



2114 West 7th Street
Tempe, AZ 85281 USA
Voice +1.480.333.2200
E-mail sales@comtechefdata.com
Web www.comtechefdata.com

Carrying Real-Time Mpeg-2 Transport Streams Over Satellite

April 2009

Introduction

The Moving Picture Experts Group (MPEG) 2 transport stream¹ is the adopted standard for transporting audio, video and data over broadcast mediums such as satellite, cable and terrestrial broadcast systems. Additionally, broadcasters have adopted the Digital Video Broadcasting (DVB) throughout the world and Advanced Television Systems Committee (ATSC) in North America for the transport of multimedia. Both DVB and ATSC have accepted the MPEG-2 as the transport medium.

The term, *transport stream* (or simply “TS”), is an industry term that refers to technology described in the ISO/IEC 13818-1 specification. This white paper will address issues that must be considered when carrying multimedia content over the TS; and more importantly, over a satellite link.

The MPEG-2 transport stream does not specify an electrical or physical format; but throughout this paper, multiple physical and electrical interfaces are introduced. The most popular interfaces for carrying an MPEG-2 transport stream are the Asynchronous Serial Interface (ASI), Synchronous Parallel Interface (SPI) and Code of Practice #3 (CoP#3) IP technologies.

The MPEG-2 Transport Stream

The MPEG-2 transport stream provides a multiplexing point where content (video, audio, data, table information, conditional access, etc.) is accepted and placed into 188-byte frames formatted into a transport structure as shown in Figure 1.

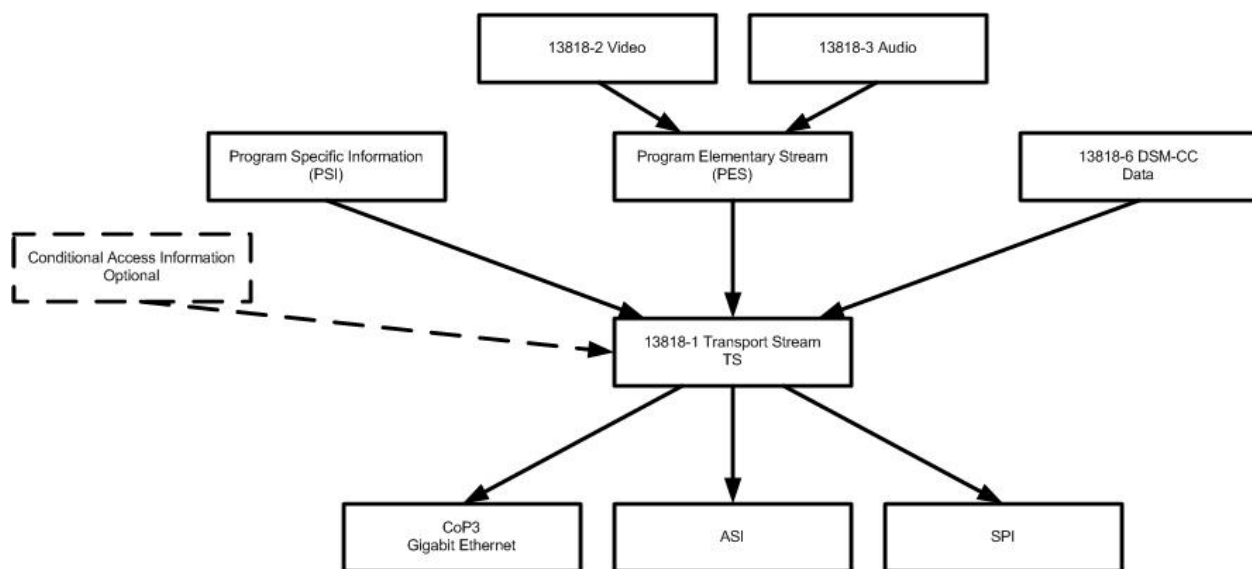


Figure 1 MPEG-2 Data Flow

The multiplexing point accepts and constructs an MPEG-2 transport stream with the following information:

- Program Specific Information (PSI) – Information about the content carried in the transport stream. The PSI tables provide information about the program identifiers (PIDs) and the programs being transported over the transport stream. The PSI is not

¹ The MPEG-2 transport stream specification is defined by the ISO/IEC committee and is specified as International Organization for Standardization/International Electrotechnical Commission ISO/IEC 13818-1. The 13818-1 specification defines a mechanism for the transport of multimedia content, information about the content and a provision to transport timing over a broadcast medium.

mandatory; but in the event a PSI is not present in the transport stream, all PIDs carried in the transport stream to the distant end without information are commonly known as “ghost PIDs.”

- Conditional Access (CA) Information – Information about conditional access information being carried over the transport stream. This information is typically carried in the Conditional Access Table (CAT). The CAT is an optional table, so it is only required if there are conditional access requirements on the content being delivered.
- Network Information - Information about any transport specific information as it pertains to the modulation or specific format for a particular medium. This information is typically carried in the Network Information Table (NIT). The NIT is an optional table.
- Content - The content (information) transported over the MPEG-2 transport stream falls into one of the following areas:
 - Program Elementary Stream (PES) contains:
 - Video defined by the Generic Coding of Moving Pictures and Associated Audio Information – part 2 - Video ISO/IEC 13818-2 specification
 - Audio defined by the Generic Coding of Moving Pictures and Associated Audio Information – part 3 - Audio ISO/IEC 13818-3 specification
 - Data defined by the Extensions for Digital Storage Media Command and Control (DSM-CC) ISO/IEC 13818-6 specification. The 13818-6 specification covers data transport via the Multi-Protocol Encapsulation (MPE) for IP datagrams, data piping for synchronous and asynchronous streamed data, and data/object carousels for repetitive data.

The MPEG-2 transport multiplex exiting the diagram in Figure 1 results in all content being parsed into 188-byte cells (MPEG-2 transport stream frames) that consist of a four-byte header as shown in Figure 2.

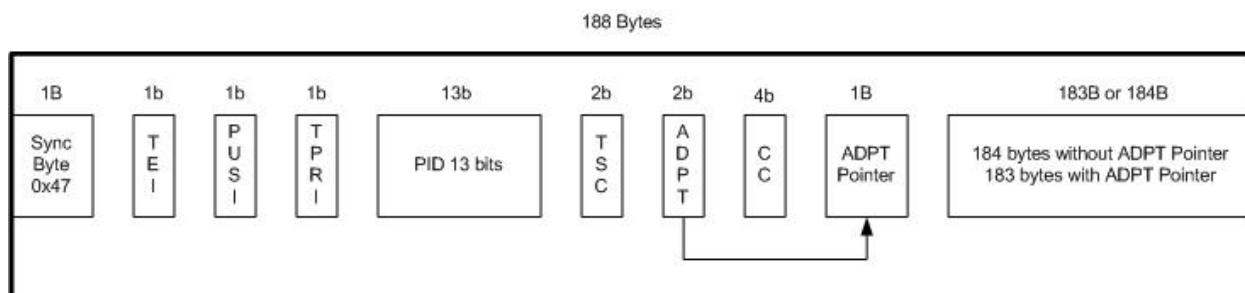


Figure 2 MPEG-2 Header

- Synchronization Byte = 0x47 – 1 byte
- Transport Error Indicator (TEI) – 1 bit
- Payload Unit Start Indicator (PUSI) – 1 bit
- Transport Priority – 1 bit
- Program Identifier (PID) – 13 bits (valid from 0x0000 to 0x1FFF)
- Transport Scrambling Control (TSC) – 2 bits
- Adaptation Field Control (ADPT) – 2 bits
- Continuity Counter (CC) – 4 bits
- ADPT Field - If the Adaptation Field Control is '01' or '11', then an adaptation byte will be present, and payload will be 183 bytes instead of 184 bytes per frame
- Payload – 183 bytes if adaptation pointer is present and 184 bytes if no adaptation pointer is present

An important note about the PIDs contained in the transport stream: certain PID values are pre-assigned and may not be used for general assignment by the user. The pre-assigned PIDs are as follows:

- PID 0x0000 – Program Association Table (PAT) – this is the first table as part of the Program Specific Information (PSI) tables
- PID 0x0001 – Conditional Access Table (CAT) – this is the first table as part of the conditional access information tables
- PID 0x1FFF – Fill PID
- It is noteworthy that DVB and ATSC standards allocate PIDs for specific use, but this is beyond the scope of this document

All other PIDs are open to be used as outlined in the MPEG-2 TS specification. However, it should be noted that both ATSC and DVB assign PIDs between the ranges of 0x0002 and 0x1FFE that are used for specific purposes for signaling, but that discussion is beyond the scope of this document.

Delivering Transport Stream

To transfer the stream from the “logical layer” to the electrical and physical layers, the following interfaces are established throughout the industry:

- Asynchronous Serial Interface (ASI) – designed for interfacing professional equipment to MPEG-2 transport streams and defined by the European Standard EN 50083-9. The interface provides a point-to-point connection with an 8B/10B Manchester encoded stream running at 270 Mbps, of which 214 Mbps is available for the transport of data.
- Synchronous Parallel Interface (SPI) – designed for interfacing professional equipment to MPEG-2 transport streams and defined by the European Standard EN 50083-9. The interface provides a point-to-point, byte-wide synchronous interface with a clock. The clock supports up to 13.5 MHz for a maximum transport rate of 108 Mbps. This interface is fading from use.
- Code of Practice #3 (CoP#3) – designed for interfacing professional video equipment over an IP enabled network as referenced in the Pro MPEG CoP#3 release 2 document and defined by SMPTE 2022-1-2007 and 2022-2-2007. Traditionally, given the higher rates video requires, the interface has been standardized as a Gigabit (1000 BaseT) interface.

To transport the MPEG-2 transport stream over any of the data interfaces requires a stream running at a constant rate. The need for a constant-rate stream requires that content (video, audio, data, tables, etc.) must be at a fixed rate, and this is typically not guaranteed. Even with constant bit rate (CBR) video, audio and data, the transport stream varies due to the addition of table, timing and content fluctuations. Therefore, to compensate for the variable nature of the incoming content, the concept of a “fill packet” is introduced. The fill packet allows variable rate content to enter the MPEG-2 transport stream multiplexing process as shown in Figure 1, and is shown leaving as a constant rate stream. The stream is kept constant by adding MPEG-2 fill frames to the transport stream and adjusting any time-sensitive frames to ensure minimal Program Clock Reference (PCR) jitter is added — PCR is discussed later in this document. The construction of a transport stream is shown in Figure 3.

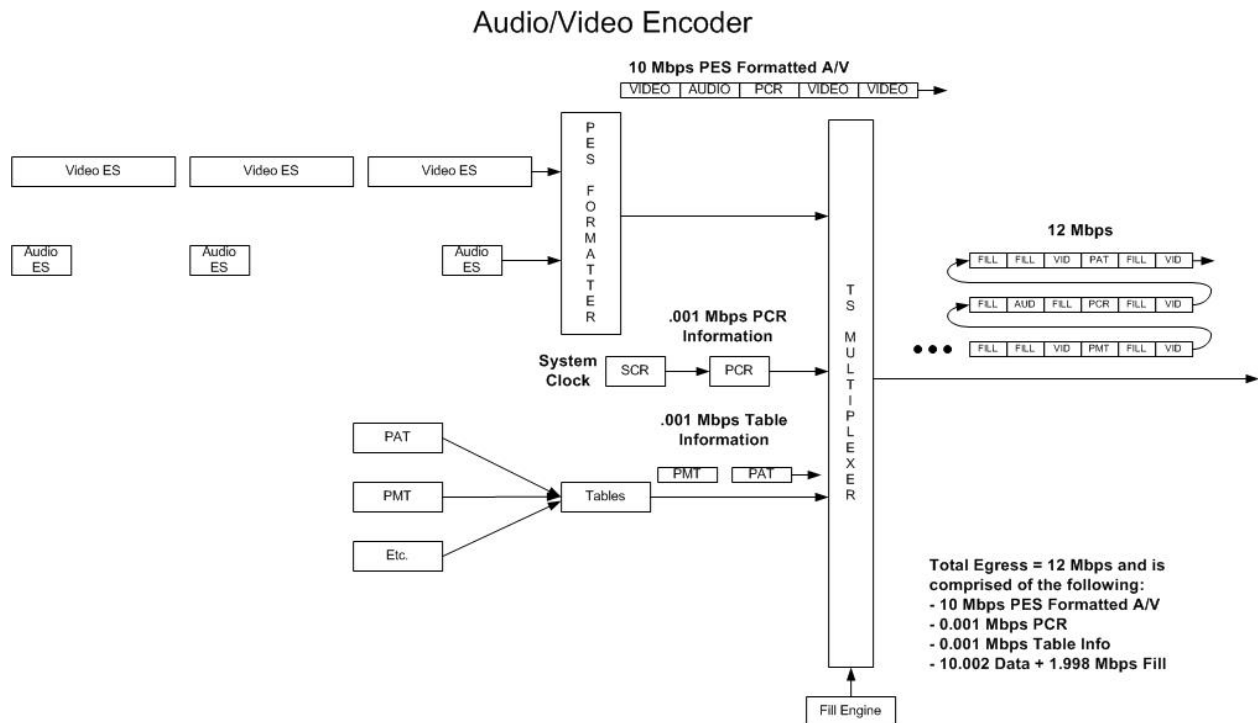


Figure 3 MPEG-2 AV MPEG-2 Example Encoder

Figure 3 shows how an MPEG-2 Audio/Video encoder creates a 12 Mbps transport stream containing the following:

- Program Specific Information (PSI) comprised of a Program Association Table (PAT) and a Program Mapping Table (PMT):
 - PAT assigned to PID 0x0000
 - PMT user defined: PID 0x0002 to 0x1FFE
- Video - Comprised of the video elements (Elementary Stream – ES) of the transport stream with a user defined PID 0x0002 to 0x1FFE
- Audio - Comprised of the audio elements (ES) of the transport with a user defined PID 0x0002 to 0x1FFE
- Program Clock Reference (PCR):
 - Case 1: PCR assigned to the video PID – in this case, the adaptation field is set in the MPEG-2 header, so the PCR is placed at the beginning of an MPEG-2 frame's payload carrying the video frame
 - Case 2: PCR assigned to a separate PID – in this case, a single MPEG-2 transport stream frame carries the PCR timing information. This results in wasted bandwidth, since 90% of the MPEG-2 frame is unused. The PID is user defined PID 0x0002 to 0x1FFE
- Fill PID is comprised of an MPEG-2 transport stream frame filled with 0xFF bytes. The PID is assigned to 0x1FFF

The key to understanding the MPEG-2 transport stream is knowing that the PSI tables provide all information about the contents of the transport stream. Figure 4 shows how the PSI tables relate to each other.




Figure 4 PSI Table Structure

The PAT and CAT are the roots of the tree, and all information about the multiplex is based on a starting point of these two tables. Example 1 shows how the PSI information can be used for a remote device to determine where available information is located in the transport stream.

Example 1: PSI Tables

Suppose the network operator configures the system as follows:

- One Multiplex
- 2 Audio/Video Programs
- PCR is NOT assigned to the Video PID
- 1 Data Program (13818-6 MPE/IP)
- No Conditional Access (no need for the CAT, EMM or ECM tables)
- PMT base PID is assigned as 0x0100
- No NIT is required, since this is a closed network

The resulting tables would be created, as depicted in Figure 5:

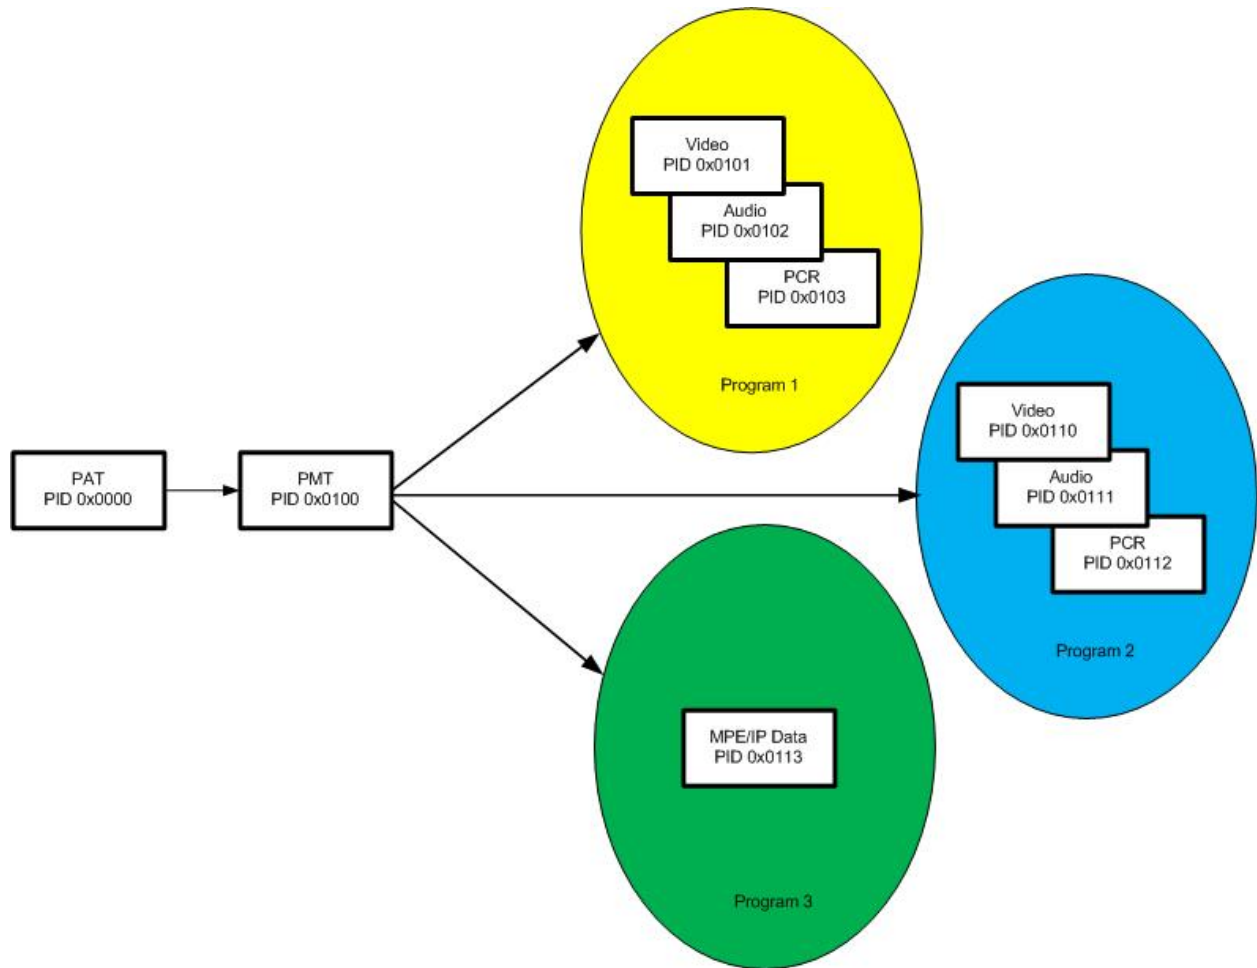


Figure 5 PSI Table Structure without Conditional Access

Manipulating a Transport Stream

Care must be exercised when supporting transport streams where the transport stream is sped up, slowed down, or merged with another transport stream. The technical issue for altering a transport stream becomes a potential problem when the transport stream contains real-time content. Real-time content is defined as carrying video and/or audio referenced to a PCR. The timing of the transport stream may not be altered such that would result in jitter by more than ± 500 nS from the encoder (origination point) to the decoder (destination point where the stream is being decoded and displayed). A Single Program Transport Stream (SPTS) implies a single stream is being transferred from the source to the destination. In the case of the SPTS, there is no need to alter the rate. However, combining one or more streams creates Multiple Program Transport Stream (MPTS). Combining transport streams, either SPTS or MPTS, results in the PCRs having to be re-stamped. Figure 6 shows the reason for providing PCR re-stamping.

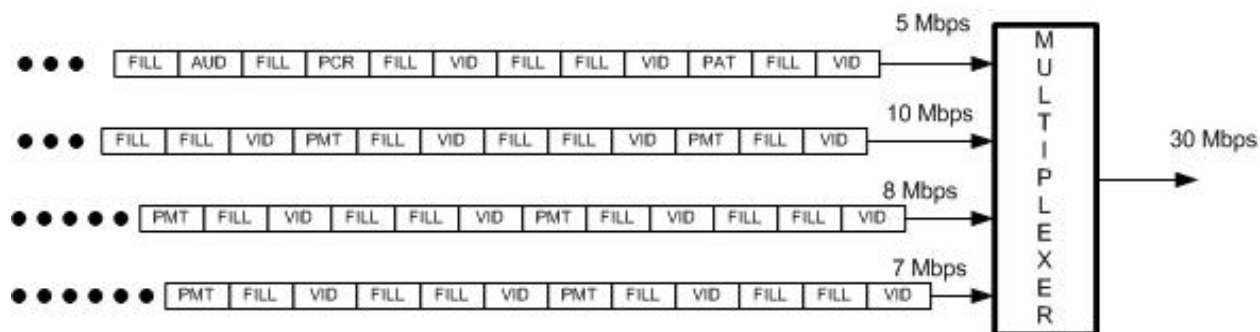


Figure 6 MPEG-2 Multiplexer

Figure 6 shows that each SPTS arrives at different rates and times, and each of the four transport streams contains real-time content (PCR frames). Because each stream is independent of the others, the location of the PCR packets for each stream is not temporally related. When combining (multiplexing), this must be taken into consideration, because the time required to merge all streams to a single stream results in jitter for all PCR packets. To address the PCR jitter, the PCR packets are re-stamped with a time correction (amount of time held in the multiplexer) and sent into the multiplexed (combined) transport stream to the egress as an MPTS. The act of combining multiple transport streams together must contribute as little jitter as possible. Multiplexing (or transport stream grooming) vendors strive to add less than 70 nS of additional jitter when merging or manipulating the speeds of a transport stream. It is not acceptable for one device to consume the entire +/- 500 nS of overall jitter. A similar issue related to timing is latency. Latency is not an issue as long as all elements (video, audio, timing, and table information) experience the same latency in the stream; no harmful affects to the stream result.

A commonly misunderstood issue is that most devices supporting transport streams have no means to recalculate PCR timing for the transport stream. When a transport stream containing a large amount of jitter is received, the jitter is not corrected. To reconstitute the jitter, e.g., correct an incoming stream with jitter, the stream would have to be completely broken down and then completely reconstructed. This is not typically done in the industry. The old adage: "garbage in, garbage out" applies to transport streams.

Delivering Transport Streams over Satellite

To carry a transport stream over a satellite system, one of the following interfaces must be present on the satellite modulator:

- ASI
- SPI
- CoP#3 (Ethernet)

ASI and SPI both provide a synchronous interface acting as a point-to-point connection from an audio/video encoder or multiplexer; and the CoP#3 interface acts as an Ethernet interface between the audio/video encoder or multiplexer for transmission. The ASI and SPI interfaces each provide a near, jitter-free connection, where the CoP#3 interface has the potential to add a small amount of jitter due to the non-deterministic nature of the Ethernet link. For this reason, CoP#3 interface (Ethernet segments) should be dedicated, and always be a Gigabit (1000 BaseT) connection.

The ASI and SPI interfaces each provide a point-to-point interface for the raw MPEG-2 transport stream, where all frames are continuous, synchronous, and arrive at fixed rate. For these interfaces, every MPEG frame is 188-bytes in length, as shown in Figure 7. There is a special configuration where the MPEG-2 transport stream may be configured for 204 bytes, where 16 additional bytes (padding) are added to the end of the MPEG-2 frame to account for the overhead for supporting forward error correction (FEC) for a modulator; but this is seldom done, because the modulator accounts for this overhead in its rate calculations.



Figure 7 MPEG-2 Synchronous Stream

For DVB-S operation, as defined by ETSI EN 300 421, the FEC inner coding is based on the Viterbi convolution coding and the outer coding is based on the Reed Solomon (RS) interleaving and framing method. The RS interleaving requires an additional 16 bytes of overhead that is post-pended to the end of the 188-byte MPEG frame. The MPEG-2 transport stream, outside of the modulator and demodulator, do not need to account for this overhead, so the resulting input and output MPEG-2 transport stream is simply configured and operates at a desired bit rate. The same is true for a DVB-S2 Constant Coding and Modulation (CCM) stream – 188-byte MPEG frames are presented for delivery over a satellite link.

Due to the non-deterministic nature of Ethernet, an additional level of timing protection is added to the transport layer. The CoP#3 interface operates by queuing up one (1), four (4) or seven (7) MPEG-2 transport stream frames of audio, video, PSI, PCR, and fill frames, and wrapping the frames into a real-time transport protocol (RTP) protected IP packet as shown in Figure 8. The added layer of RTP allows the receiving device (modulator) to receive the IP packet and correct/account for timing delays and ensure minimal jitter is added to the modulated transport stream. The CoP#3 interface also supports an added layer of FEC for the transport stream. It should be noted that both CoP#3 and RTP are optional, not mandatory. FEC is suggested, since there is no retransmit capability for the transport stream on the feed from the audio/video encoder or multiplexer to the modulator. Therefore, FEC lowers the probability a packet may be lost before transmission. It is noteworthy to mention the thrust of Pro MPEG's CoP#3 architecture is to focus more on packet loss and less on overall PCR jitter.



Figure 8. CoP#3 IP Packet Structure

Audio/Video and Transport Streams Over IP

Another common misunderstanding is that all audio/video over IP is transport stream based; however, this is not always the case. IP streams may contain audio/video content that is Program Elementary Stream (PES) based without the need of a transport stream. There are many other proprietary configurations for both MPEG-2 and MPEG-4 H.264 encoded content for both standard definition (SD) and high definition (HD) content.

Additionally, IP over Ethernet is capable of supporting audio/video transport streams based content that contains CoP#3-style packets with and without RTP. Transport streams over IP that do not contain RTP have no provision for correcting timing jitter on the transport streams delivered to the distant end, and by their very nature will contain high jitter.

Finally, another common misunderstanding is that Multiprotocol Encapsulated (MPE) and Generic Stream Encapsulation (GSE) packets that transport IP data are handled as standard audio/video transport streams, but this is not the case. There are no provisions for supporting raw IP packets that carry A/V type data – in these situations, all IP data is simply data. It is noteworthy to mention that MPE is based on DSM-CC and is carried by an MPEG-2 transport stream, but GSE has no concept of a transport stream. The raw GSE information is mapped directly into the bit stream of the carrier.

For an MPE stream, the IP data must first be encapsulated by an IP aware device such as a Comtech EF Data CME-5010 or a CMR-8500 IP Encapsulator (IPE). The function of the IPE is to accept IP datagrams and encapsulate the IP datagrams into MPE packets. Once encapsulated into MPE, the MPE packets are then assigned to a PID, and transmitted over a single MPEG-2 transport stream at a defined rate. The IP datagrams may carry Internet, file data, voice over IP (VoIP), MPEG-2 or 4 video

over IP, etc. MPE data is not timing dependent, i.e., there is no concept of a PCR, but inefficient manipulation of the data streams may result in additional jitter being introduced into the IP data.

Why Does Timing Matter?

For real-time transport streams, the PCR becomes the critical timing element. Timing issues result when supporting older set top boxes (STBs) where memory is limited. As memory has become cheaper, newer STBs have more elastic memory and more memory results in more flexibility for maintaining timing requirements. The limited memory in older STBs resulted in tighter requirements for the PCR, because limited memory allows less audio/video buffering, due to the amount of memory that is required to display the video. Newer STBs contain more memory and are able to withstand higher incoming jitter. To establish a limit on jitter, the ISO/IEC 13818-1 specification allows up to +/- 500 nS of end-to-end jitter – from origination point to reception decoding point. However, for a pure IP delivery system, the jitter requirements are more forgiving. IP systems are inherently jittery due to the non-deterministic Ethernet delivery mechanism, so the maximum IP jitter has been established at 120 mS with a goal of running significantly below this number.

Supporting an MPEG-2 transport stream over a satellite link requires synchronous transmission. For DVB-S and DVB-S2 Constant Coding and Modulation (CCM), the transport is considered to be near jitter-free, so no special conditions need to be applied to the transport stream. For DVB-S2 Variable Coding and Modulation (VCM) and Adaptive Coding and Modulation (ACM), additional timing conditions must be applied.

Conclusion

MPEG-2 transport streams are common in the broadcast communications industry and will continue to be utilized for many years to come. The ability to reliably deliver transport streams over satellite will continue to be a growing market for quite some time as well. Even though the electrical/physical layer has watched the SPI interface come and go, ASI continues to be a popular interface. However, ASI has shown signs of being replaced by A/V over IP.

With the migration to IP, all of the issues that have been discussed in this paper will continue to apply and will increase in significance due to the non-deterministic nature of an IP delivery scheme based on Ethernet.

MPEG-2 transport streams are well understood, but care must be exercised when re-rating or combining the streams for a common multiplex. Comtech EF Data's products adhere to the stringent timing requirements to ensure additional jitter continues to be kept to an absolute minimum.

Please contact Comtech EF Data Sales for more information about this innovative technology.

Voice: 480.333.2200

Fax: 480.333.2540

Email: sales@comtechefdata.com

On the Web: <http://www.comtechefdata.com>