

ETSI TS 102 831 V1.1.1 (2010-10)

Technical Specification

Digital Video Broadcasting (DVB); Implementation guidelines for a second generation digital terrestrial television broadcasting system (DVB-T2)



Reference

DTS/JTC-DVB-236

Keywords

broadcasting, digital, DVB, modulation, TV

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

Individual copies of the present document can be downloaded from:

<http://www.etsi.org>

The present document may be made available in more than one electronic version or in print. In any case of existing or perceived difference in contents between such versions, the reference version is the Portable Document Format (PDF). In case of dispute, the reference shall be the printing on ETSI printers of the 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

<http://portal.etsi.org/tb/status/status.asp>

If you find errors in the present document, please send your comment to one of the following services:

http://portal.etsi.org/chaicor/ETSI_support.asp

Copyright Notification

No part may be reproduced except as authorized by written permission.
The copyright and the foregoing restriction extend to reproduction in all media.

© European Telecommunications Standards Institute 2010.
© European Broadcasting Union 2010.
All rights reserved.

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

3GPPTM is a Trade Mark of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners.

LTETM is a Trade Mark of ETSI currently being 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 | 10 |
| Foreword..... | 10 |
| Introduction | 10 |
| 1 Scope | 12 |
| 2 References | 12 |
| 2.1 Normative references | 12 |
| 2.2 Informative references..... | 12 |
| 3 Definitions, symbols and abbreviations | 13 |
| 3.1 Definitions..... | 13 |
| 3.2 Symbols..... | 14 |
| 3.3 Abbreviations | 15 |
| 4 Overview of DVB-T2..... | 16 |
| 4.1 Commercial requirements overview..... | 16 |
| 4.2 Background design principles | 17 |
| 4.3 Benefits of DVB-T2 compared to DVB-T | 17 |
| 4.4 System overview | 18 |
| 4.4.1 Introduction..... | 18 |
| 4.4.2 Architectural model | 18 |
| 4.4.3 Specification documents describing DVB-T2 | 20 |
| 4.4.4 Physical-layer specification | 20 |
| 4.5 Key technologies | 22 |
| 4.5.1 Physical Layer Pipes (PLPs)..... | 22 |
| 4.5.1.1 Input Mode A | 22 |
| 4.5.1.2 Input Mode B | 22 |
| 4.5.2 Additional bandwidths (1,7 MHz and 10 MHz) | 23 |
| 4.5.3 Extended carrier mode (for 8 K, 16 K, 32 K) | 23 |
| 4.5.4 Alamouti-based MISO (in frequency direction) | 23 |
| 4.5.5 Preambles (P1 and P2)..... | 24 |
| 4.5.6 Pilot patterns | 24 |
| 4.5.7 256-QAM..... | 25 |
| 4.5.8 Rotated constellations | 25 |
| 4.5.9 16 K and 32 K FFT sizes and 1/128 guard-interval fraction..... | 27 |
| 4.5.10 LDPC/BCH error control coding | 27 |
| 4.5.11 Interleaving stages (time, bit, cell, frequency)..... | 28 |
| 4.5.12 Peak-to-Average Power Ratio (PAPR) reduction techniques | 28 |
| 4.5.13 Future Extension Frames (FEFs) | 28 |
| 4.6 Memory and complexity considerations..... | 28 |
| 4.7 Layer 2 signalling..... | 29 |
| 4.7.1 Overview | 29 |
| 4.7.2 T2 delivery-system descriptor..... | 29 |
| 4.8 Network Topology and Modulator Interface (T2-MI)..... | 29 |
| 4.9 Transmitter identification | 30 |
| 5 Choice of parameters..... | 30 |
| 5.1 Choice of FFT size | 30 |
| 5.2 Choice of extended or normal carrier mode | 30 |
| 5.3 Choice of Guard interval | 31 |
| 5.4 Choice of pilot pattern..... | 33 |
| 5.5 Choice of frame length | 34 |
| 5.6 Choice of superframe length | 36 |
| 5.7 Choice of input mode (A or B)..... | 36 |
| 5.8 Choice of number of sub-slices | 37 |
| 5.8.1 Avoiding undesirable interaction of subslicing with carrier mapping and frequency interleaving..... | 37 |
| 5.8.2 Subslicing with multi-frame interleaving | 38 |

| | | |
|-----------|--|----|
| 5.9 | Choice of time interleaving parameters | 38 |
| 5.10 | Choice of rotated constellations | 39 |
| 5.11 | Choice of code rate, block length and constellation | 39 |
| 5.12 | Choice of mode adaptation options | 39 |
| 6 | Anatomy of the DVB-T2 signal | 40 |
| 6.1 | Physical frame structure | 40 |
| 6.1.1 | Superframe | 40 |
| 6.1.2 | T2-frame | 41 |
| 6.1.2.1 | Frame-Closing Symbols | 43 |
| 6.1.3 | L1 signalling | 44 |
| 6.1.4 | Physical-layer pipes | 45 |
| 6.1.5 | Dummy cells | 47 |
| 6.1.6 | Auxiliary streams | 48 |
| 6.1.7 | FEFs | 48 |
| 6.1.7.1 | FEF Basics | 48 |
| 6.1.7.2 | FEF Signalling | 49 |
| 6.1.7.3 | Possible use cases | 50 |
| 6.2 | Logical frame structure | 50 |
| 6.2.1 | Packets | 50 |
| 6.2.2 | BBFRAMES, FECFrames and FEC blocks | 51 |
| 6.2.2.1 | BBFRAMES | 51 |
| 6.2.2.2 | FECFRAMES | 51 |
| 6.2.2.3 | FEC Blocks | 51 |
| 6.2.3 | Interleaving Frames | 51 |
| 6.2.4 | TI blocks | 52 |
| 6.2.5 | Mapping to physical frame structure | 52 |
| 6.2.6 | The concept of a T2-system | 52 |
| 6.3 | Capacity | 53 |
| 6.3.1 | Capacity in cells and FEC blocks per frame | 54 |
| 6.3.2 | Total bit-rates of the T2 system | 55 |
| 6.3.3 | Examples of bit-rates | 57 |
| 6.4 | Overview of interleaving | 58 |
| 7 | Modulator Interface (T2-MI) and network topology | 59 |
| 7.1 | Overview | 59 |
| 7.2 | Statistical multiplexing and service splitting | 59 |
| 7.3 | The T2 gateway and T2-Modulator Interface (T2-MI) | 60 |
| 7.3.1 | The Basic T2-Gateway | 60 |
| 7.3.2 | Practical T2-Gateways | 60 |
| 7.4 | Distribution of T2-MI packets to transmitters | 60 |
| 7.5 | Over-air distribution | 61 |
| 7.6 | Regional/local content insertion | 61 |
| 8 | T2-gateway | 62 |
| 8.1 | Input interface | 62 |
| 8.2 | Generation of one Transport Stream per data PLP | 62 |
| 8.2.1 | Separate TSs with common control | 63 |
| 8.2.2 | From one big TS | 64 |
| 8.2.3 | Discussion | 64 |
| 8.3 | Use of common PLPs | 64 |
| 8.4 | PLP grouping and statistical multiplexing | 67 |
| 8.4.1 | PLP groups with the same modulation and coding | 67 |
| 8.4.2 | T2-aware statistical multiplexing | 67 |
| 8.5 | Mode adaptation | 67 |
| 8.6 | Scheduling | 68 |
| 8.6.1 | Model for allocating number of FEC blocks to each Interleaving Frame | 69 |
| 8.6.1.1 | Mode A | 69 |
| 8.6.1.2 | Mode B | 69 |
| 8.6.1.2.1 | Allocation model 1 | 69 |
| 8.6.1.2.2 | Allocation model 2 | 73 |
| 8.6.1.2.3 | Other allocation strategies | 73 |
| 8.6.1.2.4 | Use of the allocation models with FEFs | 74 |

| | | |
|-----------|---|-----|
| 8.6.1.2.5 | Timing and alignment of collection windows | 74 |
| 8.6.2 | Scheduling of sub-slices | 75 |
| 8.7 | Padding methods: | 75 |
| 8.8 | Managing the de-jitter buffer..... | 76 |
| 8.8.1 | The receiver-buffer model | 77 |
| 8.8.2 | Principles of buffer management | 78 |
| 8.8.3 | Choice of design delay..... | 79 |
| 8.8.4 | The de-jitter buffer budget..... | 80 |
| 8.9 | Compensating Delay | 82 |
| 8.10 | Layer-1 signalling..... | 83 |
| 8.10.1 | Overview | 83 |
| 8.10.2 | Repetition of L1-dynamic signalling | 84 |
| 8.10.3 | L1 change-indication mechanism | 84 |
| 8.10.4 | Use of the L1 extension field | 84 |
| 8.10.5 | Example of L1 signalling..... | 85 |
| 8.10.6 | In-band L1-dynamic signalling..... | 86 |
| 8.10.6.1 | Generation of the in-band type A signalling message..... | 87 |
| 8.10.6.2 | Insertion of L1-dynamic information of other PLPs | 88 |
| 8.10.6.3 | Which frame is signalled for same PLP and other PLPs..... | 88 |
| 8.10.6.4 | Example of in-band signalling | 90 |
| 8.10.6.5 | In-band type B signalling | 90 |
| 8.11 | Layer-2 signalling..... | 91 |
| 8.11.1 | Transport Streams | 91 |
| 8.11.2 | Generic Streams | 93 |
| 8.12 | Generation of the T2-MI stream..... | 93 |
| 8.13 | Generation of the T2-MIP | 93 |
| 8.14 | L1 Bias Balancing | 94 |
| 8.14.1 | L1 Bias balancing Techniques | 95 |
| 8.14.1.1 | P2 Tone Reservation | 95 |
| 8.14.1.2 | L1 reserved bits and extension field..... | 95 |
| 8.14.1.3 | Bias balancing cells..... | 95 |
| 8.14.1.4 | L1-ACE..... | 95 |
| 8.14.2 | Approach to bias balancing..... | 96 |
| 9 | Modulators | 99 |
| 9.1 | Operation with T2-MI input stream | 99 |
| 9.1.1 | Overview | 99 |
| 9.1.2 | Buffering and timing relationships | 100 |
| 9.1.3 | Generation of L1-signalling from T2-MI packets..... | 101 |
| 9.1.4 | Generation of the PLP data from the T2-MI input..... | 103 |
| 9.1.5 | Start-up behaviour | 103 |
| 9.1.6 | Modulator synchronisation | 104 |
| 9.1.6.1 | Single-Frequency Network operation | 104 |
| 9.1.6.2 | Modulator test modes for SFN operation..... | 104 |
| 9.1.6.3 | Multiple-frequency network operation..... | 105 |
| 9.2 | Bit-interleaved coding and modulation | 106 |
| 9.2.1 | FEC for the PLPs | 106 |
| 9.2.1.1 | FEC overview | 106 |
| 9.2.1.2 | BCH code..... | 106 |
| 9.2.1.3 | LDPC code with multi-level constellations..... | 106 |
| 9.2.1.3.1 | LDPC code | 106 |
| 9.2.1.3.2 | Bit Interleaver and demux | 107 |
| 9.2.1.3.3 | Constellation Mapping | 109 |
| 9.2.2 | FEC for the L1 signalling | 109 |
| 9.2.2.1 | Shortening of BCH Information part | 109 |
| 9.2.2.2 | Example for shortening of BCH information..... | 110 |
| 9.2.2.3 | BCH encoding..... | 112 |
| 9.2.2.4 | LDPC encoding..... | 112 |
| 9.2.2.5 | Puncturing of LDPC parity bits..... | 112 |
| 9.2.2.6 | Example for Puncturing of LDPC parity bits | 114 |
| 9.2.2.7 | Removal of zero-padding bits | 116 |
| 9.2.2.8 | Bit Interleaving and constellation mapping after shortening and puncturing..... | 116 |

| | | |
|------------|---|-----|
| 9.2.3 | Rotated constellations | 116 |
| 9.2.4 | Cell Interleaver | 118 |
| 9.2.5 | Time Interleaving..... | 118 |
| 9.3 | Frame-building and OFDM generation | 119 |
| 9.3.1 | Frame building..... | 119 |
| 9.3.2 | Frequency Interleaving | 119 |
| 9.3.3 | MISO processing | 122 |
| 9.3.4 | Pilot insertion..... | 122 |
| 9.3.4.1 | Purpose of pilot insertion | 122 |
| 9.3.4.2 | Pilot locations..... | 123 |
| 9.3.4.3 | Number of pilot cells of each type | 124 |
| 9.3.4.3.1 | Pilots in normal symbols | 124 |
| 9.3.4.3.2 | Pilots in P2 symbols | 125 |
| 9.3.4.3.3 | Pilots in Frame-Closing Symbols | 126 |
| 9.3.4.4 | Pilot boosting | 126 |
| 9.3.4.5 | MISO pilots..... | 126 |
| 9.3.4.5.1 | Inversion of pilots..... | 127 |
| 9.3.4.5.2 | Additional P2 pilots..... | 127 |
| 9.3.4.6 | Use of the reference sequence in normal and extended carrier modes..... | 127 |
| 9.3.5 | Dummy carrier reservation | 128 |
| 9.3.6 | Inverse Fast Fourier Transform (IFFT)..... | 128 |
| 9.3.7 | P1 generation | 129 |
| 9.3.8 | PAPR reduction | 129 |
| 9.3.8.1 | Overview and comparison of the two techniques..... | 129 |
| 9.3.8.2 | ACE..... | 129 |
| 9.3.8.3 | Tone Reservation | 130 |
| 9.3.8.3.1 | Reserved carriers and kernel..... | 131 |
| 9.3.8.3.2 | Peak-cancellation algorithm | 132 |
| 9.3.8.3.3 | Selection of the clipping level | 133 |
| 9.3.8.3.4 | SFN issues | 133 |
| 10 | Receivers | 133 |
| 10.1 | Overview | 133 |
| 10.1.1 | Receiver Model..... | 134 |
| 10.1.2 | Acquisition scenarios..... | 135 |
| 10.1.2.1 | Initial scan | 136 |
| 10.1.2.2 | Service-selection and reproduction | 136 |
| 10.1.2.3 | Hand-over (portable and mobile cases)..... | 137 |
| 10.1.2.3.1 | Automatic service following | 137 |
| 10.1.2.3.2 | Announcements | 139 |
| 10.2 | Acquisition and synchronisation | 139 |
| 10.2.1 | Overview of synchronisation | 139 |
| 10.2.2 | Use of P1 | 139 |
| 10.2.2.1 | Introduction..... | 139 |
| 10.2.2.2 | Structure of P1 | 140 |
| 10.2.2.3 | P1 features..... | 140 |
| 10.2.2.4 | Overview of the use of P1 during channel scanning..... | 141 |
| 10.2.2.5 | P1 detection..... | 143 |
| 10.2.2.5.1 | Summary | 143 |
| 10.2.2.5.2 | Goals..... | 143 |
| 10.2.2.5.3 | Implementation..... | 144 |
| 10.2.2.6 | P1 validation | 145 |
| 10.2.2.6.1 | Summary | 145 |
| 10.2.2.6.2 | Goals..... | 145 |
| 10.2.2.6.3 | Implementation issues (coarse frequency correction)..... | 146 |
| 10.2.2.7 | Decoding the P1 signalling | 146 |
| 10.2.2.7.1 | Summary | 146 |
| 10.2.2.7.2 | Goals..... | 146 |
| 10.2.2.7.3 | Implementation..... | 146 |
| 10.2.2.8 | Interpretation of P1 signalling fields..... | 147 |
| 10.2.2.9 | The FEF loop | 147 |
| 10.2.3 | Determining the Guard Interval | 147 |

| | | |
|------------|--|-----|
| 10.2.3.1 | Determining the Guard Interval using P2..... | 147 |
| 10.2.3.2 | Use of Guard Interval Correlation to determine the Guard Interval..... | 148 |
| 10.2.3.3 | Use of the Guard-Interval hinting information from P1..... | 148 |
| 10.2.4 | Time and Frequency synchronisation..... | 149 |
| 10.2.5 | Use of Frame/Symbol level sequences for synchronisation..... | 149 |
| 10.2.5.1 | Introduction..... | 149 |
| 10.2.5.2 | Definition of the reference sequence..... | 149 |
| 10.2.5.3 | Synchronisation algorithm..... | 151 |
| 10.2.5.4 | Full synchronisation method..... | 151 |
| 10.2.5.5 | Evaluation of Synchronisation features..... | 152 |
| 10.2.6 | Use of P2..... | 152 |
| 10.2.6.1 | Initial Scan..... | 152 |
| 10.2.6.1.1 | Reception of P2 symbols during initial scan..... | 153 |
| 10.2.6.1.2 | TX Basic parameters verification..... | 153 |
| 10.2.6.1.3 | Determining position in the super-frame..... | 154 |
| 10.2.6.2 | Channel-zapping..... | 154 |
| 10.2.6.2.1 | Assumptions..... | 154 |
| 10.2.6.2.2 | Requirements for zapping between two channels..... | 154 |
| 10.2.7 | Receiver Synchronisation/Zapping time..... | 155 |
| 10.2.7.1 | Initial scan time..... | 155 |
| 10.2.7.2 | Zapping time..... | 155 |
| 10.3 | OFDM Demodulation..... | 156 |
| 10.3.1 | FFT..... | 156 |
| 10.3.2 | Channel equalisation..... | 156 |
| 10.3.2.1 | Common Phase-Error correction..... | 156 |
| 10.3.2.2 | Locating the CPs..... | 157 |
| 10.3.2.3 | SISO channel estimation..... | 158 |
| 10.3.2.3.1 | Overview..... | 158 |
| 10.3.2.3.2 | Fundamental limits..... | 160 |
| 10.3.2.3.3 | Interpolation..... | 160 |
| 10.3.2.3.4 | P2 symbols..... | 164 |
| 10.3.2.3.5 | Other special cases..... | 164 |
| 10.3.2.3.6 | Further observations..... | 166 |
| 10.3.2.4 | MISO channel estimation..... | 166 |
| 10.3.3 | MISO decoding..... | 167 |
| 10.3.3.1 | ZF and MMSE Alamouti matrix inversion..... | 168 |
| 10.3.3.2 | Separation of Alamouti pairs by pilot cells..... | 168 |
| 10.4 | De-interleaving and de-mapping..... | 169 |
| 10.4.1 | Frequency/L1 de-interleaver..... | 169 |
| 10.4.1.1 | Implementation of L1/Frequency de-interleaver 32 K..... | 169 |
| 10.4.1.2 | Implementation of the basic frequency de-interleaver in 1 K-16 K..... | 171 |
| 10.4.1.3 | Implementation of the L1/ Frequency De-interleaver in 1 K-16 K..... | 172 |
| 10.4.2 | Frame extraction..... | 173 |
| 10.4.2.1 | Extraction of the PLPs..... | 173 |
| 10.4.2.2 | Extraction of the L1 signalling cells..... | 173 |
| 10.4.2.3 | Timing of L1 and PLP extraction and decoding in the P2 symbols..... | 173 |
| 10.4.2.4 | Use of in-band dynamic signalling..... | 174 |
| 10.4.2.4.1 | Interleaving over one T2 frame or less..... | 174 |
| 10.4.2.4.2 | In-band signalling when interleaving over multiple T2 frames..... | 175 |
| 10.4.2.4.3 | In-band signalling for PLPs not present in every T2-frame..... | 175 |
| 10.4.2.4.4 | Reading the P2 symbols on reconfiguration..... | 175 |
| 10.4.2.4.5 | PLPs with zero FEC blocks in a given Interleaving Frame..... | 175 |
| 10.4.2.5 | Power saving..... | 176 |
| 10.4.2.6 | FEFs..... | 176 |
| 10.4.2.7 | Auxiliary streams..... | 176 |
| 10.4.3 | Time de-interleaver..... | 177 |
| 10.4.3.1 | Phase of Time de-interleaving..... | 177 |
| 10.4.3.2 | Memory-efficient implementation of time de-interleaver..... | 177 |
| 10.4.3.2.1 | De-interleaver for constant bit-rate..... | 177 |
| 10.4.3.2.2 | Implementation details..... | 180 |
| 10.4.3.2.3 | Adaptation for variable bit-rates..... | 180 |
| 10.4.3.3 | Disabled Time Interleaving..... | 181 |

| | | |
|------------|---|-----|
| 10.4.3.4 | Relationship to the receiver buffer model | 181 |
| 10.4.4 | Cell de-interleaver..... | 181 |
| 10.5 | Channel decoding | 182 |
| 10.5.1 | LLR demapping for non-rotated constellations | 182 |
| 10.5.2 | Removal of cyclic Q-delay | 182 |
| 10.5.3 | De-mapping for rotated constellations..... | 182 |
| 10.5.3.1 | 2D LLR demapping..... | 182 |
| 10.5.3.1.1 | Computation of perfect LLR | 182 |
| 10.5.3.1.2 | Computation of simplified LLR (suitable for hardware implementation) | 183 |
| 10.5.3.2 | 2D LLR demapper with iterative demapping and decoding..... | 184 |
| 10.5.3.2.1 | Computation of perfect LLR with ID | 184 |
| 10.5.3.2.2 | 'Genie-aided' demapping | 185 |
| 10.5.3.2.3 | Computation of simplified LLR with ID (suitable for hardware implementation)..... | 186 |
| 10.5.3.3 | Impact on complexity..... | 187 |
| 10.5.4 | PAPR-aware demapping | 188 |
| 10.5.5 | LDPC decoding | 188 |
| 10.5.6 | BCH decoding | 189 |
| 10.5.7 | Decoding of the L1 signalling..... | 189 |
| 10.5.7.1 | FEC decoding of the L1 signalling | 189 |
| 10.5.7.2 | Extraction of the L1-signalling fields..... | 190 |
| 10.5.7.3 | Storage required in L1 decoding | 191 |
| 10.6 | Output processing..... | 191 |
| 10.6.1 | De-/Re-multiplexing | 191 |
| 10.6.1.1 | Construction of output TS..... | 191 |
| 10.6.1.2 | Mode adaptation..... | 191 |
| 10.6.1.3 | Determination of output-TS bit-rate..... | 192 |
| 10.6.1.4 | De-jitter buffer | 193 |
| 10.6.1.5 | Re-insertion of deleted null packets | 193 |
| 10.6.1.6 | Re-combining the Common and Data PLPs..... | 194 |
| 10.6.2 | Output interface | 194 |
| 11 | Transmitters..... | 194 |
| 12 | Network planning | 194 |
| 13 | Tuners..... | 194 |
| 13.1 | Phase noise requirements | 194 |
| 13.2 | AGC, headroom and FEFs | 195 |
| 14 | Performance | 195 |
| 14.1 | Channel models for simulations | 196 |
| 14.1.1 | Introduction..... | 196 |
| 14.1.2 | Gaussian Channel | 196 |
| 14.1.3 | Ricean Channel..... | 196 |
| 14.1.4 | Rayleigh Channel..... | 197 |
| 14.1.5 | Mobile Channel, TU-6..... | 198 |
| 14.1.6 | Simple two path profile, 0 dB Echo..... | 198 |
| 14.1.7 | MISO channel..... | 198 |
| 14.1.7.1 | Ricean MISO Channel | 199 |
| 14.1.7.2 | Rayleigh MISO Channel | 200 |
| 14.1.8 | Memoryless Rayleigh Channel with erasures | 200 |
| 14.2 | Simulated system performance for 8 MHz channel | 202 |
| 14.2.1 | Correction values for pilot boosting | 204 |
| 14.3 | Simulation results for Memoryless Rayleigh Channel with erasures | 205 |
| 14.4 | Penalty for real channel estimation | 206 |
| 14.5 | P1 performance | 208 |
| 14.5.1 | Figure description | 209 |
| 14.5.2 | P1 performance results..... | 209 |
| 14.5.2.1 | Gaussian channel..... | 209 |
| 14.5.2.2 | Rayleigh channel (P1)..... | 210 |
| 14.5.2.3 | TU-6@200Hz..... | 210 |
| 14.5.2.4 | 0 dB 180° Echo, Echo with opposite sign | 210 |
| 14.6 | L1-pre- and post-signalling performance | 211 |

| | | |
|-------------------------------|---|------------|
| 15 | Examples of possible uses of the system..... | 214 |
| Annex A (informative): | Derivation of allowable and recommended sub-slicing values | 215 |
| History | | 217 |

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 (<http://webapp.etsi.org/IPR/home.asp>).

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 standardisation, interoperability and future proof specifications.

Introduction

DVB-T2 is a standard for digital terrestrial television broadcasting, offering significant benefits compared to DVB-T (EN 302 744 [i.18]).

The present document is intended to serve a number of purposes.

DVB-T2 includes many new techniques not previously used in the DVB family of standards. Some, such as the P1 preamble, are completely novel, having been invented specifically for DVB-T2 and therefore are not yet discussed in the wider literature. The present document gives further information about these techniques.

The physical layer specification ([i.1]) was written to be precise, rather than descriptive, and the present document aims to give a more extended explanation of the various elements together with some of the reasons behind the design of the DVB-T2 system. It attempts to capture as much as possible of the common understanding arrived at by the working group developing EN 302 755 [i.1]. It also gives additional information intended to make implementation easier, to act as a cross-check, and to help implementers to avoid some of the more common pitfalls.

As is conventional, [i.1] only describes the generation of the signal on the air, and the present document explains how the features of this signal are intended to be used in the receiver.

On the signal generation side, [i.1] is not always prescriptive: operations such as scheduling may be performed in any manner provided certain constraints or conditions are met. The present document suggests some methods by which these operations can be carried out.

The set of processing blocks described in [i.1] do not correspond directly to a piece of equipment, instead including elements of both a T2-gateway and a modulator. This expected topology is not obvious from [i.1] but is described in the present document.

Finally, EN 302 755 [i.1] allows a large number of options and combinations. Flexibility has been deliberately retained in [i.1] to allow optimisation as more experience and expertise is gained. However, initial implementations are expected to use a small subset of the possible combinations and the present document gives some guidance on the choice of parameters.

Readers who are new to DVB-T2 should begin with clause 4, which gives an introduction to DVB-T2, including a summary of the benefits compared to DVB-T [i.18] and an overview of the key new technologies included.

Broadcasters and other users of the DVB-T2 system requiring some guidance on the choice of parameters should see clause 5. The simulated performance results given in clause 14 will also be of interest to such readers, as will clause 12 which will describe aspects of network planning.

Readers requiring a deeper understanding of the structure of the DVB-T2 signal should see clause 6, which explains the key concepts and framing structures used in EN 302 755 [i.1]. This clause should also be studied before the clauses relating to detailed implementation, since those later clauses assume a working knowledge of the concepts described in clause 6. Many of the concepts, particularly the Interleaving Frame, are crucial to an understanding of the rest of the system.

Clause 7 describes some typical network topologies and also introduces the DVB-T2 modulator interface (T2-MI), which is designed to form the backbone of such networks; this is addressed both to network operators and also to implementers of T2-gateways and modulators.

The remaining clauses are each aimed at implementers of a particular piece of equipment: the T2-gateway (clause 8), modulator (clause 9), receiver (clause 10), transmitter (clause 11), or tuner (clause 13). Whilst implementers should pay most attention to the relevant clause, receiver implementers in particular should consult clauses 8 and 9 in addition to clause 10 since some aspects of the signal that are addressed in connection with the generation of the signal are equally applicable to the corresponding demodulation process but are not repeated in the receiver clause.

1 Scope

The present document gives guidelines for the implementation of all aspects of the DVB-T2 end-to-end chain. This includes:

- the parts of the system defined by the physical layer system specification [i.1];
- aspects of the input pre-processing, which is outside the scope of [i.1];
- the modulator interface (or T2-MI) specification;
- future developments including a possible Transmitter Identification (TxID) standard.

The scope includes guidance relevant to implementers of T2-gateways, modulators, transmitters, receivers and tuners; network planners or operators; and broadcasters.

2 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 reference 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.

2.1 Normative references

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

Not applicable.

2.2 Informative references

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 EN 302 755: "Digital Video Broadcasting (DVB); Frame structure channel coding and modulation for a second generation digital terrestrial television broadcasting system (DVB-T2)".
- [i.2] Marcel J. E. Golay "Complementary Series". IRE Transactions on Information Theory, April 1961, pp. 82-87.
- [i.3] C-C Tseng, C.L. Liu: "Complementary Sets of Sequences". IEEE Transactions on Information Theory, Vol. IT-18, N. 5, p 644-652. September 1972.
- [i.4] Popovic, B.M: "Efficient Golay Correlator", Electronic Letters 19th Aug 1999 Vol. 35 No.17.
- [i.5] X. Wang, Y. Wu, JY. Chouinard, S. Lu and B. Caron, "A Channel Characterization Technique Using Frequency Domain Pilot Time Domain Correlation Method for DVB-T Systems", IEEE Trans. On Consumer Electronics, Vol. 49, n 4, pp 949-957, November 2003.
- [i.6] A. Filippi, S. Serbetli, "OFDM Symbol Synchronization Using Frequency Domain Pilots in Time Domain", Proceedings of Vehicular Technology Conference (VTC) 2007 Fall, Baltimore, USA, October 2007.

- [i.7] B. Jahan, M. Lanoiselée, G. Degoulet and R. Rabineau, "Frame synchronization method for OFDM/QAM and OFDM/OQAM modulation", IEEE International Conference on Circuits and Systems for Communications 2008, Shanghai, China, May 2008.
- [i.8] Alamouti, S.: "A simple transmit diversity technique for wireless communications", IEEE Journal on Selected areas in Communications, Vol. 16, No. 8, pp. 1451-1458, October 1998.
- [i.9] ETSI TS 102 606: "Digital Video Broadcasting (DVB); Generic Stream Encapsulation (GSE) Protocol".
- [i.10] Stott, J. H., Summer 1998. The effects of phase noise in COFDM. EBU Technical Review, (276), pp. 12-25.
- NOTE: Available from BBC website: <http://www.bbc.co.uk/rd/pubs/papers/pdf/files/jsebu276.pdf>.
- [i.11] E.R.Berlekamp: "Algebraic Coding Theory", New York:McGraw-Hill, 1968.
- [i.12] R.G.Gallager: "Low density parity check codes", IRE Trans. Info. Theory IT-8: 21-28, 1962.
- [i.13] International Telecommunication Union: "Final Acts of the Regional Radiocommunication Conference for planning of the digital terrestrial broadcasting service in parts of Regions 1 and 3, in the frequency bands 174-230 MHz and 470-862 MHz (RRC-06)".
- [i.14] ETSI EN 300 468: "Digital Video Broadcasting (DVB); Specification for Service Information (SI) in DVB systems".
- [i.15] V. Mignone, A. Morello: "CD3 OFDM: a novel demodulation scheme for fixed and mobile receivers", IEEE Transaction on Communications, vol. 44, n. 9, September 1996.
- [i.16] ISO/IEC 13818-1: "Information technology - Generic coding of moving pictures and associated audio information: Systems".
- [i.17] ETSI EN 302 307: "Digital Video Broadcasting (DVB); Second generation framing structure, channel coding and modulation systems for Broadcasting, Interactive Services, News Gathering and other broadband satellite applications (DVB-S2)".
- [i.18] ETSI EN 300 744: "Digital Video Broadcasting (DVB); Framing structure, channel coding and modulation for digital terrestrial television".
- [i.19] ETSI TS 102 034: "Digital Video Broadcasting (DVB); Transport of MPEG-2 TS Based DVB Services over IP Based Networks".
- [i.20] ETSI TS 102 773: "Digital Video Broadcasting (DVB); Modulator Interface (T2-MI) for a second generation digital terrestrial television broadcasting system (DVB-T2)".
- [i.21] V. Mignone, A. Morello, M. Visintin: "CD3-OFDM: A new channel estimation method to improve the spectrum efficiency in digital terrestrial television systems", IBC'95 Conference, Amsterdam, 14-18 September, 1995.
- [i.22] ETSI TS 102 992: "Digital Video Broadcasting (DVB); Structure and modulation of optional transmitter signatures (T2-TX-SIG) for use with the DVB-T2 second generation digital terrestrial television broadcasting system".

3 Definitions, symbols and abbreviations

3.1 Definitions

For the purposes of the present document, the terms and definitions given in EN 302 755 [i.1] and the following apply:

allocation window: period during which the last bit of a BBFRAME can arrive and that BBFRAME be allocated to a given interleaving frame

basic T2-Gateway: device taking one input stream per data PLP and producing at its output a T2-MI stream

- BBFRAME:** set of K_{bch} bits which form the input to one FEC encoding process (BCH and LDPC encoding)
- channel extent:** duration of the channel impulse response from the first significant component to the last
- collection window:** period of time during which arriving input bits could possibly end up being carried in a given Interleaving Frame
- frame-skipping:** mechanism by which a PLP may be mapped to only a subset of frames by setting $I_{\text{jump}} > 1$
- guard-interval fraction:** ratio T_g/T_u of the guard interval duration to the active symbol period
- multi-frame interleaving:** interleaving a PLP over multiple frames by setting $P_1 > 1$
- relay (transmitter):** transmitter in a network that re-transmits a signal received off-air, either by simple frequency transposition or by regenerating the signal
- Scattered Pilot-bearing carrier:** carrier (i.e. a value of k) which is a Scattered Pilot in some symbols (i.e. for which $k \bmod D_x = 0$)
- single-profile receiver:** DVB-T2 receiver according to the single profile defined by EN 302 755 [i.1]
- slice:** set of cells for a given PLP mapped to a given T2-frame
- sub-slice:** set of cells for a given PLP mapped to a contiguous series of addresses in a T2-frame
- T2-Gateway:** device producing at T2-MI at its output, incorporating the functionality of the Basic T2-Gateway and, optionally, additional processing such as re-multiplexing
- thinning:** use of unmodulated cells in the frame closing symbol

3.2 Symbols

For the purposes of the present document, the symbols given in EN 302 755 [i.1] and the following apply:

| | |
|-----------------------------|---|
| C_{bias} | Total complex bias in the modulated cells of the L1 signalling (i.e. the complex sum of the cells) |
| $C_{\text{bias_residual}}$ | Total expected residual bias in the P2 symbols after adding any bias balancing cells |
| C_{TR} | The expected amplitude of the P2 TR cells |
| $N_{\text{L1_ACE}}$ | The number of L1 cells available for correction by the L1-ACE algorithm |
| $N_{\text{L1_ACE_IMAG}}$ | The number of L1 cells available for correction by the imaginary part of the L1-ACE algorithm |
| $N_{\text{L1_ACE_REAL}}$ | The number of L1 cells available for correction by the real part of the L1-ACE algorithm |
| N_{plp} | Number of PLPs in a T2-system |
| p | Proportion of '1's in the coded bits of the L1-post signalling |
| p_{ACE} | The probability that a constellation state is available for correction by the L1-ACE algorithm |
| P_{inc} | The expected power increase of the P2 symbols due the effects of the L1-ACE and TR cells |
| R_{total} | Total bitrate of a group of PLPs |
| T_{IF} | Duration of an Interleaving Frame: $T_{\text{IF}} = I_{\text{jump}} \times P_1 \times T_{\text{F}}$ |
| $t_{\text{post_coded}}$ | Total number of '1's in the coded fields of the L1-post signalling |
| $t_{\text{post_uncoded}}$ | Total number of '1's in the uncoded fields of the L1-post signalling |
| $t_{\text{pre_coded}}$ | Total number of '1's in the coded fields of the L1-pre signalling |
| $t_{\text{pre_uncoded}}$ | Total number of '1's in the uncoded fields of the L1-pre signalling |
| V_{peak} | The expected time domain peak voltage in each P2 symbol due to the bias |
| V_{corr} | The target correction voltage to be produced by the L1-ACE and TR algorithms |

3.3 Abbreviations

For the purposes of the present document, the abbreviations given in EN 302 755 [i.1] and the following apply:

| | |
|-----------|---|
| 1PPS | One-pulse-per-second (signal from GPS receiver or other timing reference) |
| ACE | Active Constellation Extension |
| AGC | Automatic Gain Control |
| ASI | Asynchronous Serial Interface |
| AWGN | Additive White Gaussian Noise |
| BCH | Bose-Chaudhuri-Hocquenghem multiple error correction binary block code |
| BER | Bit Error Ratio |
| BPSK | Binary Phase Shift Keying |
| BUFS | ISSY variable indicating the maximum size of the requested receiver buffer to compensate delay variations |
| CA | Conditional Access |
| CDS | Carrier-Distribution Sequence |
| COFDM | Coded Orthogonal Frequency Division Multiplexing |
| CP | Continual Pilot |
| CPE | Common Phase Error |
| CW | Continuous Wave |
| DFT | Discrete Fourier Transform |
| DJB | De-Jitter Buffer |
| EIT | Event Information Table |
| FEF | Future-Extension Frame |
| FFT | Fast Fourier Transform |
| FIFO | First-In First-Out buffer |
| FPGA | Field Programmable Gate Array |
| gcd(a, b) | greatest common divisor of a and b |

NOTE: Also known as highest common factor.

| | |
|-----------|---------------------------------------|
| GCS | Generic Continuous Stream |
| GFPS | Generic Fixed Packet size Stream |
| GIF | Guard-Interval Fraction (T_G/T_U) |
| GPS | Global Positioning System |
| GSE | Generic Stream Encapsulated |
| HEM | High Efficiency Mode |
| I/L Frame | InterLeaving Frame |
| IC | Integrated Circuit |
| ICI | Inter-Carrier Interference |
| ID | Iterative Demapping |
| IFFT | Inverse Fast Fourier Transform |
| ISI | Intersymbol Interference |
| ISSY | Input Stream SYNchronizer |
| LDPC | Low Density Parity Check (codes) |
| LLR | Log Likelihood-Ratio |
| MER | Modulation Error Ratio |
| MFN | Multiple Frequency Network |
| MIMO | Multiple Input Multiple Output |
| MIP | Megaframe Initialisation Packet |
| MODCOD | MODulation and CODing |

NOTE: This term is used to refer to a particular combination of constellation, LDPC code rate and block length.

| | |
|------|--|
| NIT | Network Information Table |
| NM | Normal Mode |
| NPD | Null-Packet Deletion |
| OFDM | Orthogonal Frequency Division Multiplexing |
| PAPR | Peak-to-Average Power Ratio |
| PAT | Program Association Table |
| PCR | Programme Clock Reference |
| PCT | Parity and Column Twist |

| | |
|----------|--|
| PLP | Physical Layer Pipe |
| PN | Pseudo Noise |
| PSI/SI | Program Specific Information / Service Information |
| QEF | Quasi-Error-Free |
| QPSK | Quaternary Phase Shift Keying |
| RF | Radio Frequency |
| RMS | root mean square |
| RS | Reed-Solomon (codes) |
| SDT | Service Description Table |
| SFN | Single-Frequency Network |
| SI | Service Information |
| SNR | Signal-to-Noise Ratio |
| SP | Scattered Pilot |
| Statmux | Statistical multiplex |
| SYNCD | the distance in bits from the beginning of the DATA FIELD of a BBFRAME to beginning of the first transmitter User Packet that starts in the DATA FIELD |
| T2dsd | DVB-T2 delivery system descriptor |
| T2-MI | DVB-T2 Modulator Interface |
| TDI | Time De-Interleaver |
| TFS | Time Frequency Slicing |
| TI-block | Time-Interleaving block |
| TR | Tone Reservation |
| TS | Transport Stream |
| UHF | Ultra High Frequency (band) |
| VBR | Variable Bit Rate |
| VHF | Very High Frequency (band) |

4 Overview of DVB-T2

4.1 Commercial requirements overview

The DVB organisation defined a set of commercial requirements which acted as a framework for the T2 developments. These commercial requirements included:

- T2 transmissions must be able use existing domestic receive antenna installations and must be able to re-use existing transmitter infrastructures. (This requirement ruled out the consideration of MIMO techniques which would involve both new receive and transmit antennas.)
- T2 should primarily target services to fixed and portable receivers.
- T2 should provide a minimum of 30 % capacity increase over DVB-T working within the same planning constraints and conditions as DVB-T.
- T2 should provide for improved single-frequency-network (SFN) performance compared with DVB-T.
- T2 should have a mechanism for providing service-specific robustness; i.e. it should be possible to give different levels of robustness to some services compared to others. For example, within a single 8 MHz channel, it should be possible to target some services for roof-top reception and target other services for reception on portables.
- T2 should provide for bandwidth and frequency flexibility.
- There should be a mechanism defined, if possible, to reduce the peak-to-average-power ratio of the transmitted signal in order to reduce transmission costs.

4.2 Background design principles

A few general principles were adopted in the design of T2. These were:

- a) The DVB organisation should aim to provide a coherent family of standards where possible.
- b) Translation between DVB standards (for example, between DVB-S2 and DVB-T2) should be as easy as possible.
- c) T2 should not re-invent solutions if they already exist within other DVB standards.

Consequently, T2 adopted two key technologies from DVB-S2. These were:

- 1) The system-layer architecture of DVB-S2; in particular, the packaging of data into "Baseband Frames" (see clause 4.4.4).
- 2) The use of the same Low Density Parity Check (LDPC) error-correcting codes as used in S2.

Most decisions in the design of T2 were directed by the requirement to maximise the data carrying capacity. Many options have been included in T2 in order that the overheads of the modulation scheme can be reduced to a minimum based on the requirements imposed by a particular transmission channel. For instance, several options have been included for FFT sizes, guard-interval fractions, and pilot patterns as described in clause 4.5.

4.3 Benefits of DVB-T2 compared to DVB-T

As a result of the technologies introduced in DVB-T2, the potential gain in capacity that could be achieved in the UK is nearly 50 % compared to the current UK mode of DVB-T (see table 1). In addition to the increased capacity, the proposed DVB-T2 mode is expected to offer greater tolerance of multipath and impulsive interference than the current DVB-T mode.

Table 1: Potential capacity increase of almost 50 % compared with current highest capacity DVB-T mode used in the UK

| | Current UK mode | T2 |
|--|------------------------|---------------|
| Modulation | 64-QAM | 256-QAM |
| FFT size | 2 K | 32 K |
| Guard Interval | 1/32 | 1/128 |
| FEC | 2/3CC + RS | 3/5LDPC + BCH |
| Scattered Pilots | 8,3 % | 1,0 % |
| Continual Pilots (see note 1) | 2,0 % | 0,53 % |
| L1 overhead (see note 2) | 1,0 % | 0,53 % |
| Carrier mode | Standard | Extended |
| Capacity | 24,1 Mbit/s | 36,1 Mbit/s |
| NOTE 1: Includes only Continual Pilot cells which are not also Scattered Pilots. | | |
| NOTE 2: TPS for DVB-T; L1-signalling, P1 and extra P2 overhead for DVB-T2. | | |

Even greater increases in capacity could be achieved in modes designed for single-frequency network (SFN) operation, because of the large fractional guard intervals used in these modes. Table 2 shows the comparison between DVB-T2 and DVB-T for a long guard interval (SFN) mode, with the same absolute guard interval in both cases. This provides a 67 % increase in capacity for DVB-T2 over DVB-T. A longer guard interval mode is also available (nearly 20 % increase), which would give improved SFN coverage for only a small loss of capacity (around 3 %).

Table 2: Potential capacity increase of 67 % for an SFN mode

| | DVB-T mode | T2 |
|-------------------------------|-------------|---------------|
| Modulation | 64-QAM | 256-QAM |
| FFT size | 8 K | 32 K |
| Guard Interval | 1/4 | 1/16 |
| FEC | 2/3CC + RS | 3/5LDPC + BCH |
| Scattered Pilots | 8,3 % | 4,2 % |
| Continual Pilots (see note 1) | 2,0 % | 0,39 % |
| L1 overhead (see note 2) | 1,0 % | 0,65 % |
| Carrier mode | Standard | Extended |
| Capacity | 19,9 Mbit/s | 33,2 Mbit/s |

NOTE 1: Includes only Continual Pilot cells which are not also Scattered Pilots.
NOTE 2: TPS for DVB-T; L1-signalling, P1, and extra overhead in P2 and frame closing symbol for DVB-T2.

4.4 System overview


4.4.1 Introduction

This clause gives a brief overview of the second-generation digital terrestrial broadcasting system (DVB-T2), designed to meet all of the commercial requirements given in clause 4.1. The complete DVB-T2 system has at its heart a draft physical-layer specification [i.1]. However, this physical-layer specification is only a part of the overall system, and is accompanied by minor modifications and additions to DVB specifications for PSI/SI, as well as some additional components needed for the complete implementation of DVB-T2 system infrastructure (see clause 4.4.3).

Outline details will be given of the key elements of the system. More detail of the specification and suggestions for its implementation will be provided in later clauses.

4.4.2 Architectural model

The top-level block diagram of a reference DVB-T2 end-to-end chain is shown in figure 1 for the Transport Stream case.

**Figure 1: Block diagram of DVB-T2 chain**

The full DVB-T2 system can be divided into three basic sub-systems on the network side (SS1, SS2 and SS3) and two sub-systems on the receiver side (SS4 and SS5). Regarding interfaces, there are two corresponding interfaces on the network side (A and B), and one receiver-internal interface (D). The RF interface (C) is common to network and receiver.

On the network side the three sub-systems are:

- **SS1: Coding and multiplexing** sub-system. This includes generation of MPEG-2 Transport Streams and/or Generic Streams, e.g. GSE. For video services this includes video/audio encoding plus associated PSI/SI, or other Layer 2 signalling. Typically the video coding (and possibly audio coding) is performed with variable bit rate with a common control ensuring a total constant bit rate (excluding NULL packets) for all streams taken together. This sub-system is largely the same as for other DVB standards, but there are some T2-specific aspects of the coding and multiplexing, which are also discussed in clause 8. The coding and multiplexing sub-system interfaces to the T2-Gateway via the A interface (typically one or more MPEG-2 TSs over ASI). When DVB-T2 uses the common PLP this sub-system is responsible for the arrangement of the output TSs in line with the requirements set out in annex D of [i.1]. When a conventional statmux system, originally generating a single TS, is used in connection with multiple PLPs carrying one TS each, this sub-system will also include some remultiplexing functionality such as PSI/SI handling and PCR restamping.

NOTE 1: If statistical multiplexing is required between PLPs of different modulation or coding, it would be necessary to replace the constant bit rate approach for statistical video (and possibly audio) multiplexing with the idea of constant data cell rate - see clause 8.4.2.

- **SS2: Basic T2-Gateway** sub-system (clause 8). The input interface to this sub-system is exactly the same as that specified in [i.1], applicable both to the basic DVB-T2 physical layer and to the extension described in annex D of [i.1]. This includes functionality for Mode adaptation and Stream adaptation for DVB-T2, together with scheduling and capacity allocation.
 - The Basic T2-Gateway delivers at its output interface (B) a "T2-MI" stream: a sequence of T2-MI packets, each containing either a Baseband frame, IQ vector data for any auxiliary streams, or signalling information (L1 or SFN). The T2-MI stream contains all the information required to describe both the content and emission timing of T2-frames, and a single T2-MI stream is fed to one or more modulators in a network. The T2-MI interface format is defined in [i.20].
 - The operations performed by the Basic T2-Gateway include all those parts of the physical-layer specification [i.1] that are not completely prescriptive, such as scheduling and allocation. These need to be done centrally in an SFN, to ensure that the same signal is generated by all modulators.
- **SS3: DVB-T2 Modulator** sub-system (clause 9). The DVB-T2 modulators use the Baseband frames and T2-frame assembly instructions carried in the incoming T2-MI stream to create DVB-T2 frames and emit them at the appropriate time for correct SFN synchronisation. The modulators interface to the receivers via the C interface (the transmitted DVB-T2 signal).

NOTE 2: In a simple network the coding and multiplexing sub-system may be connected directly to the modulators via a Transport Stream interface. In this case the modulator would include all the functionality as described in the DVB-T2 physical layer specification [i.1], including some of the operations formally performed by the Basic T2-gateway. Since these operations are not completely prescriptive, this arrangement could not be used in a Single Frequency Network (SFN) except in the case described in note 3.

NOTE 3: In an alternative arrangement (see clause 7.5), SFN stations are fed over air from a Master Station on a different frequency. In this case the modulator at the Master Station may assemble the T2 signal for broadcast either from a T2-MI stream or from simple TS inputs.

In the receiver the two sub-systems are:

- **SS4: DVB-T2 demodulator** sub-system (clause 10). This sub-system receives an RF signal from one or (in an SFN) more transmitters in the network and (in the transport-stream case) outputs one transport stream. SS4 interfaces to SS5 via the D interface, a syntactically correct transport stream carrying one or more of the services as well as any common signalling data derived from the common PLP. The streams passing the B interface are identical to those passing the D interface.
- **SS5: Stream decoder** sub-system. This sub-system receives the transport stream and outputs decoded video and audio. Since interface D is a syntactically correct transport stream, this sub-system is essentially the same as for other DVB standards, except that some new L2-signalling elements have been defined for DVB-T2 (see clause 4.7).

NOTE 4: In the case of generic streams, interface D may take other forms and the common signalling may be carried separately from the service stream.

Figure 2 shows the corresponding simplified protocol stacks. The upper horizontal red line shows that the (one or more) TSs that are generated by SS1 pass through SS2, SS3 and SS4 unaffected, i.e. exactly the same TSs are available at Interface A and D. The T2-MI layer includes all protocol layers between the MPEG-2 TS and the physical layer.

The lower horizontal black line identifies the physical layer part of the T2 signal on air (interface C) and the corresponding physical layer for interfaces A and B.




Figure 2: Reference protocol stacks for DVB-T2 (MPEG-2 TS case)

4.4.3 Specification documents describing DVB-T2

The DVB-T2 system is described by the following standards:

- EN 302 755 [i.1], Frame structure channel coding and modulation for a second generation digital terrestrial television broadcasting system (DVB-T2). EN 302 755 defines the physical layer of the DVB-T2 signal on air.
- EN 300 468 [i.14], Service Information (SI) in DVB systems, amended by the new T2 delivery system descriptor (T2dsd).
- TS 102 773 [i.20], the DVB-T2 Modulator Interface (T2-MI), which defines the interface between the Basic T2-gateway and the DVB-T2 modulator to facilitate the distribution to transmitters and the implementation of single frequency networks using the DVB-T2 system.

4.4.4 Physical-layer specification

The generic model for the DVB-T2 physical layer (T2) is represented in figure 3. The system input may be one or more MPEG Transport Stream(s) and/or one or more Generic Stream(s). These may have been modified, where appropriate, by the pre-processing in the T2 Gateway, such that by this point in the chain, the stream inputs have a one-to-one correspondence with data channels in the modulator, which are called Physical-Layer Pipes (PLPs).

The physical-layer specification [i.1] corresponds to the combination of the Basic T2-Gateway (SS2) and the modulator (SS3) of the reference architecture of figure 1. The division of the signal chain into these two sub-systems and the T2-MI interface between them are not necessary for the definition of the signal on the air and is not mentioned in [i.1]. Furthermore, in [i.1], the term "modulator" is used to refer to the combination of these two sub-systems and in many cases the relevant operation is in fact part of the Basic T2-Gateway.

On the reception side, [i.1] corresponds more or less exactly to the demodulator SS4.




Figure 3: High level T2 block diagram

The output from the T2 physical layer is an RF signal on a single RF channel. Optionally, the output may be split into a second output signal, to be conveyed to a second antenna, typically on another transmitter site, for MISO transmission mode, using a modified form of Alamouti encoding. If there is only one PLP there will be just one continuous data channel over the air. However, when there is more than one PLP the data channels may be flexibly time-sliced at the physical layer, providing a range of options that can allow parameters for time diversity and receiver power-saving to be traded against each other.

The multiple PLP and time-slicing approaches implemented by T2 allow for different levels of coding, modulation and time interleaving depth to be applied to different PLPs, to provide variable robustness on a service-by-service basis. The receiver can also concentrate its decoding resources on the one PLP containing the required data. In particular the fixed amount of memory in the receiver dedicated to time de-interleaving (for robustness against impulsive interference) can be used for a greater interleaving depth, compared with the single PLP mode, since the de-interleaver is only processing data for the required PLP. With a single PLP, the time-interleaving depth is around 70 ms, whereas with multiple PLPs this can be extended to the full frame duration (150 ms to 250 ms), or for lower data-rate services it can be extended over multiple frames.

The input-data processing and FEC have been chosen to be compatible with the equivalent mechanisms used in DVB-S2, although, where appropriate, these have been extended to implement extra features for efficiency. So the same baseband-frame and baseband-header structure from S2, together with its null-packet deletion and stream-synchronisation mechanisms, have been exactly duplicated, as have the LDPC/BCH FEC. However, to cater for the different nature of the terrestrial modulation, which is based on the guard-interval COFDM technique used by DVB-T, new bit-interleaving and constellation-mapping techniques are included in T2.

The range of standard COFDM parameters has been extended compared with DVB-T, as requested by the commercial requirements, to include:

- FFT sizes: 1 K, 2 K, 4 K, 8 K, 16 K, 32 K.
- Guard-interval fractions: 1/128, 1/32, 1/16, 19/256, 1/8, 19/128, 1/4.
- Scattered-pilot patterns: 8 different versions matched to guard intervals, providing a range of efficiencies.
- Continual pilots, similar to DVB-T but with improved optimisation to reduce overhead.
- An extended-carrier mode to allow optimum use to be made of the channel bandwidth together with the higher FFT sizes. When this option is used (supported for 8 K, 16 K and 32 K FFT) the carrier spacing is the same as when the normal carrier is used, but additional carriers are added at both ends of the spectrum.
- Extended interleaving including bit, cell, time and frequency interleavers.

The extended range of COFDM parameters allows very significant reductions in overhead to be achieved by DVB-T2 compared with DVB-T, which taken together with the improved error-correction coding allows an increase in capacity of up to nearly 50 % to be achieved for MFN operation and even higher for SFN operation.

The T2 system provides a number of new features for improved versatility:

- a frame structure which contains a special (short) identification symbol, which can be used for rapid channel scanning and signal acquisition, and which also signals some basic frame-structure parameters;

- rotated constellations, which provide a form of modulation diversity, to assist in the reception of higher-code-rate signals in demanding transmission channels;
- special techniques to reduce the peak-to-average ratio of the transmitted signal;
- an option for extending the transmitted signal by including provision for Future-Extension Frames (FEFs), which are unspecified portions of the signal that first-generation receivers will know to ignore, but which could provide a compatible route for later upgrades.

4.5 Key technologies

This clause presents an introduction to the technologies used in DVB-T2, particularly those which were not used in DVB-T.

4.5.1 Physical Layer Pipes (PLPs)

The commercial requirement for service-specific robustness together with the need for different stream types is met by the concept of fully transparent physical-layer pipes (figure 4), which enable the transport of data independently of its structure, with freely selectable, PLP-specific physical parameters. Both the allocated capacity and the robustness can be adjusted to the content/service providers' particular needs, depending on the type of receiver and the usage environment to be addressed.

The DVB-T2 specification [i.1] allows a constellation, code rate and time-interleaving depth to be assigned to each single PLP individually. Furthermore the formatting of the content follows the same baseband-frame structure as that employed by DVB-S2.




Figure 4: Different PLP's occupying different time slices

Typically a group of services will share common elements such as PSI/SI tables or CA information. To avoid the need to duplicate this information for each PLP, DVB-T2 includes the concept of common PLPs, shared by a group of PLPs. Hence, receivers need to decode up to two PLPs at the same time when receiving a single service: the data PLP and its associated common PLP.

Two general input modes are defined: input mode A uses a single PLP, whilst input mode B uses multiple PLPs.

4.5.1.1 Input Mode A

This most simple mode can be viewed as a straightforward extension of DVB-T employing the other advanced features outlined in this clause, but not being sub-divided into multiple PLPs. Here only a single PLP is used, transporting a single transport stream. Consequently the same robustness is applicable to all content, as in DVB-T.

4.5.1.2 Input Mode B

This more advanced mode of operation applies the concept of multiple physical-layer pipes (figure 5). Beside the advantage of service-specific robustness, this mode offers potentially longer time-interleaving depth as well as the option of power saving in the receiver. Therefore even in the case of identical physical-parameter settings, it might be useful to apply this mode, especially if portable and/or mobile devices are to be targeted.



Figure 5: T2-frame for single RF channel, multiple PLP mode (here: 5 PLPs)

The choice of input mode A or B is discussed in clause 5.7.

4.5.2 Additional bandwidths (1,7 MHz and 10 MHz)

In order to make DVB-T2 also suitable for professional use, e.g. transmissions between radio cameras and mobile studios, a 10 MHz option is included; consumer receivers are not expected to support the 10 MHz mode. To allow DVB-T2 to be used in narrower RF channel assignments in e.g. band III and in the L-band, the bandwidth 1,712 MHz is also included. The 1,712 MHz bandwidth is intended for mobile services.

4.5.3 Extended carrier mode (for 8 K, 16 K, 32 K)

Because the rectangular part of the spectrum rolls off more quickly for the larger FFT-sizes, the outer ends of the OFDM signal's spectrum can be extended, i.e. more sub-carriers per symbol can be used for data transport. The gain achieved is between 1,4 % (8 K) and 2,1 % (32 K). Figure 6 compares the spectrum for 2 K with that for 32 K in normal and extended carrier modes, together with the "normal" spectrum mask from [i.13] for reference. The extension is an optional feature, as it may make it more difficult to meet requirements such as spectrum masks and protection ratios. The choice of normal or extended carrier mode is discussed in more detail in clause 5.2.

NOTE: Figure 6 shows the theoretical spectrum of the signal described by [i.1]. Consideration should be given to the effect on the spectrum of any pre-processing, non-linearity in the transmitter, downstream filtering and the relevant spectrum mask for the region and band in use.




Figure 6: Power-spectral-density roll-off at the band edges for 2 K and 32 K

4.5.4 Alamouti-based MISO (in frequency direction)

Although DVB-T supports Single-Frequency Networks (SFNs), the presence of similar-strength signals from two transmitters in a network causes a significant loss of margin because the resulting channel can have deep "notches". There is evidence from the field that portable reception can suffer as a result and requires a higher transmitting power to compensate.

DVB-T2 incorporates the option of using the Alamouti technique [i.8] with a pair of transmitters (figure 7). Alamouti is an example of a Multiple Input, Single Output (MISO) system, in which every constellation point is transmitted by each transmitter, but the second transmitter (TX2 in the figure) transmits a slightly modified version of each pair of constellations, and in the reverse order in frequency. The technique gives performance equivalent to diversity reception in the sense that the operations performed by the receiver result in an optimum combination of the two signals; the resulting signal-to-noise ratio is as though the powers of the two signals had combined in the air. The extra complexity required in the receiver includes a few extra multipliers for the Alamouti processing, and also some parts of the channel estimation need to be duplicated. There is a significant overhead increase in the sense that the density of Scattered Pilots needs to be doubled for a given Guard Interval fraction.




Figure 7: MISO scheme

4.5.5 Preambles (P1 and P2)

The initial symbols of a DVB-T2 physical-layer frame are preamble symbols, transporting a limited amount of signalling data in a robust way. The frame starts with the highly robust differentially BPSK-modulated P1 symbol, with guard intervals at both ends, which carries 7 bits of information (including the FFT size of the payload symbols). The subsequent P2 symbols, whose number is fixed for a given FFT size, provide all static, configurable and dynamic layer-1 signalling (the dynamic part can also be transmitted as "in-band signalling", within the PLPs). The first few bits of signalling (L1-pre signalling) have a fixed coding and modulation; for the remainder (the L1-post signalling) the code rate is fixed to 1/2 but the modulation can be chosen from the options QPSK, 16-QAM and 64-QAM.

The P2 symbol will in general also contain data for common and/or data PLPs, which continue into the other symbols of the T2-frame.

4.5.6 Pilot patterns

Scattered pilots of pre-defined amplitude and phase are inserted into the signal at regular intervals in both time and frequency directions. They are used by the receiver to estimate changes in channel response in both time and frequency dimensions.

Whereas DVB-T applies the same static pattern independent of the FFT size and guard-interval fraction, DVB-T2 has chosen a more flexible approach by defining eight patterns that can be selected depending on the FFT size and Guard Interval fraction adopted for the particular transmission. This reduces the pilot overhead whilst assuring a sufficient channel-estimation quality. The example in figure 8 shows a corresponding overhead reduction from 8 % to 4 % when using the PP3 pattern with a 1/8 guard interval.




Figure 8: Scattered pilot patterns for DVB-T (left) and DVB-T2 (right)

For continual pilots, the percentage in DVB-T2 also depends on the FFT size, leading to a reduction in overhead from 2,5 % to 0,7 % for 8 K, 16 K and 32 K FFT modes without compromising the performance of fine-frequency-synchronisation and CPE-detection algorithms based on them.

One of the eight patterns (PP8) is intended for use with receivers that implement channel estimation based on the data rather than the pilots (see clause 10.3.2.3.5.4). Time Interleaving and PP8 should not be used together, but if a network operator chooses to do so, the receiver performance may be degraded. Similarly, PP8 should not be used in multiple PLP modes.

The choice of pilot pattern is discussed in more detail in clause 5.4.

4.5.7 256-QAM

In DVB-T the highest-order constellation is 64-QAM, delivering a gross data rate of 6 bits per symbol per carrier (i.e. 6 bits per OFDM cell). In DVB-T2, the use of 256-QAM (figure 9) increases this to 8 bits per OFDM cell, a 33 % increase in spectral efficiency and capacity transported for a given code rate. Normally this would require a significantly higher carrier-to-noise ratio (4 dB to 5 dB higher, depending on channel and code rate). This is because the Euclidean distance between two neighbouring constellation points is roughly a half that of 64-QAM, hence reception is more sensitive to noise. However, the performance of LDPC codes is much better than convolutional codes, and if a slightly stronger code rate is chosen for 256-QAM compared to the rate currently being used with 64-QAM using DVB-T, the required C/N is maintained whilst still achieving a significant bit rate increase.




Figure 9: Constellation diagram for 256-QAM

256-QAM is a promising option and is expected to be widely used in the field.

The choice of constellation and code rate is discussed in clause 5.11.

4.5.8 Rotated constellations

A new technique of constellation rotation and Q-delay has been introduced in DVB-T2. After forming a constellation, it is rotated in the complex "I-Q" plane so that each axis on its own (u_1, u_2) now carries enough information to tell which of the points was sent (figure 10).




Figure 10: Rotated 16-QAM constellation

The I and Q components are now separated by the interleaving process so that in general they travel on different frequencies, and at different times. If the channel destroys one of the components, the other component can be used to recover the information. This technique gives no loss of performance in Gaussian channels, and a gain of 0,7 dB in typical fading channels. The gain is even greater in 0 dB-echo channels (e.g. Single-Frequency Networks) and erasure channels (e.g. impulsive interference, deep selective fading). For example, figure 11 shows a 7,6 dB gain in an erasure channel. This means that higher code rates, and hence data rates, can be used. The modest increase in complexity required is not expected to be significant and an FPGA implementation has already been produced to demonstrate this. There is scope for future generations of receivers to give even better performance using Iterative Demodulation, as shown by the "ID" curve.




Figure 11: System behaviour with constellation rotation (blue) and without (black)

4.5.9 16 K and 32 K FFT sizes and 1/128 guard-interval fraction

Increasing the FFT size results in a narrower sub-carrier spacing, but longer symbol duration. The first attribute leads to greater difficulties with inter-carrier interference, hence lower Doppler frequency that can be tolerated. Whilst this is not a setting preferred for mobile reception in UHF band IV/V or higher, it could nevertheless be used in lower frequency bands. However, the second attribute, longer symbol duration, means that the guard-interval fraction is smaller for a given guard-interval duration in time (see figure 12). This reduction in overhead leads to an increase in throughput ranging from 2,3 % to 17,6 %.



Figure 12: Guard interval overhead reduction with larger FFT size

Other advantages consist of better robustness against impulsive noise, quasi-rectangular spectrum down to lower power-spectral-density levels and the option to interpolate in the frequency direction only. The memory requirements for interpolation in the receiver are in the same order for 32 K as for 8 K, but for 16 K they are doubled. The FFT-calculation complexity is only slightly increased.

The 1/128 guard interval fraction is also new in DVB-T2: it can, for example, allow 32 K to be used with an absolute guard interval duration equivalent to 8 K 1/32 with a resulting reduction in overhead.

4.5.10 LDPC/BCH error control coding

Whereas inner and outer error-control coding in DVB-T was based on convolutional and Reed-Solomon codes, ten years of technological development mean that the higher complexity of LDPC decoding can now be handled in the receiver. DVB-T2 uses concatenated LDPC/BCH coding, as for DVB-S2. These codes assure a better protection, allowing more data to be transported in a given channel; they also show a steeper behaviour in the relation of BER to C/N, i.e. they move closer to the ideally desired 'brick wall' behaviour (see figure 13).

Figure 13 shows the results before outer decoding (RS or BCH respectively), and the Reed-Solomon decoder has more powerful error-correcting capabilities than the BCH decoder. A BER of around 10^{-4} before Reed-Solomon is usually assumed to give "Quasi-Error-Free" (QEF) performance after Reed Solomon. The gain achieved here is therefore in the order of 5 dB at the QEF point.




Figure 13: Comparison of error control coding for DVB-T and DVB-T2

4.5.11 Interleaving stages (time, bit, cell, frequency)

The target of the interleaving stages is to spread the content in the time/frequency plane in such a way that neither impulsive noise (disturbance of the OFDM signal over a short time period) nor frequency-selective fading (disturbance over a limited frequency span) would erase long sequences of the original data stream. Furthermore, the interleaving is matched to the behaviour of the error-control coding, which does not protect all data equally. Lastly, the interleaving is designed such that bits carried by a given transmitted constellation point do not correspond to a sequence of consecutive bits in the original stream.

The most significant step from DVB-T to DVB-T2 is the introduction of time interleaving, typically over 70 ms, to provide protection against impulsive noise and time-selective fading.

4.5.12 Peak-to-Average Power Ratio (PAPR) reduction techniques

High PAPR in OFDM systems can reduce RF power-amplifier efficiency, i.e. RF power out vs. supply power in. Two PAPR-reduction techniques, Active Constellation Extension (ACE) and Tone Reservation (TR) are supported in DVB-T2, leading to a substantial reduction of PAPR, at the expense of a small average-power increment and/or at most 1 % reserved sub-carriers. Early practical implementations have shown a reduction of 2 dB in PAPR (at 36 dB MER).

The ACE technique reduces the PAPR by extending outer constellation points in the frequency domain, while the TR reduces the PAPR by directly cancelling out signal peaks in time domain using a set of impulse-like kernels made of the reserved sub-carriers. The two techniques are complementary, i.e. the ACE outperforms the TR in a low-order modulation while the TR outperforms the ACE in a high-order modulation. The two techniques are not mutually exclusive and a combination of them can be used. However, ACE cannot be used with rotated constellations.

4.5.13 Future Extension Frames (FEFs)

In order to build a hook into the original DVB-T2 standard for future advancements such as MIMO or a fully mobile branch of the standard, a placeholder known as Future Extension Frame (FEF) parts is included. The only currently defined attributes of FEF parts, which are inserted between T2-frames (see figure 14), are that they begin with a P1 symbol and their positions in the superframe and duration in time are signalled in the L1 signalling in the T2-frames. This enables early receivers to ignore the FEFs whilst still receiving the T2 signal, as desired.




Figure 14: Co-existence of T2-frames and future extension frames

4.6 Memory and complexity considerations

The DVB-T2 standard has been designed to provide a compromise between receiver complexity and performance of the system.

Memory requirements in the receiver are discussed in various parts of the present document. The most significant memory requirement arises in the Time De-Interleaver, which has around 500 000 cells. Assuming that a cell requires two 10-bit values (for real and imaginary) together with a few bits for channel-state information, the memory size is likely to be around 12 Mbits. Other large memory requirements include the De-jitter buffer at 2 Mbits, and the compensating delay for the channel equaliser: three 16 K symbols of complex cells or around 1 Mbit. The FFT will also need to store at least two 32 K symbols.

The amount of memory needed to decode a particular PLP can be deduced from the configurable L1 signalling. It would therefore be possible to design a simpler receiver with a smaller amount of memory which can only receive certain services. It is intended that all digital TV receivers according to the current standard should have enough memory to decode any PLP, but other applications such as radio reception may allow such lower-complexity receivers.

The computational complexity is bound to be greater than for DVB-T, not least because of the addition of LDPC coding. However, the complexity is not expected to be too great given the advances in silicon geometries since DVB-T was introduced. There are several areas in which the manufacturer can trade off complexity against performance, including the number of LDPC iterations, the type of metric calculation and the use of iterative decoding for rotated constellations.

The power consumption of DVB-T2 might be greater than for DVB-T using the same silicon geometry, although this would be significantly less than for the original DVB-T ICs. Furthermore, features have been included in DVB-T2 to make it possible to build low-power receivers. In particular, the use of Type 1 PLPs, which only appear at the beginning of the T2 frame after the L1 and L2 signalling, allow the receiver to switch off whilst the other PLPs are being transmitted. Furthermore, low data-rate PLPs need not be present in every T2 frame, allowing longer periods of power saving.

4.7 Layer 2 signalling

4.7.1 Overview

The DVB-T2 physical layer provides a number (up to 255 - depending on the FFT size and L1 coding parameters) of transparent Physical-Layer Pipes (PLPs) to the layers above. These PLPs are characterized by the following attributes:

- Constellation (QPSK, 16-QAM, 64-QAM or 256-QAM).
- Code rate (1/2, 3/5, 2/3, 3/4, 4/5, 5/6).
- Time-interleaving depth (from a fraction of a T2 frame up to 255 T2 frames).

In order to allow for multiple PLPs, the PSI/SI signalling has been extended by the addition of a new element: the T2 delivery-system descriptor (T2dsd).

Although [i.1] defines a single profile, its L2 signalling may signal possible future use of Time Frequency Slicing (TFS). Single-profile receivers should ignore any transmissions which use this technique.

4.7.2 T2 delivery-system descriptor

Following the same principles as the existing terrestrial-delivery-system descriptor for DVB-T, this new descriptor for DVB-T2 also provides the signalling for the co-existence of several Transport Streams within the same multiplex (using PLPs).

The T2 delivery-system descriptor achieves this by mapping a Transport Stream to the PLP in which it is carried and to the T2 system that the PLP belongs to. Furthermore, it describes the geographical cells and the centre frequencies used for the transmission of the particular T2 system in the specific multiplex. Signalling is also included of multiple RF channels per cell in parallel; this is for future use and single-profile receivers are not expected to be able to receive such signals.

The T2 delivery-system descriptor is docked to the NIT at the same position as the corresponding DVB-T descriptor was docked.

4.8 Network Topology and Modulator Interface (T2-MI)

Unlike in DVB-T, where a single Transport Stream is carried in a single RF channel, in DVB-T2 there is a degree of flexibility in how the data for each PLP is allocated to the T2-frame; the T2 specification merely requires that certain constraints be met. To enable Single Frequency Network (SFN) operation, decisions on allocation and scheduling are taken once in a T2-Gateway, the results of which are distributed in such a format that each modulator in the network can unambiguously create an identical on-air signal. This format, which is known as the DVB-T2 Modulator-Interface Specification (T2-MI), allows reliable networks of transmitters (in both MFN and SFN configurations) to be constructed.

The ETSI Technical Standard for the T2-MI is reference TS 102 773 [i.20]. TS 102 773 [i.20] defines the T2-MI format and additionally provides a mechanism to support the use of regenerative, off-air repeaters to feed further MFNs and SFNs.

More information on the T2-MI and possible network topologies is given in clause 7.

4.9 Transmitter identification

Proposals for techniques to assist with transmitter identification are currently being considered. More details will be given in future editions of the present document.

5 Choice of parameters

There are a large number of ways of configuring a DVB-T2 system, and this clause will discuss the choice of each of the main parameters in turn.

5.1 Choice of FFT size

The factors affecting the choice of FFT size are well known. Increasing the FFT size will give a greater delay tolerance for the same fractional guard interval, allowing larger Single Frequency Networks (SFNs) to be constructed.

Alternatively, larger FFTs allow the same delay tolerance to be achieved with a smaller overhead due to the guard interval. On the other hand, the larger FFT sizes have a greater vulnerability to fast time-varying channels, i.e. have lower Doppler performance. For a given FFT size, constellation and code-rate, the Doppler performance is roughly proportional to RF bandwidth (halving the bandwidth will halve the carrier spacing resulting in half the Doppler performance) and inversely proportional to the RF frequency.

For delivering high-bit-rate services to fixed, rooftop antennas, in the UHF IV/V bands, or lower bands, the 32 K FFT mode is recommended. In this situation the time variations in the channel are minimised, and 32 K offers the very highest bit rates achievable using DVB-T2.

For mobile reception in UHF band IV/V, or higher bands,, smaller FFT sizes should be used. For mobile use in VHF band III (about 200 MHz) it should be noted that roughly the same Doppler performance should be expected using the 32 K mode as is possible using the 8 K mode at 800 MHz, so 32 K may be an option even at VHF using 7 MHz RF bandwidth. The performance in time-varying channels can also be affected by the choice of pilot pattern - see clause 5.4.


The 1 K FFT mode offers a higher Doppler performance and is intended primarily for operation in the L-band,(about 1,5 GHz), or higher, using a nominal occupied bandwidth of 1,7 MHz. Given that the fundamental sampling rate is lower, the carrier spacing will be correspondingly smaller than it would be in an 8 MHz channel.

5.2 Choice of extended or normal carrier mode

The DVB-T2 signal has been designed so that it should not present more stringent requirements in frequency planning than DVB-T. In many cases, this even remains true when the so-called extended-carrier mode is in use.

The spectrum of a 'classical' guard-interval-OFDM signal exhibits a more-or-less rectangular spectrum for the intended band of the wanted signal, with a roll-off beyond that which depends on the OFDM carrier spacing. The larger the FFT size (for the same sampling rate and nominal bandwidth), the closer the carrier spacing and the more rapidly the spectrum rolls off outside the nominal intended band. It is this property that makes the extended-carrier option feasible for modes with FFT size of 8 K or greater. The extended-carrier option has the obvious benefit of increasing the data capacity.

Reference to figure 15 shows, for example, a comparison of the spectra for 2 K, 32 K normal and 32 K extended-carrier mode. Beyond an offset of 3,88 MHz from the centre frequency, the spectrum of the 32 K extended-carrier signal remains below that of a normal 2 K signal, and thus its propensity to cause interference in this range is actually reduced.



NOTE 1: The spectrum shown in figure 15 assumes that random data is sent in every OFDM cell. The presence of pilot information will cause some fine-detail variations in the wanted signal band. The choice of guard-interval fraction will affect the depth of the ripple in the wanted band (it increases as the guard-interval fraction increases) and also the ripples at the band edge.

NOTE 2: Figure 15 shows the theoretical spectrum of the idealised OFDM signal. Consideration should be given to the effect on the spectrum of any pre-processing, non-linear effects in transmitter amplifiers, downstream filtering and the relevant spectrum mask for the region and band in use.

Figure 15: Detail of theoretical DVB-T2 spectrum for guard-interval fraction 1/8 (for 8 MHz channels)

"Difficult" cases of adjacent-channel interference (e.g. where an adjacent-channel interferer is for some reason substantially stronger than the wanted signal) are in some instances resolved by the application of frequency offsets during the frequency-planning process. It will be clear that this should be done with special caution if an extended-carrier mode is involved. If A, B, C denote three consecutive RF channels, then potential conflict between say A and B could be resolved by offsetting B towards C. However, this will then open up the possibility of conflict between B and C. If either B or C (or both) use an extended-carrier mode then the scope for using an offset will be less than if they use normal carrier mode.

5.3 Choice of Guard interval

DVB-T2 offers a wide range of possible guard intervals in order to support a range of broadcasters' needs. We should distinguish two concepts:

- the *guard-interval duration* T_G ; and
- the *guard-interval fraction* $GIF = T_G/T_U$.

The former has the dimensions of time and the latter is dimensionless.

The simplest view is to treat the guard interval as a hard limit to the channel extent that can be tolerated by the system. Assuming the channel extent for a particular broadcast scenario is known, it is then a simple matter to choose a T_G that suffices to match or exceed it. Note that this choice will also require consideration of the FFT size as well. The greatest capacity is given by minimising the GIF, which thus implies maximising T_U . However, there are other constraints on the choice of FFT, concerning the degree of Doppler effects to be expected in the scenario of interest. These may set a limit to the FFT size (and thus T_U) that can be chosen. So, following this simple-reasoning approach, the decision process would be:

- determine maximum channel extent;
- pick next biggest T_G ;
- determine Doppler scenario of the channel;
- from this, determine largest acceptable FFT size under Doppler constraints;
- pick largest FFT size which offers desired T_G while satisfying Doppler constraints.

However, this approach is rather over-simplistic.

It is true that if the channel extent is less than the guard-interval duration, and the receiver time-synchronisation circuitry correctly places the received channel extent within the guard-interval-wide window, then no Intersymbol Interference (ISI) will occur. Once the channel extent exceeds the guard-interval duration, then some degree of ISI *will* occur. However, the significance of this ISI - the harm done - depends on:

- Does the part of the channel impulse-response that falls outside the safe window contain a significant proportion of the received signal energy?
- How far does it fall outside? (The harm done depends on the excess extent as a proportion of T_U , providing an additional reason to favour large FFTs where Doppler constraints permit).
- What is the working point? (Less ISI is acceptable if working with a mode that requires high SNR than for a more-rugged one).

A further question concerns another failure mechanism:

- How does the channel extent compare with the Nyquist limit for the pilot pattern that has been chosen?

Exceeding this Nyquist limit means that channel equalisation will be incorrect, and could cause problems even when the apparent significance of ISI was small.

We should also note that establishing a value for "maximum channel extent" is far from straightforward, whether there is a single transmitter or a single-frequency network of them.

With a single transmitter, the channel extent is brought about by natural reflections of the transmitted signal from objects in the environment. Predicting all the echoes in a mountainous area is in principle possible, given sufficient terrain data and computing power. On the other hand, a flat area might be thought to be without problems - until a reflective tall building is constructed! In other words, the necessary guard interval might be greater than expected.

Now, at the other extreme, consider a large-scale single-frequency network.

Assuming all the transmissions are co-timed (no deliberate time offsets have been introduced), then the maximum channel extent is clearly limited by the distance between the most-separated transmitters in the network. This value could be very large, and thus perhaps appear to require a guard interval even larger than the biggest available in DVB-T2. This does not mean that it is necessary to choose such a value.

The contribution of such a long-distance path is likely to be severely restricted by propagation, unless the geography is particularly unfortunate - normally it would be sufficiently beyond line-of-sight that it would be seriously attenuated. Procedures for predicting such propagation and planning radio-frequency systems are well established.

However, radio propagation is not immutable and certain. Various anomalous-propagation modes can occur, albeit with reduced probability. So it is possible that the most-distant path just discussed, which normally delivers a negligible contribution, can, during anomalous-propagation conditions (e.g. atmospheric ducting) deliver an enhanced signal that *does* cause problems. For this reason there are no certainties in broadcasting; you can plan for a system to work for a certain high % of time, but not 100 %. This is nothing new; it also occurred with analogue broadcasting and multi-frequency networks. Note also that if the same transmitter locations were used in a multi-frequency network, the same long-distance path might cause similar problems. The predicted attenuation would normally be great enough that it would be acceptable to plan for the two transmitters to share a common frequency, but in this case transmitting different signals as part of different networks. During anomalous propagation the distant transmitter would in this case cause co-channel interference instead of a delayed, out-of-guard-interval SFN contribution; the harm done in the two cases would be generally similar.

So it would not normally be necessary to choose the guard interval in a large SFN to be as great as indicated by the most-separated transmitters. Careful planning taking into account propagation predictions for real terrain is needed. It would be equally unwise to make a generalisation like "the guard interval requirement is determined by the location of the first ring of transmitters"; this statement might be true, but it would depend on the geography.

Table 61 of [i.1] lists the available guard-interval choices, in terms of supported combinations of GIF and FFT size. The GIFs are all simple fractions where the denominator is an integer power of 2; the numerator is either 1 or 19. The cases where the numerator of the GIF is 19 are for use where the maximum guard interval is wanted while using a particular FFT size and pilot pattern (e.g. to obtain the greatest guard interval under simultaneous Doppler constraints). These cases push the frequency-interpolation process (used in channel estimation in the receiver) closer to the fundamental Nyquist limit.

Table 3 lists the guard interval durations for 8 MHz channels for each allowed combination of FFT size and guard interval fraction.

Table 3: Guard-interval durations in μ s for 8 MHz channels

| FFT size | Guard-interval fraction | | | | | | |
|----------|-------------------------|------|------|--------|-----|--------|-----|
| | 1/128 | 1/32 | 1/16 | 19/256 | 1/8 | 19/128 | 1/4 |
| 32 K | 28 | 112 | 224 | 266 | 448 | 532 | N/A |
| 16 K | 14 | 56 | 112 | 133 | 224 | 266 | 448 |
| 8 K | 7 | 28 | 56 | 66.5 | 112 | 133 | 224 |
| 4 K | N/A | 14 | 28 | N/A | 56 | N/A | 112 |
| 2 K | N/A | 7 | 14 | N/A | 28 | N/A | 56 |
| 1 K | N/A | N/A | 7 | N/A | 14 | N/A | 28 |

5.4 Choice of pilot pattern

With the exception of one usage scenario (explained later), a DVB-T2 receiver makes measurements of the channel using *Scattered Pilots* (SPs) and then interpolates between these measurements to construct estimates of the channel response for every OFDM cell (see clause 10.3.2.3). The measurements need to be sufficiently dense that they can follow channel variations as a function of both frequency (carrier index) and time.

Several pilot patterns are available, named PP1 to PP8, with the intention of providing efficient options for different channel scenarios. They are defined in clause 9.2.3.1 of [i.1], with the parameters of the different patterns listed in its table 51. Each pattern can (in principle) support time and frequency variations up to corresponding Nyquist limits. These are explained in clause 10.3.2.3.2. The applicable limits do depend on certain assumptions as to how the receiver works, in particular whether it uses interpolation in both time and frequency, or in frequency only. PP8 is intended as a special case, where interpolation between measurements is not used, but instead a decision-feedback technique (CD3-OFDM) [i.15] continually updates estimates made at the start of the frame.

Frequency-only interpolation supports a greater Doppler bandwidth, at the expense of supportable channel extent. By avoiding the need for temporal interpolation it also avoids the need for the corresponding storage in the receiver. In the case of the 32 K-FFT mode, it is assumed that storage limitations only allow temporal interpolation over at most two symbols, which implies that frequency-only interpolation will be performed if PP7 is used.

The operator should choose a pilot pattern considering the expected channel for the type of usage that it is desired to support, and being aware of the trade-off between capacity and performance.

Capacity is clearly reduced as the density of inserting Scattered Pilots increases; the SP overhead can be expressed as the simple fraction $1/(D_x D_y)$. (This simple approach considers SPs only and does not take account of the presence of other types of pilot). It follows that PP1 and PP2 have the greatest overhead (8,33 %), and PP7 and PP8 the least (1,04 %). Of course, the higher-overhead cases are included since they have performance advantages. PP1 offers the highest Nyquist limit for channel extent, and PP6 and PP7 the least. PP2, PP4 and PP6 in contrast offer the highest Nyquist limit for Doppler. This is summarised in table 4. The options have to be considered in conjunction with table 52 of [i.1], which shows which guard-interval/pilot-pattern combinations are permitted.

Table 4: Comparison of Scattered Pilot patterns

| | PP1 | PP2 | PP3 | PP4 | PP5 | PP6 | PP7 | PP8 | Interpretation |
|-------------|--------|--------|--------|--------|--------|--------|--------|---------|---|
| D_x | 3 | 6 | 6 | 12 | 12 | 24 | 24 | 6 | Separation of pilot-bearing carriers |
| D_y | 4 | 2 | 4 | 2 | 4 | 2 | 4 | 16 | Length of sequence in symbols |
| $1/D_x D_y$ | 8,33 % | 8,33 % | 4,17 % | 4,17 % | 2,08 % | 2,08 % | 1,04 % | 1,04 % | SP overhead |
| $1/D_x$ | 1/3 | 1/6 | 1/6 | 1/12 | 1/12 | 1/24 | 1/24 | 1/6 | T_{Nyquist}/T_U , for f-&t interpolation |
| $1/D_x D_y$ | 1/12 | 1/12 | 1/24 | 1/24 | 1/48 | 1/48 | 1/96 | 1/96 | T_{Nyquist}/T_U , for f-only interpolation |
| $1/(2D_y)$ | 0,125 | 0,25 | 0,125 | 0,25 | 0,125 | 0,25 | 0,125 | 0,03125 | f_{NYQUIST}/f_S , for f-&t interpolation (\pm) |
| | 0,5 | 0,5 | 0,5 | 0,5 | 0,5 | 0,5 | 0,5 | 0,5 | f_{NYQUIST}/f_S , for f-only interpolation (\pm) |

Note that there may be advantages to choosing a pilot pattern that supports a greater Nyquist limit for channel extent than appears to be needed to support the chosen guard interval, despite the increase in overhead that such a choice implies. Circumstances supporting such an approach include:

- when the maximum channel extent is not well defined - while low-level signal components whose delays lie outside the nominal guard-interval window may not cause excessive ISI, they can cause aliasing in the channel measurement if they exceed the Nyquist limit;
- receiver interpolation performance will worsen as the Nyquist limit is approached;
- where the design time-width of the receiver's frequency interpolator is substantially less than the Nyquist limit, the interpolator can be designed to reduce estimation noise.

In MISO the above description has to be modified to take account of the further sub-sampling in frequency, which arises because alternate pilot-bearing carriers are used for the sum and difference estimation (see clause 10.3.2.4). The channel extent supported by each pilot pattern in MISO is one-half of that shown in table 4.

Time Interleaving and PP8 should not be used together, but if a network operator chooses to do so, the receiver performance may be degraded.

5.5 Choice of frame length

The T2-frame length, i.e. the number of symbols L_F in a T2-frame, is a configurable parameter and can be chosen as required by the broadcaster or network operator. There are three constraints imposed by [i.1]:

- The maximum T2-frame duration is 250ms. This enables the receiver to know definitively after this amount of time that no T2 signal is present on a given frequency.
- The minimum number of symbols is specified. This simplifies the process of channel estimation as described in clause 10.3.2.3.
- In 32 K, there are always an even number of symbols. This ensures that the Frequency de-Interleaver can be implemented with only one 32 K-symbol's worth of memory - see clause 10.4.1.

In addition to these normative constraints, there are a number of other factors affecting choice of frame length:

- A longer value for frame length generally decreases the percentage overhead associated with the preamble symbols P1, the L1 signalling and the higher density of pilots in the P2, thus increasing the total bit-rate (but see below).
- A shorter value for frame length means that the P1 and P2 occur more frequently, allowing faster lock-up and service acquisition.
- Longer frames may be used support longer time-interleaving depths.
- Shorter frames may be used to achieve higher peak bit-rates without using multiple time-interleaving blocks.

Although increasing the frame length reduces the preamble overhead, there is a second-order effect that should be considered in choosing the exact number of symbols. Ignoring the case of multi-frame interleaving, there are always a whole number of FEC-blocks in a T2-frame; remaining cells which do not make up a whole FEC block will generally be dummy cells and represent a loss of capacity (see clause 6.1.5). As a result, the total bit-rate depends in fine detail on the exact number of symbols, and therefore cells, in a frame. Reducing the frame length by one symbol can often *increase* the total bit-rate significantly by reducing the number of dummy cells. Figure 16 shows an example of this variation for the case of 32 K FFT with extended carrier mode and guard-interval fraction 1/128, 256-QAM with code rate 3/5 using long LDPC blocks. Note that the vertical scale for bit rate is exaggerated.

A further consideration is interleaving depth. As the number of symbols is increased, the number of cells per T2-frame increases. Since there is a limit to the number of cells in one Time-Interleaving block (TI-block), there is a point at which increasing the frame length by 1 means that the number of TI-blocks per T2-frame needs to be increased. This will have the immediate effect of reducing the time interleaving depth, since the frame duration has hardly increased, but it is now divided into more TI-blocks. Figure 16 also shows how the time-interleaving depth varies with frame length. In this example it can be seen that there is a clear best choice: 60 symbols gives almost the maximum possible bit-rate whilst also maximising the interleaving depth. For other sets of parameters the optimum might be less obvious and a degree of compromise would be required.

The exact choice of frame length will depend on the selections made for several of the other parameters, including:

- FFT size.
- Guard interval.
- The use of extended-carrier mode.
- The combinations of different types of PLPs.

Maximum capacity of T2-frame (32KE, PP7, 256-QAM, CR=3/5, GI=1/128, Nldpc=64800)




Figure 16: Variation of data rate and time-interleaving depth with frame length

5.6 Choice of superframe length

The minimum superframe length is 2 frames, and in many cases, including in Input Mode A, there will be no reason to choose a longer superframe length than this.

Longer superframes may be needed if:

- multi-frame interleaving is used (i.e. $P_1 > 1$) for some PLPs; or if
- PLPs use the "frame skipping" mechanism (i.e. $I_{\text{jump}} > 1$); or if
- FEFs are used.

EN 302 755 [i.1] requires that there are a whole number of interleaving frames for each PLP in one superframe, as well as a whole number of FEF repetitions. The minimum superframe length in this case is therefore equal to the lowest common multiple of the values of $(P_1 \times I_{\text{jump}})$ across all PLPs, and of FEF_INTERVAL.

5.7 Choice of input mode (A or B)

There are two basic input modes for single-RF operation: Input Modes A and B.

Input Mode A takes a single Transport Stream (or Generic Stream) of constant total bit-rate. The Transport Stream may contain a number of services, possibly statistically multiplexed together.

In Input Mode B, the T2 system carries more than one input stream, which may include both Transport and Generic Streams. Each stream is carried by a separate PLP; there may be statistical multiplexing between the PLPs so that the bit rate of an individual PLP is variable, and there may be one or more common PLPs, each shared by a group of PLPs.

As far as the physical layer is concerned, Input Mode A is simply a degenerate case of Input Mode B, in which the number of PLPs is equal to one, and the dynamic signalling is in fact static, with the one and only PLP always mapped to the same cell addresses in the T2-frame. An example of the signalling for Mode A is given in clause 8.10.5.

Given a number of services to convey, a broadcaster or network operator may wish to consider whether to use Mode A or B. Table 5 compares the relative advantages and disadvantages of the two modes.

Table 5: Comparison of Input Modes A and B

| Mode A | Mode B |
|---|--|
| Simple generation and distribution (single TS, T2-MI for SFN) | Distribution requires multiple TSs or T2-MI |
| All services can share all elements in the TS | Shared elements need to be in Common PLP |
| Maximum Interleaving Depth ~70 ms | Interleaving depth can be extended at the expense of peak bit-rate (e.g. to 140 ms @ 18 Mbit/s peak) |
| All services have same protection | Possibility of rugged mobile services and high-bit-rate fixed services in same multiplex |
| Receiver buffer management straightforward | Receiver buffer management (using TTO) more complicated |

5.8 Choice of number of sub-slices

Increasing the number of sub-slices increases the time diversity of each PLP. A large number of sub-slices also results in a more even delivery of data throughout the T2 frame, making management of the de-jitter buffer and time-de-interleaver memory easier. In general, therefore, the number of sub-slices should be as large as possible, if receiver power saving is not an issue.

However, some values for the number of sub-slices, though allowed by [i.1], result in unwanted interactions. For the case where all the PLPs use the same QAM modulation, it is recommended that the values given in table 6 be used. Further explanation of these choices is given in clause 5.8.1.

The derivation of the allowable sub-slicing values is explained in annex A.

Table 6: Recommended numbers of sub-slices

| FFT size | Number of sub-slices for each modulation | | | |
|-----------------|---|---------------|---------------|----------------|
| | QPSK | 16-QAM | 64-QAM | 256-QAM |
| 1 K-16 K | 1 296 | 648 | 432 | 324 |
| 32 K | 432 | 216 | 144 | 108 |

For services targeted at power-saving receivers, there might be some benefit in using a smaller number of sub-slices, but it is recommended that Type 1 PLPs be used for such services.

5.8.1 Avoiding undesirable interaction of subslicing with carrier mapping and frequency interleaving

If a poor choice of sub-slicing parameters is made, an undesirable interaction between the subslicing, the carrier mapping and the frequency interleaver can result. The potential problem is that the cells of individual FEC blocks could be consistently mapped to very few carriers, resulting in a loss of frequency diversity. A particular poor channel could then damage a high proportion of the cells of one FEC block, leading to a weakness.

The problem can only exist if the total number of FEC blocks for all type 2 PLPs shares common divisors with the number of data cells in a symbol.

A key question is how many carriers need to be mapped for the effects to be negligible. At first sight it appears that for a FEC block to be mapped to only 50 % of carriers represents a significant loss in diversity. In practice, however, this is not the case, since the mapping is random, and a given channel will "erase" a proportion of the carriers - but the resulting number of lost cells will be very similar to the case when 100 % mapping is achieved. Overall, the loss of diversity will only be significant if the absolute number of carriers mapped is relatively low - less than say about 100.

The numbers of lost cells for different numbers of mapped carriers has been studied, for a channel in which a fixed proportion of carriers is erased. A typical example is a 1 K case when every fourth carrier is erased. With 100 % mapping, 75 % of cells would remain, but with a 25 % random mapping, the proportion remaining might vary between say 73 % to 77 %. The studies used the DVB-T2 frequency-interleavers, the actual number of data carriers for every combination of FFT size and pilot pattern, and realistic erasure ratios.

With suitable choice of number of sub-slices, it has been found that, even in worst cases (where the total number of FEC blocks shares common divisors with number of data carriers per symbol), the variation per FEC block following carrier erasure is almost always within 1 % to 2 % of the erasure ratio (there are just 2 cases where it is 4 % and 5 %). This is helped by two factors:

- The extension of the repetition pattern due to the pilot pattern sequence: the pilot positions, and so the carriers to which cells are mapped, change from symbol to symbol.
- The maximum expected erasure ratio is limited for a given pilot-pattern allocation (e.g. PP2/3 are only used with max GIF 1/8, so the maximum erasure ratio due to a 0 dB echo is 1 in 8)

A suitable choice of number of sub-slices for each combination of FFT size and modulation was given in table 6 in clause 5.8. For the values in the table, all configurations will work within sensible constraints. These choices ensure that:

- all FEC blocks are mapped to at least 200 symbols (or all symbols when there are fewer than 200);
- all FEC blocks are mapped to at least 20 % of the OFDM carriers.

Other sub-slicing values are retained in the specification to allow different trade-offs to be made. They may also be useful for cases in which different PLPs use different QAM modulations. A similar analysis would need to be performed to determine a good choice for the number of subslices in such situations.

Further analysis of the interaction is given in annex A.

5.8.2 Sublicing with multi-frame interleaving

When Type 2 PLPs are interleaved over multiple frames, i.e. when $P_1 > 1$, the number of time interleaver rows per FEC block has to be divisible by $N_{\text{subslices}} \times P_1$, the total number of subslices into which the Interleaving Frame will be divided. $N_{\text{subslices}}$ has to be chosen such that, for each PLP using multi-frame interleaving, $N_{\text{subslices}} \times P_1$ appears as an allowed value in annex K of [i.1], for the constellation and FEC block size used for the PLP.

In addition to this mandatory requirement, the issue of interaction between sub-slicing, carrier mapping and frequency interleaving (described in clause 5.8.1) also needs to be analysed for the case of multi-frame interleaving.

5.9 Choice of time interleaving parameters

The choice of number of TI blocks per Interleaving Frame depends on two factors. Increasing the number of TI blocks for a given frame duration reduces the interleaving time, so reducing the time diversity and therefore the system's resistance to impulsive interference and fast time-varying channels. On the other hand, increasing the number of TI blocks increases the maximum data rate for a PLP, since the maximum number of cells in a TI block is fixed. The trade-off in mode A is discussed in clause 5.5. For mode B, using multiple TI blocks can make management of the de-jitter buffer more complicated: see clause 8.8.

An alternative option is multi-frame interleaving, where one TI block is spread over multiple T2 frames. Typically this will be used for low-data-rate PLPs, since the maximum bit-rate will be limited. It is particularly useful for Type 1 and common PLPs, which would otherwise have a very limited amount of time diversity because they only have one sub-slice in each T2 frame. It may also be used for Type 2 PLPs, perhaps to provide greater time diversity for lower-data-rate services targeted at mobile or handheld reception.

Time interleaving is optional, but it is recommended that it be used because of the time diversity it provides. The option to disable Time Interleaving is provided principally for applications where data-directed demodulation is to be used in the receiver (see clause 10.3.2.3.5.4). The delay in de-interleaving and re-interleaving the data would be too great to allow channel variations to be tracked using this method.

When common PLPs are used, the time de-interleaving resources in the receiver are shared between the data and its common PLP, and this should be taken into account when choosing the interleaving parameters.

5.10 Choice of rotated constellations

This technique does not show any known performance disadvantages, while greatly improving the robustness of the receiver to severe channel conditions. However, it cannot be used together with the ACE PAPR reduction technique.

Consequently, its use is recommended at all times except when ACE is in use.

5.11 Choice of code rate, block length and constellation

The principles affecting choice of code-rate and constellation are similar to those in DVB-T. Higher code-rates and higher-order constellations both give greater bit rates but require higher signal-to-noise ratios. However, it should be noted that the performance of LDPC codes is significantly better than convolutional codes used in DVB-T. Consequently, the SNR requirement for a given LDPC code-rate is lower than for a convolutional code with the same code rate. Conversely, at a given SNR, a higher code rate can be used and therefore a higher data rate achieved.

The performance of the short codes is some tenths of a dB worse than normal codes. However, they can be used for low-bit-rate applications requiring shorter latency. Alternatively, if long blocks are used for low-data-rate services, a user could choose either multi-frame interleaving (giving greater time diversity) or frame skipping (giving potential for power-saving); both options would entail a greater latency.

5.12 Choice of mode adaptation options

There are a number of options for mode adaptation, described in clause 5.1 of [i.1]:

- Normal Mode (NM) or High Efficiency Mode (HEM).
- Use of Input Stream Synchronisation (ISSY).
- Use of Null Packet Deletion (NPD).

It is recommended that the options be used as follows:

- Where compatibility with DVB-S2 [i.17] is required, NM should be used with the same combination of ISSY and NPD as in the original DVB-S2 system. Otherwise.
- In input mode A:
 - HEM should be used.
 - ISSY need not be used unless one of the other parameters makes it mandatory according to [i.1].
 - NPD need not be used.
- In input mode B:
 - HEM should be used unless the number of FEC blocks per Interleaving Frame can be less than three, in which case NM may be used to allow the minimum three ISSY variables (TTO, BUFS and ISCR) to be delivered in every Interleaving Frame.
 - ISSY should be used, and is mandatory for PLPs with variable bit-rates and/or using null-packet deletion. Without ISSY, the receiver will have to attempt to manage its own de-jitter buffer and this might result in under-or overflow or excessive jitter in the output Transport Stream.
 - NPD should be used if statistical multiplexing is performed between PLPs, since null packets will be used to carry the variable bit-rate services in a constant bit-rate Transport Stream.

6 Anatomy of the DVB-T2 signal

This clause will describe the various concepts involved in the DVB-T2 signal, and the relationship between them.

In terms of framing structures, there is a basic division into the *physical* frame structure of superframes, T2-frames, and symbols; and what will be termed the *logical* frame structure of BBFRAMES, Interleaving frames and TI-blocks. The physical frame structure applies to the entire T2-system, whereas the logical frame-structure is a PLP-specific concept and can have different parameters for each PLP.

Common to both the physical and logical frame-structure is the *OFDM cell*. An OFDM cell is that part of an OFDM signal corresponding to one carrier in one OFDM symbol, and can be modulated by one constellation point, pilot or reserved tone. The term "OFDM cell" is also used to refer to the modulation value for a cell, since many operations are performed on these entities before they are finally assigned to a particular symbol and carrier.

6.1 Physical frame structure

6.1.1 Superframe

The largest entity of a DVB-T2 system is a super frame. A super frame carries T2-frames and may also carry Future-Extension Frame (FEF) parts as depicted in figure 17. The maximum number of T2-frames in a super frame is 255. The maximum length of a T2-frame is 250 ms and the maximum length of a FEF part is 250 ms. The T2-frames may have a length that is different from the FEF parts. However, all T2-frames have equal lengths within a super frame. The same applies for all FEF parts. A FEF part can be inserted after every T2-frame, so the number of FEF parts in a super frame is 0 to 255. If FEF parts are used, the super frame always end with a FEF part. The number of T2-frames between two FEF parts is signalled, if FEFs are used. The maximum length of a super frame is $255 \times 250 \text{ ms} = 63,75 \text{ s}$, if not using FEFs, and $252 \times 2 \times 250 \text{ ms} = 127,5 \text{ s}$, if using FEFs.

A T2-frame is divided into OFDM symbols. Every T2-frame starts with a P1 symbol. Also, every FEF part starts with a P1 symbol. Thus, the interval between two P1 symbols is at most 250 ms.




Figure 17: The frame structure of DVB-T2 consists of super frames carrying T2-frames and possibly also Future-Extension Frames

EXAMPLE: A super frame contains 20 T2-frames. If it is wished to include one FEF part in every super frame, the FEF_interval = 20, and the FEF parts occurs after the last T2-frame in the super frame. If it is wished to include two FEF parts in every super frame, the FEF_interval = 10, and the FEF parts are inserted after the 10th T2-frame and after the 20th T2- frame.

The purpose of T2-frames is to carry the PLPs and L1 signalling. Thus, the T2-frame carries the DVB-T2 services and related signalling. The purpose of the Future-Extension Frames is to enable flexible mixing of services carried as defined in the current DVB-T2 standard and services carried as defined in a future version of EN 302 755 [i.1]. The FEF parts may also be empty or contain no data. A receiver designed for DVB-T2 reception should be able to detect and correctly handle FEF parts, so that the reception of T2-frames is not disturbed in any way.

6.1.2 T2-frame

A T2-frame consists of a P1 symbol, N_{P2} P2 symbols and a configurable number of data symbols, of which the last can be a special frame closing symbol for some parameter combinations (see clause 6.1.2.1). The number of P2 symbols N_{P2} depends on the FFT size and is defined in table 7.

Table 7: Number of P2 symbols N_{P2} for each FFT size

| FFT Size | N_{P2} |
|----------|----------|
| 1 K | 16 |
| 2 K | 8 |
| 4 K | 4 |
| 8 K | 2 |
| 16 K | 1 |
| 32 K | 1 |

The duration of the T2-frame is determined by the FFT size, guard interval and number of OFDM symbols used. The maximum allowed T2-frame length is 250 ms, which imposes a limit on the maximum number of OFDM symbols L_F for the different FFT sizes and guard intervals as presented in table 8. The T2-frame length is calculated by:

$$T_F = L_F \times T_s + T_{P1},$$

where T_s is the total OFDM symbol duration defined by $T_s = T_U + T_G = T_u \times (1 + GIF)$ and the duration of the P1 symbol is 0,224 ms. L_F includes all P2 symbols and data symbols so that $L_F = N_{P2} + L_{data}$. GIF is the guard interval fraction.

EXAMPLE: The maximum frame length for FFT size 8 K and guard-interval fractions $GIF=1/8$ is $247 \times (0,896 \text{ ms} \times (1 + 1/8)) + 0,224 \text{ ms} = 249,200 \text{ ms}$.

Table 8: Maximum frame length L_F in OFDM symbols, including P2 and data symbols, for different FFT sizes and guard-interval fractions (for 8 MHz bandwidth)

| FFT size | T_u (ms) | Guard-interval fraction | | | | | | |
|----------|------------|-------------------------|-------|-------|--------|-------|--------|-------|
| | | 1/128 | 1/32 | 1/16 | 19/256 | 1/8 | 19/128 | 1/4 |
| 32 K | 3,584 | 68 | 66 | 64 | 64 | 60 | 60 | NA |
| 16 K | 1,792 | 138 | 135 | 131 | 129 | 123 | 121 | 111 |
| 8 K | 0,896 | 276 | 270 | 262 | 259 | 247 | 242 | 223 |
| 4 K | 0,448 | NA | 540 | 524 | 519 | 495 | 485 | 446 |
| 2 K | 0,224 | NA | 1 081 | 1 049 | 1 038 | 991 | 970 | 892 |
| 1 K | 0,112 | NA | NA | 2 098 | 2 076 | 1 982 | 1 941 | 1 784 |

The P2 and data symbols carry L1 signalling, PLPs and auxiliary streams. The L1 signalling is always carried in the P2 symbols, never in the data symbols. The PLPs and auxiliary streams may be carried in P2 symbols or data symbols. The structure of the T2-frame is depicted in figure 18. The mapping of the L1 signalling, PLPs and auxiliary streams into the OFDM symbols of the T2-frame is depicted in figure 19.




Figure 18: Structure of the T2-frame




Figure 19: Mapping of the L1 signalling, PLPs and auxiliary streams into the T2-frame

The P2 symbols, data symbols and Frame-Closing Symbol (when present) contain different numbers of pilots and thus different number of active OFDM cells that can be used for carrying signalling or data. The number C_{P2} of available cells in each P2 symbol is given in table 9. The number C_{data} of available cells in normal data symbols depends on the pilot pattern and use of reserved tones and is given in table 10. Where a Frame-Closing Symbol is used, it has a number of active cells C_{FC} given in table 11 as well as a number of unmodulated cells (see clause 6.1.2.1).

Table 9: Number of available data cells C_{P2} in one P2 symbol

| FFT Size | C_{P2} | |
|----------|----------|--------|
| | SISO | MISO |
| 1 K | 558 | 546 |
| 2 K | 1 118 | 1 098 |
| 4 K | 2 236 | 2 198 |
| 8 K | 4 472 | 4 398 |
| 16 K | 8 944 | 8 814 |
| 32 K | 22 432 | 17 612 |

Table 10: Number of available data cells C_{data} in one data symbol

| FFT Size | C_{data} (no tone reservation) | | | | | | | | TR cells | |
|----------|----------------------------------|--------|--------|--------|--------|--------|--------|--------|----------|-----|
| | PP1 | PP2 | PP3 | PP4 | PP5 | PP6 | PP7 | PP8 | | |
| 1 K | 764 | 768 | 798 | 804 | 818 | | | | 10 | |
| 2 K | 1522 | 1 532 | 1 596 | 1 602 | 1 632 | | 1 646 | | 18 | |
| 4 K | 3 084 | 3 092 | 3 228 | 3 234 | 3 298 | | 3 328 | | 36 | |
| 8 K | Normal | 6 208 | 6 214 | 6 494 | 6 498 | 6 634 | | 6 698 | 6 698 | 72 |
| | Extended | 6 296 | 6 298 | 6 584 | 6 588 | 6 728 | | 6 788 | 6 788 | 72 |
| 16 K | Normal | 12 418 | 12 436 | 12 988 | 13 002 | 13 272 | 13 288 | 13 416 | 13 406 | 144 |
| | Extended | 12 678 | 12 698 | 13 262 | 13 276 | 13 552 | 13 568 | 13 698 | 13 688 | 144 |
| 32 K | Normal | | 24 886 | | 26 022 | | 26 592 | 26 836 | 26 812 | 288 |
| | Extended | | 25 412 | | 26 572 | | 27 152 | 27 404 | 27 376 | 288 |

NOTE: An empty entry indicates that the corresponding combination of FFT size and pilot pattern is never used.

Table 11: Number of available active cells C_{FC} in the frame closing symbol

| FFT Size | C_{FC} (no tone reservation) | | | | | | | | TR cells |
|----------|--------------------------------|-------|--------|-------|--------|-------|--------|--------|----------|
| | PP1 | PP2 | PP3 | PP4 | PP5 | PP6 | PP7 | PP8 | |
| 1 K | 402 | 654 | 490 | 707 | 544 | | | | 10 |
| 2 K | 804 | 1 309 | 980 | 1 415 | 1 088 | | 1 396 | | 18 |
| 4 K | 1 609 | 2 619 | 1 961 | 2 831 | 2 177 | | 2 792 | | 36 |
| 8 K | Normal | 3 218 | 5 238 | 3 922 | 5 662 | 4 354 | | 5 585 | 72 |
| | Extended | 3 264 | 5 312 | 3 978 | 5 742 | 4 416 | | 5 664 | 72 |
| 16 K | Normal | 6 437 | 10 476 | 7 845 | 11 324 | 8 709 | 11 801 | 11 170 | 144 |
| | Extended | 6 573 | 10 697 | 8 011 | 11 563 | 8 893 | 12 051 | 11 406 | 144 |
| 32 K | Normal | | 20 952 | | 22 649 | | 23 603 | | 288 |
| | Extended | | 21 395 | | 23 127 | | 24 102 | | 288 |

NOTE: An empty entry indicates that frame-closing symbols are never used for the corresponding combination of FFT size and pilot pattern. In this case the number of data cells in the last symbol is C_{data} .

6.1.2.1 Frame-Closing Symbols

In some combinations of FFT size, Scattered Pilot pattern and guard interval, the last symbol of the T2 frame is a special frame-closing symbol. This has a higher density of pilots in order to allow frequency-only interpolation to be used for the symbol itself, and temporal interpolation (as opposed to extrapolation) for the symbols immediately preceding it.

The higher pilot density would in any case result in a lower capacity for this symbol than for a normal symbol, since more pilots implies fewer data cells. This reduced capacity N_{FC} is tabulated in table 12 (reproduced from table 43 of [i.1]), which gives the number of data cells in the frame-closing symbol.

Table 12: Number of data cells N_{FC} for the frame-closing symbol

| FFT Size | N_{data} for frame-closing symbol (no tone reservation) | | | | | | | | TR cells |
|----------|---|-------|--------|--------|--------|--------|--------|--------|----------|
| | PP1 | PP2 | PP3 | PP4 | PP5 | PP6 | PP7 | PP8 | |
| 1 K | 568 | 710 | 710 | 780 | 780 | | | | 10 |
| 2 K | 1 136 | 1 420 | 1 420 | 1 562 | 1 562 | | 1 632 | | 18 |
| 4 K | 2 272 | 2 840 | 2 840 | 3 124 | 3 124 | | 3 266 | | 36 |
| 8 K | Normal | 4 544 | 5 680 | 5 680 | 6 248 | 6 248 | | 6 532 | 72 |
| | Extended | 4 608 | 5 760 | 5 760 | 6 336 | 6 336 | | 6 624 | 72 |
| 16 K | Normal | 9 088 | 11 360 | 11 360 | 12 496 | 12 496 | 13 064 | 13 064 | 144 |
| | Extended | 9 280 | 11 600 | 11 600 | 12 760 | 12 760 | 13 340 | 13 340 | 144 |
| 32 K | Normal | | 22 720 | | 24 992 | | 26 128 | | 288 |
| | Extended | | 23 200 | | 25 520 | | 26 680 | | 288 |

NOTE: An empty entry indicates that frame-closing symbols are never used for the corresponding combination of FFT size and pilot pattern.

However, there is a further reduction in capacity because some of the data cells are not modulated, i.e. their modulation value $c_{m,l,k}$ is zero. This is done in order to ensure that the frame-closing symbol has approximately the same power as the normal symbols even though there are more pilots than normal and the pilots are boosted by the same amount as in normal symbol. The number of active cells, i.e. modulated data cells, is the capacity C_{FC} given in table 11 (table 4 of [i.1]); it can be seen that these values are significantly smaller than the corresponding values in the other table.

It is probably most straightforward to insert the unmodulated cells in the frame builder as shown in figure 20. They are the last cells ($x_{m,LF-1,CFC} \dots x_{m,LF-1,NFC-1}$) of the last symbol of the frame and are set to zero, as shown by the white cells in the figure. Although the unmodulated cells are all adjacent at the frame-builder output, the frequency interleaver will distribute them throughout the spectrum as shown.




Figure 20: "thinning" in the frame-closing symbol

Figure 21 shows an example of normal symbols followed by a frame-closing symbol, after pilot insertion. This illustrates the higher density of pilots in the frame-closing symbol and the unmodulated cells spread throughout the spectrum.




Figure 21: pilots and active cells in normal and frame-closing symbols

In combinations of FFT size, pilot pattern and guard interval for which frequency-only interpolation is expected to be used, there is no frame-closing symbol and the last symbol of the frame is a normal symbol.

6.1.3 L1 signalling

The L1 signalling is divided into L1-pre signalling and L1-post signalling. The modulation and code rate of the L1-pre signalling is BPSK 1/2 and the number of signalling bits in L1-pre is constant. Thus, the L1-pre signalling always occupies 1 840 cells.

The number of signalling bits in L1-post signalling depends on the number of PLPs, number of auxiliary streams, the use of FEFs and the possible future use of TFS. The number of cells occupied by the L1-post signalling depends on the number of signalling bits and the used modulation. The L1-post signalling can be modulated using BPSK, QPSK, 16-QAM or 64-QAM. The FEC code rate is always 1/2. The modulation for the L1-post signalling is chosen so that the L1 signalling is always more robust than any PLP.

The L1 signalling is always carried in the P2 symbol(s). Thus, the number of PLPs that can be carried by the T2 system is limited by the modulation used for L1-post signalling and the number of signalling bits. The maximum number of PLPs that can be carried by a T2 system, assuming that neither FEFs nor auxiliary streams are used, is given in table 13.

Table 13: Maximum number of PLPs for different configurations

| | | BPSK | QPSK | 16-QAM | 64-QAM |
|---------------------|--------|------|------|--------|--------|
| Repetition not used | 32 K | 61 | 127 | 255 | 255 |
| | others | 19 | 43 | 87 | 132 |
| Repetition used | 32 K | 45 | 94 | 190 | 255 |
| | others | 14 | 31 | 64 | 97 |

The L1 signalling capacity may need to be limited further under some circumstances if bias in the L1 signalling causes unacceptably large peaks in the time domain signal. This is discussed further in clause 8.14.

The L1-pre signalling and L1-post signalling are interleaved and encoded separately from each other and from the PLPs as presented in clause 9.2.2. The signalling is mapped to the P2 symbols so that there is an equal amount of L1-pre and L1-post signalling in each P2 symbol as depicted in figure 22. The LDPC code is punctured and shortened based on the amount of L1-post data, so that the L1-post cells are always equally divided into the P2 symbols.




Figure 22: Mapping of L1 data into P2 symbols. The numbers are the index into the L1-pre or L1-post data field

6.1.4 Physical-layer pipes

The physical-layer pipes (PLPs) carry the service data and signalling data for other layers than L1 (physical layer). The L1 signalling is carried separately from the PLPs as described above. Every PLP can have its own modulation, coding and time-interleaving length.

There are three types of PLPs: common PLPs, data PLPs of type 1 and data PLPs of type 2.

The common PLPs are intended to carry data that is shared between several services, such as signalling, programme guides or shared service components. Several PLPs can form a group of PLPs, which share one common PLP. The memory requirements in the receiver, $2^{15} + 2^{19}$ cells for time interleaving, are set so that one common PLP and one data PLP can be received simultaneously. According to the current specification, the common PLPs are carried after the L1-post signalling (figure 18 and figure 19).

The data PLPs are intended to carry the actual T2 services. The difference between the two types of PLPs lies within the possibilities of sub-slicing and power saving.

The data PLPs of type 1 are intended for services that require good power saving and should be carried after the common PLPs (figure 18 and figure 19). Type 1 data PLPs are carried in one sub-slice per T2-frame in the T2-frames in which they occur. Time diversity can also be achieved for type 1 data PLPs by applying time interleaving over several T2-frames.

A special case of the T2-frame structure is input mode A, when the T2 system carries one big TS transmitted in one PLP. In this cases the one PLP is a data PLP of type 1 with one or (more likely) several time interleaving blocks per T2-frame.

Type 2 data PLPs are carried in multiple sub-slices per T2-frame in the T2-frames in which they occur. The number of sub-slices, which can range from 2 to 6480, is the same for all type 2 data PLPs within a super frame. Only the structure of the first sub-slice interval (SUB_SLICE_INTERVAL active cells) of a T2-frame is signalled directly. This interval contains one sub-slice of each type 2 data PLP that occurs in the particular T2-frame. The structure of the first sub-slice interval is repeated as many times as there are sub-slices per frame (SUB_SLICES_PER_FRAME). All sub-slices of a particular PLP are of equal length within a T2-frame and successive sub-slices for a given PLP are separated by SUB_SLICE_INTERVAL cells.

Figure 23 depicts scheduled data of the data PLPs of a T2-frame; the horizontal axis is the one-dimensional cell address scheme, which is explained below. There are M_1 data PLPs of type 1 and M_2 data PLPs of type 2. The type 1 data PLPs are scheduled to occur first and then the type 2 data PLPs. The first sub-slice interval contains one sub-slice of all M_2 type 2 data PLPs. The cell address of the first OFDM cell carrying the first sub-slice of a PLP is signalled in the PLP-loop of the dynamic L1-post signalling. The SUB_SLICE_INTERVAL given in cells is also signalled in the dynamic part of the L1-post. The cell address of the first common PLP is '0' (not depicted in figure 23).




Figure 23: Scheduled data PLPs for T2-frame

The addressing of the OFDM cells carrying common and data PLPs is depicted in figure 24. The first available cell after the L1-post signalling in the first P2 symbol has the cell address '0'. The cell address increases to the end of the P2 symbol and continues from the first available cell in the second P2 symbol, etc. The cell address includes only the cells available for data transmission, not the pilot or tone reservation cells.

A data PLP does not have to be mapped into every T2-frame, but instead can skip frames. The interval between successive frames to which a PLP is mapped is specified by the L1 signalling parameter FRAME_INTERVAL. This mechanism enables increased power saving and is especially useful for low-bit-rate PLPs (e.g. of type 1). The first T2 frame where the data PLP appears is determined by FIRST_FRAME_IDX. In order to have the same mapping of the data PLPs to the T2 frames in each super frame, [i.1] specifies that N_{T2} be divisible by FRAME_INTERVAL for every data PLP. Also, N_{T2} should be chosen according to P_1 so that for every data PLP there is an integer number of Interleaving Frames per super-frame.




Figure 24: Addressing of the OFDM cells for common PLPs and data PLPs
The numbers (cell addresses) are for illustration only

6.1.5 Dummy cells

The total number of active cells in a T2-frame available for PLPs depends in a complicated manner on many of the T2 system parameters and is not in general a multiple of N_{cells} , the number of cells in a FEC block. Since there have to be a whole number of FEC blocks in a T2-Frame (except for multi-frame interleaving), there will almost always be cells within the frame which do not carry cells from any FEC block.

According to the specification [i.1] these will all come at the end of the T2-frame and, unless auxiliary streams are used (see clause 6.1.6), they will be Dummy Cells.

Dummy cells are modulated using the same PRBS as defined for Baseband Frame Scrambling, using BPSK modulation with the same power as an ordinary PLP constellation. This ensures that the symbols containing the dummy cells have essentially the same power as the other OFDM symbols, to avoid problems in transmitters.

The number of symbols in a T2-frame can be adjusted in order to minimise the number of such dummy cells - see clause 5.5.

In a simple multiplex, all PLPs will use the same modulation and will therefore all have the same value for N_{cells} ; furthermore the total number of FEC blocks per T2-frame will be constant. In this case the total number of cells for the PLPs will be the same in each frame and the number of dummy cells will also be the same.

It might be possible to design a scheduler and T2-aware statistical multiplexer able to deal with different PLPs having different modulation and code rates. In this case the FEC blocks would be of different sizes and the total number of cells allocated to PLPs might vary from frame to frame. Consequently, the number of dummy cells would also change from frame to frame.

NOTE: In these circumstances the SUB_SLICE_INTERVAL would also vary dynamically.

Normally, dummy cells would only be present in the last OFDM symbol of the T2 frame. However, in the case of T2-aware statistical multiplexing the dummy cells might sometimes extend over more than one symbol in some frames.

6.1.6 Auxiliary streams

The dummy cells do not carry useful information and therefore represent wasted capacity. In the future, applications may be defined which make use of these otherwise wasted cells, and the concept of an Auxiliary Stream is defined for this purpose. One application, Transmitter Signature [i.22], has already been defined.

Applications for Auxiliary Streams could include opportunistic data, low-data-rate audio streams, or information used by network operators for communicating with professional monitoring receivers used in relay transmitters.

In order not to upset the AGC of receivers, the specification [i.1] requires that these cells be modulated with the same mean power as an ordinary PLP constellation.

The original intention was that auxiliary streams should use up the remaining cells which are not enough to accommodate a whole FEC block. However, the specification allows the auxiliary streams (and dummy cells) to occupy more than one FEC-block's worth of cells if desired. This might allow future applications to introduce services of significant bit rate.

Any future use of Auxiliary streams would need to define the coding and modulation schemes, if applicable. If LDPC were used, this would include defining a shortening and puncturing scheme. Bits are provided in the L1 signalling to indicate the auxiliary stream type, and to carry both configurable information, such as code rate and modulation, and also dynamic information, which would be necessary to signal which cells of each frame are used for a particular Auxiliary Stream. This signalling appears in a loop, so that if there are no auxiliary streams the loop is absent and overhead is minimal.

A first use of auxiliary streams has been identified with version 1.2.1 of [i.1], for the identification of transmitters. More details will be given in future editions of the present document.

6.1.7 FEFs

6.1.7.1 FEF Basics

Future-Extension Frames (FEF) provide a method of expanding the DVB-T2 standard in the future by means and ways not known at the time of writing the original standard. This is done by inserting new additional periods known as FEF-parts between the normal T2-frames. Thus the FEF parts and T2-frames are separated in time as depicted in figure 25. The use of future-extension frames is optional, and the FEF parts are inserted only if needed.

NOTE: The term "FEF part", as opposed to "FEF", is used to refer to a contiguous period of time, beginning with a P1 preamble and containing signals which should be ignored by single-profile DVB-T2 receivers. One FEF part may contain one or more frames of a future system; or the new system need not even have frames as such at all.




Figure 25: The super-frame, including T2-frames and FEF parts

A future-extension frame may carry data in way unknown to a DVB-T2 receiver based on the current standard version. This type of DVB-T2 receiver is not expected to decode the future-extension frames, but it is expected to continue receiving the normal T2-frames without being disturbed by the FEF parts. All receivers, present and future, are expected to detect the existence of FEF parts.

One or several FEF parts can be inserted in the super frame and a super frame can thus carry two different types of structure: T2-frames and FEF parts, see figure 25. The insertion scheme is configurable and is repeated in each superframe. The location of the FEF parts is described by the L1-post signalling field FEF_INTERVAL, which is the number of T2-frames at the beginning of a super frame, before the beginning of the first FEF part. The same field also describes the number of T2-frames between two FEF parts. The length of the FEF part is given by the FEF_LENGTH field of the L1 signalling. This field describes the time between two DVB-T2 frames preceding and following a FEF part as the number of elementary time periods T , i.e. samples in the receiver (see clause 9.5 of [1.1]) It is measured from the start of the P1 of the FEF part to the start of the P1 of the following T2 frame. The number of T2-frames in a super frame is given by parameter NUM_T2_FRAMES in L1-pre signalling. When FEFs are used, the super frame ends with a FEF part.

The total super frame duration T_{SF} is determined by:

$$T_{SF} = \text{NUM_T2_FRAMES} \times T_F + N_{FEF} \times T_{FEF},$$

where T_F is duration of the T2-frame, N_{FEF} is the number of FEF part in a super frame and T_{FEF} is the duration of the FEF part signalled by FEF_LENGTH.

N_{FEF} can be derived as:

$$N_{FEF} = \text{NUM_T2_FRAMES} / \text{FEF_INTERVAL}.$$

A FEF part begins with a P1 symbol that can be detected by all DVB-T2 receivers. The maximum length of a FEF part is 250 ms and is always followed by a normal T2-frame.

All other parts of the future-extension frames will be defined in future extensions of the DVB-T2 standard or elsewhere.

6.1.7.2 FEF Signalling

The use of FEFs is signalled both in P1 and in P2 symbols. The P1-symbol has two signalling fields S1 and S2. The S1 field is used to distinguish the preamble format and, hence, the frame type. If FEF parts are inserted in the T2-system, the P1 in the FEF parts will have a different S1 value than the in the T2 P1, for which S1 is either 000 or 001. This means that for any new system introduced in the FEF parts one of the reserved bit combinations should be taken into use (see table 14).

Table 14: S1 Field

| S1 | Preamble Format / P2 Type | Description |
|---------------------------------|---------------------------|--|
| 000 | T2_SISO | The preamble is a T2 preamble and the P2 part is transmitted in its SISO format |
| 001 | T2_MISO | The preamble is a T2 preamble and the P2 part is transmitted in its MISO format |
| 010 | Non-T2 preamble | |
| 011 100 101 110 111 | Reserved | These combinations may be used for future systems, including a system containing both T2-frames and FEF parts, as well as future systems not defined in the present document |

The S2 has two fields, of which field 2 (the LSB) is relevant for FEFs. If FEF parts are inserted in the superframe, this one-bit field is set both in the P1 of each T2-frame and in also in the P1 of each FEF part (see table 15).

Table 15: S2 field 2

| S1 | S2 | Meaning | Description |
|-----|------|-----------|--|
| XXX | XXX0 | Not mixed | All preambles in the current transmission are of the same type as this preamble. |
| XXX | XXX1 | Mixed | Preambles of different types are transmitted. |

The meaning of this bit is to signal to all receivers in the initial scan that two type of frames exist in the super frame. This speeds up the scan as all receivers immediately know if there is a need to wait for the other type of P1, either T2 P1 or FEF P1.

If FEFs are used, the following fields appear in the L1-post signalling:

- **FEF_TYPE:** This 4-bit field indicates the type of the associated FEF part. The FEF types are signalled according to table 27 of [i.1].
- **FEF_LENGTH:** This 22-bit field indicates the length of the associated FEF part as the number of elementary periods T (see clause 9.5 of [i.1]), from the start of the P1 symbol of the FEF part to the start of the P1 symbol of the next T2-frame.
- **FEF_INTERVAL:** This 8-bit field indicates the number of T2-frames between two FEF parts and in case of the first FEF part from the beginning of a super frame. A T2 superframe which contains both FEF parts and T2-frames always begins with a T2-frame.

The FEF_TYPE field is intended for further identification of the FEF used in the system. Currently, only one FEF type is defined in EN 302 755 [i.1], and all other bit combinations are reserved for future use. The use of the other two fields is explained above.

The one combination that is currently specified is S1=010, S2 field 1=000. This signals that the preamble corresponds to a FEF part intended for professional use. This could include measurement of the noise or interference level in a channel, or engineering trials of possible new FEF types which would eventually be allocated their own S1 and S2 code. When this code is used, the exact contents of the FEF part are not specified and may be used as required by a broadcaster or network operator. It is assumed that the reception equipment making use of the FEF would know what the FEF part contained, either by using signalling contained in the FEF part itself, or by being configured directly by the operator.

The Transmitter Signature standard [i.22] includes an method that uses FEFs for identifying transmitters in a network and measuring their relative timing.

6.1.7.3 Possible use cases

As the FEF is intended for future extensions of the system, it is impossible to describe the use cases in detail, but some possibilities will be suggested here.

The FEF-source can either be the T2-network or a different network. In any case the T2-network is responsible for inserting the P1 symbol at the beginning of the FEF part. This is important in order that all the T2 receivers in the whole T2 service area can receive the FEF P1. Otherwise the T2 receivers would not be able detect the FEF correctly. The source of the actual content in the FEF may be another network, which has different service area.

FEFs inserted by the same DVB-T2 network would probably be used for some future extension of the T2 standard itself, e.g. a new modulation scheme, a future hand-held or mobile system, MIMO, etc.

For FEFs inserted by another network, one possible use case could be a mobile network sharing the same spectrum, but having different sites and infrastructure. Another possible case would be to use the same frequency for a "T2 return channel". In this case the T2 network would leave an empty FEF, which would be used for the upstream traffic by the T2 receivers. Many other possible applications are possible.

6.2 Logical frame structure

6.2.1 Packets

PLPs may carry packets, including MPEG Transport Stream packets or other types of packets used by generic streams. They may also carry continuous, non-packetised streams. The BBFRAME header (see clause 6.2.2.1) provides a mechanism for packets to be reconstructed at the receiver, but the framing of the DVB-T2 signal itself is generally independent of any packet structure which the input streams may have.

6.2.2 BBFRAMES, FECFrames and FEC blocks

There is a 1:1:1 correspondence between these three objects, but there is a subtle difference between them and all three terms are used in [i.1].

6.2.2.1 BBFRAMES.

BBFRAMES are the basic unit in the logical framing structure of DVB-T2: allocation and scheduling are performed in whole numbers of BBFRAMES. Where packetised streams are being carried, the packets may be mapped to BBFRAMES either synchronously or asynchronously, i.e. each BBFRAME may contain a whole number of packets, or packets may be fragmented across two BBFRAMES. BBFRAMES contain a header, including the packet length and position of the first packet, which allows the original packets to be reconstructed at a receiver. BBFRAMES may contain padding, in case insufficient data is available for a whole BBFRAME (see clause 8.7), or it is desired not to fragment packets. The padding field may also be used for in-band L1 signalling. The total size of each BBFRAME, including any padding and/or in-band signalling, is constant for a given PLP and depends on the LDPC code rate (see clause 9.2.1.3) and whether short or long FEC blocks are used (see clause 6.2.2.2).

6.2.2.2 FECFRAMES

The baseband frame, complete with header, is treated as an information word to which BCH and LDPC coding are applied. The resulting codeword always contains either 64 800 bits or 16 200 bits and is known as a FECFRAME. The two different lengths correspond to the choice of long or short FEC blocks respectively. Short FEC blocks allow a finer granularity of bit-rate but incur a greater overhead and slightly worse performance than long blocks.

6.2.2.3 FEC Blocks

The bits of a FECFRAME are interleaved and arranged into words, which are then mapped onto constellation points, or OFDM cells; the set of constellation points corresponding to one FECFRAME is known as a FEC block. The number of cells in a FEC block therefore depends on both the QAM modulation and the choice of long or short blocks for a particular PLP.

6.2.3 Interleaving Frames

The Interleaving Frame is an important concept in DVB-T2, but one that causes a certain amount of confusion, and a proper understanding is essential to the understanding of many of the other parts of the specification.

To begin with, the name "Interleaving Frame" is somewhat misleading. Although the need for this concept arises from the availability of three options for time interleaving (see clause 6.5 of [i.1]), the Interleaving Frame itself also plays a role in scheduling, definition of the L1-signalling fields, insertion of in-band L1-signalling, and frame building; it is also involved in the definition of the T2-MI interface.

The Interleaving Frame is the framing structure within which allocation is performed for a particular PLP. The scheduler will allocate a dynamically variable whole number of BBFRAMES to each Interleaving Frame. An Interleaving Frame can be thought of as a period of time equal to a whole number of T2-Frames, and furthermore there will be a whole number of Interleaving Frames in a superframe.

There are two reasons why Interleaving Frames are different from T2-frames. Firstly, a PLP may use time interleaving over multiple T2-frames. It is therefore not meaningful to speak of a BBFRAME being allocated to a particular T2 frame, since each BBFRAME would be spread out over all the P_1 T2-frames in the interleaving period. Furthermore, the number of cells belonging to a given PLP in a given T2 frame depends on the total number of BBFRAMES allocated to the Interleaving Frame.

Secondly, a particular PLP need not be mapped to every T2-frame of a superframe, but instead may appear in one T2-frame in every I_{JUMP} T2-frames. The duration of the Interleaving Frame in this case corresponds to the P_1 T2-frames to which the PLP is mapped plus the $P_1 \times (I_{\text{JUMP}} - 1)$ frames which are skipped: a total of $P_1 \times I_{\text{JUMP}}$ T2-frames.

Alternatively, it is also possible to use multiple TI blocks in each T2-frame for a given PLP, in order to increase the maximum data rate for a given frame size. In this case, allocation is done over a T2-frame period, and therefore one Interleaving Frame is mapped directly to one T2-frame.

Note that although multi-frame interleaving may not be used together with multiple TI-blocks per Interleaving Frame, [i.1] does not prohibit the use of multiple TI-blocks per Interleaving Frame together with a value of $I_{\text{JUMP}} > 1$. It is difficult to see any practical application for this and it is recommended that this combination will not be used.

Since BBFRAMES are allocated to Interleaving Frames, the in-band dynamic L1 signalling appears only once per Interleaving Frame, in the first BBFRAME of the Interleaving Frame. The concept of "the first BBFRAME of a T2 frame" would be meaningless. The implications of this are discussed in clause 6.2.5.

6.2.4 TI blocks

A TI-block is the set of OFDM-cells within which Time Interleaving is performed; there is no interleaving between TI-blocks. Each TI-block contains a dynamically variable whole number of FEC blocks. An Interleaving Frame in turn contains an integer number of TI-blocks, this number being statically configurable. The total number of BBFRAMES, and therefore FEC Blocks, allocated to an Interleaving Frame need not be a multiple of the number of TI-blocks, and as a result the TI-blocks within an Interleaving Frame may contain slightly different numbers of FEC blocks. The number of FEC blocks in each TI-block of the Interleaving Frame will only differ by one, with the smaller numbers of FEC blocks occurring in the first TI-blocks of the Interleaving Frame. This division of the FEC-blocks of the Interleaving Frame into the individual TI-blocks is described by the equations in clause 6.5.2 of [i.1].

EXAMPLE: If there are three TI blocks in an Interleaving Frame and 202 FEC blocks in the Interleaving Frame, the TI-blocks will contain 67, 67 and 68 FEC blocks respectively.

Sending the smaller TI-blocks first is expected to make the design of the time de-interleaver easier (see clause 10.4.3.2).

TI-blocks are also used by the Cell Interleaver (see clause 9.2.4). The shift value $L(r)$ applied to the basic permutation function changes from one FEC block to the next, but the FEC-block index r is reset to zero at the beginning of each TI-block.

6.2.5 Mapping to physical frame structure

Each Interleaving Frame is mapped to one or more (P_1) T2-frames. Following time interleaving, each Interleaving Frame consists of a series of OFDM cell values. These cells are partitioned into equal parts to be carried in each of the corresponding T2-frames; the cells allocated to a particular T2-frame are described as a *slice*. For Type-2 PLPs, each slice may then be further partitioned into sub-slices of equal size, as described in clause 6.1.4.

The number of sub-slices is the same for all of the Type-2 PLPs. The number of sub-slices and the number of T2 frames to which the Interleaving Frame for each PLP is mapped have to be chosen such that the number of cells in a FEC block can be divided equally into slices and sub-slices as appropriate; this is described in more detail in clause 5.8.

Although the sub-slices will be the same size in all of the T2-frames to which one Interleaving Frame is mapped, the sub-slices need not necessarily be allocated in the same locations in each T2-frame.

It can therefore be seen that the PLP-specific part of the dynamic L1-signalling applies as follows:

- BBFRAMES are allocated to Interleaving Frames, so there is one value of PLP_NUM_BLOCKS which applies to the Interleaving Frame.
- The location of the sub-slices can vary from T2-frame to T2-frame and therefore PLP_START (and SUB_SLICE_INTERVAL) apply to a particular T2-frame. The sizes of the sub-slices can be derived from PLP_NUM_BLOCKS.

Further details of the dynamic signalling are given in clauses 5.2.3 and 7.2.3.2 of [i.1]. Since there is only one set of in-band signalling for each Interleaving Frame, a loop is required to give the dynamic sub-slicing parameters for each of the T2-frames to which it is mapped.

6.2.6 The concept of a T2-system

The DVB-T2 specification [i.1] uses the phrase "T2 System" in different senses. In some places it is used in a general way to refer to the DVB-T2 standard and the various elements which make it up.

However, "T2-System" is also used in a specific way to refer to a particular layer-1 concept, which is a generalisation of the concept of a multiplex in DVB-T and other DVB standards.

A T2-system in this specific sense is a set of transmissions with time-synchronised frame structure, using the same physical parameters (e.g. bandwidth, FFT size), carrying the same number and type of PLPs, and using the same physical parameters for each of the Physical-Layer Pipes (PLPs) it carries. The L1 signalling will therefore be identical for all transmissions in a T2-system, except for the cell_id which may differ.

However, DVB-T2 allows for regional content insertion, so that the same T2 System may therefore carry different sets of Transport Streams (and/or Generic Streams) and use different transmission frequencies in different geographical areas.

All of the transmissions making up a T2-system will originate from a single T2-gateway, generating an original T2-MI stream, so that the framing structure and scheduling is determined in one place and is common to all transmissions. The T2-MI will propagate through a distribution network in which some of the content may be replaced. In principle, all of the content of the PLPs might be replaced by the time a signal is transmitted, though in practice there will probably be some content which is the same for all transmissions. Figure 26 gives an example distribution network.

Since all transmissions in a T2-system come from transmitters which are ultimately connected to the same T2-gateway, the term "T2-system" can also be equivalently applied to the system of equipment and distribution networks starting from the input to the T2-gateway through to the outputs of the transmitters. It will be seen that this is consistent with the definition in [i.1].




Figure 26: Transmissions in a T2-system

6.3 Capacity

This clause explains how the total capacity of the T2 system can be calculated. The total capacity depends on a number of different factors:

- FFT size.
- Guard interval.
- Pilot pattern.
- Bandwidth and use of Normal or Extended carrier mode.
- Use of reserved tones (TR).
- Existence of Future-Extension Frames.
- Use of SISO or MISO.
- T2-frame length L_F , in OFDM symbols.

- Modulation and code rate used for PLPs and L1-post signalling.
- Number of PLPs, which affect the amount of L1-post signalling data.

The FFT size, guard interval, pilot pattern, bandwidth and use of SISO or MISO will affect the following parameters:

- Number of data cells in:
 - P2 symbols C_{P2} (table 9).
 - data symbols C_{data} (table 10).
 - in frame closing symbol C_{FC} (table 11), and whether one is used at all.
- Number of P2 symbols N_{P2} (table 9).

6.3.1 Capacity in cells and FEC blocks per frame

The total capacity in cells per frame of the T2 system can be calculated as:

$$C_{tot} = N_{P2} \times C_{P2} + (L_{data} - 1) \times C_{data} + C_{FC}$$

where $C_{FC} = C_{data}$ for cases where no frame-closing symbol is used.

To find the capacity available for transmission of PLPs and auxiliary services, the number of cells needed to carry L1 signalling data should now be subtracted from the number of cells per frame. The number of cells needed to carry L1 signalling data depends on the modulation used for L1-post signalling, the number of PLPs and the use of FEFs. The number of cells for L1 signalling per T2 frame is calculated as:

$$D_{L1} = D_{L1_pre} + D_{L1_post}$$

where $D_{L1_pre} = 1840$ and D_{L1_post} is the number of L1-post cells.

D_{L1_post} is calculated as follows:

The number of configurable L1-post signalling bits $K_{configurable}$ for one RF channel is calculated by:

$$K_{configurable} = 102 + FEF \times 34 + N_{PLP} \times 89 + N_{AUX} \times 32,$$

where $FEF = 1$ if FEF parts are present, otherwise $FEF = 0$. N_{PLP} is the number of PLPs in the system and N_{AUX} is the number of auxiliary streams. The number of dynamic L1-post signalling bits $K_{dynamic}$ is calculated by:

$$K_{dynamic} = 79 + N_{PLP} \times 48 + N_{AUX} \times 48.$$

The total number of L1-post signalling bits $K_{post_ex_pad}$ can then be calculated by:

$$K_{post_ex_pad} = K_{configurable} + X \times K_{dynamic} + 32,$$

where $X = 1$ if repetition is not used or $X = 2$ if repetition is used.

The number of FEC blocks used for the L1-post is:

$$N_{post_FEC_Block} = \left\lceil \frac{K_{post_ex_pad}}{7032} \right\rceil.$$

The number of information bits in each block is:

$$K_{sig} = \left\lceil \frac{K_{post_ex_pad}}{N_{post_FEC_Block}} \right\rceil$$

D_{L1_post} can then be calculated as follows:

$$N_{L1_mult} = \begin{cases} 2 \times \eta_{mod} & \text{if } N_{P2} = 1 \\ N_{P2} \times \eta_{mod} & \text{otherwise} \end{cases}$$

$$N_{punc_temp} = \max \left(N_{L1_mult} - 1, \left\lfloor \frac{6}{5} \times (7032 - K_{sig}) \right\rfloor \right)$$

$$N_{post_temp} = K_{sig} + 9168 - N_{punc_temp}$$

$$D_{L1_post} = \frac{N_{L1_mult}}{\eta_{mod}} \times N_{post_FEC_Block} \times \left\lceil \frac{N_{post_temp}}{N_{L1_mult}} \right\rceil$$

Thus, the available number of cells per frame for PLPs and auxiliary services can be calculated as:

$$D_{PLP} = (N_{P2} \times C_{P2} + (L_{data} - 1) \times C_{data} + C_{FC} - D_{L1}) \text{ cells/frame.}$$

Assuming multi-frame interleaving is not used, there are always a whole number of FEC blocks in a frame. The maximum number of FEC blocks per frame is therefore given by:

$$N_{B_max} = \left\lfloor \frac{D_{PLP}}{N_{cells}} \right\rfloor = \left\lfloor \frac{(N_{P2} \times C_{P2} + (L_{data} - 1) \times C_{data} + C_{FC} - D_{L1})}{N_{cells}} \right\rfloor$$

where N_{cells} is the number of cells per FEC block for the QAM modulation scheme in use. If multi-frame interleaving is used, the calculation would need to be carried out on an Interleaving-Frame basis.

In Input Mode B, this total number of blocks will need to be shared between the PLPs, either statically or dynamically. A user may choose to limit the maximum number of FEC blocks per frame for some or all PLPs to some lower value. In particular this may be done in order to allow fewer TI-blocks to be used, increasing the interleaving depth.

As described in clause 5.5, the number of OFDM symbols per T2-frame should be chosen to maximize N_{B_max}/T_F , the number of FEC blocks per second; this does not necessarily imply choosing the maximum possible value of frame length. The effect on interleaving depth should also be considered.

6.3.2 Total bit-rates of the T2 system

When a Transport Stream is carried by a data PLP together with a common PLP, the maximum useful bit rate which can be carried, assuming FEFs are not used, is given by:

$$R_{NO_FEFs} = \frac{1}{T_F} \left(\frac{188}{188 + B_{MA_data}} \times \frac{N_{B_max_data} \times (K_{bch_data} - 80) - N_{IB_data}}{P_{I_data} \times I_{jump_data}} + \frac{188}{188 + B_{MA_cplp}} \times \frac{N_{B_max_cplp} \times (K_{bch_cplp} - 80) - N_{IB_cplp}}{P_{I_cplp} \times I_{jump_cplp}} \right)$$

where:

| | |
|---|--|
| R_{NO_FEFS} | is the useful output bit-rate. |
| $N_{B_max_data}$, $N_{B_max_cplp}$ | are the values of PLP_NUM_BLOCKS_MAX for the data PLP being decoded and its associated common PLP respectively. |
| K_{bch_data} , K_{bch_cplp} | are the applicable values of K_{BCH} . |
| N_{IB_data} , N_{IB_cplp} | are the number of bits used for in-band signalling (see clause 8.10.6). |
| P_{I_data} , P_{I_cplp} | are the number of T2-frames to which each Interleaving Frame is mapped. |
| I_{jump_data} , I_{jump_cplp} | are the values of FRAME_INTERVAL. |
| T_F | is the duration of the T2-frame. |
| B_{MA_data} , B_{MA_cplp} | are values which depend on the Mode-Adaptation processing (see clause 8.5) as follows. The relevant value of B_{MA} depends on the number of bytes inserted or removed during Mode Adaptation. In HEM, this value is 0 if DNP is used or -1 if DNP is not used. For NM, the values from table 16 for each option chosen are added to obtain the overall value of B_{MA} , and this will therefore have a maximum value of 4. So for example if ISSY (Short) is used with DNP, the value of B_{MA} will be 3. If neither ISSY nor DNP are used, B_{MA} will be 0. |

Table 16: Number of bytes inserted during Mode Adaptation for TS in Normal Mode

| Option | Increment to obtain B_{MA} |
|--------------|------------------------------|
| ISSY (Short) | +2 |
| ISSY (Long) | +3 |
| DNP | +1 |

NOTE 1: The calculation assumes that $N_{B_max_data}$ FEC blocks can be allocated to the data PLP in the same Interleaving Frame as $N_{B_max_cplp}$ FEC blocks are allocated to the common PLP.

NOTE 2: If no common PLP is used, the formula should be used with $N_{B_max_cplp}=0$.

NOTE 3: For input mode A, assuming no auxiliary streams and the minimum of dummy cells, $N_{B_max_data}$ is equal to N_{B_max} calculated in clause 6.3.1.

NOTE 4: For input mode B in the case where all PLPs use the same modulation and code-rate, the calculation based on the value of N_{B_max} from the previous clause will give the total bit-rate for all PLPs. Note that N_{B_max} may be smaller because there will be more L1 signalling bits in a multiple-PLP configuration.

NOTE 5: If FEFS are used, the bit rate calculated above should be adjusted by the proportion of time that the T2-frames are on air during the super frame. This is given by $(N_{T2} \times T_F) / T_{SF}$, where N_{T2} is the number of T2-frames during one super frame and T_{SF} is the duration of the super frame, i.e. the adjusted bit-rate R_{FEFS} is given by:

$$R_{FEFS} = R_{NO_FEFS} \frac{N_{T2} T_F}{T_{SF}}$$

NOTE 6: The formula gives the useful bit-rate which the system can carry. However, the output Transport Stream rate may be higher than this owing to the re-insertion of deleted null packets; this rate can be determined at the receiver as described in clause 10.6.1.3 but may not exceed 72 Mbit/s.

6.3.3 Examples of bit-rates

In an 8 MHz channel, the highest data rate is achieved with the 32 K, GI=1/128, extended-carrier mode and no tone reservation, for which pilot-pattern PP7 is always used. The first few columns of table 17 give the maximum achievable bit rate for each combination of constellation and code-rate, together with the corresponding frame length (L_F) and the total number of FEC blocks per frame. The frame length giving this maximum bitrate varies with constellation as a result of the varying number of dummy cells (see clause 5.5). In practice it is generally recommended that a slightly shorter frame length be used, giving a slightly lower bit-rate but longer time interleaving (see clause 5.5). These recommended values are given in the rightmost columns of table 17. The bit-rates are plotted in figure 27, which shows clearly that the reduction in bit-rate for the recommended configuration is insignificant.

Table 17: Maximum bit-rate and recommended configurations for 8 MHz, 32 K 1/128, PP7

| Modulation | Code rate | Absolute maximum bit-rate | | | Recommended configuration | | |
|------------|-----------|---------------------------|--------------------|----------------------|---------------------------|--------------------|----------------------|
| | | Bitrate Mbit/s | Frame length L_F | FEC blocks per frame | Bitrate Mbit/s | Frame length L_F | FEC blocks per frame |
| QPSK | 1/2 | 7,49255 | 62 | 52 | 7,4442731 | 60 | 50 |
| | 3/5 | 9,003747 | | | 8,9457325 | | |
| | 2/3 | 10,01867 | | | 9,9541201 | | |
| | 3/4 | 11,27054 | | | 11,197922 | | |
| | 4/5 | 12,02614 | | | 11,948651 | | |
| | 5/6 | 12,53733 | | | 12,456553 | | |
| 16-QAM | 1/2 | 15,03743 | 60 | 101 | 15,037432 | 60 | 101 |
| | 3/5 | 18,07038 | | | 18,07038 | | |
| | 2/3 | 20,10732 | | | 20,107323 | | |
| | 3/4 | 22,6198 | | | 22,619802 | | |
| | 4/5 | 24,13628 | | | 24,136276 | | |
| | 5/6 | 25,16224 | | | 25,162236 | | |
| 64-QAM | 1/2 | 22,51994 | 46 | 116 | 22,481705 | 60 | 151 |
| | 3/5 | 27,06206 | | | 27,016112 | | |
| | 2/3 | 30,11257 | | | 30,061443 | | |
| | 3/4 | 33,87524 | | | 33,817724 | | |
| | 4/5 | 36,1463 | | | 36,084927 | | |
| | 5/6 | 37,68277 | | | 37,618789 | | |
| 256-QAM | 1/2 | 30,08728 | 68 | 229 | 30,074863 | 60 | 202 |
| | 3/5 | 36,15568 | | | 36,140759 | | |
| | 2/3 | 40,23124 | | | 40,214645 | | |
| | 3/4 | 45,25828 | | | 45,239604 | | |
| | 4/5 | 48,29248 | | | 48,272552 | | |
| | 5/6 | 50,34524 | | | 50,324472 | | |




Figure 27: Maximum bitrate and bitrate for recommended configuration with 8 MHz bandwidth and 32 K PP7

6.4 Overview of interleaving

There are four different interleavers in the DVB-T2 specification, occurring in the following order in the modulation chain:

- Bit Interleaver.
- Cell Interleaver.
- Time Interleaver.
- Frequency Interleaver.

The Bit Interleaver interleaves the code-bits within an LDPC codeword in order to avoid undesirable interactions between the bits carried by the same cell and the structure present in the LDPC code. This is described in clause 9.2.1.3.2.

The Cell Interleaver (see clause 9.2.4) applies a pseudo-random permutation to the cells within a FEC block, with the permutation varying from one FEC block to the next. This disrupts the structured nature of the time interleaver and prevents interactions with the structure of the LDPC code. For example, a burst of erasures caused by a channel will appear in random positions in the FEC codeword instead of being regularly spaced, which might interact with the regular structure of the code. The cell interleaver also plays a role when Rotated Constellations are used in increasing the separation between cells carrying information from the same original constellation.

The Time Interleaver (clause 9.2.5) provides the bulk of the interleaving, spreading the cells of each FEC block over many symbols and potentially over several T2-frames. This offers protection against impulsive interference as well as time-varying channels.

Finally, the Frequency Interleaver (clause 9.3.2), which is a pseudo-random block interleaver operating on OFDM symbols, has several purposes. First, it mixes up cells from different PLPs carried in the same OFDM symbol, whereas the other interleavers operate on a per-PLP basis. Secondly, it applies a pseudo-random permutation to the output of the time interleaver, thus breaking up the structured nature of the time interleaver. This is necessary because the transmission channel can cause regular bursts and patterns of errors across the symbol (i.e. across frequency), and these could interact adversely with the structure of the time interleaver without the frequency interleaver. Thirdly, the interleaver alternates between two different permutations, which helps to increase the number of OFDM carriers to which each PLP is mapped, as described in clause 5.8.1.

7 Modulator Interface (T2-MI) and network topology

The DVB-T2 Modulator-Interface Specification (T2-MI) is defined by TS 102 773 [i.20]. An outline of a typical network topology based on the use of the T2-MI specification is given below.

7.1 Overview

As described in clause 4.4.2, it is assumed that networks, particularly SFNs, will make use of a central network functionality called the *T2 Gateway*, which accepts one or more TSs (and/or GSE streams) at the input and generates, based on the input streams, a sequence of T2-MI packets at the T2-MI interface (Interface B of the reference architecture of figure 1). The T2-Gateway will include at least the functionality of the *Basic T2-Gateway* (SS2), but might also include some earlier processing stages.

The input to the Basic T2-Gateway sub-system (SS2) is one or more streams. This input is identical to the input to the DVB-T2 system (with or without the splitting-and-merging extension) as specified in [i.1]. The Basic T2-Gateway can therefore be seen as a centrally-positioned part of the T2 modulator (in the sense of [i.1]) functionality. In the case of MPEG-2 TS these streams will each have a constant bit rate. Each stream corresponds to a particular data PLP of the modulator sub-system. Parts of an MPEG-2 TS that are identical to other input TSs may be transferred to the common PLP (see clause 8.3). Each bit-stream typically contains either MPEG-2 TS or GSE packets. Correspondingly, any of these streams can be output from the DVB-T2 demodulator sub-system (SS4) at the receiver (after merging with Common PLP, where applicable).

For the MPEG-2 TS case the input TS packets of a PLP constitute a syntactically-correct MPEG-2 TS, containing typically one service (but potentially several). What is actually transmitted in a data PLP may however constitute a partial MPEG-2 TS, when the TS is split and sent also in a common PLP.

Typically the common PLP is used for layer-2 signalling, such as PSI/SI information, but it could also be used for other information or even services. When there are several groups of PLPs in a DVB-T2 signal, each such group may have one common PLP. Note that for MPEG-2 TS some PSI/SI is always sent in the data PLP.

7.2 Statistical multiplexing and service splitting

We consider the case where a number of video services are statistically multiplexed within the limits of a constant total capacity. This total capacity could be the total capacity of the DVB-T2 signal or a subset of this. In the latter case the remaining capacity would be used for some other purpose.

Depending on the way the statistical multiplexing is performed there may or may not be a need to split a big TS, containing many statistically multiplexed services, into several small syntactically correct TSs (typically with one TS per service) as described in clause 8.2. This splitting takes place *before* the DVB-T2 system and is not to be confused with the second splitting, which may take place later in the chain as part of the DVB-T2 system extension used for splitting an input TS into what is sent in a data PLP and a common PLP (see clause 8.3). For statistical multiplexing it should be noted that since the TS packets sent in the common PLP appear in all input TSs (co-timed) of a group of PLPs the sum of all input TS bit rates may actually exceed the sum of all TS bit rates transmitted in the respective data PLPs and the common PLP of the group.

The statistical multiplexing is typically performed at a central point in the network, whereas the DVB-T2 modulation naturally takes place in a distributed way, often over an entire country. This raises the question how the set of centrally generated streams can be transported to the physical modulators in a way that will ensure generation of identical RF signals at each transmitter site, which are transmitted synchronously. These are fundamental requirements to allow SFN operation. It is intended that this functionality will be achieved using a central T2 Gateway to generate a T2-MI stream which is distributed to the modulators at the transmitting sites.

7.3 The T2 gateway and T2-Modulator Interface (T2-MI)

The T2-Gateway accepts one or more TSs (and/or GSE streams) at the input and generates, based on the input streams, a sequence of T2-MI packets at the T2-MI interface (Interface B of clause 4.4.2) at the output.

7.3.1 The Basic T2-Gateway

The reference model of clause 4.4.2 defines the Basic T2-Gateway. From the point of view of this reference model any required splitting, PSI/SI handling apart from that related to the use of the common PLP, and PCR restamping of MPEG-2 TSs are assumed to have been done before the input to the Basic T2-Gateway. The Basic T2-Gateway then performs any required splitting of incoming TSs into data PLP and common PLP and performs a following pre-analysis of the first stages of the DVB-T2 modulation process, which enables it to create Baseband frames, signalling and SFN synchronisation information, all encapsulated into the sequence of T2-MI packets.

7.3.2 Practical T2-Gateways

A practical T2-Gateway might correspond exactly to the Basic T2-Gateway described in clause 7.3.1. However, it might be convenient for real implementations to include additional functionality. Possible examples are:

- Generation of one transport stream per data PLP (see clause 8.2). If a multiplex is generated by a non-T2-aware statistical multiplexer, this could be split into N conceptual TSs by the T2-gateway.
- Generation and insertion of the T2-MIP (see clause 8.13). Since this gives timing information for the current superframe, it might be difficult to generate it before the T2 gateway since the framing information might not be readily available upstream of the T2-Gateway.

7.4 Distribution of T2-MI packets to transmitters

The stream of T2-MI packets being output from the T2 gateway has to be encapsulated into appropriate lower protocol layers. The T2-MI specification [i.20] describes a way to carry the packets in an MPEG Transport Stream, which can in turn be carried over an ASI circuit as for DVB-T. The T2-MI specification also provides a method to carry the MPEG-TS using IP-based protocols already defined for IP-TV [i.19]. This enables distribution over IP networks and allows for corrections of the imperfections of such networks such as loss of packets and re-ordering of packets. Such functionality could therefore allow the splitting of a high-bit-rate stream of T2-MI packets into several lower-bit-rate streams, which are then recombined at the transmitter site. This approach would simplify satellite distribution of very-high-bit-rate T2-MI streams. The ability to correct the reordering or loss of packets ensures that the same original sequence of T2-MI frames is input to all DVB-T2 modulators. Based on the instructions within the T2-MI packets each modulator is able to construct identical sequences of DVB-T2 physical frames and transmit them synchronously in an SFN.

7.5 Over-air distribution

Where a master transmitter acts as the distribution mechanism for a separate SFN on a different frequency (see figure 28), all the information required to generate identical co-timed SFN signals needs to be carried over air.




Figure 28: Use of Main Transmitter to feed SFN Relays

This topology has been successfully deployed in DVB-T/H Networks in trial and commercial networks, and is proposed for use within the UK's post Digital Switchover networks, as it represents a low cost solution to service distribution.

In DVB-T, SFN synchronisation information is contained within Transport Streams as Megaframe Initialisation Packets (MIPs), which may then be carried over air and are ignored by consumer receivers.

This approach is retained in T2, where the necessary PLP construction and timing information is carried in dedicated TS packets, with a new Synchronisation_Id value, to differentiate them from those used in DVB-T. T2 MIPs contain all of the information required by each modulator to ensure that it builds the broadcast stream in the same way, and that it outputs the stream at the appropriate time.

In the use case discussed here the first of these issues is not a problem, since the T2 signal has already been constructed at the Master Station, and the relays merely need the timing information to ensure time synchronisation.

However, there is a special case where some of the relays are line fed from the Master Station, whilst the remainder are fed over air.

In this situation the Master Station would need to build its T2 Frame in exactly the same way as the relays in the SFN, even though it itself is not part of this network. It would therefore need to be fed with information generated by a T2 Gateway as well as the line fed relays, but should ignore the "emission time" and any frequency control information.

7.6 Regional/local content insertion

Regional or local services could be inserted in a distributed way after the T2 gateway, provided the corresponding capacity has been prepared in the T2-MI stream using one or more dedicated PLPs carrying dummy data. If e.g. a constant number (per T2 frame) of Baseband frames are reserved (e.g. by using "dummy" Baseband frames) for local/regional services then these dummy Baseband frames could simply be replaced by real Baseband frames with local/regional content, without affecting anything else in the T2-MI or DVB-T2 configuration. For single PLP use this simple replacement process may not be feasible, since a simple replacement would require PCR restamping of the regional input stream. In this case the distribution to the regional insertion point may have to be done based on conventional TS distribution followed by remultiplexing. In some scenarios this may also be applicable to the multiple PLP case.

In figure 29 a possible example of the distribution of video services to transmitters is depicted, using the T2 Gateway and the T2-MI interface.




Figure 29: Example of Distribution of video services via T2 Gateway and T2-MI interface

8 T2-gateway

This clause gives guidance on the implementation of the functions associated with a T2-gateway, as described in clause 4.4.2, or with other equipment expected to be co-sited with the T2-gateway. In simple applications or networks these functions may be carried out by the modulator itself.

8.1 Input interface

Formats for the input interface to the T2 gateway itself are expected to include:

- one or more MPEG-2-TS multiplexes (see ISO/IEC 13818-1 [i.16]); and/or
- one or more GSE streams (see TS 102 606 [i.9]).

8.2 Generation of one Transport Stream per data PLP

In input mode B a number (N) of transport streams may be input to the DVB-T2 system and allocated to N data PLPs (and a common PLP, where applicable). Two different approaches to generate these TSs are outlined below. In both approaches, when the common PLP is used a statistical multiplexer should take into account the typical difference in the MODCOD applied to the common PLP and to the data PLPs, as well as take into account the effects of the splitting of input TSs into data PLPs and a common PLP, which causes the sum of what is sent in the data PLPs and the common PLP to be lower than the sum of the input TSs.

8.2.1 Separate TSs with common control

One way of generating the N TSs is to use N variable-bit-rate encoders, each producing a constant-bit-rate TS containing one service (video+audio+other possible components), where the proportion of null packets varies depending on the instantaneous bit rate (see figure 30). A central control ensures that the total amount of data (i.e. not considering null packets) does not exceed the capacity limit of the N PLPs.




Figure 30: Example of separate TSs with common control

To allow for insertion of common data in the TSs, the N TSs need to have the same bit rate and include co-timed null packets with a bit rate at least as high as that of the common data. This common data will be inserted later in the chain, replacing the co-timed null packets, and will be carried in the common PLP. The bit rate of each TS should at least be equal to the peak of the sum of:

- 1) the service bit rate; and
- 2) the bit rate of the common data.

EXAMPLE: With peak video-bit-rate of 15 Mbps, an audio bit-rate of 0,5 Mbps and a PSI/SI bit-rate of 0,5 Mbps the TS bit-rate should be at least 16 Mbps.

A remultiplexer, which should preferably be co-sited with the encoders, could then receive these TSs and replace some co-timed null packets with common data, according to the rules for splitting of TSs into data PLP and common PLP, see clause 8.3. The remultiplexer would then output the N TSs to the T2 gateway, with the required characteristics according to annex D of [1.1]. With this method PCR values do not have to be modified anywhere in the chain. It should be noted that in a practical implementation the remux and the T2 gateway may be a single equipment.

8.2.2 From one big TS

Another way to generate the N TSs is to use a normal statmux and then split this into the N TSs. N TSs would first be generated as above, but without the requirement for same bit rate and co-timed common data, and then fed to a multiplexer. The multiplexer then assembles a big TS, containing all services of the elementary TSs and includes also all PSI/SI and other common data, according to splitting rules. The bit rate of the big TS (again not considering the null packets) should not exceed the capacity limit of the N PLPs. Using traditional TS-distribution means the big TS may be transported to another site, where the splitting takes place. The big TS is finally split into N TSs by two conceptual steps:

- 1) cloning of the big TS into N identical TSs;
- 2) for each TS, keep only the packets for one PLP in the TS by replacing all other TS packets with null packets.

The result of this split operation is N TSs, each keeping the bit rate of the big TS, including one service and including common data co-timed with other TSs. Additionally the PSI/SI of the original TS will have to be modified so that e.g. the single SDT (actual) of the input TS is separated into N new SDT (actual), one for each output TS. The PAT will also need to be modified. The operations performed at "the other site" may be performed using separate equipment for the remultiplexer and the T2 gateway or as a single equipment.

8.2.3 Discussion

The difference in end result between the two methods is only in the TS bit-rate and proportion of null packets, which is much higher in the second case. It should be noted that since Deletion of Null Packets (DNP) is limited to sequences of 255 null packets (signalling limited to 8 bits) and the bit rate of the big TS may be as high as 72 Mbps, the DNP process will not be able to eliminate all null packets when the TS bit rate is very high. In the extreme case, where a 72 Mbps TS contains only null packets, 281 kbps of null packets would have to be transmitted in the TS ($72 \text{ Mbps}/256 = 0,281 \text{ Mbps}$). This effect has to be taken into account in the choice of minimum service bit-rates in the variable-bit rate encoding.

The total complexity is lower in the first case, whereas the advantage in the second case is that conventional statmuxes can be used to generate the single big TS. However, in a system using T2 Gateways the multiplexing/remultiplexing functionality may in both cases be integrated with the T2 Gateway. In the second variant PCR values have to be restamped in the remultiplexer, since the bit rate is changed.

Another possible benefit of the second arrangement is that it would allow a professional receiver to re-generate and output the original big input TS by decoding and merging all of the PLPs, because the null packets would be in complementary positions. This might be useful for re-broadcast, head-end or archiving applications.

8.3 Use of common PLPs

The common PLP of a group of PLPs is transmitted in such a way that a receiver can receive simultaneously any data PLP of the group as well as the common PLP. Since the common PLP typically benefits from much less time interleaving than the data PLPs it may be sent with a more robust MODCOD.

When the DVB-T2 system is fed with a group of TSs, all to be sent in a group of PLPs and a common PLP, the input TSs are split into one part of a TS going into a data PLP and the other part going in the common PLP, see figure 31. The TSPSs (Transport Stream Partial Stream) and TSPSC (Transport Stream Partial Stream Common) are the streams that are carried by the data PLPs and the common PLP respectively, i.e. TS2 is split into TSPS2 and TSPSC, whereas TS3 is split into TSPS3 and *the same* TSPSC as for TS2.




Figure 31: Multiple TS input/output to/from the extended DVB-T2 PL

The input TSs follow certain rules, e.g. that they have the same bit rate and are co-timed. Data to be sent in the common PLP (common data) also has to obey certain rules in the input TSs.

There are three different categories of common data that have to be sent in the common PLP when the necessary conditions are fulfilled:

- 1) TS packets that are co-timed and identical on all input TSs of the group before the split;
- 2) Service-Description Table (SDT), not having the characteristics of category (1), i.e. SDT describing a service carried by one of the TSs in the group;
- 3) Event-Information Table (EIT) not having the characteristics of category (1), i.e. EIT describing a service carried by one of the TSs in the group.

To be carried in the common PLP, all data of category (1) have to be co-timed, in the sense that identical TS packets occur on co-timed TS packets. If this is the case, these packets are copied into the common PLP at the same time positions; the copied TS packets in the data PLPs are replaced by null packets. In the receiver a TS packet in the common PLP will always replace co-timed null packets in data PLPs.

Category (1) also includes "SDT other" and "EIT other" describing services that are not carried in any of the TSs of the group, i.e. services carried in another DVB-T or DVB-T2 transmission or another group within the same T2-system. These are treated exactly as for any other data that is common to all the TSs.

For category (2) all SDT actual tables are sent in their respective data PLP. However, TS packets carrying 'SDT other' are sent in the common PLP provided all TS packets carrying SDT other are arranged as in figure 32, where all packets in e.g. "TS3 column" describe services within TS3.




Figure 32: Arrangement of SDT other in input TSs and relationship with TSPSC

In the receiver a TS packet in the common PLP with PID=0x0011 will always replace a co-timed null packet in data PLPs.

For category (3) all EIT data is sent only in the common PLP, provided input data is arranged as in figure 33. It should be noted that it is the EIT other data that is copied into the common PLP, not the EIT actual.




Figure 33: Example of arrangement of EIT actual/other in input TSs and relationship with TSPSC

In the receiver a TS packet in the common PLP with PID=0x0012 will always replace a co-timed null packet in data PLPs. The table_id, last_table_id and CRC may also be modified in the receiver to achieve full transparency of the TS; see clause 10.6.1.6.

8.4 PLP grouping and statistical multiplexing

8.4.1 PLP groups with the same modulation and coding

The most straightforward way of performing statistical multiplexing between PLPs is to start with a conventional statmux of constant total bit-rate and allocate each of the services to a PLP, using the same coding and modulation for all the PLPs. All the PLPs would therefore have the same number of cells and the same number of net bits per FEC block. The total number of FEC blocks per Interleaving Frame would therefore be constant, although the number of FEC blocks allocated to each PLP would vary dynamically.

Even in this case, consideration needs to be given to the scheduling since allocation to Interleaving Frames is done on a BBFRAME basis whereas the statmux is done at an MPEG-packet level. See clause 8.6.1.

It may be desirable for the common PLP of a group to use different modulation and coding from the data PLPs, to compensate for the reduced time diversity of a common PLP compared to the type 2 PLPs with their multiple sub-slices. This can be accommodated if the common data is arranged to have a fixed data rate.

A T2-system may contain different groups of PLPs each with different coding and modulation, provided the total bit-rate within each group is constant.

8.4.2 T2-aware statistical multiplexing

More advanced statistical multiplexing could in principle be performed between PLPs with different code rates or modulation.

In the case of different code-rates, the statistical multiplexer would need to allocate bits to each service in proportion to the code-rate of the associated PLP, so that the total gross bit-rate remained constant. Using the "one big transport stream" model of clause 8.2.2, the big TS would have a rate corresponding to the entire capacity being given to the PLP of the highest (i.e. weakest) code-rate.

Where the PLPs in a statmux group are also modulated using different orders of QAM the situation is more complicated still, because the FEC blocks will be of different sizes. The total number of FEC blocks per Interleaving Frame would therefore not be constant. One option would be to determine the worst-case combination of numbers of FEC blocks of each size and use this as the fixed capacity. The principle outlined above, of allocating net bits to services in a weighted manner, could then be applied safely. The weighting in this case would be based on the number of net input bits per output cell rather than the code-rate alone.

NOTE: This would result in a dynamically variable number of dummy cells and a dynamically variable value of the `SUBSLICE_INTERVAL` parameter. This is one reason why this parameter appears in the dynamic rather than configurable L1-signalling.

It has also been suggested that statistical multiplexing could be carried out between PLPs with different interleaving depths. However, it is not clear whether this is possible since the allocation would need to be done on different timescales.

Further work would be needed on all of these possibilities, but the DVB-T2 specification has been kept flexible enough to allow them to be used in the future if capable equipment becomes available. They will not affect receivers as long as they are implemented according to EN 302 755 [i.1].

8.5 Mode adaptation

The choice of options for Mode Adaptation were discussed in clause 5.12.

The following points should be noted by implementers of the mode adaptation:

- Where applicable, including for Transport Streams, the sync bytes are always removed, both in HEM and NM. This means that, in HEM, a (188-byte) TS packet will occupy only 187 bytes of the Data Field and the bit rate is therefore higher than might be expected. This is taken into account in the formula of clause 6.3.2.
- The DNP field indicates the number of deleted null packets *preceding* the UP with which it is associated, even though it comes after it.

- When null packet deletion is used, implementers should ensure that the last complete packet of each Interleaving Frame is always transmitted (i.e. not deleted), even if it is a null packet. This is illustrated in figure 34. Otherwise, the DJB in the receiver could underflow when there is a run of null packets at the end of an Interleaving Frame, because it will not yet have received the DNP byte describing this run. The subsequent packet, which will be fragmented, need not be transmitted if it is null. However, to allow for the worst-case fragmentation, an extra two packet durations should be included in the design delay as described in clause 8.8.3.
- In Normal Mode, SYNCD points to the CRC-8 calculated over the preceding UP. If this CRC-8 is not entirely carried in a BBFRAME, the UP is not considered complete and SYNCD should point to the next CRC-8.
- Sync-byte removal should be done before the scheduling (clause 8.6), but BB-header insertion can only be performed after scheduling once the position of the in-band signalling is known. The implementation therefore needs to keep track of the packet start positions until the BB-header, including SYNCD, has been inserted. See clause 8.6.1.




Figure 34: Retaining the last complete packet of the interleaving frame to avoid DJB underflow

8.6 Scheduling

The scheduler has two main tasks:

- Allocating sections of the input bit-streams to Baseband Frames within particular Interleaving Frames.
- Allocating the slices and sub-slices of each PLP to cells of the T2-frame.

There is a degree of flexibility in the way that both tasks are carried out, and the T2 specification [i.1] only requires that certain constraints be met:

- The number of cells allocated does not exceed the available capacity.
- The receiver is able to locate and extract the PLPs using the L1 signalling.
- The de-jitter buffer does not underflow and the Time-Interleaver memory does not overflow.

In effect, the specification defines a receiver model for extracting the PLPs from the frame based on the dynamic signalling, decoding them, buffering them and re-combining them where appropriate. The designer of a modulator or T2-gateway may use any scheduling scheme that will work with this receiver model.

8.6.1 Model for allocating number of FEC blocks to each Interleaving Frame

8.6.1.1 Mode A

When a single PLP is used (Mode A), there will be a constant number of FEC blocks per Interleaving Frame and each Interleaving Frame will correspond to one T2 frame. This in turn will correspond to a fixed number of input bits, and the scheduler need only count these bits and map them into the BBFRAMES.

8.6.1.2 Mode B

When multiple PLPs are used (Mode B), with potentially variable bit-rates in each PLP, the situation is more complicated, because of the following:

PLP data is quantised into BBFRAMES - there have to be a whole number of FEC blocks, and therefore BBFRAMES, in each Interleaving Frame.

Even where the PLPs have constant bit-rates, the corresponding *average* number of FEC blocks per Interleaving Frame might not be an integer, and therefore the number of FEC blocks per Interleaving Frame (which *is* an integer) will vary from Interleaving Frame to Interleaving Frame.

When in-band signalling is used, the capacity of the first BBFRAME of an Interleaving Frame will be less than for the other BBFRAMES.

Different PLPs may use different modulations and code rates.

Two possible models for scheduling will be presented. As well as providing a possible implementation for scheduling, these models will also be used in the subsequent discussion of management of the de-jitter buffer.

We will assume here that there are N_{plp} PLPs, which all use the same modulation and coding and the same interleaving depth, so that their Interleaving Frames are coincident and the total capacity of the T2 frame can be expressed as a total number of FEC blocks (and therefore BBFRAMES) across all PLPs. We will also assume that the input PLPs all come from the same statistical multiplex so that the total input bit-rate R_{total} , adding all the PLPs together, is constant.

In both allocation models, the input bit streams are fed into FIFO buffers (see figure 35). The occupancy of the buffers is monitored by the scheduler.




Figure 35: Block diagram for FEC block allocation models

8.6.1.2.1 Allocation model 1

Allocation model 1 is the more efficient model, but imposes additional constraints, in particular on minimum data rate.

In the normal run of operation, the sub-system behaves as follows. Each time one of the FIFOs contains one BBFRAME's worth of data, another BBFRAME in the current Interleaving Frame is allocated to the corresponding PLP. The BBFRAME of data in the FIFO is read out, a BB Header is added, and the BBFRAME is sent on to the modulator in a T2-MI packet (see clause 7). The scheduler will keep a count of the total number of BBFRAMES allocated to all PLPs in the current Interleaving Frame.

The first BBFRAME of an Interleaving Frame may contain in-band signalling, and therefore the number of available bits will be less than in normal operation by the number of bits allocated to in-band signalling. If no BBFRAMEs have yet been allocated to a particular PLP in the current Interleaving Frame, a BBFRAME will be allocated to that PLP as soon as the buffer occupancy reaches this lower value. At this moment, the bits will be read out and assembled into a BBFRAME with zeros inserted in the PADDING field; these will later be replaced by the in-band signalling information itself.

This process will continue until the total capacity of the Interleaving Frame has been allocated, at which point all of the data for the T2-frames to which the Interleaving Frame is mapped will be available. The scheduler will also know at this stage how many FEC blocks have been allocated to each PLP and can perform the allocation of sub-slices and generate the L1 dynamic data accordingly.

The scheduler will then begin to assemble the next Interleaving Frame. Note that all of the PLPs will now be in the "first BBFRAME" state, and therefore the buffer occupancy required for a BBFRAME may be less because of in-band signalling. (This might mean that some of the PLP FIFOs already contain enough data for a BBFRAME).

As described, the instants at which a complete Interleaving Frame is ready will not be exactly regular. This is because the occupancies of the FIFOs will be continually varying and the total occupancy at the instant when the last BBFRAME is taken will also vary. In one extreme case the "race" between PLPs to fill their FIFOs will be very close and only at the last moment will one FIFO reach the required occupancy threshold before the others. Each other buffer will therefore contain very nearly one whole BBFRAME. From the constant-input-bit-rate assumption we can deduce how long it would take to fill up all the FIFOs in this way. In the other extreme case, all the FIFOs will be empty apart from the one which provides the final BBFRAME for the Interleaving Frame. Again, we can calculate how long it would take to achieve this; the difference between the two times is the amount by which the time of having a complete Interleaving Frame varies.

On the other hand, the modulator needs to generate T2-frames continuously and regularly and therefore needs to be supplied with complete Interleaving Frames regularly too. The T2-gateway should therefore generate regularly spaced instants by which each Interleaving Frame will be supplied to the modulator. The time interval between these instants will be defined as T_{IF} , the Interleaving Frame duration, and in the absence of FEFs is given by $T_{IF} = I_{jump} \times P_I \times T_F$.

We will define the *collection window* for a particular Interleaving Frame as the period of time during which arriving input bits could possibly end up being carried in that Interleaving Frame. We will also define the *allocation window*. This is the period during which a the last bit of a BBFRAME can arrive and that BBFRAME be allocated to the corresponding Interleaving Frame.

The regularly spaced instants generated by the T2-gateway correspond to the *ends* of both the collection window and the allocation window for a frame, since by definition the Interleaving Frame will be supplied as this moment and therefore no bits arriving after this could possibly be carried in it. The timing of the *beginning* of the allocation and collection windows is more complicated and will be discussed below.

The end of the first allocation window should be placed T_{IF} after the T2-gateway starts operating. At the end of the first allocation window, the T2-gateway will usually not have allocated a complete Interleaving Frame, because a number of FIFOs will each contain less than a BBFRAME's worth of data. The T2-gateway should allocate a BBFRAME to the PLP whose FIFO is the fullest, filling the unused bits using BBFRAME padding. This should be repeated as necessary until a whole Interleaving Frame's worth of BBFRAMEs have been allocated.

As a result of the padding, the buffers are now starting off partially filled, and therefore less padding is likely to be required at the end of the next allocation window. Eventually the condition will be reached where no more padding is ever needed. This will occur when the total occupancy of all the FIFOs just before the end of the allocation window is equal to $(N_{plp})(K_{bch}-80-1)+1$. This is one more than the maximum number of bits that can be in the FIFOs without any FIFO containing a whole BBFRAME, and it is therefore certain that a BBFRAME will be ready for allocation, leaving the occupancy just after the end of the window at $(N_{plp}-1)(K_{bch}-80-1)$ bits.

From this point onward, there will be no further padding, and so the number of bits carried in each Interleaving Frame will equal the number of bits arriving in this time, $R_{total}T_{IF}$, and therefore the buffer occupancy at the end of the allocation window will always be the same.

Generally, an entire Interleaving Frame will have already been allocated some time before the end of the collection window. We will now find the earliest time that the Interleaving Frame can be fully allocated. If we define that the allocation began at time $t=0$, then the latest moment at which allocation can be complete is the end of the allocation window for Interleaving Frame n and is given by the original definition as:

$$t_{aw_end}(n) = (n+1)T_{IF}$$

The earliest moment that Interleaving Frame n can be allocated occurs in the case where, after allocating the last frame, the FIFOs are all empty. This is also the beginning of the allocation window for Interleaving Frame $n+1$, since subsequent BBFRAMES can be allocated to the next Interleaving Frame.

It was shown above that the eventual state of the system has the total occupancy of the FIFOs at the end of the allocation window equal to $(N_{plp}-1)(K_{bch}-80-1)$ bits. Since the total rate of arrival of input bits is R_{total} , the time taken to collect this many bits is found by dividing by R_{total} , giving the earliest moment for complete allocation of the Interleaving Frame and hence the start of the next allocation window as:

$$t_{aw_start}(n+1) = (n+1)T_{IF} - \frac{(N_{plp}-1)(K_{bch}-81)}{R_{total}}$$

NOTE 1: We have assumed for simplicity that the scheduler will not begin allocating BBFRAMES to the next Interleaving Frame until the end of the allocation window for the current frame, even if the current frame is completed early; instead all the bits received after the Interleaving Frame is completed will remain in the FIFOs.

NOTE 2: This implies that each FIFO needs to be able to contain as many BBFRAMES as there are PLPs, although there may be scope for an efficient implementation given that the total occupancy of all the FIFOs together is limited.

We can therefore see that the allocation windows will be spaced by T_{IF} , with an overlap of $T_{overlap}$ between the allocation windows of consecutive Interleaving Frames, where:

$$T_{overlap} = T_{end}(n) - T_{start}(n+1) = \frac{(N_{plp}-1)(K_{bch}-81)}{R_{total}}$$

Figure 36 illustrates the concept: in the figure, allocation window 1 is the allocation window for T2-frame 1, and so on.




Figure 36: Allocation windows for FEC-block allocation

We can now consider the collection window. Recall that the allocation window for Interleaving Frame n is the range of times during which the *last bit* of a BBFRAME can arrive in order to be allocated to Interleaving Frame n , whereas the collection window is the range of times during which *any* bit can arrive and be allocated to Interleaving Frame n . Over what range of times can the other bits of a BBFRAME arrive?

The answer to this question reveals a problem: in allocation model 1, input bits could stay in the FIFO indefinitely. For example, the input bit-rate could fall to zero when the FIFO was nearly full, and these bits would stay there until the bit-rate increased again. This situation is clearly not acceptable, because the end-to-end chain needs to have a fixed, pre-defined delay and the bits could potentially still be stuck in the input FIFO of the T2-gateway when they need to be output from the receiver.

A solution to this is to impose a *minimum* bit-rate R_{\min} on each PLP. The maximum time for one BBFRAME of data to arrive at the input would therefore be $T_{\text{BBFRAME_max}} = (K_{\text{bch}} - 80) / R_{\min}$. The beginning of the collection window for a frame therefore becomes

$$\begin{aligned} T_{\text{cwstart}}(n) &= T_{\text{awstart}}(n) - T_{\text{bbframe_max}} \\ &= T_{\text{awstart}}(n) - \frac{K_{\text{bch}} - 80}{R_{\min}} \end{aligned}$$

This defines the earliest time that an arriving bit can be allocated to a particular interleaving frame and is the starting point for defining the latency of the end-to-end chain and managing the de-jitter buffer (clause 8.8). The collection and allocation windows are illustrated in figure 37. The input transport stream packets carried by one PLP are coloured red and those for another PLP are coloured turquoise. The corresponding baseband frames are shown below in the same colour; the duration corresponds to the time over which bits are allocated to each. The red PLP shows a typical case. The turquoise PLP is for the extreme case where the first BBFRAME allocated to a particular interleaving frame ends at the very beginning of the allocation window, and the input bit-rate for this BBFRAME was at the minimum value. The first bit of this BBFRAME is therefore at the beginning of the collection window, as shown.

Imposing a minimum bit-rate also avoids another potential problem. If $T_{\text{BBFRAME_max}}$ is less than T_{IF} , there is guaranteed to be at least one BBFRAME for each PLP in each Interleaving Frame. This ensures that the total bit-rate capacity is constant. If a PLP could have no BBFRAMEs in some Interleaving Frames, these frames would contain no in-band signalling for that PLP and therefore the total bit-rate allocated to in-band signalling would vary, changing the total capacity.




Figure 37: Collection and allocation windows for allocation model 1

The FIFOs described only provide enough buffering to enable the scheduler to determine which BBFRAME is the first in each Interleaving Frame and therefore to allow space for in-band signalling in the correct BBFRAME. The scheduler will not know at this stage how many BBFRAMEs (and therefore FEC blocks) will be allocated to each PLP in each Interleaving Frame, and so cannot yet generate the dynamic L1-post-signalling information for the corresponding T2-frame(s). The "frame delay" shown explicitly in [i.1] will serve this purpose. In addition to this, an additional Interleaving Frame's delay is needed in order to allow the in-band signalling to be generated and inserted, since this applies to the *following* Interleaving Frame. This could be provided by the Time-Interleaver memory as described in clause 9.2.5.

Note that the choice of peak bit-rate and PLP_NUM_BLOCKS_MAX should take into account the possibility that BBFRAMES can be allocated to one Interleaving Frame during the full allocation window period, $T_{IF}+T_{overlap}$, and that the first BBFRAME may be allocated just after the start of the allocation window. The maximum bit-rate should therefore be set such that at most PLP_NUM_BLOCKS_MAX-1 blocks can be received in $T_{IF}+T_{overlap}$. When using allocation model 1 in mode B, the formula for the maximum bit-rate given in clause 6.3.2 should therefore be modified to:

$$R_{\text{model1}} = \frac{1}{T_{IF} + T_{\text{overlap}}} \left(\frac{188}{188 + B_{\text{MA_data}}} \times \frac{(N_{\text{B_max_data}} - 1) \times (K_{\text{bch_data}} - 80) - N_{\text{IB_data}}}{P_{\text{L_data}} \times I_{\text{jump_data}}} + \frac{188}{188 + B_{\text{MA_cplp}}} \times \frac{(N_{\text{B_max_cplp}} - 1) \times (K_{\text{bch_cplp}} - 80) - N_{\text{IB_cplp}}}{P_{\text{L_cplp}} \times I_{\text{jump_cplp}}} \right)$$

8.6.1.2.2 Allocation model 2

Allocation model 2 is much simpler, at the expense of wasted capacity. In this model, the collection window for each interleaving frame is defined to correspond to exactly T_{IF} , with no overlap. At the end of each collection window, any data in each of the FIFOs is allocated to a baseband frame, with padding inserted if necessary.

Assuming that the total data rate corresponds to a whole number of BBFRAMES, the worst case is where the last BBFRAME for one PLP is almost full, whilst the last BBFRAMES for all the other PLPs are almost empty (perhaps containing only one bit each). In order to deal with this eventuality, the total bit rate should be set such that on average there are $N_{\text{plp}}-1$ unused baseband frames. This will typically represent a loss of capacity of around 0,5 % per PLP, although this can be reduced by using short FEC blocks.

8.6.1.2.3 Other allocation strategies

In addition to the two models described above, the following possibilities should be considered:

- Intermediate strategies between model 1 and model 2 could be used. The total input bit-rate would be set somewhat below the total capacity of the frame, and padding used if the entire frame had not been allocated at the end of the collection window. This would reduce the maximum total FIFO occupancy and therefore the size of the overlap window, at the expense of wasted capacity.
- Fixed bit-rate PLPs with a bit-rate corresponding to a whole number of BBFRAMES per Interleaving Frame can be allocated as for Input Mode A as described in clause 8.6.1.1. No BBFRAME padding is required and the considerations of allocation windows do not apply.
- Opportunistic data could be used to fill up the wasted capacity in the baseband frames of allocation model 2, or as a way of maintaining a minimum bit-rate in model 1.
- A combination of the various strategies can be used, e.g. some fixed bit-rate PLPs, a group of higher bit-rate PLPs allocated according to model 1 (with a minimum bit rate), and another group allocated according to model 2 (perhaps using short FEC blocks).

Where the PLPs being carried do not all share the same interleaving depth and MODCOD, the models could be used separately for each set of PLPs that do share the same MODCOD and interleaving depth. The capacity of the frame would need to be divided in a fixed way between each set of PLPs. The implications for the total memory size required for the FIFOs would need to be considered in this case.

8.6.1.2.4 Use of the allocation models with FEFs

As explained in clause 8.6.1.2.1, the spacing between the ends of the collection windows is equal to the Interleaving Frame duration T_{IF} . In allocation model 2 (clause 8.6.1.2.2), this is also equal to the duration of the collection windows.

Where FEFs are used, T_{IF} is longer than $I_{jump} \times P_I \times T_F$ because each Interleaving Frame duration includes a share of the duration of the FEF parts. Since the pattern of FEFs, multi-frame interleaving and frame skipping repeats at least once per superframe, T_{IF} can be found by dividing the total superframe duration by the number of Interleaving Frames in a superframe, giving:

$$T_{IF} = P_I \times I_{jump} \times \left(T_F + \frac{T_{FEF}}{I_{FEF}} \right),$$

where T_{FEF} is the FEF duration and I_{FEF} is the FEF interval.

The allocation models in the previous clauses can be used with FEFs simply by substituting this more complete formula for T_{IF} .

When FEFs are used, the first T2-frames of each Interleaving Frame will not generally be equally spaced in time. However, the ends of the collections windows should be equally spaced, in order that the Interleaving Frames of a superframe each carry the same amount of data.

The variation in the start time of Interleaving Frames will lead to varying delay in the modulator and hence the need for buffering. This varying delay is in fact complementary to the varying occupancy of the receiver's buffers, since the end-to-end delay remains constant.

8.6.1.2.5 Timing and alignment of collection windows

The alignment and timing of the collection windows requires care, particularly when FEFs, multi-frame interleaving or frame skipping are used.

We will specify the timing of the collection windows by defining the transmission delay T_{TX} as the time interval from the end of the collection window for an Interleaving Frame to the beginning of the transmission of the P1 preamble of the first T2-frame to which the Interleaving Frame is mapped.

Where one or more groups of PLPs are used, the collection windows for all PLPs within a group should be aligned such that the ends of the collection windows are co-timed. This is an underlying assumption for the allocation models involving statistical multiplexing between PLPs. This corresponds to choosing the same value of T_{TX} for each PLP.

Where different PLPs in a group have different Interleaving Frame durations T_{IF} , the ends of the collection windows should be co-timed for Interleaving Frames that are mapped starting from the same T2-frame. Again, this corresponds to choosing the same value of T_{TX} for each PLP.

Where PLPs within a group use the frame-skipping mechanism, statistical multiplexing is not possible between PLPs with different values of `FIRST_FRAME_INDEX`, since they do not share the same T2-frames. It is not therefore necessary for the collection windows to be aligned, and if T_{TX} is the same for each PLP then the ends of the collection windows will in fact be staggered.

Where FEFs are used, the value of T_{TX} can in general be different for different Interleaving Frames of a superframe, because the collection windows are evenly spaced but the T2-frames are not. T_{TX} for Interleaving Frame n is given by:

$$T_{TX}(n) = T_{TX}(0) - \frac{T_{FEF}}{I_{FEF}} \left\{ (I_{FF} + nP_I I_{jump}) \bmod I_{FEF} - I_{FF} \bmod I_{FEF} \right\}$$

The similarity of this formula to that for TToBase in clause 8.8.2 is no coincidence, since the end-to-end delay is constant.

$T_{TX}(0)$ should be chosen such that the minimum value of $T_{TX}(n)$ for all Interleaving Frames is large enough to allow the modulator to generate and output the first T2-frame in time. A fundamental limit is that the minimum T_{TX} is greater than or equal to T_{IF} for any PLPs for which in-band signalling is provided, in order for the scheduling for the next Interleaving frame to be known and inserted into the in-band signalling, as discussed in clause 9.1.2. Additional allowances should be added to allow for the distribution network and the minimum delay of the modulators used in the network.

In a single frequency network (SFN) or other network where the timestamp feature of T2-MI is used (see clause 9.1.6), the value of T_{TX} will be used directly in calculating the timestamp value. Where timestamps are not used, the absolute value is not important but the relative values for different PLPs will be needed in calculating the compensating delay (see clause 8.9). Because of the time-varying delay in the modulator, timestamps should be used when FEFs are used.

8.6.2 Scheduling of sub-slices

As explained, the second role of the scheduler is to decide where in the frame each of the sub-slices for each PLP will go. The number of sub-slices is statically configurable and is the same for all Type-2 PLPs; the factors affecting the choice of this parameter were discussed in clause 5.8. The locations of the sub-slices are constrained by what it is possible to signal, and also by the rule that the common PLPs come first, followed immediately by the Type-1 PLPs then the Type-2 PLPs.

This leaves only one aspect for the scheduler itself to choose: the order in which the PLPs are mapped. In many cases, this will have no significant impact, and implementers may map the PLPs simply in order of their PLP_ID. However, in situations which are critical in terms of management of the Time De-interleaver and De-jitter buffer memories (see clause 8.8), it might prove helpful to schedule PLPs at particular positions in the T2-frame to ensure that data is delivered as close as possible in time to when it is required.

The choice of the number of PLPs, the division into types 1 and 2, and the proportion of the T2 frame allocated to Type-1 and Type-2 data can also have an effect on the buffer management.

The scheduling of the Type-1 slices will have an effect on the amount of power-saving that a receiver can perform - receivers receiving PLPs early in the frame will need to remain powered for a shorter time to receive the L1 data, the common PLP and the data PLP slice.

8.7 Padding methods:

For a number of reasons, the capacity of the T2 system might exceed the amount of data to be carried. Possible reasons are:

- The TS input bit-rate has been deliberately set slightly below the theoretical capacity.
- Statistical multiplexing, particularly a T2-aware variant, might require some headroom to be left for safety.
- A PLP may carry opportunistic data and no such data is available at a particular time.

DVB-T2 offers three mechanisms for padding:

- Null packets (without NPD).
- BBFRAME padding.
- Dummy cells.

Although null-packet deletion is available as an option, the system can carry some null packets in their entirety. The two techniques may be mixed, with some null packets transmitted and others deleted and re-inserted.

BBFRAME padding may be used. This can be added to the scheduling model described in clause 8.6.1.2.2; when an Interleaving Frame is due for transmission, any spare space in the last incomplete BBFRAME may be filled with padding, and further BBFRAMEs may also be added containing only padding.

Finally, there will almost always be dummy cells inserted at the end of the T2-frame: as the T2-frame is very unlikely to contain a whole number of FEC blocks, the remaining space at the end of the T2-frame will be filled by dummy cells (see clause 6.1.5). However, if the scheduler provides fewer FEC blocks than normal, the number of dummy cells will increase automatically to fill the space. Note that this will cause the SUB_SLICE_INTERVAL parameter to vary dynamically. Large variations are undesirable since increasing the number of dummy cells will reduce the time period over which the sub-slices are spread, reducing time diversity and possibly making buffer management more difficult.

The first two methods are therefore preferable in most applications. The choice between them is likely to depend on the nature and source of the input stream. Using null packet stuffing only allows the modulator to set the "NO_PADDING_FLAG" in the L1-configurable signalling to '1' (clause 8.10.5), potentially allowing faster initial acquisition in the receiver.

8.8 Managing the de-jitter buffer

[i.1] provides a receiver-buffer model which includes both the time de-interleaver memory and the de-jitter buffer, as well as the frequency/L1-de-interleaver. It is the responsibility of the modulator or T2-gateway to manage these buffers and to ensure that they neither under- nor over-flow.

The buffer occupancy is managed using the TTO variable of the ISSY field. This variable is similar in intention to the BUFSTAT variable defined for DVB-S2. However, the situation in T2 is more complicated because of the presence of the time de-interleaver and also the potential use of null-packet deletion.

Whereas BUFSTAT simply indicates the buffer occupancy, this is not sufficient in T2 because the time de-interleaver memory also acts as a buffer. In fact, in the highest bit-rate modes, the TDI memory can hold twice as much data as the DJB and is therefore essential in providing buffering.

TTO is essentially defined as the delay between the beginning of a T2-frame and the time at which the first bit of one of the user packets carried in that frame is output. To allow for multi-frame interleaving, the definition is expressed in terms of the first T2-frame of an interleaving frame.

The correct use of TTO to manage the receiver's buffers depends on several parameters of the system, and is only possible if the parameters meet certain constraints. The relevant parameters are:

- Number of TI-blocks per Interleaving Frame (N_{TI}).
- Number of T2-frames to which Interleaving Frame is mapped (P_I).
- Type of PLP (Common, Type 1 or Type 2).
- Number of sub-slices (N_{sublices}) for Type-2 PLPs.
- Number of symbols per T2-frame (L_P).
- Constellation and code rate.
- Use of FEFs.
- Allocation model and minimum bit-rates.

TTO can be thought of as indicating the combined occupancy of both the de-jitter buffer and the time de-interleaver, although the units are multiples of the elementary period T rather than bits or cells.

In modes where the bit-rate does not vary and PLP cells are delivered roughly uniformly, a self-balancing method can be used for the DJB, since the DJB occupancy will be essentially constant. However, in modes where the bit-rate varies (because of statistical multiplexing), or the channel delivers cells in a non-uniform way (e.g. because of the use of FEFs), the receiver will not know what an appropriate buffer occupancy is at any given time. EN 302 755 [i.1] therefore mandates that ISSY be used in such cases.

8.8.1 The receiver-buffer model




Figure 38: receiver buffer model

The receiver-buffer model from annex C is shown in figure 38. The frequency/L1-de-interleaver, time de-interleaver and de-jitter buffer are shown, together with the processing before and between the blocks. There is assumed to be a single FEC chain, including the LDPC decoder; the FEC chain is shared between the data and any common PLP, as well as being used for decoding the L1 signalling.

In some circumstances cells could in principle be read from the TDI faster or slower whilst still delivering the output bits correctly. To avoid this uncertainty, the model states that the contents of the TDI memory will be read as fast as possible without overflowing the DJB, but subject to a maximum reading rate. The maximum reading rate of $7,6 \times 10^6$ cells/s is specified in order to limit the speed at which the FEC chain, and in particular the LDPC decoding, needs to operate. The particular rate was chosen to correspond to the maximum rate at which cells can be delivered by a T2 system over all the possible parameter settings in channel bandwidths up to 8 MHz.

The model also defines the way in which the FEC chain is shared between the data PLP, common PLP and L1. Essentially the L1 has absolute priority, to the extent that the receiver is not allowed to begin sending PLP cells to the FEC chain unless the entire FEC block can be decoded before the L1 cells become available. This ensures that the L1 can always be decoded in time to start extracting PLP cells from the frame. When the FEC chain is not needed for L1, the data and common PLPs are allowed to use it in strict rotation. The FEC chain only deals in whole FEC blocks at a time.

Since cells are only sent to the FEC chain when there is space in the DJB to receive the resulting bits, it is not possible for the DJB itself to overflow. Instead, overflow would occur at the TDI. Since the TDI is not a simple buffer, the model has to define what exactly constitutes overflow in the TDI.

The TDI is assumed to have only one block of memory containing $2^{19} + 2^{15}$ OFDM cells, and to use the algorithm described in 10.4.3.2.3. In single PLP modes the model also includes a 4 000-cell FIFO. There are three criteria for overflow:

- Overflow clearly occurs if there are more cells in the TDI than there are locations in its memory.
- The time de-interleaver of the model can only contain cells from at most two different TI-blocks at any one time: one block being written and another block being read. If there are ever cells from three or more different TI blocks in the TDI, then overflow is considered to occur.
- The third criterion is expressed as an equation relating the indices of the cells being written and read. Although the equation looks complicated, it simply represents the algorithm described in clause 10.4.3.2.3; overflow will occur if the write pointer ever overtakes the read pointer.
- The N_{FIFO} term is included to allow for the 4 000-cell FIFO in single PLP modes.
- The first argument of the "max" function in the expression makes an allowance for alternative TDI implementations.

The Time De-interleaver cannot output the first cell of a TI-block until the last cell of that block has been written into the memory. It can then begin reading subject to the rules above.

For simplicity, the FEC chain and other processing blocks are assumed to have zero delay. Note that this not only ignores practical delays such as extra registers between logic functions, but also that the processing as modelled would be non-causal. For example, the demodulator in the model starts to output the equalised cells of a symbol at the beginning of that symbol, whereas in fact a real implementation could not possibly calculate the necessary FFT until the entire symbol had been received. Similarly, the LDPC decoder in the model will correct and output each bit immediately as it is input, whereas a real implementation would need to have received all of the parity bits before it could correct the first information bit. In addition, compensating delays required in the channel equalisation process and other similar functions are also ignored. These all need to be taken into account in the receiver when interpreting the TTO variable.

8.8.2 Principles of buffer management

The general principle for managing the buffers is as follows. The end-to-end chain will have a constant delay, since the output TS will emerge at the same bit-rate as the input. The one parameter to be chosen is what this delay is. In the following discussion, we will define a number of terms. For the purpose of this clause and the following clauses, the discussion will be limited to allocation model 2 (see clause 8.6.1.2.2), since experience has been gained with this model.

TTO_{base} will be defined as the time at which the first bit carried in Interleaving Frame should be output, relative to the start of the first T2-frame to which the Interleaving Frame is mapped, in the condition where the bit arrived at the very beginning of the *collection window* for the Interleaving frame (see clause 8.6.1.2 for discussion of the collection and allocation windows).

When FEFs are not used, TTO_{base} will be the same for every Interleaving Frame, whereas when FEFs are used its value may be different in different Interleaving Frames of the superframe.

The Design Delay T_{design} will be defined as the value of TTO_{base} that applies to the first Interleaving Frame of a superframe.

TTO_{base} for Interleaving Frame n can be calculated from T_{design} using the following formula:

$$TTO_{base}(n) = T_{design} + \frac{T_{FEF}}{I_{FEF}} \left\{ (I_{FF} + nP_1 I_{jump}) \bmod I_{FEF} - I_{FF} \bmod I_{FEF} \right\},$$

where T_{FEF} is the duration of the FEF part, I_{FEF} is the FEF_INTERVAL and I_{FF} is the value of FIRST_FRAME_INDEX for the relevant PLP.

TTO for a given Interleaving Frame is carried in the first BBFRAME of the Interleaving Frame, and needs to be calculated every time from TTO_{base} , taking into account the following:

- TTO refers to the first complete user-packet, and the number of bits before the beginning of the first complete user-packet will vary.
- There could be one or more deleted null-packets to be output before the first complete user-packet, whereas the TTO (and other ISSY fields) apply to the first non-deleted packet.
- The first input bit might have arrived earlier than the end of the collection window for the previous Interleaving Frame (by up to $T_{BBFRAME_max} + T_{overlap}$). See clause 8.6.1.

The value of TTO will therefore vary from frame to frame, but it can be calculated directly by the scheduler. All three effects above can in principle be taken into account by keeping track of the arrival time of the first bit in the FIFO relative to a counter that counts down to the end of each collection window. This can be updated as bits are read from the FIFO; note that deleted null packets should also be taken into account.

TTO should be calculated as follows. When the first baseband frame is allocated to an Interleaving frame for a given PLP:

- Read the time corresponding to the first bit in the FIFO relative to beginning of the collection window.
- Adjust the time according to the position of the first complete User Packet in the FIFO. (The position will be known from the SYNC calculation - see clause 8.5).
- Add the result to the value of TTO_{base} for the current Interleaving Frame.

- Generate the ISSY TTO variable and store it in the BBFRAME in the appropriate place for the Mode Adaptation in use (see clause 8.5).

Clearly T_{design} should be chosen such that the resulting TTO value is always realisable; this will be discussed below.

Figure 39 shows an example. A series of T2-frames are shown, each carrying a number of TS packets. The TS packets are shown in the same colour as the T2-frames in which the bits are carried; some packets are fragmented, i.e. carried in two T2-frames, so have two different colours. The lower arrows indicate the design delay T_{design} between the start of the T2-frame and a bit received at the start of the collection window being output; in this example this is set to one T2-frame duration. The period during which the bits arriving in the collection window for a particular Interleaving Frame are output is referred to as the "output window".

The upper arrows show the values of TTO that are signalled, with the thin lines connecting these arrows to the TS packet with which the TTO value is associated. The variation in the TTO value, caused by the varying position of the first complete TS packet, can be seen. In the last frame, there is a deleted and re-inserted null packet; the TTO value is therefore associated with the first non-null packet of the frame and TTO has a greater value than it would have been had the null packet not been present.




Figure 39: Use of TTO

8.8.3 Choice of design delay

Having defined the terms, the next step is to choose the design delay. If the design-delay value is too great, there will not be enough memory in the de-jitter buffer to provide the delay, and overflow will occur in the TDI (recall that the DJB itself cannot overflow according to the receiver buffer model). If the design delay is too small, the output will be required before the corresponding cells have been delivered by the T2-frame structure, and the DJB will underflow. The other parameters mentioned above will determine whether it is possible to choose a design-delay value which is always in the safe range, so in practice these parameters also need to be chosen with de-jitter-buffer management in mind.

It is difficult to derive a single rule that will work for all cases. Instead, a number of useful use-cases will be presented.

For single PLP modes, there will often be more than one TI block per Interleaving Frame ($N_{\text{TI}} > 1$). ISSY is not generally mandatory in such modes, unless one or more of the other criteria for ISSY applies, e.g. FEFs are used. In such cases the design delay should be approximately $T_{\text{F}}/N_{\text{TI}}$, but some additional delay will need to be added.

For type 2 data PLPs in multiple PLP cases using null packet deletion and with $N_{TI}=P_I=I_{jump}=1$, together with allocation model 2, the recommended value is:

$$T_{design} = T_F + 2T_{packet} + \frac{N_{cells(common)}}{R_{cell}} + T_{P1} + T_{L1}$$

where $T_{packet} = \frac{1504}{R_{TS}}$, the duration of one packet in the output TS; $T_{P1} = 2048T$ is the duration of the P1 preamble, and T_{L1} is the time taken to decode the L1, as derived in clause 10.4.2.3. Other symbols are as previously defined here or in [1.1]. This formula still applies if FEFs are used.

For common PLPs for which $N_{TI}=P_I=I_{jump}=1$, the most straightforward choice is to set T_{design} to be the same as for the associated data PLPs. This is not the optimum as a shorter design delay would reduce the DJB requirements, but as long as the common PLP has a relatively low bit-rate, the impact on the DJB budget is minimal. This choice avoids the need for a compensating delay (see clause 8.9).

For PLPs which use the frame-skipping mechanism and/or multi-frame interleaving but without FEFs, the design delay should be:

$$T_{design} = \{(P_I - 1)I_{jump} + 1\}T_F + 2T_{packet} + \frac{N_{cells(common)}}{R_{cell}} + T_{P1} + T_{L1}$$

Where FEFs are used with frame skipping and/or multi-frame interleaving, the design delay needs to take into account the worst case for FEF parts occurring in between T2-frames of an interleaving frame. Deriving a general formula is difficult but for specific cases the worst case can be found by inspection.

In general it is recommended that some margin be added to the design delay, where possible.

8.8.4 The de-jitter buffer budget

Analysis of the DJB budget for multiple PLPs (with variable bit-rate) is rather involved, since it is necessary to identify the worst possible combination of incoming cells and outgoing packets. Broadly speaking, the worst case occurs when the data to be output in a particular frame is all concentrated at the end of the output window, with only deleted null packets at the beginning. Cells are "pushed out" of the TDI by the cells arriving in the T2 frame; these cells become decoded bits in the DJB, but cannot be output until later because the deleted null packets have to be output first. Finding the worst-case occupancy involves determining the worst proportion of the output window to be filled with non-null packets (known as the "critical fill ratio").

NOTE: The cells are not really pushed out of the TDI, since the TDI is always read as quickly as possible under the receiver buffer model. However, there is a latest moment that they have to be read, determined by the definition of TDI overflow, so for the purpose of analysis we can consider that the cells are pushed out at the moment that the location is about to be overwritten.

This process is illustrated in figure 40.




Figure 40: Basic model of DJB budget at the critical fill ratio

In this model, bits are forced into the DJB after a short gap, T_{gap} , which corresponds to the time for the common and type 1 PLPs to be delivered. The DJB occupancy increases at a rate R_{chan} , which is the equivalent maximum decoded bit-rate at which cells are delivered by the channel. This will be given by the total number of bits delivered in a frame for the PLP in question, divided by the time to deliver these bits, $T_{channel}$. $T_{channel}$ is the total duration of type 2 PLPs. When the data in the output window is at the critical fill ratio, filling will just stop at the moment that the output starts, and bits will be output at the maximum output rate for that PLP, R_{max} . Output will finish at exactly the end of the output window, and so it can be shown that the total time illustrated in figure 40 is given by $TTObase$.

In many simple cases, the budget for a given interleaving frame is given by the following formula, assuming that $N_{TI}=1$:

$$DJB = \frac{R_{max} \times R_{chan}}{R_{max} + R_{chan}} (TTObase - T_{gap})$$

For a given set of parameters, it can be seen that the maximum budget will occur for the maximum value of $TTObase$, which in turn is given by:

$$TTObase_{max} = T_{design} + T_{fef} \left[\frac{I_{fef} - \gcd(P_i I_{jump}, I_{fef})}{I_{fef}} \right]$$

From this, it can be seen that if the gap is small, so that R_{chan} is similar to R_{max} , and there are no FEFs, the DJB budget will be approximately $R_{max} \cdot T_{design} / 2$. If T_{design} has been set to its minimum value, this will correspond roughly to no more than 1/2 of one TI-block, and since this is guaranteed to be less than 2 Mbit, the DJB will not usually impose any constraint on the system design.

If there is a significant amount of common or type 1 data, this will have 2 effects. Firstly the duration $T_{channel}$ will reduce, and so the ratio R_{max}/R_{chan} will decrease. This would tend to increase the DJB budget, but this is likely to be more than compensated by the reduction caused by the increase in T_{gap} .

When FEFs are present, $TTObase$ may increase, and this will tend to increase the DJB budget. However, if the FEF interval is low, the maximum bit-rate will be correspondingly reduced. The worst case for the DJB, which will ultimately become a limiting case, will be when a single, infrequent, long FEF is present - this will increase the value $TTObase_{max}$, but with very little decrease in the maximum bit-rate.

The situation with type 1 and common PLPs is quite similar, but in this case the channel duration $T_{channel}$ will be just the duration for the cells of the corresponding PLP, and hence in general the ratio R_{max}/R_{chan} will be quite small. The DJB budget will then tend towards $R_{max} \cdot T_{design}$. This is unlikely to be a problem, provided that the bit-rates for type 1 and common PLPs are kept reasonably low.

In more complex cases for example multi-frame interleaving is used, similar principles to the ones described above still apply, but an additional allowance will need to be considered for the delivery of full frames of data, possibly interspersed with FEFs, as shown in figure 41. Deriving a general formula to deal with such cases is not easy, but the calculation of the budget is more amenable to a computer programme to search for the worst case.




Figure 41: The delivery of data interleaved over several T2-frames, with a FEF between frames

For single PLPs, or indeed PLPs as part of a group of multiple PLPs where the output bit-rate is known to be constant, the DJB budget calculation is generally simpler. Broadly speaking, it can be derived by calculating the bits delivered during one frame (or TI-block in the case of multiple TI-blocks), and subtracting the bits which will be known to have been output during the frame time. The minimum DJB occupancy will need to be added. Again, suitable allowance needs to be made for FEFs.

NOTE: Previous drafts of the present document included different calculations. Those calculations were thought to be conservative, and further study has confirmed that this is indeed the case. In particular, the various factors were assumed to combine additively, whereas in fact some of them act in opposite directions to others, and even those which act in the same direction tend not to create the worst-case occupancy at the same time.

8.9 Compensating Delay

A compensating delay is required if common PLPs are used and the data PLP uses different interleaving parameters to its associated common PLP. The relevant interleaving parameters are N_{TI} , P_1 and I_{jump} ; if any of these differ then a compensating delay is required.

The purpose of the compensating delay is to ensure that the streams at the output of the de-jitter buffer, operating according to the signalled values of TTO, are essentially co-timed, i.e. that they are output with the same relative timing as they arrived at the input of the basic T2-gateway.

In the receiver, the TTO variable defines exactly when the first bit of an interleaving frame should be output; this is in turn derived from the fixed value of design delay as discussed in clause 8.8.3. The design delay T_{design} is the time from the beginning of the P1 preamble of the first T2 frame of the first Interleaving Frame of the superframe to the time a bit is output that arrived at the end of the overlap period of the collection window.

At the transmitting end, the transmission delay $T_{TX}(0)$ was defined in clause 8.6.1.2.5 as the time from the end of the collection window to the beginning of the same P1 preamble. A bit arriving at the end of the overlap period will have arrived T_{IF} earlier.

Other delays in the system will be the same for all PLPs, so can be ignored.

The PLP-specific end-to-end delay is therefore given by $T_{delay} = T_{IF} + T_{TX}(0) + T_{design}$, as illustrated in figure 42.




Figure 42: End-to-end delay

The required compensating delay can therefore be found by calculating $T_{\text{delay_max}}$, the maximum value of T_{delay} for all the PLPs in a group, including the common PLP itself. Each PLP should be delayed by a period equal to the difference ($T_{\text{delay_max}} - T_{\text{delay}}$) between its own delay and the maximum.

8.10 Layer-1 signalling

8.10.1 Overview

The Layer-1 (L1) signalling has two main functions. First, it provides receiver a means for fast signal detection and hence enables fast signal scanning. Secondly, it provides all information that receiver needs to access the Layer-2 signalling and the services themselves within the current (and possibly also the next) frame. The L1-signalling structure is illustrated in figure 43, where the three main L1-signalling sections are presented: the P1 signalling, the L1-pre signalling and L1-post signalling.

The purpose of the P1 signalling, which is carried by the P1 symbol, is to indicate the transmission type and basic transmission parameters. The remaining signalling is carried by the P2 symbol(s), which may also carry PLP data. The L1-pre signalling provides basic transmission parameters including parameters required to enable the reception and decoding of the L1-post signalling. The L1-post signalling is further split into two main parts: configurable and dynamic, and these may be followed by an optional extension field. The L1-post signalling finishes with a CRC and padding (if necessary).




Figure 43: The L1 signalling structure

All L1-signalling data, except for the dynamic L1-post signalling, will remain unchanged for the entire duration of one superframe. Hence any changes implemented to the current configuration will be always done at the border of two superframes.

The P1 signalling is almost solely defined to enable fast signal detection and provide a sub-set of transmission parameters valid for the entire T2 frame. In addition to the latter, it also provides an early indication of the presence of FEF parts within the T2 frame.

The L1-pre signalling, as well as enabling the reception and decoding of the L1-post signalling, also provides information on the current superframe, relating to the network topology, configuration and to the transmission protocols used within the superframe (e.g. TS and GSE).

Finally, the L1-post signalling contains the most of the information needed for extracting and decoding the PLPs from the T2 frames.

8.10.2 Repetition of L1-dynamic signalling

To obtain increased robustness for the dynamic part of the L1-post signalling, the information may be repeated in the P2 symbols of two successive T2-frames. The use of this repetition is signalled by the L1-pre parameter `L1_REPETITION_FLAG`. If the flag is set to '1', dynamic L1-post signalling for the current and next T2-frames are present in the P2 symbol(s) as illustrated in figure 44. Thus, if repetition of L1-post dynamic data is used, the L1-post signalling consists of one configurable and two dynamic parts as depicted.




Figure 44: Repetition of L1-post dynamic information

The L1-post signalling cannot change size between the frames of one super-frame. If there is to be a configuration change at the start of super-frame j , the loops of both parts of the dynamic information of the last T2-frame of super-frame $j-1$ contain only the PLPs and `AUXILIARY_STREAMS` present in super-frame $j-1$. If a PLP or `AUXILIARY_STREAM` is not present in super-frame j , the fields of the relevant loop are set to '0' in super-frame $j-1$.

EXAMPLE: Super-frame 7 contains 4 PLPs, with `PLP_IDs` 0, 1, 2 and 3. A configuration change means that super-frame 8 will contain `PLP_IDs` 0, 1, 3 and 4 (i.e. `PLP_ID` 2 is to be dropped and replaced by `PLP_ID` 4). The last T2-frame of super-frame 7 contains 'current frame' and 'next frame' dynamic information where the PLP loop signals `PLP_IDs` 0, 1, 2 and 3 in both cases, even though this is not the correct set of `PLP_IDs` for the next frame. In this case the receiver will need to read all of the new configuration information at the start of the new super-frame.

8.10.3 L1 change-indication mechanism

The L1 signalling offers a change-indication mechanism, by means of the `L1_CHANGE_COUNTER`-field in the dynamic part of the L1-post signalling. It is used to indicate when the configuration (i.e. the contents of the L1-pre signalling or the configurable part of the L1-post-signalling) will change. The change is indicated as the number of superframes after which the change will occur. The next superframe with changes in the configuration is indicated by the value signalled within this field. If this field is set to the value '0', it means that no change is foreseen during the next 255 superframes. E.g. value '1' indicates that there will be a change in the next superframe. This counter starts counting down from an initial value of at least 2.

8.10.4 Use of the L1 extension field

The L1 extension field may be defined for future applications. Possible applications include signalling to professional receivers for control of network operation, e.g. for control and synchronisation of relay sites. The L1 extension field, if present, should be ignored by receivers.

In version 1.2.1 of [i.1], the first use for an extension field has been defined, for balancing the bias in the L1 signalling (see clause 8.14.1.2). A block structure has been given to the extension field to allow different types of extension block to be added in the future.

8.10.5 Example of L1 signalling

The details of the use of the L1 signalling are specified in [i.1]. This clause gives a full worked example of the construction of the fields of the L1 signalling for the simple case of a single PLP, together with comments to explain the chosen meaning in each case. Table 18 gives a worked example of the signalling fields of the L1-pre-signalling.

Table 18: An example of the contents of the L1-pre signalling fields

| L1-pre field | Contents | Explanation |
|--------------------|--------------------|---|
| TYPE | 00000000 | Only Transport Streams are carried in this signal |
| BWT_EXT | 1 | Extended-carrier mode is used |
| S1 | 000 | This is a T2 preamble operating in SISO mode |
| S2 | 1110 | 32 K FFT is used with one of the guard intervals 1/128, 19/256 or 19/128. This is a repeat of the corresponding field in P1; the exact guard interval is signalled in the GUARD_INTERVAL field. |
| L1_REPETITION_FLAG | 0 | Repetition of the dynamic signalling is not used |
| GUARD_INTERVAL | 100 | Guard interval 1/128 |
| PAPR | 0000 | PAPR is not used |
| L1_MOD | 0011 | L1-post signalling uses 64-QAM modulation |
| L1_COD | 00 | L1-post signalling uses code rate 1/2 |
| L1_FEC_TYPE | 00 | L1-post signalling uses LDPC 16 K |
| L1_POST_SIZE | 000000000011111010 | The number of OFDM cells for L1-post is 250 |
| L1_POST_INFO_SIZE | 000000000100111110 | The number of information bits in the L1-post signalling is 318 |
| PILOT_PATTERN | 0110 | Pilot pattern PP7 is used |
| TX_ID_AVAILABILITY | 00000000 | TX_ID is not used |
| CELL_ID | 0000000000000000 | The provision of CELL_ID is not foreseen |
| NETWORK_ID | 0011000010000101 | The DVB network ID is 0x3085. This corresponds to the value of network_id in the corresponding NIT (see clause 8.11.1) |
| T2_SYSTEM_ID | 1000000000000001 | The T2 system ID is 0x8001 |
| NUM_T2_FRAMES | 00000010 | There are two T2-frames in one super-frame |
| NUM_DATA_SYMBOLS | 000000111011 | There are 59 data symbols in one frame (plus one P1 symbol and one P2 symbol) |
| REGEN_FLAG | 000 | This T2 signal has not been regenerated |
| L1_POST_EXTENSION | 0 | There is no L1-post extension field |
| NUM_RF | 001 | The T2 signal is carried on 1 RF channel |
| CURRENT_RF_IDX | 000 | TFS is not being used |
| RESERVED | 0000000000 | Reserved for future use |
| CRC_32 | 0x99F585A2 | This is the hexadecimal value of the CRC_32 for this L1-pre signalling |

Table 19 gives a worked example of the signalling fields of the L1-post configurable signalling.

Table 19: An example of the contents of the L1-post configurable signalling fields

| L1-post configurable field | Contents legacy/recommended (see note) | Explanation |
|----------------------------|--|---|
| SUB_SLICES_PER_FRAME | 0000000000000001 | There are no type 2 PLPs |
| NUM_PLP | 00000001 | There is 1 PLP (mode A) |
| NUM_AUX | 0000 | There are no auxiliary streams |
| AUX_CONFIG_RFU | 00000000 | Reserved for future use |
| RF_IDX | 000 | This RF has index '0' (the only allowed value, since NUM_RF=1) |
| FREQUENCY | 0x2B805F75 | The centre frequency of this RF channel is 729,833333 MHz (UHF channel 53 with a negative offset) |
| PLP_ID | 00000000 | The ID of this PLP is '0' |
| PLP_TYPE | 001 | This PLP is a data PLP |
| PLP_PAYLOAD_TYPE | 00011 | This PLP carries a TS |
| FF_FLAG | 0 | TFS is not used so this has no meaning |
| FIRST_RF_IDX | 000 | TFS is not used so this has no meaning |
| FIRST_FRAME_IDX | 00000000 | This PLP first appears on the first frame of the super-frame |
| PLP_GROUP_ID | 00000001 | This PLP belongs to the PLP group with ID '1' |
| PLP_COD | 001 | This PLP uses code rate 3/5 |
| PLP_MOD | 011 | This PLP uses 256-QAM modulation |

| L1-post configurable field | Contents legacy/recommended (see note) | Explanation |
|----------------------------|---|--|
| PLP_ROTATION | 1 | This PLP uses constellation rotation |
| PLP_FEC_TYPE | 01 | This PLP uses the 64 K LPDC FEC |
| PLP_NUM_BLOCKS_MAX | 0011001010 | The maximum number of FEC blocks for one interleaving frame for this PLP is 202 |
| FRAME_INTERVAL | 00000001 | This PLP appears in every frame of the super-frame |
| TIME_IL_LENGTH | 00000011 | There are 3 TI blocks in one interleaving frame for this PLP |
| TIME_IL_TYPE | 0 | There is a 1:1 mapping of interleaving frames to T2-frames for this PLP |
| IN_BAND_A_FLAG | 0 | In-band type A signalling is not used for this PLP |
| IN_BAND_B_FLAG | 0 / 0 | In-band type B signalling is not used for this PLP (see note) |
| RESERVED_1 | 0000000000 | Reserved for future use |
| PLP_MODE | 00 / 10 | The PLP mode is signalled only in the BB-Header/This PLP uses HEM (see note) |
| STATIC_FLAG | 0 / 1 | The dynamic signalling may vary from frame to frame/The scheduling for the current PLP will only change at a configuration change (see note) |
| NO_PADDING_FLAG | 0 / 1 | BBFRAME padding may be used for this PLP/BBFRAME padding will not be used for this PLP (see note) |
| RESERVED_2 | 0x00000000 | Reserved for future use |
| NOTE: | A number of backward-compatible changes were introduced from version 1.2.1 of [i.1]. In those cases the first value and description given is that which would be transmitted by a legacy modulator that does not use these features. The second value given is the recommended value for this scenario which should be transmitted by newly designed or upgraded equipment. | |

Table 20 gives a worked example of the signalling fields of the L1-post dynamic signalling.

Table 20: An example of the contents of the L1-post dynamic signalling fields

| L1-post dynamic field | Contents | Explanation |
|-----------------------|--------------------------|--|
| FRAME_IDX | 00000000 | This frame is the first frame of the super-frame |
| SUB_SLICE_INTERVAL | 000000000000000000000000 | There are no type 2 PLPs so this field is set to '0' |
| TYPE_2_START | 000000000000000000000000 | There are no type 2 PLPs so this field is set to '0' |
| L1_CHANGE_COUNTER | 00000000 | There are no changes foreseen in L1 configurable parameters |
| START_RF_IDX | 000 | TFS is not used so this has no meaning |
| RESERVED_1 | 00000000 | Reserved for future use |
| PLP_ID | 00000000 | The ID of this PLP is '0' |
| PLP_START | 000000000000000000000000 | The cell index for the start of this PLP is '0', i.e. it starts in the P2 symbol immediately following the L1 data |
| PLP_NUM_BLOCKS | 0011001010 | The number of FEC blocks for one Interleaving Frame for this PLP is 202 |
| RESERVED_2 | 00000000 | Reserved for future use |
| RESERVED_3 | 00000000 | Reserved for future use |

The CRC₃₂ for the above combination of L1-post configurable and dynamic fields is 0x61014CB3.

Since all of the bits of the L1-post signalling fit within 1 16 K LDPC block, there is no L1 padding.

NOTE: The L1 padding is not included in the calculation of the CRC₃₂.

8.10.6 In-band L1-dynamic signalling

In Input Mode B, the PADDING field in BBFRAME may be used to carry in-band signalling.

Two types of in-band signalling are defined: type A and type B. Type A was defined in version 1.1.1 of [i.1], whilst type B was introduced from version 1.2.1 of [i.1] in a backward-compatible way. Type B is optional and receivers according to version 1.1.1 of [i.1] should ignore it because it will come after any type A block and will therefore appear to be the RESERVED₄ field of the type A block. If type A signalling is not used it will simply appear to be padding.

8.10.6.1 Generation of the in-band type A signalling message

In-band type A signalling blocks carry L1 update information and co-scheduled information. When the IN_BAND_FLAG field in the L1-post signalling, defined in clause 7.2.3.1 of [i.1], is set to '0', the in-band type A is not carried in the PADDING field. The use of in-band type A is mandatory for PLPs that appear in every T2-frame and for which one Interleaving Frame is mapped to one T2-frame (i.e. the values for P_I and I_{JUMP} for the current PLP are both equal to 1).

The L1 dynamic signalling for Interleaving Frame $n + 1$ of a PLP or multiple PLPs is inserted in the PADDING field of the first BBFRAME of Interleaving Frame n of each PLP. If NUM_OTHER_PLP_IN_BAND = 0 (see below), the relevant PLP carries only its own in-band L1 dynamic information. If NUM_OTHER_PLP_IN_BAND > 0, it carries L1 dynamic information for other PLPs as well as its own information. This can be used to reduce the time to switch between services carried in different PLPs of the same T2-system.

Figure 45 illustrates the signalling format of the PADDING field when in-band type A is delivered.




Figure 45: PADDING format at the output of the STREAM ADAPTER for in-band type A

Table 21 indicates the detailed use of fields for in-band signalling, consisting of 3 parts: a common part, the dynamic signalling for the current PLP (i.e. the PLP carrying the in-band signalling), and dynamic signalling for other PLPs (see details in clause 5.2.3 of [i.1]). The in-band signalling information is similar to that carried in the P2, (see clause 8.10.1), but there are some differences:

- The L1-change indication is only active if an upcoming change in the configuration affects the current PLP or one of the other PLPs being signalled in this in-band signalling field. By contrast, the change indication in the L1-post signalling occurs only once and indicates a change affecting *any* of the PLPs.
- There are differences in exactly which frames the dynamic signalling applies to. See clause 8.10.6.3.

Table 21: Padding-field mapping for in-band type A

| Field | Size |
|--|--|
| PADDING_TYPE | 2 bits |
| PLP_L1_CHANGE_COUNTER | 8 bits |
| RESERVED_1 | 8 bits |
| For $j=0..P_I-1$ { | |
| SUB_SLICE_INTERVAL | 22 bits |
| START_RF_IDX | 3 bits |
| CURRENT_PLP_START | 22 bits |
| RESERVED_2 | 8 bits |
| } | |
| CURRENT_PLP_NUM_BLOCKS | 10 bits |
| NUM_OTHER_PLP_IN_BAND | 8 bits |
| For $i=0..NUM_OTHER_PLP_IN_BAND-1$ { | |
| PLP_ID | 8 bits |
| PLP_START | 22 bits |
| PLP_NUM_BLOCKS | 10 bits |
| RESERVED_3 | 8 bits |
| } | |
| NOTE: | Version 1.1.1 of [i.1] included a variable-length field RESERVED_4. Version 1.2.1 made use of this reserved field to allow in-band type B signalling in a backward-compatible way. |

8.10.6.2 Insertion of L1-dynamic information of other PLPs

If $NUM_OTHER_PLP_IN_BAND > 0$, the in-band signalling carries L1-dynamic information for other PLPs as well as L1-dynamic information for its own PLP. The L1-dynamic information for other PLPs can be carried only when the values of P_I and I_{JUMP} for the current PLP are both equal to '1'. In other words, $NUM_OTHER_PLP_IN_BAND$ can be larger than '0' only when the current PLP appears in every T2-frame and one Interleaving Frame is mapped to one T2-frame. The L1-dynamic information for other PLPs can be used for fast channel switching.

$NUM_OTHER_PLP_IN_BAND$, and the particular PLPs for which this information is carried, should therefore be chosen taking into account both the trade-off between signalling overhead and the advantage from fast channel switching.

A compromise way to select other PLPs is to group all PLPs in a T2 system into one or multiple subsets, which are mutually exclusive to each other. Each PLP in a subset carries L1-dynamic information of the other PLPs within the same subset. Note that the minimum subset size is 1 (a single PLP), which corresponds to each PLP delivering only its own L1-dynamic information. By choosing to group together PLPs which have something in common (e.g. from the same service provider), the channel switching between related services could be reduced.

8.10.6.3 Which frame is signalled for same PLP and other PLPs

The following points should be noted:

- The dynamic signalling for the current PLP always applies to the *next*, rather than the current, *Interleaving Frame*, i.e. the in-band L1 dynamic information for the current PLP for Interleaving Frame $n + 1$ is carried in Interleaving Frame n .
- The L1 dynamic signalling for the current PLP contains a loop. In many cases this loop will only have one entry, but in the case of multi-frame interleaving there will be one entry giving the dynamic parameters (PLP_START and SUB_SLICE_INTERVAL) for *each* T2 frame to which the next Interleaving Frame will be mapped.
- PLP_NUM_BLOCKS is a property of the entire interleaving frame so this appears outside the loop, applying to the next interleaving frame.
- The L1 dynamic signalling for *other* PLPs applies to the next *T2-frame*, rather than the next Interleaving Frame, but can only be present if the *current* PLP appears in every frame and does not use multi-frame interleaving. The signalling for other PLPs is in effect the same as the information that will appear in the P2 symbol of the next T2 frame.

- Since PLP_NUM_BLOCKS is a property of the Interleaving Frame, the same value may be signalled multiple times for other PLPs using multi-frame interleaving; this is also true of the signalling in P2.

Therefore, in the simplest case of mapping one Interleaving Frame to one T2-frame, the in-band signalling information sent in the current T2-frame both for the current PLP and for other PLPs applies to the next T2-Frame, thereby allowing the receiver to decode the wanted PLP in the next T2-frame without reception of P2 symbols in the next T2-frame. The heavy arrows in figure 46(a) show the relationship between the in-band signalling in the first FEC block of each Interleaving Frame and the frame to which it applies.

When one Interleaving Frame is mapped to multiple T2-frames, i.e. $P_1 > 1$, the in-band L1-dynamic signalling of a PLP carries only its own L1 dynamic information. In this case, since the L1 dynamic information within one Interleaving Frame spreads over multiple T2-frames, the in-band signalling (carried over multiple T2-frames) belonging to the Interleaving Frame n applies to the multiple T2-frames to which the Interleaving Frame $n + 1$ is mapped.

For PLPs not present in every T2-frame, the in-band signalling for the current PLP applies to the T2-frame or multiple T2-frames to which the next Interleaving Frame is mapped. If each Interleaving Frame is mapped to one T2-frame, the L1 dynamic information from the T2-frame corresponding to the Interleaving Frame n therefore applies to the T2-frame to which Interleaving Frame $n + 1$ is mapped.

In the case of interleaving over multiple T2-frames, the co-scheduled L1 dynamic signalling from the T2-frames corresponding to current Interleaving Frame provides the dynamic signalling as follows:

- PLP_START and PLP_INTERVAL for each of the T2-frames to which the next interleaving frame is mapped.
- PLP_NUM_BLOCKS for the next interleaving frame.

Figure 46(b) shows an example where both P_1 and I_{jump} are greater than 1.




Figure 46(a): In-band signalling for $P_1 = 1$, $I_{\text{JUMP}} = 1$, $N_{\text{TI}} = 1$




Figure 46(b): In-band signalling for $P_1 = 2$, $I_{\text{JUMP}} = 2$, $N_{\text{TI}} = 1$

8.10.6.4 Example of in-band signalling

An example of the contents of the padding field when carrying in-band signalling is given in table 22.

Table 22: Example of the use of the in-band signalling fields

| In-band signalling field | Contents | Explanation |
|---|------------------------|--|
| PADDING_TYPE | 00 | This is in-band signalling type A |
| PLP_L1_CHANGE_COUNTER | 00000000 | There are no changes foreseen in L1 configurable parameters which would affect this PLP |
| RESERVED_1 | 00000000 | Reserved for future use |
| The next four parameters are repeated only once as $P_i=1$ for this PLP | | |
| SUB_SLICE_INTERVAL | 0000000000111001110100 | The offset from the start of one sub-slice to the start of the next for this PLP is 0x000E74 cells |
| START_RF_IDX | 000 | TFS is not used so this has no meaning |
| CURRENT_PLP_START | 0000000000000000000000 | The cell index for the start of this PLP is '0', i.e. it starts in the P2 symbol immediately following the L1 data |
| RESERVED_2 | 00000000 | Reserved for future use |
| CURRENT_PLP_NUM_BLOCKS | 0001100100 | The number of FEC blocks for one interleaving frame for this PLP is 100 |
| NUM_OTHER_PLP_IN_BAND | 00000000 | No other PLPs are signalled by this in-band signalling |
| NOTE 1: The last bit of NUM_OTHER_PLP_IN_BAND is not necessarily the last bit of the BBFRAME, since the type A block might be followed by an in-band type B block and/or padding to provide padding. The number of such bits following the type A block can be deduced from the known length of the rest of the in-band type A field, the known value of K_{bch} , and the value of DFL signalled in the BB-header. | | |

8.10.6.5 In-band type B signalling

In-band type B signalling was introduced in a backward-compatible way in version 1.2.1 of [i.1]. It is used to give extra information related to the input streams, in particular allowing several pieces of ISSY information related to the first User Packet that begins in the BBFRAME.

The intention of in-band type B is to enable faster acquisition in second-generation receivers when receiving signals from transmitting equipment that implements it. This is particularly an issue when HEM is used, because each BBFRAME only carries one ISSY variable. As explained in clause 10.6.1, the output processing needs to know the TTO, BUFS and the transport stream rate in order to start outputting data; the transport stream rate itself needs to be calculated from two successive ISCR values. In HEM, this means that four BBFRAMEs need to be decoded. A type B block provides all of this information in the first BBFRAME of an Interleaving Frame.

Table 23, reproduced from convenience from [i.1], shows the signalling fields.

Table 23: Padding field mapping for in-band type B

| Field | Size |
|---------------------|---------|
| PADDING_TYPE ('01') | 2 bits |
| TTO | 31 bits |
| FIRST_ISCR | 22 bits |
| BUFS_UNIT | 2 bits |
| BUFS | 10 bits |
| TS_RATE | 27 bits |
| RESERVED_B | 8 bits |

TTO is signalled directly in units of T, rather than using the mantissa+exponent form used in the normal ISSY fields. BBFRAMEs carrying in-band type B signalling will always contain an ISSY TTO variable, but the type B field gives additional precision because it is signalled to the nearest multiple of T. Although the ISSY format apparently allows a greater range for TTO, in practice the larger exponent (TTO_E) values will never be needed. The 31 bit TTO field allows a value up to 234,8 s; this is nearly twice the longest allowed superframe duration.

FIRST_ISCR, BUFS, BUFS_UNIT and TS_RATE allow this information to be provided in the first BBFRAME, for the reasons discussed above.

Because it allows for faster acquisition, it is recommended that type B signalling should be supported by new designs and upgrades for transmitting equipment, and that new receiver designs should make use of it where available.

8.11 Layer-2 signalling

Layer 2 signalling is provided differently for MPEG-2 Transport Streams and Generic Streams.

DVB-T2 introduces the layer-1 concept of the "T2 System" (see clause 6.2.6). The mapping of Transport Streams to PLPs and the T2-system is signalled in the T2_delivery_system_descriptor transmitted in the NIT.

8.11.1 Transport Streams

For the Transport-Stream case a single new PSI/SI signalling element is used: the T2 delivery-system descriptor (T2dsd): see figure 47. It is a mandatory element of the Network Information Table (NIT) that is provided for all Transport Streams of a T2 system. For each Transport Stream a single instance of the descriptor will be provided, see figure 48.

The T2dsd maps a Transport Stream being signalled with the NIT, and heading the TS descriptor loop, to the corresponding T2 system (T2_system_id) and the PLP (PLP_id) that carries this Transport Stream within the T2 system. This mapping is handled by the upper part (blue part of figure 47) of the T2dsd. Each TS carried requires a separate instance of the T2 delivery-system descriptor, i.e. another loop cycle of the NIT's TS loop. The T2-system is uniquely identified by the combination of network_id in the NIT and t2_system_id in the t2dsd; these correspond directly to NETWORK_ID and T2_SYSTEM_ID in the L1-pre signalling (see clause 8.10.5).

The optional lower part of T2dsd (which appears only in the T2dsd long version) describes the physical parameters of the T2 system (orange part) and the cell parameters including carrier frequencies (grey part) used within each cell of the transmitter network. Since these parameters don't change within the same T2 system, the lower part, i.e. the T2 system and cell parameters, (orange and grey) occurs only once per T2 system. This means that a receiver needs to check the instances of the T2dsd until it finds the long version (if present at all) in order to get the information about the T2 system parameters and its cells and frequencies, see figure 48. The presence of both parts, the T2 system parameters (orange) and the cell parameters (grey), is optional.

| T2_delivery_system_descriptor(){ | | Mapping of TS to PLP and T2 system |
|----------------------------------|-----------|--|
| descriptor_tag | 8 uimsbf | |
| descriptor_length | 8 uimsbf | |
| descriptor_tag_extension | 8 uimsbf | |
| plp_id | 8 uimsbf | |
| T2_system_id | 16 uimsbf | |
| ----- | | |
| if (descriptor_length > 4){ | | |
| SISO/MISO | 2 bslbf | T2 system parameters (apart from T2_system_id) |
| bandwidth | 4 bslbf | |
| reserved_future_use | 2 bslbf | |
| guard_interval | 3 bslbf | |
| transmission_mode | 3 bslbf | |
| other_frequency_flag | 1 bslbf | |
| tfs_flag | 1 bslbf | |
| for (i=0; i<N; i++){ | | |
| cell_id | 16 uimsbf | Cell parameters |
| if (tfs_flag = 1){ | | |
| frequency_loop_length | 8 uimsbf | |
| for (j=0; j<M; j++){ | | |
| centre_frequency | 32 uimsbf | |
| } | | |
| } | | |
| else{ | | |
| centre_frequency | 32 uimsbf | |
| } | | |
| subcell_info_loop_length | 8 uimsbf | |
| for (k=0; k<L; k++){ | | |
| cell_id_extension | 8 uimsbf | |
| transposer_frequency | 32 uimsbf | |
| }}} | | |

Figure 47: T2 delivery system descriptor




Figure 48: Instances of T2 delivery system descriptor (within same NIT), example

The physical parameters provided for the T2 system (orange) reflect a full description of the signal on air, apart from its centre frequency/ies; these are an attribute of the cell (grey part) rather than the T2 system. New parameters, which were not present in the corresponding DVB-T descriptor, are the SISO/MISO field (see clause 9.3.3) and the 'tfs_flag'.

All PLP-related information (modulation/constellation, coding, time interleaving, start address etc.) is provided via and derived from layer-1 signalling, see clause 8.10.

The coding structure of the cell and frequency signalling is well known already from DVB-SI descriptors, except that there can be an extension signalled by the 'tfs_flag'. Single-profile receivers should therefore ignore any T2 signal for which 'tfs_flag' is set.

The following DVB-SI descriptors are not used for DVB-T2, either because they are not suitable for DVB-T2 and the T2dsd covers the related functionality, or to avoid ambiguities resulting from two ways of providing the same information:

- a) Cell frequency link descriptor.
- b) Cell list descriptor.
- c) Frequency list descriptor.

8.11.2 Generic Streams.

Generic Streams might occur in one of the following three variants:

- a) Generic Fixed Packet size Stream (GFPS).
- b) Generic Continuous Stream (GCS), also covering the case of a Generic Variable Packet size Stream.
- c) Generic Stream Encapsulated (GSE) according to TS 102 606 [i.9].

For cases a) and b) there is no layer 2 signalling definition available, for c) refer to TS 102 606 [i.9], clause 6.

8.12 Generation of the T2-MI stream

The generation of BBFRAMEs and their allocation to Interleaving frames, together with the generation of the L1 dynamic data for each frame has been described in the previous clauses. This provides all of the information required for generation of the T2-MI packets containing BBFRAMEs and L1 signalling.

Each time a BBFRAME is allocated, it can be inserted into a T2-MI packet, with the appropriate information in the header, and sent over the T2-MI interface immediately (see figure 35). As soon as allocation of BBFRAMEs is complete for all PLPs whose Interleaving Frames are mapped to a given T2-frame, the L1 signalling describing that T2-frame should be generated and transmitted in a T2-MI packet.

The T2-gateway should then generate a T2-MI packet containing SFN synchronisation information.

Implementers should ensure that the T2-MI packets are sent in the correct order and with the correct timing as specified in the T2-MI specification [i.1].

Care should be taken to ensure that the clocks in the T2 gateway and modulators are synchronised so that the frames signalled in the T2-MI packets are generated at the same rate as the modulator generates T2-frames, even in MFN operation.

8.13 Generation of the T2-MIP

If over-air distribution is used as described in clause 7.5, the T2-Gateway should also generate the T2-MIP. This indicates the transmission time for the T2 super-frame in which it is carried. The T2-gateway is the easiest place to generate this for two reasons. Firstly, the T2-gateway generates the SFN information contained in T2-MI SFN packets, used for line-fed transmitters, so will include the relevant functionality and timing information. Secondly, the T2-gateway is responsible for determining the mapping of data packets to BBFRAMEs and Interleaving Frames; consequently, earlier pieces of equipment in the DVB-T2 chain will not know for sure which superframe will carry a particular packet.

Unless the T2-gateway performs a re-multiplexing operation, the earlier multiplexing stages will need to allocate null packets at the appropriate rate, for the T2-gateway to replace with T2-MIPs. The insertion of the SFN information in the T2-MIPs should be done after the packets have been allocated to BBFRAMEs and Interleaving Frames, since the final superframe will not be known until this time. However, a dummy T2-MIP could be inserted earlier in the T2-gateway, if necessary to avoid the null packet being deleted by the null-packet deletion process.

8.14 L1 Bias Balancing

It was realised, after version 1.1.1. of [i.1] was published and receivers had been manufactured, that the lack of energy dispersal scrambling for the L1-signalling could, under some circumstances, result in large peaks in the time domain signal, during the P2 symbols. The underlying cause of this was the fact that, almost inevitably, there is a significant 'bias' in the L1-signalling - many more of the bits end up set to '0's than to '1's. This tends to occur for many of the fields, because large bit-widths have been allowed for, but often only very low numbers are signalled - e.g. the field `FIRST_FRAME_IDX` is an 8-bit field, but is very often set to 0, or perhaps to 1 or 2. Therefore this field will have a 'bias' of perhaps 6 or 8 (where bias is defined to be the number of '0's less the number of '1's). Furthermore, all of the bits of the fields "reserved for future use", in version 1.1.1 of [i.1], had default values of '0', contributing significantly to the problem.

Since the L1-signalling contains loops whose lengths are dependent on the number of PLPs, the more PLPs that are signalled, the more the bias tends to grow - with relatively few PLPs there is no significant problem. Without any energy dispersal scrambling, this bias ends up causing a large DC component to the modulated cells, and when the FFT is applied to create the time domain waveform, the bias appears as a large peak in the first 'bin' of the FFT - i.e. the first sample after the guard interval.

Without any other mitigation, the circumstances leading to an unacceptably large peak are:

- a relatively large underlying bias (because of a large number of PLPs to be signalled);
- the use of a low order L1 modulation (BPSK or QPSK);
- the use of one of the higher FFT sizes.

The size of the time-domain peak generated by this bias is given by the following formula (but note this is just the component due to the L1-bias - it will be modified by the DC in the rest of the signal):

$$V_{peak} = \frac{5 * bias}{N_{P2} * \sqrt{27 * K_{total}}}$$

(This will give the size of the peak in volts when the RMS of the rest of the signal is 1V). The variation of this with FFT size is illustrated in figure 49 for a bias of 1 000, such as would occur when around 9 PLPs are used with BPSK modulation and with no mitigation techniques.




Figure 49: Variation of P2 PAPR with FFT size

Four mitigation techniques have been adopted in version 1.2.1 of [i.1] to reduce the effect of the bias, and hence the peak and these will now be described in a little more detail. No one technique is likely to be sufficient on its own in all cases, and so it has been agreed to adopt all four techniques.

8.14.1 L1 Bias balancing Techniques

8.14.1.1 P2 Tone Reservation

The reserved tones are always available in the P2 symbols, and so even when TR PAPR (see clause 9.3.8.3) is not being applied to the rest of the signal, it is worthwhile to apply it to the P2 symbols alone. Since no peaks in the rest of the signal are being cancelled, it only really makes sense to try to cancel the peak due the L1 bias, and for this one iteration of the algorithm is sufficient, which reduces the complexity of the implementation.

The one disadvantage of this technique is that it causes a slight increase in the power of the P2 symbol.

8.14.1.2 L1 reserved bits and extension field

There are several fields in the L1 signalling "reserved for future use". Originally, these were set to '0's, but they could equally well be set to '1's. This is doubly effective since it reduces the number of '0's and adds '1's at the same time. The L1 signalling itself can be extended by adding an extension field, and this may be used for bias balancing as well. All of the reserved fields and the extension field are taken together and their bits set, usually to '1's. However, to allow for the possibility that the bias may be small or negative, a formula based on the actual bias is used to set these bits (see clause 7.2.3.7 of [i.1]). Also, it has been noted that it is possible to optimise the bias reduction by choosing different values for some of these bits - and so influencing the parity bits that are added by the coding - an exhaustive search of possible settings of some fields could be used to improve the bias reduction.

The L1 extension field has the rate 1/2 LDPC code applied to it, which means that for each bit in the extension field, in effect another (random) parity bit is added, which is inefficient. Generally it will be more efficient to use bias balancing cells, but the L1 extension field may be more efficient with FFT sizes of 8 K or smaller.

The use of these bits has no effect on the P2 symbol power, but the use of the extension field may reduce the available capacity.

8.14.1.3 Bias balancing cells

Additional cells may be inserted between the L1-signalling and the PLPs (see clause 8.3.6.3.1 of [i.1]). The sole purpose of these cells is to reduce the bias and so they are called bias balancing cells - they can be thought of as being a little like reserved tones, although their modulation is set directly depending upon the measured bias.

These cells are limited to a maximum amplitude of 1, and so cannot increase the power of the P2 symbols, but may even reduce it slightly. They may however reduce the available capacity.

Due to the way the cell mapping is done within the frame builder, when there are two or more P2 symbols (i.e. with FFT sizes of 8 K or lower), it is necessary to 'pad' the ends of the first $N_{P2}-1$ symbols with dummy cells. This could mean that, under some circumstances, the use of bias balancing cells is inefficient, and so it may be better to use the L1 extension field.

8.14.1.4 L1-ACE

A modified form of the ACE algorithm may be applied to the modulated cells of the L1 signalling. With this technique, the modification is calculated directly from the bias, to simplify the implementation (see clause 7.3.3.3 of [i.1]).

The use of L1-ACE will cause a slight increase in the power of the P2 symbols, but has no effect on the capacity.

Note that using the L1-ACE and P2 TR-PAPR algorithms together, with a careful choice of parameters, will generally mean that the increase in the power of the P2 symbols will be less than if either technique was used alone.

For example, for one particular test case, figure 50 shows the variation in the P2 power increase as the amount of L1-ACE is varied. The different lines represent the different amounts of voltage correction of the time domain peak achieved by the combined effects of the two techniques (and so as the amount of L1-ACE is varied, so too does the amount of TR-PAPR).




Figure 50: Increase in P2 power for combined P2-TR and L1-ACE

8.14.2 Approach to bias balancing

A good approach to bias balancing would comprise the following steps:

- 1) Set all of the reserved bits in the best possible way, either using the algorithm from [i.1] or a better one if available.
- 2) If appropriate to the FFT size, and capacity is available, add an L1 extension field. However, it will usually be more efficient to use bias balancing cells instead (and always in 16 K and 32 K, see clause 8.14.1.3).
- 3) If possible, measure an average value for the bias of the modulated cells of the L1 signalling, and skip to step 7. If this is not possible, for example in the initial stages of configuration of a mode, the bias may be estimated as follows.
- 4) The bias from the L1-pre signalling is fixed with the configuration. This may be calculated by adding the total number of 1's ($t_{\text{pre_uncoded}}$) in each of the fields - the bias is then $(200 - 2t_{\text{pre_uncoded}})$. A rough estimate for the coded bias is the same as the uncoded bias, although there will be some variation. Note that the modulated cells of the L1-pre will have exactly the same bias as for the coded bits, due the use of BPSK for the L1-pre.

- 5) The bias from the L1-post is mostly fixed with the configuration, although a few of the fields vary, the most significant ones being PLP_START and PLP_NUM_BLOCKS - average values for these may be estimated. If the total number of 1's is $t_{\text{post_uncoded}}$, and the total number of bits is $N_{\text{post_uncoded}}$, the uncoded bias is $(N_{\text{post_uncoded}} - 2t_{\text{post_uncoded}})$. Again the coded bias will be roughly the same. The modulated bias for BPSK is exactly equal to the coded bias, but for the higher order modulations it can be estimated. If the proportion of 1's in the coded L1-post is given by p , where $p = t_{\text{post_coded}}/N_{\text{post_coded}}$, and $t_{\text{post_coded}}$ is the total number of 1's in the coded bits of the L1-post, and $N_{\text{post_coded}}$ is the total number of L1-post bits after coding, the modulated bias, for each of the real and imaginary axes, is given roughly by the following formulae for each modulation option:

QPSK:

$$C_{\text{bias}} = \frac{N_{\text{post_coded}}}{2\sqrt{2}} [(1-p) - p] = \frac{N_{\text{post_coded}}}{2\sqrt{2}} (1-2p)$$

16-QAM:

$$C_{\text{bias}} = \frac{N_{\text{post_coded}}}{4\sqrt{10}} [3(1-p)^2 - 2p(1-p) - p^2]$$

64-QAM:

$$C_{\text{bias}} = \frac{N_{\text{post_coded}}}{6\sqrt{42}} [7(1-p)^3 - p(1-p)^2 - 3p^2(1-p) - 3p^3]$$

(These are obtained from the probability for each combination of bits multiplied by the total bias for all states with that combination, assuming the allocation of bits to constellation states was random. So for example, in 64-QAM, the total real component of all columns with two '1's and a single '0' in the even

numbered bit positions is $\frac{-3}{\sqrt{42}}$ and the probability of obtaining each of these states is $p^2[1-p]$).

- 6) The overall bias is simply the sum of the bias for the L1-pre and the complex bias for the L1-post. Note that this means that in general, for a large number of PLPs, the bias in the real will usually be larger than the bias in the imaginary.
- 7) If capacity is available and it was not appropriate to use the L1 extension field, add bias balancing cells instead - ideally the number required is given by:

$$N_{\text{biasCellsActive}} = \frac{|C_{\text{bias}}|}{N_{\text{P2}}}$$

where $N_{\text{biasCellsActive}}$ is the number of bias balancing cells per P2 symbol.

If the actual number of bias balancing cells that can be added is less than this ideal number, there will be a residual bias, which is given by:

$$C_{\text{bias_residual}} = C_{\text{bias}} - N_{\text{P2}} N_{\text{biasCellsActive}} \frac{C_{\text{bias}}}{|C_{\text{bias}}|}$$

- 8) If the bias will not be balanced from these measures, carefully choose the parameters which control the L1-ACE and TR-PAPR to minimise the increase in power from the P2 symbols. This can be done as follows:

Decide the amount of correction V_{corr} which the L1-ACE and TR-PAPR should provide. Ideally this would be the full size of the expected peak given by:

$$V_{\text{peak}} = \frac{5C_{\text{bias_residual}}}{N_{P2}\sqrt{27K_{\text{total}}}}$$

However, the actual bias will vary from frame to frame, as will peaks due to variations in the data. In principle, the L1-ACE and TR levels should therefore be varied similarly, but this is neither practical nor necessary. Also, it is not necessary to eliminate the full peak - it simply needs to be reduced to an acceptable level, typically around 11 dB (or 3,55V). So, if $|V_{\text{peak}}| > 3,55$, the amount of correction required would be set to:

$$V_{\text{corr}} = (V_{\text{peak}} - 3,55) \frac{V_{\text{peak}}}{|V_{\text{peak}}|}$$

and the value of V_{clip} for the TR algorithm is also set to 3,55.

For BPSK, the minimum power increase is then achieved when:

$$C_{\text{L1_ACE_MAX}} = \max\left(\frac{V_{\text{corr}}\sqrt{27K_{\text{total}}} + 5N_{\text{L1_ACE}}}{5(N_{\text{TR}} + N_{\text{L1_ACE}})} - 1, 0\right), \text{ and the expected amplitude of the TR cells is}$$

$$C_{\text{TR}} = \max\left(\frac{V_{\text{corr}}\sqrt{27K_{\text{total}}} - 5N_{\text{L1_ACE}}C_{\text{L1_ACE_MAX}}}{5N_{\text{TR}}}, 0\right), \text{ and the power increase of the P2 symbols is :}$$

$$P_{\text{inc}} = \frac{N_{\text{L1_ACE}}C_{\text{L1_ACE_MAX}}(C_{\text{L1_ACE_MAX}} + 2) + N_{\text{TR}}C_{\text{TR}}^2}{K_{\text{total}}}$$

where $N_{\text{L1_ACE}}$ is the number of L1 cells available for correction by the L1-ACE algorithm, and is given by $t_{\text{pre_coded}} + t_{\text{post_coded}}$, where $t_{\text{pre_coded}} = (1840 - [200 - 2t_{\text{pre_uncoded}}])/2 = 820 + t_{\text{pre_uncoded}}$.

For higher order constellations, the above calculation appears to give a reasonable first estimate, by treating the real and imaginary components separately, and calculating an average value for $C_{\text{L1_ACE_MAX}}$ from the two components. Note that the value of $N_{\text{L1_ACE}}$ can be estimated for each component from the proportion of ones in the L1 signalling as follows:

$$N_{\text{L1_ACE_REAL}} = (820 + t_{\text{pre_uncoded}}) + p_{\text{ACE}}N_{\text{post_coded}} \quad \text{for the real component}$$

$$N_{\text{L1_ACE_IMAG}} = p_{\text{ACE}}N_{\text{post_coded}} \quad \text{for the imaginary component}$$

where

$$p_{\text{ACE}} = \begin{cases} p & \text{for QPSK} \\ p(1-p) & \text{for 16-QAM} \\ p(1-p)^2 & \text{for 64-QAM} \end{cases} \quad \text{is the probability, for each component, of a state being available for correction by the L1-ACE algorithm, assuming the bias is positive}$$

As a worked example, consider a configuration where the total number of '1's in the L1-pre is 50, and the estimated average for the L1-post is 4 185, out of a total of 12 132 uncoded bits. The total number of bits after coding is 28 160. The L1-post constellation is 16-QAM, making 7 040 L1-post cells. The reserved bits of the L1-signalling have all been set to 1's, but there is no capacity for an L1 extension field or bias balancing cells. The FFT size is 16 K, so $N_{P2} = 1$ and $K_{\text{total}} = 13 921$.

So: $t_{\text{pre}} = 50$; $t_{\text{post}} = 4 185$; $N_{\text{post_uncoded}} = 12 132$; $N_{\text{post_coded}} = 28 160$.

From step 4, the bias due the L1-pre before coding is $200 - 2 \times 50 = 100$. It can be estimated that both the coded bias and the bias of the modulated L1-pre cells will be the same.

From step 5, the uncoded L1-post bias is $12\,132 - 2 \times 4\,185 = 3\,762$.

Assuming the coded bias is roughly the same, the expected value of $t_{\text{post_coded}}$ is $(28\,160 - 3\,762)/2 = 12\,199$.

So, again from step 5, the expected value of $p \approx 12\,199/28\,160 = 0,433$.

So the expected bias from each of the real and imaginary components of the modulated cells of the L1-post is given by the formula in step 5 and is 634,6.

The expected total bias is therefore $100 + 634,6 + 634,6i = 734,6 + 634,6i$.

No bias balancing cells can be added, so from step 8, the expected voltage peak in the time domain will be $6,0 + 5,2i$. So a suitable value for V_{corr} might be $3,3 + 2,9i$. The expected number of constellation states available for L1-ACE will be $(820+50) + (0,25 \times 7\,040) = 2\,630$ for the real and $(0,25 \times 7\,040) = 1\,760$ for the imaginary (note there are some states that count twice).

An estimate for the optimum value of $C_{\text{L1_ACE_MAX}}$ for the real is therefore 0,09 and for the imaginary 0,11, and an overall average of 0,10 is chosen.

9 Modulators

This clause assumes that the modulator is fed with a T2-MI stream as defined in [i.20], generated by the T2 gateway. In some applications, some or all of the functions of the T2 gateway may be included in the modulator. See clause 8 for details of those functions.

9.1 Operation with T2-MI input stream

This clause describes how a modulator might work when operating with a T2-MI input stream. Details of the individual processing stages are given later in the later clauses of clause 9, whereas this clause discusses how the various blocks work together in terms of delays and relative timings. This clause is subject to change depending on decisions in the T2-MI standardisation process.

9.1.1 Overview

Figure 51 is a top-level block diagram of a modulator operating with T2-MI input. Firstly the headers of the incoming T2-MI packets should be decoded to determine which of the defined packet types it is.

The majority of packets will contain BBFRAMES; in this case the PLP index should also be decoded from the header and used to send the BBFRAME to the BICM chain for the appropriate PLP. The BBFRAMES belonging to one Interleaving Frame should be buffered (in the I/L frame delay) whilst awaiting the L1 signalling describing that Interleaving Frame itself (see clause 9.1.2). This is needed since some BICM stages, in particular the cell and time interleavers, need the dynamic information in order to divide the Interleaving Frame into TI-blocks.

T2-MI packets containing auxiliary stream data should be delayed by one T2-frame to compensate for the delays in the PLPs and fed directly to the Frame Builder.

For T2-MI packets containing L1 signalling, the information should be extracted and fed immediately to the BICM and Frame builder. The Frame Builder will use the L1 information to determine which PLPs should be mapped to which cells of the frame, and will also insert the coded and modulated L1 signalling into the P2 symbols.

T2-MI packets containing SFN data should be sent to the OFDM output stages to control the exact output time.




Figure 51: Processing of T2-MI input stream in the T2 modulator

9.1.2 Buffering and timing relationships

When a T2-gateway is used, BBFRAMES are formed and immediately sent to the modulators via T2-MI as soon as they are ready. Delays are required in the modulator both to hold up the data until the scheduling has been performed, and also to enable the L1 signalling, both in P2 and in-band, to be inserted at the appropriate moment (see clause 9.1.3).

Figure 52 illustrates the timing relationships for a T2-system including multiple PLPs, with time running horizontally. At the top the "Input T2-MI" stream is shown, including all the T2-MI packets. Below this, the coloured boxes show what is happening for three of the PLPs; below these the packets for L1-current signalling (cross-hatched) and L1-future (horizontal stripes) are shown, and at the very bottom are the T2-frames.


Considering first the box for PLP 1: this is a type-2 PLP with $P_i=1$; it has one or more TI-blocks per Interleaving Frame and each Interleaving Frame is mapped to one T2-frame. At the top are the incoming T2-MI packets for this PLP, i.e. at the output of the second commutator of figure 51. The T2-MI packets are coloured according to the `frame_idx` in the T2-MI header, which for this PLP indicates directly the T2-frame in which they will be carried. As these packets arrive, they are stored in the Interleaving Frame delay. Once the allocation for e.g. T2-frame 1 (red) has been completed, the T2-gateway will send an T2-MI packet containing L1-current information describing that frame (the red cross-hatched packet). The BBFRAMES for this Interleaving/T2-Frame can now be encoded by the BICM stages, based on the information in the L1 signalling packet, and written into the Time Interleaver memory as shown.

The Time Interleaver memory provides a further T2-frame of delay, and this is required to allow both in-band signalling and L1-repetition to be inserted, as will be explained below.

PLP1 is a type-2 PLP, so the Interleaving Frame is carried in multiple sub-slices per T2-frame; these are also shown in the diagram.

PLP2 uses multi-frame interleaving over two T2-frames; an Interleaving Frame therefore lasts for two T2-frame durations. The `frame_idx` field in the T2-MI packet headers indicates the *first* T2-frame to which the Interleaving Frame is mapped. This is illustrated by the colours of the T2-MI packets for PLP2: they are turquoise, green or purple corresponding to a `frame_idx` value of 0, 2 or 4. The incoming T2-MI packets for a given Interleaving Frame are stored up in the Interleaving Frame delay until the allocation is complete and the relevant L1-current T2-MI packet is received. Note carefully how the green T2-MI packets with a `frame_idx` of 2 finish at the same frame boundary as the green packets for PLP1, and just before the green L1-current signalling packet is received, and one T2-frame before T2-frame 2 is transmitted. As for PLP1, the single T2-frame of delay is provided by the Time Interleaver memory; this memory is filled up during one T2-frame duration but read out over two T2-Frames (one Interleaving Frame). Since this is a type 1 PLP, there is one slice per T2-frame as shown; each slice contains half an Interleaving Frame's worth of cells.

Finally, PLP3 uses the frame-skipping mechanism, with $I_{\text{jump}}=2$. Furthermore, `FIRST_FRAME_IDX` is equal to 1, i.e. PLP3 is mapped to T2-frames 1, 3 and 5. This corresponds to the `frame_idx` values in the T2-MI headers, as denoted by the colours of the packets in the figure (red, yellow and blue). As for PLP2, an Interleaving Frame has a duration of two T2-frames, and again allocation is complete (and L1-dynamic information available) one T2-frame before the first T2-frame carrying that Interleaving Frame.



NOTE: The PLPs in figure 52 belong to different PLP groups. The transmission delay for all three PLPs, as discussed in 8.6.1.2.5, is equal to T_F .

Figure 52: Timing of TS packets, BFRAMES and Interleaving Frames

9.1.3 Generation of L1-signalling from T2-MI packets

The L1-signalling packets serve two purposes: they are used by the modulator to provide the dynamic information needed to perform the BICM processing on the BFRAMES. They are also needed to provide the signalling bits to be inserted into the T2-frame.

The arrows at the bottom of figure 52 show how the L1 information is inserted into the P2 symbols of the T2-frames. The information describing a given T2-frame m is carried in the P2 symbols of T2-Frame m , as shown by the arrow marked "curr.". In addition, the dynamic information may optionally be "repeated" in the P2 of the previous frame, as explained in clause 7.2.3.3 of [i.1] and clause 8.10.2 of the present document. In that case, the relevant information is sent in an "L1-future" T2-MI packet. The arrow marked "next" indicates this. For example, T2-frame 1 will contain signalling for T2-frame 2, delivered by the red-striped L1-future packet; as can be seen this packet arrives just in time for the "next" frame information to be generated.

Figure 53 shows the contents of the two types of T2-MI packets containing L1 data for two consecutive T2-frames. The L1-current packet with frame_idx m contains all the information to describe the static, configurable and dynamic properties of frame m . However, this may not be all the information required to generate the L1-signalling for the P2 symbols of T2-frame m , because of the possibility of L1 repetition. The dynamic information for the next frame is not yet available and, where used, is carried in the subsequent L1-future packet, describing frame $m+1$.

As the figure shows, the full L1 signalling to be carried in the P2 symbols of T2-frame m is formed by taking the L1-pre, L1-configurable "dynamic, current frame" and extension fields from the L1-current T2-MI packet for frame_idx= m , together with the "dynamic, next frame", from the L1-future packet with frame_idx= m .

It may be noted that the "dynamic, next" in the L1-future packet for frame m and the "dynamic, current" in the L1-current packet for frame $m+1$ carry essentially the same information, because they describe the scheduling for the same frame ($m+1$). However, they are not always identical. This is because T2-frame m and T2-frame $m+1$ might belong to different superframes, and these superframes might contain a different set of PLPs. Since the entries in the PLP loops only contain PLPs present in the current superframe, the loops in the "current" and "next" dynamic parts can be of different lengths and/or contain different PLPs. Where both frames belong to the superframe, or they are in different superframes but both superframes contain the same PLPs, the dynamic information in these two locations will be identical.




Figure 53: Generation of the L1-signalling from the T2-MI packets

The modulator may also replace the fields CELL_ID (in the L1-pre) and FREQUENCY (in the L1-post) with its own specific values. The values may be configured locally or using the Individual Addressing mechanism in the T2-MI. Clearly all modulators in the same SFN should use the same values of these fields.

The CRCs for both L1-pre and L1-post signalling should be generated by the modulator, since these will depend on the values of fields that the modulator might replace.

The L1 information is also used to generate in-band signalling.

The in-band signalling describes the next *Interleaving* Frame, but in-band signalling is optional where the Interleaving Frame is more than one T2-frame. The discussion here and the timing relationships in figure 52 assume that in-band signalling is not used for PLPs 2 and 3, which use multi-frame interleaving and frame skipping respectively. If in-band signalling were used for these PLPs, there would need to be more delay than shown in the figure. Specifically, the FEC blocks would need to be stored in the Time Interleaver memory for a whole Interleaving Frame. This would not affect the memory requirements for the Time Interleaver, which is assumed to have two Interleaving Frames of memory (see clause 9.2.5). However, it would affect the end-to-end delay of the chain.

Inserting the in-band signalling is slightly complicated. The Time Interleaver stores the OFDM cells making up FEC blocks, following coding and modulation. However, since the first BBFRAME of the interleaving frame includes the in-band signalling, the corresponding FEC block cannot be generated until the dynamic information for the *next* interleaving frame is available (see clause 9.1.2). As can be seen from figure 52, this information does not arrive from the T2-gateway until just before the relevant frame is transmitted. One solution would be to reserve space in the TI memory for the first FEC-block (shown in white in figure 52), but generate it out of sequence after all the other FEC blocks, by which time the dynamic information will have been received. The rest of the data for that BBFRAME would need to be stored whilst awaiting the dynamic signalling.

In-band signalling is somewhat flexible regarding which other PLPs are signalled, and is therefore assembled in the T2-gateway to avoid different modulators from generating different in-band signalling. The in-band signalling for each PLP should therefore be taken from the T2-MI L1-future packet and inserted verbatim in the relevant BBFRAME.

9.1.4 Generation of the PLP data from the T2-MI input

Generating and encoding the PLP data for a given T2-frame from the T2-MI input packets is entirely deterministic. For a given T2-frame index, the modulator can determine from the configurable L1 signalling the following pieces of information:

- 1) Whether a given PLP is mapped to the T2-frame at all. The PLP is mapped to the frame if:

$$(FRAME_IDX - FIRST_FRAME_IDX) \bmod I_{JUMP} = 0.$$

- 2) The interleaving frame index within the superframe:

$$n = \left\lfloor \frac{FRAME_IDX}{P_I \times I_{JUMP}} \right\rfloor$$

- 3) The index of the T2 frame within the interleaving frame =

$$\frac{FRAME_IDX - FIRST_FRAME_IDX}{I_{JUMP}} \bmod P_I \text{ (provided test 1 was true).}$$

The modulator can use this information to determine which cells, if any, for a given PLP should be mapped to the T2-frame.

The modulator also needs to determine which Baseband Frames make up a particular Interleaving Frame n . These are the BBFRAMEs for which $\frac{frame_idx - FIRST_FRAME_IDX}{P_I \times I_{JUMP}} = n$, where $frame_idx$ is the field in the

T2-MI packet header. Note that because of the definition of $frame_idx$, the division will always give an integer result for n .

For a given PLP, the modulator can determine which BBFRAME is the first of an Interleaving Frame because the corresponding T2-MI packet for that PLP will have "intl_frame_start" bit set to 1. Subsequent BBFRAMEs will be received in the correct order.

9.1.5 Start-up behaviour

At start-up, or when the T2-MI input is restored following an outage, the modulator needs to wait until it has received a complete Interleaving Frame for each PLP. The worst case is that it just misses the beginning of one, so has to wait most of one Interleaving Frame to get the start of the next one, then receive all the BBFRAMEs of that Interleaving Frame. This worst-case time is just under two Interleaving Frames. The modulator can then generate start generating the slices and subslices for that PLP. The total start-up time is therefore twice the longest Interleaving Frame duration, i.e. the Frame duration times the maximum value of $(P_I \times I_{jump})$.

The following points should be noted:

- 1) The modulator does not need to wait until a T2-frame which is the start of an Interleaving Frame for all PLPs together. It can start generating slices and subslices after (at most) 2 interleaving frames, and the modulator can switch on as soon as all PLPs are up and running.
- 2) A corollary of (1) is that the modulator doesn't necessarily need to wait until the start of a superframe, unless one or more of the PLPs is interleaved over an entire superframe.
- 3) In the very worst case, the superframe duration will be 64 s (or 128 s if FEFs are used). However, in most circumstances it will be shorter. Reasons for having a long superframe are:
 - (a) some PLPs have long time interleaving;
 - (b) the PLPs have different time interleaving depths such that the superframe needs to be the lowest common multiple of all the interleaving depths;

(c) because occasional FEFs are used;

Only in case (a) will the modulator need to wait until the beginning of the superframe before it can generate the signal.

The effect of interleaving depth on startup time (or recovery time from a fault) is one of the factors that broadcasters and network operators need to consider when making their parameter choices.

9.1.6 Modulator synchronisation

9.1.6.1 Single-Frequency Network operation

When operating in Single Frequency Networks (SFNs) the need for modulators to be synchronised both in time and centre frequency is well known already from DVB-T and other COFDM systems. Typically a GPS receiver is used to provide timing signals in the form of a 10MHz frequency reference and a one-pulse-per-second (1PPS) signal.

T2-MI provides both absolute and relative timing modes (see [i.20]). In the absolute mode, a source of time-of-day is also required. In the relative mode, the 1PPS signal is sufficient.

9.1.6.2 Modulator test modes for SFN operation

A modulator test mode has been defined for operation in Single Frequency Networks (SFNs). In this mode, the first P1 preamble of each superframe is replaced by zero modulation at the modulator output (a "null P1 preamble"). This can be used in the laboratory to compare the timing of different modulators using the same T2-MI input. Figure 54 shows an example oscilloscope trace with two correctly aligned null P1 preambles.

NOTE: The test signal with null P1 preambles is not intended to be radiated on-air and receivers are not expected to be able to synchronise to or decode such a signal.




Figure 54: Modulator test mode showing correct relative timing of the null P1 symbols

The test signal alone only checks the relative timing of two or more modulators, but a method has been devised for checking the absolute timing. There are several sets of DVB-T2 parameters in which a superframe duration is exactly seven seconds or a multiple of it. These modes are particularly useful in testing modulators in an SFN configuration, because the T2-Gateway can arrange that every superframe begins on a second boundary (i.e. the relative part of the T2 timestamp is always zero).

Coupled with the test mode described, this allows the absolute timing to be checked. By triggering an oscilloscope on the null P1 preamble, the alignment relative to the 1PPS edge can be checked. An example oscilloscope trace for such a test is shown in figure 55. The small offset between the 1PPS and the start of the null P1 is due to some delay through the output filter of the modulator, Modulator manufacturers should take account of such delays when reading the T2 timestamp and computing the output timing to ensure that the alignment is correct.

Table 24 is an exhaustive list of modes with 7 s superframes, and also gives a 32 K mode with a 2 1s superframe (no 7 s mode exists in 32 K). Other modes with longer superframes exist, and 7 s modes can be created for other parameter combinations by the addition of FEF parts.

Table 24: DVB-T2 modes with a superframe duration of exactly 7 s or a multiple of 7 s

| FFT size | Guard Interval | L_F | T2-Frames per superframe | Superframe duration |
|----------|----------------|-------|--------------------------|---------------------|
| 1 K | 1/8 | 276 | 200 | 7s |
| 2 K | 1/8 | 138 | 200 | 7s |
| 2 K | 1/32 | 946 | 32 | 7s |
| 4 K | 1/8 | 69 | 200 | 7s |
| 4 K | 1/32 | 473 | 32 | 7s |
| 8 K | 1/128 | 242 | 32 | 7s |
| 16 K | 1/128 | 121 | 32 | 7s |
| 32 K | 1/16 | 22 | 250 | 21s |




Figure 55: Absolute timing test showing the zero P1 symbol in relation to the 1PPS edge

9.1.6.3 Multiple-frequency network operation

In Multiple Frequency Networks (MFNs), the need for absolute synchronisation of the transmitters in a network is less critical. In DVB-T, one method used is to lock the modulator to the incoming Transport Stream, since for a given mode this has a fixed relationship to the fundamental clock rate of the DVB-T system. The suitability of this method will depend on the stability of the incoming Transport Stream.

When operating with T2-MI input, this method is more problematic since the rate of the TS carrying the T2-MI can be chosen arbitrarily and does not have a fixed relationship to the modulator clock rate. It is recommended that modulators operating with T2-MI input should always use an external timing reference such as GPS, even when operating in an MFN.

9.2 Bit-interleaved coding and modulation

9.2.1 FEC for the PLPs

9.2.1.1 FEC overview

The Forward Error-Correction (FEC) scheme in DVB-T2 is based on the concatenated BCH and LDPC code of DVB-S2, with the addition of a bit interleaver and demux for mapping bits to QAM constellation cells (see figure 56).




Figure 56: DVB-T2 FEC scheme

The BCH outer code is included as an insurance against unwanted error floors at high C/N ratios, and is exactly the same as DVB-S2.

For the LDPC code, DVB-T2 includes two choices of code length ($N_{ldpc} = 64\,800$ for normal blocks, and 16 200 for short blocks) and 6 choices of code rate (1/2, 3/5, 2/3, 3/4, 4/5, 5/6) from DVB-S2.

NOTE: The code rate of 1/4 for short code length is used only in L1 signalling.

The codes for normal blocks are identical to DVB-S2. The codes for short blocks are based on the DVB-S2 codes except rate 3/5, for which a new short code was chosen.

The bit interleaver and demux were newly designed for 16-QAM/64-QAM/256-QAM constellations. Another new feature, called a *parity and column-twist interleaver*, is included in DVB-T2; this improves performance in terrestrial environments, such as erasures caused by a 0 dB echo.

The performance of the short codes is some tenths of a dB worse than normal codes. However they can be used for low-bit-rate applications requiring shorter latency (see clause 5.11).

9.2.1.2 BCH code

The shortened BCH codes, which are defined on $GF(2^{16})$ for the normal and $GF(2^{14})$ for the short LDPC codes, are exactly the same as in DVB-S2. The error-correction capability is 10 bits or 12 bits depending on the LDPC code rate.

9.2.1.3 LDPC code with multi-level constellations

9.2.1.3.1 LDPC code

The LDPC codes of DVB-S2 have the following structure in their parity-check matrices:

- 1) Cyclic structure in information part.
- 2) Staircase structure in parity part.

The first characteristic can be used for the implementation of encoder and decoder hardware based on partly parallel processing architecture. The second characteristic can be used for generating parity bits by an accumulator.

In DVB-T2, 6 code rates of 1/2, 3/5, 2/3, 3/4, 4/5, and 5/6 are defined for both normal and short codes. Note that the real code rates of short codes ('effective LDPC rate' in table 5b of [i.1]) are in some cases slightly lower than the nominal rate ('LDPC code identifier' in the same table) (see table 26).

9.2.1.3.2 Bit Interleaver and demux

LDPC codes in DVB-T2 are irregular LDPC codes and the error-protection level of each code bit is not uniform, but depends on the column weight of the parity-check matrix. Table 25 and table 26 give the distribution of column weights, i.e. the number of columns having each weight, for the two FEC-block sizes. The protection level among bits in a multi-level-constellation symbol is not uniform, either. The performance of an LDPC code with multi-level constellation depends highly on the correspondence between code bits and constellation bits. For the purpose of controlling this correspondence, a bit interleaver and demux are required.

Table 25: Column-weight distribution of normal ($N_{ldpc} = 64\ 800$) codes

| code rate | column weight | | | | | | |
|-----------|---------------|--------|-------|--------|--------|--------|---|
| | 13 | 12 | 11 | 8 | 3 | 2 | 1 |
| 1/2 | | | | 12 960 | 19 440 | 32 399 | 1 |
| 3/5 | | 12 960 | | | 25 920 | 25 919 | 1 |
| 2/3 | 4 320 | | | | 38 880 | 21 599 | 1 |
| 3/4 | | 5 400 | | | 43 200 | 16 199 | 1 |
| 4/5 | | | 6 480 | | 45 360 | 12 959 | 1 |
| 5/6 | 5 400 | | | | 48 600 | 10 799 | 1 |

Table 26: Column-weight distribution of short ($N_{ldpc} = 16\ 200$) codes

| code rate | | column weight | | | | | |
|--------------|-----------|---------------|-------|-------|--------|-------|---|
| nominal rate | real rate | 13 | 12 | 8 | 3 | 2 | 1 |
| 1/2 | 4/9 | | | 1 800 | 5 400 | 8 999 | 1 |
| 3/5 | 3/5 | | 3 240 | | 6 480 | 6 479 | 1 |
| 2/3 | 2/3 | 1 080 | | | 9 720 | 5 399 | 1 |
| 3/4 | 11/15 | | 360 | | 11 520 | 4 319 | 1 |
| 4/5 | 7/9 | | | | 12 600 | 3 599 | 1 |
| 5/6 | 37/45 | 360 | | | 12 960 | 2 879 | 1 |

The bit interleaver is a block interleaver applied to each LDPC code word. In DVB-T2, a bit interleaver having $N_c = 2m$ columns is used for the 2^m -QAM constellation - except for 256-QAM with the short code, which uses $m = 8$ columns. LDPC-coded bits are written column-wise (with column-twisting: see below) and read out row-wise. Figure 57 shows the basic structure without column-twisting. The demux part demultiplexes $2m$ bits in the same row to compose two constellation symbols (except 256-QAM).




Figure 57: Basic structure of 2m-columns bit interleaver and demux

The interleaving part includes a new feature: the Parity and Column Twist (PCT) interleaver. The parity check matrices of LDPC codes of DVB-T2 have the diagonal structure depicted in figure 58 (after parity interleaving). If the block interleaver were applied directly, many constellation symbols would contain multiple code bits which were connected to the same check node. This would lead to performance degradation on channels with erasures.




Figure 58: Row-and-column-permuted parity-check matrix of rate 3/4 and 8-column bit interleaver

To avoid this performance degradation, the parity interleaver interleaves parity bits in such a way that the parity part of the parity-check matrix has the same structure as the information part (figure 58). Then the start-position for writing each column is twisted so as that multiple code bits which are connected to the same check node are not contained in the same constellation symbol (figure 59). These 2 interleavers can be implemented as a single interleaver because both interleavers permute bits in a code word.




Figure 59: Column twist bit interleaver

The effect of PCT has been tested on the memory less Rayleigh+Erasure+AWGN channel (see clause 14.1.8). Significant improvement was observed for 16-QAM rate 3/4, 64-QAM rates 3/4 and 5/6, and 256-QAM rate 3/4.

The demux pattern was chosen to optimize performance in an AWGN channel for each QAM constellation.

NOTE: The demux patterns for code-rate 3/5 of the normal-length code, and for code-rate 2/3 using 256-QAM and the normal-length code are different from those of other code rates.

For QPSK, no bit interleaver is used because simulations showed no performance improvement on an AWGN channel and no significant improvement on an erasure channel.

9.2.1.3.3 Constellation Mapping

The constellation mapping for QPSK, 16-QAM, 64-QAM is the same as for DVB-T. In addition, Gray mapping of 256-QAM is added in DVB-T2. Note that this is different to the QAM mapping in DVB-C, which is not Gray.

9.2.2 FEC for the L1 signalling

This clause gives an overview on protection of L1 signalling based on BCH and LDPC codes. In particular, this clause briefly introduces the shortening of the BCH information part and puncturing of the LDPC parity part according to the length of signalling information: these are L1-specific operations which are not performed on the PLP data.

9.2.2.1 Shortening of BCH Information part

Both L1-pre- and post-signalling bits are protected by a concatenation of a BCH outer code and an LDPC inner code. The L1-signalling bits are first BCH-encoded and the encoded word is then further protected by a shortened and punctured 16 K LDPC code. For the protection of the L1-pre signalling, a BCH code with $K_{\text{bch}} = 3\,072$ and a 16 K LDPC code with $K_{\text{ldpc}} = 3\,240$ are used (see clause 9.2.1.1). For the L1-post signalling, a BCH code with $K_{\text{bch}} = 7\,032$ and a 16 K LDPC code with $K_{\text{ldpc}} = 7\,200$ are used. The BCH-encoded word corresponds to LDPC information bits. That is, the output of the BCH encoder is the input of the LDPC encoder. Note that K_{ldpc} is the same as $(K_{\text{bch}} + 168)$ for L1 signalling since the number of BCH parity bits in BCH encoding, $N_{\text{bch_parity}}$ is fixed at 168 bits.

The length of L1 signalling bits has different values case by case. The length of L1-pre signalling, K_{pre} is fixed at 200, whereas that of L1-post signalling is variable. Furthermore, the L1-post signalling bits can be segmented into multiple blocks as described in clause 7.3.1 of [i.1]. A segmented L1-post signalling has a length less than or equal to $K_{bch} = 7\,032$.

For the shortening operation, an input parameter K_{sig} is first obtained from the lengths of L1-pre signalling and the segmented post-signalling as described in clause 7.3.1 in [i.1]. If K_{sig} is less than K_{bch} , the BCH information part is shortened for encoding. More precisely, K_{bch} BCH information bits are filled with K_{sig} signalling bits and $(K_{bch} - K_{sig})$ zero-padded bits for BCH encoding. Note that the shortening of BCH information may be regarded as the shortening of LDPC information since the LDPC information consists of BCH information bits plus parity bits.

All the BCH information bits are divided into $K_{ldpc}/360 (= N_{group})$ bit-groups which are $(N_{group} - 1)$ groups of 360 bits and 1 group of 192 bits as illustrated in figure 60. For example, $K_{bch} = 7\,032$ BCH information bits are divided into $K_{ldpc}/360 = 7\,200/360 = 20$ bit-groups which are 19 groups of 360 bits and 1 group of 192 bits. Note that K_{ldpc} is always a multiple of 360, which is a parameter for LDPC encoding in clause 6.1.2 in [i.1].

The shortening of the BCH information is performed on bit-group basis, that is, the positions of the $(K_{bch} - K_{sig})$ zero-padded bits for BCH encoding are allocated according to the bit-groups. For easy understanding, a concrete example of shortening and BCH encoding is given in the next clauses.




Figure 60: Format of data after LDPC encoding of L1 signalling

9.2.2.2 Example for shortening of BCH information

Assume that the pure L1-post signalling consists of 7 971 bits, i.e. $K_{post_ex_pad} = 7\,971$ and it is transmitted using 64-QAM. $K_{post_ex_pad}$ denotes the number of information bits of the L1-post signalling excluding the padding field as described in clause 7.2.3 of [i.1]. Then, the number of LDPC blocks needed for the protection of L1-post signalling, $N_{post_FEC_Block}$ is determined by:

$$N_{post_FEC_Block} = \left\lceil \frac{K_{post_ex_pad}}{K_{bch}} \right\rceil = \left\lceil \frac{7971}{7032} \right\rceil = 2$$

where $\lceil x \rceil$ means the smallest integer larger than or equal to x . Note that BCH code with $K_{bch} = 7\,032$ and 16 K LDPC code and $K_{ldpc} = 7\,200$ are used for the protection of L1-post signalling.

Next, the length of L1_PADDING field, $K_{L1_PADDING}$ is calculated as:

$$\begin{aligned} K_{L1_PADDING} &= \left\lceil \frac{K_{post_ex_pad}}{N_{post_FEC_Block}} \right\rceil \times N_{post_FEC_Block} - K_{post_ex_pad} \\ &= \left\lceil \frac{7971}{2} \right\rceil \times 2 - 7971 = 1. \end{aligned}$$

Then the final length of the whole L1-post signalling including the padding field, K_{post} is set as follows:

$$K_{post} = K_{post_ex_pad} + K_{L1_PADDING} = 7971 + 1 = 7972$$

These K_{post} post-signalling bits are divided into 2 ($= N_{post_FEC_Block}$) blocks consisting of K_{sig} bits which is calculated as:

$$K_{sig} = \frac{K_{post}}{N_{post_FEC_Block}} = \frac{7972}{2} = 3986$$

Each block, with information size of K_{sig} , is protected by a concatenation of BCH outer codes and LDPC inner codes. Here, 7 032 BCH information bits consist of a segmented block of 3 986 ($= K_{sig}$) bits and 3 046 ($= K_{bch} - K_{sig} = 7 032 - 3 986$) zero-padded bits. The positions for zero-padding are determined by the following procedure described in clause 7.3.2.1 of [i.1].

First, all 7 032 BCH information bits denoted by $\{m_0, m_1, \dots, m_{7031}\}$ are divided into 20 ($= N_{group} = K_{ldpc}/360$) groups as follows:

$$X_0 = \{m_0, m_1, m_2, \dots, m_{359}\}, \quad (360 \text{ bits})$$

$$X_1 = \{m_{360}, m_{361}, m_{362}, \dots, m_{719}\}, \quad (360 \text{ bits})$$

...

$$X_{18} = \{m_{6480}, m_{6481}, m_{6482}, \dots, m_{6839}\}, \quad (360 \text{ bits})$$

$$X_{19} = \{m_{6840}, m_{6841}, m_{6842}, \dots, m_{7031}\}. \quad (192 \text{ bits})$$

The bits in each bit-group correspond to the data field in ascending order as illustrated in figure 61.




Figure 61: Allocation of BCH information bits in Data field

Then, the shortening procedure is as follows:

Step 1) Compute the number of groups in which all the bits will be padded, N_{pad} as follows:

$$N_{pad} = \left\lfloor \frac{7032 - 3986}{360} \right\rfloor = \left\lfloor \frac{3046}{360} \right\rfloor = 8$$

Step 2) All information bits of 8 groups, $X_{18}, X_{17}, X_{16}, X_4, X_{15}, X_{14}, X_{13}$, and X_{12} , are be padded with zeros.

Step 3) For the group X_3 , 166 ($= 7 032 - 3 986 - 360 \times 8$) information bits in the last part of X_3 are additionally padded. In Steps 2) and 3), the choice of bit-groups is determined by table 35 in clause 7.3.2.1 in [i.1].

Step 4) Finally, 3 986 ($= K_{sig}$) information bits are sequentially mapped to bit positions which are not zero-padded in 7 032 BCH information bits, $\{m_0, m_1, \dots, m_{7031}\}$, by the above procedure.

If 3 986 information bits which do not correspond to zero-padded position are denoted by $\{s_0, s_1, \dots, s_{3985}\}$, the result of above shortening procedure can be presented as follows:

$$\begin{aligned} m_i &= s_i, & 0 \leq i < 1274, \\ m_i &= 0, & 1274 \leq i < 1800, \\ m_i &= s_{i-526}, & 1800 \leq i < 4320, \\ m_i &= 0, & 4320 \leq i < 6840, \\ m_i &= s_{i-3046}, & 6840 \leq i < 7032. \end{aligned}$$

Figure 62 illustrates the shortening of the BCH information part in this case, i.e. filling the BCH information-bit positions that are not zero-padded with 3 986 information bits.




Figure 62: Example of Shortening of BCH information part

9.2.2.3 BCH encoding

The K_{bch} information bits (including the $K_{bch} - K_{sig}$ zero padding bits) are first BCH encoded according to clause 6.1.1 in [i.1] to generate $N_{bch} = K_{ldpc}$ output bits ($i_0 \dots i_{N_{bch}-1}$).

9.2.2.4 LDPC encoding

The $N_{bch} = K_{ldpc}$ output bits ($i_0 \dots i_{N_{bch}-1}$) from the BCH encoder, including the $(K_{bch} - K_{sig})$ zero padding bits and the $(K_{ldpc} - K_{bch})$ BCH parity bits, form the K_{ldpc} information bits $I = (i_0, i_1, \dots, i_{K_{ldpc}-1})$ for the LDPC encoder. Note that for the protection of L1 signalling, the number of BCH parity bits, $N_{bch_parity} (= K_{ldpc} - K_{bch})$ is fixed as 168. The LDPC encoder systematically encodes the K_{ldpc} information bits onto a codeword A of size N_{ldpc} :

$$A = (i_0, i_1, \dots, i_{K_{ldpc}-1}, p_0, p_1, \dots, p_{M_{ldpc} - K_{ldpc}-1})$$

according to clause 6.1.2 in EN 302 755 [i.1].

9.2.2.5 Puncturing of LDPC parity bits

16 K LDPC codes with $K_{ldpc} = 3\,240$ and $K_{ldpc} = 7\,200$ for L1-pre- and post-signalling have 12 960 and 9 000 parity bits, respectively. Since the code length of 16 K LDPC code, N_{ldpc} is 16 200, their effective LDPC code rates (R_{eff}) are 1/5 and 4/9, respectively.

All $N_{ldpc} - K_{ldpc}$ LDPC parity bits, denoted by $\{p_0, p_1, \dots, p_{N_{ldpc} - K_{ldpc} - 1}\}$, are divided into Q_{ldpc} parity bit groups where each parity group is formed from a sub-set of the $N_{ldpc} - K_{ldpc}$ LDPC parity bits as follows:

$$P_0 = \{ p_0, p_{Q_{ldpc}}, p_{2Q_{ldpc}}, \dots, p_{359Q_{ldpc}} \},$$

$$P_1 = \{ p_1, p_{Q_{ldpc}+1}, p_{2Q_{ldpc}+1}, \dots, p_{359Q_{ldpc}+1} \},$$

$$P_2 = \{ p_2, p_{Q_{ldpc}+2}, p_{2Q_{ldpc}+2}, \dots, p_{359Q_{ldpc}+2} \},$$

$$P_{Q_{ldpc}-1} = \{ p_{Q_{ldpc}-1}, p_{2Q_{ldpc}-1}, p_{3Q_{ldpc}-1}, \dots, p_{360Q_{ldpc}-1} \}$$

where P_j represents the j th parity group and Q_{ldpc} is given in table 7b in EN 302 755 [i.1]. Each group has $(N_{ldpc} - K_{ldpc})/Q_{ldpc} = 360$ bits, as illustrated in figure 63. Note that the number of LDPC parity bits is always a multiple of 360 and Q_{ldpc} .




Figure 63: Parity-bit groups in an FEC block

When the shortening is applied to encoding of the signalling bits, some LDPC parity bits are punctured after the LDPC encoding. The punctured LDPC parity bits are not transmitted and they should be regarded as erasures in the receiver.

Similarly to the shortening of BCH information, the puncturing of LDPC parity bits is performed on parity-bit-group basis. Specifically, according to the number of LDPC parity bits to be punctured, N_{punc} , the positions of the parity bits are allocated by unit of parity-bit-groups. However, the method to obtain the value N_{punc} for L1-pre signalling is different to that for the post-signalling.

For L1-pre signalling, N_{punc} is calculated as follows:

$$N_{punc} = (K_{bch} - K_{sig}) \times \left(\frac{1}{R_{eff}} - 1 \right) = (K_{bch} - K_{sig}) \times 4$$

where K_{bch} denotes the number of BCH information bits, 3 072, and R_{eff} denotes the effective LDPC code rate 1/5 for L1-pre signalling. In the DVB-T2 system, the value of N_{punc} is fixed at 11 488 since K_{sig} is also fixed at $K_{pre} = 200$. Therefore, for L1-pre signalling, only 1 472 (= 12 960 - 11 488) LDPC parity bits are transmitted. After the shortening and puncturing, the encoded bits of the L1-pre signalling are mapped to:

$$(K_{sig} + N_{bch_parity}) \times \frac{1}{R_{eff}} = K_{sig} + N_{bch_parity} + 1472 = 1840$$

BPSK symbols where N_{bch_parity} denotes the number of BCH parity bits, 168 for 16 K LDPC codes.

For L1-post signalling, the procedure for calculating N_{punc} is more complicated as described in clause 7.3.1.2 in EN 302 755 [i.1]. The calculation is complicated because there is a choice of modulation, the length of the L1-post signalling can vary, and the number of L1 cells is constrained to be a multiple of N_{P2} so that there are the same number of L1 cells in each of the N_{P2} P2 symbols.

For easy understanding, a concrete example of calculating N_{punc} and puncturing of LDPC parity bits is given in the next clause.

9.2.2.6 Example for Puncturing of LDPC parity bits

Assume that the pure L1-post signalling consists of 7 971 bits and is transmitted over 64-QAM. Then, K_{sig} is 3 986 as calculated in clause 9.2.2.2. Additionally, assume that the number of P2 symbols of a given FFT size, N_{P2} is 2. That is, the FFT size of the symbol in the T2 frame is 8 K as shown in table 44 in clause 8.3.2 in [i.1].

For a given $K_{sig} = 3 986$, $N_{P2} = 2$ and modulation order 64-QAM, N_{punc} is determined by the following steps:

$$\begin{aligned} \text{Step 1)} \quad N_{punc_temp} &= \left\lfloor \frac{6}{5} \times (K_{bch} - K_{sig}) \right\rfloor = \left\lfloor \frac{6}{5} \times (7032 - 3986) \right\rfloor = 3655 \\ \text{Step 2)} \quad N_{post_temp} &= 3986 + 168 + 16200 \times (1 - 4/9) - 3655 = 9499 \\ \text{Step 3)} \quad N_{post} &= \left\lceil \frac{N_{post_temp}}{\eta_{MOD} \times N_{P2}} \right\rceil \times \eta_{MOD} \times N_{P2} = \left\lceil \frac{9499}{6 \times 2} \right\rceil \times 6 \times 2 = 9504. \quad (\eta_{MOD} = 6 \text{ for 64-QAM}). \\ \text{Step 4)} \quad N_{punc} &= N_{punc_temp} - (N_{post} - N_{post_temp}) = 3655 - (9504 - 9499) = 3650. \end{aligned}$$

Step 1 makes sure that the effective LDPC code rate of the L1-post signalling, R_{eff_post} is always less than or equal to $R_{eff_16K_LDPC_1_2}$ ($= 4/9$). Since R_{eff_post} is defined by $R_{eff_post} = (K_{sig} + 168) / N_{post}$, its value varies with K_{sig} as follows:

$$\begin{aligned} R_{eff_post} &= \frac{(K_{sig} + 168)}{N_{post}} = \frac{(K_{sig} + 168)}{(K_{sig} + 168) + 16200(1 - 4/9) - N_{punc}} \\ &\approx \frac{(K_{sig} + 168)}{(K_{sig} + 168) + 9000 - \left\lfloor \frac{6}{5} (K_{bch} - K_{sig}) \right\rfloor} \approx \frac{(K_{sig} + 168)}{(K_{sig} + 168) + 9000 - \frac{6}{5} (K_{bch} - K_{sig})} \\ &= \frac{(K_{sig} + 168)}{(K_{sig} + 168) + 9000 - \frac{6}{5} (K_{ldpc} - (K_{sig} + 168))} = \frac{(K_{sig} + 168)}{360 + \frac{11}{5} (K_{sig} + 168)}. \end{aligned}$$

As shown in figure 64, R_{eff_post} tends to decrease as the information length K_{sig} decreases. Furthermore, it is easily checked that $N_{post} \approx 360 + \frac{11}{5} (K_{sig} + 168)$. Note that if the multiplicative coefficient 6/5 in Step 1) had been 5/4, i.e. $N_{punc_temp} = \left\lfloor 5/4 \times (K_{bch} - K_{sig}) \right\rfloor$, then the approximate LDPC code rate would have been:

$$\frac{K_{sig} + N_{bch_parity}}{16200 - (K_{bch} - K_{sig}) - \left\lfloor 5/4 \times (K_{bch} - K_{sig}) \right\rfloor} \approx \frac{K_{sig} + 168}{16200 - 9/4 \times (K_{bch} - K_{sig})} = \frac{4}{9}.$$

In other words, the effective LDPC code rate would be almost constant regardless of the size of K_{sig} . However, a fixed code-rate is not the best choice, as will now be explained.

In general, for a fixed code-rate, the code performance is degraded as the information length decreases since the code length also decreases. On the other hand, for a fixed information length, it is clear that the code performance is improved as the code rate decreases. From these facts, the performance variation induced by the variable information length can be reduced by code-rate control. Step 1 for systematically calculating the number of puncturing bits is a simple code-rate control technique. That is, the rate control can be easily accomplished by adjusting the multiplicative coefficient in Step 1). The multiplicative coefficient 6/5 in Step 1) is chosen as the best value among those formed of $(B+1)/B$ for an integer B. In other words, the receiving coverage for the signalling information will be approximately invariant and stable, whatever the number of bits to be signalled.




Figure 64: Effective LDPC code rate $R_{\text{eff_post}}$ with $(N_{P2} = 2, \eta_{\text{MOD}} = 6)$

In Step 3), η_{MOD} denotes the modulation order and it is 1, 2, 4 and 6 for BPSK, QPSK, 16-QAM, and 64-QAM, respectively. Step 3) guarantees that N_{post} is a multiple of the number of columns of the bit interleaver (described in clause 7.3.2.6 in EN 302 755 [i.1]) and that $N_{\text{post}}/\eta_{\text{MOD}}$ is a multiple of N_{P2} . N_{post} means the number of the encoded bits for each information block. After the shortening and puncturing, the encoded bits of each block will be mapped to:

$$N_{\text{MOD_per_Block}} = \frac{N_{\text{post}}}{\eta_{\text{MOD}}} = \frac{9504}{6} = 1584$$

modulated cells. The total number of cells in all $N_{\text{post_FEC_Block}}$ blocks, $N_{\text{MOD_Total}}$ is

$N_{\text{MOD_Total}} = N_{\text{MOD_per_Block}} \times N_{\text{post_FEC_Block}} = 1584 \times 2 = 3168$. Note that L1_POST_SIZE (a pre-signalling field) will be set to $N_{\text{MOD_Total}}$.

Finally, for $N_{\text{punc}} = 3650$ given in Step 4), the positions of bits to be punctured are determined as follows:

Step 1) Compute $N_{\text{punc_groups}}$ such that $N_{\text{punc_groups}} = \left\lfloor \frac{3650}{360} \right\rfloor = 10$.

Step 2) All parity bits of 10 parity bit-groups $P_6, P_{15}, P_{13}, P_{10}, P_3, P_{17}, P_{21}, P_8, P_5, P_{19}$ will be punctured.

Step 3) For the group P_2 , $50 (= 3650 - 3600)$ parity bits in the first part of the group will be additionally punctured. That is, the parity bits $p_2, p_{Q_{ldpc}+2}, p_{2Q_{ldpc}+2}, \dots, p_{49Q_{ldpc}+2}$ will be punctured. In Steps 2) and 3), the choice of parity bit-groups is determined by table 37 in clause 7.3.2.4 in EN 302 755 [i.1].

9.2.2.7 Removal of zero-padding bits

The $(K_{\text{bch}} - K_{\text{sig}})$ zero-padding bits are removed and are not transmitted. This leaves a word consisting of the K_{sig} information bits, followed by the 168 BCH parity bits and $9\,000 - N_{\text{punc}} (= N_{\text{ldpc}} - K_{\text{ldpc}} - N_{\text{punc}})$ LDPC parity bits without N_{punc} punctured bits, as illustrated in figure 65.




Figure 65: Example of removal of zero-padding bits

9.2.2.8 Bit Interleaving and constellation mapping after shortening and puncturing

When 16-QAM or 64-QAM modulation is used for the L1-post signalling, the LDPC codeword of length N_{post} consisting of K_{sig} information bits, 168 BCH parity bits, and $(9\,000 - N_{\text{punc}})$ LDPC parity bits, are bit-interleaved using a block interleaver as described in clause 7.3.2.6 in [i.1]. Note that bit interleaving is not applied when BPSK or QPSK is used.

Each bit-interleaved LDPC codeword is mapped onto constellations as described in clause 7.3.3 in [i.1].

The modulated constellations may be modified by the L1-ACE algorithm, which is specified in clause 7.3.3.3 in [i.1]. This is also described further in clause 8.14.1.4 of the present document.

9.2.3 Rotated constellations

The technique of rotated constellations consists of applying a rotation to the QAM constellation followed by a component-axes interleaving.

After rotating, the projections of the constellation points on the I and Q channels carry the information regarding the m mapped bits. For a 16-QAM, instead of $2^{m/2}=4$ projections on each axis, the constellation now has $2^m=16$ projections as seen in figure 66. The insertion of the interleaving between I and Q leads to the same information being sent twice over the channel in different cells, as if an inner repetition code was used. The resulting "virtual" constellation after rotation and cyclic delay in the case of a 16-QAM is shown in figure 67. It is equivalent to sending a higher-order irregular QAM while having the spectral efficiency of the original 16-QAM. This leads to additional diversity that improves the error-correcting performance when severely faded channels are encountered.




Figure 66: 16-QAM constellation after rotation




Figure 67: Resulting "virtual" 16-QAM constellation after rotation and cyclic delay

The rotation applies the angles defined in clause 6.3 of [i.1]. The component-axes (or I/Q) interleaving is performed by using a cyclic delay before the cell- and time-interleavers. The optimum choice of the angle is dependent on modulation order, channel type and mapping type. For each modulation order (constellation size), a single corresponding rotation has been chosen. This can only be strictly optimum for a particular channel type, but nevertheless the values chosen still deliver a performance improvement (compared with non-rotated constellations) for all encountered channel models ranging from the classical faded channel (Rayleigh) to the severely faded channel (Rayleigh with erasures).

In order to have a deeper understanding of the rotated constellation technique, figure 68 and figure 69 show the projections as a function of the Gray mapping for the 16-QAM constellation.




Figure 68: rotated 16-QAM constellation and its projection over components I and Q showing bit 0 and bit 1 mappings

A clear advantage with respect to non-rotated 16-QAM is observed in the extreme case where a projection over one axis is completely erased. In fact, the detection of the transmitted constellation point is still possible since the projection over the remaining axis will generally have been transmitted on a different OFDM carrier and in a different OFDM symbol.




Figure 69: rotated 16-QAM constellation and its projection over components I and Q showing bit 2 and bit 3 mappings

On each projection axis (I or Q) there are 3 pairs of points with a small separation. However, the points that are closer together on, say, the I component, are greatly separated on the Q component (and vice versa) and thus do not represent neighbouring 2D constellation points. Consequently, they will generally have a lower pairwise error probability in all but the most extreme case, where one axis is essentially erased.

The previous analysis corresponding to the 16-QAM case is also valid for the 4-QAM, 64-QAM and 256-QAM schemes adopted in DVB-T2. Note however that in 256-QAM, the angle chosen is such that the projections on each axis are uniformly spaced.

Rotated constellations cannot be used in conjunction with the ACE PAPR reduction technique (see clause 9.3.8 regarding PAPR).

This technique does not show any known disadvantages, while greatly improving the robustness of the receiver to severe channel conditions. Consequently, we recommend its use at all times unless ACE PAPR reduction is in use.

9.2.4 Cell Interleaver

The Cell Interleaver (CI) is a Pseudo-Random permutation of the cells in the FEC codeword, different for each FEC codeword of one Time-Interleaver Block, to ensure in the receiver an uncorrelated distribution of channel distortions and interference along the FEC codewords. The permutation-address-generation algorithm is similar to the one used in the frequency interleaver/de-interleaver, as described in clause 9.3.2, but without the "wires permutation". In addition, in order to have different interleaving patterns for different FEC codewords, a constant shift (modulo N_{cells}) is added to the permutation, generated as the bit-reversed value of an incrementing N_d -bit counter, where $N_d = \lceil \log_2 N_{\text{cells}} \rceil$. A new shift value is generated for each FEC-block in the TI-Block. For both the permutation-address generation and the shift, values exceeding N_{cells} are discarded and the next value in the sequence generated.

Memory requirements of the Cell Interleaver are dependent on the Time Interleaver, and thus efficient implementation of the Time Interleaver results in memory efficiency in the Cell Interleaver.

9.2.5 Time Interleaving

The concept of the Interleaving Frame was described in clause 6.2.3, and the choice of interleaving parameters was discussed in clause 5.9. This clause will consider the implementation of the Time Interleaver.

It is most straightforward to consider the Time Interleaver as dealing with Interleaving Frames. This makes the interleaving delay, and therefore the required compensating delay (see clause 8.9), straightforward to calculate, and also allows the memory to be used to provide the extra frame's delay discussed in clause 9.1.2.

Incoming cells should be written in columns into the interleaver memory as shown in figure 70. Once the whole frame has been received at the input and written into the memory, the Interleaving Frame can be divided into TI-blocks, noting that the TI-blocks within an Interleaving Frame can contain slightly different numbers of FEC blocks, with the smaller TI-blocks coming first. The Interleaving Frame can then be read out one TI-block at a time, by reading the relevant portion of each row of the interleaver memory.

If the interleaving memory is used to provide the necessary extra frame's delay, there is a slight complication arising from the use of in-band signalling. This is discussed in clause 9.1.3.




Figure 70: Use of time-interleaver memory

9.3 Frame-building and OFDM generation

9.3.1 Frame building

The function of the frame builder is to assemble the cells produced by the time interleavers for each of the PLPs and the cells of the modulated L1 signalling data into arrays of active OFDM cells corresponding to each of the OFDM symbols which make up the overall frame structure. The frame builder operates according to the dynamic information produced by the scheduler and the configuration of the frame structure.

The scheduling of the PLPs is actually done already in the scheduler, which is assumed to be part of the T2-gateway (see clause 8.6). In order to generate the required L1 dynamic signalling information, the scheduler decides exactly which cells of the final T2 signal will carry data belonging to which PLPs. Although this operation has no effect on the data stream itself at this stage, the scheduler will define the exact composition of the frame structure. The scheduler does not change the data in the PLPs whilst it is operating. Instead, the data is buffered in preparation for frame building (see clause 9.1.2).

9.3.2 Frequency Interleaving

To minimise memory use the frequency interleaver may be implemented with a single memory comprising of N_{\max} locations. N_{\max} is the maximum number of data cells in an OFDM symbol of the given FFT size. It is envisaged that practical implementations of DVB-T2 modulators should support all the FFT sizes included in [i.1]. In such implementations, N_{\max} would be the memory capacity necessary to implement odd-even interleaving for 32 K. Thus, $N_{\max} = 27\,404$, this being the maximum number of data cells (C_{data}) in 32 K mode using PP8 Scattered Pilot pattern, extended bandwidth and no tone reservation (see table 42 of [i.1]). The frequency interleaver for all modes can thus be implemented with one memory having N_{\max} locations. For a given system configuration with $C_{\max} = \max(C_{\text{data}})$ then, this memory would need C_{\max} locations (for 32 K) or $2C_{\max}$ locations (for all other modes).

The following describes one way of implementing the interleaver.

For odd-even interleaving as used in 32 K, the interleaver writes data cells from even symbols (symbol number of form $2n$) into its memory in a permuted order, the permuted-order addresses $H(p)$ being provided by the relevant pseudo-random address generator from clause 8.5 of [i.1]. For odd symbols (symbol number of form $2n+1$) the data cells are written into memory in a sequential order. As only one memory is used in the interleaver and in order to avoid losing data cells, one cell of the previous symbol should be read out from the memory location to which the cell of the current symbol is to be written. It therefore follows that the sequence of write addresses for symbol number $2n+1$ should match the sequence of read addresses for symbol number $2n$ otherwise some data cells of symbol $2n$ will be overwritten. Suppose for the moment that all symbols contained $C_{\max} = C_{\text{data}}$ data cells. Then with symbol $2n$ at the input of the interleaver, the interleaver would follow the following steps:

- 1) $p = 0$.
- 2) Generate address $H(p)$.
- 3) Read cell p of outgoing interleaved symbol $2n - 1$ from location $H(p)$ of the memory.
- 4) Write cell p of incoming un-interleaved symbol $2n$ into location $H(p)$ of the memory.
- 5) Increment p .
- 6) if ($p < C_{\max}$) go to 2.

Then with symbol $2n + 1$ at the input of the interleaver:

- 7) $p = 0$.
- 8) Read cell p of outgoing interleaved symbol $2n$ from location p of the memory.
- 9) Write cell p of incoming un-interleaved symbol $2n + 1$ into location p of the memory.
- 10) Increment p .
- 11) if ($p < C_{\max}$) go to 8.

However, the T2-frame in DVB-T2 is composed of different types of symbol with correspondingly different numbers of data cells for each type of symbol. Thus for a given system configuration (FFT size, Scattered Pilot pattern, bandwidth extension and tone-reservation status), a P2 symbol typically carries fewer data cells than a data symbol which in turn carries more data cells than a frame-closing symbol. Imagine that the de-interleaver is about to process symbol $2n + 1$ (a data symbol) for which the preceding symbol $2n$ was a P2 symbol. Then since data symbols have more data cells than P2 symbols, it follows that the number of generated addresses in the loop 7 - 11 above would exceed the number of data cells carried by the P2 symbol. Therefore, the pseudo-random address generator would produce some addresses that are valid for writing cells of symbol $2n+1$ into memory but not valid for reading cells of symbol $2n$. The case in which symbol $2n$ has more data cells than symbol $2n + 1$ also occurs when going from a data symbol to a frame-closing symbol. For a given system configuration therefore $C_{\max} = \max(C_{\text{data}}, C_{\text{P2}}, C_{\text{FC}}) = C_{\text{data}}$. The loops above should be changed as follows, where the function $\text{DataCells}(n)$ returns the number of data cells in symbol n and HoldBuffer is a small amount of storage with write pointer $wptr$ and read pointer $rptr$:

- 1) $p = 0$; $wptr = rptr = 0$;
- 2) generate address $H(i)$;
- 3) $rdEnable = (H(p) < \text{DataCells}(2n - 1))$;
- 4) $wrEnable = (H(p) < \text{DataCells}(2n))$;
- 5) if ($rdEnable$) Read cell p of output interleaved symbol $2n - 1$ from location $H(p)$ of the memory;
- 6) store cell p of incoming un-interleaved symbol $2n$ into location $wptr$ of HoldBuffer and increment $wptr$;
- 7) if ($wrEnable$):
 - a) Write cell $rptr$ of HoldBuffer into location $H(p)$ of the memory and increment $rptr$.
 - b) If ($wptr == rptr$) reset both $rptr = wptr = 0$.

- 8) increment p ;
- 9) if ($p < C_{\max}$) goto 2.

Then with symbol $2n+1$ at the input of the interleaver:

- 10) $p = 0$; $wptr = rptr = 0$;
- 11) $rdEnable = (p < DataCells(2n))$;
- 12) $wrEnable = (p < DataCells(2n+1))$;
- 13) if ($rdEnable$) Read cell p of output interleaved symbol $2n$ from location p of the memory;
- 14) store cell p of incoming un-interleaved symbol $2n + 1$ into location $wptr$ of HoldBuffer and increment $wptr$;
- 15) if ($wrEnable$):
 - a) Write cell $rptr$ of HoldBuffer into location p of the memory and increment $rptr$.
 - b) If ($wptr == rptr$) reset both $rptr = wptr = 0$.
- 16) Increment p ;
- 17) if ($p < C_{\max}$) goto 11.

For modes other than 32 K that use odd-only interleaving, with symbol $2n$ at the input the interleaver steps are as follows:

- 1) $p = 0$;
- 2) Generate address $H_1(p)$;
- 3) $rdEnable = (H_1(p) < DataCells(2n - 1))$;
- 4) $wrEnable = (p < DataCells(2n))$;
- 5) if ($rdEnable$) Read cell p of output interleaved symbol $2n - 1$ from location $C_{\max} + H_1(p)$ of the memory;
- 6) store cell p of incoming un-interleaved symbol $2n$ into location $wptr$ of HoldBuffer and increment $wptr$;
- 7) if ($wrEnable$):
 - a) Write cell $rptr$ of HoldBuffer into location p of the memory and increment $rptr$.
 - b) If ($wptr == rptr$) reset both $rptr = wptr = 0$.
- 8) Increment p ;
- 9) if ($p < C_{\max}$) goto 2.

Then with symbol $2n + 1$ at the input of the interleaver:

- 10) $p = 0$;
- 11) Generate address $H_0(p)$;
- 12) $rdEnable = (H_0(p) < DataCells(2n))$;
- 13) $wrEnable = (p < DataCells(2n + 1))$;
- 14) if ($rdEnable$) Read cell p of output interleaved symbol $2n$ from location $H_0(p)$ of the memory;
- 15) store cell p of incoming un-interleaved symbol $2n + 1$ into location $wptr$ of HoldBuffer and increment $wptr$;

- 16) if (wrEnable):
 - a) Write cell $rptr$ of HoldBuffer into location $C_{\max} + p$ of the memory and increment $rptr$.
 - b) If ($wptr == rptr$) reset both $rptr = wptr = 0$.
- 17) Increment p ;
- 18) if ($p < C_{\max}$) goto 11.

The required width for each memory location depends on the resolution with which each cell is represented after QAM mapping and constellation rotation (if used).

Care should be taken to implement the interleaving function in the correct sense. As shown in the steps detailed above, the interleaver should work as follows:

- In 1 K, 2 K, 4 K, 8 K and 16 K, the interleaver should write to the memory in normal order and read in permuted order.
- In even symbols of 32 K, the interleaver should write to the memory in permuted order and read in normal order. The P2 symbol is an even symbol in 32 K, since numbering starts from zero.
- In odd symbols of 32 K, the interleaver should write to the memory in normal order and read in permuted order. The first normal symbol is an odd symbol in 32 K.

9.3.3 MISO processing

The MISO processing in the modulator is fairly straightforward. The processing is applied to pairs of cells of adjacent cell index following the frequency interleaver, so the storage requirements are minimal. Note that the cells of a pair might not end up on adjacent carriers once pilots have been inserted.

The carrier-mapping structure has been designed such that each symbol always contains an even number of data cells, so the MISO processing can always be carried out. In Frame-Closing Symbols, some of the data cells will be unmodulated (see clause 6.1.2.1), but this does not affect the MISO processing.

It is worth remarking that the MISO processing and pilot modification is such that the signal from transmitters in MISO group 1 is essentially unmodified. However, the choice of pilot pattern will generally be different, and there are extra pilots in both the P2 symbol and in the frame closing symbol.

Due to the different nature of the cells of an auxiliary stream used for transmitter signature, MISO processing should not be applied to these cells. There are therefore special provisions given in clause 8.3.7 of [i.1] to ensure that this is the case.

9.3.4 Pilot insertion

The choice of pilot pattern is dealt with in clause 5.4.

9.3.4.1 Purpose of pilot insertion

Pilots are inserted for several reasons:

- The scattered, P2 and Frame-Closing pilots can be used for channel estimation and equalisation (see clause 10.3.2).
- The continual, P2 and Frame-Closing pilots can be used for Common-Phase-Error correction (clause 10.3.2.1).
- All of the pilots can potentially be used for Synchronisation (clause 10.2.1).
- Continual, P2 and Frame-Closing pilots are also used as a form of 'padding'.

Examples of the use of pilots for padding are:

- To meet constraints on number of data cells: this is always even (to allow MISO) and (for data symbols) not a multiple of 5 (to avoid undesirable interactions with sub-slicing - see clause 5.8.1). The exact number of continual pilots, the extra Frame-Closing pilot on carrier $K_{\max} - 1$, and the extra Continual Pilots in extended-carrier modes are used for this purpose. As a result, each normal symbol of the frame contains an equal number of active cells.
- Extra P2 pilots are used to fill up the carriers in the extended-carrier region in the P2 symbols; these cannot carry useful data since the receiver might not know whether extended-carrier mode is used or not.
- To prevent MISO pairs being split, P2 pilots are added adjacent to reserved carriers in the P2 symbols.
- Extra P2 pilots are added at the edges of P2 in MISO mode to allow both the "sum" and "difference" to be interpolated up to the edge of the spectrum.

Although these pilots are primarily added for padding purposes, they are transmitted as pilots rather than being set to zero to allow receiver developers the opportunities to use them in improved algorithms.

9.3.4.2 Pilot locations

The pilot locations are described in detail in clause 9.2 of [i.1]. Particular care should be taken to note the following points:

- The different types of symbol contain different combinations of the different pilot types. Table 27, reproduced from [i.1], summarizes this.
- The Scattered Pilot pattern effectively begins on symbol zero, even though this is always a P2 symbol and so does not itself contain Scattered Pilots.
- In extended carrier mode the Scattered Pilots in the central portion of the signal fall on the same physical carriers (i.e. values of k') in a given symbol as in normal carrier mode. As a result, the values of k for the Scattered Pilots in a given symbol are not necessarily the same in normal and extended modes. In practice, the difference is only relevant in 8 K and 16 K mode with PP7 or PP8 (because in these modes $K_{\text{ext}} \bmod D_x D_y \neq 0$). In summary:
 - In normal carrier mode, the pilot pattern starts on carrier $k=0$ of symbol 0, i.e. the same as DVB-T.
 - In extended carrier mode with PP7 in either 8 K or 16 K, there is a Scattered Pilot on carrier $k=0$ in symbol 2 (coincident with the Edge Pilot).
 - In extended carrier mode with PP8 in either 8 K and 16 K, there is a Scattered Pilot on carrier $k=0$ in symbol 8 (coincident with the Edge Pilot).
 - In all other combinations of FFT size and Scattered Pilot pattern, the pilot pattern starts on carrier $k=0$ of symbol 0 in both normal and extended carrier modes.
- The locations of the continual pilots have been designed to meet a number of constraints whilst keeping the required lookup-table sizes to a minimum. Clause 10.3.2.2 explains how the calculation of the continual-pilot locations can be done efficiently, particularly in a receiver.
- There are P2 pilots on every third carrier in all modes except for 32 K SISO. In 32 K SISO they are on every sixth carrier.
- The extended-carrier regions of the P2 signal are padded with P2 pilots.
- There is a Frame-Closing Pilot on every Scattered Pilot bearing carrier.
- There is an extra Frame-Closing Pilot on carrier $K_{\max} - 1$ in some combinations of Scattered Pilot pattern and FFT size in order to make the number of data cells even (this is needed for MISO, but is also true in SISO mode).

Table 27: Presence of the various types of pilots in each type of symbol (X=present)

| Symbol | PILOT TYPE | | | | |
|---------------|------------|-----------|------|----|---------------|
| | Scattered | Continual | Edge | P2 | Frame-Closing |
| P1 | | | | | |
| P2 | | | | X | |
| Normal | X | X | X | | |
| Frame-closing | | | X | | X |

9.3.4.3 Number of pilot cells of each type

Implementing the various types of pilot is potentially error-prone given the number of combinations and the various additional pilots added under different circumstances. This clause presents a number of tables intended to provide a cross-check for implementers.

9.3.4.3.1 Pilots in normal symbols

Table 28 gives the total number of Continual Pilots (CPs) in each combination of FFT size and Scattered Pilot pattern. This includes CPs which are sometimes replaced by Scattered Pilots. The letters 'N' and 'E' after the FFT sizes indicate normal and extended-carrier modes respectively.

Most but not all of the CPs are sited on carriers which sometimes carry Scattered Pilots. Table 29 gives the number of CPs that are not on Scattered Pilot-bearing carriers, in each combination of FFT size and Scattered Pilot pattern.

Table 30 gives the number of Scattered Pilot cells in each symbol. A '+' indicates that there is one more SP for one symbol of the sequence, whereas a '-' indicates that there is one fewer SP on one symbol; similarly '++' indicates that there is one more SP on two symbols of the sequence.

Table 28: Total number of continual pilots for each FFT size and Scattered Pilot pattern

| | PP1 | PP2 | PP3 | PP4 | PP5 | PP6 | PP7 | PP8 |
|--------|-----|-----|-----|-----|-----|-----|-----|-----|
| 1 K | 20 | 20 | 22 | 20 | 19 | | | |
| 2 K | 45 | 42 | 42 | 43 | 42 | | 45 | |
| 4 K | 45 | 44 | 43 | 44 | 45 | | 50 | |
| 8 K N | 45 | 46 | 43 | 46 | 46 | | 53 | 47 |
| 8 K E | 45 | 50 | 45 | 48 | 46 | | 58 | 52 |
| 16 K N | 89 | 87 | 87 | 90 | 90 | 88 | 88 | 86 |
| 16 K E | 93 | 89 | 89 | 92 | 92 | 90 | 91 | 89 |
| 32 K N | | 175 | | 176 | | 176 | 180 | 175 |
| 32 K E | | 177 | | 178 | | 180 | 182 | 181 |

Table 29: Number of CPs on non-SP-bearing carriers

| | PP1 | PP2 | PP3 | PP4 | PP5 | PP6 | PP7 | PP8 |
|--------|-----|-----|-----|-----|-----|-----|-----|-----|
| 1 K | 5 | 5 | 5 | 4 | 5 | | | |
| 2 K | 22 | 17 | 19 | 18 | 17 | | 23 | |
| 4 K | 22 | 19 | 20 | 19 | 18 | | 25 | |
| 8 K N | 22 | 21 | 20 | 21 | 19 | | 26 | 32 |
| 8 K E | 22 | 25 | 22 | 23 | 19 | | 31 | 37 |
| 16 K N | 42 | 32 | 40 | 33 | 31 | 31 | 29 | 39 |
| 16 K E | 46 | 34 | 42 | 35 | 33 | 33 | 32 | 42 |
| 32 K N | | 36 | | 35 | | 31 | 33 | 48 |
| 32 K E | | 38 | | 37 | | 35 | 35 | 54 |

Table 30: Number of SPs per symbol

| | PP1 | PP2 | PP3 | PP4 | PP5 | PP6 | PP7 | PP8 |
|--------|--------|--------|------|--------|------|------|------|------|
| 1 K | 71- | 70+ | 35+ | 35 | 17++ | | | |
| 2 K | 142- | 141+ | 71- | 70+ | 35+ | | 17++ | |
| 4 K | 284- | 283+ | 142- | 141+ | 71- | | 35+ | |
| 8 K N | 568- | 567+ | 284- | 283+ | 142- | | 71- | 71- |
| 8 K E | 576- | 575+ | 288- | 287+ | 144- | | 72- | 72- |
| 16 K N | 1 136- | 1 135+ | 568- | 567+ | 284- | 283+ | 142- | 142- |
| 16 K E | 1 160- | 1 159+ | 580- | 579+ | 290- | 289+ | 145- | 145- |
| 32 K N | | 2 271+ | | 1 135+ | | 567+ | 284- | 284- |
| 32 K E | | 2 319+ | | 1 159+ | | 579+ | 290- | 290- |

9.3.4.3.2 Pilots in P2 symbols

Calculating the number of pilots in P2 symbols is complicated by the extended carrier mode, MISO and tone reservation carriers. Table 31 details the various numbers of cells of different types for each FFT size, normal/extended carrier mode, and SISO/MISO where relevant.

The first row gives the number of active carriers, equal to $K_{\max} + 1$. The next row gives the number of carriers at the edges that are filled entirely with pilots: this is zero in the normal carrier mode and $2 \times K_{\text{ext}}$ in the extended carrier mode. The next row gives the regular pilot spacing in SISO and MISO modes: this is 3 except for 32 K SISO when the value is 6. The next row gives the number of P2 pilots on this regular grid: this is the same in SISO and MISO except in 32 K because of the change in pilot spacing. Following this is the number of extra pilots in the padding region which are not on the regular grid.

The next three rows relate to the addition of P2 pilots adjacent to reserved-tone locations in MISO mode, to avoid separating cells in a MISO pair. Firstly the number of TR cells is listed. Some of the TR cells already occur in adjacent pairs, and the next row gives the number of such pairs. The number of P2 pilots added to act as TR partners is then given: there is one for every TR cell that is not in a pair, i.e. the value equals the number of TR cells minus twice the number of pairs.

In MISO mode, there are always four additional pilots added in what would otherwise have been the first and last pair of data cells; the next row gives this number.

The totals are then presented. For SISO this is the number of pilots on the regular grid plus the extra pilots in the extended region. In MISO the pilots partnering TR cells and the extra pilots near the edge are also counted. Finally the resulting number of data cells in each P2 symbol is tabulated for reference: these values are exactly the same as those in table 41 of [1.1]. Note that the capacity is the same in normal and extended carrier mode; this is required since the receiver does not know whether normal or extended mode is in use when it first decodes the P2.

Table 31: Number of P2 pilots in the different parameter combinations

| | 1 K | 2 K | 4 K | 8 K N | 8 K E | 16 K N | 16 K E | 32 K N | 32 K E |
|---------------------------------------|------|-------|-------|-------|-------|--------|--------|--------|--------|
| Active carriers | 853 | 1 705 | 3 409 | 6 817 | 6 913 | 13 633 | 13 921 | 27 265 | 27 841 |
| Padded carriers in extended mode | 0 | 0 | 0 | 0 | 96 | 0 | 288 | 0 | 576 |
| P2 pilot spacing | SISO | 3 | 3 | 3 | 3 | 3 | 3 | 6 | 6 |
| | MISO | 3 | 3 | 3 | 3 | 3 | 3 | 3 | 3 |
| Regular P2 pilots | SISO | 285 | 569 | 1 137 | 2 273 | 2 305 | 4 545 | 4 641 | 4 641 |
| | MISO | 285 | 569 | 1 137 | 2 273 | 2 305 | 4 545 | 4 641 | 9 089 |
| Extra P2 pilots in extended region | SISO | 0 | 0 | 0 | 0 | 64 | 0 | 192 | 480 |
| | MISO | 0 | 0 | 0 | 0 | 64 | 0 | 192 | 384 |
| Number of TR cells | 10 | 18 | 36 | 72 | 72 | 144 | 144 | 288 | 288 |
| Number of adjacent TR pairs | 1 | 1 | 1 | 1 | 1 | 9 | 9 | 8 | 8 |
| P2 pilots partnering TR cells in MISO | 8 | 16 | 34 | 70 | 70 | 126 | 126 | 272 | 272 |
| Extra P2 pilots near edges in MISO | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 |
| Total P2 pilots | SISO | 285 | 569 | 1 137 | 2 273 | 2 369 | 4 545 | 4 833 | 5 121 |
| | MISO | 297 | 589 | 1 175 | 2 347 | 2 443 | 4 675 | 4 963 | 9 941 |
| C_{P2} | SISO | 558 | 1 118 | 2 236 | 4 472 | 4 472 | 8 944 | 8 944 | 22 432 |
| | MISO | 546 | 1 098 | 2 198 | 4 398 | 4 398 | 8 814 | 8 814 | 17 612 |

9.3.4.3.3 Pilots in Frame-Closing Symbols

The pilots in the Frame-Closing symbols are relatively straightforward. There are Frame-Closing Pilots on all carriers which are a multiple of D_x , the Scattered Pilot-bearing carrier spacing, except for the edge carriers which are Edge Pilots.

NOTE: In SISO mode there is no difference between edge, scattered and Frame Closing pilots but there is a difference in the inversion pattern for MISO. See clause 9.3.4.5.1.

In addition there is an extra Frame-Closing Pilot on carrier $K_{\max} - 1$ in three cases. Table 32 gives the total number of Frame-Closing Pilots for each of the parameter combinations. These numbers do not include the Edge Pilots, but the three entries that are shaded are the three cases where the extra pilot is added and this has been included in the number.

Table 32: Number of Frame-Closing Pilots (not including Edge Pilots)

| | K_{total} | N_{TR} | PP1 | PP2 | PP3 | PP4 | PP5 | PP6 | PP7 | PP8 |
|-------------|--------------------|-----------------|-------|-------|-------|-------|-------|-------|-----|-----|
| 1 K | 853 | 10 | 283 | 141 | 141 | 71 | 71 | N/A | | N/A |
| 2 K | 1 705 | 18 | 567 | 283 | 283 | 141 | 141 | | 71 | |
| 4 K | 3 409 | 36 | 1 135 | 567 | 567 | 283 | 283 | | 141 | |
| 8 K | 6 817 | 72 | 2 271 | 1 135 | 1 135 | 567 | 567 | | 283 | |
| 8 K | 6 913 | 72 | 2 303 | 1 151 | 1 151 | 575 | 575 | | 287 | |
| 16 K | 13 633 | 144 | 4 543 | 2 271 | 2 271 | 1 135 | 1 135 | 567 | 567 | |
| 16 K | 13 921 | 144 | 4 639 | 2 319 | 2 319 | 1 159 | 1 159 | 579 | 579 | |
| 32 K | 27 265 | 288 | N/A | 4 543 | N/A | 2 271 | N/A | 1 135 | N/A | |
| 32 K | 27 841 | 288 | | 4 639 | | 2 319 | | 1 159 | | |

NOTE: For the highlighted combinations there is an extra Frame Closing pilot on carrier $K_{\max}-1$ (included in the number).

9.3.4.4 Pilot boosting

The boosting of the Scattered Pilots is a function of the Scattered Pilot pattern, with more-sparse patterns having a boost value which is much greater than in DVB-T.

The boosting of the continual pilots is a function of the FFT size; the larger FFT sizes have greater boosting but a lower density of CPs in order to reduce the overhead. The (amplitude) boosting value in 32 K, namely 8/3, represents the largest modulation value used in the DVB-T2 system.

Care should be taken in implementations to reserve sufficient headroom to represent these boosted pilots.

Edge Pilots have the same boost as Scattered Pilots because they are intended to be used for channel estimation.

9.3.4.5 MISO pilots

Pilots in MISO are different from SISO in three ways:

- Some pilots are inverted from transmitters in MISO group 2. This allows the sum and difference of the two channel responses to be estimated and the individual responses found.
- Additional pilots are added in the P2 symbol to meet the constraints.
- The Scattered Pilot pattern chosen for a given FFT size and guard interval generally needs to be twice as dense because only half of the pilots are used for each of the sum and difference channel estimates.

The first two points will be dealt with in the following clauses. The choice of pilot pattern for SISO and MISO was discussed in clause 5.4.

9.3.4.5.1 Inversion of pilots

The following pilots are inverted from transmitters in group 2:

- Scattered, Continual and Frame-Closing Pilots which fall on alternate Scattered Pilot-bearing carriers;
- P2 pilots on carriers which are alternate multiples of 3 or 6; and
- Edge Pilots on alternate symbols.

The definition of the inversion rules and the Scattered Pilot patterns is such that, for cells which meet the definition for both a Scattered Pilot and Edge Pilot, both inversion rules will give the same result. However, in the frame closing symbol, the rule for Edge Pilots can give a different result to that for Frame Closing Pilots. This occurs when there are an even number of symbols in a T2-frame. In such cases the rule for Edge Pilots should take precedence.

9.3.4.5.2 Additional P2 pilots

The additional pilots added in the P2 symbol are as follows:

- The two carriers between the first normal and the first inverted pilot at each end of the normal carrier region of the spectrum are P2 pilots. These are padding pilots allowing the sum and difference estimates to be made without extrapolation.
- The carrier adjacent to a reserved tone is a pilot, except in cases where there are two adjacent reserved tones. This ensures that data cells remain in MISO pairs.

9.3.4.6 Use of the reference sequence in normal and extended carrier modes

In extended carrier mode, applicable for 8 K, 16 K and 32 K modes, the number of sub-carriers per OFDM symbol is increased, compared to normal mode.

In extended mode, an equal number of sub-carriers is added on each side of the used carriers available in normal mode. This number is referred to as K_{ext} . As a consequence, the first used carrier in extended mode does not align with the first used carrier in normal mode.

The PRBS sequence used to modulate pilot tones is applied in such a way that a given physical carrier in both normal and extended modes is modulated by the same PRBS sequence sample.

When in normal carrier mode, the carrier with index K_{min} is modulated with the element of the PRBS sequence whose index is K_{ext} in order to ensure that the same modulation is applied to the same physical carrier in both normal and extended carrier modes. Figure 71 illustrates this for the 8 K FFT mode; the same principle is applied to the 16 K and 32 K modes.




Figure 71: Modulation of the carriers in extended mode

9.3.5 Dummy carrier reservation

In addition to the pilots, the cells reserved for the Tone Reservation method of PAPR reduction (clause 9.3.8.3) are initially set to zero and not allocated to data. The tone reservation algorithm itself is applied later, in the time domain but in such a way that all the energy of the added signal will end up in the reserved carriers after the receiver's FFT. See clause 9.3.8.3 for more details on the locations of reserved tones.

9.3.6 Inverse Fast Fourier Transform (IFFT)

As for other COFDM systems, the Inverse Fast Fourier Transform (IFFT) algorithm may be used to implement the equations in clause 9.5 (and 9.8.2.4.1 for P1) of [i.1]. However, care is needed both in the exact IFFT implementation and in the mapping of the coefficients $c_{m,l,k}$ to the bins of the IFFT. The considerations are exactly the same as for DVB-T and are described in detail in annex D of [i.18].

In the extended carrier mode, the value of K_{\max} is increased, while K_{\min} remains at zero. However, because of the definition of k' , the centre carrier will remain at the centre frequency f_c , and an equal number of extra carriers will be added on each side. This, together with the extra P2 pilots added in the extended carrier regions, ensure that the part of the signal corresponding to the normal carrier region is identical in both normal and extended carrier modes.

9.3.7 P1 generation

The P1 signal can also be generated using an Inverse Fourier Transform. The P1 carrier of index 426 corresponds to the centre carrier ($k' = 0$) for the other symbols, and should use the IFFT bin corresponding to the same frequency.

Care should be taken with the normalisation of the IFFT coefficients and the gain of the IFFT, to ensure that the relative amplitudes of the P1 and other symbols are correct even though a different IFFT size might be used.

Implementers should also note the eccentric ordering of the FFT sizes in the S2 field of the P1 signalling.

9.3.8 PAPR reduction

9.3.8.1 Overview and comparison of the two techniques

DVB-T2 provides two complementary techniques for reduction of PAPR, viz., ACE and Tone Reservation. ACE provides significant reduction in PAPR for lower order constellations (especially QPSK) while tone reservation provides greater benefit for higher order constellations. It is possible to use both ACE and TR simultaneously; however, depending on the constellation, picking the better technique based on the constellation will provide most of the benefit. Both techniques come at a cost: ACE increases the noise level that the receiver sees while TR decreases throughput. The major trade-off is that of increased transmission power in the linear region of the power amplifier vs. the loss in throughput (whether directly, as in the case of TR, or indirectly, as in the case of decreased SNR from ACE).

From the perspective of the receiver, the presence of ACE would degrade the received SNR. The reserved tones would also be completely discarded by a typical receiver.

9.3.8.2 ACE

The Active Constellation Extension (ACE) algorithm, which is described in clause 5.13.1 of [i.1], modifies the power distribution of time-domain signal samples so as to enable an improvement in power efficiency of subsequent power amplifiers, or in out-of-band emission levels.

An additive signal is defined in such a way that, in the frequency domain, only data-cell values are affected - modified values being constrained to lie in extended constellations. Figure 72 shows the basic principle.




Figure 72: Principle of the ACE algorithm

Extended constellations are fully defined from original data-cell constellations and the knowledge of a maximal extension value L . L is a parameter of the ACE algorithm.

Various performance trade-offs can be obtained by adjusting L along with the two other ACE parameters: the clipping threshold V_{clip} and the gain G .

The ACE algorithm enables SFN operation. Indeed, multiple transmitters in an SFN configuration will transmit signals containing identical data-cell values provided all modulators are configured with identical ACE parameter sets (L, V_{clip}, G) .

The ACE algorithm as described in [i.1] is not suitable for use in conjunction with rotated constellations. Therefore, ACE processing should be disabled when rotated constellations are used.

A somewhat different form of ACE may also be applied to the L1-signalling to reduce the bias. In this case, the modulated constellations are modified directly according to the measured bias. See clause 8.14.1.4.

9.3.8.3 Tone Reservation

The basic idea of Tone Reservation is that some carriers are reserved to reduce PAPR. These reserved carriers don't carry any data information and are instead filled with a peak-reduction signal. Because data and reserved carriers are allocated in disjoint subsets of subcarriers, Tone Reservation needs no side information at the receiver other than an indication that the technique is in use, carried in the L1-pre signalling field "PAPR".

Figure 73 shows the structure of the OFDM transmitter using Tone Reservation. Reserved carriers are allocated according to predetermined carrier locations which are reserved carrier indices. After the IFFT, peak cancellation is operated to reduce PAPR by using a predetermined signal. The predetermined signal, or kernel, is generated by the reserved carriers.




Figure 73: The structure of the OFDM transmitter using Tone Reservation

The tone reservation algorithm may be applied to all of the symbols (not including P1), but even when it is not being used for the data symbols it is still applied to the P2 symbols to help reduce the L1 bias. See clause 8.14.1.1.

9.3.8.3.1 Reserved carriers and kernel

The kernel signal is defined as:

$$p_n = \frac{1}{N_{TR}} \sum_{k \in S_I} e^{j \frac{2\pi n(k-K_C)}{N_{FFT}}}$$

where N_{FFT} and N_{TR} indicate the FFT size and the number of reserved carriers, respectively. K_C is the index of the centre carrier. The kernel is the inverse Fourier transform of an $(N_{FFT}, 1)$ vector \mathbf{I}_{TR} having N_{TR} elements of ones at the positions corresponding to the reserved-carrier indices (S_I) and has $(N_{FFT} - N_{TR})$ elements of zeros at the others. The characteristic of the kernel is similar to an impulse. For example, the primary peak (p_0) of the kernel is one and the secondary peaks (p_1 to $p_{N_{FFT}-1}$) of the kernel have a value considerably smaller than p_0 . The reserved-carrier indices in annex H of [i.1] were selected for the given number of reserved carriers, N_{TR} based on obtaining a good kernel, i.e. having low secondary peaks.

When TR is activated by the L1-pre signalling field 'PAPR', the dummy carriers are reserved in the P2, data and frame-closing symbols.

NOTE: The reserved tones are always reserved in the P2 symbols, whatever the state of the 'PAPR' field, because the receiver will not know whether PAPR is used or not until it has decoded the L1 signalling which is carried in the P2 symbols themselves.

In normal carrier mode, the reserved carriers are as follows. In the P2 symbols, the carrier indices defined in table H.1 of [i.1] are reserved and the reserved-carrier indices do not change across the P2 symbols. In the data symbols excluding any frame-closing symbol, the set of carriers corresponding to carrier indices defined in table H.2 of [i.1] or their shifted equivalents are reserved depending on the OFDM symbol index l of the data symbol. In the frame-closing symbol, the set of carriers in the frame-closing symbol corresponding to the same carrier indices as for the P2 symbols, defined in table H.1, are reserved.

The reserved tones in the normal symbols are shifted from one symbol to the next in sympathy with the Scattered Pilot pattern so as to avoid collisions between Scattered Pilots and reserved tones. This was done in order to give the original designers more freedom in the choice of reserved tone locations than would have been the case if they had remained in fixed locations; this enabled a better kernel to be obtained.

For example, in the 8 K normal carrier mode with Scattered Pilot pattern PP1, assume that the number of P2 symbols (N_{P2}) and data symbols (L_{normal}) is 2 and 30, respectively. Table 33 describes the reserved-carrier indices in T2-frame excluding P1 symbol. The kernel signal will be changed according to the reserved-carrier indices.

Table 33: Example of the reserved-carrier indices in T2-frame using PP1

| Symbol index (l) | Type | Reserved-carrier indices |
|----------------------|-----------------------|--|
| 0 | P2 | {106, 109, 110, 112, ... , 6698, 6701} |
| 1 | P2 | {106, 109, 110, 112, ... , 6698, 6701} |
| 2 | Data (S_2) | {117, 121, 129, 221, ... , 6551, 6571} |
| 3 | Data (S_3) | {120, 124, 132, 224, ... , 6554, 6574} |
| 4 | Data ($S_4=S_0$) | {111, 115, 123, 215, ... , 6545, 6565} |
| 5 | Data ($S_5=S_1$) | {114, 118, 126, 218, ... , 6548, 6568} |
| 6 | Data ($S_6=S_2$) | {117, 121, 129, 221, ... , 6551, 6571} |
| 7 | Data ($S_7=S_3$) | {120, 124, 132, 224, ... , 6554, 6574} |
| ... | ... | ... |
| 30 | Data ($S_{30}=S_2$) | {117, 121, 129, 221, ... , 6551, 6571} |
| 31 | Data ($S_{31}=S_3$) | {120, 124, 132, 224, ... , 6554, 6574} |
| 32 | Frame-closing symbol | {106, 109, 110, 112, ... , 6698, 6701} |

In extended carrier mode, the positions of the reserved tones are altered as follows:

- The reserved carriers in P2 have the same physical carrier index k' as in normal mode, so the carrier indices k are increased by K_{ext} relative to the normal mode.
- The phase of the stepping pattern in the data symbols is adjusted to compensate for the change in the stepping pattern of Scattered Pilots.
- The reserved carriers in the frame-closing symbol (if any) are the same as in P2, including the addition of K_{ext} .

9.3.8.3.2 Peak-cancellation algorithm

Figure 74 shows the detailed block diagram of the peak-cancellation algorithm. The IFFT output (\mathbf{x}) is fed into the peak-cancellation block, and the peak position and value of \mathbf{x} are detected. Then the reference kernel, generated by the reserved carriers corresponding to the current OFDM symbol, is circularly shifted to the peak position, scaled and phase rotated. The resulting kernel is subtracted from \mathbf{x} and the new PAPR is calculated. The principle is shown in figure 75. If the PAPR of the resulting signal satisfies the target PAPR level, this signal is transmitted. If not, the cancellation operation is repeated iteratively, until the number of iterations reaches the predetermined maximum iteration number.

The power of each reserved tone is limited to 10 times the power of a data cell. If the cancellation signal to be added in an iteration would cause this limit to be exceeded, the cancellation signal is scaled down such the reserved tone of the largest amplitude ends up with a power equal to the limit. Note that the iterations can continue even after such limiting has occurred, since the largest remaining peak might require a correction of a different phase, which could even reduce the amplitude of the largest reserved tone.

The peak cancellation algorithm is described in more detail in clause 9.6.2 of [i.1].




Figure 74: Block diagram of the peak-cancellation algorithm




Figure 75: Principle of Tone Reservation algorithm

9.3.8.3.3 Selection of the clipping level

The desired clipping level (in magnitude), V_{clip} , should be selected taking into account the requirements of the hardware modules. If it is too low, the cancelled peak re-grows owing to the impact of the imperfect kernel. Otherwise, if it is too high, the benefit from PAPR reduction scheme becomes marginal.

9.3.8.3.4 SFN issues

Since the reserved tones are ignored by the receiver, the technique can be used in SFNs. It has been confirmed that the use of tone reservation does not exacerbate ICI problems, regardless of the SFN frequency offset.

10 Receivers

10.1 Overview

This clause mainly concerns consumer receivers, which are required to decode and output the Transport Stream or other stream carried by one PLP, together with its common PLP if appropriate.

Professional receivers may be designed which are capable of decoding more than one PLP, and possibly reproducing the original "big" transport stream (see clause 8.2.2), or alternatively re-generating a T2-MI stream corresponding to the received signal.

Some parts of a receiver correspond directly to a block or feature in the modulator, and implementers of receivers may find it useful to consult the corresponding parts of clause 9. In some cases an algorithm exists which is more efficient in terms of memory or processing on the receiver side, and some details will be given in the present clause. Other parts of the receiver, such as synchronisation, have no counterpart in the modulator.

The DVB-T2 receiver model outlined below, and the signal processing steps described with the text of this clause, are illustrative examples of how one might proceed. There are certainly other ways to realise the same functionality, hence the mechanisms outlined here are of no prescriptive character.

This clause suggests methods that may be used to implement some of the receiver functions, but should not be considered as a complete description of the implementation of a receiver.

10.1.1 Receiver Model

The top-level receiver model shown in figure 76 will be assumed for the purposes of the following clauses. More detail of the "DVB-T2 demodulator" block is shown in figure 77. Figure 78 in turn gives the details of the chains for Bit Interleaved Coding and Modulation (BICM) for the PLPs and L1 signalling. The processing chains for the data and common PLPs are the same.




Figure 76: DVB-T2 receiver model




Figure 77: DVB-T2 demodulator




Figure 78: Details of the BICM decoder chains for PLPs and L1 signalling

10.1.2 Acquisition scenarios

The tasks of acquisition and synchronisation can take slightly different forms in a range of situations, for example:

- when switched on for the very first time after purchase, or after a change of location;
- when switched on again at the same location;
- when a different service is selected which necessitates a change of multiplex from the one currently in use;

- when handing over between cells in mobile or portable reception;
- when switching to a different service for the purpose of announcements.

In the first case above, the receiver has no knowledge of the DVB-T2 services available and will have to perform a band scan to determine which (if any) RF channels contain DVB-T2 services, and then obtain more information about those services. This will be referred to as the "initial scan" scenario.

In the second and third cases, we presume that the receiver knows this information, and simply needs to acquire the signal in the known RF channel. This will be referred to as the "service selection and reproduction" scenario.

The fourth and fifth cases will be referred to as the "handover" scenario.

10.1.2.1 Initial scan

The initial scan is executed in order to find RF signals of DVB-T2 format at the current location and to determine their quality, i.e. if they are suitable for undisturbed reproduction or not. This requires determination of the modulation and coding parameters as well as synchronization to the RF signal in time and frequency.

Hence the following steps might be executed:

- 1) Select bandwidth.
- 2) Tune to first channel in RF range.
- 3) Detect RF signal.
- 4) Detect, validate and decode P1 (see clause 10.2.2). If unsuccessful, go to 1).
- 5) Identify Guard Interval (see clause 10.2.3).
- 6) Achieve fine synchronisation in frequency and time (see clause 10.2.4). If not possible, go to 1).
- 7) Extract and decode L1 signalling from P2 symbols (see clauses 10.2.6 and 10.5.7).
- 8) Select first PLP (according to L1 signalling).
- 9) Demodulator starts to extract the PLP and feed to the LDPC/BCH decoder.
- 10) BBF header decoder begins to decode and pass relevant parameters (including data-field length, ISSY, etc.) to stream regenerator.
- 11) Stream regenerator provides TS or GSE packets or bursts of Generic Continuous Stream data (incl. Layer-2 signalling) from error-corrected Baseband Frames to back-end, including reinsertion of deleted null packets and de-jitter buffering (see clause 10.6).
- 12) Back-end extracts relevant part of L2 signalling (relevant part of PSI/SI in TS case).
- 13) Modulation parameters and relevant part of L2 signalling stored for all services of current PLP.
- 14) Repeat steps 9-13 for the other PLPs (if any).
- 15) Repeat steps 2-14 for each RF channel.

10.1.2.2 Service-selection and reproduction

For service selection a programme guide listing all accessible services should be provided to the user. In order to set this up, Layer-2 signalling should also be processed.

As soon as a new service is selected, changes to the settings of that service need to be monitored by the receiver. These changes might affect the physical-layer settings or settings on higher layers.

The proposed simplified sequence of processing might look like this:

- 1) Other service selected? If so, go to 2), if not, go to 1).
- 2) Service selection is converted into RF bandwidth, channel selection and PLP ID on the basis of information stored during the initial scan (see above), information gathered during subsequent visits to each channel, and updates received via the PSI/SI (e.g. the NIT and SDT) if available.

NOTE: Receiver implementers should take care to avoid problems resulting from erroneous, incomplete, missing or infrequently transmitted PSI/SI information. In the case of a discrepancy between the PSI/SI and the information gathered from decoding a particular channel, the latter should be given precedence.

- 3) Tune to corresponding RF channel.
- 4) Apply steps 3) to 7) of initial scan procedure (see above).
- 5) Select desired PLP (according to Layer-1 signalling).
- 6) Apply steps 9) to 11) of initial scan procedure (see above).
- 7) Output of regenerated stream to back-end.
- 8) Reproduction of selected service.
- 9) Go to 1).

10.1.2.3 Hand-over (portable and mobile cases)

The word "handover" in DVB describes the change from one RF channel to another within the same delivery system for particular purposes as follows:

- a) The signal currently being received no longer gives undisturbed reproduction of the desired service, and the receiver should attempt to follow the service by switching to another signal, carrying the same service. If the same service is not available (i.e. receivable with sufficient quality for undisturbed reproduction) via another multiplex, the next level of customer satisfaction might be that the receiver changes to a service of the same or similar nature.
- b) The desired service is interrupted for a limited time period for the reproduction of one or more announcement(s).

Alternatively, the receiver may be able to follow the same service by changing to a different DVB delivery system. However, because the use cases discussed here are portable and mobile ones, DVB-S, -S2 and DVB-C are unlikely to be suitable alternative signal sources.

DVB-T, on the other hand, is in many cases an alternative source during the migration period from DVB-T to DVB-T2. Based on the fact that for the first years of regular T2 transmissions all receivers will be able to process both signals (but not at the same time), mobile and portable receivers will probably make use of this handover option, although it will often mean a change from the HD to the SD version of the content. For mobile and portable devices with smaller screens this might be an acceptable solution.

The scope regarding alternative signal sources can be extended beyond DVB. Other broadcasting as well as mobile and wireless communication systems, partly in conjunction with the Internet, will serve as sources of the same or related content.

Emphasis will be put on DVB-T2-internal mechanisms first.

10.1.2.3.1 Automatic service following

10.1.2.3.1.1 Overview

The functionality "automatic service following" is required for receivers that move, i.e. for portable and mobile devices. The purpose is to stay tuned to the same service although the multiplexes might change depending on the geographical region the terminal is located in.

A lower level of this functionality is to choose an appropriate substitute for a service that is no longer receivable at the receiver's current position. Appropriate substitute services might be those that are explicitly indicated as such (by the linkage descriptor) or those that are of the same genre.

In the absence of both the above options, the lowest level of user-friendliness would be to prevent a blank screen by reproducing one of the remaining available services.

Although feasible in the absence of any supporting information, this functionality can be supported by providers through the provision of the following information:

- Frequency information for neighbouring areas.
- Cell and sub-cell description for current and neighbouring areas.
- Information about the physical-parameter settings of neighbouring multiplexes carrying the same and/or related service.
- Information about the PLPs and Transport Streams carrying the desired service.
- Information about the receiver's current geographical position.

The following clauses group the possible levels of receiver-supporting information into three categories:

- none of the above-mentioned information available;
- frequency and cell information available; and
- all information defined in DVB for easing the handover process is available to the receiver.

10.1.2.3.1.2 Case "no supporting information"

In the absence of supporting information, a single-tuner device would benefit from a time-slicing structure of the signal on air (i.e. Input Mode B of DVB-T2). Between slices consisting of the desired data, the complete RF band(s) can be checked for alternative signals channel by channel, e.g. one channel per time interval between two relevant time-slices.

The same functionality can be realised even without time slicing in the presence of two tuners.

A single-tuner receiver would need to scan the RF band in the case of Input Mode A when the received signal becomes too weak for undisturbed reproduction of the desired service. This can be done in one go or sporadically in such a way that the reproduction of the original service would be interrupted for short time periods. In either case, the interruptions would be visible and audible.

10.1.2.3.1.3 Case "frequency, cell and service information provided"

With information about neighbouring multiplexes in terms of their centre frequency and the services carried, the receiver is able to separate multiplexes carrying the service to be followed from those not carrying it. For checking the signal quality of the desired multiplex(es) it is not necessary to initialise a complete band scan in the presence of frequency information. Of course, the receiver cannot determine the completeness of the information provided and would therefore need to start a band scan if the provided information did not lead to the reception of an undisturbed signal.

Service information can be provided with following PSI/SI table:

- SDT: This table itself maps services to Transport Streams. If applicable, the linkage descriptor can be employed for indicating services with different identifiers than those in the tuned multiplex that carry identical or related content.

Frequency information can be provided with the following PSI/SI tables:

- NIT: The T2 delivery-system descriptor, the frequency-list descriptor or the cell frequency link descriptor can be employed to provide information about alternative frequencies for the currently selected Transport Stream as well as about any other Transport Stream.

Cell Information can be provided with the following PSI/SI Table:

- NIT: The T2 delivery-system descriptor, the cell-list descriptor as well as the cell frequency link descriptor can be used to provide information about geographical cells to the receivers. The first and the last descriptors can combine the cell information with information about the RF channels (their centre frequencies, to be exact) that are carrying the TS described in the NIT.

10.1.2.3.1.4 Case "receiver's geographical position can be determined in addition"

On the basis of transmitter-identification information and/or in the presence of a GPS/Galileo receiver, the T2 terminal is able to further reduce the options described in the latter clause to those geographical cells that are in the vicinity of the receiver's current position.

10.1.2.3.2 Announcements

For the accurate provision of announcements that are provided via a different RF channel than the tuned one, a-priori information needs to be available to the widest possible extent since a time-consuming search exercise would mean missing the initial part of the announcement. Hence, beside the announcement trigger, frequency and cell information as well as information about the receiver's rough position should also be given.

10.2 Acquisition and synchronisation

10.2.1 Overview of synchronisation

The receiver needs to locate the received signal in both time and frequency before it can extract the transmitted information. Once initially achieved, this synchronisation should then be tracked continually throughout reception. The DVB-T2 signal contains many features that can be exploited for this purpose, leaving receiver manufacturers some freedom to devise their own implementations. Clauses 10.2.2 to 10.2.4 describe one combination of methods by way of example. Clause 10.2.5 explains an alternative method that can be used for some of the synchronisation tasks.

Different signal features are particularly adapted for these different tasks. These include:

- P1 symbol.
- P2 symbol(s).
- Scattered and Continual Pilots.

10.2.2 Use of P1

10.2.2.1 Introduction

The P1 symbol occurs once per frame, being inserted at the very beginning of each. It marks the beginning of each T2 frame and of each future-expansion-frame (FEF) part. The symbol has four main purposes:

- 1) To enable the receiver to determine quickly (e.g. during a band scan at a new location) whether a particular RF channel contains a DVB-T2 signal. (If it doesn't, then it can move on to try another channel, or check the same one for other services, e.g. DVB-T).
- 2) To identify the preamble itself as a T2 preamble. (Note that a signal may contain both T2 frames and FEF parts).
- 3) To signal some transmission parameters that are needed to decode the rest of the preamble and subsequently the main payload. In particular, the P1 discloses the FFT mode of the transmission. Although the guard interval of the transmission is still unknown and has to be ascertained, the reduction in the number of possibilities is useful to decrease the detection time.
- 4) To enable the receiver to detect and correct frequency and timing synchronisation.

The objective of the P1 symbol is, then, to save time during the scanning and channel set up processes.

10.2.2.2 Structure of P1.

The length of the P1 is fixed, regardless of the FFT mode and guard-interval configurations of the payload OFDM symbols. This makes it easier to detect - there is only one thing to look for.

The P1 signal is designed so that it can be detected despite the presence of a substantial frequency offset. This offset will occur for two reasons:

- tolerances in the receiver's frequency reference;
- deliberate offsets introduced at the transmitter as part of the planning process to control mutual interference between networks, especially those carrying different types of transmissions.

Detection is possible because the P1 signal contains sections which are frequency-shifted repetitions of the 'main' part of the P1, which is itself a 1 K OFDM symbol. These shifted repetitions can be distinguished despite the presence of a global frequency offset. Figure 79 illustrates the signal structure.




Figure 79: P1 structure

Note in particular that the first part, C, is a frequency-shifted copy of the first 542 samples of A, and the last part, B, is a frequency-shifted copy of the last 482 samples of A.

In principle, with a 1 K symbol, 853 carriers could be transmitted within the nominal bandwidth. However, only 384 are used, leaving others set to zero. The used carriers are centred in a block in the middle of the band, leaving 44 and 43 empty carriers at the sides. Within the middle block which is used, further carriers are set to zero in a known pattern which can be used for coarse frequency synchronisation. Assuming 8 MHz systems, the used carriers occupy roughly a 6,83 MHz block in the middle of the nominal 7,61 MHz signal bandwidth.

A modulation pattern is applied to the 384 used carriers, using DBPSK together with scrambling. The modulation patterns encode 7 bits of signalling.

10.2.2.3 P1 features.

The architecture of the P1 symbol has been designed to be robust, enabling the receiver to decode the P1 in challenging conditions, with the minimum overhead. P1 was therefore designed with the following features:

- Protection against interference:
 - It is expected that the P1 symbol can be correctly received and decoded under extremely adverse circumstances. The choice (for efficiency) of a fixed, short length for P1 means that inter-symbol interference may occur, but can be tolerated since the modulation and coding is designed to operate at very low signal-to-noise ratio. (The length of the left-portion C added to the main 1 K symbol A is not enough to absorb completely the response of the channel to the previous data symbol).
 - The presence of the two portions C and B at the beginning and end of the 1 K symbol improves robustness against both false detection and loss of detection which might otherwise occur in the presence of long-delayed echoes of the channel (even of opposite sign) or spurious signals (like CW interferers).

- Reception without any knowledge of the channel:
 - Owing to its carrier distribution, the P1 can be correctly recovered when the receiver is tuned to the nominal centre frequency. In fact, the symbol supports offsets of up to ± 500 kHz from the centre of the bandwidth in the case of 8 MHz systems.
 - The PAPR of the symbol has been optimised in order to make its reception better, even if any AGC loops are not yet stable.
- Offset-correction capabilities:
 - Within the initialisation tasks, the P1 symbol can be used to gain coarse time synchronisation of the receiver, as well as to detect (and then correct) any frequency deviation, for both fractional and whole-carrier shifts, from its nominal centre bandwidth.
- Robustness of signalling:
 - The signalling that is conveyed within the P1 is DBPSK modulated. It has been ensured that this protection is enough to recover the P1 signalling information even under negative values of SNR.
 - In addition, the signalling is encoded by using a complementary set of sequences which has the following two main theoretical properties:
 - a) The Sum of the AutoCorrelations (SoAC) of all the sequences of the set is equal to a Krönecker delta, multiplied by a factor KN , where K is the number of the sequences of each set and N is the length of each sequence. In the case of S1 $K=N=8$; in the case of S2, $K=N=16$.
 - b) The sets of sequences are mutually uncorrelated, and the sequences within any one set ("mates") are also mutually uncorrelated. The orthogonality of the sequences improves the robustness, in terms of decoding the correct pattern.

For further information on Complementary Sets of Sequences, see [i.2], [i.3] and [i.4].

10.2.2.4 Overview of the use of P1 during channel scanning

The main use of the P1, when receivers can make use of the most of its capabilities, is the "initial scan" scenario. Assuming the receiver starts from scratch, it may know nothing about which (channelised) frequencies convey a T2 signal. In this situation, the first task would be the discovery of the T2 frequencies.

As DVB-T2 signals are discovered, the receiver may go on to perform the set-up for each corresponding channel. Figure 80 shows the initialisation flowchart.




Figure 80: Initialisation Flow Chart

From figure 80, four general phases can be distinguished:

- T2 Initial Scanning:** This stage corresponds to the initial recognition of the T2 signal for a given frequency. The iterations corresponding to the bandwidth and nominal frequency selection were discussed in clause 10.1.2; as far as the P1 is concerned the key aim of this phase is to detect the presence of the P1. The receiver should perform two processes: the first one, in the time domain, aims at the detection of the presence of P1 structure and is referred to as *P1 detection* (clause 10.2.2.5). The second process is then used to validate the presence of a correct P1. This process, known as *P1 validation* (clause 10.2.2.6), is achieved by correlating the use of the carriers in the frequency domain.
- T2 Validation:** The next phase of the flowchart corresponds to the decoding of the signalling information carried by the P1 signal. The process of *P1 decoding* will be described in detail in clause 10.2.2.7. This phase is reached once the previous stage reports that a P1 was detected. Since the P1 can be defined not only for the T2 specification, but also for any other future broadcast system, it is possible that the detection corresponds to another system rather than T2. A validation step is needed in order to confirm the presence of the T2 in the selected nominal frequency, using the S1 and S2 fields.
- FEF loop:** The next stage aims to detect whether the channel is a pure T2 channel, a mixed T2 channel (i.e. it has FEF parts in it), or it does not transport any T2 signal at all. This process takes into account the S1 and S2 fields of the current and also the subsequent P1 preambles and is described in clause 10.2.2.9.

- **Set-up:** This step is not part of the initial scanning and T2 signal recognition, but part of the set-up process. Its main aim is the determination of the transmission parameters (i.e. L1 signalling).

The P1 symbol signals the FFT size as well as the SISO/MISO parameter. The rest of the parameters have to be decoded from the L1 signalling in the P2 symbols. However, the system has to extract the guard interval parameter before being able to correctly decode the L1 signalling.

The first task during the set-up is therefore the identification of the guard interval length, which is achieved by correlating the beginning and end parts of the normal OFDM symbols (see clause 10.2.3). After the GI length is obtained, the system may wait until the next T2 frame arrives.

The objectives of the initialisation tasks can be divided into two groups:

- Those objectives aimed at the detection of the T2 signal and its transmission parameters.
- Those objectives which are focused on synchronising and tuning the received signal.

Both groups of objectives can be achieved in parallel: much of the computation and processing is shared by both tasks.

The core of the initialisation tasks as far as the P1 symbol is concerned is based on the detection and/or decoding of three different correlations. Each of the correlations pursues different objectives. Two of them are located in the "Initial Scan" stage, while the other can be used to decode the signalling of the P1.

The flowchart for the "service selection and reproduction" scenario (clause 10.1.2.2) is similar, except that the iterations over different frequencies are not required, and the receiver may achieve faster acquisition by assuming initially that the parameters such as FFT size, guard interval and pilot pattern are the same as the last time the frequency was visited. Nevertheless, a receiver should behave correctly if these parameters do change, so should revert to the procedure in the acquisition flowchart if it does not succeed in acquiring the signal after a certain amount of time.

10.2.2.5 P1 detection

10.2.2.5.1 Summary

The P1 detection can be achieved by correlating the two parts added at either side of the 1 K OFDM (see figure 79). The correlation will be based on two branches running in parallel, each branch looking for the maximum similarity of the respective part of the repetition. By adapting the correlation windows and offsets to the architecture of the P1, the correlations will give their maximum peaks when the P1 is detected.

Although the P1 and the detection algorithm has been designed to cancel various spurious effects, there will sometimes be false detections. Thus, the detection has to be validated before assuming that the P1 is received.

10.2.2.5.2 Goals

The P1 detection aims to achieve the following:

- Detect the presence of a P1 symbol.
- Set a reference for coarse time synchronisation.
- Set an argument used to correct fractional ($\pm 0,5$ carrier spacing) frequency offsets.

10.2.2.5.3 Implementation

The implementation shown below performs the correlation for both parts of the P1 symbol, and has the merit of minimising the number of complex multipliers.




Figure 81: Correlation for the P1 Detection (Block diagram)

f_{SH} is the frequency shift applied to the sections C and B of the P1 symbol at the transmitter, and corresponds to the carrier spacing of the 1 K OFDM symbol that comprises section A of the P1. T_A , T_B and T_C correspond to 1 024, 482 and 542 samples respectively, the lengths of the sections A, B and C of the P1 symbol.

The delay elements T_C and T_B , together with an associated multiplier and running-average filter, each form part of a circuit (loosely called a 'correlator') which detects the frequency-shifted repetition in the signal in parts C and B respectively. The delay element T_A makes the outputs of these two 'correlators' line up in time.

T_R is chosen to be the reciprocal of f_{SH} and thus corresponds to 1024 sample periods, the same as T_A . This choice interacts with the specified f_{SH} in such a way as to eliminate unwanted complex-constant terms at the outputs of the two 'correlators' which might otherwise be caused by CW interference or certain other unwanted correlation conditions.

With this choice of T_R , the 'correlator' outputs, for a simple Gaussian channel, are complex pulses whose magnitude is a trapezoidal pulse of base width $(T_A + T_X)$, sloping sides of duration T_X , and top of width $(T_A - T_X)$, where T_X takes the value T_C or T_B respectively for the two correlators.

The argument of the 'correlator' outputs contains information about the fine frequency offset, but also the unknown arbitrary phase (with respect to the transmitter) of the down-shifter oscillator. By multiplying the two correlator pulses as shown, the effect of the unknown arbitrary phase is cancelled. The argument of the final output pulse can be shown to be proportional to the 'fine' component of the frequency offset.

The resulting waveforms are shown in figure 82. The left side shows the outputs of the two correlations, while the right side shows how they are time-aligned and multiplied.




Figure 82: Correlation for the P1 Detection

The presence of a pulse at the output may be detected by applying a suitable threshold to the magnitude and may be taken to indicate the presence of a DVB-T2 signal (or subsequent development using the same P1 structure).

10.2.2.6 P1 validation

10.2.2.6.1 Summary

The P1-validation step is needed to filter those false detections that match the previous step. The P1 symbol carries information in a subset of the available carriers of the symbol. This step should identify whether the incoming symbol matches with the distribution of the carriers. Even if the distribution is as expected, it can be found in different locations (i.e. shifted from the centre point). This step will also report that shift in order to enable the system to correct the frequency offset.

The main part A of the P1 symbol, which is itself a 1 K OFDM symbol, should be demodulated. Before doing this, a frequency-shift operation should be performed to correct for the 'fine' component of the frequency offset measured in the previous step. This will ensure that the remaining error comprises a whole number of 1 K OFDM carrier spacings (the 'coarse' error) plus a residual error which should be sufficiently small compared with the carrier spacing that OFDM demodulation can be performed to sufficient accuracy.

The timing of the pulse obtained in the previous step is used to select the 1 024-sample-long section of the received signal deemed (after suitable allowance for processing delay) to correspond to section A of the P1 symbol. This section is then demodulated in the usual way for OFDM by using an FFT. The result is a list of complex amplitudes for each of 1 024 'carrier' positions. Of course, not all of these positions correspond to a transmitted carrier, since only 384 carriers are transmitted in the P1 symbol.

The validation can be achieved by correlating the distribution across the carriers of the power of the received P1 with the expected Carrier-Distribution Sequence (CDS). In practice, further refinements may be needed to cope with the presence of CW interferers, e.g. by setting the value of any excessively-large square-magnitudes to zero before performing the correlation.

10.2.2.6.2 Goals

The P1 validation provides the following information:

- It confirms that the 1 K symbol received is a P1 symbol. This implies that the power in the frequency domain is distributed in accordance with the distribution sequence specified for the P1 symbol.
- It reports and corrects the integer number of carriers offset where the distribution of the power has been localised.

10.2.2.6.3 Implementation issues (coarse frequency correction)

Since the active carriers in which pilots are conveyed have been distributed keeping a safety area on both sides of the centred carrier, the detection of the peak of the correlation is possible when the frequency of the transmitted signal is inside ± 500 kHz range from the centre of the bandwidth (in the case of 8 MHz systems).

In this case, the peak of the correlation will report the integer number of carriers that the symbol is shifted. This measure corresponds to the coarse frequency offset (measured in 1 K-FFT carrier spacings).

After finding the location of the pilots, the system is ready to move forward to the decoding of the S1 and S2 fields.

10.2.2.7 Decoding the P1 signalling

10.2.2.7.1 Summary

Once the P1 structure is correctly detected and centred (i.e. the system knows which of the active carriers are used), the modulated and scrambled pilots are ready to be decoded.

Before reaching the decoding process the received sequence extracted from the carrier distribution has to be unscrambled and then demapped.

The signalling fields S1 and S2 are coded by using the patterns described in the specification. A simple comparison of the received sequences with those patterns expected will recover the signalling fields.

However, it is possible to improve their reception and detection by correlating the received sequence with the corresponding pattern.

10.2.2.7.2 Goals

Signalling decoding provides basic transmission parameters that are needed to receive the rest of the preamble. The main features of the decoding are:

- Firstly, the P1 identifies what type of preamble it is. Since it is possible to find FEF parts within the same channel carrying the T2 signal, the P1 will disclose whether it belongs to a FEF part or a T2 frame.
- When the P1 is at the start of a T2 frame, the signalling reports the following basic TX parameters: FFT size, MISO/SISO transmission.
- In addition, the P1 also signals whether the channel is a pure T2-channel or FEF parts are embedded.
- The robustness of the encoded and modulated signalling is such that it does not impose more stringent SNR requirements than are needed to receive and validate the P1.
- In addition to simple comparison of the received signal, the best performance is obtained by applying the properties of the complementary set of sequences theory from which the sequences were generated.

10.2.2.7.3 Implementation

In order to decode the signalling without any loss in the P1 robustness, the CSS-theory properties should be applied. The main two properties are described in 10.2.2.3, and also mentioned in [i.1]. In order to get the most from the sequence, the algorithm to decode each of the fields is the following:

- 1) Extract the received S1 or S2 sequence from the carrier distribution. We will use S2 for this example.
- 2) The S2 pattern is generated by concatenating 16 sub-sequences of 16 chips long.
- 3) For each pattern of the 16 possible patterns, do the following:
 - For each subgroup of 16 cells, perform the correlation of the subgroup with the corresponding sub-pattern, where each sub-pattern is coded with 1 and -1.
 - Add up all 16 correlation results.
 - The resulting figure, ideally, should be a Kronecker delta (with a scaling factor of $K \times N$).

- 4) At this point, 16 Sum of AutoCorrelation (SoAC) results are available corresponding to the "correlation" of the received symbol with all the possibilities.
- 5) Finally, a decision should be taken to decide which of the SoAC figures is the most probable. Basically, the decoded value will correspond to index of the pattern whose SoAC peak is the highest.

10.2.2.8 Interpretation of P1 signalling fields

In order to validate the presence of the T2 signal, the signalling bits conveyed within the P1 should be decoded. Taking both S1 and S2 fields into account, three possibilities may occur:

- a) P1 signals a T2-preamble ($S1=00X$). In this case, the algorithm confirms the presence of the T2 signal. The scanning for the current bandwidth and nominal frequency configurations has finished. The system would be ready to either scan a different frequency or enter into the set-up for that channel. In the latter case, the S1 and S2 fields also signal the FFT size and whether MISO is used, essential information for the set-up process.
- b) P1 is not a T2-preamble and there are no FEF parts mixed in the channel ($S1 \neq 00X$ & $S2 = XXX0$). In this case, the detection corresponds to a not-T2 preamble. In addition, the signalling is also instructing that the 'unknown preamble' is the only type that can be found in the channel. Thus, the channel is not conveying any T2 signal.
- c) P1 is not a T2-preamble, but there are FEF parts in the channel ($S1 \neq 00X$ & $S2 = XXX1$). In this case, the detection corresponds to a not-T2 preamble. However, the mixed bit is set: the channel conveys more than one preamble type. A second iteration (i.e. preamble detection) is needed to resolve the ambiguity.

10.2.2.9 The FEF loop

The FEF loop itself is entered only when the P1 that has been detected does not correspond to the T2 specification, but FEFs are mixed within the channel, i.e. case (c) above. In the other two cases, the FEF loop is not required.

In order to identify whether the T2 signal is transmitted, the algorithm may take into account the following three facts:

- Each FEF part in the T2 signal will be preceded by a P1 symbol.
- Only one FEF part may be found between two T2 frames.
- The FEF part has a maximum length, which is fixed in the T2 specification.

Taking these three facts together, we can deduce that if T2 frames are ever present in the signal, and a P1 which does not precede a T2 frame is detected, the next P1 will belong to a T2 frame.

The FEF loop is then in charge of looking for the next P1 and identifying if the S1 signalling matches the condition $S1 = 00X$. If the condition is true, the channel conveys a T2 signal. Otherwise, if $S1 \neq 00X$, it is not a T2 channel.

10.2.3 Determining the Guard Interval

Decoding the P1 signalling information tells the receiver what FFT size is in use and, in 8 K and 32 K a hint as to the guard interval, but the exact length of the guard interval (which can be chosen broadly independently of the FFT size) at this stage remains unknown.

There are at least two possible approaches to determining the guard interval.

10.2.3.1 Determining the Guard Interval using P2

In the 16 K and 32 K FFT modes, there is only one P2 symbol and it may therefore be possible to demodulate it and decode the L1-pre signalling (which signals the guard interval directly) even if the guard interval is not known. The L1-pre signalling can be decoded at an SNR of -4 dB in a Gaussian channel (see clause 14.6), which allows for a significant amount of inter-symbol interference resulting from imperfect sampling frequency and carrier frequency correction together with non-ideal positioning of the FFT window because of the unknown guard interval.

Nevertheless, demodulating P2 requires the coarse carrier frequency offset to be corrected to within half a carrier spacing or better. Whilst some indication of the offset will be obtained from the P1 processing (see clause 10.2.2), the residual error might still correspond to several carrier spacings at 16 K and 32 K. It may therefore be necessary to perform a search over a range of frequency offsets, which could be done during the first few normal symbols of the T2-frame. The frequency accuracy required is greater in 32 K than in 16 K.

10.2.3.2 Use of Guard Interval Correlation to determine the Guard Interval

In the FFT modes 8 K and below, there is more than one P2 symbol, making it more difficult to decode the L1-pre signalling without knowledge of the guard interval. In these modes, and as an alternative to decoding P2 in 16 K and 32 K, another way to determine the guard interval is to use so-called guard-interval correlation.

Knowing the length of the FFT defines T_U , but the total symbol length T_S also includes the guard interval, i.e. $T_S = T_U + T_G$. The guard interval can therefore be found if either the total symbol length T_S or, equivalently, the symbol rate $f_S = 1/T_S$ is measured.

Guard-interval correlation is commonly used for deriving time synchronisation in OFDM systems which (like DVB-T and DVB-T2) use a cyclic-prefix guard interval. The received signal is passed through a delay element of length T_U (remember, T_U is known by this stage), and the input and output are multiplied together after taking the complex conjugate of one of them. (This is the same process as used in the P1 detector described in clause 10.2.2.5, but without the need for the frequency shifter in this case).

To the extent that the data conveyed is essentially random, the complex output of this 'correlator' will appear to be noise-like. However, for a Gaussian channel (and in the absence of interference), it will nevertheless contain a degree of periodicity in its underlying statistics. The 'correlator' output has two distinct sections, which occur alternately:

- length T_U ; the underlying mean of the complex noise is zero;
- length T_G ; the underlying mean is non-zero (and has an argument which is related to the fine component of frequency error, although this property is not needed in order to determine the guard interval).

Note that the underlying mean cannot be exposed without some kind of averaging process, and the 'power' of the noise-like 'correlator' output is the same in the two sections. The underlying repetition rate $f_S = 1/T_S$ can nevertheless be extracted in a number of ways.

10.2.3.3 Use of the Guard-Interval hinting information from P1

In some modes, the S2 signalling field of the P1 preamble gives a hint as to the guard interval. This indicates whether the guard interval belongs to the "DVB-T" set (1/32, 1/16, 1/8 or 1/4) or is one of the "new" guard intervals introduced in DVB-T2 (1/128, 19/256 or 19/128). This is helpful when performing guard interval correlation in distinguishing between guard interval fractions which are close together, and which therefore result in similar total symbol periods T_S . Since the time required to distinguish two frequencies is roughly proportional to the reciprocal of the difference in frequency, eliminating close together guard intervals from the search set reduces the time taken to identify the guard interval.

The hinting information is provided in the following modes:

- 8 K, because there are two P2 symbols making decoding of P2 difficult without knowing the guard interval.
- 32 K, since it might be difficult to achieve enough frequency accuracy to decode the P2, and the long symbol duration might result in a long time required to distinguish the similar guard intervals if guard interval correlation is used.

The modes below 8 K do not have the "new" guard interval options, so the hinting information is not provided.

10.2.4 Time and Frequency synchronisation

The P1-symbol decoding process described in clause 10.2.2 can give an first estimate of the carrier frequency offset and the FFT window position. However, there is likely to be some error in both these estimates; in particular, the frequency-offset estimate is unlikely to be accurate enough for the 32 K-FFT mode, since the P1 is based on a 1 K FFT. Furthermore, the P1 signal was not designed specifically to allow estimation of the sampling-frequency error.

To perform fine synchronisation in frequency and time, it is expected that the receiver will use one of the many existing synchronisation algorithms developed for DVB-T. Such algorithms typically make use of the continual pilots for frequency-error measurements, together with an impulse-response estimate derived from the Scattered Pilots for fine adjustment of the FFT window position. Both continual and Scattered Pilots are available for use in DVB-T2. Synchronisation methods from DVB-T should be adapted to take account of the symbol-level PN sequence; note that the continual pilots are modulated by this sequence and will therefore change phase by either 0 degrees or 180 degrees from one symbol to the next.

In addition, the frame- and symbol-level sequences themselves may also be used for synchronisation as described in clause 10.2.5.

10.2.5 Use of Frame/Symbol level sequences for synchronisation

10.2.5.1 Introduction

Clause 10.2.2 describes how the P1 preamble symbol can be used, among other things, to estimate frequency and time synchronisation. In addition to this process (or as an alternative), and to get a more robust synchronisation (providing clock sampling, symbol, frequency and frame synchronisations), a PN sequence has been added on the top of the PRBS Scattered Pilots sequence. The synchronisation process, based on Scattered Pilots, is described in this clause.

10.2.5.2 Definition of the reference sequence

The pilots are modulated according to a reference sequence, $r_{l,k}$, where l and k are the symbol and carrier indices as defined in [i.1]. The reference sequence is derived from a symbol-level PRBS, w_k and a frame-level PN-sequence, pn_l . This reference sequence is applied to all the pilots (i.e. Scattered, Continual, Edge, P2 and Frame-Closing pilots) of each symbol of a T2-frame.

The output of the symbol-level sequence, w_k , is inverted or not inverted according to the frame-level sequence, pn_l , as shown in figure 83.

The symbol-level PRBS is mapped to the carriers such that the first output bit (w_0) from the PRBS coincides with the first active carrier ($k = K_{\min}$) in 1 K, 2 K and 4 K. In 8 K, 16 K and 32 K, bit w_0 coincides with the first active carrier ($k = K_{\min}$) in the extended carrier mode. In the normal carrier mode, carrier $k = K_{\min}$ is modulated by the output bit of the sequence whose index is K_{ext} (see table 60 in [i.1] for values of K_{ext}). This ensures that the same modulation is applied to the same physical carrier in both normal and extended carrier modes.

A new value is generated by the PRBS on every used carrier (whether or not it is a pilot).

Hence:

$$r_{l,k} = \begin{cases} w_{k+K_{\text{ext}}} \oplus pn_l & \text{normal carrier mode} \\ w_k \oplus pn_l & \text{extended carrier mode} \end{cases}$$




Figure 83: Formation of the reference sequence from the PN and PRBS sequences

The frame-level PN sequence is mapped to the OFDM symbols of the T2-frame such that the first chip (pn_0) from PN-sequence coincides with the first P2 symbol of the T2-Frame (no chip is allocated to the P1 symbol). The length of the frame-level PN sequence N_{PN} is therefore equal to the T2-frame length, i.e. the number of symbols in the T2-frame excluding P1. Table 34 shows the maximum length of PN sequence for different FFT modes in 8 MHz channels. The maximum number of symbols per frame will be different for channel bandwidths other than 8 MHz (see table 59 in [i.1]). The greatest possible value of N_{PN} is 2 624 (for 10 MHz bandwidth).

At each rising edge of the symbol clock, a new value of the PN-sequence is generated and applied to the current OFDM symbol (including P2 symbols but excluding the P1 symbol).

Table 34: Maximum lengths of PN-sequences for different FFT modes (8 MHz channel)

| FFT mode | Maximum sequence length, N_{PN} (chips) |
|----------|--|
| 1 K | 2 098 |
| 2 K | 1 081 |
| 4 K | 540 |
| 8 K | 276 |
| 16 K | 138 |
| 32 K | 69 |

The sequence ($pn_0, pn_1 \dots pn_{N_{PN}-1}$) of length $N_{PN} = L_F$, is formed by taking the first N_{PN} bits from an overall PN-sequence. The overall PN sequence is defined by table 35, and each four binary digits of the overall sequence are formed from the hexadecimal digits in table 35 taking the MSB first.

Table 35: PN-sequence Frame level (up to 2 624 chips) Hexadecimal description

| |
|--|
| 4DC2AF7BD8C3C9A1E76C9A090AF1C3114F07FCA2808E9462E9AD7B712D6F4AC8A59BB069CC50BF1149927E6B B1C9FC8C18BB949B30CD09DDD749E704F57B41DEC7E7B176E12C5657432B51B0B812DF0E14887E24D80C97F09 374AD76270E58FE1774B2781D8D3821E393F2EA0FFD4D24DE20C05D0BA1703D10E52D61E013D837AA62D007CC 2FD76D23A3E125BDE8A9A7C02A98B70251C556F6341EBDECB801AAD5D9FB8CBEA80BB619096527A8C475B3D8 DB28AF8543A00EC3480DFF1E2CDA9F985B523B879007AA5D0CE58D21B18631006617F6F769EB947F924EA5161E C2C0488B63ED7993BA8EF4E552FA32FC3F1BDB19923902BCBBE5DDABB824126E08459CA6CFA0267E5294A98C6 32569791E60EF659AEE9518CDF08D87833690C1B79183ED127E53360CD86514859A28B5494F51AA4882419A25A2 D01A5F47AA27301E79A5370CCB3E197F |
| NOTE: Table 35 defines the value of each chip (pn_l) of the PN-Sequence (for $0 \leq l < L_F$, where L_F is the T2-Frame length in use). |
| EXAMPLE: The beginning of the PN-sequence is: 010011011100..... ("4DC..."). |

10.2.5.3 Synchronisation algorithm

In the synchronisation method described below, all the synchronisations can be achieved, such as frame, symbol, frequency, and clock recovery, either from Scattered Pilots only or using all types of pilots. This synchronisation algorithm can provide an alternative or additional method to that based on the use of P1 symbol (see clause 10.2.2), which gives a means for frame synchronisation. In fact, the present method allows a very robust synchronisation, whatever the type of channel, and it can achieve synchronisation even at SNRs below -6 dB to -12 dB, depending on the FFT size and type of pilot pattern used.

NOTE: If the P1 symbol is not used, an alternative method needs to be used to determine the FFT size and use of MISO.

This method of synchronisation is entirely based on the analysis of Scattered Pilots. The basic algorithm consists in carrying out the cross-correlation between the received signal and the time representation of the Scattered Pilot sequence. The complex correlation results are successively exploited to perform, first, sampling-clock and symbol synchronisation, then frequency synchronisation and finally frame synchronisation with specific attention paid to the definition of the Scattered Pilots.

In the DVB-T2 frame structure, carriers are reserved for insertion of known pilots dedicated to channel-estimation purposes. These pilots multiplexed with the data in the frequency domain can thus be interpreted as a signal superimposed to the time-domain OFDM signal. The idea of making use of this known sequence prior to the reception FFT operation has been first published in [i.5] to provide an alternative channel-estimation technique for DVB-T systems. In [i.6], its exploitation has then been extended to the refinement of time synchronisation. Further information is given in [i.7]

10.2.5.4 Full synchronisation method

In order to minimize the complexity and to increase performance in the presence of a frequency offset, the cross-correlation is performed on N correlation fragments ($N \geq 16$), that are afterwards added to deliver a fully reliable symbol peak independent of frequency offsets (within a given frequency range) and of the other synchronisation processes. Phase variations can then be gradually calculated on the N correlation peaks to determine the frequency shift direction and used to tune the frequency synchronisation. The symbol phase law introduced by the PN sequence is finally used for frame synchronisation. The sequencing of this global synchronisation method is illustrated in figure 84.




Figure 84: Full synchronisation method sequencing.

In figure 85, we can distinguish the five main operations within the global synchronisation module: symbol correlation peak detection, sampling-clock synchronisation which is achieved by the phase comparison between peak detection and a pulse produced by a counter at the Scattered Pilot pattern frequency (e.g. if the PP spreads over 4 OFDM symbols, a pulse is produced every 4 symbols); then this counter gives the symbol synchronisation, and finally fully reliable frequency synchronisation and efficient frame synchronisation could be delivered to the baseband receiver.




Figure 85: synchronisation module block diagram

10.2.5.5 Evaluation of Synchronisation features

The main characteristics of the performance of this technique are as follows:

- Maximal acquisition time in the worst propagation conditions is very low (less than 0,2 seconds).
- Threshold for false detections is up to -12 dB of C/N in AWGN channel, 2 dB in SFN channel (Single Frequency Network, 0 dB echo) and 2 dB in TU6 channel (Typical Urban, 6 paths); the limitations of the other receiver digital processing are then already reached.
- Frequency offset up to 500 kHz for 8 MHz channel bandwidth (and even more, especially depending on channel bandwidth).
- There is no frame-synchronisation loss when the system switches from one path to another in multipath propagation conditions.

10.2.6 Use of P2

10.2.6.1 Initial Scan

After the P1 has been recovered, or frame synchronisation has been achieved by another method (see clause 10.2.5), the system may enter a set-up stage. Only basic parameters are known because either they were disclosed in the P1 symbol (i.e. FFT size, SISO/MISO) or they have been obtained from another process (i.e. guard-interval length). In any case, with these three parameters, the system is able to decode and extract the information in the P2 symbols.

It is assumed that RF and sampling-frequency offset has been corrected, and time synchronisation achieved (see clause 10.2.4). Any other control loops (such as AGC) should also be stable.

10.2.6.1.1 Reception of P2 symbols during initial scan

The P2 symbols come immediately after the reception of the P1. The number of P2 symbols is variable, but fixed for each FFT configuration. Thus, since the FFT size is already known, the number of P2 symbols that convey the L1 signalling is known.

L1 signalling is interleaved throughout all the P2 symbols. Thus, the receiver may need to store the P2 symbols before starting the decoding process.

There are many other transmission parameters that have not yet been recovered. However, this situation does not affect the reception of the P2 symbols, because they have been designed to be received without knowledge of all the L1 parameters. In particular, the receiver is not aware of the following configurations:

- **Scattered-Pilot patterns:** These patterns apply to the data symbols. The pattern of the pilots for the P2 symbols is known, when both the FFT and the MISO/SISO configurations are known, since there is a pilot every third carrier (every sixth carrier in 32 K SISO) irrespective of the Scattered Pilot pattern in the data symbols.
- **Extended or normal carrier mode:** For certain FFT sizes, the T2 signal can be sent using an extension in the number of used carriers. This information is conveyed in the L1-pre signalling and, therefore, is unknown when receiving the P2 symbols.

The P2 symbols, however, do not use the extended carriers to convey data. Instead, when the extended carrier mode is enabled, the extra carriers are padded with pilots. The central part of the signal, corresponding to the used carriers in the normal carrier mode, is mapped identically in both normal and extended carrier modes, both in terms of pilot locations and modulation, and in the location of L1 and PLP data.

During the initial scan, the receiver can therefore decode the P2 symbols assuming that normal carrier mode is used. If extended mode is used this can be deduced from the L1 signalling and the receiver can be adjusted accordingly.

Alternatively a receiver could assume that extended carrier mode is used; this has the advantage that it will be able to store the pilots in the extended carrier regions and therefore be able to equalise the subsequent data symbols immediately (see clause 10.3.2). On the other hand, this would require a wider channel filter bandwidth, which might result in worse interference from the adjacent channel if normal carrier mode is in fact in use.

Receivers may use the extra pilots in the extended carrier region of the P2 symbols to enhance their algorithms (e.g. channel estimation, frequency offset tracking, etc.).

- **PAPR - Tone reservation:** The carriers in the P2 symbols corresponding to the tones involved in the tone-reservation algorithm (to reduce the PAPR of the OFDM symbol) are always reserved and are never used for data, whether or not Tone Reservation is enabled for the data symbols. Thus, even if the PAPR method is enabled (the receiver does not yet know this information, which is disclosed in the L1-pre signalling), the reception of the P2 symbols is ensured.
- **Frame length, L1 signalling length, etc.:** Unlike other specifications where the L1 signalling is spread out over the frame, all the L1 signalling in T2 can be found in the P2 symbols. Thus, it is not needed to know in advance other configurations in addition to those already mentioned.

10.2.6.1.2 TX Basic parameters verification

The signalling carried in the P1 as well as the guard-interval length of the transmission are repeated in the L1-pre signalling conveyed in the P2 symbols. These parameters are normally needed to correctly receive the P2 symbols. However, there may be scenarios in which receivers can correctly recover the rest of the preamble even when some of these parameters were erroneously decoded, in particular the guard interval. The inclusion of the signalling for these parameters is intended to help receivers to verify transmission parameters.

10.2.6.1.3 Determining position in the super-frame

The P1 preamble sets a time reference for the beginning of every T2 frame. The P1 is also used to signal the beginning of the FEF parts embedded in a mixed-T2 channel.

In order to complete the set-up of the system, some other synchronisation tasks are still pending. In particular, the position of the T2-frame within the super-frame needs to be determined, using the current T2 frame index from the FRAME_IDX field of the dynamic L1-post signalling.

Together with the configurable post-signalling parameters FEF_INTERVAL and FEF_LENGTH, the frame index allows the position of any FEF parts to be determined and correctly dealt with.

Knowledge of the current T2 frame index is also needed to extract and decode PLPs which use multi-frame interleaving, since these have an Interleaving Frame that is multiple T2-frames in length. Similarly, PLPs which skip T2 frames (i.e. for which $I_{\text{jump}} > 1$) also have a repeating pattern of multiple T2 frames.

This frame index is not needed for the decoding of the L1-pre- and post-signalling.

10.2.6.2 Channel-zapping

This clause will briefly enumerate the steps a receiver may follow when changing channel ("zapping"). The scenario is described as follows: the receiver is decoding a given PLP on the tuned frequency, and is instructed to change to a different PLP which is transmitted on a different frequency.

10.2.6.2.1 Assumptions

It is assumed that the receiver has already visited all the frequencies, so that their initial scan and set-up process has been achieved (see clause 10.2.2). In this stage, and also assuming that L1-pre signalling will not change, the receiver is in the following state:

- Synchronisation: the receiver is perfectly synchronised to the tuned frequency. All control loops and tracking are stable. The other frequency has not been visited for a long time. Synchronisation will therefore have been lost, although the coarse frequency offset measured on the last visit may still be valid.
- Reference time: The frames on the two frequencies are not necessarily synchronised. Both the frames and super-frames on the two channels might start at different times.
- Service-to-transport-stream and transport-stream-to-PLP mapping has already been obtained and stored from the L2 signalling. Thus, the receiver knows to which frequency the tuner has to be changed.

10.2.6.2.2 Requirements for zapping between two channels

A typical flow might be as follows:

- 1) The tuner is re-tuned to the other frequency (already known). The receiver is already aware about L1-pre signalling, which implies that the receiver has all basic TX parameters backed-up.
- 2) After its stabilisation, the receiver has to synchronise to the new frequency. In general, synchronisation measurements and corrections (time, frequency, AGC, etc.) have to be detected and corrected again as for the initial scan case, i.e. using P1 or an alternative method (clauses 10.2.2 to 10.2.5). It may be possible to speed up the process by making use of values stored on a previous visit.
- 3) The next step will search for frame-reference synchronisation. The detection of the beginning of the frame can be achieved in parallel with step two, although it need not.
- 4) Then, the P2 has to be decoded in order to recover L1 signalling. The zapping time may be sped up by initially assuming that basic parameters needed to decode P2 have not changed.
- 5) The receiver should then identify the location of the target PLP:
 - If the PLP uses multi-frame interleaving, the receiver should wait until the start of the next Interleaving Frame for that PLP. This depends on the time interleaving configuration. The information can be deduced from the configurable L1-post signalling (FRAME_INTERVAL, FIRST_FRAME_IDX, TIME_IL_LENGTH and TIME_IL_TYPE).

- If the PLP is interleaved over one frame or less, the receiver may start processing the PLP in the next frame.

In both cases, the exact location of the cells for the PLP can be done according to the dynamic L1-post signalling information.

10.2.7 Receiver Synchronisation/Zapping time

Owing to the complexity of the system as well as the wide variety of use cases, it is not possible to characterise the time needed to do the initial scan or zapping with just one formula. However, acquisition times in typical scenarios will be presented here.

10.2.7.1 Initial scan time

As explained in clause 10.2.2.4, the receiver can deduce that no T2 signal is present for a given frequency and channel bandwidth if no P1 is received in 250 ms. It might be sensible for a receiver to wait longer than this, in case it missed a P1 owing to poor channel conditions. The minimum scanning time for channels which do not contain DVB-T2 is therefore 250 ms per RF channel and bandwidth.

If a P1 is received, the receiver needs to determine whether the channel contains a DVB-T2 signal, synchronise and start decoding the L1 signalling in the P2 symbol. If FEFs are not used, this will take a further frame, i.e. around 250 ms. In the worst case, in which maximum-length FEF parts are sent every T2 frame and the first P1 received corresponds to a P1 preceding a FEF part, this might take a further 3×250 ms under ideal channel conditions (assuming no loss of preambles due to any fading in the channel).

The receiver will then need to decode the L1 and L2 signalling; how long this takes will depend among other things on the repetition rates of the L2 signalling.

The tuning time and any bandwidth iterations also need to be taken into account.

Overall the scanning time will be roughly 500 ms per RF channel plus 1 s to 2 s for each channel that contains a T2 or T2-like signal.

10.2.7.2 Zapping time

The zapping time depends on many factors, some of them listed below:

- If the target PLP is in the same channel or a different one.
- The type of the target PLP and its interleaving configuration.
- The presence of FEF parts in within the T2 super-frame.
- The status of the information in the configurable L1-post signalling (whether it is obsolete or not).
- The moment within a T2 frame/super-frame in which the set-up of the target frequency ends (i.e. demodulator is synchronised, control loops stable, etc.).

The typical zapping sequence between PLPs on different frequencies, assuming there are no FEF parts, is:

- 1) Switch tuner frequency; Tuner and AGC settle (a few 10 s of ms).
- 2) Achieve fine RF-frequency, sampling-frequency and time synchronisation (in parallel with step 4).
- 3) Wait for the next P1 symbol (maximum one T2 frame; typically one half of a T2 frame).
- 4) Decode the L1 signalling in the P2 symbols following the P1.
- 5) Start to de-interleave the PLP data carried in the same frame.
- 6) Data will be output after the design delay. In input mode A the design delay will be one TI block, typically around 70 ms. In input mode B the design delay will typically be one frame for a high-data-rate PLP (see clause 8.8.3). Implementation latencies should be added to this figure.

The total worst-case zapping time for TS output from the demodulator, under these assumptions is therefore around 300 ms in input mode A and about 500 ms in input mode B. Typical values will be around 200 ms and 350 ms respectively. There will be further delays in the video decoder.

The zapping time will be greater for PLPs with longer time interleaving depths, or when FEFs are used.

The zapping time may be kept to a minimum by choice of the DVB-T2 parameters. The use of multiple TI blocks (particularly applicable in input mode A) will reduce the de-interleaving delay (at the expense of time diversity), and reducing the frame duration will reduce the acquisition time (at the expense of data rate)

10.3 OFDM Demodulation

10.3.1 FFT

The considerations for the use of the Fast Fourier Transform (FFT) algorithm in a receiver are essentially the same as those for the use of the IFFT in the modulator (see clause 9.3.6).

10.3.2 Channel equalisation

Channel equalisation comprises the processes of common phase-error correction, channel estimation and equalisation.

10.3.2.1 Common Phase-Error correction

As explained in reference [i.10], phase noise (whether introduced by local oscillators in the transmitter or the receiver) can be shown to introduce two distinct kinds of imperfection in OFDM:

- **Inter-Carrier Interference (ICI):** crosstalk between carriers of a more-or-less noise-like nature.
- **Common Phase Error (CPE):** a rotation, by a random but equal amount, of the received constellations on all the carriers of one symbol.

The degree of these imperfections depends on the phase-noise spectrum. Very loosely speaking, the low-frequency components of phase noise cause CPE whilst the remainder of the phase-noise spectrum causes ICI.

Since the random rotations of CPE are (by definition) the same for every carrier within the same OFDM symbol they can be measured (and thus corrected if desired) by using the reference information provided by the Continual Pilots (CPs).

NOTE 1: If the amount of CPE caused by the oscillators in a particular receiver is small enough, there is no need to perform CPE correction!

NOTE 2: There is some scope to exploit other cells as well as CPs, but CPs were provided with this function in mind.

CPs are specified carriers which transmit reference information in every symbol. The reference information is a function of both carrier number and symbol number through the frame (see clause 9.2.2 of [i.1]), being a real constant multiplied by two sequences taking the values ± 1 (one varies across the carriers within each symbol, the other varies from symbol to symbol through the frame).

The random change in CPE from one symbol to the next can be measured as follows. The received CP values are first corrected by being multiplied by the frame-based ± 1 sequence and then differentially demodulated on each CP carrier by multiplying the value for the current symbol by the complex conjugate of the value for the same carrier in the previous symbol. The values are summed for all the CPs in one symbol. The argument of the resulting complex number is the CPE of the current symbol with respect to the previous one. This can then be used to correct the current symbol. In effect the very first symbol received is treated as a reference of zero CPE.

The summation of the results for a large number of CPs reduces the effect of normal additive noise (e.g. receiver noise) by averaging. This is important since this additive noise will itself contribute a phase-noise component to the measurement - it is important for this to be substantially smaller than the CPE which is being corrected, otherwise the correction process will do more harm than good! It may also be necessary to exclude from the calculation any CPs on carriers that have been found to be suffering from CW interference.

The CPE-corrected values are then used in all subsequent processing.

Special care needs to be taken with the frame-closing symbol (see clause 9.2.7 of [i.1]). When present, this contains no continual pilots as such. However, when used, the frame-closing symbol contains sufficient Scattered Pilots that frequency-only interpolation can be used to produce a channel estimate for every cell in it, including those that would otherwise have been CPs. These can be used in the algorithm described above to facilitate CPE correction of the frame-closing symbol (and should themselves be corrected for CPE before being used for equalisation). Similar reasoning can also be applied at the start of the frame, where P2 symbol(s) also do not contain CPs but do contain so-called P2 pilots at a very dense spacing.

10.3.2.2 Locating the CPs

The number and locations of the Continual Pilots vary depending on FFT size, Scattered Pilot Pattern and the use of Extended Carrier mode. Receiver implementers should refer to clauses 9.3.4.2 and 9.3.4.3 which list a number of important points to be noted as well as tabulating numbers of pilots of various types as a cross-check.

The continual-pilot locations have been chosen in an attempt to reduce the size of look-up table required in the receiver. For each Scattered Pilot pattern there is a set of carrier indices $k_{i,32\text{ K}}$ which can be taken directly as the continual-pilot locations in 32 K. In smaller FFT sizes, a subset of the larger set is taken, and the $k_{i,32\text{ K}}$ are interpreted modulo K_{mod} , where K_{mod} is given in table 55 of [i.1]. The set of $k_{i,32\text{ K}}$ values for each FFT size is a subset of the values used in the next larger FFT size, so the values for 8 K are a subset of the 16 K values and so on.

One possible implementation is as follows, and the scheme is shown in figure 86.

The full set of $k_{i,32\text{ K}}$ values is partitioned into seventeen subsets according to the value of $k_{i,32\text{ K}} \text{ div } 1\ 632$, and each subset stored in numerical order in its own lookup table, together with an indication of the group to which it belongs. As the incoming carriers are processed, the current carrier index k is compared to the next value in the first lookup table (values of $k_{i,32\text{ K}}$ less than 1 632). At the same time, modified values of k are compared with the next value in each of the other lookup tables. If any value matches, and the group number means that the value is included in the current FFT size, the current carrier is a CP. Whether or not the value is used in the current FFT size, the pointer is advanced to the next value (if any) in the lookup table.

The modification consists of adding a multiple of 1 632; in practice this may be achieved by having a separate counter which is initialised to the relevant value instead of zero. In the 1 K and 2 K cases, $k_{\text{mod}} = 1\ 632$, so the initialisation values for each lookup table are 0, 1 632, 3 264, 4 896, etc. In 4 K, $k_{\text{mod}} = 3\ 264$ and the offsets are 0, 0, 3 264, 3 264, 4 896, etc. The structure of figure 86 can be used for all modes, with the preset values (in the ellipses) changed according to the FFT size and the lookup table changed according to the Scattered Pilot pattern. The full set of preset values is given in table 36.




Figure 86: Generation of CP locations (preset values shown for 1 K mode)

Table 36: Preset values for CP generation

| Table index | 1 K and 2 K | 4 K | 8 K | 16 K | 32 K |
|-------------|-------------|--------|--------|--------|------|
| 0 | 0 | 0 | 0 | 0 | 0 |
| 1 | 1 632 | | | | |
| 2 | 3 264 | 3 264 | 6 528 | 0 | |
| 3 | 4 896 | | | | |
| 4 | 6 528 | 6 528 | 6 528 | 0 | |
| 5 | 8 160 | | | | |
| 6 | 9 792 | 9 792 | 13 056 | 13 056 | |
| 7 | 11 424 | | | | |
| 8 | 13 056 | 13 056 | 13 056 | 13 056 | |
| 9 | 14 688 | | | | |
| 10 | 16 320 | 16 320 | 19 584 | 13 056 | |
| 11 | 17 952 | | | | |
| 12 | 19 584 | 19 584 | 19 584 | 13 056 | |
| 13 | 21 216 | | | | |
| 14 | 22 848 | 22 848 | 19 584 | 13 056 | |
| 15 | 24 480 | | | | |

10.3.2.3 SISO channel estimation

10.3.2.3.1 Overview

10.3.2.3.1.1 The need for channel estimation

The received carrier amplitudes output by the receiver FFT are not in general the same as transmitted - they are affected by the channel through which the signal has passed on its way from the transmitter.

Consider the *channel extent*, which could be variously described as: the duration of the impulse response from the first significant component to the last; or the shortest duration which can be chosen without excluding any significant impulse-response components; or, more practically, the shortest duration which can be chosen so that at least X % of the total signal energy is included, where X % is some substantial proportion, e.g. 99,9 %.

Provided this channel extent of the channel's impulse response does not exceed the guard interval T_G , and that correct OFDM synchronisation is maintained, the received (complex) carrier amplitudes can be given by:

$$Y_{k,l} = H_{k,l} X_{k,l} + N_{k,l}$$

where:

$X_{k,l}$ represents the complex modulation-'symbol' (constellation) applied on carrier k , symbol l ;

$Y_{k,l}$ represents the corresponding received carrier-amplitude;

$H_{k,l}$ represents the (complex) frequency response of the channel during symbol l , sampled at the carrier frequency, i.e. $H_{k,l} = H_l(f_k)$; and

$N_{k,l}$ represents the additive receiver noise.

NOTE 1: A further explanation of what happens is as follows. The received signal is the transmitted signal convolved with the channel impulse-response. The addition of a COFDM guard-interval (also known as a cyclic prefix) has the effect of converting this linear convolution into a cyclic one (provided the channel extent does not exceed the guard interval). Cyclic convolution in time corresponds to multiplication in frequency, when the two are related by a DFT operation.

NOTE 2: The relationship can also be written as a matrix equation; the effect of the channel is to multiply by a simple diagonal matrix - unless orthogonality is lost (e.g. because the channel extent is too great), whereupon more entries in the matrix become non-zero.

The receiver needs knowledge of the $H_{k,l}$ if it is to interpret the $Y_{k,l}$ in the best way. One simple way for re-scaling the received constellations is equalisation: dividing by our *best estimate* $H'_{k,l}$ of the complex channel response $H_{k,l}$ pertaining to each data cell. It will be clear that provided the estimate $H'_{k,l}$ is noise-free, this operation does not change the signal-to-noise ratio. The SNR will however already be degraded by the channel response as we have noted: weakly-received carriers have a poorer SNR than others. Detailed discussion of noise in the estimate itself is postponed to clause 14.4.

Dividing by the estimated channel response is somewhat similar to a zero-forcing equaliser, and would normally be deprecated on the grounds of aggravating the effects of noise. However, under the assumption that coded OFDM is being used, we simply weight the soft-decision values fed to the error corrector appropriately to take account of the different SNRs with which the data on the various carriers are received. (These weighted soft decisions are also known as metrics).

An alternative method simply inputs both $Y_{k,l}$ and $H'_{k,l}$ into the metric calculation without re-scaling the $Y_{k,l}$ to the standard size first - the result is equivalent but perhaps less intuitive.

10.3.2.3.1.2 Obtaining the estimates

The channel estimate can be derived from the known information inserted in certain OFDM *cells* - a term we use for the entity conveyed by a particular combination of carrier (location in frequency) and symbol (location in time). These cells containing known information are known as *pilot cells*; they are affected by the channel in exactly the same way as the data and thus - barring the effect of noise - precisely measure the $H_{k,l}$ for the cell they occupy. We calculate the measured channel estimate for this cell as:

$$H'_{k,l} = \frac{Y_{k,l}}{X_{k,l}} = H_{k,l} + \frac{N_{k,l}}{X_{k,l}}$$

In principle this can be done for any cell where we know what information $X_{k,l}$ has been transmitted. In most cases this will be the Scattered Pilot cells (SPs). However, continual pilots (CPs) are also available for a smaller proportion of cells, while another method (so-called CD3-OFDM [i.15]) is even able to use every data cell for channel measurement, by a data-directed feedback process: once the data has been decoded, the transmitted constellations $X_{k,l}$ can be reconstructed and used to measure $H_{k,l}$ retrospectively.

To obtain the estimates of the channel response for every data cell, the normal approach is to *interpolate* between the values $H'_{k,l}$ (which are only available for those $\{k, l\}$ corresponding to transmitted pilots) to provide values for every cell.

10.3.2.3.2 Fundamental limits

The channel response $H_{k,l}$ in general varies with both *time* (symbol index l) and *frequency* (carrier index k). The temporal variation corresponds to external causes such as Doppler shift and spread, but also to instantaneous error in the receiver's frequency tracking. The variation with frequency is a symptom of channel selectivity, itself caused by the channel comprising paths having different delays.

In effect, the receiver samples the channel response $H_{k,l}$ by measuring it for the cells $\{k, l\}$ within which pilot information has been transmitted. For the most part, the SPs are used for this purpose, but where other types of pilot information are present, they can also be used if desired. The SPs constitute a form of 2-D sampling grid, and in consequence there are limits, according to Nyquist's criterion, on the rates of variation of the channel response with time and frequency that can be measured using SPs.

Clause 9.2.3.1 of [i.1] defines the various SP patterns of DVB-T2, and introduces terms D_X and D_Y to characterise them. D_X is the separation between pilot-bearing carriers, so if $D_X = 3$, say, then every third carrier contains SPs - but not in general within a single symbol. This is because there is a diagonal pattern, which repeats every D_Y symbols. So, on carriers that are pilot-bearing, an SP occurs, and a measurement can be made, in every D_Y th symbol. Symbols occur at the rate $f_s = 1/T_s = 1/(T_U + T_G)$. It follows that the Nyquist limit for temporal channel variation that can be measured is $\frac{\pm 1}{2D_Y(T_U + T_G)}$ Hz.

Given suitable temporal interpolation, then we will have either a measurement or an interpolated estimate of the channel response for every cell on the pilot-bearing carriers. Estimates for the remaining cells can then be found by frequency interpolation between the pilot-bearing carriers. Since these are spaced by D_X carriers, or $D_X f_U = D_X/T_U$ Hz, it follows that the maximum Nyquist channel extent, or spread between the first and last paths in a channel that can be supported, is T_U/D_X sec.

Note that this approach is a variables-separable one leading to a rectangular Nyquist area on a diagram of Doppler versus delay. This rectangular area corresponds to a (dimensionless) timewidth-bandwidth product having the value:

$$\frac{1}{D_Y(T_U + T_G)} \cdot \frac{T_U}{D_X} = \frac{1}{D_X D_Y (1 + GIF)}$$

where $GIF = T_G/T_U$ is the guard-interval fraction.

It can be shown that the same sampling grid of channel measurements can be interpreted to produce other shapes of supportable 'area', in general non-rectangular, but whose total area remains the same.

Note that the spacing between pilots in just one particular symbol is in general greater than D_X carrier spacings, being $D_X D_Y$, the inverse of the Scattered Pilot density. It follows that if no temporal interpolation is performed, and channel estimates are obtained solely by frequency interpolation within one symbol, then the applicable Nyquist limit for channel extent is tighter than described above, being $\frac{T_U}{D_X D_Y}$. Some of the pilot patterns have nevertheless been chosen so that frequency-only interpolation is both possible and sensible in certain scenarios.

10.3.2.3.3 Interpolation

10.3.2.3.3.1 Limitations

In principle, channel variations within the Nyquist supported area described in the previous section can be measured. However, the Nyquist limit can only be very closely approached when using an interpolator having a very large number of taps. This is unattractive on grounds of cost and complexity, but is also bounded by other practical constraints.

The frequency interpolator can only make use of the finite number of pilot-bearing carriers.

Similarly, the temporal interpolator clearly cannot access measurements from before the time the receiver was switched on or the current radio-frequency channel was selected. Much more importantly, the length of the temporal interpolator is tightly limited by the fact that the main signal stream awaiting equalisation has to be delayed while the measurements to be input into the temporal interpolator are gathered. This delay has a large cost in terms of memory needed, but also in terms of the delay it introduces before the programme material can be delivered to the viewer. The temporal interpolator is thus usually much more constrained in its size than the frequency interpolator.

It follows that the full Nyquist area cannot in practice be supported.

An *ideal* interpolator, whose bandwidth (or as appropriate, 'time-width' as defined later) is chosen to be some fraction $x < 1$ of the Nyquist limit can also reduce the noise on the channel estimate by the same factor x . Note however that practical interpolators are unlikely to realise this good a result. Perversely, simple linear interpolation does achieve noticeable noise reduction, but in this case at the expense of poor interpolation accuracy through most of the wanted bandwidth.

10.3.2.3.3.2 Temporal interpolation

The presumption is that most receivers, on grounds of simplicity and minimising delay, will use at most simple linear temporal interpolation. Figure 87 shows an example pilot pattern and three different cases of temporal interpolation. To begin with we will consider the simplest case of interpolation between Scattered Pilots (blue in figure 87).




Figure 87: Various cases of temporal interpolation

In DVB-T a single pilot pattern was used, having $D_x = 3$ and $D_y = 4$. It follows that simply to perform simple linear interpolation requires three ($= D_y - 1$) symbols' worth of storage for the main data stream. In DVB-T2, the pilot patterns for each mode were originally chosen so as to limit the storage requirement for this purpose to twice that of DVB-T, rather than four times as might be expected given the presence of the 32 K mode. In 32 K, the receiver was expected to perform temporal interpolation over at most two symbols, whilst in 16 K mode interpolation over four symbols was expected. However, improved performance for echoes outside the guard interval can be achieved by performing temporal interpolation in modes for which it was not intended. For example, PP7, which has $D_y = 4$, is allowed in 32 K mode, but only in the 1/128 Guard Interval where it was intended that frequency-only interpolation will be used. Performing temporal interpolation increases the range of echo delays that can be equalised to $1/32$ of T_U . This may be necessary to meet certain performance targets.

Simple linear temporal interpolation leads to the Doppler bandwidth within which the interpolation is reasonably accurate being quite restricted, and significantly less than the Nyquist bandwidth. However, quite apart from the bandwidth of the interpolator, there are other constraints on achievable Doppler performance, especially the effect of inter-carrier interference caused by Doppler. This means that linear interpolation is not quite as bad a handicap as it at first appears.

Temporal-interpolation accuracy can be increased somewhat without increasing main-signal-path memory by using 'one-sided' interpolator designs having more taps than a simple linear interpolator.

10.3.2.3.3.3 Frequency interpolation

Frequency interpolation is illustrated in figure 88.




Figure 88: Frequency interpolation

Fortunately the frequency interpolator can use rather more taps. This is in fact necessary, since the 'time-width' (an analogous term, for frequency-domain sampling, to the common use of bandwidth in relation to time-domain sampling) of the interpolator now has to be an appreciable fraction of the Nyquist limit. Depending on the pilot pattern and guard-interval fraction that are chosen, the necessary time-width (if the full guard-interval range of delay is to be supported) can be 75 % or 89 % Nyquist; the 75 % applies for the simple guard-interval fractions whilst the 89 % is for the 19/128 and 19/256 cases. Furthermore, there is a case for making the time-width (if possible) a little greater than simply the guard-interval duration, in order to cope better with a degree of timing error, or with channels whose extent isn't strictly contained within the guard-interval duration.

Suppose N taps are used. For most carriers, the interpolator would make use of estimates/measurements from $N/2$ pilot-bearing carriers on one side plus $N/2$ pilot-bearing carriers on the other side. However, as the carrier to be estimated approaches the upper or lower limit of the OFDM spectrum, this is no longer possible, as some of the interpolator taps 'fall off the edge'. In this case it is still possible to use an N -tap interpolator, but with taps chosen so that it becomes progressively more one-sided. The performance is slightly compromised, but remains better than the alternative of using fewer and fewer interpolator taps as the edge is approached.

A further complication arises because of the diagonal pattern, in the case where frequency-only interpolation is performed. The edge carriers are chosen to be pilot bearing, so in one symbol out of D_Y , the edge carrier is a conventional SP, and all the pilot carriers in that symbol have uniform spacing $D_X D_Y$. In the other symbols, the edge carrier would not be an SP. However, the edge carriers have been designated as *Edge Pilots*, so that pilot information is nevertheless always available there. In this case, the spacing between the Edge Pilot and the next carrier having an SP is less than the normal spacing of $D_X D_Y$. As the spacing is less than normal there is no violation of the Nyquist limit, but the interpolator used near the edges will be 'irregular', having one spacing between its taps which is less than the others.

10.3.2.3.3.4 The choice whether to use frequency-only interpolation

The Nyquist limits to the channel extents that can be supported have already been explained. Table 37 is table 51 from [i.1], but with the Nyquist limits for supported channel extent added, for frequency-only and time-and-frequency interpolation cases.

Table 37: Nyquist limits for t-only and t-and-f-interpolation

| Pilot Pattern | D_X (Separation of pilot-bearing carriers) | D_Y (No. of symbols forming one Scattered Pilot sequence) | Nyquist limit as fraction of T_U , for f -only interpolation | Nyquist limit as fraction of T_U , for t -and f -interpolation |
|---------------|---|--|--|--|
| PP1 | 3 | 4 | 1/12 | 1/3 |
| PP2 | 6 | 2 | 1/12 | 1/6 |
| PP3 | 6 | 4 | 1/24 | 1/6 |
| PP4 | 12 | 2 | 1/24 | 1/12 |
| PP5 | 12 | 4 | 1/48 | 1/12 |
| PP6 | 24 | 2 | 1/48 | 1/24 |
| PP7 | 24 | 4 | 1/96 | 1/24 |
| PP8 | 6 | 16 | 1/96 | 1/6 |

By comparing the Nyquist limits with the guard intervals available in conjunction with the various combinations of FFT size, pilot pattern and guard-interval fraction, we can identify the cases where frequency-only interpolation is an option. Table 38 is a copy of table 52 from [i.1], with green highlighting added to the cases where frequency-only interpolation is an option, i.e. its Nyquist limit is sufficient to support at least the intended guard interval.

Table 38: Combinations for which frequency-only interpolation is possible (green) and not possible (black)

| FFT size | Guard-interval fraction | | | | | | |
|----------|-------------------------|-------------------|--------------------------|--------------------------|-------------------|-------------------|------------|
| | 1/128 | 1/32 | 1/16 | 19/256 | 1/8 | 19/128 | 1/4 |
| 32 K | PP7 | PP4 PP6 | PP2 PP8 PP4 | PP2 PP8 PP4 | PP2 PP8 | PP2 PP8 | NA |
| 16 K | PP7 | PP7 PP4 PP6 | PP2 PP8 PP4 PP5 | PP2 PP8 PP4 PP5 | PP2 PP3 PP8 | PP2 PP3 PP8 | PP1 PP8 |
| 8 K | PP7 | PP7 PP4 | PP8 PP4 PP5 | PP8 PP4 PP5 | PP2 PP3 PP8 | PP2 PP3 PP8 | PP1 PP8 |
| 4 K, 2 K | NA | PP7 PP4 | PP4 PP5 | NA | PP2 PP3 | NA | PP1 |
| 1 K | NA | NA | PP4 PP5 | NA | PP2 PP3 | NA | PP1 |

In all the cases where the PPx designator is rendered in green in table 38, it is *possible* to use frequency-only interpolation since channel extents up to the guard interval remain within the Nyquist limit. However, in some cases the margin is small. If temporal interpolation were used as well in these cases, then the margin between the desired time-width and the Nyquist limit would be increased. This would ease the design of the frequency-domain interpolator, and also permit the supported channel extent to exceed the guard interval, which can be advantageous in some cases (noting that there is often no hard physical limit to the channel extents that may occur). Making the design time-width appreciably less than the Nyquist limit also enables the interpolator to achieve some noise reduction.

On the other hand, note also that using frequency-only interpolation, as well as avoiding the complexity and delay of a temporal interpolator also avoids its poor frequency response - the time-domain Nyquist limit now becomes simply $f_s = 1/T_s = 1/(T_U + T_G)$, which is very substantially above the limiting factors of OFDM itself.

10.3.2.3.4 P2 symbols

The P2 symbols are different from the main payload data symbols since they always use a relatively dense pilot pattern. There are two versions: one in which every sixth carrier is a pilot, which is used with 32 K FFT in SISO mode; the other one has every third carrier as a pilot and is used for all other FFT sizes (and for 32 K MISO mode; see clause 10.3.2.4). The pattern thus depends only on whether the FFT size is 32 K, and in particular does not depend on the selected guard interval or (data-symbol) pilot pattern. The pattern is chosen to be sufficient to permit decoding with frequency-only interpolation, since in 32 K and 16 K modes P2 comprises a single symbol only.

The reason for having a P2 pilot pattern that only depends on the FFT size is simple. When performing a scan, the receiver starts with no knowledge of the options in use. By finding a P1 symbol, it can deduce the presence of a T2 signal and also determine the FFT size. Once the FFT size is known, the receiver can go on to deduce the guard interval.

For further information (including the pilot pattern used for the main payload), the receiver then needs to decode the L1 signalling in the P2 symbol(s). At this stage it knows the FFT size, can deduce which of the two P2 pilot patterns is in use in the P2, and is thus able to decode the P2 and with it the L1 data.

This explains why dense patterns are used - they have to be adequate for the largest guard interval that is available. In general, the largest guard-interval fraction is 1/4, which is adequately covered by having 1 pilot in 3. For 32 K, the largest guard-interval fraction is 19/128, for which 1 carrier in 6 is just enough.

Using dense patterns of pilots in P2 has its consequences. Obviously there is a loss of capacity, but this is considered acceptable since there are very few P2 symbols. It also means that the pilot pattern may not be boosted as heavily as the payload's Scattered Pilots are. This implies that the channel estimation for P2 symbols is noisier than for the main payload.

Some strategies can be used to improve the channel-estimate noise in P2. Once the L1 information has been decoded, so that the payload pilot pattern and guard interval are both known, if the corresponding expected channel extent is shorter than the maximum then a frequency-only interpolator having a narrower time-width can be used in decoding subsequent P2s. In most cases this should be enough to compensate for the reduced pilot boosting in the P2 symbols compared to the data symbols. When the FFT size is smaller than 16 K, in which case there is more than one P2 symbol per frame, some temporal smoothing of the channel measurements may also be considered, in order to reduce noise.

10.3.2.3.5 Other special cases

The foregoing clauses have outlined a strategy for channel estimation for the main body of payload symbols and also for the P2 symbol(s). There are some other special cases pertaining to data near the start and finish of the frame, where the sequence of Scattered Pilots is inevitably in some sense disturbed, and thus requiring special treatment.

10.3.2.3.5.1 Frame-closing symbol

A frame-closing symbol is included with certain combinations of FFT size, guard-interval fraction and pilot patterns, as listed in table 58 of [i.1]. It was introduced to ensure that when time interpolation is in use, it remains always possible, even at the end of the frame.

Without the frame-closing symbol, the diagonal pattern of Scattered Pilots would cause the problem that whilst *some* pilot-bearing carriers would have an SP in the last symbol of the frame, others would not - i.e. the last SP to occur within the frame on such pilot-bearing carriers would be in an earlier symbol than the last. It follows that time interpolation for all symbols *after* this occurrence could not be performed and would appear to require *extrapolation* instead (with the associated unwanted loss of performance).

This problem is avoided by incorporating the frame-closing symbol. It has pilots (having the same degree of pilot boosting as for 'regular' SPs) inserted on every carrier that elsewhere is pilot bearing. This ensures that for all pilot-bearing carriers the gap between successive pilots (SPs or frame-closing ones) is never greater than occurs in the 'regular' SP pattern. It thus remains possible for the receiver to perform interpolation in time along pilot-bearing carriers, before performing interpolation in frequency within each symbol, without needing any additional memory beyond that required in the main body of the frame. This is illustrated in figure 87.

Note that all pilots involved in the process just described have the same power boost (see clauses 9.2.3.2 and 9.2.7.2 of [i.1]) so no significant issues arise concerning the amount of noise on the interpolated channel estimates.

The channel estimate for the Frame-Closing Symbol itself is expected to be obtained by frequency-only interpolation.

If temporal interpolation is being performed in modes for which it was not originally intended (see clause 10.3.2.3.3.2), there will be no frame closing symbol and alternative methods will need to be used at the end of the T2-frame.

10.3.2.3.5.2 Interpolation between P2 and subsequent 'normal' payload symbols

In some ways the P2 symbol(s) can be considered to offer similar uses at the *beginning* of the frame as offered at the *end* of the frame by the frame-closing symbol. Since the P2 symbols contain pilots on at least every pilot-bearing carrier, time interpolation can always be performed along pilot-bearing carriers, without needing additional memory or resort to extrapolation. This is also illustrated in figure 87.

There is a small difference from the end-of-frame case since the power boost for pilots in the P2 symbol(s) is less than for the SPs in the main body, see clauses 9.2.3.2 and 9.2.6.2 of [i.1]. It follows that channel estimates derived using simple linear temporal interpolation between P2 pilots and SPs will be noisier than for those wholly derived from SPs. However, in most such cases the P2 pilots will have a closer spacing than the Scattered Pilot-bearing carriers, and so it should be possible to apply some smoothing of the P2 pilots in the frequency direction, to obtain less noisy estimates prior to temporal interpolation.

Where there is more than one P2 symbol, it is also possible to use a more complicated temporal interpolator (having > 2 taps) for some cells near the beginning of the frame, without needing any additional memory for the main data path. This will permit some optimisation of the interpolation with respect to noise performance.

10.3.2.3.5.3 Implications of minimum frame length

The previous clauses have explained how special variation in temporal interpolation is needed at the beginning and end of the frame. It should be noted that the minimum frame length is so specified (see clause 8.3.1 of [i.1]) that these two effects never need to be compounded. Every data cell on a pilot-bearing carrier can be treated as one and one only of the following cases:

- Main body of frame (normal simple linear interpolation between two SPs).
- At or near end of frame (simple linear interpolation between an SP and frame-closing pilot).
- At or near beginning of frame (interpolation between one or more P2 pilots and an SP).

Furthermore, in any given symbol, there will never be a mixture of data cells belonging to the second and third categories. This should make "housekeeping" in the receiver easier.

10.3.2.3.5.4 Channel estimation with PP8

When the Scattered Pilot pattern PP8 is in use, the 'conventional' application of temporal and frequency interpolation described so far cannot be used. With this pattern the spacing between SPs on a pilot-bearing carrier is considered too great to be able to perform temporal interpolation on grounds of the receiver memory that would be needed.

This pattern is really intended to be used in conjunction with the so-called CD3-OFDM technique. In this the receiver starts from an initial channel measurement for every carrier and updates this by a decision-feedback technique. Once the data for a particular cell has been decoded, it is possible to determine the constellation point that was transmitted and thus use the cell as received as if it were a pilot. In this way the channel estimate for every carrier can be updated as reception proceeds. However, if a mistake is made in decoding, the channel estimate will be updated incorrectly. Frequency and Time filtering of the channel frequency response (as proposed in [i.21]) can prevent this unwanted effect.

The use of PP8 does however permit a simple alternative to the use of CD3-OFDM demodulation: instead of temporal interpolation along the pilot-bearing carriers, the receiver simply stores the most-recently received measurement for each pilot-bearing carrier. This is sometimes described as "zero-order hold" and is a trivial example of extrapolation. Clearly, it would require an essentially static channel, e.g. reception with a roof-top antenna, in favourable circumstances where Doppler effects are not expected.

In DVB-T2, the CD3-OFDM technique can only practically be used in single PLP modes because it relies on decoding and re-encoding all the data cells. In multiple PLP modes the receiver is not expected to decode cells belonging to other PLPs and therefore cannot generate channel estimates for those cells.


10.3.2.3.6 Further observations

The procedure outlined in the previous clause derives all channel estimation for data cells within a T2 frame from measurements made within the same the T2 frame. This has the merit of being undisturbed when T2 frames are separated by FEFs (whose internal format is currently undefined, apart from the need to start with a P1 symbol).

Not all pilots are necessarily inserted with the express intention of being used for channel measurement. E.g. while continual pilots could, in principle, be included in the frequency-interpolation process, it is not necessary to incur the complication of doing so - they are mainly there for the purposes of measuring frequency offset and common phase error. The precise number of CPs, and of pilots in the frame-closing symbol, is also at least in part chosen to meet other constraints, in effect to ensure the right number of data cells, see clause 9.3.4.1.

10.3.2.4 MISO channel estimation

The basic principle of MISO channel estimation is that the pilots can be partitioned into two subsets: in one subset the pilots are transmitted in the same phase from both sets of transmitters, allowing the sum of the two channel responses to be estimated; in the other subset the pilots are inverted from transmitters in MISO group 2, allowing the difference to be estimated. Figure 89 shows an example pilot pattern; the "sum" pilots are solid whilst the "difference" pilots are shown hatched. The different types of pilot are colour-coded as in figure 87.



● "Sum" pilot: not inverted from either Tx

⊘ "Difference" pilot: inverted from Tx2

Figure 89: An example MISO pilot pattern showing the "sum" and "difference" pilots

The sum and difference should be measured from the pilots in each subset and then interpolation performed in frequency and time in order to obtain an estimate of the sum and difference of the channel response on every carrier and every symbol. The sum and difference estimates for each cell can then be added to give the response from transmitter group 1, and subtracted to give the response from group 2.

The principles of interpolation are similar to those outlined for SISO in clause 10.3.2.3, except for the following points:

- The pilots on a given symbol will be either all sum or all difference pilots. The receiver will therefore always need to carry out temporal interpolation in order to obtain sum and difference estimates for each symbol.
- Since temporal interpolation is always used in MISO, there is always a frame closing symbol in MISO mode.

- In SISO, the edge carriers are always pilots, allowing the response to be interpolated up to the edge carrier. In MISO, these carriers are also Edge Pilots, but they alternate between sum and difference on alternate symbols. The receiver can therefore obtain an estimate of one response directly and should carry out temporal interpolation between previous and following symbols to obtain the other estimate. Frequency interpolation can then be carried out up to the edge carrier.
- Estimating the response on the edge carrier in the frame closing symbol is more difficult. The Edge Pilot will give one of the estimates (sum or difference), but the other estimate cannot be made by temporal interpolation in the normal way. Possible approaches are to use zero-order hold or to interpolate into the next frame, assuming no FEF part intervenes. The number of cells affected in this way will be a relatively small proportion of the whole frame and performance is not expected to be significantly degraded.
- Since the pilot-bearing carriers across frequency are alternately of sum and difference types, the Nyquist limit for the delay range (or *timewidth* - see clause 10.3.2.3.3.3) which can be estimated is halved compared to SISO for a given pilot pattern. This is reflected in the choice of Scattered Pilot patterns available in each guard interval and FFT size (see table 53 of [i.1]).

10.3.3 MISO decoding

This clause describes MISO decoding in the receiver for a MISO processed transmission signal, according to clause 9.1 of [i.1].




Figure 90: System model for MISO diversity scheme

The diversity scheme used in DVB-T2 is based on the methodology described by Alamouti in [i.8]. Figure 90 depicts the system model for the modified Alamouti diversity scheme applied. Tx1 is a transmitter belonging to MISO group 1, whilst Tx2 belongs to MISO group 2.

When MISO is used, the respective payload cells are processed in the frequency direction. The transmitter-side signal treatment is defined in such a way that for transmitters in MISO group 1, neither buffering nor negation or complex conjugation is applied, i.e. both $a_{m,l,p}$ and $a_{m,l,p+1}$ remain unmodified regarding frequency order or arithmetic operations. However, the signal transmitted by transmitters in MISO group 2 experiences pairwise modification, i.e. Tx2 transmits $-a_{m,l,p+1}^*$ while Tx1 applies $a_{m,l,p}$, and Tx2 transmits $a_{m,l,p}^*$ while Tx1 transmits $a_{m,l,p+1}$, where $*$ denotes the complex conjugation operation. Note that the two cells $a_{m,l,p}$ and $a_{m,l,p+1}$ being treated as a pair by the Alamouti processing have no special relationship and might come from different PLPs.

Accordingly, the signal components can be recovered at the receiver side, as follows. The received complex values for the cells of a MISO pair, r_1 and r_2 are given by:

$$r_1 = H_1 a_{m,l,p} - H_2 a_{m,l,p+1}^* + n_1$$

$$r_2 = H_1 a_{m,l,p+1} - H_2 a_{m,l,p}^* + n_2$$

where H_1 and H_2 describe the channel transfer function between Tx1 and Rx and Tx2 and Rx respectively, and with n_1 and n_2 characterising the noise terms. The diversity signal may be used with an arbitrary number of receive aerials. The received pair of cells can be represented by a vector r , where:

$$r = \begin{pmatrix} r_1 \\ r_2^* \end{pmatrix} = H a + n$$

in which:

$$H = \begin{pmatrix} H_1 & -H_2 \\ H_2^* & H_1^* \end{pmatrix}, \quad a = \begin{pmatrix} a_{m,l,p} \\ a_{m,l,p+1}^* \end{pmatrix}, \quad n = \begin{pmatrix} n_1 \\ n_2^* \end{pmatrix},$$

and the recombination of the received signals may finally be performed by:

$$a' = H^H r = \begin{pmatrix} |H_1|^2 + |H_2|^2 & 0 \\ 0 & |H_1|^2 + |H_2|^2 \end{pmatrix} \begin{pmatrix} a_{m,l,p} \\ a_{m,l,p+1}^* \end{pmatrix} + n'$$

10.3.3.1 ZF and MMSE Alamouti matrix inversion

The analysis above assumes that the channel response is the same on both cells of an Alamouti pair. In real-world channels the variation of the channel response is likely to be significant, and this assumption is not valid. It is therefore recommended that a more sophisticated Alamouti decoding process be used. In the varying-channel case, the received vector \mathbf{r} is:

$$\mathbf{r} = \begin{pmatrix} r_1 \\ r_2^* \end{pmatrix} = H \mathbf{a} + \mathbf{n}$$

in which:

$$H = \begin{pmatrix} h_{00} & -h_{01} \\ h_{10}^* & h_{11}^* \end{pmatrix}, \quad \mathbf{a} = \begin{pmatrix} a_{m,l,p} \\ a_{m,l,p+1}^* \end{pmatrix}, \quad \mathbf{n} = \begin{pmatrix} n_1 \\ n_2^* \end{pmatrix},$$

where all four h 's are distinct. There are a number of approaches to solving this problem. The most obvious is to use the simple matrix inverse, in which case the Zero-Forcing (ZF) solution results:

$$\hat{\mathbf{a}} = H^{-1} \mathbf{r}$$

Other pseudo-inverse solutions also exist, including the MMSE solution:

$$\hat{\mathbf{a}} = (H^* H + \alpha^2 I)^{-1} H^* \mathbf{r}$$

where H^* represents the conjugate transpose of H and α is the noise standard deviation.

10.3.3.2 Separation of Alamouti pairs by pilot cells

In normal data and frame-closing symbols, the cells making up a MISO pair will sometimes be separated by a pilot, since no special measures are taken in those symbols to avoid this. Care should therefore be taken to ensure that the correct channel estimates are used in the Alamouti decoding process.

In frequency-selective channels, the channel variation between the cells of such a pair will be greater than for adjacent cells, so the performance will be degraded even for the full MMSE matrix inversion. However, since these pairs only represent a relatively small proportion of the data cells, it has been found that the impact on performance is minimal, particularly since the guard interval, and therefore the maximum tolerable echo delay, is limited to 1/8 in MISO.

In the P2 symbols, special pilots have been inserted where necessary in order to prevent cells from being separated. This was necessary because the pilot spacing of one-in-three would mean that roughly half of all pairs would be separated in this way.

10.4 De-interleaving and de-mapping

10.4.1 Frequency/L1 de-interleaver

To minimise memory use the frequency de-interleaver may be implemented with a single memory comprising N_{\max} locations. N_{\max} is the maximum number of data cells in an OFDM symbol of the given FFT size. It is envisaged that practical implementations of DVB-T2 receivers should support all the FFT sizes included in [i.1]. In such implementations, N_{\max} would be the memory capacity necessary to implement odd-even interleaving for 32 K. Thus, $N_{\max} = 27\,404$ since this is the maximum number of data cells (C_{data}) in 32 K mode with PP8 Scattered Pilot pattern, extended carrier mode and no tone reservation (see table 42 of [i.1]). The frequency de-interleaver for all modes can thus be implemented with one memory having C_{\max} locations (for 32 K) or $2C_{\max}$ locations (for all other modes) where $C_{\max} = \max(C_{\text{data}})$ for the given mode. For most system configurations, not all of these locations would be needed.

NOTE: This is the reason why there are always an even number of symbols in 32 K. Otherwise the last symbol of one T2-frame and the P2 symbol of the next T2-frame could both be even, and odd-even interleaving using a single memory would not be possible.

The receiver buffer model (clause 8.8.1) assumes that the L1 and frequency de-interleavers are combined, and that the combined interleaver can be used for to provide a certain amount of buffering. Since this is included in the model that the T2-gateway will used to manage the receiver's buffers, it is recommended that the same implementation be used in receivers.

The function of the combined de-interleaver is to accept data cells in carrier order, and to output first the L1 cells followed by the data cells in order of cell address. It should also be able to retain cells of the P2 symbols to allow time for the L1 signalling to be decoded. First, the L1-post cells need to be kept in the de-interleaver until the L1-pre has been decoded, so that the receiver knows how many L1-post blocks and how many cells per L1-post block there are. Second, the PLP cells need to be kept while the L1-post is decoded, so that the receiver knows which data cells belong to the current PLP or PLPs. Finally the remaining data cells in the P2 symbols should be output at a slightly higher rate than normal (R_s), in order to finish reading them before the next symbol begins.

In the case of 32 K or 16 K, there is only one P2 symbol, so the ordinary frequency de-interleaver will naturally output the L1 cells first followed by the PLP cells. In other modes, the "zig-zag" mapping needs to be undone.

10.4.1.1 Implementation of L1/Frequency de-interleaver 32 K

For odd-even interleaving as used in 32 K, the de-interleaver should read data cells from even symbols (symbol number of form $2n$) into its memory in a sequential order. For odd symbols (symbol number of form $2n + 1$) the data cells are read into memory in a permuted order, the permuted order addresses $H(p)$ being provided by the relevant pseudo-random address generator from clause 8.5 of [i.1]. As only one memory is used in the de-interleaver and in order to avoid losing data cells, one cell of the previous symbol should be read out from the memory location to which the cell of the current symbol is to be written. It therefore follows that the sequence of write addresses for symbol number $2n + 1$ should match the sequence of read addresses for symbol number $2n$ otherwise some data cells of symbol $2n$ will be overwritten. Thus, if all symbols contained the same number $C_{\max} = C_{\text{data}}$ of data cells, then starting from symbol $2n$ at the input of the de-interleaver, the de-interleaver would follow the following steps:

- 1) $p = 0$;
- 2) Read cell p of output de-interleaved symbol $2n - 1$ from location p of the memory;
- 3) Write cell p of incoming interleaved symbol $2n$ into location p of the memory;
- 4) Increment p ;
- 5) if ($p < C_{\max}$) go to 2.

Then with symbol $2n + 1$ at the input of the de-interleaver:

- 6) $p = 0$;
- 7) Generate address $H(p)$;

- 8) Read cell p of output de-interleaved symbol $2n$ from location $H(p)$ of the memory;
- 9) Write cell p of incoming interleaved symbol $2n + 1$ into location $H(p)$ of the memory;
- 10) Increment p ;
- 11) if ($p < C_{\max}$) goto 7.

Note that in 32 K, there is always an even number of symbols in a T2-frame, so the odd-even sequence continues, even across frame boundaries. In other FFT sizes there can be an odd number of symbols but this is not a problem since there is always at least two-symbols' worth of memory available.

Things are more complicated, however, because the T2-frame in DVB-T2 is composed of different types of symbol with correspondingly different number of data cells for each type of symbol. Thus for a given system configuration (FFT size, Scattered Pilot pattern, bandwidth extension and tone-reservation status), a P2 symbol carries fewer data cells than a data symbol which in turn carries more data cells than the frame-closing symbol. Imagine that the de-interleaver is about to process symbol $2n + 1$ (a data symbol) for which the preceding symbol $2n$ was a P2 symbol. Then since data symbols have more data cells than P2 symbols, it follows that the number of generated addresses in the loop 7 to 11 above would exceed the number of data cells carried by the P2 symbol. Therefore, the pseudo-random address generator would produce some addresses that are valid for writing cells of symbol $2n + 1$ into memory but not valid for reading cells of symbol $2n$. The case in which symbol $2n$ has more data cells than symbol $2n + 1$ also occurs when going from a data symbol to a frame-closing symbol. For a given system configuration therefore $C_{\max} = \max(C_{\text{data}}, C_{\text{P2}}, C_{\text{FC}}) = C_{\text{data}}$. The loops above should be changed as follows where the function $\text{DataCells}(l)$ returns the number of data cells in symbol l and HoldBuffer is a small amount of storage with write address $wptr$ and read address $rptr$:

- 1) $p = 0$; $wptr = rptr = 0$;
- 2) $\text{rdEnable} = (p < \text{DataCells}(2n - 1))$;
- 3) $\text{wrEnable} = (p < \text{DataCells}(2n))$;
- 4) if (rdEnable) Read cell p of output de-interleaved symbol $2n - 1$ from location p of the memory;
- 5) Store cell p of incoming interleaved symbol $2n$ into location $wptr$ of HoldBuffer and increment $wptr$;
- 6) If (wrEnable):
 - a) Write cell $rptr$ of HoldBuffer into location p of the memory and increment $rptr$.
 - b) If ($rptr == wptr$) then reset $rptr = wptr = 0$.
- 7) Increment p ;
- 8) if ($p < C_{\max}$) goto 2;

Then with symbol $2n + 1$ at the input of the de-interleaver:

- 9) $p = 0$; $wptr = rptr = 0$;
- 10) Generate address $H(p)$;
- 11) $\text{rdEnable} = (H(p) < \text{DataCells}(2n - 1))$;
- 12) $\text{wrEnable} = (H(p) < \text{DataCells}(2n))$;
- 13) if (rdEnable) Read cell p of output de-interleaved symbol $2n$ from location $H(p)$ of the memory;
- 14) Store cell p of incoming interleaved symbol $2n + 1$ into location $wptr$ of HoldBuffer and increment $wptr$;

- 15) If (wrEnable):
 - a) Write cell $rptr$ of HoldBuffer into location $H(p)$ of the memory and increment $rptr$.
 - b) If ($rptr == wptr$) then reset $rptr = wptr = 0$.
- 16) Increment p ;
- 17) if ($p < C_{\max}$) goto 10.

10.4.1.2 Implementation of the basic frequency de-interleaver in 1 K-16 K

Similarly for modes other than 32 K that use odd-only interleaving, with symbol $2n$ at its input the de-interleaver steps for the basic frequency de-interleaver are as follows. Modifications to implement the combined L1/Frequency de-interleaver are described in clause 10.4.1.3.

- 1) $p = 0$; $wptr = rptr = 0$;
- 2) Generate address $H_0(p)$;
- 3) Generate address $H_1(p)$;
- 4) $rdEnable = (H_1(p) < DataCells(2n - 1))$;
- 5) $wrEnable = (H_0(p) < DataCells(2n))$;
- 6) if (rdEnable) Read cell p of output de-interleaved symbol $2n - 1$ from location $C_{\max} + H_1(p)$ of the memory;
- 7) Store cell p of incoming interleaved symbol $2n$ into location $wptr$ of HoldBuffer and increment $wptr$;
- 8) If (wrEnable):
 - a) Write cell $rptr$ of HoldBuffer into location $H_0(p)$ of the memory and increment $rptr$.
 - b) If ($rptr == wptr$) then reset $rptr = wptr = 0$.
- 9) Increment p ;
- 10) if ($p < C_{\max}$) goto 2.

Then with symbol $2n+1$ at the input of the de-interleaver:

- 11) $p = 0$; $wptr = rptr = 0$;
- 12) Generate address $H_0(p)$;
- 13) Generate address $H_1(p)$;
- 14) $rdEnable = (H_0(p) < DataCells(2n))$;
- 15) $wrEnable = (H_1(p) < DataCells(2n + 1))$;
- 16) if (rdEnable) Read cell p of output de-interleaved symbol $2n$ from location $H_0(p)$ of the memory;
- 17) Store cell p of incoming interleaved symbol $2n + 1$ into location $wptr$ of HoldBuffer and increment $wptr$;

- 18) If (wrEnable):
- Write cell $rp\text{tr}$ of HoldBuffer into location $C_{\text{max}} + H_1(p)$ of the memory and increment $rp\text{tr}$.
 - If ($rp\text{tr} == wp\text{tr}$) then reset $rp\text{tr} = wp\text{tr} = 0$.
- 19) Increment p ;
- 20) if ($p < C_{\text{max}}$) goto 12.

The required width for each memory location depends on the resolution with which each cell is represented after channel equalisation. Each de-interleaver memory cell would hold at least: the complex cell information and the channel state information for the cell. If rotated constellations were used, the complex channel coefficient might also be included.

10.4.1.3 Implementation of the L1/ Frequency De-interleaver in 1 K-16 K

Figure 91 shows how the combined de-interleaver can be implemented in memory. The figure is for 8 K, with two P2 symbols, but the principle can easily be extended to other modes. The red and green cells are L1-pre- and post- cells respectively, and the blue cells are data cells.




Figure 91: Combined L1/Frequency de-interleaver

Each row of the memory is used to store one symbol. In all modes except 32 K, the size of the de-interleaver memory will always be twice the number of P2 symbols. Whilst the P2 symbols are being received, they are stored in the memory (a and b in figure 91). Because of the way the interleaving is defined, the de-interleaver should write in permuted order, as described in clause 10.4.1.2. Once all the P2 symbols have been received, they can start to be read out; meanwhile the following data symbols are written to the other half of the memory (c). The effect of writing in permuted order is that the cells are now in de-interleaved order in the memory. The figure (c) shows how the L1-pre cells can be read out in order, followed by the L1-post, followed by the data cells. Once the P2 symbols have been read out, the rows that have been freed up can be used for the next data symbols (d).

The de-interleaver as shown will always have approximately N_{P2} symbols' delay.

10.4.2 Frame extraction

Following the frequency de-interleaving, the cells will be in the correct order to reverse the function of the frame builder, and so extract the L1 signalling and the PLPs.

10.4.2.1 Extraction of the PLPs

The PLP (or PLPs) to be received should be extracted based on the PLP_START and PLP_NUM_BLOCKS parameters as described in clause 7.2.3.2 of [i.1]. These parameters are carried in the P2 symbol(s) and may also be signalled using in-band signalling (see clause 10.4.2.4).

The specification describes how the slices or sub-slices of the PLPs should be placed immediately one after the other, with the common PLPs starting from the first cell not used for L1 signalling, followed immediately by the type 1 and then type 2 PLPs.

However, this should not be assumed by the receiver, since future applications may introduce other uses for the cells, which might involve inserting cells between PLPs, between sub-slices or between the PLPs of different types. In all cases the locations as signalled by PLP_START and PLP_NUM_BLOCKS will be correct.

Any data cells which are not allocated to one of the PLPs being decoded should be discarded.

10.4.2.2 Extraction of the L1 signalling cells

The L1 signalling cells are carried in the P2 symbols as described in clause 8.5.3 of [i.1]. When the cells are extracted, the zig-zag order should be undone; this can be considered as a de-interleaving process, and should be implemented as part of the frequency de-interleaver as described in clause 10.4.1. This is both an efficient use of memory and also corresponds to the assumptions in the receiver buffer model.

Since L1-pre signalling is mapped to 1 840 BPSK symbols in N_{P2} P2 symbols as described in clause 8.3.5 of [i.1], the received signal for L1-pre signalling should be obtained from the first $1\ 840/N_{P2}$ BPSK symbols in each P2 symbol.

The process is similar for the L1-post signalling cells, except that the receiver needs to determine how many there are; this information is signalled directly by the L1-pre signalling parameter L1_POST_SIZE. In normal reception, this will already be known since it is statically configurable. However, on first acquisition, or if the value has changed since the frequency was last visited, the L1-pre signalling will need to be decoded before the L1-post cells can be extracted.

NOTE: The field L1_POST_SIZE is provided for convenience, but the same information could be deduced from the field L1_POST_INFO_SIZE as described in clause 10.5.7.1.

10.4.2.3 Timing of L1 and PLP extraction and decoding in the P2 symbols

The timing of signal flows during the P2 symbols can be quite critical, because of the way each of the L1-pre, L1-post and PLP extraction and decoding depends on the previous stage.

Figure 92 shows a timeline of the process, with the time taken for each stage according to the receiver buffer model. Some of the stages can take place in parallel; in particular the FEC chain can be receiving input cells, outputting decoded cells and decoding a block at the same time.

As can be seen, the critical path comprises extracting the L1-pre cells, decoding the L1-pre, outputting the decoded L1-pre bits, extracting the L1-post cells for the first block, $N_{\text{post_FEC_block}}-1$ intermediate stages, and finally decoding and outputting the last L1-post block.

The "extracting cells" stage includes soft demapping (see clause 10.5.1). The model assumes that this can be done at a maximum rate of $R_S=1/T$ for the L1-post, which might use up to 64-QAM. For the L1-pre, BPSK is always used, so soft-demapping is trivial. In this case the model assumes that the demapping can be done eight times faster. Implementers should therefore provide a simple 1-D demapper for BPSK capable of working at this rate, rather than using the full 2D demapper which is unlikely to be able to operate quickly enough.

In the figure, the intermediate stages are limited by the decoding time for a FEC block. However, in another configuration the time taken to extract (and soft-demap) the cells of the next block might be longer than the decoding time, particularly if a low-order constellation is used for the L1-post.




Figure 92: Timeline for extraction and FEC decoding for the P2 symbols

We can therefore calculate the complete time for L1 extraction and decoding:

$$T_{L1} = \frac{1840}{8}T + \frac{2025}{R_{cell}} + \frac{200}{8}T + T_{extract} + (N_{post_FEC_block} - 1) \times \max(T_{extract}, T_{decode}) + T_{decode} + \frac{K_{sig_post}}{8}T,$$

where $T_{extract} = N_MOD_PER_BLOCK \times T$ and $T_{decode} = \frac{2025}{R_{cell}}$ are the times to demap and decode one L1-post FEC block, respectively.

10.4.2.4 Use of in-band dynamic signalling

10.4.2.4.1 Interleaving over one T2 frame or less

The T2 receiver can obtain L1-dynamic signalling from in-band signalling as well as P2 symbols. The in-band signalling is carried in the PADDING field of the first BBFRAME of each PLP, when IN-BAND_FLAG field in L1-post signalling, defined in EN 302 755 [i.1], is set to '1' and one Interleaving Frame is mapped to one T2-frame (i.e. the values for P_I and I_{JUMP} for the current PLP are both equal to 1; see clauses 6.5 and 8.2 of [i.1]).

If NUM_OTHER_PLP_IN_BAND field in the in-band signalling is set to '0', the in-band signalling of a PLP carries only its own L1 dynamic information. On the other hand, if the NUM_OTHER_PLP_IN_BAND field is larger than '0', the in-band signalling of a PLP carries L1 dynamic information of other PLPs as well as its own information, for shorter channel switching time. The in-band signalling in T2-Frame n carries L1 dynamic information of T2-frame $n + 1$. Therefore, the T2 receiver can obtain the L1-dynamic signalling of the next T2-frame through the in-band signalling received in the current T2-frame and can decode PLPs of the next T2-frame without decoding the P2 symbols of the next frame.

10.4.2.4.2 In-band signalling when interleaving over multiple T2 frames

When multi-frame interleaving is used, i.e. $P_1 > 1$ so that one Interleaving Frame is mapped to more than one T2 frame, in-band L1-dynamic signalling can still be used. In this case the in-band signalling of a PLP carries only its own L1 dynamic information, and the information applies to the T2 frames to which the next Interleaving Frame is mapped.

The information applies to the next Interleaving Frame because the receiver cannot decode any of the FEC blocks and recover the BBFRAME containing the in-band L1 until the whole Interleaving Frame has been received; any information for the current Interleaving frame would be decoded too late to be of any use.

A loop is used to carry dynamic L1 information for each of the P_1 frames to which the next Interleaving Frame will be mapped. This is necessary because the BBFRAMES do not belong to any one T2 frame so cannot meaningfully signal only the next T2 frame's parameters. Furthermore, there might only be one BBFRAME in the whole Interleaving Frame.

There is therefore only one place in the interleaving frame in which the in-band signalling can be sent: the first BBFRAME of the Interleaving Frame.

Carrying in-band L1 dynamic information for other PLPs in this case would have made the signalling excessively complicated and was not thought to be worthwhile.

Note that in-band signalling is not mandatory in this case, since it requires extra memory in the modulator and incurs extra end-to-end delay.

10.4.2.4.3 In-band signalling for PLPs not present in every T2-frame

In-band signalling can also be used for PLPs which are not present in every T2 frame. In this case, the signalling applies to the T2 frame to which the next Interleaving Frame is mapped. If each Interleaving Frame is mapped to one T2 frame, the signalling therefore applies to the next T2 frame to which the PLP is mapped. For multi-frame interleaving, the dynamic information is signalled in a loop for each of the T2 frames to which the next interleaving frame is mapped.

Again, in-band signalling is not mandatory in this case because of the implications for memory in the modulator and end-to-end delay.

10.4.2.4.4 Reading the P2 symbols on reconfiguration

It can be signalled in the in-band signalling where the configuration (i.e. the contents of the fields in the L1-pre signalling or the L1-post signalling) will change in a way that affects the PLPs referred to by this in-band signalling field. This is indicated by **PLP_L1_CHANGE_COUNTER** field in the in-band signalling of 'In-band type A'. If **PLP_L1_CHANGE_COUNTER** field is set to the value '0', it means that no scheduled change is foreseen. For example, a value of '1' indicates that there will be a change in the next superframe. This counter will always start counting down from a minimum value of 2. At the first frame of the superframe indicated by **PLP_L1_CHANGE_COUNTER** field, the T2 receiver should decode the P2 symbols and update the L1-post signalling.

10.4.2.4.5 PLPs with zero FEC blocks in a given Interleaving Frame

EN 302 755 [i.1] explains that PLPs might be allocated no FEC blocks at all in a given Interleaving Frame, and there is therefore no BBFRAME frame to carry the in-band L1 signalling. The receiver will already know that this is the case from the in-band L1 signalling in the previous frame. There are two options for the receiver:

- It could wait until the subsequent frame and decode the P2 which always carries the dynamic L1 signalling for all PLPs.

- If the PLP carries in-band signalling for other PLPs, the receiver could use this information from the previous frame to receive one of these PLPs instead. These PLPs might in return contain in-band signalling for the wanted PLP.

10.4.2.5 Power saving

Receivers may save power by switching off some of the circuitry while signals for the PLP of interest are not being transmitted. This is most likely to be applicable to Type-1 PLPs, since they only have one sub-slice per T2 frame, and furthermore this sub-slice generally appears near the beginning of the T2 frame, close to the P2 (containing the L1 signalling) and any common PLPs (generally containing L2 signalling). All the information required is therefore close together at the beginning of the frame and the receiver need not receive the rest of the frame.

The use of in-band signalling can increase the possibility for power saving, since for most of the time even the L1 signalling in P2 will not need to be decoded: the dynamic information will already be known from the in-band signalling in the previous Interleaving Frame and any changes in the configurable information will be announced in advance in the in-band signalling.

The scope for saving power is greater still for PLPs which skip T2 frames, i.e. for which $I_{\text{jump}} > 1$. The receiver may ignore the frames which do not contain data for the PLP being received. In this case, in-band signalling is not mandatory and the receiver may need to receive the P2 signals.

The considerations for synchronisation in a power-saving receiver would be more complicated since the receiver would not be tracking the signal during the power-down period.

10.4.2.6 FEFs

The receiver can detect the presence of FEF parts by looking for P1 symbols and by decoding the L1-post signalling (in the P2 symbols). In the initial scan the receiver will look for P1-symbols. If the first P1-symbol found is from a T2-frame, P2 can be decoded normally. The receiver will also see from the P1 if FEFs are present by looking at the S2 field 2 "mixed" bit. If the first P1 symbol found belongs to a FEF part, the receiver can check S2 field 2 "mixed" bit and if it is set, then the receiver knows that there are also other (T2) frames in the multiplex and the receiver should continue looking for those. The use of P1, including in the "FEF loop" of the acquisition process is described in more detail in clause 10.2.2.

When a T2 P1 has been found, the L1-post signalling in P2 can be decoded and the position of the FEF parts in the superframe, as well as their duration, can be found. The receiver can then correctly skip the FEF parts.

The use of FEFs makes management of the receiver's de-jitter buffer more complicated, and the parameters affecting the configuration of FEFs should be chosen to ensure that, if a receiver obeys the TTO signalling (see annex C of [i.1]) and implements the model of buffer management defined in clause C.1.1 of [i.1], the receiver's de-jitter buffer and time de-interleaver memory will neither overflow nor underflow. See clause 8.8 for more information about managing the de-jitter buffer.

As the content of the FEF part is not known, receivers should be able to ignore totally the content of the FEF part except the P1-symbol. In order not to affect the reception of the T2 data signal, the receiver's automatic-gain control should be held constant for the duration of FEF part, so that it is not affected by any power variations during the FEF part. It is even possible that the FEF part is empty (see clause 6.1.7.3).

10.4.2.7 Auxiliary streams

Auxiliary streams are for future definition. A receiver may completely ignore them if they are present. However, a useful feature for a receiver chip would be the ability to output the complex cell values for a particular range of cell addresses, either on the main output or by reading a register. This could then be used by external hardware or software to implement future applications.

One application for Auxiliary Streams has been defined, as part of the Transmitter Signature standard [i.22]. This application is intended primarily for professional applications, but could be used in the future for location-aware services.

10.4.3 Time de-interleaver

10.4.3.1 Phase of Time de-interleaving

The receiver need not wait until the beginning of a superframe to start de-interleaving. Instead it should deduce, from the L1 signalling, which T2 frame is the first of the next Interleaving Frame and begin de-interleaving from that point. In many cases one Interleaving Frame will be mapped to one T2 frame so the de-interleaving can begin immediately. Only for PLPs whose interleaving length $P_I \times I_{\text{jump}}$ is equal to N_{T2} , the number of T2 frames in a superframe, will the receiver need to wait for the next superframe.

10.4.3.2 Memory-efficient implementation of time de-interleaver

The permutation function for the time interleaver is the same for each TI-block and so it might appear that two blocks of memory are required in the de-interleaver. However, a more efficient method exists using only one TI-block's worth of memory.

10.4.3.2.1 De-interleaver for constant bit-rate

The de-interleaving method for a constant bit-rate will be described, as this forms the core of the method for the variable bit-rate case. The de-interleaver for the constant bit-rate case is almost identical. The basic de-interleaver is shown in the block diagram shown in figure 93.




Figure 93: De-interleaver block diagram

During each TI-block (except the first) data cells are read out one at a time from the de-interleaver memory according to the addressing sequence produced by the address generator. For each cell read out, a new cell from the input is written into the memory at the same address, as this memory location has just been cleared by reading the output cell.

The time interleaver is defined to have data input into c columns, with each column having r rows. The de-interleaver will therefore write data into the r rows, as shown in figure 94(a), and read in the c columns. The total memory of the interleaver is therefore defined by:

$$M = r \times c$$




Figure 94: Implementation of de-interleaver memory, showing (a) the conceptual implementation with rows and columns which defines the de-interleaving sequence and (b) the actual implementation as a sequential block of memory

The actual block of memory to be used will be implemented as a sequential block of memory, as shown in figure 94(b), and the main problem to be solved is to calculate the correct address sequence. The addresses of the elements of the memory will be calculated by index i , where:

$$0 \leq i \leq M - 1$$

For the first TI-block of data, data is received and written into the memory in sequential order. During this TI-block only, when the de-interleaver memory is first being filled, no output data is produced. Once the whole TI-block has been input, data is ready to be output. The address generator should calculate the output sequence, which can be done from the definition of the interleaver.

The first column of the output sequence will be found by addressing the sequential memory in the order:

$$0, c, 2c, 3c, \dots, (r - 1).c$$

The address sequence for the second column will be:

$$1, c + 1, 2c + 1, 3c + 1, \dots, (r - 1).c + 1$$

In general, the address generator will calculate an initial value for each new address from the previous address by adding an increment, k_j , which varies with TI-block number j . Whilst the very first input TI-block was being received ($j = 0$), $k_0 = 1$. When the second TI-block of input data is being received, the first TI-block will be read out, and so the increment k_1 will be equal to c . To calculate the correct address for all columns, the initial address calculated is modified according to:

$$\left. \begin{aligned} a_{0,1} &= 0 \\ a_{i,1} &= (a_{i-1,1} + c) \bmod M + (a_{i-1,1} + c) \operatorname{div} M \end{aligned} \right\} \text{for } i > 0$$

Where:

$a_{i,j}$ is the address of the i^{th} element of the j^{th} input data TI-block.

mod is the integer modulo operator.

div is the integer division operator.

In practice, since the value $(a_{i-1,1} + c)$ will always be less than $2M$, this equation can be simplified to:

$$\left. \begin{aligned} a_{0,1} &= 0 \\ a_{i,1} &= (a_{i-1,1} + c) && \text{when } (a_{i-1,1} + c) < M \\ a_{i,1} &= (a_{i-1,1} + c) - (M - 1) && \text{when } (a_{i-1,1} + c) \geq M \end{aligned} \right\} \text{for } i > 0$$

or:

$$\left. \begin{aligned} a_{0,1} &= 0 \\ a_{i,1} &= (a_{i-1,1} + c) && \text{when } a_{i-1,1} < M - c \\ a_{i,1} &= (a_{i-1,1} - (M - c - 1)) && \text{when } a_{i-1,1} \geq M - c \end{aligned} \right\} \text{for } i > 0$$

Consider as an example the case when the de-interleaver has 8 rows and 5 columns. The address sequence generated for this TI-block would be 0, 5, 10, 15, 20, 25, 30, 35, 1, 6, 11...29, 34, 39.

This describes the addresses to be generated to read out the first TI-block of de-interleaved data. However, the same address will also be used to write the new data symbols. Therefore, when this second input TI-block has all been stored, it too will be need to be read out in de-interleaved sequence.

Once again, the addresses of the output sequence can be calculated using an increment of c , but this refers to increments through the interleaved sequence of the previous TI-block. Returning to the previous example, every 5th address is required from the sequence 0, 5, 10, 15, 20, 25, 30, 35, 1, 6, 11...29, 34, 39. In other words, the actual memory addresses will be given by the sequence 0, 25, 11, 36, 22, 8, ...28, 14, 39. This sequence can be calculated exactly as before, but this time the increment is now c^2 :

$$\left. \begin{aligned} a_{0,2} &= 0 \\ a_{i,2} &= (a_{i-1,2} + c^2) && \text{when } a_{i-1,2} < M - c^2 \\ a_{i,1} &= (a_{i-1,2} - (M - c^2 - 1)) && \text{when } a_{i-1,2} \geq M - c^2 \end{aligned} \right\} \text{for } i > 0$$

The process now continues as before with data being read and written continuously. In general, the address sequence for any input frame j of data can be calculated using the increment k_j :

$$\left. \begin{aligned} a_{0,j} &= 0 \\ a_{i,j} &= (a_{i-1,j} + k_j) && \text{when } a_{i-1,j} < M - k_j \\ a_{i,j} &= (a_{i-1,j} - (M - k_j - 1)) && \text{when } a_{i-1,j} \geq M - k_j \end{aligned} \right\} \text{for } i > 0$$

Where:

$$k_0 = 1$$

$$k_j = (k_{j-1} \cdot c) \bmod M + (k_{j-1} \cdot c) \operatorname{div} M \quad \text{which simplifies to:}$$

$$\left. \begin{aligned} k_0 &= 1 \\ k_j &= (k_{j-1} \cdot c) \bmod M + k_{j-1} \operatorname{div} r \end{aligned} \right]$$

So continuing with the previous example, $k_3 = 8$, and it can be seen that the addressing sequence 0, 8, 16, 24, 32, 1, ..., 23, 31, 39 will be correctly generated. For this example, the next frame generates $k_4 = 1$, and the interleaver returns to the initial state, where data is input in sequential order into the interleaver memory.

NOTE: It can be shown that the values of k_j produced from this last formula are always in the range $0..M-1$.

10.4.3.2.2 Implementation details

A practical implementation to generate the above address sequence is straightforward. At the start of each frame, the address generator is reset to 0. Subsequent addresses are then calculated by modulo M addition of the increment for that frame, k_j , with an extra 1 to be added if M is exceeded.

The value of k_j can easily be stored from the previous frame, since it will be equal to the address generated at the $(c+1)^{\text{th}}$ input value of the previous frame. So in the second frame of the above example, the address generator would wait until the 6th address had been generated (from the sequence 0, 5, 10, 15, 20, 25, 30, etc.), which is 25 in this case, and store this as the value of k_j for the third frame. Similarly in the third frame, the address generator stores the 6th value of the sequence 0, 25, 11, 36, 22, 8, etc. and correctly stores 8 as the value of k_3 . Finally the 6th value of the last sequence 0, 8, 16, 24, 1, etc. is stored and the generator will return to its initial state on the next frame.

10.4.3.2.3 Adaptation for variable bit-rates

The same basic structure can be used for the variable bit-rate case, dimensioned according to the maximum number of

FEC blocks per TI-block. This is simply given by $N_{\max} = \left\lceil \frac{PLP_NUM_BLOCKS_MAX}{N_{TI}} \right\rceil$. Since the

de-interleaver writes in rows, there will in general be unused columns, which the write sequence should skip. The receiver buffer model assumes a particular implementation, in which the "last" columns are filled and the first columns left empty, as shown in figure 95, in which N_{wr} FEC blocks are being written. The write pointer will skip $5(N_{\max} - N_{wr})$ locations (in the $a_{i,j}$ sequence) at the beginning of each row. At the same time, the read pointer will not skip any locations since there are no unused rows. Because the write pointer is skipping locations whilst the read pointer is not, there is a danger that the write pointer will overtake the read pointer and write to locations that have not yet been written. However, the receiver buffer model (clause 8.8.1) requires the signal to be constructed such that this will not in fact happen.

In practice this is not usually a problem, provided type 2 PLPs are used with a large number of subslices. This is because, in frames where the number of blocks is reduced, the subslices will get shorter in the same ratio. The skipping of columns will be cancelled out, on average, by the lower average arrival of cells from the channel.

The advantage of writing in the last columns is that, when the TI-block comes to be read, the cells in the unused columns will appear at the beginning of the address sequence. Consequently, the read pointer will start part-way through the sequence, by a number of locations equal to the number of cells in the unused columns, which is

$N_{cell} \times (N_{\max} - N_{wr})$. This effectively gives the read pointer a "head start" which can help in preventing the write pointer from catching it up.




Figure 95: Writing sequence in the de-interleaver when writing less than the maximum number of blocks

In single PLP modes, it can be difficult to guarantee that the TDI will not overflow near the beginning of the frame, because the FEC chain might be occupied decoding the L1. For this reason, the receiver buffer model assumes that there is a 4 000-cell FIFO associated with the TDI. This should be implemented in receivers. Either a FIFO can be placed before or after the TDI, or some additional memory could be added to the FEC decoder or cell de-interleaver.

Other implementations of the TDI are possible, but implementers using alternative implementations should ensure that they will work with the range of conditions allowed by the receiver buffer model. Implementations using two blocks of memory will normally meet this requirement.

10.4.3.3 Disabled Time Interleaving

The time de-interleaving may be disabled in cases where the receiver is expected to perform data-directed demodulation. In this case, it may still be useful to use the de-interleaver memory as a buffer, but the read and write sequences will both be identical.

10.4.3.4 Relationship to the receiver buffer model

The Time De-interleaver forms part of the receiver buffer model (see annex C of [1.1] and clause 8.8.1 of the present document). Implementers should pay attention to the assumptions described in the model, including the rate and timing of reading from the memory and the interaction with the de-jitter buffer. Implementers may assume that the T2-gateway will never require the time de-interleaver to contain cells from more than two TI-block at the same time: one being written and one being read, since this is mandated by the receiver buffer model.

10.4.4 Cell de-interleaver

The cell de-interleaver simply performs the inverse of the cell interleaver function (see clause 9.2.4). The cell de-interleaver on its own would require two FEC-blocks' worth of memory.

However, it may be possible to use the same memory as for the preceding or following stages (Time de-interleaver or bit de-interleaver).

10.5 Channel decoding

10.5.1 LLR demapping for non-rotated constellations

LDPC decoding should be performed using Log Likelihood-Ratios (LLRs), also known as soft decisions or metrics. When rotated constellations are not used, the LLRs can be derived in the normal one-dimensional way well-known from DVB-T. The derivation of LLRs in the case of rotated constellations is discussed in clause 10.5.3.

10.5.2 Removal of cyclic Q-delay

When rotated constellations are used, the cyclic Q delay needs to be removed. This can be done simply by delaying the 'I' component by one cell, to re-unite it with the corresponding 'Q' value; this requiring only one cell of memory. The 'Q' component of the first cell also needs to be stored until the end of the FEC block in order to re-unite it with its corresponding 'I' component.

10.5.3 De-mapping for rotated constellations

This clause describes the main implications on the receiver of the use of rotated constellations.

Rotated constellations require the receiver to form metrics - Log Likelihood-Ratios (LLRs) - as a function of two dimensions, rather than the conventional, variables-separable 1-D approach appropriate to conventional, unrotated QAM. Consider that the QAM constellation conveys m coded bits, i.e. it has 2^m possible states (constellation points).

10.5.3.1 2D LLR demapping

At the receiver, the classical demapper is replaced by the structure shown in figure 96:




Figure 96: 2D LLR demapper for rotated constellations

10.5.3.1.1 Computation of perfect LLR

We define as follows the LLRs that we wish the de-mapper of figure 96 to form from the received constellation I, Q ; this expression describes the LLR for the i th bit b_i :

$$\text{LLR}(b_i) = \ln \left(\frac{\Pr(b_i = 1 | I, Q)}{\Pr(b_i = 0 | I, Q)} \right)$$

A positive LLR would indicate that b_i was more probably transmitted as a 1, a negative value that b_i was more probably transmitted as a 0. How should the likelihoods, and hence the LLR be calculated?

A particular constellation point, x say, is transmitted with coordinates I_x, Q_x . Ideally the received constellation I, Q would be identical. In practice it differs for two reasons: the cells in which I and Q travelled have been subjected to amplitude-fading factors ρ_I, ρ_Q respectively and furthermore noise has been added. We can therefore write the following for the conditional probability distribution function (pdf) of receiving particular I, Q , given that we know point x was transmitted:

$$p(I, Q | x \text{ was transmitted}) = \frac{1}{2\pi\sigma^2} e^{-\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2}}$$

Now, consider reception of a particular bit b_i . If it was transmitted as a 1, then that implies that any of 2^{m-1} possible states was transmitted, which one depending of course on the state of the other $(m-1)$ bits. If it was a 0, then one of the other 2^{m-1} possible states would be transmitted. Let C_i^j denote the set of constellation points x for which the i th bit, b_i , takes the value j (0 or 1).

The conditional pdf for the received values I, Q , given that b_i was transmitted as a 1 is thus given by the following expression, under the assumption that all of the 2^{m-1} possible transmitted states for which $b_i = 1$ are transmitted with equal probability. (In other words, each of the other $(m-1)$ bits takes the values 0 and 1 with equal frequency).

$$p(I, Q | b_i = 1) = \frac{1}{2^m \pi \sigma^2} \sum_{x \in C_i^1} e^{-\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2}}$$

The conditional pdf for receiving I, Q , given that b_i was transmitted as a 0 is the same except that the summation is taken over points $x \in C_i^0$. Now if we invoke Bayes' theorem, and make the further assumption that the transmitted bit b_i is itself equally likely to be 0 or 1, then we can obtain the LLR as required:

$$\begin{aligned} \text{LLR}(b_i) &= \ln \left(\frac{\Pr(b_i = 1 | I, Q)}{\Pr(b_i = 0 | I, Q)} \right) = \ln \left(\frac{p(I, Q | b_i = 1)}{p(I, Q | b_i = 0)} \right) \\ &= \ln \left(\frac{\sum_{x \in C_i^1} e^{-\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2}}}{\sum_{x \in C_i^0} e^{-\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2}}} \right) \end{aligned}$$

10.5.3.1.2 Computation of simplified LLR (suitable for hardware implementation)

The LLR computation can be simplified by applying the Max-Log approximation:

$$\ln(e^{a_1} + \dots + e^{a_k}) \approx \max_{i=1 \dots k} (a_i)$$

The LLR becomes:

$$\text{LLR}(b_i) \approx \frac{1}{2\sigma^2} \left[\min_{x \in C_i^0} \left((I - \rho_I I_x)^2 + (Q - \rho_Q Q_x)^2 \right) - \min_{x \in C_i^1} \left((I - \rho_I I_x)^2 + (Q - \rho_Q Q_x)^2 \right) \right]$$

10.5.3.2 2D LLR demapper with iterative demapping and decoding

When Iterative Demapping (ID) is applied, the demapper has to be slightly modified in order to take extrinsic bit information coming from the LDPC decoder into account as shown in figure 97.




Figure 97: LLR demapper for iterative demapping

10.5.3.2.1 Computation of perfect LLR with ID

With iterative demapping, the metric for bit b_i should be calculated in the light of *a priori* knowledge of the likely state of the $m - 1$ other bits, obtained from the LDPC decoder in the previous iteration. Because of this knowledge, we should no longer assume that all states x of the constellation are equiprobable. To reflect this, we need a more complicated expression for the conditional pdf of the received values I, Q , given that b_i was transmitted as a 1:

$$p(I, Q | b_i = 1) = \frac{1}{2^m \pi \sigma^2} \sum_{x \in C_i^1} \left(e^{-\frac{(I - \rho_I I_x)^2 + (Q - \rho_Q Q_x)^2}{2\sigma^2}} \Pr_{\text{apriori}}(x | b_i = 1) \right)$$

This expression sums up the contributions from each of the 2^{m-1} possible transmitted points x in the half-constellation C_i^1 that is distinguished by our choosing $b_i = 1$. Each point has its own probability of having being transmitted. This probability can be expressed as a function of the probabilities that the $(m - 1)$ bits other than b_i take the value 0 or 1, as needed to match the mapping of state x . These probabilities are in effect estimated by the previous iteration of LDPC decoding. Let $x_{m-1}x_{m-2}\dots x_2x_1x_0$ represent the values of the bits $b_{m-1}b_{m-2}\dots b_2b_1b_0$ that correspond to the point x in the half-constellation C_i^1 . Then:

$$\Pr_{\text{apriori}}(x | b_i = 1) = \prod_{k \neq i} \Pr_{\text{LDPC}}(b_k = x_k)$$

(As before, the obvious changes are made to obtain the corresponding pdf for the case that b_i was transmitted as a 0).

We assume that the LDPC decoder handles soft decisions (LLRs) at its inputs and outputs, rather than bit probabilities. The extrinsic bit LLR related to bit b , $\text{LLR}_{\text{EXT}}(b)$, is related to the LDPC decoder's input and output soft-decision values $\text{LLR}_{\text{IN}}(b)$ and $\text{LLR}_{\text{OUT}}(b)$ by:

$$\text{LLR}_{\text{EXT}}(b) = \text{LLR}_{\text{OUT}}(b) - \text{LLR}_{\text{IN}}(b)$$

We obtain the extrinsic bit probabilities from the extrinsic bit LLR by the following transformation (in effect the reverse of the LLR definition):

$$\Pr_{\text{LDPC}}(b = 0) = \frac{1}{1 + e^{\text{LLR}_{\text{EXT}}(b)}}, \text{ and}$$

$$\Pr_{\text{LDPC}}(b = 1) = 1 - \Pr_{\text{LDPC}}(b = 0)$$

Finally, we can collect the previous equations, applying Bayes' theorem as before (still under the assumption that 0 and 1 are equiprobable for b_i), to obtain the LLR for bit b_i at the output of the demapper, taking into account the iterative process:

$$\text{LLR}(b_i) = \ln \left\{ \frac{\sum_{x \in C_i^1} \left(e^{-\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2}} \prod_{k \neq i} \text{Pr}_{\text{LDPC}}(b_k = x_k) \right)}{\sum_{x \in C_i^0} \left(e^{-\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2}} \prod_{k \neq i} \text{Pr}_{\text{LDPC}}(b_k = x_k) \right)} \right\}$$

10.5.3.2.2 'Genie-aided' demapping

'Genie-aided' demapping is the assumed limiting case of the benefits of iterative decoding and demapping. In this case we assume that the LDPC decoding of all the bits except b_i is error-free and known to be so. It follows that only two constellation points could possibly have been transmitted, which two points of course depending on the values of the bits except b_i . One of these two points corresponds to $b_i = 1$, and the other to $b_i = 0$. See figure 98 for some selected examples.




Figure 98: examples of the two possible rotated-64-QAM constellation states, when bit 1 is unknown, and the other 5 bits take 4 out of the possible 32 example combinations

So the numerator in the LLR expression above will contain only one state x (corresponding to the point for which $b_i = 1$) for which the probability product is non-zero and identical to unity. This is because the product is of terms $\text{Pr}_{\text{LDPC}}(b_k = x_k)$ that each (since the decoding is known to be error-free) take a value of either 0 or 1, and thus the product is non-zero, and equal to 1, only in the case of the state x that corresponds to the $b_i = 1$, together with the known values of the other bits. Let the transmitted coordinates of this state be I_1, Q_1 .

Similar reasoning also applies to the denominator, with a different single state corresponding to the point for which $b_i = 0$. Let the transmitted coordinates of this state be I_0, Q_0 . In effect, the constellation has been reduced to a special case of BPSK.

It follows that the LLR becomes a simple function of I, Q :

$$\text{LLR}(b_i) = \ln \left\{ \frac{e^{-\frac{(I-\rho_I I_1)^2 + (Q-\rho_Q Q_1)^2}{2\sigma^2}}}{e^{-\frac{(I-\rho_I I_0)^2 + (Q-\rho_Q Q_0)^2}{2\sigma^2}}} \right\}$$

This function takes the form of an inclined plane that intersects the I, Q plane in a straight line. This straight line is itself the perpendicular bisector of the line joining the two (scaled) points $(\rho_I I_1, \rho_Q Q_1)$ and $(\rho_I I_0, \rho_Q Q_0)$. See figure 99.

LLR and mappings for bit 2, given $\{b_0, b_1, b_3, b_4, b_5\} = \{0, 1, 1, 0, 1\}$
 (blue represents 0 and red represents 1)




Figure 99: an example of the 'Genie-aided' LLR calculation for rotated-64-QAM.
In this example the LLR is for bit 2, while the other bits are 'known' to take the values indicated

10.5.3.2.3 Computation of simplified LLR with ID (suitable for hardware implementation)

While we trust that iterative demapping will eventually converge on the 'Genie-aided' scenario (where the state of the other bits $b_k|_{k \neq i}$ is deemed to be known with absolute certainty), during the iterations we do need to take into account the relative certainty about these bits expressed by the probabilities $\Pr_{\text{LDPC}}(b_k = x_k)$, which form part of the LLR equation at the end of clause 10.5.3.2.1.

Once again, we would like to apply the Max-Log approximation in order to simplify this LLR computation. However, in order to do so, some additional mathematical simplifications are first necessary since all the terms in the summations within the LLR equation are exponentials.

Consider one such term $\left(e^{-\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2}} \prod_{k \neq i} \Pr_{\text{LDPC}}(b_k = x_k) \right)$. If we divide every such term, in both the

numerator and the denominator of the full expression by the same quantity, the LLR calculation will be unchanged. Suppose we divide it by $\prod_{k \neq i} \Pr_{\text{LDPC}}(b_k = 0)$, i.e. the probability of all the other bits taking the value 0. The term thus

becomes $\left(e^{-\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2}} \prod_{k \neq i} \frac{\Pr_{\text{LDPC}}(b_k = x_k)}{\Pr_{\text{LDPC}}(b_k = 0)} \right)$.

Now consider the product term $\prod_{k \neq i} \frac{\Pr_{\text{LDPC}}(b_k = x_k)}{\Pr_{\text{LDPC}}(b_k = 0)}$, a product of probability ratios for each k other than i . For those bits b_k whose value $x_k = 0$, the probabilities cancel, so that the only terms remaining in the product are those for which

$x_k = 1$. These surviving product terms take the form $\frac{\Pr_{\text{LDPC}}(b_k = 1)}{\Pr_{\text{LDPC}}(b_k = 0)} = e^{\text{LLR}_{\text{EXT}}(b_k)}$.

The LLR equation at the end of clause 10.5.3.2.1 can thus be re-written as:

$$\text{LLR}(b_i) = \ln \left\{ \frac{\sum_{x \in C_i^1} \left(e^{\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2}} \prod_{k \neq i} \text{Pr}_{\text{LDPC}}(b_k = x_k) \right)}{\sum_{x \in C_i^0} \left(e^{\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2}} \prod_{k \neq i} \text{Pr}_{\text{LDPC}}(b_k = x_k) \right)} \right\} = \ln \left\{ \frac{\sum_{x \in C_i^1} \left(e^{\left(\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2} + \sum_{k \neq i \text{ and } x_k=1} \text{LLR}_{\text{EXT}}(b_k) \right)} \right)}{\sum_{x \in C_i^0} \left(e^{\left(\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2} + \sum_{k \neq i \text{ and } x_k=1} \text{LLR}_{\text{EXT}}(b_k) \right)} \right)} \right\}$$

Note that the summation over k spans all bits b_k that are equal to one for the constellation state x (i.e. $x_k = 1$), excluding bit b_i .

This is now in a form where we can apply the Max-Log approximation, to obtain:

$$\text{LLR}(b_i) \approx \min_{x \in C_i^0} \left(\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2} - \sum_{k \neq i \text{ and } x_k=1} \text{LLR}_{\text{EXT}}(b_k) \right) - \min_{x \in C_i^1} \left(\frac{(I-\rho_I I_x)^2 + (Q-\rho_Q Q_x)^2}{2\sigma^2} - \sum_{k \neq i \text{ and } x_k=1} \text{LLR}_{\text{EXT}}(b_k) \right)$$

10.5.3.3 Impact on complexity

The main added complexity for the receiver, when rotated constellations are used, arises in the calculation of metrics, i.e. the LLRs for each coded bit carried by each OFDM constellation. Whether rotated or not, 2^m -QAM conveys m bits per cell, and thus requires the calculation of m LLRs per cell. The difference arises in the complexity of those calculations.

In the case of classical, non-rotated 2^m -QAM, each axis can be treated quite independently. Each axis conveys $m/2$ bits and thus requires calculation of $m/2$ LLRs. Each of these, if calculated fully, without approximation, requires the calculation of $2^{m/2}$ 1-D distances (which are then used in the exponents of the same number of exponential terms). So per cell, there are $2^{(1+m/2)}$ such terms to be calculated in total. However, it is common to use an approximation (equivalent to the max-log approximation) which gives a simple piecewise-linear form for the LLR as a function of position on the axis. So no distances need be explicitly calculated at all.

In the case of rotated 2^m -QAM, the axes can no longer be treated independently; each LLR is a function of both I and Q , plus the fading factors of the two axes. To perform the full calculation now requires the calculation of the 2-D distances from the received point to all of the 2^m constellation points. In effect this amounts to 2^{m+1} 1-D distances. So the number of 1-D distances that needs to be calculated has been increased by the factor $2^{m/2}$, compared with the non-rotated case. This same number of distances appears (in the absence of further simplifications) to need to be calculated even when the max-log approximation is applied, but at least the corresponding exponential terms are avoided.

In the case of the demapper with iterative demapping, some adders have to be introduced after the computation of metrics and before the comparators, in order to take extrinsic LLR values coming from the LDPC decoder into account.

10.5.4 PAPR-aware demapping

The ACE PAPR technique (see clause 9.3.8.2) makes modifications to the outer transmitted-constellation points in order to reduce the peak-to-mean ratio of the transmitted signal. This could potentially affect the LLR demapping process in the receiver, since the assumption built into the traditional LLR is that the transmitted points are all on the QAM grid and any departure of the received point from these points is the result of additive noise. Constellation-based measures such as the average distance from the nearest constellation point are also used to estimate the level of noise and interference affecting each carrier. Clearly, this could also be misled by the use of ACE.

Receiver implementers should therefore be careful to take into account the use of ACE in developing these algorithms. The receiver can tell whether ACE is in use by examining the L1 signalling parameter 'PAPR'.

10.5.5 LDPC decoding

The most general decoding algorithm of LDPC code is an iterative message passing algorithm called belief propagation [i.12]. This algorithm is naturally suitable for parallel message computation. The structure of the parity check matrices of the LDPC codes adopted in DVB-T2 implies a partly-parallel decoder architecture.

The performance of LDPC codes can be improved by increasing the number of decoding iterations. Figure 100 shows the performance sensitivity, i.e. the variation of performance with the number of iterations for rate-2/3 $N_{ldpc}=64\ 800$ code with a 256-QAM constellation on the AWGN channel.




Figure 100: Performance sensitivity for rate 2/3 $N_{ldpc} = 64\ 800$ code at 256-QAM constellation on AWGN channel

The number of decoding iterations per FECFRAME achievable at the receiver for a given hardware complexity can be estimated from the following equation:

$$I = (N_{ldpc} \cdot F_{clk} \cdot P_{dec}) / (C_{code} \cdot E_{mat} \cdot a_{dec}),$$

where:

N_{ldpc} : code length of LDPC code (16 400 or 16 200);

E_{mat} : number of '1's in parity check matrix;

C_{code} : code bit rate (bits/sec);

F_{clk} : decoder clock frequency (Hz);

P_{dec} : parallel order of decoder (typically 360);

a_{dec} : efficiency factor of decoder ($a_{dec}=2$ for variable- and check-node calculation with no overhead).

EXAMPLE: If $N_{ldpc}/E_{mat} = 0.3$ on average, $C_{code} = 64$ Mbps, $P_{dec} = 360$, $a_{dec} = 4$, and $F_{clk} = 120$ MHz, then $I = 50$ iterations/frame.

The maximum rate at which cells are delivered by the T2 system will depend on the channel bandwidth. However, the receiver buffer model defined in annex C of [i.1] assumes that cells can be read out of the time de-interleaver at up to $7,6 \times 10^6$ cells/s. The LDPC decoding hardware should therefore be capable of operating at this cell rate even in lower channel bandwidths.

10.5.6 BCH decoding

For the purpose of removing any possible error floor, use of bounded-distance decoding [i.11] is sufficient.

10.5.7 Decoding of the L1 signalling

The extraction of the L1 signalling cells was described in clause 10.4.2.2. Subsequent processing stages are similar to the corresponding stages in the PLP decoding chains, with the following differences:

- No Time or Cell de-interleaving is applied to the L1 signalling, but the zig-zag mapping constitutes a form of interleaving.
- The L1-pre signalling uses BPSK.
- The bit de-interleaver for the L1-post signalling does not include column-twisting and there is no parity de-interleaver.
- The FEC decoding is more complicated because of the shortening and puncturing as will be described in clause 10.5.7.1.

10.5.7.1 FEC decoding of the L1 signalling

This clause gives an overview of the procedure of decoding of L1 dynamic signalling. Figure 101 illustrates the L1 signalling structure, which is split into three main sections: the P1 signalling, L1-pre-, and L1-post signalling.

Then the L1-pre signalling should be decoded by using the BCH code with $K_{bch} = 3\,072$ and 16 K LDPC code with $K_{ldpc} = 3\,240$ (the coding is described in clause 9.2.2). BCH decoding is performed after LDPC decoding and removal of the 168 BCH parity bits from the decoded LDPC information bits. Note that in LDPC and BCH decoding, the shortened bits in BCH (or LDPC) information field should be regarded as 0. Furthermore, the punctured LDPC parity bits should be regarded as erasures in LDPC decoding.




Figure 101: The L1-signalling structure

For decoding of the L1-post signalling bits, several signalling fields in the L1-pre signalling are needed. The LDPC code rate, modulation order and FEC type, i.e. the length of LDPC code, should be obtained from the fields L1_COD, L1_MOD and L1_FEC_TYPE. Note that the L1-post signalling should be decoded by using BCH code with $K_{\text{bch}} = 7\ 032$ and 16 K LDPC code with $K_{\text{ldpc}} = 7\ 200$. Additionally, the size of the information part of the L1-post signalling data block excluding the CRC, $K_{\text{post_ex_pad}}$ can be obtained from L1_POST_INFO_SIZE since $K_{\text{post_ex_pad}} = \text{L1_POST_INFO_SIZE} + 32$ CRC bits. Note that if $K_{\text{post_ex_pad}}$ is larger than $K_{\text{bch}} = 7\ 032$, the L1-post signalling data block will be segmented into multiple blocks as described in clause 7.3.1 of [i.1]. The segmented blocks of L1 signalling bits are independently encoded by BCH and LDPC codes sequentially. Furthermore, if the segmented block is less than $K_{\text{bch}} = 7\ 032$, the shortening and puncturing is applied to encoding of the segmented block. The numbers and positions of the shortened bits of BCH (or LDPC) information and punctured bits of LDPC parity bits are respectively determined by $K_{\text{post_ex_pad}}$, as described in clause 9.2.2. Therefore, if the L1-pre signalling has been correctly decoded, the shortened and punctured bit-positions can be calculated from L1_POST_INFO_SIZE in the receiver. In decoding procedure, the value of the shortened and punctured bits should be regarded as 0 and erasures, respectively.

Since the coding and modulation parameters for L1-post signalling and the numbers of shortening and puncturing bits in BCH and LDPC encoding are already known from the L1-pre signalling, the code length of each segmented L1-post signalling can be calculated.

Finally, the original L1-post signalling data is reconstructed by composition of the decoded segmented blocks after their BCH and LDPC decoding. Note that the code length of each segmented L1-post signalling can be calculated since it is equal to L1_POST_SIZE divided by the number of segmented blocks.

10.5.7.2 Extraction of the L1-signalling fields

The L1-post signalling contains the configurable, dynamic and (optional) extension part as illustrated in figure 101. The configurable part of L1-post signalling has a variable length which depends not only on the numbers of RFs, PLPs, and auxiliary streams but also on the use of FEF. That is, the length of configurable part is determined by those parameters. The configurable part is followed by the dynamic part of the L1-post signalling. Therefore, the start position of the L1 dynamic part is also determined by the length of the configurable part. The length of dynamic part is determined by the numbers of PLPs and auxiliary streams. Therefore, the end position of the dynamic part of the L1-post signalling is determined since the start position is already known.

If the indicator `L1_REPETITION_FLAG` of 1-bit is set to value '1', the dynamic signalling for the next frame will also be provided within the current frame. In this case, it will follow immediately after the dynamic signalling of the current frame as illustrated in figure 102. This dynamic signalling for the next frame may be jointly used with the identical information, received in the next T2-frame, in order to increase the detection reliability.




Figure 102: Repetition of L1-post dynamic information

10.5.7.3 Storage required in L1 decoding

If the decoding of the Layer-1 signalling information carried in the P2 symbol(s) is to take more than one symbol time, then the data symbols received during the decoding time need to be stored until the Layer-1 signalling is decoded to extract the Layer-1 post-signalling parameters. The amount of storage required depends on the latency of P2-symbol demodulation and LDPC decoding of the signalling information.

In addition, in modes in which there is more than one P2 symbol, storage must be provided for the short interleaving between P2 symbols. The amount of storage required for this is at most one 32 K symbol.

10.6 Output processing

10.6.1 De-/Re-multiplexing

10.6.1.1 Construction of output TS

Construction of the output TS will make use of the following pieces of information:

- The `SYNCD`, `SYNC` and `UPL` fields of the BB-Header: Used to regenerate the user packets (where applicable). See clause 10.6.1.2.
- The `TTO` variable of `ISSY`: used to set the initial occupancy of the `DJB` (see clause 10.6.1.4).
- The `DNP` bytes in the `BBFRAME`: used to re-insert the null packets that were deleted in the modulator or T2-gateway (see clause 10.6.1.5).
- The `ISCR` variable of `ISSY`: used to calculate the output bit-rate and for fine adjustment of the relative timing of data and common PLPs (see clause 10.6.1.6).

10.6.1.2 Mode adaptation

Given the decoded baseband frames from the BCH decoder (clause 10.5.6), the receiver should partially reverse the process of mode adaptation performed in the T2-Gateway (clause 8.5) and so recover the input packets.

The receiver should first determine whether Normal or High Efficiency Mode is being used. If the `PLP_MODE` bits of the L1-configurable signalling are set to '00', then this is indicated only by performing an exclusive-OR of the CRC-8 field with 0 or 1. From a single `BBFRAME` the receiver cannot tell the difference between a correctly received BB-Header using HEM and an errored BB-Header using NM (and vice versa). Instead, the receiver should implement a confidence-count mechanism over a number of `BBFRAME`s. If the received CRC consistently differs from the calculated value only in the LSB, then HEM is in use.

If the `PLP_MODE` bits of the L1-configurable signalling are set to '01' or '10', then the mode (NM or HEM) is indicated directly and does not need to be deduced from the CRC. See clause 8.10.5.

The Data Field should then be extracted from the BBFRAME; it begins with the first bit after the BBFRAME header and its length is given by the DFL field in the header. The data field should then be processed to recover the original packets.

NOTE: The data field need not be a whole number of bytes, since although K_{bch} is always a multiple of 8, the in-band signalling and any padding introduced for any other reason might not contain a whole number of bytes. Receivers should therefore be designed to work with bytes which are not aligned to the data field.

If the NO_PADDING_FLAG in the L1-configurable signalling is set to '1', then the receiver should assume that DFL will always equal $K_{\text{bch}}-80$ in all BBFRAMEs except the first BBFRAME of each Interleaving Frame. This should be used instead of the signalled DFL value, to improve resilience to bit errors in the BB-Header. In this configuration, the receiver should also store the first DFL of the Interleaving Frame, since the first BBFRAME of subsequent Interleaving Frames will have the same DFL value.

The mode adaptation process should only be partially reversed at this stage because some information needs to be retained:

- The TTO and ISCR values should be kept until the transport stream has been regenerated.
- The DNP fields should be kept until the deleted null packets have been reinserted, which is done at the output of the de-jitter buffer (see clause 10.6.1.5).

The receiver buffer model (clause 8.8.1) assumes that the receiver will remove the BB-Headers and reconstitute the stream in a canonical form for processing by the later stages, i.e. a form which does not depend on the particular mode adaptation options in use. The canonical form is described in annex C of [i.1]; in summary:

- The sync bytes are not reinserted until the output of the DJB. The DJB should therefore keep track of the position within a packet.
- The CRC-8, if present, is removed.
- The DNP fields are retained, or inserted with a value of zero if DNP is not used.
- ISCR and TTO values stored in long (24-bit) format. For packets with no associated field, or if ISSY is not used, these could either "freewheel" (see clause 10.6.1.4) or be set to a "unknown" value. If short ISSY is used, the MSBs of ISCR could be re-generated using a freewheel mechanism.

In principle the mode adaptation in use could change from one BBFRAME to the next, although this is unlikely to happen in practice. Reconstituting the stream in a canonical form provides a simple way to deal with such changes, since each BBFRAME can be processed individually and later processing stages need not be aware of the mode adaptation options that were used.

The only exception is the use of NM or HEM, which should be assumed not to change from BBFRAME to BBFRAME, in order to allow the confidence-count mechanism described above to work correctly.

10.6.1.3 Determination of output-TS bit-rate

The receiver needs to know the exact bit-rate of the output transport stream, in order to be able to output the stream.

If ISSY is used, this can be calculated from the ISCR values. Following null-packet re-insertion, the bit stream will consist of a sequence of TS packets, some of which will have associated ISCR values.

The ISCR values are in units of the elementary time period T , which will have one of a limited number of values and will be known by the receiver based on the frequency band and channel bandwidth in use. The difference in time between the beginning of the two packets is therefore known.

The number of packets, and therefore bits, between these two packets will also be known. The TS bit-rate R_{TS} can therefore be calculated simply by dividing the number of bits N_{bits} by the time:

$$R_{\text{TS}} = \frac{N_{\text{bits}}}{(\text{ISCR}_2 - \text{ISCR}_1) \times T}$$

The time between two packets might not be a whole number of cycles of T , and consequently there will be some rounding error in this calculation. The residual error can be removed using a feedback loop of the kind described in annex I of [i.1].

If ISSY is not used, the receiver should use a self-balancing buffer as described in annex I of [i.1]. An initial estimate for the TS bit-rate can be made by calculating the maximum useful bit-rate based on `PLP_NUM_BLOCKS_MAX` as described in clause 6.3.2. However, if null-packet deletion is used, the TS bit-rate may be significantly higher and the feedback loop would need to be able to adjust for this.

Where variable bit-rate PLPs with null-packet deletion are used, the buffer occupancy will be constantly changing, making it difficult to use the occupancy to drive a feedback loop reliably. Receiver implementers should assume that ISSY will be used in such cases and need not implement sophisticated strategies to deal with the possibility that it is not.

If in-band type B signalling is used, the transport stream rate can be read directly from the `TS_RATE` field. See clause 8.10.6.5. This could enable faster initial acquisition and should be supported by new receiver designs.

10.6.1.4 De-jitter buffer

The de-jitter buffer itself can be implemented using a FIFO more or less as described in annex C of [i.1].

The data should be read out using a clock derived from the recovered modulator sampling clock taking into account the calculation of output bit rate described in clause 10.6.1.3. The TTO value should be used to determine the approximate time, after the fixed, implementation-dependent processing delays have been added.

When a common PLP is used, the common PLP has its own de-jitter buffer and time de-interleaver. Both de-jitter buffers should be read out using the same clock. Provided the TTO values are observed both for the data and common PLP, the two partial streams, from the data and common PLP, will be almost aligned. However, the TTO value itself is of limited resolution: the exponent `TTO_E` will typically be set to 15, to give a range of around 455 ms at a resolution of $128T=14\mu\text{s}$, corresponding to several hundred output bits. The alignment should be refined using the `ISCR` values associated with the packets, which have a resolution of $T\approx 109\text{ns}$. Finally the known position of the packet starts should be used to align the output packets at the bit level. These might require a small amount of additional buffering, or could be done in the de-jitter buffer itself.

Both the time de-interleaver memory and the DJB are shared between the data and common PLPs, in a statically configurable manner. The amount of memory to allocate to each PLP in the TDI memory can be deduced from the value of `PLP_NUM_BLOCKS_MAX`. The division of memory in the DJB is indicated by the `BUFS` variable of ISSY.

In principle, TTO only needs to be read once, on initial acquisition. However, receivers should continually track the value to ensure that they are operating correctly.

The TTO and `ISCR` values are carried in the BBHeader; they are therefore prone to bit errors. The CRC allows bit errors to be detected, and when this happens the receiver should provide a "flywheel" mechanism to generate the missing values.

10.6.1.5 Re-insertion of deleted null packets

Conceptually, the deleted null packets are re-inserted before the de-jitter buffer as illustrated in annex I of [i.1]. The de-jitter buffer would then consist of a simple FIFO as shown. However, the de-jitter buffer memory will generally not be large enough to accommodate all of the null packets, since the scheduler is allowed to assume that they occupy no space (see clause C.1.1 of [i.1]).

In practice, therefore, the deleted null packets should be re-inserted at the output of the de-jitter buffer. This can be done according to the `DNP` fields in the BBFRAME. On each TS clock pulse, data should either be read from the de-jitter buffer or regenerated by the null packet re-insertion process. This is the point in the chain at which the variable bit-rate stream carried in a PLP becomes a constant bit-rate TS, and the re-insertion of null packets is the mechanism which achieves this.

This implies that information about the deleted null packets should be stored either in the de-jitter buffer itself or in separate storage. If the data in the DJB is stored in the canonical format assumed by clause C.1.1 of [i.1], as discussed in clause 10.6.1.2, then the DNPs will be included as part of this format. Note that in this case the `DNP` counts will need to be acted upon before the preceding packet (in transmission order) is output and therefore some re-ordering at the DJB input or non-sequential read-access to the DJB memory may be required.

10.6.1.6 Re-combining the Common and Data PLPs

When a particular input TS has been split and is sent in a data PLP and a common PLP (see clause 8.3) the original TS can be recreated in the receiver by combining the contents of the data PLP and the common PLP.

The operations required to re-combine the common and data PLPs, together with recommendations for receiver implementations including modification of the PSI/SI, are described in annex D of [i.1].

Essentially, where the data PLP has a non-null packet, this should be output; only when the data PLP has a null packet should the packet in the common PLP be output. For this purpose, both deleted and transmitted null packets should be treated identically.

As explained in clause 8.3, the EIT is carried as 'EIT other' in the common PLP, even when it describes one of the TSs in the group, so sometimes the TS packet being output by the receiver will contain 'EIT other' describing the very same TS that is being output. In this case the `table_id` and `last_table_id` may be modified from 'other' to 'actual' to achieve full transparency of the TS. If these ids are modified, the CRC should also be recalculated accordingly. Alternatively, the receiver may leave the `table_id` and `last_table_id` as 'other', provided that this is acceptable to subsequent processing stages, to avoid the need to modify the table and recalculate the CRC. When EIT schedule is scrambled, modifying the table may be difficult and this might be another reason to leave the `table_id/last_table_id` as 'other'.

10.6.2 Output interface

For TS input, the output from a demodulator should be the relevant complete Transport Stream.

For generic streams, the output format will depend on the type of stream.

11 Transmitters

Many considerations for the implementation and use of transmitters for DVB-T2 will be similar to those for DVB-T, since the occupied bandwidth and signal spectra are broadly similar. One area in which DVB-T2 is significantly different is in the availability of two techniques for PAPR reduction (see clause 9.3.8).

Future versions of the present document will give more information on transmitter aspects, including PAPR reduction, as more experience is gained.

12 Network planning

Network planning for DVB-T2 will have much in common with planning for DVB-T, but the availability of a wider range of technologies and options in DVB-T2 gives more flexibility. In particular, the larger FFT sizes and wider choice of guard intervals will give greater scope for optimisation of networks. The use of MISO (clause 9.3.3) could also give a worthwhile benefit in some SFN configurations.

The choice of the various parameters was discussed in clause 5 and network operators should refer to that clause for guidance. Future versions of the present document will give more information on Network Planning aspects, as more experience is gained.

13 Tuners

13.1 Phase noise requirements

Generally phase-noise considerations are the same as for DVB-T.

Phase noise added to an OFDM signal causes two distinct effects: CPE and ICI. CPE is a common rotation of all the constellations transmitted in one OFDM symbol. Because it is common for all carriers, it is possible to measure it and cancel it out. ICI is a form of crosstalk between carriers and manifests itself as an additional noise term, degrading the constellation SNR.

EN 302 755 [i.1] includes continual pilots which, amongst other uses, can be used for the purpose of cancelling CPE, according to the method described in clause 10.3.2.1. Taking this and other possible cancellation measures into account, the impact of CPE is expected to be negligible.

Some practical measurements with DVB-T signals have shown that for a variety of phase noise spectra the ICI values do not vary significantly between 2 K and 8 K modes. Measurement results for DVB-T have shown no difference between 2 K or 8 K OFDM concerning the QEF limit for a range of given phase-noise spectra; these measurements also confirm that there is no relevant CPE influence. Further ICI calculations based on [i.10] using phase-noise spectra for particular available tuners have shown that, for these tuners, the ICI in 16 K and 32 K FFT modes will not be significantly worse than in 2 K and 8 K.

In practice, tuners will vary in their phase-noise spectra, and manufacturers should therefore calculate the ICI value due to the phase noise values of their tuners. Calculation may be done according to [i.9]. Because ICI behaves like AWGN noise manufacturers then can estimate the implementation loss of their tuners.

NOTE: The calculation should take into account the accumulated effect of all the oscillators in the chain; in principle this includes the modulator, although tuner manufacturers may assume that the effect of phase noise in the transmitted signal is insignificant.

EXAMPLE: For the transmission parameters selected according to table 1 a minimum C/N of 16,5 dB ([i.18], table A.1) for QEF is required (gaussian channel). Phase-noise characteristics of typical DVB-T tuners lead to ICI values which are at least 10 dB better than this minimum C/N value. Therefore in this case the implementation loss caused by ICI will be less than 0,5 dB).

13.2 AGC, headroom and FEFs

Considerations for the dynamics of the AGC, together with the setting of the backoff to allow headroom for the OFDM signal, are likely to be similar to the DVB-T case.

As was explained in clause 10.4.2.6, the AGC of a receiver should be locked during any FEF part. If the tuner has its own analogue AGC loop, this should also be locked during the FEF. This will require accurate timing control from the demodulator.

14 Performance

This clause presents simulation results indicating the performance of DVB-T2 in a range of terrestrial channels. The results are based on the simulation and validation effort carried out during and following the development of the DVB-T2 standard [i.1].

The following information will be presented:

- Description of the channel models used for simulations.
- Performance simulation results for the whole chain, showing the trade-off between performance and data capacity.
- Performance simulation results in the memoryless Rayleigh erasure channel: this channel only tests the BICM part of the chain.
- A discussion of the effect that real, as opposed to ideal, channel estimation would have on the results.
- Performance simulation results for the P1 detection and decoding.
- Performance simulation results for L1-pre and post-signalling.

14.1 Channel models for simulations

14.1.1 Introduction

During the development of the DVB-T2 specification [1.1], different channel models have been used to provide simulated performance results. In addition to the well-known channels from DVB-T, some other channels have been added, some specific to DVB-T2. The channels have been chosen to verify the performance in a wide range of reception conditions including fixed, portable, mobile and even SFN and MISO.

14.1.2 Gaussian Channel

In this channel model only white Gaussian noise (AWGN) is added to the signal, and there is only one path.

14.1.3 Ricean Channel

The same F_1 Ricean fading channel defined in the DVB-T specification is used to describe the fixed, outdoor rooftop-antenna reception conditions. The channel does not include any Doppler and should therefore be considered as a snapshot of the real time-variant channel. The model has 21-taps.

The model is defined by:

$$y(t) = \frac{\rho_0 x(t) + \sum_{i=1}^N \rho_i e^{-j\theta_i} x(t - \tau_i)}{\sqrt{\sum_{i=0}^N \rho_i^2}}$$

where:

- $x(t)$ and $y(t)$ are input and output signals respectively;
- the first term before the sum represents the line of sight ray;
- N is the number of echoes and is equal to 20;
- θ_i is the phase shift from scattering of the i th path - listed in table 39;
- ρ_i is the attenuation of the i th path - listed in table 39;
- τ_i is the relative delay of the i th path - listed in table 39.

The Ricean factor K , the ratio of the power of the direct path (the line-of-sight ray) to the reflected paths, is given as:

$$K = \frac{\rho_0^2}{\sum_{i=1}^N \rho_i^2}$$

In the simulations a Ricean factor $K = 10$ dB has been used. In this case:

$$\rho_o = \sqrt{10 \sum_{i=1}^N \rho_i^2}$$

14.1.4 Rayleigh Channel

The same P_1 Rayleigh-fading channel defined in the DVB-T specification is used to describe the portable indoor or outdoor reception conditions. The channel does not include any Doppler and should therefore be considered as a snapshot of the real time-variant Rayleigh channel. The model has 20 taps.

The model is defined by:

$$y(t) = k \sum_{i=1}^N \rho_i e^{-j\theta_i} x(t - \tau_i)$$

where:

- $x(t)$ and $y(t)$ are input and output signals respectively;
- $k = \frac{1}{\sqrt{\sum_{i=1}^N \rho_i^2}}$
- θ_i , ρ_i and τ_i are given in table 39.

Table 39: Relative power, phase and delay values for F_1 and P_1

| i | ρ_i | τ_i (μs) | θ_i (rad) |
|----|-----------|----------------------------|------------------|
| 1 | 0,057 662 | 1,003 019 | 4,855 121 |
| 2 | 0,176 809 | 5,422 091 | 3,419 109 |
| 3 | 0,407 163 | 0,518 650 | 5,864 470 |
| 4 | 0,303 585 | 2,751 772 | 2,215 894 |
| 5 | 0,258 782 | 0,602 895 | 3,758 058 |
| 6 | 0,061 831 | 1,016 585 | 5,430 202 |
| 7 | 0,150 340 | 0,143 556 | 3,952 093 |
| 8 | 0,051 534 | 0,153 832 | 1,093 586 |
| 9 | 0,185 074 | 3,324 866 | 5,775 198 |
| 10 | 0,400 967 | 1,935 570 | 0,154 459 |
| 11 | 0,295 723 | 0,429 948 | 5,928 383 |
| 12 | 0,350 825 | 3,228 872 | 3,053 023 |
| 13 | 0,262 909 | 0,848 831 | 0,628 578 |
| 14 | 0,225 894 | 0,073 883 | 2,128 544 |
| 15 | 0,170 996 | 0,203 952 | 1,099 463 |
| 16 | 0,149 723 | 0,194 207 | 3,462 951 |
| 17 | 0,240 140 | 0,924 450 | 3,664 773 |
| 18 | 0,116 587 | 1,381 320 | 2,833 799 |
| 19 | 0,221 155 | 0,640 512 | 3,334 290 |
| 20 | 0,259 730 | 1,368 671 | 0,393 889 |

14.1.5 Mobile Channel, TU-6

This profile reproduces the terrestrial propagation in an urban area. It has been defined by COST 207 as a typical urban (TU6) profile and is made of 6 paths having wide dispersion in delay and relatively strong power. The profile parameters are given in table 40.

Table 40: Typical Urban profile (TU6)

| Tap number | Delay (μs) | Power (dB) | Doppler Spectrum |
|------------|-------------------------|------------|------------------|
| 1 | 0,0 | -3 | Classical |
| 2 | 0,2 | 0 | Classical |
| 3 | 0,5 | -2 | Classical |
| 4 | 1,6 | -6 | Classical |
| 5 | 2,3 | -8 | Classical |
| 6 | 5,0 | -10 | Classical |

where the Classical Doppler spectrum is defined as:

$$K(f; f_D) = \frac{1}{\sqrt{1 - (f/f_D)^2}}$$

14.1.6 Simple two path profile, 0 dB Echo

This profile only includes two paths. Each path, 'i', is defined by α_i , τ_i , Δf_i which denote the amplitude, delay and frequency shift of the particular path 'i', respectively. The profile parameters are given in table 41.

Table 41: 0 dB Echo profile

| i | α_i (dB) | τ_i (μs) | Δf_i (Hz) |
|---|-----------------|----------------------------|-------------------|
| 1 | 0 | 0 | 0 |
| 2 | 0 | 0,9 Δ | 1 |

The parameter ' Δ ' denotes the guard interval duration in μs .

14.1.7 MISO channel

As MISO transmission is used from different transmit antennas, the simulation of MISO channels requires at least two independent fading channels, one for each transmitter. For the first transmitter, the values defined in table 39 are used, whereas the values for the second transmitter are defined in table 42.

Table 42: MISO channel: relative power, phase and delay values for second path (F_2 and P_2)

| i | $\rho_{2,i}$ | $\tau_{2,i}$ (μs) | $\theta_{2,i}$ (rad) |
|----|--------------|--------------------------------|----------------------|
| 1 | 0,057 662 | 2,003 019 | 1,855 121 |
| 2 | 0,176 809 | 2,422 091 | 2,419 109 |
| 3 | 0,407 163 | 1,518 650 | 3,864 470 |
| 4 | 0,303 585 | 0,751 772 | 1,215 894 |
| 5 | 0,258 782 | 3,602 895 | 0,758 058 |
| 6 | 0,061 831 | 0,016 585 | 2,430 202 |
| 7 | 0,150 340 | 5,143 556 | 4,952 093 |
| 8 | 0,051 534 | 1,153 832 | 0,093 586 |
| 9 | 0,185 074 | 2,324 866 | 4,775 198 |
| 10 | 0,400 967 | 4,935 570 | 6,154 459 |
| 11 | 0,295 723 | 3,429 948 | 1,928 383 |
| 12 | 0,350 825 | 1,228 872 | 2,053 023 |
| 13 | 0,262 909 | 1,848 831 | 1,628 578 |
| 14 | 0,225 894 | 3,073 883 | 2,128 544 |
| 15 | 0,170 996 | 1,203 952 | 4,099 463 |
| 16 | 0,149 723 | 4,194 207 | 3,462 951 |
| 17 | 0,240 140 | 1,924 450 | 3,664 773 |
| 18 | 0,116 587 | 2,381 320 | 3,833 799 |
| 19 | 0,221 155 | 0,640 512 | 3,334 290 |
| 20 | 0,259 730 | 3,368 671 | 0,393 889 |

14.1.7.1 Ricean MISO Channel

The Ricean MISO fading channel is very similar to the normal SISO F_1 Ricean Channel, except that a second path is added.

The Ricean MISO channel is defined by:

$$y(t) = \frac{\rho_{1,0} x_1(t) + \sum_{i=1}^N \rho_{1,i} e^{-j\theta_{1,i}} x_1(t - \tau_{1,i}) + A \left[\rho_{2,0} x_2(t - \Delta_2) + \sum_{i=1}^N \rho_{2,i} e^{-j\theta_{2,i}} x_2(t - \tau_{2,i} - \Delta_2) \right] \cdot e^{j2\pi f_{\Delta} t}}{\sqrt{\sum_{i=0}^N \rho_{1,i}^2 + A^2 \sum_{i=0}^N \rho_{2,i}^2}}$$

where:

- $x_1(t)$ and $x_2(t)$ are the input signals from the two transmitters and $y(t)$ is the output signal;
- the first term before each sum represents the line of sight ray for each path;
- N is the number of echoes and is equal to 20;
- $\theta_{1,i}$ and $\theta_{2,i}$ are the phase shifts from scattering of the i 'th path - listed in table 39 and table 42 respectively;
- $\rho_{1,i}$ and $\rho_{2,i}$ are the attenuations of the i 'th path - listed in table 39 and table 42 respectively;
- $\tau_{1,i}$ and $\tau_{2,i}$ are the relative delays of the i 'th path - listed in table 39 and table 42 respectively;
- Δ_2 is the delay of the second transmitter with respect to the first;
- f_{Δ} is the frequency offset between the two transmitters;
- A is the attenuation of the second transmitter with respect to the first one.

The Ricean factor K_j for each channel is given as:

$$K_j = \frac{\rho_{j,0}^2}{\sum_{i=1}^N \rho_{j,i}^2}$$

For the simulations, a Ricean factor $K = 10$ dB was used for both channels.

The parameters for the attenuation of the second transmitter and the relative frequency shift depend on the scenario. In case of co-located application of the MISO approach, where both transmitters are located at the same site, both transmitters will have nearly the same strengths at the reception site, arrive at the same time, and as the transmitters may use the same local oscillator, no frequency offset will occur. If distributed MISO is used, there may be a frequency offset between the two transmitters and the two paths will have different attenuation and different arrival times. The values for the three different scenarios are given in table 43.

Table 43: MISO parameters for different scenarios

| Scenario | A(dB) | $\Delta_2(\mu\text{s})$ | f_Δ (Hz) |
|---------------|-------|-------------------------|-------------------------------------|
| Co-located | 0 | 0 | 0 |
| Distributed 1 | 0 | 0,18 Δ | $\frac{1}{550} \cdot \frac{1}{T_U}$ |
| Distributed 2 | -3 | 0,23 Δ | $\frac{1}{550} \cdot \frac{1}{T_U}$ |

The parameter ' Δ ' denotes the guard interval duration in μs and $1/T_U$ is the carrier spacing in Hz.

14.1.7.2 Rayleigh MISO Channel

The Rayleigh MISO channel is very similar to the SISO P_1 Rayleigh channel, except that a second path is added.

The MISO Rayleigh channel is:

$$y(t) = \frac{\sum_{i=1}^N \rho_{1,i} e^{-j\theta_{1,i}} x_1(t - \tau_{1,i}) + A \sum_{i=1}^N \rho_{2,i} e^{-j\theta_{2,i}} x_2(t - \tau_{2,i} - \Delta_2) \cdot e^{j2\pi f_\Delta t}}{\sqrt{\sum_{i=1}^N \rho_{1,i}^2 + A^2 \sum_{i=1}^N \rho_{2,i}^2}}$$

while the same parameters as in the Rice MISO channel are applied.

14.1.8 Memoryless Rayleigh Channel with erasures

This channel model is used to simulate the behaviour of the BICM module of DVB-T2 system over terrestrial multipath channels. As described in figure 103, the principle consists in modelling the concatenation of the frame builder, OFDM generator, wireless channel, OFDM demodulator and frame extraction as an equivalent flat fading channel.




Figure 103: Insertion of memoryless Rayleigh channel with erasures model within T2 system

This channel models an OFDM-based system with perfect frequency interleaving over a multipath channel providing arbitrary probability of carrier erasures. The equivalent flat-fading coefficient h is modelled as a Rayleigh process concatenated with random erasures of probability R_e . Figure 104 depicts how to simulate this equivalent channel.

NOTE: This assumes an input complex signal $x(t)$ with zero mean and unit variance, such that the additive noise variance is simply:

$$\sigma_n^2 = 10^{-SNR/10}$$

where SNR is the signal-to-noise ratio in dB.




Figure 104: Memoryless Rayleigh channel with erasures

It should be noted that performance obtained with $R_e = 0$ provides a lower bound on the performance of the DVB-T2 system over a P1 channel.

14.2 Simulated system performance for 8 MHz channel

Table 44 and table 45 give simulated performance, assuming perfect channel-estimation, perfect synchronization and without phase noise, of channel coding and modulation combinations, and are subject to confirmation by testing.

These results are given for the Gaussian channel, Ricean channel (F1), Rayleigh channel (P1), 0 dB single-echo channel and TU6 channel, when the centre carrier of DVB-T2 signal is positioned at 32/7 MHz. The channel models are described in clause 14.1.

Results are given at a BER of 10^{-7} after LDPC, corresponding to approximately 10^{-11} after BCH.

To ensure reliable results, the simulations were run until the following two conditions had been fulfilled:

- minimum of 100 erroneous FEC blocks; and
- minimum of 1 000 erroneous bits detected.

The DVB-T2 OFDM parameters used for these simulations were chosen to be as similar as possible to those for DVB-T. These parameters are as follows: the FFT size is 8 K with a guard interval of 1/32, and the bandwidth is 8 MHz with normal carrier mode. Rotated constellations were used and PAPR techniques were not applied. The simulations assumed ideal conditions, i.e. ideal synchronisation and ideal channel estimation.

In the simulations, the transmitted signal includes no pilots at all and the special symbols at both start (i.e. P1, P2) and end (i.e. frame closing symbol) of the frame are not included. The values of $(C/N)_0$ should therefore be corrected for the FFT size and pilot pattern in use as described in clause 14.2.1.

NOTE: This correction is separate from the penalty for real channel estimation discussed in clause 14.4 and both should be applied.

The simulations include "Genie-Aided" demapping (see clause 10.5.3.2.2). Iterative demapping will approach this performance at low BERs and low-order constellations but will be optimistic at higher BERs and for high-order constellations. Iterative demapping is computationally intensive and network operators should take into account the state of technological development when determining the appropriate implementation margin to use in receiver specifications

**Table 44: Required raw $(C/N)_0$ to achieve a BER = 1×10^{-7} after LDPC decoding
LDPC block length: 64 800 bits**

| Constellation | Code rate | Spectral Efficiency (see note 2) | Required $(C/N)_0$ (dB) for BER = 1×10^{-7} after LDPC decoding | | | |
|---------------|-----------|----------------------------------|--|--------------------------|----------------------------|-----------------------------|
| | | | Gaussian Channel (AWGN) | Ricean channel (F_1) | Rayleigh channel (P_1) | 0 dB echo channel @ 90 % GI |
| QPSK | 1/2 | <i>0,99</i> | <i>1,0</i> | <i>1,2</i> | <i>2,0</i> | <i>1,7</i> |
| QPSK | 3/5 | <i>1,19</i> | <i>2,3</i> | <i>2,5</i> | <i>3,6</i> | <i>3,2</i> |
| QPSK | 2/3 | <i>1,33</i> | <i>3,1</i> | <i>3,4</i> | <i>4,9</i> | <i>4,5</i> |
| QPSK | 3/4 | <i>1,49</i> | <i>4,1</i> | <i>4,4</i> | <i>6,2</i> | <i>5,7</i> |
| QPSK | 4/5 | <i>1,59</i> | <i>4,7</i> | <i>5,1</i> | <i>7,1</i> | <i>6,6</i> |
| QPSK | 5/6 | <i>1,66</i> | <i>5,2</i> | <i>5,6</i> | <i>7,9</i> | <i>7,5</i> |
| 16-QAM | 1/2 | <i>1,99</i> | <i>6,0</i> | <i>6,2</i> | <i>7,5</i> | <i>7,2</i> |
| 16-QAM | 3/5 | <i>2,39</i> | <i>7,6</i> | <i>7,8</i> | <i>9,3</i> | <i>9,0</i> |
| 16-QAM | 2/3 | <i>2,66</i> | <i>8,9</i> | <i>9,1</i> | <i>10,8</i> | <i>10,4</i> |
| 16-QAM | 3/4 | <i>2,99</i> | <i>10,0</i> | <i>10,4</i> | <i>12,4</i> | <i>12,1</i> |
| 16-QAM | 4/5 | <i>3,19</i> | <i>10,8</i> | <i>11,2</i> | <i>13,6</i> | <i>13,4</i> |
| 16-QAM | 5/6 | <i>3,32</i> | <i>11,4</i> | <i>11,8</i> | <i>14,5</i> | <i>14,4</i> |
| 64-QAM | 1/2 | <i>2,98</i> | <i>9,9</i> | <i>10,2</i> | <i>11,9</i> | <i>11,8</i> |
| 64-QAM | 3/5 | <i>3,58</i> | <i>12,0</i> | <i>12,3</i> | <i>14,0</i> | <i>13,9</i> |
| 64-QAM | 2/3 | <i>3,99</i> | <i>13,5</i> | <i>13,8</i> | <i>15,6</i> | <i>15,5</i> |
| 64-QAM | 3/4 | <i>4,48</i> | <i>15,1</i> | <i>15,4</i> | <i>17,7</i> | <i>17,6</i> |
| 64-QAM | 4/5 | <i>4,78</i> | <i>16,1</i> | <i>16,6</i> | <i>19,2</i> | <i>19,2</i> |
| 64-QAM | 5/6 | <i>4,99</i> | <i>16,8</i> | <i>17,2</i> | <i>20,2</i> | <i>20,4</i> |
| 256-QAM | 1/2 | <i>3,98</i> | <i>13,2</i> | <i>13,6</i> | <i>15,6</i> | <i>15,7</i> |
| 256-QAM | 3/5 | <i>4,78</i> | <i>16,1</i> | <i>16,3</i> | <i>18,3</i> | <i>18,4</i> |
| 256-QAM | 2/3 | <i>5,31</i> | <i>17,8</i> | <i>18,1</i> | <i>20,1</i> | <i>20,3</i> |
| 256-QAM | 3/4 | <i>5,98</i> | <i>20,0</i> | <i>20,3</i> | <i>22,6</i> | <i>22,7</i> |
| 256-QAM | 4/5 | <i>6,38</i> | <i>21,3</i> | <i>21,7</i> | <i>24,3</i> | <i>24,5</i> |
| 256-QAM | 5/6 | <i>6,65</i> | <i>22,0</i> | <i>22,4</i> | <i>25,4</i> | <i>25,8</i> |

NOTE 1: Figures in italics are approximate values.
NOTE 2: Spectral efficiency does not take into account loss due to signalling / synchronization / sounding and Guard interval.
NOTE 3: The BER targets are discussed above.
NOTE 4: The expected implementation loss due to real channel estimation needs to be added to the above figures (see clause 14.4). This value will be significantly less than the corresponding figure for DVB-T in some cases, due to better optimisation of the boosting and pattern densities for DVB-T2.
NOTE 5: Entries shaded blue are results from a single implementation. All other results are confirmed by multiple implementations.

**Table 45: Required raw $(C/N)_0$ to achieve a BER = 1×10^{-7} before BCH decoding.
LDPC block length: 16 200 bits**

| | | | Required $(C/N)_0$ (dB) for BER = 1×10^{-7} after LPDC decoding | | | | |
|---|-----------|---------------------|--|-------------------------|--------------------------|----------------------------|-----------------------------|
| Constellation | Code rate | Effective Code Rate | Spectral Efficiency (see note 2) | Gaussian Channel (AWGN) | Ricean channel (F_1) | Rayleigh channel (P_1) | 0 dB echo channel @ 90 % GI |
| QPSK | 1/2 | 4/9 | <i>0,87</i> | <i>0,7</i> | <i>0,9</i> | <i>2,0</i> | <i>1,6</i> |
| QPSK | 3/5 | 3/5 | <i>1,18</i> | <i>2,5</i> | <i>2,7</i> | <i>4,1</i> | <i>3,7</i> |
| QPSK | 2/3 | 2/3 | <i>1,31</i> | <i>3,4</i> | <i>3,6</i> | <i>5,3</i> | <i>4,8</i> |
| QPSK | 3/4 | 11/15 | <i>1,45</i> | <i>4,3</i> | <i>4,6</i> | <i>6,6</i> | <i>6,2</i> |
| QPSK | 4/5 | 7/9 | <i>1,53</i> | <i>4,9</i> | <i>5,3</i> | <i>7,4</i> | <i>7,0</i> |
| QPSK | 5/6 | 37/45 | <i>1,62</i> | <i>5,5</i> | <i>5,9</i> | <i>8,3</i> | <i>7,9</i> |
| 16-QAM | 1/2 | 4/9 | <i>1,74</i> | <i>5,5</i> | <i>5,7</i> | <i>6,9</i> | <i>6,6</i> |
| 16-QAM | 3/5 | 3/5 | <i>2,36</i> | <i>7,9</i> | <i>8,2</i> | <i>9,6</i> | <i>9,3</i> |
| 16-QAM | 2/3 | 2/3 | <i>2,63</i> | <i>9,1</i> | <i>9,4</i> | <i>11,1</i> | <i>10,8</i> |
| 16-QAM | 3/4 | 11/15 | <i>2,89</i> | <i>10,3</i> | <i>10,7</i> | <i>12,8</i> | <i>12,5</i> |
| 16-QAM | 4/5 | 7/9 | <i>3,07</i> | <i>11,1</i> | <i>11,5</i> | <i>13,9</i> | <i>13,8</i> |
| 16-QAM | 5/6 | 37/45 | <i>3,25</i> | <i>11,7</i> | <i>12,2</i> | <i>15,0</i> | <i>15,0</i> |
| 64-QAM | 1/2 | 4/9 | <i>2,60</i> | <i>9,2</i> | <i>9,5</i> | <i>11,0</i> | <i>10,8</i> |
| 64-QAM | 3/5 | 3/5 | <i>3,54</i> | <i>12,3</i> | <i>12,6</i> | <i>14,4</i> | <i>14,3</i> |
| 64-QAM | 2/3 | 2/3 | <i>3,94</i> | <i>13,8</i> | <i>14,1</i> | <i>16,1</i> | <i>15,9</i> |
| 64-QAM | 3/4 | 11/15 | <i>4,34</i> | <i>15,5</i> | <i>15,8</i> | <i>18,2</i> | <i>18,0</i> |
| 64-QAM | 4/5 | 7/9 | <i>4,60</i> | <i>16,4</i> | <i>16,8</i> | <i>19,5</i> | <i>19,5</i> |
| 64-QAM | 5/6 | 37/45 | <i>4,87</i> | <i>17,1</i> | <i>17,6</i> | <i>20,6</i> | <i>20,9</i> |
| 256-QAM | 1/2 | 4/9 | <i>3,47</i> | <i>12,6</i> | <i>12,9</i> | <i>14,6</i> | <i>14,6</i> |
| 256-QAM | 3/5 | 3/5 | <i>4,72</i> | <i>16,9</i> | <i>17,2</i> | <i>19,0</i> | <i>19,3</i> |
| 256-QAM | 2/3 | 2/3 | <i>5,25</i> | <i>18,1</i> | <i>18,4</i> | <i>20,5</i> | <i>20,9</i> |
| 256-QAM | 3/4 | 11/15 | <i>5,78</i> | <i>20,3</i> | <i>20,6</i> | <i>22,9</i> | <i>23,3</i> |
| 256-QAM | 4/5 | 7/9 | <i>6,14</i> | <i>21,6</i> | <i>22,0</i> | <i>24,5</i> | <i>25,1</i> |
| 256-QAM | 5/6 | 37/45 | <i>6,49</i> | <i>22,4</i> | <i>22,9</i> | <i>25,8</i> | <i>26,6</i> |
| NOTE 1: Figures in italics are approximate values. | | | | | | | |
| NOTE 2: Spectral efficiency does not take into account loss due to signalling / synchronization / sounding and Guard interval. | | | | | | | |
| NOTE 3: The BER targets are discussed above. | | | | | | | |
| NOTE 4: The expected implementation loss due to real channel estimation needs to be added to the above figures (see clause 14.4). This value will be significantly less than the corresponding figure for DVB-T in some cases, due to better optimisation of the boosting and pattern densities for DVB-T2. | | | | | | | |
| NOTE 5: All the results in this table are from a single implementation and are therefore shaded blue. | | | | | | | |

14.2.1 Correction values for pilot boosting

The values of required $(C/N)_0$ given in table 44 and table 45 are raw and do not taken into account the reduction in data C/N resulting from the presence of boosted pilots, since this depends on the pilot pattern in use. Net values of C/N can be derived from the raw values $(C/N)_0$ by computing a correction factor Δ_{BP} such that:

$$\frac{C}{N} = \left(\frac{C}{N} \right)_0 + \Delta_{BP}$$

This correction factor is calculated by the following formula:

$$\Delta_{BP} = 10 \log_{10} \frac{(N_{data} + N_{NBP} + N_{BP} \cdot B_{BP} + N_{CP} B_{CP})}{N_{data} + N_{NBP} + N_{BP} + N_{CP}}$$

Where:

N_{data} Number of data cells per OFDM symbol.

N_{NBP} Number of non-boosted pilots per OFDM symbol.

N_{BP} Number of boosted pilots (i.e. scattered and Edge Pilots) per OFDM symbol.

B_{BP} Power boost of boosted pilots relative to data cells, equal to A_{SP}^2 .

N_{CP} Number of continual pilots per OFDM symbol.

B_{CP} Power boost of continual pilots relative to data cells, equal to A_{CP}^2 .

NOTE: The above formula is derived for the normal data symbols, but the P1, P2 and frame closing symbols are designed to have essentially the same power as the normal symbols (to within 0,1 dB) so the formula can be applied to the entire T2 frame.

The correction factor Δ_{BP} varies from 0,29 dB to 0,53 dB; values for each combination of FFT size and Scattered Pilot Pattern PP1-PP8 are given in table 46.

Table 46: Correction factors Δ_{BP} for pilots (dB)

| | PP1 | PP2 | PP3 | PP4 | PP5 | PP6 | PP7 | PP8 |
|------------------|------|------|------|------|------|------|------|------|
| 1 K | 0,34 | 0,32 | 0,44 | 0,42 | 0,48 | | 0,29 | |
| 2 K | 0,35 | 0,33 | 0,43 | 0,42 | 0,47 | | 0,29 | |
| 4 K | 0,39 | 0,37 | 0,47 | 0,45 | 0,51 | | 0,34 | |
| 8 K | 0,41 | 0,39 | 0,49 | 0,48 | 0,53 | | 0,37 | 0,37 |
| 8 K Ext, | 0,41 | 0,41 | 0,50 | 0,48 | 0,52 | | 0,39 | 0,38 |
| 16 K | 0,41 | 0,38 | 0,49 | 0,47 | 0,52 | 0,49 | 0,33 | 0,35 |
| 16 K Ext, | 0,42 | 0,38 | 0,49 | 0,47 | 0,52 | 0,49 | 0,34 | 0,35 |
| 32 K | | 0,37 | 0,48 | 0,45 | | 0,48 | 0,33 | 0,35 |
| 32 K Ext, | | 0,37 | 0,48 | 0,45 | | 0,48 | 0,33 | 0,35 |

14.3 Simulation results for Memoryless Rayleigh Channel with erasures

Table 47 gives the simulated performance of the BICM part of the DVB-T2 system in the memoryless Rayleigh channel with erasures (see clause 14.1.8). Results for the AWGN channel are also given for reference. Care should be taken in interpreting the results since the target BERs are different to those used in clause 14.2, and also because neither rotated constellations nor Genie-Aided demapping are used. Details are given in the notes to table 47.

**Table 47: Simulated performance in Memoryless Rayleigh channel with erasures
LDPC block length 64 800 bits**

| Constellation | Code rate | Spectral Efficiency (see note 2) | AWGN channel | Memoryless Rayleigh channel K=0 | Memoryless Rayleigh channel K=0,1 | Memoryless Rayleigh channel K=0,2 |
|--|-----------|----------------------------------|--------------|---------------------------------|-----------------------------------|-----------------------------------|
| QPSK | 1/2 | <i>0,99</i> | <i>1,0</i> | <i>2,8</i> | <i>3,5</i> | <i>4,5</i> |
| QPSK | 3/5 | <i>1,19</i> | <i>2,2</i> | <i>4,7</i> | <i>5,7</i> | <i>7,8</i> |
| QPSK | 2/3 | <i>1,33</i> | <i>3,1</i> | <i>5,9</i> | <i>7,2</i> | <i>9,9</i> |
| QPSK | 3/4 | <i>1,49</i> | <i>4,1</i> | <i>7,6</i> | <i>9,8</i> | <i>17,4</i> |
| QPSK | 4/5 | <i>1,59</i> | <i>4,7</i> | <i>8,9</i> | <i>12,2</i> | - |
| QPSK | 5/6 | <i>1,66</i> | <i>5,2</i> | <i>9,8</i> | <i>14,6</i> | - |
| | | | | | | |
| 16-QAM | 1/2 | <i>1,99</i> | <i>6,2</i> | <i>8,1</i> | <i>9,0</i> | <i>10,2</i> |
| 16-QAM | 3/5 | <i>2,39</i> | <i>7,6</i> | <i>10,1</i> | <i>11,3</i> | <i>13,4</i> |
| 16-QAM | 2/3 | <i>2,66</i> | <i>8,9</i> | <i>11,5</i> | <i>13,0</i> | <i>15,8</i> |
| 16-QAM | 3/4 | <i>2,99</i> | <i>10,0</i> | <i>13,3</i> | <i>15,7</i> | <i>24,0</i> |
| 16-QAM | 4/5 | <i>3,19</i> | <i>10,8</i> | <i>14,7</i> | <i>18,2</i> | - |
| 16-QAM | 5/6 | <i>3,32</i> | <i>11,3</i> | <i>15,6</i> | <i>20,6</i> | - |
| | | | | | | |
| 64-QAM | 1/2 | <i>2,98</i> | <i>10,5</i> | <i>12,6</i> | <i>13,6</i> | <i>15,0</i> |
| 64-QAM | 3/5 | <i>3,58</i> | <i>12,3</i> | <i>14,7</i> | <i>16,0</i> | <i>18,3</i> |
| 64-QAM | 2/3 | <i>3,99</i> | <i>13,6</i> | <i>16,2</i> | <i>17,8</i> | <i>20,9</i> |
| 64-QAM | 3/4 | <i>4,48</i> | <i>15,1</i> | <i>18,2</i> | <i>20,8</i> | <i>29,6</i> |
| 64-QAM | 4/5 | <i>4,78</i> | <i>16,1</i> | <i>19,7</i> | <i>23,4</i> | - |
| 64-QAM | 5/6 | <i>4,99</i> | <i>16,7</i> | <i>20,7</i> | <i>26,0</i> | - |
| | | | | | | |
| 256-QAM | 1/2 | <i>3,98</i> | <i>14,4</i> | <i>16,5</i> | <i>17,7</i> | <i>19,4</i> |
| 256-QAM | 3/5 | <i>4,78</i> | <i>16,7</i> | <i>19,0</i> | <i>20,6</i> | <i>23,2</i> |
| 256-QAM | 2/3 | <i>5,31</i> | <i>18,1</i> | <i>20,6</i> | <i>22,6</i> | <i>25,9</i> |
| 256-QAM | 3/4 | <i>5,98</i> | <i>20,0</i> | <i>22,9</i> | <i>25,8</i> | <i>35,3</i> |
| 256-QAM | 4/5 | <i>6,38</i> | <i>21,3</i> | <i>24,6</i> | <i>28,6</i> | - |
| 256-QAM | 5/6 | <i>6,65</i> | <i>22,0</i> | <i>25,6</i> | <i>31,2</i> | - |
| NOTE 1: Figures in italics are approximate values. | | | | | | |
| NOTE 2: Spectral efficiency does not take into account loss due to signalling/synchronization/sounding and Guard interval. | | | | | | |
| NOTE 3: AWGN channel results are at BER=10 ⁻⁶ (after BCH). BCH is just emulated, with a minimum of 50 erroneous FEC blocks required to achieve target BER. The BER of 10 ⁻⁶ has been used to keep simulation times reasonable. Rayleigh channel results are at BER=10 ⁻⁴ (after LDPC with no BCH). It is expected that the target packet error rate of the T2 system will be 10 ⁻⁷ . | | | | | | |
| NOTE 4: Rotated constellations are off, perfect CSI is assumed, conventional demapping (not Genie-aided) is used. | | | | | | |
| NOTE 5: Entries shaded blue are results from a single implementation. All other results are confirmed by multiple implementations. | | | | | | |

14.4 Penalty for real channel estimation

The simulations reported in the previous clause were made under the assumption that the receiver has perfect, noiseless knowledge of the channel. This is an ideal which cannot be realised, but at least sets a clear single system-performance criterion. Receivers may use various practical implementations, the choice of which will be a compromise between two or more aspects of performance, and so results for a practical receiver would not set a clear performance criterion.

It is nevertheless of interest to consider how closely a practical receiver might approach the ideal results, and under what circumstances.

A practical receiver uses measurements of the channel in order to 'equalise' the received constellations (i.e. to restore them to the right size and orientation so that the received signal can be compared with the possible constellation-spot locations, one of which will have been transmitted in each data cell. As explained in clause 10.3.2, the received demodulated value for the cell is divided by the estimated channel response. If the estimate were precisely correct (no estimation error and no noise) this would return the value transmitted, plus the noise added to the received data cell.

In practice, of course, the estimate does suffer from inaccuracies caused both by imperfect interpolation and by the presence of noise on the pilot cells.

NOTE 1: The 'equalisation' can be subsumed into the process of generating metrics without performing an explicit division, but the end result is the same.

NOTE 2: The CD3-OFDM method (see clause 10.3.2.3.5.4) functions very differently, but still has some similar fundamental limits.

The process of dividing by a noisy, imperfect channel estimate is clearly non-linear. The effect of imperfect estimates will be greater on outer constellation points than on inner ones. Nevertheless, treating the process as if channel-estimation noise is additive in effect, *while lacking any mathematical respectability*, can lead in practice to a reasonable indication of the loss in performance (from ideal noiseless channel estimation), provided the channel estimates are reasonably accurate and not too noisy.

Fortunately, we get off to a good start since the pilot cells in DVB-T2 are, in general, boosted in power compared with the data cells. This means that their signal-to-noise ratio at the receiver is correspondingly better than that of the data cells. The channel measurements that are made using the pilot cells are then interpolated to give channel estimates for all the data cells, see clause 10.3.2.3.3. The noisiness of the channel estimates is thus a function of the pilot boosting and the interpolator design. Let us suppose that the interpolator changes the noise-power level in the pilot-based measurements by a 'noise-gain' factor f_{INT} , so that $f_{INT} < 1$ corresponds to a reduction in the noise level. Assuming that the noise added to all cells, pilot and data, has the same level, it follows that we can relate the SNRs (expressed as linear power ratios) for data cells and interpolated channel estimates as follows:

$$SNR_{Channel-Estimate} = \frac{(A_{SP}^2) SNR_{Data}}{f_{INT}} = \frac{B_{BP} SNR_{Data}}{f_{INT}}$$

while under the approximation of noise-power addition we can estimate the SNR of *equalised* data cells as:

$$\begin{aligned} \frac{1}{SNR_{EQ-data}} &\approx \frac{1}{SNR_{Data}} + \frac{1}{SNR_{Channel-Estimate}} \\ &\approx \frac{1}{SNR_{Data}} \left(1 + \frac{f_{INT}}{B_{BP}} \right) \end{aligned}$$

Note that the previous section has already introduced a term Δ_{BP} which accounts for the effect of pilot boosting on the relationship between data-cell SNR and overall SNR. We now define a further correction term Δ_{RCE} so that net values of C/N required when using real, noisy channel estimation can be deduced from the tabulated simulation results using:

$$\frac{C}{N} = \left[\frac{C}{N} \right]_0 + \Delta_{BP} + \Delta_{RCE}$$

where:

$$\begin{aligned} \Delta_{RCE} &= 10 \log_{10} \frac{SNR_{Data}}{SNR_{EQ-data}} \\ &\approx 10 \log_{10} \left(1 + \frac{f_{INT}}{B_{BP}} \right) \end{aligned}$$

B_{BP} is determined by the choice of pilot pattern, taking the values {16/9, 49/16, 49/9} respectively for patterns {PP1 and 2, PP3 and 4, PP5 to 8}. f_{INT} depends on the interpolator. Indeed, strictly it varies from cell to cell according to their position within the two-dimensional Scattered Pilot pattern, as each corresponds to a combination of particular phases of the frequency and temporal interpolators. Some form of averaging will be needed to give a representative single value, taking account of both frequency and temporal interpolation.

In principle, as the design time-width/bandwidth (for frequency/temporal interpolation respectively) is reduced in comparison with the Nyquist limit, so the corresponding 'noise gain' factor reduces accordingly. However, the pilot patterns are so chosen in relation to the guard intervals in use as to imply operation approaching the Nyquist limit. It follows that very little noise benefit will be achieved by the frequency interpolator, unless it in some way adapts its time-width to match the channel. If this were done, then benefit could be had while the channel extent does not explore the full range of the guard-interval duration. Where the channel fully exploits the guard interval, or where fixed interpolation is used, it follows that the noise-gain contribution from frequency interpolation will not be much less than 1.

The temporal interpolator, when used, is expected to take the form of a simple linear interpolator (in the interests of saving receiver memory). This does offer some noise-gain benefit.

For those pilot patterns having $D_Y = 4$ (namely PP1, PP3, PP5 and PP7), the four phases of linear temporal interpolator will have noise-gain factors $\{1, 0,625, 0,5, 0,625\}$, with a mean value of 0,6875.

For those pilot patterns having $D_Y = 2$ (namely PP2, PP4, and PP6), the two phases of linear temporal interpolator will have noise-gain factors $\{1, 0,5\}$, with a mean value of 0,75.

These values will help to guide interpretation of figure 105, which illustrates the dependence of Δ_{RCE} on the overall noise-gain factor of the interpolation, f_{INT} , for the pilot boosts corresponding to the various pilot patterns.




Figure 105: Penalty for real channel estimation

14.5 P1 performance

The figures presented in this clause are the results of simulations of P1 performance run under critical conditions. These include not only the cases defined in clause 14.1 to test data performance, but also cases using different parameter values.

The figures show the robustness behaviour in terms of probability of match / error of the P1 detection vs. SNR.

14.5.1 Figure description

In the following results, each figure shows five curves for a particular scenario. The meaning of these curves is:

- T2 OK: This shows the probability of detecting a P1 preamble when it is expected to be received.
- Lost: This shows the probability of not detecting a P1 preamble when it should. In this case, the preamble is lost.
- Wrong: This shows the probability of detecting a P1 preamble when no preamble is present, i.e. a false detection.
- Sig 256: This curve represents the probability of recovering the S2 field correctly.
- Sig 64: This curve shows the probability of recovering S1 field correctly.

14.5.2 P1 performance results

The simulations were run for various channels, assuming a guard interval of 1/8 with FFT-size=2 K for the data OFDM symbols, and a CW-interference in place (CW-interference was not added in the 0 dB 180° channel).

14.5.2.1 Gaussian channel




Figure 106: P1 robustness in Gaussian channel

14.5.2.2 Rayleigh channel (P1)




Figure 107: P1 robustness in Rayleigh channel (P1)

14.5.2.3 TU-6@200Hz




Figure 108: P1 robustness in Tu-6@200Hz channel

14.5.2.4 0 dB 180° Echo, Echo with opposite sign

The worst case of a 0 dB-echo channel for the P1 symbol, are those in which the echo is destructive. This means that the echo is rotated 180 degrees. The following figure shows the probability of detection for every delay in this channel (from 0 to 2 000 samples delay) at an SNR of 0 dB.




Figure 109: P1 robustness in 0 dB 180° Echo channel (SNR=0 dB)

14.6 L1-pre- and post-signalling performance

Simulations were done to analyze the coded performance of L1-pre- and post-signalling over an AWGN channel. By using the shortening and puncturing procedures described in clause 9.2.2, L1 signalling bits are first BCH-encoded, and then the BCH encoded bits are LDPC-encoded. Note that the 16 K LDPC codes with the effective code-rates 1/5 and 4/9 are used for LDPC encoding of L1-pre- and post-signalling, respectively.

Figure 110 shows the coded performance of the L1-pre signalling and figure 111, figure 112, and figure 113 show those of L1-post signalling for a range of BCH information sizes and modulation orders. As shown in figure 111, figure 112, and figure 113, the performance for L1-post signalling is almost constant because of the flexible adaptation of LDPC code rates (see clause 9.2.2.6). Note that the effective LDPC code-rate for L1-pre signalling is fixed as 1/5, while those for L1-post signalling vary as a function of the size of BCH information, as described in clause 9.2.2.6.

The simulation parameters are as follows:

- The size of BCH information for L1-pre: 200 (bits).
- The size of BCH information for L1-post: 396, 628, 880, 1 132, 1 670, 2 140, 3 092, 3 986, 6 094, 7 032 (bits).
- Modulation order: BPSK (for L1-pre); QPSK, 16-QAM, 64-QAM (for L1-post).
- $N_{p2} = 2$ (8 K FFT Mode).
- Floating-point LDPC decoding (iteration = 50).
- BCH emulation (error correction capability = 12 bits).




Figure 110: Performance of L1 Pre (BPSK, AWGN channel)




Figure 111: Performance of L1 Post (QPSK, AWGN channel)




Figure 112: Performance of L1 Post (16QAM, AWGN channel)




Figure 113: Performance of L1 Post (64QAM, AWGN channel)

15 Examples of possible uses of the system

The DVB-T2 system is designed to be highly flexible, allowing different trade-offs to be made in terms of capacity and ruggedness, as well as flexibility and overhead. For example, the system may be used in a very simple configuration to carry a few HDTV services within a single PLP, intended for fixed rooftop reception. In this case, some typical parameter choices might be:

- 32 K FFT with 1/128 guard interval for an MFN configuration, which maximises the available capacity; or
- 32 K FFT with 19/128 guard interval for a national SFN configuration (providing a guard interval of 532 μ s); plus
- 256-QAM rotated constellations with code rate 3/5 or 2/3; the 256-QAM provides the maximum possible data capacity and is suitable for fixed rooftop reception, and the rotated constellations provide additional robustness for difficult reception conditions.

These configurations could provide approximately 36 Mbit/s to 40 Mbit/s for the MFN case, or 29 Mbit/s to 32 Mbit/s in the SFN case.

The use of the 32 K FFT will only be possible where the transmission channel is fairly static, and so for a network targeting portable and/or mobile reception, a lower FFT size and more robust constellation are likely to be used. For example 8 K FFT with 64-QAM and code rate 1/2 or 3/5 would provide good reception in more dynamic channels and with a lower carrier-to-noise ratio requirement. The penalty of course would be a lower available bit-rate, approximately 16 Mbit/s to 26 Mbit/s, depending on other parameter choices.

Annex A (informative): Derivation of allowable and recommended sub-slicing values

The Interleaving frames at the output of the time interleaver may be split into a multiple number of sub-slices, in order to allow for maximum flexibility in selecting time diversity characteristics. According to clause 6.5.4 of [i.1], sub-slicing may be applied in conjunction with three different time interleaving options that differ in the choice of the parameters for the number of T2 frames to which each interleaving frame is mapped (P_I), the frame interval and thus the difference in the frame indices between successive T2 frames to which a particular PLP is mapped (I_{JUMP}), and finally the number of time interleaving blocks in an interleaving frame (N_{TI}). According to the definition given in [i.1], the type 2 PLPs are transmitted immediately after the type 1 PLPs and feature at least two sub-slices per T2 frame. Clause 8.3.6.3.2 of [i.1] addresses the specific issue of type 2 PLP mapping.




Figure A.1: Source diagram of divisors for valid numbers of sub-slices

The derivation of allowable sub-slicing values for a certain constellation and LDPC-block length starts with the determination of potential candidate values for the allowable number of sub-slices. As a minimum, any valid number of sub-slices needs to be an integer divisor of one fifth of the LDPC-block length after QAM mapping, i.e. 6 480 or 1 620 for QPSK. As such, any valid number of sub-slices needs to be an integer divisor of one tenth of the respective LDPC-block length. The numbers that satisfy this condition are fully reflected in annex K of the specification document, for each constellation size.

Figure A.1 shows the entire Hasse source diagram of all divisors of 6 480 and their theoretical connectivity. These constitute the valid numbers of sub-slices for the long (64 800-bit) LDPC-block size with QPSK mapping. The subset that is represented by the grey and black cells in the bottom right represents the valid number of sub-slices for the short (16 200-bit) LDPC-block size with QPSK mapping. All of these are divisors of 1 620. For constellations exceeding QPSK, the valid number of sub-slices will be subsets of the above. In order to find out whether a certain value constitutes an allowed number of sub-slices for a given constellation size and LDPC-block length, these parameters need to be additionally taken into account. For each potential candidate N , a value is valid as an allowed number of sub-slices, if and only if the calculation of:

$$\frac{N_{ldpc}}{N \cdot \eta_{mod}}$$

results in an integer value, where the LDPC block length N_{ldpc} is 64 800 or 16 200 and η_{mod} is the number of code-bits per symbol, i.e. $\log_2(\text{QAM_size})$ where QAM_size is 4, 16, 64 or 256. Accordingly, for a certain combination of PLPs, each with its own LDPC block length and constellation size, the allowed number of sub-slices may be described by all positive integers N that satisfy the relationship:

$$N_{ldpc} \bmod(5 \cdot \eta_{mod} \cdot N) = 0$$

for each PLP. This is the relationship that is used for determining the values in table K.1 of [i.1].

Finally, aiming at the extraction of the most flexible numbers of sub-slices, the values bordered in black represent the subset of the valid numbers of sub-slices that are valid for all possible combinations of LDPC block lengths and constellation sizes, i.e. the intersection of all possible subsets. These correspond to table K.2 of EN 302 755 [i.1]. In the same way, the subset below the node '540' constitutes the list of valid values for number of sub-slices which may be used with any combination of PLPs for the 64 800-bit-long LDPC-blocks only. These correspond to table K.3 [i.1].

Clause 5.8.1 discussed a further consideration affecting the choice of the number of sub-slices, because of the potential for interaction between the sub-slicing, carrier mapping and frequency interleaving which can reduce the amount of frequency diversity. In particular, for some sub-slicing factors, depending on the actual number of LDPC blocks within each PLP, it may happen that the cells of each PLP are concentrated amongst a relatively limited number of carriers, which may result in a loss of frequency diversity. Some loss of frequency diversity always occurs whenever the resulting sub-slice interval, measured in cells, shares common factors with $2 \times 5 \times C_{DATA}$, where C_{DATA} is the number of data cells per OFDM symbol, '2' results from the odd/even frequency interleaver, and '5' is the split-parameter of the time interleaver. The sub-slice interval is given by:

$$SUB_SLICE_INTERVAL = \sum_{all\ PLPs} N_{BLOCKS_IF} \cdot \frac{N_{ldpc}}{\eta_{mod} \cdot N_{subslices} \cdot P_1}$$

Note that the sub-slice interval depends on N_{BLOCKS_IF} , the number of LDPC blocks within the Interleaving Frame for each PLP. It is therefore constant for CBR streams, and is also constant for VBR streams using the same modulation and coding and derived from a constant-bit-rate statistical multiplex, since the total number of FEC blocks will be constant in this case. It will be variable for VBR streams using different modulation and/or code rate.

The interleaving length parameter P_1 is included to take into account the use of multi-frame interleaving for some PLPs.

In conclusion, in the case of CBR streams only, it is recommended to choose the number of sub-slices so as to avoid having common factors between $2 \cdot 5 \cdot C_{DATA}$ and the sub-slice interval.

History

| Document history | | |
|-------------------------|--------------|-------------|
| V1.1.1 | October 2010 | Publication |
| | | |
| | | |
| | | |
| | | |