

**The Sliding Window Random Linear Code (RLC) Forward Erasure Correction
(FEC) Scheme for FECFRAME
draft-roca-tsvwg-rlc-fec-scheme-01**

Abstract

This document describes a fully-specified FEC scheme for the Sliding Window Random Linear Codes (RLC) over $GF(2^m)$, where m equals 1 (binary case), 4 or 8, that can be used to protect arbitrary media streams along the lines defined by FECFRAME extended to sliding window codes. These sliding window FEC codes rely on an encoding window that slides over the source symbols, generating new repair symbols whenever needed. Compared to block FEC codes, these sliding window FEC codes offer key advantages with real-time flows in terms of reduced FEC-related latency while often providing improved erasure recovery capabilities.

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 <http://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 December 29, 2017.

Copyright Notice

Copyright (c) 2017 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	3
1.1.	Limits of Block Codes with Real-Time Flows	3
1.2.	Lower Latency and Better Protection of Real-Time Flows with the Sliding Window RLC Codes	3
1.3.	Small Transmission Overheads with the Sliding Window RLC FEC Scheme	4
1.4.	Document Organization	5
2.	Definitions and Abbreviations	5
3.	Procedures	6
3.1.	Parameters Derivation	6
3.2.	ADU, ADUI and Source Symbols Mappings	7
3.3.	Encoding Window Management	9
3.4.	Pseudo-Random Number Generator	9
3.5.	Coding Coefficients Generation Function	10
4.	Sliding Window RLC FEC Scheme for Arbitrary ADU Flows	12
4.1.	Formats and Codes	12
4.1.1.	FEC Framework Configuration Information	12
4.1.2.	Explicit Source FEC Payload ID	13
4.1.3.	Repair FEC Payload ID	13
4.1.4.	Additional Procedures	15
5.	FEC Code Specification	15
5.1.	Encoding Side	15
5.2.	Decoding Side	15
6.	Implementation Status	16
7.	Security Considerations	16
7.1.	Attacks Against the Data Flow	17
7.1.1.	Access to Confidential Content	17
7.1.2.	Content Corruption	17
7.2.	Attacks Against the FEC Parameters	17
7.3.	When Several Source Flows are to be Protected Together .	18
7.4.	Baseline Secure FEC Framework Operation	18
8.	Operations and Management Considerations	18
8.1.	Operational Recommendations: Finite Field Element Size (m Parameter)	18
9.	IANA Considerations	19
10.	Acknowledgments	19
11.	References	19
11.1.	Normative References	19
11.2.	Informative References	20

Roca

Expires December 29, 2017

[Page 2]

Appendix A. Decoding Beyond Maximum Latency Optimization	22
Author's Address	22

[1. Introduction](#)

Application-Level Forward Erasure Correction (AL-FEC) codes are a key element of communication systems. They are used to recover from packet losses (or erasures) during content delivery sessions to a large number of receivers (multicast/broadcast transmissions). This is the case with the FLUTE/ALC protocol [[RFC6726](#)] in case of reliable file transfers over lossy networks, and the FECFRAME protocol for reliable continuous media transfers over lossy networks.

The present document only focusses on the FECFRAME protocol, used in multicast/broadcast delivery mode, with contents that feature stringent real-time constraints: each source packet has a maximum validity period after which it will not be considered by the destination application.

[1.1. Limits of Block Codes with Real-Time Flows](#)

With FECFRAME, there is a single FEC encoding point (either a end-host/server (source) or a middlebox) and a single FEC decoding point (either a end-host (receiver) or middlebox). In this context, currently standardized AL-FEC codes for FECFRAME like Reed-Solomon [[RFC6865](#)], LDPC-Staircase [[RFC6816](#)], or Raptor/RaptorQ, are all linear block codes: they require the data flow to be segmented into blocks of a predefined maximum size. The block size is a balance between robustness (in particular in front of long erasure bursts for which there is an incentive to increase the block size) and maximum decoding latency (for which there is an incentive to decrease the block size). Therefore, with a multicast/broadcast session, the block code is dimensioned by considering the worst communication channel one wants to support, and this choice impacts all receivers, no matter their individual channel quality.

[1.2. Lower Latency and Better Protection of Real-Time Flows with the Sliding Window RLC Codes](#)

This document introduces a fully-specified FEC scheme that follows a totally different approach: the Sliding Window Random Linear Codes (RLC) over $GF(2^m)$, where m equals 1, 4 or 8. This FEC scheme is used to protect arbitrary media streams along the lines defined by FECFRAME extended to sliding window codes [[fecframe-ext](#)]. This FEC scheme is extremely efficient for instance with media that feature real-time constraints sent within a multicast/broadcast session.

The RLC codes belong to the broad class of sliding window AL-FEC codes (A.K.A. convolutional codes). The encoding process is based on an encoding window that slides over the set of source packets (in fact source symbols as we will see in [Section 3.2](#)), and which is either of fixed or variable size (elastic window). Repair packets (symbols) are generated and sent on-the-fly, after computing a random linear combination of the source symbols present in the current encoding window.

At the receiver, a linear system is managed from the set of received source and repair packets. New variables (representing source symbols) and equations (representing the linear combination of each repair symbol received) are added upon receiving new packets. Variables are removed when they are too old with respect to their validity period (real-time constraints), as well as the associated equations they are involved in (Appendix A introduces an optimisation that extends the time a variable is considered in the system). Erased source symbols are then recovered thanks this linear system whenever its rank permits it.

With RLC codes (more generally with sliding window codes), the protection of a multicast/broadcast session also needs to be dimensioned by considering the worst communication channel one wants to support. However the receivers experiencing a good to medium channel quality observe a FEC-related latency close to zero [[Roca16](#)] since an isolated erased source packet is quickly recovered by the following repair packet. On the opposite, with a block code, recovering an isolated erased source packet always requires waiting the end of the block for the first repair packet to arrive. Additionally, under certain situations (e.g., with a limited FEC-related latency budget and with constant bit rate transmissions after FECFRAME encoding), sliding window codes achieve more easily a target transmission quality (e.g., measured by the residual loss after FEC decoding) by sending fewer repair packets (i.e., higher code rate) than block codes.

[1.3.](#) Small Transmission Overheads with the Sliding Window RLC FEC Scheme

The Sliding Window RLC FEC scheme is designed so as to reduce the transmission overhead. The main requirement is that each repair packet header must enable a receiver to reconstruct the list of source symbols and the associated random coefficients used during the encoding process. In order to minimize packet overhead, the set of symbols in the encoding window as well as the set of coefficients over $GF(2^m)$ used in the linear combination are not individually listed in the repair packet header. Instead, each FEC repair packet header contains:

- o the Encoding Symbol Identifier (ESI) of the first source symbol in the encoding window as well as the number of symbols (since this number may vary with a variable size, elastic window). These two pieces of information enable each receiver to easily reconstruct the set of source symbols considered during encoding, the only constraint being that there cannot be any gap;
- o the seed used by a coding coefficients generation function ([Section 3.5](#)). This information enables each receiver to generate the same set of coding coefficients over $GF(2^m)$ as the sender;

Therefore, no matter the number of source symbols present in the encoding window, each FEC repair packet features a fixed 64-bit long header, called Repair FEC Payload ID (Figure 7). Similarly, each FEC source packet features a fixed 32-bit long trailer, called Explicit Source FEC Payload ID (Figure 5), that contains the ESI of the first source symbol (see the ADUI and source symbol mapping, [Section 3.2](#)).

1.4. Document Organization

This fully-specified FEC scheme follows the structure required by [\[RFC6363\]](#), [section 5.6](#). "FEC Scheme Requirements", namely:

3. Procedures: This section describes procedures specific to this FEC scheme, namely: RLC parameters derivation, ADUI and source symbols mapping, pseudo-random number generator, and coding coefficients generation function;
4. Formats and Codes: This section defines the Source FEC Payload ID and Repair FEC Payload ID formats, carrying the signalling information associated to each source or repair symbol. It also defines the FEC Framework Configuration Information (FFCI) carrying signalling information for the session;
5. FEC Code Specification: Finally this section provides the code specification.

2. 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 [\[RFC2119\]](#).

This document uses the following definitions and abbreviations:

$GF(q)$ denotes a finite field (also known as the Galois Field) with q elements. We assume that $q = 2^m$ in this document

m defines the length of the elements in the finite field, in bits.

In this document, m is equal to 1, 4 or 8

ADU: Application Data Unit

ADUI: Application Data Unit Information (includes the F, L and padding fields in addition to the ADU)

E: encoding symbol size (i.e., source or repair symbol), assumed fixed (in bytes)

br_out: transmission bitrate at the output of the FECFRAME sender, assumed fixed (in bits/s)

max_lat: maximum FEC-related latency within FECFRAME (in seconds)

cr: AL-FEC coding rate

plr: packet loss rate on the erasure channel

ew_size: encoding window current size at a sender (in symbols)

ew_max_size: encoding window maximum size at a sender (in symbols)

dw_size: decoding window current size at a receiver (in symbols)

dw_max_size: decoding window maximum size at a receiver (in symbols)

ls_max_size: linear system maximum size (or width) at a receiver (in symbols)

ls_size: linear system current size (or width) at a receiver (in symbols)

PRNG: pseudo-random number generator

pmms_rand(maxv): PRNG defined in [Section 3.4](#) and used in this specification, that returns a new random integer in $[0; \text{maxv}-1]$

3. Procedures

This section introduces the procedures that are used by this FEC scheme.

3.1. Parameters Derivation

The Sliding Window RLC FEC Scheme relies on several key internal parameters:

Maximum FEC-related latency budget, max_lat (in seconds) A source ADU flow can have real-time constraints, and therefore any FECFRAME related operation must take place within the validity period of each ADU. When there are multiple flows with different real-time constraints, we consider the most stringent constraints (see [\[RFC6363\]](#), [Section 10.2](#), item 6, for recommendations when several flows are globally protected). This maximum FEC-related latency accounts for all sources of latency added by FEC encoding (sender) and FEC decoding (receiver). Other sources of latency (e.g., added by network communications) are out of scope and must be considered separately (e.g., they have already been deducted). It can be regarded as the latency budget permitted for all FEC-related operations. This is also an input parameter that enables to derive other internal parameters;

Encoding window current (resp. maximum) size, ew_size (resp. ew_max_size) (in symbols):

these parameters are used by a sender during FEC encoding. More precisely, each repair symbol is a linear combination of the `ew_size` source symbols present in the encoding window when RLC encoding took place. In all situations, we MUST have `ew_size <= ew_max_size`;

Decoding window current (resp. maximum) size, `dw_size` (resp. `dw_max_size`) (in symbols):

these parameters are used by a receiver when managing the linear system used for decoding. `dw_size` is the current size of the decoding window, i.e., the set of received or erased source symbols that are currently part of the linear system. In all situations, we MUST have `dw_size <= dw_max_size`;

In order to comply with the maximum FEC-related latency budget, assuming a constant transmission bitrate at the output of the FECFRAME sender (`br_out`), encoding symbol size (`E`), and code rate (`cr`), we have:

$$dw_max_size = (max_lat * br_out * cr) / (8 * E)$$

This `dw_max_size` defines the maximum delay after which an old source symbol may be recovered: after this delay, this old source symbol will be removed from the decoding window.

It is often good practice to choose:

$$ew_max_size = dw_max_size / 2$$

However any value `ew_max_size < dw_max_size` can be used without impact on the FEC-related latency budget. Finding the optimal value can depend on the erasure channel one wants to support and should be determined after simulations or field trials.

Note that the decoding beyond maximum latency optimisation (Appendix A) enables an old source symbol to be kept in the linear system beyond the FEC-related latency budget, but not delivered to the receiving application. Here we have: `ls_size >= dw_max_size`

3.2. ADU, ADUI and Source Symbols Mappings

An ADU, coming from the application, cannot be mapped to source symbols directly. Indeed, an erased ADU recovered at a receiver must contain enough information to be assigned to the right application flow (UDP port numbers and IP addresses cannot be used to that purpose as they are not protected by FEC encoding). This requires adding the flow identifier to each ADU before doing FEC encoding.

Additionally, since ADUs are of variable size, padding is needed so that each ADU (with its flow identifier) contribute to an integral number of source symbols. This requires adding the original ADU length to each ADU before doing FEC encoding. Because of these requirements, an intermediate format, the ADUI, or ADU Information, is considered [[RFC6363](#)].

For each incoming ADU, an ADUI is created as follows. First of all, 3 bytes are prepended: (Figure 1):

Flow ID (F) (8-bit field): this unsigned byte contains the integer identifier associated to the source ADU flow to which this ADU belongs. It is assumed that a single byte is sufficient, which implies that no more than 256 flows will be protected by a single FECFRAME instance.

Length (L) (16-bit field): this unsigned integer contains the length of this ADU, in network byte order (i.e., big endian). This length is for the ADU itself and does not include the F, L, or Pad fields.

Then, zero padding is added to the ADU if needed:

Padding (Pad) (variable size field): this field contains zero padding to align the F, L, ADU and padding up to a size that is multiple of E bytes (i.e., the source and repair symbol length).

Each ADUI contributes to an integral number of source symbols. The data unit resulting from the ADU and the F, L, and Pad fields is called ADU Information (or ADUI). Since ADUs can be of different size, this is also the case for ADUIs.




Figure 1: ADUI Creation example (here 3 source symbols are created for this ADUI).

Note that neither the initial 3 bytes nor the optional padding are sent over the network. However, they are considered during FEC encoding. It means that a receiver who lost a certain FEC source packet (e.g., the UDP datagram containing this FEC source packet) will be able to recover the ADUI if FEC decoding succeeds. Thanks to the initial 3 bytes, this receiver will get rid of the padding (if any) and identify the corresponding ADU flow.

3.3. Encoding Window Management

Source symbols and the corresponding ADUs are removed from the encoding window:

- o when the sliding encoding window has reached its maximum size, `ew_max_size`. In that case the oldest symbol MUST be removed before adding a new symbol, so that the current encoding window size always remains inferior or equal to the maximum size: `ew_size <= ew_max_size`;
- o when an ADU has reached its maximum validity duration in case of a real-time flow. When this happens, all source symbols corresponding to the ADUI that expired SHOULD be removed from the encoding window;

Source symbols are added to the sliding encoding window each time a new ADU arrives, once the ADU to ADUI and then to source symbols mapping has been performed ([Section 3.2](#)). The current size of the encoding window, `ew_size`, is updated after adding new source symbols. This process may require to remove old source symbols so that: `ew_size <= ew_max_size`.

Note that a FEC codec may feature practical limits in the number of source symbols in the encoding window (e.g., for computational complexity reasons). This factor may further limit the `ew_max_lat` value, in addition to the maximum FEC-related latency budget ([Section 3.1](#)).

3.4. Pseudo-Random Number Generator

The RLC codes rely on the following Pseudo-Random Number Generator (PRNG), identical to the PRNG used with LDPC-Staircase codes ([\[RFC5170\]](#), [section 5.7](#)).

The Park-Miller "minimal standard" PRNG [\[PM88\]](#) MUST be used. It defines a simple multiplicative congruential algorithm: $I_{j+1} = A * I_j \pmod{M}$, with the following choices: $A = 7^{^5} = 16807$ and $M = 2^{^31} - 1 = 2147483647$. A validation criteria of such a PRNG is the following: if seed = 1, then the 10,000th value returned MUST be equal to 1043618065.

Several implementations of this PRNG are known and discussed in the literature. An optimized implementation of this algorithm, using only 32-bit mathematics, and which does not require any division, can be found in [\[rand31pmc\]](#). It uses the Park and Miller algorithm [\[PM88\]](#) with the optimization suggested by D. Carta in [\[CA90\]](#). The history behind this algorithm is detailed in [\[WI08\]](#). Yet, any other implementation of the PRNG algorithm that matches the above

validation criteria, like the ones detailed in [\[PM88\]](#), is appropriate.

This PRNG produces, natively, a 31-bit value between 1 and 0x7FFFFFFE ($2^{31}-2$) inclusive. Since it is desired to scale the pseudo-random number between 0 and maxv-1 inclusive, one must keep the most significant bits of the value returned by the PRNG (the least significant bits are known to be less random, and modulo-based solutions should be avoided [\[PTVF92\]](#)). The following algorithm MUST be used:

Input:

raw_value: random integer generated by the inner PRNG algorithm, between 1 and 0x7FFFFFFE ($2^{31}-2$) inclusive.
maxv: upper bound used during the scaling operation.

Output:

scaled_value: random integer between 0 and maxv-1 inclusive.

Algorithm:

```
scaled_value = (unsigned long) ((double)maxv * (double)raw_value /  
(double)0x7FFFFFFF);  
(NB: the above C type casting to unsigned long is equivalent to  
using floor() with positive floating point values.)
```

In this document, pmms_rand(maxv) denotes the PRNG function that implements the Park-Miller "minimal standard" algorithm, defined above, and that scales the raw value between 0 and maxv-1 inclusive, using the above scaling algorithm.

Additionally, the pmms_srand(seed) function must be provided to enable the initialization of the PRNG with a seed before calling pmms_rand(maxv) the first time. The seed is a 31-bit integer between 1 and 0x7FFFFFFE inclusive. In this specification, the seed is restricted to a value between 1 and 0xFFFF inclusive, as this is the Repair_Key 16-bit field value of the Repair FEC Payload ID ([Section 4.1.3](#)).

3.5. Coding Coefficients Generation Function

The coding coefficients, used during the encoding process, are generated at the RLC encoder by the following function each time a new repair symbol needs to be produced:


```

<CODE BEGINS>
/*
 * Fills in the table of coding coefficients (of the right size)
 * provided with the appropriate number of coding coefficients to
 * use for the repair symbol key provided.
 *
 * (in) repair_key    key associated to this repair symbol
 * (in) cc_tab[]      pointer to a table of the right size to store
 *                   coding coefficients. All coefficients are
 *                   stored as bytes, regardless of the m parameter,
 *                   upon return of this function.
 * (in) cc_nb[]       number of entries in the table. This value is
 *                   equal to the current encoding window size.
 * (in) m             Finite Field GF(2^m) parameter.
 * (out)              returns an error code
 */
int generate_coding_coefficients (UINT16    repair_key,
                                UINT8      cc_tab[],
                                UINT16     cc_nb,
                                UINT8      m)
{
    UINT32    i;

    if (repair_key == 0) {
        return SOMETHING_WENT_WRONG;
    }
    pmms_srand(repair_key);
    if (m == 1) {
        /* 0 is a valid coefficient value in binary GF */
        for (i = 0 ; i < cc_nb ; i++) {
            cc_tab[i] = (UINT8) pmms_rand(2);
        }
    } else {
        /* coefficient 0 is avoided in non-binary GF to consider each
         * source symbol */
        UINT32    maxv;
        maxv = get_gf_size(); /* i.e., 16 if m=4 or 256 if m=8 */
        for (i = 0 ; i < cc_nb ; i++) {
            do {
                cc_tab[i] = (UINT8) pmms_rand(maxv);
            } while (cc_tab[i] == 0)
        }
    }
    return EVERYTHING_IS_OKAY;
}
<CODE ENDS>

```

Figure 2: Coding Coefficients Generation Function pseudo-code

4. Sliding Window RLC FEC Scheme for Arbitrary ADU Flows

4.1. Formats and Codes

4.1.1. FEC Framework Configuration Information

The FEC Framework Configuration Information (or FFCI) includes information that MUST be communicated between the sender and receiver(s). More specifically, it enables the synchronization of the FECFRAME sender and receiver instances. It includes both mandatory elements and scheme-specific elements, as detailed below.

4.1.1.1. Mandatory Information

- o FEC Encoding ID: the value assigned to this fully specified FEC scheme MUST be XXXX, as assigned by IANA ([Section 9](#)).

When SDP is used to communicate the FFCI, this FEC Encoding ID is carried in the 'encoding-id' parameter.

4.1.1.2. FEC Scheme-Specific Information

The FEC Scheme-Specific Information (FSSI) includes elements that are specific to the present FEC scheme. More precisely:

Encoding symbol size (E): a non-negative integer that indicates the size of each encoding symbol in bytes;

m parameter (m): the length of the elements in the finite field, in bits, where m is equal to 1, 4 or 8;

These elements are required both by the sender (RLC encoder) and the receiver(s) (RLC decoder).

When SDP is used to communicate the FFCI, this FEC scheme-specific information is carried in the 'fssi' parameter in textual representation as specified in [[RFC6364](#)]. For instance:

```
fssi=E:1400,m:8
```

If another mechanism requires the FSSI to be carried as an opaque octet string (for instance, after a Base64 encoding), the encoding format consists of the following 2 octets:

Encoding symbol length (E): 16-bit field.

m parameter (m): 8-bit field.




Figure 3: FSSI Encoding Format

4.1.2. Explicit Source FEC Payload ID

A FEC source packet MUST contain an Explicit Source FEC Payload ID that is appended to the end of the packet as illustrated in Figure 4.




Figure 4: Structure of an FEC Source Packet with the Explicit Source FEC Payload ID

More precisely, the Explicit Source FEC Payload ID is composed of the following field (Figure 5):

Encoding Symbol ID (ESI) (32-bit field): this unsigned integer identifies the first source symbol of the ADUI corresponding to this FEC source packet. The ESI is incremented for each new source symbol, and after reaching the maximum value ($2^{32}-1$), wrapping to zero occurs.




Figure 5: Source FEC Payload ID Encoding Format

4.1.3. Repair FEC Payload ID

A FEC repair packet MUST contain a Repair FEC Payload ID that is prepended to the repair symbol as illustrated in Figure 6. There can be one or more repair symbol per FEC repair packet. When this is the

case, the number of repair symbols within this FEC repair packet is easily deduced by comparing the known received FEC repair packet size (equal to the UDP payload size when UDP is the underlying transport protocol) and the symbol size, E , communicated in the FFCI.

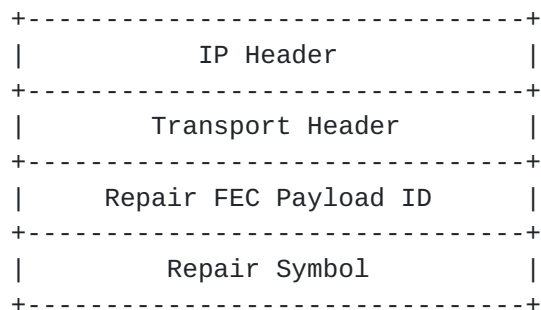


Figure 6: Structure of an FEC Repair Packet with the Repair FEC Payload ID

More precisely, the Repair FEC Payload ID is composed of the following fields (Figure 7):

Repair_Key (16-bit field): this unsigned integer is used as a seed by the coefficient generation function ([Section 3.5](#)) in order to generate the desired number of coding coefficients. Value 0 MUST NOT be used. When a FEC repair packet contains several repair packets, this repair key value is that of the first repair symbol. The remaining repair keys can be deduced by incrementing by 1 this value, up to a maximum value of 65535 after which it loops back to 1.

Number of Source Symbols in the Encoding Window, NSS (16-bit field):

this unsigned integer indicates the number of source symbols in the encoding window when this repair symbol was generated.

ESI of first source symbol in encoding window, FSS_ESI (32-bit field):

this unsigned integer indicates the ESI of the first source symbol in the encoding window when this repair symbol was generated.




Figure 7: Repair FEC Payload ID Encoding Format

4.1.4. Additional Procedures

The following procedure applies:

- o The ESI of source symbols MUST start with value 0 for the first source symbol and MUST be managed sequentially. Wrapping to zero will happen after reaching the maximum 32-bit value.

5. FEC Code Specification

5.1. Encoding Side

This section provides a high level description of a Sliding Window RLC encoder.

Whenever a new FEC repair packet is needed, the RLC encoder instance first gathers the `ew_size` source symbols currently in the sliding encoding window. Then it chooses a repair key, which can be a non zero monotonically increasing integer value, incremented for each repair symbol up to a maximum value of 65535 (as it is carried within a 16-bit field) after which it loops back to 1 (indeed, being used as a PRNG seed, value 0 is prohibited). This repair key is communicated to the coefficient generation function (Section [Section 3.5](#)) in order to generate `ew_size` coding coefficients. Finally, the FECFRAME sender computes the repair symbol as a linear combination of the `ew_size` source symbols using the `ew_size` coding coefficients. When `E` is small and when there is an incentive to pack several repair symbols within the same FEC Repair Packet, the appropriate number of repair symbols are computed. The only constraint is to increment by 1 the repair key for each of them, keeping the same `ew_size` source symbols, since only the first repair key will be carried in the Repair FEC Payload ID. The FEC repair packet can then be sent. The source versus repair FEC packet transmission order is out of scope of this document and several approaches exist that are implementation specific.

5.2. Decoding Side

This section provides a high level description of a Sliding Window RLC decoder.

A FECFRAME receiver needs to maintain a linear system whose variables are the received and lost source symbols. Upon receiving a FEC repair packet, a receiver first extracts all the repair symbols it contains (in case several repair symbols are packed together). For each repair symbol, when at least one of the corresponding source symbols it protects has been lost, the receiver adds an equation to the linear system (or no equation if this repair packet does not

change the linear system rank). This equation of course re-uses the `ew_size` coding coefficients that are computed by the same coefficient generation function (Section [Section 3.5](#)), using the repair key and encoding window descriptions carried in the Repair FEC Payload ID. Whenever possible (i.e., when a sub-system covering one or more lost source symbols is of full rank), decoding is performed in order to recover lost source symbols. Each time an ADUI can be totally recovered, it is assigned to the corresponding application flow (thanks to the Flow ID (F) field of the ADUI) and padding (if any) removed (thanks to the Length (L) field of the ADUI). This ADU is finally passed to the corresponding upper application. Received FEC source packets, containing an ADU, can be passed to the application either immediately or after some time to guaranty an ordered delivery to the application(s). This document does not mandate any approach as this is an operational and management decision.

With real-time flows, a lost ADU that is decoded after the maximum latency (or an ADU received far too late) should not be considered by the application. Instead the associated source symbols should be removed from the linear system maintained by the receiver(s). [Appendix A](#) discusses a backward compatible optimization whereby those late source symbols may still be useful to improve the global loss recovery performance.

6. Implementation Status

Editor's notes: RFC Editor, please remove this section motivated by [RFC 6982](#) before publishing the RFC. Thanks.

An implementation of the Sliding Window RLC FEC Scheme for FECFRAME exists:

- o Organisation: Inria
- o Description: This is an implementation of the Sliding Window RLC FEC Scheme. It relies on a modified version of our OpenFEC (<http://openfec.org>) FEC code library. It is integrated in our FECFRAME software (see [[fecframe-ext](#)]).
- o Maturity: prototype.
- o Coverage: this software complies with the Sliding Window RLC FEC Scheme (limited to $m=8$ as of June, 2017).
- o Lincensing: proprietary.
- o Contact: vincent.roca@inria.fr

7. Security Considerations

The FEC Framework document [[RFC6363](#)] provides a comprehensive analysis of security considerations applicable to FEC schemes.

Therefore, the present section follows the security considerations section of [[RFC6363](#)] and only discusses specific topics.

7.1. Attacks Against the Data Flow

7.1.1. Access to Confidential Content

The Sliding Window RLC FEC Scheme specified in this document does not change the recommendations of [[RFC6363](#)]. To summarize, if confidentiality is a concern, it is RECOMMENDED that one of the solutions mentioned in [[RFC6363](#)] is used with special considerations to the way this solution is applied (e.g., is encryption applied before or after FEC protection, within the end-system or in a middlebox) to the operational constraints (e.g., performing FEC decoding in a protected environment may be complicated or even impossible) and to the threat model.

7.1.2. Content Corruption

The Sliding Window RLC FEC Scheme specified in this document does not change the recommendations of [[RFC6363](#)]. To summarize, it is RECOMMENDED that one of the solutions mentioned in [[RFC6363](#)] is used on both the FEC Source and Repair Packets.

7.2. Attacks Against the FEC Parameters

The FEC Scheme specified in this document defines parameters that can be the basis of attacks. More specifically, the following parameters of the FFCI may be modified by an attacker who only targets receivers ([Section 4.1.1.2](#)):

- o FEC Encoding ID: changing this parameter leads the receivers to consider a different FEC Scheme, which enables an attacker to create a Denial of Service (DoS);
- o Encoding symbol length (E): setting this E parameter to a different value will confuse the receivers and create a DoS. More precisely, the FEC Repair Packets received will probably no longer be multiple of E, leading receivers to reject them;
- o m parameter: changing this parameter triggers a DoS since the receivers will generate a different set of coding coefficients. The recovered source symbols (and thereafter ADUs) will be corrupted.

An attacker who only targets a sender will achieve the same results. However if the attacker targets both sender and receivers at the same time (the same wrong piece of information is communicated to everybody), the results will be suboptimal but less severe.

It is therefore RECOMMENDED that security measures are taken to guarantee the FFCI integrity, as specified in [RFC6363]. How to achieve this depends on the way the FFCI is communicated from the sender to the receiver, which is not specified in this document.

Similarly, attacks are possible against the Explicit Source FEC Payload ID and Repair FEC Payload ID: by modifying the Encoding Symbol ID (ESI), or the repair key, NSS or FSS_ESI. It is therefore RECOMMENDED that security measures are taken to guarantee the FEC Source and Repair Packets as stated in [RFC6363].

7.3. When Several Source Flows are to be Protected Together

The Sliding Window RLC FEC Scheme specified in this document does not change the recommendations of [RFC6363].

7.4. Baseline Secure FEC Framework Operation

The Sliding Window RLC FEC Scheme specified in this document does not change the recommendations of [RFC6363] concerning the use of the IPsec/ESP security protocol as a mandatory to implement (but not mandatory to use) security scheme. This is well suited to situations where the only insecure domain is the one over which the FEC Framework operates.

8. Operations and Management Considerations

The FEC Framework document [RFC6363] provides a comprehensive analysis of operations and management considerations applicable to FEC schemes. Therefore, the present section only discusses specific topics.

8.1. Operational Recommendations: Finite Field Element Size (m Parameter)

The present document requires that m equals 1 (binary case), 4 or 8. It is expected that $m = 8$ will be mostly used since it warrants a high loss protection. Additionally, elements in the finite field are 8 bits long, which makes read/write memory operations aligned on bytes during encoding and decoding.

An alternative when one can accommodate a lower loss protection is $m = 4$. Elements in the finite field are 4 bits long, so if 2 elements are accessed at a time, read/write memory operations are aligned on bytes during encoding and decoding.

Finally, in particular when dealing with large encoding windows, an alternative is $m = 1$. In that case operations symbols can be

directly XORed together which warrants high bitrate encoding and decoding operations.

Since several values for the m parameter are possible, the use case SHOULD define which value or values need to be supported. In any case, any compliant implementation MUST support at least the default $m = 8$ value.

9. IANA Considerations

This document registers one value in the "FEC Framework (FECFRAME) FEC Encoding IDs" registry [[RFC6363](#)] as follows:

- o XXX refers to the Sliding Window Random Linear Codes (RLC) FEC Scheme for Arbitrary Packet Flows, as defined in Section XXX of this document.

10. Acknowledgments

The authors would like to thank Belkacem Teibi (Inria) who in particular implemented the RLC codec. The author would also like to thank Marie-Jose Montpetit for her valuable feedbacks on this document.

11. References

11.1. Normative References

[fecframe-ext]

Roca, V. and A. Begen, "Forward Error Correction (FEC) Framework Extension to Sliding Window Codes", Transport Area Working Group (TSVWG) [draft-roca-tsvwg-fecframev2](#) (Work in Progress), June 2017, <<https://tools.ietf.org/html/draft-roca-tsvwg-fecframev2>>.

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

[RFC6363] Watson, M., Begen, A., and V. Roca, "Forward Error Correction (FEC) Framework", [RFC 6363](#), DOI 10.17487/RFC6363, October 2011, <<http://www.rfc-editor.org/info/rfc6363>>.

- [RFC6364] Begen, A., "Session Description Protocol Elements for the Forward Error Correction (FEC) Framework", [RFC 6364](#), DOI 10.17487/RFC6364, October 2011, <<http://www.rfc-editor.org/info/rfc6364>>.

11.2. Informative References

- [CA90] Carta, D., "Two Fast Implementations of the Minimal Standard Random Number Generator", Communications of the ACM, Vol. 33, No. 1, pp.87-88, January 1990.
- [PM88] Park, S. and K. Miller, "Random Number Generators: Good Ones are Hard to Find", Communications of the ACM, Vol. 31, No. 10, pp.1192-1201, 1988.
- [PTVF92] Press, W., Teukolsky, S., Vetterling, W., and B. Flannery, "Numerical Recipes in C; Second Edition", Cambridge University Press, ISBN: 0-521-43108-5, 1992.
- [rand31pmc] Whittle, R., "31 bit pseudo-random number generator", September 2005, <<http://www.firstpr.com.au/dsp/rand31/rand31-park-miller-carta.cc.txt>>.
- [RFC5170] Roca, V., Neumann, C., and D. Furodet, "Low Density Parity Check (LDPC) Staircase and Triangle Forward Error Correction (FEC) Schemes", [RFC 5170](#), DOI 10.17487/RFC5170, June 2008, <<http://www.rfc-editor.org/info/rfc5170>>.
- [RFC6726] Paila, T., Walsh, R., Luby, M., Roca, V., and R. Lehtonen, "FLUTE - File Delivery over Unidirectional Transport", [RFC 6726](#), DOI 10.17487/RFC6726, November 2012, <<http://www.rfc-editor.org/info/rfc6726>>.
- [RFC6816] Roca, V., Cunche, M., and J. Lacan, "Simple Low-Density Parity Check (LDPC) Staircase Forward Error Correction (FEC) Scheme for FECFRAME", [RFC 6816](#), DOI 10.17487/RFC6816, December 2012, <<http://www.rfc-editor.org/info/rfc6816>>.
- [RFC6865] Roca, V., Cunche, M., Lacan, J., Bouabdallah, A., and K. Matsuzono, "Simple Reed-Solomon Forward Error Correction (FEC) Scheme for FECFRAME", [RFC 6865](#), DOI 10.17487/RFC6865, February 2013, <<http://www.rfc-editor.org/info/rfc6865>>.

- [Roca16] Roca, V., Teibi, B., Burdinat, C., Tran, T., and C. Thienot, "Block or Convolutional AL-FEC Codes? A Performance Comparison for Robust Low-Latency Communications", Submitted for publication <https://hal.inria.fr/hal-01395937/en/>, November 2016, <<https://hal.inria.fr/hal-01395937/en/>>.
- [WI08] Whittle, R., "Park-Miller-Carta Pseudo-Random Number Generator", <http://www.firstpr.com.au/dsp/rand31/>, January 2008, <<http://www.firstpr.com.au/dsp/rand31/>>.

Appendix A. Decoding Beyond Maximum Latency Optimization

This annex introduces non normative considerations. They are provided as suggestions, without any impact on interoperability. For more information see [[Roca16](#)].

It is possible to improve the decoding performance of sliding window codes without impacting maximum latency, at the cost of extra CPU overhead. The optimization consists, for a receiver, to extend the linear system beyond the decoding window:

$$ls_max_size > dw_max_size$$

Usually the following choice is a good trade-off between decoding performance and extra CPU overhead:

$$ls_max_size = 2 * dw_max_size$$




Figure 8: Relationship between parameters to decode beyond maximum latency.

It means that source symbols (and therefore ADUs) may be decoded even if their transport protocol added latency exceeds the maximum value permitted by the application. It follows that these source symbols SHOULD NOT be delivered to the application and SHOULD be dropped once they are no longer needed. However, decoding these late symbols significantly improves the global robustness in bad reception conditions and is therefore recommended for receivers experiencing bad channels[Roca16]. In any case whether or not to use this facility and what exact value to use for the `ls_max_size` parameter are decisions made by each receiver independently, without any impact on others, neither the other receivers nor the source.

Author's Address

Vincent Roca
INRIA
Grenoble
France

EMail: vincent.roca@inria.fr