



(51) International Patent Classification:

H04N 19/70 (2014.01) H04N 19/30 (2014.01)
H04N 19/46 (2014.01) H04N 19/68 (2014.01)

(21) International Application Number:

PCT/US2014/067363

(22) International Filing Date:

25 November 2014 (25.11.2014)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/908,671 25 November 2013 (25.11.2013) US
61/909,347 26 November 2013 (26.11.2013) US
14/552,129 24 November 2014 (24.11.2014) US

(71) Applicant: **QUALCOMM INCORPORATED** [US/US];
ATTN: International IP Administration, 5775 Morehouse
Drive, San Diego, California 92121-1714 (US).

(72) Inventors: **HENDRY, Fnu**; 5775 Morehouse Drive, San
Diego, California 92121-1714 (US). **WANG, Ye-Kui**;

5775 Morehouse Drive, San Diego, California 92121-1714
(US). **RAMASUBRAMONIAN, Adarsh Krishnan**; 5775
Morehouse Drive, San Diego, California 92121-1714 (US).
CHEN, Ying; 5775 Morehouse Drive, San Diego, Califor-
nia 92121-1714 (US).

(74) Agent: **PRIEM, David F.**; Shumaker & Sieffert, P.A.,
1625 Radio Drive, Suite 300, Woodbury, Minnesota 55125
(US).

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

[Continued on next page]

(54) Title: POC VALUE DESIGN FOR MULTI-LAYER VIDEO CODING

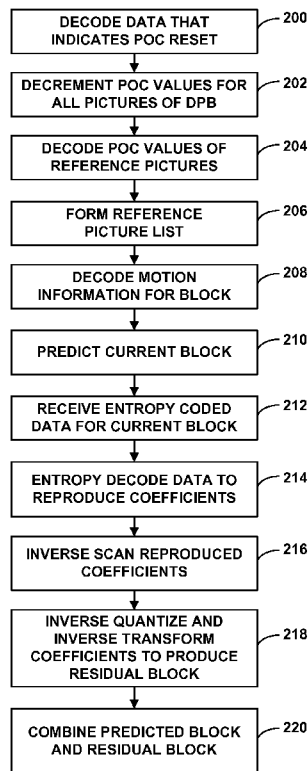


FIG. 5

(57) Abstract: In an example, a method of decoding video data includes decoding data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of multi-layer video data, wherein the first picture is included in an access unit. The example method also includes, based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrementing POC values of all pictures stored to a decoded picture buffer (DPB) that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.

WO 2015/077764 A1



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

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

Published:

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

POC VALUE DESIGN FOR MULTI-LAYER VIDEO CODING

[0001] This application claims the benefit of U.S. Provisional Application No. 61/908,671, filed November 25, 2013, and U.S. Provisional Application No. 61/909,347, filed November 26, 2013, which are hereby incorporated by reference in their entirety.

TECHNICAL FIELD

[0002] This disclosure relates to video coding.

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 coding techniques, such as those described in the standards defined by MPEG-2, MPEG-4, ITU-T H.263, ITU-T H.264/MPEG-4, Part 10, Advanced Video Coding (AVC), the High Efficiency Video Coding (HEVC) standard presently under development, and extensions of such standards. The video devices may transmit, receive, encode, decode, and/or store digital video information more efficiently by implementing such video coding techniques.

[0004] Video coding techniques include 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, which may also be referred to as treeblocks, coding units (CUs) and/or coding nodes. Video blocks in an intra-coded (I) slice of a picture are encoded using spatial prediction with respect to reference samples in neighboring blocks in the same picture. Video blocks in an inter-coded (P or B) slice of a picture may use spatial prediction with respect to reference samples in neighboring blocks in the same picture or temporal prediction with respect to reference samples in

other reference pictures. Pictures may be referred to as frames, and reference pictures may be referred to as reference frames.

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

SUMMARY

[0006] In general, this disclosure describes techniques for improving error resilience for coding video data related to reference pictures. In particular, this disclosure described a picture order count (POC) value design that may improve such error resilience. These techniques may be used when coding single layer video data or multi-layer video data. In general, these techniques include decrementing POC values to improve error resilience in a picture outputting process. For example, a video decoder may be configured to decrement POC values of all pictures (of all layers) stored to a decoded picture buffer (DPB) prior to decoding an initial picture of an access unit in which POC resetting is performed. In this manner, POC values may remain aligned in the DPB, thereby allowing the pictures of the access unit to be output in the proper order.

[0007] In one example, a method of decoding video data includes decoding data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of multi-layer video data, wherein the first picture is included in an access unit, and, based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrementing POC values of all pictures stored to a decoded picture buffer (DPB) that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.

[0008] In another example, a method of encoding video data includes encoding data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of multi-layer video data, wherein the first picture is included in an access unit, and, based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrementing POC values of all pictures stored to a decoded picture buffer (DPB) that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.

[0009] In another example, a device for coding video data includes a memory comprising a decoded picture buffer (DPB) configured to store multi-layer video data. The device also includes a video coder configured to code data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of the multi-layer video data, wherein the first picture is included in an access unit, and, based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrement POC values of all pictures stored to the DPB that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.

[0010] In another example, a device for coding video data includes means for coding data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of the multi-layer video data, wherein the first picture is included in an access unit, and, based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, means for decrementing POC values of all pictures stored to the DPB that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.

[0011] In another example, a non-transitory computer-readable storage medium has stored thereon instructions that, when executed, cause a processor of a device for coding video data to code data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of the multi-layer video data, wherein the first picture is included in an access unit, and, based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrement POC values of all pictures stored to the DPB that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.

[0012] The details of one or more examples are set forth in the accompanying drawings and the description below. Other features, objects, and advantages will be apparent from the description and drawings, and from the claims.

BRIEF DESCRIPTION OF DRAWINGS

[0013] FIG. 1 is a block diagram illustrating an example video encoding and decoding system that may utilize techniques for coding POC value information according to the techniques of this disclosure.

[0014] FIG. 2 is a block diagram illustrating an example of a video encoder that may implement techniques for coding POC value information according to the techniques of this disclosure.

[0015] FIG. 3 is a block diagram illustrating an example of a video decoder that may implement techniques for coding POC value information according to the techniques of this disclosure.

[0016] FIG. 4 is a flowchart illustrating an example method for encoding data including resetting POC values in accordance with the techniques of this disclosure.

[0017] FIG. 5 is a flowchart illustrating an example method for decoding data including resetting POC values in accordance with the techniques of this disclosure.

DETAILED DESCRIPTION

[0018] In general, this disclosure describes a picture order count (POC) value design for video coding, which may be beneficial for multi-layer video coding. This disclosure describes various design improvements for signaling and derivation of POC values in multi-layer video coding. However, while the techniques are generally described with respect to multi-layer video coding, the techniques of this disclosure may also be applied to single-layer video coding.

[0019] POC values may indicate a display order of pictures and may generally be used to identify the picture. For example, a picture having a POC value of 0 (i.e., a picture having a POC value equal to 0) is displayed prior to a picture having a POC value of 1. When a block of a current picture is inter-predicted relative to a reference picture, the reference picture may be identified using a POC value for the reference picture. More particularly, POC values for reference pictures to be included in a reference picture list may be signaled in a parameter set, such as sequence parameter set (SPS), picture parameter set (PPS), and/or a slice header corresponding to the block. In this manner, a video encoder may identify a reference picture by signaling an index into the reference picture list corresponding to the position of the reference picture in the reference picture

list, and a video decoder may identify the reference picture by constructing the reference picture list (based on the POC values) and using the reference index to identify the position of a reference picture in the reference picture list.

[0020] POC values may be composed of least significance bits (LSB) and most significant bits (MSB). Rather than increasing POC values of pictures indefinitely during coding, a POC value of a picture may be periodically reset to be equal to zero. POC values are typically reset for an Instantaneous Decoder Refresh (IDR) picture. In addition, the video coder may decrement the POC values of pictures that have been processed and stored to a decoded picture buffer (DPB) previous to the current picture associated with the POC reset. The POC values of reference pictures that have been decremented may have a negative value in order to maintain the appropriate output order of such pictures.

[0021] In multi-layer video coding (e.g., multi-view video coding, scalable video coding, or the like), a DPB may be partitioned into “sub-DPBs,” with each layer of video data assigned to a separately managed sub-DPB of the DPB. A view component may comprise an encoded picture for a particular layer or view at a particular time (in terms of display order, or output order). An access unit is, generally, a unit of data including all view components (e.g., all network abstraction layer (NAL) units) for a common temporal instance. Accordingly, an access unit for a particular temporal instance may include respective view components from respective sub-DPBs. The view components of an access unit are typically intended to be output together (i.e., output substantially simultaneously), where outputting a picture generally involves transferring pictures from the DPB (e.g., storing pictures from the DPB to an external memory, sending the pictures from the DPB to a display, removing pictures from the DPB, or the like).

[0022] A video coder may use an output and bumping process to output pictures from the DPB that are marked as “needed for output” and “unused for reference.” In some instances, the picture output process may be invoked two times for each picture, including (1) before the decoding of the picture but after parsing the slice header of the first slice of the picture, and (2) when the last decoding unit of the picture is removed from the coded picture buffer. When the bumping process is invoked, the process outputs all pictures that belong to the selected access unit.

[0023] Existing POC designs (e.g., as noted below) may not be optimized. For example, in some instances, decrementing POC values may be a layer-specific process.

Hence, the POC values of earlier pictures in a sub-DPB (e.g., pictures that have been processed and stored to the sub-DPB previously to the picture currently being coded) for a layer may not be cross-layer aligned with the POC values of earlier pictures in other sub-DPBs when POC values or POC most significant bits (MSB) of pictures of a layer are reset during the decoding of a picture of that layer, as described in the example below. Cross-layer alignment, for example, occurs when each picture of an access unit has the same POC value, such that the pictures of the access unit are output simultaneously or nearly simultaneously. In existing POC designs, POC values of the pictures in higher layers (e.g., layers having relatively larger layer identifiers) may not be updated until after decoding a picture of a base layer. Updating the POC or POC MSB in this manner may create a problem for picture output through the bumping process, which requires cross-layer alignment of POC values for each access unit that contains decoded pictures in the DPB.

[0024] In an example for purposes of illustration, assume a multi-layer sequence has three layers (e.g., layers A, B, and C). In addition, assume that a particular access unit includes pictures for which POC resetting is performed. In this example, a video decoder may reset the POC value of a layer A picture of the access unit (e.g., upon receiving an indication in a bitstream to perform a POC reset). The video decoder may also decrement pictures of layer A that were previously stored to the sub-DPB for layer A to maintain the appropriate output order. However, while decoding the first picture of layer A and prior to decoding the layer B picture of the access unit, the POC values of pictures for layers B and C that are stored to the respective sub-DPBs for layers B and C become misaligned (not cross-layer aligned) with the POC values of the pictures stored to the sub-DPB of layer A. That is, because the layer B picture and the layer C picture of the access unit have not yet been reset, the POC values of the respective sub-layers of each of layers B and C have also not yet been decremented. Such misalignment may cause the pictures to be output from the DPB in the improper order, because the pictures of an access unit should be cross-layer aligned prior to being output.

[0025] The POC value design described in this disclosure includes various aspects, any or all of which may be implemented alone or in any combination. The techniques of this disclosure may overcome the issue discussed above. In one example, video coders (e.g., video encoders and video decoders) may be configured to decrement POC values of all pictures (of all layers) stored to a DPB prior to decoding an initial picture of an access unit in which POC resetting is performed. For example, video coders may apply

a POC decrement process to earlier pictures in coding order in a DPB (i.e., all pictures that precede a current picture in decoding order for all sub-DPBs) once, at the beginning of each access unit that requires POC reset. In other words, the video coders may apply the POC decrement process a single time, at the beginning of each access unit that requires POC reset, such that POC values of all pictures that precede the current picture in coding order within all sub-DPBs are decremented based on the single invocation of the POC decrement process. The POC decrement process may be invoked before the decoding of the first picture received for the access unit but after parsing and decoding of the slice header information of the first slice of that picture. In this manner, POC values of pictures stored to the DPB (i.e., for all sub-DPBs) may remain aligned when coding the pictures of the access unit, thereby allowing the pictures of the DPB to be output in the proper order.

[0026] Other techniques of this disclosure include techniques for signaling POC values. For example, another potential issue with certain POC designs is with respect to signalling of additional POC least significant bit (LSB) information. In some instances, as described in greater detail below, the POC LSB of a POC-anchor picture is not signalled when the POC-anchor picture in a base layer is an IDR picture. Not signaling the POC LSB in such instances may create an issue, because an IDR picture in a base layer does not contain information for deriving the value used to decrement POC values of earlier pictures (e.g., pictures preceding a picture currently being coded in coding order) in the DPB.

[0027] According to aspects of this disclosure, a condition for signalling POC LSB of the POC-anchor picture may be updated. For example, according to aspects of this disclosure, a video encoder may signal (and a video decoder may decode) data that indicates the additional POC LSB for base layer IDR pictures having a `poc_reset_idc` syntax element that is equal to 2 (as described in greater detail below).

[0028] FIG. 1 is a block diagram illustrating an example video encoding and decoding system 10 that may utilize techniques for managing POC values according to the 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 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, so-called “smart” pads, televisions, cameras, display devices, digital media players, video gaming consoles, video streaming device, or the like. In some cases, source device 12 and destination device 14 may be equipped for wireless communication.

[0029] Destination device 14 may receive the 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 one example, computer-readable medium 16 may comprise a communication medium to enable source device 12 to transmit encoded video data directly to destination device 14 in real-time. The encoded video data may be modulated according to a communication standard, such as a wireless communication protocol, and transmitted to destination device 14. The communication medium may comprise any wireless or wired communication medium, such as a radio frequency (RF) spectrum or one or more physical transmission lines. The communication medium may form part of a packet-based network, such as a local area network, a wide-area network, or a global network such as the Internet. The communication medium may include routers, switches, base stations, or any other equipment that may be useful to facilitate communication from source device 12 to destination device 14.

[0030] In some examples, encoded data may be output from output interface 22 to a storage device. Similarly, encoded data may be accessed from the storage device by input interface. The storage device may include any of a variety of distributed or locally accessed data storage media such as a hard drive, Blu-ray discs, DVDs, CD-ROMs, flash memory, volatile or non-volatile memory, or any other suitable digital storage media for storing encoded video data. In a further example, the storage device may correspond to a file server or another intermediate storage device that may store the encoded video generated by source device 12. Destination device 14 may access stored video data from the storage device via streaming or download. The file server may be any type of server capable of storing encoded video data and transmitting that encoded video data to the destination device 14. Example file servers include a web server (e.g., for a website), an FTP server, network attached storage (NAS) devices, or a local disk drive. Destination device 14 may access the encoded video data through any standard data connection, including an Internet connection. This may include a wireless channel (e.g., a Wi-Fi connection), a wired connection (e.g., DSL, cable modem, etc.), or a combination of both that is suitable for accessing encoded video data stored on a file

server. The transmission of encoded video data from the storage device may be a streaming transmission, a download transmission, or a combination thereof.

[0031] The techniques of this disclosure are not necessarily limited to wireless applications or settings. The techniques may be applied to video coding in support of any of a variety of multimedia applications, such as over-the-air television broadcasts, cable television transmissions, satellite television transmissions, 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. In some examples, system 10 may be configured to support one-way or two-way video transmission to support applications such as video streaming, video playback, video broadcasting, and/or video telephony.

[0032] In the example of FIG. 1, source device 12 includes video source 18, video encoder 20, and output interface 22. Destination device 14 includes input interface 28, video decoder 30, and display device 32. In accordance with this disclosure, video encoder 20 of source device 12 may be configured to apply the techniques for coding POC value information according to the techniques of this disclosure. In other examples, a source device and a destination device may include other components or arrangements. For example, source device 12 may receive video data from an external video source 18, such as an external camera. Likewise, destination device 14 may interface with an external display device, rather than including an integrated display device.

[0033] The illustrated system 10 of FIG. 1 is merely one example. Techniques for managing POC values according to the techniques of this disclosure may be performed by any digital video encoding and/or decoding device. Although generally the techniques of this disclosure are performed by a video encoding device, the techniques may also be performed by a video encoder/decoder, typically referred to as a "CODEC." Moreover, the techniques of this disclosure may also be performed by a video preprocessor. Source device 12 and destination device 14 are merely examples of such coding devices in which source device 12 generates coded video data for transmission to destination device 14. In some examples, devices 12, 14 may operate in a substantially symmetrical manner such that each of devices 12, 14 include video encoding and decoding components. Hence, system 10 may support one-way or two-way video transmission between video devices 12, 14, e.g., for video streaming, video playback, video broadcasting, or video telephony.

[0034] Video source 18 of source device 12 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 from a video content provider. As a further alternative, video source 18 may generate computer graphics-based data as the source video, or a combination of live video, archived video, and computer-generated video. In some cases, if video source 18 is a video camera, source device 12 and destination device 14 may form so-called camera phones or video phones. As mentioned above, however, the techniques described in this disclosure may be applicable to video coding in general, and may be applied to wireless and/or wired applications. In each case, the captured, pre-captured, or computer-generated video may be encoded by video encoder 20. The encoded video information may then be output by output interface 22 onto a computer-readable medium 16.

[0035] 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, 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 various forms, in various examples.

[0036] Input interface 28 of destination device 14 receives information from computer-readable medium 16. The information of computer-readable medium 16 may include syntax information defined by video encoder 20, which is also used by video decoder 30, that includes syntax elements that describe characteristics and/or processing of blocks and other coded units, e.g., GOPs. Display device 32 displays the decoded video data to a user, and may comprise any of a variety of display devices such as a cathode ray tube (CRT), a liquid crystal display (LCD), a plasma display, an organic light emitting diode (OLED) display, or another type of display device.

[0037] Video encoder 20 and video decoder 30 each may be implemented as any of a variety of suitable encoder circuitry, such as one or more microprocessors, digital signal processors (DSPs), application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs), discrete logic, software, hardware, firmware or any combinations

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

[0038] Although not shown in FIG. 1, in some aspects, video encoder 20 and video decoder 30 may each be integrated with an audio encoder and decoder, and may include appropriate MUX-DEMUX units, or other hardware and software, to handle encoding of both audio and video in a common data stream or separate data streams. If applicable, MUX-DEMUX units may conform to the ITU H.223 multiplexer protocol, or other protocols such as the user datagram protocol (UDP).

[0039] This disclosure may generally refer to video encoder 20 “signaling” certain information to another device, such as video decoder 30. It should be understood, however, that video encoder 20 may signal information by associating certain syntax elements with various encoded portions of video data. That is, video encoder 20 may “signal” data by storing certain syntax elements to headers of various encoded portions of video data. In some cases, such syntax elements may be encoded and stored prior to being received and decoded by video decoder 30. Thus, the term “signaling” may generally refer to the communication of syntax or other data for decoding compressed video data, whether such communication occurs in real- or near-real-time or over a span of time, such as might occur when storing syntax elements to a medium at the time of encoding, which then may be retrieved by a decoding device at any time after being stored to this medium.

[0040] Video encoder 20 and video decoder 30 may operate according to a video coding standard. Example video coding standards include ITU-T H.261, ISO/IEC MPEG-1 Visual, ITU-T H.262 or ISO/IEC MPEG-2 Visual, ITU-T H.263, ISO/IEC MPEG-4 Visual and ITU-T H.264 (also known as ISO/IEC MPEG-4 AVC), including its Scalable Video Coding (SVC) and Multiview Video Coding (MVC) extensions. Recently, the design of a new video coding standard, namely High-Efficiency Video Coding (HEVC), has been finalized by the Joint Collaboration Team on Video Coding (JCT-VC) of ITU-T Video Coding Experts Group (VCEG) and ISO/IEC Motion Picture

Experts Group (MPEG). The latest HEVC draft specification, and referred to as HEVC WD hereinafter, is available from http://phenix.int-evry.fr/jct/doc_end_user/documents/15_Geneva/wg11/JCTVC-O1003-v1.zip. The multiview extension to HEVC, namely MV-HEVC, is also being developed by the JCT-3V. A recent Working Draft (WD) of MV-HEVC, referred to as MV-HEVC WD6 hereinafter, is available from http://phenix.it-sudparis.eu/jct2/doc_end_user/documents/6_Geneva/wg11/JCT3V-F1004-v3.zip. The scalable extension to HEVC, named SHVC, is also being developed by the JCT-VC. A recent Working Draft (WD) of SHVC and referred to as SHVC WD4 hereinafter, is available from http://phenix.int-evry.fr/jct/doc_end_user/documents/15_Geneva/wg11/JCTVC-O1008-v1.zip. The techniques of this disclosure, however, are not limited to any particular coding standard. Other examples of video coding standards include MPEG-2 and ITU-T H.263.

[0041] The HEVC standardization efforts are based on a model of a video coding device referred to as the HEVC Test Model (HM). The HM presumes several additional capabilities of video coding devices relative to existing devices according to, e.g., ITU-T H.264/AVC. For example, whereas H.264 provides nine intra-prediction encoding modes, the HM may provide as many as thirty-three intra-prediction encoding modes.

[0042] In general, the working model of the HM describes that a video frame or picture may be divided into a sequence of treeblocks or largest coding units (LCU) that include both luma and chroma samples. Syntax data within a bitstream may define a size for the LCU, which is a largest coding unit in terms of the number of pixels. A slice includes a number of consecutive treeblocks in coding order. A video frame or picture may be partitioned into one or more slices. Each treeblock may be split into coding units (CUs) according to a quadtree. In general, a quadtree data structure includes one node per CU, with a root node corresponding to the treeblock. If a CU is split into four sub-CUs, the node corresponding to the CU includes four leaf nodes, each of which corresponds to one of the sub-CUs.

[0043] Each node of the quadtree data structure may provide syntax data for the corresponding CU. For example, a node in the quadtree may include a split flag, indicating whether the CU corresponding to the node is split into sub-CUs. Syntax elements for a CU may be defined recursively, and may depend on whether the CU is split into sub-CUs. If a CU is not split further, it is referred to as a leaf-CU. In this disclosure, four sub-CUs of a leaf-CU will also be referred to as leaf-CUs even if there

is no explicit splitting of the original leaf-CU. For example, if a CU at 16x16 size is not split further, the four 8x8 sub-CUs will also be referred to as leaf-CUs although the 16x16 CU was never split.

[0044] A CU has a similar purpose as a macroblock of the H.264 standard, except that a CU does not have a size distinction. For example, a treeblock may be split into four child nodes (also referred to as sub-CUs), and each child node may in turn be a parent node and be split into another four child nodes. A final, unsplit child node, referred to as a leaf node of the quadtree, comprises a coding node, also referred to as a leaf-CU. Syntax data associated with a coded bitstream may define a maximum number of times a treeblock may be split, referred to as a maximum CU depth, and may also define a minimum size of the coding nodes. Accordingly, a bitstream may also define a smallest coding unit (SCU). This disclosure uses the term "block" to refer to any of a CU, PU, or TU, in the context of HEVC, or similar data structures in the context of other standards (e.g., macroblocks and sub-blocks thereof in H.264/AVC).

[0045] A CU includes a coding node and prediction units (PUs) and transform units (TUs) associated with the coding node. A size of the CU corresponds to a size of the coding node and must be square in shape. The size of the CU may range from 8x8 pixels up to the size of the treeblock with a maximum of 64x64 pixels or greater. Each CU may contain one or more PUs and one or more TUs. Syntax data associated with a CU may describe, for example, partitioning of the CU into one or more PUs.

Partitioning modes may differ between whether the CU is skip or direct mode encoded, intra-prediction mode encoded, or inter-prediction mode encoded. PUs may be partitioned to be non-square in shape. Syntax data associated with a CU may also describe, for example, partitioning of the CU into one or more TUs according to a quadtree. A TU can be square or non-square (e.g., rectangular) in shape.

[0046] The HEVC standard allows for transformations according to TUs, which may be different for different CUs. The TUs are typically sized based on the size of PUs within a given CU defined for a partitioned LCU, although this may not always be the case. The TUs are typically the same size or smaller than the PUs. In some examples, residual samples corresponding to a CU may be subdivided into smaller units using a quadtree structure known as "residual quad tree" (RQT). The leaf nodes of the RQT may be referred to as transform units (TUs). Pixel difference values associated with the TUs may be transformed to produce transform coefficients, which may be quantized.

[0047] A leaf-CU may include one or more prediction units (PUs). In general, a PU represents a spatial area corresponding to all or a portion of the corresponding CU, and may include data for retrieving a reference sample for the PU. Moreover, a PU includes data related to prediction. For example, when the PU is intra-mode encoded, data for the PU may be included in a residual quadtree (RQT), which may include data describing an intra-prediction mode for a TU corresponding to the PU. As another example, when the PU is inter-mode encoded, the PU may include data defining one or more motion vectors for the PU. The data defining the motion vector for a PU may describe, for example, a horizontal component of the motion vector, a vertical component of the motion vector, a resolution for the motion vector (e.g., one-quarter pixel precision or one-eighth pixel precision), a reference picture to which the motion vector points, and/or a reference picture list (e.g., List 0, List 1, or List C) for the motion vector.

[0048] Motion information for an inter-predicted PU (more particularly, an inter-predicted portion of a CU to which a PU corresponds) may include a reference picture list identifier and a reference index corresponding to a position of a reference picture in the reference picture list. Video encoder 20, for example, may construct the reference picture list (including any modifications made to an originally-constructed reference picture list), and perform a motion search among the reference pictures in the reference picture list to identify a closely matching block (e.g., based on sum of absolute difference (SAD) metrics, or similar metrics), also referred to as a reference block. In order to signal the location of the reference block, video encoder 20 may encode the reference picture list identifier and the reference index for the PU. Furthermore, video encoder 20 may encode data representative of reference pictures to be included in the reference picture list, e.g., by signaling data representative of POC values for the reference pictures in a slice header and/or in a parameter set, such as a PPS or an SPS.

[0049] A leaf-CU having one or more PUs may also include one or more transform units (TUs). The transform units may be specified using an RQT (also referred to as a TU quadtree structure), as discussed above. For example, a split flag may indicate whether a leaf-CU is split into four transform units. Then, each transform unit may be split further into further sub-TUs. When a TU is not split further, it may be referred to as a leaf-TU. Generally, for intra coding, all the leaf-TUs belonging to a leaf-CU share the same intra prediction mode. That is, the same intra-prediction mode is generally applied to calculate predicted values for all TUs of a leaf-CU. For intra coding, a video

encoder may calculate a residual value for each leaf-TU using the intra prediction mode, as a difference between the portion of the CU corresponding to the TU and the original block. A TU is not necessarily limited to the size of a PU. Thus, TUs may be larger or smaller than a PU. For intra coding, a PU may be collocated with a corresponding leaf-TU for the same CU. In some examples, the maximum size of a leaf-TU may correspond to the size of the corresponding leaf-CU.

[0050] Moreover, TUs of leaf-CUs may also be associated with respective quadtree data structures, referred to as residual quadtrees (RQTs). That is, a leaf-CU may include a quadtree indicating how the leaf-CU is partitioned into TUs. The root node of a TU quadtree generally corresponds to a leaf-CU, while the root node of a CU quadtree generally corresponds to a treeblock (or LCU). TUs of the RQT that are not split are referred to as leaf-TUs. In general, this disclosure uses the terms CU and TU to refer to leaf-CU and leaf-TU, respectively, unless noted otherwise.

[0051] A video sequence typically includes a series of video frames or pictures. A group of pictures (GOP) generally comprises a series of one or more of the video pictures. A GOP may include syntax data in a header of the GOP, a header of one or more of the pictures, or elsewhere, that describes a number of pictures included in the GOP. Each slice of a picture may include slice syntax data that describes an encoding mode for the respective slice. Video encoder 20 typically operates on video blocks within individual video slices in order to encode the video data. A video block may correspond to a coding node within a CU. The video blocks may have fixed or varying sizes, and may differ in size according to a specified coding standard.

[0052] As an example, the HM supports prediction in various PU sizes. Assuming that the size of a particular CU is $2N \times 2N$, the HM supports intra-prediction in PU sizes of $2N \times 2N$ or $N \times N$, and inter-prediction in symmetric PU sizes of $2N \times 2N$, $2N \times N$, $N \times 2N$, or $N \times N$. The HM also supports asymmetric partitioning for inter-prediction in PU sizes of $2N \times nU$, $2N \times nD$, $nL \times 2N$, and $nR \times 2N$. In asymmetric partitioning, one direction of a CU is not partitioned, while the other direction is partitioned into 25% and 75%. The portion of the CU corresponding to the 25% partition is indicated by an "n" followed by an indication of "Up", "Down," "Left," or "Right." Thus, for example, " $2N \times nU$ " refers to a $2N \times 2N$ CU that is partitioned horizontally with a $2N \times 0.5N$ PU on top and a $2N \times 1.5N$ PU on bottom.

[0053] In this disclosure, " $N \times N$ " and " N by N " may be used interchangeably to refer to the pixel dimensions of a video block in terms of vertical and horizontal dimensions,

e.g., 16x16 pixels or 16 by 16 pixels. In general, a 16x16 block will have 16 pixels in a vertical direction ($y = 16$) and 16 pixels in a horizontal direction ($x = 16$). Likewise, an $N \times N$ block generally has N pixels in a vertical direction and N pixels in a horizontal direction, where N represents a nonnegative integer value. The pixels in a block may be arranged in rows and columns. Moreover, blocks need not necessarily have the same number of pixels in the horizontal direction as in the vertical direction. For example, blocks may comprise $N \times M$ pixels, where M is not necessarily equal to N .

[0054] Following intra-predictive or inter-predictive coding using the PUs of a CU, video encoder 20 may calculate residual data for the TUs of the CU. The PUs may comprise syntax data describing a method or mode of generating predictive pixel data in the spatial domain (also referred to as the pixel domain) and the TUs may comprise coefficients in the transform domain following application of a transform, e.g., a discrete cosine transform (DCT), an integer transform, a wavelet transform, or a conceptually similar transform to residual video data. The residual data may correspond to pixel differences between pixels of the unencoded picture and prediction values corresponding to the PUs. Video encoder 20 may form the TUs including the residual data for the CU, and then transform the TUs to produce transform coefficients for the CU.

[0055] Following any transforms to produce transform coefficients, video encoder 20 may perform quantization of the transform coefficients. Quantization generally refers to a process in which transform coefficients are quantized to possibly reduce the amount of data used to represent the coefficients, providing further compression. The quantization process may reduce the bit depth associated with some or all of the coefficients. For example, an n -bit value may be rounded down to an m -bit value during quantization, where n is greater than m .

[0056] Following quantization, the video encoder may scan the transform coefficients, producing a one-dimensional vector from the two-dimensional matrix including the quantized transform coefficients. The scan may be designed to place higher energy (and therefore lower frequency) coefficients at the front of the array and to place lower energy (and therefore higher frequency) coefficients at the back of the array. In some examples, video encoder 20 may utilize a predefined scan order to scan the quantized transform coefficients to produce a serialized vector that can be entropy encoded. In other examples, video encoder 20 may perform an adaptive scan. After scanning the quantized transform coefficients to form a one-dimensional vector, video encoder 20

may entropy encode the one-dimensional vector, e.g., according to context-adaptive variable length coding (CAVLC), context-adaptive binary arithmetic coding (CABAC), syntax-based context-adaptive binary arithmetic coding (SBAC), Probability Interval Partitioning Entropy (PIPE) coding or another entropy encoding methodology. Video encoder 20 may also entropy encode syntax elements associated with the encoded video data for use by video decoder 30 in decoding the video data.

[0057] To perform CABAC, video encoder 20 may assign a context within a context model to a symbol to be transmitted. The context may relate to, for example, whether neighboring values of the symbol are non-zero or not. To perform CAVLC, video encoder 20 may select a variable length code for a symbol to be transmitted.

Codewords in VLC may be constructed such that relatively shorter codes correspond to more probable symbols, while longer codes correspond to less probable symbols. In this way, the use of VLC may achieve a bit savings over, for example, using equal-length codewords for each symbol to be transmitted. The probability determination may be based on a context assigned to the symbol.

[0058] Video encoder 20 may further send syntax data, such as block-based syntax data, frame-based syntax data, and GOP-based syntax data, to video decoder 30, e.g., in a frame header, a block header, a slice header, or a GOP header. The GOP syntax data may describe a number of frames in the respective GOP, and the frame syntax data may indicate an encoding/prediction mode used to encode the corresponding frame.

[0059] Coded video segments may be organized into NAL units, which provide a “network-friendly” video representation addressing applications such as video telephony, storage, broadcast, or streaming. NAL units can be categorized as Video Coding Layer (VCL) NAL units and non-VCL NAL units. VCL units may contain output from the core compression engine and may include block, macroblock, and/or slice level data. Other NAL units may be non-VCL NAL units. In some examples, a coded picture in one time instance, normally presented as a primary coded picture, may be contained in an access unit, which may include one or more NAL units.

[0060] As noted above, video encoder 20 and video decoder 30 may generally be configured to operate according to a video coding standard or an extension to a video coding standard, such as MV-HEVC or SHVC. Thus, for purposes of example, techniques of this disclosure are described below with respect to MV-HEVC, but it should be understood that these techniques may be applied to other video coding standards or extensions, such as SHVC.

[0061] In HEVC, there are a number of different picture types that can be identified by the NAL unit type. One picture type is a random access picture type, which may generally include IDR pictures, a clean random access (CRA) pictures, a temporal layer access (TLA) pictures, broken link access (BLA) pictures, STSA (Step-wise Temporal Sub-layer Access) pictures, leading pictures (i.e., RASL and RADL pictures noted below) or a number of other different picture types.

[0062] Video encoder 20 and video decoder 30 may assign each picture of a video sequence a POC value for identification of the picture. Video encoder 20 may encode (and video decoder 30 may decode) data representative of reference pictures to be included in the reference picture list, e.g., by signaling data representative of POC values for the reference pictures in a slice header and/or in a parameter set, such as a PPS or an SPS. A POC value, represented by the syntax element `PicOrderCntVal`, for a particular coded picture denotes the picture's relative order in the picture output process with respect to other pictures in the same coded video sequence.

[0063] A POC value includes least significant bits (LSB) and most significant bits (MSB). The POC may be obtained by concatenating or adding the MSB with the LSB. The LSB may be signaled in the slice header, and the MSB may be computed by video decoder 30 based on the NAL unit type of the current picture and the MSB and LSB of the previous picture in decoding order that is not of NAL unit type random access skipped leading (RASL) or random access decodable leading (RADL), or not a sub-layer non-reference picture, and that has a `temporal_id` value that is equal to 0. Such pictures that have `TemporalId` equal to 0 and `nuh_layer_id` equal to `nuh_layer_id` of the current picture and that are not RASL pictures, RADL pictures, or sub-layer non-reference pictures, are referred to as POC-anchor pictures.

[0064] When the current picture is an intra-random access point (IRAP) picture with `NoRaslOutputFlag` equal to 1, or a CRA picture that is the first picture of the bitstream, the value of POC MSB is inferred (i.e., determined by the coder) to be equal to 0. In a multi-layer bitstream (e.g., an SHVC or MV-HEVC bitstream with more than one layer), there may exist access units where one or more than one picture is an IRAP picture and one or more pictures are non-IRAP pictures. Such AUs may be referred to as non-aligned IRAP AUs. When decoding bit streams containing non-aligned IRAP AUs, it may be that the POC derived for the picture based on the POC LSB values signaled would violate the bitstream requirement that all the pictures (e.g., pictures of each layer) in an access unit should have the same value of `PicOrderCntVal`.

[0065] Accordingly, in some instances, POC values may be reset for a particular access unit in order to maintain POC alignment of the pictures in the access unit. In MV-HEVC WD5, a flag `poc_reset_flag` syntax element may be used to reset the POC of the pictures such that even when non-aligned IRAP AUs are present in the bitstream, the value of `PicOrderCntVal` of the current picture and the pictures in the DPB would be adjusted such that the POC of all the pictures in an AU are the same.

[0066] Chen et al., "CROSS-LAYER POC ALIGNMENT FOR MULTI-LAYER BITSTREAMS THAT MAY INCLUDE NON-ALIGNED IRAP PICTURES," U.S. Patent Application Serial No. 14/245,115, filed April 4, 2014, describes another method of achieving a POC reset, which uses two flags: `poc_msb_reset_flag` and `poc_lsb_reset_flag`. The former flag resets the MSB of `PicOrderCntVal`, and the latter flag resets the LSB of the `PicOrderCntVal`. Both these flags are signaled in the slice header.

[0067] U.S. Provisional No. 61/890,868, filed October 14, 2013 proposes an SEI message containing information for recovery of correct POC values when a picture that contains either POC MSB or POC resetting indication is lost.

[0068] Hannuksela et al., "MV-HEVC/SHVC HLS: on POC value derivation," Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T SG 16 WP 3 and ISO/IEC JTC 1/SC 29/WG 11, 15th Meeting: Geneva, CH, 23 Oct. – 1 Nov. 2013, JCTVC-O0275v3 (available at http://phenix.int-evry.fr/jct/doc_end_user/documents/15_Geneva/wg11/JCTVC-O0275-v3.zip) and Sjöberg et al., "HLS: Error robust POC alignment," Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T SG 16 WP 3 and ISO/IEC JTC 1/SC 29/WG 11, 15th Meeting: Geneva, CH, 23 Oct. – 1 Nov. 2013, JCTVC-O0176v3 (available at http://phenix.int-evry.fr/jct/doc_end_user/documents/15_Geneva/wg11/JCTVC-O0176-v3.zip), propose other methods for signaling and deriving POC values in multi-layer video coding.

[0069] U.S. Application No. 14/543,470, filed November 17, 2014 proposes to replace the two flags (i.e., `poc_msb_reset_flag` and `poc_lsb_reset_flag`) with a two-bit POC reset indicator and introduces the concept of POC resetting period wherein each POC resetting period is identified by a POC resetting period identifier, which is required to be different for two consecutive POC resetting periods. With such concepts, the error resiliency of POC reset mechanism is improved, such that it is possible for decoder to

keep cross-layer alignment of POC in cases of the loss of one or more pictures in the access unit where the POC reset is performed.

[0070] In general, a POC resetting period may be defined as follows, e.g., in subclause F.3 of MV-HEVC:

F.3.1 picture order count (POC) resetting period: A sequence of pictures in decoding order within a layer that all have poc_reset_idc greater than 0 and the same value of poc_reset_period_id.

Alternatively, the phrase “within a layer” in the definition above may be omitted, e.g., as follows:

F.3.2 picture order count (POC) resetting period: A sequence of pictures in decoding order that all have poc_reset_idc greater than 0 and the same value of poc_reset_period_id.

Thus, the POC resetting period is layer specific. In one example, MV-HEVC (or another applicable standard or standard extension) may require that pictures in a POC resetting period be continuous in decoding order within a layer. However, allowing these pictures to be interleaved with pictures with poc_reset_idc equal to 0 in the same layer may add flexibility and allow use of the same overhead to achieve better error resilience for clustered packet losses.

[0071] Video encoder 20 may assign slices to respective POC resetting periods. Each POC resetting period may include one or more slices. Thus, when POC values are reset for a POC resetting period (including resetting POC values of reference pictures preceding the POC resetting period in coding order), video encoder 20 may signal POC values of reference pictures to be included in a reference picture list based on the reset POC values.

[0072] Video encoder 20 and video decoder 30 may perform picture output and bumping of pictures stored to the DPB using POC values. The picture output and bumping process is the process to output pictures that are marked as “needed for output” and “unused for reference” from the DPB. After the adoption of Ramasubramonian et al., “MV-HEVC/SHVC HLS: Sub-DPB based DPB operations,” Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T SG 16 WP 3 and ISO/IEC JTC 1/SC 29/WG 11, 15th Meeting: Geneva, CH, 23 Oct. – 1 Nov. 2013, JCTVC-O0217 (available at http://phenix.int-evry.fr/jct/doc_end_user/documents/15_Geneva/wg11/JCTVC-O0217-v1.zip), the

concept for outputting and bumping pictures of the DPB for HEVC extensions has been updated to incorporate the following:

- Specify a separate DPB capacity for each layer so called sub-DPB
- There is no sharing of sub-DPB capacity across layers
- Each layer has its own parameters (max pictures, max latency, max reordering)
- Distinct parameters for each "output layer set"
- Change the definition of an operation point to be specific to an output layer set instead of a 'layer set'

After the changes, picture output and bumping process works as follows: for decoder conformance, the picture output process is invoked two times for each picture, including (1) before the decoding of the picture but after parsing the slice header of the first slice of the picture, and (2) when the last decoding unit of the picture is removed from the coded picture buffer. When the bumping process is invoked, the process outputs all pictures that belong to the selected access unit.

[0073] An example operation of the output order for a DPB, including picture output and bumping, after inclusion of the additions in JCTVC-O0217 and of Ramasubramonian et al., "MV-HEVC/SHVC HLS: On flushing of decoded pictures from DPB based on NoOutputOfPriorPicsFlag," Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T SG 16 WP 3 and ISO/IEC JTC 1/SC 29/WG 11, 15th Meeting: Geneva, CH, 23 Oct. – 1 Nov. 2013, O0266 (available at http://phenix.int-evry.fr/jct/doc_end_user/documents/15_Geneva/wg11/JCTVC-O0266-v1.zip) to MV-HEVC WD6 is as follows:

C.5.2 Operation of the output order DPB

C.5.2.1 General

The decoded picture buffer consists of sub-DPBs, and each sub-DPB contains picture storage buffers for storage of decoded pictures of one layer only. Each of the picture storage buffers of a sub-DPB contains a decoded picture that is marked as "used for reference" or is held for future output.

The process for output and removal of pictures from the DPB, the process for picture decoding, marking, additional bumping, and storage, and the "bumping" process are specified below. These processes are applied indendently for each layer, starting from the base layer, in increasing order of the nuh_layer_id values of the layers in the bitstream. When these processes are applied for a particular layer, only the sub-DPB for

the particular layer is affected except for the "bumping" process, which may crop and output pictures, mark pictures as "not needed for output" and empty picture storage buffers for any layer.

Let picture n and the current picture be the coded picture or decoded picture of the access unit n for a particular value of nuh_layer_id , wherein n is a non-negative integer number.

When these processes are applied for a layer with nuh_layer_id equal to $currLayerId$, the variables $MaxNumReorderPics$, $MaxLatencyIncreasePlus1$, $MaxLatencyPictures$, and $MaxDecPicBufferingMinus1$ are derived as follows:

- If a CVS conforming to one or more of the profiles specified in Annex G or H is decoded by applying the decoding process specified in clauses 2–10, Annex F, and Annex G or H, the following applies:
 - $MaxNumReorderPics$ is set equal to $max_vps_num_reorder_pics[TargetOptLayerSetIdx][HighestTid]$ of the active VPS.
 - $MaxLatencyIncreasePlus1$ is set equal to the value of the syntax element $max_vps_latency_increase_plus1[TargetOptLayerSetIdx][HighestTid]$ of the active VPS,.
 - $MaxLatencyPictures$ is set equal to $VpsMaxLatencyPictures[TargetOptLayerSetIdx][HighestTid]$ of the active VPS.
 - $MaxDecPicBufferingMinus1$ is set equal to the value of the syntax element $max_vps_dec_pic_buffering_minus1[TargetOptLayerSetIdx][currLayerId][HighestTid]$ of the active VPS.
- Otherwise (a CVS conforming to one or more of the profiles specified in Annex A is decoded by applying the decoding process specified in clauses 2–10), the following applies:
 - $MaxNumReorderPics$ is set equal to $sps_max_num_reorder_pics[HighestTid]$ of the active SPS for the base layer.
 - $MaxLatencyIncreasePlus1$ is set equal to $sps_max_latency_increase_plus1[HighestTid]$ of the active SPS for the base layer.
 - $MaxLatencyPictures$ is set equal to $SpsMaxLatencyPictures[HighestTid]$ of the active SPS for the base layer.

- MaxDecPicBufferingMinus1 is set equal to $\text{sps_max_dec_pic_buffering_minus1}[\text{HighestTid}]$ of the active SPS for the base layer.

C.5.2.2 Output and removal of pictures from the DPB

When the current picture is not picture 0 (i.e., the current picture does not have a POC value of 0) in the current layer, the output and removal of pictures in the current layer from the DPB before the decoding of the current picture, i.e. picture n , but after parsing the slice header of the first slice of the current picture, happens instantaneously when the first decoding unit of the current picture is removed from the CPB and proceeds as follows:

- The decoding process for RPS as specified in subclause 8.3.2 is invoked.
- If the current picture is an IRAP picture with NoRaslOutputFlag equal to 1, or the base layer picture in the current access unit is an IRAP picture with NoRaslOutputFlag equal to 1 and NoCirasOutputFlag is equal to 1, the following ordered steps are applied:
 1. The variable NoOutputOfPriorPicsFlag is derived for the decoder under test as follows:
 - If the current picture is a CRA picture with NoRaslOutputFlag equal to 1, NoOutputOfPriorPicsFlag is set equal to 1 (regardless of the value of $\text{no_output_of_prior_pics_flag}$).
 - Otherwise, if the current picture is an IRAP picture with NoRaslOutputFlag equal to 1 and the value of $\text{pic_width_in_luma_samples}$, $\text{pic_height_in_luma_samples}$, or $\text{sps_max_dec_pic_buffering_minus1}[\text{HighestTid}]$ derived from the active SPS for the current layer is different from the value of $\text{pic_width_in_luma_samples}$, $\text{pic_height_in_luma_samples}$, or $\text{sps_max_dec_pic_buffering_minus1}[\text{HighestTid}]$, respectively, derived from the SPS that was active for the current layer when decoding the preceding picture in the current layer, NoOutputOfPriorPicsFlag may (but should not) be set to 1 by the decoder under test, regardless of the value of $\text{no_output_of_prior_pics_flag}$.

NOTE – Although setting NoOutputOfPriorPicsFlag equal to $\text{no_output_of_prior_pics_flag}$ is preferred under these conditions, the

decoder under test is allowed to set NoOutputOfPriorPicsFlag to 1 in this case.

- Otherwise, if the current picture is an IRAP picture with NoRaslOutputFlag equal to 1, NoOutputOfPriorPicsFlag is set equal to no_output_of_prior_pics_flag.
 - Otherwise (the current picture is not an IRAP picture with NoRaslOutputFlag equal to 1, the base layer picture in the current access unit is an IRAP picture with NoRaslOutputFlag equal to 1, and NoCrasOutputFlag is equal to 1), NoOutputOfPriorPicsFlag is set equal to 1.
2. The value of NoOutputOfPriorPicsFlag derived for the decoder under test is applied for the HRD as follows:
- If NoOutputOfPriorPicsFlag is equal to 1, all picture storage buffers in the sub-DPB are emptied without output of the pictures they contain, and the sub-DPB fullness is set equal to 0.
 - Otherwise (NoOutputOfPriorPicsFlag is equal to 0), all picture storage buffers containing a picture that is marked as "not needed for output" and "unused for reference" are emptied (without output), and all non-empty picture storage buffers in the sub-DPB are emptied by repeatedly invoking the "bumping" process specified in subclause C.5.2.4, and the sub-DPB fullness is set equal to 0.
 - Otherwise, all picture storage buffers that contain a picture in the current layer and that are marked as "not needed for output" and "unused for reference" are emptied (without the pictures being output). For each picture storage buffer that is emptied, the sub-DPB fullness is decremented by one. When one or more of the following conditions are true, the "bumping" process specified in subclause C.5.2.4 is invoked repeatedly while further decrementing the sub-DPB fullness by one for each additional picture storage buffer that is emptied, until none of the following conditions are true:
 - The number of access units that contain at least one decoded picture in the DPB (i.e. any of the sub-DPBs) marked as "needed for output" is greater than MaxNumReorderPics.
 - MaxLatencyIncreasePlus1 is not equal to 0 and there is at least one access unit that contains at least one decoded picture in the DPB marked as "needed for

output" for which the associated variable PicLatencyCount is greater than or equal to MaxLatencyPictures.

- The number of pictures in the current layer in the sub-DPB of the current layer is greater than or equal to MaxDecPicBufferingMinus1 + 1.

C.5.2.3 Picture decoding, marking, additional bumping, and storage

The processes specified in this subclause happen instantaneously when the last decoding unit of picture n is removed from the CPB.

PicOutputFlag is updated as follows:

- If the current access unit does not contain a picture at a target output layer and alt_output_layer_flag is equal to 1, the following ordered steps apply:
 - The list nonOutputLayerPictures is the list of pictures of the access unit with PicOutputFlag equal to 1 and with nuh_layer_id values that are included in the TargetDecLayerIdList and that are not on target output layers.
 - The picture with the highest nuh_layer_id value among the list nonOutputLayerPictures is removed from the list nonOutputLayerPictures.
 - PicOutputFlag for each picture that is included in the list nonOutputLayerPictures is set equal to 0.
- Otherwise, PicOutputFlag for pictures that are not included in a target output layer is set equal to 0.

When the current picture has PicOutputFlag equal to 1, for each picture in the current layer in the sub-DPB that is marked as "needed for output" and follows the current picture in output order, the associated variable PicLatencyCount is set equal to PicLatencyCount + 1.

The current picture is considered as decoded after the last decoding unit of the picture is decoded. The current decoded picture is stored in an empty picture storage buffer in the sub-DPB, and the following applies:

- If the current decoded picture has PicOutputFlag equal to 1, it is marked as "needed for output" and its associated variable PicLatencyCount is set equal to 0.
- Otherwise (the current decoded picture has PicOutputFlag equal to 0), it is marked as "not needed for output".

The current decoded picture is marked as "used for short-term reference".

When one or more of the following conditions are true, the "bumping" process specified in subclause C.5.2.4 is invoked repeatedly until none of the following conditions are true:

- The number of that contain at least one decoded picture in the DPB (i.e. any of the sub-DPBs) marked as "needed for output" is greater than MaxNumReorderPics.
- MaxLatencyIncreasePlus1 is not equal to 0 and there is at least one access unit that contains at least one decoded picture in the DPB marked as "needed for output" for which the associated variable PicLatencyCount is greater than or equal to MaxLatencyPictures.

C.5.2.4 "Bumping" process

The "bumping" process consists of the following ordered steps:

1. The picture or pictures that are first for output are selected as the ones having the smallest value of PicOrderCntVal of all pictures in the DPB marked as "needed for output".
2. Each of these pictures is, in ascending nuh_layer_id order, cropped, using the conformance cropping window specified in the active SPS for the picture, the cropped picture is output, and the picture is marked as "not needed for output".
3. Each picture storage buffer that contains a picture marked as "unused for reference" and that was one of the pictures cropped and output is emptied.

[0074] In accordance with the techniques of this disclosure, video encoder 20 and video decoder 30 may be respectively configured to decrement POC values of all pictures (e.g., pictures of all layers) stored to a DPB prior to decoding an initial picture of an access unit in which POC resetting is performed. For example, video encoder 20 and video decoder 30 may apply a POC decrement to pictures that have been processed and stored to the DPB previously to the picture being coded (e.g., pictures that precede the picture currently being coded in coding order, which may generally be referred to as "earlier pictures" to the current picture). According to aspects of this disclosure, video encoder 20 and video decoder 30 may apply the POC decrement to such earlier pictures in a DPB (i.e., for all sub-DPBs) once, at the beginning of each access unit that requires POC reset. The POC decrement process may be invoked before coding (e.g., encoding or decoding) the initial picture for the access unit, but after (i.e., subsequent to) coding slice header information of the first slice of the initial picture. In this manner, POC values of pictures stored to the DPB of video encoder 20 and video decoder 30 may remain aligned when coding the pictures of the access unit, thereby allowing the pictures of the DPB to be output in the proper order and thus ensuring that the pictures may be displayed correctly. For example, by maintaining cross-layer alignment of pictures stored to the DPB (e.g., such that pictures of each respective access unit have

the same POC values), video decoder 30 may output all pictures from one access unit prior to outputting pictures of another access unit. This may allow the pictures of the access unit to be displayed substantially simultaneously to provide an appropriate viewing experience.

[0075] In another example, according to aspects of this disclosure, the timing of the POC decrementing may be altered. For example, video decoder 30 may invoke the POC decrementing process for pictures that have been processed (e.g., decoded) and stored to the DPB previous to the picture being decoded in each sub-DPB before decoding each picture in the access unit, but after the parsing and decoding of the slice header information of the first slice of the picture. Video decoder 30 may also perform an additional POC decrement process of earlier pictures in a sub-DPB for a layer having a picture that is not present or not received for the access unit after decoding the pictures that are present in the access unit.

[0076] In still another example, according to aspects of this disclosure, video decoder 30 may invoke the POC decrementing process for earlier pictures in each sub-DPB before the decoding of each picture in the access unit, but after the parsing and decoding of the slice header information of the first slice of the respective picture. Before the POC decrement process of earlier pictures in a sub-DPB is performed, if the same process has not yet been performed for one or more lower layer sub-DPBs, the process is performed for those sub-DPBs.

[0077] In still another example, according to aspects of this disclosure, video decoder 30 may perform the POC decrement process of earlier pictures in the DPB (i.e., for all sub-DPBs) is applied once, at the end of decoding each access unit that requires POC reset. That is, video decoder 30 may decode all pictures of an access unit for with POC resetting is performed, followed by decrementing the POC values of all pictures of the DPB.

[0078] According to aspects of this disclosure, video encoder 20 and video decoder may be respectively configured according to restrictions (e.g., bitstream constraints as determined and deemed applicable, by a coder, to a given bitstream) established by an applicable video coding standard, such as MV-HEVC. For example, according to aspects of this disclosure, the derived value for decrementing earlier pictures in a DPB (e.g., all sub-DPBs) shall be the same for all pictures in an access unit where POC or POC MSB resetting is performed.

[0079] As noted above, a layer-specific POC resetting period may be specified based on a POC resetting period identifier signaled in slice segment headers. That is, video encoder 20 and video decoder 30 may respectively code data representative of the POC resetting period identifier in slice segment headers. Each non-IRAP picture that belongs to an access unit that contains at least one IRAP picture may be the start of a POC resetting period in the layer containing the non-IRAP picture. That is, video encoder 20 may set a POC resetting type for a non-IRAP picture of an access unit containing at least one IRAP picture to indicate that the non-IRAP picture is the start of a new POC resetting period. In such an access unit, each picture would be the start of a POC resetting period in the layer containing the picture. POC resetting, either POC MSB only or both POC MSB and POC LSB, and update of POC values of same-layer pictures in the DPB may be applied only for the first picture within each POC resetting period.

[0080] Video encoder 20 may signal a POC LSB value in a slice segment header that video decoder 30 may use to derive POC values of pictures in a layer including the slice having the slice segment header. The POC LSB value signaled in the slice segment header may be used for derivation of the delta POC value, which is used for updating the POC values of the same-layer pictures in the DPB, and also for derivation of the POC MSB of the POC value of the current picture. When the current picture has such a POC LSB value signaled, and when the POC-anchor picture associated with the current picture is present in the bitstream, the POC-anchor picture has an indication of either POC resetting or POC MSB resetting. Such a POC LSB value signaled for the current picture is equal to the POC LSB value of the POC-anchor picture, which is also the first picture of the POC resetting period as the current picture. According to aspects of this disclosure, video encoder 20 signal data that indicates the additional POC LSB for base layer IDR pictures having a `poc_reset_idc` syntax element that is equal to 2.

[0081] As noted above, video encoder 20 and video decoder 30 may be respectively configured according to restrictions established by an applicable video coding standard. For example, according to aspects of this disclosure, all pictures in the same access unit shall have the same the same value of `poc_reset_idc`. In addition, as another example, one POC resetting period shall not include more than one access unit with `poc_reset_idc` equal to 1 or 2. In addition, as another example, an access unit with `poc_reset_idc` equal to 1 or 2 shall be the first access unit in a POC resetting period.

[0082] According to aspects of this disclosure, the POC resetting period may be applicable to all layers of the multi-layer video data. For example, a POC resetting period may be defined such that the POC resetting period is applied across layers of multi-layer data. As described in greater detail below, the POC resetting period may include a sequence of access units in decoding order that all have a `poc_reset_idc` that is greater than 0. The data indicating the POC value reset discussed above may be associated with a POC resetting period.

[0083] The following disclosure and tables include example syntax in accordance with the techniques of this disclosure. The syntax may be modified relative to, e.g., that of MV-HEVC as modified by documents JCTVC-O0217 and JCTVC-O0266 noted above. In the example description, syntax tables and semantics below, additions to MV-HEVC (as modified by documents JCTVC-O0217 and JCTVC-O0266) are represented using *italics* and deletions are represented using bracketed text preceded by “removed” (e.g., [removed: “removed text”]). In general, statements regarding “requirements” should be understood to form part of the text of the standard or standard extension, and not a requirement for purposes of the techniques of this disclosure. In some instances, such “requirements” may include bitstream constraints that may be determined to be applicable and then adhered to by, for example, a video coder based on the determination).

[0084] The following definition of a mathematical function may be added to MV-HEVC, e.g., to subclause 5.8 of MV-HEVC:

$$GetCurrMsb(cl, pl, pm, ml) = \begin{cases} pm + ml & ; \quad pl - cl \geq ml / 2 \\ pm - ml & ; \quad cl - pl > ml / 2 \\ pm & ; \quad otherwise \end{cases}$$

[0085] The following definition may be added to sub-clause F.3:

F.3.1 picture order count (POC) resetting period: A sequence of *access units* in decoding order that all have `poc_reset_idc` greater than 0 and the same value of `poc_reset_period_id`, *which may or may not be interleaved with access units that have `poc_reset_idc` equal to 0.*

[0086] In one alternative, pictures in a POC resetting period are required to be continuous in decoding order within a layer. However, allowing these pictures to be interleaved with pictures with `poc_reset_idc` equal to 0 in the same layer is more

flexible and allows using the same overhead to achieve better error resilience for clustered packet losses.

[0087] The following table is an example of syntax for a PPS in accordance with the techniques of this disclosure:

	Descriptor
pic_parameter_set_rbsp() {	
...	
pps_extension_flag	u(1)
if(pps_extension_flag) {	
poc_reset_info_present_flag	u(1)
pps_extension2_flag	u(1)
if(pps_extension2_flag)	
while(more_rbsp_data())	
pps_extension_data_flag	u(1)
}	
rbsp_trailing_bits()	
}	

[0088] In the example above, pps_extension_flag equal to 0 specifies that no pps_extension_data_flag syntax elements are present in the PPS RBSP syntax structure. When slice_segment_header_extension_present_flag is equal to 0, pps_extension_flag shall be equal to 0 in bit streams conforming to this version of this Specification, the value of 1 for pps_extension_flag is reserved for future use by ITU-T | ISO/IEC, and decoders shall ignore all data that follow the value 1 for pps_extension_flag in a PPS NAL unit.

[0089] In addition, poc_reset_info_present_flag equal to 0 specifies that the syntax element poc_reset_idc is not present in the slice segment headers of the slices referring to the PPS. In addition, poc_reset_info_present_flag equal to 1 specifies that the syntax element poc_reset_idc is present in the slice segment headers of the slices referring to the PPS.

[0090] In addition, pps_extension2_flag equal to 0 specifies that no pps_extension_data_flag syntax elements are present in the PPS RBSP syntax structure. pps_extension2_flag shall be equal to 0 in bit streams conforming to this version of this Specification. The value of 1 for pps_extension2_flag is reserved for future use by ITU-T | ISO/IEC. Decoders shall ignore all data that follow the value 1 for pps_extension2_flag in a PPS NAL unit.

[0091] The following table is an example of syntax for a slice header in accordance with the techniques of this disclosure:

	Descriptor
slice_segment_header() {	
...	
if(!dependent_slice_segment_flag) {	
i = 0	
if(num_extra_slice_header_bits > i) {	
i++	
discardable_flag	u(1)
}	
if(num_extra_slice_header_bits > i) {	
i++	
cross_layer_bla_flag	u(1)
}	
...	
if(slice_segment_header_extension_present_flag) {	
slice_segment_header_extension_length	ue(v)
if(poc_reset_info_present_flag)	
poc_reset_idc	u(2)
if(poc_reset_idc > 0)	
poc_reset_period_id	u(6)
if(poc_reset_idc == 3)	
full_poc_reset_flag	u(1)
if((poc_reset_idc == 3) (nuh_layer_id == 0 && poc_reset_idc == 2))	
poc_lsb_val	u(v)
if((poc_reset_info_present_flag && poc_reset_idc == 0) (poc_reset_idc == 3 && (log2_max_pic_order_cnt_lsb_minus4 + 5) % 8 != 0) (nuh_layer_id == 0 && poc_reset_idc == 2 && (log2_max_pic_order_cnt_lsb_minus4 + 4) % 8 != 0))	
slice_segment_header_extension_reserved_bits	u(v)
for(i = BytesInSliceSegmtHdrExt; i < slice_segment_header_extension_length; i++)	
slice_segment_header_extension_data_byte[i]	u(8)
}	
byte_alignment()	
}	

[0092] Alternatively, the poc_reset_period_id syntax element may be signaled using a different number of bits, e.g. coded as u(14).

[0093] When present, the value of the slice segment header syntax elements slice_pic_parameter_set_id, pic_output_flag, no_output_of_prior_pics_flag, slice_pic_order_cnt_lsb, short_term_ref_pic_set_sps_flag, short_term_ref_pic_set_idx, num_long_term_sps, num_long_term_pics, slice_temporal_mvp_enabled_flag, discardable_flag, cross_layer_bla_flag, inter_layer_pred_enabled_flag, num_inter_layer_ref_pics_minus1, poc_reset_idc, poc_reset_pic_id, full_poc_reset_flag, and poc_lsb_val shall be the same in all slice segment headers of a coded picture. When present, the value of the slice segment header syntax elements lt_idx_sps[i], poc_lsb_lt[i], used_by_curr_pic_lt_flag[i], delta_poc_msb_present_flag[i], delta_poc_msb_cycle_lt[i], and inter_layer_pred_layer_idc[i] shall be the same in all slice segment headers of a coded picture for each possible value of i.

[0094] The syntax element slice_segment_header_extension_length may specify the length of the slice segment header extension data in bytes, not including the bits used for signalling slice_segment_header_extension_length itself. If poc_reset_info_present_flag is equal to 0, the value of the syntax element slice_segment_header_extension_length shall be in the range of 0 to 256, inclusive. Otherwise, the value of the syntax element slice_segment_header_extension_length shall be in the range of 1 to 256, inclusive, when poc_reset_idc is equal to 1 or when poc_reset_idc is equal to 2 and nuh_layer_id is greater than 0, and in the range of $1 + \text{Ceil}((\log_2 \text{max_pic_order_cnt_lsb_minus4} + 4) / 8)$ to 256, inclusive, when poc_reset_idc is equal to 2 and nuh_layer_id is equal to 0, and in the range of $1 + \text{Ceil}((\log_2 \text{max_pic_order_cnt_lsb_minus4} + 5) / 8)$ to 256, inclusive, when poc_reset_idc is equal to 3.

[0095] The syntax element poc_reset_idc equal to 0 specifies that neither the most significant bits nor the least significant bits of the picture order count value for the current picture are reset. The syntax element poc_reset_idc equal to 1 specifies that only the most significant bits of the picture order count value for the current picture may be reset. The syntax element poc_reset_idc equal to 2 specifies that both the most significant bits and the least significant bits of the picture order count value for the current picture may be reset. The syntax element poc_reset_idc equal to 3 specifies that either only the most significant bits or both the most significant bits and the least

significant bits of the picture order count value for the current picture may be reset and additional picture order count information is signaled. When not present, the value of the syntax element `poc_reset_poc` is inferred (i.e., determined, by a coder, without explicit signaling) to be equal to 0.

[0096] It is a requirement of bitstream conformance that the following constraints apply:

- The value of `poc_reset_idc` shall not be equal to 1 or 2 for a RASL picture, a RADL picture, a sub-layer non-reference picture, or a picture that has `TemporalId` greater than 0.

- *The value of `poc_reset_idc` of all pictures in an access unit shall be the same.*

[removed: – When the value of `poc_reset_idc` of a picture in an access unit is equal to 1, the value of `poc_reset_idc` of all pictures in the access unit shall be equal to 1.]

[removed: – When the value of `poc_reset_idc` of a picture in an access unit is equal to 2, the value of `poc_reset_idc` of all pictures in the access unit shall be equal to 2.]

- When one picture in an access unit is an IRAP picture with a particular value of `nal_unit_type` and there is at least one other picture in the same access unit with a different value of `nal_unit_type`, the value of `poc_reset_idc` shall be equal to 1 or 2 for all pictures in the access unit.

- When the picture with `nuh_layer_id` equal to 0 in an access unit is an IDR picture and there is at least one non-IDR picture in the same access unit, the value of `poc_reset_idc` shall be equal to 2 for all pictures in the access unit.

- When the picture with `nuh_layer_id` equal to 0 in an access unit is not an IDR picture, the value of `poc_reset_idc` shall not be equal to 2 for any picture in the access unit.

[0097] *The value of `poc_reset_idc` of an access unit refers to the value of `poc_reset_idc` of the pictures in the access unit.*

[0098] The syntax element `poc_reset_period_id` identifies a POC resetting period. There shall be no two pictures consecutive in decoding order in the same layer that have the same value of the syntax element `poc_reset_period_id` and the syntax element `poc_reset_idc` equal to 1 or 2.

NOTE – It is not prohibited for multiple pictures in a layer to have the same value of `poc_reset_pic_id` and to have `poc_reset_idc` equal to 1 or 2 unless such pictures occur in two consecutive access units in decoding order. To minimize the likelihood of such two pictures appearing in the bitstream due to picture losses, bitstream

extraction, seeking, or splicing operations, encoders should set the value of `poc_reset_pic_id` to be a random value for each POC resetting period (subject to the constraints specified above).

[0099] *It is a requirement of bitstream conformance that the following constraints apply:*

- *One POC resetting period shall not include more than one access unit with `poc_reset_idc` equal to 1 or 2.*
- *An access unit with `poc_reset_idc` equal to 1 or 2 shall be the first access unit in a POC resetting period.*

[removed: -- When pictures with `poc_reset_idc` equal to 1 and 3 are present in the same POC resetting period, all pictures with `poc_reset_idc` equal 3 shall follow the picture with `poc_reset_idc` equal to 1 in decoding order.]

[removed: -- When pictures with `poc_reset_idc` equal to 2 and 3 are present in the same POC resetting period, all pictures with `poc_reset_idc` equal 3 shall follow the picture with `poc_reset_idc` equal to 2 in decoding order.]

[0100] The syntax element `full_poc_reset_flag` equal to 1 specifies that both the most significant bits and the least significant bits of the picture order count value for the current picture are reset when the previous picture in decoding order in the same layer does not belong to the same POC resetting period. The syntax element `full_poc_reset_flag` equal to 0 specifies that only the most significant bits of the picture order count value for the current picture are reset when the previous picture in decoding order in the same layer does not belong to the same POC resetting period.

[0101] The syntax element `poc_lsb_val` specifies a value that may be used to derive the picture order count of the current picture. The length of the `poc_lsb_val` syntax element is $\log_2_max_pic_order_cnt_lsb_minus4 + 4$ bits.

[0102] It is a requirement of bitstream conformance that, when `poc_reset_idc` is equal to 3, and the previous picture `picA` in decoding order that is in the same layer as the current picture, has `poc_reset_idc` equal to 1 or 2, and belongs to the same POC resetting period is present in the bitstream, `picA` shall be the same picture as the previous picture in decoding order that is in the same layer as the current picture, that is not a RASL picture, a RADL picture or a sub-layer non-reference picture, and that has `TemporalId` equal to 0, and the value of `poc_lsb_val` of the current picture shall be equal to the value of `slice_pic_order_cnt_lsb` of `picA`.

[0103] The variables numRsvBits and BytesInSliceSegmtHdrExt are derived as follows:

```

if( !poc_reset_info_present_flag ) {
    numRsvBits = 0
    BytesInSliceSegmtHdrExt = 0
} else if( poc_reset_idc == 0 ) {
    numRsvBits = 6
    BytesInSliceSegmtHdrExt = 1
}
else if( ( poc_reset_idc == 2 && nuh_layer_id > 0 ) || poc_reset_idc == 1 ) {
    numRsvBits = 0
    BytesInSliceSegmtHdrExt = 1
}
else if( poc_reset_idc == 2 && nuh_layer_id == 0 ) {
    numRsvBits = 8 - ( log2_max_pic_order_cnt_lsb_minus4 + 4 ) % 8
    bitsInSliceSegmtHdrExt = 1 + Ceil( ( log2_max_pic_order_cnt_lsb_minus4 + 4
) / 8 )
}
else { // poc_reset_idc == 3
    numRsvBits = 8 - ( log2_max_pic_order_cnt_lsb_minus4 + 5 ) % 8
    bitsInSliceSegmtHdrExt = 1 + Ceil( ( log2_max_pic_order_cnt_lsb_minus4 + 5
) / 8 )
}

```

[0104] The syntax element slice_segment_header_extension_reserved_bits may have any value. The length of the slice_segment_header_extension_reserved_bits syntax element is equal to numRsvBits bits. Decoders, such as video decoder 30, shall ignore the value of slice_segment_header_extension_reserved_bits. Its value does not affect decoder conformance to profiles specified in this version of this Specification.

[0105] Video decoder 30 may perform the following decoding process for starting the decoding of a coded picture with nuh_layer_id greater than 0. Each “picture” referred to in the process below is a complete coded picture. The decoding process operates as follows for the current picture CurrPic:

- Decode NAL units as described herein.

- The processes in subclause F.8.3 specify the following decoding processes using syntax elements in the slice segment layer and above:
 - Variables and functions relating to picture order count are derived in subclause F.8.3.1. This needs to be invoked only for the first slice segment of a picture. It is a requirement of bitstream conformance that `PicOrderCntVal` shall remain unchanged within an access unit. *It is also a requirement of bitstream conformance that `DeltaPocVal` shall remain unchanged within an access unit.*
 - The decoding process for RPS in subclause F.8.3.2 is invoked, wherein only reference pictures with a `nuh_layer_id` equal to that of `CurrPic` may be marked as "unused for reference" or "used for long-term reference" and any picture with a different value of `nuh_layer_id` is not marked. This needs to be invoked only for the first slice segment of a picture.
 - When `FirstPicInLayerDecodedFlag[nuh_layer_id]` is equal to 0, the decoding process for generating unavailable reference pictures specified in subclause F.8.1.3 is invoked, which needs to be invoked only for the first slice segment of a picture.
 - When `FirstPicInLayerDecodedFlag[nuh_layer_id]` is not equal to 0 and the current picture is an IRAP picture with `NoRaslOutputFlag` equal to 1, the decoding process for generating unavailable reference pictures specified in subclause F.8.1.3 is invoked, which needs to be invoked only for the first slice segment of a picture.

[0106] In a first example, according to aspects of this disclosure, video decoder 30 may perform the following decoding process for POC values. The first example may correspond to the first example technique for decrementing POC values described above. That is, the example below may be used to decrement POC values of all pictures (of all layers) stored to a DPB prior to decoding an initial picture of an access unit in which POC resetting is performed. The POC decrement process may be invoked before decoding the initial picture for the access unit, but after coding slice header information of the first slice of the initial picture. The first example process may be performed as follows:

Output of this process is `PicOrderCntVal`, the picture order count of the current picture.

Picture order counts are used to identify pictures, for deriving motion parameters in merge mode and motion vector prediction, and for decoder conformance checking (see subclause C.5).

Each coded picture is associated with a picture order count variable, denoted as `PicOrderCntVal`.

If `FirstPicInLayerDecodedFlag[nuh_layer_id]` is equal to 1, `poc_reset_idc` is greater than 0, and the *access unit containing the current picture is the first access unit in decoding order in a POC resetting period*, the following applies:

- The variables `pocMsbDelta`, `pocLsbDelta` and `DeltaPocVal` are derived as follows:

```

prevPicOrderCntLsb =
PrevPicOrderCnt[ nuh_layer_id ] & ( MaxPicOrderCntLsb - 1 )
prevPicOrderCntMsb = PrevPicOrderCnt[ nuh_layer_id ] - prevPicOrderCntLsb
if( poc_reset_idc == 3 || ( poc_reset_idc == 2 && nuh_layer_id == 0 ) )
    pocLsbVal = poc_lsb_val
else
    pocLsbVal = slice_pic_order_cnt_lsb
pocMsbDelta = getCurrMsb( pocLsbVal, prevPicOrderCntLsb,
prevPicOrderCntMsb, MaxPicOrderCntLsb )
if( poc_reset_idc == 2 || ( poc_reset_idc == 3 && full_poc_reset_flag ) )
    pocLsbDelta = pocLsbVal
else
    pocLsbDelta = 0
DeltaPocVal = pocMsbDelta + pocLsbDelta

```

- The `PicOrderCntVal` of the current picture is derived as follows:

```

if( poc_reset_idc == 1 )
    PicOrderCntVal = slice_pic_order_cnt_lsb
else if( poc_reset_idc == 2 )
    PicOrderCntVal = 0
else { // poc_reset_idc == 3
    PicOrderCntMsb = getCurrMsb( slice_pic_order_cnt_lsb, poc_lsb_val, 0,
MaxPicOrderCntLsb )
    PicOrderCntVal = PicOrderCntMsb + slice_pic_order_cnt_lsb
}

```

- The value of `PrevPicOrderCnt[nuh_layer_id]` is derived as follows:

- If the current picture is not a RASL picture, a RADL picture, or a sub-layer non-reference picture, and the current picture has TemporalId equal to 0, PrevPicOrderCnt[nuh_layer_id] is set equal to PicOrderCntVal.
- Otherwise when poc_reset_idc is equal to 3, PrevPicOrderCnt[nuh_layer_id] is set equal to full_poc_reset_flag ? 0 : poc_lsb_val.

Otherwise, the following applies:

- *The value of DeltaPocVal is set equal to 0.*
- The PicOrderCntVal of the current picture is derived as follows:

```

if( !FirstPicInLayerDecodedFlag[ nuh_layer_id ] ) {
    if( poc_reset_idc == 1 )
        PicOrderCntVal = slice_pic_order_cnt_lsb
    else if( poc_reset_idc == 2 )
        PicOrderCntVal = 0
    else if( poc_reset_idc == 3 ) {
        PicOrderCntMsb = getCurrMsb(slice_pic_order_cnt_lsb, poc_lsb_val, 0,
MaxPicOrderCntLsb )
        PicOrderCntVal = PicOrderCntMsb + slice_pic_order_cnt_lsb
    } else // the current picture is an IRAP picture with NoRaslOutputFlag equal to
1
        PicOrderCntVal = slice_pic_order_cnt_lsb
} else { // the POC derivation as in HEVC version 1
    if( the current picture is an IRAP picture with NoRaslOutputFlag equal to 1 )
        PicOrderCntMsb = 0
    else {
        prevPicOrderCntLsb =
PrevPicOrderCnt[ nuh_layer_id ] & ( MaxPicOrderCntLsb - 1 )
        prevPicOrderCntMsb =
PrevPicOrderCnt[ nuh_layer_id ] - prevPicOrderCntLsb
        PicOrderCntMsb = getCurrMsb( slice_pic_order_cnt_lsb,
prevPicOrderCntLsb,
prevPicOrderCntMsb, MaxPicOrderCntLsb )
    }
    PicOrderCntVal = PicOrderCntMsb + slice_pic_order_cnt_lsb
}

```

- The value of $\text{PrevPicOrderCnt}[\text{nuh_layer_id}]$ is derived as follows:
 - If the current picture is not a RASL picture, a RADL picture, or a sub-layer non-reference picture, and the current picture has TemporalId equal to 0, $\text{PrevPicOrderCnt}[\text{nuh_layer_id}]$ is set equal to PicOrderCntVal .
 - Otherwise when $\text{FirstPicInLayerDecodedFlag}[\text{nuh_layer_id}]$ is equal to 0 and poc_reset_idc is equal to 3, $\text{PrevPicOrderCnt}[\text{nuh_layer_id}]$ is set equal to $\text{full_poc_reset_flag} ? 0 : \text{poc_lsb_val}$.

If the current picture is the first picture in an access unit (i.e. the picture with the lowest value of nuh_layer_id among all pictures in the access unit) and DeltaPocVal is greater than 0, the PicOrderCntVal values of all pictures in all sub-DPBs are decremented by the value of DeltaPocVal .

The value of PicOrderCntVal shall be in the range of -2^{31} to $2^{31} - 1$, inclusive. In one CVS, the PicOrderCntVal values for any two coded pictures in the same layer shall not be the same.

The function $\text{PicOrderCnt}(\text{picX})$ is specified as follows:

$$\text{PicOrderCnt}(\text{picX}) = \text{PicOrderCntVal of the picture picX} \quad (\text{F-25})$$

The function $\text{DiffPicOrderCnt}(\text{picA}, \text{picB})$ is specified as follows:

$$\text{DiffPicOrderCnt}(\text{picA}, \text{picB}) = \text{PicOrderCnt}(\text{picA}) - \text{PicOrderCnt}(\text{picB}) \quad (\text{F-26})$$

The bitstream shall not contain data that result in values of $\text{DiffPicOrderCnt}(\text{picA}, \text{picB})$ used in the decoding process that are not in the range of -2^{15} to $2^{15} - 1$, inclusive.

NOTE – Let X be the current picture and Y and Z be two other pictures in the same sequence, Y and Z are considered to be in the same output order direction from X when both $\text{DiffPicOrderCnt}(X, Y)$ and $\text{DiffPicOrderCnt}(X, Z)$ are positive or both are negative.

[0107] In a second example, according to aspects of this disclosure, video decoder 30 may perform the following decoding process for POC values. The second example may correspond to the second example technique for decrementing POC values described above. That is, the example below may be used to invoke the POC decrementing process for earlier pictures in each sub-DPB before decoding each picture in the access unit, but after the parsing and decoding of the slice header information of the first slice of the picture. Video decoder 30 may also perform an additional POC decrement process of earlier pictures in a sub-DPB for a layer having a picture that is not present or

not received for the access unit after decoding the pictures that are present in the access unit. The second example process may be performed as follows:

Output of this process is `PicOrderCntVal`, the picture order count of the current picture.

Picture order counts are used to identify pictures, for deriving motion parameters in merge mode and motion vector prediction, and for decoder conformance checking (see subclause C.5).

Each coded picture is associated with a picture order count variable, denoted as `PicOrderCntVal`.

If current picture is the first picture in the access unit, the following applies:

- The value of `PocResetFlag` is set to 0.
- The variable `DeltaPocVal` is set to 0.
- The variable `UpdateSubDpbFlag[i]` is set equal to 0 for all values of `i` from 0 to 63, inclusive.

If `FirstPicInLayerDecodedFlag[nuh_layer_id]` is equal to 1, `poc_reset_idc` is greater than 0, and the access unit containing the current picture is the first access unit in decoding order in a POC resetting period, the following applies:

If `FirstPicInLayerDecodedFlag[nuh_layer_id]` is equal to 1, `poc_reset_idc` is greater than 0, and the access unit containing the current picture is the first access unit in decoding order in a POC resetting period, the following applies:

- The variables `pocMsbDelta`, `pocLsbDelta` and `DeltaPocVal` are derived as follows:


```

prevPicOrderCntLsb =
PrevPicOrderCnt[ nuh_layer_id ] & ( MaxPicOrderCntLsb - 1 )
prevPicOrderCntMsb = PrevPicOrderCnt[ nuh_layer_id ] - prevPicOrderCntLsb
if( poc_reset_idc == 3 || ( poc_reset_idc == 2 && nuh_layer_id == 0 ) )
    pocLsbVal = poc_lsb_val
else
    pocLsbVal = slice_pic_order_cnt_lsb
pocMsbDelta = getCurrMsb( pocLsbVal, prevPicOrderCntLsb,
prevPicOrderCntMsb, MaxPicOrderCntLsb )
if( poc_reset_idc == 2 || ( poc_reset_idc == 3 && full_poc_reset_flag ) )
    pocLsbDelta = pocLsbVal
else

```

- pocLsbDelta = 0
- DeltaPocVal* = *pocMsbDelta* + *pocLsbDelta*
- The *PicOrderCntVal* of each picture that is in the DPB and belongs to the same layer as the current picture is decremented by *DeltaPocVal*.
 - *PocResetFlag* = 1
 - *UpdateSubDpbFlag[nuh_layer_id]* = 1
 - The *PicOrderCntVal* of the current picture is derived as follows:
 - if(*poc_reset_idc* == 1)
 - PicOrderCntVal* = *slice_pic_order_cnt_lsb*
 - else if(*poc_reset_idc* == 2)
 - PicOrderCntVal* = 0
 - else { // *poc_reset_idc* == 3
 - PicOrderCntMsb* = *getCurrMsb(slice_pic_order_cnt_lsb, poc_lsb_val, 0, MaxPicOrderCntLsb)*
 - PicOrderCntVal* = *PicOrderCntMsb* + *slice_pic_order_cnt_lsb*
 - The value of *PrevPicOrderCnt[nuh_layer_id]* is derived as follows:
 - If the current picture is not a RASL picture, a RADL picture, or a sub-layer non-reference picture, and the current picture has *TemporalId* equal to 0, *PrevPicOrderCnt[nuh_layer_id]* is set equal to *PicOrderCntVal*.
 - Otherwise when *poc_reset_idc* is equal to 3, *PrevPicOrderCnt[nuh_layer_id]* is set equal to *full_poc_reset_flag ? 0 : poc_lsb_val*.

Otherwise, the following applies:

- The value of *DeltaPocVal* is set equal to 0
- The *PicOrderCntVal* of the current picture is derived as follows:
 - if(!*FirstPicInLayerDecodedFlag[nuh_layer_id]*) {
 - if(*poc_reset_idc* == 1)
 - PicOrderCntVal* = *slice_pic_order_cnt_lsb*
 - else if(*poc_reset_idc* == 2)
 - PicOrderCntVal* = 0
 - else if(*poc_reset_idc* == 3) {
 - PicOrderCntMsb* = *getCurrMsb(slice_pic_order_cnt_lsb, poc_lsb_val, 0, MaxPicOrderCntLsb)*
 - PicOrderCntVal* = *PicOrderCntMsb* + *slice_pic_order_cnt_lsb*

```

    } else // the current picture is an IRAP picture with NoRaslOutputFlag equal to
1
    PicOrderCntVal = slice_pic_order_cnt_lsb
} else { // the POC derivation as in HEVC version 1
    if( the current picture is an IRAP picture with NoRaslOutputFlag equal to 1 )
        PicOrderCntMsb = 0
    else {
        prevPicOrderCntLsb =
PrevPicOrderCnt[ nuh_layer_id ] & ( MaxPicOrderCntLsb - 1 )
        prevPicOrderCntMsb =
PrevPicOrderCnt[ nuh_layer_id ] - prevPicOrderCntLsb
        PicOrderCntMsb = getCurrMsb( slice_pic_order_cnt_lsb,
prevPicOrderCntLsb,
                                prevPicOrderCntMsb, MaxPicOrderCntLsb )
    }
    PicOrderCntVal = PicOrderCntMsb + slice_pic_order_cnt_lsb
}

```

- The value of PrevPicOrderCnt[nuh_layer_id] is derived as follows:
 - If the current picture is not a RASL picture, a RADL picture, or a sub-layer non-reference picture, and the current picture has TemporalId equal to 0, PrevPicOrderCnt[nuh_layer_id] is set equal to PicOrderCntVal.
 - Otherwise when FirstPicInLayerDecodedFlag[nuh_layer_id] is equal to 0 and poc_reset_idc is equal to 3, PrevPicOrderCnt[nuh_layer_id] is set equal to full_poc_reset_flag ? 0 : poc_lsb_val.

If the current picture is the last picture in an access unit and the value of variable PocResetFlag is equal to 1, the following applies:

- *For all values of i from 0 to 63, inclusive, if UpdateSubDpbFlag[i] is equal to 0, the PicOrderCntVal of all pictures in the DPB that have nuh_layer_id equal to i is decremented by the value of variable DeltaPocVal*

The value of PicOrderCntVal shall be in the range of -2^{31} to $2^{31} - 1$, inclusive.

In one CVS, the PicOrderCntVal values for any two coded pictures in the same layer shall not be the same.

The function PicOrderCnt(picX) is specified as follows:

$$\text{PicOrderCnt}(\text{picX}) = \text{PicOrderCntVal of the picture picX} \quad (\text{F-25})$$

The function $\text{DiffPicOrderCnt}(\text{picA}, \text{picB})$ is specified as follows:

$$\text{DiffPicOrderCnt}(\text{picA}, \text{picB}) = \text{PicOrderCnt}(\text{picA}) - \text{PicOrderCnt}(\text{picB}) \quad (\text{F-26})$$

The bitstream shall not contain data that result in values of $\text{DiffPicOrderCnt}(\text{picA}, \text{picB})$ used in the decoding process that are not in the range of -2^{15} to $2^{15} - 1$, inclusive.

NOTE – Let X be the current picture and Y and Z be two other pictures in the same sequence, Y and Z are considered to be in the same output order direction from X when both $\text{DiffPicOrderCnt}(X, Y)$ and $\text{DiffPicOrderCnt}(X, Z)$ are positive or both are negative.

[0108] In a third example, according to aspects of this disclosure, video decoder 30 may perform the following decoding process for POC values. The third example may correspond to the third example technique for decrementing POC values described above. That is, the example below may be used to invoke the POC decrementing process for earlier pictures in each sub-DPB before the decoding of each picture in the access unit, but after the parsing and decoding of the slice header information of the first slice of the respective picture. Before the POC decrement process of earlier pictures in a sub-DPB is performed, if the same process has not yet been performed for one or more lower layer sub-DPBs, the process is performed for those sub-DPBs. The third example process may be performed as follows:

Output of this process is PicOrderCntVal , the picture order count of the current picture.

Picture order counts are used to identify pictures, for deriving motion parameters in merge mode and motion vector prediction, and for decoder conformance checking (see subclause C.5).

Each coded picture is associated with a picture order count variable, denoted as PicOrderCntVal .

If current picture is the first picture in the access unit, the following applies:

- *The value of PocResetFlag is set to 0.*
- *The variable DeltaPocVal is set to 0.*
- *The variable $\text{UpdateSubDpbFlag}[i]$ is set equal to 0 for all values of i from 0 to 63, inclusive.*

If `FirstPicInLayerDecodedFlag[nuh_layer_id]` is equal to 1, `poc_reset_idc` is greater than 0, and the *access unit containing the current picture is the first access unit in decoding order in a POC resetting period*, the following applies:

- The variables `pocMsbDelta`, `pocLsbDelta` and *DeltaPocVal* are derived as follows:


```

prevPicOrderCntLsb =
PrevPicOrderCnt[ nuh_layer_id ] & ( MaxPicOrderCntLsb - 1 )
prevPicOrderCntMsb = PrevPicOrderCnt[ nuh_layer_id ] - prevPicOrderCntLsb
if( poc_reset_idc == 3 || ( poc_reset_idc == 2 && nuh_layer_id == 0 ) )
    pocLsbVal = poc_lsb_val
else
    pocLsbVal = slice_pic_order_cnt_lsb
pocMsbDelta = getCurrMsb( pocLsbVal, prevPicOrderCntLsb,
prevPicOrderCntMsb, MaxPicOrderCntLsb )
if( poc_reset_idc == 2 || ( poc_reset_idc == 3 && full_poc_reset_flag ) )
    pocLsbDelta = pocLsbVal
else
    pocLsbDelta = 0
DeltaPocVal = pocMsbDelta + pocLsbDelta

```
- *PocResetFlag = 1*
- *UpdateSubDpbFlag[nuh_layer_id] = 1*
- *For all values of i from 0 to nuh_layer_id - 1, inclusive, if UpdateSubDpbFlag[i] is equal to 0, then the PicOrderCntVal of each picture that is in the DPB and belongs to the same layer as the current picture is decremented by DeltaPocVal.*
- The `PicOrderCntVal` of each picture that is in the DPB and belongs to the same layer as the current picture is decremented by *DeltaPocVal*.
- The `PicOrderCntVal` of the current picture is derived as follows:


```

if( poc_reset_idc == 1 )
    PicOrderCntVal = slice_pic_order_cnt_lsb
else if( poc_reset_idc == 2 )
    PicOrderCntVal = 0
else { // poc_reset_idc == 3
    PicOrderCntMsb = getCurrMsb( slice_pic_order_cnt_lsb, poc_lsb_val, 0,
MaxPicOrderCntLsb )

```

```

    PicOrderCntVal = PicOrderCntMsb + slice_pic_order_cnt_lsb
  }

```

- The value of PrevPicOrderCnt[nuh_layer_id] is derived as follows:
 - If the current picture is not a RASL picture, a RADL picture, or a sub-layer non-reference picture, and the current picture has TemporalId equal to 0, PrevPicOrderCnt[nuh_layer_id] is set equal to PicOrderCntVal.
 - Otherwise when poc_reset_idc is equal to 3, PrevPicOrderCnt[nuh_layer_id] is set equal to full_poc_reset_flag ? 0 : poc_lsb_val.

Otherwise, the following applies:

- The value of DeltaPocVal is set equal to 0
- The PicOrderCntVal of the current picture is derived as follows:


```

if( !FirstPicInLayerDecodedFlag[ nuh_layer_id ] ) {
  if( poc_reset_idc == 1 )
    PicOrderCntVal = slice_pic_order_cnt_lsb
  else if( poc_reset_idc == 2 )
    PicOrderCntVal = 0
  else if( poc_reset_idc == 3 ) {
    PicOrderCntMsb = getCurrMsb(slice_pic_order_cnt_lsb, poc_lsb_val, 0,
MaxPicOrderCntLsb )
    PicOrderCntVal = PicOrderCntMsb + slice_pic_order_cnt_lsb
  } else // the current picture is an IRAP picture with NoRaslOutputFlag equal to
1
    PicOrderCntVal = slice_pic_order_cnt_lsb
} else { // the POC derivation as in HEVC version 1
  if( the current picture is an IRAP picture with NoRaslOutputFlag equal to 1 )
    PicOrderCntMsb = 0
  else {
    prevPicOrderCntLsb =
PrevPicOrderCnt[ nuh_layer_id ] & ( MaxPicOrderCntLsb - 1 )
    prevPicOrderCntMsb =
PrevPicOrderCnt[ nuh_layer_id ] - prevPicOrderCntLsb
    PicOrderCntMsb = getCurrMsb( slice_pic_order_cnt_lsb,
prevPicOrderCntLsb,
prevPicOrderCntMsb, MaxPicOrderCntLsb )

```

```

    }
    PicOrderCntVal = PicOrderCntMsb + slice_pic_order_cnt_lsb
  }

```

- The value of PrevPicOrderCnt[nuh_layer_id] is derived as follows:
 - If the current picture is not a RASL picture, a RADL picture, or a sub-layer non-reference picture, and the current picture has TemporalId equal to 0, PrevPicOrderCnt[nuh_layer_id] is set equal to PicOrderCntVal.
 - Otherwise when FirstPicInLayerDecodedFlag[nuh_layer_id] is equal to 0 and poc_reset_idc is equal to 3, PrevPicOrderCnt[nuh_layer_id] is set equal to full_poc_reset_flag ? 0 : poc_lsb_val.

The value of PicOrderCntVal shall be in the range of -2^{31} to $2^{31} - 1$, inclusive.

In one CVS, the PicOrderCntVal values for any two coded pictures in the same layer shall not be the same.

The function PicOrderCnt(picX) is specified as follows:

PicOrderCnt(picX) = PicOrderCntVal of the picture picX (F-25)

The function DiffPicOrderCnt(picA, picB) is specified as follows:

DiffPicOrderCnt(picA, picB) = PicOrderCnt(picA) – PicOrderCnt(picB)(F-26)

The bitstream shall not contain data that result in values of DiffPicOrderCnt(picA, picB) used in the decoding process that are not in the range of -2^{15} to $2^{15} - 1$, inclusive.

NOTE – Let X be the current picture and Y and Z be two other pictures in the same sequence, Y and Z are considered to be in the same output order direction from X when both DiffPicOrderCnt(X, Y) and DiffPicOrderCnt(X, Z) are positive or both are negative.

[0109] In a fourth example, according to aspects of this disclosure, video decoder 30 may perform the following decoding process for POC values. The fourth example may correspond to the fourth example technique for decrementing POC values described above. That is, the example below may be used to perform the POC decrement process of earlier pictures in the DPB (i.e., for all sub-DPBs) is applied once, at the end of decoding each access unit that requires POC reset. That is, video decoder 30 may decode all pictures of an access unit for with POC resetting is performed, followed by decrementing the POC values of all pictures of the DPB. The fourth example process may be performed as follows:

Output of this process is `PicOrderCntVal`, the picture order count of the current picture.

Picture order counts are used to identify pictures, for deriving motion parameters in merge mode and motion vector prediction, and for decoder conformance checking (see subclause C.5).

Each coded picture is associated with a picture order count variable, denoted as `PicOrderCntVal`.

If `FirstPicInLayerDecodedFlag[nuh_layer_id]` is equal to 1, `poc_reset_idc` is greater than 0, and the *access unit containing the current picture* is the first *access unit* in decoding order in a POC resetting period, the following applies:

- The variables `pocMsbDelta`, `pocLsbDelta` and `DeltaPocVal` are derived as follows:


```

prevPicOrderCntLsb =
PrevPicOrderCnt[ nuh_layer_id ] & ( MaxPicOrderCntLsb - 1 )
prevPicOrderCntMsb = PrevPicOrderCnt[ nuh_layer_id ] - prevPicOrderCntLsb
if( poc_reset_idc == 3 || ( poc_reset_idc == 2 && nuh_layer_id == 0 ) )
    pocLsbVal = poc_lsb_val
else
    pocLsbVal = slice_pic_order_cnt_lsb
pocMsbDelta = getCurrMsb( pocLsbVal, prevPicOrderCntLsb,
prevPicOrderCntMsb, MaxPicOrderCntLsb )
if( poc_reset_idc == 2 || ( poc_reset_idc == 3 && full_poc_reset_flag ) )
    pocLsbDelta = pocLsbVal
else
    pocLsbDelta = 0
DeltaPocVal = pocMsbDelta + pocLsbDelta

```
- The `PicOrderCntVal` of the current picture is derived as follows:


```

if( poc_reset_idc == 1 )
    PicOrderCntVal = slice_pic_order_cnt_lsb
else if( poc_reset_idc == 2 )
    PicOrderCntVal = 0
else { // poc_reset_idc == 3
    PicOrderCntMsb = getCurrMsb( slice_pic_order_cnt_lsb, poc_lsb_val, 0,
MaxPicOrderCntLsb )

```

```

    PicOrderCntVal = PicOrderCntMsb + slice_pic_order_cnt_lsb
}

```

- The value of PrevPicOrderCnt[nuh_layer_id] is derived as follows:
 - If the current picture is not a RASL picture, a RADL picture, or a sub-layer non-reference picture, and the current picture has TemporalId equal to 0, PrevPicOrderCnt[nuh_layer_id] is set equal to PicOrderCntVal.
 - Otherwise when poc_reset_idc is equal to 3, PrevPicOrderCnt[nuh_layer_id] is set equal to full_poc_reset_flag ? 0 : poc_lsb_val.

Otherwise, the following applies:

- *The value of DeltaPocVal is set equal to 0*
- The PicOrderCntVal of the current picture is derived as follows:


```

if( !FirstPicInLayerDecodedFlag[ nuh_layer_id ] ) {
    if( poc_reset_idc == 1 )
        PicOrderCntVal = slice_pic_order_cnt_lsb
    else if( poc_reset_idc == 2 )
        PicOrderCntVal = 0
    else if( poc_reset_idc == 3 ) {
        PicOrderCntMsb = getCurrMsb(slice_pic_order_cnt_lsb, poc_lsb_val, 0,
MaxPicOrderCntLsb )
        PicOrderCntVal = PicOrderCntMsb + slice_pic_order_cnt_lsb
    } else // the current picture is an IRAP picture with NoRaslOutputFlag equal to
1
        PicOrderCntVal = slice_pic_order_cnt_lsb
} else { // the POC derivation as in HEVC version 1
    if( the current picture is an IRAP picture with NoRaslOutputFlag equal to 1 )
        PicOrderCntMsb = 0
    else {
        prevPicOrderCntLsb =
PrevPicOrderCnt[ nuh_layer_id ] & ( MaxPicOrderCntLsb - 1 )
        prevPicOrderCntMsb =
PrevPicOrderCnt[ nuh_layer_id ] - prevPicOrderCntLsb
        PicOrderCntMsb = getCurrMsb( slice_pic_order_cnt_lsb,
prevPicOrderCntLsb,
prevPicOrderCntMsb, MaxPicOrderCntLsb )

```

```

    }
    PicOrderCntVal = PicOrderCntMsb + slice_pic_order_cnt_lsb
  }

```

- The value of PrevPicOrderCnt[nuh_layer_id] is derived as follows:
 - If the current picture is not a RASL picture, a RADL picture, or a sub-layer non-reference picture, and the current picture has TemporalId equal to 0, PrevPicOrderCnt[nuh_layer_id] is set equal to PicOrderCntVal.
 - Otherwise when FirstPicInLayerDecodedFlag[nuh_layer_id] is equal to 0 and poc_reset_idc is equal to 3, PrevPicOrderCnt[nuh_layer_id] is set equal to full_poc_reset_flag ? 0 : poc_lsb_val.

If the current picture is the last picture in an access unit and the value of variable DeltaPocVal is greater than 0, the PicOrderCntVal of all pictures in all sub-DPBs is decremented by the value of variable DeltaPocVal

The value of PicOrderCntVal shall be in the range of -2^{31} to $2^{31} - 1$, inclusive. In one CVS, the PicOrderCntVal values for any two coded pictures in the same layer shall not be the same.

The function PicOrderCnt(picX) is specified as follows:

PicOrderCnt(picX) = PicOrderCntVal of the picture picX (F-25)

The function DiffPicOrderCnt(picA, picB) is specified as follows:

DiffPicOrderCnt(picA, picB) = PicOrderCnt(picA) – PicOrderCnt(picB) F-26

The bitstream shall not contain data that result in values of DiffPicOrderCnt(picA, picB) used in the decoding process that are not in the range of -2^{15} to $2^{15} - 1$, inclusive.

NOTE – Let X be the current picture and Y and Z be two other pictures in the same sequence, Y and Z are considered to be in the same output order direction from X when both DiffPicOrderCnt(X, Y) and DiffPicOrderCnt(X, Z) are positive or both are negative.

[0110] In this manner, video encoder 20 and video decoder 30 represent examples of video coders configured to code data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of the multi-layer video data, wherein the first picture is included in an access unit, and, based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrement POC values of all pictures stored to the DPB that precede the first picture in

coding order including at least one picture of a second layer of the multi-layer video data.

[0111] FIG. 2 is a block diagram illustrating an example of video encoder 20 that may implement techniques for managing POC values according to the techniques of this disclosure. In particular, video encoder 20 may be configured to manage POC values of pictures of a multi-layer bitstream according to techniques of this disclosure. As noted above, video encoder 20 may be adapted to perform multiview and/or scalable video coding. For example, video encoder 20 may be configured to encode a bitstream that conforms to one or more video coding standard extensions, such as SHVC, MV-HEVC, or 3D-HEVC. However, while reference is made to specific coding standards, it should be understood that the techniques are not specific to any one coding standard, and may be implemented with future and/or not yet developed standards.

[0112] Video encoder 20 may perform intra- and inter-coding of video blocks within video slices. Intra-coding relies on spatial prediction to reduce or remove spatial redundancy in video within a given video frame or picture. Inter-coding relies on temporal prediction to reduce or remove temporal redundancy in video within adjacent frames or pictures of a video sequence. Intra-mode (I mode) may refer to any of several spatial based coding modes. Inter-modes, such as uni-directional prediction (P mode) or bi-prediction (B mode), may refer to any of several temporal-based coding modes.

[0113] As shown in FIG. 2, video encoder 20 receives a current video block within a video frame to be encoded. In the example of FIG. 2, video encoder 20 includes video data memory 38, mode select unit 40, reference picture memory 64, summer 50, transform processing unit 52, quantization unit 54, and entropy encoding unit 56. Mode select unit 40, in turn, includes motion compensation unit 44, motion estimation unit 42, intra-prediction unit 46, and partition unit 48. For video block reconstruction, video encoder 20 also includes inverse quantization unit 58, inverse transform unit 60, and summer 62. A deblocking filter (not shown in FIG. 2) may also be included to filter block boundaries to remove blockiness artifacts from reconstructed video. If desired, the deblocking filter would typically filter the output of summer 62. Additional filters (in loop or post loop) may also be used in addition to the deblocking filter. Such filters are not shown for brevity, but if desired, may filter the output of summer 50 (as an in-loop filter).

[0114] During the encoding process, video encoder 20 receives a video frame or slice to be coded. The frame or slice may be divided into multiple video blocks. Video data

memory 38 may store the video data to be encoded by the components of video encoder 20. The video data stored in video data memory 38 may be obtained, for example, from video source 18. Reference picture memory 64 may be referred to as a DPB 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 38 and reference picture memory 64 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 38 and reference picture memory 64 may be provided by the same memory device or separate memory devices. In various examples, video data memory 38 may be on-chip with other components of video encoder 20, or off-chip relative to those components.

[0115] Motion estimation unit 42 and motion compensation unit 44 perform inter-predictive coding of the received video block relative to one or more blocks in one or more reference frames to provide temporal prediction. Intra-prediction unit 46 may alternatively perform intra-predictive coding of the received video block relative to one or more neighboring blocks in the same frame or slice as the block to be coded to provide spatial prediction. Video encoder 20 may perform multiple coding passes, e.g., to select an appropriate coding mode for each block of video data.

[0116] Moreover, partition unit 48 may partition blocks of video data into sub-blocks, based on evaluation of previous partitioning schemes in previous coding passes. For example, partition unit 48 may initially partition a frame or slice into LCUs, and partition each of the LCUs into sub-CUs based on rate-distortion analysis (e.g., rate-distortion optimization). Mode select unit 40 may further produce a quadtree data structure indicative of partitioning of an LCU into sub-CUs. Leaf-node CUs of the quadtree may include one or more PUs and one or more TUs.

[0117] Mode select unit 40 may select one of the coding modes, intra or inter, e.g., based on error results, and provides the resulting intra- or inter-coded block to summer 50 to generate residual block data and to summer 62 to reconstruct the encoded block for use as a reference frame. Mode select unit 40 also provides syntax elements, such as motion vectors, intra-mode indicators, partition information, and other such syntax information, to entropy encoding unit 56.

[0118] Motion estimation unit 42 and motion compensation unit 44 may be highly integrated, but are illustrated separately for conceptual purposes. Motion estimation, performed by motion estimation unit 42, is the process of generating motion vectors,

which estimate motion for video blocks. A motion vector, for example, may indicate the displacement of a PU of a video block within a current video frame or picture relative to a predictive block within a reference picture (or other coded unit) relative to the current block being encoded within the current slice (or other coded unit). A predictive block is a block that is found to closely match the block to be encoded, in terms of pixel difference, which may be determined by sum of absolute difference (SAD), sum of square difference (SSD), or other difference metrics.

[0119] Motion estimation unit 42 calculates a motion vector for a PU of a video block in an inter-coded slice by comparing the position of the PU to the position of a predictive block of a reference picture. The reference picture may be selected from a first reference picture list (List 0) or a second reference picture list (List 1), each of which identify one or more reference pictures stored in reference picture memory 64. Motion estimation unit 42 sends the calculated motion vector to entropy encoding unit 56 and motion compensation unit 44.

[0120] Motion compensation, performed by motion compensation unit 44, may involve fetching or generating the predictive block based on the motion vector determined by motion estimation unit 42. Again, motion estimation unit 42 and motion compensation unit 44 may be functionally integrated, in some examples. Upon receiving the motion vector for the PU of the current video block, motion compensation unit 44 may locate the predictive block to which the motion vector points in one of the reference picture lists. Summer 50 forms a residual video block by subtracting pixel values of the predictive block from the pixel values of the current video block being encoded, forming pixel difference values, as discussed below. In general, motion estimation unit 42 performs motion estimation relative to luma components, and motion compensation unit 44 uses motion vectors calculated based on the luma components for both chroma components and luma components. Mode select unit 40 may also generate syntax elements associated with the video blocks and the video slice for use by video decoder 30 in decoding the video blocks of the video slice.

[0121] Motion estimation unit 42 searches reference pictures stored in a DPB of reference picture memory 64. Mode select unit 40 may determine which reference pictures are to be used to encode a picture, slice of picture, and/or a sequence of pictures. Video encoder 20 may encode picture order count (POC) values for reference pictures used to encode the picture, slice, or sequence. Video encoder 20 may encode

the POC values indicative of the reference pictures to be included in the reference picture list in slice headers and/or parameter sets, such as PPSs and/or SPSs.

[0122] In this manner, a video decoder, such as video decoder 30, may reproduce the reference picture list by including the reference pictures indicated in the slice header and/or parameter set(s). Furthermore, after encoding a block using a motion vector produced by motion estimation unit 42, video encoder 20 may encode motion information for the block, where the motion information may include data representative of the motion vector, an identifier for the reference picture list, and a reference index that identifies a reference picture in the reference picture list.

[0123] In some examples, video encoder 20 may calculate values for sub-integer pixel positions of reference pictures stored in reference picture memory 64. For example, video encoder 20 may interpolate values of one-quarter pixel positions, one-eighth pixel positions, or other fractional pixel positions of the reference picture. Therefore, motion estimation unit 42 may perform a motion search relative to the full pixel positions and fractional pixel positions and output a motion vector with fractional pixel precision.

[0124] Video encoder 20 may also form POC resetting periods and encode values for POC resetting period identifiers for slices of pictures to assign the slices to respective POC resetting periods. Video encoder 20 may assign slices to POC resetting periods individually within each layer of video data. Moreover, video encoder 20 may determine POC resetting types for the slices and encode values for indications (e.g., two-bit values) of the POC resetting types for the slices.

[0125] As discussed with respect to FIG. 1, video encoder 20 may be configured according to restrictions established by an applicable video coding standard. For example, video encoder 20 may ensure that the first picture in each POC resetting period has a POC resetting type indicating that POC values are to be fully or partially reset. Furthermore, video encoder 20 may encode additional information indicative of a POC LSB value in some examples, which may indicate a POC LSB value for a POC anchor picture of the POC resetting period.

[0126] Video encoder 20 may also decrement POC values of reference pictures in the DPB of reference picture memory 64. According to aspects of this disclosure, video encoder 20 may be configured to decrement POC values of all pictures (of all layers) stored to reference picture memory 64 prior to encoding an initial picture of an access unit in which POC resetting is performed. In some examples, video encoder 20 may apply a POC decrement process to earlier pictures in all sub-DPBs (where each layer of

video data is associated with a respective sub-DPB) of the reference picture memory 64 only one time when encoding the access unit for which POC resetting is being performed.

[0127] By decrementing the reference pictures stored to reference picture memory 64, video encoder 20 may use the decremented POC values when encoding POC values in slice headers and/or parameter sets to indicate reference pictures to be included in a reference picture list. The amount by which video encoder 20 decrements the POC values of pictures in the DPB may be the same amount by which the POC anchor picture is decremented. The POC reset of the POC anchor picture may be limited to only resetting POC MSB or resetting the full POC value of the POC anchor picture.

[0128] Intra-prediction unit 46 may intra-predict a current block, as an alternative to the inter-prediction performed by motion estimation unit 42 and motion compensation unit 44, as described above. In particular, intra-prediction unit 46 may determine an intra-prediction mode to use to encode a current block. In some examples, intra-prediction unit 46 may encode a current block using various intra-prediction modes, e.g., during separate encoding passes, and intra-prediction unit 46 (or mode select unit 40, in some examples) may select an appropriate intra-prediction mode to use from the tested modes.

[0129] For example, intra-prediction unit 46 may calculate rate-distortion values using a rate-distortion analysis for the various tested intra-prediction modes, and select the intra-prediction mode having the best rate-distortion characteristics among the tested modes. Rate-distortion analysis generally determines an amount of distortion (or error) between an encoded block and an original, unencoded block that was encoded to produce the encoded block, as well as a bitrate (that is, a number of bits) used to produce the encoded block. Intra-prediction unit 46 may calculate ratios from the distortions and rates for the various encoded blocks to determine which intra-prediction mode exhibits the best rate-distortion value for the block.

[0130] After selecting an intra-prediction mode for a block, intra-prediction unit 46 may provide information indicative of the selected intra-prediction mode for the block to entropy encoding unit 56. Entropy encoding unit 56 may encode the information indicating the selected intra-prediction mode. Video encoder 20 may include in the transmitted bitstream configuration data, which may include a plurality of intra-prediction mode index tables and a plurality of modified intra-prediction mode index tables (also referred to as codeword mapping tables), definitions of encoding contexts

for various blocks, and indications of a most probable intra-prediction mode, an intra-prediction mode index table, and a modified intra-prediction mode index table to use for each of the contexts.

[0131] Video encoder 20 forms a residual video block by subtracting the prediction data from mode select unit 40 from the original video block being coded. Summer 50 represents the component or components that perform this subtraction operation. Transform processing unit 52 applies a transform, such as a discrete cosine transform (DCT) or a conceptually similar transform, to the residual block, producing a video block comprising residual transform coefficient values. Transform processing unit 52 may perform other transforms which are conceptually similar to DCT. Wavelet transforms, integer transforms, sub-band transforms or other types of transforms could also be used. In any case, transform processing unit 52 applies the transform to the residual block, producing a block of residual transform coefficients. The transform may convert the residual information from a pixel value domain to a transform domain, such as a frequency domain.

[0132] Transform processing unit 52 may send the resulting transform coefficients to quantization unit 54. Quantization unit 54 quantizes the transform coefficients to further reduce bit rate. The quantization process may reduce the bit depth associated with some or all of the coefficients. The degree of quantization may be modified by adjusting a quantization parameter. In some examples, quantization unit 54 may then perform a scan of the matrix including the quantized transform coefficients. Alternatively, entropy encoding unit 56 may perform the scan.

[0133] Following quantization, entropy encoding unit 56 entropy codes the quantized transform coefficients. For example, entropy encoding unit 56 may perform context adaptive variable length coding (CAVLC), context adaptive binary arithmetic coding (CABAC), syntax-based context-adaptive binary arithmetic coding (SBAC), probability interval partitioning entropy (PIPE) coding or another entropy coding technique. In the case of context-based entropy coding, context may be based on neighboring blocks. Following the entropy coding by entropy encoding unit 56, the encoded bitstream may be transmitted to another device (e.g., video decoder 30) or archived for later transmission or retrieval.

[0134] Inverse quantization unit 58 and inverse transform unit 60 apply inverse quantization and inverse transformation, respectively, to reconstruct the residual block in the pixel domain, e.g., for later use as a reference block. Motion compensation unit

44 may calculate a reference block by adding the residual block to a predictive block of one of the frames of reference picture memory 64. Motion compensation unit 44 may also apply one or more interpolation filters to the reconstructed residual block to calculate sub-integer pixel values for use in motion estimation. Summer 62 adds the reconstructed residual block to the motion compensated prediction block produced by motion compensation unit 44 to produce a reconstructed video block for storage in reference picture memory 64. The reconstructed video block may be used by motion estimation unit 42 and motion compensation unit 44 as a reference block to inter-code a block in a subsequent video frame.

[0135] In this manner, video encoder 20 of FIG. 2 represents an example of a video encoder configured to encode data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of multi-layer video data, where the first picture is included in an access unit. Video encoder 20 is also configured to, based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrement POC values of all pictures stored to a decoded picture buffer (DPB) that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.

[0136] FIG. 3 is a block diagram illustrating an example of video decoder 30 that may implement techniques for managing POC values according to the techniques of this disclosure. As noted above, video decoder 30 may be adapted to perform multiview and/or scalable video coding. For example, video decoder 30 may be configured to decode a bitstream that conforms to one or more video coding standard extensions, such as SHVC, MV-HEVC, or 3D-HEVC. However, while reference is made to specific coding standards, it should be understood that the techniques are not specific to any one coding standard, and may be implemented with future and/or not yet developed standards.

[0137] In the example of FIG. 3, video decoder 30 includes video data memory 68, an entropy decoding unit 70, motion compensation unit 72, intra prediction unit 74, inverse quantization unit 76, inverse transformation unit 78, reference picture memory 82 and summer 80. Video decoder 30 may, in some examples, perform a decoding pass generally reciprocal to the encoding pass described with respect to video encoder 20 (FIG. 2). Motion compensation unit 72 may generate prediction data based on motion vectors received from entropy decoding unit 70, while intra-prediction unit 74 may

generate prediction data based on intra-prediction mode indicators received from entropy decoding unit 70.

[0138] Video data memory 68 may store video data, such as an encoded video bitstream, to be decoded by the components of video decoder 30. The video data stored in video data memory 68 may be obtained, for example, from computer-readable medium, 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 68 may form a coded picture buffer (CPB) that stores encoded video data from an encoded video bitstream. Reference picture memory 82 may be referred to as a DPB that stores reference video data for use in decoding video data by video decoder 30, e.g., in intra- or inter-coding modes. Video data memory 68 and reference picture memory 82 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 68 and reference picture memory 82 may be provided by the same memory device or separate memory devices. In various examples, video data memory 68 may be on-chip with other components of video decoder 30, or off-chip relative to those components.

[0139] During the decoding process, video decoder 30 receives an encoded video bitstream that represents video blocks of an encoded video slice and associated syntax elements from video encoder 20. Entropy decoding unit 70 of video decoder 30 entropy decodes the bitstream to generate quantized coefficients, motion vectors or intra-prediction mode indicators, and other syntax elements. Entropy decoding unit 70 forwards the motion vectors to and other syntax elements to motion compensation unit 72. Video decoder 30 may receive the syntax elements at the video slice level and/or the video block level.

[0140] When the video slice is coded as an intra-coded (I) slice, intra prediction unit 74 may generate prediction data for a video block of the current video slice based on a signaled intra prediction mode and data from previously decoded blocks of the current frame or picture. When the video frame is coded as an inter-coded (i.e., B, P or GPB) slice, motion compensation unit 72 produces predictive blocks for a video block of the current video slice based on the motion vectors and other syntax elements received from entropy decoding unit 70. The predictive blocks may be produced from one of the reference pictures within one of the reference picture lists. Video decoder 30 may

construct the reference frame lists, List 0 and List 1, using default construction techniques based on reference pictures stored in a DPB of reference picture memory 82.

[0141] More particularly, video decoder 30 may decode a PPS indicating whether slices of a picture corresponding to the PPS include picture order count (POC) resetting information. Assuming the slices include the POC resetting information, video decoder 30 may decode a slice segment header of a slice of the picture including the POC resetting information. The POC resetting information may include a POC resetting period identifier and a POC resetting type.

[0142] The POC resetting period identifier may indicate a POC resetting period to which the slice corresponds. The POC resetting type may indicate whether the POC value of a picture including the slice is fully reset or that only the MSB of the POC value are reset. Furthermore, the POC resetting type may indicate that no POC reset is to be performed or that additional information is signaled. If the POC resetting type indicates that additional information is signaled, video decoder 30 may decode the additional information, which may indicate a POC LSB value and whether the POC value is fully reset or only MSB are reset.

[0143] According to aspects of this disclosure, after decoding information that indicates a POC reset for a slice of video data of an initial picture of an access unit, but prior to decoding the video data of the access unit, video decoder 30 may decrement POC values of pictures stored to reference picture memory 82. According to aspects of this disclosure, in some examples, video decoder 30 may be configured to decrement POC values of all pictures (of all layers) stored to reference picture memory 82 prior to decoding video data of the initial picture of the access unit for which POC resetting is performed. Video decoder 30 may apply a POC decrement process to earlier pictures in all sub-DPBs (where each layer of video data is associated with a respective sub-DPB) of the reference picture memory 82 only one time when decoding the access unit for which POC resetting is being performed.

[0144] After performing the POC reset, video decoder 30 may decode POC values of reference pictures to be included in a reference picture list. Video decoder 30 may decode these POC values in the slice segment headers and/or from parameter sets, such as a PPS or an SPS. Video decoder 30 may then construct a reference picture list including the reference pictures identified by the decoded POC values.

[0145] After constructing the reference picture list for a slice, video decoder 30 may decode blocks of the slice. Motion compensation unit 72 determines prediction

information for a video block of the current video slice by parsing the motion vectors and other syntax elements, and uses the prediction information to produce the predictive blocks for the current video block being decoded. For example, motion compensation unit 72 uses some of the received syntax elements to determine a prediction mode (e.g., intra- or inter-prediction) used to code the video blocks of the video slice, an inter-prediction slice type (e.g., B slice, P slice, or GPB slice), motion vectors for each inter-encoded video block of the slice, inter-prediction status for each inter-coded video block of the slice, and other information to decode the video blocks in the current video slice. Motion information for an inter-predicted block may include a reference picture list identifier and a reference index to identify a reference picture in the reference picture list to use to predict the block.

[0146] Motion compensation unit 72 may also perform interpolation based on interpolation filters. Motion compensation unit 72 may use interpolation filters as used by video encoder 20 during encoding of the video blocks to calculate interpolated values for sub-integer pixels of reference blocks. In this case, motion compensation unit 72 may determine the interpolation filters used by video encoder 20 from the received syntax elements and use the interpolation filters to produce predictive blocks.

[0147] Inverse quantization unit 76 inverse quantizes, i.e., de-quantizes, the quantized transform coefficients provided in the bitstream and decoded by entropy decoding unit 70. The inverse quantization process may include use of a quantization parameter QP_Y calculated by video decoder 30 for each video block in the video slice to determine a degree of quantization and, likewise, a degree of inverse quantization that should be applied.

[0148] Inverse transform unit 78 applies an inverse transform, e.g., an inverse DCT, an inverse integer transform, or a conceptually similar inverse transform process, to the transform coefficients in order to produce residual blocks in the pixel domain.

[0149] After motion compensation unit 72 generates the predictive block for the current video block based on the motion vectors and other syntax elements, video decoder 30 forms a decoded video block by summing the residual blocks from inverse transform unit 78 with the corresponding predictive blocks generated by motion compensation unit 72. Summer 80 represents the component or components that perform this summation operation. If desired, a deblocking filter may also be applied to filter the decoded blocks in order to remove blockiness artifacts. Other loop filters (either in the coding loop or after the coding loop) may also be used to smooth pixel transitions, or

otherwise improve the video quality. The decoded video blocks in a given frame or picture are then stored in reference picture memory 82, which stores reference pictures used for subsequent motion compensation. Reference picture memory 82 also stores decoded video for later presentation on a display device, such as display device 32 of FIG. 1.

[0150] In this manner, video decoder 30 of FIG. 3 represents an example of a video decoder configured to decode data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of multi-layer video data, where the first picture is included in an access unit. Video decoder 30 is also configured to, based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrementing POC values of all pictures stored to a decoded picture buffer (DPB) that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.

[0151] FIG. 4 is a flowchart illustrating an example method for encoding data indicating a POC resetting period for a slice in accordance with the techniques of this disclosure. In addition, the method of FIG. 4 includes encoding a current block. The current block may comprise a current CU or a portion of the current CU. Although described with respect to video encoder 20 (FIGS. 1 and 2), it should be understood that other devices may be configured to perform a method similar to that of FIG. 4.

[0152] In this example, video encoder 20 may assign the current slice to a POC resetting period (148). For example, if the current slice forms part of an anchor picture, video encoder 20 may signal that the current slice forms the beginning of a new POC resetting period with a POC resetting period identifier that is different from previous POC resetting period identifiers. Alternatively, if the current slice does not form part of an anchor picture, video encoder 20 may signal that the current slice forms part of an existing POC resetting period.

[0153] Video encoder 20 then decrements POC values of reference pictures of all pictures that are currently in a DPB (150). That is, according to aspects of this disclosure, video encoder 20 may decrement the POC values of the pictures in all sub-DPBs of the DPB. Video encoder 20 may decrement the POC values after assigning the current slice to a POC resetting period (e.g., encoding the slice header), but prior to encoding video data of the first slice. Video encoder 20 may decrement the POC values of the pictures in the DPB by the same decrement value.

[0154] Video encoder 20 then forms a reference picture list including at least some of the reference pictures (152). Video encoder 20 further signals POC values of the reference pictures included in the reference picture list (154). For example, video encoder 20 may encode POC values (or portions of the POC values, e.g., POC LSB values) for the reference pictures in a parameter set such as an SPS or PPS, and/or in a slice header for the slice. Some reference pictures (e.g., long-term reference pictures) may be signaled in the parameter set, whereas other reference pictures (e.g., short-term reference pictures) may be signaled in the slice header.

[0155] It should be understood that the steps of forming the reference picture list and signaling which pictures are included in the reference picture list may be performed multiple times over several different encoding passes, in order to determine the set of reference pictures that yields the best rate-distortion characteristics for, e.g., blocks of the current slice. That is, video encoder 20 may select the set of reference pictures included in the reference picture list based on characteristics of all blocks in the current slice, and not just based on the individual characteristics of a single block.

[0156] Video encoder 20 may then predict the current block of the current slice (156). For example, video encoder 20 may calculate one or more prediction units (PUs) for the current block. More particularly, motion estimation unit 42 may perform a motion search for the current block among the reference pictures of the reference picture list to identify a matching block used as a reference block, e.g., based on SAD, SSD, MAD, MSD, or other error calculation metrics. Motion estimation unit 42 may produce a motion vector for the current block based on the motion search. Video encoder 20 may then encode motion information for the block (158), which may include motion vector difference values for the motion vector, a reference picture list identifier, and a reference index, which together identify the reference block.

[0157] Video encoder 20 may then calculate a residual block for the current block, e.g., to produce a transform unit (TU) (160). To calculate the residual block, video encoder 20 may calculate a difference between the original, uncoded block and the predicted block for the current block. Video encoder 20 may then transform and quantize coefficients of the residual block (162). Next, video encoder 20 may scan the quantized transform coefficients of the residual block (164). During the scan, or following the scan, video encoder 20 may entropy encode the coefficients (166). For example, video encoder 20 may encode the coefficients using CAVLC or CABAC. Video encoder 20 may then output the entropy coded data of the block (168).

[0158] In this manner, the method of FIG. 4 represents an example of a method including encoding, by a video encoder, data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of multi-layer video data, where the first picture is included in an access unit. The example method of FIG. 4 also includes, based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrementing POC values of all pictures stored to a decoded picture buffer (DPB) that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.

[0159] FIG. 5 is a flowchart illustrating an example method for decoding data indicating a POC resetting period for a slice in accordance with the techniques of this disclosure. In addition, the method of FIG. 4 includes encoding a current block. The current block may comprise a current CU or a portion of the current CU. Although described with respect to video decoder 30 (FIGS. 1 and 3), it should be understood that other devices may be configured to perform a method similar to that of FIG. 5.

[0160] Initially, video decoder 30 decodes data that indicates a POC reset (200). In some examples, as described herein, the data that indicates the POC reset may be a POC resetting identifier that is included in a slice segment header of the current slice. In such examples, video decoder 30 may further decode a POC resetting type indicator. The method of FIG. 5 is based on the assumption that the POC resetting period identifier indicates that the current slice is part of a new POC resetting period. Video decoder 30 may further determine that the current picture is the first picture in a particular access unit (e.g., the picture with the lowest layer identifier among all pictures of a particular access unit).

[0161] Video decoder 30 may also decrement POC values of all pictures stored to a DPB (202). For example, according to aspects of this disclosure, video decoder 30 may decrement the POC values of the pictures in all sub-DPBs of the DPB. In some examples, video decoder 30 may decrement the POC values after decoding the slice header for the current slice (e.g., as performed in step 200), but prior to decoding video data of the first slice as described below. Video decoder 30 may decrement the POC values of the pictures in the DPB by the same decrement value.

[0162] Video decoder 30 then decodes POC values of reference pictures to be included in a reference picture list (204), e.g., from the slice segment header and/or a parameter set such as a PPS or SPS. Video decoder 30 then forms the reference picture list (206).

[0163] Next, video decoder 30 decodes motion information for a current block (208). The motion information may include, for example, a reference picture list identifier and a reference index into the reference picture list. Video decoder 30 then predicts the current block (200), e.g., using inter-prediction mode to calculate a predicted block for the current block. More particularly, video decoder 30 uses the reference picture list identifier to identify which reference picture list to use, and the reference index to identify a reference index in the reference picture list. Video decoder 30 then decodes a motion vector for the current block and identifies a reference block in the identified reference picture.

[0164] Video decoder 30 may also receive entropy coded data for the current block, such as entropy coded data for coefficients of a residual block corresponding to the current block (212). Video decoder 30 may entropy decode the entropy coded data to reproduce coefficients of the residual block (214). Video decoder 30 may then inverse scan the reproduced coefficients (216), to create a block of quantized transform coefficients. That is, using the inverse scan, video decoder 30 converts a one-dimensional vector to a two-dimensional matrix. Video decoder 30 may then inverse quantize and inverse transform the coefficients to produce a residual block (218). Video decoder 30 may ultimately decode the current block by combining the predicted block and the residual block (220).

[0165] In this manner, the method of FIG. 5 represents an example of a method including decoding data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of multi-layer video data, where the first picture is included in an access unit. The example method of FIG. 5 also includes, based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrementing POC values of all pictures stored to a decoded picture buffer (DPB) that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.

[0166] 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.

[0167] 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 on a computer-readable medium and executed by a hardware-based processing unit. Computer-readable media may include computer-readable storage media, which corresponds to a tangible medium such as data storage media, or communication media including any medium that facilitates transfer of a computer program from one place to another, e.g., according to a communication protocol. In this manner, computer-readable media generally may correspond to (1) tangible computer-readable storage media which is non-transitory or (2) a communication medium such as a signal or carrier wave. Data storage media may be any available media that can be accessed by one or more computers or one or more processors to retrieve instructions, code and/or data structures for implementation of the techniques described in this disclosure. A computer program product may include a computer-readable medium.

[0168] 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 transitory media, but are instead directed to non-transitory, 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.

[0169] Instructions may be executed by one or more processors, such as one or more digital signal processors (DSPs), general purpose microprocessors, application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs), or other

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

[0170] 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.

[0171] Various examples have been described. These and other examples are within the scope of the following claims.

WHAT IS CLAIMED IS:

1. A method of decoding video data, the method comprising:
decoding data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of multi-layer video data, wherein the first picture is included in an access unit; and
based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrementing POC values of all pictures stored to a decoded picture buffer (DPB) that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.
2. The method of claim 1, wherein decrementing the POC values of all pictures stored to the DPB that temporally precede the first picture comprises decrementing the POC values only one time during decoding of the access unit.
3. The method of claim 1, wherein decrementing the POC values comprises decrementing the POC values after decoding a slice header for an initial slice of video data of the first picture.
4. The method of claim 1, wherein decoding the data that indicates the POC reset comprises decoding data representing a POC reset index for a slice of the first picture of video data.
5. The method of claim 1, wherein the data that indicates the POC reset comprises data that indicates a POC resetting period, and wherein the access unit is the first access unit in decoding order of the POC resetting period.
6. The method of claim 1, wherein the first picture is a picture associated with a lowest layer identifier that is present in the access unit, and wherein decrementing the POC values comprises decrementing the POC values using a common POC decrement value.
7. The method of claim 6, further comprising decoding a DeltaPocVal syntax element that indicates a POC decrement value for decrementing the POC values.

8. The method of claim 1, wherein the access unit includes the first picture and at least one additional picture, wherein the data indicating the POC reset comprises a POC reset index, the method further comprising decoding the same POC reset index for all other pictures of the access unit including the at least one second picture of the access unit.
9. The method of claim 1, wherein decrementing the POC values comprises one of decrementing all bits of the POC values, decrementing the most significant bits of the POC values, or decrementing the least significant bits of the POC values.
10. The method of claim 1, wherein the data that indicates the POC reset value is associated with a POC resetting period that is applicable to all layers of the multi-layer video data.
11. A method of encoding video data, the method comprising:
 - encoding data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of multi-layer video data, wherein the first picture is included in an access unit; and
 - based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrementing POC values of all pictures stored to a decoded picture buffer (DPB) that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.
12. The method of claim 11, wherein decrementing the POC values of all pictures stored to the DPB that temporally precede the first picture comprises decrementing the POC values only one time during encoding of the access unit.
13. The method of claim 11, wherein decrementing the POC values comprises decrementing the POC values after encoding a slice header for an initial slice of video data of the first picture.
14. The method of claim 11, wherein encoding the data that indicates the POC reset comprises encoding data representing a POC reset index for a slice of the first picture of video data.

15. The method of claim 11, wherein the data that indicates the POC reset comprises data that indicates a POC resetting period, and wherein the access unit is the first access unit in decoding order of the POC resetting period.

16. The method of claim 11, wherein the first picture is a picture associated with a lowest layer identifier that is present in the access unit, and wherein decrementing the POC values comprises decrementing the POC values using the same POC decrement value.

17. The method of claim 16, further comprising encoding a DeltaPocVal syntax element that indicates the POC decrement value for decrementing the POC values.

18. The method of claim 11, wherein the access unit includes the first picture and at least one second picture, wherein the data indicating the POC reset comprises a POC reset index, and further comprising encoding the same POC reset index for all other pictures of the access unit including the at least one second picture of the access unit.

19. The method of claim 11, wherein decrementing the POC values comprises one of decrementing all bits of the POC values, decrementing the most significant bits of the POC values, or decrementing the least significant bits of the POC values.

20. The method of claim 11, wherein the data that indicates the POC reset value is associated with a POC resetting period that is applicable to all layers of the multi-layer video data.

21. A device for coding video data, the device comprising:
a memory comprising a decoded picture buffer (DPB) configured to store multi-layer video data; and
a video coder configured to:
code data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of the multi-layer video data, wherein the first picture is included in an access unit; and
based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrement POC values of all pictures stored to the DPB that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.

22. The device of claim 21, wherein to decrement the POC values of all pictures stored to the DPB that temporally precede the first picture, the video coder is configured to decrement the POC values only one time during coding of the access unit.

23. The device of claim 21, wherein to decrement the POC values, the video coder is configured to decrement the POC values after encoding a slice header for an initial slice of video data of the first picture.

24. The device of claim 21, wherein to code the data that indicates the POC reset, the video coder is configured to code data representing a POC reset index for a slice of the first picture of video data.

25. The device of claim 21, wherein the data that indicates the POC reset comprises data that indicates a POC resetting period, and wherein the access unit is the first access unit in decoding order of the POC resetting period.

26. The device of claim 21, wherein the first picture is a picture associated with a lowest layer identifier that is present in the access unit, and wherein to decrement the POC values, the video coder is configured to decrement the POC values using the same POC decrement value.

27. The device of claim 26, wherein the video coder is further configured to code a DeltaPocVal syntax element that indicates the POC decrement value for decrementing the POC values.

28. The device of claim 21, wherein the access unit includes the first picture and at least one second picture, wherein the data indicating the POC reset comprises a POC reset index, and wherein the video coder is further configured to code the same POC reset index for all other pictures of the access unit including the at least one second picture of the access unit.

29. The device of claim 21, wherein decrementing the POC values comprises one of decrementing all bits of the POC values, decrementing the most significant bits of the POC values, or decrementing the least significant bits of the POC values.

30. The device of claim 21, wherein the data that indicates the POC reset value is associated with a POC resetting period that is applicable to all layers of the multi-layer video data.
31. The device of claim 21, wherein the device comprises at least one of:
an integrated circuit;
a microprocessor; or
a wireless communication device.
32. A device for coding video data, the device comprising:
means for coding data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of the multi-layer video data, wherein the first picture is included in an access unit; and
based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, means for decrementing POC values of all pictures stored to the DPB that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.
33. The device of claim 32, wherein the means for decrementing the POC values of all pictures stored to the DPB that temporally precede the first picture comprises means for decrementing the POC values only one time during coding of the access unit.
34. The device of claim 32, wherein the means for decrementing the POC values comprises decrementing the POC values after encoding a slice header for an initial slice of video data of the first picture.
35. The device of claim 32, wherein the means for coding the data that indicates the POC reset comprises means for coding data representing a POC reset index for a slice of the first picture of video data.
36. The device of claim 32, wherein the data that indicates the POC reset comprises data that indicates a POC resetting period, and wherein the access unit is the first access unit in decoding order of the POC resetting period.
37. A non-transitory computer-readable storage medium having stored thereon instructions that, when executed, cause a processor of a device for coding video data to:

code data that indicates a picture order count (POC) reset for a POC value of a first picture of a first layer of the multi-layer video data, wherein the first picture is included in an access unit; and

based on the data that indicates the POC reset for the POC value of the first picture and prior to decoding the first picture, decrement POC values of all pictures stored to the DPB that precede the first picture in coding order including at least one picture of a second layer of the multi-layer video data.

38. The non-transitory computer-readable storage medium of claim 37, wherein to decrement the POC values of all pictures stored to the DPB that temporally precede the first picture, the instructions cause the processor to decrement the POC values only one time during coding of the access unit.

39. The non-transitory computer-readable storage medium of claim 37, wherein to decrement the POC values, the instructions cause the processor to decrement the POC values after encoding a slice header for an initial slice of video data of the first picture.

40. The non-transitory computer-readable storage medium of claim 37, wherein to code the data that indicates the POC reset, the instructions cause the processor to code data representing a POC reset index for a slice of the first picture of video data.

41. The non-transitory computer-readable storage medium of claim 37, wherein the data that indicates the POC reset comprises data that indicates a POC resetting period, and wherein the access unit is the first access unit in decoding order of the POC resetting period.

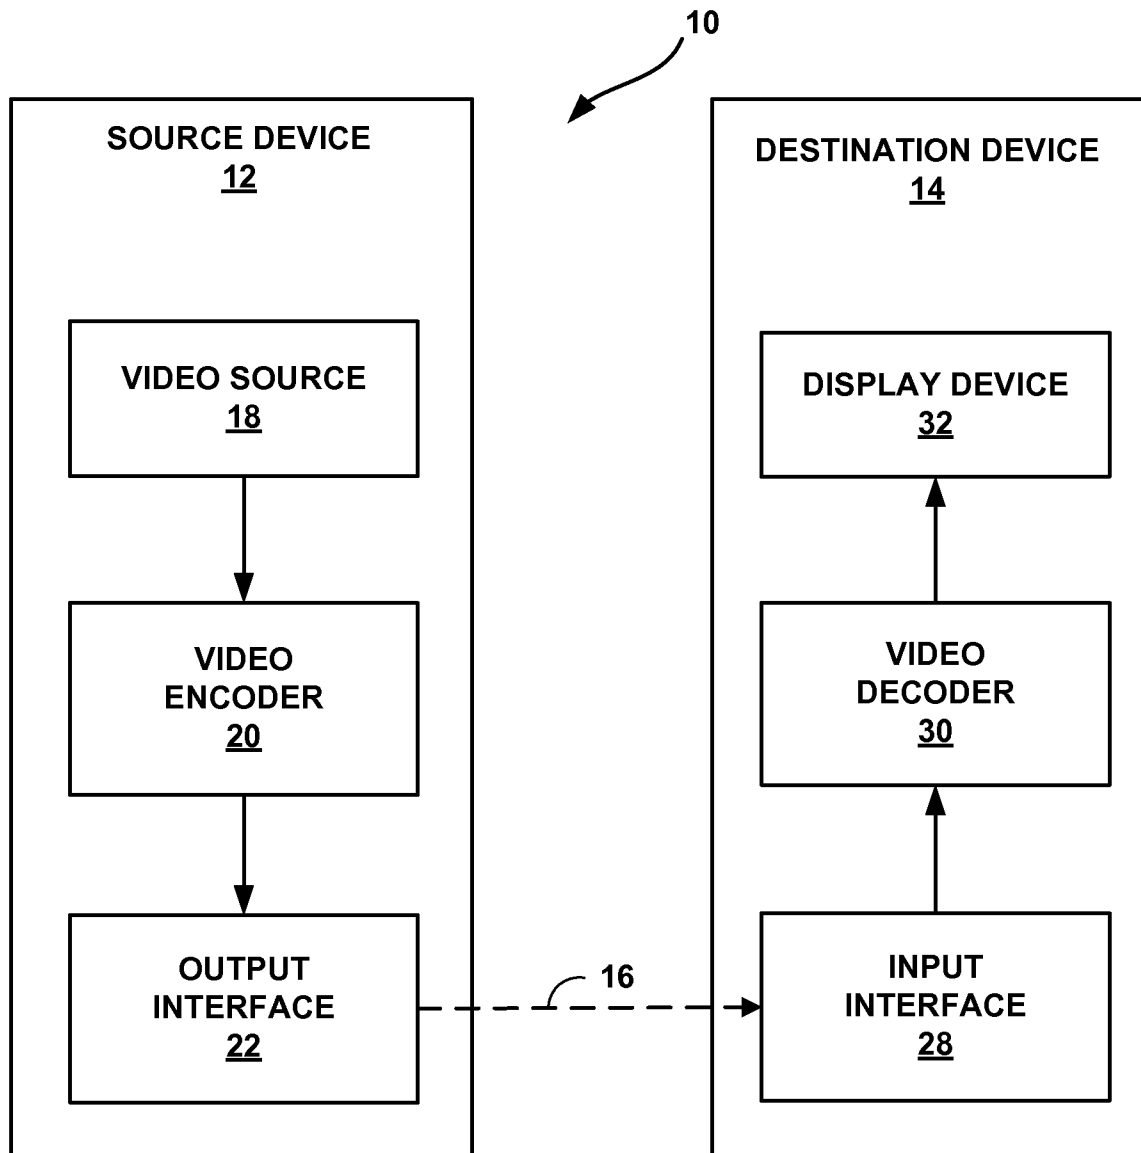


FIG. 1




FIG. 2




FIG. 3




FIG. 4

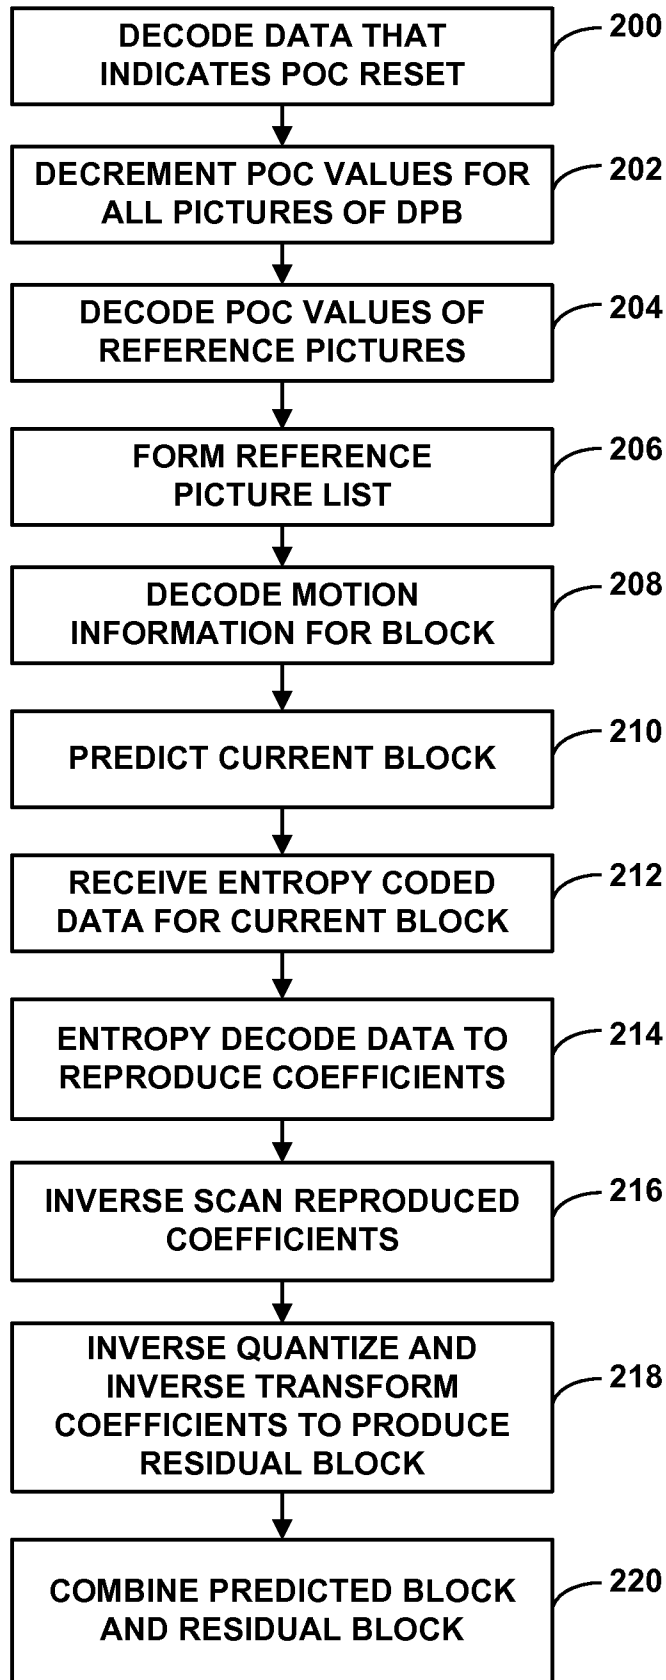


FIG. 5

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/067363

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04N19/70 H04N19/46 H04N19/30 H04N19/68
ADD.
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
Minimum documentation searched (classification system followed by classification symbols)
H04N
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	ADARSH K RAMASUBRAMONIAN ET AL: "MV-HEVC/SHVC HLS: On picture order count", JOINT COLLABORATIVE TEAM ON VIDEO CODING (JCT-VC) OF ITU-T SG 16 WP 3 AND ISO/IEC JTC 1/SC 29/WG 11. JOINT COLLABORATIVE TEAM ON 3D VIDEO CODING EXTENSIONS OF ITU-T SG 16 WP 3 AND ISO/IEC JTC 1/SC 29/WG 11,, no. JCTVC-00213 v3, 21 November 2013 (2013-11-21), pages 1-11, XP007922965, the whole document ----- -/--	1-41

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search 22 January 2015	Date of mailing of the international search report 10/02/2015
------------------------------------------------------------------------------	------------------------------------------------------------------

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

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2014/067363

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	CHEN Y ET AL: "MV-HEVC/SHVC HLS: Cross-layer POC alignment", 5. JCT-3V MEETING; 27-7-2013 - 2-8-2013; VIENNA; (THE JOINT COLLABORATIVE TEAM ON 3D VIDEO CODING EXTENSION DEVELOPMENT OF ISO/IEC JTC1/SC29/WG11 AND ITU-T SG.16); URL: HTTP://PHENIX.INT-EVRY.FR/JCT2/,, no. JCT3V-E0075, 16 July 2013 (2013-07-16) , XP030131078,	1,11,21, 32,37
A	Chapters 1-2; abstract	2-10, 12-20, 22-31, 33-36, 38-41
X,P	----- TECH G ET AL: "MV-HEVC Draft Text 9", 9. JCT-3V MEETING; 3-7-3014 - 9-7-2014; SAPPORO; (THE JOINT COLLABORATIVE TEAM ON 3D VIDEO CODING EXTENSION DEVELOPMENT OF ISO/IEC JTC1/SC29/WG11 AND ITU-T SG.16); URL: HTTP://PHENIX.INT-EVRY.FR/JCT2/,, no. JCT3V-I1002, 10 July 2014 (2014-07-10) , XP030132530, F.3.25-F.3.26 F.7.4.2.4.4 -----	1-41