Intel® Video Pro Analyzer

Version 2.3.0

LEGAL DISCLAIMER

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web Site.

MPEG is an international standard for video compression/decompression promoted by ISO. Implementations of MPEG CODECs, or MPEG enabled platforms may require licenses from various entities, including Intel Corporation.

Intel, the Intel logo, Intel Core are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

*Other names and brands may be claimed as the property of others.

Revision History:

Revision Number	Date	Description
1.0	10/2/13	Initial Draft
1.1	1/2/14	Updated text and screenshots for new features
1.2	7/11/14	Updated text for C++ version
1.2.1	8/28/14	Added descriptions for new feature, updated pictures
1.3.0	10/3/14	Updated documentation for 1.3.0 release
1.3.5	12/10/14	Updated documentation for 1.3.5 release
1.4.0	02/23/15	Added description for AVC
1.5.0	03/27/15	Added description for MPEG2, updated AVC/HEVC/VP9 descriptions
1.5.2	05/25/15	Updated index, small corrections
1.6.0	06/16/15	Updated screenshots with new GUI and new options description
2.0.0	08/18/15	Updated AVC description and screenshots with new GUI and options
2.1.0	11/11/15	Added description of dual view mode and other updates
2.2.0	03/18/16	Updated manual with new features
2.3.0	07/18/16	Added descript of indexing mode and some new features

Contents

1 Overview

Intel Video Pro Analyzer is a graphical coded video bitstream analysis tool, supporting the following coding standards:

- HEVC: (ISO/IEC 23008-2 MPEG-H Part 2 or ITU-T H.265) , 8/10-bit
- HEVC: RExt extension, 8/10/12-bit, 4:0:0/4:2:0/4:2:2/4:4:4
- HEVC: SCC extension, conform to HM + SCM 8.0 in reference code
- HEVC Scalable/Multiview Extension
- Google's VP9, profiles 0,1,2,3, 4:2:0/4:2:2/4:4:0/4:4:4, 8/10/12-bit
- AVC: (H.264/AVC, ISO/IEC 14496-10, MPEG-4 Part 10), except SVC/MVC
- MPEG2 (ISO/IEC 13818-2 Part 2), 4:2:0/4:2:2, 8-bit
- MKV container
- MP4 container
- MPEG2 TS/PS container
- AVI container
- Mpeg Media Transport container

Once a bitstream is loaded, the tool allows the user to inspect each major step of the decode process visually and numerically, and the structure of the coded image can be explored. This data can be used as a visual reference when learning about HEVC/VP9/AVC/MPEG2 or when debugging a particular encoder or decoder.

Intel Video Pro Analyzer is written in C++. There are separate packages for each supported operation system – Windows*, Linux*, Mac* OS X*.

2 Document Conventions, Symbols and Terms

Term	Definition
CABAC	Context-adaptive binary arithmetic coding
CTB	Coding Tree Block
CU	Coding Unit
CVC	Coded Video Stream
GUI/UI	(Graphical) User Interface
HEVC	High Efficiency Video Coding
NAL	Network Abstraction Layer
PPS	Picture Parameter Set
PU	Prediction Unit
QM	Quantization Matrix
SAO	Sample adaptive offset
SEI	Supplemental Enhancement Information
SPS	Sequence Parameter Set
TU	Transform Unit
VP9	Google video codec, VP9 is a successor to VP8.
VPS	Video Parameter Set
YUV	Color space (YUV)

3 HEVC

The following sections describe all available features when loading an HEVC bitstream. The supported format of a bitstream is the raw bitstream with no surrounding container. Output from the publicly available HEVC reference software HM is in this format and can be opened directly. Intel Video Pro Analyzer supports Main, Main 10, and Main Still Picture profiles, HEVC REext profiles (including SCC), SHVC profile and HEVC MultiView profiles. Sequences that go beyond the profile limits maybe be supported as well.

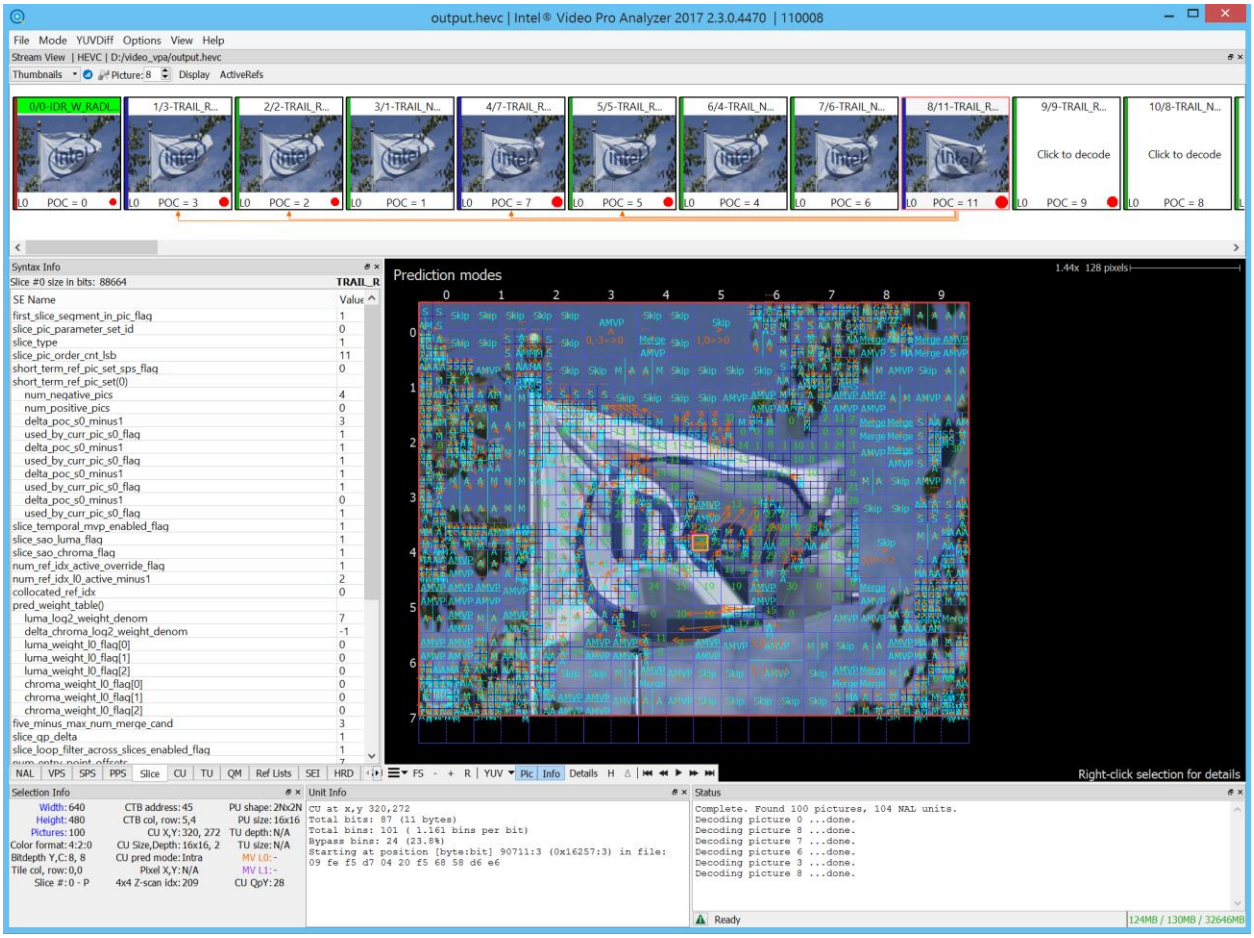

3.1 UI Components

The following sections describe the various components in the UI when an HEVC bitstream is loaded.

3.1.1 Top Filmstrip

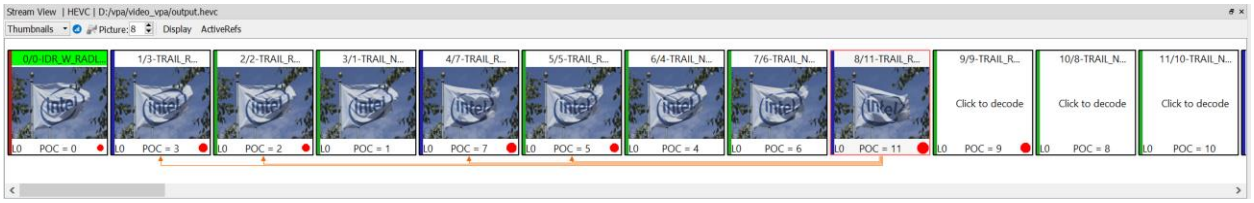

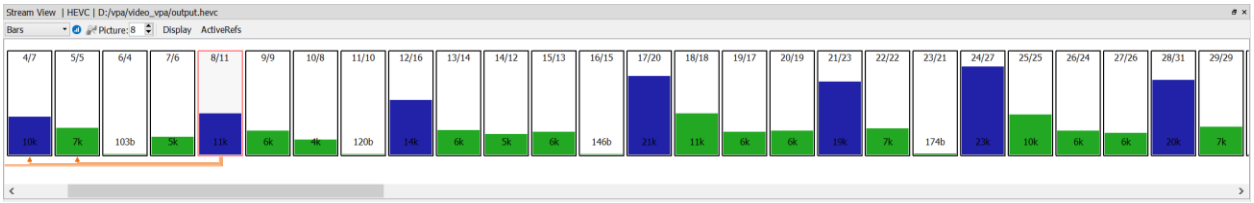

The top filmstrip is a horizontally scrolling overview of the pictures in the bitstream. When a new sequence begins (CVS from the HEVC spec), the top info label is highlighted in green. The current picture is highlighted with a red border. On top is the picture number - the decode order index in the bitstream - and the first slice in the picture type as indicated by the NAL unit types, and on bottom there are layerId shown as Lx where is x gives layerId itself and the POC value. Referenced pictures for the current picture are indicated with arrows. Green arrows indicate the long term references. The bar filmstrip has underline green band under those frames that begin new CVS. Picture type is core coded: red – intra picture (I), blue – P-picture, green – B-picture. Red dots show how often picture is used as reference picture. Size of the dot is proportional to such usage index. By using this information you have a hint witch image is mostly important for prediction other images.

Title bar shows currently loaded stream and selected codec for it.
Clicking on a picture will cause Intel Video Pro Analyzer to decode it and mark it as the current picture, updating the rest of the UI correspondingly. If the new picture requires other pictures to be available as reference that have not yet been decoded, those will automatically be decoded.

Pictures are decoded on demand since it is impractical to store all details of all pictures in memory, especially for HD sequences. The current picture can also be selected by typing in a picture number in the box on the left of the panel. When using this box, the filmstrip is scrolled such that the newly selected current picture is visible.

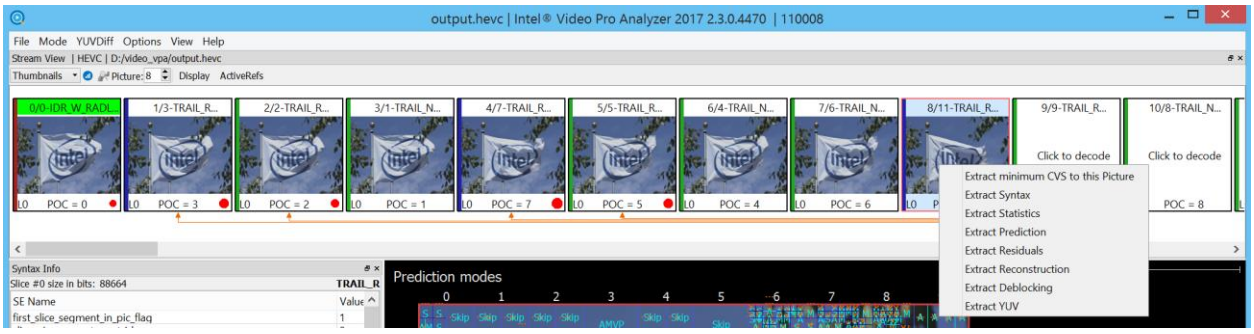

*Other names and brands may be claimed as the property of others.

Application title bar shows currently loaded stream, selected codec for it and process ID for easier identification of dock widgets. When any dock widget is detached from the main window his title bar will show process ID of the main window also.

In Thumbnails and Bar mode it is possible to choose between Display/Decode order representation by click on additional. By right clicking on selected picture one can get popup submenu with a set of extract actions. An action called "Extract minimum CVS to this Picture" will write out a bitstream that contains the minimum number of pictures needed to decode the current picture. Typically this will consist of the current picture and all previous pictures that are reference pictures up to the nearest IRAP picture. This can be useful for debugging issues in long sequences. Other actions allow to save in the output file statistics and intermediate/final YUV planes.

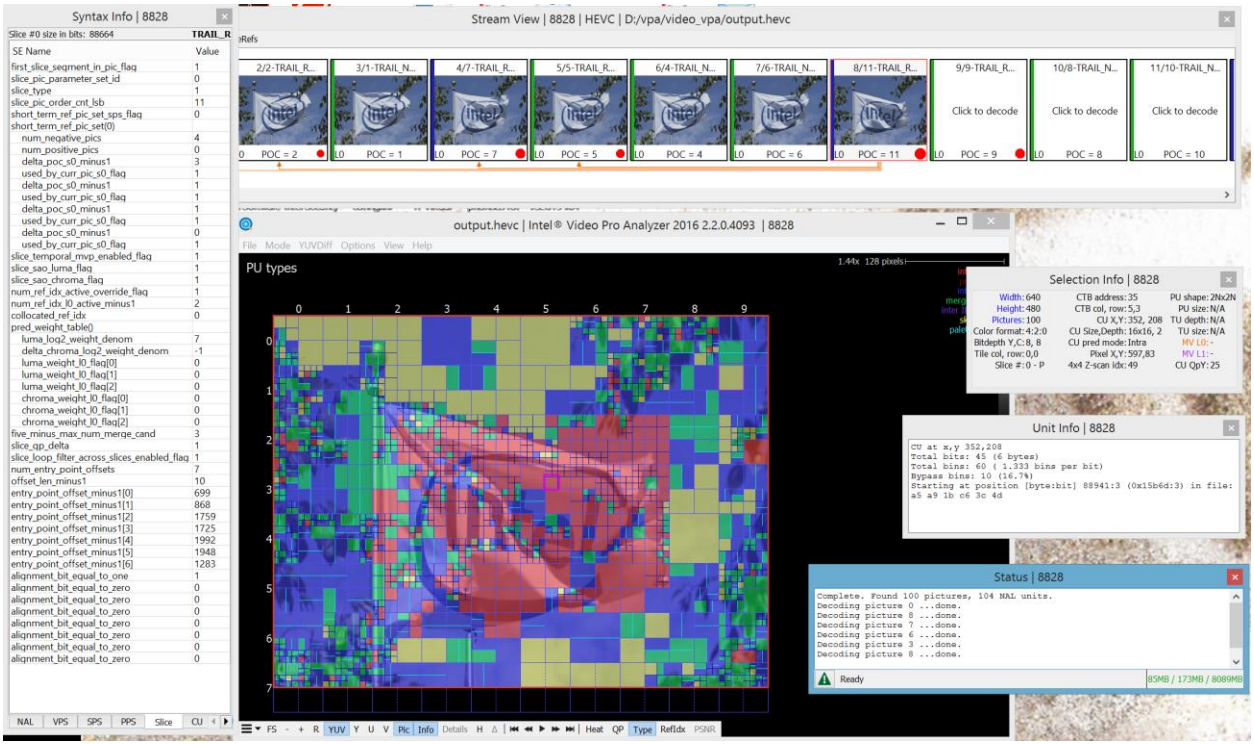

Allow separate panels can be detached from them main window or hidden. To attach panels back double click on the titlebar. To unhide panel go to View menu and check appropriate panel.

3.1.2 Main Panel

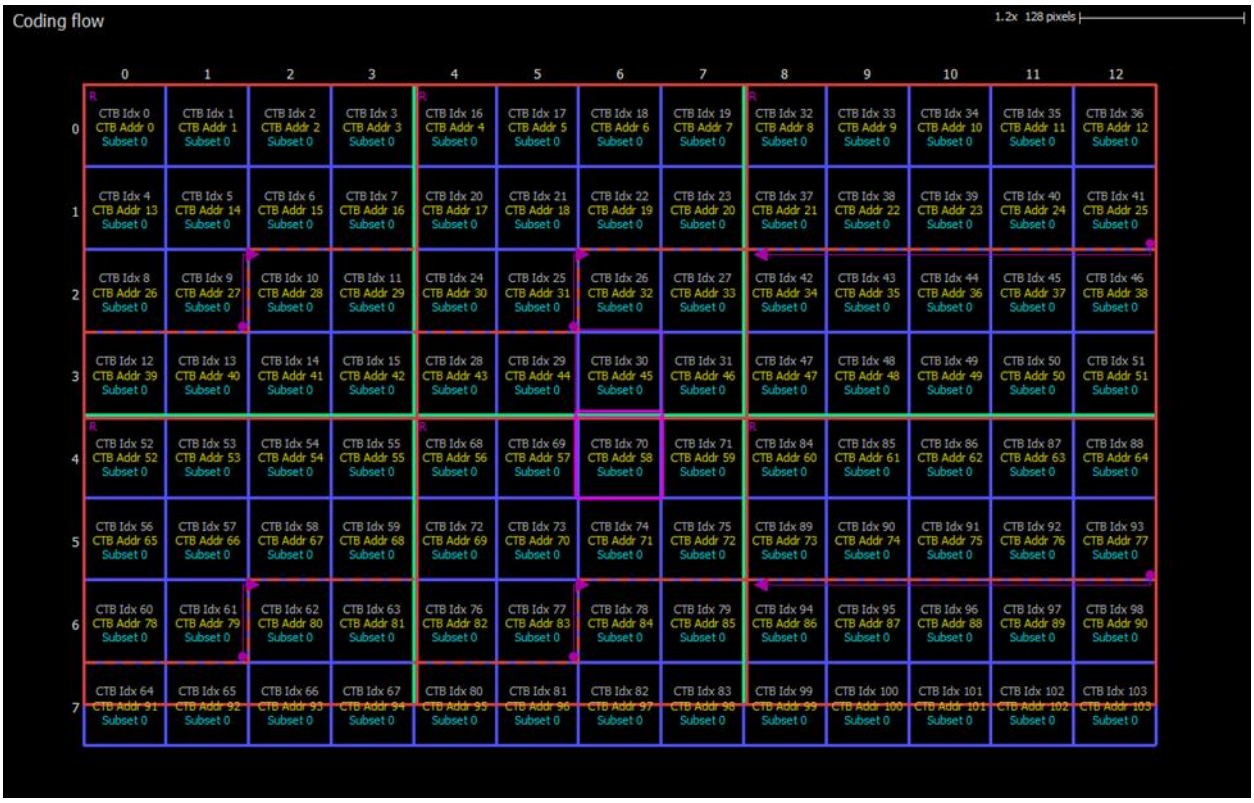

The main panel displays the selected picture with visual annotations. The type of annotations and associated interactive behavior depends on the current mode, see 3.1.8.1 HEVC Modes. Click-dragging moves the picture around, and the mouse wheel zooms in or out about the mouse cursor.

At the top left of the panel, the current mode is displayed. Along the top right the current scale is shown. In all modes, the CTB row and column values are displayed along the left and top border of the picture, slices are indicated with thick red lines, dependent slice boundaries are dashed red lines and tile boundaries are thick green lines. The screenshot above shows an example of a stream that is 13x8 CTBs, has 6 tiles and a number of slice segments, some of which are dependent. The selected CU is outlined with a pink box.

The information in the top right corner show current zoom factor and pixel length at selected zoom. By double clicking on zoom factor you can reset it to 1x factor.

The bottom left has a cluster of buttons:

- First submenu has special options to output the currently displayed information to the external image file with different scaling possibilities.
- FS: Toggles full-screen mode, which hides all UI elements except the main panel. Keyboard shortcut is F.
- +/-: Zooms in or out, centered on the center of the panel. Note that zooming with the mousewheel is much easier than using these buttons. Keyboard shortcuts are + and -.
- R: Resets the current zoom to fit the entire picture in the main panel. Keyboard shortcut is R.
- YUV, Y, U, V: Toggles between per-component image, or full-color YUV.

- Pic: Toggles the actual picture underneath the annotations. The picture is mode-dependent. For example in prediction mode, the picture is comprised of the prediction samples before the residual signal has been added in. It is often helpful to turn off the picture in order to make the annotation easier to read. Keyboard shortcut is P.
- Info: Toggles the annotations on or off. Keyboard shortcut is I.
- Details: Switches to the detail mode if it allowed of currently selected mode
- H/D: Toggles between hexadecimal and decimal display of the annotations. This applies to the values in the Left Tabs as well. Keyboard shortcuts are H and D.
- Delta sign: Switches to the difference mode with slave bitstream if it it loaded and difference mode is possible.
- There are five navigation buttons: first frame, previous frame, playback mode, next frame, last frame. The previous/last button will rewind, causing the current picture to be set to the first/last picture in the loaded bitstream. The "Play" button will start playback of loaded bitstream at a reasonable rate, limited by CPU performance. All normal mode operations are possible during playback. Note that playback mode always operates in currently selected order in the filmstrip.

3.1.3 Left Tabs

The left side of the UI is a resizable panel with several tabs that display information about the current selection in one way or another. The following subsections describe each tab.

3.1.3.1 NAL

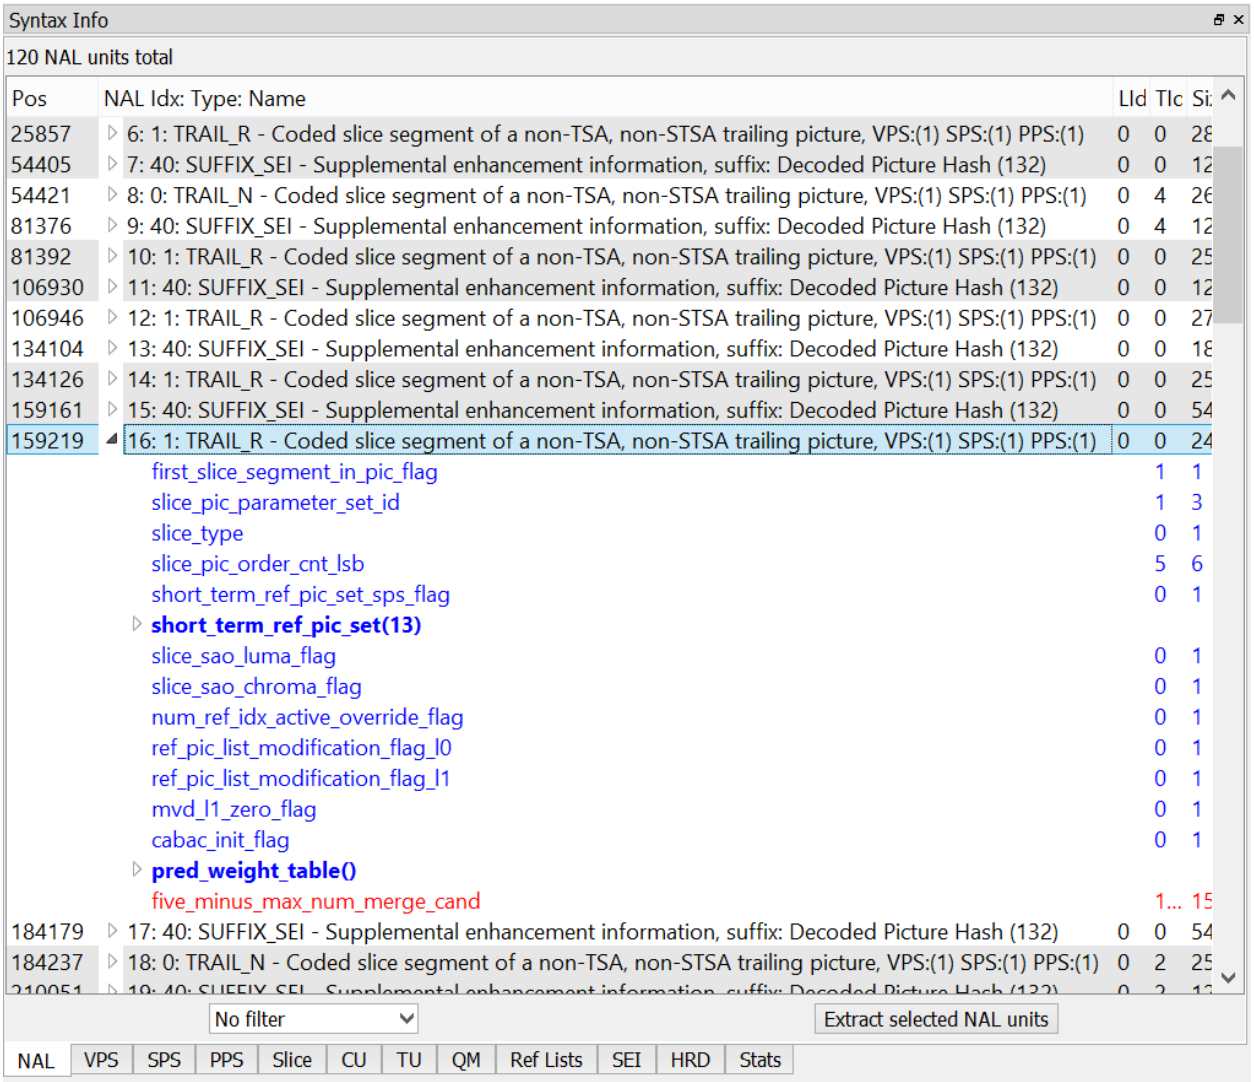

The NAL tab lists all NAL units found in the bitstream, in decode order. Above the list the total number of NAL units is displayed. The list has four columns: "Pos" is the byte position of the NAL unit in the file, "NAL Type" is the textual description of the NAL's type, "TID" is the Temporal ID of the NAL unit, and "Size" indicates the total number of bytes in the NAL unit. NAL description string contain header IDs used in decoding particular NAL unit.

Changes in coloring from white to grayscale and back show Access Unit borders. By clicking on branch sing you can roll out NAL unit syntax elements read from the bitstream. For each syntax element Intel Video Pro Analyzer shows the value of it and the number of bits it takes in the bitstream. Red colored elements will hint you about error syntax element.

Clicking on any NAL unit will display up to the first 500 raw bytes in the Raw Bytes panel. If a NAL unit is a VCL NAL unit, meaning it is a coded slice segment, it can be double-clicked in order to make the picture that it belongs to the current picture. Conversely, changing the

current picture will highlight the NAL unit that contains the first slice segment of that picture. Selecting a CU in the Main Panel will highlight the NAL unit containing the slice that contains the selected CU.

At the bottom of the NAL tab is a button called "Extract selected NAL units". Clicking this will write all selected NAL units out to a file.

Filter combo box allows you to select which NAL types will be displayed inside NAL tab. “No filter” means all NALs are shown.

3.1.3.2 VPS, SPS, PPS, Slice, SEI

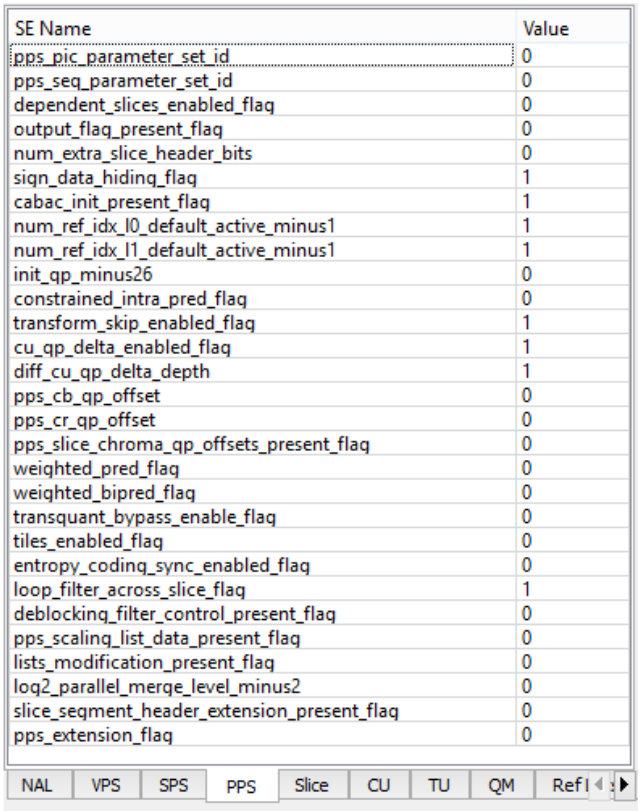

SE Name	Value
pps_pic_parameter_set_id	0
pps_seq_parameter_set_id	0
dependent_slices_enabled_flag	0
output_flag_present_flag	0
num_extra_slice_header_bits	0
sign_data_hiding_flag	1
cabac_init_present_flag	1
num_ref_idx_l0_default_active_minus1	1
num_ref_idx_l1_default_active_minus1	1
init_qp_minus26	0
constrained_intra_pred_flag	0
transform_skip_enabled_flag	1
cu_qp_delta_enabled_flag	1
diff_cu_qp_delta_depth	1
pps_cb_qp_offset	0
pps_cr_qp_offset	0
pps_slice_chroma_qp_offsets_present_flag	0
weighted_pred_flag	0
weighted_bipred_flag	0
transquant_bypass_enable_flag	0
tiles_enabled_flag	0
entropy_coding_sync_enabled_flag	0
loop_filter_across_slice_flag	1
deblocking_filter_control_present_flag	0
pps_scaling_list_data_present_flag	0
lists_modification_present_flag	0
log2_parallel_merge_level_minus2	0
slice_segment_header_extension_present_flag	0
pps_extension_flag	0

These tabs contain a list of every syntax element, in decode order, of the VPS, SPS, PPS, slice segment header and SEI messages that apply to the currently selected CU in the Main Panel. Syntax elements that come from a sub-function call are indented accordingly, and the function call itself appears in the list as a syntax element with no value. Any part of the list can be selected and copied for pasting in other programs.

3.1.3.3 CU, TU

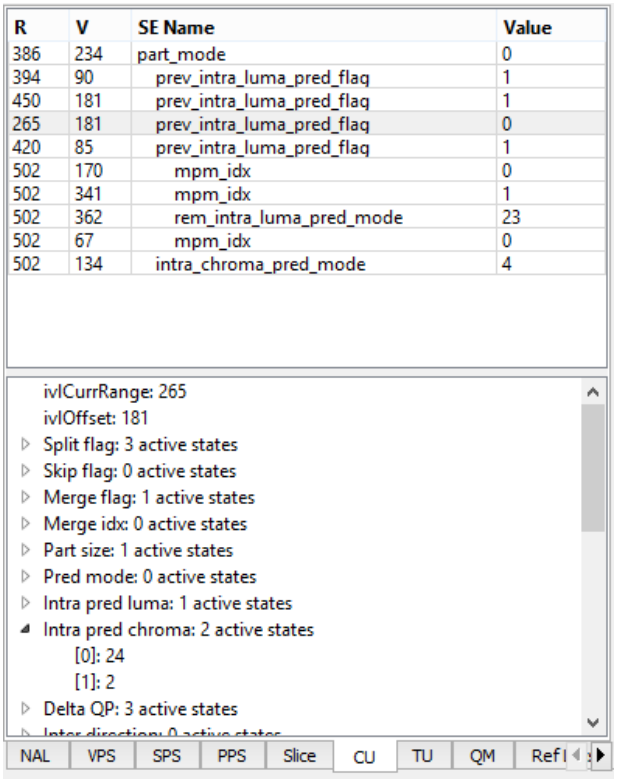

These tabs contain a list of the CABAC syntax elements decoded, in decode order, of the currently selected CU or TU in the Main Panel. Note that only the syntax elements pertaining to the currently selected CU/TU are shown, not any of the syntax elements in a higher or lower nodes in the quadtree. So if a CU is selected that is split into 4 smaller CUs, only the split flag will be shown, not the syntax elements of CUs further down the quadtree.
The left side of the list contains two columns, R and V, which denote the state of the CABAC engine's Range (ivlCurrRange) and Value (ivlOffset) prior to the syntax element decoding process. Any part of the list can be selected and copied for pasting in other programs.

Clicking a syntax element will update the tree in the lower half of this panel. This tree displays the values of all CABAC state variables prior to the decoding of the selected syntax element.

3.1.3.4 QM

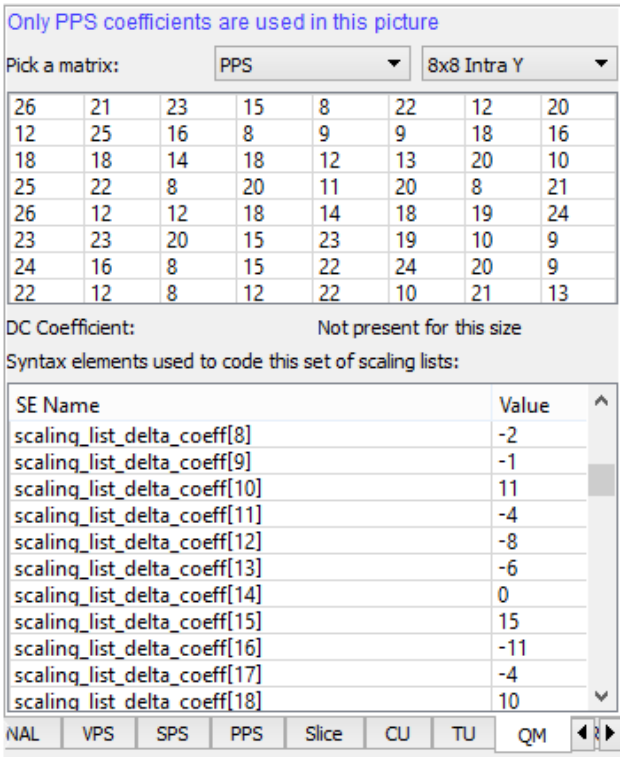

This tab displays the scaling lists, or quantizer matrices used by the current picture. Scaling lists can be present in the SPS, PPS, or both. PPS scaling lists take precedence over SPS when present. The top of the panel shows in blue text how the various scaling lists are used in this picture. Below that, one of the 20 scaling list from either the SPS or PPS can be chosen for inspection, and the chosen matrix is displayed in the grid. Matrices for 16x16 and larger TUs have a separate DC coefficient, which is displayed below the grid. Also shown in this panel is the list of syntax elements that is used to code the scaling list set, and is identical to the syntax in the corresponding SPS or PPS tab.

3.1.3.5 Ref Lists

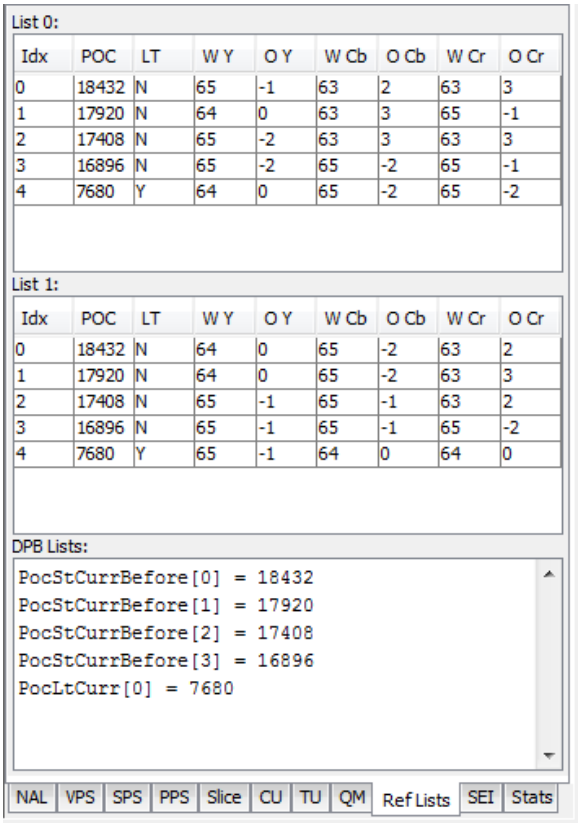

This tab displays the details of the two reference lists L0 and L1, as well as the Reference Picture Set arrays used to construct the L0 and L1 lists. The L0 and L1 lists have the following columns:

- Idx: The index associated with the reference picture
- POC: The Picture Order Count of the reference picture
- LT: A Yes/No flag indicating if the reference picture is a long-term picture or not
- W Y: The derived luma weight used in the weighted prediction process of the reference picture
- Y: The derived luma offset used in the weighted prediction process of the reference picture
- W Cb: The derived chroma Cb weight used in the weighted prediction process of the reference picture
- Cb: The derived chroma Cb offset used in the weighted prediction process of the reference picture
- W Cr: The derived chroma Cr weight used in the weighted prediction process of the reference picture
- Cr: The derived chroma Cr offset used in the weighted prediction process of the reference picture

The third panel contains in textual form the contents of the Reference Picture Set arrays from section 8.3.2 in the HEVC specification.

3.1.3.6 Stats

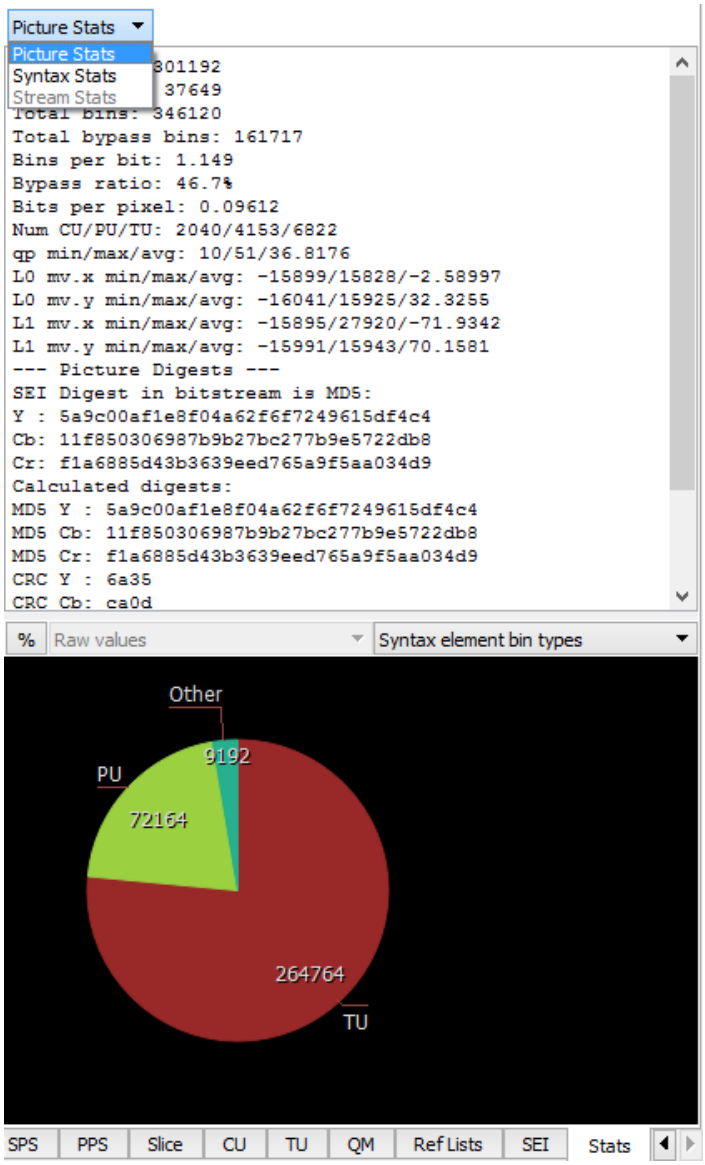

This tab displays various statistics extracted from the current picture. The top half shows some picture size and compression stats as well as the image digest information. There are also some min/max/average information about some picture parameters, such as QP value, MVs and so on, in this window.

The bottom half displays pie charts for a number of metrics. Each pie chart can be drawn normalized or un-normalized. Normalized data is weighted by area or compressed bits. For example there may be a much lower number of 64x64 CUs than 8x8 CUs in the picture (smaller un-normalized pie wedge), but they could still make up the majority of the picture area, making the 64x64 pie wedge large. Normalized numbers in the pie chart are in units of pixels or bits, and un-normalized numbers are raw counts.

The pie chart can be moved by dragging the mouse, and zoomed with the mouse scroll wheel.

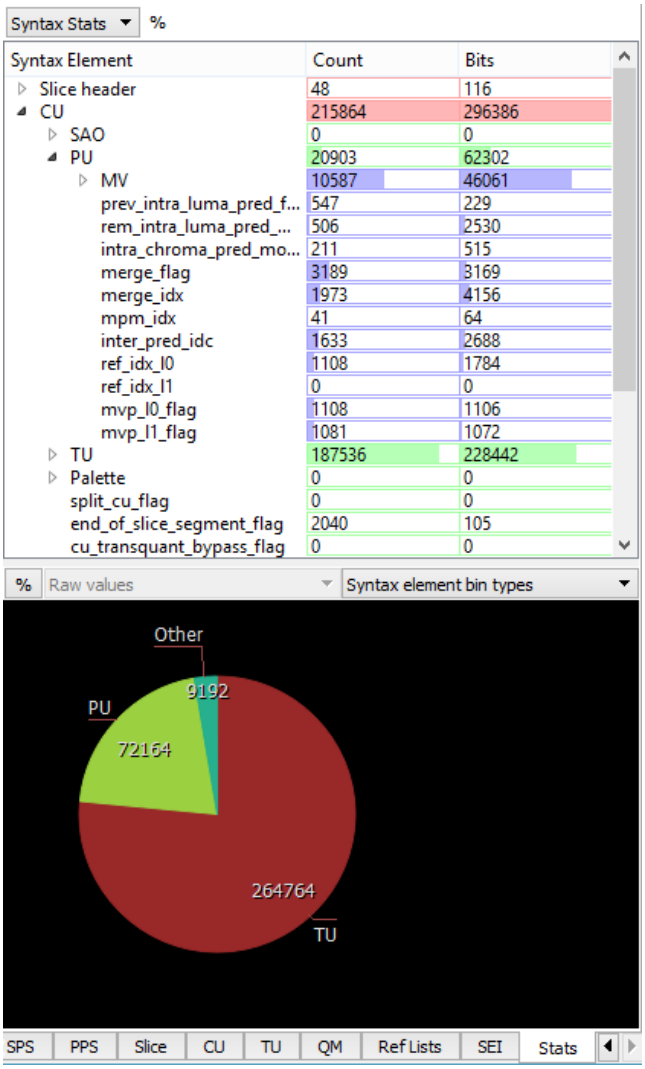

By selecting, different modes in the stats combo box you can explore different type of statistics in the bitstream. Namely, picture statistics, per syntax element statistics, overall stream statistics. Overall stream statistics is available only after full stream analysis and shows accumulated values overall the stream. Pie chart for stream statistics show accumulated value also.

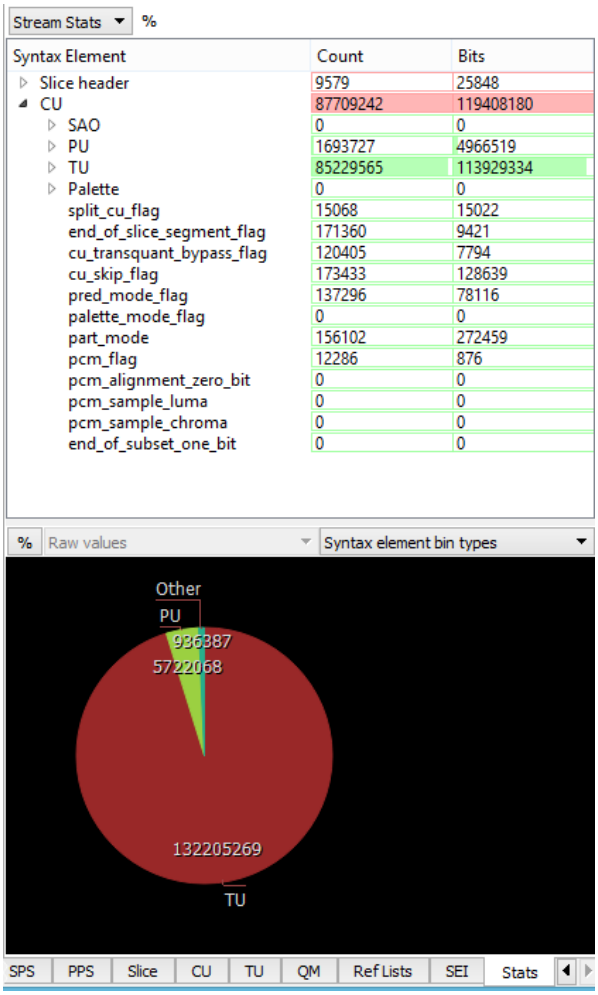

By pressing '%' buttons the values in the corresponding view changes to the percentage view.

3.1.4 Selection Info

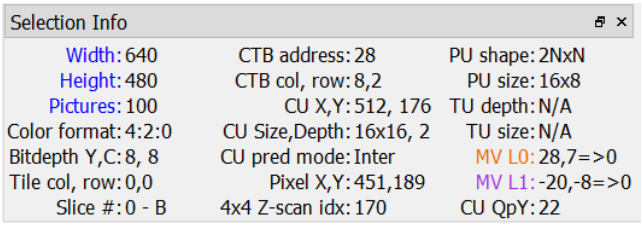

This panel shows a few details about the current selection at a glance. All values are decimal:

- Picture height & width
- Number of pictures in the bitstream
- Color format of the current picture
- Bitdepth values for Luma anc Chroma components
- Tile row/col position
- Slice segment number of the slice segment that contains the selected CU, and the slice type.
- CTB address and row/col position
- Size of the selected CU. Note it applies to the selection at any point in the quadtree, not necessarily just the leaf nodes.
- X/Y position of the selected CU in pixels
- CU prediction mode. Intra, Inter or Skip
- Depth in the quadtree of the selected CU
- Pixel X/Y position in the picture of the mouse cursor
- Z-order index at 4x4 granularity in the CTB at the mouse cursor location
- PU size and shape. Note that this will be equal to the size of the leaf CU in any mode other than Prediction.
- TU size and depth. Note that this will be equal to the size of the leaf CU (top of TU tree) in any mode other than Residuals.
- The L0 and L1 motion vectors, if any, of the selected PU. The MVs are shown with an arrow => followed by the reference index.
- The Luma QP value associated with the selected CU.

3.1.5 Raw Bytes

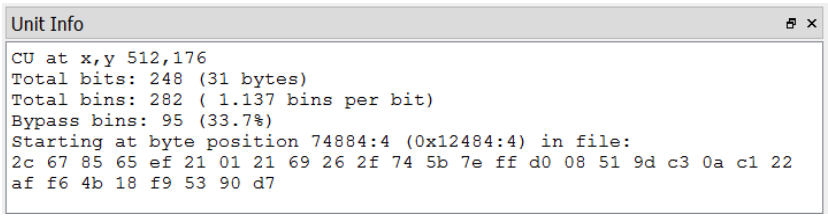

The panel in the bottom center of the UI displays the raw bytes that were used to code the selected CU, or NAL unit. A few details about the selection are included as well.

3.1.6 NAL unit detail view

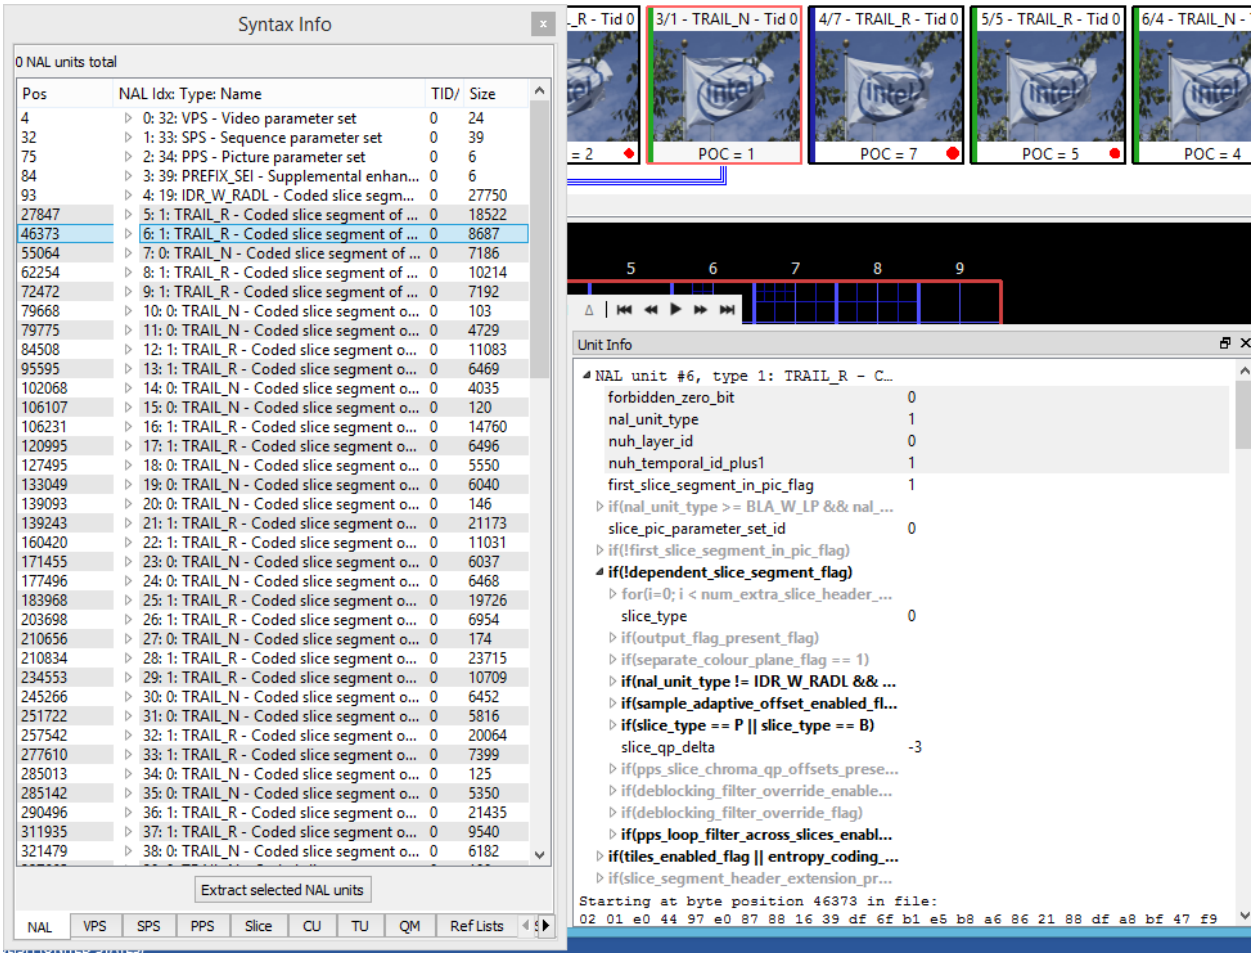

When one selects NAL unit in the NAL tab with single click Intel Video Pro Analyzer will display up to the first 500 raw bytes in the Raw Bytes panel and the controls for expanding NAL unit tree will appear from the left of "NAL unit" keyword. Control representation is system depended.

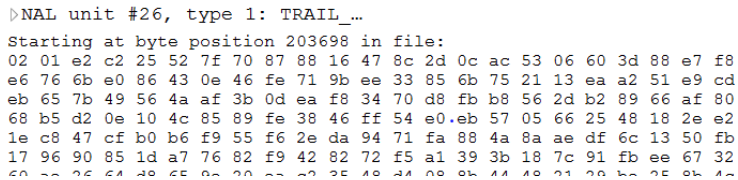

By clicking on this control one can expand or collapse NAL unit tree items. The structure of this tree mostly replicate NAL unit definitions from the standard. There are items with statements and items with conditions. The items with conditions are grey out if the condition is false. In this case only underneath structure definition from standard is shown. Otherwise (in true case), real syntax values can be explored.

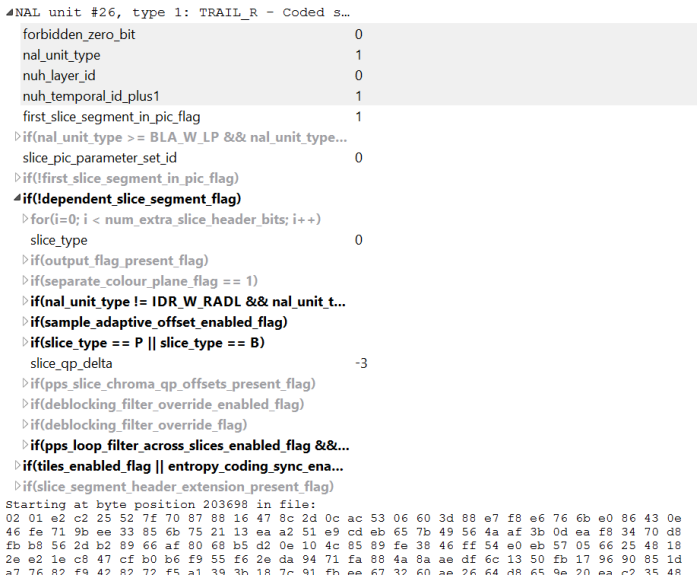

3.1.7 Messages

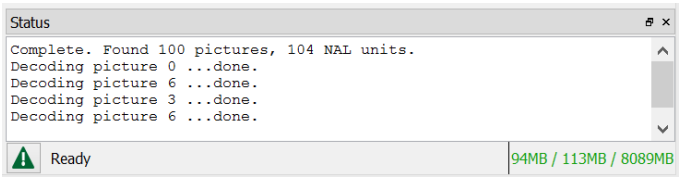

This panel in the lower right of the UI displays messages about the decoding process, and shows progress of any actions that may take a while to complete. If a decoded picture's digest does not match the SEI digest (when present), a warning will be displayed here.

The 3 numbers in the lower right of this panel represent the memory state of JVM. From left to right the 3 numbers are:

- Amount of memory used
- Peak used memory
- Maximum available memory

Error button at the bottom left of the window changes color to red if Intel Video Pro Analyzer detects any errors/warnings in the stream. By clicking on the button Intel Video Pro Analyzer switches to the list of detected problems and back. Double clicking on selected problem will navigate you the appropriate place in the header where the error occurs.

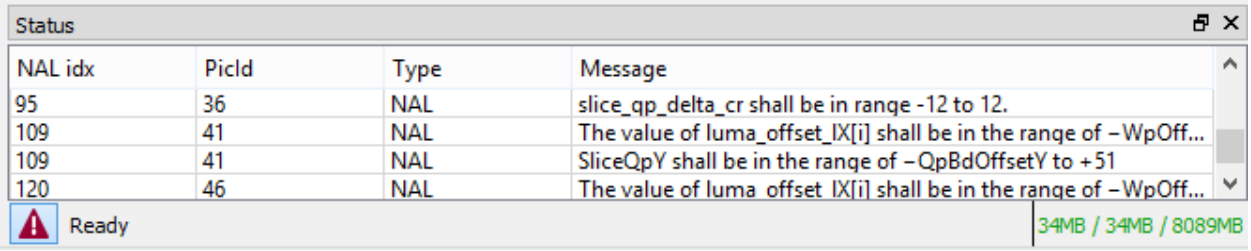

By right clicking inside error's list you can choose additional filtering of the errors

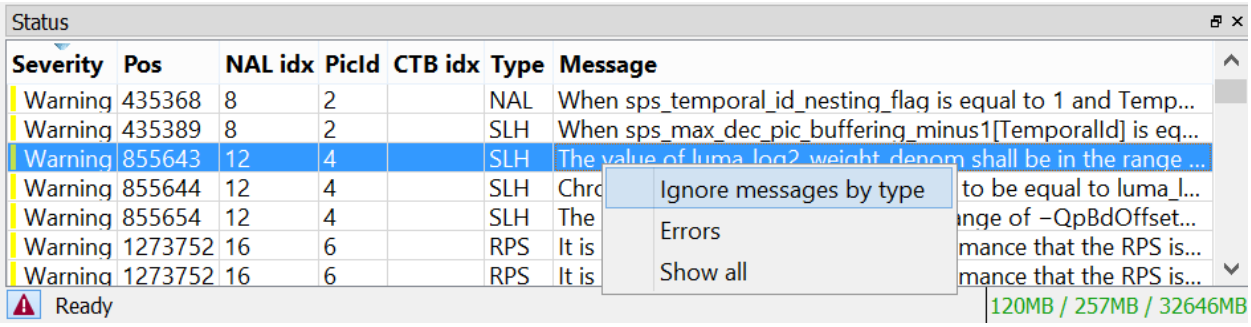

"Error" submenu runs common error dialog where it is possible to select which errors will by displayed in the error's list.

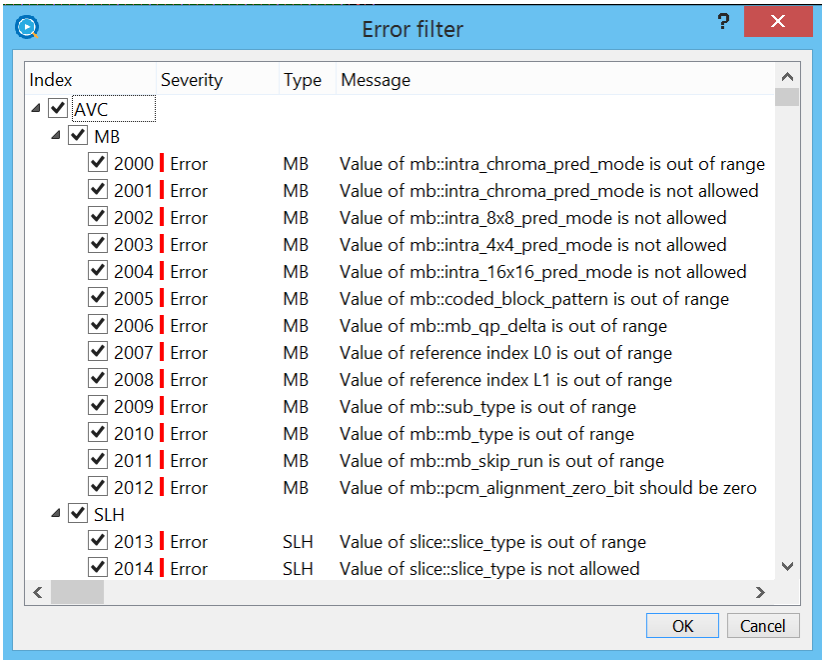

3.1.8 Extended modes

There are a set of extended modes that can be accessed via menus in the Stream View toolbar.

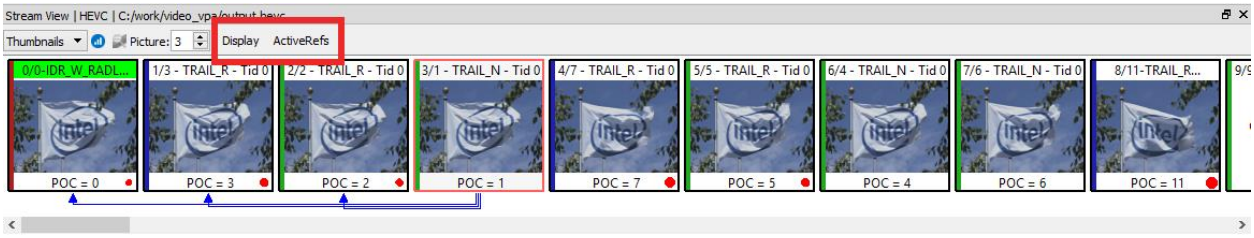

Then three possible modes with two sub-modes can be selected. Sub-modes of the currently selected mode changed dynamically depending on current mode.

3.1.8.1 Active/DPB refs

Sometimes, it may be that Intel Video Pro Analyzer shows reference on INTRA frame. It is not mistake because Intel Video Pro Analyzer shows DPB reference by default - all frames stored in DPB and marked as used for reference. But some of that frames could be inactive and as the result not used for motion estimation. For INTRA frame all such reference are inactive. To show only active references choose "ActiveRefs" sub-mode from popup toolbar.

3.1.8.2 Frame sizes view

When one choose "Buffer" menu from popup toolbar, frames size view will appear by default. This is a bar char plot with frame sizes. Each bar is colored according to the frame type. When you move the mouse along the bar current frame is highlighted and tooltip shows information about current frame. By clicking on selected frame you can decode it. When you press left mouse button and move it all plots move along both axes. Double click on canvas area resets plot into default state. You can zoom in/out with mouse wheel on both axis when the mouse inside canvas. If the mouse on some axis (below canvas, for example, x axis) you can zoom, pan or double click on only on that axis. In that case all actions applied to that axis only.
The movement of the plots is restricted to the positive values only, you can't move or zoom in negative direction. Currently decoded frame is shown with vertical line and frame number on it.

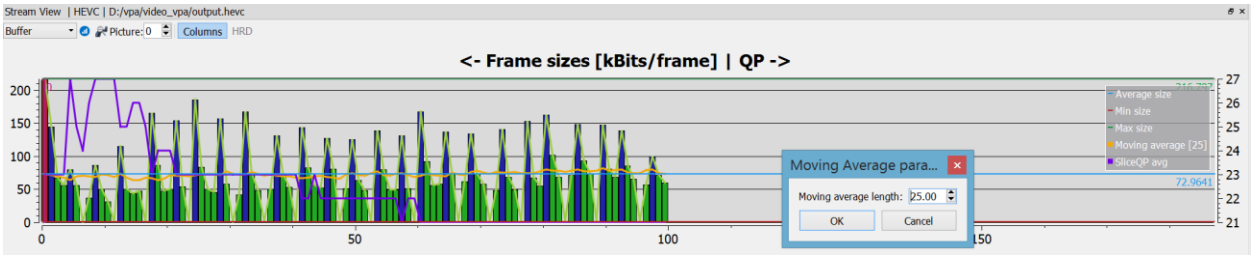

On the right axis average slice QP values are shown. When full stream analysis is complete Intel Video Pro Analyzer adds average block QP plot that is also attached to right axis. By double clicking on 'Moving average' label item sliding window size can be adjusted in the separate dialog.

3.1.8.3 HRD buffer fullness view

If you choose "Buffer" mode then sub-mode "HRD" appears in sub-modes block. It is enabled if stream has HRD parameters in it, otherwise it is disabled (grey out). The interaction with buffer fullness plot works in the same way as for frame sizes plot. It is possible to get buffer state and picture parameters in each point of the plot where the monotonic behavior changes.

With right click on central area of the plot Intel Video Pro Analyzer shows additional option for plot adjustment. In the same popup menu you can export HRD timings into .csv file.

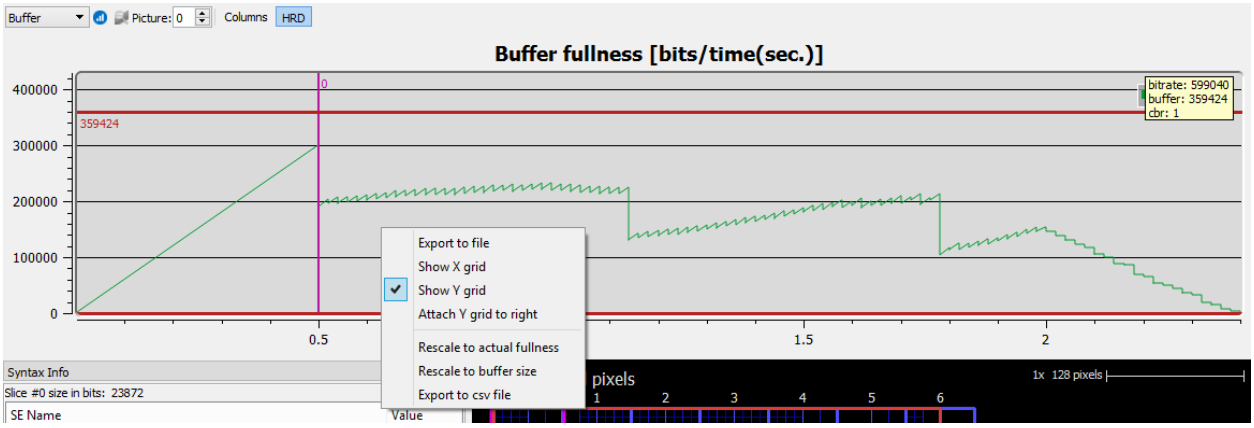

3.1.8.4 B-pyramid visualization

Intel Video Pro Anayzer show B-pyramid visualization when option Hierarchy in the Stream View is selected. It show the level of mutual B-frame references, frame numbers in decoding /display orders, frame type coded by color and frame references. By clicking on each frame you can highlight particular frame references.

3.2 Modes

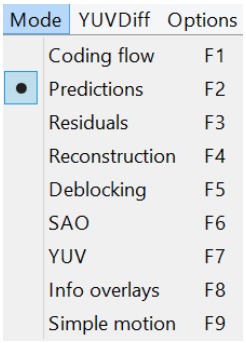

With an HEVC bitstream loaded, Intel Video Pro Analyzer can be put into one of 9 modes using either the F1-F9 keys, or using the Mode menu. The mode selection affects only what is displayed in the Main Panel. Most modes can also show details of the current selection (CU/PU/TU). This can be toggled with the right mouse button or the main panel's button strip.

3.2.1 Coding Flow

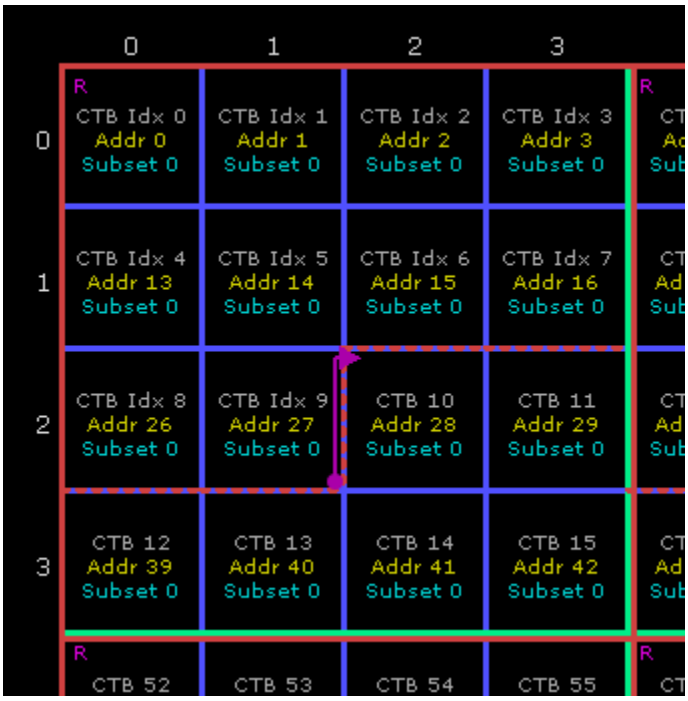

The coding flow mode gives a visual overview of the ordering of CTBs in the stream, and some information of the decoding process. The blue grid shows the boundaries of the CTBs in the picture. Each CTB contains 3 values:

- The decode index of the CTB, showing the order of decode in the picture.
- The CTB address, which is simply the raster scan index of the CTB.
- The substream that the CTB belongs to. This number will only be greater than 0 if tiles or wavefront tools are employed in the picture.

Operations on the CABAC engine state are displayed as well. A small purple R at the top left of a CTB indicates the CABAC engine is reset before decoding that CTB. A small F indicates a CABAC flush at the end of a substream that is not the end of a slice. Purple arrows indicate CABAC engine state transfer or copy to dependent slices or wavefront rows:

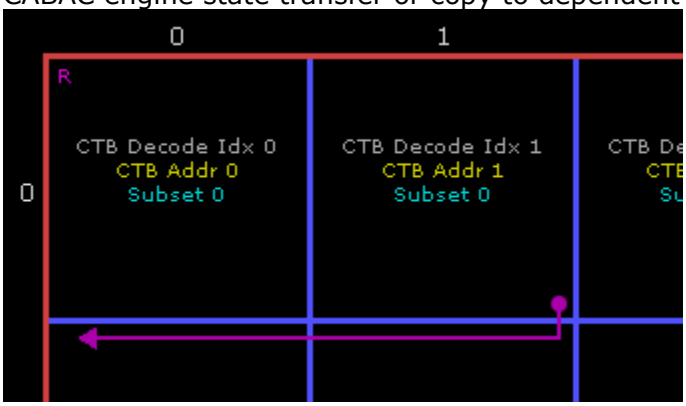

3.2.2 Predictions

This mode shows the PUs in each CU. CU quad tree splits are indicated with blue lines, which get darker as the split depth increases. When a CU split is implied by HEVC and not directly coded in the bitstream, the split lines are dashed.

Intra modes are indicated with green colors, and directional modes also show an arrow indicating the prediction direction. In the lower right corner of an intra PU the chroma mode is indicated in a darker green. When the prediction shape is not the same as the CU size, the prediction unit shapes are shown with green lines for intra blocks, cyan for inter blocks. Dashed green lines in intra blocks means the PU was split along with the transform tree.

Inter PUs are indicated with cyan colors showing the PU splits and mode. Additionally, skipped blocks are shaded with a lined texture. An inter PU is Skip, AMVP or Merge mode. The L0 motion vectors are drawn with an orange color, the L1 motion vectors (B slices only) are drawn with a purple color. The MV value is shown in the lower left corner along with the reference index.

Clicking a PU will select it, and the syntax elements used to code it and the CU it belongs to are displayed in the CU tab on the left. Clicking repeatedly on the same PU will cycle through the CU hierarchy, showing the parent and child relationship.

Below is shown Prediction mode on a zoomed-in selection along the top edge of a picture. The selected CU is surrounded with a pink box while the selected PU is surrounded with a yellow box.

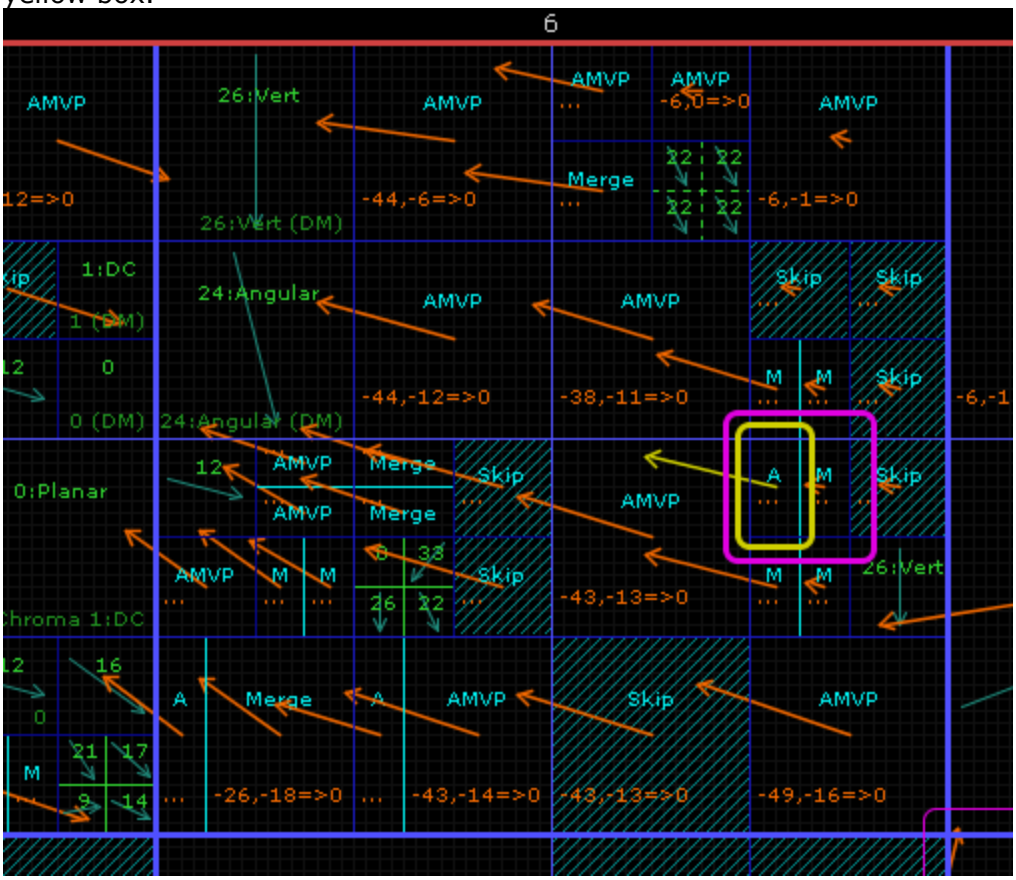

3.2.2.1 Prediction Detail Mode

In Prediction mode, the sample values of a particular PU can be viewed in detail.

Intra PUs are displayed in green, and the left, upper-left and above prediction arrays are drawn adjacent. Luma blocks show 3 versions of the predictions arrays. From furthest out to inner, they represent the arrays in 3 steps:

1. Initial neighboring samples. Unavailable samples are given with a red "X".
2. Samples after reference sample substitution process (section 8.4.4.2.2 from the HEVC spec).
3. Samples after filtering process (section 8.4.4.2.3 from the HEVC spec).

Chroma blocks only show the first two steps, since reference sample filtering is not applied to chroma intra prediction.

Inter PUs are displayed in cyan (turquoise). Each prediction sample can be clicked to display a window that contains the 8x8 or 4x4 sample array for luma or chroma respectively that are the inputs to the motion compensation filter for the selected pixel. In the case of a bi-directional PU, both arrays are shown the L0/L1 colors of orange/purple. Clicking the selected sample again or outside the PU closes this window. Underneath the prediction array values the motion vector predictor lists are displayed. In merge mode, the merge list is shown with the L0 predictor in orange and the L1 predictor in purple. An arrow indicates the merge candidate that is chosen by the bitstream (merge_idx).

When the inter PU does not code motion vectors in merge mode, it is considered AMVP mode (Advance Motion Vector Prediction). The 2-entry list for each motion vector (L0 and/or L1) is shown and the selected predictor is indicated with an arrow.

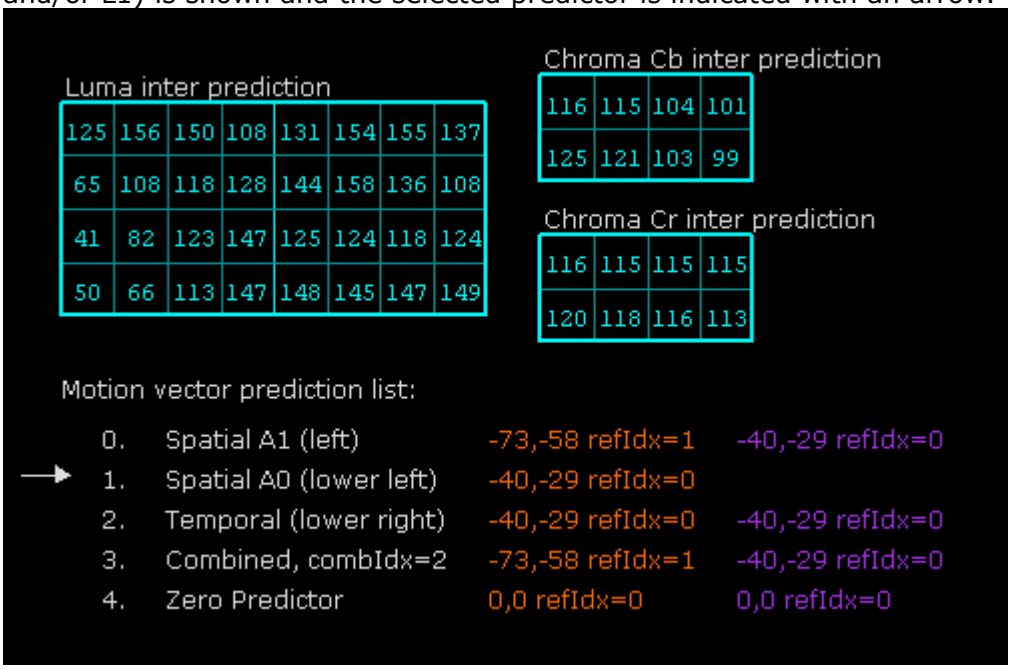

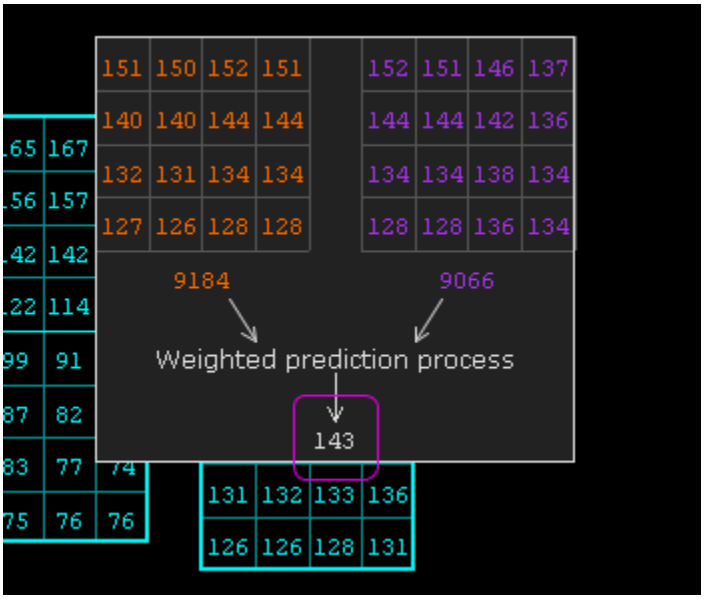

For HEVC screen content coding stream there is special detail mode exist – Palette mode. In this mode Intel Video Pro Analyzer shows palette predictor, currently used palette, coded indices, resulting indices and final pixel values.

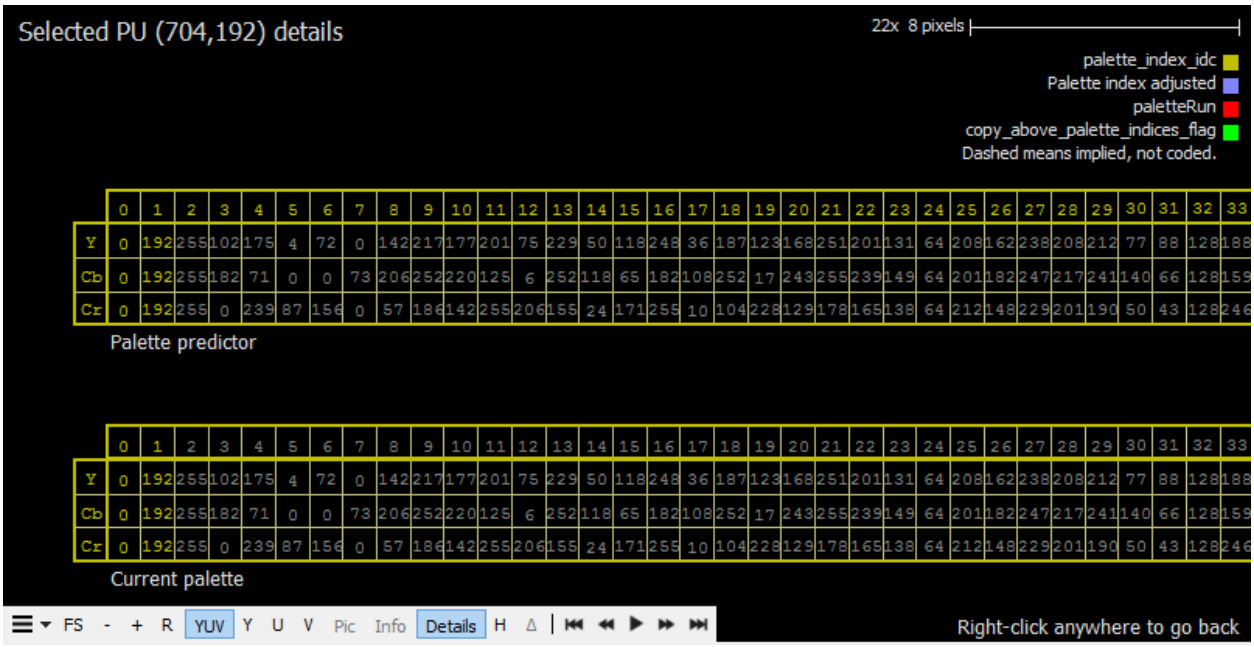

3.2.3 Residuals

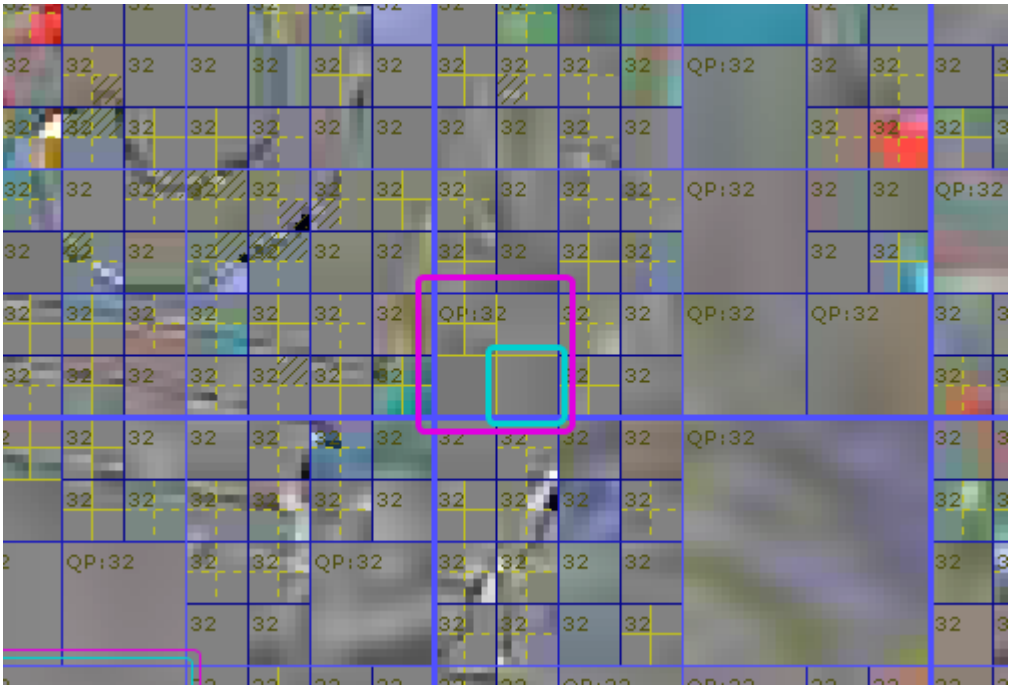

In Residual mode the transform trees and accompanying residual signal of the picture can be seen. As in Predictions mode, CU boundaries are shown with blue lines. The transform

splits, when present, are indicated with yellow lines. A dashed yellow line indicates that the TU split was implied. QP values of each CU are shown in the upper-left corner. When this value is bright yellow, it means that a delta-QP was coded in the bitstream during the decode of that particular CU. Otherwise the number is dark. When DQP is enabled by the PPS, this helps visualize the DQP coding depth (diff_cu_qp_delta_depth). TUs that have transform_skip_flag set are shaded. CUs that are PCM coded are marked with "PCM" in white, and CUs that have cu_transquant_bypass_flag set have a white X through them.

If the "Pic" button is turned on, the raw residual signal is shown in image form. Residual values of 0 result in flat grey, negative values are darker and positive values are brighter.

Clicking on TUs causes that TU to be selected with a blue surrounding box and its syntax elements (when present) to be displayed in the TU tab of the left panel. Clicking a TU repeatedly causes the selection to move up and cycle though the TU, then CU quadtree hierarchy. Note that when the selected TU is a 4x4, the chroma coefficients are coded one level up in the TU hierarchy since HEVC doesn't specify a 2x2 chroma TU. So to view chroma TU syntax when the luma TU is 4x4, click the luma TU twice to move up to the 8x8 level where the chroma syntax is decoded.

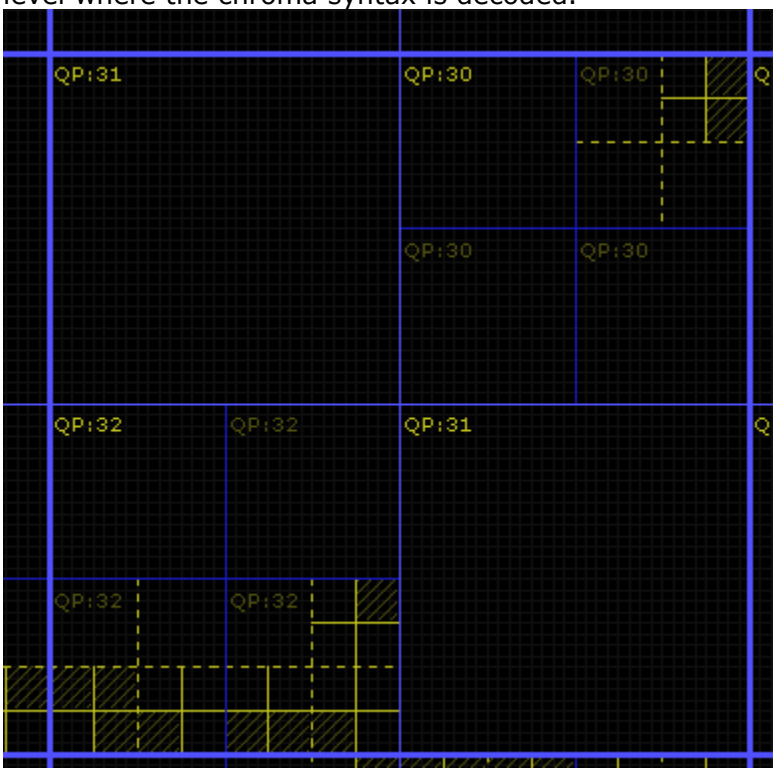

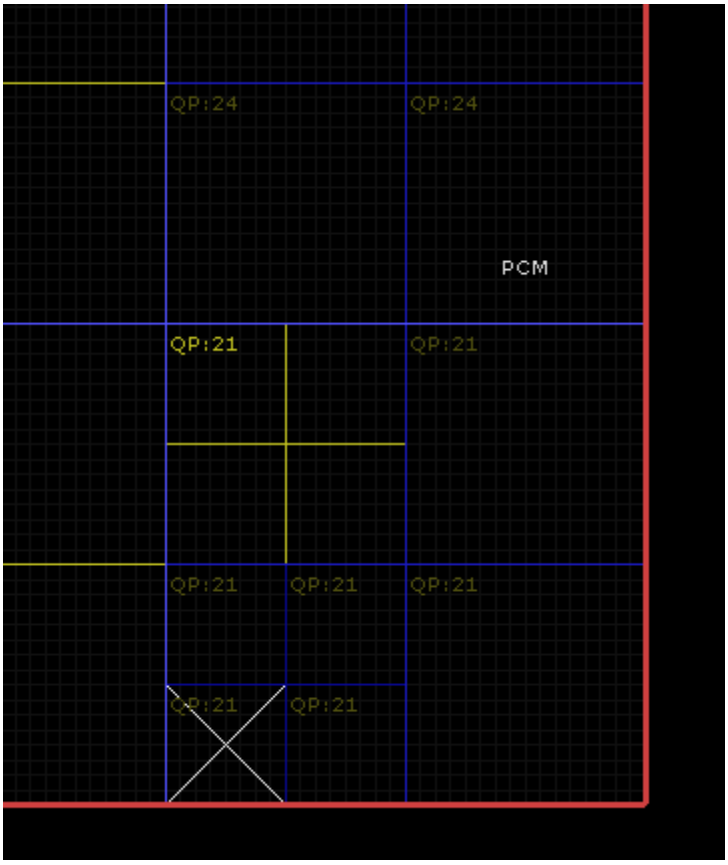

3.2.3.1 Residual Detail Mode

To view the full details of a particular TU or group of TUs, make the TU selection and enter detail mode by right clicking or using the detail mode button at the bottom of the main panel. In this mode the selected TU structure is drawn three times, arranged in a column from top to bottom, showing the 3 major steps in recovering the residual signal:

1. Coefficient decode. This diagram shows the scan order of the coefficients, the coefficient values and small icons indicating the syntax elements associated with coefficient decode. Refer to the legend shown in the top right of the main panel for the exact details. Coefficient groups are outlined with a green box. This box is bright if the coded_sub_block_flag was equal to 1, dark green if 0. A dashed line means the sub-group was implied to contain coefficients

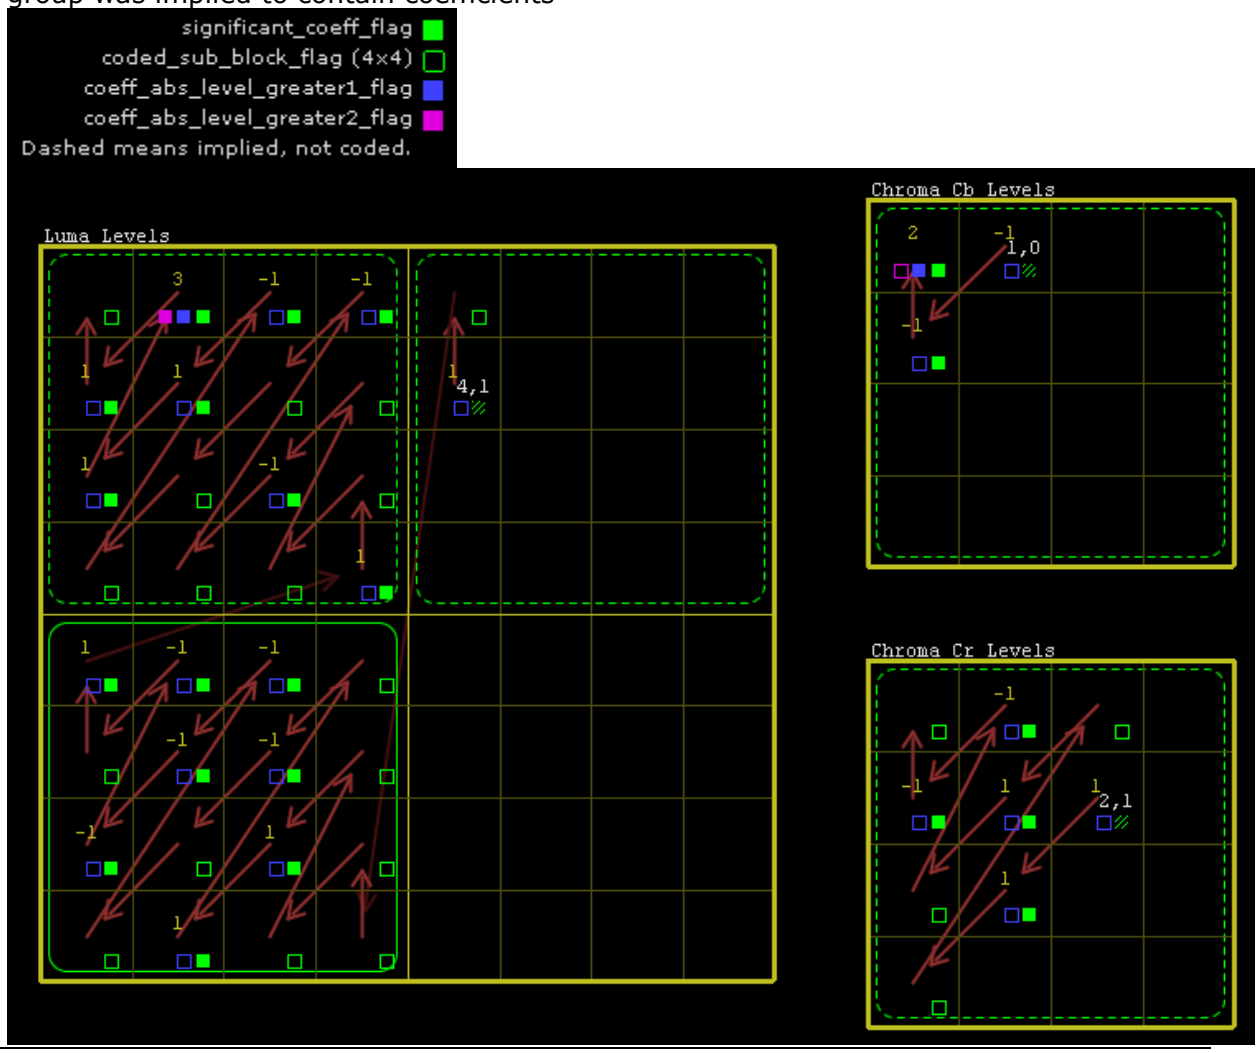

2. Inverse quantization. In this diagram the coefficient values are displayed after the inverse quantization process 8.6.3 from the HEVC spec has been applied.

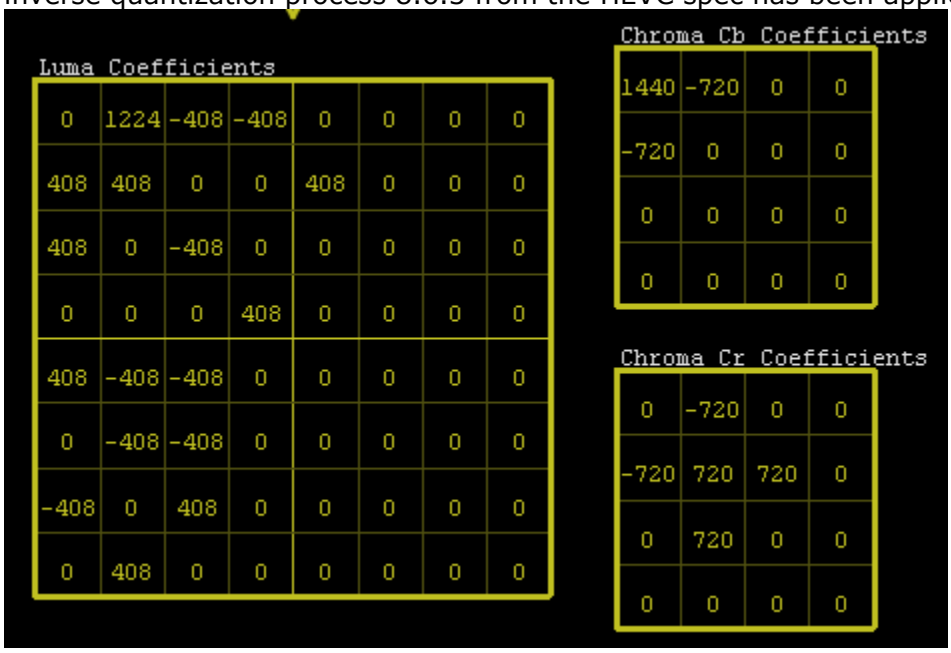

3. Inverse transformation. This diagram shows the recovered residual signal after the inverse transform process 8.6.4 from the HEVC spec.

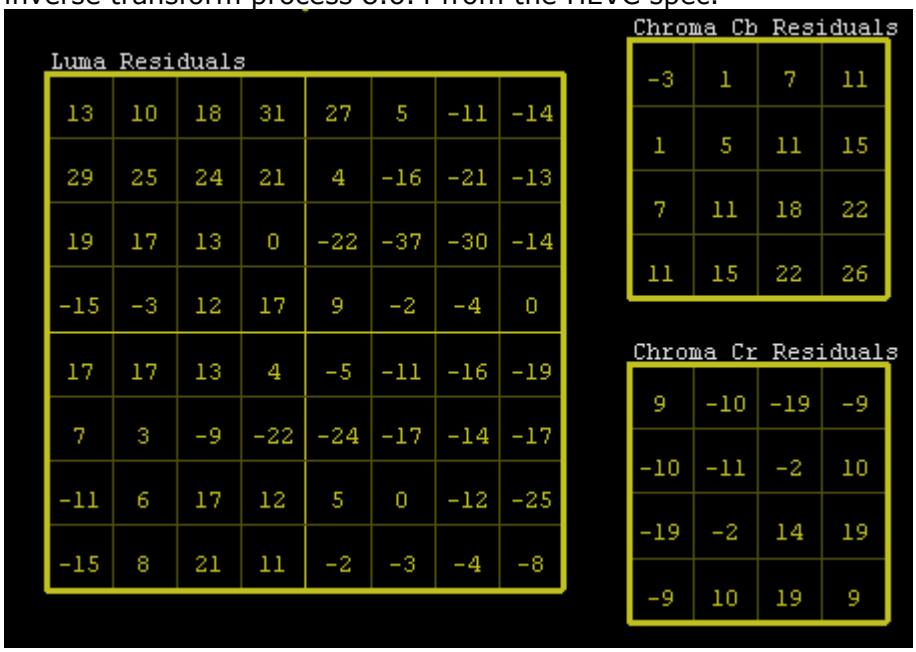

3.2.4 Reconstruction

In Reconstruction mode the reconstructed samples prior to deblocking can be inspected. As with Prediction/Residual mode, CU boundaries are shown with blue lines. CUs that have at least one TU with non-zero coefficients are marked with "Non-zero CBF". PCM-coded CUs are indicated with "PCM" in white.

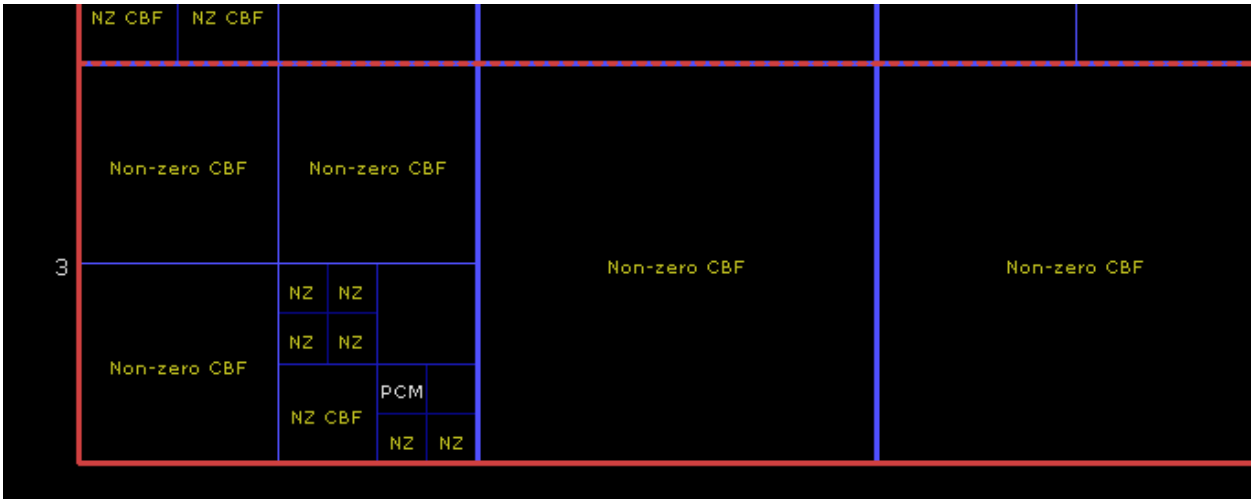

3.2.4.1 Reconstruction Detail Mode

In detail mode, the reconstructed sample values can be inspected. The selection may be a single CU, or a group of CUs up to and including the entire encompassing CTB.

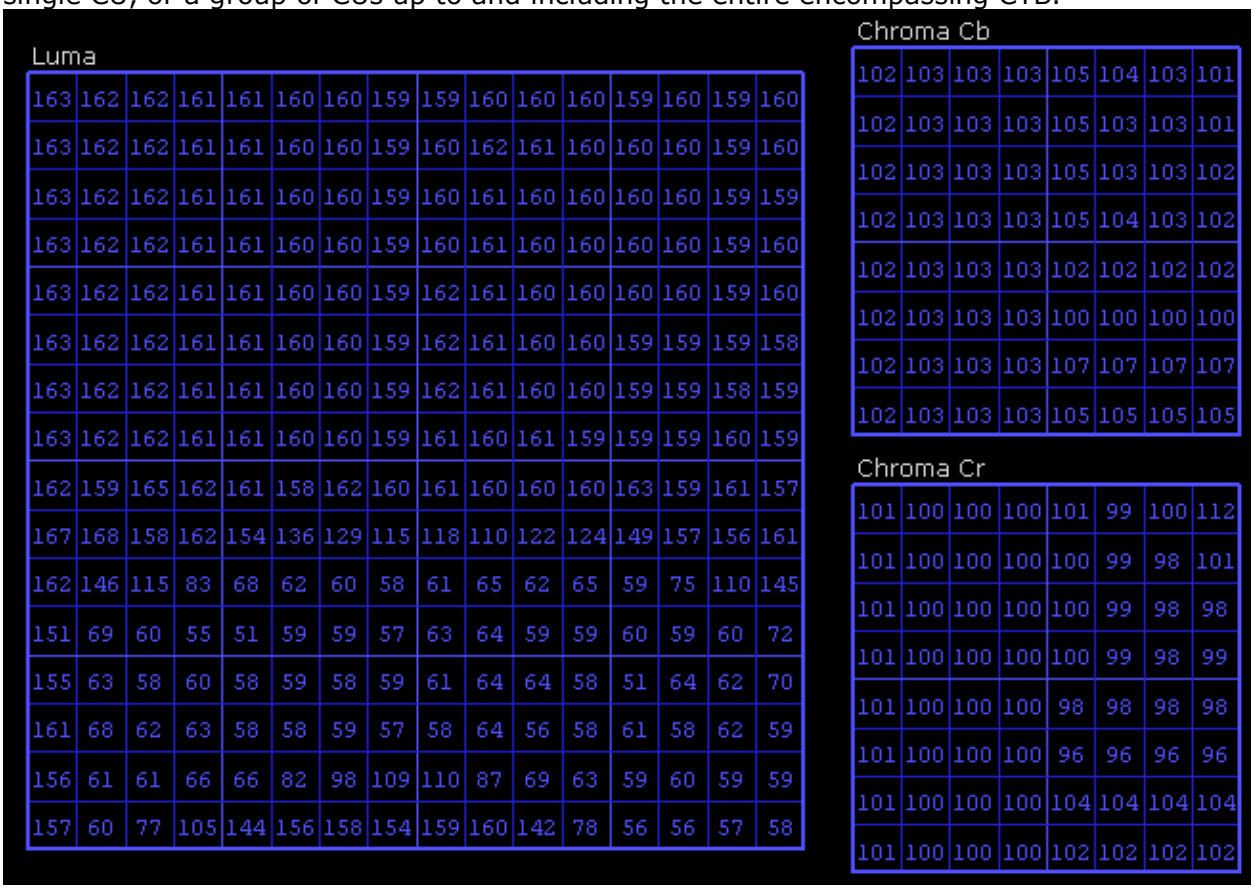

3.2.5 Deblocking

Deblocking mode shows all edges processed by the deblocking filter as described in the HEVC spec section 8.7.2. CU boundaries are not displayed in this mode, however CU selection still functions as in the other modes. Edges shown are for the luma deblocking process and are color coded in the following manner:

- Green: Strong luma filter applied
- Yellow: Weak luma filter applied
- Red: No filter applied (but edge was evaluated)

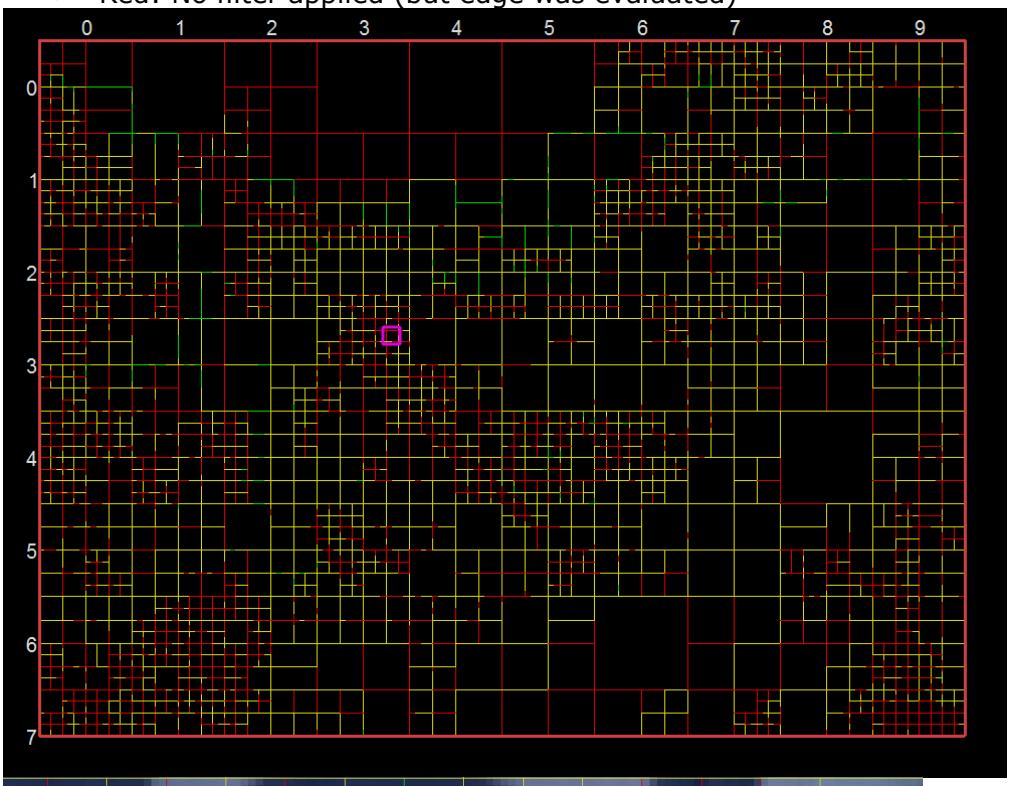

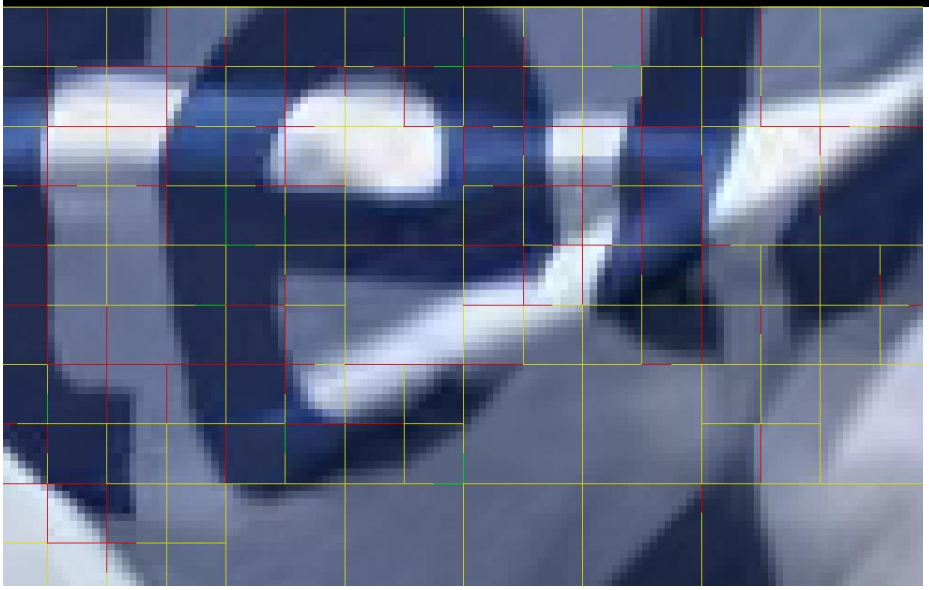

3.2.5.1 Deblocking Detail Mode

The deblocked sample values can be inspected directly in detail mode. As with Reconstruction Details, the inspected area can consist of one or more CUs. Sample values are indicated in yellow if they were modified by the deblocking process. Sample values that cannot be modified due to pcm_flag or cu_transquant_bypass_flag are indicated in red. Filtered edges belonging to the selected CU are shown in thick dashed lines, color coded the same way as in the non-detail mode. Each edge can be selected and clicked, bringing up a window showing the input values to the edge filter. Note that these values may differ from the reconstructed samples for horizontal edges, as they may have been modified by a vertical edge filter. Additionally, the boundary strength is displayed. Edges that were processed but ultimately not filtered (red colored) will indicate the reason. Typically it is due to the boundary strength being 0, the threshold condition not being met, or a slice/tile edge when those edges shouldn't be processed.

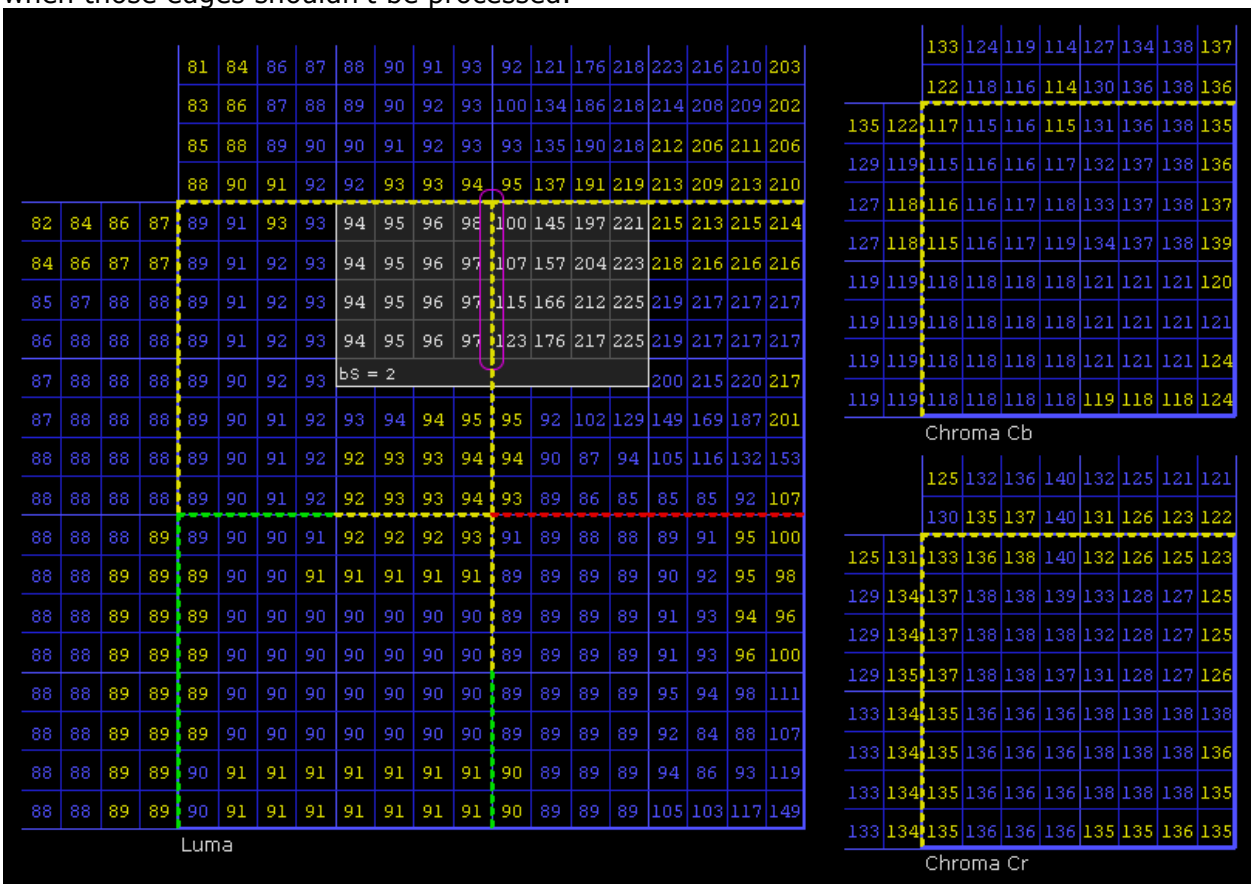

3.2.6 SAO

In SAO mode the SAO filter parameters assigned to each CTB are shown. The mode and associated four offsets for each component are indicated with yellow text, and the merge_up / merge_left flags are indicated with an arrow. The CU quadtree boundaries are shown here as well.

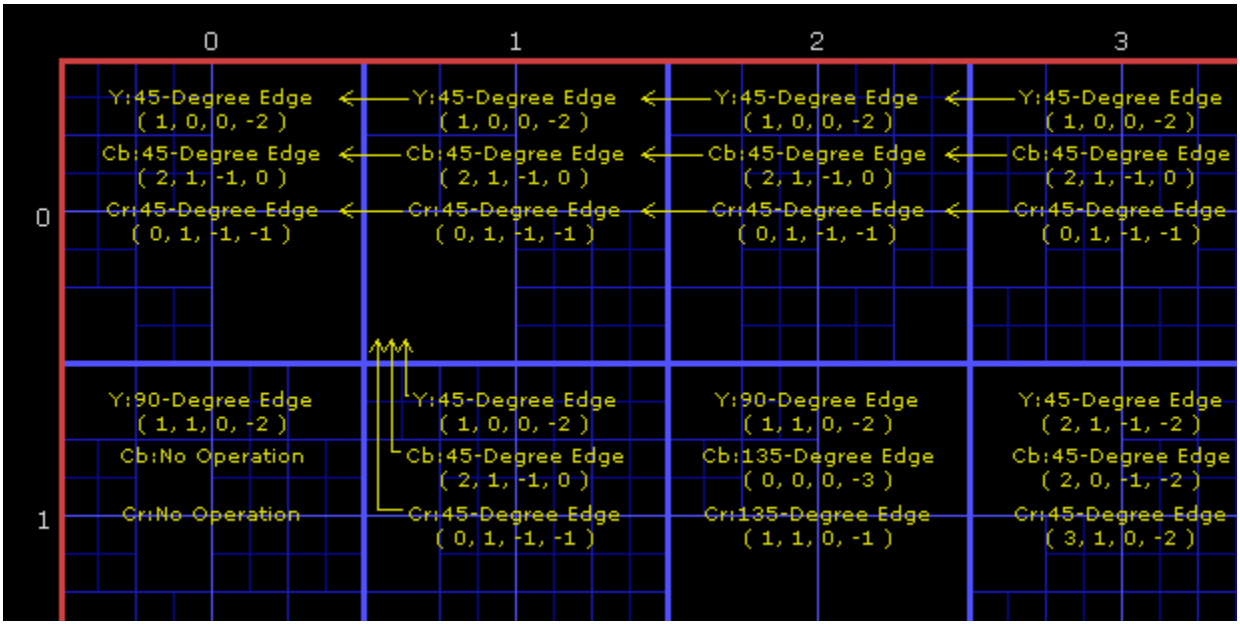

3.2.6.1.1 SAO Detail Mode

Entering detail mode on a selected CU will show all samples after being processed by the SAO filter. Samples values that were actually modified by the SAO operation are highlighted in yellow. Each sample can be clicked, bringing up a 3x3 window that shows the input values to the filter process. In the screenshot below, the 45-degree edge filter inputs are the samples to the lower left and upper right. Samples that cannot be modified by the filter are indicated in red.

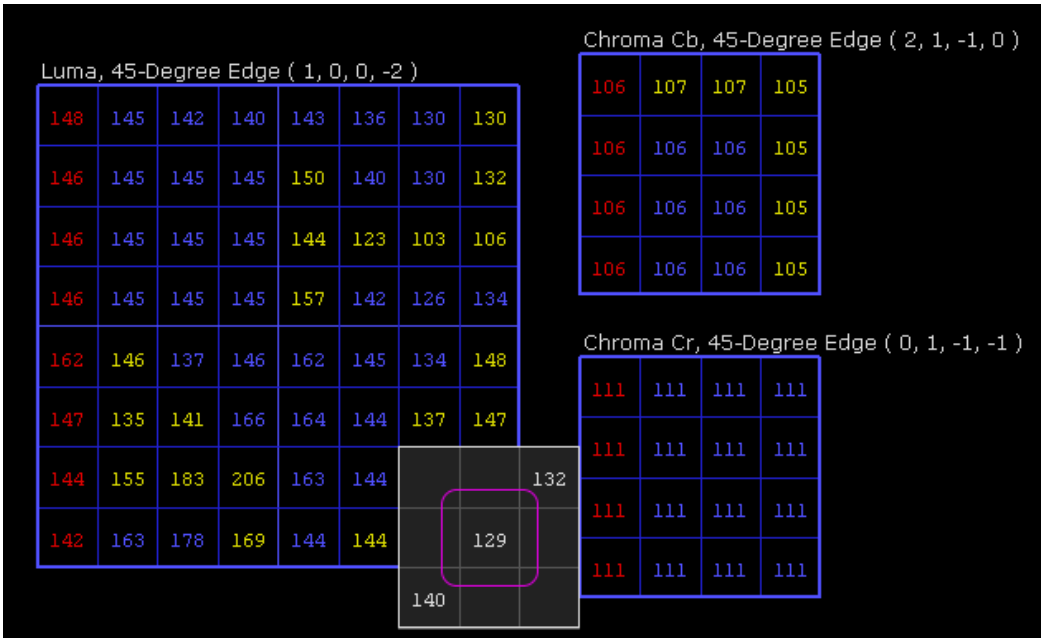

3.2.7 YUV

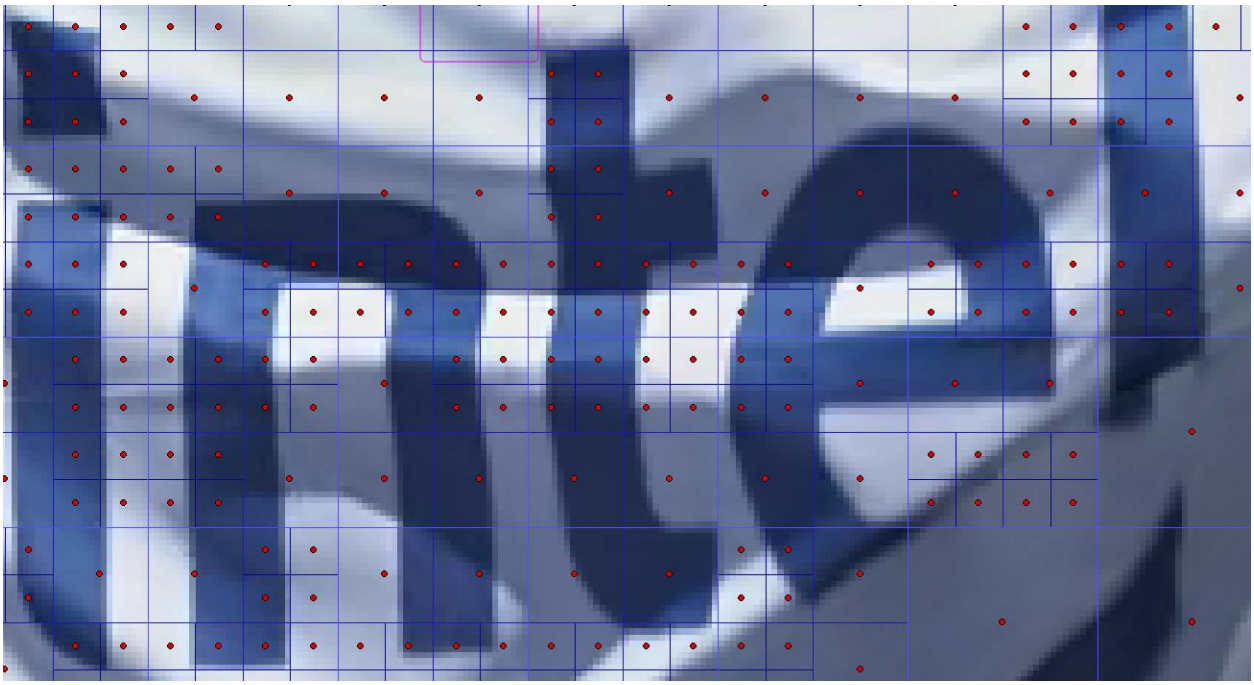

YUV mode allows inspection of the final decoded sample values, without additional overlay data. In this mode, the YUVDiff feature is exposed. This feature allows for comparison of an external decoded YUV file with the decoded bitstream. Supported formats are 4:2:0 planar and NV12. Also, YUV files that are zipped or gzipped may be opened directly without the need to decompress separately. If a zip file contains more than one YUV file, only the first one is used. When a YUV file is loaded, any mismatches will be indicated with a red dot in the CU containing the mismatch. This allows the user to quickly identify the nature of the mismatch which can assist with debug.

When a file is opened, four additional buttons become visible on the lower left of the main window:

- Ref YUV: The original, expected YUV image.
- Debug YUV: The loaded debug YUV file.
- Diff YUV: The delta image. Like in Residual mode, areas with zero delta (i.e. Original and Debug image are identical) are flat gray. Areas where the debug YUV has a lower value are darker, and areas where the debug YUV has a higher value are brighter.
- Reload: Reloads the YUV file, which can be useful as a shortcut to loading via the YUVDiff menu.

3.2.7.1 YUVDiff Menu

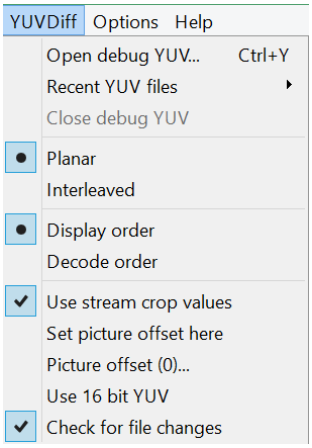

- Open: Brings up a file chooser to open a YUV file. Intel Video Pro Analyzer considers YUV stream as having the same frame size, per frame bitdepth and color format as encoded bitstream.
- Close: Closes the currently opened YUV file. Any mismatch indicators will disappear.
- Planar: Causes Intel Video Pro Analyzer to interpret the loaded YUV file as Planar.
- Interleaved: Causes Intel Video Pro Analyzer to interpret the loaded YUV file as Y values followed by interleaved U and V values. This option is chosen automatically when the loaded YUV file has the file extension ".nv12".
- Display order: Causes Intel Video Pro Analyzer to use display order picture numbers when determining how far to seek into the YUV file to extract the image to compare with the current picture.
- Decode order: Causes Intel Video Pro Analyzer to use decode order picture numbers when determining how far to seek into the YUV file to extract the image to compare with the current picture.
- Use stream crop values: When checked, YUV files are assumed to contain samples only within the cropping window as defined by the loaded bitstream. Samples outside this window are not compared, and are assumed 0 when viewing the loaded YUV image directly.
- Set picture offset here: Shortcut for setting the picture offset to the current picture number. See bullet below.
- Use 16 bit YUV: this option allow you to load YUV file where both Luma and Chrome written in 16 bit (they could be normalized from actual bitdepth to 16 bit during encoding). This is useful when you have bitdepth changing across the stream. Intel Video Pro Analyzer will shift values to appropriate bitdepth on the current frame.
- Picture offset: This brings up a dialog allowing the user to enter the picture number of the first picture in the YUV file. For example if a 100-picture bitstream is loaded but the YUV file only contains pictures 80-99, the user would enter 80 to properly line up the YUV file with the decoded bitstream.
- Check for file changes: When checked, this option causes Intel Video Pro Analyzer to periodically check if the loaded YUV file has changed on disk since it was last loaded. If the file has indeed changed, a dialog pops up offering a chance to reload the YUV file.

3.2.7.2 YUV Detail Mode

In detail mode the expected YUV values can be inspected. Note that the displayed sample values are the same as those in SAO mode since SAO filtering is the last step in the decode process. When a YUV file is loaded, the debug and delta YUV values can be inspected as well. Mismatching values are shown in red.

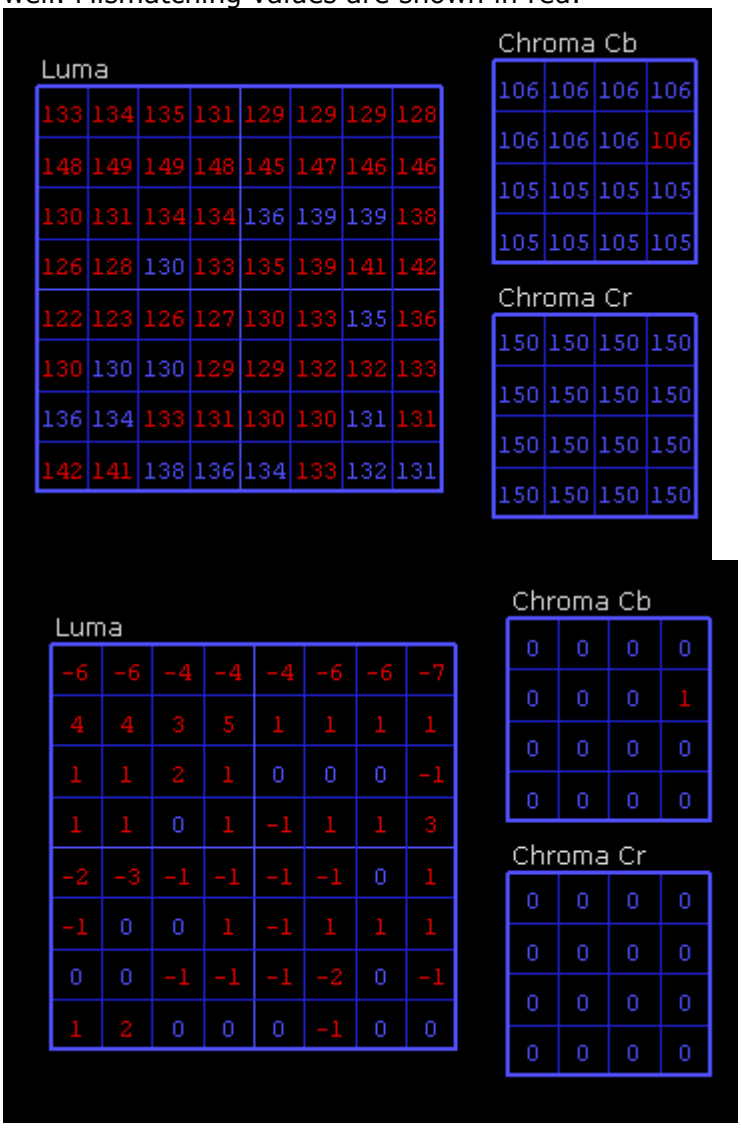

3.2.8 Info Overlays

There are a number of overlay that provides additional information about currently investigated picture. When you switch mode to "Info Overlays" mode several additional submenus will appear in the bottom-left panel of info window.

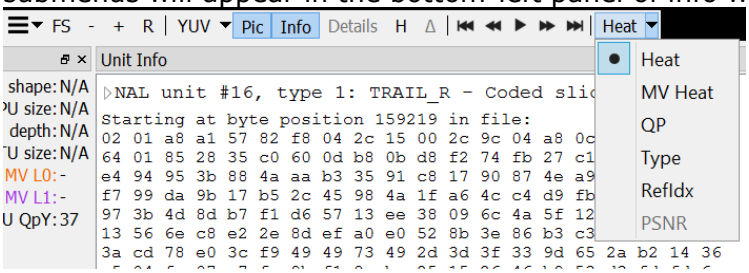

You can switch now between additional modes. PSNR mode is enabled if you have loaded reference YUV file.

3.2.8.1 Heat Map

Heat map mode shows visually how the compressed bits of the picture are distributed spatially. CUs with more bits per pixel are brighter than CUs with less. By zooming in actual CU sizes in bits will be displayed as soon as it size in the borders of the CU. By default this mode also shows the CU quadtree boundaries in blue. It may be useful to turn off this overlay using the "Info" button on the lower left of the main panel. Heat Map does not offer a detail mode.

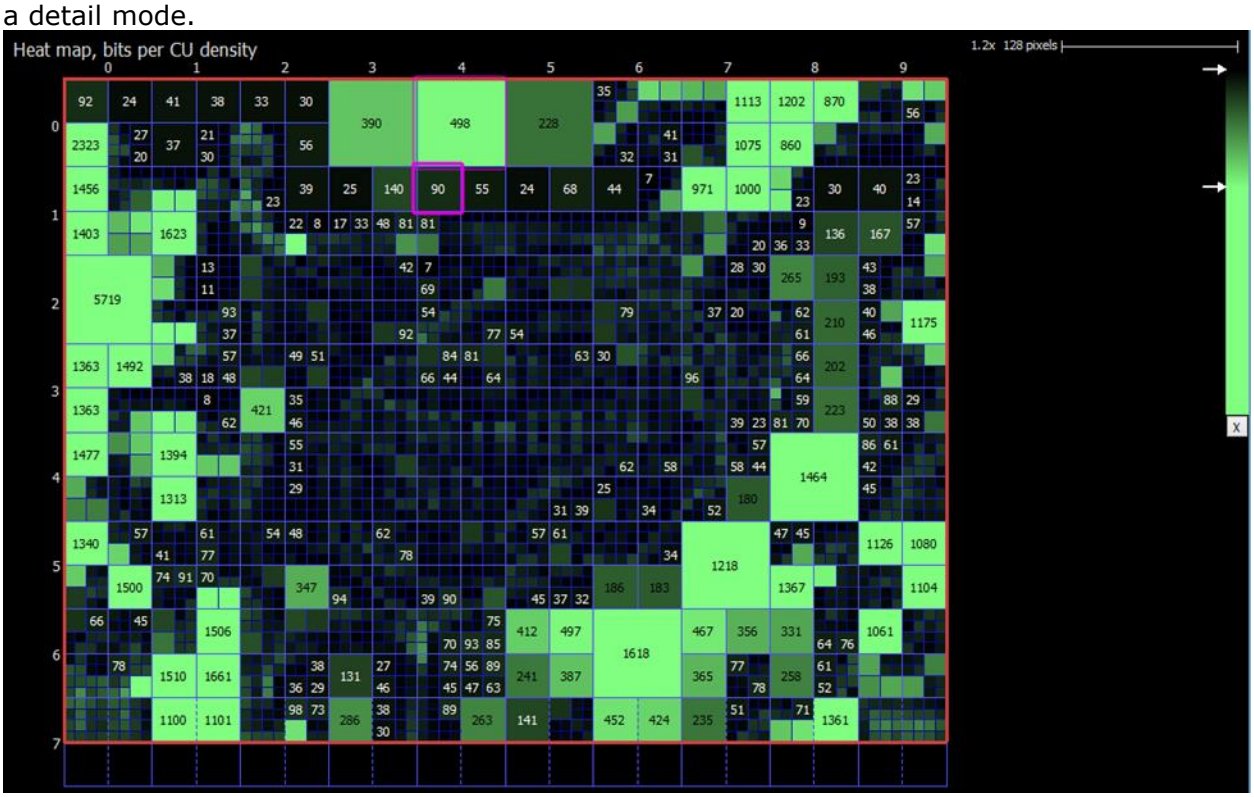

The gradient used to display the heat map may be edited:

- Drag the gradient markers up and down to reposition them.
- Drag a marker away from the gradient to remove it. A red X indicates removal will take place.
- Double-click a marker to change the color.
- Double click the gradient bar to add a new marker.

3.2.8.2 QP Map

QP Map show per CU QP. CUs with bigger QP are brighter than CUs with less. Coloring of the QP range can be adjusted with a gradient tool on the right top corner.

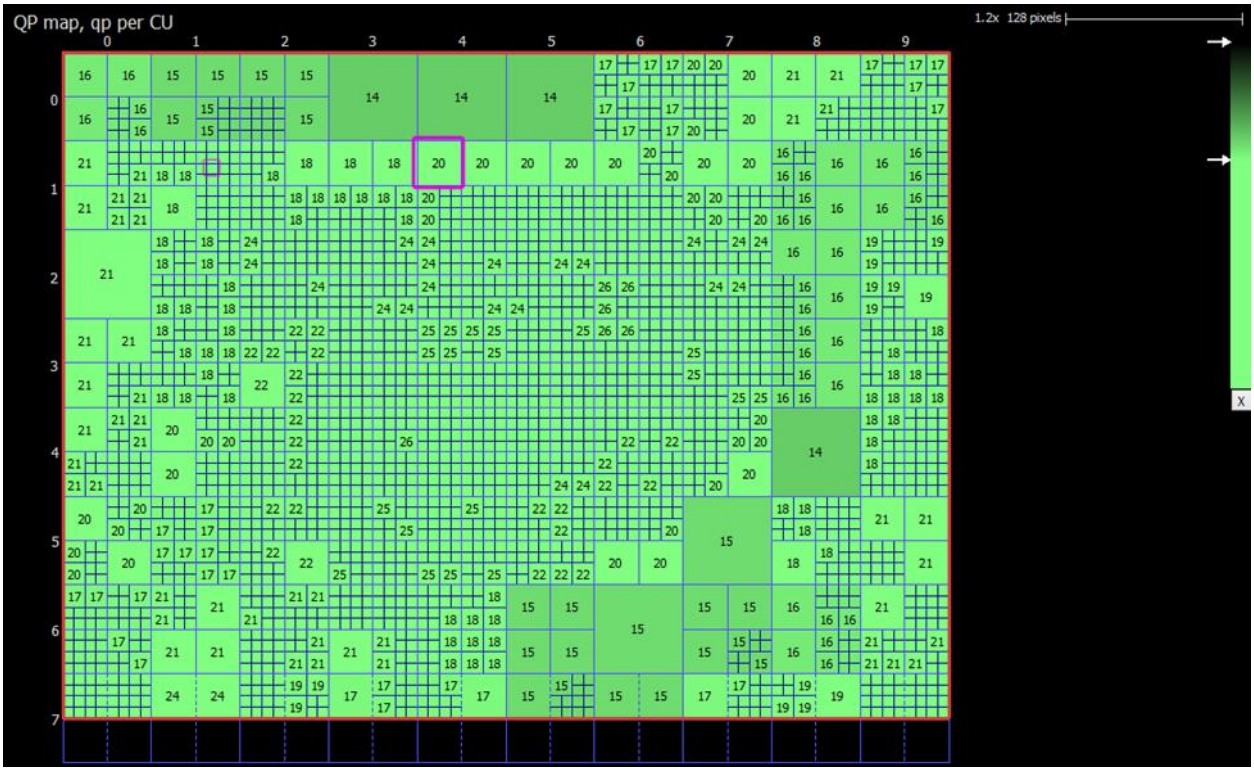

3.2.8.3 PU Type

PU Type overlay show color map of their types. Intra blocks is colored in red, inter skip is yellow colored, inter merge is in green and inter in blue. The legend on the right top corner show color types for blocks. Gradient band can adjust transparency when the Pic switch is on.

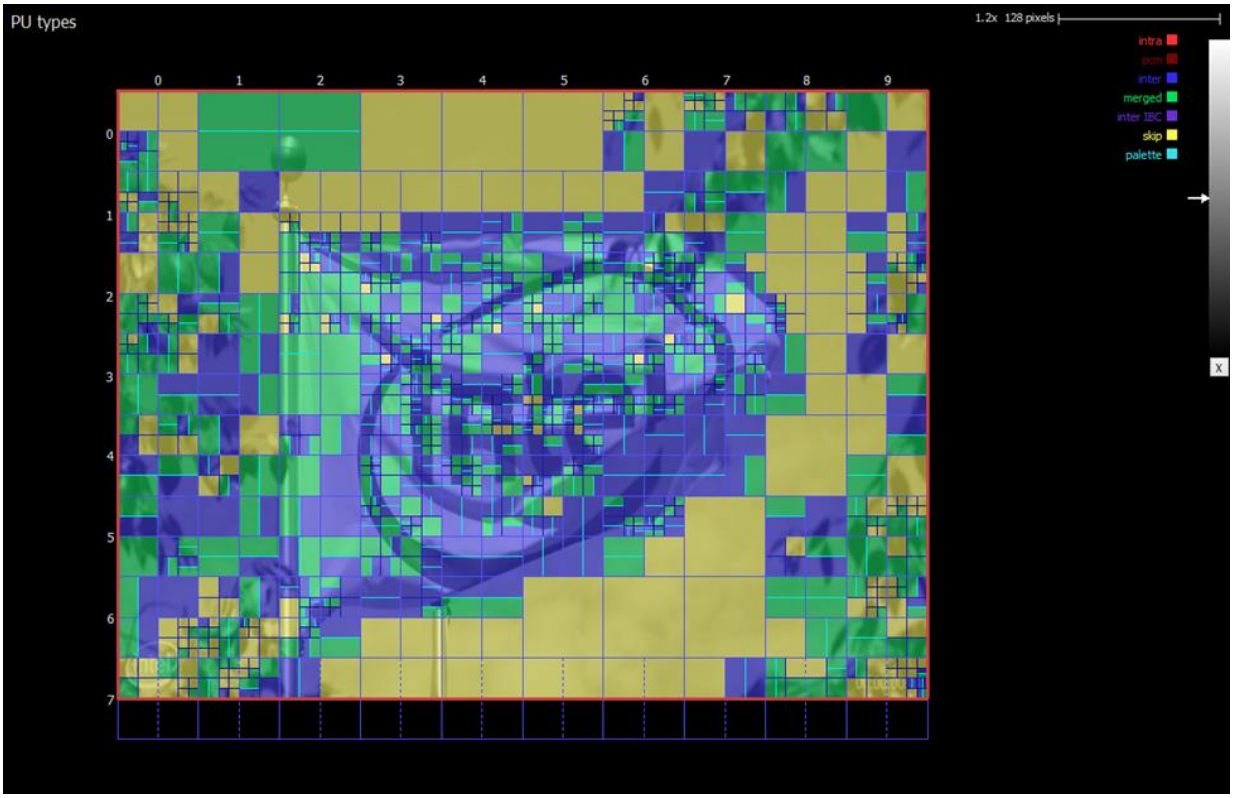

3.2.8.4 PU Reference Indices

PU reference indices shows reference indices for each PU. Orange color shows L0 reference index and purple color is for L1 reference index. When PU has only one reference missing reference index is displayed with dash. For blocks with no reference indices nothing is shown.

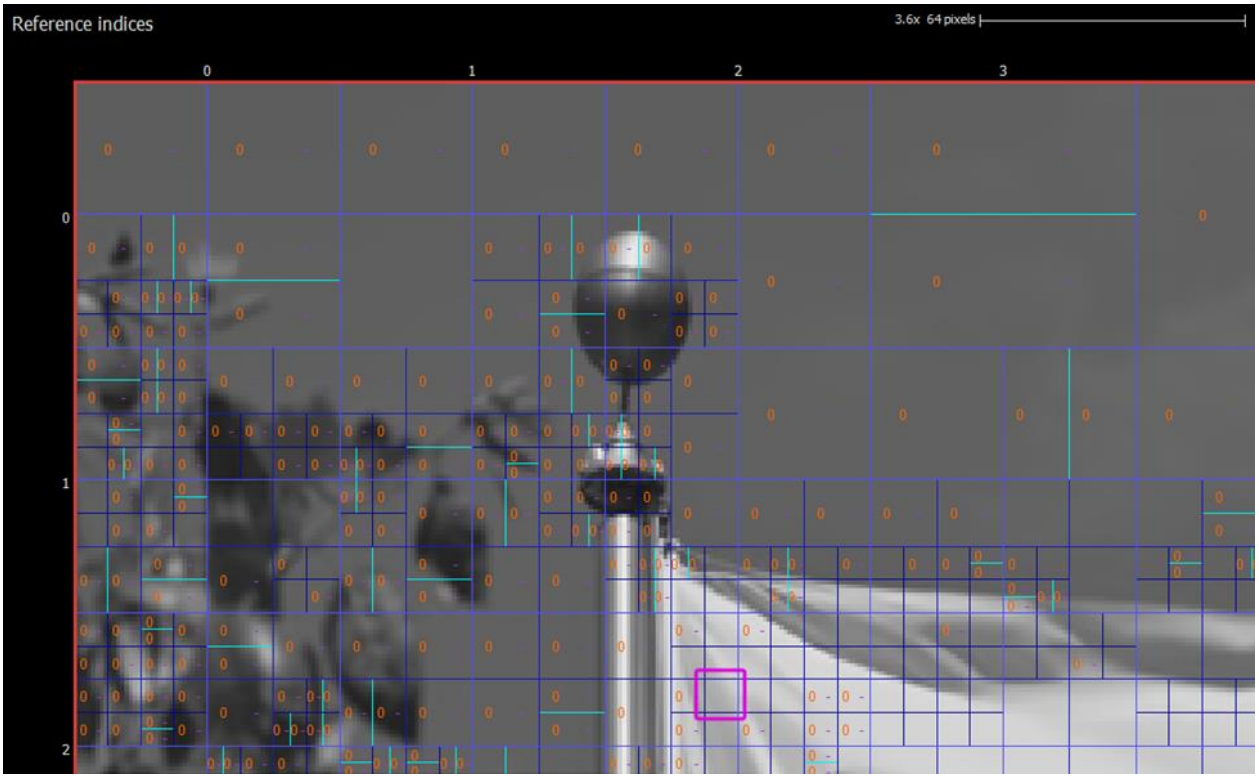

3.2.8.5 PSNR

PSNR map shows per pixel PSNR value normalized on CU size. To enable this mode you have to load reference YUV file via DebugYUV menu. The value is calculated per CU. Bigger value has brighter color. The color range can be adjust with gradient tool on the right top corner. If the block has no difference with reference YUV a latter 'e' is shown which means exact match.

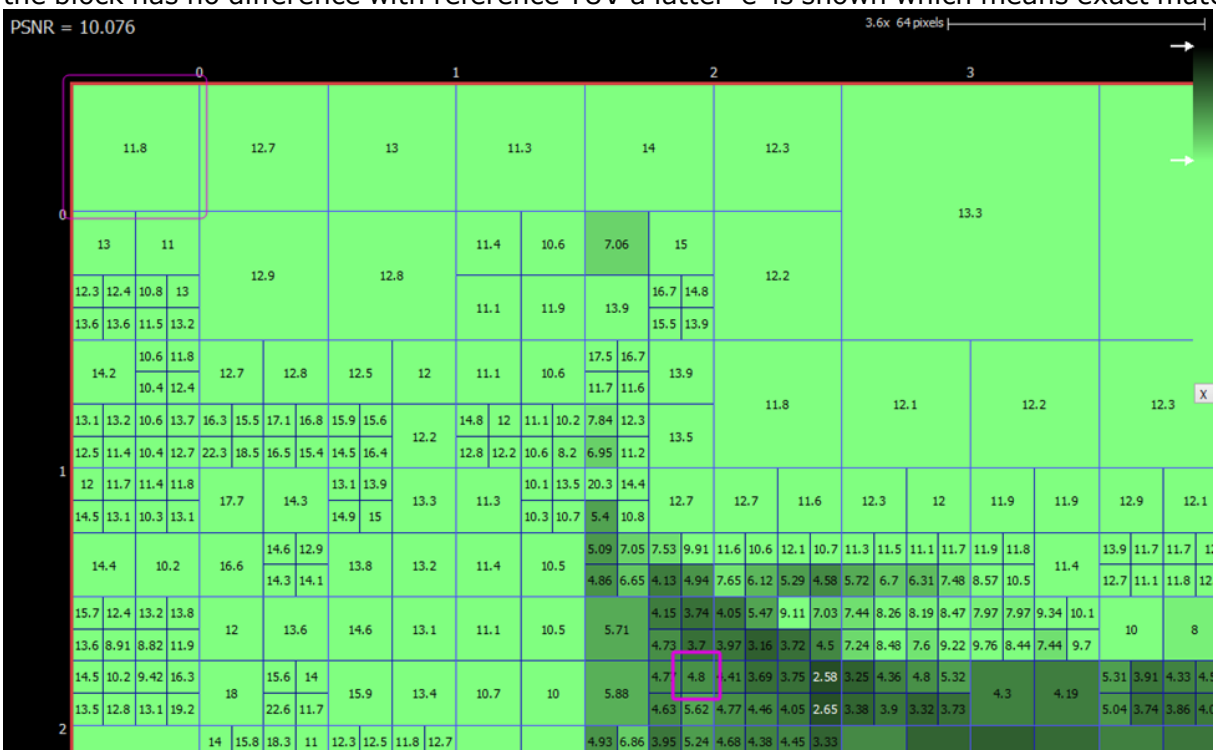

3.2.9 Simple Motion

Simple Motion mode offers a way to quickly view the modes and motion of a picture at a glance. No overlays are present except a colored dot showing information about the PU's mode. Green dots indicate intra PUs, and purple/orange L0/L1 motion vectors for inter PUs are drawn as simple lines. When the motion vector is very small, a dot is drawn instead so that the mode is still easy to identify. When a motion vector does not point to the 0-th index of its respective list, the vector is drawn using a dashed line.

4 VP9

The following sections describe all available features when loading a VP9 bitstream. The bitstream must be containerized with either the WebM or IVF container format. The publicly available VP9 reference software can output both formats. Raw uncontainerized bitstreams are presently not supported due to the nature of the VP9 standard; there is no way to determine where each frame starts without doing a full decode of the sequence.

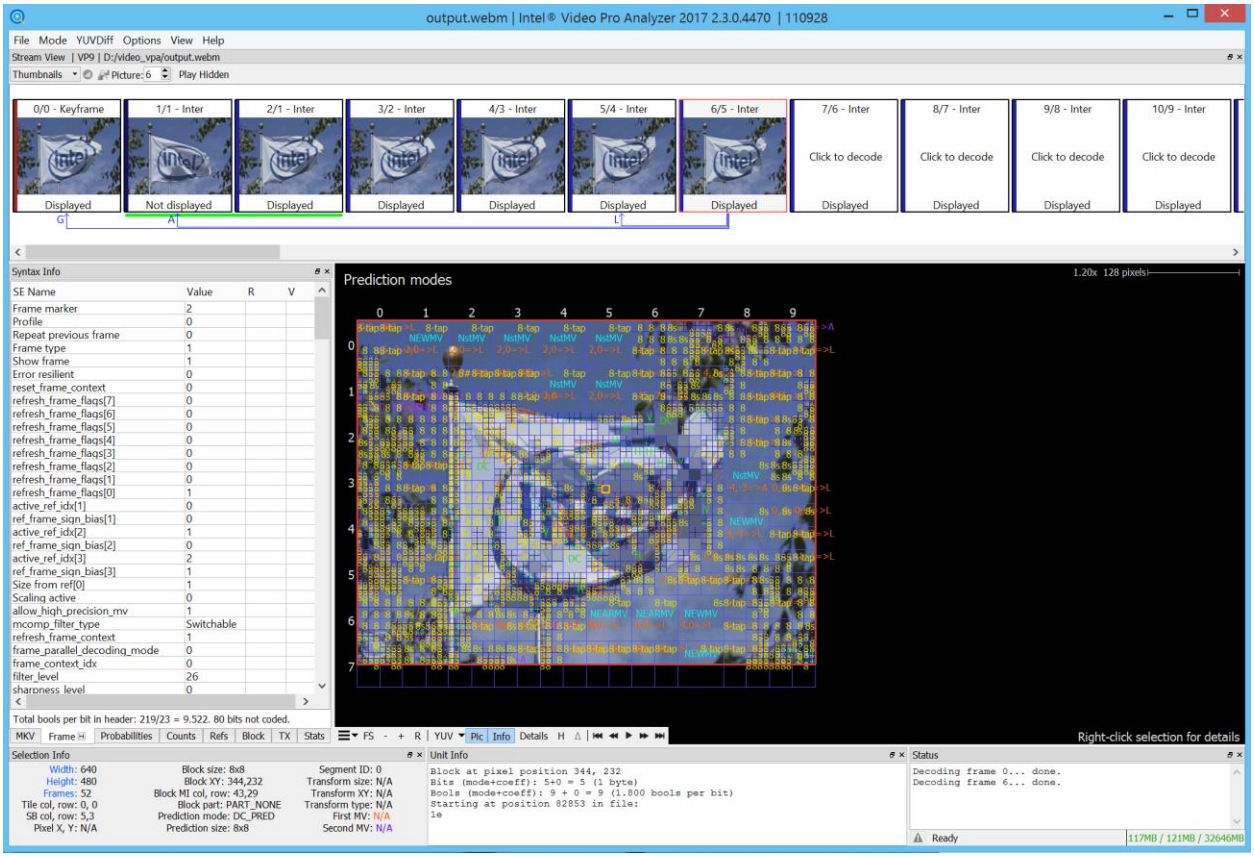

4.1 UI Components

The following sections describe the various components in the UI when an VP9 bitstream is loaded.

4.1.1 Top Filmstrip

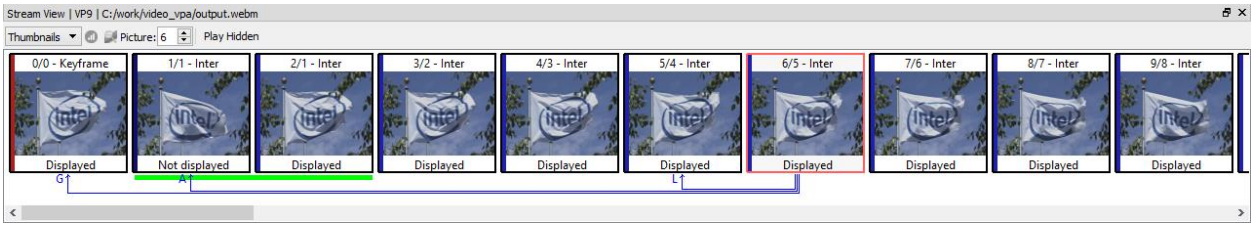

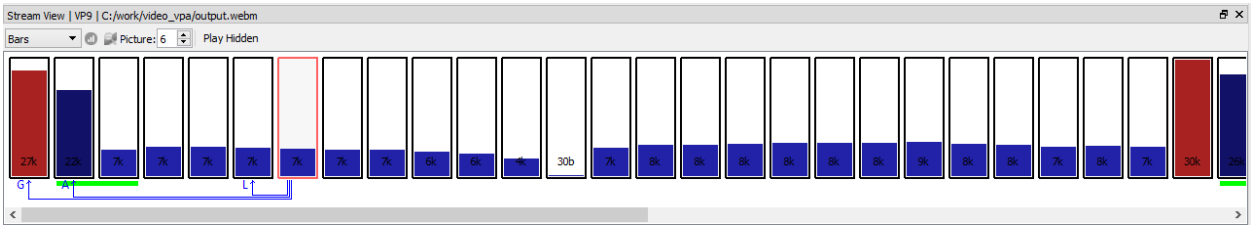

The top filmstrip is a horizontally scrolling overview of the pictures in the bitstream. The current picture is highlighted. On top is the frame number in both decode and display order respectively and the picture type as indicated by the frame header. On bottom is indicated if the picture should be displayed or not. Arrows indicate the reference frames.

Clicking on a picture will cause Intel Video Pro Analyzer to decode it and mark it as the current picture, updating the rest of the UI correspondingly. If the new picture requires other pictures to be available as reference that have not yet been decoded, those will automatically be decoded.

When a frame from the IVF or WebM/MKV container is a super frame (meaning 1 or more non-displayed picture and one displayable), this collection of pictures has superscription with line over those frames belong to the super frame. The bar filmstrip has underline green band under those frames belong to the same super frame.

Pictures are decoded on demand since it is impractical to store all details of all pictures in memory, especially for HD sequences. The current picture can also be selected by typing in a picture number in the box on the left of the panel. When using this box, the filmstrip is scrolled such that the newly selected current picture is visible.

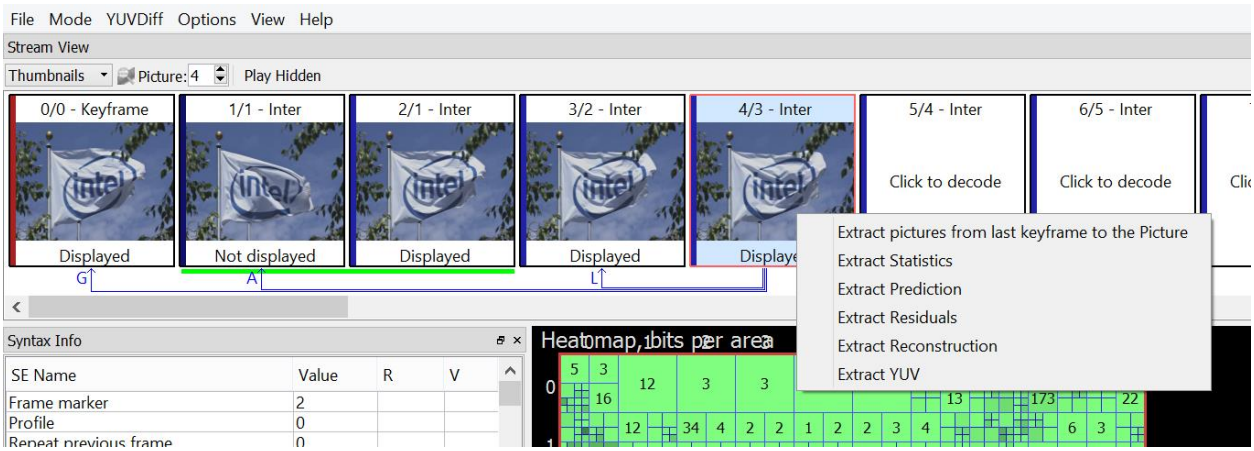

By right clicking on selected picture one can get popup submenu with a set of extract actions. An action called "Extract pictures form last keyframe to the Picture" will write out a bitstream that contains the minimum number of pictures needed to decode the current picture. Typically this will consist of the current picture and all previous pictures that are reference pictures up to the nearest keyframe picture. This can be useful for debugging issues in long sequences. Other actions allow to save in the output file statistics and intermediate/final YUV planes.

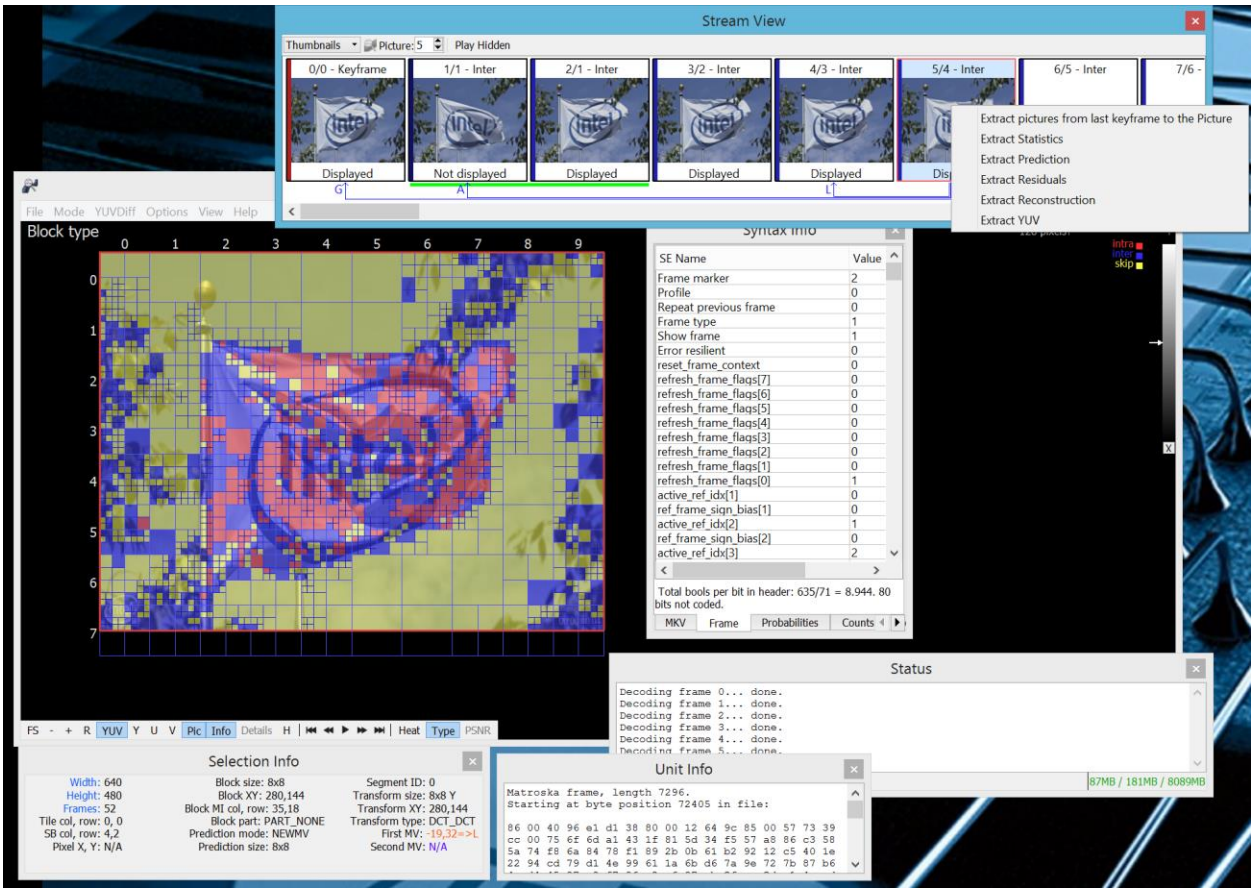

Allow separate panels can be detached from them main window or hidden. To attach panels back double click on the titlebar. To unhide panel go to View menu and check appropriate panel.

4.1.2 Main Panel

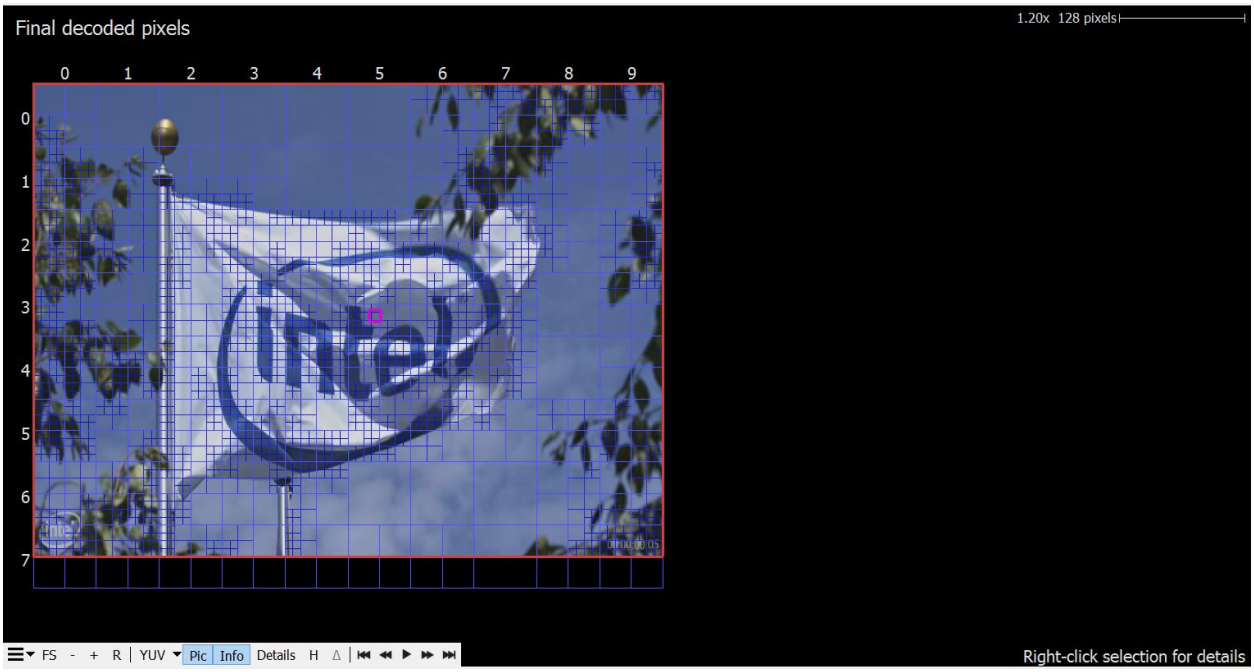

The main panel displays the selected picture with visual annotations. The type of annotations and associated interactive behavior depends on the current mode, see 4.2 VP9 Modes. Click-dragging moves the picture around, and the mouse wheel zooms in or out about the mouse cursor.

At the top left of the panel, the current mode is displayed. Along the top right the current scale is shown. In all modes, the super-block row and column values are displayed along the left and top border of the picture and tile boundaries are thick green lines. The screenshot above shows an example of a stream that is 13x8 SBs.

The bottom left has a cluster of buttons:

- First submenu has special options to output the currently displayed information to the external image file with different scaling possibilities.
- FS: Toggles full-screen mode, which hides all UI elements except the main panel. Keyboard shortcut is F.
- +/-: Zooms in or out, centered on the center of the panel. Note that zooming with the mousewheel is much easier than using these buttons. Keyboard shortcuts are + and -.
- R: Resets the current zoom to fit the entire picture in the main panel. Keyboard shortcut is R.
- YUV, Y, U, V: Toggles between per-component image, or full-color YUV.
- Pic: Toggles the actual picture underneath the annotations. The picture is mode-dependent. For example in prediction mode, the picture is comprised of the prediciton samples before the residual signal has been added in. It is often helpful to turn off the picture in order to make the annotation easier to read. Keyboard shortcut is P.
- Info: Toggles the annotations on or off. Keyboard shortcut is I.
- H/D: Toggles between hexadecimal and decimal display of the annotations. This applies to the values in the Left Tabs as well. Keyboard shortcuts are H and D.

- There are five navigation buttons: first frame, previous frame, playback mode, next frame, last frame. The previous/last button will rewind, causing the current picture to be set to the first/last picture in the loaded bitstream. The "Play" button will start playback of loaded bitstream at a reasonable rate, limited by CPU performance. All normal mode operations are possible during playback. Note that playback mode always operates in currently selected order in the filmstrip.

4.1.3 Left Tabs

The left side of the UI is a resizable panel with several tabs that display information about the current selection in one way or another. The following subsections describe each tab.

4.1.3.1 Container (IVF or MKV)

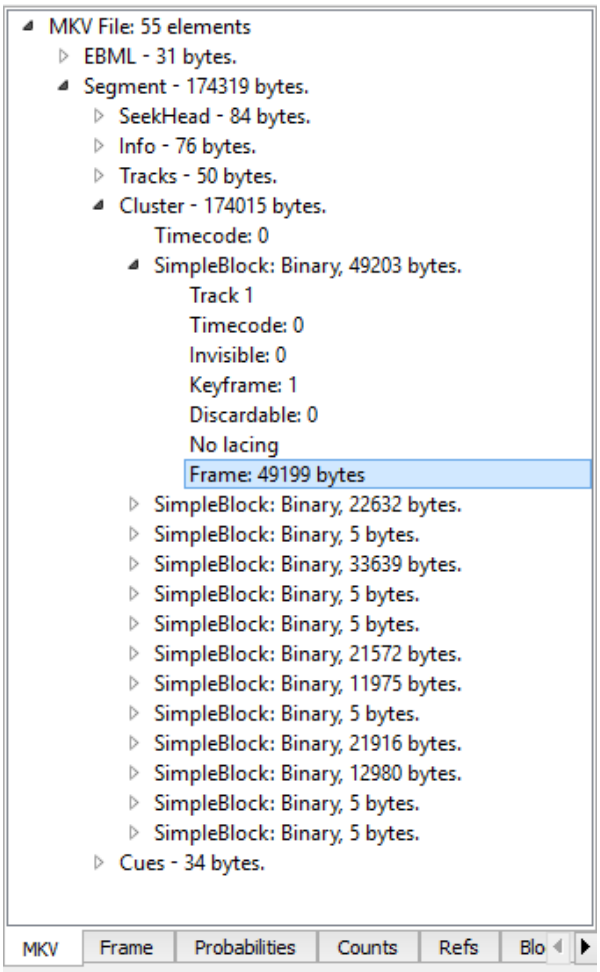

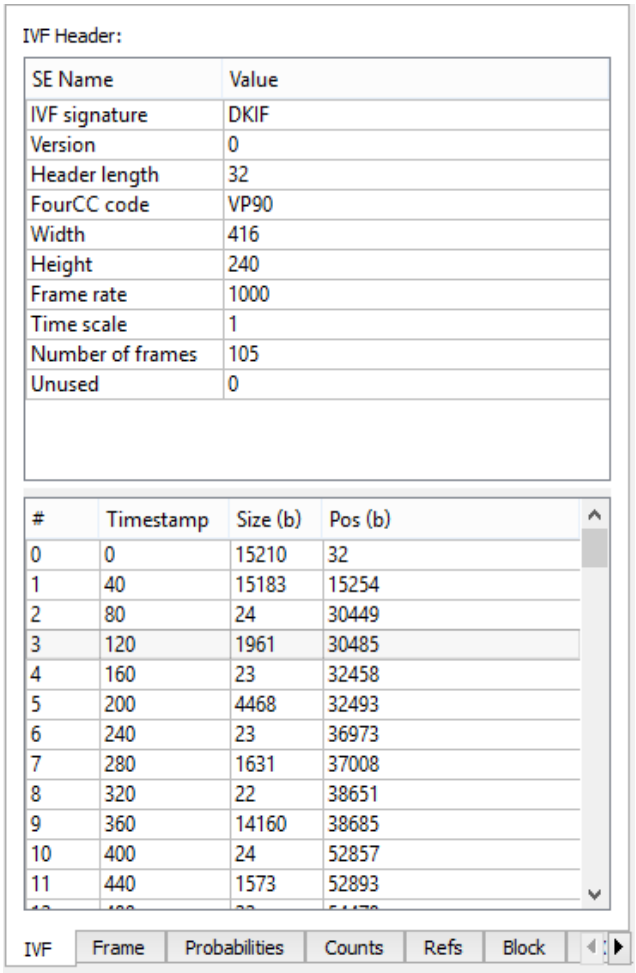

Depending on the container type of the loaded bitstream, either the MKV or IVF panel will be present in the UI.

- VP9 streams using the WebM or Matroska container will cause the UI to display the container's EBML document in tree form. Clicking an element causes the raw bytes coding that element to be displayed in the Raw Bytes panel (bottom center of the UI). When decoding a particular frame, the corresponding element will be highlighted in this panel. Other tracks may be present in the MKV file (such as audio), but they are not processed. The raw bytes can still be viewed though.

- Streams containerized with the IVF format causes the UI to display the IVF information in this left-most tab. Selecting a frame causes the Raw Bytes panel to show the raw bytes of that frame, including the IVF frame header. Clicking a frame will causes that frame to be decoded.

4.1.3.2 Frame

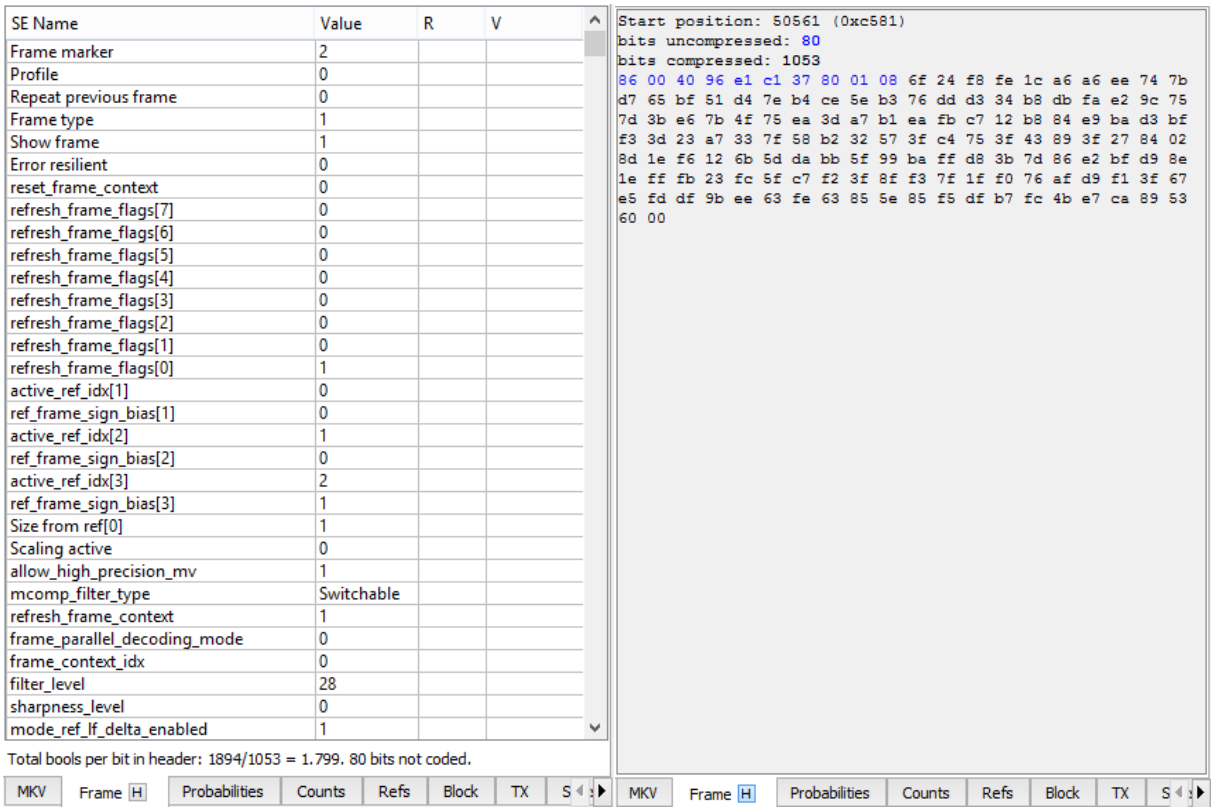

SE Name	Value	R	V
Frame marker	2		
Profile	0		
Repeat previous frame	0		
Frame type	1		
Show frame	1		
Error resilient	0		
reset_frame_context	0		
refresh_frame_flags[7]	0		
refresh_frame_flags[6]	0		
refresh_frame_flags[5]	0		
refresh_frame_flags[4]	0		
refresh_frame_flags[3]	0		
refresh_frame_flags[2]	0		
refresh_frame_flags[1]	0		
refresh_frame_flags[0]	1		
active_ref_idx[1]	0		
ref_frame_sign_bias[1]	0		
active_ref_idx[2]	1		
ref_frame_sign_bias[2]	0		
active_ref_idx[3]	2		
ref_frame_sign_bias[3]	1		
Size from ref[0]	1		
Scaling active	0		
allow_high_precision_mv	1		
mcomp_filter_type	Switchable		
refresh_frame_context	1		
frame_parallel_decoding_mode	0		
frame_context_idx	0		
filter_level	28		
sharpness_level	0		
mode_ref_lf_delta_enabled	1		

Total bools per bit in header: 1894/1053 = 1.799. 80 bits not coded.

MKV | Frame H | Probabilities | Counts | Refs | Block | TX | S

```
Start position: 50561 (0xc581)
bits uncompressed: 80
bits compressed: 1053
86 00 40 96 e1 c1 37 80 01 08 6f 24 f8 fe 1c a6 a6 ee 74 7b
d7 65 bf 51 d4 7e b4 ce 5e b3 76 dd d3 34 b8 db fa e2 9c 75
7d 3b e6 7b 4f 75 ea 3d a7 b1 ea fb c7 12 b8 84 e9 ba d3 bf
f3 3d 23 a7 33 7f 58 b2 32 57 3f c4 75 3f 43 89 3f 27 84 02
8d 1e f6 12 6b 5d da bb 5f 99 ba ff d8 3b 7d 86 e2 bf d9 8e
1e ff fb 23 fc 5f c7 f2 3f 8f f3 7f 1f f0 76 af d9 f1 3f 67
e5 fd df 9b ee 63 fe 63 85 5e 85 f5 df b7 fc 4b e7 ca 89 53
60 00
```

MKV | Frame H | Probabilities | Counts | Refs | Block | TX | S

The Frame tab shows all syntax elements decoded in the frame header, both the uncompressed and bool-coded partitions. A few stats about the size of the header are shown below the list.

Since the VP9 reference software does not always assign a variable name to each decoded token, a descriptive name has been assigned which conveys the meaning or usage of the token.

The two left columns R and V represent the value of the bool coder's range and value prior to decoding the token. For the uncompressed part of the frame header, these values are not present.

By pressing on 'H' button in the header tab one can show hexadecimal representation of the frame header. This representation show all bytes of the header, starting position, number of bytes in the compressed and uncompressed sections. Uncompressed bytes are highlighted in blue color.

4.1.3.3 Probabilities

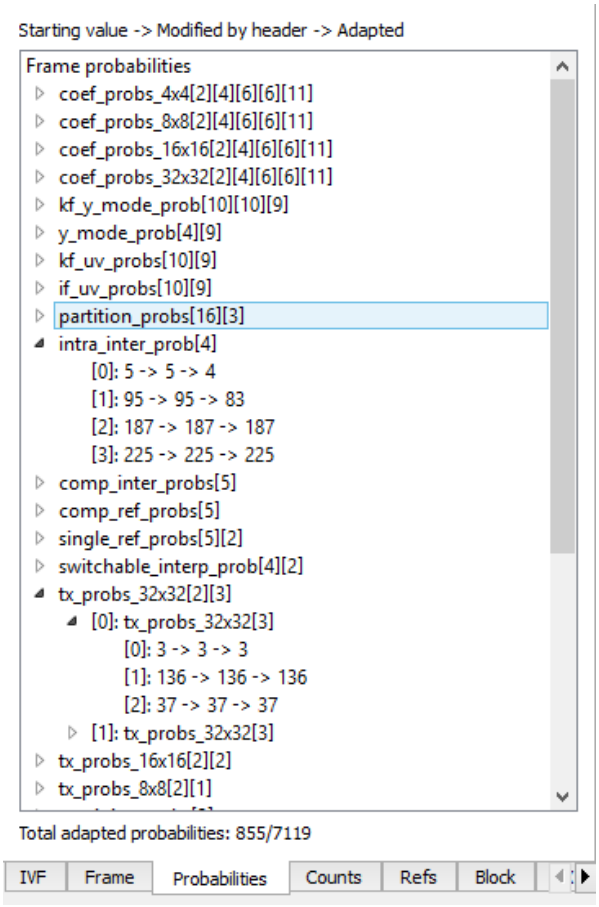

The probabilities used to decode the current frame are visible in this panel. They are organized hierarchically in the same way as the VP9 reference software. Each probability is given as three numbers. From left to right:

1. Probability as it was at the start of the frame decode process. For keyframes or when the reset_frame_context flag is equal to 1, this is the VP9 default value.
2. Probability value after modification by the bool-coded section of the frame header.
3. Probability value after modification by the adaptation process after frame decode.

When the probability adaptation process causes the value to actually change, the probability is shown in blue. The total number of probability values that were actually changed is given below the tree.

4.1.3.4 Counts

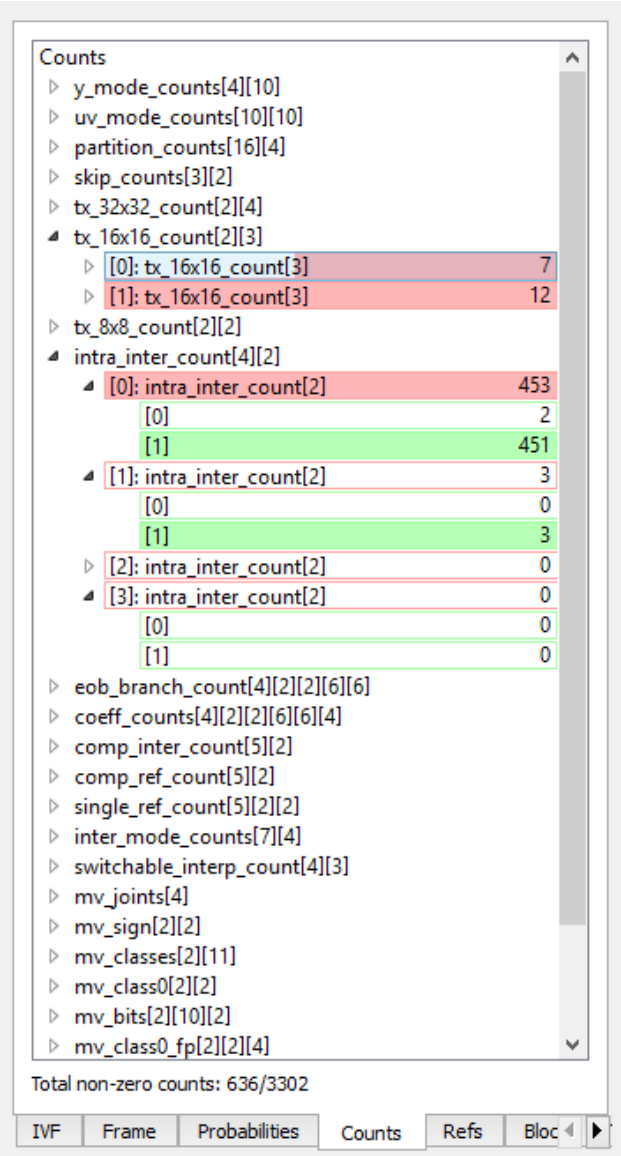

This tab shows the values of all the counters used for probability adaptation at the end of the frame decode. The counters are organized in hierarchical fashion in the same way as the VP9 reference decoder. Color bars give an indication of relative size of each counter in its group, and the parent icon contains the sum of the counts.

4.1.3.5 Refs

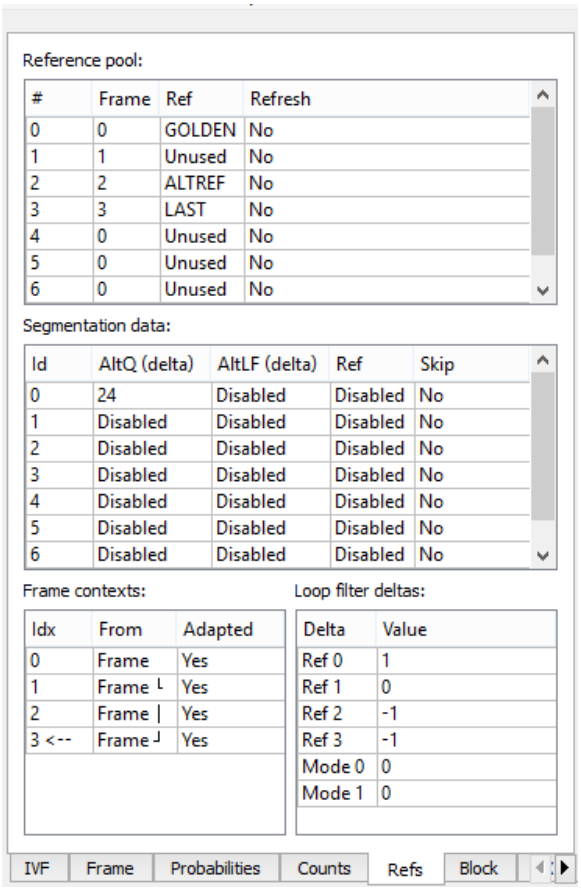

This panel displays the data that persists in the decoder between frames:

- The 8-entry pool of reference frames and which of each of these is considered the Last, Golden and AltRef frame. If any of the frames should be replaced with the current frame after decode (as specified by refresh_frame_flags), the right most column will indicate that.
- Segmentation data for each segment.
- The 4-entry frame context buffer as it applies to the current frame. The center column indicates the frame that generated that context, and the right column indicates whether or not it was the result of adaptation at that time.
- Loop filter deltas for all four references and both modes.

4.1.3.6 Block

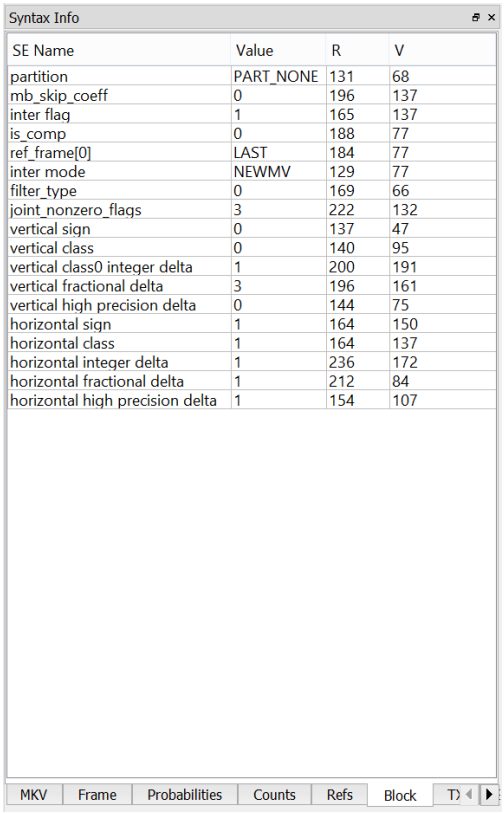

This panel displays the syntax elements associated with the currently selected block. Since the VP9 reference software does not always assign a variable name to each decoded token, a descriptive name has been assigned which conveys the meaning or usage of the token.

Note that only the syntax elements at the currently selected block hierarchy are shown. To see the PART_SPLIT tokens at higher levels in the recursive subdivision process, repeatedly click the same block in the main panel to navigate the hierarchy.

The two left columns R and V represent the value of the bool coder's range and value prior to decoding the token.

4.1.3.7 TX

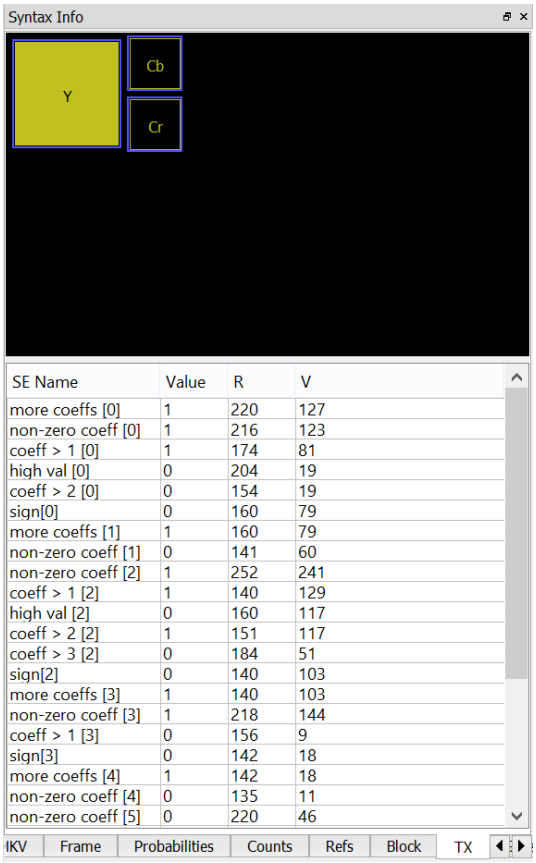

This panel displays the syntax elements associated with the currently selected transform block. Since the VP9 reference software does not always assign a variable name to each decoded token, a descriptive name has been assigned which conveys the meaning or usage of the token.

A particular transform block may be selected in the Main Panel (when in Residual mode), or by clicking in the diagram above the syntax list. Note that when not in detail mode, only the luma transform blocks are selectable from the Main Panel.

The two left columns R and V represent the value of the bool coder's range and value prior to decoding the token.

4.1.3.8 Stats

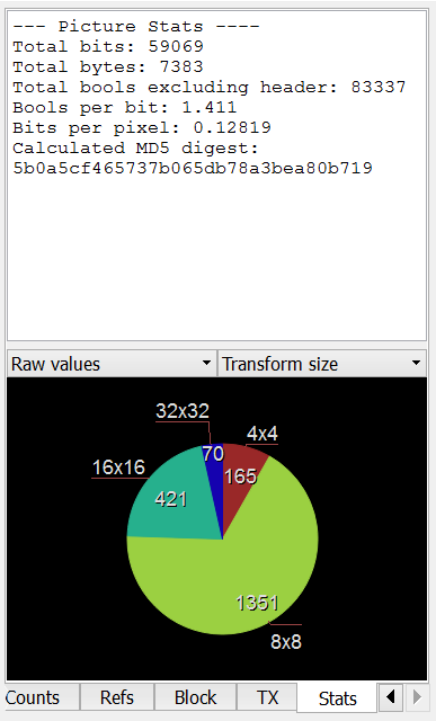

This tab displays various statistics extracted from the current picture. The top half shows some picture size and compression stats as well as the image digest information. The bottom half displays pie charts for a number of metrics. Each pie chart can be drawn normalized or un-normalized. Normalized data is weighted by area or bits. For example there may be a much lower number of 64x64 blocks than 8x8 blocks in the frame (smaller un-normalized pie wedge), but they could still make up the majority of the picture area, making the 64x64 pie wedge large. Normalized numbers in the pie chart are in units of pixels or bits, and un-normalized numbers are raw counts.

The pie chart can be moved by dragging the mouse, and zoomed with the mouse scroll wheel.

4.1.4 Selection Info

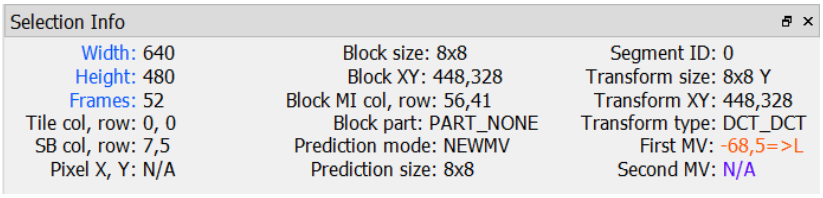

This panel shows a few details about the current selection at a glance:

- Frame height & width
- Number of frames in the bitstream
- Tile row and column number
- Current super-block row and column number
- X/Y position of the mouse cursor in the frame
- Size of the selected block
- X/Y position in pixels of the selected block's upper left corner
- MI (8x8 granularity) column and row number of the upper left of the selected block
- Selected block's partitioning mode
- Selected block's prediction mode
- Size of the selected prediction block
- Segment ID of the selected block
- Transform size used in the selected block (luma)
- X/Y position in pixels of the selected transform block (Residual Mode)
- Type of transform used on the selected transform block (Residual Mode)
- Value of the motion vector(s) used by the current prediction block in $1/8^{th}$ pel. An arrow and letter indicate the reference frame: L for Last, G for Golden and A for AltRef.

4.1.5 Raw Bytes

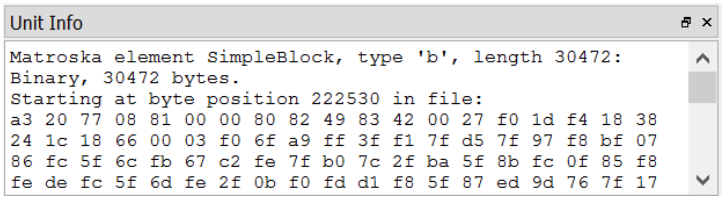

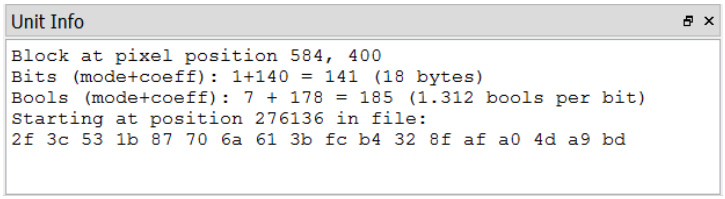

This resizable panel displays raw bytes from the bitstream as it applies the current selection, as well as few extra details. The current selection can be a particular block, a Matroska element/frame or an IVF frame (depending on the current bitstream's container).

4.1.6 Messages

This resizable panel in the lower right of the UI displays messages about the decoding process, and shows progress of any actions that may take a while to complete.
The 3 numbers in the lower right of this panel represent the memory state of the proccess. From left to right the 3 numbers are:

- Amount of memory used by process
- Peak memory used by process
- Total memory claimed by system

4.2 Modes

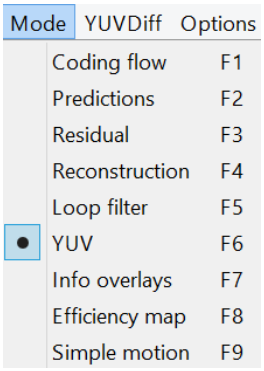

With a VP9 bitstream loaded, Intel Video Pro Analyzer can be put into one of 9 modes using either the F1-F9 keys, or using the Mode menu. The mode selection affects only what is displayed in the Main Panel.

Most modes can also show details of the current selection (block, prediction/transform block). This can be toggled with the right mouse button or the main panel's button strip.

4.2.1 Coding Flow

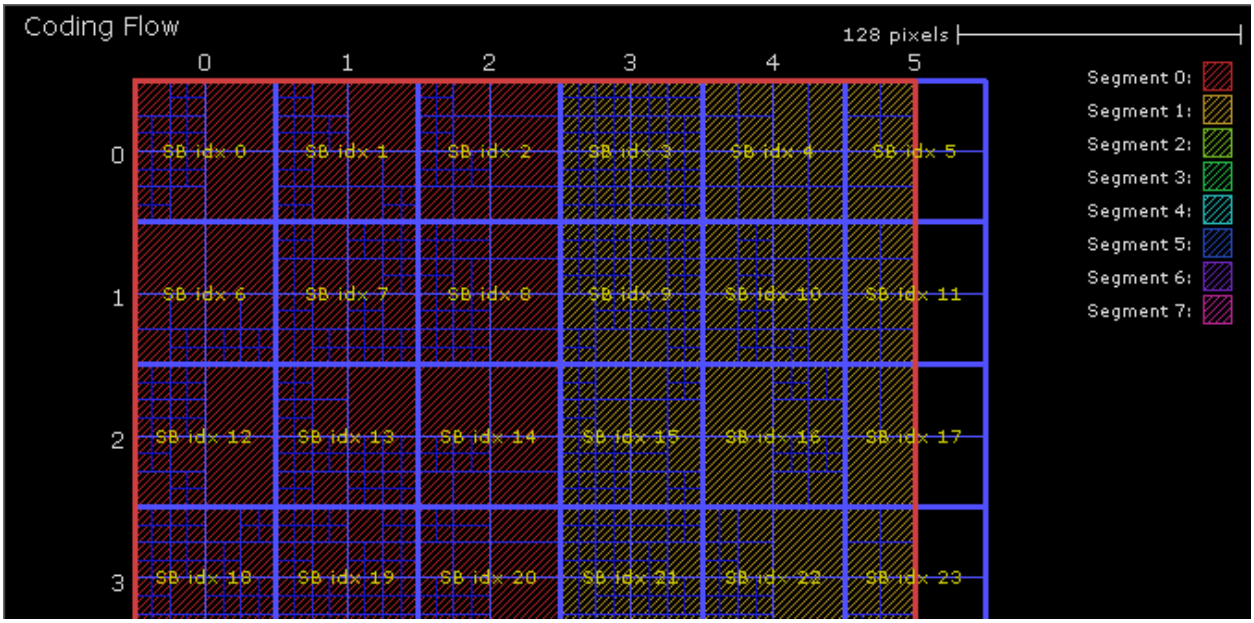

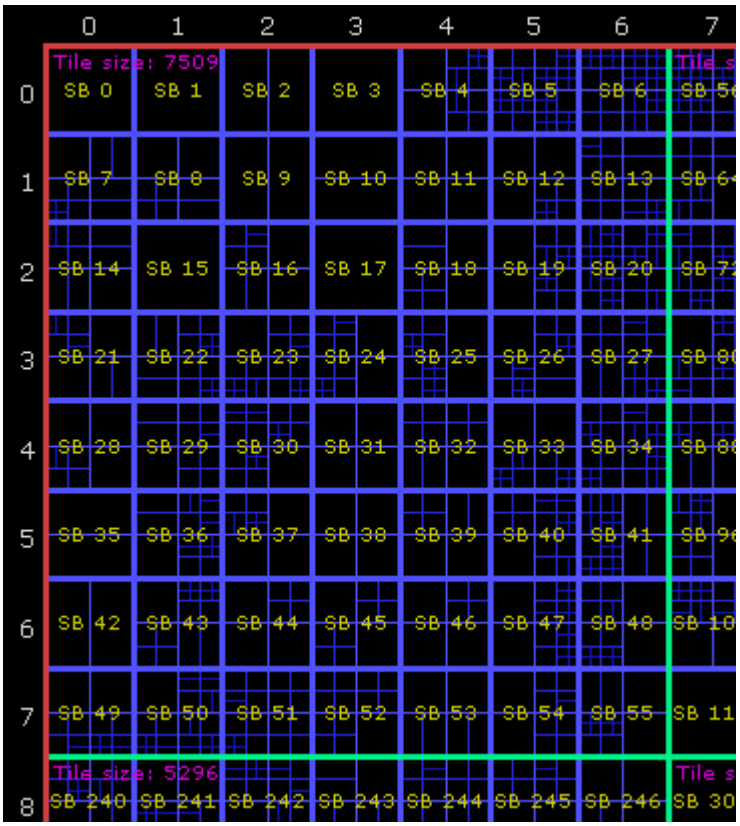

Coding flow mode gives an overview of the way the frame is constructed. Each superblock's decode index is shown in yellow and block-partioning is indicated with blue boundaries. Tile boundaries, when present, are shown with thick green lines and the tile size as coded in the bitstream is shown at the top left of each tile that has a size coded. When segmentation is enabled, each segment is shaded with a unique color. No detail mode is available.

4.2.2 Predictions

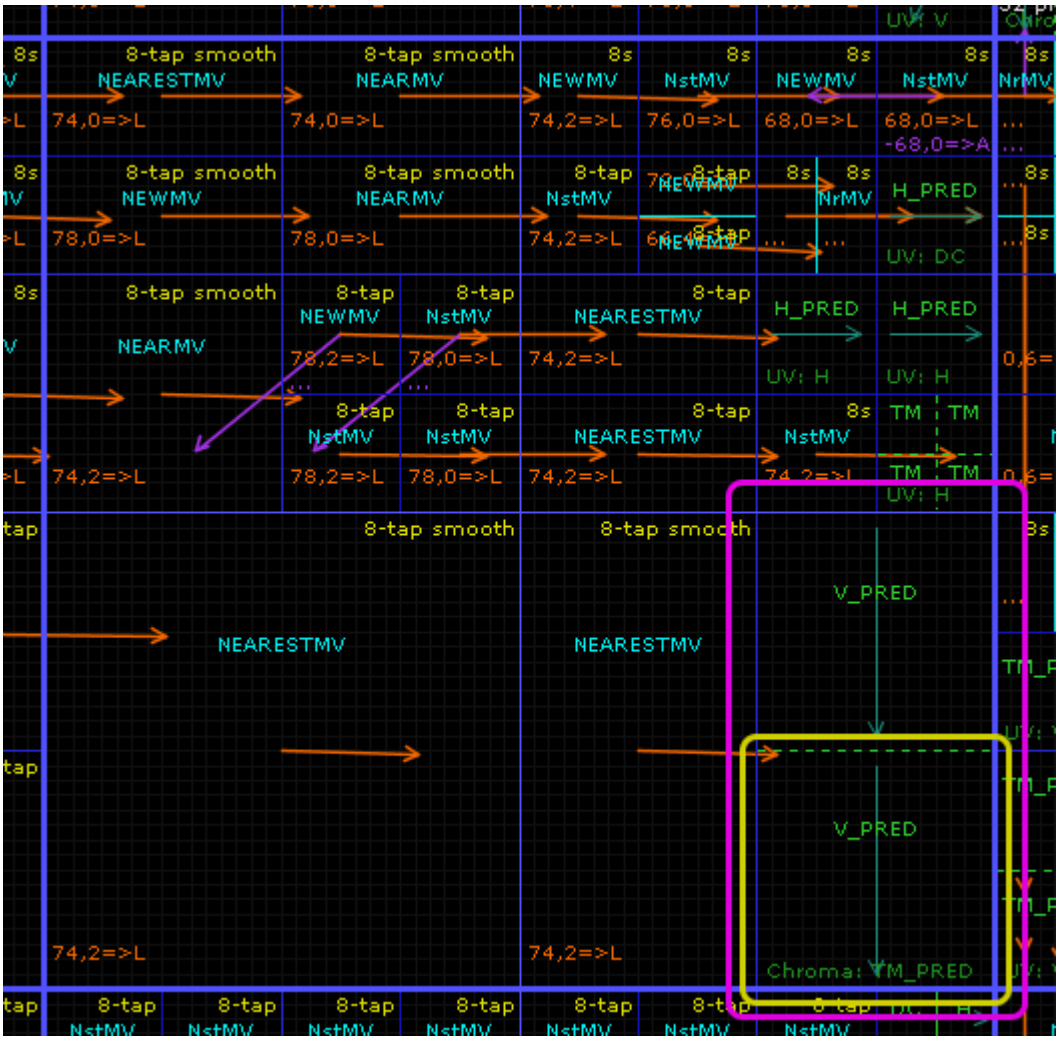

This mode shows the details of the prediction blocks in the frame. Block splits are indicated with blue lines, the selected block is surrounded with a pink box and the selected prediction block is surrounded with a yellow box. When a prediction block is selected, the syntax used to code it is displayed in the Block tab of the Left Panel. Clicking a block repeatedly will navigate the recursive block hierarchy.

Intra prediction block contain the intra luma mode in bright green, and the chroma intra mode in dark green near the bottom of the block. When the intra prediction operation is split due the size of the transform blocks, the intra split boundaries are indicated with dashed green lines. Directional intra luma modes also contain a turquoise arrow indicating the prediction direction.

Inter prediction blocks show the inter mode in cyan, the first motion vector in orange and the second motion vector in purple. The motion vector value is in units of 1/8th pixels. After the motion vector value the reference frame it points to is indicated with a single letter: L for Last, G for Golden and A for AltRef. The motion compensation filter type that applies to the prediction block is indicated in yellow near the top right corner of the block.

4.2.2.1 Prediction Detail Mode

Entering detail mode on a selected prediction block allows the sample values of the prediction process to be viewed directly.

Intra-predicted blocks display the intra prediction process of each component. The left, above-left and above reference sample arrays are shown as well.

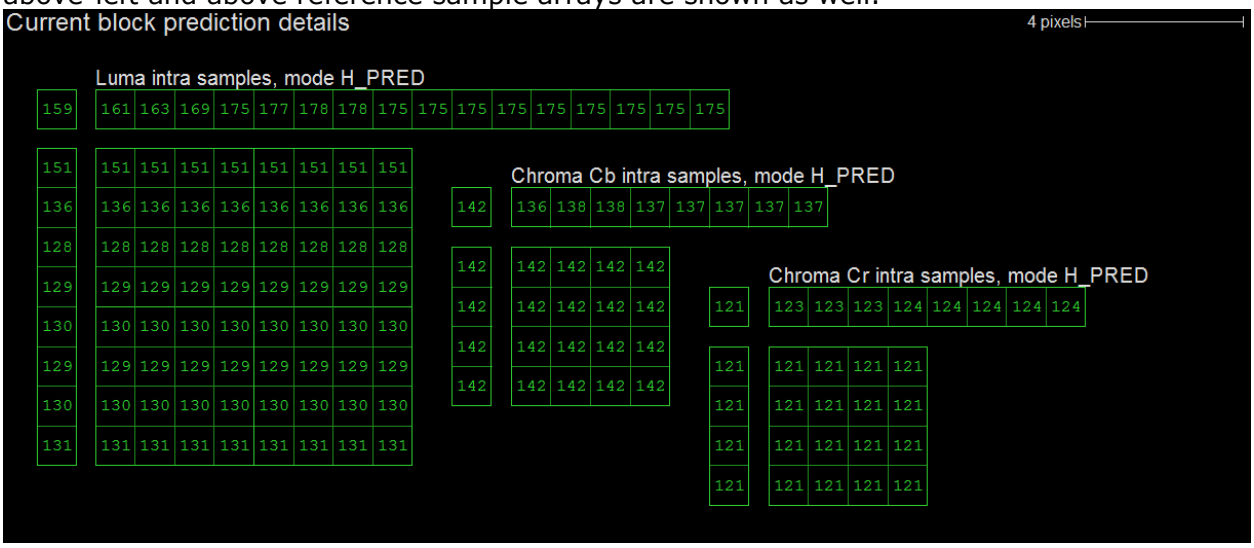

Inter-predicted blocks show the filtered prediction samples arrays of each component, the predictor list used for motion vector prediction (for each motion vector) and the filter type applied. To the left of the prediction sample arrays the impulse response of the applied filter is drawn (mostly decoratively). Each prediction sample can be clicked to display a window that contains the 8x8 or 4x4 sample array for luma or chroma respectively that are the inputs to the motion compensation filter for the selected pixel. The applicable filter phases are also indicated. In the case of compound prediction, both arrays & phase sets are shown with the first/second colors of orange/purple. Clicking the selected sample again or outside the block closes this window.

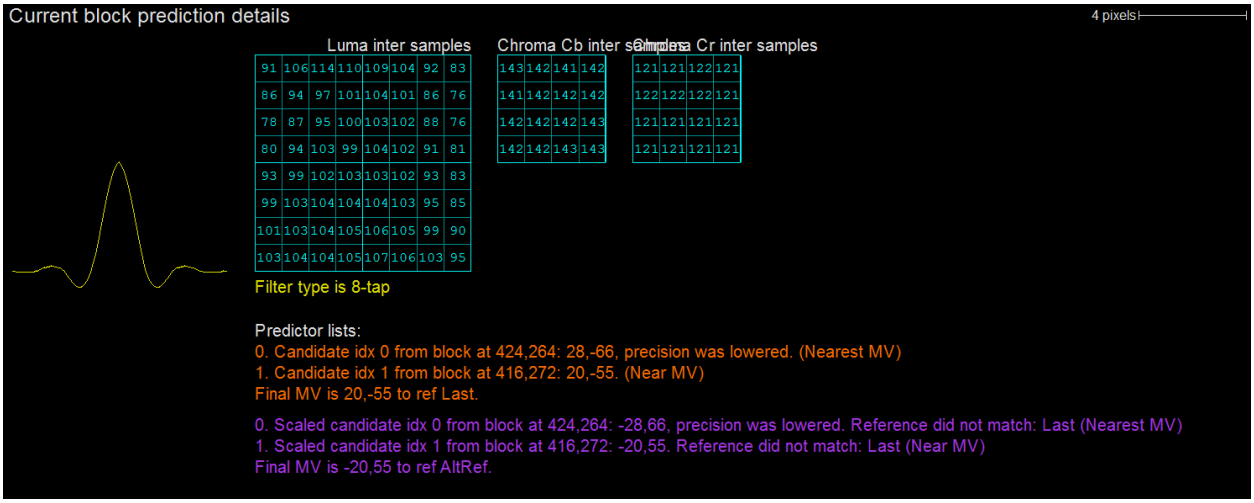

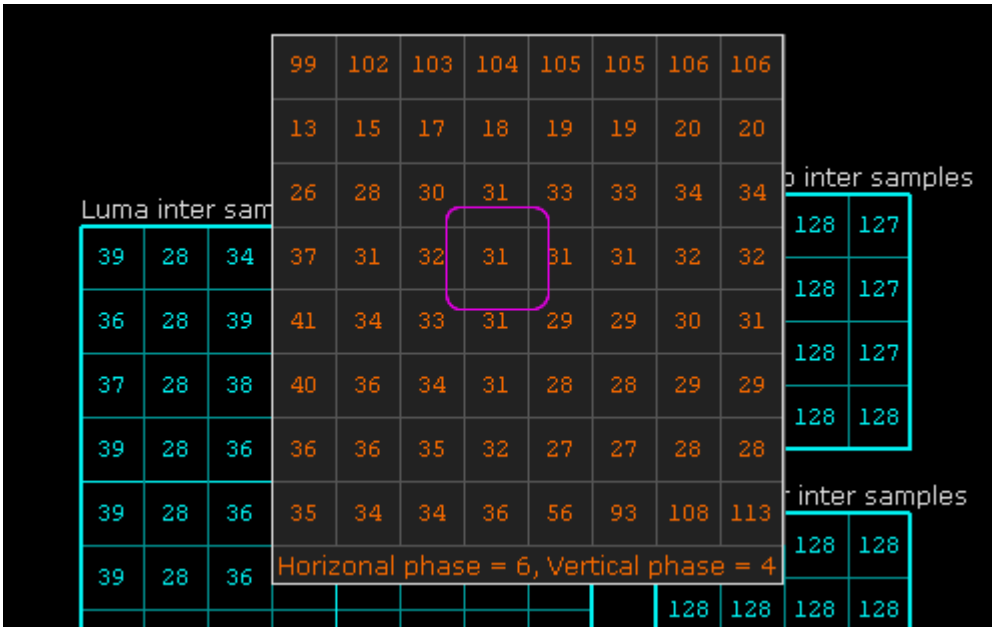

4.2.3 Residuals

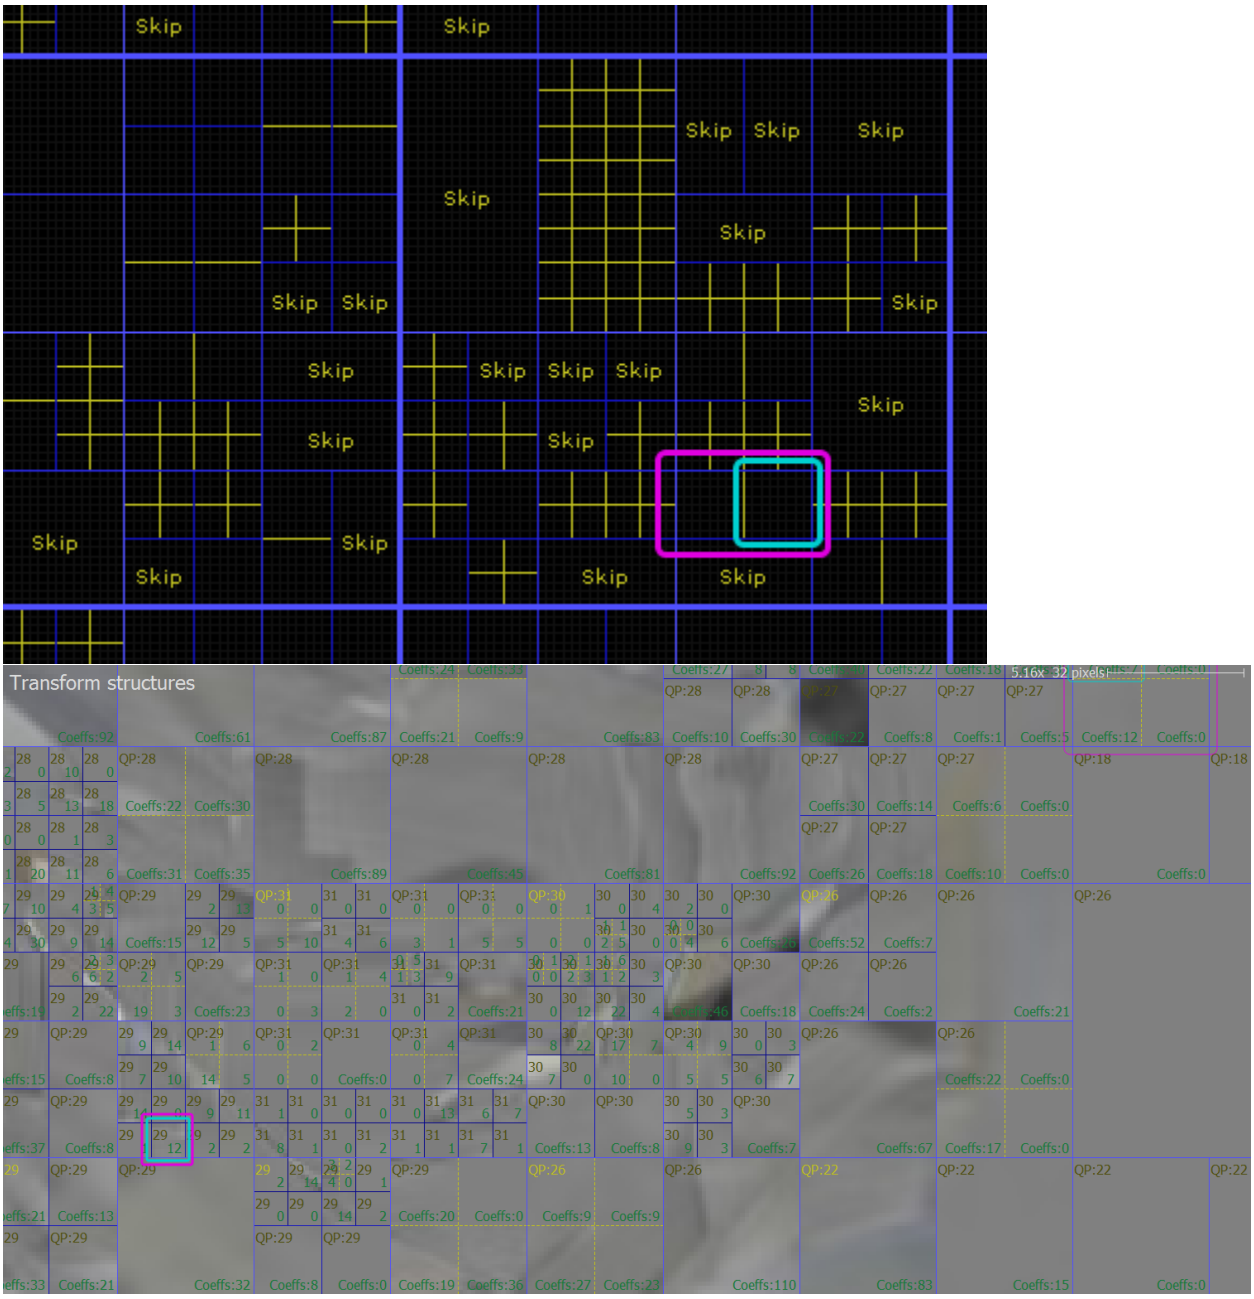

Residual mode allows inspection of the transform structure of each block. Block boundaries are displayed in blue, the currently selected block is outlined with a pink box and the selected transform block is outlined with a blue box. Blocks that code no residual signal are indicated with "Skip" in yellow. Blocks that use a transform size that is smaller than the block's own size show the transform boundaries within the block in yellow.

Clicking a particular transform block causes the syntax elements used to decode it to be visible in the "TX" tab of the Left Panel.

The residual image itself can be shown using the "Pic" button on the bottom button strip. Residual values of 0 are flat gray, negative values are darker and positive values are brighter.

Intel Video Pro Analyzer shows the number of non-zero coefficients with "Coeffs: 2" label on the corresponding block.

4.2.3.1 Residual Detail Mode

To view the full details of the transform blocks belonging to a particular block, selected a transform block and enter detail mode by right clicking or using the detail mode button at the bottom of the main panel. The currently selected transform block is highlighted with the sample grid. The transform structure of the selected block is drawn three times, arranged in a column from top to bottom, showing the 3 major steps in recovering the residual signal:

1. Decode and inverse scan process. In this diagram the scan pattern is drawn for each transform block. In an attempted to reduce clutter, larger scan pattern jumps are drawn with a darker color. The currently selected transform block is highlighted with the sample grid. Moving the mouse over a particular coefficient will highlight the previous and next coefficient in scan order.

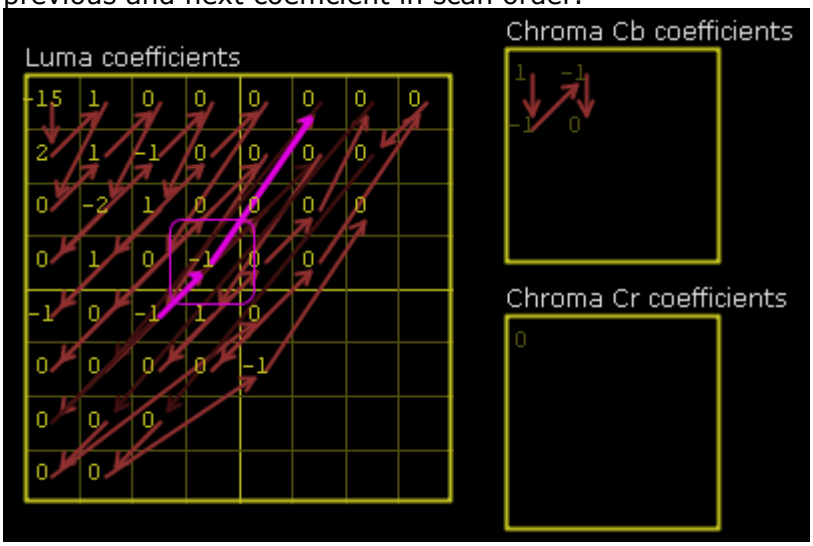

2. Inverse scaling process. The scaled values and the scaling factors used to obtain them are shown in this diagram.

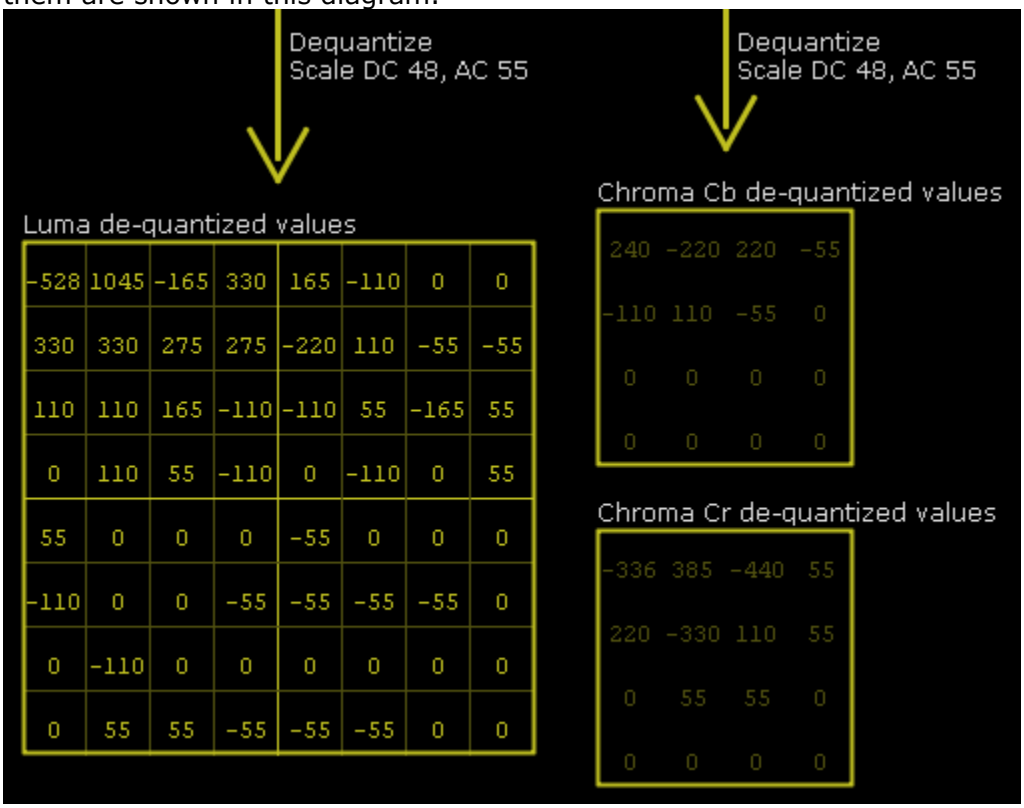

3. Inverse transform process. This recovered residual values and the transform type used are displayed in this diagram.

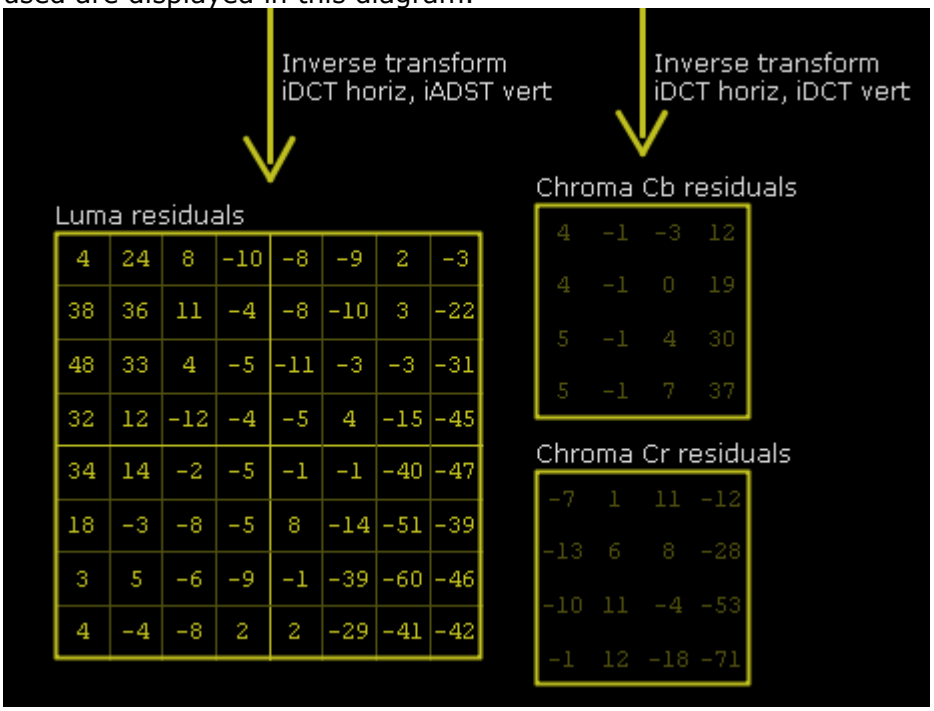

4.2.4 Reconstruction

In Reconstruction mode the reconstructed samples prior to loop filtering can be inspected. As with Prediction/Residual mode, block boundaries are shown with blue lines. Blocks with no coefficients are marked with "Skip" in white.

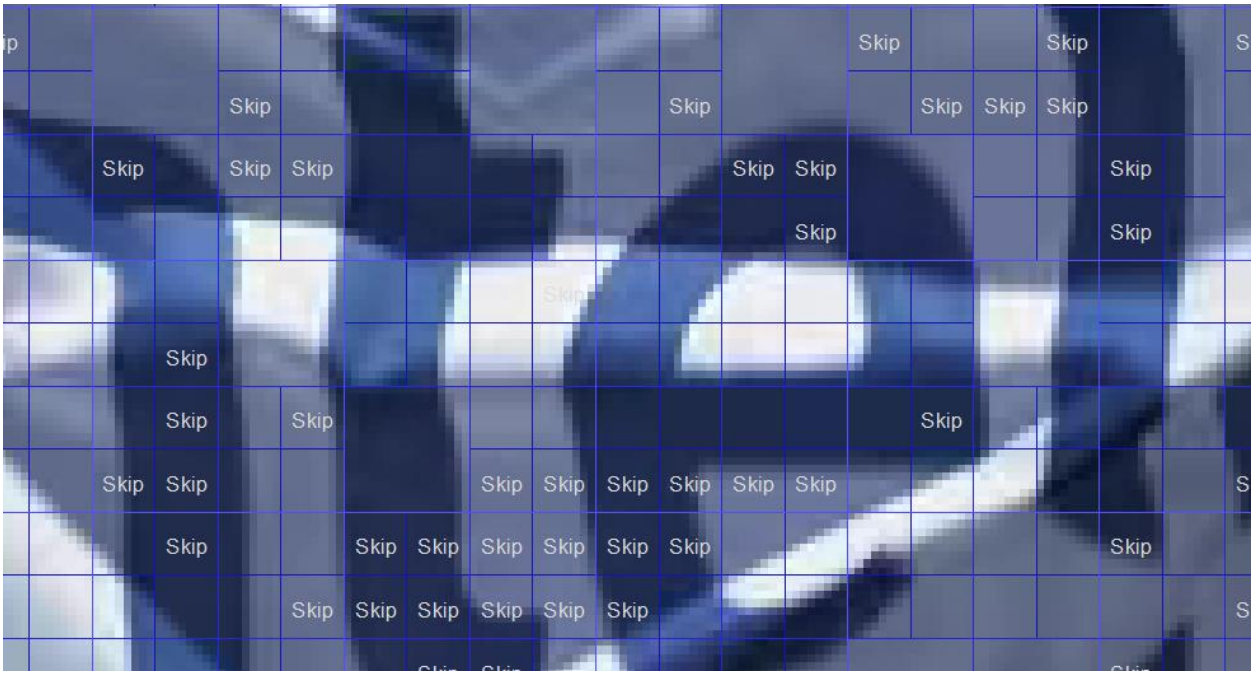

4.2.4.1.1 Reconstruction Detail Mode

In detail mode, the reconstructed sample values can be inspected. The selection may be a single block, or a group of blocks up to and including the entire encompassing super-block.

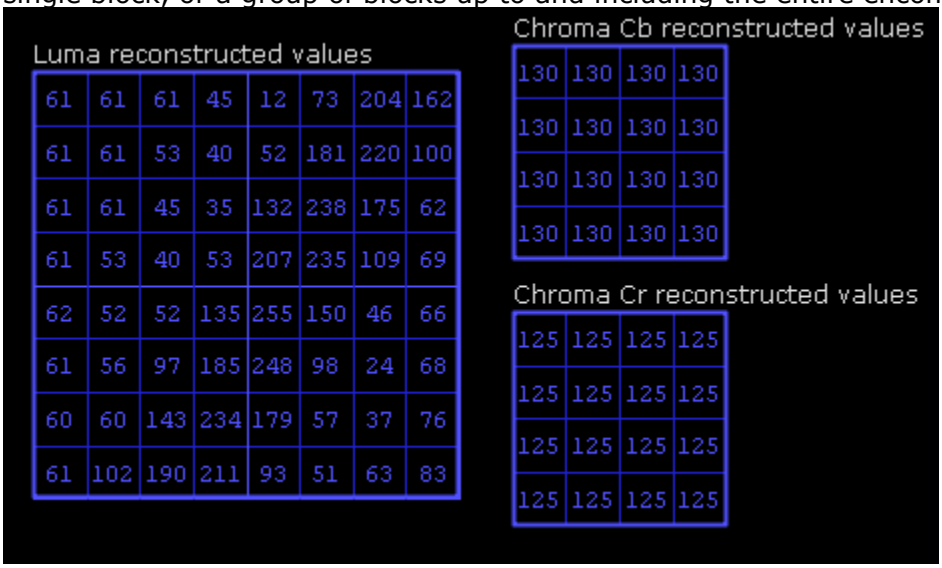

4.2.5 Loop Filter

Loop filter mode shows all processed edges in the frame. Edges are 1 sample high/tall, and as color coded as follows:

- Green – 16-tap filter
- Yellow – 8-tap filter
- Orange – 4-tap filter
- Red – Not filtered, threshold not met (but processed).

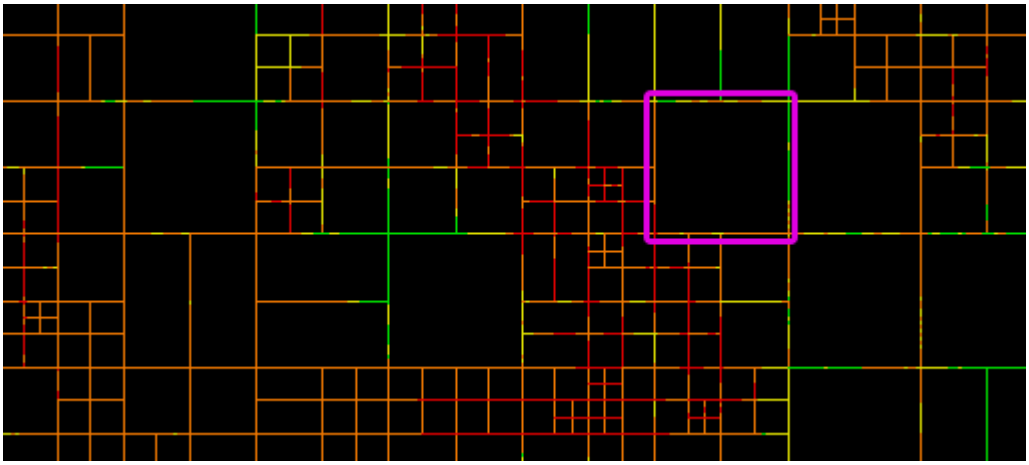

4.2.5.1 Detail Loop Filter Mode

The filtered sample values can be inspected directly in detail mode. As with Reconstruction Details, the inspected area can consist of one or more blocks. Sample values are indicated in yellow if they were modified by the deblocking process. Samples to the left and above the current block need for the filter are shown as well. Filtered edges belonging to the selected blocks are shown in thick dashed lines, color coded the same way as in non-detail mode. Each edge can be selected and clicked, bringing up a window showing the input values to the edge filter. Note that these values may differ from the reconstructed samples, as they may have been modified by an earlier edge filter.

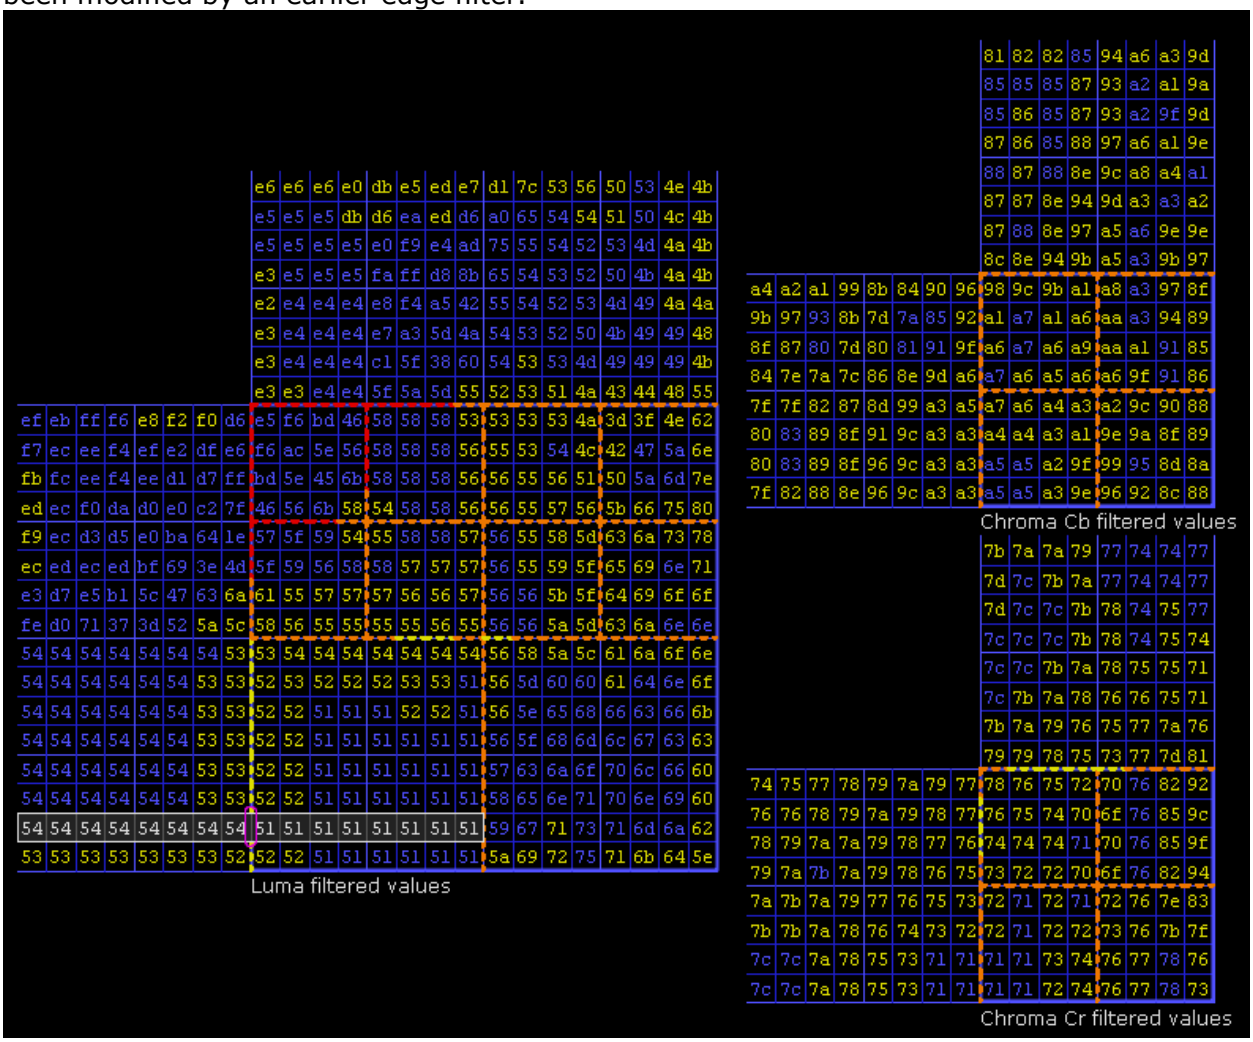

4.2.6 YUV

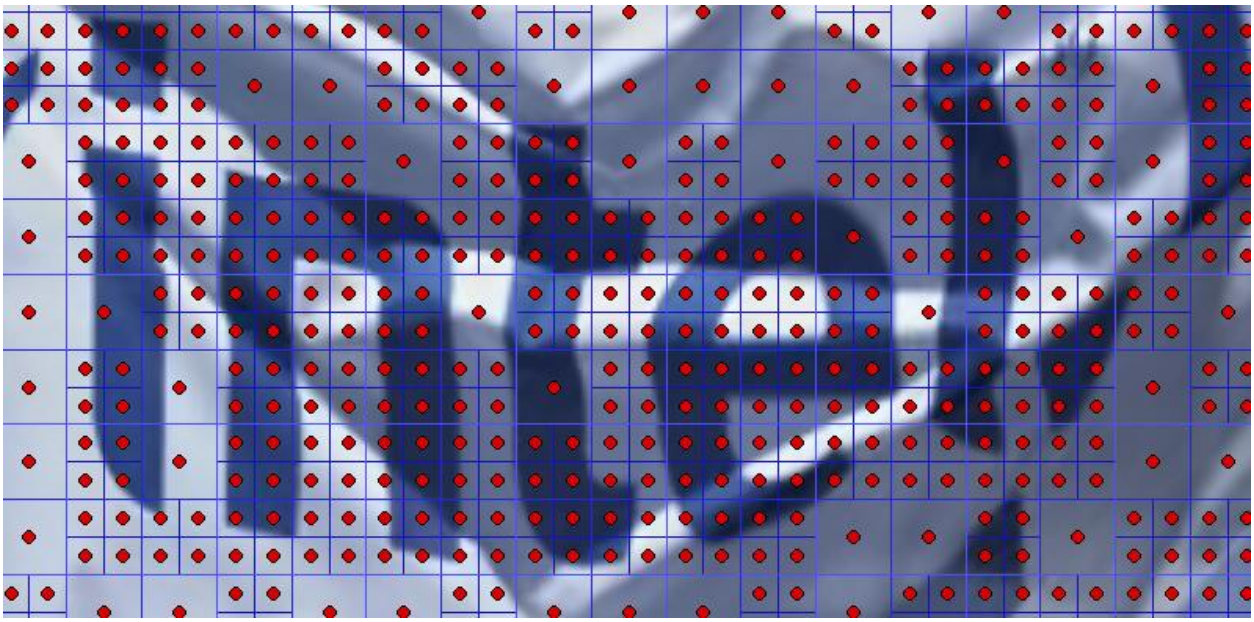

YUV mode allows inspection of the final decoded sample values, without additional overlay data. In this mode, the YUVDiff feature is exposed. This feature allows for comparison of an external decoded YUV file with the decoded bitstream. Supported formats are 4:2:0 planar and NV12. Also, YUV files that are zipped or gzipped may be opened directly without the need to decompress separately. If a zip file contains more than one YUV file, only the first one is used. When a YUV file is loaded, any mismatches will be indicated with a red dot in the CU containing the mismatch. This allows the user to quickly identify the nature of the mismatch which can assist with debug.

When a file is opened, four additional buttons become visible on the lower left of the main window:

- Ref YUV: The original, expected YUV image.
- Debug YUV: The loaded debug YUV file.
- Diff YUV: The delta image. Like in Residual mode, areas with zero delta (i.e. Original and Debug image are identical) are flat gray. Areas where the debug YUV has a lower value are darker, and areas where the debug YUV has a higher value are brighter.
- Reload: Reloads the YUV file, which can be useful as a shortcut to loading via the YUVDiff menu.

4.2.6.1 YUVDiff Menu

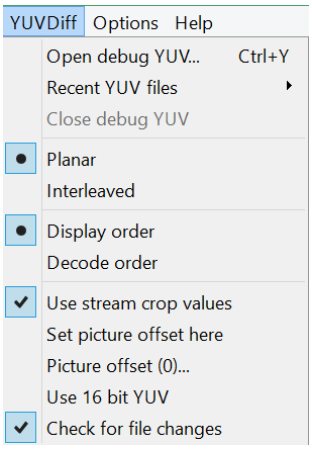

- Open: Brings up a file chooser to open a YUV file
- Close: Closes the currently opened YUV file. Any mismatch indicators will disappear.
- Planar: Causes Intel Video Pro Analyzer to interpret the loaded YUV file as Planar.
- NV12: Causes Intel Video Pro Analyzer to interpret the loaded YUV file as NV12. This option is chosen automatically when the loaded YUV file has the file extension ".nv12".
- Display order: Causes Intel Video Pro Analyzer to use display order picture numbers when determining how far to seek into the YUV file to extract the image to compare with the current picture.
- Decode order: Causes Intel Video Pro Analyzer to use decode order picture numbers when determining how far to seek into the YUV file to extract the image to compare with the current picture.
- Use stream crop values: When checked, YUV files are assumed to contain samples only within the cropping window as defined by the loaded bitstream. Samples outside this window are not compared, and are assumed 0 when viewing the loaded YUV image directly.
- Set picture offset here: Shortcut for setting the picture offset to the current picture number. See bullet below.
- Picture offset: This brings up a dialog allowing the user to enter the picture number of the first picture in the YUV file. For example if a 100-picture bitstream is loaded but the YUV file only contains pictures 80-99, the user would enter 80 to properly line up the YUV file with the decoded bitstream.
- Use 16 bit YUV: this option allow you to load YUV file where both Luma and Chrome written in 16 bit (they could be normalized from actual bitdepth to 16 bit during encoding). This is useful when you have bitdepth changing across the stream. Intel Video Pro Analyzer will shift values to appropriate bitdepth on the current frame.
- Check for file changes: When checked, this option causes Intel Video Pro Analyzer to periodically check if the loaded YUV file has changed on disk since it was last loaded. If the file has indeed changed, a dialog pops up offering a chance to reload the YUV file.

4.2.6.2 YUV Detail Mode

In detail mode the expected YUV values can be inspected. Note that the displayed sample values are the same as those in Loop Filter mode since loop filtering is the last step in the decode process. When a YUV file is loaded, the debug and delta YUV values can be inspected as well. Mismatching values are shown in red.

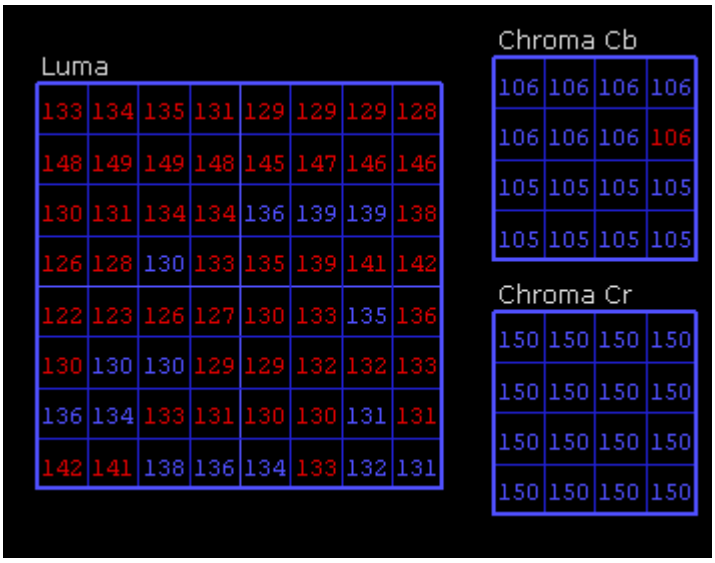

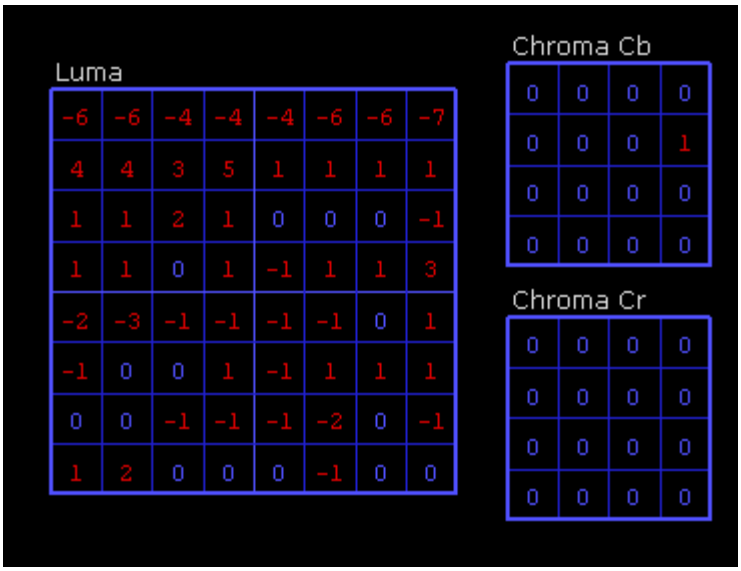

4.2.7 Info Overlays

4.2.7.1 Heat Map

Heat map mode shows visually how the compressed bits of the picture are distributed spatially. Blocks with more bits per pixel are brighter than blocks with less. By default this mode also shows the block boundaries in blue. It may be useful to turn off this overlay using the "Info" button on the lower left of the main panel. Heat Map does not offer a detail mode.

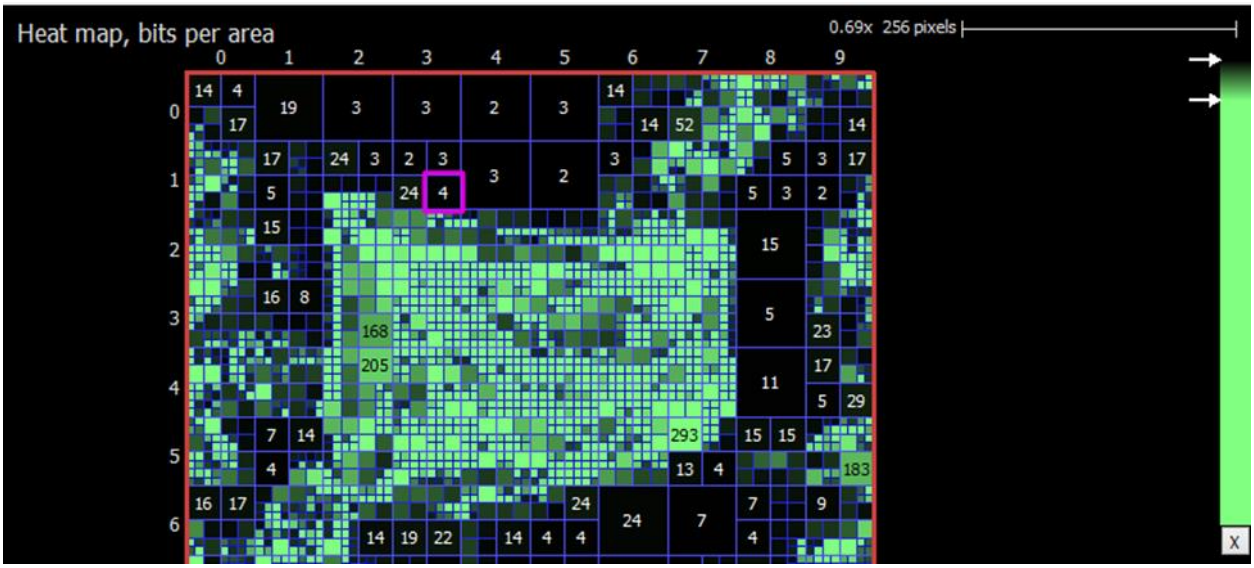

The gradient used to display the heat map may be edited:

- Drag the gradient markers up and down to reposition them.
- Drag a marker away from the gradient to remove it. A red X indicates removal will take place.
- Double-click a marker to change the color.
- Double click the gradient bar to add a new marker.

4.2.7.2 Block Type

Block Type overlay show color map of their types. Intra blocks is colored in red, inter skip is yellow colored and inter is in blue.

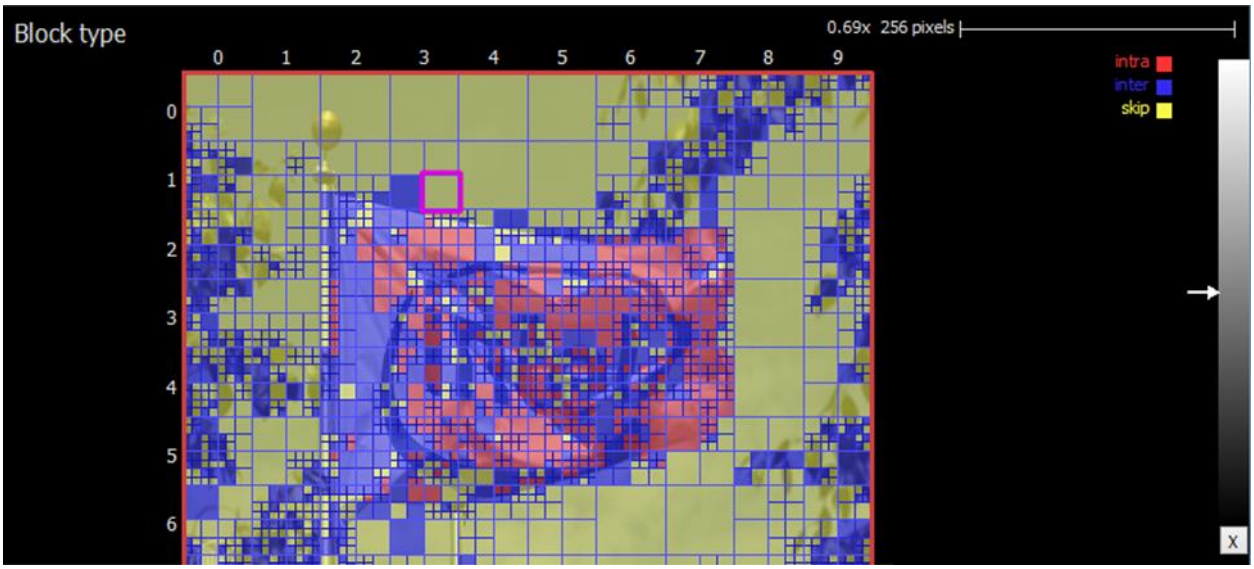

4.2.7.3 PSNR

PSNR map shows per pixel PSNR value normalized on Block size. To enable this mode you have to load reference YUV file via DebugYUV menu. The value is calculated per Block. Bigger value has brighter color. The color range can be adjust with gradient tool on the right top corner. If the block has no difference with reference YUV a latter 'e' is shown which means exact match.

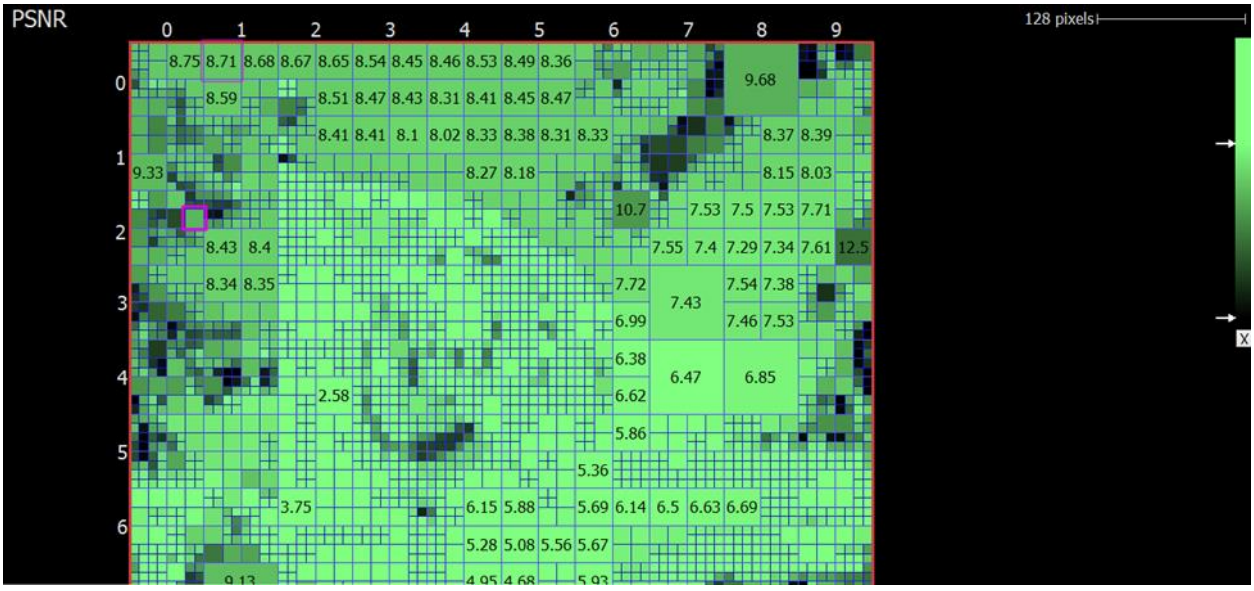

4.2.8 Efficiency Map

Efficiency map mode shows visually how efficient the arithmetic coding of each block is. Blocks with more bools per bits are brighter than blocks with less. By default this mode also shows the block boundaries in blue. It may be useful to turn off this overlay using the "Info" button on the lower left of the main panel. Efficiency Map does not offer a detail mode.

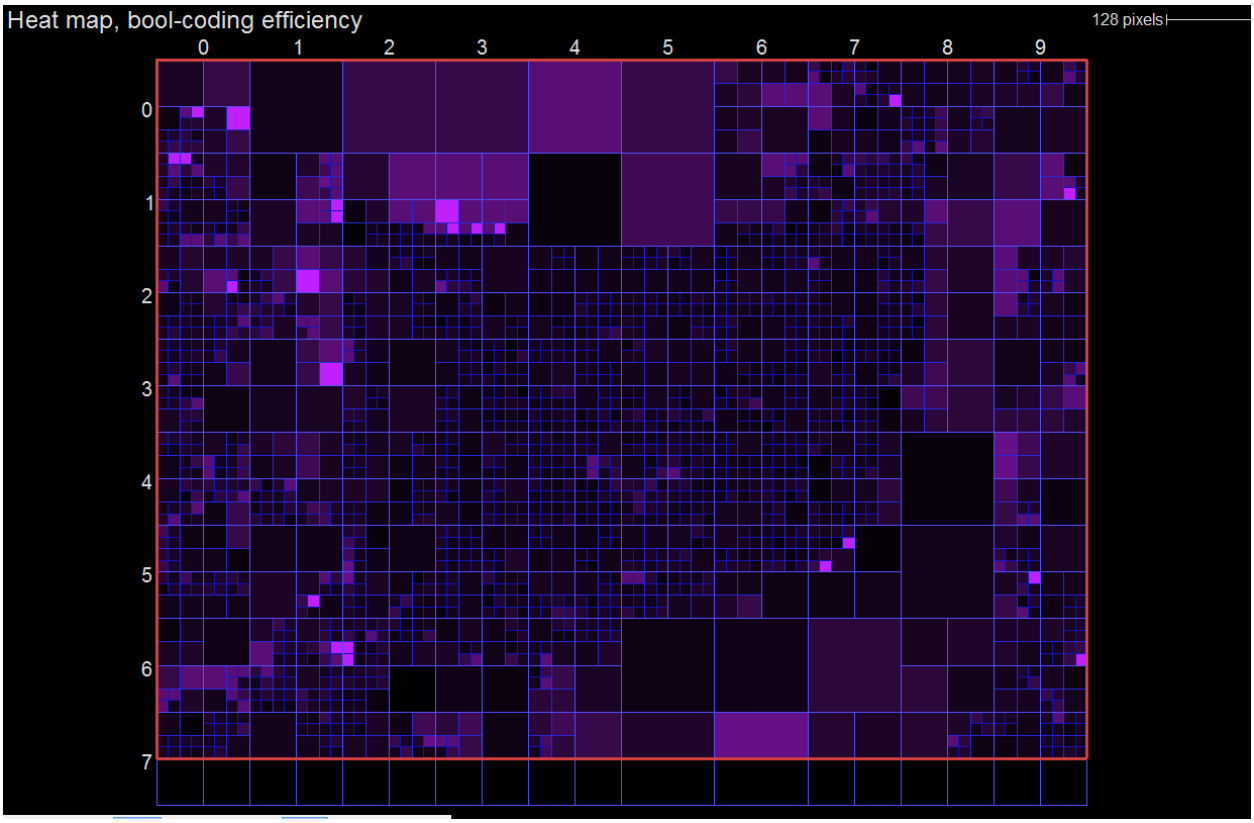

4.2.9 Simple Motion

Simple Motion mode offers a way to quickly view the modes and motion of a picture at a glance. No overlays are present except a colored dot showing information about the prediction mode. Green dots indicate intra blocks, and purple/orange motion vectors for inter blocks are drawn as simple lines. When the motion vector is very small, a dot is drawn instead so that the mode is still easy to identify.

5 AVC

The following sections describe all available features when loading an AVC bitstream. The supported format of a bitstream is the raw bitstream with no surrounding container. Output from the publicly available AVC reference software JM is in this format and can be opened directly. Intel Video Pro Analyzer supports profiles up to High profile. Sequences that go beyond the profile limits maybe be supported as well.

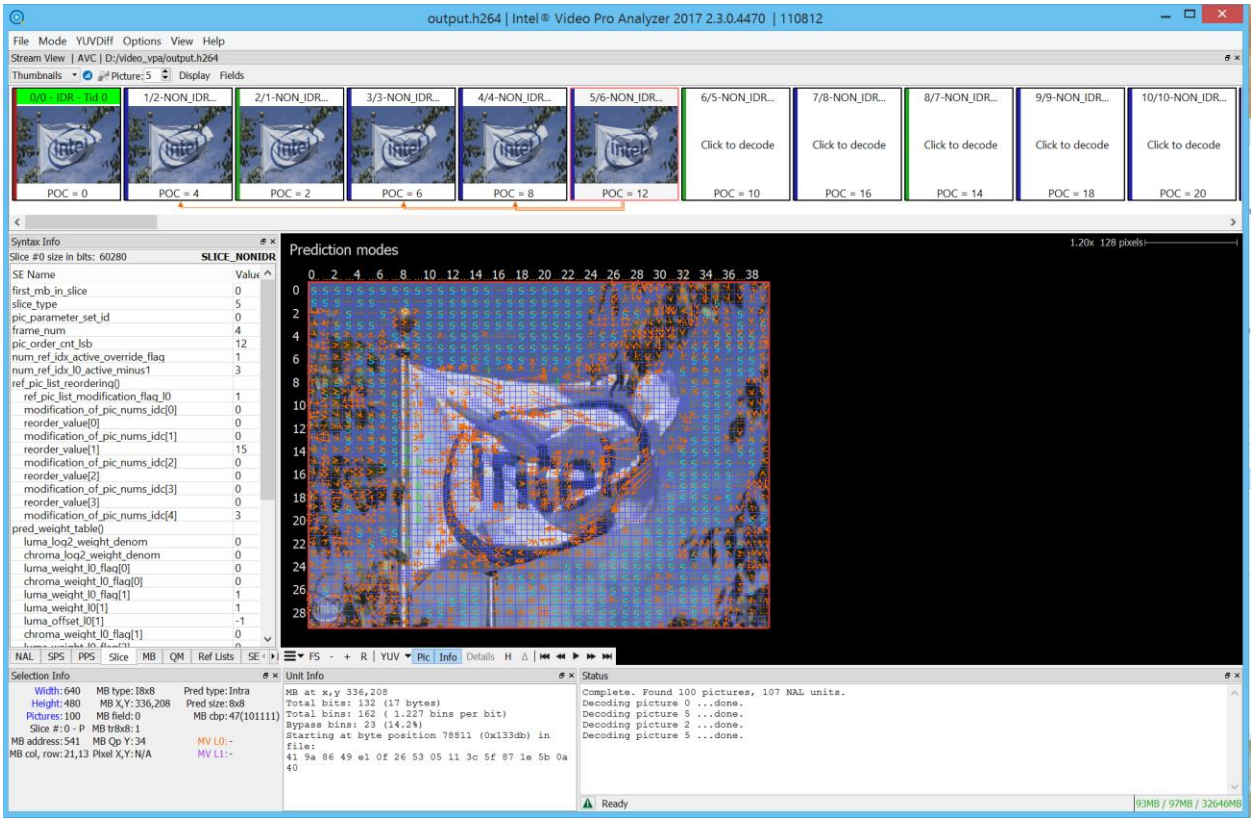

5.1 UI Components

The following sections describe the various components in the UI when an VC bitstream is loaded.

5.1.1 Top Filmstrip

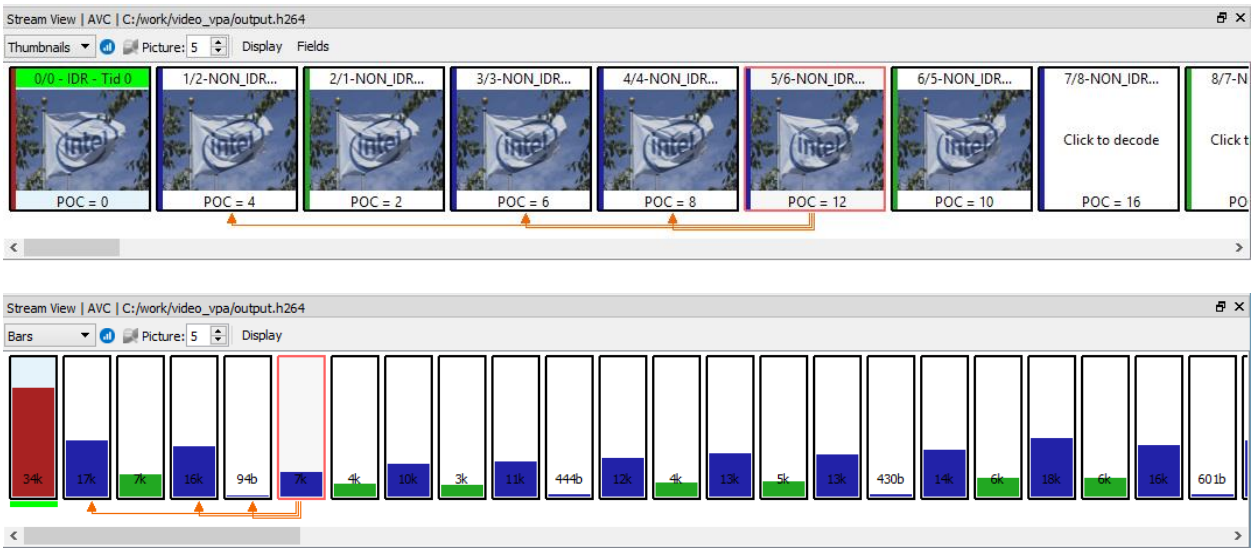

The top filmstrip is a horizontally scrolling overview of the pictures in the bitstream. When a new IDR starts new GOP, the top info label is highlighted in green. The current picture is highlighted with a red border. On top is the picture number - the decode order index in the bitstream - and the first slice in the picture type as indicated by the NAL unit types, and on bottom is the POC value. Referenced pictures for the current picture are indicated with arrows. Green arrows indicate the long term references. The bar filmstrip has underline green band on new GOP start. Picture type is core coded: red – intra picture (I), blue – P-picture, green – B-picture.

Clicking on a picture will cause Intel Video Pro Analyzer to decode it and mark it as the current picture, updating the rest of the UI correspondingly. If the new picture requires other pictures to be available as reference that have not yet been decoded, those will automatically be decoded.

Pictures are decoded on demand since it is impractical to store all details of all pictures in memory, especially for HD sequences. The current picture can also be selected by typing in a picture number in the box on the left of the panel. When using this box, the filmstrip is scrolled such that the newly selected current picture is visible.

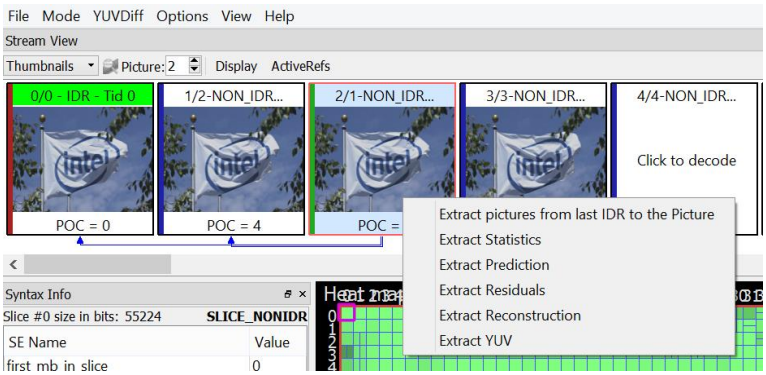

In Thumbnails and Bar mode it is possible to choose between Display/Decode order representation by click on additional. By right clicking on selected picture one can get popup submenu with a set of extract actions. An action called "Extract pictures from last IDR to the Picture" will write out a bitstream that contains the minimum number of pictures needed to decode the current picture. Typically this will consist of the current picture and all previous pictures that are reference pictures up to the nearest IDR picture. This can be useful for debugging issues in long sequences. Other actions allow to save in the output file statistics and intermediate/final YUV planes.

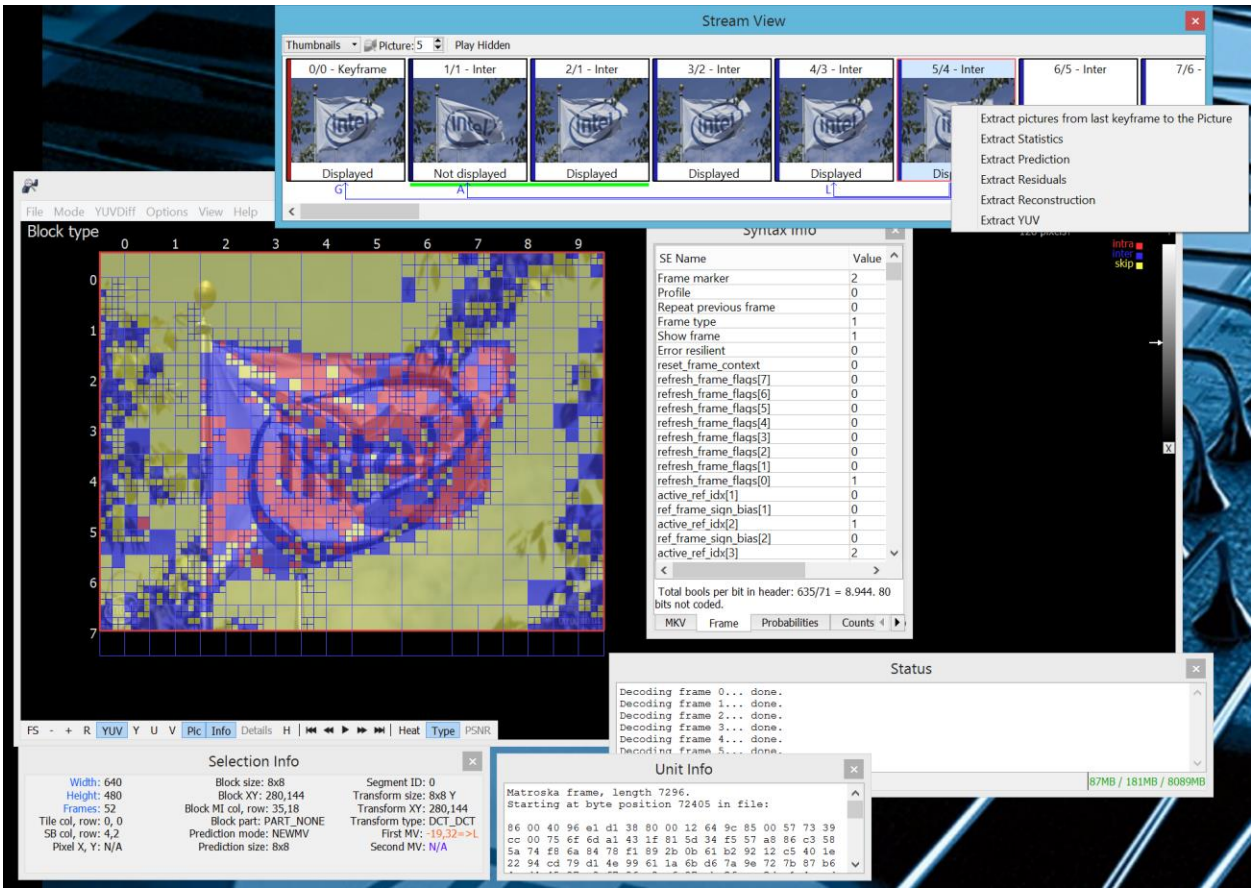

Allow separate panels can be detached from them main window or hidden. To attach panels back double click on the titlebar. To unhide panel go to View menu and check appropriate panel.

5.1.2 Main Panel

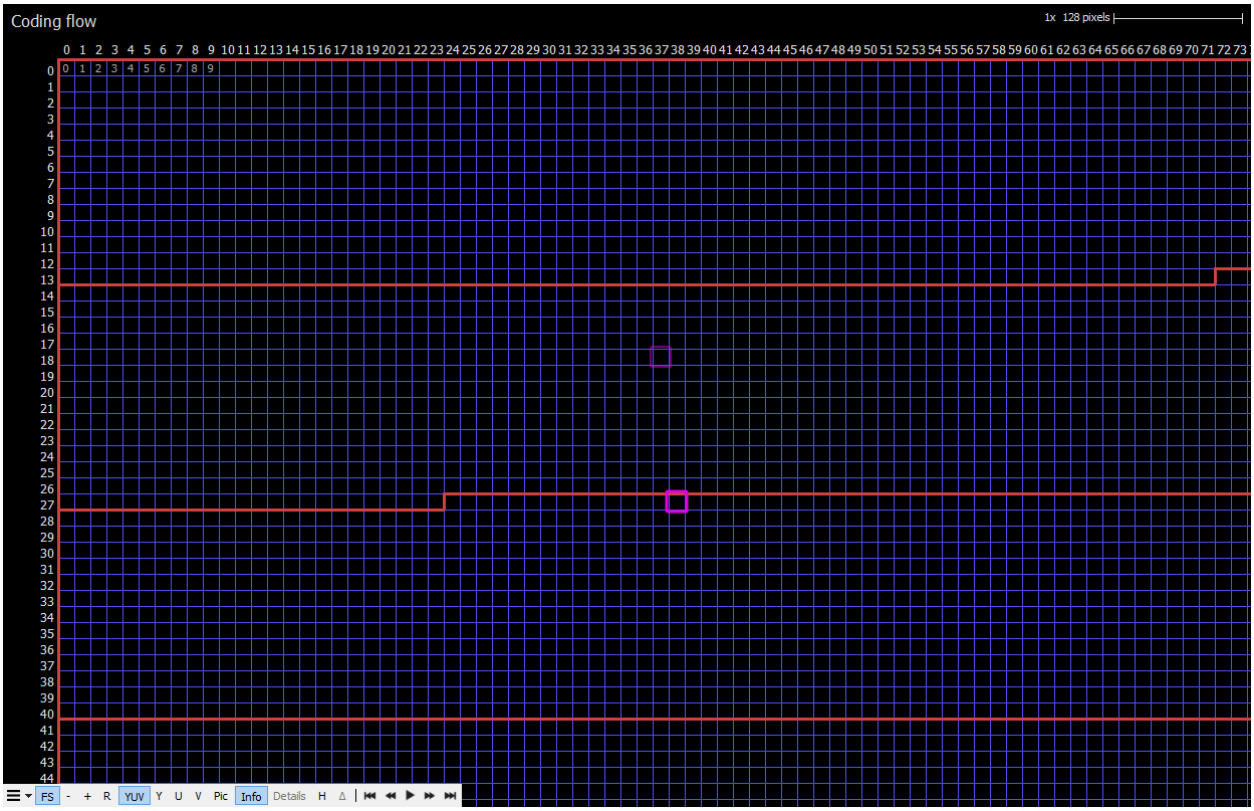

The main panel displays the selected picture with visual annotations. The type of annotations and associated interactive behavior depends on the current mode, see 3.1.8.1 AVC Modes. Click-dragging moves the picture around, and the mouse wheel zooms in or out about the mouse cursor.

At the top left of the panel, the current mode is displayed. Along the top right the current scale is shown. In all modes, the MB row and column values are displayed along the left and top border of the picture, slices are indicated with thick red lines, dependent slice boundaries are dashed red lines and tile boundaries are thick green lines. The screenshot above shows an example of a stream that is 10x8 MBs. The selected MB is outlined with a pink box.

The bottom left has a cluster of buttons:

- First submenu has special options to output the currently displayed information to the external image file with different scaling possibilities.
- FS: Toggles full-screen mode, which hides all UI elements except the main panel. Keyboard shortcut is F.
- +/-: Zooms in or out, centered on the center of the panel. Note that zooming with the mousewheel is much easier than using these buttons. Keyboard shortcuts are + and -.
- R: Resets the current zoom to fit the entire picture in the main panel. Keyboard shortcut is R.
- YUV, Y, U, V: Toggles between per-component image, or full-color YUV.
- Pic: Toggles the actual picture underneath the annotations. The picture is mode-dependent. For example in prediction mode, the picture is comprised of the

prediction samples before the residual signal has been added in. It is often helpful to turn off the picture in order to make the annotation easier to read. Keyboard shortcut is P.

- Info: Toggles the annotations on or off. Keyboard shortcut is I.
- H/D: Toggles between hexadecimal and decimal display of the annotations. This applies to the values in the Left Tabs as well. Keyboard shortcuts are H and D.
- There are five navigation buttons: first frame, previous frame, playback mode, next frame, last frame. The previous/last button will rewind, causing the current picture to be set to the first/last picture in the loaded bitstream. The "Play" button will start playback of loaded bitstream at a reasonable rate, limited by CPU performance. All normal mode operations are possible during playback. Note that playback mode always operates in currently selected order in the filmstrip.

5.1.3 Left Tabs

The left side of the UI is a resizable panel with several tabs that display information about the current selection in one way or another. The following subsections describe each tab.

5.1.3.1 NAL

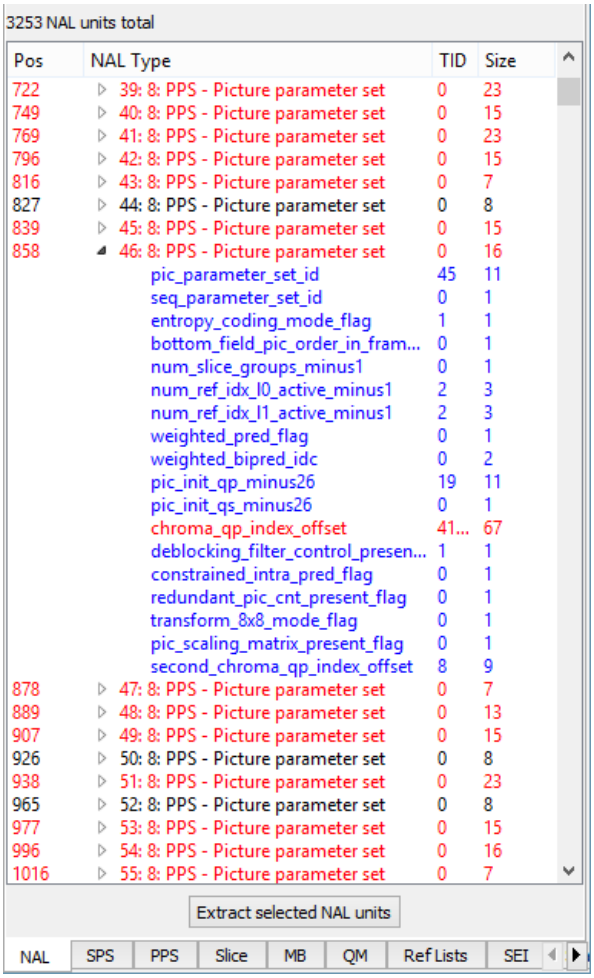

The NAL tab lists all NAL units found in the bitstream, in decode order. Above the list the total number of NAL units is displayed. The list has four columns: "Pos" is the byte position of the NAL unit in the file, "NAL Type" is the textual description of the NAL's type, "TID" is the Temporal ID of the NAL unit, and "Size" indicates the total number of bytes in the NAL unit.

Changes in coloring from white to grayscale and back show Access Unit borders. By clicking on branch sing you can roll out NAL unit syntax elements read from the bitstream. For each syntax element Intel Video Pro Analyzer shows the value of it and the number of bits it takes in the bitstream. Red colored elements will hint you about error syntax element.

Clicking on any NAL unit will display up to the first 500 raw bytes in the Raw Bytes panel. If a NAL unit is a VCL NAL unit, meaning it is a coded slice segment, it can be double-clicked in order to make the picture that it belongs to the current picture. Conversely, changing the current picture will highlight the NAL unit that contains the first slice segment of that picture. Selecting a MB in the Main Panel will highlight the NAL unit containing the slice that contains the selected MB.

At the bottom of the NAL tab is a button called "Extract selected NAL units". Clicking this will write all selected NAL units out to a file.

5.1.3.2 SPS, PPS, Slice, SEI

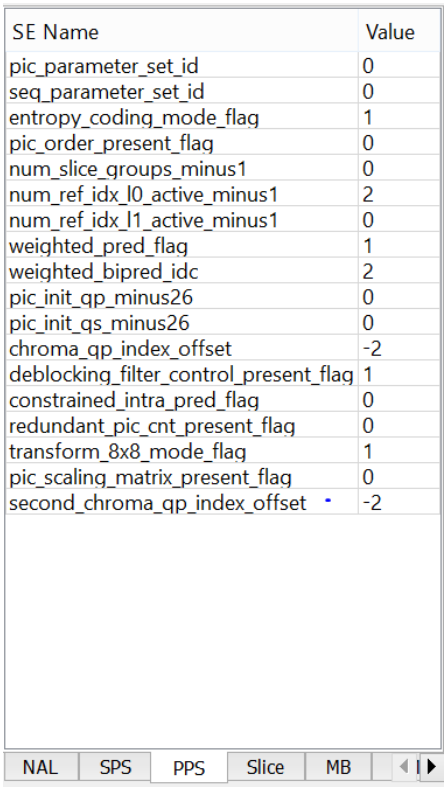

SE Name	Value
pic_parameter_set_id	0
seq_parameter_set_id	0
entropy_coding_mode_flag	1
pic_order_present_flag	0
num_slice_groups_minus1	0
num_ref_idx_l0_active_minus1	2
num_ref_idx_l1_active_minus1	0
weighted_pred_flag	1
weighted_bipred_idc	2
pic_init_qp_minus26	0
pic_init_qs_minus26	0
chroma_qp_index_offset	-2
deblocking_filter_control_present_flag	1
constrained_intra_pred_flag	0
redundant_pic_cnt_present_flag	0
transform_8x8_mode_flag	1
pic_scaling_matrix_present_flag	0
second_chroma_qp_index_offset	-2

NAL | SPS | PPS | Slice | MB

These tabs contain a list of every syntax element, in decode order, of the SPS, PPS, slice segment header and SEI messages that apply to the currently selected MB in the Main Panel. Syntax elements that come from a sub-function call are indented accordingly, and the function call itself appears in the list as a syntax element with no value. Any part of the list can be selected and copied for pasting in other programs.

5.1.3.3 QM

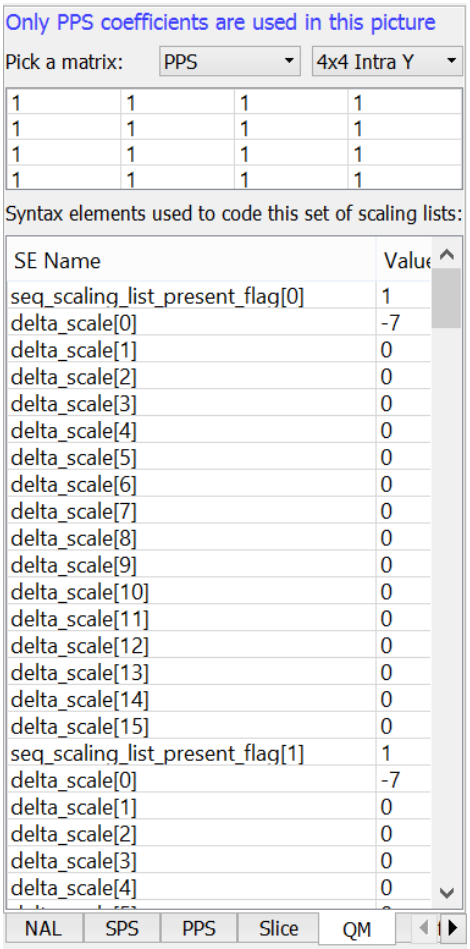

This tab displays the scaling lists, or quantizer matrices used by the current picture. Scaling lists can be present in the SPS, PPS, or both. PPS scaling lists take precedence over SPS when present. The top of the panel shows in blue text how the various scaling lists are used in this picture. Below that, one of the 20 scaling list from either the SPS or PPS can be chosen for inspection, and the chosen matrix is displayed in the grid. Also shown in this panel is the list of syntax elements that is used to code the scaling list set, and is identical to the syntax in the corresponding SPS or PPS tab.

5.1.3.4 Ref Lists

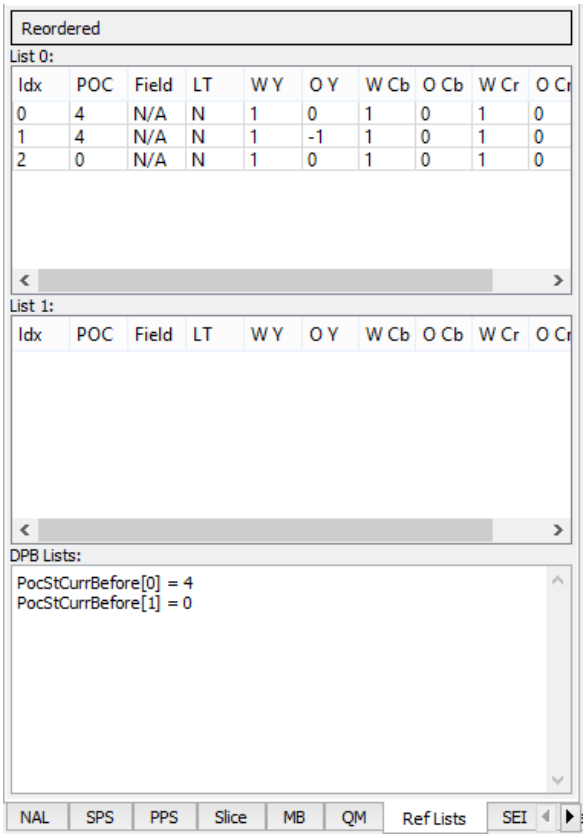

This tab displays the details of the two reference lists L0 and L1, as well as the Reference Picture Set arrays used to construct the L0 and L1 lists. The L0 and L1 lists have the following columns:

- Idx: The index associated with the reference picture
- POC: The Picture Order Count of the reference picture
- LT: A Yes/No flag indicating if the reference picture is a long-term picture or not
- W Y: The derived luma weight used in the weighted prediction process of the reference picture
- Y: The derived luma offset used in the weighted prediction process of the reference picture
- W Cb: The derived chroma Cb weight used in the weighted prediction process of the reference picture
- Cb: The derived chroma Cb offset used in the weighted prediction process of the reference picture
- W Cr: The derived chroma Cr weight used in the weighted prediction process of the reference picture
- Cr: The derived chroma Cr offset used in the weighted prediction process of the reference picture

The third panel contains in textual form the contents of the Reference Picture Set arrays from section 8.3.2 in the HEVC specification.

By clicking on 'Reordered' button you will see reference lists state before reordering.

5.1.3.5 Stats

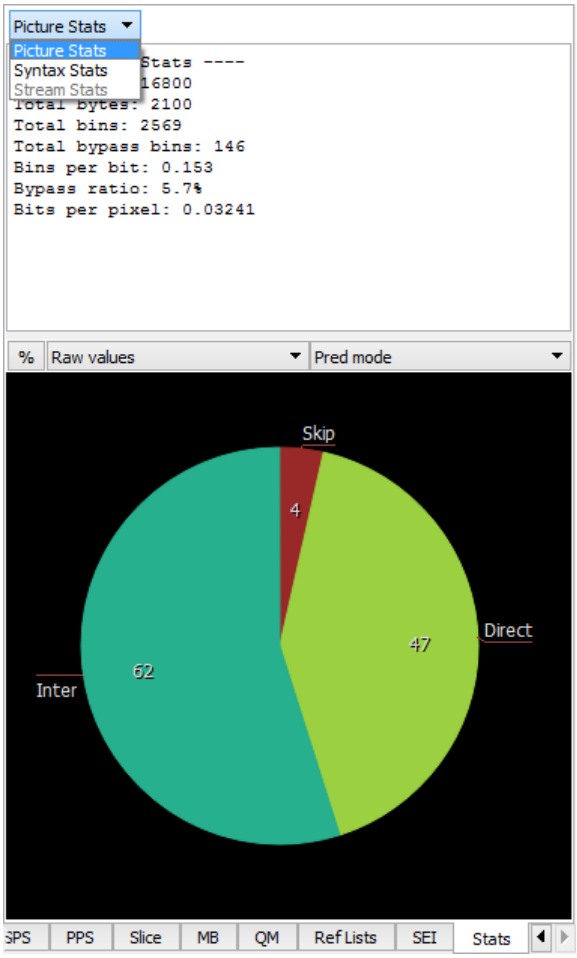

This tab displays various statistics extracted from the current picture. The top half shows some picture size and compression stats. The bottom half displays pie charts for a number of metrics. Each pie chart can be drawn normalized or un-normalized. Normalized data is weighted by area or compressed bits. Normalized numbers in the pie chart are in units of pixels or bits, and un-normalized numbers are raw counts.

The pie chart can be moved by dragging the mouse, and zoomed with the mouse scroll wheel.

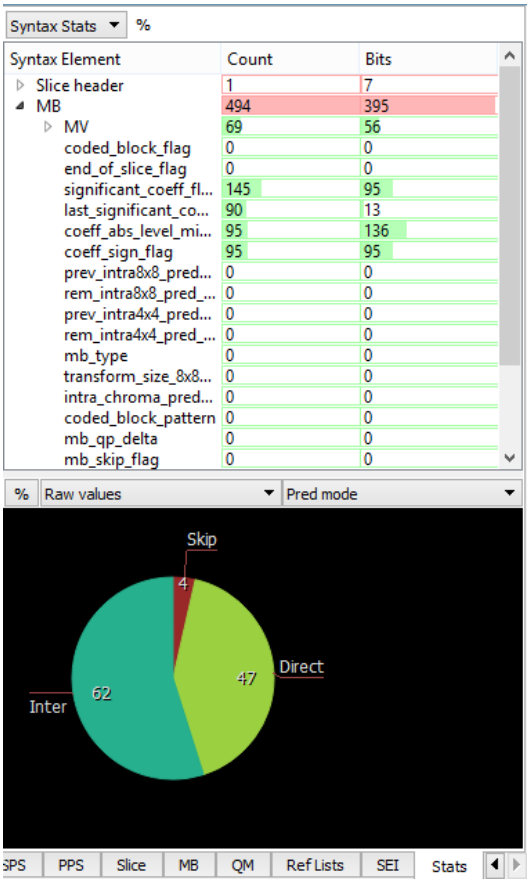

By selecting, different modes in the stats combo box you can explore different type of statistics in the bitstream. Namely, picture statistics, per syntax element statistics, overall stream statistics. Overall stream statistics is available only after full stream analysis and shows accumulated values overall the stream. Pie chart for stream statistics show accumulated value also.

5.1.4 Selection Info

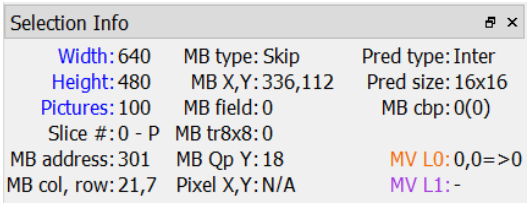

This panel shows a few details about the current selection at a glance. All values are decimal:

- Picture height & width
- Number of pictures in the bitstream
- Slice segment number of the slice segment that contains the selected CU, and the slice type.
- MB address and row/col position
- Selected MB type and flags for field MB coding and use 8x8 transform.
- X/Y position of the selected MB in pixels
  MB prediction mode. Intra, Inter or Skip

- Pixel X/Y position in the picture of the mouse cursor
- Prediction size and shape.
- The L0 and L1 motion vectors, if any, of the selected Prediction block. The MVs are shown with an arrow => followed by the reference index.
- The Luma and Chroma QP values associated with the selected MB.

5.1.5 Raw Bytes

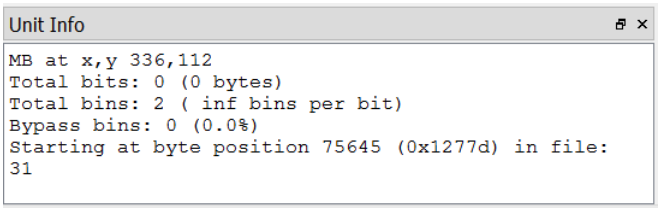

The panel in the bottom center of the UI displays the raw bytes that were used to code the selected MB, or NAL unit. A few details about the selection are included as well.

5.1.6 Messages

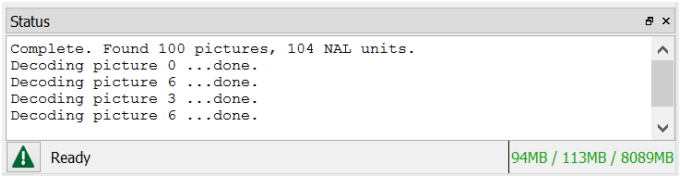

This panel in the lower right of the UI displays messages about the decoding process, and shows progress of any actions that may take a while to complete.

The 3 numbers in the lower right of this panel represent the memory state of JVM. From left to right the 3 numbers are:

- Amount of memory used
- Peak used memory
- Maximum available memory

Error button at the bottom left of the window changes color to red if Intel Video Pro Analyzer detects any errors/warnings in the stream. By clicking on the button Intel Video Pro Analyzer switches to the list of detected problems and back. By double clicking on selected error, you can highlight corresponding MB in central info window and corresponding NAL unit will be highlighted.

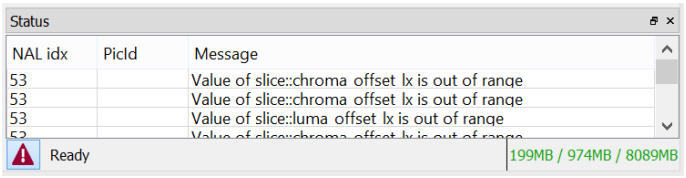

By right clicking inside error's list you can choose additional filtering of the errors

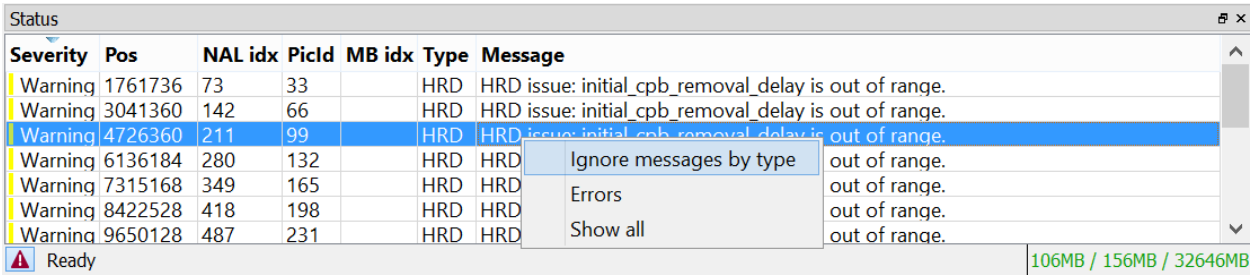

"Error" submenu runs common error dialog where it is possible to select which errors will by displayed in the error's list.

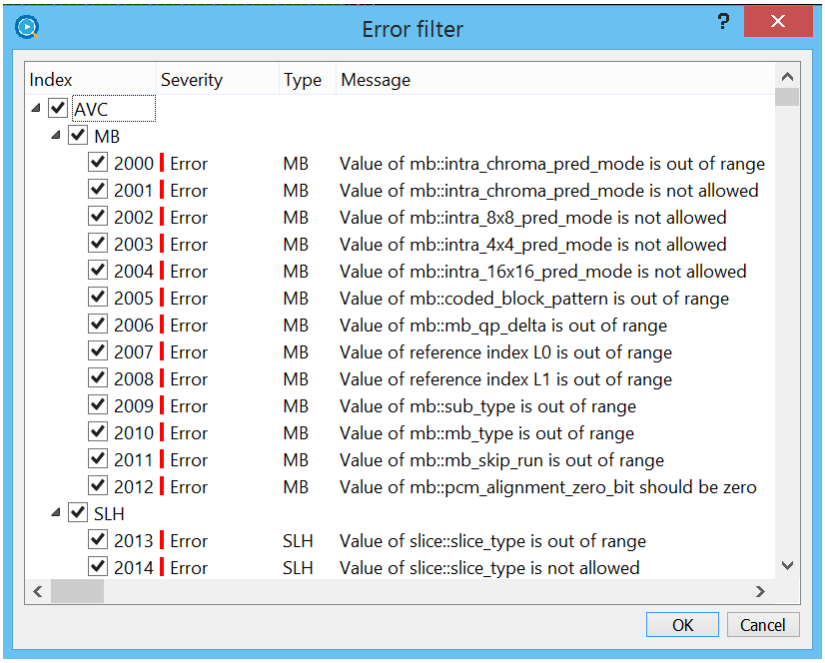

5.1.7 Extended modes

There are a set of extended modes that can be accessed via menus in the Stream View toolbar.

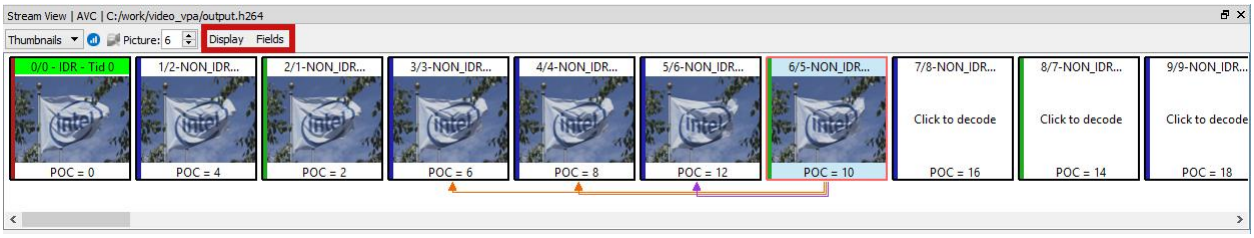

Field option gives special visualization capabilities for stream with fields. Namely, it shows references on per field base. Video Pro Analyzer shows top and bottom fields as the happens in the stream. Their position is shifted according to top or bottom flag is used for the field.

Then two possible modes with two sub-modes can be selected. Sub-modes of the currently selected mode changed dynamically depended on current mode.

5.1.7.1 Frame sizes view

When one choose "Buffer" menu from popup toolbar, frames size view will appear by default. This is a bar char plot with frame sizes. Each bar is colored according to the frame type. When you move the mouse along the bar current frame is highlighted and tooltip shows information about current frame. By clicking on selected frame you can decode it. When you press left mouse button and move it all plots move along both axes. Double click on canvas area resets plot into default state. You can zoom in/out with mouse wheel on both axis when the mouse

inside canvas. If the mouse on some axis (below canvas, for example, x axis) you can zoom, pan or double click on only on that axis. In that case all actions applied to that axis only. The movement of the plots is restricted to the positive values only, you can't move or zoom in negative direction. Currently decoded frame is shown with vertical line and frame number on it.

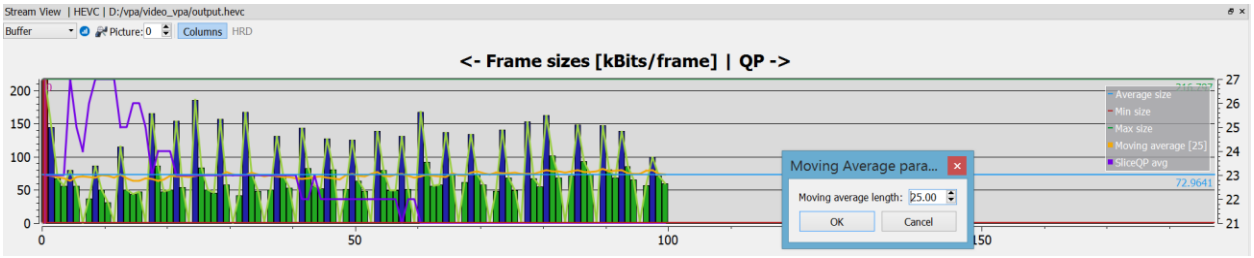

On the right axis average slice QP values are shown. When full stream analysis is complete Intel Video Pro Analyzer adds average block QP plot that is also attached to right axis. By double clicking on 'Moving average' label item sliding window size can be adjusted in the separate dialog.

5.1.7.2 HRD buffer fullness view

If you choose "Buffer" mode then sub-mode "HRD" appears in sub-modes block. It is enabled if stream has HRD parameters in it, otherwise it is disabled (grey out). The interaction with buffer fullness plot works in the same way as for frame sizes plot. It is possible to get buffer state and picture parameters in each point of the plot where the monotonic behavior changes.

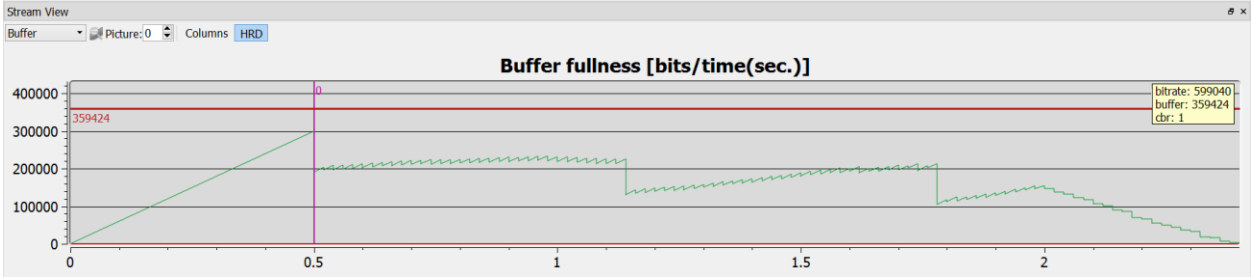

With right click on central area of the plot Intel Video Pro Analyzer shows additional option for plot adjustment. In the same popup menu you can export HRD timings into .csv file.

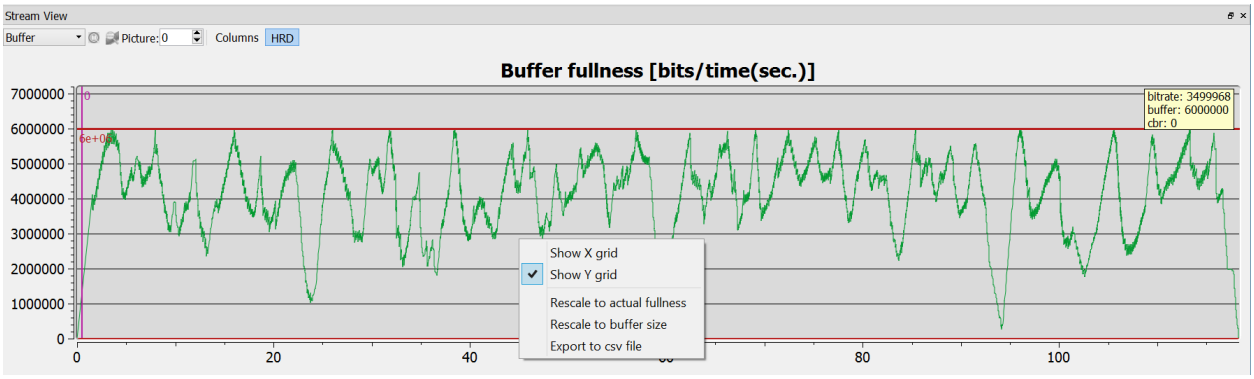

*Other names and brands may be claimed as the property of others.

5.2 Modes

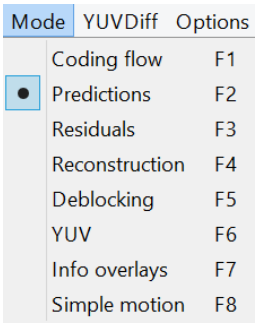

With an AVC bitstream loaded, Intel Video Pro Analyzer can be put into one of 8 modes using either the F1-F8 keys, or using the Mode menu. The mode selection affects only what is displayed in the Main Panel. Most modes can also show details of the current selection (MB,Prediction unit, Transform unit). This can be toggled with the right mouse button or the main panel's button strip.

5.2.1 Coding Flow

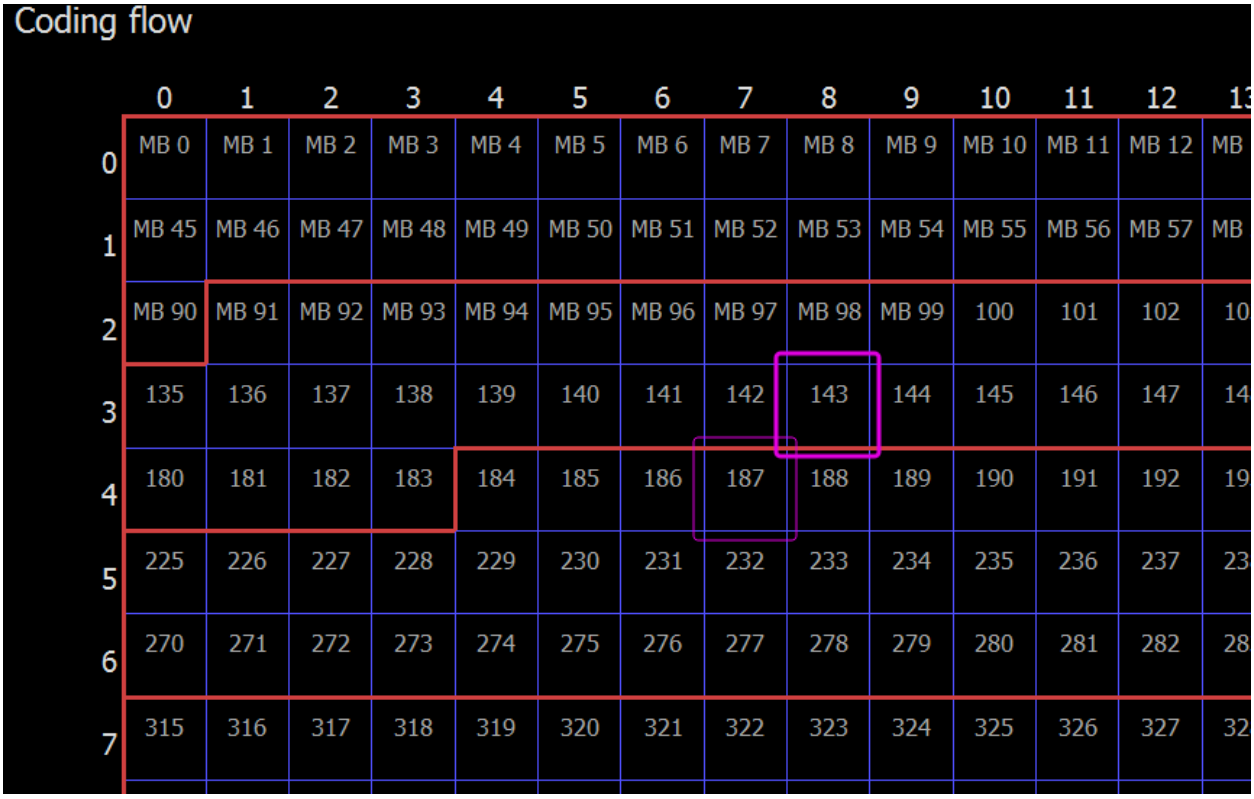

The coding flow mode gives a visual overview of the ordering of MBs in the stream, and some information of the decoding process. The blue grid shows the boundaries of the MBs in the picture. Each MB contains 3 values:

- The MB address, which is simply the raster scan index of the CTB.

5.2.2 Predictions

This mode shows the prediction units in each MB. MB splits are indicated with blue lines, which get darker as the split depth increased. When a MB split is implied by AVC (Direct mode for B slices) and not directly coded in the bitstream, the split lines are dashed.

Intra modes are indicated with green colors, and directional modes also show an arrow indicating the prediction direction. In the lower right corner of an intra MB the chroma mode is indicated in a darker green.

Inter Prediction unitss are indicated with cyan colors showing the Prediction unit splits and mode. Additionally, skipped blocks are shaded with a lined texture. The L0 motion vectors are drawn with an orange color, the L1 motion vectors (B slices only) are drawn with a purple color. The MV value is shown in the lower left corner along with the reference index.

Below is shown Prediction mode on a zoomed-in selection along the top edge of a picture. The selected MB is surrounded with a pink box while the selected Prediction Unit is surrounded with a yellow box.

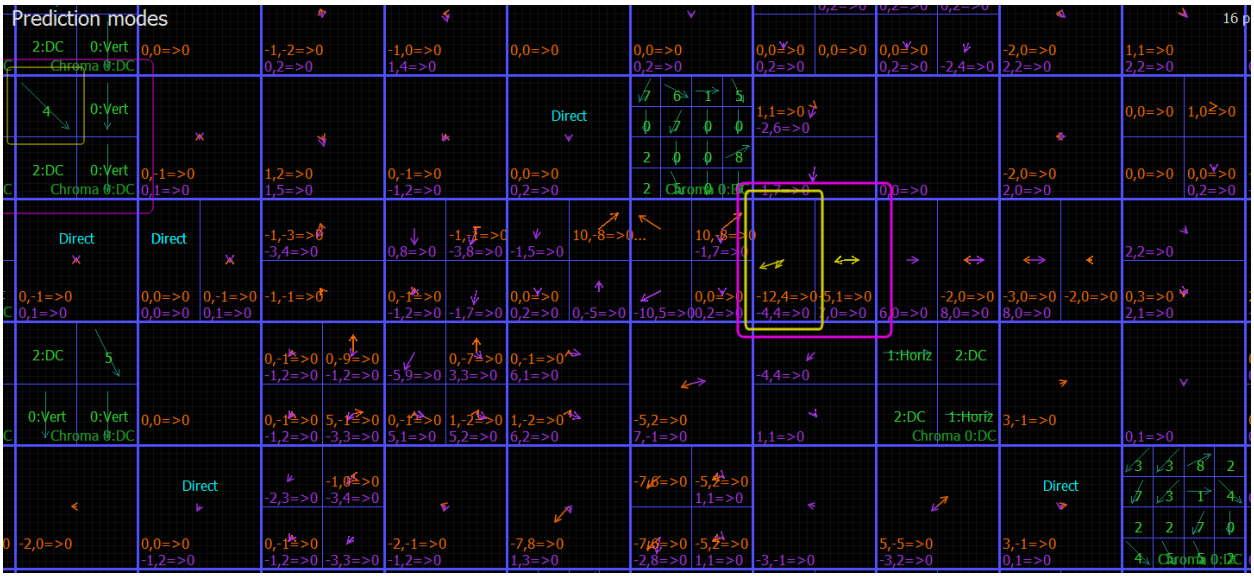

5.2.2.1 Prediction Detail Mode

In Prediction mode, the sample values of a particular MB can be viewed in detail.

Intra MBs are displayed in green, and the left, upper-left and above prediction arrays are drawn adjacent.

Inter MBs are displayed in cyan (turquoise).

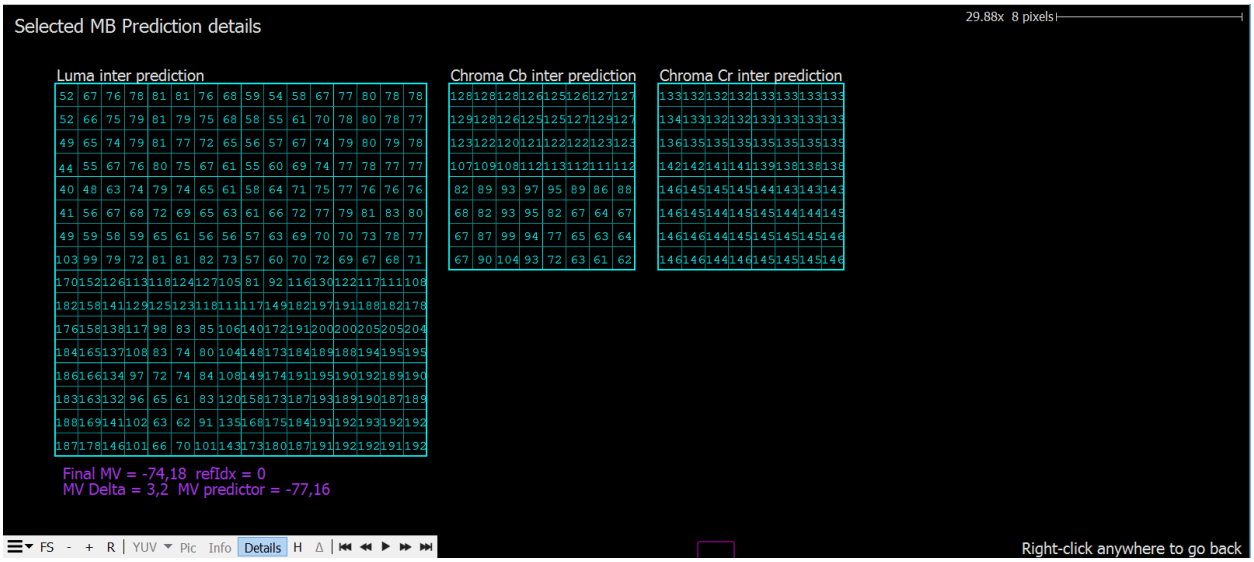

5.2.3 Residuals

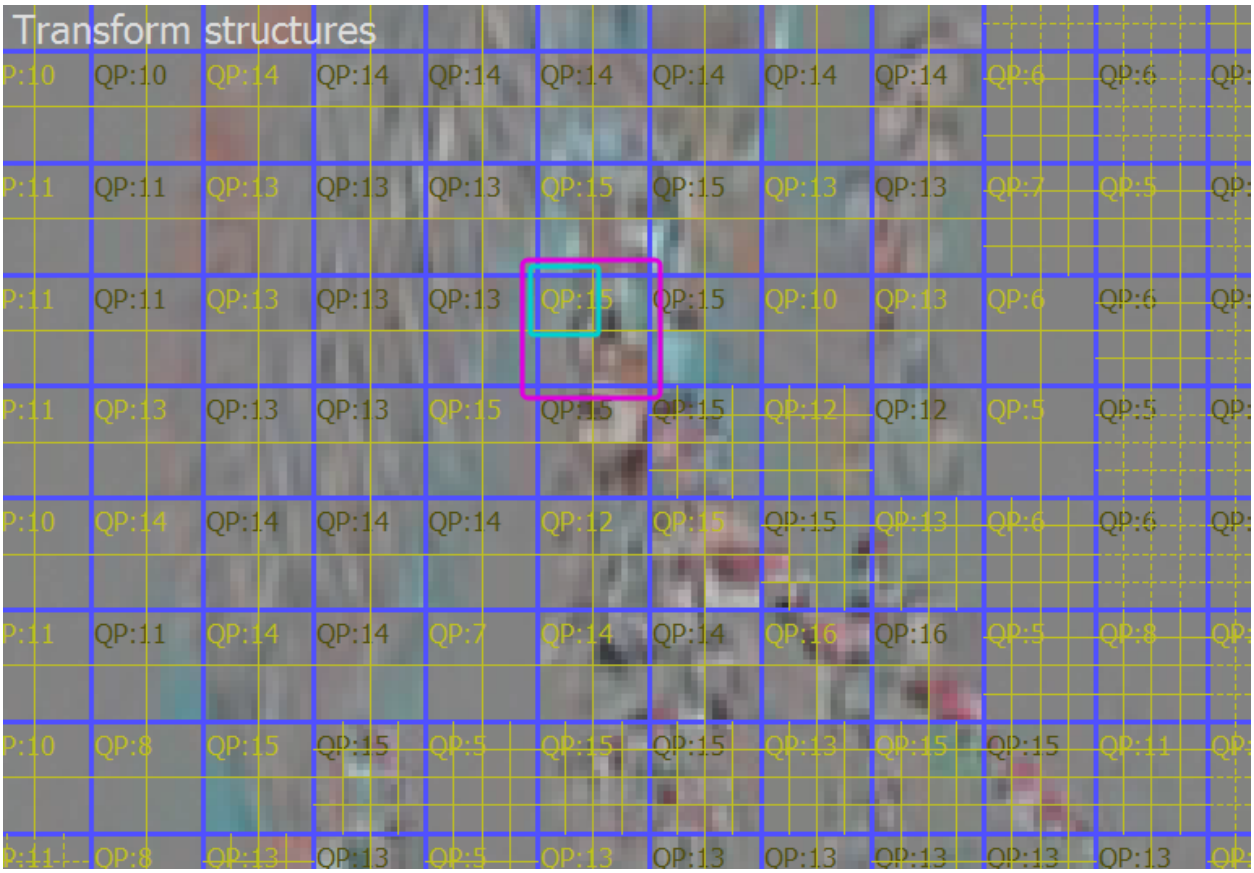

In Residual mode the transforms and accompanying residual signal of the picture can be seen. As in Predictions mode, MB boundaries are shown with blue lines. The transform splits, when present, are indicated with yellow lines. QP values of each MB are shown in the upper-left corner. When this value is bright yellow, it means that a delta-QP was coded in the bitstream during the decode of that particular MB. Otherwise the number is dark. MBs that are PCM coded are marked with "PCM" in white.

If the "Pic" button is turned on, the raw residual signal is shown in image form. Residual values of 0 result in flat grey, negative values are darker and positive values are brighter.

Clicking on MBs causes that transform unit to be selected with a blue surrounding box. Clicking a transform unit repeatedly causes the selection to move up and cycle though the transforms in MB.

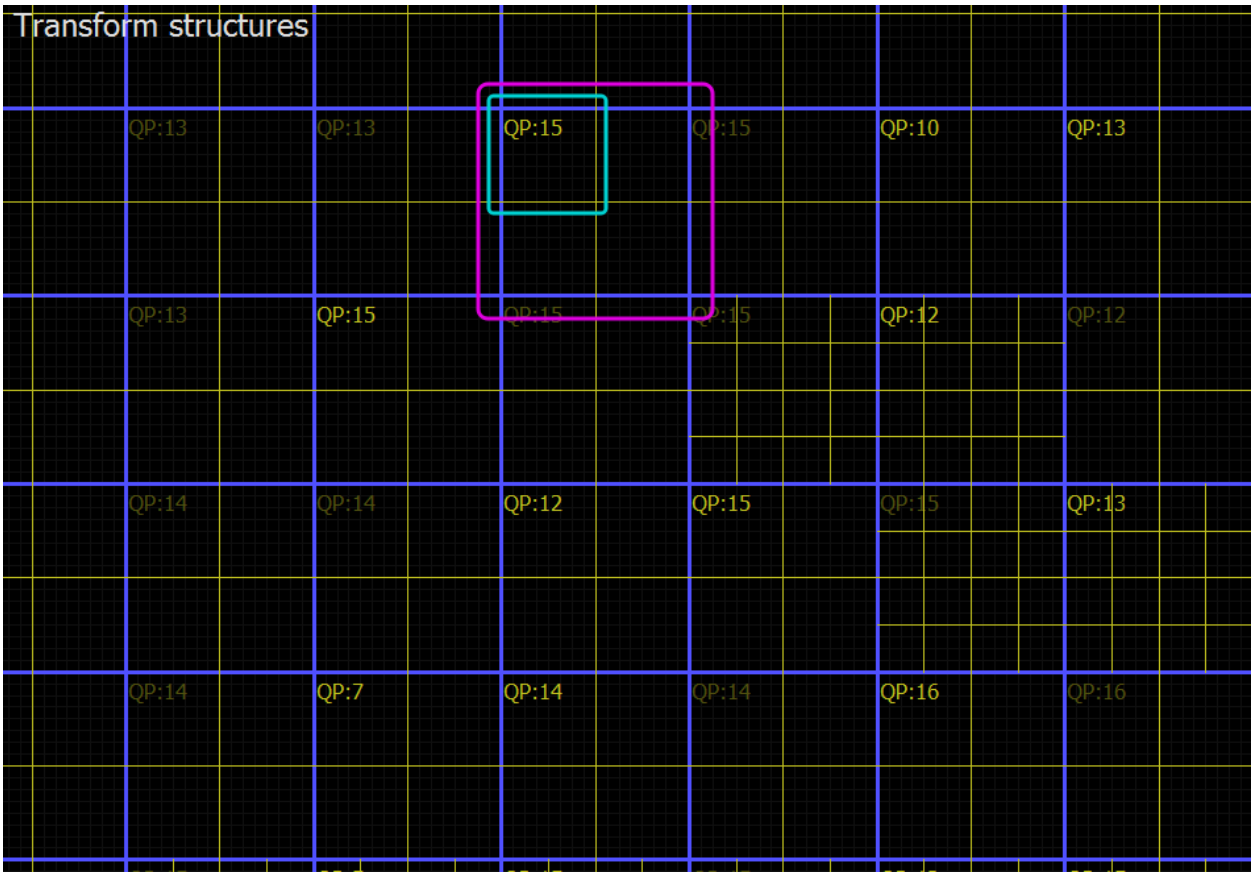

5.2.3.1 Residual Detail Mode

To view the full details of a particular MB, make the MB selection and enter detail mode by right clicking or using the detail mode button at the bottom of the main panel. In this mode the selected TU structure is drawn three times, arranged in a column from top to bottom, showing the 3 major steps in recovering the residual signal:

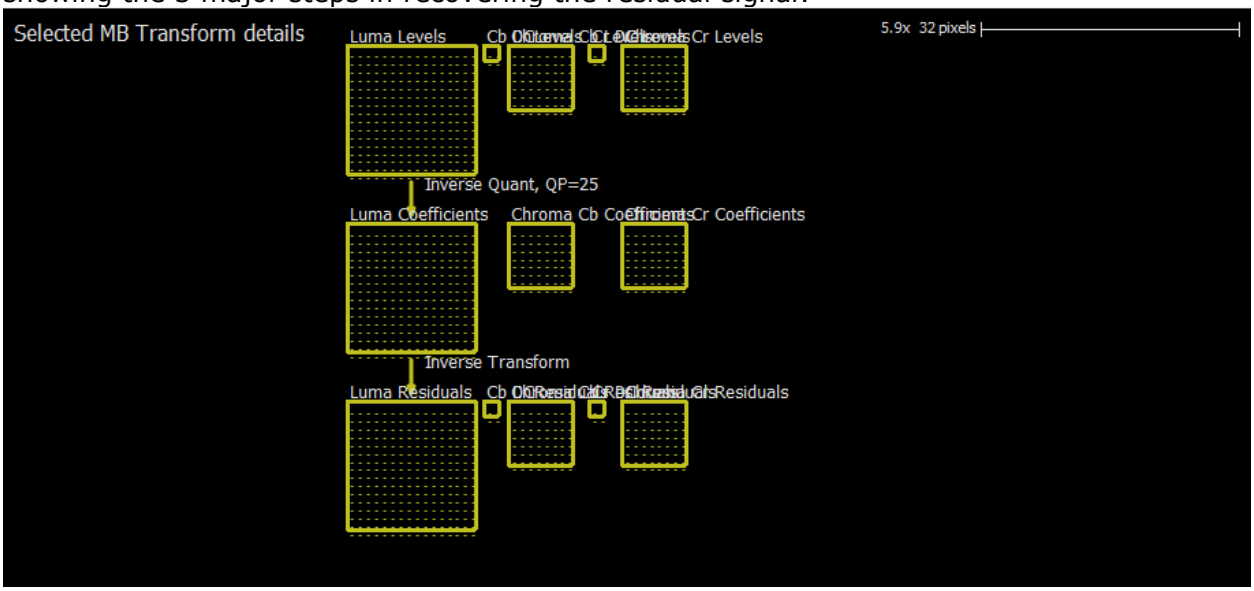

*Other names and brands may be claimed as the property of others.

5.2.4 Reconstruction

In Reconstruction mode the reconstructed samples prior to deblocking can be inspected. As with Prediction/Residual mode, MB boundaries are shown with blue lines. MBs that have at least one transform with non-zero coefficients are marked with "Non-zero CBP". PCM-coded MBs are indicated with "PCM" in white.

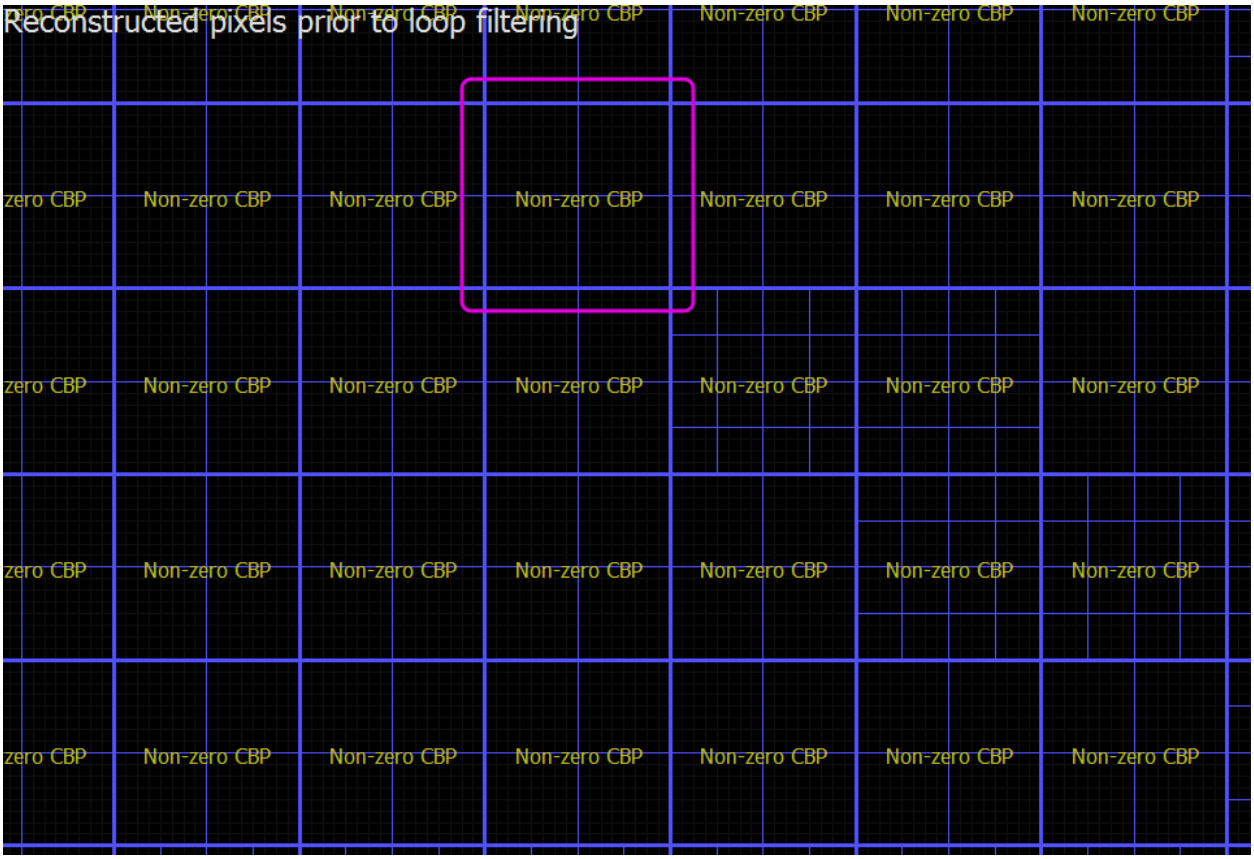

5.2.4.1 Reconstruction Detail Mode

In detail mode, the reconstructed sample values can be inspected. The selection may be a single MB.

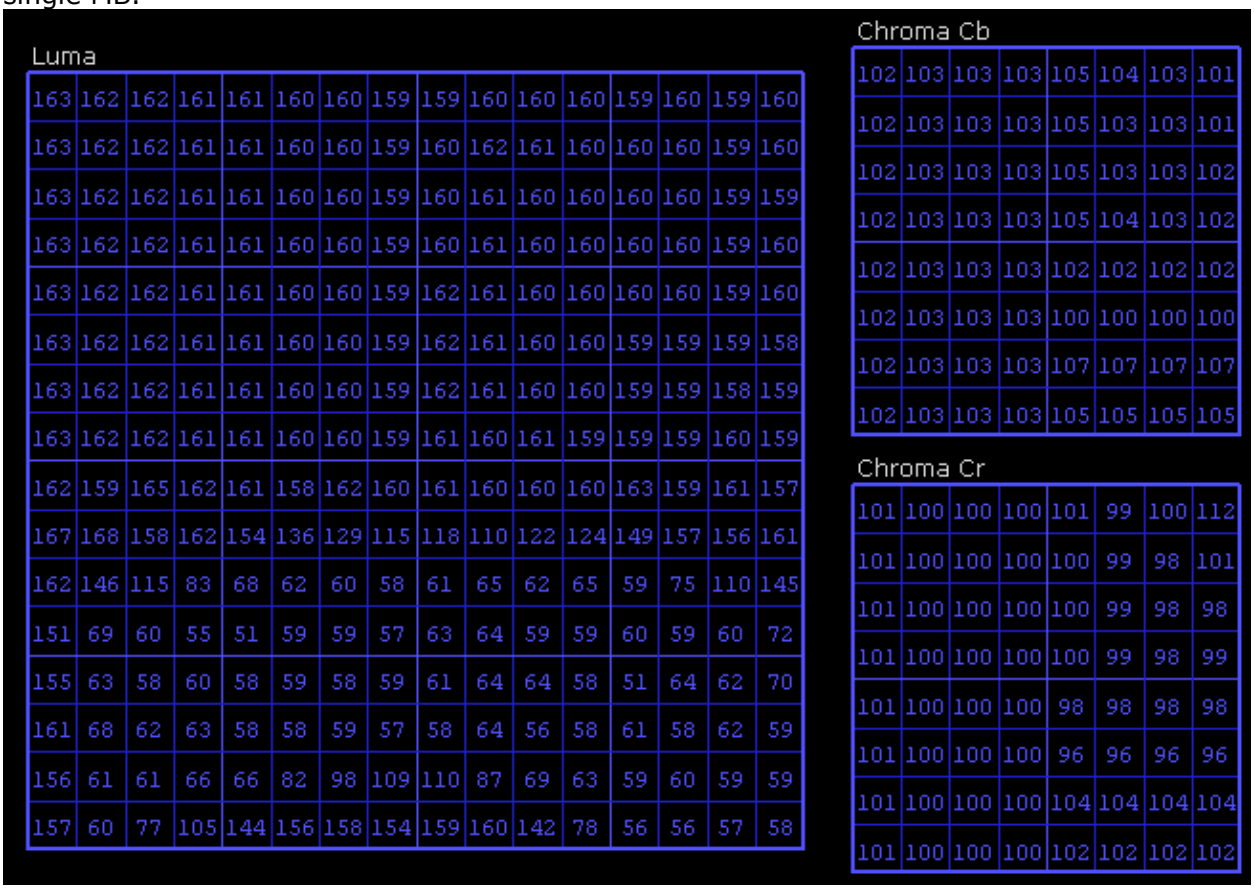

5.2.5 Deblocking

Deblocking mode shows all edges processed by the deblocking filter as described in the MB boundaries are not displayed in this mode, however MB selection still functions as in the other modes. Edges shown are for the luma deblocking process and are color coded in the following manner:

- Green: Strong luma filter applied
- Yellow: Weak luma filter applied
- Red: No filter applied (but edge was evaluated)

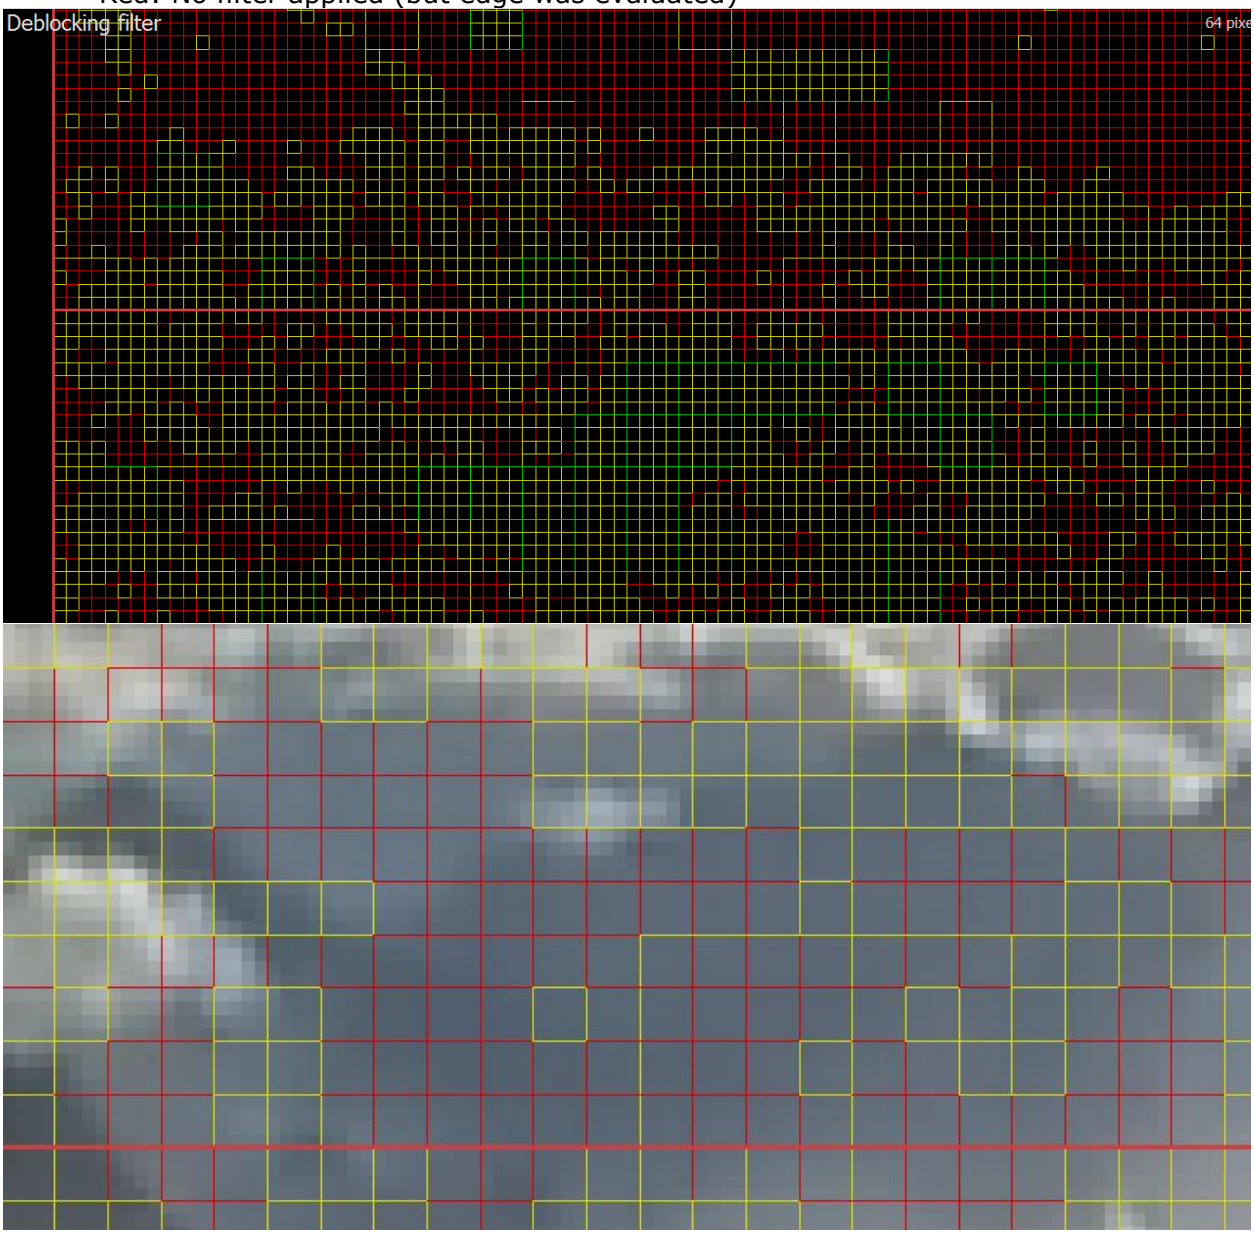

5.2.5.1 Deblocking Detail Mode

The deblocked sample values can be inspected directly in detail mode. As with Reconstruction Details. Sample values are indicated in yellow if they were modified by the deblocking process. Sample values that cannot be modified due to PCM MB or bypass_flag are indicated in red. Filtered edges belonging to the selected CU are shown in thick dashed lines, color coded the same way as in the non-detail mode. Each edge can be selected and clicked, bringing up a window showing the input values to the edge filter. Note that these values may differ from the reconstructed samples for horizontal edges, as they may have been modified by a vertical edge filter. Additionally, the boundary strength is displayed. Edges that were processed but ultimately not filtered (red colored) will indicate the reason. Typically it is due to the boundary strength being 0, the threshold condition not being met, or a slice/tile edge when those edges shouldn't be processed.

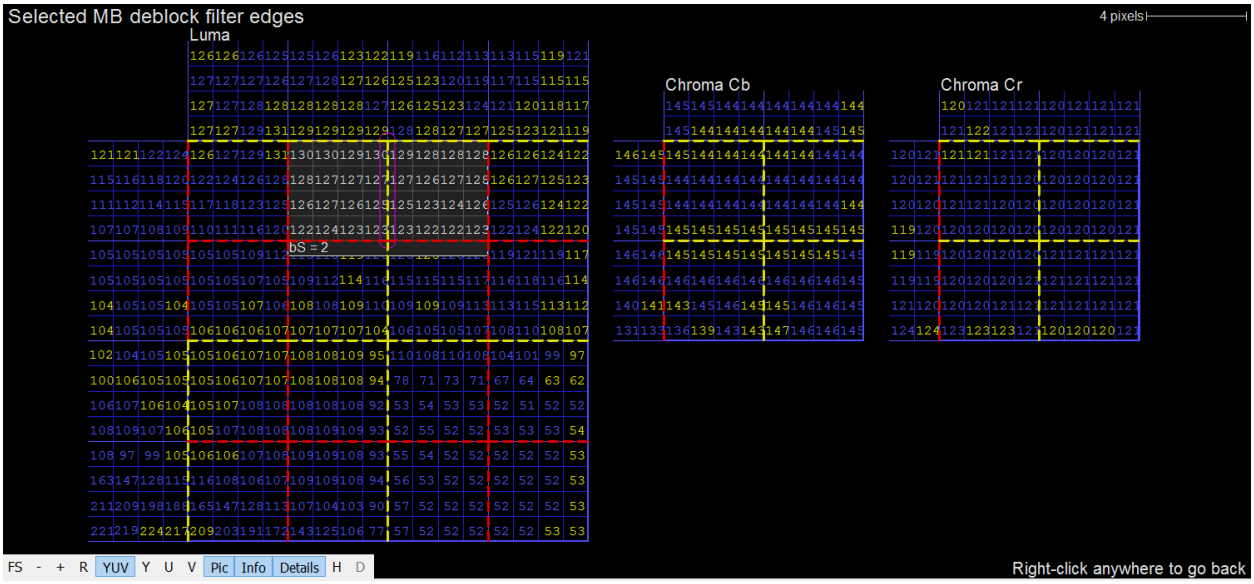

5.2.6 YUV

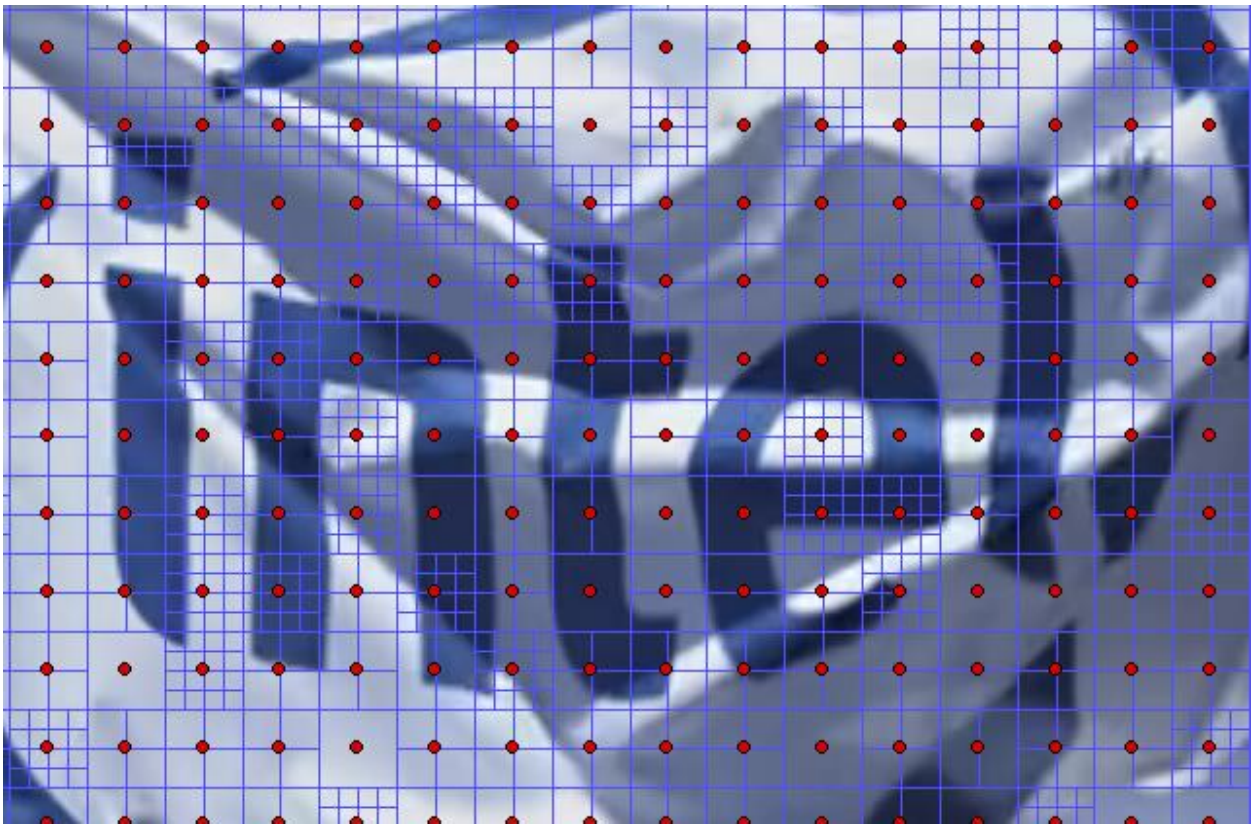

YUV mode allows inspection of the final decoded sample values, without additional overlay data. In this mode, the YUVDiff feature is exposed. This feature allows for comparison of an external decoded YUV file with the decoded bitstream. Supported formats are 4:2:0 planar and NV12. Also, YUV files that are zipped or gzipped may be opened directly without the need to decompress separately. If a zip file contains more than one YUV file, only the first one is used. When a YUV file is loaded, any mismatches will be indicated with a red dot in the CU containing the mismatch. This allows the user to quickly identify the nature of the mismatch which can assist with debug.

When a file is opened, four additional buttons become visible on the lower left of the main window:

- Ref YUV: The original, expected YUV image.
- Debug YUV: The loaded debug YUV file.
- Diff YUV: The delta image. Like in Residual mode, areas with zero delta (i.e. Original and Debug image are identical) are flat gray. Areas where the debug YUV has a lower value are darker, and areas where the debug YUV has a higher value are brighter.
- Reload: Reloads the YUV file, which can be useful as a shortcut to loading via the YUVDiff menu.

5.2.6.1 YUVDiff Menu

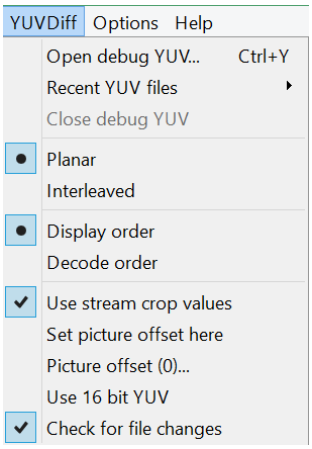

- Open: Brings up a file chooser to open a YUV file. Intel Video Pro Analyzer considers YUV stream as having the same frame size, per frame bitdepth and color format as encoded bitstream.
- Close: Closes the currently opened YUV file. Any mismatch indicators will disappear.
- Planar: Causes Intel Video Pro Analyzer to interpret the loaded YUV file as Planar.
- Interleaved: Causes Intel Video Pro Analyzer to interpret the loaded YUV file as Y values followed by interleaved U and V values. This option is chosen automatically when the loaded YUV file has the file extension ".nv12".
- Display order: Causes Intel Video Pro Analyzer to use display order picture numbers when determining how far to seek into the YUV file to extract the image to compare with the current picture.
- Decode order: Causes Intel Video Pro Analyzer to use decode order picture numbers when determining how far to seek into the YUV file to extract the image to compare with the current picture.
- Use stream crop values: When checked, YUV files are assumed to contain samples only within the cropping window as defined by the loaded bitstream. Samples outside this window are not compared, and are assumed 0 when viewing the loaded YUV image directly.
- Set picture offset here: Shortcut for setting the picture offset to the current picture number. See bullet below.
- Picture offset: This brings up a dialog allowing the user to enter the picture number of the first picture in the YUV file. For example if a 100-picture bitstream is loaded but the YUV file only contains pictures 80-99, the user would enter 80 to properly line up the YUV file with the decoded bitstream.
- Check for file changes: When checked, this option causes Intel Video Pro Analyzer to periodically check if the loaded YUV file has changed on disk since it was last loaded. If the file has indeed changed, a dialog pops up offering a chance to reload the YUV file.

5.2.6.2 YUV Detail Mode

In detail mode the expected YUV values can be inspected. Note that the displayed sample values are the same as those in Deblocking mode since loop filtering is the last step in the decode process. When a YUV file is loaded, the debug and delta YUV values can be inspected as well. Mismatching values are shown in red.

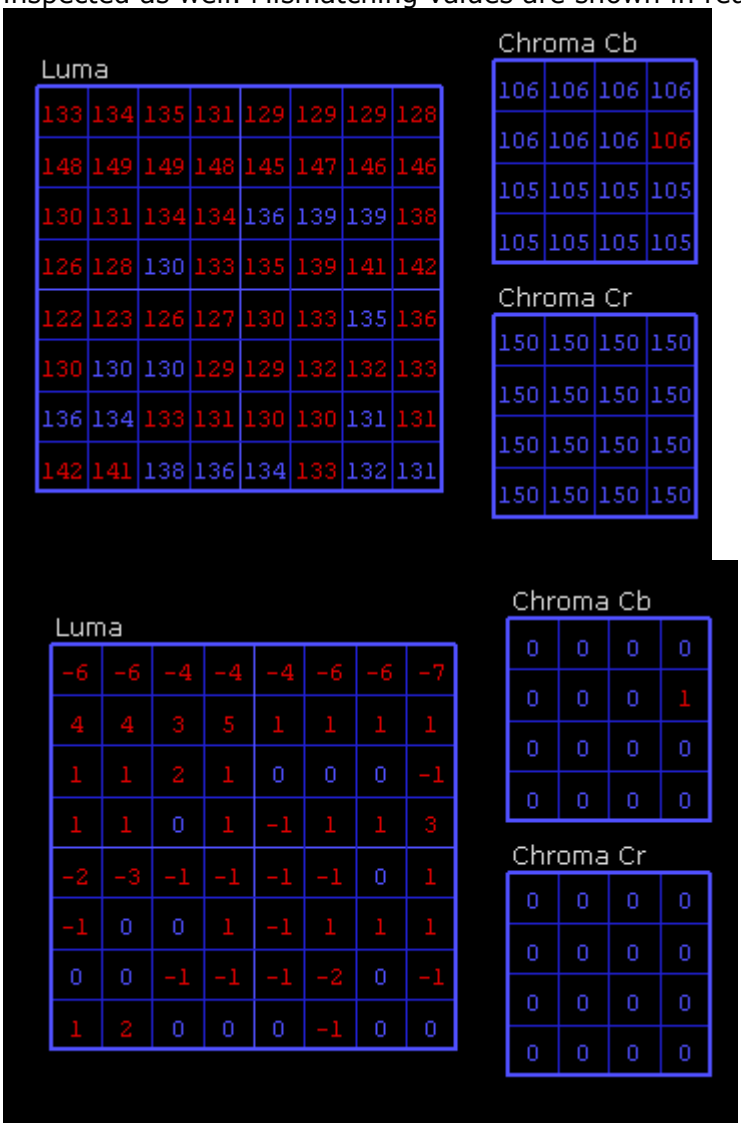

5.2.7 Info Overlays

There are a number of overlay that provides additional information about currently investigated picture. When you switch mode to "Info Overlays" mode several additional submenus will appear in the bottom-left panel of info window.

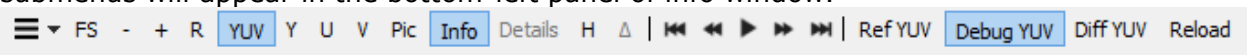

You can switch now between additional modes. PSNR mode is enabled if you have loaded reference YUV file.

5.2.7.1 Heat Map

Heat map mode shows visually how the compressed bits of the picture are distributed spatially. MBs with more bits per pixel are brighter than MBs with less. By default this mode also shows the MB boundaries in blue. It may be useful to turn off this overlay using the "Info" button on the lower left of the main panel. Heat Map does not offer a detail mode.

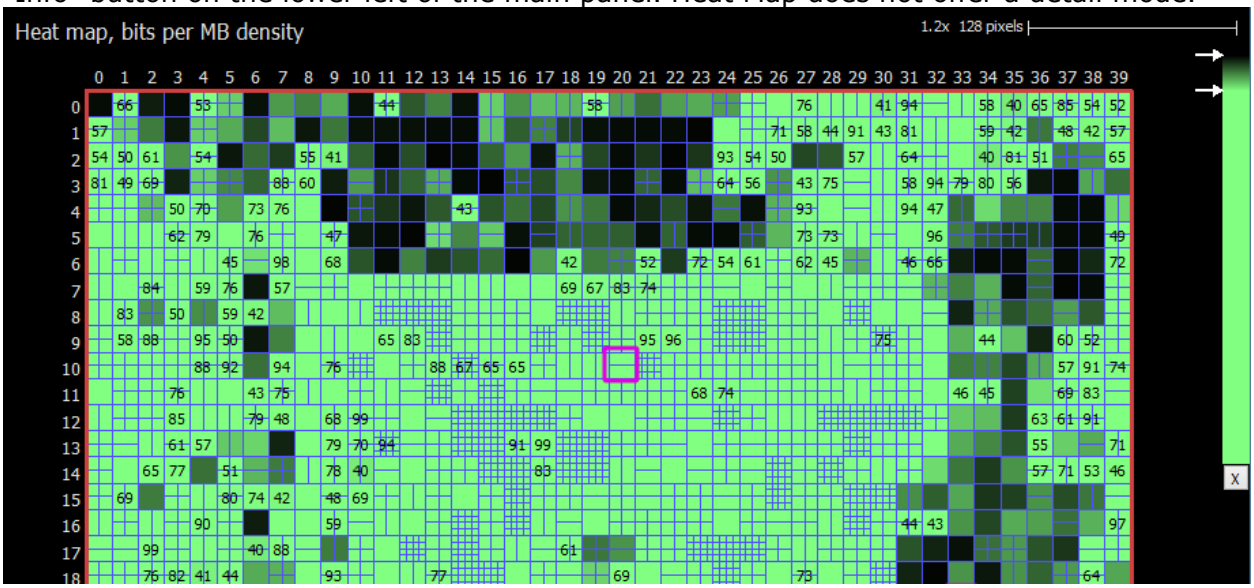

The gradient used to display the heat map may be edited:

- Drag the gradient markers up and down to reposition them.
- Drag a marker away from the gradient to remove it. A red X indicates removal will take place.
- Double-click a marker to change the color.
- Double click the gradient bar to add a new marker.

5.2.7.2 QP Map

QP Map show per MB QP. MBs with bigger QP are brighter than MBs with less. Coloring of the QP range can be adjusted with a gradient tool on the right top corner.

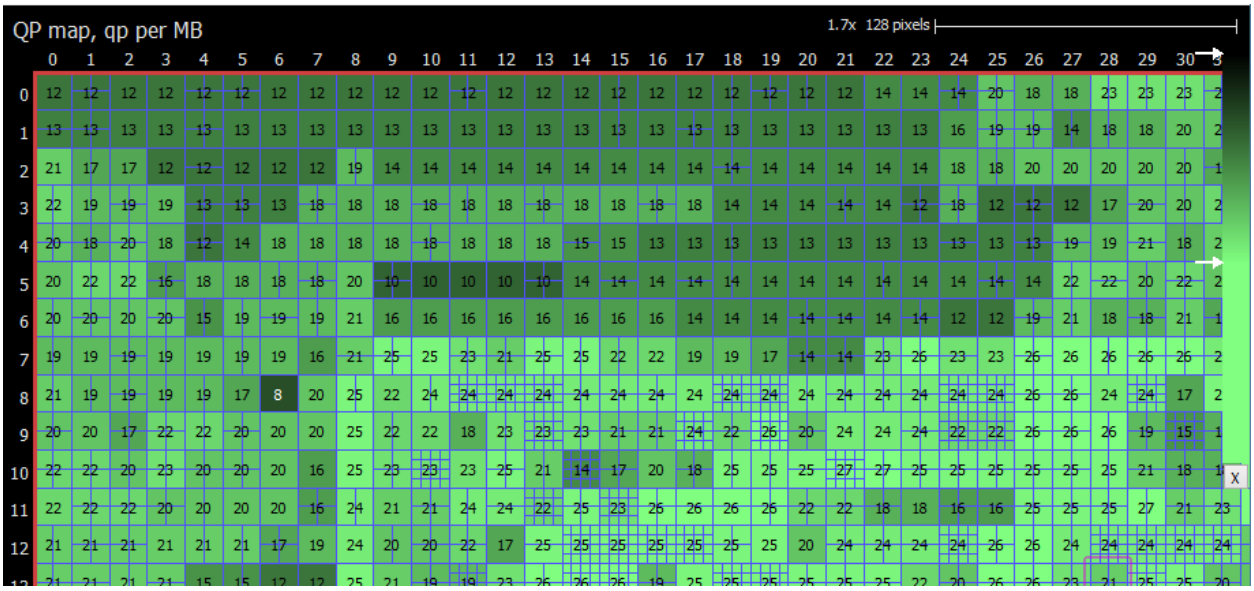

5.2.7.3 MB Type

MB Type overlay show color map of their types. Intra blocks is colored in red, inter skip is yellow colored, B-Direct is in green and inter in blue.

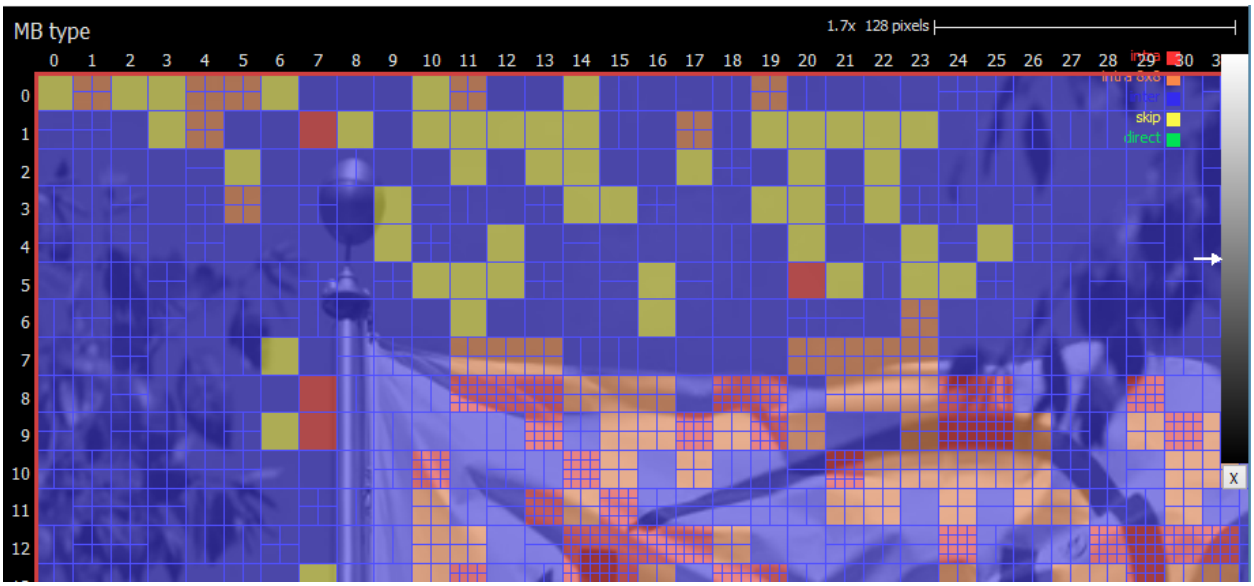

5.2.7.4 MB Reference Indices

MB reference indices shows reference indices for each subblock 8x8. Orange color shows L0 reference index and purple color is for L1 reference index. When PU has only one reference missing reference index is displayed with dash. For blocks with no reference indices nothing is shown.

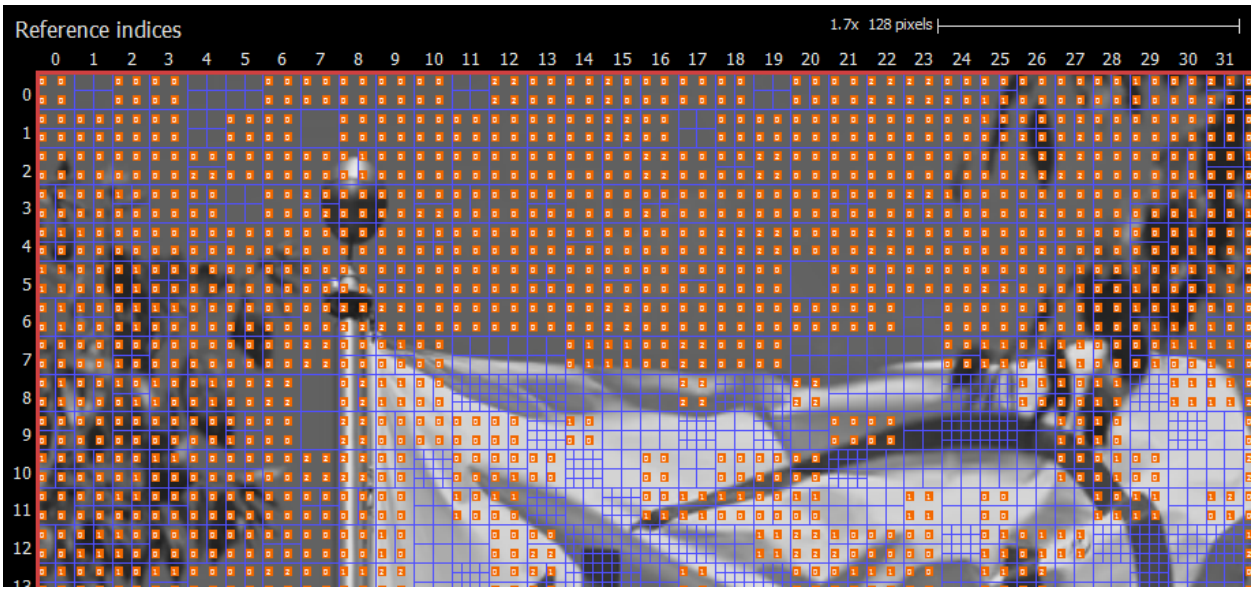

5.2.7.5 PSNR

PSNR map shows per pixel PSNR value normalized on MB size. To enable this mode you have to load reference YUV file via DebugYUV menu. The value is calculated per MB. Bigger value has brighter color. The color range can be adjust with gradient tool on the right top corner. If the block has no difference with reference YUV a latter 'e' is shown which means exact match.

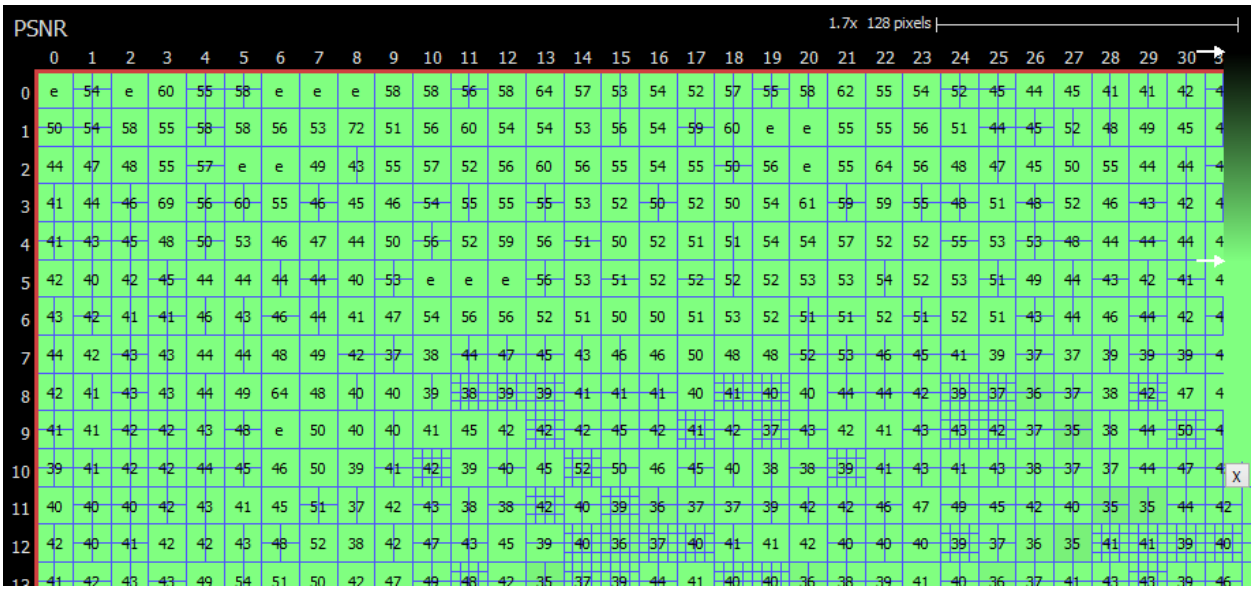

5.2.8 Simple Motion

*Other names and brands may be claimed as the property of others.

Simple Motion mode offers a way to quickly view the modes and motion of a picture at a glance. No overlays are present except a colored dot showing information about the PU's mode. Green dots indicate intra prediction, and purple/orange L0/L1 motion vectors for inter prediction are drawn as simple lines. When the motion vector is very small, a dot is drawn instead so that the mode is still easy to identify. When a motion vector does not point to the 0-th index of its respective list, the vector is drawn using a dashed line.

6 MPEG-2

The following sections describe all available features when loading a MPEG-2 bitstream. The bitstream must be containerized with mpeg-2 container format.

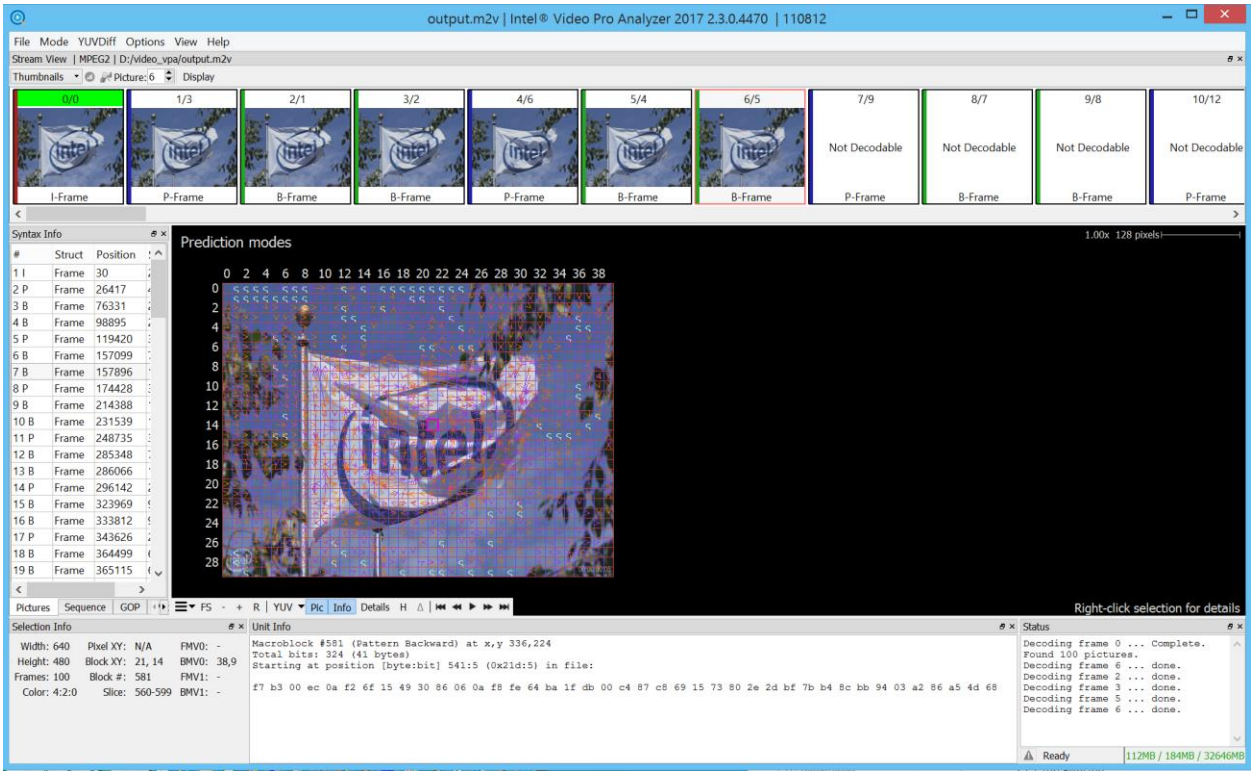

6.1 UI Components

The following sections describe the various components in the UI when an VP9 bitstream is loaded.

6.1.1 Top Filmstrip

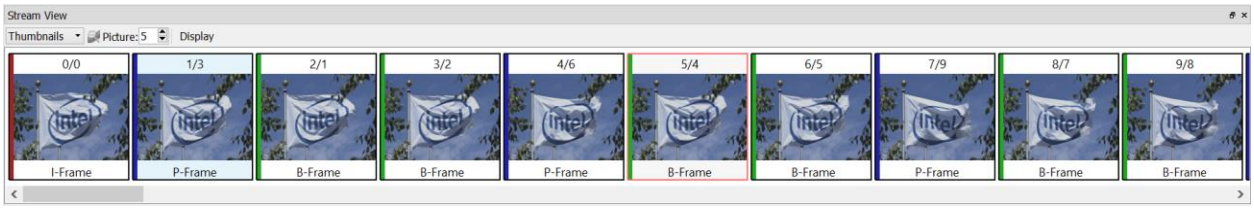

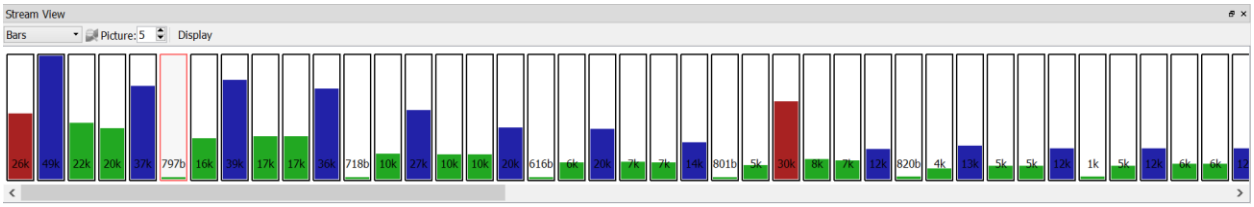

The top filmstrip is a horizontally scrolling overview of the pictures in the bitstream. The current picture is highlighted. On top is the frame number in both decode and display order respectively and the picture structure (frame- or field-picture) as indicated by the frame header. On bottom is indicated the picture coding type (I-, P- or B-frame).

Clicking on a picture will cause Intel Video Pro Analyzer to decode it and mark it as the current picture, updating the rest of the UI correspondingly. If the new picture requires other pictures to be available as reference that have not yet been decoded, those will automatically be decoded.

Pictures are decoded on demand since it is impractical to store all details of all pictures in memory, especially for HD sequences. The current picture can also be selected by typing in a picture number in the box on the left of the panel. When using this box, the filmstrip is scrolled such that the newly selected current picture is visible.

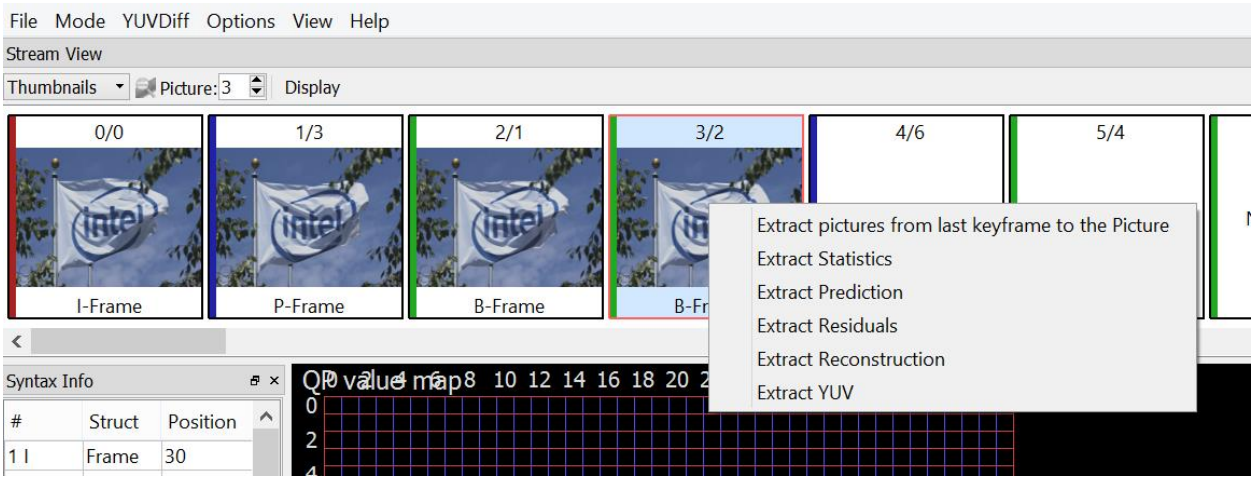

In Thumbnails and Bar mode it is possible to choose between Display/Decode order representation by click on additional. By right clicking on selected picture one can get popup

submenu with a set of extract actions. An action called "Extract pictures from last keyframe to the Picture" will write out a bitstream that contains the minimum number of pictures needed to decode the current picture. Typically this will consist of the current picture and all previous pictures that are reference pictures up to the nearest keyframe picture. This can be useful for debugging issues in long sequences. Other actions allow to save in the output file statistics and intermediate/final YUV planes.

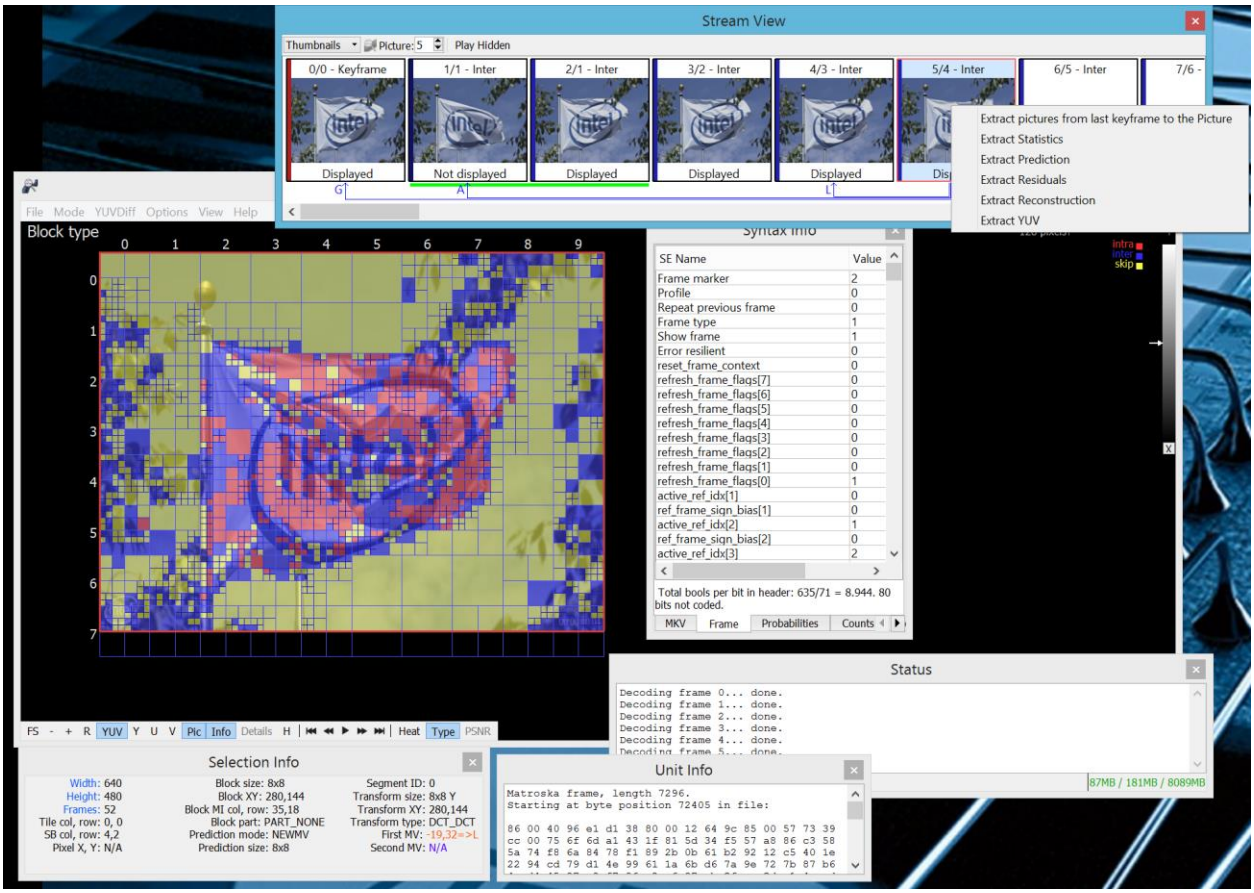

Allow separate panels can be detached from them main window or hidden. To attach panels back double click on the titlebar. To unhide panel go to View menu and check appropriate panel.

6.1.2 Main Panel

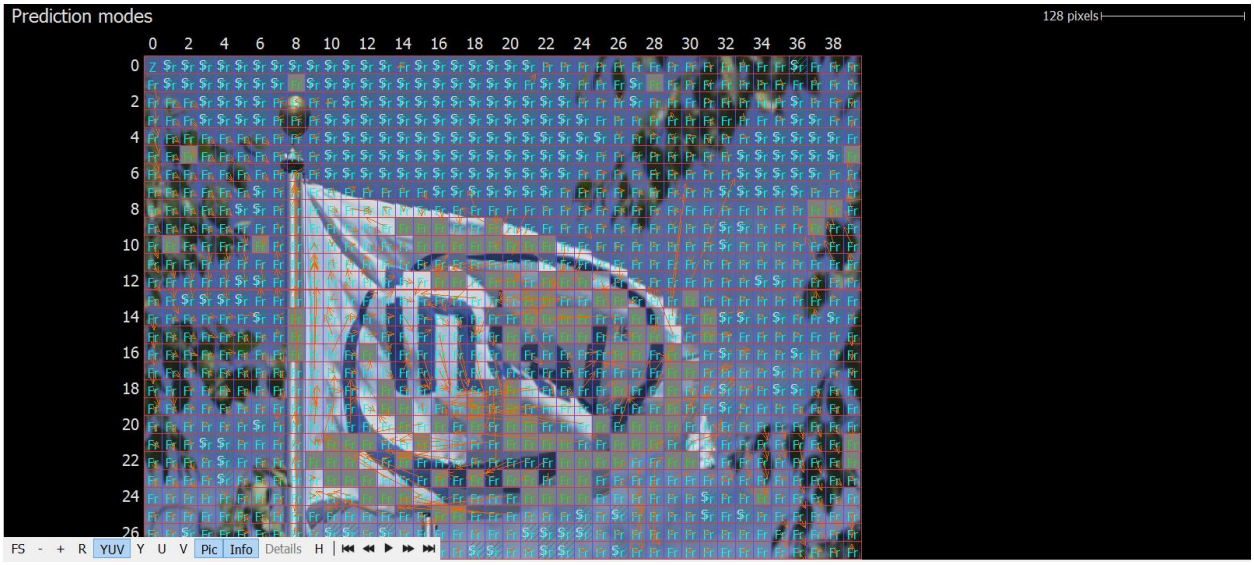

The main panel displays the selected picture with visual annotations. The type of annotations and associated interactive behavior depends on the current mode, see 4.2 MPEG-2 Modes. Click-dragging moves the picture around, and the mouse wheel zooms in or out about the mouse cursor.

At the top left of the panel, the current mode is displayed. Along the top right the current scale is shown. In all modes, the macro-block row and column values are displayed along the left and top border of the picture and tile boundaries are thick green lines. The screenshot above shows an example of a stream that is 44x14 SBs.

The bottom left has a cluster of buttons

- FS: Toggles full-screen mode, which hides all UI elements except the main panel. Keyboard shortcut is F.
- +/-: Zooms in or out, centered on the center of the panel. Note that zooming with the mousewheel is much easier than using these buttons. Keyboard shortcuts are + and -.
- R: Resets the current zoom to fit the entire picture in the main panel. Keyboard shortcut is R.
- YUV, Y, U, V: Toggles between per-component image, or full-color YUV.
- Pic: Toggles the actual picture underneath the annotations. The picture is mode-dependent. For example in prediction mode, the picture is comprised of the prediciton samples before the residual signal has been added in. It is often helpful to turn off the picture in order to make the annotation easier to read. Keyboard shortcut is P.
- Info: Toggles the annotations on or off. Keyboard shortcut is I.
- H/D: Toggles between hexadecimal and decimal display of the annotations. This applies to the values in the Left Tabs as well. Keyboard shortcuts are H and D.

- There are five navigation buttons: first frame, previous frame, playback mode, next frame, last frame. The previous/last button will rewind, causing the current picture to be set to the first/last picture in the loaded bitstream. The "Play" button will start playback of loaded bitstream at a reasonable rate, limited by CPU performance. All normal mode operations are possible during playback. Note that playback mode always operates in currently selected order in the filmstrip.

6.1.3 Left Tabs

The left side of the UI is a resizable panel with several tabs that display information about the current selection in one way or another. The following subsections describe each tab.

6.1.3.1 Pictures (bitstream)

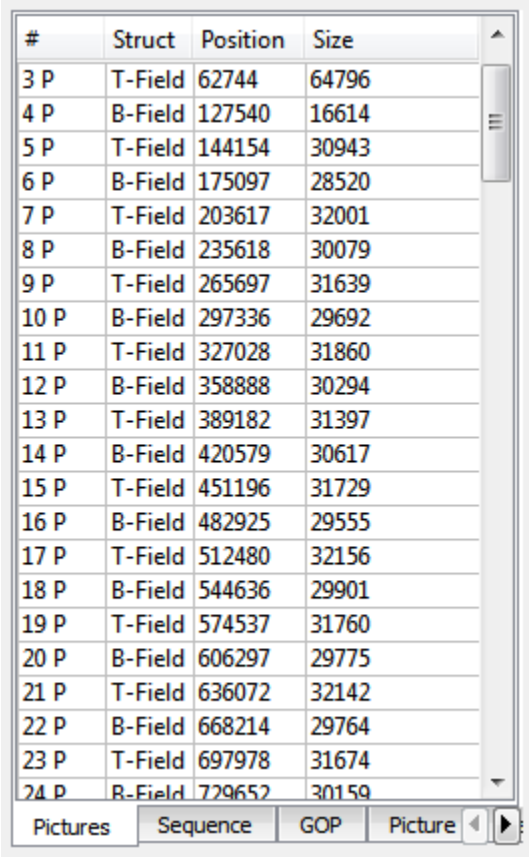

#	Struct	Position	Size
3 P	T-Field	62744	64796
4 P	B-Field	127540	16614
5 P	T-Field	144154	30943
6 P	B-Field	175097	28520
7 P	T-Field	203617	32001
8 P	B-Field	235618	30079
9 P	T-Field	265697	31639
10 P	B-Field	297336	29692
11 P	T-Field	327028	31860
12 P	B-Field	358888	30294
13 P	T-Field	389182	31397
14 P	B-Field	420579	30617
15 P	T-Field	451196	31729
16 P	B-Field	482925	29555
17 P	T-Field	512480	32156
18 P	B-Field	544636	29901
19 P	T-Field	574537	31760
20 P	B-Field	606297	29775
21 P	T-Field	636072	32142
22 P	B-Field	668214	29764
23 P	T-Field	697978	31674
24 P	B-Field	729652	30159

Pictures | Sequence | GOP | Picture

In "Pictures" tab VPA presents raw picture sequence with type, offset and size.

6.1.3.2 Sequence

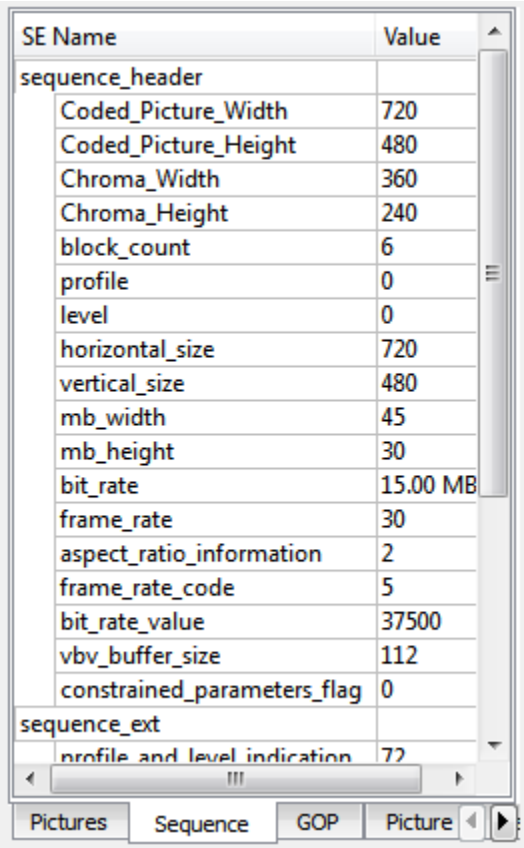

The Sequence tab shows all syntax elements decoded in the sequence header and sequence extention header, both the uncompressed and bool-coded partitions.

6.1.3.3 GOP

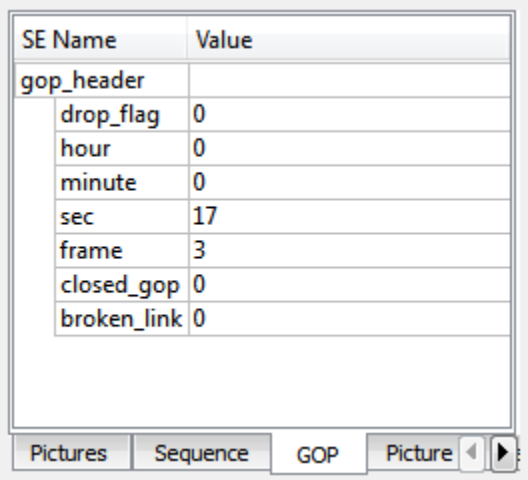

The GOP tab shows all syntax element decoded in GOP header, both the uncompressed and bool-coded partitions.

6.1.3.4 Picture

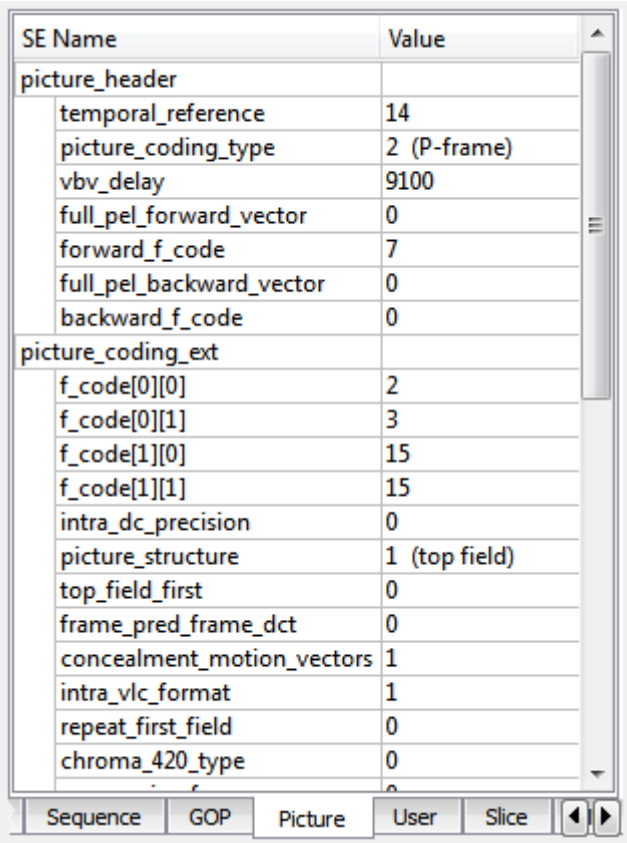

The Picture tab shows all syntax element decoded in picture header, picture coding extension, picture display extension and picture spatial scalable extension, all of them are uncompressed and bool-coded.

6.1.3.5 Refs

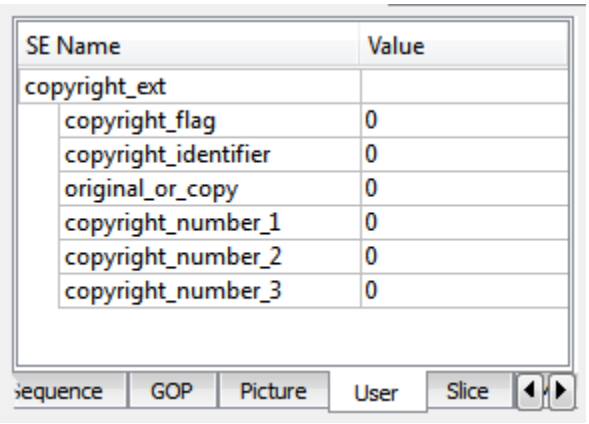

This panel displays the data that is decoded in user extension.

6.1.3.6 Slice

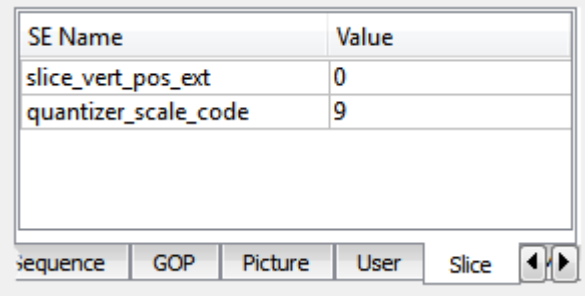

This panel displays the syntax elements associated with the currently selected slice.

6.1.3.7 TX

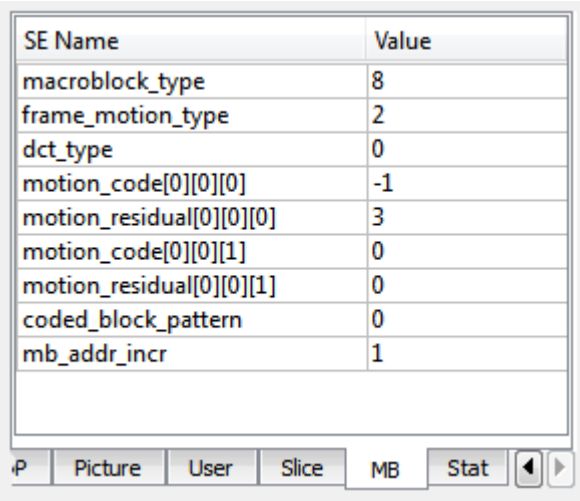

This panel displays the syntax elements associated with the currently selected macro block.

6.1.3.8 Stats

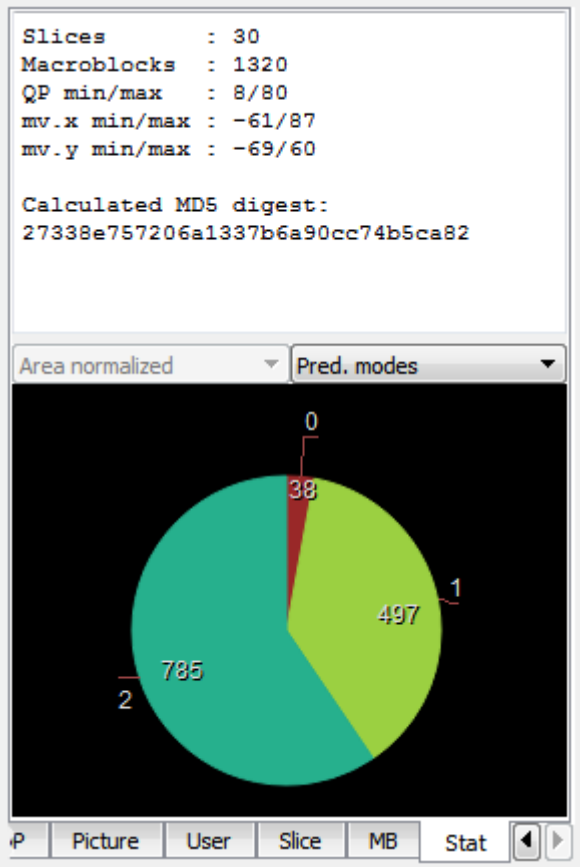

This tab displays various statistics extracted from the current picture. The top half shows some picture size and compression stats as well as the image digest information. The bottom half displays pie charts for a number of metrics. Each pie chart can be drawn normalized or un-normalized. Normalized data is weighted by area or bits.

The pie chart can be moved by dragging the mouse, and zoomed with the mouse scroll wheel.

6.1.4 Selection Info

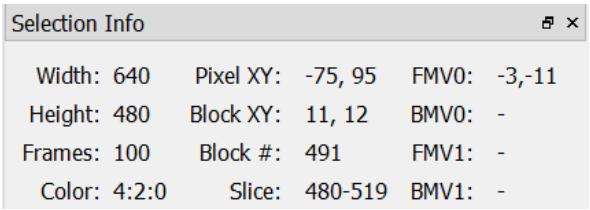

This panel shows a few details about the current selection at a glance:

- Frame height & width
- Number of frames in the bitstream
- Chroma format (4:2:0, 4:2:2 or 4:4:4)
- X/Y position of the mouse cursor in the frame
- Current macro-block row and column number
- Current macro-block number in slice
- Macro-block numbers, which belong to current slice ("1st MB number" – "last MB number")
- MV list (up to four) of currently selected MB, where FMV0, BMV0 are forward and backward MVs for the first field and FMV1, BMV1 are forward and backward MVs for the second field. If no MV is present in stream, "--" sign is displayed.

6.1.5 Raw Bytes

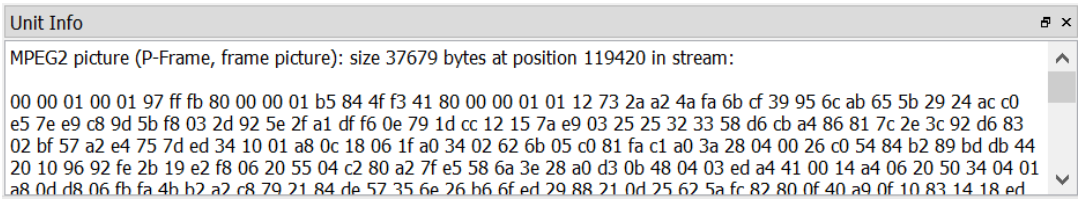

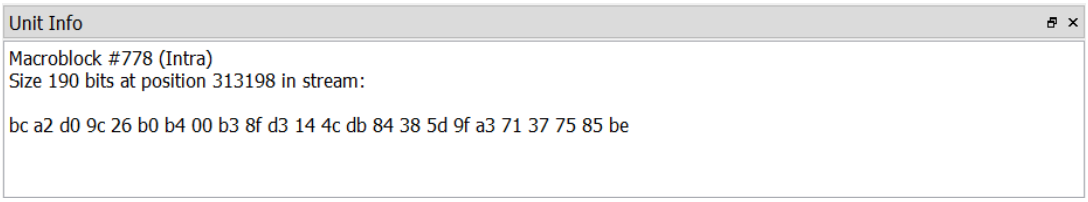

This resizable panel displays raw bytes from the bitstream as it applies the current selection, as well as few extra details. The current selection can be a particular block or MPEG-2 frame.

6.1.6 Messages

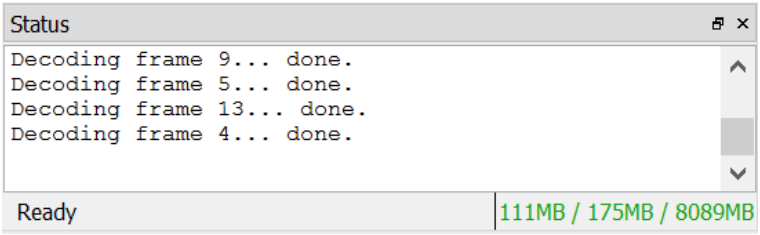

This resizable panel in the lower right of the UI displays messages about the decoding process, and shows progress of any actions that may take a while to complete.
The 3 numbers in the lower right of this panel represent the memory state of the proccess. From left to right the 3 numbers are:

- Amount of memory used by process
- Peak memory used by process
- Total memory claimed by system

6.2 Modes

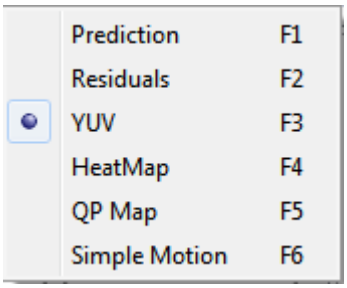

With a MPEG-2 bitstream loaded, Intel Video Pro Analyzer can be put into one of 6 modes using either the F1-F6 keys, or using the Mode menu. The mode selection affects only what is displayed in the Main Panel.

Most modes can also show details of the current selection (block, prediction/transform block). This can be toggled with the right mouse button or the main panel's button strip.

6.2.1 Predictions

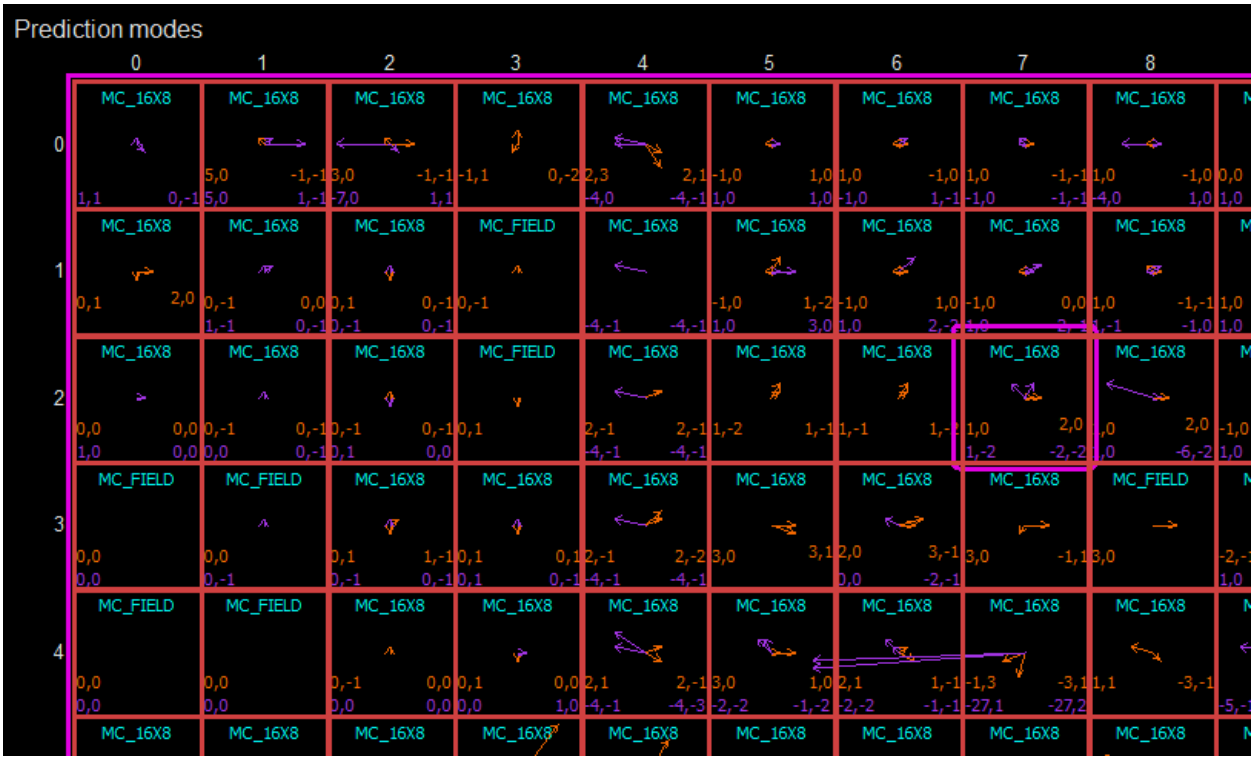

This mode shows the details of the prediction blocks in the frame. The selected block is surrounded with a pink box. When a prediction block is selected, the syntax used to code it is displayed in the Block tab of the Left Panel.

Intra prediction block contain the intra luma mode in bright green, and can contain so called "concealment" motion vectors if they are present in bitstream.

Inter prediction blocks show the inter mode in cyan. There are up to two pairs of MVs are drawn in block. Left pair corresponds with the first field? The right one – with the second one (if any). The forward motion vector in orange and the backward one in purple. The motion vector value is in units of 1/8th pixels.

6.2.1.1 Prediction Detail Mode

Entering detail mode on a selected prediction block allows the sample values of the prediction process to be viewed directly. No details mode is available for intra blocks.

Inter-predicted blocks display contains predicted values for luma and chroma components.

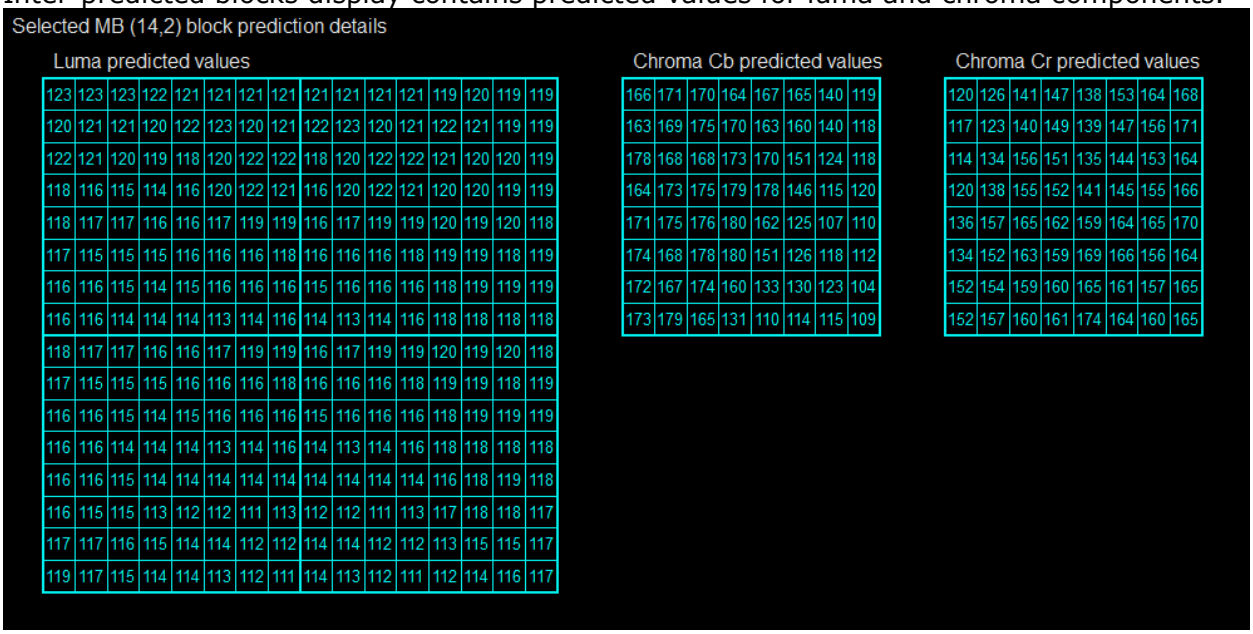

6.2.2 Residuals

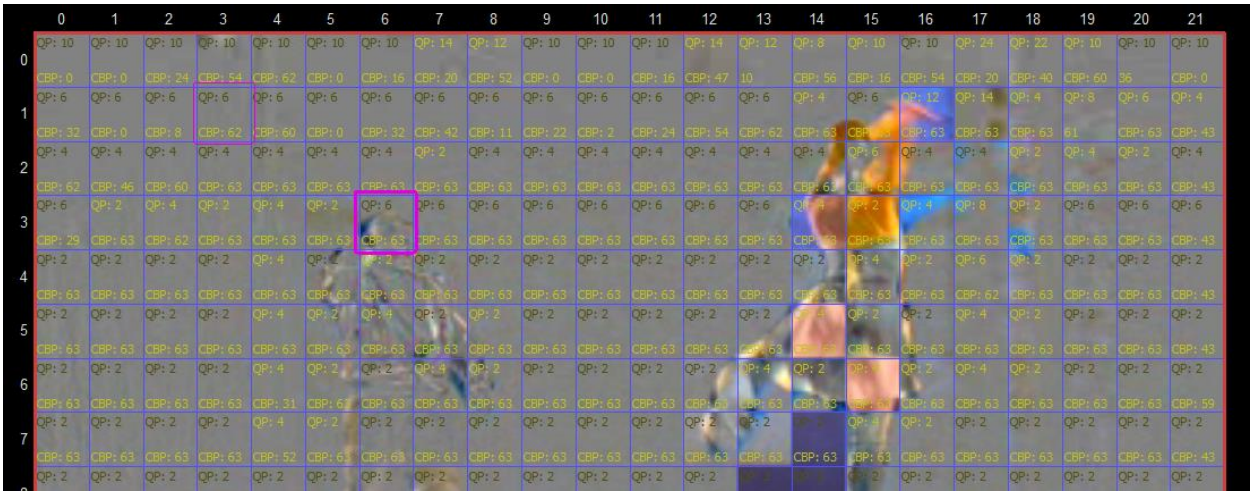

Residual mode allows inspection of the transform structure of each block. Block boundaries are displayed in blue, the currently selected block is outlined with a pink box. Blocks that code no residual signal are indicated with "Skip" label.

Clicking a particular transform block causes the syntax elements used to decode it to be visible in the "MB" tab of the Left Panel.

The residual image itself can be shown using the "Pic" button on the bottom button strip. Residual values of 0 are flat gray, negative values are darker and positive values are brighter.

6.2.2.1 Residual Detail Mode

To view the full details of the transform blocks belonging to a particular block, selected a transform block and enter detail mode by right clicking or using the detail mode button at the bottom of the main panel. The currently selected transform block is highlighted with the sample grid. The transform structure of the selected block is drawn two times, arranged in a column from top to bottom, showing the 2 major steps in recovering the residual signal:

4. Inverse scaling process. The scaled values and the scaling factors used to obtain them are shown in this diagram.

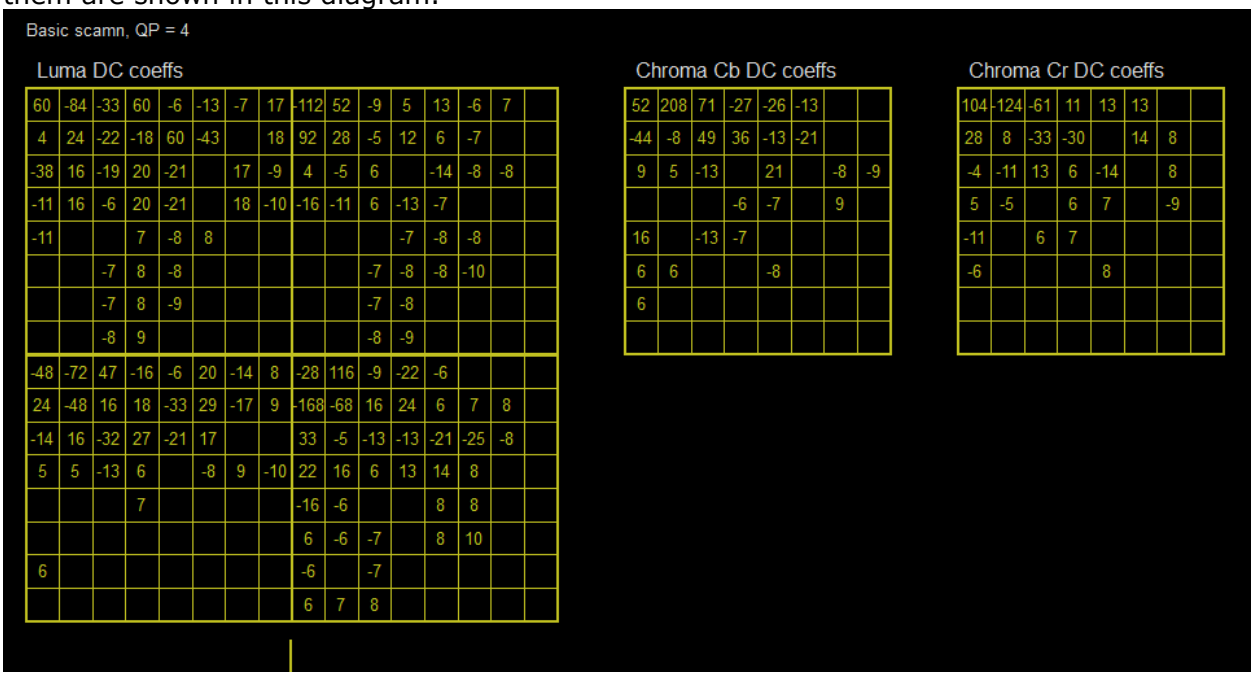

60	-84	-33	60	-6	-13	-7	17	-112	52	-9	5	13	-6	7	
4	24	-22	-18	60	-43		18	92	28	-5	12	6	-7		
-38	16	-19	20	-21		17	-9	4	-5	6		-14	-8	-8	
-11	16	-6	20	-21		18	-10	-16	-11	6	-13	-7			
-11			7	-8	8						-7	-8	-8		
		-7	8	-8						-7	-8	-8	-10		
		-7	8	-9						-7	-8				
		-8	9							-8	-9				
-48	-72	47	-16	-6	20	-14	8	-28	116	-9	-22	-6			
24	-48	16	18	-33	29	-17	9	-168	-68	16	24	6	7	8	
-14	16	-32	27	-21	17			33	-5	-13	-13	-21	-25	-8	
5	5	-13	6		-8	9	-10	22	16	6	13	14	8		
			7					-16	-6			8	8		
								6	-6	-7		8	10		
6								-6		-7					
								6	7	8					

52	208	71	-27	-26	-13		
-44	-8	49	36	-13	-21		
9	5	-13		21		-8	-9
			-6	-7		9	
16		-13	-7				
6	6			-8			
6							

104	-124	-61	11	13	13		
28	8	-33	-30		14	8	
-4	-11	13	6	-14		8	
5	-5		6	7		-9	
-11		6	7				
-6				8			

5. Inverse transform process. This recovered residual values and the transform type used are displayed in this diagram.

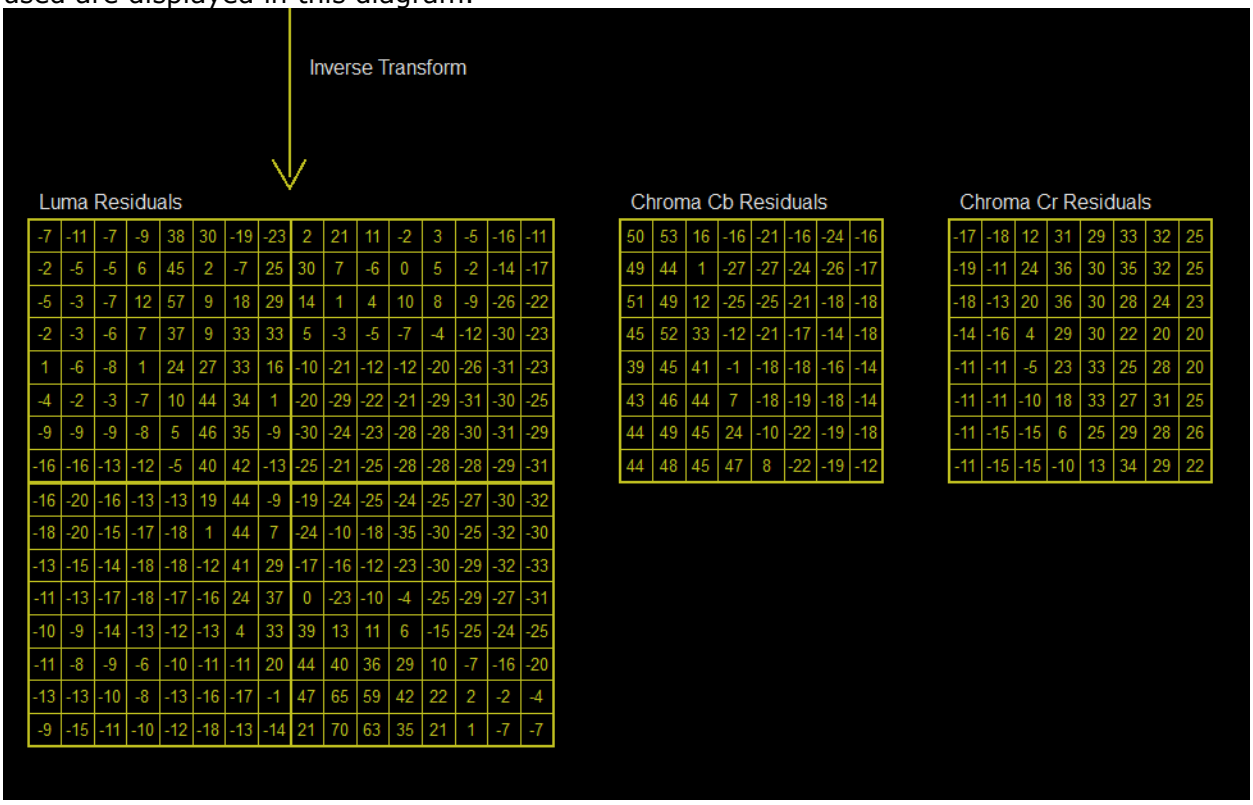

-7	-11	-7	-9	38	30	-19	-23	2	21	11	-2	3	-5	-16	-11
-2	-5	-5	6	45	2	-7	25	30	7	-6	0	5	-2	-14	-17
-5	-3	-7	12	57	9	18	29	14	1	4	10	8	-9	-26	-22
-2	-3	-6	7	37	9	33	33	5	-3	-5	-7	-4	-12	-30	-23
1	-6	-8	1	24	27	33	16	-10	-21	-12	-12	-20	-26	-31	-23
-4	-2	-3	-7	10	44	34	1	-20	-29	-22	-21	-29	-31	-30	-25
-9	-9	-9	-8	5	46	35	-9	-30	-24	-23	-28	-28	-30	-31	-29
-16	-16	-13	-12	-5	40	42	-13	-25	-21	-25	-28	-28	-28	-29	-31
-16	-20	-16	-13	-13	19	44	-9	-19	-24	-25	-24	-25	-27	-30	-32
-18	-20	-15	-17	-18	1	44	7	-24	-10	-18	-35	-30	-25	-32	-30
-13	-15	-14	-18	-18	-12	41	29	-17	-16	-12	-23	-30	-29	-32	-33
-11	-13	-17	-18	-17	-16	24	37	0	-23	-10	-4	-25	-29	-27	-31
-10	-9	-14	-13	-12	-13	4	33	39	13	11	6	-15	-25	-24	-25
-11	-8	-9	-6	-10	-11	-11	20	44	40	36	29	10	-7	-16	-20
-13	-13	-10	-8	-13	-16	-17	-1	47	65	59	42	22	2	-2	-4
-9	-15	-11	-10	-12	-18	-13	-14	21	70	63	35	21	1	-7	-7

50	53	16	-16	-21	-16	-24	-16
49	44	1	-27	-27	-24	-26	-17
51	49	12	-25	-25	-21	-18	-18
45	52	33	-12	-21	-17	-14	-18
39	45	41	-1	-18	-18	-16	-14
43	46	44	7	-18	-19	-18	-14
44	49	45	24	-10	-22	-19	-18
44	48	45	47	8	-22	-19	-12

-17	-18	12	31	29	33	32	25
-19	-11	24	36	30	35	32	25
-18	-13	20	36	30	28	24	23
-14	-16	4	29	30	22	20	20
-11	-11	-5	23	33	25	28	20
-11	-11	-10	18	33	27	31	25
-11	-15	-15	6	25	29	28	26
-11	-15	-15	-10	13	34	29	22

6.2.3 YUV

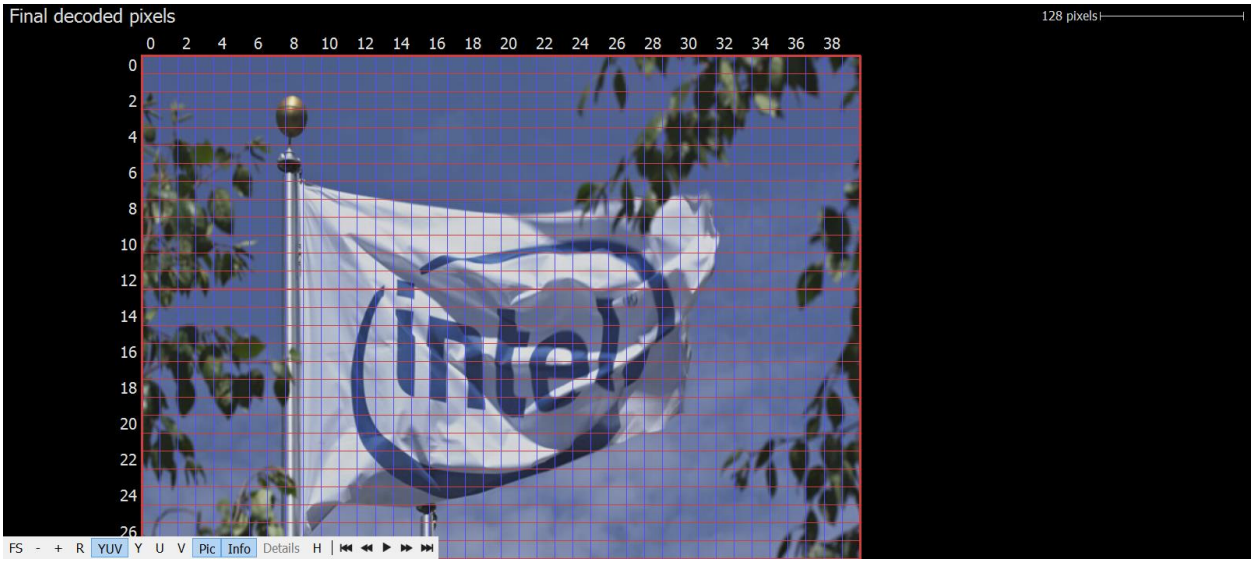

YUV mode allows inspection of the final decoded sample values, without additional overlay data. In this mode, the YUVDiff feature is exposed. This feature allows for comparison of an external decoded YUV file with the decoded bitstream. Supported formats are 4:2:0 planar and NV12. Also, YUV files that are zipped or gzipped may be opened directly without the need to decompress separately. If a zip file contains more than one YUV file, only the first one is used. When a YUV file is loaded, any mismatches will be indicated with a red dot in the CU containing the mismatch. This allows the user to quickly identify the nature of the mismatch which can assist with debug.

When a file is opened, four additional buttons become visible on the lower left of the main window:

- Ref YUV: The original, expected YUV image.
- Debug YUV: The loaded debug YUV file.
- Diff YUV: The delta image. Like in Residual mode, areas with zero delta (i.e. Original and Debug image are identical) are flat gray. Areas where the debug YUV has a lower value are darker, and areas where the debug YUV has a higher value are brighter.
- Reload: Reloads the YUV file, which can be useful as a shortcut to loading via the YUVDiff menu.

6.2.3.1 YUVDiff Menu

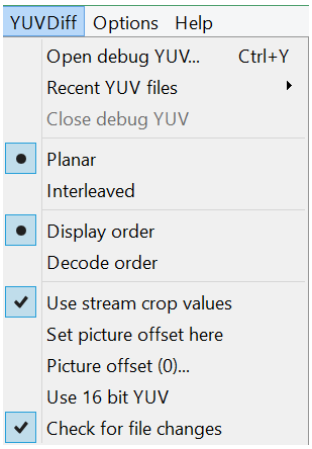

- Open: Brings up a file chooser to open a YUV file
- Close: Closes the currently opened YUV file. Any mismatch indicators will disappear.
- Planar: Causes Intel Video Pro Analyzer to interpret the loaded YUV file as Planar.
- NV12: Causes Intel Video Pro Analyzer to interpret the loaded YUV file as NV12. This option is chosen automatically when the loaded YUV file has the file extension ".nv12".
- Display order: Causes Intel Video Pro Analyzer to use display order picture numbers when determining how far to seek into the YUV file to extract the image to compare with the current picture.
- Decode order: Causes Intel Video Pro Analyzer to use decode order picture numbers when determining how far to seek into the YUV file to extract the image to compare with the current picture.
- Use stream crop values: When checked, YUV files are assumed to contain samples only within the cropping window as defined by the loaded bitstream. Samples outside this window are not compared, and are assumed 0 when viewing the loaded YUV image directly.
- Set picture offset here: Shortcut for setting the picture offset to the current picture number. See bullet below.
- Picture offset: This brings up a dialog allowing the user to enter the picture number of the first picture in the YUV file. For example if a 100-picture bitstream is loaded but the YUV file only contains pictures 80-99, the user would enter 80 to properly line up the YUV file with the decoded bitstream.
- Use 16 bit YUV: this option allow you to load YUV file where both Luma and Chrome written in 16 bit (they could be normalized from actual bitdepth to 16 bit during encoding). This is useful when you have bitdepth changing across the stream. Intel Video Pro Analyzer will shift values to appropriate bitdepth on the current frame.
- Check for file changes: When checked, this option causes Intel Video Pro Analyzer to periodically check if the loaded YUV file has changed on disk since it was last loaded. If the file has indeed changed, a dialog pops up offering a chance to reload the YUV file.

6.2.3.2 YUV Detail Mode

In detail mode the expected YUV values can be inspected. Note that the displayed sample values are the same as those in Loop Filter mode since loop filtering is the last step in the decode process. When a YUV file is loaded, the debug and delta YUV values can be inspected as well. Mismatching values are shown in red.

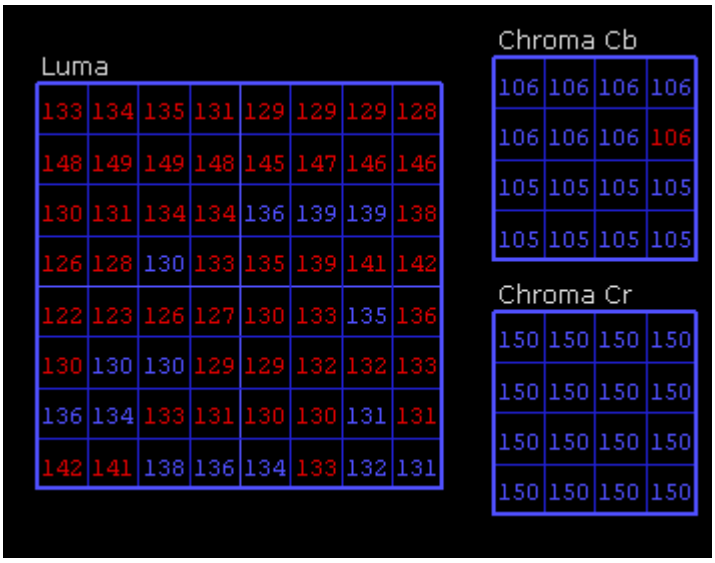

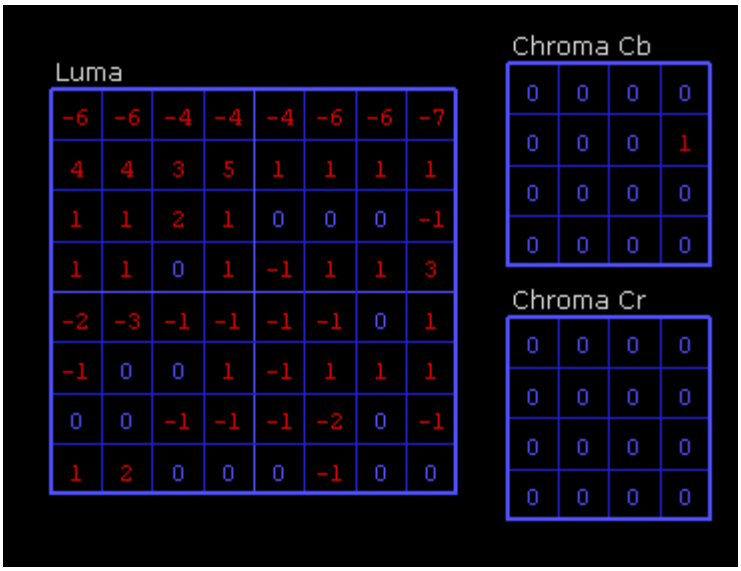

6.2.4 Heat Map

Heat map mode shows visually how the compressed bits of the picture are distributed spatially. Blocks with more bits per pixel are brighter than blocks with less. By default this mode also shows the block boundaries in blue. It may be useful to turn off this overlay using the "Info" button on the lower left of the main panel. Heat Map does not offer a detail mode.

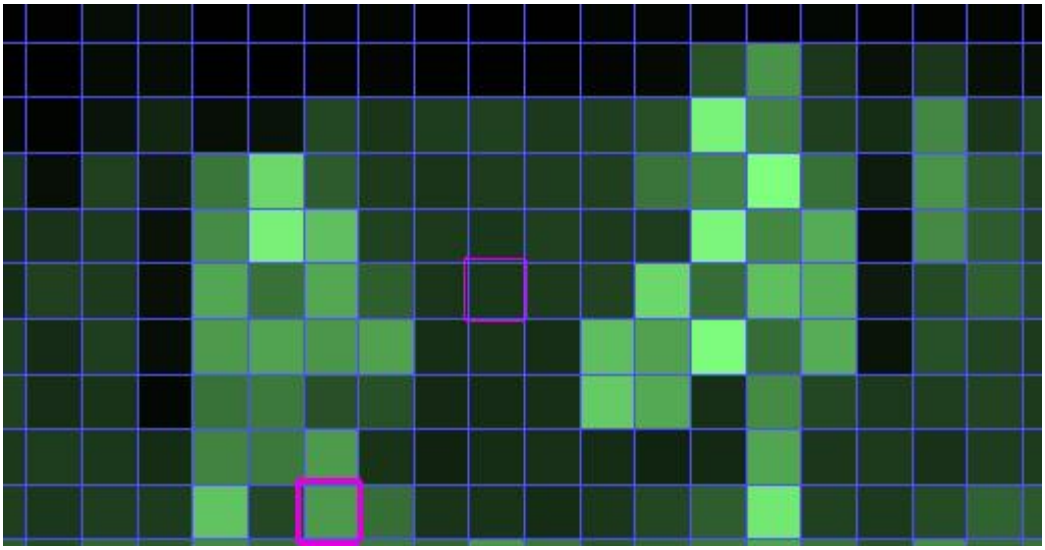

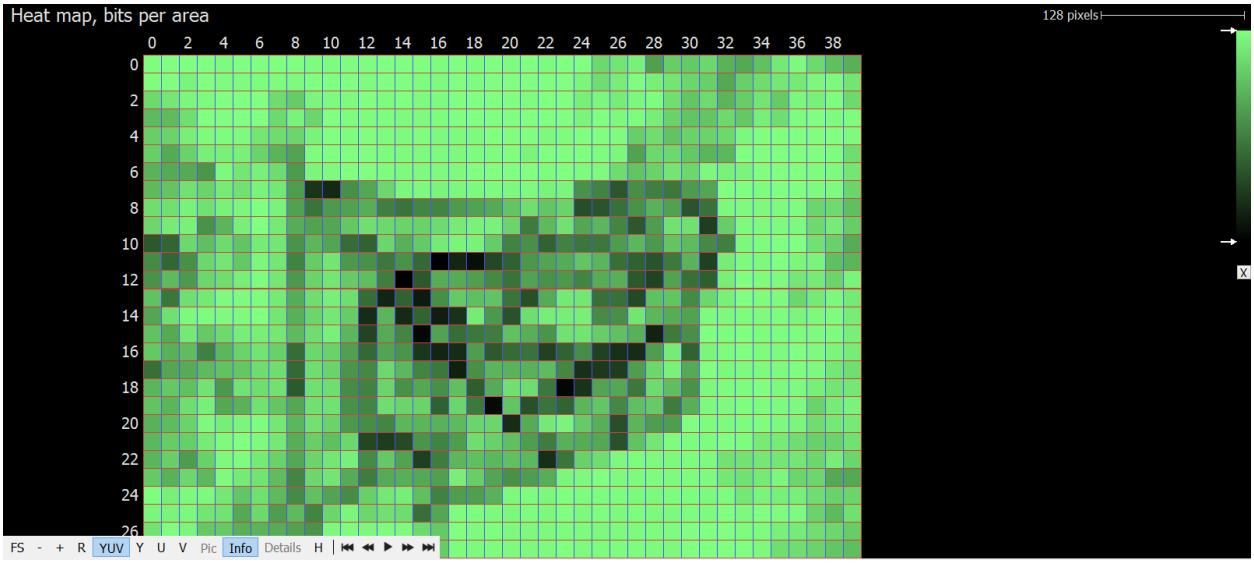

The gradient used to display the heat map may be edited:

- Drag the gradient markers up and down to reposition them.
- Drag a marker away from the gradient to remove it. A red X indicates removal will take place.
- Double-click a marker to change the color.
- Double click the gradient bar to add a new marker.

6.2.5 Qp Map

Efficiency map mode shows visually how QP values are distributed on image..

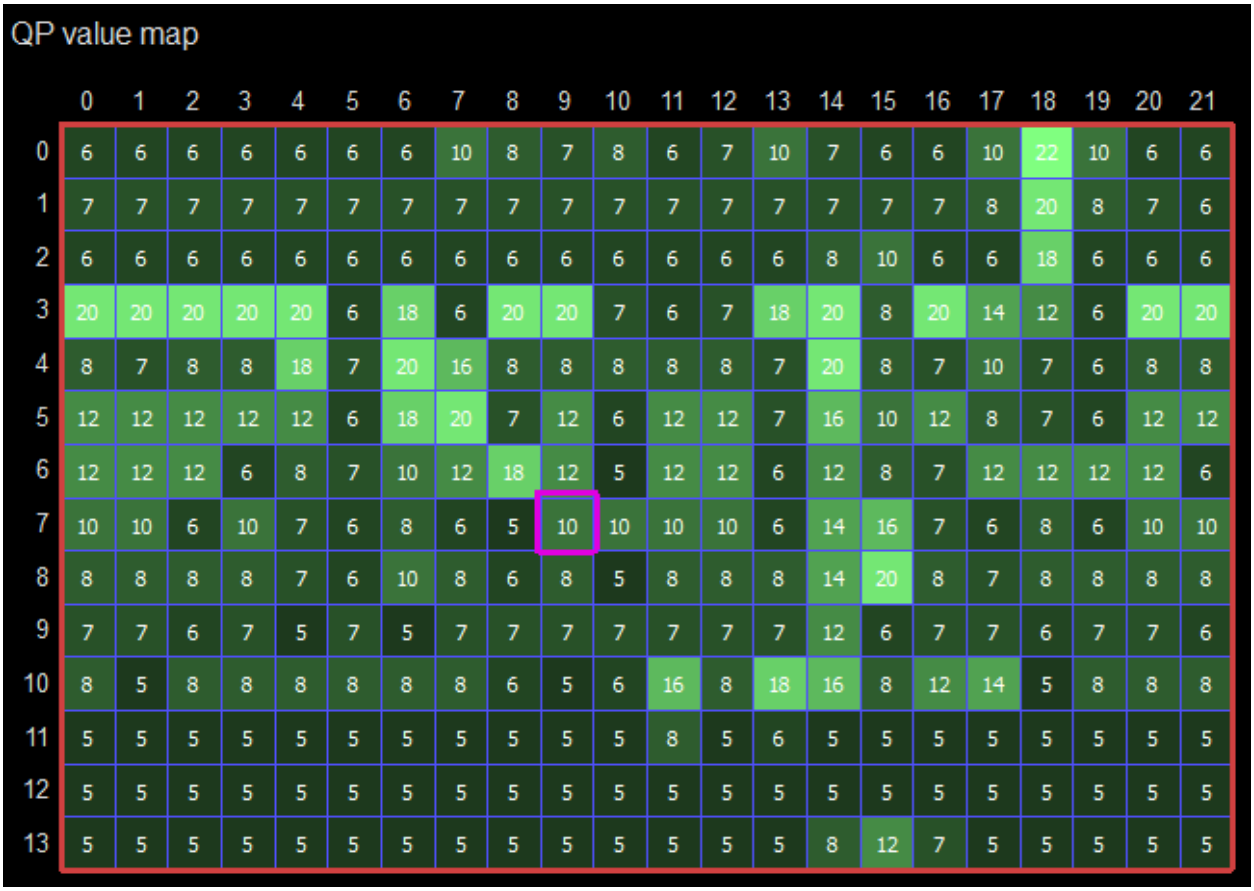

6.2.6 Simple Motion

Simple Motion mode offers a way to quickly view the modes and motion of a picture at a glance. No overlays are present except a colored dot showing information about the prediction mode. Green dots indicate intra blocks, and purple/orange motion vectors for inter blocks are drawn as simple lines. When the motion vector is very small, a dot is drawn instead so that the mode is still easy to identify.

7 Find first difference with Debug YUV

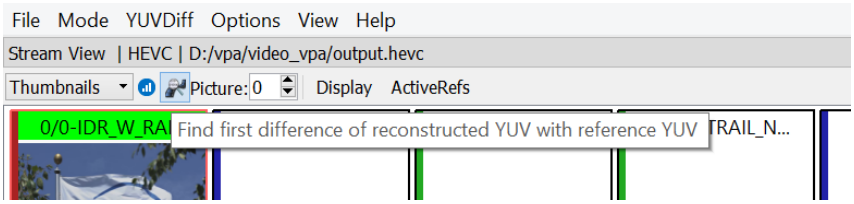

This feature "Find first difference of reconstructed YUV with reference YUV" works for all codecs VP9, HEVC, AVC, MPEG2. By pressing on the button with magnifier icon you start sequential search in decode order for difference between reconstructed YUV and reference YUV that should be preloaded via YUVDiff menu. The search runs till the first difference will be found or to the end of the stream. If you don't want to wait search completion the process can be canceled it at any time with cancel button in info modal dialog. The search always starts from the first picture in decoder order.

8 Options menu

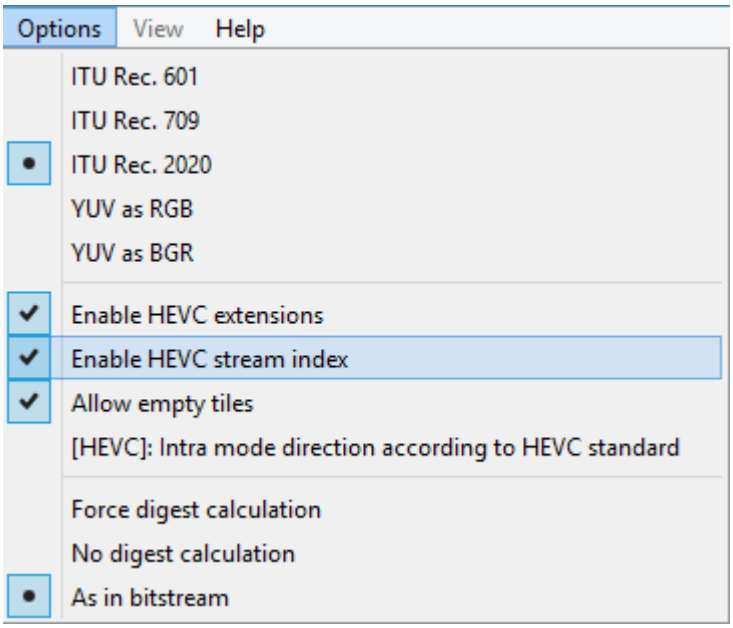

Options menu allows to choose different types of operation modes. The first two options switch between ITU Rec. 601, ITU Rec. 709 and ITU Rec. 2020 color conversion standards used to map YUV data on RGB surface. Intel Video Pro Analyzer uses chosen type of YUV conversion on next picture conversion. Conversions that already has been made will remain unchanged.

For specific coding modes like HEVC screen content coding two additional mode can be used: "YUV as RGB" and "YUV as BGR". In this case decoded YUV values are directly mapped to RGB or BGR values for monitor.

Color conversion with accomplished the following transformation (0.0 <= [Y,R,G,B] <= 1.0) ; (-1.0 < [U,V] < 1.0)), where Y,U,V are normalized accordingly:

R = Y + V*(1-Kr)
G = Y - U*(1-Kb)*Kb/Kg - V*(1-Kr)*Kr/Kg
B = Y + U*(1-Kb),

Where Kr, Kb, Kg have the following values:

	Kr	Kg	Kb
ITU Rec. 601	0.299	0.587	0.114
ITU Rec. 709	0.2126	0.7152	0.0722
ITU Rec. 2020	0.2627	0.6780	0.0593

"Enable HEVC extensions" option enables HEVC RExt extension support (by default). It might be required to disable HEVC extensions support if you decode stream in HEVC v.1 format with wrong headers.

"Enable HEVC stream index" option enables creating and reading index file for HEVC stream. If Intel Video Pro Analyzer finds ".idx" file with the same name as stream file including extension then it will use it as index file to accelerate decoding process. Intel Video Pro Analyzer creates index file during full stream analysis if the option enabled. This feature is

useful for streams with long GOP where analyzer start decoding from the nearest IDR frame and need significant time to decoder target frame. If index file is used Intel Video Pro Analyzer will skip decoding process for frame found in index.

"Allow empty tiles" option allows Intel Video Pro Analyzer decode streams with empty tiles that is not valid case according to standard.

"[HEVC]: Intra mode direction according to HEVC standard" option changes intra mode vectors direction as it stated in the HEVC standard.

There are a set of option for different cases of hash calculations. Those possibilities could save wait time for decoded frame when digest value is not important to have:
"Force digest calculation" – always calculate digest regardless of the digest presents in the bitstream
"No digest calculation" – forces VPA to skip digest calculations
"As in bitstream" – digest calculation is depended on the digest presents in the bitstreams. It only occurs when the bitstream has digest.

9 Dual view

Dual view mode allows you make direct visual comparison between two bitstreams. This feature is useful for codecs fine tuning with slight variation in coding parameters. Bitstreams itself should be coded with the same codec. Intel Video Pro Analyzer considers first opened stream as master stream. Depended stream you open just right after the master one by choosing "Open depended bitstream…" in File menu.

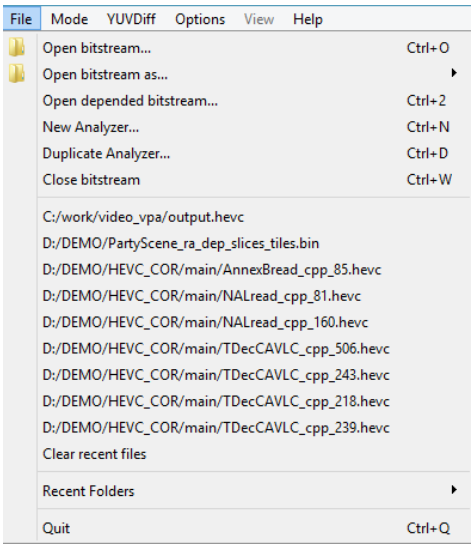

When the depended bitstream is opened it is synchronized with the master bitstream in frame position, chosen mode, offset and scaling factor in the central information window. Also, toolbar buttons in central info window work synchronously with master.

Intel Video Pro Analyzer provides new delta button  in the toolbar of central window.

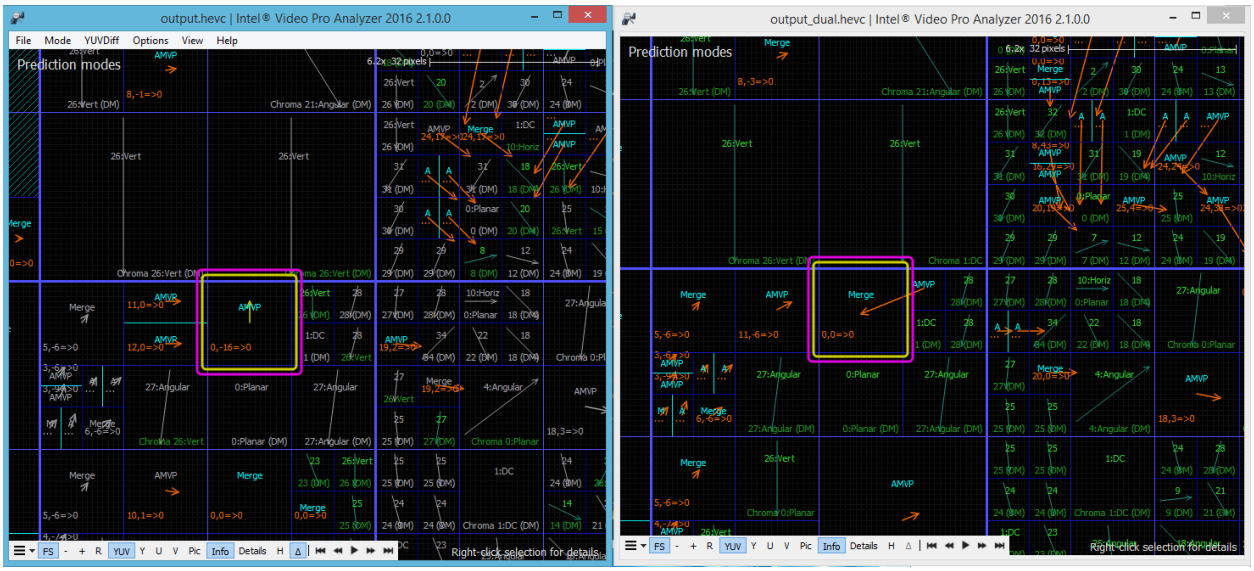

The purpose of this button is to view visual comparision between master and depended bitstreams. When this mode is activated the blocks with the same structure and the same currently visualized modes will be in gray color. While the blocks with different parameters or blocks with different structure on both view still will be colored. In this way it simple for the user to see visually difference in the coding parameters. The delta mode can be used for the following modes: Prediction, Residuals, Reconstruction, Deblocking, SAO, Info overlays.

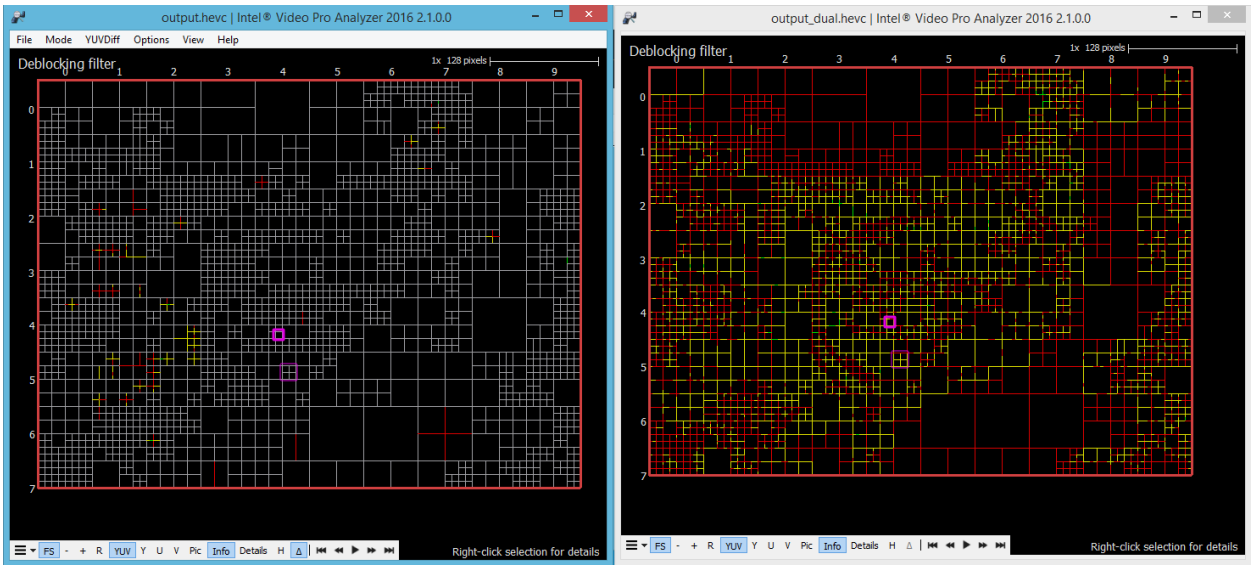

For heat map like modes that display one parameter visual difference is shown pre block in three colors: red – positive values, gray - zero difference, blue – negative difference. So you can easily locate increase or decrease in comparison with depended bitstream.

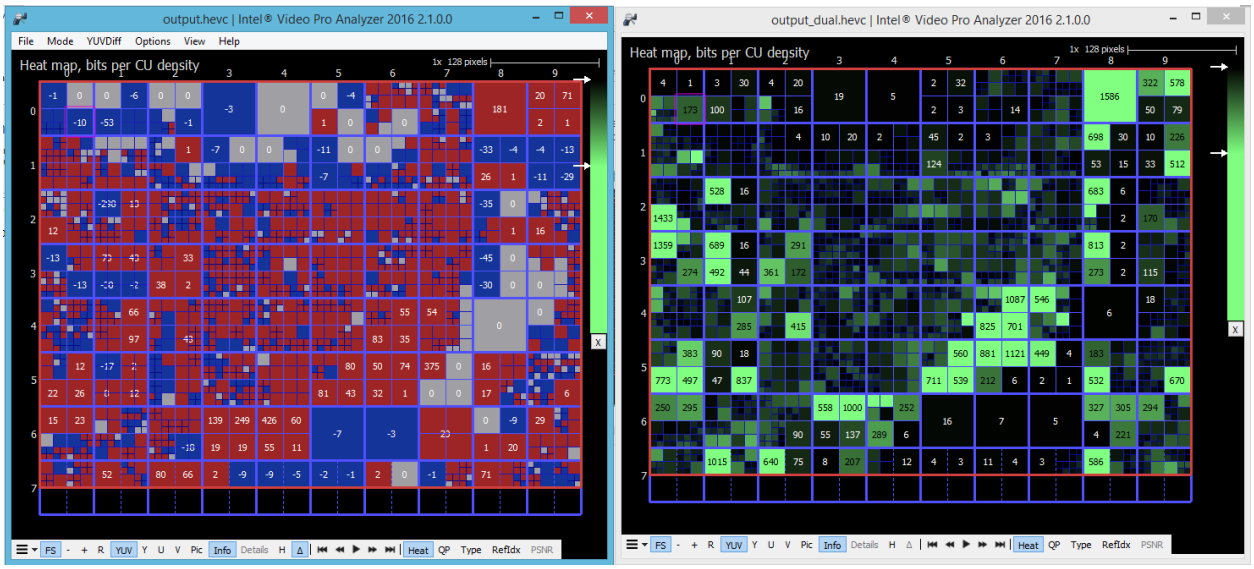

Currently, dual view mode is supported for HEVC and AVC only.

10 Command Line Features

```
Intel(R) Video Pro Analyzer 2017 - version 2.3.0.4470
<file name>: Open the indicated file immediatly after loading the UI.
-h (-help): help for options
-hevc  : force HEVC decoder
-vp9   : force VP9 decoder
-avc   : force AVC decoder
-mpeg2 : force MPEG2 decoder
-o <file name>: output file name (default is "output.yuv"). Must be used with -regress
-n <frame number>: position stream at this frame on startup
-yuv <file name>: debug yuv file name to load on startup (for GUI mode) or debug yuv to measure some metrics (for command line)
-regress: don't bring up the GUI, just decode the stream.
-frames <num_frames>: decode specified number of frames in decode order.  Must be used with -regress
-md5: output pre frame md5 sums
-psnr: calculate PSNR for the debug yuv
-dump: write the decoded sequence to output.yuv (default is decode order). Must be used with -regress
-nocrop: don't crop the output according to stream parameters. Must be used with -regress and -dump
-errors <file name>: write errors to file. Must be used with -regress
-stats: write some stats to stats.csv. Must be used with -regress
-stream_stats <file name>: write whole stream statistics information to specified file
-stream_stats_level <level>: statistics level (should be between 1, 2 or 3, default is 1), the higher the level, the more information
-norm_pix: Normalize stats by pixel area
-norm_bits: Normaluze stats by compressed bits
-percent: Output statistics in percentage
-fast: Skip reconstruction. Must be used with -regress and -stats, and cannot be used with -dump.
-display_order: if used with -dump, the output yuv is written in display order.
-dump_bitdepth <bitdepth>: set output bitdepth of yuv file to bitdepth. Should be in the range [8..16].
-disable_hevc_ext: disable HEVC extensions
-enable_demo_mode: enable demo mode option
-fullversion: output full build information.
```

Intel Video Pro Analyzer has separate console application under Windows* to run in command line mode. GUI application under Windows* also accepts the same command line options but does not provide console output.

Intel Video Pro Analyzer has a few command line options:

- Adding -help prints the available command line options.
- Adding a file name to the command line will cause Intel Video Pro Analyzer to open that file in the UI right after startup, bypassing the file menu. This can be helpful as part of a debug script.
- -regress: This causes Intel Video Pro Analyzer to run without the UI. A file must be specified. This is mainly useful only for HEVC streams with digest SEI messages, as it will display an error message and abort if the digest mismatches. For VP9 and other HEVC streams, the decoder will simply run and exit with no output.

- -dump: Must be used with -regress and a file name. Causes Intel Video Pro Analyzer to decode the specified file and write the decoded output sequence to a file called output.yuv in decode order. The format written is planar YUV. To change the output order to display order, add the command line switch -display_order. By default Intel Video Pro Analyzer crops output according to display size specified by the bitstream. Adding -nocrop will no crop the output to the display size, as specified by the bitstream.
- -stats: Must be used with -regress and a file name. Causes Intel Video Pro Analyzer to write out a file called stats.csv. This file contains the statistics of each picture/frame as comma-separated values. Adding -fast causes the decoder to skip the reconstruction process (intra prediction, motion filtering, deblocking etc.) for faster statistics dumping. -dump may be combined with this switch if -fast is not used. To normalize statistics by area or compressed bits, use -norm_pix and -norm_bits, respectively.
- -o filename: sets output YUV file name. If the name is not specified "output.yuv" will be used.
- -md5: output per frame md5 checksums in batch mode.
- -fullversion: outputs detailed information about current build
- -n frame_number: specify frame to show on startup
- -yuv filename: load specified yuv file on startup
- -dump_bitdepth bitdepth: specify number of bits in output file. For example, if you want save 10 bit output as 16 bit yuv, just put –dump_bitdepth 16 (values will be shifted accordingly). It is useful when you decode file with variable bitdepth.
- -disable_hevc_ext: disables HEVC RExt extension support. This might be required if you decode stream in HEVC v.1 format with wrong headers.
- -errors <file_name> : write list of found errors in the file. Must be used with –regress option.
- -frames : decoder specified number of frames (in stream order). Must be used with –regress option.
- -hevc : force HEVC decoder
- -avc : force avc decoder
- -vp9 : force vp9 decoder
- -mpeg2 : force mpeg2 decoder

11 Miscellaneous

A few other features to keep in mind:

- Intel Video Pro Analyzer saves various settings upon exit. Settings include window size, recent files etc., and are restored on the next launch.
- If the loaded bitstream is modified or deleted while loaded, a dialog will pop-up asking to reload or close. This protects against unintended behavior.

12 Attributions

Qt

Qt licensed under the GNU Lesser General Public License (LGPL) version 2.1. It is appropriate for the development of Qt applications provided you can comply with the terms and conditions of the GNU LGPL version 2.1.
For details, please refer to http://qt-project.org/doc/qt-5/licensing.html for specific details.

Qwt

Qwt License
Version 1.0, January 1, 2003

The Qwt library and included programs are provided under the terms of the GNU LESSER GENERAL PUBLIC LICENSE (LGPL) with the following exceptions:
1. Widgets that are subclassed from Qwt widgets do not constitute a derivative work.
2. Static linking of applications and widgets to the Qwt library does not constitute a derivative work and does not require the author to provide source code for the application or widget, use the shared Qwt libraries, or link their applications or widgets against a user-supplied version of Qwt.
If you link the application or widget to a modified version of Qwt, then the changes to Qwt must be provided under the terms of the LGPL in sections 1, 2, and 4.
3. You do not have to provide a copy of the Qwt license with programs that are linked to the Qwt library, nor do you have to identify the Qwt license in your program or documentation as required by section 6 of the LGPL.
However, programs must still identify their use of Qwt.
The following example statement can be included in user documentation to satisfy this requirement:
[program/widget] is based in part on the work of
the Qwt project (http://qwt.sf.net).