

Title: HM Software Manual
Status: Software AHG working document
Purpose: Information
Author(s): Frank Bossen
David Flynn
Karl Sharman
Karsten Sühning
Source: AHG chairs

frank@bossentech.com
dflynn@blackberry.com
karl.sharman@eu.sony.com
karsten.suehring@hhi.fraunhofer.de

Abstract

This document is a user manual describing usage of reference software for the HEVC project. It applies to version 16.13 of the software.

Contents

1	General Information	2
2	Installation and compilation	2
3	Using the encoder	2
3.1	GOP structure table	3
3.2	Encoder parameters	7
3.3	Encoder SEI parameters	20
3.4	Hardcoded encoder parameters	29
4	Using the decoder	30
4.1	General	30
4.2	Using the decoder analyser	31

List of Tables

1	Available project files	3
2	GOP structure example	5
3	File, I/O and source parameters.	7
4	Profile and level parameters	9
5	Unit definition parameters	10
6	Coding structure parameters	10
7	Motion estimation parameters	11
8	Mode decision parameters	11
9	Quantization parameters	12
10	Slice coding parameters	14
11	Deblocking filter parameters	15
12	Coding tools parameters	16
13	Rate control parameters	17
14	VUI parameters	18
15	Range Extensions (Version 2) tool parameters	19
16	List of Version 1 and RExt SEI messages	21
17	Buffering period SEI message encoder parameters	22
18	Picture timing SEI message encoder parameters	22
19	Recovery point SEI message encoder parameters	22

20	Tone mapping information SEI message encoder parameters	22
21	Frame packing arrangement SEI message encoder parameters	24
22	Display orientation SEI message encoder parameters	24
23	Green Metadata SEI message encoder parameters	25
24	Structure of pictures information SEI message encoder parameters	25
25	Active parameter sets SEI message encoder parameters	25
26	Decoding unit information SEI message encoder parameters	25
27	Temporal sub-layer zero index SEI message encoder parameters	25
28	Decoded picture hash SEI message encoder parameters	26
29	Scalable nesting SEI message encoder parameters	26
30	Region refresh information SEI message encoder parameters	26
31	No display SEI message encoder parameters	26
32	Time code SEI message encoder parameters	26
33	Mastering display colour volume SEI message encoder parameters	27
34	Segmented rectangular frame packing arrangement SEI message encoder parameters	27
35	Temporal motion-constrained tile sets SEI message encoder parameters	28
36	Chroma resampling filter hint SEI message encoder parameters	28
37	Knee function SEI message encoder parameters	28
38	Colour remapping SEI message encoder parameters	29
39	CommonDef.h constants	29
40	Decoder options	30

1 General Information

Reference software is being made available to provide a reference implementation of the HEVC standard being developed by the Joint Collaborative Team on Video Coding (JCT-VC) regrouping experts from ITU-T SG 16 and ISO/IEC SC29 WG11. One of the main goals of the reference software is to provide a basis upon which to conduct experiments in order to determine which coding tools provide desired coding performance. It is not meant to be a particularly efficient implementation of anything, and one may notice its apparent unsuitability for a particular use. It should not be construed to be a reflection of how complex a production-quality implementation of a future HEVC standard would be.

This document aims to provide guidance on the usage of the reference software. It is widely suspected to be incomplete and suggestions for improvements are welcome. Such suggestions and general inquiries may be sent to the general JCT-VC email reflector on jct-vc@lists.rwth-aachen.de (registration required).

Bug reporting

Bugs should be reported on the issue tracker set up at <http://hevc.kw.bbc.co.uk/trac/>

2 Installation and compilation

The software may be retrieved from one of the following SVN servers (mirrored):

- https://hevc.hhi.fraunhofer.de/svn/svn_HEVCSoftware/
- <svn://hevc.kw.bbc.co.uk/svn/jctvc-hm/>

Table 1 enumerates various project files that are provided for development environments.

For encoding large picture sizes (like UHD TV) it is strongly advised to build 64-bit binaries and to use a 64-bit OS. This will allow the software to use more than 2GB of RAM.

3 Using the encoder

```
TAppEncoder [--help] [-c config.cfg] [--parameter=value]
```

Table 1: Available project files

Environment	Location of project file
MS Visual Studio 2008 (VC9)	build/HM_vc9.sln
MS Visual Studio 2010 (VC10)	build/HM_vc2010.sln
MS Visual Studio 2012 (VC11)	build/HM_vc2012.sln
MS Visual Studio 2013 (VC12)	build/HM_vc2013.sln
Xcode	HM.xcodeproj
Eclipse	.project .cproject
make/gcc (e.g. Linux)	build/linux/makefile

Option	Description
<code>--help</code>	Prints parameter usage.
<code>-c</code>	Defines configuration file to use. Multiple configuration files may be used with repeated <code>-c</code> options.
<code>--parameter=value</code>	Assigns value to a given parameter as further described below. Some parameters are also supported by shorthand " <code>-opt value</code> ". These are shown in brackets after the parameter name in the tables of this document

Sample configuration files are provided in the `cfg/` folder. Parameters are defined by the last value encountered on the command line. Therefore if a setting is set via a configuration file, and then a subsequent command line parameter changes that same setting, the command line parameter value will be used.

3.1 GOP structure table

Defines the cyclic GOP structure that will be used repeatedly throughout the sequence. The table should contain `GOPSize` lines, named `Frame1`, `Frame2`, etc. The frames are listed in decoding order, so `Frame1` is the first frame in decoding order, `Frame2` is the second and so on. Among other things, the table specifies all reference pictures kept by the decoder for each frame. This includes pictures that are used for reference for the current picture as well as pictures that will be used for reference in the future. The encoder will not automatically calculate which pictures have to be kept for future references, they must be specified. Note that some specified reference frames for pictures encoded in the very first GOP after an IDR frame might not be available. This is handled automatically by the encoder, so the reference pictures can be given in the GOP structure table as if there were infinitely many identical GOPs before the current one. Each line in the table contains the parameters used for the corresponding frame, separated by whitespace:

Type: Slice type, can be either I, P or B.

POC: Display order of the frame within a GOP, ranging from 1 to `GOPSize`.

QPOffset: QP offset is added to the QP parameter to set the final QP value to use for this frame.

QPOffsetModelOff: Offset parameter to a linear model to adjust final QP based on $QP + QPOffset$.

QPOffsetModelScale: Scale parameter to a linear model to adjust final QP based on $QP + QPOffset$.

SliceCbQPOffset: The slice-level Cb QP offset.

SliceCrQPOffset: The slice-level Cr QP offset.

QPFactor: Weight used during rate distortion optimization. Higher values mean lower quality and less bits. Typical range is between 0.3 and 1.

tcOffsetDiv2: In-loop deblocking filter parameter `tcOffsetDiv2` is added to the base parameter `LoopFilterTcOffset_div2` to set the final `tc_offset_div2` parameter for this picture signalled in the slice segment header. The final value of `tc_offset_div2` shall be an integer number in the range $-6..6$.

betaOffsetDiv2: In-loop deblocking filter parameter betaOffsetDiv2 is added to the base parameter LoopFilterBetaOffset_div2 to set the final beta_offset_div2 parameter for this picture signalled in the slice segment header. The final value of beta_offset_div2 shall be an integer number in the range $-6..6$.

temporal_id: Temporal layer of the frame. A frame cannot predict from a frame with a higher temporal id. If a frame with higher temporal IDs is listed among a frame's reference pictures, it is not used, but is kept for possible use in future frames.

num_ref_pics_active: Size of reference picture lists L0 and L1, indicating how many reference pictures in each direction that are used during coding.

num_ref_pics: The number of reference pictures kept for this frame. This includes pictures that are used for reference for the current picture as well as pictures that will be used for reference in the future.

reference_pictures: A space-separated list of num_ref_pics integers, specifying the POC of the reference pictures kept, relative the POC of the current frame. The picture list shall be ordered, first with negative numbers from largest to smallest, followed by positive numbers from smallest to largest (e.g. $-1 -3 -5 1 3$). Note that any pictures not supplied in this list will be discarded and therefore not available as reference pictures later.

predict: Defines the value of the syntax element inter_ref_pic_set_prediction_flag. A value of 0 indicates that the reference picture set is encoded without inter RPS prediction and the subsequent parameters deltaRIdx-1, deltaRPS, num_ref_idcs and Reference_idcs are ignored and do not need to be present. A value of 1 indicates that the reference picture set is encoded with inter prediction RPS using the subsequent parameters deltaRIdx-1, deltaRPS, num_ref_idcs and Reference_idcs in the line. A value of 2 indicates that the reference picture set is encoded with inter RPS but only the deltaRIdx-1 parameters is needed. The deltaRPS, num_ref_idcs and Reference_idcs values are automatically derived by the encoder based on the POC and refPic values of the current line and the RPS pointed to by the deltaRIdx-1 parameters.

deltaRIdx-1: The difference between the index of the current RPS and the predictor RPS minus 1.

deltaRPS: The difference between the POC of the predictor RPS and POC the current RPS.


num_ref_idcs: The number of ref_idcs to encode for the current RPS. The value is equal to the value of num_ref_pics of the predictor RPS plus 1.

reference_idcs: A space-separated list of num_ref_idcs integers, specifying the ref_idcs of the inter RPS prediction. The value of ref_idcs may be 0, 1 or 2 indicating that the reference picture is a reference picture used by the current picture, a reference picture used for future picture or not a reference picture anymore, respectively. The first num_ref_pics of ref_idcs correspond to the Reference pictures in the predictor RPS. The last ref_idcs corresponds to the predictor picture.

For example, consider the coding structure of Figure 1. This coding structure is of size 4. The pictures are listed in decoding order. Frame1 shall therefore describe picture with POC = 4. It references picture 0, and therefore has -4 as a reference picture. Similarly, Frame2 has a POC of 2, and since it references pictures 0 and 4, its reference pictures are listed as $-2 2$. Frame3 is a special case: even though it only references pictures with POC 0 and 2, it also needs to include the picture with POC 4, which must be kept in order to be used as a reference picture in the future. The reference picture list for Frame3 therefore becomes $-1 1 3$. Frame4 has a POC of 3 and its list of reference pictures is $-1 1$.

Inter RPS prediction may be used for Frame2, Frame3 and Frame4, hence the predict parameter is set to 1 for these frames. Frame2 uses Frame1 as the predictor hence the deltaRIdx-1 is 0. Similarly for Frame3 and Frame4 which use Frame2 and Frame3 as predictors, respectively. The deltaRPS is equal to the POC of the predictor minus the POC of the current picture, therefore the deltaRPS for Frame2 is $4 - 2 = 2$, for Frame3 is $2 - 1 = 1$ and for Frame4 is $1 - 3 = -2$.

Figure 1: A GOP structure



In Frame2, reference pictures with POC 0 and 2 are used, so the reference ids for Frame2 are 1 1 indicating that the reference picture, -4, in Frame1 is still a reference picture in Frame2 and Frame1 is also a reference picture in Frame2. The reference ids for Frame3 are 1 1 1. The first and second “1”s indicating that the reference pictures “-2 2” in Frame2 are still reference pictures in Frame3 and the last “1” indicating that Frame2 is also a reference picture in Frame3. In Frame 4, the reference ids are 0 1 1 0. The first “0” indicates that the reference pictures “-1” in Frame 3 is no longer a reference picture in Frame4. The next two “1”s indicate that the reference pictures “1 3” are now reference pictures of Frame4. The final “0” indicates that Frame3 is not a reference picture.

In order to specify this to the encoder, the parameters in Table 2 could be used.

Table 2: GOP structure example

	Frame1	Frame2	Frame3	Frame4
Type	P	B	B	B
POC	4	2	1	3
QPOffset	1	2	3	3
QPOffsetModelOff	0.0	0.0	0.0	0.0
QPOffsetModelScale	0.0	0.0	0.0	0.0
SliceCbQPOffset	0	0	0	0
SliceCrQPOffset	0	0	0	0
QPfactor	0.5	0.5	0.5	0.5
tcOffsetDiv2	0	1	2	2
betaOffsetDiv2	0	0	0	0
temporal_id	0	1	2	2
num_ref_pics_active	1	1	1	1
num_ref_pics	1	2	3	2
reference_pictures	-4	-2 2	-1 1 3	-1 1
predict	0	1	1	1
deltaRIdx-1		0	0	0
deltaRPS		2	1	-2
num_ref_ids		2	3	4
reference_ids		1 1	1 1 1	0 1 1 0

Here, the frames used for prediction have been given higher quality by assigning a lower QP offset. Also, the non-reference frames have been marked as belonging to a higher temporal layer, to make it possible to decode only every other frame. Note: each line should contain information for one frame, so this configuration would be specified as:

```

Frame1: P 4 1 0 0 0.5 0 0 0 1 1 -4 0
Frame2: B 2 2 0 0 0.5 1 0 1 1 2 -2 2 1 0 2 2 1 1
Frame3: B 1 3 0 0 0.5 2 0 2 1 3 -1 1 3 1 0 1 3 1 1 1
Frame4: B 3 3 0 0 0.5 2 0 2 1 2 -1 1 1 0 -2 4 0 1 1 0
    
```

The values of deltaRIdx-1, deltaRPS, num_ref_ids and reference_ids of Frame K can be derived from the POC value of Frame K and the POC, num_ref_pics and reference_pictures values of Frame M , where K is the index of the RPS to be inter coded and the M is the index of the reference RPS, as follows.

Note: The above (automatic) generation of the inter RPS parameter values has been integrated into the encoder, and is activated by the value of predict = 2 followed by the value of deltaRIdx-1, only, as

```

deltaRIdxK - 1 ← K - M - 1 ;
deltaRPSK ← POCM - POCK ;
num_ref_idcsK ← num_ref_picsM + 1 ;
for j ← 0 to num_ref_picsM do
    for i ← 0 to num_ref_idcsK do
        if reference_picturesM,j + deltaRPSK == reference_picturesK,i then
            if reference_picturesK,i is used by the current frame then reference_idcsK,j = 1;
            ;
            else reference_idcsK,j = 2;
            ;
        else
            reference_idcsK[j] = 0 ;
        end
    end
end
/* reference_picturesM,num_ref_picsM does not exist and is assumed to be 0 */

```

described above.

3.2 Encoder parameters

Shorthand alternatives for the parameter that can be used on the command line are shown in brackets after the parameter name.

Table 3: File, I/O and source parameters.

Option	Default	Description
InputFile (-i)		Specifies the input video file. Video data must be in a raw 4:2:0, or 4:2:2 planar format, 4:4:4 planar format (Y'CbCr, RGB or GBR), or in a raw 4:0:0 format. Note: When the bit depth of samples is larger than 8, each sample is encoded in 2 bytes (little endian, LSB-justified).
BitstreamFile (-b)		Specifies the output coded bit stream file.
ReconFile (-o)		Specifies the output locally reconstructed video file.
SourceWidth (-wdt)	0	Specifies the width and height of the input video in luma samples.
SourceHeight (-hgt)	0	
InputBitDepth	8	Specifies the bit depth of the input video.
MSBExtendedBitDepth	0	Extends the input video by adding MSBs of value 0. When 0, no extension is applied and the InputBitDepth is used. The MSBExtendedBitDepth becomes the effective file InputBitDepth for subsequent processing.
InternalBitDepth	0	Specifies the bit depth used for coding. When 0, the setting defaults to the value of the MSBExtendedBitDepth. If the input video is a different bit depth to InternalBitDepth, it is automatically converted by: <div style="text-align: center;"> $\left\lfloor \frac{PeI * 2^{\text{InternalBitDepth}}}{2^{\text{MSBExtendedBitDepth}}} \right\rfloor$ </div> Note: The effect of this option is as if the input video is externally converted to the MSBExtendedBitDepth and then to the InternalBitDepth and then coded with this value as InputBitDepth. The codec has no notion of different bit depths.
OutputBitDepth	0	Specifies the bit depth of the output locally reconstructed video file. When 0, the setting defaults to the value of InternalBitDepth. Note: This option has no effect on the decoding process.
InputBitDepthC	0	Specifies the various bit-depths for chroma components. These only need to be specified if non-equal luma and chroma bit-depth processing is required. When 0, the setting defaults to the corresponding non-Chroma value.
MSBExtendedBitDepthC	0	
InternalBitDepthC	0	
OutputBitDepthC	0	
InputColourSpaceConvert		The colour space conversion to apply to input video. Permitted values are: UNCHANGED No colour space conversion is applied YCbCrToYCrCb Swap the second and third components YCbCrtoYYY Set the second and third components to the values in the first RGBtoGBR Reorder the three components If no value is specified, no colour space conversion is applied. The list may eventually also include RGB to YCbCr or YCgCo conversions.
SNRInternalColourSpace	false	When this is set true, then no colour space conversion is applied prior to PSNR calculation, otherwise the inverse of InputColourSpaceConvert is applied.
OutputInternalColourSpace	false	When this is set true, then no colour space conversion is applied to the reconstructed video, otherwise the inverse of InputColourSpaceConvert is applied.

Continued...

Table 3: File, I/O and source parameters. (Continued)

Option	Default	Description
InputChromaFormat	420	Specifies the chroma format used in the input file. Permitted values (depending on the profile) are 400, 420, 422 or 444.
ChromaFormatIDC (-cf)	0	Specifies the chroma format to use for processing. Permitted values (depending on the profile) are 400, 420, 422 or 444; the value of 0 indicates that the value of InputChromaFormat should be used instead.
MSEBasedSequencePSNR	false	When 0, the PSNR output is a linear average of the frame PSNRs; when 1, additional PSNRs are output which are formed from the average MSE of all the frames. The latter is useful when coding near-losslessly, where occasional frames become lossless.
PrintFrameMSE	false	When 1, the Mean Square Error (MSE) values of each frame will also be output alongside the default PSNR values.
PrintSequenceMSE	false	When 1, the Mean Square Error (MSE) values of the entire sequence will also be output alongside the default PSNR values.
SummaryOutFilename	false	Filename to use for producing summary output file. If empty, do not produce a file.
SummaryPicFilenameBase	false	Base filename to use for producing summary picture output files. The actual filenames used will have I.txt, P.txt and B.txt appended. If empty, do not produce a file.
SummaryVerboseness	false	Specifies the level of the verbosity of the text output.
CabacZeroWordPaddingEnabled	false	When 1, CABAC zero word padding will be enabled. This is currently not the default value for the setting.
ConformanceWindowMode	0	Specifies how the parameters related to the conformance window are interpreted (cropping/padding). The following modes are available: <ul style="list-style-type: none"> 0 No cropping / padding 1 Automatic padding to the next minimum CU size 2 Padding according to parameters HorizontalPadding and VerticalPadding 3 Cropping according to parameters ConfWinLeft, ConfWinRight, ConfWinTop and ConfWinBottom
HorizontalPadding (-pdx) VerticalPadding (-pdy)	0	Specifies the horizontal and vertical padding to be applied to the input video in luma samples when ConformanceWindowMode is 2. Must be a multiple of the chroma resolution (e.g. a multiple of two for 4:2:0).
ConfWinLeft ConfWinRight ConfWinTop ConfWinBottom	0	Specifies the horizontal and vertical cropping to be applied to the input video in luma samples when ConformanceWindowMode is 3. Must be a multiple of the chroma resolution (e.g. a multiple of two for 4:2:0).
FrameRate (-fr)	0	Specifies the frame rate of the input video. Note: This option only affects the reported bit rates.
FrameSkip (-fs)	0	Specifies a number of frames to skip at beginning of input video file.
FramesToBeEncoded (-f)	0	Specifies the number of frames to be encoded (see note regarding TemporalSubsampleRatio). When 0, all frames are coded.
TemporalSubsampleRatio (-ts)	1	Temporally subsamples the input video sequence. A value of N will skip $(N - 1)$ frames of input video after each coded input video frame. Note the FramesToBeEncoded does not account for the temporal skipping of frames, which will reduce the number of frames encoded accordingly. The reported bit rates will be reduced and VUI information is scaled so as to present the video at the correct speed. The minimum and default value is 1.

Continued...

Table 3: File, I/O and source parameters. (Continued)

Option	Default	Description
FieldCoding	false	When 1, indicates that field-based coding is to be applied.
TopFieldFirst (-Tff)	0	Indicates the order of the fields packed into the input frame. When 1, the top field is temporally first.
ClipInputVideoToRec709Range	0	If 1 then clip input video to the Rec. 709 Range on loading when InternalBitDepth is less than MSBExtendedBitDepth.
ClipOutputVideoToRec709Range	0	If 1 then clip output video to the Rec. 709 Range on saving when OutputBitDepth is less than InternalBitDepth.
EfficientFieldIRAPEnabled	1	Enable to code fields in a specific, potentially more efficient, order.
HarmonizeGopFirstFieldCoupleEnabled	1	Enables harmonization of Gop first field couple.
AccessUnitDelimiter	0	Add Access Unit Delimiter NAL units between all Access Units.

Table 4: Profile and level parameters

Option	Default	Description
Profile	none	Specifies the profile to which the encoded bitstream complies. Valid HEVC Ver. 1 values are: none, main, main10, main-still-picture Valid HEVC Ver. 2 (RExt) values are: main-RExt, high-throughput-RExt, monochrome, monochrome12, monochrome16, main12, main_422_10, main_422_12, main_444, main_444_10, main_444_12, main_444_16, main_intra, main_10_intra, main_12_intra, main_422_10_intra, main_422_12_intra, main_444_intra, main_444_10_intra, main_444_12_intra, main_444_16_intra. When main-RExt is specified, the constraint flags are either manually specified, or calculated via the other supplied settings. Compatibility flags are automatically determined according to the profile. NB: There is currently only limited validation that the encoder configuration complies with the profile, level and tier constraints.
Level	none	Specifies the level to which the encoded bitstream complies. Valid values are: none, 1, 2, 2.1, 3, 3.1, 4, 4.1, 5, 5.1, 5.2, 6, 6.1, 6.2, 8.5 NB: There is currently only limited validation that the encoder configuration complies with the profile, level and tier constraints.
Tier	main	Specifies the level tier to which the encoded bitstream complies. Valid values are: main, high. NB: There is currently only limited validation that the encoder configuration complies with the profile, level and tier constraints.
MaxBitDepthConstraint	0	For <code>-profile=main-RExt</code> , specifies the value to use to derive the <code>general_max_bit_depth</code> constraint flags for RExt profiles; when 0, use $\max(\text{InternalBitDepth}, \text{InternalBitDepthC})$
MaxChromaFormatConstraint	0	For <code>-profile=main-RExt</code> , specifies the chroma-format to use for the general profile constraints for RExt profiles; when 0, use the value of <code>ChromaFormatIDC</code> .
IntraConstraintFlag	false	For <code>-profile=main-RExt</code> , specifies the value of <code>general_intra_constraint_flag</code> to use for RExt profiles.
OnePictureOnlyConstraintFlag	false	For <code>-profile=main-RExt</code> , specifies the value of <code>general_one_picture_only_constraint_flag</code> to use for RExt profiles.

Continued...

Table 4: Profile and level parameters (Continued)

Option	Default	Description
LowerBitRateConstraintFlag	true	Specifies the value of <code>general_lower_bit_constraint_flag</code> to use for RExt profiles.
ProgressiveSource	false	Specifies the value of <code>general_progressive_source_flag</code>
InterlacedSource	false	Specifies the value of <code>general_interlaced_source_flag</code>
NonPackedSource	false	Specifies the value of <code>general_non_packed_constraint_flag</code>
FrameOnly	false	Specifies the value of <code>general_frame_only_constraint_flag</code>

Table 5: Unit definition parameters

Option	Default	Description
MaxCUWidth	64	Defines the maximum CU width.
MaxCUHeight	64	Defines the maximum CU height.
MaxCUSize (-s)	64	Defines the maximum CU size.
MaxPartitionDepth (-h)	4	Defines the depth of the CU tree.
QuadtreeTULog2MaxSize	6 (= $\log_2(64)$)	Defines the Maximum TU size in logarithm base 2.
QuadtreeTULog2MinSize	2 (= $\log_2(4)$)	Defines the Minimum TU size in logarithm base 2.
QuadtreeTUMaxDepthIntra	1	Defines the depth of the TU tree for intra CUs.
QuadtreeTUMaxDepthInter	2	Defines the depth of the TU tree for inter CUs.

Table 6: Coding structure parameters

Option	Default	Description
IntraPeriod (-ip)	-1	Specifies the intra frame period. A value of -1 implies an infinite period.
DecodingRefreshType (-dr)	0	Specifies the type of decoding refresh to apply at the intra frame period picture. <ul style="list-style-type: none"> 0 Applies an I picture (not a intra random access point). 1 Applies a CRA intra random access point (open GOP). 2 Applies an IDR intra random access point (closed GOP). 3 Use recovery point SEI messages to indicate random access.
GOPSize (-g)	1	Specifies the size of the cyclic GOP structure.
FrameN		Multiple options that define the cyclic GOP structure that will be used repeatedly throughout the sequence. The table should contain GOPSize elements. See section 3.1 for further details.

Table 7: Motion estimation parameters

Option	Default	Description
FastSearch	1	Enables or disables the use of a fast motion search. 0 Full search method 1 Fast search method - TZSearch 2 Predictive motion vector fast search method 3 Extended TZSearch method
SearchRange (-sr)	96	Specifies the search range used for motion estimation. Note: the search range is defined around a predictor. Motion vectors derived by the motion estimation may thus have values larger than the search range.
BipredSearchRange	4	Specifies the search range used for bi-prediction refinement in motion estimation.
ClipForBiPredMEEnabled	0	Enables clipping in the Bi-Pred ME, which prevents values over- or under-flowing. It is usually disabled to reduce encoder run-time.
FastMEAssumingSmootherMVEnabled	0	Enables fast ME assuming a smoother MV.
HadamardME	true	Enables or disables the use of the Hadamard transform in fractional-pel motion estimation. 0 SAD for cost estimation 1 Hadamard for cost estimation
ASR	false	Enables or disables the use of adaptive search ranges, where the motion search range is dynamically adjusted according to the POC difference between the current and the reference pictures. $\text{SearchRange}' = \text{Round} \left(\text{SearchRange} * \text{ADAPT_SR_SCALE} * \frac{\text{abs}(\text{POC}_{\text{cur}} - \text{POC}_{\text{ref}})}{\text{RateGOPSize}} \right)$
MaxNumMergeCand	5	Specifies the maximum number of merge candidates to use.
DisableIntraInInter	0	Flag to disable intra PUs in inter slices.

Table 8: Mode decision parameters

Option	Default	Description
LambdaModifierN (-LMN)	1.0	Specifies a value that is multiplied with the Lagrange multiplier λ , for use in the rate-distortion optimised cost calculation when encoding temporal layer N . If LambdaModifierI is specified, then LambdaModifierI will be used for intra pictures. N may be in the range 0 (inclusive) to 7 (exclusive).
LambdaModifierI (-LMI)		Specifies one or more of the LambdaModifiers to use intra pictures at each of the temporal layers. If not present, then the LambdaModifierN settings are used instead. If the list of values (comma or space separated) does not include enough values for each of the temporal layers, the last value is repeated as required.
IQPFactor (-IQF)	-1	Specifies the QP factor to be used for intra pictures during the lambda computation. (The values specified in the GOP structure are only used for inter pictures). If negative (default), the following equation is used to derive the value: $IQP_{factor} = 0.57 * (1.0 - \text{Max}(0.5, \text{Min}(0.0, 0.05 * s)))$ where $s = \text{Int}(\text{isField?}(GS - 1)/2 : GS - 1)$ and GS is the gop size.
ECU	false	Enables or disables the use of early CU determination. When enabled, skipped CUs will not be split further.

Continued...

Table 8: Mode decision parameters (Continued)

Option	Default	Description
CFM	false	Enables or disables the use of Cbf-based fast encoder mode. When enabled, once a 2Nx2N CU has been evaluated, if the RootCbf is 0, further PU splits will not be evaluated.
ESD	false	Enables or disables the use of early skip detection. When enabled, the skip mode will be tested before any other.
FEN	0	Controls the use of different fast encoder coding tools. The following tools are supported in different combinations: <ul style="list-style-type: none"> a In the SAD computation for blocks having size larger than 8, only the lines of even rows in the block are considered. b The number of iterations used in the bi-directional motion vector refinement in the motion estimation process is reduced from 4 to 1. Depending on the value of the parameter, the following combinations are supported: <ul style="list-style-type: none"> 0 Disable all modes 1 Use both a & b tools 2 Use only tool b 3 Use only tool a
FDM	true	Enables or disables the use of fast encoder decisions for 2Nx2N merge mode. When enabled, the RD cost for the merge mode of the current candidate is not evaluated if the merge skip mode was the best merge mode for one of the previous candidates.
RDpenalty	0	RD-penalty for 32x32 TU for intra in non-intra slices. Enabling this parameter can reduce the visibility of CU boundaries in the coded picture. <ul style="list-style-type: none"> 0 No RD-penalty 1 RD-penalty 2 Maximum RD-penalty (no 32x32 TU)

Table 9: Quantization parameters

Option	Default	Description
QP (-q)	30.0	Specifies the base value of the quantization parameter. If it is non-integer, the QP is switched once during encoding.
IntraQPOffset	0	Specifies a QP offset from the base QP value to be used for intra frames.
LambdaFromQpEnable	false	When enabled, the λ , which is used to convert a cost in bits to a cost in distortion terms, is calculated as: $\lambda = qpFactor \times 2^{qp+6*(bitDepthLuma-8)-12}$, where qp is the slice QP and $qpFactor$ is calculated as follows: <ul style="list-style-type: none"> = IQF if $IQF \geq 0$ and slice is a periodic intra slice = $0.57 \times \lambda_{scale}$ if slice is a non-periodic intra slice = value from GOP table otherwise where IQF is the value specified using the IntraQPFactor option, and where λ_{scale} is: <ul style="list-style-type: none"> 1 if LambdaFromQpEnable=true $1.0 - \max(0, \min(0.5, 0.05 * B))$ if LambdaFromQpEnable=false where B is the number of B frames. If LambdaFromQpEnable=false, then the λ is also subsequently scaled for non-top-level hierarchical depths, as follows: $\lambda = \lambda_{base} \times \max(2, \min(4, (sliceQP - 12)/6))$ In addition, independent on the IntraQPFactor, if HadamardME=false, then for an inter slice the final λ is scaled by a factor of 0.95.

Continued...

Table 9: Quantization parameters (Continued)

Option	Default	Description
CbQpOffset (-cbqpofs) CrQpOffset (-crqpofs)	0	Global offset to apply to the luma QP to derive the QP of Cb and Cr respectively. These options correspond to the values of <code>cb_qp_offset</code> and <code>cr_qp_offset</code> , that are transmitted in the PPS. Valid values are in the range $[-12, 12]$.
LumaLevelToDeltaQPMode	0	Luma-level based Delta QP modulation. 0 not used 1 Based on CTU average 2 Based on Max luma in CTU
LumaLevelToDeltaQPMaxValWeight	1.0	Weight of per block maximum luma value when LumaLevelToDeltaQPMode=2.
LumaLevelToDeltaQPMappingLuma		Specify luma values to use for the luma to delta QP mapping instead of using default values. Default values are: 0, 301, 367, 434, 501, 567, 634, 701, 767, 834.
LumaLevelToDeltaQPMappingDQP		Specify DQP values to use for the luma to delta QP mapping instead of using default values. Default values are: -3, -2, -1, 0, 1, 2, 3, 4, 5, 6.
WCGPPSEnable	0	Enable the WCG PPS modulation of the chroma QP, rather than the slice, which, unlike slice-level modulation, allows the deblocking process to consider the adjustment. To use, specify a fractional QP: the first part of the sequence will use $qpc = \text{floor}(QP)$ in the following calculation and PPS-0; the second part of the sequence will use $qpc = \text{ceil}(QP)$ and PPS-1. The <i>chromaQp</i> that is then stored in the PPS is given as: $\text{clip}(\text{round}(\text{WCGPPSXXQpScale} * \text{baseCQp}) + \text{XXQpOffset})$ where $\text{baseCQp} = (\text{WCGPPSChromaQpScale} * qpc + \text{WCGPPSChromaQpOffset})$. Note that the slices will continue to have a delta QP applied.
WCGPPSChromaQpScale	0.0	Scale parameter for the linear chroma QP offset mapping used for WCG content.
WCGPPSChromaQpOffset	0.0	Offset parameter for the linear chroma QP offset mapping used for WCG content.
WCGPPScbQpScale WCGPPScrQpScale	1.0	Per chroma component QP scale factor depending on capture and representation color space. For Cb component with BT.2020 container use 1.14; for BT.709 material and 1.04 for P3 material. For Cr component with BT.2020 container use 1.79; for BT.709 material and 1.39 for P3 material.
SliceChromaQPOffsetPeriodicity	0	Defines the periodicity for inter slices that use the slice-level chroma QP offsets, as defined by <code>SliceCbQpOffsetIntraOrPeriodic</code> and <code>SliceCrQpOffsetIntraOrPeriodic</code> . A value of 0 disables the periodicity. It is intended to be used in low-delay configurations where an regular intra period is not defined.
SliceCbQpOffsetIntraOrPeriodic SliceCrQpOffsetIntraOrPeriodic	0	Defines the slice-level QP offset to be used for intra slices, or once every 'SliceChromaQPOffsetPeriodicity' pictures.
MaxCuDQPDepth (-dqd)	0	Defines maximum depth of a minimum CuDQP for sub-LCU-level delta QP. MaxCuDQPDepth shall be greater than or equal to <code>SliceGranularity</code> .
RDOQ	true	Enables or disables rate-distortion-optimized quantization for transformed TUs.
RDOQTS	true	Enables or disables rate-distortion-optimized quantization for transform-skipped TUs.

Continued...

Table 9: Quantization parameters (Continued)

Option	Default	Description
SelectiveRDOQ	false	Enables or disables selective rate-distortion-optimized quantization. A simple quantization is used to pre-analyze, whether to bypass the RDOQ process or not. If all the coefficients are quantized to 0, the RDOQ process is bypassed. Otherwise, the RDOQ process is performed as usual.
DeltaQpRD (-dqr)	0	Specifies the maximum QP offset at slice level for multi-pass slice encoding. When encoding, each slice is tested multiple times by using slice QP values in the range $[-\text{DeltaQpRD}, \text{DeltaQpRD}]$, and the best QP value is chosen as the slice QP.
MaxDeltaQP (-d)	0	Specifies the maximum QP offset at the largest coding unit level for the block-level adaptive QP assignment scheme. In the encoder, each largest coding unit is tested multiple times by using the QP values in the range $[-\text{MaxDeltaQP}, \text{MaxDeltaQP}]$, and the best QP value is chosen as the QP value of the largest coding unit.
dQPFile (-m)		Specifies a file containing a list of QP deltas. The n -th line (where n is 0 for the first line) of this file corresponds to the QP value delta for the picture with POC value n .
AdaptiveQp (-aq)	false	Enable or disable QP adaptation based upon a psycho-visual model.
MaxQPAdaptationRange (-aqr)	6	Specifies the maximum QP adaptation range.
AdaptiveQpSelection (-aqps)	false	Specifies whether QP values for non-I frames will be calculated on the fly based on statistics of previously coded frames.
RecalculateQP... AccordingToLambda	false	Recalculate QP values according to lambda values. Do not suggest to be enabled in all intra case.
ScalingList	0	Controls the specification of scaling lists: 0 Scaling lists are disabled 1 Use default scaling lists 2 Scaling lists are specified in the file indicated by ScalingListFile
ScalingListFile		When ScalingList is set to 2, this parameter indicates the name of the file, which contains the defined scaling lists. If ScalingList is set to 2 and this parameter is an empty string, information on the format of the scaling list file is output and the encoder stops.
MaxCUChromaQpAdjustmentDepth	-1	Specifies the maximum depth for CU chroma QP adjustment; if negative, CU chroma QP adjustment is disabled.

Table 10: Slice coding parameters

Option	Default	Description
SliceMode	0	Controls the slice partitioning method in conjunction with SliceArgument. 0 Single slice 1 Maximum number of CTUs per slice 2 Maximum number of bytes per slice 3 Maximum number of tiles per slice
SliceArgument		Specifies the maximum number of CTUs, bytes or tiles in a slice depending on the SliceMode setting.

Continued...

Table 10: Slice coding parameters (Continued)

Option	Default	Description
SliceSegmentMode	0	Enables (dependent) slice segment coding in conjunction with SliceSegmentArgument. 0 Single slice 1 Maximum number of CTUs per slice segment 2 Maximum number of bytes per slice segment 3 Maximum number of tiles per slice segment
SliceSegmentArgument		Defines the maximum number of CTUs, bytes or tiles a slice segment depending on the SliceSegmentMode setting.
WaveFrontSynchro	false	Enables the use of specific CABAC probabilities synchronization at the beginning of each line of CTBs in order to produce a bitstream that can be encoded or decoded using one or more cores.
TileUniformSpacing	false	Controls the mode used to determine per row and column tile sizes. 0 Each tile column width and tile row height is explicitly set by TileColumnWidthArray and TileRowHeightArray respectively 1 Tile columns and tile rows are uniformly spaced.
NumTileColumnsMinus1 NumTileRowsMinus1	0	Specifies the tile based picture partitioning geometry as $\text{NumTileColumnsMinus1} + 1 \times \text{NumTileRowsMinus1} + 1$ columns and rows.
TileColumnWidthArray TileRowHeightArray		Specifies a space or comma separated list of widths and heights, respectively, of each tile column or tile row. The first value in the list corresponds to the leftmost tile column or topmost tile row.

Table 11: Deblocking filter parameters

Option	Default	Description
LoopFilterDisable	false	Enables or disables the in-loop deblocking filter.
LFCrossSliceBoundaryFlag	true	Enables or disables the use of in-loop filtering across slice boundaries.
LoopFilterOffsetInPPS	false	If enabled, the in-loop deblocking filter control parameters are sent in PPS. Otherwise, the in-loop deblocking filter control parameters are sent in the slice segment header. If deblocking filter parameters are sent in PPS, the same values of deblocking filter parameters are used for all pictures in the sequence (i.e. deblocking parameter = base parameter value). If deblocking filter parameters are sent in the slice segment header, varying deblocking filter parameters can be specified by setting parameters tcOffsetDiv2 and betaOffsetDiv2 in the GOP structure table. In this case, the final value of the deblocking filter parameter sent for a certain GOP picture is equal to (base parameter + GOP parameter for this picture). Intra-pictures use the base parameters values.
LoopFilterTcOffset_div2	0	Specifies the base value for the in-loop deblocking filter parameter tc_offset_div2. The final value of tc_offset_div2 shall be an integer number in the range $-6..6$.
LoopFilterBetaOffset_div2	0	Specifies the base value for the in-loop deblocking filter parameter beta_offset_div2. The final value of beta_offset_div2 shall be an integer number in the range $-6..6$.
DeblockingFilterMetric	0	Specifies the use of a deblocking filter metric to evaluate the suitability of deblocking. If non-zero then LoopFilterOffsetInPPS and LoopFilterDisable must be 0. Currently excepted values are 0, 1 and 2.

Continued...

Table 11: Deblocking filter parameters (Continued)

Option	Default	Description
LFCrossSliceBoundaryFlag	true	Enables or disables the use of a deblocking across tile boundaries.

Table 12: Coding tools parameters

Option	Default	Description
AMP	true	Enables or disables the use of asymmetric motion partitions.
SAO	true	Enables or disables the sample adaptive offset (SAO) filter.
TestSAODisableAtPictureLevel	false	Enables the testing of disabling SAO at the picture level after having analysed all blocks.
SaoEncodingRate	0.75	When >0 SAO early picture termination is enabled for luma and chroma.
SaoEncodingRateChroma	0.5	The SAO early picture termination rate to use for chroma (when m_SaoEncodingRate is >0). If <=0, use results for luma.
SAOLcuBoundary	false	Enables or disables SAO parameter estimation using non-deblocked pixels for LCU bottom and right boundary areas.
SAOResetEncoderStateAfterIRAP	false	When true, resets the encoder's SAO state after an IRAP (POC order).
ConstrainedIntraPred	false	Enables or disables constrained intra prediction. Constrained intra prediction only permits samples from intra blocks in the same slice as the current block to be used for intra prediction.
FastUDIUseMPMEnabled	true	If enabled, adapt intra direction search, accounting for MPM
FastMEForGenBLowDelayEnabled	true	If enabled use a fast ME for generalised B Low Delay slices
UseBLambdaForNonKeyLowDelayPictures	true	Enables use of B-Lambda for non-key low-delay pictures
TransquantBypassEnable	false	Enables or disables the ability to bypass the transform, quantization and filtering stages at CU level. This option corresponds to the value of transquant_bypass_enabled_flag that is transmitted in the PPS. See CUTransquantBypassFlagForce for further details.
CUTransquantBypassFlagForce	0	Controls the per CU transformation, quantization and filtering mode decision. This option controls the value of the per CU cu_transquant_bypass_flag. 0 Bypass is searched on a CU-by-CU basis and will be used if the cost is lower than not bypassing. 1 Bypass is forced for all CUs. This option has no effect if TransquantBypassEnable is disabled.
PCMEnabledFlag	false	Enables or disables the use of PCM. The encoder will use cost measures on a CU-by-CU basis to determine if PCM mode is to be applied.
PCMLog2MaxSize	5 (= $\log_2(32)$)	Specifies log ₂ of the maximum PCM block size. When PCM is enabled, the PCM mode is available for 2Nx2N intra PUs smaller than or equal to the specified maximum PCM block size
PCMLog2MinSize	3	Specifies log ₂ of the minimum PCM block size. When PCM is enabled, the PCM mode is available for 2Nx2N intra PUs larger than or equal to the specified minimum PCM block size. When larger than PCMLog2MaxSize, PCM mode is not used.
PCMInputBitDepthFlag	true	If enabled specifies that PCM sample bit-depth is set equal to InputBitDepth. Otherwise, it specifies that PCM sample bit-depth is set equal to InternalBitDepth.

Continued...

Table 12: Coding tools parameters (Continued)

Option	Default	Description
PCMFilterDisableFlag	false	If enabled specifies that loop-filtering on reconstructed samples of PCM blocks is skipped. Otherwise, it specifies that loop-filtering on reconstructed samples of PCM blocks is not skipped.
WeightedPredP (-wpP)	false	Enables the use of weighted prediction in P slices.
WeightedPredB (-wpB)	false	Enables the use of weighted prediction in B slices.
WPMMethod (-wpM)	0	Sets the Weighted Prediction method to be used. 0 Image DC based method with joint colour component decision. 1 Image DC based method with separate colour component decision. 2 DC + Histogram refinement method (no clipping). 3 DC + Histogram refinement method (with clipping). 4 DC + Dual Histogram refinement method (with clipping).
Log2ParallelMergeLevel	2	Defines the PPS-derived Log2ParMrgLevel variable.
SignHideFlag (-SBH)	true	If enabled specifies that for each 4x4 coefficient group for which the number of coefficients between the first nonzero coefficient and the last nonzero coefficient along the scanning line exceeds 4, the sign bit of the first nonzero coefficient will not be directly transmitted in the bitstream, but may be inferred from the parity of the sum of all nonzero coefficients in the current coefficient group.
StrongIntraSmoothing (-sis)	true	If enabled specifies that for 32x32 intra prediction block, the intra smoothing when applied is either the 1:2:1 smoothing filter or a stronger bi-linear interpolation filter. Key reference sample values are tested and if the criteria is satisfied, the stronger intra smoothing filter is applied. If disabled, the intra smoothing filter when applied is the 1:2:1 smoothing filter.
TMVPMode	1	Controls the temporal motion vector prediction mode. 0 Disabled for all slices. 1 Enabled for all slices. 2 Disabled only for the first picture of each GOPSize.
TransformSkip	false	Enables or disables transform-skipping mode decision.
TransformSkipFast	false	Enables or disables reduced testing of the transform-skipping mode decision for chroma TUs. When enabled, no RDO search is performed for chroma TUs, instead they are transform-skipped if the four corresponding luma TUs are also skipped. This option has no effect if TransformSkip is disabled.

Table 13: Rate control parameters

Option	Default	Description
RateControl	false	Rate control: enables rate control or not.
TargetBitrate	0	Rate control: target bitrate, in bps.
KeepHierarchicalBit	0	Rate control: 0: equal bit allocation among pictures; 1: fix ratio hierarchical bit allocation; 2: adaptive hierarchical ratio bit allocation. It is suggested to enable hierarchical bit allocation for hierarchical-B coding structure.
LCULevelRateControl	true	Rate control: true: LCU level RC; false: picture level RC.
RCLCUSeparateModel	true	Rate control: use LCU level separate R-lambda model or not. When LCULevelRateControl is equal to false, this parameter is meaningless.

Continued...

Table 13: Rate control parameters (Continued)

Option	Default	Description
InitialQP	0	Rate control: initial QP value for the first picture. 0 to auto determine the initial QP value.
RCForceIntraQP	false	Rate control: force intra QP to be equal to initial QP or not.
RCCpbSaturation	false	Rate control: enable target bits saturation to avoid CPB overflow and underflow or not.
RCCpbSize	0	Rate control: CPB size, in bps.
RCInitialCpbFullness	0.9	Rate control: ratio of initial CPB fullness per CPB size. (InitialCpbFullness/CpbSize) RCInitialCpbFullness should be smaller than or equal to 1.

Table 14: VUI parameters

Option	Default	Description
VuiParametersPresent (-vui)	false	Enable generation of vui_parameters().
AspectRatioInfoPresent	false	Signals whether aspect_ratio_idc is present.
AspectRatioIdc	0	aspect_ratio_idc
SarWidth	0	Specifies the horizontal size of the sample aspect ratio.
SarHeight	0	Specifies the vertical size of the sample aspect ratio.
OverscanInfoPresent	false	Signals whether overscan_info_present_flag is present.
OverscanAppropriate	false	Indicates whether cropped decoded pictures are suitable for display using overscan. 0 Indicates that the decoded pictures should not be displayed using overscan. 1 Indicates that the decoded pictures may be displayed using overscan.
VideoSignalTypePresent	false	Signals whether video_format, video_full_range_flag, and colour_description_present_flag are present.
VideoFormat	5	Indicates representation of pictures.
VideoFullRange	false	Indicates the black level and range of luma and chroma signals. 0 Indicates that the luma and chroma signals are to be scaled prior to display. 1 Indicates that the luma and chroma signals are not to be scaled prior to display.
ColourDescriptionPresent	false	Signals whether colour_primaries, transfer_characteristics and matrix_coefficients are present.
ColourPrimaries	2	Indicates chromaticity coordinates of the source primaries.
TransferCharacteristics	2	Indicates the opto-electronic transfer characteristics of the source.
MatrixCoefficients	2	Describes the matrix coefficients used in deriving luma and chroma from RGB primaries.
ChromaLocInfoPresent	false	Signals whether chroma_sample_loc_type_top_field and chroma_sample_loc_type_bottom_field are present.
ChromaSampleLocTypeTopField	0	Specifies the location of chroma samples for top field.
ChromaSampleLocTypeBottomField	0	Specifies the location of chroma samples for bottom field.

Continued...

Table 14: VUI parameters (Continued)

Option	Default	Description
NeutralChromaIndication	false	Indicates that the value of all decoded chroma samples is equal to $1 \ll (\text{BitDepthCr}-1)$.
DefaultDisplayWindowFlag	flag	Indicates the presence of the Default Window parameters. false Disabled true Enabled
DefDispWinLeftOffset DefDispWinRightOffset DefDispWinTopOffset DefDispWinBottomOffset	0	Specifies the horizontal and vertical offset to be applied to the input video from the conformance window in luma samples. Must be a multiple of the chroma resolution (e.g. a multiple of two for 4:2:0).
FrameFieldInfoPresentFlag	false	Specifies the value of the VUI syntax element 'frame_field_info_present_flag', which indicates that pic_struct and field coding related values are present in picture timing SEI messages.
PocProportionalToTimingFlag	false	Specifies the value of the VUI syntax element 'vui_poc_proportional_to_timing_flag', which indicates that the POC value is proportional to the output time with respect to the first picture in the CVS.
NumTicksPocDiffOneMinus	0	Specifies the value of the VUI syntax element 'vui_num_ticks_poc_diff_one_minus1', which specifies the number of clock ticks corresponding to a difference of picture order count values equal to 1, and is used only when PocProportionalToTimingFlag is true.
BitstreamRestriction	false	Signals whether bitstream restriction parameters are present.
TilesFixedStructure	false	Indicates that each active picture parameter set has the same values of the syntax elements related to tiles.
MotionVectorsOverPicBoundaries	false	Indicates that no samples outside the picture boundaries are used for inter prediction.
MaxBytesPerPicDenom	2	Indicates a number of bytes not exceeded by the sum of the sizes of the VCL NAL units associated with any coded picture.
MaxBitsPerMinCuDenom	1	Indicates an upper bound for the number of bits of coding_unit() data.
Log2MaxMvLengthHorizontal	15	Indicate the maximum absolute value of a decoded horizontal MV component in quarter-pel luma units.
Log2MaxMvLengthVertical	15	Indicate the maximum absolute value of a decoded vertical MV component in quarter-pel luma units.

Table 15: Range Extensions (Version 2) tool parameters

Option	Default	Description
CostMode	lossy	Specifies the cost mode to use. lossy $cost = distortion + \lambda \times bits$ sequence_level_lossless $cost = distortion/\lambda + bits$. lossless As with sequence_level_lossless, but QP is also set to 0 (this will be deprecated in the future) mixed_lossless_lossy As with sequence_level_lossless, but QP=4 is used for pre-estimates of transquant-bypass blocks

Continued...

Table 15: Range Extensions (Version 2) tool parameters (Continued)

Option	Default	Description
ExtendedPrecision	false	Specifies the use of <code>extended_precision_processing</code> flag. Note that unless the <code>HIGH_BIT_DEPTH_SUPPORT</code> macro in <code>TypeDef.h</code> is enabled, all internal bit depths must be 8 when the <code>ExtendedPrecision</code> setting is enabled. This setting is only valid for the 16-bit RExt profiles.
HighPrecisionPredictionWeighting	false	Specifies the value of <code>high_precision_prediction_weighting_flag</code> . This setting is only valid for the 16-bit or 4:4:4 RExt profiles.
CrossComponentPrediction	false	When true, specifies the use of the cross component prediction tool (4:4:4 processing only). Version 1 and some Version 2 (RExt) profiles require this to be false.
ReconBasedCrossCPredictionEstimate	false	If true, then when determining the alpha value for cross-component prediction, use the reconstructed residual rather than the pre-transform encoder-side residual
SaoLumaOffsetBitShift	0	Specifies the shift to apply to the SAO parameters. If negative, an estimate will be calculated based upon the initial QP. Version 1 and some Version 2 (RExt) profiles require this to be 0.
SaoChromaOffsetBitShift	0	
TransformSkipLog2MaxSize	2	Specifies the maximum TU size for which transform-skip can be used; the minimum value is 2. Version 1 and some Version 2 (RExt) profiles require this to be 2.
ImplicitResidualDPCM	false	When true, specifies the use of the implicitly signalled residual RDPCM tool (for intra). Version 1 and some Version 2 (RExt) profiles require this to be false.
ExplicitResidualDPCM	false	When true, specifies the use of the explicitly signalled residual RDPCM tool (for intra-block-copy and inter). Version 1 and some Version 2 (RExt) profiles require this to be false.
ResidualRotation	false	When true, specifies the use of the residual rotation tool. Version 1 and some Version 2 (RExt) profiles require this to be false.
SingleSignificanceMapContext	false	When true, specifies the use of a single significance map context for transform-skipped and transquant-bypassed TUs. Version 1 and some Version 2 (RExt) profiles require this to be false.
GolombRiceParameterAdaptation	false	When true, enable the adaptation of the Golomb-Rice parameter over the course of each slice. Version 1 and some Version 2 (RExt) profiles require this to be false.
AlignCABACBeforeBypass	false	When true, align the CABAC engine to a defined fraction of a bit prior to coding bypass data (including sign bits) when <code>coeff_abs_level_remaining</code> syntax elements are present in the group. This must always be true for the high-throughput-RExt profile, and false otherwise.
IntraReferenceSmoothing	true	When true, enable intra reference smoothing, otherwise disable it. Version 1 and some Version 2 (RExt) profiles require this to be true.

3.3 Encoder SEI parameters

The table below lists the SEI messages defined for Version 1 and Range-Extensions, and if available, the respective table that lists the controls within the HM Encoder to include the messages within the bit stream.

Table 16: List of Version 1 and RExt SEI messages

SEI Number	SEI Name	Table number of encoder controls, if available
0	Buffering period	Table 17
1	Picture timing	Table 18
2	Pan-scan rectangle	(Not handled)
3	Filler payload	(Not handled)
4	User data registered by Rec. ITU-T T.35	(Not handled)
5	User data unregistered	Decoded only
6	Recovery point	Table 19
9	Scene information	(Not handled)
15	Picture snapshot	(Not handled)
16	Progressive refinement segment start	(Not handled)
17	Progressive refinement segment end	(Not handled)
19	Film grain characteristics	(Not handled)
22	Post-filter hint	(Not handled)
23	Tone mapping information	Table 20
45	Frame packing arrangement	Table 21
47	Display orientation	Table 22
56	Green Metadata	Table 23
128	Structure of pictures information	Table 24
129	Active parameter sets	Table 25
130	Decoding unit information	Table 26
131	Temporal sub-layer zero index	Table 27
132	Decoded picture hash	Table 28
133	Scalable nesting	Table 29
134	Region refresh information	Table 30
135	No display	Table 31
136	Time code	Table 32
137	Mastering display colour volume	Table 33
138	Segmented rectangular frame packing arrangement	Table 34
139	Temporal motion-constrained tile sets	Table 35
140	Chroma resampling filter hint	Table 36
141	Knee function information	Table 37
142	Colour remapping information	Table 38

Continued...

Table 16: List of Version 1 and RExt SEI messages (Continued)

SEI Number	SEI Name	Table number of encoder controls, if available
143	Deinterlaced field identification	(Not handled)

Table 17: Buffering period SEI message encoder parameters

Option	Default	Description
SEIBufferingPeriod	0	Enables or disables the insertion of the Buffering period SEI messages. This option has no effect if VuiParametersPresent is disabled. SEIBufferingPeriod requires SEIActiveParameterSets to be enabled.

Table 18: Picture timing SEI message encoder parameters

Option	Default	Description
SEIPictureTiming	0	Enables or disables the insertion of the Picture timing SEI messages. This option has no effect if VuiParametersPresent is disabled.

Table 19: Recovery point SEI message encoder parameters

Option	Default	Description
SEIRecoveryPoint	0	Enables or disables the insertion of the Recovery point SEI messages.

Table 20: Tone mapping information SEI message encoder parameters

Option	Default	Description
SEIToneMappingInfo	0	Enables or disables the insertion of the Tone Mapping SEI message.
SEIToneMapId	0	Specifies Id of Tone Mapping SEI message for a given session.
SEIToneMapCancelFlag	false	Indicates that Tone Mapping SEI message cancels the persistence or follows.
SEIToneMapPersistenceFlag	true	Specifies the persistence of the Tone Mapping SEI message.
SEIToneMapCodedDataBitDepth	8	Specifies Coded Data BitDepth of Tone Mapping SEI messages.
SEIToneMapTargetBitDepth	8	Specifies Output BitDepth of Tome mapping function.
SEIToneMapModelId	0	Specifies Model utilized for mapping coded data into target_bit_depth range. 0 linear mapping with clipping 1 sigmoidal mapping 2 user-defined table mapping 3 piece-wise linear mapping 4 luminance dynamic range mapping
SEIToneMapMinValue	0	Specifies the minimum value in mode 0.
SEIToneMapMaxValue	1023	Specifies the maxmum value in mode 0.

Continued...

Table 20: Tone mapping information SEI message encoder parameters (Continued)

Option	Default	Description
SEIToneMapSigmoidMidpoint	512	Specifies the centre point in mode 1.
SEIToneMapSigmoidWidth	960	Specifies the distance between 5the target_bit_depth in mode 1.
SEIToneMapStartOfCodedInterval		Array of user-defined mapping table. Default table can be set to the following: 0 12 24 36 48 60 72 84 96 108 120 132 144 156 168 180 192 192 196 204 208 216 220 228 232 240 248 252 260 264 272 276 284 292 292 296 300 304 308 312 320 324 328 332 336 344 348 352 356 360 368 372 376 380 384 388 396 400 404 408 412 420 424 428 432 436 444 444 444 448 452 456 460 464 468 472 476 476 480 484 488 492 496 500 504 508 508 512 516 520 524 528 532 536 540 540 544 548 552 556 560 564 568 572 572 576 580 584 588 592 596 600 604 604 608 612 616 620 624 628 632 636 636 640 644 648 652 656 660 664 668 672 672 672 676 680 680 684 688 692 692 696 700 704 704 708 712 716 716 720 724 724 728 732 736 736 740 744 748 748 752 756 760 760 764 768 768 772 776 780 780 784 788 792 792 796 800 804 804 808 812 812 816 820 824 824 828 832 836 836 840 844 848 848 852 856 860 860 860 864 864 868 872 872 876 880 880 884 884 888 892 892 896 900 900 904 908 908 912 912 916 920 920 924 928 928 932 936 936 940 940 944 948 948 952 956 956 960 964 964 968 968 972 976 976 980 984 984 988 992 992 996 996 1000 1004 1004 1008 1012 1012 1016 1020 1024
SEIToneMapNumPivots	0	Specifies the number of pivot points in mode 3.
SEIToneMapCodedPivotValue		Array of coded pivot point in mode 3. A suggested table is: 64 128 256 512 768
SEIToneMapTargetPivotValue		Array of target pivot point in mode 3. A suggested table is: 48 73 111 168 215
SEIToneMap... CameraIsoSpeedIdc	0	Indicates the camera ISO speed for daylight illumination.
SEIToneMap... CameraIsoSpeedValue	400	Specifies the camera ISO speed for daylight illumination of Extended_ISO.
SEIToneMap... ExposureIndexIdc	0	Indicates the exposure index setting of the camera.
SEIToneMap... ExposureIndexValue	400	Specifies the exposure index setting of the cameran of Extended_ISO.
SEIToneMapExposure... CompensationValueSignFlag	0	Specifies the sign of ExposureCompensationValue.
SEIToneMapExposure... CompensationValueNumerator	0	Specifies the numerator of ExposureCompensationValue.
SEIToneMapExposure... CompensationValueDenomIdc	2	Specifies the denominator of ExposureCompensationValue.
SEIToneMapRef... ScreenLuminanceWhite	350	Specifies reference screen brightness setting in units of candela per square metre.
SEIToneMapExtended... RangeWhiteLevel	800	Indicates the luminance dynamic range.
SEIToneMapNominal... BlackLevelLumaCodeValue	16	Specifies luma sample value of the nominal black level assigned decoded pictures.

Continued...

Table 20: Tone mapping information SEI message encoder parameters (Continued)

Option	Default	Description
SEIToneMapNominal... WhiteLevelLumaCodeValue	235	Specifies luma sample value of the nominal white level assigned decoded pictures.
SEIToneMapExtended... WhiteLevelLumaCodeValue	300	Specifies luma sample value of the extended dynamic range assigned decoded pictures.

Table 21: Frame packing arrangement SEI message encoder parameters

Option	Default	Description
SEIFramePacking	0	Enables or disables the insertion of the Frame packing arrangement SEI messages.
SEIFramePackingType	0	Indicates the arrangement type in the Frame packing arrangement SEI message. This option has no effect if SEIFramePacking is disabled. 3 Side by Side 4 Top Bottom 5 Frame Alternate
SEIFramePackingInterpretation	0	Indicates the constituent frames relationship in the Frame packing arrangement SEI message. This option has no effect if SEIFramePacking is disabled. 0 Unspecified 1 Frame 0 is associated with the left view of a stereo pair 2 Frame 0 is associated with the right view of a stereo pair
SEIFramePackingQuincunx	0	Enables or disables the quincunx_sampling signalling in the Frame packing arrangement SEI messages. This option has no effect if SEIFramePacking is disabled.
SEIFramePackingId	0	Indicates the session number in the Frame packing arrangement SEI messages. This option has no effect if SEIFramePacking is disabled.

Table 22: Display orientation SEI message encoder parameters

Option	Default	Description
SEIDisplayOrientation	0	Enables or disables the insertion of the Display orientation SEI messages. 0 Disabled N: $0 < N < (2^{16} - 1)$ Enable display orientation SEI message with anticlockwise_rotation = N and display_orientation_repetition_period = 1

Table 23: Green Metadata SEI message encoder parameters

Option	Default	Description
SEIGreenMetadataType	0	Specifies the type of metadata that is present in the SEI message. 0 Reserved 1 Metadata enabling quality recovery after low-power encoding is present
SEIXSDMetricType	0	Indicates the type of the objective quality metric. 0 PSNR is used as objective quality metric

Table 24: Structure of pictures information SEI message encoder parameters

Option	Default	Description
SEISOPDescription	0	Enables or disables the insertion of the Structure of pictures information SEI messages.

Table 25: Active parameter sets SEI message encoder parameters

Option	Default	Description
SEIActiveParameterSets	0	Enables or disables the insertion of the Active parameter sets SEI messages.

Table 26: Decoding unit information SEI message encoder parameters

Option	Default	Description
SEIDecodingUnitInfo	0	Enables or disables the insertion of the Decoding unit information SEI messages. This option has no effect if VuiParametersPresent is disabled.

Table 27: Temporal sub-layer zero index SEI message encoder parameters

Option	Default	Description
SEITemporalLevel0Index	0	Enables or disables the insertion of the Temporal level zero index SEI messages.

Table 28: Decoded picture hash SEI message encoder parameters

Option	Default	Description
SEIDecodedPictureHash	0	Enables or disables the calculation and insertion of the Decoded picture hash SEI messages. 0 Disabled 1 Transmits MD5 in SEI message and writes the value to the encoder log 2 Transmits CRC in SEI message and writes the value to the encoder log 3 Transmits checksum in SEI message and writes the value to the encoder log

Table 29: Scalable nesting SEI message encoder parameters

Option	Default	Description
SEIScalableNesting	0	Enables or disables the use of the scalable nesting SEI messages.

Table 30: Region refresh information SEI message encoder parameters

Option	Default	Description
SEIGradualDecodingRefreshInfo	0	Enables or disables the insertion of the Gradual decoding refresh information SEI messages.

Table 31: No display SEI message encoder parameters

Option	Default	Description
SEINoDisplay	0	When non-zero, generate no-display SEI message for temporal layer N or higher.

Table 32: Time code SEI message encoder parameters

Option	Default	Description
SEITimeCodeEnabled	false	When true (non-zero), generate Time code SEI messages.
SEITimeCodeNumClockTs	0	Number of clock time sets, in the range of 0 to 3 (inclusive).
SEITimeCodeTimeStampFlag		Time stamp flag associated to each time set (comma or space separated list of entries).
SEITimeCodeFieldBasedFlag		Field based flag associated to each time set (comma or space separated list of entries).
SEITimeCodeCountingType		Counting type associated to each time set (comma or space separated list of entries).
SEITimeCodeFullTsFlag		Full time stamp flag associated to each time set (comma or space separated list of entries).

Continued...

Table 32: Time code SEI message encoder parameters (Continued)

Option	Default	Description
SEITimeCodeDiscontinuityFlag		Discontinuity flag associated to each time set (comma or space separated list of entries).
SEITimeCodeCntDroppedFlag		Counter dropped flag associated to each time set (comma or space separated list of entries).
SEITimeCodeNumFrames		Number of frames associated to each time set (comma or space separated list of entries).
SEITimeCodeSecondsFlag		Flag to signal seconds value presence in each time set (comma or space separated list of entries).
SEITimeCodeMinutesFlag		Flag to signal minutes value presence in each time set (comma or space separated list of entries).
SEITimeCodeHoursFlag		Flag to signal hours value presence in each time set (comma or space separated list of entries).
SEITimeCodeSecondsValue		Seconds value for each time set (comma or space separated list of entries).
SEITimeCodeMinutesValue		Minutes value for each time set (comma or space separated list of entries).
SEITimeCodeHoursValue		Hours value for each time set (comma or space separated list of entries).
SEITimeCodeOffsetLength		Time offset length associated to each time set (comma or space separated list of entries).
SEITimeCodeTimeOffset		Time offset associated to each time set (comma or space separated list of entries).

Table 33: Mastering display colour volume SEI message encoder parameters

Option	Default	Description
SEIMasteringDisplayColourVolume	false	When true (non-zero), generate Mastering display colour volume SEI message.
SEIMasteringDisplayMaxLuminance	10000	Specifies the mastering display maximum luminance value in units of 1/10000 candela per square metre.
SEIMasteringDisplayMinLuminance	0	Specifies the mastering display minimum luminance value in units of 1/10000 candela per square metre.
SEIMasteringDisplayPrimaries	0,50000, 0,0, 50000,0	Mastering display primaries for all three colour planes in CIE xy coordinates in increments of 1/50000 (results in the ranges 0 to 50000 inclusive).
SEIMasteringDisplayWhitePoint	16667, 16667	Mastering display white point CIE xy coordinates in normalized increments of 1/50000 (e.g. 0.333 = 16667).

Table 34: Segmented rectangular frame packing arrangement SEI message encoder parameters

Option	Default	Description
SEISegmentedRectFramePacking	0	Controls generation of segmented rectangular frame packing SEI messages.
SEISegmentedRectFramePackingCancel	false	If true, cancels the persistence of any previous SRFPA SEI message.

Continued...

Table 34: Segmented rectangular frame packing arrangement SEI message encoder parameters (Continued)

Option	Default	Description
SEISegmentedRectFramePackingType	0	Specifies the arrangement of the frames in the reconstructed picture.
SEISegmentedRectFramePackingPersistence	false	If false the SEI applies to the current frame only.

Table 35: Temporal motion-constrained tile sets SEI message encoder parameters

Option	Default	Description
SEITempMotionConstrainedTileSets	false	When true (non-zero), generates example temporal motion constrained tile sets SEI messages.

Table 36: Chroma resampling filter hint SEI message encoder parameters

Option	Default	Description
SEIChromaResamplingFilterHint	false	When true (non-zero), generates example chroma sampling filter hint SEI messages.
SEIChromaResamplingHorizontalFilterType	2	Defines the index of the chroma sampling horizontal filter: 0 Unspecified 1 Filters signalled within the SEI message 2 Filters as described by SMPTE RP 2050-1:2012
SEIChromaResamplingVerticalFilterType	2	Defines the index of the chroma sampling vertical filter: 0 Unspecified 1 Filters signalled within the SEI message 2 Filters as described in the 5/3 filter description of ITU-T Rec. T.800 ISO/IEC 15444-1

Table 37: Knee function SEI message encoder parameters

Option	Default	Description
SEIKneeFunctionInfo	false	Enables (true) or disables (false) the insertion of the Knee function SEI messages.
SEIKneeFunctionId	0	Specifies Id of Knee function SEI message for a given session.
SEIKneeFunctionCancelFlag	false	Indicates that Knee function SEI message cancels the persistence (true) or follows (false).
SEIKneeFunctionPersistenceFlag	true	Specifies the persistence of the Knee function SEI message.
SEIKneeFunctionInputDrange	1000	Specifies the peak luminance level for the input picture of Knee function SEI messages.
SEIKneeFunctionInputDispLuminance	100	Specifies the expected display brightness for the input picture of Knee function SEI messages.
SEIKneeFunctionOutputDrange	4000	Specifies the peak luminance level for the output picture of Knee function SEI messages.
SEIKneeFunctionOutputDispLuminance	800	Specifies the expected display brightness for the output picture of Knee function SEI messages.

Continued...

Table 37: Knee function SEI message encoder parameters (Continued)

Option	Default	Description
SEIKneeFunctionNumKneePointsMinus1	2	Specifies the number of knee points - 1.
SEIKneeFunctionInputKneePointValue		Array of input knee point. Default table can be set to the following: 600 800 900
SEIKneeFunctionOutputKneePointValue		Array of output knee point. Default table can be set to the following: 100 250 450

Table 38: Colour remapping SEI message encoder parameters

Option	Default	Description
SEIColourRemappingInfoFileRoot (-cri)		Specifies the prefix of input Colour Remapping Information file. Prefix is completed by “_x.txt” where x is the POC number. The contents of the file are a list of the SEI message’s syntax element names (in decoding order) immediately followed by a ‘:’ and then the associated value. An example file can be found in <code>cfg/misc/example_colour_remapping_sei_encoder_0.txt</code> .

3.4 Hardcoded encoder parameters

Table 39: CommonDef.h constants

Option	Default	Description
ADAPT_SR_SCALE	1	Defines a scaling factor used to derive the motion search range is adaptive (see ASR configuration parameter). Default value is 1.
MAX_GOP	64	maximum size of value of hierarchical GOP.
MAX_NUM_REF	4	maximum number of multiple reference frames
MAX_NUM_REF_LC	8	maximum number of combined reference frames
AMVP_MAX_NUM_CANDS	2	maximum number of final candidates
AMVP_MAX_NUM_CANDS_MEM	3	
MRG_MAX_NUM_CANDS	5	
DYN_REF_FREE	off	dynamic free of reference memories
MAX_TLAYER	8	maximum number of temporal layers
ADAPT_SR_SCALE	on	division factor for adaptive search range
EARLY_SKIP_THRES	1.5	early skip if $RD < EARLY_SKIP_THRES * avg[BestSkipRD]$
MAX_NUM_REF_PICS	16	
MAX_CHROMA_FORMAT_IDC	3	

TypeDef.h

Numerous constants that guard individual adoptions are defined within [source/Lib/TLibCommon/TypeDef.h](#).

4 Using the decoder

4.1 General

TAppDecoder -b str.bin -o dec.yuv [options]

Table 40: Decoder options

Option	Default	Description
(-help)		Prints usage information.
BitStreamFile (-b)		Defines the input bit stream file name.
ReconFile (-o)		Defines reconstructed YUV file name. If empty, no file is generated.
SkipFrames (-s)	0	Defines the number of pictures in decoding order to skip.
MaxTemporalLayer (-t)	-1	Defines the maximum temporal layer to be decoded. If -1, then all layers are decoded.
TarDecLayerIdSetFile (-l)		Specifies the targetDecLayerIdSet file name. The file would contain white-space separated LayerId values of the layers that are to be decoded. Omitting the parameter, or using a value of -1 in the file decodes all layers.
OutputBitDepth (-d)	0 (Native)	Specifies the luma bit-depth of the reconstructed YUV file (the value 0 indicates that the native bit-depth is used)
OutputBitDepthC	0 (Native)	Defines the chroma bit-depth of the reconstructed YUV file (the value 0 indicates that the native bit-depth is used)
SEIDecodedPictureHash	1	Enable or disable verification of any Picture hash SEI messages. When this parameter is set to 0, the feature is disabled and all messages are ignored. When set to 1 (default), the feature is enabled and the decoder has the following behaviour: <ul style="list-style-type: none">• If Picture hash SEI messages are included in the bit stream, the same type of hash is calculated for each decoded picture and written to the log together with an indication whether the calculated value matches the value in the SEI message. Decoding will continue even if there is a mismatch.• After decoding is complete, if any MD5sum comparison failed, a warning is printed and the decoder exits with the status EXIT_FAILURE• The per-picture MD5 log message has the following formats: [MD5:d41d8cd98f00b204e9800998ecf8427e,(OK)], [MD5:d41d8cd98f00b204e9800998ecf8427e,(unk)], [MD5:d41d8cd98f00b204e9800998ecf8427e,(***ERROR***)] [rxMD5:b9e1...] where, “(unk)” implies that no MD5 was signalled for this picture, “(OK)” implies that the decoder agrees with the signalled MD5, “(***ERROR***)” implies that the decoder disagrees with the signalled MD5. “[rxMD5:...]” is the signalled MD5 if different.
OutputDecodedSEIMessagesFilename		When a non-empty file name is specified, information regarding any decoded SEI messages will be output to the indicated file. If the file name is '-', then stdout is used instead.
SEIColourRemappingInfoFilename		Specifies that the colour remapping SEI message should be applied to the output video, with the output written to this file. If no value is specified, the SEI message is ignored and no mapping is applied.

Continued...

Table 40: Decoder options (Continued)

Option	Default	Description
RespectDefDispWindow (-w)	0	Video region to be output by the decoder. 0 Output content inside the conformance window. 1 Output content inside the default window.
OutputColourSpaceConvert		Specifies the colour space conversion to apply to 444 video. Permitted values are: UNCHANGED No colour space conversion is applied YCrCbToYCbCr Swap the second and third components GBRtoRGB Reorder the three components If no value is specified, no colour space conversion is applied. The list may eventually also include RGB to YCbCr or YCgCo conversions.
SEINoDisplay	false	When true, do not output frames for which there is an SEI NoDisplay message.
ClipOutputVideoToRec709Range	0	If 1 then clip output video to the Rec. 709 Range on saving when Output-BitDepth is less than InternalBitDepth.

4.2 Using the decoder analyser

If the decoder is compiled with the macro `RExt__DECODER_DEBUG_BIT_STATISTICS` defined as 1 (either externally, or by editing `TypeDef.h`), the decoder will gather fractional bit counts associated with the different syntax elements, producing a table of the number of bits per syntax element, and where appropriate, according to block size and colour component/channel. The Linux makefile will compile both the analyser and standard version when the ‘all’ or ‘everything’ target is used (where the latter will also build high-bit-depth executables).