ETSI TS 102 366 V1.4.1 (2017-09)

Digital Audio Compression
(AC-3, Enhanced AC-3) Standard

EBU
OPERATING EUROVISION

Reference
RTS/JTC-041

Keywords
audio, broadcasting, coding, digital

ETSI

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

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

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

Important notice

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

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

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

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

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.
The content of the PDF version shall not be modified without the written authorization of ETSI.
The copyright and the foregoing restriction extend to reproduction in all media.

© ETSI 2017.
© European Broadcasting Union 2017.
All rights reserved.

DECT™, **PLUGTESTS**™, **UMTS**™ and the ETSI logo are trademarks of ETSI registered for the benefit of its Members.
3GPP™ and **LTE**™ are trademarks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners.
oneM2M logo is protected for the benefit of its Members.
GSM® and the GSM logo are trademarks registered and owned by the GSM Association.

Contents

Intellectual Property Rights 17
Foreword 17
Modal verbs terminology 17
Introduction 18
1 Scope 21
2 References 21
2.1 Normative references 21
2.2 Informative references 21
3 Definitions and abbreviations 22
3.1 Definitions 22
3.2 Abbreviations 24
4 Bit stream syntax 27
4.1 Synchronization frame 27
4.2 Semantics of syntax specification 27
4.3 Syntax specification 28
4.3.0 AC-3_bit_stream and syncframe 28
4.3.1 syncinfo - Synchronization information 28
4.3.2 bsi - Bit stream information 28
4.3.3 audblk - Audio block 29
4.3.4 auxdata - Auxiliary data 32
4.3.5 errorcheck - Error detection code 32
4.4 Description of bit stream elements 32
4.4.0 Introduction 32
4.4.1 syncinfo - Synchronization information 32
4.4.1.1 syncword - Synchronization word - 16 bits 32
4.4.1.2 crc1 - Cyclic redundancy check 1 bit to 16 bits 32
4.4.1.3 fscod - Sample rate code - 2 bits 33
4.4.1.4 frmsizecod - Frame size code - 6 bits 33
4.4.2 bsi - Bit stream information 33
4.4.2.1 bsid - Bit stream identification - 5 bits 33
4.4.2.2 bsmod - Bit stream mode - 3 bits 33
4.4.2.3 acmod - Audio coding mode - 3 bits 33
4.4.2.4 cmixlev - Centre mix level - 2 bits 34
4.4.2.5 surmixlev - Surround mix level - 2 bits 34
4.4.2.6 dsurmod - Dolby® Surround mode - 2 bits 34
4.4.2.7 lfeon - Low frequency effects channel on - 1 bit 35
4.4.2.8 dialnorm - Dialogue normalization - 5 bits 35
4.4.2.9 compre - Compression gain word exists - 1 bit 35
4.4.2.10 compr - Compression gain word - 8 bits 35
4.4.2.11 langcode - Language code exists - 1 bit 35
4.4.2.12 langcod - Language code - 8 bits 35
4.4.2.13 audprodie - Audio production information exists - 1 bit 35
4.4.2.14 mixlevel - Mixing level - 5 bits 36
4.4.2.15 roomtyp - Room type - 2 bits 36
4.4.2.16 dialnorm2 - Dialogue normalization, Ch2 - 5 bits 36
4.4.2.17 compr2e - Compression gain word exists, Ch2 - 1 bit 36
4.4.2.18 compr2 - Compression gain word, Ch2 - 8 bits 36
4.4.2.19 langcod2e - Language code exists, Ch2 - 1 bit 36
4.4.2.20 langcod2 - Language code, Ch2 - 8 bits 36
4.4.2.21 audprodi2e - Audio production information exists, Ch2 - 1 bit 36
4.4.2.22 mixlevel2 - Mixing level, Ch2 - 5 bits 36
4.4.2.23 roomtyp2 - Room type, Ch2 - 2 bits 37
4.4.2.24 copyrightb - Copyright bit - 1 bit 37

4.4.2.25 origbs - Original bit stream - 1 bit 37
4.4.2.26 timecod1e, timecod2e - Time code (first and second) halves exists - 2 bits 37
4.4.2.27 timecod1 - Time code first half - 14 bits 37
4.4.2.28 timecod2 - Time code second half - 14 bits 37
4.4.2.29 addbsie - Additional bit stream information exists - 1 bit 37
4.4.2.30 addbsil - Additional bit stream information length - 6 bits 37
4.4.2.31 addbsi - Additional bit stream information - ((addbsil + 1) x 8) bits 37
4.4.3 audblk - Audio block 38
4.4.3.1 blksw[ch] - Block switch flag - 1 bit 38
4.4.3.2 dithflag[ch] - Dither flag - 1 bit 38
4.4.3.3 dynrnge - Dynamic range gain word exists - 1 bit 38
4.4.3.4 dynrng - Dynamic range gain word - 8 bits 38
4.4.3.5 dynrng2e - Dynamic range gain word exists, Ch2 - 1 bit 38
4.4.3.6 dynrng2 - dynamic range gain word, Ch2 - 8 bits 38
4.4.3.7 cplstre - Coupling strategy exists - 1 bit 38
4.4.3.8 cplinu - Coupling in use - 1 bit 38
4.4.3.9 chincpl[ch] - Channel in coupling - 1 bit 38
4.4.3.10 phsflginu - Phase flags in use - 1 bit 38
4.4.3.11 cplbegf - Coupling begin frequency code - 4 bits 39
4.4.3.12 cplendf - Coupling end frequency code - 4 bits 39
4.4.3.13 cplbndstrc[sbnd] - Coupling band structure - 1 bit 39
4.4.3.14 cplcoe[ch] - Coupling coordinates exist - 1 bit 39
4.4.3.15 mstrcplco[ch] - Master coupling coordinate - 2 bits 40
4.4.3.16 cplcoexp[ch][bnd] - Coupling coordinate exponent - 4 bits 40
4.4.3.17 cplcomant[ch][bnd] - Coupling coordinate mantissa - 4 bits 40
4.4.3.18 phsflg[bnd] - Phase flag - 1 bit 40
4.4.3.19 rematstr - Rematrixing strategy - 1 bit 40
4.4.3.20 rematflg[rbnd] - Rematrix flag - 1 bit 40
4.4.3.21 cplexpstr - Coupling exponent strategy - 2 bits 40
4.4.3.22 chexpstr[ch] - Channel exponent strategy - 2 bits 41
4.4.3.23 lfeexpstr - Low frequency effects channel exponent strategy - 1 bit 41
4.4.3.24 chbwcod[ch] - Channel bandwidth code - 6 bits 41
4.4.3.25 cplabsexp - Coupling absolute exponent - 4 bits 41
4.4.3.26 cplexps[grp] - Coupling exponents - 7 bits 41
4.4.3.27 exps[ch][grp] - Channel exponents - 4 bits or 7 bits 41
4.4.3.28 gainrng[ch] - Channel gain range code - 2 bits 41
4.4.3.29 lfeexps[grp] - Low frequency effects channel exponents - 4 bits or 7 bits 41
4.4.3.30 baie - Bit allocation information exists - 1 bit 41
4.4.3.31 sdcycod - Slow decay code - 2 bits 41
4.4.3.32 fdcycod - Fast decay code - 2 bits 42
4.4.3.33 sgaincod - Slow gain code - 2 bits 42
4.4.3.34 dbpbcod - dB per bit code - 2 bits 42
4.4.3.35 floorcod - Masking floor code - 3 bits 42
4.4.3.36 snroffste - SNR offset exists - 1 bit 42
4.4.3.37 csnroffst - Coarse SNR offset - 6 bits 42
4.4.3.38 cplfsnroffst - Coupling fine SNR offset - 4 bits 42
4.4.3.39 cplfgaincod - Coupling fast gain code - 3 bits 42
4.4.3.40 fsnroffst[ch] - Channel fine SNR offset - 4 bits 42
4.4.3.41 fgaincod[ch] - Channel fast gain code - 3 bits 42
4.4.3.42 lfefsnroffst - Low frequency effects channel fine SNR offset - 4 bits 42
4.4.3.43 lfefgaincod - Low frequency effects channel fast gain code - 3 bits 42
4.4.3.44 cplleake - Coupling leak initialization exists - 1 bit 42
4.4.3.45 cplfleak - Coupling fast leak initialization - 3 bits 43
4.4.3.46 cplsleak - Coupling slow leak initialization - 3 bits 43
4.4.3.47 deltbaie - Delta bit allocation information exists - 1 bit 43
4.4.3.48 cpldeltbae - Coupling delta bit allocation exists - 2 bits 43
4.4.3.49 deltbae[ch] - Delta bit allocation exists - 2 bits 43
4.4.3.50 cpldeltnseg - Coupling delta bit allocation number of segments - 3 bits 43
4.4.3.51 cpldeltoffst[seg] - Coupling delta bit allocation offset - 5 bits 43
4.4.3.52 cpldeltlen[seg] - Coupling delta bit allocation length - 4 bits 43
4.4.3.53 cpldeltba[seg] - Coupling delta bit allocation - 3 bits 43
4.4.3.54 deltnseg[ch] - Channel delta bit allocation number of segments - 3 bits 44

4.4.3.55 deltoffst[ch][seg] - Channel delta bit allocation offset - 5 bits 44
4.4.3.56 deltlen[ch][seg] - Channel delta bit allocation length - 4 bits 44
4.4.3.57 deltba[ch][seg] - Channel delta bit allocation - 3 bits 44
4.4.3.58 skiple - Skip length exists - 1 bit 44
4.4.3.59 skipl - Skip length - 9 bits 44
4.4.3.60 skipfld - Skip field - (skipl x 8) bits 44
4.4.3.61 chmant[ch][bin] - Channel mantissas - 0 bits to 16 bits 44
4.4.3.62 cplmant[bin] - Coupling mantissas - 0 bits to 16 bits 45
4.4.3.63 lfemant[bin] - Low frequency effects channel mantissas - 0 bits to 16 bits 45
4.4.4 auxdata - Auxiliary data field 45
4.4.4.0 Introduction 45
4.4.4.1 auxbits - Auxiliary data bits - nauxbits bits 45
4.4.4.2 auxdatal - Auxiliary data length - 14 bits 46
4.4.4.3 auxdatae - Auxiliary data exists - 1 bit 46
4.4.5 errorcheck - Frame error detection field 46
4.4.5.1 crcrsv - CRC reserved bit - 1 bit 46
4.4.5.2 crc2 - Cyclic redundancy check 2 - 16 bits 46
4.5 Bit stream constraints 46

5 Decoding the AC-3 bit stream 47
5.1 Introduction 47
5.2 Summary of the decoding process 47
5.2.1 Input bit stream 47
5.2.1.0 Introduction 47
5.2.1.1 Continuous or burst input 47
5.2.1.2 Byte or word alignment 47
5.2.2 Synchronization and error detection 48
5.2.3 Unpack BSI, side information 49
5.2.4 Decode exponents 50
5.2.5 Bit allocation 50
5.2.6 Process mantissas 50
5.2.7 Decoupling 50
5.2.8 Rematrixing 50
5.2.9 Dynamic range compression 50
5.2.10 Inverse transform 50
5.2.11 Window, overlap/add 51
5.2.12 Downmixing 51
5.2.13 PCM output buffer 51
5.2.14 Output PCM 51

6 Algorithmic details 51
6.0 Introduction 51
6.1 Exponent coding 51
6.1.1 Overview 51
6.1.2 Exponent strategy 52
6.1.3 Exponent decoding 53
6.2 Bit allocation 55
6.2.1 Overview 55
6.2.2 Parametric bit allocation 56
6.2.2.0 Introduction 56
6.2.2.1 Initialization 56
6.2.2.2 Exponent mapping into psd 57
6.2.2.3 psd integration 57
6.2.2.4 Compute excitation function 58
6.2.2.5 Compute masking curve 59
6.2.2.6 Apply delta bit allocation 60
6.2.2.7 Compute bit allocation 60
6.2.3 Bit allocation tables 61
6.3 Quantization and decoding of mantissas 65
6.3.1 Overview 65
6.3.2 Expansion of mantissas for asymmetric quantization ($6 \leq bap \leq 15$) 66
6.3.3 Expansion of mantissas for symmetrical quantization ($1 \leq bap \leq 5$) 66

6.3.4 Dither for zero bit mantissas (bap = 0) 66
6.3.5 Ungrouping of mantissas 68
6.4 Channel coupling 68
6.4.1 Overview 68
6.4.2 Sub-band structure for coupling 69
6.4.3 Coupling coordinate format 70
6.5 Rematrixing 70
6.5.1 Overview 70
6.5.2 Frequency band definitions 71
6.5.2.0 Introduction 71
6.5.2.1 Coupling not in use 71
6.5.2.2 Coupling in use, cplbegf > 2 71
6.5.2.3 Coupling in use, 2 ≥ cplbegf > 0 72
6.5.2.4 Coupling in use, cplbegf = 0 72
6.5.3 Encoding technique 72
6.5.4 Decoding technique 73
6.6 Dialogue normalization 73
6.7 Dynamic range compression 74
6.7.1 Overview 74
6.7.2 Dynamic range control; dynrng, dynrng2 74
6.7.2.1 Overview 74
6.7.2.2 Detailed implementation 75
6.7.3 Heavy compression; compr, compr2 76
6.7.3.1 Overview 76
6.7.3.2 Detailed implementation 76
6.8 Downmixing 77
6.8.0 Introduction 77
6.8.1 General downmix procedure 77
6.8.2 Downmixing into two channels 80
6.9 Transform equations and block switching 81
6.9.1 Overview 81
6.9.2 Technique 82
6.9.3 Decoder implementation 82
6.9.4 Transformation equations 82
6.9.4.1 512-sample IMDCT transform 82
6.9.4.2 256-sample IMDCT transforms 84
6.9.5 Channel gain range code 86
6.10 Error detection 87
6.10.0 Introduction 87
6.10.1 CRC checking 87
6.10.2 Checking bit stream consistency 88

7 Encoding the AC-3 bit stream 90
7.1 Introduction 90
7.2 Summary of the encoding process 90
7.2.1 Input PCM 90
7.2.1.1 Input word length 90
7.2.1.2 Input sample rate 90
7.2.1.3 Input filtering 90
7.2.2 Transient detection 91
7.2.3 Forward transform 93
7.2.3.1 Windowing 93
7.2.3.2 Time to frequency transformation 93
7.2.4 Coupling strategy 93
7.2.4.1 Basic encoder 93
7.2.4.2 Advanced encoder 93
7.2.5 Form coupling channel 93
7.2.5.1 Coupling channel 93
7.2.5.2 Coupling coordinates 94
7.2.6 Rematrixing 94
7.2.7 Extract exponents 94
7.2.8 Exponent strategy 94

7.2.9 Dither strategy 94
7.2.10 Encode exponents 94
7.2.11 Normalize mantissas 95
7.2.12 Core bit allocation 95
7.2.13 Quantize mantissas 95
7.2.14 Pack AC-3 syncframe 95

Annex A (normative): AC-3 bit streams in the MPEG-2 multiplex 96

A.0 Scope 96

A.1 Introduction 96

A.2 Detailed specification for System A (ATSC) 96
A.2.1 Stream_type 96
A.2.2 Stream_id 96
A.2.3 Registration_descriptor 97
A.2.4 AC-3 audio_descriptor 97
A.2.5 ISO_639_language_code 100
A.2.6 STD audio buffer size 100

A.3 Specification for System B (DVB) 101

A.4 PES constraints 101
A.4.0 Introduction 101
A.4.1 Encoding 101
A.4.2 Decoding 101
A.4.3 Byte-alignment 102

Annex B (informative): Void 103

Annex C (informative): AC-3 karaoke mode 104

C.0 Scope 104

C.1 Introduction 104

C.2 Detailed specification 104
C.2.1 Karaoke mode indication 104
C.2.2 Karaoke mode channel assignment 104
C.2.3 Reproduction of karaoke mode bit streams 105
C.2.3.0 Introduction 105
C.2.3.1 Karaoke-aware decoder 105
C.2.3.2 Karaoke-capable decoders 105

Annex D (normative): Alternate bit stream syntax 107

D.0 Scope 107

D.1 Specification 107
D.1.1 Indication of alternate bit stream syntax 107
D.1.2 Alternate bit stream syntax specification 107
D.1.3 Description of alternate syntax bit stream elements 108
D.1.3.0 Introduction 108
D.1.3.1 xbsi1e: Extra bit stream information #1 exists, 1 bit 108
D.1.3.2 dmixmod: Preferred stereo downmix mode, 2 bits 108
D.1.3.3 ltrtcmixlev: Lt/Rt centre mix level, 3 bits 108
D.1.3.4 ltrtsurmixlev: Lt/Rt surround mix level, 3 bits 109
D.1.3.5 lorocmixlev: Lo/Ro centre mix level, 3 bits 109
D.1.3.6 lorosurmixlev: Lo/Ro surround mix level, 3 bits 110
D.1.3.7 xbsi2e: Extra bit stream information #2 exists, 1 bit 110
D.1.3.8 dsurexmod: Dolby® Surround EX® mode, 2 bits 110
D.1.3.9 dheadphonmod: Dolby® Headphone mode, 2 bits 110
D.1.3.10 adconvtyp: A/D converter type, 1 bit 111
D.1.3.11 xbsi2: Extra bit stream information, 8 bits 111

D.1.3.12 encinfo: Encoder information, 1 bit 111
D.2 Decoder processing 111
D.2.0 Introduction 111
D.2.1 Compliant decoder processing 111
D.2.1.1 Two-channel downmix selection 111
D.2.1.2 Two-channel downmix processing 111
D.2.1.3 Informational parameter processing 111
D.2.2 Legacy decoder processing 112
D.3 Encoder processing 112
D.3.0 Introduction 112
D.3.1 Encoder processing steps 112
D.3.1.1 Dynamic range overload protection processing 112
D.3.2 Encoder requirements 112
D.3.2.1 Legacy decoder support 112
D.3.2.2 Original bit stream syntax support 112
Annex E (normative): Enhanced AC-3 113
E.0 Scope 113
E.1 Bit stream syntax and semantics specification 113
E.1.1 Indication of Enhanced AC-3 bit stream syntax 113
E.1.2 Syntax specification 113
E.1.2.0 E-AC-3_bit_stream and syncframe 113
E.1.2.1 syncinfo - Synchronization information 114
E.1.2.2 bsi - Bit stream information 114
E.1.2.3 audfrm - Audio frame 116
E.1.2.4 audblk - Audio block 118
E.1.2.5 auxdata - Auxiliary data 125
E.1.2.6 errorcheck - Error detection code 125
E.1.3 Description of Enhanced AC-3 bit stream elements 126
E.1.3.0 Introduction 126
E.1.3.1 bsi - Bit stream information 126
E.1.3.1.1 strmtyp - Stream type - 2 bits 126
E.1.3.1.2 substreamid - Substream identification - 3 bits 126
E.1.3.1.3 frmsiz - Frame size - 11 bits 127
E.1.3.1.4 fscod - Sample rate code - 2 bits 127
E.1.3.1.5 numblkscod - Number of audio blocks - 2 bits 127
E.1.3.1.6 bsid - Bit stream identification - 5 bits 127
E.1.3.1.7 chanmape - Custom channel map exists - 1 bit 127
E.1.3.1.8 chanmap - Custom channel map - 16 bits 128
E.1.3.1.9 mixmdate - Mixing Meta-data exists - 1 bit 128
E.1.3.1.10 lfemixlevcode - LFE mix level code exists - 1 bit 128
E.1.3.1.11 lfemixlevcod - LFE mix level code - 5 bits 128
E.1.3.1.12 pgmscle - Programme scale factor exists - 1 bit 129
E.1.3.1.13 pgmscl - Programme scale factor - 6 bits 129
E.1.3.1.14 pgmscl2e - Programme scale factor #2 exists - 1 bit 129
E.1.3.1.15 pgmscl2 - Programme scale factor #2 - 6 bits 129
E.1.3.1.16 extpgmscle - External programme scale factor exists - 1 bit 129
E.1.3.1.17 extpgmscl - External programme scale factor - 6 bits 129
E.1.3.1.18 mixdef - Mix control field length - 2 bits 129
E.1.3.1.19 premixcmpsel - Premix compression word select - 1 bit 129
E.1.3.1.20 drcsrc - Dynamic range control word source for the mixed output - 1 bit 130
E.1.3.1.21 premixcmpscl - Premix compression word scale factor - 3 bits 130
E.1.3.1.22 mixdeflen - Length of mixing parameter data field - 5 bits 130
E.1.3.1.23 mixdata - Mixing parameter data - (5 to 264) bits 130
E.1.3.1.24 mixdata2e - Mixing parameters for individual channel scaling exist - 1 bit 130
E.1.3.1.25 extpgmlscle - External programme left channel scale factor exists - 1 bit 130
E.1.3.1.26 extpgmlscl - External programme left channel scale factor - 4 bits 130
E.1.3.1.27 extpgmcscle - External programme centre channel scale factor exists - 1 bit 131
E.1.3.1.28 extpgmcscl - External programme centre channel scale factor - 4 bits 131

E.1.3.1.29 extpgmrscle - External programme right channel scale factor exists - 1 bit....131
E.1.3.1.30 extpgmrscl - External programme right channel scale factor - 4 bits....131
E.1.3.1.31 extpgmlsscle - External programme left surround channel scale factor exists - 1 bit....131
E.1.3.1.32 extpgmlsscl - External programme left surround channel scale factor - 4 bits....131
E.1.3.1.33 extpgmrsscle - External programme right surround channel scale factor exists - 1 bit....132
E.1.3.1.34 extpgmrsscl - External programme right surround channel scale factor - 4 bits....132
E.1.3.1.35 extpgmlfescle - External programme LFE channel scale factor exists - 1 bit....132
E.1.3.1.36 extpgmlfescl - External programme LFE channel scale factor - 4 bits....132
E.1.3.1.37 dmixscle - External programme downmix scale factor exists - 1 bit....132
E.1.3.1.38 dmixscl - External programme downmix scale factor - 4 bits....132
E.1.3.1.39 addche - Scale factors for additional external programme channels exist - 1 bit....132
E.1.3.1.40 extpgmaux1scle - External programme first auxiliary channel scale factor exists - 1 bit....132
E.1.3.1.41 extpgmaux1scl - External programme first auxiliary channel scale factor - 4 bits....133
E.1.3.1.42 extpgmaux2scle - External programme second auxiliary channel scale factor exists - 1 bit....133
E.1.3.1.43 extpgmaux2scl - External programme second auxiliary channel scale factor - 4 bits....133
E.1.3.1.44 mixdata3e - Mixing parameters for speech processing exist - 1 bit....133
E.1.3.1.45 spchdat - Speech enhancement processing data - 5 bits....133
E.1.3.1.46 addspchdate - Additional speech enhancement processing data exists - 1 bit....133
E.1.3.1.47 spchdat1 - Additional speech enhancement processing data - 5 bits....133
E.1.3.1.48 spchan1att - Speech enhancement processing attenuation data - 2 bits....133
E.1.3.1.49 addspchdat1e - Additional speech enhancement processing data exists - 1 bit....133
E.1.3.1.50 spchdat2 - Additional speech enhancement processing data - 5 bits....134
E.1.3.1.51 spchan2att - Speech enhancement processing attenuation data - 3 bits....134
E.1.3.1.52 mixdatafill - Mixdata field fill bits - 0 to 7 bits....134
E.1.3.1.53 paninfoe - Pan information exists - 1 bit....134
E.1.3.1.54 panmean - Pan mean direction index - 8 bits....134
E.1.3.1.55 paninfo - reserved - 6 bits....134
E.1.3.1.56 paninfo2e - Pan information exists - 1 bit....134
E.1.3.1.57 panmean2 - Pan mean direction index - 8 bits....134
E.1.3.1.58 paninfo2 - reserved - 6 bits....134
E.1.3.1.59 frmmixcfginfoe - Frame mixing configuration information exists - 1 bit....134
E.1.3.1.60 blkmixcfginfoe - Block mixing configuration information exists - 1 bit....134
E.1.3.1.61 blkmixcfginfo[blk] - Block mixing configuration information - 5 bits....135
E.1.3.1.62 infomdate - Informational metadata Exists - 1 bit....135
E.1.3.1.63 sourcefscod - Source sample rate code - 1 bit....135
E.1.3.1.64 convsync - Converter synchronization flag - 1 bit....135
E.1.3.1.65 blkid - Block identification - 1 bit....135
E.1.3.2 audfrm - Audio frame....135
E.1.3.2.1 expstre - Exponent strategy syntax enabled - 1 bit....135
E.1.3.2.2 ahte - Adaptive hybrid transform enabled - 1 bit....135
E.1.3.2.3 snroffststr - SNR offset strategy - 2 bits....135
E.1.3.2.4 transproce - Transient pre-noise processing enabled - 1 bit....136
E.1.3.2.5 blkswe - Block switch syntax enabled - 1 bit....136
E.1.3.2.6 dithflage - Dither flag syntax enabled - 1 bit....136
E.1.3.2.7 bamode - Bit allocation model syntax enabled - 1 bit....136
E.1.3.2.8 frmfgaincode - Fast gain codes enabled - 1 bit....136
E.1.3.2.9 dbaflde - Delta bit allocation syntax enabled - 1 bit....136
E.1.3.2.10 skipflde - Skip field syntax enabled - 1 bit....136
E.1.3.2.11 spxattene - Spectral extension attenuation enabled - 1 bit....136
E.1.3.2.12 frmcplexpstr - Frame based coupling exponent strategy - 5 bits....136
E.1.3.2.13 frmchexpstr[ch] - Frame based channel exponent strategy - 5 bits....136
E.1.3.2.14 convexpstre - Converter exponent strategy exists - 1 bit....137
E.1.3.2.15 convexpstr[ch] - Converter channel exponent strategy - 5 bits....137
E.1.3.2.16 cplahtinu - Coupling channel AHT in use - 1 bit....137
E.1.3.2.17 chahtinu[ch] - Channel AHT in use - 1 bit....137
E.1.3.2.18 lfeahtinu - LFE channel AHT in use - 1 bit....137
E.1.3.2.19 frmcsnroffst - Frame coarse SNR offset - 6 bits....138
E.1.3.2.20 frmfsnroffst - Frame fine SNR offset - 4 bits....138
E.1.3.2.21 chintransproc[ch] - Channel in transient pre-noise processing - 1 bit....138
E.1.3.2.22 transprocloc[ch] - Transient location relative to start of frame - 10 bits....138
E.1.3.2.23 transproclen[ch] - Transient processing length - 8 bits....138
E.1.3.2.24 chinspxatten[ch] - Channel in spectral extension attenuation processing - 1 bit....138

E.1.3.2.25 spxattencod[ch] - Spectral extension attenuation code - 5 bits....138
E.1.3.2.26 blkstrtinfoe - Block start information exists - 1 bit....138
E.1.3.2.27 blkstrtinfo - Block start information - nblkstrtbits....138
E.1.3.2.28 firstspxcos[ch] - First spectral extension coordinates states - 1 bit....139
E.1.3.2.29 firstcplcos[ch] - First coupling coordinates states - 1 bit....139
E.1.3.2.30 firstcplleak - First coupling leak state - 1 bit....139
E.1.3.3 audblk - Audio block....139
E.1.3.3.1 spxstre - Spectral extension strategy exists - 1 bit....139
E.1.3.3.2 spxinu - Spectral extension in use - 1 bit....139
E.1.3.3.3 chinspx[ch] - Channel using spectral extension - 1 bit....139
E.1.3.3.4 spxstrtf - Spectral extension start copy frequency code - 2 bits....139
E.1.3.3.5 spxbegf - Spectral extension begin frequency code - 3 bits....139
E.1.3.3.6 spxendf - Spectral extension end frequency code - 3 bits....140
E.1.3.3.7 spxbndstrce - Spectral extension band structure exist - 1 bit....140
E.1.3.3.8 spxbndstrc[bnd] - Spectral extension band structure - 1 bit to 14 bits....140
E.1.3.3.9 spxcoe[ch] - Spectral extension coordinates exist - 1 bit....140
E.1.3.3.10 spxblnd[ch] - Spectral extension blend - 5 bits....140
E.1.3.3.11 mstrspxco[ch] - Master spectral extension coordinate - 2 bits....141
E.1.3.3.12 spxcoexp[ch][bnd] - Spectral extension coordinate exponent - 4 bits....141
E.1.3.3.13 spxcomant[ch][bnd] - Spectral extension coordinate mantissa - 2 bits....141
E.1.3.3.14 ecplinu - Enhanced coupling in use - 1 bit....141
E.1.3.3.15 cplbndstrce - Coupling banding structure exist - 1 bit....141
E.1.3.3.16 ecplbegf - Enhanced coupling begin frequency code - 4 bits....142
E.1.3.3.17 ecplendf - Enhanced coupling end frequency code - 4 bits....142
E.1.3.3.18 ecplbndstrce - Enhanced coupling banding structure exists - 1 bit....142
E.1.3.3.19 ecplbndstrc[sbnd] - Enhanced coupling band (and sub-band) structure - 1 bit....142
E.1.3.3.20 ecplparam1e[ch] - Enhanced coupling parameters 1 exist - 1 bit....143
E.1.3.3.21 ecplamp[ch][bnd] - Enhanced coupling amplitude scaling - 5 bits....143
E.1.3.3.22 rsvdfieldse - reserved fields exist - 1 bit....143
E.1.3.3.23 blkfsnroffst - Block fine SNR offset - 4 bits....143
E.1.3.3.24 fgaincode - Fast gain codes exist - 1 bit....143
E.1.3.3.25 convsnroffste - Converter SNR offset exists - 1 bit....143
E.1.3.3.26 convsnroffst - Converter SNR offset - 10 bits....144
E.1.3.3.27 chgaqmod[ch] - Channel gain adaptive quantization mode - 2 bits....144
E.1.3.3.28 chgaqgain[ch][n] - Channel gain adaptive quantization gain - 1 bit or 5 bits....144
E.1.3.3.29 pre_chmant[n][ch][bin] - Pre channel mantissas - 0 to 16 bits....144
E.1.3.3.30 cplgaqmod - Coupling channel gain adaptive quantization mode - 2 bits....144
E.1.3.3.31 cplgaqgain[n] - Coupling gain adaptive quantization gain - 1 bit or 5 bits....144
E.1.3.3.32 pre_cplmant[n][bin] - Pre coupling channel mantissas - 0 bits to 16 bits....144
E.1.3.3.33 lfegaqmod - LFE channel gain adaptive quantization mode - 2 bits....144
E.1.3.3.34 lfegaqgain[n] - LFE gain adaptive quantization gain - 1 bit or 5 bits....144
E.1.3.3.35 pre_lfemant[n][bin] - Pre LFE channel mantissas - 0 bits to 16 bits....144

E.2 Decoder processing....145
E.2.0 Introduction....145
E.2.1 Glitch-free switching between different stream types....145
E.2.2 Error detection and concealment....145
E.2.3 Modifications to previously defined parameters....145
E.2.3.0 Introduction....145
E.2.3.1 cplendf - Coupling end frequency code....145
E.2.3.2 nrematbd - Number of rematrixing bands....145
E.2.3.3 endmant - End mantissa....146
E.2.3.4 nchmant - Number of fbw channel mantissas....146
E.2.3.5 ncplgrps - Number of coupled exponent groups....146
E.2.4 Adaptive Hybrid Transform processing....147
E.2.4.1 Overview....147
E.2.4.2 Bit stream helper variables....147
E.2.4.3 Bit allocation....151
E.2.4.3.0 Introduction....151
E.2.4.3.1 Parametric bit allocation....151
E.2.4.3.2 Bit allocation tables....152
E.2.4.4 Quantization....153

E.2.4.4.0 Introduction 153
E.2.4.4.1 Vector quantization 153
E.2.4.4.2 Gain adaptive quantization 154
E.2.4.5 Transform equations 156
E.2.5 Enhanced channel coupling 157
E.2.5.1 Overview 157
E.2.5.2 Sub-band structure for enhanced coupling 157
E.2.5.3 Enhanced coupling tables 158
E.2.5.4 Enhanced coupling coordinate format 159
E.2.5.5 Enhanced coupling processing 159
E.2.5.5.0 Introduction 159
E.2.5.5.1 Amplitude parameter processing 160
E.2.5.5.2 Generate channel transform coefficients 160
E.2.6 Spectral extension processing 160
E.2.6.0 Introduction 160
E.2.6.1 Overview 161
E.2.6.2 Sub-band structure for spectral extension 161
E.2.6.3 Spectral extension coordinate format 162
E.2.6.4 High frequency transform coefficient synthesis 163
E.2.6.4.0 Introduction 163
E.2.6.4.1 Transform coefficient translation 163
E.2.6.4.2 Transform coefficient noise blending 163
E.2.6.4.2.0 Introduction 163
E.2.6.4.2.1 Blending factor calculation 164
E.2.6.4.2.2 Banded RMS energy calculation 164
E.2.6.4.2.3 Noise Scaling and Transform Coefficient Blending Calculation 165
E.2.6.4.2.4 Noise scaling and transform coefficient blending calculation 166
E.2.6.4.3 Blended transform coefficient scaling 167
E.2.7 Transient pre-noise processing 167
E.2.7.0 Introduction 167
E.2.7.1 Overview 167
E.2.7.2 Application of transient pre-noise processing data 167
E.2.8 Channel and programme extensions 169
E.2.8.0 Introduction 169
E.2.8.1 Overview 169
E.2.8.2 Decoding a single programme with greater than 5.1 channels 170
E.2.8.3 Decoding multiple programmes with up to 5.1 channels 170
E.2.8.4 Decoding a mixture of programmes with up to 5.1 channels and programmes with greater than 5.1 channels 171
E.2.8.5 Dynamic range compression for programmes containing greater than 5.1 channels 171
E.2.9 LFE downmixing decoder description 171
E.2.10 Control of Programme Mixing 172
E.2.10.1 Overview 172
E.2.10.2 pgmscl 172
E.2.10.3 extpgmscl 172
E.2.10.4 mixdef 173
E.2.10.5 mixdeflen 173
E.2.10.6 mixdata2e 173
E.2.10.7 extpgm(X)scl 173
E.2.10.8 dmixscl 174
E.2.10.9 panmean 174
E.3 AHT vector quantization tables 175

Annex F (normative): AC-3 and Enhanced AC-3 bit streams in the ISO Base Media File Format 192

F.0 Scope 192
F.1 AC-3 and Enhanced AC-3 Track definition 192
F.2 AC-3 and Enhanced AC-3 Sample definition 193
F.3 AC3SampleEntry Box 196

F.3.0 Introduction 196
F.3.1 Syntax....196
F.3.2 Semantics 196
F.4 AC3SpecificBox....196
F.4.1 Syntax....196
F.4.2 Semantics 197
F.4.2.1 BoxHeader.Type - 32 bits....197
F.4.2.2 fscod - 2 bits....197
F.4.2.3 bsid - 5 bits....197
F.4.2.4 bsmod - 3 bits....197
F.4.2.5 acmod - 3 bits....197
F.4.2.6 lfeon - 1 bit....197
F.4.2.7 bit_rate_code - 5 bits....197
F.4.2.8 reserved - 5 bits....198
F.5 EC3SampleEntry Box 198
F.5.1 Syntax....198
F.5.2 Semantics 198
F.6 EC3SpecificBox....198
F.6.1 Syntax....198
F.6.2 Semantics 199
F.6.2.1 BoxHeader.Type - 32 bits....199
F.6.2.2 data_rate - 13 bits....199
F.6.2.3 num_ind_sub - 3 bits....199
F.6.2.4 fscod - 2 bits....199
F.6.2.5 bsid - 5 bits....199
F.6.2.6 reserved - 1 bit 200
F.6.2.7 asvc - 1 bit....200
F.6.2.8 bsmod - 3 bits....200
F.6.2.9 acmod - 3 bits....200
F.6.2.10 lfeon - 1 bit....200
F.6.2.11 reserved - 3 bits....200
F.6.2.12 num_dep_sub - 4 bits....200
F.6.2.13 chan_loc - 9 bits....200
F.6.2.14 reserved - variable....201
Annex G (normative): Enhanced AC-3 bit streams in the MPEG-2 multiplex 202
Annex H (normative): Extensible format for delivery of metadata....203
H.0 Scope 203
H.1 Overview 203
H.2 Container syntax and semantics specification....204
H.2.1 Syntax....204
H.2.1.1 emdf_sync - EMDF container synchronization information....204
H.2.1.2 emdf_container - EMDF container 204
H.2.1.2.0 emdf_container....204
H.2.1.2.1 variable_bits - variable bit-length field format....204
H.2.1.3 emdf_payload_config - EMDF payload configuration 205
H.2.1.4 emdf_protection - EMDF protection data....205
H.2.2 Semantics 205
H.2.2.1 emdf_sync - synchronization information 205
H.2.2.1.1 syncword - synchronization word - 16 bits 205
H.2.2.1.2 emdf_container_length - EMDF container length - 16 bits....205
H.2.2.2 emdf_container - EMDF metadata container structure....206
H.2.2.2.1 variable_bits(n) - variable-bit field encoding....206
H.2.2.2.2 emdf_version - EMDF syntax version - 2 bits / variable_bits(2) 206
H.2.2.2.3 key_id - authentication key identification - 3 bits....206
H.2.2.2.4 emdf_payload_id - EMDF payload identification - 5 bits....207
H.2.2.2.5 emdf_payload_size - size of EMDF payload - variable_bits(8)....207

H.2.2.3 emdf_payload_config - EMDF payload configuration data 207
H.2.2.3.1 smploffste - payload sample offset exists - 1 bit 207
H.2.2.3.2 smploffst - payload sample offset - 11 bits 207
H.2.2.3.3 duratione - payload duration data exists - 1 bit 207
H.2.2.3.4 duration - payload duration data - variable_bits(11) 207
H.2.2.3.5 groupide - payload group ID data exists - 1 bit 207
H.2.2.3.6 groupid - payload group ID - variable_bits(2) 208
H.2.2.3.7 codecdatae - codec-specific data exists - 1 bit 208
H.2.2.3.8 discard_unknown_payload - discard unknown payload during transcode - 1 bit 208
H.2.2.3.9 create_duplicate - create duplicate payload during transcode - 1 bit 208
H.2.2.3.10 remove_duplicate - remove duplicate payload during transcode - 1 bit 208
H.2.2.3.11 payload_frame_aligned - payload to audio data frame alignment - 1 bit 208
H.2.2.3.12 priority - payload priority - 5 bit 208
H.2.2.3.13 proc_allowed - retain payload during transcoding - 2 bits 209
H.2.2.4 emdf_protection - EMDF container protection data 209
H.2.2.4.1 protection_length_primary - length of protection_bits_primary field - 2 bits 209
H.2.2.4.2 protection_length_secondary - length of protection_bits_secondary field - 2 bits 209
H.2.2.4.3 protection_bits_primary - primary EMDF container protection data - 8 to 128 bits 209
H.2.2.4.4 protection_bits_secondary - secondary EMDF container protection data - 0 to 128 bits 210
H.3 EMDF payload types 210
H.3.1 Payload ID 0x1 - programme loudness data 210
H.3.1.1 Overview 210
H.3.1.2 Syntax 210
H.3.1.3 Semantics 211
H.3.1.3.1 version - Version of loudness payload - 2 bits 211
H.3.1.3.2 dialchane - Dialogue presence data exists 211
H.3.1.3.3 dialchan - Dialogue present in channel - 3 bits 212
H.3.1.3.4 loudpractyp - Loudness measurement practice used - 4 bits 212
H.3.1.3.5 loudcorrdialgat - Dialogue-gated loudness correction applied - 1 bit 212
H.3.1.3.6 loudcorrtyp - Type of loudness correction applied to programme - 1 bit 212
H.3.1.3.7 loudrelgate - Relative gated loudness data (ITU) exists - 1 bit 212
H.3.1.3.8 loudrelgat - Relative gated programme loudness (ITU) - 7 bits 212
H.3.1.3.9 loudspchgate - Speech-gated loudness data (ITU) exists - 1 bit 213
H.3.1.3.10 loudspchgat - Speech-gated programme loudness (ITU) - 7 bits 213
H.3.1.3.11 loudstrm3se - Short-term (3 second) loudness data exists - 1 bit 213
H.3.1.3.12 loudstrm3s - Short-term (3 second) programme loudness - 8 bits 213
H.3.1.3.13 truepke - True peak loudness data exists - 1 bit 213
H.3.1.3.14 truepk - True peak value - 8 bits 213
H.3.1.3.15 dmixloudoffste - Downmix loudness offset data exists - 1 bit 213
H.3.1.3.16 dmixloudoffst - Downmix loudness offset - 5 bits 213
H.3.1.3.17 prgmbndye - Programme boundary data exists - 1 bit 214
H.3.1.3.18 prgmbndy_bit - Programme boundary bit value - 1 bit 214
H.3.1.3.19 end_or_start - Programme boundary indicates end or start of programme - 1 bit 214
H.3.1.3.20 prgmbndyoffste - Programme boundary sample offset exists - 1 bit 214
H.3.1.3.21 prgmbndyoffst - Programme boundary sample offset - 11 bits 214
H.3.1.3.22 hrloudrelgate - High resolution relative gated loudness data (ITU) exists - 1 bit 214
H.3.1.3.23 hrloudrelgat - High resolution relative gated programme loudness (ITU) - 3 bits 215
H.3.1.3.24 hrloudspchgate - High resolution speech-gated loudness data (ITU) exists - 1 bit 215
H.3.1.3.25 hrloudspchgat - High resolution speech-gated programme loudness (ITU) - 3 bits 215
H.3.1.3.26 hrloudstrm3se - High resolution short-term (3 second) loudness data exists - 1 bit 215
H.3.1.3.27 hrloudstrm3s - High resolution short-term (3 second) loudness - 3 bits 215
H.3.1.3.28 hrtruepke - High resolution true peak loudness data exists - 1 bit 216
H.3.1.3.29 hrtruepk - High resolution true peak value - 3 bits 216
H.3.1.3.30 loudcorrdialgattyp - Dialog gating type used for loudness correction - 3 bits 216
H.3.1.3.31 extloudrelgate - Extended range relative gated loudness data (ITU) exists - 1 bit 216
H.3.1.3.32 extloudrelgat - Extended range relative gated loudness data (ITU) - 11 bits 217
H.3.1.3.33 extloudspchgate - Extended range speech-gated loudness data (ITU) exists - 1 bit 217
H.3.1.3.34 extloudspchgat - Extended range speech-gated loudness value - 11 bits 217
H.3.1.3.35 loudspchdialgattyp - Dialog gating type used for loudness measurement - 3 bits 217
H.3.1.3.36 extloudstrm3se - Extended range short-term (3 second) programme loudness exists - 1 bit 218
H.3.1.3.37 extloudstrm3s - Extended range short-term (3 second) programme loudness - 11 bits 218

H.3.1.3.38 exttruepkee - Extended true peak value exists - 1 bit 218
H.3.1.3.39 exttruepk - Extended true peak - 11 bits 218
H.3.1.3.40 maxloudstrm3se - Max short-term (3 second) loudness value exists - 1 bit 218
H.3.1.3.41 maxloudstrm3s - Max short-term (3 second) loudness value - 11 bits 218
H.3.1.3.42 maxtruepke - Max true peak value exists - 1 bit 219
H.3.1.3.43 maxtruepk - Max true peak value - 11 bits 219
H.3.1.3.44 lrae - Loudness range data exists - 1 bit 219
H.3.1.3.45 lra - Loudness range - 10 bits 219
H.3.1.3.46 lrapractyp - Loudness range dialog gating practice type - 3 bits 219
H.3.1.3.47 loudmntrye - Momentary loudness data exists - 1 bit 219
H.3.1.3.48 loudmntry - Momentary loudness data - 11 bits 219
H.3.1.3.49 maxloudmntrye - Maximum momentary loudness data exists - 1 bit 220
H.3.1.3.50 maxloudmntry - Maximum momentary loudness data - 11 bits 220
H.3.2 Payload ID 0x2 - programme information 220
H.3.2.1 Overview 220
H.3.2.2 Syntax 220
H.3.2.3 Semantics 221
H.3.2.3.1 version - Version of programme information payload - 2 bits 221
H.3.2.3.2 activechane - Active programme channels data exists - 1 bit 221
H.3.2.3.3 activechan - Active programme channels - 16 bits 221
H.3.2.3.4 dmixtype - Programme has been downmixed - 1 bit 222
H.3.2.3.5 dmixtyp - type of Downmixing applied to programme - 4 bits 222
H.3.2.3.6 upmixtype - Programme has been upmixed - 1 bit 223
H.3.2.3.7 upmixtyp - Type of upmixing applied to programme - 4 bits 223
H.3.2.3.8 preproinfoe - Preprocessing information exists - 1 bit 223
H.3.2.3.9 suratten - Surround attenuation applied - 1 bit 223
H.3.2.3.10 ph90filt - 90 degree phase shift applied - 1 bit 223
H.3.2.3.11 lfefilt - LFE filter applied - 1-bit 224
H.3.2.3.12 lfemonlevcod - LFE channel monitoring level - 2 bits 224
H.3.2.3.13 drcprofinfoe - DRC profile information exists - 1 bit 224
H.3.2.3.14 dynrngprof - DRC profile for dynrng data - 3 bits 224
H.3.2.3.15 comprprof - DRC profile for compr data - 3 bits 225
H.3.2.3.16 specprocinfoe - Spectral processing information exists - 1 bit 225
H.3.2.3.17 specprocstartf - Spectral processing start frequency - 3 bits 225
H.3.2.3.18 specprocendf - Spectral processing end frequency - 3 bits 225
H.3.2.3.19 chancplinfoe - Channel coupling information exists - 1 bit 226
H.3.2.3.20 chancplstartf - Channel coupling start frequency - 5 bits 226
H.3.2.3.21 chancplendf - Channel coupling end frequency - 5 bits 226
H.3.2.3.22 enhncrnge - Dialog enhancement adjustment range data exists - 1 bit 226
H.3.2.3.23 enhncrng - Dialog enhancement adjustment range - 2 bits 227
H.3.3 Payload ID 0x03 - Enhanced AC-3 substream structure 227
H.3.3.1 Overview 227
H.3.3.2 Syntax 227
H.3.3.3 Semantics 227
H.3.3.3.1 num_ind_sub - Number of independent substreams - 3 bits 227
H.3.3.3.2 dep_sub_exist - Dependent substreams exist- 1 bit 227
H.3.3.3.3 num_dep_sub - Number of dependent substreams - 3 bits 227
H.3.4 Payload ID 0x04 - dynamic range compression data for portable devices 228
H.3.4.1 Overview 228
H.3.4.2 Syntax 228
H.3.4.3 Semantics 228
H.3.4.3.1 version - Version of portable device DRC payload - 2 bits 228
H.3.4.3.2 portdrce - Portable device DRC data exists - 1 bit 228
H.3.4.3.3 portdrc - Portable device DRC gain word - 8 bits 228
H.3.4.3.4 drcprofinfoe - DRC profile information exists - 1 bit 228
H.3.4.3.5 portdrcprof - DRC profile for portable device DRC data - 3 bits 228
H.3.5 Payload ID 0x05 - programme language 229
H.3.5.1 Overview 229
H.3.5.2 Syntax 229
H.3.5.3 Semantics 229
H.3.5.3.1 frame_sequence - Syncframe sequence indication - 1 bit 229
H.3.5.3.2 sequence_start - Syncframe sequence start indication - 1 bit 229

H.3.5.3.3 sequence_end - Syncframe sequence end indication - 1 bit 229
H.3.5.3.4 langtag_chunk - Language tag data chunk - 16 bits 230
H.3.5.3.5 langtag_length - Language tag length - 6 bits 230
H.3.5.3.6 langtag - Language_tag - 16 to 336 bits 230
H.3.6 Payload ID 0x6 - external data 230
H.3.6.1 Overview 230
H.3.6.2 Syntax 230
H.3.6.3 Semantics 230
H.3.6.3.1 frame_sequence - Syncframe sequence indication - 1 bit 230
H.3.6.3.2 sequence_start - Syncframe sequence start flag - 1 bit 231
H.3.6.3.3 sequence_end - Syncframe sequence end flag - 1 bit 231
H.3.6.3.4 external_data_chunk - External data chunk - 16 bits 231
H.3.6.3.5 external_data_length - External data field length - variable_bits(7) 231
H.3.6.3.6 external_data - External data field - external_data_length * 8 bits 231
H.3.7 Payload ID 0x7 - headphone rendering data 231
H.3.7.1 Overview 231
H.3.7.2 Syntax 231
H.3.7.3 Semantics 232
H.3.7.3.1 chan_count - Number of full-bandwidth channels - 3 bits 232
H.3.7.3.2 chan_gain - Individual channel gain - 6 bits 232
H.3.7.3.3 lfegaine - LFE channel gain value exists - 1 bit 232
H.3.7.3.4 lfegain - LFE channel gain value - 6 bits 232
H.3.7.3.5 sequence_start - Syncframe sequence start indication - 1 bit 232
H.3.7.3.6 sequence_end - Syncframe sequence end indication - 1 bit 232
H.3.7.3.7 brir_datae - BRIR data exists - 1 bit 233
H.3.7.3.8 chan_brir_chunk - Channel BRIR data chunk - 96 bits 233
H.3.7.3.9 late_BRIR_chunk - Late brir data chunk - 1536 bits 233
H.3.7.3.10 parity_check - 8 bits 233
H.3.7.3.11 propdlye - Propagation delay data exists - 1 bit 233
H.3.7.3.12 propdly - Propagation delay data - 11 bits 233
H.3.7.3.13 rt60e - RT_{60} data exists - 1 bit 233
H.3.7.3.14 num_bnds - Number of bands of RT_{60} data - 4 bits 233
H.3.7.3.15 rt60[bnd] - RT_{60} value - 11 bits 234

Annex I (normative): Signalling in MPEG-DASH 235

I.0 Scope 235

I.1 Media Presentation Description (MPD) 235
I.1.1 General MPD requirements 235
I.1.1.1 Introduction 235
I.1.1.2 Adaptation Sets and Representations 235
I.1.1.3 AudioChannelConfiguration descriptor 235
I.1.2 Descriptors 235
I.1.2.1 AudioChannelConfiguration descriptor 235

I.2 Example Media Presentation Description 237

Annex J (normative): Bit streams in the Common Media Application Format (CMAF) 239

J.1 CMAF Tracks 239
J.1.1 Overview 239
J.1.2 *codecs* parameter signaling 239
J.1.3 Considerations for Audio Encoding 239
J.1.3.1 Overview 239
J.1.3.2 Priming and delay 239
J.1.3.3 Delay compensation using an edit list 240
J.1.3.4 Delay compensation before encapsulation 240
J.1.3.5 Loudness and dynamic range control (DRC) 240
J.1.4 Track constraints 240
J.1.4.1 Storage of ISOBMFF samples 240
J.1.4.2 ISOBMFF sample entry box 240
J.1.5 Elementary stream constraints 240

J.1.5.1 General....240
J.1.5.2 Enhanced AC-3 elementary stream constraints....241
J.1.5.2.1 General....241
J.1.5.2.2 Independent substream 0 constraints....241
J.1.5.2.3 Dependent substream 0 constraints....241
J.2 Switching Sets....241
J.3 Core Audio Media Profile....242
Annex K (informative): Bibliography....243
History....244

Intellectual Property Rights

Essential patents

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

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

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

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

Modal verbs terminology

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

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

Introduction

Motivation

In order to more efficiently broadcast or record audio signals, the amount of information required to represent the audio signals may be reduced. In the case of digital audio signals, the amount of digital information needed to accurately reproduce the original Pulse Code Modulation (PCM) samples may be reduced by applying a digital compression algorithm, resulting in a digitally compressed representation of the original signal (the term compression used in this context means the compression of the amount of digital information which is stored or recorded, and not the compression of dynamic range of the audio signal). The goal of the digital compression algorithm is to produce a digital representation of an audio signal which, when decoded and reproduced, sounds the same as the original signal, while using a minimum of digital information (bit rate) for the compressed (or encoded) representation. The AC-3 digital compression algorithm specified in the present document can encode from 1 to 5.1 channels of source audio from a PCM representation into a serial bit stream at data rates ranging from 32 kbit/s to 640 kbit/s. The ".1" channel refers to a fractional bandwidth channel intended to convey only low frequency effect signals.

A typical application of the algorithm is shown in figure 0.1. In this example, a 5.1-channel audio programme is converted from a PCM representation requiring more than 5 Mbit/s (6 channels x 48 kHz x 18 bits = 5,184 Mbit/s) into a 384 kbit/s serial bit stream by the AC-3 encoder. Satellite transmission equipment converts this bit stream to an RF transmission which is directed to a satellite transponder. The amount of bandwidth and power required by the transmission has been reduced by more than a factor of 13 by the AC-3 digital compression. The signal received from the satellite is demodulated back into the 384 kbit/s serial bit stream, and decoded by the AC-3 decoder. The result is the original 5.1-channel audio programme.

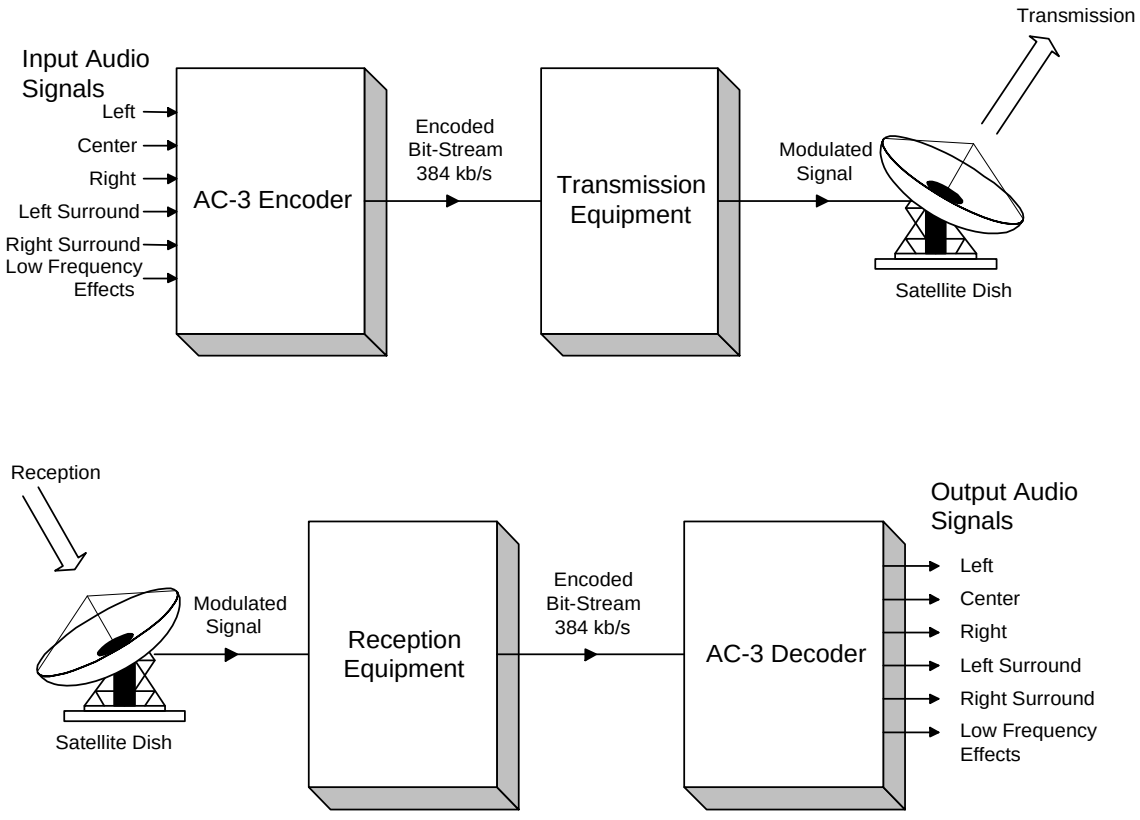

Figure 0.1: Example application of AC-3 to satellite audio transmission

Digital compression of audio is useful wherever there is an economic benefit to be obtained by reducing the amount of digital information required to represent the audio. Typical applications are in satellite or terrestrial audio broadcasting, delivery of audio over metallic or optical cables, or storage of audio on magnetic, optical, semiconductor, or other storage media.

Encoding

The AC-3 encoder accepts PCM audio and produces an encoded bit stream consistent with the present document. The specifics of the audio encoding process are not normative requirements of the present document. The encoder produces a bit stream matching the syntax described in clause 4, which, when decoded according to clauses 5 and 6, produces audio of sufficient quality for the intended application. Clause 7 contains information on the encoding process. The encoding process is briefly described below.

The AC-3 algorithm achieves high coding gain (the ratio of the input bit rate to the output bit rate) by coarsely quantizing a frequency domain representation of the audio signal. A block diagram of this process is shown in figure 0.2. The first step in the encoding process is to transform the representation of audio from a sequence of PCM time samples into a sequence of blocks of frequency coefficients. This is done in the analysis filter bank. Overlapping blocks of 512 time samples are multiplied by a time window and transformed into the frequency domain. Due to the overlapping blocks, each PCM input sample is represented in two sequential transformed blocks. The frequency domain representation may then be decimated by a factor of two so that each block contains 256 frequency coefficients. The individual frequency coefficients are represented in binary exponential notation as a binary exponent and a mantissa. The set of exponents is encoded into a coarse representation of the signal spectrum which is referred to as the spectral envelope. This spectral envelope is used by the core bit allocation routine which determines how many bits to use to encode each individual mantissa. The spectral envelope and the coarsely quantized mantissas for 6 audio blocks (1 536 audio samples) are formatted into an AC-3 syncframe. The AC-3 bit stream is a sequence of AC-3 syncframes.

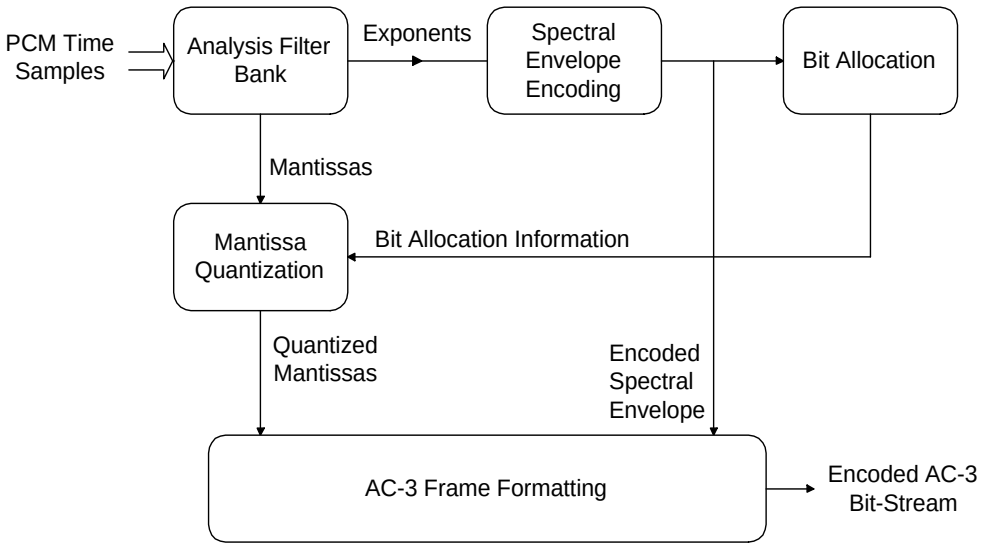

Figure 0.2: The AC-3 encoder

The actual AC-3 encoder is more complex than indicated in figure 0.2. The following functions not shown above are also included:

- a syncframe header is attached which contains information (bit rate, sample rate, number of encoded channels, etc.) required to synchronize to and decode the encoded bit stream;
- error detection codes are inserted in order to allow the decoder to verify that a received syncframe of data is error free;
- the analysis filter bank spectral resolution may be dynamically altered so as to better match the time/frequency characteristic of each audio block;
- the spectral envelope may be encoded with variable time/frequency resolution;
- a more complex bit allocation may be performed, and parameters of the core bit allocation routine modified so as to produce a more optimum bit allocation;
- the channels may be coupled together at high frequencies in order to achieve higher coding gain for operation at lower bit rates;

- in the two-channel mode a rematrixing process may be selectively performed in order to provide additional coding gain, and to allow improved results to be obtained in the event that the two-channel signal is decoded with a matrix surround decoder.

Decoding

The decoding process is basically the inverse of the encoding process. The decoder, shown in figure 0.3, synchronizes to the encoded bit stream, checks for errors, and de-formats the various types of data such as the encoded spectral envelope and the quantized mantissas. The bit allocation routine is run and the results used to unpack and de-quantize the mantissas. The spectral envelope is decoded to produce the exponents. The exponents and mantissas are transformed back into the time domain to produce the decoded PCM time samples.

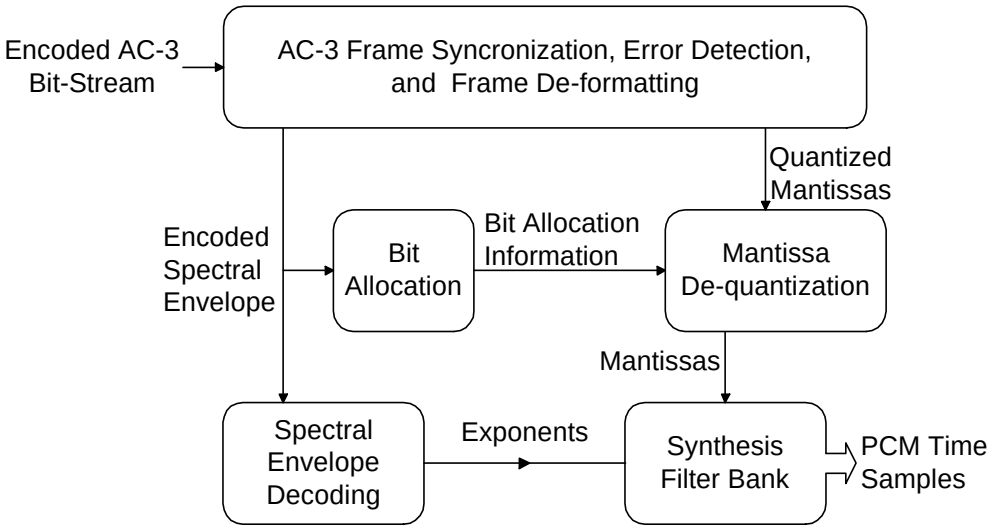

Figure 0.3: The AC-3 decoder

The actual AC-3 decoder is more complex than indicated in figure 0.3. The following functions not shown above are included:

- error concealment or muting may be applied in case a data error is detected;
- channels which have had their high-frequency content coupled together need to be de-coupled;
- dematrixing needs to be applied (in the 2-channel mode) whenever the channels have been rematrixed;
- the synthesis filter bank resolution needs to be dynamically altered in the same manner as the encoder analysis filter bank had been during the encoding process.

1 Scope

The present document specifies two coded representations of audio information, and specifies the decoding process for each coded representation. Informative information on the encoding process is included. The coded representations specified herein are suitable for use in digital audio transmission and storage applications. The coded representations may convey multiple full bandwidth audio signals, along with a low frequency enhancement signal. A wide range of encoded bit-rates is supported by the present document.

A short form designation of the audio coding algorithm specified in the body of the present document (whether or not annex D is included) is "AC-3". The short form designation of the audio coding algorithm specified in annex E is "E-AC-3".

2 References

2.1 Normative references

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

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

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

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

[1] ISO/IEC 23000-19 (2017): "Information technology -- Multimedia application format (MPEG-A) -- Part 19: Common media application format (CMAF) for segmented media".

[2] ISO/IEC 23009-1: "Information technology -- Dynamic adaptive streaming over HTTP (DASH) -- Part 1: Media presentation description and segment formats".

[3] IETF RFC 6381: "The 'Codecs' and 'Profiles' Parameters for "Bucket" Media Types".

[4] ISO/IEC 23001-8: "Information technology -- MPEG systems technologies -- Part 8: Coding-independent code points".

2.2 Informative references

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

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

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

[i.1] ISO 639-1 (2002): "Codes for the representation of names of languages -- Part 1: Alpha-2 code".

[i.2] ISO 639-2 (1998): "Codes for the representation of names of languages -- Part 2: Alpha-3 code".

[i.3] ISO/IEC 8859-1 (1998): "Information technology -- 8-bit single-byte coded graphic character sets -- Part 1: Latin alphabet No. 1".

[i.4] ISO/IEC 13818-1 (2013): "Information technology -- Generic coding of moving pictures and associated audio information: Systems".

[i.5] Recommendation ITU-R BT.1300: "Service multiplex, transport, and identification methods for digital terrestrial television broadcasting".

[i.6] ETSI TS 101 154: "Digital Video Broadcasting (DVB); Specification for the use of Video and Audio Coding in Broadcasting Applications based on the MPEG-2 Transport Stream".

[i.7] ETSI EN 300 468: "Digital Video Broadcasting (DVB); Specification for Service Information (SI) in DVB systems".

[i.8] SMPTE ST 428-3-2006: "D-Cinema Distribution Master - Audio Channel Mapping and Channel Labeling".

[i.9] ISO/IEC 14496-12 (2012): "Information technology -- Coding of audio-visual objects -- Part 12: ISO base media file format".

[i.10] ATSC Standard: "Digital Audio Compression (AC-3, E-AC-3) Standard, Document A/52:2012".

[i.11] Recommendation ITU-R BS.1771-1: "Requirements for loudness and true-peak indicating meters".

[i.12] Recommendation ITU-R BS.1770-4: "Algorithms to measure audio programme loudness and true-peak audio level".

[i.13] ATSC: "ATSC Recommended Practice: Techniques for Establishing and Maintaining Audio Loudness for Digital Television".

[i.14] EBU - Recommendation R 128: "Loudness normalisation and permitted maximum level of audio signals".

[i.15] ARIB TR-B32: "Operational Guidelines for Loudness of Digital Television Programs".

[i.16] FreeTV OP-59: "Measurement and Management of Loudness in Soundtracks for Television Broadcasting".

[i.17] Dolby® Laboratories Speech Gating Reference Code and Information.

NOTE Available at: http://www.dolby.com/in/en/professional/technology/broadcast/dialogue-intelligence.aspx.

[i.18] Recommendation EBU Tech 3341: "Loudness Metering: 'EBU Mode' metering to supplement loudness normalisation in accordance with EBU R 128".

[i.19] Recommendation EBU Tech 3342: "Loudness Range: A Measure to supplement loudness normalization in accordance with EBU R128".

[i.20] IETF BCP-47: "Tags for identifying languages".

NOTE Available at: http://tools.ietf.org/html/bcp47.

3 Definitions and abbreviations

3.1 Definitions

For the purposes of the present document, the following terms and definitions apply:

audio block: set of 512 audio samples consisting of 256 samples of the preceding audio block, and 256 new time samples

NOTE 1: A new audio block occurs every 256 audio samples

NOTE 2: Each audio sample is represented in two audio blocks.

audio frame: portion of an Enhanced AC-3 synchronization frame

NOTE: See syntax for audfrm() in clause E.1.3.2 for the precise definition.

bin: number of the frequency coefficient, as in frequency bin number *n*

NOTE: The 512 point TDAC transform produces 256 frequency coefficients or frequency bins.

coefficient: time domain samples are converted into frequency domain coefficients by the transform

coupled channel: full bandwidth channel whose high frequency information is combined into the coupling channel

coupling band: band of coupling channel transform coefficients covering one or more coupling channel sub-bands

coupling channel: channel formed by combining the high frequency information from the coupled channels

coupling sub-band: sub-band consisting of a group of 12 coupling channel transform coefficients

downmixing: combining (or mixing down) the content of n original channels to produce *m* channels, where $m<n$

exponent set: set of exponents for an independent channel, for the coupling channel, or for the low frequency portion of a coupled channel

frame: generic term used for a portion of a bit stream read in context

NOTE: See syntactical definitions for audio frame and synchronization frame.

full bandwidth (fbw) channel: audio channel capable of full audio bandwidth

NOTE: All channels (left, centre, right, left surround, right surround) except the lfe channel are fbw channels.

independent channel: channel whose high frequency information is not combined into the coupling channel

NOTE: The lfe channel is always independent.

Loudness Units relative to Full Scale (LUFS): loudness units relative to nominal full scale

NOTE: See Recommendation EBU Tech 3341 [i.18].

low frequency effects (lfe) channel: optional single channel of limited (< 120 Hz) bandwidth, which is intended to be reproduced at a level +10 dB with respect to the fbw channels

NOTE: The optional lfe channel allows high sound pressure levels to be provided for low frequency sounds.

spectral envelope: spectral estimate consisting of the set of exponents obtained by decoding the encoded exponents

NOTE: Similar (but not identical) to the original set of exponents

substream: subcomponent of the overall bit stream, specific to Enhanced AC-3, which may be either "dependent" or "independent" as specified by the associated semantics

synchronization frame: minimum portion of the audio serial bit stream capable of being fully decoded, sometimes abbreviated "syncframe"

NOTE: See the syntax for syncframe() (AC-3 synchronization frame) in clause 4.3 and the syntax for syncframe() (E-AC-3 synchronization frame) in clause E.1.2 for the precise definitions.

window: time vector which is multiplied by an audio block to provide a windowed audio block

NOTE: The window shape establishes the frequency selectivity of the filterbank, and provides for the proper overlap/add characteristic to avoid blocking artefacts.

3.2 Abbreviations

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

AC	Audio Codec
acmod	audio coding mode
addbsi	additional bit stream information
addbsie	additional bit stream information exists
addbsil	additional bit stream information length
AHT	Adaptive Hybrid Transform
ARIB	Association of Radio Industries and Businesses
ATSC	Advanced Television Systems Committee
AU	Access Unit
audblk	audio block
audprodi2e	audio production information exists, ch2
audprodie	audio production information exists
auxbits	auxiliary data bits
auxdata	auxiliary data field
auxdatae	auxiliary data exists
auxdatal	auxiliary data length
baie	bit allocation information exists
bap	bit allocation pointer
BCP	Best Current Practice
bin	frequency coefficient bin in index [bin]
blk	block in array index [blk]
blksw	block switch flag
bnd	band in array index [bnd]
BRIR	Binaural Room Impulse Response
bsi	bit stream information
bsid	bit stream identification
bsmod	bit stream mode
ch	channel in array index [ch]
chbwcod	channel bandwidth code
chexpstr	channel exponent strategy
chincpl	channel in coupling
chmant	channel mantissas
CICP	Coding-Independent Code-Points
clev	centre mixing level coefficient
CM	Complete Main
CMAF	Common Media Application Format
cmixlev	centre mix level
compr	compression gain word
compr2	compression gain word, ch2
compr2e	compression gain word exists, ch2
compre	compression gain word exists
copyrightb	copyright bit
cplabsexp	coupling absolute exponent
cplbegf	coupling begin frequency code
cplbndstrc	coupling band structure
cplco	coupling coordinate
cplcoe	coupling coordinates exist
cplcoexp	coupling coordinate exponent
cplcomant	coupling coordinate mantissa
cpldeltba	coupling dba
cpldeltbae	coupling dba exists
cpldeltlen	coupling dba length
cpldeltnseg	coupling dba number of segments
cpldeltoffst	coupling dba offset
cplendf	coupling end frequency code
cplexps	coupling exponents
cplexpstr	coupling exponent strategy

cplfgaincod	coupling fast gain code
cplfleak	coupling fast leak initialization
cplfsnroffst	coupling fine SNR offset
cplinu	coupling in use
cplleake	coupling leak initialization exists
cplmant	coupling mantissas
cplsleak	coupling slow leak initialization
cplstre	coupling strategy exists
CRC	Cyclic Redundancy Check
crc1	crc - cyclic redundancy check word 1
crc2	crc - cyclic redundancy check word 2
crcrsv	crc reserved bit
csnroffst	coarse SNR offset
d15	d15 exponent coding mode
d25	d25 exponent coding mode
d45	d45 exponent coding mode
DAC	Digital to Analogue Converter
DASH	Dynamic Adaptive Streaming over HTTP
dba	delta bit allocation
dbpbcod	dB per bit code
DC	Direct Current
DCT	Discrete Cosine Transform
deltba	channel dba
deltbae	channel dba exists
deltbaie	dba information exists
deltlen	channel dba length
deltnseg	channel dba number of segments
deltoffst	channel dba offset
dialnorm	dialogue normalization word
dialnorm2	dialogue normalization word, ch2
dithflag	dither flag
DRC	Dynamic Range Compression
dsurmod	Dolby® surround mode
DVB	Digital Video Broadcasting
dynrng	dynamic range gain word
dynrng2	dynamic range gain word, ch2
dynrng2e	dynamic range gain word exists, ch2
dynrnge	dynamic range gain word exists
EMDF	Extensible Metadata Delivery Format
exps	channel exponents
fbw	full bandwidth
fdcycod	fast decay code
FFT	Fast Fourier Transform
fgaincod	channel fast gain code
FIR	Finite Impulse Response
floorcod	masking floor code
floortab	masking floor table
frmsizecod	frame size code
fscod	sampling frequency code
fsnroffst	channel fine SNR offset
gainrng	channel gain range code
GAQ	Gain Adaptive Quantization
grp	group in index [grp]
HDCD	High Definition Compatible Digital
HI	Hearing Impaired
HTTP	Hyper Text Transfer Protocol
IDCT	Inverse DCT
IETF	Internet Engineering Task Force
IFFT	Inverse Fast Fourier Transform
IMDCT	Inverse Modified Discrete Transform
ISO	International Organization for Standardization
langcod	language code

langcod2	language code, ch2
langcod2e	language code exists, ch2
langcode	language code exists
lfe	low frequency effects
lfeexps	lfe exponents
lfeexpstr	lfe exponent strategy
lfefgaincod	lfe fast gain code
lfefsnroffst	lfe fine SNR offset
lfemant	lfe mantissas
lfeon	lfe on
LFSR	Linear Feedback Shift Register
LSB	Least Significant Bit
LU	Loudness Unit

NOTE: See Recommendation EBU Tech 3341 [i.18].

LUFS	Loudness Units relative to Full Scale
MDCT	Modified Discrete Cosine Transform
ME	Music and Effects
MIPS	Million Instructions Per Second
mixlevel	mixing level
mixlevel2	mixing level, ch2
MPD	Media Presentation Description
MPEG	Motion Pictures Experts Group
MSB	Most Significant Bit
mstrcplco	master coupling coordinate
NA	Not Applicable
nauxbits	number of auxiliary bits
nchans	number of channels
nchgrps	number of fbw channel exponent groups
nchmant	number of fbw channel mantissas
ncplbnd	number of structured coupled bands
ncplgrps	number of coupled exponent groups
ncplmant	number of coupled mantissas
ncplsubnd	number of coupling sub-bands
nfchans	number of fbw channels
nlfegrps	number of lfe channel exponent groups
nlfemant	number of lfe channel mantissas
origbs	original bit stream
PCM	Pulse Code Modulation
PES	Packetized Elementary Stream
phsflg	phase flag
phsflginu	phase flags in use
PID	Packet Identifier
PMT	Programme Map Table
PSD	Power Spectral Density
PSI	Programme Specific Information
PTS	Presentation Time Stamp
PU	Presentation Unit
RAM	Random Access Memory
rbnd	rematrix band in index [rbnd]
rematflg	rematrix flag
rematstr	rematrixing strategy
RF	Radio Frequency
RFC	Request For Comments
RMS	Root Mean Square
roomtyp	room type
roomtyp2	room type, ch2
SAP	Service Access Point
sbnd	sub-band in index [sbnd]
sdcycod	slow decay code
seg	segment in index [seg]

sgaincod	slow gain code
SI	Synchronization Information
skipfld	skip field
skipl	skip length
skiple	skip length exists
slev	surround mixing level coefficient
SMPTE	Society of Motion Picture & Television Engineers
SNR	Signal to Noise Ratio
snroffste	SNR offset exists
SPL	Sound Pressure Level
STD	System Target Decoder
surmixlev	surround mix level
syncframe	synchronization frame
syncinfo	synchronization information
syncword	synchronization word
TDAC	Time Division Aliasing Cancellation
timecod1	time code first half
timecod1e	time code first half exists
timecod2	time code second half
timecod2e	time code second half exists.
VI	Visually Impaired
VO	Voice Over
VQ	Vector Quantization
XML	eXtensible Markup Language

4 Bit stream syntax

4.1 Synchronization frame

An AC-3 serial coded audio bit stream is made up of a sequence of synchronization frames. Each synchronization frame contains 6 coded audio blocks, each of which represents 256 new audio samples. A Synchronization Information (SI) header at the beginning of each syncframe contains information needed to acquire and maintain synchronization. A bit stream information (BSI) header follows SI, and contains parameters describing the coded audio service. The coded audio blocks may be followed by an auxiliary data (Aux) field. At the end of each syncframe is an error check field that includes a CRC word for error detection. An additional CRC word is located in the SI header, the use of which is optional.

4.2 Semantics of syntax specification

The following pseudo code describes the order of arrival of information within the bit stream. This pseudo code is roughly based on C language syntax, but simplified for ease of reading. For bit stream elements which are larger than 1 bit, the order of the bits in the serial bit stream is either most-significant-bit-first (for numerical values), or left-bit-first (for bit-field values). Fields or elements contained in the bit stream are indicated with **bold** type. Syntactic elements are typographically distinguished by the use of a different font (e.g. dynrng).

Some AC-3 bit stream elements naturally form arrays. This syntax specification treats all bit stream elements individually, whether or not they would naturally be included in arrays. Arrays are thus described as multiple elements (as in blksw[ch] as opposed to simply blksw or blksw[]), and control structures such as *for* loops are employed to increment the index ([ch] for channel in this example).

4.3 Syntax specification

4.3.0 AC-3_bit_stream and syncframe

A continuous audio bit stream consists of a sequence of synchronization frames:

Syntax

```
AC-3_bit_stream()
{
   while(true)
   {
      syncframe() ;
   }
} /* end of AC-3 bit stream */
```

The syncframe consists of the **syncinfo** and **bsi** fields, the 6 coded **audblk** fields, the **auxdata** field, and the **errorcheck** field.

Syntax

```
syncframe()
{
   syncinfo() ;
   bsi() ;
   for(blk = 0; blk < 6; blk++)
   {
      audblk() ;
   }
   auxdata() ;
   errorcheck() ;
} /* end of syncframe */
```

Each of the bit stream elements, and their length, are itemized in the following pseudo code. Note that all bit stream elements arrive most significant bit first, or left bit first, in time.

4.3.1 syncinfo - Synchronization information

Syntax	Word size
syncinfo()	
{	
syncword	16
crc1	16
fscod	2
frmsizecod	6
} /* end of syncinfo */	

4.3.2 bsi - Bit stream information

Syntax	Word size
bsi()	
{	
bsid	5
bsmod	3
acmod	3
if((acmod & 0x1) && (acmod != 0x1)) /* if 3 front channels */ **{cmixlev}**	2
if(acmod & 0x4) /* if a surround channel exists */ **{surmixlev}**	2
if(acmod == 0x2) /* if in 2/0 mode */ **{dsurmod}**	2
lfeon	1
dialnorm	5
compre	1
if(compre) **{compr}**	8
langcode	1
if(langcode) **{langcod}**	8
audprodie	1
if(audprodie)	
{	
mixlevel	5
roomtyp	2

Syntax	Word size
}	
if(acmod == 0) /* if 1+1 mode (dual mono, so some items need a second value) */	
{	
dialnorm2	5
compr2e	1
if(compr2e) **{compr2}**	8
langcod2e	1
if(langcod2e) **{langcod2}**	8
audprodi2e	1
if(audprodi2e)	
{	
mixlevel2	5
roomtyp2	2
}	
}	
copyrightb	1
origbs	1
timecod1e	1
if(timecod1e) **{timecod1}**	14
timecod2e	1
if(timecod2e) **{timecod2}**	14
addbsie	1
if(addbsie)	
{	
addbsil	6
addbsi	(addbsil+1) x 8
}	
} /* end of bsi */	

4.3.3 audblk - Audio block

Syntax	Word size
audblk()	
{	
/* these fields for block switch and dither flags */	
for(ch = 0; ch < nfchans; ch++) **{blksw[ch]}**	1
for(ch = 0; ch < nfchans; ch++) **{dithflag[ch]}**	1
/* these fields for dynamic range control */	
dynrnge	1
if(dynrnge) **{dynrng}**	8
if(acmod == 0) /* if 1+1 mode */	
{	
dynrng2e	1
if(dynrng2e) **{dynrng2}**	8
}	
/* these fields for coupling strategy information */	
cplstre	1
if(cplstre)	
{	
cplinu	1
if(cplinu)	
{	
for(ch = 0; ch < nfchans; ch++) **{chincpl[ch]}**	1
if(acmod == 0x2) **{phsflginu}** /* if in 2/0 mode */	1
cplbegf	4
cplendf	4
/* ncplsubnd = 3 + cplendf - cplbegf */	
for(bnd = 1; bnd < ncplsubnd; bnd++) **{cplbndstrc[bnd]}**	1
}	
}	
/* these fields for coupling coordinates, phase flags */	
if(cplinu)	
{	
for(ch = 0; ch < nfchans; ch++)	
{	
if(chincpl[ch])	
{	
cplcoe[ch]	1
if(cplcoe[ch])	
{	
mstrcplco[ch]	2
/* ncplbnd derived from ncplsubnd, and cplbndstrc */	
for(bnd = 0; bnd < ncplbnd; bnd++)	
{	

Syntax	Word size
cplcoexp[ch][bnd]	4
cplcomant[ch][bnd]	4
}	
}	
}	
}	
if((acmod == 0x2) && phsflginu && (cplcoe[0] || cplcoe[1]))	
{	
for(bnd = 0; bnd < ncplbnd; bnd++) **{phsflg[bnd]}**	1
}	
}	
/* these fields for rematrixing operation in the 2/0 mode */	
if(acmod == 0x2) /* if in 2/0 mode */	
{	
rematstr	1
if(rematstr)	
{	
if((cplbegf > 2) || (cplinu == 0))	
{	
for(rbnd = 0; rbnd < 4; rbnd++) **{rematflg[rbnd]}**	1
}	
if((2 >= cplbegf > 0) && cplinu)	
{	
for(rbnd = 0; rbnd < 3; rbnd++) **{rematflg[rbnd]}**	1
}	
if((cplbegf == 0) && cplinu)	
{	
for(rbnd = 0; rbnd < 2; rbnd++) **{rematflg[rbnd]}**	1
}	
}	
}	
/* these fields for exponent strategy */	
if(cplinu) **{cplexpstr}**	2
for(ch = 0; ch < nfchans; ch++) **{chexpstr[ch]}**	2
if(lfeon) **{lfeexpstr}**	1
for(ch = 0; ch < nfchans; ch++)	
{	
if(chexpstr[ch] != reuse)	
{	
if(!chincpl[ch]) **{chbwcod[ch]}**	6
}	
}	
/* these fields for exponents */	
if(cplinu) /* exponents for the coupling channel */	
{	
if(cplexpstr != reuse)	
{	
cplabsexp	4
/* ncplgrps derived from ncplsubnd, cplexpstr */	
for(grp = 0; grp < ncplgrps; grp++) **{cplexps[grp]}**	7
}	
}	
for(ch = 0; ch < nfchans; ch++) /* exponents for full bandwidth channels */	
{	
if(chexpstr[ch] != reuse)	
{	
exps[ch][0]	4
/* nchgrps derived from chexpstr[ch], and cplbegf or chbwcod[ch] */	
for(grp = 1; grp <= nchgrps[ch]; grp++) **{exps[ch][grp]}**	7
gainrng[ch]	2
}	
}	
if(lfeon) /* exponents for the low frequency effects channel */	
{	
if(lfeexpstr != reuse)	
{	
lfeexps[0]	4
/* nlfegrps = 2 */	
for(grp = 1; grp <= nlfegrps; grp++) **{lfeexps[grp]}**	7
}	
}	
/* these fields for bit-allocation parametric information */	
baie	1
if(baie)	
{	
sdcycod	2
fdcycod	2

Syntax	Word size
sgaincod	2
dbpbcod	2
floorcod	3
}	
snroffste	1
if(snroffste)	
{	
csnroffst	6
if(cplinu)	
{	
cplfsnroffst	4
cplfgaincod	3
}	
for(ch = 0; ch < nfchans; ch++)	
{	
fsnroffst[ch]	4
fgaincod[ch]	3
}	
if(lfeon)	
{	
lfefsnroffst	4
lfefgaincod	3
}	
}	
if(cplinu)	
{	
cplleake	1
if(cplleake)	
{	
cplfleak	3
cplsleak	3
}	
}	
/* these fields for delta bit allocation information */	
deltbaie	1
if(deltbaie)	
{	
if(cplinu) **{cpldeltbae}**	2
for(ch = 0; ch < nfchans; ch++) **{deltbae[ch]}**	2
if(cplinu)	
{	
if(cpldeltbae==new info follows)	
{	
cpldeltnseg	3
for(seg = 0; seg <= cpldeltnseg; seg++)	
{	
cpldeltoffst[seg]	5
cpldeltlen[seg]	4
cpldeltba[seg]	3
}	
}	
}	
for(ch = 0; ch < nfchans; ch++)	
{	
if(deltbae[ch]==new info follows)	
{	
deltnseg[ch]	3
for(seg = 0; seg <= deltnseg[ch]; seg++)	
{	
deltoffst[ch][seg]	5
deltlen[ch][seg]	4
deltba[ch][seg]	3
}	
}	
}	
}	
/* these fields for inclusion of unused dummy data */	
skiple	1
if(skiple)	
{	
skipl	9
skipfld	skipl x 8
}	
/* These fields for quantized mantissa values */	
got_cplchan = 0	
for (ch = 0; ch < nfchans; ch++)	
{	

Syntax	Word size
`for (bin = 0; bin < nchmant[ch]; bin++) {chmant[ch][bin]}`	(0-16)
`if (cplinu && chincpl[ch] && !got_cplchan)`	
`{`	
`for (bin = 0; bin < ncplmant; bin++) {cplmant[bin]}`	(0-16)
`got_cplchan = 1`	
`}`	
`}`	
`if(lfeon) /* mantissas of low frequency effects channel */`	
`{`	
`for (bin = 0; bin < nlfemant; bin++) {lfemant[bin]}`	(0-16)
`}`	
`} /* end of audblk */`	

4.3.4 auxdata - Auxiliary data

Syntax	Word size
`auxdata()`	
`{`	
`auxbits`	nauxbits
`if(auxdatae)`	
`{`	
`auxdatal`	14
`}`	
`auxdatae`	1
`} /* end of auxdata */`	

4.3.5 errorcheck - Error detection code

Syntax	Word size
`errorcheck()`	
`{`	
`crcrsv`	1
`crc2`	16
`} /* end of errorcheck */`	

4.4 Description of bit stream elements

4.4.0 Introduction

A number of bit stream elements have values which may be transmitted, but whose meaning has been reserved. If a decoder receives a bit stream which contains reserved values, the decoder may or may not be able to decode and produce audio. In the description of bit stream elements which have reserved codes, there is an indication of what the decoder can do if the reserved code is received. In some cases, the decoder can not decode audio. In other cases, the decoder can still decode audio by using a default value for a parameter which was indicated by a reserved code.

4.4.1 syncinfo - Synchronization information

4.4.1.1 syncword - Synchronization word - 16 bits

The syncword is always 0x0B77, or 0000 1011 0111 0111. Transmission of the syncword, like other bit field elements, is left bit first.

4.4.1.2 crc1 - Cyclic redundancy check 1 bit to 16 bits

This 16 bit-CRC applies to the first 5/8 of the syncframe. Transmission of the CRC, like other numerical values, is most significant bit first.

4.4.1.3 fscod - Sample rate code - 2 bits

This is a 2-bit code indicating sample rate according to Table 4.1. If the reserved code is indicated, the decoder should not attempt to decode audio and should mute.

Table 4.1: Sample rate codes

fscod	Sample rate (kHz)
00	48
01	44,1
10	32
11	Reserved

4.4.1.4 frmsizecod - Frame size code - 6 bits

The frame size code is used along with the sample rate code to determine the number of (2-byte) words before the next syncword (see Table 4.13).

4.4.2 bsi - Bit stream information

4.4.2.1 bsid - Bit stream identification - 5 bits

This bit field has a value of 01000 (= 8) in this version of the present document. Future modifications of the present document may define other values. Values of bsid smaller than 8 will be used for versions of AC-3 which are backward compatible with version 8 decoders. Decoders which can decode version 8 will thus be able to decode bsid version numbers less than 8. If the present document is extended by the addition of additional elements or features that are not compatible with decoders that follow this bsid version 8 specification, a value of bsid greater than 8 will be used. Decoders built to this version of the standard will not be able to decode versions with bsid greater than 8. Thus, decoders built to the present document shall mute if the value of bsid is greater than 8, and should decode and reproduce audio if the value of bsid is less than or equal to 8.

4.4.2.2 bsmod - Bit stream mode - 3 bits

This 3-bit code indicates the type of service that the bit stream conveys as defined in Table 4.2.

Table 4.2: Bit stream mode

bsmod	acmod	Type of service
000	Any	Main audio service: complete main (CM)
001	Any	Main audio service: music and effects (ME)
010	Any	Associated service: visually impaired (VI)
011	Any	Associated service: hearing impaired (HI)
100	Any	Associated service: dialogue (D)
101	Any	Associated service: commentary (C)
110	Any	Associated service: emergency (E)
111	001	Associated service: voice over (VO)
111	010 to 111	Main audio service: karaoke

4.4.2.3 acmod - Audio coding mode - 3 bits

This 3-bit code, shown in Table 4.3, indicates which of the main service channels are in use, ranging from 3/2 to 1/0. If the MSB of acmod is a 1, surround channels are in use and surmixlev follows in the bit stream. If the MSB of acmod is a 0, the surround channels are not in use and surmixlev does not follow in the bit stream. If the LSB of acmod is a 0, the centre channel is not in use. If the LSB of acmod is a 1, the centre channel is in use.

NOTE: The state of acmod sets the number of full-bandwidth channels parameter, nfchans, (e.g. for 3/2 mode, nfchans = 5; for 2/1 mode, nfchans = 3; etc.). The total number of channels, nchans, is equal to nfchans if the lfe channel is off, and is equal to 1 + nfchans if the lfe channel is on. If acmod is 0, then two completely independent programme channels (dual mono) are encoded into the bit stream, and are referenced as Ch1, Ch2. In this case, a number of additional items are present in BSI or audblk to fully describe Ch2. Table 4.3 also indicates the channel ordering (the order in which the channels are processed) for each of the modes.

Table 4.3: Audio coding mode

acmod	Audio coding mode	Nfchans	Channel array ordering
000	1 + 1	2	Ch1, Ch2
001	1/0	1	C
010	2/0	2	L, R
011	3/0	3	L, C, R
100	2/1	3	L, R, S
101	3/1	4	L, C, R, S
110	2/2	4	L, R, Ls, Rs
111	3/2	5	L, C, R, Ls, Rs

4.4.2.4 cmixlev - Centre mix level - 2 bits

When three front channels are in use, this 2-bit code, shown in Table 4.4, indicates the nominal down mix level of the centre channel with respect to the left and right channels. If cmixlev is set to the reserved code, decoders should still reproduce audio. The intermediate value of cmixlev (-4,5 dB) may be used in this case.

Table 4.4: Centre mix level

cmixlev	clev
00	0,707 (-3,0 dB)
01	0,595 (-4,5 dB)
10	0,500 (-6,0 dB)
11	Reserved

4.4.2.5 surmixlev - Surround mix level - 2 bits

If surround channels are in use, this 2-bit code, shown in Table 4.5, indicates the nominal down mix level of the surround channels. If surmixlev is set to the reserved code, the decoder should still reproduce audio. The intermediate value of surmixlev (-6 dB) may be used in this case.

Table 4.5: Surround mix level

surmixlev	slev
00	0,707 (-3 dB)
01	0,500 (-6 dB)
10	0
11	Reserved

4.4.2.6 dsurmod - Dolby® Surround mode - 2 bits

When operating in the two channel mode, this 2-bit code, as shown in Table 4.6, indicates whether or not the programme has been encoded in Dolby® Surround. This information is not used by the AC-3 decoder, but may be used by other portions of the audio reproduction equipment. If dsurmod is set to the reserved code, the decoder should still reproduce audio. The reserved code may be interpreted as "not indicated".

NOTE: "Dolby®", "Pro Logic®", "Surround EX™" and the double -D symbol are trademarks of Dolby® Laboratories. This information is given for the convenience of users of the present document and does not constitute an endorsement by ETSI of the product named. Equivalent products may be used if they can be shown to lead to the same results."

Table 4.6: Dolby® Surround mode

dsurmod	Indication
00	Not indicated
01	NOT Dolby® Surround encoded
10	Dolby® Surround encoded
11	Reserved

4.4.2.7 lfeon - Low frequency effects channel on - 1 bit

This bit has a value of 1 if the lfe (sub woofer) channel is on, and a value of 0 if the lfe channel is off.

4.4.2.8 dialnorm - Dialogue normalization - 5 bits

This 5-bit code indicates how far the average dialogue level is below digital 100 percent. Valid values are 1 to 31. The value of 0 is reserved. The values of 1 to 31 are interpreted as -1 dB to -31 dB with respect to digital 100 %. If the reserved value of 0 is received, the decoder shall use -31 dB. The value of dialnorm shall affect the sound reproduction level. If the value is not used by the AC-3 decoder itself, the value shall be used by other parts of the audio reproduction equipment. Dialogue normalization is further explained in clause 6.6.

4.4.2.9 compre - Compression gain word exists - 1 bit

If this bit is a 1, the following 8 bits represent a compression control word.

4.4.2.10 compr - Compression gain word - 8 bits

This encoder generated gain word may be present in the bit stream. If so, it may be used to scale the reproduced audio level in order to reproduce a very narrow dynamic range, with an assured upper limit of instantaneous peak reproduced signal level in the monophonic downmix. The meaning and use of compr is described further in clause 6.7.3.

4.4.2.11 langcode - Language code exists - 1 bit

If this bit is a 1, the following 8 bits (i.e. the element langcod) shall be reserved. If this bit is a 0, the element langcod does not exist in the bit stream.

4.4.2.12 langcod - Language code - 8 bits

This is an 8 bit reserved value. (This element was originally intended to carry an 8-bit value that would, via a table lookup, indicate the language of the audio programme. Because modern delivery systems provide the ISO 639-2 [i.2] language code in the multiplexing layer, indication of language within the AC-3 bit stream was unnecessary, and so was removed from the AC-3 syntax.)

4.4.2.13 audprodie - Audio production information exists - 1 bit

If this bit is a 1, the mixlevel and roomtyp fields exist, indicating information about the audio production environment (mixing room).

4.4.2.14 mixlevel - Mixing level - 5 bits

This 5-bit code indicates the absolute acoustic sound pressure level of an individual channel during the final audio mixing session. The 5-bit code represents a value in the range 0 to 31. The peak mixing level is 80 plus the value of mixlevel dB SPL, or 80 dB to 111 dB SPL. The peak mixing level is the acoustic level of a sine wave in a single channel whose peaks reach 100 percent in the PCM representation. The absolute SPL value is typically measured by means of pink noise with an RMS value of -20 dB or -30 dB with respect to the peak RMS sine wave level. The value of mixlevel is not typically used within the AC-3 decoder, but may be used by other parts of the audio reproduction equipment.

4.4.2.15 roomtyp - Room type - 2 bits

This 2-bit code, shown in Table 4.7, indicates the type and calibration of the mixing room used for the final audio mixing session. The value of roomtyp is not typically used by the AC-3 decoder, but may be used by other parts of the audio reproduction equipment. If roomtyp is set to the reserved code, the decoder should still reproduce audio. The reserved code may be interpreted as "not indicated".

Table 4.7: Room type

roomtyp	Type of mixing room
00	Not indicated
01	Large room, X curve monitor
10	Small room, flat monitor
11	Reserved

4.4.2.16 dialnorm2 - Dialogue normalization, Ch2 - 5 bits

This 5-bit code has the same meaning as dialnorm, except that it applies to the second audio channel when acmod indicates two independent channels (dual mono 1 + 1 mode).

4.4.2.17 compr2e - Compression gain word exists, Ch2 - 1 bit

If this bit is a 1, the following 8 bits represent a compression gain word for Ch2.

4.4.2.18 compr2 - Compression gain word, Ch2 - 8 bits

This 8-bit word has the same meaning as compr, except that it applies to the second audio channel when acmod indicates two independent channels (dual mono 1 + 1 mode).

4.4.2.19 langcod2e - Language code exists, Ch2 - 1 bit

If this bit is a 1, the following 8 bits (i.e. the element langcod2) shall be reserved. If this bit is a 0, the element langcod2 does not exist in the bit stream.

4.4.2.20 langcod2 - Language code, Ch2 - 8 bits

This is an 8 bit reserved value. See langcod, clause 4.4.2.12.

4.4.2.21 audprodi2e - Audio production information exists, Ch2 - 1 bit

If this bit is a 1, the following two data fields exist indicating information about the audio production for Ch2.

4.4.2.22 mixlevel2 - Mixing level, Ch2 - 5 bits

This 5-bit code has the same meaning as mixlevel, except that it applies to the second audio channel when acmod indicates two independent channels (dual mono 1 + 1 mode).

4.4.2.23 roomtyp2 - Room type, Ch2 - 2 bits

This 2-bit code has the same meaning as roomtyp, except that it applies to the second audio channel when acmod indicates two independent channels (dual mono 1 + 1 mode).

4.4.2.24 copyrightb - Copyright bit - 1 bit

If this bit has a value of 1, the information in the bit stream is indicated as protected by copyright. It has a value of 0 if the information is not indicated as protected.

4.4.2.25 origbs - Original bit stream - 1 bit

This bit has a value of 1 if this is an original bit stream. This bit has a value of 0 if this is a copy of another bit stream.

4.4.2.26 timecod1e, timecod2e - Time code (first and second) halves exists - 2 bits

These values indicate, as shown in Table 4.8, whether time codes follow in the bit stream. The time code can have a resolution of $1/64^{th}$ of a frame (one frame = $1/30^{th}$ of a second). Since only the high resolution portion of the time code is needed for fine synchronization, the 28 bit time code is broken into two 14 bit halves. The low resolution first half represents the code in 8 second increments up to 24 hours. The high resolution second half represents the code in $1/64^{th}$ frame increments up to 8 seconds.

4.4.2.27 timecod1 - Time code first half - 14 bits

The first 5 bits of this 14 bit field represent the time in hours, with valid values of 0 to 23. The next 6 bits represent the time in minutes, with valid values of 0 to 59. The final 3 bits represents the time in 8 second increments, with valid values of 0 - 7 (representing 0, 8, 16, ..., 56 seconds).

Table 4.8: Time code exists

timecod2e, timecod1e	Time code present
0,0	Not present
0,1	First half (14 bits) present
1,0	Second half (14 bits) present
1,1	Both halves (28 bits) present

4.4.2.28 timecod2 - Time code second half - 14 bits

The first 3 bits of this 14-bit field represent the time in seconds, with valid values from 0 to 7 (representing 0 to 7 seconds). The next 5 bits represents the time in frames, with valid values from 0 to 29. The final 6 bits represents fractions of $1/64^{th}$ of a frame, with valid values from 0 to 63.

4.4.2.29 addbsie - Additional bit stream information exists - 1 bit

If this bit has a value of 1 there is additional bit stream information, the length of which is indicated by the next field. If this bit has a value of 0, there is no additional bit stream information.

4.4.2.30 addbsil - Additional bit stream information length - 6 bits

This 6-bit code, which exists only if addbsie is a 1, indicates the length in bytes of additional bit stream information. The valid range of addbsil is 0 to 63, indicating 1 to 64 additional bytes, respectively. The decoder is not required to interpret this information, and thus shall skip over this number of bytes following in the data stream.

4.4.2.31 addbsi - Additional bit stream information - ((addbsil + 1) x 8) bits

This field contains 1 to 64 bytes of any additional information included with the bit stream information structure.

4.4.3 audblk - Audio block

4.4.3.1 blksw[ch] - Block switch flag - 1 bit

This flag, for channel [ch], indicates whether the current audio block was split into 2 sub-blocks during the transformation from the time domain into the frequency domain. A value of 0 indicates that the block was not split, and that a single 512 point TDAC transform was performed. A value of 1 indicates that the block was split into 2 sub-blocks of length 256, that the TDAC transform length was switched from a length of 512 points to a length of 256 points, and that 2 transforms were performed on the audio block (one on each sub-block). Transform length switching is described in more detail in clause 6.9.

4.4.3.2 dithflag[ch] - Dither flag - 1 bit

This flag, for channel [ch], indicates that the decoder should activate dither during the current block. Dither is described in detail in clause 6.3.4.

4.4.3.3 dynrnge - Dynamic range gain word exists - 1 bit

If this bit is a 1, the dynamic range gain word follows in the bit stream. If it is 0, the gain word is not present, and the previous value is reused, except for block 0 of a syncframe where if the control word is not present the current value of dynrng is set to 0.

4.4.3.4 dynrng - Dynamic range gain word - 8 bits

This encoder-generated gain word is applied to scale the reproduced audio as described in clause 6.7.

4.4.3.5 dynrng2e - Dynamic range gain word exists, Ch2 - 1 bit

If this bit is a 1, the dynamic range gain word for channel 2 follows in the bit stream. If it is 0, the gain word is not present, and the previous value is reused, except for block 0 of a syncframe where if the control word is not present the current value of dynrng2 is set to 0.

4.4.3.6 dynrng2 - dynamic range gain word, Ch2 - 8 bits

This encoder-generated gain word is applied to scale the reproduced audio of Ch2, in the same manner as dynrng is applied to Ch1, as described in clause 6.7.

4.4.3.7 cplstre - Coupling strategy exists - 1 bit

If this bit is a 1, coupling information follows in the bit stream. If it is 0, new coupling information is not present, and coupling parameters previously sent are reused. This parameter shall not be set to 0 in block 0.

4.4.3.8 cplinu - Coupling in use - 1 bit

If this bit is a 1, coupling is currently being utilized, and coupling parameters follow. If it is 0, coupling is not being utilized (all channels are independent) and no coupling parameters follow in the bit stream.

4.4.3.9 chincpl[ch] - Channel in coupling - 1 bit

If this bit is a 1, then the channel indicated by the index [ch] is a coupled channel. If the bit is a 0, then this channel is not coupled. Since coupling is not used in the 1/0 mode, if any chincpl[] values exist there will be 2 to 5 values. Of the values present, at least two values will be 1, since coupling requires more than one coupled channel to be coupled.

4.4.3.10 phsflginu - Phase flags in use - 1 bit

If this bit (defined for 2/0 mode only) is a 1, phase flags are included with coupling coordinate information. Phase flags are described in clause 6.4.

4.4.3.11 cplbegf - Coupling begin frequency code - 4 bits

This 4-bit code is interpreted as the sub-band number (0 to 15) which indicates the lower frequency band edge of the coupling channel (or the first active sub-band) as shown in Table 6.24.

4.4.3.12 cplendf - Coupling end frequency code - 4 bits

This 4-bit code indicates the upper band edge of the coupling channel. The upper band edge (or last active sub-band) is cplendf + 2, or a value between 2 and 17 (see Table 6.24).

The number of active coupling sub-bands is equal to ncplsubnd, which is calculated as:

ncplsubnd = 3 + cplendf - cplbegf.

4.4.3.13 cplbndstrc[sbnd] - Coupling band structure - 1 bit

There are 18 coupling sub-bands defined in Table 6.24, each containing 12 frequency coefficients. The fixed 12-bin wide coupling sub-bands are converted into coupling bands, each of which may be wider than (a multiple of) 12 frequency bins. Each coupling band may contain one or more coupling sub-bands. Coupling coordinates are transmitted for each coupling band. Each band's coupling coordinate shall be applied to all the coefficients in the coupling band.

The coupling band structure indicates which coupling sub-bands are combined into wider coupling bands. When cplbndstrc[sbnd] is a 0, the sub-band number [sbnd] is not combined into the previous band to form a wider band, but starts a new 12 wide coupling band. When cplbndstrc[sbnd] is a 1, then the sub-band [sbnd] is combined with the previous band, making the previous band 12 bins wider. Each successive value of cplbndstrc which is a 1 will continue to combine sub-bands into the current band. When another cplbndstrc value of 0 is received, then a new band will be formed, beginning with the 12 bins of the current sub-band. The set of cplbndstrc[sbnd] values is typically considered an array.

Each bit in the array corresponds to a specific coupling sub-band in ascending frequency order. The first element of the array corresponds to the sub-band cplbegf, is always 0, and is not transmitted. (There is no reason to send a cplbndstrc bit for the first sub-band at cplbegf, since this bit would always be 0.) Thus, there are ncplsubnd-1 values of cplbndstrc transmitted. If there is only one coupling sub-band, then no cplbndstrc bits are sent.

The number of coupling bands, ncplbnd, may be computed from ncplsubnd and cplbndstrc:

ncplbnd = (ncplsubnd - (cplbndstrc[1] + ... + cplbndstrc[ncplsubnd - 1])).

4.4.3.14 cplcoe[ch] - Coupling coordinates exist - 1 bit

Coupling coordinates indicate, for a given channel and within a given coupling band, the fraction of the coupling channel frequency coefficients to use to re-create the individual channel frequency coefficients. Coupling coordinates are conditionally transmitted in the bit stream. If new values are not delivered, the previously sent values remain in effect. See clause 6.4 for further information on coupling.

If cplcoe[ch] is 1, the coupling coordinates for the corresponding channel [ch] exist and follow in the bit stream. If the bit is 0, the previously transmitted coupling coordinates for this channel are reused. This parameter shall not be set to 0 in block 0, or in any block for which the corresponding channel is participating in coupling but was not participating in coupling in the previous block.

4.4.3.15 mstrcplco[ch] - Master coupling coordinate - 2 bits

This per channel parameter establishes a per channel gain factor (increasing the dynamic range) for the coupling coordinates as shown in Table 4.9.

Table 4.9: Master coupling coordinate

mstrcplco[ch]	cplco[ch][bnd] gain multiplier
00	1
01	2^{-3}
10	2^{-6}
11	2^{-9}

4.4.3.16 cplcoexp[ch][bnd] - Coupling coordinate exponent - 4 bits

Each coupling coordinate is composed of a 4-bit exponent and a 4-bit mantissa. This element is the value of the coupling coordinate exponent for channel [ch] and band [bnd]. The index [ch] only will exist for those channels which are coupled. The index [bnd] will range from 0 to ncplbnds. See clause 6.4.3 for further information on how to interpret coupling coordinates.

4.4.3.17 cplcomant[ch][bnd] - Coupling coordinate mantissa - 4 bits

This element is the 4-bit coupling coordinate mantissa for channel [ch] and band [bnd].

4.4.3.18 phsflg[bnd] - Phase flag - 1 bit

This element (only used in the 2/0 mode) indicates whether the decoder should phase invert the coupling channel mantissas when reconstructing the right output channel. The index [bnd] can range from 0 to ncplbnd. Phase flags are described in clause 6.4.

4.4.3.19 rematstr - Rematrixing strategy - 1 bit

If this bit is a 1, then new rematrix flags are present in the bit stream. If it is 0, rematrix flags are not present, and the previous values should be reused. The rematstr parameter is present only in the 2/0 audio coding mode. This parameter shall not be set to 0 in block 0.

4.4.3.20 rematflg[rbnd] - Rematrix flag - 1 bit

This bit indicates whether the transform coefficients in rematrixing band [rbnd] have been rematrixed. If this bit is a 1, then the transform coefficients in [rbnd] were rematrixed into sum and difference channels. If this bit is a 0, then rematrixing has not been performed in band [rbnd]. The number of rematrixing bands (and the number of values of [rbnd]) depend on coupling parameters as shown in Table 4.10. Rematrixing is described in clause 6.5.

Table 4.10: Number of rematrixing bands

Condition	No. of rematrixing bands
cplinu == 0	4
(cplinu == 1) && (cplbegf > 2)	4
(cplinu == 1) && (2 ≥ cplbegf > 0)	3
(cplinu == 1) && (cplbegf == 0)	2

4.4.3.21 cplexpstr - Coupling exponent strategy - 2 bits

This element indicates the method of exponent coding that is used for the coupling channel as shown in Table 6.4. See clause 6.1 for explanation of each exponent strategy. This parameter shall not be set to 0 in block 0, or in any block for which coupling is enabled but was disabled in the previous block.

4.4.3.22 chexpstr[ch] - Channel exponent strategy - 2 bits

This element indicates the method of exponent coding that is used for channel [ch], as shown in Table 6.4. This element exists for each full bandwidth channel. This parameter shall not be set to 0 in block 0.

4.4.3.23 lfeexpstr - Low frequency effects channel exponent strategy - 1 bit

This element indicates the method of exponent coding that is used for the lfe channel, as shown in Table 6.5. This parameter shall not be set to 0 in block 0.

4.4.3.24 chbwcod[ch] - Channel bandwidth code - 6 bits

The chbwcod[ch] element is an unsigned integer which defines the upper band edge for full-bandwidth channel [ch]. This parameter is only included for fbw channels which are not coupled. (See clause 6.1.3 on exponents for the definition of this parameter.) Valid values are in the range of 0 - 60. If a value greater than 60 is received, the bit stream is invalid and the decoder shall cease decoding audio and mute.

4.4.3.25 cplabsexp - Coupling absolute exponent - 4 bits

This is an absolute exponent, which is used as a reference when decoding the differential exponents for the coupling channel.

4.4.3.26 cplexps[grp] - Coupling exponents - 7 bits

Each value of cplexps indicates the value of 3, 6, or 12 differentially-coded coupling channel exponents for the coupling exponent group [grp] for the case of d15, d25, or d45 coding, respectively. The number of cplexps values transmitted equals ncplgrps, which may be determined from cplbegf, cplendf, and cplexpstr. Refer to clause 6.1.3 for further information.

4.4.3.27 exps[ch][grp] - Channel exponents - 4 bits or 7 bits

These elements represent the encoded exponents for channel [ch]. The first element ([grp] = 0) is a 4-bit absolute exponent for the first (DC term) transform coefficient. The subsequent elements ([grp] > 0) are 7-bit representations of a group of 3, 6, or 12 differentially coded exponents (corresponding to d15, d25, d45 exponent strategies respectively). The number of groups for each channel, nchgrps[ch], is determined from cplbegf if the channel is coupled, or chbwcod[ch] if the channel is not coupled. Refer to clause 6.1.3 for further information.

4.4.3.28 gainrng[ch] - Channel gain range code - 2 bits

This per channel 2-bit element may be used to determine a block floating-point shift value for the inverse TDAC transform filter bank. Use of this code allows increased dynamic range to be obtained from a limited word length transform computation. For further information see clause 6.9.5.

4.4.3.29 lfeexps[grp] - Low frequency effects channel exponents - 4 bits or 7 bits

These elements represent the encoded exponents for the lfe channel. The first element ([grp] = 0) is a 4-bit absolute exponent for the first (DC term) transform coefficient. There are two additional elements (nlfegrps = 2) which are 7-bit representations of a group of 3 differentially coded exponents. The total number of lfe channel exponents (nlfemant) is 7.

4.4.3.30 baie - Bit allocation information exists - 1 bit

If this bit is a 1, then five separate fields (totalling 11 bits) follow in the bit stream. Each field indicates parameter values for the bit allocation process. If this bit is a 0, these fields do not exist. Further details on these fields may be found in clause 6.2. This parameter shall not be set to 0 in block 0.

4.4.3.31 sdcycod - Slow decay code - 2 bits

This 2-bit code specifies the slow decay parameter in the bit allocation process.

4.4.3.32 fdcycod - Fast decay code - 2 bits

This 2-bit code specifies the fast decay parameter in the decode bit allocation process.

4.4.3.33 sgaincod - Slow gain code - 2 bits

This 2-bit code specifies the slow gain parameter in the decode bit allocation process.

4.4.3.34 dbpbcod - dB per bit code - 2 bits

This 2-bit code specifies the dB per bit parameter in the bit allocation process.

4.4.3.35 floorcod - Masking floor code - 3 bits

This 3-bit code specifies the floor code parameter in the bit allocation process.

4.4.3.36 snroffste - SNR offset exists - 1 bit

If this bit has a value of 1, a number of bit allocation parameters follow in the bit stream. If this bit has a value of 0, SNR offset information does not follow, and the previously transmitted values should be used for this block. The bit allocation process and these parameters are described in clause 6.2. This parameter shall not be set to 0 in block 0.

4.4.3.37 csnroffst - Coarse SNR offset - 6 bits

This 6-bit code specifies the coarse SNR offset parameter in the bit allocation process.

4.4.3.38 cplfsnroffst - Coupling fine SNR offset - 4 bits

This 4-bit code specifies the coupling channel fine SNR offset in the bit allocation process.

4.4.3.39 cplfgaincod - Coupling fast gain code - 3 bits

This 3-bit code specifies the coupling channel fast gain code used in the bit allocation process.

4.4.3.40 fsnroffst[ch] - Channel fine SNR offset - 4 bits

This 4-bit code specifies the fine SNR offset used in the bit allocation process for channel [ch].

4.4.3.41 fgaincod[ch] - Channel fast gain code - 3 bits

This 3-bit code specifies the fast gain parameter used in the bit allocation process for channel [ch].

4.4.3.42 lfefsnroffst - Low frequency effects channel fine SNR offset - 4 bits

This 4-bit code specifies the fine SNR offset parameter used in the bit allocation process for the lfe channel.

4.4.3.43 lfefgaincod - Low frequency effects channel fast gain code - 3 bits

This 3-bit code specifies the fast gain parameter used in the bit allocation process for the lfe channel.

4.4.3.44 cplleake - Coupling leak initialization exists - 1 bit

If this bit is a 1, leak initialization parameters follow in the bit stream. If this bit is a 0, the previously transmitted values still apply. This parameter shall not be set to 0 in block 0, or in any block for which coupling is enabled but was disabled in the previous block.

4.4.3.45 cplfleak - Coupling fast leak initialization - 3 bits

This 3-bit code specifies the fast leak initialization value for the coupling channel's excitation function calculation in the bit allocation process.

4.4.3.46 cplsleak - Coupling slow leak initialization - 3 bits

This 3-bit code specifies the slow leak initialization value for the coupling channel's excitation function calculation in the bit allocation process.

4.4.3.47 deltbaie - Delta bit allocation information exists - 1 bit

If this bit is a 1, some delta bit allocation information follows in the bit stream. If this bit is a 0, the previously transmitted delta bit allocation information still applies, except for block 0. If deltbaie is 0 in block 0, then cpldeltbae and deltbae[ch] are set to the binary value "10", and no delta bit allocation is applied. Delta bit allocation is described in clause 6.2.2.

4.4.3.48 cpldeltbae - Coupling delta bit allocation exists - 2 bits

This 2-bit code indicates the delta bit allocation strategy for the coupling channel, as shown in Table 4.11. If the reserved state is received, the decoder should not decode audio, and should mute. This parameter shall not be set to "00" in block 0, or in any block for which coupling is enabled but was disabled in the previous block.

Table 4.11: Delta bit allocation exist states

cpldeltbae, deltbae	Code
00	Reuse previous state
01	New info follows
10	Perform no delta alloc
11	Reserved

4.4.3.49 deltbae[ch] - Delta bit allocation exists - 2 bits

This per full bandwidth channel 2-bit code indicates the delta bit allocation strategy for the corresponding channel, as shown in Table 4.11. This parameter shall not be set to "00" in block 0.

4.4.3.50 cpldeltnseg - Coupling delta bit allocation number of segments - 3 bits

This 3-bit code indicates the number of delta bit allocation segments that exist for the coupling channel. The value of this parameter ranges from 1 to 8, and is calculated by adding 1 to the 3-bit binary number represented by the code.

4.4.3.51 cpldeltoffst[seg] - Coupling delta bit allocation offset - 5 bits

The first 5-bit code ([seg] = 0) indicates the number of the first bit allocation band (as specified in clause 6.4.2) of the coupling channel for which delta bit allocation values are provided. Subsequent codes indicate the offset from the previous delta segment end point to the next bit allocation band for which delta bit allocation values are provided.

4.4.3.52 cpldeltlen[seg] - Coupling delta bit allocation length - 4 bits

Each 4-bit code indicates the number of bit allocation bands that the corresponding segment spans.

4.4.3.53 cpldeltba[seg] - Coupling delta bit allocation - 3 bits

This 3-bit value is used in the bit allocation process for the coupling channel.

Each 3-bit code indicates an adjustment to the default masking curve computed in the decoder. The deltas are coded as shown in Table 4.12.

Table 4.12: Bit allocation deltas

cpldeltba, deltba	Adjustment (dB)
000	-24
001	-18
010	-12
011	-6
100	+6
101	+12
110	+18
111	+24

4.4.3.54 deltnseg[ch] - Channel delta bit allocation number of segments - 3 bits

These per full bandwidth channel elements are 3-bit codes indicating the number of delta bit allocation segments that exist for the corresponding channel. The value of this parameter ranges from 1 to 8, and is calculated by adding 1 to the 3-bit binary code.

4.4.3.55 deltoffst[ch][seg] - Channel delta bit allocation offset - 5 bits

The first 5-bit code ([seg] = 0) indicates the number of the first bit allocation band (see clause 6.2.2) of the corresponding channel for which delta bit allocation values are provided. Subsequent codes indicate the offset from the previous delta segment end point to the next bit allocation band for which delta bit allocation values are provided.

4.4.3.56 deltlen[ch][seg] - Channel delta bit allocation length - 4 bits

Each 4-bit code indicates the number of bit allocation bands that the corresponding segment spans.

4.4.3.57 deltba[ch][seg] - Channel delta bit allocation - 3 bits

This 3-bit value is used in the bit allocation process for the indicated channel. Each 3-bit code indicates an adjustment to the default masking curve computed in the decoder. The deltas are coded as shown in Table 4.13.

4.4.3.58 skiple - Skip length exists - 1 bit

If this bit is a 1, then the skipl parameter follows in the bit stream. If this bit is a 0, skipl does not exist.

4.4.3.59 skipl - Skip length - 9 bits

This 9-bit code indicates the number of dummy bytes to skip (ignore) before unpacking the mantissas of the current audio block.

4.4.3.60 skipfld - Skip field - (skipl x 8) bits

This field contains the bytes of data to be skipped, as indicated by the skipl parameter.

4.4.3.61 chmant[ch][bin] - Channel mantissas - 0 bits to 16 bits

The actual quantized mantissa values for the indicated channel. Each value may contain from 0 to as many as 16 bits. The number of mantissas for the indicated channel is equal to nchmant[ch], which may be determined from chbwcod[ch] (see clause 6.1.3) if the channel is not coupled, or from cplbegf (see clause 6.4.2) if the channel is coupled. Detailed information on packed mantissa data is in clause 6.3.

4.4.3.62 cplmant[bin] - Coupling mantissas - 0 bits to 16 bits

The actual quantized mantissa values for the coupling channel. Each value may contain from 0 to as many as 16 bits. The number of mantissas for the coupling channel is equal to ncplmant, which may be determined from:

ncplmant = 12 x ncplsubnd.

4.4.3.63 lfemant[bin] - Low frequency effects channel mantissas - 0 bits to 16 bits

The actual quantized mantissa values for the lfe channel. Each value may contain from 0 to as many as 16 bits. The value of nlfemant is 7, so there are 7 mantissa values for the lfe channel.

4.4.4 auxdata - Auxiliary data field

4.4.4.0 Introduction

Unused data at the end of a syncframe will exist whenever the encoder does not utilize all available data for encoding the audio signal. This may occur if the final bit allocation falls short of using all available bits, or if the input audio signal simply does not require all available bits to be coded transparently. Or, the encoder may be instructed to intentionally leave some bits unused by audio so that they are available for use by auxiliary data. Since the number of bits required for auxiliary data may be smaller than the number of bits available (which will be time varying) in any particular syncframe, a method is provided to signal the number of actual auxiliary data bits in each syncframe.

4.4.4.1 auxbits - Auxiliary data bits - nauxbits bits

This field contains auxiliary data. The total number of bits in this field is:

nauxbits = (bits in syncframe) - (bits used by all bit stream elements except for auxbits).

The number of bits in the syncframe can be determined from the frame size code (frmsizcod) and Table 4.13. The number of bits used includes all bits used by bit stream elements with the exception of auxbits. Any dummy data which has been included with skip fields (skipfld) is included in the used bit count. The length of the auxbits field is adjusted by the encoder such that the crc2 element falls on the last 16-bit word of the syncframe.

Table 4.13: Frame size code table (1 word = 16 bits)

frmsizecod	Nominal bit rate (kbit/s)	Words/syncframe fs = 32 kHz	Words/syncframe fs = 44,1 kHz	Words/syncframe fs = 48 kHz
000000 (0)	32	96	69	64
000001 (0)	32	96	70	64
000010 (1)	40	120	87	80
000011 (1)	40	120	88	80
000100 (2)	48	144	104	96
000101 (2)	48	144	105	96
000110 (3)	56	168	121	112
000111 (3)	56	168	122	112
001000 (4)	64	192	139	128
001001 (4)	64	192	140	128
001010 (5)	80	240	174	160
001011 (5)	80	240	175	160
001100 (6)	96	288	208	192
001101 (6)	96	288	209	192
001110 (7)	112	336	243	224
001111 (7)	112	336	244	224
010000 (8)	128	384	278	256
010001 (8)	128	384	279	256
010010 (9)	160	480	348	320
010011 (9)	160	480	349	320
010100 (10)	192	576	417	384
010101 (10)	192	576	418	384
010110 (11)	224	672	487	448

frmsizecod	Nominal bit rate (kbit/s)	Words/syncframe fs = 32 kHz	Words/syncframe fs = 44,1 kHz	Words/syncframe fs = 48 kHz
010111 (11)	224	672	488	448
011000 (12)	256	768	557	512
011001 (12)	256	768	558	512
011010 (13)	320	960	696	640
011011 (13)	320	960	697	640
011100 (14)	384	1 152	835	768
011101 (14)	384	1 152	836	768
011110 (15)	448	1 344	975	896
011111 (15)	448	1 344	976	896
100000 (16)	512	1 536	1 114	1 024
100001 (16)	512	1 536	1 115	1 024
100010 (17)	576	1 728	1 253	1 152
100011 (17)	576	1 728	1 254	1 152
100100 (18)	640	1 920	1 393	1 280
100101 (18)	640	1 920	1 394	1 280
NOTE: f_s : sampling frequency.				

If the number of user bits indicated by auxdatal is smaller than the number of available aux bits nauxbits, the user data is located at the end of the auxbits field. This allows a decoder to find and unpack the auxdatal user bits without knowing the value of nauxbits (which can only be determined by decoding the audio in the entire syncframe). The order of the user data in the auxbits field is forward. Thus the aux data decoder (which may not decode any audio) may simply look to the end of the AC-3 syncframe to find auxdatal, backup auxdatal bits (from the beginning of auxdatal) in the data stream, and then unpack auxdatal bits moving forward in the data stream.

4.4.4.2 auxdatal - Auxiliary data length - 14 bits

This 14-bit integer value indicates the length, in bits, of the user data in the auxbits auxiliary field.

4.4.4.3 auxdatae - Auxiliary data exists - 1 bit

If this bit is a 1, then the auxdatal parameter precedes in the bit stream. If this bit is a 0, auxdatal does not exist, and there is no user data.

4.4.5 errorcheck - Frame error detection field

4.4.5.1 crcrsv - CRC reserved bit - 1 bit

Reserved for use in specific applications to ensure crc2 will not be equal to the sync word. Use of this bit is optional by encoders. If the crc2 calculation results in a value equal to the syncword, the crcrsv bit may be inverted. This will result in a crc2 value which is not equal to the syncword.

4.4.5.2 crc2 - Cyclic redundancy check 2 - 16 bits

The 16-bit CRC applies to the entire syncframe. The details of the CRC checking are described in clause 6.10.1.

4.5 Bit stream constraints

The following constraints shall be imposed upon the encoded bit stream by the AC-3 encoder. These constraints allow AC-3 decoders to be manufactured with smaller input memory buffers:

1) The combined size of the syncinfo fields, the bsi fields, block 0 and block 1 combined, shall not exceed 5/8 of the syncframe.

2) The combined size of the block 5 mantissa data, the auxiliary data fields, and the errorcheck fields shall not exceed the final 3/8 of the syncframe.

3) Block 0 shall contain all necessary information to begin correctly decoding the bit stream.

4) Whenever the state of cplinu changes from off to on, all coupling information shall be included in the block in which coupling is turned on. No coupling related information shall be reused from any previous blocks where coupling may have been on.

5) Coupling shall not be used in dual mono (1 + 1) or mono (1/0) modes. For blocks in which coupling is used, there shall be at least two channels in coupling.

6) Bit stream elements shall not be reused from a previous block if other bit stream parameters change the dimensions of the elements to be reused. For example, exponents shall not be reused if the start or end mantissa bin changes from the previous block.

5 Decoding the AC-3 bit stream

5.1 Introduction

Clause 4 specifies the details of the AC-3 bit stream syntax. The following clauses provide an overview of the AC-3 decoding process as diagrammed in Figure 5.1, where the decoding process flow is shown as a sequence of blocks down the centre of the page, and some of the information flow is indicated by arrowed lines at the sides of the page. More detailed information on some of the processing blocks will be found in clause 6. The decoder described in the following clauses should be considered one example of a decoder. Other methods may exist to implement decoders, and these other methods may have advantages in certain areas (such as instruction count, memory requirement, number of transforms required, etc.).

5.2 Summary of the decoding process

5.2.1 Input bit stream

5.2.1.0 Introduction

The input bit stream will typically come from a transmission or storage system. The interface between the source of AC-3 data and the AC-3 decoder is not specified in the present document. The details of the interface affect a number of decoder implementation details.

5.2.1.1 Continuous or burst input

The encoded AC-3 data may be input to the decoder as a continuous data stream at the nominal bit rate, or chunks of data may be burst into the decoder at a high rate with a low duty cycle. For burst mode operation, either the data source or the decoder may be the master controlling the burst timing. The AC-3 decoder input buffer may be smaller in size if the decoder can request bursts of data on an as-needed basis. However, the external buffer memory may be larger in this case.

5.2.1.2 Byte or word alignment

Most applications of the present document will convey the bit AC-3 bit stream with byte or (16-bit) word alignment. The syncframe is always an integral number of words in length. The decoder may receive data as a continuous serial stream of bits without any alignment. Or, the data may be input to the decoder with either byte or word (16-bit) alignment. Byte or word alignment of the input data may allow some simplification of the decoder. Alignment does reduce the probability of false detection of the sync word.

5.2.2 Synchronization and error detection

The AC-3 bit-stream format allows rapid synchronization. The 16-bit sync word has a low probability of false detection. With no input stream alignment the probability of false detection of the sync word is 0,0015 % per input stream bit position. For a bit rate of 384 kbit/s, the probability of false sync word detection is 19 % per syncframe. Byte alignment of the input stream drops this probability to 2,5 %, and word alignment drops it to 1,2 %.

When a sync pattern is detected the decoder may be estimated to be in sync and one of the CRC words (crc1 or crc2) may be checked. Since crc1 comes first and covers the first 5/8 of the syncframe, the result of a crc1 check may be available after only 5/8 of the syncframe has been received. Or, the entire syncframe size can be received and crc2 checked. If either CRC checks, the decoder may safely be presumed to be in sync and decoding and reproduction of audio may proceed. The chance of false sync in this case would be the concatenation of the probabilities of a false sync word detection and a CRC misdetection of error. The CRC check is reliable to 0,0015 %. This probability, concatenated with the probability of a false sync detection in a byte aligned input bit stream, yield a probability of false synchronization of 0,000035 % (or about once in 3 million synchronization attempts).

If this small probability of false sync is too large for an application, there are several methods which may reduce it. The decoder may only presume correct sync in the case that both CRC words check properly. The decoder may require multiple sync words to be received with the proper alignment. If the data transmission or storage system is aware that data is in error, this information may be made known to the decoder.

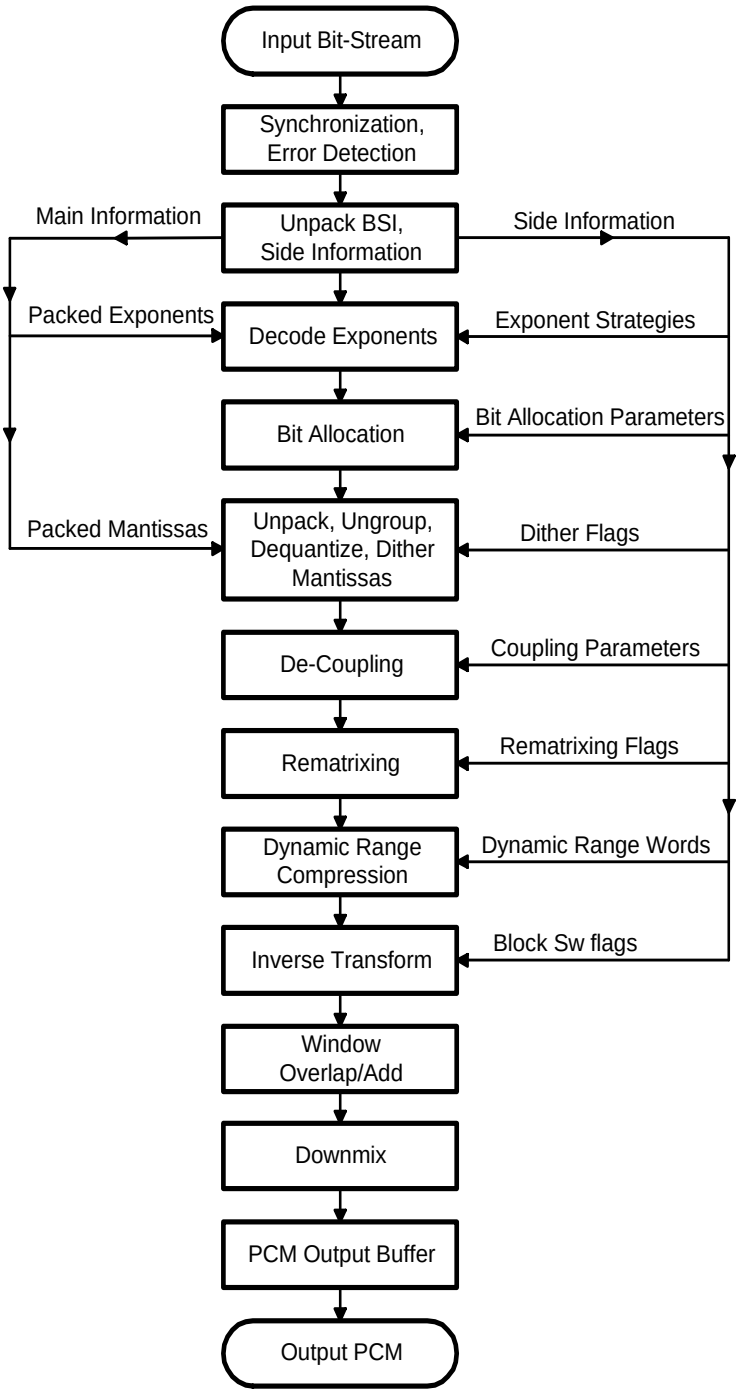

Figure 5.1: Flow diagram of the decoding process

Additional details on methods of bit stream synchronization are not provided in the present document. Details on the CRC calculation are provided in clause 6.10.

5.2.3 Unpack BSI, side information

Inherent to the decoding process is the unpacking (de-multiplexing) of the various types of information included in the bit stream. Some of these items may be copied from the input buffer to dedicated registers, some may be copied to specific working memory location, and some of the items may simply be located in the input buffer with pointers to them saved to another location for use when the information is required. The information to be unpacked is specified in detail in clause 4.3. Further details on the unpacking of BSI and side information are not provided in the present document.

5.2.4 Decode exponents

The exponents are delivered in the bit stream in an encoded form. In order to unpack and decode the exponents two types of side information are required. First, the number of exponents is determined. For fbw channels this may be determined from either chbwcod[ch] (for uncoupled channels) or from cplbegf (for coupled channels). For the coupling channel, the number of exponents may be determined from cplbegf and cplendf. For the lfe channel (when on), there are always 7 exponents. Second, the exponent strategy in use (d15, etc.) by each channel is determined. The details on how to unpack and decode exponents are provided in clause 6.1.

5.2.5 Bit allocation

The bit allocation computation reveals how many bits are used for each mantissa. The inputs to the bit allocation computation are the decoded exponents, and the bit allocation side information. The outputs of the bit allocation computation are a set of bit allocation pointers (baps), one bap for each coded mantissa. The bap indicates the quantizer used for the mantissa, and how many bits in the bit stream were used for each mantissa. The bit allocation computation is described in detail in clause 6.2.

5.2.6 Process mantissas

The coarsely quantized mantissas make up the bulk of the AC-3 data stream. Each mantissa is quantized to a level of precision indicated by the corresponding bap. In order to pack the mantissa data more efficiently, some mantissas are grouped together into a single transmitted value. For instance, two 11-level quantized values are conveyed in a single 7-bit code (3,5 bits/value) in the bit stream.

The mantissa data is unpacked by peeling off groups of bits as indicated by the baps. Grouped mantissas are ungrouped. The individual coded mantissa values are converted into a de-quantized value. Mantissas which are indicated as having zero bits may be reproduced as either zero, or by a random dither value (under control of the dither flag). The mantissa processing is described in full detail in clause 6.3.

5.2.7 Decoupling

When coupling is in use, the channels which are coupled are decoupled. Decoupling involves reconstructing the high frequency section (exponents and mantissas) of each coupled channel, from the common coupling channel and the coupling coordinates for the individual channel. Within each coupling band, the coupling channel coefficients (exponent and mantissa) are multiplied by the individual channel coupling coordinates. The coupling process is described in detail in clause 6.4.

5.2.8 Rematrixing

In the 2/0 audio coding mode rematrixing may be employed, as indicated by the rematrix flags (rematflg[rbnd]). Where the flag indicates a band is rematrixed, the coefficients encoded in the bit stream are sum and difference values instead of left and right values. Rematrixing is described in detail in clause 6.5.

5.2.9 Dynamic range compression

For each block of audio a dynamic range control value (dynrng) may be included in the bit stream. The decoder, by default, shall use this value to alter the magnitude of the coefficient (exponent and mantissa) as specified in clause 6.7.

5.2.10 Inverse transform

The decoding steps described above will result in a set of frequency coefficients for each encoded channel. The inverse transform converts the blocks of frequency coefficients into blocks of time samples. The inverse transform is detailed in clause 6.9.

5.2.11 Window, overlap/add

The individual blocks of time samples are windowed, and adjacent blocks are overlapped and added together in order to reconstruct the final continuous time output PCM audio signal. The window and overlap/add steps are described along with the inverse transform in clause 6.9.

5.2.12 Downmixing

If the number of channels required at the decoder output is smaller than the number of channels which are encoded in the bit stream, then downmixing is required. Downmixing in the time domain is shown in this example decoder. Since the inverse transform is a linear operation, it is also possible to downmix in the frequency domain prior to transformation. Clause 6.8 describes downmixing and specifies the downmix coefficients which decoders shall employ.

5.2.13 PCM output buffer

Typical decoders will provide PCM output samples at the PCM sample rate. Since blocks of samples result from the decoding process, an output buffer is typically required. The present document does not specify or describe output buffering in any further detail.

5.2.14 Output PCM

The output PCM samples may be delivered in form suitable for interconnection to a digital to analogue converter (DAC), or in any other form. The present document does not specify the output PCM format.

6 Algorithmic details

6.0 Introduction

The following clauses describe various aspects of AC-3 coding in detail.

6.1 Exponent coding

6.1.1 Overview

The actual audio information conveyed by the AC-3 bit stream consists of the quantized frequency coefficients. The coefficients are delivered in floating point form, with each coefficient consisting of an exponent and a mantissa. This clause describes how the exponents are encoded and packed into the bit stream.

Exponents are 5-bit values which indicate the number of leading zeros in the binary representation of a frequency coefficient. The exponent acts as a scale factor for each mantissa, equal to 2^{-exp}. Exponent values are allowed to range from 0 (for the largest value coefficients with no leading zeros) to 24. Exponents for coefficients which have more than 24 leading zeros are fixed at 24, and the corresponding mantissas are allowed to have leading zeros. Exponents require 5 bits in order to represent all allowed values.

AC-3 bit streams contain coded exponents for all independent channels, all coupled channels, and for the coupling and low frequency effects channels (when they are enabled). Since audio information is not shared across syncframes, block 0 of every syncframe will include new exponents for every channel. Exponent information may be shared across blocks within a syncframe, so blocks 1 through 5 may reuse exponents from previous blocks.

AC-3 exponent transmission employs differential coding, in which the exponents for a channel are differentially coded across frequency. The first exponent of a fbw or lfe channel is always sent as a 4-bit absolute value, ranging from 0 to 15. The value indicates the number of leading zeros of the first (DC term) transform coefficient. Successive (going higher in frequency) exponents are sent as differential values which shall be added to the prior exponent value in order to form the next absolute value.

The differential exponents are combined into groups in the audio block. The grouping is done by one of three methods, d15, d25, or d45, which are referred to as exponent strategies. The number of grouped differential exponents placed in the audio block for a particular channel depends on the exponent strategy and on the frequency bandwidth information for that channel. The number of exponents in each group depends only on the exponent strategy.

An AC-3 audio block contains two types of fields with exponent information. The first type defines the exponent coding strategy for each channel, and the second type contains the actual coded exponents for channels requiring new exponents. For independent channels, frequency bandwidth information is included along with the exponent strategy fields. For coupled channels, and the coupling channel, the frequency information is found in the coupling strategy fields.

6.1.2 Exponent strategy

Exponent strategy information for every channel is included in every AC-3 audio block. Information is never shared across syncframes, so block 0 will always contain a strategy indication (d15, d25, or d45) for each channel. Blocks 1 through 5 may indicate reuse of the prior (within the same syncframe) exponents. The three exponent coding strategies provide a trade-off between data rate required for exponents, and their frequency resolution. The d15 mode provides the finest frequency resolution, and the d45 mode requires the least amount of data. In all three modes, a number of differential exponents are combined into 7-bit words when coded into an audio block. The main difference between the modes is how many differential exponents are combined together.

The absolute exponents found in the bit stream at the beginning of the differentially coded exponent sets are sent as 4-bit values which have been limited in either range or resolution in order to save one bit. For fbw and lfe channels, the initial 4-bit absolute exponent represents a value from 0 to 15. Exponent values larger than 15 are limited to a value of 15. For the coupled channel, the 5-bit absolute exponent is limited to even values, and the LSB is not transmitted. The resolution has been limited to valid values of 0, 2, 4...24. Each differential exponent can take on one of five values: -2, -1, 0, +1, +2. This allows deltas of up to ±2 (±12 dB) between exponents. These five values are mapped into the values 0, 1, 2, 3, 4 before being grouped, as shown in Table 6.1.

Table 6.1: Mapping of differential exponent values, d15 mode

Differential exponent	Mapped value
+2	4
+1	3
0	2
-1	1
-2	0
NOTE 1: Mapped value: differential exponent +2. NOTE 2: Differential exponent: mapped value -2.	

In the d15 mode, the above mapping is applied to each individual differential exponent for coding into the bit stream. In the d25 mode, each *pair* of differential exponents is represented by a single mapped value in the bit stream. In this mode the second differential exponent of each pair is implied as a delta of 0 from the first element of the pair as indicated in Table 6.2.

Table 6.2: Mapping of differential exponent values, d25 mode

Differential exponent *n*	Differential exponent *n* + 1	Mapped value
+2	0	4
+1	0	3
0	0	2
-1	0	1
-2	0	0

The d45 mode is similar to the d25 mode except that *quads* of differential exponents are represented by a single mapped value, as indicated by Table 6.3.

Table 6.3: Mapping of differential exponent values, d45 mode

Differential exponent *n*	Differential exponent *n* + 1	Differential exponent *n* + 2	Differential exponent *n* + 3	Mapped value
+2	0	0	0	4
+1	0	0	0	3
0	0	0	0	2
-1	0	0	0	1
-2	0	0	0	0

Since a single exponent is effectively shared by 2 or 4 different mantissas, encoders shall ensure that the exponent chosen for the pair or quad is the minimum absolute value (corresponding to the largest exponent) needed to represent all the mantissas.

For all modes, sets of three adjacent (in frequency) mapped values (M1, M2 and M3) are grouped together and coded as a 7-bit value according to the following formula:

$$\text{Coded 7-bit grouped value} = (25 \times M1) + (5 \times M2) + M3.$$

The exponent field for a given channel in an AC-3 audio block consists of a single absolute exponent followed by a number of these grouped values.

6.1.3 Exponent decoding

The exponent strategy for each coupled and independent channel is included in a set of 2-bit fields designated chexpstr[ch]. When the coupling channel is present, a cplexpstr strategy code is also included. Table 6.4 shows the mapping from exponent strategy code into exponent strategy.

Table 6.4: Exponent strategy coding

chexpstr[ch], cplexpstr	Exponent strategy	Exponents per group
00	Reuse prior exponents	0
01	d15	3
10	d25	6
11	d45	12

When the low frequency effects channel is enabled the lfeexpstr field is present. It is decoded as shown in Table 6.5.

Table 6.5: lfe channel exponent strategy coding

lfeexpstr	Exponent strategy	Exponents per group
0	Reuse prior exponents	0
1	d15	3

Following the exponent strategy fields in the bit stream is a set of channel bandwidth codes, chbwcod[ch]. These are only present for independent channels (channels not in coupling) that have new exponents in the current block. The channel bandwidth code defines the end mantissa bin number for that channel according to the following:

```
endmant[ch] = ((chbwcod[ch] + 12) × 3) + 37;          /* (ch is not coupled) */
```

For coupled channels the end mantissa bin number is defined by the starting bin number of the coupling channel:

```
endmant[ch] = cplstrtmant;                            /* (ch is coupled) */
```

where cplstrtmant is as derived below. By definition the starting mantissa bin number for independent and coupled channels is 0:

```
strtmant[ch] = 0.
```

For the coupling channel, the frequency bandwidth information is derived from the fields cplbegf and cplendf found in the coupling strategy information. The coupling channel starting and ending mantissa bins are defined as:

$$cplstrtmant = (cplbegf \times 12) + 37;$$

$$cplendmant = ((cplendf + 3) \times 12) + 37.$$

The low frequency effects channel, when present, always starts in bin 0 and always has the same number of mantissas:

$$lfestrtmant = 0;$$

$$lfeendmant = 7.$$

The second set of fields contains coded exponents for all channels indicated to have new exponents in the current block. These fields are designated as exps[ch][grp] for independent and coupled channels, cplexps[grp] for the coupling channel, and lfeexps[grp] for the low frequency effects channel. The first element of the exps fields (exps[ch][0]) and the lfeexps field (lfeexps[0]) is always a 4-bit absolute number. For these channels the absolute exponent always contains the exponent value of the first transform coefficient (bin #0). These 4-bit values correspond to a 5-bit exponent which has been limited in range (0 to 15, instead of 0 to 24), i.e. the most significant bit is zero. The absolute exponent for the coupled channel, cplabsexp, is only used as a reference to begin decoding the differential exponents for the coupling channel (i.e. it does not represent an actual exponent). The cplabsexp is contained in the audio block as a 4-bit value, however it corresponds to a 5-bit value. The LSB of the coupled channel initial exponent is always 0, so the decoder shall take the 4-bit value which was sent, and double it (left shift by 1) in order to obtain the 5-bit starting value.

For each coded exponent set the number of grouped exponents (not including the first absolute exponent) to decode from the bit stream is derived as follows:

For independent and coupled channels:

```
nchgrps[ch]  = truncate ((endmant[ch] - 1) / 3);          /* for d15 mode */
             = truncate {(endmant[ch] - 1 + 3) / 6);      /* for d25 mode */
             = truncate {(endmant[ch] - 1 + 9) / 12);     /* for d45 mode */
```

For the coupling channel:

```
ncplgrps     = (cplendmant - cplstrtmant) 3;              /* for d15 mode */
             = (cplendmant - cplstrtmant ) 6;             /* for d15 mode */
             = (cplendmant - cplstrtmant) 12;             /* for d15 mode */
```

For the low frequency effects channel:

$$nlfegrps = 2.$$

Decoding a set of coded grouped exponents will create a set of 5-bit absolute exponents. The exponents are decoded as follows:

1) Each 7-bit grouping of mapped values (gexp) is decoded using the inverse of the encoding procedure:

   - M1 = truncate (gexp / 25);
   - M2 = truncate ((gexp % 25) / 5);
   - M3 = (gexp % 25) % 5.

2) Each mapped value is converted to a differential exponent (dexp) by subtracting the mapping offset:

   - dexp = M2.

3) The set of differential exponents is converted to absolute exponents by adding each differential exponent to the absolute exponent of the previous frequency bin:

$$exp[n] = exp[n-1] + dexp[n].$$

4) For the d25 and d45 modes each absolute exponent is copied to the remaining members of the pair or quad.

The above procedure can be summarized as follows:

Pseudo code

```
/* unpack the mapped values */
for (grp = 0; grp < ngrps; grp++)
{
    expacc = gexp[grp];
    dexp[grp * 3] = truncate (expacc / 25);
    expacc = expacc - ( 25 * dexp[grp * 3]);
    dexp[(grp * 3) + 1] = truncate (expacc / 5);
    expacc = expacc - (5 * dexp[(grp * 3) + 1]);
    dexp[(grp * 3) + 2] = expacc;
}
/* unbiased mapped values */
for (grp = 0; grp < (ngrps * 3); grp++)
{
    dexp[grp] = dexp[grp] - 2;
}
/* convert from differentials to absolutes */
prevexp = absexp;
for (i = 0; i < (ngrps * 3); i++)
{
    aexp[i] = prevexp + dexp[i];
    prevexp = aexp[i];
}
/* expand to full absolute exponent array, using grpsize */
exp[0] = absexp;
for (i = 0; i < (ngrps * 3); i++)
{
    for (j = 0; j < grpsize; j++)
    {
        exp[(i * grpsize) + j +1] = aexp[i];
    }
}
```

where:

- ngrps: number of grouped exponents (nchgrps[ch], ncplgrps, or nlfegrps).
- grpsize = 1 for d15;

  = 2 for d25;

  = 4 for d45.
- absexp: absolute exponent (exps[ch][0], (cplabsexp<<1), or lfeexps[0]).

For the coupling channel the above output array, exp[n], should be offset to correspond to the coupling start mantissa bin:

$$cplexp[n + cplstrtmant] = exp[n + 1].$$

For the remaining channels exp[n] will correspond directly to the absolute exponent array for that channel.

6.2 Bit allocation

6.2.1 Overview

The bit allocation routine analyses the spectral envelope of the audio signal being coded with respect to masking effects to determine the number of bits to assign to each transform coefficient mantissa. In the encoder, the bit allocation is performed globally on the ensemble of channels as an entity, from a common bit pool. There are no preassigned exponent or mantissa bits, allowing the routine to flexibly allocate bits across channels, frequencies, and audio blocks in accordance with signal demand.

The bit allocation contains a parametric model of human hearing for estimating a noise level threshold, expressed as a function of frequency, which separates audible from inaudible spectral components. Various parameters of the hearing model can be adjusted by the encoder depending upon signal characteristics. For example, a prototype masking curve is defined in terms of two piecewise continuous line segments, each with its own slope and y-axis intercept. One of several possible slopes and intercepts is selected by the encoder for each line segment. The encoder may iterate on one or more such parameters until an optimal result is obtained. When all parameters used to estimate the noise level threshold have been selected by the encoder, the final bit allocation is computed. The model parameters are conveyed to the decoder with other side information. The decoder executes the routine in a single pass.

The estimated noise level threshold is computed over 50 bands of non-uniform bandwidth (an approximate 1/6 octave scale). The banding structure, defined by tables in the next clause, is independent of sampling frequency. The required bit allocation for each mantissa is established by performing a table look-up based upon the difference between the input signal power spectral density (PSD) evaluated on a fine-grain uniform frequency scale, and the estimated noise level threshold evaluated on the coarse-grain (banded) frequency scale. Therefore, the bit allocation result for a particular channel has spectral granularity corresponding to the exponent strategy employed. More specifically, a separate bit allocation will be computed for each mantissa within a d15 exponent set, each pair of mantissas within a d25 exponent set, and each quadruple of mantissas within a d45 exponent set.

The bit allocation shall be computed in the decoder whenever the exponent strategy (chexpstr, cplexpstr, lfeexpstr) for one or more channels does not indicate reuse, or whenever baie, snroffste, or deltbaie = 1. Accordingly, the bit allocation can be updated at a rate ranging from once per audio block to once per 6 audio blocks, including the integral steps in between. A complete set of new bit allocation information is always transmitted in audio block 0.

Since the parametric bit allocation routine is required to generate identical results in all encoder and decoder implementations, each step is defined exactly in terms of fixed-point integer operations and table look-ups. Throughout the discussion below, signed two's complement arithmetic is employed. All additions are performed with an accumulator of 14 or more bits. All intermediate results and stored values are 8-bit values.

6.2.2 Parametric bit allocation

6.2.2.0 Introduction

This clause describes the seven-step procedure for computing the output of the parametric bit allocation routine in the decoder. The approach outlined here starts with a single uncoupled or coupled exponent set and processes all the input data for each step prior to continuing to the next one. This technique, called vertical execution, is conceptually straightforward to describe and implement. Alternatively, the seven steps can be executed horizontally, in which case multiple passes through all seven steps are made for separate subsets of the input exponent set.

The choice of vertical vs. horizontal execution depends upon the relative importance of execution time vs. memory usage in the final implementation. Vertical execution of the algorithm is usually faster due to reduced looping and context save overhead. However, horizontal execution requires less RAM to store the temporary arrays generated in each step. Hybrid horizontal/vertical implementation approaches are also possible which combine the benefits of both techniques.

6.2.2.1 Initialization

Compute start/end frequencies for the channel being decoded. These are computed from parameters in the bit stream as follows:

Pseudo code

```
/* for fbw channels */
for(ch=0; ch<nfchans; ch++)
{
    strtmant[ch] = 0;
    if(chincpl[ch]) endmant[ch] = 37 + (12 x cplbegf); /* channel is coupled */
    else endmant[ch] = 37 + (3 x (chbwcod + 12)); /* channel is not coupled */
}
/* for coupling channel */
cplstrtmant = 37 + (12 x cplbegf);
cplendmant = 37 + (12 x (cplendf + 3));
/* for lfe channel */
lfestartmant = 0;
lfeendmant = 7;
```

Special case processing step:

Before continuing with the initialization procedure, all snr offset parameters from the bit stream should be evaluated. These include csnroffst, fsnroffst[ch], cplfsnroffst, and lfefsnroffst. If they are all found to be equal to zero, then all elements of the bit allocation pointer array bap[] should be set to zero, and no other bit allocation processing is required for the current audio block.

Perform table look-ups to determine the values of sdecay, fdecay, sgain, dbknee, and floor from parameters in the bit stream as follows:

Pseudo code

```
sdecay = slowdec[sdcycod];      /* Table 6.6 */
fdecay = fastdec[fdcycod];  /* Table 6.7 */
sgain = slowgain[sgaincod]; /* Table 6.8 */
dbknee = dbpbtab[dbpbcod];      /* Table 6.9 */
floor = floortab[floorcod]; /* Table 6.10 */
```

Initialize as follows for uncoupled portion of fbw channel:

Pseudo code

```
start = strtmant[ch];
end = endmant[ch];
lowcomp = 0;
fgain = fastgain[fgaincod[ch]];  /* Table 6.11 */
snroffset[ch] = ((csnroffst - 15) << 4 + fsnroffst[ch]) << 2;
```

Initialize as follows for coupling channel:

Pseudo code

```
start = cplstrtmant;
end = cplendmant;
fgain = fastgain[cplfgaincod];  /* Table 6.11 */
snroffset = ((csnroffst - 15) << 4 + cplfsnroffst) << 2;
fastleak = (cplfleak << 8) + 768;
slowleak = (cplsleak << 8) + 768;
```

Initialize as follows for lfe channel:

Pseudo code

```
start = lfestrtmant;
end = lfeendmant;
lowcomp = 0;
fgain = fastgain[lfefgaincod];
snroffset = ((csnroffst - 15) << 4 + lfefsnroffst) << 2;
```

6.2.2.2 Exponent mapping into psd

This step maps decoded exponents into a 13-bit signed log power-spectral density function.

Pseudo code

```
for (bin=start; bin<end; bin++)
{
    psd[bin] = (3 072 - (exp[bin] << 7));
}
```

Since exp[k] assumes integral values ranging from 0 to 24, the dynamic range of the psd[] values is from 0 (for the lowest-level signal) to 3 072 for the highest-level signal. The resulting function is represented on a fine-grain, linear frequency scale.

6.2.2.3 psd integration

This step of the algorithm integrates fine-grain psd values within each of a multiplicity of 1/6th octave bands. Table 6.12 contains the 50 array values for bndtab[] and bndsz[]. The bndtab[] array gives the first mantissa number in each band. The bndsz[] array provides the width of each band in number of included mantissas.

Table 6.13 contains the 256 array values for masktab[], showing the mapping from mantissa number into the associated 1/6 octave band number. These two tables contain duplicate information, all of which need not be available in an actual implementation. They are shown here for simplicity of presentation only.

The integration of psd values in each band is performed with log-addition. The log-addition is implemented by computing the difference between the two operands and using the absolute difference divided by 2 as an address into a length 256 look-up table, latab[], shown in Table 6.14.

Pseudo code

```
j = start ;
k = masktab[start] ;
do
{
    lastbin = min(bndtab[k] + bndsz[k], end);
    bndpsd[k] = psd[j] ;
    j++ ;
    for (i = j; i < lastbin; i++)
    {
        bndpsd[k] = logadd(bndpsd[k], psd[j]) ;
        j++ ;
    }
    k++ ;
}
while (end > lastbin) ;
logadd(a, b)
{
    c = a □ b ;
    address = min((abs(c) >> 1), 255) ;
    if (c >= 0)
    {
        return(a + latab(address)) ;
    }
        else
        {
            return(b + latab(address)) ;
        }
    }
}
```

6.2.2.4 Compute excitation function

The excitation function is computed by applying the prototype masking curve selected by the encoder (and transmitted to the decoder) to the integrated psd spectrum (bndpsd[]). The result of this computation is then offset downward in amplitude by the fgain and sgain parameters, which are also obtained from the bit stream.

Pseudo code

```
bndstrt = masktab[start];
bndend = masktab[end - 1] + 1;
if (bndstrt == 0) /* For fbw and lfe channels */
{ /* note: do not call calc_lowcomp() for the last band of the lfe channel, (bin = 6) */
    lowcomp = calc_lowcomp(lowcomp, bndpsd[0], bndpsd[1], 0);
    excite[0] = bndpsd[0] - fgain - lowcomp;
    lowcomp = calc_lowcomp(lowcomp, bndpsd[1], bndpsd[2], 1);
    excite[1] = bndpsd[1] - fgain - lowcomp;
    begin = 7;
    for (bin = 2; bin < 7; bin++)
    {
        if ((bndend != 7) || (bin != 6)) /* skip for last bin of lfe channels */
        {
            lowcomp = calc_lowcomp(lowcomp, bndpsd[bin], bndpsd[bin+1], bin) ;
        }
        fastleak = bndpsd[bin] - fgain ;
        slowleak = bndpsd[bin] - sgain ;
        excite[bin] = fastleak - lowcomp ;
        if ((bndend != 7) || (bin != 6)) /* skip for last bin of lfe channel */
        {
            if (bndpsd[bin] <= bndpsd[bin+1])
            {
                begin = bin + 1 ;
                break ;
            }
        }
    }
```

Pseudo code

```
    for (bin = begin; bin < min(bndend, 22); bin++)
    {
        if ((bndend != 7) || (bin != 6)) /* skip for last bin of lfe channel */
        {
            lowcomp = calc_lowcomp(lowcomp, bndpsd[bin], bndpsd[bin+1], bin) ;
        }
        fastleak -= fdecay ;
        fastleak = max(fastleak, bndpsd[bin] - fgain) ;
        slowleak -= sdecay ;
        slowleak = max(slowleak, bndpsd[bin] - sgain) ;
        excite[bin] = max(fastleak - lowcomp, slowleak) ;
    }
    begin = 22 ;
}
else /* For coupling channel */
{
    begin = bndstrt;
}
for (bin = begin; bin < bndend; bin++)
{
    fastleak - = fdecay;
    fastleak = max(fastleak, bndpsd[bin] - fgain);
    slowleak -= sdecay;
    slowleak = max(slowleak, bndpsd[bin] - sgain);
    excite[bin] = max(fastleak, slowleak);
}
calc_lowcomp(a, b0, b1, bin)
{
    if (bin < 7)
    {
        if ((b0 + 256) == b1);
        {
            a = 384;
        }
        else if (b0 > b1)
        {
            a = max(0, a - 64);
        }
    }
    else if (bin < 20)
    {
        if ((b0 + 256) == b1)
        {
            a = 320;
        }
        else if (b0 > b1)
        {
            a = max(0, a - 64);
        }
    }
    else
    {
        a = max(0, a - 128);
    }
    return(a);
}
```

6.2.2.5 Compute masking curve

This step computes the masking (noise level threshold) curve from the excitation function, as shown below. The hearing threshold hth[][] is shown in Table 6.15. The fscod and dbpbcod variables are received by the decoder in the bit stream.

Pseudo code

```
for (bin = bndstrt; bin < bndend; bin++)
{
    if (bndpsd[bin] < dbknee)
    {
        excite[bin] += ((dbknee - bndpsd[bin]) >> 2);
    }
    mask[bin] = max(excite[bin], hth[fscod][bin]);
}
```

6.2.2.6 Apply delta bit allocation

The optional delta bit allocation, dba, information in the bit stream provides a means for the encoder to transmit side information to the decoder which directly increases or decreases the masking curve obtained by the parametric routine. Delta bit allocation can be enabled by the encoder for audio blocks which derive an improvement in audio quality when the default bit allocation is appropriately modified. The delta bit allocation option is available for each fbw channel and the coupling channel.

In the event that delta bit allocation is not being used, and no dba information is included in the bit stream, the decoder shall not modify the default allocation. One way to ensure this is to initialize the cpldeltnseg and deltnseg[ch] delta bit allocation variables to 0 at the beginning of each syncframe. This makes the dba processing (shown below) to immediately terminate, unless dba information (including cpldeltnseg and deltnseg[ch]) is included in the bit stream.

The dba information which modifies the decoder bit allocation is transmitted as side information. The allocation modifications occur in the form of adjustments to the default masking curve computed in the decoder. Adjustments can be made in multiples of ±6 dB. On the average, a masking curve adjustment of -6 dB corresponds to an increase of 1 bit of resolution for all the mantissas in the affected $1/6^{th}$ octave band. The following code indicates, for a single channel, how the modification is performed. The modification calculation is performed on the coupling channel (where deltnseg below equals cpldeltnseg) and on each fbw channel (where deltnseg equals deltnseg[ch]).

Pseudo code

```
if ((deltbae == 0) || (deltbae == 1))
{
    band = 0;
    for (seg = 0; seg < deltnseg+1; seg++)
    {
        band += deltoffst[seg];
        if (deltba[seg] >= 4)
        {
            delta = (deltba[seg] - 3) << 7;
        }
        else
        {
            delta = (deltba[seg] - 4) << 7;
        }
        for (k = 0; k < deltlen[seg]; k++)
        {
            mask[band] += delta;
            band++;
        }
    }
}
```

6.2.2.7 Compute bit allocation

The bit allocation pointer array (bap[]) is computed in this step. The masking curve, adjusted by snroffset in an earlier step and then truncated, is subtracted from the fine-grain psd[] array. The difference is right-shifted by 5 bits, thresholded, and then used as an address into baptab[] to obtain the final allocation. The baptab[] array is shown in Table 6.16.

The sum of all channel mantissa allocations in one syncframe is constrained by the encoder to be less than or equal to the total number of mantissa bits available for that syncframe. The encoder accomplishes this by iterating on the values of csnroffst and fsnroffst (or cplfsnroffst or lfefsnroffst for the coupling and low frequency effects channels) to obtain an appropriate result. The decoder is guaranteed to receive a mantissa allocation which meets the constraints of a fixed transmission bit rate.

At the end of this step, the bap[] array contains a series of 4-bit pointers. The pointers indicate how many bits are assigned to each mantissa. The correspondence between bap pointer value and quantization accuracy is shown in Table 6.17.

Pseudo code

```
i = start ;
j = masktab[start] ;
do
{
    lastbin = min(bndtab[j] + bndsz[j], end) ;
    mask[j] -= snroffset ;
    mask[j] -= floor ;
    if (mask[j] < 0)
    {
        mask[j] = 0 ;
    }
    mask[j] &= 0x1fe0 ;
    mask[j] += floor ;
    for (k = i; k < lastbin; k++)
    {
        address = (psd[i] - mask[j]) >> 5 ;
        address = min(63, max(0, address)) ;
        bap[i] = baptab[address] ;
        i++ ;
    }
    j++;
}
while (end > lastbin) ;
```

6.2.3 Bit allocation tables

Table 6.6: Slow decay table, slowdec[]

Address	slowdec[address]
0	0×0f
1	0×11
2	0×13
3	0×15

Table 6.7: Fast decay table, fastdec[]

Address	fastdec[address]
0	0×3f
1	0×53
2	0×67
3	0×7b

Table 6.8: Slow gain table, slowgain[ ]

Address	slowgain[address]
0	0×540
1	0×4d8
2	0×478
3	0×410

Table 6.9: dB/bit table, dbpbtab[]

Address	dbpbtab[address]
0	0×000
1	0×700
2	0×900
3	0×b00

Table 6.10: Floor table, floortab[]

Address	floortab[address]
0	0×2f0
1	0×2b0
2	0×270
3	0×230
4	0×1f0
5	0×170
6	0×0f0
7	0×f800

Table 6.11: Fast gain table, fastgain[]

Address	fastgain[address]
0	0×080
1	0×100
2	0×180
3	0×200
4	0×280
5	0×300
6	0×380
7	0×400

Table 6.12: Banding structure tables, bndtab[], bndsz[]

Band No.	bndtab[band]	bndsz[band]	Band No.	bndtab[band]	bndsz[band]
0	0	1	25	25	1
1	1	1	26	26	1
2	2	1	27	27	1
3	3	1	28	28	3
4	4	1	29	31	3
5	5	1	30	34	3
6	6	1	31	37	3
7	7	1	32	40	3
8	8	1	33	43	3
9	9	1	34	46	3
10	10	1	35	49	6
11	11	1	36	55	6
12	12	1	37	61	6
13	13	1	38	67	6
14	14	1	39	73	6
15	15	1	40	79	6
16	16	1	41	85	12
17	17	1	42	97	12
18	18	1	43	109	12
19	19	1	44	121	12
20	20	1	45	133	24
21	21	1	46	157	24
22	22	1	47	181	24
23	23	1	48	205	24
24	24	1	49	229	24

Table 6.13: Bin number to band number table, masktab[bin]
bin = (10 × A) + B

	B = 0	B = 1	B = 2	B = 3	B = 4	B = 5	B = 6	B = 7	B = 8	B = 9
A = 0	0	1	2	3	4	5	6	7	8	9
A = 1	10	11	12	13	14	15	16	17	18	19
A = 2	20	21	22	23	24	25	26	27	28	28

	B = 0	B = 1	B = 2	B = 3	B = 4	B = 5	B = 6	B = 7	B = 8	B = 9
A = 3	28	29	29	29	30	30	30	31	31	31
A = 4	32	32	32	33	33	33	34	34	34	35
A = 5	35	35	35	35	35	36	36	36	36	36
A = 6	36	37	37	37	37	37	37	38	38	38
A = 7	38	38	38	39	39	39	39	39	39	40
A = 8	40	40	40	40	40	41	41	41	41	41
A = 9	41	41	41	41	41	41	41	42	42	42
A = 10	42	42	42	42	42	42	42	42	42	43
A = 11	43	43	43	43	43	43	43	43	43	43
A = 12	43	44	44	44	44	44	44	44	44	44
A = 13	44	44	44	45	45	45	45	45	45	45
A = 14	45	45	45	45	45	45	45	45	45	45
A = 15	45	45	45	45	45	45	45	46	46	46
A = 16	46	46	46	46	46	46	46	46	46	46
A = 17	46	46	46	46	46	46	46	46	46	46
A = 18	46	47	47	47	47	47	47	47	47	47
A = 19	47	47	47	47	47	47	47	47	47	47
A = 20	47	47	47	47	47	48	48	48	48	48
A = 21	48	48	48	48	48	48	48	48	48	48
A = 22	48	48	48	48	48	48	48	48	48	49
A = 23	49	49	49	49	49	49	49	49	49	49
A = 24	49	49	49	49	49	49	49	49	49	49
A = 25	49	49	49	0	0	0				

Table 6.14: Log-addition table, latab[val]
val = (10 × A) + B

	B = 0	B = 1	B = 2	B = 3	B = 4	B = 5	B = 6	B = 7	B = 8	B = 9
A = 0	0×0040	0×003f	0×003e	0×003d	0×003c	0×003b	0×003a	0×0039	0×0038	0×0037
A = 1	0×0036	0×0035	0×0034	0×0034	0×0033	0×0032	0×0031	0×0030	0×002f	0×002f
A = 2	0×002e	0×002d	0×002c	0×002c	0×002b	0×002a	0×0029	0×0029	0×0028	0×0027
A = 3	0×0026	0×0026	0×0025	0×0024	0×0024	0×0023	0×0023	0×0022	0×0021	0×0021
A = 4	0×0020	0×0020	0×001f	0×001e	0×001e	0×001d	0×001d	0×001c	0×001c	0×001b
A = 5	0×001b	0×001a	0×001a	0×0019	0×0019	0×0018	0×0018	0×0017	0×0017	0×0016
A = 6	0×0016	0×0015	0×0015	0×0015	0×0014	0×0014	0×0013	0×0013	0×0013	0×0012
A = 7	0×0012	0×0012	0×0011	0×0011	0×0011	0×0010	0×0010	0×0010	0×000f	0×000f
A = 8	0×000f	0×000e	0×000e	0×000e	0×000d	0×000d	0×000d	0×000d	0×000c	0×000c
A = 9	0×000c	0×000c	0×000b	0×000b	0×000b	0×000b	0×000a	0×000a	0×000a	0×000a
A = 10	0×000a	0×0009	0×0009	0×0009	0×0009	0×0009	0×0008	0×0008	0×0008	0×0008
A = 11	0×0008	0×0008	0×0007	0×0007	0×0007	0×0007	0×0007	0×0007	0×0006	0×0006
A = 12	0×0006	0×0006	0×0006	0×0006	0×0006	0×0006	0×0005	0×0005	0×0005	0×0005
A = 13	0×0005	0×0005	0×0005	0×0005	0×0004	0×0004	0×0004	0×0004	0×0004	0×0004
A = 14	0×0004	0×0004	0×0004	0×0004	0×0004	0×0003	0×0003	0×0003	0×0003	0×0003
A = 15	0×0003	0×0003	0×0003	0×0003	0×0003	0×0003	0×0003	0×0003	0×0003	0×0002
A = 16	0×0002	0×0002	0×0002	0×0002	0×0002	0×0002	0×0002	0×0002	0×0002	0×0002
A = 17	0×0002	0×0002	0×0002	0×0002	0×0002	0×0002	0×0002	0×0002	0×0001	0×0001
A = 18	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001
A = 19	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001
A = 20	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001	0×0001
A = 21	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000
A = 22	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000
A = 23	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000
A = 24	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000
A = 25	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000	0×0000

Table 6.15: Hearing threshold table, hth[fscod][band]

Band No.	hth[0][band] (f_s = 48 kHz)	hth[1][band] (f_s = 44,1 kHz)	hth[2][band] (f_s = 32 kHz)	Band No.	hth[0][band] (f_s = 48 kHz)	hth[1][band] (f_s = 44,1 kHz)	hth[2][band] (f_s = 32 kHz)
0	0×04d0	0×04f0	0×0580	25	0×0340	0×0350	0×0380
1	0×04d0	0×04f0	0×0580	26	0×0330	0×0340	0×0380
2	0×0440	0×0460	0×04b0	27	0×0320	0×0340	0×0370
3	0×0400	0×0410	0×0450	28	0×0310	0×0320	0×0360
4	0×03e0	0×03e0	0×0420	29	0×0300	0×0310	0×0350
5	0×03c0	0×03d0	0×03f0	30	0×02f0	0×0300	0×0340
6	0×03b0	0×03c0	0×03e0	31	0×02f0	0×02f0	0×0330
7	0×03b0	0×03b0	0×03d0	32	0×02f0	0×02f0	0×0320
8	0×03a0	0×03b0	0×03c0	33	0×02f0	0×02f0	0×0310
9	0×03a0	0×03a0	0×03b0	34	0×0300	0×02f0	0×0300
10	0×03a0	0×03a0	0×03b0	35	0×0310	0×0300	0×02f0
11	0×03a0	0×03a0	0×03b0	36	0×0340	0×0320	0×02f0
12	0×03a0	0×03a0	0×03a0	37	0×0390	0×0350	0×02f0
13	0×0390	0×03a0	0×03a0	38	0×03e0	0×0390	0×0300
14	0×0390	0×0390	0×03a0	39	0×0420	0×03e0	0×0310
15	0×0390	0×0390	0×03a0	40	0×0460	0×0420	0×0330
16	0×0380	0×0390	0×03a0	41	0×0490	0×0450	0×0350
17	0×0380	0×0380	0×03a0	42	0×04a0	0×04a0	0×03c0
18	0×0370	0×0380	0×03a0	43	0×0460	0×0490	0×0410
19	0×0370	0×0380	0×03a0	44	0×0440	0×0460	0×0470
20	0×0360	0×0370	0×0390	45	0×0440	0×0440	0×04a0
21	0×0360	0×0370	0×0390	46	0×0520	0×0480	0×0460
22	0×0350	0×0360	0×0390	47	0×0800	0×0630	0×0440
23	0×0350	0×0360	0×0390	48	0×0840	0×0840	0×0450
24	0×0340	0×0350	0×0380	49	0×0840	0×0840	0×04e0

Table 6.16: Bit allocation pointer table, baptab[]

Address	Baptab[address]	Address	baptab[address]
0	0	32	10
1	1	33	10
2	1	34	10
3	1	35	11
4	1	36	11
5	1	37	11
6	2	38	11
7	2	39	12
8	3	40	12
9	3	41	12
10	3	42	12
11	4	43	13
12	4	44	13
13	5	45	13
14	5	46	13
15	6	47	14
16	6	48	14
17	6	49	14
18	6	50	14
19	7	51	14
20	7	52	14
21	7	53	14
22	7	54	14
23	8	55	15
24	8	56	15
25	8	57	15
26	8	58	15
27	9	59	15

Address	Baptab[address]
28	9
29	9
30	9
31	10

Address	baptab[address]
60	15
61	15
62	15
63	15

Table 6.17: Quantizer levels and mantissa bits vs. bap

bap	Quantizer levels	Mantissa bits (group bits/number in group)
0	0	0
1	3	1,67 (5/3)
2	5	2,33 (7/3)
3	7	3
4	11	3,5 (7/2)
5	15	4
6	32	5
7	64	6
8	128	7
9	256	8
10	512	9
11	1 024	10
12	2 048	11
13	4 096	12
14	16 384	14
15	65 536	16

6.3 Quantization and decoding of mantissas

6.3.1 Overview

All mantissas are quantized to a fixed level of precision indicated by the corresponding bap. Mantissas quantized to 15 or fewer levels use symmetric quantization. Mantissas quantized to more than 15 levels use asymmetric quantization which is a conventional two's complement representation.

Some quantized mantissa values are grouped together and encoded into a common codeword. In the case of the 3-level quantizer, 3 quantized values are grouped together and represented by a 5-bit codeword in the data stream. In the case of the 5-level quantizer, 3 quantized values are grouped and represented by a 7-bit codeword. For the 11-level quantizer, 2 quantized values are grouped and represented by a 7-bit codeword.

In the encoder, each transform coefficient (which is always < 1,0) is left justified by shifting its binary representation left the number of times indicated by its exponent (0 to 24 left shifts). The amplified coefficient is then quantized to a number of levels indicated by the corresponding bap.

Table 6.18 indicates which quantizer to use for each bap. If a bap equals 0, no bits are sent for the mantissa. Grouping is used for baps of 1, 2 and 4 (3, 5, and 11 level quantizers.)

Table 6.18: Mapping of bap to quantizer

bap	Quantizer levels	Quantization type	Mantissa bits (qntztab[bap]) (group bits/number in group)
0	0	None	0
1	3	Symmetric	1,67 (5/3)
2	5	Symmetric	2,33 (7/3)
3	7	Symmetric	3
4	11	Symmetric	3,5 (7/2)
5	15	Symmetric	4
6	32	Asymmetric	5
7	64	Asymmetric	6

bap	Quantizer levels	Quantization type	Mantissa bits (qntztab[bap]) (group bits/number in group)
8	128	Asymmetric	7
9	256	Asymmetric	8
10	512	Asymmetric	9
11	1 024	Asymmetric	10
12	2 048	Asymmetric	11
13	4 096	Asymmetric	12
14	16 384	Asymmetric	14
15	65 536	Asymmetric	16

During the decode process, the mantissa data stream is parsed up into single mantissas of varying length, interspersed with groups representing combined coding of either triplets or pairs of mantissas. In the bit stream, the mantissas in each exponent set are arranged in frequency ascending order. However, groups occur at the position of the first mantissa contained in the group. Nothing is unpacked from the bit stream for the subsequent mantissas in the group.

6.3.2 Expansion of mantissas for asymmetric quantization (6 ≤ bap ≤ 15)

For bit allocation pointer array values, $6 \leq bap \leq 15$, asymmetric fractional two's complement quantization is used. Each mantissa, along with its exponent, are the floating point representation of a transform coefficient. The decimal point is considered to be to the left of the MSB; therefore the mantissa word represents the range of:

$$(1,0 - 2^{-qntztab[bap]-1}) \text{ to } -1,0.$$

The mantissa number k, of length qntztab[bap[k]], is extracted from the bit stream. Conversion back to a fixed point representation is achieved by right shifting the mantissa by its exponent. This process is represented by the following formula:

$$transform\_coefficient[k] = mantissa[k] >> exponent[k].$$

No grouping is done for asymmetrically quantized mantissas.

6.3.3 Expansion of mantissas for symmetrical quantization (1 ≤ bap ≤ 5)

For bap values of 1 through 5 ($1 \leq bap \leq 5$), the mantissas are represented by coded values. The coded values are converted to standard 2's complement fractional binary words by a table look-up. The number of bits indicated by a mantissa's bap are extracted from the bit stream and right justified. This coded value is treated as a table index and is used to look up the mantissa value. The resulting mantissa value is right shifted by the corresponding exponent to generate the transform coefficient value.

$$transform\_coefficient[k] = quantization\_table[mantissa\_code[k]] >> exponent[k].$$

The mapping of coded mantissa value into the actual mantissa value is shown in Table 6.19 to Table 6.23.

6.3.4 Dither for zero bit mantissas (bap = 0)

The AC-3 decoder uses random noise (dither) values instead of quantized values when the number of bits allocated to a mantissa is zero (bap = 0). The use of the random value is conditional on the value of dithflag. When the value of dithflag is 1, the random noise value is used. When the value of dithflag is 0, a true zero value is used. There is a dithflag variable for each channel. Dither is applied after the individual channels are extracted from the coupling channel. In this way, the dither applied to each channel's upper frequencies is uncorrelated.

Any reasonably random sequence may be used to generate the dither values. The word length of the dither values is not critical. Eight bits is sufficient. The optimum scaling for the dither words is to take a uniform distribution of values between -1 and +1, and scale this by 0,707, resulting in a uniform distribution between - 0,707 and + 0,707. A scalar of 0,75 is close enough to also be considered optimum. A scalar of 0,5 (uniform distribution between - 0,5 and + 0,5) is also acceptable.

Once a dither value is assigned to a mantissa, the mantissa is right shifted according to its exponent to generate the corresponding transform coefficient.

transform_coefficient[k] = scaled_dither_value >> exponent[k].

Table 6.19: bap = 1 (3-level) quantization

Mantissa code	Mantissa value
0	-2/3
1	0
2	2/3

Table 6.20: bap = 2 (5-level) quantization

Mantissa code	Mantissa value
0	-4/5
1	-2/5
2	0
3	2/5
4	4/5

Table 6.21: bap = 3 (7-level) quantization

Mantissa code	Mantissa value
0	-6/7
1	-4/7
2	-2/7
3	0
4	2/7
5	4/7
6	6/7

Table 6.22: bap = 4 (11-level) quantization

Mantissa code	Mantissa value
0	-10/11
1	-8/11
2	-6/11
3	-4/11
4	-2/11
5	0
6	2/11
7	4/11
8	6/11
9	8/11
10	10/11

Table 6.23: bap = 5 (15-level) quantization

Mantissa code	Mantissa value
0	-14/15
1	-12/15
2	-10/15
3	-8/15
4	-6/15
5	-4/15
6	-2/15
7	0
8	2/15

Mantissa code	Mantissa value
9	4/15
10	6/15
11	8/15
12	10/15
13	12/15
14	14/15

6.3.5 Ungrouping of mantissas

In the case when bap = 1, 2, or 4, the coded mantissa values are compressed further by combining 3 level words and 5 level words into separate groups representing triplets of mantissas, and 11 level words into groups representing pairs of mantissas. Groups are filled in the order that the mantissas are processed. If the number of mantissas in an exponent set does not fill an integral number of groups, the groups are shared across exponent sets. The next exponent set in the block continues filling the partial groups. If the total number of 3 or 5 level quantized transform coefficient derived words are not each divisible by 3, or if the 11 level words are not divisible by 2, the final groups of a block are padded with dummy mantissas to complete the composite group. Dummies are ignored by the decoder. Groups are extracted from the bit stream using the length derived from bap. Three level quantized mantissas (bap = 1) are grouped into triples each of 5 bits. Five level quantized mantissas (bap = 2) are grouped into triples each of 7 bits. Eleven level quantized mantissas (bap = 4) are grouped into pairs each of 7 bits.

Encoder equations

bap = 1:
group_code = 9 × mantissa_code[a] + 3 × mantissa_code[b] + mantissa_code[c];
bap = 2:
group_code = 25 × mantissa_code[a] + 5 × mantissa_code[b] + mantissa_code[c];
bap = 4:
group_code = 11 × mantissa_code[a] + mantissa_code[b].

Decoder equations

bap = 1:
mantissa_code[a] = truncate (group_code / 9);
mantissa_code[b] = truncate ((group_code % 9) / 3);
mantissa_code[c] = (group_code % 9) % 3.
bap = 2:
mantissa_code[a] = truncate (group_code / 25);
mantissa_code[b] = truncate ((group_code % 25) / 5);
mantissa_code[c] = (group_code % 25) % 5.
bap = 4:
mantissa_code[a] = truncate (group_code / 11);
mantissa_code[b] = group_code % 11.

where mantissa a comes before mantissa b, which comes before mantissa c.

6.4 Channel coupling

6.4.1 Overview

If enabled, channel coupling is performed on encode by averaging the transform coefficients across channels that are included in the coupling channel. Each coupled channel has a unique set of coupling coordinates which are used to preserve the high frequency envelopes of the original channels. The coupling process is performed above a coupling frequency that is defined by the cplbegf value.

The decoder converts the coupling channel back into individual channels by multiplying the coupled channel transform coefficient values by the coupling coordinate for that channel and frequency sub-band. An additional processing step occurs for the 2/0 mode. If the phsflginu bit = 1 or the equivalent state is continued from a previous block, then phase restoration bits are sent in the bit stream via phase flag bits. The phase flag bits represent the coupling sub-bands in a frequency ascending order. If a phase flag bit = 1 for a particular sub-band, all the right channel transform coefficients within that coupled sub-band are negated after modification by the coupling coordinate, but before inverse transformation.

6.4.2 Sub-band structure for coupling

Transform coefficients (tc) numbers 37 through 252 are grouped into 18 sub-bands of 12 coefficients each, as shown in Table 6.24. The parameter cplbegf indicates the number of the coupling sub-band which is the first to be included in the coupling process. Below the frequency (or transform coefficient number) indicated by cplbegf all channels are independently coded. Above the frequency indicated by cplbegf, channels included in the coupling process (chincpl[ch] = 1) share the common coupling channel up to the frequency (or tc #) indicated by cplendf. The coupling channel is coded up to the frequency (or tc #) indicated by cplendf, which indicates the last coupling sub-band which is coded. The parameter cplendf is interpreted by adding 2 to its value, so the last coupling sub-band which is coded can range from 2 to 17.

The coupling sub-bands are combined into coupling bands for which coupling coordinates are generated (and included in the bit stream). The coupling band structure is indicated by cplbndstrc[sbnd]. Each bit of the cplbndstrc[] array indicates whether the sub-band indicated by the index is combined into the previous (lower in frequency) coupling band. Coupling bands are thus made from integral numbers of coupling sub-bands (see clause 4.4.3).

Table 6.24: Coupling sub-bands

Coupling sub-band No.	Low tc No.	High tc No.	lf cut-off (kHz) at f_s = 48 kHz	hf cut-off (kHz) at f_s = 48 kHz	lf cut-off (kHz) at f_s = 44,1 kHz	hf cut-off (kHz) at f_s = 44,1 kHz
0	37	48	3,42	4,55	3,14	4,18
1	49	60	4,55	5,67	4,18	5,21
2	61	72	5,67	6,80	5,21	6,24
3	73	84	6,80	7,92	6,24	7,28
4	85	96	7,92	9,05	7,28	8,31
5	97	108	9,05	10,17	8,31	9,35
6	109	120	10,17	11,30	9,35	10,38
7	121	132	11,30	12,42	10,38	11,41
8	133	144	12,42	13,55	11,41	12,45
9	145	156	13,55	14,67	12,45	13,48
10	157	168	14,67	15,80	13,48	14,51
11	169	180	15,80	16,92	14,51	15,55
12	181	192	16,92	18,05	15,55	16,58
13	193	204	18,05	19,17	16,58	17,61
14	205	216	19,17	20,30	17,61	18,65
15	217	228	20,30	21,42	18,65	19,68
16	229	240	21,42	22,55	19,68	20,71
17	241	252	22,55	23,67	20,71	21,75

NOTE 1: f_s: sampling frequency.
NOTE 2: At 32 kHz sample rate the sub-band frequency ranges are 2/3 the values of those for 48 kHz.

6.4.3 Coupling coordinate format

Coupling coordinates exist for each coupling band [bnd] in each channel [ch] which is coupled (chincp[ch]==1). Coupling coordinates are sent in a floating point format. The exponent is sent as a 4-bit value (cplcoexp[ch][bnd]) indicating the number of right shifts which should be applied to the fractional mantissa value. The mantissas are transmitted as 4-bit values (cplcomant[ch][bnd]) which shall be properly scaled before use. Mantissas are unsigned values so a sign bit is not used. Except for the limiting case where the exponent value = 15, the mantissa value is known to be between 0,5 and 1,0. Therefore, when the exponent value < 15, the MSB of the mantissa is always equal to "1" and is not transmitted; the next 4 bits of the mantissa are transmitted. This provides one additional bit of resolution. When the exponent value = 15 the mantissa value is generated by dividing the 4-bit value of cplcomant by 16. When the exponent value is < 15 the mantissa value is generated by adding 16 to the 4-bit value of cplcomant and then dividing the sum by 32.

Coupling coordinate dynamic range is increased beyond what the 4-bit exponent can provide by the use of a per channel 2-bit master coupling coordinate (mstrcplco[ch]) which is used to range all of the coupling coordinates within that channel. The exponent values for each channel are increased by 3 times the value of mstrcplco which applies to that channel. This increases the dynamic range of the coupling coordinates by an additional 54 dB.

The following pseudo code indicates how to generate the coupling coordinate (cplco) for each coupling band [bnd] in each channel [ch].

Pseudo code

```
if (cplcoexp[ch, bnd] == 15)
{
    cplco_temp[ch, bnd] = cplcomant[ch, bnd] / 16 ;
}
else
{
    cplco_temp[ch, bnd] = (cplcomant[ch, bnd] + 16) / 32 ;
}
cplco[ch, bnd] = cplco_temp[ch, bnd]>> (cplcoexp[ch, bnd] + 3 * mstrcplco[ch]) ;
```

Using the cplbndstrc[] array, the values of coupling coordinates which apply to coupling bands are converted (by duplicating values as indicated by values of "1" in cplbandstrc[]) to values which apply to coupling sub-bands.

Individual channel mantissas are then reconstructed from the coupled channel as follows:

Pseudo code

```
for (sbnd = cplbegf; sbnd < 3 + cplendf; sbnd++)
{
    for (bin = 0; bin < 12; bin++)
    {
        chmant[ch, sbnd*12+bin+37] = cplmant[sbnd*12+bin+37] * cplco[ch, sbnd] * 8;
    }
}
```

6.5 Rematrixing

6.5.1 Overview

Rematrixing in AC-3 is a channel combining technique in which sums and differences of highly correlated channels are coded rather than the original channels themselves. That is, rather than code and pack left and right in a two channel coder, one can construct:

$$\text{left}' = 0{,}5 \times (\text{left} + \text{right});$$
$$\text{right}' = 0{,}5 \times (\text{left} - \text{right}).$$

The usual quantization and data packing operations are then performed on left' and right'. Clearly, if the original stereo signal were identical in both channels (i.e. two-channel mono), this technique will result in a left' signal that is identical to the original left and right channels, and a right' signal that is identically zero. As a result, one can code the right' channel with very few bits, and increase accuracy in the more important left' channel.

This technique is especially important for preserving Dolby® Surround compatibility. To see this, consider a two channel mono source signal such as that described above. A Dolby® Pro Logic® decoder will try to steer all in-phase information to the centre channel, and all out-of-phase information to the surround channel. If rematrixing is not active, the Pro Logic® decoder will receive the following signals:

```
Received left  = left + QN1;
Received right = right + QN2.
```

where QN1 and QN2 are independent (i.e. uncorrelated) quantization noise sequences, which correspond to the AC-3 coding algorithm quantization, and are programme dependent. The Pro Logic® decoder will then construct centre and surround channels as:

```
centre   = 0,5 × (left + QN1) + 0,5 × (right + QN2);
surround = 0,5 × (left + QN1) - 0,5 × (right + QN2);  /* ignoring the 90° phase shift */.
```

In the case of the centre channel, QN1 and QN2 add, but remain masked by the dominant signal left + right. In the surround channel, however, left - right cancels to zero, and the surround speakers are left to reproduce the difference in the quantization noise sequences (QN1 - QN2).

If channel rematrixing is active, the centre and surround channels will be more easily reproduced as:

```
centre   = left' + QN1;
surround = right' + QN2.
```

In this case, the quantization noise in the surround channel QN2 is much lower in level, and it is masked by the difference signal, right'.

6.5.2 Frequency band definitions

6.5.2.0 Introduction

In AC-3, rematrixing is performed independently in separate frequency bands. There are four bands with boundary locations dependent on coupling information. The boundary locations are by coefficient bin number, and the corresponding rematrixing band frequency boundaries change with sampling frequency. The tables below indicate the rematrixing band frequencies for sampling rates of 48 kHz and 44,1 kHz. At 32 kHz sampling rate the rematrixing band frequencies are 2/3 the values of those shown for 48 kHz.

6.5.2.1 Coupling not in use

If coupling is not in use (cplinu = 0), then there are 4 rematrixing bands, (nrematbd = 4).

Table 6.25: Rematrix banding Table A

Band No.	Low coefficient No.	High coefficient No.	Low frequency (kHz) f_s = 48 kHz	High frequency (kHz) f_s = 48 kHz	Low frequency (kHz) f_s = 44,1 kHz	High frequency (kHz) f_s = 44,1 kHz
0	13	24	1,17	2,30	1,08	2,11
1	25	36	2,30	3,42	2,11	3,14
2	37	60	3,42	5,67	3,14	5,21
3	61	252	5,67	23,67	5,21	21,75

6.5.2.2 Coupling in use, cplbegf > 2

If coupling is in use (cplinu = 1), and cplbegf > 2, there are 4 rematrixing bands (nrematbd = 4). The last (fourth) rematrixing band ends at the point where coupling begins.

Table 6.26: Rematrixing band Table B

Band No.	Low coefficient No.	High coefficient No.	Low frequency (kHz) f_s = 48 kHz	High frequency (kHz) f_s = 48 kHz	Low frequency (kHz) f_s = 44,1 kHz	High frequency (kHz) f_s = 44,1 kHz
0	13	24	1,17	2,30	1,08	2,11
1	25	36	2,30	3,42	2,11	3,14
2	37	60	3,42	5,67	3,14	5,21
3	61	A	5,67	B	5,21	C
A = 36 + cplbegf × 12			B = (A + 1/2) × 0,09375 kHz		C = (A + 1/2) × 0,08613 kHz	

6.5.2.3 Coupling in use, 2 ≥ cplbegf > 0

If coupling is in use (cplinu = 1), and 2 ≥ cplbegf > 0, there are 3 rematrixing bands (nrematbd = 3). The last (third) rematrixing band ends at the point where coupling begins.

Table 6.27: Rematrixing band Table C

Band No.	Low coefficient No.	High coefficient No.	Low frequency (kHz) f_s = 48 kHz	High frequency (kHz) f_s = 48 kHz	Low frequency (kHz) f_s = 44,1 kHz	High frequency (kHz) f_s = 44,1 kHz
0	13	24	1,17	2,30	1,08	2,11
1	25	36	2,30	3,42	2,11	3,14
2	37	A	3,42	B	3,14	C
A = 36 + cplbegf × 12			B = (A + 1/2) × 0,09375 kHz		C = (A + 1/2) × 0,08613 kHz	

6.5.2.4 Coupling in use, cplbegf = 0

If coupling is in use (cplinu = 1), and cplbegf = 0, there are 2 rematrixing bands (nrematbd = 2).

Table 6.28: Rematrixing band Table D

Band No.	Low coefficient No.	High coefficient No.	Low frequency (kHz) f_s = 48 kHz	High frequency (kHz) f_s = 48 kHz	Low frequency (kHz) f_s = 44,1 kHz	High frequency (kHz) f_s = 44,1 kHz
0	13	24	1,17	2,30	1,08	2,11
1	25	36	2,30	3,42	2,11	3,14

6.5.3 Encoding technique

If the 2/0 mode is selected, then rematrixing is employed by the encoder. The squares of the transform coefficients are summed up over the previously defined rematrixing frequency bands for the following combinations:
L, R, L + R, L - R.

Pseudo code

```
if (minimum sum for a rematrixing sub-band n is L or R)
{
    the variable rematflg[n] = 0;
    transmitted left = input L;
    transmitted right = input R;
}
if (minimum sum for a rematrixing sub-band n is L+R or L-R)
{
    the variable rematflg[n] = L;
    transmitted left = 0,5* input (L+R);
    transmitted right = 0,5* input (L-R);
}
```

This selection of matrix combination is done on a block by block basis. The remaining encoder processing of the transmitted left and right channels is identical whether or not the rematrixing flags are 0 or 1.

6.5.4 Decoding technique

For each rematrixing band, a single bit (the rematrix flag) is sent in the data stream, indicating whether or not the two channels have been rematrixed for that band. If the bit is clear, no further operation is required. If the bit is set, the AC-3 decoder performs the following operation to restore the individual channels:

left(band n) = received left(band n) + received right(band n);

right(band n) = received left(band n) - received right(band n).

Note that if coupling is not in use, the two channels may have different bandwidths. As such, rematrixing is only applied up to the lower bandwidth of the two channels. Regardless of the actual bandwidth, all four rematrixing flags are sent in the data stream (assuming the rematrixing strategy bit is set).

6.6 Dialogue normalization

The AC-3 syntax provides elements which allow the encoded bit stream to satisfy listeners in many different situations. The dialnorm element allows for uniform reproduction of spoken dialogue when decoding any AC-3 bit stream.

When audio from different sources is reproduced, the apparent loudness often varies from source to source. The different sources of audio might be different programme segments during a broadcast (i.e. the movie vs. a commercial message); different broadcast channels; or different media (disc vs. tape). The AC-3 coding technology solves this problem by explicitly coding an indication of loudness into the AC-3 bit stream.

The subjective level of normal spoken dialogue is used as a reference. The 5-bit dialogue normalization word which is contained in BSI, dialnorm, is an indication of the subjective loudness of normal spoken dialogue compared to digital 100 %. The 5-bit value is interpreted as an unsigned integer (most significant bit transmitted first) with a range of possible values from 1 to 31. The unsigned integer indicates the headroom in dB above the subjective dialogue level. This value can also be interpreted as an indication of how many dB the subjective dialogue level is below digital 100 %.

The dialnorm value may be used directly by the AC-3 decoder, or the value may be used by the section of the sound reproduction system responsible for setting the reproduction volume, e.g. the system volume control. The system volume control is generally set based on listener input as to the desired loudness, or sound pressure level (SPL). The listener adjusts a volume control which generally directly adjusts the reproduction system gain. With AC-3 and the dialnorm value, the reproduction system gain becomes a function of both the listeners desired reproduction sound pressure level for dialogue, and the dialnorm value which indicates the level of dialogue in the audio signal. The listener is thus able to reliably set the volume level of dialogue, and the subjective level of dialogue will remain uniform no matter which AC-3 programme is decoded.

EXAMPLE: The listener adjusts the volume control to 67 dB. (With AC-3 dialogue normalization, it is possible to calibrate a system volume control directly in sound pressure level, and the indication will be accurate for any AC-3 encoded audio source). A high quality entertainment programme is being received, and the AC-3 bit stream indicates that dialogue level is 25 dB below 100 % digital level. The reproduction system automatically sets the reproduction system gain so that full scale digital signals reproduce at a sound pressure level of 92 dB. The spoken dialogue (down 25 dB) will thus reproduce at 67 dB SPL.

The broadcast programme cuts to a commercial message, which has dialogue level at -15 dB with respect to 100 % digital level. The system level gain automatically drops, so that digital 100 % is now reproduced at 82 dB SPL. The dialogue of the commercial (down 15 dB) reproduces at a 67 dB SPL, as desired.

In order for the application of dialogue normalization by the system volume control to work, the dialnorm value needs to be communicated from the AC-3 decoder to the system gain controller so that dialnorm can interact with the listener adjusted volume control. The listener-selected volume setting and the dialnorm value shall be combined in order to adjust the final reproduction system gain.

Adjustment of the system volume control is not an AC-3 function. The AC-3 bit stream simply conveys useful information which allows the system volume control to be implemented in a way which automatically removes undesirable level variations between programme sources. It is mandatory that the dialnorm value and the user selected volume setting both be used to set the reproduction system gain.

6.7 Dynamic range compression

6.7.1 Overview

A consistent problem in the delivery of audio programming is that different members of the audience wish to enjoy different amounts of dynamic range. Original high quality programming (such as feature films) are typically mixed with quite a wide dynamic range. Using dialogue as a reference, loud sounds like explosions are often 20 dB or more louder, and faint sounds like leaves rustling may be 50 dB quieter. In many listening situations it is objectionable to allow the sound to become very loud, and thus the loudest sounds need to be compressed downwards in level. Similarly, in many listening situations the very quiet sounds would be inaudible, and need to be brought upwards in level to be heard. Since most of the audience will benefit from a limited programme dynamic range, soundtracks which have been mixed with a wide dynamic range are generally compressed: the dynamic range is reduced by bringing down the level of the loud sounds and bringing up the level of the quiet sounds. While this satisfies the needs of much of the audience, it removes the ability of some in the audience to experience the original sound programme in its intended form. The AC-3 audio coding technology solves this conflict by allowing dynamic range control values to be placed into the AC-3 bit stream.

6.7.2 Dynamic range control; dynrng, dynrng2

6.7.2.1 Overview

The dynrng element allows the programme provider to implement subjectively pleasing dynamic range reduction for most of the intended audience, while allowing individual members of the audience the option to experience more (or all) of the original dynamic range. The dynamic range control values, dynrng, indicate a gain change to be applied in the decoder in order to implement dynamic range compression. Each dynrng value can indicate a gain change of ±24 dB. The sequence of dynrng values is a compression control signal. An AC-3 encoder (or a bit stream processor) will generate the sequence of dynrng values. Each value is used by the AC-3 decoder to alter the gain of one or more audio blocks. The dynrng values typically indicate gain reduction during the loudest signal passages, and gain increases during the quiet passages. For the listener, it is desirable to bring the loudest sounds down in level towards dialogue level, and the quiet sounds up in level, again towards dialogue level. Sounds which are at the same loudness as the normal spoken dialogue will typically not have their gain changed.

The compression is actually applied to the audio in the AC-3 decoder. The encoded audio has full dynamic range. It is permissible for the AC-3 decoder to (optionally, under listener control) ignore the dynrng values in the bit stream. This will result in the full dynamic range of the audio being reproduced. It is also permissible (again under listener control) for the decoder to use some fraction of the dynrng control value, and to use a different fraction of positive or negative values. The AC-3 decoder can thus reproduce either fully compressed audio (as intended by the compression control circuit in the AC-3 encoder); full dynamic range audio; or audio with partially compressed dynamic range, with different amounts of compression for high level signals and low level signals.

EXAMPLE: A feature film soundtrack is encoded into AC-3. The original programme mix has dialogue level at -25 dB. Explosions reach full scale peak level of 0 dB. Some quiet sounds which are intended to be heard by all listeners are 50 dB below dialogue level (or -75 dB). A compression control signal (sequence of dynrng values) is generated by the AC-3 encoder. During those portions of the audio programme where the audio level is higher than dialogue level the dynrng values indicate negative gain, or gain reduction. For full scale 0 dB signals (the loudest explosions), gain reduction of -15 dB is encoded into dynrng. For very quiet signals, a gain increase of 20 dB is encoded into dynrng.

A listener wishes to reproduce this soundtrack quietly so as not to disturb anyone, but wishes to hear all of the intended programme content. The AC-3 decoder is allowed to reproduce the default, which is full compression. The listener adjusts dialogue level to 60 dB SPL. The explosions will only go as loud as 70 dB (they are 25 dB louder than dialogue but get -15 dB of gain applied), and the quiet sounds will reproduce at 30 dB SPL (20 dB of gain is applied to their original level of 50 dB below dialogue level). The reproduced dynamic range will be 70 dB - 30 dB = 40 dB.

The listening situation changes, and the listener now wishes to raise the reproduction level of dialogue to 70 dB SPL, but still wishes to limit how loud the programme plays. Quiet sounds may be allowed to play as quietly as before. The listener instructs the AC-3 decoder to continue using the dynrng values which indicate gain reduction, but to attenuate the values which indicate gain increases by a factor of 1/2. The explosions will still reproduce 10 dB above dialogue level, which is now 80 dB SPL. The quiet sounds are now increased in level by 20 dB / 2 = 10 dB. They will now be reproduced 40 dB below dialogue level, at 30 dB SPL. The reproduced dynamic range is now 80 dB - 30 dB = 50 dB.

Another listener wishes the full original dynamic range of the audio. This listener adjusts the reproduced dialogue level to 75 dB SPL, and instructs the AC-3 decoder to ignore the dynamic range control signal. For this listener the quiet sounds reproduce at 25 dB SPL, and the explosions hit 100 dB SPL. The reproduced dynamic range is 100 dB - 25 dB = 75 dB. This reproduction is exactly as intended by the original programme producer.

In order for this dynamic range control method to be effective, it should be used by all programme providers. Since all broadcasters wish to supply programming in the form that is most usable by their audience, nearly all broadcasters will apply dynamic range compression to any audio programme which has a wide dynamic range. This compression is not reversible unless it is implemented by the technique embedded in AC-3. If broadcasters make use of the embedded AC-3 dynamic range control system, then listeners can have some control over their reproduced dynamic range. Broadcasters need to be confident that the compression characteristic that they introduce into AC-3 will, by default, be heard by the listeners. Therefore, the AC-3 decoder shall, by default, implement the compression characteristic indicated by the dynrng values in the data stream. AC-3 decoders may optionally allow listener control over the use of the dynrng values, so that the listener may select full or partial dynamic range reproduction.

6.7.2.2 Detailed implementation

The dynrng field in the AC-3 bit stream is 8 bits in length. In the case that acmod = 0 (1 + 1 mode, or 2 completely independent channels) dynrng applies to the first channel (Ch1), and dynrng2 applies to the second channel (Ch2). While dynrng is described below, dynrng2 is handled identically. The dynrng value may be present in any audio block. When the value is not present, the value from the previous block is used, except for block 0. In the case of block 0, if a new value of dynrng is not present, then a value of 0000 0000 should be used. The most significant bit of dynrng (and of dynrng2) is transmitted first. The first three bits indicate gain changes in 6,02 dB increments which can be implemented with an arithmetic shift operation. The following five bits indicate linear gain changes, and require a 6-bit multiply. The 3 and 5 bit fields of dynrng can be represented as follows:

$$X_0\ X_1\ X_2\ .\ Y_3\ Y_4\ Y_5\ Y_6\ Y_7$$

The meaning of the X values is most simply described by considering X to represent a 3-bit signed integer with values from - 4 to 3. The gain indicated by X is then $(X + 1) \times 6{,}02$ dB. Table 6.29 shows this in detail.

Table 6.29: Meaning of 3 MSB of dynrng

X_0	X_1	X_2	Integer value	Gain indicated (dB)	Arithmetic shifts
0	1	1	3	+24,08	4 left
0	1	0	2	+18,06	3 left
0	0	1	1	+12,04	2 left
0	0	0	0	+6,02	1 left
1	1	1	-1	0	None
1	1	0	-2	-6,02	1 right
1	0	1	-3	-12,04	2 right
1	0	0	-4	-18,06	3 right

The value of Y is a linear representation of a gain change of up to -6 dB. Y is considered to be an unsigned fractional integer, with a leading value of 1, or: 0,1 $Y_3\ Y_4\ Y_5\ Y_6\ Y_7$ (base 2). Y can represent values between $0{,}111111_2$ (or 63/64) and $0{,}100000_2$ (or 1/2). Thus, Y can represent gain changes from -0,14 dB to -6,02 dB.

The combination of X and Y values allows dynrng to indicate gain changes from 24,08 - 0,14 = +23,94 dB, to -18,06 - 6 = -24,06 dB. The bit code of 0000 0000 indicates 0 dB (unity) gain.

The dynrng value may be operated on in order to make it represent a gain change which is a fraction of the original value. In order to alter the amount of compression which will be applied, consider the dynrng to represent a signed fractional number, or:

$$X_0 \cdot X_1\, X_2\, Y_3\, Y_4\, Y_5\, Y_6\, Y_7$$

where X_0 is the sign bit and $X_1\, X_2\, Y_3\, Y_4\, Y_5\, Y_6\, Y_7$ are a 7-bit fraction. This 8 bit signed fractional number may be multiplied by a scale factor indicating the fraction of the original compression to apply. If this value is multiplied by 1/2, then the compression range of ±24 dB will be reduced to ±12 dB. After the multiplicative scaling, the 8-bit result is once again considered to be of the original form $X_0\, X_1\, X_2\, .\, Y_3\, Y_4\, Y_5\, Y_6\, Y_7$ and used normally.

6.7.3 Heavy compression; compr, compr2

6.7.3.1 Overview

Some products which decode the AC-3 bit stream will need to deliver the resulting audio via a link with very restricted dynamic range. One example is the case of a television signal decoder which modulates the received picture and sound onto a RF channel in order to deliver a signal usable by a low cost television receiver. In this situation, it is necessary to restrict the maximum peak output level to a known value with respect to dialogue level, in order to prevent overmodulation. Most of the time, the dynamic range control signal, dynrng, will produce adequate gain reduction so that the absolute peak level will be constrained. However, since the dynamic range control system is intended to implement a subjectively pleasing reduction in the range of perceived loudness, there is no assurance that it will control instantaneous signal peaks adequately to prevent overmodulation.

In order to allow the decoded AC-3 signal to be constrained in peak level, a second control signal, compr, (compr2 for Ch2 in 1 + 1 mode) may be present in the AC-3 data stream. The compr element allows the programme provider (or broadcaster) to implement a large dynamic range reduction (heavy compression) in a way which assures that a monophonic downmix will not exceed a certain peak level. The heavily compressed audio programme may be desirable for certain listening situations such as movie delivery to a hotel room, or to an airline seat. The peak level limitation is useful when, for instance, a monophonic downmix will feed an RF modulator and overmodulation needs to be avoided. This control signal should be present in all bit streams which are intended to be receivable by, for instance, a television set top decoder. The compr control signal is similar to the dynrng control signal in that it is used by the decoder to alter the reproduced audio level. The compr control signal has twice the control range as dynrng (±48 dB compared to ±24 dB) with 1/2 the resolution (0,5 dB vs. 0,25 dB). Also, since the compr control signal lives in BSI, it only has a time resolution of an AC-3 syncframe (32 ms @ 48 kHz) instead of a block (5,3 ms @ 48 kHz).

Products which require peak audio level to be constrained should use compr instead of dynrng when compr is present in BSI. Since most of the time the use of dynrng will prevent large peak levels, the AC-3 encoder may only need to insert compr occasionally, i.e. during those instants when the use of dynrng would lead to excessive peak level. If the decoder has been instructed to use compr, and compr is not present for a particular syncframe, then the dynrng control signal shall be used for that syncframe.

In some applications of AC-3, some receivers may wish to reproduce a very restricted dynamic range. In this case, the compr control signal may be present at all times. Then, the use of compr instead of dynrng will allow the reproduction of audio with very limited dynamic range. This might be useful, for instance, in the case of audio delivery to a hotel room or an airplane seat.

6.7.3.2 Detailed implementation

The compr field in the AC-3 bit stream is 8 bits in length. In the case that acmod = 0 (1 + 1 mode, or 2 completely independent channels) compr applies to the first channel (Ch1), and compr2 applies to the second channel (Ch2). While compr is described below (for Ch1), compr2 is handled identically (but for Ch2).

The most significant bit is transmitted first. The first four bits indicate gain changes in 6,02 dB increments which can be implemented with an arithmetic shift operation. The following four bits indicate linear gain changes, and require a 5-bit multiply. The two 4-bit fields of compr can be represented as follows:

$$X_0 \cdot X_1\, X_2\, X_3\, Y_4\, Y_5\, Y_6\, Y_7$$

The meaning of the X values is most simply described by considering X to represent a 4-bit signed integer with values from -8 to +7. The gain indicated by X is then $(X + 1) \times 6{,}02$ dB. Table 6.30 shows this in detail.

The value of Y is a linear representation of a gain change of up to -6 dB. Y is considered to be an unsigned fractional integer, with a leading value of 1, or: 0,1 Y_4 Y_5 Y_6 Y_7 (base 2). Y can represent values between $0{,}11111_2$ (or 31/32) and $0{,}10000_2$ (or 1/2). Thus, Y can represent gain changes from -0,28 dB to -6,02 dB.

The combination of X and Y values allows compr to indicate gain changes from 48,16 - 0,28 = + 47,88 dB, to - 42,14 - 6 = -48,14 dB.

Table 6.30: Meaning of 4 MSB of compr

X_0	X_1	X_2	X_3	Integer value	Gain indicated (dB)	Arithmetic shifts
0	1	1	1	7	+48,16	8 left
0	1	1	0	6	+42,14	7 left
0	1	0	1	5	+36,12	6 left
0	1	0	0	4	+30,10	5 left
0	0	1	1	3	+24,08	4 left
0	0	1	0	2	+18,06	3 left
0	0	0	1	1	+12,04	2 left
0	0	0	0	0	+6,02	1 left
1	1	1	1	-1	0	None
1	1	1	0	-2	-6,02	1 right
1	1	0	1	-3	-12,04	2 right
1	1	0	0	-4	-18,06	3 right
1	0	1	1	-5	-24,08	4 right
1	0	1	0	-6	-30,10	5 right
1	0	0	1	-7	-36,12	6 right
1	0	0	0	-8	-42,14	7 right

6.8 Downmixing

6.8.0 Introduction

In many reproduction systems the number of loudspeakers will not match the number of encoded audio channels. In order to reproduce the complete audio programme downmixing is required. It is important that downmixing be standardized, so that programme providers can be confident of how their programme will be reproduced over systems with various numbers of loudspeakers. With standardized downmixing equations, programme producers can monitor how the downmixed version will sound and make any alterations necessary so that acceptable results are achieved for all listeners. The programme provider can make use of the cmixlev and smixlev syntactical elements in order to affect the relative balance of centre and surround channels with respect to the left and right channels.

Downmixing of the lfe channel is optional. An ideal downmix would have the lfe channel reproduce at an acoustic level of +10 dB with respect to the left and right channels. Since the inclusion of this channel is optional, any downmix coefficient may be used in practice. Care should be taken to assure that loudspeakers are not overdriven by the full scale low frequency content of the lfe channel.

6.8.1 General downmix procedure

The following pseudo code describes how to arrive at un-normalized downmix coefficients. In a practical implementation it may be necessary to then normalize the downmix coefficients in order to prevent any possibility of overload. Normalization is achieved by attenuating all downmix coefficients equally, such that the sum of coefficients used to create any single output channel never exceeds 1.

Pseudo code

```
downmix()
{
    if (acmod == 0) /* 1+1 mode, dual independent mono channels present */
    {
        if (output_nfront == 1) /* 1 front loudspeaker (center) */
        {
            if (dualmode == chan 1) /* Ch1 output requested */
            {
                route left into center;
            }
            else if (dualmode == chan 2) /* Ch2 output requested */
            {
                route right into center;
            }
            else
            {
                mix left into center with -6 dB gain;
                mix right into center with -6 dB gain;
            }
        }
        else if (output_nfront == 2) /* 2 front loudspeakers (left, right) */
        {
            if (dualmode == stereo) /* output of both mono channels requested */
            {
                route left into left;
                route right into right;
            }
            else if (dualmode == chan 1)
            {
                mix left into left with -3 dB gain;
                mix left into right with -3 dB gain;
            }
            else if (dualmode == chan 2)
            {
                mix right into left with -3 dB gain;
                mix right into right with -3 dB gain;
            }
            else /* mono sum of both mono channels requested */
            {
                mix left into left with -6 dB gain;
                mix right into left with -6 dB gain;
                mix left into right with -6 dB gain;
                mix right into right with -6 dB gain;
            }
        }
        else /* output_nfront==3 */
        {
            if (dualmode == stereo)
            {
                route left into left;
                route right into right;
            }
            else if (dualmode == chan 1)
            {
                route left into center;
            }
            else if (dualmode == chan 2)
            {
                route right into center;
            }
            else
            {
                mix left into center with -6 dB gain;
                mix right into center with -6 dB gain;
            }
        }
    }
    else /* acmod > 0 */
    {
        for i = { left, center, right, leftsur/monosur, rightsur }
        {
            if (exists(input_chan[i])) and (exists(output_chan[i]))
            {
                route input_chan[i] into output_chan[i];
            }
        }
```

Pseudo code

```
        if (output_mode == 2/0 Dolby Surround compatible)
        /* 2 ch matrix encoded output requested */
        {
            if (input_nfront != 2)
            {
                mix center into left with -3 dB gain;
                mix center into right with -3 dB gain;
            }
            if (input_nrear == 1)
            {
                mix -mono surround into left with -3 dB gain;
                mix mono surround into right with -3 dB gain;
            }
            else if (input_nrear == 2)
            {
                mix -left surround into left with -3 dB gain;
                mix -right surround into left with -3 dB gain;
                mix left surround into right with -3 dB gain;
                mix right surround into right with -3 dB gain;
            }
        }
        else if (output_mode == 1/0) /* center only */
        {
            if (input_nfront != 1)
            {
                mix left into center with -3 dB gain;
                mix right into center with -3 dB gain;
            }
            if (input_nfront == 3)
            {
                mix center into center using clev and -3 dB gain;
            }
            if (input_nrear == 1)
            {
                mix mono surround into center using slev and -3 dB gain;
            }
            else if (input_nrear == 2)
            {
                mix left surround into center using slev and -3 dB gain;
                mix right surround into center using slev and -3 dB gain;
            }
        }
        else /* more than center output requested */
        {
            if (output_nfront == 2)
            {
                if (input_nfront == 1)
                {
                    mix center into left with -3 dB gain;
                    mix center into right with -3 dB gain;
                }
                else if (input_nfront == 3)
                {
                    mix center into left using clev;
                    mix center into right using clev;
                }
            }
            if (input_nrear == 1) /* single surround channel coded */
            {
                if (output_nrear == 0) /* no surround loudspeakers */
                {
                    mix mono surround into left with slev and -3 dB gain;
                    mix mono surround into right with slev and -3 dB gain;
                }
                else if (output_nrear == 2) /* two surround loudspeaker channels */
                {
                    mix mono srnd into left surround with -3 dB gain;
                    mix mono srnd into right surround with -3 dB gain;
                }
            }
            else if (input_nrear == 2) /* two surround channels encoded */
            {
                if (output_nrear == 0)
                {
                    mix left surround into left using slev;
                    mix right surround into right using slev;
                }
```

Pseudo code

```
                else if (output_nrear == 1) .
                {
                    mix left srnd into mono surround with -3 dB gain;
                    mix right srnd into mono surround with -3 dB gain;
                }
            }
        }
    }
}
```

The actual coefficients used for downmixing will affect the absolute level of the centre channel. If dialogue level is to be established with absolute SPL calibration, this should be taken into account.

6.8.2 Downmixing into two channels

Let L, C, R, L_s, R_s refer to the 5 discrete channels which are to be mixed down to 2 channels. In the case of a single surround channel (*n*/1 modes), S refers to the single surround channel. Two types of downmix should be provided: downmix to a L_tR_t matrix surround-encoded stereo pair; and downmix to a conventional stereo signal, L_0R_0. The downmixed stereo signal (L_0R_0, or L_tR_t) may be further mixed to mono, M, by a simple summation of the 2 channels. If the L_tR_t downmix is combined to mono, the surround information will be lost. The L_0R_0 downmix is preferred when a mono signal is desired. Downmix coefficients shall have relative accuracy of at least ±0,25 dB.

Prior to the scaling needed to prevent overflow, the general 3/2 downmix equations for an L_0R_0 stereo signal are:

$$L_0 = 1{,}0 \times L + clev \times C + slev \times L_S;$$
$$R_0 = 1{,}0 \times R + clev \times C + slev \times R_S.$$

If L_0R_0 are subsequently combined for monophonic reproduction, the effective mono downmix equation becomes:

$$M = 1{,}0 \times L + 2{,}0 \times clev \times C + 1{,}0 \times R + slev \times L_s + slev \times R_s.$$

If only a single surround channel, S, is present (3/1 mode) the downmix equations are:

$$L_0 = 1{,}0 \times L + clev \times C + 0{,}7 \times slev \times S;$$
$$R_0 = 1{,}0 \times R + clev \times C + 0{,}7 \times slev \times S;$$
$$M = 1{,}0 \times L + 2{,}0 \times clev \times C + 1{,}0 \times R + 1{,}4 \times slev \times S.$$

The values of clev and slev are indicated by the cmixlev and surmixlev bit fields in the BSI data, as shown in Table 4.4 and Table 4.5 respectively.

If the cmixlev or surmixlev bit fields indicate the reserved state (value of 1 1), the decoder should use the intermediate coefficient values indicated by the bit field value of 0 1. If the centre channel is missing (2/1 or 2/2 mode), the same equations may be used without the C term. If the surround channels are missing, the same equations may be used without the L_s, R_s, or S terms.

Prior to the scaling needed to prevent overflow, the 3/2 downmix equations for an L_tR_t stereo signal are:

$$L_t = 1{,}0 \times L + 0{,}707 \times C - 0{,}707 \times L_s - 0{,}707 \times R_s;$$
$$R_t = 1{,}0 \times R + 0{,}707 \times C + 0{,}707 \times L_s + 0{,}707 \times R_s.$$

If only a single surround channel, S, is present (3/1 mode) these equations become:

$$L_t = 1{,}0\ L + 0{,}707\ C - 0{,}707\ S;$$
$$R_t = 1{,}0\ R + 0{,}707\ C + 0{,}707\ S.$$

If the centre channel is missing (2/2 or 2/1 mode) the C term is dropped.

The actual coefficients used may need to be scaled downwards so that arithmetic overflow does not occur if all channels contributing to a downmix signal happen to be at full scale. For each audio coding mode, a different number of channels contributes to the downmix, and a different scaling could be used to prevent overflow. For simplicity, the scaling for the worst case may be used in all cases. This minimizes the number of coefficients required. The worst-case scaling occurs when clev and slev are both 0,707. In the case of the L_0R_0 downmix, the sum of the unscaled coefficients is 1 + 0,707 + 0,707 = 2,414, so all coefficients would need to be multiplied by 1/2,414 = 0,4143 (downwards scaling by 7,65 dB). In the case of the L_tR_t downmix, the sum of the unscaled coefficients is 1 + 0,707 + 0,707 + 0,707 = 3,121, so all coefficients would need to be multiplied by 1/3,121 = 0,3204 (downwards scaling by 9,89 dB). The scaled coefficients will typically be converted to binary values with limited wordlength. The 6-bit coefficients shown below have sufficient accuracy.

If a scaled L_0R_0 downmix is desired, scaled coefficient values which correspond to the values of 1,0, 0,707, 0,596, 0,500, 0,354 are needed. These coefficients are specified in Table 6.31.

Table 6.31: L_0R_0 scaled downmix coefficients

Unscaled coefficient	Scaled coefficient	6-bit quantized coefficient	Gain (dB)	Relative gain (dB)	Coefficient error (dB)
1,0	0,414	26/64	-7,8	0,0	-
0,707	0,293	18/64	-11,0	-3,2	-0,2
0,596	0,247	15/64	-12,6	-4,8	+0,3
0,500	0,207	13/64	-13,8	-6,0	0,0
0,354	0,147	9/64	-17,0	-9,2	-0,2

If a scaled L_tR_t downmix is desired, scaled coefficient values which correspond to the values of 1,0 and 0,707 are needed. These coefficients are specified in Table 6.32.

Table 6.32: L_tR_t scaled downmix coefficients

Unscaled coefficient	Scaled coefficient	6-bit quantized coefficient	Gain (dB)	Relative gain (dB)	Coefficient error (dB)
1,0	0,3204	20/64	-10,1	0,0	-
0,707	0,2265	14/64	-13,20	-3,1	-0,10

If it is necessary to implement a mixdown to mono, a further scaling of 1/2 will have to be applied to the L_0R_0 downmix coefficients to prevent overload of the mono sum of $L_0 + R_0$.

6.9 Transform equations and block switching

6.9.1 Overview

The choice of analysis block length is fundamental to any transform-based audio coding system. A long transform length is most suitable for input signals whose spectrum remains stationary, or varies only slowly, with time. A long transform length provides greater frequency resolution, and hence improved coding performance for such signals. On the other hand, a shorter transform length, possessing greater time resolution, is more desirable for signals which change rapidly in time. Therefore, the time vs. frequency resolution trade-off should be considered when selecting a transform block length.

The traditional approach to solving this dilemma is to select a single transform length which provides the best trade-off of coding quality for both stationary and dynamic signals. AC-3 employs a more optimal approach, which is to adapt the frequency/time resolution of the transform depending upon spectral and temporal characteristics of the signal being processed. This approach is very similar to behaviour known to occur in human hearing. In transform coding, the adaptation occurs by switching the block length in a signal dependent manner.

6.9.2 Technique

In the AC-3 transform block switching procedure, a block length of either 512 or 256 samples (time resolution of 10,7 or 5,3 ms for sampling frequency of 48 kHz) can be employed. Normal blocks are of length 512 samples. When a normal windowed block is transformed, the result is 256 unique frequency domain transform coefficients. Shorter blocks are constructed by taking the usual 512 sample windowed audio segment and splitting it into two segments containing 256 samples each. The first half of an MDCT block is transformed separately but identically to the second half of that block. Each half of the block produces 128 unique non-zero transform coefficients representing frequencies from 0 to $f_s/2$, for a total of 256. This is identical to the number of coefficients produced by a single 512 sample block, but with two times improved temporal resolution. Transform coefficients from the two half-blocks are interleaved together on a coefficient-by-coefficient basis to form a single block of 256 values. This block is quantized and transmitted identically to a single long block. A similar, mirror image procedure is applied in the decoder during signal reconstruction.

Transform coefficients for the two 256 length transforms arrive in the decoder interleaved together bin-by-bin. This interleaved sequence contains the same number of transform coefficients as generated by a single 512-sample transform. The decoder processes interleaved sequences identically to non-interleaved sequences, except during the inverse transformation described below.

Prior to transforming the audio signal from time to frequency domain, the encoder performs an analysis of the spectral and/or temporal nature of the input signal and selects the appropriate block length. This analysis occurs in the encoder only, and therefore can be upgraded and improved without altering the existing base of decoders. A one bit code per channel per transform block (blksw[ch]) is embedded in the bit stream which conveys length information: (blksw[ch] = 0 or 1 for 512 or 256 samples, respectively). The decoder uses this information to deformat the bit stream, reconstruct the mantissa data, and apply the appropriate inverse transform equations.

6.9.3 Decoder implementation

TDAC transform block switching is accomplished in AC-3 by making an adjustment to the conventional forward and inverse transformation equations for the 256 length transform. The same window and FFT sine/cosine tables used for 512 sample blocks can be reused for inverse transforming the 256 sample blocks; however, the pre- and post-FFT complex multiplication phase-shift requires an additional 128 table values for the block-switched transform.

Since the input and output arrays for blksw[ch] = 1 are exactly one half of the length of those for blksw = 0, the size of the inverse transform RAM and associated buffers is the same with block switching as without.

The adjustments required for inverse transforming the 256 sample blocks are:

- The input array contains 128 instead of 256 coefficients.
- The IFFT pre- and post-phase-shift use a different cosine table, requiring an additional 128 table values (64 cosine, 64 sine).
- The complex IFFT employs 64 points instead of 128. The same FFT cosine table can be used with subsampling to retrieve only the even numbered entries.
- The input pointers to the IFFT post-windowing operation are initialized to different start addresses, and operate modulo 128 instead of modulo 256.

6.9.4 Transformation equations

6.9.4.1 512-sample IMDCT transform

The following procedure describes the technique used for computing the IMDCT for a single N = 512 length real data block using a single N/4 point complex IFFT with simple pre- and post-phase-shift operations. These are the inverse transform equations used when the blksw flag is set to zero (indicating absence of a transient, and 512 sample transforms).

Step 1: Define the MDCT transform coefficients = X[k], k = 0, 1, ... N/2-1.

Step 2: Pre-IFFT complex multiply step.

Compute N/4-point complex multiplication product Z[k], k = 0, 1, ... N/4-1:

Pseudo code

```
for(k=0; k<N/4; k++)
{
    /* Z[k] = (X[N/2-2*k-1] + j * X[2*k]) * (xcos1[k] + j * xsin1[k]) ; */
    Z[k]=(X[N/2-2*k-1]*xcos1[k]-X[2*k]*xsin1[k])+j*(X[2*k]*xcos1[k]+X[N/2-2*k-1]*xsin1[k]);
}
```

where:

$$xcos1[k] = -\cos(2\pi \times (8 \times k + 1) / (8 \times N));$$
$$xsin1[k] = -\sin(2\pi \times (8 \times k + 1) / (8 \times N)).$$

Step 3: Complex IFFT step.

Compute N/4-point complex IFFT of Z[k] to generate complex-valued sequence z[n]:

Pseudo code

```
for(n=0; n<N/4; n++)
{
    z[n] = 0 ;
    for(k=0; k<N/4; k++)
    {
        z[n] + = Z[k] * (cos(8*pi*k*n/N) + j * sin(8*pi*k*n/N)) ;
    }
}
```

Step 4: Post-IFFT complex multiply step.

Compute N/4-point complex multiplication product y[n], n = 0,1,...N/4-1 as:

Pseudo code

```
for(n=0; n<N/4; n++)
{
    /* y[n] = z[n] * (xcos1[n] + j * xsin1[n]) ; */
    y[n] = (zr[n] * xcos1[n] - zi[n] * xsin1[n]) + j * (zi[n] * xcos1[n] + zr[n] * xsin1[n]) ;
}
```

where:

$$zr[n] = \text{real}(z[n]);$$
$$zi[n] = \text{imag}(z[n]);$$

and $xcos1[n]$ and $xsin1[n]$ are as defined in step 2 above.

Step 5: Windowing and de-interleaving step.

Compute windowed time-domain samples x[n]:

Pseudo code

```
for(n=0; n<N/8; n++)
{
    x[2*n] = -yi[N/8+n] * w[2*n] ;
    x[2*n+1] = yr[N/8-n-1] * w[2*n+1] ;
    x[N/4+2*n] = -yr[n] * w[N/4+2*n] ;
    x[N/4+2*n+1] = yi[N/4-n-1] * w[N/4+2*n+1] ;
    x[N/2+2*n] = -yr[N/8+n] * w[N/2-2*n-1] ;
    x[N/2+2*n+1] = yi[N/8-n-1] * w[N/2-2*n-2] ;
    x[3*N/4+2*n] = yi[n] * w[N/4-2*n-1] ;
    x[3*N/4+2*n+1] = -yr[N/4-n-1] * w[N/4-2*n-2] ;
}
```

where:

$$yr[n] = \text{real}(y[n]);$$

yi[n] = imag(y[n]);

w[n] is the transform window sequence (see Table 6.33).

Step 6: Overlap and add step.

The first half of the windowed block is overlapped with the second half of the previous block to produce PCM samples (the factor of 2 scaling undoes headroom scaling performed in the encoder):

Pseudo code

```
for(n=0; n<N/2; n++)
{
    pcm[n] = 2 * (x[n] + delay[n]) ;
    delay[n] = x[N/2+n) ;
}
```

The arithmetic processing in the overlap/add processing shall use saturation arithmetic to prevent overflow (wraparound). Since the output signal consists of the original signal plus coding error, it is possible for the output signal to exceed 100 % level even though the original input signal was less than or equal to 100 % level.

6.9.4.2 256-sample IMDCT transforms

The following equations should be used for computing the inverse transforms in the case of blksw = 1, indicating the presence of a transient and two 256 sample transforms (N below still equals 512).

Step 1: Define the MDCT transform coefficients = X[k], k = 0,1,...N/2.

Pseudo code

```
for(k=0; k<n/4; k++)
{
    x1[k] = x[2*k];
    x2[k] = x[2*k+1];
}
```

Step 2: Pre-IFFT complex multiply step.

Compute N/8-point complex multiplication products Z1[k] and Z2[k], k = 0,1,...N/8-1.

Pseudo code

```
for(k=0; k<n/8; k++)
{
    /* z1[k] = (x1[n/4-2*k-1] + j * x1[2*k]) * (xcos2[k] + j * xsin2[k]); */
    z1[k]=(x1[n/4-2*k-1]*xcos2[k]-x1[2k]*xsin2[k])+j*(x1[2*k]*xcos2[k]+x1[n/4-2*k-1]*xsin2[k]);
    /* z2[k] = (x2[n/4-2*k-1] + j * x2[2*k]) * (xcos2[k] + j * xsin2[k]); */
    z2[k]=(x2[n/4-2*k-1]*xcos2[k]-x2[2*k]*xsin2[k])+j*(x2[2*k]*xcos2[k]+x2[n/4-2*k-1]*xsin2[k]);
}
```

where:

xcos2[k] = -cos(2pi × (8 × k + 1) / (4 × N)), xsin2(k) = -sin(2pi × (8 × k + 1) / (4 × N)).

Step 3: Complex IFFT step.

Compute N/8-point complex IFFTs of Z1[k] and Z2[k] to generate complex-valued sequences z1[n] and z2[n].

Pseudo code

```
for(n=0; n<n/8; n++)
{
    z1[n] = 0.;
    z2[n] = 0.;
    for(k=0; k<n/8; k++)
    {
        z1[n] + = z1[k] * (cos(16*pi*k*n/n) + j * sin(16*pi*k*n/n));
        z2[n] + = z2[k] * (cos(16*pi*k*n/n) + j * sin(16*pi*k*n/n));
    }
}
```

Step 4: Post-IFFT complex multiply step.

Compute N/8-point complex multiplication products y1[n] and y2[n], n = 0,1,...N/8-1.

Pseudo code

```
for(n=0; n<n/8; n++)
{
    /* y1[n] = z1[n] * (xcos2[n] + j * xsin2[n]); */
    y1[n] = (zr1[n] * xcos2[n] - zi1[n] * xsin2[n]) + j * (zi1[n] * xcos2[n] + zr1[n] * xsin2[n]);
    /* y2[n] = z2[n] * (xcos2[n] + j * xsin2[n]); */
    y2[n] = (zr2[n] * xcos2[n] - zi2[n] * xsin2[n]) + j * (zi2[n] * xcos2[n] + zr2[n] * xsin2[n]);
}
```

where:

zr1[n] = real(z1[n]);

zi1[n] = imag(z1[n]);

zr2[n] = real(z2[n]);

zi2[n] = imag(z2[n]);

and xcos2[n] and xsin2[n] are as defined in Step 2 above.

Step 5: Windowing and de-interleaving step.

Compute windowed time-domain samples x[n].

Pseudo code

```
for(n=0; n<n/8; n++)
{
    x[2*n] = -yi1[n] * w[2*n];
    x[2*n+1] = yr1[n/8-n-1] * w[2*n+1];
    x[n/4+2*n] = -yr1[n] * w[n/4+2*n];
    x[n/4+2*n+1] = yi1[n/8-n-1] * w[n/4+2*n+1];
    x[n/2+2*n] = -yr2[n] * w[n/2-2*n-1];
    x[n/2+2*n+1] = yi2[n/8-n-1] * w[n/2-2*n-2];
    x[3n/4+2*n] = yi2[n] * w[n/4-2*n-1];
    x[3n/4+2*n+1] = -yr2[n/8-n-1] * w[n/4-2*n-2];
}
```

where:

yr1[n] = real(y1[n]);
yi1[n] = imag(y1[n]);
yr2[n] = real(y2[n]);
yi2[n] = imag(y2[n]);

and w[n] is the transform window sequence (see Table 6.33).

Step 6: Overlap and add step.

The first half of the windowed block is overlapped with the second half of the previous block to produce PCM samples (the factor of 2 scaling undoes headroom scaling performed in the encoder):

Pseudo code

```
for(n=0; n<n/2; n++)
{
    pcm[n] = 2 * (x[n] + delay[n]);
    delay[n] = x[n/2+n];
}
```

The arithmetic processing in the overlap/add processing shall use saturation arithmetic to prevent overflow (wrap around). Since the output signal consists of the original signal plus coding error, it is possible for the output signal to exceed 100 % level even though the original input signal was less than or equal to 100 % level.

Table 6.33: Transform window sequence (w[addr]), with addr = (10 × A) + B

	B = 0	B = 1	B = 2	B = 3	B = 4	B = 5	B = 6	B = 7	B = 8	B = 9
A = 0	0,00014	0,00024	0,00037	0,00051	0,00067	0,00086	0,00107	0,00130	0,00157	0,00187
A = 1	0,00220	0,00256	0,00297	0,00341	0,00390	0,00443	0,00501	0,00564	0,00632	0,00706
A = 2	0,00785	0,00871	0,00962	0,01061	0,01166	0,01279	0,01399	0,01526	0,01662	0,01806
A = 3	0,01959	0,02121	0,02292	0,02472	0,02662	0,02863	0,03073	0,03294	0,03527	0,03770
A = 4	0,04025	0,04292	0,04571	0,04862	0,05165	0,05481	0,05810	0,06153	0,06508	0,06878
A = 5	0,07261	0,07658	0,08069	0,08495	0,08935	0,09389	0,09859	0,10343	0,10842	0,11356
A = 6	0,11885	0,12429	0,12988	0,13563	0,14152	0,14757	0,15376	0,16011	0,16661	0,17325
A = 7	0,18005	0,18699	0,19407	0,20130	0,20867	0,21618	0,22382	0,23161	0,23952	0,24757
A = 8	0,25574	0,26404	0,27246	0,28100	0,28965	0,29841	0,30729	0,31626	0,32533	0,33450
A = 9	0,34376	0,35311	0,36253	0,37204	0,38161	0,39126	0,40096	0,41072	0,42054	0,43040
A = 10	0,44030	0,45023	0,46020	0,47019	0,48020	0,49022	0,50025	0,51028	0,52031	0,53033
A = 11	0,54033	0,55031	0,56026	0,57019	0,58007	0,58991	0,59970	0,60944	0,61912	0,62873
A = 12	0,63827	0,64774	0,65713	0,66643	0,67564	0,68476	0,69377	0,70269	0,71150	0,72019
A = 13	0,72877	0,73723	0,74557	0,75378	0,76186	0,76981	0,77762	0,78530	0,79283	0,80022
A = 14	0,80747	0,81457	0,82151	0,82831	0,83496	0,84145	0,84779	0,85398	0,86001	0,86588
A = 15	0,87160	0,87716	0,88257	0,88782	0,89291	0,89785	0,90264	0,90728	0,91176	0,91610
A = 16	0,92028	0,92432	0,92822	0,93197	0,93558	0,93906	0,94240	0,94560	0,94867	0,95162
A = 17	0,95444	0,95713	0,95971	0,96217	0,96451	0,96674	0,96887	0,97089	0,97281	0,97463
A = 18	0,97635	0,97799	0,97953	0,98099	0,98236	0,98366	0,98488	0,98602	0,98710	0,98811
A = 19	0,98905	0,98994	0,99076	0,99153	0,99225	0,99291	0,99353	0,99411	0,99464	0,99513
A = 20	0,99558	0,99600	0,99639	0,99674	0,99706	0,99736	0,99763	0,99788	0,99811	0,99831
A = 21	0,99850	0,99867	0,99882	0,99895	0,99908	0,99919	0,99929	0,99938	0,99946	0,99953
A = 22	0,99959	0,99965	0,99969	0,99974	0,99978	0,99981	0,99984	0,99986	0,99988	0,99990
A = 23	0,99992	0,99993	0,99994	0,99995	0,99996	0,99997	0,99998	0,99998	0,99998	0,99999
A = 24	0,99999	0,99999	0,99999	1,00000	1,00000	1,00000	1,00000	1,00000	1,00000	1,00000
A = 25	1,00000	1,00000	1,00000	1,00000	1,00000	1,00000				

6.9.5 Channel gain range code

When the signal level is low, the dynamic range of the decoded audio is typically limited by the wordlength used in the transform computation. The use of longer wordlength improves dynamic range but increases cost, as the wordlength of both the arithmetic units and the working RAM needs to be increased. In order to allow the wordlength of the transform computation to be reduced, the AC-3 bit stream includes a syntactic element gainrng[ch]. This 2-bit element exists for each encoded block for each channel.

The gainrng element is a value in the range of 0 - 3. The value is an indication of the maximum sample level within the coded block. Each block represents 256 new audio samples and 256 previous audio samples. Prior to the application of the 512 point window, the maximum absolute value of the 512 PCM values is determined. Based on the maximum value within the block, the value of gainrng is set as indicated in table 6.33a.

Table 6.33a

Maximum absolute value (max)	gainrng
max ≥ 0,5	0
0,5 > max ≥ 0,25	1
0,25 > max ≥ 0,125	2
0,125 > max	3

If the encoder does not perform the step of finding the maximum absolute value within each block then the value of gainrng should be set to 0.

The decoder may use the value of gainrng to pre-scale the transform coefficients prior to the transform and to post-scale the values after the transform. With careful design, the post-scaling process can be performed right at the PCM output stage allowing a 16-bit output buffer RAM to provide 18-bit dynamic range audio.

6.10 Error detection

6.10.0 Introduction

There are several ways in which the AC-3 data may determine that errors are contained within a syncframe of data. The decoder may be informed of that fact by the transport system which has delivered the data. The data integrity may be checked using the embedded CRCs. Also, some simple consistency checks on the received data can indicate that errors are present. The decoder strategy when errors are detected is user definable. Possible responses include muting, block repeats, or syncframe repeats. The amount of error checking performed, and the behaviour in the presence of errors are not specified in the present document, but are left to the application and implementation.

6.10.1 CRC checking

Each AC-3 syncframe contains two 16-bit CRC words. crc1 is the second 16-bit word of the syncframe, immediately following the sync word. crc2 is the last 16-bit word of the syncframe, immediately preceding the sync word of the following syncframe. crc1 applies to the first 5/8 of the syncframe, not including the sync word. crc2 provides coverage for the last 3/8 of the syncframe as well as for the entire syncframe (not including the sync word). Decoding of CRC word(s) allows errors to be detected.

The following generator polynomial is used to generate each of the 16-bit CRC words: $x^{16} + x^{15} + x^2 + 1$.

The 5/8 of a syncframe is defined in Table 6.34, and may be calculated by:

5/8_framesize = truncate(framesize ÷ 2) + truncate(framesize / 8); or

5/8_framesize = (int) (framesize>>1) + (int) (framesize>>3);

where framesize is in units of 16-bit words. Table 6.34 shows the value of 5/8 of the frame size as a function of AC-3 bit rate and audio sample rate.

The CRC calculation may be implemented by one of several standard techniques. A convenient hardware implementation is a linear feedback shift register (LFSR). An example of an LFSR circuit for the above generator polynomial is given in Figure 6.1.

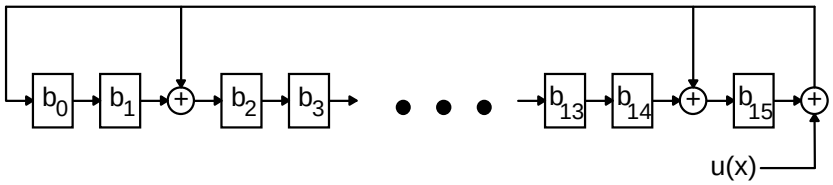

Figure 6.1: LFSR Circuit for Generator Polynomial

Checking for valid CRC with the above circuit consists of resetting all registers to zero, and then shifting the AC-3 data bits serially into the circuit in the order in which they appear in the data stream. The sync word is not covered by either CRC (but is included in the indicated 5/8_framesize) so it should not be included in the CRC calculation. crc1 is considered valid if the above register contains all zeros after the first 5/8 of the syncframe has been shifted in. If the calculation is continued until all data in the syncframe has been shifted through, and the value is again equal to zero, then crc2 is considered valid. Some decoders may choose to only check crc2, and not check for a valid crc1 at the 5/8 point in the syncframe. If crc1 is invalid, it is possible to reset the registers to zero and then check crc2. If crc2 then checks, then the last 3/8 of the syncframe is probably error free. This is of little utility however, since if errors are present in the initial 5/8 of a syncframe it is not possible to decode any audio from the syncframe even if the final 3/8 is error free.

Note that crc1 is generated by encoders such that the CRC calculation will produce zero at the 5/8 point in the syncframe. It is *not* the value generated by calculating the CRC of the first 5/8 of the syncframe using the above generator polynomial. Therefore, decoders should not attempt to save crc1, calculate the CRC for the first 5/8 of the syncframe, and then compare the two.

Syntactical block size restrictions within each syncframe (enforced by encoders), guarantee that blocks 0 and 1 are completely covered by crc1. Therefore, decoders may immediately begin processing block 0 when the 5/8 point in the data syncframe is reached. This may allow smaller input buffers in some applications. Decoders that are able to store an entire syncframe may choose to process only crc2. These decoders would not begin processing block 0 of a syncframe until the entire syncframe is received.

Table 6.34: 5/8_framesize table; number of words in the first 5/8 of the syncframe

frmsizecod	Nominal bit-rate (kbit/s)	fs = 32 kHz 5/8_framesize	fs = 44,1 kHz 5/8_framesize	fs = 48 kHz 5/8_framesize
000000 (0)	32	60	42	40
000001 (0)	32	60	43	40
000010 (1)	40	75	53	50
000011 (1)	40	75	55	50
000100 (2)	48	90	65	60
000101 (2)	48	90	65	60
000110 (3)	56	105	75	70
000111 (3)	56	105	76	70
001000 (4)	64	120	86	80
001001 (4)	64	120	87	80
001010 (5)	80	150	108	100
001011 (5)	80	150	108	100
001100 (6)	96	180	130	120
001101 (6)	96	180	130	120
001110 (7)	112	210	151	140
001111 (7)	112	210	152	140
010000 (8)	128	240	173	160
010001 (8)	128	240	173	160
010010 (9)	160	300	217	200
010011 (9)	160	300	217	200
010100 (10)	192	360	260	240
010101 (10)	192	360	261	240
010110 (11)	224	420	303	280
010111 (11)	224	420	305	280
011000 (12)	256	480	347	320
011001 (12)	256	480	348	320
011010 (13)	320	600	435	400
011011 (13)	320	600	435	400
011100 (14)	384	720	521	480
011101 (14)	384	720	522	480
011110 (15)	448	840	608	560
011111 (15)	448	840	610	560
100000 (16)	512	960	696	640
100001 (16)	512	960	696	640
100010 (17)	576	1 080	782	720
100011 (17)	576	1 080	783	720
100100 (18)	640	1 200	870	800
100101 (18)	640	1 200	871	800

6.10.2 Checking bit stream consistency

It is always possible that an AC-3 syncframe could have valid sync information and valid CRCs, but otherwise be undecodable. This condition may arise if a syncframe is corrupted such that the CRC word is nonetheless valid, or in the case of an encoder error (bug). One safeguard against this is to perform some error checking tests within the AC-3 decoder and bit stream parser. Despite its coding efficiency, there are some redundancies inherent in the AC-3 bit stream. If the AC-3 bit stream contains errors, a number of illegal syntactical constructions are likely to arise. Performing checks for these illegal constructs will detect a great many significant error conditions.

The following is a list of known bit stream error conditions. In some implementations it may be important that the decoder be able to benignly deal with these errors. Specifically, decoders may wish to ensure that these errors do not cause reserved memory to be overwritten with invalid data, and do not cause processing delays by looping with illegal loop counts. Invalid audio reproduction may be allowable, so long as system stability is preserved:

1) (blknum == 0) &&
(cplstre == 0);

2) (cplinu == 1) &&
(fewer than two channels in coupling);

3) (cplinu == 1) &&
(cplbegf > (cplendf+2));

4) (cplinu == 1) &&
((blknum == 0) || (previous cplinu == 0)) &&
(chincpl[n] == 1) &&
(cplcoe[n] == 0);

5) (blknum == 0) &&
(acmod == 2) &&
(rematstr == 0);

6) (cplinu == 1) &&
((blknum == 0) || (previous cplinu == 0)) &&
(cplexpstr == 0);

7) (cplinu == 1) &&
((cplbegf != previous cplbegf) || (cplendf != previous cplendf)) &&
(cplexpstr == 0);

8) (blknum == 0) &&
(chexpstr[n] == 0);

9) (nchmant[n] != previous nchmant[n]) &&
(chexpstr[n] == 0);

10) (blknum == 0) &&
(lfeon == 1) &&
(lfeexpstr == 0);

11) (chincpl[n] == 0) &&
(chbwcod[n] > 60);

12) (blknum == 0) &&
(baie == 0);

13) (blknum == 0) &&
(snroffste == 0);

14) (blknum == 0) &&
(cplinu == 1) &&
(cplleake == 0);

15) (cplinu == 1) &&
(expanded length of cpl delta bit allocation > 50);

16) expanded length of delta bit allocation[n] > 50;

17) compositely coded 5-level exponent value > 124;

18) compositely coded 3-level mantissa value > 26;

19) compositely coded 5-level mantissa value > 124;

20) compositely coded 11-level mantissa value > 120;

21) bit stream unpacking continues past the end of the syncframe.

Note that some of these conditions (such as numbers 17 to 20) can only be tested for at low-levels within the decoder software, resulting in a potentially significant MIPS impact. So long as these conditions do not affect system stability, they do not need to be specifically prevented.

7 Encoding the AC-3 bit stream

7.1 Introduction

The following clauses provides some guidance on AC-3 encoding. Since AC-3 is specified by the syntax and decoder processing, the encoder is not precisely specified. The only normative requirement on the encoder is that the output bit stream follows the AC-3 syntax. Encoders of varying levels of sophistication may be produced. More sophisticated encoders may offer superior audio performance, and may make operation at lower bit rates acceptable. Encoders are expected to improve over time. All decoders will benefit from encoder improvements. The encoder described in the following clauses, while basic in operation, provides good performance. The description which follows indicates several avenues of potential improvement. A flow diagram of the encoding process is shown in Figure 7.1.

7.2 Summary of the encoding process

7.2.1 Input PCM

7.2.1.1 Input word length

The AC-3 encoder accepts audio in the form of PCM words. The internal dynamic range of AC-3 allows input word lengths of up to 24 bits to be useful.

7.2.1.2 Input sample rate

The input sample rate is locked to the output bit rate so that each AC-3 syncframe contains 1 536 samples of audio per channel. If the input audio is available in a PCM format at a different sample rate than that required, sample rate conversion is performed to conform the sample rate.

7.2.1.3 Input filtering

Individual input channels may be high-pass filtered. Removal of DC components of signals can allow more efficient coding since data rate is not used up encoding DC. However, there is the risk that signals which do not reach 100 % PCM level before high-pass filtering will exceed 100 % level after filtering, and thus be clipped. A typical encoder would high-pass filter the input signals with a single pole filter at 3 Hz.

The lfe channel should be low-pass filtered at 120 Hz. A typical encoder would filter the lfe channel with an 8th order elliptic filter with a cut-off frequency of 120 Hz.

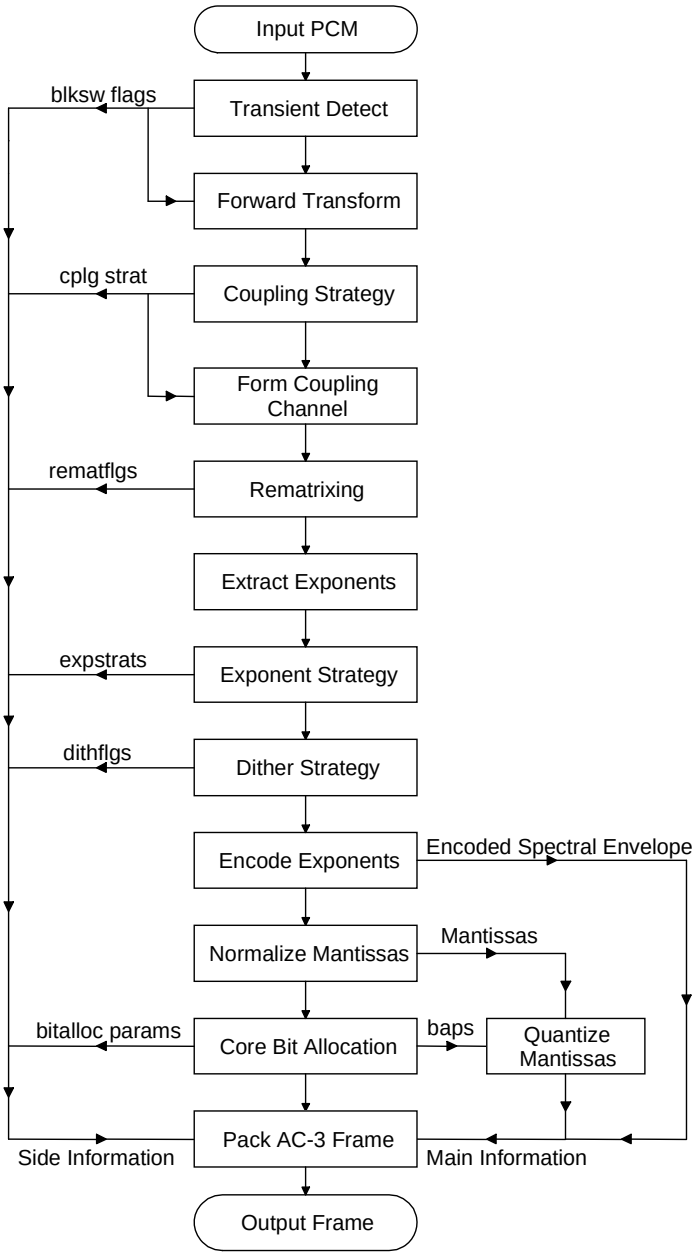

Figure 7.1: Flow diagram of the encoding process

7.2.2 Transient detection

Transients are detected in the full-bandwidth channels in order to decide when to switch to short length audio blocks to improve pre-echo performance. High-pass filtered versions of the signals are examined for an increase in energy from one sub-block time-segment to the next. Sub-blocks are examined at different time scales. If a transient is detected in the second half of an audio block in a channel, that channel switches to a short block. A channel that is block-switched uses the d45 exponent strategy.

The transient detector is used to determine when to switch from a long transform block (length 512), to the short block (length 256). It operates on 512 samples for every audio block. This is done in two passes, with each pass processing 256 samples. Transient detection is broken down into four steps:

- *Step 1:* high-pass filtering;
- *Step 2:* segmentation of the block into submultiples;
- *Step 3:* peak amplitude detection within each sub-block segment; and

- *Step 4:* threshold comparison.

The transient detector outputs a flag blksw[n] for each full-bandwidth channel, which when set to "one" indicates the presence of a transient in the second half of the 512 length input block for the corresponding channel.

Step 1: High-pass filtering:

The high-pass filter is implemented as a cascaded biquad direct form II IIR filter with a cut-off of 8 kHz.

Step 2: Block segmentation:

The block of 256 high-pass filtered samples are segmented into a hierarchical tree of levels in which level 1 represents the 256 length block, level 2 is two segments of length 128, and level 3 is four segments of length 64.

Step 3: Peak detection

The sample with the largest magnitude is identified for each segment on every level of the hierarchical tree. The peaks for a single level are found as follows:

$$P[j][k] = \max(x(n))$$
$$\text{for } n = (512 \times (k-1) / 2j), (512 \times (k-1) / 2j) + 1, \ldots (512 \times k / 2j) - 1$$
$$\text{and } k = 1, \ldots 2^{j}-1;$$

where:

x(n) = the nth sample in the 256 length block
j = 1, 2, 3 is the hierarchical level number
k = the segment number within level j

Note that P[j][0], (i.e. k = 0) is defined to be the peak of the last segment on level j of the tree calculated immediately prior to the current tree. For example, P[3][4] in the preceding tree is P[3][0] in the current tree.

Step 4: Threshold comparison:

The first stage of the threshold comparator checks to see if there is significant signal level in the current block. This is done by comparing the overall peak value P[1][1] of the current block to a "silence threshold". If P[1][1] is below this threshold then a long block is forced. The silence threshold value is 100/32 768. The next stage of the comparator checks the relative peak levels of adjacent segments on each level of the hierarchical tree. If the peak ratio of any two adjacent segments on a particular level exceeds a pre-defined threshold for that level, then a flag is set to indicate the presence of a transient in the current 256 length block. The ratios are compared as follows:

$$\text{mag}(P[j][k]) \times T[j] > \text{mag}(P[j][(k-1)])$$

where:

T[j] is the pre-defined threshold for level j, defined as:
T[1] = 0,1
T[2] = 0,075
T[3] = 0,05

If this inequality is true for any two segment peaks on any level, then a transient is indicated for the first half of the 512 length input block. The second pass through this process determines the presence of transients in the second half of the 512 length input block.

7.2.3 Forward transform

7.2.3.1 Windowing

The audio block is multiplied by a window function to reduce transform boundary effects and to improve frequency selectivity in the filter bank. The values of the window function are included in Table 6.33. Note that the 256 coefficients given are used back-to-back to form a 512-point symmetrical window.

7.2.3.2 Time to frequency transformation

Based on the block switch flags, each audio block is transformed into the frequency domain by performing one $N = 512$ point transform, or two $N = 256$ point transforms. Let $x[n]$ represent the windowed input time sequence. The output frequency sequence, $X_D[k]$ is defined by:

$$X_D[k] = \frac{-2}{N}\sum_{n=0}^{N-1} x[n]\cos\left(\frac{2\pi}{4N}(2n+1)(2k+1) + \frac{\pi}{4}(2k+1)(1+\alpha)\right) \quad \text{for } 0 \le k < N/2$$

where:

- $\alpha = -1$ for the first short transform.
- 0 for the long transform.
- +1 for the second short transform.

7.2.4 Coupling strategy

7.2.4.1 Basic encoder

For a basic encoder, a static coupling strategy may be employed. Suitable coupling parameters are:

```
cplbegf = 6; /*  coupling starts at 10,2 kHz */
cplendf = 12; /*  coupling channel ends at 20,3 kHz */
cplbndstrc  = 0, 0, 1, 1, 0, 1, 1, 1;
cplinu  = 1; /*  coupling always on */
/* all non-block switched channels are coupled */
for(ch=0; ch<nfchans; ch++) if(blksw[ch]) chincpl[ch] = 0; else chincpl[ch] = 1.
```

Coupling coordinates for all channels may be transmitted for every other block, i.e. blocks 0, 2, and 4. During blocks 1, 3, and 5, coupling coordinates are reused.

7.2.4.2 Advanced encoder

More advanced encoders may make use of dynamically variable coupling parameters. The coupling frequencies may be made variable based on bit demand and on a psychoacoustic model which compares the audibility of artefacts caused by bit starvation versus those caused by the coupling process. Channels with a rapidly time varying power level may be removed from coupling. Channels with slowly varying power levels may have their coupling coordinates sent less often. The coupling band structure may be made dynamic.

7.2.5 Form coupling channel

7.2.5.1 Coupling channel

The most basic encoder can form the coupling channel by simply adding all of the individual channel coefficients together, and dividing by 8. The division by 8 prevents the coupling channel from exceeding a value of 1. Slightly more sophisticated encoders can alter the sign of individual channels before adding them into the sum so as to avoid phase cancellations.

7.2.5.2 Coupling coordinates

Coupling coordinates are formed by taking magnitude ratios within of each coupling band. The power in the original channel within a coupling band is divided by the power in the coupling channel within the coupling band, and the square root of this is then computed. This magnitude ratio becomes the coupling coordinate. The coupling coordinates are converted to floating point format and quantized. The exponents for each channel are examined to see if they can be further scaled by 3, 6, or 9. This generates the 2-bit master coupling coordinate for that channel. (The master coupling coordinates allow the dynamic range represented by the coupling coordinate to be increased.)

7.2.6 Rematrixing

Rematrixing is active only in the 2/0 mode. Within each rematrixing band, power measurements are made on the L, R, L+R, and L-R signals. If the maximum power is found in the L or R channels, the rematrix flag is not set for that band. If the maximum power is found in the L+R or L-R signal, then the rematrix flag is set. When the rematrix flag for a band is set, the encoder codes L+R and L-R instead of L and R. Rematrixing is described in clause 6.5.

7.2.7 Extract exponents

The binary representation of each frequency coefficient is examined to determine the number of leading zeros. The number of leading zeros (up to a maximum of 24) becomes the initial exponent value. These exponents are extracted and the exponent sets (one for each block for each channel, including the coupling channel) are used to determine the appropriate exponent strategies.

7.2.8 Exponent strategy

For each channel, the variation in exponents over frequency and time is examined. There is a trade-off between fine frequency resolution, fine time resolution, and the number of bits required to send exponents. In general, when operating at very low bit rates, it is necessary to trade-off time vs. frequency resolution.

In a basic encoder a simple algorithm may be employed. First, look at the variation of exponents over time. When the variation exceeds a threshold new exponents will be sent. The exponent strategy used is made dependent on how many blocks the new exponent set is used for. If the exponents will be used for only a single block, then use strategy d45. If the new exponents will be used for 2 or 3 blocks, then use strategy d25. If the new exponents will be used for 4, 5, or 6 blocks, use strategy d15.

7.2.9 Dither strategy

The encoder controls, on a per channel basis, whether coefficients which will be quantized to zero bits will be reproduced with dither. The intent is to maintain approximately the same energy in the reproduced spectrum even if no bits are allocated to portions of the spectrum. Depending on the exponent strategy, and the accuracy of the encoded exponents, it may be beneficial to defeat dither for some blocks.

A basic encoder can implement a simple dither strategy on a per channel basis. When blksw[ch] is 1, defeat dither for that block and for the following block.

7.2.10 Encode exponents

Based on the selected exponent strategy, the exponents of each exponent set are pre-processed. d25 and d45 exponent strategies require that a single exponent be shared over more than one mantissa. The exponents will be differentially encoded for transmission in the bit stream. The difference between successive raw exponents does not necessarily produce legal differential codes (maximum value of ±2) if the slew rate of the raw exponents is greater than that allowed by the exponent strategy. Pre-processing adjusts exponents so that transform coefficients that share an exponent have the same exponent and so that differentials are legal values. The result of this processing is that some exponents will have their values decreased, and the corresponding mantissas will have some leading zeros.

The exponents are differentially encoded to generate the encoded spectral envelope. As part of the encoder processing, a set of exponents is generated which is equal to the set of exponents which the decoder will have when it decodes the encoded spectral envelope.

7.2.11 Normalize mantissas

Each channel's transform coefficients are normalized by left shifting each coefficient the number of times given by its corresponding exponent to create normalized mantissas. The original binary frequency coefficients are left shifted according to the exponents which the decoder will use. Some of the normalized mantissas will have leading zeros. The normalized mantissas are what are quantized.

7.2.12 Core bit allocation

A basic encoder may use the core bit allocation routine with all parameters fixed at nominal default values.

sdcycod = 2;
fdcycod = 1;
sgaincod = 1;
dbpbcod = 2;
floorcod = 4;
cplfgaincod = 4;
fgaincod[ch] = 4;
lfegaincod = 4;
cplsnroffst = fsnroffst[ch] = lfesnroffst = fineoffset.

Since the bit allocation parameters are static, they are only sent during block 0. Delta bit allocation is not used, so deltbaie = 0. The core bit allocation routine (described in clause 6.2) is run, and the coarse and fine SNR offsets are adjusted until all available bits in the syncframe are used up. The coarse SNR offset adjusts in 3 dB increments, and the fine offset adjusts in 3/16 dB increments. Bits are allocated globally from a common bit pool to all channels. The combination of csnroffst and fineoffset which uses the largest number of bits without exceeding the syncframe size is chosen. This involves an iterative process. When, for a given iteration, the number of bits exceeds the pool, the SNR offset is decreased for the next iteration. On the other hand, if the allocation is less than the pool, the SNR offset is increased for the next iteration. When the SNR offset is at its maximum without causing the allocation to exceed the pool, the iterating is complete. The results of the bit allocation routine are the final values of csnroffst and fineoffset, and the set of bit allocation pointers (baps). The SNR offset values are included in the bit stream so that the decoder does not need to iterate.

7.2.13 Quantize mantissas

The baps are used by the mantissa quantization block. There is a bap for each individual transform coefficient. Each normalized mantissa is quantized by the quantizer indicated by the corresponding bap. Asymmetrically quantized mantissas are quantized by rounding to the number of bits indicated by the corresponding bap. Symmetrically quantized mantissas are quantized through the use of a table look-up. Mantissas with baps of 1, 2, and 4 are grouped into triples or duples.

7.2.14 Pack AC-3 syncframe

All of the data is packed into the encoded AC-3 syncframe. Some of the quantized mantissas are grouped together and coded by a single codeword. The output format is dependent on the application. The syncframe may be output in a burst, or delivered as a serial data stream at a constant rate.

Annex A (normative): AC-3 bit streams in the MPEG-2 multiplex

A.0 Scope

This annex contains specifications on how to combine one or more AC-3 bit streams into the ATSC (Recommendation ITU-R BT.1300 [i.5], System A) or DVB (Recommendation ITU-R BT.1300 [i.5], System B) MPEG-2 transport stream (ISO/IEC 13818-1 [i.4]).

A.1 Introduction

The AC-3 bit stream is included in an MPEG-2 multiplex bit stream in much the same way an MPEG-1 audio stream would be included. The AC-3 bit stream is packetized into PES packets. An MPEG-2 multiplex bit stream containing AC-3 bit streams shall meet all constraints described in the STD model in clause A.2.6. It is necessary to unambiguously indicate that an AC-3 stream is, in fact, an AC-3 stream (and not an MPEG audio stream). The MPEG-2 standard does not explicitly indicate codes to be used to indicate an AC-3 stream. Also, the MPEG-2 standard does not have an audio descriptor adequate to describe the contents of the AC-3 bit stream in the PSI tables.

The AC-3 audio access unit (AU) or presentation unit (PU) is an AC-3 syncframe. The AC-3 syncframe contains 1 536 audio samples. The duration of an AC-3 access (or presentation) unit is 32 ms for audio sampled at 48 kHz, approximately 34,83 ms for audio sampled at 44,1 kHz, and 48 ms for audio sampled at 32 kHz.

The items which need to be specified in order to include AC-3 within the MPEG-2 bit stream are: stream_type, stream_id, AC-3 audio descriptor, and, for system A only, registration descriptor. The registration descriptor is not required in System B since the AC-3_descriptor is regarded as a public descriptor in this system. The ISO 639-1 [i.1] language descriptor may be employed to indicate language. Some constraints are placed on the PES layer for the case of multiple audio streams intended to be reproduced in exact sample synchronism. In System A (ATSC) the AC-3 audio descriptor is titled "AC-3_audio_stream_descriptor" while in System B (DVB) the AC-3 audio descriptor is titled "AC-3_descriptor". It should be noted that the syntax of these descriptors differs significantly between the two systems.

A.2 Detailed specification for System A (ATSC)

A.2.1 Stream_type

The value of stream_type for AC-3 shall be 0×81.

A.2.2 Stream_id

The value of stream_id in the PES header shall be 0×BD (indicating private_stream_1). Multiple AC-3 streams may share the same value of stream_id since each stream is carried with a unique PID value. The mapping of values of PID to stream_type is indicated in the transport stream programme map table (PMT).

A.2.3 Registration_descriptor

The syntax of the AC-3 registration_descriptor is shown below. The AC-3 registration_descriptor shall be included in the TS_programme_map_section.

Table A.2.0a

Syntax	No. of bits	Mnemonic
registration_descriptor()		
{		
descriptor_tag	8	uimsbf
descriptor_length	8	uimsbf
format_identifier	32	uimsbf
}		
descriptor_tag - 0×05.		
descriptor_length - 0×04.		
format_identifier - The AC-3 format_identifier is 0×41432D33 ("AC-3").		

A.2.4 AC-3 audio_descriptor

The AC-3 audio_stream_descriptor allows information about individual AC-3 bit streams to be included in the programme specific information (PSI) tables. This information is useful to enable decision making as to the appropriate AC-3 stream(s) that are present in the current broadcast to be directed to the audio decoder, and also to enable the announcement of characteristics of audio streams that will be included in future broadcasts. Note that horizontal lines in table A.2.0b indicate allowable termination points for the descriptor.

Table A.2.0b

Syntax	No. of bits	Mnemonic
AC-3_audio_stream_descriptor()		
{		
descriptor_tag	8	uimsbf
descriptor_length	8	uimsbf
sample_rate_code	3	bslbf
bsid	5	bslbf
bit_rate_code	6	bslbf
surround_mode	2	bslbf
bsmod	3	bslbf
num_channels	4	bslbf
full_svc	1	bslbf
langcod	8	bslbf
if(num_channels == 0) /* 1+1 mode */		
{		
langcod2	8	bslbf
}		
if(bsmod < 2)		
{		
mainid	3	uimsbf
priority	2	bslbf
reserved	3	'111'
}		
else **asvcflags**	8	bslbf
textlen	7	uimsbf
text_code	1	bslbf
for(i = 0; i < m; i++)		
{		
text[i]	8	bslbf
}		
language_flag	1	bslbf
language_flag_2	1	bslbf
reserved	6	'111111'
if(language_flag == 1) **{language}**	3x8	uimsbf
if(language_flag_2 == 1) **{language_2}**	3x8	uimbsf
for(i = 0; i < n; i++)		
{		
additional_info[i]	nx8	bslbf
}		
}		

descriptor_tag - The value for the AC-3 descriptor_tag is 0×81.

descriptor_length - This is an 8-bit field specifying the number of bytes of the descriptor immediately following descriptor_length field.

sample_rate_code - This is a 3-bit field which indicates the sample rate of the encoded audio. The indication may be of one specific sample rate, or may be of a set of values which include the sample rate of the encoded audio (see Table A.2.1).

Table A.2.1: Sample_rate_code table

sample_rate_code	Sample rate (kHz)
000	48
001	44,1
010	32
011	Reserved
100	48 or 44,1
101	48 or 32
110	44,1 or 32
111	48 or 44,1 or 32

bsid - This is a 5-bit field which is set to the same value as the bsid field in the AC-3 bit stream.

bit_rate_code - This is a 6-bit field. The lower 5 bits indicate a nominal bit rate. The MSB indicates whether the indicated bit rate is exact (MSB = 0) or an upper limit (MSB = 1) (see Table A.2.2).

Table A.2.2: Bit_rate_code table

bit_rate_code	Exact bit rate (kbit/s)	bit_rate_code	Bit rate upper limit (kbit/s)
000000 (0)	32	100000 (32)	32
000001 (1)	40	100001 (33)	40
000010 (2)	48	100010 (34)	48
000011 (3)	56	100011 (35)	56
000100 (4)	64	100100 (36)	64
000101 (5)	80	100101 (37)	80
000110 (6)	96	100110 (38)	96
000111 (7)	112	100111 (39)	112
001000 (8)	128	101000 (40)	128
001001 (9)	160	101001 (41)	160
001010 (10)	192	101010 (42)	192
001011 (11)	224	101011 (43)	224
001100 (12)	256	101100 (44)	256
001101 (13)	320	101101 (45)	320
001110 (14)	384	101110 (46)	384
001111 (15)	448	101111 (47)	448
010000 (16)	512	110000 (48)	512
010001 (17)	576	110001 (49)	576
010010 (18)	640	110010 (50)	640

dsurmod - This is a 2-bit field which may be set to the same value as the dsurmod field in the AC-3 bit stream, or which may be set to "00" (not indicated) (see Table A.2.3).

Table A.2.3: dsurmod table

surround_mode	Meaning
00	Not indicated
01	NOT Dolby® surround encoded
10	Dolby® surround encoded
11	Reserved

bsmod - This is a 3-bit field which is set to the same value as the bsmod field in the AC-3 bit stream.

num_channels - This is a 4-bit field which indicates the number of channels in the AC-3 bit stream. When the MSB is 0, the lower 3 bits are set to the same value as the acmod field in the AC-3 bit stream. When the MSB field is 1, the lower 3 bits indicate the maximum number of encoded audio channels (counting the lfe channel as 1). If the value of acmod in the AC-3 bit stream is "000" (1 + 1 mode), then the value of num_channels shall be set to "0000" (see Table A.2.4).

Table A.2.4: Num_channels table

num_channels	Audio coding mode (acmod)
0000	1 + 1
0001	1/0
0010	2/0
0011	3/0
0100	2/1
0101	3/1
0110	2/2
0111	3/2

num_channels	Number of encoded channels
1000	1
1001	≤ 2
1010	≤ 3
1011	≤ 4
1100	≤ 5
1101	≤ 6
1110	Reserved
1111	Reserved

full_svc - This is a 1-bit field which indicates whether or not this audio service is a full service suitable for presentation, or whether this audio service is only a partial service which should be combined with another audio service before presentation. This bit should be set to a "1" if this audio service is sufficiently complete to be presented to the listener without being combined with another audio service (for example, a visually impaired service which contains all elements of the programme; music, effects, dialogue, and the visual content descriptive narrative). This bit should be set to a "0" if the service is not sufficiently complete to be presented without being combined with another audio service (e.g. a visually impaired service which only contains a narrative description of the visual programme content and which needs to be combined with another audio service which contains music, effects, and dialogue).

langcod - This is an 8-bit field which is set to the same value as the langcod field in the AC-3 bit stream. If the AC-3 bit stream langcod field is not present, then this 8-bit field shall be set to 0xFF if present.

langcod2 - This is an 8-bit field which is set to the value of the langcod2 field in the AC-3 bit stream. If the AC-3 bit stream langcod2 field is not present, then this 8-bit field shall be set to 0xFF if present.

NOTE 1: The langcod and langcod2 fields are not (that is, are no longer) used to indicate language. The ISO 639 [i.1] language descriptor is used to indicate language. However, the AC-3 audio descriptor may optionally include the ISO_639_language_code, see below "language" and "language_2" fields.

mainid - This is a 3-bit field which contains a number in the range 0 - 7 which identifies a main audio service. Each main service should be tagged with a unique number. This value is used as an identifier to link associated services with particular main services.

priority - This is a 2-bit field that indicates the priority of the audio service. This field allows a Main audio service (bsmod equal to 0 or 1) to be marked as the primary audio service. Other audio services may be explicitly marked or not specified. Table A.2.5 shows how this field is encoded.

Table A.2.5: Priority Field Coding

Bit Field	Meaning
00	reserved
01	Primary Audio
10	Other Audio
11	Not specified

asvcflags - This is an 8-bit field. Each bit (0 - 7) indicates with which main service(s) this associated service is associated. The left most bit, bit 7, indicates whether this associated service may be reproduced along with main service number 7. If the bit has a value of 1, the service is associated with main service number 7. If the bit has a value of 0, the service is not associated with main service number 7.

textlen - This is an unsigned integer which indicates the length, in bytes, of a descriptive text field which follows.

text_code - This is a 1-bit field which indicates how the following text field is encoded. If this bit is a "1", the text is encoded as 1-byte characters using the ISO Latin-1 alphabet (ISO 8859-1 [i.3]). If this bit is a "0", the text is encoded with 2-byte unicode characters.

text[i] - The text field may contain a brief textual description of the audio service.

language_flag - This is a 1-bit flag that indicates whether or not the 3-byte language field is present in the descriptor. If this bit is set to "1", then the 3-byte language field is present. If this bit is set to "0", then the language field is not present.

language_flag_2 - This is a 1-bit flag that indicates whether or not the 3-byte language_2 field is present in the descriptor. If this bit is set to "1", then the 3-byte language_2 field is present. If this bit is set to "0", then the language_2 field is not present. This bit shall always be set to "0", unless the num_channels field is set to "0000" indicating the audio coding mode is 1+1 (dual mono). If the num_channels field is set to "0000" then this bit may be set to "1" and the language_2 field may be included in this descriptor.

language - This field is a 3-byte language code per ISO 639-2/B [i.2] defining the language of this audio service. If the AC-3 stream audio coding mode is 1+1 (dual mono), this field indicates the language of the first channel (channel 1, or "left" channel).The language field shall contain a three-character code as specified by ISO 639-2/B [i.2]. Each character is coded into 8 bits according to ISO 8859-1 [i.3] (ISO Latin-1) and inserted in order into the 24-bit field. The coding is identical to that used in the ISO_639_language_code value in the ISO_639_language_descriptor specified in ISO/IEC 13818-1 [i.4].

language_2 - This field is only present if the AC-3 stream audio coding mode is 1+1 (dual mono). This field is a 3-byte language code per ISO 639-2/B [i.2] defining the language of the second channel (channel 2, or "right" channel) in the AC-3 bit stream. The language_2 field shall contain a three-character code as specified by ISO n order into the 24-bit field. The coding is identical to that used in the ISO_639_language_code value in the ISO_639_language_descriptor specified in ISO/IEC 13818-1 [i.4].

additional_info[j] - This is a set of additional bytes filling out the remainder of the descriptor. The purpose of these bytes is not currently defined. This field is provided to allow the ATSC to extend this descriptor. No other use is permitted.

NOTE 2: In the event that there is a single Main service that alternates between different languages, the ISO 639-1 [i.1] Language descriptor may be used to communicate that additional information.

A.2.5 ISO_639_language_code

The ISO_639_language_code descriptor allows a stream to be tagged with the 24-bit ISO 639-1 [i.1] language code.

A.2.6 STD audio buffer size

For an MPEG-2 transport stream, the T-STD model defines the main audio buffer size BS_n as:

$$BS_n = BS_{mux} + BS_{dec} + BS_{oh}$$

where:

BS_{mux} = 736 bytes
BS_{oh}: PES header overhead
BS_{dec}: access unit buffer

MPEG-2 specifies a fixed value for BS_n (3 584 bytes) and indicates that any excess buffer may be used for additional multiplexing.

When an AC-3 bit stream is carried by an MPEG-2 transport stream, the transport stream shall be compliant with a main audio buffer size of:

$$BS_n = BS_{mux} + BS_{pad} + BS_{dec}$$

where:

BS_{mux} = 736 bytes
BS_{pad} = 64 bytes

The value of BS_{dec} employed shall be that of the highest bit rate supported by the system (i.e. the buffer size is not decreased when the audio bit rate is less than the maximum value allowed by a specific system). The 64 bytes in BS_{pad} are available for BS_{oh} and additional multiplexing. This constraint makes it possible to implement decoders with the minimum possible memory buffer.

A.3 Specification for System B (DVB)

NOTE: Transport of AC-3 bit streams in an MPEG-2 Transport Stream that conforms to System B is defined in ETSI TS 101 154 [i.6] and ETSI EN 300 468 [i.7].

A.4 PES constraints

A.4.0 Introduction

The following clauses apply to both System A and System B.

A.4.1 Encoding

In some applications, the audio decoder may be capable of simultaneously decoding two bit streams containing different programme elements, and then combining the programme elements into a complete programme.

Most of the programme elements are found in the *main audio service*. Another programme element (such as a narration of the picture content intended for the visually impaired listener) may be found in the *associated audio service*.

In order to have the audio from the two bit streams reproduced in exact sample synchronism, it is necessary for the original audio bit stream encoders to have encoded the two audio programme elements frame synchronously; i.e. if audio stream 1 has sample 0 of frame *n* taken at time *t*0, then audio stream 2 should also have frame *n* beginning with its sample 0 taken the identical time *t*0. If the encoding of multiple audio services is done frame and sample synchronous, and decoding is intended to be frame and sample synchronous, then the PES packets of these audio services shall contain identical values of PTS which refer to the audio access units intended for synchronous decoding.

Audio services intended to be combined together for reproduction shall be encoded at an identical sample rate.

A.4.2 Decoding

If audio access units from two audio services which are to be simultaneously decoded have identical values of PTS indicated in their corresponding PES headers, then the corresponding audio access units shall be presented to the audio decoder for simultaneous synchronous decoding. Synchronous decoding means that for corresponding audio frames (access units), corresponding audio samples are presented at the identical time.

If the PTS values do not match (indicating that the audio encoding was not frame synchronous) then the audio frames (access units) of the main audio service may be presented to the audio decoder for decoding and presentation at the time indicated by the PTS. An associated service which is being simultaneously decoded may have its audio frames (access units), which are in closest time alignment (as indicated by the PTS) to those of the main service being decoded, presented to the audio decoder for simultaneous decoding. In this case the associated service may be reproduced out of sync by as much as 1/2 of a frame time. (This is typically satisfactory; a visually impaired narration does not require highly precise timing).

A.4.3 Byte-alignment

The AC-3 bit stream shall be byte-aligned within the MPEG-2 data stream. This means that the initial 8 bits of an AC-3 syncframe shall reside in a single byte which is carried by the MPEG-2 data stream.

Annex B (informative): Void

Annex C (informative): AC-3 karaoke mode

C.0 Scope

This annex contains specifications for how *karaoke-aware* and *karaoke-capable* AC-3 decoders should reproduce *karaoke* AC-3 bit streams. A minimum level of functionality is defined which allows a *karaoke-aware* decoder to produce an appropriate 2/0 or 3/0 default output when presented with a karaoke mode AC-3 bit stream. An additional level of functionality is defined for the *karaoke-capable* decoder so that the listener may optionally control the reproduction of the karaoke bit stream.

C.1 Introduction

The AC-3 karaoke mode has been defined in order to allow the multi-channel AC-3 bit stream to convey audio channels designated as L, R (e.g. 2-channel stereo music), M (e.g. guide melody), and V1, V2 (e.g. one or two vocal tracks). This annex does not specify the contents of L, R, M, V1, and V2, but does specify the behaviour of AC-3 decoding equipment when receiving a karaoke bit stream containing these channels. An AC-3 decoder which is karaoke-capable will allow the listener to optionally reproduce the V1 and V2 channels, and may allow the listener to adjust the relative levels (mixing balance) of the M, V1, and V2 channels. An AC-3 decoder which is karaoke-aware will reproduce the L, R, and M channels, and will reproduce the V1 and V2 channels at a level indicated by the encoded bit stream. The 2-channel karaoke-aware decoder will decode the karaoke bit stream using the Lo, Ro downmix. The L and R channels will be reproduced out of the left and right outputs, and the M channel will appear as a phantom centre. The precise level of the M channel is determined by cmixlev which is under control of the programme provider. The level of the V1 and V2 channels which will appear in the downmix is determined by surmixlev, which is under control of the programme provider. A single V channel (V1 only) will appear as a phantom centre. A pair of V channels (V1 and V2) will be reproduced with V1 in left output and V2 in right output. The 5-channel karaoke-aware decoder will reproduce the L, R channels out of the left and right outputs, and the M channel out of the centre output. A single V channel (V1 only) will be reproduced in the centre channel output. A pair of V implementation of the flexible karaoke-capable decoder is not specified; it is left up to the implementation as to the degree of adjustability to be offered to the listener.

C.2 Detailed specification

C.2.1 Karaoke mode indication

AC-3 bit streams are indicated as karaoke type when bsmod = "111" and acmod ≥ 0 x 2.

C.2.2 Karaoke mode channel assignment

The channel assignments for both the normal mode and the karaoke mode are shown in Table C.2.1.

Table C.2.1: Channel Array Ordering

acmod	Audio Coding Mode	Normal Channel Assignment (bsmod != "111")	Karaoke Channel Assignment (bsmod ="111")
010	2/0	L, R	L, R
011	3/0	L, C, R	L, M, R
100	2/1	L, R, S	L, R, V1
101	3/1	L, C, R, S	L, M, R, V1
110	2/2	L, R, Ls, Rs	L, R, V1, V2
111	3/2	L, C, R, Ls, Rs	L, M, R, V1, V2

C.2.3 Reproduction of karaoke mode bit streams

C.2.3.0 Introduction

This clause contains the specifications which should be met by decoders which are designated as karaoke-aware or karaoke-capable. The following general equations indicate how the AC-3 decoder's output channels, L_k, C_k, R_k, are formed from the encoded channels L, M, R, V1, V2.

Typically, the surround loudspeakers are not used when reproducing karaoke bit streams.

$$L_k = L + a \times V1 + b \times V2 + c \times M$$

$$C_k = d \times V1 + e \times V2 + f \times M$$

$$R_k = R + g \times V1 + h \times V2 + i \times M$$

C.2.3.1 Karaoke-aware decoder

The values of the coefficients a-i, which are used by karaoke-aware decoders, are given in Table C.2.2. Values are shown for both 2-channel (2/0) and multi-channel (3/0) reproduction. For each of these situations, a coefficient set is shown for the case of a single encoded V channel (V1 only) or two encoded V channels (V1, V2). The actual coefficients used need to be scaled downwards so that arithmetic overflow does not occur if all channels contributing to an output channel happen to be at full scale. Monophonic reproduction would be obtained by summing the left and right output channels of the 2/0 reproduction. Any AC-3 decoder will produce the appropriate output if it is set to perform an Lo, Ro 2-channel downmix.

Table C.2.2: Coefficient values for karaoke-aware decoders

Coefficient	2/0 Reproduction		3/0 Reproduction	
	1 Vocal	2 Vocals	1 Vocal	2 Vocals
a	0,7 × slev	slev	0,0	slev
b	---	0,0	---	0,0
c	clev	clev	0,0	0,0
d	---	---	slev	0,0
e	---	---	---	0,0
f	---	---	1,0	1,0
g	0,7 × slev	0,0	0,0	0,0
h	---	slev	---	slev
i	clev	clev	0,0	0,0

C.2.3.2 Karaoke-capable decoders

Karaoke-capable decoders allow the user to choose to have the decoder reproduce none, one, or both of the V channels. The default coefficient values for the karaoke-capable decoder are given in Table C.2.3. When the listener selects to have none, one, or both of the V channels reproduced, the default coefficients are given in Table C.2.3. Values are shown for both 2-channel (2/0) and multi-channel (3/0) reproduction, and for the cases of user selected reproduction of no V channel (None), one V channel (either V1 or V2), or both V channels (V1+V2). The M channel and a single V channel are reproduced out of the centre output (phantom centre in 2/0 reproduction), and a pair of V channels are reproduced out of the left (V1) and right (V2) outputs. The actual coefficients used need to be scaled downwards so that arithmetic overflow does not occur if all channels contributing to an output happen to be at full scale.

Table C.2.3: Default coefficient values for karaoke-capable decoders

Coefficient	2/0 Reproduction				3/0 Reproduction			
	None	V1	V2	V1+V2	None	V1	V2	V1+V2
a	0,0	0,7	0,0	1,0	0,0	0,0	0,0	1,0
b	0,0	0,0	0,7	0,0	0,0	0,0	0,0	0,0
c	clev	clev	clev	clev	0,0	0,0	0,0	0,0
d	---	---	---	---	0,0	1,0	0,0	0,0

Coefficient	2/0 Reproduction				3/0 Reproduction			
	None	V1	V2	V1+V2	None	V1	V2	V1+V2
e	---	---	---	---	0,0	0,0	1,0	0,0
f	---	---	---	---	1,0	1,0	1,0	1,0
g	0,0	0,7	0,0	0,0	0,0	0,0	0,0	0,0
h	0,0	0,0	0,7	1,0	0,0	0,0	0,0	1,0
i	clev	clev	clev	clev	0,0	0,0	0,0	0,0

Additional flexibility may be offered optionally to the user of the karaoke decoder. For instance, the coefficients a, d, and g might be adjusted to allow the V1 channel to be reproduced in a different location and with a different level. Similarly the level and location of the V2 and M channels could be adjusted. The details of these additional optional user controls are not specified and are left up to the implementation. Also left up to the implementation is what use might be made of the L_s, R_s outputs of the 5-channel decoder, which would naturally reproduce the V1, V2 channels.

Annex D (normative): Alternate bit stream syntax

D.0 Scope

This annex contains specifications for an alternate bit stream syntax that may be implemented by some AC-3 encoders and interpreted by some AC-3 decoders. The new syntax redefines certain bit stream information (bsi) fields to carry new meanings. It is not necessary for decoders to be aware of this alternate syntax in order to properly reconstruct an audio soundfield; however those decoders that are aware of this syntax will be able to take advantage of the new system features described in this annex. This alternate bit stream syntax is identified by setting the bsid to a value of 6. This annex is normative to the extent that when bsid is set to the value of 6, the alternate syntax elements shall have the meaning described in this annex. Thus this annex may be considered normative on encoders that set bsid to 6. This annex is informative for decoders. Interpretation and use of the new syntactical elements is optional for decoders. The new syntactical elements defined in this annex are placed in the two 14-bit fields that are defined as timecod1 and timecod2 in clause 4 (these fields have never been applied for their originally anticipated purpose).

D.1 Specification

D.1.1 Indication of alternate bit stream syntax

An AC-3 bit stream shall have the alternate bit stream syntax described in this annex when the bit stream identification (bsid) field is set to 6.

D.1.2 Alternate bit stream syntax specification

Syntax	Word Size
bsi()	
{	
bsid	5
bsmod	3
acmod	3
if((acmod & 0x1) && (acmod != 0x1)) /* if 3 front channels */ **{cmixlev}**	2
if(acmod & 0x4) /* if a surround channel exists */ **{surmixlev}**	2
if(acmod == 0x2) /* if in 2/0 mode */ **{dsurmod}**	2
lfeon	1
dialnorm	5
compre	1
if(compre) **{compr}**	8
langcode	1
if(langcode) **{langcod}**	8
audprodie	1
if(audprodie)	
{	
mixlevel	5
roomtyp	2
}	
if(acmod == 0) /* if 1+1 mode (dual mono, so some items need a second value) */	
{	
dialnorm2	5
compr2e	1
if(compr2e) **{compr2}**	8
langcod2e	1
if(langcod2e) **{langcod2}**	8
audprodi2e	1
if(audprodi2e)	
{	
mixlevel2	5
roomtyp2	2
}	
}	
copyrightb	1
origbs	1

```
Syntax                                                                    Word Size
   xbsi1e      ..................................................................1
   if(xbsi1e)
   {
      dmixmod  ..................................................................2
      ltrtcmixlev  ..............................................................3
      ltrtsurmixlev  ............................................................3
      lorocmixlev  ..............................................................3
      lorosurmixlev  ............................................................3
   }
   xbsi2e      ..................................................................1
   if(xbsi2e)
   {
      dsurexmod  ................................................................2
      dheadphonmod  .............................................................2
      adconvtyp  ................................................................1
      xbsi2    ..................................................................8
      encinfo  ..................................................................1
   }
   addbsie  ..................................................................1
   if(addbsie)
   {
      addbsil  ..................................................................6
      addbsi (addbsil+1) x 8
   }
} /* end of bsi */
```

D.1.3 Description of alternate syntax bit stream elements

D.1.3.0 Introduction

The following clauses describe the meaning of the alternate syntax bit stream elements. Elements not specifically described retain the same meaning as specified in clause 4, except as noted in the alternate bit stream constraints clause above.

D.1.3.1 xbsi1e: Extra bit stream information #1 exists, 1 bit

If this bit is a 1, the following 14 bits contain extra bit stream information.

D.1.3.2 dmixmod: Preferred stereo downmix mode, 2 bits

This 2-bit code, as shown in Table D.1.1, indicates the type of stereo downmix preferred by the mastering engineer. This information may be used by the decoder to automatically configure the type of stereo downmix, but may also be overridden or ignored. If dmixmod is set to the reserved code, the decoder should still reproduce audio. The reserved code may be interpreted as "not indicated".

Table D.1.1: Preferred Stereo Downmix Mode

dmixmod	Indication
00	Not indicated
01	Lt/Rt downmix preferred
10	Lo/Ro downmix preferred
11	Reserved
NOTE: The meaning of this field is only defined as described if the audio coding mode is 3/0, 2/1, 3/1, 2/2 or 3/2. If the audio coding mode is 1+1, 1/0 or 2/0 then the meaning of this field is reserved.	

D.1.3.3 ltrtcmixlev: Lt/Rt centre mix level, 3 bits

This 3-bit code, shown in Table D.1.2, indicates the nominal down mix level of the centre channel with respect to the left and right channels in an Lt/Rt downmix.

Table D.1.2: Lt/Rt Centre Mix Level

ltrtcmixlev	Clev
000	1,414 (+3,0 dB)
001	1,189 (+1,5 dB)
010	1,000 (0,0 dB)
011	0,841 (-1,5 dB)
100	0,707 (-3,0 dB)
101	0,595 (-4,5 dB)
110	0,500 (-6,0 dB)
111	0,000 (-inf dB)
NOTE: The meaning of this field is only defined as described if the audio coding mode is 3/0, 3/1 or 3/2. If the audio coding mode is 1+1, 1/0, 2/0, 2/1 or 2/2 then the meaning of this field is reserved.	

D.1.3.4 ltrtsurmixlev: Lt/Rt surround mix level, 3 bits

This 3-bit code, shown in Table D.1.3, indicates the nominal down mix level of the surround channels with respect to the left and right channels in an Lt/Rt downmix. If one of the reserved values is received, the decoder should use a value of 0,841 for slev.

Table D.1.3: Lt/Rt Surround Mix Level

ltrtsurmixlev	Slev
000	Reserved
001	Reserved
010	Reserved
011	0,841 (-1,5 dB)
100	0,707 (-3,0 dB)
101	0,595 (-4,5 dB)
110	0,500 (-6,0 dB)
111	0,000 (-inf dB)
NOTE: The meaning of this field is only defined as described if the audio coding mode is 2/1, 3/1, 2/2 or 3/2. If the audio coding mode is 1+1, 1/0, 2/0 or 3/0 then the meaning of this field is reserved.	

D.1.3.5 lorocmixlev: Lo/Ro centre mix level, 3 bits

This 3-bit code, shown in Table D.1.4, indicates the nominal down mix level of the centre channel with respect to the left and right channels in an Lo/Ro downmix.

Table D.1.4: Lo/Ro Centre Mix Level

lorocmixlev	Clev
000	1,414 (+3,0 dB)
001	1,189 (+1,5 dB)
010	1,000 (0,0 dB)
011	0,841 (-1,5 dB)
100	0,707 (-3,0 dB)
101	0,595 (-4,5 dB)
110	0,500 (-6,0 dB)
111	0,000 (-inf dB)
NOTE: The meaning of this field is only defined as described if the audio coding mode is 3/0, 3/1 or 3/2. If the audio coding mode is 1+1, 1/0, 2/0, 2/1 or 2/2 then the meaning of this field is reserved.	

D.1.3.6 lorosurmixlev: Lo/Ro surround mix level, 3 bits

This 3-bit code, shown in Table D.1.5, indicates the nominal down mix level of the surround channels with respect to the left and right channels in an Lo/Ro downmix. If one of the reserved values is received, the decoder should use a value of 0,841 for slev.

Table D.1.5: Lo/Ro Surround Mix Level

lorosurmixlev	Slev
000	Reserved
001	Reserved
010	Reserved
011	0,841 (-1,5 dB)
100	0,707 (-3,0 dB)
101	0,595 (-4,5 dB)
110	0,500 (-6,0 dB)
111	0,000 (-inf dB)
NOTE: The meaning of this field is only defined as described if the audio coding mode is 2/1, 3/1, 2/2 or 3/2. If the audio coding mode is 1+1, 1/0, 2/0 or 3/0 then the meaning of this field is reserved.	

D.1.3.7 xbsi2e: Extra bit stream information #2 exists, 1 bit

If this bit is a 1, the following 14 bits contain extra bit stream information.

D.1.3.8 dsurexmod: Dolby® Surround EX® mode, 2 bits

This 2-bit code, as shown in Table D.1.6, indicates whether or not the programme has been encoded in Dolby Surround EX®, Dolby Pro Logic® IIx or Dolby Pro Logic® IIz. This information is not used by the decoder, but may be used by other portions of the audio reproduction equipment.

Table D.1.6: Dolby Surround EX® Mode

dsurexmod	Indication
00	Not indicated
01	Not Dolby Surround EX®, Dolby Pro Logic® IIx or Dolby Pro Logic® IIz encoded
10	Dolby Surround EX® or Dolby Pro Logic® IIx encoded
11	Dolby Pro Logic® IIz encoded
NOTE: The meaning of this field is only defined as described if the audio coding mode is 2/2 or 3/2. If the audio coding mode is 1+1, 1/0, 2/0, 3/0, 2/1 or 3/1 then the meaning of this field is reserved.	

D.1.3.9 dheadphonmod: Dolby® Headphone mode, 2 bits

This 2-bit code, as shown in Table D.1.7, indicates whether or not the programme has been Dolby® Headphone-encoded. This information is not used by the decoder, but may be used by other portions of the audio reproduction equipment. If dheadphonmod is set to the reserved code, the decoder should still reproduce audio. The reserved code may be interpreted as "not indicated".

Table D.1.7: Dolby® Headphone Mode

dheadphonmod	Indication
00	Not indicated
01	Not Dolby® Headphone encoded
10	Dolby® Headphone encoded
11	Reserved
NOTE: The meaning of this field is only defined as described if the audio coding mode is 2/0. If the audio coding mode is 1+1, 1/0, 3/0, 2/1, 3/1, 2/2 or 3/2 then the meaning of this field is reserved.	

D.1.3.10 adconvtyp: A/D converter type, 1 bit

This 1-bit code, as shown in Table D.1.8, indicates the type of A/D converter technology used to capture the PCM audio. This information is not used by the AC-3 decoder, but may be used by other portions of the audio reproduction equipment. If the type of A/D converter used is not known, the "Standard" setting should be chosen.

Table D.1.8: A/D Converter Type

Adconvtyp	Indication
"0"	Standard
"1"	HDCD

D.1.3.11 xbsi2: Extra bit stream information, 8 bits

This field is reserved for future assignment. Encoders shall set these bits to all 0's.

D.1.3.12 encinfo: Encoder information, 1 bit

This field is reserved for use by the encoder, and is not used by the decoder.

D.2 Decoder processing

D.2.0 Introduction

There are two types of decoders: those that recognize the alternate syntax (*compliant decoders*), and those that do not (*legacy decoders*). This clause specifies how each type of decoder will process bit streams that use the alternate bit stream syntax. Implementation of compliant decoding is optional.

D.2.1 Compliant decoder processing

D.2.1.1 Two-channel downmix selection

In the case of a two-channel downmix, compliant decoders should allow the end user to specify which two-channel downmix is chosen. Three separate options should be allowed: Lt/Rt downmix, Lo/Ro downmix, or automatic selection of either Lt/Rt or Lo/Ro based on the preferred downmix mode parameter dmixmod.

D.2.1.2 Two-channel downmix processing

Once a particular two-channel downmix has been selected, compliant decoders should use the new centre mix level and surround mix level parameters associated with the selected downmix type (assuming they are included in the bit stream). If Lt/Rt downmix is selected, compliant decoders should use the ltrtcmixlev and ltrtsurmixlev parameters (if included). If Lo/Ro downmix is selected, compliant decoders should use the lorocmixlev and lorosurmixlev parameters (if included). If these parameters are not included in the bit stream, then downmixing should be performed as defined in the original specification.

D.2.1.3 Informational parameter processing

Compliant decoders should provide a means for informational parameters (e.g. dsurexmod, dheadphonmod, etc.) to be accessed by external system components. Note that these parameters do not otherwise affect decoder processing.

D.2.2 Legacy decoder processing

Legacy decoders do not recognize the alternate bit stream syntax, but rather interpret these bit fields according to their original definitions in clause 4. The extra bit stream information words (xbsi1e, xbsi2e, dmixmod, etc.) are interpreted as time code words (timecod1e, timecod1, timecod2e, and timecod2).

As described in clause 4, the time code words do not affect the decoding process in legacy decoders. As a result, the alternate bit stream syntax can be safely decoded without causing incorrect decoder processing. However, legacy decoders will not be able to take advantage of new functionality provided by the alternate syntax.

D.3 Encoder processing

D.3.0 Introduction

This clause describes processing steps and requirements associated with encoders that create bits streams according to the alternate bit stream syntax.

D.3.1 Encoder processing steps

D.3.1.1 Dynamic range overload protection processing

If the alternate bit stream syntax is used, the dynamic range overload protection function within the encoder shall account for potential overload in either legacy or compliant decoders, using any downmix mode. No assumption should be made that compliant decoders will necessarily use the preferred downmix mode.

D.3.2 Encoder requirements

D.3.2.1 Legacy decoder support

In order to support legacy decoder operations, it is necessary to continue to specify valid values for bit stream information parameters that are made obsolete by the alternate bit stream syntax. For example, the new ltrtcmixlev, ltrtsurmixlev, lorocmixlev, and lorosurmixlev fields (if included in the alternate bit stream) override the functionality of the previously defined cmixlev and surmixlev fields. Nonetheless, alternate bit stream syntax encoders shall continue to specify valid values for the cmixlev and surmixlev fields.

D.3.2.2 Original bit stream syntax support

Encoding equipment that is capable of creating bit streams according to the alternate bit stream syntax shall also provide an option that allows for creation of bit streams according to the present document not including this annex.

Annex E (normative): Enhanced AC-3

E.0 Scope

This annex defines the audio coding algorithm denoted as Enhanced AC-3 ("E-AC-3") and the alterations to the AC-3 bit stream necessary to convey E-AC-3 data along with a reference decoding process. E-AC-3 bit streams are similar in nature to standard AC-3 bit streams but are not backwardly compatible (i.e., they are not decodable by standard AC-3 decoders). This annex specifies either directly or by reference the bit stream syntax of E-AC-3. When an AC-3 bit stream carries E-AC-3 bit stream syntax, it is referred herein to as an E-AC-3 bit stream.

E.1 Bit stream syntax and semantics specification

E.1.1 Indication of Enhanced AC-3 bit stream syntax

An AC-3 bit stream is indicated as using the Enhanced AC-3 bit stream syntax described in this annex when the bit stream identification (bsid) field is set to 16. To enable differentiation between an AC-3 bit stream and an E-AC-3 bit stream, the bsid field is placed the same number of bits from the beginning of the syncframe as defined in the syntax below.

E.1.2 Syntax specification

E.1.2.0 E-AC-3_bit_stream and syncframe

Unless otherwise specified, all bit stream elements shall have the same meaning and purpose as described in the body and Annex D of the present document. Single bit boolean values shall be treated as '1' equals TRUE. A continuous audio bit stream consists of a sequence of synchronization frames.

Syntax

```
E-AC-3_bit_stream()
{
   while(true)
   {
      syncframe() ;
   }
} /* end of bit stream */
```

The syncframe consists of the syncinfo, bsi and audfrm fields, up to 6 coded audblk fields, the auxdata field, and the errorcheck field.

Syntax

```
syncframe()
{
   syncinfo() ;
   bsi() ;
   audfrm() ;
   for(blk = 0; blk < number_of_blocks_per_syncframe; blk++)
   {
      audblk() ;
   }
   auxdata() ;
   errorcheck() ;
} /* end of syncframe */
```

Each of the bit stream elements, and their length, are itemized in the following pseudo code. Note that all bit stream elements arrive most significant bit first, or left bit first, in time.

E.1.2.1 syncinfo - Synchronization information

Syntax	Word size
syncinfo()	
{	
syncword	**16**
} /* end of syncinfo */	

E.1.2.2 bsi - Bit stream information

Syntax	Word size
bsi()	
{	
strmtyp	2
substreamid	3
frmsiz	11
fscod	2
numblkscod	2
acmod	3
lfeon	1
bsid	5
dialnorm	5
compre	1
if(compre) **{compr}**	8
if(acmod == 0x0) /* if 1+1 mode (dual mono, so some items need a second value) */	
{	
dialnorm2	5
compr2e	1
if(compr2e) **{compr2}**	8
}	
if(strmtyp == 0x1) /* if dependent stream */	
{	
chanmape	1
if(chanmape) **{chanmap}**	16
}	
mixmdate	1
if(mixmdate) /* mixing metadata */	
{	
if(acmod > 0x2) /* if more than 2 channels */ **{dmixmod}**	2
if((acmod & 0x1) && (acmod > 0x2)) /* if three front channels exist */	
{	
ltrtcmixlev	3
lorocmixlev	3
}	
if(acmod & 0x4) /* if a surround channel exists */	
{	
ltrtsurmixlev	3
lorosurmixlev	3
}	
if(lfeon) /* if the LFE channel exists */	
{	
lfemixlevcode	1
if(lfemixlevcode) **{lfemixlevcod}**	5
}	
if(strmtyp == 0x0) /* if independent stream */	
{	
pgmscle	1
if(pgmscle) **{pgmscl}**	6
if(acmod == 0x0) /* if 1+1 mode (dual mono, so some items need a second value) */	
{	
pgmscl2e	1
if(pgmscl2e) **{pgmscl2}**	6
}	
extpgmscle	1
if(extpgmscle) **{extpgmscl}**	6
mixdef	2
if(mixdef == 0x1) /* mixing option 2 */	
{	
premixcmpsel	1
drcsrc	1
premixcmpscl	3
}	
else if(mixdef == 0x2) /* mixing option 3 */ **{mixdata}**	12

Syntax	Word size
else if(mixdef == 0x3) /* mixing option 4 */	
{	
mixdeflen	5
mixdata2e	1
if (mixdata2e)	
{	
premixcmpsel	1
drcsrc	1
premixcmpscl	3
extpgmlscle	1
if(extpgmlscle) **{extpgmlscl}**	4
extpgmcscle	1
if(extpgmcscle) **{extpgmcscl}**	4
extpgmrscle	1
if(extpgmrscle) **{extpgmrscl}**	4
extpgmlsscle	1
if(extpgmlsscle) **{extpgmlsscl}**	4
extpgmrsscle	1
if(extpgmrsscle) **{extpgmrsscl}**	4
extpgmlfescle	1
if(extpgmlfescle) **{extpgmlfescl}**	4
dmixscle	1
if(dmixscle) **{dmixscl}**	4
addche	1
if (addche)	
{	
extpgmaux1scle	1
if(extpgmaux1scle) **{extpgmaux1scl}**	4
extpgmaux2scle	1
if(extpgmaux2scle) **{extpgmaux2scl}**	4
}	
}	
mixdata3e	1
if(mixdata3e)	
{	
spchdat	5
addspchdate	1
if(addspchdate)	
{	
spchdat1	5
spchan1att	2
addspchdat1e	1
if(addspdat1e)	
{	
spchdat2	5
spchan2att	3
}	
}	
}	
mixdata	(8*(mixdeflen+2)) - no. mixdata bits
mixdatafill	0 - 7
}	
if(acmod < 0x2) /* if mono or dual mono source */	
{	
paninfoe	1
if(paninfoe)	
{	
panmean	8
paninfo	6
}	
if(acmod == 0x0) /* if 1+1 mode (dual mono - some items need a second value) */	
{	
paninfo2e	1
if(paninfo2e)	
{	
panmean2	8
paninfo2	6
}	
}	
}	
frmmixcfginfoe	1
if(frmmixcfginfoe) /* mixing configuration information */	
{	
if(numblkscod == 0x0) **{blkmixcfginfo[0]}**	5
else	
{	
for(blk = 0; blk < number_of_blocks_per_sync_frame; blk++)	

Syntax	Word size
{	
blkmixcfginfoe	1
if(blkmixcfginfoe)**{blkmixcfginfo[blk]}**	5
}	
}	
}	
}	
}	
infomdate	1
if(infomdate) /* informational metadata */	
{	
bsmod	3
copyrightb	1
origbs	1
if(acmod == 0x2) /* if in 2/0 mode */	
{	
dsurmod	2
dheadphonmod	2
}	
if(acmod >= 0x6) /* if both surround channels exist */ **{dsurexmod}**	2
audprodie	1
if(audprodie)	
{	
mixlevel	5
roomtyp	2
adconvtyp	1
}	
if(acmod == 0x0) /* if 1+1 mode (dual mono, so some items need a second value) */	
{	
audprodi2e	1
if(audprodi2e)	
{	
mixlevel2	5
roomtyp2	2
adconvtyp2	1
}	
}	
if(fscod < 0x3) /* if not half sample rate */ **{sourcefscod}**	1
}	
if((strmtyp == 0x0) && (numblkscod != 0x3)) **{convsync}**	1
if(strmtyp == 0x2) /* if bit stream converted from AC-3 */	
{	
if(numblkscod == 0x3) /* 6 blocks per syncframe */ {blkid = 1}	
else **{blkid}**	1
if(blkid) **{frmsizecod}**	6
}	
addbsie	1
if(addbsie)	
{	
addbsil	6
addbsi	(addbsil+1)×8
}	
} /* end of bsi */	

E.1.2.3 audfrm - Audio frame

Syntax	Word size
audfrm()	
{	
/* these fields for audio frame exist flags and strategy data */	
if(numblkscod == 0x3) /* six blocks per syncframe */	
{	
expstre	1
ahte	1
}	
else	
{	
expstre = 1	
ahte = 0	
}	
snroffststr	2
transproce	1
blkswe	1
dithflage	1

Syntax	Word size
bamode	1
frmfgaincode	1
dbaflde	1
skipflde	1
spxattene	1
/* these fields for coupling data */	
if(acmod > 0x1)	
{	
cplstre[0] = 1	
cplinu[0]	1
for(blk = 1; blk < number_of_blocks_per_sync_frame; blk++)	
{	
cplstre[blk]	1
if(cplstre[blk] == 1) **{cplinu[blk]}**	1
else {cplinu[blk] = cplinu[blk-1]}	
}	
}	
else	
{	
for(blk = 0; blk < number_of_blocks_per_sync_frame; blk++) {cplinu[blk] = 0}	
}	
/* these fields for exponent strategy data */	
if(expstre)	
{	
for(blk = 0; blk < number_of_blocks_per_sync_frame; blk++)	
{	
if(cplinu[blk] == 1) **{cplexpstr[blk]}**	2
for(ch = 0; ch < nfchans; ch++) **{chexpstr[blk][ch]}**	2
}	
}	
else	
{	
ncplblks = 0	
for(blk = 0; blk < number_of_blocks_per_sync_frame; blk++) {ncplblks += cplinu[blk]}	
if((acmod > 0x1) && (ncplblks > 0)) **{frmcplexpstr}**	5
for(ch = 0; ch < nfchans; ch++) **{frmchexpstr[ch]}**	5
/* cplexpstr[blk] and chexpstr[blk][ch] derived from table lookups - see Table E.1.8*/	
}	
if(lfeon)	
{	
for(blk = 0; blk < number_of_blocks_per_sync_frame; blk++) **{lfeexpstr[blk]}**	1
}	
/* These fields for converter exponent strategy data */	
if(strmtyp == 0x0)	
{	
if(numblkscod != 0x3) **{convexpstre}**	1
else {convexpstre = 1}	
if(convexpstre == 1)	
{	
for(ch = 0; ch < nfchans; ch++) **{convexpstr[ch]}**	5
}	
}	
/* these fields for AHT data */	
if(ahte)	
{	
/* coupling can use AHT only when coupling in use for all blocks */	
/* ncplregs derived from cplstre and cplexpstr - see clause E.2.4.2 */	
if((ncplblks == 6) && (ncplregs ==1)) **{cplahtinu}**	1
else {cplahtinu = 0}	
for(ch = 0; ch < nfchans; ch++)	
{	
/* nchregs derived from chexpstr - see clause E.2.4.2 */	
if(nchregs[ch] == 1) **{chahtinu[ch]}**	1
else {chahtinu[ch] = 0}	
}	
if(lfeon)	
{	
/* nlferegs derived from lfeexpstr - see clause E.2.4.2 */	
if(nlferegs == 1) **{lfeahtinu}**	1
else {lfeahtinu = 0}	
}	
}	
/* these fields for audio frame SNR offset data */	
if(snroffststr == 0x0)	
{	
frmcsnroffst	6
frmfsnroffst	4

```
Syntax                                                                    Word size
   }
/* these fields for audio frame transient pre-noise processing data */
   if(transproce)
   {
      for(ch = 0; ch < nfchans; ch++)
      {
         chintransproc[ch] ..................................................... 1
         if(chintransproc[ch])
         {
            transprocloc[ch] ................................................... 10
            transproclen[ch] ..................................................... 8
         }
      }
   }
/* These fields for spectral extension attenuation data */
   if(spxattene)
   {
      for(ch = 0; ch < nfchans; ch++)
      {
         chinspxatten[ch] ..................................................... 1
         if(chinspxatten[ch])
         {
            spxattencod[ch] ..................................................... 5
         }
      }
   }
/* these fields for block start information */
   if (numblkscod != 0x0) {blkstrtinfoe} ..................................... 1
   else {blkstrtinfoe = 0}
   if(blkstrtinfoe)
   {
      /* nblkstrtbits determined from frmsiz (see clause E.1.3.2.27) */
      blkstrtinfo ..................................................... nblkstrtbits
   }
/* these fields for syntax state initialization */
   for(ch = 0; ch < nfchans; ch++)
   {
      firstspxcos[ch] = 1
      firstcplcos[ch] = 1
   }
   firstcplleak = 1
} /* end of audfrm */
```

E.1.2.4 audblk - Audio block

```
Syntax                                                                    Word size
audblk()
{
/* these fields for block switch and dither flags */
   if(blkswe)
   {
      for(ch = 0; ch < nfchans; ch++) {blksw[ch]} ........................... 1
   }
   else
   {
      for(ch = 0; ch < nfchans; ch++) {blksw[ch] = 0}
   }
   if(dithflage)
   {
      for(ch = 0; ch < nfchans; ch++) {dithflag[ch]} ........................ 1
   }
   else
   {
      for(ch = 0; ch < nfchans; ch++) {dithflag[ch] = 1} /* dither on */
   }
/* these fields for dynamic range control */
   dynrnge  1
   if(dynrnge) {dynrng} ..................................................... 8
   if(acmod == 0x0) /* if 1+1 mode */
   {
      dynrng2e ............................................................. 1
      if(dynrng2e) {dynrng2} ............................................... 8
   }
/* these fields for spectral extension strategy information */
```

Syntax	Word size

```
   if(blk == 0) {spxstre = 1}
   else {spxstre} ..... 1
   if(spxstre)
   {
      spxinu ..... 1
      if(spxinu)
      {
         if(acmod == 0x1)
         {
            chinspx[0] = 1
         }
         else
         {
            for(ch = 0; ch < nfchans; ch++) {chinspx[ch]} ..... 1
         }
         spxstrtf ..... 2
         spxbegf ..... 3
         spxendf ..... 3
         if(spxbegf < 6)  {spx_begin_subbnd = spxbegf + 2}
         else  {spx_begin_subbnd = spxbegf * 2 - 3}
         if(spxendf < 3)  {spx_end_subbnd = spxendf + 5}
         else  {spx_end_subbnd = spxendf * 2 + 3}
         spxbndstrce ..... 1
         if(spxbndstrce)
         {
            for(bnd = spx_begin_subbnd+1; bnd < spx_end_subbnd ; bnd++) {spxbndstrc[bnd]} ..... 1
         }

      }
      else /* !spxinu */
      {
         for(ch = 0; ch < nfchans; ch++)
         {
            chinspx[ch] = 0
            firstspxcos[ch] = 1
         }
      }
   }
/* these fields for spectral extension coordinates */
   if(spxinu)
   {
      for(ch = 0; ch < nfchans; ch++)
      {
         if(chinspx[ch])
         {
            if(firstspxcos[ch])
            {
               spxcoe[ch] = 1
               firstspxcos[ch] = 0
            }
            else /* !firstspxcos[ch] */ {spxcoe[ch]} ..... 1
            if(spxcoe[ch])
            {
               spxblnd[ch] ..... 5
               mstrspxco[ch] ..... 2
               /* nspxbnds determined from spx_begin_subbnd, spx_end_subbnd, and spxbndstrc[ ] */
               for(bnd = 0; bnd < nspxbnds; bnd++)
               {
                  spxcoexp[ch][bnd] ..... 4
                  spxcomant[ch][bnd] ..... 2
               }
            }
         }
         else /* !chinspx[ch] */
         {
            firstspxcos[ch] = 1
         }
      }
   }
/* These fields for coupling strategy and enhanced coupling strategy information */
   if(cplstre[blk])
   {
      if (cplinu[blk])
      {
         ecplinu ..... 1
         if (acmod == 0x2)
         {
```

Syntax	Word size

```
            chincpl[0] = 1
            chincpl[1] = 1
         }
         else
         {
            for(ch = 0; ch < nfchans; ch++) {chincpl[ch]} ..... 1
         }
         if (ecplinu == 0) /* standard coupling in use */
         {
            if(acmod == 0x2) {phsflginu} /* if in 2/0 mode */ ..... 1
            cplbegf ..... 4
            if (spxinu == 0) /* if SPX not in use */
            {
               cplendf ..... 4
            }
            else /* SPX in use */
            {
               if (spxbegf < 6)
               {
                  /* note that in this case the value of cplendf may be negative */
                  cplendf = spxbegf - 2
               }
               else
               {
                  cplendf = (spxbegf * 2) - 7
               }

            }
            /* ncplsubnd = 3 + cplendf - cplbegf */
            cplbndstrce ..... 1
            if(cplbndstrce)
            {
               for(bnd = 1; bnd < ncplsubnd; bnd++) {cplbndstrc[bnd]} ..... 1
            }
         }
         else /* enhanced coupling in use */
         {
            ecplbegf ..... 4
            if(ecplbegf < 3) {ecpl_begin_subbnd = ecplbegf * 2}
            else if(ecplbegf < 13) {ecpl_begin_subbnd = ecplbegf + 2}
            else {ecpl_begin_subbnd = ecplbegf * 2 - 10}
            if (spxinu == 0) /* if SPX not in use */
            {
               ecplendf ..... 4
               ecpl_end_subbnd = ecplendf + 7
            }
            else /* SPX in use */
            {
               if (spxbegf < 6)
               {
                  ecpl_end_subbnd = spxbegf + 5
               }
               else
               {
                  ecpl_end_subbnd = spxbegf * 2
               }
            }
            ecplbndstrce ..... 1
            if (ecplbndstrce)
            {
               for(sbnd = max(9, ecpl_begin_subbnd+1); sbnd < ecpl_end_subbnd; sbnd++)
               {
                  ecplbndstrc[sbnd] ..... 1
               }
            }
         } /* ecplinu[blk] */
      }
      else /* !cplinu[blk] */
      {
         for(ch = 0; ch < nfchans; ch++)
         {
            chincpl[ch] = 0
            firstcplcos[ch] = 1
         }
         firstcplleak = 1
         phsflginu = 0
         ecplinu = 0;
```

Syntax	Word size

```
      }
   } /* cplstre[blk] */
/* These fields for coupling coordinates */
   if(cplinu[blk])
   {
      if(ecplinu == 0) /* standard coupling in use */
      {
         for(ch = 0; ch < nfchans; ch++)
         {
            if(chincpl[ch])
            {
               if (firstcplcos[ch])
               {
                  cplcoe[ch] = 1
                  firstcplcos[ch] = 0
               }
               else /* !firstcplcos[ch] */ {cplcoe[ch]} ..... 1
               if(cplcoe[ch])
               {
                  mstrcplco[ch] ..... 2
                  /* ncplbnd derived from ncplsubnd and cplbndstrc */
                  for(bnd = 0; bnd < ncplbnd; bnd++)
                  {
                     cplcoexp[ch][bnd] ..... 4
                     cplcomant[ch][bnd] ..... 4
                  }
               }  /* cplcoe[ch] */
            }
            else /* !chincpl[ch] */
            {
               firstcplcos[ch] = 1;
            }
         }  /* ch */
         if((acmod == 0x2) && phsflginu && (cplcoe[0] || cplcoe[1]))
         {
            for(bnd = 0; bnd < ncplbnd; bnd++) {phsflg[bnd]} ..... 1
         }
      }
      else /* enhanced coupling in use */
      {
         firstchincpl = -1
         reserved 1
         for(ch = 0; ch < nfchans; ch++)
         {
            if(chincpl[ch])
            {
               if(firstchincpl == -1) {firstchincpl = ch}
               if(firstcplcos[ch])
               {
                  ecplparam1e[ch] = 1
                  if (ch > firstchincpl) {rsvdfieldse[ch] = 1}
                  else {rsvdfieldse[ch] = 0}
                  firstcplcos[ch] = 0
               }
               else /* !firstcplcos[ch] */
               {
                  ecplparam1e[ch] ..... 1
                  if(ch > firstchincpl) {rsvdfieldse[ch]}  1
                  else {rsvdfieldse[ch] = 0}
               }
               if(ecplparam1e[ch])
               {
                  /* necplbnd derived from ecpl_begin_subbnd, ecpl_end_subbnd, and ecplbndstrc */
                  for(bnd = 0; bnd < necplbnd; bnd++) {ecplamp[ch][bnd]} ..... 5
               }
               if(rsvdfieldse[ch])
               {
                  reserved ..... 9 x (necplbnd - 1)
               }
               if(ch > firstchincpl) {reserved} 1
            }
            else /* !chincpl[ch] */
            {
               firstcplcos[ch] = 1
            }
         } /* ch */
      } /* ecplinu[blk] */
```

Syntax	Word size

```
   } /* cplinu[blk] */
/* these fields for rematrixing operation in the 2/0 mode */
   if(acmod == 0x2) /* if in 2/0 mode */
   {
      if (blk == 0) {rematstr = 1}
      else {rematstr} ..... 1
      if(rematstr)
      {
         /* nrematbd determined from cplinu, ecplinu, spxinu, cplbegf, ecplbegf and spxbegf */
         for(bnd = 0; bnd < nrematbd; bnd++) {rematflg[bnd]} ..... 1
      }
   }
/* this field for channel bandwidth code */
   for(ch = 0; ch < nfchans; ch++)
   {
      if(chexpstr[blk][ch] != reuse)
      {
         if((!chincpl[ch]) && (!chinspx[ch])) {chbwcod[ch]} ..... 6
      }
   }
/* these fields for exponents */
   if(cplinu[blk]) /* exponents for the coupling channel */
   {
      if(cplexpstr[blk] != reuse)
      {
         cplabsexp ..... 4
         /* ncplgrps derived from cplexpstr, cplbegf, cplendf, ecplinu, ecpl_begin_subbnd, and ecpl_end_subbnd */
         for(grp = 0; grp< ncplgrps; grp++) {cplexps[grp]} ..... 7
      }
   }
   for(ch = 0; ch < nfchans; ch++) /* exponents for full bandwidth channels */
   {
      if(chexpstr[blk][ch] != reuse)
      {
         exps[ch][0] ..... 4
         /* nchgrps derived from chexpstr[ch], and endmant[ch] */
         for(grp = 1; grp <= nchgrps[ch]; grp++) {exps[ch][grp]} ..... 7
         gainrng[ch] ..... 2
      }
   }
   if(lfeon) /* exponents for the low frequency effects channel */
   {
      if(lfeexpstr[blk] != reuse)
      {
         lfeexps[0] ..... 4
         nlfegrps = 2
         for(grp = 1; grp <= nlfegrps; grp++) {lfeexps[grp]} ..... 7
      }
   }
/* these fields for bit-allocation parametric information */
   if(bamode)
   {
      baie  1
      if(baie)
      {
         sdcycod ..... 2
         fdcycod ..... 2
         sgaincod ..... 2
         dbpbcod ..... 2
         floorcod ..... 3
      }
   }
   else
   {
      sdcycod = 0x2
      fdcycod = 0x1
      sgaincod = 0x1
      dbpbcod = 0x2
      floorcod = 0x7
   }
   if(snroffststr == 0x0)
   {
      if(cplinu[blk]) {cplfsnroffst = frmfsnroffst}
      for(ch = 0; ch < nfchans; ch++) {fsnroffst[ch] = frmfsnroffst}
      if(lfeon) {lfefsnroffst = frmfsnroffst}
   }
   else
```

Syntax	Word size

```
   {
      if(blk == 0) {snroffste = 1}
      else {snroffste} ..... 1
      if(snroffste)
      {
         csnroffst ..... 6
         if(snroffststr == 0x1)
         {
            blkfsnroffst ..... 4
            cplfsnroffst = blkfsnroffst
            for(ch = 0; ch < nfchans; ch++) {fsnroffst[ch] = blkfsnroffst}
            lfefsnroffst = blkfsnroffst
         }
         else if(snroffststr == 0x2)
         {
            if(cplinu[blk]) cplfsnroffst ..... 4
            for(ch = 0; ch < nfchans; ch++) {fsnroffst[ch]} ..... 4
            if(lfeon) lfefsnroffst ..... 4
         }
      }
   }
   if(frmfgaincode) {fgaincode} ..... 1
   else {fgaincode = 0}
   if(fgaincode)
   {
      if(cplinu[blk]) {cplfgaincod} ..... 3
      for(ch = 0; ch < nfchans; ch++) {fgaincod[ch]} ..... 3
      if(lfeon) {lfefgaincod} ..... 3
   }
   else
   {
      if(cplinu[blk]) {cplfgaincod = 0x4}
      for(ch= 0; ch < nfchans; ch++) {fgaincod[ch] = 0x4}
      if(lfeon) {lfefgaincod = 0x4}
   }
   if(strmtyp == 0x0)
   {
      convsnroffste ..... 1
      if(convsnroffste) {convsnroffst} ..... 10
   }
   if(cplinu[blk])
   {
      if (firstcplleak)
      {
         cplleake = 1
         firstcplleak = 0
      }
      else /* !firstcplleak */
      {
         cplleake ..... 1
      }
      if(cplleake)
      {
         cplfleak ..... 3
         cplsleak ..... 3
      }
   }
/* these fields for delta bit allocation information */
   if(dbaflde)
   {
      deltbaie ..... 1
      if(deltbaie)
      {
         if(cplinu[blk]) {cpldeltbae} ..... 2
         for(ch = 0; ch < nfchans; ch++) {deltbae[ch]} ..... 2
         if(cplinu[blk])
         {
            if(cpldeltbae==new info follows)
            {
               cpldeltnseg ..... 3
               for(seg = 0; seg <= cpldeltnseg; seg++)
               {
                  cpldeltoffst[seg] ..... 5
                  cpldeltlen[seg] ..... 4
                  cpldeltba[seg] ..... 3
               }
            }
```

Syntax	Word size

```
        }
        for(ch = 0; ch < nfchans; ch++)
        {
           if(deltbae[ch]==new info follows)
           {
              deltnseg[ch] ..... 3
              for(seg = 0; seg <= deltnseg[ch]; seg++)
              {
                 deltoffst[ch][seg] ..... 5
                 deltlen[ch][seg] ..... 4
                 deltba[ch][seg] ..... 3
              }
           }
        }
     } /* if(deltbaie) */
  }/* if(dbaflde) */
/* these fields for inclusion of unused dummy data */
  if(skipflde)
  {
     skiple ..... 1
     if(skiple)
     {
        skipl ..... 9
        skipfld ..... skipl × 8
     }
  }
/* these fields for quantized mantissa values */
  got_cplchan = 0
  for(ch = 0; ch < nfchans; ch++)
  {
     if(chahtinu[ch] == 0)
     {
        for(bin = 0; bin < nchmant[ch]; bin++) {chmant[ch][bin]} ..... (0-16)
     }
     else if(chahtinu[ch] == 1)
     {
        chgaqmod[ch] ..... 2
        if( (chgaqmod[ch] > 0x0) && (chgaqmod[ch] < 0x3) )
        {
           for(n = 0; n < chgaqsections[ch]; n++) {chgaqgain[ch][n]} ..... 1
        }
        else if(chgaqmod[ch] == 0x3)
        {
           for(n = 0; n < chgaqsections[ch]; n++) {chgaqgain[ch][n]} ..... 5
        }
        for(bin = 0; bin < nchmant[ch]; bin++)
        {
           if(chgaqbin[ch][bin])
           {
              for(n = 0; n < 6; n++) {pre_chmant[n][ch][bin]} ..... (0-16)
           }
           else {pre_chmant[0][ch][bin]} ..... (0-9)
        }
        chahtinu[ch] = -1 /* AHT info for this audio frame has been read - do not read again */
     }
     if(cplinu[blk] && chincpl[ch] && !got_cplchan)
     {
        if(cplahtinu == 0)
        {
           for(bin = 0; bin < ncplmant; bin++) {cplmant[bin]} ..... (0-16)
           got_cplchan = 1
        }
        else if(cplahtinu == 1)
        {
           cplgaqmod ..... 2
           if( (cplgaqmod > 0x0) && (cplgaqmod < 0x3) )
           {
              for(n = 0; n < cplgaqsections; n++) {cplgaqgain[n]} ..... 1
           }
           else if(cplgaqmod == 0x3)
           {
              for(n = 0; n < cplgaqsections; n++) {cplgaqgain[n]} ..... 5
           }
           for(bin = 0; bin < ncplmant; bin++)
           {
              if(cplgaqbin[bin])
              {
```

Syntax	Word size
for(n = 0; n < 6; n++) **{pre_cplmant[n][bin]}**	(0-16)
}	
else **{pre_cplmant[0][bin]}**	(0-9)
}	
got_cplchan = 1	
cplahtinu = -1 /* AHT info for this audio frame has been read - do not read again */	
}	
else {got_cplchan = 1}	
}	
}	
if(lfeon) /* mantissas of low frequency effects channel */	
{	
if(lfeahtinu == 0)	
{	
for(bin = 0; bin < nlfemant; bin++) **{lfemant[bin]}**	(0-16)
}	
else if(lfeahtinu == 1)	
{	
lfegaqmod	2
if((lfegaqmod > 0x0) && (lfegaqmod < 0x3))	
{	
for(n = 0; n < lfegaqsections; n++) **{lfegaqgain[n]}**	1
}	
else if(lfegaqmod == 0x3)	
{	
for(n = 0; n < lfegaqsections; n++) **{lfegaqgain[n]}**	5
}	
for(bin = 0; bin < nlfemant; bin++)	
{	
if(lfegaqbin[bin])	
{	
for(n = 0; n < 6; n++) **{pre_lfemant[n][bin]}**	(0-16)
}	
else **{pre_lfemant[0][bin]}**	(0-9)
}	
lfeahtinu = -1 /* AHT info for this audio frame has been read - do not read again */	
}	
}	
} /* end of audblk */	

E.1.2.5 auxdata - Auxiliary data

Syntax	Word size
auxdata()	
{	
auxbits	nauxbits
if(auxdatae)	
{	
auxdatal	14
}	
auxdatae	1
} /* end of auxdata */	

E.1.2.6 errorcheck - Error detection code

Syntax	Word size
errorcheck()	
{	
encinfo	1
crc2	16
} /* end of errorcheck */	

E.1.3 Description of Enhanced AC-3 bit stream elements

E.1.3.0 Introduction

In the definition of some semantic elements relationships with other elements are described for clarity.

E.1.3.1 bsi - Bit stream information

E.1.3.1.1 strmtyp - Stream type - 2 bits

The 2-bit strmtyp code, as shown in Table E.1.1, indicates the stream type.

Table E.1.1: Stream type

strmtyp	Indication
00	Type 0
01	Type 1
10	Type 2
11	Type 3

The stream types are defined as follows:

Type 0: These frames comprise an independent stream or substream. The programme may be decoded independently of any other substreams that might exist in the bit stream.

Type 1: These frames comprise a dependent substream. The programme shall be decoded in conjunction with the independent substream with which it is associated.

Type 2: These frames comprise an independent stream or substream that was previously coded in AC-3. Type 2 streams shall be independently decodable, and may not have any dependent streams associated with them.

Type 3: Reserved.

E.1.3.1.2 substreamid - Substream identification - 3 bits

The 3-bit substreamid field indicates the substream identification parameter. The substream identification parameter can be used, in conjunction with additional bit stream metadata, to enable carriage of a single programme of more than 5.1 channels, multiple programmes of up to 5.1 channels, or a mixture of programmes with up to 5.1 channels and programmes with greater than 5.1 channels.

All Enhanced AC-3 bit streams shall contain an independent substream assigned substream ID 0. The independent substream assigned substream ID 0 shall be the first substream present in the bit stream. If an AC-3 bit stream is present in the Enhanced AC-3 bit stream, then the AC-3 bit stream shall be treated as an independent substream assigned substream ID 0.

Enhanced AC-3 bit streams may also contain up to 7 additional independent substreams assigned substream ID's 1 to 7. Independent substream ID's shall be assigned sequentially in the order the independent substreams are present in the bit stream. Independent substreams 1 to 7 shall contain the same number of blocks per syncframe and be encoded at the same sample rate as independent substream 0.

Each independent substream may have up to 8 dependent substreams associated with it. Dependent substreams shall immediately follow the independent substream with which they are associated. Dependent substreams are assigned substream ID's 0 to 7, which shall be assigned sequentially according to the order the dependent substreams are present in the bit stream. Dependent substreams 0 to 7 shall contain the same number of blocks per syncframe and be encoded at the same sample rate as the independent substream with which they are associated.

For more information about usage of the substreamid parameter, please refer to clause E.2.8.

E.1.3.1.3 frmsiz - Frame size - 11 bits

The frmsiz field indicates a value one less than the overall size of the coded syncframe in 16-bit words. That is, this field may assume a value ranging from 0 to 2 047, and these values correspond to syncframe sizes ranging from 1 to 2 048. Note that some values at the lower end of this range may not be valid, as they may not represent enough words to convey a complete syncframe. It is the responsibility of the encoder to ensure that this does not occur in practice.

E.1.3.1.4 fscod - Sample rate code - 2 bits

The fscod field contains a 2-bit code indicating sample rate according to Table E.1.2.

Table E.1.2: Sample rate codes

fscod	Sample Rate (kHz)
00	48
01	44,1
10	32
11	reserved

E.1.3.1.5 numblkscod - Number of audio blocks - 2 bits

numblkscod: The 2-bit numblkscod code, as shown in Table E.1.3, indicates the number of audio blocks per syncframe.

Table E.1.3: Number of audio blocks per syncframe

numblkscod	Indication
00	1 block per syncframe
01	2 blocks per syncframe
10	3 blocks per syncframe
11	6 blocks per syncframe

E.1.3.1.6 bsid - Bit stream identification - 5 bits

The bsid field has a value of "10000" (=16) in this version of the present document. Future modifications of the present document may define other values. Values of bsid smaller than 16 and greater than 10 shall be used for versions of E-AC-3 which are backward compatible with version 16 decoders. Decoders which can decode version 16 shall be able to decode version numbers less than 16 and greater than 10. Additionally, because E-AC-3 decoders shall also be able to decode AC-3 bit streams, E-AC-3 decoders shall be able to decode versions 0 through 8. If the present document is extended by the addition of additional elements or features that are not compatible with decoders that can decode version 16, a value of bsid greater than 16 will be used. Decoders built to this version of the standard would likely not be able to decode versions with bsid greater than 16. Additionally, decoders built to this version of the standard shall not decode bit streams with bsid = 9 or 10. Thus, decoders built to the present document shall mute if the value of bsid is 9, 10, or greater than 16, and shall decode and reproduce audio if the value of bsid is 0 to 8, or 11 to 16.

E.1.3.1.7 chanmape - Custom channel map exists - 1 bit

If the chanmape bit is set to '0', the channel map for a dependent substream is defined by the audio coding mode (acmod) and LFE on (lfeon) parameters. If this bit is a 1, the following 16 bits define the custom channel map for this dependent substream.

Only dependent substreams can have a custom channel map.

E.1.3.1.8 chanmap - Custom channel map - 16 bits

The 16-bit chanmap field shall specify the custom channel map for a dependent substream. The channel locations supported by the custom channel map are as defined in Table E.1.4. Shaded entries in Table E.1.4 represent channel locations present in the independent substream with which the dependent substream is associated. Non-shaded entries in Table E.1.4 represent channel locations not present in the independent substream with which the dependent substream is associated. These channel locations are defined in SMPTE 428-3 [i.8].

Table E.1.4: Custom channel map locations

chanmap bit	Location
0 (MSB)	Left
1	Centre
2	Right
3	Left Surround
4	Right Surround
5	Lc/Rc pair
6	Lrs/Rrs pair
7	Cs
8	Ts
9	Lsd/Rsd pair
10	Lw/Rw pair
11	Vhl/Vhr pair
12	Vhc
13	Lts/Rts pair
14	LFE2
15	LFE

The custom channel map indicates which coded channels are present in the dependent substream and the order of the coded channels in the dependent substream. Bit 0, which indicates the presence of the left channel, is stored in the most significant bit of the chanmap field. For each channel present in the dependent substream, the corresponding location bit in the chanmap is set to 1. The order of the coded channels in the dependent substream is the same as the order of the enabled location bits in the chanmap. For example, if bits 0, 3, and 4 of the chanmap field are set to 1, and the dependent stream is coded with acmod = 3 and lfeon = 0, the first coded channel in the dependent stream is the Left channel, the second coded channel is the Left Surround channel, and the third coded channel is the Right Surround channel. When the enabled location bit in the chanmap field refers to a pair of channels, this defines the channel location of two adjacent channels in the dependent substream. For example, if bits 3, 4 and 6 of the chanmap field are set to 1, and the dependent stream is coded with acmod = 6 and lfeon = 0, the first coded channel in the dependent stream is the Left Surround channel, the second coded channel is the Right Surround channel, and the third and fourth channels are the Left Rear Surround and Right Rear Surround channels. The number of channel locations indicated by the chanmap field shall equal the total number of coded channels present in the dependent substream, as indicated by the acmod and lfeon bit stream parameters.

For more information about usage of the chanmap parameter, please refer to clause E.2.8.

E.1.3.1.9 mixmdate - Mixing Meta-data exists - 1 bit

If the mixdatae bit is set to '1', mixing and mapping information follows in the bit stream.

E.1.3.1.10 lfemixlevcode - LFE mix level code exists - 1 bit

If the lfemixlevcode bit is set to '1', the LFE mix level code shall follow in the bit stream. If lfemixcode is set to '0', the LFE mix level code shall not be present in the bit stream, and LFE downmixing shall be disabled.

E.1.3.1.11 lfemixlevcod - LFE mix level code - 5 bits

the 5-bit lfemixlevcod field specifies the level at which the LFE channel is mixed into the Left and Right channels during downmixing. The LFE mix level (in dB) shall be derived from the LFE mix level code according to the following formula:

LFE mix level (dB) = 10 - lfemixlevcod

As the valid values for the LFE mix level code are 0 to 31, the valid values for the LFE mix level are therefore +10 to -21 dB. For more information on LFE mixing, please refer to clause E.2.9.

E.1.3.1.12 pgmscle - Programme scale factor exists - 1 bit

If the pgmscle bit is set to '1', the programme scale factor word shall follow in the bit stream. If pgmscle is set to '0', the programme scale factor shall be 0 dB (no scaling).

E.1.3.1.13 pgmscl - Programme scale factor - 6 bits

The pgmscl field specifies a scale factor that shall be applied to the programme during decoding. Valid values are 0 to 63. The value 0 shall be interpreted as mute, and the values 1-63 interpreted as -50 dB to +12 dB of scaling in 1 dB steps.

E.1.3.1.14 pgmscl2e - Programme scale factor #2 exists - 1 bit

The pgmscl2e field shall have the same meaning as pgmscle, except that it applies to the second audio channel when acmod indicates two independent channels (dual mono 1+1 mode).

E.1.3.1.15 pgmscl2 - Programme scale factor #2 - 6 bits

The pgmscl2 field shall have the same meaning as pgmscl, except that it applies to the second audio channel when acmod indicates two independent channels (dual mono 1+1 mode).

E.1.3.1.16 extpgmscle - External programme scale factor exists - 1 bit

If the extpgmscle bit is set to '1', the external programme scale factor word shall follow in the bit stream. If it is 0, the external programme scale factor word shall be 0 dB (no scaling).

E.1.3.1.17 extpgmscl - External programme scale factor - 6 bits

In some applications, two bit streams or independent substreams may be decoded and mixed together. The 6-bit extpgmscl field specifies a scale factor that shall be applied to an external programme during decoding of the external programme. An external programme is defined as a programme that is carried in a separate bit stream or independent substream from the bit stream or independent substream carrying this instance of extpgmscl. The extpgmscl field shall use the same scale as pgmscl.

E.1.3.1.18 mixdef - Mix control field length - 2 bits

The 2-bit mixdef code, as shown in Table E.1.5, shall indicate the mode and parameter field lengths for additional mixing control data carried in the syncframe.

Table E.1.5: Mix control field length

mixdef	Indication
00	mixing option 1, no additional bits
01	mixing option 2, 5 bits
10	mixing option 3, 12 bits reserved
11	mixing option 4, 16-264 bits reserved by mixdeflen

E.1.3.1.19 premixcmpsel - Premix compression word select - 1 bit

If the premixcmpsel bit is set to '0', the dynrng field is used in the premix compression process, otherwise the compr field is used in the premix compression process. The premixcmpsel bit should be set to '0'.

E.1.3.1.20 drcsrc - Dynamic range control word source for the mixed output - 1 bit

If the drcsrc field is set to '0', the dynrng and compr fields of the external programme (i.e. a programme that is carried in a separate bit stream or independent substream) are used to provide dynamic range control data that may be applied to the output of the programm mixing process, otherwise the dynrng and compr fields from the current substream are used. The drcsrc field should be set to '0'.

E.1.3.1.21 premixcmpscl - Premix compression word scale factor - 3 bits

The 3-bit premixcmpscl field indicates the amount of scaling, as shown in Table E.1.6, to be applied to the dynrng and compr fields selected for application to the Main Programme audio before mixing. The premixcmpscl field should be set to '000'.

Table E.1.6: Premix compression word scale factor

premixcmpscl	Scale Factor
000	0 % (no compression)
001	16,7 %
010	33,3 %
011	50 %
100	66,7 %
101	83,3 %
111	100 % (maximum compression)

NOTE: When the mixdef field is set to '01', '10' or '11', the drcsrc, premixcmpsel and premixcmpscl fields are present in the bitstream. However they should be set to the recommended values, as decoders are not required to use them. These fields were originally defined to support a mixing model that was capable of using DRC as well as gain adjustment to control the mixing process.

E.1.3.1.22 mixdeflen - Length of mixing parameter data field - 5 bits

The 5-bit mixdeflen field defines the mixing data field size for the most flexible mixing control mode. mixdeflen = {0, 1, 2, 3 … 31) corresponds to mixdata field length = {2, 3 , 4, 5 … 33} bytes.

E.1.3.1.23 mixdata - Mixing parameter data - (5 to 264) bits

The mixdata field contains control parameters for mixing the programme and external programme streams during decoding.

E.1.3.1.24 mixdata2e - Mixing parameters for individual channel scaling exist - 1 bit

If the mixdata2e field is set to 1, mixing parameters to scale individual channels in an external programme containing up to 7.1 audio channels shall follow in the bit stream.

E.1.3.1.25 extpgmlscle - External programme left channel scale factor exists - 1 bit

If the extpgmlscle field is set to '1', the extpgmlscl field shall follow in the bit stream. If the external programme does not contain a left channel, this field shall be set to 0.

E.1.3.1.26 extpgmlscl - External programme left channel scale factor - 4 bits

The 4-bit extpgmlscl field specifies a scale factor that shall be applied to the left channel of the external programme during mixing. If the extpgmscl field is present in the bit stream, the total gain applied to the left channel of the external programme shall be equal to the sum of the gain values indicated by the extpgmscl and extpgmlscl fields. The extpgmlscl field shall be interpreted as shown in Table E.1.7.

Table E.1.7: external programme left channel scale factor

extpgmlscl	Left channel scale factor(dB)
0	-1
1	-2
2	-3
3	-4
4	-5
5	-6
6	-8
7	-10
8	-12
9	-14
10	-16
11	-19
12	-22
13	-25
14	-28
15	-infinity (mute)

E.1.3.1.27 extpgmcscle - External programme centre channel scale factor exists - 1 bit

If the extpgmcscle field is set to '1', the extpgmcscl field shall follow in the bit stream. If the external programme does not contain a centre channel, this field shall be set to '0'.

E.1.3.1.28 extpgmcscl - External programme centre channel scale factor - 4 bits

The 4-bit extpgmcscl field specifies a scale factor that shall be applied to the centre channel of the external programme during mixing. If the extpgmscl field is present in the bit stream, the total gain applied to the centre channel of the external programme shall be equal to the sum of the gain values indicated by the extpgmscl and extpgmcscl fields. This field shall be coded and interpreted in the same way as extpgmlscl (per Table E.1.7).

E.1.3.1.29 extpgmrscle - External programme right channel scale factor exists - 1 bit

If the extpgmrscle field is set to '1', the extpgmrscl fieldshall follow in the bit stream. If the external programme does not contain a right channel, this field shall be set to '0'.

E.1.3.1.30 extpgmrscl - External programme right channel scale factor - 4 bits

The 4-bit extpgmrscl field specifies a scale factor that shall be applied to the right channel of the external programme during mixing. If the extpgmscl field is present in the bit stream, the total gain applied to the right channel of the external programme shall be equal to the sum of the gain values indicated by the extpgmscl and extpgmrscl fields. This field is coded and interpreted in the same way as extpgmlscl (per Table E.1.7).

E.1.3.1.31 extpgmlsscle - External programme left surround channel scale factor exists - 1 bit

If the extpgmlsscle field is set to '1', the extpgmlsscl field shall follow in the bit stream. If the external programme does not contain a left surround or mono surround channel, this field shall be set to '0'.

E.1.3.1.32 extpgmlsscl - External programme left surround channel scale factor - 4 bits

The 4-bit extpgmlsscl field specifies a scale factor that shall be applied to the left surround or mono surround channel of the external programme during mixing. If the extpgmscl field is present in the bit stream, the total gain applied to the left surround channel of the external programme shall be equal to the sum of the gain values indicated by the extpgmscl and extpgmlsscl fields. This field shall be coded and interpreted in the same way as extpgmlscl (per Table E.1.7).

E.1.3.1.33 extpgmrsscle - External programme right surround channel scale factor exists - 1 bit

If the extpgmrsscle field is set to '1', the extpgmrsscl field shall follow in the bit stream. If the external programme does not contain a right surround channel, this field shall be set to '0'.

E.1.3.1.34 extpgmrsscl - External programme right surround channel scale factor - 4 bits

The 4-bit extpgmrsscl field specifies a scale factor that shall be applied to the right surround channel of the external programme during mixing. If the extpgmscl field is present in the bit stream, the total gain applied to the right surround channel of the external programme shall be the sum of the gain values indicated by the extpgmscl and extpgmrsscl fields. This field shall be coded and interpreted in the same way as extpgmlscl (per Table E.1.7).

E.1.3.1.35 extpgmlfescle - External programme LFE channel scale factor exists - 1 bit

If the extpgmlfescle field is set to '1', the extpgmlfescl field shall follow in the bit stream. If the external programme does not contain a LFE channel, this field shall be set to '0'.

E.1.3.1.36 extpgmlfescl - External programme LFE channel scale factor - 4 bits

The 4-bit extpgmlfescl field specifies a scale factor that shall be applied to the LFE channel of the external programme during mixing. If the extpgmscl field is present in the bit stream, the total gain applied to the LFE channel of the external programme shall be equal to the sum of the gain values indicated by the extpgmscl and extpgmlfescl fields. This field is coded and interpreted in the same way as extpgmlscl (per Table E.1.7).

E.1.3.1.37 dmixscle - External programme downmix scale factor exists - 1 bit

If the dmixscle field is set to '1', the dmixscl field shall follow in the bit stream.

E.1.3.1.38 dmixscl - External programme downmix scale factor - 4 bits

The dmixscl field specifies a scale factor that can be applied to a multichannel external programme that has been downmixed to two channels before individual channel scale factors could be applied. If the extpgmscl field is present in the bit stream, the total gain that shall be applied to the downmixed external programme shall be the sum of the gain values indicated by the extpgmscl and dmixscl fields. This scale factor shall only be applied when the external programme has been downmixed before individual channel scale factors could be applied. This field shall be coded and interpreted in the same way as extpgmlscl (per Table E.1.7).

NOTE: The dmixscl parameter is useful in the case where the main audio has been downmixed to 2-channels inside the AC-3/E-AC-3 decoder, and only two channels are being delivered to the audio mixer. In this situation, individual channel scaling factors that are carried in the E-AC-3 bit stream for the purpose of scaling the decoded multichannel main audio can no longer be applied to their corresponding channels in the audio mixer, as these channels have already been downmixed. The dmixscl parameter allows for additional attenuation to be applied to the downmixed main audio when necessary.

E.1.3.1.39 addche - Scale factors for additional external programme channels exist - 1 bit

When the addche field is set to '1', the extpgmaux1scle and extpgmaux2scle fields shall follow in the bit stream. If the external programme does not contain channel locations other than L, C, R, Ls, Rs and LFE, this field shall be set to '0'.

E.1.3.1.40 extpgmaux1scle - External programme first auxiliary channel scale factor exists - 1 bit

If the extpgmaux1scle field is set to '1', the extpgmaux1scl field shall follow in the stream.

E.1.3.1.41 extpgmaux1scl - External programme first auxiliary channel scale factor - 4 bits

The 4-bit extpgmaux1scl field specifies a scale factor that shall be applied to the first auxiliary channel of the external programme during mixing. If the extpgmscl field is present in the bit stream, the total gain applied to the first auxiliary channel of the external programme shall be equal to the sum of the gain values indicated by the extpgmscl and extpgmaux1scl fields. This field shall be coded and interpreted in the same way as extpgmlscl (per Table E.1.7).

E.1.3.1.42 extpgmaux2scle - External programme second auxiliary channel scale factor exists - 1 bit

If the extpgmaux2scle field is set to '1', the extpgmaux1scl field shall follow in the bit stream. If the external programme does not contain more than 6.1 channels of audio, the extpgmaux2scle field shall be set to '0' when present in the bit stream.

E.1.3.1.43 extpgmaux2scl - External programme second auxiliary channel scale factor - 4 bits

The 4-bit extpgmaux2scl field specifies a scale factor that shall be applied to the second auxiliary channel of the external programme during mixing. If the extpgmscl field is present in the bit stream, the total gain applied to the second auxiliary channel of the external programme shall be equal to the sum of the gain values indicated by the extpgmscl and extpgmaux2scl fields. This field shall be coded and interpreted in the same way as extpgmlscl (per Table E.1.7).

E.1.3.1.44 mixdata3e - Mixing parameters for speech processing exist - 1 bit

If the mixdata3e field is set to '1', information for controlling speech enhancement processing follows in the bit stream.

NOTE: The following fields are placeholders for as yet undefined data to enhance speech intelligibility.

E.1.3.1.45 spchdat - Speech enhancement processing data - 5 bits

The 5-bit spchdat field contains speech enhancement processing parameters, the values of which are determined by the degree to which a first channel or pair of channels of the programme are dominated by speech.

E.1.3.1.46 addspchdate - Additional speech enhancement processing data exists - 1 bit

If the addspchdate field is set to '1', additional information for controlling speech enhancement processing follows in the bit stream.

E.1.3.1.47 spchdat1 - Additional speech enhancement processing data - 5 bits

The 5-bit spchdat1 field contains speech enhancement processing parameters, the values of which are determined by the degree to which a second channel or pair of channels of the programme are dominated by speech.

E.1.3.1.48 spchan1att - Speech enhancement processing attenuation data - 2 bits

The 2-bit spchan1att field shall define which channels in the programme are designated as containing speech information and whether channels not containing speech information may be attenuated.

E.1.3.1.49 addspchdat1e - Additional speech enhancement processing data exists - 1 bit

If the addspchdat1e field is set to '1', additional information for controlling speech enhancement processing follows in the bit stream.

E.1.3.1.50 spchdat2 - Additional speech enhancement processing data - 5 bits

The 5-bit spchdat2 field contains speech enhancement processing parameters, the values of which are determined by the degree to which a third channel or pair of channels of the programme are dominated by speech.

E.1.3.1.51 spchan2att - Speech enhancement processing attenuation data - 3 bits

The 3-bit spchan2att field shall define which additional channels in the programme are designated as containing speech information and whether channels not containing speech information may be attenuated.

E.1.3.1.52 mixdatafill - Mixdata field fill bits - 0 to 7 bits

This mixdatafill field is a variable length field that shall be used to round up the size of the mixdata field to the nearest byte. All bits within mixdatafill shall be set to '0'.

E.1.3.1.53 paninfoe - Pan information exists - 1 bit

If the paninfoe bit is set to '1', panning information shall follow in the bit stream. If paninfoe is set to '0', the panned position of the programme shall be interpreted as "centre".

E.1.3.1.54 panmean - Pan mean direction index - 8 bits

The 8-bit panmean field shall define the mean angle of rotation index relative to the centre position for a panned mono programm source in a two dimensional sound field. A value of 0 indicates the panned mono programme points toward the centre speaker location (defined as 0 degrees). The index indicates 1,5 degree increments in a clockwise rotation. Values 0 to 239 represent 0 to 358,5 degrees, while values 240 to 255 are reserved.

E.1.3.1.55 paninfo - reserved - 6 bits

The 6-bit paninfo field is reserved for future mixing applications.

E.1.3.1.56 paninfo2e - Pan information exists - 1 bit

If the paninfo2e bit is set to '1', panning information #2 shall follow in the bit stream. If paninfo2e is set to '0', the panned position of the programme shall be interpreted as "centre".

E.1.3.1.57 panmean2 - Pan mean direction index - 8 bits

The 8-bit panmean2 field shall have the same meaning as panmean, except that it applies to the second audio channel when acmod indicates two independent channels (dual mono 1+1 mode).

E.1.3.1.58 paninfo2 - reserved - 6 bits

The 6-bit paninfo2 field is reserved for future mixing applications.

E.1.3.1.59 frmmixcfginfoe - Frame mixing configuration information exists - 1 bit

The frmmixcfginfoe bit indicates whether mixing configuration information that applies to the entire syncframe follows in the bit stream. If the frmmixcfginfoe bit is set to '0', no frame mixing configuration information shall follow in the bit stream. If the frmmixcfginfoe bit is set to '1', frame mixing configuration information shall follow in the bit stream.

E.1.3.1.60 blkmixcfginfoe - Block mixing configuration information exists - 1 bit

The blkmixcfginfoe bit indicates whether block mixing configuration information follows in the bit stream. If the flag is set to 0, no block mixing configuration information follows in the bit stream. If the flag is set to 1, block mixing configuration information follows in the bit stream. In the case where the number of blocks per syncframe is 1, the value of this bit shall be inferred as '1' and the bit shall not be present in the bit stream.

E.1.3.1.61 blkmixcfginfo[blk] - Block mixing configuration information - 5 bits

The blkmixcfginfo[blk] field contains block mixing configuration information for the designated audio block.

E.1.3.1.62 infomdate - Informational metadata Exists - 1 bit

If the infomdate bit is set to '1', informational metadata shall follow in the bit stream. The semantics for bsmod, copyrightb, origbs, dsurexmod, audprodie, mixlevel, roomtyp, adconvtyp, audprodi2e, roomtyp2, and adconvtyp2 fields are given in clauses 4.4.2 and D.1.3 above, and the semantics for sourcefscod are given in clause E.1.3.1.63.

E.1.3.1.63 sourcefscod - Source sample rate code - 1 bit

If the sourcefscod bit is a 1, the source material was sampled at twice the sample rate indicated by fscod.

E.1.3.1.64 convsync - Converter synchronization flag - 1 bit

The convsync bit shall be used for synchronization by a device that converts an Enhanced AC-3 bit stream to an AC-3 bit stream.

E.1.3.1.65 blkid - Block identification - 1 bit

If strmtyp indicates a Type 2 bit stream, the blkid bit shall be set to '1' to indicate that the first block in this Enhanced AC-3 syncframe was the first block in the original AC-3 syncframe.

E.1.3.2 audfrm - Audio frame

E.1.3.2.1 expstre - Exponent strategy syntax enabled - 1 bit

If the expstre bit is set to '1', the fields that carry the full exponent strategy syntax shall be present in each audio block. If this bit is set to '0', then the exponent strategy shall be as specified by the frame-based exponent strategy defined in clauses E.1.3.2.12 and E.1.3.2.13.

E.1.3.2.2 ahte - Adaptive hybrid transform enabled - 1 bit

If an Adaptive Hybrid Transform is used to code at least one of the independent channels, the coupling channel, or the LFE channel in the current audio frame, the ahte bit shall be set to '1'. If the entire audio frame is coded using the bit allocation and quantization model as described in clause 6, this bit shall be set to '0'.

E.1.3.2.3 snroffststr - SNR offset strategy - 2 bits

The snroffstr field shall indicate the SNR offset strategy using one of the values defined in Table E.1.8.

Table E.1.8: SNR offset strategy

snroffststr	Indication
00	SNR offset strategy 1
01	SNR offset strategy 2
10	SNR offset strategy 3
11	Reserved

SNR Offset Strategy 1: When SNR Offset Strategy 1 is used, one coarse SNR offset value and one fine SNR offset value are transmitted in the bit stream. These SNR offset values are used for every channel of every block in the audio frame, including the coupling and LFE channels.

SNR Offset Strategy 2: When SNR Offset Strategy 2 is used, one coarse SNR offset value and one fine SNR offset value are transmitted in the bit stream as often as once per block. The fine SNR offset value is used for every channel in the block, including the coupling and LFE channels. For blocks in which coarse and fine SNR offset values are not transmitted in the bit stream, the decoder shall reuse the coarse and fine SNR offset values from the previous block. One coarse and one fine SNR offset value shall be transmitted in block 0.

SNR Offset Strategy 3: When SNR Offset Strategy 3 is used, coarse and fine SNR offset values are transmitted in the bit stream as often as once per block. Separate fine SNR offset values are transmitted for each channel, including the coupling and LFE channels. For blocks in which coarse and fine SNR offset values are not transmitted in the bit stream, the decoder shall reuse the coarse and fine SNR offset values from the previous block. Coarse and fine SNR offset values shall be transmitted in block 0.

E.1.3.2.4 transproce - Transient pre-noise processing enabled - 1 bit

If at least one channel in the current audio frame contains transient pre-noise processing data, the transproce bit shall be set to '1'. If transient pre-noise processing is not being utilized in this audio frame, this bit shall be set to '0'.

E.1.3.2.5 blkswe - Block switch syntax enabled - 1 bit

If the blkswe bit is set to '1', full block switch syntax shall be present in each audio block.

E.1.3.2.6 dithflage - Dither flag syntax enabled - 1 bit

If the dithflage bit is set to '1', full dither flag syntax shall be present in each audio block.

E.1.3.2.7 bamode - Bit allocation model syntax enabled - 1 bit

If the bamode bit is set to '1', full bit allocation syntax shall be present in each audio block.

E.1.3.2.8 frmfgaincode - Fast gain codes enabled - 1 bit

If fast gain codes (per clauses 4.4.3.39, 4.4.3.41 and 4.4.3.43) are transmitted in the bit stream, frmfgaincode shall be set to '1'. If no fast gain codes are transmitted in the bit stream, this bit shall be set to '0'. and default fast gain code values (see clause 7.2.12) are used for every channel of every block in the audio frame.

E.1.3.2.9 dbaflde - Delta bit allocation syntax enabled - 1 bit

If the dbaflde bit is set to '1', full delta bit allocation syntax shall be present in each audio block.

E.1.3.2.10 skipflde - Skip field syntax enabled - 1 bit

If the skipflde bit is set to '1', full skip field syntax shall be present in each audio block.

E.1.3.2.11 spxattene - Spectral extension attenuation enabled - 1 bit

If at least one channel in the current audio frame contains spectral extension attenuation data, the spxattene bit shall be set to '1'. If spectral extension attenuation processing is not being utilized in the audio frame, this bit shall be set to '0'.

E.1.3.2.12 frmcplexpstr - Frame based coupling exponent strategy - 5 bits

The 5-bit frmcplexpstr field shall specify the coupling channel exponent strategy for all audio blocks, as defined in Table E.1.9. Note that exponent strategies D15, D25 and D45 are as defined in clause 6.1.2. "R" indicates that exponents from the previous block shall be reused.

E.1.3.2.13 frmchexpstr[ch] - Frame based channel exponent strategy - 5 bits

The 5-bit frmchexpstr[ch] field specifies the channel exponent strategy for all audio blocks, as defined in Table E.1.9. Note that exponent strategies D15, D25 and D45 are as defined in clause 6.1.2. "R" indicates that exponents from the previous block shall be reused.

E.1.3.2.14 convexpstre - Converter exponent strategy exists - 1 bit

If the convexpstre bit is set to '1', exponent strategy data used by the E-AC-3 to AC-3 converter shall follow in the bit stream. Exponent strategy is required to be provided once every 6 blocks.

E.1.3.2.15 convexpstr[ch] - Converter channel exponent strategy - 5 bits

The 5-bit convexpstr[ch] field shall specify the exponent strategy, as defined in Table E.1.9, for each full bandwidth channel of each block of an AC-3 syncframe converted from a set of 1 or more E-AC-3 frames.

Table E.1.9: Frame exponent strategy combinations

frmcplexpstr	Audio Block Number					
	0	1	2	3	4	5
0	D15	R	R	R	R	R
1	D15	R	R	R	R	D45
2	D15	R	R	R	D25	R
3	D15	R	R	R	D45	D45
4	D25	R	R	D25	R	R
5	D25	R	R	D25	R	D45
6	D25	R	R	D45	D25	R
7	D25	R	R	D45	D45	D45
8	D25	R	D15	R	R	R
9	D25	R	D25	R	R	D45
10	D25	R	D25	R	D25	R
11	D25	R	D25	R	D45	D45
12	D25	R	D45	D25	R	R
13	D25	R	D45	D25	R	D45
14	D25	R	D45	D45	D25	R
15	D25	R	D45	D45	D45	D45
16	D45	D15	R	R	R	R
17	D45	D15	R	R	R	D45
18	D45	D25	R	R	D25	R
19	D45	D25	R	R	D45	D45
20	D45	D25	R	D25	R	R
21	D45	D25	R	D25	R	D45
22	D45	D25	R	D45	D25	R
23	D45	D25	R	D45	D45	D45
24	D45	D45	D15	R	R	R
25	D45	D45	D25	R	R	D45
26	D45	D45	D25	R	D25	R
27	D45	D45	D25	R	D45	D45
28	D45	D45	D45	D25	R	R
29	D45	D45	D45	D25	R	D45
30	D45	D45	D45	D45	D25	R
31	D45	D45	D45	D45	D45	D45

E.1.3.2.16 cplahtinu - Coupling channel AHT in use - 1 bit

If the cplahtinu bit is set to '1', the coupling channel shall be coded using an Adaptive Hybrid Transform. If this bit is set to '0', conventional coupling channel coding shall be used.

E.1.3.2.17 chahtinu[ch] - Channel AHT in use - 1 bit

If the chahtinu[ch] bit is set to '1', channel ch shall be coded using an Adaptive Hybrid Transform. If this bit is set to '0', conventional channel coding shall be used.

E.1.3.2.18 lfeahtinu - LFE channel AHT in use - 1 bit

If the lfeahtinu bit is set to '1', the LFE channel shall be coded using an Adaptive Hybrid Transform. If this bit is set to '0', conventional LFE channel coding shall be used.

E.1.3.2.19 frmcsnroffst - Frame coarse SNR offset - 6 bits

The frmcsnroffst field shall contain the frame coarse SNR offset value. Valid values are 0-63, which shall be interpreted as an offset value of -45 dB to 144 dB in 3 dB steps. This coarse SNR offset value shall be used for every channel of every block in the audio frame, including the coupling and LFE channels.

E.1.3.2.20 frmfsnroffst - Frame fine SNR offset - 4 bits

The frmfsnroffst field shall contain the frame fine SNR offset value. Valid values are 0-15, which shall be interpreted as an offset value of 0 dB to 2,8125 dB in 0,1875 dB steps. This fine SNR offset value shall be used for every channel of every block in the audio frame, including the coupling and LFE channels.

E.1.3.2.21 chintransproc[ch] - Channel in transient pre-noise processing - 1 bit

If the chintransproc[ch] bit is '1', then the corresponding full bandwidth audio channel is required to have associated transient pre-noise processing data.

E.1.3.2.22 transprocloc[ch] - Transient location relative to start of frame - 10 bits

The transprocloc[ch] field shall provide the location of the transient relative to the start of the current audio frame. The transient location (in samples) shall be calculated by multiplying this value by 4. It is possible for the transient to be located in a later audio frame and therefore this number can exceed the number of PCM samples contained within the current audio frame.

E.1.3.2.23 transproclen[ch] - Transient processing length - 8 bits

The transproclen[ch] field shall provide the transient pre-noise processing length in samples, relative to the location of the transient provided by the value of transprocloc[ch].

E.1.3.2.24 chinspxatten[ch] - Channel in spectral extension attenuation processing - 1 bit

If the chinspxatten[ch] bit is set to '1', the channel indicated by the index ch shall be coded using spectral extension attenuation processing. If this bit is set to '0', the channel indicated by the index ch shall not be coded using spectral extension attenuation processing.

E.1.3.2.25 spxattencod[ch] - Spectral extension attenuation code - 5 bits

The 5-bit spxattencod[ch] field shall specify the index into Table E.2.12 from which spectral extension attenuation values for the channel indicated by the index ch are to be derived.

E.1.3.2.26 blkstrtinfoe - Block start information exists - 1 bit

If the blkstrtinfoe bit is set to '1', block start information shall follow in the bit stream. If this bit is set to '0', block start information shall not follow in the bit stream.

E.1.3.2.27 blkstrtinfo - Block start information - nblkstrtbits

The blkstrtinfo field shall contain the block start information. The number of bits of block start information shall be given by the formula:

$$nblkstrtbits = (numblks - 1) \times (4 + ceiling\ (\log_2\ (words\_per\_syncframe)))$$

where:

- numblks is equal to the number of blocks in the frame as indicated by the value of numblkscod per Table E.1.3;
- ceiling(n) is a function that rounds the fractional number n up to the next higher integer.

For example:

- ceiling(2,1) = 3
- $\log_2(n)$ is the base 2 logarithm of n
- words_per_syncframe = frmsiz + 1

E.1.3.2.28 firstspxcos[ch] - First spectral extension coordinates states - 1 bit

The firstspxcos[ch] field determines the state of when new spectral extension coordinates shall be present in the bit stream. If firstspxcos[ch] is set to '1', the spxcoe[ch] bit is assumed to be '1' for the current block and is not transmitted in the bit stream.

E.1.3.2.29 firstcplcos[ch] - First coupling coordinates states - 1 bit

The firstcplcos[ch] field determines the state of when new coupling coordinates shall be present in the bit stream. If firstcplcos[ch] is set to '1', the cplcoe[ch] bit is assumed to be '1' for the current block and is not transmitted in the bit stream.

E.1.3.2.30 firstcplleak - First coupling leak state - 1 bit

The firstcplleak field determines the state of when new coupling leak values shall be present in the bit stream. If firstcplleak is set to '1', the cplleake bit is assumed to be '1' for the current block and is not transmitted in the bit stream.

E.1.3.3 audblk - Audio block

E.1.3.3.1 spxstre - Spectral extension strategy exists - 1 bit

If the spxstre bit is set to '1', spectral extension information shall follow in the bit stream. If it is set to '0', new spectral extension information shall not be present, and spectral extension parameters previously sent shall be reused.

E.1.3.3.2 spxinu - Spectral extension in use - 1 bit

If spxinu bit is set to '1', then the spectral extension technique shall be used in this block. If this bit is set to '0', then the spectral extension technique shall not be used in this block.

E.1.3.3.3 chinspx[ch] - Channel using spectral extension - 1 bit

If the chinspx[ch] bit is set to '1', the channel indicated by the index ch shall utilize spectral extension. If the bit is set to '0', then this channel shall not utilize spectral extension.

E.1.3.3.4 spxstrtf - Spectral extension start copy frequency code - 2 bits

The 2-bit spxstrtf field shall be used to derive the number of the lowest frequency sub-band of the spectral extension copy region. See Table E.2.11 for the definition of the spectral extension sub-bands.

E.1.3.3.5 spxbegf - Spectral extension begin frequency code - 3 bits

The 3-bit spxbegf field shall be used to derive the number of the lowest frequency sub-band of the spectral extension region. The index of the first active spectral extension sub-band shall be equal to spx_begin_subbnd and shall be calculated as shown in the folloiwng pseudo-code:

if (spxbegf < 6) {spx_begin_subbnd = spxbegf + 2}

else {spx_begin_subbnd = spxbegf × 2 - 3}

E.1.3.3.6 spxendf - Spectral extension end frequency code - 3 bits

The 3-bit spxendf field shall be used to derive a number one greater than the highest frequency sub-band of the spectral extension region. The index of one greater than the highest active spectral extension sub-band shall be equal to spx_end_subbnd and shall be calculated as shown in the following pseudo-code:

if (spxendf < 3) {spx_end_subbnd = spxendf + 5}

else {spx_end_subbnd = spxendf × 2 + 3}

E.1.3.3.7 spxbndstrce - Spectral extension band structure exist - 1 bit

If the spxbndstrce bit is set to '1', the spectral extension band structure shall follow. If it is set to '0' in the first block using spectral extension, the default spectral extension band structure shall be used. If it is set to '0' in any other block, the band structure from the previous block shall be reused. The default banding structure defspxbndstrc[] is shown in Table E.1.10.

Table E.1.10: Default spectral extension banding structure

spx sub-band #	defspxbndstrc[]
0	0
1	0
2	0
3	0
4	0
5	0
6	0
7	0
8	1
9	0
10	1
11	0
12	1
13	0
14	1
15	0
16	1

E.1.3.3.8 spxbndstrc[bnd] - Spectral extension band structure - 1 bit to 14 bits

The spxbndstrc[bnd] data structure shall determine the grouping of subbands in spectral extension, and operates in the same fashion as the coupling band structure. For each subband:

- A '0' shall represent the beginning of a new band.
- A '1' shall indicate that the subband shall be combined into the previous band.

Note that it is assumed that the first band begins at the first subband. Therefore, the first band is assumed to be '0' and not sent. The first band in the structure corresponds to the second subband.

E.1.3.3.9 spxcoe[ch] - Spectral extension coordinates exist - 1 bit

If the spxcoe bit is set to '1', spectral extension coordinate information shall follow. If it is set to '0', the spectral extension coordinates from the previous block shall be used.

E.1.3.3.10 spxblnd[ch] - Spectral extension blend - 5 bits

The 5-bit spxblnd[ch] field shall be used to determine, for the channel indicated by the index ch, the noise blending factor (translated signal mixed with random noise) for the spectral extension process.

E.1.3.3.11 mstrspxco[ch] - Master spectral extension coordinate - 2 bits

The 2-bit mstrspxco[ch] field shall be used to establish, for the channel indicated by the index ch, a gain factor (increasing the dynamic range) for the spectral extension coordinates as shown in Table E.1.11.

Table E.1.11: Master Spectral Extension Coordinate

mstrspxco[ch]	spxco[ch][bnd] gain multiplier
00	1
01	2^{-3}
10	2^{-6}
11	2^{-9}

E.1.3.3.12 spxcoexp[ch][bnd] - Spectral extension coordinate exponent - 4 bits

Each spectral extension coordinate is composed of a 4-bit exponent and a 2-bit mantissa. The spxcoexp[ch][bnd] field shall be set to the value of the 4-bit spectral extension coordinate exponent for the channel [ch] and band [bnd]. The index [ch] shall only exist for those channels that are in spectral extension. The index [bnd] will range from zero to nspxbnds.

E.1.3.3.13 spxcomant[ch][bnd] - Spectral extension coordinate mantissa - 2 bits

The spxcomant[ch][bnd] field shall be set to the value of the 2-bit spectral extension coordinate mantissa for the channel [ch] and band [bnd]. The index [ch] shall only exist for those channels that are in spectral extension. The index [bnd] will range from zero to nspxbnds.

E.1.3.3.14 ecplinu - Enhanced coupling in use - 1 bit

If the ecplinu bit is set to '1', enhanced coupling shall be used for the current block. If this bit is set to '0', standard coupling shall be used for the current block.

E.1.3.3.15 cplbndstrce - Coupling banding structure exist - 1 bit

If the cplbndstrce bit is set to '1', the coupling banding structure shall follow. If it is set to '0' in the first block using coupling, the default coupling banding structure shall be used. If it is set to '0' in any other block, the banding structure from the previous block is reused. The default coupling banding structure defcplbndstrc[] shall be as shown in Table E.1.12.

Table E.1.12: Default coupling banding structure

couple sub-band #	defcplbndstrc[]
0	
1	0
2	0
3	0
4	0
5	0
6	0
7	0
8	1
9	0
10	1
11	1
12	0
13	1
14	1
15	1
16	1
17	1

E.1.3.3.16 ecplbegf - Enhanced coupling begin frequency code - 4 bits

The 4-bit ecplbegf field shall be used to derive the number of the first (lowest frequency) active enhanced coupling sub-band) as shown in Table E.2.8. The index of the first active enhanced coupling sub-band shall be equal to ecpl_begin_subbnd and shall be calculated as shown in the following pseudo-code:

if (ecplbegf < 3) {ecpl_begin_subbnd = ecplbegf × 2}

Else if (ecplbegf < 13) {ecpl_begin_subbnd = ecplbegf + 2}

Else {ecpl_begin_subbnd = ecplbegf × 2 - 10}

E.1.3.3.17 ecplendf - Enhanced coupling end frequency code - 4 bits

The 4-bit ecplendf field shall be used to derive a number one greater than the highest frequency sub-band of the enhanced coupling region as shown in Table E.2.8. The index of one greater than the highest active enhanced coupling sub-band shall be equal to ecpl_end_subbnd and shall be calculated as shown in the following pseudo-code:

if (spxinu == 0) {ecpl_end_subbnd = ecplendf + 7}

Else if (spxbegf < 6) {ecpl_end_subbnd = spxbegf + 5}

Else {ecpl_end_subbnd = spxbegf × 2}

E.1.3.3.18 ecplbndstrce - Enhanced coupling banding structure exists - 1 bit

If the ecplbndstrce bit is set to '1', the enhanced coupling banding structure shall follow. If it is set to '0' in the first block using enhanced coupling, the default enhanced coupling band structure shall be used. If it is set to '0' in any other block, the banding structure from the previous block shall be reused. The default enhanced coupling banding structure defecplbndstrc[] shall be as shown in Table E.1.13.

Table E.1.13: Default enhanced coupling banding structure

Enhanced Coupling Sub-Band #	defecplbndstrc[]
0 to 8	0
9	1
10	0
11	1
12	0
13	1
14	0
15	1
16	1
17	1
18	0
19	1
20	1
21	1

E.1.3.3.19 ecplbndstrc[sbnd] - Enhanced coupling band (and sub-band) structure - 1 bit

There are 22 enhanced coupling sub-bands defined in Table E.2.7, each containing either 6 or 12 frequency coefficients. The fixed 12-bin wide enhanced coupling sub-bands 8 and above are converted into enhanced coupling bands, each of which may be wider than (a multiple of) 12 frequency bins. Sub-bands 0 through 7 are never grouped together to form larger enhanced coupling bands, and are thus each considered enhanced coupling bands. Each enhanced coupling band may contain one or more enhanced coupling sub-bands. Enhanced coupling coordinates are transmitted for each enhanced coupling band. Each band's enhanced coupling coordinate shall be applied to all the coefficients in the enhanced coupling band.

The enhanced coupling band structure shall indicate which enhanced coupling sub-bands shall be combined into wider enhanced coupling bands. When ecplbndstrc[sbnd] is set to '0', the sub-band number [sbnd] shall not be combined into the previous band to form a wider band, but shall start a new 12-bin wide enhanced coupling band. When ecplbndstrc[sbnd] is set to '1', then the sub-band [sbnd] shall be combined with the previous band, making the previous band 12 bins wider. For each successive value of ecplbndstrc that is set to '1', the indexed sub-band [sbnd] shall be combined into the current band. When another ecplbndstrc value of 0 is received, then a new band shall be formed, beginning with the 12 bins of the indexed sub-band.

The set of ecplbndstrc[sbnd] values can be considered as an array. Each bit in the array corresponds to a specific enhanced coupling sub-band in ascending frequency order. The elements of the array corresponding to the sub-bands up to and including ecpl_begin_subbnd or 8 (whichever is greater), are always zero, and as the ecplbndstrc bits for these sub-bands are known to be zero, they are not transmitted. Furthermore, if there is only one enhanced coupling sub-band above sub-band 7, then no ecplbndstrc bits are sent.

The total number of enhanced coupling bands, necplbnd, may be computed as shown in the following pseudo-code:

necplbnd = ecpl_end_subbnd - ecpl_begin_subbnd;

necplbnd = ecplbndstrc[ecpl_begin_subbnd] + … + ecplbndstrc[ecpl_end_subbnd -1]

E.1.3.3.20 ecplparam1e[ch] - Enhanced coupling parameters 1 exist - 1 bit

Enhanced coupling parameters are used to derive the enhanced coupling coordinates which indicate, for a given channel and within a given enhanced coupling band, the fraction of the enhanced coupling channel frequency coefficients to use to re-create the individual channel frequency coefficients. Enhanced coupling parameters are conditionally transmitted in the bit stream. If new values are not delivered, the previously sent values remain in effect. See clause E.2.5 for further information on enhanced coupling.

Each enhanced coupling coordinate is derived from a 5-bit amplitude value. If ecplparam1e[ch] is set to '1', the amplitudes for the corresponding channel [ch] exist and shall follow in the bit stream. If the bit is set to '0', the previously transmitted amplitudes for this channel shall be reused. Amplitudes for all channels that are in enhanced coupling shall always be transmitted in the first block in which enhanced coupling is enabled.

E.1.3.3.21 ecplamp[ch][bnd] - Enhanced coupling amplitude scaling - 5 bits

The ecplamp[ch][bnd] field shall be set to the value of the enhanced coupling amplitude for channel [ch] and band [bnd]. The index [ch] shall exist only for those channels that are in enhanced coupling. The index [bnd] shall range from 0 to necplbnds - 1. See clause E.2.5.5 for more information on how to interpret enhanced coupling parameters.

E.1.3.3.22 rsvdfieldse - reserved fields exist - 1 bit

If the rsvdfieldse bit is set to '1', a number of reserved bits follow in the bit stream. The number of reserved bits shall be equal to 9 × (necplbnd - 1), where necplbnd is calculated as shown in clause E.1.3.3.19.

E.1.3.3.23 blkfsnroffst - Block fine SNR offset - 4 bits

The 4-bit blkfsnroffst field shall specify the fine SNR offset value that shall be used by all channels, including the coupling and LFE channels.

E.1.3.3.24 fgaincode - Fast gain codes exist - 1 bit

If the fgaincode bit is set to '1', fast gain codes for each channel shall be transmitted in the bit stream. If this parameter is set to '0' in the first block of the frame, no fast gain codes shall be transmitted in the bit stream, and default fast gain codes (as specified in clause E.1.2.4) shall be used for all blocks in the frame.

E.1.3.3.25 convsnroffste - Converter SNR offset exists - 1 bit

If the convsnroffste bit is set to '1', a SNR offset value for use by an Enhanced AC-3 to AC-3 conversion process shall follow.

E.1.3.3.26 convsnroffst - Converter SNR offset - 10 bits

The 10-bit convsnroffst field shall specify the SNR offset required to convert the current syncframe to a compliant AC-3 syncframe.

E.1.3.3.27 chgaqmod[ch] - Channel gain adaptive quantization mode - 2 bits

The 2-bit chgaqmod[ch] field shall specify which one of four possible quantization modes is used for mantissas in the channel indicated by ch. If the value of chgaqmod[ch] is set to 0, conventional scalar quantization shall be used for channel ch. Otherwise, gain adaptive quantization shall be used and chgaqgain[ch][n] fields shall follow in the bit stream.

E.1.3.3.28 chgaqgain[ch][n] - Channel gain adaptive quantization gain - 1 bit or 5 bits

The chgaqgain[ch][n] field shall signal the adaptive quantizer gain value or values associated with one or more exponents. If the value of chgaqmod[ch] is either 1 or 2, chgaqgain[ch][n] shall be 1 bit in length, signalling two possible gain states. If the value of chgaqmod[ch] is 3, chgaqgain[ch][n] shall be 5 bits in length, representing a triplet of gains coded compositely. In this case, each gain shall signal three possible gain states.

E.1.3.3.29 pre_chmant[n][ch][bin] - Pre channel mantissas - 0 to 16 bits

The pre_chmant[n][ch][bin] field values shall represent the channel mantissas coded either with scalar, vector or gain adaptive quantization.

E.1.3.3.30 cplgaqmod - Coupling channel gain adaptive quantization mode - 2 bits

The 2-bit cplgaqmod field shall specify which one of four possible quantization modes is used for mantissas in the coupling channel. If the value of cplgaqmod is 0, conventional scalar quantization is used. Otherwise, gain adaptive quantization is used and cplgaqgain[n] fields shall follow in the bit stream.

E.1.3.3.31 cplgaqgain[n] - Coupling gain adaptive quantization gain - 1 bit or 5 bits

The cplgaqgain[n] field shall signal the adaptive quantizer gain value or values associated with one or more exponents. If the value of cplgaqmod is either 1 or 2, cplgaqgain[n] is 1 bit in length, signalling two possible gain states. If the value of cplgaqmod is 3, cplgaqgain[n] is 5 bits in length, representing a triplet of gains coded compositely. In this case, each gain shall signal three possible gain states.

E.1.3.3.32 pre_cplmant[n][bin] - Pre coupling channel mantissas - 0 bits to 16 bits

The pre_cplmant[n][bin] field values shall represent the coupling channel mantissas coded either with scalar, vector or gain adaptive quantization.

E.1.3.3.33 lfegaqmod - LFE channel gain adaptive quantization mode - 2 bits

The 2-bit lfegaqmod field shall specify which one of four possible quantization modes is used for mantissas in the LFE channel. If the value of lfegaqmod is set to 0, conventional scalar quantization is used. Otherwise, gain adaptive quantization shall be used and lfegaqgain[n] fields shall follow in the bit stream.

E.1.3.3.34 lfegaqgain[n] - LFE gain adaptive quantization gain - 1 bit or 5 bits

The lfegaqgain[n] field shall signal the adaptive quantizer gain value or values associated with one or more exponents. If the value of lfegaqmod is set to either 1 or 2, lfegaqgain[n] shall be 1 bit in length, signalling two possible gain states. If the value of lfegaqmod is set to 3, lfegaqgain[n] shall be 5 bits in length, representing a triplet of gains coded compositely. In this case, each gain signals three possible gain states.

E.1.3.3.35 pre_lfemant[n][bin] - Pre LFE channel mantissas - 0 bits to 16 bits

The pre_lfemant[n][bin] field values shall represent the LFE channel mantissas coded either with scalar, vector or gain adaptive quantization.

E.2 Decoder processing

E.2.0 Introduction

This clause specifies how Enhanced AC-3 decoders shall process bit streams that use the Enhanced AC-3 bit stream syntax.

E.2.1 Glitch-free switching between different stream types

Enhanced AC-3 decoders should be designed to switch between all supported bit stream types without introducing audible clicks, pops or mutes.

E.2.2 Error detection and concealment

Enhanced AC-3 decoders are required to implement error detection based on the bit stream CRC word. Enhanced AC-3 bit streams contain only one CRC word, which covers the entire syncframe. When decoding bit streams that use the Enhanced AC-3 bit stream syntax, Enhanced AC-3 decoders shall verify the CRC word prior to decoding any of the blocks in the syncframe.

If the CRC word for an Enhanced AC-3 bit stream is found to be invalid, all blocks in the syncframe shall be substituted with an appropriate error concealment signal. For most applications, this can be easily accomplished by simply repeating the last known-good block (before the overlap-add window process).

E.2.3 Modifications to previously defined parameters

E.2.3.0 Introduction

A number of previously defined parameters are utilized differently in this annex than as previously specified. The following modifications apply to devices decoding bit streams adhering to the syntax specified in this annex.

E.2.3.1 cplendf - Coupling end frequency code

With the addition of the spectral extension tool, the determination of the coupling end frequency code has changed, as shown in clause E.1.2.4. When spectral extension processing is in use (spxinu == 1), the coupling end frequency code parameter (cplendf) is not transmitted in the bit stream. Instead, it is derived from the spectral extension begin frequency code parameter (spxbegf). It should be noted that when spectral extension processing is in use, the range of allowable values of the coupling end frequency code parameter changes from 0 to 15 to -2 to 7. All subsequent operations referencing the coupling end frequency code parameter remain valid and are unchanged.

E.2.3.2 nrematbd - Number of rematrixing bands

With the addition of the spectral extension and enhanced coupling tools, the determination of the number of rematrixing bands had changed. The following pseudo code demonstrates how to determine nrematbd.

Pseudo Code

```
if (cplinu)
{
    if (ecplinu)
    {
        if (ecplbegf == 0) {nrematbd = 0}
        else if (ecplbegf == 1) {nrematbd = 1}
        else if (ecplbegf == 2) {nrematbd = 2}
        else if (ecplbegf < 5) {nrematbd = 3}
        else {nrematbd = 4}
    }
    else /* standard coupling */
    {
        if (cplbegf == 0) {nrematbd = 2}
```

Pseudo Code

```
        else if (cplbegf < 3) {nrematbd = 3}
        else {nrematbd = 4}
    }
}
else if (spxinu)
{
    if (spxbegf < 2) {nrematbd = 3}
    else {nrematbd = 4}
}
else
{
    nrematbd = 4
}
```

E.2.3.3 endmant - End mantissa

With the addition of the spectral extension and enhanced coupling tools, the determination of the end mantissa bin number has changed. The following pseudo-code demonstrates how to determine endmant[ch].

Pseudo Code

```
if (ecplinu)
{
    endmant[ch] = ecplsubbndtab[ecpl_begin_subbnd]
}
else
{
    if ((spxinu) && (cplinu == 0)) {endmant[ch] = spxbandtable[spx_begin_subbnd]
}
else
{
    /* see clause 6.1.3*/
}
```

E.2.3.4 nchmant - Number of fbw channel mantissas

Although not previously stated in any version of the present document, the parameter nchmant[ch] is equivalent to the parameter endmant[ch].

E.2.3.5 ncplgrps - Number of coupled exponent groups

With the addition of the enhanced coupling tool, the determination of the number of coupled exponent groups has changed. The following pseudo-code demonstrates how to determine ncplgrps.

Pseudo Code

```
if (ecplinu)
{
    ecplstartmant = ecplsubbndtab[ecpl_begin_subbnd];
    ecplendmant = ecplsubbndtab[ecpl_end_subbnd];
    if (cplexpstr == D15) {ncplgrps = (ecplendmant - ecplstartmant) / 3}
    else if (cplexpstr == D25) {ncplgrps = (ecplendmant - ecplstartmant) / 6}
    else if (cplexpstr == D45) {ncplgrps = (ecplendmant - ecplstartmant) / 12}
}
else /* standard coupling */
{
    /* see clause 6.1.3 */
}
```

E.2.4 Adaptive Hybrid Transform processing

E.2.4.1 Overview

The AHT is composed of two linear transforms connected in cascade. The first transform is identical to that employed in AC-3 - a windowed Modified Discrete Cosine Transform (MDCT) of length 128 or 256 frequency samples. This feature provides compatibility with legacy AC-3 decoders without the need to return to the time domain in the decoder. For frames containing audio signals which are not time-varying in nature (stationary), a second transform can optionally be applied by the encoder, and inverted by the decoder. The second transform is composed of a non-windowed, non-overlapped Discrete Cosine Transform (DCT Type II). When this DCT is employed, the effective audio transform length increases from 256 to 1 536 audio samples. This results in significantly improved coding gain and perceptual coding performance for stationary signals.

The AHT transform is enabled by setting the ahte bit stream parameter to 1. If ahte is 1, at least one of the independent channels, the coupling channel, or the LFE channel has been coded with AHT. The chahtinu[ch], cplahtinu, and lfeahtinu bit stream parameters are used to indicate which channels are channels coded with AHT.

In order to realize gains made available by the AHT, the AC-3 scalar quantizers have been augmented with two new coding tools. When AHT is in use, both 6-dimensional vector quantization (VQ) and gain-adaptive quantization (GAQ) are employed. VQ is employed for the largest step sizes (coarsest quantization), and GAQ is employed for the smallest stepsizes (finest quantization). The selection of quantizer step size is performed using the same parametric bit allocation method as AC-3, except the conventional bit allocation pointer (bap) table is replaced with a high-efficiency bap table (hebap[]). The hebap[] table employs finer-granularity than the conventional bap table, enabling more efficient allocation of bits.

E.2.4.2 Bit stream helper variables

Several helper variables need to be computed during the decode process in order to decode a syncframe containing at least one channel using AHT (ahte = 1). These variables are not transmitted in the bit stream itself, but are computed from other bit stream parameters. The first helper variables of this type are denoted in the bit stream syntax as ncplregs, nchregs[ch], and nlferegs. The method for computing these variables is presented in the following three sections of pseudo code. Generally speaking, the nregs variables are set equal to the number of times exponents are transmitted in the syncframe.

Pseudo Code

```
/* only compute ncplregs if coupling in use for all 6 blocks */
ncplregs = 0;
/* AHT is only available in 6 block mode (numblkscod ==0x3) */
for(blk = 0; blk < 6; blk++)
{
    if( (cplstre[blk] == 1) || (cplexpstr[blk] != reuse) )
    {
        ncplregs++;
    }
}
```

Pseudo Code

```
for(ch = 0; ch < nfchans; ch++)
{
    nchregs[ch] = 0;
    /* AHT is only available in 6 block mode (numblkscod ==0x3) */
    for(blk = 0; blk < 6; blk++)
    {
        if(chexpstr[blk][ch] != reuse)
        {
            nchregs[ch]++;
        }
    }
}
```

Pseudo Code

```
nlferegs = 0;
/* AHT is only available in 6 block mode (numblkscod ==0x3) */
for(blk = 0; blk < 6; blk++)
{
    if( lfeexpstr[blk] != reuse)
    {
        nlferegs++;
    }
}
```

A second set of helper variables are required for identifying which and how many mantissas employ GAQ. The arrays identifying which bins are GAQ coded are called chgaqbin[ch][bin], cplgaqbin[bin], and lfegaqbin[bin]. Since the number and position of GAQ-coded mantissas varies from syncframe to syncframe, these variables need to be computed after the corresponding hebap[] array is available, but prior to mantissa unpacking. This procedure is shown in pseudo-code below.

Pseudo Code

```
if(cplahtinu == 0)
{
    for(bin = cplstrtmant; bin < cplendmant; bin++)
    {
        cplgaqbin[bin] = 0;
    }
}
else
{
    if (cplgaqmod < 2)
    {
        endbap = 12;
    }
    else
    {
        endbap = 17;
    }
    cplactivegaqbins = 0;
    for(bin = cplstrtmant; bin < cplendmant; bin++)
    {
        if(cplhebap[bin] > 7 && cplhebap[bin] < endbap)
        {
            cplgaqbin[bin] = 1;     /* Gain word is present */
            cplactivegaqbins++;
        }
        else if (cplhebap[bin] >= endbap)
        {
            cplgaqbin[bin] = -1;    /* Gain word is not present */
        }
        else
        {
            cplgaqbin[bin] = 0;
        }
    }
}
```

Pseudo Code

```
for(ch = 0; ch < nfchans; ch++)
{
    if(chahtinu[ch] == 0)
    {
        for(bin = 0; bin < endmant[ch]; bin++)
        {
            chgaqbin[ch][bin] = 0;
        }
    }
    else
    {
        if (chgaqmod < 2)
        {
            endbap = 12;
        }
        else
        {
            endbap = 17;
        }
```

Pseudo Code

```
        chactivegaqbins[ch] = 0;
        for(bin = 0; bin < endmant[ch]; bin++)
        {
            if(chhebap[ch][bin] > 7 && chhebap[ch][bin] < endbap)
            {
                chgaqbin[ch][bin] = 1;  /* Gain word is present */
                chactivegaqbins[ch]++;
            }
            else if (chhebap[ch][bin] >= endbap)
            {
                chgaqbin[ch][bin] = -1; /* Gain word not present */
            }
            else
            {
                chgaqbin[ch][bin] = 0;
            }
        }
    }
}
```

Pseudo Code

```
if(lfeahtinu == 0)
{
    for(bin = 0; bin < lfeendmant; bin++)
    {
        lfegaqbin[bin] = 0;
    }
}
else
{
    if (lfegaqmod < 2)
    {
        endbap = 12;
    }
    else
    {
        endbap = 17;
    }
    lfeactivegaqbins = 0;
    for(bin = 0; bin < lfeendmant; bin++)
    {
        if(lfehebap[bin] > 7 && lfehebap[bin] < endbap)
        {
            lfegaqbin[bin] = 1;     /* Gain word is present */
            lfeactivegaqbins++;
        }
        else if (lfehebap[bin] >= endbap)
        {
            lfegaqbin[bin] = -1;    /* Gain word is not present */
        }
        else
        {
            lfegaqbin[bin] = 0;
        }
    }
}
```

In a final set of helper variables, the number of gain words to be read from the bit stream is computed. These variables are called chgaqsections[ch], cplgaqsections, and lfegaqsections for the independent channels, coupling channel, and LFE channel, respectively. They denote the number of GAQ gain words transmitted in the bit stream, and are computed as shown in the following pseudo code.

Pseudo Code

```
if(cplahtinu == 0)
{
    cplgaqsections = 0;
}
else
{
switch(cplgaqmod)
    {
        case 0: /* No GAQ gains present */
        {
```

Pseudo Code

```
            cplgaqsections = 0;
            break;
        }
        case 1: /* GAQ gains 1 and 2 */
        case 2: /* GAQ gains 1 and 4 */
        {
            cplgaqsections = cplactivegaqbins;  /* cplactivegaqbins was computed earlier */
            break;
        }
        case 3: /* GAQ gains 1, 2, and 4 */
        {
            cplgaqsections = cplactivegaqbins / 3;
if (cplactivegaqbins % 3) cplgaqsections++;
            break;
        }
    }
}
```

Pseudo Code

```
for(ch = 0; ch <nfchans; ch ++)
{
    if(chahtinu[ch] == 0)
    {
        chgaqsections[ch] = 0;
    }
    else
    {
        switch(chgaqmod[ch])
        {
            case 0: /* No GAQ gains present */
            {
                chgaqsections[ch] = 0;
                break;
            }
            case 1: /* GAQ gains 1 and 2 */
            case 2: /* GAQ gains 1 and 4 */
            {
                chgaqsections[ch] = chactivegaqbins[ch];
                /* chactivegaqbins[ch] was computed earlier */
                break;
            }
            case 3: /* GAQ gains 1, 2, and 4 */
            {
                chgaqsections[ch] = chactivegaqbins[ch] / 3;
                if (chactivegaqbins[ch] % 3) chgaqsections[ch]++;
                break;
            }
        }
    }
}
```

Pseudo Code

```
if(lfeahtinu == 0)
{
    lfegaqsections = 0;
}
else
{
    sumgaqbins = 0;
    for(bin = 0; bin < lfeendmant; bin++)
    {
        sumgaqbins += lfegaqbin[bin];
    }
    switch(lfegaqmod)
    {
        case 0: /* No GAQ gains present */
        {
            lfegaqsections = 0;
            break;
        }
        case 1: /* GAQ gains 1 and 2 */
        case 2: /* GAQ gains 1 and 4 */
        {
            lfegaqsections = lfeactivegaqbins;  /* lfeactivegaqbins was computed earlier */
```

Pseudo Code

```
            break;
        }
        case 3: /* GAQ gains 1, 2, and 4 */
        {
            lfegaqsections = lfeactivegaqbins / 3;
            if(lfeactivegaqbins % 3) lfegaqsections++;
            break;
        }
    }
}
```

If the gaqmod bit stream parameter bits are set to 0, conventional scalar quantization is used in place of GAQ coding. If the gaqmod bits are set to 1 or 2, a 1-bit gain is present for each mantissa coded with GAQ. If the gaqmod bits are set to 3, the GAQ gains for three individual mantissas are compositely coded as a 5-bit word.

E.2.4.3 Bit allocation

E.2.4.3.0 Introduction

When AHT is in use for any independent channel, the coupling channel, or the LFE channel, higher coding efficiency is achieved by allowing quantization noise to be allocated with higher precision. The higher precision allocation is achieved using a combination of a new bit allocation pointer look up table and vector quantization. The following clauses describe the changes to the bit allocation routines defined in clause 6.2 in order to achieve higher precision allocation.

E.2.4.3.1 Parametric bit allocation

If the ahtinu flag is set for any independent channel, the coupling channel, or the LFE channel then the bit allocation routine (defined in clause 6.2.2) for that channel is modified to incorporate the new high efficiency bit allocation pointers. When AHT is in use, the exponents are first decoded and the PSD, excitation function, and masking curve are calculated. The delta bit allocation, if present in the bit stream, is then applied (in accordance with clause 6.2.2.6). The final computation of the bit allocation, however, is modified as follows:

The high efficiency bit allocation array (hebap[]) is now computed. The masking curve, adjusted by the snroffset and then truncated, is subtracted from the fine-grain psd[] array. The difference is right shifted by 5 bits, limited, and then used as an address into the hebaptab[] to find the final bit allocation and quantizer type applied to the mantissas. The hebaptab[] array is shown in Table E.2.1.

At the end of the bit allocation procedure, shown in the following pseudo-code, the hebap[] array contains a series of 5-bit pointers. The pointers indicate how many bits have been allocated to each mantissa and the type of quantizer applied to the mantissas. The correspondence between the hebap pointer and quantizer type and quantizer levels is shown in Table E.2.2.

Pseudo Code

```
if(ahtinu == 1) /* cplAHTinu, chAHTinu[ch], or lfeAHTinu */
{
    i = start ;
    j = masktab[start] ;
    do
    {
        lastbin = min(bndtab[j] + bndsz[j]), end);
        mask[j] -= snroffset ;
        mask[j] -= floor ;
        if (mask[j] < 0)
        {
            mask[j] = 0 ;
        }
        mask[j] &= 0x1fe0 ;
        mask[j] += floor ;
        for (k = i; k < lastbin; k++)
        {
            address = (psd[i] - mask[j]) >> 5 ;
            address = min(63, max(0, address)) ;
            hebap[i] = hebaptab[address] ;
            i++ ;
        }
        j++;
    }
    while (end > lastbin) ;
}
else
{
    i = start ;
    j = masktab[start] ;
    do
    {
        lastbin = min(bndtab[j] + bndsz[j], end);
        mask[j] -= snroffset ;
        mask[j] -= floor ;
        if (mask[j] < 0)
        {
            mask[j] = 0 ;
        }
        mask[j] &= 0x1fe0 ;
        mask[j] += floor ;
        for (k = i; k < lastbin; k++)
        {
            address = (psd[i] - mask[j]) >> 5 ;
            address = min(63, max(0, address)) ;
            bap[i] = baptab[address] ;
            i++ ;
        }
        j++;
    }
    while (end > lastbin) ;
}
```

Note that if AHT is not in use for a given independent channel, the coupling channel, or the LFE channel, then the bit allocation procedure and resulting bap[] arrays for that channel are the same as described in clause 6.2.

E.2.4.3.2 Bit allocation tables

Table E.2.1: High efficiency bit allocation pointers, hebaptab[]

Address	hebaptab[address]	Address	hebaptab[address]
0	0	32	14
1	1	33	14
2	2	34	14
3	3	35	15
4	4	36	15
5	5	37	15
6	6	38	15
7	7	39	16
8	8	40	16
9	8	41	16

Address	hebaptab[address]	Address	hebaptab[address]
10	8	42	16
11	8	43	17
12	9	44	17
13	9	45	17
14	9	46	17
15	10	47	18
16	10	48	18
17	10	49	18
18	10	50	18
19	11	51	18
20	11	52	18
21	11	53	18
22	11	54	18
23	12	55	19
24	12	56	19
25	12	57	19
26	12	58	19
27	13	59	19
28	13	60	19
29	13	61	19
30	13	62	19
31	14	63	19

Table E.2.2: Quantizer type, quantizer level, and mantissa bits vs. hebap

hebap	Quantizer Type	Levels	Mantissa Bits
0	NA	NA	0
1	VQ	NA	(2/6)
2	VQ	NA	(3/6)
3	VQ	NA	(4/6)
4	VQ	NA	(5/6)
5	VQ	NA	(7/6)
6	VQ	NA	(8/6)
7	VQ	NA	(9/6)
8	symmetric + GAQ	7	3
9	symmetric + GAQ	15	4
10	symmetric + GAQ	31	5
11	symmetric + GAQ	63	6
12	symmetric + GAQ	127	7
13	symmetric + GAQ	255	8
14	symmetric + GAQ	511	9
15	symmetric + GAQ	1 023	10
16	symmetric + GAQ	2 047	11
17	symmetric + GAQ	4 095	12
18	symmetric + GAQ	16,383	14
19	symmetric + GAQ	65,535	16

E.2.4.4 Quantization

E.2.4.4.0 Introduction

Depending on the bit allocation pointer (hebap) calculated in clause E.2.4.2, the mantissa values are either coded using vector quantization or gain adaptive quantization. The following clauses describe both of these coding techniques.

E.2.4.4.1 Vector quantization

Vector quantization is a quantization technique that takes advantage of similarities and patterns in an ordered series of values, or vector, to reduce redundancy and hence improve coding efficiency. For AHT processing, 6 mantissa values across blocks within a single spectral bin are grouped together to create a 6-dimensional Euclidean space.

If AHT is in use and the bit allocation pointer is between 1 and 7 inclusive, then vector quantization (VQ) is used to encode the mantissas. The range of hebap values that use VQ are shown in Table E.2.2. If VQ is applied to a set of 6 mantissa values then the data in the bit stream represents an N bit index into a 6-dimensional look up table, where N is dependent on the hebap value as defined in Table E.2.2. The vector tables are shown in clause E.3; the values in the vector tables are represented as 16-bit, signed values.

If a hebap value is within the VQ range, the encoder selects the best vector to transmit to the decoder by locating the vector which minimizes the Euclidean distance between the actual mantissa vector and the table vector. The index of the closest matching vector is then transmitted to the decoder.

In the decoder, the index is read from the bit stream and the mant values are replaced with the values from the appropriate vector table.

E.2.4.4.2 Gain adaptive quantization

Gain-Adaptive Quantization (GAQ) is a method for quantizing mantissas using variable-length codewords. In the encoder, the technique is based upon conditionally amplifying one or more of the smaller and typically more frequently occurring transform coefficient mantissas in one DCT block, and representing these with a shorter length code. Larger transform coefficients are not gain amplified, but are transmitted using longer codes since these occur relatively infrequently for typical audio signals. The gain words selected by the encoder, one per GAQ-coded DCT block of length six, are packed together with the mantissa codewords and transmitted as side information. With this system, the encoder can adapt to changing local signal statistics from audio frame to audio frame, and/or from channel to channel. Since a coding mode using constant-length output symbols is included as a subset, gain-adaptive quantization cannot cause a noticeable coding loss compared to the fixed-length codes used in AC-3.

In the decoder, the individual gain words are unpacked first, followed by a bit stream parsing operation (using the gains) to reconstruct the individual transform coefficient mantissas. To compensate for amplification applied in the encoder, the decoder applies an attenuation factor to the small mantissas. The level of large mantissas is unaffected by these gain factors in both the encoder and decoder.

The decoder structure for gain-adaptive quantization is presented in Figure E.2.1. Decoder processing consists of a bit stream deformatter connected in cascade with the switched gain attenuation element, labelled as $1/G_k$ in the figure. The three inputs to the deformatter are the packed mantissa bit stream, the hebap[] output from the parametric bit allocation, and the gaqgain[] array received from the encoder. The hebap[] array is used by the deformatter to determine if the current (k^{th}) DCT block of six mantissas to be unpacked is coded with GAQ, and if so, what the small and large mantissa bit lengths are. The gaqgain[] array is processed by the deformatter to produce the gain attenuation element corresponding to each DCT mantissa block identified in the bit stream. The switch position is also derived by the deformatter for each GAQ-coded mantissa. The switch position is determined from the presence or absence of a unique bit stream tag, as discussed in the next paragraph. When the deformatting operation is complete, the dequantized and level-adjusted mantissas are available for the next stage of processing.

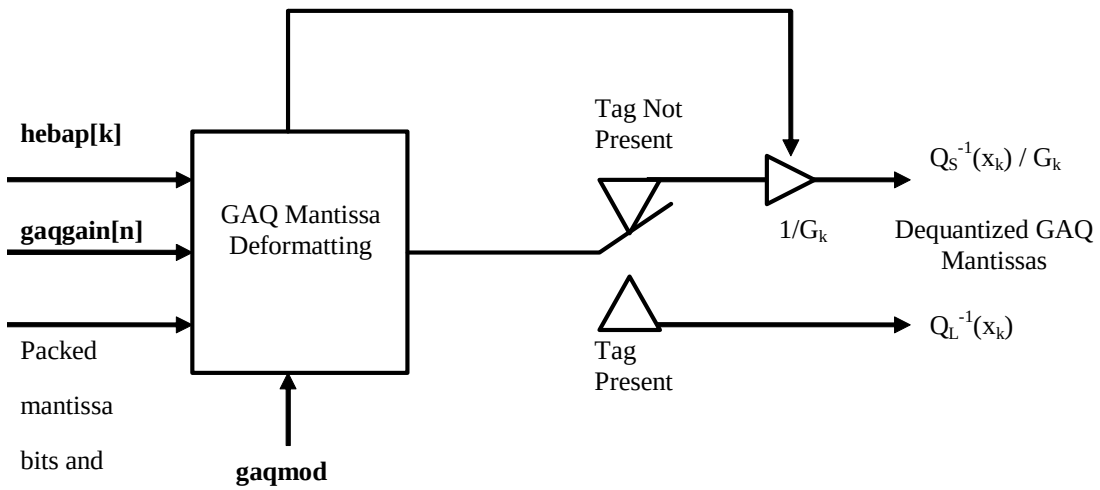

Figure E.2.1: Flow diagram for GAQ mantissa dequantization

As a means for signalling the two mantissa lengths to the decoder, quantizer output symbols for large mantissas are flagged in the bit stream using a unique identifier tag. In Enhanced AC-3, the identifier tag is the quantizer symbol representing a full-scale negative output (e.g. the "100" symbol for a 3-bit two's complement quantizer). In a conventional mid-tread quantizer, this symbol is often deliberately unused since it results in an asymmetric quantizer characteristic. In gain-adaptive quantization, this symbol is employed to indicate the presence of a large mantissa. The tag length is equal to the length of the small mantissa codeword (computed from hebap[] and gaqgain[]), allowing unique bit stream decoding. If an identifier tag is found, additional bits immediately following the tag (also of known length) convey the quantizer output level for the corresponding large mantissas.

Four different gain transmission modes are available for use in the encoder. The different modes employ switched 0, 1 or 1,67-bit gains. For each independent, coupling, and LFE channel in which AHT is in use, a 2-bit parameter called gaqmod is transmitted once per syncframe to the decoder. The bit stream parameters, values, and active hebap range are shown for each mode in Table E.2.3. If gaqmod = 0x0, GAQ is not in use and no gains are present in the bit stream. If gaqmod = 0x1, a 1-bit gain value is present for each block of DCT coefficients having a hebap value between 8 and 11, inclusive. Coefficients with hebap higher than 11 are decoded using the same quantizer as gaqmod 0x0. If gaqmod = 0x2 or 0x3, gain values are present for each block of DCT coefficients having a hebap value between 8 and 16, inclusive. Coefficients with hebap higher than 16 are decoded using the same quantizer as gaqmod 0x0. The difference between the two last modes lies in the gain word length, as shown in Table E.2.3.

Table E.2.3: Gain adaptive quantization modes

chgaqmod[ch], cplgaqmod, and lfegaqmod	GAQ Mode for Audio Frame	Active hebap Range (for which gains are transmitted)
0x0	GAQ not in use	None
0x1	1-bit gains (G_k = 1 or 2)	8 ≤ hebap ≤ 11
0x2	1-bit gains (G_k = 1 or 4)	8 ≤ hebap ≤ 16
0x3	1,67 bit gains (G_k = 1, 2, or 4)	8 ≤ hebap ≤ 16

For the case of gaqmod = 0x1 and 0x2, the gains are coded using binary 0 to signal G_k = 1, and binary 1 to signal G_k = 2 or 4. For the case of gaqmod = 0x3, the gains are composite-coded in triplets (three 3-state gains packed into 5-bit words). The gains are unpacked in a manner similar to exponent unpacking as described in clause 6.1.3. For example, for a 5-bit composite gain triplet grpgain:

M1 = truncate (grpgain / 9)

M2 = truncate ((grpgain % 9) / 3)

M3 = (grpgain % 9) % 3

In this example, M1, M2, and M3 correspond to mapped values derived from consecutive gains in three ascending frequency blocks, respectively, each ranging in value from 0 to 2 inclusive as shown in Table E.2.4.

Table E.2.4: Mapping of gain elements, gaqmod = 0x3

Gain, G_k	Mapped Value
1	0
2	1
4	2

Details of the gain adaptive quantizer characteristics are shown in Table E.2.5. If the received gain is 1, or no gain was received at all, a single quantizer with no tag is used. If the received gain is either 2 or 4, both the small and large mantissas (and associated tags) shall be decoded using the quantizer characteristics shown. Both small and large mantissas are decoded by interpreting them as signed two's complement fractional values. The variable m in the table represents the number of mantissa bits associated with a given hebap value as shown in Table E.2.2.

Table E.2.5: Gain adaptive quantizer characteristics

	$G_k = 1$	$G_k = 2$		$G_k = 4$	
	Quantizer	Small Quantizer	Large Quantizer	Small Quantizer	Large Quantizer
Length of quantizer codeword	m	m - 1	m - 1	m - 2	m
Number of reconstruction (output) points	$2^m - 1$	$2^{m-1} - 1$	2^{m-1}	$2^{m-2} - 1$	2^m
Step size	$2 / (2^m - 1)$	$1 / (2^{m-1})$	$1 / (2^{m-1} - 1)$	$1 / (2^{m-1})$	$3 / (2^{m+1} - 2)$

Since the large mantissas are coded using a dead-zone quantizer, a post-processing step is required to transform (remap) large mantissa codewords received by the decoder into a reconstructed mantissa. This remapping is applied when $G_k = 2$ or 4. An identical post-processing step is required to implement a symmetric quantizer characteristic when $G_k = 1$, and for all gaqmod = 0x0 quantizers. The post-process is a computation of the form y = x + ax + b. In this equation, x represents a mantissa codeword (interpreted as a signed two's complement fractional value), and the constants a and b are provided in Table E.2.6. The constants are also interpreted as 16-bit signed two's complement fractional values. The expression for y was arranged for implementation convenience so that all constants will have magnitude less than one. For decoders where this is not a concern, the remapping can be implemented as y = a'x + b, where the new coefficient a' = 1 + a. The sign of x shall be tested prior to retrieving b from the table. Remapping is not applicable to the table entries marked N/A.

Table E.2.6: Large mantissa inverse quantization (remapping) constants

hebap		$G_k = 1$		$G_k = 2$		$G_k = 4$	
		a	b	a	b	a	b
8	x ≥ 0	0x1249	0x0000	0xd555	0x4000	0xedb7	0x2000
	x < 0	0x1249	0x0000	0xd555	0xeaab	0xedb7	0xfb6e
9	x ≥ 0	0x0889	0x0000	0xc925	0x4000	0xe666	0x2000
	x < 0	0x0889	0x0000	0xc925	0xd249	0xe666	0xeccd
10	x ≥ 0	0x0421	0x0000	0xc444	0x4000	0xe319	0x2000
	x < 0	0x0421	0x0000	0xc444	0xc889	0xe319	0xe632
11	x ≥ 0	0x0208	0x0000	0xc211	0x4000	0xe186	0x2000
	x < 0	0x0208	0x0000	0xc211	0xc421	0xe186	0xe30c
12	x ≥ 0	0x0102	0x0000	0xc104	0x4000	0xe0c2	0x2000
	x < 0	0x0102	0x0000	0xc104	0xc208	0xe0c2	0xe183
13	x ≥ 0	0x0081	0x0000	0xc081	0x4000	0xe060	0x2000
	x < 0	0x0081	0x0000	0xc081	0xc102	0xe060	0xe0c1
14	x ≥ 0	0x0040	0x0000	0xc040	0x4000	0xe030	0x2000
	x < 0	0x0040	0x0000	0xc040	0xc081	0xe030	0xe060
15	x ≥ 0	0x0020	0x0000	0xc020	0x4000	0xe018	0x2000
	x < 0	0x0020	0x0000	0xc020	0xc040	0xe018	0xe030
16	x ≥ 0	0x0010	0x0000	0xc010	0x4000	0xe00c	0x2000
	x < 0	0x0010	0x0000	0xc010	0xc020	0xe00c	0xe018
17	x ≥ 0	0x0008	0x0000	N/A	N/A	N/A	N/A
	x < 0	0x0008	0x0000	N/A	N/A	N/A	N/A
18	x ≥ 0	0x0002	0x0000	N/A	N/A	N/A	N/A
	x < 0	0x0002	0x0000	N/A	N/A	N/A	N/A
19	x ≥ 0	0x0000	0x0000	N/A	N/A	N/A	N/A
	x < 0	0x0000	0x0000	N/A	N/A	N/A	N/A

E.2.4.5 Transform equations

The AHT processing uses a DCT to achieve higher coding efficiency. Hence, if AHT is in use, the DCT shall be inverted prior to applying the exponents. The inverse DCT (IDCT) for AHT is given in the following equation. Any fast technique may be used to invert the DCT in Enhanced AC-3 decoders. In the following equation, $C(k,m)$ is the MDCT spectrum for the kth bin and mth block, and $X(k,j)$ is the AHT spectrum for the kth bin and jth block.

$$C(k,m) = \sqrt{2}\sum_{j=0}^{5} R_j X(k,j)\cos\left(\frac{j(2m+1)\pi}{12}\right) \quad m = 0,1,\ldots,5$$

Where:

$$R_j = \begin{cases} 1 & j \neq 0 \\ 1/\sqrt{2} & j = 0 \end{cases}$$

and *k* is the bin index, *m* is the block index, and *j* is the AHT transform index.

E.2.5 Enhanced channel coupling

E.2.5.1 Overview

Enhanced channel coupling is a coding technique that elaborates on conventional channel coupling, principally through the use of a more compact amplitude representation.The process may be used at lower frequencies than conventional channel coupling.

The decoder converts the enhanced coupling channel back into individual channels by applying amplitude scaling for each channel and frequency sub-band.

E.2.5.2 Sub-band structure for enhanced coupling

Enhanced coupling transform coefficients are transmitted in exactly the same manner as conventional coupling. That is, coefficients are reconstructed from exponents and quantized mantissas. Transform coefficients # 13 through # 252 are grouped into 22 sub-bands of either 6 or 12 coefficients each, as shown in Table E.2.7. The parameter ecplbegf is used to derive the value ecpl_begin_subbnd which indicates the number of the enhanced coupling sub-band which is the first to be included in the enhanced coupling process. Below the frequency (or transform coefficient number) indicated by ecplbegf, all channels are independently coded. Above the frequency indicated by ecplbegf, channels included in the enhanced coupling process (chincpl[ch] = 1) share the common enhanced coupling channel up to the frequency (or tc #) indicated by ecplendf. The enhanced coupling channel is coded up to the frequency (or tc #) indicated by ecplendf, which is used to derive ecpl_end_subbnd. The value of ecpl_end_subbnd is one greater than the last coupling sub-band which is coded.

Table E.2.7: Enhanced coupling sub-bands

enhanced coupling sub-band #	low tc #	high tc #	lf cutoff (kHz) @ fs = 48 kHz	hf cutoff (kHz) @ fs = 48 kHz	lf cutoff (kHz) @ fs = 44,1 kHz	hf cutoff (kHz) @ fs = 44,1 kHz
0	13	18	1,17	1,73	1,08	1,59
1	19	24	1,73	2,30	1,59	2,11
2	25	30	2,30	2,86	2,11	2,63
3	31	36	2,86	3,42	2,63	3,14
4	37	48	3,42	4,55	3,14	4,18
5	49	60	4,55	5,67	4,18	5,21
6	61	72	5,67	6,80	5,21	6,24
7	73	84	6,80	7,92	6,24	7,28
8	85	96	7,92	9,05	7,28	8,31
9	97	108	9,05	10,17	8,31	9,35
10	109	120	10,17	11,30	9,35	10,38
11	121	132	11,30	12,42	10,38	11,41
12	133	144	12,42	13,55	11,41	12,45
13	145	156	13,55	14,67	12,45	13,48
14	157	168	14,67	15,80	13,48	14,51
15	169	180	15,80	16,92	14,51	15,55
16	181	192	16,92	18,05	15,55	16,58
17	193	204	18,05	19,17	16,58	17,61
18	205	216	19,17	20,30	17,61	18,65
19	217	228	20,30	21,42	18,65	19,68
20	229	240	21,42	22,55	19,68	20,71
21	241	252	22,55	23,67	20,71	21,75
NOTE: At 32 kHz sampling rate the sub-band frequency ranges are 2/3 the values of those for 48 kHz.						

Table E.2.8: Enhanced coupling start and end indexes

ecpl sub-band #	low tc #	high tc #	ecplbegf	ecplendf
0	13	18	0	
1	19	24		
2	25	30	1	
3	31	36		
4	37	48	2	
5	49	60	3	
6	61	72	4	
7	73	84	5	0
8	85	96	6	1
9	97	108	7	2
10	109	120	8	3
11	121	132	9	4
12	133	144	10	5
13	145	156	11	6
14	157	168	12	7
15	169	180		8
16	181	192	13	9
17	193	204		10
18	205	216	14	11
19	217	228		12
20	229	240	15	13
21	241	252		14
22	253			15

The enhanced coupling sub-bands are combined into enhanced coupling bands for which coupling coordinates are generated (and included in the bit stream). The coupling band structure is indicated by ecplbndstrc[sbnd]. Each bit of the ecplbndstrc[] array indicates whether the sub-band indicated by the index is combined into the previous (lower in frequency) enhanced coupling band. Enhanced coupling bands are thus made from integral numbers of enhanced coupling sub-bands. (See clause E.1.3.3.19.)

E.2.5.3 Enhanced coupling tables

Table E.2.9 and Table E.2.10 are used to lookup various parameter values used by the enhanced coupling process.

Table E.2.9: Sub-band transform start coefficients: ecplsubbndtab[]

sbnd	ecplsubbndtab[sbnd]
0	13
1	19
2	25
3	31
4	37
5	49
6	61
7	73
8	85
9	97
10	109
11	121
12	133
13	145
14	157
15	169
16	181
17	193
18	205
19	217
20	229
21	241

sbnd	ecplsubbndtab[sbnd]
22	253

Table E.2.10: Amplitudes: ecplampexptab[], ecplampmanttab[]

ecplamp	ecplampexptab[ecplamp]	ecplampmanttab[ecplamp]
0	0	0x20
1	0	0x1b
2	0	0x17
3	0	0x13
4	0	0x10
5	1	0x1b
6	1	0x17
7	1	0x13
8	1	0x10
9	2	0x1b
10	2	0x17
11	2	0x13
12	2	0x10
13	3	0x1b
14	3	0x17
15	3	0x13
16	3	0x10
17	4	0x1b
18	4	0x17
19	4	0x13
20	4	0x10
21	5	0x1b
22	5	0x17
23	5	0x13
24	5	0x10
25	6	0x1b
26	6	0x17
27	6	0x13
28	6	0x10
29	7	0x1b
30	7	0x17
31	-	0x00

E.2.5.4 Enhanced coupling coordinate format

Enhanced coupling coordinates exist for each enhanced coupling band [bnd] in each channel [ch] which is coupled (chincp[ch]==1). Enhanced coupling coordinates are derived from a 5-bit amplitude scaling value (ecplamp[ch][bnd]). These values will always be transmitted in the first block containing a coupled channel and are optionally transmitted in subsequent blocks, as indicated by the enhanced coupling parameter exists flag (ecplparam1e[ch]). If ecplparam1e[ch] is set to 0, corresponding coordinate values from the previous block are reused.

The ecplamp values 0 to 30 represent gains between 0 dB and -45,01 dB quantized to increments of approximately 1,5 dB, and the value 31 represents minus infinity dB.

E.2.5.5 Enhanced coupling processing

E.2.5.5.0 Introduction

This clause describes the processing steps required to recover transform coefficients for each coupled channel from the enhanced coupling data. The following steps are performed for each block:

- Process the enhanced coupling channel.
- Prepare amplitudes for each channel and band.

- Generate transform coefficients for each channel from the processed enhanced coupling channel and amplitudes.

E.2.5.5.1 Amplitude parameter processing

Amplitude values for each enhanced coupling band [bnd] in each channel [ch] are obtained from the ecplamp parameters as:

Pseudo Code

```
if (ecplamp[ch][bnd] == 31)
{
    ampbnd[ch][bnd] = 0;
}
else
{
    ampbnd[ch][bnd] = (ecplampmanttab[ecplamp[ch][bnd]] / 32) >> ecplampexptab[ecplamp[ch][bnd]];
}
```

Using the ecplbndstrc[] array, an array indicating the number of bins in each enhanced coupling band is populated. Additionally, the amplitude values ampbnd[ch][bnd] which apply to enhanced coupling bands are converted to values which apply to enhanced coupling sub-bands ampbnd[ch][sbnd] by duplicating values as indicated by values of "1" in ecplbndstrc[]. Amplitude values for individual enhanced coupling transform coefficients ampbin[ch][bin] are then reconstructed as follows:

Pseudo Code

```
necplbnds = -1;
for(sbnd=ecpl_begin_subbnd; sbnd<ecpl_end_subbnd; sbnd++)
{
    if(ecplbndstrc[sbnd] == 0)
    {
        necplbnds++;
        nbins_per_bnd_array[necplbnds] = 0;
    }
    for(bin=ecplsubbndtab[sbnd]; bin<ecplsubbndtab[sbnd+1]; bin++)
    {
        ampbin[ch][bin] = ampbnd[ch][necplbnds];
        nbins_per_bnd_array[necplbnds]++;
    }
}
```

E.2.5.5.2 Generate channel transform coefficients

Individual channel transform coefficients are then reconstructed from the coupling channel using the following process.

Pseudo Code

```
for(sbnd=ecpl_begin_subbnd; sbnd<ecpl_end_subbnd; sbnd++)
{
    for(bin=ecplsubbndtab[sbnd]; bin<ecplsubbndtab[sbnd+1]; bin++)
    {
        chmant[ch][bin] = ecplmant[bin] * ampbnd[ch][sbnd];
    }
}
```

E.2.6 Spectral extension processing

E.2.6.0 Introduction

Enhanced AC-3 supports a coding technique, based on high frequency regeneration, called spectral extension. A detailed description of the spectral extension process follows.

E.2.6.1 Overview

When spectral extension is in use, high frequency transform coefficients of the channels that are participating in spectral extension are synthesized. Transform coefficient synthesis involves copying low frequency transform coefficients, inserting them as high frequency transform coefficients, blending the inserted transform coefficients with pseudo-random noise, and scaling the blended transform coefficients to match the coarse (banded) spectral envelope of the original signal. To enable the decoder to scale the blended transform coefficients to match the spectral envelope of the original signal, scale factors are computed by the encoder and transmitted to the decoder on a banded basis for all channels participating in the spectral extension process. For a given channel and spectral extension band, the blended transform coefficients for that channel and band are multiplied by the scale factor associated with that channel and band.

The spectral extension process is performed beginning at the spectral extension begin frequency, and ending at the spectral extension end frequency. The spectral extension begin frequency is derived from the spxbegf bit stream parameter. The spectral extension end frequency is derived from the spxendf bit stream parameter.

In some cases, it may be desirable to use channel coupling for a mid-range portion of the frequency spectrum and spectral extension for the higher-range portion of the frequency spectrum. In this configuration, the highest coupled transform coefficient number shall be 1 less than the lowest transform coefficient number generated by spectral extension.

E.2.6.2 Sub-band structure for spectral extension

Transform coefficients #25 through #228 are grouped into 17 sub-bands of 12 coefficients each, as shown in Table E.2.11. The final table entry does not represent an actual sub-band, but is included for the case when the spxendf parameter is 17. The spectral extension sub-bands containing transform coefficients #37 through #228 coincide with coupling sub-bands. The parameter spx_begin_subbnd, derived from the spxbegf bit stream parameter, indicates the number of the first spectral extension sub-band. The parameter spx_end_subbnd, derived from the spxendf bit stream parameter, indicates a number one greater than the last spectral extension sub-band. From the sub-band indicated by spx_begin_subbnd to the sub-band indicated by spx_end_subbnd, transform coefficients are synthesized for all channels participating in the spectral extension process (chinspx[ch] == 1). Below the sub-band indicated by spx_begin_subbnd, channels may be independently coded. Alternatively, channels may be coded independently below the coupling begin frequency, and coupled from the coupling begin frequency to the spectral extension begin frequency.

Spectral extension sub-bands are combined into spectral extension bands for which spectral extension coordinates are generated (and included in the bit stream). Like channel coupling, each spectral extension band is made up of one or more consecutive spectral extension sub-bands. The number of spectral extension bands and the size of each band are determined from the spectral extension band structure array (spxbndstrc[]). Upon audio frame initialization, the default spectral extension banding structure is copied into the spxbndstrc[] array. If (spxbndstrce == 1), the spxbndstrc[sbnd] bit stream parameters are present in the bit stream and are used to fill the spxbndstrc[] array. If (spxbndstrce == 0), the existing values in the spxbndstrc[] array are used to compute the number of spectral extension bands and the size of each band.

The following pseudo code indicates how to determine the number of spectral extension bands and the size of each band.

Pseudo Code

```
nspxbnds = 1;
spxbndsztab[0] = 12;

for (bnd = spx_begin_subbnd+1; bnd < spx_end_subbnd; bnd ++)
{
    if (spxbndstrc[bnd] == 0)
    {
        spxbndsztab[nspxbnds] = 12;
        nspxbnds++;
    }
    else
    {
        spxbndsztab[nspxbnds - 1] += 12;
    }
}
```

Table E.2.11: Spectral extension band table

spx sub-band #	low tc #	high tc #	spxbegf	spxendf
0	25	36		
1	37	48		
2	49	60	0	
3	61	72	1	
4	73	84	2	
5	85	96	3	0
6	97	108	4	1
7	109	120	5	2
8	121	132		
9	133	144	6	3
10	145	156		
11	157	168	7	4
12	169	180		
13	181	192		5
14	193	204		
15	205	216		6
16	217	228		
17	229			7

E.2.6.3 Spectral extension coordinate format

Spectral extension coordinates exist for each spectral extension band [bnd] of each channel [ch] that is using spectral extension (chinspx[ch] ==1). Spectral extension coordinates shall be sent at least once per audio frame, and may be sent as often as once per block. The spxcoe[ch] bit stream parameter informs the decoder when spectral extension coordinates are present in the bit stream. If (spxcoe[ch] == 0), no spectral extension coordinates for channel [ch] are present in the bit stream, and the previous spectral extension coordinates should be reused. If (spxcoe[ch] == 1), spectral extension coordinates are present in the bit stream for channel [ch].

When present in the bit stream, spectral extension coordinates are transmitted in a floating point format. The exponent is sent as a 4-bit value (spxcoexp[ch][bnd]) indicating the number of right shifts which should be applied to the fractional mantissa value. The mantissas are sent as 2-bit values (spxcomant[ch][bnd]) which shall be properly scaled before use. Mantissas are unsigned values so a sign bit is not used. Except for the limiting case where the exponent value = 15, the mantissa value is known to be between 0,5 and 1,0. Therefore, when the exponent value < 15, the MSB of the mantissa is always equal to "1" and is not transmitted; the next 2 bits of the mantissa are transmitted. This provides one additional bit of resolution. When the exponent value = 15 the mantissa value is generated by dividing the 2-bit value of spxcomant by 4. When the exponent value is < 15 the mantissa value is generated by adding 4 to the 2-bit value of spxcomant and then dividing the sum by 8.

Spectral extension coordinate dynamic range is increased beyond what the 4-bit exponent can provide by the use of a per channel 2-bit master spectral extension coordinate (mstrspxco[ch]) which is used to scale all of the spectral extension coordinates within that channel. The exponent values for each channel are increased by 3 times the value of mstrspxco which applies to that channel. This increases the dynamic range of the spectral extension coordinates by an additional 54 dB.

The following pseudo code indicates how to generate the spectral extension coordinate (spxco) for each spectral extension band [bnd] in each channel [ch].

Pseudo code

```
if (spxcoexp[ch][bnd] == 15)
{
    spxco_temp[ch][bnd] = spxcomant[ch][bnd] / 4;
}
else
{
    spxco_temp[ch][bnd] = (spxcomant[ch][bnd] + 4) / 8;
}
spxco[ch][bnd] = spxco_temp[ch][bnd] >> (spxcoexp[ch][bnd] + 3*mstrspxco[ch]);
```

E.2.6.4 High frequency transform coefficient synthesis

E.2.6.4.0 Introduction

This process synthesizes transform coefficients above the spectral extension begin frequency. The synthesis process consists of a number of different steps, described in the following clauses.

E.2.6.4.1 Transform coefficient translation

The first step of the high frequency transform coefficient synthesis process is transform coefficient translation. Transform coefficient translation consists of making copies of a channel's low frequency transform coefficients and inserting them as the channel's high frequency transform coefficients. The parameter spxstrtf, derived from the bit stream parameter of the same name, is used as the index into a table to determine the first transform coefficient to be copied. The parameter spx_begin_subbnd, derived from the spxbegf bit stream parameter, is used as the index into a table to determine the first transform coefficient to be inserted. The parameter spx_end_subbnd, derived from the spxendf bit stream parameter, is used as the index into a table to determine the last transform coefficient to be inserted.

Transform coefficient translation is performed on a banded basis. For each spectral extension band, coefficients are copied sequentially starting with the transform coefficient at copyindex and ending with the transform coefficient at (copyindex + bandsize - 1). Transform coefficients are inserted sequentially starting with the transform coefficient at insertindex and ending with the transform coefficient at (insertindex + bandsize - 1).

Prior to beginning the translation process for each band, the value of (copyindex + bandsize - 1) is compared to the copyendmant parameter. If (copyindex + bandsize - 1) is greater than or equal to the copyendmant parameter, the copyindex parameter is reset to the copystartmant parameter and wrapflag[bnd] is set to 1. Otherwise, wrapflag[bnd] is set to 0.

Pseudo Code

```
copystartmant = spxbandtable[spxstrtf];
copyendmant = spxbandtable[spx_begin_subbnd];
copyindex = copystartmant;
insertindex = copyendmant;
for (bnd = 0; bnd < nspxbnds; band++)
{
    bandsize = spxbndsztab[bnd];
    if ((copyindex + bandsize) > copyendmant)
    {
        copyindex = copystartmant;
        wrapflag[bnd] = 1;
    }
    else
    {
        wrapflag[bnd] = 0;
    }
    for (bin = 0; bin < bandsize; bin++)
    {
        if (copyindex == copyendmant)
        {
            copyindex = copystartmant;
        }
        tc[ch][insertindex] = tc[ch][copyindex];
        insertindex++;
        copyindex++;
    }
}
```

E.2.6.4.2 Transform coefficient noise blending

E.2.6.4.2.0 Introduction

The next step of the high frequency transform coefficient synthesis process is transform coefficient noise blending. In this step, the translated transform coefficients are blended with pseudo-random noise in order to create a more natural sounding signal.

E.2.6.4.2.1 Blending factor calculation

The first step of the transform coefficient noise blending process is to determine blending factors for the pseudo-random noise and the translated transform coefficients. The blending factor calculation for each band is based on both the spxblend bit stream parameter and the frequency mid-point of the band. This enables unique blending factors to be computed for each band from a single bit stream parameter. Because the spxblnd parameter exists in the bit stream only when new spectral extension coordinates exist in the bit stream, the blending factors can be reused for all blocks in which spectral extension coordinates are reused.

The following pseudo code indicates how the blending factors for a channel [ch] are determined.

Pseudo Code

```
noffset[ch] = spxblend[ch] / 32,0;
spxmant = spxbandtable[spx_begin_subbnd];
if  (spxcoe[ch])
{
    for (bnd = 0; bnd < nspxbnds; bnd++)
    {
        bandsize = spxbndsztab[bnd];
        nratio = ((spxmant + 0,5*bandsize) / spxbandtable[spx_end_subbnd]) - noffset[ch];
        if (nratio < 0,0)
        {
            nratio = 0,0;
        }
        else if (nratio > 1,0)
        {
            nratio = 1,0;
        }
        nblendfact[ch][bnd] = squareroot(nratio);
        sblendfact[ch][bnd] = squareroot(1 - nratio);
        spxmant += bandsize;
    }
}
```

E.2.6.4.2.2 Banded RMS energy calculation

The next step is to compute the banded RMS energy of the translated transform coefficients. The banded RMS energy measures are needed to properly scale the pseudo-random noise samples prior to blending.

The following pseudo code indicates how to compute the banded RMS energy of the translated transform coefficients for channel [ch].

Pseudo Code

```
spxmant = spxbandtab[spx_begin_subbnd];
for (bnd = 0; bnd < nspxbnds; bnd++)
{
    bandsize = spxbndsztab[bnd];
    accum = 0;
    for (bin = 0; bin < bandsize; bin++)
    {
        accum = accum + (tc[ch][spxmant] * tc[ch][spxmant]);
        spxmant++;
    }
    rmsenergy[ch][band] = squareroot(accum / bandsize);
}
```

E.2.6.4.2.3 Noise Scaling and Transform Coefficient Blending Calculation

When spectral extension attenuation is enabled for channel [ch], a notch filter is applied to the transform coefficients surrounding the border between the baseband and extension region. The filter is symmetric about the first bin of the extension region, and covers a total of 5 bins. The first 3 attenuation values of the filter are determined by lookup into Table E.2.12 with index spxattencod[ch]. The last two attenuation values of the filter are determined by symmetry and are not explicitly stored in the table. The filter is also applied to the transform coefficients surrounding each border between bands where wrapping occurs during the transform coefficient translation operation, as indicated by wrapflag[bnd]. It is important that filtering occurs after the transform coefficient translation and banded RMS energy calculation but prior to the noise scaling and transform coefficient blending calculation. The following pseudo code demonstrates the application of the notch filter at the border between the baseband and extension region and all wrap points for each channel [ch].

Pseudo Code

```
if (chinspxatten[ch])
{
    /* apply notch filter at baseband / extension region border */
    filtbin = spxbandtable[spx_begin_subbnd] - 2;

    for (bin = 0; bin < 3; bin++)
    {
        tc[ch][filtbin] *= spxattentab[spxattencod[ch]][binindex];
        filtbin++;
    }
    for (bin = 1; bin >= 0; bin--)
    {
        tc[ch][filtbin] *= spxattentab[spxattencod[ch]][binindex];
        filtbin++;
    }
    filtbin += spxbndsztab[0];

    /* apply notch at all other wrap points */
    for (bnd = 1; bnd < nspxbnds; bnd++)
    {
        if (wrapflag[bnd])      /* wrapflag[bnd] set during transform coefficient translation */
        {
            filtbin = filtbin - 5;
            for (binindex = 0; binindex < 3; bin++)
            {
                tc[ch][filtbin] *= spxattentab[spxattencod[ch]][binindex];
                filtbin++;
            }
            for (bin = 1; bin >= 0; bin--)
            {
                tc[ch][filtbin] *= spxattentab[spxattencod[ch]][binindex];
                filtbin++;
            }
        }
        filtbin += spxbndsztab[bnd];
    }
}
```

Table E.2.12: Spectral extension attenuation table: spxattentab[][]

spxattencod	binindex		
	0	1	2
0	0,954841604	0,911722489	0,870550563
1	0,911722489	0,831237896	0,757858283
2	0,870550563	0,757858283	0,659753955
3	0,831237896	0,690956440	0,574349177
4	0,793700526	0,629960525	0,500000000
5	0,757858283	0,574349177	0,435275282
6	0,723634619	0,523647061	0,378929142
7	0,690956440	0,477420802	0,329876978
8	0,659753955	0,435275282	0,287174589
9	0,629960525	0,396850263	0,250000000
10	0,601512518	0,361817309	0,217637641
11	0,574349177	0,329876978	0,189464571
12	0,548412490	0,300756259	0,164938489
13	0,523647061	0,274206245	0,143587294

spxattencod	binindex		
	0	1	2
14	0,500000000	0,250000000	0,125000000
15	0,477420802	0,227930622	0,108818820
16	0,455861244	0,207809474	0,094732285
17	0,435275282	0,189464571	0,082469244
18	0,415618948	0,172739110	0,071793647
19	0,396850263	0,157490131	0,062500000
20	0,378929142	0,143587294	0,054409410
21	0,361817309	0,130911765	0,047366143
22	0,345478220	0,119355200	0,041234622
23	0,329876978	0,108818820	0,035896824
24	0,314980262	0,099212566	0,031250000
25	0,300756259	0,090454327	0,027204705
26	0,287174589	0,082469244	0,023683071
27	0,274206245	0,075189065	0,020617311
28	0,261823531	0,068551561	0,017948412
29	0,250000000	0,062500000	0,015625000
30	0,238710401	0,056982656	0,013602353
31	0,227930622	0,051952369	0,011841536

E.2.6.4.2.4 Noise scaling and transform coefficient blending calculation

In order to properly blend the translated transform coefficients with pseudo-random noise, the noise components for each band shall be scaled to match the energy of the translated transform coefficients in the band. The energy matching can be achieved by scaling all the noise components in a given band by the RMS energy of the translated transform coefficients in that band, provided the noise components are generated by a zero-mean, unity-variance noise generator. Once the zero-mean, unity-variance noise components for each band have been scaled by the RMS energy for that band, the scaled noise components can be blended with the translated transform coefficients.

The following pseudo code indicates how the translated transform coefficients and pseudo-random noise for a channel [ch] are blended. The function noise() returns a pseudo-random number generated from a zero-mean, unity-variance noise generator.

Pseudo Code

```
spxmant = spxbandtable[spx_begin_subbnd];
for (bnd = 0; bnd < nspxbnds; bnd++)
{
    bandsize = spxbndsztab[bnd];
    nscale = rmsenergy[ch][bnd] * nblendfact[ch][bnd];
    sscale = sblendfact[ch][bnd];
    for (bin = 0; bin < bandsize; bin++)
    {
        tctemp = tc[ch][spxmant];
        ntemp = noise();
        tc[ch][spxmant] = tctemp * sscale + ntemp * nscale;
        spxmant++;
    }
}
```

E.2.6.4.3 Blended transform coefficient scaling

The final step of the high frequency transform coefficient synthesis process is blended transform coefficient scaling. In this step, blended transform coefficients are scaled by the spectral extension coordinates to form the final synthesized high frequency transform coefficients. After this step, the banded energy of the synthesized high frequency transform coefficients should match the banded energy of the high frequency transform coefficients of the original signal.

The blended transform coefficient scaling process for channel [ch] is shown in the following pseudo code.

Pseudo Code

```
spxmant = spxbandtable[spx_begin_subbnd];

for (bnd = 0; bnd < nspxbnds; bnd++)
{
    bandsize = spxbndsztab[bnd];
    spxcotemp = spxco[ch][bnd];
    for (bin = 0; bin < bandsize; bin++)
    {
        tctemp = tc[ch][spxmant];
        tc[ch][spxmant] = tctemp * spxcotemp * 32;
        spxmant++;
    }
}
```

E.2.7 Transient pre-noise processing

E.2.7.0 Introduction

Transient pre-noise processing is an audio coding improvement technique that reduces the duration of pre-noise introduced by low-bit rate audio coding of transient material. A detailed description of transient pre-noise processing follows.

E.2.7.1 Overview

When transient pre-noise processing is used, decoded PCM audio located prior to transient material is used to overwrite the transient pre-noise, thereby improving the perceived quality of low-bit rate audio coded transient material. To enable the decoder to efficiently perform transient pre-noise processing with minimal decoding complexity, transient location detection and time scaling synthesis analysis is performed by the encoder and the information transmitted to the decoder. The encoder performs transient pre-noise processing for each full bandwidth audio channel and transmits the information once per audio frame. The transmitted transient location and time scaling synthesis information are relative to the first decoded PCM sample contained in the audio frame containing the bit stream information. It should be noted that it is possible for the time scaling synthesis parameters contained in audio frame N, to reference PCM samples and transients located in audio frame N+1, but this does not create a requirement for multi-frame decoding.

E.2.7.2 Application of transient pre-noise processing data

The bit stream syntax and high level description of the transient pre-noise parameters contained in the audio frame field are outlined in clauses E.1.2.3 and E.1.3.2, respectively. The parameter transproce indicates whether any of the full bandwidth channels in the current audio frame have associated transient pre-noise time scaling synthesis processing information. If transproce is set to a value of "1", then the parameter chintransproc[ch] can be set for each full bandwidth channel. For each full bandwidth channel where chintransproc[ch] is set to a value of "1", the transient location parameter transprocloc[ch] and time scaling length parameter transproclen[ch] are each set to values that have been calculated by the encoder.

Figure E.2.2 provides an overview of how the transient pre-noise parameters that are computed and transmitted by the encoder are applied in the decoder. As shown in Figure E.2.2 a), the parameter transprocloc[ch] identifies the location of the transient relative to the first sample of decoded PCM channel data in the audio frame that contains the transient pre-noise processing parameters. As defined, transprocloc[ch] has four sample resolution to reduce the data rate required to transmit the transient location and shall be multiplied by 4 to get the location of the transient in samples. As also shown in Figure E.2.2 a), the parameter transproclen[ch] provides the time scaling length, in samples, relative to the leading edge of the audio coding block prior to the block in which the transient is located. As shown in Figure E.2.2 b), the location of the leading edge of the audio coding block prior to the block containing the transient indicates the start of the transient pre-noise. The start of the previous audio coding block and location of the transient provide the total length of the transient pre-noise in samples, PN. As part of the normal decoding operation, the decoder inherently knows the starting location of the audio coding block that contains the transient and this does not need to be transmitted.

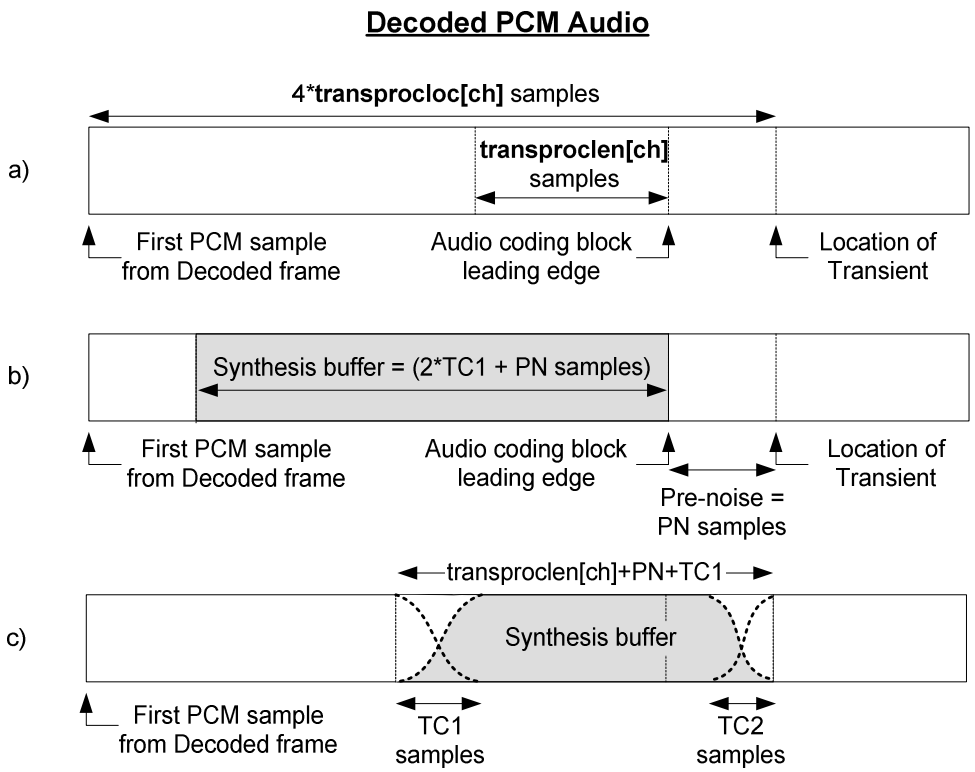

Figure E.2.2: Transient pre-noise time scaling synthesis summary

Also shown in Figure E.2.2 b) is how the time scaling synthesis audio buffer, which is used to modify the transient pre-noise, is defined relative to the decoded audio frame. The time scaling synthesis buffer is ($2 \times TC1 + PN$) PCM samples in length, where TC1 is a time scaling synthesis system parameter equal to 256 samples. The first sample of the time scaling synthesis buffer is located ($2 \times TC1 + 2 \times PN$) samples before the location of the transient.

Figure E.2.2 c) outlines how the time scaling synthesis buffer is used along with the transproclen[ch] parameter to remove the transient pre-noise. As shown in Figure E.2.2 c) the original decoded audio data is cross-faded with the time scaling synthesis buffer starting at the sample located (PN + TC1 + transproclen[ch]) samples before the location of the transient. The length of the cross-fade is TC1 or 256 samples. Nearly any pair of constant amplitude cross-fade windows may be used to perform the overlap-add between the original data and the synthesis buffer, although standard Hanning windows have been shown to provide good results. The time scaling synthesis buffer is then used to overwrite the decoded PCM audio data that is located before the transient, including the transient pre-noise. This overwriting continues until TC2 samples before the transient where TC2 is another time scaling synthesis system parameter equal to 128 samples. At TC2 samples before the transient, the time scaling synthesis audio buffer is cross-faded with the original decoded PCM data using a set of constant amplitude cross-fade windows.

The following pseudo code outlines how to implement the transient pre-noise time scaling synthesis functionality in the decoder for a single full bandwidth channel, [ch].

Where:

win_fade_out1 = TC1 sample length cross-fade out window (unity to zero in value);

win_fade_in1 = TC1 sample length cross-fade in window (zero to unity in value);

win_fade_out2 = TC2 sample length cross-fade out window (unity to zero in value);

win_fade_in2 = TC2 sample length cross-fade in window (zero to unity in value).

Pseudo Code

```
/* unpack the transient location relative to first decoded pcm sample. */
transloc = transprocloc[ch];
/* unpack time scaling length relative to first decoded pcm sample. */
translen = transproclen[ch];
/* compute the transient pre-noise length using audio coding block first sample,
aud_blk_samp_loc. */
pnlen = (transloc - aud_blk_samp_loc);
/* compute the total number of samples corrected in the output buffer. */
tot_corr_len = (pnlen + translen + TC1);
/* create time scaling synthesis buffer from decoded output pcm buffer, pcm_out[ ]. */
for (samp = 0;  samp < (2*TC1 + pnlen); samp++)
    synth_buf[samp] = pcm_out[(transloc - (2*tc + 2*pnlen) + samp)];
end
/* use time scaling synthesis buffer to overwrite and correct pre-noise in output pcm buffer. */
start_samp = (transloc - tot_corr_len);
for (samp = 0; samp < TC1; samp++)
{
    pcm_out[start_samp + samp] =    (pcm_out[start_samp + samp] * win_fade_out1[samp]) +
        (synth_buf[samp] * win_fade_in1[samp]);
}
for (samp = TC1; samp < (tot_corr_len - TC2); samp++)
{
    pcm_out[start_samp + samp] = synth_buf[samp];
}
for (samp = (tot_corr_len - TC2); samp < tot_corr_len; samp++)
{
    pcm_out[start_samp + samp] =    (pcm_out[start_samp + samp] * win_fade_in2[samp]) +
        (synth_buf[samp] * win_fade_out2[samp]);
}
```

E.2.8 Channel and programme extensions

E.2.8.0 Introduction

The Enhanced AC-3 bit stream syntax allows for time-multiplexed substreams to be present in a single bit stream. By allowing time-multiplexed substreams, the Enhanced AC-3 bit stream syntax enables a single programme with greater than 5.1 channels, multiple programmes of up to 5.1 channels, or a mixture of programmes with up to 5.1 channels and programmes with greater than 5.1 channels, to be carried in a single bit stream.

E.2.8.1 Overview

An Enhanced AC-3 bit stream shall consist of at least one independently decodable stream (type 0 or 2). Optionally, Enhanced AC-3 bit streams may consist of multiple independent substreams (type 0 or 2) or a combination of multiple independent (type 0 and 2) and multiple dependent (type 1) substreams.

All Enhanced AC-3 decoders shall be able to decode independent substream 0, and skip over any additional independent and dependent substreams present in the bit stream.

Optionally, Enhanced AC-3 decoders may use the information present in the acmod, lfeon, strmtyp, substreamid, chanmape, and chanmap bit stream parameters to decode bit streams with a single programme with greater than 5.1 channels, multiple programmes of up to 5.1 channels, or a mixture of programmes with up to 5.1 channels and programmes with greater than 5.1 channels.

E.2.8.2 Decoding a single programme with greater than 5.1 channels

When a bit stream contains a single programme with greater than 5.1 channels, independent substream 0 contains a 5.1 channel downmix of the programme for compatibility with playback systems containing 5.1 speakers. The audio in independent substream 0 can also be downmixed for compatibility with playback systems with fewer than 5.1 speakers. Decoders reproducing 5.1 or fewer channels from a programme containing greater than 5.1 channels shall decode only independent substream 0 and skip all associated dependent substreams.

In order to accommodate playback by systems with greater than 5.1 speakers, the Enhanced AC-3 bit stream will carry one or more dependent substreams that contain channels that either replace or supplement the 5.1-channel audio data carried in independent substream 0.

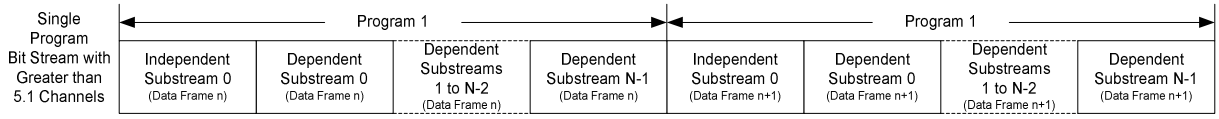

Figure E.2.3: Bit stream with a single programme of greater than 5.1 channels

If the chanmape parameter of a dependent substream is set to 0, then the acmod and lfeon parameters of the dependent substream are used to identify the channels present in the dependent substream, and the corresponding audio channels in the independent substream are overwritten with the dependent audio channel data. For example, if the dependent substream uses acmod 1/0 (centre channel only) and has lfeon set to 1, then the centre channel audio data carried in the dependent stream will replace the centre channel audio data carried in the independent stream, and the LFE audio data carried in the dependent stream will replace the LFE data carried in the independent stream.

If the chanmape parameter of a dependent substream is set to 1, then the chanmap parameter is used to determine the channel mapping for all channels contained in the dependent stream. Each bit of the chanmap parameter corresponds to a particular channel location or pair of channel locations. Audio data is contained in the dependent substream for each chanmap bit that is set to 1. The order of the coded channels in the dependent substream is the same as the order of the bits set to 1 in the chanmap parameter. For example, if the Left channel bit is set to 1 in the channel map field, then Left channel audio data will be contained in the first coded channel of data in the dependent substream. If channels are present in the dependent substream that correspond to channels in the associated independent substream, then the dependent substream data for those channels replaces the independent substream data for the corresponding channels. All channels present in the dependent substream that do not correspond to channels in the independent substream are used to enable output for speaker configurations with greater than 5.1 channels.

The maximum number of channels rendered for a single programme is 16.

E.2.8.3 Decoding multiple programmes with up to 5.1 channels

When an Enhanced AC-3 bit stream contains multiple independent substreams, each independent substream corresponds to an independent audio programme. The application interface may inform the decoder which independent audio programme should be decoded by selecting a specific independent substream ID. The decoder should then only decode substreams with the desired independent substream ID, and skip over any other programmes present in the bit stream with different substream ID's. The default programme selection should always be Programme 1.

In some cases, it may be desirable to decode multiple independent audio programmes. In these cases, the application interface should inform the decoder which independent audio programmes to decode by selecting specific independent substream ID's. The decoder should then decode all substreams with the desired independent substream ID's, and skip over any other programmes present in the bit stream with different substream ID's.

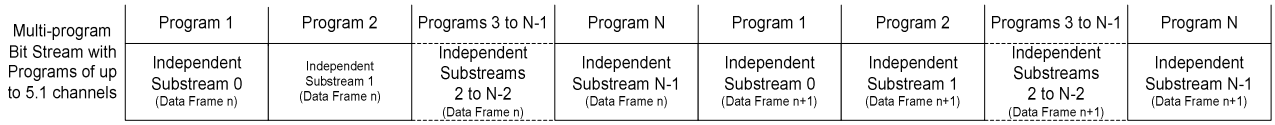

Figure E.2.4: Bit stream with multiple programmes of up to 5.1 channels

E.2.8.4 Decoding a mixture of programmes with up to 5.1 channels and programmes with greater than 5.1 channels

When an Enhanced AC-3 bit stream contains multiple independent and dependent substreams, each independent substream and its associated dependent substreams correspond to an independent audio programme. The application interface may inform the decoder which independent audio programme should be decoded by selecting a specific independent substream ID. The decoder should then only decode the desired independent substream and all its associated dependent substreams, and skip over all other independent substreams and their associated dependent substreams. If the selected independent audio programme contains greater than 5.1 channels, the decoder should decode the selected independent audio programme as explained in clause E.2.8.2. The default programme selection should always be Programme 1.

In some cases, it may be desirable to decode multiple independent audio programmes. In these cases, the application interface should inform the decoder which independent audio programmes to decode by selecting specific independent substream ID's. The decoder should then decode the desired independent substreams and their associated dependent substreams, and skip over all other independent substreams and associated dependent substreams present in the bit stream.

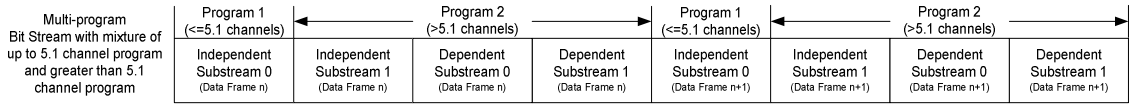

Figure E.2.5: Bit stream with mixture of programmes of up to 5.1 channels and programmes of greater than 5.1 channels

E.2.8.5 Dynamic range compression for programmes containing greater than 5.1 channels

A programme using channel extensions to convey greater than 5.1 channels may require two different sets of compr and dynrng metadata words: one set for the 5.1-channel downmix carried by independent substream 0 and a separate set for the complete (greater than 5.1 channel) mix. If a decoder is reproducing the complete mix, the compr and dynrng metadata words carried in independent substream 0 shall be ignored. The decoder shall instead use the compr and dynrng metadata words carried by the associated dependent substream. If multiple associated dependent substreams are present, only the last dependent substream may carry compr and dynrng metadata words, and these metadata words shall apply to all substreams in the programme, including the independent substream.

The compre bit is used by the decoder to determine which dependent substream in a programme is the last dependent substream of the programme. Therefore, the compre bit in the last dependent substream of a programme shall be set to 1, and the compre bit in all other dependent substreams of the programme shall be set to 0. Additionally, the compr2e, dynrnge, and dynrng2e bits for all but the last dependent substream of a programme shall be set to 0. The compr2e, dynrnge, and dynrng2e bits for the last dependent substream shall be set as required to transmit the proper compr2, dynrng, and dynrng2 words for the programme.

Note that the compr2e, compr2, dynrng2e, and dynrng2 metadata words are only present in the bit stream when acmod = 0.

E.2.9 LFE downmixing decoder description

For decoders with only 2-channel or mono outputs, where a dedicated LFE/Subwoofer output is not available, Enhanced AC-3 enables the LFE channel audio to be mixed into the Left and Right channels at a level indicated by the lfemixcodlev bit stream parameter.

LFE downmixing shall occur only if the lfemixlevcod parameter is present in the bit stream and the decoder is operating in 1/0 (C only) or 2/0 (L/R) output modes with the LFE channel output disabled. For all other output modes, the lfemixlevcod parameter, if present, shall be ignored. Note that lfemixlevcode shall be assumed to be 0 when it is not transmitted in the bit stream. For the 1/0 case, the decoder should perform a standard 2/0 downmix with the LFE mixed into the Left and Right channels, followed by a subsequent mix of the L/R channels to a mono C channel. The following pseudo code indicates how the decoder should perform the LFE downmix.

Pseudo Code

```
if (output mode == 1/0 or 2/0) && (lfeoutput == disabled) && (lfemixlevcode == 1))
{
    mix LFE into left with (LFE mix level - 4,5) dB gain
    mix LFE into right with (LFE mix level - 4,5) dB gain
}
if (output mode == 1/0)
{
    mix left into centre with -6 dB gain
    mix right into centre with -6 dB gain
}
```

E.2.10 Control of Programme Mixing

E.2.10.1 Overview

The Enhanced AC-3 bitstream syntax includes parameters that can be used to control the mixing of two audio programs after simultaneous decoding by a device containing an Enhanced AC-3 decoder. Typically these two programs are (1) a main audio component, which contains the majority of the audio and is sufficiently complete that it can be decoded on its own to deliver a full audio presentation to the listener, and (2) an associated audio component, which contains supplementary audio content (for example a commentary or video description track) that is intended to be combined with the main audio service before presentation to the listener.

These services should be delivered using one of the two following methods:

1) As two separate Enhanced AC-3 streams (with one programme carried in independent substream 0 of the first Enhanced AC-3 stream and the second carried in independent substream 0 of the second stream).

2) As a single Enhanced AC-3 stream with two (or more) independent substreams.

If case number 2 is used, then the main audio component shall be carried in independent substream 0 and dependent substreams associated with independent substream 0, if any; and the associated audio component shall be carried in an independent substream with a non-zero substreamid value.

The mixing metadata parameters are carried within the Bit Stream Information (BSI) field of each Enhanced AC-3 syncframe, and are defined in clause E.1.2.2 and clause E.1.3.1. The following clauses provide information on the intended usage of each mixing metadata parameter when control of a mixing process is required.

E.2.10.2 pgmscl

The pgmscl (programme scale factor) parameter is defined in clause E.1.3.1.13. This parameter specifies a gain value used to adjust the level of audio service that is being carried in the same substream as the pgmscl parameter. For example, if the pgmscl parameter is present in independent substream 0 of an Enhanced AC-3 stream that is carrying a main audio service, and the pgmscl parameter specifies a gain of -3 dB, all audio channels of the main audio service carried in independent substream 0 will be attenuated by 3 dB during the mixing process.

E.2.10.3 extpgmscl

The extpgmscl (external programme scale factor) parameter is defined in clause E.1.3.1.17. This parameter specifies a gain value used to adjust the level of an audio service that is being carried in a different Enhanced AC-3 bitstream or substream from the bitstream or substream that contains the extpgmscl parameter. For example, if independent substream 1 of an Enhanced AC-3 stream that is carrying an associated audio service contains extpgmscl data that specifies a gain value of -10 dB, and independent substream 0 of the same Enhanced AC-3 stream contains the main audio service, all audio channels of the main audio service carried in independent substream 0 will be attenuated by 10 dB during the mixing process.

E.2.10.4 mixdef

The mixdef (mix control field length) parameter is defined in clause E.1.3.1.18. This parameter defines the length of the mixdata field, which is a variable length container used to store a range of mixing metadata parameters that supplement the pgmscl, extpgmscl and panmean parameters, providing additional control of the mixing process.

When the mixdef parameter is set to '00', the mixdata field is not present in the syncframe, and only the pgmscl, extpgmscl and panmean parameters may be present in the Enhanced AC-3 syncframe.

When the mixdef parameter is set to '01', the mixdata field is 5 bits long, and contains the premixcmpsel, drcsrc and premixcmpscl parameters. These parameters were originally defined to enable dynamic range compression to be applied to the main audio service as part of the mixing process, but this functionality is not supported by the Enhanced AC-3 mixing model, so these parameters should be set to the values that are recommended in clause E.1.3.1 by the encoder.

When the mixdef parameter is set to '11', the mixdata field can be between 2 and 33 bytes long, and the actual length of the mixdata field is defined by the mixdeflen parameter.

E.2.10.5 mixdeflen

When the mixdef parameter is set to '11', the mixdeflen parameter specifies the length of the mixdata field in bytes. The range of the mixdeflen parameter is 0 to 31, which specifies a mixdata field length of between 2 and 33 bytes in one byte increments. In this case, the mixdata field is required, at a minimum, to contain the mixdeflen, mixdata2e and mixdata3e parameters, and if the mixdata2e and mixdata3e flags are set to '0', the remaining bits in mixdata are required to be set to '0'.

E.2.10.6 mixdata2e

The mixdata2e flag is set to '1' when an additional set of mixing metadata parameters is included in the syncframe. These parameters enable control over individual channels of an external programme.

E.2.10.7 extpgm(X)scl

Up to eight individual channel scaling parameters (extpgmXscl, where X is l, c, r, ls, rs, lfe, aux1 or aux2) are available to adjust the level of each individual channel in an external programme containing up to 7.1 channels. These parameters are defined in clauses:

- E.1.3.1.26
- E.1.3.1.28
- E.1.3.1.30
- E.1.3.1.32
- E.1.3.1.34
- E.1.3.1.36
- E.1.3.1.41
- E.1.3.1.43

The parameters are named to match the corresponding channels - e.g. extpgmlscl adjusts the gain of the left channel of the programme, and extpgmrsscl adjusts the level of the right surround channel of the external programme.

The extpgmaux1scl and extpgmaux2scl parameters are used to adjust the level of channels with channel locations that can be specified only by using the chanmap parameter (e.g. the Vhc channel). The use of "auxiliary" rather than assigning fixed channel location labels is because Enhanced AC-3 can assign a number of different channel locations to these coded channels through use of the chanmap parameter. Up to two of these auxiliary channels may be present in a programme.

The gain indicated by each of the individual channel scaling parameters is combined with the gain indicated by the extpgmscl parameter (which applies to all channels) to specify the total gain that is to be applied to that channel of the external programme. For example, if independent substream 1 of an Enhanced AC-3 stream that is carrying an associated audio service contains extpgmscl data that specifies a gain value of -10 dB, and also contains extpgmlsscl data that specifies a gain value of -10 dB, and independent substream 0 of the same Enhanced AC-3 stream contains the main audio service, all audio channels of the main audio service carried in independent substream 0 will be attenuated by 10 dB during the mixing process (as specified by the value of extpgmscl), and the left surround channel of the main audio service will be attenuated by a further 10 dB (as specified by the value of extpgmlsscl).

E.2.10.8 dmixscl

When a multichannel audio programme is downmixed to 2 channels within the Enhanced AC-3 decoder, it is no longer possible to apply the individual channel scaling parameters to each individual channel of the multichannel audio programme in the mixer as these channels have been combined during the downmixing process. In this situation it may still be desirable to apply additional attenuation to the downmixed audio that is output by the Enhanced AC-3 decoder, and the dmixscl parameter is used for this purpose. Similarly to the individual channel scaling parameters, the gain indicated by the dmixscl parameter is combined with the gain indicated by the extpgmscl parameter to specify the total gain that is to be applied to the downmixed multichannel audio programme. The dmixscl parameter should only be used when a multichannel audio programme has been downmixed to 2 channels within the Enhanced AC-3 decoder, preventing the use of individual channel scaling parameters. If the individual channels of the multichannel audio programme are available to the mixer, then the individual channel scaling parameters are used, and the dmixscl parameter should be ignored.

E.2.10.9 panmean

The panmean parameter allows a mono associated audio stream to be panned to any of the channels of the main audio service. When the value of the panmean parameter is 0, this indicates the panned virtual source points toward the center speaker location (defined as 0 degrees). The index indicates 1.5 degree increments in a clockwise rotation. Values 0 to 239 represent 0 to 358,5 degrees, while values 240 to 255 are reserved.

The proportion of associated audio in each channel of the main audio is dependent on the output configuration of the main audio decoder, and the number of channels in the main audio service.

For mixing a mono associated audio service with a stereo (or downmixed) main audio service, the associated audio service is split into two channels, A_L and A_R, to be mixed with the Left and Right channels of the main audio respectively. Table E.2.13 shows, for each value of panmean, the scale factors to be applied to A_L and A_R prior to mixing with the corresponding main audio channel.

Table E.2.13: Associated Audio scale factors for stereo output panning

panmean range	A_L scale factor	A_R scale factor
0 - 19	$= \cos\left[\frac{\pi}{2}\frac{(panmean + 20)}{40}\right]$	$= \sin\left[\frac{\pi}{2}\frac{(panmean + 20)}{40}\right]$
20 - 99	0	1
100 - 139	$= \sin\left[\frac{\pi}{2}\frac{(panmean - 100)}{40}\right]$	$= \cos\left[\frac{\pi}{2}\frac{(panmean - 100)}{40}\right]$
140 - 219	1	0
220 - 239	$= \cos\left[\frac{\pi}{2}\frac{(panmean - 220)}{40}\right]$	$= \sin\left[\frac{\pi}{2}\frac{(panmean - 220)}{40}\right]$

For mixing a mono associated audio service with a 5.1-channel main audio service, the associated audio service is split into five channels (the LFE channel is not included), A_L, A_C, A_R, A_{LS} and A_{RS}, to be mixed with the Left, Center, Right, Left Surround and Right Surround channels of the main audio respectively. Table E.2.14 shows the scale factors to be applied to A_L, A_C and A_R prior to mixing with the corresponding main audio channel for each value of panmean. Table E.2.15 shows the scale factor to be applied to A_{LS} and A_{RS}, prior to mixing with the corresponding main audio channel for each value of panmean.

Table E.2.14: Associated Audio scale factors for 5.1-channel output panning: L, C, and R channels

panmean value	A_L scale factor	A_C scale factor	A_R scale factor
0-19	0	$=\cos\left[\frac{\pi}{2}\frac{panmean}{20}\right]$	$=\sin\left[\frac{\pi}{2}\frac{panmean}{20}\right]$
20-72	0	0	$=\cos\left[\frac{\pi}{2}\frac{(panmean-20)}{53}\right]$
73-166	0	0	0
167-219	$=\sin\left[\frac{\pi}{2}\frac{(panmean-167)}{53}\right]$	0	0
220-239	$=\cos\left[\frac{\pi}{2}\frac{(panmean-220)}{20}\right]$	$=\sin\left[\frac{\pi}{2}\frac{(panmean-220)}{20}\right]$	0

Table E.2.15: Associated Audio scale factors for 5.1-channel output panning: Ls and Rs channels

panmean value	A_{LS} scale factor	A_{RS} scale factor
0-19	0	0
20-72	0	$=\sin\left[\frac{\pi}{2}\frac{(panmean-20)}{53}\right]$
73-166	$=\sin\left[\frac{\pi}{2}\frac{(panmean-73)}{94}\right]$	$=\cos\left[\frac{\pi}{2}\frac{(panmean-73)}{94}\right]$
167-219	$=\cos\left[\frac{\pi}{2}\frac{panmean-167}{53}\right]$	0
220-239	0	0

E.3 AHT vector quantization tables

Table E.3.1: VQ table for hebap 1 (16-bit two's complement)

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
0	0x1bff	0x1283	0x0452	0x10ad	0x28ac	0x12d4
1	0xe9ba	0xf38d	0xc76d	0xfa90	0xf815	0x0351
2	0x0279	0x1837	0x1b61	0xce15	0xf6fe	0xf5b4
3	0xfa44	0xe489	0x1da8	0x2979	0xe8c6	0xf40a

Table E.3.2: VQ table for hebap 2 (16-bit two's complement)

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
0	0xd0d7	0x0260	0xe495	0x024e	0x0fa0	0x0365
1	0x1a24	0x3d49	0xe7de	0xdbe9	0xffb6	0x0085
2	0x073f	0xfc23	0x5074	0xf498	0xee85	0x00e1
3	0xfb56	0xf0c3	0xfccb	0xe65a	0xfc95	0xb0b6
4	0xf536	0xf393	0xf002	0xea09	0xbdcf	0x2625
5	0x060b	0x1ab7	0x07bc	0x4f09	0xfbd1	0xec86
6	0x184d	0xba05	0xea74	0x187a	0x0166	0x048a
7	0x0ea9	0xfbd6	0x10bb	0xf365	0x3e38	0x27ca

Table E.3.3: VQ table for hebap 3 (16-bit two's complement)

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
0	0xd8d4	0x512b	0x2ae6	0xee30	0x031e	0xffbc
1	0x2b2a	0x500a	0xe627	0xeb22	0xf8fb	0xf9a1
2	0x0f89	0xfde2	0x1bce	0xfb72	0x499c	0x3956
3	0xef20	0xffa0	0xe381	0xfe14	0xa9de	0xef4b
4	0x0a84	0x16e0	0x159a	0x5566	0xe3d4	0xeb33
5	0xff79	0xa4a1	0x03c2	0x1fb3	0xfd7c	0x017e
6	0xf9e5	0x0d48	0xf31d	0x1255	0xe514	0x577e
7	0x0dcf	0x0bd6	0x1c80	0x1846	0x4ffc	0xd0bd
8	0x0039	0xe559	0x0738	0xa8b3	0xe8e1	0x1aa7
9	0xfccb	0xf1b9	0xfe7d	0xe793	0xf939	0xa89b
10	0xe862	0x0632	0xb636	0xc7c8	0x23fe	0x02c1
11	0xe9ac	0x0108	0xb9d4	0x391a	0x1ef1	0xfeaf
12	0xff92	0x006c	0x0008	0x004a	0xffa7	0xffce
13	0x19d4	0xfa13	0x54b7	0xf986	0xe0f3	0xff0a
14	0x54a3	0xe741	0xdf9e	0xff9b	0xfabb	0xffea
15	0xaa0d	0xe6b4	0x1f26	0x0288	0x0806	0xfeb5

Table E.3.4: VQ table for hebap 4 (16-bit two's complement)

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
0	0x5903	0x15c0	0xe9e6	0xff64	0xfe06	0xffdf
1	0x19ec	0xee0f	0x375d	0xbc6f	0xbf75	0x0360
2	0x0e4a	0x580c	0x0068	0xf91d	0xffac	0x0006
3	0x544c	0xba69	0xe38e	0xf9d9	0xf7e2	0xfec0
4	0xf747	0x2721	0xf558	0x3a5a	0xcab8	0xbb05
5	0xf9e4	0xbab6	0xb527	0x35a7	0x0ac5	0x0b87
6	0x11a8	0x1586	0x1ce1	0x2a2f	0x4b1f	0xca36
7	0x018f	0x0ba0	0xfbb5	0x1395	0xfb79	0x564f
8	0x0e28	0xf6c9	0x1248	0xf742	0x58ae	0x0eb5
9	0xef97	0xdfa3	0xe566	0xcf9a	0xb812	0x3c16
10	0xebb2	0xe52b	0xd8c1	0xdf54	0xc16a	0xafae
11	0xff72	0xa771	0xfe90	0x1127	0xfe30	0xfff3
12	0x032e	0xfba2	0xfbbf	0xa9fd	0x004a	0x0611
13	0xf9ae	0x4b16	0xbb16	0xcb4e	0x034a	0xf6fb
14	0x1251	0x406a	0x514d	0xc3e5	0xefbc	0xf080
15	0xf314	0x2bce	0xcb1a	0x351f	0xb3ef	0x35ca
16	0x0719	0x0356	0x52e9	0xfc3a	0xf995	0xfef4
17	0xf5e5	0xff95	0xb146	0x0178	0x0496	0xfed0
18	0xf499	0x01c5	0xeaf2	0x02ee	0xa9ee	0xfc2e
19	0xb5bc	0x41c7	0x2710	0xf204	0x08a3	0x05b3
20	0x0553	0xf59e	0xffdf	0xf01d	0x048d	0xaa1f
21	0xde70	0xf538	0xbb90	0xc18f	0x3a31	0x052b
22	0x028c	0xdb8d	0x0cb5	0xc6e2	0x2f95	0x4cec
23	0xe727	0x168d	0xc3dd	0x438b	0x40ce	0xf496
24	0xfd6b	0xfda7	0x0649	0x5852	0x03e0	0xfbeb
25	0x1361	0x2393	0x2bd9	0x1e95	0x3fc0	0x48c3
26	0xaa90	0xfa67	0x008a	0x05be	0xf89d	0xff3c
27	0xb3d5	0xb8e5	0x2b30	0xfdfc	0x09ef	0xf737
28	0xfb54	0xbb5a	0x4eb6	0x2cc6	0xfe6f	0x0a3b
29	0x121e	0xe026	0x2e73	0xc271	0x44cf	0xc595
30	0xffad	0x0116	0x0143	0x0037	0xff66	0x00e8
31	0x1e6c	0x05b6	0x47db	0x3bc0	0xc26d	0xfb95

Table E.3.5: VQ table for hebap 5 (16-bit two's complement)

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
0	0xf2be	0xb2ee	0x0b93	0x2576	0x1234	0x4cd9
1	0xc2cf	0xe6fb	0x4507	0x0f14	0xdfd8	0xb41b
2	0x1173	0x019c	0xba2f	0xe09b	0x02b3	0xbc65
3	0x0dfc	0x093b	0x1ae6	0x0eb3	0x18eb	0xafd6
4	0xbcb2	0xc8cb	0xfa8c	0xa27d	0x20b5	0xcf07
5	0xe077	0xac23	0xc2ea	0x0c8e	0x1fa9	0xe8b3
6	0x10f7	0x1431	0x0a7b	0xbe4a	0xebe6	0xbf52
7	0x18cc	0xd654	0x32c3	0x9c64	0xa9b6	0x0ffb
8	0xf4c0	0xdf52	0xe8b0	0xbcfa	0xf5b2	0x5a5c
9	0xec19	0xc837	0xa89d	0x54ed	0x0e69	0x0b91
10	0xf675	0xbab5	0x6243	0x0a93	0x063a	0x0007
11	0xb835	0x2332	0x10ae	0x02db	0xfe56	0xfd80
12	0xa371	0x609c	0x160a	0x0264	0xfecc	0xfc3c
13	0xfd01	0x04f4	0x00e1	0x0663	0x00ad	0x0394
14	0x154f	0x195d	0x1326	0x2940	0x5a01	0xbd0c
15	0x4343	0xadc2	0xb6e4	0x1348	0xf2a4	0x0d1d
16	0xfa92	0x3c80	0xaa46	0xc6ed	0x053b	0x021e
17	0xe52e	0xf732	0xd0da	0xf3fd	0xb1f3	0xaf72
18	0xf8f5	0x2dff	0x053f	0x22d5	0x02b5	0x5fb1
19	0xab96	0x24f6	0x1249	0x2426	0xe179	0x3e20
20	0xea49	0xf436	0xdc2f	0xfabd	0xa7ed	0x3244
21	0xfe92	0x13d4	0xf941	0x4fcb	0xfee5	0xf495
22	0xf8a2	0xe757	0xfc55	0xf7df	0xfa89	0x0db9
23	0xf3a7	0xfde7	0xec2d	0x2c04	0x4bc4	0x03dd
24	0x0929	0x1039	0x1689	0xef4f	0x00e9	0xfe71
25	0xaa7a	0xfb8e	0xbfa6	0x170e	0x1570	0xf375
26	0x2717	0xcf0e	0x498d	0x51c4	0xfb7a	0x06fe
27	0xfb73	0x1396	0xfb51	0x190f	0xdf1e	0xadd2
28	0x0764	0xf232	0x0ee7	0xe92a	0x402b	0x4f40
29	0xf598	0xd295	0xefcd	0xb879	0xa74a	0x3a00
30	0x4368	0x28b3	0x1e54	0x2f08	0x4a0c	0x09dd
31	0xac55	0xb703	0xd56f	0x1110	0xe475	0x11bb
32	0xf9da	0x0802	0x1680	0x60b4	0x3e6f	0x450e
33	0xfde6	0xa6ad	0x2b3b	0x283d	0x0181	0x0210
34	0xdeef	0xf42f	0xc01b	0xa53b	0x406b	0x0e46
35	0x16d0	0x023f	0x2e72	0x079b	0x6245	0x19fd
36	0x19e1	0xf244	0xf854	0x0f0a	0xfe7a	0xff8c
37	0x4655	0x51a4	0x37f3	0xe23b	0xd556	0x2e1a
38	0xed07	0xf48c	0xcbea	0xe179	0x5476	0x08db
39	0xfdbd	0xdb29	0xfd14	0xacb7	0x304f	0x2049
40	0xdf83	0x055f	0xba49	0x0b69	0x2366	0x561e
41	0x47de	0x21bb	0xfa21	0xf68e	0xb889	0xc672
42	0xf455	0x3b19	0xf2fd	0x571c	0x3636	0xcab9
43	0x16f2	0xb5ae	0x3ce3	0x2c56	0xaefe	0x07b3
44	0x062d	0xe4d5	0xac40	0x0997	0x0041	0x019e
45	0x0203	0xee8c	0xfd67	0xedc0	0x007d	0xb4ea
46	0x53f7	0xb0b3	0xf8b0	0xf87a	0xff2d	0xfc02
47	0x1445	0xd026	0xf911	0xa402	0xee3e	0x16b5
48	0x0141	0xe745	0x3936	0x1b3f	0xf913	0x0363
49	0xca0a	0x0c6c	0x1ef7	0x01bc	0x4c60	0x0c4a
50	0xe5fc	0x2fdc	0xf84c	0x4400	0xa128	0xcd64
51	0xfd17	0x3814	0xfbad	0x5cbe	0xda61	0xb858
52	0x476c	0xe11b	0xe295	0x4aae	0x1e29	0xce8d
53	0x0786	0x3afd	0xcdd0	0x0869	0x547f	0x0748
54	0xf7ae	0x5b78	0x42a0	0xc313	0xf9f8	0x0057
55	0x207a	0xd1d0	0x38f5	0xaf91	0x1ed3	0xf7cd
56	0x4c90	0x591e	0xbc68	0xf808	0x011d	0xf0e9
57	0xdfea	0xb86e	0x29e4	0xca50	0xcb63	0xf97e
58	0x380f	0x1310	0xb1be	0x03c4	0xef83	0xff4c
59	0x9fea	0xbf05	0x4d0c	0x1725	0x12a9	0x113e
60	0xf641	0x5bc5	0xc0f3	0x0b66	0xfbf2	0xf826

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
61	0x4a1e	0xf614	0x341f	0x057d	0xe7ce	0xfb90
62	0x09b9	0x3566	0x586e	0xe371	0xff7f	0xf518
63	0xc976	0x4187	0x59e4	0x02d8	0x0d45	0x00a2
64	0x0bde	0x03e1	0x2246	0xaa2f	0xe62f	0x038e
65	0xcf64	0xa88e	0xf5be	0xeb51	0x4c40	0x2690
66	0xf889	0xb89e	0xb7b6	0xc58e	0x1298	0x1bcf
67	0x206a	0xf45e	0x651e	0x1dec	0xe127	0x03fc
68	0x17f4	0x3b17	0x4945	0xa0ce	0xe67f	0xe61d
69	0x1ef4	0x2f5d	0xdb0d	0xa266	0x157e	0x03a9
70	0xbd60	0xeb03	0x09da	0x0147	0x0469	0xfe7a
71	0x3d9e	0x4df3	0xd774	0x2ba4	0xf3dd	0x3a05
72	0xd180	0xe065	0xb9f8	0xa8f1	0xbcab	0xe56d
73	0xcdc2	0xf784	0xe693	0x1727	0x30aa	0xa8ad
74	0xfe0f	0x0142	0x040e	0xe60d	0xeae4	0x4f57
75	0x043b	0xa638	0xded2	0x2f62	0xfd06	0x0a3f
76	0x13cb	0x4d00	0xf893	0xffe2	0xfebb	0x0055
77	0x03db	0xe93a	0x1074	0xdcba	0x23a1	0x9e32
78	0xe144	0x1c74	0xcffc	0x3272	0xaba8	0x51cd
79	0xf9a2	0xe1f2	0xf775	0xdeb3	0xea1c	0x9d94
80	0xe5f4	0x0184	0xa7f9	0x05f6	0x237a	0x00c1
81	0xe145	0xa8dc	0x142b	0x016a	0x03b0	0xfefd
82	0x0ef0	0xd1b6	0x1da7	0xa578	0x62fe	0x5cdb
83	0xd6f8	0x101b	0xad89	0x52b5	0x57a7	0xfcba
84	0xed8d	0x5523	0x1828	0xff86	0x066a	0xfd33
85	0x5fb8	0x4daf	0xf805	0x03da	0x0007	0xffc9
86	0x954f	0xff79	0x0985	0x0103	0x0059	0x0133
87	0x5f7e	0xf0df	0xeaf1	0xfccc	0xf6ad	0x0169
88	0x1599	0x1698	0x48fa	0x00f2	0xaa78	0xf05d
89	0x5720	0x1183	0x02d2	0xd02e	0x1d92	0x3c58
90	0x21e1	0x0bc1	0x4fd5	0x5274	0xad94	0xf3f8
91	0xfb94	0x0a91	0xf8df	0x152c	0xfcef	0x4864
92	0x4224	0xcb33	0xbf83	0xc5f6	0xb099	0xc873
93	0x08ab	0x0564	0x53e2	0xfb98	0x0147	0x0053
94	0xf77f	0x540d	0xf0f0	0xc89c	0xff34	0xf771
95	0x03b9	0xdb2e	0x3e02	0xd62a	0xf361	0x5226
96	0xfe5b	0xfa9f	0x0280	0xdfd1	0xae10	0x087e
97	0x10d5	0x4852	0xdc74	0xb871	0xc362	0x0e78
98	0xe8c9	0x01c1	0xdf3d	0x0433	0xa93e	0xec80
99	0x0b89	0x31f4	0x476d	0x0596	0x3a59	0x54e3
100	0xf49f	0x0191	0xed7d	0xb177	0x06a3	0xfb85
101	0x0d79	0x1479	0x2295	0x5676	0xe285	0x05ab
102	0xf796	0x2188	0x46c8	0xc302	0x4b77	0xe899
103	0x0dad	0x0b19	0x1709	0x18fd	0x21b6	0x59ea
104	0x09a3	0x0b8c	0x017b	0x1647	0xa9e1	0xf773
105	0xbdbd	0xfdae	0x4986	0xeb51	0x0668	0x0306
106	0x0b50	0xfa70	0x0e02	0xf70c	0x4dc6	0xf8e2
107	0xb771	0x52e3	0xc94f	0xcee3	0x4052	0x027b
108	0xf832	0xb48e	0xbf71	0x2fb0	0xbf40	0xe152
109	0xda36	0x03f4	0xab73	0x0b43	0xce58	0x0911
110	0xfc13	0x01d7	0xf1d3	0x1f6d	0xd4b1	0x63bd
111	0x102d	0xac20	0xf58f	0x02f4	0xfd69	0xfdf5
112	0x195a	0x2153	0x4b59	0x4a05	0x17cc	0xdb7d
113	0x4245	0x6017	0x36c8	0x2758	0xfde8	0xd73a
114	0xe02d	0x0861	0xa60c	0xbd4f	0x154b	0xeecf
115	0xc5e7	0x5028	0xb881	0xda0b	0xd193	0xba59
116	0xf70e	0xc8d8	0x0816	0x57c3	0x0687	0x02d5
117	0xdea6	0x3925	0x0dc1	0xaf9f	0x1a11	0x2008
118	0x4f18	0x113a	0xfaaa	0xfdb7	0x04cd	0xf66f
119	0x1d2b	0xe414	0x3563	0xdfca	0x5778	0xbc58
120	0xf874	0x0f23	0xdc98	0xf11c	0x03be	0x0109
121	0xeed1	0x0b8f	0xc1d9	0x4c8e	0x135a	0xfbaf
122	0x4659	0xd93d	0xb927	0xf0ea	0x2baa	0x16bd

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
123	0xc6fc	0xfb35	0x25bc	0x5473	0x2bdc	0xd237
124	0xfd2f	0xf95c	0x006d	0xf7a2	0x003d	0xe58c
125	0x9fd5	0xa808	0x15e8	0xf85b	0xf91f	0xfc0c
126	0xa350	0xee9d	0xf580	0xc6a9	0xef56	0x26bf
127	0x212f	0xfc82	0x4fd6	0xca04	0xbc8d	0x008b

Table E.3.6: VQ table for hebap 6 (16-bit two's complement)

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
0	0x27aa	0x1cc5	0x41dd	0x48f9	0xa693	0xf1cc
1	0xf5c5	0xf134	0xea67	0xebb8	0xdccf	0xb0b6
2	0xea31	0xa6f0	0x5331	0x1b64	0x02e9	0x02d0
3	0x01ac	0xfa4d	0x006d	0xf3f6	0x0169	0xdf2d
4	0x1fe1	0x5781	0x00f1	0x06db	0xfc96	0xf4f8
5	0x0474	0x3163	0x0902	0x56f7	0x9dc6	0xbb6b
6	0xf5cf	0x0d33	0x2861	0xb2ee	0xc394	0xa278
7	0xf038	0xce04	0x9b54	0x3328	0x0f41	0x0523
8	0x1210	0xa309	0x3528	0x62e5	0xfb69	0x087d
9	0xff9f	0x35b3	0xebfe	0x5ad7	0x1076	0xa97f
10	0x1ade	0xfebe	0x4758	0xfcaa	0xd174	0xfd23
11	0x4380	0xce83	0xda23	0x5c0b	0xc090	0xfae3
12	0x16aa	0xec98	0x4c46	0xad36	0x9fd2	0xff49
13	0x16db	0xc0f7	0x3b7d	0xdae8	0xf9fe	0x0179
14	0x3710	0x61e1	0x346b	0x2062	0x5b18	0x41b2
15	0xe3a3	0x0076	0xc205	0x4a99	0x2635	0xfeeb
16	0xef40	0x5455	0xcc18	0xc07d	0x40f9	0xed02
17	0x132d	0xb4ef	0x5b73	0x3971	0xfd2e	0x007d
18	0x4c06	0xed84	0xf878	0xd2f9	0x5122	0x1531
19	0x9456	0xf4bf	0xef15	0x0180	0xf7c9	0x0557
20	0xfef6	0xdc29	0x1541	0x66dd	0xf87c	0x107d
21	0xf466	0xb136	0xaac8	0x154a	0xe2fe	0x14e0
22	0xff23	0xe5d8	0x025b	0xdc4c	0x051c	0x948e
23	0x2595	0xdf44	0xf851	0x24bb	0xf98d	0x5921
24	0x1d8e	0xeb7e	0xefbb	0x0569	0xfc22	0x0230
25	0xfb12	0x60a2	0xb58f	0x29f5	0x1da1	0xe446
26	0x01c3	0x4ea2	0xd923	0xe881	0xf774	0xfa4e
27	0x56e9	0x24a4	0x2388	0x2acf	0xf6c3	0xf174
28	0x48ec	0xfd76	0xfb2e	0x2b54	0x1dfe	0x1751
29	0x4b07	0xfa33	0xfbcc	0xfd25	0xfd54	0x002b
30	0xec93	0x3476	0x4eab	0x003c	0x01dc	0xfc59
31	0xb1c3	0x2206	0x09c3	0x03f8	0xfb7a	0x014f
32	0x98d3	0x48a6	0xf767	0xfd63	0x0d51	0x0319
33	0xed8a	0x22ab	0x9fe1	0xda52	0x0e3b	0xfee5
34	0x33f7	0xac64	0xf195	0xfb60	0xf84e	0x064c
35	0x00ad	0x003c	0x0397	0x04cd	0x1b1e	0xfd67
36	0x3ff9	0x425f	0x14dd	0xc941	0xf700	0xb05a
37	0x62f6	0xd68f	0x2eab	0xe21b	0xe725	0x36ea
38	0x5d79	0xcc35	0xe3c6	0xfc57	0x00ea	0xff45
39	0x18a7	0xf8ab	0x30da	0xf8a9	0x493f	0xa4d3
40	0x026d	0x192d	0x0d1a	0xa12e	0x20d6	0x14c3
41	0xf31f	0xec56	0xeda0	0xec28	0x9b7e	0x14e3
42	0xfb05	0xcc11	0xfc3b	0xa4ea	0x04be	0x6693
43	0xe794	0x2733	0xb177	0x3bc5	0xc137	0x1410
44	0x255a	0xf0b9	0xb3ca	0x1289	0x56fe	0xefb5
45	0x1f2a	0xb370	0x36c8	0xe98f	0xae89	0x22eb
46	0x0007	0xf039	0x03df	0xe84f	0x0034	0xb421
47	0x0d9d	0x0b99	0x1e34	0x1e6a	0x62e0	0x183e
48	0xfc41	0xcdf4	0xf8d0	0xa729	0x1c9c	0x2a4e
49	0xedb2	0x068e	0xd844	0xebab	0x10c6	0xfb09

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
50	0x0f31	0x0516	0x1d1a	0x027e	0x4f96	0xf3c3
51	0xcf30	0xdc5d	0x481f	0xcfc9	0xe3ba	0x4878
52	0xe7d7	0x21c9	0xe509	0xfc81	0x42d5	0x40dc
53	0xd958	0x6fa3	0x0b1d	0x0668	0x0b6d	0xfed6
54	0x3a78	0x9a7c	0x3a1e	0xa234	0x0717	0xe6b6
55	0x65fb	0x142e	0x52e9	0x3e01	0x5471	0x39e9
56	0xab4c	0x4036	0x5018	0xc7f6	0xe436	0xefbe
57	0x6fe7	0x005a	0xf9dc	0x0315	0xfc7a	0xffb5
58	0xfa39	0x09a7	0xf023	0x0e1c	0xf740	0x2aa2
59	0x21a8	0x4453	0x4367	0xbbd0	0x427e	0xc01b
60	0xaf0e	0xb75b	0x62ba	0x4538	0xf20b	0x069f
61	0xfc1b	0x17f1	0xe761	0x2bf2	0xd3a1	0xb2e5
62	0xffb6	0xf05f	0xf9d0	0x3448	0x00a2	0xff70
63	0xfdef	0x524c	0x1ef3	0xd37c	0x01a6	0xffe6
64	0x1bbe	0xcb25	0xb1a9	0x0a45	0xff4e	0xfe53
65	0x23f1	0x0558	0xa922	0x0a3f	0xafed	0x6139
66	0xfe50	0x1a13	0xfef6	0x2213	0x0050	0x6d78
67	0x4c25	0xf3dc	0xdbd3	0x0776	0xaaef	0x14e1
68	0x36ff	0xd31f	0x313c	0x17bf	0x4da5	0x0523
69	0x2ac3	0x266d	0xb74c	0x3d7e	0x12b8	0x025d
70	0xf90f	0x0eae	0xf009	0x54c0	0x1788	0xfdc0
71	0x0def	0xf206	0x3ffb	0x0a78	0xf928	0x02cc
72	0xec47	0xfa89	0xee3a	0xfd74	0xbac7	0xf2da
73	0xf1cd	0xeeec	0xe686	0xa978	0x1cd6	0x05b2
74	0x2fd2	0x4af6	0x160e	0xe179	0xb0bf	0x5360
75	0xe2ac	0x4df0	0x5bf6	0xd9e7	0x1625	0xf83a
76	0xf71d	0x3c4e	0x2a9b	0xba29	0x1961	0x350e
77	0xc1ea	0xc2e2	0xed94	0x1783	0x5eba	0xe7dd
78	0xf7ff	0xe538	0xfb48	0x0396	0x4547	0xffbb
79	0xf177	0x238b	0xc13f	0xa3bb	0x175d	0xf6d8
80	0x1eb6	0xdd2a	0x5de1	0x63a4	0xd4e7	0xfd1b
81	0xced4	0x4c0c	0x3939	0x3c5b	0xad16	0x0493
82	0x0836	0x047b	0x0ae5	0x1000	0x0883	0x222e
83	0xb8da	0xbaa2	0xd782	0xebad	0xfbd6	0xf22b
84	0xf4fd	0xb20a	0xd16f	0x1790	0x207b	0x2886
85	0xdc8a	0xf7cc	0x4be7	0xffef	0x02dc	0xfd4f
86	0xc750	0xb4e8	0xe449	0x4927	0x074e	0x597a
87	0x0f48	0x0293	0x63fd	0xf05a	0x2593	0x036d
88	0x0a38	0x58a7	0xe976	0x4600	0x0ee4	0x4efc
89	0x0a01	0x68df	0xeb83	0xd564	0x08d0	0xfdfb
90	0xec92	0x00c6	0xaa21	0xf1e8	0x569e	0xb614
91	0x533c	0xfb45	0x4ac8	0x4133	0xf9cc	0x2c7e
92	0xf902	0x0f77	0xf260	0x1b5b	0xe43d	0x518d
93	0xe824	0xb9dd	0xb6de	0x60bb	0x407c	0x0c8b
94	0x4fee	0x9e65	0xb077	0x1184	0xec09	0xfe22
95	0xe716	0xf832	0xd80b	0xfdcf	0xa9e9	0xa8bd
96	0x0be7	0xb65e	0x1da2	0x3997	0xb26a	0x18cf
97	0xec49	0x057d	0xda38	0x041f	0xaa87	0x2ba2
98	0x0d99	0xda1d	0x197e	0xbef1	0x591d	0x5593
99	0xb776	0x445d	0x3948	0x050b	0x13a2	0x4cdc
100	0x3f06	0xb29e	0xbdc4	0xb9ed	0xbddb	0x16a8
101	0xdfe0	0x132c	0x22e7	0x08e0	0xfb8c	0xa54f
102	0x0624	0x0ac1	0xf9c2	0x085f	0xf2ee	0xaa5a
103	0xd998	0xfbdc	0x9356	0x04be	0x1c79	0x0096
104	0x0062	0x0602	0x0217	0x4415	0xa562	0xfc7b
105	0x535c	0xb14e	0x0ce1	0xf930	0xdff1	0xac2a
106	0xefba	0xede7	0xba12	0x1566	0x0505	0x0088
107	0x4919	0x520b	0x60f2	0x2c9d	0x0502	0xedf6
108	0xf231	0x1dd4	0xfef7	0x085d	0xfcc3	0xf80d
109	0xf390	0x4d01	0x0ad7	0xfffe	0x0442	0x0068
110	0xe58d	0xb127	0x0b7a	0xf7b3	0xffdc	0x04f4
111	0x2558	0x24d6	0x2572	0x5654	0x3603	0x1898

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
112	0xfde9	0xb1ce	0x10b4	0xf8b4	0xfe40	0xbce1
113	0xa0e0	0x37a4	0xcab1	0xadd0	0x08df	0x2d23
114	0xf5aa	0x3c4d	0xee13	0x48ce	0xef35	0xfd92
115	0xb1a0	0x1049	0x46c3	0xfa84	0x359a	0xf8df
116	0xc019	0x2378	0x02e8	0x5605	0x007d	0x2a2a
117	0x25ac	0xc6f1	0xb7d1	0xc686	0x2ba6	0xaeee
118	0xfeba	0xa32e	0x1800	0x1ee5	0x025a	0x0604
119	0xe606	0x19ea	0xce75	0x5394	0x5131	0xe549
120	0x109c	0xadcd	0x15fc	0x48ff	0x5d34	0x2088
121	0x4642	0x1648	0xeb83	0xb953	0xfdd5	0x0c93
122	0x17cb	0x3798	0xec03	0xbbd0	0xb404	0xd27f
123	0xabae	0x2c26	0x3c4a	0x63f6	0x1a79	0x97c5
124	0x536b	0xdfcc	0x16f5	0xf22c	0x17bf	0xf5f9
125	0x0a2b	0xf669	0x152d	0xd002	0xb564	0x15c6
126	0xf947	0x98e7	0xa390	0x5978	0xfea3	0x0ecb
127	0x088d	0xfb4d	0x14dc	0x0cb1	0xa7a7	0x0068
128	0xf980	0xd4f4	0xf4d7	0xaf0d	0xa20f	0x4dbc
129	0x5959	0xe34f	0xb7cf	0xc6e8	0xdf30	0xcd5b
130	0x0ec1	0x0f76	0x202f	0x500e	0xe4b1	0xfb4f
131	0xff60	0xf9b3	0xfce7	0xde17	0x023d	0x0308
132	0x10c9	0xf136	0x4f95	0x17c2	0xeb37	0xb820
133	0x4939	0x099f	0x3102	0xe1bb	0xe1ca	0xf779
134	0x2b42	0xed90	0x5667	0x0721	0xa0e1	0x0ff0
135	0x05df	0xb516	0xf9df	0x000d	0xfec7	0x0177
136	0x013e	0xfdc1	0x09f0	0x00b2	0x0066	0x0028
137	0xc184	0x96ef	0x1390	0x0cf8	0x02ae	0x0487
138	0x649b	0x6906	0x023e	0xe8d6	0xf0b4	0x057f
139	0xdc44	0xe20f	0xf4c5	0xdf40	0xb719	0x6720
140	0xe2eb	0xb988	0xb824	0x2262	0xf734	0xaa82
141	0x1eab	0x2dfd	0x6b5d	0xcdd1	0xfa7e	0x4c86
142	0x08c0	0x173b	0x2bef	0x3e6c	0xe69d	0x5ed8
143	0x54a9	0xb7ad	0x262b	0x1996	0xf556	0x014e
144	0xefcb	0x0628	0xd4fe	0x0059	0xa093	0xe9b2
145	0x1e28	0x05c6	0x53a4	0x9e3f	0xdf3f	0x0009
146	0xf670	0x27ea	0xce2c	0xc131	0x0489	0xacdc
147	0xddcb	0xc7a3	0xa67a	0xc624	0x0a45	0x3614
148	0xe3ac	0x0b1b	0xda59	0x0b42	0xc6df	0x5fb1
149	0xfd5e	0xe67e	0x019e	0xa4db	0xaca1	0x01c6
150	0x0838	0xe758	0x2a87	0x46a7	0xfb51	0x00af
151	0xfe13	0xfdce	0xf54d	0x0076	0xfbce	0x005d
152	0xd8e5	0xf015	0x9259	0x56a4	0x3ae5	0xfd84
153	0xede3	0xbfe8	0xdcd5	0xb03e	0xd2a8	0xae3c
154	0x12cf	0x3e14	0x5eae	0xcabe	0xf3fe	0xfbdd
155	0xe5bc	0x1202	0xb6ac	0xc44d	0xbed3	0x5db4
156	0x3bf5	0xfd5e	0xf19e	0x54af	0x117b	0xd0c8
157	0x1294	0x0a21	0x14ea	0x1771	0x3ad7	0x677a
158	0xa2f9	0xbc9d	0x1b20	0x017a	0x02b6	0x029e
159	0x5b60	0xdd79	0xc687	0x1d78	0xfc94	0x2b50
160	0x0e38	0x0d08	0x5841	0xf259	0xf6e8	0xff8f
161	0x011c	0x1b02	0x0c19	0x27bb	0x19ee	0xb743
162	0x09a8	0x1758	0x2b2e	0xd160	0xfda5	0xfd69
163	0x3f2f	0x4039	0x336c	0xf035	0x123b	0x1d07
164	0x4b8a	0x3cae	0xe67b	0x0691	0xed07	0x4298
165	0x4283	0x0214	0xb588	0xfa5f	0xebf6	0x043d
166	0xceb7	0xbb37	0x080e	0x9d0c	0x4a41	0xc107
167	0x2748	0xadf8	0xcabe	0xf47b	0x3c07	0x4dde
168	0xfd78	0xf9bb	0x273e	0xf9c8	0x33f0	0x4d60
169	0xfbe2	0x29f8	0x021a	0x616a	0x259e	0xdca4
170	0xd88d	0x0be2	0x9e0c	0xa20c	0x3693	0x0064
171	0x1993	0x1afb	0x1b77	0x286c	0x5cdf	0xba22
172	0xa6f7	0xf840	0xfa8f	0xf2fe	0x2433	0x37ed
173	0xc7f6	0xf081	0x0be2	0x3f7e	0xbc69	0x25ae

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
174	0xac6f	0x5c4c	0x4185	0x02cc	0x0a67	0x0072
175	0xb5b8	0xf422	0x0626	0xff0b	0x05b7	0xfce7
176	0x578a	0x5b91	0xc6d3	0xfdee	0x439e	0x3531
177	0xd2c2	0x1eff	0xc97e	0x5ba9	0x9fcc	0x67b6
178	0xfbeb	0x0e5f	0xf756	0x294c	0x5207	0xf18a
179	0xc367	0x00c5	0x414e	0x9fe5	0x1351	0x0005
180	0x2a1d	0x10ef	0x68a6	0xdc9d	0xc0e8	0xf4e8
181	0x3ecb	0xa1dc	0xf0a3	0xe54f	0x3165	0xe48b
182	0x0830	0x9c1c	0xdf4e	0x1a9e	0x000b	0x049a
183	0xd1b8	0xfdb9	0xdd47	0xafc1	0xd719	0xfe84
184	0xf649	0x60c9	0xab79	0xb473	0x067c	0xfd24
185	0x0909	0x356f	0x0ff5	0x5fe5	0x6073	0xad45
186	0xf6c2	0xfe08	0xefde	0xd6b6	0x5c74	0x07a9
187	0x4f9b	0x4591	0xdade	0x0e95	0xb5f6	0xe76c
188	0xf0f0	0x41a2	0xfc5f	0xb0aa	0xbab5	0x1a8d
189	0x308f	0x17be	0xd3f8	0xc78e	0x1b01	0x5bb4
190	0x1dd4	0xf989	0x59e9	0x29df	0xdf9c	0x0346
191	0xde91	0xfb2d	0xb950	0x0f39	0x3edd	0x05d2
192	0xf1fe	0x2054	0x3b3d	0xf131	0xad63	0x06cd
193	0xee6f	0x54eb	0x093e	0xfeea	0xed48	0x3cbd
194	0xa5ae	0xca74	0x1df4	0x3f68	0x5e38	0x3ab1
195	0xb1b5	0x3215	0xb140	0x4133	0xd279	0xc12f
196	0xcec7	0x4f0f	0x0da8	0xf60b	0xe5a7	0xd1b6
197	0x1159	0x1e84	0x512f	0x42b8	0x2d03	0xda55
198	0x60be	0x212e	0xa4fe	0xf342	0x2b5d	0xe430
199	0xd885	0xe239	0xa978	0xb881	0x6815	0x254e
200	0x9c33	0x01dd	0x1ec2	0xf9fe	0x0463	0xff58
201	0x01d6	0x266a	0xfea5	0x5d89	0xd773	0xdb05
202	0xf000	0xda1a	0xe538	0xabd8	0x516d	0x1c06
203	0x14fa	0x2614	0xa32b	0xfb5a	0x0200	0xf9fe
204	0xfc12	0xd8c2	0xce97	0x4b22	0xf902	0xfc86
205	0x3b04	0x5c44	0xc2e2	0xf626	0xfb4d	0xfad3
206	0xe312	0xf5d3	0x0447	0xff09	0xfe27	0x00b1
207	0x1f99	0x0004	0x3088	0xa8f4	0x28a5	0xe1d0
208	0x56b4	0x2a17	0xec4d	0x02b2	0x0216	0xff2c
209	0xf3af	0xfa76	0xbe3d	0x47fa	0x3dcd	0x59ac
210	0x1631	0xf74b	0x0c7c	0xf2aa	0xaac7	0xc629
211	0x0013	0x0313	0x0408	0x00aa	0xdf99	0xfd7b
212	0xfc8e	0xf6f1	0x961f	0x01b0	0xeed8	0x05db
213	0xfab6	0xd1d5	0xffb4	0xb064	0xd7cb	0x2c40
214	0x00d3	0xed6f	0xedbd	0xe4eb	0xcb1e	0x388f
215	0x179b	0x148c	0xfe35	0xfe32	0x008f	0xffbf
216	0xf5f4	0x1c58	0xf30b	0x23fc	0xa570	0xd8fa
217	0x9ece	0xdac4	0x49ba	0x17d5	0x097d	0xc76e
218	0x207a	0x08e5	0x3770	0x0db8	0x6519	0x55f0
219	0x00d0	0x4efa	0xfee7	0x9f36	0xffc1	0xfb61
220	0x0447	0xe86e	0x0a92	0xaa51	0xf5a1	0x0233
221	0x0017	0xe8d6	0x00f3	0xdce3	0x14e1	0x504e
222	0xc396	0x319b	0x1040	0x2b4f	0x508d	0xd750
223	0x5203	0xffab	0xdeec	0x00c2	0x03eb	0xdad5
224	0xb34b	0xf2f9	0xc8ff	0x0df6	0xa4ab	0xfd65
225	0xf7e4	0x0da1	0xf388	0xb459	0x021b	0xfa06
226	0x1cb8	0xc493	0x5844	0x4ba9	0x0413	0x40f3
227	0xf8b0	0xfe63	0x04d3	0xeb64	0xf222	0x558f
228	0x1efb	0xf828	0x4248	0xe571	0x72d1	0xf655
229	0xcaeb	0x20c5	0xa3ac	0xa9b5	0xc89e	0xc827
230	0xd2c9	0xb186	0x3eb8	0xf8c8	0x3d69	0x1194
231	0x0f09	0xbf3b	0x4ec1	0xad5d	0x1e62	0x2e58
232	0xe66d	0xfb07	0xb66b	0xd42e	0x2d74	0x0414
233	0x09e0	0xe5dd	0xba03	0xd39e	0xece2	0xfc10
234	0x04d9	0x10a4	0x090f	0x17df	0x0d9d	0x4ef1
235	0x0bc6	0xf418	0x14c4	0xee45	0x515f	0x21fe

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
236	0xf902	0xc6a5	0x0116	0x3684	0xd8af	0xd6cd
237	0xa734	0xe0eb	0xfb7e	0x35fd	0xfa34	0xfb21
238	0xe36b	0xfd99	0x3326	0x49ef	0x26a9	0x05ac
239	0x09f8	0xf6de	0x0d60	0xedea	0x2b74	0xb380
240	0xd48b	0xafb7	0xd599	0xd5e1	0xaee1	0x1ab1
241	0x03d8	0xc509	0x168f	0x6225	0x1501	0xb2a9
242	0x0205	0x33d8	0xe2de	0xf951	0x5084	0xe883
243	0xac57	0x33c3	0xaec5	0x3489	0x4381	0x3330
244	0xc23d	0xc088	0x5a35	0xf03b	0xdffd	0x0367
245	0x0246	0x311b	0xad77	0xc652	0xdc1d	0x1635
246	0x10de	0xf910	0x2ca1	0xba9d	0xd93f	0x0241
247	0x177d	0x41be	0x44f7	0x9b5a	0xeed0	0xf222
248	0xca50	0xbf63	0x0e34	0xf2fe	0xad9d	0xc1f2
249	0x19a5	0xd475	0x21c9	0xccc6	0x5b31	0xcb03
250	0xf612	0xdcaa	0xe27a	0x7238	0x0e75	0xfe81
251	0xd68c	0x61a3	0x0765	0xdfee	0x51b8	0xc0ae
252	0x149c	0x4156	0x29a3	0x4de4	0xed41	0xb484
253	0xfdec	0xdbac	0x6cd0	0x1365	0xff0f	0x0218
254	0xfd03	0xaf1e	0xf2ac	0x49b6	0x0acd	0x058c
255	0xf40d	0x0a94	0xb5b2	0xfeb5	0x0dd1	0x0074

Table E.3.7: VQ table for hebap 7 (16-bit two's complement)

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
0	0xad4b	0x5585	0x2896	0x354e	0x29de	0xdc27
1	0xa809	0xdfff	0x4798	0xe61b	0x63ae	0xd5a0
2	0x1a90	0xca42	0xcc22	0x5792	0x394b	0xae36
3	0x092b	0x2914	0x0465	0xf281	0x15c1	0x6a00
4	0xe627	0x2e4b	0xa034	0x5999	0x4f8a	0xe87d
5	0xaaf8	0x29b9	0xb361	0xc553	0xdee2	0xf7df
6	0xf547	0x21ec	0xece1	0x6c5b	0x1d82	0xd147
7	0xfc04	0x099c	0xfc46	0x1292	0xfd8d	0xc010
8	0xb30a	0x5a39	0x004b	0xca8c	0xf5ac	0x083c
9	0x0fd1	0xf4c8	0x16db	0xee95	0x5686	0x3110
10	0xacc8	0xbd17	0xfd26	0x1cfb	0xd276	0xd6e5
11	0x2c44	0x0700	0x682a	0x5bde	0xb397	0xfe15
12	0xba5d	0xbe77	0xcada	0xc7cb	0xa9f3	0xf660
13	0x0443	0xe8b1	0xe0d9	0xbdaf	0xb030	0xaa55
14	0x47d4	0xfbb1	0x078d	0x341e	0xbbc9	0x46c2
15	0x5876	0x43c1	0xd912	0x45ff	0x4762	0x02ba
16	0x05cc	0x4f49	0xe986	0x986d	0x134d	0xa909
17	0xf5d5	0x11eb	0xe92e	0x4820	0x223f	0xf5f8
18	0xf513	0xf9be	0x54d1	0x0c1b	0x9bad	0x0c98
19	0xb5ad	0x1255	0xec71	0x17ac	0x07b4	0xc509
20	0xf773	0x252c	0xfdee	0x50bd	0xedca	0xdf93
21	0xa8cb	0xdd49	0x09e1	0xd3a8	0x1564	0x03e6
22	0x5654	0xec44	0x0673	0xf59f	0x1207	0x090f
23	0x5177	0xf3fa	0xf2fe	0x1009	0x349e	0x0bfd
24	0x0055	0x4389	0x2818	0xc660	0x00d6	0x005a
25	0x9903	0xb65f	0xb468	0x4b2c	0xd816	0x26b5
26	0xd9f5	0x5011	0xe64d	0xe4b9	0x0b4b	0xfd1e
27	0x505f	0xc20c	0xa67f	0x1ad6	0x004c	0x0147
28	0x2228	0xcdb3	0xa65f	0xf6bc	0xb420	0xd9d5
29	0xcdaa	0x3f37	0x525c	0x0eed	0x02ed	0xca20
30	0xc185	0x47df	0x0957	0xbb03	0x4c1c	0xe87e
31	0x058f	0x2dd6	0x0fd3	0x4b5a	0x1ac9	0xb31f
32	0xebb0	0x2626	0x4746	0x099f	0x494c	0xed0c
33	0xfdab	0x4c2a	0x052b	0xdc78	0xfecc	0xfbb0
34	0xf3e5	0x9b7d	0xc2cf	0x62f4	0x121a	0x0a4b

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
35	0x4ca7	0xf6b0	0xe117	0x2e14	0xdb8b	0xc8fc
36	0x0a52	0x6755	0xad9d	0xd78e	0xf963	0xf951
37	0x560f	0x5479	0x2d3c	0xa67d	0xefd3	0x0081
38	0xe816	0x0dd6	0x0393	0xfefb	0xffef	0xfe81
39	0x06a0	0x1a30	0xfa6f	0x5166	0x0359	0xeec0
40	0x058f	0xc450	0xde9a	0xda3d	0x145a	0x1637
41	0xee58	0xfd9b	0xd25d	0x15f2	0x1086	0x026b
42	0x03a9	0xec9d	0xc8ea	0xbd30	0xe506	0xe8c0
43	0xc524	0xfe1f	0xe3bb	0xc5d2	0x49bc	0x54a9
44	0x9bc6	0x0b5e	0x0477	0xfeb9	0xfe36	0xfc1d
45	0xda48	0xfccd	0x9ebc	0x0af4	0x4f01	0x043b
46	0xfba9	0xf19e	0xf904	0xb3dc	0x03c6	0x0335
47	0x1c7d	0xab01	0x2a26	0xe46d	0xa503	0xf945
48	0xfee6	0xd4ab	0x00aa	0xae2a	0x8f02	0x3147
49	0x4612	0x0e81	0xf9e5	0x0375	0x0005	0x0234
50	0x17e4	0x58a8	0x08c2	0xe4d9	0x26f7	0xe80c
51	0x10f2	0x68b8	0xf187	0x07b8	0xfbc9	0xf5f6
52	0xfd6b	0xe123	0xf594	0xc4a6	0x453a	0x1117
53	0xefb2	0xd4d3	0x02cd	0xa816	0x061a	0x2fdc
54	0xe6fb	0x479e	0x17d7	0x1b47	0x1744	0x4713
55	0x267b	0x14fa	0x5c34	0xe533	0xe657	0xffc2
56	0x55cc	0x342f	0xfd55	0x0ec9	0x0878	0x00d1
57	0xf20f	0xfb99	0xb2f4	0xf9f8	0x051c	0xfcdd
58	0xf3f5	0x3eb1	0xca21	0xf3fb	0x10c6	0x5ca1
59	0xd8f1	0x26d7	0xc200	0x3286	0xa3b1	0x54c3
60	0x25fa	0x5935	0x2fa0	0x3af3	0x159d	0x12e5
61	0x08c3	0x0833	0x04db	0x0ff9	0x128c	0x329c
62	0x0fa7	0xf65c	0x0d19	0xf3ec	0x228b	0x4280
63	0x10ea	0x17ef	0x15ad	0x2421	0x2bda	0x6fb0
64	0xda8e	0xdd87	0x00ec	0x03f1	0x01c7	0xfc3c
65	0x1aad	0x4b5a	0xfc06	0x00c8	0x071d	0x0242
66	0x144c	0x03bd	0x2884	0x0d02	0xce00	0xff81
67	0xf432	0xdfff	0xc723	0x562d	0x1720	0x041d
68	0x2ae6	0x6556	0xa01e	0xa512	0xd17f	0xe57b
69	0x588b	0xd4fe	0x1668	0x0a07	0x5c99	0xd7f3
70	0xf2f1	0xef77	0xeaae	0x50bb	0xd5a5	0xf1eb
71	0xefdd	0xf1e4	0x11df	0xfcc3	0xfea2	0xfcb1
72	0xf319	0x0d7b	0xe31a	0xd2ac	0x0bcf	0x01c7
73	0x0c80	0xdab5	0x0c82	0xa693	0x2bb0	0x989e
74	0xc8f3	0xefeb	0x3c16	0x37d7	0xd55d	0xb067
75	0x0edf	0xd4f8	0x5624	0x3822	0xc420	0xe1cb
76	0xe76d	0xbac9	0xf9e8	0x2f10	0xb2a3	0xfe45
77	0xe7fd	0xef76	0xff60	0x20ab	0x586e	0x2e87
78	0x4afd	0x0497	0x1cfe	0xd96d	0xefd8	0x1260
79	0xffb8	0xe21c	0xff90	0xd14d	0xf362	0x6a27
80	0x0cca	0x174b	0x1d4d	0xbd85	0x0362	0x9c94
81	0x02e1	0x0745	0x0729	0x07e6	0x0950	0x1293
82	0xeb9f	0x1d58	0x0cfb	0x0a9b	0x0bf9	0xf9ba
83	0x1097	0x0235	0x15fd	0x09c1	0x4663	0xecc8
84	0xf4ef	0xba00	0xe082	0x3d7a	0xfc06	0x0858
85	0x0bea	0xb3e5	0x4222	0x748b	0xd812	0x3b31
86	0xd3ae	0x0076	0x9b3c	0xca3f	0x3bd8	0xfe2c
87	0xed28	0x1360	0xef59	0x0627	0xd69f	0x4c69
88	0xdff5	0xfa7f	0xfd05	0xfb8d	0xfda1	0x0580
89	0xf765	0xd369	0x07e5	0xe70c	0xf5d8	0x02c7
90	0xfe63	0xf631	0xff28	0xf241	0x9195	0x06b7
91	0xc792	0x429a	0x365d	0x34bb	0x9b5e	0xc107
92	0x4b1c	0x1cad	0xcff3	0x02aa	0xf131	0xff39
93	0xef9f	0x510a	0xc2dd	0x2c55	0x16e4	0xfcc8
94	0xac0e	0xf226	0xfffd	0xf957	0xf089	0x23fd
95	0x3c55	0xf8ac	0x07dc	0xb355	0x3f64	0xed13
96	0xf4cd	0xf16b	0xe346	0xff51	0xb169	0x2ba6

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
97	0xf20d	0x9ff5	0x4cf4	0x19fe	0x03d3	0xfd72
98	0x553c	0xe2fa	0xe611	0xd5f1	0xdf56	0x3cb7
99	0x39eb	0x4639	0xe3dc	0xf2af	0x0772	0xbc78
100	0x0dc5	0xf095	0xfa79	0xf512	0x44f0	0x0822
101	0xe64c	0xc469	0xba07	0x0539	0x3bea	0x52a6
102	0x1842	0x25e2	0x3b33	0x9fa6	0xa815	0xf061
103	0xf934	0xfdaf	0x0447	0xe19d	0x61e2	0x15e1
104	0x53a7	0xfe50	0xf986	0xe50e	0xfa62	0xc78a
105	0xe4e1	0x02bc	0xd095	0xfd17	0xa185	0x57c2
106	0x188f	0x0cd3	0x2afe	0x0f04	0x4af0	0x39bd
107	0xa81a	0x3baa	0x1543	0xf508	0xfc36	0xf2f1
108	0x0cb9	0xf184	0x1288	0xdf93	0x591e	0xd820
109	0x5f1a	0xae16	0x4d86	0x03db	0xd14a	0xe77b
110	0x0f42	0xb30b	0x3304	0xf9b7	0x48d1	0x1d2a
111	0x98d7	0xa7eb	0x3fb1	0x07de	0x2adf	0x4670
112	0xe481	0x122f	0xc61e	0x4933	0x3dad	0x0510
113	0x245e	0xf96f	0x394b	0xf302	0x67a7	0xd1b3
114	0x1660	0x171d	0x3458	0x2724	0xf744	0x9f80
115	0x06cd	0xe5b9	0x3197	0xaa07	0x0ff0	0x154a
116	0xf5c3	0x24b1	0x5297	0x9aae	0xf3a6	0xf61f
117	0x50c0	0x49ce	0xc98d	0x1b4e	0xdfbc	0x3dc3
118	0xa2f6	0x2baf	0xcab9	0x2e5c	0x3ead	0x0a46
119	0x47b9	0xd814	0x033d	0x0358	0xfc0e	0x009d
120	0x3840	0xedba	0x1421	0xcc16	0x94d6	0xd4ec
121	0x546d	0x2bf8	0x442d	0x1db4	0x334a	0xfe1c
122	0x0007	0x04d4	0x023d	0x1076	0x15c8	0xf3f7
123	0x0394	0xdc7c	0x0505	0xdd02	0x04a1	0x8fe5
124	0x5453	0x5c8f	0x4aac	0xf4bb	0xc836	0xdf0a
125	0x5b76	0xe7ef	0x32b2	0x0bf5	0xdb79	0x08bc
126	0xf402	0xe350	0xb154	0x169c	0x0246	0xfdd9
127	0xf067	0x013b	0xe1a3	0x2020	0x924e	0xcf4f
128	0x35c6	0xc403	0x4b05	0xaf70	0x32f3	0xb4d1
129	0x0ec1	0xff4f	0x1f5d	0xfc17	0x4594	0x142a
130	0xe374	0xef19	0xb950	0xfd94	0xfaba	0x3a54
131	0x39a4	0xfb3b	0xcded	0xc5b6	0xfddd	0x69f5
132	0x08ba	0x06ac	0x0acc	0x1528	0x1f32	0x9db5
133	0x0b39	0x0e34	0x0f98	0x14e0	0x279e	0x530b
134	0x0486	0x1503	0x01fc	0xd6ee	0x0122	0xf9b1
135	0x045a	0x60d5	0x40bf	0x9db0	0xfed6	0xf4f0
136	0xfbad	0xe800	0xf882	0xe191	0xf465	0xa514
137	0x0fb0	0x2a29	0x43a5	0xef0a	0xae0a	0xf2c9
138	0xee72	0xff31	0xd921	0xf209	0x1f0b	0x0482
139	0xe268	0x1fb5	0xc921	0x4256	0x98a7	0x946c
140	0xc4c4	0x3ee0	0xbe34	0xdd4a	0xa358	0x3e22
141	0x615a	0x1630	0xf8ae	0x01a4	0x0084	0x0075
142	0xfe06	0xb492	0xff3a	0x019c	0xfec9	0x02f0
143	0xf88e	0x0f8d	0xe1f8	0x40b6	0xb4a5	0xc67e
144	0xfe71	0xfd27	0xf121	0xef9c	0xcf95	0x1dd7
145	0x0d28	0x091a	0x2384	0x5c86	0xd7ce	0xf957
146	0xf3b4	0x0a24	0xe0ce	0x390a	0xed39	0x40f3
147	0x1f79	0x05c9	0x0031	0x4335	0x6125	0x1d32
148	0xb498	0xfdff	0x2e81	0x092a	0x15d4	0x0d25
149	0xec39	0xaacc	0x2c6a	0x2a90	0x1311	0x0105
150	0x12ba	0x5061	0x13f5	0xe877	0xe08f	0xfa0f
151	0x1fbd	0xc65c	0x509f	0xc5ba	0x5d85	0xf1be
152	0x30a3	0x0565	0x0e1d	0x21ef	0xa23e	0x12f0
153	0x1a46	0x2993	0x2766	0x6281	0x9db9	0xfbd7
154	0x19a1	0x3699	0x0b5f	0x54e9	0x4051	0x9a3e
155	0xf910	0x0a0f	0xb36a	0xbe60	0x0bd8	0x1a17
156	0x3aa4	0xba0a	0xdf0a	0xabce	0x9619	0x2e20
157	0x0d78	0xfc64	0xc1d7	0xfb91	0x1406	0xaf7b
158	0x1e28	0x08b2	0x4437	0x153a	0x710e	0x4490

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
159	0x04de	0x3cfe	0xd221	0x602a	0xbb7d	0x0cc8
160	0x0c8f	0x461e	0x0adf	0xfd2e	0xa770	0x175b
161	0xe9d2	0xf390	0x9a19	0x65b2	0x19b7	0x0ce6
162	0x4f56	0xf21d	0xf565	0xfe44	0xfa31	0x05f6
163	0xaf60	0xaa2e	0xd051	0x9b3f	0x229f	0xfbf4
164	0x45e0	0x023a	0xc11a	0x2089	0xf607	0x3bab
165	0xf58b	0x26de	0xf8a9	0x405d	0xce26	0x8eb1
166	0xff88	0xf753	0x00db	0x0061	0x016d	0x0023
167	0x04f6	0xfd32	0x05c8	0xf57f	0x078a	0xe299
168	0x0768	0x222e	0x0772	0x473b	0xce6c	0xe7e2
169	0xf16b	0x3591	0xd966	0xc1a8	0xfaa0	0xe416
170	0xd698	0x2130	0x3e5f	0xdda8	0x1d6c	0x4fd7
171	0x0be1	0xcb6f	0x0408	0x96b8	0x169b	0x6198
172	0xee12	0xdfe4	0xdb96	0xe820	0xbca5	0x6491
173	0xba70	0x1b3a	0x0ea8	0x0272	0xff8e	0x0882
174	0x1161	0xed02	0x1b8e	0xeae4	0x1282	0xf4f5
175	0x133a	0xfd75	0x49fb	0xd976	0x0350	0x075e
176	0xfeb0	0xeade	0x1c42	0x4fdc	0xda91	0xfda8
177	0x030d	0xb3ee	0xce98	0x19ea	0x0586	0x01c2
178	0xf2b9	0xbe7e	0x2b63	0x3390	0xea86	0x549f
179	0xf33f	0x12fb	0xe8b7	0x1d6a	0xd5ab	0x6db6
180	0x286e	0xcd9b	0x6463	0x6428	0xfd81	0x015f
181	0x048b	0x494b	0xeaa6	0xc511	0xff6f	0xfa9f
182	0xc773	0x6a5d	0x8569	0x8073	0x53bf	0xf4b2
183	0x3c3c	0x4987	0x5670	0x4bc6	0x5837	0x3513
184	0xd64e	0x29d6	0x13e1	0xed6c	0x038d	0xaee8
185	0xcd6c	0xaf4c	0x1cf2	0x0aa2	0x0d63	0x2d41
186	0xfbaf	0x47c6	0x4d13	0xda4e	0x57aa	0x4cb2
187	0xfed8	0xe572	0xc6ab	0x5463	0x4d54	0x5397
188	0xb46a	0xe2b3	0x6366	0x3358	0x218c	0x9d2e
189	0x0c14	0xd686	0x51a0	0x2421	0xf302	0x0704
190	0xfcd5	0x05a9	0x0c22	0x128c	0x2f29	0xc84a
191	0xaf10	0x37c3	0xef14	0x9b12	0xe96b	0xad63
192	0xebf4	0x293a	0xc93c	0xa97a	0x0b18	0xfdd6
193	0x63bd	0x44f0	0x3a26	0xadae	0x099b	0x6236
194	0xdb66	0xf904	0xcdc2	0xe912	0x9b2d	0xd4f1
195	0x1a2a	0x0333	0x2849	0x00a6	0x6bbd	0x020b
196	0x0065	0xb444	0x0d55	0x25a6	0x0040	0x0326
197	0xf54a	0xb9f5	0xf5f0	0x5922	0x2169	0x0466
198	0x0b9c	0x3b63	0x0700	0x635a	0xe9a0	0xbc8f
199	0xfa75	0x0644	0x112e	0x2cbc	0x06c3	0x5ceb
200	0xebf0	0x1211	0xd663	0x6d4d	0x26a9	0xf632
201	0xd6e0	0x927f	0x0bb7	0xfa06	0xfcc0	0xfcc2
202	0xd483	0xcf21	0x56be	0xe3b5	0xa3e6	0xab3e
203	0x4227	0xaa7c	0x0745	0xda7a	0x24d8	0x4a52
204	0x2825	0x252c	0x68bf	0x07da	0xecb1	0xdc88
205	0x15ab	0xf75e	0x37be	0xc43c	0xb48c	0x071e
206	0xed0e	0xfcf1	0xdd01	0xf3fc	0xb1a8	0xf383
207	0x2028	0xf516	0xbaa8	0x33fc	0x0c9d	0xfc21
208	0xd033	0xe64b	0x284b	0xdab0	0x08d4	0xaf58
209	0xe4a8	0x1599	0xe27f	0xe2be	0xd79a	0xd7e6
210	0x0e39	0x4c17	0xe8ac	0xb567	0xb776	0x3205
211	0x0503	0xefbc	0x1066	0x90c7	0xf63e	0x074a
212	0x3eaf	0x68ca	0xcd03	0xe754	0x03d9	0xf9c3
213	0xfe6d	0x3570	0x1939	0x61ee	0x69f4	0xaf1a
214	0xb96a	0xf902	0x9e66	0x1741	0xfc46	0x67e8
215	0xa160	0xc3e9	0x60d4	0x07a1	0xfb90	0x00bb
216	0xf70f	0x30d9	0xaefe	0xfc78	0x4794	0x530a
217	0x0a62	0xe804	0x3f33	0x5704	0xfdd4	0x086a
218	0xe839	0x367e	0x9bae	0x93bf	0x0fd1	0xed45
219	0xac34	0x6769	0x4beb	0xdc5c	0x037f	0x012f
220	0xa948	0x99bf	0xe876	0x6099	0xa672	0xdcba

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
221	0xc83c	0xc192	0x5cb4	0xa6bd	0x2434	0xf0ff
222	0x732a	0x55a3	0xe7bb	0x068f	0xf7f5	0xfba0
223	0xfe4d	0x264a	0xf0cd	0x3047	0xef40	0xb5e5
224	0x4d38	0xffaa	0x09a3	0x07c6	0xfc03	0xeb16
225	0x51fa	0xddb1	0xeb2f	0xa3f6	0xed86	0x0a71
226	0xec19	0x15e5	0xedeb	0x4ace	0x65b5	0xd01d
227	0x03cc	0x1aca	0x11c7	0x6d2d	0xf047	0xf720
228	0x17bb	0xf344	0xec83	0xfe8b	0xf9dd	0xf16e
229	0xe3a8	0xcd40	0xdd8c	0xec0b	0x5a0e	0x13be
230	0x0398	0x0a37	0x1ee8	0xe347	0xecd7	0x4eda
231	0xff06	0x154e	0x0c44	0x1b10	0xb6dd	0xf7fd
232	0xd7c5	0xeeec	0x4c98	0x130f	0xfd6b	0xf8a3
233	0x39e0	0xde65	0xb299	0x17f7	0xad26	0x371c
234	0xd157	0xf751	0x139a	0x2e74	0x58d5	0x0196
235	0xcc80	0xf5cb	0xcc38	0xa85f	0xcf3e	0xdf44
236	0x42aa	0x62b3	0xf71f	0x13c0	0xfeaa	0x0091
237	0x20d1	0xbaed	0x4aa8	0x2977	0xb403	0x42bb
238	0x5155	0xe9bc	0x300a	0x9c02	0x2897	0x1e0c
239	0x11c6	0x3da3	0x43ba	0xb44d	0xed60	0x04b6
240	0xe1d5	0x2a54	0x95e4	0xd351	0x1ab3	0xf910
241	0x09ee	0x0c7f	0x115a	0x4469	0xf181	0xfc6e
242	0x51e0	0xbe7a	0xe94a	0x2b4f	0xffba	0x59b1
243	0x0ce9	0x0b67	0x1870	0xed40	0xae1a	0xf362
244	0x1724	0xbf5d	0x0887	0x0aad	0x0d76	0xa4f6
245	0xe853	0xff3e	0xc9e4	0xd525	0x4c20	0x0405
246	0x1173	0xe8b4	0xb5c4	0x05ef	0xfe99	0x0357
247	0xf9d3	0xe249	0x5636	0xd2c4	0xd8d0	0x42ce
248	0xcf84	0x09f9	0x10e4	0x57e4	0x1677	0x2f8a
249	0x9dd9	0x464c	0xe710	0x049c	0x049e	0x2596
250	0x5ba6	0xdee9	0xeed8	0xf593	0x1dd6	0xbe3d
251	0xea79	0xf4b9	0xd5fb	0xae6d	0x1c4e	0x041d
252	0x0a8f	0xaf86	0xe27e	0x1d5c	0xe1c4	0x16ec
253	0x50be	0x558d	0x01c9	0x3a79	0xbb07	0xd16f
254	0x0e13	0xf9c5	0xf77f	0xff63	0xffd5	0x025d
255	0x09d1	0x22fa	0x291f	0x581f	0xc11c	0xc157
256	0x1772	0x1357	0x1a8b	0xed02	0xa880	0x49a1
257	0x1da6	0xf963	0x9f90	0xf2b4	0x3759	0x04be
258	0xeed2	0xe5f9	0xe52a	0xd89d	0x9fec	0x2425
259	0x28e4	0x4557	0xe1bc	0x0093	0xe756	0x1143
260	0x3f3b	0xbf53	0xefe9	0x10ce	0x1dc9	0x1521
261	0x0ce7	0x0aaf	0x1d22	0xb242	0xf732	0xf18a
262	0xf7e3	0x5469	0x3a16	0x3101	0xe83f	0xf91c
263	0x1246	0x2ddc	0x0b2b	0x1b29	0x077f	0xf0e1
264	0x0dc2	0xaaa3	0xf65b	0xd72b	0x49cd	0xd60a
265	0x0eaf	0xd831	0xecfe	0xf59d	0xba59	0xfb26
266	0x3a8f	0x2487	0x2e5e	0xf9db	0xed10	0x5815
267	0x2525	0x95f0	0x29ee	0x5173	0x99b7	0xba2a
268	0xe3fe	0xfa90	0xb3d9	0x31ca	0x2006	0xf83c
269	0x075b	0x6dfe	0xfcb2	0xe3bd	0x00f9	0x00e9
270	0xe3e0	0x029d	0xfe8d	0xf47c	0x5ac2	0xe9fd
271	0x0c45	0x0120	0x0c97	0xfb16	0xff9e	0x9429
272	0x43dd	0xa53d	0x13f6	0xd441	0xf5f2	0xd321
273	0xecc0	0x05ee	0xeab0	0x029e	0xb89a	0x079f
274	0x285e	0xb267	0xedd7	0x0169	0xff60	0xfc65
275	0x492c	0x37b8	0xf3ad	0xe2c3	0xf300	0x1747
276	0xf1e2	0x5255	0x1c6c	0x0dd0	0x1fb9	0xfa08
277	0xdf1a	0x01f4	0xb512	0x49f1	0x6718	0xfbf1
278	0x3d17	0x6444	0x20b7	0x076f	0x0799	0xd135
279	0xf564	0x0d3d	0x68e2	0xee15	0x070b	0x0016
280	0x0499	0xfd71	0x04d1	0xf7b0	0x1ea4	0x06e7
281	0xfd07	0x2011	0xb4a6	0xee0f	0x0783	0xfea9
282	0xfd4f	0xf236	0xf33d	0xf124	0xf53f	0x4886

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
283	0xf7c2	0x07aa	0xfab7	0x4103	0x0acd	0xa5c2
284	0xfe4f	0x1329	0x012e	0x32d8	0x3e3d	0xe8ef
285	0x0c83	0x101e	0x2bad	0xea88	0xf61f	0xfb78
286	0xfbbd	0xe6bb	0xfa79	0x1632	0xfef4	0x0247
287	0xdb43	0xb38c	0x1848	0x067a	0x03e1	0xffb5
288	0xf961	0xee68	0xf70f	0xf008	0xe664	0xbf3f
289	0x1298	0xfc84	0xd56a	0x1974	0x5e87	0xe885
290	0xff03	0x03e8	0x003f	0xffaf	0xff8d	0xfe82
291	0xfacb	0x5ea0	0xfd46	0xedc5	0xf50f	0xb538
292	0xfc94	0x8f3e	0xaa8f	0x3185	0xe738	0x0ca3
293	0x41cf	0x5299	0x99c4	0xf391	0xfe74	0x00e6
294	0x4778	0xe192	0xcdc7	0xfd59	0xfa3f	0x0005
295	0xd708	0x2ca5	0x64cd	0xfb9e	0x0579	0xfe4a
296	0x0ec6	0xe2fb	0x6860	0x449f	0x4b39	0x30fe
297	0x18bc	0xfd16	0x31f5	0x2464	0xa7f2	0xeb16
298	0x0d5a	0xa738	0x6962	0x477f	0x0434	0x03bc
299	0x954d	0xf454	0x0398	0x00eb	0x08b9	0x0051
300	0x1837	0x14b0	0x3edd	0x39b0	0xdf13	0xfba8
301	0xe6e0	0x4b2c	0x26c1	0xf34b	0x04fe	0xfc46
302	0x5e95	0x0801	0xa66d	0x0a19	0xf696	0xef88
303	0x2446	0x37ca	0xb2e9	0xf06f	0xf6d8	0x0404
304	0xb160	0x4649	0xdb0e	0x59e4	0xbda9	0x21b1
305	0xe510	0xaf06	0x0eb2	0x4407	0x5745	0x4a50
306	0x034a	0x5e75	0x61e6	0xe931	0xffb2	0x03a9
307	0xfd93	0x4d0a	0xa174	0xf856	0xc5fa	0xffc8
308	0x58ee	0xec01	0x43d5	0x5d3c	0xb3e8	0xe662
309	0xf792	0x4452	0xac45	0x0d0c	0xcded	0xb0b9
310	0xda6b	0x43ad	0x02cb	0x08d9	0xefe5	0xfe14
311	0x23c4	0x3293	0x6aa7	0xad49	0xe848	0xdb0f
312	0xcc94	0xa51b	0xc94a	0xefa8	0x1b42	0x0002
313	0x03aa	0xcbbb	0x0dc0	0xa117	0x5976	0x4c85
314	0xecd1	0xb2c2	0x4d34	0xdba2	0xce96	0x0eeb
315	0xeaaa	0xef67	0xe4b5	0xe78c	0xc989	0x9dc2
316	0x247d	0x2881	0xc9da	0xe5d0	0x581c	0xfdf9
317	0x19fb	0x4950	0xed09	0x311a	0x398a	0xd81f
318	0xfcc9	0x46c7	0x018e	0xf9d2	0xff8c	0xfe95
319	0xe4e9	0xce6a	0x9118	0x2168	0x1b31	0xff11
320	0xf5d6	0xeda0	0xfc03	0x07df	0x1409	0x5c76
321	0xcef1	0xe002	0x9e3c	0x4870	0x3763	0x067f
322	0x0ee5	0x522c	0xda6c	0xec45	0xf8f8	0xfbc1
323	0xa9d7	0x4123	0x3a70	0x24f3	0x0ae2	0x425f
324	0x9a48	0xb48a	0xe6f2	0x0450	0x16a6	0xb989
325	0xf937	0x60f9	0x28b1	0xd4b1	0x0380	0xeb67
326	0xf8c1	0x2d8d	0xf50d	0x60e9	0xac45	0xb2b0
327	0xa44f	0xbea7	0xe96a	0x160b	0x0a4c	0x134c
328	0xf944	0x1124	0x97cf	0xca81	0x294a	0x9ad9
329	0x3bfe	0xb3d8	0x6682	0xb7c3	0x06c8	0x1f76
330	0x1634	0x519a	0x0ffb	0xb564	0xc704	0xd71c
331	0x436c	0xc05d	0x3a0b	0xbad1	0xb51a	0x3093
332	0x95cf	0xcee3	0x1a57	0xfdce	0x03d0	0xfeff
333	0x306b	0xde56	0xa918	0xb27d	0x2b05	0x1e52
334	0x0ed7	0x2e4d	0x941a	0xdee7	0x0441	0xfa29
335	0x102d	0xf77a	0x97a0	0xfd21	0xfcfa	0x05bd
336	0x0c35	0x35c2	0x11fe	0x7249	0x4953	0xd91a
337	0xbbc7	0xdb1b	0xbb66	0xf61e	0xe6dd	0xf172
338	0xf932	0x10ff	0xe547	0xb2c3	0x259b	0xd662
339	0x1c53	0x0dc5	0x2a53	0x15e1	0x626e	0xa4cc
340	0xd7c4	0xba5a	0x0277	0x2d78	0x07fc	0xae72
341	0xfc97	0xdeca	0xfbd9	0xc2c6	0xd63b	0x3a56
342	0xc1ab	0x6de9	0x1494	0x01dd	0xfbe3	0x0486
343	0xfa29	0xdd92	0xe97c	0x9e7b	0x6584	0x1ee3
344	0xfbf2	0xff8e	0xf6fc	0xfad9	0xe6b0	0x05c0

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
345	0x131f	0xba17	0x9b06	0x14b5	0xff44	0x062d
346	0x0c80	0x4349	0x10fa	0x5655	0xb791	0x560c
347	0xd7f6	0x0221	0xd54c	0x08e4	0x925a	0x1fb6
348	0x3bef	0x0919	0x2464	0x5039	0x3a3c	0x521d
349	0x18b9	0x17f2	0xa044	0x0345	0xde43	0xe92c
350	0x1cda	0xfe0b	0x2907	0x4ea3	0x2cab	0xed6d
351	0xf547	0x5e6e	0xdbc6	0x3ba9	0xdf3b	0xe935
352	0x0bb0	0xf4d0	0x17a0	0xe2cf	0x2da7	0xb1e4
353	0xfc8d	0xd14e	0xd908	0xaabb	0xeeac	0x95d6
354	0x0d82	0x4caa	0x0500	0x0a25	0x4d89	0x1487
355	0xeb3d	0x4abd	0xc74a	0xdd0e	0x35b5	0xfab8
356	0x48d2	0x44f7	0x2af9	0x1aa1	0xb80e	0x18c0
357	0xf95f	0x08c4	0xede0	0x0f6c	0xcda6	0xeb67
358	0x4fcc	0x292e	0x104a	0xfc0c	0x4bef	0x54bb
359	0xf481	0xb2e9	0xef90	0x0528	0x038d	0xdd3f
360	0x2487	0xe07e	0xf5c6	0xcd7b	0x67d6	0x0db3
361	0x25e9	0xa79c	0x2077	0x1fe7	0xcc13	0x15e8
362	0x0c96	0x0ea5	0xfa1c	0x00a5	0xffcc	0xff3c
363	0x0066	0xa728	0xdd80	0x0387	0xd363	0xc6ba
364	0xff88	0x176e	0x4d35	0x3459	0x0e2c	0x144d
365	0x2150	0x16c3	0xfbd6	0x0306	0xffd9	0xff5a
366	0x24c3	0xdafc	0x256d	0xcd34	0x5f88	0x6144
367	0x45d6	0x08bb	0xab79	0x4ffe	0x126c	0xe3ea
368	0xf64e	0x2527	0x064b	0xaa49	0x3796	0xfaf7
369	0x2448	0xf70d	0x5aaf	0xf284	0xd5a6	0x000b
370	0x2518	0x0be1	0x140a	0xf0ce	0xad1d	0xa7c3
371	0x37b6	0xd992	0x4ee3	0x36c3	0x005b	0xbcd0
372	0xb761	0x03d4	0x0011	0x0335	0x0078	0xfdc2
373	0x2ffd	0xb4bb	0x35ae	0x3ff5	0xff5f	0x1789
374	0xf2dc	0x05fa	0xf05b	0x0996	0xd588	0xa2e1
375	0x0069	0x13dd	0xfefc	0x169e	0xfdb4	0x4ae2
376	0x1019	0x1049	0x347f	0x3934	0x51a3	0x1d0a
377	0xff51	0x332d	0xf188	0x5ac1	0x0f43	0x277a
378	0xe82b	0x5bab	0x1454	0xfac3	0x063f	0x3376
379	0xf36f	0xf25a	0x3b0d	0xdf3d	0xd20e	0xed72
380	0x047a	0x1243	0xb44e	0x3a45	0xec1d	0x00f9
381	0xabfe	0x2798	0xbfa7	0xcc07	0x47ce	0xde67
382	0x0274	0x098f	0x0d10	0x0c3a	0xec05	0x0077
383	0x45ec	0xa86a	0xbb1f	0x55cf	0xc05b	0xe204
384	0x41df	0x5e96	0x15ec	0xf0ee	0xfcd7	0x0eee
385	0xf70d	0x276b	0xf6c8	0x9deb	0xfb36	0x0138
386	0x0b8d	0x2bf8	0x6879	0xcc2e	0xf281	0xfb98
387	0xb2ce	0xf56c	0x11fc	0x18d3	0x0666	0x639d
388	0xb377	0xe1b7	0x0c57	0xffab	0xfe17	0xf8c1
389	0x032e	0x30de	0x4a85	0xedb7	0xf5ce	0xfa3e
390	0xa490	0xb5ad	0x1fc9	0x4da6	0x1ee8	0xfee6
391	0x0347	0xb33c	0x2e97	0x6a8e	0xf375	0x08da
392	0x0fb4	0xfbaa	0x2022	0xfb06	0x51ba	0x61e4
393	0x67d0	0x0145	0xde0b	0xff18	0xf756	0xfd45
394	0xd3e3	0xef98	0x070d	0xe5ef	0xa664	0xfac5
395	0xf82b	0xc1f2	0xfbe9	0x93d9	0xcc4d	0x3822
396	0xa9c7	0x079d	0x3377	0xc2d8	0xf8ca	0x1f77
397	0x0bdf	0x2ef9	0x1bdc	0x9fc8	0x019d	0xf6d5
398	0xa210	0xff32	0x30ab	0xe602	0xfe5f	0xd895
399	0x4703	0xa378	0xafdd	0xbff4	0x1c3e	0x02fb
400	0x161b	0xec23	0x3636	0xa34f	0xd4bb	0xb37d
401	0x2c4c	0x01f5	0x61d0	0x1dc0	0xb336	0x0645
402	0x97e6	0x22ae	0x2930	0x01a1	0x0513	0x0105
403	0x387c	0x2c69	0xf341	0x2706	0x2002	0x46bf
404	0x054b	0xae9a	0xdc14	0xc144	0xde91	0xfd26
405	0xf882	0xae37	0xb88b	0xf663	0xf5a5	0x10dc
406	0xf506	0x5fc9	0xd50c	0x9b87	0x0134	0xfb2e

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
407	0xdc8d	0xbc80	0xf8d7	0x8d62	0xa16b	0xbf09
408	0xf4e5	0x27b5	0xeb25	0xf4b8	0x5562	0xaca4
409	0xc228	0x3a01	0xa31c	0x1440	0x2781	0xaf61
410	0xb3b1	0xd39f	0x20dd	0x05ce	0xa396	0xe981
411	0xe2a8	0x0403	0xaec6	0x35a4	0x4db4	0xaa52
412	0xd09c	0xe492	0xb519	0xdd78	0x566d	0xbf96
413	0x0791	0x145a	0xe752	0xa314	0x3355	0x2b4a
414	0xff33	0x1794	0xfe84	0x21d2	0xff17	0x6d74
415	0xea6d	0x1d35	0x1dd3	0x5c2b	0x2623	0xf5e2
416	0x549a	0x9167	0xf3f2	0xfed4	0xfbf8	0x06d0
417	0xa8b0	0x4106	0x00d0	0x1a09	0xbc08	0xf42c
418	0x4832	0x2478	0xf54f	0xb454	0x0197	0xeedb
419	0xeccf	0xbc26	0x4983	0xbb0a	0x3468	0x3b80
420	0x1e45	0x18e0	0x5a5f	0xb902	0x1da0	0xef68
421	0xfa2f	0xe685	0x024a	0xd853	0x3a74	0x63e0
422	0x0f04	0xe7f4	0x1321	0xcd0b	0xa802	0x160f
423	0xded5	0xf7c7	0x9f3a	0x0389	0xdb92	0x05b0
424	0xf420	0xfa3c	0x048e	0xeeb4	0x2be4	0x23f4
425	0x0d45	0xfa55	0x351e	0xc21f	0x5fdc	0x16bb
426	0x2123	0xf44f	0x542b	0xbdec	0x1e3d	0x5dd2
427	0xc5ac	0xa332	0xeb2c	0xe5f8	0xee6f	0x33d3
428	0x4bb3	0x3274	0xf7a2	0xfd1f	0x526c	0xa9ab
429	0x0d41	0xedeb	0x1667	0xb61f	0xe4c7	0x0a7f
430	0x047c	0xc0ed	0xac47	0x9241	0xfd8a	0xc78f
431	0x1c84	0x02a0	0x4862	0xbbd4	0xd85b	0x015f
432	0x2c5c	0xd522	0x433c	0x1210	0x0091	0x457f
433	0xfd39	0xf269	0xf742	0x3e0f	0x07eb	0x0000
434	0x9270	0x0702	0xfdaf	0xf53a	0xaaa4	0x2d0f
435	0xb31d	0x1349	0x55f4	0x5413	0xf3b4	0x06fe
436	0x032d	0x2027	0x0a49	0x2ecd	0xf41d	0x56b9
437	0x22f8	0x9f48	0xfd4e	0x3a19	0xf6c2	0xeb04
438	0x20d6	0xeac1	0xfeee	0xfd7e	0xff6f	0x030a
439	0xe633	0x1c5a	0x512c	0xa42d	0xb73f	0x58fe
440	0xa690	0x9c70	0x2724	0xf9b2	0x05e4	0xfa8f
441	0x1db7	0x0197	0x9f9a	0xbfff	0xf8f4	0xeda5
442	0xd6a0	0xb53d	0x28de	0xf15d	0x2211	0xe4f9
443	0x32d2	0x14ac	0xe7aa	0xecec	0xae58	0xf8fb
444	0x41fb	0xca36	0xfe2f	0x4b8f	0xd60b	0xcd61
445	0x6269	0xc631	0xe9cf	0xfdf7	0xfebf	0xfb45
446	0x1b05	0xf3eb	0x4ed7	0x96e9	0xd106	0x050f
447	0x0131	0x07c8	0x4c01	0xfc27	0x0019	0xfdf7
448	0x1a33	0xf18e	0x20ad	0xde11	0x55a1	0x95e2
449	0x123c	0x176d	0x1bcd	0x2db0	0x5f51	0xd5d6
450	0x02e8	0xdb38	0x4db5	0x07ab	0x1ef2	0xd9a0
451	0x0d66	0x5322	0xf938	0x2a5c	0x2275	0x6987
452	0xdd93	0x05f1	0xa21a	0x0673	0x1e9e	0xfb48
453	0x0f47	0xd42b	0x0cc9	0xcf03	0x1c00	0x47e2
454	0x548a	0x239d	0xd2f0	0xeb78	0x1ba5	0x094e
455	0x0064	0x0ee9	0xe5c7	0x04dc	0x05ee	0xfebf
456	0x1f0a	0xb712	0x29ab	0xecfe	0x02d7	0x0308
457	0xc1f5	0xe02a	0xf7d9	0x58d3	0x061f	0xf266
458	0x111c	0xf551	0x2115	0xe48f	0xd360	0x0525
459	0x695a	0x1129	0x1df1	0x4499	0xfd36	0x028a
460	0xc0c1	0xfcbd	0x20ad	0x0703	0xc816	0x3fa9
461	0x1198	0xd8c0	0x1dee	0x97be	0xbbec	0x0a14
462	0x0030	0xf070	0x0234	0xe911	0x0a62	0xb71e
463	0x3123	0x9a60	0xc2e6	0x0a70	0xfabd	0xfc89
464	0xecaa	0x1070	0xe565	0x0a09	0xaf73	0xde2e
465	0xf8d4	0xc61e	0xea3d	0xa4e6	0xc630	0x650b
466	0x153a	0x9215	0xf6cb	0xf4bd	0xfdc6	0x097f
467	0x3328	0xf52d	0x61a2	0xcf30	0x9f6d	0xfbff
468	0xe9d4	0xef0d	0x0774	0x48c4	0xacb5	0x43d6

index	val[index][0] (16-bit two's complement)	val[index][1] (16-bit two's complement)	val[index][2] (16-bit two's complement)	val[index][3] (16-bit two's complement)	val[index][4] (16-bit two's complement)	val[index][5] (16-bit two's complement)
469	0x6c0c	0x9307	0xc3cf	0x059c	0xe438	0xf73f
470	0x1f53	0x0f07	0x5ff8	0xfe2b	0x25ca	0x29bb
471	0xfc79	0xd85b	0x0709	0xacf4	0x12bb	0xddd1
472	0x0462	0xda92	0x0a41	0x5907	0x03bc	0x0372
473	0x1ec4	0x4a83	0xd954	0xa136	0x1d48	0x243d
474	0x03d4	0x9774	0xeaf6	0x1514	0x043e	0x0670
475	0x70a6	0xfb0a	0xfe41	0x0005	0xfe53	0xffec
476	0xc44d	0x17f4	0x591c	0x04e4	0xd915	0x01ff
477	0x0353	0x1ef5	0xfe37	0xd04e	0x10a5	0x1d9b
478	0xee4e	0x2104	0xfb22	0x38a5	0x9e89	0xe980
479	0xba6a	0xd619	0x269f	0xa287	0xcb88	0x0756
480	0xc537	0x27b5	0x406b	0xc6f5	0xd240	0xad5c
481	0xf30b	0x0348	0xe9cd	0x578d	0x07ca	0x024a
482	0x5a76	0xe964	0xc53d	0xd77c	0xdbc9	0xcb2d
483	0xfcfb	0xdadb	0xf067	0xa138	0x210f	0x16ac
484	0xde9f	0xfd41	0xcf68	0xf06f	0x9dde	0x920d
485	0xbeed	0x3e81	0x0aba	0x064b	0x13e9	0xfbed
486	0x0029	0xe3f3	0x4dbf	0x7b43	0x8213	0x3667
487	0xe9e6	0x034d	0xce1a	0x1649	0x4137	0xffaa
488	0x14a2	0x3a1b	0x6992	0x5284	0x3da0	0xd713
489	0x3978	0x4cc0	0xd321	0xcbcf	0xb11c	0xc483
490	0x2195	0xdc4e	0xfd8e	0x2a8b	0xe881	0x18ca
491	0xfa30	0xfb08	0xfa39	0xfae9	0xf188	0xea93
492	0xf2d6	0x45cf	0xe634	0x6162	0x651e	0xf3c9
493	0x20e0	0x6c87	0xfa97	0x14e6	0xef5c	0x4e19
494	0x1638	0x016a	0x435e	0x0ee1	0xf352	0x0440
495	0xff97	0x8c59	0x0abb	0x3b77	0xff59	0x0e8a
496	0x0dae	0xf385	0x219a	0x1e5c	0xf9e2	0xfc6d
497	0xfe15	0x0cb9	0xf689	0x1592	0x507e	0xff9c
498	0xc984	0xd398	0xc3f1	0xaa96	0xdeb8	0x2fbd
499	0xfd98	0x0978	0xf819	0x112e	0xf123	0x1fac
500	0xe3dc	0x5233	0x52db	0xdb4d	0xb441	0x0380
501	0xe997	0xc4c8	0xacce	0x42aa	0xfc12	0xfe92
502	0x1875	0x0ca8	0xd15f	0xc0ab	0xc234	0x19b5
503	0xf3ad	0x60dc	0x0aad	0xfb17	0xfc95	0xf9c3
504	0xb00b	0x2b56	0x5e07	0xdce5	0x3738	0x08ac
505	0xc8e6	0x2eb7	0xa821	0x1027	0xfbe1	0xead4
506	0x0321	0xf5a1	0x003c	0xeb34	0xfcea	0x1731
507	0xe334	0xf91c	0xa85f	0x9a34	0x54cb	0x1052
508	0xe9ad	0xe608	0xc5c4	0x052d	0xa214	0x05d5
509	0xe878	0xcf38	0x5d7a	0x0b86	0x0641	0x0495
510	0x4a7b	0x44de	0x4609	0xd662	0x2ab0	0xeca2
511	0x0c9f	0xf32c	0x6ac8	0x104e	0xf96d	0x01f1

Annex F (normative): AC-3 and Enhanced AC-3 bit streams in the ISO Base Media File Format

F.0 Scope

The purpose of this annex is to define the necessary structures for the storage and identification of AC-3 and Enhanced AC-3 bit streams in a file format that is compliant with the ISO Base Media File Format. Examples of file formats that are derived from the ISO Base Media File Format include the MP4 file format and the 3GPP file format.

F.1 AC-3 and Enhanced AC-3 Track definition

In the terminology of the ISO Base Media File Format specification [i.9], AC-3 and Enhanced AC-3 tracks are both audio tracks. Therefore the following rules shall apply to the media box in the AC-3 or Enhanced AC-3 track:

- In the Handler Reference Box, the handler_type field shall be set to "soun".
- The Media Information Header Box shall contain a Sound Media Header Box.
- The Sample Description Box shall contain a box derived from AudioSampleEntry. For AC-3 tracks, this box is called AC3SampleEntry and is defined in clause F.3. For Enhanced AC-3 tracks this box is called EC3SampleEntry and is defined in clause F.5.
- The value of the timescale parameter in the Media Header Box, and the value of the SamplingRate parameter in the AC3SampleEntry Box or EC3SampleEntry Box shall be equal to the sample rate (in Hz) of the AC-3 or Enhanced AC-3 bit stream respectively.
- AC-3 bit streams shall always be identified by using the AC3SampleEntry Box. The use of the MP4AudioSampleEntry Box in combination with the Object Type Indicator value assigned for AC-3 to identify AC-3 bit streams is prohibited.
- Enhanced AC-3 bit streams shall always be identified by using the EC3SampleEntry Box. The use of the MP4AudioSampleEntry Box in combination with the Object Type Indicator value assigned for Enhanced AC-3 to identify Enhanced AC-3 bit streams is prohibited.
- The following bit stream parameters shall remain constant within an AC-3 bit stream that is identified by the AC3SampleEntry Box:
  - bsid
  - bsmod
  - acmod
  - lfeon
  - fscod
  - frmsizcod
- The following bit stream parameters shall remain constant within an Enhanced AC-3 bit stream that is identified by the EC3SampleEntry Box:
  - Number of independent substreams
  - Number of dependent substreams

- Within each independent substream:
  - bsid
  - bsmod
  - acmod
  - lfeon
  - fscod
- Within each dependent substream:
  - bsid
  - acmod
  - lfeon
  - fscod
  - chanmap

While the present document defines the AC-3 and Enhanced AC-3 bit stream syntax in big-endian byte order, it should be noted that little-endian byte order streams are valid and may be present within files that are compliant with the ISO Base Media File Format. The byte order of the AC-3 or Enhanced AC-3 bit stream can be determined from the order of the bytes in the syncword field at the start of each AC-3 or Enhanced AC-3 syncframe. A bit stream stored in big-endian byte order has a syncword value of 0x0B77, and a bit stream stored in little-endian byte order has a syncword value of 0x770B. It is strongly recommended that AC-3 and Enhanced AC-3 bit streams are stored within ISO Base Media Files in big-endian byte order.

F.2 AC-3 and Enhanced AC-3 Sample definition

AC-3 or Enhanced AC-3 Samples contain six audio blocks, equivalent in duration to 1 536 contiguous samples of PCM audio data. Consequently, the value of the sample_delta field in the Decoding Time to Sample Box shall be 1 536.

An AC-3 Sample shall be defined as follows:

- Exactly one AC-3 syncframe, as defined in clause 4.1 of the present document.

An Enhanced AC-3 Sample shall be defined as follows:

The number of Enhanced AC-3 syncframes required to deliver six blocks of audio data from every substream present in the Enhanced AC-3 bit stream, beginning with independent substream 0.

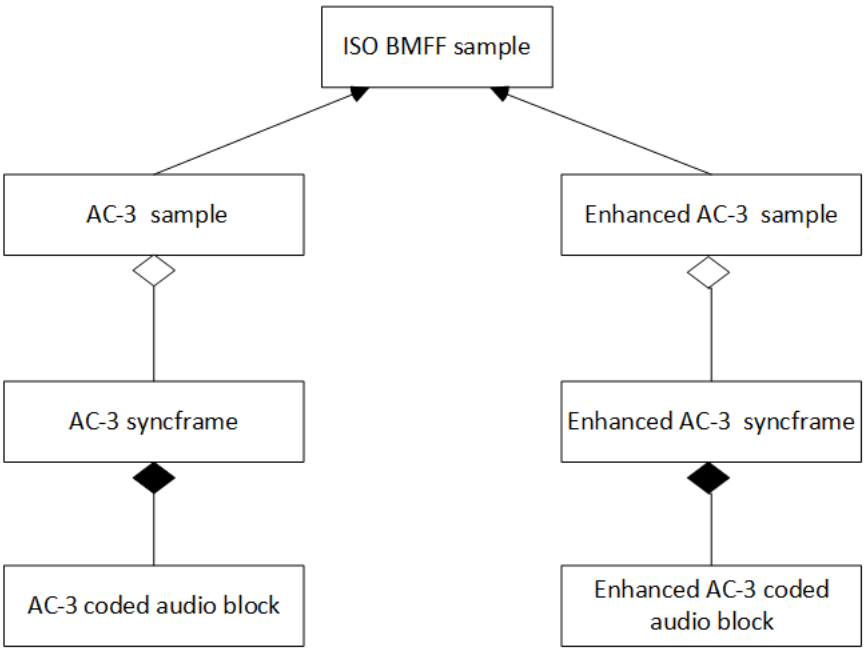

Figure F.2.1: Composition of ISOBMFF samples

How data is structured within an ISOBMFF sample depends on the configuration of the bitstream. Figures F.2.2 through F.2.4 provide different examples of ISOBMFF sample makeup.

Figure F.2.2 shows the construction of an ISOBMFF sample that contains a single Enhanced AC-3 Sample consisting of six audio blocks.

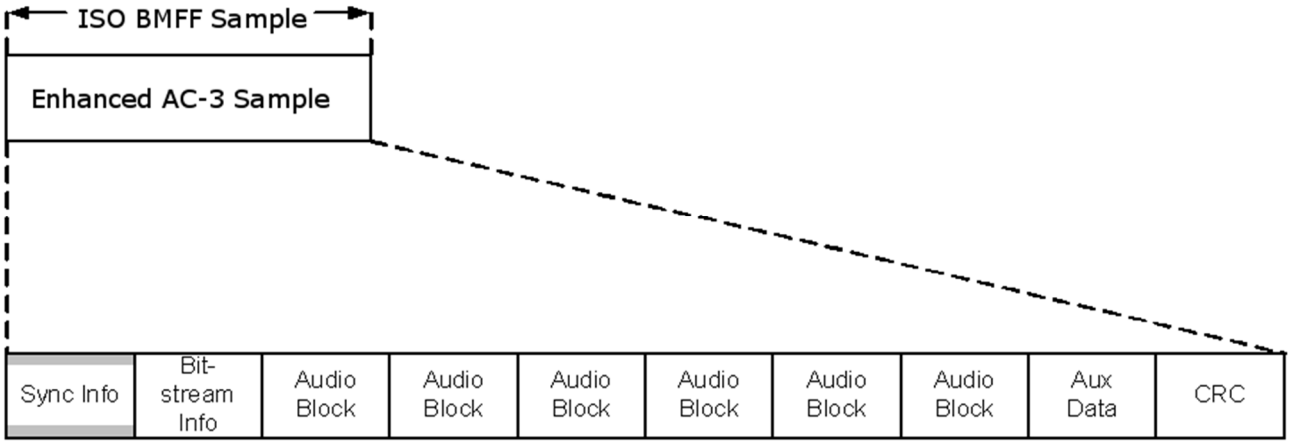

Figure F.2.2: ISOBMFF sample with a single substream with six blocks per frame

The six audio blocks represent 1, 536 PCM samples of audio from a single substream (substream 0).

Figure F.2.3 shows an ISOBMFF sample that contains a single Enhanced AC-3 Sample consisting of four frames.

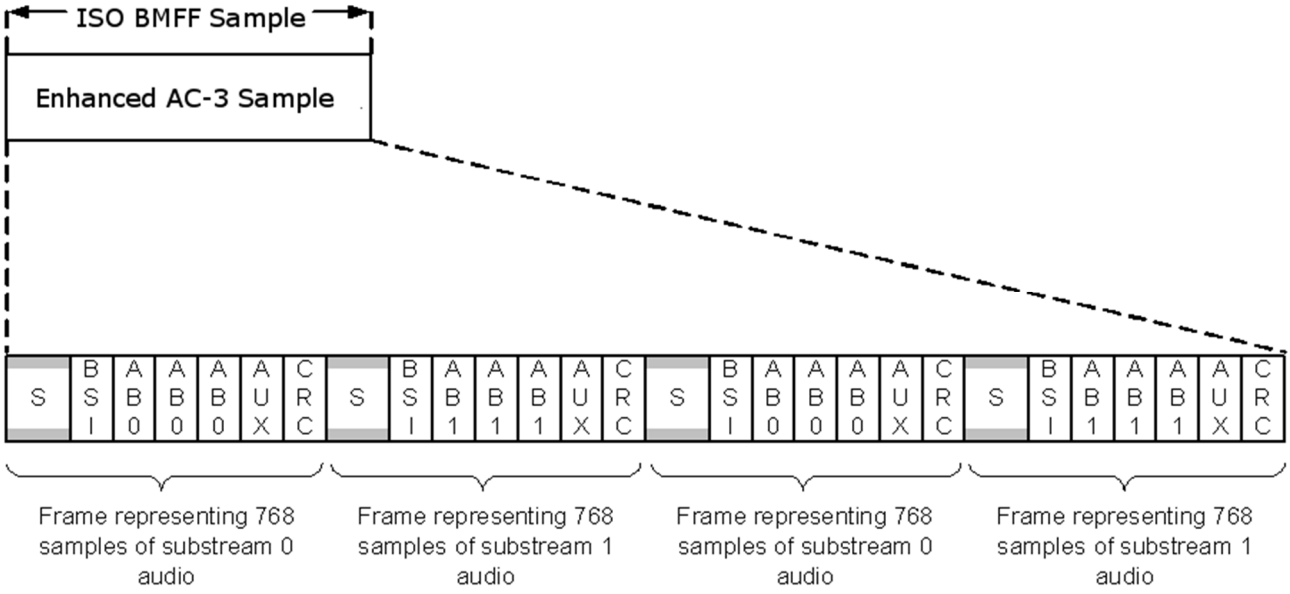

Figure F.2.3: ISOBMFF sample with two substreams with three blocks per frame

Each frame contains three audio blocks (denoted 'AB0' for substream 0 and 'AB1' for substream 1), each representing 256 PCM samples of audio from all channels in each substream.

Figure F.2.4 shows an ISOBMFF sample that contains a single Enhanced AC-3 Sample consisting of six frames. Each frame contains one audio block, each representing 256 PCM samples of audio from every channel in the substream.

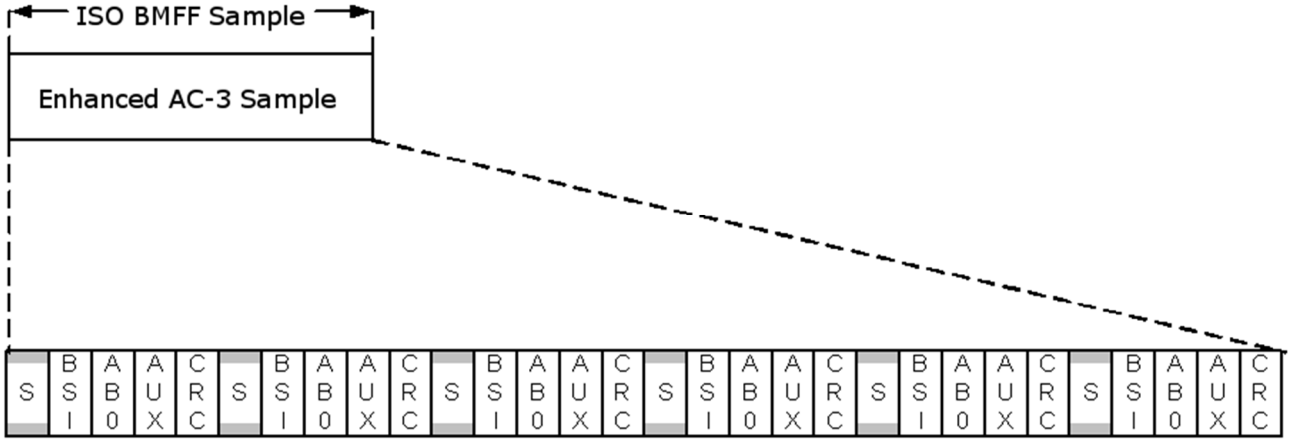

Figure F.2.4: ISOBMFF sample with a single substream with one block per frame

AC-3 and Enhanced AC-3 Samples shall be byte-aligned. If necessary, up to 7 zero-valued padding bits shall be added to the end of an AC-3 or Enhanced AC-3 Sample to achieve byte-alignment. The Padding Bits Box (defined in section 8.23 of ISO/IEC 14496-12 [i.9] need not be used to record padding bits that are added to a Sample to align its size to the nearest byte boundary.

F.3 AC3SampleEntry Box

F.3.0 Introduction

The bitstream format does not require external metadata to set up the decoder, as it is fully contained in that regard. Descriptor data is present, however, to provide information to the system without requiring access to the elementary stream, as the bitstream can be encrypted.

F.3.1 Syntax

Syntax	No. of bits	Identifier	Value
AC3SampleEntry()			
{			
BoxHeader.Size	32	uimsbf	
BoxHeader.Type	32	uimsbf	
Reserved [6]	8	uimsbf	0
Data-reference-index	16	uimsbf	
Reserved [2]	32	uimsbf	0
ChannelCount	16	uimsbf	2
SampleSize	16	uimsbf	16
Reserved	32	uimsbf	
SamplingRate	16	uimsbf	0
Reserved	16	uimsbf	0
AC3SpecificBox			
}			

F.3.2 Semantics

The layout of the AC3SampleEntry box is identical to that of AudioSampleEntry defined in ISO/IEC 14496-12 [i.9] (including the reserved fields and their values), except that AC3SampleEntry ends with a box containing AC-3 bit stream information called AC3SpecificBox. The AC3SpecificBox field structure for AC-3 is defined in clause F.4.

For unencrypted tracks, The value of the BoxHeader.Type field shall be set to 0x61632d33 ('ac-3'). For encrypted tracks, the value should be set according to the encryption scheme.

The values of the ChannelCount and SampleSize fields within the AC3SampleEntry Box shall be ignored.

F.4 AC3SpecificBox

F.4.1 Syntax

Syntax	No. of bits	Identifier
AC3SpecificBox ()		
{		
BoxHeader.Size	32	uimsbf
BoxHeader.Type	32	uimsbf
fscod	2	uimsbf
bsid	5	uimsbf
bsmod	3	uimsbf
acmod	3	uimbsf
lfeon	1	bslbf
bit_rate_code	5	uimsbf
reserved	5	uimbsf
}		

F.4.2 Semantics

F.4.2.1 BoxHeader.Type - 32 bits

The value of the 32-bit BoxHeader.Type field shall be set to 0x64616333 ('dac3').

F.4.2.2 fscod - 2 bits

The 2-bit fscod field shall have the same meaning and shall be set to the same value as the fscod field in the AC-3 bit stream.

F.4.2.3 bsid - 5 bits

The 5-bit bsid field shall have the same meaning and shall be set to the same value as the bsid field in the AC-3 bit stream.

F.4.2.4 bsmod - 3 bits

The 3-bit bsmod field shall have the same meaning and shall be set to the same value as the bsmod field in the AC-3 bit stream.

F.4.2.5 acmod - 3 bits

The 3-bit acmod field shall have the same meaning and shall be set to the same value as the acmod field in the AC-3 bit stream.

F.4.2.6 lfeon - 1 bit

The 1-bit lfeon field has the same meaning and is set to the same value as the lfeon field in the AC-3 bit stream.

F.4.2.7 bit_rate_code - 5 bits

The 5-bit bit_rate_code field shall indicate the data rate of the AC-3 bit stream in kbit/s, as shown in Table F.4.1. The value of this field shall be derived from the value of the frmsizcod parameter (see Table 4.13).

Table F.4.1: bit_rate_code

Bit_rate_code	Nominal bit rate (kbit/s)
00000	32
00001	40
00010	48
00011	56
00100	64
00101	80
00110	96
00111	112
01000	128
01001	160
01010	192
01011	224
01100	256
01101	320
01110	384
01111	448
10000	512
10001	576
10010	640

F.4.2.8 reserved - 5 bits

These bits are reserved, and shall be set to 0.

F.5 EC3SampleEntry Box

F.5.1 Syntax

Syntax	No. of bits	Identifier	Value
EC3SampleEntry()			
{			
BoxHeader.Size	32	uimsbf	
BoxHeader.Type	32	uimsbf	
Reserved [6]	8	uimsbf	0
Data-reference-index	16	uimsbf	
Reserved [2]	32	uimsbf	0
ChannelCount	16	uimsbf	2
SampleSize	16	uimsbf	16
Reserved	32	uimsbf	0
SamplingRate	16	uimsbf	
Reserved	16	uimsbf	0
EC3SpecificBox			
}			

F.5.2 Semantics

The layout of the EC3SampleEntry box is identical to that of AudioSampleEntry defined in ISO/IEC 14496-12 [i.9] (including the reserved fields and their values), except that EC3SampleEntry ends with a box containing Enhanced AC-3 bit stream information called EC3SpecificBox. The EC3SpecificBox field structure for Enhanced AC-3 is defined in clause F.6.

For unencrypted tracks, the value of the BoxHeader.Type field shall be set to 0x65632d33 ('ec-3'). For encrypted tracks, the value should be set according to the encryption scheme.

The values of the ChannelCount and SampleSize fields within the EC3SampleEntry Box shall be ignored.

F.6 EC3SpecificBox

F.6.1 Syntax

Syntax	No. of bits	Identifier
EC3SpecificBox()		
{		
BoxHeader.Size	32	uimsbf
BoxHeader.Type	32	uimsbf
data_rate	13	uimsbf
num_ind_sub	3	uimsbf
for(i = 0; i < num_ind_sub + 1; i++)		
{		
fscod	2	uimsbf
bsid	5	uimsbf
reserved	1	bslbf
asvc	1	bslbf
bsmod	3	uimsbf
acmod	3	uimsbf
lfeon	1	bslbf
reserved	3	uimbsf
num_dep_sub	4	uimsbf
if num_dep_sub > 0		
{		
chan_loc	9	uimsbf

Syntax	No. of bits	Identifier

}
else
{
reserved .....
	1	bslbf

} } **reserved**	variable	bslbf
}		

F.6.2 Semantics

F.6.2.1 BoxHeader.Type - 32 bits

The value of the 32-bit BoxHeader.Type field shall be set to 0x64656333 ('dec3').

F.6.2.2 data_rate - 13 bits

The 13-bit data_rate field indicates the data rate (in kbps) of the entire bitstream. The value is the sum of the data rates of all the substreams. When a bitstream uses variable data-rate encoding, data_rate indicates the maximum data rate of the bitstream.

The data rate of each substream is calculated using this equation:

$$\mathbf{data\_rate\_sub} = \frac{(\mathbf{frmsiz}+1)*\mathbf{fs}}{\mathbf{numblks}*16}$$

In this equation:

- frmsiz is the value of the frmsiz field in the frame as defined in clause E.1.3.1.3.
- fs is the sample frequency of the bitstream (in kHz) as defined in clause E.1.3.1.4. (The fs value is derived from the fscod parameter in the frame.)
- numblks is the number of audio blocks per frame as defined in clause E.1.3.1.5. (The numblks value is derived from the numblkscod parameter in the frame.)

F.6.2.3 num_ind_sub - 3 bits

The 3-bit num_ind_sub field shall indicate the number of independent substreams that are present in the Enhanced AC-3 bit stream. The value of this field shall be equal to the substreamID value of the last independent substream of the bit stream.

NOTE: This is the frame with a strmtyp value of 0 that precedes the frame with both a strmtyp value of 0 and a substreamid value of 0 (indicating that this frame belongs to the first independent substream of the bitstream).

F.6.2.4 fscod - 2 bits

The 2-bit fscod field shall have the same meaning and shall be set to the same value as the fscod field in the independent substream.

F.6.2.5 bsid - 5 bits

The 5-bit bsid field shall have the same meaning and shall be set to the same value as the bsid field in the independent substream.

F.6.2.6 reserved - 1 bit

This bit is reserved, and shall be set to 0.

F.6.2.7 asvc - 1 bit

The asvc field is used to signal whether the audio service carried by the independent substream (and dependent substreams associated with the independent substream, if present) is a main audio service intended to be presented on its own, or is an associated audio service that is intended to be decoded and mixed with a main audio service prior to presentation. The asvc parameter shall be set to '1' if the audio service is an associated audio service, and is set to '0' if the audio service is a main audio service.

Decoding devices that are not capable of simultaneously decoding two independent substreams and mixing the decoded audio together shall ignore independent substreams with a substreamID value greater than 0.

If the asvc parameter is set to '1' for independent substream 0, then the audio service carried in independent substream 0 is intended to be mixed with a main audio service carried in a separate Enhanced AC-3 bit stream stored in a different track in the ISO base media file. Decoding devices that are not capable of simultaneously decoding two Enhanced AC-3 bit streams and mixing the decoded audio together shall ignore Enhanced AC-3 audio tracks that have the asvc parameter set to '1' for independent substream 0.

F.6.2.8 bsmod - 3 bits

The 3-bit bsmod field has the same meaning and is set to the same value as the bsmod field in the independent substream. If the bsmod field is not present in the independent substream (i.e. when the infomdate field in the independent substream is set to '0'), this field shall be set to 0.

F.6.2.9 acmod - 3 bits

The 3-bit acmod field shall have the same meaning and shall be set to the same value as the acmod field in the independent substream.

F.6.2.10 lfeon - 1 bit

The 1-bit lfeon field shall have the same meaning and shall be set to the same value as the lfeon field in the independent substream.

F.6.2.11 reserved - 3 bits

These bits are reserved, and shall be set to 0.

F.6.2.12 num_dep_sub - 4 bits

The 4-bit num_dep_sub field shall be set to the value of the substreamid field found in the frame with a strmtyp value of 1 (that is, in the dependent substream) immediately preceding a frame with a strmtyp value of 0 (that is, in the independent substream).

F.6.2.13 chan_loc - 9 bits

The chan_loc field indicates channel locations (beyond the standard 5.1 channels) that are carried by dependent substreams associated with an independent substream. The contents of the chan_loc field are determined by parsing the chanmap bit field in every dependent substream associated with a particular independent substream, and setting the corresponding channel locations in the chan_loc field to a value of 1.

Because this field is used by the system only to indicate the unique channel locations present in the bitstream, it is not necessary to reflect replacement channels in this field. Therefore, duplicate channel locations in the chanmap field indicate replacement channels and can be ignored.

Table F.6.1: chan_loc field bit assignments

Bit	Location
0	Lc/Rc pair
1	Lrs/Rrs pair
2	Cs
3	Ts
4	Lsd/Rsd pair
5	Lw/Rw pair
6	Lvh/Rvh pair
7	Cvh
8 (MSB)	LFE2

F.6.2.14 reserved - variable

Additional reserved bytes may follow at the end of the EC3SpecificBox. The number of reserved bytes present is determined by subtracting the number of bytes used by the EC3SpecificBox fields specified above from the total length of the EC3SpecificBox as specified by the value of the BoxHeader.Size field of the EC3SpecificBox. Decoders conformant with this annex should ignore any reserved bytes that are present at the end of the EC3SpecificBox.

Annex G (normative): Enhanced AC-3 bit streams in the MPEG-2 multiplex

NOTE: Transport of Enhanced AC-3 bit streams in an MPEG-2 Transport Stream that conforms to System A is defined in ATSC Standard A/52:2012 Annex G [i.10]. Transport of Enhanced AC-3 bit streams in an MPEG-2 Transport Stream that conforms to System B is defined in ETSI TS 101 154 [i.6] and ETSI EN 300 468 [i.7].

Annex H (normative): Extensible format for delivery of metadata

H.0 Scope

This annex defines an extensible metadata delivery format (EMDF) to enable delivery of metadata that is not natively supported by the AC-3 and Enhanced AC-3 bit stream syntax. Extensibility is enabled through the use of versioning and identifiers for specific payload types. The metadata payloads defined in this Annex may be used to describe the nature or configuration of the audio programme being delivered in an AC-3 or Enhanced AC-3 bit stream, or may be used to control audio processing algorithms that are designed to further process the output of the AC-3 or Enhanced AC-3 decoder.

H.1 Overview

The extensible metadata delivery format consists of a container that contains one or more data payloads. Each payload is identified in the EMDF container using a unique payload identifier value to provide an unambiguous indication of the type of data present in the payload. The order of payloads within the EMDF container is undefined - payloads can be stored in any order, and a parser shall be able to parse the entire EMDF container to extract relevant payloads, and ignore payloads that are either not relevant or are unsupported. Protection data follows the final payload in the EMDF container. This protection data can be used by a decoding device to verify that the EMDF container and the payloads within the EMDF container are error-free. Figure H.1.1 illustrates the structure of the EMDF container and payloads within the EMDF container.

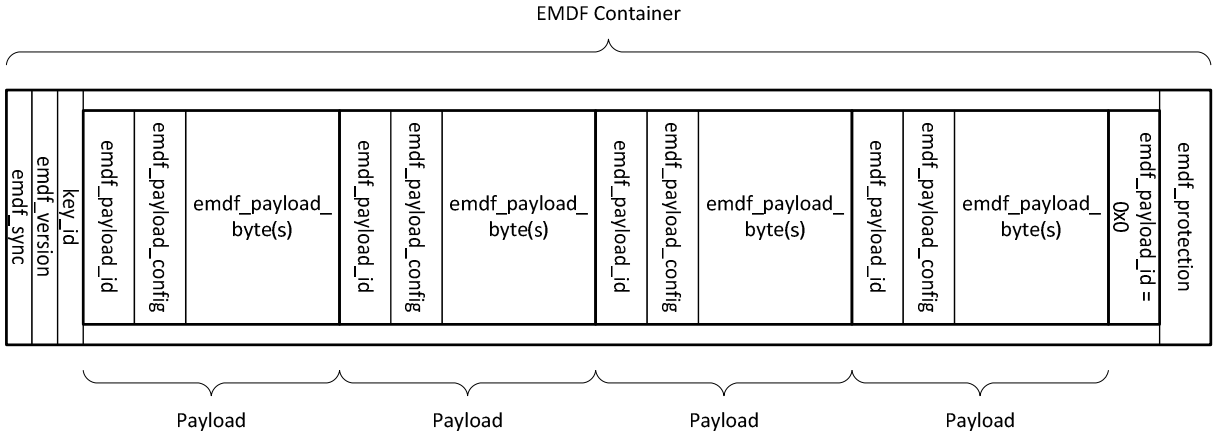

Figure H.1.1: Diagram of EMDF container structure

The EMDF container may be carried within an AC-3 or Enhanced AC-3 syncframe using one of the reserved data spaces supported by the AC-3 and Enhanced AC-3 bitstream syntaxes, for example the auxdata field located at the end of the syncframe (see clause 4.4.4), or the skip fields (skipfld, see clause (see clause 4.4.3.60)) present at the end of each audio block. Care should be taken to ensure that the addition of the EMDF container and the data payloads within the EMDF container do not adversely affect the sound quality of the AC-3 or Enhanced AC-3 bitstream.

H.2 Container syntax and semantics specification

H.2.1 Syntax

H.2.1.1 emdf_sync - EMDF container synchronization information

Syntax	No. of bits
emdf_sync()	
{	
syncword	16
emdf_container_length	16
}	

H.2.1.2 emdf_container - EMDF container

H.2.1.2.0 emdf_container

Syntax
emdf_container()
{
emdf_version 2
if (version == 3)
{
version += variable_bits(2) ..... variable_bits(2)
}
key_id 3
if (key_id == 7)
{
key_id += variable_bits(3) ..... variable_bits(3)
}
while (**emdf_payload_id**) != 0x0 5
{
if (emdf_payload_id == 0x1F)
{
emdf_payload_id += variable_bits(5) ..... variable_bits(5)
}
emdf_payload_config()
emdf_payload_size ..... variable_bits(8)
for (i = 0; i < payload_size; i++)
{
emdf_payload_byte 8
}
}
emdf_protection()
}

H.2.1.2.1 variable_bits - variable bit-length field format

Syntax
variable_bits (n_bits)
{
value = 0;
do {
value += **read** ..... n_bits
read_more 1
if (read_more)
{
value <<= n_bits;
value += (1<<n_bits);
}
} while (read_more);
return value;
}

H.2.1.3 emdf_payload_config - EMDF payload configuration

Syntax	No. of bits
emdf_payload_config()	
{	
smploffste	1
if(smploffste)	
{	
smploffst	11
reserved	1
}	
duratione	1
if(duratione) **{duration}**	variable_bits(11)
groupide	1
if (groupide) **{groupid}**	variable_bits(2)
codecdatae	1
if (codecdatae) **{reserved}**	8
discard_unknown_payload	1
if(discard_unknown_payload == 0)	
{	
if (smploffste == 0)	
{	
payload_frame_aligned	1
if (payload_frame_aligned)	
{	
create_duplicate	1
remove_duplicate	1
}	
}	
if (smpoffste == 1 || payload_frame_aligned == 1)	
{	
priority	5
proc_allowed	2
}	
}	
}	

H.2.1.4 emdf_protection - EMDF protection data

Syntax	No. of bits
emdf_protection()	
{	
protection_length_primary	2
protection_length_secondary	2
protection_bits_primary	8/32/128
protection_bits_secondary	0/8/32/128
}	

H.2.2 Semantics

H.2.2.1 emdf_sync - synchronization information

H.2.2.1.1 syncword - synchronization word - 16 bits

The 16-bit syncword field identifies the start of the EMDF container. The value of the syncword shall be set to 0x5838, or 0101 1000 0011 1000.

H.2.2.1.2 emdf_container_length - EMDF container length - 16 bits

The 16-bit emdf_container_length field specifies the length of the EMDF container in bytes. If the EMDF container does not occupy an integer number of bytes, the remaining bits needed to reach the byte boundary indicated by the value of the emdf_container_length shall be set to zero.

H.2.2.2 emdf_container - EMDF metadata container structure

H.2.2.2.1 variable_bits(n) - variable-bit field encoding

A number of fields within the EMDF container structure, payload_config, and payload data, are encoded using a method known as variable_bits. This method enables efficient coding of small field values with extensibility to be able to express arbitrarily large field values.

A field coded using variable_bits coding consists of one or more groups of n bits, with each group followed by a 1-bit read_more field. At a minimum, coding of n bits requires n + 1 bits to be transmitted. All fields coded using variable_bits shall be interpreted as unsigned integers. The number of groups required is dependent on the value to be encoded, and the value of n. The read_more bit that follows each group of n bits shall be set to 1 for all groups except for the last group. A read_more value of 0 indicates that no additional groups follow in the bitstream. Table H.2.1 shows examples of variable_bits encoding for different value ranges.

Table H.2.1: Examples of variable_bits encoding for different value ranges

Range of values to be encoded	Number of n-bit groups required to encode value	group_offset value	Encoded value stored in variable_bits field
$0 \le value \le (2^n - 1)$	1	0	= value
$2^n \le value \le (2^n + 2^{2n} - 1)$	2	2^n	= value - 2^n
$(2^n + 2^{2n}) \le value \le (2^n + 2^{2n} + 2^{3n} - 1)$	3	$2^n + 2^{2n}$	= value - $(2^n + 2^{2n})$

In general terms, the encoded value to be stored in the variable_bits field is always equal to the original value to be encoded, minus the group_offset value (as shown in Table H.2.1). Once the encoded value has been calculated, it is split into the required number of groups of n-bits, and a read_more bit is added each group to create the final variable_bits-encoded field to be written to the bit stream.

To decode a value from a field encoded using variable_bits, the encoded value is read from the bitstream by concatenating each group of n bits, ignoring the read_more bits between each group. To calculate the decoded value, the encoded value is then added to the group_offset value that corresponds to the number of groups of n-bits present in the variable_bits field. Table H.2.2 shows examples of variable_bits decoding for different numbers of groups of n-bits.

Table H.2.2: Examples of variable_bits decoding for different numbers of groups

Number of groups	group_offset value	Decoded value
1	0	= encoded value
2	2^n	= encoded value + (2^n)
3	$2^n + 2^{2n}$	= encoded value + $(2^n + 2^{2n})$
4	$2^n + 2^{2n} + 2^{3n}$	= encoded value + $(2^n + 2^{2n} + 2^{3n})$

H.2.2.2.2 emdf_version - EMDF syntax version - 2 bits / variable_bits(2)

The 2-bit emdf_version field shall indicate the syntax version that the EMDF container conforms to. For EMDF containers that conforms to the syntax defined in this Annex, the emdf_version field shall be set to '0'.

H.2.2.2.3 key_id - authentication key identification - 3 bits

The value of the 3-bit key_id field identifies the authentication key used to calculate the value of the protection_bits_primary and protection_bits_secondary fields of the EMDF protection data. The values of key_id are implementation dependent and are not defined in the present document.

H.2.2.2.4 emdf_payload_id - EMDF payload identification - 5 bits

The 5-bit emdf_payload_id field identifies the type of payload that follows in the EMDF container. The assignment of emdf_payload_id field values to specific payload types shall be specified in Table H.2.3. The final payload in the EMDF container shall have a emdf_payload_id value of 0x0, indicating that no additional payloads following the EMDF container. When the value of the emdf_payload_id field is 0x0, all fields in the emdf_payload_config (see clause H.2.2.3) shall be set to 0, and the value of the emdf_payload_size field (see clause H.2.2.2.5) shall be set to 0. The EMDF payload types defined in this annex are specified in clause H.3.

Table H.2.3: Defined EMDF payload types

Payload_ID value	EMDF Payload Type	Definition
0x0	Container End	clause H.2.2.2.4
0x1	Programme loudness data	clause H.3.1
0x2	Programme information	clause H.3.2
0x3	E-AC-3 substream structure	clause H.3.3
0x4	Dynamic range compression data for portable devices	clause H.3.4
0x5	Programme Language	clause H.3.5
0x6	External Data	clause H.3.6
0x7	Headphone rendering data	clause H.3.7

H.2.2.2.5 emdf_payload_size - size of EMDF payload - variable_bits(8)

The emdf_payload_size field indicates the size, in bytes, of the following payload. The value of the emdf_payload_size field shall be equal to the number of bytes in the following payload, excluding the emdf_payload_size field and all fields in emdf_payload_config().

H.2.2.3 emdf_payload_config - EMDF payload configuration data

H.2.2.3.1 smploffste - payload sample offset exists - 1 bit

If the 1-bit smploffste field is set to '1', the smploffst field shall follow in the bit stream. If the current payload applies to the first sample of the AC-3 or Enhanced AC-3 syncframe, this field shall be set to '0'.

H.2.2.3.2 smploffst - payload sample offset - 11 bits

The 11-bit smploffst field shall indicate the offset, in units of PCM audio samples, from the beginning of the AC-3 or Enhanced AC-3 syncframe to the first PCM audio sample that the data in the payload shall apply to. Note that this field may indicate a sample number that extends beyond the current AC-3 or Enhanced AC-3 syncframe.

H.2.2.3.3 duratione - payload duration data exists - 1 bit

If the 1-bit duratione field is set to '1', the duration field shall follow in the bit stream. If the current payload applies to all samples up to and including the final sample of the AC-3 or Enhanced AC-3 syncframe, this field shall be set to '0'.

H.2.2.3.4 duration - payload duration data - variable_bits(11)

The 11-bit duration field shall indicate the time period in units of PCM audio samples that the data in the payload applies to. The duration field shall be limited in size to a maximum of two groups of 11 bits (i.e. the read_more bit of the second group shall be set to '0'). Note that this field may indicate a sample number that extends beyond the current AC-3 or Enhanced AC-3 syncframe.

H.2.2.3.5 groupide - payload group ID data exists - 1 bit

If the 1-bit groupide field is set to '1', the groupid field shall follow in the bit stream.

H.2.2.3.6 groupid - payload group ID - variable_bits(2)

The groupid field provides a mechanism to indicate to a decoding device that specific payloads within the EMDF container are associated with one another and that the payload data are related. For payloads that are associated with each other, the groupid field for each payload shall be set to the same value.

H.2.2.3.7 codecdatae - codec-specific data exists - 1 bit

The 1-bit codecdatae field enables codec-specific data to be included in the payload configuration. For payload configuration data that conforms to this Annex, this field shall be set to '0'.

H.2.2.3.8 discard_unknown_payload - discard unknown payload during transcode - 1 bit

The 1-bit discard_unknown_payload field indicates whether the current payload is retained or discarded during a transcoding process if the transcoder does not recognize the specific emdf_payload_id value of the payload. If the discard_unknown_payload field is set to '1', the current payload shall be discarded during a transcoding process if the emdf_payload_id value is unrecognized. If the discard_unknown_payload field is set to '0', the current payload shall be retained during a transcoding process unless otherwise indicated by the proc_allowed field, as specified in clause H.2.2.3.13.

H.2.2.3.9 create_duplicate - create duplicate payload during transcode - 1 bit

In some transcode operations, the frame length of the output format may be shorter than the length of an AC-3 or Enhanced AC-3 syncframe. The 1-bit create_duplicate field shall be set to '1' to indicate that the payload shall be duplicated in all frames of the output format that correspond to the same period of time as the current AC-3 or Enhanced AC-3 syncframe. The create_duplicate field shall be set to '0' if the payload does not need to be repeated in all frames of the output format that correspond to the same period of time as the current AC-3 or Enhanced AC-3 syncframe.

H.2.2.3.10 remove_duplicate - remove duplicate payload during transcode - 1 bit

The 1-bit remove_duplicate field shall be set to '1' to indicate that payloads with the same emdf_payload_id value as the current payload need to be included only once in every output frame of the transcoding process when the output frame duration of a transcoding process is longer than the length of an AC-3 or Enhanced AC-3 syncframe. Subsequent payloads carried in AC-3 or Enhanced AC-3 syncframes that correspond to the same period of time as the current frame of the output format shall be discarded. The remove_duplicate field shall be set to '0' to indicate that payloads carried in all AC-3 or Enhanced AC-3 syncframes that correspond to the same time period as the current output frame of the transcoding process shall be included in this output frame.

H.2.2.3.11 payload_frame_aligned - payload to audio data frame alignment - 1 bit

The 1-bit payload_frame_aligned field is used to indicate how closely the payload data and the audio data within the AC-3 or Enhanced AC-3 syncframe need to be synchronized after decoding and/or during transcoding. If the payload_frame_aligned field is set to '0', the application of the payload data to the decoded audio from the AC-3 or Enhanced AC-3 syncframe is not required to be frame aligned. If the field is set to '1', the payload data shall be applied to the decoded audio from that AC-3 or Enhanced AC-3 syncframe.

H.2.2.3.12 priority - payload priority - 5 bit

The 5-bit priority field is used to indicate the priority of this payload in relation to other payloads in the EMDF container. This field may be used to indicate whether a payload can be discarded during a transcoding process to a lower data rate output format that is not able to support all payloads. A priority field value of '0'shall indicate that the payload has the highest priority of all payloads within the EMDF container. Higher values of the priority field shall indicate a lower payload priority. Multiple payloads within the EMDF container may have the same priority field value.

H.2.2.3.13 proc_allowed - retain payload during transcoding - 2 bits

The 2-bit proc_allowed field is used to indicate whether a payload should be retained or discarded during transcoding if processing is applied to the metadata and/or audio data decoded from the AC-3 or Enhanced AC-3 syncframe during transcoding. The value of the proc_allowed field shall be set as shown in Table H.2.4.

Table H.2.4: Meaning of proc_allowed values

proc_allowed field value	Meaning
00	payload may be retained only if no processing and no changes to any metadata in the AC-3 or Enhanced AC-3 syncframe occur during transcoding
01	payload may be retained if no processing other than sample rate conversion of the PCM audio is applied during transcoding
10	payload may be retained if one or more of the following processes, but no other processing, occur during transcoding: • sample rate conversion • momentary sound elements are mixed with the decoded PCM from the syncframe • informational metadata related to the audio data is modified • the channel configuration of the audio is changed
11	payload may be retained regardless of any processing performed during transcoding

H.2.2.4 emdf_protection - EMDF container protection data

H.2.2.4.1 protection_length_primary - length of protection_bits_primary field - 2 bits

The 2-bit protection_length_primary field is used to specify the length of the protection_bits_primary field. The value of the protection_length_primary field shall be set as shown in Table H.2.5.

Table H.2.5: Meaning of protection_length_primary field values

protection_length_primary field value	length of protection_bits_primary field
00	Reserved
01	8 bits
10	32 bits
11	128 bits

H.2.2.4.2 protection_length_secondary - length of protection_bits_secondary field - 2 bits

The 2-bit protection_length_secondary field is used to specify the length of the protection_bits_secondary field. The value of the protection_length_secondary field shall be set as shown in Table H.2.6.

Table H.2.6: Meaning of protection_length_secondary field values

protection_length_secondary field value	length of protection_bits_secondary field
00	0 bits
01	8 bits
10	32 bits
11	128 bits

H.2.2.4.3 protection_bits_primary - primary EMDF container protection data - 8 to 128 bits

The protection_bits_primary field contains protection data that may be used to verify the integrity of the EMDF container and the payload data within the EMDF container. Calculation of the value of the protection_bits_primary field is implementation dependent and is not defined in the present document.

H.2.2.4.4 protection_bits_secondary - secondary EMDF container protection data - 0 to 128 bits

The protection_bits_secondary field, if present, contains additional optional protection data that may be used to check the integrity of the EMDF container and the payload data within the EMDF container. Calculation of the value of the protection_bits_secondary field is implementation dependent and is not defined in the present document.

H.3 EMDF payload types

H.3.1 Payload ID 0x1 - programme loudness data

H.3.1.1 Overview

The programme loudness data payload conveys data about programme loudness, including measured loudness values for the programme, and indication as to whether the loudness of the programme has been corrected prior to AC-3 or Enhanced AC-3 encoding. This data may be used to control the operation of loudness processing downstream of the AC-3 or Enhanced AC-3 decoder.

H.3.1.2 Syntax

Syntax	No. of bits
loudness_data()	
{	
version	2
if (version == 0x3)	
{	
version += **extended _version**	4
}	
dialchane	1
if (dialchane)**{dialchan}**	3
loudpractyp	4
if (loudpractyp != 0x0)	
{	
loudcorrdialgat	1
loudcorrtyp	1
}	
loudrelgate	1
if (loudrelgate) **{loudrelgat}**	7
loudspchgate	1
if (loudspchgate) **{loudspchgat}**	7
loudstrm3se	1
if (loudstrm3se) **{loudstrm3s}**	8
truepke	1
if (truepke) **{truepk}**	8
dmixloudoffste	1
if (dmixloudoffste) **{dmixloudoffst}**	5
prgmbndye	1
if (prgmbndye)	
{	
prgmbndy = 1	
prgmbndy_bit = 0	
while (prgmbndy_bit == 0)	
{	
prgmbndy <<=1	
prgmbndy_bit	1
}	
end_or_start	1
prgmbndyoffste	1
if (prgmbndyoffste) **{prgmbndyoffst}**	11
}	
fill bits = 0 /* up to byte boundary */	0 to 7
if (loudrelgate)	
{	
hrloudrelgate	1
if (hrloudrelgate) **{hrloudrelgat}**	3
}	
if (loudspchgate)	

Syntax	No. of bits
{	
hrloudspchgate	1
if (hrloudspchgate) **{hrloudspchgat}**	3
}	
if (loudstrm3se)	
{	
hrloudstrm3se	1
if (hrloudstrm3se) **{hrloudstrm3s}**	3
}	
if (truepke)	
{	
hrtruepke	1
if (hrtruepke) **{hrtruepk}**	3
}	
if (loudcorrdialgat) **{loudcorrdialgattyp}**	3
if (loudrelgate == 0)	
{	
extloudrelgate	1
if (extloudrelgate) **{extloudrelgat}**	11
}	
if (loudspchgate == 0)	
{	
extloudspchgate	1
if (extloudspchgate) **{extloudspchgat}**	11
}	
if (loudspchgate || extloudspchgate) **{loudspchdialgattyp}**	3
if (loudstrm3se == 0)	
{	
ext**loudstrm3se**	1
if (extloudstrm3se) **{extloudstrm3s}**	11
}	
if (truepke == 0)	
{	
exttruepke	1
if (exttruepke) **{exttruepk}**	11
}	
maxloudstrm3se	1
if (maxloudstrm3se) **{maxloudstrm3s}**	11
maxtruepke	1
if (maxtruepke) **{maxtruepk}**	11
lrae	1
if (lrae)	
{	
lra	10
lrapractyp	3
}	
loudmntrye	1
if (loudmntrye) **{loudmntry}**	11
maxloudmntrye	1
if (maxloudmntrye) **{maxloudmntry}**	11
fill bits = 0 /* up to byte boundary */	0 to 7
}	

H.3.1.3 Semantics

H.3.1.3.1 version - Version of loudness payload - 2 bits

The 2-bit version field shall indicate the version of the loudness payload. For loudness data payloads that conform to this Annex, the version field shall be set to '01', and the extended_version field shall not be present in the payload.

H.3.1.3.2 dialchane - Dialogue presence data exists

If the 1-bit dialchane field is set to '1', the 3-bit dialchan field shall follow in the payload.

H.3.1.3.3 dialchan - Dialogue present in channel - 3 bits

The 3-bit dialchan field shall be used to indicate whether the Left, Right and/or Center channels contain spoken dialogue. The bit allocation of the dialchan field is shown in Table H.3.1. Bit 0, which indicates the presence of the left channel, shall be stored in the most significant bit of the dialchan field. Each bit of the dialchan field shall be set to '1' if the corresponding channel contains spoken dialogue during the preceding 0,5 seconds of the programme.

Table H.3.1: channel assignment of dialchan bits

Bit	Channel
0	L
1	R
2	C

H.3.1.3.4 loudpractyp - Loudness measurement practice used - 4 bits

The 4-bit loudpractyp field shall be used to indicate which recommended practice for programme loudness measurement has been followed to generate the programme loudness data that follows in the payload and to calculate the value of the dialnorm field (see clause 4.4.2.8) in the AC-3 or Enhanced AC-3 bit stream. The interpretation of the loudpractyp field shall be as defined in Table H.3.2. If the loudpractyp field is set to any value other than '0000', the loudcorrdialgat and loudcorrtyp fields shall follow in the payload.

Table H.3.2: loudpractyp bit values

loudpractyp value	Meaning
0000	Programme loudness measurement method not indicated
0001	Programme loudness measured according to ATSC A/85 [i.13]
0010	Programme loudness measured according to EBU R128 [i.14]
0011	Programme loudness measured according to ARIB TR-B32 [i.15]
0100	Programme loudness measured according to FreeTV OP-59 [i.16]
0101 - 1101	Reserved
1110	Manual
1111	Consumer leveller

H.3.1.3.5 loudcorrdialgat - Dialogue-gated loudness correction applied - 1 bit

If the loudness of the programme has been corrected using dialogue gating, the value of the loudcorrdialgat field shall be set to '1'. Otherwise this field shall be set to '0'.

H.3.1.3.6 loudcorrtyp - Type of loudness correction applied to programme - 1 bit

If the loudness of the programme has been corrected with an infinite look-ahead (file-based) loudness correction process, the value of the loudcorrtyp field shall be set to '0'. If the loudness of the programme has been corrected using a combination of real-time loudness measurement and dynamic range compression, the value of this field shall be set to '1'.

H.3.1.3.7 loudrelgate - Relative gated loudness data (ITU) exists - 1 bit

If the 1-bit loudrelgate field is set to '1', the 7-bit ituloudrelgat field shall follow in the payload.

H.3.1.3.8 loudrelgat - Relative gated programme loudness (ITU) - 7 bits

The 7-bit loudrelgat field shall indicate the integrated loudness of the audio programme, measured according to Recommendation ITU-R BS.1770-3 [i.12] without any gain adjustments due to dialogue normalization and dynamic range compression being applied. The values of 0 to 127 shall be interpreted as -58 LUFS to +5,5 LUFS, in 0,5 LU steps. If the loudrelgat field is present in the payload and the loudcorrdialgat field is set to '0', the value of the dialnorm field (see clause 4.4.2.8) in the AC-3 or Enhanced AC-3 bit stream shall match the loudness indicated by the value of the loudrelgat field.

H.3.1.3.9 loudspchgate - Speech-gated loudness data (ITU) exists - 1 bit

If the 1-bit loudspchgate field is set to '1', the 7-bit loudspchgat field shall follow in the payload.

H.3.1.3.10 loudspchgat - Speech-gated programme loudness (ITU) - 7 bits

The 7-bit loudspchgat field shall indicate the integrated speech-gated loudness of the entire audio programme, measured according to Recommendation ITU-R BS.1770-3 [i.12], and without any gain adjustments due to dialogue normalization and dynamic range compression being applied. The values of 0 to 127 shall be interpreted as -58 to +5,5 LUFS, in 0,5 LU steps. If the loudspchgat field is present in the payload and the loudcorrdialgat field is set to '1', the value of the dialnorm field (see clause 4.4.2.8) in the AC-3 or Enhanced AC-3 bit stream shall match the loudness indicated by the value of the loudspchgat field.

H.3.1.3.11 loudstrm3se - Short-term (3 second) loudness data exists - 1 bit

If the 1-bit loudstrm3se field is set to '1', the 8-bit loudstrm3s field shall follow in the payload.

H.3.1.3.12 loudstrm3s - Short-term (3 second) programme loudness - 8 bits

The 7-bit loudstrm3s field shall indicate the ungated loudness of the preceding 3 seconds of the audio programme, measured according to Recommendation ITU-R BS.1771-1 [i.11] and without any gain adjustments due to dialogue normalization and dynamic range compression being applied. The values of 0 to 255 shall be interpreted as -116 LUFS to +11,5 LUFS in 0,5 LUFS steps.

H.3.1.3.13 truepke - True peak loudness data exists - 1 bit

If the 1-bit truepke field is set to '1', the 8-bit truepk field shall follow in the payload.

H.3.1.3.14 truepk - True peak value - 8 bits

The 8-bit truepk field shall indicate the true peak value measured between the preceding occurrence of the truepk or exttruepk field and the present occurrence of the truepk field in the bit stream. The measurement is according to Annex 2 of Recommendation ITU-R BS.1770-3 [i.12], without any gain adjustments due to dialogue normalization and dynamic range compression being applied. The values of 0 to 255 shall be interpreted as -116 dBTP to +11,5 dBTP in 0,5 dB steps.

H.3.1.3.15 dmixloudoffste - Downmix loudness offset data exists - 1 bit

If the 1-bit dmixloudoffste field is set to '1', the 5-bit dmixloudoffst field shall follow in the bit stream. The dmixloudoffste field may only be set to '1' when the acmod value of the AC-3 or Enhanced AC-3 bit stream is greater than 2.

H.3.1.3.16 dmixloudoffst - Downmix loudness offset - 5 bits

The 5-bit dmixloudoffst field shall be used to indicate whether there is a difference between the expected loudness of a 2-channel downmix output from an AC-3 or Enhanced AC-3 decoder, and the actual measured loudness, with all gain adjustments due to dialogue normalization, dynamic range compression and/or fixed attenuation to protect against downmix overload having been applied prior to measurement. For example, an Enhanced AC-3 decoder that is decoding a bit stream with a dialnorm field value of 24 (-24 dBFS) will apply 7 dB of attenuation during decoding to align the loudness of the decoder output to a reference level of -31 LUFS. -31 LUFS is therefore the expected loudness of the 2-channel downmix output. The dmixloudoffst field is used to indicate any deviation between this expected downmix output loudness and the measured downmix output loudness.

Values of 0 to 30 shall be interpreted as -7,5 LU to + 7,5 LU, in 0,5 LU steps. A value of 31 is reserved, and if present shall be interpreted as an offset of 0 LU. When the value of the dmixloudoffst field indicates a positive LU value, this shall indicate that the measured downmix loudness is louder than the expected downmix loudness by the indicated LU value. When a negative LU value is indicated, this shall indicate that the measured downmix loudness is quieter than the expected downmix loudness by the indicated LU value. Note that if a negative LU value is indicated, the system should ensure that any positive gain applied to the 2-channel downmix to compensate for the loudness offset does not introduce clipping of the audio signal.

H.3.1.3.17 prgmbndye - Programme boundary data exists - 1 bit

If the 1-bit prgmbndye field is set to '1', programme boundary data shall follow in the payload.

H.3.1.3.18 prgmbndy_bit - Programme boundary bit value - 1 bit

The multi-bit prgmbndy variable is delivered in the bitstream as a sequence of 1-bit prgmbndy_bit fields. The value of the prgmbndy variable shall be equal to the number of syncframes between the current syncframe and the syncframe that contains the boundary between two different programmes. This data may be used to determine when to begin and end the measurement of the loudness parameters specified in this payload.

The prgmbndy variable is limited to positive integer values that are a power of 2, and is stored in the bit stream as a bit-reversed integer. For example, if the sequence of prgmbndy_bit values read from the bit stream is '001b', the prgmbndy value shall be interpreted as '100b' (decimal value 8). To decode the prgmbndy variable value from the bit stream, the following process is followed:

- The prgmbndy variable is initialized to a value of '1'.
- The decoder reads the prgmbndy_bit field that follows a prgmbndye field value of '1'.
- If the value of the prgmbndy_bit field is '0', this bit is added to the prgmbndy variable as the least significant bit (i.e. a left shift is performed on the initial prgmbndy variable value of '1'), and the decoder repeats the reading process.
- Each subsequent prgmbndy_bit field with a value of '0' that is read from the bitstream is added to the prgmbndy variable as the least significant bit, increasing the size of the prgmbndy variable by one bit.
- The reading process is repeated until a prgmbndy_bit field with a value of '1' is read from the bit stream. A prgmbndy_bit field value of '1' shall indicate that all bits of the prgmbndy variable have been read from the bit stream.

The range of the prgmbndy variable shall be limited to a minimum value of '10' and a maximum value of '1000000000'.

H.3.1.3.19 end_or_start - Programme boundary indicates end or start of programme - 1 bit

If the 1-bit end_or_start field is set to '1', then the value of the prgmbndy variable shall indicate the number of syncframes between the current syncframe and the syncframe that contains the next programme boundary. If the end_or_start field is set to '0', then the value of the prgmbndy variable shall indicate the number of syncframes between the current syncframe and the syncframe that contains the previous programme boundary.

H.3.1.3.20 prgmbndyoffste - Programme boundary sample offset exists - 1 bit

If the 1-bit prgmbndyoffste field is set to '1', the 11-bit prgmbndyoffst field shall follow in the payload.

H.3.1.3.21 prgmbndyoffst - Programme boundary sample offset - 11 bits

The 11-bit prgmbndyoffst field shall indicate the offset in audio samples from the first sample of the syncframe indicated by the prgmbndy variable to the actual audio sample in that syncframe that corresponds to the programme boundary.

H.3.1.3.22 hrloudrelgate - High resolution relative gated loudness data (ITU) exists - 1 bit

If the 1-bit hrloudrelgate field is set to '1', the 3-bit hrloudrelgat field shall follow in the payload.

H.3.1.3.23 hrloudrelgat - High resolution relative gated programme loudness (ITU) - 3 bits

In applications where transmission of high resolution relative gated loudness data is required, the 3-bit hrloudrelgat field is used in conjunction with the loudrelgat field to allow the relative gated programme loudness to be specified in 0,1 LU steps by adding or subtracting up to 0,4 LU from the loudness value indicated by the loudrelgat field. The interpretation of the hrloudrelgat field shall be as indicated in Table H.3.3.

Table H.3.3: Meaning of hrloudrelgat values

hrloudrelgat value	Meaning
000	-0,4 LU
001	-0,3 LU
010	-0,2 LU
011	-0,1 LU
100	+0,1 LU
101	+0,2 LU
110	+0.3 LU
111	+0,4 LU

When the hrloudrelgat field is present in the payload, the integrated loudness of the audio programme, measured according to Recommendation ITU-R BS.1770-3 [i.12] without any gain adjustments due to dialogue normalization and dynamic range compression being applied, is equal to the LUFS value indicated by the loudrelgat field summed with the LU value indicated by the hrloudrelgat field.

H.3.1.3.24 hrloudspchgate - High resolution speech-gated loudness data (ITU) exists - 1 bit

If the 1-bit hrloudrelgat field is set to '1', the 3-bit hrloudrelgat field shall follow in the payload.

H.3.1.3.25 hrloudspchgat - High resolution speech-gated programme loudness (ITU) - 3 bits

In applications where transmission of high resolution speech-gated loudness data is required, the 3-bit hrloudspchgat field is used in conjunction with the loudspchgat field to allow the speech-gated programme loudness to be specified in 0,1 LU steps by adding or subtracting up to 0,4 LU from the loudness value indicated by the loudspchgat field. The hrloudspchgat field shall be coded and interpreted in the same way as the hrloudrelgat field, as specified in Table H.3.3.

When the hrloudspchgat field is present in the payload, the integrated speech-gated loudness of the audio programme, measured according to Recommendation ITU-R BS.1770-3 [i.12], and without any gain adjustments due to dialogue normalization and dynamic range compression being applied, is equal to the LUFS value indicated by the loudspchgat field summed with the LU value indicated by the hrloudspchgat field.

H.3.1.3.26 hrloudstrm3se - High resolution short-term (3 second) loudness data exists - 1 bit

If the 1-bit hrloudstrm3se field is set to '1', the 3-bit hrloudstrm3s field shall follow in the payload.

H.3.1.3.27 hrloudstrm3s - High resolution short-term (3 second) loudness - 3 bits

In applications where transmission of high resolution short-term (3second) loudness data is required, the 3-bit hrloudstrm3s field is used in conjunction with the loudstrm3s field to allow the short-term (3-second) loudness to be specified in 0,1 LU steps by adding or subtracting up to 0,4 LU from the loudness value indicated by the loudstrm3s field. The hrloudstrm3s field shall be coded and interpreted in the same way as the hrloudrelgat field, as specified in Table H.3.3.

When the hrloudstrm3s field is present in the payload, the ungated loudness of the preceding 3 seconds of the audio programme, measured according to Recommendation ITU-R BS.1771-1 [i.11] and without any gain adjustments due to dialogue normalization and dynamic range compression being applied, is equal to the LUFS value indicated by the loudstrm3s field summed with the LU value indicated by the hrloudstrm3s field.

H.3.1.3.28 hrtruepke - High resolution true peak loudness data exists - 1 bit

If the 1-bit hrtruepke field is set to '1', the 3-bit hrtruepk field shall follow in the payload.

H.3.1.3.29 hrtruepk - High resolution true peak value - 3 bits

In applications where transmission of high resolution true peak loudness data is required, the 3-bit hrtruepk field is used in conjunction with the truepk field to allow the true peak value to be specified in 0,1 dB steps by adding or subtracting up to 0,4 dB from the true peak value indicated by the truepk field. The hrtruepk field shall be coded and interpreted as specified in Table H.3.4.

Table H.3.4: Meaning of hrtruepk values

hrtruepk value	Meaning
000	-0,4 dB
001	-0,3 dB
010	-0,2 dB
011	-0,1 dB
100	+0,1 dB
101	+0,2 dB
110	+0.3 dB
111	+0,4 dB

When the hrtruepk field is present in the payload, the true peak value of the audio programme, measured since the previous occurrence of the truepk field in the bit stream and according to Annex 2 of Recommendation ITU-R BS.1770-3 [i.12], and without any gain adjustments due to dialogue normalization and dynamic range compression being applied, is equal to the dBTP value indicated by the truepk field summed with the dB value indicated by the hrtruepk field.

H.3.1.3.30 loudcorrdialgattyp - Dialog gating type used for loudness correction - 3 bits

This 3-bit loudcorrdialgattyp field shall indicate which dialog gating method has been utilized for loudness correction of the programme and for calculation of the value of the dialnorm field (see clause 4.4.2.8) in the AC-3 or Enhanced AC-3 bit stream. The loudcorrdialgattyp field shall be interpreted as shown in Table H.3.5.

Table H.3.5: Meaning of loudcorrdialgattyp values

loudcorrdialgattyp value	Meaning
000	Dialog gating method - Not Indicated
001	Dialog gating method - Automated Center or Left+Right Channel(s)
010	Dialog gating method - Automated Left, Center, and/or Right Channel(s)
011	Dialog gating method - Manual Selection
100 - 111	Reserved
NOTE 1: loudcorrdialgattyp methods '001' and '010' are described in [] For method '001', dialog gating is applied to the Center channel of a multichannel programme OR to the power sum of Left and Right channels in a stereo programme. For method '010', dialog gating is applied individually to all front (main) channels of the programme.	
NOTE 2: A value of '011' indicates that a manual process was utilized to select representative portions of the programme containing dialog for measurement.	

H.3.1.3.31 extloudrelgate - Extended range relative gated loudness data (ITU) exists - 1 bit

If the 1-bit extloudrelgate field is set to '1', the 11-bit extloudrelgat field shall follow in the payload.

H.3.1.3.32 extloudrelgat - Extended range relative gated loudness data (ITU) - 11 bits

If the measured relative gated loudness exceeds the range supported by the loudrelgat parameter specified in clause H.3.1.3.8, the 11-bit extloudrelgat field should be used instead of the loudrelgat field. The 11-bit extloudrelgat field shall indicate the integrated loudness of the audio programme, measured according to Recommendation ITU-R BS.1770 [i.12] and without any gain adjustments due to dialogue normalization and dynamic range compression being applied.

The extloudrelgat field shall be encoded according to the following equation:

$$\text{extloudrelgat} = (\text{loudness\_value\_relative\_gated\_LUFS} \times 10 + 0{,}5) + 1\,024$$

and shall be decoded according to the following equation:

$$\text{loudness\_value\_relative\_gated\_LUFS} = (\text{extloudrelgat} - 1\,024) / 10$$

If the extloudrelgat field is present in the payload, and the loudcorrdialgat field is set to '0', the value of the dialnorm field (see clause 4.4.2.8) in the AC-3 or Enhanced AC-3 bit stream shall match the loudness indicated by the value of the extloudrelgat field.

H.3.1.3.33 extloudspchgate - Extended range speech-gated loudness data (ITU) exists - 1 bit

If the 1-bit extloudspchgate field is set to '1', the 11-bit extloudspchgat field shall follow in the payload.

H.3.1.3.34 extloudspchgat - Extended range speech-gated loudness value - 11 bits

If the measured dialog-based loudness exceeds the range supported by the loudspchgat parameter specified in clause H.3.1.3.10, the 11-bit extloudspchgat field should be used instead of the loudspchgat field. The extloudspchgat field shall indicate the integrated dialog-based loudness of the entire audio programme, measured according to Recommendation ITU-R BS.1770-3 [i.12] with dialog-gating. The extloudspchgat value represents the dialog-based loudness without any gain adjustments due to dialogue normalization and dynamic range compression being applied.

The extloudspchgat field shall be encoded according to the following equation:

$$\text{extloudspchgat} = (\text{loudness\_value\_speech\_gated\_LUFS} \times 10 + 0{,}5) + 1\,024$$

and shall be decoded according to the following equation:

$$\text{loudness\_value\_speech\_data\_LUFS} = (\text{extloudspchgat} - 1\,024) / 10$$

If the extloudspchgat field is present in the payload and the loudcorrdialgat field is set to '1', the value of the dialnorm field (see clause 4.4.2.8) in the AC-3 or Enhanced AC-3 bit stream shall match the loudness indicated by the value of the extloudspchgat field.

H.3.1.3.35 loudspchdialgattyp - Dialog gating type used for loudness measurement - 3 bits

This 3-bit loudspchdialgattyp field indicates which dialog gating method has been utilized for speech-based loudness measurement of the programme. The loudspchdialgattyp field shall be interpreted in the same way as the loudcorrdialgattyp field, as specified in clause H.3.1.3.30.

Typically the value of the loudspchdialgattyp should be equal to the value of the loudcorrdialgattyp field. However, in some applications it may be desirable to re-measure the loudness of the programme using a different dialog gating method than that used for loudness correction of the programme. In this case, the loudspchdialgattyp parameter may be used to identify the type of dialog gating used for this additional measurement step. The loudspchgat or extloudspchgat field values calculated by this additional measurement step shall not be used to set the value of the dialnorm field (see clause 4.4.2.8) in the AC-3 or Enhanced AC-3 bit stream.

H.3.1.3.36 extloudstrm3se - Extended range short-term (3 second) programme loudness exists - 1 bit

If the 1-bit extloudstrm3se field is set to '1', the 11-bit extloudstrm3s field shall follow in the payload.

H.3.1.3.37 extloudstrm3s - Extended range short-term (3 second) programme loudness - 11 bits

If the measured short-term (3 second) loudness exceeds the range supported by the loudstrm3s field specified in clause H.3.1.3.12, the 11-bit extloudstrm3s field should be used instead of the loudstrm3s field. The extloudstrm3s field shall indicate the ungated loudness of the preceding 3 seconds of the audio programme, measured according to Recommendation ITU-R BS.1771-1 [i.11] and without any gain adjustments due to dialogue normalization and dynamic range compression being applied.

The extloudstrm3s field shall be encoded according to the following equation:

$$\text{extloudstrm3s} = \text{loudness\_value\_short\_term\_3s\_LUFS} \times 10 + 0{,}5) + 1\,024$$

and shall be decoded according to the following equation:

$$\text{loudness\_value\_short\_term\_3s\_LUFS} = (\text{extloudstrm3s} - 1\,024) / 10$$

H.3.1.3.38 exttruepkee - Extended true peak value exists - 1 bit

If the 1-bit exttruepke field is set to '1', the 11-bit exttruepke field shall follow in the payload.

H.3.1.3.39 exttruepk - Extended true peak - 11 bits

If the measured true peak value exceeds the range supported by the truepk field specified in clause H.3.1.3.14, the 11-bit exttruepk field should be used instead of the truepk field. The exttruepk field shall indicate the true peak value measured between the preceding occurrence of the truepk or exttruepk field and the present occurrence of the exttruepk field in the bit stream. The measurement is according to Annex 2 of Recommendation ITU-R BS.1770-3 [i.12], without any gain adjustments due to dialnorm and dynamic range compression being applied.

The exttruepk field shall be encoded according to the following equation:

$$\text{exttruepk} = (\text{true\_peak\_value\_dBTP} \times 10 + 0{,}5) + 1\,024$$

and shall be decoded according to the following equation:

$$\text{true\_peak\_value\_dBTP} = (\text{exttruepk} - 1\,024) / 10$$

H.3.1.3.40 maxloudstrm3se - Max short-term (3 second) loudness value exists - 1 bit

If the 1-bit maxloudstrm3se field is set to '1', the 11-bit maxloudstrm3s field shall follow in the bitstream.

H.3.1.3.41 maxloudstrm3s - Max short-term (3 second) loudness value - 11 bits

The 11-bit maxloudstrm3s field shall indicate the maximum short-term ungated loudness of the audio programme, measured according to Recommendation ITU-R BS.1771-1 [i.11] and without any gain adjustments due to dialogue normalization and dynamic range compression being applied.

The maxloudstrm3s field shall be encoded according to the following equation:

$$\text{maxloudstrm3s} = (\text{max\_loudness\_value\_short\_term\_3s\_LUFS} \times 10 + 0{,}5) + 1\,024$$

and shall be decoded according to the following equation:

$$\text{max\_loudness\_value\_short\_term\_3s\_LUFS} = (\text{maxloudstrm3s} - 1\,024) / 10$$

H.3.1.3.42 maxtruepke - Max true peak value exists - 1 bit

If the 1-bit maxtruepke field is set to '1', the 11-bit maxtruepk field shall follow in the bitstream.

H.3.1.3.43 maxtruepk - Max true peak value - 11 bits

The 11-bit maxtruepk field shall indicate the maximum true peak value of the programme measured according to Annex 2 of Recommendation ITU-R BS.1770-3 [i.12].

The maxtruepk field shall be encoded according to the following equation:

$$maxtruepk = (max\_true\_peak\_value\_dBTP) \times 10 + 0,5) + 1\,024$$

and shall be decoded according to the following equation:

$$max\_true\_peak\_value\_dBTP = (maxtruepk - 1\,024) / 10$$

H.3.1.3.44 lrae - Loudness range data exists - 1 bit

If the 1-bit lrae field is set to '1', the 10-bit lra field shall follow in the bitstream.

H.3.1.3.45 lra - Loudness range - 10 bits

The 10-bit lra field shall indicate the loudness range of the programme as specified in EBU Tech Document 3342 [i.19].

The lra field shall be encoded according to the following equation:

$$lra = (loudness\_range\_LU \times 10) + 0,5$$

and shall be decoded according to the following equation:

$$loudness\_range\_LU = lra / 10$$

H.3.1.3.46 lrapractyp - Loudness range dialog gating practice type - 3 bits

The 3-bit lrapractyp field shall indicate which method has been utilized to compute the loudness range of the audio programme, as shown in Table H.3.6.

Table H.3.6: Meaning of lrapractyp values

lrapractyp value	Meaning
000	Loudness range computed per EBU Tech 3342 [i.19] v1
001	Loudness range computed per EBU Tech 3342 [i.19] v2
010 - 111	Reserved

H.3.1.3.47 loudmntrye - Momentary loudness data exists - 1 bit

If the 1-bit loudmntrye field is set to '1', the 11-bit loudmntry field shall follow in the bitstream.

H.3.1.3.48 loudmntry - Momentary loudness data - 11 bits

The 11-bit loudmntry field shall indicate the momentary loudness of the programme measured since the preceding occurrence of the loudmntry field in the bitstream, measured according to ITU-R BS.1771-1 [i.11] (or EBU Tech 3341 [i.18]) without any gain adjustments due to dialogue normalization and dynamic range compression being applied.

The loudmntry field shall be encoded according to the following equation:

$$loudmntry = (momentary\_loudness\_LUFS \times 10 + 0,5) + 1\,024$$

and shall be decoded according to the following equation:

momentary_loudness_LUFS = (loudmntry - 1 024) / 10

H.3.1.3.49 maxloudmntrye - Maximum momentary loudness data exists - 1 bit

If the 1-bit maxloudmntrye field is set to '1', the 11-bit maxloudmntry field shall follow in the bitstream.

H.3.1.3.50 maxloudmntry - Maximum momentary loudness data - 11 bits

The 11-bit maxloudmntry field shall indicate the maximum momentary loudness of the programme measured as specified in Recommendation ITU-R BS.1771-1 [i.11] (or EBU Tech 3341 [i.18]) without any gain adjustments due to dialogue normalization and dynamic range compression being applied.

The maxloudmntry parameter shall be encoded according to the following equation:

maxloudmntry = (maximum_momentary_loudness_LUFS × 10 + 0,5) + 1 024

and shall be decoded according to the following equation:

maximum_momentary_loudness_LUFS = (maxloudmntry - 1 024) / 10

H.3.2 Payload ID 0x2 - programme information

H.3.2.1 Overview

The programme information payload conveys additional information about the audio programme that cannot be carried in other portions of the AC-3 or Enhanced AC-3 bit stream. This payload contains information about processing applied to the PCM audio prior to AC-3 or Enhanced AC-3 encoding, which portions of the audio programme bandwidth have been encoded using specific audio coding techniques, and the compression profile used to create DRC data in the bit stream.

H.3.2.2 Syntax

Syntax	No. of bits
programme_information()	
{	
version	2
if (version == 0x3)	
{	
version += **extended _version**	4
}	
activechane	1
if (activechane) **{activechan}**	16
dmixtype	1
if (dmixtype == 0)	
{	
upmixtype	1
if (upmixtype) **{upmixtyp}**	4
}	
else if (dmixtype) **{dmixtyp}**	4
preproinfoe	1
if (preproinfoe)	
{	
suratten	1
ph90filt	1
lfefilt	1
lfemonlevcod	2
}	
drcprofinfoe	1
if (drcprofinfoe)	
{	
dynrngprof	3
comprprof	3
}	
specprocinfoe	1
if (specprocinfoe)	

Syntax	No. of bits
{	
specprocstartf	3
specprocendf	3
}	
chancplinfoe	1
if (chancplinfoe)	
{	
chancplstartf	5
chancplendf	5
}	
enhncrnge	1
if (enhncrnge) **{enhncrng}**	2
fill bits = 0 /* up to byte boundary */	0 to 7
}	

H.3.2.3 Semantics

H.3.2.3.1 version - Version of programme information payload - 2 bits

The 2-bit version field shall indicate the version of the programme information payload. For programme information payloads that conform to this Annex, the version field shall be set to '00', and the "extended_version" field shall not be present in the payload.

H.3.2.3.2 activechane - Active programme channels data exists - 1 bit

If the activechane bit is set to '1', the activechan field shall follow in the payload. If the strmtyp value of the Enhanced AC-3 syncframe is set to '01' (dependent substream syncframe), this field shall be set to '0'.

H.3.2.3.3 activechan - Active programme channels - 16 bits

The 16-bit activechan field is used to indicate which channels within the programme contain audio information, and which (if any) contain only silence for the duration of the corresponding AC-3 or Enhanced AC-3 syncframe. This field shall be used in conjunction with the acmod field in the current syncframe (and, if present, the chanmap field in the associated dependent substream syncframe(s)) to determine which channels within the programme contain audio information and which contain silence. This information may be useful for upmixing processes implemented downstream of the AC-3 or Enhanced AC-3 decoder, for example to add audio to channels that contain silence at the output of the decoder. The bit allocation of the activechan field is shown in Table H.3.7, and is the same as the bit allocation of the chanmap field as specified in clause E.1.3.1.8. Bit 0 of the activechan field is the most significant bit of the field, and corresponds to the Left channel. Each bit of the activechan field shall be set to '1' if the corresponding channel location or pair of channel locations are present in the programme and contain audio content.

Table H.3.7: channel assignment of activechan bits

Bit	Location
0 (MSB)	Left
1	Centre
2	Right
3	Left Surround
4	Right Surround
5	Lc/Rc pair
6	Lrs/Rrs pair
7	Cs
8	Ts
9	Lsd/Rsd pair
10	Lw/Rw pair
11	Vhl/Vhr pair
12	Vhc
13	Lts/Rts pair
14	LFE2
15	LFE

H.3.2.3.4 dmixtype - Programme has been downmixed - 1 bit

If the audio programme was downmixed from a greater number of channels prior to AC-3 or Enhanced AC-3 encoding, the 1-bit dmixtype field shall be set to '1' and the 5-bit dmixtyp field shall follow in the bit stream. If the acmod field in the AC-3 or Enhanced AC-3 syncframe is set to a value other than 2 or 7, or the strmtyp value of the Enhanced AC-3 syncframe is '01' (dependent substream syncframe), this field shall be set to '0'.

H.3.2.3.5 dmixtyp - type of Downmixing applied to programme - 4 bits

The 5-bit dmixtyp field indicates the type of downmixing applied to the audio programme prior to AC-3 or Enhanced AC-3 encoding. This information may be useful to upmixing processes implemented downstream of the AC-3 or Enhanced AC-3 decoder, for example to upmix the programme audio using parameters that most closely match the downmixing method used. The interpretation of the dmixtyp field is dependent on the value of the acmod field in the AC-3 or Enhanced AC-3 syncframe. The details of each downmix are specified in Table H.3.8 and Table H.3.9.

Table H.3.8: Meaning of dmixtyp when acmod = 2

dmixtyp value	Downmix Equation
0	Reserved
1	L_o = L + (-3 dB) × C + (-3 dB) × Ls R_o = R + (-3 dB) × C + (-3 dB) × Rs
2	L_t = L + (-3 dB) × C - (-3 dB) × Ls - (-3 dB) × Rs R_t = R = R + (-3 dB) × C + (-3 dB) × Ls + (-3 dB) × Rs
3	L_t = L + (-3 dB) × C - (-1,2 dB) × Ls - (-6,2 dB) × Rs R_t = R + (-3 dB) × C + (-6,2 dB) × Ls + (-1,2 dB) × Rs
4 - 15	Reserved

Table H.3.9: Meaning of dmixtyp when acmod = 7

dmixtyp value	Downmix Equation
0	L = L R = R C = C LFE = LFE Ls' = Ls + (-3 dB) × Cs Rs' = Rs + (-3 dB) × Cs
1	L = L R = R C = C LFE = LFE Ls' = (-3 dB) × Ls + (-3 dB) × Lrs Rs' = (-3 dB) × Rs + (-3 dB) × Rrs
2	L = L R = R C = C LFE = LFE Ls' = Ls + (-1,2 dB) × Lrs + (-6,2 dB) × Rrs Rs' = Rs + (-1,2 dB) × Rrs + (-6,2 dB) × Lrs
3	L = L R = R C = C LFE = LFE Ls' = Ls - (-1,2 dB) × Lvh - (-6,2 dB) × Rvh Rs' = Rs + (-1,2 dB) × Rvh + (-6,2 dB) × Lvh
4	L' = L + Lvh R' = R + Lvh C = C LFE = LFE Ls = Ls Rs = Rs
5 - 15	Reserved

H.3.2.3.6 upmixtype - Programme has been upmixed - 1 bit

The 1-bit upmixtype field may be used to indicate whether the audio programme was upmixed from a smaller number of channels prior to AC-3 or Enhanced AC-3 encoding. If the upmixtype field is set to '1', the 5-bit upmixtyp field shall follow in the bit stream. If the acmod field in the AC-3 syncframe or in an Enhanced AC-3 syncframe is set to a value other than 7, this field shall be set to '0'.

H.3.2.3.7 upmixtyp - Type of upmixing applied to programme - 4 bits

The 5-bit upmixtyp field indicates the type of upmixing applied to the audio programme prior to AC-3 or Enhanced AC-3 encoding. This information may be useful to downmixing processes implemented downstream of the AC-3 or Enhanced AC-3 decoder, for example to ensure that the downmix method used is compatible with the upmixing method applied to the audio programme. If the programme information payload is present in an AC-3 syncframe, or an Enhanced AC-3 syncframe with a strmtyp value of '00', the upmixtyp field shall be interpreted as shown in Table H.3.10.

Table H.3.10: Meaning of upmixtyp when acmod = 7

upmixtyp value	Upmix Type
0	Dolby Pro Logic®
1	Dolby Pro Logic® II Movie Mode
2	Dolby Pro Logic® II Music Mode
3	Dolby® Professional Upmixer
4 - 15	Reserved

If the current programme information payload is present in an Enhanced AC-3 syncframe with a strmtyp value of '01'and a substreamid value of '000' (D0 syncframe), the upmixtyp field shall be interpreted as shown in Table H.3.11, and the interpretation shall be also dependent on the value of the chanmap field in the syncframe. The value of the upmixtyp field shall only be considered valid if the conditions identified in the "Restrictions" column are satisfied.

Table H.3.11: Meaning of upmixtyp when present in D0 syncframe

upmixtyp value	Upmix Type	Restrictions
0	Reserved	
1	Dolby Pro Logic® IIx	Valid only when bit 6 (Lrs/Rrs pair) of chanmap = '1'
2	Dolby Pro Logic® IIz Height	Valid only when bit 11 (Lvh/Rvh pair) of chanmap = '1'
3 - 15	Reserved	

H.3.2.3.8 preproinfoe - Preprocessing information exists - 1 bit

If the 1-bit preproinfoe field is set to '1', information about processing that was applied to the PCM audio before AC-3 or Enhanced AC-3 encoding shall follow in the payload. If the audio programme does not contain surround channels or an LFE channel, or if the current programme information payload is present in an Enhanced AC-3 syncframe with a strmtyp value of '01', this field shall be set to '0'.

H.3.2.3.9 suratten - Surround attenuation applied - 1 bit

If the surround channels of the audio programme were attenuated by 3 dB prior to encoding, the 1-bit suratten field shall be set to '1'. If the audio programme does not contain surround channels, this field shall be set to '0'.

H.3.2.3.10 ph90filt - 90 degree phase shift applied - 1 bit

If a 90 degree phase shift was applied to the Ls and Rs channels of the audio programme prior to AC-3 or Enhanced AC-3 encoding, the 1-bit ph90filt field shall be set to '1'. If the audio programme does not contain Ls and Rs channels, this field shall be set to '0'.

H.3.2.3.11 lfefilt - LFE filter applied - 1-bit

If a low-pass filter was applied to the LFE channel of the audio programme prior to encoding, the 1-bit lfefilt field shall be set to '1'. If the audio programme does not contain an LFE channel, this field shall be set to '0'.

H.3.2.3.12 lfemonlevcod - LFE channel monitoring level - 2 bits

The 2-bit lfemonlevcod field is used to indicate the level at which the LFE channel was monitored during production relative to the full range audio channels in the programme. Table H.3.12 shows the interpretation of the lfemonlevcod field. The value of this field may be used by the system to align the reproduction level of the LFE channel relative to the reproduction level of the full range audio channels in the programme. If the audio programme does not contain an LFE channel, this field shall be set to '00'.

Table H.3.12: Meaning of lfemonlevcod values

lfemonlevcod value	LFE channel monitoring level
00	Not indicated
01	+10 dB relative to full range audio programme channels
10	0 dB relative to full range audio programme channels
11	Reserved

H.3.2.3.13 drcprofinfoe - DRC profile information exists - 1 bit

If the 1-bit drcprofinfoe field is set to '1', fields describing the compression profile used to compute the values of DRC data in the AC-3 or Enhanced AC-3 syncframe shall follow in the payload.

H.3.2.3.14 dynrngprof - DRC profile for dynrng data - 3 bits

The 3-bit dynrngprof field indicates the compression profile that was used by the AC-3 or Enhanced AC-3 encoder to compute the values of the dynrng DRC data in the current AC-3 or Enhanced AC-3 syncframe. Table H.3.13 shows the interpretation of dynrngprof values.

Table H.3.13: DRC profiles indicated by dynrngprof

dynrngprof value	DRC profile
000	None
001	Film Standard
010	Film Light
011	Music Standard
100	Music Light
101	Speech
110 - 111	Reserved

The compression curves and time constants of each compression profile are specified in Table H.3.14. All gain boundary values have a reference level of -31 dB.

Table H.3.14: Compression curves and time constants of DRC profiles indicated by dynrngprof

Parameter \ Profile	None	Film standard	Film light	Music standard	Music light	Speech
Compression Curve						
Maximum boost gain [dB]	0	6	6	12	12	15
Maximum boost gain boundary [dB]	0	-12	-22	-24	-34	-19
Null-band lower boundary [dB]	0	0	-10	0	-10	0
Null-band upper boundary [dB]	0	5	10	5	10	5
Cut section gain [dB]	0	-5	-5	-5		-5
Cut section gain boundary [dB]	0	15	20	15		15
Maximum cut gain [dB]	0	-24	-24	-24	-15	-24
Maximum cut gain boundary [dB]	0	35	40	35	40	35
Time smoothing						
Slow Attack Time Constant [ms]	N/A	100	100	100	100	100

Parameter \ Profile	None	Film standard	Film light	Music standard	Music light	Speech
Slow Release Time Constant [ms]	N/A	3 000	3 000	10 000	3 000	1 000
Fast Attack Time Constant [ms]	N/A	10	10	10	10	10
Fast Release Time Constant [ms]	N/A	1 000	1 000	1 000	1 000	200
Attack threshold [dB]	N/A	15	15	15	15	10
Release threshold [dB]	N/A	20	20	20	20	10
NOTE: Profile **None** has a constant gain of 0 dB (i.e. it effectively disables dynamic range compression). Therefore, the time constants are not applicable.						

H.3.2.3.15 comprprof - DRC profile for compr data - 3 bits

The 3-bit comprprof field indicates the compression profile that was used by the AC-3 or Enhanced AC-3 encoder to compute the values of the compr DRC data in the current AC-3 or Enhanced AC-3 syncframe. The comprprof field is interpreted in the same way as dynrngprof, per clause H.3.2.3.14.

H.3.2.3.16 specprocinfoe - Spectral processing information exists - 1 bit

If the 1-bit specprocinfoe field is set to '1', fields describing the portion of audio bandwidth that is coded using spectral processing techniques shall follow in the payload. If the current programme information payload is present in an AC-3 syncframe, or if the Enhanced AC-3 syncframe is not coded using spectral processing techniques, this field shall be set to '0'.

H.3.2.3.17 specprocstartf - Spectral processing start frequency - 3 bits

The 3-bit specprocstartf field shall indicate the lowest frequency at which spectral processing is being used to code the programme audio. This information may be useful to equalization processes implemented downstream of the Enhanced AC-3 decoder. The values of 0 to 7 shall be interpreted as shown in Table H.3.15.

Table H.3.15: Frequencies indicated by specprocstartf

specprocstartf value	Frequency (kHz)
000	4,55
001	5,67
010	6,80
011	7,92
100	9,05
101	10,17
110	12,42
111	14,67

The value of specprocstartf shall be lower than the value of the following specprocendf field.

H.3.2.3.18 specprocendf - Spectral processing end frequency - 3 bits

The 3-bit specprocendf field shall indicate the highest frequency at which spectral processing is being used to code the programme audio. This information may be useful to equalization processes implemented downstream of the Enhanced AC-3 decoder. The values of 0 to 7 shall be interpreted as shown in Table H.3.16.

Table H.3.16: Frequencies indicated by specprocendf

specprocendf value	Frequency (kHz)
000	7,92
001	9,05
010	10,17
011	12,42
100	14,67
101	16,92
110	19,17
111	21,42

The value of specprocendf shall be higher than the value of the preceding specprocstartf field.

H.3.2.3.19 chancplinfoe - Channel coupling information exists - 1 bit

If the 1-bit chancplinfoe field is set to '1', fields describing the portion of audio bandwidth that is coded using channel coupling shall follow in the payload. If the AC-3 or Enhanced AC-3 is not coded using channel coupling, this field shall be set to '0'.

H.3.2.3.20 chancplstartf - Channel coupling start frequency - 5 bits

The 8-bit chancplstartf field shall indicate the lowest frequency at which channel coupling is being used to code the programme audio. This information may be useful to upmixing processes implemented downstream of the AC-3 or Enhanced AC-3 decoder. Values of 0 to 26 shall be interpreted as shown in Table H.3.17. Values of 27 to 31 shall be reserved and, if present, shall be interpreted as a frequency of 0 kHz. The value of chancplstartf shall be lower than the value of the following chancplendf field.

Table H.3.17: Frequencies indicated by chancplstartf/chancplendf

chancplstartf	Frequency (kHz)
0	0
1	0,29
2	0,58
3	0,88
4	1,17
5	1,73
6	2,30
7	2,86
8	3,42
9	4,55
10	5,67
11	6,80
12	7,92
13	9,05
14	10,17
15	11,30
16	12,42
17	13,55
18	14,67
19	15,80
20	16,92
21	18,05
22	19,17
23	20,30
24	21,42
25	22,55
26	23,67
27 - 31	reserved

H.3.2.3.21 chancplendf - Channel coupling end frequency - 5 bits

The 8-bit chancplstendf field shall indicate the lowest frequency at which channel coupling is being used to code the programme audio. This information may be useful to upmixing processes implemented downstream of the AC-3 or Enhanced AC-3 decoder. Values of 0 to 26 shall be interpreted as shown in Table H.3.17. Values of 27 to 31 shall be reserved and, if present, shall be interpreted as a frequency of 23.67 kHz. The value of chancplstartf shall be lower than the value of the following chancplendf field. The value of chancplendf shall be higher than the value of the preceding chancplstartf field.

H.3.2.3.22 enhncrnge - Dialog enhancement adjustment range data exists - 1 bit

If the 1-bit enhncrnge field is set to '1' the enhcrng field shall follow in the payload.

H.3.2.3.23 enhncrng - Dialog enhancement adjustment range - 2 bits

The 2-bit enhncrng field shall specify the range of adjustment available to the user of a dialog enhancement process when the user wishes to adjust the level of dialogue content relative to the level of non-dialogue content in the audio programme. The values of the enhncrng field shall be interpreted as shown in Table H.3.18.

Table H.3.18: Adjustment range indicated by enhncrng

enhncrng value	Adjustment range permitted
00	± 3 dB
01	± 6 dB
10	± 9 dB
11	± 12 dB

H.3.3 Payload ID 0x03 - Enhanced AC-3 substream structure

H.3.3.1 Overview

The Enhanced AC-3 substream structure payload provides an overall description of the substream structure of an Enhanced AC-3 bit stream. The Enhanced AC-3 substream structure payload shall be stored only in an Enhanced AC-3 syncframe that has a strmtyp value of '0' and a substreamid value of '0'.

H.3.3.2 Syntax

Syntax	No. of bits
eac3_substream_structure()	
{	
num_ind_sub	3
for (i = 0; i < num_ind_sub + 1, i++)	
{	
dep_sub_exists[i]	1
if (dep_sub_exists[i]) **{num_dep_sub[i]}**	3
}	
fill bits = 0 /* up to byte boundary */	0 to 7
}	

H.3.3.3 Semantics

H.3.3.3.1 num_ind_sub - Number of independent substreams - 3 bits

The 3-bit num_ind_sub field shall be set to a value of one less than the number of independent substreams in the E-AC-3 bit stream. This value is equal to the value of the substreamid field of the last independent substream in the bit stream.

H.3.3.3.2 dep_sub_exist - Dependent substreams exist- 1 bit

The 1-bit dep_sub_exists field shall be set to '1' if one or more dependent substreams are associated with independent substream [i]. If the dep_sub_exists field is set to '1', the num_dep_sub field shall follow in the payload.

H.3.3.3.3 num_dep_sub - Number of dependent substreams - 3 bits

The 3-bit num_dep_sub field shall be set to a value of one less than the number of dependent substreams associated with the corresponding independent substream[i]. This value is equal to the value of the substreamid field of the last dependent substream in the bit stream.

H.3.4 Payload ID 0x04 - dynamic range compression data for portable devices

H.3.4.1 Overview

This payload provides dynamic range compression data intended to be used by portable devices. Portable devices typically have limited loudspeaker output capabilities, and are primarily used in environments with high levels of ambient noise, requiring the dynamic range of the programme audio to be significantly restricted to ensure audibility.

H.3.4.2 Syntax

Syntax	No. of bits
portable_device_drc()	
{	
version	2
if (version == 0x3)	
{	
version += **extended _version**	4
}	
portdrce	1
if (portdrce)**{portdrc}**	8
drcprofinfoe	1
if (drcprofinfoe)**{portdrcprof}**	3
fill bits = 0 /* up to byte boundary */	0 to 7
}	

H.3.4.3 Semantics

H.3.4.3.1 version - Version of portable device DRC payload - 2 bits

The 2-bit version field shall indicate the version of the portable device DRC payload. For portable device DRC payloads that conform to this Annex, the version field shall be set to '00', and the "extended_version" field shall not be present in the payload.

H.3.4.3.2 portdrce - Portable device DRC data exists - 1 bit

If the 1-bit portdrce field is set to '1', the portdrc field follows in the payload.

H.3.4.3.3 portdrc - Portable device DRC gain word - 8 bits

The 8-bit portdrc field is generated by the AC-3 or Enhanced AC-3 encoder, and may be used by the AC-3 or Enhanced AC-3 decoder to scale the reproduced audio level in order to reproduce an extremely narrow dynamic range. The portdrc field is interpreted in the same way as the compr field, as specified in clause 6.7.3.

H.3.4.3.4 drcprofinfoe - DRC profile information exists - 1 bit

If the 1-bit drcprofinfoe field is set to '1', fields describing the compression profile used to compute the values of the portdrc field shall follow in the payload.

H.3.4.3.5 portdrcprof - DRC profile for portable device DRC data - 3 bits

The 3-bit portdrcprof field indicates the compression profile that was used by the AC-3 or Enhanced AC-3 encoder to compute the values of the portdrc DRC data in the current payload. The interpretation of the portdrcprof field shall be the same as the dynrngprof field, as shown in Table H.3.13.

H.3.5 Payload ID 0x05 - programme language

H.3.5.1 Overview

The programme language payload indicates the language of the audio programme. This payload shall use a language tag and subtag structure that conforms to [i.20]. For Enhanced AC-3 bit streams, this payload shall only be carried in independent substreams (strmtyp = 0x0).

H.3.5.2 Syntax

Syntax	No. of bits
programme_language()	
{	
frame_sequence	1
if (frame_sequence)	
{	
sequence_start	1
sequence_end	1
langtag_chunk	16
}	
else	
{	
langtag_length	6
langtag	8 x (langtag_length + 2)
}	
fill bits = 0 /* up to byte boundary */	0 to 7
}	

H.3.5.3 Semantics

H.3.5.3.1 frame_sequence - Syncframe sequence indication - 1 bit

The 1-bit frame_sequence field is used to indicate whether the language tag is delivered using multiple payloads in a sequence of AC-3 or Enhanced AC-3 syncframes, or whether the complete language tag is delivered in the current payload. A value of '1' shall indicate that the language tag is delivered using multiple payloads delivered in a sequence of AC-3 or Enhanced AC-3 syncframes, and information describing the sequence of syncframes shall follow in the payload. A value of '0' shall indicate that the complete language tag is stored in the current payload. The value of frame_sequence shall be set to '1' in all programme language payloads that follow the programme language payload in which both the value of the frame_sequence field and the value of the sequence_start field (see clause H.3.5.3.2) is set to '1', up to and including the payload that contains the final byte of the language tag.

H.3.5.3.2 sequence_start - Syncframe sequence start indication - 1 bit

The 1-bit sequence_start field is used to indicate the start of a multi-frame sequence of language tag data. A value of '1' shall indicate that this payload contains the first chunk of the language tag, and decoders shall start decoding the language tag beginning with the data in this payload. A value of '0' shall indicate that this payload does not contain the start of the multi-frame sequence.

H.3.5.3.3 sequence_end - Syncframe sequence end indication - 1 bit

The 1-bit sequence_end field is used to indicate the end of a multi-frame sequence of language tag data. A value of '1' shall indicate that this payload contains the last chunk of the language tag, and decoders shall finish decoding the language tag using the data in this payload. A value of '0' shall indicate that this payload does not contain the end of the multi-frame sequence.

H.3.5.3.4 langtag_chunk - Language tag data chunk - 16 bits

The 16-bit langtag_chunk field shall contain a two byte section of a language tag that conforms to the syntax and semantics defined in [i.20]. The most significant byte of the language tag shall be stored in the most significant byte of the langtag_chunk in a programme language payload that has a frame_sequence field value of '1'. Subsequent bytes of the language tag shall be stored in the second byte of the current langtag_chunk and in the langtag_chunk of each of the subsequent programme language payloads required to deliver the complete language tag. If the length of the language tag does not correspond to an even number of bytes, the value of the extra byte that follows the end of the language tag shall be set to 0x0.

H.3.5.3.5 langtag_length - Language tag length - 6 bits

The 6-bit langtag_length field shall indicate the total length, in bytes, of the language tag. Values of 0 to 40 shall be interpreted as 2 to 42 bytes. Values of 41 to 63 are reserved.

H.3.5.3.6 langtag - Language_tag - 16 to 336 bits

The variable length langtag field shall contain a language tag that conforms to the syntax and semantics defined in IETF BCP-47 [i.20]. The minimum length of the langtag field shall be two bytes, supporting a two-character language tag as specified in ISO 639-1 [i.1]. The maximum length of the language tag supported by this field shall be 336 bits or 42 bytes.

H.3.6 Payload ID 0x6 - external data

H.3.6.1 Overview

The external data payload is intended to carry user-defined data types that do not have an assigned payload ID.

H.3.6.2 Syntax

Syntax	No. of bits
external_data()	
{	
frame_sequence	1
if (frame_sequence)	
{	
sequence_start	1
sequence_end	1
external_data_chunk	16
}	
else	
{	
external_data_length	variable_bits(7)
external_data	8 x (external_data_length)
}	
fill bits = 0 /* up to byte boundary */	0 to 7
}	

H.3.6.3 Semantics

H.3.6.3.1 frame_sequence - Syncframe sequence indication - 1 bit

The 1-bit frame_sequence field is used to indicate whether the user-defined data are delivered using multiple payloads in a sequence of AC-3 or Enhanced AC-3 syncframes, or whether the user-defined data are delivered in the current payload only. A value of '1' shall indicate that the user-defined data is delivered using multiple payloads delivered in a sequence of AC-3 or Enhanced AC-3 syncframes, and information describing the sequence of syncframes shall follow in the payload. A value of '0' shall indicate that the current payload contains the complete user-defined data. The value of the frame_sequence field shall be set to '1' in all external data payloads that follow the external data payload in which both the value of the frame_sequence field and the value of the sequence_start field (see clause H.3.6.3.2) is set to '1', up to and including the payload that contains the final byte of the user-defined data.

H.3.6.3.2 sequence_start - Syncframe sequence start flag - 1 bit

The 1-bit sequence_start field is used to indicate the start of a multi-syncframe sequence of user-defined data. A value of '1' shall indicate that this payload contains the first 16 bits of the user-defined data, and decoders shall start decoding the user-defined data beginning with the data in this payload. A value of '0' shall indicate that this payload does not contain the start of the multi-syncframe sequence.

H.3.6.3.3 sequence_end - Syncframe sequence end flag - 1 bit

The 1-bit sequence_end field is used to indicate the end of a multi-syncframe sequence of user-defined data. A value of '1' shall indicate that this payload contains the last 16 bits of the user-defined data, and decoders shall complete decoding the user-defined data using the data in this payload. A value of '0' shall indicate that this payload does not contain the end of the multi-syncframe sequence.

H.3.6.3.4 external_data_chunk - External data chunk - 16 bits

If the frame_sequence field is set to '1', the 16-bit external_data_chunk field shall contain a 16-bit chunk of the user-defined data. The most significant byte of the user-defined data shall be stored in the most significant byte of the external_data_chunk field. Subsequent bytes of the user-defined data shall be stored in the second byte of the current external_data_chunk field, and in the external_data_chunk field of each of the subsequent external data payloads required to deliver the user-defined data. If the length of the user-defined data does not correspond to an even number of bytes, all bits that follow the end of the user-defined data shall be set to '0'.

H.3.6.3.5 external_data_length - External data field length - variable_bits(7)

If the frame_sequence field is set to '0', the variable length external_data_length field shall indicate the total length, in bytes, of the external_data field.

H.3.6.3.6 external_data - External data field - external_data_length * 8 bits

If the multi-syncframe sequence is not used to deliver the user-defined data, the variable length external_data field shall contain the user-defined data.

H.3.7 Payload ID 0x7 - headphone rendering data

H.3.7.1 Overview

The headphone rendering data payload provides data intended for use by a headphone virtualization algorithm.

H.3.7.2 Syntax

Syntax	No. of bits
headphone_rendering_data()	
{	
chan_count	3
for (ch = 0; ch < chan_count; ch++)	
{	
chan_gain[ch]	6
}	
lfechangaine	1
if (lfechangaine) **{lfechangain}**	6
sequence_start	1
sequence_end	1
brir_datae	1
if (brir_datae)	
{	
for (ch = 0; ch < chan_count; ch++)	
{	
chan_brir_chunk[ch]	96
}	
late_BRIR_chunk	1536
}	
if (sequence_end) **{parity_check}**	8

Syntax	No. of bits
`if (sequence_start)`	
`{`	
`  propdlye`	1
`  if (propdlye)`	
`  {`	
`    for (ch = 0; ch < chan_count; ch++)`	
`    {`	
`      propdly[ch]`	11
`    }`	
`  }`	
`  rt60date`	1
`  if (rt60date)`	
`  {`	
`    num_bnds`	4
`    for (bnd = 0; bnd < num_bnds; bnd++)`	
`    {`	
`      RT60[bnd]`	11
`    }`	
`  }`	
`}`	
`fill bits = 0 /* up to byte boundary */`	0 to 7
`}`	

H.3.7.3 Semantics

H.3.7.3.1 chan_count - Number of full-bandwidth channels - 3 bits

The 3-bit chan_count field shall indicate the number of full-bandwidth audio channels in the AC-3 or Enhanced AC-3 bitstream to be processed for binaural output. The value of the chan_count field shall be between 1 and 7, indicating between 1 and 7 full bandwidth channels.

H.3.7.3.2 chan_gain - Individual channel gain - 6 bits

The 8-bit chan_gain[i] field shall indicate the gain to be applied to each full bandwidth channel of the AC-3 or Enhanced AC-3 bitstream before binaural processing. Values of 0 to 63 shall be interpreted as gains between 0 and -63 dB in 1 dB steps.

H.3.7.3.3 lfegaine - LFE channel gain value exists - 1 bit

If the 1-bit lfegaine field is set to 1, LFE channel gain data shall follow in the bitstream.

H.3.7.3.4 lfegain - LFE channel gain value - 6 bits

The 8-bit lfegain field shall indicate the gain to be applied to the LFE channel of the AC-3 or Enhanced AC-3 bitstream before binaural processing. Values of 0 to 673 shall be interpreted as gains between 0 and -63 dB in 1 dB steps.

H.3.7.3.5 sequence_start - Syncframe sequence start indication - 1 bit

To minimize the data rate impact of the headphone rendering payload, some parameters are serialized across multiple headphone rendering payloads, delivered in a sequence of AC-3 or Enhanced AC-3 syncframes. The 1-bit sequence_start field is used to indicate the start of the multi-syncframe sequence of headphone rendering data. A value of '1' shall indicate that this payload contains the start of the serialized headphone rendering data, and decoders shall start decoding the headphone rendering data beginning with the data in this payload. A value of '0' shall indicate that this payload does not contain the start of the multi-syncframe sequence.

H.3.7.3.6 sequence_end - Syncframe sequence end indication - 1 bit

The 1-bit sequence_end field is used to indicate the end of the multi-syncframe sequence of headphone rendering data. A value of '1' shall indicate that this payload contains the final set of serialized headphone rendering data for this sequence, and decoders shall complete decoding the headphone rendering data using the data in this payload. A value of '0' shall indicate that this payload does not contain the end of the multi-syncframe sequence.

H.3.7.3.7 brir_datae - BRIR data exists - 1 bit

If the 1-bit brir_datae field is set to '1', data specifying the binaural room impulse response shall follow in the bitstream.

H.3.7.3.8 chan_brir_chunk - Channel BRIR data chunk - 96 bits

The chan_brir[ch] variable is a multi-dimensional array used to specify 24-bit FIR filter coefficients that represent the early BRIRs for channel ch to be used by the binaural rendering process. The size of each vector in the chan_BRIR[ch] array is defined as follows: 24-bit filter coefficients × 2 binaural output channels × 64 samples. To minimize the data rate overhead of the array, the data is delivered in a series of 96-bit chan_brir_chunk fields in a sequence of 32 AC-3 or six-block Enhanced AC-3 syncframes, resulting in an update rate for the chan_brir array data of once every 1,024 seconds. The first 96 bit chunk of the chan_brir[ch] variable shall be delivered in the syncframe that has a sequence_start field value of '1'. The last 96 bit chunk of the chan_brir[ch] variable shall be delivered in the syncframe that has a sequence_end field value of '1'.

H.3.7.3.9 late_BRIR_chunk - Late brir data chunk - 1536 bits

The late_brir variable is a multi-dimensional arrayused to specify 24-bit FIR filter coefficients that represent the combined set of late BRIRs for each output channel of the binaural rendering process.The size of each vector in the late_BRIR array is defined as follows: 24-bit filter coefficients × 2 binaural output channels × 1 024 samples. To minimize the data rate overhead of the array, the data is delivered in a series of 1 536-bit late_brir_chunk fields in a sequence of 32 AC-3 or six-block Enhanced AC-3 syncframes, resulting in an update rate for the late_brir array data of once every 1,024 seconds. The first 1 536 bit chunk of the late_brir variable shall be delivered in the syncframe that has a sequence_start field value of '1'. The last 1 536 bit chunk of the late_brir variable shall be delivered in the syncframe that has a sequence_end field value of '1'

H.3.7.3.10 parity_check - 8 bits

The 8-bit parity_check field is used to check that the set of chan_brir_chunk and late_brir_chunk data chunks delivered in the sequence of syncframes that start with sequence_start = '1' and ends with sequence_end = '1' is complete and valid. The parity check shall be calculated by performing an exclusive OR operation of all the bytes of the chan_brir_chunk and late_brir_chunk data chunks from all frames of a single sequence of chunks, beginning with the syncframe which has a sequence_start field value of '1', and ending with the syncframe which has a sequence_end field value of '1', exclusive OR'd with the value 0xA9. The purpose of the latter is to force a failure in the case that all data chunks are zeroes.

H.3.7.3.11 propdlye - Propagation delay data exists - 1 bit

If the 1-bit propdlye field is set to '1', per-channel propagation delay data shall follow in the bitstream.

H.3.7.3.12 propdly - Propagation delay data - 11 bits

The 11-bit propdly field shall specify the propagation delay of channel ch in units of samples. Valid values shall be between 0 and 2 047 samples.

H.3.7.3.13 rt60e - RT_{60} data exists - 1 bit

If the 1-bit rt60e field is set to '1', RT_{60} data shall follow in the bitstream.

H.3.7.3.14 num_bnds - Number of bands of RT_{60} data - 4 bits

The 4-bit num_bnds field shall indicate the number of octave bands for which RT_{60} data is present. Valid values shall be between 1 and 11. Values of 0 and 12 to 15 are reserved. The centre frequency of each octave band shall be as specified in Table H.3.19.

Table H.3.19: Centre frequencies of each octave band indicated by num_bnds

num_bnds value	Band number	Centre Frequency (Hz)
0	N/A	N/A
1	1	16
2	2	31.5
3	3	63
4	4	125
5	5	250
6	6	500
7	7	1 000
8	8	2 000
9	9	4 000
10	10	8 000
11	11	16 000
12 - 15	N/A	N/A

H.3.7.3.15 rt60[bnd] - RT_{60} value - 11 bits

The 11-bit rt60[bnd] field specifies the RT_{60} value for band bnd of the binaural rendering environment. The value of the rt60[bnd] parameter may be used in the rendering device to update the room model or synthesize the reverb tail during binaural rendering. Values 1 to 2 047 shall be interpreted as RT_{60} values of 0 to 2 047 milliseconds.

Annex I (normative): Signalling in MPEG-DASH

I.0 Scope

This annex describes the requirements and recommendations for delivering streams using the MPEG Dynamic Adaptive Streaming over HTTP (DASH) standard in conjunction with the ISO base media file format, specifically referencing audio streams within the MPEG-DASH media presentation description (MPD) file.

I.1 Media Presentation Description (MPD)

I.1.1 General MPD requirements

I.1.1.1 Introduction

The syntax of the MPD can consist of common XML elements to describe almost any media format, the encoding type and the configuration of a media stream that is part of a larger media presentation. This clause defines the values for selected attributes and descriptors that are used in an MPD to properly describe a media stream.

The MPD and the contained XML elements shall conform to ISO/IEC 23009-1 [2].

I.1.1.2 Adaptation Sets and Representations

The information conveyed by an Adaptation Set XML element describes the presentation of one media component. An Adaptation Set typically consists of multiple instances (Representations) of the same audio content, with each instance encoded at a different data rate:

- The @codecs value is required. It specifies the codecs used to encode the Representations within one Adaptation Set. The value of @codecs shall be the value of the codecs parameter as specified in clause J.2.
- The @mimeType value indicates the format used for encapsulation of the elementary streams present in the adaptation set. The value of @mimeType shall be set to 'audio/mp4'.
- The @startWithSAP value shall be set to '1'.

Each Adaptation Set carries one or more Representations.

I.1.1.3 AudioChannelConfiguration descriptor

The XML Element that references the elementary stream shall include an AudioChannelConfiguration descriptor, which unambiguously describes the channel configuration of the referenced elementary stream. Refer to clause I.1.2.1 for details.

I.1.2 Descriptors

I.1.2.1 AudioChannelConfiguration descriptor

For the channel configurations listed in table I.1 column 1 (indicated by the acmod and chanmap field), the audio channel configuration descriptor should use the scheme described in the *schemeIdUri* `urn:mpeg:mpegB:cicp:ChannelConfiguration.` The value should be as listed in column 3, and interpreted as index to ISO/IEC 23001-8 [4], Table 9. Valid values are 1-7, 9-12, 14, 16-17, and 19. Column 2 in Table I.1.1 provides additional information how individual channels relate for the specified mapping.

The audio channel configuration descriptor may use the scheme described in the *schemeIdUri* `tag:dolby.com,2014:dash:audio_channel_configuration:2011`.

The *value* element shall contain a four-digit hexadecimal representation of the 16-bit field, which describes the channel assignment of the referenced stream according to Table E.1.4, clause E.1.3.1.8, where left channel is MSB.

Table I.1.1: Mapping of channels indicated by acmod and chanmap to MPEG channel configuration scheme

Present document	ISO/IEC 23001-8 [4] CICP Table 3	
Channels indicated by acmod and chanmap	**Loudspeaker names in Loudspeaker Layout**	**Value**
Centre	Centre Front	1
Left	Left Front	2
Right	Right Front	
Centre	Centre Front	3
Left	Left Front	
Right	Right Front	
Centre	Centre Front	4
Left	Left Front	
Right	Right Front	
Cs	Rear Centre	
Centre	Centre Front	5
Left	Left Front	
Right	Right Front	
Left Surround	Left Surround	
Right Surround	Right Surround	
Centre	Centre Front	6
Left	Left Front	
Right	Right Front	
Left Surround	Left Surround	
Right Surround	Right Surround	
LFE	LFE	
Centre	Centre Front	7
Left	Left Front Centre	
Right	Right Front Centre	
Lw/Rw pair	Left Front	
	Right Front	
Left Surround	Left Surround	
Right Surround	Right Surround	
LFE	LFE	
Left	Left Front	9
Right	Right Front	
Cs	Rear Centre	
Left	Left Front	10
Right	Right Front	
Left Surround	Left Surround	
Right Surround	Right Surround	
Centre	Centre Front	11
Left	Left Front	
Right	Right Front	
Lrs/Rrs pair	Left Surround	
	Right Surround	
Cs	Rear Centre	
LFE	LFE	
Centre	Centre Front	12
Left	Left Front	
Right	Right Front	
Left Surround	Left Surround	
Right Surround	Right Surround	
Lrs/Rrs pair	Rear Surround Left	
	Rear Surround Right	
LFE	LFE	
Centre	Centre Front	14
Left	Left Front	
Right	Right Front	
Left Surround	Left Surround	
Right Surround	Right Surround	

Present document	ISO/IEC 23001-8 [4] CICP Table 3	
Channels indicated by acmod and chanmap	**Loudspeaker names in Loudspeaker Layout**	**Value**
LFE	LFE	
Vhl/Vhr pair	Left Front Vertical Height	
	Right Front Vertical Height	
Centre	Centre Front	16
Left	Left Front	
Right	Right Front	
Left Surround	Left Surround	
Right Surround	Right Surround	
LFE	LFE	
Vhl/Vhr pair	Left Front Vertical Height	
	Right Front Vertical Height	
Lts/Rts pair	Left Vertical Height Surround	
	Right Vertical Height Surround	
Centre	Centre Front	17
Left	Left Front	
Right	Right Front	
Left Surround	Left Surround	
Right Surround	Right Surround	
LFE	LFE	
Vhl/Vhr pair	Left Front Vertical Height	
	Right Front Vertical Height	
Vhc	Centre Front Vertical Height	
Lts/Rts pair	Left Vertical Height Surround	
	Right Vertical Height Surround	
Ts	Top Centre Surround	
Centre	Centre Front	19
Left	Left Front	
Right	Right Front	
Lsd/Rsd pair	Left Side Surround	
	Right Side Surround	
Lrs/Rrs pair	Rear Surround Left	
	Rear Surround Right	
LFE	LFE	
Vhl/Vhr pair	Left Front Vertical Height	
	Right Front Vertical Height	
Lts/Rts pair	Left Vertical Height Surround	
	Right Vertical Height Surround	

EXAMPLE 1: For an AC-3 stream with acmod = 2 (Stereo) the value of the audio channel configuration descriptor for the scheme `urn:mpeg:mpegB:cicp:ChannelConfiguration` should be 2.

EXAMPLE 2: For an E-AC-3 stream with acmod = 7 for independent stream 0, and chanmap = 0x0058 for dependent stream 0, the value of the audio channel configuration descriptor for the scheme `urn:mpeg:mpegB:cicp:ChannelConfiguration` should be 12.

EXAMPLE 3: For a stream with L, C, R, Ls, Rs, LFE, the value shall be set to "F801" (i.e. the hexadecimal equivalent of the binary value 1111 1000 0000 0001).

I.2 Example Media Presentation Description

This is a simple example of a dynamic presentation, with multiple languages and multiple base URLs.

The MPD document in Table I.2.1 describes content available from two sources with audio available in two different English-language presentations: main audio service only, or a visually impaired receiver-mix service. The visually impaired service is enabled by simultaneously delivering the Enhanced AC-3 bitstream containing the main audio service and an additional Enhanced AC-3 bitstream containing the associated audio service for visually impaired listeners.

Three versions of the video are provided at bit rates between 250 kbps and 1 Mbps in different spatial resolutions.

Table I.2.1: Example Media Presentation Description

```
<?xml version="1.0" encoding="utf-8"?>
<MPD xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:dolby="http://www.dolby.com/ns/online/DASH" xmlns="urn:mpeg:dash:schema:mpd:2011"
xsi:schemaLocation="urn:mpeg:dash:schema:mpd:2011 DASH-MPD.xsd"
type="dynamic"
minimumUpdatePeriod="PT2S"
timeShiftBufferDepth="PT30M"
availabilityStartTime="2011-12-25T12:30:00"
minBufferTime="PT4S"
profiles="urn:mpeg:dash:profile:isoff-live:2011">
   <BaseURL>http://cdn1.example.com/</BaseURL>
   <BaseURL>http://cdn2.example.com/</BaseURL>
   <Period>
      <!-- Video -->
      <AdaptationSet mimeType="video/mp4" codecs="avc1.4D401F" frameRate="30000/1001"
       segmentAlignment="true" startWithSAP="1">
         <BaseURL>video/</BaseURL>
         <SegmentTemplate timescale="90000" media="$Bandwidth$/$Number$.m4s"
initialization="$Bandwidth$/0.mp4">
            <SegmentTimeline>
               <S t="0" d="180180" r="12"/>
            </SegmentTimeline>
         </SegmentTemplate>
         <Representation id="v0" width="320" height="240" bandwidth="250000" />
         <Representation id="v1" width="640" height="480" bandwidth="500000" />
         <Representation id="v2" width="960" height="720" bandwidth="1000000" />
      </AdaptationSet>
      <!-- English Audio -->
      <AdaptationSet mimeType="audio/mp4" codecs="ec-3" lang="en" segmentAlignment="0"
startWithSAP="1">
         <Accessibility schemeIdUri="urn:tva:metadata:cs:AudioPurposeCS:2007" value="6"/>
         <Role schemeIdUri="urn:mpeg:dash:role:2011" value="main" />
         <SegmentTemplate timescale="48000" media="audio/en_main/$Bandwidth$/$Number$.m4s"
           initialization="audio/en_main/$Bandwidth$/0.mp4">
            <SegmentTimeline>
               <S t="0" d="96000" r="11"/>
            </SegmentTimeline>
         </SegmentTemplate>
         <Representation id="a0" bandwidth="256000">
            <AudioChannelConfiguration
schemeIdUri="tag:dolby.com,2014:dash:audio_channel_configuration:2011" value="F801"/>
         </Representation>
      </AdaptationSet>
      <!-- English Audio for visually impaired listeners -->
      <AdaptationSet mimeType="audio/mp4" codecs="ec-3" lang="en" segmentAlignment="true"
startWithSAP="1">
         <Accessibility schemeIdUri="urn:tva:metadata:cs:AudioPurposeCS:2007" value="1"/>
         <Role schemeIdUri="urn:mpeg:dash:role:2011" value="commentary" />
         <SegmentTemplate timescale="48000" media="audio/en_vi/$Bandwidth$/$Number$.m4s"
          initialization="audio/en_vi/$Bandwidth$/0.mp4">
            <SegmentTimeline>
               <S t="0" d="96000" r="11"/>
            </SegmentTimeline>
         </SegmentTemplate>
         <Representation id="a1" dependencyId="a0" bandwidth="64000">
            <AudioChannelConfiguration
schemeIdUri="tag:dolby.com,2014:dash:audio_channel_configuration:2011" value="4000"/>
         </Representation>
      </AdaptationSet>
      <!-- French Audio -->
      <AdaptationSet mimeType="audio/mp4" codecs="ec-3" lang="fr" segmentAlignment="0"
startWithSAP="1">
            <SegmentTemplate timescale="48000" media="audio/fr/$Bandwidth$/$Number$.m4s"
             initialization="audio/fr/$Bandwidth$/0.mp4">
               <SegmentTimeline>
                  <S t="0" d="96000" r="11"/>
               </SegmentTimeline>
            </SegmentTemplate>
         <Representation id="a2" bandwidth="192000">
            <AudioChannelConfiguration
schemeIdUri="tag:dolby.com,2014:dash:audio_channel_configuration:2011" value="F801"/>
         </Representation>
      </AdaptationSet>
   </Period>
</MPD>
```

Annex J (normative): Bit streams in the Common Media Application Format (CMAF)

J.1 CMAF Tracks

J.1.1 Overview

This section specifies constraints that apply to all audio CMAF Tracks containing audio as defined in the present document. These constraints are used to derive the Core Media Profile specified in clause J.3.

CMAF Tracks shall conform to Annex F as constrained and extended here; and shall conform to MPEG CMAF [1], sections 6, 7, 8, 10.1 and 10.2.

J.1.2 *codecs* parameter signaling

The value of the codecs parameter shall be as specified in Table J.1.1. This is the MIME type codecs parameter as specified in IETF RFC 6381 [3].

Table J.1.1: codecs parameter

Audio coding algorithm	Value of codecs parameter
AC-3	ac-3
Enhanced AC-3	ec-3

J.1.3 Considerations for Audio Encoding

J.1.3.1 Overview

Synchronized Enhanced AC-3 frames represent a contiguous audio stream where each audio frame represents an equal size block of PCM samples at a given sampling frequency.

Additionally, the Enhanced AC-3 format does not require external metadata to set up the decoder, as it is fully contained in that regard. Descriptor data is present, however, to provide information to the system without requiring access to the elementary stream, as the bitstream is often encrypted.

J.1.3.2 Priming and delay

The codec uses audio blocks of a fixed length of 256 samples, and a transform which applies over two audio blocks. To obtain the correct audio from a block, both blocks in the transform are needed, and hence both the prior encoded block and the current encoded block need to be decoded to output the first frame. This is sometimes called "priming" and may be signaled using the 'roll' sample group.

Thus, a full reconstruction of the first 256 audio samples is sometimes not possible since there is no previous access unit. If it is desired to achieve full reconstruction of these samples, it is possible to add silence to the beginning of the audio signal.

In practice, an encoder might prepend an arbitrary amount of silent audio waveform samples to the signal. This portion of the audio signal is sometimes called "encoder delay" and varies depending on the implementation. This can be compensated using one of the following delay compensation approaches.

J.1.3.3 Delay compensation using an edit list

An edit list can be used to compensate any codec delay. When an edit list is used, a single EditListBox shall be recorded in the CMAF Header, as specified in MPEG CMAF [1], section 7.5.13.

J.1.3.4 Delay compensation before encapsulation

Delay can be compensated before stream encapsulation using the techniques specified in MPEG CMAF [1], annex G.5. No EditListBox is required and thus no further delay compensation is required in the receiver.

J.1.3.5 Loudness and dynamic range control (DRC)

The audio stream should contain DRC and loudness metadata according to annex E and annex H. The audio encoder should include the EMDF programme loudness data according to clause H.3.1. The audio encoder should include an appropriate dynamic range compression word according to clause 4.3.3.

J.1.4 Track constraints

J.1.4.1 Storage of ISOBMFF samples

Storage of bitstream access units within a CMAF Track shall conform to the CMAF audio Track format as specified in MPEG CMAF [1], section 10.2. Storage shall also conform to clause F.2. The following additional constraints apply:

- An ISOBMFF sample shall consist of either:
  - exactly one AC-3 syncframe; or
  - the number of Enhanced AC-3 syncframes required to deliver six blocks of audio data from every substream present, beginning with independent substream 0.
- The first syncframe of an Enhanced AC-3 sample shall have:
  - a stream type value of 0 (independent); and
  - a substream ID value of 0; and
  - the convsync flag set to 1.

J.1.4.2 ISOBMFF sample entry box

The ISOBMFF sample entry box is specified in annex F.

J.1.5 Elementary stream constraints

J.1.5.1 General

The bitstream shall be constrained as specified in clause F.1. The following additional constraints apply:

- The syncword parameter shall be 0x0B77, indicating "big endian" byte order.
- The value of the acmod parameter shall not be 0 (1+1) in any substream.
- The values of the parameters bsid, bsmod, acmod, lfeon, fscod, frmsizcod, and chanmap shall remain constant over the duration of one CMAF Track.

J.1.5.2 Enhanced AC-3 elementary stream constraints

J.1.5.2.1 General

The following additional constraints apply to Enhanced AC-3 elementary streams:

- The value of the numblkscod parameter shall be the same for all substreams, indicating the number of blocks per frame.
- The value of the strmtyp parameter shall be neither 2 (transcoded) nor 3 (reserved).
- The maximum number of substreams shall be 2, where the first substream shall be an independent substream and, if present, the second substream shall be a dependent substream.

J.1.5.2.2 Independent substream 0 constraints

Independent substream 0 consists of a sequence of Enhanced AC-3 synchronization frames. These synchronization frames shall comply with the following constraints:

- bsid - bit-stream identification: If independent substream 0 is the only substream in the bitstream, this field shall be set to 10000b (16). If the Enhanced AC-3 bitstream contains both independent substream 0 and dependent substream 0, this field shall be set to 00110 (6), 01000 (8) or 10000 (16).
- When bsid is 10000b (16), then:
  - strmtyp - stream type: This field shall be set to 00b (Stream Type 0 - independent substream); and
  - substreamid - substream identification: This field shall be set to 000b (substream ID = 0).

J.1.5.2.3 Dependent substream 0 constraints

Dependent substream 0 consists of a sequence of Enhanced AC-3 synchronization frames. These synchronization frames shall comply with the following constraints:

- bsid - bit-stream identification: This field shall be set to 10000b (16).
- strmtyp - stream type: This field shall be set to 01b (Stream Type 1 - dependent substream).
- substreamid - substream identification: This field shall be set to 000b (substream ID = 0).

J.2 Switching Sets

Since all ISOBMFF samples consist of exactly one complete Access Unit (see Annex F), all ISOBMFF samples are Stream Access Points (SAPs) of SAP type 1.

If a Switching Set contains more than one CMAF Track, the following requirements apply:

For time alignment between CMAF Tracks, either an encoder shall time align each SAP across all concerned CMAF Tracks, or proper use of Edit Lists shall be employed, such that the decoded PCM samples align.

NOTE 1: Sample presentation times have to align with the tfdt, since that time is used for the switching alignment.

NOTE 2: Switching Sets can contain one or more CMAF Tracks.

To enable seamless switching between CMAF Tracks in a CMAF Switching Set, the following bitstream parameter, if present, shall have identical values in all CMAF Tracks in the same CMAF Switching Set:

- bsid
- bsmod
- acmod

- lfeon
- fscod
- chanmap

J.3 Core Audio Media Profile

The Core CMAF Media Profile shall conform to the CMAF Track Format specified in clause J.1.

The following additional constraints apply.

- The fscod parameter shall always be 0, indicating 48 kHz.
- The bitrate as calculated in accordance with clause F.6.2.2 shall be less than or equal to 3,024 kbps.

The FileTypeBox compatibility brand shall be 'ceac' and should be used to indicate CMAF Tracks that conform to this Media Profile.

Annex K (informative): Bibliography

- Todd, C. et. al.: "AC-3: Flexible Perceptual Coding for Audio Transmission and Storage", AES 96th Convention, Preprint 3796, February 1994.
- Fielder L. D., Bosi M. A., Davidson G. A., Davis M. F., Todd C. and Vernon S.: "AC-2 and AC-3: Low-Complexity Transform-Based Audio Coding", Collected Papers on Digital Audio Bit-Rate Reduction, Neil Gilchrist and Christer Grewin, Eds. (Audio Eng. Soc., New York, NY, 1996), pp. 54-72.
- Davidson G. A.: "The Digital Signal Processing Handbook", Madisetti V. K. and Williams D. B. Eds. (CRC Press LLC, 1997), pp. 41-1 - 41-21.
- Princen J., Bradley A.: "Analysis/synthesis filter bank design based on time domain aliasing cancellation", IEEE Trans. Acoust. Speech and Signal Processing, vol. ASSP-34, pp. 1153-1161, October 1986.
- Davidson G. A, Fielder L. D., Link B. D.: "Parametric Bit Allocation in Perceptual Audio Coder", Presented at the 97th Convention of the Audio Engineering Society, Preprint 3921, November 1994.
- Vernon Steve: "Dolby Digital: Audio Coding for Digital Television and Storage Applications", Presented at the AES 17th International Conference: High-Quality Audio Coding; August 1999.
- Vernon Steve; Fruchter Vlad; Kusevitzky Sergio: "A Single-Chip DSP Implementation of a High-Quality Low Bit-Rate Multichannel Audio Coder", Presented at the 95th Convention of the Audio Engineering Society, Preprint 3775, Sept. 1993.
- R. Rao, P. Yip: "Discrete Cosine Transform", Academic Press, Boston 1990, pp. 11.
- Cover T. M., Thomas J. A.: "Elements of Information Theory", Wiley Series in Telecommunications, New York, 1991, pp. 13.
- Gersho A., Gray R. M.: "Vector Quantization and Signal Compression", Kluwer Academic Publisher, Boston, 1992, pp. 309.
- Truman M. M., Davidson G. A., Ubale A., Fielder L D.: "Efficient Bit Allocation, Quantization, and Coding in an Audio Distribution System", Presented at the 107th Audio Engineering Convention, Preprint 5068, August 1999.
- Fielder Louis D. and Davidson Grant A.: "Audio Coding Tools for Digital Television Distribution", Presented at the 108th Audio Engineering Convention, Preprint 5104, January 2000.
- Crockett B.: "High Quality Multi-Channel Time-Scaling and Pitch-Shifting using Auditory Scene Analysis," Presented at the 115th Audio Engineering Convention, Preprint 5948, October 2003.
- Crockett B.: "Improved Transient Pre-Noise Performance of Low Bit Rate Audio Coders Using Time Scaling Synthesis", Presented at the 117th Audio Engineering Convention, October 2004.
- Fielder L. D., Andersen, R. L., Crockett B. G., Davidson G. A., Davis M. F., Turner S. C., Vinton M. S., and Williams P. A.: "Introduction to Dolby Digital Plus, an Enhancement to the Dolby Digital Coding System", Presented at the 117th Audio Engineering Convention, October 2004.

History

Document history		
V1.1.1	February 2005	Publication
V1.2.1	August 2008	Publication
V1.3.1	August 2014	Publication
V1.4.1	September 2017	Publication