

| | | |
|-------------------|--|---|
| <i>Title:</i> | VTM Software Manual | |
| <i>Status:</i> | Software AHG working document | |
| <i>Purpose:</i> | Information | |
| <i>Author(s):</i> | Frank Bossen David Flynn Xiang Li Karl Sharman Karsten Sühning | <i>frank@bossentech.com</i> <i>dflynn@blackberry.com</i> <i>xlxiangli@tencent.com</i> <i>karl.sharman@eu.sony.com</i> <i>karsten.suehring@hhi.fraunhofer.de</i> |
| <i>Source:</i> | AHG chairs | |

Abstract

This document is a user manual describing usage of the VTM reference software for the VVC project. It applies to version 5.2 of the software.

Contents

| | | |
|----------|--|-----------|
| 1 | General Information | 2 |
| 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 | 4 |
| 3 | Using the encoder | 5 |
| 3.1 | GOP structure table | 5 |
| 3.2 | Encoder parameters | 8 |
| 3.3 | Encoder SEI parameters | 20 |
| 3.4 | Hardcoded encoder parameters | 27 |
| 4 | Using the decoder | 29 |
| 4.1 | General | 29 |
| 4.2 | Using the decoder analyser | 30 |
| 5 | Block statistics extension | 30 |
| 5.1 | Usage | 30 |
| 5.2 | Block statistics file formats | 31 |
| 5.3 | Visualization | 32 |
| 5.4 | Adding statistics | 32 |
| 6 | Using the stream merge tool | 34 |
| 6.1 | Usage | 34 |

List of Tables

| | | |
|---|-----------------------|---|
| 1 | Supported compilers | 3 |
| 2 | GOP structure example | 7 |

| | | |
|----|--|----|
| 3 | File, I/O and source parameters. | 8 |
| 4 | Profile and level parameters | 10 |
| 5 | Unit definition parameters | 10 |
| 6 | Coding structure parameters | 11 |
| 7 | Motion estimation parameters | 11 |
| 8 | Mode decision parameters | 12 |
| 9 | Quantization parameters | 12 |
| 10 | Slice, tile and brick coding parameters | 14 |
| 11 | In-loop filtering parameters | 15 |
| 12 | Coding tools parameters | 16 |
| 13 | Rate control parameters | 18 |
| 14 | Encoder debug parameters | 18 |
| 15 | VUI parameters | 18 |
| 16 | Range Extensions (Version 2) tool parameters | 20 |
| 17 | List of Version 1 and RExt SEI messages | 20 |
| 18 | Buffering period SEI message encoder parameters | 21 |
| 19 | Picture timing SEI message encoder parameters | 22 |
| 20 | Recovery point SEI message encoder parameters | 22 |
| 21 | Tone mapping information SEI message encoder parameters | 22 |
| 22 | Frame packing arrangement SEI message encoder parameters | 23 |
| 23 | Display orientation SEI message encoder parameters | 24 |
| 24 | Green Metadata SEI message encoder parameters | 24 |
| 25 | Structure of pictures information SEI message encoder parameters | 24 |
| 26 | Active parameter sets SEI message encoder parameters | 24 |
| 27 | Decoding unit information SEI message encoder parameters | 24 |
| 28 | Temporal sub-layer zero index SEI message encoder parameters | 24 |
| 29 | Decoded picture hash SEI message encoder parameters | 25 |
| 30 | Scalable nesting SEI message encoder parameters | 25 |
| 31 | Region refresh information SEI message encoder parameters | 25 |
| 32 | No display SEI message encoder parameters | 25 |
| 33 | Time code SEI message encoder parameters | 25 |
| 34 | Mastering display colour volume SEI message encoder parameters | 26 |
| 35 | Segmented rectangular frame packing arrangement SEI message encoder parameters | 26 |
| 36 | Temporal motion-constrained tile sets SEI message encoder parameters | 26 |
| 37 | Chroma resampling filter hint SEI message encoder parameters | 26 |
| 38 | Knee function SEI message encoder parameters | 27 |
| 39 | Colour remapping SEI message encoder parameters | 27 |
| 40 | CommonDef.h constants | 27 |
| 41 | Decoder options | 29 |
| 42 | Decoder options | 30 |

1 General Information

Reference software is being made available to provide a reference implementation of the HEVC standard being developed by the Joint 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 JNET 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/VVCSSoftware_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 | 2015 and 2017 |
| GCC | 5.4 and 7.3 |
| 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

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

| Option | Description |
|--------------------------------|--|
| <code>--help</code> | Prints parameter usage. |
| <code>-c</code> | Defines configuration file to use. Multiple configuration files may be used with repeated <code>-c</code> options. |
| <code>--parameter=value</code> | Assigns value to a given parameter as further described below. Some parameters are also supported by shorthand “ <i>-opt</i> 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: In-loop deblocking filter parameter tcOffsetDiv2 is added to the base parameter LoopFilterTcOffset_div2 to set the final tc_offset_div2 parameter for this picture signalled in the slice segment header. The final value of tc_offset_div2 shall be an integer number in the range $-6..6$.

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

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

num_ref_pics_active_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 intended 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.

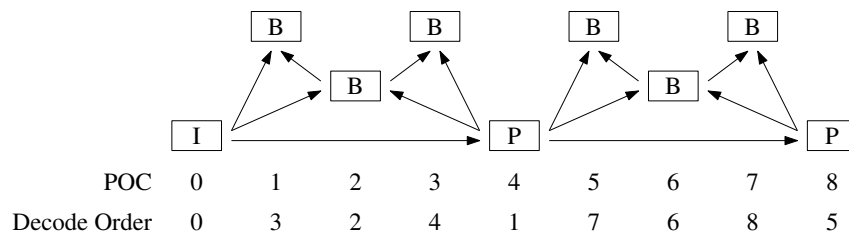
num_ref_pics_active_L1: Number of reference pictures in lists L1 that are used during coding.

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 intended 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.

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:

Table 2: GOP structure example

| | Frame1 | Frame2 | Frame3 | Frame4 |
|------------------------|--------|--------|--------|--------|
| Type | P | B | B | B |
| POC | 4 | 2 | 1 | 3 |
| QPOffset | 1 | 2 | 3 | 3 |
| QPOffsetModelOff | 0.0 | 0.0 | 0.0 | 0.0 |
| QPOffsetModelScale | 0.0 | 0.0 | 0.0 | 0.0 |
| SliceCbQPOffset | 0 | 0 | 0 | 0 |
| SliceCrQPOffset | 0 | 0 | 0 | 0 |
| QPfactor | 0.5 | 0.5 | 0.5 | 0.5 |
| tcOffsetDiv2 | 0 | 1 | 2 | 2 |
| betaOffsetDiv2 | 0 | 0 | 0 | 0 |
| temporal_id | 0 | 1 | 2 | 2 |
| num_ref_pics_active_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 |

Frame1: P 4 1 0 0 0.5 0 0 0 1 1 4 1 1 4
 Frame2: B 2 2 0 0 0.5 1 0 1 1 1 2 1 1 -2
 Frame3: B 1 3 0 0 0.5 2 0 2 1 1 1 1 2 -1 -3
 Frame4: B 3 3 0 0 0.5 2 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. |
| 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: $\left\lfloor \frac{Pel * 2^{InternalBitDepth}}{2^{MSBExtendedBitDepth}} \right\rfloor$ |
| OutputBitDepth | 0 | Specifies the bit depth of the output locally reconstructed video file. When 0, the setting defaults to the value of InternalBitDepth. Note: This option has no effect on the decoding process. |
| InputBitDepthC | 0 | Specifies the various bit-depths for chroma components. These only need to be specified if non-equal luma and chroma bit-depth processing is required. When 0, the setting defaults to the corresponding non-Chroma value. |
| MSBExtendedBitDepthC | 0 | |
| InternalBitDepthC | 0 | |
| OutputBitDepthC | 0 | |
| InputColourSpaceConvert | | The colour space conversion to apply to input video. Permitted values are: UNCHANGED No colour space conversion is applied YCbCrToYCrCb Swap the second and third components YCbCrToYYY Set the second and third components to the values in the first RGBtoGBR Reorder the three components If no value is specified, no colour space conversion is applied. The list may eventually also include RGB to YCbCr or YCgCo conversions. |
| SNRInternalColourSpace | false | When this is set true, then no colour space conversion is applied prior to PSNR calculation, otherwise the inverse of InputColourSpaceConvert is applied. |
| OutputInternalColourSpace | false | When this is set true, then no colour space conversion is applied to the reconstructed video, otherwise the inverse of InputColourSpaceConvert is applied. |
| InputChromaFormat | 420 | Specifies the chroma format used in the input file. Permitted values (depending on the profile) are 400, 420, 422 or 444. |
| ChromaFormatIDC (-cf) | 0 | Specifies the chroma format to use for processing. Permitted values (depending on the profile) are 400, 420, 422 or 444; the value of 0 indicates that the value of InputChromaFormat should be used instead. |
| MSEBasedSequencePSNR | false | When 0, the PSNR output is a linear average of the frame PSNRs; when 1, additional PSNRs are output which are formed from the average MSE of all the frames. The latter is useful when coding near-losslessly, where occasional frames become lossless. |

Continued...

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

| Option | Default | Description |
|--|---------|---|
| 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 | 0 | Add Access Unit Delimiter NAL units between all Access Units. |

Table 4: Profile and level parameters

| Option | Default | Description |
|------------------------------|---------|---|
| Profile | none | Specifies the profile to which the encoded bitstream complies. Valid HEVC Ver. 1 values are: none, main, main10, main-still-picture Valid HEVC Ver. 2 (RExt) values are: main-RExt, high-throughput-RExt, monochrome, monochrome12, monochrome16, main12, main_422_10, main_422_12, main_444, main_444_10, main_444_12, main_444_16, main_intra, main_10_intra, main_12_intra, main_422_10_intra, main_422_12_intra, main_444_intra, main_444_10_intra, main_444_12_intra, main_444_16_intra. When main-RExt is specified, the constraint flags are either manually specified, or calculated via the other supplied settings. Compatibility flags are automatically determined according to the profile. NB: There is currently only limited validation that the encoder configuration complies with the profile, level and tier constraints. |
| Level | none | Specifies the level to which the encoded bitstream complies. Valid values are: none, 1, 2, 2.1, 3, 3.1, 4, 4.1, 5, 5.1, 5.2, 6, 6.1, 6.2, 8.5 NB: There is currently only limited validation that the encoder configuration complies with the profile, level and tier constraints. |
| Tier | main | Specifies the level tier to which the encoded bitstream complies. Valid values are: main, high. NB: There is currently only limited validation that the encoder configuration complies with the profile, level and tier constraints. |
| SubProfile | 0 | Indicates interoperability metadata registered as specified by X Recommendation ITU-T T.35. |
| EnableDecodingParameterSet | false | Enables writing of a decoding parameter set. If disabled, no parameter set will be written and the special reserved ID zero will be used in the SPS indicating no constraint. |
| MaxBitDepthConstraint | 0 | For <code>-profile=main-RExt</code> , specifies the value to use to derive the <code>general_max_bit_depth</code> constraint flags for RExt profiles; when 0, use $\max(\text{InternalBitDepth}, \text{InternalBitDepth}C)$ |
| MaxChromaFormatConstraint | 0 | For <code>-profile=main-RExt</code> , specifies the chroma-format to use for the general profile constraints for RExt profiles; when 0, use the value of <code>ChromaFormatIDC</code> . |
| IntraConstraintFlag | false | For <code>-profile=main-RExt</code> , specifies the value of <code>general_intra_constraint_flag</code> to use for RExt profiles. |
| OnePictureOnlyConstraintFlag | false | For <code>-profile=main-RExt</code> , specifies the value of <code>general_one_picture_only_constraint_flag</code> to use for RExt profiles. |
| LowerBitRateConstraintFlag | true | Specifies the value of <code>general_lower_bit_constraint_flag</code> to use for RExt profiles. |
| ProgressiveSource | false | Specifies the value of <code>general_progressive_source_flag</code> |
| InterlacedSource | false | Specifies the value of <code>general_interlaced_source_flag</code> |
| NonPackedSource | false | Specifies the value of <code>general_non_packed_constraint_flag</code> |
| FrameOnly | false | Specifies the value of <code>general_frame_only_constraint_flag</code> |

Table 5: Unit definition parameters

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

Continued...

Table 5: Unit definition parameters (Continued)

| Option | Default | Description |
|-------------------------|---------|---|
| QuadtreeTUMaxDepthIntra | 1 | Defines the depth of the TU tree for intra CUs. |
| QuadtreeTUMaxDepthInter | 2 | Defines the depth of the TU tree for inter CUs. |

Table 6: Coding structure parameters

| Option | Default | Description |
|---------------------------|---------|--|
| IntraPeriod (-ip) | -1 | Specifies the intra frame period. A value of -1 implies an infinite period. |
| DecodingRefreshType (-dr) | 0 | Specifies the type of decoding refresh to apply at the intra frame period picture. 0 Applies an I picture (not a intra random access point). 1 Applies a CRA intra random access point (open GOP). 2 Applies an IDR intra random access point (closed GOP). 3 Use recovery point SEI messages to indicate random access. |
| GOPSize (-g) | 1 | Specifies the size of the cyclic GOP structure. |
| FrameN | | Multiple options that define the cyclic GOP structure that will be used repeatedly throughout the sequence. The table should contain GOPSize elements. See section 3.1 for further details. |
| ReWriteParamSets | 0 | Enable writing of parameter sets (SPS, PPS, etc.) before every (intra) random access point to enable true random access. |

Table 7: Motion estimation parameters

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

Continued...

Table 7: Motion estimation parameters (Continued)

| Option | Default | Description |
|------------|---------|--|
| MmvdDisNum | 6 | Specifies the number of MMVD distance entries used from the distance table at encoder. |

Table 8: Mode decision parameters

| Option | Default | Description |
|---|---------|---|
| LambdaModifier <i>N</i> (-LM <i>N</i>) | 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 <i>N</i> . If LambdaModifierI is specified, then LambdaModifierI will be used for intra pictures. <i>N</i> may be in the range 0 (inclusive) to 7 (exclusive). |
| LambdaModifierI (-LMI) | | Specifies one or more of the LambdaModifiers to use intra pictures at each of the temporal layers. If not present, then the LambdaModifier <i>N</i> 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 <i>GS</i> 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. |
| RDpenalty | 0 | RD-penalty for 32x32 TU for intra in non-intra slices. Enabling this parameter can reduce the visibility of CU boundaries in the coded picture. <ul style="list-style-type: none"> 0 No RD-penalty 1 RD-penalty 2 Maximum RD-penalty (no 32x32 TU) |

Table 9: Quantization parameters

| Option | Default | Description |
|---------------|---------|--|
| QP (-q) | 30.0 | Specifies the base value of the quantization parameter. If it is non-integer, the QP is switched once during encoding. |
| IntraQPOffset | 0 | Specifies a QP offset from the base QP value to be used for intra frames. |

Continued...

Table 9: Quantization parameters (Continued)

| Option | Default | Description |
|--------------------------------|---------|--|
| LambdaFromQpEnable | false | <p>When enabled, the λ, which is used to convert a cost in bits to a cost in distortion terms, is calculated as:</p> $\lambda = qpFactor \times 2^{qp+6*(bitDepthLuma-8)-12}$ <p>where qp is the slice QP and $qpFactor$ is calculated as follows:</p> $= IQF \quad \text{if } IQF \geq 0 \text{ and slice is a periodic intra slice}$ $= 0.57 \times \lambda_{scale} \quad \text{if slice is a non-periodic intra slice}$ $= \text{value from GOP table} \quad \text{otherwise}$ <p>where IQF is the value specified using the IntraQPFactor option, and where λ_{scale} is:</p> $1 \quad \text{if LambdaFromQpEnable=true}$ $1.0 - \max(0, \min(0.5, 0.05 * B)) \quad \text{if LambdaFromQpEnable=false}$ <p>where B is the number of B frames.</p> <p>If LambdaFromQpEnable=false, then the λ is also subsequently scaled for non-top-level hierarchical depths, as follows:</p> $\lambda = \lambda_{base} \times \max(2, \min(4, (sliceQP - 12)/6))$ <p>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.</p> |
| 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 $[-12, 12]$. |
| LumaLevelToDeltaQPMode | 0 | <p>Luma-level based Delta QP modulation.</p> <ul style="list-style-type: none"> 0 not used 1 Based on CTU average 2 Based on Max luma in CTU |
| LumaLevelToDeltaQPMaxValWeight | 1.0 | Weight of per block maximum luma value when LumaLevelToDeltaQPMode=2. |
| LumaLevelToDeltaQPMappingLuma | | Specify luma values to use for the luma to delta QP mapping instead of using default values. Default values are: 0, 301, 367, 434, 501, 567, 634, 701, 767, 834. |
| LumaLevelToDeltaQPMappingDQP | | Specify DQP values to use for the luma to delta QP mapping instead of using default values. Default values are: -3, -2, -1, 0, 1, 2, 3, 4, 5, 6. |
| WCGPPSEnable | 0 | <p>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 $chromaQp$ 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.</p> |
| 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 (-dqdp) | 0 | Defines maximum depth of a minimum CuDQP for sub-LCU-level delta QP. Max-CuDQPDepth shall be greater than or equal to SliceGranularity. |

Continued...

Table 9: Quantization parameters (Continued)

| Option | Default | Description |
|---------------------------------------|---------|---|
| 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 . |
| AdaptiveQp (-aq) | false | Enable or disable QP adaptation based upon a psycho-visual model. |
| MaxQPAdaptationRange (-aqr) | 6 | Specifies the maximum QP adaptation range. |
| AdaptiveQpSelection (-aqps) | false | Specifies whether QP values for non-I frames will be calculated on the fly based on statistics of previously coded frames. |
| RecalculateQP... AccordingToLambda | false | Recalculate QP values according to lambda values. Do not suggest to be enabled in all intra case. |
| ScalingList | 0 | Controls the specification of scaling lists: 0 Scaling lists are disabled 1 Use default scaling lists 2 Scaling lists are specified in the file indicated by ScalingListFile |
| ScalingListFile | | When ScalingList is set to 2, this parameter indicates the name of the file, which contains the defined scaling lists. If ScalingList is set to 2 and this parameter is an empty string, information on the format of the scaling list file is output and the encoder stops. |
| MaxCUChromaQpAdjustmentDepth | -1 | Specifies the maximum depth for CU chroma QP adjustment; if negative, CU chroma QP adjustment is disabled. |

Table 10: Slice, tile and brick coding parameters

| Option | Default | Description |
|-------------------------|---------|--|
| SliceMode | 0 | Controls the slice partitioning method in conjunction with SliceArgument. 0 Single slice 1 (deprecated) Maximum number of CTUs per slice 2 (deprecated) Maximum number of bytes per slice 3 Maximum number of tiles per slice 4 One slice per brick |
| SliceArgument | | (deprecated) Specifies the maximum number of CTUs, bytes or tiles in a slice depending on the SliceMode setting. |
| RectSliceFlag | 1 | Controls the slice shape method in conjunction with SliceMode, SliceArgument, and Tile/Brick configurations. 0 1 |
| NumRecSlicesInPicMinus1 | 0 | NOTE: When SliceMode is equal to 3, RectSliceFlag is equal to 1, and there is more than one tiles/bricks. Specifies the number of rectangular slices in the picture. |

Continued...

Table 10: Slice, tile and brick coding parameters (Continued)

| Option | Default | Description |
|--|---------|---|
| RectSlicesBoundaryArray | | <p>Specifies a space or comma separated list of top-left brick index and bottom-right brick index of rectangular slices. The top-left brick index and bottom-right brick index corresponds to the top left rectangular slice in the picture. The rest of indices corresponds to the each rectangular slices in the picture in the rectangular slice raster scan in the picture, respectively.</p> <p>For example, when the picture is partitioned into 16 tiles (4 tile columns and 4 tile rows), each tile is not further partitioned into bricks, SliceMode is equal to 3, SliceArgument is equal to 4, and NumRecSlicesInPicMinus1 is equal to 3, the values of RectSlicesBoundaryArray shall be as followss: 0 5 2 7 8 13 10 15.</p> <p>First slice has top-left brick index 0 and bottom-right brick index 5 Second slice has top-left brick index 2 and bottom-right brick index 7 Third slice has top-left brick index 8 and bottom-right brick index 13 Fourth slice has top-left brick index 10 and bottom-right brick index 15</p> |
| 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. |
| TileUniformSpacing | false | <p>Controls the mode used to determine per row and column tile sizes.</p> <p>0 Each tile column width and tile row height is explicitly set by TileColumnWidthArray and TileRowHeightArray respectively</p> <p>1 Tile columns and tile rows are uniformly spaced.</p> |
| NumTileColumnsMinus1 NumTileRowsMinus1 | 0 | Specifies the tile based picture partitioning geometry as $\text{NumTileColumnsMinus1} + 1 \times \text{NumTileRowsMinus1} + 1$ columns and rows. |
| TileColumnWidthArray TileRowHeightArray | | Specifies a space or comma separated list of widths and heights, respectively, of each tile column or tile row. The first value in the list corresponds to the leftmost tile column or topmost tile row. |
| BrickSplit1 BrickSplit2 BrickSplitN | | <p>Specifies the splitting of tiles into bricks, N can be in the range of 1 to 128 and has no further meaning except to distinguish different splitting parameters.</p> <p>The following options are supported:</p> <p>BrickSplitN tileIdx uniform (numplits) heights</p> <p>tileIdx the index of the tile to be split in raster scan order starting with zero</p> <p>uniform specifies whether uniform splitting shall be used (1 for uniform, 0 non-uniform)</p> <p>numplits specifies the number of splits, if uniform is not used</p> <p>heights specifies the height of the uniform splits or a list of heights for non-uniform splits</p> |
| TileRowHeightArray | | Specifies a space or comma separated list of widths and heights, respectively, of each tile column or tile row. The first value in the list corresponds to the leftmost tile column or topmost tile row. |

Table 11: In-loop filtering parameters

| Option | Default | Description |
|--------------------------|---------|--|
| LoopFilterDisable | false | Enables or disables the in-loop deblocking filter. |
| LFCrossSliceBoundaryFlag | true | Enables or disables the use of in-loop filtering across slice boundaries. |
| LoopFilterOffsetInPPS | false | <p>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.</p> <p>If deblocking filter parameters are sent in PPS, the same values of deblocking filter parameters are used for all pictures in the sequence (i.e. deblocking parameter = base parameter value). If deblocking filter parameters are sent in the slice segment header, varying deblocking filter parameters can be specified by setting parameters tcOffsetDiv2 and betaOffsetDiv2 in the GOP structure table. In this case, the final value of the deblocking filter parameter sent for a certain GOP picture is equal to (base parameter + GOP parameter for this picture). Intra-pictures use the base parameters values.</p> |
| LoopFilterTcOffset_div2 | 0 | Specifies the base value for the in-loop deblocking filter parameter tc_offset_div2. The final value of tc_offset_div2 shall be an integer number in the range $-6..6$. |

Continued...

Table 11: In-loop filtering parameters (Continued)

| Option | Default | Description |
|---|---------|--|
| LoopFilterBetaOffset_div2 | 0 | Specifies the base value for the in-loop deblocking filter parameter beta_offset_div2. The final value of beta_offset_div2 shall be an integer number in the range $-6..6$. |
| DeblockingFilterMetric | 0 | Specifies the use of a deblocking filter metric to evaluate the suitability of deblocking. If non-zero then LoopFilterOffsetInPPS and LoopFilterDisable must be 0. Currently excepted values are 0, 1 and 2. |
| LFCrossSliceBoundaryFlag | true | Enables or disables the use of a deblocking across tile boundaries. |
| LoopFilterAcrossVirtualBoundariesDisabledFlag | false | Disables in-loop filtering operations across the virtual boundaries. |
| 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 |

Table 12: Coding tools parameters

| Option | Default | Description |
|-------------------------------------|---------|--|
| MIP | true | Enables or disables the use of matrix-based intra prediction (MIP). |
| AMP | true | Enables or disables the use of asymmetric motion partitions. |
| ISP | false | Enables or disables the Intra Sub-Partitions coding mode. |
| ISPFast | false | Enables or disables reduced testing of non-DCT-II transforms if ISP is likely to become the best mode for a given CU. This option has no effect if either ISP or MTS are disabled. |
| SAO | true | Enables or disables the sample adaptive offset (SAO) filter. |
| TestSAODisableAtPictureLevel | false | Enables the testing of disabling SAO at the picture level after having analysed all blocks. |
| SaoEncodingRate | 0.75 | When >0 SAO early picture termination is enabled for luma and chroma. |
| SaoEncodingRateChroma | 0.5 | The SAO early picture termination rate to use for chroma (when m_SaoEncodingRate is >0). If ≤ 0 , use results for luma. |
| SAOLcuBoundary | false | Enables or disables SAO parameter estimation using non-deblocked pixels for LCU bottom and right boundary areas. |
| SAOResetEncoderStateAfterIRAP | false | When true, resets the encoder's SAO state after an IRAP (POC order). |
| ConstrainedIntraPred | false | Enables or disables constrained intra prediction. Constrained intra prediction only permits samples from intra blocks in the same slice as the current block to be used for intra prediction. |
| FastUDIUseMPMEnabled | true | If enabled, adapt intra direction search, accounting for MPM |
| FastMEForGenBLowDelayEnabled | true | If enabled use a fast ME for generalised B Low Delay slices |
| UseBLambdaForNonKeyLowDelayPictures | true | Enables use of B-Lambda for non-key low-delay pictures |
| TransquantBypassEnable | false | Enables or disables the ability to bypass the transform, quantization and filtering stages at CU level. This option corresponds to the value of transquant_bypass_enabled_flag that is transmitted in the PPS. See CUTransquantBypassFlagForce for further details. |
| CUTransquantBypassFlagForce | 0 | Controls the per CU transformation, quantization and filtering mode decision. This option controls the value of the per CU cu_transquant_bypass_flag. 0 Bypass is searched on a CU-by-CU basis and will be used if the cost is lower than not bypassing. 1 Bypass is forced for all CUs. This option has no effect if TransquantBypassEnable is disabled. |

Continued...

Table 12: Coding tools parameters (Continued)

| Option | Default | Description |
|-----------------------------|------------------------|--|
| 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. |
| WPMMethod (-wpM) | 0 | Sets the Weighted Prediction method to be used. 0 Image DC based method with joint colour component decision. 1 Image DC based method with separate colour component decision. 2 DC + Histogram refinement method (no clipping). 3 DC + Histogram refinement method (with clipping). 4 DC + Dual Histogram refinement method (with clipping). |
| Log2ParallelMergeLevel | 2 | Defines the PPS-derived Log2ParMrgLevel variable. |
| SignHideFlag (-SBH) | true | If enabled specifies that for each 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. |
| StrongIntraSmoothing (-sis) | true | If enabled specifies that for 32×32 intra prediction block, the intra smoothing when applied is either the 1:2:1 smoothing filter or a stronger bi-linear interpolation filter. Key reference sample values are tested and if the criteria is satisfied, the stronger intra smoothing filter is applied. If disabled, the intra smoothing filter when applied is the 1:2:1 smoothing filter. |
| TMVPMODE | 1 | Controls the temporal motion vector prediction mode. 0 Disabled for all slices. 1 Enabled for all slices. 2 Disabled only for the first picture of each GOPSize. |
| TransformSkip | false | Enables or disables transform-skipping mode decision. |
| TransformSkipFast | false | Enables or disables reduced testing of the transform-skipping mode decision for chroma TUs. When enabled, no RDO search is performed for chroma TUs, instead they are transform-skipped if the four corresponding luma TUs are also skipped. This option has no effect if TransformSkip is disabled. |
| 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. |
| SMVD | false | Enables or disables symmetric MVD mode. |
| RDPCM | false | Enables or disables RDPCM coding 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). |

Table 13: Rate control parameters

| Option | Default | Description |
|----------------------|---------|---|
| RateControl | false | Rate control: enables rate control or not. |
| TargetBitrate | 0 | Rate control: target bitrate, in bps. |
| KeepHierarchicalBit | 0 | Rate control: 0: equal bit allocation among pictures; 1: fix ratio hierarchical bit allocation; 2: adaptive hierarchical ratio bit allocation. It is suggested to enable hierarchical bit allocation for hierarchical-B coding structure. |
| LCULevelRateControl | true | Rate control: true: LCU level RC; false: picture level RC. |
| RCLCUSeparateModel | true | Rate control: use LCU level separate R-lambda model or not. When LCULevelRateControl is equal to false, this parameter is meaningless. |
| InitialQP | 0 | Rate control: initial QP value for the first picture. 0 to auto determine the initial QP value. |
| RCForceIntraQP | false | Rate control: force intra QP to be equal to initial QP or not. |
| RCCpbSaturation | false | Rate control: enable target bits saturation to avoid CPB overflow and underflow or not. |
| RCCpbSize | 0 | Rate control: CPB size, in bps. |
| RCInitialCpbFullness | 0.9 | Rate control: ratio of initial CPB fullness per CPB size. (InitialCpbFullness/CpbSize) RCInitialCpbFullness should be smaller than or equal to 1. |

Table 14: 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. |
| 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 15: 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. |

Continued...

Table 15: VUI parameters (Continued)

| Option | Default | Description |
|--|---------|---|
| SarHeight | 0 | Specifies the vertical size of the sample aspect ratio. |
| OverscanInfoPresent | false | Signals whether <code>overscan_info_present_flag</code> is present. |
| OverscanAppropriate | false | Indicates whether cropped decoded pictures are suitable for display using overscan. 0 Indicates that the decoded pictures should not be displayed using overscan. 1 Indicates that the decoded pictures may be displayed using overscan. |
| VideoSignalTypePresent | false | Signals whether <code>video_format</code> , <code>video_full_range_flag</code> , and <code>colour_description_present_flag</code> are present. |
| VideoFormat | 5 | Indicates representation of pictures. |
| VideoFullRange | false | Indicates the black level and range of luma and chroma signals. 0 Indicates that the luma and chroma signals are to be scaled prior to display. 1 Indicates that the luma and chroma signals are not to be scaled prior to display. |
| ColourDescriptionPresent | false | Signals whether <code>colour_primaries</code> , <code>transfer_characteristics</code> and <code>matrix_coefficients</code> are present. |
| ColourPrimaries | 2 | Indicates chromaticity coordinates of the source primaries. |
| TransferCharacteristics | 2 | Indicates the opto-electronic transfer characteristics of the source. |
| MatrixCoefficients | 2 | Describes the matrix coefficients used in deriving luma and chroma from RGB primaries. |
| ChromaLocInfoPresent | false | Signals whether <code>chroma_sample_loc_type_top_field</code> and <code>chroma_sample_loc_type_bottom_field</code> are present. |
| ChromaSampleLocTypeTopField | 0 | Specifies the location of chroma samples for top field. |
| ChromaSampleLocTypeBottomField | 0 | Specifies the location of chroma samples for bottom field. |
| NeutralChromaIndication | false | Indicates that the value of all decoded chroma samples is equal to $1 \ll (\text{BitDepthCr} - 1)$. |
| DefaultDisplayWindowFlag | flag | Indicates the presence of the Default Window parameters. false Disabled true Enabled |
| DefDispWinLeftOffset DefDispWinRightOffset DefDispWinTopOffset DefDispWinBottomOffset | 0 | Specifies the horizontal and vertical offset to be applied to the input video from the conformance window in luma samples. Must be a multiple of the chroma resolution (e.g. a multiple of two for 4:2:0). |
| FrameFieldInfoPresentFlag | false | Specifies the value of the VUI syntax element ‘ <code>frame_field_info_present_flag</code> ’, which indicates that <code>pic_struct</code> and field coding related values are present in picture timing SEI messages. |
| PocProportionalToTimingFlag | false | Specifies the value of the VUI syntax element ‘ <code>vui_poc_proportional_to_timing_flag</code> ’, which indicates that the POC value is proportional to the output time with respect to the first picture in the CVS. |
| NumTicksPocDiffOneMinus | 0 | Specifies the value of the VUI syntax element ‘ <code>vui_num_ticks_poc_diff_one_minus1</code> ’, which specifies the number of clock ticks corresponding to a difference of picture order count values equal to 1, and is used only when <code>PocProportionalToTimingFlag</code> is true. |
| BitstreamRestriction | false | Signals whether bitstream restriction parameters are present. |
| TilesFixedStructure | false | Indicates that each active picture parameter set has the same values of the syntax elements related to tiles. |
| MotionVectorsOverPicBoundaries | false | Indicates that no samples outside the picture boundaries are used for inter prediction. |
| MaxBytesPerPicDenom | 2 | Indicates a number of bytes not exceeded by the sum of the sizes of the VCL NAL units associated with any coded picture. |
| MaxBitsPerMinCuDenom | 1 | Indicates an upper bound for the number of bits of <code>coding_unit()</code> data. |
| Log2MaxMvLengthHorizontal | 15 | Indicate the maximum absolute value of a decoded horizontal MV component in quarter-pel luma units. |
| Log2MaxMvLengthVertical | 15 | Indicate the maximum absolute value of a decoded vertical MV component in quarter-pel luma units. |

Table 16: Range Extensions (Version 2) tool parameters

| Option | Default | Description |
|------------------------------------|---------|---|
| CostMode | lossy | Specifies the cost mode to use. lossy $cost = distortion + \lambda \times bits$ sequence_level_lossless $cost = distortion / \lambda + bits$ lossless As with sequence_level_lossless, but QP is also set to 0 (this will be deprecated in the future) mixed_lossless_lossy As with sequence_level_lossless, but QP=4 is used for pre-estimates of transquant-bypass blocks |
| 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. |
| CrossComponentPrediction | false | When true, specifies the use of the cross component prediction tool (4:4:4 processing only). Version 1 and some Version 2 (RExt) profiles require this to be false. |
| ReconBasedCrossCPredictionEstimate | false | If true, then when determining the alpha value for cross-component prediction, use the reconstructed residual rather than the pre-transform encoder-side residual |
| SaoLumaOffsetBitShift | 0 | Specifies the shift to apply to the SAO parameters. If negative, an estimate will be calculated based upon the initial QP. Version 1 and some Version 2 (RExt) profiles require this to be 0. |
| SaoChromaOffsetBitShift | 0 | |
| TransformSkipLog2MaxSize | 2 | Specifies the maximum TU size for which transform-skip can be used; the minimum value is 2. Version 1 and some Version 2 (RExt) profiles require this to be 2. |
| ImplicitResidualDPCM | false | When true, specifies the use of the implicitly signalled residual RDPCM tool (for intra). Version 1 and some Version 2 (RExt) profiles require this to be false. |
| ExplicitResidualDPCM | false | When true, specifies the use of the explicitly signalled residual RDPCM tool (for intra-block-copy and inter). Version 1 and some Version 2 (RExt) profiles require this to be false. |
| ResidualRotation | false | When true, specifies the use of the residual rotation tool. Version 1 and some Version 2 (RExt) profiles require this to be false. |
| SingleSignificanceMapContext | false | When true, specifies the use of a single significance map context for transform-skipped and transquant-bypassed TUs. Version 1 and some Version 2 (RExt) profiles require this to be false. |
| GolombRiceParameterAdaptation | false | When true, enable the adaptation of the Golomb-Rice parameter over the course of each slice. Version 1 and some Version 2 (RExt) profiles require this to be false. |
| AlignCABACBeforeBypass | false | When true, align the CABAC engine to a defined fraction of a bit prior to coding bypass data (including sign bits) when 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 17: List of Version 1 and RExt SEI messages

| SEI Number | SEI Name | Table number of encoder controls, if available |
|--------------|------------------|--|
| 0 | Buffering period | Table 18 |
| 1 | Picture timing | Table 19 |
| Continued... | | |

Table 17: List of Version 1 and RExt SEI messages (Continued)

| SEI Number | SEI Name | Table number of encoder controls, if available |
|------------|---|--|
| 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 20 |
| 9 | Scene information | (Not handled) |
| 15 | Picture snapshot | (Not handled) |
| 16 | Progressive refinement segment start | (Not handled) |
| 17 | Progressive refinement segment end | (Not handled) |
| 19 | Film grain characteristics | (Not handled) |
| 22 | Post-filter hint | (Not handled) |
| 23 | Tone mapping information | Table 21 |
| 45 | Frame packing arrangement | Table 22 |
| 47 | Display orientation | Table 23 |
| 56 | Green Metadata | Table 24 |
| 128 | Structure of pictures information | Table 25 |
| 129 | Active parameter sets | Table 26 |
| 130 | Decoding unit information | Table 27 |
| 131 | Temporal sub-layer zero index | Table 28 |
| 132 | Decoded picture hash | Table 29 |
| 133 | Scalable nesting | Table 30 |
| 134 | Region refresh information | Table 31 |
| 135 | No display | Table 32 |
| 136 | Time code | Table 33 |
| 137 | Mastering display colour volume | Table 34 |
| 138 | Segmented rectangular frame packing arrangement | Table 35 |
| 139 | Temporal motion-constrained tile sets | Table 36 |
| 140 | Chroma resampling filter hint | Table 37 |
| 141 | Knee function information | Table 38 |
| 142 | Colour remapping information | Table 39 |
| 143 | Deinterlaced field identification | (Not handled) |

Table 18: 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 19: 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 20: Recovery point SEI message encoder parameters

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

Table 21: Tone mapping information SEI message encoder parameters

| Option | Default | Description |
|--------------------------------|---------|--|
| SEIToneMappingInfo | 0 | Enables or disables the insertion of the Tone Mapping SEI message. |
| SEIToneMapId | 0 | Specifies Id of Tone Mapping SEI message for a given session. |
| SEIToneMapCancelFlag | false | Indicates that Tone Mapping SEI message cancels the persistence or follows. |
| SEIToneMapPersistenceFlag | true | Specifies the persistence of the Tone Mapping SEI message. |
| SEIToneMapCodedDataBitDepth | 8 | Specifies Coded Data BitDepth of Tone Mapping SEI messages. |
| SEIToneMapTargetBitDepth | 8 | Specifies Output BitDepth of Tome mapping function. |
| SEIToneMapModelId | 0 | Specifies Model utilized for mapping coded data into target_bit_depth range. 0 linear mapping with clipping 1 sigmoidal mapping 2 user-defined table mapping 3 piece-wise linear mapping 4 luminance dynamic range mapping |
| SEIToneMapMinValue | 0 | Specifies the minimum value in mode 0. |
| SEIToneMapMaxValue | 1023 | Specifies the 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 96 108 120 132 144 156 168 180 192 192 196 204 208 216 220 228 232 240 248 252 260 264 272 276 284 292 292 296 300 304 308 312 320 324 328 332 336 344 348 352 356 360 368 372 376 380 384 388 396 400 404 408 412 420 424 428 432 436 444 444 444 448 452 456 460 464 468 472 476 476 480 484 488 492 496 500 504 508 508 512 516 520 524 528 532 536 540 540 544 548 552 556 560 564 568 572 572 576 580 584 588 592 596 600 604 604 608 612 616 620 624 628 632 636 636 640 644 648 652 656 660 664 668 672 672 672 676 680 680 684 688 692 692 696 700 704 704 708 712 716 716 720 724 724 728 732 736 736 740 744 748 748 752 756 760 760 764 768 768 772 776 780 780 784 788 792 792 796 800 804 804 808 812 812 816 820 824 824 828 832 836 836 840 844 848 848 852 856 860 860 860 864 864 868 872 872 876 880 880 884 884 888 892 892 896 900 900 904 908 908 912 912 916 920 920 924 928 928 932 936 936 940 940 944 948 948 952 956 956 960 964 964 968 968 972 976 976 980 984 984 988 992 992 996 996 1000 1004 1004 1008 1012 1012 1016 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 |

Continued...

Table 21: Tone mapping information SEI message encoder parameters (Continued)

| Option | Default | Description |
|---|---------|--|
| SEIToneMapTargetPivotValue | | Array of target pivot point in mode 3. A suggested table is: 48 73 111 168 215 |
| SEIToneMap... CameraIsoSpeedIdc | 0 | Indicates the camera ISO speed for daylight illumination. |
| SEIToneMap... CameraIsoSpeedValue | 400 | Specifies the camera ISO speed for daylight illumination of Extended_ISO. |
| SEIToneMap... ExposureIndexIdc | 0 | Indicates the exposure index setting of the camera. |
| SEIToneMap... ExposureIndexValue | 400 | Specifies the exposure index setting of the cameran of Extended_ISO. |
| SEIToneMapExposure... CompensationValueSignFlag | 0 | Specifies the sign of ExposureCompensationValue. |
| SEIToneMapExposure... CompensationValueNumerator | 0 | Specifies the numerator of ExposureCompensationValue. |
| SEIToneMapExposure... CompensationValueDenomIdc | 2 | Specifies the denominator of ExposureCompensationValue. |
| SEIToneMapRef... ScreenLuminanceWhite | 350 | Specifies reference screen brightness setting in units of candela per square metre. |
| SEIToneMapExtended... RangeWhiteLevel | 800 | Indicates the luminance dynamic range. |
| SEIToneMapNominal... BlackLevelLumaCodeValue | 16 | Specifies luma sample value of the nominal black level assigned decoded pictures. |
| 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 22: Frame packing arrangement SEI message encoder parameters

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

Table 23: 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 24: 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 25: 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 26: 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 27: 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 28: 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 29: 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 30: Scalable nesting SEI message encoder parameters

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

Table 31: 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 32: 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 33: 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). |

Continued...

Table 33: Time code SEI message encoder parameters (Continued)

| Option | Default | Description |
|-------------------------|---------|--|
| 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 34: 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 35: 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 SRPFA 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 36: 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 37: 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. |

Continued...

Table 37: Chroma resampling filter hint SEI message encoder parameters (Continued)

| Option | Default | Description |
|---|---------|--|
| 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 38: 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 39: 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. |

3.4 Hardcoded encoder parameters

Table 40: 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. |

Continued...

Table 40: CommonDef.h constants (Continued)

| Option | Default | Description |
|------------------------|---------|---|
| MAX_GOP | 64 | maximum size of value of hierarchical GOP. |
| MAX_NUM_REF | 4 | maximum number of multiple reference frames |
| MAX_NUM_REF_LC | 8 | maximum number of combined reference frames |
| AMVP_MAX_NUM_CANDS | 2 | maximum number of final candidates |
| AMVP_MAX_NUM_CANDS_MEM | 3 | |
| MRG_MAX_NUM_CANDS | 5 | |
| DYN_REF_FREE | off | dynamic free of reference memories |
| MAX_TLAYER | 8 | maximum number of temporal layers |
| ADAPT_SR_SCALE | on | division factor for adaptive search range |
| EARLY_SKIP_THRES | 1.5 | early skip if $RD < EARLY_SKIP_THRES * avg[BestSkipRD]$ |
| MAX_NUM_REF_PICS | 16 | |
| MAX_CHROMA_FORMAT_IDC | 3 | |

TypeDef.h

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

4 Using the decoder

4.1 General

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

Table 41: Decoder options

| Option | Default | Description |
|----------------------------------|----------------|--|
| (-help) | | Prints usage information. |
| BitStreamFile (-b) | | Defines the input bit stream file name. |
| ReconFile (-o) | | Defines reconstructed YUV file name. If empty, no file is generated. |
| SkipFrames (-s) | 0 | Defines the number of pictures in decoding order to skip. |
| MaxTemporalLayer (-t) | -1 | Defines the maximum temporal layer to be decoded. If -1, then all layers are decoded. |
| TarDecLayerIdSetFile (-l) | | Specifies the targetDecLayerIdSet file name. The file would contain white-space separated LayerId values of the layers that are to be decoded. Omitting the parameter, or using a value of -1 in the file decodes all layers. |
| OutputBitDepth (-d) | 0 (Native) | Specifies the luma bit-depth of the reconstructed YUV file (the value 0 indicates that the native bit-depth is used) |
| OutputBitDepthC | 0 (Native) | Defines the chroma bit-depth of the reconstructed YUV file (the value 0 indicates that the native bit-depth is used) |
| TargetLayer (-p) | -1 (Native) | Specifies the target bitstream Layer to be decoded. (the value -1 indicates that decoding the whole bitstream) |
| SEIDecodedPictureHash | 1 | Enable or disable verification of any Picture hash SEI messages. When this parameter is set to 0, the feature is disabled and all messages are ignored. When set to 1 (default), the feature is enabled and the decoder has the following behaviour: <ul style="list-style-type: none">• If Picture hash SEI messages are included in the bit stream, the same type of hash is calculated for each decoded picture and written to the log together with an indication whether the calculated value matches the value in the SEI message. Decoding will continue even if there is a mismatch.• After decoding is complete, if any MD5sum comparison failed, a warning is printed and the decoder exits with the status EXIT_FAILURE• The per-picture MD5 log message has the following formats: [MD5:d41d8cd98f00b204e9800998ecf8427e,(OK)], [MD5:d41d8cd98f00b204e9800998ecf8427e,(unk)], [MD5:d41d8cd98f00b204e9800998ecf8427e,(***ERROR***)] [rxMD5:b9e1...] where, “(unk)” implies that no MD5 was signalled for this picture, “(OK)” implies that the decoder agrees with the signalled MD5, “(***ERROR***)” implies that the decoder disagrees with the signalled MD5. “[rxMD5:...]” is the signalled MD5 if different. |
| OutputDecodedSEIMessagesFilename | | When a non-empty file name is specified, information regarding any decoded SEI messages will be output to the indicated file. If the file name is '-', then stdout is used instead. |
| SEIColourRemappingInfoFilename | | Specifies that the colour remapping SEI message should be applied to the output video, with the output written to this file. If no value is specified, the SEI message is ignored and no mapping is applied. |
| RespectDefDispWindow (-w) | 0 | Video region to be output by the decoder. 0 Output content inside the conformance window. 1 Output content inside the default window. |
| OutputColourSpaceConvert | | Specifies the colour space conversion to apply to 444 video. Permitted values are: UNCHANGED No colour space conversion is applied YCrCbToYCbCr Swap the second and third components GBRtoRGB Reorder the three components If no value is specified, no colour space conversion is applied. The list may eventually also include RGB to YCbCr or YCbCo conversions. |

Continued...

Table 41: Decoder options (Continued)

| Option | Default | Description |
|------------------------------|---------|--|
| 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 42: 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:

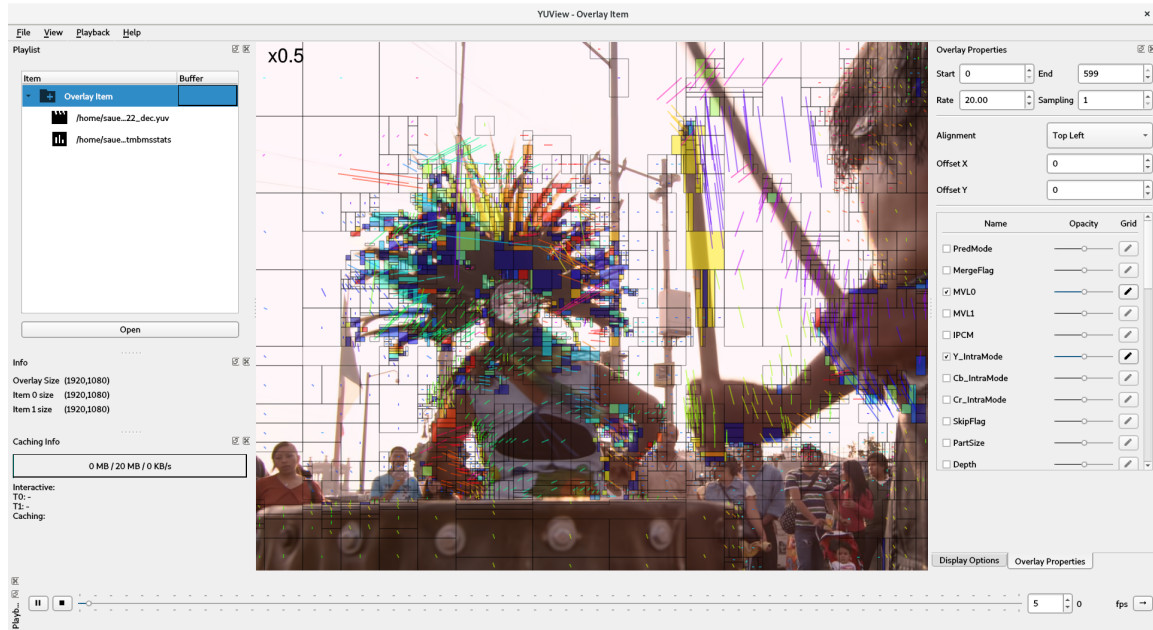


Figure 2: YUView



Figure 3: Motion vectors

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:

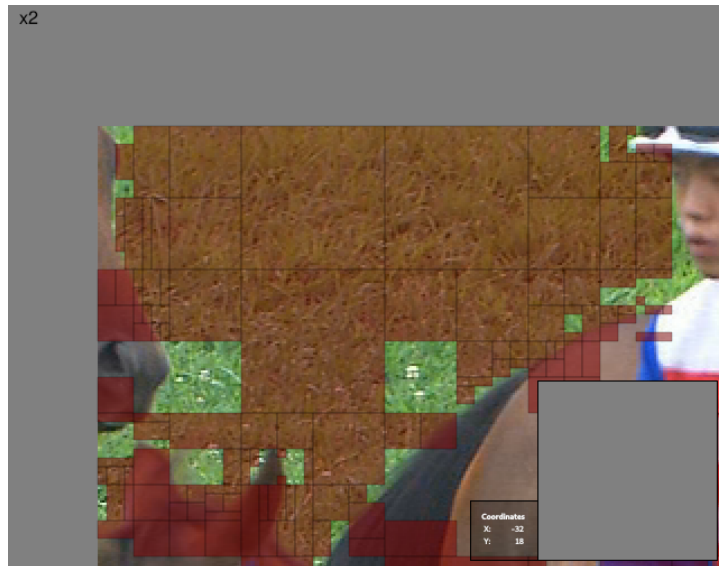


Figure 4: Skip flag

```
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 =
{
    { 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 bistream 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 bistream will be stored into the outfile.