OPERATIONAL MANAGEMENT OF DIGITAL CONTENT

Yvette M. Gordon
Time Warner Cable

Abstract

In the development of digital services such as high speed data access, near video on demand, and video on demand, the management of digital content becomes critical. Most content, consisting of compressed images, MPEG video/audio, text, and data will require frequent turnover. This paper addresses hardware and software solutions to enable successful content manipulation in a digital environment.

THE DIGITAL ENVIRONMENT

Content Management Today

Management of content in an analog environment has been centered around the hardware for receiving, encoding, scrambling, amplifying, etc. video signals as depicted in Figure 1. With the addition of digital services, our focus is changing. No longer is content just transmitted, but now it must be stored and transported on demand. We now have to ask questions such as how much disk storage do we need, what are the bandwidth expectations at peak usage hours, and what are the staffing impacts on operational management of the digital hardware and software.

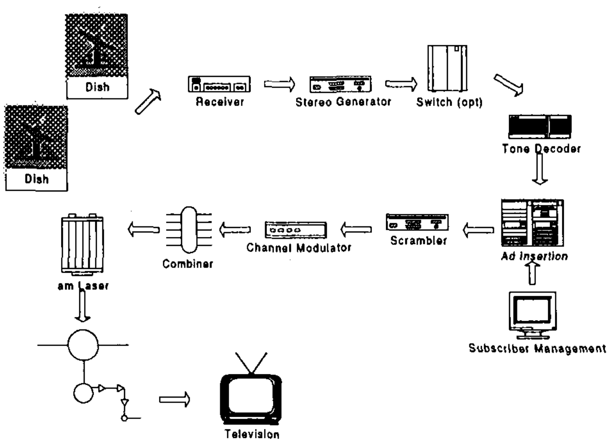

Figure 1. General Headend Flow Diagram

Digital Content Management

As digital services such as high speed data access through cable modems and video on demand are implemented, headend hardware may tend to reflect Figure 2 and Figure 3 respectively. In addition to traditional headend hardware, we are now faced with the management and administration of multimedia servers, routers, and large quantities of disk and/or tape drives.

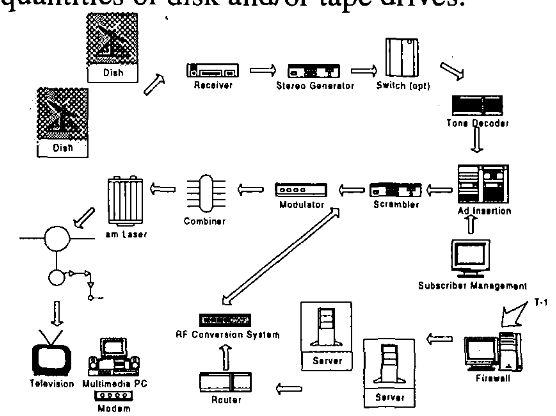

Figure 2. Potential Cable Modem Headend Flow

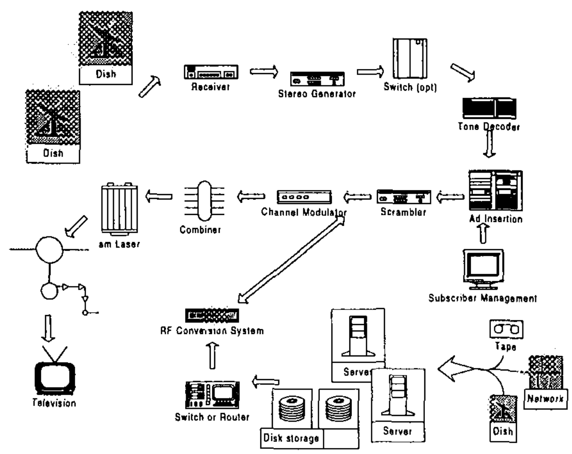

Figure 3. Potential VOD Headend Flow

OPERATIONAL MANAGEMENT

Operational management of digital content is currently very manual; monitoring of content usage, deletion and addition of content onto digital servers must be maintained at that server and, depending on the robustness of the operating system, may require experienced system administrative personnel. The main areas affected: disk space management, failover recovery, bandwidth management, and supporting viewing percentages can be automated in software. We will look at each of these areas and the operational impacts. As software develops to address digital content turnover, these are the key areas that will need to be focused on.

Disk Space Management

With the addition of digital content we are now faced with storage of movies, sporting events, ads, HyperText Markup Language (HTML) files, graphics, and more. Since Video on Demand (VOD) and cable modem services are both on-demand based, we have to store all available content on digital storage media. Considering a full library of movies at approximately 4 gigabytes (1 billion bytes) each, it is easy to imagine a headend consisting mostly of large storage vaults containing terabytes (1,000 billion bytes) of digital storage.

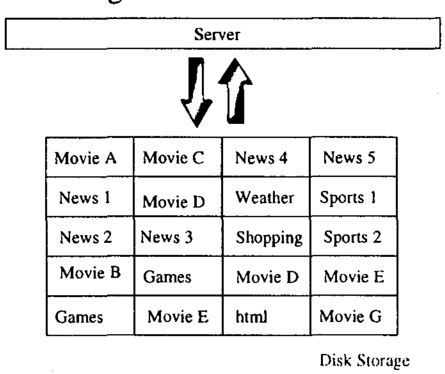

Figure 4. Sample VOD Disk Layout

Figure 4 represents a sample shelf of disks and their stored content whereas Figure 5 represents a similar setup for providing cable modem services. For demonstration purposes, we will analyze the VOD disk layout. What would a typical VOD content turnover look like in one day? A potential schedule may be:

- Remove Movie
- Capture sporting event C at 10am
- Add Movie G from tape
- Update 50 shopping TIFF files and data to display new catalog items
- Add new movie trailers (ads)
- Update NEWS every 1/2 hour

We must also note that digital content turnover is still a very manual process today and each of these items would involve a headend staff member working at a server to complete the tasks. Software solutions are, however, possible and being developed for current trials such as Orlando's Full Service Network.

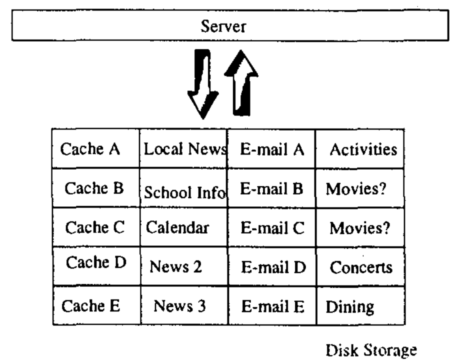

Figure 5. Sample Cable Modem Disk Layout

All networks have many areas for potential throughput bottlenecks. One of these areas is disk and SCSI (Small Computer System Interface) bandwidth. Consider that a single view of a MPEG compressed movie may take up 3.5 Mbps (megabits per second), but the SCSI controller connecting the server to the disks may only support 20 Mbps. Even though a common configuration involves sharing, or striping across many disks, there are still limitations to the number of simultaneous views of a single copy of a movie or other digital content. Depending on the settop or PC application, the result of having no disk bandwidth may be screen blocking or denial of content to the subscriber. What this means to the operational staff is the need to monitor content requests and to make copies of digital content as required by viewing percentages. A recommended threshold may be to copy a digital asset when the bandwidth is 70% used. Note that copying of content also uses bandwidth, as is addressed in the bandwidth section of this paper.

Considering the addition of monitoring and copying, the daily operational schedule may now look like:

- Remove Movie
- Capture sporting event C at 10am
- Add Movie G from tape
- Update 50 shopping TIFF files and data to display new catalog items
- Add new movie trailers (ads)
- Monitor all viewing of content 24 hours a day
- Copy movies whenever viewing percentage exceeds 70% of digital bandwidth

We can now see that the operational impact of digital content management could be significant; however, if the digital system is designed with failover and smooth error recovery in mind, all of the above tasks can be automated in software. More about failover will be addressed in the failover section of this paper.

Hierarchical Storage Management

Although the cost of digital memory and disk drives has been exhibiting typical "Moore's Law" behavior, the cost of storing terabytes of digital content on this media is not yet cost effective. However, access times to other media devices, such as tape drives, do not lend themselves to true video on demand. One potential solution: Hierarchical Storage Management (HSM). HSM has traditionally been designed to allow NVOD or access to digital data on a "wait" basis as the data is moved from tape drives to on-line media.

Assuming that a 90 minute MPEG compressed movie requires 4 gigabytes of storage space and that a suggested file size maximum is 500 megabytes, a digital movie would consist of approximately 8 files representing 11 minutes each. As opposed to storing all 8 files on expensive, fast-access storage, Hierarchical Storage Management would also allow us to store only the first 11 minute file on disk and stream the remaining files as the first file is playing. The only disadvantage of this would be the lack of fast-forward capability for some period of time, until the complete title can be moved to disk. To maximize space with this storage method, the streaming software could delete viewed files from disk; therefore using a maximum of 1.5 gigabytes of disk storage per movie as displayed in Figure 6.

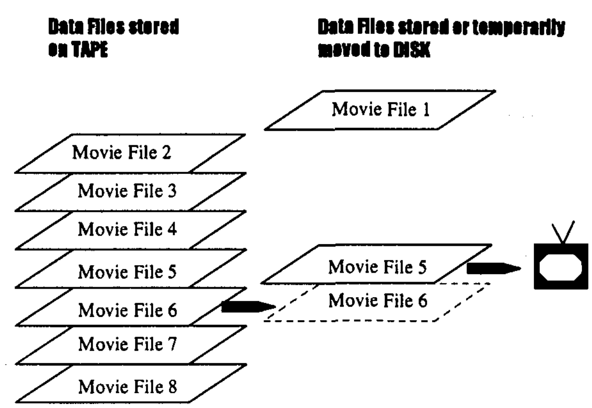

Figure 6. Hierarchical Storage Management Implementation

Network Management Systems

Considering the quantity and complexity of hardware and software to support digital services, monitoring can quickly become overwhelming. Network Management Systems are designed to:

a) reduce downtime by ensuring all components are monitored constantly and consistently thus allowing problems to be addressed immediately and,

b) reduce cost by automating the monitoring of network components, and
c) reduce mean time to repair.

Simple Network Management Protocol (SNMP) was designed as the standard network management protocol for TCP/IP (transmission control protocol/internet protocol) based networks. Management information bases (MIBs) are created by vendors that describe the network objects and SNMP allows the *monitoring* and *control* of each object through a management station containing a user interface for operational staff.

Network management systems have long been used in elaborate monitoring of telephony networks; however, the unique combination of hardware and new, customized software for digital cable services do not allow for easy, complete management with existing management software and MIBs. The investment in customized MIBs and agents to support software and hardware monitoring are worth it for day to day operations considering the monitoring of :

- digital servers
- headend hardware (routers, switches, hubs, bridges, etc.)
- server network traffic
- SCSI and RF bandwidth
- viewing percentages
- custom operating software
- server operating system
- client and field devices

The automation of this monitoring will allow headend staff to isolate problems quickly and focus on maintenance and repair as opposed to elaborate monitoring.

Traditional network management systems will monitor components, determine where a fault occurs and notify the user through a graphical user interface. As expert software develops, we have the additional capability of having the system automatically fix the problem and reconfigure hardware and software to isolate the troubled area.

<u>Failover Recovery</u>

Another area critical to the operational management of digital content is how to handle failover. Assuming we have implemented a network management system that finds network problems early, what do we do about it? Proper management of failures begins in the design of the digital system. Few cost models for digital failover exist; however, we must keep in mind that subscriber downtime must be minimized and recovery of disk or server failures can be very time consuming.

Server Failures, or commonly called server crashes, should be accounted for in the design of the operational software. All digital services should maintain a shared memory area with a backup server so when one server fails, another can continue serving it's subscribers and functions.

Software Failures can be difficult to quantify. Applications for a PC or settop should be developed so software failures have user friendly recovery programs for the subscriber and allow continued navigation. When possible, software should utilize common libraries to ensure standards with failover are properly met.

Network Failures should be minimized through duplicate network paths between servers and configurable RF paths. If, for example, a modulator were to fail or be taken out of service, the system routing/connection scheme should be configurable to easily send that data through another modulator for the same node or area. Accounting for periodic maintenance of network hardware must be built into the operating software design.

Disk Drive Failures can be worked around by implementing existing technologies such as redundant array of independent disks (RAID) which uses a predetermined amount of disk space as overhead for redundancy; therefore allowing disk failures with minimal impact to streaming of data.

BANDWIDTH MANAGEMENT

Figure 3 shows a network topology which includes two media servers with disks attached that contain digital content. Data is switched and converted to MPEG2 for transport to the home. Digital bottlenecks in this environment shift continuously. We can have bottlenecks in moving data through the SCSI controllers from disks to the servers, in moving from the servers to the switch, and in sending burst data through the network, upstream and downstream. It is critical that all available customer bandwidth must be allocated on an as needed basis, creating a dynamic utilization environment.

However, even by having a dynamic bandwidth topology, what happens when no bandwidth is available? For example, as depicted in Figure 7, the same SCSI bandwidth may be used to populate unused drives as well as to stream video to subscribers. When the device capacity is maximized, a bandwidth manager has to be able to prioritize based on business issues. Each piece of digital content should be associated with meta data such as the application this content is associated with, time of day weighting factors, type of media, usage expiration date, and priority of use. When bandwidth contention occurs, these business factors can be used to prioritize bandwidth access. In the evening, for example, news content may be prioritized higher than children's programming and vice versa in the morning hours. Newly released movie titles may take precedence over the placement of older titles onto disks. This type of prioritization can only be accomplished with a software-based, automated content management system. Development of such systems is currently in its infancy.

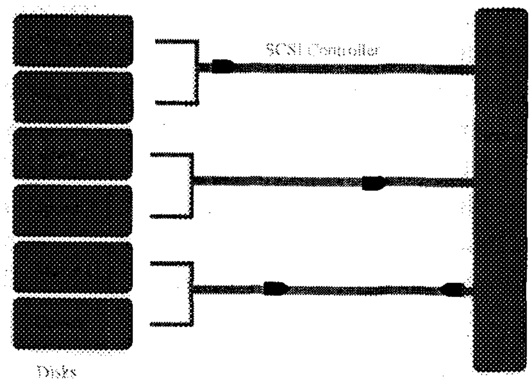

Figure 7. Sample SCSI Controller Assignment

Supporting Viewer Percentages

In addition to having an automated system managing bandwidth based on business and contractual factors, we must consider the use of bandwidth to support viewer percentages. If all bandwidth is being used to stream data to customers, we have no server bandwidth left to make additional copies of the content or to populate more content to disks using the same SCSI controllers.

There are several things we can implement to plan ahead for peak usage times. Figure 8 shows us sample VOD bandwidth usage. We can see that the average bandwidth usage is low over each SCSI controller, however, the peak use is quite high. When bandwidth usage (read and write) exceeds 70% or some pre-determined mark, bandwidth should be allocated for copying the popular content to another disk. By having an automated system for content placement which also monitors and allocates bandwidth for disk writes, we can be guaranteed not to create poor video quality (colored blocks on the television set) for subscribers when populating new content.

Another effective method of maximizing bandwidth is to offset the bandwidth and storage needs of popular content with less viewed content. In Figure 7 we see that SCSI bandwidth is shared amongst many disks. If news is popular in the evenings and children's movies are popular in the mornings, these two types of content would be well suited to utilize the same controller, therefore maximizing the available bandwidth for each at various times of day.

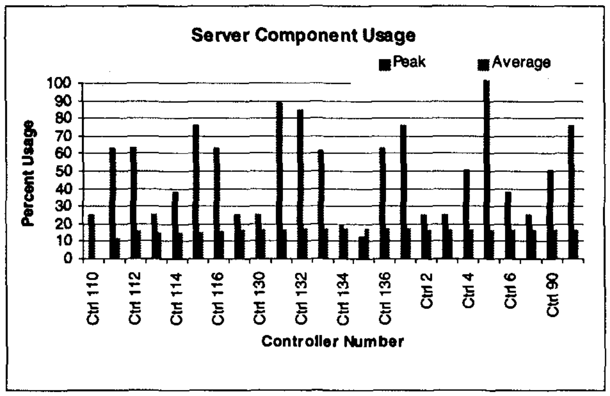

Figure 8. SCSI Bandwidth Usage

SUMMARY

Supporting the operational management of digital content poses new challenges. There is currently no digital content management software that tracks content access, predicts content access based on past and current usage, tracks and manages content distribution and storage, tracks and manages storage availability, manages scheduling of asset distribution based on predetermined criteria such as licensing constraints, manages optimal bandwidth utilization, and supports hierarchical storage management. By implementing the provided solutions for disk, network, and bandwidth management we can manage digital content today, maintain control of the system operations, and ensure quick failover recovery.

Subjective Effects of Bit Error on MPEG-2 Video

Mukta Kar, Majid Chelehmal, Rhonda Hilton, David Eng, and Richard Prodan
Cable Television Laboratories, Inc.

Abstract

The technology of digital communication systems allows more robust transmission than analog systems, but it is not completely error-free. During transmission, bit errors are introduced into the signal stream by the terminal equipment or the transmission media or both. Bit errors are unwanted but often unavoidable. Burst noise distributes bit errors due to the process of data interleaving before radio frequency (RF) modulation and transmission.

An attempt has been made to study the subjective effects of bit errors on the MPEG-2 (Moving Picture Experts Group) digital video stream. After studying syntax elements at the system layer and the video layer of the MPEG-2 standard, subjective effects due to bit errors have been estimated. The experiment at the CableLabs laboratory introduced errors into the bitstream at random locations using a software bitstream editing tool. The resulting bitstream was demultiplexed and decoded. Visual impairments were subsequently observed on a TV monitor and bit errors were analyzed using CableLabs' conformance tools. In most cases, the experimental results matched estimated ones.

INTRODUCTION

The MPEG-2 standard has been established to facilitate the large scale delivery and exchange of compressed audio/video information using communication networks and digital storage media. Another implicit goal of this standardization is to achieve interoperability among the equipment that will be used in this process of delivery and exchange. In the MPEG-2 standardization, syntax and semantics have been developed to represent compressed audio/video signals so that after transmission they can be recovered uniquely. Some errors are unavoidable and are present as channel noise in transmission system or as bit/byte dropouts in storage media. These errors are often termed "burst errors."

BACKGROUND

If the samples of a signal are sent sequentially in a transmission system or stored sequentially in a storage system, the effects of burst errors are overwhelming. The majority of the samples are lost due to burst errors and are not recoverable. To minimize the effect of burst errors, a scheme known as data interleaving is employed. The process of interleaving is shown in Figure 1.

In interleaving, a number of sequential symbols are assembled into code words. A number of sequential code words are ordered along rows in memory. When the memory is full, the symbol sequence is reordered over the medium by reading down columns. If burst errors occur on such interleaved data, full samples are not necessarily lost. The net effect of interleaving is to spread out the burst errors as bit errors over a large number of symbols instead of affecting one or more contiguous symbols. The original symbols are recovered by the inverse process of interleaving known as de-interleaving.

For the above reason, the authors investigated the effect of bit errors in the MPEG-2 bitstream on viewing or subjective quality.

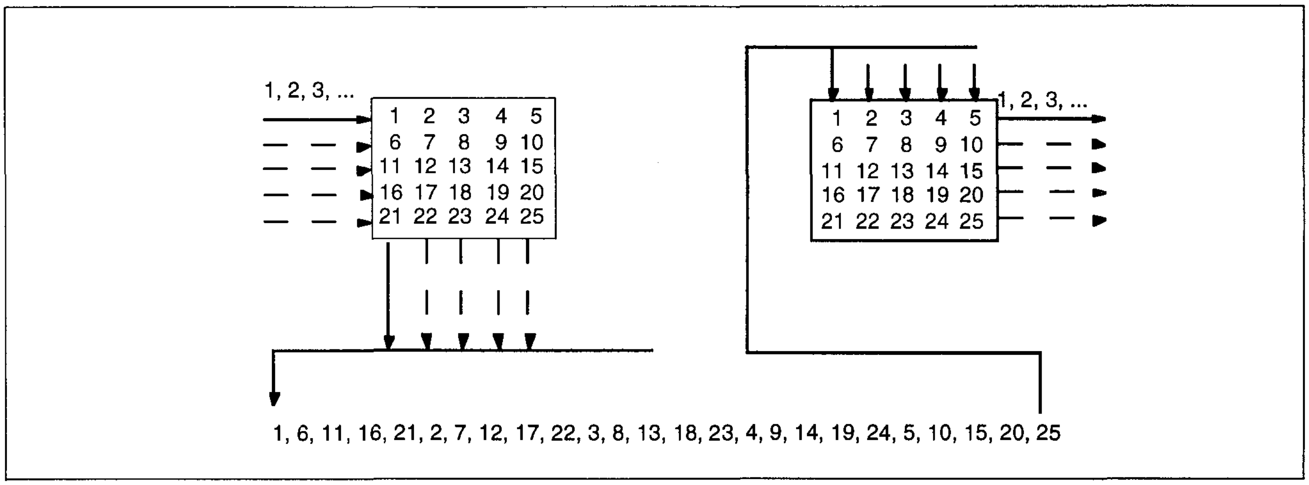

Figure 1: Interleaving and De-interleaving

MPEG-2 STRUCTURE

In order to understand better the effect of bit errors on picture quality, it is worthwhile to review the fundamental structure of MPEG-2 bitstream syntax. The MPEG-2 standard is divided into two distinct layers - the transport layer and the compression layer as shown in Figures 2 and 3.

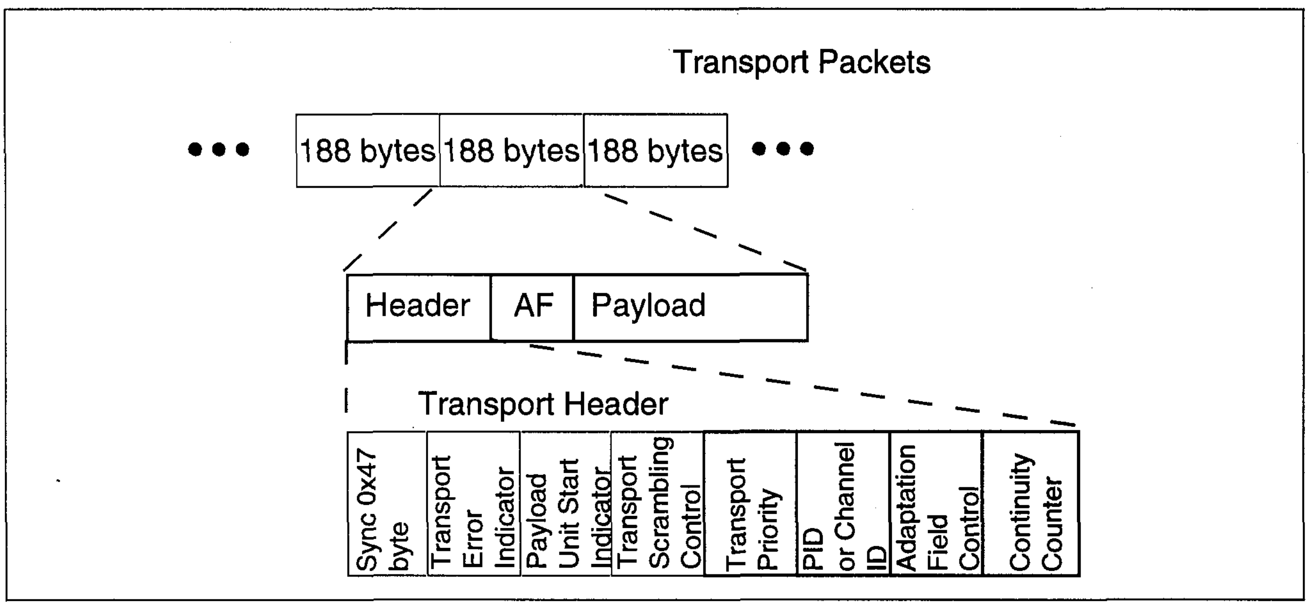

Figure 2: The Basic Content of TS Packet