

ETSI TS 102 542-3-2 V2.1.1 (2016-04)



**Digital Video Broadcasting (DVB);
Guidelines for the implementation of
DVB-IPTV Phase 1 specifications;
Part 3: Error Recovery;
Sub-part 2: Application Layer -
Forward Error Correction (AL-FEC)**

EBU
OPERATING EUROVISION

DVB[®]
Digital Video
Broadcasting

ReferenceRTS/JTC-DVB-360-3-2

Keywordsbroadcasting, digital, DVB, IP, TV, video

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

The present document can be downloaded from:
<http://www.etsi.org/standards-search>

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the only prevailing document is the print of the Portable Document Format (PDF) version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at
<https://portal.etsi.org/TB/ETSIDeliverableStatus.aspx>

If you find errors in the present document, please send your comment to one of the following services:
<https://portal.etsi.org/People/CommiteeSupportStaff.aspx>

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.

The content of the PDF version shall not be modified without the written authorization of ETSI.
The copyright and the foregoing restriction extend to reproduction in all media.

© European Telecommunications Standards Institute 2016.

© European Broadcasting Union 2016.

All rights reserved.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are Trade Marks of ETSI registered for the benefit of its Members.

3GPP™ and **LTE™** are Trade Marks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners.

GSM® and the GSM logo are Trade Marks registered and owned by the GSM Association.

Contents

Intellectual Property Rights	5
Foreword.....	5
Modal verbs terminology.....	5
1 Scope	6
2 References	6
2.1 Normative references	6
2.2 Informative references.....	6
3 Abbreviations	7
4 Overview of DVB AL-FEC implementation guidelines	7
5 Configuring FEC protection	8
5.0 General Preparations	8
5.1 Correcting for burst losses.....	8
5.2 Correcting for random losses.....	8
6 FEC sending arrangement considerations	9
6.1 Introduction	9
6.2 Client considerations	9
6.3 FEC Sending Arrangements	10
6.3.0 General.....	10
6.3.1 Constant rate, non-interleaved sending	10
6.3.2 Fully interleaved sending.....	11
6.3.3 Partially interleaved sending.....	11
6.3.4 Faststart sending for stored/buffered content.....	12
7 Layered multicast sending.....	12
8 Criterion for selection of Forward Error Correction for the protection of audiovisual streams delivered over IP Network Infrastructure.....	13
8.1 Requirements.....	13
8.2 System description	14
8.3 Packet loss characteristics	14
8.4 FEC Scheme Evaluation Criteria.....	15
9 AL-FEC evaluation report for DVB-TM IPI.....	16
9.0 General	16
9.1 Introduction	16
9.2 Sending arrangement considerations	16
9.3 Bandwidth costs	17
9.3.0 Objectives	17
9.3.1 Loss models	17
9.3.2 Multicast case	17
9.3.2.0 Simulation results.....	17
9.3.2.1 Results with constant sending arrangement	18
9.3.2.2 Results with burst sending arrangement.....	19
9.3.3 Unicast case	22
9.3.3.1 Stored/buffered content	22
9.3.3.2 Live content.....	24
9.3.3.2.0 Simulation results for unicast delivery	24
9.3.3.2.1 Constant sending arrangement.....	25
9.3.3.2.2 Burst sending.....	26
9.3.4 A note on latency, jitter and traffic shaping.....	29
9.3.5 Summary of simulation results	30
9.4 Flexibility	31
9.5 Processing and Memory requirements	31
9.6 Additional criteria.....	32
9.7 Content Download.....	33

9.8	Raptor vs. Pro-MPEG Summary	33
9.9	Conclusions	34
10	Sending arrangements used for simulations	35
10.1	DF Raptor™ default sending arrangement	35
10.2	Pro-MPEG COP3 fully interleaved sending arrangement	36
10.3	Pro-MPEG COP3 burst sending arrangement	36
10.4	Concurrent Interleaved sending	37
10.5	DF Raptor™ faststart sending for stored/buffered content	38
11	Concurrent interleaving results	39
12	Hybrid code	43
12.0	General	43
12.1	Hybrid code results	44
	History	48

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: "*Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards*", which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (<https://ipr.etsi.org>).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Foreword

This Technical Specification (TS) has been produced by Joint Technical Committee (JTC) Broadcast of the European Broadcasting Union (EBU), Comité Européen de Normalisation ELECtrotechnique (CENELEC) and the European Telecommunications Standards Institute (ETSI).

NOTE: The EBU/ETSI JTC Broadcast was established in 1990 to co-ordinate the drafting of standards in the specific field of broadcasting and related fields. Since 1995 the JTC Broadcast became a tripartite body by including in the Memorandum of Understanding also CENELEC, which is responsible for the standardization of radio and television receivers. The EBU is a professional association of broadcasting organizations whose work includes the co-ordination of its members' activities in the technical, legal, programme-making and programme-exchange domains. The EBU has active members in about 60 countries in the European broadcasting area; its headquarters is in Geneva.

European Broadcasting Union
CH-1218 GRAND SACONNEX (Geneva)
Switzerland
Tel: +41 22 717 21 11
Fax: +41 22 717 24 81

The Digital Video Broadcasting Project (DVB) is an industry-led consortium of broadcasters, manufacturers, network operators, software developers, regulatory bodies, content owners and others committed to designing global standards for the delivery of digital television and data services. DVB fosters market driven solutions that meet the needs and economic circumstances of broadcast industry stakeholders and consumers. DVB standards cover all aspects of digital television from transmission through interfacing, conditional access and interactivity for digital video, audio and data. The consortium came together in 1993 to provide global standardization, interoperability and future proof specifications.

The present document is part 3, sub-part 2 of a multi-part deliverable. Full details of the entire series can be found in part 1 [i.1].

Modal verbs terminology

In the present document "**shall**", "**shall not**", "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the [ETSI Drafting Rules](#) (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

1 Scope

The present document is designed as a companion document to help implement the DVB-IPTV Phase 1 specification ETSI TS 102 034 [1], which is referred to as the Handbook.

Part 3 of this multi-part deliverable deals with Error recovery technologies [i.2] and [i.3]. The present document provides guidelines on the Application Layer - Forward Error Correction (AL-FEC) technology.

2 References

2.1 Normative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at <http://docbox.etsi.org/Reference>.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are necessary for the application of the present document.

- [1] ETSI TS 102 034: "Digital Video Broadcasting (DVB); Transport of MPEG-2 TS Based DVB Services over IP Based Networks".
- [2] ETSI TS 126 346: "Universal Mobile Telecommunications System (UMTS); Multimedia Broadcast/Multicast Service (MBMS); Protocols and codecs (3GPP TS 26.346 Release 7)".
- [3] SMPTE Specification 2022-1 (2007): "Forward Error Correction for Real-time Video/Audio Transport Over IP Networks".
- [4] ETSI TS 102 005: "Digital Video Broadcasting (DVB); Specification for the use of Video and Audio Coding in DVB services delivered directly over IP protocols".

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

- [i.1] ETSI TS 102 542-1: "Digital Video Broadcasting (DVB); Guidelines for the implementation of DVB-IPTV Phase 1 specifications; Part 1: Core IPTV Functions".
- [i.2] ETSI TS 102 542-3-1: "Digital Video Broadcasting (DVB); Guidelines for the implementation of DVB-IPTV Phase 1 specifications; Part 3: Error Recovery; Sub-part 1: Overview of DVB-IPTV Error Recovery".
- [i.3] ETSI TS 102 542-3-3: "Digital Video Broadcasting (DVB); Guidelines for the implementation of DVB-IPTV Phase 1 specifications; Part 3: Error Recovery; Sub-part 3: Retransmission (RET); IPTV Implementation Guidelines".

3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

AL-FEC	Application Layer - Forward Error Correction
AVC	Advanced Video Coding
CBMS	Convergence of Broadcast and Mobile Services
DF	Digital Fountain
DMA	Direct Memory Access
DSL	Digital Subscriber Line
DVB	Digital Video Broadcasting
DVB-TM	Digital Video Broadcasting - Technical Module
FEC	Forward Error Correction
HD	High Definition
IP	Internet Protocol
IPI	IP Infrastructure
IPTV	IP TeleVision
KB	Kilo Byte
MIPS	Mega Instructions Per Second
MPEG	Moving Picture Experts Group
PLR	Packet Loss Ratio
REIN	Repetitive Electrical Impulse Noise
RTCP	Real-time Transport Control Protocol
RTP	Real-time Transport Protocol
SD	Standard Definition
SD&S	Service Discovery and Selection
SMPTE	Society of Motion Picture and Television Engineers
STB	Set Top Box
TM	Technical Module
TS	Transport Stream
UDP	User Datagram Protocol
VoD	Video on Demand
XOR	eXclusive OR

4 Overview of DVB AL-FEC implementation guidelines

The present document is intended to help people implement the DVB-IPTV Phase 1 Handbook ETSI TS 102 034 [1] which includes the DVB application layer forward error correction (AL-FEC) specification.

The main part of the present document is intended to provide guidance for those intending to use the DVB AL-FEC specification. Clause 5 discusses issues around configuration of AL-FEC and clause 6 summarizes the options for sending. Clause 7 describes how layered multicast sending can be used to allow the amount of FEC overhead to be varied to suite the packet error rates experienced on individual connections for multicast delivery.

The present document includes two documents that were created by DVB as part of the evaluation process as clauses 8, 9 to 12. Clause 8 presents the evaluation criterion that were agreed before the selection process started and 9 to 12 are from the report on the evaluation process and give the rationale for the choice of the hybrid approach used in the DVB AL-FEC specification.

Some parts of the present document (mainly from clauses 8 to 12) have been included in a separate DVB bluebook on AL-FEC evaluations because a number of standards organizations and others requested sight of it before this version of the guidelines was published by ETSI.

The DVB AL-FEC code is defined only for the case of RTPover-UDP transport. The defined UDP transport cannot support AL-FEC in a backwards compatible manner.

5 Configuring FEC protection

5.0 General Preparations

This clause provides a brief overview of the issues and parameters which have to be considered when configuring FEC protection using the DVB AL-FEC code.

The two principle parameters that shall be determined are as follows:

- the AL-FEC block size, or "protection period";
- the AL-FEC overhead.

The AL-FEC block size is the number of packets which are protected together as a block, or equivalently the time required to send those packets at the stream rate ("protection period"). The AL-FEC overhead is the amount of additional FEC data ("repair packets") that are sent as a fraction of the original data. For example, if the AL-FEC block size is 100 packets and 10 repair packets are sent for each block, then the AL-FEC overhead is 10 %.

In order to determine the AL-FEC protection required, a good understanding of the loss characteristics of the target network is required. In general, the objective of AL-FEC is to provide error-free reception over long periods of time (several hours) - loss events which cannot be corrected by the AL-FEC should therefore be very rare. This means that loss characteristics have to be understood over long time periods.

The AL-FEC code operates on each FEC block independently. This means that loss characteristics also have to be understood at a timescale equal to the FEC block size: averages over long time periods are not sufficient. For example, the average packet loss over a one hour period may be very low, but if many of the losses are concentrated in a short period of time they may still overwhelm the AL-FEC code.

Clauses 5.1 and 5.2 provide some general guidelines on the effect of configuration parameters targeted at correcting for burst and random loss respectively. In practice, losses are a combination of these two.

5.1 Correcting for burst losses

In order to correct for isolated burst losses, a number of repair packets greater than or equal to the worst expected burst loss has to be provided for each FEC block. If only the AL-FEC base layer is used, then a number of repair packets equal to the worst expected burst loss has to be provided. Note that for the base layer, the number of repair packets per block has to be a divisor of the block size (in packets). This configuration will correct for isolated burst losses, but will often not correct randomly distributed losses and generally cannot correct for cases where multiple bursts occur within a block.

When the AL-FEC base and enhancement layers are used, it is generally sufficient to provide a number of repair packets one greater than the worst expected burst loss per block. This configuration generally can also correct for randomly distributed losses or multiple bursts per block provided sufficient enhancement layer packets are sent.

Note that the number of repair packets required per block is fixed independent of the block size. As a result longer block sizes will result in a lower relative AL-FEC overhead. However, longer block sizes will also contribute additional latency, affecting channel change times.

5.2 Correcting for random losses

In order to correct for randomly distributed losses, it is necessary to understand the "worst case" number of lost packets within an FEC block. If losses were truly random and independent, then the statistical probability of losing 1, 2, 3, etc. packets in a block could be calculated and from these probabilities the expected frequency of such events could also be calculated. The "worst case" of interest is then the worst case occurring frequently enough to be an issue from a quality perspective (which is a judgement issue on the part of the service provider). If this "worst case" can be corrected by AL-FEC, then although there may remain uncorrected events these will be so rare that they can be ignored (for example once a day or once a week).

In practice, losses are not independent and random and so the worst case cannot be calculated statistically. Network measurements need to be used to determine the worst case that the AL-FEC has to correct.

When only the base AL-FEC layer is used, then certain levels of random packet loss can be corrected. Clause 9.3 provides some simulated examples.

When base and enhancement layer AL-FEC is used, then the minimum number of repair packets needed to correct a worst case of n randomly distributed lost packets per block is $n+1$, where one of the packets is a base layer packet and the remainder are enhancement layer packets. It should be noted that this configuration will not provide very much protection for end devices which support only the base layer. If more than one base layer packet is provided, for example to provide burst loss protection to devices which do not support the enhancement layer, then this reduces the effectiveness of the overall code in the face of random losses and more than $n+1$ repair packets may be needed in total.

Finally, although in this case the number of repair packets required is not independent of the block size, a similar trade-off exists between additional latency and bandwidth. For example, if the block size is 100 packets and the worst case loss is 10 packets, then it is highly unlikely that two 100 packet blocks, each with 10 lost packets should occur in sequence. In fact the worst case loss for 200 packet blocks may be only slightly larger than 10, meaning that the bandwidth overhead can still be roughly halved by increasing the block size from 100 packets to 200 packets. Note that because measurements are taken over very long time periods, and thus millions of blocks, then even if the average packet loss is very low, it may still occur that occasionally events such as 10 lost packets in a block of 100 occur.

6 FEC sending arrangement considerations

6.1 Introduction

Another important issue in the determination of FEC performance is the arrangement of data packets (source and FEC "repair" packets) in time for sending. The sending arrangement impacts FEC performance in three ways:

- The additional latency introduced by the use of FEC.
- The data rate profile (constant vs. bursty) of the resulting stream.
- The FEC overhead required to overcome packet loss with given characteristics.

The additional latency is impacted because it is necessary for the receiver to wait long enough for reception of all packets (source and repair) of the first source block before beginning presentation of the stream to the user. This is true even if there is no loss in the first block because once presentation has begun, and assuming freezing of the video is unacceptable, the presentation schedule for the whole stream is set by the initial start time. If presentation of the first block begins before all packets have arrived then presentation of every block will have to begin before all packets have arrived and this will prevent the FEC operation from being applied to recover losses when they do occur.

A sending arrangement which sends the FEC repair packets as soon as possible after the source packets minimizes this additional latency. Sending arrangements which interleave repair packets with source packets from the subsequent block increase the latency according to the amount of interleaving.

The sending arrangement clearly impacts the data rate profile. If FEC repair packets are sent in a burst immediately after each source block then the overall data rate profile will be very bursty.

Finally, the FEC overhead required may also be impacted. For example, when packets from a given block are sent in quick succession, then a burst outage may cause the loss of many packets. If they are spread out over time then fewer packets (from that block) will be lost. The FEC overhead is often dimensioned based on anticipated worst case burst outages and thus the sending arrangement can impact the required overhead.

6.2 Client considerations

The present document does not prescribe a particular sending arrangement: sending devices are free to use whatever sending arrangement they choose, subject to certain constraints. Receivers should be able to process incoming packets whatever arrangement they arrive in.

The service discovery signalling for FEC protected streams may provide information about the stream which may be used by receivers to determine the amount of buffering required for FEC purposes. The "FEC Maximum Block Size (Packets)" indicates the maximum number of stream source packets that will occur between the first packet of a source block and the last packet for that source block (source or repair). A receiver may keep a count of the number of source packets which have been received since the first packet of a particular source block. This count should include packets from any blocks, not just the particular one of interest. Once this count reached the signalled Maximum Block Size (Packets), the receiver may assume that no further packets (source or repair) for the particular source block will be received. It is then safe to begin presentation of the block. This approach is applicable in cases where the stream is constant bit-rate and the source blocks of constant duration.

The SD&S signalling may alternatively or additionally indicate the "FEC Maximum Block Size (Time)". This indicates the maximum sending duration of any FEC block. A receiver may measure the elapsed time from the receipt of the first packet of a particular block. Once this time exceeds the Maximum Block Size (Time) the receiver may assume that no further packets (source or repair) for the particular source block will be received. It is then safe to begin presentation of the block. This approach is applicable in cases where the stream is constant or variable bit-rate and the source blocks are of constant or variable duration.

Since IP networks introduce jitter, receivers should not make assumptions based on short-term measurements of packet arrival times. Long-term measurements can yield reliable information about clock drift between sender and receiver, but otherwise, clock recovery at receivers should be based on RTP timestamps and MPEG-2 Program Clock References, not on the absolute packet arrival times.

6.3 FEC Sending Arrangements

6.3.0 General

This clause describes some possible FEC Sending Arrangements. Other arrangements are possible. Receiver implementations should not make assumptions about the sending arrangement in used, except that it will conform to the signalled Maximum Block Size.

6.3.1 Constant rate, non-interleaved sending

In this sending arrangement, depicted in figure 6.3.1.1, the overall sending rate is kept constant and the source packets of each block are sent before any of the repair packets of the block. This approach requires that the sending rate of the source packets be increased marginally to make space for the repair packets at the end of the block.

It is important to note that the sequencing of packets is determined by the FEC procedures which operate "below" the RTP layer. The contents of the packets, in particular the RTP timestamps, should not be modified compared to the case in which FEC is not applied so that the correct timing for the packets can be reconstructed with the usual procedures.




Figure 6.3.1.1: Constant-rate, non-interleaved sending arrangement

Advantages of this sending arrangement are that the total data rate remains constant and the additional latency due to FEC is minimized. However, insertion of repair packets introduces small amount of jitter on all source packets.

6.3.2 Fully interleaved sending

In this sending arrangement, depicted in figure 6.3.2.1, the overall sending rate is kept roughly constant and the sending rate of source packets is also kept constant.

Because this sending arrangement distributes repair packets for one block over the entire duration of the next block, then the additional latency due to FEC is equal to the duration of two blocks. When working with a fixed latency budget, this implies that the block size for the sending arrangement described here would be half that for the sending arrangement described in clause 6.3.1. As a result, the overhead required by the code is increased.

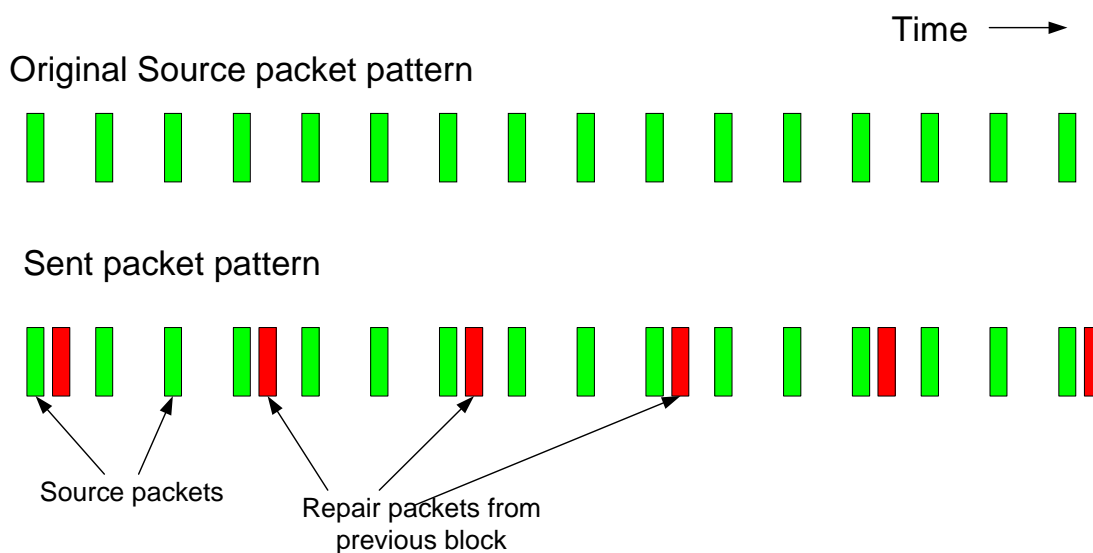


Figure 6.3.2.1: Fully interleaved sending arrangement

An advantage of this approach is that, except for small perturbations caused by the introduction of the repair packets, then the arrival times of the source packets are similar to the arrival times when FEC is not used. Additionally, the overall data rate is roughly constant. However, the high latency with respect to the block size is a significant issue.

6.3.3 Partially interleaved sending

In this sending arrangement, depicted in figure 6.3.3.1, repair packets for one block are interleaved with the first few packets of the next block. As a result, the instantaneous sending rate during these first few packets is significantly increased. However, the block size may now be set almost as large as the latency budget, which reduces the required overhead.




Figure 6.3.3.1: Partial interleaving sending arrangement

An advantage of this arrangement is that the source block size may be almost as large as the latency budget. However, the sending rate is extremely bursty - with double the bandwidth used at the beginning of each block. Note that if traffic shaping is used to return the stream to constant bit-rate, this will introduce jitter similar to that introduced by the constant, non-interleaved sending arrangement of clause 6.3.1. However, the additional latency will still be higher than in the constant non-interleaved case.

6.3.4 Faststart sending for stored/buffered content

This sending arrangement, applicable to stored or buffered content (i.e. VoD and trick modes on live content) is illustrated in figure 6.3.4.1. In this arrangement, source data is sent slightly faster than the nominal stream rate at the start of the session or when trick modes are used. This allows the buffering period to be gradually increased without introducing additional latency.

Two variants of this approach are described here:

- "faststart with constant rate sending": in this approach the additional source data bandwidth is obtained by reducing the FEC bandwidth at the beginning of the stream. As a result the total stream rate remains constant, but stream quality is reduced for these few initial seconds.
- "faststart with variable rate sending": in this approach the overall stream rate at the beginning of the stream is somewhat higher than the nominal stream rate (e.g. 20 % higher) for the initial few seconds of the stream, but as a result the stream quality is maintained.

During the DVB FEC evaluation exercise, the second approach provided the best results.




Figure 6.3.4.1: Faststart sending arrangement

An advantage of this approach is that the source block size can quickly be increased without introducing additional latency. For example, the additional latency budget for FEC might be set at 100 ms, but by using the arrangement above the protection period can be increased to 500 ms over the first few seconds of the stream.

7 Layered multicast sending

The AL-FEC code defined by DVB supports layered sending in the multicast case. When layered multicast is used, then multiple multicast groups are used for a single DVB-IPTV stream. Each multicast group introduces incrementally more FEC protection so that receivers can adjust the amount and type of FEC data received according to their capabilities and requirements by joining and leaving the appropriate multicast group.

Source and repair packets within a DVB IP stream are "self identifying", meaning that the type and meaning of each packet can be identified from the packet contents (and in particular the UDP destination port number), without reference to the multicast group on which the packet was received.

DVB AL-FEC transmissions consist of a base layer and optionally one or more enhancement layers. Each layer may be provided on a different multicast group. The IP multicast group and destination UDP port number for each layer are provided within SD&S signalling.

Receivers which do not support or do not require FEC data should join only the multicast group associated with the original stream. Those receivers supporting or requiring only the FEC base layer should additionally join the multicast group associated with the base layer and those receiver supporting or requiring the enhancement layer or layers should join the multicast group or groups associated with the enhancement layer. Where multiple groups are advertised, receivers should join them "incrementally", i.e. they should join multiple groups rather than choosing a single group.

Receivers may determine the amount of AL-FEC required based on measurements of packet loss. However, since the AL-FEC is designed to deliver a broadcast-quality stream the protection has to be sufficient to handle even relatively rare packet loss events and so any such measurements shall be done over a long period of time. Alternatively, the number of layers to receive may be determined by operator configuration possibly linked to remote management.

The AL-FEC standard does not prescribe how much FEC overhead is allocated to each layer, nor the number of layers or the allocation of layers to multicast groups. In fact, all FEC data (base and enhancement layers) may be sent on the same multicast group as the original data or there may be one multicast group for original data and one for FEC data (base and enhancement layers).

8 Criterion for selection of Forward Error Correction for the protection of audiovisual streams delivered over IP Network Infrastructure

8.1 Requirements

Audiovisual services delivered over networks are subjected to the inherent properties of those networks including latency and errors. DVB commercial requirement is quoted as:

"Inclusion of suitable error protection strategies such as an FEC mechanism to enable DVB services to be carried over typical IP access networks with an acceptable quality of service (maximum 1 visible artefact/hour).

- *The selected solution shall be in line with work of other standards bodies such as DSL-Forum. If necessary, DVB should liaise with relevant other bodies.*
- *The selected solution shall provide flexibility so that it covers a reasonable range of networks and a variety of business models (trade-off versus payload). Furthermore, the selected solution shall be extensible to cover likely future streaming requirements.*
- *The selected solution shall be implementable on a range of HNEDs without significantly increasing product cost."*

The DVB TM agreed that the IP Infrastructure group should recommend an (optional) application layer FEC. It is agreed that it should work end to end including the core and home network where required.

The FEC scheme selection process should take into account:

- 1) Packet loss characteristics of practical IP access network implementations e.g. DSL. These might include the use of interleaving at the physical layer to improve transport performance.
- 2) Further packet losses that could occur in the core network due to congestion and/or the home environment e.g. wireless technologies.
- 3) Sensitivity of A/V coding to errors.
- 4) Practical viability and flexibility of FEC scheme (encoding and decoding) to meet the min and max correction at minimal cost (processing, memory) for large numbers of simultaneous streams.
- 5) Ongoing cost of bandwidth inefficiency inherent in the code - i.e. difference between the bandwidth required by the code and the theoretical minimum bandwidth needed for service in the given loss conditions.
- 6) Pre-computation of the FEC to enable later usage when the content is streamed.
- 7) Carriage directly over RTP in the future i.e. without an MPEG2 transport stream.
- 8) Dynamically varying length of IP packets carrying A/V content.

8.2 System description

Figure 8.2.1 is an example of video service delivery over DSL network from source (top left) to set top box (top right). It highlights the components through which the service is delivered and the logical position of the Application Layer FEC. Key points brought out by this diagram are:

- There are other possible mechanisms that affect the delivery of acceptable quality of service (maximum 1 visible artefact/hour). These are DSL layer FEC/interleave, video/audio coding type and any error concealment at the decoder. The application layer FEC performance should provide adequate protection from errors **with and without** these mechanisms present (shown as min and max correction in figure 8.2.1).
- When these other mechanisms are present, the application layer FEC should take into account the effect of failure of these other mechanisms under severe error conditions.
- When these other mechanisms are present, the 'load' on the application layer FEC is reduced under normal error conditions, leading to possible 'cost' reductions in terms of latency, memory, processor, etc.
- Gaps in the core network domain and home network domain highlight the possible presence of other network types that could introduce service affecting packet loss. These networks should ideally be taken into account in the specification of application layer FEC performance, though will vary between implementations.

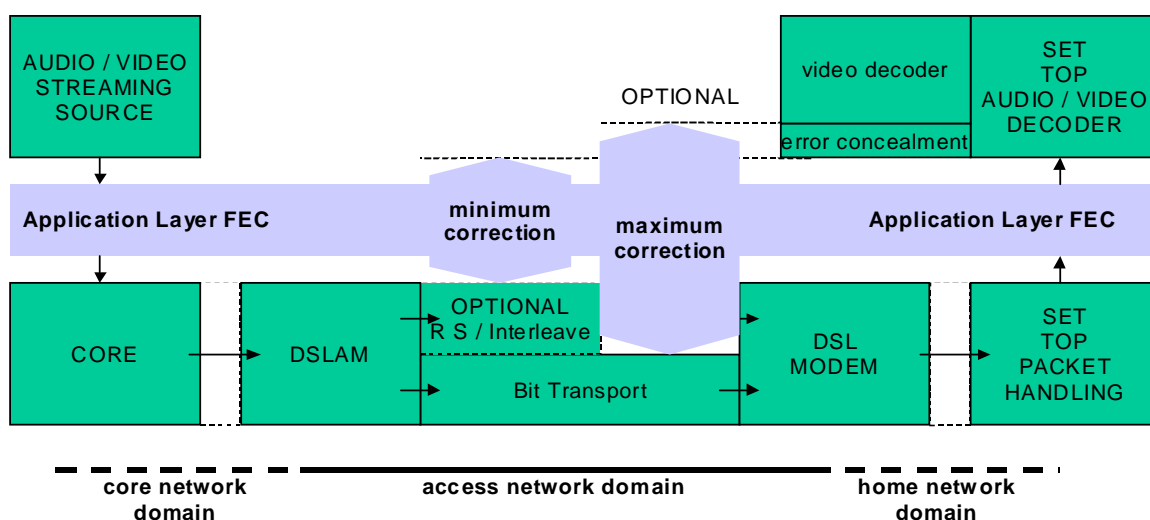


Figure 8.2.1: Minimum and maximum correction requirement for DSL access network domain

8.3 Packet loss characteristics

The packet loss characteristics should be provided by network operators and DSL chip vendors, ideally in the form of data collected from implementations or (if this is too commercially sensitive) in the form of a statement on what level of errors should be corrected by the application layer i.e. the requirements.

Worst case end-to-end packet loss metrics can be provided in terms of average loss rate, and loss distribution (independent random vs. bursty) for the IP packets, independent of bit rate.

NOTE: Methods for characterization of the loss distribution need further discussion.

Results for impulsive noise in DSL networks are available from the ITU and (until other information becomes available) they will be used as the basis of the evaluations. Although DSL is clearly an important case (where the results may vary widely), it is desirable to allow for other core, access and home networks also.

8.4 FEC Scheme Evaluation Criteria

Assume the following criteria:

- 1) Consider 3 error distributions:
 - a) random losses (PLR $1e-3$ to $1e-5$);
 - b) burst losses (PLR $1e-3$ to $1e-5$ with distributions based on ITU DSL results);
 - c) better than $1e-5$.
- 2) Additional latency due to FEC depending on applications (VoD = 100 ms, Broadcast = 400 ms).
- 3) Bit-rate for VoD = 2 Mbit/s, Broadcast = 2 Mbits and 6 Mbits (both based on H264/AVC).
- 4) Target mean time between FEC blocks that contain uncorrectable errors = 4 hours.

Data should be provided for each FEC proposal, specifying the performance **for each set of parameters employed** to illustrate range of performance available in terms of:

- Overhead required by the FEC to achieve the target performance in each of the given loss conditions (FEC data)/(protected data) (%).
- Flexibility:
 - Changing the overhead or/and the block size dynamically (within or between FEC blocks).
 - Range of protection periods.
 - Suitability for use with a wide variety of FEC sending strategies.
- STB memory requirement for buffering/processing (bytes).
- STB processing requirement measured as:
 - Maximum and average number of XOR operations.
 - Maximum and average number of conditional statements (IF..THEN).
 - Maximum and average number of context switching.
 - Maximum and average size of additional temporary memory needed.
 - Maximum and average number of threads (if threaded).
- Headend memory requirements for buffering (bytes).
- Headend processing requirement measured as:
 - Maximum and average number of XOR operations.
 - Maximum and average number of conditional statements (IF..THEN).
 - Maximum and average number of context switching.
 - Maximum and average size of additional temporary memory needed.
 - Maximum and average number of threads (if threaded).
 - Maximum memory bandwidth.
- Scalability, e.g. suitability for hardware implementation and cost.
- How much data is lost when the FEC fails? Visibility of artefacts when FEC fails.
- Ability to discard the FEC flow and process only the original packets as normal.

- Ability to add or remove FEC correction packets.

Additionally, systems considerations should be addressed including:

- Continued functioning of existing STB products in presence of FEC data.
- Option for new STB products to use or ignore FEC data.
- Confirmation of FEC scheme IPR compliance with DVB rules.
- Support of combined protection of audio and video packets.

9 AL-FEC evaluation report for DVB-TM IPI

9.0 General

This clause contains the evaluation report of the DVB-TM IPI group on the proposed AL-FEC codes. Note that the two codes originally proposed were the Pro-MPEG Code of Practice 3 code as now specified in SMPTE 2022-1 [3] and the Digital Fountain Raptor™ code essentially as specified in ETSI TS 126 346 [2]. The eventually standardized code was a hybrid of these two original proposals.

NOTE: DF Raptor™ is the trade name of a technology supplied by Qualcomm Inc. This information is given for the convenience of users of the present document and does not constitute an endorsement by ETSI of the product named. Equivalent products may be used if they can be shown to lead to the same results.

9.1 Introduction

The report provides results of the DVB-TM IPI evaluation process for forward error correction for IPTV. Two candidate FEC codes have been considered, the Digital Fountain Raptor™ code, and the Pro-MPEG Code of Practice 3 based proposal.

Clause 8 provides the agreed evaluation criteria, with the exception that it was later agreed to consider "additional latency due to FEC" of 100 ms and 400 ms (rather than "protection periods") and "mean time between packet loss" (rather than "mean time between FEC blocks with errors").

During the evaluation process, it was realized that a key issue in determining the FEC performance is the sequencing and timing of the sending of source and FEC packets. This issue is discussed further in clause 9.2. Examples of sending arrangements are described in clause 6.

This clause also includes simulation results for the following cases:

- "concurrent interleaved sending": in which FEC packets are interleaved with the source packets they protect, these results are included in clause 11.
- "hybrid code": in which a mixture of Pro-MPEG and Raptor packets are sent, these results are included in clause 12.

9.2 Sending arrangement considerations

An important issue in the evaluations was the way the different codes arrange data packets (source and FEC "repair" packets) for sending. Many different arrangements are possible for both codes. Since the arrangement can slightly impact the latency introduced by the FEC code with particular settings, and since these evaluations considered fixed latency budgets, the choice of sending arrangement affects the choice of parameters which are possible within the latency budget and therefore affects the bandwidth requirements of the codes.

An additional consideration with respect to sending arrangements is whether the resulting data stream has a constant bit-rate.

9.3 Bandwidth costs

9.3.0 Objectives

A primary objective of the simulations performed as part of this evaluation exercise was to measure the bandwidth overhead required to achieve a target quality of service. Although not the only evaluation criteria for AL-FEC, bandwidth consumption represents an ongoing cost of the solution for the operator: excessive bandwidth consumption may translate into lower service quality, fewer services or a smaller target market.

In order to assess bandwidth requirements, simulations were performed according to the agreed cases. For each case, the simulated time was 96 hours and the *mean time between packet loss* was measured. The minimum bandwidth required was assessed by performing repeated simulations, gradually increasing the FEC overhead until the target mean time between packet loss was achieved. Note that in the case of the Pro-MPEG code, increasing the bandwidth required that a different code was used - i.e. change in the L and D parameters and possibly change in the type of parity packets sent: row, column or both.

9.3.1 Loss models

Two loss models were used in the simulations, independent random packet loss and a loss model based on DSL Repetitive Electrical Impulse Noise (REIN).

The REIN model results in fixed length (8 ms) burst losses which are randomly placed in order to achieve an overall loss rate within the 10^{-6} to 10^{-3} loss range of interest. As such, the results below for the REIN case give a good indication of the code performance in the presence of burst losses.

9.3.2 Multicast case

9.3.2.0 Simulation results

For the multicast case, a maximum additional latency of 400 ms was used. The graphs below show the FEC overhead required to achieve a mean time between packet loss of four hours, plotted against packet loss for both independent random packet loss and Repetitive Electrical Impulse Noise simulated. The overhead calculation is based on the actual number of bytes sent, including IP and other headers, not just the ration of repair packets to source packets.

The figures also include a plot for an "Ideal Block Code" - this represents the theoretical lowest overhead which could achieve the target quality within the maximum latency using a block FEC code and gives a useful guide as to how much of the bandwidth dedicated to FEC is actually needed to provide the required FEC protection and how much is overhead due to inefficiency in the FEC code itself.

Note that the overhead scale in each graph may be different, to show the range of interest.

9.3.2.1 Results with constant sending arrangement




Figure 9.3.2.1.1: 2 Mbit/s MPEG-2 Transport Stream, 400 ms latency, random loss, constant sending




Figure 9.3.2.1.2: 6 Mbit/s MPEG-2 Transport Stream, 400 ms latency, random loss, constant sending




Figure 9.3.2.1.3: 2 Mbit/s MPEG-2 Transport Stream, 400 ms latency, REIN, constant sending

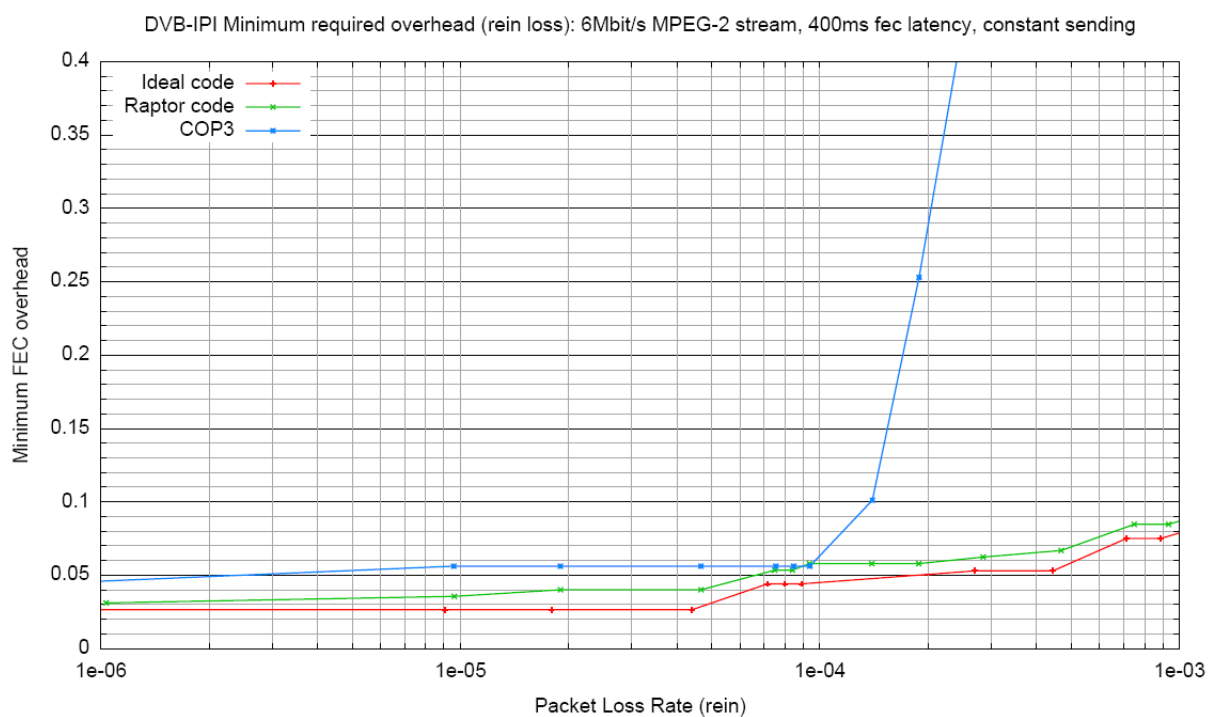


Figure 9.3.2.1.4: 6 Mbit/s MPEG-2 Transport Stream, 400 ms latency, REIN, constant sending

9.3.2.2 Results with burst sending arrangement

NOTE: Curves for the "Ideal" block code and Raptor below are for constant rate sending, compared with burst sending for Pro-MPEG.




Figure 9.3.2.2.1: 2 Mbit/s MPEG-2 Transport Stream, 400 ms latency, random loss, burst sending




Figure 9.3.2.2.2: 6 Mbit/s MPEG-2 Transport Stream, 400 ms latency, random loss, burst sending




Figure 9.3.2.2.3: 2 Mbit/s MPEG-2 Transport Stream, 400 ms latency, REIN loss, burst sending

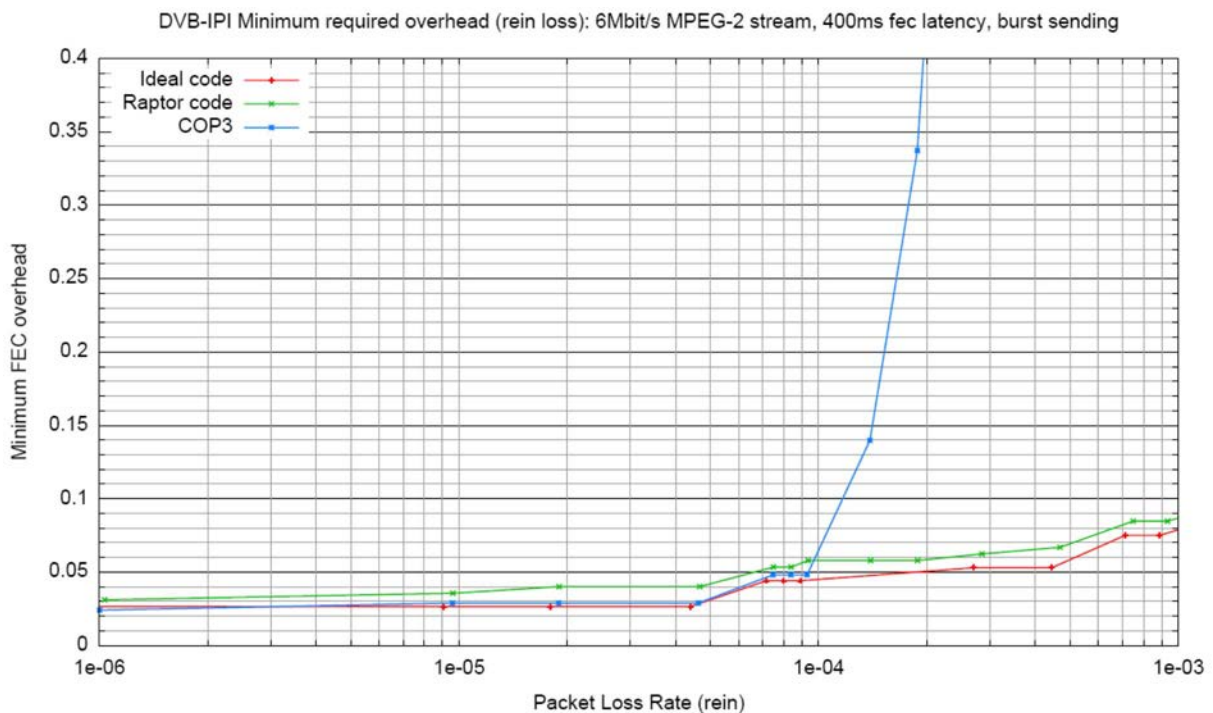


Figure 9.3.2.2.4: 6 Mbit/s MPEG-2 Transport Stream, 400 ms latency, REIN loss, burst sending

These simulations results lead to the following conclusions:

- The Raptor code consistently requires close to the minimum possible overhead for a block code (as illustrated by the red "ideal" plots).
- The overhead required for the Raptor code increases smoothly as the loss rate increases.

- A modest Raptor overhead of 9 % provides for FEC protection up to above 10^{-3} packet loss in both the random and REIN loss models.
- The Pro-MPEG COP3 code with constant sending rate performs close to the ideal code whenever PLR remains under a threshold value around 10^{-4} and only in the case of random loss this is the case since the Pro-MPEG row code is a simple parity code, which is optimal when only one packet of protection data is needed per block.
- Around 10^{-4} packet loss rate for the random loss case, the Pro-MPEG code requires higher overhead - around 34 % for the 2 Mbit/s stream and 20 % for the 6 Mbit/s stream.
- Depending on the sending arrangement, above around 3×10^{-4} packet loss for the REIN case, no settings for the Pro-MPEG code which supported the required quality target (measured in mean time between packet losses) could be found. Nevertheless, when using a slightly lowest quality target (same time but measured in mean time between FEC blocks with errors), it is possible to find Pro-MPEG settings to support the required quality target.
- The burst arrangement for the Pro-MPEG code requires somewhat less overhead at high loss rates, although still significantly more than Raptor.
- The burst sending arrangement for the Pro-MPEG code offers significant improvements in the REIN case - in fact improving on the ideal block code (which uses a constant sending arrangement).
- The choice of burst or constant sending arrangement for Raptor makes little difference in the required overhead.
- The burst sending arrangement for Pro-MPEG does not allow the quality target to be achieved in the REIN case across the whole loss range. It should be noted that simulations based on a lower quality target can be met by Pro-MPEG.

It should be noted that in the above cases the parameters for the Pro-MPEG code were selected to provide the best performance for each particular loss rate and pattern through a wide search of the possible parameter set. In practice, loss rates and error patterns are expected to be mostly unknown in advance.

In particular, for the REIN cases, the Pro-MPEG column code with a number of columns equal to the burst length provides adequate protection so long as events with two error bursts within a protection period happen only once every four hours or less.

This may happen when the overall loss rate is high or when there is strong correlation between bursts. Moreover if random single loss errors happen very close to a burst, they may not be corrected neither.

9.3.3 Unicast case

9.3.3.1 Stored/buffered content

In these cases, content is available at the server in advance of sending to the user: for VoD services the content is stored in its entirety and for live broadcast in trick modes the content is buffered for at least a few hundred ms when the user activates the trick mode by pausing the multicast broadcast.

In these cases the Raptor code incorporates a fast buffer fill technique (called "faststart" in this paper) which allows the protected block size to be gradually increased over the first few seconds of transmission. Note that this technique is possible only because of the independence of block size and overhead supported by Raptor and the possibility to flexibly vary the overhead in single packet increments without impacting the error correction performance of the code.

As above, repeated 96 hours simulations were performed with the FEC overhead again increased for each simulation until the target quality was achieved. The fast-start procedure is repeated every 10 minutes during the simulation to model the impact of repeated channel change or use of trick-modes.




Figure 9.3.3.1.1: 2 Mbit/s MPEG-2 Transport Stream, 100 ms latency (stored/buffered content), random loss




Figure 9.3.3.1.2: 6 Mbit/s MPEG-2 Transport Stream, 100 ms latency (stored/buffered content), random loss




Figure 9.3.3.1.3: 2 Mbit/s MPEG-2 Transport Stream, 100 ms latency (stored/buffered content), REIN

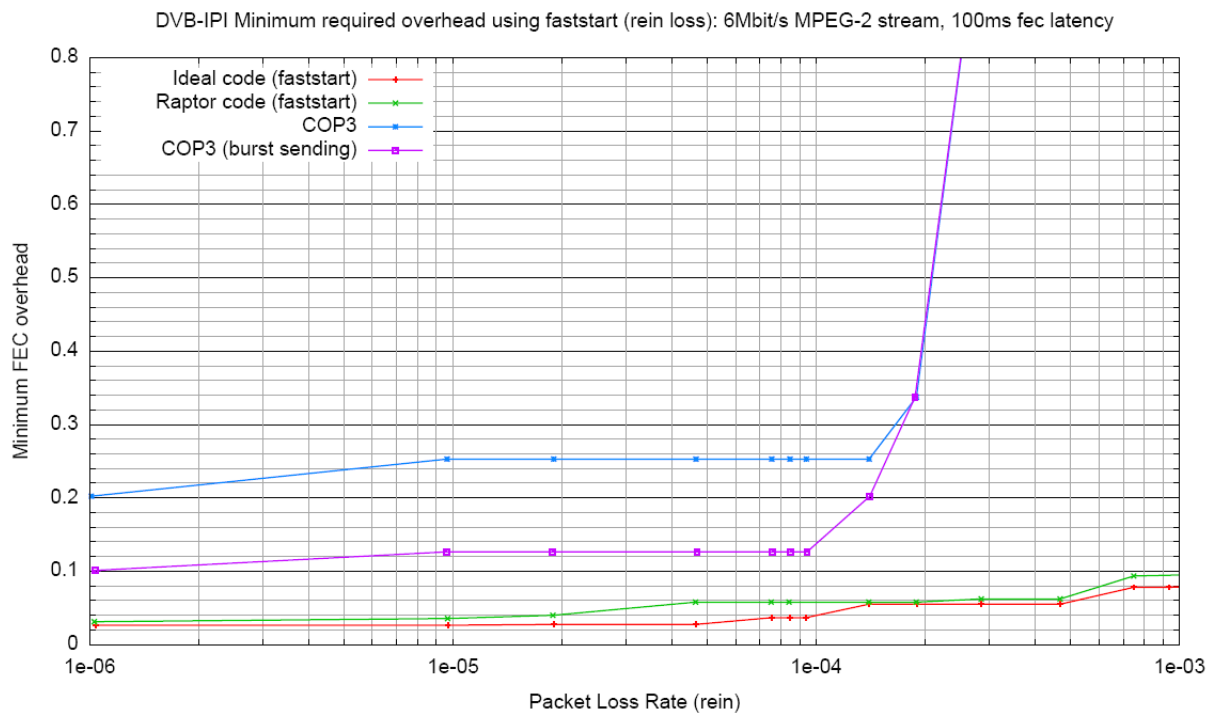


Figure 9.3.3.1.4: 6 Mbit/s MPEG-2 Transport Stream, 100 ms latency (stored/buffered content), REIN

9.3.3.2 Live content

9.3.3.2.0 Simulation results for unicast delivery

In the case of unicast delivery of live content (for example in networks which do not support multicast) then the block size for the Raptor code is limited by the requirement of a maximum latency due to FEC of 100 ms. The figures 9.3.3.2.1.1, 9.3.3.2.1.2, 9.3.3.2.1.3 and 9.3.3.2.1.4 show simulation results for this case.

9.3.3.2.1 Constant sending arrangement

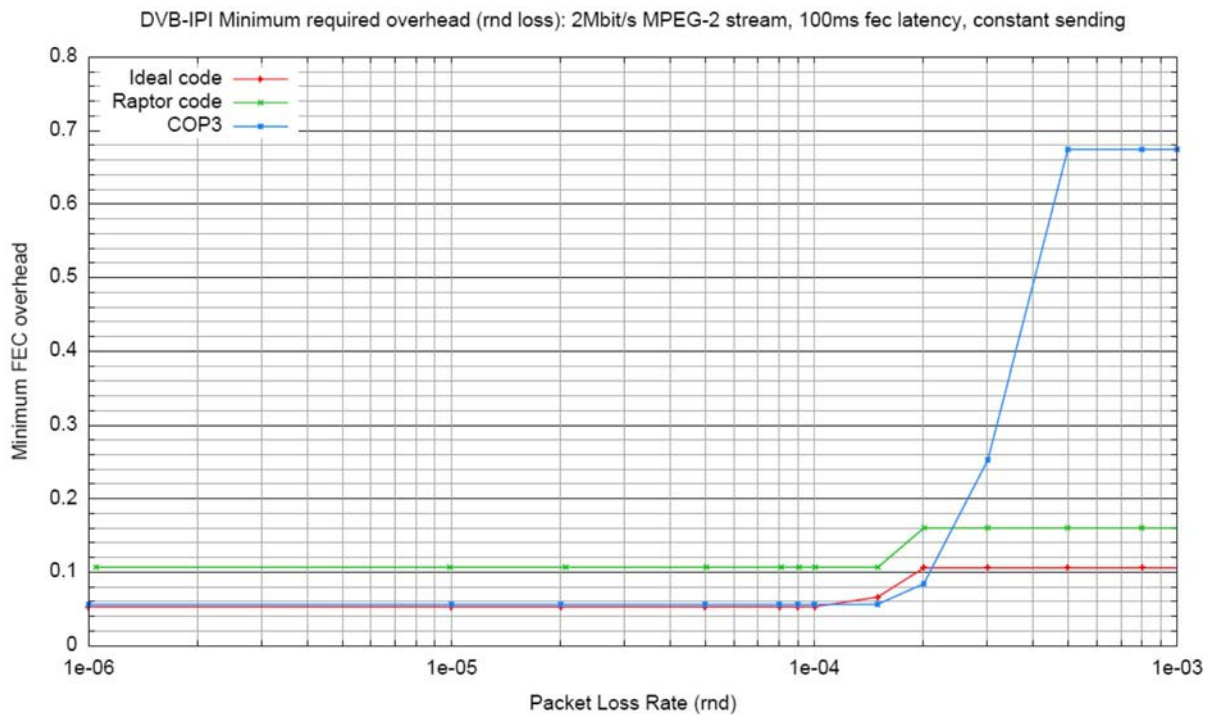


Figure 9.3.3.2.1.1: 2 Mbit/s MPEG-2 Transport Stream, 100 ms latency (live content), random loss, constant sending




Figure 9.3.3.2.1.2: 6 Mbit/s MPEG-2 Transport Stream, 100 ms latency (live content), random loss, constant sending




Figure 9.3.3.2.1.3: 2 Mbit/s MPEG-2 Transport Stream, 100 ms latency (live content), REIN, constant sending



Figure 9.3.3.2.1.4: 6 Mbit/s MPEG-2 Transport Stream, 100 ms latency (live content), REIN, constant sending

9.3.3.2.2 Burst sending

NOTE: Curves for the "Ideal" block code and Raptor below are for constant rate sending, compared with burst sending for Pro-MPEG.




Figure 9.3.3.2.2.1: 2 Mbit/s MPEG-2 Transport Stream, 100 ms latency (live content), random loss, burst sending




Figure 9.3.3.2.2.2: 6 Mbit/s MPEG-2 Transport Stream, 100 ms latency (live content), random loss, burst sending




Figure 9.3.3.2.2.3: 2 Mbit/s MPEG-2 Transport Stream, 100 ms latency (live content), REIN loss, burst sending

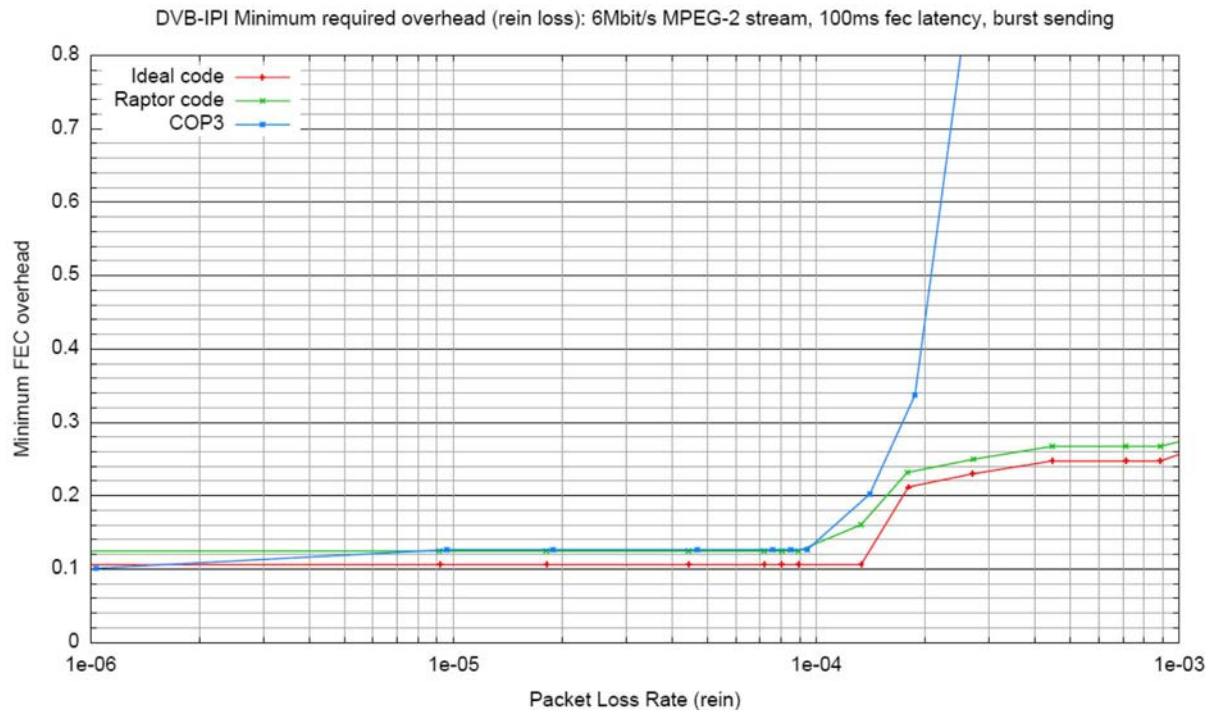


Figure 9.3.3.2.2.4: 6 Mbit/s MPEG-2 Transport Stream, 100 ms latency (live content), REIN loss, burst sending

As in previous cases, the Raptor code meets the quality target at all error rates with overhead close to the minimum possible. The Pro-MPEG code meets the quality target with minimum overhead only in cases where the loss rate is below a threshold which is around 10^{-4} packet loss rate.

With the constant sending arrangement, and REIN losses, the Raptor codes requires an overhead which is less than or (approximately) equal to the Pro-MPEG overhead for all loss rates. For other cases (burst sending and/or random loss) the Pro-MPEG code requires marginally less overhead for the loss rates which are below the threshold.

For low loss rates and in the presence of random loss, the Pro-MPEG code is simple a 1D parity code, which is well known to be ideal. In these cases Pro-MPEG achieves lower overhead than Raptor.

9.3.4 A note on latency, jitter and traffic shaping

All the above simulations assume that the sent traffic should maintain a constant bit-rate (although it is accepted that the constant-bitrate Pro-MPEG scheme actually doubles the instantaneous bit-rate each time a repair packet is sent, this is only visible as a variation in bit-rate over very short time periods. However for the burst sending arrangement, the variation is significant and over a longer period of time).

In order to support legacy receivers in the case of multicast, whenever this is feasible, the use of FEC should not introduce significant additional jitter in the source packets. Using the sending arrangement proposed for Raptor codes does introduce a small amount of additional jitter to the arrival of source packets at the receiver. Using the constant sending arrangement proposed for Pro-MPEG avoids such jitter, however using the burst sending arrangement proposed for Pro-MPEG will introduce a small amount of additional jitter as the bursts are traffic shaped on the access link. Sending arrangements are interchangeable between the codes, so there are many possibilities. See clause 6 for more details. Clause 10 gives details of the sending arrangements used in the simulations.

In the simulations above, the maximum additional jitter in the case of Raptor is around 40 ms for the 400 ms latency cases and in most cases significantly less. Finally, "latency" in these simulations has been interpreted as the additional latency introduced between the source and the playout due to the use of FEC. This is equivalent to the size of the FEC data buffer assumed to exist at the receiver. This latency adds directly to the response time for user actions, such as channel change, re-wind, forward-wind, etc.

In the case of live content, the Raptor scheme as proposed adds a small additional amount to the time between the event actually occurring at the sender and the presentation to the user (distinct from the response time for user actions, referred to above). In the cases above this is at most around 40 ms and in general considerably less. Since the overall end-to-end delay is general much higher than 40 ms, this additional delay is not considered significant, especially since it does not contribute to the response time for user actions. The Raptor scheme is sufficiently flexible that this delay could be reduced if required. Targets on this end-to-end delivery time have not been discussed and again could be included in a further phase of this evaluation if necessary, but again it is unlikely to significantly affect the results.

Finally, the only two latency figures (100 ms and 400 ms) were tested in these evaluations. It is instructive to consider the trade-off involved in selection of an FEC latency figure. Lower latency results in shorter channel change time but has a cost in that a higher FEC overhead is required for a given level of protection. Conversely, a longer latency budget results in longer channel change time in return for a lower FEC overhead. Figure 9.3.4.1 illustrates this trade-off for an "ideal" code and for several quality targets ("Mean Time Between Artifacts"). Figure 9.3.4.1 suggests that a significant bandwidth saving is available if the latency budget is increased from 100 ms to (say) 200 ms, but that there is little to be gained by increasing the latency above 400 ms. In particular, figure 9.3.4.1 throws doubt on the practical validity of the 2 MBit/s, 100 ms case evaluated above: an operator who was sufficiently bandwidth-constrained to use 2 Mbit/s encoding would surely also take advantage of the FEC bandwidth savings that could be achieved with a 200 ms latency budget.




Figure 9.3.4.1: Latency/FEC bandwidth trade-off

9.3.5 Summary of simulation results

The results above according to the sending arrangement and type of loss can be summarized as follow:

Summary for multicast and unicast live video:

- There is a "loss rate threshold" in each case: below this threshold, the Pro-MPEG overhead is very low and close to Raptor (sometimes higher, sometimes lower) and above this threshold, the Pro-MPEG overhead is significant (always much higher than Raptor overhead).
- The threshold is around $1e-4$ Packet Loss Rate (actually between $5e-5$ and $2e-4$), depending on the case.

Constant sending arrangement, random loss:

- Below the threshold, the Pro-MPEG overhead is slightly less than the Raptor overhead and above this threshold, the Raptor overhead is much less than the Pro-MPEG overhead.

Constant sending arrangement burst (REIN) loss:

- Below the threshold, the Raptor overhead is slightly less than the Pro-MPEG overhead and above this threshold, the Raptor overhead is much less than the Pro-MPEG overhead. Please note that in this case, Raptor overhead is always the lowest.

Burst sending arrangement, random loss:

- Burst sending does not have much effect on results below the threshold.
- The Pro-MPEG overhead is reduced above the "threshold" compared to constant sending arrangement, but is still much greater than the Raptor threshold.
- Burst sending does not have much effect on the Raptor overhead.

Burst sending arrangement, burst (REIN) loss:

- The Pro-MPEG overhead is reduced both above and below the "threshold", but above the threshold is still much greater than the Raptor threshold.
- Below the threshold the Pro-MPEG overhead is slightly less than the Raptor overhead.
- Burst sending does not have much effect on the Raptor overhead.

Summary for unicast stored or buffered content:

- In the particular case of unicast stored or buffered content, Raptor code can use the faststart sending arrangement so as to use significantly less bandwidth than Pro-MPEG in all cases.
- When faststart mechanism is not used, results are the same as multicast and unicast live video.

In all cases, the results plotted above show the overhead required by the "best" configuration parameters for the Pro-MPEG COP3 code according to guidelines for setting Pro-MPEG parameters and the specification in ETSI TS 126 346 [2]. These were chosen by searching through the various possible configurations (including row packets only, column packets only, both row and column packets and different matrix sizes) and reporting only the lowest overhead which achieved the required quality. This means that the choice of code was based implicitly on complete knowledge of the loss rates and patterns in each case.

In summary, the requirements on network quality (target end-to-end loss rates) depend significantly on the choice of FEC code (Pro-MPEG or Raptor): network quality requirements are much more stringent if Pro-MPEG is chosen since it works well only as long as the packet loss rate remains under the previously defined threshold (around $1e-4$).

9.4 Flexibility

The FEC evaluation criteria for flexibility states:

"Flexibility:

- Changing the overhead or/and the block size dynamically (within or between FEC blocks).
- Range of protection periods.
- Suitability for use with a wide variety of FEC sending strategies'.

The Raptor code provides complete flexibility in terms of overhead (protection amount) and block size (protection period). These parameters can be set independently according to application requirements and the error correction performance of the code remains just as close to "ideal" whatever the parameter settings. Parameter settings can easily be changed dynamically and protection periods from 10 s to 1 000 s of milliseconds can be efficiently supported.

For the Pro-MPEG code, the protection period and protection amount are related and constrained and in practice only certain combinations are supported. Nevertheless, the possible number of combinations is large enough to offer many different levels of protections.

9.5 Processing and Memory requirements

The Raptor code has been designed to have very modest computational complexity such that it is easy to implement in software on resource constrained devices such as Set-Top Boxes and mobile devices. Techniques for efficient hardware implementation for high capacity encoders have also been presented and many options exist for hardware-assisted implementations for decoders.

The Pro-MPEG code has been designed to have very low computational complexity such that it is easy to implement it in software or in hardware.

For both Raptor and Pro-MPEG, the complexity of encoding is comparable with the complexity of decoding. For Raptor, both scale linearly with the volume of data to be encoded/decoded, making the overall computational requirements proportional the service bit-rate and to a large extent independent of the losses or level of protection.

Raptor encoding complexity for the scenarios considered here is in the region of 2 MIPS per Mbit/s - so a 6 Mbit/s stream would require ~12 MIPS of processing power to encode, although in practice the encode time is also dependent on memory bus speed and cache/DMA availability. For example, Digital Fountain has demonstrated an off-the-shelf rack-mounted server with a Pentium processor running at 3 GHz performing Raptor encoding at 2 Gbit/s - the equivalent of one thousand 2 Mbit/s video streams. Further optimizations for the specific case of video stream encoding and platform-specific optimization could be expected to increase this encoding speed significantly. Leading Pro-MPEG COP3 processing cards encode at around 400 Mbit/s and so similar performance could be easily achieved with Raptor with modest processing requirements.

Hardware optimizations of Raptor codes in the form of hardware assist for XOR operations or complete implementation of the code in hardware are also possible and can further improve capacity. The application of the Raptor code for streaming has been designed so that for a given stream rate/latency the block size and structure from the encoders point of view is the same for every block. Thus the sequence of operations required to encode repair packets for a block can be calculated or stored in advance and executed quickly (in software or hardware) for each block. This is true even if the actual block size (in terms of packets) differs between protection periods.

The number of primitive symbol XOR operations required for Raptor encoding or decoding for the scenarios considered here is around 12 to 14 operations for each source symbol.

The number of primitive symbol XOR operations required for Pro-MPEG encoding or decoding for the scenarios considered here is 1 operation for each source symbol in Pro-MPEG 1D and 2 operations for each source symbol in Pro-MPEG 2D.

Nevertheless, in practice, for each symbol, these operations are performed on-chip (in cache) and so the bottleneck is the speed with which data can be moved between memory and the processor, rather than the precise number of XOR operations. All modern processors employ pipelining and so can perform the XOR operations on-chip concurrently with moving data for future operations between off- and on-chip memory. This means a reduction in XOR operations does not necessarily translate into a significant increase in speed of encoding or decoding.

With Raptor, minimum memory requirements for data to be encoded/decoded at both encoder and decoder are slightly greater than the source block size. At the decoder, received data (which is a mix of source data and repair data) may be transformed "in-place" into the recovered source block. Thus, these memory requirements are less than 350 KB for the largest block size considered in this evaluation.

With Pro-MPEG, the encoder only needs to have buffers so as to store the repair packets of a protection block. Since amount of protection is always much lower than the amount of data, it means a Pro-MPEG encoder requires memory much smaller than the source block size. On the decoder, Pro-MPEG only requires enough memory to store the current protection block and its repair packets. Therefore it means a Pro-MPEG decoder requires memory slightly greater than the source block size. Note also, that depending on the sequencing arrangement used, the decoder may need more memory. For instance, when repair packets are arranged within the block after the one they protect, the decoder would need twice as much memory to store the current and following protection blocks.

Note that for decoders, this memory requirement is still very modest compared to the memory required, for example, for storing a single HD frame after decoding.

9.6 Additional criteria

The following additional criteria are included in the evaluation criteria document:

- Continued functioning of existing STB products in presence of FEC data.
- Option for new STB products to use or ignore FEC data.
- Confirmation of FEC scheme IPR compliance with DVB rules.
- Support of combined protection of different streams (such as when audio and video packets are sent in two separate streams).

Raptor is compliant to all these criteria.

Pro-MPEG is compliant to the first two criteria and believed to be compliant to the third (IPR compliance is currently being clarified by SMPTE).

The Pro-MPEG code does not support combined protection of different streams - separate protection streams are required for each RTP flow. Specifically in the case of audio streams, which have much lower bandwidth than the video streams, then high quality protection will be extremely difficult to achieve if latency needs to be kept very small.

In general, combined protection is more efficient than separate protection and in particular separate protection of the relatively low bit-rate audio stream can be extremely inefficient.

Combined protection can also encompass the RTCP packets that provide time synchronization information between the audio and video streams.

9.7 Content Download

It has been suggested that the FEC solution chosen for streaming services should also be suitable for use in content download applications. It should be noted that it has not yet been agreed, (or even discussed in detail), that Forward Error Correction is required for Content download - other solutions do exist.

However, solutions based on forward error correction have a number of significant advantages over other solutions in the multicast case. The Raptor code proposed for DVB-IPTV streaming applications is highly suitable for content download applications as well (and has been adopted for such applications by 3GPP and DVB CBMS). The same code could therefore be used for both streaming and content download.

No description is available of whether and how the Pro-MPEG code could be applied to content downloading: it was clearly designed for streaming services in extremely low packet loss cases only. The Pro-MPEG code is by nature a short block code and for content downloading a large block code is much more efficient if FEC is to be used.

9.8 Raptor vs. Pro-MPEG Summary

Table 9.8.1 summarizes the results described above. The green font identifies the best result while the red font identifies the worst result. When the result between codes is very close, an orange font is used to identify the code that only performs slightly less well.

Table 9.8.1

Criteria	Pro-MPEG Constant	Pro-MPEG Burst	Raptor	Comments
Bandwidth cost - loss rates > ~1e-4				
- SD MPEG-2 TS broadcast (400 ms)	High	High	Low	
- HD MPEG-2 TS broadcast (400 ms)	High	High	Low	
- SD MPEG-2 TS unicast (100 ms)	High	High	Low	Thanks to its fast-start mechanism, Raptor achieves very low overhead in case of stored/buffered content.
- HD MPEG-2 TS unicast (100 ms)	High	High	Low	Thanks to its fast-start mechanism, Raptor achieves very low overhead in case of stored/buffered content.
Bandwidth cost - loss rates < ~1e-4				
- SD MPEG-2 TS broadcast (400 ms)	Low	Lowest	Low	
- HD MPEG-2 TS broadcast (400 ms)	Low	Lowest	Low	
- SD MPEG-2 TS unicast (100 ms)	Modest	Lowest	Low	Thanks to its fast-start mechanism, Raptor achieves very low overhead achieved in case of stored/buffered content.
- HD MPEG-2 TS unicast (100 ms)	Modest	Lowest	Low	Thanks to its fast-start mechanism, Raptor achieves very low overhead achieved in case of stored/buffered content.

Criteria	Pro-MPEG Constant	Pro-MPEG Burst	Raptor	Comments
Support of target quality for evaluated packet loss range/patterns	See corresponding comment		Yes	Pro-MPEG COP3 could not provide a Mean Time Between Packet Loss of 4 hours for a number of the burst loss cases. However, a slightly weaker target of Mean Time Between Artifacts (visible errors) of 4 hours could be achieved.
Further packet losses that could occur in the core network due to congestion and/or the home environment e.g. wireless technologies	-	-	-	Not yet evaluated.
Flexible engineering of code parameters	Yes (but fixed number of combinations and direct correlation between overhead and protection block size)		Yes (fully)	
Computational complexity	Lowest		Modest	
Scalability (e.g. encoding of 1 000 s of streams)	Yes		Yes	
Memory requirements (encoder)	Lowest		Modest	
Memory requirements (decoder)	Modest		Modest	
Visibility of artifacts after FEC decoding	-	-	-	Both codes could perform partial correction.
Continued functioning of existing STB products in presence of FEC data	Yes		Yes	
Option for new STB products to use or ignore FEC data	Yes		Yes	
Confirmation of FEC scheme IPR compliance with DVB rules	Yes		Yes	Pro-MPEG IPR compliance is currently under SMPTE process.
Efficient support of direct encapsulation of audio/video in RTP (as defined in ETSI TS 102 005 [4]): Support of combined protection of audio and video packets	No		Yes	Raptor can protect several RTP and RTCP streams together whereas Pro-MPEG has to consider each RTP and RTCP streams separately.
Efficient support of direct encapsulation of audio/video in RTP (as defined in ETSI TS 102 005 [4]): support of variable length packets	Yes (but less efficient)		Yes	
Suitable for Content Download Service	No (much less efficient)		Yes	

9.9 Conclusions

The sending arrangement chosen has a significant impact on the performance / bandwidth cost.

The comparison of the two codes also differs depending on the packet loss rate.

In the case that burst sending is used and for loss rates below a threshold (between $5e-5$ and $2e-4$), the Pro-MPEG code requires slightly less bandwidth than Raptor code.

In the case that burst sending is not used and for loss rates below a threshold (between $5e-5$ and $2e-4$), both Pro-MPEG and Raptor codes requires similar bandwidth overhead although there are differences depending on the precise case (see clause 9.3.5).

For loss rates above a threshold (between $5e-5$ and $2e-4$), Raptor code requires much less bandwidth than Pro-MPEG code.

The threshold identified through these simulations depends on quality target, source stream bitrate, latency budget and loss patterns.

When the Raptor fast-start mechanism is used for unicast/buffered content, Raptor requires less overhead than Pro-MPEG.

Regarding implementation aspects (complexity, memory requirements, etc.), though there are differences between codes (see clause 9.8), no significant issues were identified with either code.

Both codes meet the requirement for backward compatibility with existing equipments.

The Raptor code supports various future requirements which the Pro-MPEG does not (see clause 9.8).

Since neither of these two codes is optimal in all cases, an hybrid code with performance similar to the best of either was defined (see clause 12 for simulation results).

10 Sending arrangements used for simulations

10.1 DF Raptor™ default sending arrangement

The sending arrangement proposed for the DF Raptor™ code is illustrated in figure 10.1.1. In this sending arrangement the overall sending rate is kept constant and the source packets of each block are sent before any of the repair packets of the block. This approach requires that the sending rate of the source packets be increased marginally to make space for the repair packets at the end of the block.

It is important to note that the sequencing of packets is determined by the FEC procedures which operate "below" the RTP layer. The contents of the packets, in particular the RTP timestamps, are not modified compared to the contents in the case in which FEC is not applied and therefore the correct timing for the packets can be reconstructed with the usual procedures.




Figure 10.1.1: DF Raptor™ sending arrangement

Note that while this arrangement ensures a global constant bitrate, it actually modifies the rate at which source packets are sent and consequently creates a small amount of additional jitter on the transmission.

Other sending arrangements are also possible for DF Raptor™ but were not investigated.

Pros and cons:

- + Global sending rate is constant.
- + Full latency budget available for FEC protection.
- Source data sending rate is different from original source data sending rate.

- Insertion of repair packets introduces small amount of jitter on all source packets.

10.2 Pro-MPEG COP3 fully interleaved sending arrangement

Annex C of the Pro-MPEG specification [3] proposes a sending arrangement as illustrated in figure 10.2.1. In this sending arrangement the overall sending rate is kept constant and the sending rate of source packets is also kept constant.

Because this sending arrangement distributes repair packets for one block over the entire duration of the next block, then the maximum block size is limited to one half of the latency budget. As a result, the overhead required by the code is increased. This is illustrated in the "constant sending arrangement" is illustrated in figure 10.1.1.




Figure 10.2.1: Pro-MPEG COP3 fully interleaved sending arrangement

Pros and cons:

- + Source data sending rate is the same as original source data sending rate.
- + Global sending rate is kept constant.
- Only half of latency budget is available for FEC protection.
- Insertion of repair packets introduces very small amount of jitter at the beginning when total stream bandwidth is close to available channel bandwidth.

10.3 Pro-MPEG COP3 burst sending arrangement

This arrangement is illustrated in figure 10.3.1. In this case, repair packets for one block are interleaved with the first few packets of the next block. As a result, the instantaneous sending rate during these first few packets is significantly increased. However, the block size may now be set almost as large as the latency budget, which reduces the required overhead. This is illustrated in the "burst sending" as illustrated in figure 10.2.1.

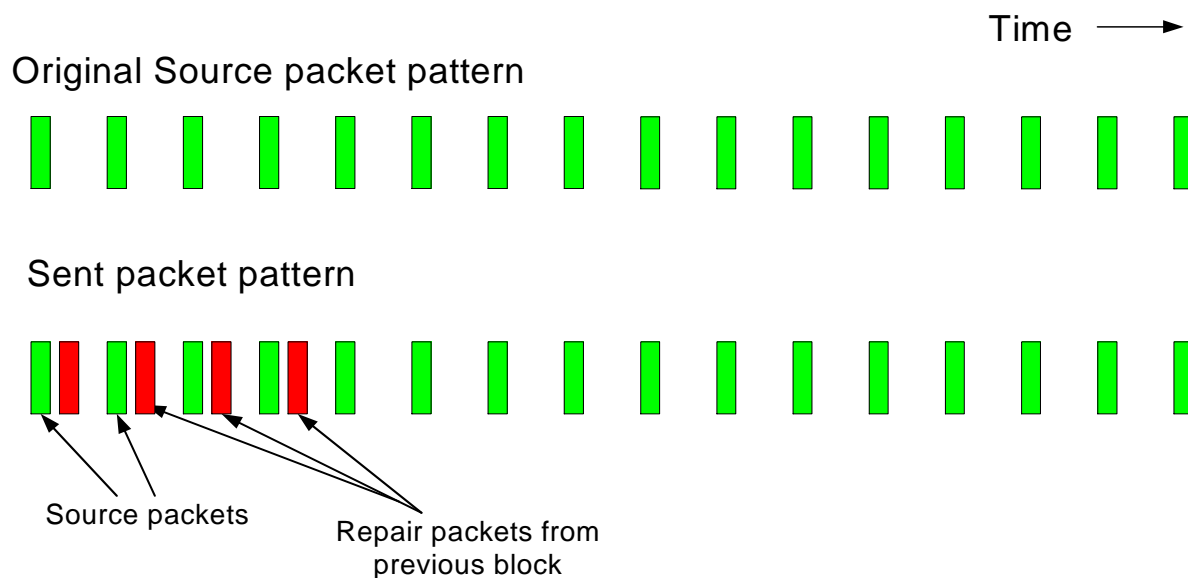


Figure 10.3.1: Pro-MPEG COP3 burst sending arrangement

Pros and cons:

- + source data sending rate is the same as original source data sending rate.
- + Almost all of latency budget is available for FEC protection.
- Global sending rate is very bursty (and therefore not constant).
- Insertion of repair packets introduces small amount of jitter at the beginning when total stream bandwidth is close to available channel bandwidth.

10.4 Concurrent Interleaved sending

In the case of Video on Demand, or if additional latency at the encoder is acceptable, a sending arrangement as depicted in figure 10.4.1 is possible. In this case, repair packets are interleaved within the block that they protect. This is possible in the Video on demand case because the data to be protected is available for FEC calculations to be performed slightly in advance of sending the data. Alternatively, a live stream can be buffered at the encoder for long enough for the FEC calculations to be performed before beginning to send the source packets of the block.

This sending arrangement could also be used for live content with a penalty that buffering equal to the block size would be required at the sender. This buffering contributes additional end-to-end delay to the playout of live streams i.e. the delay between a live event occurring and being presented on the user's screen. However it would not contribute additional channel change delay. This option may be important if there is existing equipment which is affected by changes in the timing of source packets. The procedures for timing recovery specified in ETSI TS 102 034 [1], annex A allow MPEG 2 timing to be recovered even in the presence of significant IP packet arrival jitter - however, if these procedures have not been correctly implemented then equipment may be adversely affected by the additional jitter introduced by some of the other sending arrangements described here.

This sending arrangement has the desirable properties that both the source packet data rate and the total data rate are constant. However, in the Pro-MPEG case, unlike the constant data rate arrangement in clause 10.2, the whole latency budget can be used for a single source block.

New simulation results are presented for this sending arrangement in clause 11. Note that only the Pro-MPEG column code was tested, not the 2D code.

For random loss, the results are similar to the comparison between Raptor with constant sending and Pro-MPEG with burst sending - i.e. Pro-MPEG uses slightly less overhead below the loss rate threshold than Raptor does. However, for burst loss, the Pro-MPEG code is significantly affected by interleaving of repair packets with the source packets they protect. For the 2 Mbit/s stream, this pushes the threshold where Pro-MPEG performs well down to $1e-5$ or below. For the 6 Mbit/s stream, the quality target was not achievable: it is easy to see why, since a burst loss of 6 source packets will often hit a repair packet as well, and it is not possible with only 6 repair packets per block to avoid that the burst hits a source packet that is protected by that repair packet.




Figure 10.4.1: Interleaved sending for VoD

Pros and cons:

- + Source data sending rate is the same as original source data sending rate.
- + All of latency budget is available for FEC protection.
- Global sending rate is kept constant.
- Insertion of repair packets introduces very small amount of jitter at the beginning when total stream bandwidth is close to available channel bandwidth.
- Not resilient to burst losses for the Pro-MPEG FEC.

10.5 DF Raptor™ faststart sending for stored/buffered content

An additional sending arrangement for stored or buffered content (i.e. VoD and trick modes on live content) was proposed and simulated for DF Raptor™. This sending arrangement is illustrated in figure 10.5.1. In this arrangement, source data is sent slightly faster than the nominal stream rate at the start of the session or when trick modes are used. This allows the buffering period to be gradually increased without introducing additional channel change latency.

Two variants of this approach were simulated:

- "faststart with constant rate sending" - in which the additional source data bandwidth is obtained by reducing the FEC bandwidth at the beginning of the stream. As a result the total stream rate remains constant, but stream quality is reduced for these few initial seconds.
- "faststart with variable rate sending" - in which the overall stream rate at the beginning of the stream is somewhat higher than the nominal stream rate (e.g. 20 % higher) for the initial few seconds of the stream, but as a result the stream quality is maintained.

The second variant provided the best results.

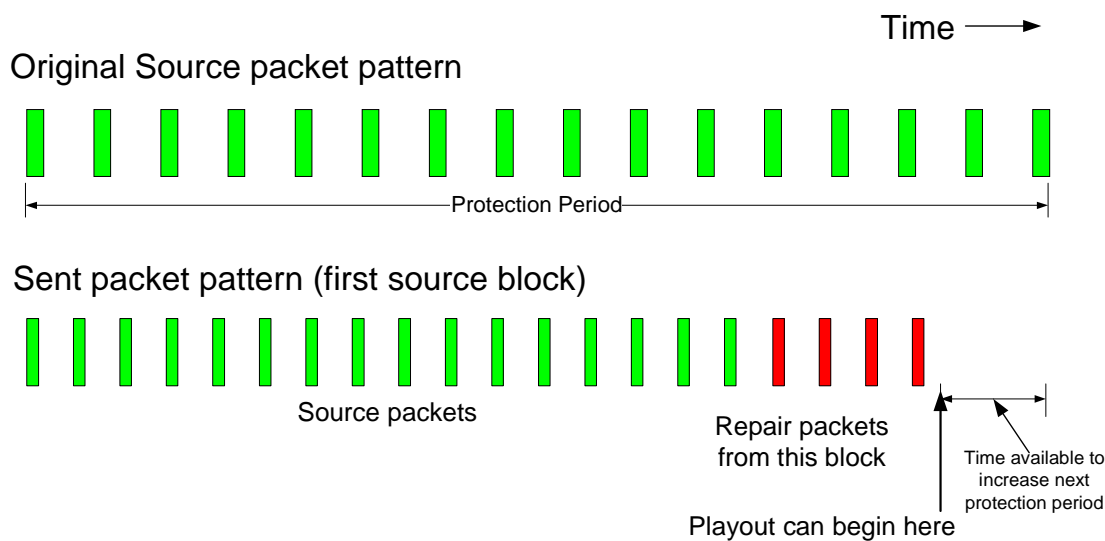


Figure 10.5.1: DF Raptor™ faststart sending arrangement

Pros and cons:

- + FEC protection period can be increased to much greater than the latency budget.
- Only applicable to unicast/buffered content for Raptor.

11 Concurrent interleaving results

This clause presents simulation results for the sending arrangement described in clause 10.4 in which both the source packet rate and the total stream rate are kept constant, whilst also allowing the full latency budget to be used for the FEC block.

Note that, due to lack of time, these results do not include the Pro-MPEG 2D code. It might be expected that in some of the cases where a result is not shown with the 1D code then the 2D code could provide the target quality, but at a relatively high overhead.

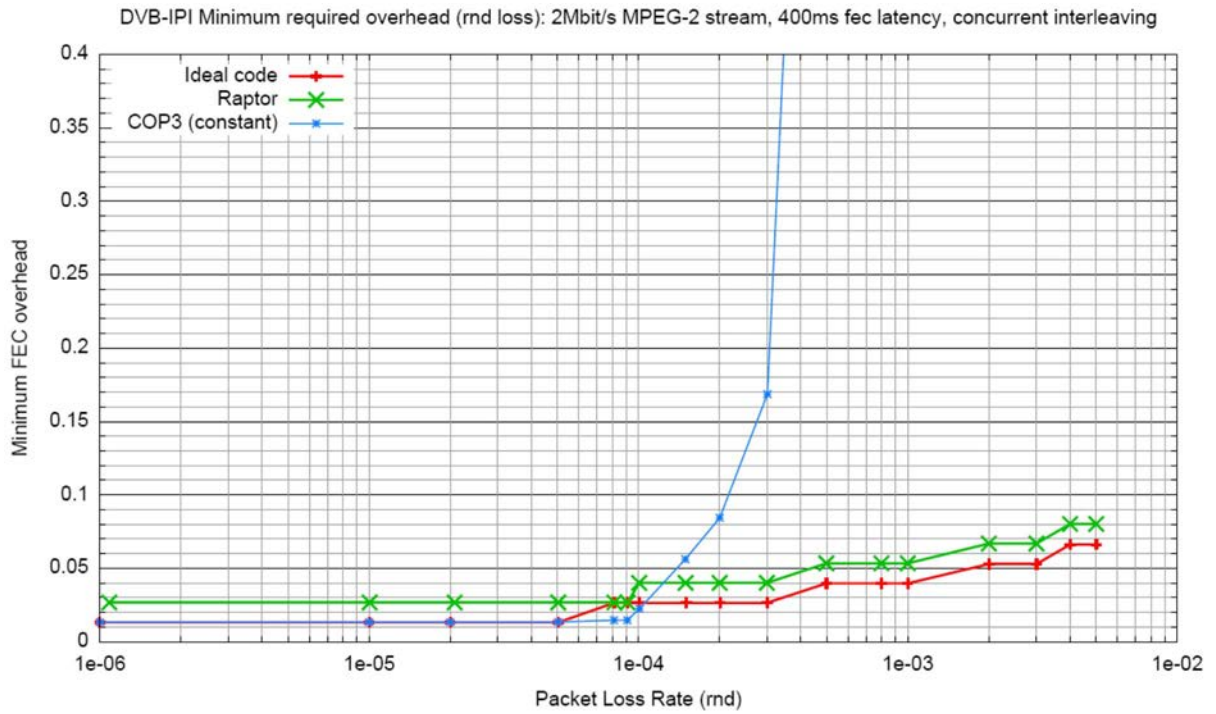


Figure 11.1: 2 Mbit/s MPEG-2 Transport Stream, 400 ms latency, random loss, concurrent interleaving




Figure 11.2: 6 Mbit/s MPEG-2 Transport Stream, 400 ms latency, random loss, concurrent interleaving




Figure 11.3: 2 Mbit/s MPEG-2 Transport Stream, 400 ms latency, REIN Loss, concurrent interleaving



Figure 11.4: 6 Mbit/s MPEG-2 Transport Stream, 400 ms latency, REIN Loss, concurrent interleaving




Figure 11.5: 2 Mbit/s MPEG-2 Transport Stream, 100 ms latency, random loss, concurrent interleaving

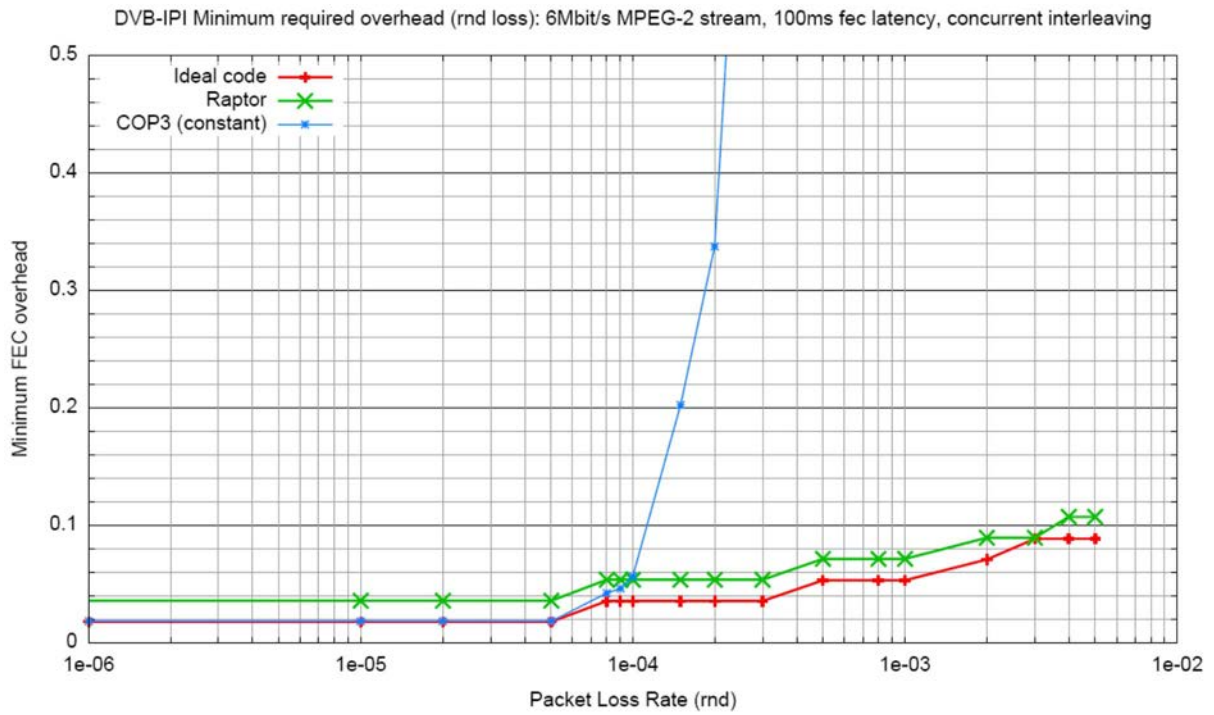


Figure 11.6: 6 Mbit/s MPEG-2 Transport Stream, 100 ms latency, random loss, concurrent interleaving




Figure 11.7: 2 Mbit/s MPEG-2 Transport Stream, 100 ms latency, REIN Loss, concurrent interleaving




Figure 11.8: 6 Mbit/s MPEG-2 Transport Stream, 100 ms latency, REIN Loss, concurrent interleaving

12 Hybrid code

12.0 General

A hybrid of the Pro-MPEG 1D column code and the Raptor code was proposed in order to provide a single scalable FEC solution with performance similar to the best of either the Pro-MPEG or Raptor codes in any given case.

12.1 Hybrid code results

This clause presents results for the Hybrid code. The hybrid cases are denoted "Raptor P<n>" where <n> is the number of parity packets used. The value of <n> chosen in each case is the smallest such that the quality target can be achieved with Pro-MPEG packets alone at loss rates of 1e-5 and lower.

The sending arrangement of clause 10.1 was used for these simulations.




Figure 12.1.1: 2 Mbit/s MPEG-2 Transport Stream, 400 ms latency, random loss, constant sending




Figure 12.1.2: 6 Mbit/s MPEG-2 Transport Stream, 400 ms latency, random loss, constant sending




Figure 12.1.3: 2 Mbit/s MPEG-2 Transport Stream, 400 ms latency, REIN Loss, constant sending




Figure 12.1.4: 6 Mbit/s MPEG-2 Transport Stream, 400 ms latency, REIN Loss, constant sending




Figure 12.1.5: 2 Mbit/s MPEG-2 Transport Stream, 100 ms latency, random loss, constant sending




Figure 12.1.6: 6 Mbit/s MPEG-2 Transport Stream, 100 ms latency, random loss, constant sending




Figure 12.1.7: 2 Mbit/s MPEG-2 Transport Stream, 100 ms latency, REIN Loss, constant sending




Figure 12.1.8: 6 Mbit/s MPEG-2 Transport Stream, 100 ms latency, REIN Loss, constant sending

History

Document history		
V1.1.1	November 2006	Published as ETSI TR 102 542
V1.2.1	April 2008	Published as ETSI TS 102 542
V1.3.1	January 2010	Publication
V1.3.2	May 2011	Publication
V2.1.1	April 2016	Publication