

SMPTE STANDARD

Forward Error Correction for Transport of High Bit Rate Media Signals over IP Networks (HBRMT)



Page 1 of 22 pages

Table of Contents	Page
Foreword	2
Intellectual Property	2
Introduction.....	2
1 Scope	4
2 Conformance Notation	4
3 Normative References	4
4 Acronyms (Informative).....	5
5 Definitions (Normative).....	5
6 FEC Datagram Structure.....	6
6.1 Datagram Structure.....	6
6.2 RTP Header for FEC Datagrams	7
7 FEC Scheme.....	8
7.1 FEC Datagram Arrangement	8
7.2 FEC Buffer Overhead and Latency Implications.....	10
7.3 FEC Header Format.....	10
7.4 Construction of FEC Datagrams	12
7.5 FEC Traffic Shaping Issues	14
7.6 Re-Order Tolerance (Informative).....	14
Annex A Bibliography (Informative)	15
Annex B Non Block-Aligned FEC Arrangement (Informative)	16
Annex C Block-Aligned FEC Arrangement (Informative).....	17
Annex D Delay (Informative).....	18
Annex E Conceptual Diagram (Informative)	19
Annex F Datagram Recovery (Informative)	20
Annex G Revision Notes (Informative).....	22

Foreword

SMPTE (the Society of Motion Picture and Television Engineers) is an internationally-recognized standards developing organization. Headquartered and incorporated in the United States of America, SMPTE has members in over 80 countries on six continents. SMPTE's Engineering Documents, including Standards, Recommended Practices, and Engineering Guidelines, are prepared by SMPTE's Technology Committees. Participation in these Committees is open to all with a bona fide interest in their work. SMPTE cooperates closely with other standards-developing organizations, including ISO, IEC and ITU.

SMPTE Engineering Documents are drafted in accordance with the rules given in Part XIII of its Operations Manual.

SMPTE ST 2022-5 was prepared by Technology Committee 32NF.

Intellectual Property

At the time of publication no notice had been received by SMPTE claiming patent rights essential to the implementation of this Standard. However, attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. SMPTE shall not be held responsible for identifying any or all such patent rights.

Introduction

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

IP-based networks have become increasingly important for delivery of Standard Definition (SD) and High-Definition (HD) television signals at high bit rates (without compression). However, existing transport protocols do not fully meet user requirements, especially for a standardized forward error correction (FEC) technique required for interoperability between different manufacturers. This standard defines a FEC technique to be used to correct for errors induced when real time video/audio payloads are transported over IP networks. This standard is designed to support FEC for SMPTE ST 2022-6, Transport of High Bit Rate Media Signals over IP (HBRMT).

This document defines a structured mechanism for FEC. Nothing in this standard defines when this FEC mechanism should be used. It is up to the implementer to determine whether FEC is required, and whether the mechanism presented here meets their application requirements.

This standard is intended for real-time audio/video applications such as contribution, primary distribution and digital cinema. This standard is designed to be applied to television transport for broadcast production and is not intended for emission purposes. Typically a connection will be set up and torn down as a managed configuration of transmitting and receiving equipment. A connection may be unicast or multicast mode. The support of Session Description Protocol (SDP) and Real Time Control Protocol (RTCP) are not required for equipment supporting this standard.

This standard defines two levels of FEC – Level A and Level B. Level A uses one FEC stream and Level B uses two FEC streams. These levels are different from the Uneven Level Protection levels of RFC 5109.

On IP networks, datagram losses typically come from three sources — gross reordering, bit-error induced datagram drops and burst losses/drops. For any FEC scheme to operate properly, errors from these sources need to be low enough so that the FEC scheme can correct enough of these errors to meet the application requirements. Implementers should be aware of limitations of any FEC scheme and take steps to ensure that application of this standard will meet their objectives, given uncorrected link performance.

Because, when using the IP protocol, the underlying transport technologies such as Ethernet will ensure datagrams received with bit errors will be discarded, there is no requirement for an error correction scheme that can handle datagrams containing errors – every datagram will either arrive correctly or not at all.

An RTP payload format for Generic Forward Error Correction Datagrams has been defined in RFC 2733 which was used in SMPTE ST 2022-1 to enable regeneration of lost IP datagrams with MPEG-2 video/audio Transport Streams. However, in RFC 2733 a few fields (the P, X, and CC fields) in the RTP header are specified in ways that are not consistent with the way that they are designed in RFC 3550. This prevents payload-independent validity check of the RTP Datagrams. This has been corrected in RFC 5109. This standard is based on SMPTE ST 2022-1 and includes the corrections contained in RFC 5109.

The RFC 5109 FEC Datagram structure consists of a FEC Header followed by one or several sections of FEC Level Header and FEC Level Payload. RFC 5109 provides for protection of up to 16 or 48 Media Datagrams and includes a mechanism for providing uneven correction of essential media content and for exclusion of less important datagrams by use of masking. Uneven protection is not included in this standard.

To recover from burst loss at high bit rates, the SMPTE ST 2022-1 FEC column size and row size is extended and an RFC 5109 style header is used. The same traditional error correcting codes are applied to non-consecutive Media Datagrams. Each FEC Datagram is Associated with Media Datagrams periodically selected. Therefore, consecutive RTP Datagrams can be recovered from consecutive FEC Datagrams.

This standard is intended for constant bit rate (CBR) and anticipates variable bit rate (VBR) applications. CBR applications are assumed to use constant datagram size and constant datagram rate as this provides the lowest FEC overheads. VBR applications can use variable datagram rates and/or variable datagram lengths as defined in SMPTE ST 2022-3.

The method described herein has been in the process of standardization for several years; the informal name “High Bit-Rate Media Transport” and the acronym “HBRMT” have been associated with the work throughout that time. The mention of HBRMT in the title of this document references that heritage.

1 Scope

This Standard defines a forward error correction technique for the carriage of High Bit Rate real-time media over IP networks.

The term High Bit Rate is used herein to distinguish from other Media-over-IP applications in which primarily compressed signals are transported; the FEC system described herein is differentiated by an ability to accommodate media payloads at rates of 270 Mbits/sec and higher.

2 Conformance Notation

Normative text is text that describes elements of the design that are indispensable or contains the conformance language keywords: "shall", "should", or "may". Informative text is text that is potentially helpful to the user, but not indispensable, and can be removed, changed, or added editorially without affecting interoperability. Informative text does not contain any conformance keywords.

All text in this document is, by default, normative, except: the Introduction, any section explicitly labeled as "Informative" or individual paragraphs that start with "Note:"

The keywords "shall" and "shall not" indicate requirements strictly to be followed in order to conform to the document and from which no deviation is permitted.

The keywords, "should" and "should not" indicate that, among several possibilities, one is recommended as particularly suitable, without mentioning or excluding others; or that a certain course of action is preferred but not necessarily required; or that (in the negative form) a certain possibility or course of action is deprecated but not prohibited.

The keywords "may" and "need not" indicate courses of action permissible within the limits of the document.

The keyword "reserved" indicates a provision that is not defined at this time, shall not be used, and may be defined in the future. The keyword "forbidden" indicates "reserved" and in addition indicates that the provision will never be defined in the future.

A conformant implementation according to this document is one that includes all mandatory provisions ("shall") and, if implemented, all recommended provisions ("should") as described. A conformant implementation need not implement optional provisions ("may") and need not implement them as described.

Unless otherwise specified, the order of precedence of the types of normative information in this document shall be as follows: Normative prose shall be the authoritative definition; Tables shall be next; followed by formal languages; then figures; and then any other language forms.

3 Normative Reference

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

IETF RFC 3550, RTP: A Transport Protocol for Real-Time Applications

4 Acronyms (Informative)

CBR: Constant Bit Rate

CSRC: Contributing Source

D: Number of Media Datagrams protected by a row FEC Datagram

FEC: Forward Error Correction

IANA: Internet Assigned Numbers Authority

IP: Internet Protocol (see IETF STD05)

L: Number of Media Datagrams protected by a column FEC Datagram

RTCP: Real Time Control Protocol (see IETF RFC 3550)

RTP: Real Time Protocol (see IETF RFC 3550)

SDI: Serial Digital Interface

SDP: Session Description Protocol (see IETF RFC 4566)

SDTI: Serial Digital Transport Interface

SNbase: Base RTP Sequence Number

SSRC: Synchronization Source

TS: Timestamp

UDP: User Datagram Protocol (see IETF STD06)

VBR: Variable Bit Rate

XOR: Exclusive OR

5 Definitions (Normative)

Associated: A FEC Datagram is said to be "associated" with one or more Media Datagrams (or vice versa) when those Media Datagrams are used to generate the FEC Datagram (by use of the XOR operation).

Block-Aligned FEC Operation: Block-Aligned FEC operation is a correction scheme using a two dimensional matrix where the Media Datagrams are a contiguous group of $L \times D$ datagrams. The Media Datagrams are protected as follows. Level A protection is achieved by L FEC Datagrams derived from each column for the FEC Matrix. Optional level B protection is achieved by D FEC Datagrams derived from each row for the FEC Matrix. Level A FEC stream shall protect all media packets exactly once. Optional Level B FEC stream shall protect all media packets exactly once. See Figure 3.

Delay: Delay is time consumed by all sources such as processing delay, accumulation and de-accumulation delay, FEC processing delay, IP transfer delay, IP delay variation buffering and FEC Latency delay.

FEC Datagram: An FEC Datagram is an RTP Datagram consisting of an RTP header and FEC data payload. The FEC data payload is composed of an FEC Header and the FEC Payload. The FEC Datagrams are formatted according to the rules contained this standard.

FEC Header: The FEC Header is the header information contained in an FEC Datagram.

FEC Matrix: An FEC Matrix is a set of Media Datagrams ordered in a matrix with L columns and D rows. The datagrams are entered into the matrix to fill each row sequentially with incremented RTP numbers.

FEC Payload: The FEC Payload is the payload of an FEC Datagram.

High Bit Rate: At the time of development of this Standard, High Bit Rate commonly refers to transmission rates higher than 270 Mb/s (SDI, HD SDI, 3G), although the techniques specified herein may have continued utility at higher rates (e.g. 40 Gb/s, 100 Gb/s) as technology allows.

Latency: Latency is the time for data to become available for a subsequent process.

Media Datagram: A Media Datagram is an RTP Datagram consisting of an RTP header and media data payload. The media data payload is composed of a Payload Header and a Media Payload. Media Datagrams are formatted according to another standard (e.g. SMPTE ST 2022-6).

Media Header: A Media Header is the RTP header for the datagram containing the Media Payload.

Media Payload: The Media Payload is the raw data (including video, audio, and ancillary data) that are transmitted from the sender. The Payload Header and the Media Payload are placed inside of a Media Datagram.

Non Block-Aligned FEC Operation: A Non Block-Aligned FEC Operation is a correction scheme similar to a Block-Aligned FEC Operation, but performed over non-contiguous Media Datagrams. In a Non Block-aligned FEC Operation, the FEC columns are offset from each other to facilitate Traffic Shaping and/or lower Latency. The FEC rows of a Non Block-Aligned FEC Operation cover contiguous Media Datagrams. See Annex B and in particular Figure B1.

RTP Datagram: RTP Datagrams are defined in IETF RFC 3550.

Traffic Shaping: Traffic Shaping is the process of reducing peak variations in instantaneous bit rate so as to produce a more constant data flow.

Wrapped Media: Media that is in a container such as that described in SMPTE ST 377-1:2009.

6 FEC Datagram Structure

This section defines the construction of RTP Datagrams carrying FEC information.

6.1 Datagram Structure

A FEC Datagram is constructed by placing an FEC Header and FEC processed payload into the RTP payload, as shown in Figure 1.




Figure 1 – FEC Datagram Header Format

6.2 RTP Header for FEC Datagrams

All the fields in the RTP header of FEC Datagrams are used according to RFC 3550 with some of them further clarified below.

Version (V): 2 bits

The version field shall be set to 2.

Padding (P): 1 bit

Set to '0'. Padding shall not be used for FEC Datagram.

Extension (X): 1 bit

This field is not used for this payload type, and shall be set to 0

CSRC Count (CC): 4 bits

This field shall be set to 0.

Note: There are no CSRC lists present in the FEC Datagrams.

Marker (M): 1 bit

This field is not used for this payload type, and shall be set to 0

Payload Type (PT): 7 bits

If a Dynamic Payload Type is used, then the value should be set to 99 decimal. Alternatively, payload types may be set by other means in accordance with IETF RFC 3550.

Sequence Number (SN): 16 bits

The sequence number shall be one higher than the sequence number in the previously transmitted FEC Datagram for the same FEC stream.

Timestamp (TS): 32 bits

The timestamp shall be set to the value of the media RTP clock at the instant the FEC Datagram is transmitted. For FEC senders which lack access to the media RTP clock, an approximate value may be generated based on the RTP timestamps of the Media Datagrams, providing that this generated value is no earlier than the RTP timestamp of the last Media Datagram which is protected by that FEC Datagram.

Synchronization Source (SSRC): 32 bits

The SSRC value shall be the same as the SSRC value of the media stream it protects.

7 FEC scheme

7.1 FEC Datagram Arrangement

The FEC Datagrams shall be sent by the sender as separate RTP streams. The RTP header of the FEC Datagrams shall be according to RFC 3550 and consequently shall not require the use of any additional communication. This Standard defines two levels of FEC protection — Level A and Level B. Level A devices shall support only the column FEC stream. Level B devices shall support both column and row FEC streams. The column and row FEC streams shall be carried on separate UDP ports to allow them to have separate sequence number handling.

For all senders of FEC (Level A or B), Media Datagrams shall be sent to destination port N (where N is an even integer per RFC 3550), and the column FEC Datagrams shall be sent to destination port N+2. For senders of Level B FEC, the additional row FEC stream shall be sent to destination port N+4.

The column and row FEC streams shall use the same IP source address and the same UDP source port as the Associated media stream. Senders and receivers shall support all port numbers from the registered port range, and the Dynamic and/or private ports range as identified in the IANA Ports Registry. The column and row FEC streams shall use the same IP destination address (either unicast or multicast) as the Associated media stream.

Level A and Level B receivers shall process column FEC stream. Level A receivers shall function nominally in the presence of column and row FEC streams. Level B receivers shall be able to process both column and row FEC streams. Level A and Level B receivers shall be able to continue to function nominally without any FEC stream.

Each FEC Datagram is Associated with Media Datagrams periodically selected. Therefore, consecutive Media Datagrams can be recovered using consecutive FEC Datagrams. The process is detailed in Figure 2. In this Standard, L shall designate the number of columns of protected datagrams and D shall designate the number of rows.

Note: The main advantage of this scheme over other mechanisms is the burst error correction capacity. The error correcting function chosen is XOR which has the ability to recover any one lost datagram. If a one dimensional error correction scheme based on XOR is used (i.e. applied to D consecutive datagrams), a burst error of two or more lost datagrams is not recoverable. By using the two dimensional error correction scheme, the recoverability is greatly improved, since it can recover up to L consecutive lost datagrams.

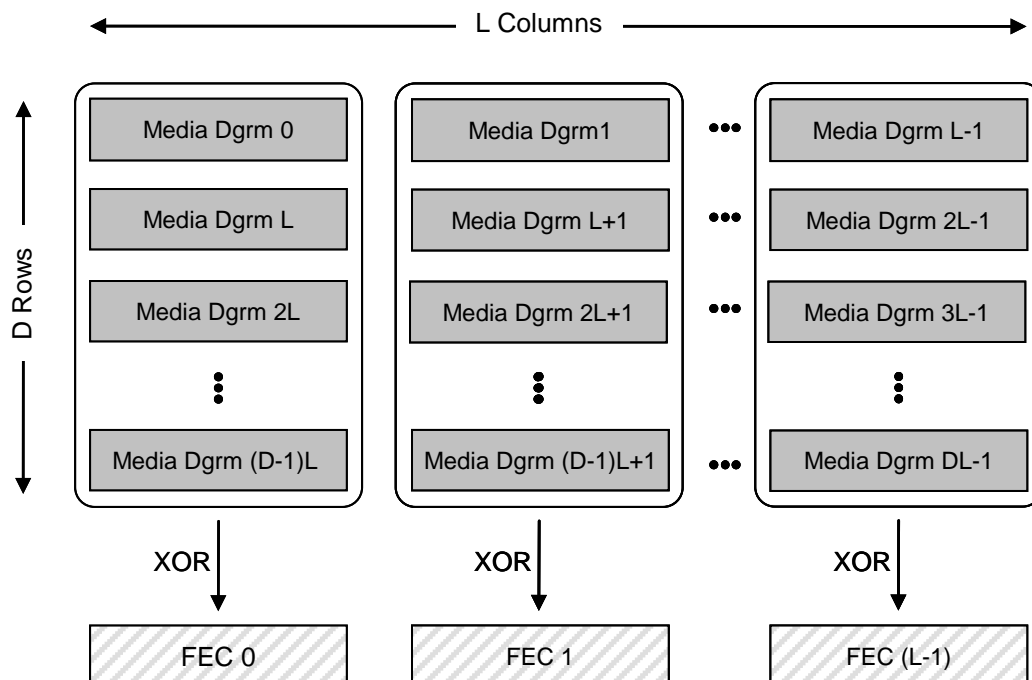


Figure 2 – FEC Column-only Encoding Scheme

In Figure 2, the encoding scheme is schematized for $L \cdot D$ Media Datagrams. The period chosen between Media Datagrams covered by a given column FEC Datagram is L . Thus the payload of the k^{th} column FEC Datagram is computed based on the D datagrams numbered $nL+k$ ($0 \leq n \leq D-1$). For the row FEC stream there is no interleaving in the FEC calculation. The payload of the j^{th} row FEC Datagram is computed based on the L datagrams numbered $jL+n$ ($0 \leq n \leq L-1$). The alignment of the columns is for illustration. Implementations may use this alignment for simplicity, but there are some advantages to be gained by offsetting the columns — see Annex B. This means that receiving devices shall not make any assumptions about the relationship between FEC Datagrams beyond those that are explicitly specified.

The row FEC stream shall be applied to a run of consecutive datagrams the length of the L parameter of the first stream. If the column FEC is block-aligned, the row FEC shall be aligned with the column FEC to produce an FEC structure as shown in Figure 3. In Figure 3 the datagrams labeled Media Dgrm are the Media Datagrams, the datagrams labeled FEC_{col} are the column FEC stream datagrams, and the datagrams labeled FEC_{row} are the row FEC stream datagrams.

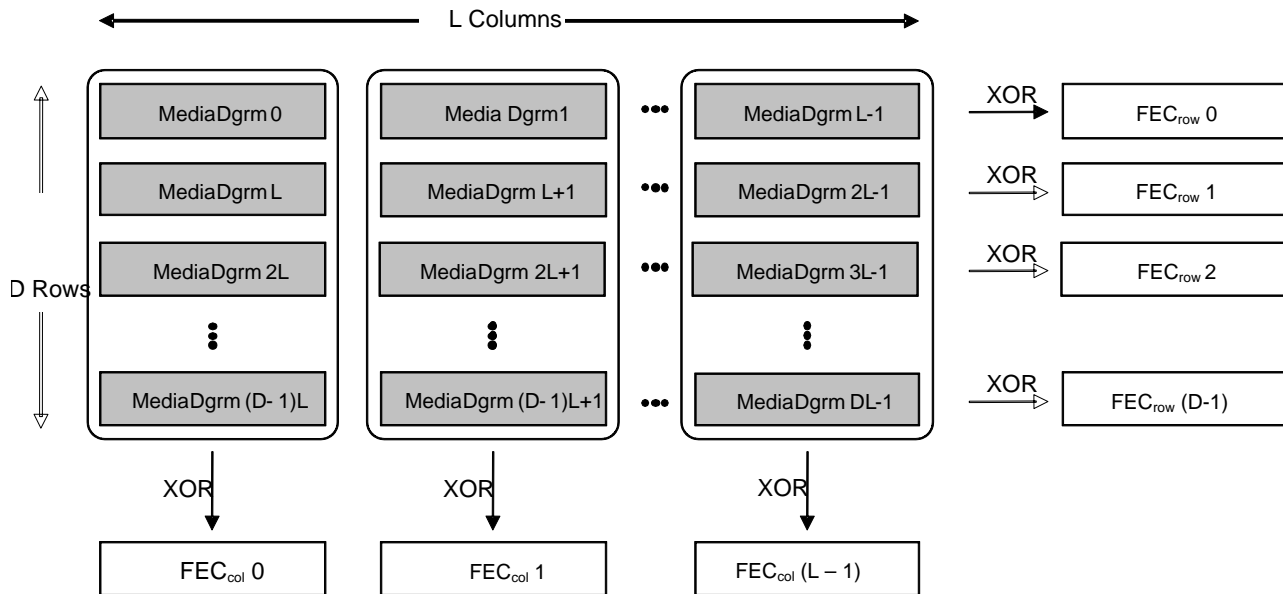


Figure 3 – The Column and Row Encoding Scheme

Note: The row FEC stream can cope with any single datagram loss, and the column FEC stream can cope with burst losses up to 'L' in length.

7.2 FEC Buffer Overhead and Latency Implications

To promote interoperability and simplify implementation, limits shall be specified for values of the L and D parameters. These limits are media type specific and are defined in the appropriate mapping documents (e.g. SMPTE ST 2022-6). Devices shall support all combinations of values of L and D that comply with the limits specified in the appropriate mapping document. Devices may extend beyond these values if desired.

In addition to the constraints of the mapping document, a device shall only support two FEC streams in the case where $L \geq 4$. Due to constraints in the FEC Header format defined in this document, L and D values are each limited to 1020.

7.3 FEC Header Format

The FEC Header shall be 16 octets. The FEC Header contains data that are used to identify the Media Datagrams that are protected by this FEC Datagram as well as data for protection of RTP header for the same Media Datagrams. The P recovery field, X recovery field, CC recovery field, M recovery field, PT recovery field, TS recovery field and length recovery field, are computed by XOR operation. The SN base field, offset field and the NA field are used to identify the datagrams that are protected by this FEC Datagram.

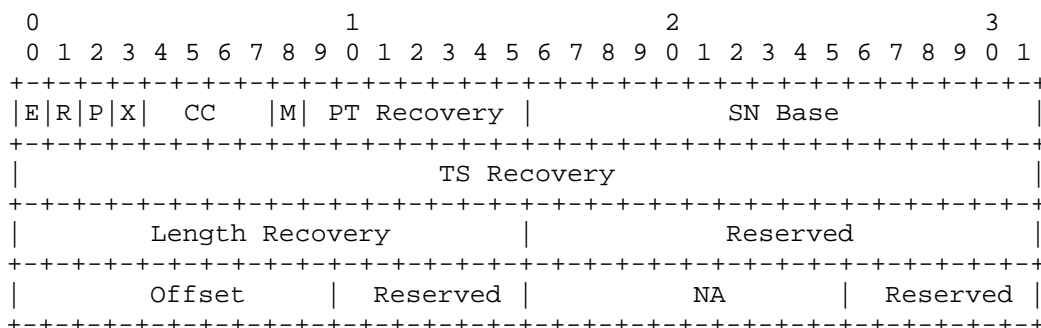


Figure 4 – Definition of the FEC Header

Extension flag (E) 1 bit:

The E bit is the extension flag reserved to indicate any future extension to this specification. It shall be set to 0.

Reserved (R) 1 bit:

Shall be set to zero by the sender.

Padding Recovery field (P) 1 bit:

The Padding Recovery field shall be computed via the XOR operation applied to the corresponding P values from the RTP headers of the Media Datagrams Associated with the FEC Datagram.

Extension Recovery field (X) 1 bit:

The Extension Recovery field shall be computed via the XOR operation applied to the corresponding X values from the RTP headers of the Media Datagrams Associated with the FEC Datagram.

CSRC Count Recovery field (CC) 4 bits:

The CSRC Count Recovery field shall be computed via the XOR operation applied to the corresponding CC values from the RTP headers of the Media Datagrams Associated with the FEC Datagram.

Marker Recovery field (M) 1 bit:

The Marker Recovery field shall be computed via the XOR operation applied to the corresponding M values from the RTP headers of the Media Datagrams Associated with the FEC Datagram.

Payload Type Recovery field (PT Recovery) 7 bits:

The Payload Type Recovery field shall be computed via the XOR operation applied to the corresponding PT values from the RTP headers of the Media Datagrams Associated with the FEC Datagram.

Sequence Number Base (SN Base) 2 octets:

The Sequence Number Base field shall be set to the lowest sequence number, taking wrap around into account, of those Media Datagrams protected by FEC.

Time Stamp Recovery (TS) 4 octets:

The Time Stamp Recovery field shall be computed via the XOR operation applied to the timestamps of the Media Datagrams Associated with this FEC Datagram. This allows the timestamps to be completely recovered.

Length Recovery (LR) 2 octets:

This field is used to determine the length of any recovered datagrams. The Length Recovery field shall be computed via the XOR operation applied to the unsigned network-ordered 16-bit

representation of the adjusted lengths (in octets) of each of the Media Datagrams Associated with this FEC Datagram. The adjusted length of each Media Datagram shall be computed as the unsigned network-ordered 16-bit sum of the length (in octets) of the Media Payload, CSRC list, extension and padding fields that are contained within the datagram. This allows the FEC procedure to be applied even when the lengths of the protected Media Datagrams are not identical.

The CSRC list, RTP extension, and padding of the Media Payload datagrams, if present, shall be counted as part of the payload. For example, assume that an FEC Datagram is being generated by XORing two Media Datagrams together. The length of the payload of two Media Datagrams is 3 (0b011) and 5 (0b101) octets, respectively. The length recovery field is then encoded as 0b011 XOR 0b101 = 0b110.

Reserved (Reserved) 2 octets:
Shall be set to zero by the sender.

Offset: (Offset) 10 bits:
This field is the period used to select the Media Datagrams Associated with this FEC Datagram, and shall be the L parameter as defined in this Standard for datagrams computed over columns (belonging to the first FEC stream). For datagrams computed over rows (belonging to the second FEC stream), this parameter shall always be one. This field shall be kept constant by the sender during a transmission for each FEC stream.

Reserved (Reserved) 6 bits:
Shall be set to zero by the sender.

Number of Media Datagrams Associated (NA): 10 bits.
This field indicates the number of Media Datagrams Associated with this FEC Datagram, and shall be the D parameter as defined in this standard for datagrams computed over columns (belonging to the first FEC stream), and shall correspond to the L parameter as defined in this standard for datagrams computed over rows (belonging to the second FEC stream). This field shall be kept constant by the sender during a transmission for each FEC stream.

Reserved (Reserved) 6 bits:
Shall be set to zero by the sender.

The Media Datagrams protected by any given FEC Datagram shall be defined as those with sequence numbers given by the formula:

$$SNBase + j \times Offset, \dots \text{where } 0 \leq j < NA$$

7.4 Construction of FEC Datagrams

Senders shall construct FEC Datagrams by performing an XOR operation over protected Media Datagrams. The process is illustrated in Figure 5. This section describes one process that senders can use to construct FEC Datagrams. Other methods may be used to create FEC Datagrams, however regardless of the process the sender uses, the result shall be identical to the result produced by the method described below.

The column and row process are identical. Column FEC uses an offset that is equal to the row length, L, thereby providing interleaving for burst loss protection. Row FEC uses an offset of 1 and protects L consecutive datagrams. The payload of the FEC Datagram consists of the FEC Header and FEC Payload. The FEC Header contains data that associates the FEC Datagram to the packets it protects, as well as data that can be used to reconstruct the RTP header of lost packets. The FEC Payload carries data that can be used to reconstruct the Media Datagram payload.

7.4.1 Construction of FEC Header

The fields SN Base, Offset and NA in the FEC Header define the RTP sequence numbers of the Media Datagrams that are protected by this FEC Datagram. The fields shall be set such that all RTP Media Datagrams are protected once, and only once, by the consecutive flow of FEC Datagrams with this value of Offset.

The recovery fields in the FEC Header, Figure 4, P, X, CC, M, PT and TS Recovery are computed by XOR processing of the same fields of the RTP header for the Media Datagrams that are Associated with this FEC Datagram. This is illustrated in Figure 5 where the RTP headers for the protected datagrams are XORed and the relevant fields are then inserted in the FEC Header.

The Length Recovery field is computed by XOR processing of the combined length of the CSRC list, payload header, Media Payload and padding. The result of the XOR computation is inserted in the Length recovery field.

Note: In the case of media carried according to SMPTE ST 2022-6, the CSRC field is empty and its length is 0 octets, the length of the payload header is 8 octets unless extension is signaled and the length of Media Payload is 1376 octets.

7.4.2 Construction of FEC Payload

The FEC Payload shall be computed by XOR processing of the RTP payloads of the Associated Media Datagrams. The RTP payload of the Media Datagrams consists of the payload header and Media payload.

Note: The length of the Payload Header and the Media Payload are for the case of SMPTE ST 2022-6. For other applications, the length fields could be different.

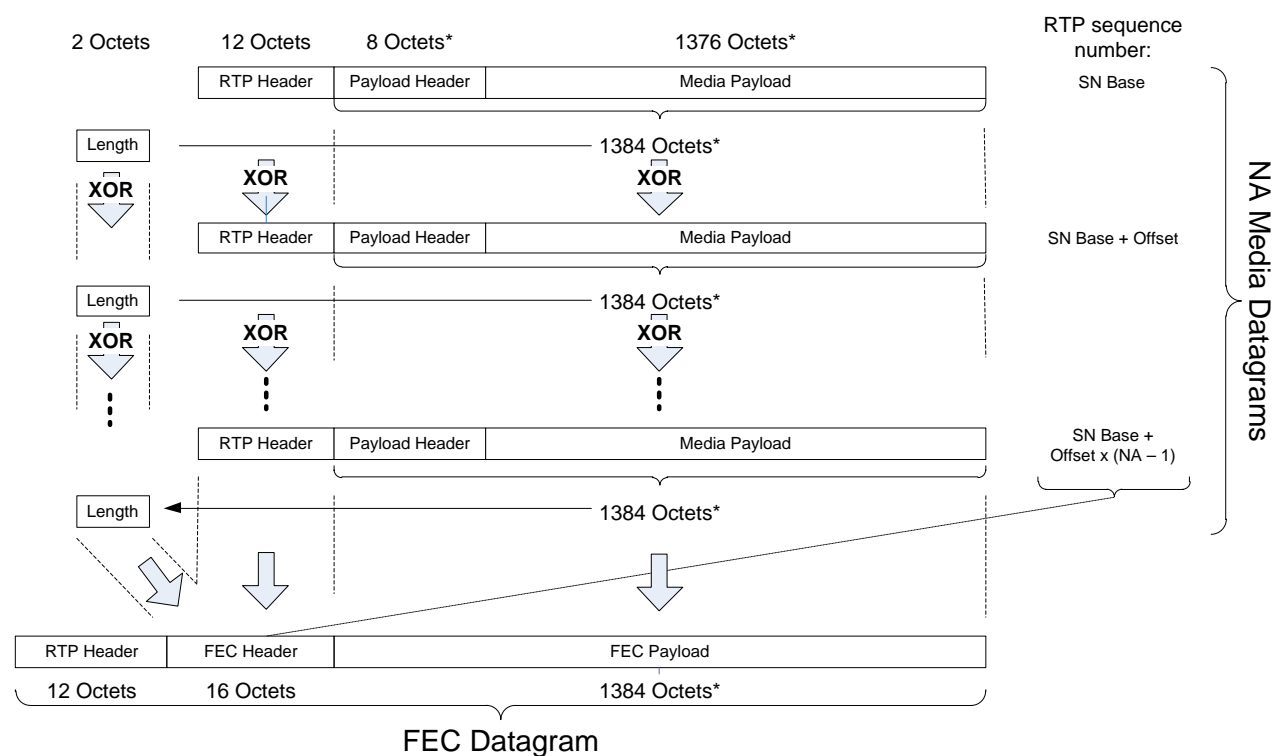


Figure 5 – Generation of 2022-5 FEC Datagram

The FEC process results in Latency, as the receiver is not able to perform datagram reconstruction until all FEC Datagrams (column and row) which contribute to the protection have been received.

7.5 FEC Traffic Shaping Issues

Senders should ensure that FEC Datagrams are interleaved with Media Datagrams in a way that avoids large changes in sending rate over time. The following constraints shall be applied in order to shape the output IP stream of the FEC system:

The row FEC Datagrams shall be sent no sooner than the last Media Datagram protected, and no later than the L^{th} Media Datagram after the last Media Datagram protected.

Column FEC Datagrams shall not be sent sooner than the L^{th} Media Datagram after the last Media Datagram protected in order to ensure that the burst loss tolerance of the system is not compromised.

Column FEC Datagrams shall be sent no later than the $(LxD)^{\text{th}}$ Media Datagram after the last Media Datagram protected to constrain the level of buffering required at the receiver.

Column FEC Datagrams shall be evenly distributed within the Media Datagram stream conforming to these constraints.

Example schemes which comply with these constraints and which intrinsically resolve Traffic Shaping issues are presented in Annexes B and C. These are examples only — any interleaving scheme which meets the requirements above shall be permitted by this Standard. There are a wide variety of valid methods for organizing the Media Datagrams to create the column FEC stream which meet the normative requirements of this standard.

It should be noted that in every case, each individual FEC Datagram indicates the base sequence number (SN-base), the offset (L) and a number of data datagrams (NA). Receivers shall observe these transmitted values in each FEC Datagram to correctly associate the FEC Datagram with the original Media Datagrams.

7.6 Re-Order Tolerance (Informative)

Datagrams traveling over IP networks are not guaranteed to arrive in the order sent. Sequence numbering is provided for each stream by RTP, which allows the receiving end equipment to correct for this effect.

If a datagram is grossly out of order, then it can be discarded, and might be corrected by the FEC scheme as if it were a lost datagram. If a system holds multiple FEC matrices, then it can tolerate re-ordering within this group of matrices.

Annex A Bibliography (Informative)

Note: All references in this document to other SMPTE documents use the current numbering style (e.g. SMPTE ST 2022-1:2007) although, during a transitional phase, the document as published (printed or PDF) may bear an older designation (such as SMPTE 2022-1-2007). Documents with the same root number (e.g. 2022-1) and publication year (e.g. 2007) are functionally identical.

IETF Standard 6, User Datagram Protocol

IETF RFC 4566, Session Description Protocol

IETF RFC 5109, RTP Payload Format for Generic Forward Error Correction

SMPTE ST 2022-1:2007, Forward Error Correction for Real-time Video/Audio Transport over IP Networks

SMPTE ST 2022-3:2010, Unidirectional Transport of Variable Bit Rate MPEG-2 Transport Streams on IP Networks

SMPTE ST 2022-6:2012, Transport of High Bit Media Signals over IP Networks (HBRMT)

IANA Ports Registry <http://www.iana.org/assignments/port-numbers>

Annex B Non Block-Aligned FEC Arrangement (Informative)

Readers are cautioned that this is an example only – any interleaving scheme which meets the requirements in this standard is allowed. There are a wide variety of valid methods for organizing the Media Datagrams and FEC Datagrams to create the column FEC stream, which meet the normative requirements of this standard.

Figure B.1 shows an example alternative arrangement for the case of (L=5, D=3). In this example, the Column FEC Datagram C0 protects Media Datagrams [0, 5, 10] while FEC Datagram C6 protects Media Datagrams [6, 11,16]. For the row FEC, FEC Datagram R0 protects Media Datagrams [0, 1, 2, 3, 4].

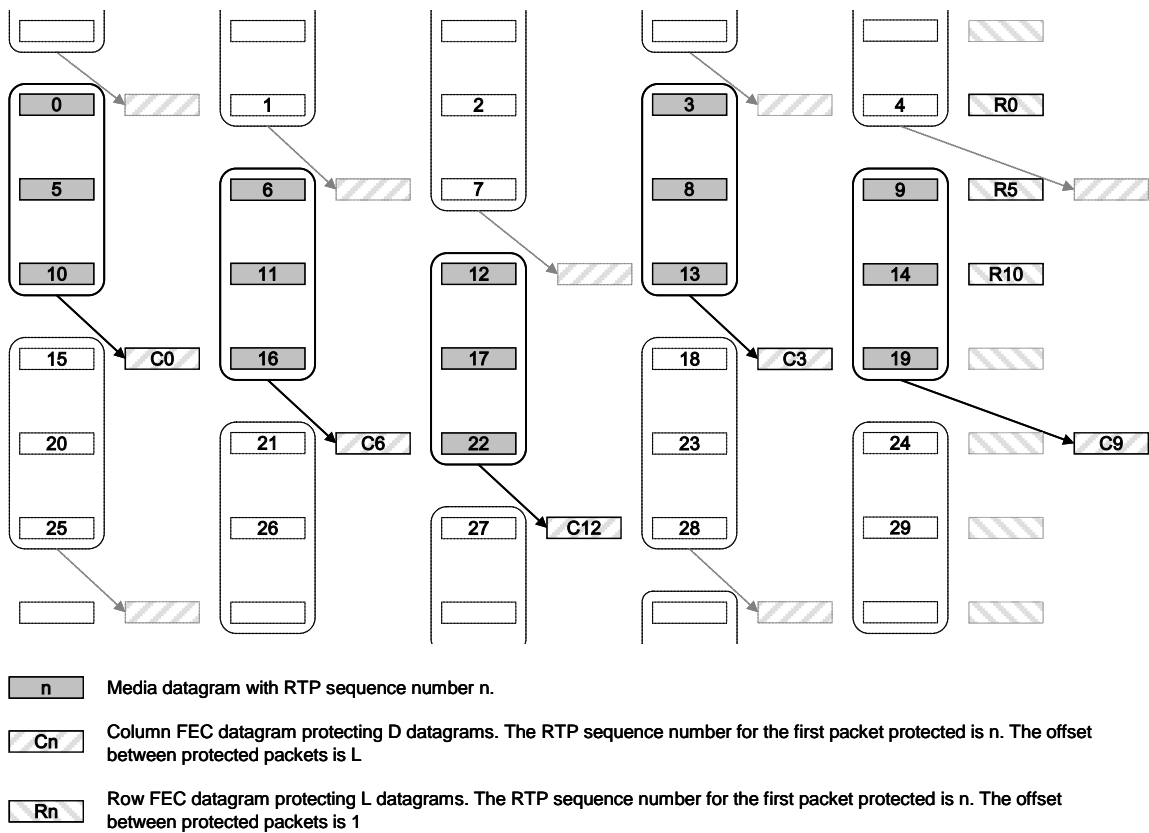


Figure B.1 – Non Block-Aligned FEC arrangement. Column and row protection shown

The goal of this arrangement is to avoid large changes in sending rate over time by ensuring that FEC Datagrams are regularly inserted among the subsequent Media Datagrams, and at the same time minimize the Latency. Using this arrangement, the Latency due to FEC is $L \times D + D$. However, it should be noted that in some cases with combined row and column FEC, it will be possible by iterative FEC process to correct lost Media Datagrams well prior to the Media Datagrams protected by the most recent FEC Datagrams.

Each column FEC Datagram is transmitted L datagram times after the last data datagram it pertains to creating a highly time-linear datagram flow on the FEC stream.

It should be noted that in every case each individual column-FEC Datagram indicates the base sequence number (SN-base), the offset (L) and a number of data datagrams (NA). Receivers can observe these transmitted values in each FEC Datagram to correctly associate the FEC Datagram with the original data-stream datagrams.

Annex C Block-Aligned FEC Arrangement (Informative)

Readers are cautioned that this is an example only – any interleaving scheme which meets the requirements in this standard is allowed. There are a wide variety of valid methods for organizing the Media Datagrams and FEC Datagrams to create the column FEC stream, which meet the normative requirements of this standard.

Figure C.1 illustrates an arrangement with $L=5$ and $D=3$. R_n datagram refers to the row FEC Datagram computed over row $\#n$ and C_n datagram refers to the column FEC Datagram computed over column $\#n$. Datagrams are sequenced in the reading direction (from left to right then from top to bottom).

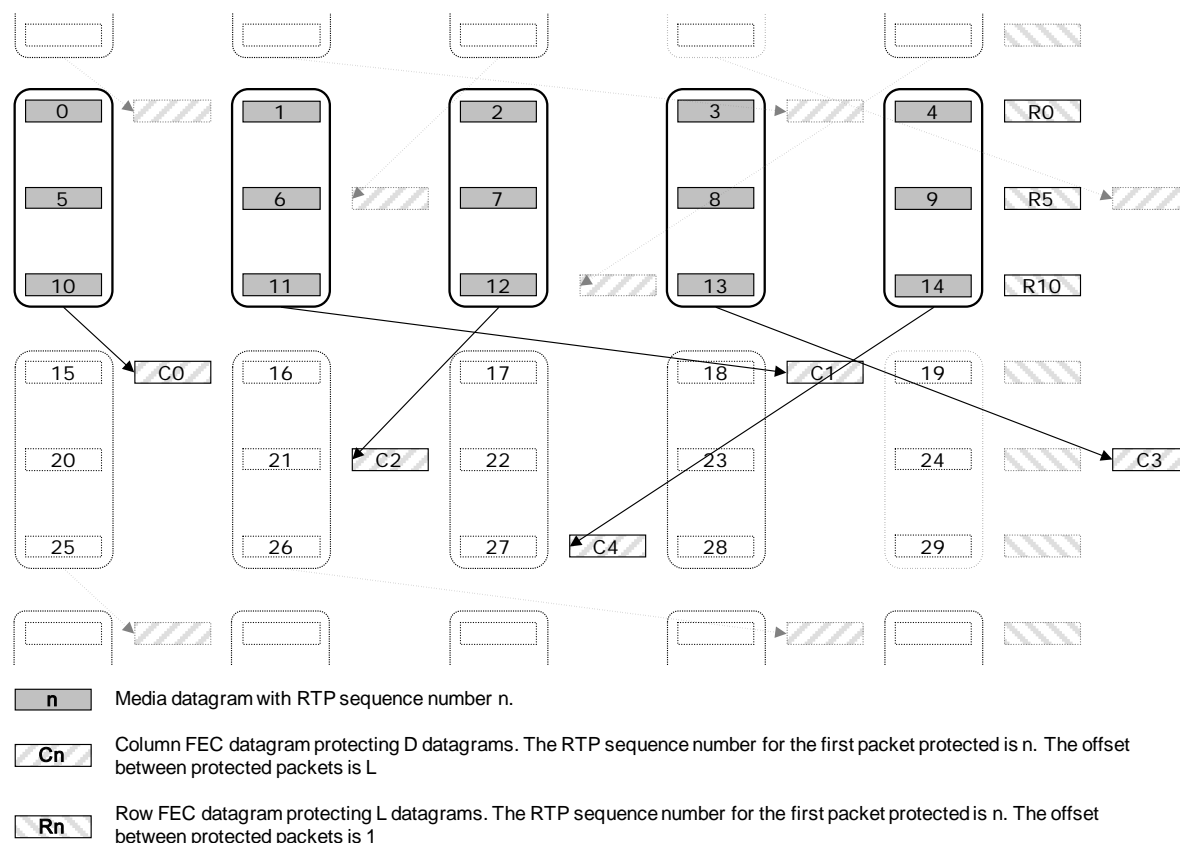


Figure C.1 – Block-Aligned FEC arrangement — Column and row protection shown

The goal of this example arrangement is to avoid large changes in sending rate over time by ensuring that FEC Datagrams are regularly inserted among data datagrams of the next matrix. This arrangement means that the Latency is approximately $2 \times L \times D - L$.

Column FEC Datagrams are sent interleaved with the Media Datagrams for the next matrix. The depth of the interleaver is chosen to be D so as to make sure that the L column FEC Datagrams will be regularly inserted among the $L \times D$ data datagrams of the next matrix. In this example, we choose to send the first column FEC immediately following the first data datagram of the next matrix so a receiver can easily detect it.

It should be noted that in every case each individual column-FEC Datagram indicates the base sequence number (SN-base), the offset (L) and a number of data datagrams (NA). Receivers can observe these transmitted values in each FEC Datagram to correctly associate the FEC Datagram with the original data-stream datagrams.

Annex D Delay (Informative)

Delay within an IP adaptation unit is bounded by a combination of the jitter tolerance buffer, the FEC system Latency, FEC processing Delay (implementation dependent) and the clock recovery mechanism (not covered by this standard). The Latency required for handling of jitter is independent of the FEC system and not included in the examples below. There are additional Delays caused by the IP network transmission, jitter tolerance buffer and IP encapsulation/de-encapsulation process.

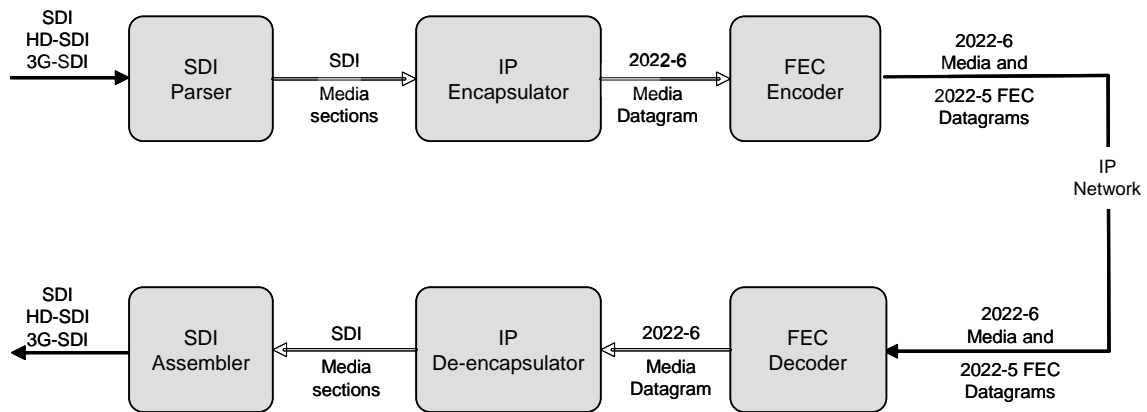
A number of professional applications have demanding round-trip Delay requirements (for example live interviews) which employ video and audio. It is possible to use this standard in live applications, depending upon the user's Latency requirements and the maximum protection required. The table below illustrates different latencies introduced by various FEC configurations.

The table that follows gives the Latency that will be introduced by the FEC process for different bit rates and FEC configurations using various combinations of both the 1D FEC code sending only column datagrams and of the non Block-Aligned FEC arrangement shown in Annex B. In this case, the FEC Latency is LxD .

The FEC Latency for the Block-Aligned FEC arrangement, Annex C, is implementation dependent on the Traffic Shaping applied. A typical value may be derived by using the formula $(2xLxD)-D$ for both 1D column only FEC as well as for 2D FEC code sending both column and row datagrams. In the case of the 2D code, it might be possible to recover multiple lost Media Datagrams by iterative FEC processing. Datagrams belonging to previous blocks may be recoverable, especially in the case of Non-Block-Aligned FEC, provided sufficient Latency is configured. Note that the number of 2D recovery iterations possible are limited by the size of the FEC Latency buffer and this may limit the recovery capability of the code.

Table D.1 – Overhead and Latency due to forward error correction only for column FEC with non-block-aligned (Annex B) and block-aligned (Annex C) FEC. The FEC processing Delay is not included, but in the general case this is negligible (Informative)

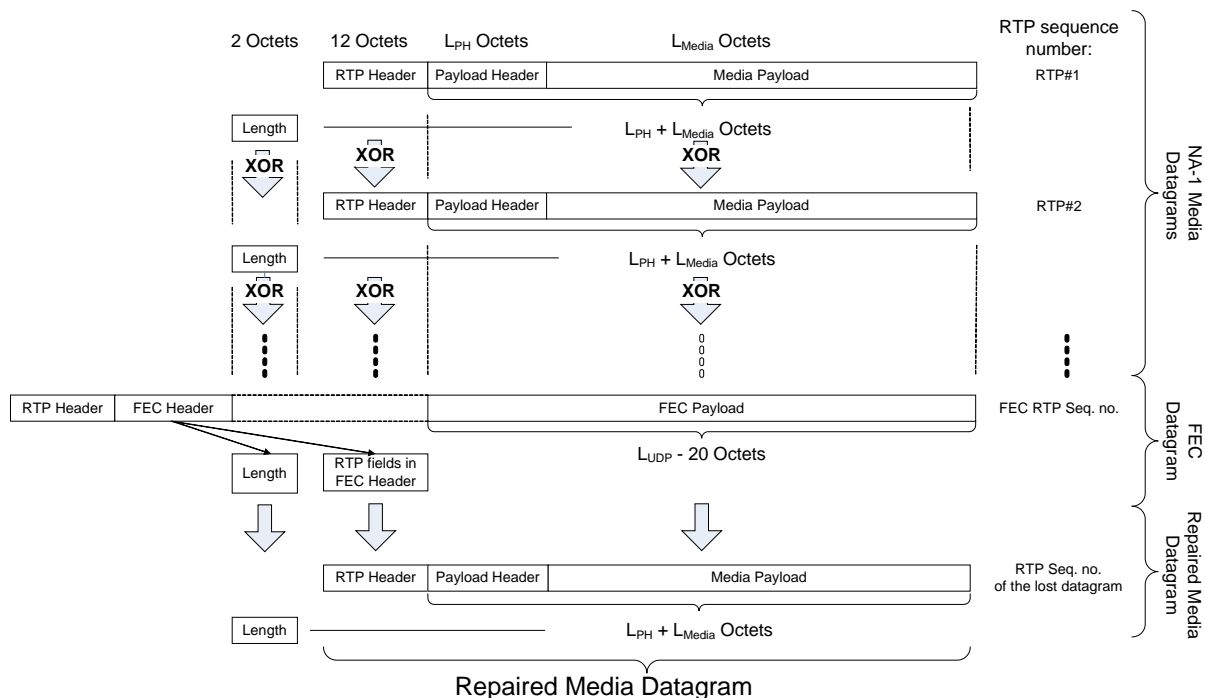
Matrix (L,D) Annex B or Annex C	Overhead	Latency			Maximum recovery
		270Mbps	1485 Mbps	2970 Mbps	
XOR (16,16) Annex B	6,25%	10.4 ms	1.9 ms	0.9 ms	16 IP datagrams
XOR (32,32) Annex B	3.1%	41.7 ms	7.6 ms	3.8 ms	32 IP datagrams
XOR (128,16) Annex B	6.25%	83.4 ms	15.2 ms	7.6 ms	128 IP datagrams
XOR (256,16) Annex B	6.25%	166.9 ms	30.3 ms	15.2 ms	256 IP datagrams
XOR (1020,8) Annex B	12.5%	332 ms	60.4 ms	30.2 ms	1020 IP datagrams
XOR (16,16) Annex C	6,25%	20.2 ms	3.7 ms	1.8 ms	16 IP datagrams
XOR (32,32) Annex C	3.1%	82.1 ms	14.9 ms	7.5 ms	32 IP datagrams
XOR (128,16) Annex C	6.25%	162 ms	29 ms	14.7 ms	128 IP datagrams
XOR (256,16) Annex C	6.25%	323 ms	58.8 ms	29.4 ms	256 IP datagrams
XOR (1020,8) Annex C	12.5%	623 ms	113 ms	56.7 ms	1020 IP datagrams

Annex E Conceptual Diagram (Informative)**Figure E.1 – Conceptual Diagram**

Annex F Datagram Recovery (Informative)

This annex illustrates one way an implementer might choose to recover lost datagrams.

Once a set of Media Datagrams and their Associated FEC Datagrams have been delivered to a receiving device, the receiving device can recover some or all of the datagrams that have been lost in transmission using the FEC data to reconstruct the missing Media Datagram within any row or any column that has lost only one datagram. Recovery of a single missing datagram can be accomplished by using the following calculation employing data from all of the available (non-missing) Media Datagrams in the same row or column and the Associated FEC Datagram. The result of this calculation will be equal to the value of the missing Media Datagram.



F.1 – An example of reconstruction of a lost datagram

Step 1. Produce a bit-wise XOR of the RTP headers of all of the available (non-missing) Media Datagrams in the desired row or column. Take this result and apply a bit-wise XOR with the most significant 96 bits of the FEC Payload. The result of this calculation is a twelve-octet recovery string.

Step 2. RTP header field for the recovered datagram is constructed as follows:

- Version (V): 2 bits
Set to 2
- Padding (P): 1 bit
Set to the value of the third bit of the recovery string
- Extension (X): 1 bit
Set to 0
- CSRC Count (CC): four bits
Set to 0

Marker (M): 1 bit

Set to the value of the eighth bit of the recovery string

Payload Type (PT): 7 bits

Set to the value of the ninth – fifteenth bits of the recovery string

Sequence number (low bits): 16 bits

Set to the value of the sequence number that would apply to the missing datagram (i.e. 1 greater than the value of the sequence number of the datagram transmitted immediately prior to the missing datagram)

Timestamp: 32 bits

Set to the value of bits 32-63 of the recovery string

SSRC: 32 bits

Copy the SSRC of the media stream that is being protected (i.e. the SSRC of the media stream to which the FEC stream is Associated)

This resulting twelve-octet field will be the header of the recovered datagram.

Step 3. Calculate the Media Payload portion of the recovered datagram taking the bit-wise XOR of the combined payload header and payload of all the available (non-missing) datagrams. Take the resulting value and XOR it with the contents of the FEC Payload. Take this result and append it to the header generated in step 2 above to complete the formation of the recovered datagram.

If more than one datagram is lost in any row or column, then any intersecting columns or rows that have only a single lost datagram may be corrected first, until there is only one remaining uncorrected datagram, which may then be corrected. This process may be applied recursively in the receiver to correct a variety of datagram loss patterns.

Note: There are no restrictions on the algorithms that can be employed by the receiver to optimize the quantity of datagrams to be corrected or the order in which datagrams are corrected.

Figure F.2 illustrates one method in which FEC data can be used to correct rows and columns that have multiple lost datagrams. Consider the following loss pattern:




Figure F.2 – Example loss pattern

This illustration shows a Block-Aligned FEC arrangement with $L=5$ and $D=4$, and a number of missing datagrams indicated with “X” marks. In this example, the receiver could begin by recovering datagram 15 by performing the datagram recovery operation on Column 0 using the values of datagrams 0, 5, 10 and C0. Next, datagrams 3, 13 and 18 could be recovered using the information in Rows 0, 10 and 15, respectively. Next, datagrams 6, 7 and 9 could be recovered using the data in Columns 1, 2 and 4, respectively. Finally, datagram 8 could then be recovered by using datagrams 5, R5, as well as the recovered datagrams 6, 7, and 9, which are all in Row 5.

Annex G Revision Notes (Informative)

The following sections have been revised:

1. Title
2. Introduction
3. Section 1 Scope
4. Section 3 Normative References
5. Section 6.2 RTP Header for FEC Datagrams
6. Section 7.1 FEC Datagram Arrangement
7. Section 7.2 FEC Buffer Overhead and Latency Implications
8. Section 7.4.1 Construction of FEC Header
9. Section 7.4.2 Construction of FEC Payload
10. Annex A Bibliography (updated)