SMPTE STANDARD

Professional Media Over Managed IP Networks: Traffic Shaping and Delivery Timing for Video

Table of Contents

445 Hamilton Avenue., White Plains, NY 10601
(914) 761-1100

Approved November 22, 2017

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 its Standards Operations Manual. 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 Engineering Document. 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

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

The capability and capacity of IP networking equipment has improved steadily, enabling the use of IP switching and routing technology to transport and switch video, audio, and metadata essence within television facilities. Existing standards such as SMPTE ST 2022-6 have gained some amount of use in this application, but there was a desire in the industry to switch different essence elements separately.

This family of SMPTE engineering documents builds on the work of Video Services Forum (VSF) Technical Recommendations TR03 and TR04, and on AES67, documenting a system for transporting various essence streams over IP networks, capturing the timing relationships between those streams. The system is designed to be extensible to a variety of essence types .

Part 10 covers the system as a whole, the timing model, and common requirements across all essence types. Other documents will cover specific media essence formats. Part 20 documents the transport of uncompressed active video in such systems, using an RTP format based on IETF RFC 4175.

Part 21 (this part) specifies the timing model for senders and receivers of video RTP streams.

1 Scope

This standard specifies a timing model for SMPTE ST 2110-10 video RTP streams as measured leaving the RTP sender, and defines the sender SDP parameters used to signal the timing properties of such streams.

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; then formal languages; then figures; and then any other language forms.

3 Normative References

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

Recommendation ITU-R BT.656-5 Interfaces for digital component video signals in 525-line and 625-line television systems operating at the 4:2:2 level of Recommendation ITU-R BT.601

Recommendation ITU-R BT.709-6 Parameter values for the HDTV standards for production and international programme exchange

Recommendation ITU-R BT.1543-1 1280×720, 16:9 progressively-captured image format for production and international programme exchange in the 60 Hz environment

Recommendation ITU-R BT.1847-1 1280×720, 16:9 progressively-captured image format for production and international programme exchange in the 50 Hz environment

Recommendation ITU-R BT.2020-2 (10/2015) Parameter values for ultra-high definition television systems for production and international programme exchange

SMPTE ST 2059-1:2015 Generation and Alignment of Interface Signals to the SMPTE Epoch

SMPTE ST 2110-10:2017 Professional Media over IP Networks: System Timing and Definitions

4 Terms and Definitions

For the purposes of this document, the terms and definitions of SMPTE ST 2110-10 apply.

5 Textual Conventions

5.1 Mathematical Functions

The following functions are defined for one or more numerical arguments.

Maximum	MAX(a, b) shall return the largest value of a and b.
Integer	INT(a) shall return the largest integer not greater than a.

6 Stream Timing Characteristics

6.1 General

This section specifies two parametric models for the packet delivery timing characteristics of frame-based video RTP streams as they leave the transmission interface of the Sender. A network compatibility model is defined in section 6.6.1, which regulates the burst characteristics of senders in order to promote compatibility with a wide variety of switches with varied buffer sizes. Additionally, a Virtual Receiver Buffer model is defined in section 6.6.2.

In the Virtual Receiver Buffer model, packets are deposited at the actual moment of transmission, and removed on a specific schedule. Two fundamental Packet Read Schedules (PRS) are defined, a Gapped schedule and a Linear schedule. The number of outstanding packets in the Virtual Receiver Buffer cannot exceed a parametric maximum for each sender class.

This Virtual Receiver Buffer model provides important information about the stream characteristics useful in the design of receivers – however a practical receiver needs to also accommodate the network-induced packet jitter and latency which accumulates along the path from sender to receiver. Design of receivers is outside the scope of this standard.

6.2 Virtual Receiver Buffer Packet Read Schedule (PRS) Parameters

The following parameters, illustrated for the case of the gapped model in Figure 1 and the linear model in Figure 2, are used in specifying the read schedule for the Virtual Receiver Buffer model of section 6.6. The model defines a Packet Read Time instant TPR_j for each packet j within the video frame of an RTP stream, relative to the Video Transmission Datum T_{VD}. TPR_j is the time at which the packet j will be removed from the Virtual Receiver Buffer in the model.

T_{FRAME} is the time period between consecutive frames of video at the prevailing frame rate

$N_{PACKETS}$ is the number of packets per frame of video (depends on mapping details)

T_{VD} a time point given by $(N \times T_{FRAME}) + TR_{OFFSET}$, where N is an integer and the time scale has its origin at the SMPTE Epoch as defined in SMPTE ST 2059-1.

TR_{OFFSET} is the difference between the most recent integer multiple of T_{FRAME} and T_{VD}. TR_{OFFSET} shall be a positive number or zero, such that $T_{VD} = (N \times T_{FRAME}) + TR_{OFFSET}$ for each frame

$TRO_{DEFAULT}$ is the model-specific default value for the TR_{OFFSET} parameter.

T_{RS} is the time between removing adjacent packets from the Virtual Receiver Buffer during the frame/field (Time-Read-Spacing). Packet removal shall be modeled as an instantaneous event with zero time duration.

TPR_0 is the time when the first packet of the frame is removed from the Virtual Receiver Buffer. TPR_0 is coincident with T_{VD}. $(TPR_0 = T_{VD})$ (Time-Packet-Read-Zero).

TPR_j is the time when packet j will be removed from the Virtual Receiver Buffer (Time-Packet-Read-j).

Video RTP Senders whose streams utilize a value of TR_{OFFSET} which differs from $TRO_{DEFAULT}$ shall signal the prevailing value of TR_{OFFSET} in the Session Description Protocol (SDP) with a Media Type Parameter `TROFF` of the prevailing TR_{OFFSET} value, in microseconds, expressed as a positive integer decimal value (rounded if necessary). If this parameter is not present, receivers shall assume the default value $TRO_{DEFAULT}$.

6.3 Gapped Packet Read Schedule (PRS)

6.3.1 Overview

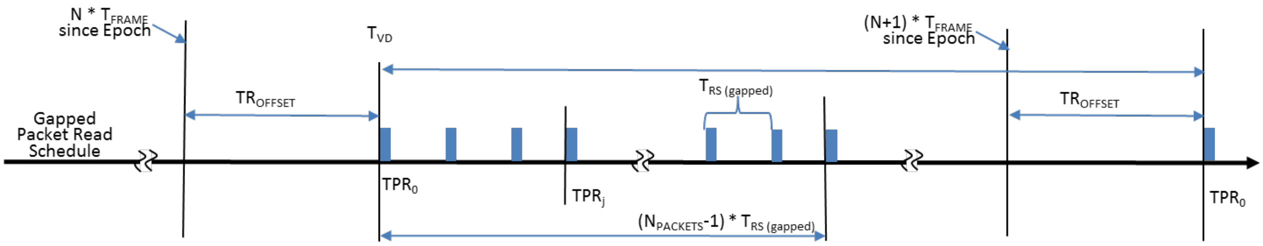

Figure 1 -- Packet Read Schedule (Gapped)

The gapped PRS, illustrated in Figure 1, consists of a sequence of TPR_j values which loosely approximates the delivery of samples in the SDI signal, including a gap in the transmission corresponding to the vertical blanking time of SDI. The TPR_j instants are spaced uniformly throughout the active field (or segment) or frame interval. The progressive scan signal has one gap after each frame, and an interlaced scan signal (and PsF signal) has two gaps, one after each field (or segment in the case of PsF). The following variable is used in the equations below:

R_{ACTIVE} is the ratio of active time to total time within the frame period.

Although RTP puts no specific constraints on the image height, width, or frame rate, the gapped PRS shall only apply to streams with image dimensions and frame rates specified in Recommendation ITU-R BT.656-5, Recommendation ITU-R BT.1543-1, Recommendation ITU-R BT.1847-1, Recommendation ITU-R BT.709-6, or Recommendation ITU-R BT.2020-2.

6.3.2 Gapped PRS – Progressive Images

For progressively scanned images (but excluding Progressive segmented Frame (PsF) images) using the gapped model, TPR_j shall be defined as follows:

$$R_{ACTIVE} = (1080/1125)$$

$$T_{RS} = (T_{FRAME} \times R_{ACTIVE})/N_{PACKETS}$$

$$T_{VD} = (N \times T_{FRAME}) + TR_{OFFSET} \qquad (for\ an\ integer\ value\ of\ N)$$

$$TPR_j = (j \times T_{RS}) + T_{VD}$$

$$TRO_{DEFAULT} = \begin{cases} (43/1125) \times T_{FRAME} & \text{Image Height} \geq 1080 \text{ lines} \\ (28/\ 750\ ) \times T_{FRAME} & \text{Image Height} < 1080 \text{ lines} \end{cases}$$

Note: this gapped model defines the inter-frame gap and TR_{OFFSET} value for all (non-PsF) progressive formats, including 720p, 1080p, 2160p, and 4320p. The $TRO_{DEFAULT}$ values are chosen slightly later than the beginning of the active video in order to allow sufficient time to assemble packets and buffer the samples for transmission when working from a timed SDI signal and packing into the Standard UDP Size Limit.

6.3.3 Gapped PRS – Interlaced and PsF Images

For interlaced images and Progressive segmented Frame (PsF) images using the gapped model, the TPR_j values depend on the interlace standard in use. The values TPR_j for interlaced images shall be defined as follows:

$$R_{ACTIVE} = \text{(see Table 1)}$$

$$TRO_{DEFAULT} = \text{(see Table 1)}$$

$$T_{RS} = (T_{FRAME} \times R_{ACTIVE})/N_{PACKETS}$$

$$T_{VD} = (N \times T_{FRAME}) + TR_{OFFSET} \qquad (for\ an\ integer\ value\ of\ N)$$

$$TPR_j = \begin{cases} (j \times T_{RS}) + T_{VD}, & 0 \le j < (N_{PACKETS}/2) \\ (T_{FRAME}/2) + \left((j - (N_{PACKETS}/2)) \times T_{RS}\right) + T_{VD}, & (N_{PACKETS}/2) \le j < N_{PACKETS} \end{cases}$$

Table 1 – Ratio of active to total time for interlaced systems

System	R_{ACTIVE}	$TRO_{DEFAULT}$
525 line interlaced system as specified in Recommendation ITU-R BT.656-5	$(487/525)$	$(20/525) \times T_{FRAME}$
625 line interlaced system as specified in Recommendation ITU-R BT.656-5	$(576/625)$	$(26/625) \times T_{FRAME}$
1125 line interlaced systems and Progressive segmented Frame (PsF) systems as specified in Recommendation ITU-R BT.709-6	$(1080/1125)$	$(22/1125) \times T_{FRAME}$

Note: The $TRO_{DEFAULT}$ values are chosen slightly later than the start of the active video in order to allow sufficient time to assemble packets and buffer the samples for transmission when working from a timed SDI signal and packing into the Standard UDP Size Limit.

6.4 Linear Packet Read Schedule (PRS)

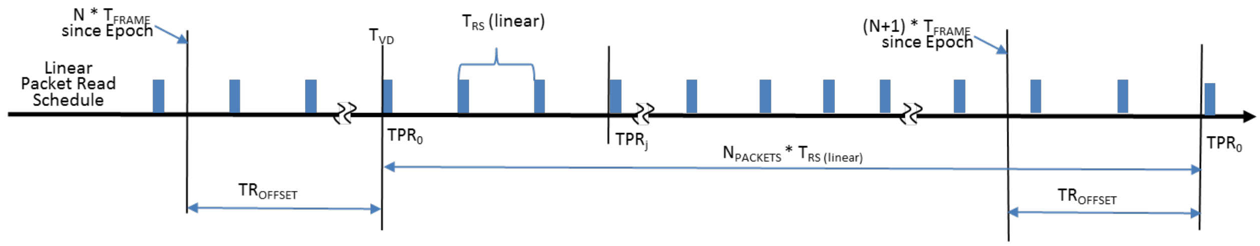

Figure 2 -- Packet Read Schedule (Linear)

The linear PRS illustrated in Figure 2, consists of a sequence of TPR_j values which are evenly spaced throughout the frame period T_{FRAME}.

For all images utilizing the linear PRS model, TPR_j shall be defined as follows:

$$T_{RS} = T_{FRAME}/N_{PACKETS}$$

$$T_{VD} = (N \times T_{FRAME}) + TR_{OFFSET} \qquad (for\ an\ integer\ value\ of\ N)$$

$$TPR_j = (j \times T_{RS}) + T_{VD}$$

$$TRO_{DEFAULT} = TRO_{DEFAULT}\ as\ defined\ in\ the\ gapped\ model\ of\ section\ 6.3$$

6.5 Relationship between Linear and Gapped PRS (Informative)

The Video Transmission Datum, T_{VD}, is common to both models. By using the same value of TR_{OFFSET} in both models, given a sufficient VRXfull value, a receiver reading at the linear PRS can also accommodate a signal meeting the requirements of the gapped PRS. Figure 3 shows the relationship between the linear and gapped PRS for progressive signals.

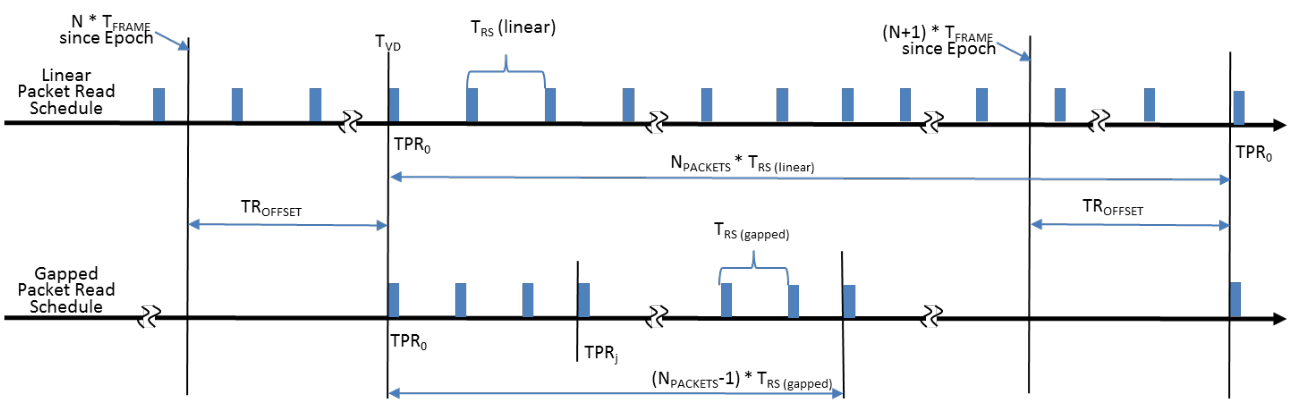

Figure 3 - Relationship between the Linear and Gapped PRS

6.6 Transmission Traffic Shape Models

6.6.1 Network Compatibility Model

Video RTP Senders shall ensure that their sequence of actual transmission instants, as measured on their network egress interface, passes the Network Compatibility Model shown in Figure 4 at all times and in all operating configurations. The model is tested at the output of the Sender, prior to any network-induced delivery impairments.

Packets from the sender shall enter a leaky bucket of infinite capacity at the instant they are emitted from the sender. The bucket drains a packet every T_{DRAIN} seconds, if a packet is available. For the purpose of accurate modeling, the specific instant of the bucket draining the packet is N * T_{DRAIN} seconds since the SMPTE Epoch as defined in SMPTE ST 2059-1. C_{INST} represents the instantaneous number of packets in the bucket at any time. The value of C_{INST} shall never exceed the sender-type-specific C_{MAX} value specified in section 7.1.

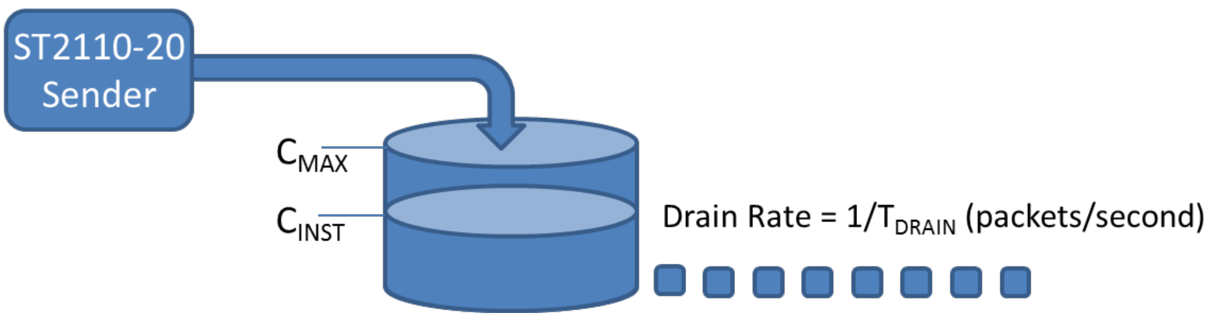

Figure 4 -- Network Compatibility Model

The parameters of the Network Compatibility Model are as follows:

$R_{NOMINAL}$ The long-term average packet rate generated by a sender, in units of packets per second.
$R_{NOMINAL} = N_{PACKETS} / T_{FRAME}$

β A scaling factor applied to $R_{NOMINAL}$, specified in section 7.1

T_{DRAIN} The time interval (in seconds) between packets being drained in the Network Compatibility buffer model.
$T_{DRAIN} = (T_{FRAME} / N_{PACKETS}) * (1 / \beta)$

C_{INST} = The instantaneous fullness (in packets) of the bucket at any time

6.6.2 Virtual Receiver Buffer Model

Video RTP Senders shall ensure that their sequence of actual transmission instants, as measured on their network egress interface, passes the Virtual Receiver Buffer Model shown in Figure 5 at all times and in all operating configurations. The model is tested at the output of the Sender, prior to any network-induced delivery impairments.

Packets from the sender shall enter a leaky bucket of capacity VRX_{FULL} at the instant they are emitted from the sender. For purposes of the model packets shall be assumed to enter and exit the bucket instantaneously. The VRX_{FULL} bucket drains packet j at the Packet Read Schedule instant TPR_j. The sender shall ensure that this bucket does not overflow, and the sender shall ensure that packet j is available in the bucket (emitted onto the network) no later than time instant TPR_j, so that the VRX_{FULL} bucket shall not underflow.

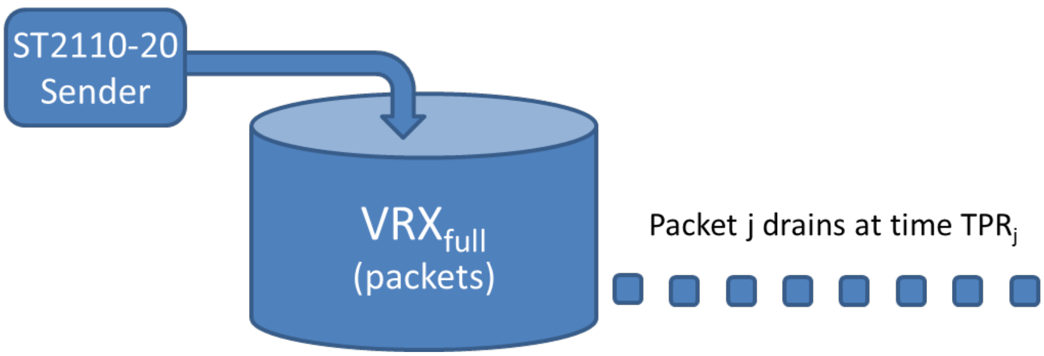

Figure 5 -- Virtual Receiver Buffer Model

The parameters of the VRX_{FULL} buffer are shown below

TPR_j = as defined in section 6.2 for the prevailing PRS model.

VRX_{FULL} = (units of packets) represents the capacity of the Virtual Receiver model's buffer; it constrains the number of outstanding packets between the sender's actual emission times and the Virtual Receiver's drain schedule TPR_j. See section 7 for the VRX_{FULL} values for different types of senders.

7 Compliance Definitions

7.1 Senders

7.1.1 Narrow Senders (Type N)

Type N Senders shall employ the gapped PRS defined in section 6.3, and shall comply with the Virtual Receiver Buffer Model specified in section 6.6.2 with the VRX_{FULL} value specified below:

$$VRX_{FULL} = MAX\left(INT\left(\frac{1500 \times 8}{MAXIP}\right),\ INT\left(\frac{N_{PACKETS}}{27000 \times T_{FRAME}}\right)\right)$$

The value MAXIP above shall be assumed to be 1500 if the Standard UDP Size Limit is in use, and shall be as defined in SMPTE ST 2110-10 when the Extended UDP Size Limit is in use.

Type N senders shall signal compliance with a Media Type Parameter `TP` of value `2110TPN`.

Type N senders shall comply to the Network Compatibility Model specified in section 6.6.1, with the following parameters:

$$C_{MAX} = \mathrm{MAX}\left(4,\ INT\left(\frac{N_{PACKETS}}{43200 \times R_{ACTIVE} \times T_{FRAME}}\right)\right)$$

β = 1.10

Note 1: the VRX_{FULL} specification above provides a minimum 8 packet Virtual Receiver Buffer, which scales up for higher-rate signals. 8 packets is, by example, just over 2 lines of 4:2:2/10 samples in a 1920 sample width format. For larger UDP packet sizes the values are scaled proportionally by the reduction in the value of $N_{PACKETS}$.

Note 2: A sender which packs the video payload of a (locked and phased) SDI signal using the maximum standard-sized packets, and transmitting the packets when they become full, can be compliant to this type.

Note 3: See Annex A for more details regarding the parameters in the C_{MAX} expression above. A minimum value is ensured for the case of small images or low frame rates. In the Extended UDP Packet Size, the number of packets per frame $N_{PACKETS}$ is correspondingly reduced leading to a proportionally smaller calculated C_{MAX} value. For type N senders with data rates such that the minimum C_{MAX} value (4 for type N) is in effect, care needs to be taken over switch buffer and loading if simultaneous bursts on all streams passing through the switch are to be accommodated.

7.1.2 Narrow Linear Senders (Type NL)

Senders of this type (NL) shall employ the linear PRS defined in section 6.4, and shall comply with the Virtual Receiver Buffer Model specified in section 6.6.2 with the VRX_{FULL} value specified below:

$$VRX_{FULL} = MAX\left(INT\left(\frac{1500 \times 8}{MAXIP}\right),\ INT\left(\frac{N_{PACKETS}}{27000 \times T_{FRAME}}\right)\right)$$

The value MAXIP above shall be assumed to be 1500 if the Standard UDP Size Limit is in use, and shall be as defined in SMPTE ST 2110-10 when the Extended UDP Size Limit is in use.

Type NL senders shall signal compliance with a Media Type Parameter `TP` of value `2110TPNL`.

Type N senders shall comply to the Network Compatibility Model specified in section 6.6.1, with the following parameters:

$$C_{MAX} = \mathrm{MAX}\left(4,\ INT\left(\frac{N_{PACKETS}}{43200\ \times T_{FRAME}}\right)\right)$$

β = 1.10

Note: For type NL senders with data rates such that the minimum C_{MAX} value (4 for type NL) is in effect, care needs to be taken over switch buffer and loading if simultaneous bursts on all streams passing through the switch are to be accommodated. Similar care needs to be taken if using the Extended UDP Size Limit. Refer to Annex A for details of the method used to derive the C_{MAX} expressions.

7.1.3 Wide Senders (Type W)

Type W Senders shall employ the linear PRS defined in section 6.4, and shall comply with the Virtual Receiver Buffer Model specified in section 6.6.2 with the VRX_{FULL} value specified below:

$$VRX_{FULL} = MAX\left(INT\left(\frac{1500 \times 720}{MAXIP}\right),\ INT\left(\frac{N_{PACKETS}}{300 \times T_{FRAME}}\right)\right)$$

Type W senders shall signal compliance with a Media Type Parameter `TP` of value `2110TPW`.

Type W senders shall comply to the Network Compatibility Model specified in section 6.6.1, with the following parameters:

$$C_{MAX} = \mathrm{MAX}\left(16,\ INT\left(\frac{N_{PACKETS}}{21600\ \times T_{FRAME}}\right)\right)$$

β = 1.10

The C_{MAX} definition above shall only apply to streams of less than 900000 packets/second.

The value MAXIP above in the VRX_{FULL} definition shall be assumed to be 1500 if the Standard UDP Size Limit is in use, and shall be as defined in SMPTE ST 2110-10 when the Extended UDP Size Limit is in use.

Note 1: The CMAX specification above remains a topic of study, particularly in regards to scaling at higher stream rates.

Note 2: The VRX_{FULL} specification above provides a minimum 720 packet Virtual Receiver Buffer, which scales up proportionally for higher-rate signals. 720 packets is approximately 20 percent of a 1920x1080 4:2:2/10 frame when using

Standard UDP Size packets. Using larger UDP packets proportionally reduces the allowed buffer size in larger formats due to the change in $N_{PACKETS}$.

Note 3: This type W is defined in order to support present and future software-based signal sources which might exhibit a wider degree of packet transmission timing variation than Type N or NL. This larger VRX_{FULL} value accommodates both increased packet delay variation, and also mis-alignment between the sender's packet generation process and its PTP reference.

Note 4: The minimum C_{MAX} value is larger than that allowed for type N or NL. For W senders with data rates such that the minimum C_{MAX} value (16 for type W) is in effect, care needs to be taken over switch buffer and loading if simultaneous bursts on all streams passing through the switch are to be accommodated. Similar care needs to be taken if using the Extended UDP Size Limit. Refer to Annex A for details of the method used to derive the C_{MAX} expressions.

7.2 Receivers

7.2.1 Overview (Informative)

Senders of SMPTE ST 2110-10 RTP streams signal their clock source in the SDP using the “ts-refclk clksrc” parameter as defined in ST 2110-10, and receivers can compare the value of this parameter with their own reference clock source to determine if the sender is synchronized to the same clock source. In addition, senders signal their relative alignment to that clock source using the optional `TROFF` parameter, so that each receiver can assess the suitability of the signal’s phase against its own requirements.

This standard defines three types of receiver capabilities as enumerated in the following sections.

While the sender traffic models of section 6.6 provide some bound on the traffic shape leaving the sender, packet jitter and delay might accumulate during network transit. Practical receiver designs ought to accommodate a reasonable amount of accumulated packet arrival jitter and delay over and above the specification in the traffic profile.

7.2.2 Narrow, Synchronous Receivers (Type N)

Type N receivers shall be capable of receiving signals from a type N sender when the following conditions are satisfied:

a) The receiver is locked to the same clock source as indicated in the sender’s `ts-refclk clksrc` SDP clause.
b) The sender’s `TROFF` parameter is equal to the default, or absent (indicating the default)

A Type N receiver should support alternative values of `TROFF`.

A Type N receiver should support reception of signals from a type NL sender subject to the conditions above.

7.2.3 Wide, Synchronous Receivers (Type W)

Type W receivers shall be capable of receiving signals from a type N, NL, or W sender when the receiver is locked to the same clock source as indicated in the sender’s `ts-refclk clksrc` SDP clause.

A Type W receiver shall support alternative values of TR_{OFFSET} as signaled by the sender using the `TROFF` parameter.

Note: the default value of TR_{OFFSET} is common between the type N and type W senders – see section 6.5.

7.2.4 Asynchronous Receivers (Type A)

Type A receivers shall be capable of receiving signals from a type N, NL, or W sender, regardless of the value of the sender’s `ts-refclk clksrc` or the sender’s `TROFF` parameter value.

8 Session Description Considerations

8.1 Required Parameters

Senders shall include the following additional payload-format-specific Media Type parameters in the a=fmtp clause of the SDP for all video RTP streams conforming to this standard.

`TP` Signals the type of the sender as defined in section 7.1. See section 7.1 for permitted values of this parameter under this standard.

8.2 Optional Parameters

The following payload-format-specific Media Type parameters may be added to the a=fmtp clause of the SDP for video RTP streams conforming to this standard.

`TROFF` If this parameter name is present, it indicates the value of TR_{OFFSET} in use by the sender. If this parameter is not present, the default value $TRO_{DEFAULT}$ shall be assumed. See section 6.2 for details of when this parameter is mandatory. This value is expressed as an integer number of microseconds.

`CMAX` If this parameter name is present, it indicates the largest value of C_{INST} which will be present in the signal as measured at the output of the sender. If this parameter is not present, then the value of C_{MAX} defined for the specific sender class shall be assumed. This value is expressed as an integer number.

Annex A Regarding the C_{MAX} Expression and Network Compatibility Parameters for Senders (Informative)

Network Switches available in the marketplace today contain buffer memory which is used to queue packets waiting to egress the switch on a given interface. Switch architectures vary significantly across the industry, and the amount of memory available to queue packets for a given egress interface is different in different switches. Typically the presence of more memory for packet queuing is a feature associated with the more expensive switches within the Commercial Off The Shelf (COTS) marketplace.

The value of C_{MAX} in the transmission packet model of sections 6.4 and 7.1 scales with the nominal rate of the stream.

We use a model of a simple switch with a certain total buffer which is shared across egress ports according to demand. The model represents optimum usage of buffer and real implementations might do less well.

The following variables are used:

$\boldsymbol{E_{total}}$	**Total egress capacity of the switch**
E_{used}	Used egress capacity of the switch
U	E_{used}/E_{total} (Aggregate Utilization factor)
B_{total}	Total buffer capacity of the switch
BB	Buffer per unit of used bandwidth
R_{stream}	Data rate for a single stream
B_{stream}	Buffer for a single stream

Demand for buffer occurs when the rate of packets arriving for egress at the port exceeds the maximum egress rate of the port.

We assume a worst-case scenario where all ports are loaded to maximum egress utilization and demand for buffer occurs simultaneously at all egress ports.

The buffer per unit of used bandwidth:

$$BB = B_{total}/E_{used}$$

$$BB = (B\_total\ )\ /\ (U \times E\_total\ )$$

As the demand for buffer is the same on all ports, we assume buffer is allocated evenly across all ports and across all streams egressing from the port.

The buffer for a single stream:

$$B_{stream} = R_{stream} \times BB$$

$$B_{stream} = R_{stream} \times B_{total} / (U \times E_{total})$$

$$B_{stream} = \frac{R_{stream}}{(U \times E_{total}) / (B_{total})}$$

Using figures for a typical open-market switch ASIC and a maximum utilization of 90 %:

$\boldsymbol{E_{total}}$	**3.2 Tbit/s**
B_{total}	16 Mbytes
U	0.90 (90% utilization)

The buffer for a single stream can be found as:

$$B_{stream} = \frac{R_{stream}}{(3.2 \times 10^{12}) / (1.1 \times 16 \times 2^{20} \times 8)}$$

$$B_{stream} = \frac{R_{stream}}{21674}$$

To be even more cautious, we modify the value further based on the relative likelihood of synchronicity between bursting events - reducing the value by approximately half (for type N) from that above, and rounding for convenience:

$$B_{Nstream} = \frac{R_{stream}}{43200}$$

$$B_{Wstream} = \frac{R_{stream}}{21600}$$

The C_{MAX} values specified for type N and W senders in section 7.1 are based on the derivation above. These values are based on assumptions of highly loaded networks ($U = 90\%$), with a modest likelihood of coupling of the burst behaviors between streams. The marketplace also contains switches with larger buffer capacity (more buffer per stream) should these assumptions be inadequate for a specific situation.

In a simple switch, the buffer on an egress port drains at the line rate of that port. In the scenario above, each port is loaded to the maximum utilization U, i.e. the long-term average data-rate through the port is equal to the line rate of the port times U. Since the drain rate for the egress buffer is equal to the line rate, by extension on a per-stream basis, the stream's effective buffer drain rate is equal to the stream data rate divided by U. We used the value U = 90% in the model above and this corresponds to a value of $\beta = 1.1$ in the network compatibility specifications above.