Payload Working Group S. Lugan
Internet-Draft G. Rouvroy
Intended status: Standards Track A. Descampe
Expires: November 26, 2018 intoPIX
T. Richter
IIS
A. Willeme
UCL/ICTEAM
May 25, 2018

RTP Payload Format for ISO/IEC 21122 (JPEG XS)
draft-lugan-rtp-jpegxs-00

Abstract

This document specifies a Real-Time Transport Protocol (RTP) payload format to be used for transporting ISO/IEC 21122 (JPEG XS) encoded video. ISO/IEC 21122 (JPEG XS) is a low-latency, lightweight image coding system allowing for an increased resolution and frame rate, while offering visually lossless quality with reduced amount of ressources such as power and bandwidth.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

This document specifies a payload format for packetization of ISO/IEC 21122 (JPEG XS) [ISO21122-1] encoded video signals into the Real-time Transport Protocol (RTP) [RFC3550].

2. Conventions, Definitions, and Abbreviations

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2.1. Application Dependent Unit

See Real-time Transport Protocol (RTP) [RFC3550], though for the purpose of this document identical to a JPEG XS frame.

2.2. JPEG XS codestream

A sequence of bytes representing compressed images formatted according to ISO/IEC 21122-1.

2.3. JPEG XS frame

Concatenation of the Video Essence box and a JPEG XS codestream

2.4. Marker

A two-byte functional sequence that is part of a JPEG XS codestream starting with a 0xff byte and a subsequent byte defining its function.

2.5. Marker Sequence

A marker along with a 16-bit marker size and payload data following the size.

2.6. JPEG XS Header

A sequence of bytes at the beginning of each JPEG XS codestream encoded in multiple markers and marker sequences that does not carry entropy coded data, but metadata such as the frame dimension and component precision.

2.7. Video Essence Box

A ISO super box in the sense of ISO/IEC 15444-1 defined in ISO/IEC 21122-3 that includes metadata required to play back a JPEG XS video stream, such as its color space, its buffer model and its frame rate.

2.8. JPEG XS Header Segment

The concatenation of the Video Essence Box and the JPEG XS Header.

2.9. Slice

The smallest independently decodable unit of a JPEG XS stream.

2.10. Slice group

A contiguous sequence of slices belonging to a fragment.

2.11. Fragment

A slice group along with the metadata immediately preceeding and/or following it sized such that the first byte of the fragment and the byte following the last byte of the fragment are in two distinct packets, except for the last fragment of a ADU.

3. Media Format Description

3.1. Wavelet decomposition

JPEG XS is a low-latency lightweight image coding system for coding continuous-tone grayscale or continuous-tone color digital images.

This coding system provides an efficient representation of image signals through the mathematical tool of wavelet analysis. The wavelet filter process separates each component into multiple bands, where each band consists of multiple coefficients describing the image signal of a given component within a frequency domain specific to the wavelet filter type, i.e. the particular filter corresponding to the band.

Wavelet coefficients are grouped into precincts, where each precinct includes all coefficients over all bands that contribute to a spatial region of the image.

One or multiple Precincts are furthermore combined into slices consisting of an integral number of precincts. Precincts do not cross slice boundaries, and wavelet coefficients in precincts that are part of different slices can be decoded independently from each other. Note, however, that the wavelet transformation runs across site boundaries. A slice always extends over the full width of the image, but may only cover parts of its height.

A slice the the smallest indepedently decodable unit of a JPEG XS codestream, bearing in mind that it decodes to wavelet coefficients which still require inverse wavelet filtering to give an image.

3.2. Codestream

The codestream is a linear stream of bits from the first bit to the last bit representing the sample values of a single frame, bare any interpretation relative to a colorspace. It can be divided into (8-bit) bytes, starting with the first bit of the codestream. Bits within bytes are enumerated from the least significant bit (LSB) to the most significant bit (MSB), with the least significant bit having the index zero. Bits within bytes are transmitted in decreasing magnitude order, with the MSB of a byte transmitted first and the LSB transmitted last. This implies, in particular, that fields that are longer than 8 bits are transmitted with the most significant byte first. This is also denoted as "big endian" format.

The codestream consists of multiple syntax elements: markers, marker segments and entropy coded data.

Markers inidicate syntactical elements of the codestreams. They consist of an 0xff-byte and a second byte defining the nature of the marker. The SOC marker (hex 0xff10) indicates the start of the codestream, the EOC marker (hex 0xff11) its end.

Marker segments are markers along with a length field and payload data following the length field. Marker segments define control information necessary to steer the decoding process. The JPEG XS specification ISO/IEC 21122-1 [ISO21122-1] defines additional markers beyond SOC and EOC.

The sequence of bytes made up by all markers that precede the entropy coded data is also denoted as JPEG XS Header in the following.

Entropy coded data represents the image data itself. The data is organized in slices, where each slice consists of a slice header that starts with the SLC marker (hex 0xff20) and payload data, consisting of encoded wavelet coefficients.

The overall codestream format, including the definition of all markers, is further defined in ISO/IEC 21122-1 [ISO21122-1].

3.3. Video Essence Box

While the information defined in the codestream is sufficient to reconstruct the sample values of one video frame, the interpretation of the samples remains undefined by the codestream itself. This interpretation, including the color space, frame rate and other information significant to play a JPEG XS stream are contained the Video Essence box, which preceeds each JPEG XS codestream. The syntax of the Video Essence box follows ISO/IEC 15444-1 [ISO15444-1]; it consists of multiple subboxes, each with a particular meaning. Its contents, in particular its subboxes are defined in ISO/IEC 21122-3 [ISO21122-3].

3.4. JPEG XS Stream

A JPEG XS stream is a sequence of frames, where each frame is coded independently of each other. For the purpose of RTP transport, each frame forms an Application Dependent Unit (ADU).

A JPEG XS frame consists the concatenation of a Video Essence box (as defined in ISO/IEC 21122-3 [ISO21122-3]) and a JPEG XS codestream (as defined in ISO/IEC 21122-1 [ISO21122-1]). As defined above, the codestream consists of a JPEG XS header, one or multiple slice groups, and an EOC marker.

3.5. Fragments

For the purpose of transport, JPEG XS frames are separated into one or multiple fragments such that the start of the fragment and the byte following the last byte of a fragment are in two distinct packets used for RTP transport, except for the last fragment of a JPEG XS frame which may be contained in only a single packet.

A fragment consists of all metadata preceeding its first slice, one or multiple slices, and potentially the EOC marker following the last slice.

The collection of slices in a fragment is also denoted as slice group, and slice groups within a frame are enumerated from top to bottom by the slice group counter. That is, the first slice group of a frame is slice group #0, and the slice group counter increments by 1 from top to bottom for each slice group, and by that for each fragment.

NOTE: By this definition, the first fragment consists of at least the Video Essence Box, the JPEG XS header, and the first slice group. The last fragment consists of at least the last slice group and the

EOC marker. In case the frame consists of only a single fragment, this fragment contains both the JPEG XS header segment and the EOC.

4. Payload Format

This section specifies the payload format for JPEG XS video streams over the Real-time Transport Protocol (RTP) [RFC3550].

In order to be transported over RTP, each JPEG XS stream is transported in a distinct RTP stream, identified by a distinct SSRC.

Each of those RTP streams is divided into Application Data Units (ADUs). Each ADU shall correspond to a single JPEG XS frame.

Each ADU is split into packets, depending e.g. on the Maximum Transmission unit (MTU) of the network. Every packet shall have same size, except the last packet of every ADU which could be shorter. Packet boundaries shall coincide with ADU boundaries, i.e. the first byte of an ADU shall be the first byte of payload data within a JPEG XS segment.

A JPEG XS frame, and by that each ADU, shall consist of a Video Essence box defining the meta information required for playback, concatenated to the JPEG XS codestream, defining the sample values of the picture.

The JPEG XS stream, as defined in ISO/IEC 21122-1 [ISO21122-1] itself consists of a JPEG XS header that defines picture parameters, and one or multiple slices that contain the entropy coded picture data and an EOC marker. A slice is the smallest independently decodable unit of a JPEG XS codestream.

JPEG XS frames are separated into fragments such that the first byte of a fragment and the byte following the last byte of a fragment are in two disinct packets, except for the last fragment of the frame. Fragments are enumerated by the slice group index of the slice group contained within.

4.1. RTP Header Usage

The SSRC RTP field is used to discriminate each separate JPEG XS video stream from others. Within a specific JPEG XS video stream, identified by its SSRC, the picture counter field is used to identify to which picture a packet corresponds to.

4.2. Payload Header

The following figure illustrates the RTP payload header used in order to transport each JPEG XS video stream (identified by a distinct SSRC).

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+-+-+-------+-+-------------+-------------------------------+
| V |P|X|  CC   |M|     PT      |       Sequence number         |
+---+-+-+-------+-+-------------+-------------------------------+
|                           Timestamp                           |
+---------------------------------------------------------------+
|           Synchronization source (SSRC) identifier            |
+-----+-+-+-------+-----------------------+-+-------------------+
| Ver |f|l| SlcGp |      SliceGrpStart    |L| Picture Counter   |
+-----+-+-+-------+-----------------------+-+-------------------+
|                             Data                              |
+---------------------------------------------------------------+
```

Figure 1: RTP and payload headers

The version (V), padding (P), extension (X), CSRC count (CC), sequence number and synchronization source (SSRC) fields follow their respective definitions in RFC 3550 [RFC3550].

The timestamp should be based on a globally synchronized 90 kHz clock reference, and should correspond to the number of cycles since the SMPTE Epoch (as per defined in SMPTE ST 2059-1:2015 [SMPTE-ST2059]) modulo 2^32:

```
timestamp = floor((now - epoch)*90000) % 2^32
```

where now and epoch are real numbers expressed in seconds, now being the current timestamp and epoch the reference timestamp and floor indicates rounding to the next lower integer.

As per specified in RFC 3550 [RFC3550] and RFC 4175 [RFC4175], the RTP timestamp designates the sampling instant of the first octet of the picture to which the RTP packet belongs. Packets shall not include data from multiple frames, and all packets belonging to the same frame shall have the same timestamp. Several successive RTP packets will consequently have equal timestamps if they belong to the same picture (that is until the marker bit is set to 1, marking the last packet of the frame), and the timestamp is only increased when a new frame begins.

If the sampling instant does not correspond to an integer value of the clock, the value shall be truncated to the next lowest integer, with no ambiguity.

The remaining fields are defined as follows:

Field	Width	Description
Marker (M)	1 bit	The marker bit is used to indicate the last packet of a frame. This enables a decoder to finish decoding the picture, where it otherwise may need to wait for the next packet to explicitly know that the frame is finished.
Payload Type (PT)	7 bits	A dynamically allocated payload type field that designates the payload as JPEG XS video.
Vers	3 bits	This field indicates the version number of the payload header. The value of this field shall be 0 for the purpose of this edition of the RFC.
f	1 bit	The f field shall be set if a new fragment is started within this packet, i.e. if this packet contains the first byte of a fragment. NOTE: The first slice group of a frame and the JPEG XS header segment form a fragment. For that reason, the f-bit remains unset in the packet that contains the first byte of slice group 0 but does not also contains the first byte of the Video Essence box. All other slice groups form framgents of their own. The f bit allows a quick identificaiton of packets that start a fragment. The SliceGrpStart field (see below) can be used to identify the start of a slice group.
l	1 bit	The l field is a one-bit field that is cleared if the fragment to which the first byte of the packet belongs extends througout a subsequent packet. It is set if the fragment to which the first byte of the packet

		belongs ends in this packet.
SlcGp	5 bits	The SlcGp (Slice Group) field contains the slice group index modulo 64 that is contained in the fragment that is started in this packet. If no fragment starts in this packet, it contains the slice group index modulo 64 of the slice group that is contained in the fragment to which the first byte of the payload data of this packet belongs.
SliceGrpStart	11 bits	This field indicates the byte offset of the slice header marker (SLH, hex 0xff20, see ISO/IEC 21122-1 [ISO21122-1]) of the slice group that starts in this packet, relative to the start of the packet. If no slice group starts in this packet, this field shall be 0. NOTE: Since the payload data has a non-zero offset within a packet, this field can also be used to identify whether a slice group starts in a packet. If 0, no slice group starts in this packet. Consequently, for slice groups with a non-zero slice group index, this field will be non-zero if and only if the f-field is set. For the first slice gorup of a frame, however, the f bit indiecate sthe start of the fragment. whereas this field indicates the start of the slice group. Due to the non-zero size of the JPEG XS header segment, this need not to happen in the same packet.
F	1 bit	The F flag in the JPEG XS payload header shall be set if the packet contains the first byte of the JPEG XS Header Segment, and hence includes the first bytes of the Video Essence box. Readers may use this flag to extract information easily from the video essence box.
Picture number	10 bits	Counter indicating the current picture number modulo 2^11. The picture number is incremented by one

		at the beginning of each frame, and stays constant throuout all packets that contribute to to the same frame.

Table 1: Payload header fields description

4.3. Payload Data

The payload data of a JPEG XS transport stream consists of a concatenation of multiple JPEG XS Frames.

Each JPEG XS frame is the concatenation of multiple fragments where each fragment contains one and only one slice group. The first fragment of a frame also contains the Video Essence box and the JPEG XS header, the last fragment also contains the EOC marker. Figure Figure 2 depicts this layout.

Fragments may extend over multiple RTP packets. In particular, slice groups and by that fragments have to be sized such that the first byte of a fragment and the byte following the last byte of a fragment are in two distinct packets.

The start of a fragment can be identified by the "f" bit in the Payload header, the start of a slice group within a packet and its location in the packet by the SliceGrpStart field in the same Payload header.

```
  ^      +--------------------------------------------+     ^
  |      |             Video Essence Box              |     |
  |      |  +--------------------------------------+  |     |
  |      |  |  Sub boxes of the Video Essence Box  |  |     |
Frag-    |  +--------------------------------------+  |    JPEG
ment     |  : additional sub-boxes of the VE-Box   :  |     XS
 #0      |  +--------------------------------------+  |    Header
  |      |                                            |     Seg-
  |      +--------------------------------------------+     ment
  |      |              JPEG XS Header                |     |
  |      |  +--------------------------------------+  |     |
  |      |  |              SOC Marker              |  |     |
  |      |  +--------------------------------------+  |     |
  |      |  :      Additional Marker Segments      :  |     |
  |      |  +--------------------------------------+  |     |
  |      |                                            |     |
  |      +--------------------------------------------+     v
  |      |             Slice Group #0                 |
  |      |  +--------------------------------------+  |
```

```
   |     |  |   Slice #0 of Slice Group #0      |  |
   |     |  |  +--------------------------------+  |  |
   |     |  |  |            SLH Marker          |  |  |
   |     |  |  +--------------------------------+  |  |
   |     |  |  :     Entropy Coded Data         :  |  |
   |     |  |  +--------------------------------+  |  |
   |     |  +--------------------------------------+  |
   |     |  |   Slice #1 of Slice Group #0         |  |
   |     |  :                                      :  |
   |     |  +--------------------------------------+  |
   |     |  |   Slice #n-1 of Slice Group #0       |  |
   |     |  :                                      :  |
   v     |  +--------------------------------------+  |
   ^     +--------------------------------------------+
   |     |            Slice Group #1                  |
 Frag-   :                                            :
 ment    :                                            :
  #1     :                                            :
   |     :                                            :
   v     +--------------------------------------------+
         :                                            :
   ^     +--------------------------------------------+
   |     |            Slice Group #n-1                |
 Frag-   :                                            :
 ment    :                                            :
 #n-1    +--------------------------------------------+
   |     |             EOC Marker                     |
   v     +--------------------------------------------+
```

Figure 2: JPEG XS Payload Data

4.4. Traffic Shaping and Delivery Timing

The traffic shaping and delivery timing shall be in accordance with the Network Compatibility Model compliance definitions specified in SMPTE ST 2110-21 [SMPTE-ST2110-21] for either Narrow Linear Senders (Type NL) or Wide Senders (Type W).

Note: The Virtual Receiver Buffer Model compliance definitions of ST 2110-21 do not apply.

5. Congestion Control Considerations

Congestion control for RTP SHALL be used in accordance with RFC 3550 [RFC3550], and with any applicable RTP profile: e.g., RFC 3551 [RFC3551]. An additional requirement if best-effort service is being used is users of this payload format MUST monitor packet loss to ensure that the packet loss rate is within acceptable parameters.

Circuit Breakers [RFC8083] is an update to RTP [RFC3550] that defines criteria for when one is required to stop sending RTP Packet Streams. The circuit breakers is to be implemented and followed.

6. Payload Format Parameters

6.1. Media Type Definition

Type name: video

Subtype name: jpeg-xs

Encoding considerations:

This media type is framed and binary; see Section 4.8 in RFC 6838 [RFC6838].

Security considerations:

Please see the Security Considerations section in RFC XXXX

6.2. Mapping to SDP

6.2.1. General

A Session Description Protocol (SDP) object shall be created for each RTP stream and it shall be in accordance with the provisions of SMPTE ST 2110-10 [SMPTE-ST2110-10].

The information carried in the media type specification has a specific mapping to fields in the Session Description Protocol (SDP), which is commonly used to describe RTP sessions. When SDP is used to specify sessions employing the DV encoding, the mapping is as follows:

6.2.2. Media type and subtype

The media type ("video") goes in SDP "m=" as the media name.

The media subtype ("jpeg-xs") goes in SDP "a=rtpmap" as the encoding name. The RTP clock rate in "a=rtpmap" MUST be 90000, which for the payload format defined in this document is a 90 kHz clock.

6.2.3. Traffic shaping

The SDP object shall include the TP parameter and may include the CMAX parameter as specified in SMPTE ST 2110-21 [SMPTE-ST2110-21].

6.2.4. Other parameters

The SDP object shall include the following payload-format-specific parameter in the a=fmtp line:

SSN SMPTE Standard Number in the format: ST<number>-<part>:<year>
e.g. ST2110-20:2017
The number shall be that of the JPEG XS standard

Any remaining parameters go in the SDP "a=fmtp" attribute by copying them directly from the media type string as a semicolon-separated list of parameter=value pairs.

6.2.5. Offer/Answer Considerations

The following considerations apply when using SDP offer/answer procedures [RFC3264] to negotiate the use of the JPEG XS payload in RTP:

- o The "encode" parameter can be used for sendrecv, sendonly, and recvonly streams. Each encode type MUST use a separate payload type number.
- o Any unknown parameter in an offer MUST be ignored by the receiver and MUST NOT be included in the answer.

7. IANA Considerations

This memo requests that IANA registers video/jpeg-xs as specified in Section 6.1. The media type is also requested to be added to the IANA registry for "RTP Payload Format MIME types" [1].

8. Security Considerations

[FIXME: imported from RFC 7587]

RTP packets using the payload format defined in this specification are subject to the security considerations discussed in the RTP specification [RFC3550] and in any applicable RTP profile such as RTP/AVP [RFC3551], RTP/AVPF [RFC4585], RTP/SAVP [RFC3711], or RTP/SAVPF [RFC5124]. This implies that confidentiality of the media streams is achieved by encryption.

However, as "Securing the RTP Framework: Why RTP Does Not Mandate a Single Media Security Solution" [RFC7202] discusses, it is not an RTP payload format's responsibility to discuss or mandate what solutions are used to meet the basic security goals like confidentiality, integrity, and source authenticity for RTP in general. This responsibility lies on anyone using RTP in an application. They can find guidance on available security mechanisms and important considerations in "Options for Securing RTP Sessions" [RFC7201]. Applications SHOULD use one or more appropriate strong security mechanisms.

This payload format and the JPEG XS encoding do not exhibit any substantial non-uniformity, either in output or in complexity to perform the decoding operation and thus are unlikely to pose a denial-of-service threat due to the receipt of pathological datagrams.

[FIXME: imported from RFC 4175]

It is important to note that HD or UHDTV JPEG XS-encoded video can have significant bandwidth requirements (typically more than 1 Gbps for ultra high-definition video, especially if using high framerate). This is sufficient to cause potential for denial-of-service if transmitted onto most currently available Internet paths.

Accordingly, if best-effort service is being used, users of this payload format MUST monitor packet loss to ensure that the packet loss rate is within acceptable parameters. Packet loss is considered acceptable if a TCP flow across the same network path, and experiencing the same network conditions, would achieve an average throughput, measured on a reasonable timescale, that is not less than the RTP flow is achieving. This condition can be satisfied by implementing congestion control mechanisms to adapt the transmission rate (or the number of layers subscribed for a layered multicast session), or by arranging for a receiver to leave the session if the loss rate is unacceptably high.

This payload format may also be used in networks that provide quality-of-service guarantees. If enhanced service is being used, receivers SHOULD monitor packet loss to ensure that the service that was requested is actually being delivered. If it is not, then they SHOULD assume that they are receiving best-effort service and behave accordingly.

9. RFC Editor Considerations

Note to RFC Editor: This section may be removed after carrying out all the instructions of this section.

RFC XXXX is to be replaced by the RFC number this specification receives when published.

10. References

10.1. Normative References

[ISO15444-1]
International Organization for Standardization (ISO) - International Electrotechnical Commission (IEC), "Information technology - JPEG 2000 image coding system: Core coding system", ISO/IEC IS 15444-1, 2016, <https://www.iso.org/standard/70018.html>.

[ISO18477-3]
International Organization for Standardization (ISO) - International Electrotechnical Commission (IEC), "Information technology - Scalable compression and coding of continuous-tone still images - Part 3: Box file format", ISO/IEC 18477-3:2015, 2015, <https://www.iso.org/standard/66071.html>.

[ISO21122-1]
International Organization for Standardization (ISO) - International Electrotechnical Commission (IEC), "Information technology - Low-latency lightweight image coding system - Part 1: Core coding system", ISO/IEC DIS 21122-1, under development, <https://www.iso.org/standard/74535.html>.

[ISO21122-3]
International Organization for Standardization (ISO) - International Electrotechnical Commission (IEC), "Information technology - Low-latency lightweight image coding system - Part 3: Transport and container formats", ISO/IEC NP 21122-3, under development, <https://www.iso.org/standard/74537.html>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, DOI 10.17487/RFC3264, June 2002, <https://www.rfc-editor.org/info/rfc3264>.

[RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <https://www.rfc-editor.org/info/rfc3550>.

[RFC3551] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, RFC 3551, DOI 10.17487/RFC3551, July 2003, <https://www.rfc-editor.org/info/rfc3551>.

[RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, DOI 10.17487/RFC3711, March 2004, <https://www.rfc-editor.org/info/rfc3711>.

[RFC6838] Freed, N., Klensin, J., and T. Hansen, "Media Type Specifications and Registration Procedures", BCP 13, RFC 6838, DOI 10.17487/RFC6838, January 2013, <https://www.rfc-editor.org/info/rfc6838>.

[RFC8083] Perkins, C. and V. Singh, "Multimedia Congestion Control: Circuit Breakers for Unicast RTP Sessions", RFC 8083, DOI 10.17487/RFC8083, March 2017, <https://www.rfc-editor.org/info/rfc8083>.

[SMPTE-ST2110-10] Society of Motion Picture and Television Engineers, "SMPTE Standard - Professional Media Over Managed IP Networks: System Timing and Definitions", SMPTE ST 2110-10:2017, 2017, <https://doi.org/10.5594/SMPTE.ST2110-10.2017>.

[SMPTE-ST2110-21] Society of Motion Picture and Television Engineers, "SMPTE Standard - Professional Media Over Managed IP Networks: Traffic Shaping and Delivery Timing for Video", SMPTE ST 2110-21:2017, 2017, <https://doi.org/10.5594/SMPTE.ST2110-21.2017>.

10.2. Informative References

[RFC4175] Gharai, L. and C. Perkins, "RTP Payload Format for Uncompressed Video", RFC 4175, DOI 10.17487/RFC4175, September 2005, <https://www.rfc-editor.org/info/rfc4175>.

[RFC4585] Ott, J., Wenger, S., Sato, N., Burmeister, C., and J. Rey, "Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)", RFC 4585, DOI 10.17487/RFC4585, July 2006, <https://www.rfc-editor.org/info/rfc4585>.

[RFC5124] Ott, J. and E. Carrara, "Extended Secure RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/SAVPF)", RFC 5124, DOI 10.17487/RFC5124, February 2008, <https://www.rfc-editor.org/info/rfc5124>.

[RFC7201] Westerlund, M. and C. Perkins, "Options for Securing RTP Sessions", RFC 7201, DOI 10.17487/RFC7201, April 2014, <https://www.rfc-editor.org/info/rfc7201>.

[RFC7202] Perkins, C. and M. Westerlund, "Securing the RTP Framework: Why RTP Does Not Mandate a Single Media Security Solution", RFC 7202, DOI 10.17487/RFC7202, April 2014, <https://www.rfc-editor.org/info/rfc7202>.

[SMPTE-ST2059] Society of Motion Picture and Television Engineers, "SMPTE Standard - Generation and Alignment of Interface Signals to the SMPTE Epoch", SMPTE ST 2059-1:2015, 2015, <https://doi.org/10.5594/SMPTE.ST2059-1.2015>.

10.3. URIs

[1] http://www.iana.org/assignments/rtp-parameters

Authors' Addresses

Sebastien Lugan
intoPIX S.A.
Rue Emile Francqui, 9
1435 Mont-Saint-Guibert
Belgium

Phone: +32 10 23 84 70
Email: s.lugan@intopix.com
URI: http://www.intopix.com

Gael Rouvroy
intoPIX S.A.
Rue Emile Francqui, 9
1435 Mont-Saint-Guibert
Belgium

Phone: +32 10 23 84 70
Email: g.rouvroy@intopix.com
URI: http://www.intopix.com

Antonin Descampe
intoPIX S.A.
Rue Emile Francqui, 9
1435 Mont-Saint-Guibert
Belgium

Phone: +32 10 23 84 70
Email: a.descampe@intopix.com
URI: http://www.intopix.com

Thomas Richter
Fraunhofer IIS
Am Wolfsmantel 33
91048 Erlangen
Germany

Phone: +49 9131 776 5126
Email: thomas.richter@iis.fraunhofer.de
URI: https://www.iis.fraunhofer.de/

Alexandre Willeme
Universite catholique de Louvain
Place du Levant, 2 - bte L5.04.04
1348 Louvain-la-Neuve
Belgium

Phone: +32 10 47 80 82
Email: alexandre.willeme@uclouvain.be
URI: https://uclouvain.be/en/icteam