

LINE BUFFER REDUCTION FOR ADAPTIVE LOOP FILTERING IN VIDEO CODING

[0001] This application claims the benefit of U.S. Provisional Application No. 62/591,504, filed November 28, 2017, and claims priority to U.S. Application Number 16/200,197, filed November 26, 2018, the entire content of which is incorporated by reference herein.

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 ITU-T H.265, High Efficiency Video Coding (HEVC) standard, and extensions of such standards. The video devices may transmit, receive, encode, decode, and/or store digital video information more efficiently by implementing such video compression techniques.

[0004] Video compression techniques may 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 (e.g., a video frame or a portion of a video frame) may be partitioned into video blocks, such as coding tree blocks and coding blocks. 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. For further compression, the residual data may be transformed from the pixel domain to a transform domain, resulting in residual transform coefficients, which then may be quantized.

SUMMARY

[0005] In general, this disclosure describes techniques related to filters for use in video coding. In some examples, this disclosure describes techniques for complexity reduction of adaptive loop filters (ALFs) and bilateral filters used in video coding. The techniques of this disclosure may include enabling or disabling a filter (e.g., ALF, bilateral filter, or other filters) for certain samples of a block based on positions of the samples within the block. Other techniques of this disclosure include using padded samples instead of samples outside of largest coding unit rows to perform filtering. In this way, the need for line buffers to store samples from neighboring blocks may be eliminated and/or reduced. The techniques of this disclosure may be used in the context of advanced video codecs, such as extensions of HEVC or the next generation of video coding standards, such as H.266/VVC (versatile video coding).

[0006] In one example, this disclosure describes a method of coding video data, the method comprising coding a block of video data to obtain a coded block of video data, applying an adaptive loop filter defined by a filter support to samples of the coded block of video data, including using padded samples for portions of the filter support that are outside of the block, and outputting the coded block of video data after applying adaptive loop filtering to the coded block of video data.

[0007] In another example, this disclosure describes an apparatus configured to code video data, the apparatus comprising a memory configured to store a coded block of video data, and one or more processors implemented in circuitry in communication with the memory, the one or more processors configured to code a block of video data to obtain a coded block of video data, apply an adaptive loop filter defined by a filter support to samples of the coded block of video data, including using padded samples for portions of the filter support that are outside of the block, and output the coded block of video data after applying adaptive loop filtering to the coded block of video data.

[0008] In another example, this disclosure describes an apparatus configured to code video data, the apparatus comprising means for coding a block of video data to obtain a coded block of video data, means for applying an adaptive loop filter defined by a filter support to samples of the coded block of video data, including using padded samples for portions of the filter support that are outside of the block, and means for outputting the coded block of video data after applying adaptive loop filtering to the coded block of video data.

[0009] In another example, this disclosure describes a non-transitory computer-readable medium storing instructions that, when executed, cause one or more processors to code a block of video data to obtain a coded block of video data, apply an adaptive loop filter defined by a filter support to samples of the coded block of video data, including using padded samples for portions of the filter support that are outside of the block, and output the coded block of video data after applying adaptive loop filtering to the coded block of video data.

[0010] The details of one or more aspects of the disclosure are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the techniques described in this disclosure will be apparent from the description, drawings, and claims.

BRIEF DESCRIPTION OF DRAWINGS

[0011] FIG. 1 is a block diagram illustrating an example video encoding and decoding system that may use one or more techniques described in this disclosure.

[0012] FIGS. 2A and 2B are conceptual diagrams illustrating an example quadtree binary tree (QTBT) structure, and a corresponding coding tree unit (CTU).

[0013] FIG. 3 illustrates three different example Adaptive Loop Filter (ALF) filter supports.

[0014] FIG. 4 is a conceptual diagram that illustrates an example of class index denoted by C_i based on matrix results (activity value Act and directionality D).

[0015] FIG. 5 is a conceptual diagram illustrating a 5x5 diamond-shaped filter support.

[0016] FIG. 6 is a conceptual diagram illustrating examples of geometry transformations.

[0017] FIG. 7 illustrates an example 9x9 diamond filter support for GALF.

[0018] FIG. 8 is a conceptual diagram illustrating a line buffer for ALF.

[0019] FIG. 9 is a conceptual diagram illustrating padded samples with motion compensation.

[0020] FIG. 10A is a conceptual diagram illustrating an example use of padding with mapping for ALF.

[0021] FIG. 10B is a conceptual diagram illustrating another example use of padding with mapping for ALF.

[0022] FIG. 11 is a conceptual diagram illustrating a modified filter support.

[0023] FIG. 12 is a block diagram illustrating an example video encoder that may implement one or more techniques described in this disclosure.

[0024] FIG. 13 is a block diagram illustrating an example video decoder that may implement one or more techniques described in this disclosure.

[0025] FIG. 14 is a flowchart illustrating an example coding method of the disclosure.

DETAILED DESCRIPTION

[0026] In general, this disclosure describes techniques related to adaptive loop filters (ALFs) and bilateral filters. As will be explained in more detail below, this disclosure describes techniques that may reduce the amount of memory needed for use as a line buffer when performing ALF, bilateral filtering or other types of filtering that may use samples from neighboring blocks.

[0027] FIG. 1 is a block diagram illustrating an example video encoding and decoding system 10 that may use techniques of this disclosure. As shown in FIG. 1, system 10 includes a source device 12 that provides encoded video data to be decoded at a later time by a destination device 14. In particular, source device 12 provides the encoded video data to destination device 14 via a computer-readable medium 16. Source device 12 and destination device 14 may comprise any of a wide range of devices, including desktop computers, notebook (i.e., laptop) computers, tablet computers, set-top boxes, telephone handsets such as so-called “smart” phones, tablet computers, televisions, cameras, display devices, digital media players, video gaming consoles, video streaming devices, or the like. In some cases, source device 12 and destination device 14 are equipped for wireless communication. Thus, source device 12 and destination device 14 may be wireless communication devices. The techniques described in this disclosure may be applied to wireless and/or wired applications. Source device 12 is an example video encoding device (i.e., a device for encoding video data). Destination device 14 is an example video decoding device (i.e., a device for decoding video data).

[0028] The illustrated system 10 of FIG. 1 is merely one example. Techniques for encoding, decoding, and processing video data may be performed by any digital video encoding and/or decoding device. In some examples, the techniques may be performed by a video encoder/decoder, typically referred to as a “CODEC.” Source device 12 and destination device 14 are examples of such coding devices in which source device 12 generates coded video data for transmission to destination device 14. In some

examples, source device 12 and destination device 14 operate in a substantially symmetrical manner such that each of source device 12 and destination device 14 include video encoding and decoding components. Hence, system 10 may support one-way or two-way video transmission between source device 12 and destination device 14, e.g., for video streaming, video playback, video broadcasting, or video telephony.

[0029] In the example of FIG. 1, source device 12 includes video source 18, storage media 19 configured to store video data, video encoder 20, and output interface 22. Destination device 14 includes input interface 26, storage media 28 configured to store encoded video data, video decoder 30, and display device 32. In other examples, source device 12 and destination device 14 include other components or arrangements. For example, source device 12 may receive video data from an external video source, such as an external camera. Likewise, destination device 14 may interface with an external display device, rather than including an integrated display device.

[0030] Video source 18 is a source of video data. The video data may comprise a series of pictures. Video source 18 may include a video capture device, such as a video camera, a video archive containing previously captured video, and/or a video feed interface to receive video data from a video content provider. In some examples, video source 18 generates computer graphics-based video data, or a combination of live video, archived video, and computer-generated video. Storage media 19 may be configured to store the video data. In each case, the captured, pre-captured, or computer-generated video may be encoded by video encoder 20.

[0031] Output interface 22 may output the encoded video information to a computer-readable medium 16. Output interface 22 may comprise various types of components or devices. For example, output interface 22 may comprise a wireless transmitter, a modem, a wired networking component (e.g., an Ethernet card), or another physical component. In examples where output interface 22 comprises a wireless transmitter, output interface 22 may be configured to transmit data, such as encoded video data, modulated according to a cellular communication standard, such as 4G, 4G-LTE, LTE Advanced, 5G, and the like. In some examples where output interface 22 comprises a wireless transmitter, output interface 22 may be configured to transmit data, such as encoded video data, modulated according to other wireless standards, such as an IEEE 802.11 specification, an IEEE 802.15 specification (e.g., ZigBee™), a Bluetooth™ standard, and the like. In some examples, circuitry of output interface 22 is integrated into circuitry of video encoder 20 and/or other components of source device 12. For

example, video encoder 20 and output interface 22 may be parts of a system on a chip (SoC). The SoC may also include other components, such as a general-purpose microprocessor, a graphics processing unit, and so on.

[0032] Destination device 14 may receive encoded video data to be decoded via computer-readable medium 16. Computer-readable medium 16 may comprise any type of medium or device capable of moving the encoded video data from source device 12 to destination device 14. In some examples, computer-readable medium 16 comprises a communication medium to enable source device 12 to transmit encoded video data directly to destination device 14 in real-time. The communication medium may comprise any wireless or wired communication medium, such as a radio frequency (RF) spectrum or one or more physical transmission lines. The communication medium may form part of a packet-based network, such as a local area network, a wide-area network, or a global network such as the Internet. The communication medium may include routers, switches, base stations, or any other equipment that may be useful to facilitate communication from source device 12 to destination device 14. Destination device 14 may comprise one or more data storage media configured to store encoded video data and decoded video data.

[0033] Computer-readable medium 16 may include transient media, such as a wireless broadcast or wired network transmission, or storage media (that is, non-transitory storage media), such as a hard disk, solid state drive, flash drive, compact disc, digital video disc, Blu-ray disc, or other computer-readable media. In some examples, a network server (not shown) may receive encoded video data from source device 12 and provide the encoded video data to destination device 14, e.g., via network transmission. Similarly, a computing device of a medium production facility, such as a disc stamping facility, may receive encoded video data from source device 12 and produce a disc containing the encoded video data. Therefore, computer-readable medium 16 may be understood to include one or more computer-readable media of any of various forms, in various examples.

[0034] In some examples, output interface 22 may output data, such as encoded video data, to an intermediate device, such as a storage device. Similarly, input interface 26 of destination device 12 may receive encoded data from the intermediate device. The intermediate device may include any of a variety of distributed or locally accessed data storage media such as a hard drive, solid state drive, Blu-ray discs, DVDs, CD-ROMs, flash memory, volatile or non-volatile memory, or any other suitable digital storage

media for storing encoded video data. In some examples, the intermediate device corresponds to a file server. Example file servers include web servers, FTP servers, network attached storage (NAS) devices, or local disk drives.

[0035] Destination device 14 may access the encoded video data through any standard data connection, including an Internet connection. This may include a wireless channel (e.g., a Wi-Fi connection), a wired connection (e.g., DSL, cable modem, optical fiber, etc.), or a combination of both that is suitable for accessing encoded video data stored on a file server. The transmission of encoded video data from the storage device may be a streaming transmission, a download transmission, or a combination thereof.

[0036] Input interface 26 of destination device 14 receives data from computer-readable medium 16. Input interface 26 may comprise various types of components or devices. For example, input interface 26 may comprise a wireless receiver, a modem, a wired networking component (e.g., an Ethernet card), an optical networking card, or another physical component. In examples where input interface 26 comprises a wireless receiver, input interface 26 may be configured to receive data, such as the bitstream, modulated according to a cellular communication standard, such as 4G, 4G-LTE, LTE Advanced, 5G, and the like. In some examples where input interface 26 comprises a wireless receiver, input interface 26 may be configured to receive data, such as the bitstream, modulated according to other wireless standards, such as an IEEE 802.11 specification, an IEEE 802.15 specification (e.g., ZigBee™), a Bluetooth™ standard, and the like. In some examples, circuitry of input interface 26 may be integrated into circuitry of video decoder 30 and/or other components of destination device 14. For example, video decoder 30 and input interface 26 may be parts of a SoC. The SoC may also include other components, such as a general-purpose microprocessor, a graphics processing unit, and so on.

[0037] Storage media 28 may be configured to store encoded video data, such as encoded video data (e.g., a bitstream) received by input interface 26. Display device 32 displays the decoded video data to a user. Display device 32 may comprise any of a variety of display devices such as a liquid crystal display (LCD), a plasma display, an organic light emitting diode (OLED) display, or another type of display device.

[0038] Video encoder 20 and video decoder unit 30 each may be implemented as any of a variety of suitable processing circuitry, such as one or more microprocessors, digital signal processors (DSPs), application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs), discrete logic, software, hardware, firmware or any

combinations thereof. When the techniques are implemented partially in software, a device may store instructions for the software in a suitable, non-transitory computer-readable medium and may execute the instructions in hardware using one or more processors to perform the techniques of this disclosure. Each of video encoder 20 and video decoder 30 may be included in one or more encoders or decoders, either of which may be integrated as part of a combined encoder/decoder (CODEC) in a respective device.

[0039] In some examples, video encoder 20 and video decoder 30 encode and decode video data according to one or more video coding standards or specifications. For example, video encoder 20 and video decoder 30 may encode and decode video data according to ITU-T H.261, ISO/IEC MPEG-1 Visual, ITU-T H.262 or ISO/IEC MPEG-2 Visual, ITU-T H.263, ISO/IEC MPEG-4 Visual and ITU-T H.264 (also known as ISO/IEC MPEG-4 AVC), including its Scalable Video Coding (SVC) and Multi-View Video Coding (MVC) extensions, or another video coding standard or specification. In some examples, video encoder 20 and video decoder 30 encode and decode video data according to the High Efficiency Video Coding (HEVC) standard, which is known as ITU-T H.265, its range and screen content coding extensions, its 3D video coding extension (3D-HEVC), its multiview extension (MV-HEVC), or its scalable extension (SHVC).

[0040] ITU-T VCEG (Q6/16) and ISO/IEC MPEG (JTC 1/SC 29/WG 11) are now studying the potential need for standardization of future video coding technology with a compression capability that exceeds that of the current HEVC standard (including its current extensions and near-term extensions for screen content coding and high-dynamic-range coding). The groups are working together on this exploration activity in a joint collaboration effort known as the Joint Video Exploration Team (JVET) to evaluate compression technology designs proposed by their experts in this area. The JVET first met during 19–21 October 2015. The latest version of reference software, i.e., Joint Exploration Model 7 (JEM7) can be downloaded from:

https://jvet.hhi.fraunhofer.de/svn/svn_HMJEMSoftware/tags/HM-16.6-JEM-7.0/. This algorithm description for JEM7 can be referred to as J. Chen, E. Alshina, G. J. Sullivan, J.-R. Ohm, J. Boyce "Algorithm description of Joint Exploration Test Model 7 (JEM7)", JVET-G1001, Torino, July 2017.

[0041] An early draft for a new video coding standard, referred to as the H.266/Versatile Video Coding (VVC) standard, is available in the document JVET-

J1001 “Versatile Video Coding (Draft 1)” by Benjamin Bross, and its algorithm description is available in the document JVET-J1002 “Algorithm description for Versatile Video Coding and Test Model 1 (VTM 1)” by Jianle Chen and Elena Alshina.

[0042] The techniques of this disclosure may be used in the context of advanced video codecs, such as extensions of HEVC or next generation video coding standards (e.g., H.266/VVC). While the techniques of this disclosure are generally described with reference to HEVC and next generation video coding standards (e.g., VVC/JEM), it should be understood that the techniques of this disclosure may be used in conjunction with any video coding techniques that use loop filters, including ALFs and deblocking filters.

[0043] As will be explained in more detail below, video encoder 20 and video decoder 30 may be configured to apply an adaptive loop filter (ALF) to decoded/reconstructed blocks of video data in a manner that reduces the amount of line buffer memory needed to store samples outside decoded/reconstructed blocks for use with particular filter supports of the ALF. For example, video encoder 20 and video decoder 30 may be configured to code (e.g., encode or decode) a block of video data to obtain a coded block of video data, apply an adaptive loop filter defined by a filter support to samples of the coded block of video data, including using padded samples for portions of the filter support that are outside of the block, and output the coded block of video data after applying adaptive loop filtering to the coded block of video data.

[0044] In HEVC and other video coding specifications, video data includes a series of pictures. Pictures may also be referred to as “frames.” A picture may include one or more sample arrays. Each respective sample array of a picture may comprise an array of samples for a respective color component. 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 chroma samples. S_{Cr} is a two-dimensional array of Cr chroma samples. In other instances, a picture may be monochrome and may only include an array of luma samples.

[0045] As part of encoding video data, video encoder 20 may encode pictures of the video data. In other words, video encoder 20 may generate encoded representations of the pictures of the video data. An encoded representation of a picture may be referred to herein as a “coded picture” or an “encoded picture.”

[0046] To generate an encoded representation of a picture, video encoder 20 may encode blocks of the picture. Video encoder 20 may include, in a bitstream, an encoded

representation of the video block. In some examples, to encode a block of the picture, video encoder 20 performs intra prediction or inter prediction to generate one or more predictive blocks. Additionally, video encoder 20 may generate residual data for the block. The residual block comprises residual samples. Each residual sample may indicate a difference between a sample of one of the generated predictive blocks and a corresponding sample of the block. Video encoder 20 may apply a transform to blocks of residual samples to generate transform coefficients. Furthermore, video encoder 20 may quantize the transform coefficients. In some examples, video encoder 20 may generate one or more syntax elements to represent a transform coefficient. Video encoder 20 may entropy encode one or more of the syntax elements representing the transform coefficient.

[0047] More specifically, when encoding video data according to HEVC or other video coding specifications, to generate an encoded representation of a picture, video encoder 20 may partition each sample array of the picture into coding tree blocks (CTBs) and encode the CTBs. A CTB may be an NxN block of samples in a sample array of a picture. In the HEVC main profile, the size of a CTB can range from 16x16 to 64x64, although technically 8x8 CTB sizes can be supported.

[0048] A coding tree unit (CTU) of a picture may comprise one or more CTBs and may comprise syntax structures used to encode the samples of the one or more CTBs. For instance, each a CTU may comprise a CTB of luma samples, two corresponding CTBs of chroma samples, and syntax structures used to encode the samples of the CTBs. In monochrome pictures or pictures having three separate color planes, a CTU may comprise a single CTB and syntax structures used to encode the samples of the CTB. A CTU may also be referred to as a “tree block” or a “largest coding unit” (LCU). In this disclosure, a “syntax structure” may be defined as one or more syntax elements present together in a bitstream in a specified order. In some codecs, an encoded picture is an encoded representation containing all CTUs of the picture.

[0049] To encode a CTU of a picture, video encoder 20 may partition the CTBs of the CTU into one or more coding blocks. A coding block is an NxN block of samples. In some codecs, to encode a CTU of a picture, video encoder 20 may recursively perform quad-tree partitioning on the coding tree blocks of a CTU to partition the CTBs into coding blocks, hence the name “coding tree units.” A coding unit (CU) may comprise one or more coding blocks and syntax structures used to encode samples of the one or more coding blocks. For example, a CU may comprise a coding block of luma samples

and two corresponding coding blocks of chroma samples of a picture that has a luma sample array, a Cb sample array, and a Cr sample array, and syntax structures used to encode the samples of the coding blocks. In monochrome pictures or pictures having three separate color planes, a CU may comprise a single coding block and syntax structures used to code the samples of the coding block.

[0050] Furthermore, video encoder 20 may encode CUs of a picture of the video data. In some codecs, as part of encoding a CU, video encoder 20 may partition a coding block of the CU into one or more prediction blocks. A prediction block is a rectangular (i.e., square or non-square) block of samples on which the same prediction is applied. A prediction unit (PU) of a CU may comprise one or more prediction blocks of a CU and syntax structures used to predict the one or more prediction blocks. For example, a PU may comprise a prediction block of luma samples, two corresponding prediction blocks of chroma samples, and syntax structures used to predict the prediction blocks. In monochrome pictures or pictures having three separate color planes, a PU may comprise a single prediction block and syntax structures used to predict the prediction block.

[0051] Video encoder 20 may generate a predictive block (e.g., a luma, Cb, and Cr predictive block) for a prediction block (e.g., luma, Cb, and Cr prediction block) of a PU of a CU. Video encoder 20 may use intra prediction or inter prediction to generate a predictive block. If video encoder 20 uses intra prediction to generate a predictive block, video encoder 20 may generate the predictive block based on decoded samples of the picture that includes the CU. If video encoder 20 uses inter prediction to generate a predictive block of a PU of a current picture, video encoder 20 may generate the predictive block of the PU based on decoded samples of a reference picture (i.e., a picture other than the current picture). In HEVC, video encoder 20 generates a “prediction_unit” syntax structure within a “coding_unit” syntax structure for inter predicted PUs, but does not generate a “prediction_unit” syntax structure within a “coding_unit” syntax structure for intra predicted PUs. Rather, in HEVC, syntax elements related to intra predicted PUs are included directly in the “coding_unit” syntax structure.

[0052] JEM/VVC also provides an affine motion compensation mode, which may be considered an inter-prediction mode. In affine motion compensation mode, video encoder 20 may determine two or more motion vectors that represent non-translational motion, such as zoom in or out, rotation, perspective motion, or other irregular motion types.

[0053] Video encoder 20 may generate one or more residual blocks for a CU. For instance, 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 Cb residual block of a CU 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 decompose the residual blocks of a CU into one or more transform blocks. For instance, video encoder 20 may use quad-tree partitioning to decompose the residual blocks of a CU into one or more transform blocks. A transform block is a rectangular (e.g., square or non-square) block of samples on which the same transform is applied. A transform unit (TU) of a CU may comprise one or more transform blocks. For example, a TU may comprise a transform block of luma samples, two corresponding transform blocks of chroma samples, and syntax structures used to transform the transform block samples. Thus, each TU of a CU may have a luma transform block, a Cb transform block, and a Cr transform block. The luma transform block of the TU may be a sub-block of the CU's luma residual block. The Cb transform block may be a sub-block of the CU's Cb residual block. The Cr transform block may be a sub-block of the CU's Cr residual block. In monochrome pictures or pictures having three separate color planes, a TU may comprise a single transform block and syntax structures used to transform the samples of the transform block.

[0055] In JEM7, rather than using the quadtree partitioning structure of HEVC described above, a quadtree binary tree (QTBT) partitioning structure may be used. The QTBT structure removes the concepts of multiple partition types. That is, the QTBT structure removes the separation of the CU, PU, and TU concepts, and supports more flexibility for CU partition shapes. In the QTBT block structure, a CU can have either a square or rectangular shape. In one example, a CU is first partitioned by a quadtree structure. The quadtree leaf nodes are further partitioned by a binary tree structure.

[0056] In some examples, there are two splitting types: symmetric horizontal splitting and symmetric vertical splitting. The binary tree leaf nodes are called CUs, and that

segmentation (i.e., the CU) is used for prediction and transform processing without any further partitioning. This means that the CU, PU, and TU have the same block size in the QTBT coding block structure. In JEM, a CU sometimes consists of coding blocks (CBs) of different color components. For example, one CU contains one luma CB and two chroma CBs in the case of P and B slices of the 4:2:0 chroma format and sometimes consists of a CB of a single component. For example, one CU contains only one luma CB or just two chroma CBs in the case of I slices.

[0057] FIGS. 2A and 2B are conceptual diagrams illustrating an example quadtree binary tree (QTBT) structure 130, and a corresponding coding tree unit (CTU) 132. The solid lines represent quadtree splitting, and dotted lines indicate binary tree splitting. In each split (i.e., non-leaf) node of the binary tree, one flag is signaled to indicate which splitting type (i.e., horizontal or vertical) is used, where 0 indicates horizontal splitting and 1 indicates vertical splitting in this example. For the quadtree splitting, there is no need to indicate the splitting type, since quadtree nodes split a block horizontally and vertically into 4 sub-blocks with equal size. Accordingly, video encoder 20 may encode, and video decoder 30 may decode, syntax elements (such as splitting information) for a region tree level of QTBT structure 130 (i.e., the solid lines) and syntax elements (such as splitting information) for a prediction tree level of QTBT structure 130 (i.e., the dashed lines). Video encoder 20 may encode, and video decoder 30 may decode, video data, such as prediction and transform data, for CUs represented by terminal leaf nodes of QTBT structure 130.

[0058] In general, CTU 132 of FIG. 2B may be associated with parameters defining sizes of blocks corresponding to nodes of QTBT structure 130 at the first and second levels. These parameters may include a CTU size (representing a size of CTU 132 in samples), a minimum quadtree size (MinQTSIZE, representing a minimum allowed quadtree leaf node size), a maximum binary tree size (MaxBTSize, representing a maximum allowed binary tree root node size), a maximum binary tree depth (MaxBTDepth, representing a maximum allowed binary tree depth), and a minimum binary tree size (MinBTSize, representing the minimum allowed binary tree leaf node size).

[0059] The root node of a QTBT structure corresponding to a CTU may have four child nodes at the first level of the QTBT structure, each of which may be partitioned according to quadtree partitioning. That is, nodes of the first level are either leaf nodes (having no child nodes) or have four child nodes. The example of QTBT structure 130

represents such nodes as including the parent node and child nodes having solid lines for branches. If nodes of the first level are not larger than the maximum allowed binary tree root node size (MaxBTSize), they can be further partitioned by respective binary trees. The binary tree splitting of one node can be iterated until the nodes resulting from the split reach the minimum allowed binary tree leaf node size (MinBTSize) or the maximum allowed binary tree depth (MaxBTDepth). The example of QTBT structure 130 represents such nodes as having dashed lines for branches. The binary tree leaf node is referred to as a coding unit (CU), which is used for prediction (e.g., intra-picture or inter-picture prediction) and transform, without any further partitioning. As discussed above, CUs may also be referred to as “video blocks” or “blocks.”

[0060] In one example of the QTBT partitioning structure, the CTU size is set as 128x128 (luma samples and two corresponding 64x64 chroma samples), the MinQTSIZE is set as 16x16, the MaxBTSize is set as 64x64, the MinBTSize (for both width and height) is set as 4, and the MaxBTDepth is set as 4. The quadtree partitioning is applied to the CTU first to generate quad-tree leaf nodes. The quadtree leaf nodes may have a size from 16x16 (i.e., the MinQTSIZE) to 128x128 (i.e., the CTU size). If the leaf quadtree node is 128x128, it will not be further split by the binary tree, since the size exceeds the MaxBTSize (i.e., 64x64, in this example). Otherwise, the leaf quadtree node will be further partitioned by the binary tree. Therefore, the quadtree leaf node is also the root node for the binary tree and has the binary tree depth as 0. When the binary tree depth reaches MaxBTDepth (4, in this example), no further splitting is permitted. When the binary tree node has width equal to MinBTSize (4, in this example), it implies no further horizontal splitting is permitted. Similarly, a binary tree node having a height equal to MinBTSize implies no further vertical splitting is permitted for that binary tree node. As noted above, leaf nodes of the binary tree are referred to as CUs, and are further processed according to prediction and transform without further partitioning.

[0061] Video encoder 20 may apply one or more transforms to a transform block of a TU to generate a coefficient block for the TU. A coefficient block may be a two-dimensional array of transform coefficients. A transform coefficient may be a scalar quantity. In some examples, the one or more transforms convert the transform block from a pixel domain to a frequency domain. Thus, in such examples, a transform coefficient may be a scalar quantity considered to be in a frequency domain. A transform coefficient level is an integer quantity representing a value associated with a

particular 2-dimensional frequency index in a decoding process prior to scaling for computation of a transform coefficient value.

[0062] In some examples, video encoder 20 skips application of the transforms to the transform block. In such examples, video encoder 20 may treat residual sample values in the same way as transform coefficients. Thus, in examples where video encoder 20 skips application of the transforms, the following discussion of transform coefficients and coefficient blocks may be applicable to transform blocks of residual samples.

[0063] After generating a coefficient block, video encoder 20 may quantize the coefficient block to possibly reduce the amount of data used to represent the coefficient block, potentially providing further compression. Quantization generally refers to a process in which a range of values is compressed to a single value. For example, quantization may be done by dividing a value by a constant, and then rounding to the nearest integer. To quantize the coefficient block, video encoder 20 may quantize transform coefficients of the coefficient block. In some examples, video encoder 20 skips quantization.

[0064] Video encoder 20 may generate syntax elements indicating some or all the potentially quantized transform coefficients. Video encoder 20 may entropy encode one or more of the syntax elements indicating a quantized transform coefficient. For example, video encoder 20 may perform Context-Adaptive Binary Arithmetic Coding (CABAC) on the syntax elements indicating the quantized transform coefficients. Thus, an encoded block (e.g., an encoded CU) may include the entropy encoded syntax elements indicating the quantized transform coefficients.

[0065] Video encoder 20 may output a bitstream that includes encoded video data. In other words, video encoder 20 may output a bitstream that includes an encoded representation of video data. The encoded representation of the video data may include an encoded representation of pictures of the video data. For example, the bitstream may comprise a sequence of bits that forms a representation of encoded pictures of the video data and associated data. In some examples, a representation of an encoded picture may include encoded representations of blocks of the picture.

[0066] Video decoder 30 may receive a bitstream generated by video encoder 20. As noted above, the bitstream may comprise an encoded representation of video data. Video decoder 30 may decode the bitstream to reconstruct pictures of the video data. As part of decoding the bitstream, video decoder 30 may obtain syntax elements from the bitstream. Video decoder 30 may reconstruct pictures of the video data based at

least in part on the syntax elements obtained from the bitstream. The process to reconstruct pictures of the video data may be generally reciprocal to the process performed by video encoder 20 to encode the pictures.

[0067] For instance, as part of decoding a picture of the video data, video decoder 30 may use inter prediction or intra prediction to generate predictive blocks. Additionally, video decoder 30 may determine transform coefficients based on syntax elements obtained from the bitstream. In some examples, video decoder 30 inverse quantizes the determined transform coefficients. Furthermore, video decoder 30 may apply an inverse transform on the determined transform coefficients to determine values of residual samples. Video decoder 30 may reconstruct a block of the picture based on the residual samples and corresponding samples of the generated predictive blocks. For instance, video decoder 30 may add residual samples to corresponding samples of the generated predictive blocks to determine reconstructed samples of the block.

[0068] More specifically, in HEVC and other video coding specifications, video decoder 30 may use inter prediction or intra prediction to generate one or more predictive blocks for each PU of a current CU. In addition, video decoder 30 may inverse quantize coefficient blocks of TUs of the current CU. Video decoder 30 may perform inverse transforms on the coefficient blocks to reconstruct transform blocks of the TUs of the current CU. Video decoder 30 may reconstruct a coding block of the current CU based on samples of the predictive blocks of the PUs of the current CU and residual samples of the transform blocks of the TUs of the current CU. In some examples, video decoder 30 may reconstruct the coding blocks of the current CU by adding the samples of the predictive blocks for PUs of the current CU to corresponding decoded 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.

[0069] A slice of a picture may include an integer number of blocks of the picture. For example, in HEVC and other video coding specifications, a slice of a picture may include an integer number of CTUs of the picture. The CTUs of a slice may be ordered consecutively in a scan order, such as a raster scan order. In HEVC, a slice is defined as an integer number of CTUs contained in one independent slice segment and all subsequent dependent slice segments (if any) that precede the next independent slice segment (if any) within the same access unit. Furthermore, in HEVC, a slice segment is defined as an integer number of CTUs ordered consecutively in the tile scan and

contained in a single NAL unit. A tile scan is a specific sequential ordering of CTBs partitioning a picture in which the CTBs are ordered consecutively in CTB raster scan in a tile, whereas tiles in a picture are ordered consecutively in a raster scan of the tiles of the picture. A tile is a rectangular region of CTBs within a particular tile column and a particular tile row in a picture.

[0070] In the field of video coding, it is common to apply filtering in order to enhance the quality of a decoded video signal. Filtering may also be applied in the reconstruction loop of video encoder 20 and video decoder 30. A filter may be applied as a post-filter, where the filtered frame is not used for prediction of future frames, or a filter may be applied as an in-loop filter, where the filtered frame is used to predict future frames. A filter can be designed, for example, by minimizing the error between the original signal and the decoded filtered signal. Similarly to transform coefficients, video encoder 20 may quantize code the coefficients of the filter $h(k, l)$, $k = -K, \dots, K, l = -K, \dots, K$ by:

$$f(k, l) = \text{round}(\text{normFactor} \cdot h(k, l))$$

Video encoder 20 may send the coded quantized coefficients to video decoder 30. The *normFactor* may be equal to 2^n . The larger the value of *normFactor*, the more precise is quantization and the quantized filter coefficients $f(k, l)$ provide better performance. However, larger values of *normFactor* produce coefficients $f(k, l)$ requiring more bits to transmit.

[0071] In video decoder 30, the decoded filter coefficients $f(k, l)$ are applied to the reconstructed image $R(i, j)$ as follows

$$\tilde{R}(i, j) = \sum_{k=-K}^K \sum_{l=-K}^K f(k, l) R(i + k, j + l) / \sum_{k=-K}^K \sum_{l=-K}^K f(k, l), \quad (1)$$

where i and j are the coordinates of the pixels within the frame.

[0072] The in-loop adaptive loop filter (ALF) was evaluated during the HEVC development stage, but ALF was not included in the final version of the HEVC standard.

[0073] One example in-loop ALF employed in JEM was proposed in J. Chen et al., “Coding tools investigation for next generation video coding,” SG16-Geneva-C806, Jan. 2015. The basic idea is the same as the ALF with block-based adaptation in HM-3.

(See T. Wiegand et al., “WD3: Working Draft 3 of High-Efficiency Video Coding,” Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T SG16 WP3 and ISO/IEC JTC1/SC29/WG11, JCTVC-E603, 5th Meeting: Geneva, CH, 16-23 March 2011, hereinafter, “JCTVC-E603”).

[0074] For the luma component, video encoder 20 and video decoder 30 may classify 4x4 blocks in the picture based on a 1D Laplacian direction (up to 3 directions) and a 2D Laplacian activity (up to 5 activity values). The calculation of direction Dir_b and unquantized activity Act_b is shown in equations (2) through (5), where $\hat{l}_{i,j}$ indicates a reconstructed pixel with relative coordinate (i, j) to the top-left of a 4x4 block. Act_b is further quantized to the range of 0 to 4 inclusively as described in T. Wiegand et al., “WD3: Working Draft 3 of High-Efficiency Video Coding,” Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T SG16 WP3 and ISO/IEC JTC1/SC29/WG11, JCTVC-E603, 5th Meeting: Geneva, CH, 16-23 March 2011.

$$V_{i,j} = |\hat{l}_{i,j} \times 2 - \hat{l}_{i,j-1} - \hat{l}_{i,j+1}| \quad (2)$$

$$H_{i,j} = |\hat{l}_{i,j} \times 2 - \hat{l}_{i-1,j} - \hat{l}_{i+1,j}| \quad (3)$$

$$Dir_b = \begin{cases} 1, & \text{if } (\sum_{i=0}^3 \sum_{j=0}^3 H_{i,j} > 2 \times \sum_{i=0}^3 \sum_{j=0}^3 V_{i,j}) \\ 2, & \text{if } (\sum_{i=0}^3 \sum_{j=0}^3 V_{i,j} > 2 \times \sum_{i=0}^3 \sum_{j=0}^3 H_{i,j}) \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

$$Act_b = \sum_{i=0}^3 \sum_{j=0}^3 (\sum_{m=i-1}^{i+1} \sum_{n=j-1}^{j+1} (V_{m,n} + H_{m,n})) \quad (5)$$

[0075] In total, video encoder 20 and video decoder 30 may be configured to categorize each block into one out of 15 (5x3) groups and an index is assigned to each 4x4 block according to the value of Dir_b and Act_b of the block. Denote the group index by C and the categorization is set equal to $5Dir_b + \hat{A}$, wherein $\hat{A}$ is the quantized value of Act_b .

[0076] The quantization process to obtain activities value Act_b to activity index $\hat{A}$ may be performed as follows. Basically, this process defines the rule regarding how to merge blocks with different activities to one class if Dir_b is the same. The quantization process of Act_b is defined as follows:

$$\text{avg_var} = \text{Clip_post}((\text{NUM_ENTRY}-1), (Act_b * \text{ScaleFactor}) >> \text{shift});$$

$$\hat{A} = \text{ActivityToIndex}[\text{avg_var}]$$

wherein NUM_ENTRY is set to 16, ScaleFactor is set to 114, shift is equal to (3 + internal coded bit-depth), ActivityToIndex[NUM_ENTRY] = {0, 1, 2, 2, 2, 3, 3, 3, 3, 3,

4, 4, 4, 4, 4, 4}}, and function Clip_post (a, b) returns the smaller value between a and b.

[0077] Therefore, up to 15 sets of ALF parameters could be signalled for the luma component of a picture. To save the signaling cost, video encoder 20 may be configured to merge the groups along a group index value. For each merged group, a set of ALF coefficients is signaled. In one example of JEM2.0, up to three circular symmetric filter shapes (e.g., as shown in FIG. 3) are supported. FIG. 3 shows a 5x5 diamond filter support 300, a 7x7 diamond filter support 302, and a truncated 9x9 diamond filter support 304. In one example, for both chroma components in a picture, a single set of ALF coefficients is applied and the 5x5 diamond shape filter is always used. The symbol C_i (e.g., C_0 , C_1 , C_2 , etc.) in FIG. 3 represents filter coefficients.

[0078] In some examples, at video decoder 30, each pixel sample $\hat{I}_{i,j}$ is filtered, resulting in pixel value $I'_{i,j}$ as shown in equation (6), where L denotes filter length, $f_{m,n}$ represents filter coefficient and o indicates filter offset,

$$I'_{i,j} = \sum_{m=-L}^L \sum_{n=-L}^L f_{m,n} \times \hat{I}_{i+m,j+n} + o \quad (6)$$

wherein $(1 \ll (BD_F - 1)) = \sum_{m=-L}^L \sum_{n=-L}^L f(m,n)$ and $(o = (1 \ll (BD_F - 2)))$. In JEM2, the bit-depth, denoted by BD_F is set to 9, which means the filter coefficient could be in the range of $[-256, 256]$. Note that for some examples, only up to one filter is supported for two chroma components.

[0079] The temporal prediction of filter coefficients will now be discussed. Video encoder 20 and/or video decoder 30 may be configured to store the ALF coefficients of previously coded pictures and may be configured to reuse such coefficients as ALF coefficients of a current picture. For the current picture, video encoder 20 and/or video decoder 30 may be configured to choose to use ALF coefficients stored for the reference pictures and bypass the ALF coefficients signalling. In this case, only an index to one of the reference pictures is signalled, and the stored ALF coefficients of the indicated reference picture are simply inherited (e.g., reused) for the current picture. To indicate the usage of temporal prediction, video encoder 20 may be configured to first code one flag before sending the index.

[0080] Geometry transformations-based ALF will now be discussed. In M. Karczewicz, L. Zhang, W.-J. Chien, X. Li, "EE2.5: Improvements on adaptive loop filter", Exploration Team (JVET) of ITU-T SG 16 WP 3 and ISO/IEC JTC 1/SC 29/WG 11, Doc. JVET-B0060, 2nd Meeting: San Diego, USA, 20 Feb. – 26 Feb. 2016, and in

M. Karczewicz, L. Zhang, W.-J. Chien, X. Li, “EE2.5: Improvements on adaptive loop filter”, Exploration Team (JVET) of ITU-T SG 16 WP 3 and ISO/IEC JTC 1/SC 29/WG 11, Doc. JVET-C0038, 3rd Meeting: Geneva, CH, 26 May – 1 June 2016, the Geometric transformations-based ALF (GALF) is proposed. GALF was adopted to JEM3.0. In GALF, the classification is modified with the diagonal gradients taken into consideration and geometric transformations could be applied to filter coefficients. Each 2×2 block is categorized into one out of 25 classes based on its directionality and quantized value of activity. In other examples, video encoder 20 and/or video decoder 30 may be configured to perform classification on 4x4 subblocks. The details are described in the following sub-sections.

[0081] Classifications in GALF are discussed in this section. Similar to the design of example ALF implementations, the classification for GALF is based on the 1D Laplacian direction and 2D Laplacian activity of each N×N luma block. However, the definitions of both direction and activity have been modified to better capture local characteristics. First, values of two diagonal gradients, in addition to the horizontal and vertical gradients used in the existing ALF, are calculated using a 1-D Laplacian. As can be seen from equations (7) to (10) below, the sum of gradients of all pixels within a 6×6 window that covers a target pixel is employed as the represented gradient of the target pixel. In other examples, video encoder 20 and/or video decoder 30 may be configured to use an 8x8 window that covers a target pixel. According to experiments, the window size, i.e., 6×6, provides a good trade-off between complexity and coding performance. Each pixel is associated with four gradient values, with vertical gradient denoted by g_v , horizontal gradient denoted by g_h , 135-degree diagonal gradient denoted by g_{d1} and 45-degree diagonal gradient denoted by g_{d2} .

[0082]

$$g_v = \sum_{k=i-2}^{i+3} \sum_{l=j-2}^{j+3} V_{k,l}, \quad (7)$$

$$V_{k,l} = |2R(k, l) - R(k, l - 1) - R(k, l + 1)|$$

$$g_h = \sum_{k=i-2}^{i+3} \sum_{l=j-2}^{j+3} H_{k,l}, \quad (8)$$

$$H_{k,l} = |2R(k, l) - R(k - 1, l) - R(k + 1, l)|$$

$$g_{d1} = \sum_{k=i-2}^{i+3} \sum_{l=j-3}^{j+3} D1_{k,l}, \quad (9)$$

$$D1_{k,l} = |2R(k, l) - R(k-1, l-1) - R(k+1, l+1)|$$

$$g_{d2} = \sum_{k=i-2}^{i+3} \sum_{j=j-2}^{j+3} D2_{k,l}, \quad (10)$$

$$D2_{k,l} = |2R(k, l) - R(k-1, l+1) - R(k+1, l-1)|$$

Here, indices i and j refer to the coordinates of the upper left pixel in the 2×2 block.

[0083] Table 1. Values of Direction and Its Physical Meaning

Direction values	physical meaning
0	Texture
1	Strong horizontal/vertical
2	horizontal/vertical
3	strong diagonal
4	diagonal

[0084] To assign the directionality D , a ratio of the maximum and the minimum of the horizontal and vertical gradients, denoted by $R_{h,v}$ in (10), and the ratio of the maximum and the minimum of two diagonal gradients, denoted by $R_{d1,d2}$ in (11), are compared against each other with two thresholds, t_1 and t_2 .

$$R_{h,v} = g_{h,v}^{max} / g_{h,v}^{min} \quad (11)$$

$$\text{wherein } g_{h,v}^{max} = \max(g_h, g_v), g_{h,v}^{min} = \min(g_h, g_v),$$

$$R_{d0,d1} = g_{d0,d1}^{max} / g_{d0,d1}^{min} \quad (12)$$

$$\text{wherein } g_{d0,d1}^{max} = \max(g_{d0}, g_{d1}), g_{d0,d1}^{min} = \min(g_{d0}, g_{d1})$$

[0085] By comparing the detected ratios of horizontal/vertical and diagonal gradients, five direction modes, i.e., D within the range of $[0, 4]$ inclusive, are defined in (12) and (13). The values of D and its physical meaning are described in Table I.

$$D = \begin{cases} 0 & R_{h,v} \leq t_1 \ \&\& R_{d0,d1} \leq t_1 \\ 1 & R_{h,v} > t_1 \ \&\& R_{h,v} > R_{d0,d1} \ \&\& R_{h,v} > t_2 \\ 2 & R_{h,v} > t_1 \ \&\& R_{h,v} > R_{d0,d1} \ \&\& R_{h,v} \leq t_2 \\ 3 & R_{d0,d1} > t_1 \ \&\& R_{h,v} \leq R_{d0,d1} \ \&\& R_{d0,d1} > t_2 \\ 4 & R_{d0,d1} > t_1 \ \&\& R_{h,v} \leq R_{d0,d1} \ \&\& R_{d0,d1} \leq t_2 \end{cases} \quad (13)$$

[0086] The activity value Act is calculated as:

$$Act = \sum_{k=i-2}^{i+3} \sum_{l=j-2}^{j+3} (V_{k,l} + H_{k,l}). \quad (14)$$

[0087] Act is further quantized to the range of 0 to 4 inclusive, and the quantized value is denoted as $\hat{A}$.

[0088] Quantization process from activity value A to activity index $\hat{A}$

[0089] The quantization process is defined as follows:

$$\begin{aligned} \text{avg_var} &= \text{Clip_post}(\text{NUM_ENTRY}-1, (Act * \text{ScaleFactor}) \gg \text{shift}); \\ \hat{A} &= \text{ActivityToIndex}[\text{avg_var}] \end{aligned}$$

wherein NUM_ENTRY is set to 16, ScaleFactor is set to 24, shift is (3 + internal coded-bitdepth), ActivityToIndex[NUM_ENTRY] = {0, 1, 2, 2, 2, 2, 2, 3, 3, 3, 3, 3, 3, 3, 4}, and function Clip_post(a, b) returns the smaller value between a and b.

[0090] Please note that due to different ways of calculating the activity value, the ScaleFactor and ActivityToIndex are both modified compared to the ALF design in JEM2.0.

[0091] Therefore, in the proposed GALF scheme, each $N \times N$ block is categorized into one of 25 classes based on its directionality D and quantized value of activity $\hat{A}$:

$$C = 5D + \hat{A}. \quad (15)$$

[0092] An example of class index according to D and quantized value of activity $\hat{A}$ is depicted in FIG. 4. Note that the value of $\hat{A}$ is set to 0 ... 4 for each column which is derived from the variable Act . The smallest Act for a new value of $\hat{A}$ is marked in along the top line (e.g., 0, 8192, 16384, 57344, 122880). For example, Act with values within [16384, 57344-1] will fall in $\hat{A}$ equal to 2.

[0093] Geometry transformations will now be discussed. For each category, one set of filter coefficients may be signalled. To better distinguish different directions of blocks marked with the same category index, four geometry transformations, including no transformation, diagonal, vertical flip and rotation, are introduced. FIG. 5 is a conceptual diagram illustrating a 5×5 diamond-shaped filter support 500. An example of a 5×5 filter support with the three geometric transformations is depicted in FIG. 6. FIG. 6 shows a diagonal 5×5 diamond-shaped filter support 600, a vertical flip 5×5 diamond-shaped filter support 602, and a rotation 5×5 diamond-shaped filter support 604. The symbol C_i (e.g., C_0 , C_1 , C_2 , etc.) in FIG. 5 and FIG. 6 represents filter

coefficients. Comparing FIG. 5 and FIG. 6, the formula forms of the three additional geometry transformations are:

$$\begin{aligned} \text{Diagonal: } f_D(k, l) &= f(l, k), \\ \text{Vertical flip: } f_V(k, l) &= f(k, K - l - 1), \\ \text{Rotation: } f_R(k, l) &= f(K - l - 1, k). \end{aligned} \tag{16}$$

where K is the size of the filter and $0 \leq k, l \leq K - 1$ are coefficients coordinates, such that location $(0,0)$ is at the upper left corner and location $(K - 1, K - 1)$ is at the lower right corner.

[0094] Note that when the diamond filter support is used, such as in the existing examples of ALF, the coefficients with coordinates out of the filter support will be always set to 0. One technique for indicating the geometry transformation index is to derive geometry transformation index implicitly to avoid additional overhead. In GALF, the transformations are applied to the filter coefficients $f(k, l)$ depending on gradient values calculated for that block. The relationship between the transformation and the four gradients calculated using (6)-(9) is described in Table 2. To summarize, the transformations are based on which one of two gradients (horizontal and vertical, or 45 degree and 135 degree gradients) is larger. Based on the comparison, more accurate direction information can be extracted. Therefore, different filtering results could be obtained due to transformation while the overhead of filter coefficients is not increased.

[0095] TABLE 2. MAPPING OF GRADIENT AND TRANSFORMATIONS.

Gradient values	Transformation
$g_{d2} < g_{d1}$ and $g_h < g_v$	No transformation
$g_{d2} < g_{d1}$ and $g_v < g_h$	Diagonal
$g_{d1} < g_{d2}$ and $g_h < g_v$	Vertical flip
$g_{d1} < g_{d2}$ and $g_v < g_h$	Rotation

[0096] Similar to the ALF in HM3, GALF also adopts the 5x5 and 7x7 diamond filter supports. In addition, the 9x7 filter support (truncated 9x9 diamond of FIG. 3) is replaced by the 9x9 diamond filter support. FIG. 7 illustrates an example 9x9 diamond filter support 700 for GALF. In some examples of VVC, a 7x7 filter support is used instead of a 9x9 diamond filter support. The techniques of this disclosure are applicable for any size of filter support.

[0097] Prediction from fixed filters will now be discussed. To improve coding efficiency when temporal prediction is not available (e.g., for intra frames), a set of fixed filters (e.g., 16 fixed filters) may be assigned to each class. To indicate the usage of the fixed filter, a flag for each class is signaled. In some examples, the index of the chosen fixed filter may also be signaled. Even when the fixed filter is selected for a given class, the coefficients of the adaptive filter $f(k, l)$ can still be sent for this class in which case the coefficients of the filter which will be applied to the reconstructed image are a sum of both sets of coefficients. A number of classes can share the same coefficients $f(k, l)$ signaled in the bitstream even if different fixed filters were chosen for them. U.S. Patent Application No. 15/432,839, filed February 14, 2017 (U.S. Publication No. 2017/0238020, published August 17, 2017), describes that the fixed filters could also be applied to inter-coded frames.

[0098] Signalling of filter coefficients will now be discussed, including a prediction pattern and prediction index from fixed filters. Three cases are defined: case 1: whether none of the filters of the classes (e.g., 25 classes) are predicted from the fixed filter; case 2: all filters of the classes are predicted from the fixed filter; and case 3: filters associated with some classes are predicted from fixed filters and filters associated with the rest of the classes are not predicted from the fixed filters.

[0099] Video encoder 20 may first code an index to indicate one of the three cases. In addition, the following applies:

- If it is case 1, there is no need to further signal the index of fixed filter.
- Otherwise, if it is case 2, an index of the selected fixed filter for each class is signaled
- Otherwise (if it is case 3), one bit for each class is firstly signaled, and if a fixed filter is used, the index is further signaled.

[0100] Skipping of DC filter coefficient

[0101] Since the sum of all filter coefficients have to be equal to 2^K (wherein K denotes the bit-depth of filter coefficient), the DC filter coefficient which is applied to a current pixel (center pixel within a filter support, such as C_6 in FIG. 5) could be derived without signaling.

[0102] Filter index

[0103] To reduce the number of bits required to represent the filter coefficients, different classes can be merged. However, unlike in T. Wiegand, B. Bross, W.-J. Han, J.-R. Ohm and G. J. Sullivan, “WD3: Working Draft 3 of High-Efficiency Video Coding,” Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T SG16 WP3 and ISO/IEC JTC1/SC29/WG11, JCTVC-E603, 5th Meeting: Geneva, CH, 16-23 March 2011, any set of classes can be merged, even classes having non-consecutive values of C which denotes the class index as defined in (15). The information indicating which classes are merged is provided by sending for each of the 25 classes an index i_C . Classes having the same index i_C share the same filter coefficients that are coded. The index i_C is coded with truncated binary binarization method. Other information, such as coefficients, are coded in the same way as in JEM2.0.

[0104] Existing ALF/GALF designs exhibit several drawbacks. For example, some existing ALF designs use an additional eight lines of line buffer memory, as depicted by the gray samples 800 in FIG. 8. For samples located at the top-line of LCUs (like the top-left sample marked by X) in FIG. 8, the above four lines of samples of a neighboring LCU are used for the 9x9 diamond filter support. Similarly, for filtering samples located at the bottom line of LCUs (like the bottom-left sample marked by X) in FIG. 8, the below four lines of samples of a neighboring LCU are used for the 9x9 diamond filter support. Video encoder 20 and video decoder 30 may be configured with eight line buffers to store these eight lines of samples of the neighboring LCUs. Other

types of filters (such as deblocking filter and sample adaptive offset (SAO)) may use fewer lines of neighboring samples, and thus fewer and/or smaller line buffers. Some example deblocking filters use four lines of line buffer memory, while some SAO filters only use one line of line buffer memory. Compared to deblocking filters and SAO filters, ALF uses more line buffers, which increases cost and chip size.

[0105] In view of these drawbacks, this disclosure describes the following techniques for performing ALF with a reduced number of line buffers. Both video encoder 20 and video decoder 30 may be configured to perform the techniques described below. The following itemized techniques may be applied individually. Alternatively, any combination of the techniques described below may be applied together.

[0106] In one example of the disclosure, for samples located at certain positions of a block (e.g., largest coding unit, LCU), video encoder 20 and video decoder 30 may be configured to disable GALF/ALF (i.e., GALF and/or ALF) without any signaling. For example, video decoder 30 may be configured to disable GALF/ALF for certain samples of a block based on their position on the block, and not based on any signalled syntax elements that indicate whether or not ALF/GALF are disabled. In this way, a smaller line buffer (or no line buffer) may be needed, since samples from neighboring blocks will not be needed since ALF/GALF is disabled.

[0107] In one example, for samples located at the first N (non-zero variable) rows within one LCU, video encoder 20 and video decoder 30 may be configured to disable ALF/GALF. For example, video encoder 20 and video decoder 30 may be configured to mark the samples located within the first N rows as 'inapplicable' for ALF/GALF.

[0108] In another example, for samples located at the last M (non-zero variable) rows within one LCU, video encoder 20 and video decoder 30 may be configured to disable ALF/GALF. For example, video encoder 20 and video decoder 30 may be configured to mark the samples located within the last M rows as 'inapplicable' for ALF/GALF.

[0109] In one example, the variable N or/and M may be pre-defined (e.g., stored at both video encoder 20 and video decoder 30). In another example, video encoder 20 may be configured to signal syntax elements indicating the values of N and/or M in an encoded video bitstream for use by video decoder 30.

[0110] In other examples, the variables N or/and M may be different for different sequences/pictures/slices. That is, the variables N and M need not be same for all sequences/pictures/slices.

[0111] In another example, the variable N or/and M may be dependent on the allowed filter support. For example, if video encoder 20 and/or video decoder 30 determines to use the 5x5 diamond filter support for one slice, video encoder 20 and/or video decoder 30 may set the variables N and M to 0. Similarly, if video encoder 20 and/or video decoder 30 determines to use the 9x9 diamond filter support for one slice, video encoder 20 and/or video decoder 30 may set the variables N and M to 2.

[0112] In another example, for blocks with all samples (or a majority of samples) marked as 'inapplicable,' video encoder 20 and/or video decoder 30 may be configured to skip the signaling of GALF/ALF on/off flags. Instead, video encoder 20 and/or video decoder 30 may be configured to infer the GALF/ALF on/off flags to be false (i.e., ALF/GALF is turned off).

[0113] In another example of the disclosure, instead of accessing samples outside the LCU rows, thus using one or more line buffers, video encoder 20 and/or video decoder 30 may be configured to use padded samples for the filtering process. Padded samples are samples that are generated by video encoder 20 and/or video decoder 30, as opposed to samples from neighboring blocks stored in line buffers. In this way, the amount of line buffers may be reduced, thus decreasing the complexity and cost of hardware implementations.

[0114] In one example, video encoder 20 and/or video decoder 30 may be configured to generate the padded samples from blocks in one or more reference pictures. An example is given in FIG. 9. For the current sample 900 to be filtered (located at the top row of one LCU 902 in frame N-1), the neighboring samples (marked with A-C, 1-9) may be padded from the motion compensated block in a different reference picture 904 (frame N-2). That is, samples A-C and 1-9 from picture 904 may be copied and used as samples for filtering current sample 900 of current LCU 902.

[0115] In another example, video encoder 20 and/or video decoder 30 may be configured to copy the padded samples from block(s) in one or more reference pictures. When the padded samples are from multiple reference pictures, video encoder 20 and/or video decoder 30 may be configured to apply a linear function to multiple samples from the same position to generate one padded sample. In some example, the block(s) may be identified by motion vectors associated with a current block covering the sample to be filtered, or associated with the neighboring block covering the sample to be padded. In some examples, the motion vector(s) may be rounded to integer positions. In other examples, the motion vectors may be set to zero motion vectors.

[0116] In another example, instead of generating padded samples by duplicating existing samples from the first row or the column of a block, video encoder 20 and/or video decoder 30 may be configured to pad samples with a mapping function. In some examples, the mapping function used may depend on the filter support used for GALF/ALF. In some examples, the mapping function could include, but is not limited to, a swap, and/or a flipping, and/or a rotation function. Swap functions, flipping functions, and rotation functions may be any function, weighted or otherwise) where samples within an LCU boundary are mapped to locations outside of the LCU boundary. FIG. 10A is a conceptual diagram illustrating an example use of padding with mapping for ALF. In FIG. 10A, samples indicated by the symbol S_i are padded (e.g., reused) as filter supports for samples outside of an LCU upper boundary. The samples outside the LCU upper boundary that have the same S_i designation (e.g., S_0) as samples inside the LCU boundary are copied as filter supports. FIG. 10A shows an example of diagonal padding. For example, samples within the LCU upper boundary may be flipped about a diagonal axis to locations outside the LCU upper boundary. FIG. 10B is a conceptual diagram illustrating another example use of padding with mapping for ALF. For example, samples within the LCU upper boundary may be flipped about a horizontal axis to locations outside the LCU upper boundary.

[0117] In another example, the filter coefficients for GALF/ALF that are associated with samples outside the LCU row may be forced to be 0. In another example, multiple filter supports may be applied within one slice. In one example, only a partial number of the filter supports (e.g., coefficients) are signaled, and the remaining number of filter supports may be derived from the signaled supports. For example, a 9x9 diamond filter support may be signaled by an index. For samples which require additional line buffers, the filter support may be modified to be, for example, the filter support 1100 shown in FIG. 11.

[0118] Whether to enable or disable the above methods may be signaled in a sequence parameter set (SPS), picture parameter set (PPS), view parameter set and/or slice header.

[0119] FIG. 12 is a block diagram illustrating an example video encoder 20 that may implement the techniques of this disclosure. FIG. 12 is provided for purposes of explanation and should not be considered limiting of the techniques as broadly exemplified and described in this disclosure. The techniques of this disclosure may be applicable to various coding standards or methods.

[0120] Processing circuitry includes video encoder 20, and video encoder 20 is configured to perform one or more of the example techniques described in this disclosure. For instance, video encoder 20 includes integrated circuitry, and the various units illustrated in FIG. 12 may be formed as hardware circuit blocks that are interconnected with a circuit bus. These hardware circuit blocks may be separate circuit blocks or two or more of the units may be combined into a common hardware circuit block. The hardware circuit blocks may be formed as combination of electric circuit components that form operation blocks such as arithmetic logic units (ALUs), elementary function units (EFUs), as well as logic blocks such as AND, OR, NAND, NOR, XOR, XNOR, and other similar logic blocks.

[0121] In some examples, one or more of the units illustrated in FIG. 12 may be software units executing on the processing circuitry. In such examples, the object code for these software units is stored in memory. An operating system may cause video encoder 20 to retrieve the object code and execute the object code, which causes video encoder 20 to perform operations to implement the example techniques. In some examples, the software units may be firmware that video encoder 20 executes at startup. Accordingly, video encoder 20 is a structural component having hardware that performs the example techniques or has software and/or firmware executing on the hardware to specialize the hardware to perform the example techniques.

[0122] In the example of FIG. 12, video encoder 20 includes a prediction processing unit 100, video data memory 101, 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 may include a motion estimation unit and a motion compensation unit (not shown).

[0123] Video data memory 101 may be configured to store video data to be encoded by the components of video encoder 20. The video data stored in video data memory 101 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 101 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 101 and decoded picture buffer 116 may be provided by the same memory device or separate memory devices. In various examples, video data memory 101 may be on-chip with other components of video encoder 20, or off-chip relative to those components. Video data memory 101 may be the same as, part of, or separate from storage media 19 of FIG. 1.

[0124] Video encoder 20 receives 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 partitioning to divide the CTBs of the CTU into progressively-smaller blocks. The smaller blocks may be coding blocks of CUs. For example, prediction processing unit 100 may partition a CTB associated with a CTU according to a tree structure.

[0125] 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.

[0126] Inter-prediction processing unit 120 may generate predictive data for a PU. As part of generating the predictive data for a PU, inter-prediction processing unit 120 performs inter prediction on the PU. The predictive data for the PU may include predictive blocks of the PU and motion information for the PU. Inter-prediction processing unit 120 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 processing unit 120 does not perform inter prediction on the PU. Thus, for blocks encoded in I-mode, the predicted

block is formed using spatial prediction from previously-encoded neighboring blocks within the same frame. If a PU is in a P slice, inter-prediction processing unit 120 may use uni-directional inter prediction to generate a predictive block of the PU. If a PU is in a B slice, inter-prediction processing unit 120 may use uni-directional or bi-directional inter prediction to generate a predictive block of the PU.

[0127] 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 blocks of 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.

[0128] 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. Intra-prediction processing unit 126 may use samples from sample blocks of neighboring PUs to generate a predictive block for a PU. 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.

[0129] 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 blocks of the selected predictive data may be referred to herein as the selected predictive blocks.

[0130] Residual generation unit 102 may generate, based on the coding blocks (e.g., luma, Cb and Cr coding blocks) for a CU and the selected predictive blocks (e.g., predictive luma, Cb and Cr blocks) for the PUs of the CU, residual blocks (e.g., luma, Cb and Cr residual blocks) for 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 block of a PU of the CU.

[0131] Transform processing unit 104 may perform partition the residual blocks of a CU into transform blocks of TUs of the CU. For instance, transform processing unit

104 may perform quad-tree partitioning to partition the residual blocks of the CU into transform blocks of TUs of the CU. Thus, 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.

[0132] 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.

[0133] 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.

[0134] 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 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.

[0135] Filter unit 114 may perform one or more SAO, GALF, ALF, and/or deblocking operations on the coding blocks associated with a CU. Filter unit 114 may include one

or more line buffers to store neighboring samples in order to perform the ALF and/or GAL filtering operations. Filter unit 114 may perform the ALF and GALF techniques of this disclosure that are described above. Decoded picture buffer 116 may store the reconstructed coding blocks after filter unit 114 performs the one or more filtering 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. In some examples, video encoder 20 may be configured to not process samples used by intra-prediction processing unit 126 with filter unit 114. That is, in some examples, intra-prediction processing unit 126 does not use sample values that have been filtered using GALF or ALF.

[0136] 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 values of transform coefficients for a CU.

[0137] FIG. 13 is a block diagram illustrating an example video decoder 30 that is configured to implement the techniques of this disclosure. FIG. 13 is provided for purposes of explanation and is not limiting on the techniques as broadly exemplified and described in this disclosure.

[0138] Processing circuitry includes video decoder 30, and video decoder 30 is configured to perform one or more of the example techniques described in this disclosure. For instance, video decoder 30 includes integrated circuitry, and the various units illustrated in FIG. 13 may be formed as hardware circuit blocks that are

interconnected with a circuit bus. These hardware circuit blocks may be separate circuit blocks or two or more of the units may be combined into a common hardware circuit block. The hardware circuit blocks may be formed as combination of electric circuit components that form operation blocks such as arithmetic logic units (ALUs), elementary function units (EFUs), as well as logic blocks such as AND, OR, NAND, NOR, XOR, XNOR, and other similar logic blocks.

[0139] In some examples, one or more of the units illustrated in FIG. 13 may be software units executing on the processing circuitry. In such examples, the object code for these software units is stored in memory. An operating system may cause video decoder 30 to retrieve the object code and execute the object code, which causes video decoder 30 to perform operations to implement the example techniques. In some examples, the software units may be firmware that video decoder 30 executes at startup. Accordingly, video decoder 30 is a structural component having hardware that performs the example techniques or has software and/or firmware executing on the hardware to specialize the hardware to perform the example techniques.

[0140] In the example of FIG. 13, video decoder 30 includes an entropy decoding unit 150, video data memory 151, 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. In other examples, video decoder 30 may include more, fewer, or different functional components.

[0141] Video data memory 151 may store encoded 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 151 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 151 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, or for output. Video data memory 151 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 151 and decoded picture buffer 162 may be provided by the same memory device or separate memory devices. In various examples, video data memory 151 may be on-chip with other components of video decoder 30, or off-chip relative to those components. Video data memory 151 may be the same as or part of storage media 28 of FIG. 1.

[0142] Video data memory 151 receives and stores 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 151 and may parse the NAL units to obtain 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 extracted from the bitstream. Entropy decoding unit 150 may perform a process generally reciprocal to that of entropy encoding unit 118.

[0143] In addition to obtaining 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 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.

[0144] 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. 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.

[0145] Inverse quantization unit 154 may perform particular techniques of this disclosure. For example, for at least one respective quantization group of a plurality of quantization groups within a CTB of a CTU of a picture of the video data, inverse quantization unit 154 may derive, based at least in part on local quantization information signaled in the bitstream, a respective quantization parameter for the respective quantization group. Additionally, in this example, inverse quantization unit

154 may inverse quantize, based on the respective quantization parameter for the respective quantization group, at least one transform coefficient of a transform block of a TU of a CU of the CTU. In this example, the respective quantization group is defined as a group of successive, in coding order, CUs or coding blocks so that boundaries of the respective quantization group must be boundaries of the CUs or coding blocks and a size of the respective quantization group is greater than or equal to a threshold. Video decoder 30 (e.g., inverse transform processing unit 156, reconstruction unit 158, and filter unit 160) may reconstruct, based on inverse quantized transform coefficients of the transform block, a coding block of the CU.

[0146] If a PU is encoded using intra prediction, intra-prediction processing unit 166 may perform intra prediction to generate predictive blocks of the PU. Intra-prediction processing unit 166 may use an intra prediction mode to generate the predictive blocks of the PU based on samples of spatially-neighboring blocks. Intra-prediction processing unit 166 may determine the intra prediction mode for the PU based on one or more syntax elements obtained from the bitstream.

[0147] If a PU is encoded using inter prediction, entropy decoding unit 150 may determine motion information for the PU. Motion compensation unit 164 may determine, based on the motion information of the PU, one or more reference blocks. Motion compensation unit 164 may generate, based on the one or more reference blocks, predictive blocks (e.g., predictive luma, Cb and Cr blocks) for the PU.

[0148] Reconstruction unit 158 may use transform blocks (e.g., luma, Cb and Cr transform blocks) for 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) for 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., luma, Cb and Cr predictive blocks) to reconstruct the coding blocks (e.g., luma, Cb and Cr coding blocks) of the CU.

[0149] Filter unit 160 may perform one or more SAO, GALF, ALF, and/or deblocking operations on the coding blocks associated with a CU. Filter unit 160 may include one or more line buffers to store neighboring samples in order to perform the ALF and/or GAL filtering operations. Filter unit 160 may perform the ALF and GALF techniques of this disclosure that are described above. Decoded picture buffer 162 may store the reconstructed coding blocks after filter unit 160 performs the one or more filtering

operations on the reconstructed coding blocks. 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 in decoded picture buffer 162, intra prediction or inter prediction operations for PUs of other CUs. In some examples, video decoder 30 may be configured to not process samples used by intra-prediction processing unit 166 with filter unit 160. That is, in some examples, intra-prediction processing unit 166 does not use sample values that have been filtered using GALF or ALF.

[0150] FIG. 14 is a flowchart illustrating an example coding method of the disclosure. Both video encoder 20 and video decoder 30, including filter unit 114 and filter unit 160, may be configured to perform the techniques of FIG. 14.

[0151] In one example of the disclosure, video encoder 20 and/or video decoder 30 may be configured to code a block of video data to obtain a coded block of video data (1400), apply an adaptive loop filter defined by a filter support to samples of the coded block of video data, including using padded samples for portions of the filter support that are outside of the block (1402), and output the coded block of video data after applying adaptive loop filtering to the coded block of video data (1404).

[0152] In one example, video encoder 20 and video decoder 30 may be further configured to generate the padded samples from samples of a reference picture pointed to by a motion vector associated with the block of video data.

[0153] In another example, video encoder 20 and video decoder 30 may be further configured to generate the padded samples from a first row or a last row of the coded block of video data.

[0154] In another example, video encoder 20 and video decoder 30 may be further configured to generate the padded samples by applying a linear function to samples of one or more reference pictures.

[0155] In another example, video encoder 20 and video decoder 30 may be further configured to identify the samples of the one or more reference pictures using one or more motion vectors.

[0156] In another example, video encoder 20 and video decoder 30 may be further configured to round the one or more motion vectors to integer values.

[0157] In another example, video encoder 20 and video decoder 30 may be further configured to set the one or more motion vector to zero motion vectors.

[0158] In another example, video encoder 20 and video decoder 30 may be further configured to generate the padded samples using a mapping function and samples of the decoded block of video data. The mapping function may include one or more of a swap function, a flipping function, or a rotation function.

[0159] In another example, video encoder 20 and video decoder 30 may be further configured to set coefficients of the portions of the filter support of the adaptive loop filter that are outside the block of video data to zero.

[0160] In another example, video encoder 20 and video decoder 30 may be further configured to generate/receive a syntax element indicating the use of the padded samples for portions of the filter support that are outside of the block.

[0161] Certain aspects of this disclosure have been described with respect to HEVC, extensions of the HEVC, and the developing VVC 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.

[0162] A video coder, as described in this disclosure, may refer to a video encoder or a video decoder. Similarly, a video coding unit may refer to a video encoder or a video decoder. Likewise, video coding may refer to video encoding or video decoding, as applicable. In this disclosure, the phrase “based on” may indicate based only on, based at least in part on, or based in some way on. 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 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. Example types of video blocks may include coding tree blocks, coding blocks, and other types of blocks of video data.

[0163] The techniques of this disclosure may be applied to video coding in support of any of a variety of multimedia applications, such as over-the-air television broadcasts, cable television transmissions, satellite television transmissions, Internet streaming video transmissions, such as dynamic adaptive streaming over HTTP (DASH), digital video that is encoded onto a data storage medium, decoding of digital video stored on a data storage medium, or other applications.

[0164] 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.

[0165] 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 processing circuits 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.

[0166] 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.

[0167] Functionality described in this disclosure may be performed by fixed function and/or programmable processing circuitry. For instance, instructions may be executed by fixed function and/or programmable processing circuitry. Such processing circuitry may include one or more processors, such as one or more digital signal processors (DSPs), general purpose microprocessors, application specific integrated circuits (ASICs), field programmable logic 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 examples, 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. Processing circuits may be coupled to other components in various ways. For example, a processing circuit may be coupled to other components via an internal device interconnect, a wired or wireless network connection, or another communication medium.

[0168] 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.

[0169] Various examples have been described. These and other examples are within the scope of the following claims.

CLAIMS

1. A method of coding video data, the method comprising: coding a block of video data to obtain a coded block of video data; generating padded samples using a mapping function and samples of the decoded block of video data; applying an adaptive loop filter defined by a filter support to samples of the coded block of video data, including using the padded samples for portions of the filter support that are outside of the block; and outputting the coded block of video data after applying adaptive loop filtering to the coded block of video data.
2. The method of claim 1, further comprising: generating the padded samples from samples of a reference picture pointed to by a motion vector associated with the block of video data.
3. The method of claim 1, further comprising: generating the padded samples from a first row or a last row of the coded block of video data.
4. The method of claim 1, further comprising: generating the padded samples by applying a linear function to samples of one or more reference pictures.
5. The method of claim 4, further comprising: identifying the samples of the one or more reference pictures using one or more motion vectors.
6. The method of claim 5, further comprising: rounding the one or more motion vectors to integer values.
7. The method of claim 5, further comprising: setting the one or more motion vector to zero motion vectors.
8. The method of claim 1, wherein the mapping function includes one or more of a swap function, a flipping function, or a rotation function.
9. The method of claim 1, further comprising: setting coefficients of the portions of the filter support of the adaptive loop filter that are outside the block of video data to zero.
10. The method of claim 1, further comprising: receiving a syntax element indicating the use of the padded samples for portions of the filter support that are outside of the block.
11. The method of claim 1, wherein coding is encoding.
12. The method of claim 1, wherein coding is decoding.
13. An apparatus configured to code video data, the apparatus comprising: a memory configured to store a block of video data; and one or more processors implemented in circuitry in

communication with the memory, the one or more processors configured to: code the block of video data to obtain a coded block of video data; generate padded samples using a mapping function and samples of the decoded block of video data; apply an adaptive loop filter defined by a filter support to samples of the coded block of video data, including using the padded samples for portions of the filter support that are outside of the block; and output the coded block of video data after applying adaptive loop filtering to the coded block of video data.

14. The apparatus of claim 13, wherein the one or more processors are further configured to: generate the padded samples from samples of a reference picture pointed to by a motion vector associated with the block of video data.

15. The apparatus of claim 13, wherein the one or more processors are further configured to: generate the padded samples from a first row or a last row of the coded block of video data.

16. The apparatus of claim 13, wherein the one or more processors are further configured to: generate the padded samples by applying a linear function to samples of one or more reference pictures.

17. The apparatus of claim 16, wherein the one or more processors are further configured to: identify the samples of the one or more reference pictures using one or more motion vectors.

18. The apparatus of claim 17, wherein the one or more processors are further configured to: round the one or more motion vectors to integer values.

19. The apparatus of claim 17, wherein the one or more processors are further configured to: set the one or more motion vector to zero motion vectors.

20. The apparatus of claim 13, wherein the mapping function includes one or more of a swap function, a flipping function, or a rotation function.

21. The apparatus of claim 13, wherein the one or more processors are further configured to: set coefficients of the portions of the filter support of the adaptive loop filter that are outside the block of video data to zero.

22. The apparatus of claim 13, wherein the one or more processors are further configured to: receive a syntax element indicating the use of the padded samples for portions of the filter support that are outside of the block.

23. The apparatus of claim 13, wherein to code the block of video data to obtain the coded block of video, the one or more processors are configured to encode the block of video data to obtain an encoded block of video.

24. The apparatus of claim 13, wherein to code the block of video data to obtain the coded block

of video, the one or more processors are configured to decode the block of video data to obtain a decoded block of video.

25. An apparatus configured to code video data, the apparatus comprising: means for coding a block of video data to obtain a coded block of video data; means for generating padded samples using a mapping function and samples of the decoded block of video data; means for applying an adaptive loop filter defined by a filter support to samples of the coded block of video data, including using the padded samples for portions of the filter support that are outside of the block; and means for outputting the coded block of video data after applying adaptive loop filtering to the coded block of video data.

26. The apparatus of claim 25, further comprising: means for generating the padded samples from samples of a reference picture pointed to by a motion vector associated with the block of video data.

27. A non-transitory computer-readable medium storing instructions that, when executed, cause one or more processors to: code a block of video data to obtain a coded block of video data; generate padded samples using a mapping function and samples of the decoded block of video data; apply an adaptive loop filter defined by a filter support to samples of the coded block of video data, including using the padded samples for portions of the filter support that are outside of the block; and output the coded block of video data after applying adaptive loop filtering to the coded block of video data.

28. The non-transitory computer-readable medium of claim 27, wherein the instructions further cause the one or more processors to: generate the padded samples from samples of a reference picture pointed to by a motion vector associated with the block of video data.