Local Commercial Insertion in the Digital Headend

Mukta L. Kar, Cable Television Laboratories, Inc.
Sam Narasimhan, Motorola Broadband Communication
Richard S. Prodan, Terayon Communication Systems

Abstract

Existing ad insertion systems enable cable headends and broadcast affiliates to insert locally generated commercials and short programs in a channel seamlessly before delivery to the home. The revenue generated by local ads and short programs is very significant. Current ad insertion systems are hybrid systems, where commercials are stored in the ad server in digitally compressed form, while splicing is accomplished using analog techniques. The compressed commercial is decompressed and converted to analog format before insertion into an analog channel. The complete digital technology for delivery of compressed digital audio, video, and data is superior to existing analog methods. However, dealing with compressed video poses several challenges. One such challenge is splicing into a compressed digital bitstream compliant to the MPEG-2 standard seamlessly, that is to say, without adversely affecting the display due to decoding discontinuities at the insertion point. These applications involve insertion of locally generated compressed digital commercials and short programs into a digital channel containing previously multiplexed and digitally compressed programs.

The Society of Motion Picture and Television Engineers (SMPTE) has defined standard SMPTE 312M “Splice Points and Cue Messages for MPEG-2 Transport Streams” to promote interoperability of compressed digital video splicing equipment. SMPTE 312M is aimed at studio broadcast and requires a constrained (a synonym for pre-conditioned) MPEG-2 transport stream. For use in deployed digital cable systems, this will require significant changes to encoding methodologies to enforce predetermined conditions at the splice points.

The Society of Cable Telecommunications Engineers (SCTE) Digital Video Subcommittee (DVS) has established a splicing standard DVS 253 “Digital Program Insertion Cueing Message for Cable” to provide cue messaging and splicing in a more digital headend friendly manner, which does not require restrictions or constraints on MPEG-2-compliant transport streams. This paper will present a brief overview of existing analog/hybrid program insertion systems. The limitations of SMPTE 312M will be discussed. The solutions employed in DVS 253 will be described including digital headend friendly features. The implementation of DVS 253 to insert compressed commercials at the headend, including the issue of invisibility from commercial killers, will be addressed. As cost is related to complexity and flexibility of implementation, complexity issues and network operational constraints will also be examined for various implementations of digital program insertion systems and constituent components.

Introduction

Traditionally, the cable industry has been a popular conduit for the delivery of premium analog television content to the home. The access to the broadcast premium content was controlled using an analog set top box at the consumer's premises. The signal scrambled at the headend could only be descrambled using proprietary hardware embedded in the set top box. The advanced audio/video compression technology and subsequent standardization along with advances in digital modulation has made delivery of digital audio/video to the consumer's home a reality. The digital technology is very bandwidth efficient and provides superior video quality compared to its analog counterpart. In the last few years the cable industry has been deploying delivery of digital signals using a hybrid analog/digital set top box mainly for premium and pay-per-view (PPV) services.

Over the years, cable systems have grown significantly to become a primary conduit for the delivery of television contents to the majority of the homes. The industry also started inserting commercials both at national and local levels in some channels to enhance their revenues. Currently, the revenue generated from local ads and short programs contribute significantly to the total revenue of the cable industry. The revenue generated from local ads and short programs has been growing significantly.

Typically when a content is created, a few avails are included to insert commercials. A certain percentage of the avails are used for national ads and the rest for the insertion of local commercials. However, the networks (such as CNN, TBS, etc.) fill all the avails with the national ads before broadcasting to the headends using satellite. The headends can replace a percentage of these national ads with locally inserted commercials using the network cue signals. Basically the cue signal has three components, the Pre-roll (setup time), the Start (beginning of an avail) and the Stop (end of the avail). The Pre-roll is sent in advance to allow downstream devices at the headend (such as the ad-server and the splicer) time to prepare for the insertion. When the Start signal is received, the splicer switches from the network program to the commercial. After receiving the Stop signal the splicer leaves the commercial and returns to the network program.

Until a few years ago most of the ad-insertion systems were analog where the ad-insertion equipment is analog and analog storage devices were used to store the ads. However, the majority of the ad-insertion systems have been replaced by a hybrid one. In a hybrid system, the ads are stored in digital compressed format and the ad insertion equipment is analog in nature. In the hybrid system, the digitally compressed commercial is being decoded and converted to analog before insertion.

<u>Analog Ad Insertion</u>

In the analog domain, splicing between two NTSC video sources or TV programs is relatively easy. A switch between two synchronized video sources at any frame boundary (i.e. the vertical blanking interval) can be accomplished, since the frames are equal in size and time duration and are synchronized to one another. The resultant video, with the switched video source at the insertion or splice point, will appear seamless when displayed.

Figure 1 shows the block diagram of a typical commercial insertion system in cable headends and broadcast stations. In existing analog systems, networks distribute their content to cable headends and broadcast affiliates via satellite modulating a baseband bandwidth close to 10 MHz. As shown in Figure 2, the lower portion of the transmitted signal spectrum is used for NTSC video. The upper portion of the spectrum above 4.2 MHz contains a few FM subcarriers, which are used to transmit mono, stereo, multi-language audio, and data. One of the subcarriers is used to transmit cue-tone signals. The subcarrier is modulated with a cue-tone signal using frequency modulation (FM) or frequency shift keying (FSK). Some networks send cue-tones using a pair of FSK tones (e.g., 25 Hz and 35 Hz) within one of the audio channels. The integrated receiver-decoder (IRD) receives the RF-modulated signal, demodulates the cue-tone subcarrier signal, and provides an output signal. Either a dual-tone multiple frequency (DTMF) audio cue-tone sequence and/or a relay contact closure signal are provided to the insertion equipment. The networks send their schedules and spot availability (avail) in advance for local advertising insertion in a program. The precise location of the spot is signaled by a cue-tone during program broadcasts. The cue-tone consists of a sequence of numbers indicating the start and the end of an insertion opportunity. For example, the Discovery Channel uses a DTMF local avail cue-tone with 826* indicating the start, and 826# indicating the end of the spot.

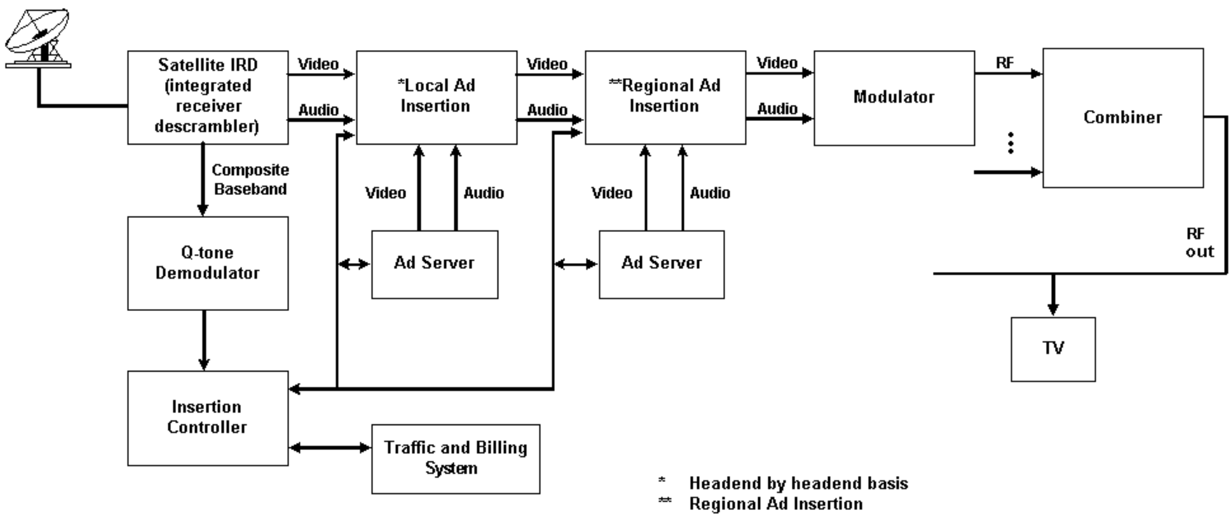

Figure 1. Typical Commercial Insertion System

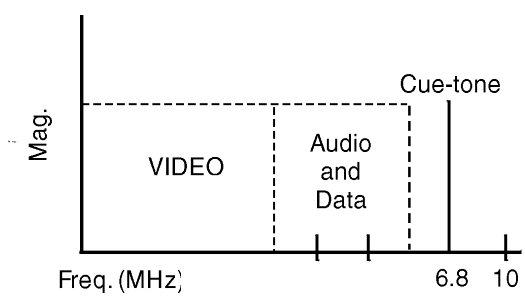

Figure 2. Analog Baseband Used by Networks

Digital Program Insertion

At the time of standardization of the MPEG-2 compression and system/transport layers [1], it had been expected that switching and splicing of compressed elementary streams and transport streams will be needed for editing commercial insertions, etc. Subsequently some hooks were provided for this purpose. After MPEG-2 standards were completed, the Society of Motion Picture and Television Engineers (SMPTE) decided to create a standard for splicing compressed digital streams. SMPTE completed the splicing standard in 1998 and is designated as SMPTE 312M [2]. SMPTE 312M is more suited to post production studios. Some of the major difficulties with SMPTE 312M are:

The splicer requires a constrained input bitstream.

Cue Signaling is not flexible enough. Cue signaling can only be sent to perform splicing on all the component streams of a program but not on a single component if needed.

No encryption mechanism for cue signaling was provided.

The impact of the requirement of constrained bitstreams for splicing alone is significant. For example, existing MPEG-2 compliant TV contents have to be pre-conditioned before any local commercials can be inserted into them. Also the installed encoders either have to be modified or replaced by the new ones. In essence, the standard is not very cable friendly for insertion of local compressed commercial in a cable centric statistical multiplexing environment [3]. The Society of Cable Tele-communications Engineers (SCTE) Digital Video Subcommittee (DVS) reviewed the SMPTE standard in detail and decided to work out a separate, more cable-friendly standard. SCTE established the digital program insertion ad-hoc group to work out such a standard in 1998. After working for more than a year, the ad-hoc completed their work, which has become DVS 253 titled "Digital Program Insertion Cueing Message for Cable" [5].

In creating such a standard no precondition or constraint for the MPEG-2 transport stream is assumed. The standard does not control the behavior of the splicing device in any way. The cue message carries timing information, which the splicing device may use to perform the splice. In other words, it may be treated as an opportunity for splicing but not a command. Cue messaging, if required, may be passed on to the downstream equipment such as a pass-through via a remultiplexer to a set-top box. The timing correction needed after remultiplexing may also be transmitted to maintain timing accuracy of the cue signal.

A cue message has basically two components – the first schedules the insertion points far in advance (such as the previous week). The second defines frame accurate insertion points. These messages are carried in an MPEG-2 private section syntax and called a splice_information section. A splice schedule defines the insertion time using coordinated universal time (GPS UTC) accurate to one second, while the splice insert uses MPEG PTS time, which enables frame accurate insertion.

Another important feature added to DVS 253 is the capability to encrypt cue messages. This will enable selective authorization of splicing devices in multiple headends. If a headend is not authorized, its splicing device will not be able to identify the insertion spot. The headends, which are authorized, will get the correct information about an avail.

In addition to these two message components, each splice command enables splicing of complete programs or individual components of a program (such as Video or Audio or Data). Today's systems use only the program level splicing where all the components of a program are replaced at the splice point. In the future, we will see programs and advertisements that are 'enhanced' using data broadcast as well as programs that offer interactivity. These systems will use the component level splicing where only selected components of a program are replaced at the splice point. The component level splicing enables pre-loading of data streams that are part of the same program by splicing the data component ahead of video/audio content. This may be done to load and run the data enhancements in the receiver's application engine. In the future, we will also see part of the bandwidth being used for 'opportunistic data' carriage and these components of the program are not spliced over during the insertion.

The components of a program are identified using the component_tag defined in the program map table using a new descriptor called the stream_id_descriptor. In this command mode, each component of the program can signal a different PTS time for insertion and the components that are not signaled are continued in the stream. When the inserted streams (such as enhanced commercials) have more components than the network program, the program map table will typically include the definition of these additional components even though they are not present in the original program. At the splice point, these components are enabled or disabled. Typical enhanced commercials could include delivery of discount coupons, prize drawings, or free software.

The use of the splice event_id and equipment interfaces such as between the splicer and ad_server are not defined. The DPI group has started working on the definition of APIs between the splicer and compressed commercial servers. DVS 253 along with standardized APIs will allow splicers and commercial servers to be interoperable.

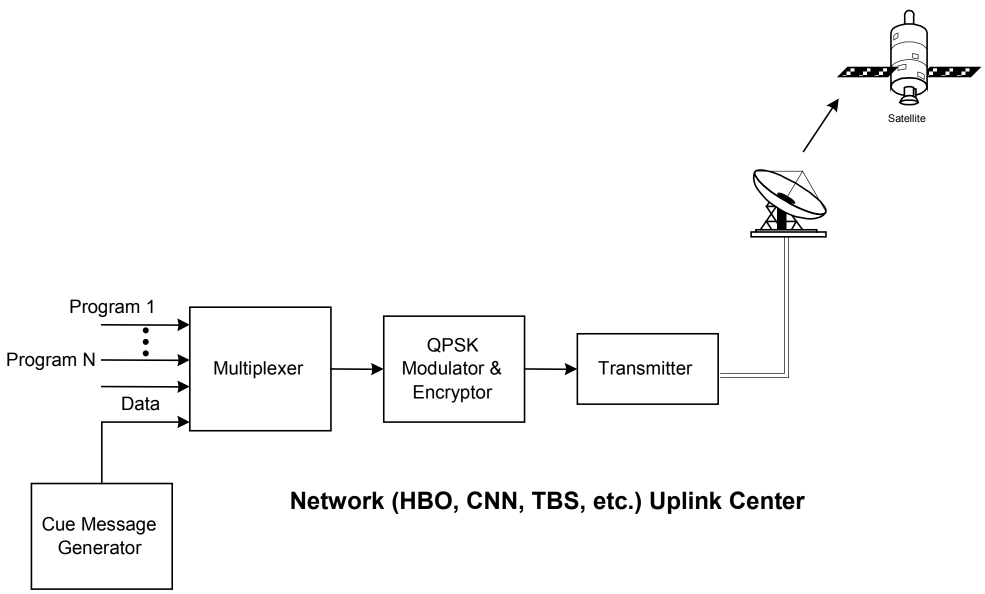

Figure 3. Digital Cue Messaging System (per DVS 253)

Figure 3 shows a typical block diagram depicting how cue messages will be multiplexed with a program before transmission to the headend via satellite. A Cue Message Generator generates a cue-signal per DVS 253. The digital cue message gets multiplexed with the other elementary streams to form an MPEG-2 compliant transport stream (TS). This baseband TS then can digitally modulate a RF carrier and be distributed to headends.

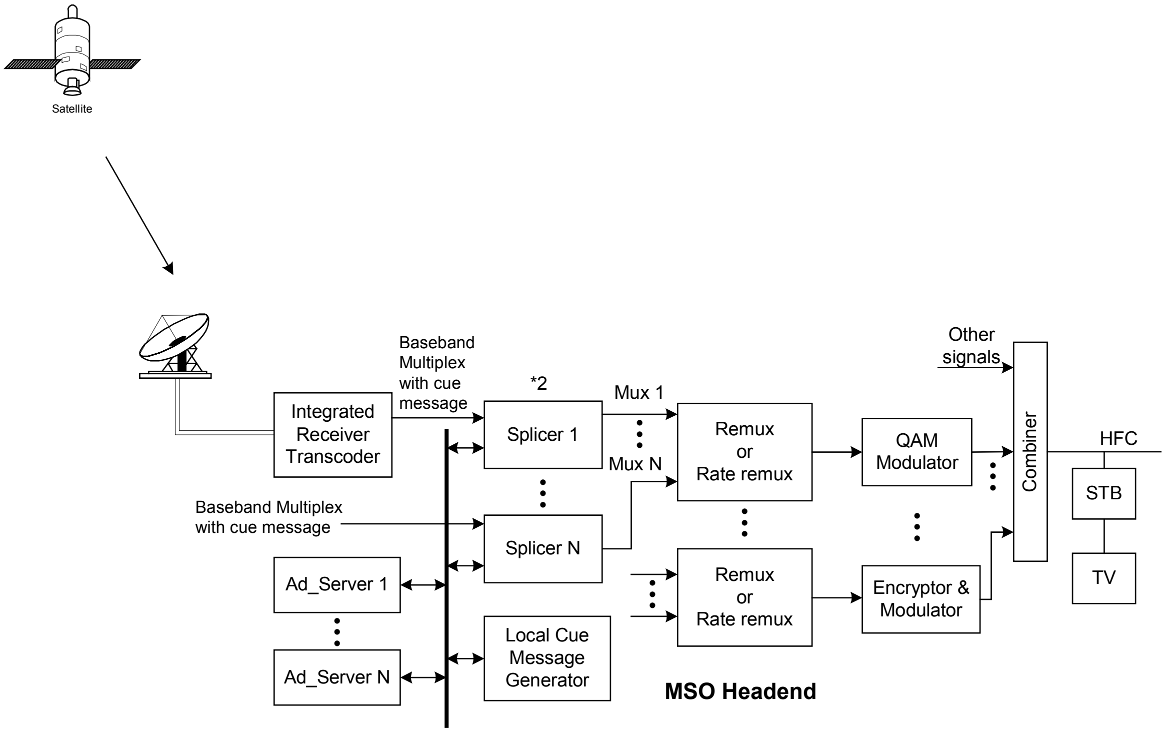

Figure 4. Typical Commercial Insertion System per DVS 253

Figure 4 shows the block diagram of a digital headend where local commercials are inserted into a program utilizing the cue message multiplexed in the transmitted program. The satellite integrated receiver demodulates the RF signal to its baseband MPEG-2 transport stream and decrypts it. The cue message, multiplexed in the transport stream is detected and decoded by the splicer. If it is a "pre-roll" it communicates with the ad_server to prepare. When cue message with a "splice_insert" is detected, the splicer switches from the network program to the local ad. At the end of the avail, the splicer switches back to the network program. A headend may have multiple splicers and ad_servers interconnected among them so that the same commercial need not be stored in more than one server. If for any reason a network cue signal is not detected or decoded properly, a local cue signal generator may be used instead. The local cue signal generator may be used as a backup system.

The cue message standard DVS 253 does not specify how to splice between two bitstreams. The techniques and resultant constraints for splicing compressed streams have been left to the innovation of the splicing equipment manufacturer. Splicing may be designed to splice seamlessly, near-seamlessly, or in a non-seamless way. Also, the compressed bitrate of the inserted ad may need to be reduced to the original program compressed bitrate to provide a seamless transition (i.e. the spliced stream cannot violate the MPEG-2 TSTD buffer model). Figure 5 shows the typical trajectory of a decoder buffer fullness. The constant transmission rate of the compressed video bitstream fills the buffer and the decoder removes variable size compressed frames (such as I, P or B) every display period. For the MPEG-2 main profile at main level (MP@ML), the minimum size of the decoder buffer must be 1.8 Mbits. Decoder buffer overflow or underflow should be avoided so that no audible or visual artifact is created.

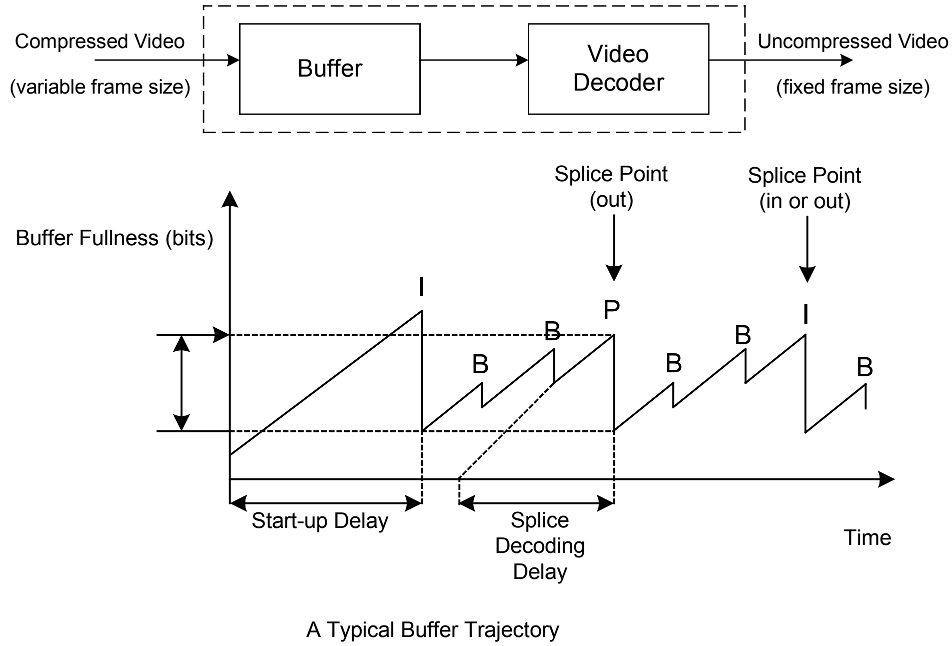

Figure 5. Simplified Representation of a Video Decoder

Splicing for Unconstrained, Seamless Insertion

Statistical multiplexing of compressed video streams helps to utilize digital channels better than constant bitrate encoded streams, assuming that the peak demand for bits from all video encoders does not coincide. The important constraint here is that all the videos need to be encoded while multiplexing. A previously compressed stream typically had to be decompressed and re-encoded before it can be inserted in a statistical multiplex. An alternative to such an approach to splicing compressed video streams is to use remultiplexing or "grooming" so that new services such as digital ad insertion may be added more efficiently without compromising visual quality due to cascaded decompression and recompression.

A *remultiplexer* receives one or more multiplexed streams as input and creates a new output multiplexed stream from local operator selected programs such as local ads. A few of the important functions performed by a remultiplexer are:

Demultiplexes the input multiplexed streams (unbundles the individual programs);

Creates a new multiplex out of the operator-selected programs and includes PSI for the new multiplex;

Maintains the bitrate constraint such that sum of all elementary stream bitrates, including PSI does not exceed the transmission channel rate;

Removes jitter in Program Clock Reference (PCR) time stamp values and maintains A/V synchronization within the programs;

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

Nominally, a remultiplexer does not alter the bitrates while constructing a new multiplex out of the input multiplexed streams. The technology that deals with the multiplexing of compressed video streams along with other streams (audio and data), and trims the resulting multiplex to match an assigned constant total transmission channel bitrate, is known as *rate-remultiplexing* [4]. Rate-remultiplexing meets the latter constraint by transcoding individual video stream within the output multiplex. Transcoding is the technique by which a compressed video stream is translated to a lower bitrate strictly within the compressed domain. Thus without cascaded compression, degradation in picture quality is not noticeable with occasional or moderate reductions in the average bitrate of individual video streams.

The above advantages of rate-remultiplexing make it a very attractive choice over standard remultiplexing. Rate-remultiplexing combines existing remultiplexing technology with a new capability known as transcoding. Transcoding can reduce the bitrate of MPEG-2 compressed video without fully decoding and re-encoding a bitstream with the attendant loss of picture quality inherent in such cascaded compression.

In creating compressed commercials, content producers can produce one high-quality version and store it in a server. But, based on the availability of digital bandwidth in the multiplex at various times and at various systems, the commercial's bitrate will have to be reduced during insertion. It is also possible to create different versions of commercials with different bitrates. However, storing different versions of the same commercial could be redundant if rate-remultiplexing is employed. Rate-remultiplexing technology provides the capability to insert compressed digital commercials into digital channels at the headend and removes the need to match the bitrate of the local compressed commercial with that of the remotely transmitted program, or creating and storing different bitrate versions of the commercials in the ad server.

<u>Related Standards Development</u>

Currently two work items are in progress in the Digital Program Insertion Ad-Hoc Group of the SCTE DVS Subcommittee – development of a standard splicing API and a splicing Guide for the use of DVS 253.

The goal of the API is to create a common interface for communication between the Ad Servers and Splicers for inserting ads into any multiplex. The API will be flexible enough to support multiple Ad Servers attached to one Splicer and one Ad Server attached to multiple Splicers.

The Guide will explain briefly the features of DVS 253. This will help not only to minimize the chance of misunderstanding of the features of the standard but will also reduce the miscommunication between the creators of DVS 253 messages (the networks) and the ad insertion equipment manufacturers.

ATSC specialist group T3/S8 has just completed the specifications for 'Directed Channel Change' (DCC). This is an extension to PSIP standard and specifies the signaling to automatically switch a 'virtual channel' based on the user preference at appropriate times. The specification also signals the end of this event so that normal programming can resume. This feature is expected to be used for applications such as targeted advertisements as well as enhanced advertisements. The time reference in this specification is based on UTC/GPS time and is only accurate to 1 second and the specification is very similar to the splice schedule command. Acquisition of the new 'virtual channel' as well seamless transition to the new channel requires splicer architectures that are very similar to what is presented in this paper. In this model, the streams for advertisement are expected to be within the same transport multiplex identified by a different virtual channel number so that the channel change can occur at the receiver. In the cable headend model the server is resident at the headend and the switching occurs in the headend instead of at the end user equipment.

Conclusion

In summary, the problems and solutions for flexible, unconstrained insertion of local digital compressed programming such as local ads into remotely downlinked digital video multiplexes has been described. The evolution of technology from previous analog and hybrid ad insertion systems to this new digital insertion/remultiplexing technology has been described. The features of the new digital program insertion cueing message standard for cable were summarized. Splicing technology using rate-remultiplexing for unconstrained, seamless local insertion and grooming of remotely compressed digital video programming was described. Related standards developments were noted.

References

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

2. *Splice Points for MPEG-2 Transport Streams,* SMPTE Standard 312M, December 1997.

3. *Digital Program Insertion for Local Advertising,* by Mukta Kar, Majid Chelehmal and Richard S. Prodan, NCTA Technical Papers, 1998.

4. *Rate-Remultiplexing: An Optimum Bandwidth Utilization Technology,* by Richard S. Prodan, Mukta Kar, and Majid Chelehmal, NCTA Technical Papers, 1999.

5. *Digital Program Insertion Cueing Message for Cable,* SCTE Standard DVS 253, December, 1999.

Acknowledgements

The authors would like to thank Carol Derr of AT&T, Joe Davis of Atlantek Electronic Development Corp. and other members of the Digital Program Insertion (DPI) Adhoc group of SCTE for their participation in the various digital program insertion technology development and standardization activities.