HEADEND MANAGEMENT OF MPEG TRANSPORT STREAMS FROM MULTIPLE SOURCES

Caitlin Bestler, Manager Architecture & Advance Development
Zenith Electronics, Network Systems

Abstract

Traditionally the CATV Operator has had the option of grooming the signals carried over their plants. Applications have included arranging the RF channel line-up, the logical channel line-up, Conditional Access tiering and local ad insertion.

Digital material is carried within MPEG Transport Streams, each with multiple services. Without the ability to separate those services the CATV Operator is faced with a choice between two undesirable options: surrendering control of many of these options to a remote satellite content provider, or paying for their own real-time MPEG encoding.

However there is a third option. MPEG Transport Streams can be managed in the headend, without decompression or recompression. With the techniques described the CATV Operator will once again have the tools to exercise per-service control over packaging the material carried on their network.

Techniques covered include selecting material from incoming MPEG Transport Streams, creating new MPEG Transport Streams from a mixture of remote and/or local sources, adding data and other services using the same downstream RF channels, and handling of local ad insertion in the digital domain.

BASIC TRANSPORT STREAM MIXING

The ability to obtain combine services from multiple MPEG sources is vital if CATV Operators are to maintain control over the services provided over their network.

This section will review the need to obtain material from multiple sources, the possibility granularities when mixing them, and the business implications of less than per-service mixing.

The Need For Multiple Sources

MPEG Material will be available from several sources to each CATV headend:

Satellite feeds will be available. Because of their wide delivery base satellites can be the most cost effective means of obtaining most digital programming. However for those very reasons it cannot be the source of all material. Satellites cannot carry local material, and certainly not local ads. They cannot even adapt schedules of nationally distributed material to local tastes.

Public SONET and/or ATM networks may be carrying MPEG Transport material primarily intended for Switched Digital Video (SDV) deployments.

Local Video Servers can be cost effective if their capacity is kept under control. When dealing with local ad insertion or the top titles for NVOD (Near Video On Demand) services they could prove to be very effective. However, the cost of storing hundreds of movies at **each** headend will simply not compare with the cost advantages of a satellite feed that is distributing its signal to a much wider market. Further, almost by definition, Video Servers cannot provide MPEG material for live events.

Real-time MPEG encoding is required to provide digital service when the material is received in analog format, for example with current off-air channels and existing local production equipment.

The need for multiple MPEG sources can be limited In a fully hybrid analog/digital network, supporting hybrid analog/digital set tops, along with analog only set tops and cable ready TVs and VCRs. Such a network has numerous advantages in preserving current equipment investments and efficient spectrum utilization.

However, there are several reasons why the CATV Operator may wish to provide services received in Analog format via MPEG to those customers with digital set tops:

- DBS and DVD (Digital Video Disc) marketing may succeed in creating consumer demand for "digital video" quality *per se*, not merely as a means of accessing more material.
- A digital-only set top box costs less than a hybrid analog/digital set top box.
- Digital conditional access and security is inherently superior to analog security.

Once they are being broadcast, off-air HDTV and/or SDTV signals will be MPEG Transport streams. Whether or not any portion of these channels are "must carry" services is not known. However they will be the easiest way to provide local broadcast services to customers with digital-only set tops.

Pure Data Sources will be available to supply data feeds independent of MPEG A/V programming on serial, ethernet or other data networking interfaces. Applications would include utility data for the set top Navigation and EPG programs as well as full Cable modems.

Selection Granularity

Given that the CATV Operator will be mixing and matching material from multiple sources the question is at what granularity: full 6 MHz channels or on a per service basis?

Material arriving from a satellite feed or other network will typically be pre-bundled into what the Satellite Distributor believes should be 6 MHz Transport Streams. This will typically be 28 Mbits worth of material, and could easily represent seven movie channels.

This is a fairly large chunk of material. Without the ability to separate these externally supplied feeds the CATV Operator is faced with a take-it or leave-it choice on the material.

Should CATV Operators' technology choices leave themselves in such a "all or nothing" position the Satellite Providers would be foolish if they did not take advantage of the situation and bundle weaker material with strong material.

The CATV Operator should retain the option of picking which services will be carried on their network. This can apply to whole services, portions of a multiplex, or even NVOD theaters. The supplied satellite feed may be showing movies every 15 minutes. If that is overkill in this market wouldn't it make sense to have the technological option of dropping every other showing to provide 30 minute NVOD?

Requiring local material and data services to be an even multiple of 6 MHz, carrying either 28 or 38 Mbits, would be especially wasteful.

Consider local community access programming. Will you **ever** need a full 28 Mbits of local community access programming per community?

Data services may be a valuable mix, but will they always consume an even 28 or 38 Mbits capacity? This is especially true for HFC plants where data services can be carried on a per-node basis.

LOCAL CONDITIONAL ACCESS

The CATV Operator should not only retain control over which of their subscribers is authorized for what material, but also on how that material is packaged. What services are part of what tiers?

Without these capabilities the CATV Operator is reducing their role to that of a pipe provider, or at most an order taker. Traditional leading roles in marketing material and building a customer base will be lost.

Obviously these issues must be negotiated with the Copyright Holders of the source material, but without the capability of controlling tiering locally there isn't anything to negotiate about.

Local Ad Insertion

Local Ad Insertion should remain a major part of the business plan even as material shifts from analog to digital.

CRITICAL ISSUES

Given that the ability to control and manipulate MPEG Transport streams is desirable, the question then becomes what types of equipment should be considered for these jobs.

This section will introduce some of the technical issues in handling MPEG Transport streams as a lead-in to discussing possible solutions.

Split, Merge or NxM Multiplexing

The next question is how much mixing of Services is required. The basic options are:

- **Split Only**: (1 to N) each incoming Source Transport Stream is split into N outgoing Transport Streams on the basis of services. Each outgoing Transport Stream contains material solely from a single incoming stream, but not all of it. Additional data services can be fill the RF channels capacity.
- **Merge Only** (N to 1) each outgoing Transport Stream combines material from N incoming Transport Streams, dropping specific services as configured to fit within the channel's capacity.
- **Full Mix** (NxM) any service from any incoming Transport Stream can be placed within any outgoing Transport Stream, subject solely to the capacity of the outgoing Transport Stream.

The Merge and Full mixing options require the ability to relabel MPEG Program Numbers and PIDs. Further, since these numbers are themselves referenced within System Information tables, those will have to edited and/or replaced to carry the modified references.

Any downloaded applications that reference MPEG Program Numbers or PIDs directly, without going through the proper directory access procedures may pose a problem. This can be solved by restoring all MPEG Program numbers and PIDs for the selected MPEG service either while demultiplexing or through software illusions, or preferably just by simply refusing to support the practice.

When splitting an MPEG Transport Stream based upon service it may be necessary to duplicate certain elementary streams. Of course this should be avoided whenever possible, and never done if it requires duplicating a video elementary stream.

The likely conditions when it would occur are for data streams associated with multiple services. A data stream associated with an NVOD service would be an example. Each showing might want access to the data stream, but in the bandwidth allocation chosen the NVOD service may have been split between two RF channels.

Far more common will be sharing of the PCR tracking PID. Depending on what PID was selected, the selected set of services may or may not have use of the data on the designated PCR PID. If not, the ideal split would copy only the adaptation headers containing the PCR values and reduce the total bandwidth required.

The choice between Split, Merge and NxM mixing has a direct impact on redundancy planning. Split and Merge modes can support N+1 redundancy schemes only at the level of the entire set of equipment supporting an RF channel.

If one real-time MPEG encoder fails either the entire set of equipment supporting an RF channel must be switched to the hot standby equipment set, or there must be full (N+N) redundancy on the MPEG encoding equipment.

With NxM mixing each component of the system can have N+1 redundancy. Whenever work is switched to the hot standby it's inputs/outputs can be routed as required.

Mixing Is More Than Packet Switching

As you have probably already read many times each MPEG Transport Stream is a multiplex of several Elementary Streams. At first glance that makes the problem of manipulating MPEG Transport Streams appear to be very similar to other packet-switching applications such as ATM.

However MPEG Transport Streams cannot be "switched" in the way that ATM can. That is because each Transport Stream is not only a collection of Elementary Streams, it is a **self described** and **self-timed** collection of Elementary Streams.

Transport Stream Self-Description: System Information

Not only are there Elementary Streams carrying Video, Audio and Private Data for applications, there are other Elementary Streams that describe what Elementary Streams are present. Further that self description can take place at two different levels: a here-and-now level to support basic user navigation and another to support Electronic Program Guides.

All of this self-description is carried in a set of tables on various elementary streams. The MPEG Transport standard defines a standard format for download of tables that allows for updating of tables in sections, identification of new versions of tables and even pre-loading of 'next' versions of tables before they are required.

This self-description of an MPEG Transport Stream is contained in tables such as these:

- The Program Association Table (PAT), which is always carried in PID 0, identifies the Program Map Table (PMT) stream for each MPEG Program (or service) carried in this Transport Stream. It also identifies the PID where to Network tables will be found.
- The Conditional Access Table (CAT), always carried in PID 1, identifies where EMM messages are to be obtained for each service in this Transport Stream.
- Program Map Tables (PMTs) are used to specify what elementary streams are used for what portions of an identified service. For example what PID is the video in, etc.
- Program Information and Program Name tables are placed in the same stream as the PMT and supply further display and navigation information about the current program.
- The Network table stream is used to provide information about services available on other Transport Streams in the system and maintenance functions such as the System Time. Virtual Channel lineups would be included here.
- Electronic Program Guide data will usually be supplied as a Private Data feed as part of a pre-defined service.

Dynamic vs. Canned System Information

When creating a new Transport Stream from multiple sources, new System Information must be created for the new Transport Stream.

The old System Information cannot just be simply merged. If the multiple sources are truly independent they may have chosen the same PIDs or MPEG Program Ids. These numbers must be re-assigned to avoid conflicts.

Most current equipment solves this problem by the following steps:

- Each incoming packet is either thrown away or relabeled to the correct PID. All incoming System Information is thrown away.
- New System Information is inserted with the correct PIDs.

This approach is dependent on having correct information delivered off-line **in advance**. It requires a highly co-operative relationship with the source of the MPEG Transport stream.

Should the DBS system adjust its schedule dynamically in anyway, even for something as trivial as the PID selection, then the inserted System Information will be incorrect.

This approach also requires a cumbersome distribution of Program Guide information. It can be argued that as long as viewers rely upon printed schedules that the data cannot be updated that rapidly. However, even while maintaining compatibility with printed schedules the frequency of NVOD showings can be adjusted transparently to the user. An unexpectedly popular title could be bumped up to every 15 minutes, while another that is proving to be more of a cult favorite could be dropped to every hour.

There are ultimately only a few choices:

- Contractually lock your source to adhere to a precise schedule that must be provided off-line well in advance. Since this may restrict their flexibility in supporting other customers, such as DBS, expect negotiations to be difficult and/or expensive.
- Provide a separate real-time communication path from the Satellite headend to your headend and contractually obligate the supplier to provide prompt notification of all changes via this link. Not only will this link be expensive, but expect the contract negotiations to be even worse. The DBS Operator is now agreeing to maintain and

keep operational extra equipment that is not relevant to their DBS customers.

- Parse the System Information already being provided to the DBS boxes within the incoming Transport Streams.

Requiring CATV headend equipment to be able to parse the incoming System Information is clearly the only viable long-term solution.

Dynamic Source System Information

Real-time parsing of incoming System Information can be used in two different ways.

Canned System Information can still be used, only all material is identified relative to the System Information. The source of the output Transport Stream is no longer Port X PID Y, but Port X Virtual Channel Y's video PID. Should the mapping change, the new relabellling will adjust automatically, and hence no new system information is required.

However the ultimate solution is to merge the System Information base on CATV Operator specific rules. For example the CATV Operator would specify which MPEG services, or Virtual Channels were desired from each source on each output Transport Stream. The MPEG handling equipment would have to create the merged System Information on the fly from the source material.

Divergent System Information Formats

The MPEG Transport standard itself only deals with a portion of the System Information tables.

Other standards, such as DVB or the proposed ATV standards are required to define the exact set of tables in use. The earlier examples of System Information tables was based upon General Instrument's proposed System Information for Digital Television.

Dealing with multiple sources will mean dealing with multiple formats for System Information. With luck, only the ATV and DVB formats will be required, but there is always the chance of encountering proprietary formats, particularly for Electronic Program Guide data.

No matter how many formats System Information arrives in, the system set tops should only deal with a single format. The logic to deal with multiple formats would be error prone and drive up set top Flash and/or DRAM requirements. It is far preferable to require the information to translated in the headend CATV equipment to one of the standards based formats.

Variable Bit Rates

Putting together a Transport Stream from multiple sources requires knowledge that the proposed streams will fit within the capacity of the 6 MHz channel, whether it is 28 or 38 Mbits.

When the Source Transport Stream is statistically multiplexed there are some special considerations.

For each set of programs that are being switched as a group the following must be known:

- The **Sustained Bit Rate** for these service: what bit-rate must be available on a sustainable long-term basis for these services.
- The **Peak Bit Rate**: what is the maximum bit rate that these services will ever reach. Also associated with any peak rate is a defined period over which the actual rate must fall back down to the Sustained Rate.

For variable delay services, such as data, the two rates are a **Guaranteed Rate** and a **Sustainable Rate**. The Guaranteed Rate is lower than the Sustainable rate, meaning that for any short period the data service may be restricted to this lower rate, but it can count on catching up to the Sustainable rate within the defined time window. The headend equipment must offer sufficient buffering to hold the extra material until it can catch up.

In order to plan an output Transport Stream the component services must be picked such that the available capacity is not exceed by either.

- the sum of all Guaranteed and Peak Bit Rates.
- the sum of all Sustained Rates.

This solution complements the traffic requirements of Data and MPEG services naturally. MPEG Transport benefits from variable bit rate encoding to deal with more difficult scenes. The Peak Bit Rate is always greater than or equal to the Sustained Bit Rate. Data services, by contrast can survive with occasionally delayed delivery as long as the aggregate Sustained rate can be guaranteed.

Conditional Access

MPEG Transport Streams received from other sources may contain their own Conditional Access logic. According to the MPEG Transport standards Conditional Access is enforced in two types of elementary streams: ECM and EMM.

EMM (Entitlement Management Messages) streams are used to entitle individual set tops to receive (or not receive) specific services and elementary streams. This typically involves editing authorization bitmaps and/or lists and distributing decryption keys. EMMs typically individually addressed.

ECM (Entitlement Control Message) streams are used to control the decryption of other elementary streams. They enable set tops that have been previously authorized by EMMs to correctly decrypt the controlled material.

How these two stream types are handled will determine how much control the CATV Operator will have over the Conditional Access packaging of material on their network.

The ability to substitute EMM streams is required to have control over which set tops are authorized. However the EMM stream can be delivered out-of-band, even though doing so makes things easier for any cryptographic attackers.

ECM streams, however, must be delivered in-band within the MPEG Transport stream. Unless the ECM streams are **locally generated** the tiering of that Transport Streams material will be dictated by the source.

The CATV Operator would loose control over decisions such as what was part of an Extended Basic tier.

Transport Stream Self-Timing

MPEG Transport Streams are self-timed. Material that is jitter-intolerant, such as the Video and Audio, carry PCR (Program Clock Reference) values. PCRs allow the receiving MPEG decoder to lock its 27 MHz clock to the source clock.

An MPEG decoding subsystem can only lock to a single source clock at a time (of course a set top could theoretically contain multiple MPEG decoding subsystems, but each would still only be capable of locking to a single clock). All PCR sensitive streams for the selected service must be based upon the same clock.

When a single Transport Stream contains services from multiple sources there will be more than a single clock within the Transport Stream. This is allowed in MPEG Transport syntax, although it is seldom implemented that way when a single back of MPEG encoding equipment generates the entire Transport Stream.

Any form of multiplexing inherently causes at least a minor amount of jitter. Ideally the headend equipment should minimize the jitter it causes, if possible dejitter any jitter caused by ATM or other networked delivery of material, and properly not any jitter it caused by adjust PCR values in the packets it handles.

Audio Format Mixing

Sources carrying fully legitimate standards-compliant MPEG-2 Transport Stream can still have incompatible audio formats: some may be in MPEG or Musicam format, others will be in Dolby AC-3 format.

There are only two sure solutions to this problem:

- Restrict material selection to sources that includes the audio format you have selected. Some material may be available with dual audio streams to address this problem.
- Specify set tops that are capable of either. The cost premium of this option will continue to decline, and will eventually be the unbeatable solution.

It is **theoretically** possible to convert between the two formats without altering the matching video stream. Such a product would be highly desirable, if anyone were working on it. However I have yet to encounter even one engineer who claimed that it would be easy. Given the general optimism of engineers it would be prudent to rely on one of the first two options.

POSSIBLE SOLUTIONS

Full mixing of MPEG Transport Stream requires a vast amount of processing and communications horsepower.

Each incoming Transport Stream must:

- be captured according to the framing and error detection/correction of the delivery protocols. This could be a simple CRC as from a Satellite, or it could be ATM AAL-5 from an ATM network.
- packets must be relabelled and routed to the correct output Transport Streams.
- be tracked as to maintain lock with its 27 MHz clock and ensure that the rates supplied are within the configured range.
- have all incoming System Information tables parsed.
- dejittered as necessary and possible. This function is dependent on the type of network port material is being received from.

Each outgoing Transport Stream must:

- merge portions of the incoming Transport Streams routed to it.
- make PCR value adjustments to reflect any unavoidable jitter it creates.

Additional processing is required to update relabelling and routing instructions based on the parsed incoming System Information and/or to rebuild the new System Information to be output.

The solution should also be scaleable. The number of Transport Streams requiring mixing may start low and be expanded as time goes on. The solution should not require buying equipment capable of remixing 70 RF channels when only two will be required.

Add/Drop Remultiplexers

The first generation of solutions are variants on the equipment used to generate the original Transport Streams after real-time MPEG encoding.

This equipment typically is capable of relabelling PIDs and re-timestamping PCRs. Management plane interfaces must specify how each PID is to be handled, and supply new System Information for insertion.

To support dynamic tracking of sources these devices would have to learn how to parse and kick-out updated tables for other headend equipment. That equipment would then interpret the new tables and issue modified instructions to the Mux.

Each remultiplexer generates a single output Transport Stream. Implementation of Merge mixing is relatively easy. Each Remultiplexer is simply overfed, and drops PIDs as configured.

Splitting can also be performed on a limited basis. Each remultiplexer simply drops a different streams.

To accomplish full NxM mixing extensive splitters and wiring is required. Each remultiplexer would typically be overfed several times the material it required and drop virtually all of it. Such solutions are scaleable, but only to a certain degree before full ATM switching equipment is required in addition to the remultiplexers for each RF channel.

Backplane base NxM Mixers

An alternate approach is to perform the Transport Stream mixing function in a cardrack with processor cards connected on a backplane.

Modular processing cards communicate with each other over the backplane, eliminating the need for splitters and special wiring between each source and each multiplexer.

Each processing card deals with the processing requirements of a single input or output transport stream (or if very powerful possible two). Processing power need only be added as required.

Use of an industry standard backplane, such as PCI or VME-64, would allows easy addition of more processing power later. Network connection cards can be obtained that are not developed solely for this project. The processor cards can be replaced with later more powerful cards as new applications are found and/or costs per MIP drop.

Two types of backplane architectures are worth consideration: switched and shared. Switched backplanes allow multiple processor cards to feeding material to other cards concurrently. This is preferable, but currently results in far higher costs on a per port basis.

The shared backplane approach requires blocking backplane transfers, resulting in a fixed delay on the order of 100 msecs on all switchable services. With careful design portions of the multiplex which are not switched can be passed through without this fixed delay.

ADDITIONAL POSSIBILITIES

Local Ad Insertion

A processor based Transport Stream handler would naturally be capable of handling functions such as Ad Insertion in the digital domain.

However, it can also do so in a way that could radically reduce the cost of interfacing with an Ad Server.

Storing local Ads in digital format is already desirable for random access to them. More ads can be available for insertion, with no restrictions on which can be played in what order or when.

A Transport Stream Mixer as described, would be capable of handling a more flexible and cost effective interface. After loading a very short portion of the Ad into its memory, the Transport Stream Mixer could **pull** the remaining material at the clock rate of the overlaid material. This material could now be pulled over 100 Base T or switched 10 Base T network lines from a plain File Server. No special Video Server or OC-3 ports would be required.

Data Routing

Data received from network ports can be encapsulated within MPEG Transport packets and sent downstream as opportunistic data.

Additionally a Transport Stream Mixer could perform several routing functions:

- Tracking acknowledgments.
- Routing packets to an in-band Transport Stream when it knows the set top is tuned there and there is available capacity, rather than using up the limited capacity on any out-of-and Transport Stream.

Optimized Data Circulation

Many tables downloaded via MPEG Transport are circulated as a data carousel allowing the set tops to collect the data as required and/or able.

This type of data repetition is required to deal with plant noise, the set top being busy while changing channel, loss of power to set tops and to minimize the amount of DRAM required in each set top.

However it is hardly desirable to provision network inputs to accept these data feeds at their full speed. Accepting only the content of each data carousel table, its priority and duration would greatly simplify the interfacing between external servers, management systems and the MPEG Transport handling sub-systems.

Feedback to Video Server

The output of a Transport Stream Mixer could be routed back to a Local Video Server, rather than just being delivered.

This could allow for delayed rebroadcast of local material, capture of satellite distributed ads, and capture of new material for locally controlled NVOD re-broadcast.

CONCLUSION

The ability to handle MPEG Transport Streams within a CATV headend can provide valuable options to the CATV Operator. A single system based upon standard interfaces can provide an upgradable cost and space effective solution.