SMPTE STANDARD

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

Table of Contents

Page

Foreword 2
Introduction 2
1 Scope 3
2 Conformance Notation 3
3 Normative References 3
4 Acronyms (Informative) 4
5 Definition (Normative) 4
6 Transmission Protocols 4
6.1 Transmitter Configuration 4
6.2 RTP/UDP/IP Mapping 4
6.3 TS Packet per IP Packet 5
6.4 TS Packet Length 5
6.5 MPEG-2 Timing (Informative) 5
6.6 Timing Recovery 5
7 FEC Buffer Overhead and Latency Implications 6
7.1 FEC Buffer Overhead and Latency Implications 6
7.2 Latency Calculations 6
8 System Configuration 7
Annex A Bibliography (Informative) 8
Annex B Jitter, Latency, Reorder Tolerance and Encryption (Informative) 9
B.1 Scope of Performance 9
B.2 Jitter Tolerance 9
B.3 Reorder Tolerance 9
B.4 Encryption 9

595 W. Hartsdale Ave., White Plains, NY 10607
(914) 761-1100

Approved
May 24, 2007

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.

SMPTE 2022-2 was prepared by Technology Committee N26 on File Management and Networking Technology.

Introduction

This section is entirely informative and does not form an integral part of this document.

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

This standard is intended for real-time audio/video applications such as contribution, distribution, and film. The applications addressed by this standard may employ any compression scheme that is supported by the CBR MPEG-2 Transport Stream. This standard defines two classes of devices. Class 1 supports 188 byte Transport Stream packets, class 2 supports 188 byte and 204 byte Transport Stream packets.

1 Scope

This standard defines a unidirectional transport protocol for the carriage of real-time Constant Bitrate (CBR) MPEG-2 Transport Streams over IP networks. For professional applications, MPEG-2 using the 4:2:2P@ML profile is currently the normal practice. However, Transport Streams containing other forms of MPEG-2 and newer MPEG standards encapsulated as MPEG-2 Transport Streams are also supported by this Standard.

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.

3 Normative References

The following standards contain provisions which, through reference in this text, constitute provisions of this standard. 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.

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

ISO/IEC 13818-1:2000, Generic Coding of Moving Pictures and Associated Audio Information: Systems

IETF RFC 2250, RTP Payload Format for MPEG1 / MPEG2 Video User Performance Requirements

IETF RFC 2236, Internet Group Management Protocol, Version 2

IETF RFC 3376, Internet Group Management Protocol, Version 3

4 Acronyms (Informative)

CBR: Constant Bitrate

CSRC: Contributing Sources List

FEC: Forward Error Correction

IGMP: Internet Group Management Protocol

IP: Internet Protocol

IPDV: IP Delay Variation

IPLR: IP Loss Ratio

IPTD: IP Total Delay

MPTS: Multi-Program Transport Stream

MTU: Maximum Transmission Unit

PCR: Program Clock Reference

RTCP: Real Time Control Protocol

RTP: Real Time Protocol

SSRC: Synchronization Sources List

TOS: Type Of Service

UDP: User Datagram Protocol

XOR: Exclusive OR

5 Definition (Normative)

CBR Transport Stream: A Constant Bitrate ("CBR") Transport Stream as used in this document shall mean a MPEG-2 compliant Transport Stream such that the rate of departure of packets from a hypothetical transmitter is constant with time. In the case of Ethernet-style networks, a hypothetical transmitter shall never experience a packet collision and all packets will be drained by the network at the rate sent by the transmitter.

6 Transmission and Protocols

6.1 Transmitter Configuration

The size of the output IP packet from a transmitting device shall be limited so that IP fragmentation does not occur at the output of the device. The IP 'don't fragment' bit shall be set to '1'. As end-point devices will typically be connected to Ethernet style networks, this limits the maximum transmission unit (MTU) to 1500 bytes.

6.2 RTP/UDP/IP Mapping

The use of RTP shall be required. The RFC2250 mapping shall be used as it provides a suitable mapping for MPEG-2 Transport Streams. Issues for the carriage of 204 byte packets are considered later in this document.

The following additional restrictions on RFC3550 and RFC2250 shall be adopted:

The Padding (P) bit shall be set to zero. This means there will be no padding bytes in the payload.

The Extension (X) bit shall be set to zero. This means there will be no header extension(s) present.

The Marker (M) bit shall always be set to zero. This means there are no discontinuities in the stream during a session.

6.3 TS Packet per IP Packet

The range of possible MPEG Transport Stream (TS) packets per IP packet is from 1 to 7. Long-length packets are undesirable due to the excessive impact (lost data) from losing each IP packet. Short packets cause a high overhead, so the value chosen will be a compromise between these factors. For simplicity, the value chosen shall be kept constant for the duration of a send-receive session.

Senders and receivers shall use 1, 4 and 7 Transport Stream packets.

6.4 TS Packet Length

This standard defines two TS packet length operating points. At the first operating point, the TS packet length shall be 188 bytes. At the second operating point, the TS packet length shall be 204 bytes. There are two classes of devices. The first class shall be identified as Class 1, and shall only support the first operating point (188 byte TS packets). The second class shall be identified as Class 2, and shall support both operating points (188 byte and 204 byte TS packets).

The TS packet length shall be kept constant for the duration of a send-receive session.

NOTE – In more complex network designs the support of the transparent carriage of 204 byte TS packets might be required for end to end integrity checking of the whole network. Currently RFC2250 does not explicitly mention 204 byte packets; so many existing implementations only support 188 byte packets. A Class 2 receiver that can support both 188 and 204 byte TS packets can use the received IP packet length to determine whether 188 or 204 byte packets are present. If 188 byte packets are present then the RTP payload length divides exactly by 188 and not by 204, and vice versa for 204 byte packets.

6.5 **MPEG-2 Timing** (Informative)

Systems based on MPEG-2 Transport Streams have timing recovery information present in the stream (PCR information). This only provides precise timing information for some Transport Stream packets, which means that in the IP domain not every IP packet will contain a timestamp. RFC2250 has a timing recovery mechanism, though the clock for this only has a 90kHz resolution. Under certain limitations, this may be sufficient to allow clock recovery of CBR streams.

RFC2250 requires that the RTP clock be derived from the PCR clock, but this is not a realistic requirement for systems handling multi-program Transport Streams (MPTS), where there may be more than one PCR present, and the PCRs present can change over time. For CBR streams, it is not a requirement that sending systems lock their RTP timestamps to any PCR.

6.6 Timing Recovery

Receiving systems shall not assume that the RTP timestamp will be locked to a PCR.

PCR shall not be modified (same value and same position) by the source device.

Null packets present in the stream shall be kept by the source device.

To comply with the current specification, source devices shall at least implement the compatibility mode defined above. Manufacturers may also implement additional modes on their source devices, but such additional modes shall not jeopardize interoperability.

NOTE – There are several ways to achieve clock recovery in receiving equipment. The method chosen usually depends on the environment and the foreseen link performance (see informative annex B.2 Jitter Tolerance). Receiver clock recovery is out of the scope of this document.

7 FEC Buffer Overhead and Latency Implications

7.1 FEC Buffer Overhead and Latency Implications

As a minimum, senders and receivers shall support all combinations of values of L and D that comply with the limits below:

$$L \times D \leq 100$$
$$1 \leq L \leq 20$$
$$4 \leq D \leq 20$$

These limitations apply to both FEC streams. A device shall only support two FEC streams in the case where $L \geq 4$.

7.2 Latency Calculations (Informative)

Table 1 summarizes the trade-off for different values of L and D between the overhead, the latency implied (for the case of 7 TS packets per IP packet) and the recovery capacity.

Table 1 – Overhead and Latency (Informative)

	Overhead	Latency			Recovery	Buffer size
		3Mbps	30 Mbps	100 Mbps		
XOR (5,10)	10%	175.5 ms	17.5 ms	5.3 ms	5 IP packets	66400 Bytes
XOR (10,10)	10%	350.9 ms	35.1 ms	10.5 ms	10 IP packets	132800 Bytes
XOR (20,5)	20%	350.9 ms	35.1 ms	10.5 ms	20 IP packets	132800 Bytes
XOR (8,8)	12.5%	224.6 ms	22.5 ms	6.7 ms	8 IP packets	84992 Bytes
XOR (10,5)	20%	175.5 ms	17.5 ms	5.3 ms	10 IP packets	66400 Bytes
XOR (8,5)	20%	140.4 ms	14.0 ms	4.2 ms	8 IP packets	53120 Bytes
XOR (5,5)	20%	87.7 ms	8.8 ms	2.7 ms	5 IP packets	33200 Bytes
XOR (4,6)	16.7%	84.2 ms	8.4 ms	2.5 ms	4 IP packets	31872 Bytes
XOR (6,4)	25%	84.2 ms	8.4 ms	2.5 ms	6 IP packets	31872 Bytes

8 System Configuration

Terminal equipment should allow the IP header TOS byte to be fully configured. This will allow both traditional TOS values, and DiffServ markings.

Senders and receivers shall be capable of configuring the destination port number and IP address (which may be a multicast group).

Senders and receivers shall support IGMP version 2. Senders and receivers may support IGMP version 3.

NOTE – IGMP version 3 is backward compatible with IGMP version 2.

Annex A (Informative)
Bibliography

IETF Standard 5, Internet Protocol

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

Annex B (Informative)
Jitter, Latency, Reorder Tolerance and Encryption

B.1 Scope of Performance

Applications targeted by the mechanisms described in this document are professional transport applications in the context of contribution (point to point) and distribution (point to multipoint). These classes of applications have to provide defined service quality. This can only be achieved with knowledge of the network bearer service quality. A reference is the ITU-T Rec. Y.1541 (02/2006). This document defines a reference network model and a number of criteria (IPTD,IPDV,IPLR) with bounded network performance objectives for different classes of network Quality of Service. Network QoS classes 0 & 1 defined there are the target of the present document.

B.2 Jitter Tolerance

Network jitter can be absorbed by buffering at the receiver. There are two components in a typical IP network jitter issue. There is a first high-frequency component, caused by load spikes in the network. These tend to be quite small in value, of the order of ± 10-15ms. On networks carrying Internet or other data traffic there is a 'wander/drift' component, as the loading of the network varies over a 24-hour period. This will typically be larger, of at least ±30-40ms. For the benefit of simplicity, these can be treated as one by providing a 'jitter budget' buffer of 120ms. This buffer should be run half full on average, providing a 60ms latency. For flexibility, it is recommended that system designs make it possible to modify the size of this buffer, as networks can have either significantly better or worse jitter performance.

The jitter absorption needs to be handled carefully, to ensure that the re-generated MPEG stream is still legal in terms of the PCR accuracy etc.

B.3 Latency

Latency within an IP adaptation unit is bounded by a combination of the jitter tolerance buffer, the FEC system and the clock recovery mechanism (not covered by this recommendation).

There are additional latencies caused by the MPEG encode/decode process, and the IP network transmission.

A number of professional applications have demanding round-trip delay requirements. A round-trip delay of 400ms is widely accepted as the worst that would be acceptable for live interview applications for video and audio. Camera control by remote telemetry has been shown to require as low as 200ms. Buffering requirements as part of the error-correction/protection mechanism make these difficult targets to attain for a system following this Standard.

B.4 Reorder Tolerance

Packets traveling over IP networks are not guaranteed to arrive in the order sent. Sequence numbering is provided by RTP, which should allow this effect to be corrected within the receiving end equipment. Any re-ordering that is present is likely to be of a small order, less than 10 packets out of place.

B.5 Encryption

As the initial consideration is for the carriage of MPEG2 Transport Streams, then it will be possible to use standard MPEG Conditional Access systems before the IP encapsulation step. BISS is the EBU proposed encryption system for use in this area (EBUTech/ITU 3290 Rev2). For users wishing to stream content without using MPEG conditional access, IPsec provides a means of encryption at the IP level.