ENCODING AND DELIVERY OF MPEG-2 COMPRESSED CONTENT TO CABLE HEADENDS

John Vartanian
iN Demand

Abstract

Many technological advances are making video-on-demand (VOD) economically feasible for cable operators. Fiber rich architectures, the deployment of digital settops and dropping file server costs have resulted in virtually all of the major MSOs planning to either test or deploy VOD on their systems in the near future. Although capital costs have been dropping rapidly, operational expenses are still significant. The cost of encoding, distributing and loading content onto VOD file servers is a major expense if borne by a single operator.

One way to reduce these operational expenses is to spread costs across many systems offering video-on-demand. To this end, iN Demand has embarked on an effort to encode content to a single, open specification that works with a variety of file servers and settops, and then distribute that programming through a nationwide satellite network.

Encoding

Encoding content is a significant expense, especially when the number of potential VOD subscribers is small or the programming is not expected to generate many VOD purchases. It is a labor-intensive process that requires the expertise of a compressionist. By spreading the fixed encoding cost across a wide base, this expense becomes less burdensome to individual operators.

An encoding specification is necessary to ensure that content is compatible with the various file servers and settops that will be used to deliver VOD. Many of the details of this specification deal with packet identifier (PID) definitions, packet lengths, program clock reference (PCR) information and group-of-picture (GOP) structure. All of these are important and necessary to ensure the interoperability of content. However, I will focus on three areas that are more discretionary: video bit rate, resolution and audio.

Many encoding decisions affect the economics of offering VOD. Bit rates result in tradeoffs that determine the number of simultaneous users on a VOD network. Video resolution ensures high quality by properly allocating bits between detail and motion needs. Provisions for high quality audio mean that VOD delivered by cable operators will not be at a competitive disadvantage to DVDs.

The first step in encoding a movie is to acquire a dub of the movie master. In order to minimize video artifacts, dubs for digital video compression are usually provided in a Digital Betacam or D1 format. Both of these digital component formats allow the compressionist to start with the best possible source material.

Bit Rate

iN Demand encodes content that is MPEG-2 compliant, but with additional parameters defined to meet the needs of VOD on cable systems. One of the most important parameters in the specification is the video bit rate assigned to each VOD stream. Bit rates that are too high sacrifice VOD availability without a visible improvement in quality. Bit rates that are too low result in quality degradation that can be perceived by VOD users.

The video bit rate in the specification was determined by holding subjective tests with a dozen viewers, including representatives from the technical, marketing and business development areas. A five-minute section of The Matrix, which was a challenging movie to encode, was compressed with varying parameters. The clip was encoded at video bit rates that changed in 100 kbps increments from 2.6 Mbps to 3.4 Mbps. Compression was done for each bit rate at full, ¾ and ½ resolution. In addition, compression for each bit rate and resolution was done twice: once at a constant bit rate and then at a variable bit rate, but with the CBR set as the maximum allowable bit rate. In all, the bit rate, resolution and CBR/VBR permutations resulted in 54 versions of the test clip.

The consensus of the evaluators was that a video bit rate around 3.1 Mbps was needed to produce good video quality with material that is challenging to encode. Bit rates below that resulted in perceived degradation of the video quality. Bit rates above that did not increase the video quality.

We found that CBR produced better video quality than VBR within the constraints we set on the VBR clips. While the advantages that true VBR bring are desirable (as seen with DVD encoding), it is not economically feasible at this time to process VBR streams in a headend without reducing the number of streams carried in a cable channel. It is anticipated that once rate remultiplexing and other video processing technologies are further developed, VBR will lead to average bit rate reductions without degradation of video quality or a reduction in VOD system capacity.

Resolution

One result of the evaluation that was somewhat surprising occurred in the area of video resolution. With full resolution, all the picture elements, or pixels, are evaluated while encoding a movie. In this mode, the encoder uses a frame of video that is typically 720 pixels in a horizontal line and 480 active vertical lines. In the ¾ resolution mode, the encoder reduces the number of horizontal pixels to 544 or 528, but still keeps the 480 vertical lines. With ½ resolution, 352 pixels across and 480 lines down are used. Although it may seem that using full resolution would provide the best quality, this was not the case. Since a consumer television cannot display all the detail encoded at full resolution, it is inefficient to encode and transmit this extra data. Instead, either ¾ or ½ resolution provides higher quality. In these modes, bits that would be used to encode detail in the full resolution mode are instead used to encode motion. The result is no loss of detail on a consumer television and a reduction in motion artifacts that would occur at full resolution. In general, we found that ¾ resolution usually provided the highest quality, but ½ resolution was better in some scenes that contained a lot of action.

Audio

In addition to the video requirements, a portion of the bit stream must be reserved for several other components. Dolby Digital audio, which provides state-of-the-art surround sound for home theaters, requires .384 Mbps. This type of audio is the standard for DVDs and some DBS channels and must be included in order for VOD to be competitive with these other technologies.

We also looked at secondary language support. If alternate language tracks are multiplexed with the video and English audio, additional bandwidth needs to be reserved. Instead, the specification calls for additional audio services to be assigned their own audio PIDs. The desired audio is multiplexed with the video at the file server and only that one language is delivered to the consumer settop. This method of multiplexing allows for a virtually unlimited number of language choices, with no effect on bandwidth requirements.

QAM Loading

Once we determined the minimum video bit rate and the audio requirements, we looked at QAM loading. Program specific information (PSI) and overhead needs were added to the 3.1 Mbps video and the .384 Mbps audio. This resulted in a minimum need of approximately 3.67 Mbps per stream. At this rate, seven streams can fit into the 26.97 Mbps payload of a 64 QAM channel. The maximum number in the 38.8 Mbps payload of a 256 QAM channel is 10 streams. Once the QAM channel loading was determined, the video bit rate was increased from 3.1 Mbps to 3.18 Mbps. This allowed for a higher video bit rate, without a decrease in the number of streams in a 64 or 256 QAM modulator. After some other small requirements are added, a total multiplexed rate of 3.75 Mbps per stream became the bit rate specification.

Meta Data

Meta data is any information that is associated with a movie. Common meta data elements for a movie include its length, a summary, the actors, a rating and the movie's availability dates. Meta data is necessary in the operation of a VOD service. It allows a customer to do such things as set parental lockouts or sort a movie by genre.

As with content encoding, it is efficient to add meta data to a movie once, with the meta data being interoperable with different VOD applications developed by various file server and settop vendors. This removes the need for insertion of meta data at each VOD system. In addition to reducing labor in the headend, centralizing meta data functions reduces the opportunity for errors that could occur if meta data had to be entered at each VOD system. iN Demand is using a meta data standard that defines each meta data field, its type and its length.

Efforts are underway to expand the meta data standard to incorporate all meta data needs by the major VOD file server and settop vendors. By creating a superset of meta data and sending the same version to everyone, systems get the advantages of centralized meta data creation. Except for movie promotion, which includes an MPEG-2 video trailer and a JPEG or GIF box cover or poster art image, meta data is text information that does not require significant additional bandwidth. There is little disadvantage to sending a superset of meta data to all systems. Each VOD application at a cable system can pick the meta data fields it needs and ignore the others.

Distribution

Once a movie has been encoded and meta data has been added, the next challenge is to distribute the asset package to VOD systems. The easiest way to distribute compressed movies, at least in the early stages, is by sending digital linear tapes (DLTs) via an overnight shipping service. However, this method of distribution becomes inefficient when the number of deployed VOD sites grows beyond a few. Satellite delivery is a natural means for sending content to cable headends. Once the satellite distribution network has been built, it is easily scalable. Additional VOD systems can be added at a relatively low incremental cost.

Satellite Link

iN Demand distributes pay-per-view movies, events and sports on eight 36 MHz C-band transponders. To maintain compatibility with cable systems that are using 64 QAM to deliver digital signals, 24 MHz of each transponder is used for MCPC transmission. The other 12 MHz in the transponder creates an opportunity for VOD file transmission.

The first step in our satellite analysis was to perform a link budget to determine the power level of each carrier and their center frequencies. When multiple carriers occupy a transponder, the power of each carrier must be reduced to minimize intermodulation products. Once the power levels and carrier frequencies were determined, the MCPC carrier was moved from the center of the transponder to its edge, allowing the maximum bandwidth for the new VOD carrier. An 8 Mbps carrier was then added to the unoccupied transponder bandwidth.

IP Multicast System

On the following page is a diagram of the IP multicast system that iN Demand is using deliver VOD files to cable systems. The pitcher is at the heart of the IP multicast system. It is the component, along with external hard drives, that stores the assets and sends them to systems using an IP multicast transmission. It also performs a triple DES encryption, ensuring the security of the VOD file.

The pitcher is controlled by the manager, which schedules the delivery of assets to headends. It analyzes available satellite bandwidth, asset size and required delivery dates to determine when assets need to be transmitted by the pitcher. The manager is operated through a user interface on a remote terminal, with information sent via an Internet connection from the terminal to the manager. The manager also provides monitoring and management functions, both of which can also be accessed at the remote terminal.

An asset creation tool (ACT) adds meta data in an XML format to the movie. It then transfers the content and associated meta data to either the pitcher's internal or external storage, where it resides until the pitcher receives transfer instructions from the manager. As with the manager, the ACT is also accessed through a user interface on a remote terminal, with information sent to the ACT over the Internet.

When the manager instructs the pitcher to transmit a file to a group of systems, the compressed content is sent to an IP encapsulator. The IP encapsulator creates an MPEG transport stream from the IP data and translates the data into a form suitable for satellite delivery. It also applies a ¾ rate forward error correction to compensate for transmission errors in the satellite link. The MPEG transport stream is then sent in an ASI format to a quadrature phase shift key modulator and uplinked to the satellite.

The VOD carrier is downlinked at cable headends and sent via L-band from the satellite antenna to a catcher, which is a Pentium III device operating on Linux. The first function of the catcher is to demodulate the carrier and apply forward error correction to fix transmission errors. It extracts the IP data from the MPEG transport stream and analyzes the file to verify the transmission. If any packets are missing or corrupted, a request for retransmission is sent over the Internet to the pitcher. The pitcher collects these requests from all the catchers and resends any corrupted packets. Once a catcher verifies receipt of a file, it sends an acknowledgement to the pitcher and the transfer to the catcher is complete.

The catcher then decrypts the content and passes it over a 100 baseT network to the VOD file server, where it is ready to be played once the movie availability date is reached.

Conclusion

A significant way to minimize the costs of operating a video-on-demand system is to centralize the encoding and distribution functions. iN Demand has developed such a

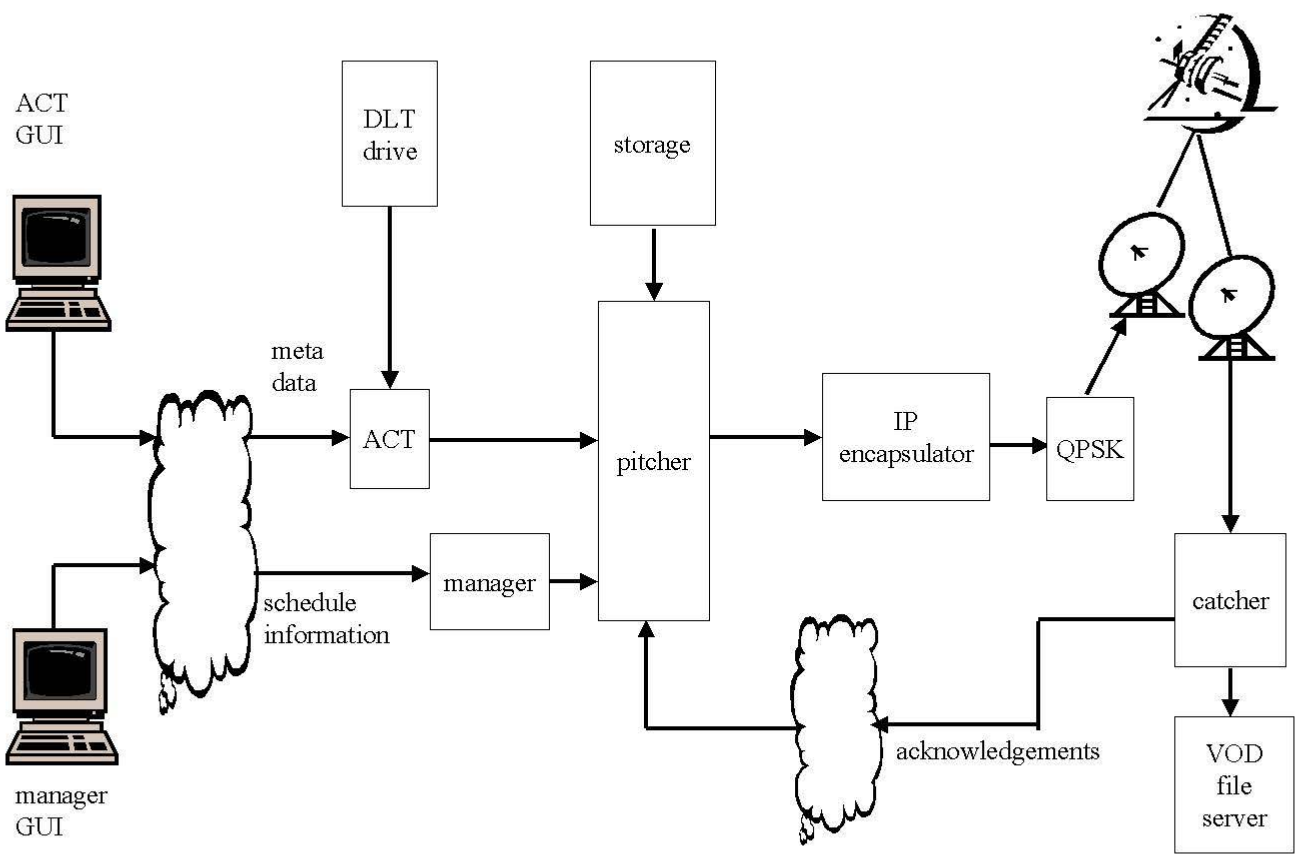

system by using an open encoding specification that supports delivery of content through a variety of file servers and digital settops vendors. Adding meta data to the movie saves labor costs at cable systems and reduces the chance for errors. Finally, an IP multicast system with triple DES encryption is used for secure delivery of content. A dedicated satellite carrier is used for VOD, allowing for a scalable, cost efficient system for transmission.

The use of a common encoding specification and distribution system allows VOD to become a viable additional revenue stream for cable operators.