

RTP Payload Format for 1-D Interleaved Parity
Forward Error Correction (FEC)

Abstract

This document defines a new RTP payload format for the Forward Error Correction (FEC) that is generated by the 1-D interleaved parity code from a source media encapsulated in RTP. The 1-D interleaved parity code is a systematic code, where a number of repair symbols are generated from a set of source symbols and sent in a repair flow separate from the source flow that carries the source symbols. The 1-D interleaved parity code offers a good protection against bursty packet losses at a cost of reasonable complexity. The new payload format defined in this document should only be used (with some exceptions) as a part of the Digital Video Broadcasting-IPTV (DVB-IPTV) Application-layer FEC specification.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6015>.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
1.1. Use Cases	6
1.2. Overhead Computation	8
1.3. Relation to Existing Specifications	8
1.3.1. RFCs 2733 and 3009	8
1.3.2. SMPTE 2022-1	8
1.3.3. ETSI TS 102 034	9
1.4. Scope of the Payload Format	10
2. Requirements Notation	10
3. Definitions, Notations, and Abbreviations	10
3.1. Definitions	10
3.2. Notations	11
4. Packet Formats	11
4.1. Source Packets	11
4.2. Repair Packets	11
5. Payload Format Parameters	15
5.1. Media Type Registration	15
5.1.1. Registration of audio/ld-interleaved-parityfec	15
5.1.2. Registration of video/ld-interleaved-parityfec	16
5.1.3. Registration of text/ld-interleaved-parityfec	18
5.1.4. Registration of application/ld-interleaved-parityfec	19
5.2. Mapping to SDP Parameters	20
5.2.1. Offer-Answer Model Considerations	21
5.2.2. Declarative Considerations	22
6. Protection and Recovery Procedures	22
6.1. Overview	22
6.2. Repair Packet Construction	22
6.3. Source Packet Reconstruction	24
6.3.1. Associating the Source and Repair Packets	25
6.3.2. Recovering the RTP Header and Payload	25
7. Session Description Protocol (SDP) Signaling	27
8. Congestion Control Considerations	27
9. Security Considerations	28
10. IANA Considerations	29
11. Acknowledgments	29
12. References	29
12.1. Normative References	29
12.2. Informative References	30

1. Introduction

This document extends the Forward Error Correction (FEC) header defined in [RFC2733] and uses this new FEC header for the FEC that is generated by the 1-D interleaved parity code from a source media encapsulated in RTP [RFC3550]. The resulting new RTP payload format is registered by this document.

The type of the source media protected by the 1-D interleaved parity code can be audio, video, text, or application. The FEC data are generated according to the media type parameters that are communicated through out-of-band means. The associations/relationships between the source and repair flows are also communicated through out-of-band means.

The 1-D interleaved parity FEC uses the exclusive OR (XOR) operation to generate the repair symbols. In a nutshell, the following steps take place:

1. The sender determines a set of source packets to be protected together based on the media type parameters.
2. The sender applies the XOR operation on the source symbols to generate the required number of repair symbols.
3. The sender packetizes the repair symbols and sends the repair packet(s) along with the source packets to the receiver(s) (in different flows). The repair packets may be sent proactively or on demand.

Note that the source and repair packets belong to different source and repair flows, and the sender needs to provide a way for the receivers to demultiplex them, even in the case in which they are sent in the same transport flow (i.e., same source/destination address/port with UDP). This is required to offer backward compatibility (see Section 4). At the receiver side, if all of the source packets are successfully received, there is no need for FEC recovery and the repair packets are discarded. However, if there are missing source packets, the repair packets can be used to recover the missing information. Block diagrams for the systematic parity FEC encoder and decoder are sketched in Figures 1 and 2, respectively.

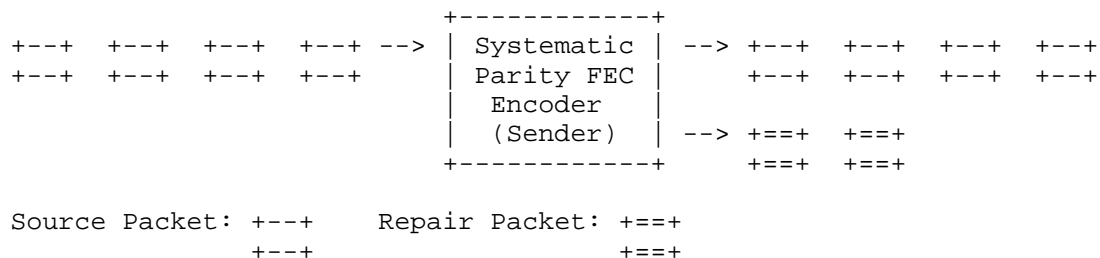


Figure 1: Block diagram for systematic parity FEC encoder

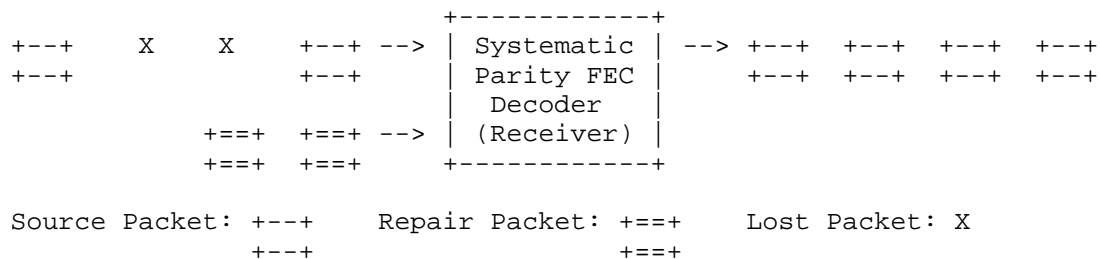


Figure 2: Block diagram for systematic parity FEC decoder

Suppose that we have a group of $D \times L$ source packets that have sequence numbers starting from 1 running to $D \times L$. If we apply the XOR operation to the group of the source packets whose sequence numbers are L apart from each other as sketched in Figure 3, we generate L repair packets. This process is referred to as 1-D interleaved FEC protection, and the resulting L repair packets are referred to as interleaved (or column) FEC packets.




Figure 3: Generating interleaved (column) FEC packets

In Figure 3, S_n and C_m denote the source packet with a sequence number n and the interleaved (column) FEC packet with a sequence number m, respectively.

1.1. Use Cases

We generate one interleaved FEC packet out of D non-consecutive source packets. This repair packet can provide a full recovery of the missing information if there is only one packet missing among the corresponding source packets. This implies that 1-D interleaved FEC protection performs well under bursty loss conditions provided that a large enough value is chosen for L, i.e., L packet duration should not be shorter than the duration of the burst that is intended to be repaired.

For example, consider the scenario depicted in Figure 4 in which the sender generates interleaved FEC packets and a bursty loss hits the source packets. Since the number of columns is larger than the number of packets lost due to the bursty loss, the repair operation succeeds.

```

+----+
| 1 |   X       X       X
+----+

+----+ +----+ +----+ +----+
| 5 | | 6 | | 7 | | 8 |
+----+ +----+ +----+ +----+

+----+ +----+ +----+ +----+
| 9 | | 10| | 11| | 12|
+----+ +----+ +----+ +----+

+====+ +====+ +====+ +====+
|C_1| |C_2| |C_3| |C_4|
+====+ +====+ +====+ +====+

```

Figure 4: Example scenario where 1-D interleaved FEC protection succeeds error recovery

The sender may generate interleaved FEC packets to combat the bursty packet losses. However, two or more random packet losses may hit the source and repair packets in the same column. In that case, the repair operation fails. This is illustrated in Figure 5. Note that it is possible that two or more bursty losses may occur in the same source block, in which case interleaved FEC packets may still fail to recover the lost data.

```

+----+       +----+ +----+
| 1 |   X   | 3 | | 4 |
+----+       +----+ +----+

+----+       +----+ +----+
| 5 |   X   | 7 | | 8 |
+----+       +----+ +----+

+----+ +----+ +----+ +----+
| 9 | | 10| | 11| | 12|
+----+ +----+ +----+ +----+

+====+ +====+ +====+ +====+
|C_1| |C_2| |C_3| |C_4|
+====+ +====+ +====+ +====+

```

Figure 5: Example scenario where 1-D interleaved FEC protection fails error recovery

1.2. Overhead Computation

The overhead is defined as the ratio of the number of bytes that belong to the repair packets to the number of bytes that belong to the protected source packets.

Assuming that each repair packet carries an equal number of bytes carried by a source packet and ignoring the size of the FEC header, we can compute the overhead as follows:

$$\text{Overhead} = 1/D$$

where D is the number of rows in the source block.

1.3. Relation to Existing Specifications

This section discusses the relation of the current specification to other existing specifications.

1.3.1. RFCs 2733 and 3009

The current specification extends the FEC header defined in [RFC2733] and registers a new RTP payload format. This new payload format is not backward compatible with the payload format that was registered by [RFC3009].

1.3.2. SMPTE 2022-1

In 2007, the Society of Motion Picture and Television Engineers (SMPTE) - Technology Committee N26 on File Management and Networking Technology - decided to revise the Pro-MPEG Code of Practice (CoP) #3 Release 2 specification (initially produced by the Pro-MPEG Forum in 2004), which discussed several aspects of the transmission of MPEG-2 transport streams over IP networks. The new SMPTE specification is referred to as [SMPTE2022-1].

The Pro-MPEG CoP #3 Release 2 document was originally based on [RFC2733]. SMPTE revised the document by extending the FEC header proposed in [RFC2733] (by setting the E bit). This extended header offers some improvements.

For example, instead of utilizing the bitmap field used in [RFC2733], [SMPTE2022-1] introduces separate fields to convey the number of rows (D) and columns (L) of the source block as well as the type of the repair packet (i.e., whether the repair packet is an interleaved FEC packet computed over a column or a non-interleaved FEC packet computed over a row). These fields, plus the base sequence number, allow the receiver side to establish associations between the source

and repair packets. Note that although the bitmap field is not utilized, the FEC header of [SMPTE2022-1] inherently carries over the bitmap field from [RFC2733].

On the other hand, some parts of [SMPTE2022-1] are not in compliance with RTP [RFC3550]. For example, [SMPTE2022-1] sets the Synchronization Source (SSRC) field to zero and does not use the timestamp field in the RTP headers of the repair packets (receivers ignore the timestamps of the repair packets). Furthermore, [SMPTE2022-1] also sets the CSRC Count (CC) field in the RTP header to zero and does not allow any Contributing Source (CSRC) entry in the RTP header.

The current document adopts the extended FEC header of [SMPTE2022-1] and registers a new RTP payload format. At the same time, this document fixes the parts of [SMPTE2022-1] that are not compliant with RTP [RFC3550], except the one discussed below.

The baseline header format first proposed in [RFC2733] does not have fields to protect the P and X bits and the CC fields of the source packets associated with a repair packet. Rather, the P bit, X bit, and CC field in the RTP header of the repair packet are used to protect those bits and fields. This, however, may sometimes result in failures when doing the RTP header validity checks as specified in [RFC3550]. While this behavior has been fixed in [RFC5109], which obsoleted [RFC2733], the RTP payload format defined in this document still allows this behavior for legacy purposes. Implementations following this specification must be aware of this potential issue when RTP header validity checks are applied.

1.3.3. ETSI TS 102 034

In 2009, the Digital Video Broadcasting (DVB) consortium published a technical specification [ETSI-TS-102-034] through the European Telecommunications Standards Institute (ETSI). This specification covers several areas related to the transmission of MPEG-2 transport stream-based services over IP networks.

Annex E of [ETSI-TS-102-034] defines an optional protocol for Application-layer FEC (AL-FEC) protection of streaming media for DVB-IP services carried over RTP [RFC3550] transport. The DVB-IPTV AL-FEC protocol uses two layers for protection: a base layer that is produced by a packet-based interleaved parity code, and an enhancement layer that is produced by a Raptor code [DVB-AL-FEC]. While the use of the enhancement layer is optional, the use of the base layer is mandatory wherever AL-FEC is used. The DVB-IPTV AL-FEC protocol is also described in [DVB-AL-FEC].

The interleaved parity code that is used in the base layer is a subset of [SMPTE2022-1]. In particular, the AL-FEC base layer uses only the 1-D interleaved FEC protection from [SMPTE2022-1]. The new RTP payload format that is defined and registered in this document (with some exceptions listed in [DVB-AL-FEC]) is used as the AL-FEC base layer.

1.4. Scope of the Payload Format

The payload format specified in this document must only be used in legacy applications where the limitations explained in Section 1.3.2 are known not to impact any system components or other RTP elements. Whenever possible, a payload format that is fully compliant with [RFC3550], such as [RFC5109] or other newer payload formats, must be used.

2. Requirements Notation

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 [RFC2119].

3. Definitions, Notations, and Abbreviations

The definitions and notations commonly used in this document are summarized in this section.

3.1. Definitions

This document uses the following definitions:

Source Flow: The packet flow(s) carrying the source data to which FEC protection is to be applied.

Repair Flow: The packet flow(s) carrying the repair data.

Symbol: A unit of data. Its size, in bytes, is referred to as the symbol size.

Source Symbol: The smallest unit of data used during the encoding process.

Repair Symbol: Repair symbols are generated from the source symbols.

Source Packet: Data packets that contain only source symbols.

Repair Packet: Data packets that contain only repair symbols.

Source Block: A block of source symbols that are considered together in the encoding process.

3.2. Notations

- o L: Number of columns of the source block.
- o D: Number of rows of the source block.

4. Packet Formats

This section defines the formats of the source and repair packets.

4.1. Source Packets

The source packets need to contain information that identifies the source block and the position within the source block occupied by the packet. Since the source packets that are carried within an RTP stream already contain unique sequence numbers in their RTP headers [RFC3550], we can identify the source packets in a straightforward manner, and there is no need to append additional field(s). The primary advantage of not modifying the source packets in any way is that it provides backward compatibility for the receivers that do not support FEC at all. In multicast scenarios, this backward compatibility becomes quite useful as it allows the non-FEC-capable and FEC-capable receivers to receive and interpret the same source packets sent in the same multicast session.

4.2. Repair Packets

The repair packets MUST contain information that identifies the source block to which they pertain and the relationship between the contained repair symbols and the original source block. For this purpose, we use the RTP header of the repair packets as well as another header within the RTP payload, which we refer to as the FEC header, as shown in Figure 6.




Figure 6: Format of repair packets

The RTP header is formatted according to [RFC3550] with some further clarifications listed below:

- o Version: The version field is set to 2.
- o Padding (P) Bit: This bit is equal to the XOR sum of the corresponding P bits from the RTP headers of the source packets protected by this repair packet. However, padding octets are never present in a repair packet, independent of the value of the P bit.
- o Extension (X) Bit: This bit is equal to the XOR sum of the corresponding X bits from the RTP headers of the source packets protected by this repair packet. However, an RTP header extension is never present in a repair packet, independent of the value of the X bit.
- o CSRC Count (CC): This field is equal to the XOR sum of the corresponding CC values from the RTP headers of the source packets protected by this repair packet. However, a CSRC list is never present in a repair packet, independent of the value of the CC field.
- o Marker (M) Bit: This bit is equal to the XOR sum of the corresponding M bits from the RTP headers of the source packets protected by this repair packet.
- o Payload Type: The (dynamic) payload type for the repair packets is determined through out-of-band means. Note that this document registers a new payload format for the repair packets (refer to Section 5 for details). According to [RFC3550], an RTP receiver that cannot recognize a payload type must discard it. This action provides backward compatibility. The FEC mechanisms can then be used in a multicast group with mixed FEC-capable and non-FEC-

capable receivers. If a non-FEC-capable receiver receives a repair packet, it will not recognize the payload type, and hence, discards the repair packet.

- o Sequence Number (SN): The sequence number has the standard definition. It MUST be one higher than the sequence number in the previously transmitted repair packet. The initial value of the sequence number SHOULD be random (unpredictable) [RFC3550].
- o Timestamp (TS): The timestamp SHALL be set to a time corresponding to the repair packet's transmission time. Note that the timestamp value has no use in the actual FEC protection process and is usually useful for jitter calculations.
- o Synchronization Source (SSRC): The SSRC value SHALL be randomly assigned as suggested by [RFC3550]. This allows the sender to multiplex the source and repair flows on the same port or multiplex multiple repair flows on a single port. The repair flows SHOULD use the RTP Control Protocol (RTCP) CNAME field to associate themselves with the source flow.

In some networks, the RTP Source (which produces the source packets) and the FEC Source (which generates the repair packets from the source packets) may not be the same host. In such scenarios, using the same CNAME for the source and repair flows means that the RTP Source and the FEC Source MUST share the same CNAME (for this specific source-repair flow association). A common CNAME may be produced based on an algorithm that is known both to the RTP and FEC Source. This usage is compliant with [RFC3550].

Note that due to the randomness of the SSRC assignments, there is a possibility of SSRC collision. In such cases, the collisions MUST be resolved as described in [RFC3550].

Note that the P bit, X bit, CC field, and M bit of the source packets are protected by the corresponding bits/fields in the RTP header of the repair packet. On the other hand, the payload of a repair packet protects the concatenation of (if present) the CSRC list, RTP extension, payload, and padding of the source RTP packets associated with this repair packet.

The FEC header is 16 octets. The format of the FEC header is shown in Figure 7.

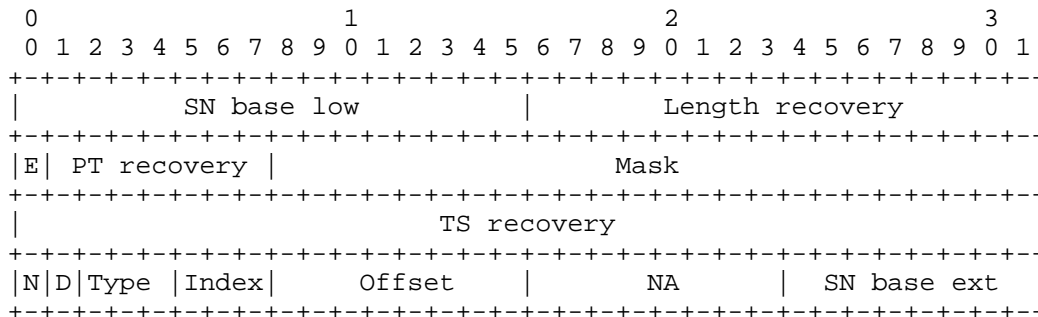


Figure 7: Format of the FEC header

The FEC header consists of the following fields:

- o The SN base low field is used to indicate the lowest sequence number, taking wraparound into account, of those source packets protected by this repair packet.
- o The Length recovery field is used to determine the length of any recovered packets.
- o The E bit is the extension flag introduced in [RFC2733] and used to extend the [RFC2733] FEC header.
- o The PT recovery field is used to determine the payload type of the recovered packets.
- o The Mask field is not used.
- o The TS recovery field is used to determine the timestamp of the recovered packets.
- o The N bit is the extension flag that is reserved for future use.
- o The D bit is not used.
- o The Type field indicates the type of the error-correcting code used. This document defines only one error-correcting code.
- o The Index field is not used.
- o The Offset and NA fields are used to indicate the number of columns (L) and rows (D) of the source block, respectively.
- o The SN base ext field is not used.

The details on setting the fields in the FEC header are provided in Section 6.2.

It should be noted that a Mask-based approach (similar to the one specified in [RFC2733]) may not be very efficient to indicate which source packets in the current source block are associated with a given repair packet. In particular, for the applications that would like to use large source block sizes, the size of the Mask that is required to describe the source-repair packet associations may be prohibitively large. Instead, a systematized approach is inherently more efficient.

5. Payload Format Parameters

This section provides the media subtype registration for the 1-D interleaved parity FEC. The parameters that are required to configure the FEC encoding and decoding operations are also defined in this section.

5.1. Media Type Registration

This registration is done using the template defined in [RFC4288] and following the guidance provided in [RFC4855].

5.1.1. Registration of audio/1d-interleaved-parityfec

Type name: audio

Subtype name: 1d-interleaved-parityfec

Required parameters:

- o rate: The RTP timestamp (clock) rate in Hz. The (integer) rate SHALL be larger than 1000 to provide sufficient resolution to RTCP operations. However, it is RECOMMENDED to select the rate that matches the rate of the protected source RTP stream.
- o L: Number of columns of the source block. L is a positive integer that is less than or equal to 255.
- o D: Number of rows of the source block. D is a positive integer that is less than or equal to 255.
- o repair-window: The time that spans the FEC block (i.e., source packets and the corresponding repair packets). An FEC encoder processes a block of source packets and generates a number of repair packets, which are then transmitted within a certain duration not larger than the value of the repair window. At the

receiver side, the FEC decoder should wait at least for the duration of the repair window after getting the first packet in an FEC block to allow all the repair packets to arrive (the waiting time can be adjusted if there are missing packets at the beginning of the FEC block). The FEC decoder can start decoding the already received packets sooner; however, it SHOULD NOT register an FEC decoding failure until it waits at least for the repair-window duration. The size of the repair window is specified in microseconds.

Optional parameters: None.

Encoding considerations: This media type is framed (see Section 4.8 in the template document [RFC4288]) and contains binary data.

Security considerations: See Section 9 of [RFC6015].

Interoperability considerations: None.

Published specification: [RFC6015].

Applications that use this media type: Multimedia applications that want to improve resiliency against packet loss by sending redundant data in addition to the source media.

Additional information: None.

Person & email address to contact for further information: Ali Begen <abegen@cisco.com> and the IETF Audio/Video Transport Working Group.

Intended usage: COMMON.

Restriction on usage: This media type depends on RTP framing, and hence, is only defined for transport via RTP [RFC3550].

Author: Ali Begen <abegen@cisco.com>.

Change controller: IETF Audio/Video Transport Working Group delegated from the IESG.

5.1.2. Registration of video/ld-interleaved-parityfec

Type name: video

Subtype name: ld-interleaved-parityfec

Required parameters:

- o rate: The RTP timestamp (clock) rate in Hz. The (integer) rate SHALL be larger than 1000 to provide sufficient resolution to RTCP operations. However, it is RECOMMENDED to select the rate that matches the rate of the protected source RTP stream.
- o L: Number of columns of the source block. L is a positive integer that is less than or equal to 255.
- o D: Number of rows of the source block. D is a positive integer that is less than or equal to 255.
- o repair-window: The time that spans the FEC block (i.e., source packets and the corresponding repair packets). An FEC encoder processes a block of source packets and generates a number of repair packets, which are then transmitted within a certain duration not larger than the value of the repair window. At the receiver side, the FEC decoder should wait at least for the duration of the repair window after getting the first packet in an FEC block to allow all the repair packets to arrive (the waiting time can be adjusted if there are missing packets at the beginning of the FEC block). The FEC decoder can start decoding the already received packets sooner; however, it SHOULD NOT register an FEC decoding failure until it waits at least for the repair-window duration. The size of the repair window is specified in microseconds.

Optional parameters: None.

Encoding considerations: This media type is framed (see Section 4.8 in the template document [RFC4288]) and contains binary data.

Security considerations: See Section 9 of [RFC6015].

Interoperability considerations: None.

Published specification: [RFC6015].

Applications that use this media type: Multimedia applications that want to improve resiliency against packet loss by sending redundant data in addition to the source media.

Additional information: None.

Person & email address to contact for further information: Ali Begen <abegen@cisco.com> and the IETF Audio/Video Transport Working Group.

Intended usage: COMMON.

Restriction on usage: This media type depends on RTP framing, and hence, is only defined for transport via RTP [RFC3550].

Author: Ali Begen <abegen@cisco.com>.

Change controller: IETF Audio/Video Transport Working Group delegated from the IESG.

5.1.3. Registration of text/ld-interleaved-parityfec

Type name: text

Subtype name: ld-interleaved-parityfec

Required parameters:

- o rate: The RTP timestamp (clock) rate in Hz. The (integer) rate SHALL be larger than 1000 to provide sufficient resolution to RTCP operations. However, it is RECOMMENDED to select the rate that matches the rate of the protected source RTP stream.
- o L: Number of columns of the source block. L is a positive integer that is less than or equal to 255.
- o D: Number of rows of the source block. D is a positive integer that is less than or equal to 255.
- o repair-window: The time that spans the FEC block (i.e., source packets and the corresponding repair packets). An FEC encoder processes a block of source packets and generates a number of repair packets, which are then transmitted within a certain duration not larger than the value of the repair window. At the receiver side, the FEC decoder should wait at least for the duration of the repair window after getting the first packet in an FEC block to allow all the repair packets to arrive (the waiting time can be adjusted if there are missing packets at the beginning of the FEC block). The FEC decoder can start decoding the already received packets sooner; however, it SHOULD NOT register an FEC decoding failure until it waits at least for the repair-window duration. The size of the repair window is specified in microseconds.

Optional parameters: None.

Encoding considerations: This media type is framed (see Section 4.8 in the template document [RFC4288]) and contains binary data.

Security considerations: See Section 9 of [RFC6015].

Interoperability considerations: None.

Published specification: [RFC6015].

Applications that use this media type: Multimedia applications that want to improve resiliency against packet loss by sending redundant data in addition to the source media.

Additional information: None.

Person & email address to contact for further information: Ali Begen <abegen@cisco.com> and the IETF Audio/Video Transport Working Group.

Intended usage: COMMON.

Restriction on usage: This media type depends on RTP framing, and hence, is only defined for transport via RTP [RFC3550].

Author: Ali Begen <abegen@cisco.com>.

Change controller: IETF Audio/Video Transport Working Group delegated from the IESG.

5.1.4. Registration of application/ld-interleaved-parityfec

Type name: application

Subtype name: ld-interleaved-parityfec

Required parameters:

- o rate: The RTP timestamp (clock) rate in Hz. The (integer) rate SHALL be larger than 1000 to provide sufficient resolution to RTCP operations. However, it is RECOMMENDED to select the rate that matches the rate of the protected source RTP stream.
- o L: Number of columns of the source block. L is a positive integer that is less than or equal to 255.
- o D: Number of rows of the source block. D is a positive integer that is less than or equal to 255.
- o repair-window: The time that spans the FEC block (i.e., source packets and the corresponding repair packets). An FEC encoder processes a block of source packets and generates a number of repair packets, which are then transmitted within a certain duration not larger than the value of the repair window. At the receiver side, the FEC decoder should wait at least for the

duration of the repair window after getting the first packet in an FEC block to allow all the repair packets to arrive (the waiting time can be adjusted if there are missing packets at the beginning of the FEC block). The FEC decoder can start decoding the already received packets sooner; however, it SHOULD NOT register an FEC decoding failure until it waits at least for the repair-window duration. The size of the repair window is specified in microseconds.

Optional parameters: None.

Encoding considerations: This media type is framed (see Section 4.8 in the template document [RFC4288]) and contains binary data.

Security considerations: See Section 9 of [RFC6015].

Interoperability considerations: None.

Published specification: [RFC6015].

Applications that use this media type: Multimedia applications that want to improve resiliency against packet loss by sending redundant data in addition to the source media.

Additional information: None.

Person & email address to contact for further information: Ali Begen <abegen@cisco.com> and the IETF Audio/Video Transport Working Group.

Intended usage: COMMON.

Restriction on usage: This media type depends on RTP framing, and hence, is only defined for transport via RTP [RFC3550].

Author: Ali Begen <abegen@cisco.com>.

Change controller: IETF Audio/Video Transport Working Group delegated from the IESG.

5.2. Mapping to SDP Parameters

Applications that use RTP transport commonly use Session Description Protocol (SDP) [RFC4566] to describe their RTP sessions. The information that is used to specify the media types in an RTP session has specific mappings to the fields in an SDP description. In this section, we provide these mappings for the media subtype registered by this document ("ld-interleaved-parityfec"). Note that if an application does not use SDP to describe the RTP sessions, an

appropriate mapping must be defined and used to specify the media types and their parameters for the control/description protocol employed by the application.

The mapping of the media type specification for "ld-interleaved-parityfec" and its parameters in SDP is as follows:

- o The media type (e.g., "application") goes into the "m=" line as the media name.
- o The media subtype ("ld-interleaved-parityfec") goes into the "a=rtpmap" line as the encoding name. The RTP clock rate parameter ("rate") also goes into the "a=rtpmap" line as the clock rate.
- o The remaining required payload-format-specific parameters go into the "a=fmtp" line by copying them directly from the media type string as a semicolon-separated list of parameter=value pairs.

SDP examples are provided in Section 7.

5.2.1. Offer-Answer Model Considerations

When offering 1-D interleaved parity FEC over RTP using SDP in an Offer/Answer model [RFC3264], the following considerations apply:

- o Each combination of the L and D parameters produces a different FEC data and is not compatible with any other combination. A sender application may desire to offer multiple offers with different sets of L and D values as long as the parameter values are valid. The receiver SHOULD normally choose the offer that has a sufficient amount of interleaving. If multiple such offers exist, the receiver may choose the offer that has the lowest overhead or the one that requires the smallest amount of buffering. The selection depends on the application requirements.
- o The value for the repair-window parameter depends on the L and D values and cannot be chosen arbitrarily. More specifically, L and D values determine the lower limit for the repair-window size. The upper limit of the repair-window size does not depend on the L and D values.
- o Although combinations with the same L and D values but with different repair-window sizes produce the same FEC data, such combinations are still considered different offers. The size of the repair-window is related to the maximum delay between the

transmission of a source packet and the associated repair packet. This directly impacts the buffering requirement on the receiver side, and the receiver must consider this when choosing an offer.

- o There are no optional format parameters defined for this payload. Any unknown option in the offer MUST be ignored and deleted from the answer. If FEC is not desired by the receiver, it can be deleted from the answer.

5.2.2. Declarative Considerations

In declarative usage, like SDP in the Real-time Streaming Protocol (RTSP) [RFC2326] or the Session Announcement Protocol (SAP) [RFC2974], the following considerations apply:

- o The payload format configuration parameters are all declarative and a participant MUST use the configuration that is provided for the session.
- o More than one configuration may be provided (if desired) by declaring multiple RTP payload types. In that case, the receivers should choose the repair flow that is best for them.

6. Protection and Recovery Procedures

This section provides a complete specification of the 1-D interleaved parity code and its RTP payload format.

6.1. Overview

The following sections specify the steps involved in generating the repair packets and reconstructing the missing source packets from the repair packets.

6.2. Repair Packet Construction

The RTP header of a repair packet is formed based on the guidelines given in Section 4.2.

The FEC header includes 16 octets. It is constructed by applying the XOR operation on the bit strings that are generated from the individual source packets protected by this particular repair packet. The set of the source packets that are associated with a given repair packet can be computed by the formula given in Section 6.3.1.

The bit string is formed for each source packet by concatenating the following fields together in the order specified:

- o Padding bit (1 bit) (This is the most significant bit of the bit string.)
- o Extension bit (1 bit)
- o CC field (4 bits)
- o Marker bit (1 bit)
- o PT field (7 bits)
- o Timestamp (32 bits)
- o Unsigned network-ordered 16-bit representation of the source packet length in bytes minus 12 (for the fixed RTP header), i.e., the sum of the lengths of all the following if present: the CSRC list, header extension, RTP payload, and RTP padding (16 bits).
- o If CC is nonzero, the CSRC list (variable length)
- o If X is 1, the header extension (variable length)
- o Payload (variable length)
- o Padding, if present (variable length)

Note that if the lengths of the source packets are not equal, each shorter packet MUST be padded to the length of the longest packet by adding octet(s) of 0 at the end. Due to this possible padding and mandatory FEC header, a repair packet has a larger size than the source packets it protects. This may cause problems if the resulting repair packet size exceeds the Maximum Transmission Unit (MTU) size of the path over which the repair flow is sent.

By applying the parity operation on the bit strings produced from the source packets, we generate the FEC bit string. Some parts of the RTP header and the FEC header of the repair packet are generated from the FEC bit string as follows:

- o The first (most significant) bit in the FEC bit string is written into the Padding bit in the RTP header of the repair packet.
- o The next bit in the FEC bit string is written into the Extension bit in the RTP header of the repair packet.
- o The next 4 bits of the FEC bit string are written into the CC field in the RTP header of the repair packet.

- o The next bit of the FEC bit string is written into the Marker bit in the RTP header of the repair packet.
- o The next 7 bits of the FEC bit string are written into the PT recovery field in the FEC header.
- o The next 32 bits of the FEC bit string are written into the TS recovery field in the FEC header.
- o The next 16 bits are written into the Length recovery field in the FEC header. This allows the FEC procedure to be applied even when the lengths of the protected source packets are not identical.
- o The remaining bits are set to be the payload of the repair packet.

The remaining parts of the FEC header are set as follows:

- o The SN base low field MUST be set to the lowest sequence number, taking wraparound into account, of those source packets protected by this repair packet.
- o The E bit MUST be set to 1 to extend the [RFC2733] FEC header.
- o The Mask field SHALL be set to 0 and ignored by the receiver.
- o The N bit SHALL be set to 0 and ignored by the receiver.
- o The D bit SHALL be set to 0 and ignored by the receiver.
- o The Type field MUST be set to 0 and ignored by the receiver.
- o The Index field SHALL be set to 0 and ignored by the receiver.
- o The Offset field MUST be set to the number of columns of the source block (L).
- o The NA field MUST be set to the number of rows of the source block (D).
- o The SN base ext field SHALL be set to 0 and ignored by the receiver.

6.3. Source Packet Reconstruction

This section describes the recovery procedures that are required to reconstruct the missing source packets. The recovery process has two steps. In the first step, the FEC decoder determines which source

and repair packets should be used in order to recover a missing packet. In the second step, the decoder recovers the missing packet, which consists of an RTP header and RTP payload.

In the following, we describe the RECOMMENDED algorithms for the first and second steps. Based on the implementation, different algorithms MAY be adopted. However, the end result MUST be identical to the one produced by the algorithms described below.

6.3.1. Associating the Source and Repair Packets

The first step is to associate the source and repair packets. The SN base low field in the FEC header shows the lowest sequence number of the source packets that form the particular column. In addition, the information of how many source packets are available in each column and row is available from the media type parameters specified in the SDP description. This set of information uniquely identifies all of the source packets associated with a given repair packet.

Mathematically, for any received repair packet, p^* , we can determine the sequence numbers of the source packets that are protected by this repair packet as follows:

$$p^*_{\text{snb}} + i * L \text{ (modulo 65536)}$$

where p^*_{snb} denotes the value in the SN base low field of the FEC header of the p^* , L is the number of columns of the source block and

$$0 \leq i < D$$

where D is the number of rows of the source block.

We denote the set of the source packets associated with repair packet p^* by set $T(p^*)$. Note that in a source block whose size is L columns by D rows, set T includes D source packets. Recall that 1- D interleaved FEC protection can fully recover the missing information if there is only one source packet missing in set T . If the repair packet that protects the source packets in set T is missing, or the repair packet is available but two or more source packets are missing, then missing source packets in set T cannot be recovered by 1- D interleaved FEC protection.

6.3.2. Recovering the RTP Header and Payload

For a given set T , the procedure for the recovery of the RTP header of the missing packet, whose sequence number is denoted by $SEQNUM$, is as follows:

1. For each of the source packets that are successfully received in set T, compute the bit string as described in Section 6.2.
2. For the repair packet associated with set T, compute the bit string in the same fashion except use the PT recovery field instead of the PT field and TS recovery field instead of the Timestamp field, and set the CSRC list, header extension and padding to null regardless of the values of the CC field, X bit, and P bit.
3. If any of the bit strings generated from the source packets are shorter than the bit string generated from the repair packet, pad them to be the same length as the bit string generated from the repair packet. For padding, the padding of octet 0 MUST be added at the end of the bit string.
4. Calculate the recovered bit string as the XOR of the bit strings generated from all source packets in set T and the FEC bit string generated from the repair packet associated with set T.
5. Create a new packet with the standard 12-byte RTP header and no payload.
6. Set the version of the new packet to 2.
7. Set the Padding bit in the new packet to the first bit in the recovered bit string.
8. Set the Extension bit in the new packet to the next bit in the recovered bit string.
9. Set the CC field to the next 4 bits in the recovered bit string.
10. Set the Marker bit in the new packet to the next bit in the recovered bit string.
11. Set the Payload type in the new packet to the next 7 bits in the recovered bit string.
12. Set the SN field in the new packet to SEQNUM.
13. Set the TS field in the new packet to the next 32 bits in the recovered bit string.
14. Take the next 16 bits of the recovered bit string and set the new variable Y to whatever unsigned integer this represents (assuming network order). Convert Y to host order and then take Y bytes from the recovered bit string and append them to the new

packet. Y represents the length of the new packet in bytes minus 12 (for the fixed RTP header), i.e., the sum of the lengths of all the following if present: the CSRC list, header extension, RTP payload, and RTP padding.

15. Set the SSRC of the new packet to the SSRC of the source RTP stream.

This procedure completely recovers both the header and payload of an RTP packet.

7. Session Description Protocol (SDP) Signaling

This section provides an SDP [RFC4566] example. The following example uses the FEC grouping semantics [RFC5956].

In this example, we have one source video stream (mid:S1) and one FEC repair stream (mid:R1). We form one FEC group with the "a=group:FEC-FR S1 R1" line. The source and repair streams are sent to the same port on different multicast groups. The repair window is set to 200 ms.

```
v=0
o=ali 1122334455 1122334466 IN IP4 fec.example.com
s=Interleaved Parity FEC Example
t=0 0
a=group:FEC-FR S1 R1
m=video 30000 RTP/AVP 100
c=IN IP4 233.252.0.1/127
a=rtpmap:100 MP2T/90000
a=mid:S1
m=application 30000 RTP/AVP 110
c=IN IP4 233.252.0.2/127
a=rtpmap:110 ld-interleaved-parityfec/90000
a=fmtp:110 L=5; D=10; repair-window=200000
a=mid:R1
```

8. Congestion Control Considerations

FEC is an effective approach to provide applications with resiliency against packet losses. However, in networks where the congestion is a major contributor to the packet loss, the potential impacts of using FEC SHOULD be considered carefully before injecting the repair flows into the network. In particular, in bandwidth-limited networks, FEC repair flows may consume most or all of the available bandwidth and may consequently congest the network. In such cases, the applications MUST NOT arbitrarily increase the amount of FEC

protection since doing so may lead to a congestion collapse. If desired, stronger FEC protection MAY be applied only after the source rate has been reduced.

In a network-friendly implementation, an application SHOULD NOT send/receive FEC repair flows if it knows that sending/receiving those FEC repair flows would not help at all in recovering the missing packets. Such a practice helps reduce the amount of wasted bandwidth. It is RECOMMENDED that the amount of FEC protection is adjusted dynamically based on the packet loss rate observed by the applications.

In multicast scenarios, it may be difficult to optimize the FEC protection per receiver. If there is a large variation among the levels of FEC protection needed by different receivers, it is RECOMMENDED that the sender offers multiple repair flows with different levels of FEC protection and the receivers join the corresponding multicast sessions to receive the repair flow(s) that is best for them.

9. 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.

The main security considerations for the RTP packet carrying the RTP payload format defined within this memo are confidentiality, integrity, and source authenticity. Confidentiality is achieved by encrypting the RTP payload. Altering the FEC packets can have a big impact on the reconstruction operation. An attack that changes some bits in the FEC packets can have a significant effect on the calculation and the recovery of the source packets. For example, changing the length recovery field can result in the recovery of a packet that is too long. Depending on the application, it may be helpful to perform a sanity check on the received source and FEC packets before performing the recovery operation and to determine the validity of the recovered packets before using them.

The integrity of the RTP packets is achieved through a suitable cryptographic integrity protection mechanism. Such a cryptographic system may also allow the authentication of the source of the payload. A suitable security mechanism for this RTP payload format should provide source authentication capable of determining if an RTP packet is from a member of the RTP session.

Note that the appropriate mechanism to provide security to RTP and payloads following this memo may vary. It is dependent on the application, transport and signaling protocol employed. Therefore, a

single mechanism is not sufficient, although if suitable, using the Secure Real-time Transport Protocol (SRTP) [RFC3711] is RECOMMENDED. Other mechanisms that may be used are IPsec [RFC4301] and Transport Layer Security (TLS) [RFC5246]; other alternatives may exist.

If FEC protection is applied on already encrypted source packets, there is no need for additional encryption. However, if the source packets are encrypted after FEC protection is applied, the FEC packets should be cryptographically as secure as the source packets. Failure to provide an equal level of confidentiality, integrity, and authentication to the FEC packets can compromise the source packets' confidentiality, integrity or authentication since the FEC packets are generated by applying XOR operation across the source packets.

10. IANA Considerations

New media subtypes are subject to IANA registration. For the registration of the payload format and its parameters introduced in this document, refer to Section 5.

11. Acknowledgments

A major part of this document is borrowed from [RFC2733], [RFC5109], and [SMPTE2022-1]. Thus, the author would like to thank the authors and editors of these earlier specifications. The author also thanks Colin Perkins for his constructive suggestions for this document.

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC5956] Begen, A., "Forward Error Correction Grouping Semantics in Session Description Protocol", RFC 5956, September 2010.

- [RFC4288] Freed, N. and J. Klensin, "Media Type Specifications and Registration Procedures", BCP 13, RFC 4288, December 2005.
- [RFC4855] Casner, S., "Media Type Registration of RTP Payload Formats", RFC 4855, February 2007.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.

12.2. Informative References

- [DVB-AL-FEC] Begen, A. and T. Stockhammer, "Guidelines for Implementing DVB-IPTV Application-Layer Hybrid FEC Protection", Work in Progress, December 2009.
- [RFC2733] Rosenberg, J. and H. Schulzrinne, "An RTP Payload Format for Generic Forward Error Correction", RFC 2733, December 1999.
- [RFC3009] Rosenberg, J. and H. Schulzrinne, "Registration of parityfec MIME types", RFC 3009, November 2000.
- [RFC5109] Li, A., "RTP Payload Format for Generic Forward Error Correction", RFC 5109, December 2007.
- [ETSI-TS-102-034] ETSI TS 102 034 V1.4.1, "Transport of MPEG 2 TS Based DVB Services over IP Based Networks", August 2009.
- [RFC2326] Schulzrinne, H., Rao, A., and R. Lanphier, "Real Time Streaming Protocol (RTSP)", RFC 2326, April 1998.
- [RFC2974] Handley, M., Perkins, C., and E. Whelan, "Session Announcement Protocol", RFC 2974, October 2000.
- [SMPTE2022-1] SMPTE 2022-1-2007, "Forward Error Correction for Real-Time Video/Audio Transport over IP Networks", 2007.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC5246]

Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

Author's Address

Ali Begen
Cisco
181 Bay Street
Toronto, ON M5J 2T3
Canada

EMail: abegen@cisco.com