

PALETTE PREDICTION IN PALETTE-BASED VIDEO CODING

[0001] This application claims the benefit of U.S. Provisional Application No. 61/845,824, filed July 12, 2013, U.S. Provisional Application No. 61/899,048, filed November 1, 2013, and U.S. Provisional Application No. 61/913,040, filed December 6, 2013, the entire contents of each of which is incorporated herein by reference.

TECHNICAL FIELD

[0002] This disclosure relates to video encoding and 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 High Efficiency Video Coding (HEVC) standard presently under development, 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. 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 indicates 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 coefficients, which then may be quantized. The quantized coefficients, initially arranged in a two-dimensional array, may be scanned in order to produce a one-dimensional vector of coefficients, and entropy coding may be applied to achieve even more compression.

[0006] A multiview coding bitstream may be generated by encoding views, e.g., from multiple perspectives. Some three-dimensional (3D) video standards have been developed that make use of multiview coding aspects. For example, different views may transmit left and right eye views to support 3D video. Alternatively, some 3D video coding processes may apply so-called multiview plus depth coding. In multiview plus depth coding, a 3D video bitstream may contain not only texture view components, but also depth view components. For example, each view may comprise one texture view component and one depth view component.

SUMMARY

[0007] Techniques of this disclosure relate to palette-based video coding. In palette-based coding, a video coder (e.g., a video encoder or a video decoder) may form a so-called “palette” as a table of colors or pixel values representing the video data of a particular area (e.g., a given block). In this way, rather than coding actual pixel values or their residuals for a current block of video data, the video coder may code index values for one or more of the pixels values of the current block, where the index values indicate entries in the palette that are used to represent the pixel values of the current block. A current palette for a current block of video data may be explicitly encoded and sent to the video decoder, predicted from previous palette entries, predicted from previous pixel values, or a combination thereof.

[0008] According to the techniques described in this disclosure for generating a current palette for a current block, the video decoder first determines one or more palette entries in a predictive palette that are copied to the current palette, and then determines a

number of new palette entries that are not in the predictive palette but that are included in the current palette. Based on this information, the video decoder calculates a size of the current palette to be equal to the sum of the number of the copied palette entries and the number of the new palette entries, and generates the current palette of the determined size including the copied palette entries and the new palette entries. A video encoder may perform similar techniques to generate the current palette for the current block. In addition, the video encoder may explicitly encode and send pixel values for the new palette entries to the video decoder. The techniques described in this disclosure may also include techniques for various combinations of one or more of signaling palette-based coding modes, transmitting palettes, predicting palettes, deriving palettes, or transmitting palette-based coding maps and other syntax elements.

[0009] In one example, this disclosure is directed toward a method of coding video data, the method comprising generating a predictive palette including palette entries that indicate pixel values, determining one or more of the palette entries in the predictive palette that are copied to a current palette for a current block of the video data, determining a number of new palette entries not in the predictive palette that are included in the current palette for the current block, calculating a size of the current palette equal to the sum of a number of the copied palette entries and the number of the new palette entries, and generating the current palette including the copied palette entries and the new palette entries. The method further comprises determining index values for one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.

[0010] In another example, this disclosure is directed toward an apparatus for coding video data, the apparatus comprising a memory storing video data, and one or more processors configured to generate a predictive palette including palette entries that indicate pixel values, determine one or more of the palette entries in the predictive palette that are copied to a current palette for a current block of the video data, determine a number of new palette entries not in the predictive palette that are included in the current palette for the current block, calculate a size of the current palette equal to the sum of a number of the copied palette entries and the number of the new palette entries, and generate the current palette including the copied palette entries and the new palette entries. The processors are further configured to determine index values for one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.

[0011] In another example, this disclosure is directed toward an apparatus for coding video data, the apparatus comprising means for generating a predictive palette including palette entries that indicate pixel values, means for determining one or more of the palette entries in the predictive palette that are copied to a current palette for a current block of the video data, means for determining a number of new palette entries not in the predictive palette that are included in the current palette for the current block, means for calculating a size of the current palette equal to the sum of a number of the copied palette entries and the number of the new palette entries, means for generating the current palette including the copied palette entries and the new palette entries, and means for determining index values for one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.

[0012] In a further example, this disclosure is directed toward a non-transitory computer-readable medium storing instructions thereon that, when executed, cause one or more processors to generate a predictive palette including palette entries that indicate pixel values, determine one or more of the palette entries in the predictive palette that are copied to a current palette for a current block of the video data, determine a number of new palette entries not in the predictive palette that are included in the current palette for the current block, calculate a size of the current palette equal to the sum of a number of the copied palette entries and the number of the new palette entries, generate the current palette including the copied palette entries and the new palette entries, and determine index values for one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.

[0013] 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

[0014] FIG. 1 is a block diagram illustrating an example video coding system that may utilize the techniques described in this disclosure.

[0015] FIG. 2 is a block diagram illustrating an example video encoder that may implement the techniques described in this disclosure.

[0016] FIG. 3 is a block diagram illustrating an example video decoder that may implement the techniques described in this disclosure.

[0017] FIG. 4 is a conceptual diagram illustrating an example of determining a palette for coding video data, consistent with techniques of this disclosure.

[0018] FIG. 5 is a conceptual diagram illustrating examples of determining indices to a palette for a video block, consistent with techniques of this disclosure.

[0019] FIG. 6 is a conceptual diagram illustrating examples of determining a geometric edge of a video block using a run of palette indices for the luma component adaptively downsampled for the chroma components, consistent with techniques of this disclosure.

[0020] FIG. 7 is a flowchart illustrating an example process for encoding prediction residual video data using a palette-based coding mode, consistent with techniques of this disclosure.

[0021] FIG. 8 is a flowchart illustrating an example process for decoding prediction residual video data using a palette-based coding mode, consistent with techniques of this disclosure.

[0022] FIG. 9 is a flowchart illustrating an example process for generating a palette for palette-based coding, consistent with techniques of this disclosure.

DETAILED DESCRIPTION

[0023] This disclosure includes techniques for video coding and compression. In particular, this disclosure describes techniques for palette-based coding of video data. In traditional video coding, images are assumed to be continuous-tone and spatially smooth. Based on these assumptions, various tools have been developed such as block-based transform, filtering, etc., and such tools have shown good performance for natural content videos.

[0024] In applications like remote desktop, collaborative work and wireless display, however, computer generated screen content (e.g., such as text or computer graphics) may be the dominant content to be compressed. This type of content tends to have discrete-tone and feature sharp lines, and high contrast object boundaries. The assumption of continuous-tone and smoothness may no longer apply for screen content, and thus traditional video coding techniques may not be efficient ways to compress video data including screen content.

[0025] This disclosure describes palette-based coding, which may be particularly suitable for screen generated content coding. For example, assuming a particular area of

video data has a relatively small number of colors, a video coder (a video encoder or video decoder) may form a so-called “palette” as a table of colors or pixel values representing the video data of the particular area (e.g., a given block). For example, the palette may include the most dominant pixel values in the given block. In some cases, the most dominant pixel values may include the one or more pixel values that occur most frequently within the block. In addition, in some cases a threshold value may be applied to define whether a pixel value is included as one of the most dominant pixel values in the block. According to this disclosure, rather than coding actual pixel values or their residuals for a current block of video data, the video coder may code index values indicative of one or more of the pixels values of the current block, where the index values indicate entries in the palette that are used to represent the pixel values of the current block.

[0026] For example, the video encoder may encode a block of video data by determining the palette for the block (e.g., coding the palette explicitly, predicting the palette, or a combination thereof), locating an entry in the palette to represent one or more of the pixel values, and encoding the block with index values that indicate the entry in the palette used to represent the pixel values of the block. In some examples, the video encoder may signal the index values in an encoded bitstream. A video decoder may obtain, from an encoded bitstream, a palette for a block, as well as index values for the pixels of the block. The video decoder may relate the index values of the pixels to entries of the palette to reconstruct the pixel values of the block.

[0027] The examples above are intended to provide a general description of palette-based coding. In various examples, the techniques described in this disclosure may include techniques for various combinations of one or more of signaling palette-based coding modes, transmitting palettes, predicting palettes, deriving palettes, or transmitting palette-based coding maps and other syntax elements. Such techniques may improve video coding efficiency, e.g., requiring fewer bits to represent screen generated content.

[0028] For example, a current palette for a current block of video data may be explicitly encoded and sent to the video decoder, predicted from previous palette entries, predicted from previous pixel values, or a combination thereof. According to the techniques described in this disclosure for generating a current palette for a current block, the video decoder first determines one or more palette entries in a predictive palette that are copied to the current palette, and then determines a number of new palette entries that

are not in the predictive palette but that are included in the current palette. Based on this information, the video decoder calculates a size of the current palette to be equal to the sum of the number of the copied palette entries and the number of the new palette entries, and generates the current palette of the determined size including the copied palette entries and the new palette entries. A video encoder may perform similar techniques to generate the current palette for the current block. In addition, the video encoder may explicitly encode and send pixel values for the new palette entries to the video decoder.

[0029] In some examples of this disclosure, the techniques for palette-based coding of video data may be used with one or more other coding techniques, such as techniques for inter-predictive coding or intra-predictive coding of video data. For example, as described in greater detail below, an encoder or decoder, or combined encoder-decoder (codec), may be configured to perform inter- and intra-predictive coding, as well as palette-based coding. In some examples, the palette-based coding techniques may be configured for use in one or more coding unit (CU) modes of High Efficiency Video Coding (HEVC). In other examples, the palette-based coding techniques can be used independently or as part of other existing or future systems or standards.

[0030] High Efficiency Video Coding (HEVC) is a new video coding standard developed by the Joint Collaboration Team on Video Coding (JCT-VC) of ITU-T Video Coding Experts Group (VCEG) and ISO/IEC Motion Picture Experts Group (MPEG). A recent draft of the HEVC standard, referred to as "HEVC Draft 10" or "WD10," is described in document JCTVC-L1003v34, Bross et al., "High Efficiency Video Coding (HEVC) Text Specification Draft 10 (for FDIS & Last Call)," Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T SG16 WP3 and ISO/IEC JTC1/SC29/WG11, 12th Meeting: Geneva, CH, 14-23 January 2013, available from:

http://phenix.int-evry.fr/jct/doc_end_user/documents/12_Geneva/wg11/JCTVC-L1003-v34.zip.

[0031] With respect to the HEVC framework, as an example, the palette-based coding techniques may be configured to be used as a CU mode. In other examples, the palette-based coding techniques may be configured to be used as a PU mode in the framework of HEVC. Accordingly, all of the following disclosed processes described in the context of a CU mode may, additionally or alternatively, apply to PU. However, these HEVC-based examples should not be considered a restriction or limitation of the palette-based coding techniques described herein, as such techniques may be applied to

work independently or as part of other existing or yet to be developed systems/standards. In these cases, the unit for palette coding can be square blocks, rectangular blocks or even regions of non-rectangular shape.

[0032] FIG. 1 is a block diagram illustrating an example video coding system 10 that may utilize the techniques of this disclosure. As used herein, the term “video coder” refers generically to both video encoders and video decoders. In this disclosure, the terms “video coding” or “coding” may refer generically to video encoding or video decoding. Video encoder 20 and video decoder 30 of video coding system 10 represent examples of devices that may be configured to perform techniques for palette-based video coding in accordance with various examples described in this disclosure. For example, video encoder 20 and video decoder 30 may be configured to selectively code various blocks of video data, such as CUs or PUs in HEVC coding, using either palette-based coding or non-palette based coding. Non-palette based coding modes may refer to various inter-predictive temporal coding modes or intra-predictive spatial coding modes, such as the various coding modes specified by HEVC Draft 10.

[0033] As shown in FIG. 1, video coding system 10 includes a source device 12 and a destination device 14. Source device 12 generates encoded video data. Accordingly, source device 12 may be referred to as a video encoding device or a video encoding apparatus. Destination device 14 may decode the encoded video data generated by source device 12. Accordingly, destination device 14 may be referred to as a video decoding device or a video decoding apparatus. Source device 12 and destination device 14 may be examples of video coding devices or video coding apparatuses.

[0034] Source device 12 and destination device 14 may comprise a wide range of devices, including desktop computers, mobile computing devices, notebook (e.g., laptop) computers, tablet computers, set-top boxes, telephone handsets such as so-called “smart” phones, televisions, cameras, display devices, digital media players, video gaming consoles, in-car computers, or the like.

[0035] Destination device 14 may receive encoded video data from source device 12 via a channel 16. Channel 16 may comprise one or more media or devices capable of moving the encoded video data from source device 12 to destination device 14. In one example, channel 16 may comprise one or more communication media that enable source device 12 to transmit encoded video data directly to destination device 14 in real-time. In this example, source device 12 may modulate the encoded video data according to a communication standard, such as a wireless communication protocol, and

may transmit the modulated video data to destination device 14. The one or more communication media may include wireless and/or wired communication media, such as a radio frequency (RF) spectrum or one or more physical transmission lines. The one or more communication media may form part of a packet-based network, such as a local area network, a wide-area network, or a global network (e.g., the Internet). The one or more communication media may include routers, switches, base stations, or other equipment that facilitate communication from source device 12 to destination device 14.

[0036] In another example, channel 16 may include a storage medium that stores encoded video data generated by source device 12. In this example, destination device 14 may access the storage medium via disk access or card access. The storage medium may include a variety of locally-accessed data storage media such as Blu-ray discs, DVDs, CD-ROMs, flash memory, or other suitable digital storage media for storing encoded video data.

[0037] In a further example, channel 16 may include a file server or another intermediate storage device that stores encoded video data generated by source device 12. In this example, destination device 14 may access encoded video data stored at the file server or other intermediate storage device via streaming or download. The file server may be a type of server capable of storing encoded video data and transmitting the encoded video data to destination device 14. Example file servers include web servers (e.g., for a website), file transfer protocol (FTP) servers, network attached storage (NAS) devices, and local disk drives.

[0038] Destination device 14 may access the encoded video data through a standard data connection, such as an Internet connection. Example types of data connections may include wireless channels (e.g., Wi-Fi connections), wired connections (e.g., DSL, cable modem, etc.), or combinations of both that are suitable for accessing encoded video data stored on a file server. The transmission of encoded video data from the file server may be a streaming transmission, a download transmission, or a combination of both.

[0039] The techniques of this disclosure are not limited to wireless applications or settings. The techniques may be applied to video coding in support 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 video data for storage on a data storage medium, decoding of video data stored on a data storage medium, or other applications. In some examples,

video coding 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.

[0040] Video coding system 10 illustrated in FIG. 1 is merely an example and the techniques of this disclosure may apply to video coding settings (e.g., video encoding or video decoding) that do not necessarily include any data communication between the encoding and decoding devices. In other examples, data is retrieved from a local memory, streamed over a network, or the like. A video encoding device may encode and store data to memory, and/or a video decoding device may retrieve and decode data from memory. In many examples, the encoding and decoding is performed by devices that do not communicate with one another, but simply encode data to memory and/or retrieve and decode data from memory.

[0041] In the example of FIG. 1, source device 12 includes a video source 18, a video encoder 20, and an output interface 22. In some examples, output interface 22 may include a modulator/demodulator (modem) and/or a transmitter. Video source 18 may include a video capture device, e.g., a video camera, a video archive containing previously-captured video data, a video feed interface to receive video data from a video content provider, and/or a computer graphics system for generating video data, or a combination of such sources of video data.

[0042] Video encoder 20 may encode video data from video source 18. In some examples, source device 12 directly transmits the encoded video data to destination device 14 via output interface 22. In other examples, the encoded video data may also be stored onto a storage medium or a file server for later access by destination device 14 for decoding and/or playback.

[0043] In the example of FIG. 1, destination device 14 includes an input interface 28, a video decoder 30, and a display device 32. In some examples, input interface 28 includes a receiver and/or a modem. Input interface 28 may receive encoded video data over channel 16. Display device 32 may be integrated with or may be external to destination device 14. In general, display device 32 displays decoded video data. Display device 32 may comprise 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.

[0044] This disclosure may generally refer to video encoder 20 “signaling” or “transmitting” certain information to another device, such as video decoder 30. The

term “signaling” or “transmitting” may generally refer to the communication of syntax elements and/or other data used to decode the compressed video data. Such communication may occur in real- or near-real-time. Alternately, such communication may occur over a span of time, such as might occur when storing syntax elements to a computer-readable storage medium in an encoded bitstream at the time of encoding, which then may be retrieved by a decoding device at any time after being stored to this medium. Thus, while video decoder 30 may be referred to as “receiving” certain information, the receiving of information does not necessarily occur in real- or near-real-time and may be retrieved from a medium at some time after storage.

[0045] Video encoder 20 and video decoder 30 each may be implemented as any of a variety of suitable circuitry, such as one or more microprocessors, digital signal processors (DSPs), application-specific integrated circuits (ASICs), field-programmable gate arrays (FPGAs), discrete logic, hardware, or any combinations thereof. If the techniques are implemented partially in software, a device may store instructions for the software in a suitable, non-transitory computer-readable storage medium and may execute the instructions in hardware using one or more processors to perform the techniques of this disclosure. Any of the foregoing (including hardware, software, a combination of hardware and software, etc.) may be considered to be one or more processors. 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.

[0046] In some examples, video encoder 20 and video decoder 30 operate according to a video compression standard, such as HEVC standard mentioned above, and described in HEVC Draft 10. In addition to the base HEVC standard, there are ongoing efforts to produce scalable video coding, multiview video coding, and 3D coding extensions for HEVC. In addition, palette-based coding modes, e.g., as described in this disclosure, may be provided for extension of the HEVC standard. In some examples, the techniques described in this disclosure for palette-based coding may be applied to encoders and decoders configured to operation according to other video coding standards, such as the ITU-T-H.264/AVC standard or future standards. Accordingly, application of a palette-based coding mode for coding of coding units (CUs) or prediction units (PUs) in an HEVC codec is described for purposes of example.

[0047] In HEVC and other video coding standards, 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.

[0048] 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 be 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. 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 the raster scan. A coded slice may comprise a slice header and slice data. The slice header of a slice may be a syntax structure that includes syntax elements that provide information about the slice. The slice data may include coded CTUs of the slice.

[0049] This disclosure may use the term “video unit” or “video block” or “block” to refer to one or more sample blocks and syntax structures used to code samples of the one or more blocks of samples. Example types of video units or blocks may include CTUs, CUs, PUs, transform units (TUs), macroblocks, macroblock partitions, and so on. In some contexts, discussion of PUs may be interchanged with discussion of macroblocks or macroblock partitions.

[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 is an $N \times N$ block of samples. A CU may be 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. Video encoder 20 may partition a coding block of a CU into one or more prediction blocks. A prediction block may be 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 be a prediction block of luma samples, two corresponding prediction blocks of chroma samples of a picture, and syntax structures used to predict the prediction block

samples. Video encoder 20 may generate predictive luma, Cb and Cr blocks for luma, Cb and Cr prediction blocks of each PU of the CU.

[0051] 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.

[0052] 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. Video encoder 20 may use uni-prediction or bi-prediction to generate the predictive blocks of a PU. When video encoder 20 uses uni-prediction to generate the predictive blocks for a PU, the PU may have a single motion vector (MV). When video encoder 20 uses bi-prediction to generate the predictive blocks for a PU, the PU may have two MVs.

[0053] After video encoder 20 generates predictive blocks (e.g., predictive luma, Cb and Cr blocks) for one or more PUs of a CU, video encoder 20 may generate residual blocks for the CU. Each sample in a residual block of the CU may indicate a difference between a sample in a predictive block of a PU of the CU and a corresponding sample in a coding block of the CU. For example, 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 residual blocks (e.g., luma, Cb and Cr residual blocks) of a CU into one or more transform blocks (e.g., luma, Cb and Cr transform blocks). A transform block may be a rectangular block of samples on which the same transform is applied. A transform unit (TU) of a CU may be 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.

[0055] Video encoder 20 may apply one or more transforms to a transform block to generate a coefficient block for a TU. A coefficient block may be a two-dimensional array of transform coefficients. A transform coefficient may be a scalar quantity. For example, 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. 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 encoding 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. Video encoder 20 may output the entropy-encoded syntax elements in a bitstream. The bitstream may also include syntax elements that are not entropy encoded.

[0057] Video encoder 20 may output a bitstream that includes the entropy-encoded syntax elements. The bitstream may include a sequence of bits that forms a representation of coded pictures and associated data. The bitstream may comprise a sequence of network abstraction layer (NAL) units. Each of the NAL units includes a NAL unit header and encapsulates a raw byte sequence payload (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 picture parameter set (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 supplemental enhancement information (SEI), 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 video coding layer (VCL) NAL units.

[0059] Video decoder 30 may receive a bitstream generated by video encoder 20. In addition, video decoder 30 may obtain syntax elements from the bitstream. For example, video decoder 30 may parse the bitstream to decode 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 (e.g., decoded) from the bitstream. The process to reconstruct the video data may be generally reciprocal to the process performed by video encoder 20. For instance, video decoder 30 may use MVs of PUs to determine predictive sample blocks (i.e., predictive blocks) for the PUs of a current CU. In addition, video decoder 30 may inverse quantize transform coefficient blocks associated with TUs of the current CU. Video decoder 30 may perform inverse transforms on the transform 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 sample 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] In some examples, video encoder 20 and video decoder 30 may be configured to perform palette-based coding. For example, in palette based coding, rather than performing the intra-predictive or inter-predictive coding techniques described above, video encoder 20 and video decoder 30 may code a so-called palette as a table of colors or pixel values representing the video data of a particular area (e.g., a given block). In this way, rather than coding actual pixel values or their residuals for a current block of video data, the video coder may code index values for one or more of the pixels values of the current block, where the index values indicate entries in the palette that are used to represent the pixel values of the current block.

[0061] In one example, video encoder 20 may encode a block of video data by determining a palette for the block, locating an entry in the palette having a value

representative of the value of one or more pixel of the block, and encoding the block with index values that indicate the entry in the palette used to represent the one or more pixel values of the block. In some examples, video encoder 20 may signal the index values in an encoded bitstream. A video decoder may obtain, from an encoded bitstream, a palette for a block, as well as index values for the pixels of the block. The video decoder may relate the index values of the pixels to entries of the palette to reconstruct the pixel values of the block.

[0062] In another example, video encoder 20 may encode a block of video data by determining prediction residual values for the block, determining a palette for the block, locating an entry in the palette having a value representative of the value of one or more of the prediction residual values, and encoding the block with index values that indicate the entry in the palette used to represent the prediction residual values for the block. Video decoder 30 may obtain, from an encoded bitstream, a palette for a block, as well as index values for the prediction residual values of the block. Video decoder 30 may relate the index values of the prediction residual values to entries of the palette to reconstruct the prediction residual values of the block. The prediction residual values may be added to the prediction values (for example, obtained using intra or inter prediction) to reconstruct the pixel values of the block.

[0063] As described in more detail below, the basic idea of palette-based coding is that, for a given block of video data to be coded, a palette is derived that includes the most dominant pixel values in the current block. For instance, the palette may refer to a number of pixel values which are assumed to be dominant and/or representative for the current CU. Video encoder 20 may first transmit the size and the elements of the palette to video decoder 30. Video encoder 20 may encode the pixel values in the given block according to a certain scanning order. For each pixel location in the given block, video encoder 20 may transmit a flag or other syntax element to indicate whether the pixel value at the pixel location is included in the palette or not. If the pixel value is in the palette (i.e., a palette entry exists that specifies the pixel value), video encoder 20 may signal the index value associated with the pixel value for the pixel location in the given block followed by a “run” of like-valued consecutive pixel values in the given block. In this case, video encoder 20 does not transmit the flag or the palette index for the following pixel locations that are covered by the “run” as they all have the same pixel value.

[0064] If the pixel value is not in the palette (i.e., no palette entry exists that specifies the pixel value), video encoder 20 may transmit the pixel value or a residual value (or quantized versions thereof) for the given pixel location in the given block. Video decoder 30 may first determine the palette based on the information received from video encoder 20. Video decoder 30 may then map the received index values associated with the pixel locations in the given block to entries of the palette to reconstruct the pixel values of the given block.

[0065] Palette-based coding may have a certain amount of signaling overhead. For example, a number of bits may be needed to signal characteristics of a palette, such as a size of the palette, as well as the palette itself. In addition, a number of bits may be needed to signal index values for the pixels of the block. The techniques of this disclosure may, in some examples, reduce the number of bits needed to signal such information. For example, the techniques described in this disclosure may include techniques for various combinations of one or more of signaling palette-based coding modes, transmitting palettes, predicting palettes, deriving palettes, or transmitting palette-based coding maps and other syntax elements. Particular techniques of this disclosure may be implemented in video encoder 20 and/or video decoder 30.

[0066] Aspects of this disclosure are directed to palette prediction. For example, according to aspects of this disclosure, video encoder 20 and/or video decoder 30 may determine a first palette having a first set of entries indicative of first pixel values. Video encoder 20 and/or video decoder 30 may then determine, based on the first set of entries of the first palette, a second set of entries indicative of second pixel values of a second palette. Video encoder 20 and/or video decoder 30 may also code pixels of a block of video data using the second palette (i.e., using the second set of pixel values).

[0067] When determining the second set of entries of the second palette based on the first set of entries, video encoder 20 may encode a variety of syntax elements, which may be used by video decoder 30 to reconstruct the second palette. For example, video encoder 20 may encode one or more syntax elements in a bitstream to indicate that an entire palette (or palettes, in the case of each color component, e.g., Y, Cb, Cr, or Y, U, V, or R, G, B, of the video data having a separate palette) is predicted from (e.g., copied from) one or more neighboring blocks of the block currently being coded.

[0068] The palette from which entries of the current palette of the current block are predicted (e.g., copied) may be referred to as a predictive palette. The predictive palette may contain palette entries from one or more neighboring blocks including spatially

neighboring blocks and/or neighboring blocks in a particular scan order of the blocks. For example, the neighboring blocks may be spatially located to the left (left neighboring block) or above (upper neighboring block) the block currently being coded. In another example, video encoder 20 may determine predictive palette entries using the most frequent sample values in a causal neighborhood of the current block. In another example, the neighboring blocks may neighbor the block currently being coded according to a particular scan order used to code the blocks. That is, the neighboring blocks may be one or more blocks coded prior to the current block in the scan order. Video encoder 20 may encode one or more syntax elements to indicate the location of the neighboring blocks from which the palettes are copied.

[0069] In some examples, palette prediction may be performed entry-wise. For example, video encoder 20 may encode one or more syntax elements to indicate, for each entry of a predictive palette, whether the given palette entry is included in the current palette for the current block. If video encoder 20 does not use prediction to populate an entry of the current palette for the current block, video encoder 20 may encode one or more additional syntax elements to specify the non-predicted entries, as well as the number of such entries, in the current palette for the current block.

[0070] As described above, for a current block, e.g., a CU or PU, the entries in its palette may be predicted from entries in a predictive palette including palette entries from one or more previously coded neighboring blocks. This disclosure describes several alternative techniques to predict the palette for the current block.

[0071] In one example, a predictive palette includes a number of entries, N . In this example, video encoder 20 first transmits a binary vector, V , having the same size as the predictive palette, i.e., a vector of 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 a current palette for a current block. For example, video encoder 20 may encode one or more syntax elements including the binary vector. In some cases, video encoder 20 encodes the binary vector including a one-bit flag for each of the palette entries in the predictive palette that indicates whether a respective palette entry is copied to the current palette. In other cases, video encoder 20 encodes a losslessly compressed binary vector in which the indications for the entries in the binary vector are compressed or combined together instead of being sent individually as one-bit flags. In this way, video decoder 30 determines the one or more of the palette entries in the predictive palette that are copied to the current palette.

[0072] In addition, video encoder 20 transmits a number, M , that indicates how many new entries will be included in the palette for the current block, and then transmits pixel values for the new entries to video decoder 30. For example, video encoder 20 may encode one or more syntax elements indicating the number of the new palette entries that are included in the current palette using one of unary codes, truncated unary codes, Exponential-Golomb codes, or Golomb-Rice codes. In this way, video decoder 30 determines the number of new palette entries not in the predictive palette that are included in the current palette for the current block;

[0073] In this example, the final size of the current 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 are reused in the palette for current block. Video decoder 30 may calculate a size of the current palette to be equal to the sum of a number of the copied palette entries and the number of the new palette entries. Once the size of the current palette is determined, video decoder 30 generates the current palette including the copied palette entries from the predictive palette and the new palette entries explicitly signaled from video encoder 20.

[0074] To generate the palette for the current block, video decoder 30 may merge the received M new palette entries and the S copied palette entries that are being 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, for example, when separate palette is used for each component. In other cases, the merge may be a concatenation of the two sets of entries, i.e., the copied palette entries and the new palette entries.

[0075] 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., a vector of 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 the entries that are copied from the predictive palette, 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 indicating 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.

[0076] In the two above examples, the one or more previously coded neighboring blocks, from which the predictive palette used for the prediction of the current palette for the current block is formed, may be spatially neighboring blocks of the current block and/or neighboring blocks of the current block in a particular scan order of the blocks. For example, the neighboring blocks may be spatially located above (i.e., top-neighboring blocks) or to the left (i.e., left-neighboring blocks) the current block. In some examples, a candidate list of neighboring blocks may be constructed, and video encoder 20 transmits an index to indicate one or more of the candidate neighboring blocks and associated palettes are used to form the predictive palette.

[0077] For certain blocks, e.g., CUs at a beginning of a slice or at other slice boundaries or leftmost CUs of the slice or a picture of video data, palette prediction may be disabled. For example, when the current block of video data comprises one or a first block in a slice of video data or a leftmost block of the slice or a picture of the video data, video encoder 20 and/or video decoder 30 may disable copying of palette entries in the prediction palette to the current palette for the current block.

[0078] 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 of the palette for the current block is explicitly transmitted by video encoder 20 or whether the palette entry is derived from a previously reconstructed pixel. For each of the palette entries of the palette for the current block 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 pixel location of the reconstructed pixel in 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.

[0079] In some examples, techniques for predicting an entire palette may be combined with techniques for predicting one or more entries of a palette. For example, video

encoder 20 may encode one or more syntax elements in a bitstream to indicate whether the current palette is entirely copied from the predictive palette (for example, the palette for the last palette-coded block). If this is not the case, video encoder 20 may encode one or more syntax elements in a bitstream to indicate whether each entry in the predictive palette is copied.

[0080] In some instances, the size of the palette may be a fixed value specified in the video coding standard applied by video encoder 20 and video decoder 30, or may be signaled from video encoder 20 to video decoder 30. In the case where each of the color components has a separate palette, video encoder 20 may separately signal the sizes for the different palettes. In the case of a single palette for all the color components, video encoder 20 may encode a single size for the single palette. In another example, instead of signaling the number of entries and the palette values, video encoder 20 may signal, after signaling each palette value, a flag to indicate whether the signaled palette value is the final palette entry for the palette. Video encoder 20 may not signal such an “end of palette” flag if the palette has already reached a certain maximum size.

[0081] Video encoder 20 may encode one or more syntax elements to indicate whether palette prediction is enabled and/or active. In an example for purposes of illustration, video encoder 20 may encode a `pred_palette_flag` to indicate, for each block (e.g., CU or PU), whether video encoder 20 uses palette prediction to predict the palette for the respective block. In some examples, video encoder may signal a separate flag for each color component (e.g., three flags for each block). In other examples, video encoder 20 may signal a single flag that is applicable to all color components of a block.

[0082] Video decoder 30 may obtain the above-identified information from an encoded bitstream and may use the information to reconstruct the palette. For example, video decoder 30 may receive data indicating whether a particular palette is predicted from another palette, as well as information that allows video decoder 30 to use the appropriate predictive palette entries.

[0083] In some instances, additionally or alternatively, video encoder 20 and/or video decoder 30 may construct a palette “on-the-fly,” i.e., dynamically. For example, video encoder 20 and/or video decoder 30 may add entries to an empty palette during coding. That is, video encoder 20 may add pixel values to a palette as the pixel values are generated and transmitted for positions in a block. Pixels (e.g., pixels having values that have previously been added and indexed within the palette) that are coded relatively later

in the block may refer to earlier added entries of the palette, e.g., with index values associated with pixel values, instead of transmitting the pixel values. Likewise, upon receiving a new pixel value for a position in a block, video decoder 30 may follow the same process as video encoder 20 and include the pixel value in a palette. In this way, video decoder 30 constructs the same palette as video encoder 20. Video decoder 30 may receive, for pixels having values that are already included in the palette, index values that identify the pixel values. Video decoder 30 may use the received information, e.g., pixel values for the palette and/or index values, to reconstruct the pixels of a block.

[0084] In some instances, video encoder 20 and video decoder 30 may maintain a palette of a fixed size. For example, video encoder 20 and video decoder 30 may add the most recent reconstructed pixel values to the palette. For each entry that is added to the palette, the entry that was added to the palette the earliest is discarded. This is also sometimes referred to as First-in-First-out (FIFO). This process of updating the palette may be applied only to blocks that are coded using the palette mode or to all the blocks irrespective of the coding mode.

[0085] The techniques described above generally relate to video encoder 20 and video decoder 30 constructing and/or transmitting a palette for palette-based coding. Other aspects of this disclosure relate to constructing and/or transmitting a map that allows video encoder 20 and/or video decoder 30 to determine pixel values. For example, other aspects of this disclosure relate constructing and/or transmitting a map of indices that indicate entries in a palette that specify pixel values of a block of video data.

[0086] In some examples, video encoder 20 may indicate whether pixels of a block have a corresponding value in a palette. In an example for purposes of illustration, assume that an (i, j) entry of a map corresponds to an (i, j) pixel position in a block of video data. In this example, video encoder 20 may encode a flag for each pixel position of a block. Video encoder 20 may set the flag equal to one for the (i, j) entry to indicate that the pixel value at the (i, j) location is one of the values in the palette. When a pixel value is included in the palette (i.e., the flag is equal to one), video encoder 20 may also encode data indicating a palette index for the (i, j) entry that identifies the corresponding entry in the palette that specifies the pixel value. When a pixel value is not included in the palette (i.e., the flag is equal to zero), video encoder 20 may also encode data indicating a sample value (possibly quantized) for the pixel. In some cases, the pixel that is not included in the palette is referred to as an “escape pixel.”

[0087] Video decoder 30 may obtain the above-described data from an encoded bitstream and use the data to determine a palette index and/or pixel value for a particular location in a block. For example, video decoder 30 may decode one or more syntax elements indicating whether each of the pixel values of the current block has a corresponding pixel value in the current palette, decode one or more syntax elements indicating the index values for the one or more pixel values of the current block that have corresponding pixel values in the current palette, and decode one or more syntax element indicating the pixel values for the one or more pixel values of the current block that do not have a corresponding pixel value in the current palette.

[0088] In some instances, there may be a correlation between the palette index to which a pixel at a given position is mapped and the probability of a neighboring pixel being mapped to the same palette index. That is, when a pixel is mapped to a particular palette index, the probability may be relatively high that one or more neighboring pixels (in terms of spatial location) are mapped to the same palette index.

[0089] According to aspects of this disclosure, video encoder 20 and/or video decoder 30 may determine and code one or more indices of a block of video data relative to one or more indices of the same block of video data. For example, video encoder 20 and/or video decoder 30 may be configured to determine a first index value associated with a first pixel in a block of video data, where the first index value relates a value of the first pixel to an entry of a palette. Video encoder 20 and/or video decoder 30 may also be configured to determine, based on the first index value, one or more second index values associated with one or more second pixels in the block of video data, and to code the first and the one or more second pixels of the block of video data. Thus, in this example, indices of a map may be coded relative to one or more other indices of the map.

[0090] In some examples, video encoder 20 may encode one or more syntax elements indicating a number of consecutive pixels in a given scan order that are mapped to the same index value. The string of like-valued index values may be referred to herein as a “run.” In some examples, a pixel value may be associated with exactly one index value in a palette. Accordingly, in some instances, a run of values may also refer to a string of like-valued pixel values. In other examples, as described with respect to lossy coding below, more than one pixel value may map to the same index value in a palette. In such examples, a run of values refers to like-valued index values. In this scenario, on the

decoder side, runs of like-valued index values may correspond to runs of pixel values that correspond to the index values.

[0091] In an example for purposes of illustration, if two consecutive indices in a given scan order have different values, the run is equal to zero. If two consecutive indices in a given scan order have the same value but the third index in the scan order has a different value, the run is equal to one. Video decoder 30 may obtain the syntax elements indicating a run from an encoded bitstream and may use the data indicated by the syntax elements to determine the number of consecutive pixel locations that have the same index value.

[0092] 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 when 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 calculate a prediction error or residual value between the original pixel value and the prediction pixel value included in the palette. Video encoder 20 may quantize, encode and transmit the residual value for the pixel location to video decoder 30.

[0093] Video decoder 30 may then derive a pixel value at the pixel location based on the corresponding received palette index. The derived pixel value and the residual value (received from video encoder 20) are then used to predict the 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 residual quad-tree (RQT) to transform the residual value, quantize the transform coefficients, and entropy encode the quantized transform coefficients. In some cases, the residual values may be quantized directly without applying a transform. As an example, video decoder 30 may decode one or more syntax elements indicating the index values for the one or more pixel values of the current block, where the index values identify corresponding pixel values in the current palette as prediction pixel values, and decode one or more syntax element indicating residual values between the one or more pixel values of the current block and the identified prediction pixel values in the current palette. In some cases, the above examples may be referred to as lossy coding.

[0094] Additionally or alternatively, according to aspects of this disclosure, video encoder 20 and video decoder 30 may perform line copying for one or more entries of a map. The entries may also be referred to as “positions” due to the relationship between entries of the map and pixel positions of a block. The line copying may depend, in some examples, on the scan direction. For example, video encoder 20 may indicate that a pixel value or index map value for a particular position in a block is equal to the pixel or index value in a line above (e.g., preceding) the particular position (for a horizontally oriented scan) or the column to the left of (e.g., preceding) the particular position (for a vertically oriented scan). Video encoder 20 may also indicate, as a run, the number of pixel values or indices in the scan order that are equal to the pixel values or indices in the line above or the column to the left of the particular position. In this example, video encoder 20 and or video decoder 30 may copy pixel or index values from the specified neighboring line (or column for vertical scan) and for the specified number of entries for the line (or column for vertical scan) of the block currently being coded.

[0095] In some instances, the line (or column for vertical scan) from which values are copied may be directly adjacent to, e.g., above or to the left of, the line (or column for vertical scan) of the position currently being coded. In other examples, a number of lines of the block 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 values for a line of the map currently being coded. Similar techniques may be applied to previous columns for a vertical scan. In an example for purposes of illustration, video encoder 20 and/or video decoder 30 may be configured to store the previous four rows of indices or pixel values prior to coding the current row of pixels. In this example, the predictive row (the row from which indices or pixel values are copied) may be indicated in a bitstream with a truncated unary code or other codes such as unary codes. With respect to a truncated unary code, video encoder 20 and/or video decoder 30 may determine a maximum value for the truncated unary code based on a maximum row calculation (e.g., $\text{row_index}-1$) for horizontal scans or a maximum column calculation (e.g., $\text{column_index}-1$) for vertical scans. In addition, an indication of the number of positions from the predictive row that are copied may also be included in the bitstream. In some instances, if the line (or column in the case of vertical scans) from which a current position is being predicted belongs to another block (e.g., CU or CTU) such prediction may be disabled.

[0096] 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. Thus, any scan may be used to convert the two dimensional array of pixel or index values into a one dimensional array.

[0097] 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 the map. 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 use the information to reconstruct the map and pixel values for a block.

[0098] 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 a constant value or 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.

[0099] In some examples, video encoder 20 may not transmit runs of like-valued index values in a given scan order to video decoder 30. Instead, video encoder 20 and/or video decoder 30 may implicitly derive the values of the runs in order to determine the entries of the map. In this case, video encoder 20 may signal to video decoder 30 that a run of a given index value occurs, but may not signal a value of the run. For example, the value of a run may be a constant value or may be derived based on side information for the current block of video data being coded such as, for example, the block size. 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 some examples, video encoder 20 may signal the value of a run to video decoder 30 using high level syntax. In some examples, the phrase “high-level syntax” refers to syntax in parameter sets, e.g., picture parameter sets (PPSs), sequence parameter sets (SPSs), and video parameter sets (VPSs), and slice headers.

[0100] Additionally or alternatively, video encoder 20 may not even need to transmit the map 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 the map. 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 the current block. 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. In another example, the start location may be derived depending on side 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 of the current block, or a fraction of each row and/or column of the current block. In some examples, video encoder 20 may signal the start position to video decoder 30 using high level syntax.

[0101] 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 in 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 the 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.

[0102] In some cases, described in more detail below, one palette is generated and shared for multiple color components in the current block. For example, for each pixel location in the current block, the pixel values in three color components (e.g., Y luma and both U and V chroma components) may form a vector (i.e., a color vector). Then, a palette may be formed by selecting a certain number of vectors to represent the current block. It may be possible to have one palette of pixel values for the luma component,

and another palette of pixel values for the chroma components. The line copying described in more detail above may also work with a single palette. With a shared palette, a palette entry may be a triplet of (Y, U, V) or (Y, Cb, Cr) or (R, G, B). In this case, the palette index for each pixel location is signaled as being equal to the palette index of the row above, if the scan is horizontal, or the column on the left, if the scan is vertical, and then the associated number of palette indices is also copied from the previous row or column based on the run.

[0103] In the case of either a shared palette for two or more color components or of separate palettes for each of the color components, geometric information may be shared between the color components. Usually there is high correlation between edge locations of collocated blocks in different color components because the chroma components may have been downsampled from the luma components in a pre-defined way, such as 4:2:2 or 4:2:0 sampling.

[0104] For example, in palette-based coding, run coding may be used to indicate geometry information for the current block because an edge of the current block will break the run. In case of the 4:4:4 chroma format, the run may be generated once and used for all color components. The run may be generated based on one of the color components, or the run may be generated using more than one of the color components. In case of the 4:2:2 chroma format or the 4:2:0 chroma format, the run used for the luma component may be downsampled for application to the chroma components

[0105] The techniques of this disclosure also include other aspects of palette-based coding. For example, according to aspects of this disclosure, video encoder 20 and/or video decoder 30 may code one or more syntax elements for each block to indicate that the block is coded using a palette coding mode. For example, video encoder 20 and/or video decoder 30 may code a palette mode flag (PLT_Mode_flag) to indicate whether a palette-based coding mode is to be used for coding a particular block. In this example, video encoder 20 may encode a PLT_Mode_flag that is equal to one to specify that the block currently being encoded ("current block") is encoded using a palette mode. A value of the PLT_Mode_flag equal to zero specifies that the current block is not encoded using palette mode. In this case, video decoder 30 may obtain the PLT_Mode_flag from the encoded bitstream and apply the palette-based coding mode to decode the block. In instances in which there is more than one palette-based coding mode available (e.g., there is more than one palette-based technique available for

coding) one or more syntax elements may indicate one of a plurality of different palette modes for the block.

[0106] In some instances, video encoder 20 may encode a `PLT_Mode_flag` that is equal to zero to specify that the current block is not encoded using a palette mode. In such instances, video encoder 20 may encode the block using any of a variety of inter-predictive, intra-predictive, or other coding modes. When the `PLT_Mode_flag` is equal to zero, video encoder 20 may transmit additional information (e.g., syntax elements) to indicate the specific mode that is used for encoding the respective block. In some examples, as described below, the mode may be an HEVC coding mode, e.g., a regular inter-predictive mode or intra-predictive mode in the HEVC standard. The use of the `PLT_Mode_flag` is described for purposes of example. In other examples, other syntax elements such as multi-bit codes may be used to indicate whether the palette-based coding mode is to be used for one or more blocks, or to indicate which of a plurality of modes are to be used.

[0107] When a palette-based coding mode is used, a palette is transmitted by video encoder 20, e.g., using one or more of the techniques described herein, in the encoded video data bitstream for use by video decoder 30. A palette may be transmitted for each block or may be shared among a number of blocks. The palette may refer to a number of pixel values that are dominant and/or representative for the block.

[0108] The size of the palette, 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. As described in greater detail below, a pixel value may be composed of a number of samples, e.g., depending on the color space used for coding. For example, a pixel value may include luma and chrominance samples (e.g., luma, U chrominance and V chrominance (YUV) or luma, Cb chrominance, and Cr chrominance (YCbCr) samples). In another example, a pixel value may include Red, Green, and Blue (RGB) samples. As described herein, the term pixel value may generally refer to one or more of the samples contributing to a pixel. That is, the term pixel value does not necessarily refer to all samples contributing to a pixel, and may be used to describe a single sample value contributing to a pixel.

[0109] In some examples, a palette may be transmitted separately for each color component of a particular block. For example, in the YUV color space, there may be a palette for the Y component (representing Y values), another palette for the U component (representing U values), and yet another palette for the V component

(representing V values). In another example, a palette may include all components of a particular block. In this example, the i -th entry in the palette may include three values (e.g., Y_i , U_i , V_i). According to aspects of this disclosure, one or more syntax elements may separately indicate the size of the palette for each component (e.g., Y, U, V, or the like). In other examples, a single size may be used for all components, such that one or more syntax elements indicate the size of all components.

[0110] Video encoder 20 and/or video decoder 30 may perform palette-based coding in a lossy or lossless manner. That is, in some examples, video encoder 20 and/or video decoder 30 may losslessly code video data for a block using palette entries that match the pixel values of the block (or by sending the actual pixel values if the pixel value is not included in the palette). In other examples, as described in greater detail with respect to FIG. 5 below, video encoder 20 and/or video decoder 30 may code video data for a block using palette entries that do not exactly match the pixel values of the block (lossy coding). Similarly, if the actual pixel value is not included in the palette, the actual pixel value may be quantized in a lossy manner.

[0111] According to techniques described in this disclosure, video encoder 20 and/or video decoder 30 may perform palette-based coding of predicted video blocks. In one example, video encoder 20 first derives a palette for a current block based on the pixel values in the current block, and then maps the pixel values in the current block to palette indices for encoding. The mapping may be one to one (i.e., for lossless coding) or multiple to one (i.e., for lossy coding). Video encoder 20 also maps reference pixel values in a previously coded block that will be used to predict the pixel values in the current block. Once the pixel values of the current block have been mapped to palette indices, video encoder 20 may encode the current block with palette indices using regular encoding methods, e.g., regular intra coding in the HEVC standard.

[0112] In the above example, the current block with palette indices is treated as if the current block were an original block with pixel values. Similarly, the palette indices of the reference pixels are used for performing regular intra prediction on the current block with palette indices. Video encoder 20 transmits the prediction error or residual values to video decoder 30. After encoding the current block, video encoder 20 converts the indices of the reference pixels, prediction pixels, and the residual values back to the pixel values for reconstruction of the current block and the normal prediction of future blocks. Video decoder 30 may obtain the encoded residual values for the current block from the bitstream, and decode the current block using regular decoding method to

obtain the current block with palette indices. Video decoder 30 may then determine the pixel values of the current block based on the pixel values in the palette that are associated with the palette indices.

[0113] In another example, video encoder 20 may generate a palette for a current block where the palette includes entries that indicate prediction residual values for the given block. The prediction residual values for the given block may be generated using any prediction mode, e.g., regular inter-prediction or intra-prediction in the HEVC standard. The prediction residual values for the given block may be residual pixel values or residual transform coefficient values. In either case, the prediction residual values may be quantized. In this example, video encoder 20 maps the prediction residual values for the current block to the index values that indicate entries in the palette for the current block used to represent the prediction residual values for the current block, and encodes prediction residual values using the index values. Video decoder 30 may obtain the block of index values from the bitstream, and determine the prediction residual values for the current block based on the corresponding prediction residual values in the palette that are identified by the index values. Video decoder 30 may then reconstruct the pixel values of the current block using regular decoding methods based on the prediction residual values and previously coded reference pixel values.

[0114] In some examples, video encoder 20 and/or video decoder 30 may perform the palette-based video coding with video block prediction by applying the intra prediction mode (i.e., the prediction only uses previously coded pixel information in the current picture). In other cases, video encoder 20 and/or video decoder 30 may apply the inter prediction mode (i.e., the prediction is from pixels in a previously coded picture). In some cases, video encoder 20 and/or video decoder 30 may determine the prediction residual values for the current block using only a subset of prediction mode processes for either the inter prediction mode or the intra prediction mode.

[0115] In another example, video encoder 20 and/or video decoder 30 may perform no prediction for the current block. In this case, video encoder 20 instead maps the pixel values to palette indices, and encodes the indices using entropy coding without prediction. In an additional example, video encoder 20 and/or video decoder 30 may perform residual differential pulse code modulation (RDPCM) using pixel values of the current block that are mapped to palette index values. In this example, no prediction from pixels outside the current block is used, and horizontal or vertical prediction may be used for line copying index values within the current CU.

[0116] In some examples, the techniques for palette-based coding of video data may be used with one or more other coding techniques, such as techniques for inter- or intra-predictive coding. For example, as described in greater detail below, an encoder or decoder, or combined encoder-decoder (codec), may be configured to perform inter- and intra-predictive coding, as well as palette-based coding.

[0117] FIG. 2 is a block diagram illustrating an example video encoder 20 that may implement the techniques of this disclosure. FIG. 2 is provided for purposes of explanation and should not be considered limiting of the techniques as broadly exemplified and described in this disclosure. For purposes of explanation, this disclosure describes video encoder 20 in the context of HEVC coding. However, the techniques of this disclosure may be applicable to other coding standards or methods.

[0118] Video encoder 20 represents an example of a device that may be configured to perform techniques for palette-based video coding in accordance with various examples described in this disclosure. For example, video encoder 20 may be configured to selectively code various blocks of video data, such as CUs or PUs in HEVC coding, using either palette-based coding or non-palette based coding. Non-palette based coding modes may refer to various inter-predictive temporal coding modes or intra-predictive spatial coding modes, such as the various coding modes specified by HEVC Draft 10. Video encoder 20, in one example, may be configured to generate a palette having entries indicating pixel values. Furthermore, in this example, video encoder 20 may select pixel values in a palette to represent pixel values of at least some positions of a block of video data. In this example, video encoder 20 may 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. Video decoder 30 may use the signaled information to decode video data.

[0119] In the example of FIG. 2, video encoder 20 includes a video data memory 98, a prediction processing unit 100, a residual generation unit 102, a transform processing unit 104, a quantization unit 106, an inverse quantization unit 108, an inverse transform processing unit 110, a reconstruction unit 112, a filter unit 114, a decoded picture buffer 116, and an entropy encoding unit 118. Prediction processing unit 100 includes an inter-prediction processing unit 120 and an intra-prediction processing unit 126. Inter-prediction processing unit 120 includes a motion estimation unit and a motion compensation unit (not shown). Video encoder 20 also includes a palette-based encoding unit 122 configured to perform various aspects of the palette-based coding

techniques described in this disclosure. In other examples, video encoder 20 may include more, fewer, or different functional components.

[0120] Video data memory 98 may store video data to be encoded by the components of video encoder 20. The video data stored in video data memory 98 may be obtained, for example, from video source 18. Decoded picture buffer 116 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 98 and decoded picture buffer 116 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 98 and decoded picture buffer 116 may be provided by the same memory device or separate memory devices. In various examples, video data memory 98 may be on-chip with other components of video encoder 20, or off-chip relative to those components.

[0121] Video encoder 20 may receive video data. Video encoder 20 may encode each CTU in a slice of a picture of the video data. Each of the CTUs may be associated with equally-sized luma coding tree blocks (CTBs) and corresponding CTBs of the picture. As part of encoding a CTU, prediction processing unit 100 may perform quad-tree partitioning to divide the CTBs of the CTU into progressively-smaller blocks. The smaller block may be coding blocks of CUs. For example, prediction processing unit 100 may partition a CTB associated with a CTU into four equally-sized sub-blocks, partition one or more of the sub-blocks into four equally-sized sub-sub-blocks, and so on.

[0122] Video encoder 20 may encode CUs of a CTU to generate encoded representations of the CUs (i.e., coded CUs). As part of encoding a CU, prediction processing unit 100 may partition the coding blocks associated with the CU among one or more PUs of the CU. Thus, each PU may be associated with a luma prediction block and corresponding chroma prediction blocks. Video encoder 20 and video decoder 30 may support PUs having various sizes. As indicated above, the size of a CU may refer to the size of the luma coding block of the CU and the size of a PU may refer to the size of a luma prediction block of the PU. Assuming that the size of a particular CU is $2N \times 2N$, video encoder 20 and video decoder 30 may support PU sizes of $2N \times 2N$ or $N \times N$ for intra prediction, and symmetric PU sizes of $2N \times 2N$, $2N \times N$, $N \times 2N$, $N \times N$, or similar for inter prediction. Video encoder 20 and video decoder 30 may also support

asymmetric partitioning for PU sizes of $2N \times nU$, $2N \times nD$, $nL \times 2N$, and $nR \times 2N$ for inter prediction.

[0123] Inter-prediction processing unit 120 may generate predictive data for a PU by performing inter prediction on each PU of a CU. The predictive data for the PU may include one or more predictive sample blocks of the PU and motion information for the PU. Inter-prediction unit 121 may perform different operations for a PU of a CU depending on whether the PU is in an I slice, a P slice, or a B slice. In an I slice, all PUs are intra predicted. Hence, if the PU is in an I slice, inter-prediction unit 121 does not perform inter prediction on the PU. Thus, for blocks encoded in I-mode, the predictive block is formed using spatial prediction from previously-encoded neighboring blocks within the same frame.

[0124] If a PU is in a P slice, the motion estimation unit of inter-prediction processing unit 120 may search the reference pictures in a list of reference pictures (e.g., "RefPicList0") for a reference region for the PU. The reference region for the PU may be a region, within a reference picture, that contains sample blocks that most closely correspond to the sample blocks of the PU. The motion estimation unit may generate a reference index that indicates a position in RefPicList0 of the reference picture containing the reference region for the PU. In addition, the motion estimation unit may generate an MV that indicates a spatial displacement between a coding block of the PU and a reference location associated with the reference region. For instance, the MV may be a two-dimensional vector that provides an offset from the coordinates in the current decoded picture to coordinates in a reference picture. The motion estimation unit may output the reference index and the MV as the motion information of the PU. The motion compensation unit of inter-prediction processing unit 120 may generate the predictive sample blocks of the PU based on actual or interpolated samples at the reference location indicated by the motion vector of the PU.

[0125] If a PU is in a B slice, the motion estimation unit may perform uni-prediction or bi-prediction for the PU. To perform uni-prediction for the PU, the motion estimation unit may search the reference pictures of RefPicList0 or a second reference picture list ("RefPicList1") for a reference region for the PU. The motion estimation unit may output, as the motion information of the PU, a reference index that indicates a position in RefPicList0 or RefPicList1 of the reference picture that contains the reference region, an MV that indicates a spatial displacement between a sample block of the PU and a reference location associated with the reference region, and one or more prediction

direction indicators that indicate whether the reference picture is in RefPicList0 or RefPicList1. The motion compensation unit of inter-prediction processing unit 120 may generate the predictive sample blocks of the PU based at least in part on actual or interpolated samples at the reference region indicated by the motion vector of the PU.

[0126] To perform bi-directional inter prediction for a PU, the motion estimation unit may search the reference pictures in RefPicList0 for a reference region for the PU and may also search the reference pictures in RefPicList1 for another reference region for the PU. The motion estimation unit may generate reference picture indexes that indicate positions in RefPicList0 and RefPicList1 of the reference pictures that contain the reference regions. In addition, the motion estimation unit may generate MVs that indicate spatial displacements between the reference location associated with the reference regions and a sample block of the PU. The motion information of the PU may include the reference indexes and the MVs of the PU. The motion compensation unit may generate the predictive sample blocks of the PU based at least in part on actual or interpolated samples at the reference region indicated by the motion vector of the PU.

[0127] In accordance with various examples of this disclosure, video encoder 20 may be configured to perform palette-based coding. With respect to the HEVC framework, as an example, the palette-based coding techniques may be configured to be used as a CU mode. In other examples, the palette-based coding techniques may be configured to be used as a PU mode in the framework of HEVC. Accordingly, all of the disclosed processes described herein (throughout this disclosure) in the context of a CU mode may, additionally or alternatively, apply to a PU mode. However, these HEVC-based examples should not be considered a restriction or limitation of the palette-based coding techniques described herein, as such techniques may be applied to work independently or as part of other existing or yet to be developed systems/standards. In these cases, the unit for palette coding can be square blocks, rectangular blocks or even regions of non-rectangular shape.

[0128] Palette-based encoding unit 122, for example, may perform palette-based decoding when a palette-based encoding mode is selected, e.g., for a CU or PU. For example, palette-based encoding unit 122 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 122, some or all of such functions may be performed by other processing units, or a combination of different processing units.

[0129] Palette-based encoding unit 122 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 as 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.

[0130] Intra-prediction processing unit 126 may generate predictive data for a PU by performing intra prediction on the PU. The predictive data for the PU may include predictive sample blocks for the PU and various syntax elements. Intra-prediction processing unit 126 may perform intra prediction on PUs in I slices, P slices, and B slices.

[0131] To perform intra prediction on a PU, intra-prediction processing unit 126 may use multiple intra prediction modes to generate multiple sets of predictive data for the PU. When using some intra prediction modes to generate a set of predictive data for the PU, intra-prediction processing unit 126 may extend values of samples from sample blocks of neighboring PUs across the predictive blocks of the PU in directions associated with the intra prediction modes. The neighboring PUs may be above, above and to the right, above and to the left, or to the left of the PU, assuming a left-to-right, top-to-bottom encoding order for PUs, CUs, and CTUs. Intra-prediction processing unit 126 may use various numbers of intra prediction modes, e.g., 33 directional intra prediction modes. In some examples, the number of intra prediction modes may depend on the size of the region associated with the PU.

[0132] Prediction processing unit 100 may select the predictive data for PUs of a CU from among the predictive data generated by inter-prediction processing unit 120 for the PUs or the predictive data generated by intra-prediction processing unit 126 for the PUs. In some examples, prediction processing unit 100 selects the predictive data for the PUs of the CU based on rate/distortion metrics of the sets of predictive data. The predictive

sample blocks of the selected predictive data may be referred to herein as the selected predictive sample blocks.

[0133] Residual generation unit 102 may generate, based on the coding blocks (e.g., luma, Cb and Cr coding blocks) of a CU and the selected predictive sample blocks (e.g., predictive luma, Cb and Cr blocks) of the PUs of the CU, residual blocks (e.g., luma, Cb and Cr residual blocks) of the CU. For instance, residual generation unit 102 may generate the residual blocks of the CU such that each sample in the residual blocks has a value equal to a difference between a sample in a coding block of the CU and a corresponding sample in a corresponding selected predictive sample block of a PU of the CU.

[0134] Transform processing unit 104 may perform quad-tree partitioning to partition the residual blocks associated with a CU into transform blocks associated with TUs of the CU. Thus, in some examples, a TU may be associated with a luma transform block and two chroma transform blocks. The sizes and positions of the luma and chroma transform blocks of TUs of a CU may or may not be based on the sizes and positions of prediction blocks of the PUs of the CU. A quad-tree structure known as a “residual quad-tree” (RQT) may include nodes associated with each of the regions. The TUs of a CU may correspond to leaf nodes of the RQT.

[0135] Transform processing unit 104 may generate transform coefficient blocks for each TU of a CU by applying one or more transforms to the transform blocks of the TU. Transform processing unit 104 may apply various transforms to a transform block associated with a TU. For example, transform processing unit 104 may apply a discrete cosine transform (DCT), a directional transform, or a conceptually similar transform to a transform block. In some examples, transform processing unit 104 does not apply transforms to a transform block. In such examples, the transform block may be treated as a transform coefficient block.

[0136] Quantization unit 106 may quantize the transform coefficients in a coefficient block. The quantization process may reduce the bit depth associated with some or all of the transform coefficients. For example, an n -bit transform coefficient may be rounded down to an m -bit transform coefficient during quantization, where n is greater than m . Quantization unit 106 may quantize a coefficient block associated with a TU of a CU based on a quantization parameter (QP) value associated with the CU. Video encoder 20 may adjust the degree of quantization applied to the coefficient blocks associated with a CU by adjusting the QP value associated with the CU. Quantization may

introduce loss of information, thus quantized transform coefficients may have lower precision than the original ones.

[0137] Inverse quantization unit 108 and inverse transform processing unit 110 may apply inverse quantization and inverse transforms to a coefficient block, respectively, to reconstruct a residual block from the coefficient block. Reconstruction unit 112 may add the reconstructed residual block to corresponding samples from one or more predictive sample blocks generated by prediction processing unit 100 to produce a reconstructed transform block associated with a TU. By reconstructing transform blocks for each TU of a CU in this way, video encoder 20 may reconstruct the coding blocks of the CU.

[0138] Filter unit 114 may perform one or more deblocking operations to reduce blocking artifacts in the coding blocks associated with a CU. Decoded picture buffer 116 may store the reconstructed coding blocks after filter unit 114 performs the one or more deblocking operations on the reconstructed coding blocks. Inter-prediction processing unit 120 may use a reference picture that contains the reconstructed coding blocks to perform inter prediction on PUs of other pictures. In addition, intra-prediction processing unit 126 may use reconstructed coding blocks in decoded picture buffer 116 to perform intra prediction on other PUs in the same picture as the CU.

[0139] Entropy encoding unit 118 may receive data from other functional components of video encoder 20. For example, entropy encoding unit 118 may receive coefficient blocks from quantization unit 106 and may receive syntax elements from prediction processing unit 100. Entropy encoding unit 118 may perform one or more entropy encoding operations on the data to generate entropy-encoded data. For example, entropy encoding unit 118 may perform a CABAC operation, a context-adaptive variable length coding (CAVLC) operation, a variable-to-variable (V2V) length coding operation, a syntax-based context-adaptive binary arithmetic coding (SBAC) operation, a Probability Interval Partitioning Entropy (PIPE) coding operation, an Exponential-Golomb encoding operation, or another type of entropy encoding operation on the data. Video encoder 20 may output a bitstream that includes entropy-encoded data generated by entropy encoding unit 118. For instance, the bitstream may include data that represents a RQT for a CU.

[0140] In some examples, residual coding is not performed with palette coding. Accordingly, video encoder 20 may not perform transformation or quantization when

coding using a palette coding mode. In addition, video encoder 20 may entropy encode data generated using a palette coding mode separately from residual data.

[0141] According to one or more of the techniques of this disclosure, video encoder 20, and specifically palette-based encoding unit 122, may perform palette-based video coding of predicted video blocks. As described above, a palette generated by video encoder 20 may be explicitly encoded and sent to video decoder 30, predicted from previous palette entries, predicted from previous pixel values, or a combination thereof.

[0142] In one example, palette-based encoding unit 122 of video encoder 20 determines one or more palette entries in a predictive palette that are copied to a current palette for a current block of video data, and determines a number of new palette entries that are not in the predictor palette but that are included in the current palette. Based on this information, the palette-based video encoder 20 calculates a size of the current palette to be equal to the sum of the number of the copied palette entries and the number of the new palette entries, and generates the current palette of the determined size including the copied palette entries and the new palette entries. Video encoder 20 may transmit the determined information regarding the copied palette entries and the new palette entries to video decoder 30. In addition, video encoder 20 may explicitly encode and transmit pixel values for the new palette entries to video decoder 30. Palette-based encoding unit 122 of video encoder 20 may then encode the current block by determining index values for one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.

[0143] The techniques described in this disclosure may also include techniques for various combinations of one or more of signaling palette-based coding modes, transmitting palettes, predicting palettes, deriving palettes, or transmitting palette-based coding maps and other syntax elements.

[0144] FIG. 3 is a block diagram illustrating an example video decoder 30 that is configured to implement the techniques of this disclosure. FIG. 3 is provided for purposes of explanation and is not limiting on the techniques as broadly exemplified and described in this disclosure. For purposes of explanation, this disclosure describes video decoder 30 in the context of HEVC coding. However, the techniques of this disclosure may be applicable to other coding standards or methods.

[0145] Video decoder 30 represents an example of a device that may be configured to perform techniques for palette-based video coding in accordance with various examples

described in this disclosure. For example, video decoder 30 may be configured to selectively decode various blocks of video data, such as CUs or PUs in HEVC coding, using either palette-based coding or non-palette based coding. Non-palette based coding modes may refer to various inter-predictive temporal coding modes or intra-predictive spatial coding modes, such as the various coding modes specified by HEVC Draft 10. In one example, video decoder 30 may be configured to generate a palette having entries indicating pixel values. Furthermore, in this example, video decoder 30 may receive information associating at least some positions of a block of video data with entries in the palette. In this example, video decoder 30 may select pixel values in the palette based on the information and reconstruct pixel values of the block based on the selected pixel values.

[0146] In the example of FIG. 3, video decoder 30 includes a video data memory 148, an entropy decoding unit 150, a prediction processing unit 152, an inverse quantization unit 154, an inverse transform processing unit 156, a reconstruction unit 158, a filter unit 160, and a decoded picture buffer 162. Prediction processing unit 152 includes a motion compensation unit 164 and an intra-prediction processing unit 166. Video decoder 30 also includes a palette-based decoding unit 165 configured to perform various aspects of the palette-based coding techniques described in this disclosure. In other examples, video decoder 30 may include more, fewer, or different functional components.

[0147] Video data memory 148 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 148 may be obtained, for example, from computer-readable medium 16, e.g., from a local video source, such as a camera, via wired or wireless network communication of video data, or by accessing physical data storage media. Video data memory 148 may form a coded picture buffer (CPB) that stores encoded video data from an encoded video bitstream. Decoded picture buffer 162 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 148 and decoded picture buffer 162 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 148 and decoded picture buffer 162 may be provided by the same memory device or separate memory devices. In various examples,

video data memory 148 may be on-chip with other components of video decoder 30, or off-chip relative to those components.

[0148] Video data memory 148, i.e., a CPB, may receive and store encoded video data (e.g., NAL units) of a bitstream. Entropy decoding unit 150 may receive encoded video data (e.g., NAL units) from video data memory 148 and may parse the NAL units to decode syntax elements. Entropy decoding unit 150 may entropy decode entropy-encoded syntax elements in the NAL units. Prediction processing unit 152, inverse quantization unit 154, inverse transform processing unit 156, reconstruction unit 158, and filter unit 160 may generate decoded video data based on the syntax elements obtained (e.g., extracted) from the bitstream.

[0149] The NAL units of the bitstream may include coded slice NAL units. As part of decoding the bitstream, entropy decoding unit 150 may extract and entropy decode syntax elements from the coded slice NAL units. Each of the coded slices may include a slice header and slice data. The slice header may contain syntax elements pertaining to a slice. The syntax elements in the slice header may include a syntax element that identifies a PPS associated with a picture that contains the slice.

[0150] In addition to decoding syntax elements from the bitstream, video decoder 30 may perform a reconstruction operation on a non-partitioned CU. To perform the reconstruction operation on a non-partitioned CU, video decoder 30 may perform a reconstruction operation on each TU of the CU. By performing the reconstruction operation for each TU of the CU, video decoder 30 may reconstruct residual blocks of the CU.

[0151] As part of performing a reconstruction operation on a TU of a CU, inverse quantization unit 154 may inverse quantize, i.e., de-quantize, coefficient blocks associated with the TU. Inverse quantization unit 154 may use a QP value associated with the CU of the TU to determine a degree of quantization and, likewise, a degree of inverse quantization for inverse quantization unit 154 to apply. That is, the compression ratio, i.e., the ratio of the number of bits used to represent original sequence and the compressed one, may be controlled by adjusting the value of the QP used when quantizing transform coefficients. The compression ratio may also depend on the method of entropy coding employed.

[0152] After inverse quantization unit 154 inverse quantizes a coefficient block, inverse transform processing unit 156 may apply one or more inverse transforms to the coefficient block in order to generate a residual block associated with the TU. For

example, inverse transform processing unit 156 may apply an inverse DCT, an inverse integer transform, an inverse Karhunen-Loeve transform (KLT), an inverse rotational transform, an inverse directional transform, or another inverse transform to the coefficient block.

[0153] If a PU is encoded using intra prediction, intra-prediction processing unit 166 may perform intra prediction to generate predictive blocks for the PU. Intra-prediction processing unit 166 may use an intra prediction mode to generate the predictive luma, Cb and Cr blocks for the PU based on the prediction blocks of spatially-neighboring PUs. Intra-prediction processing unit 166 may determine the intra prediction mode for the PU based on one or more syntax elements decoded from the bitstream.

[0154] Prediction processing unit 152 may construct a first reference picture list (RefPicList0) and a second reference picture list (RefPicList1) based on syntax elements extracted from the bitstream. Furthermore, if a PU is encoded using inter prediction, entropy decoding unit 150 may extract motion information for the PU. Motion compensation unit 164 may determine, based on the motion information of the PU, one or more reference regions for the PU. Motion compensation unit 164 may generate, based on samples blocks at the one or more reference blocks for the PU, predictive blocks (e.g., predictive luma, Cb and Cr blocks) for the PU.

[0155] Reconstruction unit 158 may use the transform blocks (e.g., luma, Cb and Cr transform blocks) associated with TUs of a CU and the predictive blocks (e.g., luma, Cb and Cr blocks) of the PUs of the CU, i.e., either intra-prediction data or inter-prediction data, as applicable, to reconstruct the coding blocks (e.g., luma, Cb and Cr coding blocks) of the CU. For example, reconstruction unit 158 may add samples of the transform blocks (e.g., luma, Cb and Cr transform blocks) to corresponding samples of the predictive blocks (e.g., predictive luma, Cb and Cr blocks) to reconstruct the coding blocks (e.g., luma, Cb and Cr coding blocks) of the CU.

[0156] Filter unit 160 may perform a deblocking operation to reduce blocking artifacts associated with the coding blocks (e.g., luma, Cb and Cr coding blocks) of the CU. Video decoder 30 may store the coding blocks (e.g., luma, Cb and Cr coding blocks) of the CU in decoded picture buffer 162. Decoded picture buffer 162 may provide reference pictures for subsequent motion compensation, intra prediction, and presentation on a display device, such as display device 32 of FIG. 1. For instance, video decoder 30 may perform, based on the blocks (e.g., luma, Cb and Cr blocks) in decoded picture buffer 162, intra prediction or inter prediction operations on PUs of

other CUs. In this way, video decoder 30 may extract, from the bitstream, transform coefficient levels of a significant coefficient block, inverse quantize the transform coefficient levels, apply a transform to the transform coefficient levels to generate a transform block, generate, based at least in part on the transform block, a coding block, and output the coding block for display.

[0157] In accordance with various examples of this disclosure, video decoder 30 may be configured to perform palette-based coding. Palette-based decoding unit 165, for example, may perform palette-based decoding when a palette-based decoding mode is selected, e.g., for a CU or PU. For example, palette-based decoding unit 165 may be configured to generate a palette having entries indicating pixel values. Furthermore, in this example, palette-based decoding unit 165 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 165 may select pixel values in the palette based on the information. Additionally, in this example, palette-based decoding unit 165 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 165, some or all of such functions may be performed by other processing units, or a combination of different processing units.

[0158] Palette-based decoding unit 165 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 165 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, when the palette coding mode information indicates that the palette coding mode does not apply to the block. 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.

[0159] According to one or more of the techniques of this disclosure, video decoder 30, and specifically palette-based decoding unit 165, may perform palette-based video

decoding of predicted video blocks. As described above, a palette generated by video decoder 30 may be explicitly encoded by video encoder 20, predicted from previous palette entries, predicted from previous pixel values, or a combination thereof.

[0160] In one example, palette-based decoding unit 165 of video decoder 30 determines one or more palette entries in a predictive palette that are copied to a current palette for a current block of video data, and determines a number of new palette entries that are not in the predictor palette but that are included in the current palette. Video decoder 30 may receive the information regarding the copied palette entries and the new palette entries from video encoder 20. In addition, video decoder 30 may receive explicitly encoded pixel values for the new palette entries transmitted from video encoder 20. Based on this information, palette-based decoding unit 165 calculates a size of the current palette to be equal to the sum of the number of the copied palette entries and the number of the new palette entries, and generates the current palette of the determined size including the copied palette entries and the new palette entries. Palette-based decoding unit 165 of video decoder 30 may then decode the current block by determining index values for one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.

[0161] The techniques described in this disclosure may also include techniques for various combinations of one or more of signaling palette-based coding modes, transmitting palettes, predicting palettes, deriving palettes, or transmitting palette-based coding maps and other syntax elements.

[0162] As described above, in some examples, video encoder 20 and/or video decoder 30 may perform palette-based coding of predicted video blocks. In one example, video encoder 20 first derives a palette for a current block based on the pixel values in the current block, and then maps the pixel values in the current block to palette indices for encoding. The mapping may be one to one (i.e., for lossless coding) or multiple to one (i.e., for lossy coding). Video encoder 20 may also map reference pixel values in a previously coded block that may be used to predict the pixel values in the current block. Once video encoder 20 maps the pixel values of the current block to palette indices, video encoder 20 encodes the current block using regular encoding methods, e.g., regular intra coding in the HEVC standard.

[0163] In the above example, the current block with palette indices is treated as if it were an original block with pixel values. Similarly, the palette indices of the reference

pixels are used for performing regular intra prediction on the current block with palette indices. Video encoder 20 may transmit the prediction error or residual values to video decoder 30. In some cases, the prediction error or residual values may be transformed, quantized and entropy encoded into the bitstream. In other cases, it is also possible that the transform and quantization are disabled for the palette coding mode. After encoding the current block, video encoder 20 may convert the indices of the reference pixels, prediction pixels, and/or the residual values back to the pixel values for reconstruction of the current block and the normal prediction of future blocks. Video decoder 30 may obtain the encoded residual values for the current block from the bitstream.

Furthermore, video decoder 30 may decode the current block using a regular decoding method to obtain the current block with palette indices. Video decoder 30 may then determine the pixel values of the current block based on the pixel values in the palette that are associated with the palette indices.

[0164] In another example, video encoder 20 may generate a palette for a current block. The palette may include entries that indicate prediction residual values for the given block. The prediction residual values for the given block may be generated using any prediction mode, e.g., regular inter-prediction or intra-prediction in the HEVC standard. The prediction residual values for the given block may be residual pixel values (possibly quantized) or residual transform coefficient values (possibly quantized). In this example, video encoder 20 maps the prediction residual values for the current block to index values that indicate entries in the palette for the current block that are used to represent the prediction residual values for the current block. In this example, video encoder 20 may encode the index values for one or more positions in the current block, where the index values indicate the entries in the palette for the current block that specify the prediction residual values for the current block. Video decoder 30 may obtain the encoded block of index values from the bitstream, and determine the prediction residual values for the current block based on the corresponding prediction residual values in the palette identified by the index values. Video decoder 30 may then reconstruct the pixel values of the current block using regular decoding methods based on the prediction residual values and previously coded reference pixel values.

[0165] In some examples, video encoder 20 and/or video decoder 30 may perform the palette-based video coding with video block prediction by applying the intra prediction mode (i.e., the prediction only uses previously coded pixel information in the current picture). In other examples, video encoder 20 and/or video decoder 30 may apply the

inter prediction mode (i.e., the prediction is from pixels in a previously coded pictures). In one example, the prediction residual values for the current block may be residual pixel values for the current block calculated from the pixel values of the current block and the previously coded reference pixel values. The residual pixel values may be quantized. In another example, the prediction residual values for the current block may be residual transform coefficient values for the current block calculated from the pixel values of the current block and the previously coded reference pixel values, and then transformed and possibly quantized.

[0166] In some cases, video encoder 20 and/or video decoder 30 may determine the prediction residual values for the current block using only a subset of prediction mode processes for either the inter prediction mode or the intra prediction mode. For example, in the case of the intra prediction mode, the DC, horizontal, and/or vertical prediction processes may be enabled, but other intra prediction mode processes may be disabled. The disabled processes may include the filtering in the intra prediction mode, e.g., one or more of mode-dependent intra-smoothing (MDIS), 1/32-pel bilinear interpolation, edge filter and/or DC filter (a background introduction can be found in U.S. Provisional Application No. 61/890,844, filed October 14, 2013, entitled “Adaptive Filter Control for Intra Prediction in Video Coding,” Applicant reference number 1212-671USP3/134960P3), is disabled in this example. As a further example, in the case of the inter prediction mode, the average of pixels process, e.g., one or more of the weighted prediction, the bi-prediction, or the sub-pel interpolation, may be disabled.

[0167] In another example, video encoder 20 and/or video decoder 30 may perform no prediction for the current block. In this case, video encoder 20 instead maps the pixel values to palette indices, and encodes the indices using entropy coding without prediction. In an additional example, video encoder 20 and/or video decoder 30 may perform residual differential pulse code modulation (RDPCM) using pixel values of the current block that are mapped to palette index values. In this case, no prediction from pixels outside the current block is used, and horizontal or vertical prediction may be used for line copying index values within the current CU. For example, when using the vertical prediction, the locations in the first row of the current block are not predicted, and the locations in the subsequent rows may be predicted using values in the previous rows, e.g., values in row i ($i > 0$) equal to $x(i, j)$ are predicted using $x(i-1, j)$. When using the horizontal prediction, the locations in the first column of the current block are not

predicted, and the locations in the subsequent columns may be predicted using values in the previous columns.

[0168] In some examples, the techniques for palette-based coding of video data may be used with one or more other coding techniques, such as techniques for inter- or intra-predictive coding. For example, as described in greater detail below, an encoder or decoder, or combined encoder-decoder (codec), may be configured to perform inter- and intra-predictive coding, as well as palette-based coding.

[0169] FIG. 4 is a conceptual diagram illustrating an example of determining a palette for coding video data, consistent with techniques of this disclosure. The example of FIG. 4 includes a picture 178 having a first coding unit (CU) 180 that is associated with a first set of palettes (i.e., first palettes 184) and a second CU 188 that is associated with a second set of palettes (i.e., second palettes 192). As described in greater detail below and in accordance with one or more of the techniques of this disclosure, second palettes 192 are 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.

[0170] The techniques of FIG. 4 are described in the context of video encoder 20 (FIG. 1 and FIG. 2) and video decoder 30 (FIG. 1 and FIG. 3) and with respect to the HEVC video coding standard for purposes of explanation. With respect to the HEVC framework, as an example, the palette-based coding techniques may be configured to be used as a CU mode. In other examples, the palette-based coding techniques may be configured to be used as a PU mode or a TU mode in the framework of HEVC.

Accordingly, all of the following disclosed processes described in the context of a CU mode may, additionally or alternatively, apply to a PU or a TU. 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.

[0171] In general, a palette refers to a number of pixel values that are dominant and/or representative for a CU currently being coded (e.g., CU 188 in the example of FIG. 4). First palettes 184 and second palettes 192 are shown as including multiple palettes. In some examples, according to aspects of this disclosure, a video coder (such as video encoder 20 or video decoder 30) may code palettes separately for each color component of a CU. For example, video encoder 20 may encode a palette for a luma (Y) component of a CU, another palette for a chroma (U) component of the CU, and yet

another palette for the chroma (V) component of the CU. In this example, entries of the Y palette may represent Y values of pixels of the CU, entries of the U palette may represent U values of pixels of the CU, and entries of the V palette may represent V values of pixels of the CU.

[0172] In other examples, video encoder 20 may encode a single palette for all color components of a CU. 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 palettes 184 and 192 as a set of palettes having multiple individual palettes is merely one example and not intended to be limiting.

[0173] In the example of FIG. 4, 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.

[0174] As described herein, rather than coding the actual pixel values of first CU 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 CU 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 CU or may be shared among different CUs.

[0175] According to aspects of this disclosure, video encoder 20 and video decoder 30 may determine second palettes 192 based on first palettes 184. For example, video encoder 20 may encode a `pred_palette_flag` for each CU (including, as an example,

second CU 188) to indicate whether the palette for the CU is predicted from one or more palettes associated with one or more other CUs, such as neighboring CUs (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 CU 188 are predicted from one or more already decoded palettes and therefore no new palettes for second CU 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 palettes 192 for second CU 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 CU (e.g., three flags, one for Y, one for U, and one for V, for a CU in YUV video). In other examples, a single `pred_palette_flag` may be coded for all color components of a CU.

[0176] In the example above, the `pred_palette_flag` is signaled per-CU 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.

[0177] 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 CUs of the CU currently being coded (e.g., such as neighboring CUs (spatially or based on scan order) or the most frequent samples of a causal neighbor), i.e., second CU 188. The palettes of the one or more neighboring CUs may be associated with a predictive palette. In some examples, such as the example illustrated in FIG. 4, video encoder 20 and/or video decoder 30 may locate a left neighboring CU, first CU 180, when determining a predictive palette for second CU 188. In other examples, video encoder 20 and/or video decoder 30 may locate one or more CUs in other positions relative to second CU 188, such as an upper CU, CU 196. In another

example, the palette for the last CU in scan order that used the palette mode may be used as a predictive palette.

[0178] Video encoder 20 and/or video decoder 30 may determine a CU for palette prediction based on a hierarchy. For example, video encoder 20 and/or video decoder 30 may initially identify the left neighboring CU, first CU 180, for palette prediction. If the left neighboring CU is not available for prediction (e.g., the left neighboring CU 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 CU, CU 196. Video encoder 20 and/or video decoder 30 may continue searching for an available CU according to a predetermined order of locations until locating a CU 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.

[0179] While the example of FIG. 4 illustrates first palettes 184 as predictive palettes from a single CU, (i.e., first CU 180), in other examples, video encoder 20 and/or video decoder 30 may locate palettes for prediction from a combination of neighboring CUs. 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 CUs (spatially or in scan order).

[0180] 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 CU in the list from which the current CU 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 CU for use with the current CU. In another example, the palette of the indicated candidate CU in the list may be used as a predictive palette for per-entry prediction of a current palette for the current CU.

[0181] In an example for purposes of illustration, video encoder 20 and video decoder 30 may construct a candidate list that includes one CU that is positioned above the CU currently being coded and one CU that is positioned to the left of the CU 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 CU is copied from the CU positioned to the left of the current CU. Video encoder 20 may encode the flag having a value of one to indicate that the palette for the current CU is copied from the CU positioned above the current CU. Video decoder 30 decodes the flag and selects the appropriate CU for palette prediction. In another example, the flag may indicate whether the palette of the top or left neighboring CU 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 CU.

[0182] In still other examples, video encoder 20 and/or video decoder 30 determine the palette for the CU currently being coded based on the frequency with which sample values included in one or more other palettes occur in one or more neighboring CUs. 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 CUs. Video encoder 20 and/or video decoder 30 may include the most frequently used colors in the palette for the CU currently being coded.

[0183] As noted above, in some examples, video encoder 20 and/or video decoder may copy an entire palette from a neighboring CU for coding a current CU. 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 CU). 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 CU). 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 CU. In this example, video encoder 20 may also encode additional data indicating the value of the non-predicted palette entry.

[0184] This disclosure describes several alternative techniques for predicting a palette for a current CU. In one example, a predictive palette that includes palette entries from one or more previously coded neighboring CUs 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 CU. For example, $V(i) = 1$ means that the i -th entry

in the predictive palette for the neighboring CU will be reused or copied to the palette for the current CU, which may have a different index in the current CU.

[0185] 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 CU, 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 CU 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 CU (i.e., $V(i) = 1$). To generate the palette for the current CU, 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 CU 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.

[0186] In another example, video encoder 20 first transmits an indication of a size of a palette, N , for a current CU to video decoder 30. Video encoder 20 then transmits a vector, V , having the same size as the palette for the current CU, i.e., size N , to video decoder 30. Each entry in the vector indicates whether the corresponding entry in the palette for the current CU 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 CU. 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 CU. In other cases, video encoder 20 may signal an index offset, which is the difference between the index in the palette for the current CU and the index in the predictive palette.

[0187] In the two above examples, the one or more previously coded neighboring CUs used to generate the predictive palette used for prediction of the palette for the current CU may be a top-neighboring (i.e., upper) CU or a left-neighboring CU with respect to the current CU. In some examples, a candidate list of neighboring CUs may be constructed, and video encoder 20 transmits an index to indicate which candidate neighboring CUs and associated palettes are used for palette prediction for the current

CU. For certain CUs, e.g., CUs that are positioned at a beginning of a slice or at other slice boundaries or leftmost CUs in the slice or a picture of video data, palette prediction may be disabled.

[0188] In an additional example, video encoder 20 transmits an indication of a number of entries included in a palette for a current CU 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 CU or a neighboring CU 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 CU. 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 CU. For example, this list may include all the reference pixels that may be used for normal intra prediction in HEVC.

[0189] In the example of FIG. 4, 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 CU 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 CU 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.

[0190] 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. 4 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.

[0191] 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.

[0192] 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 CU 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. That is, with respect to the size of the CU 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 CUs 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.

[0193] 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 CU 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.

[0194] 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 CU 188, there are no entries in palettes 192. As video encoder 20 and video decoder 30 code new values for pixels of second CU 188, each new value is included in palettes 192. That is, for example, video encoder 20 adds pixel values to palettes 192 as the pixel values are generated and signaled for positions in CU 188. As video encoder 20 encodes pixels relatively later in the CU, 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 CU 188, video decoder 30

includes the pixel value in palettes 192. When pixel positions decoded relatively later in second CU 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 CU 188.

[0195] In some examples, as described in greater detail below, video encoder 20 and/or video decoder 30 may maintain palettes 184 and 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.

[0196] 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 CU 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 CU.

[0197] 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 CU 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 CU 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 CU are scanned.

[0198] Video encoder 20 may, in some examples, signal whether the current CU 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.

[0199] 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.

[0200] 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 CU 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).

[0201] 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 CU without exactly matching the pixel values of

palettes exactly to the actual pixel values in the CU. 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.

[0202] 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.

[0203] 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.

[0204] 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.

[0205] 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.

[0206] 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).

[0207] 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.

[0208] 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.

[0209] In some examples, the delta value may be different for different CU 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 CU size, color space, color component, or the like, as noted above).

[0210] 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.

[0211] 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 CU 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 CU, 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.

[0212] 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.

[0213] 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.

[0214] 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.

[0215] 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 CU, 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 CU (or a PU in other examples). For example, this flag may be signaled in an encoded video bitstream at the CU level, and then received by video decoder 30 upon decoding the encoded video bitstream.

[0216] In this example, a value of this PLT_Mode_flag equal to 1 may specify that the current CU is encoded using a palette-based coding mode. In this case, video decoder 30 may apply the palette-based coding mode to decode the CU. In some examples, a syntax element may indicate one of a plurality of different palette-based coding modes for the CU (e.g., lossy or lossless). A value of this PLT_Mode_flag equal to 0 may specify that the current CU 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 CU (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 CU (or PU in other examples) or to indicate which of a plurality of modes are to be used for coding.

[0217] In some examples, the above-described flag or other syntax elements may be transmitted at a higher level than the CU (or PU) 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 CUs 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 CU level. In another example, video encoder 20 may signal such a flag in a PPS, SPS or VPS.

[0218] 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 CU. 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 CU. 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 CU. Again, as mentioned above, application of these techniques for indicating palette-based coding of a CU could additionally or alternatively be used to indicate palette-based coding of a PU.

[0219] 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 CU, 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.

[0220] 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 CU, 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.

[0221] While the techniques of FIG. 4 are described above in the context of CUs (HEVC), it should be understood that the techniques may also be applied to prediction units (PUs) or in other video coding processes and/or standards.

[0222] FIG. 5 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. 5 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. 4.

[0223] Again, the techniques of FIG. 5 are described in the context of video encoder 20 (FIG. 1 and FIG. 2) and video decoder 30 (FIG. 1 and FIG. 3) 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.

[0224] While map 240 is illustrated in the example of FIG. 5 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.

[0225] 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 CU. 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 CU 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. 5 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 CU 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.

[0226] The value of a pixel in one position of a CU may provide an indication of values of one or more other pixels in other positions of the CU. For example, there may be a relatively high probability that neighboring pixel positions of a CU 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).

[0227] 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.

[0228] In some examples, all pixel locations in the current CU having pixel values that are in the palette for the current CU 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 CU. In the case where the pixel value at one of the pixel locations in the current CU 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 CU. 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.

[0229] 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).

[0230] 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.

[0231] 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 line 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).

[0232] 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.

[0233] 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.

[0234] 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.

[0235] 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.

[0236] 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.

[0237] 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.

[0238] 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. 5). 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.

[0239] 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.

[0240] 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.

[0241] 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 picture parameter set (PPS), a sequence parameter set (SPS), a video parameter set (VPS) and/or a slice header.

[0242] 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.

[0243] 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.

[0244] 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.

[0245] FIG. 6 is a conceptual diagram illustrating examples of determining a geometric edge 270, 272 or 274 of a video block using a run of palette indices for the luma component adaptively downsampled for the chroma components, consistent with techniques of this disclosure. In FIG. 6, the luma samples are illustrated as un-filled circles, and the chroma samples are illustrated as one of the luma samples overlaid with an x-symbol. FIG. 6 illustrates examples of different run values for luma and chroma components based on a location of geometric edge 270, 272 or 274 of the video block.

[0246] In some cases, one palette is generated and shared for multiple color components in the current block, and in other cases, separate palettes are generated for one or more of the color components. In one case, one palette may be generated for the luma component and another palette may be generated for both the chroma components. In either case, the geometric information may be shared between the color components. Usually there is high correlation between edge locations of collocated blocks in different color components because the chroma components may have been downsampled from the luma components in a pre-defined way, such as 4:2:2 or 4:2:0 sampling.

[0247] For example, in palette-based coding, run coding may be used to indicate geometry information for the current block because an edge of the current block will break the run. In case of the 4:4:4 chroma format, the run may be generated once and used for all color components. The run may be generated based on one of the color components, or the run may be generated using more than one of the color components. In case of the 4:2:2 chroma format, the run used for the luma component may be horizontally downsampled by a factor of two for application to the chroma components. In the case of the 4:2:0 chroma format, the run used for the luma component may be horizontally and vertically downsampled by a factor of two for application to the chroma components.

[0248] In some cases, the run downsampling method can be adaptive to a chroma downsampling method. In this case, the downsampled run value for the chroma components may be differently calculated according to the location of the edge, e.g., edge 270, 272 or 274, of the video block as shown in FIG. 6. In a first example, FIG. 6 illustrates a geometric edge 270 between two neighboring video blocks that is positioned such that a run for the luma component has a value "1" in the left-hand block and a value of "3" in the right-hand block. In this case, the downsampled run for the chroma components has a value of "1" in both the left-hand block and the right-hand block. In a second example, FIG. 6 illustrates a geometric edge 272 between two neighboring video blocks that is positioned such that a run for the luma component has a value "2" in both the left-hand block and the right-hand block. In this case, the downsampled run for the chroma components has a value of "1" in both the left-hand block and the right-hand block. In a third example, FIG. 6 illustrates a geometric edge 274 between two neighboring video blocks that is positioned such that a run for the luma component has a value "3" in the left-hand block and a value of "1" in the right-

hand block. In this case, the downsampled run for the chroma components has a value of “2” in the left-hand block and a value of “0” in the right-hand block.

[0249] In addition to the geometric information, it may also be possible to have a single palette for pixel value of all color components. For example, for each pixel location in the current block, the pixel values in three color components (e.g., Y luma and both U and V chroma components) may form a vector (i.e., a color vector). Then, a palette may be formed by selecting a certain number of vectors to represent the current block. It may be possible to have one palette of pixel values for the luma component, and another palette of pixel values for the chroma components. In some cases, it may also be possible to combine the two methods of sharing geometry information and having a single palette of pixel values using a color vector.

[0250] In some examples, the line copying described in more detail elsewhere in this disclosure may also work with a single palette. In this case, the palette index for each pixel location is signaled as being equal to the palette index of the row above, if the scan is horizontal, or the column on the left, if the scan is vertical, and then the associated run of palette indices is also copied from the previous row or column. With a shared palette, a palette entry may be a triplet of (Y, U, V), so that later Y, U, V values may be reconstructed from the palette index. The reconstructed values may serve as the decoded pixel values or may serve as prediction values that are combined with residual values to derive the decoded pixel values. In the 4:2:2 chroma format and the 4:2:0 chroma format, the chroma components have been downsampled compared to the luma components. In the example of a 2:1 downsampling, the luma positions may be at 0, 1, 2, ..., and the chroma positions may be at 1, 3, 5, ... or may be at 0, 2, 4, For positions where chroma components do not exist, the U and V components in the palette entry may be discarded.

[0251] FIG. 7 is a flowchart illustrating an example process for encoding prediction residual video data using a palette-based coding mode, consistent with techniques of this disclosure. The example process illustrated in FIG. 7 is described herein with respect to palette-based encoding unit 122 of video encoder 20 from FIG. 2. In other examples, one or more other or additional components may perform the example process illustrated in FIG. 7.

[0252] Video encoder 20 receives video data of a current video block to be encoded using palette-based video coding of predicted video blocks, and sends the video data to palette-based encoding unit 122. Palette-based encoding unit 122 determines prediction

residual values for the current block based on pixel values of the current block and previously coded reference pixel values (280).

[0253] Palette-based encoding unit 122 may calculate the prediction residual values using any prediction mode, e.g., an inter-prediction mode or an intra-prediction mode of the HEVC standard. In one example, palette-based encoding unit 122 may use inter-prediction processing unit 120 to predict pixel values of the current block using previously coded pixel values in a reference block. In another example, palette-based encoding unit 122 may use intra-prediction processing unit 126 to predict pixel values of the current block using previously coded pixel values in the current block.

[0254] In some cases, palette-based encoding unit 122 may determine the prediction residual values for the current block using only a subset of prediction mode processes. For example, in the case of the intra prediction mode, the DC, horizontal, and/or vertical prediction processes may be enabled, but other intra prediction mode processes may be disabled. The disabled processes may include the filtering in the intra prediction mode, e.g., one or more of MDIS, 1/32 pel bilinear interpolation, edge filter or DC filter. As a further example, in the case of the inter prediction mode, the average of pixels process, e.g., one or more of the weighted prediction, the bi-prediction, or the sub-pel interpolation, may be disabled.

[0255] In one example, the prediction residual values for the current block may be residual pixel values for the current block. In this example, palette-based encoding unit 122 calculates the residual pixel values for the current block from the pixel values of the current block and the previously coded reference pixel values. Palette-based encoding unit 122 then proceeds to encode the residual pixel values for the current block using palette-based video coding as described in the following steps.

[0256] In another example, the prediction residual values for the current block may be residual quantized transform coefficient values for the current block. In this example, palette-based encoding unit 122 calculates residual pixel values for the current block from the pixel values of the current block and the previously coded reference pixel values, and then sends the residual pixel values to transform processing unit 104 and quantization unit 106 to be transformed and quantized into residual quantized transform coefficient values for the current block. Palette-based encoding unit 122 then proceeds to encode the residual quantized transform coefficients values for the current block using palette-based video coding as described in the following steps.

[0257] Palette-based encoding unit 122 generates a palette for the current block including one or more entries that indicate the prediction residual values for the current block (282). Palette-based encoding unit 122 maps one or more of the prediction residual values for the current block to index values that identify entries in the palette used to represent the prediction residual values in the palette for the current block (284). Palette-based encoding unit 122 encodes the index values for one or more positions in the current block (286). The encoded index values indicate the prediction residual values included in the palette for the current block that are used to represent the prediction residual values for the current block. Video encoder 20 then transmits the index values for the one or more positions in the current block.

[0258] FIG. 8 is a flowchart illustrating an example process for decoding prediction residual video data using a palette-based coding mode, consistent with techniques of this disclosure. The example process illustrated in FIG. 8 is described herein with respect to palette-based decoding unit 165 of video decoder 30 from FIG. 3. In other examples, one or more other or additional components may perform the example process illustrated in FIG. 8.

[0259] Video decoder 30 receives a bitstream representing coded video data using palette-based coding, and sends the entropy decoded video data to palette-based decoding unit 165. Based on one or more syntax elements included in the decoded bitstream, palette-based decoding unit 165 generates a palette for a current block of video data including one or more entries that indicate prediction residual values for the current block (290). Palette-based decoding unit 165 then decodes index values for one or more positions in the current block (292). The decoded index values indicate the prediction residual values included in the palette for the current block that are used to represent the prediction residual values for the current block.

[0260] Palette-based decoding unit 165 determines one or more of the prediction residual values for the current block based on the index values that identify entries in the palette that represent the prediction residual values for the current block (294). Palette-based decoding unit 165 may determine the prediction residual values using any prediction mode, e.g., an inter-prediction mode or an intra-prediction mode of the HEVC standard. In one example, palette-based decoding unit 165 may use motion compensation unit 164 to predict pixel values of the current block using previously coded pixel values in a reference block. In another example, palette-based decoding

unit 165 may use intra-prediction processing unit 166 to predict pixel values of the current block using previously coded pixel values in the current block.

[0261] In some cases, palette-based decoding unit 165 may determine the prediction residual values for the current block using only a subset of prediction mode processes. For example, in the case of the intra prediction mode, the DC, horizontal, and/or vertical prediction processes may be enabled, but other intra prediction mode processes may be disabled. The disabled processes may include the filtering in the intra prediction mode, e.g., one or more of MDIS, 1/32 pel bilinear interpolation, edge filter or DC filter. As a further example, in the case of the inter prediction mode, the average of pixels process, e.g., one or more of the weighted prediction, the bi-prediction, or the sub-pel interpolation, may be disabled.

[0262] Video decoder 30 then determines pixel values of the current block based on the prediction residual values for the current block and previously coded reference pixel values (296). In one example, the prediction residual values for the current block may be residual pixel values for the current block. In this case, video decoder 30 reconstructs the pixel values of the current block using the residual pixel values and the previously coded reference pixel values. In another example, the prediction residual values for the current block may be residual quantized transform coefficient values for the current block. In this case, palette-based decoding unit 165 first sends the residual quantized transform coefficient values to inverse quantization unit 154 and inverse transform processing unit 156 to be inverse quantized and inverse transformed into residual pixel values for the current block. Video decoder 30 then reconstructs the pixel values of the current block using the residual pixel values and the previously coded reference pixel values.

[0263] FIG. 9 is a flowchart illustrating an example process for generating a palette for palette-based coding, consistent with techniques of this disclosure. The example process illustrated in FIG. 9 is described herein with respect to palette-based decoding unit 165 of video decoder 30 from FIG. 3. In other examples, the process may also be performed by palette-based encoding unit 122 of video encoder 20 from FIG. 2. The example process for generating a palette for palette-based coding may be used to generate a palette including palette entries that indicate pixel values. In other examples, a similar process may be used to generate a palette including palette entries that indicate prediction residual values.

[0264] Video decoder 30 receives a bitstream representing coded video data using palette-based coding, and sends the entropy decoded video data to palette-based decoding unit 165. Palette-based decoding unit 165 generates a predictive palette including palette entries that indicate pixel values (300). In some examples, palette-based decoding unit 165 generates the predictive palette to include palette entries from one or more previously coded blocks of the video data. The previously coded blocks may include neighboring blocks of a current block including spatially neighboring blocks and/or neighboring blocks in a particular scan order of the blocks.

[0265] Palette-based decoding unit 165 next determines, from the entropy decoded video data, one or more of the palette entries in the predictive palette that are copied to a current palette for the current block (302). More specifically, palette-based decoding unit 165 may decode one or more syntax elements indicating whether each of the palette entries in the predictive palette is copied to the current palette. In one example, the one or more syntax elements comprise a binary vector including a flag for each of the palette entries in the predictive palette that indicates whether a respective palette entry is copied to the current palette. In another example, the one or more syntax elements comprise a losslessly compressed version of the binary vector, where an uncompressed version of the binary vector includes a flag for each of the palette entries in the predictive palette that indicates whether a respective palette entry is copied to the current palette.

[0266] Palette-based decoding unit 165 also determines, from the entropy decoded video data, a number of new palette entries not in the predictive palette that are included in the current palette for the current block (304). Palette-based decoding unit 165 may decode one or more syntax elements indicating the number of the new palette entries that are included in the current palette. In some examples, palette-based decoding unit 165 decodes the syntax elements using one of unary codes, truncated unary codes, Exponential-Golomb codes, or Golomb-Rice codes. After determining the number of new palette entries that are included in the current palette, palette-based decoding unit 165 decodes one or more syntax elements indicating a pixel value for each of the new palette entries.

[0267] Based on the information determined from the entropy decoded video data, palette-based decoding unit 165 calculates a size of the current palette to be equal to the sum of the number of the copied palette entries and the number of the new palette entries (306). After determining the size of the current palette, palette-based decoding unit 165 generates the current palette to include the copied palette entries and the new

palette entries (308). In one example, the palette-based decoding unit 165 generates the current palette by concatenating the copied palette entries and the new palette entries.

[0268] Palette-based decoding unit 165 is then able to perform palette-based coding of the current block using the current palette. For example, palette-based decoding unit 165 determines index values for one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block (310). In the case where one or more pixel values of the current block do not have a corresponding pixel value in the current palette, video encoder 20 may use the escape pixel concept to indicate which of the pixel values are not included in the current palette, and explicitly transmit these pixel values. Palette-based decoding unit 165 in video decoder 30 may then decode one or more syntax elements indicating the pixel values for the one or more pixel values that do not have a corresponding pixel value in the current palette.

[0269] In another example, video encoder 20 may not use the escape pixel concept, but instead may identify pixel values included in the current palette as prediction pixel values for each of the one or more pixel values of the current block, and transmit residual values between the pixel values of the current block and the prediction pixel values in the current palette. Palette-based decoding unit 165 in video decoder 30 may then decode one or more syntax elements indicating the index values that identify the corresponding prediction pixel values included in the current palette, and the residual values between the one or more pixel values of the current block and the identified prediction pixel values in the current palette.

[0270] This disclosure also describes several alternative techniques for generating a palette for palette-based coding, which may be used to generate a palette having entries that associate index values with the either pixel values or prediction residual values for a current block. In one example, palette-based decoding unit 165 decodes an indication of a size of the palette for the current block, decodes a vector having the same size as the palette for the current block, where each entry in the vector indicates whether an associated palette entry is transmitted or copied from a predictive palette, and, for the one or more palette entries copied from the predictive palette, decodes an indication of a position of the entry in the predictive palette. In another example, palette-based decoding unit 165 decodes an indication of a number of entries in the palette for the current block, decodes a one-bit flag for each of the palette entries that indicates whether the palette entry is sent explicitly or derived from a previously reconstructed

pixel, and, for each of the one or more palette entries derived from the previously reconstructed pixel, decodes an indication of a position of the reconstructed pixel that corresponds to the respective palette entry. In that example, the indication of the position of the reconstructed pixel may be a displacement vector with respect to the top-left position of the current block or it may be an index into a list of reconstructed pixels that may include all the reference pixels used for normal intra prediction.

[0271] In another example, palette-based decoding unit 165 starts with a predictive palette for a neighboring block having a given size, and decodes a binary vector having the same size as the predictive palette, where each entry in the vector indicates whether an associated palette entry is reused from the predictive palette. Palette-based decoding unit 165 also decodes an indication of the number of new entries to be transmitted, and receives the new entries from video encoder 20. Palette-based decoding unit 165 then merges the reused entries and the new entries to generate the new palette for the current block.

[0272] It is to be recognized that depending on the example, certain acts or events of any of the techniques described herein can be performed in a different sequence, may be added, merged, or left out altogether (e.g., not all described acts or events are necessary for the practice of the techniques). Moreover, in certain examples, acts or events may be performed concurrently, e.g., through multi-threaded processing, interrupt processing, or multiple processors, rather than sequentially. In addition, while certain aspects of this disclosure are described as being performed by a single module or unit for purposes of clarity, it should be understood that the techniques of this disclosure may be performed by a combination of units or modules associated with a video coder.

[0273] Certain aspects of this disclosure have been described with respect to the developing HEVC standard for purposes of illustration. However, the techniques described in this disclosure may be useful for other video coding processes, including other standard or proprietary video coding processes not yet developed.

[0274] The techniques described above may be performed by video encoder 20 (FIGS. 1 and 2) and/or video decoder 30 (FIGS. 1 and 3), both of which may be generally referred to as a video coder. Likewise, video coding may refer to video encoding or video decoding, as applicable.

[0275] While particular combinations of various aspects of the techniques are described above, these combinations are provided merely to illustrate examples of the techniques described in this disclosure. Accordingly, the techniques of this disclosure should not

be limited to these example combinations and may encompass any conceivable combination of the various aspects of the techniques described in this disclosure.

[0276] 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.

[0277] 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.

[0278] Instructions may be executed by one or more processors, such as one or more digital signal processors (DSPs), general purpose microprocessors, application specific integrated circuits (ASICs), field programmable gate arrays (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.

[0279] 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.

[0280] 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 coding video data, the method comprising:
 - generating a predictive palette including palette entries that indicate pixel values;
 - determining one or more of the palette entries in the predictive palette that are copied to a current palette for a current block of the video data;
 - determining a number of new palette entries that are included in the current palette for the current block, wherein the new palette entries are not in the predictive palette;
 - calculating a size of the current palette equal to the sum of a number of the copied palette entries and the number of the new palette entries;
 - generating the current palette including the copied palette entries and the new palette entries; and
 - determining index values for one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.
2. The method of claim 1, wherein generating the predictive palette comprises generating the predictive palette including palette entries from one or more previously coded blocks of the video data.
3. The method of claim 2, wherein the one or more previously coded blocks comprise neighboring blocks of the current block including at least one of spatially neighboring blocks or neighboring blocks in a particular scan order of the blocks.
4. The method of claim 1, wherein determining the one or more palette entries in the predictive palette that are copied to the current palette comprises coding one or more syntax elements indicating whether each of the palette entries in the predictive palette is copied to the current palette.
5. The method of claim 4, wherein the one or more syntax elements comprise a binary vector including a flag for each of the palette entries in the predictive palette that indicates whether a respective palette entry is copied to the current palette.

6. The method of claim 4, wherein the one or more syntax elements comprise a losslessly compressed version of a binary vector, wherein an uncompressed version of the binary vector includes a flag for each of the palette entries in the predictive palette that indicates whether a respective palette entry is copied to the current palette.
7. The method of claim 1, wherein determining the number of the new palette entries that are included in the current palette comprises coding one or more syntax elements indicating the number of the new palette entries that are included in the current palette, wherein the one or more syntax elements are coded using one of unary codes, truncated unary codes, Exponential-Golomb codes, or Golomb-Rice codes.
8. The method of claim 1, further comprising coding one or more syntax elements indicating a pixel value for each of the new palette entries included in the current palette.
9. The method of claim 1, wherein generating the current palette including the copied palette entries and the new palette entries comprises concatenating the copied palette entries and the new palette entries.
10. The method of claim 1, wherein the current block comprises one of a first block in a slice of the video data or a leftmost block of the slice or a picture of the video data, further comprising disabling copying of the one or more palette entries in the prediction palette to the current palette for the current block.
11. The method of claim 1, wherein determining the index values for the one or more pixel values of the current block comprises:
 - coding one or more syntax elements indicating whether each of the pixel values of the current block has a corresponding pixel value in the current palette;
 - coding one or more syntax elements indicating the index values for the one or more pixel values of the current block that have corresponding pixel values in the current palette; and
 - coding one or more syntax element indicating the pixel values for the one or more pixel values of the current block that do not have a corresponding pixel value in the current palette.

12. The method of claim 1, wherein determining the index values for the one or more pixel values of the current block comprises:

coding one or more syntax elements indicating the index values for the one or more pixel values of the current block, wherein the index values identify corresponding pixel values in the current palette as prediction pixel values; and

coding one or more syntax elements indicating residual values between the one or more pixel values of the current block and the identified prediction pixel values in the current palette.

13. The method of claim 1, wherein the method of coding video data comprises a method of decoding video data, the method comprising:

receiving an indication of the one or more of the palette entries in the predictive palette that are copied to the current palette;

receiving the number of the new palette entries that are included in the current palette;

receiving a pixel value for each of the new palette entries;

generating the current palette including the copied palette entries and the received new palette entries; and

receiving the index values for the one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.

14. The method of claim 1, wherein the method of coding video data comprises a method of encoding video data, the method comprising:

generating the current palette including the copied palette entries and the new palette entries;

transmitting an indication of the one or more of the palette entries in the predictive palette that are copied to the current palette;

transmitting the number of the new palette entries that are included in the current palette;

transmitting a pixel value for each of the new palette entries; and

transmitting the index values for the one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.

15. An apparatus for coding video data, the apparatus comprising:
a memory storing video data; and
one or more processors configured to:
generate a predictive palette including palette entries that indicate pixel values;
determine one or more of the palette entries in the predictive palette that are copied to a current palette for a current block of the video data;
determine a number of new palette entries that are included in the current palette for the current block, wherein the new palette entries are not in the predictive palette;
calculate a size of the current palette equal to the sum of a number of the copied palette entries and the number of the new palette entries;
generate the current palette including the copied palette entries and the new palette entries; and
determine index values for one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.
16. The apparatus of claim 15, wherein the processors are configured to generate the predictive palette including palette entries from one or more previously coded blocks of the video data.
17. The apparatus of claim 16, wherein the one or more previously coded blocks comprise neighboring blocks of the current block including at least one of spatially neighboring blocks or neighboring blocks in a particular scan order of the blocks.
18. The apparatus of claim 15, wherein the processors are configured to code one or more syntax elements indicating whether each of the palette entries in the predictive palette is copied to the current palette.
19. The apparatus of claim 18, wherein the one or more syntax elements comprise a binary vector including a flag for each of the palette entries in the predictive palette that indicates whether a respective palette entry is copied to the current palette.

20. The apparatus of claim 18, wherein the one or more syntax elements comprise a losslessly compressed version of a binary vector, wherein the uncompressed version of the binary vector includes a flag for each of the palette entries in the predictive palette that indicates whether a respective palette entry is copied to the current palette.

21. The apparatus of claim 15, wherein the processors are configured to code one or more syntax elements indicating the number of the new palette entries that are included in the current palette, wherein the one or more syntax elements are coded using one of unary codes, truncated unary codes, Exponential-Golomb codes, or Golomb-Rice codes.

22. The apparatus of claim 15, wherein the processors are configured to code one or more syntax elements indicating a pixel value for each of the new palette entries included in the current palette.

23. The apparatus of claim 15, wherein the processors are configured to concatenate the copied palette entries and the new palette entries to generate the current palette.

24. The apparatus of claim 15, wherein the current block comprises one of a first block in a slice of the video data or a leftmost block of the slice or a picture of the video data, and wherein the processors are configured to disable copying of the one or more palette entries in the prediction palette to the current palette for the current block.

25. The apparatus of claim 15, wherein the processors are configured to:

- code one or more syntax elements indicating whether each of the pixel values of the current block has a corresponding pixel value in the current palette;
- code one or more syntax elements indicating the index values for the one or more pixel values of the current block that have corresponding pixel values in the current palette; and
- code one or more syntax element indicating the pixel values for the one or more pixel values of the current block that do not have a corresponding pixel value in the current palette.

26. The apparatus of claim 15, wherein the processors are configured to:
- code one or more syntax elements indicating the index values for the one or more pixel values of the current block, wherein the index values identify corresponding pixel values in the current palette as prediction pixel values; and
 - code one or more syntax elements indicating residual values between the one or more pixel values of the current block and the identified prediction pixel values in the current palette.
27. The apparatus of claim 15, wherein the apparatus comprises a video decoding device for decoding video data, and wherein the processors are configured to:
- receive an indication of the one or more of the palette entries in the predictive palette that are copied to the current palette;
 - receive the number of the new palette entries that are included in the current palette;
 - receive a pixel value for each of the new palette entries;
 - generate the current palette including the copied palette entries and the received new palette entries; and
 - receive the index values for the one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.
28. The apparatus of claim 15, wherein the apparatus comprises a video encoding device for encoding video data, and wherein the processors are configured to:
- generate the current palette including the copied palette entries and the new palette entries;
 - transmit an indication of the one or more of the palette entries in the predictive palette that are copied to the current palette;
 - transmit the number of the new palette entries that are included in the current palette;
 - transmit a pixel value for each of the new palette entries; and
 - transmit the index values for the one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.

29. An apparatus for coding video data, the apparatus comprising:
- means for generating a predictive palette including palette entries that indicate pixel values;
 - means for determining one or more of the palette entries in the predictive palette that are copied to a current palette for a current block of the video data;
 - means for determining a number of new palette entries that are included in the current palette for the current block, wherein the new palette entries are not in the predictive palette;
 - means for calculating a size of the current palette equal to the sum of a number of the copied palette entries and the number of the new palette entries;
 - means for generating the current palette including the copied palette entries and the new palette entries; and
 - means for determining index values for one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.
30. A non-transitory computer-readable medium storing instructions thereon that, when executed, cause one or more processors to:
- generate a predictive palette including palette entries that indicate pixel values;
 - determine one or more of the palette entries in the predictive palette that are copied to a current palette for a current block of the video data;
 - determine a number of new palette entries that are included in the current palette for the current block, wherein the new palette entries are not in the predictive palette;
 - calculate a size of the current palette equal to the sum of a number of the copied palette entries and the number of the new palette entries;
 - generate the current palette including the copied palette entries and the new palette entries; and
 - determine index values for one or more pixel values of the current block that identify the palette entries in the current palette used to represent the pixel values of the current block.