Cable Headend Architecture for Delivery of Multimedia Services

Mukta Kar, Majid Chelehmal, and Richard S. Prodan
Cable Television Laboratories

Chezhian Renganathan
HCL Technologies Private India Pvt. Ltd.

Since cable headends traditionally deal with analog audio/video (AV) television signals, analog transmission service opportunities are limited. With analog technology, it is difficult to maintain interoperability and, as such, headend equipment can be proprietary in nature, as well as expensive and sometimes difficult to replace. Digital technology, combining standards-based compressed audio/video/data and digital transmissions, is more efficient and flexible. By implementing digital technology, opportunities to deliver new services over hybrid fiber coaxial (HFC) cable networks greatly increase, along with a proportional increase in revenue.

Additional services will add complexity at the headend. One of the benefits of standardized digital technology is that, through interoperability, distinct services can be integrated across system components from multiple suppliers. ISO/MPEG has provided standards for compression, storage, and digital AV transport; SMPTE has established standards for interfaces used in baseband digital AV equipment; and IETF has provided standard protocols for IP-based applications. Headend equipment may be designed with an integrated, flexible analog/digital architecture so that new standardized equipment may be added incrementally to enhance existing services or to add new services.

An integrated, flexible architecture with coexisting analog and digital services is presented in this paper. This architecture allows the enhancement of existing services or the addition of new services with minimal disruption to headend functions. Interface design issues relating to analog AV signals and multimedia signals, such as DTV, HDTV, and IP-based streaming is discussed. Additionally, information is included on digital program insertion systems for adding locally compressed content (e.g., advertising) to remotely distributed compressed programming, as well as processing of IP/Internet traffic at the headend for delivery over HFC cable networks using cable data modems.

INTRODUCTION

Until a few years ago, the broadcast industry (cable, off-air, and satellite) was limited to the broadcast of television programming. Computers were used for scientific and industrial purposes, and the telecommunication industry provided end-to-end delivery of voice and data. Emerging technologies are blurring the boundaries between the telecommunications, television/film, and computer industries. Video, sound, and communications have migrated into computers; television is incorporating the Internet; and video, audio, and interactivity are being integrated into telecommunications. All this leads to the convergence of many different applications into what is known as "multimedia." The shaded region in Figure 1 is a graphical representation of the area covered by multimedia applications.

Multimedia has become a household word. What is multimedia? Today, the media we use to transmit information is in the form of data (e.g., text), digitized video, digitized audio, still digitized images, and/or graphics. Multimedia (multiple media) refers to the combination of two or more of these media, at least one of which is discrete (e.g., text, images) and one is continuous (e.g., video, audio).

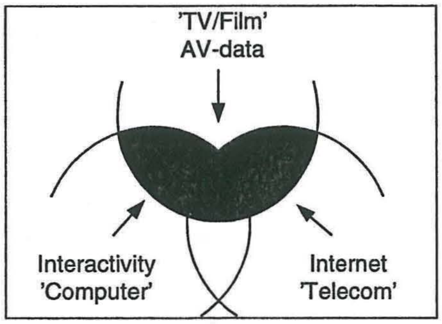

Figure 1: Multimedia Application Area (shaded)

Rapid advances in technology are producing new multimedia applications every day, opening up opportunities for traditional industries to offer new applications and services and, thus, increase revenue. The same is true for the broadcast industry. Digital technology presents a significant opportunity. Advances in video compression technology, coupled with digital modulation techniques, have opened the door to new multimedia services and applications. Using digital modulation techniques such as Quadrature Amplitude Modulation (QAM), one 6-MHz analog channel can be expanded to a digital bandwidth of 27 Mbits/sec (64 QAM) or 38 Mbits/sec (256 QAM). This QAM-modulated channel can be used to deliver either a multiplex of a few compressed TV programs or a number of other multimedia services. An analog TV program carried over a 6 MHz channel can be compressed to a 2- to 4-Mbits/sec (average) stream. Using one 6-MHz spectrum, a few analog programs can be delivered to homes with similar or better picture quality. With the size of the HFC network's total digital bandwidth, there exists the potential to deliver, in addition to analog and digital TV services, a large number of multimedia services and applications.

SIGNAL CATEGORIES

Cable systems began compressed digital programming services in 1996. Off-air broadcast of SDTV and HDTV signals began in November 1998. Per FCC guidelines, broadcasters will have to retire analog television broadcasting and completely switch to digital broadcasting by the year 2006. To be competitive, as well as to take advantage of digital bandwidth, the cable industry will add HDTV programming as well. The industry also is preparing to utilize this new potential to deliver multimedia services—IP-based services, such as e-commerce, world-wide-web, networked interactive games, video-on-demand (VOD), and others—over the HFC network. Adding all of these new services in a short period of time could be quite expensive and may not be justified economically. However, the headend may be upgraded with an integrated, yet flexible, analog-digital architecture so that new standardized equipment may be added incrementally to enhance existing services, or to add new services.

To deliver services, such as SDTV, HDTV, or any other multimedia service, the headend, the HFC cable network, and consumer premises equipment (CPE) will require different upgrade levels. For example, interactive services need a higher degree of plant upgrade than the one-way delivery of digital television programming. In particular, the integrated analog/digital headend architecture must be modular so that the enhancement of existing services, or the addition of new services, can be achieved incrementally as the resource/revenue stream permits. In this paper, we deal primarily with flexible and scalable headend architecture, which will allow incremental addition of services and, in particular, with signal interfacing and signal processing at the headend before delivery over the HFC cable network.

Signals interfaced at the headend may be categorized broadly into two types: existing analog signals and digital signals. Analog signals are already interfaced and integrated. Upgrading cable headends to interface new digital services should cause minimal or no disruption to the existing analog services. All digital signals or services can be divided into two types: MPEG-2-based digital television services and IP-based services. Some headends have been delivering digital television programming for several years. HDTV programming and the IP-based services now need to be integrated into the existing analog-digital headend. Below, we explore the tools and equipment necessary to interface these signals at the headend, as well as methods used to process them before delivery over the HFC network.

INTERFACING DTV SIGNAL SOURCES

Remultiplexers and rate-remultiplexers may be used in interfacing MPEG-2-based digital AV signals at the headend. A multiplexer is a digital device that multiplexes more than one input sequence to create a single serial digital signal (a multiplex). The multiplex bitrate will be greater than, or equal to, the sum of all the input sequences' bitrates.

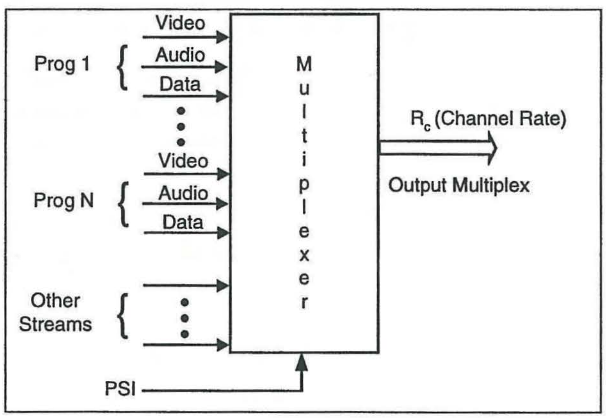

Figure 2: Multiplexer Block Diagram

Figure 2 shows a typical MPEG-2 multiplexer where MPEG-2-based elementary streams (ES) of one or more programs are packetized (188 bytes/packet) and multiplexed together to form a baseband transport multiplex compliant to ISO/IEC 13818-1 specifications. It may be noted that the multiplexer also includes PSI along with the multiplex so that the downstream equipment can demultiplex the baseband transport streams to retrieve the original constituent ES. The technology of remultiplexing is similar to that of multiplexing, but it takes one or more transport stream baseband multiplexes as input and constructs a new multiplex with the selected programs from the input multiplexes, thereby locally producing a single multiplex from multiple remote sources.

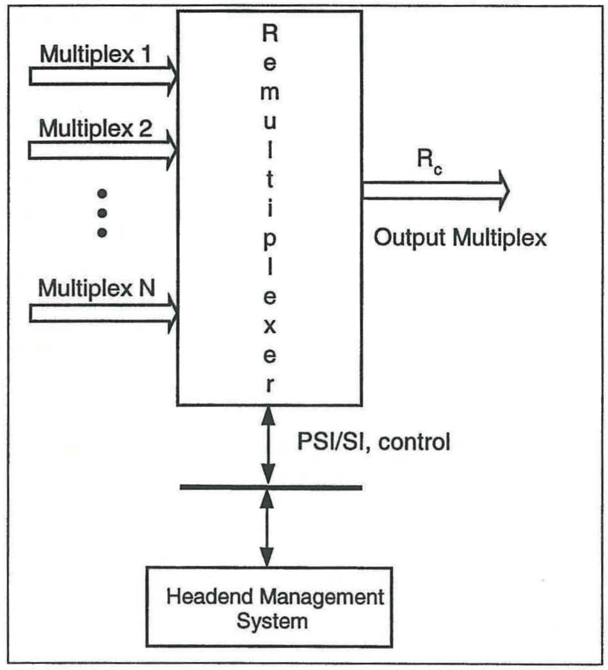

Figure 3: A Simple Remultiplexer

Figure 3 shows a typical remultiplexer. Again, the remultiplexer has the same bitrate constraints as those of a multiplexer—the sum of the ES bitrates in the selected programs cannot exceed the given bitrate of the output multiplex or channel rate. In brief, a remultiplexer does the following:

- It demultiplexes the input multiplexes and extracts the service information (SI), including program-specific information (PSI), from the incoming multiplexes;
- Constructs a new multiplex out of user-selected programs and obeys the bitrate constraint;
- It also includes appropriate PSI/SI with the output multiplex;
- Removes jitter in program clock reference (PCR) values and ensures AV synchronization within the programs;

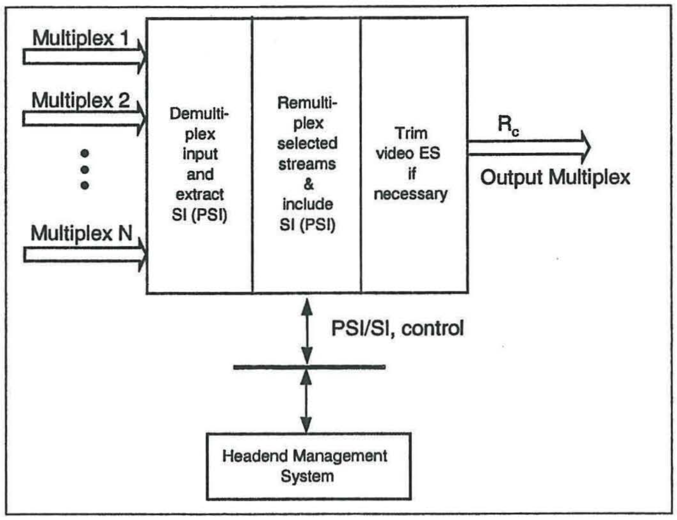

Figure 4: A Typical Rate-remultiplexer

- It provides perceptually seamless switching capability from one program to the other without any audible or visual artifacts.

An important difference between a multiplexer and a remultiplexer is that a multiplexer takes ES as input, whereas multiplexes of compressed ES are the input to a remultiplexer. In either case, ES may be of constant bitrate (CBR) or of variable bit rate (VBR).

RATE REMULTIPLEXER AND TRANSCODING TECHNOLOGY

A rate-remultiplexer (Figure 4) is a state-of-the-art piece of equipment that combines the existing remultiplexing technique with a new technology called transcoding. A standard remultiplexer does not change the bitrate of the ES, but the rate-remultiplexer can. If the sum of the bitrates of all the ES in the selected programs exceeds the given bitrate of the output multiplex, the rate-remultiplexer will reduce the bitrate of the video ES until the bitrate of the output multiplex is less than, or equal to, the available channel rate (27 Mbps for 64 QAM and 38 Mbps for 256 QAM). The rate-remultiplexer achieves this capability by implementing transcoding together with joint rate control and statmux scheduling. The latter ensures a MPEG-2 transport-compliant multiplex where video ES comply with the T-STD buffer model—no overflow or underflow occurs at the decoder buffer. Program packaging from input multiplexes for delivery over a 6-MHz channel is no longer limited by the bitrate constraint of a standard multiplexer or remultiplexer. Hence, the real advantage of a rate-remultiplexer is that it provides wider freedom in packaging a number of programs from input multiplexes into a new multiplex with a given channel rate. It is an important tool in maximizing the utilization of digital bandwidth in delivering a desired number of TV programs and, thus, may help free up some channels for delivery of other services. This rate-reducing capability provides opportunities to increase revenue with a given number of 6-MHz channels capable of carrying digital signals. The rate reducing capability is also useful in digital program insertion (DPI) and video-on-demand (VOD). The contents for DPI and VOD can be encoded with higher bitrates (e.g., 6 Mbps) than average delivery bitrates (e.g., 4 Mbps), and can be stored in the server. However, to meet the bitrate constraint, the bitrate of the stored stream will be trimmed along with other video streams when multiplexed. The advantage here is that the DPI and VOD contents do not have to be stored in different bitrate versions and it is not necessary to find a close match between the bitrate of the stored content and that of the channel [1].

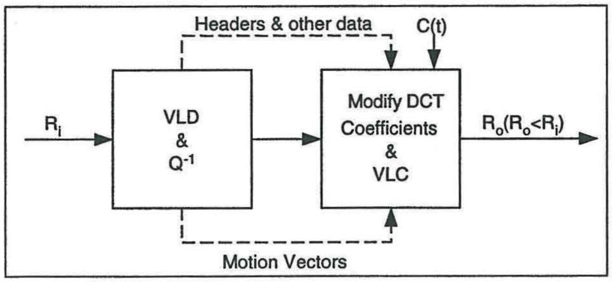

Figure 5: Open-loop Transcoder

Figure 5 is a typical open-loop rate reducer that may be implemented in the compressed domain with minimal decoding and computation. R_i is the bitrate of the input compressed video, R_o, the rate of transcoded video and c(t) may be considered as the user's input constraints. In MPEG-2 compression, the majority of the bits are used in representing the DCT coefficients of an 8 x 8 block of pixels (and there are many such blocks in a picture). The saving in bits may be achieved by reducing the number of these non-zero DCT coefficients or coarsely quantizing them. The disadvantage of the open loop scheme is that it provides little control over the picture quality of the output stream. In that respect, the closed-loop method of transcoding is significantly better where error is estimated between the picture of the input stream and that of the transcoded stream. If the error is larger than the desired threshold, the picture is re-encoded to reduce error. However, closed loop implemetation is more complicated than open loop. Irrespective of the implementation method—open loop or closed loop, pixel domain or DCT domain—the motion vectors (MV) available with the input stream should be used as much as possible in transcoding a picture (calculation of MVs is a very computationally intensive process). Transcoding of a video ES is usually associated with some degradation in picture quality. The degradation is related to the percentage in bitrate reduction. See [2] and [3] for more details on the technology of transcoding.

INTERFACING IP-BASED SERVICES

In providing digital TV services, remultiplexers (or rate-remultiplexers) are important pieces of equipment for interfacing network signals with the headend. Similarly, in providing IP-based services, a cable modem termination system (CMTS) is equally important. A CMTS and a cable modem (CM) make the entire cable network transparent to the user (Figure 6). The CMTS' and the CM's principle function is to transmit IP packets between the external IP network and the subscriber location (user), transparent of the cable network and the headend. For example, when a user uses an IP-based service via a PC (connected to a CM), the user is not aware of the cable network. The CMTS serves as the interface between the external IP network and the headend and the CM, between the cable network and the Internet appliance (e.g., a PC).

To provide IP-based services, the CMTS works as the interface between the headend and the external IP network, and may operate as a router or a bridge.

Routing and bridging are methods of moving information across an inter-network between a source and a destination. The primary difference between the two is that bridging occurs at Layer 2 (the link layer) of the open systems interconnection (OSI) reference model, while routing occurs at Layer 3 (the network layer). This distinction provides routing and bridging with different information to use in the process of moving information from source to destination. As a result, routing and bridging accomplish their tasks in different ways.

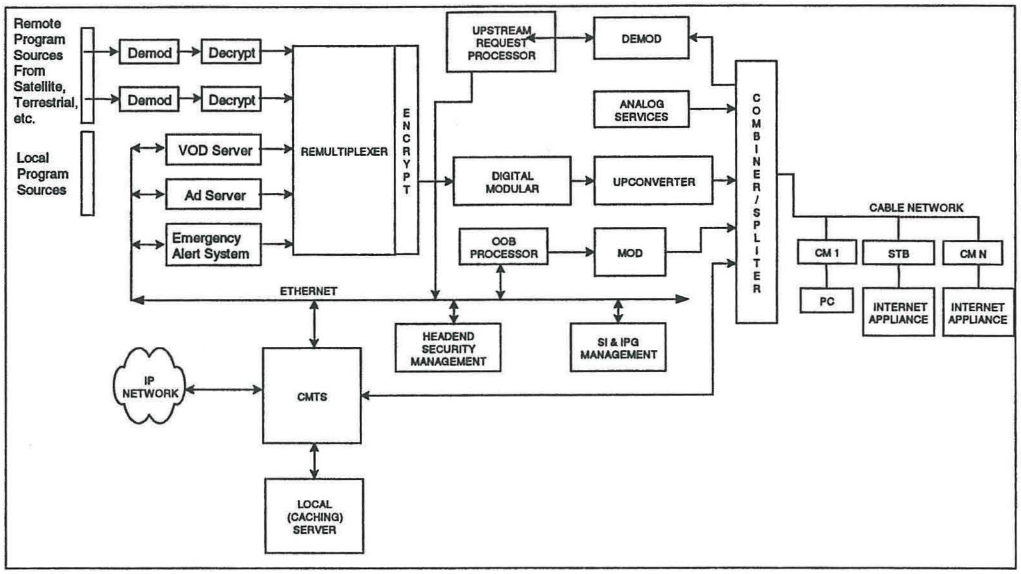

Figure 6: Headend Architecture Block Diagram

Routing involves two basic activities: determination of optimal routing paths and the transport of information groups (typically called packets) through an inter-network, the latter usually referred to as switching. Bridging and switching occur at the link layer, which controls data flow, handles transmission errors, provides physical (as opposed to logical) addressing, and manages access to the physical medium. In general, bridges and switches are not complicated devices. They analyze incoming frames, make forwarding decisions based on information contained in the frames, and forward the frames to the destination. A router looks at the logical address (e.g., the destination address in an IP packet) to decide on the routing path, while the bridge looks at the physical address (e.g., the destination MAC address) to make the decision. We consider only IP-based cable networks here.

As a bridge, the CMTS does not need to open the IP header, it simply forwards the incoming frame from the network side interface (NSI) to the radio frequency interface (RFI) on the downstream. It refers the destination media access control (MAC) address (Figure 7) of the incoming frame when deciding the outgoing interface (called a port) in the case of a CMTS with more than one interface (shown in dotted lines in Figure 9). The bridge CMTS has to maintain a MAC address versus a port mapping table to make the decision. It could build such a table by gleaning the source MAC address of frames (Figure 8) coming into a port, thus associating that port with the MAC address for future reference, a processes called "learning."

In the case of a router CMTS, packets go up to the IP layer (Figure 8). Routers in the Internet run a host of "routing protocols" which gather information about the network topology and maintain "routing tables." A routing table contains information for reaching a particular network. Typically, it could be the address of the next hop to reach a particular network. The last hop router (the router connected to the destination device) will just look at the network part of the IP address to decide on the outgoing interface. A router CMTS could just happen to be the last hop router for downstream, CPE destined packets (refer to Figure 10). Routers replace the incoming frame's destination MAC address with that of the next hop interface, while a bridge does not modify the MAC or IP addresses on transit.

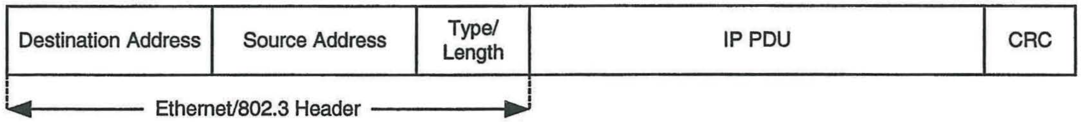

Figure 7: MAC Frame

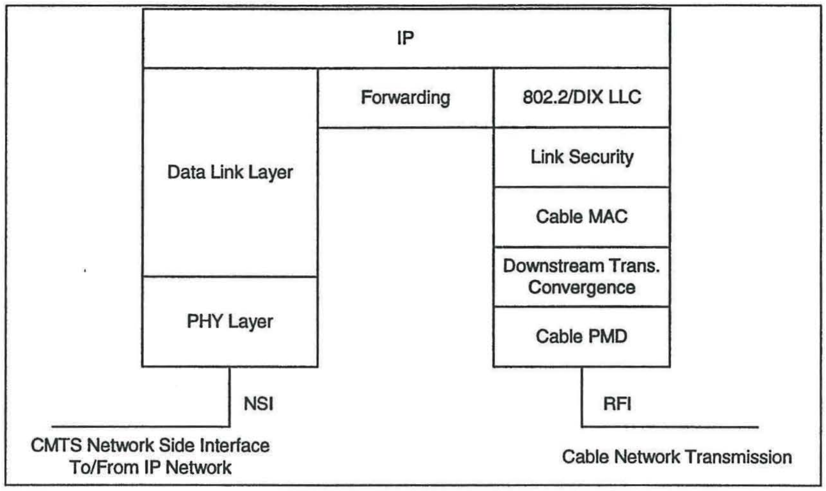

Figure 8: CMTS Protocol Stack

This decision-making process is necessary even in the case of a CMTS with two interfaces (one NSI and one RFI) if the CMTS is to stop packets not intended for the cable network from coming in from the data network (Internet).

As a router, the CMTS opens the IP packet header to check the destination IP address. Based on the network portion of the IP address, the CMTS decides on the interface in which to forward the incoming data. In such an implementation, each of the CMTS interfaces will have to be different IP networks (the network portion of the IP address will be different), while this need not be the case for the bridge CMTS.

If the CMTS finds that the network address portion of the IP destination address (present in the IP header) of the incoming data packet matches the network portion of the interface IP address of any of the CMTS interfaces, the packet is forwarded onto that interface. When the CMTS acts as a router, the incoming frame's destination address (DA) would be that of the network interface card (NIC) attached to the data network. Before forwarding on to the RFI interface, the DA address field in the MAC header has to be replaced with that of the destination's CPE (such as a PC) MAC address. For example, referring to the stack implementation in Figure 7, after stripping the incoming encapsulation, a new 802.2/DIX LLC encapsulation is performed on the RF interface.

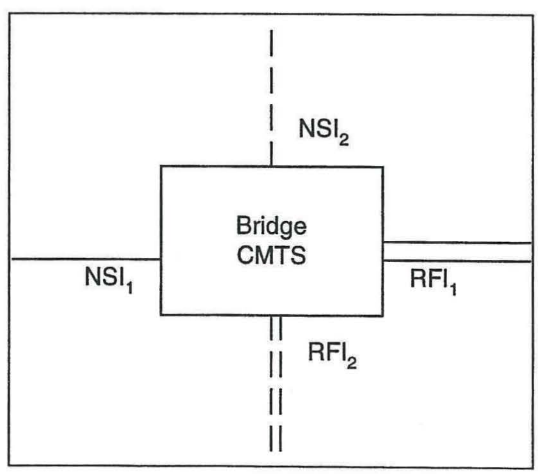

Figure 9: A Bridge CMTS

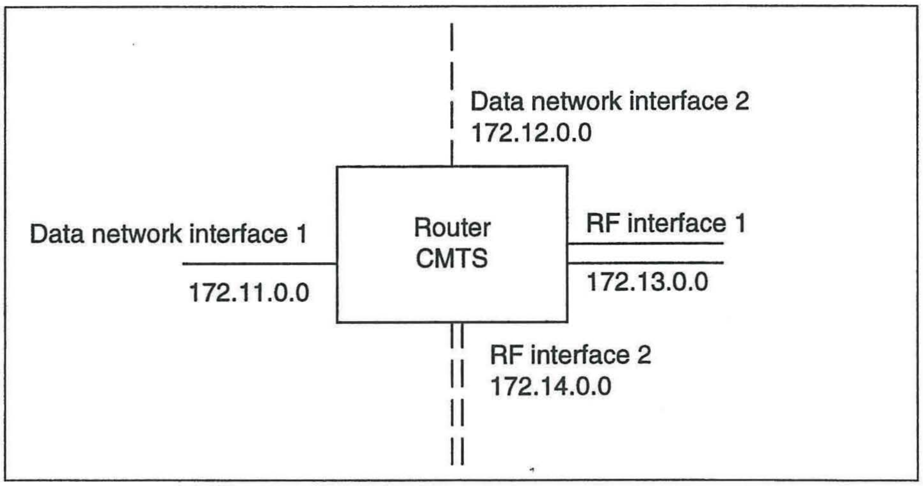

Figure 10: A Router CMTS

As a router, the CMTS could have network implementation with each of the supported interfaces attached to different networks. Therefore, handling IP packets at the CMTS entirely depends on the overall design plan of the network.

FORWARDING TO THE RFI ON THE DOWNSTREAM

Handling IP packets (router CMTS) or Ethernet/802.3 frames (bridge CMTS) on the RF interface before transmission over the cable network depends on the standards to be followed for interaction between the CMTS and the CM. Figure 8 shows the implementation recommended by the data over cable service interface specification (DOCSIS) 1.0 SP-RFI specification. In DOCSIS networks, the CPE is any device with an Ethernet NIC card. Typically, it is a PC used to connect to the IP network for services such as email, web access, etc. The functions of each layer on the RF interface are given below. Each of these layers will encapsulate incoming data.

Logical Link Control (LLC) Sublayer: The LLC conforms to Ethernet standards to provide the data link framing that the CPE's NIC needs. The following layers implement the security, modulation, spectrum management, and analog signal transmission. The functionality of these layers could be mapped as needed to other standards that define communication between the CMTS and the CM.

Link Security: This sublayer is needed to support the basic privacy needs, authorization, and authentication.

Cable MAC Layer: In this layer, the CM removes encapsulation before forwarding data packets to the CPE. It has a set of predefined messages (packets) which direct the behavior of the CMs on the upstream channel, as well as messaging from the CMs to the CMTS.

Some of the main functions of this layer include:

- CM bandwidth allocation for upstream transmission in terms of number of mini slots;
- Class-of-service support using dynamic SID allocation;
- Wide-range data rate support;
- Other features (refer to DOCSIS RFI specification).

Downstream Transmission Convergence Sublayer: This layer exists in the downstream direction only. It provides the means for delivery of additional services (e.g., digital video) over the same physical layer bitstream. This sublayer is defined as a continuous series of 188-byte MPEG transport stream packets. The PID (packet identifier), one of the header fields, identifies the type of payload it is carrying.

Cable Physical Media (PMD) Sublayer: This layer involves digitally modulated RF carriers on the analog cable network. Thus, a data packet from the Internet goes through a series of encapsulations and de-encapsulations to carry control information for each of the layers as overhead, before the IP data finally reaches the destination CPE.

For example, when a user accessing an Internet web-server from a PC (CPE) is connected to a CM by Ethernet, Hyper Text Transfer Protocol (HTTP) packets are transparently taken through the cable network. The CM encapsulates the packets inside a cable MAC header and sends them to the CMTS where they are de-encapsulated and routed out through the correct interface onto the Internet. Similarly, incoming HTTP packets are forwarded onto the cable network after cable MAC encapsulation at the CMTS RF interface. Again, the CM de-encapsulates the cable MAC header at its RF interface and bridges (the DOCSIS CM acts like a bridge) the packets to the attached CPE device. The CMTS replaces the incoming frame's destination MAC address with that of the CPE's. The CMTS could obtain the CPE's MAC address by maintaining the Address Resolution Protocol (ARP) table, which is simply a mapping of IP addresses versus NSI MAC addresses.

A CM can support more than one IP-based CPE, but will have a limit on the number of CPEs to which it can be connected. In that case, the CM should have some mechanism that should filter out packets not destined to any of its attached CPEs.

INTEGRATED ARCHITECTURE

Figure 6 shows a headend architecture where analog TV services, digital TV services, IP-based services and other interactive services have been integrated. The digital baseband signals, after appropriate processing, will be digitally modulated on a 6-MHz radio frequency (RF) carrier, then will be upconverted to desired cable channels. Finally, all the RF channels will be combined for delivery over HFC cable networks. For interactive services or IP-based services, a return path, an upstream channel or a telephone return to the headend will be required. The return signals have to be processed at the headend and user requests have to be parsed and sent to the appropriate servers or other headend equipment. The ad server is used for storing and playing local commercials; the VOD server is used for VOD services. To achieve flexibility in packaging user-selected programming and to minimize the storage requirement of compressed commercials and VOD content, a rate re-multiplexer is a better choice over a standard re-multiplexer.

FUTURE INTERFACE ISSUES

Since the MPEG-2 standard is more suited to the broadcast industry, it does not cover other application areas. To compliment MPEG-2, the MPEG-4 standard is under development and will cover wider application areas. The MPEG-4 standard will be completed in two phases; version I is already under final distribution international standard (FDIS), and version II is scheduled to be completed by the year 2000. MPEG-4 is known as the multimedia standard, and it is anticipated that new content will be created using the MPEG-4 standard in the coming years. The cable industry could deliver that content over HFC cable networks. The MPEG-2 standard has a transport protocol suitable to carry MPEG-2 compressed video, audio and data. MPEG-4 does not specify any such transport protocol. Efforts are underway to deliver MPEG-4 content using other protocols (e.g., IP, MPEG-2, etc.). The MPEG-4 standard will not be compatible with MPEG-2. An MPEG-2 set-top box will not be able to decode MPEG-4 content, but can ignore it in a backward compatible way using MPEG-2 transport encapsulation. Initially, most of the MPEG-4 content is expected to be designed for the Internet and, hence, can be delivered like any IP-based service. Consumers can decode this content using their own MPEG-4 terminals. In the future, it may be necessary to enhance MPEG-2 set-top boxes so that they can decode MPEG-4 content.

CONCLUSION

We have explored the various services and applications that may be delivered over the HFC cable network. These services include existing analog TV services, digital television services, multimedia services, and interactive applications. It was shown that by using digital technology, the HFC cable network can deliver a majority of these services and applications. These services and signals may be classified into three broad categories: analog TV signals, digital TV signals, and Internet (IP-based) signals. Analog TV signals are well understood and exist in all headends. Interfacing the other two signals at the headend should cause minimal or no disruption to analog services.

The equipment used to interface digital TV signals at the headend is known as a remultiplexer. A remultiplexer is not capable of altering the video ES bitrate of the input multiplexes, and video ES require a major portion of the multiplex' bandwidth. An enhanced remultiplexer, known as rate-remultiplexer, incorporates transcoding technology. Transcoding is a technique to reduce the bitrate of the MPEG-2 compressed video elementary streams without fully decoding and re-encoding, which is a complex and potentially perceptually degrading process.

Similar to the remultiplexer, the CMTS plays a major role in interfacing Internet signals to the headend. The CMTS and the CM make the HFC cable network user transparent. How the CMTS handles IP packets between the external IP network and the HFC cable network was explained. The protocol stack used by the CMTS as a router or as a bridge was discussed. How a CMTS may be used depends on individual plant design.

Finally, an integrated headend architecture was presented where DTV services and multimedia services coexist with analog TV services. Interface issues related to content developed using MPEG-4, the multimedia standard, also was discussed.

LIST OF REFERENCES

1. *Rate-remultiplexing: An Optimum Spectrum Utilization Technology*, by Richard S. Prodan, Mukta L. Kar, and Majid Chelehmal, presented at National Cable Television Association Conference, June 1999.

2. *Video Transcoding for Matching Network Bandwidth Constraints*, by Pedro A. A. Assuncão, a Ph.D. thesis, Department of Electrical Engineering, University of Essex, U.K., 1998.

3. Data Over Cable System Interface Specifications – SP-RFI-I04-980724 (http://www.cablemodem.com/).

4. *Internetworking with TCP/IP – Vol 1,* by Douglas E. Comer, Prentice Hall, 1995.

5. *Digital Video: An Introduction to MPEG-2*, by Barry G. Haskell, Atul Puri and Arun N. Netravali, Chapman and Hall, 1997.

6. CODING OF MOVING PICTURES AND ASSOCIATED AUDIO, ISO/IEC 13818-1 (Systems), ISO/IEC 13818-2 (Video).