

Title: HM Software Manual
Status: Software AHG working document
Purpose: Information
Author(s): Frank Bossen *frank@bossentech.com*
 David Flynn *dflynn@blackberry.com*
 Karsten Sühring *Karsten.Suehring@hhi.fraunhofer.de*
Source: AHG chairs

Abstract

This document is a user manual describing usage of reference software for the HEVC project. It applies to version 12.1 of the software.

Contents

1	General Information	1
2	Installation and compilation	2
3	Using the encoder	2
3.1	GOP structure table	2
3.2	Encoder parameters	6
3.3	Hardcoded encoder parameters	16
4	Using the decoder	18

List of Tables

1	Available project files	2
2	GOP structure example	4
3	File, I/O and source parameters	6
4	Profile and level parameters	7
5	Unit definition parameters	7
6	Coding structure parameters	7
7	Motion estimation parameters	8
8	Mode decision parameters	8
9	Quantization parameters	9
10	Entropy coding parameters	10
11	Slice coding parameters	10
12	Deblocking filter parameters	10
13	Coding tools parameters	11
14	Rate control parameters	12
15	VUI parameters	13
16	SEI messages	14
17	CommonDef.h constants	16
18	Decoder options	18

1 General Information

Reference software is being made available to provide a reference implementation of the draft 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.

Table 1: Available project files

Environment	Location of project file
MS Visual Studio 8	build/HM_vc8.sln
MS Visual Studio 9	build/HM_vc9.sln
Xcode	HM.xcodeproj
Linux	build/linux/makefile

3 Using the encoder

TAppEncoder [-h] [-c config.cfg] [--parameter=value]

Option	Description
-h	Prints parameter usage.
-c	Defines configuration file to use. Multiple configuration files may be used with repeated -c options.
--parameter=value	Assigns value to a given parameter as further described below. Some parameters are also supported by shorthand “-opt value”.

Sample configuration files are provided in the cfg/ folder.

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 what pictures that has to be kept for future references, they have to 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.

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 $-1 1$.

Figure 1: A GOP structure



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 $\text{deltaRIdx}-1$ 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 $1 - 3 = -2$.

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

The values of $\text{deltaRIdx}-1$, 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 $\text{reference_pictures}$ 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.

```

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 */

```

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 described above.

3.2 Encoder parameters

Table 3: File, I/O and source parameters

Option	Shorthand	Default	Description
InputFile	-i		Specifies the input video file. Video data must be in a raw 4:2:0 planar format (Y'CbCr). 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.
InternalBitDepth		0 (InputBitDepth)	Specifies the bit depth used for coding. If the input video is a different bit depth to InternalBitDepth, it is automatically converted by: $\left\lfloor \frac{Pel * 2^{InternalBitDepth}}{2^{InputBitDepth}} \right\rfloor$ Note: The effect of this option is as if the input video is externally converted to the InternalBitDepth and then coded with this value as InputBitDepth. The codec has no notion of two different bit depths.
OutputBitDepth		0 (InternalBitDepth)	Specifies the bit depth of the output locally reconstructed video file. Note: This option has no effect on the decoding process.
InputBitDepthC		(InputBitDepth)	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.
InternalBitDepthC		(InternalBitDepth)	
OutputBitDepthC		(InternalBitDepthC)	
ConformanceMode		0	Specifies the conformance mode (cropping/padding parameters) to be applied to the input video. The following modes are available: 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 ConfLeft, ConfRight, ConfTop and ConfBottom
HorizontalPadding	-pdx	0	Specifies the horizontal and vertical padding to be applied to the input video in luma samples. Must be a multiple of the chroma resolution (e.g. a multiple of two for 4:2:0).
VerticalPadding	-pdy		
ConfLeft		0	Specifies the horizontal and vertical cropping to be applied to the input video in luma samples. Must be a multiple of the chroma resolution (e.g. a multiple of two for 4:2:0).
ConfRight			
ConfTop			
ConfBottom			
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 (all)	Specifies the number of frames to be encoded.

Table 4: Profile and level parameters

Option	Shorthand	Default	Description
Profile		none	Specifies the profile to which the encoded bitstream complies. Valid values are: none, main, main10, main-still-picture. Compatibility flags are automatically determined according to the profile. If <code>-Profile=main</code> , then main10 will always be signalled as compatible. If <code>-Profile=main10</code> , then main will be signalled as compatible if the bit-depth is 8-bit. NB: There is currently no validation that the encoder configuration complies with the profile and level 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. NB: There is currently no validation that the encoder configuration complies with the profile and level constraints.
Tier		main	Specifies the level tier to which the encoded bitstream complies. Valid values are: main, high. NB: There is currently no validation that the encoder configuration complies with the profile and level constraints.
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	Shorthand	Default	Description
MaxCUWidth		64	Defines the maximum CU width.
MaxCUHeight		64	Defines the maximum CU height.
MaxCUSize		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	Shorthand	Default	Description
IntraPeriod	-ip	-1	Specifies the intra frame period. A value of -1 implies an infinite period.

Continued...

Table 6: Coding structure parameters (Continued)

Option	Shorthand	Default	Description
DecodingRefreshType	-dr	0	Specifies the type of decoding refresh to apply at the intra frame period picture. 0 Applies an I picture (not a clean random access point). 1 Applies a non-IDR clean random access point (open GOP). 2 Applies an IDR random access point (closed GOP).
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.
ListCombination	-lc	true	Enables or disables the use of the combined reference list for uni-prediction in B-slices. 0 Reference list 0 and reference list 1 are identical and reference list 0 is used as the combined reference list. 1 The combined reference list is derived from reference list 0 and reference list 1. NB: LComb can only be 0 in low delay coding (more precisely, when list 0 and list 1 are the same)

Table 7: Motion estimation parameters

Option	Shorthand	Default	Description
FastSearch		true	Enables or disables the use of a fast motion search. 0 Full search method 1 Fast search 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.
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. $\text{SearchRange}' = \text{Round} \left(\text{SearchRange} * \text{ADAPT_SR_SCALE} * \frac{\text{abs}(\text{POC}_{\text{cur}} - \text{POC}_{\text{ref}})}{\text{RateGOPSize}} \right)$

Table 8: Mode decision parameters

Option	Shorthand	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 . N may be in the range 0–7.

Continued...

Table 8: Mode decision parameters (Continued)

Option	Shorthand	Default	Description
FEN		false	Enables or disables the use of fast encoder mode. When enabled, the following occurs: <ul style="list-style-type: none"> In the SAD computation for blocks having size larger than 8, only the lines of even rows in the block are considered. The number of iterations used in the bi-directional motion vector refinement in the motion estimation process is reduced from 4 to 1.
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	Shorthand	Default	Description
QP	-q	30.0	Specifies the base value of the quantization parameter.
CbQpOffset	-cbqpofs	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 cb_qp_offset and cr_qp_offset, that are transmitted in the PPS. Valid values are in the range [-12, 12].
CrQpOffset	-crqpofs	0	
MaxCuDQPDepth		0	Defines maximum depth of a minimum CuDQP for sub-LCU-level delta QP. MaxCuDQPDepth shall be greater than or equal to SliceGranularity.
RDOQ		true	Enables or disables rate-distortion-optimized quantization.
RDOQTS		true	Enables or disables rate-distortion-optimized quantization for transform-skipped TUs.
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 [-DeltaQpRD, DeptaQpRD], 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 [-MaxDeltaQP, 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 .
AdaptiveQpSelection	-aqps	false	Specifies whether QP values for non-I frames will be calculated on the fly based on statistics of previously coded frames.
RecalculateQPAccordingToLambda		false	Recalculate QP values according to lambda values. Do not suggest to be enabled in all intra case.

Table 10: Entropy coding parameters

Option	Shorthand	Default	Description
SBACRD		true	Enables or disables the use of bit counts from arithmetic coder in rate-distortion decisions.

Table 11: Slice coding parameters

Option	Shorthand	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.
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.
NumTileColumnsMinus1 NumTileRowsMinus1		0	Specifies the tile based picture partitioning geometry as $\text{NumTileColumnsMinus1} + 1 \times \text{NumTileRowsMinus1} + 1$ columns and rows.
UniformSpacingIdc		0	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 ColumnWidthArray and RowHeightArray respectively 1 Tile columns and tile rows are uniformly spaced.
ColumnWidthArray RowHeightArray			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 12: Deblocking filter parameters

Option	Shorthand	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.
DeblockingFilterControlPresent		false	Enables or disables the presence of the deblocking filter control parameters in the picture parameter set and in the slice segment header. When disabled, the default deblocking filter parameters are used.

Continued...

Table 12: Deblocking filter parameters (Continued)

Option	Shorthand	Default	Description
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 $-6..6$.

Table 13: Coding tools parameters

Option	Shorthand	Default	Description
SAO		true	Enables or disables the sample adaptive offset (SAO) filter.
SAOLcuBoundary		false	Enables or disables SAO parameter estimation using non-deblocked pixels for LCU bottom and right boundary areas.
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.
TransquantBypassEnableFlag		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_enable_flag that is transmitted in the PPS. See CUTransquantBypassFlagValue for further details.
CUTransquantBypassFlagValue		0	Controls the per CU transformation, quantization and filtering mode decision. This option corresponds to the value of the per CU cu_transquant_bypass_flag. 0 Bypass is not performed on any CU 1 Bypass is performed on all CUs This option has no effect if TransquantBypassEnableFlag is disabled.
PCMEnabledFlag		false	Enables or disables the use of PCM.
PCMLog2MaxSize		5 (= $\log_2(32)$)	Specifies \log_2 of the maximum PCM block size. When PCM is enabled, the PCM mode is available for $2N \times 2N$ 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 $2N \times 2N$ intra PUs larger than or equal to the specified minimum PCM block size. When larger than PCMLog2MaxSize, PCM mode is not used.
PCMInputBitDepthFlag		1	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 13: Coding tools parameters (Continued)

Option	Shorthand	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.
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 for 4x4 TUs ¹ .
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 14: Rate control parameters

Option	Shorthand	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.
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.

¹Enables transform_skip_enabled and per 4x4 TU tests

Table 15: VUI parameters

Option	Shorthand	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.
NeutralChromaIndication		false	Indicates that the value of all decoded chroma samples is equal to $1 \ll (\text{BitDepthCr} - 1)$.
DefaultDisplayWindowFlag		0	Indicates the presence of the Default Window parameters. 0 Disabled 1 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).
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.

Continued...

Table 15: VUI parameters (Continued)

Option	Shorthand	Default	Description
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 16: SEI messages

Option	Shorthand	Default	Description
SEIDecodedPictureHash		0	Enables or disables the calculation and insertion of the Decoded SEI messages. 0 Disabled 1 Transmits MD5 in SEI message and writes the value to log 2 Transmits CRC in SEI message and writes the value to log 3 Transmits checksum in SEI message and writes the value to coder log
SEIpictureDigest		0	Deprecated alias for SEIDecodedPictureHash. Do not use any
SEIRecoveryPoint		0	Enables or disables the insertion of the Recovery point SEI me
SEIActiveParameterSets		0	Enables or disables the insertion of the Active parameter sets SEI
SEIBufferingPeriod		0	Enables or disables the insertion of the Buffering period SEI me option has no effect if VuiParametersPresent is disabled. SEI riod requires SEIActiveParameterSets to be enabled.
SEIPictureTiming		0	Enables or disables the insertion of the Picture timing SEI me option has no effect if VuiParametersPresent is disabled.
SEIDecodingUnitInfo		0	Enables or disables the insertion of the Decoding unit informat sages. This option has no effect if VuiParametersPresent is dis
SEIGradualDecodingRefreshInfo		0	Enables or disables the insertion of the Gradual decoding refr tion SEI messages.
SEITemporalLevel0Index		0	Enables or disables the insertion of the Temporal level zero inc sages.
SEIDisplayOrientation		0	Enables or disables the insertion of the Display orientation SEI 0 Disabled N: $0 < N < (2^{16} - 1)$ Enable display orientation SEI with anticlockwise_rotation = display_orientation_repetition_per
SEIFramePacking		0	Enables or disables the insertion of the Frame packing arrangem sages.

Table 16: SEI messages (Continued)

Option	Shorthand	Default	Description
SEIFramePackingType		0	Indicates the arrangement type in the Frame packing arrangement SEI message. This option has no effect if SEIFramePacking is disabled. 0 Checkerboard 1 Line Alternate 2 Column Alternate 3 Side by Side 4 Top Bottom 5 Frame Alternate 6 2D Image 7 Tile Format
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.
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		0	Indicates that Tone Mapping SEI message cancels the persistence of the previous SEI message.
SEIToneMapPersistenceFlag		1	Specifies the persistence of the Tone Mapping SEI message.
SEIToneMapCodedDataBitDepth		10	Specifies Coded Data BitDepth of Tone Mapping SEI message.
SEIToneMapTargetBitDepth		8	Specifies Output BitDepth of Tone mapping function.
SEIToneMapModelId		0	Specifies Model utilized for mapping coded data into target 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 maximum value in mode 0.
SEIToneMapSigmoidMidpoint		512	Specifies the centre point in mode 1.
SEIToneMapSigmoidWidth		960	Specifies the distance between 5th and 95th percentile 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
SEIToneMapNumPivots		5	Specifies the number of pivot points in mode 3.
SEIToneMapCodedPivotValue			Array of coded pivot point in mode 3. Default table can be set to the following: 64 128 256 512 768

Table 16: SEI messages (Continued)

Option	Shorthand	Default	Description
SEIToneMapTargetPivotValue			Array of target pivot point in mode 3. Default table can be set to the following: 48 73 111 168 215
SEIToneMapCameraIsoSpeedIdc		0	Indicates the camera ISO speed for daylight illumination.
SEIToneMapCameraIsoSpeedValue		420	Specifies the camera ISO speed for daylight illumination of Ex
SEIToneMapExposureCompensationValueSignFlag		0	Specifies the sign of ExposureCompensationValue.
SEIToneMapExposureCompensationValueNumerator		0	Specifies the numerator of ExposureCompensationValue.
SEIToneMapExposureCompensationValueDenomIdc		2	Specifies the denominator of ExposureCompensationValue.
SEIToneMapRefScreenLuminanceWhite		350	Specifies reference screen brightness setting in units of candela per metre.
SEIToneMapExtendedRangeWhiteLevel		800	Indicates the luminance dynamic range.
SEIToneMapNominalBlackLevelLumaCodeValue		16	Specifies luma sample value of the nominal black level assigned to pictures.
SEIToneMapNominalWhiteLevelLumaCodeValue		235	Specifies luma sample value of the nominal white level assigned to pictures.
SEIToneMapExtendedWhiteLevelLumaCodeValue		300	Specifies luma sample value of the extended dynamic range coded pictures.

3.3 Hardcoded encoder parameters

Table 17: 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
HB_LAMBDA_FOR_LDC	on	use of B-style lambda for non-key pictures in low-delay mode
GPB_SIMPLE	on	Fast estimation of generalized B in low-delay mode
GPB_SIMPLE_UNI	on	Fast estimation of generalized B in low-delay mode for uni-direction
FASTME_SMOOTHER_MV	on	Fast ME using smoother MV assumption
ADAPT_SR_SCALE	on	division factor for adaptive search range
CLIP_TO_709_RANGE	off	
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

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

Table 18: Decoder options

Option	Shorthand	Default	Description
	-h		Prints usage information.
	-o		Defines reconstructed YUV file name.
	-s	0	Defines the number of pictures in decoding order to skip.
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)
SEIPictureDigest		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 bitstream, 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.
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.