

Title: VTM Software Manual
Status: Software AHG working document
Purpose: Information
Editors: Frank Bossen frank@bossentech.com
David Flynn xlxiangli@tencent.com
Xiang Li karl.sharman@eu.sony.com
Karl Sharman karsten.suehring@hhi.fraunhofer.de
Karsten Sühning
Source: AHG chairs

Abstract

This document is a user manual describing usage of the VTM reference software for the VVC project. It applies to version 9.0 of the software.

Contents

1	General Information	3
2	Installation and compilation	3
2.1	Build instructions for plain CMake (suggested)	3
2.2	Build instructions for make	4
2.3	Tool Installation on Windows	5
3	Using the encoder	5
3.1	GOP structure table	5
3.2	Encoder parameters	9
3.3	Encoder SEI parameters	25
3.4	Hardcoded encoder parameters	38
4	Using the decoder	39
4.1	General	39
4.2	Using the decoder analyser	40
5	Block statistics extension	40
5.1	Usage	40
5.2	Block statistics file formats	41
5.3	Visualization	42
5.4	Adding statistics	42
6	Using the stream merge tool	44
6.1	Usage	44

List of Tables

1	Supported compilers	3
2	GOP structure example	7

3	File, I/O and source parameters.	9
4	GOP based temporal filter paramters	11
5	Profile and level parameters	11
6	Layer parameters	12
7	Unit definition parameters	13
8	Coding structure parameters	13
9	Motion estimation parameters	14
10	Mode decision parameters	14
11	Quantization parameters	15
12	Slice and tile coding parameters	18
13	Subpicture coding parameters	19
14	In-loop filtering parameters	19
15	Coding tools parameters	20
16	Rate control parameters	23
17	Encoder debug parameters	23
18	VUI parameters	24
19	Range Extensions (Version 2) tool parameters	24
20	List of Version 1 and RExt SEI messages	25
21	Buffering period SEI message encoder parameters	27
22	Picture timing SEI message encoder parameters	27
23	Recovery point SEI message encoder parameters	27
24	Film grain characteristics SEI message encoder parameters	27
25	Tone mapping information SEI message encoder parameters	27
26	Frame packing arrangement SEI message encoder parameters	29
27	Display orientation SEI message encoder parameters	29
28	Green Metadata SEI message encoder parameters	29
29	Structure of pictures information SEI message encoder parameters	30
30	Active parameter sets SEI message encoder parameters	30
31	Decoding unit information SEI message encoder parameters	30
32	Temporal sub-layer zero index SEI message encoder parameters	30
33	Decoded picture hash SEI message encoder parameters	30
34	Scalable nesting SEI message encoder parameters	30
35	Region refresh information SEI message encoder parameters	30
36	No display SEI message encoder parameters	31
37	Time code SEI message encoder parameters	31
38	Mastering display colour volume SEI message encoder parameters	31
39	Segmented rectangular frame packing arrangement SEI message encoder parameters	32
40	Temporal motion-constrained tile sets SEI message encoder parameters	32
41	Chroma resampling filter hint SEI message encoder parameters	32
42	Knee function SEI message encoder parameters	32
43	Colour remapping SEI message encoder parameters	33
44	Equiangular Projection SEI message encoder parameters	33
45	Generalized Cubemap Projection SEI message encoder parameters	33
46	Sphere Rotation SEI message encoder parameters	34
47	Region-wise packing SEI message encoder parameters	34
48	Omni Viewport SEI message encoder parameters	35
49	Sample Aspect Ratio Information SEI message encoder parameters	36
50	Frame-Field Information SEI message encoder parameters	36
51	Subpicture Level Information SEI message encoder parameters	36
52	Content light level info SEI message encoder parameters	36
53	Alternative transfer characteristics SEI message encoder parameters	37
54	Ambient viewing environment SEI message encoder parameters	37
55	Content colour volume SEI message encoder parameters	37

56	CommonDef.h constants	38
57	Decoder options	39
58	Decoder options	40

1 General Information

Reference software is being made available to provide a reference implementation of the HEVC standard being developed by the Joint Video Experts Team (JVET) 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 VVC 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 JVET email reflector on <https://lists.rwth-aachen.de/postorius/lists/jvet.lists.rwth-aachen.de/> (registration required).

Bug reporting

Bugs should be reported on the issue tracker set up at:

<https://jvet.hhi.fraunhofer.de/trac/vvc/>

2 Installation and compilation

The software may be retrieved from the GitLab server located at:

https://vcgit.hhi.fraunhofer.de/jvet/VVCSoftware_VTM

Table 1 lists the compiler environments and versions for which building the software is tested.

Note that the software makes use of C++11 language features, which may not be available in older compilers.

Table 1: Supported compilers

Compiler environment	Versions
MS Visual Studio	2017 and 2019
GCC	5.4, 7.3 and 8.3
Xcode/clang	latest

By default the software is built as 64-bit binaries to be used on a 64-bit OS. This allows the software to use more than 2GB of RAM.

The software uses CMake to create platform-specific build files.

2.1 Build instructions for plain CMake (suggested)

Note: A working CMake installation is required for building the software.

CMake generates configuration files for the compiler environment/development environment on each platform. The following is a list of examples for Windows (MS Visual Studio), macOS (Xcode) and

Linux (make).

Open a command prompt on your system and change into the root directory of this project.

Create a build directory in the root directory:

```
mkdir build
```

Use one of the following CMake commands, based on your platform. Feel free to change the commands to satisfy your needs.

Windows Visual Studio 2015 64 Bit:

```
cd build  
cmake .. -G "Visual Studio 14 2015 Win64"
```

Then open the generated solution file in MS Visual Studio.

macOS Xcode:

```
cd build  
cmake .. -G "Xcode"
```

Then open the generated work space in Xcode.

Linux

For generating Linux Release Makefile:

```
cd build  
cmake .. -DCMAKE_BUILD_TYPE=Release
```

For generating Linux Debug Makefile:

```
cd build  
cmake .. -DCMAKE_BUILD_TYPE=Debug
```

Then type

```
make -j
```

to build the software.

For more details, refer to the CMake documentation: <https://cmake.org/cmake/help/latest/>

2.2 Build instructions for make

Note: The build instructions in this section require the make tool and Python to be installed, which are part of usual Linux and macOS environments. See section 2.3 for installation instruction for Python and GnuWin32 on Windows.

Open a command prompt on your system and change into the root directory of this project.

To use the default system compiler simply call:

```
make all
```

For MSYS2 and MinGW: Open an MSYS MinGW 64-Bit terminal and change into the root directory of this project.

Call:

```
make all toolset=gcc
```

2.3 Tool Installation on Windows

Download CMake: <http://www.cmake.org/> and install it.

Python and GnuWin32 are not mandatory, but they simplify the build process for the user.

Python	https://www.python.org/downloads/release/python-371/
GnuWin32	https://sourceforge.net/projects/getgnuwin32/files/getgnuwin32/0.6.30/GetGnuWin32-0.6.3.exe/download

To use MinGW, install MSYS2: http://repo.msys2.org/distrib/msys2-x86_64-latest.exe

Installation instructions: <https://www.msys2.org/>

Install the needed toolchains:

```
pacman -S --needed base-devel mingw-w64-i686-toolchain mingw-w64-x86_64-toolchain git  
↪ subversion mingw-w64-i686-cmake mingw-w64-x86_64-cmake
```

3 Using the encoder

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

Option	Description
--help	Prints parameter usage.
-li	Applies to its next config file or command line parameter only to define i-th layer encoding option. If empty, the configuration file applies to all layers
-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”. 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: An in-loop deblocking filter parameter for luma component, 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 $-12..12$.

betaOffsetDiv2: An in-loop deblocking filter parameter for luma component, 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 $-12..12$.

CbTcOffsetDiv2: An in-loop deblocking filter parameter for Cb component, CbTcOffsetDiv2 is added to the base parameter LoopFilterCbTcOffset_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 $-12..12$.

CbBetaOffsetDiv2: An in-loop deblocking filter parameter for Cb component, CbBetaOffsetDiv2 is added to the base parameter LoopFilterCbBetaOffset_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 $-12..12$.

CrTcOffsetDiv2: An in-loop deblocking filter parameter for Cr component, CrTcOffsetDiv2 is added to the base parameter LoopFilterCrTcOffset_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 $-12..12$.

CrBetaOffsetDiv2: An in-loop deblocking filter parameter for Cr component, CrBetaOffsetDiv2 is added to the base parameter LoopFilterCrBetaOffset_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 $-12..12$.

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_L0: Number of reference pictures in lists L0 that are used during coding.

num_ref_pics_L0: Size of reference picture list L0. 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_L0: 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 as their intendend order in the L0. Note that any pictures not supplied in this list and in the list of L1 will be discarded and therefore not available as reference pictures later.


num_ref_pics_active_L1: Number of reference pictures in lists L1 that are used during coding.

num_ref_pics_L1: Size of reference picture list L1. 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_L1: 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 as their intendend order in the L1. Note that any pictures not supplied in this list and in the list of L0 will be discarded and therefore not available as reference pictures later.

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. Note that picture with POC 4 can be included in the L0 or L1. 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



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
CbTcOffsetDiv2	0	0	0	0
CbBetaOffsetDiv2	0	0	0	0
CrTcOffsetDiv2	0	0	0	0
CrBetaOffsetDiv2	0	0	0	0
temporal_id	0	1	2	2
num_ref_pics_active_L0	1	1	1	1
num_ref_pics_L0	1	1	1	1
reference_pictures_L0	4	2	1	1
num_ref_pics_active_L1	0	1	1	1
num_ref_pics_L1	0	1	2	1
reference_pictures_L1		-2	-1 -3	-1

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


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. If more than one layer is encoded (i.e. MaxLayers > 1), a reconstructed file is written for each layer and the layer index is added as suffix to ReconFile. If one or more dots exist in the file name, the layer id is added before the last dot, e.g. 'reconst.yuv' becomes 'reconst0.yuv' for layer id 0, 'reconst' becomes 'reconst0'.
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: $\left\lfloor \frac{PeI * 2^{InternalBitDepth}}{2^{MSBExtendedBitDepth}} \right\rfloor$
OutputBitDepth	0	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. 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	
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.
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.

Continued...

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

Option	Default	Description
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 verboseness 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: 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.
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	1	Add Access Unit Delimiter NAL units between all Access Units.
EnablePictureHeaderInSliceHeader	1	Enable Picture Header to be signalled in Slice Header when encoding with single slice per picture.
ScalingRatioHor	1.0	Scaling ratio in horizontal direction for reference picture resampling.

Continued...

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

Option	Default	Description
ScalingRatioVer	1.0	Scaling ratio in vertical direction for reference picture resampling.
FractionNumFrames	1.0	Encode a fraction of the specified in FramesToBeEncoded frames.
SwitchPocPeriod	0	POC period at which resolution is changed.
UpscaledOutput	0	Picture output options: output upscaled (2), decoded but in full resolution buffer (1) or decoded cropped (0, default) picture for reference picture resampling.

Table 4: GOP based temporal filter paramters

Option	Default	Description
TemporalFilter	false	Enables or disables GOP based temporal filter.
TemporalFilterFutureReference	true	Enables or disable referencing future frames in the GOP based temporal filter. Can be used to disable future referencing for low delay configurations.
TemporalFilterStrengthFrame*		Strength for every * frame in GOP based temporal filter, where * is an integer. E.g. -TemporalFilterStrengthFrame8 0.95 will enable GOP based temporal filter at every 8th frame with strength 0.95. Longer intervals overrides shorter when there are multiple matches.

Table 5: 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.
SubProfile	0	Indicates interoperability metadata registered as specified by X Recommendation ITU-T T.35.
EnableDecodingCapabilityInformation	false	Enables writing of a decoding capability information (DCI). If disabled, no DCI will be written.
MaxBitDepthConstraint	0	For -profile=main-RExt, specifies the value to use to derive the general_max_bit_depth constraint flags for RExt profiles; when 0, use InternalBitDepth.
MaxChromaFormatConstraint	0	For -profile=main-RExt, specifies the chroma-format to use for the general profile constraints for RExt profiles; when 0, use the value of ChromaFormatIDC.
IntraConstraintFlag	false	For -profile=main-RExt, specifies the value of general_intra_constraint_flag to use for RExt profiles.

Continued...

Table 5: Profile and level parameters (Continued)

Option	Default	Description
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.
LowerBitRateConstraintFlag	true	Specifies the value of <code>general_lower_bit_constraint_flag</code> to use for RExt profiles.
NonPackedSource	false	Specifies the value of <code>general_non_packed_constraint_flag</code>
NonProjectedConstraintFlag	false	Specifies the value of <code>general_non_projected_constraint_flag</code>
NoResChangeInClvsConstraintFlag	false	Specifies the value of <code>no_res_change_in_clvs_constraint_flag</code>
SingleLayerConstraintFlag	false	Specifies the value of <code>single_layer_constraint_flag</code>
AllLayersIndependentConstraintFlag	false	Specifies the value of <code>all_layers_independent_constraint_flag</code>
OneTilePerPicConstraintFlag	false	Specifies the value of <code>one_tile_per_pic_constraint_flag</code>
OneSlicePerPicConstraintFlag	false	Specifies the value of <code>one_slice_per_pic_constraint_flag</code>
OneSubpicPerPicConstraintFlag	false	Specifies the value of <code>one_subpic_per_pic_constraint_flag</code>
FrameOnly	false	Specifies the value of <code>general_frame_only_constraint_flag</code>

Table 6: Layer parameters

Option	Default	Description
MaxLayers	1	Specifies the value to use to derive the <code>vps_max_layers_minus1</code> for layered coding
MaxSubLayers	1	Specifies the maximum number of temporal sublayers to signal in the VPS
AllLayersSameNumSublayersFlag	true	Specifies the value of <code>vps_all_layers_same_num_sublayers_flag</code> in the VPS
AllowablePredDirection	""	Specifies a list of values of the allowable prediction directions for dependent layers. The number of entries is equal to the number of temporal layers. 0 Both inter-layer and intra-layer predictions are allowed for the specified temporal layer. 1 Only inter-layer prediction is allowed for the specified temporal layer. 2 Only intra-layer prediction is allowed for the specified temporal layer.
LayerId <i>i</i>	0	Specifies the <code>nuh_layer_id</code> of the <i>i</i> -th layer (with <i>i</i> an integer greater than 0)
NumRefLayers <i>i</i>	0	Specifies the number of direct reference layers of the <i>i</i> -th layer (with <i>i</i> an integer greater than 0)
RefLayerIdx <i>i</i>	""	Specifies a list of indexes of the reference layers of the <i>i</i> -th layer (with <i>i</i> an integer greater than 0)
EachLayerIsAnOlsFlag	true	Specifies the value of <code>each_layer_is_an_ols_flag</code> in the VPS
OlsModeIdc	0	Specifies the value of <code>ols_mode_idc</code> in the VPS
NumOutputLayerSets	1	Specifies the number of output layer sets (OLS) signalled in the VPS
OlsOutputLayer <i>i</i>	""	Specifies a list of indexes of the output layers of the <i>i</i> -th OLS (with <i>i</i> an integer greater than 0)
NumPTLsInVPS	1	Specifies the number of profile_tier_level (PTL) syntax structures signalled in the VPS
LevelPTL <i>i</i>	Level::NONE	Specifies the level to signal in the <i>i</i> -th PTL of the VPS (with <i>i</i> an integer greater than 0)
OlsPTLIdx <i>i</i>	0	Specifies the index of the PTL that applies to the <i>i</i> -th OLS (with <i>i</i> an integer greater than 0)
SamePicTimingInAllOLS	1	Indicates that all OLSs are using the same (not nested) picture timing SEI message, i.e. picture timing SEI will not be included in scalable nesting SEI messages (if scalable nesting SEI is enabled).

Table 7: Unit definition parameters

Option	Default	Description
CTUSize	128	Defines the CTU size (width and height).
MaxCUWidth	64	Defines the maximum CU width.
MaxCUHeight	64	Defines the maximum CU height.
MaxCUSize (-s)	64	Defines the maximum CU size.
Log2MinCuSize	2	Defines the minimum CU size in logarithm base 2.
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.
MaxMTTHierarchyDepth	3	Defines the initial maximum depth of the multi-type tree for inter slices.
MaxMTTHierarchyDepthI	3	Defines the initial maximum depth of the multi-type tree for intra slices.
MaxMTTHierarchyDepthISliceC	3	Defines the initial maximum depth of the multi-type tree in dual tree for chroma components.
MaxMTTHierarchyDepthISliceL	3	Defines the initial maximum depth of the multi-type tree in dual tree for luma component.
MinQTChromaISliceInChromaSamples	4	Defines the initial minimum size of the quad tree in dual tree for chroma components. Note: this size is defined in chroma sample unit in configuration, and it is converted into luma sample unit according to the horizontal chroma subsampling ratio when applied in the software. In chroma format 4:2:2 case, this value shall be set to the value of the height of minimum chroma QT node in chroma samples.
MinQTISlice	8	Defines the initial minimum size of the quad tree for intra slices.
MinQTLumaISlice	8	Defines the initial minimum size of the quad tree in dual tree for luma component.
MinQTNonISlice	8	Defines the initial minimum size of the quad tree for inter slices.
MaxBTLumaISlice	32	Defines the initial maximum size of the binary tree in dual tree for luma component.
MaxBTChromaISlice	64	Defines the initial maximum size of the binary tree in dual tree for chroma components.
MaxBTNonISlice	128	Defines the initial maximum size of the binary tree for inter slices.
MaxTTLumaISlice	32	Defines the initial maximum size of the tenary tree in dual tree for luma component.
MaxTTChromaISlice	32	Defines the initial maximum size of the tenary tree in dual tree for chroma components.
MaxTTNonISlice	64	Defines the initial maximum size of the tenary tree for inter slices.

Table 8: 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. 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.
DRAPPeriod	0	Specifies the DRAP period in frames. Dependent RAP indication SEI messages are disabled if DRAPPeriod is 0.
GOPSize (-g)	1	Specifies the size of the cyclic GOP structure.

Continued...

Table 8: Coding structure parameters (Continued)

Option	Default	Description
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.
ReWriteParamSets	0	Enable writing of parameter sets (SPS, PPS, etc.) before every (intra) random access point to enable true random access.

Table 9: 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.
MaxNumGeoCand	5	Specifies the maximum number of geometric partitioning mode candidates to use.
MaxNumIBCMergeCand	6	Specifies the maximum number of IBC merge candidates to use.
DisableIntraInInter	0	Flag to disable intra PUs in inter slices.
MMVD	1	Enables or disables the merge mode with motion vector difference (MMVD).
MmvdDisNum	6	Specifies the number of MMVD distance entries used from the distance table at encoder.
CIIP	1	Enables or disables the merge mode with combined inter merge and intra prediction (CIIP).

Table 10: 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).

Continued...

Table 10: Mode decision parameters (Continued)

Option	Default	Description
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.
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.
SBTFast64WidthTh	1920	Picture width threshold for testing size-64 SBT in RDO (now for HD and above sequences).
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)
FastLocalDualTreeMode	0	Controls intra coding speedup introduced with local dual tree mode. <ul style="list-style-type: none"> 0 Disabled 1 Stop testing intra modes in inter slices, if best cost is more that 1.5 times inter cost. 2 Test only one intra mode in inter slices

Table 11: 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.
DepQuant	true	Enables or disables the usage of dependent quantization.

Continued...

Table 11: Quantization parameters (Continued)

Option	Default	Description
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: = 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: 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.
UseIdentityTableForNon420Chroma	1	Specifies whether identity chroma QP mapping tables are used for 4:2:2 and 4:4:4 content. When set to 1, the identity chroma QP mapping table is used for all the three chroma components for 4:2:2 or 4:4:4 content. When set to 0, chroma QP mapping table may be specified by other parameters in the configuration.
SameCQPTablesForAllChroma	1	Specifies that the Cb, Cr and joint Cb-Cr components all use the same chroma mapping table. When set to 1, the values of QpInValCr, QpOutValCr, QpInValCbCr and QpOutValCbCr are ignored. When set to 0, all Cb, Cr and joint Cb-Cr components may have different chroma QP mapping tables specified in the configuration file. Note that SameCQPTablesForAllChroma is ignored when UseIdentityTableForNon420Chroma is set to 1 for 4:2:2 and 4:4:4 content.
QpInValCb QpOutValCb		Specifies the input and coordinates of the pivot points used to specify the chroma QP mapping tables for the Cb component. Default values are as follows: QpInValCb 25, 33, 43 QpOutValCb 25, 32, 37 The values specify the pivot points for the chroma QP mapping table, the unspecified QP values are interpolated from the remaining values. E.g., the default values above specify that the pivot points for the chroma QP mapping table for the Cb component are (25, 25), (33, 32), (43, 37). Note that that QpInValCr and QpOutValCr are ignored when UseIdentityTableForNon420Chroma is set to 1 for 4:2:2 and 4:4:4 content.
QpInValCr QpOutValCr		Specifies the input and coordinates of the pivot points used to specify the chroma QP mapping tables for the Cr component. Default values are as follows: QpInValCr 0 QpOutValCr 0 The default values specify a pivot point of (0,0) which corresponds to an identity chroma QP mapping table. Note that that QpInValCr and QpOutValCr are ignored when SameCQPTablesForAllChroma is set to 1 or when UseIdentityTableForNon420Chroma is set to 1 for 4:2:2 and 4:4:4 content.
QpInValCbCr QpOutValCbCr		Specifies the input and coordinates of the pivot points used to specify the chroma QP mapping tables for the joint Cb-Cr component. Default values are as follows: QpInValCr 0 QpOutValCbCr 0 The default values specify a pivot point of (0,0) which corresponds to a identity chroma QP mapping table. Note that that QpInValCbCr and QpOutValCbCr are ignored when SameCQPTablesForAllChroma is set to 1 or when UseIdentityTableForNon420Chroma is set to 1 for 4:2:2 and 4:4:4 content.
CbQpOffset (-cbqpofs)	0	Global offset to apply to the luma QP to derive the QP of Cb and Cr respectively.
CrQpOffset (-crqpofs)	0	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].
CbCrQpOffset (-cbcrqpofs)	-1	Global offset to apply to the luma QP to derive the QP for joint Cb-Cr residual coding mode. This option corresponds to the value of cb_cr_qp_offset transmitted in the PPS. Valid values are in the range [-12, 12].
CbCrQpOffsetDualTree	0	Tile group QP offset for joint Cb-Cr residual coding mode when separate luma and chroma trees are used. This option corresponds to the value of tile_group_cb_cr_qp_offset transmitted in the tile group header. Valid values are in the range [-12, 12].

Continued...

Continued...

Table 11: Quantization parameters (Continued)

Option	Default	Description
LumaLevelToDeltaQPMMode	0	Luma-level based Delta QP modulation. 0 not used 1 Based on CTU average 2 Based on Max luma in CTU
LumaLevelToDeltaQPMMaxValWeight	1.0	Weight of per block maximum luma value when LumaLevelToDeltaQPMMode=2.
LumaLevelToDeltaQPMMappingLuma		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.
LumaLevelToDeltaQPMMappingDQP		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: $\text{clip}(\text{round}(WCGPPSXXQpScale * baseCQp) + XXQpOffset)$ where $baseCQp = (WCGPPSChromaQpScale * qpc + WCGPPSChromaQpOffset)$. 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	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.
WCGPPSCrQpScale		
SliceChromaQPOffsetPeriodicity	0	Defines the periodicity for inter slices that use the slice-level chroma QP offsets, as defined by SliceCbQpOffsetIntraOrPeriodic and SliceCrQpOffsetIntraOrPeriodic. 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	0	Defines the slice-level QP offset to be used for intra slices, or once every 'SliceChromaQPOffsetPeriodicity' pictures.
SliceCrQpOffsetIntraOrPeriodic		
MaxCuDQPDepth (-dqd)	0	Defines maximum depth of a minimum CuDQP for sub-LCU-level delta QP. Max-CuDQPDepth shall be greater than or equal to SliceGranularity.
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.
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 .
PerceptQPA (-qpa)	false	Enables or disables the perceptually optimized QP adaptation (QPA) method described in JVET-H0047, JVET-K0206, and JVET-M0091. Use this together with 'SliceChromaQPOffsetPeriodicity=1' and, in case of HDR input, 'LumaLevelToDeltaQPMMode=1' for best subjective quality. Cannot be used together with 'SelectiveRDOQ' (see above) or 'AdaptiveQP' (see below).
AdaptiveQP (-aq)	false	Enables or disables the legacy QP adaptation method based upon a psycho-visual model.

Continued...

Table 11: Quantization parameters (Continued)

Option	Default	Description
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.
DisableScalingMatrixForLFNST	true	Specifies whether scaling matrices are to be applied to blocks coded with LFNST.
DisableScalingMatrixForAlternativeColourSpace	true	Specifies whether scaling matrices are disabled to blocks when the colour space is not equal to the designated colour space of scaling matrices.
ScalingMatrixDesignatedColourSpace	true	Indicates if the designated colour space of scaling matrices is equal to the original colour space.
MaxCUCromaQpAdjustmentDepth	-1	Specifies the maximum depth for CU chroma QP adjustment; if negative, CU chroma QP adjustment is disabled.

Table 12: Slice and tile coding parameters

Option	Default	Description
EnablePicPartitioning	0	Enable picture partitioning (0: single tile, single slice, 1: multiple tiles/slices can be used).
TileColumnWidthArray		Tile column widths in units of CTUs. Last column width in list will be repeated uniformly to cover any remaining picture width.
TileRowHeightArray		Tile row heights in units of CTUs. Last row height in list will be repeated uniformly to cover any remaining picture height.
RasterScanSlices	0	Use raster-scan or rectangular slices (0: rectangular, 1: raster-scan).
SingleSlicePerSubpic	false	Enables slice layout derivation from subpicture layout. Requires more than one subpicture to be enabled. If enabled, all other slice layout parameters will be ignored.
RectSlicePositions		Rectangular slice positions. List containing pairs of top-left CTU RS address followed by bottom-right CTU RS address.
RectSliceFixedWidth	0	Fixed rectangular slice width in units of tiles (0: disable this feature and use RectSlicePositions instead).
RectSliceFixedHeight	0	Fixed rectangular slice height in units of tiles (0: disable this feature and use RectSlicePositions instead).
RasterSliceSizes		Raster-scan slice sizes in units of tiles. Last size in list will be repeated uniformly to cover any remaining tiles in the picture.
DisableLoopFilterAcrossTiles	0	Loop filtering applied across tile boundaries or not (0: filter across tile boundaries 1: do not filter across tile boundaries).
DisableLoopFilterAcrossSlices	0	Loop filtering applied across slice boundaries or not (0: filter across slice boundaries 1: do not filter across slice boundaries).
IDRRefParamList	false	Enables the signalling of reference picture list syntax elements in slice headers of IDR pictures
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.

Continued...

Table 12: Slice and tile coding parameters (Continued)

Option	Default	Description
WaveFrontEntryPointsPresent	false	Allow signalling of entry points for WPP in slice header. Note that when a slice contains more than one tile, entry point offsets for tile are always present in the slice header.
MixedLossyLossless	0	Enable or disable mixed lossy/lossless coding. 0 means disable; 1 means enable. Mixed lossy/lossless can only be enable if CostMode is set to lossless.
SliceLosslessArray		Slice index array of lossless slices. Example: 1 5 6 means slices with index of 1, 5, and 6 are lossless coded. The rest of the slices are lossy coded. If MixedLossyLossless is disabled, the values are ignored.

Table 13: Subpicture coding parameters

Option	Default	Description
SubPicInfoPresentFlag	false	Enables coding of subpictures.
NumSubPics	0	Number of subpictures. Must be greater than zero, if SubPicInfoPresentFlag is enabled.
SubPicCtuTopLeftX		Array of subpicture top left horizontal (x) coordinates. The number of entries must be equal to NumSubPics.
SubPicCtuTopLeftY		Array of subpicture top left vertical (y) coordinates. The number of entries must be equal to NumSubPics.
SubPicWidth		Array of subpicture widths. The number of entries must be equal to NumSubPics.
SubPicHeight		Array of subpicture heights. The number of entries must be equal to NumSubPics.
SubPicTreatedAsPicFlag		Setting of subpic_treated_as_pic_flag for each subpicture. If enabled subpicture boundaries will be treated as picture boundaries. The number of entries must be equal to NumSubPics.
LoopFilterAcrossSubpicEnabledFlag		Enables loop filtering across subpicture boundaries for each subpicture. The number of entries must be equal to NumSubPics.
SubPicIdMappingExplicitlySignalledFlag	false	Enables explicit signalling of a subpicture ID map. If disabled, a default map will be derived.
SubPicIdMappingInSpsFlag	false	Specifies whether to signal the subpicture ID map in SPS or PPS. If SubPicIdMappingInSpsFlag is enabled subpicture IDs are signalled in SPS, otherwise in PPS.
SubPicIdLen	0	Length of the subpicture IDs in bits. (1«SubPicIdLen) must be bigger than the number of subpictures and the highest subpicture ID specified in SubPicId. If the value “0” is used, the encoder tries to determine the number of required bits from the number of subpictures or the highest subpicture ID. This mode should not be used, if merging of bistreams is intended.
SubPicId		Array of subpicture IDs. The number of entries must be equal to NumSubPics.

Table 14: In-loop filtering parameters

Option	Default	Description
LoopFilterDisable	false	Enables or disables the in-loop deblocking filter.

Continued...

Table 14: In-loop filtering parameters (Continued)

Option	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, betaOffsetDiv2 for luma; CbTcOffsetDiv2, CbBetaOffsetDiv2 for Cb and CrTcOffsetDiv2, CrBetaOffsetDiv2 for Cr 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 for luma component. The final value of tc_offset_div2 shall be an integer number in the range $-12..12$.
LoopFilterBetaOffset_div2	0	Specifies the base value for the in-loop deblocking filter parameter beta_offset_div2 for luma component. The final value of beta_offset_div2 shall be an integer number in the range $-12..12$.
LoopFilterCbTcOffset_div2	0	Specifies the base value for the in-loop deblocking filter parameter tc_offset_div2 for Cb component. The final value of tc_offset_div2 shall be an integer number in the range $-12..12$.
LoopFilterCbBetaOffset_div2	0	Specifies the base value for the in-loop deblocking filter parameter beta_offset_div2 for Cb component. The final value of beta_offset_div2 shall be an integer number in the range $-12..12$.
LoopFilterCrTcOffset_div2	0	Specifies the base value for the in-loop deblocking filter parameter tc_offset_div2 for Cr component. The final value of tc_offset_div2 shall be an integer number in the range $-12..12$.
LoopFilterCrBetaOffset_div2	0	Specifies the base value for the in-loop deblocking filter parameter beta_offset_div2 for Cr component. The final value of beta_offset_div2 shall be an integer number in the range $-12..12$.
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.
VirtualBoundariesPresentInSPSFlag	false	In-loop filtering operations across the virtual boundaries information present in the SPS when VirtualBoundariesPresentFlagInSPS = 1, otherwise present in the Picture Header when VirtualBoundariesPresentFlagInSPS = 0.
NumVerVirtualBoundaries	0	Specifies the number of vertical virtual boundaries. The value of NumVerVirtualBoundaries shall be in the range of 0 to 3, inclusive.
NumHorVirtualBoundaries	0	Specifies the number of horizontal virtual boundaries. The value of NumHorVirtualBoundaries shall be in the range of 0 to 3, inclusive.
VirtualBoundariesPosX		Specifies the locations of the vertical virtual boundaries in units of luma samples
VirtualBoundariesPosY		Specifies the locations of the horizontal virtual boundaries in units of luma samples
EncDbOpt	false	Enables or disables encoder-side deblocking optimization. When it is enabled, deblocking filter is applied during mode decision.

Table 15: Coding tools parameters

Option	Default	Description
MRL	false	Enables or disables the use of multiple reference line intra prediction (MRL).
DualTree	false	Enables or disables the use of separate QTBT trees for intra slice luma and chroma channel types.
MIP	true	Enables or disables the use of matrix-based intra prediction (MIP).
AMP	true	Enables or disables the use of asymmetric motion partitions.

Continued...

Table 15: Coding tools parameters (Continued)

Option	Default	Description
ISP	false	Enables or disables the Intra Sub-Partitions coding mode.
ISPFast	false	Enables or disables fast encoder methods for ISP.
JointCbCr	false	Enables or disables the joint coding of chroma residuals.
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).
SAOGreedyEnc	false	Enables or disables the SAO greedy merge encoding algorithm.
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
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 $2N \times 2N$ 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	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.
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.
WeightedPredMethod (-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).
SignHideFlag (-SBH)	true	If enabled specifies that for each 4×4 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.
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.
SliceLevelRpl	true	Code reference picture lists in slice headers rather than picture header.
SliceLevelDbk	true	Code deblocking filter parameters in slice headers rather than picture header.
SliceLevelSao	true	Code SAO parameters in slice headers rather than picture header.

Continued...

Table 15: Coding tools parameters (Continued)

Option	Default	Description
SliceLevelWeightedPrediction	true	Code Weighted Prediction parameters in slice headers rather than picture header.
SliceLevelDeltaQp	true	Code delta Qp in slice headers rather than picture header.
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.
ChromaTS	false	Enables or disables reduced testing of the transform-skipping mode decision for chroma TUs. When disabled, no RDO search is performed for chroma TUs. This option has no effect if TransformSkip is disabled.
ALF	true	Enables or disables adaptive loop filter.
UseNonLinearAlfLuma	true	Enables optimization of non-linear filters for ALF on Luma channel.
UseNonLinearAlfChroma	true	Enables optimization of non-linear filters for ALF on Chroma channels.
MaxNumAlfAlternativesChroma	8	Specified the maximum number of alternative chroma filters that can be switched at CTB level. Set to 1 to disable alternative chroma filters. Value shall be in the range 1..8.
CCALF	true	Enables cross-component ALF.
CCALFQpTh	37	QP threshold above which the encoder reduces cross-component ALF usage.
SMVD	false	Enables or disables symmetric MVD mode.
Geo	false	Enables or disables geometric partitioning mode.
PLT	false	Enables or disables palette mode coding.
BDPCM	false	Enables or disables the use of intra block differential pulse code modulation mode.
LFNST	false	Enables or disables the use of low frequency non-separable transform (LFNST).
FastLFNST	false	Enables or disables the fast encoding of low frequency non-separable transform (LFNST).
BCW	false	Enables or disables the use of Bi-prediction with CU-level Weights (BCW).
BcwFast	false	Enables or disables the fast encoding of Bi-prediction with CU-level Weights (BCW).
BDOF	false	Enables or disables the use of bi-directional optical flow (BDOF).
LMCSEnable	true	Enables or disables the use of LMCS (luma mapping with chroma scaling).
LMCSSignalType	0	LMCS signal type: 0:SDR, 1:HDR-PQ, 2:HDR-HLG.
LMCSUpdateCtrl	0	LMCS model update control: 0:RA, 1:AI, 2:LDB/LDP. 0 Random access: derive a new LMCS model at each IRAP. 1 All intra: derive a new LMCS model at each intra slice. 2 Low delay: derive a new LMCS model every second.
LMCSAdpOption	0	Adaptive LMCS mapping derivation options: Options 1 to 4 are for experimental testing purposes and need to set parameter LMCSInitialCW. 0 Automatic adaptive algorithm (default). 1 Derives LMCS mapping with input LMCSInitialCW and enables LMCS for all slices. Uses a static LMCS mapping for low QP ($QP \leq 22$). 2 Derives LMCS mapping with input LMCSInitialCW and enables LMCS only for slices in lowest temporal layer. 3 In addition to 1, disables LMCS for intra slices. 4 Derives LMCS mapping with input LMCSInitialCW and enables LMCS only for inter slices.
LMCSInitialCW	0	LMCS initial total codeword (valid values [0 – 1023]) to be used in LMCS mapping derivation when LMCSAdpOption is not equal to 0.
LMCSOffset	0	Specifies the LMCS chroma residual scaling offset. This parameter corresponds to the value of <code>lmcsDeltaCrs</code> , derived from <code>lmcs_delta_sign_crs_flag</code> and <code>lmcs_delta_abs_crs</code> , that are transmitted in the APS. Valid values are in the range [-7;7].

Continued...

Table 15: Coding tools parameters (Continued)

Option	Default	Description
ColorTransform	false	Enables or disables the use of adaptive color transform (ACT).
HorCollocatedChroma	true	Specifies location of a chroma sample relatively to the luma sample in horizontal direction in the reference picture resampling. 0 horizontally shifted by 0.5 units of luma samples. 1 collocated (default).
VerCollocatedChroma	false	Specifies location of a chroma sample relatively to the luma sample in vertical direction in the cross-component linear model intra prediction and the reference picture resampling. 0 vertically shifted by 0.5 units of luma samples (default). 1 collocated.
TSRCdisableLL	1	Enables or disables the use of Transform Skip Residual Coding for lossless compression.

Table 16: 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.
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 17: Encoder debug parameters

Option	Default	Description
DebugBitstream/DecodeBitstream1		Specifies the first bit stream to be read until a pre-defined switch point is encountered.
DecodeBitstream2		Specifies the second bit stream, to be read after the first random access point after a QP switch point (specified using SwitchPOC and SwitchQP).
DebugPOC	-1	Specifies a POC, at which a bit stream specified using DebugBitstream or DecodeBitstream1 is no longer read, but rather normal encoding is started.
DebugCTU	-1	When the POC is encountered at which normal encoding is to be resumed, if set, this option specifies that CTUs up to the specified CTU(in raster scan addressing order) are to be read from the specified bit stream, after which normal encoding is started the specified CTU.
SwitchPOC	-1	Specifies a POC, at which the specified bit stream is no longer read, but rather normal encoding is started.

Continued...

Table 17: Encoder debug parameters (Continued)

Option	Default	Description
SwitchDQP	0	Specifies a QP offset to be applied when normal encoding is started as specified by SwitchPOC.
FastForwardToPOC	0	When encoding a bit streams, all frames that are not references including transitive references to the specified POC are skipped.
StopAfterFFtoPOC	false	If enabled, causes the encoder to not encode any frame after the frame specified by FastForwardToPOC option, in encoding order.

Table 18: 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.
ColourDescriptionPresent	false	Signals whether colour_primaries, transfer_characteristics, matrix_coefficients and video_full_range_flag 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.
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.
ProgressiveSource	false	Specifies the value of general_progressive_source_flag
InterlacedSource	false	Specifies the value of general_interlaced_source_flag
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.

Table 19: Range Extensions (Version 2) tool parameters

Option	Default	Description
CostMode	lossy	Specifies the cost mode to use. lossy $cost = distortion + \lambda \times bits$ lossless $cost = bits$, QP'=0 is used for all transform blocks and the only allowed encoder result is either an empty transform block or an transform skipped block.

Continued...

Table 19: Range Extensions (Version 2) tool parameters (Continued)

Option	Default	Description
ExtendedPrecision	false	Specifies the use of extended_precision_processing flag. Note that unless the HIGH_BIT_DEPTH_SUPPORT macro in TypeDef.h is enabled, all internal bit depths must be 8 when the ExtendedPrecision setting is enabled. This setting is only valid for the 16-bit RExt profiles.
HighPrecisionPredictionWeighting	false	Specifies the value of high_precision_prediction_weighting_flag. This setting is only valid for the 16-bit or 4:4:4 RExt profiles.
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
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 coeff_abs_level_remaining 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 20: List of Version 1 and RExt SEI messages

SEI Number	SEI Name	Table number of encoder controls, if available
0	Buffering period	Table 21
1	Picture timing	Table 22
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 23
9	Scene information	(Not handled)
15	Picture snapshot	(Not handled)
16	Progressive refinement segment start	(Not handled)

Continued...

Table 20: List of Version 1 and RExt SEI messages (Continued)

SEI Number	SEI Name	Table number of encoder controls, if available
17	Progressive refinement segment end	(Not handled)
19	Film grain characteristics	Table 24
22	Post-filter hint	(Not handled)
23	Tone mapping information	Table 25
45	Frame packing arrangement	Table 26
47	Display orientation	Table 27
56	Green Metadata	Table 28
128	Structure of pictures information	Table 29
129	Active parameter sets	Table 30
130	Decoding unit information	Table 31
131	Temporal sub-layer zero index	Table 32
132	Decoded picture hash	Table 33
133	Scalable nesting	Table 34
134	Region refresh information	Table 35
135	No display	Table 36
136	Time code	Table 37
137	Mastering display colour volume	Table 38
138	Segmented rectangular frame packing arrangement	Table 39
139	Temporal motion-constrained tile sets	Table 40
140	Chroma resampling filter hint	Table 41
141	Knee function information	Table 42
142	Colour remapping information	Table 43
143	Deinterlaced field identification	(Not handled)
144	Content light level info	Table 52
147	Alternative transfer characteristics	Table 53
148	Ambient viewing environment	Table 54
149	Content colour volume	Table 55
150	Equirectangular projection	Table 44
153	Generalized cubemap projection	Table 45
154	Sphere rotation	Table 46
155	Region-wise packing	Table 47
156	Omni viewport	Table 48
168	Frame-field information	Table 50
203	Subpicture Level Information	Table 51
204	Sample Aspect Ratio Information	Table 49

Table 21: 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 22: 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 23: Recovery point SEI message encoder parameters

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

Table 24: Film grain characteristics SEI message encoder parameters

Option	Default	Description
SEIFGCEnabled	0	Enables or disables the insertion of the film grain characteristics SEI message.
SEIFGCCancelFlag	0	Specifies the persistence of any previous film grain characteristics SEI message in output order.
SEIFGCPersistenceFlag	1	Specifies the persistence of the film grain characteristics SEI message for the current layer.
SEIFGCModelID	0	Specifies the film grain simulation model. 0 frequency filtering 1 auto-regression
SEIFGCSepColourDescPresentFlag	0	Specifies the presence of a distinct colour space description for the film grain characteristics specified in the SEI message.
SEIFGCBlendingModelID	0	Specifies the blending mode used to blend the simulated film grain with the decoded images. 0 additive 1 multiplicative
SEIFGCLog2ScaleFactor	0	Specifies a scale factor used in the film grain characterization equations.
SEIFGCCompModelPresentComp0	0	Specifies the presence of film grain modelling on colour component 0.
SEIFGCCompModelPresentComp1	0	Specifies the presence of film grain modelling on colour component 1.
SEIFGCCompModelPresentComp2	0	Specifies the presence of film grain modelling on colour component 2.

Table 25: 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.

Continued...

Table 25: Tone mapping information SEI message encoder parameters (Continued)

Option	Default	Description
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 maximum value in mode 0.
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 96 108 120 132 144 156 168 180 192 192 196 204 208 216 220 228 232 240 248 252 260 264 272 276 284 292 296 300 304 308 312 320 324 328 332 336 344 348 352 356 360 368 372 376 380 384 388 396 400 404 408 412 420 424 428 432 436 444 444 444 448 452 456 460 464 468 472 476 476 480 484 488 492 496 500 504 508 508 512 516 520 524 528 532 536 540 540 544 548 552 556 560 564 568 572 572 576 580 584 588 592 596 600 604 604 608 612 616 620 624 628 632 636 636 640 644 648 652 656 660 664 668 672 672 672 676 680 680 684 688 692 692 696 700 704 704 708 712 716 716 720 724 724 728 732 736 736 740 744 748 748 752 756 760 760 764 768 768 772 776 780 780 784 788 792 792 796 800 804 804 808 812 812 816 820 824 824 828 832 836 836 840 844 848 848 852 856 860 860 860 864 864 868 872 872 876 880 880 884 884 888 892 892 896 900 900 904 908 908 912 912 916 920 920 924 928 928 932 936 936 940 940 944 948 948 952 956 956 960 964 964 968 968 972 976 976 980 984 984 988 992 992 996 996 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.

Continued...

Table 25: Tone mapping information SEI message encoder parameters (Continued)

Option	Default	Description
SEIToneMapNominal... BlackLevelLumaCodeValue	16	Specifies luma sample value of the nominal black level assigned decoded pictures.
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 26: 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	3	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	1	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 27: 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 28: 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 29: 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 30: 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 31: 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 32: 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 33: 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 34: Scalable nesting SEI message encoder parameters

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

Table 35: 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 36: 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 37: 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).
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 38: 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 39: 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.
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 40: 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 41: 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 42: 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.
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 43: 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 cfg/misc/example_colour_remapping_sei_encoder_0.txt.

Table 44: Equirectangular Projection SEI message encoder parameters

Option	Default	Description
SEIerpEnabled	false	Enables (true) or disables (false) the insertion of equirectangular projection SEI message.
SEIerpCancelFlag	true	Indicates that equirectangular projection SEI message cancels the persistence (true) or follows (false).
SEIerpPersistenceFlag	false	Specifies the persistence of the equirectangular projection SEI message.
SEIerpGuardBandFlag	false	Indicates the existence of guard band areas in the constituent picture.
SEIerpGuardBandType	0	Indicates the type of the guard bands.
SEIerpLeftGuardBandWidth	0	Indicates the width of the guard band on the left side of the constituent picture.
SEIerpRightGuardBandWidth	0	Indicates the width of the guard band on the right side of the constituent picture.

Table 45: Generalized Cubemap Projection SEI message encoder parameters

Option	Default	Description
SEIGcmpEnabled	false	Enables (true) or disables (false) the insertion of generalized cubemap projection SEI message.
SEIGcmpCancelFlag	true	Indicates that generalized cubemap projection SEI message cancels the persistence (true) or follows (false).
SEIGcmpPersistenceFlag	false	Specifies the persistence of the generalized cubemap projection SEI message.
SEIGcmpPackingType	0	Specifies the packing type. 0 6 rows and 1 columns 1 3 rows and 2 columns 2 2 rows and 3 columns 3 1 rows and 6 columns 4 1 rows and 5 columns (hemisphere cubemap) 5 5 rows and 1 columns (hemisphere cubemap)
SEIGcmpMappingFunctionType	0	Specifies the mapping function used to adjust the sample locations. 0 Disabled (conventional cubemap projection) 1 Equi-angular mapping function 2 Defined by SEIGcmpFunctionCoeffU, SEIGcmpFunctionUAffectedByVFlag, SEIGcmpFunctionCoeffV, and SEIGcmpFunctionVAffectedByUFlag
SEIGcmpFaceIndex		An array that specifies the face index for the faces packed in the cubemap projected picture. 0 Front face 1 Back face 2 Top face 3 Bottom face 4 Right face 5 Left face
SEIGcmpFaceRotation		An array that specifies the rotation to be applied to the faces. 0 No rotation 1 90 degree anticlockwise 2 180 degree anticlockwise 3 270 degree anticlockwise

Continued...

Table 45: Generalized Cubemap Projection SEI message encoder parameters (Continued)

Option	Default	Description
SEIGcmpFunctionCoeffU		An array that specifies the coefficients used in the cubemap mapping function of the u-axis for the faces when SEIGcmpMappingFunctionType is set to 2.
SEIGcmpFunctionUAffectedByVFlag		An array that specifies whether the cubemap mapping function of the u-axis refers to the v position of the sample location for the faces when SEIGcmpMappingFunctionType is set to 2.
SEIGcmpFunctionCoeffV		An array that specifies the coefficients used in the cubemap mapping function of the v-axis for the faces when SEIGcmpMappingFunctionType is set to 2.
SEIGcmpFunctionVAffectedByUFlag		An array that specifies whether the cubemap mapping function of the v-axis refers to the u position of the sample location for the faces when SEIGcmpMappingFunctionType is set to 2.
SEIGcmpGuardBandFlag	false	Indicates the existence of guard band areas in the picture.
SEIGcmpGuardBandType	0	Indicates the type of the guard bands. 0 Unspecified 1 Suffice for interpolation of sample values at sub-pel sample fractional locations within the coded face. 2 Represent actual picture content that is spherically adjacent to the content in the coded face at quality that gradually changes from the picture quality of the coded face to that of the spherically adjacent region. 3 Represent actual picture content that is spherically adjacent to the content in the coded face at a similar picture quality as within the coded face.
SEIGcmpGuardBandBoundaryExteriorFlag	false	Enables (true) or disables (false) the boundary guard bands.
SEIGcmpGuardBandSamplesMinus1	0	Specifies the number of guard band samples minus 1 used in the cubemap projected picture.

Table 46: Sphere Rotation SEI message encoder parameters

Option	Default	Description
SEISphereRotationEnabled	false	Enables (true) or disables (false) the insertion of sphere rotation SEI message.
SEISphereRotationCancelFlag	true	Indicates that the sphere rotation SEI message cancels the persistence (true) or follows (false).
SEISphereRotationPersistenceFlag	false	Specifies the persistence of the sphere rotation SEI message.
SEISphereRotationYaw	0	Specifies the value of the yaw rotation angle.
SEISphereRotationPitch	0	Specifies the value of the pitch rotation angle.
SEISphereRotationRoll	0	Specifies the value of the roll rotation angle.

Table 47: Region-wise packing SEI message encoder parameters

Option	Default	Description
SEIRwpEnabled	false	Enables (true) or disables (false) the insertion of region-wise packing SEI message.
SEIRwpCancelFlag	true	Indicates that RWP SEI message cancels the persistence (true) or follows (false).
SEIRwpPersistenceFlag	false	Specifies the persistence of the RWP SEI message.
SEIRwpConstituentPictureMatchingFlag	false	Specifies the RWP SEI message applies individually to each constituent picture (true) or to the projected picture (false).
SEIRwpNumPackedRegions	0	Specifies the number of packed regions when constituent picture matching flag is equal to 0.
SEIRwpProjPictureWidth	0	Specifies the width of the projected picture.

Continued...

Table 47: Region-wise packing SEI message encoder parameters (Continued)

Option	Default	Description
SEIRwpProjPictureHeight	0	Specifies the height of the projected picture.
SEIRwpPackedPictureWidth	0	Specifies the width of the packed picture.
SEIRwpPackedPictureHeight	0	Specifies the height of the packed picture.
SEIRwpTransformType		An array that specifies the rotation and mirroring to be applied to the packed regions.
SEIRwpGuardBandFlag		An array that specifies the existence of guard band in the packed regions.
SEIRwpProjRegionWidth		An array that specifies the width of the projected regions.
SEIRwpProjRegionHeight		An array that specifies the height of the projected regions.
SEIRwpGuardBandFlag		An array that specifies the existence of guard band in the packed regions.
SEIRwpProjRegionTop		An array that specifies the top sample row of the projected regions.
SEIRwpProjRegionLeft		An array that specifies the left-most sample column of the projected regions.
SEIRwpPackedRegionWidth		An array that specifies the width of the packed regions.
SEIRwpPackedRegionHeight		An array that specifies the height of the packed regions.
SEIRwpPackedRegionTop		An array that specifies the top luma sample row of the packed regions.
SEIRwpPackedRegionLeft		An array that specifies the left-most luma sample column of the packed regions.
SEIRwpLeftGuardBandWidth		An array that specifies the width of the guard band on the left side of the packed regions.
SEIRwpRightGuardBandWidth		An array that specifies the width of the guard band on the right side of the packed regions.
SEIRwpTopGuardBandHeight		An array that specifies the height of the guard band above the packed regions.
SEIRwpBottomGuardBandHeight		An array that specifies the height of the guard band below the packed regions.
SEIRwpGuardBandNotUsedForPredFlag		An array that specifies if the guard bands is used in the inter prediction process.
SEIRwpGuardBandType		An array that specifies the type of the guard bands for the packed regions.

Table 48: Omni Viewport SEI message encoder parameters

Option	Default	Description
SEIOmniViewportEnabled	false	Enables (true) or disables (false) the insertion of omni viewport SEI message.
SEIOmniViewportId	0	Contains an identifying number that may be used to identify the purpose of the one or more recommended viewport regions.
SEIOmniViewportCancelFlag	true	Indicates that the omni viewport SEI message cancels the persistence (true) or follows (false).
SEIOmniViewportPersistenceFlag	false	Specifies the persistence of the omni viewport SEI message.
SEIOmniViewportCntMinus1	0	Specifies the number of recommended viewport regions minus 1.
SEIOmniViewportAzimuthCentre		An array that indicates the centre of the i-th recommended viewport region.
SEIOmniViewportElevationCentre		An array that indicates the centre of the i-th recommended viewport region.
SEIOmniViewportTiltCentre		An array that indicates the tilt angle of the i-th recommended viewport region.
SEIOmniViewportHorRange		An array that indicates the azimuth range of the i-th recommended viewport region.
SEIOmniViewportVerRange		An array that indicates the elevation range of the i-th recommended viewport region.

Table 49: Sample Aspect Ratio Information SEI message encoder parameters

Option	Default	Description
SEISampleAspectRatioInfo	false	Enables (true) or disables (false) the insertion of Sample Aspect Ratio Information SEI message.
SEISARICancelFlag	true	Indicates that the Sample Aspect Ratio Information SEI message cancels the persistence (true) or follows (false).
SEISARIPersistenceFlag	false	Specifies the persistence of the Sample Aspect Ratio Information SEI message.
SEISARIAspectRatioIdc	0	Specifies aspect ratio IDC as defined in the standard.
SEISARISarWidth	0	Specifies the horizontal size of the sample aspect ratio, if SEISARIAspectRatioIdc is equal to 255.
SEISARISarHeight	0	Specifies the vertical size of the sample aspect ratio, if SEISARIAspectRatioIdc is equal to 255.

Table 50: Frame-Field Information SEI message encoder parameters

Option	Default	Description
SEIFrameFieldInfo	false	Enables (true) or disables (false) the insertion of Frame-Field Information SEI message.

Table 51: Subpicture Level Information SEI message encoder parameters

Option	Default	Description
SEISubpicLevelInfoEnabled	false	Enables (true) or disables (false) the insertion of Subpicture Level Information SEI message. Note, currently no other configuration options are available, because this depends on the number of subpictures, which are still not supported in the software. An example SEI with dummy values is generated, when the option is enabled.
SEISubpicLevelInfoRefLevels	none	List of reference levels to be signalled.
SEISubpicLevelInfoExplicitFraction	false	Enable signalling of explicit fraction for each level and subpicture
SEISubpicLevelInfoNumSubpics	1	Number of subpictures in context of the SEI. Has to be equal to NumSubpics.
SEISubpicLevelInfoRefLevelFractions	none	List of fractions of levels to be signalled. List is ordered by subpicture then level.

Table 52: Content light level info SEI message encoder parameters

Option	Default	Description
SEICLLEnabled	false	Enables or disables the insertion of the content light level SEI message.
SEICLLMaxContentLightLevel	4000	When not equal to 0, specifies an upper bound on the maximum light level among all individual samples in a 4:4:4 representation of red, green, and blue colour primary intensities in the linear light domain for the pictures of the CLVS, in units of candelas per square metre. When equal to 0, no such upper bound is indicated.
SEICLLMaxPicAvgLightLevel	0	When not equal to 0, specifies an upper bound on the maximum average light level among the samples in a 4:4:4 representation of red, green, and blue colour primary intensities in the linear light domain for any individual picture of the CLVS, in units of candelas per square metre. When equal to 0, no such upper bound is indicated.

Table 53: Alternative transfer characteristics SEI message encoder parameters

Option	Default	Description
SEIPreferredTransferCharacteristics	18	Indicates a preferred alternative value for the transfer_characteristics syntax element that is indicated by the colour description syntax of VUI parameters.

Table 54: Ambient viewing environment SEI message encoder parameters

Option	Default	Description
SEIAVEEnabled	false	Enables or disables the insertion of the ambient viewing environment SEI message.
SEIAVEAmbientIlluminance	100000	Specifies the environmental illuminance of the ambient viewing environment in units of 1/10000 lux. The value shall not be 0.
SEIAVEAmbientLightX	15635	Specifies the x chromaticity coordinate, according to the CIE 1931 definition, of the environmental ambient light in the nominal viewing environment in normalized increments of 1/50000. The value shall be in the range of 0 to 50,000, inclusive.
SEIAVEAmbientLightY	16450	Specifies the y chromaticity coordinate, according to the CIE 1931 definition, of the environmental ambient light in the nominal viewing environment in normalized increments of 1/50000. The value shall be in the range of 0 to 50,000, inclusive.

Table 55: Content colour volume SEI message encoder parameters

Option	Default	Description
SEICCVEnabled	false	Enables or disables the insertion of the content colour volume SEI message.
SEICCVCancelFlag	0	Specifies the persistence of any previous content colour volume SEI message in output order.
SEICCVPersistenceFlag	1	Specifies the persistence of the content colour volume SEI message for the current layer.
SEICCVPrimariesPresent	1	Specifies whether the CCV primaries are present in the content colour volume SEI message.
m_ccvSEIPrimariesX0	0.300	Specifies the x coordinate, according to the CIE 1931 definition, of the first (green) colour primary component in normalized increments of 1/50000.
m_ccvSEIPrimariesY0	0.600	Specifies the y coordinate, according to the CIE 1931 definition, of the first (green) colour primary component in normalized increments of 1/50000.
m_ccvSEIPrimariesX1	0.150	Specifies the x coordinate, according to the CIE 1931 definition, of the second (blue) colour primary component in normalized increments of 1/50000.
m_ccvSEIPrimariesY1	0.060	Specifies the y coordinate, according to the CIE 1931 definition, of the second (blue) colour primary component in normalized increments of 1/50000.
m_ccvSEIPrimariesX2	0.640	Specifies the x coordinate, according to the CIE 1931 definition, of the third (red) colour primary component in normalized increments of 1/50000.
m_ccvSEIPrimariesY2	0.330	Specifies the y coordinate, according to the CIE 1931 definition, of the third (red) colour primary component in normalized increments of 1/50000.
SEICCVMinLuminanceValuePresent	1	Specifies whether the CCV min luminance value is present in the content colour volume SEI message.
SEICCVMinLuminanceValue	0.0	specifies the CCV min luminance value in the content colour volume SEI message.
SEICCVMaxLuminanceValuePresent	1	Specifies whether the CCV max luminance value is present in the content colour volume SEI message.
SEICCVMaxLuminanceValue	0.1	specifies the CCV max luminance value in the content colour volume SEI message.
SEICCVAvgLuminanceValuePresent	1	Specifies whether the CCV avg luminance value is present in the content colour volume SEI message.
SEICCVAvgLuminanceValue	0.01	specifies the CCV avg luminance value in the content colour volume SEI message.

3.4 Hardcoded encoder parameters

Table 56: 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

```
DecoderApp -b str.bin -o dec.yuv [options]
```

Table 57: Decoder options

Option	Default	Description
(-help)		Prints usage information.
BitStreamFile (-b)		Defines the input bit stream file name.
ReconFile (-o)		Defines the reconstructed video file name. If empty, no file is generated. If the bitstream contains multiple layer and no single target layer is specified (i.e. TargetLayer=-1), a reconstructed file is written for each layer and the layer index is added as suffix to ReconFile. If one or more dots exist in the file name, the layer id is added before the last dot, e.g. 'decoded.yuv' becomes 'decoded0.yuv' for layer id 0, 'decoded' becomes 'decoded0'.
OplFile (-opl)		Defines the output log file name (*.opl file). If empty, no file is generated. Each output picture log file contains one row for each output picture in the bitstream, in output order. Each row contains the following information, as CSV: PicOrderCntVal, pic_width_max_in_luma_samples, pic_height_max_in_luma_samples, MD5 checksum for the Y component, MD5 checksum for the U component, MD5 checksum for the V component. The format of output log file is specified in JVET-P2008.
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)
TargetLayer (-p)	-1 (Native)	Specifies the target bitstream Layer to be decoded. (the value -1 indicates that decoding the whole bitstream)
SEIDecodedPictureHash	1	<p>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:</p> <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 57: Decoder options (Continued)

Option	Default	Description
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.
PYUV	false	When true, output 10-bit and 12-bit YUV data as 5-byte and 3-byte (respectively) packed YUV data. See doc/pyuv_format.pdf for details. Ignored for interlaced output.
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 OutputBitDepth 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).

5 Block statistics extension

The block statistics extension enables straightforward visualization and statistical analysis of coding tool usage in encoded bitstreams. The extension enables the reference software encoder and decoder to write out statistics files in a configurable way, which in turn can be loaded into a suitable YUV player for overlay of the reconstructed YUV sequence, or can be used for statistical analysis at a selectable scope (e.g. block/picture/sequence level). An example implementation for such visualization is available with the open-source YUView player (<https://github.com/IENT/YUView>).

5.1 Usage

The software has to be compiled with the macros `ENABLE_TRACING` and `K0149_BLOCK_STATISTICS` defined as 1. The statistics can be written by either encoder or decoder.

The extension adds additional trace channels to the “dtrace” functionality of the software. The following trace channels were added:

D_BLOCK_STATISTICS_ALL All syntax elements are written, no matter whether they are actually encoded or derived.

D_BLOCK_STATISTICS_CODED Tries to write only syntax elements, which have also been encoded.

The following additional encoder options are available (part of “dtrace”). See the file `dtrace_next.h` for more details.

Table 58: Decoder options

Option	Default	Description
TraceFile		File name of the produced trace file.
TraceRule		Specifies which traces should be saved, and for which POCs.

Concrete examples of calls for generating a block statistics file are:

```
bin/DecoderAppStatic -b str/BasketballDrive_1920x1080_QP37.vvc \  
--TraceFile="stats/BasketballDrive_1920x1080_QP37_coded.vtmbmsstats" \  
--TraceRule="D_BLOCK_STATISTICS_CODED:poc>=0"  
  
bin/DecoderAppStatic -b str/BasketballDrive_1920x1080_QP37.vvc \  
--TraceFile="stats/BasketballDrive_1920x1080_QP37_all.vtmbmsstats" \  
--TraceRule="D_BLOCK_STATISTICS_ALL:poc>=0"
```

5.2 Block statistics file formats

The trace file will contain a header listing information of all available block statistics. For each statistic it lists a type and a scale for vectors or range for integers if applicable:

```
# VTMBMS Block Statistics  
# Sequence size: [832x 480]  
# Block Statistic Type: PredMode; Flag;  
# Block Statistic Type: MergeFlag; Flag;  
# Block Statistic Type: MVL0; Vector; Scale: 4  
# Block Statistic Type: MVL1; Vector; Scale: 4  
# Block Statistic Type: IPCM; Flag;  
# Block Statistic Type: Y_IntraMode; Integer; [0, 73]  
# Block Statistic Type: Cb_IntraMode; Integer; [0, 73]
```

Two formats are available for the statistics for each block, a human readable format and a CSV based format. The header remains the same for both cases.

For both formats each row contains the information for one block statistic. The order of the data is: picture order count (POC), location of top left corner of the block, size of the block, name of the statistic, and value of the statistic. The macro BLOCK_STATS_AS_CSV is available in order to choose the required format. The human readable format can also be easily processed with other software, for example YUView, using regular expressions. The CSV based formats provides the universal interface required by spreadsheet applications.

The human readable format is based on the format used for the other dtrace statistics. Some examples for this format are:

```
BlockStat: POC 16 @( 112, 0) [ 8x 8] SkipFlag=1  
BlockStat: POC 16 @( 112, 0) [ 8x 8] InterDir=1  
BlockStat: POC 16 @( 112, 0) [ 8x 8] MergeFlag=1  
BlockStat: POC 16 @( 112, 0) [ 8x 8] MergeIdx=0  
BlockStat: POC 16 @( 112, 0) [ 8x 8] MergeType=0  
BlockStat: POC 16 @( 112, 0) [ 8x 8] MVPIdxL0=255  
BlockStat: POC 16 @( 112, 0) [ 8x 8] MVPNumL0=255  
BlockStat: POC 16 @( 112, 0) [ 8x 8] RefIdxL0=0  
BlockStat: POC 16 @( 112, 0) [ 8x 8] MVDL0={ 0, 0}  
BlockStat: POC 16 @( 112, 0) [ 8x 8] MVL0={ -70, 18}  
BlockStat: POC 16 @( 112, 8) [ 8x 8] PredMode=0  
BlockStat: POC 16 @( 112, 8) [ 8x 8] PartSize=0
```

Some examples of the CSV based format are:

```
BlockStat;16; 112; 0; 8; 8;SkipFlag;1  
BlockStat;16; 112; 0; 8; 8;InterDir;1  
BlockStat;16; 112; 0; 8; 8;MergeFlag;1  
BlockStat;16; 112; 0; 8; 8;MergeIdx;0  
BlockStat;16; 112; 0; 8; 8;MergeType;0  
BlockStat;16; 112; 0; 8; 8;MVPIdxL0;255
```

```

BlockStat;16; 112;    0; 8; 8;MVPNumL0;255
BlockStat;16; 112;    0; 8; 8;RefIdxL0;0
BlockStat;16; 112;    0; 8; 8;MVDL0;    0;    0
BlockStat;16; 112;    0; 8; 8;MVL0; -70;  18
BlockStat;16; 112;    8; 8; 8;PredMode;0
BlockStat;16; 112;    8; 8; 8;PartSize;0

```

5.3 Visualization

The block statistics can be viewed with YUView, which is freely available under GPLv3: <https://github.com/IENT/YUView>. The latest releases and the master branch have the functionality required for viewing the block statistics. YUView assumes that the file extension of block statistics file is “.vtmbmsstats”. However, if a file is not recognized you can choose from a list of supported file formats.

Statistics can be overlaid with YUV sequences. Some example snapshots are:



Figure 2: YUView

5.4 Adding statistics

In order to add further block statistics, do the following:

source/Lib/CommonLib/dtrace_blockstatistics.h Add your statistic to the BlockStatistic enum:

```

enum class BlockStatistic {
    // general
    PredMode,
    PartSize,
    Depth,

```

Further, add your statistic to the map blockstatistic2description:

```

static const std::map<BlockStatistic,
    std::tuple<std::string, BlockStatisticType, std::string>>
    blockstatistic2description =
{

```



Figure 3: Motion vectors




Figure 4: Skip flag

```

{ BlockStatistic::PredMode,
  std::tuple<std::string, BlockStatisticType, std::string>
  {"PredMode", BlockStatisticType::Flag, ""}},
{ BlockStatistic::MergeFlag,
  std::tuple<std::string, BlockStatisticType, std::string>
  {"MergeFlag", BlockStatisticType::Flag, ""}},
{ BlockStatistic::MVL0,
  std::tuple<std::string, BlockStatisticType, std::string>
  {"MVL0", BlockStatisticType::Vector, "Scale: 4"}},
YOURS

```

source/Lib/CommonLib/dtrace_blockstatistics.cpp All code for writing syntax elements is kept in this file in `getAndStoreBlockStatistics`. This function is called once for each CTU, after it has been en/decoded. The following macros have been defined to facilitate writing of block statistics:

```

DTRACE_BLOCK_SCALAR(ctx, channel, cs_cu_pu, stat_type, val)
DTRACE_BLOCK_SCALAR_CHROMA(ctx, channel, cs_cu_pu, stat_type, val)
DTRACE_BLOCK_VECTOR(ctx, channel, cu_pu, stat_type, v_x, v_y)
DTRACE_BLOCK_AFFINETF(ctx, channel, pu, stat_type, v_x0, v_y0, v_x1, v_y1, v_x2, v_y2)

```

An example:

```

DTRACE_BLOCK_SCALAR(g_trace_ctx, D_BLOCK_STATISTICS_ALL,
  cu, GetBlockStatisticName(BlockStatistic::PredMode), cu.predMode);

```

Block statistics for debugging The statistics can also be used to write out other data, not just syntax elements. Add your statistics to `dtrace_blockstatistics.h`. Where it should be used the following headers have to be included:

```

#include "dtrace_next.h"
#include "dtrace_blockstatistics.h"

```

6 Using the stream merge tool

The `StreamMergeApp` tool takes multiple single-layer (single `nuh_layer_id`) bitstreams as inputs and merge them into a multi-layer bitstream by interleaving the NALUs from the input single layer bitstreams. During the merge, the tool assigns a new unique `nuh_layer_id` for each input bitstream. Then the decoder could specify which layer bitstream to be decoded through the command line option “-p `nuh_layer_id`”.

6.1 Usage

```

StreamMergeApp ^~I<bitstream1> <bitstream2> [<bitstream3> ...] <outfile>

```

The command line options `bistreamX` specify the file names of the input single-layer bitstreams. At least two input bitstreams need to be specified. The merged multi-layer bitstream will be stored into the outfile.