MPEG-4 TRANSITION USING SWITCHED DIGITAL VIDEO

John Schlack
Motorola, Inc.

Abstract

During the transition from MPEG-2 to MPEG-4, it will not be possible, due to bandwidth constraints, to broadcast all channels in both MPEG-2 and MPEG-4 format. However, the switched digital video (SDV) system can be used to manage the transition from MPEG-2 to MPEG-4 and to minimize the bandwidth needed by the system until the transition completes.

The switched digital video system delivers streams with an encoding type based on the decoding capabilities of the settop boxes tuned to that channel. The SDV system may need to force tune settop boxes or use transcoding to handle transitions between different encoding types.

TERMINOLOGY

The terms H.264, AVC (advanced video coding), and MPEG-4 Part 10 all refer to the same standard for video compression. These terms may be used interchangeably in this document. Additionally, this document uses the term MPEG-4 to refer to MPEG-4 Part 10. MPEG-4 provides high quality video at substantially lower bit rates than its predecessor, MPEG-2.

BACKGROUND

Cable operators currently deliver a mix of analog video and digital video on the cable plant. The digital video includes a large number of standard definition (SD) channels as well a much smaller number of high definition channels (HD). The operators also provide other services such as high speed data, video on demand (VOD), and voice over IP (VOIP).

The cable industry standardized delivery of digital video content using the MPEG-2 format. This format requires about 3.75 Mbps of bandwidth for SD programs and 15-19 Mbps for HD programs. A 6 MHz QAM channel modulated using QAM256 can carry 10 SD or 2 HD programs. Cable operators may use statistical multiplexing to groom the channels. This allows a 6 MHz QAM channel carry 11-12 SD programs or 2-3 HD programs.

Figure 1 shows an example 860 MHz HFC plant carrying these different services. This example system delivers 50 HD channels. However, many cable plants are 750 MHz or less. These cable systems deliver much less HD programming.

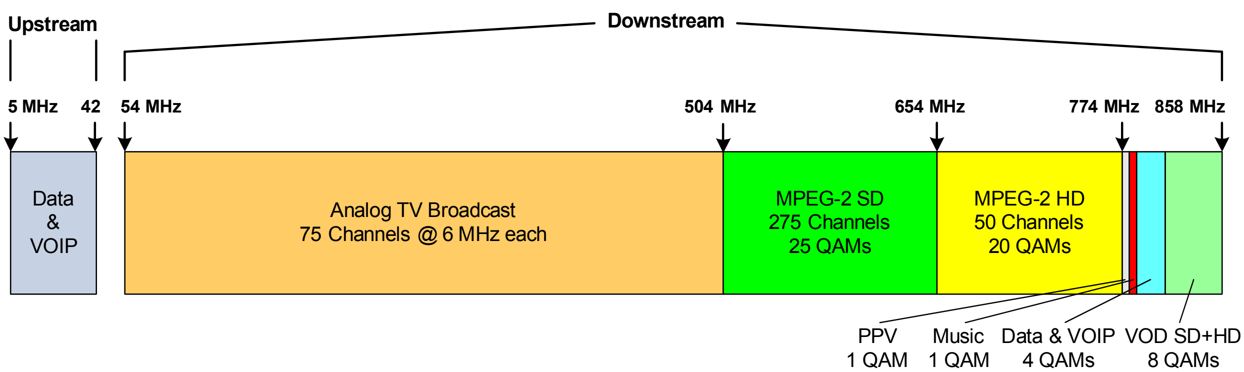

Figure 1: Example 860 MHz HFC Plant

Cable operators are being pressured to deploy additional HD channels. Subscribers are replacing standard definition television with HD sets. This is one source driving the need for additional HD content. Figure 2 shows the current and projected growth of HD sets [1].

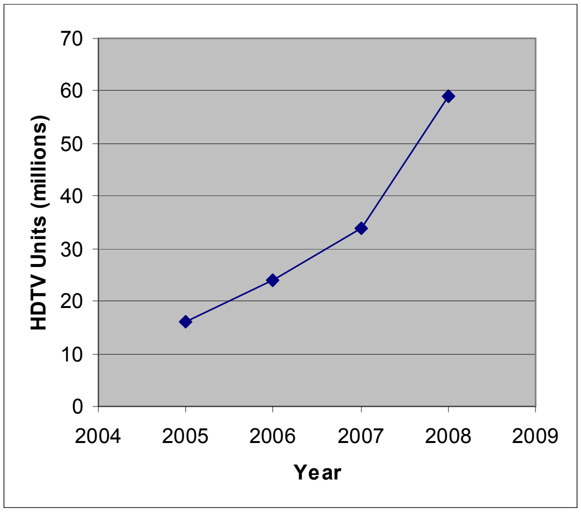

Figure 2: HDTV Growth

Competition is also driving the need to deploy more HD content. Satellite providers currently deliver 70-90 HD channels and plan to increase the number to over 100 channels by the end of 2008. Telco operators plan to deliver a similar number by the end of 2008. Cable operators currently deliver 20-50 HD channels [2]. The cable operators need to deploy more HD channels to keep pace.

The industry has begun a shift towards encoding video using MPEG-4 as a method to reduce bandwidth for delivering HD content via satellite. As an example, a major content provider has announced that it will deliver the HDTV versions of all 26 channels to cable headends using an MPEG-4 encoding [3]. MPEG-4 video requires about 30-50% less bandwidth than comparable MPEG-2 content.

The settop box manufacturers are also beginning to deliver settop boxes that can decode MPEG-4 video. Using MPEG-4 instead of MPEG-2 to encode all SD and HD programs will produce significant bandwidth savings. It will be possible to deliver 100 channels of HD programming and 275 channels of SD programming on an 860 MHz cable plant while still keeping a large analog channel lineup for subscribers without digital settop boxes. The bandwidth calculations are shown in Figure 3.

TODAY – MPEG-2

	Format	Channels	BW
Analog	Analog	75	450
Digital SD	MPEG-2	275	150
Digital HD	MPEG-2	50	120
VOD SD + HD	MPEG-2		48
Data + VOIP			24
Music	MPEG-2		6
PPV	MPEG-2		6
Total			804

TRANSITION PERIOD

	Format	Channels	BW
Analog	Analog	75	450
Digital SD	MPEG-2	275	150
Digital SD	AVC	275	78
Digital HD	MPEG-2	100	240
Digital HD	AVC	100	150
VOD SD + HD	MPEG-2		48
VOD SD + HD	AVC		36
Data + VOIP			24
Music	MPEG-2		6
PPC	MPEG-2		6
Total			**1188**

FUTURE - AVC

	Format	Channels	BW
Analog	Analog	75	450
Digital SD	AVC	275	78
Digital HD	AVC	**100**	150
VOD SD + HD	AVC		36
Data + VOIP			24
Music	AVC		6
PPV	AVC		6
Total			750

Figure 3: MPEG-4 Transition without SDV

However, most digital settop boxes deployed today in a cable system can only decode MPEG-2 encoded content. It is not realistic to replace millions of digital settop boxes with MPEG-4 capable settop boxes in a short time frame. A transition period is required where the cable system delivers both MPEG-2 and MPEG-4 content. During the transition, it will not be possible, due to bandwidth constraints, to broadcast all channels in both MPEG-2 and MPEG-4 format. The "Transition Period" table in Figure 3 shows the bandwidth required to dual carry MPEG-2 and MPEG-4 versions of each channel as well as expands the HD channel lineup to 100 channels.

Many of the cable operators are turning to switched digital video (SDV) as a way to reclaim bandwidth in order to deliver new services, particularly additional high definition content. Switched digital video replaces traditional broadcast programs with a system that only transmits a channel to a service group when requested by a subscriber. The system realizes bandwidth savings since only a subset of the available channels is being watched by subscribers at any given time.

The operator usually places niche or low take rate channels on the SDV tier. When using this "long tail" content, the SDV system can effectively offer at least twice the number of channels than can actually be delivered in a given QAM. The reclaimed QAMs from digital broadcast can be used to offer additional channels or can be assigned to other services.

One can also view MPEG-4 encoded channels as long tail content while the number of MPEG-4 capable settop boxes is less than the number of MPEG-2 capable settop boxes. Thus, moving the MPEG-4 channels onto the switched tier will prove to be a very effective way to introduce MPEG-4 onto the cable plant.

From the example in Figure 1, an operator may choose to move the least watched 100 standard definition (SD) channels and 20 high definition (HD) channels from digital broadcast to the switched digital video tier. These channels originally required 16 QAMs to broadcast. SDV requires approximately 8 QAMs to deliver this content, meaning that 8 QAMs have been reclaimed.

Assume that the 8 reclaimed QAMs will be used for delivering more HD content on the SDV tier. These 8 QAMs can be used to deliver 30-40 additional "long tail" HD channels. Transmitting both MPEG-2 and MPEG-4 format for all channels on the switched tier provides additional bandwidth savings as MPEG-4 capable settop boxes are deployed. Figure 4 shows the example cable plant using SDV to deliver MPEG-2 and MPEG-4 encoded content. The actual mix of SD and HD to deliver on the SDV tier will vary by operator and region based on popularity of the content.

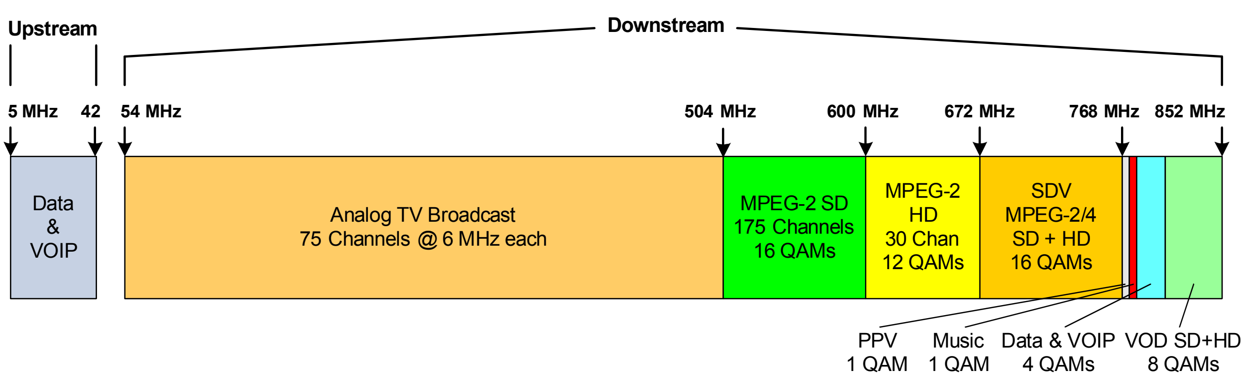

Figure 4: 860 MHz HFC Plant with SDV

The switched digital video system is an ideal platform for launching MPEG-4 channels. However, challenges exist. The SDV system must deliver multiple encoding formats for the same content, ensure that the settop boxes tune to the correct format, and minimize bandwidth usage when there are demands for competing formats of the same content.

MANAGING MULTIPLE ENCODINGS

Architecture

The SDV system has the ability to deliver a stream to the settop box based on the settop box capabilities. For example, if a settop supports decoding only MPEG-2 signals, the SDV system will ensure that the settop is only directed to tune to streams that are encoded using MPEG-2.

The SDV system may be pre-configured with the settop capabilities. For example, the capabilities may be statically tracked by settop model. The operator may track the settop model and settop ID for each fielded settop box. The channel change request from the settop includes the settop ID, which allows the SDV system to discover the settop capabilities and assign the correct stream. Alternatively, when the settop box registers with the SDV system, it may send a message to the SDV system to report its capabilities. The capabilities may include the number of tuners, the video and audio codecs supported, and the communications methods, to name a few parameters.

As the content providers begin delivering MPEG-4 content, the cable system will need to transcode these signals into an MPEG-2 version so that legacy settop boxes will be able to decode the stream. The system may also transcode MPEG-2 content into MPEG-4 content so that additional lower bandwidth content will be available for the MPEG-4 capable settops. This will reduce the overall QAM bandwidth usage on the system if only MPEG-4 capable settops are tuned to a specific channel.

Figure 5 shows an example architecture of a cable plant designed to deliver a stream with multiple encoding types. Transcoders are used to create MPEG-2 and MPEG-4 encodings for the various channels.

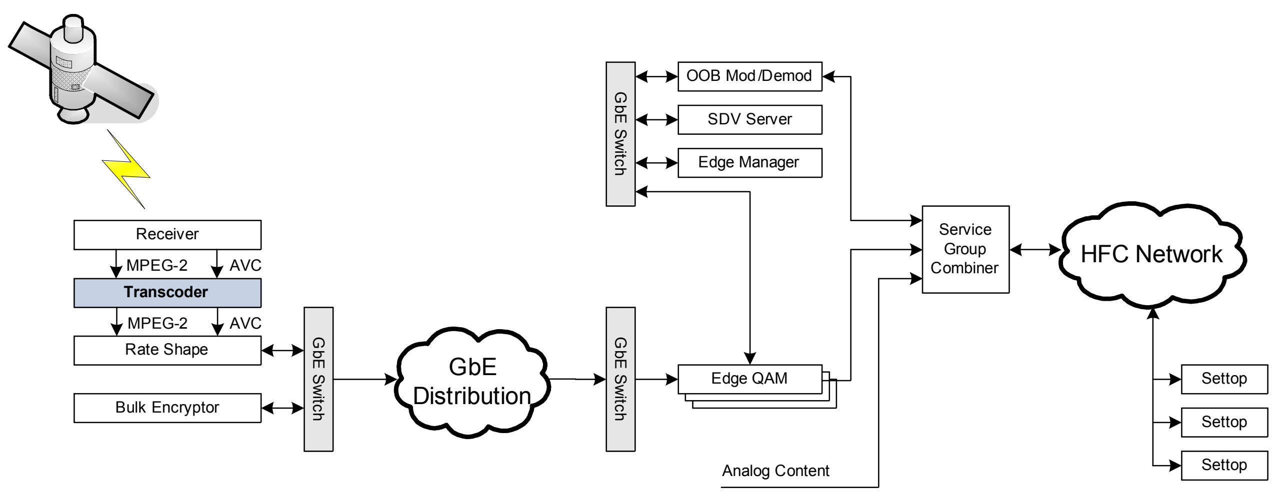

Figure 5: SDV System with MPEG-4 and MPEG-2 Transcoders

In this example, the transcoders are placed in the headend with the other grooming equipment. This increases the load on the distribution network since it will carry multiple copies of some stream to the hubs.

The transcode function is shown performed by a stand alone server. However, it may also be integrated in the edge device, VOD server, Network Encryptor, or any system that ingests and outputs GigE content. The key is that both the MPEG-4 and MPEG-2 versions of the content are available on the plant so that the SDV system can deliver the appropriate version to settop boxes based on their capabilities.

Stream Delivery

The SDV System directs channels to a service group based on requests from the settop box. The SDV system uses the settop capabilities and policy information when determining the encoding format of the stream to deliver to that settop box.

Although it is possible to transcode all content encoded in MPEG-2 into an MPEG-4 encoding to save bandwidth, it is likely that the SDV system will have some content that is only encoded using MPEG-2. In this case, the SDV system always delivers the MPEG-2 version of the stream when requested by an MPEG-2 or MPEG-4 capable settop box. This assumes that the MPEG-4 settop box can decode MPEG-2 encoded content.

The SDV System may have content that is only encoded using MPEG-4. In this case, the SDV system can only deliver such content to MPEG-4 capable settop boxes. MPEG-2 capable settop boxes will be denied service. However, the channel map for the MPEG-2 capable settop box should prevent subscribers from accessing the content if the settop box cannot view it. As an example, the cable operator may have created a special HD tier that is available as a premium package. A subscriber receives an MPEG-4 capable settop box when subscribing to that package so that the subscriber can view that HD content.

For content that is encoded using MPEG-2 and MPEG-4, the SDV system will likely be configured to attempt to deliver only a single encoding of that channel in order to conserve bandwidth.

When a settop box capable of only decoding MPEG-2 content requests an SDV channel that is currently not being delivered to the service group, the SDV system directs the MPEG-2 stream of that SDV channel to that service group and it returns the channel tuning information to the settop box. If an MPEG-4 capable settop box from the same service group subsequently requests the same SDV channel, the SDV system simply directs the settop box to tune the MPEG-2 stream that is already being delivered to the service group. This assumes that the MPEG-4 capable settop can also decode MPEG-2 content.

When a settop box capable of decoding MPEG-4 content requests an SDV channel that is not currently being delivered to the service group, the SDV system directs the MPEG-4 stream of that SDV channel to that service group and it returns the channel tuning information to the settop box. If a settop box capable of only decoding MPEG-2 content that is from the same service group subsequently requests the same SDV channel, the SDV system must direct the MPEG-2 version of the SDV channel to that service group and return the tuning information for the MPEG-2 stream to that settop box. Thus, the channel is being delivered twice to the same service group using different encoding formats.

Carrying both MPEG-2 and MPEG-4 versions of the same SDV channel is inefficient, since it requires at least 50% more bandwidth than carrying the MPEG-2 version of the channel alone.

Managing Multiple Copies

The SDV system can employ several strategies for handling delivery of multiple encodings of the same SDV channel to the same service group. The basic choices are to leave multiple copies of the stream or to force tune all settop boxes onto a single copy of the stream. The chosen strategy depends on the bandwidth available to the system and the impact of reducing the number of copies of the stream.

The main reason that the SDV system might not always choose to deliver a single encoding of the stream is the force tune operation can be disruptive to the viewing experience. Directing a settop box from one stream to another requires the settop box to tune to a different frequency and/or MPEG program number.

Tuning to a different MPEG program number on the same frequency may take several hundred milliseconds while the tuner waits for the video and audio data to arrive. This can cause jitter or blocking on the display. Tuning to a different frequency requires possibly 1-2 seconds while the tuner waits for information describing the stream contents and then waits for the video and audio data to arrive. This causes either a frozen image or a black screen for the duration of the tune. The tune operation may also impact the DVR causing recordings to fail.

If the SDV system has sufficient unused QAM bandwidth for that service group, the SDV system may simply allow multiple encodings of the channel to remain on the cable plant. The SDV system will still have the ability to service new channel requests, therefore it would be best not to disrupt the viewing experience by merging the streams. The SDV system can defer action until available bandwidth for that service group drops below a "consolidation threshold".

In the case where the SDV system elects to leave both streams on the plant, the SDV system may direct all new tune requests for that channel onto the MPEG-2 version of the stream in hopes of recovering resources when the settop boxes originally tuned to the MPEG-4 version of the channel tune off. This mechanism allows the SDV system to gradually recover resources without force tuning settop boxes to the MPEG-2 version of the channel.

In another case, the SDV System may use recent channel change activity to detect that the MPEG-2 settop is channel surfing. Thus, subsequent tune requests from MPEG-4 capable settops may be placed into the MPEG-4 encoded stream in anticipation of the MPEG-2 settop channeling off shortly. If the MPEG-2 settop stays on the channel or other MPEG-2 settops tune to that channel, the SDV system may revert to directing all new requests to the MPEG-2 channel as discussed above.

If resource availability in a service group is low and both MPEG-4 and MPEG-2 versions of the same channels are being delivered to that service group, the SDV system may choose to force tune viewers to the MPEG-2 stream and recover the resources assigned to the MPEG-4 stream. This force tune operation may be disruptive, but it is the fastest way to reduce resource usage in a service group.

The SDV system executes the force tune operation by sending a message to each settop box currently tuned to the MPEG-4 version of the stream to tune a different frequency and/or program number. If the settop box is still tuned to the stream in question, it will tune to the MPEG-2 version of the stream using a frequency and MPEG program number embedded in the request message. The settop box returns an acknowledgement indicating whether the tune operation was performed. If the settop was already tuned to a different channel, the settop box returns an acknowledgment indicating that the tuner is no longer on the channel in question. The SDV system will retry the force tune operation for settop boxes that fail to acknowledge its request.

The SDV system has the option of either immediately recovering the resources assigned to the MPEG-4 stream or waiting until all settop boxes have been tuned to the MPEG-2 stream. Recovering the resources immediately is even more disruptive than simply the force tune operation since the SDV server directs the edge QAM to stop streaming the MPEG-4 encoded content. This likely occurs before the settop boxes begin their tuning operation to the

MPEG-2 channel. This results in additional time with a frozen picture on the screen or a blank screen.

The more graceful method of handling the resource recovery would be to wait until all active tuners have been directed off the MPEG-4 stream. Although this delays resource recovery by several seconds, it minimizes the impact on the subscribers and provides an opportunity for the settop box to use some advanced features in an attempt to minimize the impact of the force tune.

The SDV system can further attempt to minimize the impact of force tuning by following these guidelines. The SDV system makes every effort to place both the MPEG-4 and MPEG-2 versions of the same channel on the same QAM. This greatly reduces force tune times. The SDV system may schedule force tunes to coincide with transitions between content and advertising. These force tunes would not only require knowledge of the expected time for the advertising, the SDV system would need to receive a trigger from the ad server indicating the exact time of the advertising since the goal would be to place the force tune on the transition into or out of the ad pod. This would minimize the impact on both the subscriber and advertiser since the transition typically includes a black frame or a fade to black sequence.

Recent developments in settop box-based ad splicing are providing the capability that the settop box can perform a seamless splice between programs carried on the same frequency. If the SDV system is able to keep both the MPEG-2 and MPEG-4 stream on the same QAM and it does not recover resources until all settop boxes have acknowledged the force tune operation, then the settop box may be able to jump from the MPEG-4 stream to the MPEG-2 stream with little or no noticeable impact on the subscriber viewing experience.

DYNAMIC TRANSCODING

Overview

The previous section described a static transcoding method in which the headend network carries both MPEG-2 and MPEG-4 copies of the same SDV channel to support both legacy settop boxes and MPEG-4 capable settop boxes.

Another method would be dynamic transcoding. In dynamic transcoding, the SDV System directs a single MPEG stream onto the service group for each requested SDV channel. The MPEG stream is dynamically transitioned between MPEG-2 and MPEG-4 based on the types of settop tuned to that SDV channel.

Dynamic transcoding should be possible to implement. The transport stream structure for MPEG-4 encoded content can be similar to that of MPEG-2. However, the video and audio packets would carry MPEG-4 encoded data instead of MPEG-2 encoded data. Additionally, the stream type identifier will signal that the content carried in the stream is MPEG-4 [4].

MPEG decoders currently exist that can decode both MPEG-2 and MPEG-4 content. These are being deployed into the new MPEG-4 capable settop boxes. However, it is likely that these decoders cannot dynamically transition between the two encoding types. New decoders may need to be deployed.

Decoders capable of handling the dynamic transition between encoding types will be alerted to the transition by a change in the PMT and a change in the stream type identifier. The transition will occur on an I-frame boundary for MPEG-2 and IDR frame for MPEG-4.

Architecture

The transcoding device itself can be a stand alone server, part of an edge QAM, embedded in a video server, or embedded in a network encrypting device. Note that the transcoding device cannot work on encrypted streams. Therefore, encryption must be applied after the transcode. This can be done by a bulk encryptor in the headend or encryption logic in the edge QAM in the hub.

Figure 6 shows an example system where dynamic transcoding occurs in the edge QAM. The headend and GbE distribution network carry a single copy of each channel. Each channel may be encoded in MPEG-2 or MPEG-4 based on how it was received from the satellite. Then, on a service group by service group basis (i.e. QAM by QAM basis), the edge QAM can be directed to output either the MPEG-2 version of that channel or the MPEG-4 version of that channel.

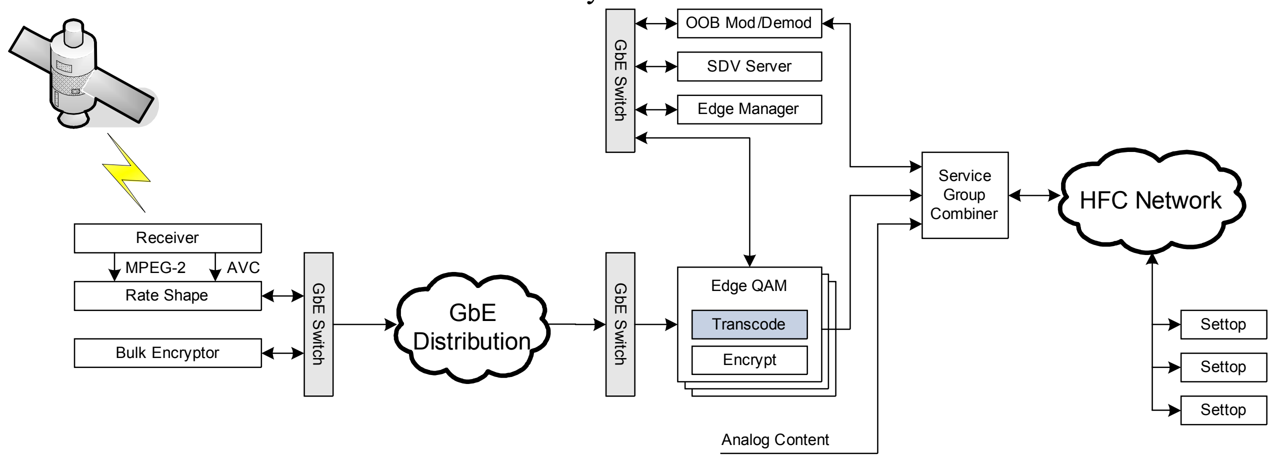

Figure 6: Dynamic Transcoding in the Edge QAM

A drawback of transcoding in the edge QAM is that the stream cannot be encrypted until after the transcode. This means that the stream must be both transcoded and encrypted at the edge. Performing both transcoding and encryption operations at the edge may be costly as it requires dedicated processing power for each service group.

Dynamic transcoding may also occur in a video server. The video server ingests the MPEG-4 or MPEG-2 source content. It can then delivers an MPEG-2 or MPEG-4 version of the stream for each service group with a settop box tuned to that channel. Figure 7 shows dynamic transcoding using a video server.

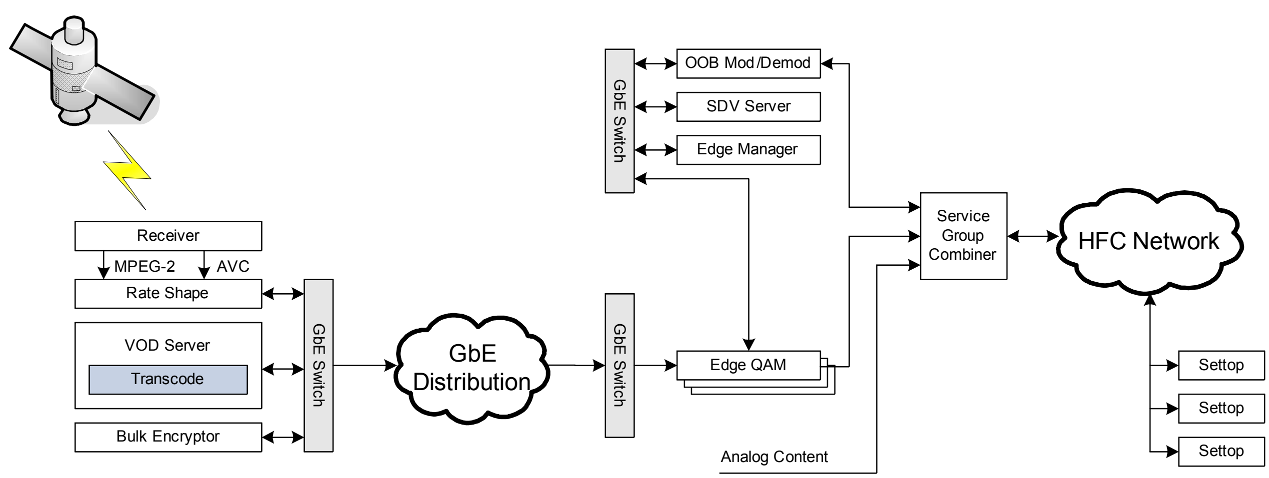

Figure 7: Dynamic Transcoding in the Video Server

By delivering a unicast stream for each service group, the Video Server is free to transition the stream between the MPEG-2 and MPEG-4 versions as requested by the SDV system. Note that this method requires significantly more bandwidth on the GbE distribution network than edge QAM based transcoding.

Stream Management

Stream management when using dynamic transcoding is greatly simplified. The SDV system simply manages the transitions between the MPEG-2 encoding and the MPEG-4 encoding based on the capabilities of the settop boxes that tune to the channel. In most cases, no stream merging concerns exist. That said, there are still issues that need to be handled.

In one case, a particular SDV channel is being delivered in MPEG-2 format to a mix of legacy and MPEG-4 capable settop boxes for a particular service group. If all legacy settop boxes tune to other channels, the SDV Server has the opportunity to transition the stream from an MPEG-2 encoding to an MPEG-4 encoding. The SDV Server sends a message to the transcoding device directing it to transition from the MPEG-2 version of the stream to the MPEG-4 version of the stream for the service group in question.

The SDV System may not immediately transition an MPEG-2 stream to an MPEG-4 stream the instant all legacy settop boxes tune off. The SDV Server may use heuristics to determine the appropriate transition time, since viewer behavior may indicate that some legacy settop boxes are only channel surfing and will return to the original SDV channel shortly. For example, the SDV System may have a database of advertising times for that SDV channel. If legacy boxes tune off during commercial breaks, the SDV System may wait until the commercial break ends before determining whether to transition to the MPEG-4 encoding.

In another case, a particular SDV channel is being delivered in MPEG-4 format to a set of MPEG-4 capable settop boxes. When a legacy settop box requests that SDV channel, the SDV Server must immediately direct the transcoding device to transcode the channel to MPEG-2 format. The transcoding device will begin transcoding at the next I-frame / IDR-frame boundary. The transcoding device then sends a response to the SDV Server when the MPEG-2 transition has completed. The SDV Server returns the tuning information to the legacy settop box.

The SDV Server must wait for confirmation that the transcoding device transitioned to the MPEG-2 version of the stream before the SDV Server returns tuning information to the legacy settop box. This may take several hundred milliseconds or more. By waiting for confirmation, the legacy settop box will be guaranteed it is attempting to tune an MPEG-2 stream and not an MPEG-4 stream.

Note that the SDV Server may deny the request of the legacy settop box to tune to the SDV channel if there is insufficient output bandwidth on that QAM to switch from MPEG-4 to the MPEG-2. Alternatively, the SDV System may need to establish the MPEG-2 version of the stream on a different QAM that has sufficient bandwidth. In this case, one returns to the case where the SDV System delivers multiple streams for the same channel.

COMPLETING THE TRANSITION

Initially with a small number of MPEG-4 capable settop boxes, SDV will be an enabling technology for deploying MPEG-4 into the cable plant, allowing delivery of additional HD content and potential bandwidth savings. As the number of MPEG-4 capable settop boxes grows beyond the deployed legacy settop box count, the cable plant will transition to broadcasting more MPEG-4 content while delivering the MPEG-2 content on the switched tier. This will provide further bandwidth savings. As the number of legacy settop boxes dwindle due to replacements and upgrades, the cable operator may find it advantageous to replace the remaining legacy settop boxes and remove MPEG-2 content from the cable plant. This would complete the transition to MPEG-4 encoded content.

REFERENCES

[1] Goetzl, D. (2007, November 26) HDTV Set Sales to Soar, *Broadcasting & Cable*; and B&C Staff (2007, December 28) CEA: DTV Penetration Tops 50% Mark in U.S., *Broadcasting & Cable*.

[2] Patterson, T. (2008, February) Fighting for HD bragging rights, *CED Magazine*, 34.

[3] Ellis, E. (2007, August 1) The MPEG Transition, *CEDMagazine.com*.

[4] Robuck, M. (2006, July 1) MPEG-2 Stays Center Stage While MPEG-4 Waits in the Wings, *Cable360.net*.

CONTACT

John Schlack
Distinguished Member of Technical Staff
john.schlack@motorola.com

MOTOROLA

101 Tournament Drive
Horsham, PA 19044
Phone: (215) 323-1000

MOTOROLA and the Stylized M Logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners. © Motorola, Inc. 2008. All rights reserved.