SMPTE STANDARD

Unidirectional Transport of Non-Piecewise Constant Variable Bit Rate MPEG-2 Streams on IP Networks

Table of Contents

Approved
June 8, 2011

Foreword

SMPTE (the Society of Motion Picture and Television Engineers) is an internationally-recognized standards developing organization. Headquartered and incorporated in the United States of America, SMPTE has members in over 80 countries on six continents. SMPTE's Engineering Documents, including Standards, Recommended Practices, and Engineering Guidelines, are prepared by SMPTE's Technology Committees. Participation in these Committees is open to all with a bona fide interest in their work. SMPTE cooperates closely with other standards-developing organizations, including ISO, IEC and ITU.

SMPTE Engineering Documents are drafted in accordance with the rules given in Part XIII of its Administrative Practices. This SMPTE Engineering Document was prepared by Technology Committee 32NF.

Intellectual Property

At the time of publication no notice had been received by SMPTE claiming patent rights essential to the implementation of this Standard. However, attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. SMPTE shall not be held responsible for identifying any or all such patent rights.

Introduction

IP-based networks have become increasingly important for delivery of compressed content contained within MPEG-2 Transport Streams. However, existing transport protocols do not fully meet all user requirements. This standard describes modifications to existing transport protocols which can be used for the unidirectional carriage of non-piecewise constant variable bit rate MPEG-2 Transport Streams over IP networks.

This standard is intended for real-time audio/video transport applications used for contribution and distribution services between professional broadcast equipment over an IP network. The applications addressed by this standard may employ any compression scheme that is supported by the MPEG-2 Transport Stream.

This standard also provides for signal recovery from limited network errors through a forward error correction scheme.

1 Scope

This standard defines a transport protocol for the carriage of real-time non-piecewise constant variable bit rate (VBR) MPEG-2 Transport Streams over IP networks, either with or without Forward Error Correction for recovery from network transmission errors. A non-piecewise constant VBR transport stream has no predictable time base to reconstruct the signal if the inter-packet timing is altered in transit through the network. This standard defines two methods to maintain the inter-packet timing through the non-synchronous IP network transport.

This standard covers the encapsulation and transmission of MPEG-2 transport streams but does not cover other processes such as MPEG-2 encoding or multiplexing. The scope of this standard is shown in Figure 1.

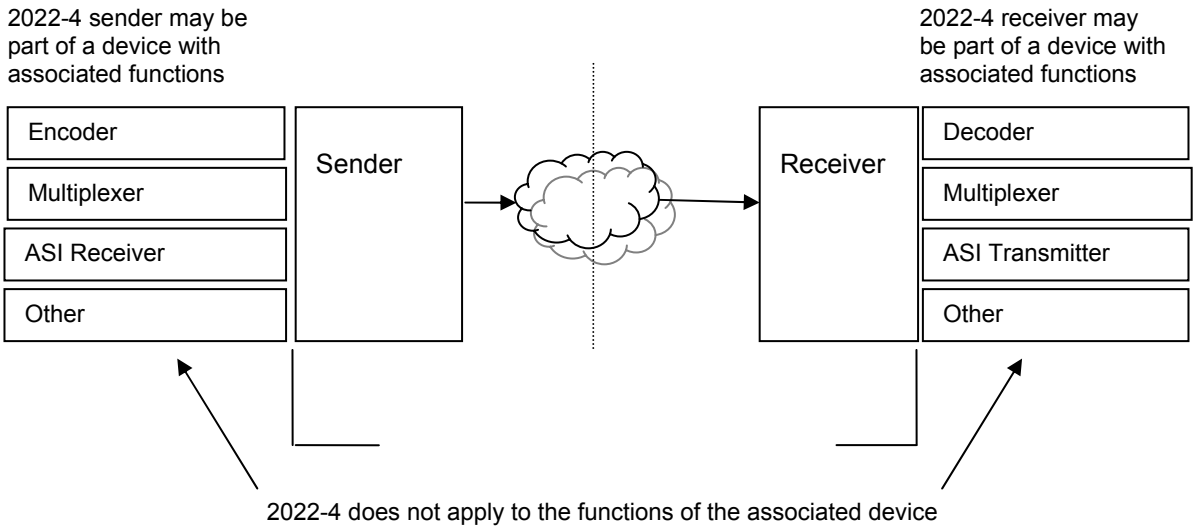

Figure 1 – SMPTE ST 2022-4 Scope Boundaries

2 Conformance Notation

Normative text is text that describes elements of the design that are indispensable or contains the conformance language keywords: "shall", "should", or "may". Informative text is text that is potentially helpful to the user, but not indispensable, and can be removed, changed, or added editorially without affecting interoperability. Informative text does not contain any conformance keywords.

All text in this document is, by default, normative, except: the Introduction, any section explicitly labeled as "Informative" or individual paragraphs that start with "Note:"

The keywords "shall" and "shall not" indicate requirements strictly to be followed in order to conform to the document and from which no deviation is permitted.

The keywords, "should" and "should not" indicate that, among several possibilities, one is recommended as particularly suitable, without mentioning or excluding others; or that a certain course of action is preferred but not necessarily required; or that (in the negative form) a certain possibility or course of action is deprecated but not prohibited.

The keywords "may" and "need not" indicate courses of action permissible within the limits of the document.

The keyword “reserved” indicates a provision that is not defined at this time, shall not be used, and may be defined in the future. The keyword “forbidden” indicates “reserved” and in addition indicates that the provision will never be defined in the future.

A conformant implementation according to this document is one that includes all mandatory provisions ("shall") and, if implemented, all recommended provisions ("should") as described. A conformant implementation need not implement optional provisions ("may") and need not implement them as described.

Unless otherwise specified, the order of precedence of the types of normative information in this document shall be as follows: Normative prose shall be the authoritative definition; Tables shall be next; followed by formal languages; then figures; and then any other language forms.

3 Normative References

Note: All references in this document to other SMPTE documents use the current numbering style (e.g. SMPTE ST 2022-1:2007) although, during a transitional phase, the document as published (printed or PDF) may bear an older designation (such as SMPTE 2022-1-2007). Documents with the same root number (e.g. 2022-1) and publication year (e.g. 2007) are functionally identical.

The following standards contain provisions which, through reference in this text, constitute provisions of this recommended practice. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this standard are encouraged to investigate the possibility of applying the most recent edition of the standards indicated below.

IETF RFC 3550 – RTP: A Transport Protocol for Real Time Applications

ISO/IEC 13818-1:2007 – Generic coding of moving pictures and associated audio information: Systems

SMPTE ST 2022-1:2007, Forward Error Correction for Real-time Video/Audio Transport Over IP Networks

SMPTE ST 2022-2:2007, Unidirectional Transport of Constant Bit Rate MPEG-2 Transport Streams on IP Networks

SMPTE ST 2022-3:2010, Unidirectional Transport of Variable Bit Rate MPEG-2 Transport Streams on IP Networks

4 Acronyms (Informative)

CBR: Constant Bit Rate

FEC: Forward Error Correction

IP: Internet Protocol

PCR: Program Clock Reference

RTP: Real Time Protocol

TS: Transport Stream

VBR: Variable Bit Rate

5 Definitions (Normative)

CBR Transport Stream: An MPEG-2 compliant Transport Stream which is constructed such that the rate of departure of packets from a sender is constant over time.

Contribution Services: Unidirectional transmission of high quality media content to a media processing facility. These services require high quality transmission such that the signal maintains sufficient quality to support further processing prior to final distribution.

Distribution Services: Unidirectional transmission of media content from a media processing facility. These services require robust transmission such that the signal is reliably delivered from one originator to potentially many destinations.

Media Datagram: An RTP Datagram consisting of a header and data payload composed of an integer number of MPEG-2 TS packets

Null Packet: An MPEG-2 Transport Stream packet consisting of a PID value of 0x1FFF and an undefined payload.

Piecewise Constant: A VBR transport stream that can change rate at packets containing PCRs such that the rate of departure of packets from a sender is constant between successive PCR packets, according to ISO/IEC 13818-1, Section 2.4.2.2.

RTP Datagram: An RTP Packet as defined in RFC 3550.

TS Packet: An MPEG-2 Transport Stream packet of 188 or 204 bytes in length as defined in ISO/IEC 13818-1.

VBR Transport Stream: An MPEG-2 compliant Transport Stream such that the rate of departure of packets from a sender is not constant.

6 Transmission Protocols (Normative)

6.1 TS packets per Media Datagram

The number of Transport Stream packets per datagram shall not vary throughout the session. This number of Transport Stream packets per Media datagram is the Packet_per_Datagram_max. The sender shall always include only the Packet _per_Datagram_max MPEG-2 Transport Stream packets per Media datagram. Senders must support the use of 1, 4 and 7 TS packets as Packet_per_Datagram_max. Senders and receivers may use any number of 1 -7 Transport Stream packets as the Packet_per_Datagram_max. Whatever number is chosen must remain constant for the duration of a session as defined in Section 6.7 below. This standard is defined for use with only Mode 1 as defined in SMPTE ST 2022-3.

Note: Long-length Media datagrams can cause longer latency in building FEC Matrices. A larger number of short Media datagrams results in more network encapsulation overhead, and a subsequent higher bit rate. Therefore, the value chosen for Packet_per_Datagram_max will be a compromise between these factors.

FEC datagrams shall always contain Packet_per_Datagram_max transport packets per RTP datagram for a session.

6.2 Signal Time Base Protection and Recovery

Note: An MPEG2 Transport Stream is deemed non piecewise constant when the bit rate varies between PCRs. In such conditions, the time base is depend ent on the inter-packet timing of each TS packet. Sin ce IP networks are non-synchronous the timing relationship between packets will be lost in transit. In this case, the inter-packet timing ought to be signaled to the receiver in the data stream.

To provide the inter-packet timing data, the payload of each datagram shall contain a number of 32-bit data fields equal to the Pa cket_per_Datagram_Max defined for a session as shown below in Figure 2. An additional 32-bit data field, called the Payload Extension Field, shall also be provided to define the number of additional inter-packet timing data fields (Packet_per_Datagram_Max) and the type of inter-packet timing data contained in these fields. The inter-packet timing data fields shall be called Packet Timing Fields.

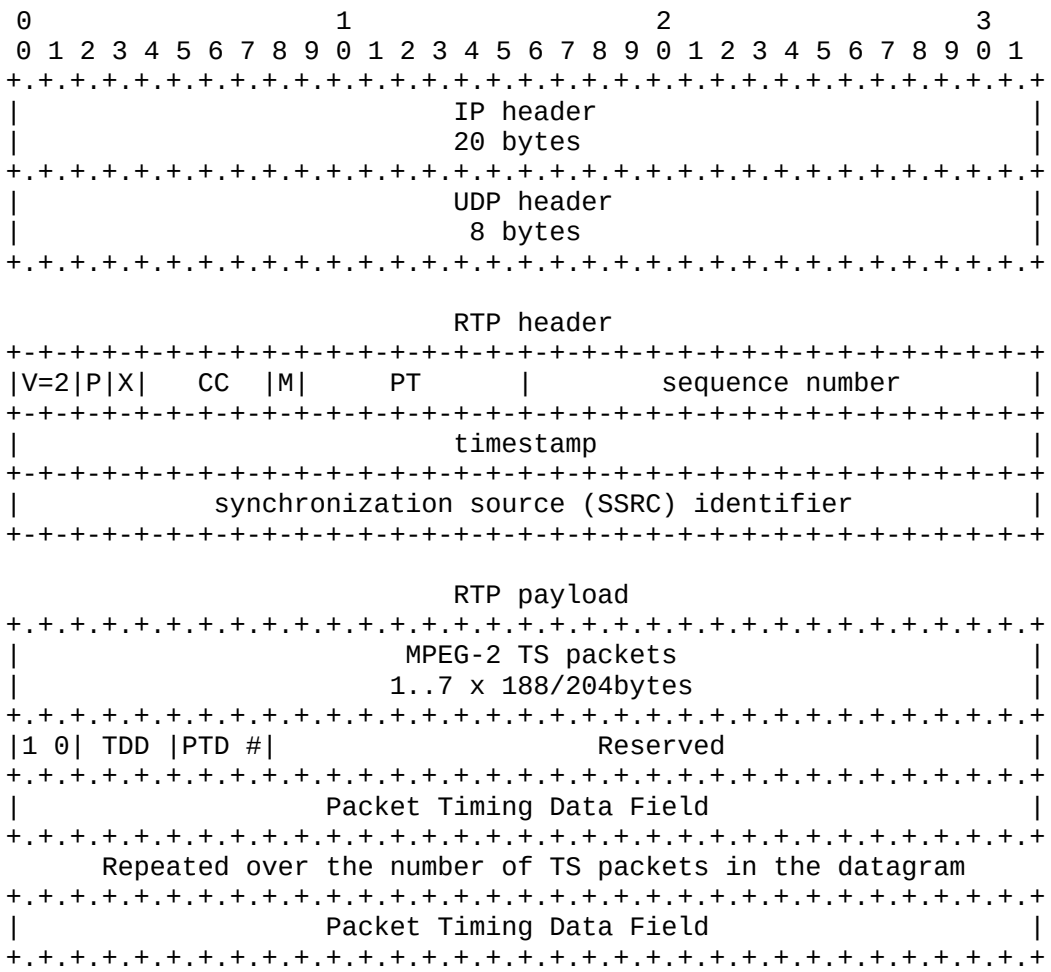

Figure 2 – SMPTE ST 2022-4 Data Mapping Format for the Packet Timing Field

10: Bit positions zero and one of each Packet Timing Field are set to “1” and “0” respectively. This is used to signal the end of TS packets, and that the remaining payload is packet timing data.

TDD: Timing Data Descriptor. This is a 3-bit field u sed to describe the conte nt of the Packet Timing Data Field. The following Table describes the Timing Data Descriptor values for each Packet Timing Data type.

000 – Reserved
001 – Running TS Packet Counter
010 – 27MHz Clock Time Stamp
011 – Reserved
100 – Reserved
101 – Reserved
110 – Reserved
111 – Reserved

PDT #: Packet Timing Data Field Number. This 3-bit field defines the number of Packet Timing Data Fields sent in the datagram.

Packet Timing Data Field: A 32-bit field that is used to carry either a 27-MHz TS time stamp or a running TS packet counter.

This standard defines two methods of timing recovery signaling. Both methods shall be supported in both sender and receiver implementations.

The sender shall provide a 32-bit TS packet counter to support the running TS packet counter timing method. The sender shall provide a 32-bit 27-MHz clock with accuracy of 10 ppm with a maximum drift rate of 22 ppb/sec to support the time stamp timing method.

Note: Of the two timing recovery methods defined in this standard, one method could be optimum for CBR input Transport Streams (that exhibit a significant number of null packets); and the other could be optimum for VBR input Transport Streams (for which specific inter-packet timing information is unknown). However, this standard does not make that distinction. It is left to the sender to determine which method is selected for a specific input signal. This can be either a manual setting or an input signal detection mechanism. It is outside of the scope of this document to define that process.

When a session is started, the method of inter-packet timing recovery shall be defined by the setting of the Timing Data Descriptor field. The receiver shall determine the method of inter-packet timing recovery from the value of the Timing Data Descriptor setting. This value shall not change for the duration of a session.

The Sender operation is to be as follows:

When the Timing Data Descriptor field is set to 001, a 32-bit running TS packet counter is started. When a full media TS packet is detected, it is placed in the first available packet location of the datagram and the value of the running counter is placed in the first Packet Timing Data field. When the next full media TS packet is detected, it is placed in the next available packet location of the datagram and the value of the running counter is placed in the next Packet Timing Data field. If any null TS packets are detected at the sender input between these two full packets, they shall be discarded but counted in the running counter, such that the counter value listed for the second full media TS packet indicates how many null TS packets occurred during that time period. This process is repeated continually throughout the session. The receiver shall use this running counter to define both the time base of the original CBR stream plus the number and location of the null TS packets, such that they are placed back in the output as a CBR stream. This mode of operation will not support a VBR input stream.

When the Timing Data Descriptor is set to 010, the data in the Packet Timing Field will contain the value of the 27-MHz clock for each TS packet in the media datagram. When the first TS packet is detected, it is placed in the first available packet location of the datagram and the value of the 27-MHz clock at that time is placed in the first Packet Timing Data field. The value of the 27-MHz clock shall be that instantaneous value when the sender detects the TS packet sync byte for that packet. When the next TS packet is detected, it is placed

in the next available packet location of the datagram and the value of the 27-MHz clock at that time is placed in the next Packet Timing Data field. If any null TS packets are detected at the sender input, they may be treated as full TS packets, time stamped and sent. This process is repeated continually throughout the session. The receiver shall use these time stamps to define the exact time base of the original VBR stream such that each TS packet is placed back in its original timing location at the output as a VBR stream with the original inter-packet timing.

6.3 FEC Operation

Use of FEC is optional, but when FEC is used, it shall be as defined in SMPTE ST 2022-3 for Mode 1 operation.

Note: Use of FEC is optional.

6.4 Media Datagram RTP Numbering

Media Datagram number shall be as described in SMPTE ST 2022-3.

6.5 Class of Service

Use of 188 or 204 byte TS packets shall be as described in SMPTE ST 2022-2.

6.6 Session

A session is defined as a continuing data stream while the parameters defined in SMPTE ST 2022-3 remain constant.

7 Compliance

An SMPTE ST 2022-4 sender and receiver shall comply with all parameters defined in SMPTE ST 2022-3.

Annex A Non-Piecewise Constant Data Streams (Informative)

There are at least two cases where a signal can become non piecewise constant. In the first case a user with a piecewise constant signal which contains many null TS packets might wish to have the null packets removed for transit. In this case the input to the sender is CBR, but the signal sent through the network to the receiver is VBR once the nulls are removed. The second case is when the input to the sender is VBR with a non standard bit rate to PCR relation ship, such as the original signal with the null packets removed prior to entry at the sender.

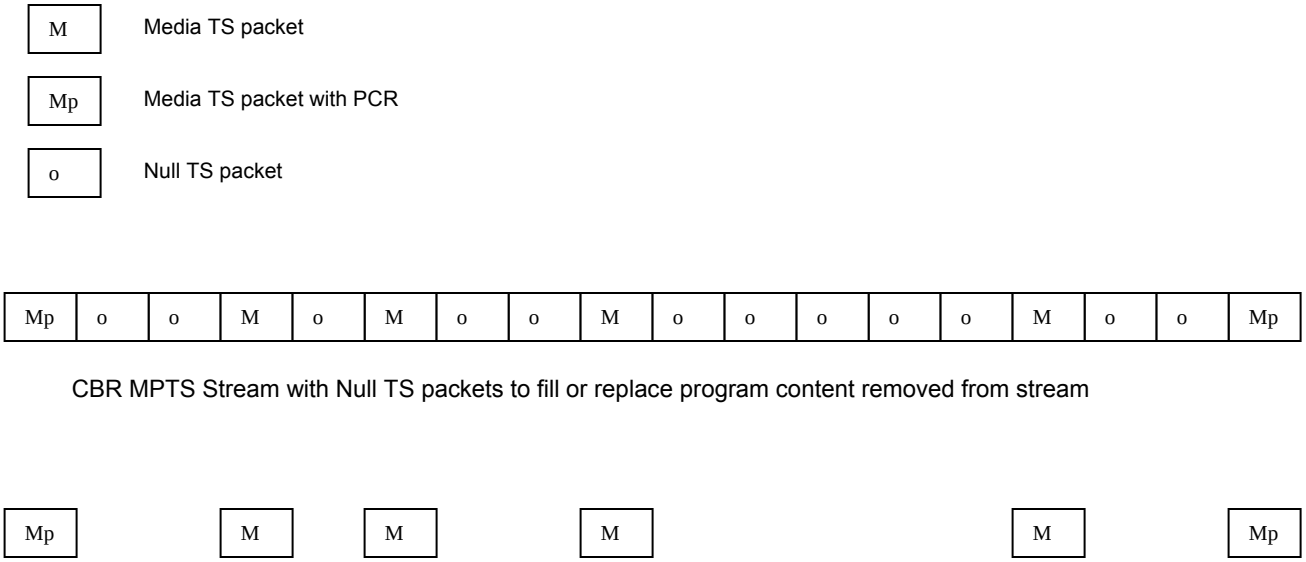

Figure A.1 – Non-Piecewise Constant MEPG Transport Stream