



(51) International Patent Classification:

H04N 19/186 (2014.01) H04N 19/463 (2014.01)
H04N 19/593 (2014.01) H04N 19/93 (2014.01)
H04N 19/70 (2014.01) H04N 19/105 (2014.01)
H04N 19/176 (2014.01)

(21) International Application Number:

PCT/US2016/051745

(22) International Filing Date:

14 September 2016 (14.09.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/218,497 14 September 2015 (14.09.2015) US
62/237,952 6 October 2015 (06.10.2015) US
15/263,686 13 September 2016 (13.09.2016) US

(71) Applicant: QUALCOMM INCORPORATED [US/US];

ATTN: International IP Administration, 5775 Morehouse Drive, San Diego, California 92121-1714 (US).

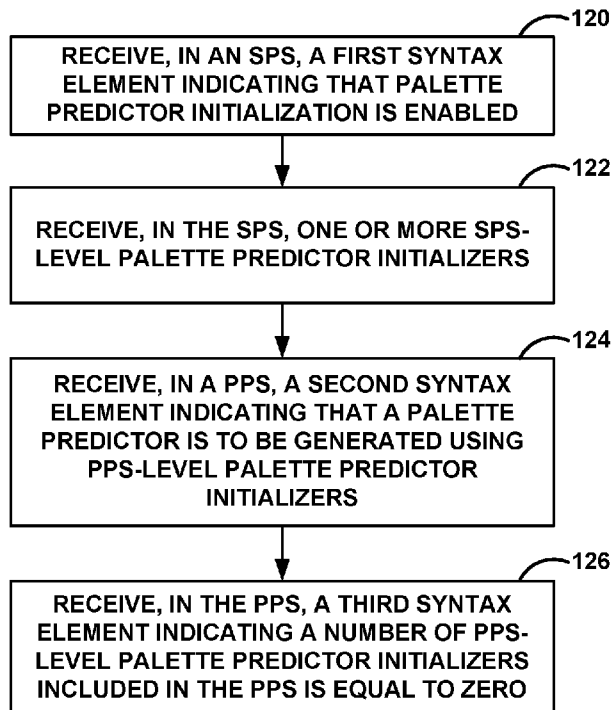
(72) Inventors: SEREGIN, Vadim; 5775 Morehouse Drive, San Diego, California 92121-1714 (US). JOSHI, Rajan Laxman; 5775 Morehouse Drive, San Diego, California 92121-1714 (US). KARCZEWICZ, Marta; 5775 Morehouse Drive, San Diego, California 55125 (US).

(74) Agent: JOSEPH, Jeffrey, R.; Shumaker & Sieffert, P.A., 1625 Radio Drive, Suite 300, Woodbury, Minnesota 55125 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

[Continued on next page]

(54) Title: PALETTE PREDICTOR INITIALIZATION AND MERGE FOR VIDEO CODING



(57) Abstract: A device for decoding video data includes a memory configured to store the video data; and one or more processors configured to receive, in a picture parameter set (PPS), a first syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor initializers; receive, in the PPS, a second syntax element indicating a number of the PPS-level palette predictor initializers included in the PPS is equal to zero; and decode a block of video data based on the first syntax element and the second syntax element.

FIG. 7

WO 2017/048849 A1



(84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE,

SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

PALETTE PREDICTOR INITIALIZATION AND MERGE FOR VIDEO CODING

[0001] This disclosure claims the benefit of U.S. Provisional Patent Application No. 62/218,497 filed 14 September 2015 and U.S. Provisional Patent Application No. 62/237,952 filed 6 October 2015, the entire content each of which is hereby incorporated by reference.

TECHNICAL FIELD

[0002] This disclosure relates to video encoding and video decoding.

BACKGROUND

[0003] Digital video capabilities can be incorporated into a wide range of devices, including digital televisions, digital direct broadcast systems, wireless broadcast systems, personal digital assistants (PDAs), laptop or desktop computers, tablet computers, e-book readers, digital cameras, digital recording devices, digital media players, video gaming devices, video game consoles, cellular or satellite radio telephones, so-called “smart phones,” video teleconferencing devices, video streaming devices, and the like. Digital video devices implement video compression techniques, such as those described in the standards defined by MPEG-2, MPEG-4, ITU-T H.263, ITU-T H.264/MPEG-4, Part 10, Advanced Video Coding (AVC), the recently finalized ITU-T H.265, High Efficiency Video Coding (HEVC) standard, and extensions of such standards. The video devices may transmit, receive, encode, decode, and/or store digital video information more efficiently by implementing such video compression techniques.

[0004] Video compression techniques perform spatial (intra-picture) prediction and/or temporal (inter-picture) prediction to reduce or remove redundancy inherent in video sequences. For block-based video coding, a video slice (i.e., a video frame or a portion of a video frame) may be partitioned into video blocks, which may also be referred to as treeblocks, coding units (CUs) and/or coding nodes. Video blocks in an intra-coded (I) slice of a picture are encoded using spatial prediction with respect to reference samples in neighboring blocks in the same picture. Video blocks in an inter-coded (P or B) slice of a picture may use spatial prediction with respect to reference samples in neighboring

blocks in the same picture or temporal prediction with respect to reference samples in other reference pictures. Pictures may be referred to as frames, and reference pictures may be referred to as reference frames.

[0005] Spatial or temporal prediction results in a predictive block for a block to be coded. Residual data represents pixel differences between the original block to be coded and the predictive block. An inter-coded block is encoded according to a motion vector that points to a block of reference samples forming the predictive block, and the residual data indicating the difference between the coded block and the predictive block. An intra-coded block is encoded according to an intra-coding mode and the residual data. For further compression, the residual data may be transformed from the pixel domain to a transform domain, resulting in residual transform coefficients, which then may be quantized. The quantized transform coefficients, initially arranged in a two-dimensional array, may be scanned in order to produce a one-dimensional vector of transform coefficients, and entropy coding may be applied to achieve even more compression.

SUMMARY

[0006] This disclosure describes techniques for video coding (e.g. video encoding and video decoding) and compression. In particular, this disclosure describes techniques for palette-based coding of video data, and more specifically, this disclosure describes techniques for initializing palette predictors.

[0007] In one example, a method of decoding video data includes receiving, in a picture parameter set (PPS), a first syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor initializers; receiving, in the PPS, a second syntax element indicating a number of the PPS-level palette predictor initializers included in the PPS is equal to zero; and decoding a block of video data based on the first syntax element and the second syntax element.

[0008] In another example, a method of encoding video data includes generating, for inclusion in a picture parameter set (PPS) syntax structure in the bitstream of encoded video data, a first syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor entries; generating, for inclusion in the PPS syntax structure, a second syntax element indicating a number of PPS-level palette predictor

entries included in the PPS is equal to zero; and outputting a bitstream of encoded video data including the first syntax element and the second syntax element.

[0009] According to another example, a device for decoding video data includes a memory configured to store the video data; and one or more processors configured to receive, in a picture parameter set (PPS), a first syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor initializers; receive, in the PPS, a second syntax element indicating a number of the PPS-level palette predictor initializers included in the PPS is equal to zero; and decode a block of video data based on the first syntax element and the second syntax element.

[0010] In another example, a device for encoding video data includes a memory configured to store the video data and one or more processors configured to: generate, for inclusion in a picture parameter set (PPS) syntax structure in the bitstream of encoded video data, a first syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor entries; generate, for inclusion in the PPS syntax structure, a second syntax element indicating a number of PPS-level palette predictor entries included in the PPS is equal to zero; and output a bitstream of encoded video data including the first syntax element and the second syntax element.

[0011] In another example, an apparatus for decoding video data includes means for receiving, in a picture parameter set (PPS), a first syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor initializers; means receiving, in the PPS, a second syntax element indicating a number of the PPS-level palette predictor initializers included in the PPS is equal to zero; and means decoding a block of video data based on the first syntax element and the second syntax element.

[0012] In another example, an apparatus for encoding video data includes means for generating, for inclusion in a picture parameter set (PPS) syntax structure in the bitstream of encoded video data, a first syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor entries; means for generating, for inclusion in the PPS syntax structure, a second syntax element indicating a number of PPS-level palette predictor entries included in the PPS is equal to zero; and means for outputting a bitstream of encoded video data including the first syntax element and the second syntax element.

[0013] In another example, a computer readable storage medium stores instructions that when executed by one or more processors cause the one or more processors to receive, in a picture parameter set (PPS), a first syntax element indicating that a palette predictor

is to be generated using PPS-level palette predictor initializers; receive, in the PPS, a second syntax element indicating a number of the PPS-level palette predictor initializers included in the PPS is equal to zero; and decode a block of video data based on the first syntax element and the second syntax element.

[0014] In another example, a computer readable storage medium stores instructions that when executed by one or more processors cause the one or more processors to generate, for inclusion in a picture parameter set (PPS) syntax structure in the bitstream of encoded video data, a first syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor entries; generate, for inclusion in the PPS syntax structure, a second syntax element indicating a number of PPS-level palette predictor entries included in the PPS is equal to zero; and output a bitstream of encoded video data including the first syntax element and the second syntax element.

[0015] The details of one or more examples of the disclosure are set forth in the accompanying drawings and the description below. Other features, objects, and advantages will be apparent from the description, drawings, and claims.

BRIEF DESCRIPTION OF DRAWINGS

[0016] FIG. 1 is a block diagram illustrating an example video encoding and decoding system that may utilize the techniques described in this disclosure.

[0017] FIG. 2 is a conceptual diagram illustrating an example of determining a palette for coding video data, consistent with techniques of this disclosure.

[0018] FIG. 3 is a conceptual diagram illustrating examples of determining indices to a palette for a video block, consistent with techniques of this disclosure.

[0019] FIG. 4 is a block diagram illustrating an example video encoder that may implement the techniques described in this disclosure.

[0020] FIG. 5 is a block diagram illustrating an example video decoder that may implement the techniques described in this disclosure.

[0021] FIG. 6 is a flowchart illustrating an example operation of a video encoder, in accordance with one or more techniques of this disclosure.

[0022] FIG. 7 is a flowchart illustrating an example operation of a video decoder, in accordance with one or more techniques of this disclosure.

[0023] FIG. 8 is a flowchart illustrating an example operation of a video decoder, in accordance with one or more techniques of this disclosure.

DETAILED DESCRIPTION

[0024] This disclosure describes techniques for video coding (e.g. video encoding and video decoding) and compression. In particular, this disclosure describes techniques for palette-based coding of video data. In palette-based video coding, a video coder (e.g., video encoder or video decoder) derives a palette table (also referred to simply as a “palette”) for a block of pixels, where each entry in the palette table includes color values that are identified by indices into the palette table.

[0025] As part of coding a block in palette mode, a video coder first determines the palette table to be used for the block. Then, palette indices for each pixel (or sample) of the block may be coded to indicate which entry from the palette is to be used to predict or reconstruct the pixel (sample). This disclosure is generally directed to techniques for generating a palette, i.e., determining the palette entries of a palette. More specifically, this disclosure is directed to techniques for initializing a palette predictor. Initializing a palette predictor generally refers to the process of generating the palette predictor for a first block of a picture, slice, tile, or other such grouping of video blocks. As will be explained in greater detail below, the palette predictor for subsequent blocks is typically generated by updating a previously used palette predictor. For a first block, however, a previously used palette predictor is not available, and therefore, the palette predictor needs to be initialized.

[0026] In order to reduce the bits needed for signaling palette entries, a video decoder may utilize a palette predictor for determining the new palette entries. The palette predictor may, for example, include entries from a previously used palette. In some implementations, the palette predictor may be set equal to a most recently used palette and include all entries of the most recently used palette, while in other implementations the palette predictor may include fewer than all entries from the most recently used palette. The palette predictor may also include entries from other previously used palettes besides just the most recently used palette. The palette predictor may be the same size as the palettes used for coding blocks or may be larger or smaller than the palettes used for coding blocks.

[0027] To generate a palette for a block of video data utilizing the palette predictor, a video decoder may, for example, receive a one-bit flag for each entry of the palette predictor, with a first value for the flag indicating the associated entry of the palette predictor is included in the palette and a second value for the flag indicating the

associated entry of the palette predictor is not included in the palette. In instances where the palette predictor is larger than the palette used for a block of video data, then the video decoder may stop receiving flags once a maximum size for the palette is reached.

[0028] Some entries of the palette may be directly signaled instead of determined using the palette predictor. For such entries, the video decoder may, for example, receive three separate m-bit values indicating the color values for the luma and two chroma components associated with the entry, where m represents the bit depth of the video data. As those palette entries derived from the palette predictor only require a one-bit flag compared to the multiple m-bit values needed for directly signaled entries, signaling some or all palette entries using the palette predictor may reduce the number of bits needed to signal the entries of a new palette, and as a result, improve the overall coding efficiency of palette mode coding.

[0029] As explained above, in many instances, the palette predictor for one block is determined based on the palette used to code one or more already coded blocks. In some coding scenarios, such as when coding the first coding tree unit in a slice or tile, however, the palette of a previously coded block may be unavailable, and thus, a palette predictor cannot be generated using entries of previously used palettes. In such an instance, palette predictor initializers may be signaled in a sequence parameter set (SPS) and/or a picture parameter set (PPS). The palette predictor initializers are values that may be used to generate a palette predictor when a previously used palette is not available.

[0030] A PPS generally refers to a syntax structure containing syntax elements that apply to zero or more entire coded pictures as determined by a syntax element found in each slice segment header. An SPS generally refers to a syntax structure containing syntax elements that apply to zero or more entire CVSs as determined by the content of a syntax element found in the PPS referred to by a syntax element found in each slice segment header. Thus, an SPS is generally considered to be a higher level syntax structure than a PPS, meaning the syntax elements included in the SPS generally change less frequently and apply to a larger portion of video compared to the syntax elements included in the PPS.

[0031] In existing implementations of palette mode coding, if palette predictor initialization is enabled at the SPS-level, then PPS-level syntax only enables a video decoder to select between generating a palette predictor using SPS-level palette

predictor entries (also referred to herein as SPS-level palette predictor initializers) and generating a palette predictor using PPS-level palette predictor entries (also referred to herein as PPS-level palette predictor initializers). Existing implementations of palette mode coding do not allow palette predictor initialization to be disabled at the PPS level once enabled at the SPS level. This disclosure introduces techniques for disabling palette predictor initialization at the PPS level by setting the size of a palette predictor to zero.

[0032] Furthermore, in existing implementations of palette mode coding, a video decoder initializes a palette predictor based on either the SPS palette predictor initializers or the PPS palette predictor initializers. According to the techniques of this disclosure, a video decoder may initialize a palette predictor based on both the SPS palette predictor initializers and the PPS palette predictor initializers. For example, a palette predictor may be initialized by including the SPS palette predictor entries followed by the PPS palette predictor entries, or vice versa. In other example implementations, a video decoder may use only some and not all of the SPS palette predictor initializers. For instance, when initializing a palette predictor, a reuse flag may be signaled per entry to indicate which entries of the SPS palette predictor entries are to be reused for palette predictor initialization.

[0033] A video coder may determine, e.g., initialize, a palette predictor for a first block (e.g, a coding unit or a coding tree unit in HEVC or different type of block in a different coding standard) in a slice or tile. A palette table for the block may then be determined from the palette predictor. After the palette predictor is initialized for the first block, then the palette predictor for the first block may be updated to determine a new palette predictor for a second block.

[0034] FIG. 1 is a block diagram illustrating an example video encoding and decoding system 10 that may utilize the techniques described in this disclosure. As shown in FIG. 1, system 10 includes a source device 12 that generates encoded video data to be decoded at a later time by a destination device 14. Source device 12 and destination device 14 may comprise any of a wide range of devices, including desktop computers, notebook (i.e., laptop) computers, tablet computers, set-top boxes, telephone handsets such as so-called “smart” phones, so-called “smart” pads, televisions, cameras, display devices, digital media players, video gaming consoles, video streaming device, or the like. In some cases, source device 12 and destination device 14 may be equipped for wireless communication.

[0035] Destination device 14 may receive the encoded video data to be decoded via a link 16. Link 16 may comprise any type of medium or device capable of moving the encoded video data from source device 12 to destination device 14. In one example, link 16 may comprise a communication medium to enable source device 12 to transmit encoded video data directly to destination device 14 in real-time. The encoded video data may be modulated according to a communication standard, such as a wireless communication protocol, and transmitted to destination device 14. The communication medium may comprise any wireless or wired communication medium, such as a radio frequency (RF) spectrum or one or more physical transmission lines. The communication medium may form part of a packet-based network, such as a local area network, a wide-area network, or a global network such as the Internet. The communication medium may include routers, switches, base stations, or any other equipment that may be useful to facilitate communication from source device 12 to destination device 14.

[0036] Alternatively, encoded data may be output from output interface 22 to a storage device 26. Similarly, encoded data may be accessed from storage device 26 by input interface. Storage device 26 may include any of a variety of distributed or locally accessed data storage media such as a hard drive, Blu-ray discs, DVDs, CD-ROMs, flash memory, volatile or non-volatile memory, or any other suitable digital storage media for storing encoded video data. In a further example, storage device 26 may correspond to a file server or another intermediate storage device that may hold the encoded video generated by source device 12. Destination device 14 may access stored video data from storage device 26 via streaming or download. The file server may be any type of server capable of storing encoded video data and transmitting that encoded video data to the destination device 14. Example file servers include a web server (e.g., for a website), an FTP server, network attached storage (NAS) devices, or a local disk drive. Destination device 14 may access the encoded video data through any standard data connection, including an Internet connection. This may include a wireless channel (e.g., a Wi-Fi connection), a wired connection (e.g., DSL, cable modem, etc.), or a combination of both that is suitable for accessing encoded video data stored on a file server. The transmission of encoded video data from storage device 26 may be a streaming transmission, a download transmission, or a combination of both.

[0037] The techniques of this disclosure are not necessarily limited to wireless applications or settings. The techniques may be applied to video coding in support of

any of a variety of multimedia applications, such as over-the-air television broadcasts, cable television transmissions, satellite television transmissions, streaming video transmissions, e.g., via the Internet, encoding of digital video for storage on a data storage medium, decoding of digital video stored on a data storage medium, or other applications. In some examples, system 10 may be configured to support one-way or two-way video transmission to support applications such as video streaming, video playback, video broadcasting, and/or video telephony.

[0038] In the example of FIG. 1, source device 12 includes a video source 18, video encoder 20 and an output interface 22. In some cases, output interface 22 may include a modulator/demodulator (modem) and/or a transmitter. In source device 12, video source 18 may include a source such as a video capture device, e.g., a video camera, a video archive containing previously captured video, a video feed interface to receive video from a video content provider, and/or a computer graphics system for generating computer graphics data as the source video, or a combination of such sources. As one example, if video source 18 is a video camera, source device 12 and destination device 14 may form so-called camera phones or video phones. However, the techniques described in this disclosure may be applicable to video coding in general, and may be applied to wireless and/or wired applications.

[0039] The captured, pre-captured, or computer-generated video may be encoded by video encoder 20. The encoded video data may be transmitted directly to destination device 14 via output interface 22 of source device 12. The encoded video data may also (or alternatively) be stored onto storage device 26 for later access by destination device 14 or other devices, for decoding and/or playback.

[0040] Destination device 14 includes an input interface 28, a video decoder 30, and a display device 32. In some cases, input interface 28 may include a receiver and/or a modem. Input interface 28 of destination device 14 receives the encoded video data over link 16. The encoded video data communicated over link 16, or provided on storage device 26, may include a variety of syntax elements generated by video encoder 20 for use by a video decoder, such as video decoder 30, in decoding the video data. Such syntax elements may be included with the encoded video data transmitted on a communication medium, stored on a storage medium, or stored a file server.

[0041] Display device 32 may be integrated with, or external to, destination device 14. In some examples, destination device 14 may include an integrated display device and also be configured to interface with an external display device. In other examples,

destination device 14 may be a display device. In general, display device 32 displays the decoded video data to a user, and may comprise any of a variety of display devices such as a liquid crystal display (LCD), a plasma display, an organic light emitting diode (OLED) display, or another type of display device.

[0042] Video encoder 20 and video decoder 30 may operate according to a video compression standard, such as the recently finalized High Efficiency Video Coding (HEVC) standard, and may conform to the HEVC Test Model (HM). Alternatively, video encoder 20 and video decoder 30 may operate according to other proprietary or industry standards, such as the ITU-T H.264 standard, alternatively referred to as MPEG-4, Part 10, Advanced Video Coding (AVC), or extensions of such standards. The techniques of this disclosure, however, are not limited to any particular coding standard. Other examples of video compression standards include MPEG-2 and ITU-T H.263. Video coding standards include ITU-T H.261, ISO/IEC MPEG-1 Visual, ITU-T H.262 or ISO/IEC MPEG-2 Visual, ITU-T H.263, ISO/IEC MPEG-4 Visual, ITU-T H.264 (also known as ISO/IEC MPEG-4 AVC), and ITU-T H.265/HEVC including its Scalable Video Coding (SVC) and Multiview Video Coding (MVC) extensions.

[0043] Recently, investigation of new coding tools for screen-content material such as text and graphics with motion was requested, and technologies that improve the coding efficiency for screen content have been proposed. This activity is currently ongoing and is expected to be completed as the H.265/HEVC screen content coding (SCC) extension. JCTVC-U1005, Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T SG 16 WP 3 and ISO/IEC JTC 1/SC 29/WG 11, 21st Meeting: Warsaw, PL, 19 June – 26 June 2015, sets forth HEVC SCC Working Draft 4 (WD4) JCTVC-U1005 and is available at:

http://phenix.int-evry.fr/jct/doc_end_user/documents/21_Warsaw/wg11/JCTVC-U1005-v2.zip, the entire content of which is hereby incorporated by reference.

[0044] Techniques of this disclosure may utilize HEVC terminology for ease of explanation. It should not be assumed, however, that the techniques of this disclosure are limited to HEVC, and in fact, it is explicitly contemplated that the techniques of this disclosure may be implemented in successor standards to HEVC and its extensions, and other future standards or processes.

[0045] Although not shown in FIG. 1, in some aspects, video encoder 20 and video decoder 30 may each be integrated with an audio encoder and decoder, and may include appropriate MUX-DEMUX units, or other hardware and software, to handle encoding

of both audio and video in a common data stream or separate data streams. If applicable, in some examples, MUX-DEMUX units may conform to the ITU H.223 multiplexer protocol, or other protocols such as the user datagram protocol (UDP).

[0046] Video encoder 20 and video decoder 30 each may be implemented as any of a variety of suitable encoder circuitry, such as one or more microprocessors, digital signal processors (DSPs), application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs), discrete logic, software, hardware, firmware or any combinations thereof. When the techniques are implemented partially in software, a device may store instructions for the software in a suitable, non-transitory computer-readable medium and execute the instructions in hardware using one or more processors to perform the techniques of this disclosure. Each of video encoder 20 and video decoder 30 may be included in one or more encoders or decoders, either of which may be integrated as part of a combined encoder/decoder (CODEC) in a respective device.

[0047] As introduced above, the JCT-VC has recently finalized development of the HEVC standard. The HEVC standardization efforts were based on an evolving model of a video coding device referred to as the HEVC Test Model (HM). The HM presumes several additional capabilities of video coding devices relative to existing devices according to, e.g., ITU-T H.264/AVC. For example, whereas H.264 provides nine intra-prediction encoding modes, the HM may provide as many as thirty-five intra-prediction encoding modes.

[0048] In HEVC and other video coding specifications, a video sequence typically includes a series of pictures. Pictures may also be referred to as “frames.” A picture may include three sample arrays, denoted S_L , S_{Cb} , and S_{Cr} . S_L is a two-dimensional array (i.e., a block) of luma samples. S_{Cb} is a two-dimensional array of Cb chrominance samples. S_{Cr} is a two-dimensional array of Cr chrominance samples. Chrominance samples may also be referred to herein as “chroma” samples. In other instances, a picture may be monochrome and may only include an array of luma samples.

[0049] To generate an encoded representation of a picture, video encoder 20 may generate a set of coding tree units (CTUs). Each of the CTUs may comprise a coding tree block of luma samples, two corresponding coding tree blocks of chroma samples, and syntax structures used to code the samples of the coding tree blocks. In monochrome pictures or pictures having three separate color planes, a CTU may comprise a single coding tree block and syntax structures used to code the samples of the coding tree block. A coding tree block may be an $N \times N$ block of samples. A CTU

may also be referred to as a “tree block” or a “largest coding unit” (LCU). The CTUs of HEVC may be broadly analogous to the macroblocks of other standards, such as H.264/AVC. However, a CTU is not necessarily limited to a particular size and may include one or more coding units (CUs). A slice may include an integer number of CTUs ordered consecutively in a raster scan order.

[0050] To generate a coded CTU, video encoder 20 may recursively perform quad-tree partitioning on the coding tree blocks of a CTU to divide the coding tree blocks into coding blocks, hence the name “coding tree units.” A coding block may be an NxN block of samples. A CU may comprise a coding block of luma samples and two corresponding coding blocks of chroma samples of a picture that has a luma sample array, a Cb sample array, and a Cr sample array, and syntax structures used to code the samples of the coding blocks. In monochrome pictures or pictures having three separate color planes, a CU may comprise a single coding block and syntax structures used to code the samples of the coding block.

[0051] Video encoder 20 may partition a coding block of a CU into one or more prediction blocks. A prediction block is a rectangular (i.e., square or non-square) block of samples on which the same prediction is applied. A prediction unit (PU) of a CU may comprise a prediction block of luma samples, two corresponding prediction blocks of chroma samples, and syntax structures used to predict the prediction blocks. In monochrome pictures or pictures having three separate color planes, a PU may comprise a single prediction block and syntax structures used to predict the prediction block. Video encoder 20 may generate predictive luma, Cb, and Cr blocks for luma, Cb, and Cr prediction blocks of each PU of the CU.

[0052] Video encoder 20 may use intra prediction or inter prediction to generate the predictive blocks for a PU. If video encoder 20 uses intra prediction to generate the predictive blocks of a PU, video encoder 20 may generate the predictive blocks of the PU based on decoded samples of the picture associated with the PU. If video encoder 20 uses inter prediction to generate the predictive blocks of a PU, video encoder 20 may generate the predictive blocks of the PU based on decoded samples of one or more pictures other than the picture associated with the PU.

[0053] After video encoder 20 generates predictive luma, Cb, and Cr blocks for one or more PUs of a CU, video encoder 20 may generate a luma residual block for the CU. Each sample in the CU’s luma residual block indicates a difference between a luma sample in one of the CU’s predictive luma blocks and a corresponding sample in the

CU's original luma coding block. In addition, video encoder 20 may generate a Cb residual block for the CU. Each sample in the CU's Cb residual block may indicate a difference between a Cb sample in one of the CU's predictive Cb blocks and a corresponding sample in the CU's original Cb coding block. Video encoder 20 may also generate a Cr residual block for the CU. Each sample in the CU's Cr residual block may indicate a difference between a Cr sample in one of the CU's predictive Cr blocks and a corresponding sample in the CU's original Cr coding block.

[0054] Furthermore, video encoder 20 may use quad-tree partitioning to decompose the luma, Cb, and Cr residual blocks of a CU into one or more luma, Cb, and Cr transform blocks. A transform block is a rectangular (e.g., square or non-square) block of samples on which the same transform is applied. A transform unit (TU) of a CU may comprise a transform block of luma samples, two corresponding transform blocks of chroma samples, and syntax structures used to transform the transform block samples. Thus, each TU of a CU may be associated with a luma transform block, a Cb transform block, and a Cr transform block. The luma transform block associated with the TU may be a sub-block of the CU's luma residual block. The Cb transform block may be a sub-block of the CU's Cb residual block. The Cr transform block may be a sub-block of the CU's Cr residual block. In monochrome pictures or pictures having three separate color planes, a TU may comprise a single transform block and syntax structures used to transform the samples of the transform block.

[0055] Video encoder 20 may apply one or more transforms to a luma transform block of a TU to generate a luma coefficient block for the TU. A coefficient block may be a two-dimensional array of transform coefficients. A transform coefficient may be a scalar quantity. Video encoder 20 may apply one or more transforms to a Cb transform block of a TU to generate a Cb coefficient block for the TU. Video encoder 20 may apply one or more transforms to a Cr transform block of a TU to generate a Cr coefficient block for the TU.

[0056] After generating a coefficient block (e.g., a luma coefficient block, a Cb coefficient block or a Cr coefficient block), video encoder 20 may quantize the coefficient block. Quantization generally refers to a process in which transform coefficients are quantized to possibly reduce the amount of data used to represent the transform coefficients, providing further compression. After video encoder 20 quantizes a coefficient block, video encoder 20 may entropy encode syntax elements indicating the quantized transform coefficients. For example, video encoder 20 may perform

Context-Adaptive Binary Arithmetic Coding (CABAC) on the syntax elements indicating the quantized transform coefficients.

[0057] Video encoder 20 may output a bitstream that includes a sequence of bits that forms a representation of coded pictures and associated data. The bitstream may comprise a sequence of NAL units. A NAL unit is a syntax structure containing an indication of the type of data in the NAL unit and bytes containing that data in the form of a RBSP interspersed as necessary with emulation prevention bits. Each of the NAL units includes a NAL unit header and encapsulates a RBSP. The NAL unit header may include a syntax element that indicates a NAL unit type code. The NAL unit type code specified by the NAL unit header of a NAL unit indicates the type of the NAL unit. A RBSP may be a syntax structure containing an integer number of bytes that is encapsulated within a NAL unit. In some instances, an RBSP includes zero bits.

[0058] Different types of NAL units may encapsulate different types of RBSPs. For example, a first type of NAL unit may encapsulate an RBSP for a PPS, a second type of NAL unit may encapsulate an RBSP for a coded slice, a third type of NAL unit may encapsulate an RBSP for SEI messages, and so on. NAL units that encapsulate RBSPs for video coding data (as opposed to RBSPs for parameter sets and SEI messages) may be referred to as VCL NAL units.

[0059] Video decoder 30 may receive a bitstream generated by video encoder 20. In addition, video decoder 30 may parse the bitstream to obtain syntax elements from the bitstream. Video decoder 30 may reconstruct the pictures of the video data based at least in part on the syntax elements obtained from the bitstream. The process to reconstruct the video data may be generally reciprocal to the process performed by video encoder 20. In addition, video decoder 30 may inverse quantize coefficient blocks associated with TUs of a current CU. Video decoder 30 may perform inverse transforms on the coefficient blocks to reconstruct transform blocks associated with the TUs of the current CU. Video decoder 30 may reconstruct the coding blocks of the current CU by adding the samples of the predictive blocks for PUs of the current CU to corresponding samples of the transform blocks of the TUs of the current CU. By reconstructing the coding blocks for each CU of a picture, video decoder 30 may reconstruct the picture.

[0060] FIG. 2 is a conceptual diagram illustrating an example of determining a palette for coding video data. The example of FIG. 2 includes a picture 178 having a first block 180 that is associated with a first set of palettes (i.e., first palettes 184) and a second

block 188 that is associated with a second set of palettes (i.e., second palettes 192). As described in greater detail below, second palettes 192 may be determined based on first palettes 184. Picture 178 also includes block 196 coded with an intra-prediction coding mode and block 200 that is coded with an inter-prediction coding mode.

[0061] The techniques of FIG. 2 are described in the context of video encoder 20 (FIG. 1 and FIG. 2) and video decoder 30 (FIG. 1 and FIG. 3). The palette-based coding techniques may, for example, be used as a coding mode for a block of video data. The various blocks described with respect to FIG. 2 may, for example, correspond to CTUs, CUs, PUs, or TUs as understood in the context of HEVC, but the blocks described in FIG. 2 are not limited to the block structure of any particular coding standard and may be compatible with future coding standards that utilize block formats different than HEVC.

[0062] In general, a palette refers to a number of pixel values that are dominant and/or representative for a block currently being coded (e.g., block 188 in the example of FIG. 2). First palettes 184 and second palettes 192 are shown as including multiple palettes. In some examples, a video coder (such as video encoder 20 or video decoder 30) may code palettes separately for each color component of a block. For example, video encoder 20 may encode a palette for a luma (Y) component of a block, another palette for a chroma (U) component of the block, and yet another palette for the chroma (V) component of the block. In this example, entries of the Y palette may represent Y values of pixels of the block, entries of the U palette may represent U values of pixels of the block, and entries of the V palette may represent V values of pixels of the block.

[0063] In other examples, video encoder 20 may encode a single palette for all color components of a block. In this example, video encoder 20 may encode a palette having an i -th entry that is a triple value, including Y_i , U_i , and V_i . In this case, the palette includes values for each of the components of the pixels. Accordingly, the representation of first palettes 184 and second palettes 192 as a set of palettes having multiple individual palettes is merely one example and not intended to be limiting.

[0064] In the example of FIG. 2, each of first palettes 184 includes three entries 202-206 having entry index value 1, entry index value 2, and entry index value 3, respectively. Entries 202-206 relate the index values to pixel values including pixel value A, pixel value B, and pixel value C, respectively. It should be noted that each of first palettes 184 do not actually include the indices and column headers, but only

include the pixel values A, B and C and the indices are used to identify the entries in the palette.

[0065] As described herein, rather than directly coding the actual pixel values of first block 180, a video coder (such as video encoder 20 or video decoder 30) may use palette-based coding to code the pixels of the block using the indices 1-3. That is, for each pixel position of first block 180, video encoder 20 may encode an index value for the pixel, where the index value is associated with a pixel value in one or more of first palettes 184. Video decoder 30 may obtain the index values from a bitstream and may reconstruct the pixel values using the index values and one or more of first palettes 184. In other words, for each respective index value for a block, video decoder 30 may determine an entry in one of first palettes 184. Video decoder 30 may replace the respective index value in the block with the pixel value specified by the determined entry in the palette. Video encoder 20 may transmit first palettes 184 in an encoded video data bitstream for use by video decoder 30 in palette-based decoding. In general, one or more palettes may be transmitted for each block or may be shared among different blocks.

[0066] According to aspects of this disclosure, video encoder 20 and video decoder 30 may determine second palettes 192 based on one or more palette predictors. The palette predictors may, for example, includes some or all of the entries of first palettes 184 and possibly include entries from other palettes as well. In some examples, the palette predictors may be generated using a first-in-first-out methodology, where upon adding the entries of first palettes 184 to the palette predictors, the oldest entries currently in the palette predictors are expunged in order to keep the palette predictors at or below a maximum size. In other examples, the palette predictor may be updated and/or maintained using different techniques. The techniques of this disclosure are generally directed to initializing a palette predictor and may be used in conjunction with numerous different techniques for updating and maintaining palette predictors.

[0067] Although the techniques of this disclosure are generally directed to generating a palette based on a palette predictor, and more specifically, directed to techniques for controlling the initializing of palette predictors, the techniques of this disclosure may be used in conjunction with other techniques for generating palettes. As one example, video encoder 20 may encode a `pred_palette_flag` for each block (including, as an example, second block 188) to indicate whether the palette for the block is predicted from one or more palettes associated with one or more other blocks, such as neighboring

block (spatially or based on scan order) or the most frequent samples of a causal neighbor. For example, when the value of such a flag is equal to one, video decoder 30 may determine that second palettes 192 for second block 188 are predicted from one or more already decoded palettes and therefore no new palettes for second block 188 are included in a bitstream containing the `pred_palette_flag`. When such a flag is equal to zero, video decoder 30 may determine that second palettes 192 for second block 188 are included in the bitstream as a new palette. In some examples, `pred_palette_flag` may be separately coded for each different color component of a block (e.g., three flags, one for Y, one for U, and one for V, for a block in YUV video). In other examples, a single `pred_palette_flag` may be coded for all color components of a block.

[0068] In the example above, the `pred_palette_flag` is signaled per-block to indicate whether any of the entries of the palette for the current block are predicted. This means that second palettes 192 are identical to first palettes 184 and no additional information is signaled. In other examples, one or more syntax elements may be signaled on a per-entry basis. That is a flag may be signaled for each entry of a palette predictor to indicate whether that entry is present in the current palette. As noted above, if a palette entry is not predicted, the palette entry may be explicitly signaled. In other examples, these two methods could be combined. For example, first the `pred_palette_flag` is signaled. If the flag is 0, a per-entry prediction flag may be signaled. In addition, the number of new entries and their explicit values may be signaled.

[0069] When determining second palettes 192 relative to first palettes 184 (e.g., `pred_palette_flag` is equal to one), video encoder 20 and/or video decoder 30 may locate one or more blocks from which the predictive palettes, in this example first palettes 184, are determined. The predictive palettes may be associated with one or more neighboring blocks of the block currently being coded (e.g., such as neighboring blocks (spatially or based on scan order) or the most frequent samples of a causal neighbor), i.e., second block 188. The palettes of the one or more neighboring blocks may be associated with a predictive palette. In some examples, such as the example illustrated in FIG. 2, video encoder 20 and/or video decoder 30 may locate a left neighboring block, first block 180, when determining a predictive palette for second block 188. In other examples, video encoder 20 and/or video decoder 30 may locate one or more blocks in other positions relative to second block 188, such as an upper block, block 196. In another example, the palette for the last block in scan order that used the palette mode may be used as a predictive palette.

[0070] Video encoder 20 and/or video decoder 30 may determine a block for palette prediction based on a hierarchy. For example, video encoder 20 and/or video decoder 30 may initially identify the left neighboring block, first block 180, for palette prediction. If the left neighboring block is not available for prediction (e.g., the left neighboring block is coded with a mode other than a palette-based coding mode, such as an intra-prediction mode or intra-prediction mode, or is located at the left-most edge of a picture or slice) video encoder 20 and/or video decoder 30 may identify the upper neighboring block, block 196. Video encoder 20 and/or video decoder 30 may continue searching for an available block according to a predetermined order of locations until locating a block having a palette available for palette prediction. In some examples, video encoder 20 and/or video decoder 30 may determine a predictive palette based on multiple blocks and/or reconstructed samples of a neighboring block.

[0071] While the example of FIG. 2 illustrates first palettes 184 as predictive palettes from a single block, (i.e., first block 180), in other examples, video encoder 20 and/or video decoder 30 may locate palettes for prediction from a combination of neighboring blocks. For example, video encoder 20 and/or video decoder may apply one or more formulas, functions, rules or the like to generate a predictive palette based on palettes of one or a combination of a plurality of neighboring blocks (spatially or in scan order).

[0072] In still other examples, video encoder 20 and/or video decoder 30 may construct a candidate list including a number of potential candidates for palette prediction. In such examples, video encoder 20 may encode an index to the candidate list to indicate the candidate block in the list from which the current block used for palette prediction is selected (e.g., copies the palette). Video decoder 30 may construct the candidate list in the same manner, decode the index, and use the decoded index to select the palette of the corresponding block for use with the current block. In another example, the palette of the indicated candidate block in the list may be used as a predictive palette for per-entry prediction of a current palette for the current block.

[0073] In an example for purposes of illustration, video encoder 20 and video decoder 30 may construct a candidate list that includes one block that is positioned above the block currently being coded and one block that is positioned to the left of the block currently being coded. In this example, video encoder 20 may encode one or more syntax elements to indicate the candidate selection. For example, video encoder 20 may encode a flag having a value of zero to indicate that the palette for the current block is copied from the block positioned to the left of the current block. Video encoder 20 may

encode the flag having a value of one to indicate that the palette for the current block is copied from the block positioned above the current block. Video decoder 30 decodes the flag and selects the appropriate block for palette prediction. In another example, the flag may indicate whether the palette of the top or left neighboring block is used as a predictive palette. Then, for each entry in the predictive palette, it may be indicated whether that entry is used in the palette for the current block.

[0074] In still other examples, video encoder 20 and/or video decoder 30 determine the palette for the block currently being coded based on the frequency with which sample values included in one or more other palettes occur in one or more neighboring blocks. For example, video encoder 20 and/or video decoder 30 may track the colors associated with the most frequently used index values during coding of a predetermined number of blocks. Video encoder 20 and/or video decoder 30 may include the most frequently used colors in the palette for the block currently being coded.

[0075] As noted above, in some examples, video encoder 20 and/or video decoder may copy an entire palette from a neighboring block for coding a current block.

Additionally, or alternatively, video encoder 20 and/or video decoder 30 may perform entry-wise based palette prediction. For example, video encoder 20 may encode one or more syntax elements for each entry of a palette indicating whether the respective entries are predicted based on a predictive palette (e.g., a palette of another block). In this example, video encoder 20 may encode a flag having a value of one for a given entry when the entry is a predicted value from a predictive palette (e.g., a corresponding entry of a palette associated with a neighboring block). Video encoder 20 may encode a flag having a value of zero for a particular entry to indicate that the particular entry is not predicted from a palette of another block. In this example, video encoder 20 may also encode additional data indicating the value of the non-predicted palette entry.

[0076] This disclosure describes several alternative techniques for predicting a palette for a current block. In one example, a predictive palette that includes palette entries from one or more previously coded neighboring blocks includes a number of entries, N . In this case, video encoder 20 first transmits a binary vector, V , having the same size as the predictive palette, i.e., size N , to video decoder 30. Each entry in the binary vector indicates whether the corresponding entry in the predictive palette will be reused or copied to the palette for the current block. For example, $V(i) = 1$ means that the i -th entry in the predictive palette for the neighboring block will be reused or copied to the palette for the current block, which may have a different index in the current block.

[0077] In addition, video encoder 20 may transmit a number, M , that indicates how many new palette entries are included in the palette for the current block, and then transmits a pixel value for each of the new palette entries to video decoder 30. In this example, the final size of the palette for the current block may be derived as equal to $M + S$, where S is the number of entries in the predictive palette that may be reused or copied to the palette for the current block (i.e., $V(i) = 1$). To generate the palette for the current block, video decoder 30 may merge the transmitted new palette entries and the copied palette entries reused from the predictive palette. In some cases, the merge may be based on the pixel values, such that the entries in the palette for the current block may increase (or decrease) with the palette index. In other cases, the merge may be a concatenation of the two sets of entries, i.e., the new palette entries and the copied palette entries.

[0078] In another example, video encoder 20 first transmits an indication of a size of a palette, N , for a current block to video decoder 30. Video encoder 20 then transmits a vector, V , having the same size as the palette for the current block, i.e., size N , to video decoder 30. Each entry in the vector indicates whether the corresponding entry in the palette for the current block is explicitly transmitted by video encoder 20 or copied from a predictive palette. For example, $V(i) = 1$ means that video encoder 20 transmits the i -th entry in the palette to video decoder 30, and $V(i) = 0$ means that the i -th entry in the palette is copied from the predictive palette. For the entries that are copied from the predictive palette (i.e., $V(i) = 0$), video encoder 20 may use different methods to signal which entry in the predictive palette is used in the palette for the current block. In some cases, video encoder 20 may signal the palette index of the entry to be copied from the predictive palette to the palette for the current block. In other cases, video encoder 20 may signal an index offset, which is the difference between the index in the palette for the current block and the index in the predictive palette.

[0079] In the two above examples, the one or more previously coded neighboring blocks used to generate the predictive palette used for prediction of the palette for the current block may be a top-neighboring (i.e., upper) block or a left-neighboring block with respect to the current block. In some examples, a candidate list of neighboring blocks may be constructed, and video encoder 20 transmits an index to indicate which candidate neighboring blocks and associated palettes are used for palette prediction for the current block. For certain blocks, e.g., blocks that are positioned at a beginning of a

slice or at other slice boundaries or leftmost blocks in the slice or a picture of video data, palette prediction may be disabled.

[0080] In an additional example, video encoder 20 transmits an indication of a number of entries included in a palette for a current block to video decoder 30. Then, for each of the palette entries, video encoder 20 transmits a flag or other syntax element to indicate whether the palette entry is explicitly transmitted by video encoder 20 or whether it is derived from a previously reconstructed pixel. For example, a one-bit flag set equal to 1 may mean that video encoder 20 explicitly sends the palette entry, and the one-bit flag set equal to 0 may mean that the palette entry is derived from a previously reconstructed pixel. For each of the palette entries that are derived from a previously reconstructed pixel, video encoder 20 transmits another indication regarding a pixel location of the reconstructed pixel in the current block or a neighboring block that corresponds to the palette entry. In some cases, the reconstructed pixel location indication may be a displacement vector with respect to the top-left position of the current block. In other cases, the reconstructed pixel location indication may be an index into a list of reconstructed pixels that can be used for specifying the palette entry for the current block. For example, this list may include all the reference pixels that may be used for normal intra prediction in HEVC.

[0081] In the example of FIG. 2, second palettes 192 includes four entries 208-214 having entry index value 1, entry index value 2, entry index value 3, and entry index 4, respectively. Entries 208-214 relate the index values to pixel values including pixel value A, pixel value B, pixel value C, and pixel value D, respectively. According to one or more aspects of this disclosure, video encoder 20 and/or video decoder 30 may use any of the above-described techniques to locate first block 180 for purposes of palette prediction and copy entries 1-3 of first palettes 184 to entries 1-3 of second palettes 192 for coding second block 188. In this way, video encoder 20 and/or video decoder 30 may determine second palettes 192 based on first palettes 184. In addition, video encoder 20 and/or video decoder 30 may code data for entry 4 to be included with second palettes 192. Such information may include the number of palette entries not predicted from a predictive palette and the pixel values corresponding to those palette entries.

[0082] In some examples, according to aspects of this disclosure, one or more syntax elements may indicate whether palettes, such as second palettes 192, are predicted entirely from a predictive palette (shown in FIG. 2 as first palettes 184, but which may

be composed of entries from one or more blocks) or whether particular entries of second palettes 192 are predicted. For example, an initial syntax element may indicate whether all of the entries are predicted. If the initial syntax element indicates that not all of the entries are predicted (e.g., a flag having a value of 0), one or more additional syntax elements may indicate which entries of second palettes 192 are predicted from the predictive palette.

[0083] According to some aspects of this disclosure, certain information associated with palette prediction may be inferred from one or more characteristics of the data being coded. That is, rather than video encoder 20 encoding syntax elements (and video decoder 30 decoding such syntax elements), video encoder 20 and video decoder 30 may perform palette prediction based on one or more characteristics of the data being coded.

[0084] In an example, for purposes of illustration, the value of `pred_palette_flag`, described above, may be inferred from one or more of, as examples, the size of the block being coded, the frame type, the color space, the color component, the frame size, the frame rate, the layer id in scalable video coding, and/or the view id in multi-view coding. That is, with respect to the size of the block as an example, video encoder 20 and/or video decoder 30 may determine that the above-described `pred_palette_flag` is equal to one for any blocks that exceed or are less than a predetermined size. In this example, the `pred_palette_flag` does not need to be signaled in the encoded bitstream.

[0085] While described above with respect to the `pred_palette_flag`, video encoder 20 and/or video decoder 30 may also or alternatively infer other information associated with palette prediction, such as the candidate block from which the palette is used for prediction, or rules for constructing palette prediction candidates, based on one or more characteristics of the data being coded.

[0086] According to other aspects of this disclosure, video encoder 20 and/or video decoder 30 may construct a palette on-the-fly. For example, when initially coding second block 188, there are no entries in second palettes 192. As video encoder 20 and video decoder 30 code new values for pixels of second block 188, each new value is included in second palettes 192. That is, for example, video encoder 20 adds pixel values to second palettes 192 as the pixel values are generated and signaled for positions in block 188. As video encoder 20 encodes pixels relatively later in the block, video encoder 20 may encode pixels having the same values as those already included in the palette using index values rather than signaling the pixel values. Similarly, when video

decoder 30 receives a new pixel value (e.g., signaled by video encoder 20) for a position in second block 188, video decoder 30 includes the pixel value in second palettes 192.

When pixel positions decoded relatively later in second block 188 have pixel values that have been added to second palettes 192, video decoder 30 may receive information such as, e.g., index values, that identify the corresponding pixel values in second palettes 192 for reconstruction of the pixel values of second block 188.

[0087] In some examples, as described in greater detail below, video encoder 20 and/or video decoder 30 may maintain first palettes 184 and second palettes 192 at or below a maximum palette size. According to aspects of this disclosure, if a maximum palette size is reached, e.g., as second palettes 192 are constructed dynamically on-the-fly, then video encoder 20 and/or video decoder 30 perform the same process to remove an entry of second palettes 192. One example process for removing palette entries is a first-in-first-out (FIFO) technique in which video encoder 20 and video decoder 30 remove the oldest entry of a palette. In another example, video encoder 20 and video decoder 30 may remove the least frequently used palette entry from the palette. In still another example, video encoder 20 and video decoder 30 may weight both FIFO and frequency of use processes to determine which entry to remove. That is, removal of an entry may be based on how old the entry is and how frequently it is used.

[0088] According to some aspects, if an entry (pixel value) is removed from a palette and the pixel value occurs again at a later position in the block being coded, video encoder 20 may encode the pixel value instead of including an entry in the palette and encoding an index. Additionally, or alternatively, video encoder 20 may re-enter palette entries into the palette after having been removed, e.g., as video encoder 20 and video decoder 30 scan the positions in the block.

[0089] In some examples, the techniques for deriving a palette on-the-fly may be combined with one or more other techniques for determining a palette. In particular, as an example, video encoder 20 and video decoder 30 may initially code second palettes 192 (e.g., using palette prediction to predict second palettes 192 from first palettes 184) and may update second palettes 192 when coding pixels of second block 188. For example, upon transmitting the initial palette, video encoder 20 may add values to the initial palette or change values in the initial palette as pixel values of additional locations in the block are scanned. Likewise, upon receiving an initial palette, video decoder 30 may add (i.e., include) values to the initial palette or change values in the initial palette as pixel values of additional locations in the block are scanned.

[0090] Video encoder 20 may, in some examples, signal whether the current block uses transmission of an entire palette, or on-the-fly palette generation, or a combination of transmission of an initial palette with updating of the initial palette by on-the-fly derivation. In some examples, the initial palette may be a full palette at maximum palette size, in which case values in the initial palette may be changed. In other examples, the initial palette may be smaller than the maximum palette size, in which cases video encoder 20 and video decoder 30 may add values to and/or change values of the initial palette.

[0091] According to one or more aspects of this disclosure, the size of palettes, such as first palettes 184 and second palettes 192, e.g., in terms of the number of pixel values that are included in the palette may be fixed or may be signaled using one or more syntax elements in an encoded bitstream. For example, according to some aspects, video encoder 20 and video decoder 30 may use unary codes or truncated unary codes (e.g., codes that truncate at a maximum limit of the palette size) to code the palette size. According to other aspects, video encoder 20 and video decoder 30 may use Exponential-Golomb or Rice-Golomb codes to code the palette size.

[0092] According to still other aspects, video encoder 20 and video decoder 30 may code data indicating the size of the palette after each entry of the palette. With respect to second palettes 192 as an example, video encoder 20 may encode a stop flag after each of entries 208-214. In this example, a stop flag equal to one may specify that the entry currently being coded is the final entry of second palettes 192, while a stop flag equal to zero may indicate that there are additional entries in second palettes 192. Accordingly, video encoder 20 may encode stop flags having a value of zero after each of entries 208-212 and a stop flag having a value of one after entry 214. In some instances, the stop flag may not be included in the bitstream upon the constructed palette reaching a maximum palette size limit. While the examples above disclose techniques for explicitly signaling the size of palettes, in other examples, the size of palettes may also be conditionally transmitted or inferred based on so-called side information (e.g., characteristic information such as the size of the block being coded, the frame type, the color space, the color component, the frame size, the frame rate, the layer id in scalable video coding or the view id in multi-view coding, as noted above).

[0093] The techniques of this disclosure include coding data losslessly, or, alternatively, with some losses (lossy coding). For example, with respect to lossy coding, video encoder 20 may code the pixels of a block without exactly matching the pixel values of

palettes exactly to the actual pixel values in the block. When the techniques of this disclosure are applied to lossy coding, some restrictions may be applied to the palette. For example, video encoder 20 and video decoder 30 may quantize palettes, such as first palettes 184 and second palettes 192. That is, video encoder 20 and video decoder 30 may merge or combine (i.e., quantize) entries of a palette when the pixel values of the entries are within a predetermined range of each other. In other words, if there is already a palette value that is within an error margin of a new palette value, the new palette value is not added to the palette. In another example, a plurality of different pixel values in a block may be mapped to a single palette entry, or, equivalently, to a single palette pixel value.

[0094] Video decoder 30 may decode pixel values in the same manner, regardless of whether a particular palette is lossless or lossy. As one example, video decoder 30 may use an index value transmitted by video encoder 20 for a given pixel position in a coded block to select an entry in the palette for the pixel position, without regard to whether the palette is lossless or lossy. In this example, the pixel value of the palette entry is used as the pixel value in the coded block, whether it matches the original pixel value exactly or not.

[0095] In an example of lossy coding, for purposes of illustration, video encoder 20 may determine an error bound, referred to as a delta value. A candidate pixel value entry `Plt_cand` may correspond to a pixel value at a position in a block to be coded, such as CU or PU. During construction of the palette, video encoder 20 determines the absolute difference between the candidate pixel value entry `Plt_cand` and all of the existing pixel value entries in the palette. If all of the absolute differences between the candidate pixel value entry `Plt_cand` and the existing pixel value entries in the palette are larger than the delta value, video encoder 20 may add the pixel value candidate to the palette as an entry. If an absolute difference between the pixel value entry `Plt_cand` and at least one existing pixel value entry in the palette is equal to or smaller than the delta value, video encoder 20 may not add the candidate pixel value entry `Plt_cand` to the palette. Thus, when coding the pixel value entry `Plt_cand`, video encoder 20 may select the entry with the pixel value that is the closest to the pixel value entry `Plt_cand`, thereby introducing some loss into the system. When a palette consists of multiple components (e.g. three color components), the sum of absolute difference of individual component values may be used for comparison against the delta value. Alternatively, or

additionally, the absolute difference for each component value may be compared against a second delta value.

[0096] In some examples, the existing pixel value entries in the palette noted above may have been added using a similar delta comparison process. In other examples, the existing pixel values in the palette may have been added using other processes. For example, one or more initial pixel value entries may be added to a palette (without a delta comparison) to start the delta comparison process of constructing the palette. The process described above may be implemented by video encoder 20 and/or video decoder 30 to produce luma and/or chroma palettes.

[0097] The techniques described above with respect to palette construction may also be used by video encoder 20 and video decoder 30 during pixel coding. For example, when encoding of a pixel value, video encoder 20 may compare the value of the pixel with the pixel values of entries in the palette. If the absolute pixel value difference between the value of the pixel and one of the entries in the palette is equal to or smaller than a delta value, video encoder 20 may encode the pixel value as the entry of the palette. That is, in this example, video encoder 20 encodes the pixel value using one of the entries of the palette when the pixel value produces a sufficiently small (e.g., within a predetermined range) absolute difference versus the palette entry.

[0098] In some examples, video encoder 20 may select the palette entry that yields the smallest absolute pixel value difference (compared to the pixel value being coded) to encode the pixel value. As an example, video encoder 20 may encode an index to indicate a palette entry that will be used for the pixel value, e.g., the palette pixel value entry that will be used to reconstruct the coded pixel value at video decoder 30. If the absolute pixel value difference between the value of the pixel and all of the entries in the palette is greater than delta, the encoder may not use one of the palette entries to encode the pixel value, and instead may transmit the pixel value of the pixel (possibly after quantization) to video decoder 30 (and possibly add the pixel value as an entry to the palette).

[0099] In another example, video encoder 20 may select an entry of a palette for encoding a pixel value. Video encoder 20 may use the selected entry as a predictive pixel value. That is, video encoder 20 may determine a residual value representing a difference between the actual pixel value and the selected entry and encode the residue. Video encoder 20 may generate residual values for pixels in a block that are predicted by entries of a palette, and may generate a residue block including respective residual

pixel values for the block of pixels. Video encoder 20 may subsequently apply transformation and quantization (as noted above with respect to FIG. 2) to the residue block. In this manner, video encoder 20 may generate quantized residual transform coefficients. In another example, the residue may be coded losslessly (without transform and quantization) or without transform.

[0100] Video decoder 30 may inverse transform and inverse quantize the transform coefficients to reproduce the residual block. Video decoder 30 may then reconstruct a pixel value using the predictive palette entry value and the residual value for the pixel value. For example, video decoder 30 may combine the residual value with the palette entry value to reconstruct the coded pixel value.

[0101] In some examples, the delta value may be different for different block sizes, picture sizes, color spaces or different color components. The delta value may be predetermined or determined based on various coding conditions. For example, video encoder 20 may signal the delta value to video decoder 30 using high level syntax, such as syntax in PPS, SPS, VPS and/or slice header. In other examples, video encoder 20 and video decoder 30 may be preconfigured to use the same, fixed delta value. In still other examples, video encoder 20 and/or video decoder 30 may adaptively derive the delta value based on side information (e.g., such as block size, color space, color component, or the like, as noted above).

[0102] In some examples, a lossy coding palette mode may be included as an HEVC coding mode. For example, coding modes may include an intra-prediction mode, an inter-prediction mode, a lossless coding palette mode, and a lossy coding palette mode. In HEVC coding, as noted above with respect to FIGS. 2 and 3, a quantization parameter (QP) is used to control the allowed distortion. The value of delta for palette-based coding may be calculated or otherwise determined as a function of QP. For example, the above-described delta value may be $1 \ll (QP/6)$ or $1 \ll ((QP+d)/6)$ where d is a constant, and “ \ll ” represents the bitwise left-shift operator.

[0103] Generation of a palette using the lossy coding techniques described in this disclosure may be performed by video encoder 20, video decoder 30 or both. For example, video encoder 20 may generate entries in a palette for a block using the delta comparison techniques described above and signal information for construction of the palette for use by video decoder 30. That is, video encoder 20 may be configured to signal information indicating pixel values for entries in a palette for a block, and then encode pixel values using the pixel values associated with such palette entries. Video

decoder 30 may construct a palette using such information, and may then use the entries to decode pixel values of a coded block. In some examples, video encoder 20 may signal index values that identify palette entries for one or more pixel positions of the coded block, and video decoder 30 may use the index values to retrieve the pertinent pixel value entries from the palette.

[0104] In other examples, video decoder 30 may be configured to construct a palette by applying the delta comparison techniques described above. For example, video decoder 30 may receive pixel values for positions within a coded block and may determine whether absolute differences between the pixel values and the existing pixel value entries in the palette are larger than a delta value. If so, video decoder 30 may add the pixel values as entries in the palette, e.g., for later use in palette-based decoding of pixel values for other pixel positions of the block using corresponding index values signaled by video encoder 20. In this case, video encoder 20 and video decoder 30 apply the same or similar processes to generate the palette. If not, video decoder 30 may not add the pixel values to the palette.

[0105] In an example for purposes of illustration, video decoder 30 may receive index values or pixel values for various pixel positions in a block. If an index value is received for a pixel position, for example, video decoder 30 may use the index value to identify an entry in the palette, and use the pixel value of the palette entry for the pixel position. If a pixel value is received for the pixel position, video decoder 30 may use the received pixel value for the pixel position, and may also apply the delta comparison to determine whether the pixel value should be added to the palette and then later used for palette coding.

[0106] On the encoder side, if a pixel value for a position in a block produces an absolute difference between the pixel value and an existing pixel value entry in the palette that is less than or equal to the delta value, video encoder 20 may send an index value to identify the entry in the palette for use in reconstructing the pixel value for that position. If a pixel value for a position in a block produces absolute difference values between the pixel value and the existing pixel value entries in the palette that are all greater than the delta value, video encoder 20 may send the pixel value and may add the pixel value as a new entry in the palette. To construct the palette, video decoder 30 may use delta values signaled by the encoder, rely on a fixed or known delta value, or infer or derive a delta value, e.g., as described above.

[0107] As noted above, video encoder 20 and/or video decoder 30 may use coding modes including an intra-prediction mode, an inter-prediction mode, a lossless coding palette mode, and a lossy coding palette mode when coding video data. According to some aspects of this disclosure, video encoder 20 and video decoder 30 may code one or more syntax elements indicating whether palette-based coding is enabled. For example, at each block, video encoder 20 may encode a syntax element, such as a flag `PLT_Mode_flag`. The `PLT_Mode_flag` or other syntax element may indicate whether a palette-based coding mode is to be used for a given block (e.g., a CU or a PU). For example, this flag may be signaled in an encoded video bitstream at the block level (e.g., a CU level), and then received by video decoder 30 upon decoding the encoded video bitstream.

[0108] In this example, a value of this `PLT_Mode_flag` equal to 1 may specify that the current block is encoded using a palette-based coding mode. In this case, video decoder 30 may apply the palette-based coding mode to decode the block. In some examples, a syntax element may indicate one of a plurality of different palette-based coding modes for the block (e.g., lossy or lossless). A value of this `PLT_Mode_flag` equal to 0 may specify that the current block is encoded using a mode other than palette mode. For example, any of a variety of inter-predictive, intra-predictive, or other coding modes may be used. When a value of `PLT_Mode_flag` is 0, video encoder 20 may also encode additional data to indicate the specific mode used for encoding the respective block (e.g., an HEVC coding mode). The use of the `PLT_Mode_flag` is described for purposes of example. In other examples, however, other syntax elements such as multi-bit codes may be used to indicate whether the palette-based coding mode is to be used for a block (e.g., a CU or PU) or to indicate which of a plurality of modes are to be used for coding.

[0109] In some examples, the above-described flag or other syntax elements may be transmitted at a higher level than the block level. For example, video encoder 20 may signal such a flag at a slice level. In this case, a value equal to 1 indicates that all of the blocks in the slice are encoded using palette mode. In this example, no additional mode information, e.g., for palette mode or other modes, is signaled at the block level. In another example, video encoder 20 may signal such a flag in a PPS, SPS or VPS.

[0110] According to some aspects of this disclosure, video encoder 20 and/or video decoder 30 may code one or more syntax elements (e.g., such as the above-described flag) at one of the slice, PPS, SPS, or VPS levels specifying whether the palette mode is

enabled or disabled for the particular slice, picture, sequence or the like, while the `PLT_Mode_flag` indicates whether the palette-based coding mode is used for each block. In this case, if a flag or other syntax element sent at the slice, PPS, SPS or VPS level indicates that palette coding mode is disabled, in some examples, there may be no need to signal the `PLT_Mode_flag` for each block. Alternatively, if a flag or other syntax element sent at the slice, PPS, SPS or VPS level indicates that palette coding mode is enabled, the `PLT_Mode_flag` may be further signaled to indicate whether the palette-based coding mode is to be used for each block.

[0111] In some examples, the above-described syntax elements may be conditionally signaled in the bitstream. For example, video encoder 20 and video decoder 30 may only encode or decode, respectively, the syntax elements based on the size of the block, the frame type, the color space, the color component, the frame size, the frame rate, the layer id in scalable video coding, and/or the view id in multi-view coding.

[0112] While the examples described above relate to explicit signaling, e.g., with one or more syntax elements in a bitstream, in other examples, video encoder 20 and/or video decoder 30 may implicitly determine whether a palette coding mode is active and/or used for coding a particular block. Video encoder 20 and video decoder 30 may determine whether palette-based coding is used for a block based on, for example, the size of the block, the frame type, the color space, the color component, the frame size, the frame rate, the layer id in scalable video coding and/or the view id in multi-view coding.

[0113] FIG. 3 is a conceptual diagram illustrating examples of determining indices to a palette for a video block, consistent with techniques of this disclosure. For example, FIG. 3 includes a map 240 of index values (values 1, 2, and 3) that relate respective positions of pixels associated with the index values to an entry of palettes 244. Palettes 244 may be determined in a similar manner as first palettes 184 and second palettes 192 described above with respect to FIG. 2.

[0114] Again, the techniques of FIG. 3 are described in the context of video encoder 20 (FIG. 1 and FIG. 4) and video decoder 30 (FIG. 1 and FIG. 5) and with respect to the HEVC video coding standard for purposes of explanation. However, it should be understood that the techniques of this disclosure are not limited in this way, and may be applied by other video coding processors and/or devices in other video coding processes and/or standards.

[0115] While map 240 is illustrated in the example of FIG. 3 as including an index value for each pixel position, it should be understood that in other examples, not all pixel positions may be associated with an index value that indicates an entry of palettes 244 that specify the pixel value of the block. That is, as noted above, in some examples, video encoder 20 may encode (and video decoder 30 may obtain, from an encoded bitstream) an indication of an actual pixel value (or its quantized version) for a position in map 240 if the pixel value is not included in palettes 244.

[0116] In some examples, video encoder 20 and video decoder 30 may be configured to code an additional map indicating which pixel positions are associated with index values. For example, assume that the (i, j) entry in the map corresponds to the (i, j) position of a block. Video encoder 20 may encode one or more syntax elements for each entry of the map (i.e., each pixel position) indicating whether the entry has an associated index value. For example, video encoder 20 may encode a flag having a value of one to indicate that the pixel value at the (i, j) location in the block is one of the values in palettes 244. Video encoder 20 may, in such an example, also encode a palette index (shown in the example of FIG. 3 as values 1-3) to indicate that pixel value in the palette and to allow video decoder to reconstruct the pixel value. In instances in which palettes 244 include a single entry and associated pixel value, video encoder 20 may skip the signaling of the index value. Video encoder 20 may encode the flag to have a value of zero to indicate that the pixel value at the (i, j) location in the block is not one of the values in palettes 244. In this example, video encoder 20 may also encode an indication of the pixel value for use by video decoder 30 in reconstructing the pixel value. In some instances, the pixel value may be coded in a lossy manner.

[0117] The value of a pixel in one position of a block may provide an indication of values of one or more other pixels in other positions of the block. For example, there may be a relatively high probability that neighboring pixel positions of a block will have the same pixel value or may be mapped to the same index value (in the case of lossy coding, in which more than one pixel value may be mapped to a single index value).

[0118] Accordingly, according to aspects of this disclosure, video encoder 20 may encode one or more syntax elements indicating a number of consecutive pixels or index values in a given scan order that have the same pixel value or index value. As noted above, the string of like-valued pixel or index values may be referred to herein as a run. In an example for purposes of illustration, if two consecutive pixels or indices in a given scan order have different values, the run is equal to zero. If two consecutive pixels or

indices in a given scan order have the same value but the third pixel or index in the scan order has a different value, the run is equal to one. For three consecutive indices or pixels with the same value, the run is two, and so forth. Video decoder 30 may obtain the syntax elements indicating a run from an encoded bitstream and use the data to determine the number of consecutive locations that have the same pixel or index value.

[0119] In some examples, all pixel locations in the current block having pixel values that are in the palette for the current block are encoded with a palette index followed by a “run” of the pixel value at consecutive pixel locations. In the case where there is only one entry in the palette, the transmission of the palette index or the “run” may be skipped for the current block. In the case where the pixel value at one of the pixel locations in the current block does not have an exact match to a pixel value in the palette, video encoder 20 may select one of the palette entries having the closest pixel value and may calculate a prediction error or residual value between the original pixel value and the prediction pixel value included in the palette. Video encoder 20 encodes and transmits the residual value for the pixel location to the video decoder. Video decoder 30 may then derive a pixel value at the pixel location based on the corresponding received palette index, and the derived pixel value and the residual value are then used to predict the original pixel value at the pixel location in the current block. In one example, the residual value is encoded using an HEVC method specified by HEVC draft 10, such as applying a RQT to transform the residual value, quantize the transform coefficients, and entropy encode the quantized transform coefficients. In some cases, the above example may be referred to as lossy coding.

[0120] In an example for purposes of illustration, consider line 248 of map 240. Assuming a horizontal, left to right scan direction, line 248 includes five index values of “2” and three index values of “3.” According to aspects of this disclosure, video encoder 20 may encode an index value of 2 for the first position of line 248 in the scan direction. In addition, video encoder 20 may encode one or more syntax elements indicating the run of consecutive values in the scan direction that have the same index value as the signaled index value. In the example of line 248, video encoder 20 may signal a run of 4, thereby indicating that the index values of the following four positions in the scan direction share the same index value as the signaled index value. Video encoder 20 may perform the same process for the next different index value in line 248. That is, video encoder 20 may encode an index value of 3 and one or more syntax elements indicating a run of two. Video decoder 30 may obtain the syntax elements

indicating the index value and the number of consecutive indices in the scan direction having the same index value (the run).

[0121] As noted above, the indices of a map are scanned in a particular order.

According to aspects of this disclosure, the scan direction may be vertical, horizontal, or at a diagonal (e.g., 45 degrees or 135 degrees diagonally in block). In some examples, video encoder 20 may encode one or more syntax elements for each block indicating a scan direction for scanning the indices of the block. Additionally or alternatively, the scan direction may be signaled or inferred based on so-called side information such as, for example, block size, color space, and/or color component. Video encoder 20 may specify scans for each color component of a block. Alternatively, a specified scan may apply to all color components of a block.

[0122] For example, with respect to a column based scan, consider column 252 of map 240. Assuming a vertical, top to bottom scan direction, column 252 includes one index value of “1,” five index values of “2” and two index values of “3.” According to aspects of this disclosure, video encoder 20 may encode an index value of 1 for the first position of column 252 in the scan direction (at the relative top of column 252). In addition, video encoder 20 may signal a run of zero, thereby indicating that the index value of the following position in the scan direction is different. Video encoder 20 may then encode an index value of 2 for the next position in the scan direction and one or more syntax elements indicating a run of four, i.e., that the index values of the following four positions in the scan direction share the same index value as the signaled index value. Video encoder 20 may then encode an index value of 3 for the next different index value in the scan direction and one or more syntax elements indicating a run of one. Video decoder 30 may obtain the syntax elements indicating the index value and the number of consecutive indices in the scan direction having the same index value (the run).

[0123] According to aspects of this disclosure, video encoder 20 and video decoder 30 may additionally or alternatively perform line copying for one or more entries of map 240. The line copying may depend, in some examples, on the scan direction. For example, video encoder 20 may indicate that a pixel or index value for a particular entry in a map is equal to a pixel or index value in a line above the particular entry (for a horizontal scan) or the column to the left of the particular entry (for a vertical scan). Video encoder 20 may also indicate, as a run, the number of pixel or index values in the scan order that are equal to the entry in the line above or the column to the left of the

particular entry. In this example, video encoder 20 and or video decoder 30 may copy pixel or index values from the specified neighboring line and from the specified number of entries for the line of the map currently being coded.

[0124] In an example for purposes of illustration, consider columns 256 and 260 of map 240. Assuming a vertical, top to bottom scan direction, column 256 includes three index values of “1,” three index values of “2,” and two index values of “3.” Column 260 includes the same index values having the same order in the scan direction. According to aspects of this disclosure, video encoder 20 may encode one or more syntax elements for column 260 indicating that the entire column 260 is copied from column 256. The one or more syntax elements may be associated with a first entry of column 260 at the relative top of map 240. Video decoder 30 may obtain the syntax elements indicating the line copying and copy the index values of column 256 for column 260 when decoding column 260.

[0125] According to aspects of this disclosure, the techniques for coding so-called runs of entries may be used in conjunction with the techniques for line copying described above. For example, video encoder 20 may encode one or more syntax elements (e.g., a flag) indicating whether the value of an entry in a map is obtained from a palette or the value of an entry in the map is obtained from a previously coded line in map 240. Video encoder 20 may also encode one or more syntax elements indicating an index value of a palette or the location of the entry in the line (the row or column). Video encoder 20 may also encode one or more syntax elements indicating a number of consecutive entries that share the same value. Video decoder 30 may obtain such information from an encoded bitstream and may use the information to reconstruct the map and pixel values for a block.

[0126] In an example for purposes of illustration, consider rows 264 and 268 of map 240. Assuming a horizontal, left to right scan direction, row 264 includes five index values of “1” and three index values of “3.” Row 268 includes three index values of “1,” two index values of “2,” and three index values of “3.” In this example, video encoder 20 may identify particular entries of row 264 followed by a run when encoding data for row 268. For example, video encoder 20 may encode one or more syntax elements indicating that the first position of row 268 (the left most position of row 268) is the same as the first position of row 264. Video encoder 20 may also encode one or more syntax elements indicating that the next run of two consecutive entries in the scan direction in row 268 is the same as the first position of row 264.

[0127] In some examples, video encoder 20 may also determine whether to code the current pixel or index value relative to a position in another row (or column) or to code the current pixel or index value using a run syntax element. For example, after encoding the one or more syntax elements indicating the first position of row 264 and the run of two entries (noted above), video encoder 20 may encode, for the fourth and fifth positions in line 268 (from left to right), one or more syntax elements indicating a value of 2 for the fourth position and one or more syntax elements indicating a run of 1. Hence, video encoder 20 encodes these two positions without reference to another line (or column). Video encoder 20 may then code the first position having an index value of 3 in row 268 relative to upper row 264 (e.g., indicating a copy from upper row 264 and the run of consecutive positions in the scan order having the same index value). Hence, according to aspects of this disclosure, video encoder 20 may select between coding pixel or index values of a line (or column) relative to other values of the line (or column), e.g., using a run, coding pixel or index values of a line (or column) relative to values of another line (or column), or a combination thereof. In some examples, video encoder 20 may perform a rate/distortion optimization to make the selection.

[0128] Video decoder 30 may receive the syntax elements described above and may reconstruct row 268. For example, video decoder 30 may obtain data indicating a particular location in a neighboring row from which to copy the associated index value for the position of map 240 currently being coded. Video decoder 30 may also obtain data indicating the number of consecutive positions in the scan order having the same index value.

[0129] In some instances, the line from which entries are copied may be directly adjacent to the entry of the line currently being coded (as illustrated in the examples of FIG. 3). However, in other examples, a number of lines may be buffered by video encoder 20 and/or video decoder 30, such that any of the number of lines of the map may be used as predictive entries for a line of the map currently being coded. Hence, in some examples, the pixel value for an entry may be signaled to be equal to a pixel value of an entry in a row immediately above (or column to the left of) or two or more rows above (or column to the left of) the current row.

[0130] In an example for purposes of illustration, video encoder 20 and/or video decoder 30 may be configured to store the previous n rows of entries prior to coding a current row of entries. In this example, video encoder 20 may indicate the predictive row (the row from which entries are copied) in a bitstream with a truncated unary code

or other codes. In another example, video encoder 20 may encode (and video decoder 30 may decode) a displacement value between the current line and the predictive line of map 240 used as a reference for coding the current line. That is, video encoder 20 may encode an indication of a particular line (or column) from which an index value is copied. In some examples, the displacement value may be a displacement vector. That is, let $c[0]$, $c[1]$, ..., denote the indices of the current line of map 240 and let $u[0]$, $u[1]$, $u[2]$, ..., denote the indices of a predictive line of map 240, such as an upper neighboring line. In this example, given a displacement vector is d , the index value for $c[i]$ may be predicted from $u[i+d]$, or $u[i-d]$ to avoid d taking negative values. The value of d may be coded using unary, truncated unary, exponential Golomb or Golomb-Rice codes.

[0131] As another example, video encoder 20 may signal an instruction, such as “copy from up line left half” or “copy from up line right half,” indicating the neighboring line and the number or portion of entries of the neighboring line to copy to the line of the map currently being coded. As an additional example, the map of index values may be re-ordered before coding. For example, the map of index values may be rotated by 90, 180 or 270 degrees, or flipped upside down or left-side right to improve coding efficiency.

[0132] In other examples, video encoder 20 may not transmit runs of like-valued index values of map 240 to video decoder 30. In this case, video encoder 20 and/or video decoder 30 may implicitly derive the values of the runs. In one example, the value of a run may be a constant value, e.g., 4, 8, 16, or the like. In another example, the value of a run may be dependent on side information for the current block of video data being coded such as, for example, the block size, the quantization parameter (QP), the frame type, the color component, the color format (e.g., 4:4:4, 4:2:2, or 4:2:0), the color space (e.g., YUV or RGB), the scan direction and/or other types of characteristic information for the current block. In the case where the value of a run depends on the block size, the run may be equal to the width of the current block, the height of the current block, the half-width (or half-height) of the current block, a fraction of the width and/or the height of the current block, or a multiple of the width and/or the height of the current block. In another example, video encoder 20 may signal the value of a run to video decoder 30 using high level syntax, such as syntax in a VPS, SPS, PPS, and/or a slice header.

[0133] Additionally, or alternatively, video encoder 20 may not even need to transmit map 240 to video decoder 30. Instead, video encoder 20 and/or video decoder 30 may

implicitly derive a start position or location of each run of index values included in map 240. In one example, the video coding standard applied by video encoder 20 and/or video decoder 30 may determine that a run can only start at certain locations. For example, the run may only start at the beginning of each row, or the beginning of every N rows of a current block being coded. The start location may be different for different scan directions. For example, if the vertical scan is used, the run may only start at the beginning of a column or the beginning of every N columns of the current block.

[0134] In another example, the start location may be derived depending on side information for the current block such as, for example, the block size, the QP, the frame type, the color component, the color format (e.g., 4:4:4, 4:2:2, or 4:2:0), the color space (e.g., YUV or RGB), the scan direction and/or other types of characteristic information for the current block. In the case where the start location of a run depends on the block size, the start location may be the mid-point of each row and/or each column, or a fraction (e.g., $1/n$, $2/n$, ... $(n-1)/n$) of each row and/or column. In another example, video encoder 20 may signal the start position to video decoder 30 using high level syntax, such as syntax in a PPS, a SPS, a VPS and/or a slice header.

[0135] In some examples, the implicit start position derivation and the implicit run derivation, each described above, may be combined. For example, video encoder 20 and/or video decoder 30 may determine that a run of like-valued index values of the map is equal to a distance between two neighboring start positions. In the case where the start position is the beginning (i.e., the first position) of every row of a current block, then video encoder 20 and/or video decoder 30 may determine that the length of the run is equal to the length of an entire row of the current block.

[0136] Various proposals for the HEVC SCC extension include a palette mode, which uses a palette predictor to derive the palette for a current block. Video encoder 20 and video decoder 30 may update the palette predictor prior to determining the palette table for a block currently being coded. For a first block where a previously used palette predictor is not available, video encoder 20 and video decoder 30 may initialize a palette predictor either at a PPS level or at a SPS level with the use of two flags:

`pps_palette_predictor_initializer_present_flag` and

`sps_palette_predictor_initializer_present_flag` to enable this initialization from the

palette predictor entries signaled in the PPS and SPS, respectively. The size of the PPS and SPS palette predictor is defined by the syntax elements

`pps_num_palette_predictor_initializer_minus1` and

sps_num_palette_predictor_initializer_minus1, respectively. Under some proposed signaling mechanisms, the minimum size of the palette predictor signaled in the PPS or SPS is 1.

[0137] One example order in which the palette predictor may be initialized is described in WD4 section 9.3.2.3.

If pps_palette_predictor_initializer_present_flag is equal to 1, palette predictor is initialized based on the PPS predictor,
otherwise if (pps_palette_predictor_initializer_present_flag is equal to 0),
if sps_palette_predictor_initializer_present_flag is equal to 1, palette predictor is initialized from the SPS predictor,
otherwise (when both flags are equal to 0), palette predictor size is set to 0.

[0138] As can be seen from the initialization rule described above, if sps_palette_predictor_initializer_present_flag is enabled, there is no way to disable the palette predictor, which is to set the palette predictor size to 0, on a per picture basis. The ability to set the palette predictor size to 0 may be desirable if the palette predictor signaled in the SPS is not relevant to predict the palette entries in a certain picture and sending them in the PPS may consume a lot of rate. Another potential problem is that when both the PPS and SPS palette predictors are present, the current signaling might be redundant because the PPS and SPS palette predictors may have many common palette entries.

[0139] This disclosure introduces various techniques related to palette predictor initialization control. To potentially address the first problem introduced above, this disclosure describes techniques for adding a new syntax element palette_predictor_initialization_flag which can be used to control palette predictor initialization. When the value of this flag is equal to 1, palette predictor initialization is used. Otherwise, a palette predictor size is set equal to 0. This flag can be signalled in any parameter set, such as VPS, SPS, PPS, slice header or its extensions or elsewhere. In one example, the flag may be signalled in the PPS. A potential advantage of using such a flag is that the flag provides the ability to disable palette predictor (set palette predictor size to 0) when SPS palette predictor initializers are signalled.

[0140] As there can be two palette initializers (e.g., one in the PPS and another one in an SPS), and when palette_predictor_initialization_flag is signalled only in one parameter set, the bitstream conformance rule can be introduced to avoid parameter set

parsing dependency and for the case when the `palette_predictor_initialization_flag` is set equal to 1, and no palette predictors are signalled in any parameter set. The bitstream conformance restriction may be expressed as follows: Conformant bitstream shall contain at least one of the `pps_palette_predictor_initializer_present_flag` and `sps_palette_predictor_initializer_present_flag` equal to 1 when `palette_predictor_initialization_flag` is equal to 1. This rule is not meant to be restricted to the existing and suggested syntax elements, as it may be also applicable to solve similar problems when the current syntax elements can be modified or replaced by some other syntax elements, but playing similar roles.

[0141] In an alternative solution, video encoder 20 may be allowed to signal the palette predictor initializer equal to 0 in a PPS. In this case, the suffix “minus1” may be removed from the syntax element `pps_num_palette_predictor_initializer_minus1`. In addition, the semantics may be changed to reflect that the palette predictor initializer size can be zero. Thus, the existing initialization rule can be changed to allow having a zero palette predictor in the PPS. In such a case, setting the palette predictor size in a PPS equal to 0 can be used to switch off the palette predictor initialization even when SPS palette predictor initializer is signalled, for example, by signalling `pps_palette_predictor_initializer_present_flag` equal to 1 with `pps_num_palette_predictor_initializer` equal to 0. That is, if video decoder 30 receives, in a PPS, a value for `pps_num_palette_predictor_initializer` that indicates a number of PPS-level palette predictor initializers included in the PPS is equal to zero, then video decoder 30 may be configured to determine that palette predictor initialization is disabled for video blocks associated with the PPS, even though the SPS with which the block is associated has palette predictor initialization enabled.

[0142] Additionally, a newly introduced syntax element, `pps_num_palette_predictor_initializer`, value can be used as a new gating flag for the syntax elements specific to palette predictor initializer entries, such as `monochrome_palette_flag`, `luma_bit_depth_entry_minus8`, and `chroma_bit_depth_entry_minus8` since those syntax elements are not needed when the initializer size is equal to 0. Thus, based on the syntax element `pps_num_palette_predictor_initializer` indicating the number of PPS-level palette predictor initializers included in the PPS is equal to zero, video decoder 30 may determine the syntax elements `monochrome_palette_flag`,

luma_bit_depth_entry_minus8, and chroma_bit_depth_entry_minus8 are not present in the PPS.

[0143] For example, PPS signalling can be modified as follows where the changes are shown below with underline and strikethrough:

pps_palette_predictor_initializer_present_flag	u(1)
if(pps_palette_predictor_initializer_present_flag) {	
<u>pps_num_palette_predictor_initializer</u>	ue(v)
if(pps_num_palette_predictor_initializer > 0) {	
monochrome_palette_flag	u(1)
luma_bit_depth_entry_minus8	ue(v)
if(!monochrome_palette_flag)	
chroma_bit_depth_entry_minus8	ue(v)
<u>pps_num_palette_predictor_initializer_minus1</u>	ue(v)
numComps = monochrome_palette_flag ? 1 : 3	
for(comp = 0; comp < numComps; comp++)	
for(i = 0; i < pps_num_palette_predictor_initializer_minus1 ; i++)	
pps_palette_predictor_initializers[comp][i]	u(v)
}	
}	

[0144] The syntax element pps_num_palette_predictor_initializer specifies the number of entries in the picture palette predictor initializer. If video decoder 30 determines that a value for pps_num_palette_predictor_initializer indicates that the number of PPS-level palette predictor initializers included in the PPS is equal to zero, then video decoder 30 may determine that palette predictor initialization is disabled for the PPS.

[0145] It may be a requirement of bitstream conformance that the value of pps_num_palette_predictor_initializer is in the range of 0 to PaletteMaxPredictorSize, inclusive.

[0146] The syntax element pps_palette_predictor_initializers[comp][i] specifies the value of the comp-th component of the i-th palette entry that is used to initialize the array PredictorPaletteEntries. For values of i in the range of 0 to pps_num_palette_predictor_initializer - 1, inclusive, the value of the pps_palette_predictor_initializers[0][i] shall be in the range of 0 to (1 << BitDepthEntryY) - 1, inclusive, and the values of pps_palette_predictor_initializers[1][i] and pps_palette_predictor_initializers[2][i] shall be in the range of 0 to (1 << BitDepthEntryC) - 1, inclusive.

[0147] The initialization rule can be changed accordingly:

Outputs of this process are the initialized palette predictor variables PredictorPaletteSize and PredictorPaletteEntries.

The variable numComps is derived as follows:

$$\text{numComps} = (\text{ChromaArrayType} == 0) ? 1 : 3$$

- If pps_palette_predictor_initializer_present_flag is equal to 1, the following applies:
 - PredictorPaletteSize is set equal to pps_num_palette_predictor_initializer
 - The array PredictorPaletteEntries is derived as follows:


```
for( comp = 0; comp < numComps; comp++ )
  for( i = 0; i < PredictorPaletteSize; i++ )
    PredictorPaletteEntries[ comp ][ i ] =
      pps_palette_predictor_initializers[ comp ][ i ]
```
- Otherwise (pps_palette_predictor_initializer_present_flag is equal to 0), if sps_palette_predictor_initializer_present_flag is equal to 1, the following applies:
 - PredictorPaletteSize is set equal to sps_num_palette_predictor_initializer_minus1 plus 1.
 - The array PredictorPaletteEntries is derived as follows:


```
for( comp = 0; comp < numComps; comp++ )
  for( i = 0; i < PredictorPaletteSize; i++ )
    PredictorPaletteEntries[ comp ][ i ] =
      sps_palette_predictor_initializers][ comp ][ i ]
```
- Otherwise (pps_palette_predictor_initializer_present_flag is equal to 0 and sps_palette_predictor_initializer_present_flag is equal to 0), PredictorPaletteSize is set equal to 0 and PredictorPaletteSize is set equal to 0.

[0148] This disclosure also describes techniques related to palette predictor initialization prediction. To potentially solve the second problem introduced above with possible redundant entries, this disclosure describes techniques for using both SPS and PPS palette predictor initializers for the palette predictor initialization instead of the existing techniques, where only one predictor can be used at a time. Thus, when initializing a palette predictor, video decoder 30 may include in the palette predictor both SPS-level palette predictor initializers and PPS-level palette predictor initializers, instead of only one of SPS-level palette predictor initializers or PPS-level palette predictor initializers

[0149] In one example, a palette predictor can be initialized by including the SPS palette predictor entries followed by the PPS palette predictor entries, or vice versa. In a

more generic approach, instead of using the entire SPS palette predictor, the reused flags can be signalled per entry to indicate which entries will be used for palette predictor initialization.

[0150] When both SPS and PPS initializers are used, an additional constraint can be imposed on the total signalled size of the combined initializer, such as the size of SPS and PPS initializers together is less than or equal to the maximum palette predictor size, since if the size of the combined initializer is greater than maximum palette predictor size, some entries have to be discarded and bits associated with this extra signalling can be wasted. In another simpler alternative, PPS and SPS initializer size can be capped to half the maximum palette predictor size.

[0151] In one example implementation, when the syntax element `pps_palette_predictor_initializer_present_flag` is equal to 1, for each entry in the SPS palette predictor, one flag is sent to signal whether that entry is used to initialize the palette predictor. The signalling method may be the same as used in Section 7.3.8.13 of WD4 to signal `PalettePredictorEntryReuseFlags`. After that, the number of additional palette predictor initializer entries is signalled. This may take the value 0. This is followed by the actual PPS palette predictor entries.

[0152] To avoid dependency between SPS and PPS, when the syntax element `pps_palette_predictor_initializer_present_flag` is equal to 1, a `sps_palette_initializer_reuse` flag may be sent. When the value of this flag is equal to 1, the number of SPS palette predictor entries is signalled. This is followed by signalling of a reuse flag for each SPS palette predictor entry, the number of additional palette predictor entries and finally, the actual PPS palette predictor entries. The number of SPS palette predictor entries signalled may be less than or equal to $\text{sps_num_palette_predictor_initializer_minus1} + 1$. For example, if `sps_num_palette_predictor_initializer_minus1` is equal to 10 and the number of SPS palette predictor entries signalled in the PPS is 6, then reuse flags are sent only for the first 6 SPS palette predictor entries and the last 4 SPS palette predictor entries may not be used in the palette predictor initialization for a picture pointing to the current PPS.

[0153] Following the last example, the palette predictor initialization reuse flags may be signalled in the slice header to indicate which entries are reused from the default palette predictor initializer which can be composed of at least one palette predictor initializer from the PPS or SPS, or both.

[0154] For example, a SPS palette predictor initializer may contain commonly used palette entries, and a PPS palette predictor initializer may have only the entries, which are specific to the picture and are different from the SPS entries. Initializing palette predictor from both PPS and SPS initializers may avoid additional signalling of the initializer entries in the PPS which may be the same across all signalled PPSs, contrary to the WD4 design, where such common entries are repeatedly signalled with each PPS palette predictor initializer.

[0155] In a similar manner, a palette predictor initializer can be combined from earlier sent PPSs and/or SPSs. For instance, instead of an SPS as described above, other palette predictor entries from an earlier sent PPS can be added to the palette predictor initializer. In one example, the palette predictor initializer may inherit one or more entries from previously-signaled PPS initializers, instead of or in addition to inheriting one or more entries from an SPS initializer. Although described above with respect to previously-signaled PPS initializers as an example, it will be appreciated that the techniques of this disclosure may also enable the inheritance of one or more entries from one or more previously-signaled SPS initializers. For instance, entries may be inherited from previously-signaled SPS initializers, from previously-signaled PPS initializers, or from a combination of both previously-signaled SPS initializers and previously-signaled PPS initializers.

[0156] In another example, a palette predictor initializer can just be updated. For example, it may initially be composed from first signaled SPS, PPS, or both initializers, and this palette predictor may be updated by inserting the entries of the newly signaled PPS or SPS into the palette predictor initializer. The entries can be inserted into the beginning of the palette predictor initializer, at the end, or basically at any location. The entries, which exceed the predictor size while inserting, can be ignored. For instance, in cases where such updates cause the initializer to exceed the predictor size limit, the updated entries that are processed after meeting the size limit can be disregarded, discarded, or excluded.

[0157] In one particular example of the described approach, a palette predictor initializer may contain newly signaled entries from SPS, or PPS, or both sets, plus the old entries of the previous palette predictor initializer following the new entries.

[0158] This disclosure also describes techniques related to scaling list update. Similar idea of combined palette predictor initializer can be applied to the scaling list or quantization matrix for transform. There can be SPS and PPS scaling lists signalled and

instead of using only either PPS or SPS list, the scaling list used for a particular picture can be composed from the SPS scaling list updated or appended by the lists signalled in the PPS. For example, consider that scaling lists for 4x4, 8x8, 16x16, and 32x32 transform sizes are signalled in the SPS. The PPS may signal only certain scaling lists, for example, the ones associated with 8x8 and 16x16 transform sizes. Then, the final scaling list used for a picture can be as follows: 4x4 from SPS, 8x8 from the PPS, 16x16 from the PPS, and 32x32 from the SPS.

[0159] In another example, the SPS scaling lists may not contain the whole set, for example, when transform depth or sizes are defined in the PPS. In this case, the SPS may have only the basic or common scaling lists, not necessarily overlapped with the PPS, and the specific ones for a picture may be signalled in the PPS and later combined set of the SPS and PPS scaling lists are used for a picture.

[0160] FIG. 4 is a block diagram illustrating an example video encoder 20 that may implement the techniques described in this disclosure. Video encoder 20 may perform intra- and inter-coding of video blocks within video slices. Intra-coding relies on spatial prediction to reduce or remove spatial redundancy in video within a given video frame or picture. Inter-coding relies on temporal prediction to reduce or remove temporal redundancy in video within adjacent frames or pictures of a video sequence. Intra-mode (I mode) may refer to any of several spatial based compression modes. Inter-modes, such as uni-directional prediction (P mode) or bi-prediction (B mode), may refer to any of several temporal-based compression modes.

[0161] In the example of FIG. 4, video encoder 20 includes a video data memory 33, partitioning unit 35, prediction processing unit 41, summer 50, transform processing unit 52, quantization unit 54, entropy encoding unit 56. Prediction processing unit 41 includes palette-based encoding unit 42, inter-prediction processing unit 44, and intra prediction unit 46. For video block reconstruction, video encoder 20 also includes inverse quantization unit 58, inverse transform processing unit 60, summer 62, filter unit 64, and decoded picture buffer (DPB) 66.

[0162] As shown in FIG. 4, video encoder 20 receives video data and stores the received video data in video data memory 33. Video data memory 33 may store video data to be encoded by the components of video encoder 20. The video data stored in video data memory 33 may be obtained, for example, from video source 18. DPB 66 may be a reference picture memory that stores reference video data for use in encoding video data by video encoder 20, e.g., in intra- or inter-coding modes. Video data

memory 33 and DPB 66 may be formed by any of a variety of memory devices, such as dynamic random access memory (DRAM), including synchronous DRAM (SDRAM), magnetoresistive RAM (MRAM), resistive RAM (RRAM), or other types of memory devices. Video data memory 33 and DPB 66 may be provided by the same memory device or separate memory devices. In various examples, video data memory 33 may be on-chip with other components of video encoder 20, or off-chip relative to those components.

[0163] Partitioning unit 35 retrieves the video data from video data memory 33 and partitions the video data into video blocks. This partitioning may also include partitioning into slices, tiles, or other larger units, as well as video block partitioning, e.g., according to a quadtree structure of LCUs and CUs. Video encoder 20 generally illustrates the components that encode video blocks within a video slice to be encoded. The slice may be divided into multiple video blocks (and possibly into sets of video blocks referred to as tiles). Prediction processing unit 41 may select one of a plurality of possible coding modes, such as one of a plurality of intra coding modes or one of a plurality of inter coding modes, for the current video block based on error results (e.g., coding rate and the level of distortion). Prediction processing unit 41 may provide the resulting intra- or inter-coded block to summer 50 to generate residual block data and to summer 62 to reconstruct the encoded block for use as a reference picture.

[0164] Intra prediction unit 46 within prediction processing unit 41 may perform intra-predictive coding of the current video block relative to one or more neighboring blocks in the same frame or slice as the current block to be coded to provide spatial compression. Inter-prediction processing unit 44 within prediction processing unit 41 perform inter-predictive coding of the current video block relative to one or more predictive blocks in one or more reference pictures to provide temporal compression.

[0165] Inter-prediction processing unit 44 may be configured to determine the inter-prediction mode for a video slice according to a predetermined pattern for a video sequence. The predetermined pattern may designate video slices in the sequence as P slices or B slices. Motion estimation, performed by inter-prediction processing unit 44, is the process of generating motion vectors, which estimate motion for video blocks. A motion vector, for example, may indicate the displacement of a PU of a video block within a current video frame or picture relative to a predictive block within a reference picture.

[0166] A predictive block is a block that is found to closely match the PU of the video block to be coded in terms of pixel difference, which may be determined by sum of absolute difference (SAD), sum of square difference (SSD), or other difference metrics. In some examples, video encoder 20 may calculate values for sub-integer pixel positions of reference pictures stored in DPB 66. For example, video encoder 20 may interpolate values of one-quarter pixel positions, one-eighth pixel positions, or other fractional pixel positions of the reference picture. Therefore, inter-prediction processing unit 44 may perform a motion search relative to the full pixel positions and fractional pixel positions and output a motion vector with fractional pixel precision.

[0167] Inter-prediction processing unit 44 calculates a motion vector for a PU of a video block in an inter-coded slice by comparing the position of the PU to the position of a predictive block of a reference picture. The reference picture may be selected from a first reference picture list (List 0) or a second reference picture list (List 1), each of which identify one or more reference pictures stored in DPB 66. Inter-prediction processing unit 44 sends the calculated motion vector to entropy encoding unit 56.

[0168] Motion compensation, performed by inter-prediction processing unit 44, may involve fetching or generating the predictive block based on the motion vector determined by motion estimation, possibly performing interpolations to sub-pixel precision. Upon determining the motion vector for the PU of the current video block, inter-prediction processing unit 44 may locate the predictive block to which the motion vector points in one of the reference picture lists. Video encoder 20 forms a residual video block by subtracting pixel values of the predictive block from the pixel values of the current video block being coded, forming pixel difference values. The pixel difference values form residual data for the block, and may include both luma and chroma difference components. Summer 50 represents the component or components that perform this subtraction operation. Inter-prediction processing unit 44 may also generate syntax elements associated with the video blocks and the video slice for use by video decoder 30 in decoding the video blocks of the video slice.

[0169] Prediction processing unit 41 also includes a palette-based encoding unit 42 configured to perform various aspects of the palette-based coding techniques described in this disclosure. Palette-based encoding unit 42, for example, may perform palette-based decoding when a palette-based encoding mode is selected, e.g., for a block such as a CU or PU. For example, palette-based encoding unit 42 may be configured to generate a palette having entries indicating pixel values, select pixel values in a palette

to represent pixel values of at least some positions of a block of video data, and signal information associating at least some of the positions of the block of video data with entries in the palette corresponding, respectively, to the selected pixel values. Although various functions are described as being performed by palette-based encoding unit 42, some or all of such functions may be performed by other processing units, or a combination of different processing units.

[0170] Palette-based encoding unit 42 may be configured to generate any of the various syntax elements described herein. Accordingly, video encoder 20 may be configured to encode blocks of video data using palette-based code modes as described in this disclosure. Video encoder 20 may selectively encode a block of video data using a palette coding mode, or encode a block of video data using a different mode, e.g., such an HEVC inter-predictive or intra-predictive coding mode. The block of video data may be, for example, a CU or PU generated according to an HEVC coding process. A video encoder 20 may encode some blocks with inter-predictive temporal prediction or intra-predictive spatial coding modes and decode other blocks with the palette-based coding mode.

[0171] After prediction processing unit 41 generates the predictive block for the current video block, either via intra prediction, inter prediction, or palette mode coding, video encoder 20 forms a residual video block by subtracting the predictive block from the current video block. For some coding modes, however, video encoder 20 may not encode a residual block. When generated, the residual video data in the residual block may be included in one or more TUs and applied to transform processing unit 52.

Transform processing unit 52 transforms the residual video data into residual transform coefficients using a transform, such as a discrete cosine transform (DCT) or a conceptually similar transform. Transform processing unit 52 may convert the residual video data from a pixel domain to a transform domain, such as a frequency domain.

[0172] Transform processing unit 52 may send the resulting transform coefficients to quantization unit 54. Quantization unit 54 quantizes the transform coefficients to further reduce bit rate. The quantization process may reduce the bit depth associated with some or all of the coefficients. The degree of quantization may be modified by adjusting a quantization parameter. In some examples, quantization unit 54 may then perform a scan of the matrix including the quantized transform coefficients.

Alternatively, entropy encoding unit 56 may perform the scan.

[0173] Following quantization, entropy encoding unit 56 entropy encodes the quantized transform coefficients. For example, entropy encoding unit 56 may perform context adaptive variable length coding (CAVLC), context adaptive binary arithmetic coding (CABAC), syntax-based context-adaptive binary arithmetic coding (SBAC), probability interval partitioning entropy (PIPE) coding or another entropy encoding methodology or technique. Following the entropy encoding by entropy encoding unit 56, the encoded bitstream may be transmitted to video decoder 30, or archived for later transmission or retrieval by video decoder 30. Entropy encoding unit 56 may also entropy encode the motion vectors and the other syntax elements, including the various palette mode syntax elements discussed above, for the current video slice being coded.

[0174] Inverse quantization unit 58 and inverse transform processing unit 60 apply inverse quantization and inverse transformation, respectively, to reconstruct the residual block in the pixel domain for later use as a reference block of a reference picture. Inter-prediction processing unit 44 may calculate a reference block by adding the residual block to a predictive block of one of the reference pictures within one of the reference picture lists. Inter-prediction processing unit 44 may also apply one or more interpolation filters to the reconstructed residual block to calculate sub-integer pixel values for use in motion estimation. Summer 62 adds the reconstructed residual block to the motion compensated prediction block produced by inter-prediction processing unit 44 to produce a reconstructed block.

[0175] Filter unit 64 filters the reconstructed block (e.g. the output of summer 62) and stores the filtered reconstructed block in DPB 66 for uses as a reference block. The reference block may be used by inter-prediction processing unit 44 as a reference block to inter-predict a block in a subsequent video frame or picture. Filter unit 64 may apply one or more of deblocking filtering, sample adaptive offset (SAO) filtering, adaptive loop filtering (ALF), or other types of loop filtering. Filter unit 64 may apply deblocking filtering to filter block boundaries to remove blockiness artifacts from reconstructed video and may apply other types of filtering to improve overall coding quality. Additional loop filters (in loop or post loop) may also be used.

[0176] FIG. 5 is a block diagram illustrating an example video decoder 30 that may implement the techniques described in this disclosure. In the example of FIG. 5, video decoder 30 includes video data memory 78, entropy decoding unit 80, prediction processing unit 81, inverse quantization unit 86, inverse transform processing unit 88, summer 90, and DPB 94. Prediction processing unit 81 includes palette-based decoding

unit 82, motion compensation unit (MCU) 83, and intra prediction unit 84. Video decoder 30 may, in some examples, perform a decoding pass generally reciprocal to the encoding pass described with respect to video encoder 20 from FIG. 4.

[0177] During the decoding process, video decoder 30 receives an encoded video bitstream that represents video blocks of an encoded video slice and associated syntax elements from video encoder 20. Video decoder 30 stores the received encoded video bitstream in video data memory 78. Video data memory 78 may store video data, such as an encoded video bitstream, to be decoded by the components of video decoder 30. The video data stored in video data memory 78 may be obtained, for example, via link 16, from storage device 26, or from a local video source, such as a camera, or by accessing physical data storage media. Video data memory 78 may form a coded picture buffer (CPB) that stores encoded video data from an encoded video bitstream. DPB 94 may be a reference picture memory that stores reference video data for use in decoding video data by video decoder 30, e.g., in intra- or inter-coding modes. Video data memory 78 and DPB 94 may be formed by any of a variety of memory devices, such as DRAM, SDRAM, MRAM, RRAM, or other types of memory devices. Video data memory 78 and DPB 94 may be provided by the same memory device or separate memory devices. In various examples, video data memory 78 may be on-chip with other components of video decoder 30, or off-chip relative to those components.

[0178] Entropy decoding unit 80 of video decoder 30 entropy decodes the video data stored in video data memory 78 to generate quantized coefficients, motion vectors, and other syntax elements. Entropy decoding unit 80 forwards the motion vectors and other syntax elements, including the various palette mode syntax elements discussed above, to prediction processing unit 81. Video decoder 30 may receive the syntax elements at the video slice level and/or the video block level.

[0179] When the video slice is coded as an intra-coded (I) slice, intra prediction unit 84 of prediction processing unit 81 may generate prediction data for a video block of the current video slice based on a signaled intra prediction mode and data from previously decoded blocks of the current frame or picture. When the video frame is coded as an inter-coded (i.e., B or P) slice, motion compensation unit 83 of prediction processing unit 81 produces predictive blocks for a video block of the current video slice based on the motion vectors and other syntax elements received from entropy decoding unit 80. The predictive blocks may be produced from one of the reference pictures within one of

the reference picture lists. Video decoder 30 may construct the reference frame lists, List 0 and List 1, using default construction techniques based on reference pictures stored in DPB 94.

[0180] Motion compensation unit 83 determines prediction information for a video block of the current video slice by parsing the motion vectors and other syntax elements, and uses the prediction information to produce the predictive blocks for the current video block being decoded. For example, motion compensation unit 83 uses some of the received syntax elements to determine a prediction mode (e.g., intra- or inter-prediction) used to code the video blocks of the video slice, an inter-prediction slice type (e.g., B slice, P slice, or GPB slice), construction information for one or more of the reference picture lists for the slice, motion vectors for each inter-encoded video block of the slice, inter-prediction status for each inter-coded video block of the slice, and other information to decode the video blocks in the current video slice.

[0181] Motion compensation unit 83 may also perform interpolation based on interpolation filters. Motion compensation unit 83 may use interpolation filters as used by video encoder 20 during encoding of the video blocks to calculate interpolated values for sub-integer pixels of reference blocks. In this case, motion compensation unit 83 may determine the interpolation filters used by video encoder 20 from the received syntax elements and use the interpolation filters to produce predictive blocks.

[0182] Inverse quantization unit 86 inverse quantizes, i.e., de-quantizes, the quantized transform coefficients provided in the bitstream and decoded by entropy decoding unit 80. The inverse quantization process may include use of a quantization parameter calculated by video encoder 20 for each video block in the video slice to determine a degree of quantization and, likewise, a degree of inverse quantization that should be applied. Inverse transform processing unit 88 applies an inverse transform, e.g., an inverse DCT, an inverse integer transform, or a conceptually similar inverse transform process, to the transform coefficients in order to produce residual blocks in the pixel domain.

[0183] In accordance with various examples of this disclosure, video decoder 30 may be configured to perform palette-based coding. Palette-based decoding unit 82, for example, may perform palette-based decoding when a palette-based decoding mode is selected for a block such as a CU or PU. For example, palette-based decoding unit 82 may be configured to generate a palette having entries indicating pixel values. Furthermore, in this example, palette-based decoding unit 82 may receive information

associating at least some positions of a block of video data with entries in the palette. In this example, palette-based decoding unit 82 may select pixel values in the palette based on the information. Additionally, in this example, palette-based decoding unit 82 may reconstruct pixel values of the block based on the selected pixel values. Although various functions are described as being performed by palette-based decoding unit 82, some or all of such functions may be performed by other processing units, or a combination of different processing units.

[0184] Palette-based decoding unit 82 may receive palette coding mode information, and perform the above operations when the palette coding mode information indicates that the palette coding mode applies to the block. When the palette coding mode information indicates that the palette coding mode does not apply to the block, or when other mode information indicates the use of a different mode, palette-based decoding unit 82 decodes the block of video data using a non-palette based coding mode, e.g., such an HEVC inter-predictive or intra-predictive coding mode. The block of video data may be, for example, a CU or PU generated according to an HEVC coding process. A video decoder 30 may decode some blocks with inter-predictive temporal prediction or intra-predictive spatial coding modes and decode other blocks with the palette-based coding mode. The palette-based coding mode may comprise one of a plurality of different palette-based coding modes, or there may be a single palette-based coding mode.

[0185] After prediction processing unit generates the predictive block for the current video block using, for example, intra or inter prediction, video decoder 30 forms a reconstructed video block by summing the residual blocks from inverse transform processing unit 88 with the corresponding predictive blocks generated by motion compensation unit 83. Summer 90 represents the component or components that perform this summation operation. Filter unit 92 filters the reconstructed video block using, for example, one or more of deblocking filtering, SAO filtering, ALF filtering, or other types of filtering. Other loop filters (either in the coding loop or after the coding loop) may also be used to smooth pixel transitions or otherwise improve the video quality. The decoded video blocks in a given frame or picture are then stored in DPB 94, which stores reference pictures used for subsequent motion compensation. DPB 94 may be part of or separate from additional memory that stores decoded video for presentation on a display device, such as display device 32 of FIG. 1.

[0186] Video decoder 30 may, for example, be configured to receive a syntax element and perform a palette predictor initialization process according to the syntax element. A first value of the syntax element may, for example, indicate palette predictor initialization is used, and a second value of the syntax element indicates a palette predictor size is set equal to 0. The syntax element may be signaled, for example, in an SPS or PPS.

[0187] Additionally, or alternatively, video decoder 30 may be configured to receive in a PPS, a first palette predictor initializer syntax element and receive in an SPS, a second palette predictor initializer syntax element. Based on the first palette predictor initializer syntax element and the second palette predictor initializer syntax element, video decoder 30 may generate a palette predictor. To generate the palette predictor, video decoder 30 may include, in the palette predictor, palette entries identified in the SPS followed by palette entries identified in the PPS, or in the vice versa ordering. To generate the palette predictor, video decoder 30 may additionally or alternatively receive one or more reuse flags per entry of the palette predictor.

[0188] Additionally, or alternatively, video decoder 30 may be configured to receive in a PPS, a first scaling list and receive in an SPS, a second scaling list. Based on the first scaling list and the second scaling list, video decoder 30 may generate a scaling list for a picture.

[0189] Additionally, or alternatively, video decoder 30 may be configured to receive, in a first parameter set, a first set of palette predictor initializer entries. Subsequent to receiving the first parameter set, video decoder 30 may receive, in a second parameter set, a second set of palette predictor initializer entries. Video decoder 30 may perform palette predictor initialization using at least a portion of the first set of palette predictor initializer entries. To perform palette predictor initialization, video decoder 30 may, for example, use at least a portion of the second set of palette predictor initializer entries in addition to the portion of the first set of palette predictor initializer entries. In some examples, the first parameter set may be a PPS and the second parameter set may be an SPS.

[0190] Additionally, or alternatively, video decoder 30 may also be configured to form a palette predictor initializer using the portion of the first set of palette predictor initializer entries and add the portion of the second set of palette predictor initializer entries to the palette predictor initializer to form an updated palette predictor initializer. To add the portion of the second set of palette predictor initializer entries to the palette

predictor initializer to form the updated palette predictor initializer, video decoder 30 may be configured to add the portion of the second set of palette predictor initializer entries at a beginning of the palette predictor initializer to form the updated palette predictor initializer.

[0191] To add the portion of the second set of palette predictor initializer entries to the palette predictor initializer to form the updated palette predictor initializer comprises, video decoder 30 may be configured to add the portion of the second set of palette predictor initializer entries at an end of the palette predictor initializer to form the updated palette predictor initializer. Video decoder 30 may also be configured to determine that adding an entirety of the portion of the second set of palette predictor entries exceeds an initializer size limit, and based on the determination, exclude one or more entries of the portion of the second set of palette predictor entries from the updated palette predictor initializer. The updated palette predictor initializer may include both a sub-portion of the portion of the first set of palette predictor entries and a sub-portion of the portion of the second set of palette predictor entries.

[0192] FIG. 6 is a flowchart illustrating an example operation of a video encoder, in accordance with the techniques of this disclosure. The flowchart of FIG. 6 is provided as an example. In other examples, the flowcharts may include more, fewer, or different steps. The operation described in FIG. 6 may be performed by a video encoder such as video encoder 20 described above, but the operation of FIG. 6 is not limited to any particular type of video encoder.

[0193] In the example of FIG. 6, the video encoder encodes video data. The video encoder generates, for inclusion in an SPS syntax structure in a bitstream of encoded video data, a first syntax element (100). The first syntax element may, for example, be a one-bit flag such as the `sps_palette_predictor_initializer_present_flag` indicating that palette predictor initialization is enabled. The video encoder generates, for inclusion in the SPS syntax structure, one or more SPS-level palette predictor initializers (102). The video encoder generates, for inclusion in a PPS syntax structure in the bitstream of encoded video data, a second syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor initializers (104). The second syntax element may, for example, be a one-bit flag such as the `pps_palette_predictor_initializer_present_flag` described above. The second syntax element may further indicate that the palette predictor is to be generated without using any of the SPS-level palette predictor initializers. The video encoder generates, for

inclusion in the PPS syntax structure, a third syntax element indicating a number of PPS-level palette predictor initializers included in the PPS is equal to zero (106). The third syntax element may, for example, be the `pps_num_palette_predictor_initializer` syntax element described above. The video encoder may also generate, for inclusion in the SPS syntax structure, a fourth syntax element indicating a number of SPS-level palette predictor initializers included in the SPS. The fourth syntax element may, for example, be the syntax element `sps_num_palette_predictor_initializer_minus1` described above.

[0194] FIG. 7 is a flowchart illustrating an example operation of a video decoder, in accordance with the techniques of this disclosure. The flowchart of FIG. 7 is provided as an example. In other examples, the flowcharts may include more, fewer, or different steps. The operation described in FIG. 7 may be performed by a video decoder such as video decoder 30 described above, but the operation of FIG. 7 is not limited to any particular type of video decoder.

[0195] In the example of FIG. 7, the video decoder decodes video data. The video decoder receives in an SPS, a first syntax element indicating that palette predictor initialization is enabled (120). The first syntax element may, for example, be a one-bit flag, such as the syntax element `sps_palette_predictor_initializer_present_flag` described above. The video decoder receives, in the SPS, one or more SPS-level palette predictor initializers (122). The video decoder receives, in a PPS, a second syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor initializers (124). The second syntax element may, for example, be a one-bit flag such as the syntax element `pps_palette_predictor_initializer_present_flag` described above. The second syntax element may further indicate that the palette predictor is to be generated without using any of the SPS-level palette predictor initializers. The video decoder receives, in the PPS, a third syntax element indicating a number of PPS-level palette predictor initializers included in the PPS is equal to zero (126). The third syntax element may, for example, be the syntax element `pps_num_palette_predictor_initializer` described above.

[0196] To generate a palette table when the third syntax element indicates the number of PPS-level palette predictor initializers included in the PPS is equal to zero, the video decoder may receive one or more palette table entry values and generate the palette table by including the one or more received palette table entry values in the palette table. Thus, instead of generating the palette using a palette predictor to determine

some or all of the entries for the palette table, the video decoder directly receives all palette entries for the palette table. For a sample of a block of video data, the video decoder receives an index and reconstructs the block of video data using a palette table entry corresponding to the index for the sample.

[0197] FIG. 8 is a flowchart illustrating an example operation of a video decoder, in accordance with the techniques of this disclosure. The flowchart of FIG. 8 is provided as an example. In other examples, the flowcharts may include more, fewer, or different steps. The operation described in FIG. 8 may be performed by a video decoder such as video decoder 30 described above, but the operation of FIG. 8 is not limited to any particular type of video decoder.

[0198] In the example of FIG. 8, the video decoder decodes video data. The video decoder receives first palette predictor initialization data in an SPS (140). Based on the first palette predictor initialization data, the video decoder determines SPS palette predictor initializers (142). The video decoder receives second palette predictor initialization data in a PPS (144). Based on the second palette predictor initialization data, the video decoder determines PPS palette predictor initializers (146). The video decoder generates a palette predictor based on the SPS palette predictor initializers and the PPS palette predictor initializers (148). To generate the palette predictor, the video decoder may, for example, include at least one PPS palette predictor entry in the palette predictor and include at least one SPS palette predictor entry in the palette predictor.

[0199] The video decoder may also receive residue data indicating differences between the palette predictor and a palette to be used for a block of video data. Thus, based on the palette predictor and the residue data, the video decoder can determine the palette to be used for the block of video data.

[0200] In one or more examples, the functions described may be implemented in hardware, software, firmware, or any combination thereof. If implemented in software, the functions may be stored on or transmitted over, as one or more instructions or code, a computer-readable medium and executed by a hardware-based processing unit. Computer-readable media may include computer-readable storage media, which corresponds to a tangible medium such as data storage media, or communication media including any medium that facilitates transfer of a computer program from one place to another, e.g., according to a communication protocol. In this manner, computer-readable media generally may correspond to (1) tangible computer-readable storage media which is non-transitory or (2) a communication medium such as a signal or

carrier wave. Data storage media may be any available media that can be accessed by one or more computers or one or more processors to retrieve instructions, code and/or data structures for implementation of the techniques described in this disclosure. A computer program product may include a computer-readable medium.

[0201] By way of example, and not limitation, such computer-readable storage media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage, or other magnetic storage devices, flash memory, or any other medium that can be used to store desired program code in the form of instructions or data structures and that can be accessed by a computer. Also, any connection is properly termed a computer-readable medium. For example, if instructions are transmitted from a website, server, or other remote source using a coaxial cable, fiber optic cable, twisted pair, digital subscriber line (DSL), or wireless technologies such as infrared, radio, and microwave, then the coaxial cable, fiber optic cable, twisted pair, DSL, or wireless technologies such as infrared, radio, and microwave are included in the definition of medium. It should be understood, however, that computer-readable storage media and data storage media do not include connections, carrier waves, signals, or other transient media, but are instead directed to non-transient, tangible storage media. Disk and disc, as used herein, includes compact disc (CD), laser disc, optical disc, digital versatile disc (DVD), floppy disk and Blu-ray disc, where disks usually reproduce data magnetically, while discs reproduce data optically with lasers. Combinations of the above should also be included within the scope of computer-readable media.

[0202] Instructions may be executed by one or more processors, such as one or more DSPs, general purpose microprocessors, ASICs, FPGAs, or other equivalent integrated or discrete logic circuitry. Accordingly, the term “processor,” as used herein may refer to any of the foregoing structure or any other structure suitable for implementation of the techniques described herein. In addition, in some aspects, the functionality described herein may be provided within dedicated hardware and/or software modules configured for encoding and decoding, or incorporated in a combined codec. Also, the techniques could be fully implemented in one or more circuits or logic elements.

[0203] The techniques of this disclosure may be implemented in a wide variety of devices or apparatuses, including a wireless handset, an integrated circuit (IC) or a set of ICs (e.g., a chip set). Various components, modules, or units are described in this disclosure to emphasize functional aspects of devices configured to perform the disclosed techniques, but do not necessarily require realization by different hardware

units. Rather, as described above, various units may be combined in a codec hardware unit or provided by a collection of interoperative hardware units, including one or more processors as described above, in conjunction with suitable software and/or firmware.

[0204] Various examples have been described. These and other examples are within the scope of the following claims.

WHAT IS CLAIMED IS:

1. A method of decoding video data, the method comprising:
receiving, in a picture parameter set (PPS), a first syntax element indicating that
a palette predictor is to be generated using PPS-level palette predictor
initializers;
receiving, in the PPS, a second syntax element indicating a number of the PPS-
level palette predictor initializers included in the PPS is equal to zero;
and
decoding a block of video data based on the first syntax element and the second
syntax element.
2. The method of claim 1, further comprising:
receiving, in a sequence parameter set (SPS), a third syntax element indicating
that palette predictor initialization is enabled; and
receiving, in the SPS, one or more SPS-level palette predictor initializers.
3. The method of claim 2, further comprising:
receiving, in the SPS, a fourth syntax element indicating a number of SPS-level
palette predictor initializers included in the SPS.
4. The method of claim 2, wherein the first syntax element further indicates that the
palette predictor is to be generated without using any of the SPS-level palette predictor
entries.
5. The method of claim 2, wherein the third syntax element comprises a one-bit
flag.
6. The method of claim 1, wherein the first syntax element comprises a one-bit
flag.

7. The method of claim 1, further comprising:
in response to the second syntax element indicating the number of PPS-level palette predictor initializers included in the PPS is equal to zero,
receiving one or more palette table entry values; and
generating a palette table that includes the one or more palette table entry values.
8. The method of claim 7, further comprising:
for a sample of the block of video data, receiving an index;
reconstructing the block of video data using one of the palette table entries
corresponding to the index for the sample.
9. The method of claim 1, further comprising:
based on the second syntax element indicating the number of PPS-level palette predictor initializers included in the PPS is equal to zero, determining that one or more of (1) an indication of whether a palette predictor includes monochrome entries, (2) an indication of a bit depth for luma palette predictor entries, or (3) an indication of a bit depth for chroma palette predictor entries are not present in the PPS.
10. A method of encoding video data, the method comprising:
generating, for inclusion in a picture parameter set (PPS) syntax structure in the bitstream of encoded video data, a first syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor entries;
generating, for inclusion in the PPS syntax structure, a second syntax element indicating a number of PPS-level palette predictor entries included in the PPS is equal to zero; and
outputting a bitstream of encoded video data including the first syntax element and the second syntax element.

11. The method of claim 10, further comprising:
generating, for inclusion in a sequence parameter set (SPS) syntax structure in a bitstream of encoded video data, a third syntax element indicating that palette predictor initialization is enabled; and
generating, for inclusion in the SPS syntax structure, one or more SPS-level palette predictor entries.
12. The method of claim 11, wherein the first syntax element further indicates that the palette predictor is to be generated without using any of the SPS-level palette predictor entries.
13. The method of claim 11, further comprising:
generating, for inclusion in the SPS syntax structure, a fourth syntax element indicating a number of SPS-level palette predictor entries included in the SPS.
14. The method of claim 11, wherein the third syntax element comprises a one-bit flag.
15. The method of claim 10, wherein the first syntax element comprises a one-bit flag.
16. A device for decoding video data, the device comprising:
a memory configured to store the video data; and
one or more processors configured to:
receive, in a picture parameter set (PPS), a first syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor initializers;
receive, in the PPS, a second syntax element indicating a number of the PPS-level palette predictor initializers included in the PPS is equal to zero; and
decode a block of video data based on the first syntax element and the second syntax element.

17. The device of claim 16, wherein the one or more processors are further configured to:
- receive, in a sequence parameter set (SPS), a third syntax element indicating that palette predictor initialization is enabled; and
 - receive, in the SPS, one or more SPS-level palette predictor initializers.
18. The device of claim 17, wherein the one or more processors are further configured to:
- receive, in the SPS, a fourth syntax element indicating a number of SPS-level palette predictor initializers included in the SPS.
19. The device of claim 17, wherein the first syntax element further indicates that the palette predictor is to be generated without using any of the SPS-level palette predictor entries.
20. The device of claim 17, wherein the third syntax element comprises a one-bit flag.
21. The device of claim 16, wherein the first syntax element comprises a one-bit flag.
22. The device of claim 16, wherein the one or more processors are further configured to:
- in response to the second syntax element indicating the number of PPS-level palette predictor initializers included in the PPS is equal to zero, receive one or more palette table entry values; and
 - generate a palette table that includes the one or more palette table entry values.
23. The device of claim 22, wherein the one or more processors are further configured to:
- for a sample of the block of video data, receive an index; and
 - reconstructing the block of video data using one of the palette table entries corresponding to the index for the sample.

24. The device of claim 16, wherein the one or more processors are further configured to:

based on the second syntax element indicating the number of PPS-level palette predictor initializers included in the PPS is equal to zero, determine that one or more of (1) an indication of whether a palette predictor includes monochrome entries, (2) an indication of a bit depth for luma palette predictor entries, or (3) an indication of a bit depth for chroma palette predictor entries are not present in the PPS.

25. A device for encoding video data, the device comprising:
a memory configured to store the video data; and
one or more processors configured to:

generate, for inclusion in a picture parameter set (PPS) syntax structure in the bitstream of encoded video data, a first syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor entries;

generate, for inclusion in the PPS syntax structure, a second syntax element indicating a number of PPS-level palette predictor entries included in the PPS is equal to zero; and

output a bitstream of encoded video data including the first syntax element and the second syntax element.

26. The device of claim 25, wherein the one or more processors are further configured to:

generate, for inclusion in a sequence parameter set (SPS) syntax structure in a bitstream of encoded video data, a third syntax element indicating that palette predictor initialization is enabled; and

generate, for inclusion in the SPS syntax structure, one or more SPS-level palette predictor entries.

27. The device of claim 26, wherein the first syntax element further indicates that the palette predictor is to be generated without using any of the SPS-level palette predictor entries.

28. The device of claim 26, further comprising:
generate, for inclusion in the SPS syntax structure, a fourth syntax element
indicating a number of SPS-level palette predictor entries included in the
SPS.
29. The device of claim 26, wherein the third syntax element comprises a one-bit
flag.
30. The device of claim 25, wherein the first syntax element comprises a one-bit
flag.
31. An apparatus for decoding video data, the apparatus comprising:
means for receiving, in a picture parameter set (PPS), a first syntax element
indicating that a palette predictor is to be generated using PPS-level
palette predictor initializers;
means receiving, in the PPS, a second syntax element indicating a number of the
PPS-level palette predictor initializers included in the PPS is equal to
zero; and
means decoding a block of video data based on the first syntax element and the
second syntax element.
32. The apparatus of claim 31, further comprising:
means receiving, in a sequence parameter set (SPS), a third syntax element
indicating that palette predictor initialization is enabled; and
means receiving, in the SPS, one or more SPS-level palette predictor initializers.
33. The apparatus of claim 32, further comprising:
means receiving, in the SPS, a fourth syntax element indicating a number of
SPS-level palette predictor initializers included in the SPS.
34. The apparatus of claim 32, wherein the first syntax element further indicates that
the palette predictor is to be generated without using any of the SPS-level palette
predictor entries.

35. The apparatus of claim 32, wherein the third syntax element comprises a one-bit flag.
36. The apparatus of claim 31, wherein the first syntax element comprises a one-bit flag.
37. The apparatus of claim 31, further comprising:
means receiving one or more palette table entry values in response to the second syntax element indicating the number of PPS-level palette predictor initializers included in the PPS is equal to zero; and
means generating a palette table that includes the one or more palette table entry values.
38. The apparatus of claim 37, further comprising:
means receiving an index for a sample of the block of video data; and
means reconstructing the block of video data using one of the palette table entries corresponding to the index for the sample.
39. The apparatus of claim 31, further comprising:
means for determining that one or more of (1) an indication of whether a palette predictor includes monochrome entries, (2) an indication of a bit depth for luma palette predictor entries, or (3) an indication of a bit depth for chroma palette predictor entries are not present in the PPS based on the second syntax element indicating the number of PPS-level palette predictor initializers included in the PPS is equal to zero.

40. A computer readable storage medium storing instructions that when executed by one or more processors cause the one or more processors to:

receive, in a picture parameter set (PPS), a first syntax element indicating that a palette predictor is to be generated using PPS-level palette predictor initializers;

receive, in the PPS, a second syntax element indicating a number of the PPS-level palette predictor initializers included in the PPS is equal to zero; and

decode a block of video data based on the first syntax element and the second syntax element.

41. The computer readable storage medium of claim 40, storing further instructions that when executed cause the one or more processors to:

receive, in a sequence parameter set (SPS), a third syntax element indicating that palette predictor initialization is enabled; and

receive, in the SPS, one or more SPS-level palette predictor initializers.

42. The computer readable storage medium of claim 41, storing further instructions that when executed cause the one or more processors to:

receive, in the SPS, a fourth syntax element indicating a number of SPS-level palette predictor initializers included in the SPS.

43. The computer readable storage medium of claim 41, wherein the first syntax element further indicates that the palette predictor is to be generated without using any of the SPS-level palette predictor entries.

44. The computer readable storage medium of claim 41, wherein the third syntax element comprises a one-bit flag.

45. The computer readable storage medium of claim 40, wherein the first syntax element comprises a one-bit flag.

46. The computer readable storage medium of claim 40, storing further instructions that when executed cause the one or more processors to:
- in response to the second syntax element indicating the number of PPS-level palette predictor initializers included in the PPS is equal to zero, receive one or more palette table entry values; and
 - generate a palette table that includes the one or more palette table entry values.
47. The computer readable storage medium of claim 46, further comprising:
- for a sample of the block of video data, receive an index; and
 - reconstruct the block of video data using one of the palette table entries corresponding to the index for the sample.
48. The computer readable storage medium of claim 40, further comprising:
- based on the second syntax element indicating the number of PPS-level palette predictor initializers included in the PPS is equal to zero, determine that one or more of (1) an indication of whether a palette predictor includes monochrome entries, (2) an indication of a bit depth for luma palette predictor entries, or (3) an indication of a bit depth for chroma palette predictor entries are not present in the PPS.

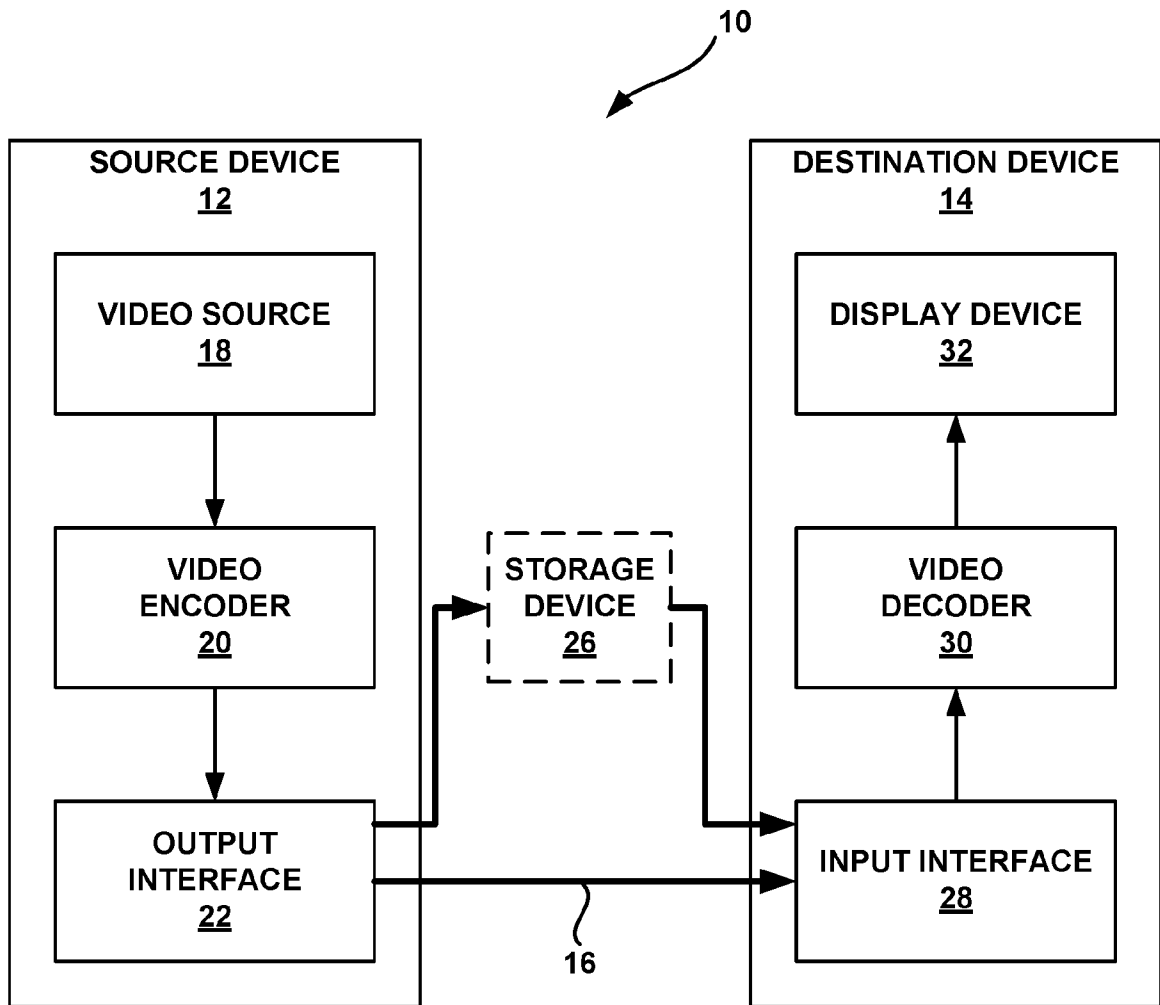


FIG. 1

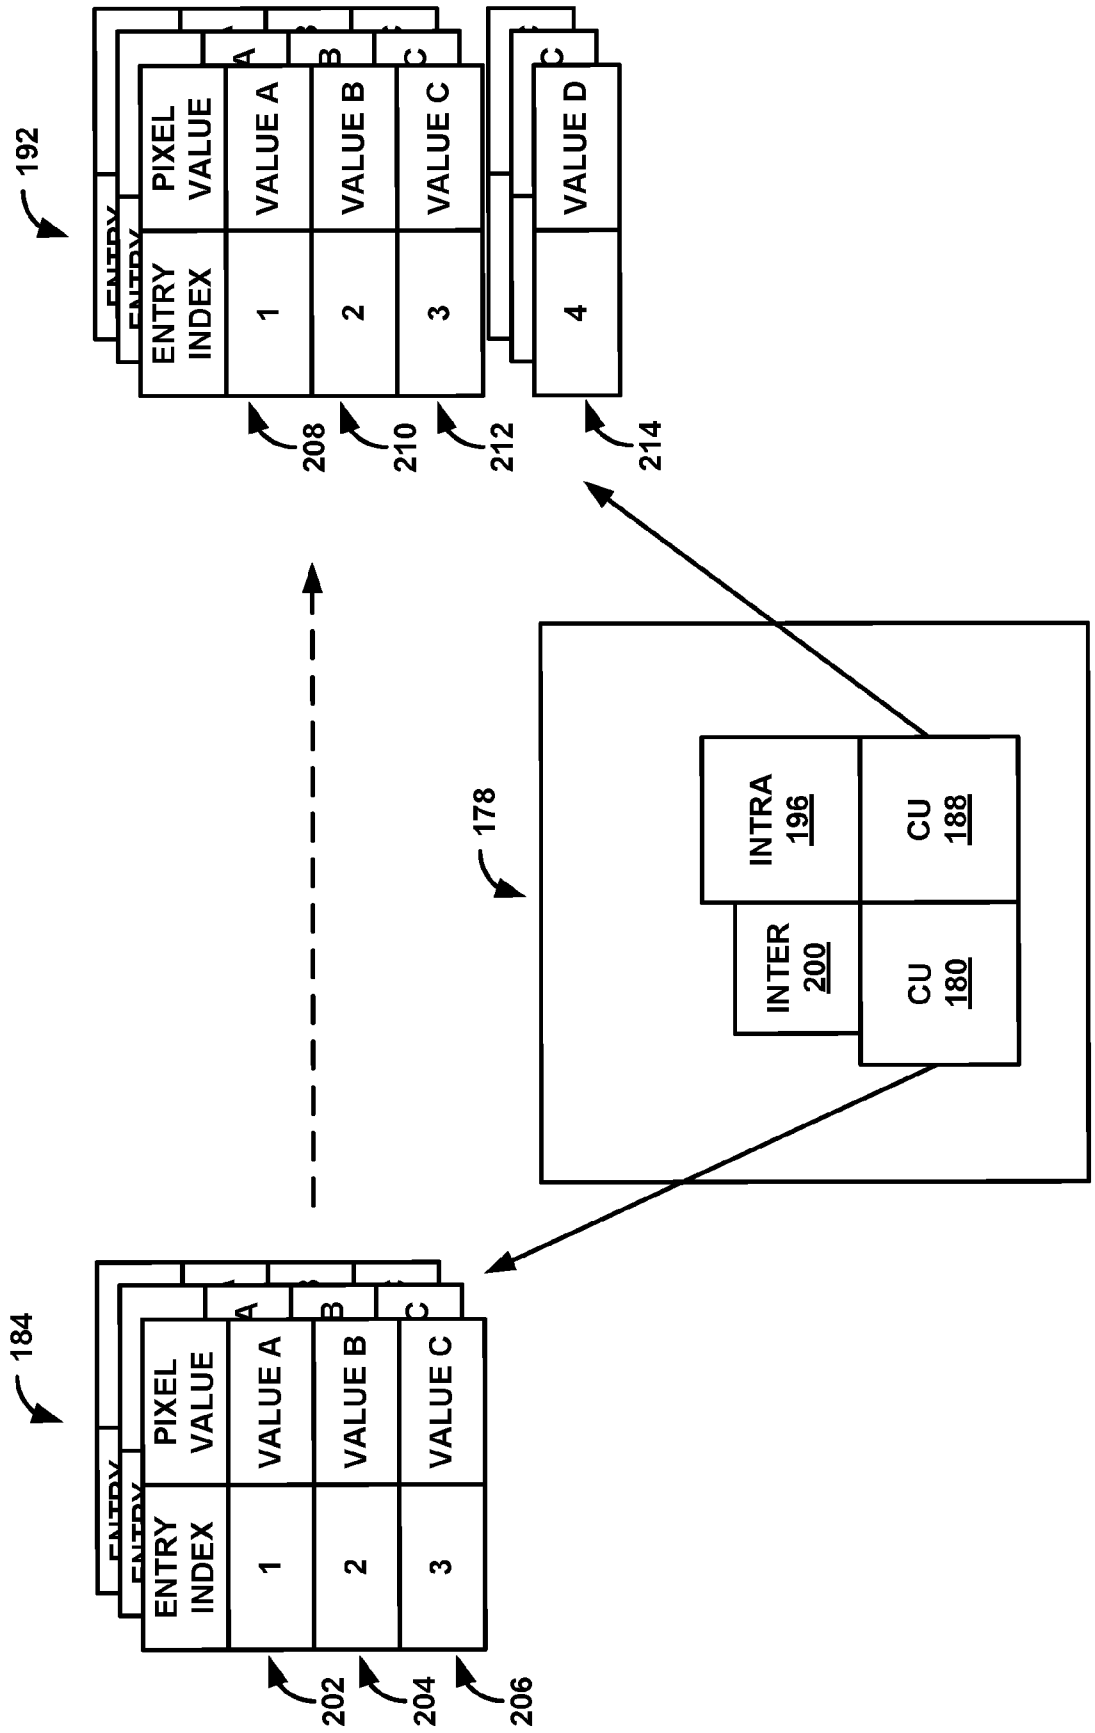


FIG. 2

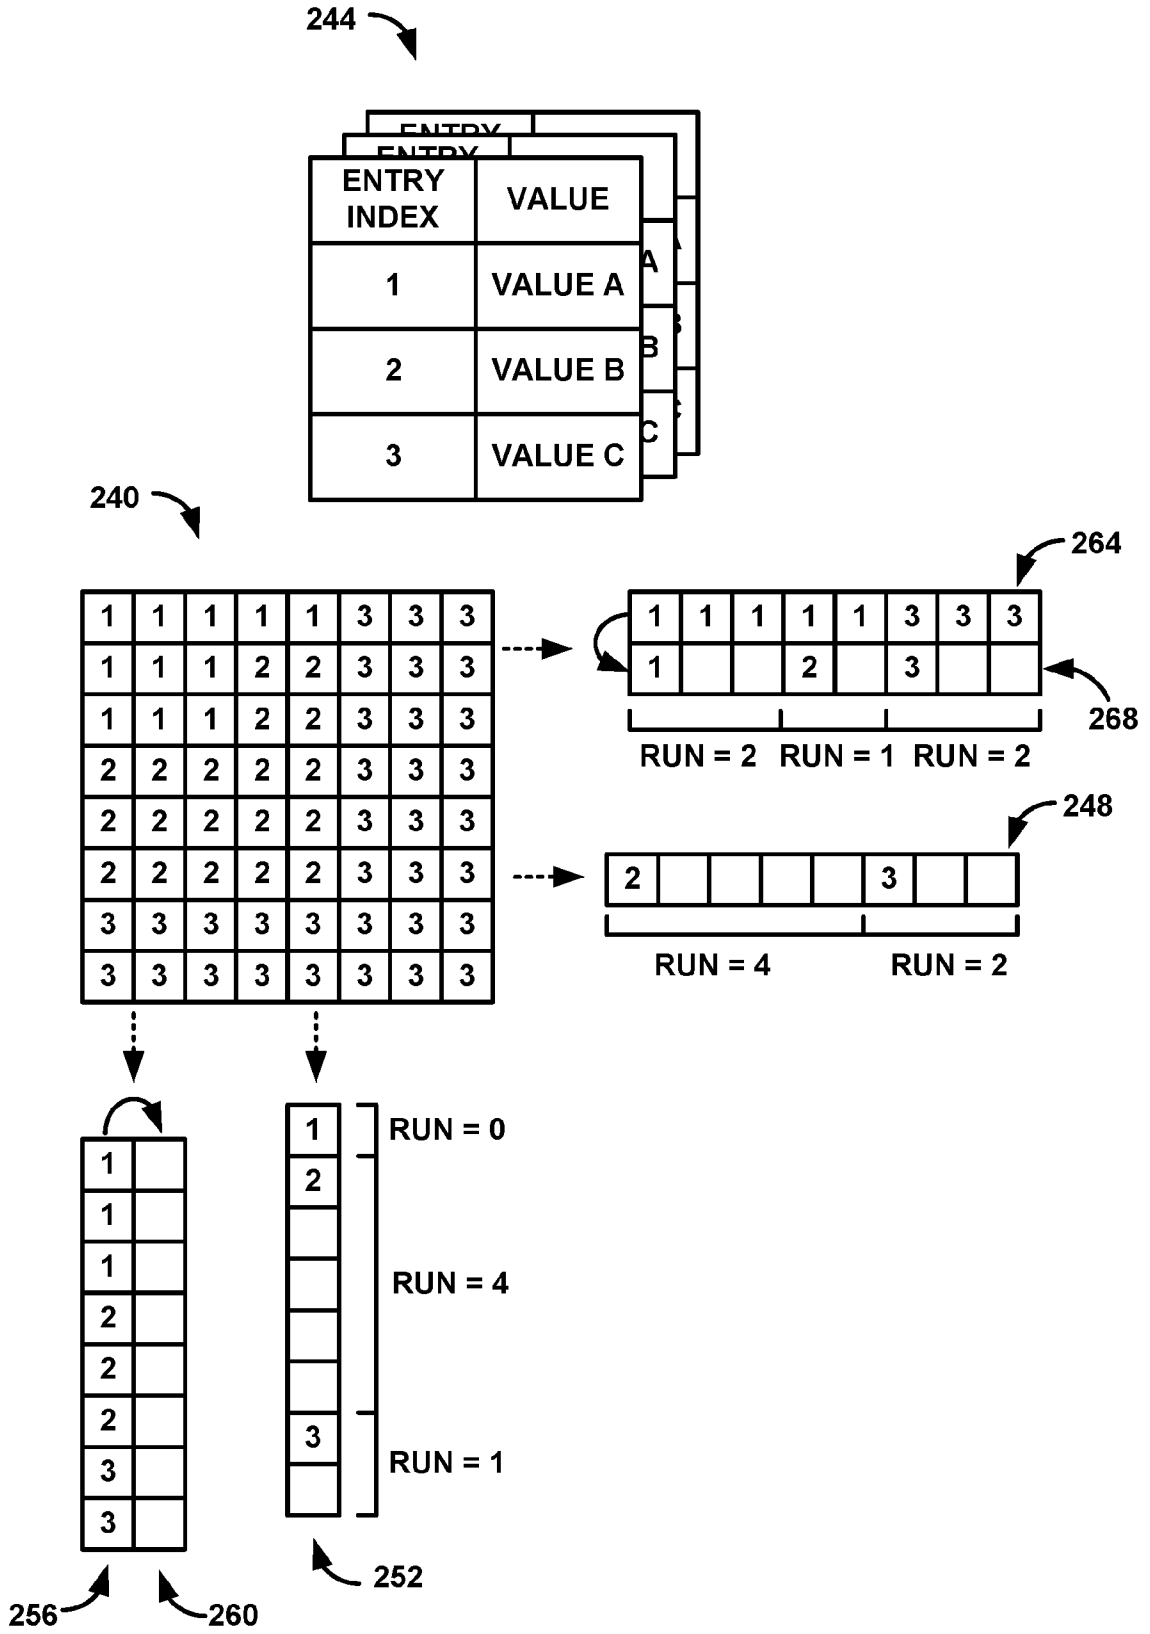


FIG. 3

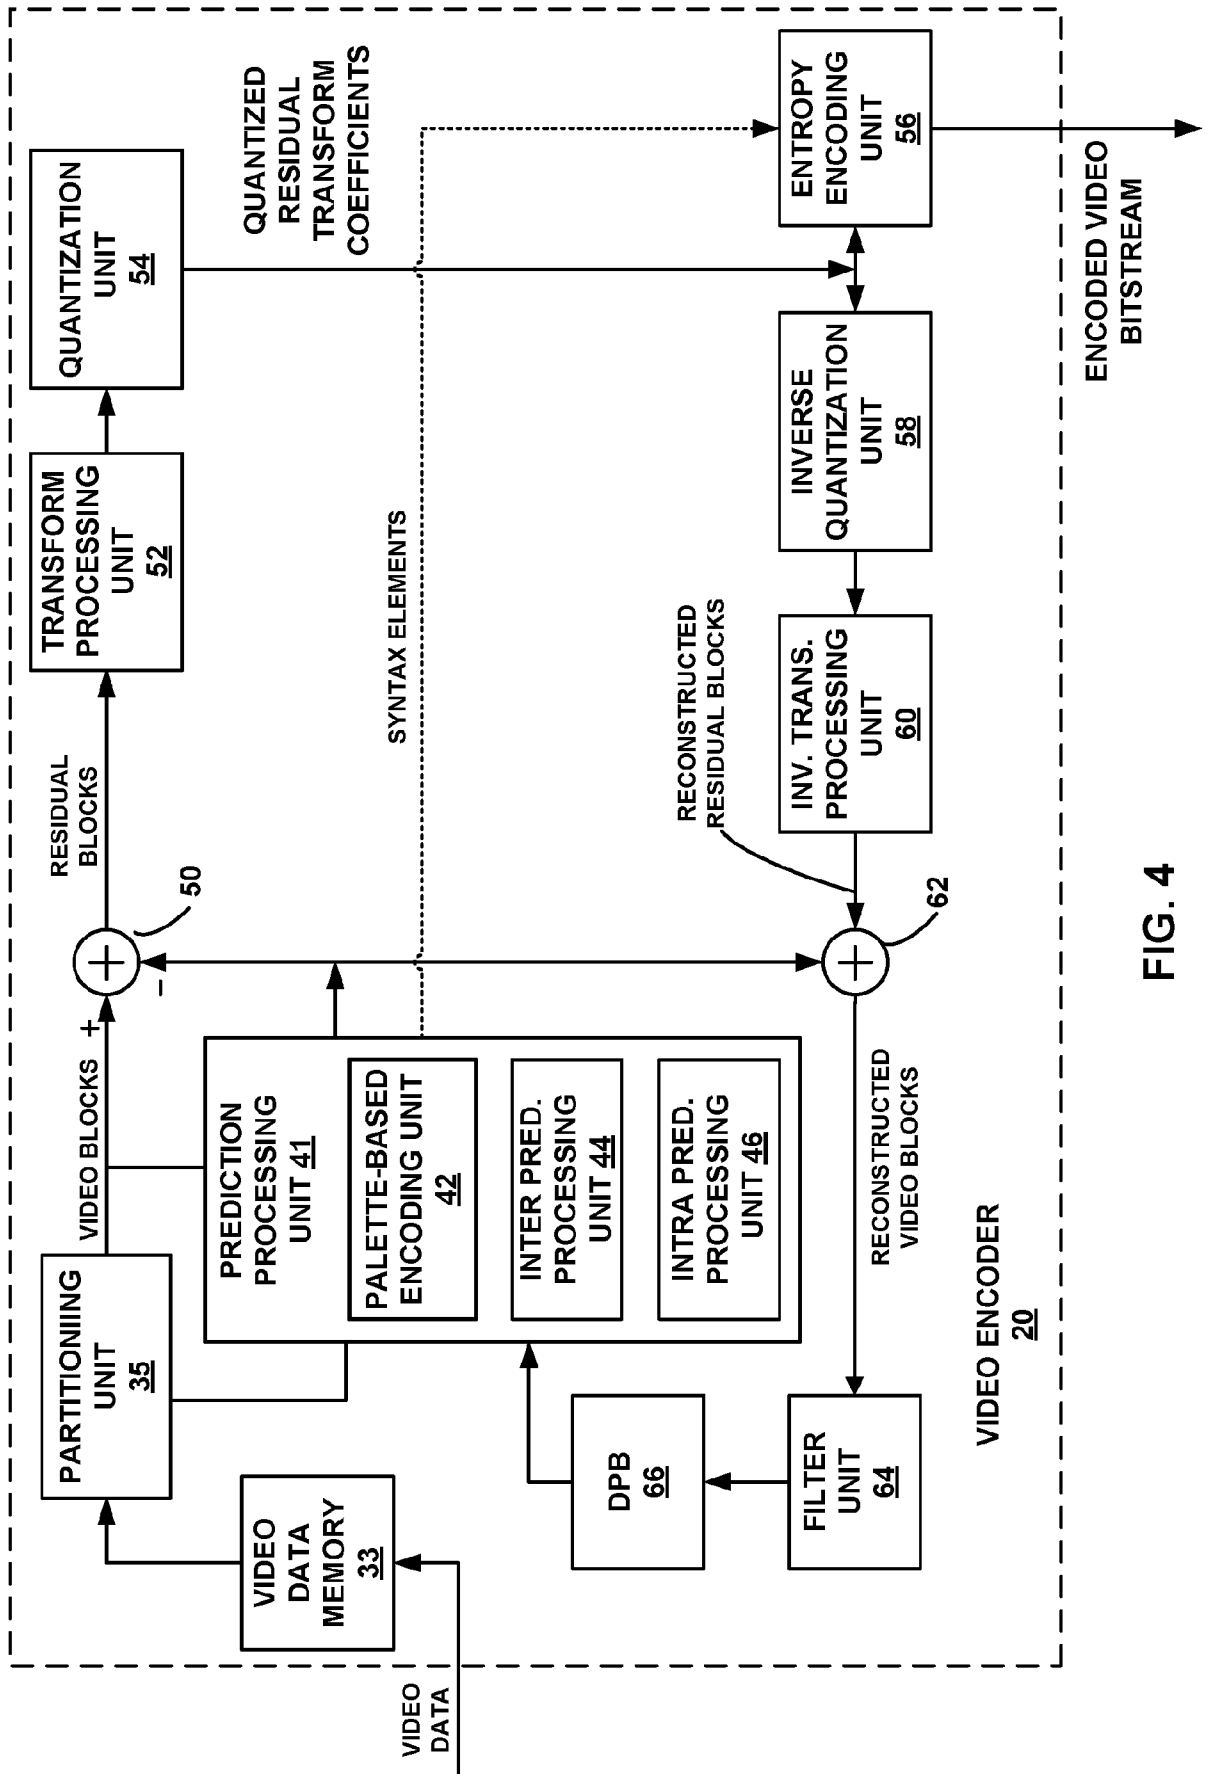


FIG. 4

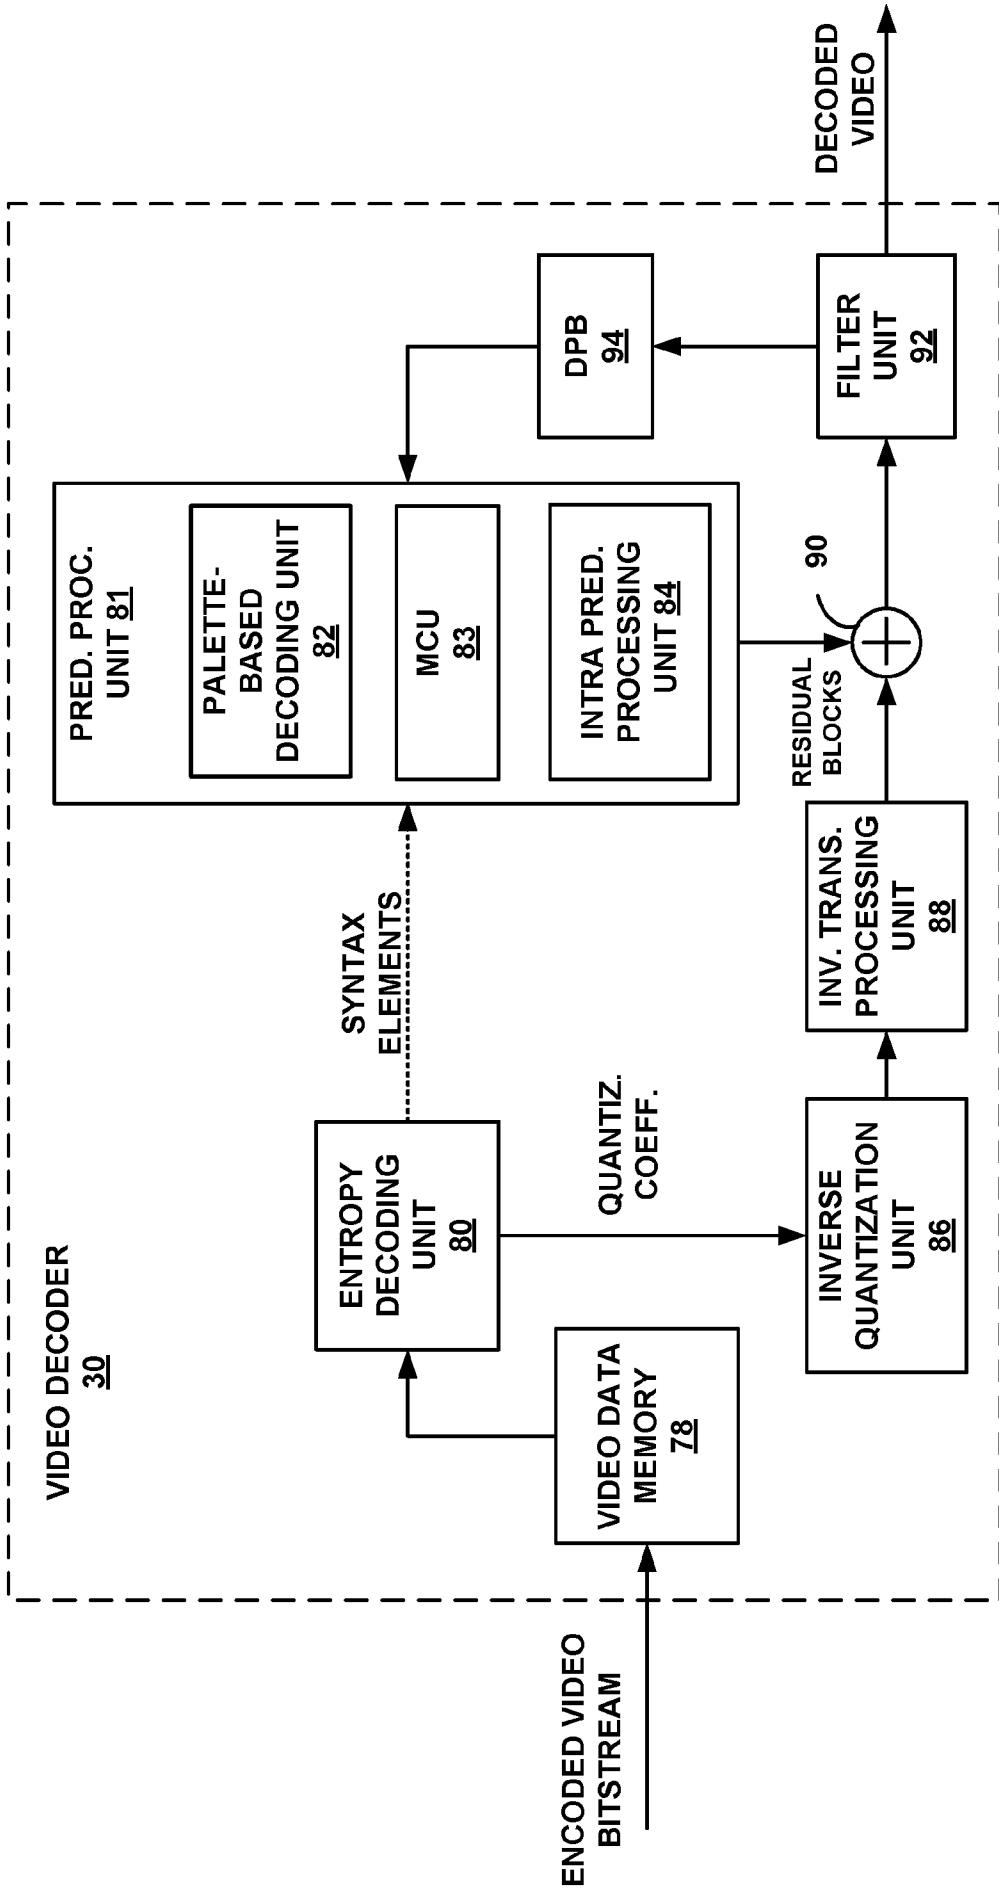


FIG. 5

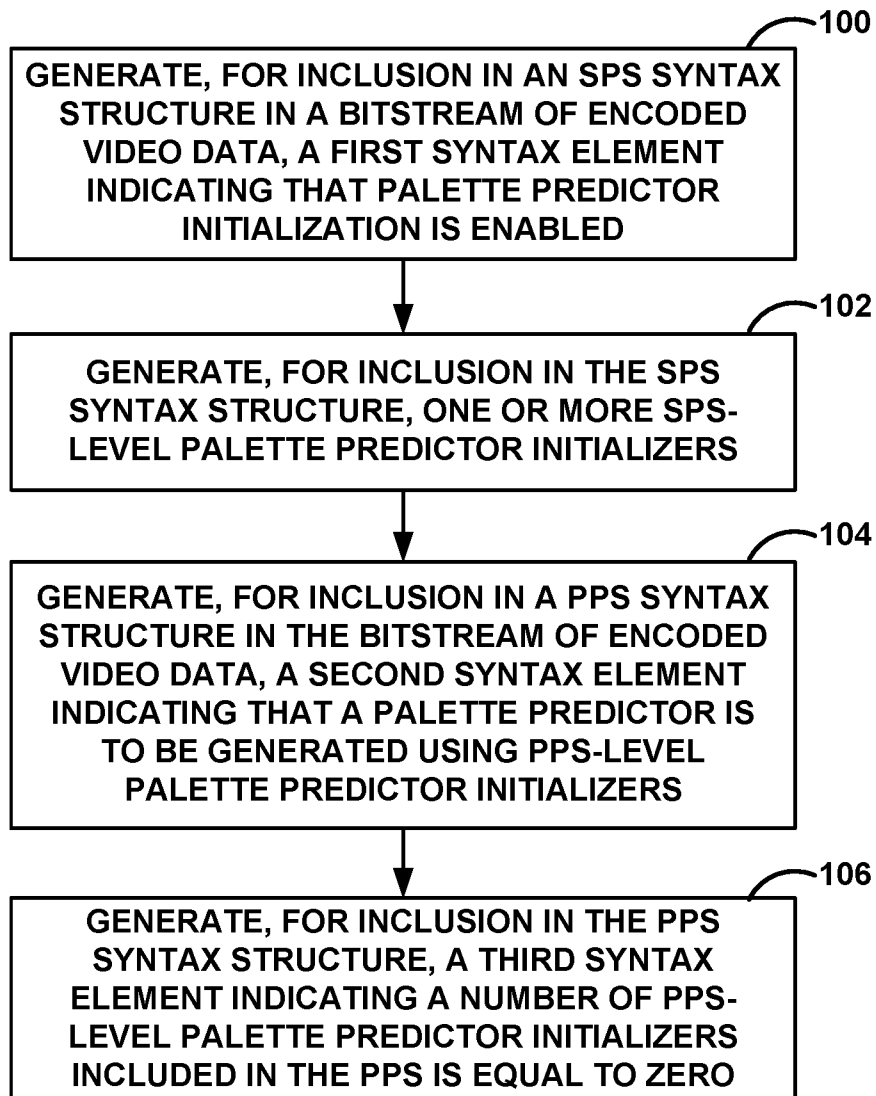


FIG. 6

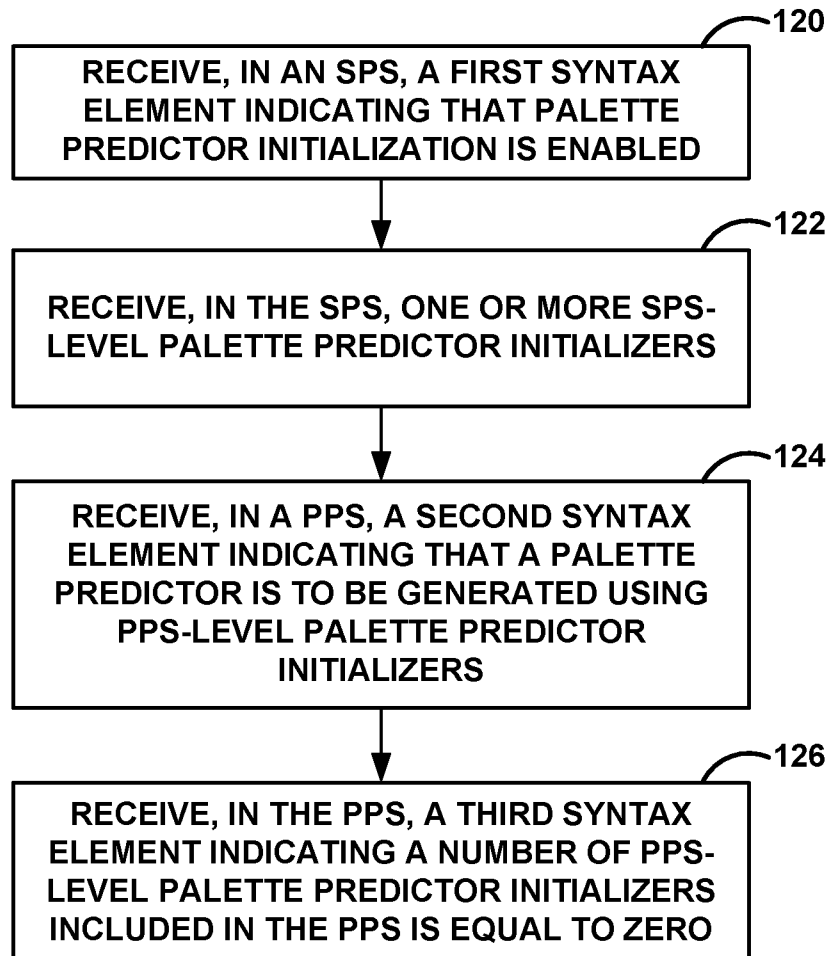


FIG. 7

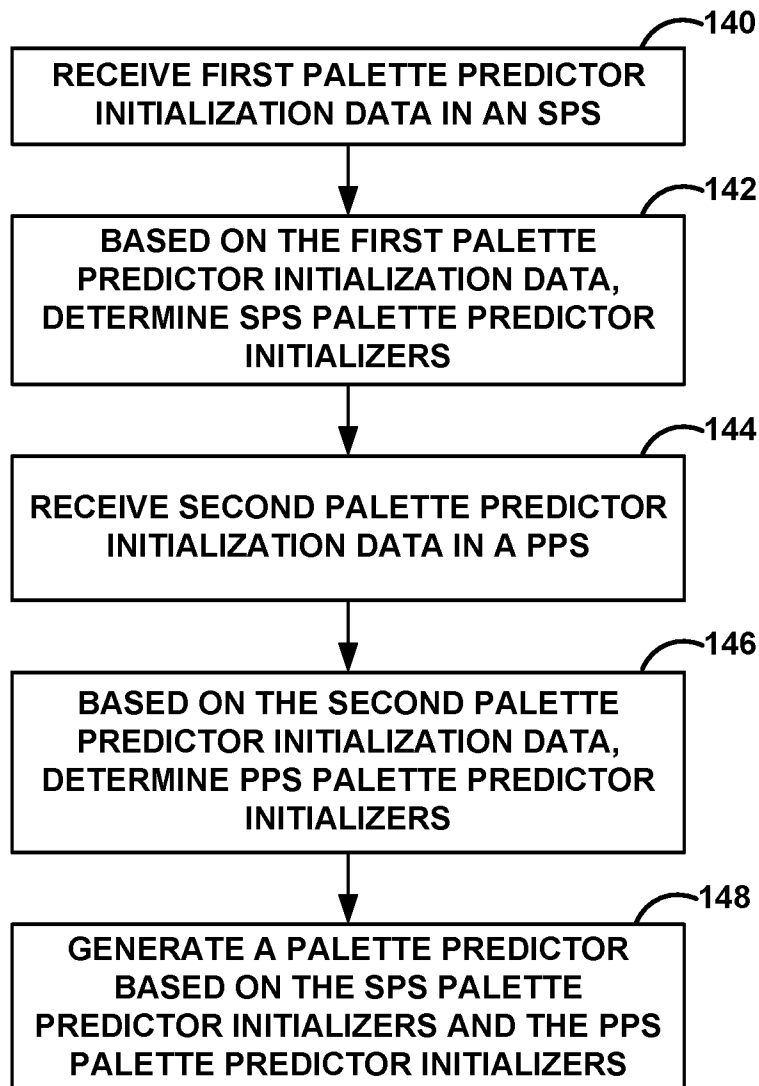


FIG. 8

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/051745

A. CLASSIFICATION OF SUBJECT MATTER
 INV. H04N19/186 H04N19/593 H04N19/70 H04N19/176 H04N19/463
 H04N19/93 H04N19/105
 ADD.
 According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
 Minimum documentation searched (classification system followed by classification symbols)
 H04N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
 EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	YE J ET AL: "On palette predictor initialization for Screen Content Coding", 21. JCT-VC MEETING; 19-6-2015 - 26-6-2015; WARSAW; (JOINT COLLABORATIVE TEAM ON VIDEO CODING OF ISO/IEC JTC1/SC29/WG11 AND ITU-T SG.16); URL: HTTP://WFTP3.ITU.INT/AV-ARCH/JCTVC-SITE/, , no. JCTVC-U0084, 10 June 2015 (2015-06-10), XP030117509, the whole document ----- -/--	1-48

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

<p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier application or patent but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p>	<p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art</p> <p>"&" document member of the same patent family</p>
---	---

Date of the actual completion of the international search 9 November 2016	Date of mailing of the international search report 21/11/2016
---	---

Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Ernst, Jens
--	--

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/051745

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>GISQUET C ET AL: "Non-CE1: On palette prediction for slices", 20. JCT-VC MEETING; 10-2-2105 - 18-2-2015; GENEVA; (JOINT COLLABORATIVE TEAM ON VIDEO CODING OF ISO/IEC JTC1/SC29/WG11 AND ITU-T SG.16); URL: HTTP://WFTP3.ITU.INT/AV-ARCH/JCTVC-SITE/,, no. JCTVC-T0048-v4, 15 February 2015 (2015-02-15), XP030117159, the whole document</p> <p style="text-align: center;">-----</p>	1-48
X	<p>JOSHI R ET AL: "HEVC Screen Content Coding Draft Text 4", 21. JCT-VC MEETING; 19-6-2015 - 26-6-2015; WARSAW; (JOINT COLLABORATIVE TEAM ON VIDEO CODING OF ISO/IEC JTC1/SC29/WG11 AND ITU-T SG.16); URL: HTTP://WFTP3.ITU.INT/AV-ARCH/JCTVC-SITE/,, no. JCTVC-U1005-v2, 5 September 2015 (2015-09-05), XP030117647, Sections 7.3.2.3.3, 7.4.3.2.3, 7.4.3.3.1, 7.4.3.3.3, 7.4.9.13 , 9.3.1 and 9.3.2</p> <p style="text-align: center;">-----</p>	1-48
X,P	<p>SEREGIN V ET AL: "SPS and PPS palette predictor initialization", 22. JCT-VC MEETING; 15-10-2015 - 21-10-2015; GENEVA; (JOINT COLLABORATIVE TEAM ON VIDEO CODING OF ISO/IEC JTC1/SC29/WG11 AND ITU-T SG.16); URL: HTTP://WFTP3.ITU.INT/AV-ARCH/JCTVC-SITE/,, no. JCTVC-V0042, 6 October 2015 (2015-10-06), XP030117689, the whole document</p> <p style="text-align: center;">-----</p>	1-48