Versions: Standards Track
Avtcore Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 7, 2022

L. Ploumhans
Silex Insight

November 8, 2021

RTP Payload Format for Colibri Video
draft-ploumhans-avtcore-rtp-colibri-00

Abstract

This memo describes an RTP Payload format for the Colibri Video "Standard". This document describes the transport of Colibri video in RTP packets and has applications for low-complexity, high-bandwidth streaming of lossy compressed video.

The Colibri Video is intended for low latency video compression (with latency potentially on the order of lines of video) at high to medium data rates (with compression ratios betwen 2:1 and 20:1).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 25, 2022.

Copyright Notice

Table of Contents

1. Introduction

This memo specifies an RTP payload format for the Colibri video coding.

The Colibri codec is a wavelet-based codec intended primarily for professional video use with high bit-rates and low to medium levels of compression. It has been designed to be low-complexity, and potentially have a very low latency through both encoder and decoder: with some choices of parameters this latency may be as low as a few lines of video.

The low level of complexity in the Colibri codec allows for this low latency operation but also means that it lacks many of the more powerful compression techniques used in other codecs. As such it is suitable for low compression ratios that produce coded data rates around half to a quarter of that of uncompressed video, at a similar visual quality. In some applications, the compression ratio can be increased up to 20.

The primary use for Colibri is likely to be in professional video production environments and "screen content".

2. Conventions, Definitions and Acronyms

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Media Format Description

The Colibri specification defines a Colibri stream as being composed of one or more Frames. Each Frame contains all of the needed parameters and metadata for configuring the decoder.

Each Frame consists of a Colibri header, a Colibri Parameters segment and a sequence of slices, describing a picture. Each picture represents a frame in a progressively scanned video Sequence or a field in an interlaced video Sequence.

Each slice in a Colibri frame can be coded either as an intra slice (not using information from a previous frame) or as an inter slice (using information on the same slice position of previously coded Colibri frames). A Colibri encoder will normally force the encoding of each slice to intra regularly, so that a decoder can accumulate information over several frames and be able to display a complete valid frame after some time.

At time of writing there is no specific definition of padding or auxiliary data for the Colibri codec. Padding is always possible in the Colibri Parameters segement or in a Colibri slice, but is not specifically seperated from the useful payload.

The Colibri Parameters segment contains all the parameters required to set up the Colibri decoder for the next picture.

Each Slice represents a rectangular region of the transformed picture. Slices within a picture may vary in coded length, but all represent the same shape and size of rectangle in the picture.

4. Payload format

This section specifies the payload format for Colibri streams over the Real-Time Transport Protocol (RTP) [RFC3550].

4.1. RTP Packetization

This specification provides two different transmission modes of the Colibri pictures over RTP.

The first transmission mode is called the "picture packetization" mode. In this mode, one Colibri picture is transmitted as a single packetization unit, split over several RTP packets, which are the Picture packets. This is illustrated in Figure 1.

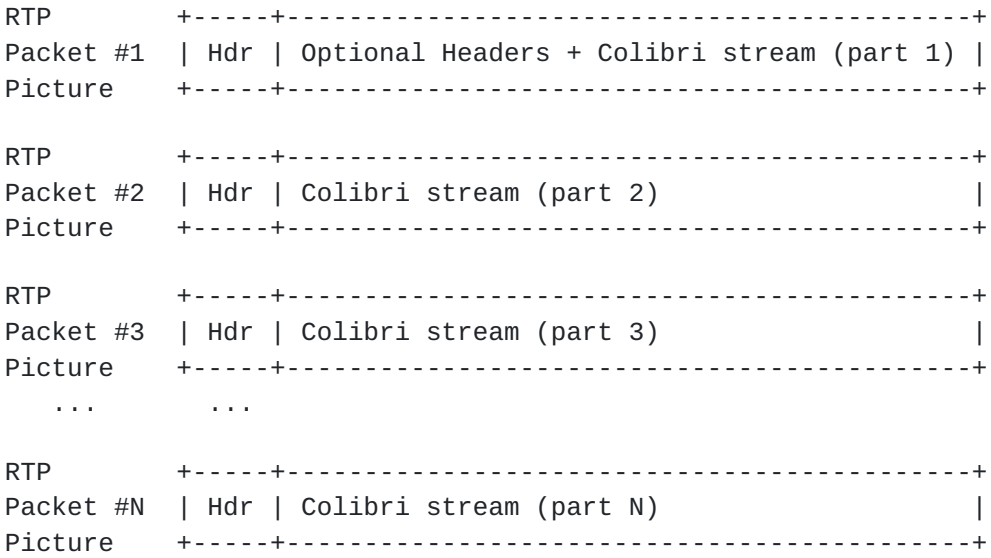

Figure 1: Example of picture packetization mode

The second transmission mode is called the "slice packetization" mode. In this mode, one Colibri picture is split into several RTP packets. The first RTP packet for the Colibri picture is the Headers Packet and contains all the header and parameters information. The remaining RTP packets are the Slice Packets and will contain one or more Colibri slices. A Colibri slice may not be split accross more than one RTP packet. Colibri slices that are not horizontally aligned MAY NOT be part of the same RTP packet.

This is illustrated in Figure 2, where NS(x) is the number of slices of the Colibri picture transmitted over RTP packets 1 to x.

In addition to the Headers Packets and the Slice Packets, this RTP specification defines the Padding Packets and the Auxiliary Packets, which allow the transport of padding and auxiliary data, for the Slice packetization mode only.

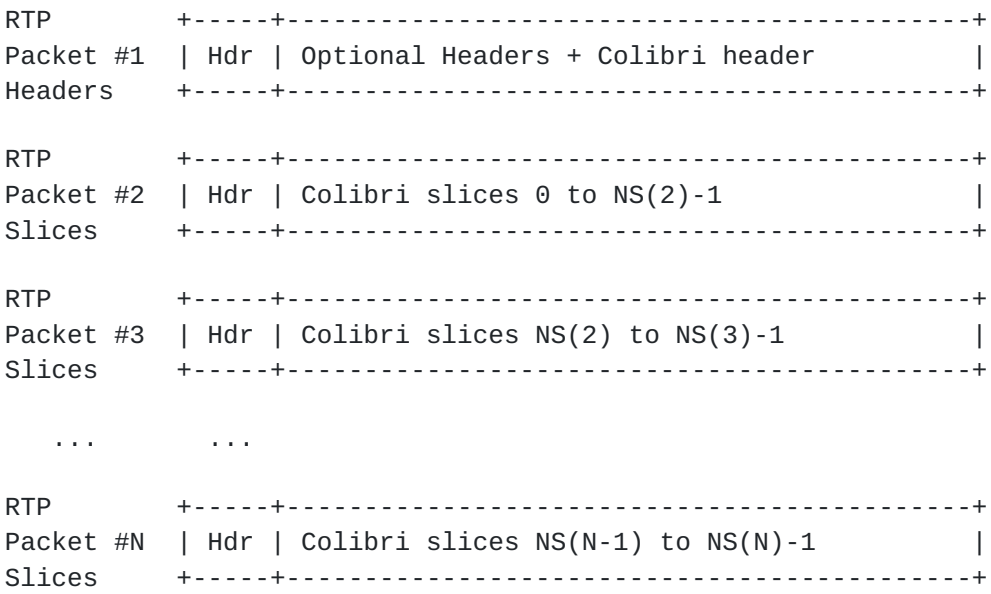

Figure 2: Example of slice packetization mode

4.2. RTP Packets for Colibri

The structure of an RTP Packet for Colibri is illustrated in Figure 3.

All fields in the headers longer than a single bit are interpreted as unsigned integers in network byte order.

The fields of the RTP header have the following additional notes on their usage:

Marker Bit (M): 1 bit The marker bit MUST be set on any packet which contains the final slice in a coded picture and MUST NOT be set otherwise.

Payload Type (PT): 7 bits A dynamically allocated payload type field that designates the payload as Colibri coded video.

Timestamp: 32 bits The timestamp corresponds to the sampling instant of the coded picture. A 90kHz clock SHOULD be used. A single RTP packet MUST NOT contain coded data for more than one coded picture, so there is no ambiguity here.

The remaining RTP header fields are used as specified in RTP [RFC3550].

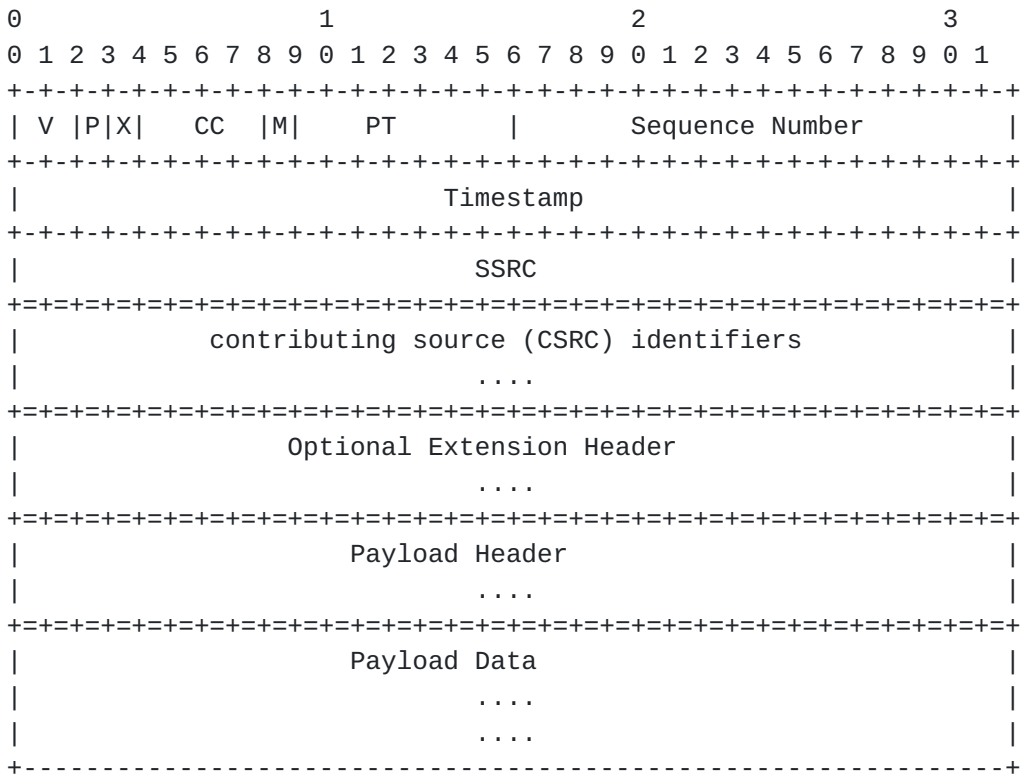

Figure 3: RTP Payload Format For Colibri Parameters

4.3. Payload Header

This document defines five types of RTP packets in a Colibri media stream:

- A Colibri Picture Packet, containing a portion of a Colibri Picture in Picture Packetization mode
- A Colibri Headers Packet, containing optional headers and the Colibri header in Slice Packetization mode
- A Colibri Slices Packet, containg at least one Colibri slice in Slice Packetization mode
- An Auxiliary data Packet, containing data auxiliary to a Colibri Picture, but not actually part of the Colibri stream, in Slice Packetization mode.
- A Padding Packet, containing irrelevant data, which is not part of the Colibri stream, in Slice Packetization mode.

4.3.1. Payload Header for a Colibri Picture Packet

The Payload Header for a Colibri Picture Packet is illustrated in Figure 4.

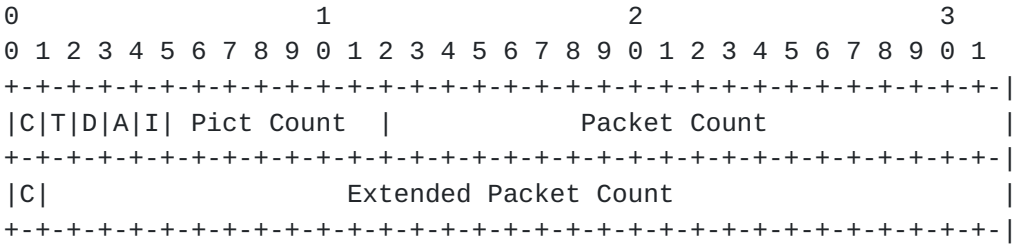

Figure 4: RTP Payload Header for a Colibri Picture Packet

The fields of the first 4 bytes of the Payload header are defined as follows:

C : 1 bit MUST be set to 1 if the Payload header must be extended by 32 additional bits.

T : 1 bit Indicates the packTization mode. MUST be set to 0 for the Picture packetization mode.

D : 1 bit Indicates that the payload contains an optional Video Definition header. When the Packet counter is 0, the Video Definition header may be included prior to the optional Color Specification header and the Colibri Header. When the Packet counter is not 0, the D bit MUST be set to 0.

A : 1 bit Indicates that the payload contains an optional Color Specification header. When the Packet counter is 0, the Color Specification header may be included prior to the Colibri Header. When the Packet counter is not 0, the A bit MUST be set to 0.

I : 1 bit Indicates whether the Video transmitted over RTP is interlaced or progressive. The value of I must be set to 1 when the video is interlaced. The value of P must be set to 0 when the video is progressive.

Pict Count : 7 bits Indicates the picture number modulo 128. In the case of interlaced video, the LSbit of the Pict Count field indicates the field number. The first field of a frame always has an even Picture Count, and the second field of a frame always has an odd Picture Count.

Packet Count : 20 bits Indicates the Packet number within the current Colibri Picture. It is set to 0 at the start of the Colibri Picture, and is incremented by 1 for every Colibri Picture Packet that belongs to the same Colibri Picture.

A Colibri Picture Packet is identified by the value of the T field, which must be set to 0.

If the value of the C field is 1, the payload header is extended by 32 bits :

C : 1 bit MUST be set to 1 if the Payload header must be extended by 32 additional bits.

Extended Packet Count : 31 bits The Packet Count is extended by 31 bits.

The Payload header (and the Packet Count) may be extended indefinitely.

4.3.2. Payload Header for a Colibri Headers Packet

The Payload Header for a Colibri Headers Packet is illustrated in Figure 5.

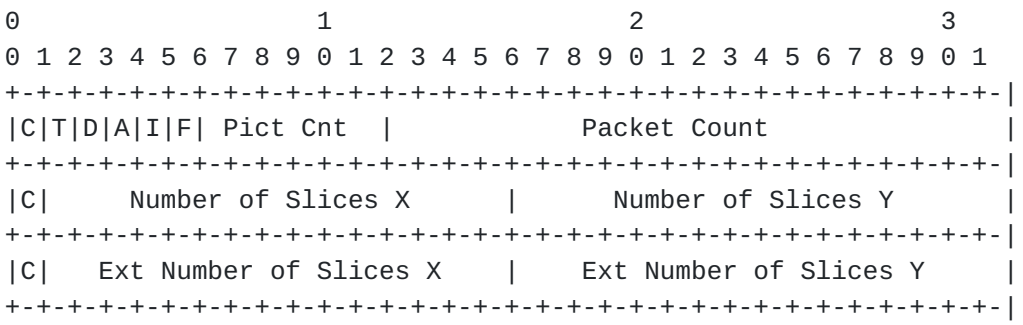

Figure 5: RTP Payload Header for a Colibri Headers Packet

The fields of the first 4 bytes of the Payload header are defined as follows:

C : 1 bit MUST be set to 1 if the Payload header must be extended by 32 additional bits. It must be set to 1 when T is 1.

T : 1 bit Indicates the packTization mode. MUST be set to 1 for the Slice packetization mode.

D : 1 bit Indicates that the payload contains an optional Video Definition header. When F is 1, the Video Definition header may be included prior to the optional Color Specification header and the Colibri Header. When F is 0, the D bit has another meaning.

A : 1 bit Indicates that the payload contains an optional Color Specification header. When F is 1, the Color Specification header may be included prior to the Colibri Header. When F is 0, the A bit has another meaning.

I : 1 bit Indicates whether the Video transmitted over RTP is interlaced or progressive. The value of I must be set to 1 when the video is interlaced. The value of P must be set to 0 when the video is progressive.

F : 1 bit Indicates the First packet of the Colibri picture. F Must be set to 1 for the Colibri Headers Packet, and must be set to 0 for any other packet in Slice packetization mode.

Pict Count : 6 bits Indicates the picture number modulo 64. In the case of interlaced video, the LSbit of the Pict Count field indicates the field number. The first field of a frame always has an even Picture Count, and the second field of a frame always has an odd Picture Count.

Packet Count : 20 bits Indicates the Packet number within the current Colibri Picture. It is set to 0 at the start of the Colibri Picture. Packet Count must be 0 for the Colibri Headers packet.

A Colibri Headers Packet is identified by the value of the T field (1), and the F field (1).

With the value of the C field being 1, the payload header is extended by 32 bits.

The fields of the next 4 bytes of the Payload header are defined as follows:

C : 1 bit MUST be set to 1 if the Payload header must be extended by 32 additional bits.

Number of Slices X: 15 bits MUST contain the horizontal size, in slices, of the Colibri picture.

Number of Slices Y: 16 bits MUST contain the vertical size, in slices, of the Colibri picture.

If the value of the C field is 1 in byte 5 of the payload header, the payload header is further extended by 32 bits.

C : 1 bit MUST be set to 1 if the Payload header must be extended by 32 additional bits.

Ext Number of Slices X : 15 bits The Number of Slices X is extended by 15 bits.

Ext Number of Slices Y : 16 bits The Number of Slices Y is extended by 16 bits.

The Payload header (and the Number of Slices X and Number of Slices Y) may be extended indefinitely.

4.3.3. Payload Header for a Colibri Slices Packet

The Payload Header for a Colibri Slices Packet is illustrated in Figure 6.

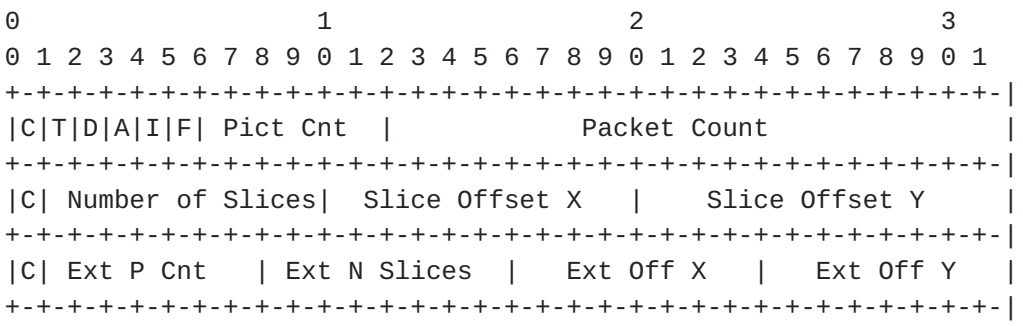

Figure 6: RTP Payload Header for a Colibri Slices Packet

The fields of the first 4 bytes of the Payload header are defined as follows:

C : 1 bit MUST be set to 1 if the Payload header must be extended by 32 additional bits. It must be set to 1 when T is 1.

T : 1 bit Indicates the packTization mode. MUST be set to 1 for the Slice packetization mode.

D : 1 bit Indicates that the payload contains Padding. Must be set to 0 for a Colibri Slices Packet.

A : 1 bit Indicates that the payload contains Auxiliary data. Must be set to 0 for a Colibri Slices Packet.

I : 1 bit Indicates whether the Video transmitted over RTP is interlaced or progressive. The value of I must be set to 1 when the video is interlaced. The value of P must be set to 0 when the video is progressive.

F : 1 bit Indicates the First packet of the Colibri picture. F Must be set to 0 for the Colibri Slices Packet.

Pict Count : 6 bits Indicates the picture number modulo 64. In the case of interlaced video, the LSbit of the Pict Count field indicates the field number. The first field of a frame always has an even Picture Count, and the second field of a frame always has an odd Picture Count.

Packet Count : 20 bits Indicates the Packet number within the current Colibri Picture. It is set to 0 at the start of the Colibri Picture, and is incremented by 1 for every Colibri Slices Packet that belongs to the same Colibri Picture.

A Colibri Slices Packet is identified by the value of the T field (1), the F field (0), the D field (0) and the A field (0).

With the value of the C field being 1, the payload header is extended by 32 bits.

The fields of the next 4 bytes of the Payload header are defined as follows:

C : 1 bit MUST be set to 1 if the Payload header must be extended by 32 additional bits.

Number of Slices : 9 bits Indicates the number of slices included in the Colibri Slices Packet.

Slices Offset X: 10 bits Indicates the horizontal position of the first slice in the Colibri Slices packet.

Slices Offset Y: 12 bits Indicates the vertical position of the slices in the Colibri Slices packet.

If the value of the C field is 1 in byte 5 of the payload header, the payload header is further extended by 32 bits.

C : 1 bit MUST be set to 1 if the Payload header must be extended by 32 additional bits.

Ext P Cnt : 7 bits The Packet Count is extended by 7 bits.

Ext N Slices : 8 bits The Number of Slices is extended by 8 bits.

Ext Off X : 8 bits The Slices Offset X is extended by 8 bits.

Ext Off Y : 8 bits The Slices Offset Y is extended by 8 bits.

The Payload header (and the Packet Count, the Number of Slices, the Slice Offset X and the Slice Offset Y) may be extended indefinitely.

4.3.4. Payload Header for a Colibri Padding Packet

The Payload Header for a Colibri Padding Packet is illustrated in Figure 7.

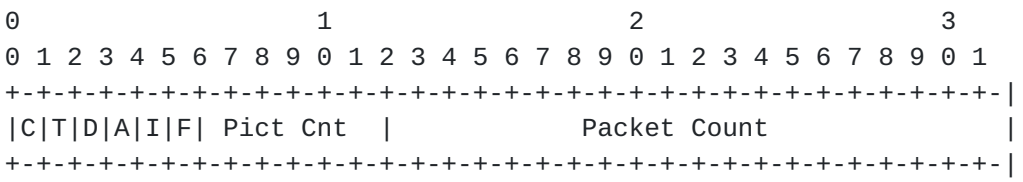

Figure 7: RTP Payload Header for a Colibri Padding Packet

The fields of the Payload header are defined as follows:

C : 1 bit MUST be set to 0.

T : 1 bit Indicates the packTization mode. MUST be set to 1 for the Slice packetization mode.

D : 1 bit Indicates that the payload contains Padding. Must be set to 1 for a Colibri Padding Packet.

A : 1 bit Indicates that the payload contains Auxiliary data. Must be set to 0 for a Colibri Padding Packet.

I : 1 bit Indicates whether the Video transmitted over RTP is interlaced or progressive. The value of I must be set to 1 when the video is interlaced. The value of P must be set to 0 when the video is progressive.

F : 1 bit Indicates the First packet of the Colibri picture. F Must be set to 0 for a Colibri Padding Packet.

Pict Count : 6 bits Indicates the picture number modulo 64. In the case of interlaced video, the LSbit of the Pict Count field indicates the field number. The first field of a frame always has an even Picture Count, and the second field of a frame always has an odd Picture Count.

Packet Count : 20 bits This field is ignored for a Colibri Padding Packet.

A Colibri Padding Packet is identified by the value of the T field (1), the F field (0), the D field (1) and the A field (0).

4.3.5. Payload Header for a Colibri Auxiliary Packet

The Payload Header for a Colibri Auxiliary Packet is illustrated in Figure 8.

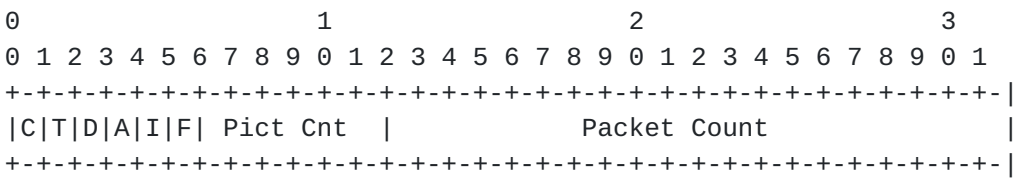

Figure 8: RTP Payload Header for a Colibri Auxiliary Packet

The fields of the Payload header are defined as follows:

C : 1 bit MUST be set to 0.

T : 1 bit Indicates the packTization mode. MUST be set to 1 for the Slice packetization mode.

D : 1 bit Indicates that the payload contains Padding. Must be set to 0 for a Colibri Auxiliary Packet.

A : 1 bit Indicates that the payload contains Auxiliary data. Must be set to 1 for a Colibri Auxiliary Packet.

I : 1 bit Indicates whether the Video transmitted over RTP is interlaced or progressive. The value of I must be set to 1 when the video is interlaced. The value of P must be set to 0 when the video is progressive.

F : 1 bit Indicates the First packet of the Colibri picture. F Must be set to 0 for a Colibri Auxiliary Packet.

Pict Count : 6 bits Indicates the picture number modulo 64. In the case of interlaced video, the LSbit of the Pict Count field indicates the field number. The first field of a frame always has an even Picture Count, and the second field of a frame always has an odd Picture Count.

Packet Count : 20 bits This field is ignored for a Colibri Auxiliary Packet.

A Colibri Auxiliary Packet is identified by the value of the T field (1), the F field (0), the D field (0) and the A field (1).

4.3.6. Optional Video Definition header

The optional Video Definition header contains information relative to the video format and the compressed video bitrate. It can optionally appear directly after the Payload Header for a Colibri Picture Packet, when the D field is set to 1 ; or directly after the Payload Header for a Colibri Headers Packet, when the D field is set to 1.

The optional Video Definition header is illustrated in Figure 9.

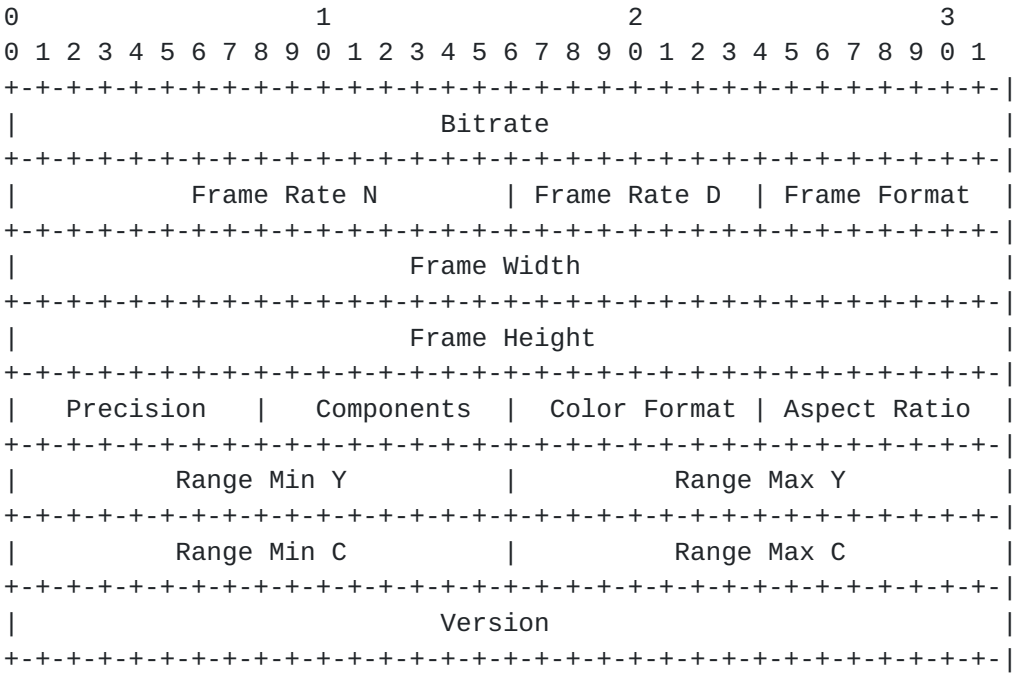

Figure 9: Optional Video Definition Header description

The fields of the optional Video Definition Header are defined as follows:

Bitrate : 32 bits The maximum bitrate of the video stream.

Frame Rate N : 16 bits Frame rate numerator.

Frame Rate D : 8 bits Frame rate denominator. When the value is set to 0, the denominator value is replaced by 1.001.

Frame Format : 8 bits A value of 0 indicates that the video sequence is progressive. A value of 1 indicates that the video sequence is interlaced, and that the first field is the bottom field (even-number pictures contain the bottom field of the video sequence). A value of 2 indicates the the video sequence is interlaced, and that the first field is the top field (even-number pictures contain the top field of the video sequence). Other values should not be used.

Frame Width : 32 bits Indicates the width of the video frames.

Frame Height : 32 bits Indicates the height of the video frames.

Precision : 8 bits Indicates the number of bits per component of the video samples.

Components : 8 bits Indicates the number of video components in the video frames.

Color format : 8 bits Indicates the color format of the video frames. The different possible values are indicated in Figure 10.

Value	Definition
0	4:4:4 YCbCr
1	4:2:2 YCbCr
2	4:2:0 YCbCr
3	4:4:4 RGB
Other	Reserved

Figure 10: Color Format value

Aspect ratio : 8 bits Indicates the aspect ratio of the pixels. The different possible values are indicated in Figure 11.

Value	Definition
0	1:1
1	4:3
Other	Reserved

Figure 11: Aspect Ratio value

Range Min Y : 16 bits For color formats that include a luminance component, indicates the minimum value of the luminance, when represented as an unsigned integer. For other color formats, indicates the minimum value of the color components, when represented as an unsigned integer.

Range Max Y : 16 bits For color formats that include a luminance component, indicates the maximum value of the luminance, when represented as an unsigned integer. For other color formats, indicates the maximum value of the color components, when represented as an unsigned integer. When the value of Range Max Y is 0, this indicates that the maximum value is limited only by the component precision, as indicated in the Precision field.

Range Min C : 16 bits For color formats that include a luminance component, indicates the minimum value of the chrominances, when represented as an unsigned integer. For other color formats, this field is not used, and should have the value 0.

Range Max C : 16 bits For color formats that include a luminance component, indicates the maximum value of the chrominances, when represented as an unsigned integer. For other color formats, this field is not used, and should have the value 0.

Version : 32 bits This field indicates the Colibri version required by a decoder to be able to decode the Colibri Video stream.

4.3.7. Optional Color Specification header

The optional Color Specification header contains information relative to the colour space of the decoded video sequence. It can optionally appear :

- In a Colibri Picture Packet, directly after the Payload Header, when the D field is set to 0 and the A field is set to 1
- In a Colibri Picture Packet, directly after the optional Video Definition header, when the D field is set to 1 and the A field is set to 1
- In a Colibri Headers Packet, directly after the Payload Header, when the D field is set to 0 and the A field is set to 1
- In a Colibri Headers Packet, directly after the optional Video Definition header, when the D field is set to 1 and the A field is set to 1

Figure 12: Optional Color Specification Header description

The fields of the optional Color Specification Header are defined as follows:

Color Primaries : 16 bits Indicates the standard to use for color primaries of the video frames. The different possible values are indicated in Figure 13.

Color Matrix : 16 bits Indicates the standard to use for the coefficients of the color matrix. The different possible values are indicated in Figure 14.

Transfer Function : 16 bits Indicates the standard to use for the transfer function on the video frames. The different possible values are indicated in Figure 15.

Reserved : 80 bits This field is not used and should be set to 0.

Value	Definition
0	ITU-R BT.709 [ITU709]
1	ITU-R BT.601 525 [ITU601]
2	ITU-R BT.601 625 [ITU601]
3	SMPTE 428.1 [SMPTE428]
4	ITU-R BT.2020 [ITU2020]
5	ITU-R BT.2100 [ITU2100]
Other	Reserved

Figure 13: Color Primaries value

Value	Definition
0	ITU-R BT.709 [ITU709]
1	ITU-R BT.601 525 [ITU601]
2	ITU-R BT.601 625 [ITU601]
3	SMPTE 428.1 [SMPTE428]
4	ITU-R BT.2020 [ITU2020]
5	ITU-R BT.2100 [ITU2100]
Other	Reserved

Figure 14: Color Matrix value

Value	Definition
0	ITU-R BT.709 [ITU709]
1	ITU-R BT.601 [ITU601]
2	SMPTE 428.1 [SMPTE428]
3	ITU-R BT.2020 [ITU2020]
4	ITU-R BT.2100 [ITU2100]
Other	Reserved

Figure 15: Transfer Function value

4.4. Stream Constraints

There are some constraints which a Sequence needs to conform to in order to be transmissible with this specification.

- In Picture packetization mode, the optional Video Definition header and the optional Color Specification headers combined SHOULD be small enough that the RTP packet carrying it will fit within the network MTU size.
- In Slice packetization mode, the optional Video Definition header, the optional Color Specification headers and the Colibri Parameters segment SHOULD be small enough that the RTP packet carrying it will fit within the network MTU size.
- In Slice packetization mode, every coded slice SHOULD be small enough that the RTP packet carrying it will fit within the network MTU size.

Sending a Stream which does not meet the above requirements is not possible unless the stream is re-encoded by a Colibri Encoder to meet them.

4.5. Payload Data

In Picture packetization mode, the payload data for the first packet must be the optional Video Definition header (if indicated in the payload header), followed by the optional Color Specification header (if indicated in the payload header), followed by the beginning of the encoded Colibri Picture. The payload data for the other packets in Picture packetization mode contain, in order, the next segment of the encoded Colibri Picture. Segments must be sent in order.

For the Colibri Headers Packet, the payload data must be the optional Video Definition header (if indicated in the payload header), followed by the optional Color Specification header (if indicated in the payload header), followed by the Colibri Headers and the Colibri Parameters exactly as it appears in the Colibri picture.

For the Colibri Slices packet type the payload data MUST be a specified number of coded slices in the same order that they appear in the Colibri slice sequence. All Colibri Slices packets may only contain slices which have the same vertical position. As a consequence, the first Colibri Slices packet with a new value of Slice Offset Y MUST have a value of Slice Offset X equal to 0.

4.5.1. Reassembling the Data

To reassemble the data in the RTP packets into a valid Colibri video stream with Picture packetization mode:

- The receiver SHOULD take the payload data from each Colibri Picture packet and remove the optional Video Definition header (if present) and the optional Color Specification header (if present).

To reassemble the data in the RTP packets into a valid Colibri video stream with Slice packetization mode:

- The receiver SHOULD take the data from each Colibri Headers packet and remove the optional Video Definition header (if present) and the optional Color Specification header (if present). The resulting Colibri Header and Colibri Parameters MUST be included in the output stream before any coded slice which followed it in the RTP stream.
- The receiver SHOULD take the data from each Colibri Slices packet. The Colibri slices must be sent in the right order. If a Colibri Slices packet is missing, the RTP receiver MUST provide a sequence of replacement slices corresponding to the number of missing slices. The choice of what replacement slice will be is left to the implementer to decide, but it is recommended to use either the empty slice, as shown in Figure 16, or the reuse slice, as shown in Figure 17. In the case of the empty slice, the decoder will erase the content of the previous slices and provide all 0 wavelet coefficients, which will result in a grey area in the output picture. In the case of the reuse slice, the decoder will reuse all the available content from the previous frame, as if the slice was not modified between the previous and the current frame.
- The receiver can either discard or transmit to an auxiliary channel the Colibri Auxiliray packets.
- The receiver SHOULD drop the Colibri Padding packets.

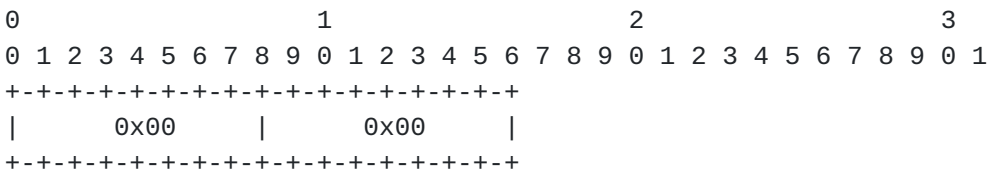

Figure 16: Colibri Empty slice

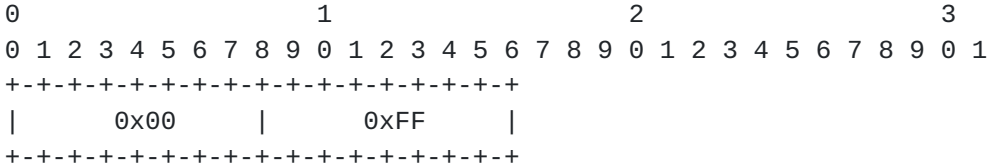

Figure 17: Colibri Reuse slice

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, and applications implementing this standard MUST comply with it. RFC 8085 [RFC8085] provides additional information on the best practices for applying congestion control to UDP streams.

In particular it should be noted that the expected data rate for RTP sessions which use this profile is likely to be close to gigabits per second. If used on a closed network which has been correctly provisioned for the expected data rates this might not pose a problem, but there is always the risk of data getting out onto the open internet.

6. Payload Format Parameters

This RTP payload format is identified using the video/colibri media type which is registered in accordance with RFC 4855 [RFC4855] and using the template of RFC 6838 [RFC6838].

6.1. Media Type Definition

Type name:

video

Subtype name:

colibri

Required parameters:

rate: The RTP timestamp clock rate. Applications using this payload format SHOULD use a value of 90000.

Optional parameters:

version: the Colibri specification version in use. The only currently allowed value is "1".

level: The Colibri level in use. Any integer may be used.

Encoding considerations:

This media type is framed and binary, see section 4.8 in RFC6838 [RFC6838].

Security considerations:

Please see security consideration in RFCXXXX

Interoperability considerations: N/A

Published specification:

RFC XXXX

Applications that use this media type:

Video Communication.

Fragment Identifier Considerations: N/A

Additional information: N/A

Person & email address to contact for further information:

luc.ploumhans@silexinsight.com

Intended usage:

COMMON

Restrictions on usage:

This media type depends on RTP framing, and hence is only defined for transfer via RTP [RFC3550]. Transport within other framing protocols is not defined at this time.

Author:

Change controller:

IETF Payload working group delegated from the IESG.

Provisional registration? (standards tree only):

No

6.2. Mapping to SDP

The mapping of the above defined payload format media type and its parameters SHALL be done according to Section 3 of RFC 4855 [RFC4855].

- The type name ("video") goes in SDP "m=" as the media name.
- The subtype name ("colibri") goes in SDP "a=rtpmap" as the encoding name, followed by a slash ("/") and the rate parameter.
- The required parameter profile and the optional parameters version and level, when present, are included in the "a=fmtp" attribute line of SDP as a semicolon-separated list of parameter=value pairs.

Version and level SHALL be specified in decimal when present.

For example, a sample SDP mapping for Colibri could be as follows:

```
m=video 30000 RTP/AVP 112
a=rtpmap:112 colibri/90000
a=fmtp:112 profile=1;version=1;level=0
```

In this example, a dynamic payload type 112 is used for colibri data. The 90 kHz RTP timestamp rate is specified in the "a=rtpmap" line after the subtype. In the "a=fmtp:" line, profile 1, version 1, and level 0 (unknown or non-standard level) are specified.

6.3. Offer/Answer Considerations

All parameters are declarative.

7. IANA Considerations

This memo requests that IANA registers video/colibri as specified in Section 7.1. The media type is also requested to be added to the IANA registry for "RTP Payload Format MIME types" (http://www.iana.org/assignments/rtp-parameters).

8. Security Considerations

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]. However, as "Securing the RTP Protocol 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 with 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. The rest of this security consideration section discusses the security impacting properties of the payload format itself.

This RTP payload format and its media decoder do not exhibit any significant non-uniformity in the receiver-side computational complexity for packet processing, and thus are unlikely to pose a denial-of-service threat due to the receipt of pathological data. Nor does the RTP payload format contain any active content.

To avoid buffer overruns when processing these packets the receiver MUST consider both the reported fragment length and the actual received size of a packet containing slice data.

In some cases the transmitter may need to decode fixed length coded headers in order to extract some data from the Colibri bitstream before assembling packets. This process is potentially subject to buffer overruns if not implemented carefully.

9. RFC Editor Considerations

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

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

10. References

10.1. Normative References

[ITU709] Recommendation ITU-R BT.709: Parameter values for high definition television systems for production and international programme exchange.

[ITU601] Recommendation ITU-R BT.601: Studio encoding parameters of digital television for standard 4:3 and wide screen 16:9 aspect ratios.

[ITU2020] Recommendation ITU-R BT.2020: Parameter values for ultra-high definition television systems for production and international programme exchange.

[ITU2100] Recommendation ITU-R BT.2100: Image parameter values for high dynamic range television for use in production and international programme exchange.

[SMPTE428] SMPTE ST428-1:2006 - SMPTE Standard - D-Cinema Distribution Master Image Characteristics.

[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>.

[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>.

[RFC4855] Casner, S., "Media Type Registration of RTP Payload Formats", RFC 4855, DOI 10.17487/RFC4855, February 2007, <https://www.rfc-editor.org/info/rfc4855>.

[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>.

[RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <https://www.rfc-editor.org/info/rfc8085>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[COLIBRI] ?????

10.2. Informative References

[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>.

[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>.

[RFC5109] Li, A., Ed., "RTP Payload Format for Generic Forward Error Correction", RFC 5109, DOI 10.17487/RFC5109, December 2007, <https://www.rfc-editor.org/info/rfc5109>.

[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>.

Author's Address

Luc Ploumhans
Silex Insight

Email: luc.ploumhans@silexinsight.com