Video Services Forum (VSF)
Technical Recommendation TR 10-06

Internet Protocol Media Experience (IPMX)
Forward Error Correction (FEC)

August 7, 2023

INTELLECTUAL PROPERTY RIGHTS

THIS RECOMMENDATION IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NONINFRINGEMENT IS EXPRESSLY DISCLAIMED. ANY USE OF THIS RECOMMENDATION SHALL BE MADE ENTIRELY AT THE IMPLEMENTER'S OWN RISK, AND NEITHER THE FORUM, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY MPLEMENTER OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS RECOMMENDATION.

LIMITATION OF LIABILITY

VSF SHALL NOT BE LIABLE FOR ANY AND ALL DAMAGES, DIRECT OR INDIRECT, ARISING FROM OR RELATING TO ANY USE OF THE CONTENTS CONTAINED HEREIN, INCLUDING WITHOUT LIMITATION ANY AND ALL INDIRECT, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES (INCLUDING DAMAGES FOR LOSS OF BUSINESS, LOSS OF PROFITS, LITIGATION, OR THE LIKE), WHETHER BASED UPON BREACH OF CONTRACT, BREACH OF WARRANTY, TORT (INCLUDING NEGLIGENCE), PRODUCT LIABILITY OR OTHERWISE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE FOREGOING NEGATION OF DAMAGES IS A FUNDAMENTAL ELEMENT OF THE USE OF THE CONTENTS HEREOF, AND THESE CONTENTS WOULD NOT BE PUBLISHED BY VSF WITHOUT SUCH LIMITATIONS.

Executive Summary

IPMX was created to foster the adoption of open standards-based protocols for interoperability over IP in the media and entertainment (M&E) and professional audio/video industries.

This Technical Recommendation (TR) describes a Forward Error Correction profile of the SMPTE ST2022-5 standard for use with IPMX.

1 Table of Contents

2 Introduction (Informative)

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.

Some ST2110 environments rely on full redundancy 2022-7 to provide increase robustness. While this protection method can be very robust, it comes at the expense of typically requiring double the bandwidth of an unprotected stream. This is not always realistic in an IPMX environment. Therefore, even though IPMX supports 2022-7 it also supports an FEC scheme which does not require a second parallel network path.

The FEC scheme is designed to correct for individual or short (depending on profile) burst packet losses.

The FEC profile defined in this Technical Recommendation (TR) complies with the ST 2022-5 standard.

2.1 Contributors

The following individuals participated in the Video Services Forum IPMX working group that developed this TR.

Aaron Doughten (Sencore)	Charles Buysschaert (Intopix)	Jean-Baptiste Lorent (IntoPIX)	Phil Nguyen (Nextera)
Alain Bouchard (Matrox)	Chris Lapp (Cisco)	Jed Deame (Nextera Video)	Prinyar Boon (Phabrix)
Albert Faust (Arista)	Clark Williams (Christie Digital)	JJ Eynon (CNN)	Raul Diaz (Intel)
Andre Testa (Matrox)	Daniel BOUQUET (Analogway)	John Belstner (Intel)	Raymond Hermans (Adeas)
Andreas Hildebrand (ALC NetworX)	Danny Pierini (Matrox)	John Dale (Media Links)	Robert Welch (Arista)
Andrew Starks (Macnica)	David Chiappini (Matrox)	John Fletcher (BBC)	Ron Stites (Macnica)
Antoine Hermans (Adeas)	David Mitchinson (Appear TV)	Karl Johnson (Christie Digital)	Tadahiro Watanabe (Macnica)
Arnaud Germain (Intopix)	Gerard Phillips (Arista)	Karl Paulsen (Diversified)	Teiji Kubota (Macnica)
Ben Cope (Intel)	Greg Schlechter (Intel)	Marc Levy (Macnica)	Thomas True (NVIDIA)
Brad Gilmer (VSF)	Greg Stigall (Warner Media)	Mike Boucke (AJA)	Tim Bruylants (intopix)
Bob Ruhl (VSF)	Jack Douglass (PacketStorm)	Paulo Francisco (Evertz)	Wes Simpson (LearnIPvideo)
Cassidy Phillips (Imagine)	Jean Lapierre (Matrox)	Peter Brightwell (BBC)	

2.2 About the Video Services Forum

The Video Services Forum, Inc. (www.videoservicesforum.org) is an international association dedicated to video transport technologies, interoperability, quality metrics and education. The VSF is composed of service providers, users and manufacturers. The organization's activities include:

- providing forums to identify issues involving the development, engineering, installation, testing and maintenance of audio and video services;
- exchanging non-proprietary information to promote the development of video transport service technology and to foster resolution of issues common to the video services industry;
- identification of video services applications and educational services utilizing video transport services;
- promoting interoperability and encouraging technical standards for national and international standards bodies.

The VSF is an association incorporated under the Not For Profit Corporation Law of the State of New York. Membership is open to businesses, public sector organizations and individuals worldwide. For more information on the Video Services Forum or this document, please call +1 929-279-1995 or e-mail opsmgr@videoservicesforum.org.

3 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 and 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.

4 Normative References

The following standards contain provisions which, through reference in this text, constitute provisions of this standard. At the time of publication, the editions indicated were 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 standards indicated below.

- IETF RFC 3550, RTP: A Transport Protocol for Real-Time Applications
- SMPTE ST 2022-5-2013 Forward Error Correction for Transport of High Bit Rate Media Signals over IP Networks (HBRMT)
- SMPTE ST 2110-22:2022 Professional Media over Managed IP Networks: Constant Bit-rate Compressed Video

5 Abbreviations

See the ST2022-5 Standard for a list of abbreviations.

6 Definitions

For the purposes of this document, the definitions in the ST2022-5 standard apply to this document in addition to the definitions listed in this section.

6.1 High Bandwidth Flow

A High Bandwidth Flow is defined as having a packet rate of at least 32 packets per millisecond.

6.2 Low Bandwidth Flow

A Low Bandwidth Flow is defined as having a bandwidth below 32 packets per millisecond.

6.3 High Bandwidth FEC Matrix

A 16x2 (L=2, D=16) FEC Matrix used for high bandwidth media flows.

6.4 Low Bandwidth FEC Matrix

A 1x1 (L=1, D=1) FEC Matrix used for low bandwidth media flows.

7 IPMX Forward Error Correction

The ST 2022-5 standard defines a FEC scheme with a broad possible range of parameters. The profiles defined in this TR constrains the ST 2022-5 FEC scheme to a limited set of those parameters.

IPMX senders and receivers are not required to support FEC, but if they do, they shall implement FEC in compliance with this TR.

IPMX Receivers shall tolerate the presence of FEC streams, whether they implement FEC or not.

FEC packet flows shall use the same IP address as the media flow they are protecting.

The UDP port for column FEC shall be two port numbers higher than the port used for the media flow.

The UDP port for row FEC shall be four port numbers higher than the port used for the media flow.

This TR defines profiles that shall be used for High Bandwidth Flows such as video, and Low Bandwidth Flows such as audio. Additional profiles may be added in future revisions.

7.1 IPMX FEC Profile A

High Bandwidth Flows using IPMX FEC Profile A shall use a fixed, block aligned FEC matrix that has 2 columns (L=2) and 16 rows (D=16)

Low Bandwidth Flows using IPMX FEC Profile A shall use a fixed, block aligned FEC matrix that has 1 column (L=1) by 1 row (D=1)

IPMX FEC Profile A uses Level A (column only) FEC.

Media datagrams used to calculate a particular FEC datagram are considered to be associated with that FEC datagram. IPMX Senders and Receivers shall ensure that the media and FEC datagrams are associated with each other.

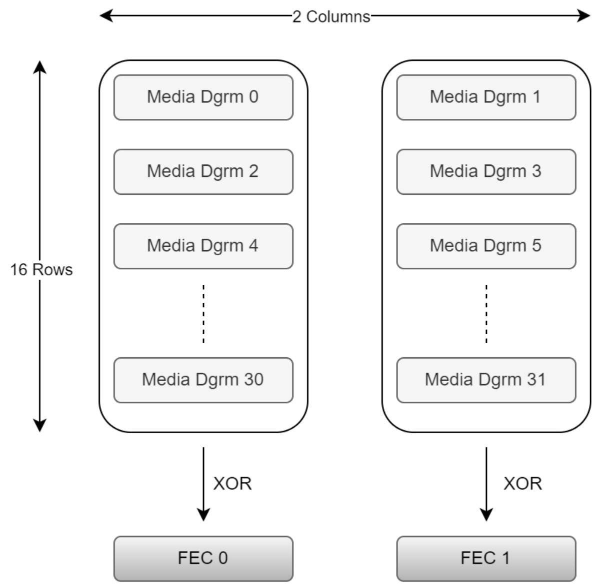

Figure 1: IPMX FEC Profile A unterminated High Bandwidth FEC matrix

In Figure 1, the FEC 0 datagram shall be calculated by XORing Media Datagrams 0, 2, 4, …, 30. The FEC 1 datagram shall be calculated by XORing Media Datagrams 1, 3, 5, …, 31.

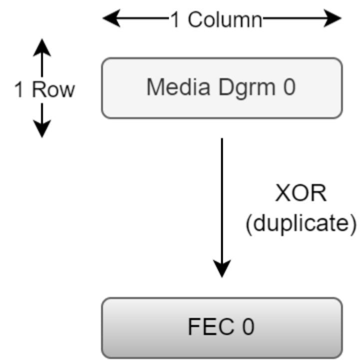

Figure 2: IPMX FEC Profile A Low Bandwidth FEC matrix

In Figure 2, the only FEC datagram (FEC 0) shall be an XOR with 0 of the only media datagram in the 1x1 matrix – which is effectively a redundant packet.

7.2 FEC Packet Traffic Shaping

FEC packets shall be evenly spaced and delayed from the associated media datagrams. The reasons for this are:

- A network dropout is less likely to cause the FEC datagram and its associated media datagrams to be lost.
- A constant packet rate on the FEC traffic flow through the IP network is desirable.

Traffic shaping for FEC matrices shall be optimized for low latency operation.

- For the IPMX FEC Profile A High Bandwidth matrix, FEC packet 0 shall be transmitted after datagram 2, and shall be transmitted before datagram 5 of the following matrix. FEC packet 1 shall be transmitted after datagram 18, and shall be transmitted before datagram 21 of the following the matrix.
- For the IPMX FEC Profile A Low Bandwidth matrix, FEC packet 0 shall be transmitted at least 100us after datagram 0 of the current matrix, and before datagram 0 of the following matrix.

FEC packet timing implementations should be based on counting media clocks from the actual transmit time of the first datagram, not from the scheduled transmission of subsequent media datagrams. This implementation method ensures that FEC packet scheduling is consistent during gaps in media packet transmission, for example during vertical blanking intervals.

Implementers may include a parameter in the SDP file specifying additional latency as a positive integer number of microseconds. (Default is 0). The relevant parameters are shown below.

FEC_ADD_LATENCY_VIDEO <time in microseconds>

FEC_ADD_LATENCY_AUDIO <time in microseconds>

7.3 FEC Latency (Informative)

A recovered datagram 1 (second datagram) in an IPMX FEC Profile A high bandwidth FEC matrix will only be available sometime after receipt of the FEC 1 datagram. The FEC 1 datagram would arrive after 50 media datagrams plus any specified additional latency, so IPMX FEC Profile A high bandwidth FEC adds at least 50 media datagrams of system latency.

The latency time for high bandwidth IPMX FEC Profile A FEC is calculated as follows:

Let B=bandwidth of media in bits per second

Let S=media payload size per packet in octets

Let P=processing time for FEC correction

Let AL = Additional latency specified in the SDP file

$$latency = s * 8 * \frac{1}{b} * 50 + P + AL$$

The latency time for low bandwidth IPMX FEC Profile A FEC is 100us plus any additional latency specified in the SDP file:

$$latency = 100us + P + AL$$

7.4 Partial FEC Matrix

High Bandwidth Flows with timing that follows the gapped packet read schedule (as defined by ST2110-21) shall terminate the FEC matrix at the end of a frame or field, that is, at the start of the discontinuity. Upon termination, a partial FEC matrix with the available data shall be generated.

If FEC is used with variable rate media data (for example a variable rate video codec), discontinuities may occur during the video frame or field. Whenever a timeout (as defined in section 7.4.1) occurs, a partial matrix shall be generated.

- When an FEC matrix is terminated, all FEC packets for the partial matrix shall be transmitted, even if a field contains no datagrams. Terminated matrices are only applicable to High Bandwidth FEC Matrices.
- When a matrix contains at least 1 datagram, both FEC 0 and FEC 1 shall be generated and transmitted. It is possible that FEC 1 will have zero associated datagrams.
- When there is a gap in datagrams that exceeds a matrix time, FEC datagrams for empty matrices shall not be transmitted. In that case, there will be gaps in the FEC flow that correspond to gaps in the media flow.
- The FEC Receiver shall determine matrix size using the “Number of Media Datagrams Associated” (NA) fields in the FEC datagram. For partial FEC matrices, the NA field

shall indicate the number of media datagrams in the matrix column associated with the received FEC datagram. Whenever the NA field indicates a smaller number of datagrams than the maximum specified by the profile, the FEC packet is covering a column of a partial matrix.

- Senders shall populate the SNBase field with the sequence number of the first RTP packet of the FEC field.

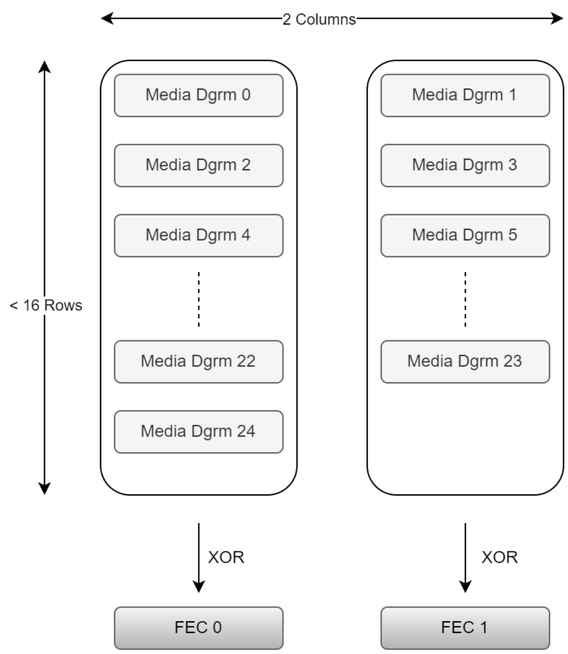

Figure 3: Example of a Partial FEC Matrix

In the example shown in Figure 3, a matrix is terminated at media datagram 24. FEC 0 and FEC 1 datagrams are generated from a smaller number of associated media datagrams and transmitted at the time they would have been transmitted for full matrices.

7.4.1 Partial FEC due to Timeout

In variable rate high bitrate media flows, large gaps in data flow may occur. In this case, a timeout shall be implemented to terminate FEC matrices.

The timeout period shall begin at the start of the first datagram of the matrix and be equal to the time period of a complete matrix (32 datagrams for IPMX FEC Profile A). If a matrix is not completed during the timeout period, a partial matrix shall be sent at the end of the timeout period as described in paragraph 7.4.

The receiver shall also implement FEC timeout, so that received datagrams are not held up for further processing should an FEC datagram associated with those media datagrams be lost by the switched network.

7.5 FEC Packets

FEC Packets are RTP packets carrying the Parity Vectors, and are sent as an independent RTP stream. Refer to the ST 2022-5 standard for definitions of the FEC Header Format and FEC Payload Formats as well as information about construction of FEC data.

For the IPMX FEC Profile A High Bitrate case, the Offset field shall be 2, since the row length (L field) is 2. The NA field shall be 16 except for the last FEC matrix of a video frame or field, or the last FEC matrix before a timeout occurs (the partial FEC matrix cases).

7.6 SDP Field for FEC stream

The following SDP field shall be used to indicate that a FEC flow accompanies the media flow, and to identify the profile of the FEC flow (currently only IPMX FEC Profile A is specified).

This is a parameter within the <format specific parameters> section

FECPROFILE=profile-a;

See Appendix B for an example.

7.7 Packet Sizes

The ST 2022-5 Standard specifies a 16 octet FEC Header. This FEC Header results in FEC datagrams being 16 octets larger than their associated media datagrams.

Note: Implementations may need to use a reduced packet payload size due to the additional 16 octet header. This will ensure that the maximum allowed IP packet size is not exceeded.

Appendix A (Informative)

ST 2022-5 Forward Error Correction

The ST 2022-5 Forward Error Correction algorithm provides error correction for a combination of burst and single packet losses. FEC can recover a single lost packet within its field.

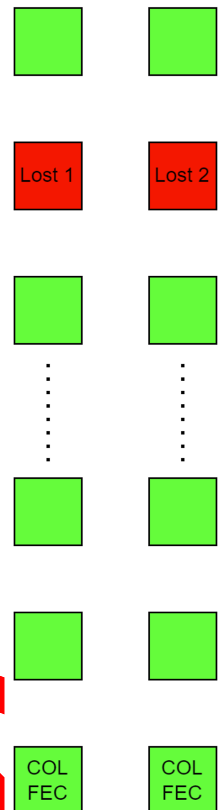

Figure 4 IPMX Profile A FEC Example

In this example:

- Lost Packets 1-2 are corrected by Column FEC

The FEC example shows how a burst loss of 2 packets can be recovered.

It should be noted that since Column FEC packets are sent during the matrix time following the datagrams, column FEC adds at least 2 matrix times worth of latency.

Partial FEC Matrix

The description of the matric size fields ("Number of Media Datagrams Associated (NA)" and "Offset") in the FEC header describe matrix size parameter values. For partial matrices, the FEC packets are populated with NA fields that correspond to the actual number of datagrams associated with that FEC datagram.

The following example shows a terminated matrix at the end of a video frame:

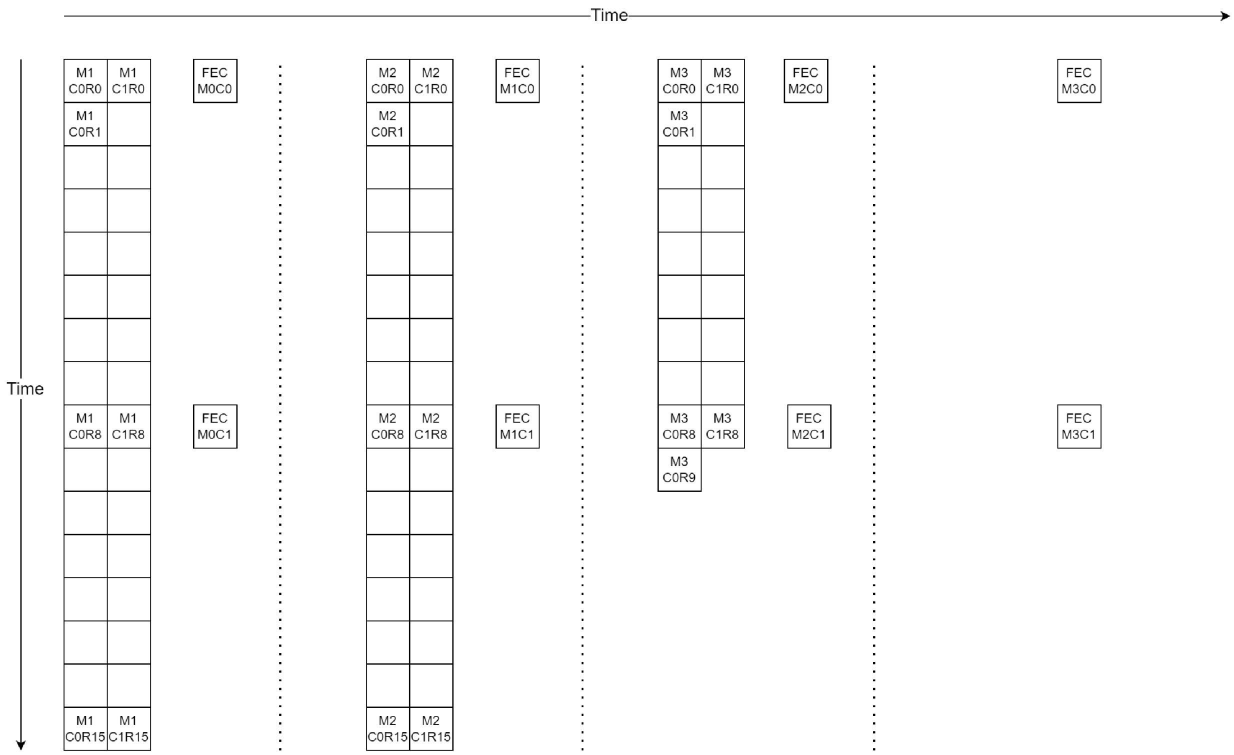

Figure 5: Partial FEC Example

Figure 6 shows video datagrams with interleaved column FEC packets. The relative time of packets follows the convention of a written page, with columns separated by dotted vertical lines. Packet ordering is {Matrix 1, Column 0, Row 0}, {Matrix 1, Column 1, Row 0}…{Matrix 1, Column 0, row 7}, {Matrix 1, Column 1, Row 7}, {FEC: Matrix 0, Column 0}…

Since datagrams stop after {Matrix 2, Column 0, Row 9}, FEC packets will stop after {FEC: Matrix 2, Column 1} until datagram transmission resumes.

Where there is vertical blanking, no video datagrams are present – but timed FEC packets that cover datagrams that were transmitted are still present. The packet labels show how datagrams are arranged into matrices. Column FEC packets are calculated over columns of datagrams to protect that column and can correct a single lost packet within the column they protect.

Note that Column FEC packet headers identify the first packet sequence number of the column they protect, and the number of columns they protect. The number of columns in the matrix is also a field in the FEC header. There is no matrix position information in the datagrams – only an RTP sequence number.

Compressed 4K Video Example:

1. Video mode is 3840*2160 at 60 frames per second.
2. Video is compressed to 1.6 bits per pixel
3. Compressed video frame size is 1,658,880 bytes
4. Assuming a packet payload size of 1280 bytes, there are 1,296 packets per video frame
5. The matrix size is 2 rows by 16 columns (32 packets per matrix)
6. There are 40 full 32 packet matrices, and one partial matrix with 16 packets per video frame

Per section 7.3, latency is calculated as follows:

Let B=bandwidth of media in bits per second (3840 * 2250[a] * 1.6 * 60 = 829.44 Mbps)

Let S=media payload size per packet in octets (1280)

Let P=processing time for FEC correction (assume 20us)

Let AL = Additional latency specified in the SDP file (assume 50us)

Note a: The number of lines in the full raster rather than active video is used to calculate bandwidth during active video.

$$latency = s * 8 * \frac{1}{b} * 50 + P + AL$$

Latency = 687.28usAppendix B (Informative)

The following is an example of a SDP file segment showing IPMX FEC Profile A:

```
v=0
o=- 1618351493884125000 1618351539175204201 IN IP4 25.25.30.151
s=IP video OUT 1
t=0 0
m=video 10000 RTP/AVP 96
c=IN IP4 239.20.0.1/128
a=source-filter: incl IN IP4 239.20.0.1 25.25.30.151
a=rtpmap:96 raw/90000
a=fmtp:96 sampling=YCbCr-4:2:2; width=1920; height=1080;
exactframerate=60000/1001; depth=10; TCS=SDR; colorimetry=BT709;
PM=2110GPM; SSN=ST2110-20:2017; TP=2110TPN; FECPROFILE=profile-
a; IPMX;  measuredpixclk=1485501040; vtotal=1125; htotal=2200
a=ts-refclk:localmac=00-20-FC-32-2F-40
a=mediaclk:sender
```