Over Cable Service Interface Specifications) and the resilience of Transmission Control Protocol (TCP), these errors are unobserved by and unaffecting to most Broadband customers. However, advanced, IP based services such as Voice over Internet Protocol (VoIP), video-conferencing and streaming audio/video are rendered inoperable by such errors and inefficiencies.

By using the capabilities built into DOCSIS, the inefficiencies can be minimized through configuration optimization, thereby allowing MSOs to defer capital expenditures, reduce operational expenses, and ready their infrastructure for advanced IP-based services.

SOURCES OF INEFFICIENCY IN THE LAST MILE

Congestion and latency have unnecessarily existed in some network segments since the first cable modem deployment. Conversely, due to a lack of visibility into real network issues and fast subscriber growth, many of today's networks have been over-engineered from a bandwidth capacity point of view, resulting in inefficient capital spending. Subscriber growth will accelerate with the addition of multiple Internet Service Providers (ISPs) and service tiers. Making optimal use of MSOs competitively superior resources of spectral bandwidth and a potentially high quality signal environment will be key to minimizing capital expenses for additional infrastructure and reducing operational expenses.

Error levels often get worse over time and are typically associated with transient noise and interference due to HFC plant problems. While email and web surfing mask these low error levels, these latent and worsening HFC plant problems are often undetected for months until customers become dissatisfied with the performance of their applications. A degraded subscriber experience results in operating inefficiencies such as subscriber churn, as well as heightened customer care and network maintenance costs.

MOTIVATIONS FOR OPTIMIZING THE LAST MILE

There are many motivations for optimization of the last mile network-

- **Harvesting bandwidth —** Models mapping packet size distributions to operational configurations and channel quality show that capacity can be increased 2 to 6 times over current commonly used DOCSIS 1.0 CMTS configurations. This additional capacity can be used to increase revenue by offering higher margin tiers of service and supporting more customers.

- **Deferring capital—**Capital expenditures for network elements such as CMTS equipment and transmission infrastructure are significant. These expenses can be postponed until absolutely needed through increased visibility and efficient use of current infrastructure and spectral resources.

- **Reducing operational costs—**Labor expenses associated with installing additional CMTS infrastructure, re-building RF combining and optical splitting networks, and transmission facility construction can also be avoided and delayed. Providing focused maintenance can result in reduced trouble call rate, reduced call handle time, and reduced truck rolls.

- **Increasing customer satisfaction—** Customer satisfaction is the result of managing capacity well and providing stable, reliable services.
    - Visibility of network issues and tracking of useful performance metrics can be used to continually manage and improve the quality of the product offering.
    - HFC networks change over time and faults will always occur. User experience can be optimized while fix agents expeditiously address problems thus repairing many faults before they become customer-critical.
    - As new services are deployed that require guaranteed bandwidth and knowledge of capacity resources, efficient management will be crucial
    - Errors due to degraded network quality will impact these advanced, IP based services and degrade them as much if not more than capacity or Quality of Service (QOS) issues. The error performance and the capacity of the network are inter-dependent and must be managed for a quality customer experience.

- **Enabling new services—**Technology is constantly changing. DOCSIS 1.1 and DOCSIS 2.0 contain many more methods for optimization than DOCSIS 1.0. DOCSIS 1.0 has many features that are not yet fully utilized by MSOs. These methods are exposed and should be used by MSOs for the purpose of optimizing the use of their spectral bandwidth and high signal quality network resources. To take advantage of the ability to offer QOS and symmetrical bandwidth, enabling new services, these methods for optimization need to be utilized efficiently. The underlying technology is becoming increasingly more complex and needs to be abstracted and automatically managed to avoid an increasingly expensive operations workforce.

A simple conservative cost model for an MSO considering only the savings in capital expenses as shown in row 1 Table 1 can defer significant capital infrastructure expenses over a 2-year period of $3.05 per CM per year. The operational expenses to perform the infrastructure upgrades or the unnecessary recombining and reconfiguration of networks before the current resources are actually fully utilized will easily outweigh the capital expense and increase these savings.

Some additional modeling based on quality optimization that improves customer experience reducing churn and call center costs along with more efficient use of truck rolls and fix agents is also shown in Table 1. The total savings per CM per year can be greater than $10.

Table 1: Operational Savings Through Capacity and Quality Optimization

Optimization Benefit	**Assumptions**	**Annual Value (subscriber/yr.)**
Deferred capital based on Capacity Optimization	• CMTS infrastructure retail cost: $32 to $36 per cable modem* • Large MSO purchase discount: 40% • Customer growth: 20% per year, 2 year average • Capacity optimization: 2X	$3.05
Churn reduction Improved customer satisfaction	• 5% of controllable churn • (.4% x $20/mo) value compounds • Number is 3 year average	$1.92
Truck roll reduction	• 50% call rate • 15% yield truck roll • 5% reduction x $50/truck roll	$2.25
Call center - Reduced calling rate	• 50% call rate • 5% reduction x $5/call	$1.50
Call center - Reduce call time duration	• 5% efficiency gain • (45% t/c rate x 5% x $5/call)	$1.35
Total (per data sub)		$10.07

*Based on recent vendor list price and approved re-seller quotes.

METHODS OF OPTIMIZATION IN THE LAST MILE

DOCSIS Cable Modems (CM), Multi-media Terminal Adaptors (MTA), Advanced Set Top Boxes (ASTB), and Cable Modem Termination Systems (CMTS) can all be utilized to detect and manage errors while providing bandwidth intelligence data, enabling the capacity optimization discussed in Method 1. The DOCSIS network can be configured to "four wheel drive" through most service affecting errors—errors that otherwise result in degraded or complete loss of service to subscribers--while maintaining optimal bandwidth capacity. This can be accomplished while notifying operators of degraded network quality before it becomes service impacting, as it occurs, and also providing isolation and identification of the faults.

The second complementary method to reduce inefficiency, discussed as method II below, addresses service quality optimization. In this case, we combat inefficiencies impacting broadband experience by considering all elements and conditions degrading service within the DOCSIS transport network. This method describes how user applications and services are mapped to all of these underlying network conditions to provide performance metrics that map directly to user experience.

Method I – Capacity Optimization

Communication Systems and the Capacity Optimization Trade Space

When designing or operating digital data communication systems, there are several goals that help drive system optimization:

Goal 1: Transmit as much data in the shortest amount of time possible through the system.
Goal 2: Transmit this high rate of data using as little of the physical resources (spectral bandwidth and power) as possible.
Goal 3: Transmit this data reliably at a much lower rate of errors than will impact the performance or reliability of any of the services.
Goal 4: Develop and operate this system with as little expense and complexity as possible

The challenge is that these four goals are not completely independent. To accomplish high transmission rate, it often requires more physical resources. Error performance can be traded off against physical resource utilization, but is best left fixed at very low error performance. High transmission rates at low error with less use of physical resources can be accomplished, but typically at much higher costs and complexity. [1]

Fortunately, the HFC network architecture has significant competitive advantage in its physical resources. The physical resources that an MSO has available to utilize are significant spectral bandwidth and power. The spectral bandwidth in the network can be up to 1 GHz and beyond in some cases. The power resources include the potentially high quality low noise power environment enabled by the use of optical transport and shielded coaxial transmission lines. These power resources also include the ability to transmit high-powered signals as long as they are within the dynamic range of the active components and the network meets FCC emission standards. Many of the other competitive transport technologies do not have anywhere near the advantage in these physical resources that MSOs have with HFC architectures.

A 3-dimensional model of the trade space that is used to optimize capacity and take advantage of the physical resources available for DOCSIS 1.0/1.1/2.0 is shown in Figure 1. The performance of the communication system is based on the trade off between error performance, channel quality, efficiency and capacity. The ability to send a high rate of data using as little spectral bandwidth as necessary is indicated by the Bandwidth Efficiency/Capacity axis. Sending this high rate of data and its relationship to signal and noise power is expressed on the channel quality axis. Sending this high rate of data reliably without errors can be expressed by the error performance axis. Choosing operating points within this space can often be determined by the cost and complexity of the communication system. [2]

In the middle of the last century, some Bell Laboratories information theory scientists determined the theoretical outer bound of what is physically possible in this trade space. [3] This outer bound of a digital communication system's capacity can be expressed simply using the Shannon-Hartley Capacity Theorem.

$$C = W * Log_2\left(1 + \frac{S}{N}\right)$$

$$Where,$$
$$C = Capacity\,(bps)$$
$$S = SignalPower$$
$$N = NoisePower$$
$$W = Bandwidth$$

$$W = \frac{1}{H^2(0)}\int_0^\infty H(F)df$$

$$H(F) = Frequency\ \mathrm{Re}\,sponseofCh\,annel$$

This theorem for the maximum capacity bound stipulates that there exists a Forward Error Correction (FEC) and modulation technique to allow information to be transmitted up to the capacity C, at an arbitrarily small error rate approaching 0. It also stipulates FEC or modulation technique does not exist that will allow information to be sent at a higher rate than C without errors.

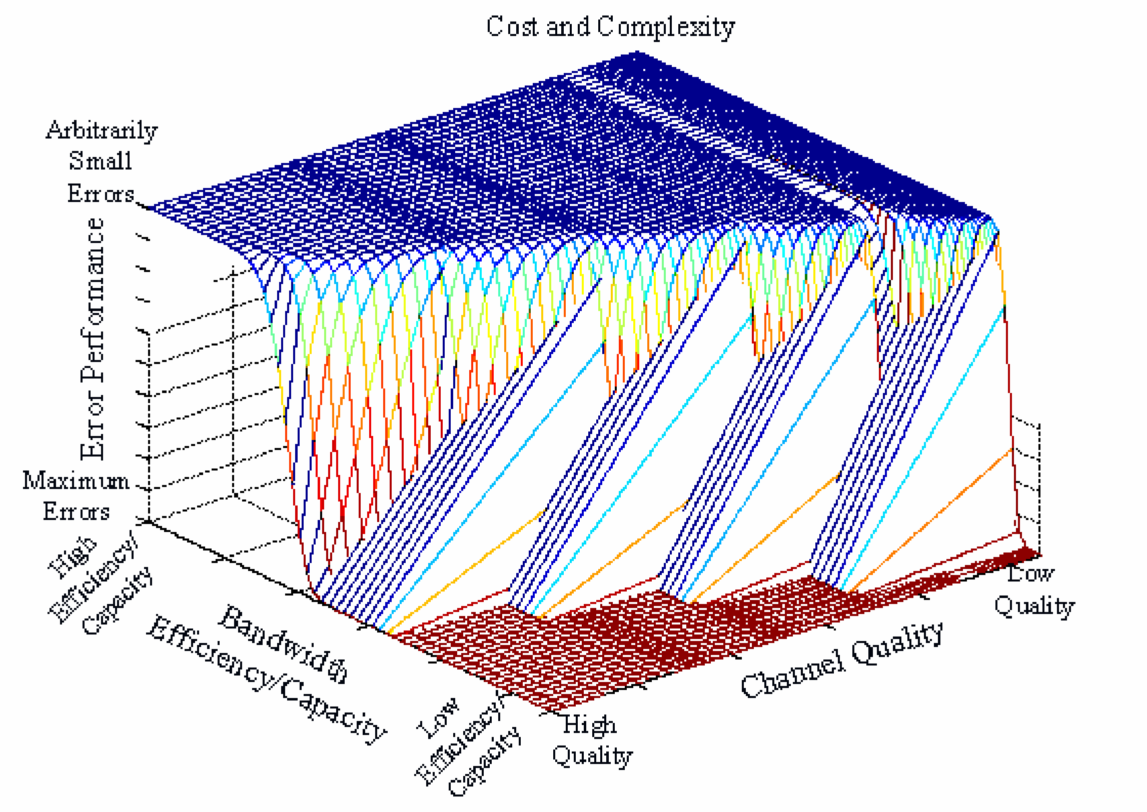

Figure 1: Capacity Optimization Trade Space

Mapping to DOCSIS

The DOCSIS 1.0/1.1/2.0 specifications developed at Cable Television Laboratories™ comprise a progression of techniques that allow operators to optimize capacity within the trade space of Figure 1 based on the services offered and the quality of the network. [4] If we refer back to the goals of designing a communication system, they are consistent with the goals of the DOCSIS program. The maximum data rates available have increased through time. The optimization against the use of physical resources has seen increasing efficiency. The complexity of operating the technology has increased significantly, while conversely, the cost of network equipment such as cable modems (CMs) and CMTSs have decreased due to production volume and competition. Due to this increasingly complex technology, capacity optimization methods will become important to deploy new IP-based services successfully.

Optimally, the network must operate at an error level small enough not to degrade the different applications and services. The communication system trade space concept can be simplified by bisecting this space in Figure 1 near the top of the Error Performance axis into a two-dimensional plane as shown in Figure 2. Based on the techniques for optimization in this space that are dictated by the DOCSIS specifications, a system can operate in the areas show in Figure 2 for the different versions of DOCSIS. Using the “knobs” provided by DOCSIS the capacity and efficiency can be increased for a given channel quality as shown by the vertical arrow. If the channel quality is sufficient it can be further increased to the maximum provided by DOCSIS. If this plane were a map of actual capacity, instead of efficiency (bits per second per Hz), the increases would appear much more pronounced since DOCSIS allows for use of increasingly wide frequency channels in the upstream to increase capacity. As DOCSIS evolves and the cost of more complex technology continues to drop, it will continue to extend closer to the theoretical boundaries described by Shannon and shown in Figure 2. This evolution will allow MSOs to make optimal use of the infrastructure they have in place as technology improves.

The ideal operating points within this space are those along the edge of the space yielding the highest capacity and efficiency relative to the quality of the channel, while maintaining errors at no-impact to performance levels.

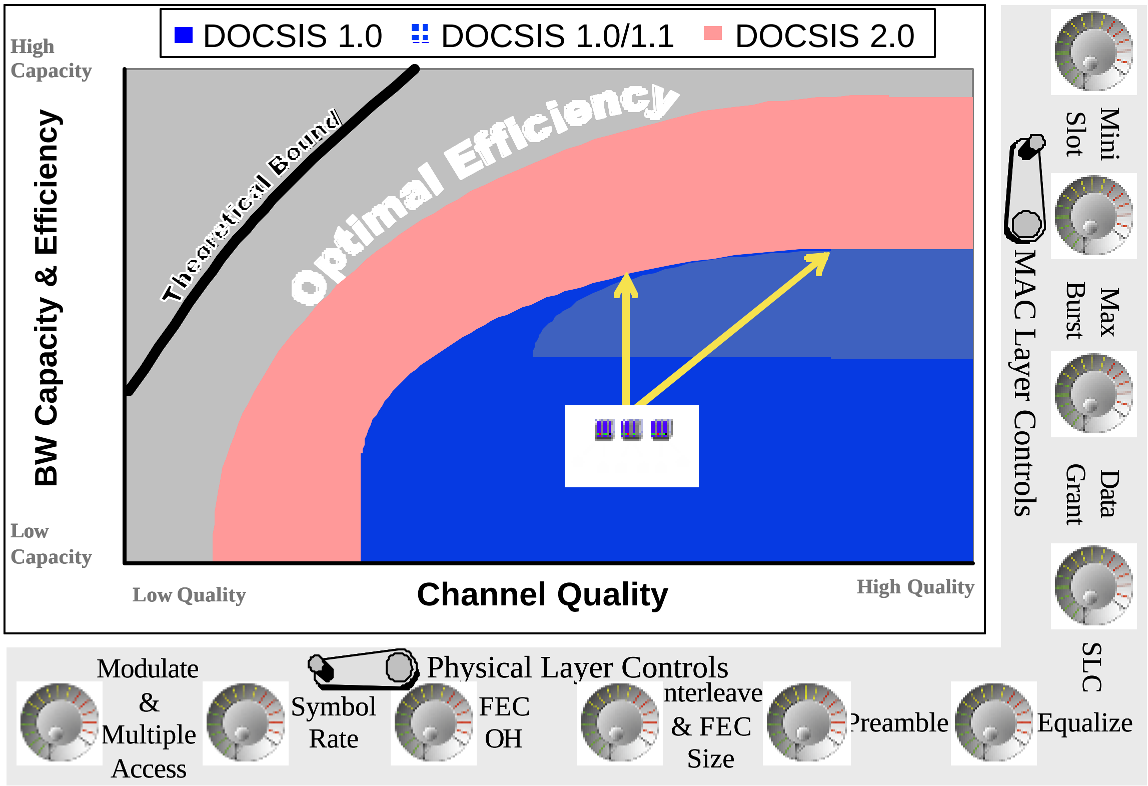

Figure 2: Efficiency and Capacity vs. Channel Quality

There are many different “knobs” and “levers” that are available in DOCSIS systems that can be tuned to allow this capacity optimization and movement in the plane show in Figure 2. They can be divided into three categories: Physical layer Capacity, MAC layer efficiency, and traffic scheduling. All three of these categories are related to each other and optimizing all aspects can result in the significant gains in capacity discussed above. While short-term traffic scheduling and power adjustments are the domain and responsibility of the network infrastructure, longer-term quality control through optimal selection of network operating parameters based on performance over longer time periods is the domain and responsibility of MSOs' management systems.

The Physical Layer and MAC Layer knobs which should be adjusted by operators to maximize the performance and provide a solid basis for the novel scheduling algorithms being developed by CMTS infrastructure vendors are listed below. Different knobs are available for the downstream and upstream direction and setting of these parameters needs to be considered as a collective. For example, increasing the symbol rate without optimally setting the mini-slot size will result in much less capacity gain than would be expected. Additionally, setting mini-slot size incorrectly can make large PDUs unable to be transmitted.

- Physical Layer capacity
  - Spectrum Bandwidth/Symbol Rate, Modulation order, Multiple access type, FEC Type, FEC Overhead, FEC Codeword size, Interleaver mode and depth, preamble length and unique word, Equalization
- MAC Layer Efficiency
  - Mini-slot Size, Maximum Burst Size, Long Data Grant, Short Data Grant, Shortened Last Codeword, Extended header size, Upstream Channel Change
- Scheduler tools
  - Unsolicited Grant Service, Dynamic Service methods, Polling Services, Fragmentation, Concatenation, Payload Header Suppression, and contention vs. unicast periods.

There are also many inputs listed below that must be considered and utilized when optimizing capacity. DOCSIS provides a rich set of standard management information through the Simple Network Management Protocol (SNMP) and a management information base (MIB). Other methods of network information collection also exist and are provided for in the DOCSIS specifications.

A rich set of vendor proprietary information also exists for harvesting. This information on network quality and performance along with current configuration of the knobs discussed above needs to be normalized across the network elements and analyzed to determine the optimal configuration for the network elements. With this analysis, MSOs can maximize capacity while maintaining the arbitrarily small error performance.

This analysis should consist of real time data collection from the network, complex analytical computer models that can be parameterized for real time use, empirical evidence from laboratory experimentation, and field experience all based on vendor implementations. All of these techniques should be used to determine the optimal configuration for maximum capacity.

- Physical Layer Inputs
  - Channel Quality –
    - Noise Type and statistics
  - Channel frequency magnitude and phase response
  - Power Levels
  - Available Spectrum
  - Error Performance
    - Packets, channel symbols, and codewords
  - CMTS Configuration

- MAC layer efficiency Inputs
    - MAC Frame Structure
    - MAC Header
        - EHDR including BPI, Req, zero pad, PHS, etc
    - Packet Size Distribution
        - VOIP, WEB, FTP, PHS
    - Channel Utilization
    - Total customers and active customers
    - CMTS Configuration

As shown in **Figure 3**, the inputs listed above and statistical analysis of the data, along with the creation of performance metrics, can be utilized to optimize Capacity. These analyses and calculations are based on the physical layer inputs described above, using the appropriate knobs to optimize for channel quality along with efficiently packing the data into the MAC layer based on the type of traffic and the mapping of DOCSIS MAC frames to the optimal physical layer, security, and MAC layer configuration. Using a capacity optimization system such as this will ensure excellent customer experience and will also ensure that the MSO is efficiently utilizing their physical resources and efficiently investing capital.

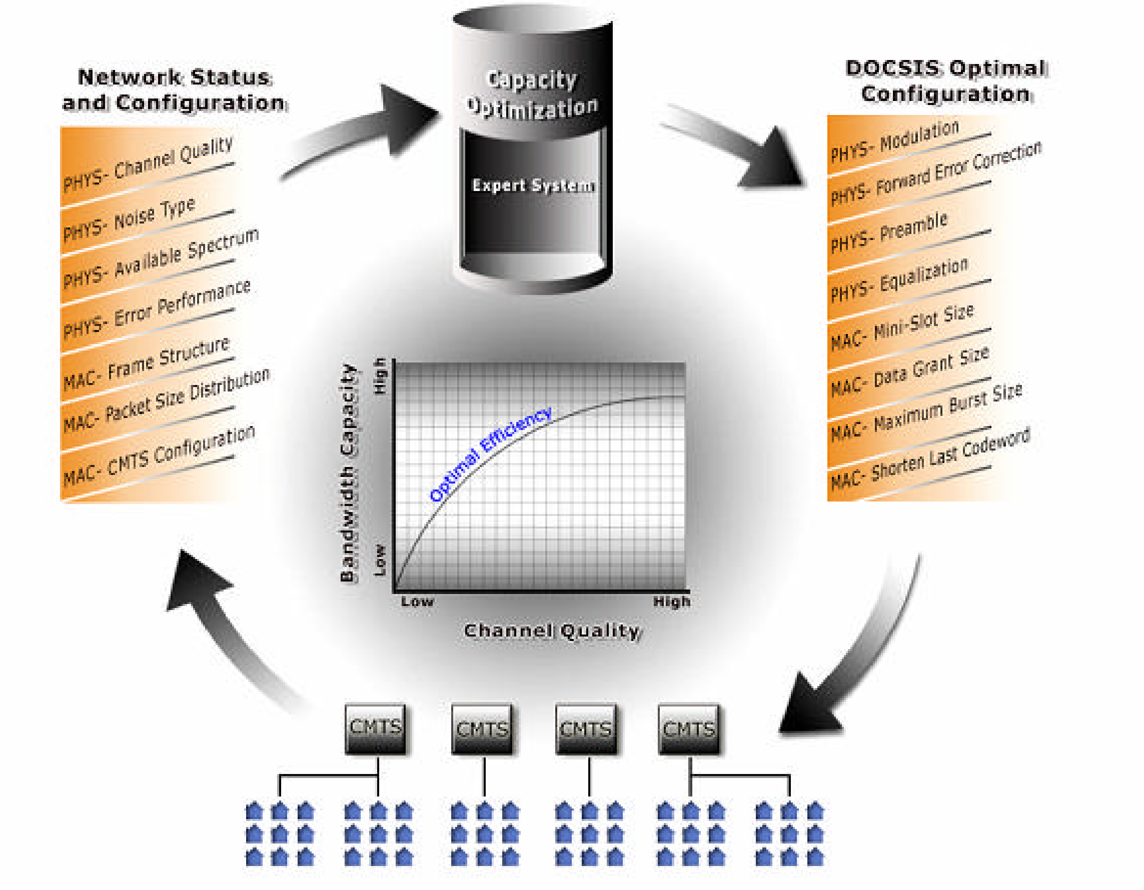

Figure 3: Capacity Optimization System.

Method II – Quality Optimization

Complementary to capacity optimization in the last mile, the second method to reduce inefficiency addresses service quality optimization. In this case, we combat inefficiencies impacting broadband experience by considering all elements and conditions degrading end users' perception of the service within the DOCSIS transport network.

Current approaches to monitoring quality using E/NMS (Element/Network Management System) technologies are limited in the following ways-

- Network element centric—Visibility into the condition of infrastructure elements provides a network-layer perspective of quality. Missing is a correlation between overall network status and its impact on the service experience of the subscribers.
- Undetected problems–Issues caused by complex or composite impairments affect the user experience but remain undetected by the monitoring system.
- Fault and performance centric--Problems detected by the monitoring system may have no discernable impact on subscriber experience.

To address problems unique to optimizing service quality management in DOCSIS networks a novel approach is required. The methodology described uses passive measurement to gather specific data related to network and infrastructure condition then correlates the value of each point of measurement to service quality. In turn, each correlation of network measurement to service quality is combined into a single composite figure that indicates the quality of the subscriber's experience. In this way, the root or composite cause of service degradation in the DOCSIS transport network can be quickly identified, topologically isolated, then cured thus optimizing service quality.

Quality optimization is ultimately a function of many points of measurement made in the DOCSIS network. The two general mechanisms for gathering of network measurements are active and passive-

- Passive monitoring relies on management instrumentation furnished by SNMP (Simple Network Management Protocol) agents embedded in network elements (CM and CMTS). These agents implement the suite of Management Information Bases (MIBs) required by the DOCSIS Operations Support System Interface (OSSI) standards [4].
- Active monitoring, the competing approach, is based on the introduction of synthetic traffic in an attempt to emulate service behavior. In this way the results of test traffic provides a sample of application and network quality.

Because the addition of traffic to the network through active measurement potentially exacerbates degraded conditions, passive measurement has been selected as the preferred mechanism for data collection. In addition, due to the "bursty" nature of traffic and impairments in DOCSIS networks, stochastic events may go undetected by scheduled active measurement.

The quality of the overall broadband experience is heavily weighted by the condition of the last mile infrastructure supporting it. Where DOCSIS provides this last mile transport of IP based services, two general areas reflect overall network quality-

- Connectivity—The physical condition of the DOCSIS connection between each CM and the CMTS. Contributors affecting connectivity include the health and configuration of the HFC plant as well as the state of the hardware resources terminating the connection.

- Capacity—The impact of traffic over the connection on latency and resources sensitive to network scheduling and loading. Contributors include downstream and upstream interface utilization, Media Access Control (MAC) domain loading, NSI utilization, queue depths, processing resources, and framing efficiency.

There exists no single point of measurement within a DOCSIS network that can be used as a proxy for overall network condition or service quality. Instead, each of the measurement points becomes a contributor to an estimation of connectivity or capacity. For each contributor, the opportunity exists to correlate the value of the raw network measurement to the effect it has on higher protocol layers. If a relationship between the measurement and overall service quality can be described, this correlation can be applied to the monitoring of an operational network transporting the service.

To illustrate, consider the quality of IP telephony service as a function of DOCSIS 1.1 Physical (PHY) layer response to HFC signal impairments. In terms of the protocol stack, upstream telephony data is packetized by first encapsulating it in a Real-Time Protocol (RTP) based UDP datagram. In turn, the datagram is wrapped in an IP packet and encapsulated in a DOCSIS frame. The DOCSIS frame is then embedded in a Reed-Solomon (RS) codeword and transmitted over the HFC network.

Packet loss is the most damaging form of impairment for IP telephony [5]. In the DOCSIS segment this loss is a function of either degraded capacity or connectivity. An operationally common cause of degraded connectivity is DOCSIS frame loss resulting from HFC impairments. This is expressed through an SNMP agent in either the CM (downstream) or CMTS (upstream) and reported as uncorrectable Reed-Solomon (RS) FEC. The theoretical relationship assuming ergodicity between Codeword Error Ratio (CER) and DOCSIS Frame Loss Ratio (FLR) is illustrated in Figure 4. The example VOIP Packet would reside in these bounds based on things such as EHDR length and PHS that impact packet size.

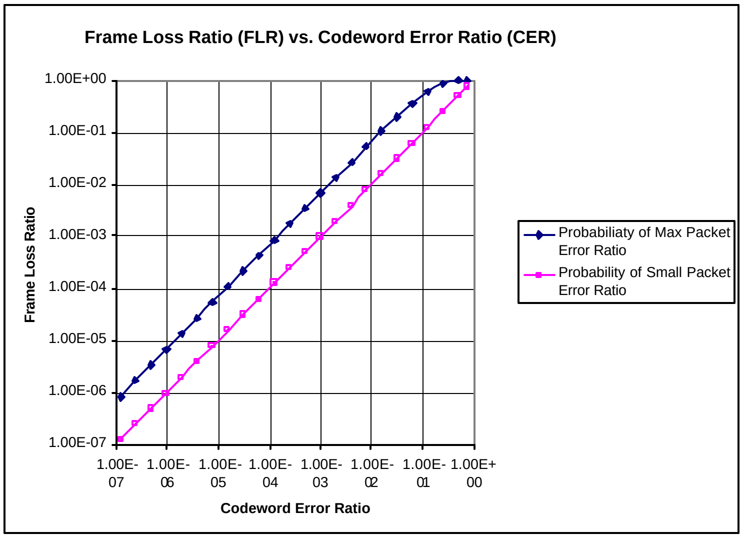

Figure 4: Frame Loss Ratio vs. Codeword Error Ratio

Holmes, Aarhus, and Maus [6] provide experimental results comparing frame loss to MOS (Mean Opinion Score) in an effort to describe a relationship between underlying network transport and subjective assessment of IP telephony service quality. Their findings suggest that with as little as 3% packet loss, audio quality degrades to an estimate of “fair”.

Assuming that through capacity optimization, IP telephony packets in the upstream will occupy short-data DOCSIS frames, a one-to-one relationship between layer 3 IP packets and layer 2 DOCSIS frames exists. Through understanding this relationship, the correlation between an SNMP-based network measurement of physical layer impairment (CER) and the subjective quality of an IP telephony service on the application layer can be made.

In order to quantify the result of this correlation, we introduce the Degraded Modem (DM) metric. A DM event occurs whenever the value of a contributor is correlated to an instance of degraded service quality provided by a single CM. In the example provided, a degraded IP telephony event (of “fair”) was defined at a contributed CER value of 3.0x10E-2.

We capture this relationship between the contributor and Degraded Modem event using the correlation function illustrated in Figure 5. In this example, the correlation function coupling CER and DM is based on both theoretical relationships between CER and DOCSIS frame loss, as well as experimental data correlating IP packet loss with a subjective score of IP telephony audio quality.

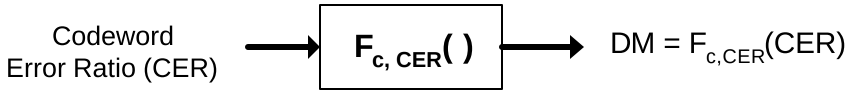

Figure 5: Correlation Function for CER Contributor

Likewise, for each contributor collected in the network a function can be described that correlates the state or value of the contributor to the degraded modem figure. In this way, all contributors and their associated DM figure can be combined using a logical OR operation to provide a composite estimate of service quality provided by the DOCSIS transport.

Figure 6 illustrates a simple combiner structure as an aggregation of outputs from an array of contributor correlation functions. In this case, the estimation of service quality (DM) is based on input from a total of six contributors. The combiner structure is extensible in order to accommodate contributors and correlation functions associated with the evolution of DOCSIS technology and the introduction of new IP based applications and services.

The combiner structure can be imposed on network topology in order to isolate the areas of the network introducing degradation. DM can be calculated and applied for any physical, logical, or organizational node within a MSO’s DOCSIS network. The combiner provides a holistic approach to analysis of DOCSIS network quality by providing a method to estimate the health of the subscriber experience while abstracting the complexity of the underlying DOCSIS infrastructure.

If applied within the operational environment, the combiner provides a promising approach to quickly isolating and identifying conditions within the last mile infrastructure that is driven by the subscriber’s experience.

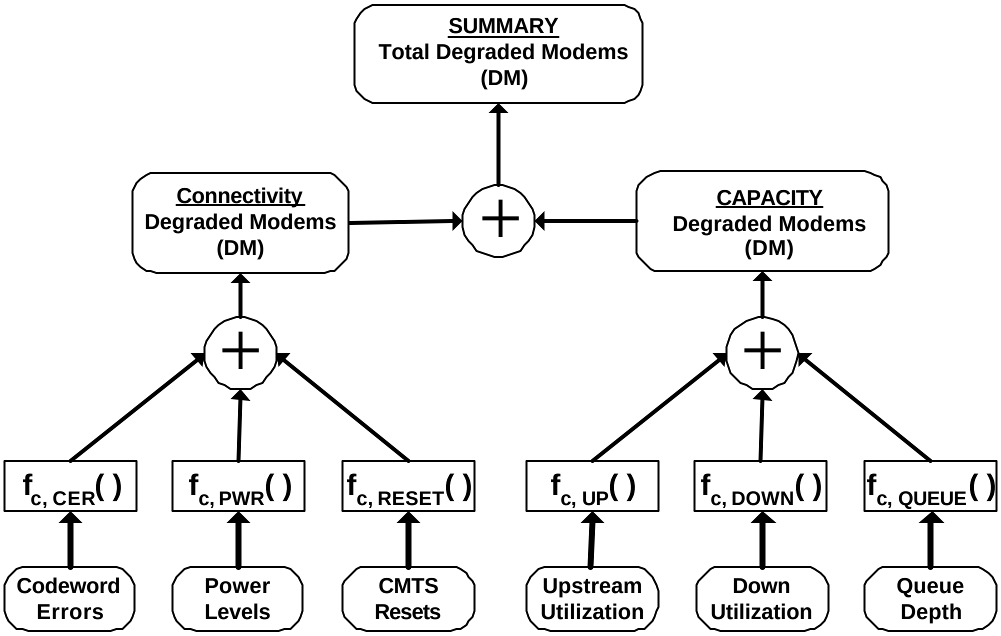

Figure 6: Simple Combiner Structure

SUMMARY AND CONCLUSIONS

Utilizing DOCSIS 1.0 and existing infrastructure, significant gains in capacity enabling new revenue and significant reduction in capex and opex can be obtained with a system that automates capacity and quality optimization. As DOCSIS technology evolves with versions 1.1 and 2.0 along with new DOCSIS devices, the customer base will experience exponential growth in the number of intelligent network elements. A scalable and reliable system that provides visibility, performance metrics and automated capacity optimization will become critical to deploying new IP-based services over the HFC network. Technology will continue to evolve, allowing new services and higher capacity and efficiencies. Costs of infrastructure will drop with volume, but complexity of operations will continue to increase and methods such as those discussed here will be essential to continued IP services product evolution.

REFERENCES

[1] B. Sklar, Digital Communications Fundamentals and Applications, Prentice Hall, 2001

[2] Performance Measures for Cable Data Transmission, NCTA technical papers, Rich Prodan, Dan Rice, Majid Chelehmal, 1999

[3] Shannon, C.E., "A Mathematical Theory of Communication," BSTJ Vol. 27, 1948, pp. 379-423,623-657

[4] http://www.cablemodem.com/specifications.html

[5] "Modeling the Effects of Burst Packet loss and Recency on Subjective Voice Quality", A.D Clark, Ph.D, 2001 IP-Telephony Workshop

[6] "Network Tolerance of Highly Degraded Conditions for an H.323-based VoIP Service", Peter Holmes, Lars Aarhus, Eirik Maus, Norwegian Computing Center, P.O. Box 114, Blindern, Oslo, Norway

The authors can be contacted at Stargus, Inc. Andover, MA 01810, or {dan,jason}@Stargus.com

OPTIMIZING TRANSMISSION PARAMETERS IN DOCSIS 2.0 WITH A DIGITAL UPSTREAM CHANNEL ANALYZER (DUCA)

Noam Geri, Itay Lusky
Texas Instruments, Broadband Communications Group

Abstract

A new generation of data over cable service interface specification (DOCSIS) 2.0 cable modem and cable modem termination systems (CMTS) offer cable operators the promise of increased upstream capacity and greater robustness to common channel impairments such as ingress and impulse noise. It is already clear that the many tools in the new DOCSIS 2.0 standard that allow for efficient use of the upstream spectrum and mitigation of impairments also make the task of optimizing transmission parameters increasingly difficult. In fact, the performance of a DOCSIS 2.0 based CMTS will greatly depend on its ability to dynamically assess upstream channel conditions and set the transmission parameters accordingly.

In this paper we present digital upstream channel analyzer (DUCA) – a set of functions running on a DOCSIS 2.0 CMTS that implements algorithms for optimal channel allocation and selection of transmission parameters. DUCA analyzes the entire upstream spectrum, measures and records noise and impairment conditions, and sets the parameters of the various noise mitigating tools in DOCSIS 2.0 optimally for maximum upstream throughput.

We will show how proper selection of parameters, using DUCA, ensures that operators will benefit significantly from the new improved upstream PHY.

INTRODUCTION

After several years of ongoing debate, cable operators have selected advanced time division multiple access (A-TDMA) and synchronous code division multiple access (S-CDMA) as the upstream modulations in the new DOCSIS 2.0 specification. These technologies offer cable operators the opportunity to better utilize their cable infrastructure and to generate more revenue from increased use of the cable network upstream spectrum. DOCSIS 2.0 offers operators powerful tools to mitigate common channel impairments and spectrally efficient modulations to maximize the throughput in the bandwidth-limited upstream channel. However, the many tools in DOCSIS 2.0 make the selection of transmission parameters extremely difficult in comparison to DOCSIS 1.0, with the performance of DOCSIS 2.0 systems greatly depending on the choice of these parameters. In fact, DOCSIS 2.0 will only provide significant benefits to operators if and when CMTS systems make proper use of the many tools in this standard by implementing technology that dynamically sets transmission parameters for optimal performance. In this paper we will present such a technology – DUCA, which measures the impairments in the upstream channel and sets the transmission parameters for maximum throughput based on time-domain and frequency-domain analysis.

THE CABLE UPSTREAM CHANNEL

The cable network upstream channel has always been the weakest link in the cable network infrastructure. Given the tree-and-branch topology of the cable network, noise and interferences from the entire network are accumulated at the headend. Common upstream impairments include the following noise sources:

1) White noise generated by active components in the network.
2) Narrowband ingress noise, typically generated by other transmitters such as

amateur radio signals or resulting from Common Path Distortion [2].

3) High rate impulse noise originating from electric current. These impulses are short, typically less than one microsecond duration, and have a repetition rate of between several hundred to a few thousand occurrences per second.
4) Low rate wideband burst noise originating from several sources including electrical appliances in homes and laser clipping. These bursts could occur as frequently as every 10-20 seconds and could last as long as 10-50 microseconds.

In addition to the noise sources described above, the upstream signal is subject to multi-path reflections due to impedance mismatch of the plant's components and unterminated cables. For a more detailed description of cable upstream impairments see [1][2].

DOCSIS 2.0 BACKGROUND

In August 2001, cable operators decided that the new DOCSIS 2.0 upstream physical layer specification would include both A-TDMA, based on a proposal by Texas Instruments and Broadcom [1][3] and S-CDMA based on a proposal by Terayon. Both of these technologies were also included in the IEEE 802.14a specification, which was never finalized [4].

A-TDMA is essentially an evolution of DOCSIS 1.0. It extends the physical layer of DOCSIS 1.0/1.1 with the following enhancements:

1) Additional constellations: 8-QAM, 32-QAM and 64-QAM. This allows an increase in spectral efficiency by as much as 50 percent in good quality channels and provides more increments in spectral efficiency for finer matching of data rate with existing channel SNR.
2) Additional Symbol Rate 5.12 MB. This reduces the number of receivers required at the headend for a given plant by a factor or two and improves network efficiency due to statistical multiplexing of more users in an upstream channel.
3) Byte Interleaver to spread the effect of impulse and burst noise over time.
4) Improved error correction code. DOCSIS 2.0 extends the maximum error protection ability of DOCSIS 1.0's Reed-Solomon FEC from 10 byte errors to 16 byte errors, providing greater robustness to burst and impulse noise.
5) Improved Pre-Equalizer for mitigating multipath distortions.

S-CDMA adds to the above enhancements a spreader that provides greater immunity to severe cases of impulse noises, and Trellis Coded Modulation, which improves performance for white noise. When in S-CDMA mode, there is no byte interleaver as described above. Instead, an S-CDMA framer introduces time (as well as code) diversity. S-CDMA calls for much stricter timing requirements in order to maintain code diversity, allowing for the elimination of guard time between data packets. For more details on S-CDMA see [4].

SETTING TRANSMISSION PARAMETERS

DOCSIS 2.0 provides a new challenge in setting transmission parameters. While DOCSIS 1.0 provided operators with some limited flexibility in setting transmission parameters to match the varying channel conditions, in practice, parameters remained relatively static. Without the ability to track dynamic changes in the plant, operators had no choice but to set transmission parameters to the most robust mode (QPSK, Reed Solomon T=10) to accommodate worst-case scenarios. With penetration still low, such inefficient use of the upstream spectrum could be tolerated. Without ingress cancellation available to them, operators would typically set the frequency manually to ensure the transmission signals are within a region with little or no ingress. The more sophisticated CMTSs could automatically identify that ingress is interfering with the data signal and automatically shift modems to a different upstream frequency with no interference.

DOCSIS 2.0 requires a much more sophisticated setting mechanism. First, there are many more parameters to play with, such as modulation type (A-TDMA or S-CDMA), constellation, baud-rate, transmission power, preamble length and type, center frequency, error correction capability, interleaver parameters, spreader parameters and number of active codes in S-CDMA mode. Second, the premise of DOCSIS 2.0 is that the upstream traffic is significantly higher, with upstream channel throughput closer to capacity, leaving less room in the spectrum to avoid interferences, and making it crucial to efficiently utilize the channel spectrum.

FREQUENCY-DOMAIN ANALYSIS

The first step in setting optimal parameters is measuring channel conditions and detecting interferences. The most common tool in current CMTSs is upstream spectral analysis. Using wideband sampling and FFT, or alternatively using a frequency-sweeping filter, the upstream spectrum can be measured, identifying frequencies with ingress. This spectrum measurement is typically used to find ingress free regions for the data signals. However, with ingress cancellation technology, first introduced by Texas Instruments in the TNETC4521 INCA burst receiver (see also [5]), avoiding the ingresses is no longer necessary. While transmitting in an ingress free region is always desirable, a clean spectrum block, which is wide enough to accommodate the highest baud-rate, is not always available. In such cases a CMTS needs to make a decision on whether to reduce baud rate, allowing the signal to fit between other signals and interferences or to maintain the high baud rate and to cancel the interference with ingress cancellation technology. Given that ingress cancellation techniques allow for operation in negative C/I ratios (i.e. ingress that is stronger than the data signals), it is foreseeable that in many cases the parameter setting mechanism will determine that maintaining the higher baud rate while overlapping the ingress will result in higher throughput than if the baud rate were reduced and the ingress avoided. Ingress cancellation technology and the new modes of operation in DOCSIS 2.0 have transformed the traditional spectrum analysis of finding ingress free regions into a more complex optimization problem of setting baud rate, center frequency, constellation, coding and other parameters to maximize upstream throughput given the constraints of available spectrum, detected ingress and the performance of the ingress cancellation technology. Furthermore, as channel conditions change, these transmission parameters need to be adapted to the new environment. Tracking spectrum changes in an upstream channel densely occupied with data signals, and having to change in some cases the center frequency, baud rate and other transmission parameters of multiple upstream data signals concurrently in order to achieve higher throughput makes this ongoing optimization problem particularly challenging.

TIME-DOMAIN ANALYSIS

DOCSIS 2.0 provides new tools for mitigating impulse and burst noise: Byte Interleaver, stronger Reed Solomon error correction, S-CDMA spreading. In order to avoid unnecessary waste of bandwidth on a spectrally inefficient constellation or on coding overhead, the DOCSIS 2.0 CMTS needs to dynamically track impulse levels, and to optimally set the relevant parameters accordingly. Impulse strength, as well as impulse frequency and arrival statistics can be determined by employing various power detectors, which measure the signal level during quiet periods or in adjacent unoccupied frequencies. Finding quiet periods of time or unoccupied frequencies for measuring impulses may not be easy when operating close to channel capacity. In such cases the CMTS may have to regularly block time slots for impulse detection. To avoid wasting bandwidth on impulse detection, impulses can also be detected by analyzing decision errors, however this method is problematic since error measurements will be erroneous during impulse occurrences (because the error measurement relies on an incorrect decision). To overcome this problem, transmitted symbols and decision errors can be estimated by re-encoding corrected data bits after the Reed Solomon decoder. However, this results in a relatively complex algorithm.

A DOCSIS 2.0 CMTS has multiple tools for impulse mitigation. The spreading function of S-CDMA spreads the effect of the impulse over time and over the code space. This is a useful tool when impulse levels are limited, however if the impulse is very strong, spreading may actually decrease performance by causing multiple errors from every impulse (due to spreading) instead of taking the hit only once. In addition to spreading, Reed Solomon parameters are also the obvious candidates for adjusting based on measured impulse rates. Less intuitive, is the choice of baud-rate and constellation. Traditionally, the most common reaction to impulse noise in the channel is reducing baud-rate and reducing the constellation size, which indeed makes the signal more robust to moderate impulses. However, this comes at the expense of upstream throughput. A better approach may actually be to transmit at a high baud-rate using one of the larger constellations, thereby allowing more coding information, which will enable impulse mitigation with Reed Solomon coding. Various factors such as impulse power, impulse frequency and upstream channel utilization will affect the choice of these transmission parameters.

The tools for mitigating burst noise are the same ones used for impulse noise. Spreading provides good immunity to long bursts of noise. Reducing baud-rate can provide very strong immunity to very long burst noise even without spreading. However, given that long bursts (over 10 microsecond) are relatively rare, it may be better to transmit at high spectral efficiency with little coding overhead and sacrifice the occasional data packet instead of using a more robust mode with lower throughput. These are the types of trade-offs that the channel analysis function in a DOCSIS 2.0 CMTS needs to consider when setting transmission parameters.

MITIGATING OTHER IMPAIRMENTS

The most common impairment is the added white noise. Dealing with white noise is rather straightforward – setting the constellation size based on SNR measured with the upstream spectral analysis or by averaging the decision errors. Reed Solomon coding parameters also need to be set according to the measured SNR. When the SNR is low, a DOCSIS 2.0 CMTS may choose also to reduce the number of active codes in S-CDMA mode, or equivalently, to reduce the baud rate in A-TDMA mode and to allocate higher spectral density to the reduced baud-rate signal. In both of these cases, modems can operate at very low SNRs.

A MIX OF IMPAIRMENTS

A greater challenge for the DOCSIS 2.0 CMTS is when it is faced with the task of mitigating different types of noise simultaneously, especially when the optimal choice of parameters for each impairment are very different. For example, when ingress is combined with burst noise, the DOCSIS 2.0 needs to choose between a higher baud-rate that will improve the performance of the ingress cancellation, or a lower baud rate for greater immunity to long bursts. It needs to decide whether spreading will be used, providing greater immunity to bursts, but at the same time making ingress cancellation very difficult. Analyzing the mix of impairments, understanding the trade-offs and selecting the compromise set of parameters, which will provide optimal robustness to the measured impairment, while at the same time maximizing throughput, is essentially the role of the parameter decision function that a DOCSIS 2.0 CMTS needs to implement.

DUCA

We have presented the challenge that the cable industry faces in taking advantage of the new state-of-the-art DOSCIS 2.0 standard. It is likely to take several years before system vendors implement sophisticated detection and analysis tools, which make optimal use of the many tools provided by the new standard. Today, several years after the first DOCSIS 1.0 systems were certified, systems are still not realizing the full capability of this standard. DOCSIS 2.0 is likely to follow a similar path. To accelerate this process, Texas Instruments has introduced the concept of the DUCA, which is a functional block in a DOCSIS 2.0 CMTS dedicated to upstream channel measurement and analysis and optimal parameter selection. DUCA, which is best implemented using a dedicated Digital Signal Processor (DSP), performs time-domain and frequency-domain analysis as described in this paper, and dynamically sets the transmission parameters for optimal use of the upstream channel.

The following simulation is an example of upstream channel analysis performed by DUCA:

We simulated an upstream channel with multiple ingress as illustrated in Figure 1.

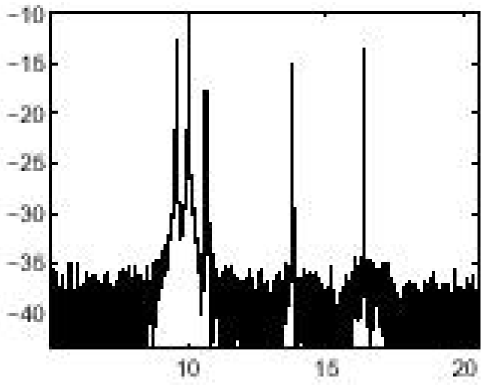

Figure 1: Upstream Channel Spectrum

In addition to ingress, this simulated channel is also corrupted by time-domain impairments, such as burst noise, which cannot be seen in the frequency-domain analysis.

Without DUCA capabilities in the CMTS, dynamic changes of the channel cannot be tracked, and therefore a robust mode, which can operate in worst-case scenarios, needs to be used. A typical choice of parameters for such a channel would include QPSK constellation, strong RS code and medium/low baud rate (2.56 Mbaud or even 1.28 Mbaud) to avoid in-band ingress noise. This results in upstream throughput of ~2.5-5 Mbit/sec, far below the optimum. Therefore, DUCA enables higher throughput in this channel. The DUCA algorithms identify and characterize the channel impairments (WGN, burst and impulse noises, ingress noise etc.), while taking into consideration ingress cancellation and other noise mitigation capabilities of the receiver. The impairment characterization is followed by an optimal channel allocation algorithm. Figure 2 shows the output of the DUCA channel allocation algorithm for one and two upstream channels. Note that for one upstream channel, the channel allocation algorithm determines that the highest throughput can be achieved by using the highest baud-rate and a 16-QAM constellation while overlapping two ingresses. The channel allocation algorithm determines that avoiding the ingress by reducing the baud-rate would not result in higher throughput even if a more spectrally efficient constellation can consequently be used. The transmission parameters selected result in upstream throughput of ~20Mbit/sec, a 4X-8X improvement compared to the over-robust transmission in the CMTS without DUCA.

SUMMARY

DOCSIS 2.0 gives cable operators a multitude of transmission parameters to define, such as modulation type (A-TDMA or S-CDMA), constellation, baud-rate, transmission power, preamble length and type, center frequency, error correction capability, interleaver parameters, spreader parameters and number of active codes in S-CDMA mode. Maximal channel throughput can only be achieved by using sophisticated mechanisms that optimally track and analyze the varying channel conditions and set the transmission parameters for optimal performance.

We have presented the concept of DUCA for optimal selection of those parameters in DOCSIS 2.0. We believe that channel analysis and parameter setting tools like DUCA will become more and more important as upstream data traffic increases, and over time more and more channel analysis algorithms will be developed and improved, enabling operators to realize the full potential of DOCSIS 2.0 and the cable upstream channel.

REFERENCES

[1] "HI PHY LITE – A Pragmatic Approach to Advanced PHY" Ofir Shalvi, Noam Geri et al, NCTA 2001 Conference Proceedings

[2] "HFC Channel Model Submission", Thomas J. Kolze, IEEE 802.14a/012, May 26 1998

[3] "Advanced TDMA Proposal for HFC Upstream Transmission", Texas Instruments-Cable Broadband Communications and Broadcom Corp., Oct 1999

[4] "S-CDMA as a High-Capacity Upstream Physical Layer", Mike Grimwood and Paul Richardson (Terayon Communications Systems), IEEE 802.14a/98-017, July 7, 1998

[5] "Advanced Modulation Schemes For Cable TV Upstream Channel", Ofir Shalvi and Noam Geri, ICCE Conference Proceedings, June 2000

[6] INCA Technology White Paper, Texas Instruments –CBC, www.ti.com/sc/docs/innovate/cable

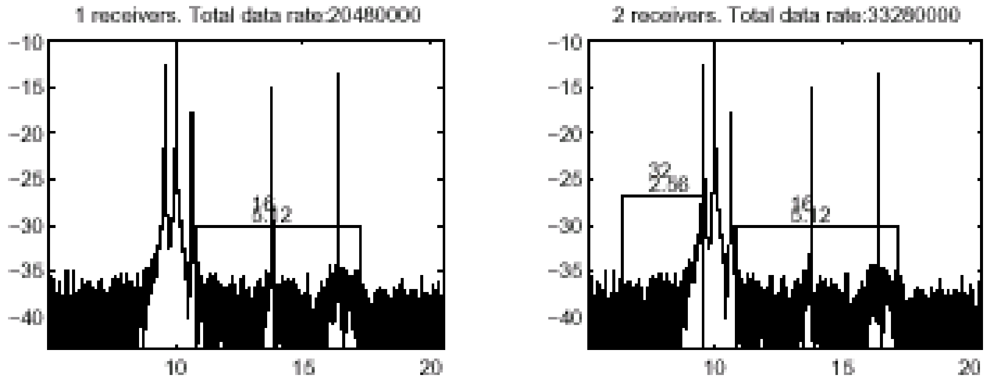

Figure 2: DUCA Channel Allocation Output

PER-FLOW QoS ENABLES TRANSACTIONAL BANDWIDTH MANAGEMENT

Gerry White
Chief Technologist
Network Infrastructure Group
Motorola Broadband Communications Sector

Abstract

Transactional bandwidth management holds great promise for the cable industry by allowing improved use of resources through statistical multiplexing and dynamic bandwidth assignment. This must occur in an environment in which Quality of Service (QoS)-enabled applications such as voice place very specific requirements on the network equipment and communication links. Any transactional changes must take place such that these applications continue to function correctly. The price to achieve this is increased complexity in management and control because bandwidth and QoS changes occur in real time. Per- flow queuing can be used to simplify the management and control associated with these changes and provide a deterministic mechanism to implement transactional bandwidth management. This paper explains the importance of per-flow QoS to the successful deployment of transactional bandwidth services and demonstrates how QoS can be provided end-to-end for transaction-based services

INTRODUCTION

As applied in a next generation cable network, transactional bandwidth management is more than simply changing the bandwidth required by a user or application. In many cases, services such as voice, video and streaming audio requires all the attributes of QoS for an application to be changed dynamically. These attributes include not only basic bandwidth but potentially latency, jitter, service interval, and both minimum guaranteed and maximum allowed data rates.

The ability to dynamically apply QoS treatments will allow cable operators and their revenue-sharing partners to offer exciting transactional bandwidth management services where subscribers can automatically provision increased bandwidth for applications such as video conferencing, interactive gaming, or Video on Demand (VOD). Operators can allow the applications themselves to trigger requests for more access bandwidth and application-specific QoS configuration. The self-provisioning of high-margin services is dependent on the ability to automate QoS control across the access network, metropolitan network and the core network of multiple providers.

Bandwidth is and will continue to be a scarce resource, but the Internet Protocol (IP) can be used to support the QoS requirements of multiple services. Operators can ensure that bandwidth is successfully allocated across multiple applications while maintaining the ability to monitor and account for each traffic flow.

Transactional bandwidth management implies that the QoS requirements on the network systems are constantly changing as new applications start, stop or change their QoS needs. In order to manage this potentially chaotic situation it is critical that operators be able to inspect individual traffic flows to maintain the end-to-end QoS requirements of diverse applications. Per-flow queuing provides operators a mechanism to manage bandwidth and network resources both at the individual flow level and also as an aggregated resource. Bandwidth transactions can therefore be processed according to network policies and according to the possible impacts they will have on existing service flows. Per-flow queuing also provides the fine-grain monitoring needed so that operators can bill for network resources automatically provisioned by the transactional processes.

Per-flow QoS assigns each packet stream its own queue and provides a guaranteed rate for flows with QoS reservations. It is implemented at the headend using high-performance routers that can classify and treat packet flows in real-time. Operators can track individual flows and ensure that users are within their Service Level Agreements (SLAs) for each session, and they can implement the fine-grain metering required to provide back-end billing systems with metering information on each session.

Since cable networks are based on equipment from multiple vendors, a standards-based approach is required to allow applications to dynamically configure network devices to enable automated provisioning. Operators can treat traffic flows using DOCSIS 1.1 and PacketCable standards on the Hybrid Fiber Coax (HFC) access network and MultiProtocol Label Switching (MPLS) across their own core networks and across the core networks of revenue-sharing partners.

WHY TRANSACTION BASED?

Transaction-based services automate bandwidth management. They eliminate operator intervention and allow transactions to request network resources according to policies defined by the operator or its partners.

Transaction-based services can therefore scale efficiently because they can request additional network resources without operator intervention. Operators do not need to pre-provision every option but can establish network policies that define resource requirements for transactions and trigger the automatic billing for incremental resources.

Effective bandwidth management is the best argument for transaction-based services, since operators do not have to reserve bandwidth that is not in use. Operators can take advantage of statistical multiplexing to more effectively deploy billable bandwidth and support high-value services without creating huge reserves of unused capacity.

The possibilities for transaction-based services are virtually unlimited. Any type of service with QoS requirements is a candidate. This includes streaming audio, streaming video, Video on Demand, telephony, interactive gaming, business applications, video conferencing, and enhanced content delivery.

The core technologies used are DOCSIS 1.1 and PacketCable standards for applying QoS on the access network. PacketCable not only supports telephony but will also support any IP traffic flow with QoS requirements. MPLS is used to provide end-to-end QoS across the core networks of one or more providers, and the ReSource ReserVation Protocol (RSVP) is used to reserve network resources.

IMPACTED NETWORK RESOURCES

Operators must be able to automate the provisioning of resources by autoconfiguring network devices. This includes the cable modem, Multimedia Terminal Adaptor (MTA), or set-top box at the subscriber location.

It also includes the autoprovisioning of the Cable Modem Termination System (CMTS) at the headend. This requires an intelligent, high-performance edge router/CMTS that can deliver QoS on the access network using DOCSIS 1.1 and/or PacketCable standards. The edge router must also be able to deliver end-to-end QoS by serving as an MPLS Label Edge Router (LER) that creates a Label Switched Path (LSP) to the destination. It applies the appropriate MPLS labels that carry the QoS configuration requirements to each intermediate device traveled throughout the metro and core networks.

There are other devices impacted as well. For example, in telephony applications the Call Management Server (CMS) and other Point of Presence (PoP) resources must be configured appropriately to ensure end-to-end QoS.

GAINING A PER-FLOW PERSPECTIVE

The DOCSIS standards are based on the concept of traffic flows and encourage operators to think first about flows and then think about how bandwidth can be adjusted to support flow requirements.

Given the DOCSIS emphasis on flows and diverse application QoS requirements it is critical to look at transactional bandwidth management from a flow perspective, with each data session, application, or voice call treated as an individual traffic flow. Individual flow QoS requirements must then be aggregated to ensure that the sum of their requirements can be accomodated within the available global resource set. Viewing only the aggregated requirements from a network resource allocation perspective can be misleading because this view will not identify potential interactions between individual flows. Thus the transaction processing must consider the requirements for classifying, isolating, policing and enforcing the QoS of the individual traffic flows in addition to looking at the aggregate requirements. In order to make the transaction processing manageable it is essential to provide isolation between the individual flows so that it is simple to calculate the impact of a new flow on existing flows. Per-flow queuing provides an efficient mechanism to implement this isolation.

For applications to dynamically configure bandwidth allocations and device configurations, operators must be able to isolate each flow and have the flexibility throughout the network to treat each flow with the proper QoS parameters. When a telephone goes off hook, an interactive gaming session is launched from the desktop, or a click on a web link launches a business application, a series of steps must be implemented to ensure that the proper QoS treatments are applied on a per-flow basis.

Operators implementing per-flow queuing can offer transactional bandwidth management as elements of new services so that when the subscriber selects the applications the appropriate network resources are allocated automatically. On-demand services can benefit from just-in-time bandwidth provisioning so the customer only pays for resources used. Operators can monitor and meter traffic so that application triggers also feed information into billing and charging systems so that they reap premium payments for premium services. They can also successfully develop relationships with third-party providers of content, applications and services based on fine-grain metering and monitoring to ensure that wholesale partners in turn pay operators for bandwidth used by their subscribers.

Transactional bandwidth management and per-flow QoS therefore open up new opportunities for both retail and wholesale revenue streams. With per-flow QoS, operators can successfully deploy transactional bandwidth services that unleash the broadband potential of HFC infrastructure. They can deploy new pay-for-use services and build closer bonds with subscribers based on increased subscriber abilities to self-select service levels based on their own unique bandwidth requirements.

UNDERSTANDING QOS REQUIREMENTS

QoS control is critical for optimizing the productive use of shared bandwidth on the cable access network. Operators need to be able to allow applications to automatically assign bandwidth appropriately according to service requirements and guaranteed commitment levels. The ability to manage QoS involves four key functions:

- Classification of packets to determine the appropriate service level for each traffic flow
- Policing of traffic to prevent flows from getting higher than agreed upon service levels
- Buffering to ensure that queues are created to contain packets during periods of congestion
- Scheduling to enforce packet handling and actually deliver service end-to-end across access, metropolitan and core networks using Internet standards such as MPLS.

The power and flexibility of policy-based QoS control and measurable QoS levels can support incremental revenue streams from transaction-based services.

BRINGING QOS TO THE ACCESS NETWORK

Network operators are now able to offer transaction-based services via shared cable infrastructure while providing guaranteed QoS levels to each service and user. The DOCSIS 1.1 specifications were developed to define enhancements to the Media Access Control (MAC) protocol of DOCSIS 1.0 to enable more sophisticated access methods over HFC access networks by adding the following:

- Packets are classified into service flows based on their content. Thus each application can be mapped to a unique service flow.
- Network access (upstream and downstream) is scheduled per service flow using one of a number of defined scheduling mechanisms including constant bit rate, real-time polling, non real-time polling and best effort.
- Service flows may be configured through management applications or created and deleted dynamically in response to the starting and stopping of applications.
- Fragmentation of large packets is required to allow low latency services to operate on lower-bandwidth upstream channels.

These features provide the basic tools for transactional bandwidth management. They allow applications to request QoS changes dynamically and allow providers to isolate multiple data streams from each cable modem, set-top box or MTA. DOCSIS 1.1-based systems can therefore potentially deliver the ability to allow dynamic application-specific QoS treatment within the HFC access network for each traffic flow.

END-TO-END FLOW CONTROL

The DOCSIS 1.1 specifications provide QoS for the upstream cable access network. In order for applications to see real benefits, QoS must be provided on an end-to-end basis. Thus the QoS-enabled traffic flows from the access network must be mapped to the QoS mechanism(s) used in the regional or backbone networks.

MPLS can provide the QoS mechanism for the regional network. In an MPLS network a number of paths are established between the end points of the network. Each path can be traffic engineered to provide a defined level of QoS.

The successful combination of these two QoS mechanisms requires that the CMTS must also act as an MPLS edge router to map packets from the DOCSIS flows into the appropriate MPLS paths and vice versa. The mechanisms employed within the CMTS/edge router must maintain the QoS during this transition. The addition of dynamic QoS changes as a result of transactional bandwidth management complicate this problem. Fortunately techniques such as per-flow queuing used in combination with a congestion control scheme such as Longest Queue Pushout Pushout (LQP) can simplify this problem.

Hierarchical Per-Flow Queuing and Longest Queue Pushout for End-to-End QoS

Per-flow queuing assigns each packet stream its own queue and provides a guaranteed service rate for flows with QoS reservations. Those traffic flows that are not assigned prioritization are forwarded in a round robin or fair-share manner. To assign flows without reservations to a queue, a per-flow method known as Stochastic Queuing can be used. The parts of the packet header that are the same for all packets of a flow—such as the source and destination IP addresses and source and destination port numbers—are fed to a hash function that is used to map the packet to a queue. This simplifies the management complexity in a dynamic transaction-based environment because it eliminates the need to pre-configure parameters required by other systems such as the bandwidth shares per-class. Consequently, it also avoids the miss-allocation of resources caused by the varying usage patterns inherent to dynamic bandwidth managed systems and the need to continuously update packet classifiers.

If the system can support more queues than there are flows, than most flows either have their own queues or share them with a small number of other flows. This provides isolation between the different flows in contrast to a class based system where all flows of a particular type (e.g. voice) share common resources.

A suitable congestion control scheme must be selected in addition to the per-flow queuing and scheduling to maintain isolation between flows. LQP is the mechanism that best meets the congestion control requirements in a dynamic system resulting from transactional bandwidth operation. It allocates buffers to the individual flows as required until 100% of the buffer pool is used. When no buffers remain and a new packet is received, LQP discards traffic from the flows which are the longest queues. The scheduling system is transmitting from these per-flow queues at a rate which matches the QoS assigned for each flow. By definition the longest queue is that which is exceeding its allocation by the greatest amount. Therefore, traffic is automatically discarded from those applications which are non-compliant with their SLAs. This occurs without the need to configure congestion control parameters and thus works well in a transaction-based bandwidth management scenario.

The combination of per-flow queuing and LQP congestion control enables a practical system to be built which is capable of operating correctly in a dynamic QoS environment without complex management and monitoring.

The ability to track and schedule based on individual flows can be used to create a system in which scheduling decisions are based on the application flows and the subscribers and service providers with which these flows are associated. Operators can implement transactional bandwidth management services within clearly defined boundaries. With hierarchical per-flow queuing and LQP, operators can schedule packet transmissions and selectively discard packets during congestion based on application needs and based on the SLA conformance of applications, subscribers and service providers.

TREATING FLOWS ACROSS ACCESS, BACKBONE, AND CORE NETWORKS

Hierarchical per-flow queuing allows operators to classify and treat traffic across access, core and metropolitan networks. A next-generation edge router with hierarchical per-flow queuing can inspect multiple fields within packets to determine the appropriate routing and QoS requirements. This requires a powerful QoS routing engine that can inspect packets in real-time and route them across multiple networks according to established network policies.

Transaction-based QoS is invoked by an action by the user. For example, the user could click on a web link or pick up the phone. The cable modem, set-top box, or MTA requests the local resources and QoS-based services using DOCSIS 1.1/PacketCable standards.

A Scenario for Transaction-Based Services

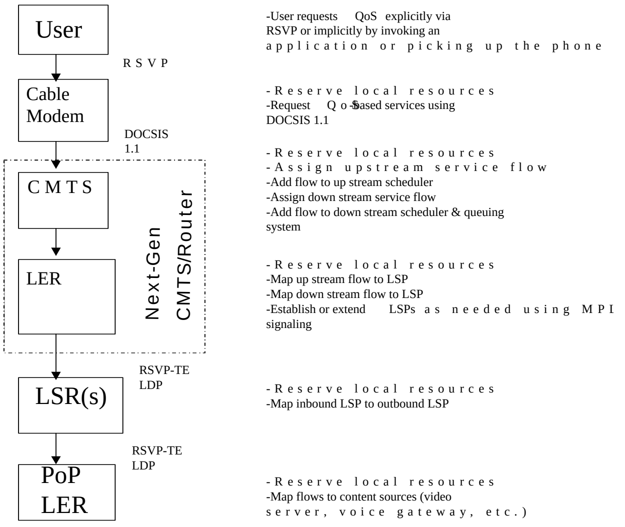
User
RSVP
Cable Modem
DOCSIS 1.1
CMTS
LER
Next-Gen CMTS/Router
RSVP-TE LDP
LSR(s)
RSVP-TE LDP
PoP LER
-User requests QoS explicitly via RSVP or implicitly by invoking an application or picking up the phone
-Reserve local resources
-Request QoS-based services using DOCSIS 1.1
-Reserve local resources
-Assign upstream service flow
-Add flow to up stream scheduler
-Assign down stream service flow
-Add flow to down stream scheduler & queuing system
-Reserve local resources
-Map up stream flow to LSP
-Map down stream flow to LSP
-Establish or extend LSPs as needed using MPI signaling
-Reserve local resources
-Map inbound LSP to outbound LSP
-Reserve local resources
-Map flows to content sources (video server, voice gateway, etc.)

The CMTS portion of an integrated CMTS/edge router at the headend reserves the resources and schedules the traffic flows in the DOCSIS network The LSR portion of the system inspects each flow in real time, maps the packets to an MPLS LSP, applies the appropriate MPLS label and routes the traffic to an MPLS Label Switched Router (LSR) in the core network. The LSRs then switch the packets across the MPLS network to the corresponding LERs at the points of presence for the content sources.

If a flow requires a new LSP (or changes to an existing LSP) then RSVP Traffic Engineering (TE) extensions or Label Distribution Protocol (LDP) messages are used to signal the setup of the MPLS path.

Authorization and Billing

Allowing multiple service providers to deliver transaction-based services over a shared access network requires extensive features for authorization, reconciliation and billing.

Operators need the flexibility to develop policies that determine whether a transaction is allowed at each point in the path and to implement admission control policies that determine which applications are prioritized during times of congestion.

They also need to determine what data to collect, where to collect it, and how to implement reconciliation between service providers. Operators need to maintain detailed accounting information on QoS usage to ensure that application commitments are enforced.

Per-flow queuing allows operators to observe and manage individual IP service flows. Each session can be carefully tracked, and per-flow metering information can be automatically exported to third-party account or mediation applications using the IP Detail Record (IPDR) format or Call Detail Record (CDR) formats.

Granular observability into traffic allows operators to implement enhanced billing applications that reach far beyond traditional flat-rate billing to allow operators to deliver—and accurately bill for—premium dynamic services that automatically increase bandwidth flows in response to application requirements. They can therefore deliver highly granular bandwidth management capabilities with detailed reporting and accounting.

TRANSACTIONAL BANDWIDTH MANAGEMENT WITH PER-FLOW QOS

Operators can now implement transactional bandwidth with per-flow QoS by deploying intelligent edge routers that can classify and treat traffic flows to ensure end-to-end QoS across cable access, metro, and core networks.

A standards-based approach allows operators to build infrastructure that can support high-value, transaction-based services that help them increase revenues, market share, and profits.

###

About the Author

Gerry White serves as Chief Technologist for the Network Infrastructure Group of Motorola Broadband Communications Sector, and he was previously CTO of RiverDelta Networks, Arris Interactive and LanCity. He has presented papers and spoken at dozens of cable industry events worldwide and has published many trade press articles on cable industry issues. White is the co-author of several patents and articles on data communications technology, and he holds a BSc. (Honors) from University College in London.

Author Contact Information

Gerry White
Chief Technologist
Network Infrastructure Group
Motorola Broadband Communications Sector
3 Highwood Drive
Tewksbury, MA 01876
(978) 858-2300
gwhite@motorola.com

PUBLIC KEY INFRASTRUCTURE - USING x.509 CERTIFICATES FOR DEVICE AUTHENICATION HERE A CERT, THERE A CERT, EVERYWHERE A CERT

Doug Jones
YAS Broadband Ventures, LLC.

Abstract

The Public Key Infrastructure (PKI) is a security standard designed to bring "trust" to Internet services. PKI is based on public key cryptography and uses digital certificates to authenticate entities to service providers. It is this authentication that provides trust for services offered over the Internet.

PKI presents a solution for several issues facing cable, and has been included in several CableLabs® projects. PKI is based not only on technology, but also policies. All of which will be discussed briefly in this paper.

INTRODUCTION

Cloning, false identities, theft of service etc., are all issues that Cable has faced before and will continue to face as it advances further into offering Internet-based services. And just like the Internet is a new technology, new security technologies will be needed. The Public Key Infrastructure (PKI) is a standard designed to bring trust to Internet-based services. Not only can PKI protect cable from cloned boxes, PKI will also protect consumers and service providers from fraud as they engage in E-commerce.

The purpose of any security technology is to protect value, whether a revenue stream or a purchasable information asset of some type. Threats to this revenue stream exist when a user of the network perceives the value, expends effort and money, and invents a technique to get around the necessary payments. Some users will go to extreme lengths to steal when they perceive extreme value. The addition of security technology to protect value has an associated cost; the more expended, the more secure the service can be. The proper engineering task is to design a reasonable costing security technology to force any user with the intent to steal or disrupt network services to spend an unreasonable amount of money to circumvent it. Security effectiveness is thus basic economics. Deploying PKI does require additional expense; however, the promise is are both a secure network and secure services that can be the basis for broadband interactive services going forward.

PKI protocols and technologies are developed within the Internet Engineering Task Force (IETF), an international standards body. PKI is a combination of encryption technologies and usage policies that enable the security of digital communications and business transactions on the Internet. PKIs integrate digital certificates, public key cryptography, and certificate authorities into a network security infrastructure.

PKI is included in most CableLabs® projects including PacketCable™, DOCSIS™, and CableHome™. CableLabs is taking an approach such that the PKI infrastructures for the projects will align.

The information in this paper is relevant to both engineers and managers. The material includes:

- introducing the concepts of PKI
- why it was developed
- what it protects against

The beginning discussion is an overview of how PKI works, and the later part of the paper goes into specific technologies.

Why PKI

For security purposes, it is important to both keep communications private and to know with whom you're communicating before exchanging any meaningful information. For electronic business to work, agencies and individuals must be convinced that transactions can be carried out both privately and securely and that both users and documents are authentic. The paper world relies on signatures. The computer world needed an electronic equivalent. PKI is a system for encrypting, decrypting, signing and verifying the authenticity of information that is transmitted over the Internet.

In a PKI, each entity on the network is issued a digital certificate. Hence the statement "Everywhere a cert." The information in certificates can be used for both encrypting and authenticating digital communications and transactions. An entity can be a device, a piece of software, or a user.

WHAT IS PKI

PKI is a group of protocols and techniques that, when put together with a group of policies, allow for secure, authenticated information exchange. The technical parts of PKI are standards and are widely available. The PKI is customized for specific applications based on administrative policies those enterprises set based on needs. It is the set of policies the truly defines how secure the PKI will be.

The center of the PKI is the Certificate Authority (CA). It is the CA that issues the digital certificates to parties within the PKI. The policy for issuing the certificates is what builds the trust in the PKI. Some CAs require rigid identification before issuing a key pair and certificate, other CAs may only require a phone call. While two CAs may use the same technologies, they may have completely different levels of trust based on how the certificates are issued.

PKI works by providing each user with two "keys" — one that is public and one that is private. The keys are represented as binary numbers, i.e., long strings of 1's and 0's. The user must keep the private key secret. The public key should be made available for anyone to use. A document encrypted with a public key can only be decrypted with the corresponding private key. The key generation algorithm is based in very complex mathematics such that the private key should not be able to be derived from analyzing the public key.

The digital certificate, issued by the CA, contains the public key of a member of the PKI. These certificates can be kept in any public place or directory as use of the public key is encouraged for encrypting data.

When a document intended to remain private is transmitted, the sender encrypts it with the public key of the recipient. Once encrypted with the public key, it can only be decrypted with the private key, which should be stored securely by the intended recipient of the message. The integrity of the PKI depends on keeping the private key secure.

In the event a private key is compromised, the certificate for the corresponding public key will be place on what's called a Certificate Revocation List (CRL). The CA will maintain the CRL, and it's the responsibility of the user of the public key to verify that key is not listed on the CRL.

PKI also includes functions that allow recipients to ensure the original message has not been tampered with. This feature uses a digital signature, based on the sender's private key, to "watermark" the message. Using the sender's public key (available in the public certificate) to decrypt the watermark, the recipient can verify that the received message is authentic. Also, by using the sender's private key to create the watermark, the sender of the message is positively identified. Thus, a PKI can ensure that messages are both authentic and that the sender is who they say they are.

There is no single PKI today, rather, various enterprises have created PKIs using the base technologies and incorporating specific policies for their enterprise. As the use of PKI continues to grow, the enterprise PKIs will most likely gradually grow together, thereby increasing interoperability. Policy decisions to be made for the PKI include:

- how certificates are issued,
- who certificates are issued to,
- how certificates are revoked,
- core technologies used,
- how keys are generated, etc.

While base technologies are the same, the policies are what define the PKI. For instance, an informal PKI could be used for the recreational securing of email. A certificate could be issued based only on asking for one. A more controlled PKI would be created for use with e-commerce applications, requiring the users to identify themselves perhaps with birth certificates, drivers license, or credit card information with rigorous, regular ongoing audits of the certificates and their holders.

PKI certificates have a standard format that will be described later in the paper. The important contents of the certificate include material that can be used to uniquely identify the owner of that certificate, including the public key associated with the owner. It is this operation that provides the "trust" associated with the certificate. When a digital certificate is received, it can be used to verify the identification of the sender.

PKI Architecture

The PKI is a layered hierarchy of digital certificates, with the root certificate at the top. The top-most certificate belongs to what is called the root Certificate Authority (CA). The root CA private key is the absolute trusted source of information within the PKI and should be stored using considerable physical security. Generally the root private keys are stored in locked bunkers with many layers of physical security, including voice prints, retinal scans, and security cameras. The physical security of this key is critical because members of the hierarchy use the trust in the CA to authenticate transactions with other members. In economic sense, the physical security of the root CA private key should be configured based on the dollar amount of business that that the CA is designed to protect. This is serious business.

The CA issues key pairs to members of the PKI. The public key is contained in a certificate that is digitally signed by the root CA private key. Being "signed" means the certificate contains a construct that is encrypted with the root CAs private key. Anyone can use the root CAs public key, which is freely available, to verify the authenticity of the certificate, and hence the public key that it contains. As mentioned, revoked certificates will be listed on a CRL maintained by the CA.

For practical purposes of management, a single root CA may be divided into several subordinate root CAs. This is the case for how the Cable CAs are being managed. Figure 1 represents a possible CA hierarchy. The root CA issues certificates for three additional subordinate CAs: 1) software, 2) devices, and 3) service provider systems. The first two subordinate CAs issue certificates to authenticate software loads and the particular devices those software loads go into. The third subordinate CA is used to issue certificates to devices within the service provider's networks, such as various servers.

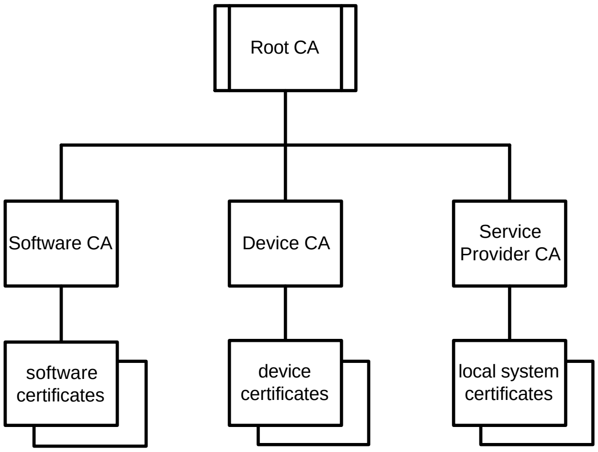

Figure 1

Note, in the specific hierarchy shown in Figure 1 no certificates are issued to users. Users might get certificates based on the e-commerce activities they would be engaged with. That is, the cable operator would issue certificates to ensure their network is trusted, and users would get certificates from their banks or brokerage houses to engage in secure authenticated e-commerce communications. There can be multiple dimensions of CA going on concurrently.

Based on Figure 1 is should be clear it is possible to have a chain of multiple certificates comprising a certificate of the public key owner (the end entity) signed by one CA, and zero or more additional certificates of CAs signed by other CAs. Such chains, also called certification paths, are required because a public key user is only initialized with a limited number of assured CA public keys. These chains are used to segment a CA into manageable entities.

THE X.509 DIGITAL CERTIFICATE

A CA issues digital certificates, and based on the policies of that CA, the certificate will have a certain level of trust associated with it. This is the basis of the operation of a PKI. The digital certificate is simply a computer file that contains information both to identify the holder and some of the policies of the CA. Most PKIs use digital certificates based on the X.509 standard, however, two CAs could have different polices and have incompatible certificates, even though all the certificates are based on X.509.

Important fields on the certificate include the public key of the issuing entity, how long that certificate is valid, and the technologies used by the CA. There can be many more additional fields in the certificate and these are defined in [1]. Because certificates contain only public keys, they can be distributed via untrusted communications systems, and can be stored in unsecured areas.

Another important field in a certificate is the digital signature placed on it by the issuing CA. This signature verifies the entire contents of the certificate. The technology associated with digital signatures is discussed later in the paper, but this is a method to verify the certificate is authentic. A person wanting secure communication with a PKI user need only verify the authenticity of that user’s public key by checking both the digital signature on the certificate and that the certificate has not been revoked by the issuing CA. If both of these steps check, then trust is based on the policies of the underlying CA.

PUBLIC KEY CRYPTOGRAPHY

Cryptography, in general, has been with cable for years and is central to the Conditional Access (CA) systems used for digital video. A “key” is a string of binary bits that is used along with a core cipher to encrypt digital content. In the case of digital video, a core cipher is used to encrypt video at the headend and to decrypt it at the set top box. The core cipher is generally based on the Data Encryption Standard (DES), though there are several ciphers in use. The interesting part of a conditional access system deals with the keys and how they are shared between the headend (to encrypt) and the set top box (to decrypt). There are CA systems based on both symmetric and public key cryptography, as described in the following paragraphs.

One type of cryptographic system uses what's called symmetric or "secret" keys. In this case, the message is encrypted and decrypted using the same key. While this certainly works, there are issues with how to manage and distribute the key that is used. Both parties involved in the transaction must have the same key, and that key must always be kept secret. But a symmetric key has to be distributed to both the encrypter and the decrypter, which means there is an opportunity to steal the key while it is being shared. Anyone stealing that one key can not only decrypt messages, but they can also change the message and forward it on as if it were the real thing. Symmetric key systems are not viable for Internet use because of the issues with key distribution.

A second type of cryptographic system uses asymmetric or "public" keys, and this is what the PKI is based on. Public key cryptography was introduced in 1976 by Whitfield Diffie and Martin Hellman, and it addresses several of the issues associated with symmetric key cryptography. Public key cryptosystems use two keys, one to encrypt and one to decrypt. The key used to encrypt data is called a public key, and as the name suggests, this key is distributed freely, in the form of a digital certificate, available for anyone to use. The key used to decrypt data is called the private key, and as the name implies, this key is to be kept secret by the holder.

Public Key Cryptography for Encryption

Figure 2 shows how public key cryptography for message privacy is used. The sender locates the receiver's public key (freely available in the form of a digital certificate) and uses it to encrypt the message. This message is sent to the receiver who uses their private key to decrypt it.

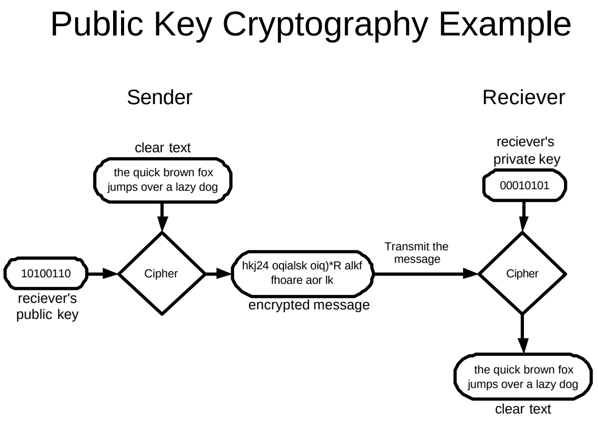
Public Key Cryptography Example
Sender
Reciever
clear text
the quick brown fox jumps over a lazy dog
reciever's private key
00010101
10100110
reciever's public key
Cipher
hkj24 oqialsk oiq)*R alkf fhoare aor lk
encrypted message
Transmit the message
Cipher
the quick brown fox jumps over a lazy dog
clear text

Figure 2

Public Key Cryptography for Authentication

In many cases of secure communication it is critical to not only encrypt the message, but also to verify the originator of a communication (or transaction). This is known as authentication, where the receiver of a message can verify and trust the source of a message. In some cases it's even permissible to send the message unencrypted, as long as there is a mechanism included with that message to verify its authenticity.

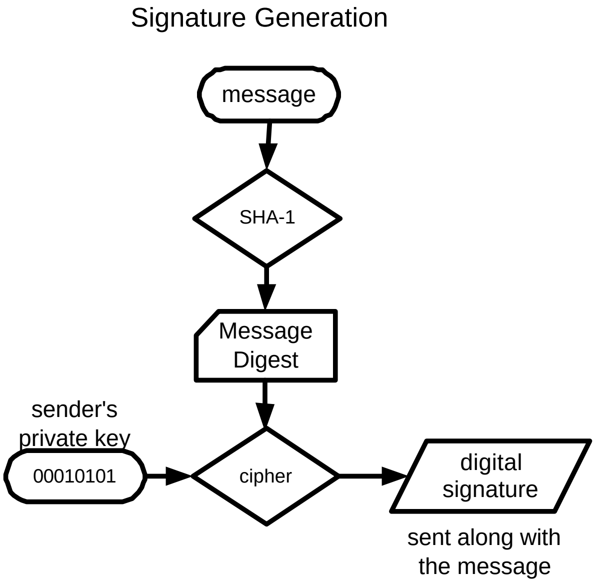

Figure 3

Adding authentication to a message is referred to as "signing" the message. To sign a message, the sender does a computation over that message and the result is called a message digest. The message digest is encrypted with the sender's private key and the result is called a digital signature. The process is shown in Figure 3. The digital signature can be included with the message when it is sent as a form of authentication.

To verify the signature, the recipient first uses the sender's freely available public key to decrypt the digital signature. The result is the message digest over the original message. The recipient then computes a message digest over the received message and compares it to the digest received with the message. If the two digests match, the signature is verified to be genuine and the message is authentic; otherwise, the message and/or the sender is fraudulent. Signature verification is shown in Figure 4.

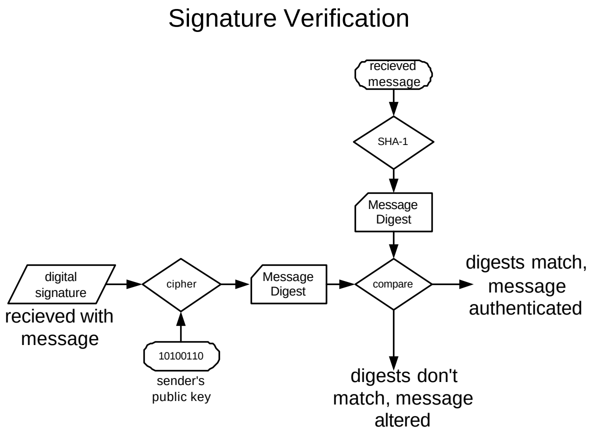

Figure 4

Message Digests

Authentication is accomplished using digital signatures. A digital signature is created by running a message digest algorithm, also known as a hashing algorithm, over the message and then encrypting the resultant message digest with the sender's private key. The input to the message digest algorithm is the original digital file. The output, the message digest, is a binary string, generally of fixed length, that is unique based on the original file. The message digest can be thought of as a mathematical summary of the original file.

A common algorithm used to create message digests is SHA-1 [4], the Secure Hash Algorithm. With SHA-1, the input file size can be up to 2^{64} bits long (a very, very large file, 2 million million megabytes). The output of SHA-1 is always a fixed-length, 160 bit (20 byte), string of digits that is unique for that file.

The SHA-1 is called secure for two reasons. First because it is mathematically infeasible to find two different messages that would produce the same message digest. Secondly, any change to a message will, with very high probability, result in a different message digest. Hence, the message digest provides a method to determine if the original message has been received with integrity. Rather than verifying the entire file, which can be millions of megabytes long, the receiver needs to only recomputed the digest over the message and compare that relatively short byte string with the original digest to determine if the message has changed during transit.

To provide authentication, the message digest is usually encrypted with the sender's private key and the result is a digital signature. Since the sender's public key is freely available, the recipient of the message can decrypt the digital signature and compare the resulting message digest to one computed locally. If the two message digests match, then the message has been received without being altered.

SUMMARY

PKI is an infrastructure that provides encrypted, authenticated communications over unsecured channels. Such a system was designed with the Internet in mind.

The trust in the PKI is based not only in the technology used by that PKI, but also in the policies and procedures instituted. PKIs with similar technologies but different policies can have very different levels of trust.

PKI technologies are fairly straightforward, and are based on public key cryptography, which allows both encryption and authentication. PKIs can be used to secure hardware, software, and e-commerce users.

REFERENCES

1. RFC-2459, Internet X.509 Public Key Infrastructure Certificate and CRL Profile, IETF, January 1999.

2. RFC-2527, Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework, IETF, March 1999.

3. Frequently Asked Questions about Today's Cryptography, RSA Laboratories, April 2000.

4. RFC-3174, US Secure Hash Algorithm 1 (SHA1), IETF, September 2001.

5. How PKI works, William Mathews, Federal Computer Week, June 2000.

Author Contact:
Doug Jones, Chief Architect
YAS Broadband Ventures
(303) 661-3823
doug@yas.com

PUSH FOR MORE MONEY:
REASSESSING THE BANDWIDTH CONSUMPTION MODEL FOR HIGH SPEED DATA

Walter Boyles, SVP, Business Development
Wavexpress, Inc.

Abstract

Streaming can provide an acceptable viewing experience for broadband video depending on the monitor size, subject matter, and bandwidth allocated to the stream. Typically, the bandwidth allocated for streaming video is insufficient to rival broadcast television and certainly will be of less than DVD-quality. In addition, for video and other on-line entertainment content, buffering and interruptions can occur during peak periods, degrading the consumer experience. It is beneficial to both the consumer and the cable operator to adopt a more efficient delivery model for the most commonly requested types of content. The current unicast model allocates bandwidth per subscriber, thus resulting in higher distribution costs and decreasing image quality. Content that is delivered via IP Multicast with secure caching combines broadcast efficiency with a substantially improved consumer experience.

Overview:

Streaming has been the prevalent mechanism that is used for broadband video delivery. Streaming is an appropriate technology for providing content that has low common usage among the subscribers on a particular system. For example, if someone living in New Orleans wants to watch some video highlights of a regatta held in Boston, this content will have a low common usage and is therefore well suited to streaming distribution. Conversely, video coverage of an NCAA Final Four game, distributed in Raleigh, NC, would have a high common usage level. Therefore IP Multicast, a one: many delivery mechanism, would be a more suitable distribution method. Moreover, it would be possible to offer a longer, higher quality video than is typically available via streaming.

The ideal distribution model for this kind of video combines IP Multicast with caching content on a local storage device. This method allows viewers to watch higher quality, longer video segments on demand, without requiring the cable operator to repeatedly carousel broadcasts of the video for subscribers who were not online during the previous delivery of the video segment.

In terms of quality, streaming can provide an acceptable quality signal, sometimes even capable of delivering a reasonable full-screen experience, depending on the monitor size, subject matter, and bandwidth allocated to the stream. We can expect that some continued improvements in codecs and compression algorithms will occur.

However, streaming video is not often considered to be any more than an adequate viewing experience when compared with higher bandwidth, higher resolution video that viewers experience from broadcast and from stored media. Higher bandwidth and/or Quality of Service (QoS) devoted to a single stream improves the image quality of that stream, but this increases costs and reduces the bandwidth available to other streams.

The willingness of subscribers to pay a fee for on-demand content will increase if the provider improves the quality of the viewing experience, eliminating buffering delays and interruption.

Storage or Caching:

In any discussion of caching, the question of storage space is always a primary consideration. In fact, available home storage space is rapidly becoming an underutilized, cost-free "asset" for the bandwidth provider. The cost of storage has continued to fall, and the amount of storage in the home has grown rapidly.

Hard drive storage is now doubling approximately every 9-12 months such that 60 GB hard drives are now found on PCs that cost under $900. In fewer than 5 years, terabyte (1,000 GB) drives will become available on consumer PCs.

In fact, not only is hard drive capacity increasing rapidly but portable media storage is increasing at an even greater pace. It was not so long ago that the only portable storage medium most consumers had access to was a 1.44 MB Floppy Drive.

Now, R/W CD drives are common in new desktop PCs and DVD drives that allow consumers to write DVDs are available. DVD drive manufacturers already have prototypes available of 50 GB capacity R/W DVDs. The drives that write these DVDs that use different wavelength lasers from those used in current DVD drives. This means that a chassis that currently holds 10-20 DVDs could be adapted, with these new DVD technologies, to store 5 – 10 terabytes of content.

Before long, the amount of storage available to the average consumer will dwarf the bandwidth available to place content on these storage devices. The efficient delivery of content to both fixed and portable storage devices will become essential.

Broadband Content Delivery Landscape:

Today third party providers delivering content over the broadband connection to cable modem subscribers are limited by the fact that only unicast delivery mechanisms are available to them. This forces these third party content delivery companies to rely on: (1) Streaming and (2) Downloading of large compressed files (e.g. movies).

Streaming takes advantage of the bandwidth of the broadband connection. Downloading takes advantage of both the speed of the broadband connection and the increasingly large amounts of storage becoming available to consumers.

The relationship between cable operators and the business model(s) of these third party content providers depends on three critical factors. First, consumers using these services use proportionately more of the available bandwidth on a system than other subscribers. Watching streamed content or downloading large files such as a compressed movie is much more bandwidth-intensive than simply viewing Web pages. Second, third party content providers are competitive, in many cases, with the cable operator's core video business. Third, these new services typically generate no revenue for the cable operator but directly increase operating costs.

As High Speed Data (HSD) penetration grows and third party video, music, and games distributors proliferate, they will capture increasing attention from consumers and use increasing amounts of bandwidth. The increasing storage available to consumers will only increase the opportunities available to these third party providers.

Detailed Discussion:

Streaming video at 400-500 kb/s will normally not be perceived as comparable to television quality. Further, the buffering delays of streaming are not only an unwelcome aspect of the experience but also an inhibitor to use and, more importantly, to monetization. Finally, and most importantly, the use of unicast streaming is an inefficient use of bandwidth that does not scale well as demand grows.

Consider a node size of 600 homes with HSD penetration of 20 percent. If, during primetime, half of those HSD subscribers view streaming video, that 50% of subscribers will demand virtually all of the available downstream bandwidth of a 6 MHz channel, using 256 QAM modulation. This is clearly an inefficient usage of bandwidth. This inefficiency is especially apparent when many of these subscribers are either viewing the same streaming event or accessing a few streaming events of predominant interest.

The alternative to unicast streaming is IP Multicast. IP Multicast allows the delivery of content in a broadcast, one-to-many format. The operator now has a rational justification for delivering a higher resolution signal (i.e., using higher bandwidth) to deliver high demand content. However, since IP Multicast is a "push" delivery, the user must be ready to view the video when it is sent or (1) it will have to be sent repeatedly, or (2) it will have to be stored. The problem with sending content multiple times is obvious: the more times a stream is sent the less advantage it offers compared to multiple on-demand streams.

The storing (caching) and time-shifting of a single IP Multicast stream allows the capture of one IP Multicast delivery by all interested subscribers, with time-shifted viewing at the consumer's discretion. However, capture or caching of content requires more robust methods of copy protection to prevent unauthorized secondary distribution.

A comparison of the streaming, IP Multicast, and IP Multicast with caching as broadband delivery options is shown in Table 1.

Streaming	IP Multicast	IP Multicast with Caching
Low/Medium Quality Video	High Quality Video	High Quality Video
Significant Latency Issues	Low Latency Issues	Low Latency Issues
High Bandwidth Application	Medium/High Bandwidth Application	Low Bandwidth Application
Low Content Protection Requirements	Medium Content Protection Requirements	High Content Protection Requirements

Table 1. A Comparison of Distribution Mechanisms for Broadband Content

The quality of the experience, convenience and cost of the delivery of broadband video via IP Multicast to a local cache is superior both to streaming and real-time IP Multicast (essentially one to many streaming).

Further, by delivering such a service, the cable operator is creating a new variable revenue stream for consumers who would otherwise look to third party providers for music, video, and games consumed on the PC.

By using IP Multicast with caching, the cable operator not only provides a superior consumer experience as compared to third party providers, but also realizes the most efficient model for bandwidth usage. Content is delivered once to every customer who is interested in that content.

The only additional requirement for the use of broadband IP Multicast with caching is the need to provide an adequate content protection mechanism that must be more robust than current minimalist solutions used on the PC. This raises two considerations:

Firstly, the content must be locked in a secured space prior to consumption. Secondly, the content is of a higher resolution than streamed content, which means that protecting it after consumption is a higher priority.

The first issue can be readily resolved by dedicating a portion of the subscriber's PC's hard drive for storing the content. This cache functions as a secure content server, conveniently located on the subscriber's home PC.

The cache must be enabled with strong security mechanisms, such as key storage in a protected hardware device to prevent unauthorized unlocking (access) of content, and a distinct real-time clock (not the PC's clock) that the cannot be tampered with, which reliably measures (meters) the withdrawal of content from the cached according to permitted usage models.

In fact, to secure high value cached content a platform that provides for conditional access (CA) as well as "hardened" digital rights management (DRM) is ideal. The platform should utilize strong encryption and authentication, a true (non-deterministic) random number generator, and secured memory spaces to load applications that allow CA and DRM to be run securely on the subscriber's PC.

Summary:

Streaming is the best choice for low-commonality demand content; the optimal delivery method for high value, high demand content is IP Multicast to a secure cache.

IP Multicast with secure caching combines the best attributes of content-on-demand and bandwidth efficiency and allows cable operators to deliver a higher quality service to consumers than third parties competitors employing unicast delivery mechanisms.

A hardware security platform, with adequate content security mechanisms, provides the key to launching an IP Multicast with secure caching in the cable broadband network.

References:

Huitema C., Routing in the Internet, Prentice Hall, 2000, 1995.

Toigo, Jon W, Avoiding a Data Crunch, Scientific American, May 2000.

QUALITY OF SERVICE - IT IS WHAT IT IS, AND IT AIN'T WHAT IT AIN'T

Doug Jones
YAS Broadband Ventures, LLC.

Abstract

What is Quality of Service (QoS) really? It impacts the network and the business. It promises the ability to offer new revenue-generating services. But QoS is not a panacea. To offer QoS services, more than just the DOCSIS network will need to be upgraded to QoS technology.

The specific QoS technology is not so important as knowing why and where QoS should be deployed in the first place. This talk defines what QoS is, and then how an operator can use it in the network.

INTRODUCTION

As with any new business, it's important to both understand the technical issues and have a business strategy for deployment. It's generally not good to get excited by a new technology without fully understanding the issues.

From an engineering perspective, QoS is a set of technologies that can be used to control the delivery of services over various types of connections. Services based on QoS could be categorized as follows:

- Guaranteed throughput
- Guaranteed latency
- Guaranteed packet jitter

QoS is what it is, which is a way to offer new revenue generating services. QoS is not a panacea to all things. For example, deploying a QoS technology is not a substitute for proper network engineering. Deploying QoS will require new expenditures for technology and operations, and the promise is new services revenue will recover these costs.

These are valuable services and subscribers may be willing to pay for them. Deploying QoS technology holds the promise of additional revenues. But only in the presence of a realistic business plan does it make sense to begin upgrading the network for QoS. However, the promise is a QoS-enabled network and services are the basis for broadband interactive services going forward.

Standard Internet QoS protocols and technologies are developed within the Internet Engineering Task Force (IETF). However, many of these protocols are new and multivendor interoperability is not always guaranteed. There are also many proprietary QoS technologies offered by suppliers as a means of holding market share

QoS technology is specific to the underlying type of connection. Cable networks use many types of connections, from DVB-ASI hops, to 100Base-T hops, to a trunk connection to a telephone network. Each type of hop can have its own QoS technology. One size does not fit all. In order to offer end-to-end QoS services, the network may need to deploy more than one QoS technology in order to have QoS on all the different hop types. Having a single network for all types of traffic would make deploying QoS simpler as fewer technologies may be needed.

QoS technology is included in most CableLabs™ projects including DOCSIS™, PacketCable™, and CableHome™. It's well known that DOCIS™ 1.1 provides QoS, and it is one intent of this paper to describe how this fits into the big picture of a QoS-enabled network.

DOCSIS QoS is only available over one hop of a network, the cable connection between a CMTS and a CM. To offer an end-to-end QoS-based service, DOCSIS 1.1 will not be enough. Operators will need to consider adding QoS technology to other elements of their network to provide true end-to-end QoS.

WHAT IS QoS

For the purpose of this paper, QoS is the ability to guarantee network resources on a particular hop through a network for a sample of data. Network resources generally center on bandwidth, latency, and jitter guarantees.

A QoS technology is generally specific for the particular technology used for the hop on the network. In this case, a hop is a point-to-point connection. Examples of these include:

- a 100Base-T connection between a CMTS and a router on the network
- a downstream connection between a QAM modulator and a set top box
- The DOCSIS connection between a CM and a CMTS

In some cases, the same data may be subject to multiple QoS technologies because it is carried over multiple hop technologies in the network. Figure 1 shows an example voice call that uses three QoS technologies. The voice connection between the CM and CMTS uses DOCSIS QoS. That same voice call between a CMTS and a gateway still needs QoS, but this QoS will be based on the technology of that particular hop, probably a 100Base-T. From the gateway to the PSTN, that phone call is generally placed on a constant bit rate 64 kbps connection, either on a GR-303 or a T1. To offer one phone call, each of these three hops needs underlying QoS that guarantees throughput, low latency, and low jitter. Without QoS on all three hops, the voice quality could not be guaranteed. In this example, only having DOCSIS QoS may not be enough to ensure the call would have acceptable voice quality.

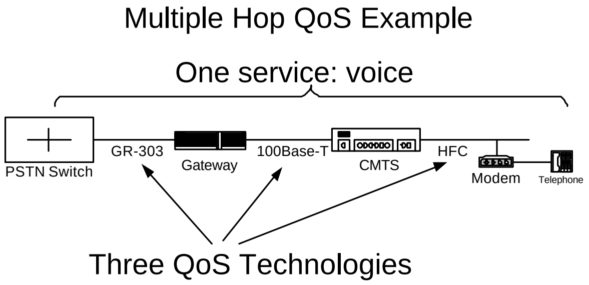

Figure 1

Up to this point, QoS has been described as a particular technology to guarantee a network resource on a network hop. Not only is this "low layer" QoS technology needed, but a higher layer technology is also needed to signal QoS on an end-to-end basis. An example is the Resource reSerVation Protocol (RSVP).

RSVP is used to signal QoS from one end of a connection to the other end of the connection. RSVP is not a "hop" QoS technology, but a method to signal what QoS is needed to allow QoS to be configured on each hop of the connection. Figure 2 gives an example of this.

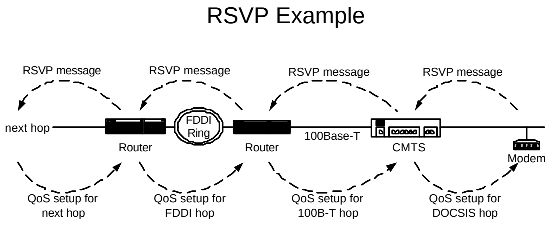

Figure 2

An RSVP message contains a high level description of what QoS is needed for that session. On each hop across the network, that high level description of QoS is translated into the lower layer QoS parameters needed to actually guarantee throughput, latency, and jitter on that particular hop. Since each hop has its own specific QoS, a higher layer messaging is needed to indicate what is needed, and then the individual hops translate that to the actual QoS parameters needed.

QoS On Specific Hops

As introduced, there is generally a specific type of QoS for each hop through the network. This section will look specifically at several hops in a cable network as illustrative examples. In some cases QoS is easier that might be thought.

Digital Video over QAM

Consider a 6 MHz wide 64 QAM channel that has a nominal capacity of 27 Mbps. Digital video streams can be added to this hop one at a time, until the 27 Mbps is used up. QAM hops are generally operated at or below total capacity. If the hop were filled to capacity, some of the digital video would have to be dropped. This, of course, would have a devastating affect on the service.

This illustrates an interesting point about QoS. If the hop is not planned to be overloaded, QoS may not be needed. If the hop is planned to be overloaded, then some bits are going to get through and some others are not. This is when a QoS technology is needed. A 27 Mbps hop cannot carry 28 Mbps, or even 27.01 Mbps. QoS technologies are needed on hops where the hop can be overloaded to guarantee that some bits get preferential treatment. When the hop is intended to be overloaded, a QoS technology can guarantee service for some of the bits but not all of them.

While always having enough (or too much) bandwidth and capacity available on the network would solve many issues, it would probably be economically prohibitive to over-engineer a network like this. Faced with that reality, the engineers developed QoS technologies to ensure that some bits can be guaranteed better service even when the hop is overloaded.

Data over DOCSIS

The DOCSIS network is defined on the hop between a CM and CMTS. DOCSIS 1.0 offers Class of Service (CoS), but not QoS. DOCSIS 1.1 does offer QoS.

DOCSIS 1.0 CoS is essentially a best-effort service that offers bandwidth limits. When the hop is full, the bandwidth limits won’t come into play and the service is just best-effort. The network will do the best it can to be fair to the users, but there are no guarantees. When the hop is full, the network has no choice except to drop packets. For DOCSIS, it’s the CMTS that decides which packets will be transmitted, and which will be dropped. There are many algorithms designed to drop packets randomly, such as Random Early Detection (RED) and Weighted Fair Queuing (WFQ). These could be considered QoS protocols, but they are really a means of choosing packets to drop as opposed deciding which packets will get through.

The DOCSIS hop is generally expected to be overloaded at peak usage hours and packets will get dropped. However, data services such as email and web browsing are tolerant to packet loss, unlike digital video or voice. In the case of email and web, if a packet gets dropped there are “higher layer protocols,” such as TCP, that cause the dropped packet to be retransmitted. So while there may be a delay of a few tens or hundreds of milliseconds, the packet will eventually get through. Internet protocols are designed to be very forgiving and can recover gracefully from packet loss.

DOCSIS 1.0 offers what’s known as Class of Service (CoS), which is different than QoS. DOCSIS 1.0 CoS can be used to implement bandwidth limits, both upstream and downstream, on modems. These are bandwidth limits, not guarantees, hence the difference between CoS and QoS. With CoS, there is no guarantee.

With bandwidth limits, the user is throttled to that amount of bandwidth even if there is additional capacity on the network that could be used. Conversely, users are not guaranteed that amount of bandwidth. Bandwidth limits are particularly useful on the upstream because this lightens the load on the hops that connect from the CMTS to the Internet. For downstream rate limits, there is an argument that if a packet comes all the way to the CMTS, it should be delivered regardless of any rate cap that may be in place. If that packet is dropped at the CMTS, then higher layer protocols will probably cause it to be retransmitted, which means additional load on the backbone. This is arguable unnecessary traffic and an issue where engineering and marketing can respectfully disagree.

DOCSIS 1.1 QoS it is possible to provide QoS guarantees for bandwidth, latency, and jitter. However, if the DOCSIS 1.1 hop is oversold, it cannot guarantee the delivery of all QoS intended to be sent over it. If the DOCSIS 1.1 upstream channel is configured to be 5 Mbps, then only 5 Mbps can be carried. The services guaranteed for delivery over that hop should not total more than 5 Mbps; it's simply not possible. Rather, QoS for guaranteed services can only be offered to a level below the total capacity of the hop. There should be some "wiggle room" on the hop, probably best-effort services that can be degraded in order to meet peak QoS loads.

DOCSIS 1.1 only offers QoS between the CM and CMTS. Just because there is a guarantee to deliver a packet on the DOCSIS hop does not mean there is a guarantee to carry that packet over the rest of the network with QoS.

Voice over DOCSIS

Voice service using DOCSIS can be offered using either of two methods. One is through a gateway to the PSTN, and the other is through a softswitch.

In the first case, there is a DOCSIS connection between the CM and CMTS, and then an IP connection to a voice gateway, and finally a constant bit rate connection to the PSTN.

In this case, there is DOCSIS QoS to the CMTS, and then generally an underutilized 100Base-T hop between the CMTS and voice gateway. If this hop is underutilized (that is, never over subscribed), then there may be no need for using a QoS technology on this hop. In cases like this, just a lot of available bandwidth is a viable way to provide QoS. If there is more available bandwidth (duplex 100 Mbps) than phone call traffic, then all the traffic fits nicely on the 100Base-T hop, and no QoS is needed. Once at the gateway, the voice traffic is converted onto a normal telco hop, like a GR-303 hop, that provides constant bit rate QoS for the voice call. So one phone call, while going from CM to PSTN switch, will traverse at least 3 different types of hops, will get 3 different types of QoS, but all those hops will offer the QoS necessary to maintain good voice quality.

In the second case, there is DOCSIS QoS between the CM and CMTS, but from the CMTS into the soft-switch network, the connection is all IP. The IP connection between the CMTS and the softswitch will probably require a QoS technology, and here there are many choices. If it's a 100Base-T connection, or Gigabit Ethernet (GigE), the QoS technology could be either Differentiated Services (DiffServ), which provides for priority forwarding through router hops, or Multiprotocol Label Switching (MPLS), which provides for a switching through router hops. If the connection from the CMTS to the softswitch is ATM-based (but still IP), the ATM link layer will provide QoS. There are many technologies available, and operators should make choices with strong input from the business case.

QoS Over All Hops

To provide an end-to-end session with QoS, all hops along that path need to offer QoS. It's a simple statement, but it has big implications.

As mentioned before, DOCSIS QoS is only between the CM and CMTS. This is just one hop in the network where there will be multiple hops to get end-to-end. If the particular data stream needs QoS from say a computer on the east coast to a computer on the west coast, then more than just DOCSIS QoS is needed. The connection between the computer and the CM can use DOCSIS QoS, but the connections between the CMTSs will entail several router hops and these will need QoS. If these router hops occur on different networks, for instance the cable operator network and a 3rd party network, then QoS will need to be coordinated through a peering agreement, which is a business issue.

If the connection requires a hop over the public Internet, then it may not be possible to offer end-to-end QoS. The public Internet provides only best effort-service, and that's probably all it will ever offer.

Operators will have to study their networks closely to decide where to deploy QoS technology. In some cases, maybe all the hops in their network will need QoS, but maybe only certain paths through the network will need QoS. These decisions will be key as operators continue to build their networks for both IP and MPEG services. Having a single network for voice, data, and video services will have benefits of using fewer technologies.

QoS Peering Agreements

Once the operator owned network is enabled for QoS, the next decision will be to peer (interconnect) with additional 3rd party networks that support QoS. Such arrangements would make it possible to offer more services to a larger number of subscribers.

Peering requires more of both technology and business. First, business agreements will be needed to monitor the connection, to ensure that subscribers get the QoS for which they are paying, etc. Also the network being peered with may use a different QoS technology, therefore, gateways and translators will be needed to ensure that QoS on one network gets carried with the proper QoS on the other network. Eventually, there should be coast-to-coast and worldwide networks that offer QoS. These QoS networks can run in parallel with the public Internet and will provide an alternative for new services.

QoS Over DOCSIS

In order to provide QoS, the technology on the hop has to provide an underlying mechanism that allows for controlling throughput, latency, and jitter. Not all hop technologies or configurations support QoS.

This final section of the paper is intended to provide technical detail on how DOCSIS provides QoS. Again, this QoS is only on the hop between the CM and CMTS, but it's interesting to know how it works. DOCSIS QoS is also compared to DSL QoS.

Upstream DOCSIS QoS

The cable data return path has the unique property of having many attached CMs that may all want to transmit at the same time. Clearly this needs to be controlled, and in fact the CMTS is the "traffic cop" on the hop that tells each CM when it can transmit, and for how long.

On the return path, the time is divided into periods called minislots. The minislot size is based on a number of parameters, but can be any of 1, 2, 4, 8, 16, 32, 64, 128, or 256 bytes of data. The size of the minislot is fixed on the return path, but will be one of the above.

Several of the minislot sizes just do not make practical sense and are rarely used, e.g., the smallest and largest sizes. For example, if the minislot size were 256 bytes and the data to be transmitted were only a 64 byte TCP ACK, then a lot of bandwidth would be wasted.

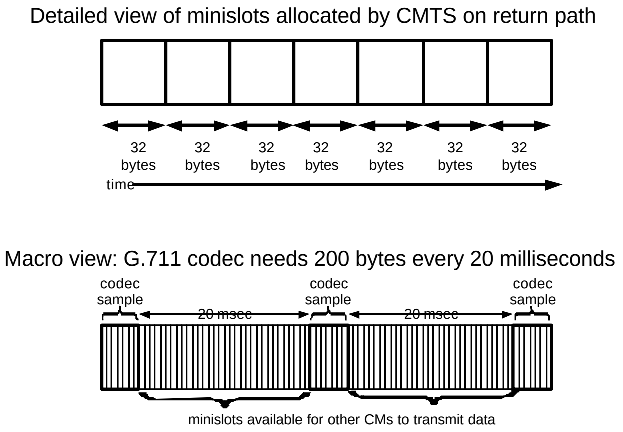
Detailed view of minislots allocated by CMTS on return path
32 bytes
32 bytes
32 bytes
32 bytes
32 bytes
32 bytes
32 bytes
time
Macro view: G.711 codec needs 200 bytes every 20 milliseconds
codec sample
20 msec
codec sample
20 msec
codec sample
minislots available for other CMs to transmit data

Figure 3

The CMTS is generally configured to use a minislot size in the "sweet spot" of from 8 to 64 bytes. Note that the chief competitor to cable data, DSL, is based on ATM which uses a 53-byte cell. DOCSIS provides comparable QoS as compared to ATM; in fact, DOCSIS can provide even finer granularity by using the 8, 16, or 32 byte minislots.

Figure 3 gives an example of how minislots can be allocated on a return path. In the top part of the figure, the conceptual view of 32 byte minislots is shown. In the lower part of the figure, the example shows how groups of minislots can be allocated for a single phone call. The CMTS assigns groups of minislots to the CM for upstream data transmission. The remaining minislots could be assigned to other CMs for upstream data transmission.

Downstream DOCSIS QoS

The cable data forward path is different from the return in that only one device transmits, the CMTS. Therefore, on the forward path the CMTS internally queues and transmits packets as necessary to meet all the QoS guarantees. There are no minislots on the forward path; these are available only on the return path. But because there is only one transmitter on the forward path, there is no need for minislots to allocate bandwidth. All the bandwidth is allocated to the CMTS to transmit packets.

CONCLUSION

QoS provides opportunity, but only if the network is upgraded to support it. Upgrading the network probably means more than just upgrading to DOCSIS 1.1. In addition, new operations and business systems will probably be needed too. But once the network is QoS enabled, there is the opportunity to claim additional revenue.

QoS is not a substitute for good network design. Once services offer real guarantees for throughput, latency, and jitter, how the network is configured and monitored becomes even more important to be sure subscribers are receiving the service being paid for.

REFERENCES

1. Quality of Service: Delivering QoS on the Internet and in Corporate Networks, Paul Ferguson and Geoff Huston, Wiley Computer Books, 1998.

2. Data Over Cable Service Interface Specification Radio Frequency Interface specification version 1.1 (DOCSIS 1.1), SP-RFIv1.1-I07-010829, CableLabs, August 2000.

3. GR-303, Integrated Digital Loop Carrier System Generic Requirements, Telcordia, December 2000.

Author Contact:
Doug Jones, Chief Architect
YAS Broadband Ventures
voice: (303) 661-3823
doug@yas.com

REQUIREMENTS FOR A CABLEHOME RESIDENTIAL GATEWAY[1]

Lior Storfer
Texas Instruments, Cable Broadband Communications

Abstract

Cable operators wish to deliver additional services to the end user in a reliable and controllable way. The IP-based home network is the natural infrastructure for delivery of these services in the home.

The CableHome initiative in CableLabs addresses the fundamental problems associated with having multiple IP devices within the home network: how addresses are allocated, how the home network is secured and how the multiple system operators (MSO) can manage this to ensure service delivery.

This paper presents CableHome from the MSO, end user and vendor's perspective.

THE POTENTIAL OF THE HOME NETWORK FOR DELIVERY OF SERVICES

Potential and Problems:

The home network connects multiple devices to enable the transfer of data to end-points within the home. This creates the opportunity to deliver new services to the end user.

The fundamental problem, however, is that home networks must share some common "look and feel" features to enable delivery of services in a reliable way. These require answers to questions such as:

- How are IP addresses allocated within the home?
- How does the MSO support problems within the home network?
- How do we ensure that the home network is secure?

The discussions of the CableHome[1] forum covered a wide range of potential home network architectures and reviewed the problems arising from them. The home network was divided into logical domains, and devices were categorized into several types.

Following the analysis phase, the CableHome forum focused on defining the requirements for the cable access device. CableHome 1.0 creates a foundation that guarantees home networks will have a consistent "look and feel" to enable services to heterogeneous devices.

Current Market Situation:

DOCSIS 1.0 and DOCSIS 1.1 cable modems are bridging devices. This means that in order for several home devices to share an Internet connection, each device must receive an IP address from the MSO. This creates a provisioning burden on the MSO and the end user alike.

More and more users today buy a "home router" type of device to augment the functionality of their cable modem. These "home router" boxes provide functionality such as a dynamic host configuration protocol (DHCP) server and Address Translation (NAT/NAPT), which allow IP communication between home devices and the global Internet, while requiring only a single IP address from the MSO.

These boxes also provide the end user with some level of additional network security. Most provide NAT functionality; some provide a packet-filtering or stateful firewall. The end user is not always aware of the significant difference in the level of security provided by these two.

Once a user installs a home router, the MSO's ability to support problems in the home network becomes very limited. The behavior of the boxes varies between vendors,

[1] Prepared in collaboration with Jeff Mandin of TI

and the likelihood of a MSO support call (or even a service visit) rises significantly.

In order to lay the foundation for delivery of services into the home network, these basic problems must be resolved. CableHome 1.0 addresses these issues.

CableHome 1.0 OVERVIEW

Devices Affected by CableHome 1.0:

CableHome 1.0 is focused on the Home Access (HA) entity. The HA is composed of two components: the DOCSIS cable modem and a new logical entity called the Portal Services (PS). The PS may reside with the cable modem or in an external box connected via Ethernet/USB. Following similar naming convention to PacketCable[2], CableHome designates these two types as "embedded PS" and "standalone PS." An embedded PS is, thus, a cable modem with home router functionality.

In the case of "standalone PS," the additional device is an external "home router" box that complies with CableHome. "Home router" boxes that will comply with CableHome will be very similar to the ones available in the market today but will share the same "look and feel." This enables support from the MSO, optimizes utilization of the MSO's hybride fiber/coaxial (HFC) plant and provides the infrastructure for the delivery of future advanced services.

Figure 1 illustrates these two options.

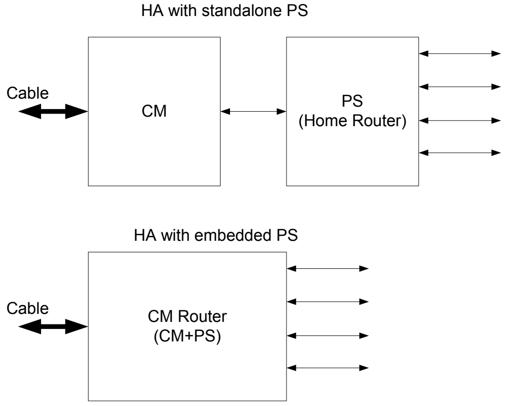

CableHome 1.0 Functionality Definition:

CableHome 1.0 covers the following major areas:

- **Addresses:** How IP addresses are managed within the home network.
- **Management**: With what tools, and to what extent, the MSO supports and controls the HA device.
- **Security:** How the home network is protected with a firewall, as well as the means by which the HA receives/ authenticates with the Head End (HE) and receives keying material.
- **QoS:** CableHome 1.0 defines very limited Quality of Service (QoS) for proper function of PacketCable devices.

Addresses in the Home:

CableHome 1.0 defines a flexible means by which multiple devices in the home share a single connection to the external network and are allocated an external IP address used to communicate with the external network.

The HA is responsible for assigning addresses within the home. It performs that by incorporating a DHCP server, which assigns IP addresses within the home. The IP addresses in the home are local to the home network. When accessing the external network, address translation (NAPT) is performed.

Address translation is an effective mechanism for IP address sharing but can cause complications with peer-to-peer applications like telephony. As a remedy, CableHome allows specific devices to receive addresses in the external address domain, or with NAT address translation, thus providing maximum flexibility.

Securing the Home Network:

CableHome requires a stateful inspection firewall to secure the home network. The firewall must be manageable by the MSO and enable the MSO to upgrade the firewall functionality via download of a new configuration file.

To ensure the integrity of the management messages delivered to the HA, CableHome defines additional security measures. The identity of the HA is authenticated using Kerberos/PKINIT (the same mechanism used in PacketCable).

Management:

The MSO manages the HA with SNMP.v3 for protection against snooping and spoofing.

Figure 2 below presents a typical home network.

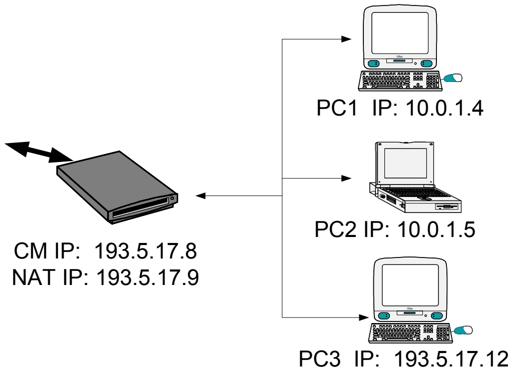

Figure 2 presents a home network with three PCs. It assumes an HA with embedded PS (i.e. a cable modem with firewall/NAT functionality). PC1 and PC2 both have local IP addresses that were received from the DHCP server located inside the CM. PC3 gets its IP address from the DHCP server at the Head End. It is an address outside the scope of the home.

IMPLEMENTING CableHome 1.0

The layer 2 bridging architecture of the DOCSIS 1.1 cable modem can be naturally extended (with straightforward APIs) to support CableHome 1.0. The CableHome component performs firewall and NAT processing on packets after the preliminary bridging decisions by DOCSIS are completed.

A security engine (similar to the one in PacketCable) implements the work of HA authentication and key exchange.

Once authenticated, the simple network management protocol (SNMP) agent supports the CableHome MIBs for the management of the HA. Figure 3 illustrates the software architecture of an HA with embedded PS. The HA contains two logical entities: the DOCSIS cable modem and the PS.

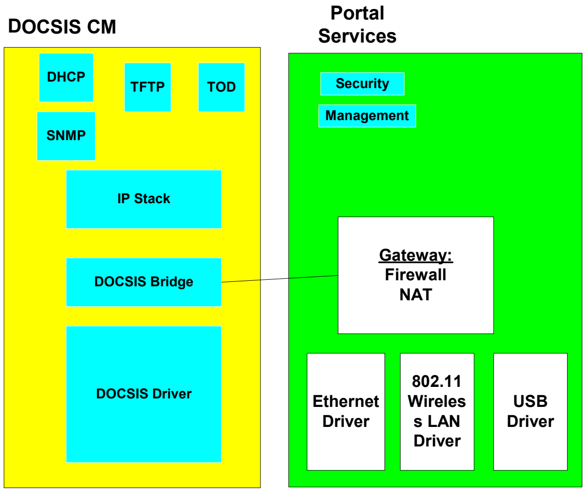

Differentiation:

CableHome 1.0 defines the basic functionality that enables multiple devices in the home network to share a secure connection to the external network. For that purpose the HA device must implement the generic DHCP server, firewall, NAT/NAPT and security features.

Areas for vendor customization include such features as VPN support and parental control, as well as the look-and-feel for the end user. Individual vendors will present different content in local Web pages that the user can browse in order to view and modify the box's configuration. To enable this, boxes will likely include DNS and HTTP servers. As well, CableHome does not require any specific physical networking media within the home. Some vendors might offer Ethernet within the home; others might offer a wireless 802.11 access point. All can be CableHome compliant.

Today a cable modem with a single Ethernet port is a simple bridging device. In the future, even these “low end” cable modems may include CableHome functionality, providing the user with firewall security as well as a simple means to enable more than one PC in the home. The user just needs to buy a hub and connect it to the cable modem. Sharing the Internet connection will already be taken care of by the HA component inside the cable modem box.

End User Benefits:

A box displaying the CableHome 1.0 sticker will provide the end user with the following benefits:

- **Internet connection sharing**: Ability to share a single Internet connection for multiple devices in the home.
- **Enhanced and well-defined security**: A firewall that can be managed by the MSO and upgraded with new policies.
- **Simple installation**: Devices are easy to install in the home.
- **Support** Support is available from the MSO if problems arise.

External CableHome Devices:

Today there are already millions of cable modems out in the field doing simple bridging only. Users today solve the problem of Internet connection sharing by buying a home router device.

It is expected that in the future, retail users will be able to buy home routers that bear a CableHome sticker. Once CableHome routers are readily available, users will naturally prefer to buy one that has the CableHome sticker.

NEXT STEPS

QoS:

CableHome 1.0 addresses QoS in a very limited manner, ensuring PacketCable traffic is handled correctly. It does not provide infrastructure for applications that require well-defined committed throughput and latency. Such an infrastructure will be needed for time critical applications such as voice delivery, video delivery and gaming.

Managing IP Devices:

The CableHome 1.0 spec manages only the HA element. It does not specify how IP devices within the home will be managed to ensure they receive the data that is intended for them. Note that the current CableHome 1.0 does provide some visibility to the IP device performance via the CableHome Test Portal.

CONCLUSION

From an MSO, end user and vendor’s perspective, the IP-based home network is the natural infrastructure for delivery of additional services in a reliable and controllable way.

REFERENCES

[1] http://www.cablelabs.com/cablehome/
[2] http://www.packetcable.com

SECURING DOCSIS CABLE NETWORKS

Annie Phan
Cisco Systems

Abstract

It is a fact that within all shared networks; there is inherent security risk. Allowing multiple entities access to shared network resources carries with it a number of known vulnerabilities. DOCSIS cable IP networks, because they are based on a shared network architecture are prone to violation of data privacy, theft-of-service attacks, and denial-of-service attacks. However, cable IP networks can be easily protected from these attacks and can supply a level of security roughly equivalent to that of an unshared access medium.

The intended audience for this paper is cable operators who are deploying or wish to deploy data over their networks today.

INTRODUCTION

Network attacks are the result of vulnerable spots within a network subjugated to the advantage of a malicious user. Types of typical network attacks are violation of data privacy, Theft-of-Service, and Denial-of-Service.

All edge and aggregation devices, including IP gateways, billing agents, file servers, and provisioning systems need to be protected accordingly. Standard procedure network protection such as locking down devices, OS hardening, firewalls, and intrusion detection are viewed as best practice security for all shared networks.

This paper will assume that best practice network security has been employed at the cable operator site and will focus primarily on the last mile of the broadband cable data network.

TYPES OF NETWORK ATTACK

Violation of Data Privacy

Violation of data privacy happens when a user gains unauthorized access to data that is sensitive in nature. The data can become at risk for being tampered with, destroyed, or distributed. The information found in a user's email, Internet browser, online banking, or any other personal software application is considered confidential. When a violation of data privacy happens, this confidentiality becomes compromised.

Theft-of-Service

A Theft-of-service attack is when a user illegitimately uses a service, which should be paid for. In a cable network, a Theft-of-Service attack is when a cable modem gains network access and services without a legitimate subscription to the cable operator.

Denial-of-Service

When a network or services supplied by the network partially or entirely fail to function due to malicious activity, it is considered a denial of service. A user or several users are usually unable to access the network or services supplied by the network in this type of attack. A subset of Denial-of-Service attacks is distributed Denial-of-Service attacks. Man-in-the-middle attacks are also known as a type of distributed Denial-of-Service attack. In a distributed Denial-of-Service attack, a host or several hosts are unknowingly utilized in a schema that prevents the network to function normally.

VULNERABILITIES RELATED TO DATA PRIVACY

Static IP Addressing

Cable modems are frequently provisioned with fixed IP addresses. Once the customer premise device (PC or modem) receives its IP address dynamically, it will keep that IP address persistently. This is a convenience for cable operators which simplifies provisioning and billing

Static IP addresses imply that the layer three identity of a subscriber stays fixed and an individual subscriber is easily identifiable at all times. A "sitting duck" scenario is created for the subscribers. Unlike dial-up Internet services, where the session is ephemeral, broadband cable is always on.

The threat is that once a malicious entity learns the IP address of a host, that host is susceptible to repeated attempts at compromise.

IP addresses of customers can be easily discovered using a port listening device called a "port scanner" connected to the shared network. These devices are commonly used in networks to monitor network traffic types and levels. More commonly known as packet sniffers, they are typically used for legitimate purposes and are readily available to any interested party.

Unauthorized Access Using Netbios

If enabled, NetBios can allow unauthorized access to a host whose IP address was learned from a packet sniffer. The NetBios protocol, a file and print sharing protocol common to Microsoft Windows Operating Systems, is a common mechanism by which user data is compromised. In a trusted network, where all users on the LAN are actively managed,

NetBios can be very useful. However, since cable IP networks are a shared medium with unmanaged hosts, NetBios if not properly enabled, is a relatively easy vulnerability to exploit.

The most frequently identified security holes for subscribers on cable IP networks are NetBios vulnerabilities. Broadband subscribers may unknowingly have ports enabled for file sharing with untrusted users on the same local cable segment. The opportunity exists for any files that contain sensitive information on the customer's hardrive to be exposed to untrusted users. Because many PC vendors install operating systems with file and print sharing enabled by default, many users are unknowing exposed to violation of their data privacy due to vulnerabilities associated with "sharing" services.

Sensitive information such as tax returns, address books, and password files that may exist on a user's PC may be viewed or even tampered. Since the vulnerability exists at the user's PC, a cable operator should encourage their subscribers to disable features such as NetBios when it is not in use for legitimate purposes.

SECURING DATA PRIVACY

Baseline Privacy Interface 1.0

Securing the subscriber's identity and sensitive data should begin at the cable modem. The Data-Over-Cable Service Interfaces Specification (DOCSIS) version 1.0 features a Baseline Privacy Interface (BPI) Specification to improve the security of data privacy over cable networks. The purpose of BPI is to provide a fundamental level of protection for all devices that attach to the cable modem network. When BPI is enabled, it helps prevent a user from passively listening on the cable network to learn sensitive information that was passed in the clear from neighboring modems.

The primary goal of BPI is to secure data privacy. BPI encrypts all traffic flows between the cable modem and cable modem termination system (CMTS).[1] This is helpful for protecting sensitive data that is exchanged across the RF portion of the cable network.

BPI uses 56-bit DES encryption to encrypt traffic in both the upstream and downstream directions. This affords cable modems roughly the same security found at the last mile in point-to-point circuit based networks. The Baseline Privacy Key Management (BPKM) protocol outlines the encryption algorithm used to exchange public-key information for the traffic exchanges between cable modems and CMTS.

Baseline Privacy is not enabled by default on cable modems, but it can be easily enabled through the DOCSIS configuration file. Once BPI is enabled, minimal utilization of the cable modem's CPU is used because BPI encryption and decryption occurs within cable modem chipsets. In other words, there is negligible change in performance when BPI is enabled on all devices of the cable modem network.

VULNERABILITIES RELATED TO THEFT-OF-SERVICE ATTACKS

Theft-of-Service attacks are typically committed through device cloning or device spoofing. These methods can be used in Denial-of-Service attacks as well. This section of the paper will discuss MAC address, IP address, and software spoofing in relation to Theft-of-Service attacks.

Weak User Authentication In BPI 1.0

In a DOCSIS 1.0 network, BPI mainly protects against unauthorized access to personal data using strong data encryption. BPI 1.0 does not have any type of authentication distribution protocol between the cable modem and CMTS; hence it does not provide strong protection from theft of service. MAC address spoofing can bypass BPI in this case, despite the encryption between the CMTS and cable modem, since there lacks authentication between them. In a "best practice" security model, strong protection is constructed upon not only strong encryption, but also strong authentication. Authenticating users in a cable environment becomes critical to protection against cloned devices.

IP Address Cloning and Spoofing

In networks where BPI is not enabled, a user could easily have the ability to learn the IP addresses that are in use. Using a packet sniffer, a user can learn the entire network structure, if all the relevant IP address information is left unencrypted. The threat is that a user can clone an IP address or spoof an unused IP address in the fixed range of addresses provisioned by the cable operator and then commit a theft of service.

Even with BPI enabled, a poor IP addressing scheme can still fall at risk to theft of service. Suppose the cable operator took a simple approach to IP address provisioning, where all cable modems reside in a single class of addresses. The entire cable modem network would become a flat structure with a single autonomous class of IP addresses, with no classification between different groups of users. As a feature, all DOCSIS cable modems have their IP addresses provisioned via DHCP. Essentially, DOCSIS cable modems are self provisioning, similar to "plug and play".

The threat is that a user who is not a current broadband subscriber will have the ability to place their modem anywhere on the network, automatically be provisioned with an IP address via DHCP, and thus gain network access illegitimately. It is for this reason it is advisable to configure multiple scopes within the DHCP provisioning server.

Multiple scopes allow the flexibility to differentiate between new (untrusted) cable modems and existing subscriber (trusted) cable modems. A default scope can be created such that it will have limited or no public Internet access for devices that have not yet been subscribed. Additional configuration in the network on routing or switching gateways will be needed to accommodate the different scopes and limit access control.

Once a subscriber contacts their cable operator to properly register their cable modem, the MAC address of the device is manually put into the provisioning system. From there the subscriber's cable modem will be tracked for billing and placed in a new IP address scope. The new subscriber will have to reset their cable modem so that it will request a new address from the new scope during initialization. Then the cable modem will be granted complete network access with the new IP address.

Again, this method does not include strong authentication for cable modem users. If a user had the ability to learn a legitimate subscriber IP address, bypassed the DHCP server and spoofed the address; there would be nothing to prevent network access to this user. This is despite if the network was configured within a hierarchal IP address scheme. There still remains the threat of theft of service or even denial of service once a malicious user clones or spoofs an IP address.

Software Spoofing

In a DOCSIS cable network, a cable modem will download via TFTP a binary configuration file after it has completed DHCP negotiation[2]. This is one of the final steps before the modem can become completely registered. Pertinent provisioning information about the cable modem including IP gateway, upstream bandwidth, downstream bandwidth, quality of service, software image and privacy can be specified by the DOCSIS configuration file.

The threat is that a user could access the DOCSIS configuration files by either compromising the existing TFTP server on the network or copying an existing file traveling across the wire, and tamper with the existing parameters of the file. Users can even compile their own DOCSIS configuration files for their own modems.

A user could commit a theft of service by using the DOCSIS configuration file to provision more upstream or downstream bandwidth for their cable modem. A user could also illegitimately trigger a new software image update through the DOCSIS configuration file and gain access to later feature sets in later software image updates if available on the TFTP server. The threat is that a user could upgrade the software image of their cable modem to an image meant for a higher-class user (i.e. business class with commercial services) without having to pay for that upgrade.

PROTECTION AGAINST THEFT-OF-SERVICE ATTACKS

Baseline Privacy Plus Interface

Authentication of cable modems on a DOCSIS network was addressed in DOCSIS version 1.1[3]. In DOCSIS 1.1 the new and improved BPI was created. It is called the

Baseline Privacy Plus Interface (BPI+). BPI+ uses RSA encrypted digital certificates to authenticate cable modems on the network.[4] It is substantially more difficult for a user to fake an embedded digital certificate than it is to spoof or clone MAC addresses. In the BPI+ specification, each cable modem uses a unique X.509 digital certificate that is issued by the cable modem manufacturer. This gives the CMTS the ability to strongly authenticate cable modems before they can successfully register online.

BPI+ inherited all the strengths that came with BPI. BPI+ protects data privacy using the same 56-bit DES encryption. Like BPI, all BPI+ encryption and decryption occurs at the hardware level, which means the performance impact is negligible when enabled across all devices in the cable network. BPI+ is not enabled by default, but it can be easily enabled through the DOCSIS 1.1 configuration file in the same fashion BPI was. It is advisable to always have BPI+ enabled to protect against theft of service. The authentication found in BPI+ helps protect against violation of data privacy, MAC address spoofing/cloning, and IP address spoofing/cloning.

IP Address Provisioning

To fend against users that are sophisticated enough to spoof an IP address by bypassing the cable operator's DHCP server, some CMTS manufacturers have decided to leverage the DHCP process for an additional level of cable modem monitoring. For example, Cisco Systems' CMTS line has the ability to verify MAC address and IP address pairs by reading the option 82 relay-agent option of DHCP.

Cisco Systems' CMTS will compare the IP address that the MAC address is trying to register with against what the DHCP server has already recorded within its scope(s). If for any reason, a cable modem and IP address pair exists that does not match against the pair found on the DHCP server, the CMTS can deny access to the mismatched MAC address. This provides authentication between the cable modems and CMTS at a layer 3 level. This prevents a user from registering online with a modem that has a cloned or spoofed IP address.

Although all DOCSIS 1.0 and DOCSIS 1.1 certified cable modems are required to receive their IP address via DHCP[5], not all DOCSIS certified CMTS's leverage the option 82 relay-agent information in DHCP. Cable operators should make themselves familiar with the IP security features that are available on their CMTS by checking with the manufacturer.

DOCSIS 1.1 Named Service Class

Quality-of-Service (QoS) profiles, that include upstream and downstream bandwidth specifications, can be created on DOCSIS 1.1 CMTS's. QoS profiles combined with named service classes eliminate the need to distribute QoS information in the DOCSIS configs files via TFTP. Cable operators with DOCSIS 1.1 CMTS's can provision cable modems with a named service class and then associate those modems with the QoS parameters configured on the CMTS. This moves sensitive information off of the TFTP server.

This provides a level of security mainly because it alleviates the need to protect a targeted TFTP server, also the parameters that a malicious user would typically tamper with, such as allocated bandwidth, no longer have to remain on the TFTP server. The added benefit of having the QoS parameters specified by the CMTS is that all interaction between the CMTS and the cable modem can be additionally secured with BPI enabled.

Common Open Policy Server (COPS)

In the DOCSIS 1.1 specification, cable modems authenticate with a Common Open Policy Server (COPS). This affords the cable operator greater authentication that extends beyond BPI+. This protocol uses a client/server model that maintains message integrity and reliability.[6] COPS is a stateful protocol in that it allows the server to push configuration information to the client, and then allows the server to remove that information from the client when it is no longer applicable. This helps prevent modems from unauthorized access on the network, thus curtailing theft of service.

TFTP Server/ DHCP Server Hardening

The DHCP server and TFTP server portion of the cable operator network is critical. To prevent compromise of these servers, best practice network security measures need to be employed. Firewalls should be implemented to keep all unwanted or unexpected traffic out. For security that extends beyond packet filtering, Intrusion Detection Systems (IDS) can be used. An IDS can listen to all traffic and monitor for any inconsistencies in the traffic. An IDS has the ability to flag interesting traffic and in some cases police against attack. These measures are not only used to ward against Theft-of-Service attacks but also against Denial-of-Service attacks as well.

VULNERABILITIES RELATED TO DENIAL-OF-SERVICE ATTACKS

Denial-of-Service (DoS) attacks are generally initiated by the exploitation of vulnerabilities within the cable operator site and not within the last mile of the cable network. This is mostly because critical devices whose failure would cripple the network and services supplied by the network reside at the cable operator headend. Distributed DoS attacks do usually utilize the last mile of the cable network in order to gain anonymity during attack as well as to launch attacks from many more hosts in order to overwhelm the network or network services.

IP Redirects

The IP address information of default gateways, time servers, TFTP servers, and name servers supplied by the DHCP policy for DOCSIS cable modems. . If a user were able to compromise the DHCP server they could redirect all IP traffic to a bogus IP address or their own server. This could create a large-scale Denial-of-Service attack, since all connected cable modems would be unable to register properly or be denied network service.

Man-in-the-Middle (MitM)

A Man in the Middle is an exploit that targets the victims TCP based applications like Telnet, rlogin, ftp, mail application, web browser, etc. Without BPI enabled across the cable network, any connected user could become prone to having their system compromised. A malicious user could sniff packets from the network, modify them, and then insert them back into the network. An attacker can grab unencrypted confidential information from a victim's network based TCP application. The authenticity and integrity of the data would then be compromised, at which point the victim could be denied service.

Distributed Denial-of-Service

If a malicious user were able to compromise multiple systems, they could attack the network on a wide scale. Assuming control of multiple customer premise devices throughout the cable modem network, a distributed DoS attack is possible. Broadcast storms could be initiated throughout the network, overloading CPU utilization of critical gateway routers and switches. The result would be large numbers of subscribers unable to reach the network.

PROTECTION AGAINST DENIAL-OF-SERVICE ATTACKS

Perimeter Security Measures

As stated earlier when protecting critical devices from Theft of Service, the cable operator site should employ perimeter security measures such as firewalls and intrusion detection systems to decrease the incidence of attack. Best practice security measures can defend against IP redirects, broadcast storms, and other types of DoS attacks. The main idea is to have strong access control to all critical gateway systems.

Firewalls are stateful packet filters, which can police access to and from a host. Firewalls have the ability to filter by source, destination, and message type. Access control lists (ACLs) have the ability to limit telnet access, web browsing, and FTP usage. Firewalls can be employed to control access to critical gateway systems.

As mentioned earlier Intrusion detection systems (IDS) listen to all traffic and monitor for any suspicious activity. The IDS has the ability to identify specific signatures that indicate a host may intend to initiate a broadcast storm or transmit a packet with tampered headers. When suspicious activity is identified, the IDS will send an alarm. If configured properly, it can also react against particular types of attack. The IDS can block offending traffic with dynamic Access Control Lists (ACLs) or reset the offending connection.

Securing Data Privacy

Enabling BPI or BPI+ will lessen the likelihood of a user taking control of its neighboring modems, since the 56-bit DES encryption discourages compromise. With the additional authentication found with BPI+, users will be discouraged from unauthorized access on the cable network as well.

CONCLUSION

Broadband cable subscribers need to trust the cable operator. However, due to security risks, cable operators must always be wary of subscriber behavior. Security is an ongoing process as opposed to a one-time event. The threats on any given network are not necessarily representative of all the vulnerabilities within the network, as much as it is the human factor, which exists there. In the future, stronger encryption for user authentication and data privacy should be expected.

[1] Baseline Privacy Interface Specification SP-BPI-I03-010829

[2] Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification (Released v1.0 SP-RFIC01-011119

[3] Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification (Interim v1.1 SP-RFIv1.1-I07-010829)

[4] Baseline Privacy Interface Plus Specification SP-BPI+-I07-010829

[5] Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification (Released v1.0 SP-RFIC01-011119) (Interim v1.1 SP-RFIv1.1-I07-010829)

[6] Common Open Policy Server Specification RFC2748

SECURITY FOR NEXT GENERATION ACCESS NETWORKS

Stephen Thomas
Wave7 Optics

Abstract

Broadband access networks—including those built with advanced HFC, wireless, and fiber technologies—have unique network security concerns. This paper analyzes the security threats present in such networks, and it develops a general security threat model for broadband access networks. Significant threats include masquerade and eavesdropping.

The paper then examines cryptographic techniques appropriate for countering these threats, focusing on authentication and encryption. It considers the strengths and costs of various alternatives. The final section discusses practical implementation issues, particularly those that arise because of the potential size of access networks and because of the data rates at which they operate. These characteristics demand security measures that are very scalable and capable of very high speed operation.

SECURITY THREATS

Like any other network infrastructure, broadband access networks face an array of threats to their correct operation. Malicious parties may attempt to steal service; they may try to deny service to legitimate users, and they may attempt to compromise the confidentiality of network users. Compounding these traditional security threats, access networks face a unique and challenging problem: customer premise equipment.

Unlike enterprise and backbone networks, significant components of the access network are not under the physical control of the network operator. In fact, customers, who may have easy physical access to these network elements, may well be the most likely attackers. This situation exacerbates two broad security threats, masquerade and eavesdropping.

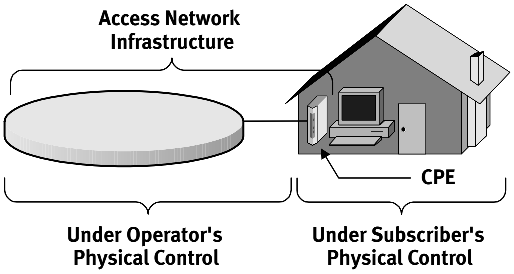
Access Network Infrastructure
CPE
Under Operator's Physical Control
Under Subscriber's Physical Control

Figure 1 CPE is not controlled by operator.

Masquerade

In most communication networks, especially those provided as a commercial service, the operator must know the identity of its users. Collecting revenue, for example, usually depends on knowing who to bill. When attackers masquerade, they disguise this very information.

> ***masquerade**: The pretense by an entity to be a different entity in order to gain unauthorized access.*[1]

Customer premise equipment can make the threat of masquerade particularly acute. The equipment is sitting inside (or just outside) the potential attacker's home, waiting to be reverse-engineered, modified, or even relocated. Some may find the temptation irresistible. "For instance, right now you can type in 'TiVo hack' on Google, and you'll get a thousand sites of hard drives that are compromised at the user's premises.'[2]

Successful masquerades can have many consequences. Attackers may pirate service by pretending to be a legitimate user, or they may intercept key exchange messages so as to decipher encrypted communications. Masquerade is also one step in a more wide-scale attack such as device cloning.

Eavesdropping

Customers of communication networks often presume that the information they exchange using those networks remains confidential. Attackers that eavesdrop compromise that confidentiality.

> ***eavesdropping**: The unauthorized interception of information-bearing emanations.*[1]

Access customers are often particularly sensitive about their privacy. Customers have objected strenuously when the network operator has apparently violated their privacy,[3] even to the point of involving senior members of the US Congress.[4] The expected repercussions would be significantly more severe if private information was exposed to an unauthorized third party.

On most access networks, customer premise equipment heightens the threat of eavesdropping. Access networks frequently rely on shared media, where the physical media for information transfer—coaxial cable, fiber optics, or wireless spectrum—is shared by many users. This characteristic means that information transmitted to one user is inherently available for reception by other users. In fact, with early cable modem deployments, it was quite easy to eavesdrop on your neighbor unintentionally.[5]

Countermeasures

Fortunately, the science of cryptography has developed countermeasures to combat these security threats. Authentication protects against masquerade, and encryption can prevent eavesdropping.

Authentication

Authentication protects access networks against masquerade attacks by giving users or devices a way to prove their identity.

> ***authentication**: A security measure designed to protect a communications system against acceptance of a fraudulent transmission or simulation by establishing the validity of a transmission, message, or originator.*[1]

In access networks, authentication allows *untrusted* network elements, such as customer premise equipment, to prove their identities to *trusted* equipment physically controlled by the network operator. Table **1** lists the trusted and untrusted elements for common access network technologies.

Table 1 Parties to Authentication

	Trusted	Untrusted
HFC	CMTS	Cable Modem
802.11	Access Point	Wireless Client
PON	OLT	ONU

To prove their identities, untrusted elements demonstrate their knowledge of a secret value. The most straightforward approach relies on a shared secret such as a password. The trusted element knows the passwords of elements that it must authenticate; the authentication process requires that untrusted elements prove they know the same shared secret.

One obvious way to authenticate using shared secrets is for the untrusted element to simply send the password to the trusted element, as in figure 2. A disadvantage to this approach is that the shared secret is transmitted, in the clear, across the access network. If an adversary can intercept those communications, the adversary can learn the shared secret and impersonate the untrusted element.

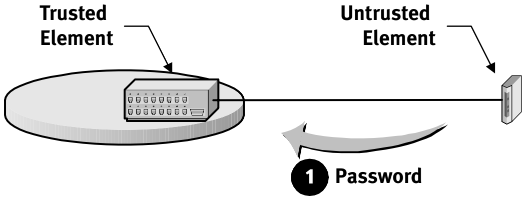

Figure 2 Simple password authentication.

Fortunately, there are simple cryptographic techniques that can significantly improve the security of shared secret authentication. A common approach relies on special mathematical functions known as *hash functions* or *message digests*. A message digest is a one-way function: it is easy to compute but extremely difficult to reverse. For example, figure 3 shows the result of computing the Secure Hash Algorithm[6] on a block of input data. The mathematical properties of the algorithm are such that, given only the output from figure 3, an adversary cannot deduce any information about the original input.

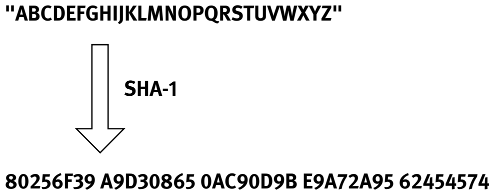
"ABCDEFGHIJKLMNOPQRSTUVWXYZ"
SHA-1
80256F39 A9D30865 0AC90D9B E9A72A95 62454574

Figure 3 A message digest algorithm.

Figure 4 illustrates how message digests improve the security of authentication exchanges. Both the trusted and untrusted elements share a secret, but the value of that secret never crosses the access network. Instead, the trusted element sends the untrusted element a *challenge*. The untrusted element combines that challenge with the shared secret and computes the message digest of the combination. It only sends the result of this digest computation across the network. Even if an adversary is able to intercept this message, the adversary will not be able to derive the shared secret.

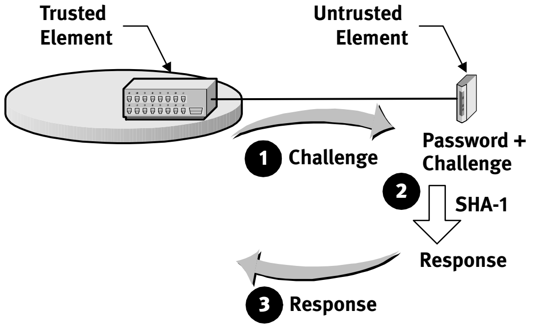

Figure 4 Digest authentication.

Although message digests solve many of the technical security problems with shared secrets, they can't eliminate the operational burden that shared secrets impose. A Hybrid Fiber-Coax (HFC) network serving 100,000 subscribers, for example, would require that the operator maintain 100,000 different shared secrets, provisioning the appropriate values in each Cable Modem Termination System (CMTS), managing additions and deletions, protecting their values from theft, and so on.

Some technologies minimize the operational burden by adopting a common shared secret for all network elements.[7] Although this approach may make operational issues more manageable, it provides substantially less security. With a common secret, the trusted network element cannot verify the individual identity of an untrusted element; *all* the untrusted elements know the secret value. For access networks, this vulnerability may allow one subscriber to impersonate another, an unacceptable deficiency in many environments.

For a more secure approach to the operational problems of shared secrets, access networks can use asymmetric encryption. Asymmetric encryption relies on a pair of related keys. One key, known as the *public key*, can be made public, even to potential attackers. The other *private key* is known by only one party. Asymmetric encryption algorithms use the mathematical properties of these keys so that information enciphered with a public key can only be deciphered with the corresponding private key, and vice versa.

Figure 5 shows a typical authentication sequence based on asymmetric encryption. The very first step is the critical one. In that step the untrusted network element sends its public key to the trusted element. This communication can safely take place even if adversaries are able to intercept it; there is no danger in having an attacker learn a public key.

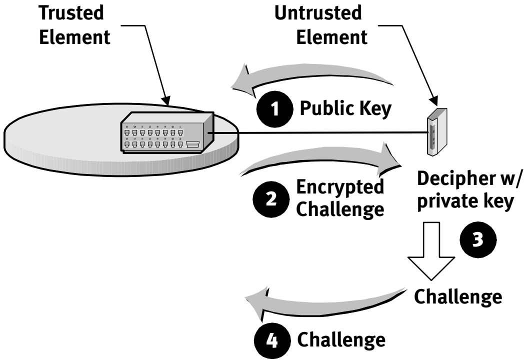

Figure 5 Authentication with public keys.

After the first step, the process is very similar to figure 4. The trusted element generates a random challenge, enciphers that challenge using the public key, and sends the result to the untrusted element. The untrusted element must then decipher the communications to recover the original challenge. Since the untrusted element is the only party that knows the private key, it is the only element that can recover the original challenge. By doing so, and by returning that challenge to the trusted element, it proves possession of that private key.

Even though authentication using asymmetric encryption resembles authentication with message digests, the extra step at the beginning is very significant. With asymmetric encryption the trusted element does not have to know the secret value, it simply learns the public key directly from the element it is authenticating.

Asymmetric encryption does introduce one additional factor in the authentication process. Figure 5 shows how a trusted element can verify that an untrusted element possesses a specific private key (the one corresponding to the exchanged public key). But how does the trusted element ensure that the keys are the right ones? Certificate authorities provide that assurance.

A certificate authority (CA) vouches for the authenticity of a public key. It creates a digital certificate that includes the public key, a distinguishing feature of the party possessing the public key, and the certificate authority's digital signature. As long as the trusted element in the access network believes the CA, it can verify the untrusted element's public key.

Encryption

Although cryptography is often critical to authentication, its more glamorous function is encryption.

> ***encrypt**: To convert plain text into unintelligible forms by means of a cryptosystem.* [1]

In access networks, encryption protects against eavesdropping. An attacker may be able to intercept a network's communications, especially if the network relies on a shared media. But if that communications is encrypted, the information content will remain unintelligible to the attacker.

The effectiveness of encryption as a security measure depends on several factors, including the particular cipher algorithm, its implementation, the size of cryptographic keys, and their generation and management. The most common measure of encryption strength is key size, as measured in bits. Key size allows objective comparisons between different encryption approaches.

In July of 1998, the Electronic Frontier Foundation demonstrated a special-purpose (but relatively inexpensive) hardware system that could exhaustively search for the cryptographic key used to encipher given ciphertext. To demonstrate its effectiveness, the EFF was able to discover a 56-bit key in 56 hours, although a full search of all possible keys would have taken 9 days.[8] Table 2 shows how long that same (1998) technology would take to exhaustively search the key space for various key sizes.

Table 2 Strength of Key Sizes

Key Size	Network	Search Time
24 bits	APON	180 microseconds
40 bits	802.11	12 seconds
56 bits	DOCSIS	9 days
128 bits	*future*	116 × 10^{18} years

Despite the popular focus on key sizes, most deployed encryption systems are actually broken because of implementation flaws. Those flaws have included poor random number generation,[9] poorly designed algorithms,[10] weakness in how keys are scheduled for use,[11] and weaknesses in how keys are derived.[12] Unfortunately, there is no simple way to assess the strength of these factors in an encryption system prior to actual deployment.

The utility of encryption in an access network link is sometimes questioned because the access network is typically only part of an end-to-end communication path. Users that desire real security will have to adopt their own measures to protect the end-to-end path, and these security measures could potentially make access network encryption redundant. This argument is false. However, because it does not consider all the information that can be obtained through eavesdropping.

Figure 6 shows a typical use of end-to-end encryption: the user is accessing a Web site with the Secure Sockets Layer (SSL) protocol.[13] If the user purchases an item from the Web site, SSL encrypts the contents of the transaction to protect the user's credit card number. The SSL protocol, however, does not obscure the identity of the Web site. The eavesdropper may not be able to intercept the credit card number, but he can certainly discover that the user made some purchase from a specific Web server. Most users would consider that to be a violation of their privacy, and only link encryption within the access network itself can prevent it.

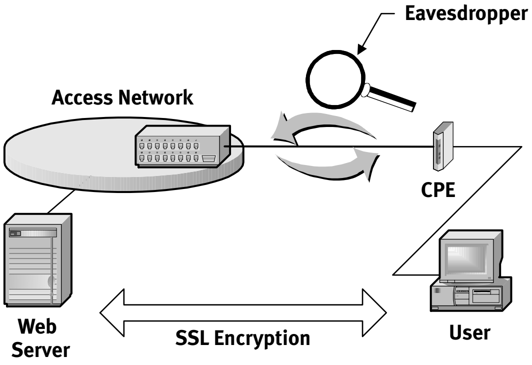

Figure 6 End-to-End Encryption.

Relying exclusively on end-to-end encryption also places a considerable burden on the access network's users. Not only must they employ appropriate security measures themselves, they have to recognize when those measures are needed. The need for security may be obvious in applications such as Web browsing, but it may be quite obscure for services such as the transport of PBX traffic to the operator's central office.

Implementation Issues

The science of cryptography has provided network designers the principle tools required to make access networks secure: authentication and encryption. Effective engineering of access network security requires applying this cryptographic theory to practical systems. Most implementations rely on combinations of public key based authentication, shared secret authentication, and traditional cryptographic ciphers.

Authentication with Public Keys

Because of the operational burdens that shared secrets impose, authentication based on asymmetric encryption is generally considered the most effective, practical technique for authenticating access devices. Creative attempts to use shared secrets in an access network, including common secrets (e.g. IEEE 802.11) and automatically learned passwords (e.g. ITU G.983[14]) have proven to be ineffective.[11][12]

When implementing public key authentication, network designers must chose an appropriate public key infrastructure (PKI). The important components of a PKI include the particular asymmetric encryption algorithm, the format of public key certificates, and the certificate hierarchy.

The asymmetric encryption algorithm of choice for most applications is the RSA cipher invented by Rivest, Shamir, and Adleman.[15] RSA is the algorithm of choice for Web security[13] and for the Cable Labs Data-Over-Cable Service Interface Specifications (DOCSIS).[16]

Opposition to RSA based on intellectual property concerns were addressed in September of 2000 when RSA Security, Inc. released the technology to the public domain (a few days before their US patents would have expired).[17] The technical merits of other algorithms, such as those based on elliptic curve cryptography (ECC),[18] appear to be limited to special environments not typical to access network. (ECC calculations require substantially fewer computational resources than RSA for equivalent levels of security, but, because authentication in access networks is generally very infrequent, the occasional requirement for lengthy calculations is not normally a problem.)

Public key certificates are almost exclusively formatted according to the ITU's X.509 standard.[19] Most of the recent work on enhancing and extending X.509 has taken place within the Internet Engineering Task Force.[20]

Shared Secret Authentication

One domain in which public key authentication is not effective is authenticating humans. Most people find it very difficult to remember lengthy random numbers, and very few are able to perform the complex calculations of asymmetric encryption. Despite their weakness, passwords have proven to be the most effective authentication tool for human users.

Access networks that need to authenticate individual users as well as customer premise equipment, may use both public key and shared secret authentication. Figure 7, for example, shows a CPE that includes an embedded 802.11 wireless access point.

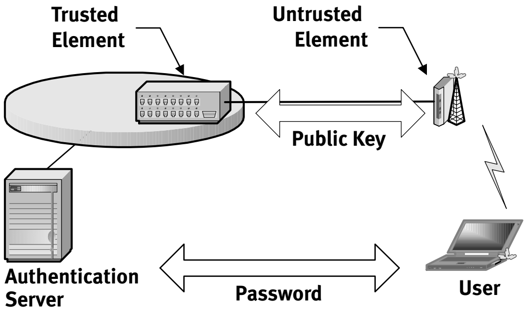

Figure 7 Combined authentication.

The network may use asymmetric encryption to authenticate the CPE and a password to authenticate individual wireless clients that attempt to connect to the access point. Note that the initial version of the IEEE 802 protocol for password-based authentication, IEEE 802.1.x,[21] has been shown to have significant vulnerabilities.[22]

Ciphers

One of the more interesting challenges facing developers of next generation access networks is the selection of an appropriate encryption algorithm. Clearly, as table 2 demonstrates, the algorithm must be capable of support 128-bit keys, but there are many competent ciphers with that capability. The challenge lies in finding a cipher that can be implemented economically and still operate at the high speeds demanded by next generation networks.

Much of the private sector research into encryption ciphers has focused on algorithms that can be efficiently implemented in software. The data rates of next generation access networks, however, will likely require hardware implementations. Considering the large number of subscribers that a single trusted element may support, efficient hardware implementation is critical for economically viable products.

This problem has led some network technologies to create their own confidentiality algorithms, without the benefit of cryptographic professionals. The results have been predictably poor.[12]

Fortunately, recent cryptographic research has begun to consider hardware implementation efficiency. Table 3, partially adapted from material presented as part of the National Institute for Standards and Technology's competition for the Advanced Encryption Standard,[23] lists representative hardware implementations for a few important ciphers.

The final algorithm in the table, W7, is a byte-wise stream cipher developed specifically for hardware implementation in high speed access networks. It has been published as an Internet Draft.[24]

Table 3 Hardware Implementations

Algorithm	Throughput	Area (gates)
3DES	407 Mbps	148,147
Rijndael	1.95 Gbps	612,834
W7	2 Gbps	20,375

Conclusions

As in all commercial data networks, security is a critical component of next generation access networks. Security for access networks, however, is particularly challenging because of customer premise equipment. The fact that elements of the network may be physically located on the premises of potential attackers requires that network designers take great care in the security of their designs. In particular, access networks must protect against masquerade and eavesdropping attacks. Fortunately, modern cryptography provides the tools necessary to defeat these attacks. Strong authentication, typically based on asymmetric encryption, assures operators of the identity of communicating network elements, and advanced encryption, particularly using ciphers optimized for high speed operation in hardware, prevents eavesdropping.

References

[1] *Telecom Glossary 2000*. American National Standard T1.523-2001.

[2] Michael Lee, Vice President and General Manager of Interactive Services for Rogers Communications, as quoted by James Careless in "Rogers, Bell Canada Square Off," *CED Magazine*, September 2001.

[3] John Rendleman. "Comcast Backs Down Over Privacy Concerns," *Information Week*, 18 February 2002.

[4] Edward J. Markey, Ranking Democrat on the House Subcommittee on Telecommunications and the Internet, in a 13 February 2002 letter to Brian

Roberts, President of Comcast Corporation.

[5] Steve Bass. “Opinion: What Life with a Cable Modem is Really Like.” *PC World*, 30 April 1999.

[6] *Secure Hash Standard*. Federal Information Processing Standards Publication 180-1, 17 April 1995.

[7] *IEEE Standard for* Information *Technology—Telecommunications and Information Exchange between Systems—Local and Metropolitan Area Networks—Specific Requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. ANSI/IEEE Standard 802.11, 1999 Edition.

[8] Michael Fitzgerald. “EFF quickly cracks Data Encryption Standard.” *PC Week*, 17 July 1998.

[9] Ian Goldberg and David Wagner. “Randomness and the Netscape Browser.” Dr. Dobb’s Journal, January 1996.

[10] Alex Biryukov, Adi Shamir, and David Wagner. “Real Time Cryptanalysis of A5/1 on a PC.” Fast Software Encryption Workshop, April 2000.

[11] Nikita Borisov, Ian Goldberg, and David Wagner. “Intercepting Mobile Communications: The Insecurity of 802.11.” *ACM Conference on Mobile Computing and Networking*, July 2001.

[12] Stephen Thomas and David Wagner. “Insecurity in ATM-based Passive Optical networks.” *IEEE International Conference on Communications*, April 2002.

[13] Stephen Thomas. *SSL and TLS Essentials: Securing the Web*. John Wiley & Sons, 2000.

[14] *Digital transmission systems —Digital sections and digital line system — Optical line systems for local and access networks — Broadband optical access systems based on Passive Optical Networks (PON)*. ITU specification G.983.1, October 1998.

[15] R. Rivest, A. Shamir, and L. Adleman. “A Method for Obtaining Digital Signatures and Public-Key Cryptosystems.” *Communications of the ACM*, February 1978.

[16] *Data-Over-Cable Service Interface Specifications — Baseline Privacy Plus Interface Specification*. Cable Labs Interim Specification SP-BPI+-I07-010829, 29 August 2001.

[17] George V. Hulme. “RSA's Patent Release will Fuel Security Competition.” *Information Week*, 11 September 2000.

[18] A. Menezes. *Elliptic Curve Public Key Cryptosystems*, Kluwer Academic Publishers, 1993.

[19] *Information Technology — Open Systems Interconnection — The Directory: Authentication Framework*. The International Telecommunications Union Recommendation X.509, August 1997.

[20] http://www.ietf.org/html.charters/pkix-charter.html

[21] *IEEE Standard for Local and Metropolitan Area Networks—Port-Based Network Access Control*. IEEE Standard 802.1X-2001.

[22] Arunesh Mishra and William A. Arbaugh. “An Initial Security Analysis of the IEEE 802.1X Standard.” University of Maryland Technical Report UMLACS-TR-2002-10, 6 February 2002.

[23] Tetsuya Ichikawa, Tomomi Kasuya, and Mitsuru Matsui. “Hardware Evaluation of the AES Finalists.” *Third*

AES Candidate Conference. April 2000.

[24] S. Thomas, D. Anthony, T. Berson, and G. Gong. “The W7 Stream Cipher Algorithm.” [draft-thomas-w7cipher-00.txt] October 2001.

Stephen Thomas is the Chief Architect for Wave7 Optics. He can be reached at stephen.thomas@wave7optics.com.

SMALL SYSTEM ARCHITECTURE FOR DIGITAL CABLE

Christopher Poli, PE and Michael Hicks
Motorola Broadband Communications Sector

Abstract

This paper focuses on the basic architectural elements required for digital cable service with an expandable architecture that makes it economically advantageous to deploy in a cable headend with as few as 1000 basic subscribers. This assumes 20% digital penetration or approximately 200 digital subscribers.

INTRODUCTION

A truly successful architecture must, of course, be economically sound, but further must allow for anticipated growth – both in number of subscribers and in variety of service offerings. The architecture described here accomplishes both. Set-top terminal cost becomes a major consideration to support a low-cost headend architecture with anticipation for growth; the architecture outlined in this paper supports a refurbished to like-new DCT1000 set-top terminal in the $100 range. Of course, any of Motorola's digital set-tops can be deployed in this system. These set-tops allow the operators to offer their customers many choices in their digital experience such as 256QAM reception, expandable application memory, integrated Hi Definition decoding, Dolby 5.1 Surround Sound support, integrated Tivo and REPLAY like digital recording support, and a variety of I/O options such as IEEE1394, SPDIF, TOSSlink optical, and USB are some of the exciting features that can be deployed on any system. System services can be grown to support VOD, internet access and interactive services (games, weather, etc.).

WHY DIGITAL?

According to Kagan Media (June 5, 2001), digital cable subscribers will increase more than 7 fold from about 10 million at the end of 2000 to over 70 million in 2010, DBS subscribers will almost double from under 15 million at the end of 2000 to about 27 million while analog-only subscribers will decline from almost 60 million to a mere 3 million subs. With these projections, it is clear that digital will play a part in almost every subscriber's suite of services delivered to the home. The lion's share of video service revenue will go to whoever is the provider of digital - whether it is the cable system or a DBS service. With growing availability of local programming from DBS services in many areas, subscriber defection could increase substantially.

Revenue is a major reason for deploying digital services. Digital generates revenue because it gives subscribers three things they are willing to pay for – higher quality, better choice and greater convenience. Digital television provides far superior picture quality compared to analog. With digital cable, up to 10 or 12 video programs occupy the same 6 MHz channel bandwidth that currently provides one analog service. The virtual expansion of an existing plant provides a more compelling package to consumers. Digital music, integrated guide, digital off air, and enhanced PPV offerings are just several of the most basic services available. It has been shown that digital systems enjoy increased Pay-Per-View® (PPV) revenue and premium-tier penetration.

With a low cost digital deployment, a cable system can protect and grow subscriber base and reduce subscriber defection. Small cable systems can offer services every bit as compelling as the competition and eliminate the reasons for subscriber defection to DBS. The architecture takes full advantage of *proven components* in a *proven system* that is servicing over 6 million set-top terminals today. The headend architecture is *fully expandable* taking advantage of every bit of capital equipment investment. Growth of the service offerings to exceed the competition is readily available.

Lastly, by protecting subscriber base, small cable systems can protect the value / net worth of the system. Cable system values have been demonstrated by recent consolidations. At $3000 to $4500 per subscriber, a small system of 1000 basic subscribers could be valued between $3 and $5 million dollars. The bottom line is that by ignoring digital cable deployments in smaller markets, there is a substantial risk of erosion of the system's market value.

BASIC ARCHITECTURE

Basic Elements of a Digital Cable Television System

The basic elementary services in a digital cable television system include analog video services, digital video services, pay-per-view (call ahead and impulse) and an interactive electronic program guide (IPG). Enhanced digital services could also include Video on Demand (VOD), Internet Based Services, Interactive Games, etc. Services, including data services, can also be encoded locally or delivered via satellite. Since this paper is focused on the ability deploy to smaller cable systems, it will primarily focus on digital cable services distributed via satellite, as they are the most economical available today. However, the architecture detailed below is readily expandable to include virtually any service including web access, VOD, home shopping, and interactive games.

As illustrated in figure 1, functions needed to control and distribute the basic services include access control, configuration, collection, service reception, decryption, encryption and the associated RF elements required to transmit the selected services on a cable plant. These functions are required to securely receive a service (or group of services) via satellite and distribute and control reception on a cable plant to selected subscribers. The return paths illustrated are strictly for the collection of impulse pay per view (IPPV) events.

Signal Flow – Video Services are digitally encoded and uplinked for satellite distribution. Each encoder output constitutes a transport stream and will consume 6 MHz of bandwidth on the cable plant. The encoder converts an analog signal to an MPEG-2 format; it also encrypts the digital signal for secure delivery to designated cable systems. Data for an IPG can also be carried on one of the satellite transport streams; the data can be extracted after the downconversion of the RF signal in the headend.

On the downlink, the video service is downconverted and decrypted. The services on the transport are then re-encrypted to control access on a subscriber-by-subscriber basis. They are then QAM modulated and upconverted for distribution on the cable plant. The program service providers (who run the encoders) control the decryption function in the headend while the cable system conditional access system controls the encryption (can be one in the same as the program service provider).

Access to specific services is controlled by the access control function. Subscriber equipment configuration and upgrade services are also required to support the system. Configuration sets the set-top filters and controls to receive specific signals, functionality, and code objects. Upgrade services allow the quick, inexpensive deployment of code for new services as well as enhancements to existing functionality. Code Upgrades can be targeted to a single set-top terminal, a group of terminals for test, or to an entire population.

All the specific information required to provision both services and set-top terminals are contained in the access control and data base functions. The business office and billing system connect with the data base interface to allow the cable operator to assign specific configurations and authorizations to a particular subscriber and bill for the services received. The access control function provides the authorizations (keys) to decrypt the subset of services a particular subscriber is entitled to receive.

A given digital PPV service can be locally configured for impulse with later collection via RF return, impulse with later collection via telco-return or call-ahead-PPV. The functions at the bottom of the figure illustrate two-way communication (call ahead not illustrated as it is a pre-authorized event similar to a subscription service). Telco takes advantage of an existing infrastructure and is therefore more economical to deploy. Either method provides reliable timely collection of IPPV revenue. The Return Demodulator and Modem functions provide the physical transport; the Data Collection function is common although there are slight differences in the collection process due entirely to the difference in the transport. For enhanced interactive services such as VOD, RF return is required. Enhanced services will be briefly discussed as part of the expandable architecture at the end of the paper. The small plant architecture either uses telco-return or is configured for CA-PPV.

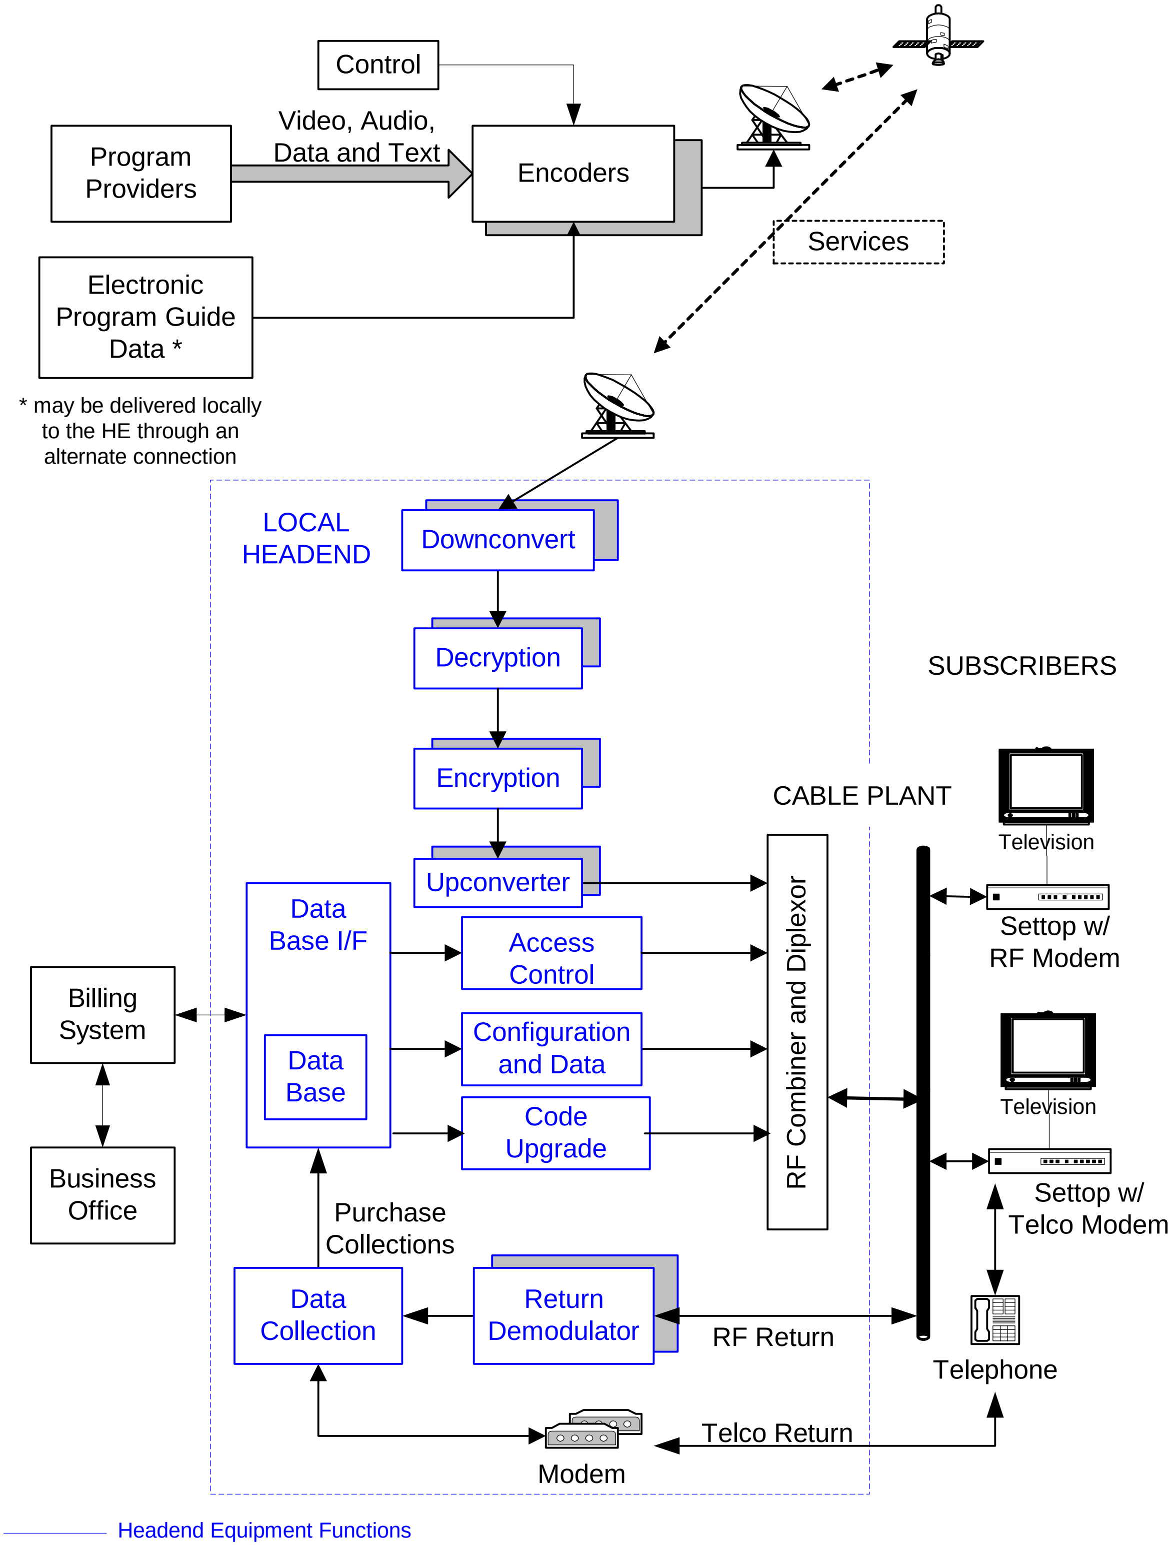

Figure 1 Functional Block Diagram for Basic Digital Services

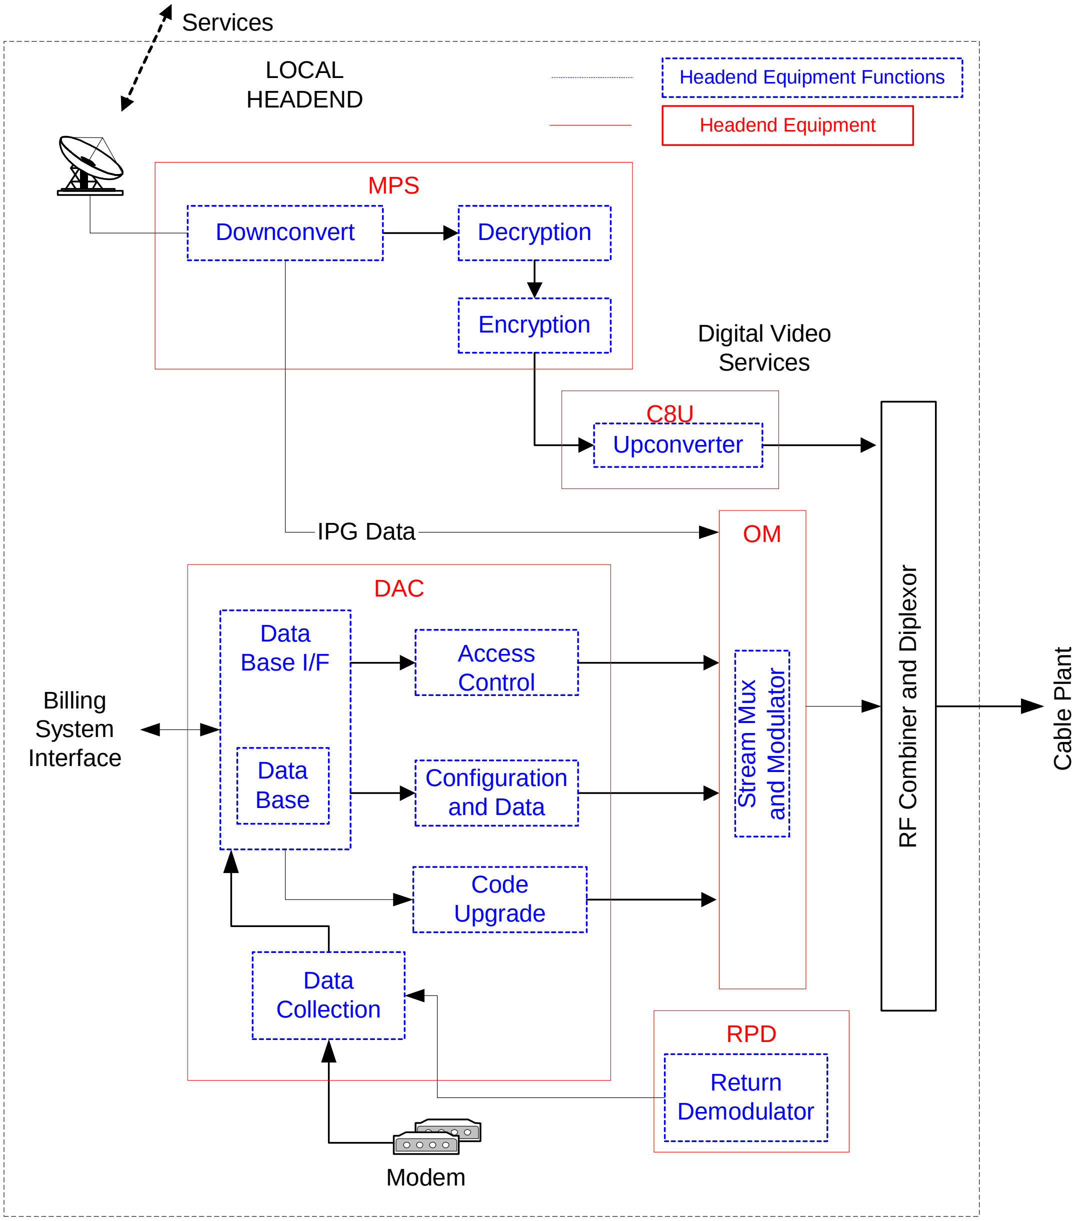

Figure 2 Basic Equipment Functions in a Local Headend

Functions Mapping to Equipment

Figure 2 illustrates how the functions fit into the equipment required for the most basic, locally controlled headend. There would normally be multiple Modular Processing Systems (MPS) and C8U's to provide more programming. Return Path Demodulators (RPD) provide for RF return (RPD's are completely eliminated with telco-return). Depending on the number of headends being controlled from the Digital Addressable Controller (DAC), the size of the plant and total number of set-tops, there might also be multiple Out of band Modulators (OM). All functions, definition of services to be delivered, equipment configuration, code upgrade, and collection responsibility are coordinated and controlled by the local operator through the DAC. Only signal flows are shown. The Operations, Administration, Maintenance & Provisioning (OAM&P) 10baseT network is not shown above.

Figure 3 illustrates the basic equipment in an end-to-end local digital cable headend. Again, typical installations include multiple MPS's, C8U's, and RPD's. Depending on the cable plant topology and number of headends being serviced by the DAC, multiple OM's may also be controlled off of the single controller. In the figure below, the headend OAM&P network connections are shown (10baseT Ethernet).

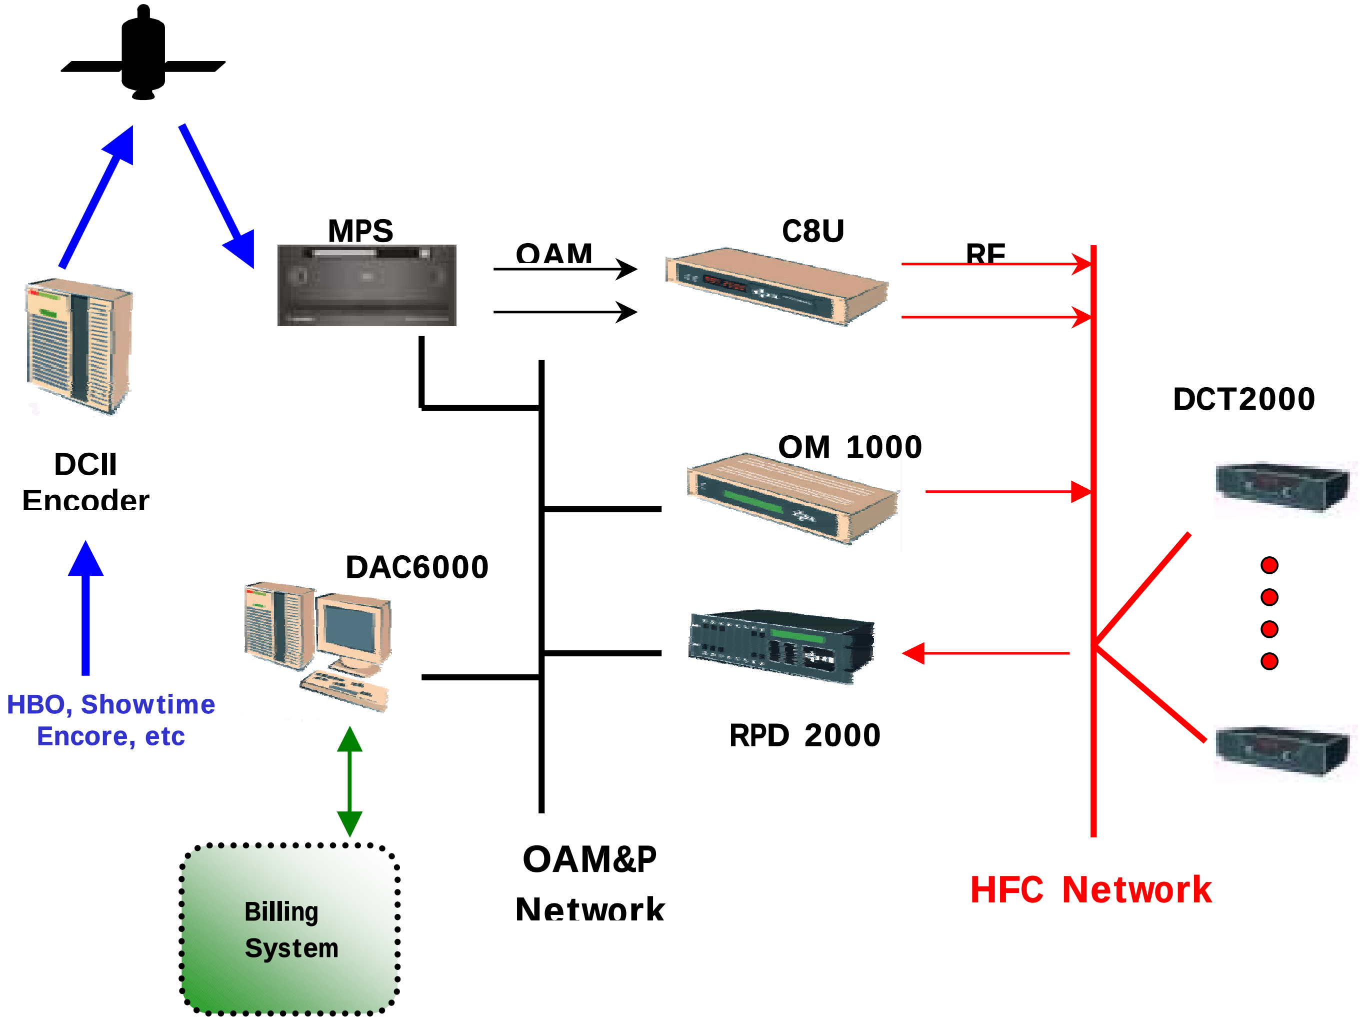

Figure 3 Local Control End to End Digital Cable Equipment

Taking the Cost Out

It is possible to centralize several of these functions in a national center to gain economies of scale. It also provides a central controlling point so that extensive in depth operational knowledge is not required of an operator to deploy digital cable in an existing analog system. Centralizing the most complex functions (the access control, set-top and headend equipment configuration and upgrade functions) provides the most benefit to the local operator.

By removing these functions, not only is a significant up front capital cost burden removed, but also significant ongoing operational costs are shifted to a central organization. Those central operational costs are shared among more than one thousand headends and millions of set-top terminals. The actual cost burden isn't based on an individual headend, but rather upon a per-set-top terminal cost so that an operator is only burdened with the operational cost associated with a given revenue-producing unit (in this case, the set-top terminal). Headend cost would otherwise make this architecture economically unacceptable for a smaller system.

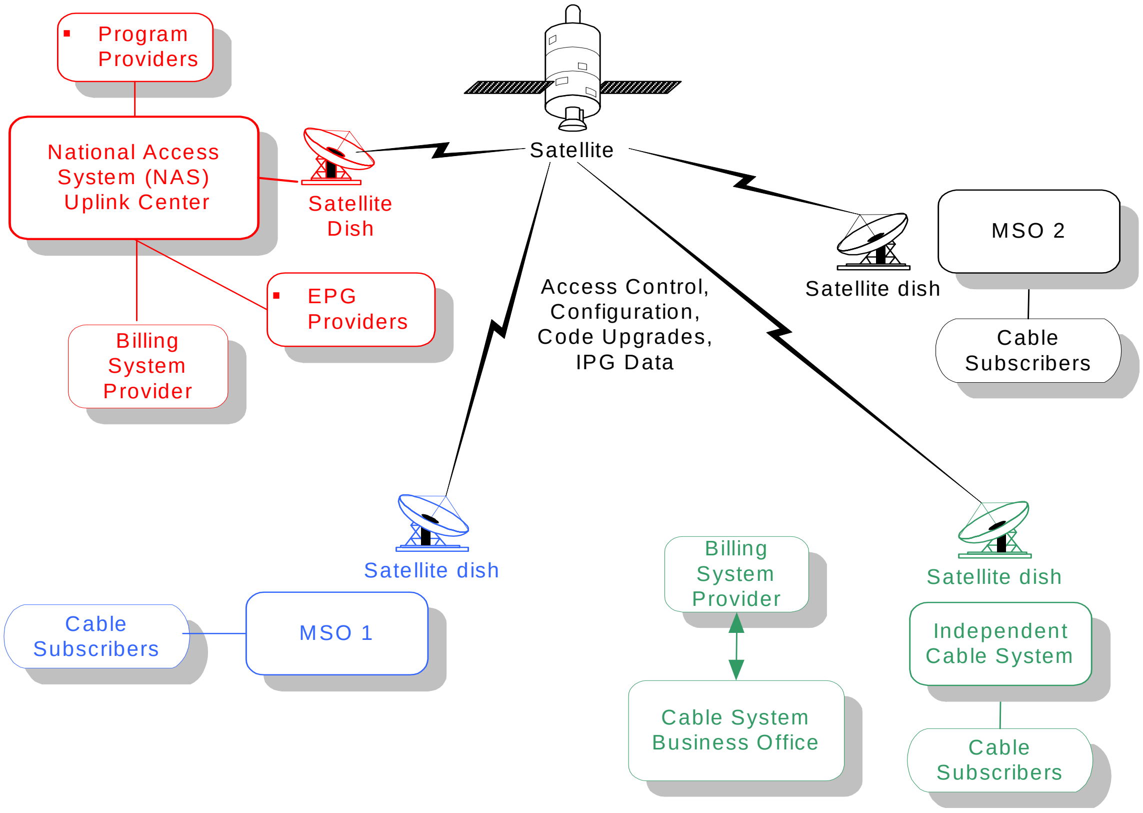

Figure 4 National Control Operational Costs Shared Across Multiple Systems (per Set-top Charge)

The access control function provides for configuration of the MPS encryption function as well as the authorizations for the cable subscribers (set-top terminals). The configuration items outside of the set-top terminal access control functions is also controlled and distributed from the national center. Lastly, the code upgrades are largely driven from the national center although it is also possible to drive the upgrade from a local source as well for early deployments or other MSO/Cable System desired activities.

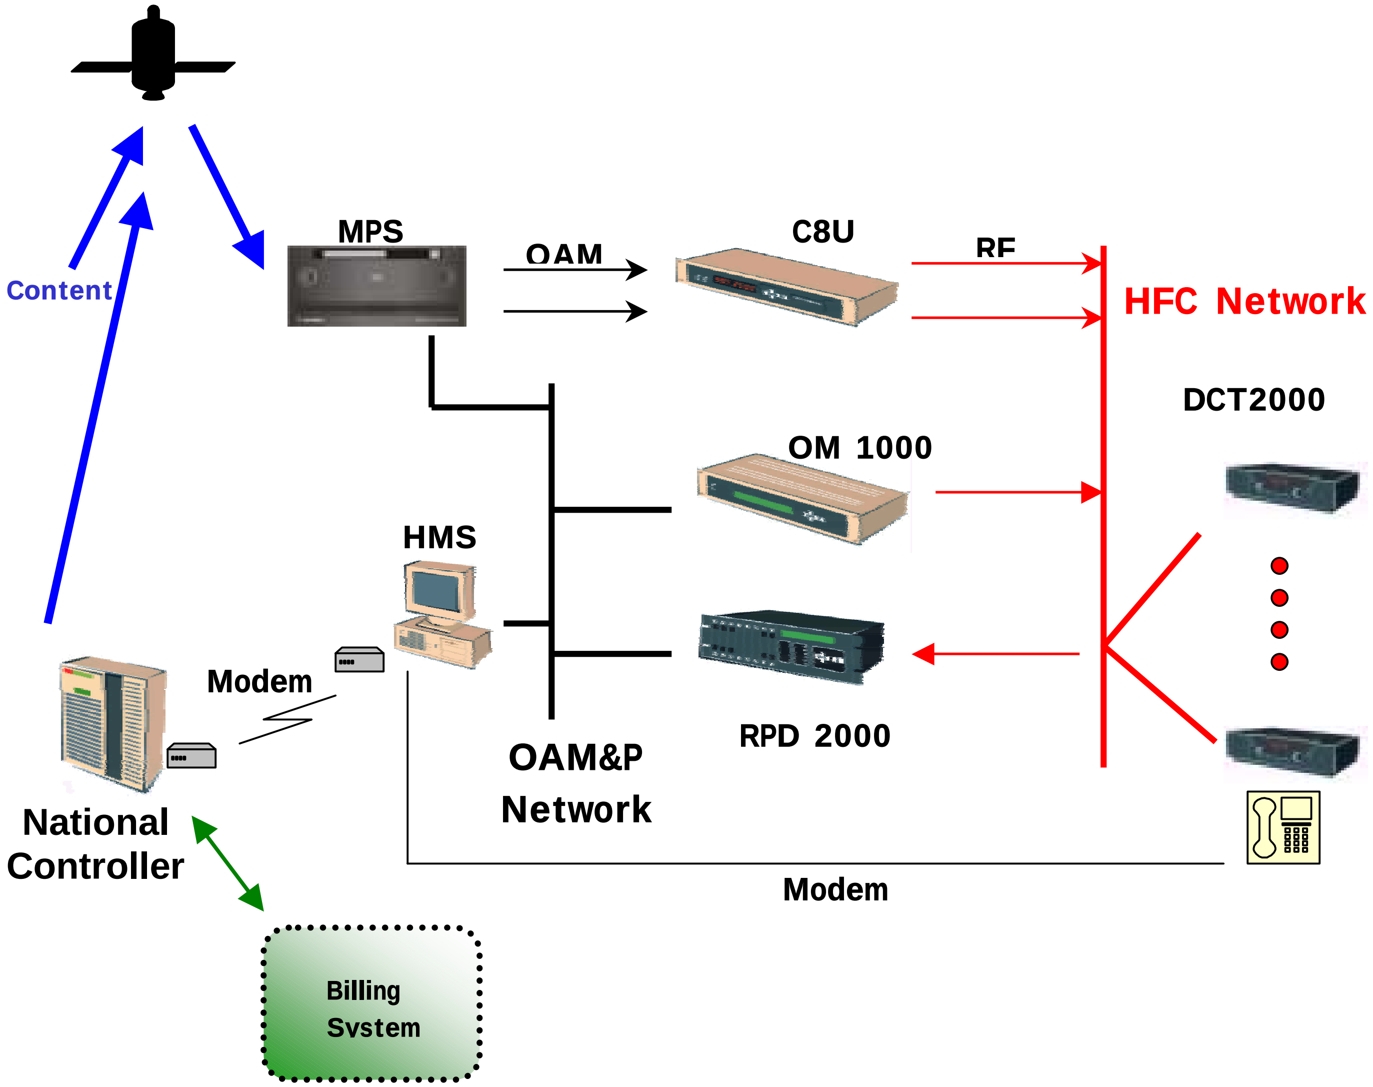

Figure 5 Standard Headend Architecture with National Control

The standard architecture for a National Control Headend includes a Headend Management System (HMS) device. The device has three functions…bootp server, headend device configuration control, and poll collection for IPPV.

Through a dial-up connection to the HMS, a national control specialist is able to check, verify or correct configuration settings with equipment on the OAM&P network including configuration of new equipment. Expanding service offerings becomes simple and low cost.

Headend equipment can be added or firmware updated through the bootp functionality in the HMS. Specific equipment configuration follows.

The billing system interface is through the national controller – not the local headend. In the standard configuration, purchase collection messages originate in the HMS; the return can be RF or telco-return. If telco is employed, the call-in number can be a local phone number or a toll-free phone number. Collections are stored locally and forwarded to the national controller periodically. From the national controller, purchases are authenticated and uploaded to the billing system. Purchases are not cleared from the terminal until the purchase is authenticated which prevents loss in the data collection path (DCT → HMS → National Controller).

SMALL SYSTEM ARCHITECTURE

Taking More Cost Out

With an HMS in the headend, the headend cost can be prohibitively expensive for smaller cable systems. Small system architecture shown in figure 7 further removes cost from the headend by centralizing the HMS functions. The equipment represents the minimum headend functionality required to provide digital cable TV. The only way to minimize the headend cost further would be to put the functionality in the set-top. Such a move would price the low-end set-top too high and not allow for growth of services within the given architecture.

A simple bridge is used to allow the same configuration and bootp functions to be hosted by a centrally located HMS. The specific headend configuration is loaded on a central configuration HMS prior to an attempt to connect to the local headend network. The central architecture includes a (second) dedicated central purchase collection HMS. Purchase collections are slightly different in that the actual poll messages originate from the national controller (rather than the HMS) and the terminals are directed to call a toll free number connected to the central purchase collection HMS. Purchases are handled the same way as a standard configuration (stored in the HMS and forwarded to the national controller; authenticated at the national controller, uploaded to the billing system and cleared from the terminal,).

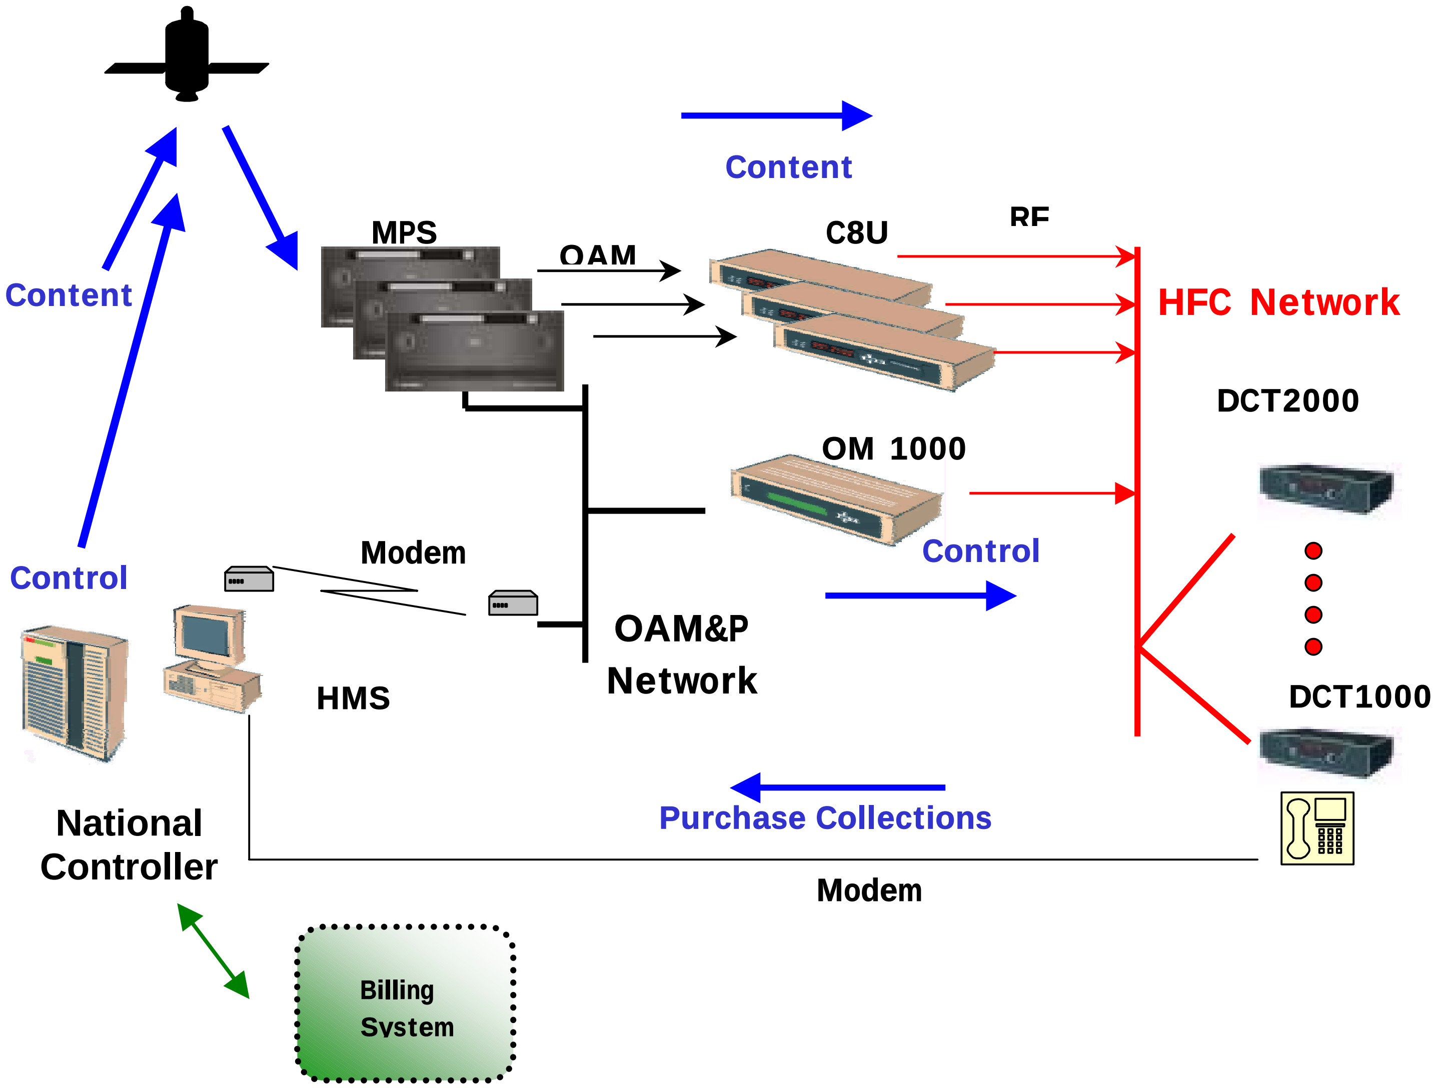

Figure 7 Small System Architecture

Configuration Support

With the small system architecture, the operator has a very flexible system for delivering digital programming. The system can be upgraded to support many different configurations that allow for the most effective use of system bandwidth. Any system can be upgraded to support service grooming, digital ad insertion, digital off-air reception, digital encoding (locally) of analog services, and 256QAM. These capabilities allow the operator to deliver more types of services in a more Bandwidth efficient manner.

Broadcast iTV Support

With the small system architecture, deployed systems need not upgrade to two-way in order to enjoy compelling interactive applications. Broadcast applications are being deployed for the DCT set-tops that will allow one-way systems to launch applications such as gaming, stock tickers, and local weather updates on demand. These applications can be delivered via local servers but will also be delivered via the National Control stream. These applications are enabled by Motorola's Horizons developer program and currently testing is underway with applications from Liberate and OpenTV. These broadcast applications offer the operator a formidable weapon in combating Dish in one-way systems.

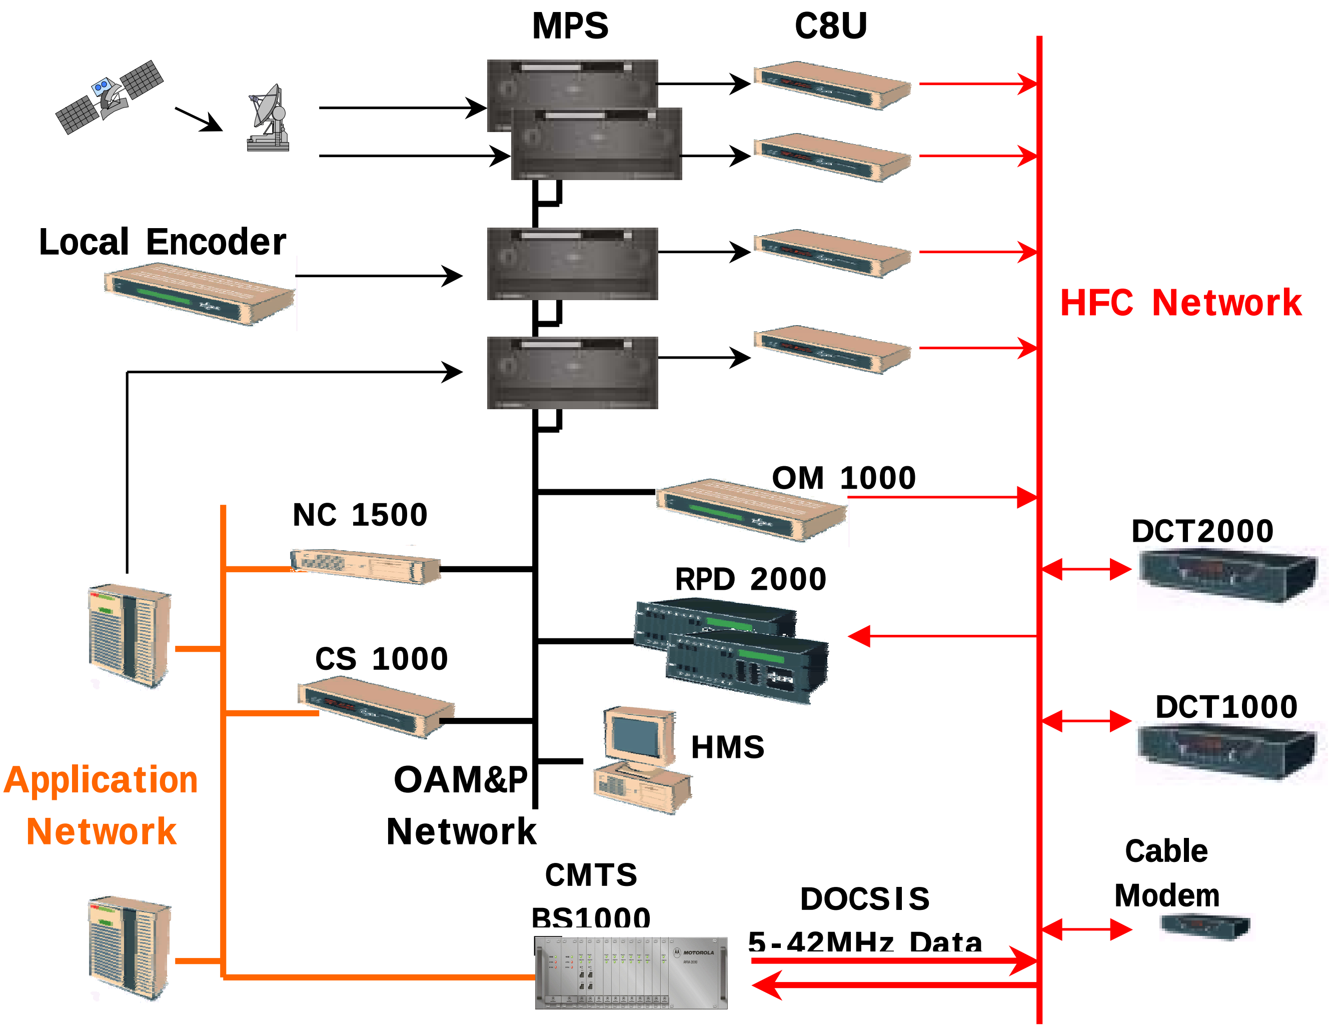

Figure 8 Headend Expanded for Multiple Services

Upgrading Services – Accommodating Service Growth

VOD, Internet Access, and High Speed Data Services (HSD) require an expanded capability and additional equipment. The final figure above illustrates the potential to grow the small system architecture to include much more comprehensive service offerings. The expansion details are beyond the scope of this article, but it is clear that none of the equipment from the small system architecture is retired as the headend expands.

CONCLUSION

A digital headend can have a small enough footprint to be self-contained in a single rack or can be integrated into existing racks. There are no compatibility issues since the digital headend is completely independent of the analog headend and also no dependency on the vendor for the analog set-tops. Most billing systems in North America are already supported.

The approach for the small system architecture is to take the cost out of the headend and provide affordable options for the set-top. Functions such as conditional access, maintenance and configuration, polling and collections are removed to a central function. Along with the functions, significant capital equipment costs are removed and the incremental operational costs are spread over a large number of set-tops. The National Control Architecture provides affordable revenue generation, competition to DBS service offerings, and service expansion options through deployment of digital cable to systems with as few as 1000 basic subscribers. As digital penetration increases, the system size where this architecture makes sense will shrink. Accomplished with a simple, reliable and well-proven architecture.

Contact information:
mhicks@motorola.com
cpoli@motorola.com

STATISTICAL PROPERTIES OF COMPOSITE DISTORTIONS IN HFC SYSTEMS AND THEIR EFFECTS ON DIGITAL CHANNELS

By Dr. Ron D. Katznelson, CTO
Broadband Innovations, Inc., San Diego CA

ABSTRACT

The statistical properties of CTB and CSO distortion terms associated with the analog channels carried on Hybrid Fiber Coax (HFC) are presented. Both simulation and measurement results show that these distortion components falling on individual channels have amplitude Probability Density Function that is nearly Rayleigh distributed (having a Standard Deviation of 5.7 dB). It is shown that both CTB and CSO components have significant likelihood of having peak envelope power fluctuations that exceed their average (measured) power levels by more than 15 dB. The temporal statistical properties of these distortion components are also examined and evidence for peak envelope power fluctuations with characteristic times on the order of 100 microseconds is presented. The implications of these statistical properties for 256 QAM digital downstream channels on which such CTB and CSO components fall are discussed. It is shown that some currently prevailing link budget design practices do not provide sufficient margin for reliable 256 QAM operation, particularly in systems that are rich with Narrowcast combining of digital channels such as cable modem and Video on Demand (VOD) applications. Several mitigation strategies and improved design practices are subsequently reviewed. These involve physical layer choices for longer digital interleaver depth, establishing QAM frequency offset relative to the analog channel grid, improved CTB and CSO specifications for cable modems and digital set-top tuners and tighter aggregate noise floor specifications for head-end RF transmission gear.

1 Introduction

Nonlinear distortions and limited dynamic range in multichannel carrier systems has received renewed attention with the proliferation of digital devices and the industry's move to launch 256 QAM digital services. Now, composite distortion components often become the bottleneck that it never was for the predominant 64 QAM service. Although the installed base of digital subscriber set-tops and cable modems is predominantly 256 QAM capable, many cable operators have yet to migrate to 256 QAM transmissions successfully. As it turns out, the 256 QAM's whopping 44% increase in channel information payload capacity over that of 64 QAM comes at a considerable cost in noise and interference margin requirements. Operating reliable 256 QAM downstream links is more challenging due to the factors described in Section 5. Suffice it to say here that as a result of the cumulative effect of these factors, the noise margin left for further noise degradations in 256 QAM all but disappears [1].

2 Composite Distortions Effect on Digital Channels

Composite Triple Beats (CTB) and Composite Second Order (CSO) distortion components produced by intermodulation of analog TV carriers can become the dominant degrading factor for digital cable QAM channels [1],[2],[3],[4],[5]. It is generally accepted that in thermal noise environments, a Carrier to Noise Ratio (*C/N*) of 30-35 dB provides adequate operating margins for 256 QAM [6]. It is also generally known that cable systems normally operate with CTB or CSO levels that meet or exceed a Carrier-to-Interference (*C/I*) ratio of 53 dB, as required in Reference [6]. These levels correspond to a *C/I* of 47 dB for a QAM signal carried with levels that are 6 dB below the analog carriers. The question that may arise is *why should distortion components having average power levels that are 12-15 dB below the noise level have any appreciable degradation effects on the QAM channel performance?* The short answer is that the average envelope power of composite distortion terms (which is what one observes in a spectrum analyzer measurement) is deceptively low in comparison to its occasional peak envelope power.

3 Description of Composite Distortion Terms

In a multicarrier CATV system one often represents the signal to be communicated as

$$\textbf{(1)} \qquad S(t) = \sum_{n=1}^{N} A_n \cos(\omega_n t + \theta_n),$$

where A_n, ω_n and θ_n are the amplitude, the angular frequency and phase of the n$^{\text{th}}$ carrier signal. We shall have occasion to investigate the unmodulated case, $A_n = A$ for all n and the modulated case, where the modulations in A_n are sufficiently slow compared to the essential RF period of $S(t)$ so as to treat them as constants (or as random variables) during such RF period discussed below. This assumption is justified since statistically, the baseband video modulation power spectrum has negligible energy at 6 MHz or 1.25 MHz, which are possible nominal periodicities of the RF signal. Furthermore, for television signals, we shall assume that other subcarriers such as the single sided audio and chroma subcarriers associated with each television carrier are sufficiently low in amplitude in comparison with the visual carrier so as to treat them as a separate additive noise process. The above assumptions permit us to treat Equation (1) as a reasonable representation of the signal in our problem.

Upon subjecting $S(t)$ to a memoryless nonlinear distortion device having a transfer function given by $y = x + \alpha x^2 + \beta x^3$, where x is the input voltage and y the output voltage, we obtain intermodulation products of second and third order and in this work we have assumed no higher order distortion exist in the nonlinear transfer function. The type and relative amplitudes of the intermodulation components that are so obtained are well known and are tabulated elsewhere [7]. We shall be interested in components of the following form:

(1-A) Second-Order

$$\frac{1}{2}\alpha A_n A_k \cos[(\omega_n \pm \omega_k)t + \theta_n \pm \theta_k]$$

(1-B) Third-Order

$$\frac{1}{4}\beta A_n A_k A_m \cos[(\omega_n \pm \omega_k \pm \omega_m)t + \theta_n \pm \theta_k \pm \theta_m]$$

wherein for a frequency plan in which virtually all carriers are nominally 6 MHz apart, many distortion components fall on the same nominal frequency. The way these terms combine and the amplitude excursions of the resultant composite waveform falling on each channel depend on the number of carriers N, the specific arrangement of the amplitudes, the frequencies and the phases (A_n, ω_n and θ_n). It is important to note that even in the simplest case in which we assume no modulation (i.e. $A_n = A$ for all n), a particular choice of the remaining parameters may submit to a solution that will not apply to any other combination. Furthermore, the exact values of these parameters may never be known for any particular head-end. Of particular significance may be identifying any frequency combinations that may yield coherent combining (for which $\omega_p = \omega_n \pm \omega_k \pm \omega_m$, for example) and taking account of the special results it entails.

Moreover, if the whole set of channels is coherent and locked to a comb (IRC or HRC), then one would need to specify only each of the phase values θ_n and the result will be specific to that phase constellation but not necessarily to any other[1]. It now becomes clear that countless combinations would have to be considered as special cases and the usefulness of any of them for making general statements is at best doubtful.

An alternative approach we adopt here is to make the (realistic) assumption that all the frequencies ω_n are linearly independent over the rational field, or stated another way, that there exists no set of integers Q_n's (positive, zero or negative) such that

$$\sum_{n=1}^{N} Q_n \, \omega_n = 0 \, .$$

It is only under such an assumption that one may obtain results that apply across many different frequency realizations within the tolerance of the nominal frequency plan. For each such realization (a given collection of ω_n and θ_n) in our model, an *unmodulated* signal $S(t)$ of Equation (1) is deterministic and nothing about it is random. Because the precise frequencies are assumed to be linearly independent, the period of the signal $S(t)$ is infinite. Consequently, the period of any of its distortion components is also infinite. With this assumption, it is mathematically possible to show that a characterization of *any* feature, measure, frequency of occurrence or distribution over a long enough observation period T from t_0 to t_0+T, would be repeatable regardless of when t_0 is. This, in general, cannot be assured for choices of exact frequencies that are linearly dependent over the rational field.

We note that, strictly speaking, the nominal frequencies of the Standard CATV channel plan are a linearly dependent set of frequency values. However, in non-coherent headends, their actual values are given by[2]

$$\text{(2)} \qquad \begin{aligned} \omega_n &= (n+8)\cdot\omega_C + \omega_0 + \Delta\omega_n \\ &= 2\pi\,[(n+8)f_C + f_0 + \Delta f_n] \end{aligned}$$

where f_C = 6 MHz, f_0 =1.25 MHz (the visual carrier offset from the channel edge) and Δf_n is the individual actual deviation from the nominal frequency due to tolerance. These deviations can be on the order of a few hundred Hz to a few kHz, and are the basis and justification for the assumption we make about their linear independence over the rationales[3]. When one uses the values in Equation (2) for the frequencies one obtains composite distortion frequencies that have linear combinations of the deviation values. For example, third order components falling on frequencies $\omega_n + \omega_k - \omega_m$ have frequency values of $p\omega_C + \omega_0 + \Delta\omega_n + \Delta\omega_k - \Delta\omega_m$ with p an integer designating the channel. Similarly second order terms falling on frequencies $\omega_n + \omega_k$ have frequency values of $q\omega_C + 2\omega_0 + \Delta\omega_n + \Delta\omega_k$ with q another channel designating integer. Because the deviation frequencies $\Delta\omega_n$ are linearly

[1] In fact, there are very atypical (but deliberate) phase constellations that yield very low distortion values that are otherwise extremely unlikely to be found at random. A method to arrive at such constellations is described in [8].

[2] We are ignoring the special case of Channels 5 and 6 that are offset by 2 MHz. Throughout this study, these channels and all others up to and including A4 were omitted, i.e. A_n=0 for n=4 through n=9 in Equations (1) and (2) above. The last channel in the array was at n=81, making up a total of 75 carriers.

[3] Because there is no relationship between the deviations Δf_n for each channel in a non-coherent system, these values can be thought of as drawn at random from the real line near the origin. Because the rationals form a set of measure zero on the real line, the probability that any random pair of numbers Δf_n and Δf_k drawn from the real line are linearly dependent over the rationals (i.e. that $\Delta f_n / \Delta f_k$ is rational) is zero.

independent over the rationales for all values of *n,* we conclude that *all possible beat frequencies produced by any combination of such deviation terms must all be distinct frequencies*.

The observation above provides certain clarity and resolution to predictions of the average power levels of composite distortion terms, as there are no terms that can combine coherently. In this case, counting the number of composite terms on any given frequency range would yield the exact power level of such components. Thus, precise derivations for such numbers [9] or tight approximations [10] can be very useful and the theoretical prediction of the average power in the unmodulated case is therefore straightforward. It is other characteristics of the composite distortion terms such as the envelope fluctuations and their temporal behavior that are of great interest. To that end, this author has embarked on a theoretical study comprising an analytical effort [11] and an empirical simulation effort, parts of which are reported in the next sections.

4 Mathematical Background and the Simulation Approach

By rewriting Equation (1) using the defined frequencies of Equation (2) we find that the simulation problem at hand is that of forming representations of, and operations with, *S*(*t*) given by:

(3)

$$S(t) = \\ = \sum_{n=1}^{N} A_n \cos\{[(n+8)\cdot\omega_C + \omega_0]t + \Delta\omega_n t + \theta_n\}$$

wherein we have grouped separately the frequency terms that correspond to the (linearly dependent) nominal frequencies. In studying the distortion components we would need to select the frequencies and apply the expression for *S*(*t*) from Equation (3) into the nonlinear transfer function $y = x + \alpha x^2 + \beta x^3$ by setting $x = S(t)$ and process the results to obtain a measure of interest that is representative on average over the observation time interval. If, for example, we wish to study a composite second order frequency-addition beat of the form given in Equation (1-A) above we would have a collection of like-frequency terms that fall on the sum-frequency of interest designated by the frequency index p:

(3-A)

$$y_{CSO}(p) = \\ = \sum_n A_n A_{p-n} \cos[(\omega_n + \omega_{p-n})t + \theta_n + \theta_{p-n}] = \\ = \sum_n A_n A_{p-n} \cos[(p\omega_C + 2\omega_0 + \Delta\omega_n + \Delta\omega_{p-n})t + \theta_n + \theta_{p-n}]$$

These terms form a narrow band signal about the frequency $p\omega_C + 2\omega_0$ with an envelope function given by

(3-B)

$$\text{Env}[y_{CSO}(p),(t)] = \\ = \sum_n \sum_m A_n A_{p-n} A_m A_{p-m} \times \\ \times \cos[(\Delta\omega_n + \Delta\omega_{p-n} - \Delta\omega_m - \Delta\omega_{p-m})t + \\ + \theta_n + \theta_{p-n} - \theta_m - \theta_{p-m}]$$

where the meaning of a signal's envelope is that ascribed to it in the classic treatise by Dugundji [12]. In this work we are focusing on the statistical characteristics of the envelope (amplitudes) of these composite distortion terms rather than their RF statistics because it is more closely related to the demodulated baseband disturbances that affect BER performance of digital set-tops and cable modems. When we analyze such envelope terms, we obtain functions that depend on the amplitudes, the deviation frequencies and phases $\Delta\omega_n, \theta_n$, but not explicitly on the actual frequencies. Thus, for a given signal *S*(*t*) and a given set of amplitudes, forming a statistical measure over time associated with such envelope terms essentially amounts to time-averaging of some (nonlinear) function $F_S(\Delta\omega_n t, \theta_n; n=1 \text{ to } N)$. That function might be the envelope's value occurrence rate or joint correlation products of envelope values, etc. The function $F_S(\bullet)$ is representative of the kind of processing, data collection and classification that would be used. Fortunately, the assumption that the frequencies of $S(t)$ are linearly independent provides a powerful tool to simplify the simulations of these time-averages. It is based on a fundamental theorem in Ergodic Theory known as the Kronecker-Weyl theorem [13]:

4.1 Kronecker-Weyl Theorem.

The Kronecker-Weyl (K-W) theorem for multiply periodic functions $F(\psi_1, \psi_2, \cdots, \psi_N) \equiv F(\mathbf{\Psi})$ defined on the N dimensional torus, on which each of the ψ_j ranges from 0 to 2π and on which $F(\mathbf{\Psi}') = F(\mathbf{\Psi}'')$ if $\psi_j' = \psi_j''$ (modulo 2π) for all j, states the following:
If

(4) $$\psi_j = \Omega_j t + \theta_j$$

and if the frequencies Ω_j are linearly independent over the rational field, then
(5)
$$\frac{1}{T}\int_0^T F[\mathbf{\Psi}(t)]dt \quad \rightarrow$$
$$\rightarrow \frac{1}{(2\pi)^N}\int_0^{2\pi}\cdots\int_0^{2\pi} F(\psi_1, \psi_2, \cdots, \psi_N)\, d\psi_1 d\psi_2 \cdots d\psi_N$$
as $T \rightarrow \infty$
.

Stated another way, under the conditions that satisfy the K-W theorem, we can avoid having to select the frequency deviations $\Delta\omega_j$ and replace time averages by averages over all the phases. In reference to the CSO envelope term expressed above (as also applies for CTB or other distortion terms), for our simulations, we identify Ω_j of the K-W theorem with $\Delta\omega_j$ and by using Equation (4) we rewrite Equation (3):

(6)
$$S(t) = \sum_{n=1}^{N} A_n \cos\{[(n+8)\cdot\omega_C + \omega_0]t + \psi_n\}$$

In our computer simulations we thus form an ensemble of uniformly distributed random phase vectors $\mathbf{\Psi} \equiv (\psi_1, \psi_2, \cdots, \psi_N)$ with independent components on the N dimensional torus. For each such randomly drawn phase vector we compute $S(t)$ in accordance with Equation (6) and obtain a function of t that becomes a member of the analysis ensemble. Since we have selected the Standard frequency plan, the period of every $S(t)$ of the ensemble is 1.25 MHz and after appropriate model distortions are introduced, an FFT routine is employed to obtain envelopes of composite distortion terms on specific frequencies. It should be noted that the dependence on time remaining in $S(t)$ of Equation (6) is used only to obtain the spectral components of interest, and not for time averaging.

4.2 Modulation Statistics Assumed

In addition to simulations of the CW case, we simulated independent modulations on each of the carriers. Because in these simulations, we were focused only on first-order statistics (i.e. amplitude distributions and joint densities), time correlation effects (i.e. second-order statistics) associated with the analog video signals were not needed and only first-order statistics of the amplitudes were used. These distributions were computationally constructed from first-order statistics of NTSC luminance values, as they control the amplitude of the visual carrier. NTSC Luminance distribution values were taken from the work done by Philips researchers [14] in their work that led to the PAL system.

Appropriate account was made for blanking (75% level) and synch (100% level) periods and the resultant distribution is shown in Figure 1. Effects of amplitude modulation were simulated by forming the ensemble of functions of Equation (6) by the use of an expanded ensemble of random variables including the amplitudes, $\{\psi_1, \psi_2, \cdots, \psi_N, A_1, A_2, \cdots, A_N\}$, where the phases are drawn independently from a uniform distribution over [0, 2π] and the amplitudes are drawn independently from the distribution in Figure 1. That way, an ensemble of millions of realizations of $S(t)$, and consequently all related CSO and CTB distortion terms on every channel, was computed and analyzed. The results for the modulated and the unmodulated cases are shown in subsequent sections.

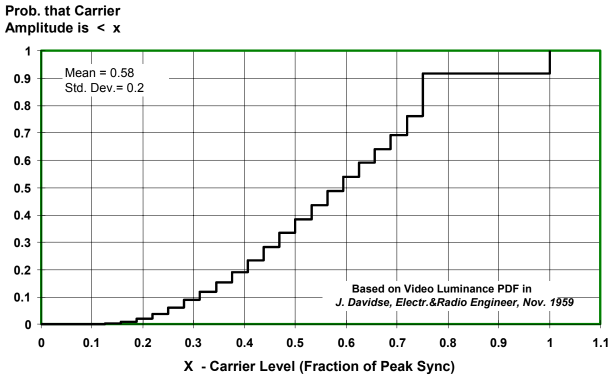

Figure 1. Cumulative probability density of RF amplitudes of a video modulated NTSC television signal. This function was used for simulating the independent modulation of 75 carriers. The function was constructed based on video luminance distributions provided in Reference [14].

4.3 Composite Distortions – Simulation Results and Discussion

4.3.1 Amplitude Statistics of Composite Distortions

Very few studies were found related to the probability density functions of CTB/CSO distortion components. One histogram was reported for a CW HRC system [15], although it was not clear what the source of phase fluctuations was. A more comprehensive landmark study involving real headend in a controlled laboratory was reported in Reference [4]. Both our simulations and the above referenced measurement results show that these distortion components falling on individual channels have amplitude Probability Density Functions (PDF) that are nearly Rayleigh. The Rayleigh probability density function is given by [16]:

$$f(x) = \frac{x}{a^2} \exp\left(\frac{x^2}{2a^2}\right)$$

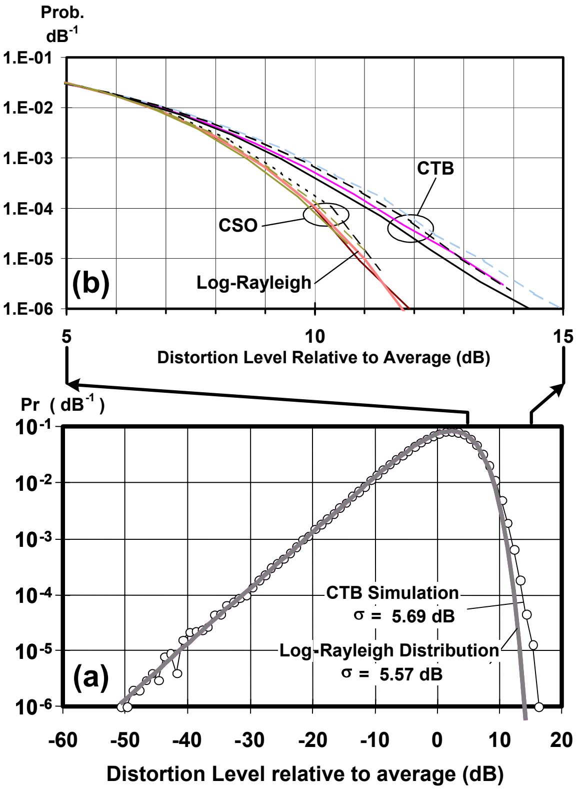

Figure 2. Probability density function of the amplitude of simulated CTB and CSO terms due to third order and second order distortions in a CATV system using the Standard Frequency plan with 75 analog video modulated carriers. On the ordinate, the amplitude squared is expressed in dB relative to the average envelope value and the corresponding probability per dB is plotted in the co-ordinate. For comparison, a Rayleigh density transformed to the dB scale (thus called Log-Rayleigh) is also shown in solid line. An expanded view from 5 to 15 dB is provided in (b), where curves for the CW case and the modulated case are shown with solid lines and broken lines respectively.

Figure 2a shows the simulation results for the probability density function of the amplitude of a simulated CTB term due to third order distortion in a CATV system employing the Standard Frequency plan with 75 analog video modulated carriers. For comparison, a Rayleigh density transformed to the dB scale (thus called Log-Rayleigh) is also shown. The Rayleigh density is relevant in this comparison because it describes the envelope probability density of a Gaussian process, which is the limit case for a linear combination of a large number of arbitrarily distributed, statistically independent random processes. One would expect that because the CTB and CSO terms encountered in CATV are comprised respectively from thousands and hundreds of beat components, the Rayleigh limit would be nearly approached. However, while there appears to be a good match at lower amplitude levels, relevant deviations at large relative amplitudes are seen for CTB in Figure 2b. This slight deviation is due to the statistical dependence among these thousands of components. It is argued that, in this case, because there are only 2×75 = 150 *independent degrees of freedom* (random phase and amplitude of each of the 75 carriers) affecting the phase and amplitude values of the individual *thousands* of beat components, these must be statistically dependent. Because the number of components in the CSO term is of the order of less than 100, these component values depend less on common degrees of freedom and are less correlated. The degree and differences of the correlations for CSO and CTB terms and their respective effects on envelope statistics are further addressed in [11].

Noteworthy is the observation that as one inspects distortion terms on different channels, the absolute levels of the envelopes differ, but when normalized to their respective average envelope power, the PDFs of the CTB and CSO components on every channel essentially follow closely the results for the CTB term shown in Figure 2(a). Moreover, each of these distributions, like the Rayleigh distribution, has a variance that is proportional to its mean. When used on a 'dB relative to average' scale, this means that the standard deviation σ in dB is constant. For the Log-Rayleigh, case it can be shown that σ = 5.7 dB and we note that the results for CSO/CTB (as seen in Figure 2) terms do not differ appreciably from this value, although at the tail of the distribution CTB peak values are slightly higher than those of the CSO and Log-Rayleigh terms (Figure 2b). In that figure, the modulated case (broken lines) is seen to have higher relative envelope values than the CW case for the same probabilities with a difference of less than half a dB at the level of interest. Although the absolute values for the modulated case are some 2.5 dB lower than the CW case (having peak sync amplitudes), we are showing here the results normalized to average power.

The trends observed in these simulations are fairly consistent with measurement results of a real headend as reported in Reference [4] (see Figures 8 and 9 in that reference), although the difference between the modulated case and the CW case in the measurement results seems to be slightly larger. It is conjectured that because the video sources in that system were received from a satellite feed, as would be the case for many cable systems, some were in fact correlated, as they carry program material that is frame synchronized across various channels[4]. Hence, more of the variations attributable to amplitude modulation were based on even fewer degrees of freedom, increasing the relative statistical dependence among these components.

[4] Many of the services such as the HBO group of channels come from a single uplink facility wherein the practice of video 'house sync' is used. Other examples of groups that are likewise synchronized are the Viacom feeds and the USA Network feeds.

Turning back to the question raised at the outset, tying these statistics to the impact on digital channels is rather straightforward. As can be seen in Figure 2b, both CTB and CSO components have significant likelihood of having peak envelope power fluctuations that exceed their average power levels by more than 12-18 dB. Unfortunately, when such fluctuations occur in a composite distortion term having an *average* level of –47 dBc, levels up to –29 dBc can be experienced. Thus, even in situations in which the prevailing noise floors are rather intuitively low, significant impairments can still be observed. To relate these results to 256 QAM Bit Error Rates (BER) and consequential video impairments we turn next to the temporal characteristics of these composite distortion components.

4.3.2 Time Domain Statistics of Composite Distortions

While we have not simulated the second-order properties of these distortion components, several observations can be made. When high envelope fluctuations of composite distortion terms occur, bursts of bit errors can be generated on the 256 QAM link. Degradations due to bursts that are relatively short can be mitigated in part by the interleaving used on the digital link. Such interleaving is part of the Forward Error Correction (FEC) system used in QAM transmission. Clearly, there would be more severe degradations for burst errors with durations that exceed the interleaver depth capability. Figure 3 shows experimental results that demonstrate that fact for 256 QAM. As can be seen in this figure, substantial reduction in coded error rate is achieved with interleavers that have depths longer than 50-100 microseconds. Improvements start to level off above 200 μs, indicating that there are much fewer burst events that are longer than 200 μs. This observation shows that the characteristic duration of CSO peak envelope fluctuations is on the order of 100 μs. These results appear consistent with other studies on CSO/CTB distortion effects. See for example the similar average durations (60 μs) reported by Germanov [2].

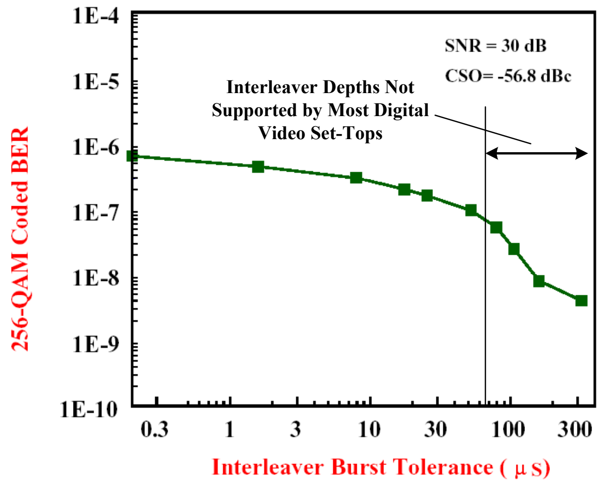

Figure 3. Coded bit error rate of 256 QAM link perturbed by CSO distortion as a function of interleaver burst span. The 256 QAM signal was carried 5.6 dB below the analog carriers, making the CSO average power level shown in the figure -51.2 dBc interference for the QAM channel. Source: Ref. [3].

While both authors in References [2] and [3] report the phenomenal fact of this characteristic duration times, neither suggests a mechanism, a reason or a cause for the observed 100 μs characteristic times. As explained below, these characteristic times are simply related to the correlation times of the individual composite distortion components. These, in turn, are inversely related to the effective bandwidth occupied by such distortion components.

Figure 4 depicts a conceptual rendition of the power spectrum of a composite distortion term produced in a non-coherent frequency plan. Because most of the power of these distortion terms is distributed around frequencies that are a linear combination of the nominal frequencies of the contributing carriers [i.e. $f_1 \pm f_2 \pm f_3$ (for CTB) or $f_1 \pm f_2$ (for CSO)], the spectral spreading for most of the energy will be on the order of several times the *frequency tolerance* of all transmitters. Some energy due to video modulation in the first (dominant 15.7 kHz horizontal rate) sidebands will widen this somewhat to a power bandwidth on the order of 10 kHz. For such narrow-band processes, the correlation times are approximately inversely proportional to the power bandwidth Δf and for a 10 kHz wide CTB or CSO spectral "clump", one obtains $\tau = 1/\Delta f = 1/10^4 = 100\mu s$ for the characteristic time. Thus, for these narrow-band composite distortion components, the envelope value cannot change appreciably in less than 50-100 μs, thereby giving rise to durations of peaks on that order. Ironically, modern (and more frequency accurate) transmitters, or for that matter, for carrier frequencies produced in a coherent headend, wherein all distortion terms fall on the same frequency, the effective power bandwidth could be substantially smaller (as shown schematically in Figure 4). This will result in longer characteristic times for burst errors associated with composite distortions. Unfortunately, as labeled in Figure 3, most digital video set-tops do not have sufficient interleaver memory to adequately protect against the *most likely* burst durations, as they are limited to a protection depth of 66 μs in the 256 QAM mode[5]. In contrast, that same memory limit provides these set-tops a longer (95μs) interleaver span in the 64 QAM mode, which further enhances the error tolerance of 64 QAM links as compared to 256 QAM.

It is in the context of this CTB/CSO burst vulnerable regime of 256 QAM that we examine the effect of additive noise in the channel. A *C/N* value of 30 dB, a value some would have judged as sufficient for 256 QAM, is shown in Figure 3 to be dramatically inadequate for 256 QAM operation under CSO distortion values that more than meet most MSO's operating standards.

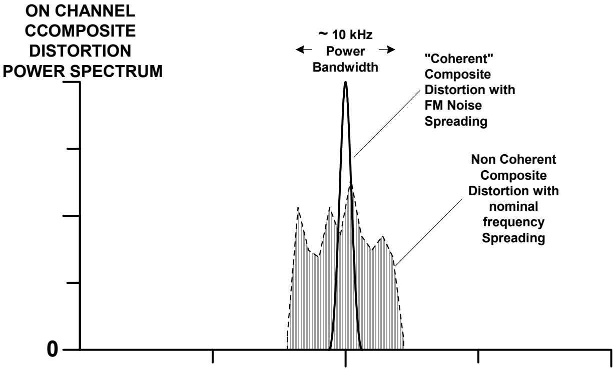

Figure 4. Conceptual rendition of the power spectrum of a composite distortion term produced in a non-coherent frequency plan and that produced in a coherent plan, wherein all distortion terms fall on the same frequency. Because most of the power is distributed around frequencies $f_1 \pm f_2 \pm f_3$ (for CTB) or $f_1 \pm f_2$ (for CSO), the spectral spreading for most of the energy will be on the order of several times the frequency tolerance of all transmitters. Some energy due to video modulation in the first 15.7 kHz (horizontal rate) sidebands will widen this somewhat to a power bandwidth on the order of 10 kHz.

Figure 5 shows measurement results reported in Reference [4] using real set-tops and real head-end live video sources to modulate each channel in the analog tier. Uncoded bit error rates were recorded as a function of total interference and noise power. Note that at target BER values of 10^{-8}, up to an additional 15 dB(!) of Noise+Interference margin is required for 256 QAM over that of 64 QAM. This increment is significantly larger than the 5.6 dB theoretical increment expected in pure random noise channels [17] and is related to the stacking of unfavorable parameters used in 256 QAM as enumerated in Section 5. Consequently, as further elaborated below and in reference to the results shown in Figure 5, this author suggests that under CSO/CTB levels that are within current operational CATV specifications, a 40 dB *C*/(*N*+*I*), corresponding to a 256 QAM BER of 10^{-6}, would *not* provide adequate margin for reliable 256 QAM operation and that rather,

[5] The 66μs protection mode in 256 QAM can be set by selecting I=128,J=1 for the convolutional interleaver.

C/(*N*+*I*) values in the range of 43-45 dB (approaching uncoded BER of 10^{-8}) would be required for a Quasi-Error-Free (QEF) 256 QAM operation[6].

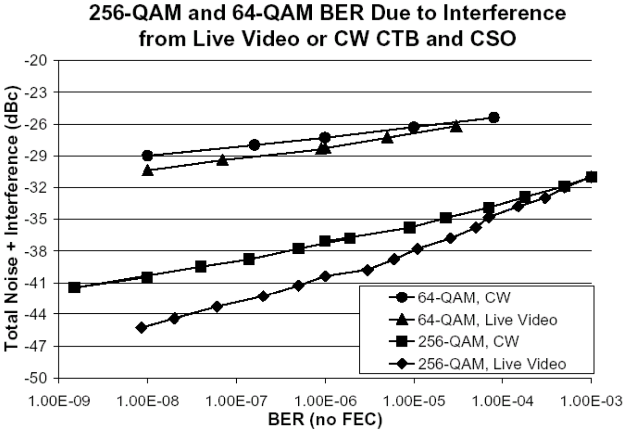

Figure 5. The measured effect of composite distortions on uncoded BER of 256 QAM link as compared with 64 QAM. CTB and CSO distortions were from analog carriers that were either modulated (Live Video) or unmodulated (CW). Note that at target BER values of 10^{-8}, an additional 11 dB of Noise+Interference margin is required for 256 QAM over that of 64 QAM. Source: Ref. [4].

Finally, it should be appreciated that many of the digital 256 QAM channels are likely to be situated well above the 550 MHz boundary and, as such, will be subject to both CTB and CSO. In this case, two distortion contributions should be accounted for. For our purposes below, we shall assume two composite distortion components each at a level of –55 dBc. This corresponds to -49 dBc with respect to the 256 QAM signal, resulting in an average composite distortion level of –46 dBc.

4.3.3 How large can CSO/CTB envelope fluctuations become?

The answer to this question depends on the observation interval. For practical purposes, we shall first present an intuitive approach to this problem, as it will establish better understanding of the reasons for a substantial rethinking of the required margins for reliable 256 QAM operation.

First, we pose the question slightly differently by noting that the problem is essentially that of Level Crossing Rate. Next, we note that according to the discussion above, the envelope of the composite distortion term is essentially uncorrelated with its values separated by more than the characteristic time. Thus, for approximation purposes, we shall make the assumption that values of the envelope at time instances that are more than 300 μs apart are essentially statistically independent. This means that on average, every 300 μs the envelope has an independent ability to achieve any value in accordance with the probability density shown in Figure 2. Stated another way, every 300 μs, we get to make an independent experiment and draw at random a new envelope value with *a-priory* probability shown in Figure 2. If one uses the North American Quasi-Error Free rate as the permissible error event rate, then on average, we have 15 minutes of many 300 μs experiments to encounter one error event. Assuming for simplicity that the average distortion level is at such a level that only the highest excursion causes an uncorrected error, this means that we have $15\times60/(300\times10^{-6}) = 3\cdot10^{6}$ independent experiments of which only one needs to produce the largest envelope value causing an error event. Stated another way, the level crossing probability need only be one in three million (3.33×10^{-7}). By inspecting Figure 2b for CTB and roughly extrapolating down below the 10^{-6} probability density line, one can estimate that by integrating its tail over

[6] Quasi Error Free (QEF) reception should be distinguished from reception at the impairment Threshold Of Visibility (TOV). The Europeans set the QEF standard for DVB at levels achieving less than one uncorrected error event per hour [18]. In North America, J-83 Annex B defines QEF at less than one uncorrected error event per 15 minutes [21],[22]. The latter event uncorrected error rate corresponds to a BER = 4.0E-11 at the input of the MPEG-2 demultiplexer.

the last few dB bins, we would pick up a probability mass of 3.33×10^{-7} for all values above 16 dB. Hence, we thus estimated that the probability that the envelope exceeds the value of 16 dB is 3.33×10^{-7}. This means that the CTB envelope will exceed a level that is 16 dB above its mean at an average rate of once per 15 minutes.

4.3.4 CSO/CTB Relationship to peak amplitudes of multicarrier signals.

Over the years, an intuition seemed to have developed among CATV engineers with respect to the causal relationship of composite distortion fluctuations and excessive amplitude excursions of the composite multicarrier signal. This intuitive view envisions situations in which occurrences of high peak amplitudes of the multicarrier system are the cause of, and are therefore time-coincident with, interference transients associated with excessive distortion components. This view appears to be further supported by observations that a rapid increase in distortion-induced impairment rates are seen when nonlinear devices are driven beyond certain high levels. The intuitive theory often advanced in relation to that observation is that at those high signal levels, the composite amplitudes of the multicarrier signal reaches clipping (saturation) levels which *simultaneously* produces transient distortion products. Many measurements were attributed to such mechanism and many models based on this theory were published and analyzed (see [19] and the extensive reference list in Reference [20]). As we shall se below, peak amplitude fluctuations of CSO or CTB terms produced by *pure parabolic and cubic* nonlinear characteristics can produce sharp degradation effects that are often unnecessarily attributed to the proverbial clipping mechanism at multicarrier signal levels that are significantly lower than actual clipping levels. To that end, simulation results for the Joint Probability Density of the envelope of a multicarrier signal and the envelope of resultant composite distortion components falling out of the analog channel tier (on channel 105), where digital signals might be carried are shown with contour plots in Figure 6. Note that very little correlation exists between the instantaneous levels of distortion terms and the composite signal. It is seen that the highest peak values of CTB (a) or CSO (b) will be encountered even when the composite multicarrier envelope is around its most likely value. Although the CTB plot shows mild correlation at high levels, virtually all occurrences of large peak values for CSO and CTB are found with composite multicarrier amplitudes that are within 2 dB of their average.

In this regard, it is worth noting that an abrupt degradation can result by simply overdriving a *smooth* second-order or third-order nonlinearity. This can readily be seen in Figure 2, where it can be appreciated that at the tail of the density function, where peak values can reach an observable degradation event, there is roughly one order of magnitude increase in occurrence rate for every dB increase of average distortion power. For CTB distortions, this means that if a drive level is just before the Threshold Of Visibility (TOV), overdriving the device by only one more dB, causes CTB to increase by 2 dBc, which corresponds to 100-fold(!) increase in the occurrence rate of visible degradations. At an initial TOV condition of, say, once every 10 seconds, the degradation rate would rise to 10 per second. That dramatic change over a one dB level change would sure look like a “clipping” phenomenon to many folks…but it is not. It is simply the result of *overdriving* the nonlinear device.

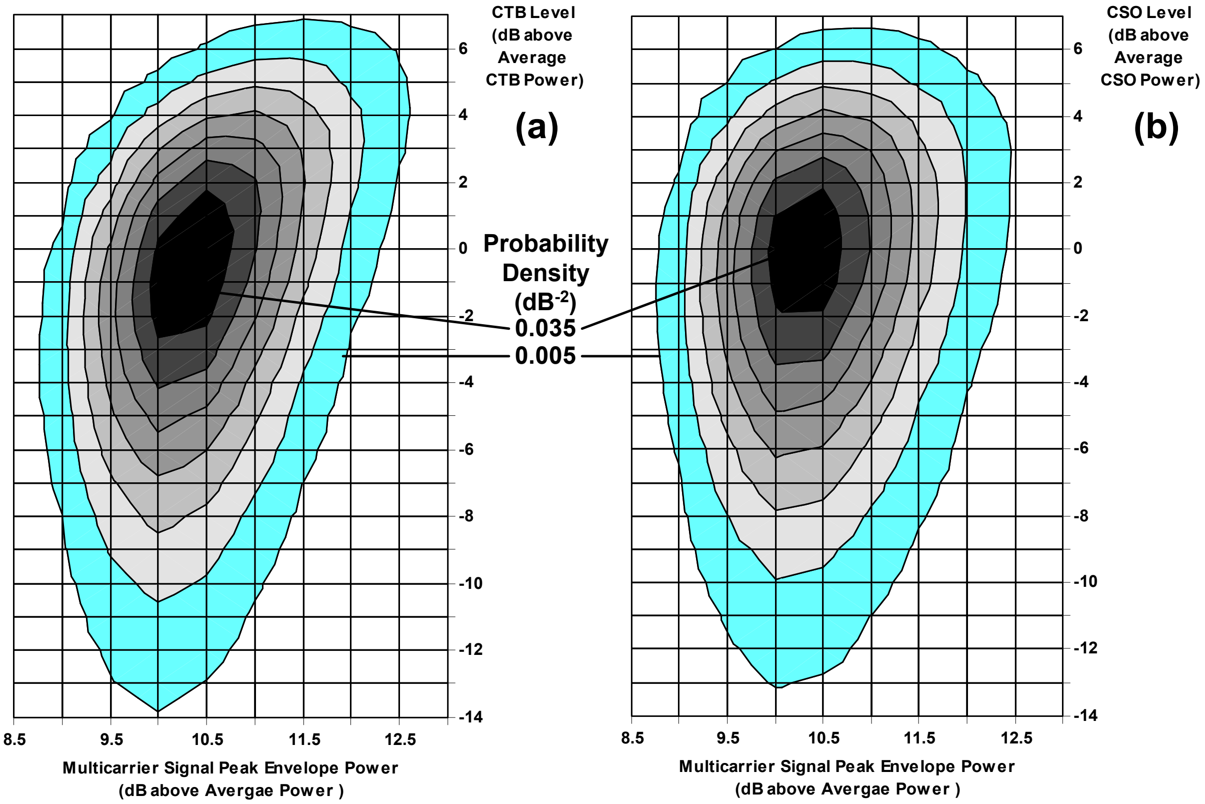

Figure 6. Simulation results for the Joint Probability Density of the envelope of a multicarrier signal with 75 carriers and the envelope of resultant composite distortion components on individual high channels. Note that very little correlation exists between the instantaneous levels of distortion terms and the composite signal. It is seen that the highest peak values of CTB (a) or CSO (b) will be encountered even when the composite multicarrier envelope is around its most likely value. In other words, "clipping" is not the most likely cause of CTB/CSO peak value excursions. Probability density contours are 0.005 per dB^2 apart**.**

In fact, in some devices, clipping values would be far from the lower levels that would otherwise render *smooth* nonlinear devices inoperable due to overdrive. There is no doubt that situations are being encountered wherein clipping is a factor in observed degradations, but as we have seen, "clipping" is not the most likely cause of excessive CTB/CSO peak value excursions. By definition, it is the higher order distortion terms that must be associated with clipping. On this likelihood balance, perhaps a more fitting title for the article in Reference [19] would have been "Don't Get *Overdriven* on the Information Highway"….

5 Moving Forward with 256 QAM

5.1 Where we are

The previous sections illustrated how the challenges associated with operating a reliable 256 QAM service far exceed those associated with the 64 QAM digital services that operators have grown accustomed to. The reasons for the substantial difference in immunity from interference between the two formats can be found by examining their parameters. A summary of relevant factors is shown in Table 1.

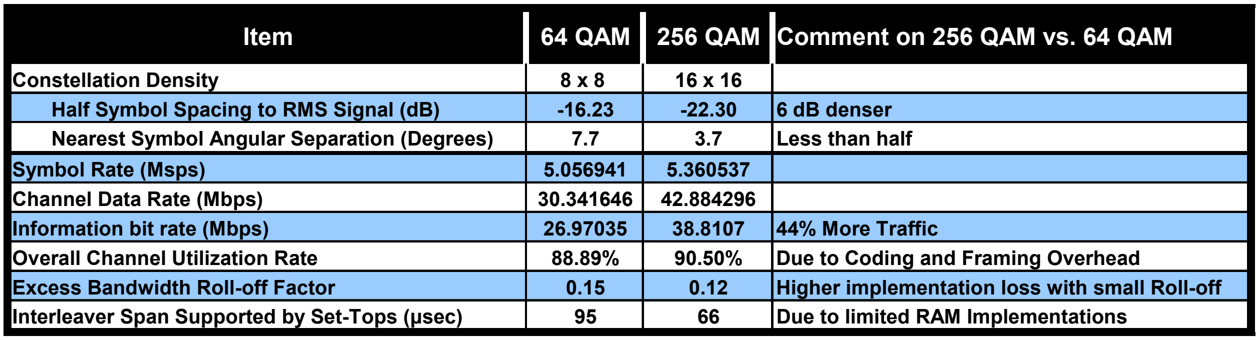

Item	64 QAM	256 QAM	Comment on 256 QAM vs. 64 QAM
Constellation Density	8 x 8	16 x 16	
Half Symbol Spacing to RMS Signal (dB)	-16.23	-22.30	6 dB denser
Nearest Symbol Angular Separation (Degrees)	7.7	3.7	Less than half
Symbol Rate (Msps)	5.056941	5.360537	
Channel Data Rate (Mbps)	30.341646	42.884296	
Information bit rate (Mbps)	26.97035	38.8107	44% More Traffic
Overall Channel Utilization Rate	88.89%	90.50%	Due to Coding and Framing Overhead
Excess Bandwidth Roll-off Factor	0.15	0.12	Higher implementation loss with small Roll-off
Interleaver Span Supported by Set-Tops (µsec)	95	66	Due to limited RAM Implementations

Table 1. Comparison of various transmission parameters of 64 QAM and 256 QAM in the North American standard. See [21],[22].

These factors and others are addressed below:

(a) *Symbol constellation density* is two times that of 64 QAM in each dimension. This means that relative to the signal level, decision regions in the 256 QAM signal space occupy an area that is one quarter of that of 64 QAM.

(b) *Burst error interleaver depth* is limited in most set tops to 66 µs in 256 QAM mode as opposed to 95 µs in 64 QAM mode.

(c) *Coding Gain* is lower for the 256 QAM format which has a trellis code rate of 19/20, dedicating one bit out of 20 for forward error correction versus one out of 15 bits used in the trellis code for 64 QAM.

(d) *Spectral Roll-Off Factor* of 0.12 for 256 QAM (vs. 0.15 in 64 QAM) further reduces demodulator "eye opening" margin, which increases demodulation implementation losses.

(e) *Reduced Immunity to CSO and CTB* distortion products (as discussed above).

(f) *Higher Susceptibility to phase noise* [3] and 'microphonics'.

(g) *Increased Tracking Time of the Adaptive Equalizer.* Due the finer constellation of 256 QAM, channel adaptation time to sufficiently low symbol interference can take longer than in 64 QAM especially if the demodulator employs a smaller step size for the coefficient updates in 256 QAM. A longer Equalizer convergence time can degrade its ability to effectively track and null out a discrete interferer with level and spectral fluctuations at rates that exceed the equalizer convergence time constant[7].

These factors are the realities that cannot be changed by operators. We turn next to steps that can be taken to improve the viability of 256 QAM in more and more systems.

[7] Convergence times of the order of 2,000-10,000 symbols are not uncommon. This means that for a 5 Msps QAM signal, fluctuations faster than 1 kHz would not be tracked well. Unfortunately, as seen above, the characteristic times of the CSO/CTB fluctuations are an order of magnitude faster. However, in the CW mode, these CSO and CTB components have a narrower spectral spreading and consequently might be better tracked by the equalizer. This may explain the data in Figure 5 showing the significant improved immunity of demodulators subject to CW CSO/CTB interference as compared to modulated CSO/CTB interference.

5.2 Where we need to go

In order to successfully employ 256 QAM services we would be well advised to consider the following mitigation measures and improved practices:

5.2.1 Control and improve CSO/CTB levels

This is an obvious item, but operators should consider reviewing their plant distortion budgets. A –53 dBc at the subscriber tap would not leave any margin for subscriber equipment degradations. In addition, CSO/CTB specifications of subscriber tuner distortion warrant further scrutiny. Several cable modem tuner suppliers offer specifications for CSO and CTB that are in the –50 dBc range. As explained above, if such average levels of distortions are encountered (-44 dBc for the digital signal), 256 QAM operation would be unreliable[8]

5.2.2 Employ judicious choice of channel frequency offsets

As described above, CSO and CTB components have most of their energy clustered in a spectral region spanning no more than 30 kHz. The most powerful CTB component consists of a combination of terms with frequencies $2f_1$-f_2, which falls on what would be the NTSC visual carrier frequency for that channel. Thus, it is narrowly concentrated 1.25 MHz above the channel edge. In a paper presented to the Joint EIA-NCTA Engineering Committee a decade ago [24], this author suggested the use of a frequency offset scheme that can virtually eliminate interference from these CTB components by moving the digital channels up by 1.25 MHz with respect to the analog channel boundaries. This way, the CTB components are situated at the edge (in between) of the digital channels where there is considerable rejection by the receiver's SAW filter and the digital matched Nyquist filter. As discussed in [24], this relative offset is also optimal for minimizing third-order noise distortion emanating from the digital channels and falling on the analog channels.

Unfortunately, not all distortion terms are being rejected that way. Second order distortion terms of the $f_1 + f_2$ class situated 2.5 MHz above the channel edge are not rejected and since their level may increase at higher frequencies in the channel lineup (accompanied by a fortuitist decrease in the levels of CTB terms), one may wish to institute a new offset of an *additional* 1.25 MHz for digital channels above a certain cross-over frequency, where CTB levels are lower than CSO levels. Thus, proper channel offsets becomes a balancing act, since one cannot avoid both the CSO and CTB terms simultaneously.

5.2.3 Maintain head-end transmitter aggregate noise power at low values.

In Reference [1], this author shows the significant impact of noise aggregation from many transmitter sources at the headend and hubs. It is shown that under currently prevailing headend and HFC signal transport practices, mass deployment of certain classes of QAM transmitters and upconverters used (or proposed to be used) for VOD, may not scale well because of excessive noise accumulation from adjacent channel modulated distortion terms as well as from broadband noise. It is further shown in Reference [1] that as operators expand and carry 50 digital channels on their downstream lineup, the use of lower performance QAM

[8] For 256 QAM, the cable modem multicarrier loading test in the DOCSIS ATP [23] permits a test under which the total power of 30 dBmV is distributed across 5 carriers (one signal of interest, four other) not on image or adjacent channels. However, such test does not guarantee that any of the intermodulation components fall on the channel of interest. Furthermore, CSO and CTB from only 4 carriers produce much lower peak envelope power fluctuations than 80-channel source having the same total power.

transmitters can result in a 2.4 dB C/(N+I) loss at the subscriber tap compared to the levels that can be realized with transmitters that meet the DOCSIS downstream QAM transmission requirements as provided in Table 6-15 of its Radio Frequency Interface Specification [25].

5.2.4 Secure next generation transmitters and set-tops with specific improvements

From the deficiency list compiled in Section 5.1 above, we turn to two items:

5.2.4.1 Increased interleaver depth

As seen in our discussion above and as presented in some of the test results, the interleaver depth supported by most set-tops today is decidedly inadequate for the type of channels found in CATV. A further concern is the fact that even some transmitters are incapable of operating at the full depths provided in the standard. Figure 3 shows that an interleaver depth of at least 300 µs would be required to handle CSO and CTB transients. Table 2 shows the various depth options available under J-83 Annex B. It becomes obvious that the earlier one starts to deploy full depth interleaving ,the better. In any event, there is no reason why operators should continue to invest in transmitters that are not fully compliant with all the modes in Table 2.

I **(# of taps)**	***J*** **(increment)**	**Burst protection 64-QAM/256-QAM**	**Latency 64-QAM/256-QAM**	**Comments**
8	16	5.9 µs /4.1 µs	0.22 ms/0.15 ms	DOCSIS Requirement
16	8	12 µs /8.2 µs	0.48 ms/0.33 ms	
32	4	24 µs /16 µs	0.98 ms/0.68 ms	
64	2	47 µs /33 µs	2.0 ms/1.4 ms	
128	1	95 µs /66 µs	4.0 ms/2.8 ms	
128	2	190 µs /132 µs	8.0 ms/5.6 ms	Not supported by most digital video set-tops
128	3	285 µs /198 µs	12 ms/8.4 ms	
128	4	379 µs /264 µs	16 ms/11 ms	
128	5	474 µs /330 µs	20 ms/14 ms	
128	6	569 µs /396 µs	24 ms/17 ms	
128	7	664 µs /462 µs	28 ms/19 ms	
128	8	759 µs /528 µs	32 ms/22 ms	

Table 2. Interleaver Depth Settings available in J-83 Annex B. It is argued that the use of the maximum length associated with (I,J) = (128,8) would be a highly desirable upgrade for new transmitters and set-tops.

5.2.4.2 Improved Adaptive Equalizers

A particularly important feature of adaptive equalization systems employed in demodulators is their ability not only to equalize the channel frequency response by minimizing Inter-Symbol Interference (ISI), but also "equalizing" out interference signals. In the presence of discrete interference signals, many equalizer systems are able to automatically converge to a state that forms a sharp notch at the frequency of the interferer. The degree of rejection and the speed at which these equalizers converge becomes a key attribute of the demodulator. There is no doubt that demodulator chip makers are now looking at improved equalization techniques with a particular goal of improving the narrow-band rejection capability. The author envisions industry efforts to better characterize and standardize CSO/CTB rejection capability of demodulators so that these attributes become part of a qualification or certification process.

6 Summary and Conclusions

There is reason to review and go over a substantial rethinking of the required margins for reliable 256 QAM operations. It is argued that composite distortion terms having envelope peaks that can reach up to 16 dB above their average level leave very little room for further noise degradations at *any* level of the signal distribution. First-order statistics for CSO and CTB envelopes were simulated and specific degradations due various factors including shorter interleaver spans were analyzed. Finally various mitigation steps have been recommended in order to enable more successful rollout of 256 QAM.

7 Acknowledgement

The author wishes to thank Dr. Andre Basovich of Broadband Innovations, Inc. who programmed the simulations and ran them over many days.

8 References

1. Ron D. Katznelson, "Delivering on the 256 QAM Promise. - Adopting Aggregate Degradation Measures for QAM transmitters". Submitted to *Cable-Tec Expo® Proceedings*, SCTE, (June 2002).

2. Vitaly Germanov, "The Impact of CSO/CTB Distortion on BER Characteristics by Hybrid Multichannel Analog/QAM Transmission Systems". *IEEE Transactions on Broadcasting* **Vol 45**. No. 3. pp. 348-352, (September 1999).

3. 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 Technical Papers*, (1998).

4. Dean Stoneback, Robert Howald, Timothy Brophy and Oleh Sniezko, "Distortion Beat Characterization and the Impact on QAM BER Performance". *NCTA Technical Papers*, (1999) (Downloadable at http://www.gi.com/whitepaper/CTBNCTA99Final.pdf).

5. Robert Howald, "QAM Bulks Up Once Again: Modulation to the Power of Ten". Submitted to *Cable-Tec Expo® Proceedings*, SCTE, (2002).

6. *Digital Cable Network Interface Standard*. SCTE 40 2001 (Formerly DVS 313), (2001).

7. Ken Simons, "A mathematical Analysis of Distortion as it Occurs in CATV Amplifiers" Ch. V in *Technical Handbook for CATV Systems*. General Instrument, 3rd Edition (1968).

8. Ron D. Katznelson, "Optimal Signal Synthesis for Distortion Canceling Multicarrier Systems". U.S. Patent No. 5,125,100, (June 23, 1992).

9. W. R. Bennett, "Cross-Modulation Requirements on Multichannel Amplifiers Below Overload", *Bell System Technical Journal*, **Vol. 19**, pp. 587-605, (1940).

10. Matrix Technical Notes MTN-108, "Some Notes On Composite Second and Third Order Intermodulation Distortions", (Dec15, 1998). (downloadable at http://www.matrixtest.com/Literat/mtn108.htm)

11. Ron D. Katznelson, "Higher Order Statistics of Composite Distortion Components in Multicarrier Systems". In preparation (2002).

12. J. Dugundji, "Envelopes and Pre-Envelopes of Real Waveforms", *IRE Transactions on Information Theory*, **Vol. IT-4**, pp. 53-57, March (1958).

13. P. Mazur and E. Montroll, "Poincare Cycles, Ergodicity and Irreversibility in Assemblies of Coupled Harmonic Oscillators". *Journal of Mathematical Physics*, **Vol. 1**. No. 1. pp. 70-84. (1960). See Appendix III.

14. J. Davidse, "NTSC colour-television signals – evaluation of measurements", *Electronic & Radio Engineer*, pp. 416-419, November (1959).

15. Darryl Schick, "Characterization of Peak Factor Effects in HFC Systems Using Phase Controlled Carriers", *NCTA Technical Papers*. pp. 44-47. (1996).

16. A. Papoulis, *Probability, Random Variables and Stochastic Processes*, McGraw-Hill (1965).

17. Dojun Rhee and Robert H. Morelos-Zaragoza. "Error Performance Analysis of a Concatenated Coding Scheme with 64/256-QAM Trellis Coded Modulation for the North American Cable Modem Standard, *Proceedings of the 1998 IEEE International Symposium on Information Theory (ISIT '98*), p. 61, MIT, Cambridge, August 17-21, (1998).

18. European Telecommunications Standards Institute, "Digital Video Broadcasting (DVB) Framing Structure, Channel Coding and Modulation for Digital Terrestrial Television" *EN 300 744* (January 2001).

19. D. Raskin, D. Stoneback, J. Chrostowski and R. Menna, "Don't Get Clipped on the Information Highway".). *NCTA Technical Papers*. pp. 294-301. (1996).

20. Qun Shi and Robert S. Burroughs, "Hybrid Multichannel Analog/Digital CATV Transmission via Fiber Optic Link: Performance Limits and Trade-Offs", *NCTA Technical Papers*. pp. 182-193. (1994).

21. American National Standard, *Digital Video Transmission Standard for Cable Television,* ANSI/SCTE 07 2000 (Formerly SCTE DVS 031). (Downloadable at http://www.scte.org/standards/pdf/standardsavail/ANSI-SCTE%2007%202000.pdf).

22. ITU-T Recommendation J.83 Annex B. "*Digital Multi-Programme Systems for Television Sound and Data Services for Cable Distribution*". (April 1997).

23. Section 2.1.7 of "DOCSIS Acceptance Test Plan", TP-RFI-ATPv1.1-I02-010919. CableLabs, September 19, (2001). (Downloadable at http://www.cablemodem.com/specifications.html).

24. Ron D. Katznelson, "Common Architectures for digital and analog receivers" Presented at the EIA/NCTA Joint Engineering Committee. Washington D.C., Fall (1992).

25. "DOCSIS Radio Frequency Interface Specification", SP-RFIv2.0-I01-11231. December 31, (2001). (Downloadable at http://www.cablemodem.com/specifications.html).

***

Author contact:
Dr. Ron D. Katznelson, CTO
Broadband Innovations, Inc.,
3550 General Atomics Ct. Bldg. 15
San Diego CA, 92121

ron@broadbandinnovations.com
(858) 395-1440 Mobile
(858) 713-8505 Office

TECHNICAL ANALYSIS OF DOCSIS 2.0

Hal Roberts,
Benoit Legault, Mike Rude ADC

Abstract

The purpose of CableLabs® first Data Over Cable System Interface Specifications (DOCSIS) - DOCSIS 1.0 and DOCSIS 1.1, were to respectively enable residential data services and voice services over a single Internet Protocol (IP) cable infrastructure. The 1.0 specification defined the upstream and downstream physical and data link layers necessary to transmit over shared multiple-access cable IP networks. DOCSIS 1.0 specified the basic Quality of Service (QoS) features required to offer tiered services based on rate-limits, and was later enhanced to support minimum guaranteed rates. The DOCSIS 1.1 specification introduced support for constant bit rate services, which greatly enhanced the QoS feature set, and somewhat improved the robustness of the return path, which allowed twice the bandwidth, while providing full backward compatibility with the 1.0 specification.

In December 2001, CableLabs® released the first version of the DOCSIS 2.0 specification. The primary objective of DOCSIS 2.0 is to enhance upstream spectral efficiency, which requires additional robustness. This paper's objective is to investigate the new features introduced in DOCSIS 2.0, by closely examining its benefits to legacy 1.0 and 1.1 cable modems (CMs), and 2.0 CMs.

DOCSIS 2.0 achieves the goal of increasing upstream spectral efficiency and robustness by enhancing the 1.x (DOCSIS 1.0 and 1.1) TDMA modulation encoding method, renaming it Advanced-TDMA (A-TDMA), and by introducing a new upstream modulation encoding method known as S-CDMA. DOCSIS 2.0 requires that CMs and cable modem termination systems (CMTSs) support both A-TDMA and S-CDMA, thereby leaving the choice of enabling either or both methods on the DOCSIS channels to the operator. This paper analyses both encoding schemes to help the reader better understand the how they should be enabled in the network.

THE EVOLUTION TO DOCSIS 2.0

Introduction

This paper presents facts about the benefits of advanced time division multiple access (A-TDMA) and synchronized code division multiple access (S-CDMA) (as embodied in the DOCSIS 2.0 RFI), and the relative advantages of one vs. the other. It was written to help multiple system operators (MSOs) better understand these technologies so that decisions may be made about whether and how to deploy them.

Document Overview

This paper first describes some of the improvements that an A-TDMA product will have, compared to the existing DOCSIS 1.1 product. All of the improvements described apply to legacy DOCSIS 1.X cable modems. Many of these improvements go beyond DOCSIS 2.0 requirements.

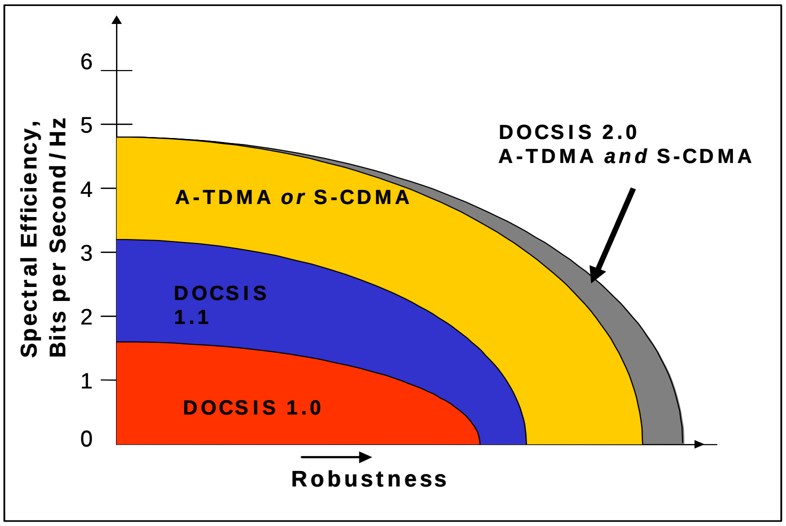

Figure 1: Evolution of DOCSIS robustness & spectral efficiency.

- Ingress Cancellation
- Improved Receive equalization
- Improved Burst acquisition
- Impulse Noise mitigation (improved FEC)

The paper then goes on to compare the two modulations (A-TDMA and S-CDMA) in the following areas:

- Dynamic Range and Timing Sensitivity
- Impulse Noise and Ingress
- Backward Compatibility and Interoperability of A-TDMA and S-CDMA
- Scheduling Efficiencies

Lastly, a conclusion is presented as to how an operator may use these facts in evolving towards DOCSIS 2.0.

Figure 1 shows a graphical illustration of the evolution of robustness and spectral efficiency. Note that spectral efficiency is quantifiable in terms of bits per second per Hertz. Robustness is essential for *allowing* operation at higher bits per second but cannot be quantified without a fully defined channel model and is therefore more subjective. The intent of the illustration is to show that there is a large step up in robustness by evolving to S-CDMA *or* A-TDMA and a smaller incremental benefit results from switching between modes according to plant conditions.

THE EVOLUTION OF ADVANCED PHY

Work on advanced PHY techniques began in March of 1998 under the auspices of IEEE 802.14a and culminated with the release of the DOCSIS 2.0 specification by CableLabs®. The fundamental improvements that were desired by MSOs from an advanced standard were: 1) increased capacity, 2) increased robustness to RF upstream impairments, and 3) no degradation to existing DOCSIS deployed systems. Increased robustness was desired both for improving the robustness of existing networks, but also to make possible the higher orders of modulation and symbol rates of advanced PHY, since the larger constellations and rates require either cleaner RF channels or greater robustness to be reliably deployed. Some of the robustness enhancements of advanced PHY are laid out in the specification, such as increased FEC, while other robustness enhancements are implemented in the receivers such as ingress cancellation, and thus are proprietary in nature, but also apply to existing DOCSIS networks. Hence, these improvements can also be used to increase the reliability of medium bandwidth channels, such as the 16-QAM, 3.2MHz channels, available in existing legacy 1.X DOCSIS modems, as long as an advanced PHY CMTS is deployed which supports the new robustness features.

On the other hand, the capacity increase provided by the larger constellations and rates requires that *both* the CMTS and the cable modem (CM) have Advanced PHY features. It should be noted that even partial Advanced PHY deployments improve the entire network capacity since the Advanced PHY modems are using less resources per modem for the same provisioned level of service. Lastly, the requirement for Advanced PHY to not degrade existing services means that the new technology must integrate seamlessly with existing networks and not create additional overhead or other bandwidth consumption that reduces network capacity.

A-TDMA 'Single Ended' Features

While many of the features of DOCSIS 2.0 (such as S-CDMA) do not provide any improvements to legacy modems, there are four robustness improvements that are 'single ended', i.e. may be obtained only by upgrading the CMTS and apply to all DOCSIS cable modems. These provide benefits long before DOCSIS 2.0 cable modems will be ubiquitous simply by upgrading the CMTS.

- Ingress Cancellation Filter
- Improved Receive Equalization
- Improved Burst Acquisition
- Improved Error Correction for Impulses

These improvements will not only improve reliability and capacity in existing upstream channels, but will also make new RF spectrum available to existing DOCSIS 1.x modems, thereby greatly expanding the upstream capacity available to *existing DOCSIS systems*. Each robustness improvement is detailed below.

Ingress Cancellation Filter

The most significant improvement in robustness, ingress cancellation, is actually not part of the DOCSIS 2.0 specification, but should be found in some form in all Advanced PHY CMTSs to support the higher order modulations. This is a digital filter that adaptively responds to narrow band and wide band ingress or common path distortion (CPD) and filters it out (ref. Figure 2).

The ingress cancellation filter (ICF) determines the nature of the upstream ingress by analyzing the channel in-between packet bursts. Cancellation coefficients are computed by a digital signal processor (DSP). These coefficients are updated up to 200 times per second to handle time-varying ingress. The ICF is integrated with the receive equalizer to optimize overall channel response.

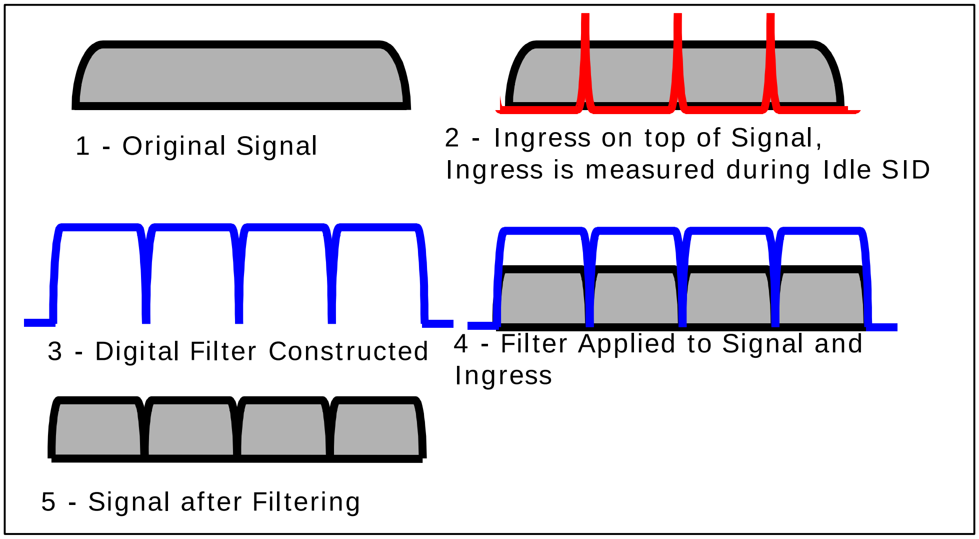
1 - Original Signal
2 - Ingress on top of Signal,
Ingress is measured during Idle SID
3 - Digital Filter Constructed
4 - Filter Applied to Signal and Ingress
5 - Signal after Filtering

Figure 2: The Process of the ICF filtering out three CW tones.

The ICF is able to cancel or mitigate the effects of narrow band and wide band interference as well as CPD (a form of wide band interference due to non-linear components in the plant). An example of this is shown in Figure 3. An example of multiple carrier wave (CW) ingress that can be effectively cancelled by a commercially available A-TDMA burst receiver is shown in Figure 4 along with the 16-QAM constellation after filtering (Figure 5.) Table 1 below contains test results from a variety of ingress types and levels.

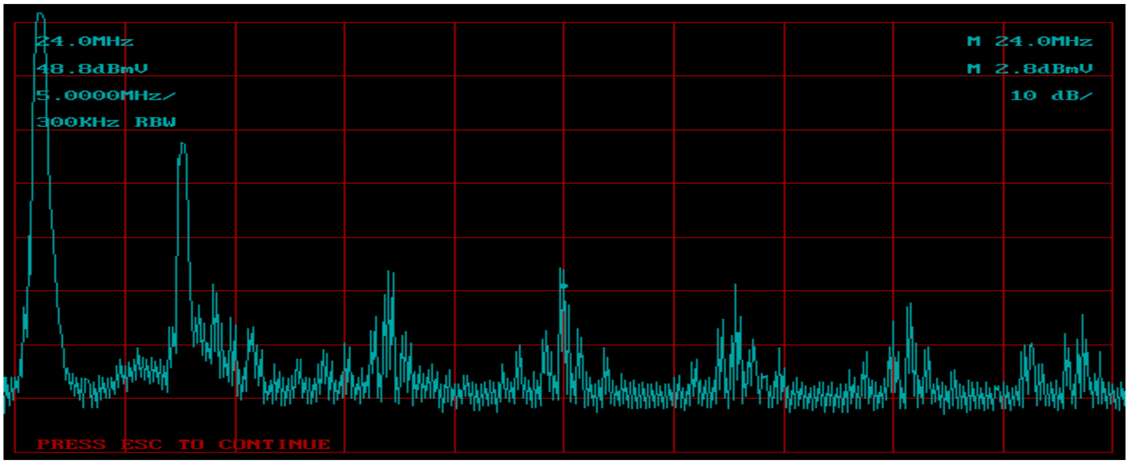

Figure 3: CPD Ingress Example.

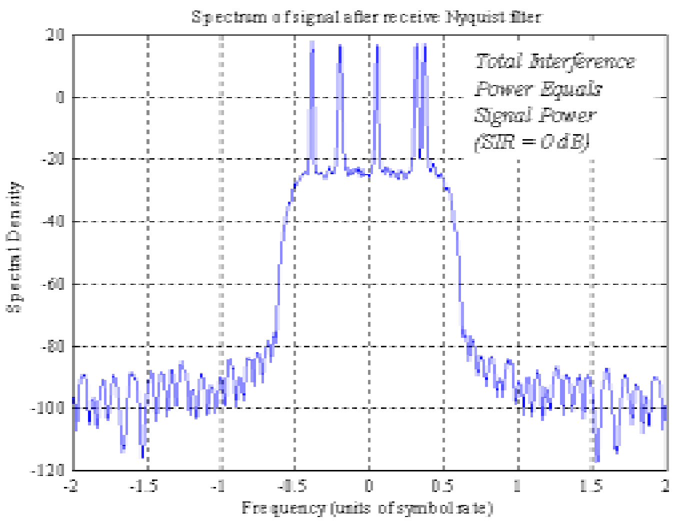

Figure 4: Signal Spectrum with 5 ingressors.

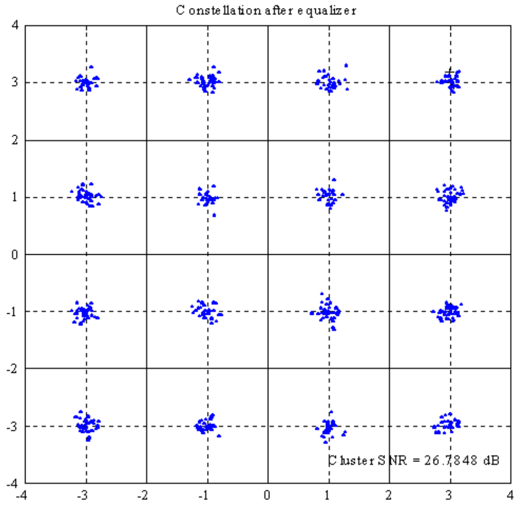

Figure 5: Constellation after ingress removal.

Impairment Type	CIR without ICF	CIR with ICF
Narrow band ingress	20 dB	-10 dB
5 narrow band ingressors	20 dB	0 dB
2 wideband ingressors	20 dB	9 dB
CPD	20 dB	5 dB

Table 1: Carrier/ingress ratio required for 16-QAM, with and without ICF.

Improved Receive Equalization

Equalization is used to mitigate the effects of frequency dependent attenuation, delay and multipath on both the upstream and downstream in DOCSIS. In the downstream equalization is performed at the cable modem receiver. Since there is a continuous downstream signal that is coming from a single source (the CMTS) the equalization can be accomplished completely within the cable modem receiver.

DOCSIS 1.1 and 2.0 implement *upstream* equalization using pre-distortion rather than relying on *receive* equalization. The reason for this is that a large number of preamble symbols are necessary in order to train the equalizer in receive-only equalization[1]. Since each CM burst will experience different upstream distortions, if each data burst was preceded by a preamble long enough to train the equalizer, then much of the upstream bandwidth would be wasted, especially for small packets.

Pre-distortion Equalization

On the upstream, the CMTS sees sequential bursts coming from potentially thousands of cable modems. The hybrid fiber coax (HFC) plant (see Figure 6) distorts each burst differently since the bursts travel through different paths and plant elements. To equalize the burst, the receiver must use different equalizer coefficients. It is impractical to store all of these coefficients at the CMTS and load them burst-by-burst into the equalizer. Instead the coefficients are sent to the CMs so that they may *pre-distort* the upstream bursts (see Figure 7). Once they arrive at the CMTS, the bursts will be undistorted, as the HFC plant will reverse the effect of the pre-distortion. Therefore the primary purpose of the CMTS receiver equalizer is to *measure* the distortion and calculate the necessary coefficients for the cable modem based on ranging bursts. In addition, the receive equalizer is still active at the CMTS on data bursts and can help mitigate transient channel distortion.

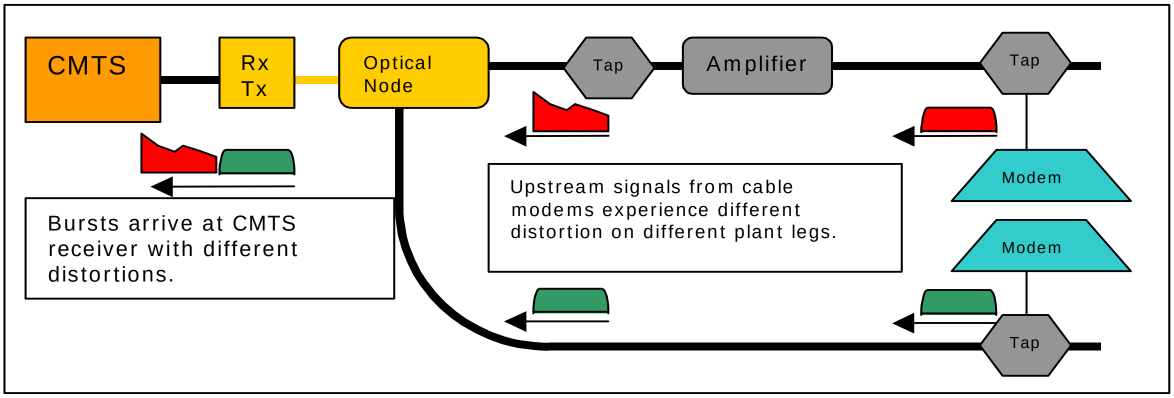
CMTS
Rx
Tx
Optical
Node
Tap
Amplifier
Tap
Modem
Modem
Tap
Bursts arrive at CMTS receiver with different distortions.
Upstream signals from cable modems experience different distortion on different plant legs.

Figure 6: No Pre-Distortion. Bursts must be equalized at the receiver by using a training pattern of adequate length

[1] In the downstream equalization is accomplished during initial synchronization and then changes slowly thereafter. Therefore the overhead for equalizer training is insignificant.

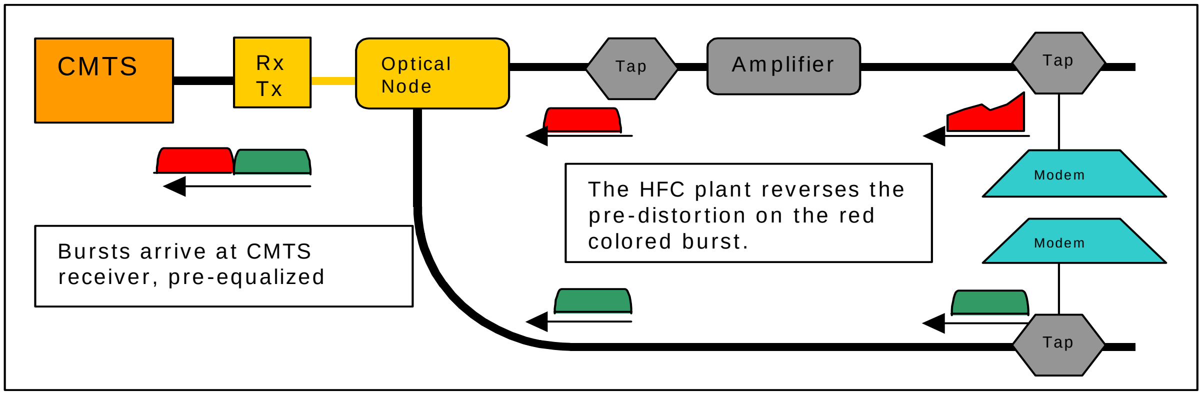
CMTS
Rx
Tx
Optical
Node
Tap
Amplifier
Tap
Modem
Modem
Tap
Bursts arrive at CMTS receiver, pre-equalized
The HFC plant reverses the pre-distortion on the red colored burst.

Figure 7: No Pre-Distortion. Bursts must be equalized at the receiver by using a training pattern of adequate length

"24-Tap" Equalization

In DOCSIS 1.1, the pre-distortion is defined for 8 taps. The A-TDMA CMTS has a 24 tap receive equalizer. It would seem that without an A-TDMA CM with 24 tap pre-distortion, that the receiver 24 tap equalizer is useless. In fact, 24 tap receive equalization may be used on a burst by burst basis, due to the improvements that have been made on the A-TDMA CMTS burst acquisition (see section on burst acquisition below). Although maximum pre-equalization is enabled when both the CMTS and CM have a matching number of taps, a higher order receive equalizer will enhance performance in a single ended fashion.

Consequently, an A-TDMA receive equalizer has the capability of compensating for multipath that is *4 times the duration*[2] of the multipath that can be handled by DOCSIS 1.1 CMTSs (see Figure 8).

Improved Burst Acquisition

The new A-TDMA burst receiver has a greatly improved burst acquisition capability. This was necessitated by higher order constellations, which requires an increased precision in the estimation of acquisition parameters. Increased precision acquisition may be accomplished by long preambles at the expense of efficiency. Instead the A-TDMA burst receiver has a new robust method of acquisition, which is accomplished in a minimum number of symbols. This applies to low order modulation, such as QPSK and 16-QAM, allowing acquisition in the presence of impulse noise and shorter preambles. In addition, the equalizer will train on the entire preamble, which allows the receiver-only equalization described in the previous section (see Table 2 below for a summary).

[2] The burst receiver has the main tap offset from the center at tap 8, which allows 16 trailing taps vs. 4 trailing taps with the standard burst receiver, i.e. 4 times the number of taps.

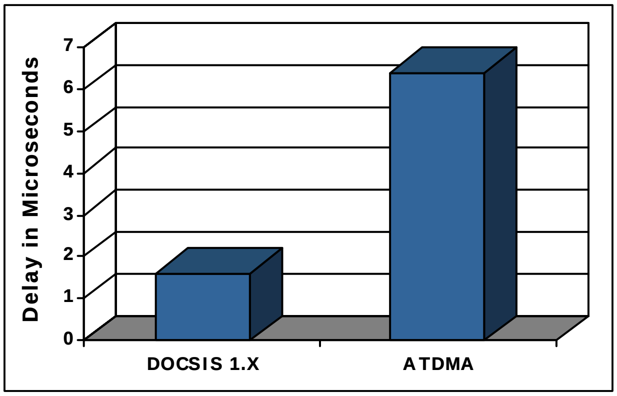

Figure 8: Enhanced ability to tolerate long delay multipath.

STANDARD PHY BURST ACQUISTION

- Standard burst receivers acquired signal parameters off of sections of the preamble in *series*. If any section of the preamble is hit by impulses, *acquisition fails*.
- Equalizer training was done *after* the preamble. If multipath is bad, *acquisition fails*.
- In short, acquisition was the weakest link- not the data forward error correction (FEC).

ADVANCED PHY BURST ACQUISTION

- Carrier and timing lock, power estimates, equalizer training and constellation phase lock are all done simultaneously. This allows shorter preambles (20 symbols) and/or robust acquisition with impairments.
- Reduction in the implementation loss to a fraction of a dB from theoretical, which means that legacy modems will be able to operate in higher additive white Gaussian noise (AWGN) noise levels than previously possible. Depending on the constellation size, symbol rate, and packet error rate being measured, up to 2.3dB improvement in AWGN performance is available with new A-TDMA burst receiver.

Table 2: Standard vs. Advanced Burst Acquisition.

Improved Forward Error Correction for Impulse Noise

DOCSIS 1.X will allow the correction of 10 errored bytes per Reed Solomon (RS) block (T=10). DOCSIS 2.0 allows correction of 16 bytes per Reed Solomon block (T=16). To obtain T=16 performance requires that a DOCSIS 2.0 cable modem is used in conjunction with a DOCSIS 2.0 CMTS.

However, the A-TDMA receiver has a new capability that offers a single ended improvement to FEC. This capability is called *erasure* correction. Erasure correction is possible when the location of errors within the Reed Solomon block is known.[1] Using erasure correction, up to 20 bytes with errors may be corrected for DOCSIS 1.X modems (up to twice as much correction power). Importantly, performance gains due to erasure correction *do not require Advanced PHY cable modems.* The technique works most effectively with impulse or burst noise where the location of errors can be inferred from detection of the impulse event at the demodulator.

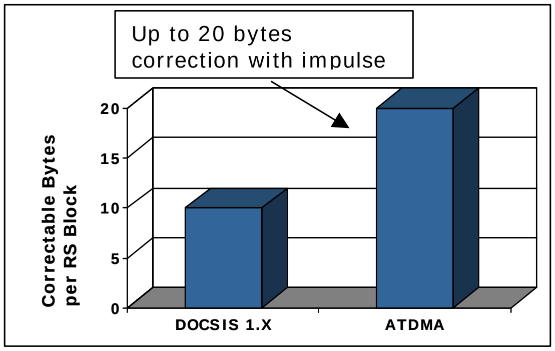

Figure 9: Erasure correction improves burst noise performance.

TECHNICAL COMPARISON BETWEEN S-CDMA AND A-TDMA

When examining modulation technologies, it is important to realize that there is not any technology that can exceed theoretical maximums such as the Nyquist and Shannon limits. In addition, if the technologies have been architected and implemented well, there is a tendency for the performance of each approach to converge towards these theoretical limits. In particular, it has been shown that if all other system parameters and coding are equal, then all modulation technologies will have identical performance in AWGN.

Early History of the Advanced PHY Standard

Some perspective may be gained by a short discussion of the early history of the Advanced PHY standard that culminated in DOCSIS 2.0.

As stated earlier, in March of 1998, IEEE 802.14a began work on improving the DOCSIS upstream modulation for increased robustness and bandwidth. There was early agreement that the DOCSIS downstream[3] performance was adequate and should remain untouched. At that time, the major battle was between selection of the S-CDMA proposal by Terayon and the variable constellation orthogonal frequency division multiplexing (VCOFDM) proposed by Ultracom (A-TDMA was assumed to be included). The battle hinged, in part, on the disadvantage of S-CDMA requiring tight timing requirements vs. the disadvantage of VCOFDM requiring dynamic constellation adjustments. Other factors, such as Terayon having a deployed HFC S-CDMA system vs. Ultracom's mostly theoretical proposal tipped the scales in favor of S-CDMA.

Technologically agnostic members of the committee demonstrated that all three proposals could provide roughly equivalent performance with each having advantages under specific operating conditions [2]. The same equivalence of performance existed between the Advanced TDMA proposal (by Broadcom and Texas Instruments) when compared to the modulation approaches of S-CDMA and VCOFDM. The main advantage of A-TDMA was that it was an *incremental* enhancement over the existing DOCSIS TDMA approach. Despite the similarity of performance between A-TDMA and S-CDMA, there are some advantages and disadvantages that result, partly because of the tradeoffs made in the specifications details.

S-CDMA ADVANTAGES

The following section assumes some knowledge of how S-CDMA works as specified in the DOCSIS 2.0 RFI. The RFI specification may be found at:

[3] Based on ITU-T J.83, Digital multi-programme systems for television, sound and data services for cable distribution.

http://www.cablemodem.com/specifications.html.

Impulse Noise

The long duration symbols of S-CDMA provide robustness in the presence of impulse noise. Long duration symbols lead to a requirement that multiple symbols be simultaneously sent to maintain aggregate bandwidth. As will be seen, this fact leads to a *crossover point* where S-CDMA outperforms TDMA on one side of this crossover and TDMA outperforms S-CDMA on the other side. The underlying mechanism of the effect is illustrated graphically in Figure 10 and Figure 11. In a sense, the S-CDMA code space is "overloaded" with a high power impulse and transmission must move to a lower order

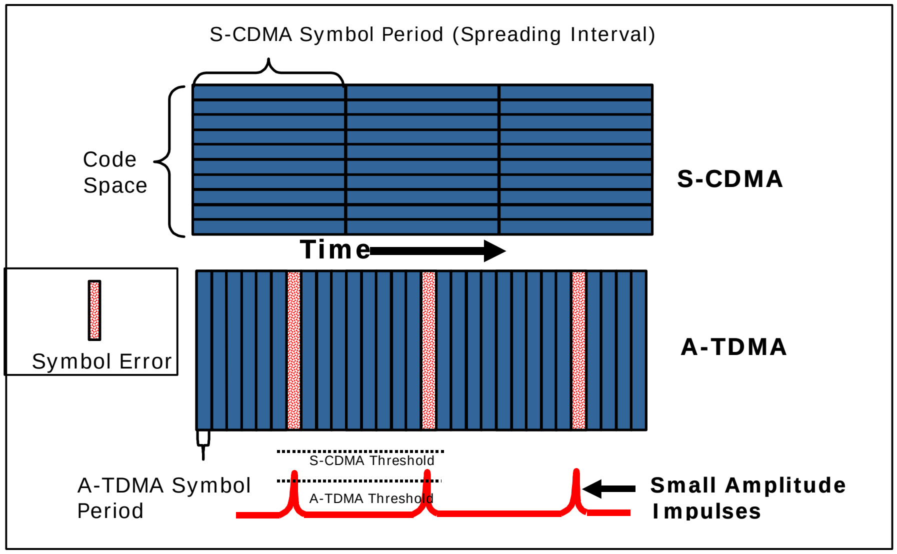

Figure 10: S-CDMA is robust vs. short duration and medium level impulses.

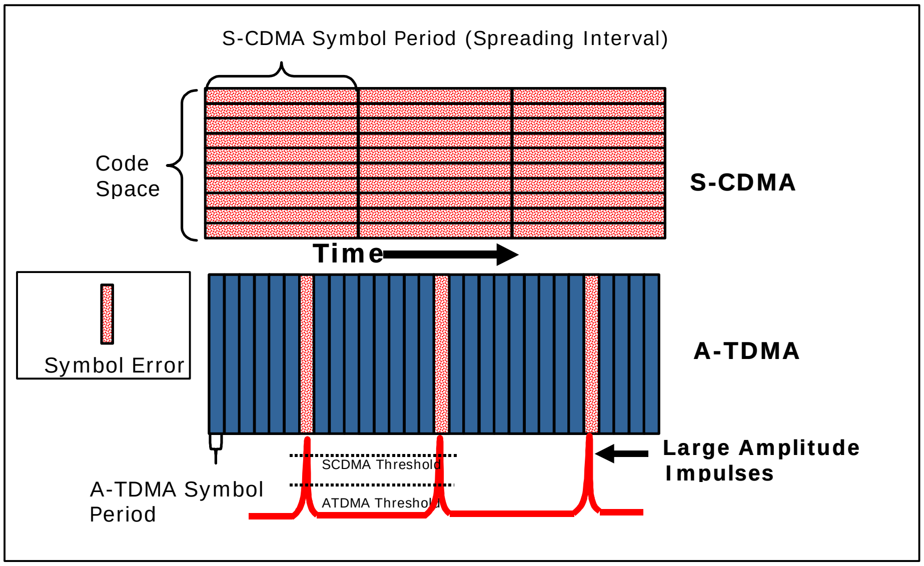

Figure 11: TDMA is robust against high power impulses.

constellation (e.g. QPSK), which will reduce throughput. A-TDMA, on the other hand, will experience errors at lower impulse power. This crossover is shown in Figure 12.

S-CDMA excels in short duration (around 1us) medium amplitude, high repetition rate impulses, and with small packets. S-CDMA tolerates these impulses (as shown in the Figure 10) due to the long duration of the S-CDMA symbols that spread the short impulse energy over the whole symbol[4]. On the other hand, spreading does not help if the impulse energy is high enough to corrupt the symbol (as shown in the Figure 11), despite the spreading effect. In general, most or all of the symbols that are simultaneously sent will be simultaneously corrupted if impulses are high in power. At that point, TDMA has an advantage because the number of corrupted symbols is limited as only one symbol per unit time.

Long Impulses

S-CDMA also has an advantage against long duration (greater than 5-10us), large impulses and with small packets. For long impulses, an important factor in S-CDMA impulse noise performance is the spreading interval factor 'K' and the number of codes per minislot (CPMS). If K is large and the CPMS are small, the number of codes simultaneously used for any burst may be small. Another way of looking at this is that the burst is stretched out as far in time as possible, making the impulse duration a small fraction of the burst duration. Since a large impulse will destroy almost all codes that are simultaneously sent, then it is best to send as few codes simultaneously as possible (2-4 codes per minislot). The few symbols that are corrupted can easily be corrected by FEC. As an example[5], see Figure 13. Assuming that the channel is operating in 64-QAM at 2.56Msps, if the S-CDMA mode is at K=32 (the maximum allowed) spreading intervals per frame and the number of codes per minislot is two, a 64 byte packet may be transmitted in two minislots (excluding preamble). The duration of the frame is 0.39us/chip x 128 x 32 = 1.6ms. The maximum allowable

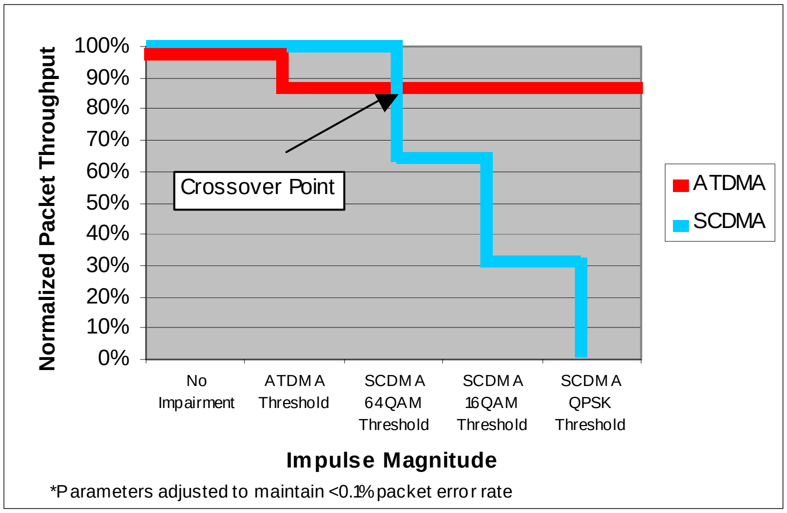

*Parameters adjusted to maintain <0.1% packet error rate

Figure 12: Crossover of A-TDMA & S-CDMA for Impulse Noise Rejection.

[4] As the impulse approaches the symbol duration this spreading advantage declines. Once the impulse is as long as the symbol period there is no spreading.

[5] In this example the MAC overhead and preamble is ignored. Interleaving is off since it won't improve performance.

duration of a burst with the maximum Reed Solomon correction of T=16 is 250 us (5 spreading intervals). The sensitivity of S-CDMA burst performance is dependent on the codes per minislot and the K factor (ref. graphic of S-CDMA burst performance vs. K in Figure 15).

On the other hand, with A-TDMA, a small packet is sent in a short amount of time (see Figure 14). For a channel operating in 64-QAM at 2.56Msps, a 64 byte packet is sent (excluding preamble) in 33us. With the maximum Reed Solomon correction of T=16, the impulse will corrupt 16 bytes and break the error correction if it is 8.3us in length or greater. Therefore, for the case with S-CDMA optimized for maximum tolerance to burst noise, S-CDMA can handle bursts that are **30** (250/8.3) times longer than A-TDMA. However, there is an impact to dynamic range as a result of operation in the most burst tolerant SCDMA mode, as will be seen below.

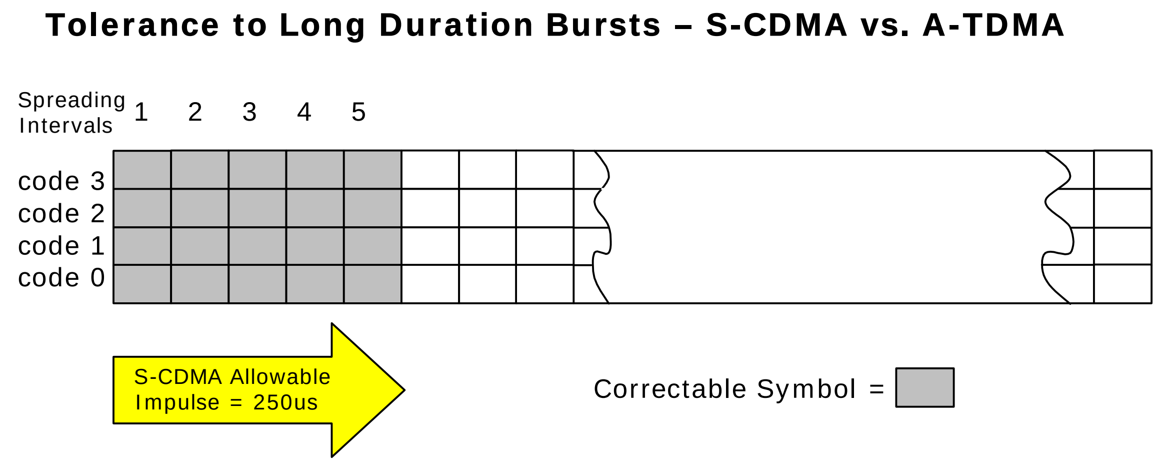
Tolerance to Long Duration Bursts – S-CDMA vs. A-TDMA
Spreading Intervals
1
2
3
4
5
code 3
code 2
code 1
code 0
S-CDMA Allowable Impulse = 250us
Correctable Symbol =

Figure 13: S-CDMA will tolerate a 250us[6] burst when configured for maximum duration burst handling, i.e. T=16, K=32, CPMS=2.

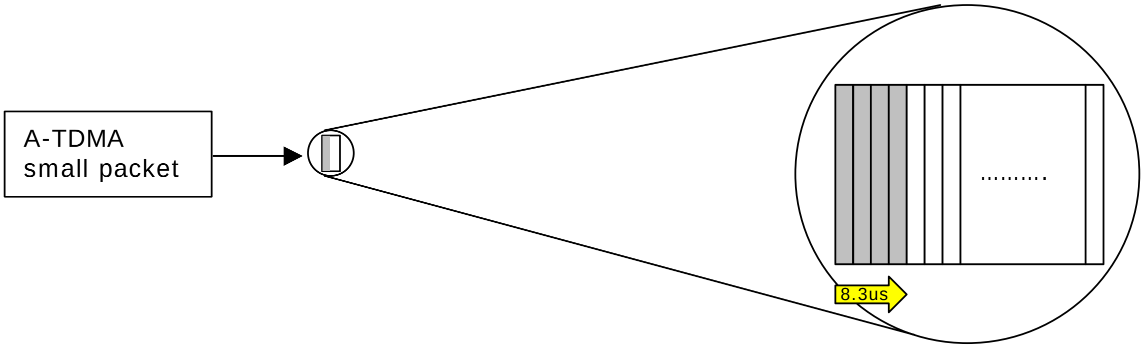

Figure 14: A-TDMA will only tolerate an 8.3us burst and will not suffer any loss of system dynamic range.

[6] If the MAC and preamble overheads are included then only a 150us impulse may be handled.

Short Preamble

S-CDMA is a synchronized system. As a result, packets arrive at the CMTS from multiple CMs pre-synchronized to the CMTS receiver clock. All that remains for the receiver is to obtain gain estimates for optimal slicer operation. Therefore, the S-CDMA preamble may only be a few symbols in length vs. 20 for A-TDMA. This matters in short packet transmission and may account for approximately a 30% reduction[7] in bandwidth for short packets.

Note that this behavior is not dependent upon whether the system is using S-CDMA or A-TDMA; it is dependent on the use of upstream synchronization. A-TDMA may also operate in a synchronized fashion, however this mode of operation is not included in the DOCSIS 2.0 specification.

A-TDMA ADVANTAGES

Short Duration High Amplitude Impulses.

Short duration high amplitude impulses are handled better by A-TDMA, as once the impulse is large enough, it will corrupt all S-CDMA codes in a spreading interval (Figure 11). At a high enough repetition rate, error correction will not be able to compensate for these errors, regardless of interleaving. A-TDMA can handle 10 to 100 times the repetition rate since only a single symbol is sent per unit time, causing only one symbol can be corrupted per unit time.

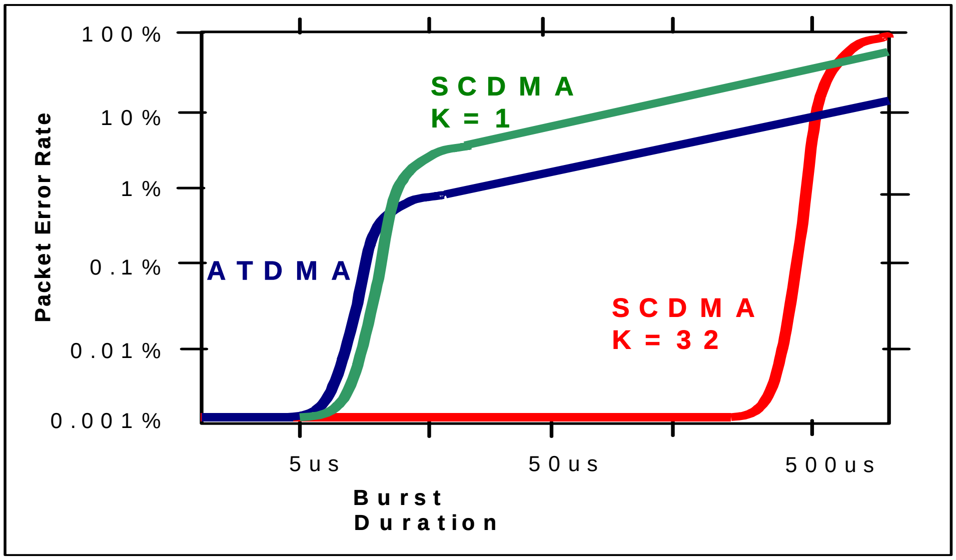

Figure 15: Long Burst Performance of Small Packets (S-CDMA K=32 vs. S-CDMA K=1 vs. A-TDMA)

[7] Broadcom estimate, November 2001.

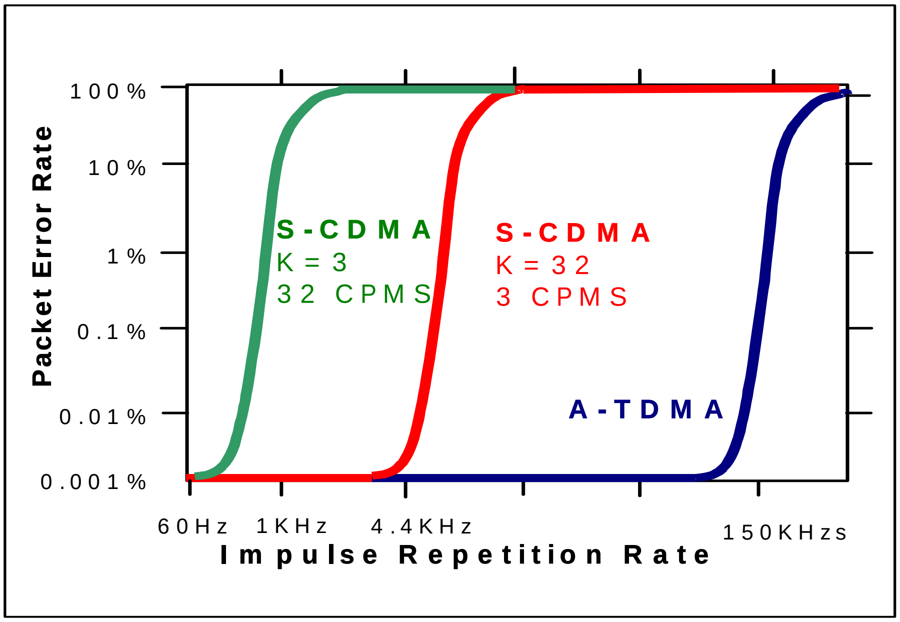

Figure 16: High amplitude impulses are handled effectively by A-TDMA.

The way this works is that after decorrelation at the S-CDMA demodulator the impulse energy is spread over all 128 CDMA 'chips' which make up a single CDMA 'symbol', so that the impulse energy is 1/128th as large. This reduced energy degrades the SNR of each of the 128 simultaneously demodulated orthogonal sequences that make up all the symbols. This works well at mitigating the impulse until the amplitude becomes so great that even 1/128th of the impulse degrades the symbol SNR at the demodulator such that the signal cannot be recovered. All 128 sequences, or 'symbols', suffer the same fate and one large impulse destroys them all.

As an example consider the same channel as in the previous example, 64-QAM at 2.56Msps and T=16. The packet size will not be critical in this case. S-CDMA will be set to maximum impulse noise tolerance settings of 2 codes per minislot and K=32. Assuming 1us impulses at a level high enough to corrupt both A-TDMA and S-CDMA symbols, at low repetition rates both S-CDMA and A-TDMA FEC will be able to correct the errors. As the impulse rate increases, the S-CDMA system will experience errors when 7 impulses occur within one frame of 1.6ms, i.e. at 4.4KHz. Alternatively, A-TDMA will start to experience errors at 150KHz (Figure 16.)

TDMA DYNAMIC RANGE

S-CDMA Low Power Limitation

As we have seen, operation in S-CDMA mode has better impulse noise handling if the number of codes per minislot is low and the K factor is high. The lowest number of codes per minislot allowed is 2. Therefore, this is also the best for impulse noise immunity. However the DOCSIS 2.0 specification requires that the dynamic range of the S-CDMA modem be from 8dBmV to 53dBmV *independent of modulation order and the number of codes per minislot.* This is a critical specification that results in a low power limit that varies according to the number of codes per minislot. The DOCSIS 2.0 specification defines the minimum power with all codes active as:

Minimum Upstream Power$_{2cpms}$ = 8dBmV + 10*log$_{10}$(128/2) = 26dBmV

At the minimum power limit, this results in a reduction of dynamic range by *18dB*.

S-CDMA High Power Case

Due to the high peak to average nature of S-CDMA signals, there is a required power backoff from maximum power. Note this is similar to the power backoff that was needed for 16-QAM and higher order QAM in DOCSIS 1.X. However S-CDMA requires the same power backoff for all QAM modes including QPSK (also called 4-QAM). Therefore, a S-CDMA modem that is operating in QPSK must operate at a maximum power level of 53dBmV. A TDMA modem operating in QPSK may operate at 58dBmV.

S-CDMA Dynamic Range vs. TDMA Dynamic Range (see dynamic range graph below)

The analysis below will assume the case of 2 codes per minislot for S-CDMA. This is the most robust case for impulse noise, but the most limited in dynamic range. Increased dynamic range may be had at the expense of S-CDMA's advantage in impulse noise resistance. This trade-off will be examined later.

There is a detailed analysis of the required dynamic range in the return path (upstream) in the book, "Broadband Return Systems for HFC CATV Networks" by Donald Raskin and Dean Stoneback. The range is extremely large (49dB) until some techniques are applied to reduce the range. One of the techniques is feeder equalization, which is not universally applied to HFC systems. If feeder equalization is applied (among other techniques), then the required dynamic range is **34dB**. Most of the remaining variance is found in the customer in-house wiring and splitters, which is not easily subject to control by the MSO. It is worthwhile to compare this range requirement to the actual dynamic range of cable modems in DOCSIS 2.0:

- The total dynamic range of TDMA modems is 58dBmV-8dBmV = **50dB**.
- The total dynamic range for S-CDMA modems is 53dBmV-26dBmV=**27dB**.

Note that in DOCSIS 1.X the dynamic range was made large to accommodate a wide range of channel bandwidths, from 2560 ksym/sec to 160 ksym/sec. The reason channel bandwidth matters is that if *uniform spectral density* is desired in the upstream and if multiple bandwidth channels are to coexist, the low bandwidth channels must operate at lower powers. For each doubling of bandwidth, an effective loss of 3dB is experienced for the *system* dynamic range. Over the channel bandwidth range of DOCSIS 1.X, there are four doublings in bandwidth, resulting in 3x4=12dB effective reduction in system bandwidth. Since A-TDMA allows operation at 5120 ksym/sec, there are 5 bandwidth doublings, equating to a 15dB reduction. Therefore the *effective* dynamic range (vs. total dynamic range) of TDMA is 50dB-15dB = 35dB, which is 1dB greater than the Raskin and Stoneback recommendations.

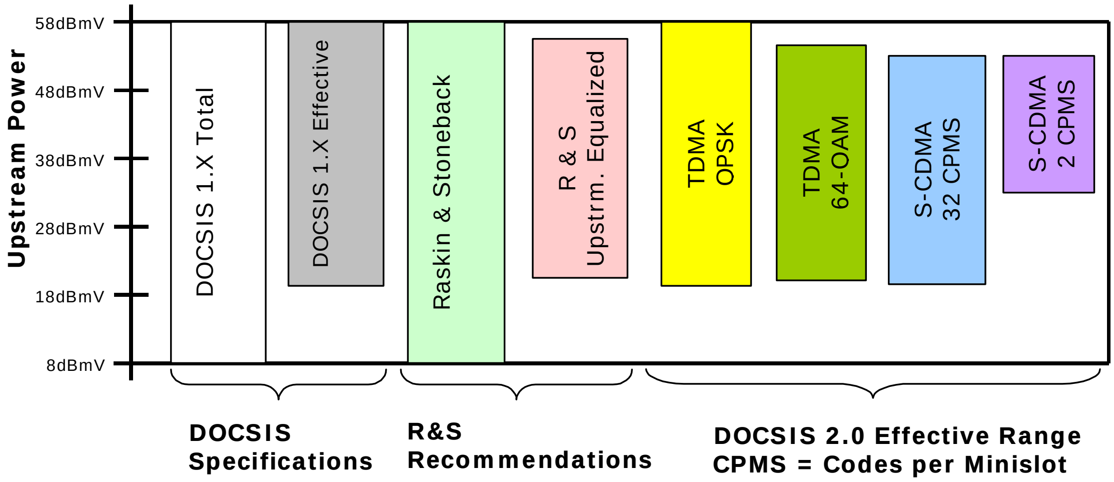

Figure 17: Dynamic Range Implications of S-CDMA

It was understood during the creation of the DOCSIS 2.0 specification, that S-CDMA dynamic range is limited. Therefore operation below 1280 ksym/sec is prohibited when in S-CDMA mode. The argument is that the Advanced PHY benefits are needed only for wider band channels. Therefore the *effective* dynamic range of S-CDMA is reduced from the *total* dynamic range by only 6dB, since only two doublings of bandwidth are allowed. The *effective* dynamic range of S-CDMA is 27dB-6dB=21dB.

Increasing the number of codes per minislot above 2 will increase the dynamic range on the low power end of S-CDMA modems. With each doubling of codes per minislot, the low power range is increased by 3dB. However, this improvement is obtained at the expense of a tradeoff in robustness against impulse noise. Because the maximum number of codes is 32, which is 4 doublings, the low-end dynamic range is improved by 12dB. The result of using the maximum codes per minislot is an effective dynamic range of 21dB + 12dB = **33dB**.

<u>SCDMA Differential Code Power and Intercode Interference</u>

With S-CDMA it is important that all modems transmit at the power required by the CMTS. Differential code powers will cause degradation in the codes with low power due to minor timing differences in the codes with high power (relative to the CMTS). The S-CDMA timing budget assumes equal power in all the codes. Therefore if the dynamic range limitations cause different modems to have unequal code power at the CMTS, the timing induced non-orthogonality will cause the lower power codes to experience intercode interference from the higher power codes.

<u>Effect of Dynamic Range in a Cable Plant</u>

A cable modem should transmit at a high enough power as required by the CMTS, otherwise the signal will end up closer to the noise floor. In this case a single modem may force the entire plant to switch to a lower QAM level in order for the 'challenged' modem to operate. In the worst case the CM will be out of the CMTS receive range. At the other extreme, the CM is unable to turn down it's transmitter to the CMTS receive level. This will cause too much RF power directed to the upstream laser and may also cause clipping, depending on the laser operating margin. In the extreme, the power will be outside of the CMTS receive range.

Operators should evaluate their plant characteristics and determine the optimal operation as a tradeoff between length and amplitude of impulse noise and dynamic range requirements of the cable plant.

One proposed solution is to switch all modems operating in S-CDMA mode to A-TDMA mode when the dynamic range is needed. This requires the CMTS system to obtain the cable modem RF power levels, monitor these levels and switch individual modems to the alternate logical channel for A-TDMA when needed.

Quantization Noise Funneling

Noise funneling in SCDMA has been addressed by the DOCSIS 2.0 specification, but it is worth noting. SCDMA allows a variation in the number of simultaneously transmitting modems from 1 to 64. (TDMA only allows a single modem transmitting at one time). As a result, in the 64 simultaneously transmitting case, there is additional upstream noise caused by the large number of modems transmitting 2 codes apiece. The most fundamental source of noise is due to the quantization noise added by the DAC (digital to analog converter). This noise is created by the fact that a DAC has a finite number of bits of resolution. The LSB (least significant bit) will cause a change in the output signal that is a step change with an error compared to the 'ideal' signal. This noise is usually identical to AWGN and can be treated as such. In DOCSIS 2.0 the requirements on the DAC were effectively tightened by 2 to 3 bits of resolution to handle this effect. Nonetheless, operating in A-TDMA mode or in S-CDMA mode with a high number of codes per minislot will lower the funneling noise and decrease the implementation loss[8].

Statistical Multiplexing Advantages

It is possible to operate all modems, from 1.0 to 2.0 in TDMA mode. It is not possible to operate 1.X modems in S-CDMA mode. Therefore, if any 1.x legacy modems coexist on the same channel with modems operating in S-CDMA mode, the channel must operate in 'dual' mode. This requires a minimum of two 'logical'[9] channels. Losses in statistical multiplexing efficiency are experienced when a channel is split into sub-channels. This effect may be reduced but not eliminated by intelligent scheduling. In addition, dual MAP and upstream channel descriptor (UCD) sets must be sent and upstream contention regions must be segregated.

TDMA Relaxed Timing Requirements (ref. Appendix B)

SCDMA requires accurate timing due to the upstream and downstream synchronization requirements. Timing variations can be caused by temperature shifts and possible wind loading effects in the plant. S-CDMA requires no more than a 2ns timing error. Station ranging must therefore be used to adjust timing before a 2ns error can build up. Appendix B discusses the timing changes that may exist in HFC systems and how they may impact

[8] Quantization noise effects combined with limited dynamic range may make low numbers of minislots in SCDMA a less desirable mode.

[9] Logical channels were created in DOCSIS 2.0 to allow coexistence between different modulation modes.

the frequency of station maintenance in S-CDMA.

Sources of HFC Impairments

It is instructive to understand the sources of the various upstream impairments described above, and also to understand how A-TDMA and S-CDMA will handle them [3].

Impulse Noise (time varying noise):

- Long Duration (~10ms) Large Impulses (0dBc) - One source of this type of noise is Impulse Pay-per-View Polling. Older upstream signaling from set-top boxes may be unbalanced. These devices were not designed with automatic gain control (AGC) and were usually set at high levels to ensure the signals were above other noise sources. This high signal level causes cross-compression of the upstream lasers and causes high magnitude impulse noise over all frequencies simultaneously (with lower frequencies degraded more than higher frequencies).
- Medium Duration (~100us) - May be caused by load switching using mechanical contacts. Another source is a loose F-connector making intermittent contact, usually due to wind loading.
- Short Duration Impulses (<1us) at 60Hz or Harmonic - These are caused by power line related sources, such as arcing on transformers and thyristors. Motors and appliances like hair dryers can also generate these impairments.

S-CDMA will spread very short impulses and perform better than TDMA unless the impulses are large in magnitude, then TDMA will perform better. For long bursts, S-CDMA (in general) will spread packets out further in time, thereby making FEC more effective. For long bursts (temporal dispersal) S-CDMA must use low numbers of codes per minislot at a concurrent loss of dynamic range.

Ingress (frequency varying noise):

- CB, Ham Radio and Spurs - These are relatively narrow in bandwidth (<20kHz). These spurs are usually caused by leakage from local oscillators, etc. The CB and Ham Radio are fixed in frequency but keyed on and off over time.
- CPD and Impulse Noise with Harmonic content - This noise is wider in bandwidth but still has repeating peaks in the frequency domain. Although impulse noise should be in the time varying noise category, some impulse noise has a harmonic signature in the frequency domain. This noise can benefit from frequency domain filtering.

A-TDMA has well known and relatively easy to implement frequency domain filtering. Specifically, one filter used in an A-TDMA design will improve performance from narrow band interference by 30dB. S-CDMA can theoretically have similar performance but the implementation is more complex. The MSO will be well advised to examine this non-specified but critical feature in selection of a CMTS solution.

SUMMARY OF S-CDMA AND A-TDMA COMPARISON

S-CDMA and A-TDMA have very similar performance in terms of robustness and bandwidth. The most significant differences between the two are:

- S-CDMA may be configured to have better impulse noise resistance for most impulse types, at the expense of dynamic range.

- S-CDMA has shorter preamble requirements, providing an advantage in small packet conditions.
- S-CDMA has tight timing requirements, which may require more frequent station ranging.
- A-TDMA has a better dynamic range under all modulation settings.
- A-TDMA is backwards compatible with all legacy modems, which results in better statistical multiplexing.
- A-TDMA has implementation advantages in ingress cancellation performance.

EVOLUTION TO DOCSIS 2.0 - CONCLUSION

DOCSIS 2.0 clearly adds a number of enhancements to improve the robustness and capacity of the upstream. *It is important to note that DOCSIS 2.0 only provides the tools for obtaining these benefits.* Given this, a crucial component to improving robustness and bandwidth is the software that utilizes DOCSIS 2.0 features to intelligently adapt to plant conditions. Finely tuned ingress filter DSP software, look-ahead channel hopping, ingress categorization and adaptation, and optimized fallback algorithms are just some of the requirements of intelligent PHY control.

Assuming intelligent PHY software[4], both A-TDMA and S-CDMA excel under somewhat different ingress environments. As was seen in section 2 of this paper, an A-TDMA CMTS provides a variety of robustness improvements to deployed DOCSIS 1.x modems.

The major advantages of DOCSIS 2.0 may be obtained with CMTS systems that comply with *either* the A-TDMA *or* the S-CDMA requirements. If a cable operator has exclusively deployed DOCSIS 2.0 cable modems and the HFC plant dynamic range complies with the S-CDMA requirements, it may well be desirable to operate in S-CDMA mode or A-TDMA mode depending on the plant conditions. On the other hand, if the HFC plant has a mixture of 1.x and 2.0 cable modems, it may be desireable that an MSO operate in A-TDMA mode.

Ultimately the choice of a CMTS should be based on a wide variety of considerations:

- Intelligent PHY Control Software
- Redundancy and Reliability
- Capacity and Throughput
- Integration with Voice Services
- Integration with Provisioning and Network Management Applications
- Advanced PHY Capability
- Carrier-Class Edge Routing Capability
- Advanced QoS Capability

Appendix A - Acronyms

ACG	Automatic Gain Control
A-TDMA	Advanced Time Division Multiple Access
ATP	Acceptance Test Procedure
AWGN	Additive White Gaussian Noise
CATV	Community Access TeleVision
CIR	Carrier to Ingress Ratio
CM	Cable Modem
CMTS	Cable Modem Terminating System
CPD	Common Path Distortion (non-linear mixing products)
CPMS	Codes Per Mini-Slot
CTE	Coefficient of Thermal Expansion
DAC	Digital to Analog Converter
DOCSIS	Data Over Cable System Interface Specification
DSP	Digital Signal Processor
FEC	Forward Error Correction
HFC	Hybrid Fiber-Coax
IEEE	Institute of Electrical and Electronics Engineers
ITU	International Telecommunication Union
LSB	Least Significant Bit
MAC	Media Access Control
MAP	Map of minislots (abbreviation)
MSO	Multiple System Operator
OPL	Optical Path Length
PICS	Protocol Implementation Conformance Statement
PHY	Physical layer (abbreviation)
QAM	Quadrature Amplitude Multiplexing
QPSK	Quadrature Phase Shift Keying
RF	Radio Frequency
RS	Reed Solomon (type of forward error correction)
S-CDMA	Synchronized Code Division Multiple Access
TEP	Test Execution Procedure
UCD	Upstream Channel Descriptor
VCOFDM	Variable Constellation Orthogonal Frequency Division Multiplexing

APPENDIX B - CALCULATIONS ON HFC OPTICAL PATH LENGTH CHANGES AND THE EFFECT ON STATION MAINTENANCE FOR S-CDMA IN DOCSIS 2.0

Optical Path Length Change with Temperature

Optical Path Length Change Calculations

In an S-CDMA system it is critical that the upstream codes from the CM be precisely timed so that they arrive at the CMTS with other codes from other CMs. The CMs must add delay such that all codes are aligned to within +/-2ns. Due to changes in the OPL of the HFC system over temperature, periodic ranging is required to keep the packets aligned. Periodic maintenance must be done depending on how quickly the optical path length (OPL) (or delay) changes with time.

The OPL is related to fiber index (n) and length (l) as: OPL = l n

According to Corning Glass Works, the change in refractive index of fiber over temperature is approximately the same as Δn/T of fused silica (optical fiber is fused silica doped with Germanium). In addition the change in physical length of fiber over temperature is approximately the same as the Δl/T of fused silica.

The effective refractive index of fiber is: n = 1.47

The change in refractive index of fused silica over temperature is,
$\Delta n/T = 1.28\text{x}10^{-5}/°\text{C}$ (from Corning)
$\Delta n/T = 1\text{x}10^{-5}/°\text{C}$ @ 589nm (from Oriel Instruments)

The Coefficient of Thermal Expansion of fused silica is:
$\text{CTE} = \Delta l/l\text{-}T = 5.5\text{x}10^{-7}/°\text{C}$

The optical path length *change* will be:
$\Delta\text{OPL} = \Delta l{*}n + l{*}\Delta n$
Normalizing for length and delta temperature results in:
$\Delta\text{OPL}/l\Delta T = (\Delta l/l\text{-}T){*}n + \Delta n = \text{CTE}{*}n + \Delta n = 5.5\text{x}10^{-7}/°\text{C} * 1.47 + 1.28\text{x}10^{-5}/°\text{C} = 1.36\text{x}10^{-5}/°\text{C}$ or 13.6 millimeters per kilometer degree centigrade using Corning's Δn/T. Using Oriel's Δn/T we obtain 10.6 mm/km or 10.6 ppm/°C (parts per million per degree C). It can be seen that the majority of the change is due to the refractive index change and not the physical change of the fiber length. Given the extraordinary low coefficient of thermal expansion (CTE) of fused silica, this is not surprising. (It turns out that the Oriel number is closer to reality therefore, this will be used.)

Mach-Zender Interferometer Test

To obtain experimental values for the change in OPL over temperature, a Mach-Zender fiber interferometer was used. To measure the change in OPL vs. T, the change in temperature must be known, the length of fiber, the wavelength of light and the number of fringes or beats that occur over the temperature change. This set-up was extremely sensitive to vibration (as most interferometers are) and the temperature chamber had to be shut down and allowed to freely cool from a high temperature. Without a beat counter the most convenient approach turned out to be to measure the beat cycles per second and the change in temperature vs. time. The temperature slope of the fiber was assumed to be the same as the air in the chamber since the temperature change was slow.

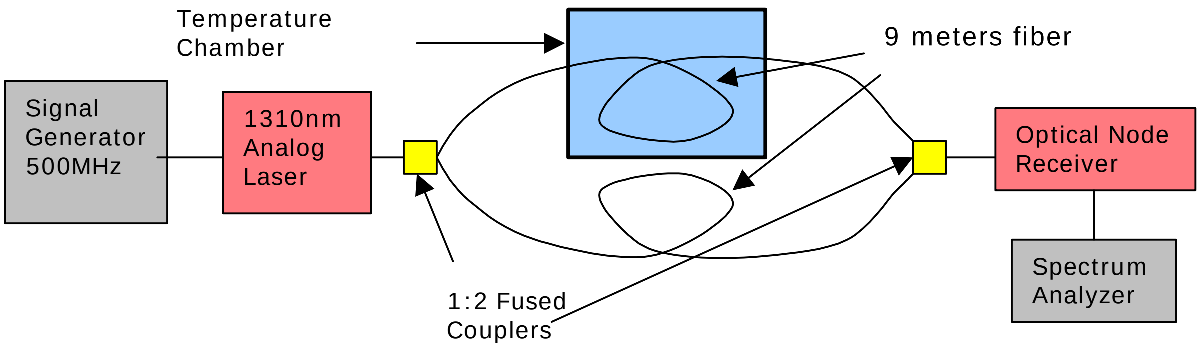

Figure 18: Experimental set-up for optical path length variation

Results

Seven measurements were made: 9.6, 8.3, 7.2, 10.2, 10.0, 7.2, 8.0 parts per million per degree centigrade. These average to 8.6 ppm/°C. This is closer to the ~7ppm/°C quoted by Passave Networks in a May 2001 IEEE 802.3ah presentation. The difference between these numbers and the Corning and Oriel numbers may be due to selection of the refractive index change at a wavelength other than 1310nm.

ANALYSIS OF DELAY CHANGE EFFECTS ON UPSTREAM SYNCHRONIZATION

Estimate of Temperature Induced Delay per unit Time

We need to know the delay changes per unit time to understand synchronization implications. Light in optical fiber traverses 1 kilometer in 5 microseconds. Therefore the delay change is:

$$\Delta T/°C = 8.6x10^{-6}/°C * 5us/km = \mathbf{43ps/°C}$$

For the DOCSIS maximum length of 200miles (320km) of fiber, the temperature-induced jitter is:

$$\Delta T/°C = 43ps/°C * 320km = \mathbf{13.7ns/°C}$$

Finally, in order to understand the impact of this jitter on the need for periodic ranging in a fully synchronized system, we have to obtain a value for the maximum rate of temperature change for the fiber. From previous field data, a value of one degree centigrade per minute has been used. This should be conservative and includes temperature changes due to environment (solar load change from clouds to sun) and rapid cooling from rain on a solar heated cable. If we use this value we obtain:

$$\Delta T/T = 13.7ns/60s = \mathbf{0.23ns/s.}$$

Delay Change due to Wind Loading

Aerial cable does stretch with wind. The construction of optical cable makes it tolerant of wind loading due to the loose tube construction that isolates the fiber from cable loading. Wind loading will affect aerial coaxial cables.

Due to the complexity and randomness of wind loading over long spans of coaxial cable, it is not clear if the loading effects will average out, reducing the possible peak values. It is difficult to analytically model this effect. Ideally, field measurements are a much more reliable method of investigating this. The following information on wind loading was an excerpt from the first published version of the DOCSIS 2.0 specification, version SP-RFIv2.0-I01-011231.

Excerpt from Appendix VIII in the DOCSIS 2.0 RFI specification, v. SP-RFIv2.0-I01-011231

Wind loading is a difficult to deal with analytically because it is unlikely to be uniform along the cable A delay model using a significant body of measured data is needed to investigate this further. Wind loading may be a source of fast delay variation and the ranging mechanism during station maintenance at the CMTS may not occur at intervals small enough to reduce this variation sufficiently.

The effects of wind loading on typical cable were investigated with a publicly available program from a coaxial cable manufacturer. These calculations showed that length changes in the range 0.01% and 0.05% are possible for various amounts of wind loading. This converts to significant propagation delay variation. As an example, with 5 miles (8 km) and 0.02% length variation, the change in propagation delay is:

(8/3e5)*(1/0.87)*2e-4 seconds = 6 nanoseconds.

This is a peak value, but the length of coax is quite short and the wind load is moderate. While the time duration over which this delay variation occurs is unspecified, it may be noted that wind gust data is readily available for most cities, and wind gust will be the primary mechanism for wind based timing changes on cable plants. For example, in New York City at the time of this writing, wind gusts of up to 40 mph are reported while average wind speed is about 10 mph. Hence, over a period of 1 to 4 seconds (the typical wind gust measurement interval), the wind speed changed by 30 mph. Much stronger wind gusts are frequently measured in locations prone to windy conditions.

FREQUENCY OF STATION MAINTENANCE

For Fiber-Temperature Induced Changes

If we wish to periodically range such that we adjust no more than 10% of the jitter requirement, then we must do periodic maintenance on each CM approximately once per second. If we are willing to allow the full movement of jitter allowance to the HFC plant, then periodic maintenance may be done once per 10 seconds. (A typical periodic ranging for an HFC system is once per 15 seconds. DOCSIS requires station maintenance about every 30 seconds (T4 time out has a maximum value of 35 seconds).

For Coax-Wind Loading Induced Changes

The wind loading is potentially much more severe than the fiber temperature effects. There is not a time constant associated with the excerpt on wind loading above. However wind induced changes can be much faster than temperature induced changes. If the 6ns change mentioned above occurs on the order of a few seconds, then station maintenance may become a large portion of the upstream traffic.

CONCLUSIONS

It would seem that HFC delay variations due to temperature swings on very long aerial fiber lengths may be accommodated by performing station maintenance at a frequency around 1Hz. However there may be physical changes to the fiber length due to wind loading as well as temperature drift that may increase the frequency of station maintenance. It is most likely to show up in high QAM/Bandwidth channels. It will be important to obtain field measurements over a variety of HFC plants over a long time period to determine the impact of this problem.

REFERENCES

[1] “Error Control Systems for Digital Communication and Storage”, Steven Wicker

[2] "Comparison of Single-Carrier, Multi-Carrier, and Spread Spectrum Modulations for Upstream PHY Layer in HFC CATV Networks", by Jeyhan Karaoguz, John Yu, Vedat Eyuboglu, Motorola, IEEE 802.14a/98-018, July 13, 1998.

[3] Prodan, R., et al., "Analysis of Two-Way Cable System Transient Impairments," CableLabs®, NCTA Proceedings, 1996.

[4] Howard, D. and Roberts, H., "Dynamic Adaptation to Impaired RF Upstream Channels using Advanced PHY", Broadcom/ADC, NCTA Proceedings, 2002.

TRANSPARENT TRANSPORT OF IP AND MPEG

Narisa N. Y. Chu
Motorola Broadband Communications Sector

Abstract

Convergence of IP and MPEG technology has gained momentum in many sectors of the Cable, Satellite, and Terrestrial Broadcast industry worldwide. With increasing cross-flows between IP and MPEG traffic, the Set-Top Terminal now has to recognize IP Multicast addresses, in addition to other essential information. New descriptors, new tables and IP-Control Channel protocol have been contemplated. The DVB Multi-Protocol Encapsulation standard is extended for implementation efficiency. Other much-proliferated new protocols for the MPEG stream are also investigated, treating IP data transport as an equal, if not the primary service to video.

The issues at stake are compatibility of IP and MPEG, network pre-determined conditions, scalability, and future evolution in anticipation that Internet streams might play a major role in the broadcast video channel space.

This paper concerns the end-to-end signal flow from the program server to user receiver. Critical issues for IP and MPEG interworking are reviewed from an STT perspective.

INTRODUCTION

The scope of transparent transport of IP and MPEG streams involve at least 3 aspects:

(1) IP over MPEG,
(2) MPEG over IP, and
(3) IP over MPEG with Terminal Mobility.

Figure 1 illustrates the concept of these various transport streams and the receiving terminal devices.

FIGURE 1. IP AND MPEG TRANSPORT

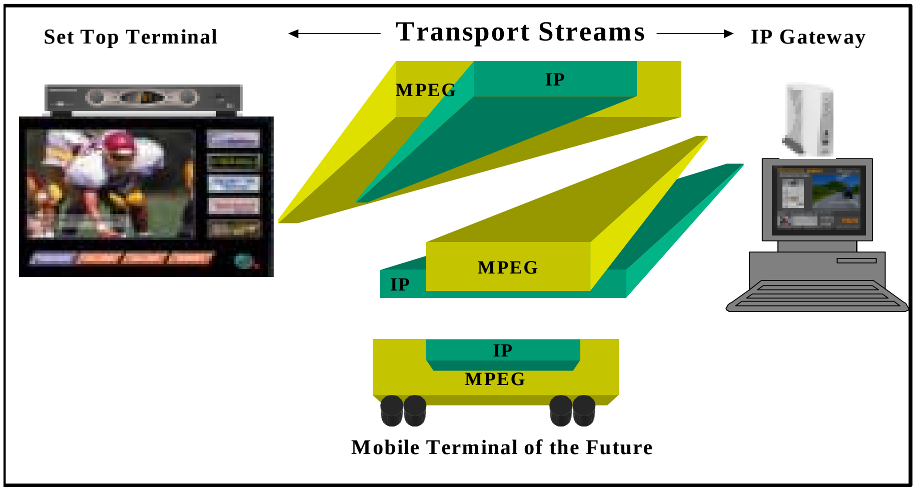

This paper provides an overview of the critical issues for IP over MPEG interworking (shown in the top part of Figure 1) for transparent transport of data and video streams. Other possible transport streams, also under active development, are only briefly mentioned with references provided at the end.

Market forces from not only America, but also Europe, are reviewed for services based on an integrated transport of MPEG and IP streams. Technical proposals associated with IP transported over MPEG streams are investigated. These proposals reflect on-going discussions within the Advanced Television Systems Committee (ATSC), Digital Video Broadcast (DVB) Forum, Internet Engineering Task Force (IETF), International Telecommunications Union (ITU) and Society of Cable Telecommunications Engineers (SCTE). Some evaluations are described and future directions are suggested regarding applications, implementations, transport stream management and services.

This paper does not intend to address all network essential issues with respect to MPEG over IP, nor MPEG over Mobile transport. This paper concentrates on IP over MPEG transport impact to the STT from the perspective of required Service Information signalling; acknowledging that future network transformation could generate other impact on the STT. These future factors are identified for further exploration along with directions that technical solutions might eventually converge.

MARKET BACKGROUND

The Internet access has been enabled via the evolving telecommunications infrastructure since the 1950's, supported by the TCP/IP transport flow. However, the vast majority (80%) of IP transport has been via narrowband connections based on the traditional telephony networks [1].

Almost fifty-four million households (or 50.5% of all households in the US) had Internet access as of September 2001 [1].

The introduction of cable modem since 1998, in conjunction with Digital Subscriber Loop (DSL) upgrades in the copper wire networks, provides high-speed Internet access in a leap-frog fashion to a range in megabits per second.

Digital transport streams carried over cable and satellite networks have been conditioned with MPEG streams since 1998. Terrestrial broadcast networks have followed suit in parallel.

Cable Modem Market Penetration

The rapid emergence of cable modems allows major influx of IP streams onto the MPEG facilities, as cable networks have demonstrated. The latest FCC report [1] has highlighted the demand of high-speed Internet delivered over Hybrid Fiber Coaxial (HFC) systems in the last mile. Fifty-four percent of the total high-speed lines were carried on cable by the end of June 2001. Cable companies report almost 5.2 million high-speed lines in service using cable modem technology at the end of June 2001, compared to 1.4 million at the end of 1999 [1].

High-speed Digital Network Readiness

Regarding the network readiness, i.e., about the availability of cable modem-ready plant, publicly available sources estimate that cable modem service is now available to about 70% of US homes.

Other broadcast networks are also shown enthusiasm to join the cable modem success story in receiving both data and video with STT(s), particularly in areas outside of the US. Even within the US, satellite technologies account for between 50,000 and 150,000 high-speed lines as of June 2001.

High-speed satellite services are now available in all 50 states, and multipoint microwave data distribution systems currently reach 55% of the population [1].

MPEG and IP Cross-flow

The typical MPEG network and the IP network start to converge as applications converge between TV STT and computer network gateway controllers. The underlining delivery system has employed sophisticated protocols, in order to present a transparent experience to the viewer for entertainment access and information retrieval. As the number of channels and bandwidths increases, the form to mix MPEG and IP streams has increased as well. The push for standardization in conjunction with implementation efficiency has received significant attention recently.

Some market data suggest, without the consideration of the content involved, that the residential high-speed subscription will increase from 1.9 million in 2000 to 40 million in 2005. By 2004, 29 % of households will access the Internet through cable modem services, 21% through DSL and 5.7 % through wireless and satellite technologies [1].

Forecasts also said that in 2005, the average Broadband household would download about 70 Mbits of files, consume more than 20 minutes of streaming per day, and download 32-hour long movies per month [1].

Cable modem subscription reached 3.9 million in 2000 with a projected rate of double the current increase to reach 28-30 million by 2006 [1]. This calls for more than a quarter of IP streams ready to be transported via MPEG networks.

INTERWORKING TECHNOLOGY ACCELERATION

Based on the above mentioned market trends, convergence of IP and MPEG technology has since gained momentum in many sectors of the Cable, Satellite and Terrestrial Broadcast industry; be it the International Standards Body, Manufacturers' Development Planning Process, Operators' Forum, or New Service Offering.

With increasing cross-flows between IP and MPEG traffic, the Set-Top Terminal (STT) that receives MPEG streams, now needs to recognize IP Multicast addresses, in addition to other essential information. Realizing the shortcoming of the broadcast data standards, new descriptors (e.g., MAC_address_list descriptor, multiprotocol_encapsulation_broadcast_descriptor), new tables (e.g., IP Map Table) and new IP Control Channel (IP-CC) have been designed to transport IP over MPEG in a manner that calls for innovative services, efficient implementation and backward compatibility. Some make use of the DVB Multi-Protocol Encapsulation (MPE) standard to achieve efficient tuning in the STT, within the existing MPEG framework. Others express a much-proliferated scope for the MPEG stream, treating IP data transport as an equal, if not the primary service.

It is un-disputable that IP streams traverse through MPEG established transport streams would find numerous applications beyond simple video association. The industry is far from a lack of imagination, as recent mobile demonstrations about digital MPEG over IP [17] and Multimedia Car Platform [18] applications suggest.

However, one has to:

(1) understand the service requirements first,
(2) analyze the scope of the technology challenge,
(3) evaluate, step-by-step, available tools and frameworks, e.g., descriptors and tables, that can facilitate an effective solution,
(4) compromise for standards to achieve economics of scale, and
(5) implement efficiently.

Therefore, the issues at stake are:

1. How do cable, satellite and terrestrial broadcast networks differ in transport of IP within the MPEG framework?
2. How does transport of IP over MPEG or MPEG over IP differ in protocol architecture?
3. Can one afford a clean slate, abandoning the prior implementation of MPEG-2 in the field?
4. Can one ignore, initially, the network pre-determined conditions, scalability, and future evolution in anticipation that Internet traffic could out-number the broadcast video traffic?
5. What is the protocol overhead involved in converging MPEG and IP transport? Or is this a new Protocol?
6. What is the viability of MPEG-Mobile applications?

These concerns suggest that one begin with a solid base on realistic requirements which can give proper guidance to the evaluation of the appropriate technical alternatives and their extensions to identify those immediate incremental revenue streams for the near term; with open and migrateable standards in mind for future realization of grandeur applications, without placing unnecessary technical obstacles along an evolutionary development process. In this capacity, IP over MPEG transport appears more familiar for the broadcast industry to tackle than MPEG over IP, Home Networking or MPEG over Universal Mobile Telecommunications System (UMTS). Therefore, with focus on IP over MPEG, an evaluation of alternatives is described below.

ALTERNATIVES FOR TRANSPARENT TRANSPORT

Within the years of 2001-2002, significant momentum has been established in addressing IP and MPEG (the latter being the baseline transport stream to the framework of ATSC, DVB, SCTE and ITU-T/SG9) interworking and interoperability. The technical conventions considering MPEG applications include ATSC and SCTE in North America and DVB in Europe, while that for IP extensions is typically IETF. Example technical discussions are shown in 15 references cited at the end of this paper.

Some of these references address requirements and others, technical designs. Requirements issues are first discussed as follows.

REQUIREMENTS ASSESSMENT

References [2] to [5] identify requirements for carrying IP over MPEG. Requirements can be put in 2 categories:

1. Commercial requirements brought forth by Service Providers.

2. Technical requirements brought forth by Manufacturers.

Businesses have emerged based on Advanced Television Enhanced Forum (ATVEF) applications where IP data were carried within the MPEG-2 video component for linking specific video to a corresponding Internet Web access [2]. Some applications were introduced as early as in 1999. In 2000, commercial requirements were identified by the European broadcast operators of the DVB project to offer Internet services via satellite broadcast channels [6]. These cross-continental requirements were compared in Reference [2].

References [5] and [7] provide a broader brush of the scope involving IP and DVB. As these references continue in development, it becomes clear that IP over MPEG can easily enlarge the traditional video bound services into Internet value-added services.

The above identified requirements have been carried out in both America and Europe in business offers from cable and satellite broadcast operators with limited success to date. A few technical alternatives are discussed with foreseeable enhancements below, to improve the service prospect for the future.

IP OVER MPEG ALTERNATIVES

Essential information for the broadcast network and the Set-Top Terminal to access IP Multicast addresses are described in References [6] to [13] for carrying IP data over MPEG streams (The DVB stream and the MPEG stream can be used interchangeably here.) Many of these references start with the DVB Multi-Protocol Encapsulation (MPE). An all-encompassing table has been created to cover network specific parameters, the IP Map Table (IMT)[6].

IP Map Table

When IP data are carried within MPEG Transport Streams (TS), the transmitted network addresses have to be made known to the STT. Currently many proprietary mechanisms exist to signal this information. “The situation becomes confusing especially when the MPEG network carries split IP data across Packet IDentifiers (PIDs) in the same TS and sometimes even across transport streams” [6]. This is a familiar case in the satellite and terrestrial broadcast environment, not so much needed in the cable environment.

The IMT mechanism signals to the STT for a mapping of IP/MAC unicast and multicast addresses to the MPEG parameters of network ID, original network ID, transport stream ID, service ID and PID. This mechanism allows a fully automated tuning for DVB-DATA receivers.

Table 1 contains an IP Platform ID and IP_platform_descriptor to track IP/MAC addresses for unicast and multicast, IPv4 as well as IPv6 addressing schemes [6]. The advantage here is that all network and address resolution information is consolidated and clearly specified with the PID, original network ID, transport stream ID and service ID by the IMT.

TABLE 1. IP_MAP_TABLE (IMT) SECTION [6]
(Extracted from Reference [6] for ease of comparison)

Syntax	No. of bits	Mnemonic
IP_MAP_TABLE{		
table_id	8	uimsbf
section_syntax_indicator	1	bslbf
private_indicator	1	bslbf
Reserved	2	bslbf
section_length	12	uimsbf
IP_platform_id	16	uimsbf
Reserved	2	bslbf
version_number	5	uimsbf
current_next_indicator	1	bslbf
section_number	8	uimsbf
last_section_number	8	uimsbf
reserved_future_use	4	bslbf
IP_platform_descriptor_length	12	uimsbf
for(i=0;i<N;i++){		
Descriptor()		
}		
network_loop_count	8	uimsbf
for(i=0; i< network_loop_count; i++){		
network_id	16	uimsbf
transport_stream_loop_count	8	uimsbf
for(i=0; i< transport_stream_loop_count; i++){		
original_network_id	16	uimsbf
transport_stream_id	16	uimsbf
service_loop_count	8	uimsbf
for (i=0; i< Service_loop_count; i++){		
service_id	16	uimsbf
PID_loop_count	8	uimsbf
for (i=0; i< PID_loop_count; i++){		
elementary_PID	13	uimsbf
address_type	2	bslbf
Reserved	1	bslbf
address_loop_count	16	uimsbf
for (i=0; i< address_loop_count; i++){		
If address_type == 0x00 {		
IPv4_address	32	
IPv4_slash_mask	8	
}		
If address_type == 0x01 {		
IPv6_address	128	
Ipv6_slash_mask	8	
}		
If address_type == 0x02 {		
MAC_address_range	1	bslbf
Reserved	2	bslbf
Reserved	5	bslbf
If MAC_address_range == 0 {		
MAC_address	48	
}		

```
                              Else {
                                  Lowest_MAC_address            48
                                  Highest_MAC_address           48
                              }
                          }
                      }
                  }
              }
          }
      }
      CRC_32                                                    32        Rpchof
}
```

IMT is compact and this method avoids advertising MAC addresses on the broadcast link. It also avoids frequent recompiling of the tables and is transparent to the transition of IPv4 to IPv6 [6].

However, MPEG fundamentally recommends that no elementary PID should appear in any table other than the Program Map Table (PMT). The IMT thus deviates from the traditional MPEG rule. One cannot deny the advantage of introducing the IMT where pure IP traffic is allowed in an independent, unused video channel. Up to 24 digital channels, depending on the modulation scheme, can be made available at the operator’s disposal, after the analog channel is transformed.

As the IMT is made optional, the operator will have a choice to either use existing SI tables with special descriptors or to render the IMT for complete IP routing information.

When dealing with the cable network, it is a more coordinated and well-informed network due to the knowledge stored in the Headend with respect to addressing to the individual STT. The network controller generally knows the configurations of the STT and their Headend association in a large network. With applications such as ATVEF, most of the addressing issues would be resolved by the applications before the STT has to get involved in the consideration of which IP router or bridge it has to communicate with.

Therefore, cable STT would direct its resource for fast tuning as its primary responsibility, leaving aside IP stream management function for the IP protocol to resolve. Reference [10] takes advantage of this well-informed cable environment.

Use of Descriptors

Reference [10] makes use of DVB MPE standard and created a MAC_address_list descriptor for multicasting. This descriptor has recently become an SCTE standard [11]. It has also being considered by the ATSC [12]. The largest cable service operator in North America has endorsed it based on its expected implementation efficiency. Table 2 shows its current form for such IP data multicasting.

TABLE 2. MAC_ADDRESS_LIST_DESCRIPTOR [10]

Syntax	No. of bits	Mnemonic
MAC_Address_List_descriptor() {		
descriptor_tag	8	Uimsbf
descriptor_length	8	Uimsbf (L)
mac_addr_list	1	Uimsbf
mac_addr_range	1	Uimsbf
pdu_size	2	Uimsbf {1024 bytes, reserved1, reserved2, 4096 bytes}
encapsulation_type	2	Uimsbf {DVB, reserved1, reserved2, ATSC }
reserved	2	Uimsbf
if (mac_addr_list == 1) {		
num_in_mac_list	8	Uimsbf (m)
M = m*sizeof(mac_address)		
L = L – M		
For (i=0; i < m; i++) {		
mac_address	48	Uimsbf
}		
}		
if (mac_addr_range == 1) {		
Num_of_mac_ranges	8	Uimsbf (n)
N = (n*sizeof(mac_address)*2)		
L = L – N		
for (i=0; i< n; i++) {		
Highest_mac_address	48	Uimsbf
Lowest_mac_address	48	Uimsbf
}		
}		
for (i=0; i< L – 1; i++) {		
Private_data_byte	8	Uimsbf
}		
}		

Although not specifically mentioned, extending this descriptor for unicasting should be straight-forward. It is possible to extend for coverage of IPv6 address scheme as well, if necessary. These additional capabilities do not impose un-stoppable obstacles when using the MAC_Address_List descriptor.

Another use of descriptor [13] creates a Multiprotocol Encapsulation Broadcast descriptor to be used associated with the Service Definition Table (SDT). This is different from use in a Network Information Table (NIT) coupled with the IMT as per Reference [6], or in the Program Map Table (PMT) as suggested in Reference [10]. Table 3 shows the current proposal of the multiportocol_encaspsulation_broadcast descriptor.

TABLE 3. MULTIPROTOCAL ENCAPSULATION BROADCAST DESCRIPTOR [13]
(Extracted from Reference [13] for ease of comparison)

Syntax	No. of bits	Mnemonic
multiprotocol_encapsulation_broadcast_descriptor(){		
descriptor tag	8	uimsbf
descriptor_length	8	uimsbf
data_broadcast_id	16	uimsbf
component_tag	8	uimsbf
service_id	16	uimsbf
for(i=0; i<N; i++){		
address_type	8	uimsbf
address_length	8	uimsbf
for(j=0; j<N; j++){		
address_byte	8	uimsbf
}		
}		
}		

"Using the descriptor in the SDT, it is possible to associate any type of unicast/multicast address to any service within a network. Using the optional component_tag it is further possible to define the elementary stream to which the MAC/IP address references. This descriptor is the same as the data_broadcast_id_descriptor, with the id_selector_bytes being utilized from the *component_tag* field onwards" [16].

Comparison of Table and Descriptor Usage

Use of the IMT [6] starts with a clean slat for IP network addressing. Use of the MAC Descriptor [10] in the PMT advocates a fundamental adherence to the MPEG principle, thus fully taking advantage of the MPEG efficiency, leading to ease of implementation and tuning efficiency. Use of the Multiprotocol Descriptor [13] in the SDT appears elegant, however it can restrict services to be only within one transponder. Table 4 compares the characteristics of these transport-enabling mechanisms.

TABLE 4. COMPARISON OF ALTERNATIVE PROTOCOLS FOR IP OVER MPEG TRANSPORT AS OF MARCH, 2002

(Information might change, as these techniques are being consolidated and harmonized.)

	IMT [6]	MAC Descriptor [10]	Multiprotocol Descriptor [13]	IP- CC [4]
Application	Broadcast/ Internet Service Provider	ATVEF	None specified	Broadcast Mobile
Network Focus	Satellite with multiple transponders	Cable	Satellite with single transponder	Mobile
MPEG Efficiency	Ignored	Adherent	Adherent	Unknown
MPEG Table Association	New	PMT	SDT[1] and PMT	In development
IPv4/ IPv6	Yes	Extendable	Yes	Yes
IP Address Resolution	Required	Not required	Required	Required
Multicast/Unicast	Yes	Extendable	Yes	Yes
Receiver Implementation Efficiency	No	Yes	No	Unknown

The trade-off among these various mechanisms seems to be between the tuning performance and the ability to manage the IP streams within the MPEG streams across all networks.

Reference [9] calculated the overhead for IP and MPEG streams. In a pure engineering fashion, Reference [9] brings up an objective comparison of the protocol overhead involved in MPEG and IP transport. The overhead and performance penalty warrants more elaborate examinations with respect to various transport alternatives.

The IP network management is yet another area requiring further studies. Does the MPEG Service Information need to track router and bridge addresses while IP data are carried in and out of the broadcast networks?

Other Extensions

Mobile requirements are also generated from the perspective of carrying IP over MPEG streams [4].

A recent proposal, IP-Control Channel (IP-CC), has taken into account merits from both the IMT and the two Descriptors described above. It also acknowledges network pre-determined conditions, scalability, and future evolution in anticipation of the vast Internet traffic on Mobile networks with video services.

The subject of IP over MPEG is not just addressed by the DVB group alone. IETF also attempted a Birds-Of-the-same-Feather (BOF) session creating IP over MPEG functional requirements [3].

[1] This table is not available for North American SI practice.

From the other perspective, the DVB-IPI group on Internet Protocol Infrastructure has actively pursued protocols for MPEG transport over IP [14]. ITU-T, Study Group 9 has published "Webcasting", J.120 [15], which defines the transmission protocol and system configuration for distributing sound and television programs over the Internet. It concerns the end-to-end signal flow from the program server to user's receiver. This transmission chain contains the signal encoding/decoding, packet mapping as well as session control and network transmission.

SUMMARY AND FUTURE DIRECTIONS

Future directions are suggested:

1. IP data can be carried in conjunction with video components.
2. MPEG channel can carry pure and full IP data.
3. IP stream can be systematically encapsulated within the MPEG stream.
4. IP and MPEG flow can be transported and intermixed, with peer relationship.

Technical solutions to enable any of the above directions to flourish are summarized as follows.

Recognition, resolution, and routing of IP addresses within the MPEG video component or in an independent data component have been investigated. The MAC_Address_List descriptor could fulfill the video related IP data broadcast. To meet the expanded Internet traffic, the IMT, with or without the MAC and/or the Multiprotocol descriptor can be employed to introduce full-fledged high-speed data services in a ground-breaking mode. The establishment of the IMT implied some IP network management capability would be established in the MPEG SI for extended utilization of freed-up digital channels. Conceivably, video and data can be intermixed, linked, or transported side-by-side. Applications based on transport of IP over MPEG, or vice versa have been demonstrated [17], and services are in active pursuit by some operators.

This trend of service expansion begs the question about combining control for signalling of video and data transport: Would the future form of Service Information (SI) [16] be carried through in either MPEG or DOCSIS format?

As Mobile begins receiving video services, the IP-CC represents one method to share MPEG Service Information with the complex mobility network; in particular, handling of roaming can be a major challenge. The advent of IP and MPEG network interoperability for services beyond the traditional video boundary, to be enhanced with mobility capability, opens up enormous potential for the communication, entertainment and information industry.

It behooves service providers to cooperate and direct the technical community in the latest standardization of IP versus MPEG interoperability.

REFERENCES

1. High-Speed Internet Report, FCC Document 02-33, Feb. 6, 2002.
2. Narisa Chu, et. al., "IP Over DVB System Requirements Issues for SI-DAT to Consider", Motorola's contribution to DVB, http://www.dvb.org, (same for references below whenever DVB is mentioned,) SI-DAT 603, July 23, 2001.
3. Gorry Fairhurst, et. al., "Requirements for Transmission of IP Datagrams over DVB Networks", IETF Draft, December 2001.

4. “IP-CC Requirements Specification,” Juha-Pekka Luoma, Nokia, DVB SI-DAT 621, Nov. 16, 2001.
5. “Response to RFI on IP over DVB,” Rod Ragland and Andrew Valentine, Hughes Network Systems – Europe, DVB, SI-DAT, 2001.
6. “Joint Proposal for an Extension to the MPE Mechanism of EN 301 192 ,” SES/Astra and France Telecom, DVB SI-DAT 585r1, Nov. 21, 2001.
7. “IPI Network Architecture Proposal,” Georges Martinez, Motorola – Paris Lab, DVB-IPI, 2001.
8. “Stages in receiving DVB Services for an IP-enabled STB,” Jeffy Goldberg, Cisco and Paul Stallard, Tandberg, DVB-IPI, 2001.
9. “A Comparison between MPEG and Ethernet Overhead for Transporting IP Packets,” Ciro Noronha, SkyStream Networks, DVB, SI-DAT, 2001.
10. Narisa Chu, et. al., “Proposal for Automatic Extraction of the IP Services Carried in DVB Streams,” Motorola’s contribution to DVB, SI-DAT 605, September 7, 2001.
11. “IP Multicast for Digital MPEG Networks”, DVS/311r5, SCTE, Oct. 19, 2001.
12. “Delivery of IP Multicast Sessions over ATSC Data Broadcast”, Draft, ATSC, Doc. S13-12r16, July 12, 2001.
13. “Recommendation for the Enhancement of the DVB MPE Standard”, DVB, SI-DAT 602, Adrian Jelffs, SkyStream Networks, 2001.
14. “Architectural Framework for the Delivery of DVB-Services over IP-based Networks”, DVB-IPI2001-012R08, Nov. 19, 2001.
15. “Distribution of Sound and Television Programme over Internet,” Miyaji, KDD, J.120, 2001, ITU–T, SG9, Geneva; Telephone +41 22 730 6071, Facsimile: +41 22 730 5853; E-mail: sa-miyaji@kdd.co.jp or masamichi.niiya@itu.int.
16. DVB-SI Specification, ETSI EN 300 468 V1.4.1 – 6.2.11.
17. DVB over IP demo at the 45th DVB-TM Meeting on Feb. 25 & 26, 2002, in EBU, Geneva, led by CISCO, participated by Pace, Philips, Scopus, Tandberg and Thales.
18. MCP demo at DVB-MHP Meeting on Feb. 20, 2002, in EBU, Geneva, presented by Braunschweig Technical University. http://mcp.fantastic.ch

TRANSPORT OF CBR TRAFFIC ON IP NETWORKS

Tom Carr, Bruce Roe
Wave7 Optics

Abstract

The success of future distribution networks will depend on their ability to support legacy services including committed bit rate traffic. Most of this traffic is transported by the PSTNs over T1 facilities.

This paper describes a technology, Time Division Multiplexing over IP, which is capable of providing T1 circuit emulation over IP networks.

INTRODUCTION

Industry visionaries foresee Next Generation Networks that offer hundreds of megabits of bandwidth to the consumer, extensively or exclusively using IP as the network transport protocol. This is an easy vision to believe in, and one that we are convinced will evolve into reality.

While we can envision our final destination, the question remains, how do we get there? Clearly the distribution network cannot, and will not, be replaced en masse with a new IP-based architecture. Pockets of new high-speed IP networks will be deployed and, over a very long period of time, finally consign the old copper telecommunications plant to the pages of history. This slow evolution means that some legacy services must be supported on the new architectures.

One likely architecture for future distribution networks is that shown in Figure 1. An Ethernet network of one or more gigabits extends over fiber from a Head-End to an active bandwidth management element. From there, the Ethernet is extended over multiple fibers to serve a pocket of customers. By using

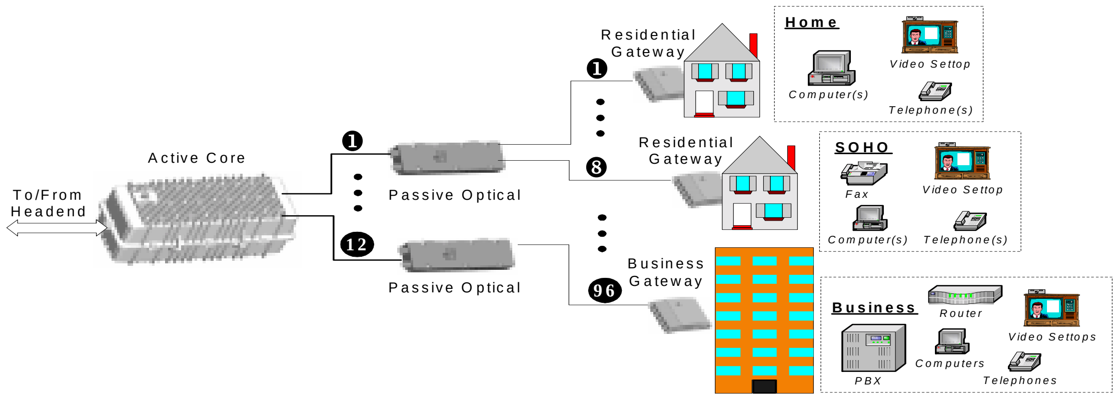

Figure 1: Next Generation IP-based Distribution Network

WDM, RF digital and analog video can be distributed over the same fibers. Plain old telephone service (POTS) would be supported by using voice over IP (VoIP) for transport. In this manner voice, video and very high-speed data could be offered over a single efficient network. Unfortunately, this alone will not provide for a significant and lucrative portion of traditional legacy telecommunication services, committed bit rate services.

Obituaries have been written for committed bit rate traffic such as private line service, international frame relay networks, and every other non-IP protocol, but the double-digit compounded growth rate for these services, particularly in the international markets, continues. Some forecasts[1] of international frame relay, for example, predict a compounded annual growth rate of 14-16% at least through 2004. Even X.25 networks still exist and continue to grow. The inertia of migrating these networks to IP will be fueled by sluggish economies, and the falling prices of both T1 service and old technology equipment. Any new distribution network, particularly those limited in geographical scope, must either accommodate these legacy services or exclude large, profitable markets. It is simply not economical for end-users to convert national or international non-IP networks to IP in a piecemeal fashion.

PBXs present another problem: signaling. Signaling consists of basic features such as recognizing that the phone is off-hook, or needs to ring; the more advanced properties required for reaching the proper destination and billing; and still more sophisticated characteristics, such as caller identification, call forwarding, and conference calls. There are literally thousands of such telephony features, with dozens of national and local variations. Available VoIP integrated circuits can handle some, but not all, of the PBX signaling in use in the U.S.. Converting PBX voice circuits to VoIP could require the end-user to give up some useful or much needed features.

The one common element of committed bit rate traffic carried by PSTNs is that they are primarily transported via T1s. Having the ability to transport T1s over IP, regardless of the data or signaling protocol, would be an ideal solution for supporting legacy services in a new IP environment. Such a technology does exist and is called TDMoIP, Time Division Multiplexing over IP. TDMoIP is a technology that combines features from Time Division Multiplexing and IP to deliver synchronous T1 circuits ***transparently*** over IP networks. An individual channel within the T1 stream is not changed in any way, nor is there any signaling conversion. This technology would be used point-to-point, from the customer's premises to the Head-End.

TDMoIP OVERVIEW

A T1 frame is composed of 24, single byte time slots plus a single synchronization bit, for a total of 193 bits. Frames are transmitted at a rate of 8000 per second, resulting in a data stream of 1.54 megabits per second. In principle, the simplest implementation of TDMoIP simply encapsulates a number of T1 frames in an IP packet by tacking on the appropriate IP header. At the destination, the stream is then recreated by stripping away the headers and reassembling the segments. *It is important to note that TDMoIP transports the T1 circuit without any attempt at interpreting the data*. This process is oblivious to signaling, time slots, or whether voice or data are being transmitted. This also implies that a data bit-stream using the entire 193 bit frame can be supported.

Standards

TDMoIP is essentially the IP counterpart of the same service in ATM referred to as "Circuit Emulation Service," (CES). While there are, as yet, no standards[2] for TDMoIP, such standards do exist for ATM-CES. Furthermore, since the performance requirements for TDM are independent of the method of transport, it is clear that the performance requirements for ATM-CES should be adhered to as closely as possible in TDMoIP. Nonetheless, *how* to achieve that has not been standardized, and will likely vary between TDMoIP platform vendors. Thus, the interoperability of equipment from different vendors should not be expected.

The TDM performance guidelines to follow primarily relate to clocking. The clock rate of the TDM stream should be stable to within +/- 32 ppm[3] and wander should not exceed 80µsec per day[4]. Performance standards directly related to IP networks, such as the maximum allowable packet loss, do not yet exist. These standards are far stricter than what is required when terminating TDM on end-user equipment and it is entirely possible that new standards may be formulated for this specific purpose.

Packetization

Primary issues to resolve include which IP protocol to use and how many T1 frames should be placed in each IP packet. Since there is no standard, any IP protocol could be used. Some, however, would clearly be inappropriate. The end-to-end reliability offered by TCP, for example, is not useful for voice packets, since re-transmitted voice packets will reach the receiving side out of order, only to be dropped anyway due to delay constraints. A good choice of protocol could be RTP and the associated RTCP, which in certain networks would offer better clocking functions. Ultimately, for Ethernet networks, only UDP is fundamentally needed.

There are tradeoffs to be considered with selecting the number of T1 frames per IP packet. The fewer the frames, the greater the IP overhead, which will increase the amount of bandwidth needed per T1. The greater the number of frames, the greater the end-to-end delay, packet loss becomes more onerous, and larger buffers are required. Larger number of frames per packet could also exacerbate adaptive clock wander. QoS demands that the number of frames per packet be kept small despite the overhead penalty. At four frames per packet, this penalty is about 50%, meaning a 1.54 Mb/s T1 would require more than 2.3 Mb/s bandwidth in IP. Even with this overhead, a 1 Gb Ethernet network is capable of supporting several hundred T1s.

Signaling

There are three primary types of signaling: in-band signaling, channel associated signaling, and common channel signaling. None of these are impacted by TDMoIP. In -band, as the name suggests, is signaling in the audio band of speech. The ubiquitous ‘touch tone’, or Dual Tone, Multiple Frequency (DTMF) is an example of in-band signaling. Since these tones are encoded in the T1 frame time slots, they are automatically carried over TDMoIP.

Channel associated signaling is also carried within the T1 frame time slots. Specific voice bits are ‘robbed’ and the signaling bits are substituted. TDMoIP does not distinguish between bits used for voice and data bits, thus this signaling is carried transparently.

Primary Rate ISDN signaling, PRI, is a popular type of common channel signaling. The twenty-fourth time slot of the T1 frame is used to carry the signaling data for the other twenty-three time slots. Again, since TDMoIP does not distinguish between voice and data, the signaling is carried transparently.

Clocking

Clocking is the most difficult problem to solve in deploying TDMoIP. There are several methods of clocking in any type network. These are: independent Stratum 1 clocks at each end-point; a synchronous network in which the primary reference source (PRS) clock is distributed throughout; an asynchronous network in which a network clock is distributed through out; and adaptive clocking in an asynchronous network with no distributed clock.

Timing is not provided in IP networks, thus synchronization must be achieved from an external source. This can be accomplished by: a Stratum 1 external master clock at each end of the TDMoIP circuit; clocking from an external clocking distribution network; or in-band clock recovery and regeneration, i.e., adaptive clocking. Stratum 1 clocks are so precise that T1 streams timed by separate Stratum 1s will be synchronized. This is a relatively costly solution, although the prices, particularly GPS-based Stratum 1s, have been declining recently.

An external distribution network for clocking is also an expensive solution, and severely compromises the entire concept of having a single network to maintain. In this scenario, a separate network would be maintained just to send clocking signals to every end-point.

In-band clock recovery and regeneration, or adaptive clocking, is the most cost effective and will meet the requirements of customer premise equipment such as PBXs. In adaptive clocking, the source TDMoIP unit, which is clocked to a Primary Reference Source, simply sends the data to the customer TDMoIP unit. The customer unit writes data to the segmentation and re-assembly (SAR) buffer and reads it with the local clock. The level of the SAR buffer controls the output frequency of the local clock by continuously measuring the fill level around the median position and feeding this measurement to drive a Phase Lock Loop (PLL), which in turn drives the local clock. Thus, the local clock frequency is modified to keep the re-assembly buffer depth constant. When the TDMoIP unit senses that its SAR buffer is filling up, it increases the clock rate. When the unit senses that the SAR buffer is emptying, it decreases the clock rate. Since the packet arrival rate is directly dependent upon the packet transmission rate established by the PRS at the head-end, synchronization of the TDM stream is maintained.

The proper choice of buffer size can prevent buffer overflow and underflow, and at the same time, control delay (greater buffer sized implies greater delay). The buffer size is proportional to the maximum packet delay variation. This variation should be determined by summing the delay variation of each network device in the circuit path. The sum of the measured delay variations that each piece of equipment introduces must be smaller than the maximum packet delay variation configured on the TDMoIP unit. If not, underflows and overflows will occur. This buffering will also remove any jitter encountered by packets arriving at slightly different time intervals.

In the event a TDMoIP packet is lost during transport, a dummy packet is transmitted by the customer TDMoIP unit in order to maintain clocking in the T1 output stream. Since the packet will contain no customer data, it must still be considered a frame-slip, however, timing problems will be minimized.

TDMoIP and Network Delay

A T1 frame represents .125 milliseconds of real time. Processing time for packetization and recreation of the T1 is less than five milliseconds. Delay is not an issue in TDMoIP unless each TDMoIP packet contains a great many T1 frames.

End-to-end round-trip network delays greater than 30 milliseconds could necessitate the need for echo cancellation on voice circuits. Round-trip delays of more that 300 milliseconds will result in unacceptable QoS for conversational speech. This is the same as for VoIP service. In both TDMoIP and VoIP, quality can only be provided and assured on tightly managed networks with well-executed prioritization procedures.

Packet Loss

Sequence bits are used to determine a lost packet condition. In the event of lost packets, timing is maintained through the insertion of dummy packets carrying appropriate framing bits. It is possible to mitigate voice quality impairments by repeating the frames that preceded the lost packet. However, the total loss of several contiguous T1 frames would not significantly degrade a voice circuit.

VoIP packet loss could be used as a guideline. Unfortunately, various studies show that while some deployments can sustain a 5% packet loss before realizing a significant degradation of QoS, other deployments can suffer less than .2% packet loss. Obviously packet loss must be minimized. As with network delay, this implies that well conceived prioritization methods must be used within the IP network, with TDMoIP given the highest possible priority.

Prioritization

Properly prioritizing the TDMoIP packets will minimize network delay and packet loss. This is critical for maintaining satisfactory QoS. By marking TDMoIP packets they may be easily identified and prioritized. This is done through proper marking of the Type of Service (ToS) bits, and VLAN tagging and priority labeling according to IEEE 802.1 p&q. Additionally, there is an assigned, IANA-registered UDP socket number for TDMoIP. These features simplify flow classification through switches and routers. In a tightly managed network, the QoS of TDMoIP should be equal to that of a traditional T1 circuit.

Supported Features

TDMoIP is capable of supporting unframed T1, Super Framing (SF), Extended Super Framing (ESF), as well as Channel Associated Signaling (CAS) and Common Channel Signaling (CCS), including Primary Rate ISDN (PRI).

There are two solutions for supporting fractional T1s. The first is to put a large multiple of the individual time-slots in the same TDMoIP packet. This would reduce the overhead penalty. However, as discussed earlier, this would also increase delay as well as create major QoS problems in the event of a lost packet. The second method is to transport the fractional T1 as though it were a full T1, filling the unused time-slots with idle code. This would require using the same amount of network resources for a fractional T1 as a full T1.

Deployment

A single, small end-user integrated access device can be configured to house a video port, Ethernet port, and multiple POTS and TDMoIP ports. The TDMoIP packets would be routed through the network to a Head-End TDMoIP unit. A T1 circuit, identical to the end-user's original T1, would be generated by the Head-End unit.

TDMoIP may be deployed in overbuilds by network service providers or in a greenfield environment. There are three primary methods of provisioning the T1 service. The first is to simply port the T1 directly to a TDM service provider, essentially just leasing T1 service from another carrier and reselling the service. If the Ethernet network provider has a Class 5 circuit switch, then the T1s would be terminated on that switch.

In a greenfield environment in which a softswitch and PSTN media gateway are deployed for VoIP service, the T1s may be terminated on the PSTN media gateway.

The Figure 2 illustrates how a Head-End configuration of such a network.

SUMMARY

Future distribution networks must support legacy services such as T1 to be viable in the marketplace. If these networks are packet based, as is expected, technology must be deployed that will emulate a TDM T1 circuit. We have presented a straightforward method of providing such a service through the use of Time Division Multiplexing over IP (TDMoIP). Since the T1 stream is carried transparently in TDMoIP, any PBX signaling protocol or data format could be accommodated.

Careful control of prioritization through the use of existing standards can minimize delay and packet loss. The resulting T1 circuit emulation service can equal that of existing TDM technology, and easily meet the requirements of end-users.

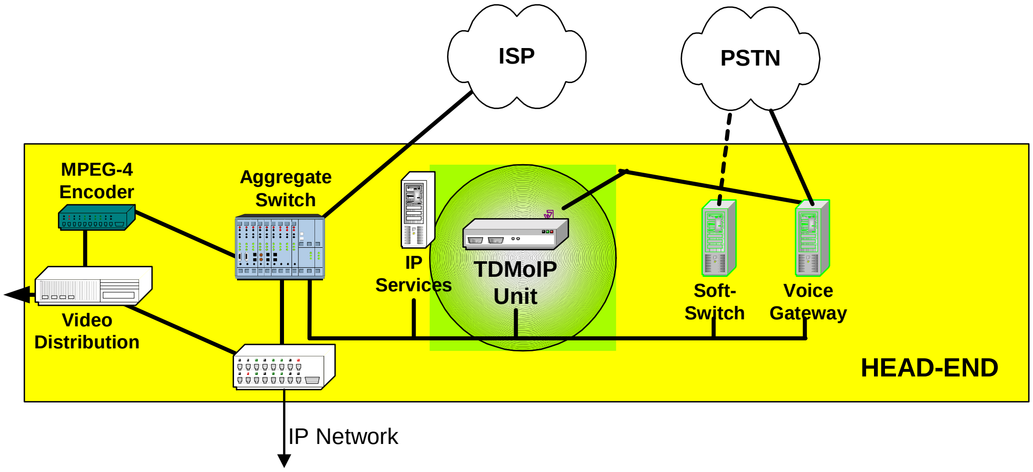

Figure 2: Greenfield Deployment Head-End Architecture

[1] 2000 Perspective on the Telecom Marketplace, Provisioning of Private Line and Frame Relay Services: A Global Perspective 1999-2004; Insight Corp.

[2] TDM over IP, Internet Draft August, 2001

[3] American National Standards Institute (ANSI) T1.403.1995

[4] ITU (CCITT) G.823 and G.824

VoIP SERVICES:PacketCable^{TM} DELIVERS A COMPREHENSIVE SYSTEM

David McIntosh
Maria Stachelek
CableLabs

Abstract

This paper provides a high-level overview of the PacketCable IP-based services delivery architecture as well as a more detailed discussion of the PacketCable Event Messages framework used to track access network resources used by those services.

The cable plant has experienced significant upgrades in the past several years allowing for the delivery of advanced broadband services. The DOCSIS^{TM} 1.1 specification defines the robust, highly-reliable, and highly-efficient broadband transport mechanism necessary to support time-critical services such as voice.

The PacketCable architecture is a multimedia services delivery platform, layered over the DOCSIS 1.1 access network, designed to support a wide variety of quality-of-service (QoS) enabled IP-based services. The end-to-end architecture as designed[1][6], offers a complete system that includes: device provisioning, signaling, event messaging, configuration management, QoS, and security. These services are managed by specific servers and network endpoints that collectively create a PacketCable network.

Voice over IP (VoIP) is the first service identified for delivery over the PacketCable architecture. Additional non-voice services are also being analyzed as candidate services for delivery over the PacketCable architecture. Examples of these beyond-voice services include multi-player gaming, videoconferencing, and unified messaging.

The PacketCable Event Messages framework supports collection of information necessary to create a PSTN-style call detail record (CDR) that may be used for purposes of customer billing, settlements, traffic analysis, and other back office functions.

ARCHITECTURAL OVERVIEW

One of the fundamental PacketCable objectives is to define a QoS-enabled, IP-based services delivery platform that extends the capabilities of the highly efficient DOCSIS 1.1 access network so as to allow cable operators to deploy a variety of IP-based services.

The initial service offering identified for the PacketCable architecture is residential Voice Over IP (VoIP). While the PacketCable architecture doesn’t preclude the delivery of small-office-home-office (SOHO) and business VoIP services, the focus has been on residential services allowing cable operators to provide value and to leverage relationships with their residential cable subscribers.

Several factors differentiate PacketCable VoIP services from traditional “IP telephony” services. For example:

- PacketCable VoIP is a phone-to-phone service rather than a personal computer-based telephony service.
- PacketCable services are guaranteed priority delivery on the DOCSIS access network ensuring a consistent, high-quality service.
- PacketCable services are not delivered over the public Internet. PacketCable mandates the use of a managed IP backbone that provides service delivery consistent with that of the DOCSIS access network.

The PacketCable architecture, pictured in Figure 1, provides the comprehensive system necessary to deliver VoIP services. When describing the architecture, we often talk about the three networks involved in the delivery of VoIP services: the access network, the managed IP backbone, and the public switched telephone network (PSTN).

Access Network - the HFC network connecting the subscriber to the MSO. The MTA and CM reside on the access network. The CMTS connects the access network to the managed IP network.

Managed IP Network - a high bandwidth IP network used to connect the MSOs headend servers. This network is often called the "Managed IP backbone" when it is used to interconnect several managed IP networks, DOCSIS HFC networks, or connect PSTN gateways to the PSTN.

PSTN - interconnects with the PacketCable Managed IP Network via a PSTN gateway.

FUNCTIONAL COMPONENTS

As part of the comprehensive end-to-end system necessary to deliver VoIP services, the PacketCable architecture requires several network elements with well-defined interfaces between those elements. This section describes several key functional components in the PacketCable architecture.

Multimedia Terminal Adapter

In the home, a standard phone plugs into the multimedia terminal adapter (MTA) allowing voice to be converted into IP packets. An MTA may be designed to be either a separate standalone device or to be embedded within the cable modem.

Cable Modem Termination System

The CMTS is responsible for managing access network resources for PacketCable services. Access network resources are first reserved when service is requested, then committed when service is delivered, and finally released when the service has completed.

Call Management Server

The CMS manages and maintains call state for VoIP services. The CMS is composed of a call agent (CA) and a gate controller (GC). The CA manages the call state and controls the MTA. The GC performs QoS admission control and communicates with the CMTS to allow services to obtain access network resources.

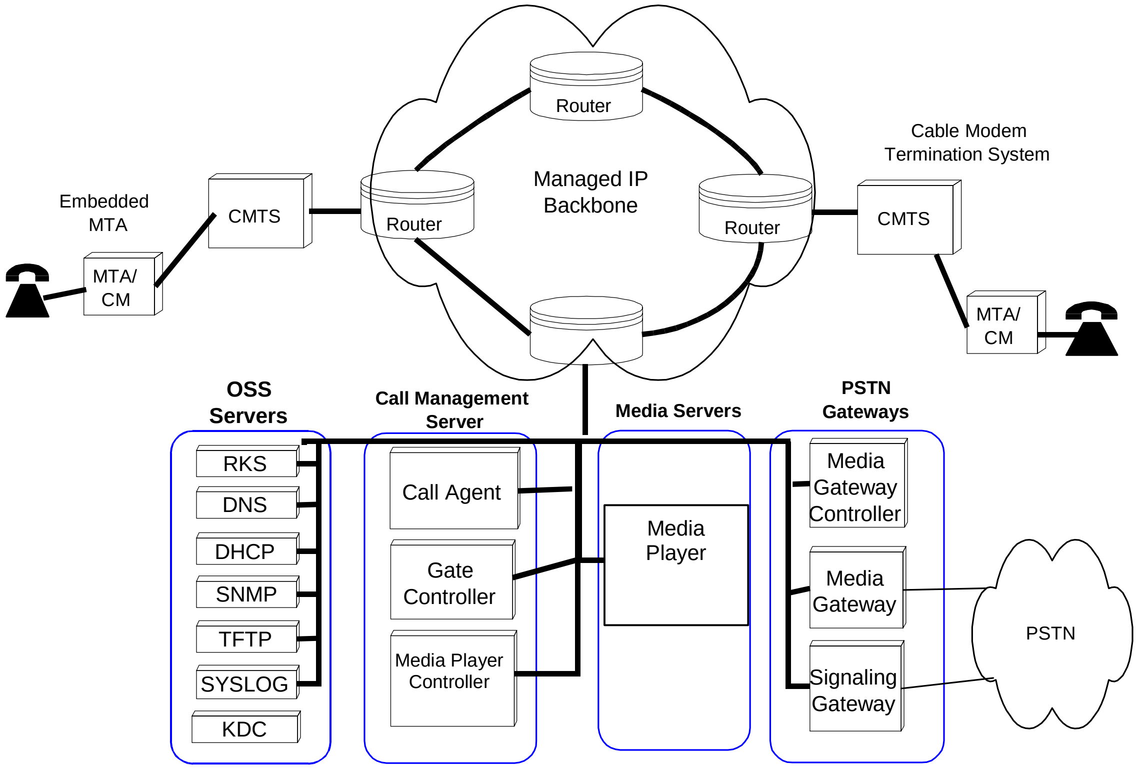

Figure 1: PacketCable 1.0 Network Component Reference Model

Record Keeping Server

The RKS is the short-term repository for PacketCable Event Messages. It receives messages from the CMS, CMTS and MGC then collates them into coherent sets for designated back office systems for additional applications processing. It serves as the interface between the PacketCable network and relevant back office servers.

Operational Support Systems

Operational support systems (OSS) contain a variety of supporting servers and infrastructure functions for such activities as provisioning, record keeping for billing, key distribution for security, domain name service (DNS) for name resolution, etc.

The OSS back office contains service and network management components supporting the core business processes.

Media Server

The media server provides network announcements and information, e.g. *"The number you have called is no longer in service…"* . There are two logical pieces to the Media Server, the Media Player Controller (MPC) and the Media Player (MP). The MPC requests the MP to play announcements based on call state as determined by the CMS. The MP is responsible for delivering the appropriate announcement to the MTA or to the MG.

PSTN Gateway

A PacketCable PSTN Gateway can be decomposed into a Media Gateway Controller (MGC), Media Gateway (MG), and Signaling Gateway (SG). The MGC manages the interconnection with the PSTN by controlling the MG and SG. The MGC is responsible for maintaining the call state for calls requiring PSTN interconnection.

The PacketCable architecture also supports a hybrid gateway solution that takes advantage of legacy PSTN switches already owned by cable operators while providing a migration path to the fully IP-based PacketCable solution. This hybrid architecture uses PacketCable IP-based components on the access network and circuit switch call control derived from a PSTN local digital switch location.

ARCHITECTURAL CAPABILITIES

Several core capabilities are fundamental to the delivery of VoIP services on the PacketCable architecture. This section provides a high-level discussion of Dynamic QoS and Security. Event Messaging is also a core capability that is discussed in more detail in a later section.

Dynamic QoS

An IP-based network, by definition of the underlying TCP and UDP transport mechanisms, delivers packets in a best-effort manner. Dropped or delayed packets result in unpredictable end-to-end throughput.

PacketCable and DOCSIS 1.1 provide a comprehensive, integrated QoS delivery mechanism [4] that ensures PacketCable packets are delivered in a guaranteed manner, not a best-effort manner.

PacketCable splits the management of QoS resources into access network segments and backbone network segments. This approach allows for different bandwidth provisioning and signaling mechanisms for different network segments: the origination side, the far end, and the backbone network. Additionally it allows for resource-constrained segments to manage resource usage and maintain per-flow reservations carefully. The PacketCable DQoS Specification details this design [3].

Security

PacketCable security spans all interfaces in the PacketCable architecture [5]. It provides confidentiality for media packets and for signaling communication across the network via authentication, encryption, and key management. It ensures that unauthorized message modification, insertion, deletion and replays anywhere in the network are easily detectable without affecting network operation. Security is interface specific, but the majority of signaling interfaces are secured using IP security (IPSec). The media stream is secured by encrypting and authenticating the payload directly.

In addition to defining the security protocol that will be applied to each interface, PacketCable also defines a corresponding key management mechanism. There are three basic key management mechanisms defined for use in PacketCable: Kerberized Key Management, internet key exchange (IKE) with either pre-shared keys or X.509 digital certificates, and randomly generated keys exchanged within secured signaling messages.

OPERATOR OBSERVATIONS

The multiple vendor and service provider environment into which a cable operator enters with the delivery of VoIP services through a distributed architecture requires on-going attention. It has been observed that this is not necessarily one-stop shopping and each cable operator will be required to develop their own management strategies to navigate the variety of vendor, network provider, and service provider relationships. Take for example, resolving a fault management issue. These can arise within any number of scenarios, each with their own respective source and solution. Having the expertise in place, whether in-house, or supported by a third party, to deal with such an issue will be important.

Another example of managing the multi-service provider environment can be found in the issue of local number portability (LNP). In those instances where a PSTN customer would like to keep their existing telephone number as part of their cable VoIP service, LNP comes into play. The interval of porting a subscriber's telephone number has to be managed closely in order for the subscriber to experience a seamless service transition.

PACKETCABLE EVENT MESSAGES

The PacketCable architecture provides a QoS-enabled IP-based service delivery platform for voice and other multimedia services. The PacketCable Event Messages framework provides a mechanism for tracking access network resources that have been requested and consumed by these services. This information can be used by back office systems for many purposes including billing, settlements, network usage monitoring, and fraud management [2].

The PacketCable Event Messages framework has been designed to be flexible and extensible enough to support the initial suite of PacketCable voice services, as well as accommodate beyond-voice services in the future.

A single PacketCable Event Messages framework has been defined to support a variety of service-delivery scenarios and network topologies. For example, tracking information for services that either originate or terminate on the PSTN, as well as services that stay on the MSOs network are supported by the framework.

Event Message Information

An Event Message is a data record containing information about usage and service activities. Telephone number is an example of the type of information carried in an Event Message. An event-based format is necessary to accommodate the distributed architecture where complete “session state” no longer resides in one or two network elements, but is instead spread across any of these, i.e. CMS, CMTS, and MGC.

A single event message may contain a complete set of data regarding usage or it may only contain part of the total usage information. When correlated by the RKS, information contained in multiple event messages provides a complete record of the service. Event messages are collected and are sent to one or more back office applications such as a billing system, a fraud detection system, or a pre-paid services processor.

Originating/Terminating Model - PacketCable makes use of an "originating/terminating model" based on the PSTN "half-call model." In this model, the originating party's service provider is responsible for tracking information sufficient to bill the originating party for service, and to settle with the terminating provider. The terminating party's service provider has the same responsibility for the terminating party. This "originating/ terminating model" supports the various PacketCable network topologies.

Batch vs. real time - PacketCable allows Event Messages to be sent to the RKS as they are generated. Alternatively, once generated, Event Messages may be stored on the CMS/CMTS/MGC and sent to the RKS in a single file.

Call Detail Records - Using the unique billing correlation ID (BCID) assigned to a given call, the RKS collects all the individual Event Messages for that call, and assembles them into a single call detail record. The format of the CDR may be AMA, BAF, IPDR, or any format appropriate for the billing and other back-office servers that will make use of the information.

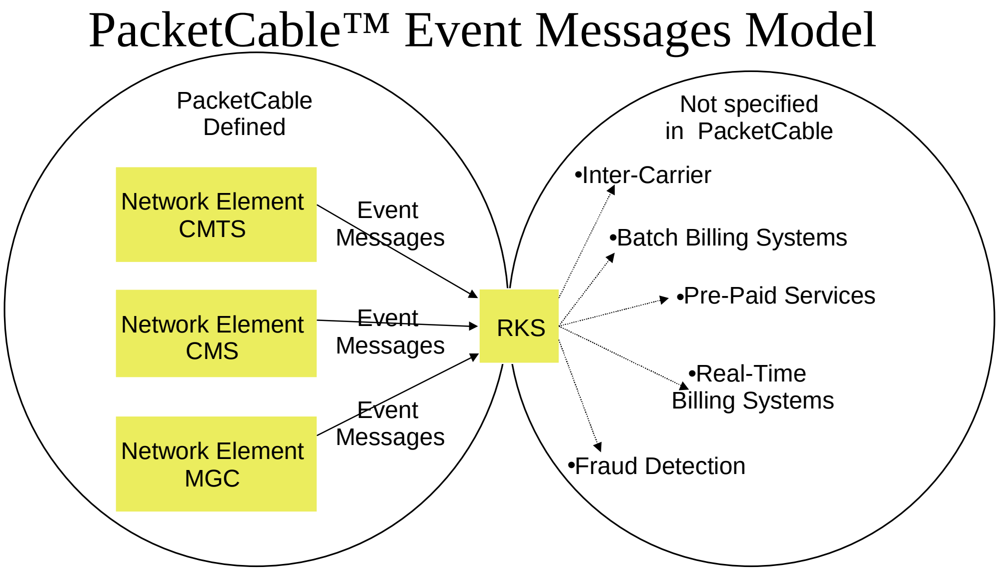

Figure 2: PacketCable Event Message Model

Figure 2 depicts the PacketCable Event Message architecture. By standardizing the transport, syntax, and collection of appropriate Event Message attributes from a distributed set of network elements (CMS, CMTS, MGC), this architecture provides a single repository (RKS) to interface with billing, settlement, reconciliation, and other systems.

The CMS, MGC, and CMTS generate Event Messages for the portion of the communication pertaining to them. For example, the CMTS generates a "start of

QoS" message, when the CMTS commits access network resources to a PacketCable service.

SPECIFICATIONS AND STANDARDS

The comprehensive nature of the PacketCable architecture is the result of a suite of Technical Reports and Specifications that delineate the end-to-end architecture and associated interfaces for a complete IP-based services delivery platform. These Technical Reports and Specifications (available at www.PacketCable.com) have been accepted as standards by several North American and International standards organizations including the Society of Cable Telecommunications Engineers (SCTE), American National Standards Institute (ANSI), and the International Telecommunications Union (ITU).

CONCLUSION

The PacketCable architecture described in this paper is a comprehensive end-to-end system necessary to deliver VoIP and other IP-based multimedia services.

For the delivery of VoIP, the PacketCable architecture can be thought of as three networks coordinated through a collection of functional components and servers. The PacketCable architecture supports several core capabilities, such as dynamic QoS and security, that are fundamental to the efficient, reliable deliver of IP-based services.

Efforts are underway to develop extensions to the PacketCable architecture to support a wide range of IP-based multimedia services.

The PacketCable Event Messages framework is a flexible and extensible model that supports subscriber billing, settlements, and other back office functions. Going forward, the PacketCable Event Messages framework will be expanded to keep pace with a wide variety of IP-based services beyond voice that will be delivered over the PacketCable service delivery platform.

REFERENCES

[1] *"PacketCable 1.0 Architecture Framework Technical Report*", PKT-TR-ARCH-V01-991201, December 1, 1999, CableLabs, www.packetcable.com

[2] *"PacketCable Event Messages Specification*," PKT-SP-EM-I03-011221, December 21, 2001, CableLabs, www.packetcable.com

[3] "*PacketCable Dynamic Quality-of-Service Specification*,"PKT-SP-DQOS-I03-020116, January 16, 2002, CableLabs, www.packetcable.com

[4] *"Quality-of-Service: A DOCSIS/PacketCable™ Perspective"*, Venkatesh Sunkad and Majid Chelehmal, Proceedings of SPIE Volume: 4522, pgs. 87-98, July 2001. www.spie.org

[5] *"PacketCable Security Specification,"* PKT-SP-SEC-I05-020116 January 16, 2002, CableLabs, www.packetcable.com

[6] *"The PacketCable Architecture"*, Ed Miller, Flemming Andreasen, and Glenn Russell, IEEE Communications Interactive, June 2001, www.ieee.org

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