Transmission of Professional MPEG-2 Transport Streams over IP Networks

1 Background

The Pro-MPEG Forum has already published guidelines in Code of Practice #2 to aid interoperability for the transmission of professional video over a variety of networks. The Operating Points described there have proven useful targets for equipment manufacturers and service providers, as well as providing a set of test-points for formal interoperability tests. Further information is available on the Pro-MPEG web site www.pro-mpeg.org.

In recent years it has become apparent that IP-based networks will become increasingly important for delivery of professional content. In the interests of interoperability, common approaches to the issues presented by such networks are desirable. End devices created by various manufacturers need to operate correctly with each other, and with networks using equipment from various vendors.

Pro-MPEG Forum Wide Area Networking Working has defined a suitable set of solutions to the issues for compressed video. The approach taken by standards bodies and industry organisations has been included in the discussions leading to this publication.

2 Scope

The application space being addressed is limited to Contribution and Primary Distribution applications. A standard for the distribution of Video over IP networks for end user distribution has been produced by the DVB-IPI group, and it is recommended that their practices are followed for such applications.

This code of practice is for MPEG-2 Transport Streams only. The Working Group will look at producing further codes of practice to cover additional formats, for example, uncompressed Serial Digital video over IP, the various forms of SDTI, all DV and derivative formats and compressed High Definition formats.

With these future needs in mind, this Code of Practice will follow practices that can be re-used for future formats where practicable.

To aid clarity to the code of practice the following words have been used:

- "shall" – mandatory to be compliant with this code of practice
- "should" – optional to be compliant but improves interoperability if adopted

This code of practice is released under the Intellectual Property Terms of the Pro-MPEG Forum WAN Group published on the internet at www.pro-mpeg.org.

2.1 Input Format

For professional applications, MPEG-2 using the 4:2:2P@ML profile is currently the normal practice. However, transport streams containing other forms of MPEG-2 and newer MPEG standards encapsulated as MPEG-2 transport streams are also supported by this Code of Practice.

2.2 System model

Not all parameters for the sender, network and the receiver are bounded by the recommendations made in this document. Further study is required to more clearly bound the overall performance of the complete system.
Models already exist for the MPEG Encoder and Decoder, so there is no requirement to repeat these, beyond ensuring that no additional limitations are imposed on the Encoder or Decoder design.

2.3 Multicast support

Multicast needs to be supported both for a transmitting edge device sending to multiple endpoints, and for a receiving edge device to be able to receive a multicast transmission.
Detail on the multicast scenarios is given in section 5.

3 User Performance Requirements

The performance which is needed, and which must be provided by systems carrying professional video over IP, is inextricably linked to the definition of professional video.

3.1 Baseline network requirements

In order for a system supporting this code of practice to function correctly it is necessary for the bandwidth available in the network to always meet or exceed that required by the IP stream generated by the system.
The total error rate on the network must be low enough to allow the error correction system referenced from section 4.5 to produce an output stream that can be successfully decoded.
The network jitter must be low enough to satisfy the requirements in section 4.7.
The network latency performance is described in detail in section 4.8.

3.2 Transport Stream bitrate

There are going to be limits on the upper and lower bound of bit rates that will be transportable by equipment following this Code of Practice. As a minimum though, it is recommended that equipment is able to support some of the transport stream rates from the operating points listed in the Code of Practice #2. The full choice of which operating points to support will depend upon the application that the particular equipment is designed to support.

For interoperability purposes all equipment shall be capable of supporting the transport stream rates required by the 5.5-A, 7.5-1:0:0, 14-0:1:0, 30-0:0:1 and 50-0:0:2 points defined in Code of Practice #2 revision 2.

3.3 System output error rate

The permissible system output error rate for professional applications is usually more stringent than that for domestic or industrial applications. The precise value is a topic for negotiation between customer (broadcaster) and network provider.
The FEC scheme proposed in the Code of Practice is configurable, so can support a range of un-corrected error rates across a range of network conditions.

4 Transmission Protocols

4.1 IP carriage issues

The size of the output IP packet from a transmitting device shall be limited so that IP fragmentation does not occur at the output of the device. The IP ‘don’t fragment’ bit shall be set. As end-point devices will typically be connected to Ethernet style networks, this limits the maximum transmission unit (MTU) to 1500 bytes. The MTU on links between intermediate nodes in the network may be lower than this, so care shall be taken to ensure that IP packets are not lost due to the ‘don’t fragment’ requirement.

4.2 RTP/UDP/IP Mapping

The use of RTP is required, as it provides a standard header for the packets. The RFC2250 mapping shall be used as it provides a suitable mapping for MPEG-2 transport streams. Issues for the carriage of 204 byte packets are considered later in this document.
For interoperability, it is required that equipment only depends on the main RTP communication between sending and receiving units, and does not require any additional communication, including but not limited to RTCP.

If equipment manufacturers wish to make use of such additional communication to improve the operation of their units then they are not forbidden from adding this, provided that the equipment can function correctly in its absence.

The following additional restrictions on RFC3550 and RFC2250 will be adopted:

- The Padding (P) bit shall be set to zero. This defines that there will be no padding bytes in the payload.
- The Extension (X) bit shall be set to zero. This defines that there will be no header extension(s) present.
- The Marker (M) bit shall always be set to zero. This defines that there are no discontinuities in the stream during a session. For the purposes of the testing intended this is not a major limitation.
- The CSRC count (CC) field shall be set to zero. This defines that there are no entries in the CSRC (Contributing SouRCes list)
- The value of the SSRC field will not be used at the receiver, so the sender is free to assign this according to their current practice.
- There is no requirement for the initial sequence number to be randomly assigned, as suggested in RFC3550.

4.3 TS Packets per IP Packet

Given the considerations in section 4.1, the range of possible MPEG packets per IP packet is from 1 to 7. Long-length packets are undesirable due to the excessive impact (lost data) from losing each IP packet. Short packets cause a high overhead, so the value chosen will be a compromise between these factors. For simplicity, the value chosen should be kept constant for the duration of a send-receive session.

As a minimum, equipment should support 1, 4 and 7 transport stream packets, but may support other values.

4.4 TS Packet length (188/204)

The required minimum is that all equipment shall be capable of operating in 188 byte mode.
In more complex network designs the support of the transparent carriage of 204 byte TS packets may be required for end to end integrity checking of the whole network. Currently RFC2250 does not explicitly mention 204 byte packets, so many existing implementations will only support 188 byte packets.
A receiver that can support both 188 and 204 byte TS packets will use the received IP packet length to determine whether 188 or 204 byte packets are present – if 188 byte packets are present then the RTP payload will divide exactly by 188 and not by 204, and vice versa for 204 byte packets.
The TS packet length shall be kept constant.

4.5 FEC Scheme

4.5.1 Background

Errors are generally not acceptable, therefore support for some sort of FEC scheme is mandatory. The use of that scheme is recommended, but there are applications where occasional errors are preferable to the overhead of the FEC, so manufacturers may support a non FEC mode.

The biggest issue with FEC systems on IP networks is that, because of the UDP checksums, channel bit errors get translated into packet losses. In addition to this, buffer and re-route issues cause burst packet losses. The combination of packet losses from the three sources – gross reordering, bit-error induced losses and burst losses needs to be low enough so that the FEC scheme is not broken more than the negotiated error rate. Because any bit errors cause the packet to be discarded there is no requirement for an error correction scheme that can handle errored packets – every packet will either arrive correct or not at all.

An RTP payload format for Generic Forward Error Correction Packets has been defined in the RFC 2733 to enable error correction of real time media. This standard allows the use of traditional error correcting codes. A major advantage of this scheme is that it can be used with any video formats standards (MPEG, SDI, SDTI, ...) as long as it is encapsulated in a RTP packet. However, this standard limits the scope of packets used to generate the Forward Error Correction payload, to 24 consecutive packets.

4.5.2 FEC packet arrangement

To recover burst loss, an extension to the existing RFC is proposed. The same traditional error correcting codes are applied to non-consecutive media packets that can be spaced of more than 24 packets. Each FEC packet is associated to packets periodically selected. Therefore, consecutive RTP packets can be recovered from consecutive FEC packets. The process is detailed in Figure 1.

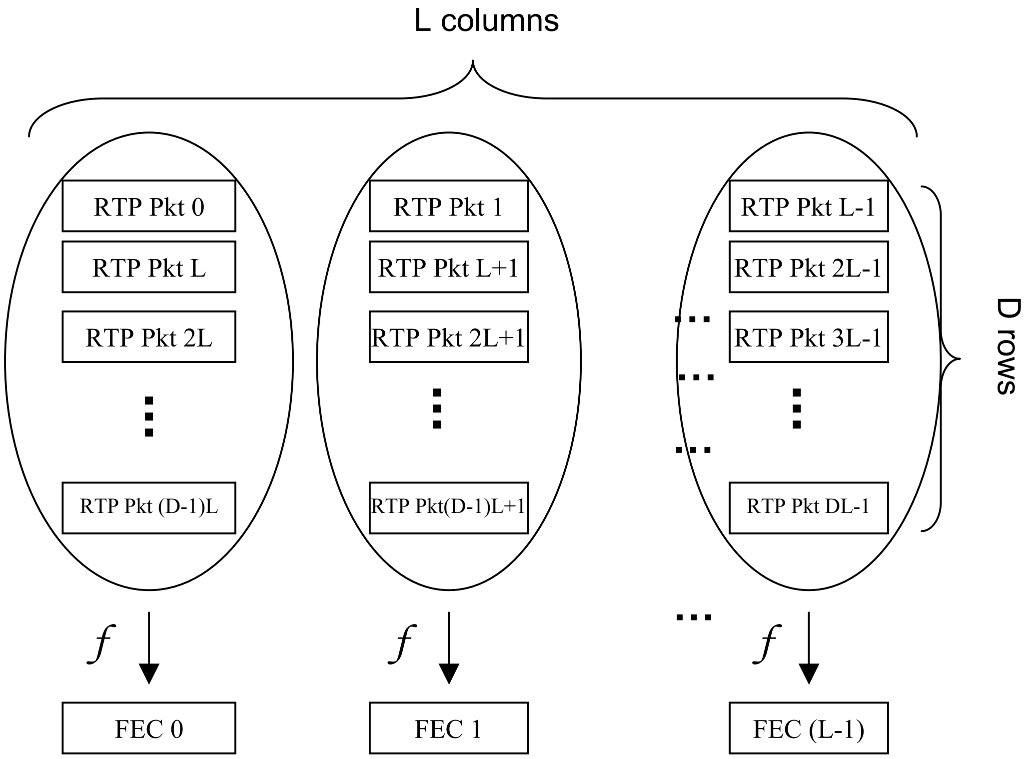

Fig.1 Encoding scheme

In Figure 1, the encoding scheme is schematized for $L*D$ media packets. The period chosen is L. Thus the payload of the k^{th} FEC packet is computed based on the D packets numbered $nL+k$ $(0 \leq n \leq D - 1)$.

The alignment of the columns is for illustration. Implementations may use this alignment for simplicity, but there are some potential advantages to be gained by offsetting the columns - see Informative Annex A. This means that receiving devices shall not make any assumptions about the relationship between FEC packets beyond those that are explicitly specified.

The main advantage of this scheme is the burst error correction capacity. The error correcting function chosen is XOR which has the ability to recover any one lost packet . If a one dimensional scheme based on XOR is used (i.e. applied to *D* consecutive packets), a burst error of two or more lost packets is not recoverable. However, if the two dimensional scheme is used, the recoverability is greatly improved, since it can recover up to *L* consecutive packets.

Though this scheme is very robust to bursts of packet losses (it corrects *L* consecutive packets lost), if only *2* packets are lost and these packets are in the same column, there is no way to correct these losses. It is therefore recommended that two simultaneous FEC streams should be supported, which will allow for a far higher correction capability, at the expense of increased overhead. These FEC streams shall be carried on separate UDP ports, to allow them to have separate sequence number handling, and to maintain backward compatibility with previous implementations that only supported a single (column) FEC stream.

The first port shall carry the column FEC stream and the second port shall carry the row FEC stream.

Obviously, for the second stream to be useful, it must have different dimensions from the first. The structure for the second FEC stream shall have OFFSET set to 1, and should have NA equal to the L parameter of the first stream. This will effectively produce an FEC structure as shown in figure 2, where the packets labeled RTP are the media packets, the packets labeled FEC are the first FEC stream packets, and the packets labeled FEC' are the second FEC stream packets.

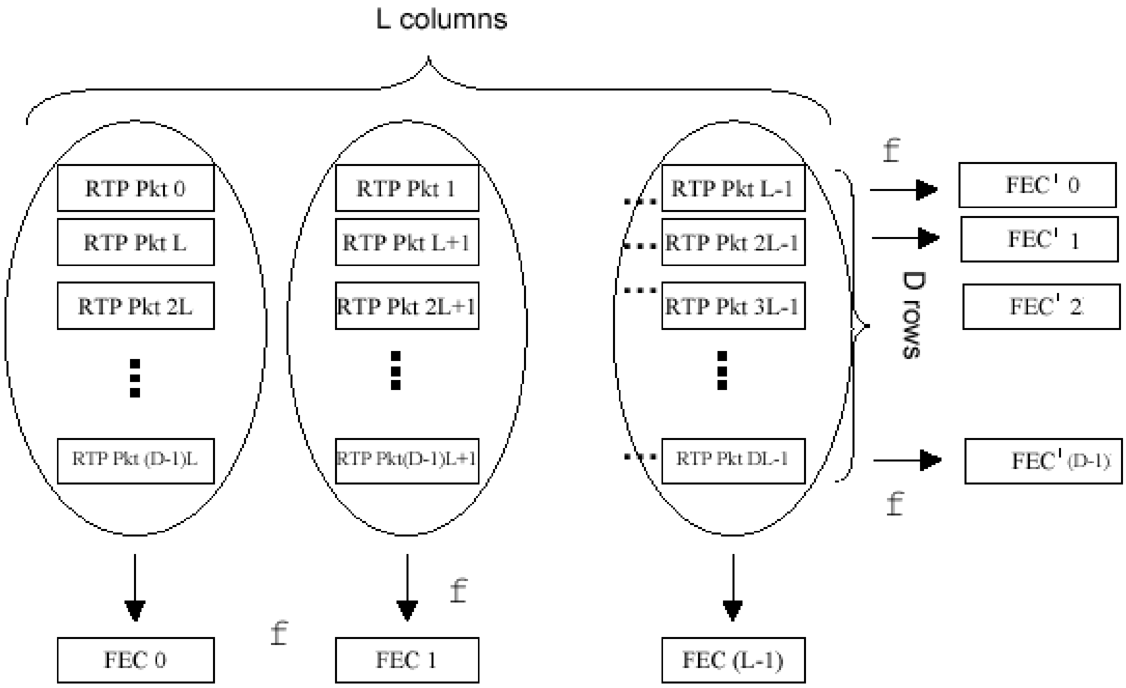

Figure 2 – The dual FEC mode structure

The second FEC stream can cope with any single packet loss, and the first FEC stream can cope with burst losses up to 'L' in length'.
The combined effect of the two FEC streams can also cope with more loss permutations than either FEC stream alone, though there are situations where recovering the maximum number of packets possible requires the iterative checking of both FEC streams until no more packets can be recovered.

4.5.3 FEC buffer, overhead and latency implications

To promote interoperability and simplify implementation, limits shall be specified for values of the L and D parameters. Manufacturers shall support all combinations of values of L and D within these limits. Manufacturers are free to extend beyond these values if desired. The specified limits are:

$$L * D \leq 100$$
$$1 \leq L \leq 20$$
$$4 \leq D \leq 20$$

These limitations apply to both FEC streams. It should be noted from this that it is only legal to transport a second FEC stream in the case where L ≥ 4.

The above limits only apply to the carriage of MPEG-2 Transport Streams, other video formats may use the same FEC scheme, but with different recommended constraints.

The following table summarizes the trade-off for different values of L and D between the overhead, the latency implied (for the case of 7 TS packets per IP packet) and the recovery capacity.

	Overhead	Latency			Recovery	Buffer size
		3Mbps	30 Mbps	100 Mbps		
XOR (5,10)	10%	175.5 ms	17.5 ms	5.3 ms	5 IP packets	66400 Bytes
XOR (10,10)	10%	350.9 ms	35.1 ms	10.5 ms	10 IP packets	132800 Bytes
XOR (20,5)	20%	350.9 ms	35.1 ms	10.5 ms	20 IP packets	132800 Bytes
XOR (8,8)	12.5%	224.6 ms	22.5 ms	6.7 ms	8 IP packets	84992 Bytes
XOR (10,5)	20%	175.5 ms	17.5 ms	5.3 ms	10 IP packets	66400 Bytes
XOR (8,5)	20%	140.4 ms	14.0 ms	4.2 ms	8 IP packets	53120 Bytes
XOR (5,5)	20%	87.7 ms	8.8 ms	2.7 ms	5 IP packets	33200 Bytes
XOR (4,6)	16.7%	84.2 ms	8.4 ms	2.5 ms	4 IP packets	31872 Bytes
XOR (6,4)	25%	84.2 ms	8.4 ms	2.5 ms	6 IP packets	31872 Bytes

4.5.4 FEC packet RTP header format

RFC2733 places constraints on the values of the fields in the RTP header. It specifies that the P, X, M, and CC fields are computed from the media packets, but because of the restrictions in section 4.2 the values of these fields will all be zero.

The static payload type mappings were exhausted early in the development of RTP, and there is now no method of registering static payload type numbers for new protocols.
In RFC2733 the payload type is defined with the symbolic name “parityfec” which is resolved to one of the RTP dynamic payload types using an external mechanism. For the proposed FEC the symbolic name recommended is “2dparityfec”.
As there is no simple mechanism available for resolving a dynamic payload type, and the only RTP traffic being sent and received by units corresponding to this code of practice should be the FEC data and the MPEG data (which has a fixed payload type), the FEC data shall be sent using the first available dynamic payload type number, which is 96 decimal.

The following additional restrictions will be adopted:

- The value of the SSRC field will not be used by the receiver.
- There is no requirement for the initial sequence number to be randomly assigned, as suggested in RFC3550.
- The time stamp field will not be used by the receiver.

4.5.5 FEC header format

The FEC header described in the RFC 2733 is originally 12 bytes. To allow for the extension to the error correction scheme, the FEC header needs to be modified as detailed in Figure 3.

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       SNBase low bits         |        Length Recovery        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|E| PT recovery |                    Mask                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          TS recovery                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|X|D|type |index|    Offset     |      NA       |SNBase ext bits|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Fig.3 Definition of the FEC header

The following fields are as defined in RFC2733:

- SNBase low bits: minimum sequence number of the packets associated to the FEC packet. For MPEG2 transport streams 16 bit sequence numbers are sufficient, so this parameter shall contain the entire sequence number. For

protocols with longer sequence numbers this field will contain the least significant 16 bits of the sequence number.
- Length Recovery: this field is used to determine the length of any media packets associated with the FEC packet.
- PT recovery: this field is used to determine the Payload Type of any media packets associated with the FEC packet.
- TS Recovery: this field is used to recover the timestamp of any media packets associated with the FEC packet.

The additional fields have been modified from what was in RFC2733, or are new. The definition of these is:

- E: In RFC2733 this shall be set to ‘0’, in this code of practice this shall be set to ‘1’ to indicate that the header is extended.
- Mask: In RFC2733 this is used to select which packets the FEC packet is applied to. The definition of the mask allows for a complex relationship between data packets and FEC packets, but this adds to implementation complexity. For simplicity, the mask field will be set to zero for implementations supporting this code of practice, and the NA field will be used instead. Handling of Mask requires special care due to the change of use from CoP #3 January 2003.
- X: This bit is reserved for future header extensions and must be set to zero to conform to this version of the FEC header.
- D: This bit is provided as an additional means of determining which FEC stream the packets belong. It must be set to 0 for FEC packets computed on columns and set to 1 for FEC packets computed on rows.
- Type: This field indicates which error-correcting code is chosen. It can be XOR (type=0), Hamming (type=1), Reed-Solomon (type=2). More encoding techniques can be used. For this version of the Code of Practice equipment shall only use the XOR type.
- Index: This field is used for more complex error protection codes. For the XOR method, only one FEC packet protects each group of media packets and hence the index field will always contain 0.
- Offset: This 1-byte field is the period chosen to select the media packets associated with this FEC packet, and corresponds exactly to the L parameter above for packets computed over columns (the first FEC stream). For packets computed over rows (the second FEC stream) this parameter shall always be one. This field should be kept constant during a session for each FEC stream.
- NA: This 1-byte field indicates the number of media packets associated with this FEC packet, and corresponds exactly to the D parameter above for packets belonging to the first FEC stream, and should correspond to the L parameter for packets belonging to the second FEC stream. This field should be kept constant during a session for each FEC stream.
- SNBase ext bits: This field is reserved for use with protocols which require extended sequence numbers longer than 16 bits. For MPEG2 transport streams 16 bit sequence numbers are sufficient, so this parameter shall be set to zero.

For protocols with longer sequence numbers this field will contain the next eight bits of the sequence number, beyond that which is in SNBase low bits.

For information, it should be noted that the media packets protected by any given FEC packet are defined as those with sequence numbers given by the formula SN-Base + j * Offset where $0 \leq j < NA$ are protected by any given FEC packet.

4.5.6 FEC linearity issues

The design of the FEC scheme can cause issues with the linearity of the output of a sending system. A simple FEC implementation, which keeps the column alignment, will have the column FEC packets available in a bursty fashion.

As for the high bitrate streams that are potentially used with this system such a burst is often undesirable, it may be necessary to buffer the data at the sender to ensure that the data is output in a more linear manner. The following constraints shall be applied when linearising the output of the FEC system:

- No linearising is required for the row FEC packets
- Column FEC packets shall be sent a minimum of L packets after the last media packet protected, to ensure that the burst loss tolerance of the system is not compromised.
- Column FEC packets shall be sent a maximum of L * D packets after the last media packet protected, to constrain the level of buffering required at the receiver.

It should be noted that the offsetting process presented in Informative Annex A intrinsically resolves the linearity issue.

An example linearising scheme is presented in Informative Annex B.

4.6 Timing recovery

Systems based on MPEG-2 transport streams already have timing recovery information present in the stream. This only provides precise timing information in some transport stream packets, which means that in the IP domain every IP packet will not contain a timestamp. The current standard RFC2250 has a timing recovery mechanism, though the clock for this only has a 90kHz resolution, this is sufficient to allow clock recovery of CBR streams.

RFC2250 requires that the RTP clock is derived from the PCR clock, but this is not a realistic requirement for systems handling multi-programme transport streams (MPTS) where there may be more than one PCR present, and the PCRs present can change over time. For CBR streams it is not required that sending systems lock their

RTP timestamps to any PCR. Because of this receiving systems shall not assume that the RTP timestamp will be locked to a PCR.

4.7 Jitter Tolerance

Network jitter can be absorbed by buffering at the receiver. There are two components in a typical IP network jitter issue. There is a first high-frequency component, caused by load spikes in the network. These tend to be quite small in value, of the order of ±10-15ms. On networks carrying Internet or other data traffic there is a 'wander/drift' component, as the loading of the network varies over a 24-hour period. This will typically be larger, of at least ±30-40ms. For the benefit of simplicity, these can be treated as one by providing a 'jitter budget' buffer of 120ms. This buffer should be run half full on average, providing a 60ms latency. For flexibility, it is recommended that system designs make it possible to modify the size of this buffer, as networks can have either significantly better or worse jitter performance.

The jitter absorption needs to be handled carefully, to ensure that the re-generated MPEG stream is still legal in terms of the PCR accuracy etc.

4.8 Latency

Latency within an IP adaptation unit is bounded by a combination of the jitter tolerance buffer, the FEC system and the clock recovery mechanism (not covered by this recommendation).

There are additional latencies caused by the MPEG encode/decode process, and the IP network transmission.

A number of professional applications have demanding round-trip delay requirements. . A round-trip delay of 400ms is widely accepted as the worst that would be acceptable for live interview applications. Camera control by remote telemetry has been shown to require as low as 200ms. Buffering requirements as part of the error-correction/protection mechanism make these difficult targets to attain for a system following this Code of Practice.

The table that follows gives the latency that will be incurred by the FEC process and jitter buffer for various bit rates and FEC configurations. This assumes the jitter buffer latency is kept constant at 60ms for all configurations. The table does not include the latency caused by the clock recovery process, or latencies caused by the MPEG encode and decode process.

	FEC + Jitter buffer latency		
	3Mbps	30 Mbps	100 Mbps
XOR (5,10)	235.5 ms	77.5 ms	65.3 ms
XOR (10,10)	410.9 ms	95.1 ms	70.5 ms
XOR (20,5)	410.9 ms	95.1 ms	70.5 ms
XOR (8,8)	284.6 ms	82.5 ms	66.7 ms
XOR (10,5)	235.5 ms	77.5 ms	65.3 ms
XOR (8,5)	200.4 ms	74.0 ms	64.2 ms
XOR (5,5)	147.7 ms	68.8 ms	62.6 ms
XOR (4,6)	144.2 ms	68.4 ms	62.5 ms
XOR (6,4)	144.2 ms	68.4 ms	62.5 ms

4.9 Re-order Tolerance

Packets travelling over IP networks are not guaranteed to arrive in the order sent. Sequence numbering is provided by RTP, which should allow this effect to be corrected within the receiving end equipment. Any re-ordering that is present is likely to be of a small order, less than 10 packets out of place.

The FEC system will correct the levels of packet re-ordering likely to be encountered. If a packet is grossly out of order then it is discarded, and will be corrected by the FEC scheme as a lost packet when it would have been expected. If a system holds multiple FEC matrices then it can tolerate re-ordering within this group of matrices. Systems should provide a means to be able to handle packets which are only out of order by a very small amount, but which occur the wrong side of their FEC matrix boundary as a result.

4.10 Encryption

Applications trying to adopt new encryption schemes may have difficulties getting the buy in of content providers for its use.
As the initial consideration is for the carriage of MPEG2 transport streams, then it will be possible to use standard MPEG Conditional Access systems before the IP encapsulation step. BISS is the EBU proposed encryption system for use in this area (EBU/ITU 3290Rev2). For users wishing to stream content without using MPEG conditional access, IPsec provides a means of encryption at the IP level.

4.11 CBR/VBR

RTP systems can be designed to support the carriage of VBR streams, but there are difficulties to be overcome. This version of the Code of Practice is intended for carrying CBR streams only. Future versions of this code of practice will be extended to support the carriage of VBR streams.

5 Signalling Protocols

5.1 Network QoS protocols (RSVP/MPLS/DiffServ etc)

It is recommended that terminal equipment allows the IP header TOS byte to be fully configured. This will allow both traditional TOS values, and DiffServ markings. Bandwidth reservation will be handled using out of band mechanisms - the end devices are not required to support a bandwidth reservation protocol.

5.2 Session control protocols (SAP/SIP/SDP/RTSP/RTCP etc)

Further work is required to define the signalling systems to be used, and the parameters for these systems.

As such the initial baseline is that systems can have the required parameters manually configured. The following parameters list is the required minimum:

- As the system is UDP based, and there is no commonly accepted standard UDP port number or range for RTP systems, it is recommended that the sender allows the destination port number and IP address (which may be a multicast group) to be configured. There is an RTP requirement that any port number chosen be an even number.
- It is required that the FEC on a sending device can be enabled or disable - though it is recommended that FEC be used for all contribution applications, there may be some applications where broadcasters are trying to maximise the use of the bandwidth, and are happy to accept the extra errors in the signal in exchange for the extra bandwidth available. Various FEC profiles should be offered on sending devices, to allow an appropriate FEC level to be selected for the network being used.
- The receiver shall be able to cope if no FEC stream is received, because the sender may or may not be transmitting one. The receiver should be configurable to expect zero, one or two FEC streams to remove the potentially time consuming process of determining how many FEC streams it is receiving.
- For configuration simplicity, where FEC is present, the first stream shall be sent on the UDP port number two higher than the main data port number configured (The port number one higher is reserved for use by RTCP), and the second stream shall be sent on the UDP port number four higher than the main data port.
- To support multicast, there should be an option on receiving units to allow a multicast group address to be configured.
- The receiver should use IGMP version 3 to join and leave multicast groups when required. As some network equipment may not yet support IGMP version 3, there may be applications where earlier versions of IGMP are required.
- The port number the receiving unit will listen on should also be configurable. Again, there is a requirement that the RTP port number shall be even.

- As mentioned earlier, the value for the IP TOS byte should be configurable.

6 Management Protocol

As applications mature there will be a desire to adopt a common management interface, this requires further study.

END OF CODE OF PRACTICE

Informative Annex A - Non block aligned FEC arrangement

Figure 4 shows an example alternative arrangement for the case of (L=4, D=5). In this example, the FEC packet F1 protects data packets [1, 5, 9, 13, 17] while FEC packet F6 protects data packets [6, 10, 14, 18, 22].

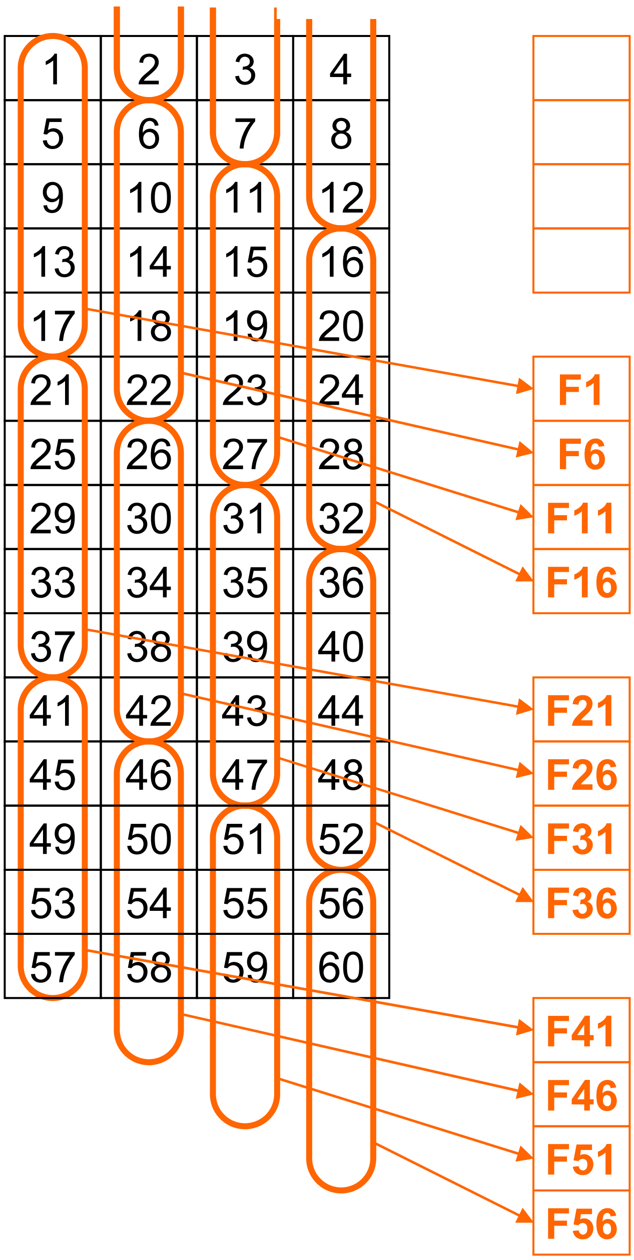

Figure 4 – alternative FEC arrangement

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

There are a wide variety of other valid methods for organising the data packets to create the column FEC stream, which also meet the normative requirements of this practice.

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

Informative Annex B - Block aligned FEC linearisation

The goal of this example arrangement is to allow the most linear use of the bandwidth by ensuring that FEC packets are regularly inserted among data packets of the next matrix.

Column FEC packets are sent by using an interleaver. The depth of the interleaver is chosen to be D so as to make sure that the L column FEC packets will be regularly inserted among the L*D data packets of the next matrix. It is chosen to send the first column FEC at the "same time" as the first data packet of the next matrix so a receiver can easily detect it.

Figure 5 illustrates such an arrangement with L=4 and D=5. FEC'n packet refers to the row FEC packet computed over row #n and FECn packet refers to the column FEC packet computed over column #n. Packets are sequenced in the reading direction (from left to right then from top to bottom).

Packets in the same cell are sent at the "same time" (by different senders and using different RTP sequences). Receivers shall not depend on the packet ordering from the same cell of a matrix.

0	1	2	3
4,FEC'0	5	6	7
8,FEC'1	9	10	11
12,FEC'2	13	14	15
16,FEC'3	17	18	19
20,FEC'4,FEC0	21	22	23
24,FEC'5	25,FEC1	26	27
28,FEC'6	29	30,FEC2	31
32,FEC'7	33	34	35,FEC3
36,FEC'8	37	38	39

Figure 5 – Linearising example for L = 4 and D = 5

As a second example, Figure 6 shows this arrangement for the case where L=4 and D=2.

0	1	2	3
4,FEC'0	5	6	7
8,FEC'1,FEC0	9	10,FEC1	11
12,FEC'2, FEC2	13,FEC2	14,FEC3	15,FEC3

Figure 6 – Linearising example for L = 4 and D = 2

END OF ANNEX

